The 2013 Guide to Network Virtualization and SDN

By Dr. Jim Metzler, Ashton Metzler & Associates
Distinguished Research Fellow and Co-Founder
Webtorials Analyst Division

Platinum Sponsors:


Gold Sponsors:

Produced by:

Table of Contents

Executive Summary 1

Chapter 1: The what, why and how of Network Virtualization 4

Introduction 4
Traditional NV & The NV Use Case 6
Network Overlays via Tunneling: Benefits & Limitations 7
Cloud Orchestration 11
Controller Based NV Solution Architecture 13
Criteria to Evaluate Overlay NV Solutions 14
Tunnel Encapsulation 16
Tunnel Control 17
Comparison of Network Overlay Virtualization Solutions 19
Software Defined NV via Flow Table Segmentation 20
Enterprise Plans for NV Adoption 21

Chapter 2: The what, why and how of SDN 23

Background 23
Potential SDN Use Cases 24
A Working Definition of SDN 27
The SDN Solution Architecture 29
Criteria to Evaluate SDN Solution Architectures 31
The Inhibitors to SDN Adoption 32
The Overlay/Underlay Model 33
Network Function Virtualization 33
The Open Networking Foundation and OpenFlow 36
Potential Use Cases and Benefits of OpenFlow 38
The OpenDaylight Consortium 41
Security 43
Management 44
Appendix – Chapter 2 46

Chapter 3: The NV and SDN Ecosystem 48

Overview of the NV and SDN Ecosystem 48
Representative Vendors 53

Chapter 4: Planning for NV and SDN 82

Introduction 82
Market Research: The Current State of Planning 83
Crafting an NV and/or SDN Plan 86
Summary and Conclusions 94

Advertorials *(please click on the sponsor's name to view their advertorial)*

A10
Alcatel-Lucent
Avaya – Software-Defined Data Center Architecture
Avaya – Ten things to know about Fabric Connect
Ciena – Software Defined Networking
Ciena – The Future is OP^n
Cisco
Dell
EMC^2
Extreme
HP
NEC
Netsocket
NuageNetworks
Packet Design
Pertino
Pica8
QualiSystems
Radware

Executive Summary

Over the last year, the hottest topics in networking have been Network Virtualization (NV) and Software Defined Networking (SDN). There is, however, considerable confusion amongst enterprise IT organizations relative to these topics. There are many sources of that confusion, including the sheer number of vendors who have solutions that solve different problems using different solution architectures and technologies, all of whom claim to be offering SDN and/or NV solutions.

The primary goal of the ***2013 Guide to Software Defined Networking & Network Virtualization*** (The Guide) is to eliminate that confusion and accelerate the adoption of NV and/or SDN. The guide will achieve that goal by walking the readers through the following set of topics:

1. What are the problems and opportunities that NV and SDN help to address?
2. What are the primary characteristics of NV and SDN solutions?
3. How does NV and SDN help IT organizations respond to problems and opportunities?
4. How are IT organizations approaching the evaluation and deployment of NV and/or SDN?
5. What is the role of organizations such as the ONF and the OpenDayLight consortium?
6. What approach are the key vendors taking relative to NV and SDN?
7. What should IT organizations do to get ready for NV and SDN?

The Guide was published both in its entirety and in a chapter-by-chapter fashion.

Chapter one of The Guide is focused on NV and it discusses:

- The primary NV use cases;
- Various ways that overlay NV solutions can be implemented and the benefits and limitations of those solutions;
- The drivers and inhibitors of NV adoption;
- The role of orchestration and service chaining;
- A canonical architecture for a controller based NV solution;
- A comparison of current NV solutions;
- How to create NV solutions by manipulating OpenFlow tables.

The second chapter focuses on SDN and it discusses:

- Potential SDN use cases;
- The ONF definition of SDN and the ONF solution architecture;

- Criteria to evaluate SDN solution architectures;
- The drivers and inhibitors of SDN adoption;
- The overlay/underlay model;
- Network function virtualization;
- Potential use cases and benefits of OpenFlow;
- The OpenDaylight consortium;
- The impact of SDN on security and on management.

The third chapter focuses on the NV and SDN ecosystem. The chapter identifies the primary classes of vendors in the ecosystem, their value proposition and also identifies some of the key players in each class of vendor. The classes of vendors discussed in chapter 3 are:

- Merchant silicon/chip vendors;
- Hyperscale data centers;
- Telecom service providers;
- Switch vendors;
- Network management and automation vendors;
- Providers of network services;
- Providers of test equipment or testing services;
- Standards bodies;
- Providers of controllers;
- Providers of Telecom Service Provider's infrastructure/optical equipment;
- Networking vendors;
- Server virtualization vendors.

Chapter 3 also profiles some of the major players in the NV and SDN ecosystem. Included in each profile is the focus each vendor is taking, a discussion of their value proposition and the identification of some proof points that demonstrate their creditability. The following vendors are profiled in chapter 3:

- Avaya;
- QualiSystems;

- Cisco;
- NEC;
- Nuage Networks;
- HP;
- Ciena;
- A10;
- Pica8;
- Packet Design;
- Netsocket;
- EMC.

Chapter 4 presents some market research that describes the current state of planning for NV and SDN. In addition, the chapter presents an outline that IT organizations can modify to use in their environment to plan for NV and SDN. That outline suggests that IT organizations that are interested in NV and SDN should:

- Create definitions of NV and SDN that are agreed to and well understood within their organization;
- Identify the primary opportunities that the organization is hoping to address and identify the key business metrics associated with each opportunity;
- Evaluate viable solutions;
- Determine how to integrate the solutions into the current environment;
- Educate the organization;
- Evaluate professional services;
- Eliminate organizational resistance;
- Perform a POC;
- Obtain management buy-in.

Chapter 1: The what, why and how of Network Virtualization

Introduction

Over the last couple of years a number of approaches to NV have emerged that are focused on addressing the limitations of the traditional techniques for network virtualization (e.g., 802.1Q VLANs and Virtual Routing and Forwarding (VRFs)). All of these approaches are based on creating a number of virtual Layer 2 or Layer 3 networks that are supported by a common physical infrastructure. The basic idea is to virtualize the network in a manner analogous to compute server virtualization. As a result of these developments, network designers will have the opportunity to choose among the following NV alternatives.

1. Traditional NV
2. Overlay Network Virtualization via Tunneling
3. Software Defined NV via Flow Table Segmentation
4. A combination of the above alternatives

The Survey Respondents were asked to indicate how their organization defines network virtualization and multiple answers were allowed. The survey question focused on the emerging forms of network virtualization – bullets 2 and 3 in the preceding list. As indicated in **Table 1**, some of the emerging forms of network virtualization are based on a device referred to as a controller. As is described below, one of the key roles of a controller is to serve as a central repository of address mappings.

The responses to this question are shown in **Table 1**.

Table 1: Characterization of NV Solutions

Definition of Network Virtualization	Percentage of Respondents
It is based on overlays using protocols such as VXLAN, NVGRE or STT but it does not involve a controller	21.0%
It is based on overlays and a controller. It may or may not use protocols such as VXLAN, NVGRE or STT	39.1%
It is part of a software defined network and may be based on segregating traffic flows	36.2%
Don't know	17.7%
Other	4.5%

The data in **Table 1** indicates that of the emerging forms of network virtualization, the controller-based approaches to NV are by a wide margin the most popular.

VXLAN, NVGRE and STT are all draft IETF standards. To understand the role that standards play in the selection of NV solutions, The Survey Respondents were asked how important it was to their organization that NV solutions are based on open standards. Their responses are shown in **Table 2**.

Table 2: Importance of Open Standards	
Level of Importance	**Percentage of Respondents**
Extremely important	16.0%
Very important	32.1%
Moderately important	24.7%
Somewhat important	14.4%
Not important	7.4%
Don't know	5.3%

The data in Table 2 indicates that NV solutions that are built on open standards are either very or extremely important to roughly half of The Survey Respondents.

Traditional NV & The NV Use Case

One-to-many virtualization of network entities is not a new concept. The most common traditional applications of the virtualization concept to networks are VRF instances and VLANs.

VRF is a form of Layer 3 network virtualization in which a physical router supports multiple virtual router (VR) instances, each running its own routing protocol instance and maintaining its own forwarding table. Unlike VLANs, VRF does not use a tag in the packet header to designate the specific VRF to which a packet belongs. The appropriate VRF is derived at each hop based on the incoming interface and information in the frame. An additional requirement is that each intermediate router on the end-to-end path traversed by a packet needs to be configured with a VRF instance that can forward that packet.

VLANs partition the standard Ethernet network into as many as 4,096 broadcast domains as designated by a 12 bit VLAN ID tag in the Ethernet header. VLANs have been a convenient means of isolating different types of traffic that share a common switched LAN infrastructure. In data centers making extensive use of server virtualization, the limited number of available VLAN IDs can present problems, especially in cases where a large number of tenants need to be supported, each of whom requires multiple VLANs. In contrast to this limitation of VLANs, part of the use case for the NV approaches that are described in The Guide is that these approaches enable IT organizations to establish virtual Ethernet networks without being constrained to only having 4,096 VLAN IDs.

Server virtualization is another factor that is driving the adoption of the approaches to NV that are described in this sub-section of The Guide. Due to server virtualization, virtual machines (VMs) can be dynamically created and moved, both within a data center and between data centers. Extending VLANs across a data center via 802.1Q trunks to support VM mobility adds operational cost and complexity due to the fact that each switch in end-to-end path has to be manually reconfigured. In data centers based on Layer 2 server-to-server connectivity, large numbers of VMs, each with its own MAC address, can also place a burden on the forwarding tables capacities of Layer 2 switches. A major component of the value proposition for the NV approaches that are described in The Guide is that they support the dynamic movement, replication and allocation of virtual resources without manual intervention. Another component of the value proposition for these approaches is that they avoid the issue of needing more MAC addresses than data center LAN switches can typically support.

The value proposition of network overlay solutions is expanded upon in the following sub-section. As is also described below, one characteristic of NV solutions that IT organizations need to understand is whether the solution enables the dynamic movement of virtual resources within a data center; between data centers; or between a data center and a branch or campus facility. A related characteristic that IT organizations need to understand is whether the solution leverages standards based protocols to federate with other NV solutions.

Network Overlays via Tunneling: Benefits & Limitations

A number of approaches to network virtualization leverage tunneling and encapsulation techniques to construct multiple virtual network topologies overlaid on a common physical network. A virtual network (VN) can be a Layer 2 network or a Layer 3 network, while the physical network can be Layer 2, Layer 3 or a combination depending on the overlay technology. With overlays, the outer (encapsulating) header includes a field (generally up to 24 bits wide) that carries a virtual network instance ID (VNID) that specifies the virtual network designated to forward the packet.

Virtual network overlays can provide a wide range of benefits, including:

- Virtualization is performed at the network edge, while the remainder of the L2/L3 network remains unchanged and doesn't need any configuration change in order to support the virtualization of the network. The most common approach is to perform the encapsulation at the hypervisor vSwitch, which acts as the virtual tunnel endpoint (VTEP) or network virtualization edge (NVE). As a result, overlay NV solutions can generally be implemented over existing networks as either an enhancement to the conventional distributed network architecture, or as a step toward an SDN architecture.

- Support for essentially unlimited numbers of VNs as the 24 bits that are typically used by network overlays to identify VNs can identify slightly more than 16 million VN IDs. While theoretically NV solutions can support 16 million VNs, practical limits are often in the range of 16,000 to 32,000 VNs.

- Decoupling of the virtual network topology from the physical network Infrastructure and decoupling of the "virtual" MAC and/or IP addresses used by VMs from the infrastructure IP addresses used by the physical data center core network. The decoupling 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 be provisioned entirely at the edge of the network.

- Ability to manage overlapping IP addresses between multiple tenants.

- Support for multi-path forwarding within virtual networks.

- Ease of provisioning virtual appliances in the data path. Network services resident on VMs can be chained together (a.k.a., service chaining) with point-and-click simplicity under the control of NV software.

- For controller-based NV solutions, the controller is not in the data path, and so it does not present a potential bottleneck.

The Survey Respondents were given a set of 15 possible challenges and opportunities and were asked to indicate which challenges and opportunities they thought that NV solutions could help them to respond to. The Survey Respondents were allowed to indicate multiple challenges and opportunities. The top 5 challenges and opportunities are shown in **Table 3.**

Table 3: Use Cases for NV Solutions	
Challenge/Opportunity	**Percentage of Respondents**
Better utilize network resources	44.0%
Support the dynamic movement, replication and allocation of virtual resources	39.1%
Establish virtual Ethernet networks without the limit and configuration burden of VLANs	32.5%
More easily scale network functionality	31.7%
Reduce OPEX	30.5%

Given the similarity of the second and third entries in **Table 3**, it follows that the primary value that IT organizations see in NV solutions is the ability to dynamically implement virtual Ethernet networks that can support the dynamic movement, replication and allocation of virtual resources.

Some of the limitations of overlay NV solutions include:

- 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 NV requires a lightly oversubscribed or non-oversubscribed physical underlay network.
- Gateways between the virtual network and systems and network service points on the physical network may need to pass high volumes of traffic. If a software gateway running on a VM or a dedicated appliance has insufficient processing power, hardware support for the gateway functionality may be required in physical switches or network service appliances. Some of the more recent merchant silicon switching chips support gateway functionality for VXLAN which is the most popular encapsulation protocol.
- Some value-added features in existing networks cannot be leveraged due to encapsulation. For example, the physical network loses its ability to provide differentiated services based on the content of the packet header.

NV solutions also create some management challenges. For example, one of the primary benefits of overlay solutions is the ability to support multiple VNs running on top of the physical network. Effective operations management requires that IT organizations have tools that give them clear visibility into the relationships between virtual and physical networks and their component devices. When performance or availability problems occur, both root cause analysis and impact analysis require bilateral mapping between the physical and virtual infrastructures.

Both increasing and complicating the need for the visibility described in the preceding paragraph is the ability of NV solutions to do service chaining. The phrase *service chaining* refers to the ability to steer VM-VM traffic flows through a sequence of physical or virtual servers that provide network services, such as firewalls, IPS/IDS, DPI, or load balancers. The primary focus of service chaining is on services provided by virtual appliances. Most SDN or NV solutions

provide service chaining. For SDN, the controller configures the forwarding plane switches to direct the flows along the desired paths, For NV, the controller adjust the FIBs of the vSwitches/vRouters to force the traffic through the right sequence of VMs. Network Function Virtualization, discussed in the next section of The Guide, is basically service chaining that focuses on network services/functions provided by virtual appliances, but isn't necessarily dependent on SDN or NV.

The bottom line is that IT organizations need visibility not just into the overlay NV solution but into the complete solution and all of its components; e.g., firewalls, load balancers.

The Survey Respondents were given a set of 12 inhibitors to the adoption of NV and were asked to indicate the two biggest inhibitors to their company adopting NV sometime in the next two years. The top 5 inhibitors are shown in **Table 4**.

Table 4: Inhibitors to the Adoption of NV Solutions

Inhibitor	% of Respondents
The immaturity of the current products	29.6%
The lack of resources to evaluate NV	29.2%
Other technology and/or business priorities	28.8%
The immaturity of the enabling technologies	29.6%
The confusion and lack of definition in terms of vendors' strategies	18.1%

One interesting observation that can be drawn from the data in **Table 4** is that IT organizations are not avoiding implementing NV solutions because they don't see value in them. Rather, the key factors inhibiting the adoption of NV solutions are the same factors that typically inhibit the adoption of any new technology or way of implementing technology: Immaturity of products and strategies; confusion; and lack of resources.

The Survey Respondents were asked to indicate the impact they thought that NV would have on security and network management. Their responses are shown in **Table 5** and **Table 6**.

Table 5: Impact of NV on Security

Impact on Security	% of Respondents
Networks will be much more secure	6.2%
Networks will be somewhat more secure	33.7%
NV will have no impact on network security	23.5%
Networks will be somewhat less secure	14.0%
Networks will be much less secure	2.5%
Don't know	20.2%

Table 6: Impact of NV on Management	
Impact on Management	**% of Respondents**
Networks will be much easier to manage	21.8%
Networks will be somewhat easier to manage	52.3%
NV will have no impact on management	4.5%
Networks will be somewhat more difficult to manage	9.9%
Networks will be much more difficult to manage	4.5%
Don't know	7.0%

One conclusion that can be drawn from the data in **Table 5** and **Table 6** is that The Survey Respondents generally think that implementing NV solutions will make their networks more secure and easier to manage. As such, security and ease of management can potentially be looked at as benefits of implementing NV solutions.

Cloud Orchestration

Cloud Orchestration platforms have evolved as a means of automating and facilitating the process of configuring pools of data center resources in order to provide a range of cloud or cloud-like services, such as Infrastructure as a Service (IaaS) solutions. The Orchestrator's role is to manipulate the basic resources of the data center (i.e., VMs, networks, storage, and applications) at a very high level of abstraction to create the service. Orchestration is most effective when the data center is fully virtualized, facilitating software control/reconfiguration and automation. As a result, there is a naturally affinity between Orchestration and software-based network controllers, such as NV controllers or SDN controllers.

OpenStack is a cloud computing orchestration project offering free open source software released under the terms of the Apache License. The project is managed by the OpenStack Foundation, a non-profit corporate entity established in September 2012 to promote OpenStack software and its community. Apache CloudStack is another open source Apache Licensed orchestration system. Eucalyptus is a third open source orchestrator with tight technical ties to Amazon Web Services (AWS).

In addition, there are a number of proprietary orchestrators that offer open APIs to allow integration across vendor boundaries. These include VMware's vCloud Director and IBM's SmartCloud Orchestrator.

Figure 1 shows a block diagram of the OpenStack system, including the OpenStack modules that are used to control resource pools in the data center. Horizon is the OpenStack Dashboard that provides administrators and users a graphical interface to access, provision and automate cloud-based resources.

Figure 1: OpenStack

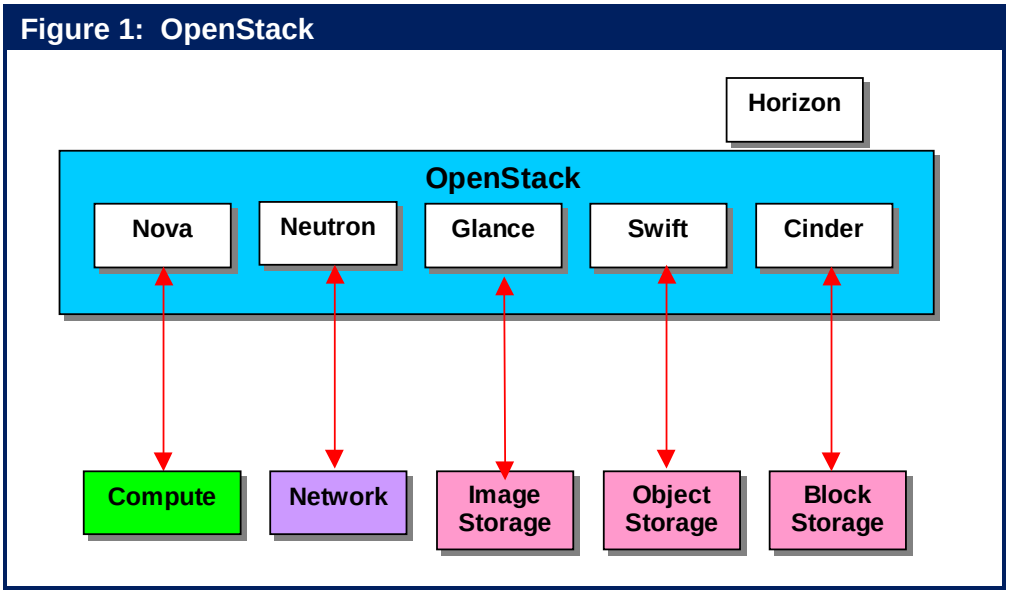

Neutron (formerly called Quantum) allows users to create their own networks, provide connectivity for servers and devices, and control traffic. With appropriate Neutron plug-ins, administrators can take advantage of various NV and SDN solutions to allow for multi-tenancy and scalability. OpenStack networking also has an extension framework allowing additional network services, such as intrusion detection systems (IDS), load balancing, firewalls and virtual private networks (VPN) to be deployed and managed.

In conjunction with the Orchestrator, the role of the SDN or NV controller is to translate the abstract model created on the Orchestrator into the appropriate configuration of the virtual and physical resources that will deliver the desired service. For example, the orchestrator can instruct the controller to perform a variety of workflows, including:

- Create a VM
- Assign a VM to a Virtual Network (VN)
- Connect a VM to an external network
- Apply a security policy to a group of VMs or a Virtual Network
- Attach Network Services to a VM or chain Network Services between VMs

Figure 2 provides a high level depiction of how an orchestrator (OpenStack) and a NV controller might interact to place a VM into service within a VN.

Figure 2: VM Creation Workflow with OpenStack

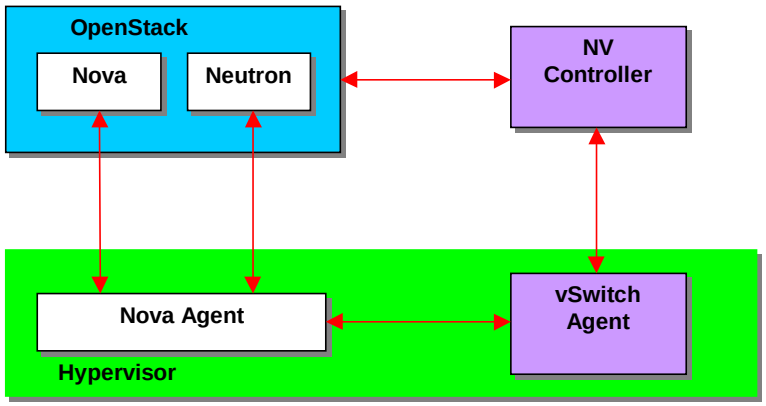

The Nova module in OpenStack instructs the Nova Agent in the hypervisor to create the VM. The Nova agent communicates with the Neutron module in OpenStack to learn the network attributes of the VM. The Nova agent then informs the vSwitch agent to configure the virtual network for the VM and then the controller provides the route table entries needed by the vSwitch.

Controller Based NV Solution Architecture

A Network Virtualization Solution typically has an architecture similar to the one shown in **Figure 3**. The main components are typically the NV Controller, hypervisor-resident vSwitches/vRouters, and gateways that provide connectivity from virtual networks to traditional network segments; e.g., VLANs, non-virtualized servers, or Internet routers. The controller function is generally supported by a high availability (HA) cluster or another HA configuration. Controller functionality may be comprised of a number of sub-functions running on different servers. Cloud Management/Orchestration is typically obtained from a third party and network services may be integrated with the controller, integrated via virtual appliances, or possibly integrated via physical appliances through the gateway.

Figure 3: Network Virtualization Architecture

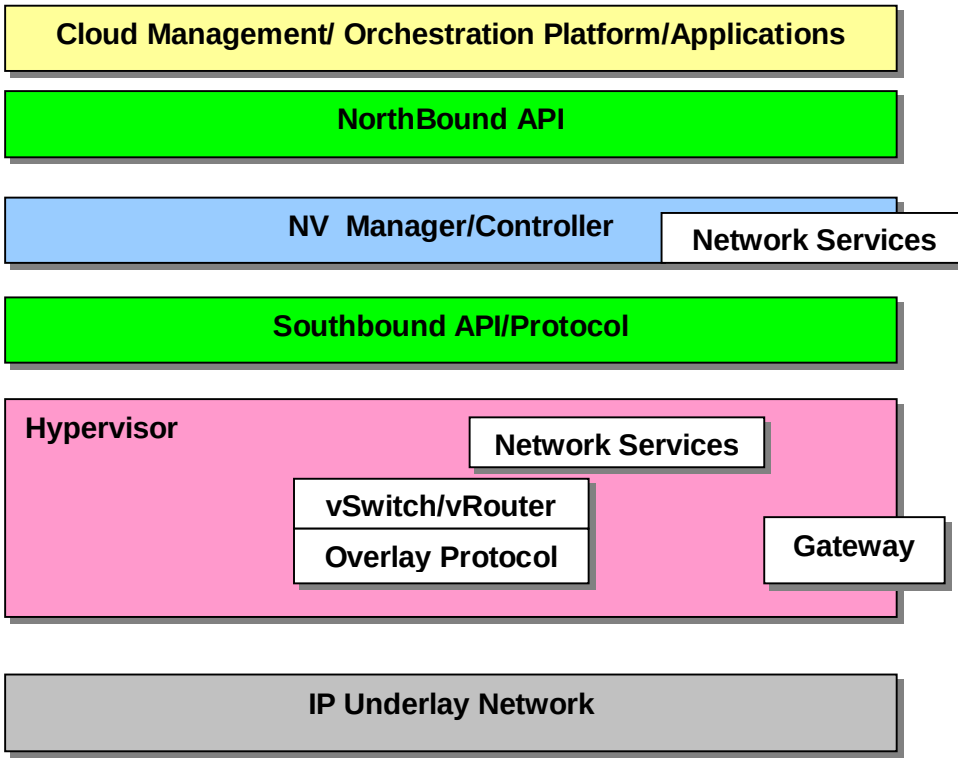

Criteria to Evaluate Overlay NV Solutions

One of the primary criterion that IT organizations should use relative to evaluating overlay network virtualization solutions is how well it solves the problem(s) that the IT organization is looking to solve. For example, can the solution enable the IT organization to move workloads between data centers? Between a data center and a branch office?

Other solution level criteria that IT organizations should evaluate include:

- Does the solution federate and hence interoperate with other solutions?
- What interaction, if any, is there between the virtual networks and the physical networks?
- What management functionality is provided into both the virtual and physical networks?
- Does the solution support service chaining?

The main technical differences between the various overlay NV solutions that IT organizations should evaluate fall into the following categories:

- **Encapsulation formats**. Some of the tunneling/encapsulation protocols that provide network virtualization of the data center include VXLAN, NVGRE, STT, and SPB MAC-in-MAC (SPBM). Both the IEEE and the IETF have already standardized SPB. It is unclear as to whether or not all of the other proposals will become standards.
- **Tunnel control plane functionality** that allows ingress (encapsulating) devices to map a frame to the appropriate egress (decapsulating) device. The first-hop overlay device implements a mapping operation that determines where the encapsulated packet should be sent to reach its intended destination VM. Specifically, the mapping function maps the destination address (either L2 or L3) of a packet received from a VM into the corresponding destination address of the egress NVE device. The main differences here are whether a controller is used and the functionality of the controller.

  Some of the initial, controller-less approaches to network virtualization relied on IP multicast as a way to disseminate address mappings. A more common solution is based on a central repository of address mappings housed in a controller. Vendors frequently refer to controller-based overlay NV solutions as SDN, while a more descriptive terminology might be Software Defined Overlay Network Virtualization.
- **vSwitches supported.** A number of vSwitches are based to some degree on the open source Open vSwitch (OVS)[1], while other vSwitches are of proprietary design. Another point of differentiation is whether the vSwitch is a virtual router as well as being an encapsulating Layer 2 switch. With Layer 3 functionality, a vSwitch can forward traffic between VMs on the same hypervisor that are in different subnets and can be used to implement Layer 3 VNs. Where the tunneling vSwitch has full Layer 3 functionality, the majority of intelligence can be implemented at the edge of network, allowing the underlay network to be implemented as a simple Layer 2 fabric.

[1] While based on OVS, many vSwitches have implemented proprietary extensions to OVS.

- **Broadcast/Multicast delivery** within a given virtual network. NVEs need a way to deliver multi-destination packets to other NVEs with destination VMs. There are three different approaches that can be taken:

    - The multicast capabilities of the underlay network can be used
    - The NVEs can replicate the packets and unicast a copy across the underlay network to each NVE currently participating in the VN.
    - The NVE can send the packet to a distribution server which replicates and unicasts the packets on the behalf of the NVEs.

- **Protocols.** Another characteristic of centralized controller solutions is the choice of Southbound protocols/APIs employed between the NV controller and the NVE and the choice of Northbound protocols/APIs used between the NV controller and cloud management systems and hypervisor management systems. If the southbound protocols are standardized, the NVE can potentially communicate with different types of NV controllers or controllers from different vendors. Some the alternatives here include OpenFlow, BGP, and CLI shell scripts.

  If the northbound protocols are standardized, the controller can be integrated with network services from ISVs or different types of third party orchestration systems. Most overlay NV controllers support a RESTful Web API for integration with cloud management and orchestration systems. With both southbound and northbound APIs the most important question becomes which third party switches, applications, virtual appliances, and orchestration systems have been certified and are supported by the overlay NV vendor.

- **VN Extension over the WAN.** VN extension over the WAN can generally be accomplished with most NV solutions. However, in some cases the encapsulation used over the wide area may differ from that used within the data center. Some of the encapsulation techniques used for VN extension over the WAN include MPLS VPNs and two proprietary protocols from Cisco: Overlay Transport Virtualization (OTV) and Locator/ID Separation Protocol (LISP). OTV is optimized for inter-data center VLAN extension over the WAN or Internet using MAC-in-IP encapsulation. It prevents flooding of unknown destinations across the WAN by advertising MAC address reachability using IS-IS routing protocol extensions. LISP is an encapsulating IP-in-IP technology that allows end systems to keep their IP address (ID) even as they move to a different subnet within the network (Location). By using LISP VM-Mobility, IP endpoints such as VMs can be relocated anywhere regardless of their IP addresses while maintaining direct path routing of client traffic. LISP also supports multi-tenant environments with Layer 3 virtual networks created by mapping VRFs to LISP instance-IDs. Inter-data center network virtualization could also potentially be based on Layer 3 vSwitches that support MPLS VPNS and implement network virtualization using RFC 4023 MPLS over IP/GRE tunnels through an IP enterprise network to connect to an MPLS VPN service. SPBM is unique in that it offers extensions over the WAN natively without requiring additional protocols such as OTV or MPLS VPNs.

The remainder of this sub-section of The Guide focuses on the primary differentiating features of Overlay NV solutions: tunnel encapsulation and tunnel control.

Tunnel Encapsulation

VXLAN: Virtual eXtensible LAN (VXLAN)[2] virtualizes the network by creating a Layer 2 overlay on a Layer 3 network via MAC-in-UDP encapsulation. The VXLAN segment is a Layer 3 construct that replaces the VLAN as the mechanism that segments the data center LAN for VMs. Therefore, a VM can only communicate or migrate within a VXLAN segment. The VXLAN segment has a 24-bit VXLAN Network identifier. VXLAN is transparent to the VM, which still communicates using MAC addresses. The VXLAN encapsulation is performed through a function known as the VXLAN Tunnel End Point (VTEP), typically a hypervisor vSwitch or a possibly a physical access switch. The encapsulation allows Layer 2 communications with any end points that are within the same VXLAN segment even if these end points are in a different IP subnet. This allows live migrations to transcend Layer 3 boundaries. Since MAC frames are encapsulated within IP packets, there is no need for the individual Layer 2 physical switches to learn MAC addresses. This alleviates MAC table hardware capacity issues on these switches. Overlapping IP and MAC addresses are handled by the VXLAN ID, which acts as a qualifier/identifier for the specific VXLAN segment within which those addresses are valid.

As noted, VXLANs use a MAC-in-UDP encapsulation. One of the reasons for this is that modern Layer 3 devices parse the 5-tuple (including Layer 4 source and destination ports). While VXLAN uses a well-known destination UDP port, 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 even in the case of multiple flows between just two VMs.

Where VXLAN nodes on a VXLAN overlay network need to communicate with nodes on a legacy (i.e., VLAN) portion of the network, a VXLAN gateway can be used to perform the required tunnel termination functions including encapsulation/decapsulation. The gateway functionality could be implemented in either hardware or software.

VXLAN is supported by a number of vendors including Cisco Systems, VMware, IBM, and Nuage Networks. Avaya's SPBM implementation (Fabric Connect) can also support a VXLAN deployment, acting as a transport layer providing optimized IP Routing and Multicast for VXLAN-attached services.

STT: Stateless Transport Tunneling (STT)[3] is a second overlay technology for creating Layer 2 virtual networks over a Layer 2/3 physical network within the data center. Conceptually, there are a number of similarities between VXLAN and STT. The tunnel endpoints are typically provided by hypervisor vSwitches, the VNID is 24 bits wide, and the transport source header is manipulated to take advantage of multipathing. STT encapsulation differs from VXLAN in two ways. First, it uses a 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 TOE server NICs. The benefits to the host include lower CPU utilization and higher utilization of 10 Gigabit Ethernet access links. STT generates a source port number based on hashing the header fields of the inner packet to ensure efficient load balancing over LAGs and multi-pathing fabrics. STT also allocates more header space to the per-packet metadata, which provides added flexibility for the virtual network tunnel control plane. With these features, STT is

[2] http://searchservervirtualization.techtarget.com/news/2240074318/VMware-Cisco-propose-VXLAN-for-VM-mobility

[3] http://tools.ietf.org/html/draft-davie-stt-01

optimized for hypervisor vSwitches as the encapsulation/decapsulation tunnel endpoints. The initial implementations of Network Virtualization using STT from Nicira Networks are based on OpenFlow-like hypervisor vSwitches (Open vSwitches) and a centralized control plane for tunnel management via downloading mapping tables to the vSwitches.

NVGRE: Network Virtualization using Generic Router Encapsulation (NVGRE)[4] uses the GRE tunneling protocol defined by RFC 2784 and RFC 2890. NVGRE is similar in most respects to VXLAN with two major exceptions. While GRE encapsulation is not new, most network devices do not parse GRE headers in hardware, which may lead to performance issues and issues with 5-tuple hashes for traffic distribution in multi-path data center LANs. With GRE hashing generally involves the GRE key. One initial implementation of NVGRE from Microsoft relies on Layer 3 vSwitches whose mapping tables and routing tables are downloaded from the vSwitch manager. Downloads are performed via a command-line shell and associated scripting language.

SPBM[5]: IEEE 802.1aq/IETF 6329 Shortest Path Bridging MAC-in-MAC uses IEEE 802.1ah MAC-in-MAC encapsulation and the IS-IS routing protocol to provide Layer 2 network virtualization and VLAN extension in addition to a loop-free equal cost multi-path Layer 2 forwarding functionality. VLAN extension is enabled by the 24-bit Service IDs (I-SIDs) that are part of the outer MAC encapsulation. Unlike other network virtualization solutions, no changes are required in the hypervisor vSwitches or NICs and switching hardware already exists that supports IEEE 802.1ah MAC-in-MAC encapsulation. For SPBM, the control plane is provided by the IS-IS routing protocol.

SPBM can also be extended to support Layer 3 forwarding and Layer 3 virtualization as described in the IP/SPB IETF draft using IP encapsulated in the outer SPBM header. This specification identifies how SPBM nodes can perform Inter-ISID or inter-VLAN routing. IP/SPB also provides for Layer 3 VSNs by extending VRF instances at the edge of the network across the SPBM network without requiring that the core switches also support VRF instances. VLAN-extensions and VRF-extensions can run in parallel on the same SPB network to provide isolation of both Layer 2 and Layer 3 traffic for multi-tenant environments. With SPBM, only those Switches that define the SPBM boundary need to be SPBM-capable. Switches not directly involved in mapping services to SPB service IDs don't require special hardware or software capabilities. SPBM isn't based on special vSwitches, data/control plane separation, or centralized controllers. SPBM hardware Switches are currently available from several vendors, including Avaya and Alcatel-Lucent.

Tunnel Control

As previously mentioned, initial implementations of VXLAN by Cisco and VMware use flooding as a distributed control solution based on Any Source Multicast (ASM) to disseminate end system location information. Because flooding requires processing by all the vSwitches in the multicast group, this type of control solution will not scale to support very large networks.
.
A more recent approach is to implement tunnel control as a centralized controller function. A control plane protocol that carries both MAC and IP addresses can eliminate the need for ARP. One controller-based solution for VXLAN control, championed by IBM's Distributed Overlay

[4] http://datatracker.ietf.org/doc/draft-sridharan-virtualization-nvgre/
[5] http://tools.ietf.org/html/draft-allan-l2vpn-spbm-evpn-00

Virtual Ethernet (DOVE) initiative, is to use a DNS-like network service to map the VM's IP address to the egress VTEP's IP address. IBM's solution does not require Multi Cast enablement in the physical network. IBM's Controller based solution has built-in IP routing capability.

In another controller-based approach, used by Nicira Networks, the controller maintains a data base of Open vSwitches (OVS) in the network and proactively updates OVS mapping tables via OpenFlow to create new tunnels when VMs are created or moved. The Nicira controller focuses on the virtual network topology and is oblivious to the topology of the core physical network. The controller is integrated with hypervisor and cloud management systems to learn of changes in the population of VMs.

A third controller approach, used by Nuage Networks and Netsocket, involves the controller maintaining a full topology of the virtual and physical network and maintaining the full address mapping and routing tables derived from standard routing protocols, such as OSPF, IS-IS, or BGP. The portion of the table needed by the vSwitch is disseminated from the controller to the vSwitches via the OpenFlow protocol. The Nuage Networks' vSwitches use VXLAN to encapsulate L2 traffic and GRE to encapsulate L3 traffic.

Comparison of Network Overlay Virtualization Solutions

The following table (**Table 7**) provides a high level summary of the primary features of some of the Network Virtualization solutions that are available or have been recently announced. Note that the solutions described in columns two and three (Cisco, VMware) are not based on a controller.

Table 7: Network Overlay Virtualization features

	Cisco	VMware	IBM	VMware/ Nicira	Nuage Networks	Avaya	Netsocket	Juniper
Product	**Nexus 1000v**	**VSphere DS**	**SDN-VE**	**NSX**	**VSP**	**Fabric Connect**	**NVN**	**Contrail**
Overlay	VXLAN	VXLAN	VXLAN	VXLAN STT?	VXLAN	SPBM	GRE	MPLS/GRE MPLS/UDP VXLAN
VM-NVE Address Learning	VTEP Multicast flooding	VTEP Multicast flooding	Pull From Controller's Directory	Push From Controller's Data Base	Push From Controller's Map Table	IS-IS SPB on physical switch	Push From Controller's Map Table	Push From Controller's Map Table
Broadcast / Multicast within VN	via underlay Multicast	via underlay Multicast	distribution server replication	distribution server replication	dVRS packet replication	via SPB multicast		VRouter packet replication or proxy
Controller Topology Awareness	na	na	Virtual Networks	Virtual Networks	Entire Network	Entire Network	Entire Network	Entire Network
Controller to NVE Protocol	NX-OS CLI	VMware API	Open source submitted to OpenDaylight	OpenFlow NSX API	OpenFlow	IS-IS	vFlow or OpenFlow	XMPP
vSwitch	Nexus 1000v	VDS	SDN-VE vSwitch	VDS, Open vSwitch**	dVRS (Open vSwitch**)	Native to Hypervisor	vFlowSwitch	v Contrail vRouter
vSwitch L3	no	no	yes	yes	yes	na	yes	yes
Gateway Support in Physical Switches				Arista 7150s Brocade ADX	Nuage Networks 7850 VSG	na		
Hypervisors	ESXi, Hyper-V, XEN, KVM	ESXi	ESXi KVM	vSphere. ESXi, XEN, KVM	ESXi, Hyper-V, XEN, KVM	ESXi, Hyper-V, XEN, KVM	Hyper-V ESXi Xen, KVM	KVM, XEN
Controller Federation					via MP-BGP			BGP
DC-DC encapsulation	OTV	OTV	VXLAN	GRE	MPLS over GRE to PE router	Over an SPBM WAN	GRE	MPLS/GRE
	OpenStack vCloud	OpenStack vCloud	OpenStack	OpenStack CloudStack vCloud	OpenStack CloudStack vCloud	OpenStack Integration in controller	OpenStack System Ctr.	OpenStack.

na = not applicable ***** = with proprietary extensions***

Software Defined NV via Flow Table Segmentation

Network virtualization can also be implemented as an application that runs on an SDN controller. Virtual networks are defined by policies that map flows to the appropriate virtual network based on L1-L4 portions of the header. With this type of SDN-based NV, there is no need for tunnels and encapsulation protocols. One example of an NV application is the Big Virtual Switch that runs on the Big Network Controller from Big Switch Networks. The Big Network Controller implements VNs by configuring forwarding tables in OpenFlow physical and virtual switches. The OpenFlow switches can be from a variety of traditional switch vendors. Another alternative is to use Big Switch Switch Light OpenFlow thin software agent running on bare metal Ethernet switches based on Broadcom merchant silicon or on virtual switches.

By exploiting the capability of OpenFlow to deal with encapsulation and de-encapsulation, the SDN controller NV application can also be used to implement overlay VNs running over a conventional L2/L3 network, or a hybrid network based partially on pure SDN VNs and partially on SDN NVs with OpenFLow virtual switches and a conventional core network.

Another slightly different approach to an NV application for SDN controllers is the Virtual Tenant Network (VTN) application developed by NEC and recently accepted as an application by the OpenDaylight consortium. The VTN solution provides a layer of abstraction between the virtual network and the physical network. In the event of a failed link, the VTN can detect and redirect the affected flows within milliseconds. This avoids the re-convergence delay associated with traditional network protocols. The VTN also supports redirection, which enables use cases related to traffic steering and service chaining. In addition, the VTN physical control of the network supports flow based traffic engineering as well as 8-way ECMP.

VTN is based on a logical abstraction that decouples the VTN from the physical network. A virtual network can be designed and deployed using the following set of logical network elements:

- vBridge L2 switch function.
- vRouter router function.
- vTEP virtual Tunnel End Point.
- vTunnel Tunnel.
- vBypass connectivity between controlled networks.
- vInterface end point on the virtual node.
- vLink L1 connectivity between virtual interfaces.

Using these elements allows the user can define a logical network with the look and feel of conventional L2/L3 network. VTN can also by used to implement an overlay network, an OpenFlow network, or a hybrid overlay/OpenFlow network. Once the network is designed on VTN, it can automatically be mapped onto the underlying physical network, and configured on the individual switches leveraging an SDN control protocol, Typically this would be OpenFlow. Mapping is used to identify the VTN to which each packet transmitted or received by an OpenFlow switch belongs, as well as which interfaces on the OpenFlow switch can transmit or receive that packet. Flows are mapped to a VTN vBridge based on the ingress port on the OpenFlow switch, the source MAC address or the VLAN ID.

Enterprise Plans for NV Adoption

The Survey Respondents were asked a series of questions about their current position relative to evaluating and adopting NV solutions and how that position might change over the next two to three years. In the first of those questions, The Survey Respondents were given a set of alternatives and were asked to indicate the alternatives that described their company's current approach to implementing NV solutions. Their responses are shown in **Table 8**.

Table 8: Current Approaches to Adopting NV Solutions

Approach to Adoption NV Solutions	% of Respondents
We have not made any analysis of NV	25.5%
We will likely analyze NV sometime in the next year	25.5%
We are currently actively analyzing the potential value that NV offers	24.7%
We expect that within a year that we will be running NV either in a lab or in a limited trial	13.6%
We are currently actively analyzing vendors' NV strategies and offerings	11.5%
We currently are running NV either in a lab or in a limited trial	9.9%
We currently are running NV somewhere in our production network	7.4%
We looked at NV and decided to not do anything with NV over the next year	6.2%
We expect that within a year that we will be running NV somewhere in our production network	5.8%
Don't know	4.9%

The data in **Table 8** indicates that while there is currently little deployment of NV, there is a lot of activity and interest relative to analyzing NV solutions. The data in Table 8 also suggests that over the next year the percentage of IT organizations that are either running NV somewhere in their production network, or in a lab or limited trial, will double.

The Survey Respondents were given a two-year time frame and were asked to indicate where in their infrastructure their organization was likely to implement NV solutions. (Multiple responses were allowed) Their responses are shown in **Table 9**.

Table 9: Likely Deployment of NV Solutions

Focus of Future NV Implementation	% of Respondents
Data Center	58.0%
Branch and/or Campus	25.1%
WAN	18.5%
We are unlikely to implement NV in the next two years	15.6%
Don't know	10.7%
We are likely to acquire a WAN service that is based on NV	9.5%

The data in **Table 9** indicates that IT organizations will primarily implement NV solutions within a data center. However, the data also indicates that a sizeable percentage of IT organizations want to extend their NV solutions over the WAN and to also implement NV solutions in their branch and campus networks.

In the final question about their potential future use of NV solutions, The Survey Respondents were asked to indicate how broadly their data center networks will be based on NV three years from now. Their responses are shown in **Table 10**.

Table 10: Data Center Design in Three Years

Balance of NV and Traditional Approach	% of Respondents
Exclusively based on NV	3.3%
Mostly based on NV	25.1%
NV and traditional networking coexisting about equally	37.9%
Mostly traditional	16.9%
Exclusively traditional	4.1%
Don't know	12.8%

The data in **Table 10** indicates that the vast majority of The Survey Respondents expect that in three years that at least half of their data center networks will be based on NV.

Chapter 2: The what, why and how of SDN

Background

In the traditional approach to networking, most network functionality is implemented in a dedicated appliance; i.e., switch, router, application delivery controller. In addition, within the dedicated appliance, most of the functionality is implemented in dedicated hardware such as an ASIC (Application Specific Integrated Circuit).

Some of the key characteristics of this approach to developing network appliances are:

- The ASICs that provide the network functionality evolve slowly;
- The evolution of ASIC functionality is under the control of the provider of the appliance;
- The appliances are proprietary;
- Each appliance is configured individually;
- Tasks such as provisioning, change management and de-provisioning are very time consuming and error prone.

Networking organizations are under increasing pressure to be more efficient and agile. One source of that pressure results from the widespread adoption of server virtualization. As part of server virtualization, virtual machines (VMs) are dynamically moved between servers in a matter of seconds or minutes. However, if the movement of a VM crosses a Layer 3 boundary, it can take days or weeks to reconfigure the network to support the VM in its new location. It can sometimes be difficult to define exactly what it means for a network to be agile. That said, if it takes weeks to reconfigure the network to support the movement of a VM, that network isn't agile.

The bottom line is that a traditional network evolves slowly; is limited in functionality by what is provided by the vendors of the network appliances; has a relatively high level of OPEX and is relatively static in nature. The majority of the potential SDN use cases (see below) are intended to overcome those characteristics of traditional networks.

Potential SDN Use Cases

There is scene in the novel *Alice in Wonderland* that is directly relevant to the adoption of NV and SDN solutions. That scene is comprised of the following dialogue between Alice and the Cheshire cat.

Alice: "Would you tell me, please, which way I ought to go from here?"

Cheshire Cat: 'That depends a good deal on where you want to get to."

Alice: "I don't much care where."

Cheshire Cat: "Then it doesn't matter which way you go."

The relevance of that dialogue to SDN is that an analysis of SDN solution architectures and subtending protocols is totally irrelevant until IT organizations identify which use cases they are hoping to address with SDN.

The left hand column of **Table 11** contains some of the primary challenges & opportunities facing the typical IT organization. The Survey Respondents were shown those challenges & opportunities and were asked to indicate which of them they thought that SDN could help them to respond to and they were allowed to check all that applied. Each row of the right hand column of **Table 11** contains the percentage of The Survey Respondents that indicated that they thought that SDN could help them to respond to the challenge or opportunity in the corresponding left hand column.

Table 11: Opportunities & Challenges that SDN Can Address	
Challenge or Opportunity	**Percentage**
Better utilize network resources	51%
Ease the administrative burden of configuration and provisioning QoS and Security	47%
Perform traffic engineering with an end-to-end view of the network	44%
More easily scale network functionality	39%
Support the dynamic movement, replication and allocation of virtual resources	38%
Establish virtual Ethernet networks without the limitations and configuration burden of VLANs	35%
Reduce Complexity	34%
Enable applications to dynamically request services from the network	32%
Reduce OPEX	30%
Have network functionality evolve more rapidly based on a software development lifecycle	27%
More easily implement QoS	27%
Implement more effective security functionality	26%
Reduce CAPEX	25%
We don't see any challenges or opportunities that SDN can help us with	3%
Don't know	3%
Other	3%

One observation that can be drawn from the data in **Table 11** is that there is a wide range of challenges and opportunities that The Survey Respondents believe that SDN can help with and conversely very few IT organizations believe that SDN won't be beneficial. Having a wide range of potential challenges and opportunities to respond to bodes well for the long-term adoption of SDN. However, having so many challenges and opportunities to respond to can create confusion in the short term and can possibly delay SDN adoption.

To exemplify the relationship between the opportunities & challenges and the two types of solutions analyzed in The Guide (i.e., NV and SDN), assume that the opportunity that a hypothetical IT organization is attempting to respond to is the need to support the dynamic movement, replication and allocation of virtual workloads. The hypothetical IT organization can respond to this challenge using any of the NV solutions that were discussed in the preceding chapter; e.g., solutions from Nuage Networks, Netsocket, Avaya and NEC. As a reminder to the reader, the NV solutions from Nuage Networks, Netsocket and Avaya are based on overlay technologies and the NV solution from NEC is based on manipulating the flow tables in NEC's SDN solution.

The situation is quite different if the opportunity that the hypothetical IT organization is trying to respond to is the need to make it easier to implement QoS or the need to enable applications to dynamically request services from the network. The hypothetical IT organization can potentially

respond to both of these challenges by implementing an SDN solution whereas that organization couldn't respond to those challenges by just implementing one of the controller based NV solutions that were discussed in the preceding chapter. As will be pointed out in the following discussion of a federated overlay/underlay model, it would potentially be possible for the hypothetical IT organization to respond to those challenges using a federation of NV overlay solutions and SDN solutions.

The challenges and opportunities that are identified in **Table 11** aren't dependent on any particular technology. For example, there are a number of technologies that can be implemented in order to ease the burden of configuration management. That said, a subsequent sub-section of this document identifies some of the specific use cases and benefits that are associated with the OpenFlow protocol.

While the use of SDN in data centers receives the majority of attention, it is also possible to implement SDN in branch and campus networks as well as in wide area networks (WANs). In order to understand where SDN will likely be implemented, The Survey Respondents were asked "If your organization is likely to implement SDN sometime over the next two years, where are you likely to implement it?" Their responses are summarized in **Table 12**.

Table 12: Focus of SDN Deployment

Focus of SDN Deployment	Percentage
Data Center	54%
Branch and/or Campus	26%
WAN	23%
We are likely to implement a service from a WAN service provider that is based on SDN	12%
We are unlikely to implement SDN within the next two years	11%
Don't know	11%
Other	7%

One observation that can be made from the data in **Table 12** is that while the primary interest in deploying SDN is focused on the data center, there is strong interest in deploying SDN broadly throughout an organization's entire network.

A Working Definition of SDN

Within the IT industry, there is not a universally agreed to definition of SDN. While **The Guide** will identify the primary characteristics of an SDN, it won't make any attempt to define SDN. It is, however, helpful to have a working definition of SDN. The working definition of SDN that will be used in this publication is the one created by the Open Networking Foundation (ONF).

The ONF is the group that is most associated with the development and standardization of SDN. According to the ONF[6], "Software-Defined Networking (SDN) is an emerging architecture that is dynamic, manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. This architecture decouples the network control and forwarding functions enabling the network control to become directly programmable and the underlying infrastructure to be abstracted for applications and network services. The OpenFlow™ protocol is a foundational element for building SDN solutions."

According to the ONF, the SDN architecture is:

- **Directly programmable**: Network control is directly programmable because it is decoupled from forwarding functions.
- **Agile**: Abstracting control from forwarding lets administrators dynamically adjust network-wide traffic flow to meet changing needs.
- **Centrally managed**: Network intelligence is (logically) centralized in software-based SDN controllers that maintain a global view of the network, which appears to applications and policy engines as a single, logical switch.
- **Programmatically configured**: SDN lets network managers configure, manage, secure, and optimize network resources very quickly via dynamic, automated SDN programs, which they can write themselves because the programs do not depend on proprietary software.
- **Open standards-based and vendor-neutral**: When implemented through open standards, SDN simplifies network design and operation because instructions are provided by SDN controllers instead of multiple, vendor-specific devices and protocols.

Part of the confusion that surrounds SDN is that many vendors don't buy in totally to the ONF definition of SDN. For example, while the vast majority of vendors do include the centralization of control in their definition of SDN, there isn't agreement as to how much control should be centralized. In addition, while some vendors are viewing OpenFlow as a foundational element of their SDN solutions, other vendors are taking a wait and see approach to OpenFlow.

Another source of confusion is the relationship between NV and SDN. It's possible to implement an SDN that resembles the ONF definition of SDN and use that SDN to implement network virtualization. For example, the OpenDayLight foundation recently accepted a contribution from NEC, referred to as Virtual Tenant Networking (VTN), which enables an SDN to implement network virtualization by manipulating the flow tables that are associated with the OpenFlow protocol. It is also possible, however, to implement network virtualization without

[6] https://www.opennetworking.org/sdn-resources/sdn-definition

implementing an SDN as defined by the ONF. For example, as described in the previous section of The Guide, Avaya offers an NV solution that doesn't rely on a controller. In addition, both Nuage Networks and VMware/Nicira implement network virtualization using an overlay model and a controller. To add to the confusion, Nuage Networks refers to their solution as SDN while VMware is adamant that their solution is network virtualization and not SDN.

The Survey Respondents were given a set of characteristics that are often associated with SDN and were asked to indicate which two characteristics would provide the most value to their company's network. Their responses are shown in **Table 13**.

Table 13: Value of SDN Characteristics

Characteristic	Percentage
Centralization of configuration and policy management	45%
Programmability of network elements	31%
Automation of administrative tasks	28%
Centralization of control	28%
The development of network functionality on a software development cycle vs. a hardware cycle	27%
Open up the network to innovation by the entire ISV community	17%
The use of open protocols	10%
The use of open source solutions	8%
Other	2%
Don't Know	1%

One observation that can be drawn from **Table 13** is that the characteristic of SDN that offers the most value to The Survey Respondents is tactical: The centralization of configuration and policy management. However, the second most important characteristic, the programmability of network elements, is strategic. That characteristic is strategic because the programmability of network elements is a key component of the overall functionality that is required in order to enable applications to dynamically request the network services they need.

Another observation that can be drawn from **Table 13** is that in spite of all of the discussion in the industry about open networking, The Survey Respondents were not very enthusiastic about the value that open protocols would bring to their networks.

The SDN Solution Architecture

Figure 4 contains a graphical representation of the SDN architecture as envisioned by the ONF. One key component of a complete SDN solution that is missing from **Figure 4** is cloud orchestration platforms such as OpenStack. The role that these platforms play in both NV and SDN solutions was described in the preceding section of The Guide.

Figure 4: ONF's SDN Architecture

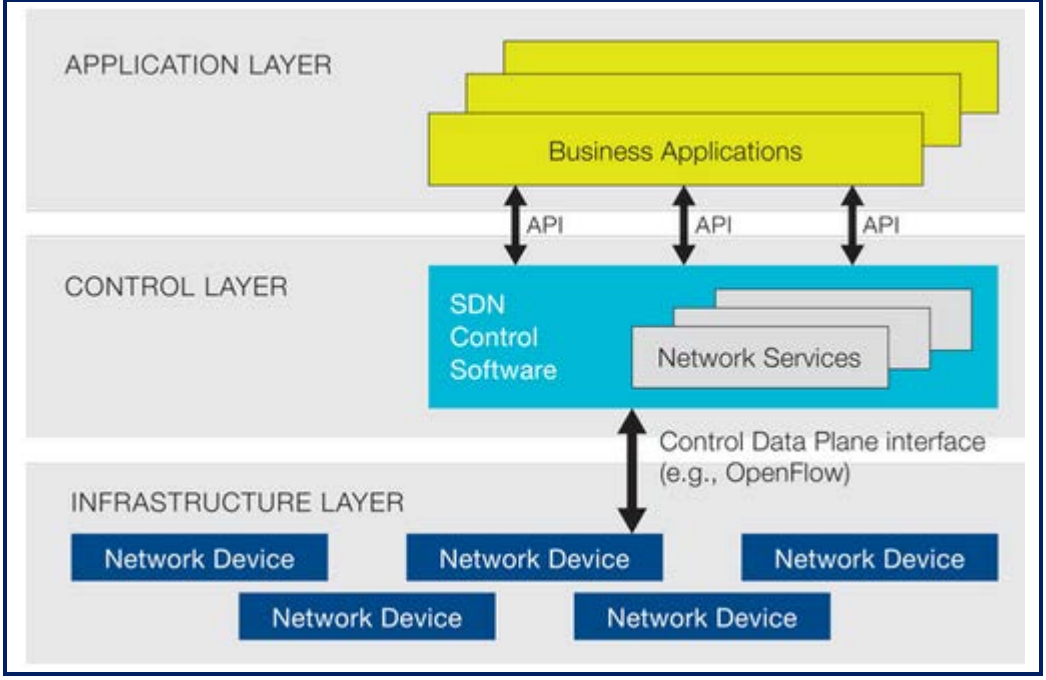

Below are definitions of some terms that are commonly associated with SDN, some of which appear in **Figure 4**.

- **Business Applications**
  This refers to applications that are directly consumable by end users. Possibilities include video conferencing, supply chain management and customer relationship management.

- **Network Services**
  This refers to functionality that enables business applications to perform efficiently and securely. Possibilities include a wide range of L4 – L7 functionality including load balancing and security capabilities such as firewalls, IDS/IPS and DDoS protection.

- **Open Protocol**
  An open protocol is a protocol whose specification a company, or group of companies, has made public.

- **Standards Based Protocol**
  A standards based protocol is an open protocol that was created by a recognized standards body such as the ONF, the IEEE or the IETF.

- **Pure SDN Switch**
  In a pure SDN switch, all of the control functions of a traditional switch (i.e., routing protocols that are used to build forwarding information bases) are run in the central controller. The functionality in the switch is restricted entirely to the data plane.

- **Hybrid Switch**
  In a hybrid switch, SDN technologies and traditional switching protocols run simultaneously on a given switch. A network manager can configure the SDN controller to discover and control certain traffic flows while traditional, distributed networking protocols continue to direct the rest of the traffic on the network.

- **Hybrid Network**
  A hybrid network is a network in which traditional switches and SDN switches, whether they are pure SDN switches or hybrid switches, operate in the same environment.

- **Southbound API**
  Relative to **Figure 4**, the southbound API is the API that enables communications between the control layer and the infrastructure layer.

- **Service Chaining**[7]
  Service chaining is the ability to steer VM-VM traffic flows through a sequence of physical or virtual servers that provide network services, such as firewalls, IPS/IDS, DPI, or load balancers.

Figure 4 shows an API between the SDN control layer and the business applications. This API is commonly referred to as *the Northbound API*. The role of the northbound API is to enable communications between the control layer and the application layer. Currently there isn't a standard for the Northbound API, although the ONF has recently begun a process that could lead to a standards based API. While it isn't possible to state how the development of the northbound API will evolve, it is likely that there won't be a single northbound API, but multiple northbound APIs. One viable alternative is that there will be a northbound API between the SDN control software and each of the following entities:

- Network services
- Business applications
- Cloud management/orchestration systems

7 Service chaining was described in greater detail in the preceding section of The Guide.

Criteria to Evaluate SDN Solution Architectures

Below is a set of 7 questions that IT organizations should ask vendors who provide all or the majority of the SDN solution architecture that is shown in **Figure 4**. These questions focus on key criteria that IT organizations should use relative to evaluating alternative SDN solutions. A more complete set of criteria can be found in *A Mock RFI for SDN Solutions*[8].

As highlighted in the preceding discussion of Alice in Wonderland, SDN solutions need to be evaluated relative to their ability to respond to the specific challenges and opportunities facing an IT organization. For the sake of example, assume that one of the opportunities that an IT organization is hoping to respond to is enabling applications to dynamically request services from the network. Given that, then one question that the IT organization should ask vendors of SDN solutions is:

1. How does your SDN solution enable applications to dynamically request services from the network?

Other questions that IT organizations should ask SDN solution vendors include:

2. Describe the SDN solution that you are proposing and include in that description how the SDN architecture for the solution you are proposing is similar to the architecture shown in **Figure 4** and also describe how it is different. In your answer, identify the southbound protocols that you support and provide the rationale for supporting those protocols.

3. Identify the aspects of your solution architecture that enable high availability; that enable scalability of performance; that enable extensibility of functionality.

4. Which components of the solution architecture do you provide yourself? Which components do partners provide? If the solutions you are proposing includes components from partners, is there a single point of accountability for the solutions?

5. In your SDN solution, what control functions reside in the control layer and which control functions reside in the infrastructure layer?

6. Describe the Northbound protocol(s)/API(s) you support between the control layer and:

    - Network services
    - Enterprise applications
    - Cloud management/orchestration systems

7. How does your proposed solution implement network virtualization? Include in your answer whether overlays are used; what protocols are supported; how the tunneling control function is implemented. If virtual networks are defined by flow partitioning, describe which header fields are used and how the partitioning is accomplished.

8 Will be published at: webtorials.com/Metzler

The Inhibitors to SDN Adoption

The left hand column of **Table 14** contains some of the primary impediments to the adoption of SDN. The Survey Respondents were shown these impediments and were asked to indicate the two impediments that would be the biggest inhibitors to their company adopting SDN sometime in the next two years. Each row of the right hand column of **Table 14** contains the percentage of The Survey Respondents that indicated that the impediment in the corresponding left column was one of the two primary inhibitors.

Table 14: Inhibitors to the Adoption of SDN

Impediment	Percentage
The immaturity of the current products	30%
The immaturity of the enabling technologies	29%
Other technology and/or business priorities	24%
The confusion and lack of definition in terms of vendors' strategies	22%
The lack of resources to evaluate SDN	21%
The lack of a critical mass of organizations that have deployed SDN	14%
Concerns that the technology will not scale to support enterprise sized networks	12%
We don't see a compelling value proposition	7%
Concern that this is just a passing fad	7%
Other	5%
The confusion around the impact of consortiums such as OpenDayLight	4%
We don't see any inhibitors to implementing SDN	3%
Don't know	3%

One clear observation that can be drawn from **Table 14** is that immaturity, broadly defined, is the primary inhibitor to the adoption of SDN. That includes the immaturity of the current products, the immaturity of the enabling technologies and the confusion and lack of definition in terms of vendor strategies.

The role that a compelling business case plays relative to driving and inhibiting the adoption of SDN is somewhat subtle. As shown in **Table 14**, only 7% of *The Survey Respondents* indicated that the lack of a compelling value proposition was an inhibitor to their adoption of SDN. It would be easy to conclude from that metric that business cases that demonstrate the compelling value of SDN exist and that these business cases are widely understood. Drawing the conclusion would be a mistake.

Arguing against that conclusion is the fact that 24% of *The Survey Respondents* indicated that "other technology and/or business priorities" was an inhibitor and that 21% of *The Survey Respondents* indicated that "the lack of resources to evaluate SDN" was an inhibitor. If indeed,

there were compelling, well-understood SDN business cases, these organizations would rearrange their priorities and find the resources to evaluate SDN solutions.

The Overlay/Underlay Model

The preceding chapter of The Guide discussed ways to implement multiple virtual network topologies overlaid on a common physical network; a.k.a., an overlay model. That chapter also discussed some of the benefits and limitations of an overlay model. Some of those limitations were:

- 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 NV requires a lightly oversubscribed or non-oversubscribed physical underlay network.
- Some value-added features in existing networks cannot be leveraged due to encapsulation. For example, the physical network loses its ability to provide differentiated services based on the content of the packet header.

An emerging approach to overcome the limitations of the overlay model is referred to as an overlay/underlay model. The cornerstone of this approach is a federation between the overlay network virtualization controller and the underlay SDN controller. In August 2013, HP and VMware announced their intention to work together to create an overlay/underlay solution[9]. As part of that announcement, HP stated their intention to develop a new application called ConvergedControl that will enable HP's Intelligent Management Center (IMC) to share information about the network with both the HP and the VMware controllers. As part of the announced solution, VMware's NSX controller will continue to provision the virtual network overlay and HP's SDN controller will continue to provision physical network flows on its switches. The solution is intended to enable the two controllers to work together to ensure that the virtual network gets the physical flows it needs. The solution is also intended to provide visibility across the virtual and physical environment so that, for example, if there is congestion or a failure on the physical network, the virtual environment is aware of the issue and can respond accordingly.

Network Function Virtualization

A concept that often gets discussed in conjunction with SDN is Network Function Virtualization (NFV). Strictly speaking, NFV is being driven primarily by telecommunications service providers to meet their specific requirements. Their interest in NFV stems from the fact that in the current environment, telecommunications and networking software is being run on three types of platforms:

- Industry standard servers running Linux or Windows;
- Virtual appliances running over hypervisors on industry standard hardware servers;
- Proprietary hardware appliances.

9 http://searchsdn.techtarget.com/news/2240204281/HP-and-VMware-NSX-Joint-management-for-virtual-and-physical-networks

Telecommunications service providers feel that they can greatly simplify their operations and reduce capital expense if all network functions were available as virtual appliances that can be easily provisioned and integrated regardless of the vendor who provided the appliance or the hypervisor(s) on which it runs.

In order to bring this vision to fruition, an Industry Specifications Group for Network Functions Virtualization (NFV ISG) has been formed under the auspices of the European Telecommunications Standards Institute (ETSI). Their vision for the transition from hardware appliances of today to a fully virtualized appliance environment is depicted in **Figure 5**.

Figure 5: The Virtualization of Network Appliances

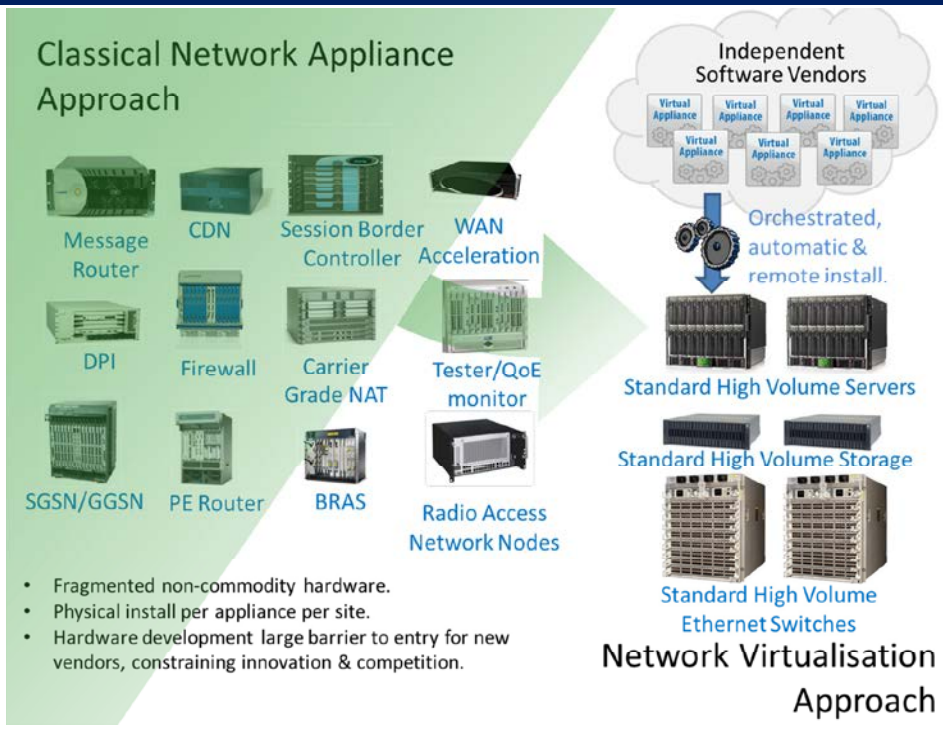

Source: NFV ISG

The approach that the NFV ISG is taking is that the virtualization of network functionality is applicable to any data plane packet processing and control plane function in both fixed and mobile networks. As shown in **Figure 5**, examples of these functions include:

- Switching elements;
- Tunneling gateway elements: IPSec/SSL VPN gateways;
- Traffic analysis: DPI, QoE measurement;
- Service Assurance, SLA monitoring, Test and Diagnostics;
- Application-level optimization: ADCs, WOCs;
- Security functions: Firewalls, virus scanners, intrusion detection systems;
- Multi-function home routers and set top boxes;
- Mobile network nodes.

The initial members of the NFV ISG were service providers such as AT&T, Deutsche Telekom and NTT. Its membership[10] has since grown and now includes a number of equipment vendors, but currently relatively few of the top vendors of virtual appliances are members.

The first meeting of the group was held in January 2013 and a number of smaller working groups were created in April 2013. In October 2013, ETSI published the first five specifications relative to NFV[11]. According to ETSI[12], "The five published documents include four ETSI Group Specifications (GSs) designed to align understanding about NFV across the industry. They cover NFV use cases, requirements, the architectural framework, and terminology. The fifth GS defines a framework for co-ordinating and promoting public demonstrations of Proof of Concept (PoC) platforms illustrating key aspects of NFV. Its objective is to encourage the development of an open ecosystem by integrating components from different players."

While the development of SDN and the development of NFV can proceed independently, there are some areas of possible overlap and cooperation. For example, one of the primary challenges the NFV group is facing is that the Operational Support Systems/Business Support Systems (OSS/BSS) that telecommunications service providers use must be able to automate the orchestration and provisioning of NFV appliances. While the NFV group believes its goals can be achieved using non-SDN mechanisms, the group is looking closely to see if standards coming from SDN consortia, such as the ONF and the OpenDaylight consortium, apply to NFV. As such, one possibility is that standards coming from the development of NV and SDN may facilitate the development of NFV. Alternatively, the development of NFV may result in technologies that facilitate the provisioning of virtual appliances in a NV or SDN solution.

[10] http://portal.etsi.org/NFV/NFV_List_members.asp
[11] http://www.etsi.org/technologies-clusters/technologies/nfv
[12] http://www.etsi.org/index.php/news-events/news/700-2013-10-etsi-publishes-first-nfv-specifications

The Open Networking Foundation and OpenFlow

The Open Networking Foundation

The Open Networking Foundation was launched in 2011 and its vision is to make OpenFlow-based SDN the new norm for networks. To help achieve that vision, the ONF has taken on the responsibility of driving the standardization of the OpenFlow protocol. Unlike most IT standards groups or industry consortiums, the ONF was not by founded by suppliers of the underlying technologies, but by Deutsche Telekom, Facebook, Google, Microsoft, Verizon, and Yahoo! As such, the ONF is one of the very few IT standards groups or industry consortiums that was launched by potential users of the technologies on which the consortium focused.

Figure 6 shows the ONF organizational model. More information on the ONF working and study groups as well as the activities that the ONF is sponsoring can be found at the ONF web site[13].

Figure 6: The ONF Organizational Model

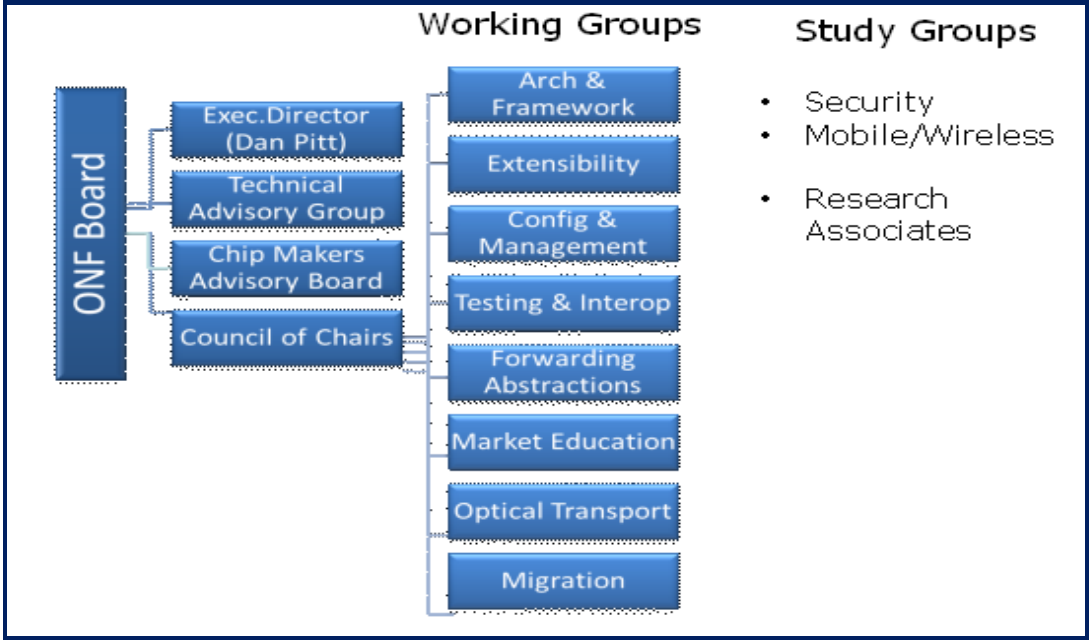

The OpenFlow Protocol

Referring back to **Figure 4** (ONF's SDN Architecture), OpenFlow is a standards-based protocol that enables an SDN controller to program the behavior of an OpenFlow-enabled switch. OpenFlow V1.0 was developed by Stanford University and was published in December 2009. The basic elements of an OpenFlow V1.0 network are shown on the left hand side of **Figure 7**. The central controller communicates with the switch's OpenFlow agent over a secure TLS (Transport Layer Security) channel. This channel could be either in-band or out-of-band. The OpenFlow agent on the switch populates the flow table as directed by the controller. Note that within **Figure 7**, OpenFlow is referred to as OF.

[13] https://www.opennetworking.org/

Subsequent to the publication of OpenFlow V1.0, the development of OpenFlow became the responsibility of the ONF. This OpenFlow specification has been enhanced three times. Version 1.1 was published in February 2011; V1.2 was published in December of 2011 and V1.3 was published in June of 2012. While few vendors adopted v1.1 or v1.2 of OpenFlow, many vendors have either already adopted v1.3 or have indicated that they will. In addition, V1.4 of OpenFlow is currently awaiting ratification.

Figure 7: The Major Components of an OpenFlow Switch V1.0-V1.3

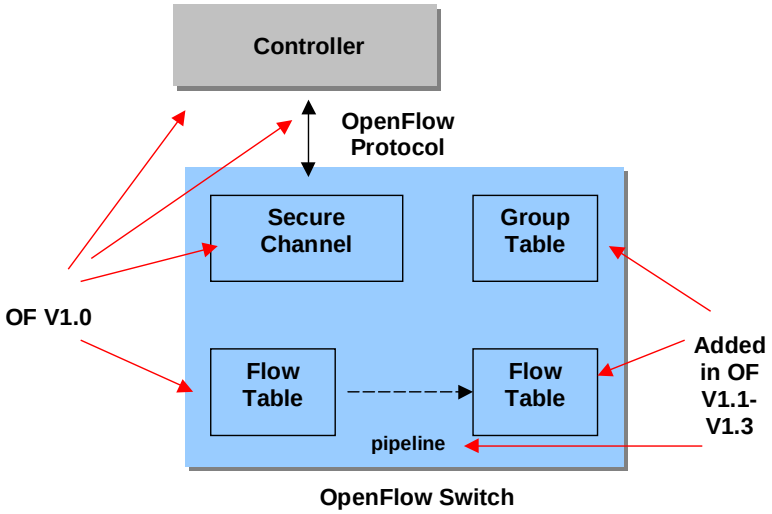

Throughout most of 2012, SDN and OpenFlow were tightly linked in the trade press as if they were either the same thing, or as if OpenFlow was required in order to implement an SDN. Neither statement is true. OpenFlow is one possible protocol that can be used to implement an SDN. In order to understand how IT organizations currently view OpenFlow, The Survey Respondents were given a set of options and were asked to indicate which option best describes the role that the OpenFlow protocol will play in their company's implementation of SDN. Their possible options and the percentage of the respondents who indicated that option are shown in **Table 15**.

Table 15: Planned Use of OpenFlow

Planned use of OpenFlow	Percentage
Will definitely include OpenFlow	16%
Will likely include OpenFlow	27%
Might include OpenFlow	31%
Will not include OpenFlow	3%
Don't know	24%
Other	1%

The data in Table 15 indicates that there is strong interest in using the OpenFlow protocol as part of implementing an SDN. The data also shows, however, that there is still a high level of uncertainty and whether or not OpenFlow will be used.

Potential Use Cases and Benefits of OpenFlow

There are a number of possible ways for the control centralization, programmability, and flow forwarding characteristics of OpenFlow to be exploited by innovative users and vendors of network devices and software. This includes the following examples.

Centralized FIB/Traffic Engineering
One of the primary benefits of OpenFlow is the centralized nature of the Forwarding Information Base (FIB). Centralization allows optimum routes to be calculated deterministically for each flow by leveraging a complete model of the end-to-end topology of the network. This model can be build using a discovery protocol, such as the Link Layer Discovery Protocol (LLDP). Based on an understanding of the service levels required for each type of flow, the centralized OpenFlow controller can apply traffic engineering principles to ensure each flow is properly serviced. Bandwidth allocations can be controlled dynamically to provide bandwidth on demand with changing traffic patterns. The result can be much better utilization of the network without sacrificing service quality. Centralized route processing also allows the pre-computation of a set of fail-over routes for each possible link or node failure. Centralized processing also can take advantage of the virtually unlimited processing power of multi-core processors and cluster computing for calculating routes and processing new flows. As shown in **Table 11**, being able to do end-to-end traffic engineering is one of the top three opportunities that The Survey Respondents associate with SDN.

The Google G-Scale WAN backbone links its various global data centers. G-Scale is a prime example of a production OpenFlow Layer 3 network that is realizing the benefits of FIB centralization. The G-Scale control plane is based on BGP and IS-to-IS and the OpenFlow-only switches are very simple 128 port 10 GbE switches built by Google using merchant silicon (when Google built these switches, 128 port 10 GbE switches had not yet been introduced in the commercial market). Google has identified a number of benefits that are associated with its G-Scale WAN backbone including that Google can run the network at utilization levels up to 95%[14]. As shown in **Table 11**, being able to increase resource utilization is the primary opportunity that The Survey Respondents associate with SDN.

Other WAN Optimizations
WAN traffic can be dynamically rerouted to reduce/control latency for VoIP and other latency sensitive applications. Traffic can also be load balanced over parallel paths of differing costs.

QoS Optimization
With OpenFlow V 1.3, per flow meters can be used for rate limiting or to provide real time visibility of application performance allowing the controller to modify forwarding behavior to maximize application performance. For example, the controller can configure an OpenFlow switch to modify the QoS markings to change the priority received over the remainder of the end-to-end path.

OpenFlow-Based Virtual Networking
With OpenFlow V1.3 virtual ports, an OpenFlow switch can be programmed to perform tunnel encapsulation and de-capsulation. Therefore, an OpenFlow switch can be programmed to be a overlay NV VTEP/NVE or gateway, as described in the section on overlay NV. As also

[14] https://www.opennetworking.org/images/stories/downloads/sdn-resources/customer-case-studies/cs-googlesdn.pdf

described in that section, OpenFlow can provide another type of network virtualization for isolating network traffic based on flows segregation or segmentation. Flows are separated based on a subset of the match fields listed earlier in the section.

OpenFlow-Based Multi-Pathing

Most networking vendors offer data center fabric solutions featuring some form of Layer 2 multi-pathing to improve the networks capacity to handle "east-west" traffic flow characteristic of server virtualization, converged storage networking, and cluster computing. OpenFlow offers another approach to multi-pathing that does not rely on standards such as TRILL or SPB. As noted earlier, the OpenFlow Controller (OFC) can use LLDP to discover the entire network topology via discovering switches and switch adjacencies. Using this topological model, OFC can compute all the parallel physical paths, including paths that share some network nodes and other paths that are entirely disjoint (and therefore offer higher reliability). OFC can then assign each flow across the network fabric to a specific path and configure the OpenFlow switches' flow tables accordingly. The OFC can then offer shared and disjoint multi-pathing as network services that can be delivered to applications. With appropriate processing power, the OFC can support very large scale networks and high availability via path redundancy and fast convergence following link or node failures.

OpenFlow Security Services and Load Balancer Services

By virtue of Layer 2-4 flow matching capability, OpenFlow access switches can perform filtering of packets as they enter the network, acting as simple firewalls at the edge. With OpenFlow switches that support modification of packet headers, the OpenFlow Controller will also be able to have the switch redirect certain suspicious traffic flows to higher-layer security controls, such as IDS/IPS systems, application firewalls, and Data Loss Prevention (DLP) devices. Other security applications built on an OpenFlow controller can match suspicious flows to databases of malware signatures or divert DDoS attacks. Another possible security application of OpenFlow would be in Network Access Control (NAC). Examples of security-oriented services that have already been announced are included in the security sub-section of this document.

OpenFlow with packet header modification will also allow the switch to function as a simple, cost-effective load-balancing device. With modification functionality, a new flow can result in a new flow table entry that includes an action to modify the destination MAC and IP addresses. The modified address can be used to direct traffic to the server selected by the controller load balancing application.

Indiana University has developed an OpenFlow-based, load-balancing application called FlowScale. According to the University[15], "FlowScale provides complex, distributed load balancing of network traffic using an OpenFlow-capable Top of Rack (ToR) switch. IU deployed the application into its Intrusion Detection System (IDS) to distribute traffic evenly to sensors. When fully deployed, the system will span the IU Bloomington and IUPUI networks and have the capability to distribute traffic at rates exceeding 500Gb/s."

Network Taps

With OpenFlow virtual ports, the functionality of a network tap can be programmed into the OpenFlow switch, allowing selected traffic to be monitored without deploying physical taps. Traffic can also be replicated and redirected to any monitoring device in the network. Big Switch networks has announced such a network monitoring application referred to as Big Tap[16].

[15] http://incntre.iu.edu/research/flowscale

16 http://www.bigswitch.com/blog/2013/07/26/network-monitoring-with-big-tap-your-first-sdn-application

Service Insertion/Chaining
OpenFlow's ability to dynamically reroute flows allows network services provided by physical or virtual appliances (e.g., firewalls, NATs, load balancers, and WOCs) to be inserted in the path of the flow. Redirecting the flow to the next service can be based on encapsulation or rewrite of the destination MAC address.

Circuit Provisioning
With extensions in V1.3 and V1.4, OpenFlow can support circuit-switched paradigms, including CWDM, DWDM, and MPLS with specific path selection and requested levels of CBR and priority. Circuits can be provisioned on a dynamic, scheduled, or permanent basis. Recovery from failed circuits can be via predetermined backup paths or by dynamic path selection. Circuit provisioning can take into account performance metrics, port states, and endpoint utilization.

The OpenDaylight Consortium

The OpenDaylight Consortium[17] was founded in April 2013. The consortium's stated mission is to facilitate a community-led, industry-supported open source framework, including code and architecture, to accelerate and advance a common, robust Software-Defined Networking platform.

As shown in **Figure 5** the consortium currently has eight platinum members, two gold members and seventeen silver members. Platinum members pay an annual fee of $500K and provide at least ten developers for a period of two years. While the commitment of the gold members and silver members is less, with the current membership the Open Daylight consortium has significant resources including annual revenues of roughly five million dollars and the full time equivalent of over eighty developers.

Figure 5: Members of the OpenDaylight Consortium

The approach that the consortium is taking to the base architecture for the OpenDaylight controller is to combine two code bases that were brought together through a collaborative proposal by Colin Dixon of IBM and David Erickson of Stanford. In addition, while the expectation is that the platinum members will make significant contributions of intellectual property, anybody can contribute code and a lot of code that has already been contributed. For example, Radware has contributed code that can be used for the detection and mitigation of Distributed Denial of Service (DDoS) attacks and IBM has contributed a version of its established network virtualization technology, called Distributed Overlay Virtual Ethernet (DOVE). Plexxi contributed code that allows both the Open Daylight controller and higher-level applications to create and share an abstract, topology and implementation independent description of the infrastructure needs, preferences and behaviors of workloads. NEC has contributed software that enables network virtualization.

[17] http://www.opendaylight.org/

The OpenDaylight Consortium has announced its intention for the first release of code. That code release is called *Hydrogen* and is expected to occur in December 2014. **Figure 6** depicts the OpenDaylight SDN Architecture and indicates some of the functionality that will be included in the first code release.

Figure 6: OpenDaylight SDN Architecture

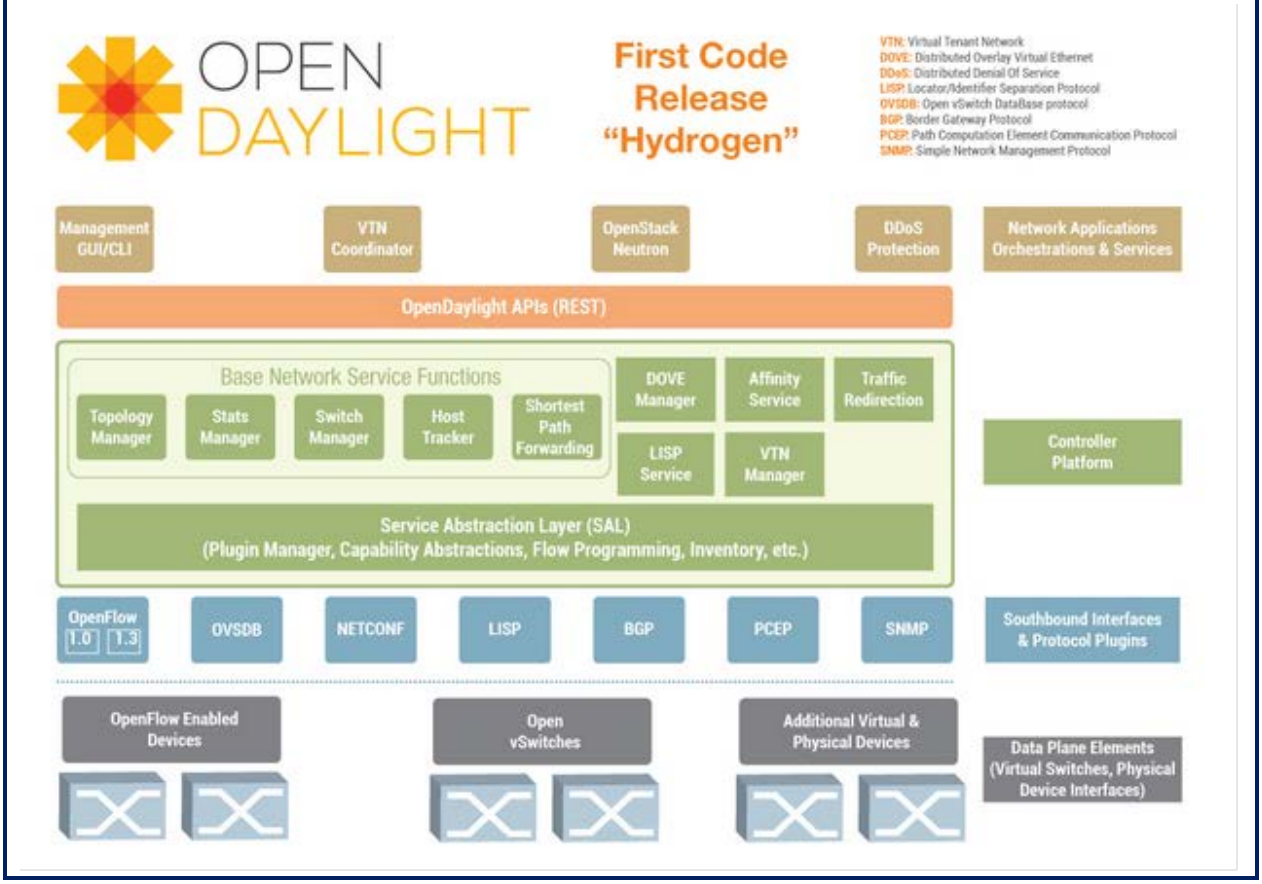

Some vendors, such as Cisco, have announced that they will use the code produced by the OpenDaylight Consortium as the basis for their SDN controller. Other vendors are taking a wait and see attitude.

Security

SDN poses both security challenges and security opportunities. The primary security challenge is to ensure that an attacker cannot compromise the central controller and hence have access to all of the subtending network elements. In addition to securing the controller itself, all communication between the controller and other devices including switches, network services platforms and management systems must be secured.

A preceding sub-section of this document contained a set of 7 questions that IT organizations should ask SDN vendors relative to their overall SDN solution architecture. Below is a set of 5 questions that IT organizations should ask SDN vendors relative to the security of their proposed SDN solutions.

1. For the controller, describe the measures that have been taken to harden its operating system and to ensure availability of the controller function.

2. Describe the authentication and authorization procedures that govern operator access to the controller. What additional physical and logical security measures are recommended?

3. Describe how communications between the controller and other devices is secured by authentication and encryption (e.g., SSL/TLS).

4. What measures are available to deal with possible control flow saturation (controller DDOS) attacks?

5. What tests have been run to verify the effectiveness of the security measures that have been taken? Is it possible to see those test results?

As noted, in addition to creating security challenges, SDN also presents opportunities to improve security by implementing security related applications that leverage the control information that has been centralized in the SDN controller. One example of such an application is DefenseFlow that was recently announced by Radware[18]. Relative to the terminology of **Figure 4**, DefenseFlow is a network service that provides DDoS protection. Another such example is HP's Sentinel application[19] that was designed to combat the security challenges that are associated with BYOD by leveraging the HP TippingPoint Repudiation Digital Vaccine data base.

To quantify the concern that IT organization have relative to security, The Survey Respondents were given the following question. "Some in the industry suggest that the implementation of SDN will make organizations less secure because if the SDN controller is hacked, the hacker has access to all of the subtending switches. Others argue that new security-oriented applications will be developed that take advantage of the SDN controller and make organizations more secure. What is the overall impact that you believe that SDN will have on network security? (Choose only one.)" Their responses are shown in **Table 16**.

18 http://www.radware.com/Products/DefenseFlow/

19 http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA4-7496ENW.pdf

Table 16: Perceived Impact of SDN on Security	
Impact	**Percentage**
Networks will be much more secure	7%
Networks will be somewhat more secure	31%
It will have no impact on network security	20%
Networks will be somewhat less secure	20%
Networks will be much less secure	3%
Don't know	19%

One observation that can be drawn from the data in **Table 16** is that overall The Survey Respondents believe that SDN will have a positive impact on security.

Management

As is the case with security, SDN presents both management opportunities and management challenges. One of the primary opportunities was highlighted in **Table 13**. That table showed the characteristic of SDN that The Survey Respondents stated would provide the most value to their company's network was the centralization of configuration and policy management. In addition, as previously described, new network management applications, such as network taps, that leverage SDN functionality are now coming to market.

SDN does, however, create some new management challenges. For example, one of the primary benefits of both the overlay NV solutions that were described in the preceding chapter of The Guide and the SDN solutions that were described in this chapter of The Guide is the ability to support multiple virtual isolated networks that run on top of the physical network. Effective operations management requires tools that give operators clear visibility into the relationships between the virtual and physical networks and their component devices. In particular, when performance or availability problems occur, both root cause analysis and impact analysis require bilateral mapping between the physical and virtual infrastructures.

With SDN, the flows between a pair of VMs can be distributed among a number of alternate paths through the network. Mapping a flow to the physical path it takes can be a challenge unless the flow monitoring solution can involve the controller's end-to-end view of the network

With SDN solutions, the controller is in the data path for new flows entering the network. During periods when many new flows are being created, the controller can potentially become a performance bottleneck adding significant latency for flow initiation. Performance management systems need visibility not only into application performance but also controller performance in processing flows. With overlay NV solutions, the controller, even if one is present, is not in the data path and does not represent a potential bottleneck. However, the overlay forwarding table must be updated frequently as VMs are created or moved

As was previously mentioned, one of the characteristics of NV and SDN is that network functions such as load balancing and firewalls are increasingly implemented in software as network services that can integrated with virtual networks or SDN flows under programmatic control; a.k.a., service chaining. Implementing these functions in software both increases the

delay associated with performing these functions and it also increases the variability of that delay. The result is an increased need for insight into the performance of each component of the overall SDN solution.

Preceding sub-sections of this document contained questions that IT organizations should ask SDN vendors relative to their overall SDN solution architecture as well as questions that IT organizations should ask SDN vendors relative to the security of their proposed SDN solutions. Below is a set of 5 questions that IT organizations should ask SDN vendors relative to SDN management.

1. Describe the extent of your management solution. For example, does it manage just the SDN solution you provide? Does the same tool also manage any traditional network components that you also provide? To what degree will it manage networks (SDN or traditional) that are provided by other vendors?

2. Describe the ability of your solution to monitor the SDN controller. Include in that description your ability to monitor functionality such as CPU utilization as well as flow throughput and latency. Also describe the statistics you collect on ports, queues, groups and meters; and the error types, codes and descriptors you report on. Also, does your solution monitor the number of flow set-ups being performed by the SDN controller?

3. How does your SDN management solution learn the end-to-end physical topology of the network? Is it possible for service assurance solutions, such root cause analysis to access this topology? Can virtual networks that have been defined be mapped to the underlying physical network elements for root cause analysis and performance analysis?

4. Describe how your SDN management solution can monitor the messages that go between the SDN controller and the SDN switches.

5. Describe the visualization functionality that your solution provides for a hybrid SDN network that is comprised of both physical network elements and virtual network elements.

The Survey Respondents were asked to indicate how much of an impact they thought that SDN will have on network management. Their responses are shown in **Table 17**.

Table 17: Perceived Impact of SDN on Management

Impact	Percentage
Networks will be much easier to manage	30%
Networks will be somewhat easier to manage	52%
SDN will have no impact on management	3%
Networks will be somewhat more difficult to manage	7%
Networks will be much more difficult to manage	4%
Don't know	4%

One observation that can be drawn from the data in Table 17 is that the vast majority of The Survey Respondents believe that SDN will have a positive impact on management.

Appendix – Chapter 2

The data path of an OpenFlow V1.0 switch is comprised of a single Flow Table that includes the rules for matching flows to table entries, an action associated with each flow entry, and counters recording the number of packets and bytes received per flow and other port and table statistics, as shown in **Figure 7.**

Figure 7: The OpenFlow V1.0 Flow Table Fields		
Header Fields	**Counters**	**Actions**

Figure 8 shows the 12-tuple of header fields that are used to match flows in the flow table,

Figure 8: The OpenFlow V1.0 Header Fields											
Ingress Port	Ether Source	Ether Dest	Ether Type	VLAN ID	VLAN Prior	IP Source	IP Dest	IP Proto	IP TOS	Source Port	Dest Port

OpenFlow V1.0 switches are required to support two basic types of actions: Forward and Drop. Forwarding is either directed to a physical port or to one of the following virtual ports:

- ALL: Send the packet out all interfaces, not including the incoming interface.
- CONTROLLER: Encapsulate and send the packet to the controller.
- LOCAL: Send the packet to the switch's local networking stack.
- TABLE: Perform actions in the flow table. Applies for only packet-out messages.
- IN PORT: Send the packet out the input port.

For OpenFlow V1.0 there are also a number of optional/recommended actions:

- NORMAL: Process the packet using the traditional forwarding path supported by the switch (for OpenFlow-hybrid switches)
- FLOOD: Flood the packet along the spanning tree
- ENQUEUE: Forward a packet through a specific port queue to provide QoS
- MODIFY FIELD: Change the content of header fields, including set VLAN ID and priority, strip VLAN, modify Ethernet or IPV4 source and destination addresses, modify IPV4 TOS, modify transport souce and destination ports

When a packet arrives at the OpenFlow V1.0 switch, the header fields are compared to flow table entries. If a match is found, the packet is either forwarded to specified port(s) or dropped depending on the action stored in the flow table. When an OpenFlow Switch receives a packet that does not match the flow table entries, it encapsulates the packet and sends it to the controller. The controller then decides how the packet should be handled and notifies the switch to either drop the packet or make a new entry in the flow table to support the new flow.

Over the last year and a half extensive enhancements have been made to the OpenFlow specification under of the auspices of the ONF. A complete listing of the enhancements included in OpenFlow V1.1-V1.3 is well beyond the scope of this document. However, some of the major changes include:

- Additional components of a flow entry in the flow table as shown below. In addition to the match and counter fields, the following fields are included in the entry:
  - Instructions to execute actions or to modify the action set or pipeline processing

- Priority: matching precedence of the flow entry
- Timeouts: maximum amount of time or idle time before flow expiration
- Cookie: opaque data value chosen and used by the controller to process flows

Error! Objects cannot be created from editing field codes.

- Flexible pipeline processing through multiple flow tables, as shown in the right hand side of **Figure 7**. As a packet is processed through the pipeline, it is associated with a set of accumulating actions and metadata. The action set is resolved and applied at the end of the pipeline. The metadata allows a limited amount of state to be passed down the pipeline.

- The new group table abstraction and group action enable OpenFlow to represent a set of ports as a single entity for forwarding packets. Different types of groups are provided, to represent different forwarding abstractions, such as multicasting or multi-pathing.

- Support for virtual ports, which can represent complex forwarding abstractions such as Link Aggregation Groups (LAGs) or tunnels. Encapsulation/Decapsulation of packets supports Network Virtualization tunnels, including PBB, QinQ VLAN stacking, and Push/Pop/Rewrite of MPLS headers.

- OpenFlow Extensible Match (OXM) uses a TLV (Type Link Value) structure to give a unique type to each header field increasing the flexibility of the match process.

- Basic support for IPv6 match and header rewrite has been added, via OXM.

- Routing emulation (Time to Live (TTL) decrement)

- Per flow meters which can be used to measure and control the rate of packet forwarding—including rate limiting packets sent to controller

- Support for multiple controllers to improve reliability

With V 1.4, OpenFlow will provide enhanced extensibility of the OpenFlow wire protocol and a new set of port properties to provide support for optical ports. This will allow Ethernet optical ports or optical ports on circuit switches to be configured and monitored.

Chapter 3: The NV and SDN Ecosystem

Overview of the NV and SDN Ecosystem

One measure of the extent of the NV and SDN ecosystem is that there are currently more than 100 members of the Open Networking Foundation[20] (ONF). This subsection of The Guide identifies the major categories of organizations that are part of the NV and SDN ecosystem and briefly discusses the value proposition of each of the categories.

This subsection of The Guide also identifies representative members of each category of organizations that are part of the NV and SDN ecosystem. The representative members that are identified either currently provide the indicated functionality or can be expected to provide the indicated functionality in the near term. As is explained below, in some instances there can be a very wide range in terms of the functionality provided by the members of a given category.

Merchant Silicon/Chip Vendors

Value Proposition: These vendors are in a position to provide hardware support in switching chips for protocols such as OpenFlow and VXLAN. This will have the effect of increasing the speed and scalability of solutions. Longer term there is also the possibility of at least some of these vendors developing cost-effective switch silicon that is optimized for OpenFlow and other controller/switch protocols.

Representative Members:

- Broadcom
- Intel
- Marvell
- Mellanox

HyperScale Data Centers

Value Proposition: Part of their value proposition is that these high-profile vendors either already are or are likely to be early adopters of SDN. As a result, these vendors are having a significant indirect impact on the development of SDN. In addition, vendors such as Google, Yahoo and Facebook are board members of the ONF. As such, these vendors directly influence the work of the ONF in general and of the evolution of the OpenFlow protocol and the northbound API in particular.

It is possible that some of these vendors will also influence the development of NV. However, some of the major players in this segment of vendors, such as Facebook and Google, currently make little use of NV.

Representative Members:

- Yahoo
- Google
- Facebook

20 https://www.opennetworking.org/blog/tag/open-networking-foundation

Telecom Service Providers

Value Proposition: Part of the value proposition of this class of vendors is similar to the value proposition of hyper-scale data center providers. For example, these vendors either already are, or are likely to be early adopters of SDN and/or NV in order to support their cloud offerings. In addition, vendors such as Deutsche Telekom, NTT Communications and Verizon are also board members of the ONF.

The preceding chapter of The Guide discussed the interest that IT organizations have in either using SDN in the WAN or in acquiring a service from a WAN service provider that is based on SDN. Responding to that interest, vendors like Pertino[21] are currently using SDN and Network Function Virtualization (NFV)[22] to enable them to offer a new generation of WAN services and Verizon[23] has announced a trial based on using SDN to enable a new generation of data center to data center WAN services.

Representative Members:

- Pertino
- Deutsche Telekom
- NTT Communications
- Verizon

Switch Vendors

Value Proposition: Relative to SDN, the majority of these vendors takes at least some of the control functionality that has typically resided in their switches and now relies on that functionality being provided by an SDN controller. In addition, these vendors implement protocols in their switches that enable those switches to communicate with an SDN controller. These vendors are increasing reliant on merchant silicon as the basis for major portions of their switching product lines.

Most of the vendors in this category represent traditional switch vendors. An exception to that is Pica8. Pica8 provides a switch that is comprised of its network operating system loaded onto commodity white box, bare-metal switches.

Representative Members:

- Alcatel-Lucent
- Avaya
- Cisco
- Dell
- Extreme Networks
- HP
- NEC
- PICA8
- IBM

[21] http://www.pcmag.com/article2/0,2817,2415354,00.asp
[22] NFV was explained in the preceding chapter of The Guide
[23] http://searchsdn.techtarget.com/news/2240182264/Intel-DPDK-switch-and-server-ref-designs-push-SDN-ecosystem-forward

Network Management and Automation

Value Proposition: Most, if not all of the providers of NV and SDN solutions will provide at least some ability for the consumers of those solutions to manage the solutions that they provide. The members of this category of the ecosystem don't provide NV and/or SDN solutions themselves. The vendors listed below either currently provide, or soon will provide management functionality that isn't offered by the providers of the NV or SDN and solutions and/or they integrate the management of these solutions into a broader management structure. The breadth of management functionality provided by the members of this category is illustrated in the next sub-section of The Guide - the sub-section entitled *Representative Vendors*.

Representative Members:

- Packet Design
- QualiSystems
- EMC
- NetScout
- CA

Providers of Network Services

Value Proposition: The members of this category provide network services such as security and optimization that are part of NV and SDN solutions[24]. Some of these services were described in the preceding section of this report. There is the possibility that over time that a large number of independent software vendors (ISVs) will also provide these services.

Representative Members:

- Embrane
- A10
- Radware
- HP
- Riverbed
- Citrix
- Cisco
- Extreme Networks
- NEC

[24] The preceding section of The Guide discussed service chaining/Insertion

Testing

Value Proposition: The members of this category either provide products that enable equipment manufacturers and others to test NV and SDN solutions or they provide the testing themselves.

Representative Members:

- QualiSystems
- InCNTRE
- Ixia
- Spirent

Standards Bodies

Value Proposition: The members of this category create standards for protocols such as OpenFlow or VXLAN. These standards form the basis for enabling products from disparate vendors to interoperate.

Representative Members:

- ONF[25]
- IEEE
- IETF
- Network Function Virtualization (NFV) – under the auspices of ETSI[26]

Providers of SDN or Network Virtualization Controllers

Value Proposition: These vendors provide the controllers that are part of any SDN solution and which are part of many NV and SDN solutions.

Representative Members:

- Big Switch Networks
- NEC
- Nuage Networks
- Netsocket
- HP
- Cisco
- Open Daylight Consortium[27]
- VMware/Nicira

[25] The role of the ONF was discussed in the preceding section of The Guide
[26] The relationship between SDN and NFV was discussed in the preceding section of The Guide
[27] The Open Daylight Consortium was discussed in the preceding section of The Guide.

Providers of Telcom Service Provider's Infrastructure/ Optical Networking

Value Proposition: These vendors are providing the infrastructure that enables telecom providers to leverage SDN in their service offerings.

Representative Members:

- ADVA Optical Networking
- Ciena
- Cyan
- Infinera
- ZTE Corporation

Server Virtualization Vendors

Value Proposition: These vendors provide the vSwitches and the hypervisor vSwitch APIs for third party vSwitches that are a key component of NV and SDN solutions.

Representative Members:

- Citrix
- Microsoft
- VMware

Representative Vendors

Avaya

The Opportunity that Avaya is Targeting

Avaya's SDN objective is to automate service delivery for applications and users across any combination of physical and virtual components – taking both human-induced error and delay out of the process.

Avaya's SDN Strategy

The key components of Avaya's SDN strategy are to:

1. Leverage OpenStack to enable rapid service creation via a common orchestration interface

- OpenStack provides an integration layer that sits above the virtualized components within the Data Center and orchestrates those resources to deliver a service through a set of APIs and a common dashboard.
- An Avaya OpenStack Horizon-based Management Platform, used to deliver, via a common GUI interface, orchestration for compute (Nova), storage (Cinder/Swift) and Avaya VENA Fabric Connect network virtualization technology (Neutron).

2. Deploy Avaya Fabric Connect (an enhanced implementation of Shortest Path Bridging) to link virtual/physical infrastructure and enable flexible network services at any scale

- Eliminates protocol overlays to deliver all services with a single protocol – making it much easier to design, manage, provision, and troubleshoot the network.
- Replaces complex network-wide provisioning practices with simple edge-only provisioning.
- Simplifies virtual machine mobility. Virtual LANs and connectivity can be extended anywhere - across Layer 3 domains and geographically dispersed Data Centers.
- Automates the provisioning of Fabric Connect through an OpenStack Neutron interface.

3. Provide public access (APIs) allowing customized interaction and integration with Avaya Fabric Connect

- Avaya is developing public access APIs directly into its Fabric Connect technology to allow for customized interaction directly with the virtualized network.

4. Extend orchestration and Fabric Connect to deliver end-to-end service creation and delivery from Data Center to Desktop

- Avaya is extending its service orchestration and network virtualization technology to the network edge in order to extend the service chain from Data Center to Desktop.
- This allows for new levels of network simplicity as services are driven top-down, by the end-point, with provisioning automated where the applications and users connect to the network.

5. Integrate policy control to automate service delivery through interaction with the application layer
- Policy controller detects users / applications and then coordinates with the orchestration system to allocate the necessary resources to support that application.

Avaya Customer Deployments

Avaya VENA Fabric Connect technology has been widely adopted by businesses across the globe. Many of these customers are also evaluating the OpenStack cloud orchestration platform as a means to enable automation and coordinated orchestration of Data Center resources. Specific customer examples include:

Sochi 2014 Olympic Winter Games

- Will be the first "Fabric-enabled" Games; providing ease of management, provisioning and deployment, simplified adds, moves and changes, and increased stability / robustness.
- This will be the first Olympics to combine video with voice and data on a single, IP-based network. Sharing the same fabric architecture will reduce the costs for video, simplify network administration and significantly boost throughput and reliability.

InteropNet 2013

- First time that the Interop organizers deployed a network fabric.
- Backbone for the entire event was staged by only three individuals, and in only a matter of days (1/10th the burden of previous years).
- Ran flawlessly for both North American Interop 2013 events (Las Vegas and NYC).

Leeds Metropolitan University

- Leverages Fabric Connect first to provide seamless L2 extensions between geographically dispersed Data Centers.
- Migrated OSPF core to Fabric Connect's Shortest Path Bridging to reduce inter-site failover times from seconds to ~20 milliseconds.
- Set up an isolated IP network across the corporate backbone to secure credit card transactions to help meet PCI DSS compliance for the banks. Provisioning at the edge only and done without adding overlay protocols or complexity.

Oslo University Hospital

- Using Fabric Connect to interconnect 40 locations.
- Leveraging the integrated VRF functionality of Fabric Connect to create secure zones for important traffic like imaging from medical devices.
- Ability to do adds, moves, and changes to the network easily without risk.
- Zero downtime environment; Fabric Connect offers a streamlined network solution – fully load balanced – with lightning fast recoveries.

Franciscan Alliance

- Leveraging Fabric Connect to provide a simplified, higher capacity network that supports increased, imaging-driven traffic requirements.
- Carrier issues no longer impact the entire network – where previously topology changes (planned or unplanned) would cause network-wide re-convergence delays – Fabric Connect instantaneously converges, delivering superior experience for both end-users and IT.

QualiSystems

The Opportunity that QualiSystems is Targeting

Software Defined Networking introduces the notion of network programmability from applications that interact with centralized SDN controllers via northbound APIs. This API-driven network paradigm opens opportunities for agile SDN application development, service creation and more seamless OSS/BSS integration. The SDN paradigm shift also creates challenges—namely the need for networking engineering organizations to adopt to a software-centric business model. Practically speaking, this means that these organizations need the ability to deliver rapid access to end-to-end network environments, in order to enable application delivery stakeholders to follow an agile dev/test process for SDN applications and OSS/BSS integration. Without access to end-to-end, production-like network environments, testing becomes a waterfall process, and quality impacts multiply as application defects are found only after they are deployed into the production network.

QualiSystems's Value Proposition

QualiSystems's value proposition in the SDN Ecosystem is to offer self-service automation for the SDN dev/test process to SDN adopters such as enterprise IT and service providers, as well as SDN eco-system vendors such as switch manufacturers. QualiSystems-empowered self-service automation provides:

- A platform for network engineering teams to build an agile dev/test process that delivers high quality SDN applications and solid network reliability
- A QA, tech support, online training test lab automation solution for SDN product vendors
- A self-service automation platform for offering cloud-based SDN app certification

Functionality Provided by QualiSystems

QualiSystems automation platform offers a number of capabilities that enable SDN dev/test self-service automation:

- **Centralized inventory management:** Engineers gain visibility to any component needed to design and publish network topologies required by developers and testers.

- **Packaged driver libraries plus open device driver creation:** QualiSystems provides driver libraries with its products, but also enable engineers to create device drivers to integrate with multi-generational legacy network devices through record and capture tools or through object-library integration of existing scripts or code. This capability ensures that networking teams maintain agility in the face of rapid changes.

- **Object-oriented automation paired with GUI tools:** QualiSystems implements an object-oriented approach to automation which contrasts with creating long, monolithic automation documents such as scripts. All automation elements including network element resources, device drivers, provisioning actions (such as loading an OS image) and automation tasks (such as running a traffic test) are captured as small-scope objects. QualiSystems' object-oriented approach offers a number of advantages:
  - The limited scope of automation objects means that they are easy to capture, maintain, and refactor to meet the requirements of a changing network environment.

- A shared library of resource, provisioning and testing objects can be maintained in a systematic fashion.
- Automation objects can be tagged with arbitrary labels so that they can be easily searched and leveraged by many users from a shared library.
- An object library optimizes the skills of programmers, who can maintain the shared library as a high quality service to the rest of the network engineering team
- The object library can be leveraged by non-programmers using GUI-based, drag and drop-style network topology design and automation workflow tools. This maximizes the productivity of the whole network engineering team, especially as topologies and workflows are shared and reused by multiple users.

- Self-Service Portal: QualiSystems enables network engineers to publish heterogeneous network topologies (including SDN and non-SDN elements) to a self-service portal catalog for dev/test users to access.

QualiSystems Proof Points

QualiSystems automation technology is being used to create a cloud-based SDN app certification self-service by one of the industry's leading networking vendors. When app developers want to certify their SDN apps, they can go to a web-based self-service catalog, reserve and activate a SDN network "sandbox" that provides them with a live environment consisting of a topology of real networking switches and a SDN controller. They can connect live to this network and run API tests and view the resulting behavior in order to ensure that their SDN application will work correctly. This cloud-based dev/test environment illustrates the type of capability that enterprise and service provider adopters of SDN technology can build as a practice within their network engineering teams to support SDN application lifecycles.

Cisco

The Opportunity that Cisco is Targeting

Cloud, mobility, and big data applications are causing a shift in the data center model. New applications are placing demands on the infrastructure in new ways. Distributed applications (for example, Big Data and Hadoop), database applications (such as those from Oracle and SAP) that run on bare metal, virtualized applications running in multi-hypervisor environments, and cloud-based applications that are available on demand all impose different demands on infrastructure. These demands include:

- Infrastructure must become application aware and more agile to support dynamic application instantiation and removal
- The non-virtual nature of new emerging applications means that the infrastructure must support physical, virtual, and cloud integration with full visibility
- Infrastructure-independent applications treat the data center as a dynamic shared resource pool
- Scale-out models promote more east-west traffic, with a need for greater network performance and scalability
- Multi-cloud models require the infrastructure to be secure and multi-tenant aware

The Cisco Value Proposition

With an open, systems-based approach, Cisco addresses the needs of diverse customer networks to enable automation, visibility, and optimization of infrastructure for applications and enhanced services. The Cisco value proposition is embodied in its Cisco Open Network Environment and Application Centric Infrastructure initiatives.

Cisco's Open Network Environment (Cisco ONE: www.cisco.com/go/one) is the industry's broadest approach to make networks open, programmable and application-aware. It is cross-architectural and supports Service Provider, branch, campus and data center deployments. It advocates open standards; open APIs and open source, for a variety of network deployment options including SDN and network virtualization models, bringing together elements of orchestration, automation policy and analytics to expose the value of networks.

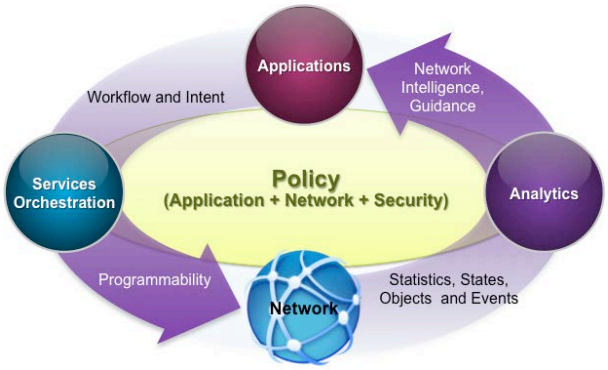

Figure 8: Cisco ONE: Applications Define the Network

Cisco ONE offers a full layer of open networking capabilities for the data center and cloud. The API helps in building custom SDN Applications to analyze fast moving networking data and to respond to application demands in real-time. Figure 8 depicts the closed loop model where open networking applications harvest network intelligence to make policy decisions in real time achieving optimized user experience across physical and virtual networks. In addition to providing consistency across physical and virtual network infrastructures with a richer set of network services,

Application Centric Infrastructure (Cisco ACI: www.cisco.com/go/aci) supports all aspects of Open Networking and delivers on the Cisco Open Network Environment strategy, embracing open APIs, open source and open standards. The vision of ACI extends beyond the network to include other infrastructure elements like compute and storage, while supporting an open ecosystem of technology and developer partners. It also goes beyond traditional SDN and overlay network virtualization models, with application centricity designed ground-up integrating the flexibility of software with the scalability of hardware. Software automation and programmability simplifies provisioning, resource scaling, and decommissioning. Hardware innovation delivers scale and performance for line rate encapsulation and de-capsulation, overlay normalization, fabric load balancing, secure multi-tenancy, and real time visibility for application or VM health.

ACI is designed around an application centric policy model, allowing the entire data center infrastructure to better align itself with application delivery requirements and the business policies of the organization. These policies automatically adapt the infrastructure (network, security, application, compute, and storage) to the needs of the business to drive shorter application deployment cycles.

Figure 9 Application Centric Infrastructure

As shown in Figure 9, the common policy driven operational model allows for both traditional enterprise applications and internally developed applications to run side by side on a network infrastructure designed to support them in a dynamic and scalable fashion. The network policies and logical topologies which have traditionally dictated application design are instead applied based on the application needs.

The key components of the ACI Fabric Include:

- **Cisco Application Policy Infrastructure Controller (Cisco APIC):** A centralized clustered controller that optimizes performance supports any application anywhere, and unified operation of physical and virtual infrastructure based on the application requirements and policies.
- **Application Network Profiles:** A collection of the end-point groups (a logical grouping of similar end-points representing an application tier or set of services that require a similar policy) their connections, and the policies that define those connections.
- **ACI-ready Nexus 9000 Switches:** Nexus 9000 Series offers modular and fixed 1/10/40 Gigabit Ethernet switch configurations to construct ACI fabrics that are designed to take full advantage of ACI's application policy driven services and infrastructure automation features.

Other elements from ecosystem partners, as well as compute, and storage will be added to make the ACI fabric richer and more inclusive.

The ACI fabric decouples application and policy from IP infrastructure through the ACI object model. The IP infrastructure retains its own distributed control architecture for performance, scale and resiliency managing the network forwarding planes. The APIC simplifies management, operations, and control with its centralized and abstracted view of the virtual and physical infrastructure. Within the APIC, software applications are defined logically using constructs that are application centric, rather than network centric. For example a group of physical and virtual web servers may be grouped in a single tier of a three-tier application. The communication between these tiers and the network and security policies that define the communication make up the complete application, which is defined as an application network profile within the APIC.

Application network profiles are used by the APIC to push the logical topology and policy definitions down to stateless network hardware in the fabric. Figure 10 shows this approach which is the reverse of traditional architectures, in which VLANs, subnets, firewall rules, etc. dictate where and how an application can run.

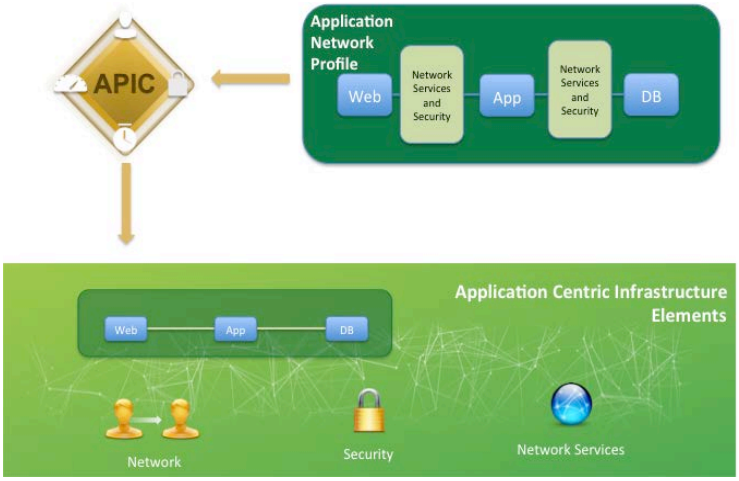

Figure 10: Application Network Profiles are used by the APIC which pushes the required Policy to the Fabric

Principles of the Cisco Open Network Environment and ACI will benefit Cisco customers with shorter application deployment cycles, driving faster business processes and quicker time to market, resulting in a sustainable competitive advantage. Agility in application delivery will define competitiveness in the future.

Cisco ONE: www.cisco.com/go/one
Cisco ACI: www.cisco.com/go/aci

NEC

The Opportunity that NEC is Targeting

NEC is focused on delivering a secure, multi-tenant software-defined network to enterprises and data centers interested in new levels of network flexibility (automation and control) and security that streamline IT management for dramatically reduced operating costs and time-to-deliver business services. NEC's ProgrammableFlow OpenFlow network fabric is also targeting solution providers who are positioning themselves for network innovation.

The NEC ProgrammableFlow Networking Suite Value Proposition

NEC Corporation is a global technology leader, with over $32 billion in annual revenues, and a 100 year history of communications innovation. The NEC ProgrammableFlow® Networking Suite was the first production ready OpenFlow-based SDN solution, generally available in May 2011. In production now within enterprises and solution providers around the world, it delivers better utilization of all IT assets, enabling network-wide virtualization and allowing customers to easily deploy, control, monitor and manage multi-tenant network infrastructure.

Key benefits delivered by the ProgrammableFlow SDN solution include:

- ***Greater business agility***: Integrating virtual and physical switches into a single control point network-wide for unified management with significantly reduced routine network maintenance and complex network protocols. NEC ProgrammableFlow customers have improved service delivery time from weeks to minutes.
- ***Improved performance***: Traditional network protocols such as spanning tree limit the use of multiple paths through the network. Traditional network designs can lead to network chokepoints. ProgrammableFlow supports multi-path layer 2 and layer 3 networks and any network topology.
- ***Reduced network operational expense***: Traditional networks require device-by-device configuration that is complex and error prone. The ProgrammableFlow Suite provides advanced automation of network management and policy, enabling centralized control that reduces routine network configuration. NEC customers have reduced network operational cost by up to 80%.
- ***Multivendor interoperability:*** Northbound and southbound APIs enable customers to choose best of breed network solutions and build unique capabilities to gain competitive advantage.

How NEC Delivers on its Value Proposition

Technology

NEC's fabric-based virtualization separates the control plane from the data plane, with network-level virtualization that provides control and management of both virtual and physical networks, delivering a resilient network with unprecedented flexibility.

Key features of ProgrammableFlow include:

- **Network Services:** Multi-tenant network virtualization, IP Mobility, L2 switching, IPv4/IPv6 switching, end-to-end QoS, policy-based routing, service chaining, metering, end-to-end flow and network visualization
- **Performance:** Throughput up to 1TB, L2+L3 Multi-path, MC-LAG, LAG

- **Scalability:** 200 Switches, 1000 virtual routers, 1 million flows
- **Reliability:** End to end reliability, High Speed Link Failure Detection, VRRP, Redundant Controller
- **Security:** Advanced ACL, Control Network Protection, Authorization
- **Standards Compliance:** PF5200 series switches are certified by Open Network Foundation OpenFlow 1.0 compliant.

Products

- **PF6800**: The award-winning ProgrammableFlow Controller, the PF6800, centralizes the control of enterprise and cloud networks, streamlining management via automation and dramatically reducing operating costs and time to deliver business services.
- **PF5200:** PF240: 48 x 1GbE + 4 x 10GbE. PF5248: 8 x 10GbE + 2 x 1GbE. OpenFlow 1.0 and 1.3.1 support. Up to 160,000 flow entries. Hardware matching and forwarding.
- **PF5820:** 48 x 10GbE + 4 x 40GbE. Up to 80,000 flow entries. Hardware matching and forwarding.

Ecosystem

Recognizing the value of an open SDN ecosystem, NEC has demonstrated interoperability with Alcatel Lucent, Arista, Brocade, Centec, Dell, Extreme Networks, IBM, Intel, Juniper Networks, and Noviflow. NEC has participated in numerous PlugFests and was the first vendor certified as compliant by the Open Networking Foundation (ONF) with OpenFlow 1.0.

In May, 2012, NEC again demonstrated SDN leadership by delivering the first northbound SDN applications, utilizing the RESTful ProgrammableFlow interfaces, provided through the NEC SDN Application Center. Today, this ecosystem includes management, optimization and security applications. Technology partners include A10 Networks, Radware (security), Real Status (management), Red Hat (cloud orchestration), Silver Peak (optimization), Telefonica Global and vArmour (security).

ProgrammableFlow Networking Suite also integrates with server and network orchestration. Now with Version 5, the ProgrammableFlow fabric can be integrated with OpenStack Grizzly, for seamless control and management. The NEC OpenStack plugin is certified by Red Hat. ProgrammableFlow Networking Suite is fully integrated with Microsoft System Center Virtual Machine Manager, further enabling automated network provisioning and management from a single point. NEC PF1000 is certified for Window Server 2012.

NEC ProgrammableFlow Networking Suite Proof Points

NEC's leadership in this disruptive technology has been widely recognized by industry watchers, the media and customers. Multiple awards have been bestowed upon the networking solution, including Best of Interop 2011 in Infrastructure, Grand Prize at Interop 2012 in Las Vegas, and the first Networking Innovation Award in SDN from Tech Target in 2012.

NEC customers, operating ProgrammableFlow networks in production today, include Nippon Express, one of the 50 largest companies in Japan, Genesis Hosting Solutions, a virtual infrastructure hosting company in Chicago, IL, NTT Communications, with the first global SDN powering their enterprise cloud solutions, Stanford University, Marist College and Georgia

Institute of Technology. These companies and institutions have discovered unprecedented flexibility and agility with ProgrammableFlow Networking Suite.

Specific examples include Kanazawa University Hospital, who had deployed ProgrammableFlow and gained significant new mobility of critical devices on the network. Kanazawa also cut as much as 80% of their operating costs with a networking solution that costs roughly the same in up front capital costs.

Nippon Express found the automation and streamlined configuration of their ProgrammableFlow solution no longer necessitated investing in systems integrators to do network configurations, saving the company a minimum of $75K annually. Network ticket turnaround time was reduced from 2 months to 10 days for increased agility, and operating expenses at Nippon Express were also reduced dramatically with a 50% reduction in footprint, and 80% reduction in power.

By implementing ProgrammableFlow within its own data centers, NEC enabled rapid provisioning of new, secure virtual networks for NEC developers – with the same speed they could previously deliver server and storage virtualization. The company has been able to delay new server investments because of the network pooling and resource management capabilities of ProgrammableFlow, enabling a higher return of over IT investment.

For more NEC ProgrammableFlow case studies and product information, visit the NEC SDN website at www.necam.com/sdn.

Nuage Networks

The Opportunity That Nuage Networks is Targeting

Nuage Networks sees the opportunity to make networks as fluid and responsive as the compute infrastructure has already become, and cloud applications need them to be. Today, data center networks have fallen significantly behind what is required and what is possible. They are operationally complex, restricted and inefficient. Configuration is highly manual, and performed device by device. As they look to turn up new applications, CIOs and IT administrators are encumbered by various tedious networking details that are irrelevant to their broader mission. It all adds up to delays, errors and frustration for users who need their applications turned up. It is simply not an IT-friendly paradigm. In one customer engagement after another, enterprise CIOs demand more agility and higher efficiency. They want to turn up applications and workloads at will anywhere, do it instantaneously and cost-effectively, and do so without losing control & visibility.

Nuage Networks sees value in SDN-enabled automation and virtualization that remove the constraints holding back the network from being as dynamic as the cloud requires. Today, it takes weeks of elapsed time and numerous iterations of work orders between manual processes in order to establish the network connectivity required by virtual machines that come up in seconds in support of application requirements. That is simply not the right thinking for the cloud era. What's needed is reflexive and instantaneous network establishment, in tune with the needs of applications and their administrators and users.

Further, broad-based migration of business-critical applications to the cloud requires more than what we have seen to date with consumer cloud offerings and early public clouds. That is because control and visibility are paramount to IT departments who are committed to ensuring application performance for their workgroups while respecting the security and compliance realities that underpin their business.

The Nuage Networks Value Proposition

With key pillars of programmability through **abstraction** and efficiency through **automation**, The Nuage Networks value proposition is to offer SDN solutions that change the current environment and deliver truly business-grade and hybrid cloud services that pave the way for a true hybrid cloud era. As part of their value proposition, Nuage Networks helps cloud service providers and enterprises make their networks as fluid and dynamic as cloud applications need them to be. Nuage Networks also offers the proper **abstraction** of networking capabilities in a more open environment and the elegant **automation** that makes network connectivity instantaneous in response to application needs, in a policy-based manner.

Functionality Provided by Nuage Networks

To deliver against that value proposition, the Nuage Networks Virtualized Services Platform (VSP) enables programmable and automated network services infrastructure in support of the most demanding virtualized applications across multi-tenant environments.

The Nuage Networks VSP is comprised of three key modules, each of which run as virtual machines (VMs) on standard compute platforms of choice, and participate in one of the three key tiers of the network hierarchy. Collectively they ensure that the Nuage Networks VSP offers

enterprises and cloud service providers IT-friendly abstraction of network services needed by applications and policy-based network automation, without compromising control and visibility.

- Within the cloud management plane, the Nuage Networks Virtualized Services Directory (VSD) serves as an advanced policy and analytics engine through which network operators can define the "rules of the game" across slices and sub-slices of network resources offered to tenants or user groups. Through the VSD, permissions and policy can be defined and assigned in a hierarchical fashion, using IT-friendly language and constructs. Once defined, policies can be templated so that they can be easily used many times. In this way, each tier of the role-based hierarchy has full visibility and control within the bounds of their defined scope. This includes access to granular analytics powered by a hadoop engine as part of the Nuage Networks VSD.

- In the control plane, the Nuage Networks Virtualized Services Controller (VSC) serves as a robust SDN controller. Leveraging the principles that underpin scaling of the Internet, instances of the Nuage Networks VSC federate using standard IP protocols to ensure boundless scaling and global network visibility. By peering with DC WAN routers and existing networks, the Nuage Networks SDN controller (VSC) discovers topology and reachability information that enables seamless connectivity within and across datacenters as well as to private datacenters and enterprise locations.

- In the data plane, the Nuage Networks Virtual Routing and Switching (VRS) element extends network endpoint control all the way out to the servers. The Nuage Networks VRS is a hypervisor-resident implementation that provides full layer 2 (L2) through layer 4 (L4) capability for virtualized or bare metal servers, making them fully integrated extensions of a massively distributed virtual routing and switching system under SDN control.

The Nuage Networks SDN approach makes data center and networks more readily consumable, programmable and scalable. It virtualizes and automates any data center network infrastructure, and extends the reach of cloud services to enterprise locations and private datacenters. In that way, cloud services are securely accessible by their users operating in branch or headquarters facilities, and seamlessly integrated across private data centers that house critical data. While eliminating network boundaries, the Nuage Networks solution has been designed to operate seamlessly across operational and organizational boundaries as well.

To deliver the benefits of SDN automation and abstraction to any cloud datacenter, the Nuage Networks SDN implementation accepts the datacenter infrastructure as it stands. Nuage Networks VSP is agnostic to hypervisors, with support for leading hypervisors including KVM, Xen, ESXi and Hyper-V. It is also agnostic to cloud management platforms of choice, including OpenStack, CloudStack, and vCloud Director. Lastly, the Nuage Networks approach is agnostic to networking hardware that is in place such as Top of Rack switches and aggregation/distribution switches. Nuage Networks simply serves to fully virtualize and automate that infrastructure within and across datacenters, and provide seamless connectivity of those assets to enterprise locations, which are already served by VPN services today.

In many cases, incorporating bare metal assets seamlessly into the SDN automation scheme is also an area of great interest and benefit. To that end, in the past quarter Nuage Networks announced further enhancements to the VSP that extend the network automation benefits of SDN to include the full breadth of datacenter assets. In addition to software gateways that have been shipping since Q2 2013 (Nuage VRS-G) and support of 3rd party Virtual Tunnel Endpoint

(VTEP) devices through the Nuage Networks **ecosystem of partners** such as Cumulus Networks, we announced the Nuage Networks 7850 Virtualized Services Gateway (VSG) platform. The 7850 VSG delivers a terabit of switching & routing capacity in a single rack unit, an innovative alternative for large datacenters in which the proportion of bare metal assets demands higher performance.

In being among the first of the global network equipment suppliers to appreciate the full potential of the cloud as a transformative technology, Alcatel-Lucent invested over two years ago in key ventures like Nuage Networks and the CloudBand NFV platform that are at the heart of making more agile and programmable cloud networks a reality.

Nuage Networks' Proof Points

Since the launch of Nuage Networks in April 2013, over a dozen trials have been successfully completed with large enterprises as well as cloud service providers and network operators.

Trial customers of the Nuage Networks VSP solution to date include enterprises for whom IT is a critical asset, in key verticals such as financial services, healthcare, or manufacturing. The University of Pittsburgh Medical Center (UPMC) is a representative example of this category of customers. In these trials, enterprises are eager to minimize delays and errors that result from highly manual network provisioning, and to accelerate application delivery to their user groups without sacrificing control and visibility.

Likewise Service Providers with datacenter assets and cloud ambitions are aggressively trialing the Nuage Networks VSP solution, in many cases to develop offers that incorporate datacenter assets as a natural extension of L2 and L3 VPN services they already offer today. Telus (Canada), SFR (France) & exponential-e (UK) are publicly disclosed Nuage Networks trials representative of this category.

Amidst the strong interest in the Nuage Networks SDN solution across all regions of the world, the need for a more fluid and automated network infrastructure that is dynamic but policy-driven in support of multi-tenant and hybrid cloud environments is a common denominator.

More info:
www.nuagenetworks.net
Delivering Effortless Connections in the Data Center Network and Beyond
View a demonstration of Nuage Networks VSP
Blog
Follow us on Twitter
Like us on Facebook
Follow us on LinkedIn

HP

The Opportunity that HP is Targeting

Many enterprises are unable to create business innovation because of aging networking environments. Network design and architectures have remained largely unchanged for more than a decade. While applications and systems have evolved to meet the demands of a world where real time rules, the underlying network infrastructure has not kept pace.

Software-defined networking (SDN) redefines the way IT organizations think about the network and removes the barriers to innovation by giving service providers and enterprises complete programmatic control of a dynamic, abstracted view of the network. With SDN, IT can become more agile by orchestrating network services and automatically controlling the network according to high-level policies, rather than low-level network device configurations.

HP's Value Proposition

HP's networking vision is centered on simplification. HP has established itself as the clear #2 vendor in the networking market. By providing a complete, open SDN solution to automate the network from data center to campus and branch, HP is providing organizations a networking solution that is simpler and can deliver more business value than traditional networks.

HP's SDN solution delivers:

- A complete solution across the infrastructure, control, and application layers
- Solutions for providers & enterprises spanning the data center, campus & branch - including network virtualization, security, UC&C
- An open-standards and open API based approach
- An open SDN ecosystem with strong partnerships and developer tools

Functionality Provided by HP

HP has been investing in SDN technologies since 2007 with its support of OpenFlow. HP has proven itself to be a leader in SDN technologies and continues to deliver milestones including:

- 2007 – HP collaborates with Stanford on Ethane, the predecessor to OpenFlow
- 2008 – HP demos first hardware OpenFlow-enabled switch on a commercially shipping platform
- 2009 – HP earns 10 first OpenFlow lighthouse customers
- 2010 – HP scales lighthouse customers to 60
- 2011- HP is a founding member of the ONF and delivers generally available OpenFlow software on 16 switch models
- 2012 – HP develops beta customers for SDN apps, SDN controller
- 2013 – HP launches more apps and creates open SDN ecosystem & SDN app store

HP currently has over 50 switch models and 10 router models that are OpenFlow-enabled. This represents over 25 million installed ports that can support OpenFlow with a simple software update.

In addition to OpenFlow, HP is committed to supporting other industry standards & protocols that help enable SDN. For example, HP currently supports overlays in the data center via the VXLAN protocol.

HP has commercially released the HP Virtual Application Networks (VAN) SDN Controller. The VAN SDN Controller features:

- Availability as either software or an appliance
- Full support for the OpenFlow protocol
- Open APIs to enable third-party SDN application development
- An extensible, scalable, and resilient controller architecture

HP has also announced several key applications that are currently in development or beta phases including:

- HP Virtual Cloud Networks – Delivers Network Virtualization
  - Automates cloud network provisioning
  - Integrates with OpenStack
- HP Sentinel Security SDN Application – Provides real-time security across SDN-enabled networks
  - Consumes reputation security intelligence from the HP TippingPoint DVLabs cloud
  - Protects from over 1M botnets, malware, and spyware malicious sites
  - Innovative features such as time-of-day whitelisting & blacklisting
- HP UC&C SDN Application for Microsoft Lync – Optimizes user experience for UC&C applications
  - Dynamic network policy on per call basis
  - Sets traffic priorities and routing
  - Integrates Lync Server with network controller for improved intelligence
- HP ConvergedControl – Unifies physical & virtual networking in the data center
  - Unifies underlay & overlay visibility and control
  - Integrates with VMware NSX through controller federation

HP also continues to work with partners and customers to develop use cases and SDN applications including:

- Verizon & Intel – Dynamic WAN bandwidth provisioning use case
- CERN OpenLab – Load balancing SDN App

HP is focused on delivering mainstream SDN and has launched a new open ecosystem intended to make that possible. The HP SDN ecosystem delivers resources to develop and create a market place for SDN applications. The HP SDN ecosystem delivers the following benefits:

- Simple - Extending simplicity of programmability across the network with OpenFlow-enabled devices
- Open - Raising the value of SDN with an open environment delivered by SDN Software Development Kit (SDK).
- Enterprise ready - Fostering innovations with industry's first SDN App Store market place for SDN applications.

Learn more about HP SDN solutions by visiting hp.com/sdn

HP Proof Points

HP customers are excited about the possibilities that SDN offers and have been working closely with HP on developing our SDN use cases and solutions. Once such customer is Ballarat

Grammar, a school in Australia. Ballarat Grammar wanted to increase student and faculty productivity and educational opportunities by allowing users to bring their own devices and securely connect to the network

Ballarat deployed the HP Sentinel Security SDN application in their production network to enable secure BYOD connectivity. They leveraged OpenFlow-enabled HP switches, the HP VAN SDN Controller along with the Sentinel app.

Through their deployment of the HP SDN solution, they were able to:

- Increase student and faculty productivity
- Enable secure BYOD connectivity
- Provide real-time protection from one million threats
- Allow proactive IT management of threats
- Decrease the time IT spends on security problems, from days or weeks to hours
- Enhance network security—thousands of threats easily detected and blocked
- Realize investment protection through free software update to OpenFlow

Read the full case study here.

Ciena

The Opportunities that Ciena is Targeting

Ciena believes that the SDN paradigm opens up new possibilities to capitalize upon multi-layer—combined packet and transport—infrastructures. Ciena also believes that by simultaneously considering both service demands and capacity across the Ethernet, OTN, and optical domains, services may be placed on the optimum technology based on the application's requirements. In addition, the network may be periodically re-optimized to ensure the most efficient deployment of network resources.

Ciena also sees the opportunity to utilize analytics to systematically match application service demands against available capacity in real time. Ciena believes that through the use of predictive analytics, dynamic pricing models are feasible to not only optimize network utilization, but more importantly, to maximize revenues by adjusting spot pricing based on real-time demands.

Ciena's Value Proposition

Ciena offers intelligent network infrastructure solutions and extensive software expertise, complemented by comprehensive services. Ciena leverages its deep expertise and proven market leadership in packet and optical networking and distributed software to deliver solutions in alignment with OP^{n}, Ciena's architecture for building next-generation networks. OP^{n} represents a highly scalable and programmable infrastructure that can be controlled by network-level applications. The company's solutions form the foundation of many of the largest, most reliable and sophisticated service provider, enterprise, government, and research and education networks across the globe.

Functionality Provided by Ciena

Ciena offers a broad range of transport solutions that are deployed in the Metro Access through the Core. Their solutions leverage industry leading optical, packet, and converged packet-optical technologies, in conjunction with their OneControl Unified Management System.

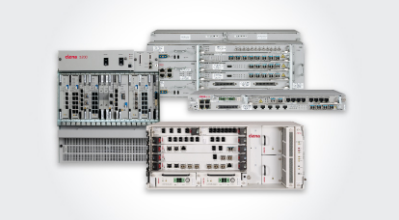		
Optical	**Packet**	**Converged Packet-Optical**

Ciena is recognized as a world leader in distributed transport control plane technology, and as a major contributor to the Metro Ethernet Forum (MEF), Optical Internetworking Forum (OIF), and other leading telecommunications and networking standards bodies.

Ciena's SDN Strategy

Ciena is embracing Software-Defined Networking (SDN) and leading the charge toward multi-layer, carrier-scale SDN in the Open Networking Foundation (ONF), where Ciena is a founding member and a leading contributor.

SDN is a key component of Ciena's OP^n architecture, which drives down the networking cost curve with a converged packet-optical architecture and highly intelligent software functionality.

Ciena's SDN strategy consists of four key pillars that enable carriers to realize the vision of SDN:

- **Autonomics Operations Intelligence** – Streamline network operations through automation and highly intelligent network control to unleash innovation and differentiation, while dramatically improving time to market.
- **Expansive Openness** – Embrace openness above, below, and within the logically centralized control layer to leverage the emerging SDN ecosystem to achieve multi-vendor interoperability.
- **Multi-Layer Control –** Achieve exponential scalability at the lowest cost by optimizing end-to-end service delivery across the traditionally separate and distinct service-layer boundaries.
- **Carrier Grand and Scale –** Augment the SDN architecture for adoption into the rigorous carrier-grade environment at significantly higher scale.

Ciena's Proof Points

Ciena's initial SDN offering is the V-WAN Network Services Module, which provides an SDN control layer for efficient data center interconnection, in conjunction with their OneControl network management system. V-WAN and its companion, the IT WAN Orchestration Application Services Module, automates the allocation of network resources across cloud data centers to provide performance on demand and to enable multi-tenancy and seamless VM mobility. Ciena is expanding its offering through its flagship controller and OpenFlow-standard agents that will be embedded into select hardware platforms throughout their portfolio.

Ciena has also teamed with Research & Engineering (R&E) networking leaders to build an international Software-Defined WAN to accelerate the SDN ecosystem. Ciena is collaborating with CANARIE, Internet2 and StarLight to build the industry's first Open SDN WAN for the R&E community. The industry's first 100G R&E WAN based on OpenFlow-based SDN, spans 2,500 Km and links Ciena laboratories in Ottawa and Hanover, Maryland with Starlight's facility in Chicago. Ciena's longstanding presence in the Internet2 and R&E community renders Ciena uniquely qualified to design, operate, and deploy the SDN testbed.

The research network is designed to help spur innovation in the telecommunications industry by giving R&E institutions and other network operators a platform to experiment with SDN and other advanced technologies like agile photonics and real-time analytics software applications.

A10

A10's Value Proposition

A10 Networks' Advanced Core Operating System is a 100% software-based platform for high-performance Application Service Gateways that deliver application availability, optimization and security. Because the system is entirely software-based, it supports a variety of virtualized and physical form factors and deployment models to meet a growing array of IT consumption models, including managed hosting providers and Cloud IaaS services. Given the software nature of ACOS, A10 is in the process of integrating a variety of emerging Cloud orchestration and SDN protocols in order to support form factors and usage models in Cloud, SDN and Network Virtualization architectures. Specific network virtualization tunneling protocols and cloud orchestration management platform support will be announced imminently (and can not be divulged publicly at this point).

Pica8

The Opportunity that Pica8 is Targeting

Many IT decision makers are uneasy about potentially disrupting their businesses by introducing new technology into their data center. Pica8 believes SDN adoption can occur with benefits realized at a pace that is best for each organization. As a result, Pica8 is delivering integrated solutions that help take the guesswork out of compiling the proper components while minimizing the resources required making SDN application implementations easier and faster.

Pica8's Value Proposition

Since 2009, Pica8 has challenged the traditional networking premise that hardware and software need to be tightly coupled. Decoupling hardware and software helps data centers enjoy greater flexibility, automation and personalization, with reduced implementation and operational costs. In a single package, PicOS™, a hardware-agnostic, Debian-based and OpenFlow-supporting switching OS, is loaded onto commoditized bare-metal switches to best leverage white box economics. Pica8 currently ships both 1G and 10G Ethernet systems leveraging PicOS on switches from Accton, Quanta and Celestica. Pica8 also offers its OS to its customers without hardware if requested.

Functionality Provided by Pica8

Pica8 provides a turnkey open networking tool in a complete stack that is designed to integrate seamlessly and quickly into an existing network infrastructure. Pica8's open switches run a high-performance L2/L3 protocol stack that has OpenFlow 1.3 integration. Pica8 leverages Nicira's Open-vSwitch (OVS) v1.9[28] as the OpenFlow interface within PicOS. OVS runs as a process within PicOS, and is interoperable with any OpenFlow device, including leading OpenFlow controllers such as Ryu, Floodlight, and NOX. The Starter Kits ship with Ryu[29] as the primary controller. Pica8 collaborates closely with both the OVS and Ryu open-source projects.

SDN Use Cases Targeted by Pica8

Use Case: Traffic Engineering – Traffic engineering is a method of optimizing network performance by dynamically analyzing, predicting and regulating the behavior of data transmitted over the network. SDN delivers the ability to externally program network devices in real time, through emerging standards like OpenFlow. For latency-sensitive applications, such as a Hadoop Cluster, traffic engineering increases performance by dynamically programming the ''fast'' path for that application's traffic flows to traverse. OpenFlow 1.3 supports statistics that give visibility into application performance. Subsequently, the best path can be externally programmed into the physical network based on real-time information, and this concept can be extended to manage WAN transit costs between data centers.

Use Case: Tunneling – Generic Routing Encapsulation (GRE) is a tunneling protocol that encapsulates a wide variety of network layer protocols inside virtual point-to-point links over an Internet Protocol inter-network. For example, it is possible to transport multicast traffic and IPv6 through a GRE tunnel. More recently, the idea of using SDN to help orchestrate the movement

[28] http://openvswitch.org/
[29] http://osrg.github.io/ryu/

of virtual machines (VM) provides protection and dedicated paths for a specific VM domain. OpenFlow 1.3-based tunnels can be externally programmed into the physical network to connect logical domains, and protect the traffic traversing between them.

Use Cases: Dynamic Network Taps – Traditionally, a network tap is a purpose-built hardware device that provides a way to access the data flowing across an IP network. Network taps are commonly used for network intrusion detection systems, VoIP recording, network probes, RMON probes, packet sniffers and other monitoring and collection devices. OpenFlow 1.3 provides the means to externally program network tap-like functionality into any OpenFlow-compliant physical switch. This SDN-driven capability reduces CapEx by dynamically adjusting the tap's characteristics, thereby increasing flexibility and avoiding dedicated devices.

Pica8 Case Study

Pica8 has more than 200 customers worldwide, including web services companies, global carriers and leading research labs. One such research lab is the Ocean Cluster for Experimental Architectures in Networks (OCEAN) at the University of Illinois, Urbana-Champaign. This lab enables networked systems research from low-level physical wiring to network protocols and applications, via as SDN-capable network testbed. OCEAN is currently composed of 176 server ports and 676 switch ports, using Pica8 Pronto 3290 switches via TAM Networks, NIAGARA 32066 NICs from Interface Masters, and servers from Dell.

OCEAN is supporting research in software-defined networks, security, cloud computing, low-latency networking, and more. Projects include:

VeriFlow – VeriFlow helps monitor and report upon network operations by verifying network-wide correctness and security properties that operates with millisecond-level latency as each forwarding rule is modified by the network controller.

Jellyfish – Unorthodox data center network architecture designed to answer the formidable demands placed on data center networks by big data analytics and cloud computing. Jellyfish is based on a random interconnect among switches that yields flexible incremental expansion and 25-40% higher bandwidth than state-of-the-art, fat-tree designs using the same equipment.

LIME – LIME consistently migrates virtual networks that clones the data plane state to the new location, then incrementally migrates traffic source in a manner that is both consistent (i.e., transparent to applications running at endpoints and the network controller) and efficient (i.e., fast and low overhead).

To conduct this research, OCEAN researchers designed new, networked systems from low-level physical wiring up to network protocols and applications, via an SDN-capable network test bed. They selected Pica8 switches to be included in their overall system because of their ability to easily integrate through its Debian-based OS while leveraging commoditized hardware and SDN.

“Incorporating support of OpenFlow was a major consideration for deploying Pica8 switches here,” said Brighten Godfrey, assistant professor of computer science at the University of Illinois and collaborator on the project. “This meant that we could have a software-agnostic way to externally program the network as needed with far less effort. That capability, in addition to leverage the economics of white box switches provided us with significant CapEx savings.”

Packet Design

The Opportunity that Packet Design is Targeting

To realize the promise of full programmability, enabling networks to adapt dynamically to workload demands, SDN controllers must have always-accurate intelligence of network topology and traffic profiles. Without this information, provisioning network resources on demand to satisfy one application request may impinge on the needs and performance of others.

Part of the SDN opportunity that Packet Design in targeting is that traditional management tools are incapable of:

- Providing real-time visibility into traffic paths across the network
- Showing how routing [mis]configurations impact service delivery
- Monitoring traffic flows across both customer and service provider networks to give a complete view

As a result, network managers face challenges like these:

- Finger pointing between service providers and customers over SLA breaches
- Little visibility into MPLS VPN routing
- Unintended consequences from routing configuration changes
- Failure to detect new devices and configurations
- Lack of accurate data to optimize peering relationships
- Inability to accurately model and predict the impact of new workloads

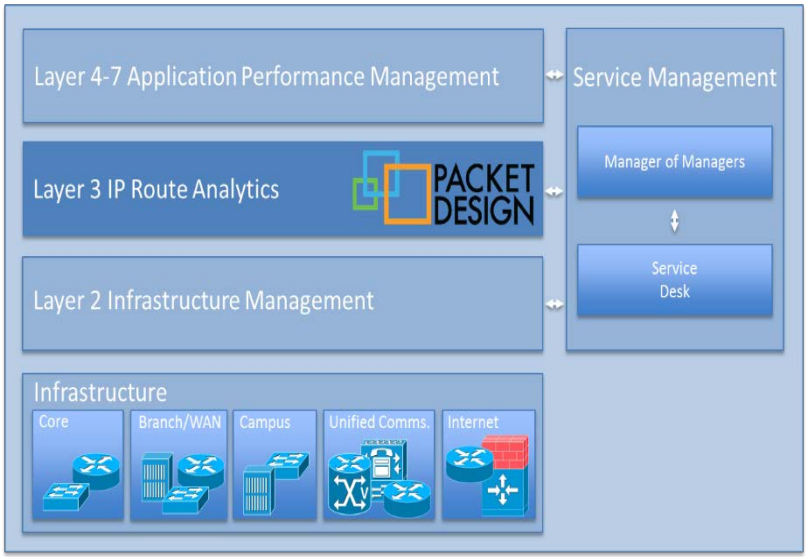

Packet Design's Route Analytics Technology Fills the Network Management Gap

Packet Design's Value Proposition

To address the opportunity discussed above, Packet Design is currently adapting its analytics for SDNs and using open APIs to create a Network Access Broker (NAB) that, based on its real-time models, historical data and business policies, informs the Controller of the impact of network change requests before they are made.

The Route Explorer system provides visibility into the network's routing topology and events that are invisible to other tools. Network managers can see exactly how

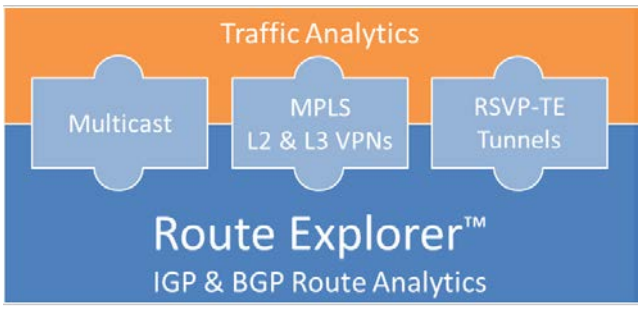

The Route Explorer System

traffic traverses the entire network and quickly identify sub-optimal routing metrics, flapping, loops, black holes and a host of other conditions that can cause service delivery issues and inefficient use of network resources. All major IGP and BGP routing protocols are supported.

Functionality Provided by Packet Design

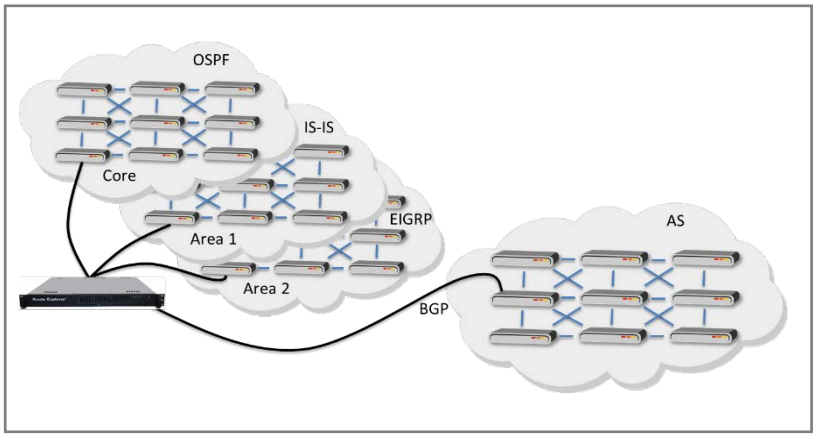

The Route Explorer System Sees the Network as the Routers See It

Using Packet Design's patented technology, the Route Explorer system participates passively in the network and subscribes to all routing announcements. It records these messages and uses them to calculate and maintain a real-time model of how the network forwards traffic. It maintains the changes in a time-indexed data base so that the network forwarding model and events for a time period in the past can be retrieved, analyzed and played back using animation.

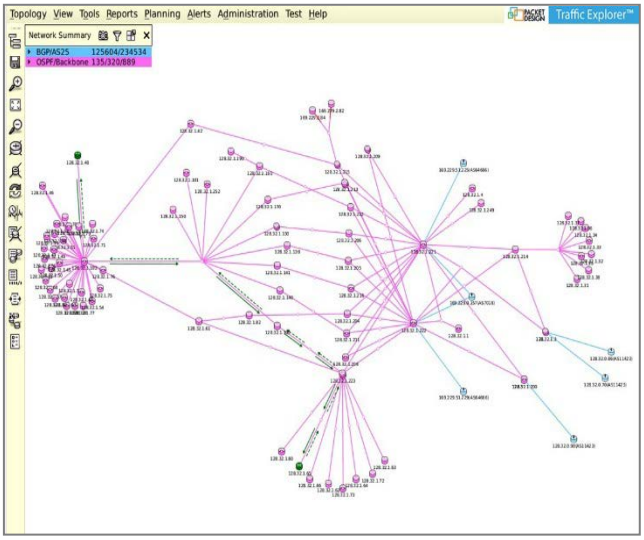

Maintaining an Accurate Routing Topology and Traffic Profiles

The Route Explorer system enables proactive service management with real-time monitoring of routing and traffic behavior, anomaly reports and alerts to deviations from baseline. Troubleshooting intermittent and hard-to-find problems is made easier with an intuitive history navigator that enables users to select and drill down into any time period for forensic analysis of routing events and traffic paths. In addition, an interactive what-if modeling capability allows engineers to see the impact of planned and unplanned network changes and failure conditions before they occur. This is invaluable prior to maintenance windows, before adding new workloads, and for assessing network resiliency.

Packet Design Proof Points

Since its founding in 2003, Packet Design has pioneered the complex science of route analytics to address these management challenges. By leveraging the distributed intelligence of the Internet Protocol, Packet Design products restore visibility into the behavior of complex, mission-critical networks. The Route Explorer™ system, a unique combination of routing and traffic analytics, delivers unmatched "path-aware" visibility, analysis and diagnosis capabilities that help network managers improve network availability and performance while reducing operating costs and delivering a strong ROI. With real-time, historical, summary and detail-level data, they can model the impact of network changes accurately and troubleshoot problems fast.

Netsocket

The Opportunity that Netsocket is Targeting

Although software-defined networks on the market today have offered a dramatic improvement over rigid and inflexible legacy networks that are costly and complex to own and manage, Netsocket believes that current SDN offerings do not adequately address the needs of enterprises and their service providers.

Netsocket's Value Proposition

Enterprise environments need a unified network architecture that supports both data center and distributed remote office networking requirements. Netsocket's mission has been to develop an SDN-based, virtualized networking infrastructure that provides:

- End-to-end completely virtualized networking for enterprises and service providers
- Automated orchestration and provisioning for deployment at scale
- A complete solution for automated low cost edge/branch networking
- Commoditization of hardware, eliminating the need for proprietary routers
- Seamless legacy network interoperability that allows an 'at-your-pace' network migration
- Automated networks that can be optimized according to network usage and events

Functionality Provided by Netsocket

Netsocket Virtual Network (NVN) is a fully optimized, automated and cost-effective virtual network specifically optimized for LAN and WAN edge network deployment. Some of the key features of NVN are:

SDN Architecture functionality for
- End-to-end virtual networking that is independent of physical infrastructure
- Virtual Layer 2/3 Switching
- Virtual Carrier-Grade Routing, Firewall and VPN/Tunneling

Centralized Network Automation Management that provides
- Unified Network Management
- Real-Time Network Service Analytics
- Intelligent Network Remediation

Interoperability and Integration with
- Legacy routed networks
- OpenStack™ & Microsoft System Center

The SDN framework of the NVN is comprised of several components within the application, controller, and infrastructure Layers of the SDN framework. All NVN applications are virtualized to run on commodity x86 server platforms.

The Netsocket Virtual Network solution's infrastructure layer component is the vFlowSwitch™. Responsible for Layer 3 packet forwarding, this virtualized application is also hosted in a hypervisor virtual machine. One or more vFlowSwitch modules can be associated with a parent vFlowController. Within a network deployment, the vFlowSwitch is coupled to Layer 2 virtual switches, a native component of the hypervisor environment.

The NVN component of the controller layer is the vFlowController™. This component runs in a virtual machine and provides control for Layer 3 packet flow. The vFlowController includes intrinsic virtualized routing, firewall and tunneling applications. A primary advantage of the intrinsic nature of these applications is their close coupling with flow control functions. This pairing delivers performance benefits above and beyond what can be delivered through an overlay methodology, where networking functions interface through a traditional SDN northbound API. The vFlowController supports critical routing protocols (BGP, OSPF) as well as a full range of edge routing and LAN features including network address translation (NAT), port forwarding and network access control (NAC).

The NVN application layer hosts Netsocket's powerful applications for orchestration and automation: vNetCommander™, VNetOptimizer™ and plug-in's that integrate NVN workflows to vendor environments. The integration mechanism is an open, feature-rich web service called vSocket. Key application capabilities derived through the vSocket northbound web service include network workflow automation for installation, bulk provisioning and software upgrade orchestration; real-time network health and state visibility; and dynamic network self-optimization based on changing flow and application environments.

The Business Case for a 50-Site Distributed Enterprise

Netsocket Virtual Network provides a cost effective, streamlined solution for distributed enterprise networks. NVN centralizes and automates enterprise-wide network management workflows through its network management application, vNetCommander.

NVN eliminates the need for special-purpose Cisco or Juniper router hardware. All networking functions are completely virtualized, allowing them to be hosted on commodity X86 hypervisor-enabled servers at each site. These servers may be used to host other applications for the business, such as session border control, WAN optimization or automated backups.

Whether the enterprise network is managed by internal staff or through a Managed Service Provider (MSP), the Netsocket approach reduces both networking capex and opex, resulting in a 4:1 savings in total cost of ownership (TCO) over legacy edge network solutions.

NVN's virtualized components can be distributed at any number of enterprise remote sites. A single vFlowController is hosted on a commodity server at a head-end facility, while the vFlowSwitch, is hosted on an inexpensive server at each site. vNetCommander, the web-based lifecycle management application, is hosted for authenticated access through a browser. The management UI contains no CLI. All network workflows are administered through an intuitive web GUI.

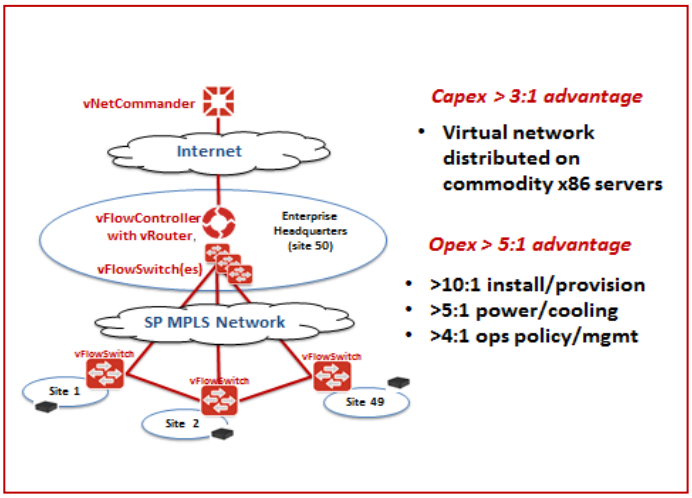

Inter-site security and privacy are provided through a set of fully-meshed IPSec tunnels between sites as well as stateful firewall access control lists (ACLs). End-to-end QoS is enabled through policing, priority queuing and rate limiting. Interoperability with legacy routers allows an at-your-pace, site-to-site migration.

The benefits of NVN over a legacy routed network are substantial and compelling:

- 3:1 Capex savings by hosting NVN on commodity X86 server platforms

- 5:1 Opex savings over the lifecycle of the network through:

    - Automation of site activation, software installation and initial provisioning — 10:1 savings
    - Reduction of power and cooling requirements — 5:1 savings
    - Reduction of daily operational and administrative costs, including move/add/change and system management costs — 4:1 savings through the simplified workflows of the intuitive web GUI of vNetCommander and the fact that no CLI expertise is required

EMC

The Opportunity that EMC is Targeting

EMC is targeting the management and monitoring challenges that are associated with network virtualization-based cloud infrastructures.

EMC's Value Proposition

Designed to manage physical, virtual and software-defined data center environments, the EMC Service Assurance Suite helps maximize availability, performance, and efficiency—providing the critical elements needed to effectively manage heterogeneous infrastructures for efficient IT service delivery.

EMC's Service Assurance Suite is differentiated in its ability to deliver automated root-cause and impact analysis to IT operations. This is due in large part from its extensible object model which provides the visibility, analysis, and dynamic updating of the relationships among elements of the service-delivery infrastructure. The extensible, dynamic nature of EMC's software suite enables EMC customer's flexibility in how they monitor the physical and virtual infrastructure used to deliver applications and services to the business.

Current and Planned Functionality Provided by EMC

By being able to easily add new classes and objects to this model, EMC has proven its ability to provide detailed operational insights into large scale virtualized server environments. Through integration with server hypervisors such as VMware vCenter Server, the EMC Service Assurance Suite, provides a complete, end-to-end view of the physical and virtual infrastructure used for the delivery of applications and services. IT operations teams can use the suite to monitor from virtualized applications/processes, to VMs, from virtual switches through the virtual network, to the physical switch, and out to the rest of the physical infrastructure. The Service Assurance Suite also provides awareness and insight into all the storage connectivity used by these VMs and their applications and processes, bringing the network, server and storage components into a single management view for complete visibility and management.

In the same way EMC has provided end-to-end IT operations monitoring and management across physical and virtual infrastructure environments, EMC will extend the power of its model-based management technology to include the software-defined networking elements. This will provide IT operations with full management visibility, root-cause analysis, and service-level monitoring for network virtualization – eliminating the management complexity associated with monitoring a new layer of networking abstraction and accelerating deployment of next generation network architectures.

EMC Proof Points

EMC has been working extensively with VMware on integration between the EMC Service Assurance Suite and VMware NSX network virtualization platform. This integration, which was demonstrated at VMworld 2013, illustrates the detailed monitoring and management needed to effectively manage the software-defined networking infrastructure. In addition, EMC customers in industries such as financial

services and telecommunications have already deployed customized versions of the Service Assurance Suite. Currently most of these customers are using the Service Assurance Suite in limited environments as part of the process of refining and enhancing their strategies and plans for bringing network virtualization into production use.

Targeted for delivery in 2014, EMC's network virtualization management capabilities will enhance the Service Assurance Suite by enabling customers to have full operational management and monitoring of network virtualization-based cloud infrastructures

Chapter 4: Planning for NV and SDN

Introduction

As noted in the executive summary, there is considerable confusion in the industry relative to NV and SDN. This confusion is understandable given the breadth of problems that these solutions are supposed to address combined with the variety of approaches that vendors are proposing. Further adding to the confusion is the embryonic nature of most of the solutions that are available in the marketplace, the current limited adoption of these solutions and the overall level of hype associated with NV and SDN.

Given this confusion, it would be understandable if an enterprise IT organization decided to take a wait and see attitude about NV and SDN. While that response would be understandable, it isn't the right approach to take either from the perspective of the IT organization or the IT professional. That follows in part because even though no reasonable person would claim to know in detail how SDN and network virtualization will evolve over the next several years, there is no doubt that:

- IT organizations need to solve the problems (e.g., support the dynamic movement of virtual machines, reduce operational complexity) that NV and SDN are designed to solve.

- Many of the characteristics of NV and SDN solutions (e.g., more reliance on software, increased use of automation) are already being broadly adopted within IT organizations.

There is also no doubt that implementing NV and/or SDN presents risk, but that ignoring NV and SDN presents significant risk to both enterprise IT organizations and to IT professionals. The risk to enterprise IT organizations is that by ignoring NV and SDN they remain unable to solve the problems that NV and SDN are designed to solve and this puts their company at a competitive disadvantage. The risk to IT professionals is that ignoring NV and SDN delays their coming up the learning curve on these new approaches which would result in a diminishment of the value they could provide either to their current employer or to a future employer.

Market Research: The Current State of Planning

This subsection of The Guide presents recent market research that quantifies:

- How IT organizations are approaching analyzing and implementing NV and SDN;
- The plans that IT organizations have for open source solutions and open protocols;
- The expectations that IT organizations have for how broadly they will implement NV and SDN.

How IT organizations are Approaching NV and SDN

The Survey Respondents were asked to indicate the approach that their company is taking relative to adopting NV and SDN. The Survey Respondents were allowed to indicate multiple approaches and a summary of their responses is shown in **Table 18**.

Table 18: Approach to Implementing NV and SDN

Approach	NV	SDN
We have not made any analysis of it	26%	19%
We will likely analyze it sometime in the next year	26%	26%
We looked at it and decided to not do anything with it over the next year	6%	5%
We are currently actively analyzing the potential value that it offers	25%	36%
We are currently actively analyzing vendors' strategies and offerings for it	12%	20%
We expect that within a year that we will be running it either in a lab or in a limited trial	14%	19%
We currently are running it either in a lab or in a limited trial	10%	13%
We expect that within a year that we will be running it somewhere in our production network	6%	10%
We currently are running it somewhere in our production network	7%	6%
Don't know	5%	4%
Other	1%	4%

The way to read the data in **Table 18** is that 26% of The Survey Respondents work for companies that haven't made any analysis of NV and 19% of The Survey Respondents work for companies that haven't made any analysis of SDN.

While there are some differences between the overall approach that IT organizations are taking to NV and the overall approach that IT organizations are taking to SDN, there are more similarities than there are differences. The high level story told by the data in **Table 18** is that today there is a lot of interest in both NV and SDN, but very little deployment of either in production networks. If The Survey Respondents are correct, there will be a modest increase in the production use of both NV and SDN in 2014.

The Expected Role of Open Source and Open Protocols

Given that the phrase *open networking* is often associated with SDN, The Survey Respondents were asked to indicate the type of SDN solution that their organization was likely to implement within the next two years and the possible responses focused on open protocols and open source solutions. The respondents were only allowed one choice. Their responses are shown in **Table 19**.

Table 19: Likely SDN Solutions

Selection	% of Respondents
Open source, open protocols, multiple vendors	21%
Open source, open protocols, single vendor with an ecosystem of partners	12%
A mix of open and proprietary protocols based on a single vendor with its ecosystem of partners	22%
Most attractive solution regardless of openness or number of vendors	17%
It is unlikely that we will implement SDN within two years	15%
Don't know	12%

One conclusion that can be drawn from **Table 19** is that in spite of all of the discussion of open networking on the part of press and analysts, less that half of The Survey Respondents who work for a company that will likely implement SDN within the next two years are currently committed to SDN solutions based on open source and open protocols[30].

[30] The Survey Respondents who answered "don't know" were excluded from the calculation.

Anticipated NV and SDN Deployment

The Survey Respondents were asked to indicate how broadly they expected that their campus, WAN and data center networks would be based on SDN and/or NV three years from now. The question didn't make any attempt to define how network virtualization would be implemented; e.g., as part of an overlay solution or through manipulating OpenFlow tables. Their responses are shown in **Table 20** and **Table 21**.

Table 20: Expected Deployment of NV

	Exclusively Based on NV	Mostly NV	Hybrid – NV and Traditional about Equal	Mostly Traditional	Exclusively Traditional	Don't know
Campus Networks	1%	12%	40%	24%	10%	14%
WAN	0%	7%	35%	32%	11%	14%
Data Center Networks	3%	25%	38%	17%	4%	13%

Table 21: Expected Deployment of SDN

	Exclusively Based on SDN	Mostly SDN	Hybrid – SDN and Traditional about Equal	Mostly Traditional	Exclusively Traditional	Don't know
Campus Networks	2%	14%	42%	26%	7%	10%
WAN	2%	10%	35%	35%	10%	8%
Data Center Networks	3%	28%	39%	18%	5%	7%

Some of the conclusions that can be drawn from the data in **Table 20** and **Table 21** include:

- Although The Survey Respondents expressed the strongest interest in deploying NV and SDN in their data centers, they also expressed significant interest in deploying NV and SDN in both campus and wide area networks.

- Only a small percentage of The Survey Respondents indicated that in three years that their networks would be either based exclusively on traditional techniques or exclusively on NV and/or SDN.

- The Survey Respondents' most common response was that three years from now that each type of network would be comprised roughly equally of a traditional approach and an approach based on NV and/or SDN.

Crafting an NV and/or SDN Plan

This section of The Guide outlines a process that a hypothetical company, that will be referred to in this section as *GottaChange,* can use to plan for the implementation of NV and/or SDN. The intention is that IT organizations will customize this process for use in their environments.

Define NV and SDN

As described in previous chapters of The Guide, there isn't uniform agreement in the industry as to the precise definition of NV and/or SDN. *GottaChange* can't wait for the brouhaha surrounding the definition of NV and SDN to sort itself out. As part of developing an implementation plan, *GottaChange* must develop a definition of NV and/or SDN that is well understood and agreed to within their organization.

Identify the Primary Opportunities

In order to intelligently choose vendors, architectures and enabling technologies, *GottaChange* needs to first identify the primary opportunities that they are hoping to address by implementing NV and/or SDN. To assist with this process, Chapter 1 of The Guide identified the primary use cases for NV and also presented market research that showed the interest that The Survey Respondents had in each of the use cases. Chapter 2 did the same for SDN.

To exemplify the relationship between the opportunities and the various solutions being proposed by vendors, consider the fact that if the primary opportunity that is driving an IT organization is the need to support the dynamic movement, replication and allocation of virtual workloads, then an overlay-based NV solution from a vendor such as Nuage Networks or Netsocket is a viable candidate, as is a solution from a company such as NEC that implements NV by manipulating OpenFlow tables. An overlay-based NV solution unto itself, however, doesn't make it easier to respond to other opportunities such as making it easier to implement QoS, nor does it enable applications to dynamically request services from the network[31].

Identify the Key Metrics

Having identified the primary opportunities, *GottaChange* needs to identify the key business-related metrics that are associated with each opportunity. The principal use of these metrics is to enable the IT organization to create a business case for implementing NV and/or SDN. However, *GottaChange* should use these metrics throughout the evaluation process; i.e., evaluating solution architectures and performing a proof of concept.

In some cases the key business metrics may be obvious. For example, if one of the primary opportunities that *GottaChange* is trying to address is the centralization of configuration management and provisioning, then one of the key business metrics associated with that opportunity is likely to be labor savings. In contrast, if one of the primary opportunities is to enable business agility, it may be more difficult for *GottaChange* to identify one or more IT-related metrics that, if NV and/or SDN improve them, lead to measurable business value.

[31] Chapter 2 of The Guide discussed an overlay/underlay model that can address these opportunities.

Define the Scope of Possible Solutions

As some point in the planning process *GottaChange* needs to define how broad of an NV and/or SDN solution they are seriously considering implementing. This could come after the first phase of the evaluation process (see below). It should come prior to *GottaChange* moving forward with performing a proof of concept (POC). As described below, the broader *GottaChange* defines the potential solution, the more risk and the more organizational resistance they will encounter.

In addition, based on how *GottaChange* defines what they mean by a NV and/or SDN solution, it may or may not be possible for them to acquire a complete solution from a single vendor. For example, it is reasonable to consider a NV solution based on overlays to be a complete solution unto itself. Analogously, it is reasonable to think of one or more SDN controllers and the underlying network elements as being a complete solution. If *GottaChange* uses one or both of these approaches as their definition of an NV and/or SDN solution, then it is possible for *GottaChange* to buy a complete solution from a single vendor.

However, if *GottaChange* has an expanded definition of *solution*, it is less likely that they will be able to acquire a complete solution from a single vendor. An expanded definition of what *GottaChange* means by solution could include functionality such as orchestration; the L4 to L7 functions that are inserted into the service that is consumed by users; and the business applications that access the control information in the SDN controller.

Decide: Best of Breed vs. Systems Solution

As described above, based on how *GottaChange* defines what they mean by an NV and/or SDN solution, it may be possible for them to acquire a complete NV and/or SDN solution from a single vendor; a.k.a., a systems solution. However, even if it is possible for GottaChange to buy a systems solution they may decide to at least explore the option of buying best of breed components from varying vendors. If *GottaChange* determines that they are willing to acquire components from varying vendors, *GottaChange* must evaluate the testing that was done on both the individual components as well as the complete solution; how the solution will be updated and tested over time; and whether or not there is a *single throat to choke*.

It's reasonable for *GottaChange* to think that if they are acquiring a complete NV and/or SDN solution from a single vendor, that the solution won't have interoperability issues. While that is a reasonable thought, IT organizations still need to request details of the testing that was performed by the vendor themselves, as well as the results of any third party testing that was performed. This testing is important both to demonstrate interoperability of the components of the solution as well as to identify the performance limits of the solution.

Evaluate NV and/or SDN Solutions

The process that *GottaChange* uses to evaluate NV and/or SDN solutions should be cyclical. As part of the first stage of the evaluation process, *GottaChange* should perform a cursory evaluation of numerous vendors. The primary goal of the first stage of the evaluation process is to enable *GottaChange* to determine which solutions correspond to the opportunities that they are seeking to respond to and it also makes *GottaChange* aware of the varying approaches to SDN that the vendors have, each with their own value add. Upon completion of the first stage of the evaluation process, *GottaChange* is in a position to eliminate vendors from consideration

and to begin a more detailed analysis on a small set of vendors. As described below, the result of this detailed analysis may well be the recommendation to go forward with a POC.

When evaluating a vendor's SDN solution, IT organizations need to understand the following aspects of those solutions.

- **The Solution Architecture**

  This includes topics such as which components of the solution are provided by the vendor and which are provided by a partner; what functionality is done in hardware vs. in software; how much control is centralized in the SDN controller; what protocols are used within the solution; how the solution supports high availability and the level of abstraction that is provided by the controller's northbound API.

  In addition, *GottaChange* must evaluate the various NV and/or SDN solutions based on their ability to respond to the opportunities that the IT organization has identified. For example, assume that one of the opportunities that the *GottaChange* has identified is being able to support the dynamic movement of VMs. Given that, then as part of the evaluation of solution architectures, *GottaChange* has to identify how each solution accomplishes this.

  Chapter 2 of The Guide contains a set of 7 key questions that *GottaChange* can ask vendors about the architecture of their SDN solutions.

- **The Controller**

  GottaChange must evaluate the architecture of a number of NV and/or SDN controllers. For example, does the controller have a modular architecture that will enable the addition of new functionality over time? *GottaChange* also needs to understand how the controller's architecture enables scalability, high availability and performance. At the author's web site[32] is a white paper that discusses ten criteria that IT organization should use to evaluate SDN controllers[33].

- **The Network Elements**

  Most overlay-based NV solutions are network agnostic. If that is the type of solution that *GottaChange* is evaluating, then it is highly likely that there isn't a need for them to evaluate the network elements on which the potential NV solutions run.

  However, if *GottaChange* is evaluating solutions that closely resemble the ONF definition of SDN that was presented in Chapter 2, then *GottaChange* should ask the vendors questions such as:

  1. Which switches, both virtual and physical, support your SDN solution? For OpenFlow-enabled switches, identify whether the switch is a pure OpenFlow switch or a hybrid OpenFlow switch.

[32] www.ashtonmetzler.com
[33] Ibid.

2. What protocols do you support between the control layer and the infrastructure layer of your proposed solution? What network behaviors are enabled by these protocols and what types of services can be constructed using those behaviors?

3. If Open Flow is supported, what versions have been implemented? What required features of the supported version are not included in the implementation? Indicate which of the optional features it supports. Describe any significant vendor-specific extensions that have been made.

4. If one of the switches in your proposed solution is in SDN mode, are there any types of traffic that must be processed partially in software before being forwarded?

5. If one of the switches in your proposed solution is in hybrid mode, does that have any impact on the behavior of the traditional component of the switch? If yes, explain.

- **Management**

  There are two aspects of NV and/or SDN management that *GottaChange* needs to evaluate. One aspect is the ability of the vendor's solution to alleviate the management challenges created by NV and/or SDN. Based on the type of solution that *GottaChange* is considering, this may include monitoring the performance of the controller; providing end-to-end visualization of the virtual networks; configuring the SDN switches and monitoring the physical and logical networks between switches. The second aspect of management that *GottaChange* needs to evaluate is the integration of the management of NV and/or SDN into a broader management solution.

  Chapter 2 of The Guide contains a set of 5 key questions that *GottaChange* can ask vendors about the management of their SDN solutions.

- **Security**

  There are also two aspects of security that GottaChange needs to evaluate. One aspect is what functionality the vendor provides in order to secure their NV and/or SDN solution. One of the reasons this is important is because the NV and SDN controllers are new attack surfaces. The other aspect of security that needs to be evaluated is the ability of the solution to enhance the overall security of the IT infrastructure. An example of how SDN can potentially improve security is Radware's recent contribution to the Open DayLight consortium's SDN controller of a toolset that can be used for the detection and mitigation of DDoS attacks.

  Chapter 2 of The Guide contains a set of 5 key questions that GottaChange can ask vendors about the security of their SDN solutions.

- **Additional Functionality**

  There are two approaches that an IT organization can take relative to implementing network functions that ride on the SDN controller. One approach is to acquire the network functions from a vendor. Two examples of vendor provided network functions were already discussed. One is Radware's DDoS application and the other is NEC's Virtual Tenant Networking functionality. Since most IT organizations will acquire network

functions from vendors, evaluating vendor supplied network functions is a key component of the overall process of evaluating SDN solutions.

The second approach is for the IT organization to develop some or all of the required network functionality itself. The primary advantage of this approach is that it enables the IT organization to customize the network functions to meet the organization's specific requirements. One of the disadvantages of this approach is that it requires the IT organization to have the base of skills that are necessary both to develop the network functions and to maintain those functions over their life cycle.

GottaChange should use the process of evaluating NV and/or SDN solutions to determine if it can acquire all of the network functions it needs to respond to the opportunities that it has identified or if it has to develop some or all of those functions itself.

Test and Certify Solutions

As previously discussed, even if all of the components of an NV or SDN solution come from a single vendor, as part of evaluating those solutions *GottaChange* needs to understand the testing that was done to ensure both the smooth operation and the performance of the solution. Particularly in those situations in which the components of the SDN solution come from multiple vendors, *GottaChange* needs to understand if the solution is certified. By that is meant, if *GottaChange* implements the solution, will it have a single point of contact to resolve any problems that develop.

There may be instances in which *GottaChange* has to either do testing itself or to commission a third party to do testing on its behalf. For example, if *GottaChange* were to develop one or more network functions, it would need to test the operation of those functions on the controller(s) that it had selected and it would need to redo that testing prior to implementing new versions of the controller or new versions of the network functions. If *GottaChange* anticipates facing a situation like this then as part of the evaluation process, *GottaChange* needs to evaluate both the tools that are available to enable the organization to do the testing itself as well as the functionality provided by external test labs.

Integrate with the Existing Environment

It is certainly possible for *GottaChange* to evaluate NV and/or SDN solutions in isolation from the IT organization's current environment. However, given that the NV and/or SDN solution might at some time be implemented in *GottaChange's* production network, then as part of the evaluation process *GottaChange* should examine how the SDN solution would fit into the existing infrastructure. For example, what mechanisms exist to enable traffic to flow between the SDN solution and the traditional network? Is it possible to extend the SDN solution so that it operates both in a data center and in a branch office? So that the solution operates in multiple data centers?

Educate the Organization

Both NV and SDN are both embryonic and rapidly evolving. Hence, in order to create and update a plan to potentially implement one or both of these architectures, *GottaChange* must continually educate itself as to what is happening in the broad NV and/or SDN ecosystem. This certainly includes analyzing what is being said in the industry about the relevant use cases and

the techniques that can be used to justify deployment. It also includes reviewing product announcements; the announcement of enabling technologies that are either new or have evolved; the results of plugfests that are intended to test the interoperability of SDN solutions; and the work of organizations such as the Open DayLight consortium.

Much of the education discussed in the preceding paragraph can be accomplished by reading articles and white papers and by attending seminars and workshops. *GottaChange* should also consider downloading some of the open source products that are readily available and playing with those solutions to gain deeper insight into their capabilities and weaknesses. In addition, by yearend 2013 the author will publish a mock RFI for SDN solutions that will be hosted at the author's web site (www.ashtonmetzler.com). *GottaChange* can use this document to structure a dialogue with selected vendors.

Evaluate Professional Services

Given that SDN is a new way of implementing networking, *GottaChange* may choose to use a professional services organization to help with one or more stages in the overall Plan, Design, Implement and Operations (PDIO) lifecycle. The relevant services that *GottaChange* might use could be technology centric (e.g., developing SDN designs, testing SDN solutions), organization centric (e.g., evaluating the skills of the current organization, identifying the skills that are needed and creating a way to develop those skills) or process centric; e.g., evaluating the current processes and developing new ones. These services could be light-weight (i.e., the professional services organization provides limited support) or heavy-weight. They may also be consumed just as part of an initial rollout of NV and/or SDN or they could be consumed over an extended period of time as The Company extends its deployment of NV and/or SDN.

If *GottaChange* is considering leveraging professional services from a third party, then as part of the overall evaluation process, *GottaChange* needs to evaluate the professional services that are provided, both by the potential providers of the NV and/or SDN solution as well as from independent providers of professional services.

Eliminate Organizational Resistance

Organizations tend to resist change and typically the amount of resistance is directly proportional to the extent of the change. Hence, if *GottaChange* is looking at a narrowly defined SDN solution, such as one that implements a network tap application, it can expect minimum organizational resistance. Conversely, if *GottaChange* is looking at a broadly defined SDN solution, then it must anticipate significant organizational resistance.

Organizations are particularly resistant to change if that change is likely to have a significant impact on jobs. Both NV and SDN have the potential to impact the jobs of network professionals. For example, the deployment of NV and/or SDN is likely to reduce the amount of manual labor that *GottaChange* has to perform and is likely to increase the amount of programming that *GottaChange* chooses to perform. As part of planning for NV and/or SDN, *GottaChange* needs to anticipate resistance from the network organization and respond accordingly. For example, *GottaChange* may sponsor members of its network organization achieving some of the new certifications that various NV and/or SDN vendors have recently announced.

However, a number of other factors are also impacting the jobs of IT professionals. This includes mobility, the virtualization of servers and desktops, the convergence of technologies (i.e. networks, servers, compute) and the broad and growing adoption of varying forms of cloud computing. As a result, *GottaChange's* VN and/or SDN initiatives may be just one more factor contributing to the need for *GottaChange's* IT organization to take a broad look at the skills it will need on a going forward basis and to implement a plan to develop those skills. As previously noted, *GottaChange* has the option of leveraging a professional services provider to perform a skills assessment of *GottaChange*'s IT organization.

Perform a POC

Assuming that the previous steps in their plan have produced positive results, *GottaChange* may well elect to perform a POC. The breadth of the POC is directly related to how *GottaChange* has scoped the proposed NV and/or SDN solution and the length of the POC is directly related to the criticality of the tasks that the solution is intended to support.

One goal of a POC is to determine if indeed the proposed solution works and if so, how well it performs. Another goal is to quantify the previously defined key metrics that are associated with each opportunity that *GottaChange* is hoping to address.

Obtain Management Buy-In

GottaChange's network organization needs varying levels of management buy-in at the various stages of their NV and/or SDN plan. For example, little if any management buy-in is needed just for members of *GottaChange's* network organization to attend a seminar or workshop and in many cases, little buy-in is needed in order for them to download open source solutions and to spend a modest amount of time coming to understand the functionality and the limitations of those solutions. Increasing levels of management buy-in are typically needed to engage vendors in detailed discussions of NV and/or SDN, to conduct a POC or to implement an NV or SDN solution.

GottaChange is more likely to get management buy-in if the members of the project team that is evaluating NV and/or SDN anticipate management's concerns and work to resolve those concerns over the entire planning cycle. For example, like virtually all organizations, *GottaChange* will likely face management resistance to implementing any technology or new way of delivering technology if the associated security and compliance concerns are not thoroughly addressed. In addition, *GottaChange* will likely face management resistance if any of *GottaChange's* key processes are impacted.

Like virtually all IT organizations, *GottaChange* will need to develop some form of business case to justify implementing NV and/or SDN. There are three primary components to the business case that *GottaChange* has to develop. One component is the identification and quantification of the benefits that will occur if *GottaChange* implements the proposed NV and/or SDN solution. As noted, one of the primary reasons for performing a POC is to quantify those benefits. Another component of the business case is a multi-year financial analysis that details all of the costs as well as the benefits that are associated with implementing the proposed solution. The third component of the business case is an analysis of what *GottaChange's* IT organization will do to mitigate the risk that is associated with implementing the proposed solution. In addition to mitigating the risk associated with the solution not performing well, this includes mitigating the

previously mentioned concerns that management has about issues such as security, compliance and existing processes.

Summary and Conclusions

There is no doubt that over the next few years that NV and SDN will have a significant impact both on enterprise networks and on the role of network professionals. Because of that, IT organizations and IT professionals need to develop a plan to evaluate and potentially implement NV and/or SDN.

Given the embryonic and rapidly changing nature of NV and SDN, any implementation plan will likely evolve over time. The process that a company such as *GottaChange* should take to evaluate solutions and possibly implement one or more solutions should include most if not all of the following steps:

1. Define NV and SDN
2. Identify the Primary Opportunities
3. Identify the Key Metrics
4. Define the Scope of Possible Solutions
5. Evaluate NV and/or SDN Solutions
6. Test and Certify Solutions
7. Integrate with the Existing Environment
8. Educate the Organization
9. Evaluate Professional Services
10. Eliminate Organizational Resistance
11. Perform a POC
12. Obtain Management Buy-In

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

Jim Metzler has a broad background in the IT industry. This includes being a software engineer, an engineering manager for high-speed data services for a major network service provider, a product manager for network hardware, a network manager at two Fortune 500 companies, and the principal of a consulting organization. In addition, he has created software tools for designing customer networks for a major network service provider and directed and performed market research at a major industry analyst firm. Jim's current interests include cloud networking and application delivery.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at *jim@webtorials.com* or Steven Taylor at *taylor@webtorials.com*.

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Customer Driven Innovation

Application Delivery

- SLB
- ADP
- GSLB
- ADC

Security

- Web App Firewall
- DNS App Firewall
- SSL Intercept
- DDoS
- AAM

Cloud, SDN & Next Gen Networking

- SDN
- aCloud
- CGNAT
- IPv6

Thunder Series

Application Service Gateways

Next-generation Application Delivery Controllers

Powered by ACOS

www.a10networks.com

The Application Fluent Data Center Fabric

Introduction

The rise of virtualization and cloud computing requires the selection of a best-of-breed data center switching solution as part of an enterprise's overall data center strategy. And at the heart of this strategy is the need to deliver a high quality user experience with new virtualized applications, including video, on new devices such as smart phones and tablets. However, the traditional 3-layer networks designed for a client/server communication model cannot meet the requirements of these new applications and devices, nor can it address the new requirements of virtualized servers and desktops.

Application Fluency for the Data Center

Resilient Architecture

- Simplified 10 & 40 GigE network with low latency and ready for 100 GigE
- Multi-path data center network extends between data center sites and to public cloud
- Supports definition of virtual data centers
- Ready for storage convergence with lossless Ethernet

Automatic Controls

- Application profiles ensure that the network is aware of application provisioning, security and QoS requirements
- The network will automatically sense virtual machine location and movement
- The network will automatically adjust to VM motion within and between data center sites

Streamlined Operations

- Applications are automatically provisioned
- Core switches automatically configure top of rack switches
- Converged management for data center network and virtual machine mobility
- Low power consumption

The Alcatel-Lucent Mesh

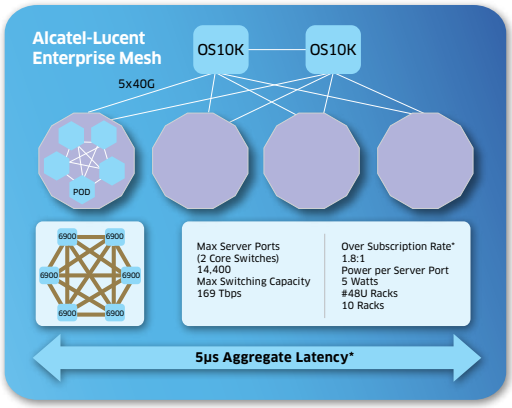

Alcatel-Lucent provides a unique Application Fluent approach to maximize the benefit from virtualization technologies for servers, the desktop, as well as the network. Alcatel-Lucent's application fluent data center fabric can scale from several hundred to over 14,000 server facing ports while keeping aggregate latency at 5ms, and can automatically adapt to virtual machine movement no matter which server virtualization platform is used.

The Alcatel-Lucent Virtual Network Profile (vNP), embedded in the Alcatel-Lucent Mesh, includes the critical information the fabric needs to understand each application, including provisioning requirements, security profiles, and expected quality of service levels. With this knowledge, the network can manage applications as services, including automatically discovering the location of each virtual machine, modifying the network configuration to follow virtual machine moves and providing an integrated view on visibility on VM movement and current location from a network perspective.

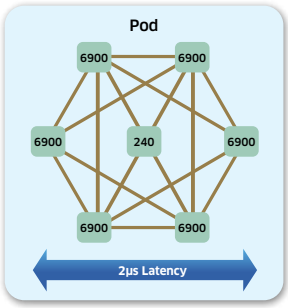

Application fluency in the corporate data center includes its transformation into a multi-site private cloud by extending layer 2 connectivity between data center sites and allowing for seamless delivery of public cloud-based services on the corporate network.

The Alcatel-Lucent Mesh enables enterprises to provide a high quality user experience with mission critical, real-time applications, and to improve agility in deploying new applications while significantly reducing data center costs.

Open Ecosystems and Market Success

Alcatel-Lucent Enterprise is committed to open standards, allowing enterprises to select best-of-breed suppliers for their complete data center solution: servers, storage, data center fabric, and data center interconnect.

- Winner: Best of Interop 2011 for Data Center Switching and Storage
- Data center ecosystem partners include Emulex, NetApp, VMware, Citrix, and QLogic
- Participant in IEEE sponsored Shortest Path Bridging interoperability test with Avaya, Huawei, Solana and Spirent
- Over 20 million Ethernet ports shipped

For More Information

Alcatel-Lucent Data Center Switching Solution
Alcatel-Lucent Application Fluent Networks
Alcatel-Lucent Enterprise

**Assuming Server to Server Traffic 70% within a Pod, 20% between Pods and 10% Via Core*

Alcatel·Lucent
Enterprise

Agile, Automated Cloud Services

Avaya's Software-Defined Data Center (SDDC) framework offers a simple five-step process for deploying cloud-based services in a matter of minutes. This framework breaks-down the frustration, complexity, and lack of agility that's typically been the norm when building and deploying business applications. Avaya replaces the complicated, independent provisioning steps between the compute, storage, and networking teams with our simplified, orchestrated, and automated workflow. With the SDDC, compute, storage, and network components are automatically combined, customized, and commissioned through a common orchestration layer.

The Avaya SDDC framework is based on the following components:

- **Avaya Fabric Connect technology** as the virtual backbone to interconnect resource pools within and between Data Centers with increased flexibility and scale
- **An Avaya OpenStack Horizon-based Management Platform**, delivering orchestration for compute (Nova), storage (Cinder/Swift) and Avaya Fabric Connect networking (Neutron)
- **Open APIs into Avaya Fabric Connect** for ease of integration, customization and interoperability with other SDN architectures

Traditional methods of configuring network, storage, and virtualized servers could take months and involve several complicated independent steps. Avaya's SDDC framework leverages OpenStack, an open-source cloud operating system. Now Data Center administrations can spin up virtual machines, assign storage, and configure networks through a single GUI. OpenStack provides a control layer that sits above all the virtualized resources within the Data Center, allowing these to be orchestrated – as a single service entity – through a set of common interfaces and a common dashboard.

Avaya Fabric Connect enhances and complements the OpenStack environment by removing the restrictions of traditional Ethernet Virtual LAN/Spanning Tree-based networks. Fabric Connect turns a complex, rigid, and un-scalable model of building network services into a dynamic, flexible, and scalable one. It facilitates the unrestricted movement of virtual machines inside the OpenStack orchestration environment, within and between Data Centers. It also enables the interconnection of old and new resources across the service chain with greater speed and agility.

In summary, with a combination of its Fabric Connect and intelligent orchestration software, based on OpenStack, Avaya is enabling simple and agile **automated** service delivery for applications and users across any combination of physical and virtual components in an evolutionary manner.

Learn more at avaya.com/sdn

Advantages of the Avaya Software-Defined Data Center Architecture

- **Reduced Time-to-Service**: Cloud services enabled in minutes, in a few simple steps.
- **Simplified Virtual Machine Mobility**: End-point provisioning to enable Virtual Machine mobility within and between geographically dispersed Data Centers.
- **Multi-Vendor Orchestration**: Coordinated allocation of compute, storage, and networking resources via a single interface to streamline the deployment of applications.
- **Openness**: APIs ease integration and customization with Fabric Connect, and interoperability with other Software-Defined Networking architectures.
- **Scale-Out Connectivity**: Services scale to more than 16 million unique services, up from the four thousand limitation of traditional Ethernet networks.
- **Improved Network Flexibility**: Overcomes the current Virtual LAN challenges to deliver a load-balanced, loop-free network where any logical topology can be built with simple end-point provisioning.

Top 10 things you **need to know** about Avaya Fabric Connect

(An enhanced implementation of Shortest Path Bridging)

A completely new way to build networks, Avaya Fabric Connect delivers a simplified, agile and resilient infrastructure that makes network configuration and deployment of new services faster and easier. A standards-based network virtualization technology based on an enhanced implementation of IEEE 802.1aq Shortest Path Bridging and IETF RFC 6329, Avaya Fabric Connect combines decades of experience with Ethernet and Intermediate System-to-Intermediate System (IS-IS) to deliver a next-generation technology that combines the best of Ethernet with the best of IP. Avaya Fabric Connect creates a multi-path Ethernet network that leverages IS-IS routing to build a topology between nodes dynamically. Traffic always takes the shortest path from source to destination, increasing performance and efficiency.

Avaya Fabric Connect is an industry unique solution that offers a number of characteristics that set it apart from competing offers. The following Top 10 list below will give you a sneak peek of the advantages Fabric Connect offers:

1 It is more than just a Spanning Tree Replacement

Avaya's dynamic, real-time, service-based Fabric Connect technology is one of the most advanced network virtualization solution on the market today. Going beyond simple L2 multi-pathing capabilities, Avaya Fabric Connect delivers the full breadth of desired integrated services including Layer 2 virtualized services, Layer 3 virtualized services (with multiple Virtual Routing and Forwarding instances), and fully optimized routing and multicast services.

As a result, Fabric Connect enables businesses to gradually migrate away from a host of legacy overlay technologies (such as STP, OSPF, RIP, BGP and PIM) and to enable all services with a single technology – delivering unprecedented levels of network simplification.

2 It's for more than just the Data Center

While many network virtualization technologies are designed exclusively as Data Center technologies, Avaya Fabric Connect extends network-wide, providing a single service end-to-end delivery model. With Fabric Connect you can extend the power of virtualization into the campus and into geographically dispersed branch offices. Services can then easily be deployed via simple end-point provisioning where servers attach and where users attach, thereby increasing speed and agility.

3 It accelerates time-to-service through edge-only provisioning

Fabric Connect allows new services or changes to services to be implemented at the edge of the network – eliminating error-prone and time-consuming network wide configuration practices. Now, add new services or make changes to existing services in days rather than weeks or months. Fabric Connect also offers new levels of flexibility in network design. It allows any logical topology to be built, whether it is Layer 2, Layer 3, or a combination of the two – anywhere where there is Ethernet connectivity. Eliminate design constraints and have the freedom to build services wherever and whenever needed on demand.

4 It offers inherent Data Center Interconnect capabilities

Customers are demanding network virtualization solutions that are not confined to the four walls of the Data Center. Avaya Fabric Connect offers a single end-to-end service construct that can extend between multiple geographically dispersed Data Centers without requiring any overlay protocols or complex protocol stitching. This allows for resource sharing, seamless VM mobility and true active, active connectivity between Data Centers and any other Ethernet-connected enterprise location.

5 It delivers PIM-free IP Multicast that is scalable, resilient and easy to manage

IP Multicast is making a come-back. Many technologies such as next-generation video surveillance, IPTV, digital signage, desktop imaging, financial applications and some network overlays are reliant on Multicast protocols. Avaya Fabric Connect offers a scalable, reliable and efficient way of supporting IP Multicast Routing, without the onerous requirement of configuring, deploying, and maintaining a complex PIM overlay.

Imagine a Multicast network without RPF checks, rendezvous points and complex configuration. Enable Multicast at the edge of the network only, while offering increased scale and performance of the multicast applications. Eliminate your PIM induced headaches forever!

6 It offers inherent multi-tenant capabilities

Avaya Fabric Connect offers integrated Virtual Routing and Forwarding Instances. This allows for private IP networks to be set up quickly and easily across the fabric-enabled network without requiring any overlay protocols. These IP networks can reflect anything from different departments or entities in a traditional multi-tenant environment to separating different types of users (wireless guests, executive access) and even isolating traffic types for security and/or regulatory compliance (i.e. banking transactions for PCI DSS compliance, medical imaging devices in a hospital). The best part is rather than complex configuration, these isolated networks can be deployed quickly and easy at the network edges with just a couple of lines of configuration.

7 It offers "lightening fast" recovergence times (sub-second)

The elimination of overlay protocols has a profound impact on the ability for the network to reconverge. Avaya Fabric Connect customers are experiencing recovery times of less than 50 milliseconds - network-wide - for core, link, or node failures. This represents a vast improvement over large OSPF routed cores and massive improvement when compared to average recovery times in PIM-based Multicast networks.

8 It scales to 16 million unique services

Many network virtualization technologies are based on VLAN virtualization which limits them to the 4096 ceiling. Avaya Fabric Connect, based on the Shortest Path Bridging standard, utilizes a 24-bit header allowing it to scale up to 16 million unique services.

9 It offers proven interoperability with other vendors SPB implementations

Avaya is committed to delivering an open and interoperable solution to market. We have been actively participating with other vendors to demonstrate Shortest Path Bridging interoperability through a series of public tests. The most recent interoperability test was conducted at Interop 2013 in Las Vegas with major industry vendors Alcatel Lucent, HP, and Spirent.

10 It is an important foundation to your SDN strategy

When it comes to SDN, Avaya's strategy is to first eliminate network complexity in order to provide a simple and flexible network foundation. Rather than adding overlays or additional protocols, and creating even more complexity than what we have today, Fabric Connect first streamlines the network then automates it though OpenStack-based orchestration functionality (via a Neutron plug-in). It provides a simplified and proven way to automate the service delivery process and evolve to the Software Defined Network of the future.

Learn more about Avaya Fabric Connect:

Avaya Fabric Connect - video on YouTube, Considerations for turning your network into a Fabric - Packet Pushers podcast, Network Virtualization Using Shortest Path Bridging and IP/SPB – White Paper

SOFTWARE-DEFINED NETWORKING

Software-Defined Networking (SDN) is a transformative network architecture that is reshaping the telecommunications landscape. SDN offers network operators the opportunity to better **monetize** and **optimize** their networks, simplify and automate network operations to reduce OPEX, improve agility to rapidly introduce and differentiate new service offerings to prevail in the increasingly competitive landscape.

Figure 1 depicts the SDN architecture, which is characterized by:

- **Programmability –** Enable unprecedented network control
- **Centralized Intelligence –** Logically centralize network state to optimize resources and construct end-to-end services under granular policy control
- **Abstraction –** Decouple business applications from the underlying network infrastructure, while allowing intelligent software to operate across multiple hardware platforms
- **Openness –** Standard interfaces (including OpenFlow™) achieve multi-vendor interoperability and software

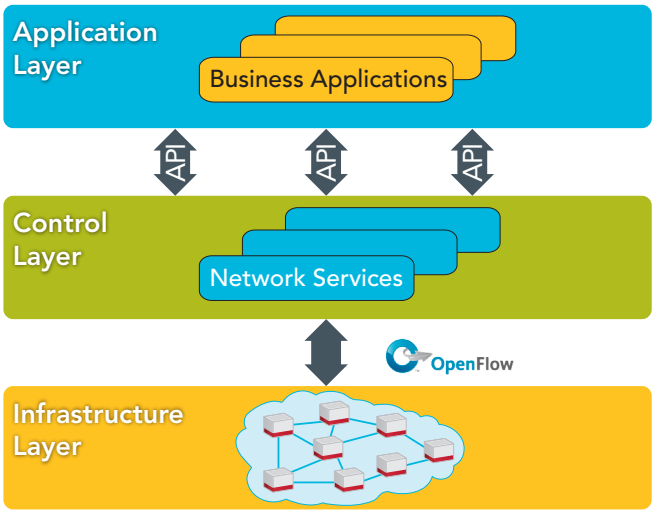

Figure 1. ONF SDN Architecture

Ciena is embracing SDN and leading the charge toward multi-layer, carrier-scale SDN in the Open Networking Foundation (ONF), where Ciena is a founding member and leading contributor. SDN is a key component of Ciena's OPn architecture, which drives down the networking cost curve with converged packet-optical architecture and highly intelligent software functionality.

For more information:

OPn: ciena.com/technology

ONF: opennetworking.org

Ciena's view of SDN emphasizes two key concepts:

- **Autonomic Operations Intelligence –** Streamline operations through automation, resource optimization, and end-to-end service delivery. Grow profit and revenue with real-time analytics: capitalizing on Ciena's experience powering the most intelligent large networks on the globe
- **Expansive Openness –** Embrace open standards and software architectures to enable network operators to innovate and differentiate their businesses

An initial step toward SDN is available today with Ciena's V-WAN Network Services Module, delivering performance on demand to optimize data center interconnection. In concert with our customers and Research & Education partners, we are introducing an ambitious carrier-scale WAN test bed to validate and demonstrate autonomic operations intelligence and expansive openness. Through these efforts—along with our leading role in the ONF, MEF, and related standardization activities—Ciena is shaping the future of multi-layer, carrier-scale SDN.

Learn more at ciena.com/technology/sdn and stay tuned for exciting announcements from Ciena in the months to come!

Cisco Network Virtualization Platform Designed to Automate Application Provisioning and Deployment

Cisco Overlay Approach Focuses on Simplifying and Automating IT Tasks

Network Virtualization (NV) has rapidly emerged as a fundamental enabler for cloud networks and highly virtualized, multi-tenant data centers. NV helps overcome many of the initial obstacles to cloud networking, including addressing network complexity, scalability issues and constraints on workload mobility. But the real promise of NV and SDN leads to orders of magnitude improvements in the automation of IT tasks focused on application deployment, provisioning, optimization and service delivery. The end result will be applications that scale on-demand, vastly improved resource utilization, and much more agile enterprises whose IT organizations respond to changing business requirements in minutes or less.

From Virtual Networks to an Application Centric Infrastructure

The Cisco Nexus 1000V virtual networking platform is a complete overlay/cloud networking solution that includes virtual switching, routing, integrated virtual security services, application delivery services, VXLAN overlay tunneling, network monitoring and analysis, and hybrid cloud integration. Cisco now takes advantage of the simplified, more flexible virtual network by integrating with a range of network automation and orchestration tools running on all major cloud and server platforms, from VMware vCloud Director, to Microsoft System Center, OpenStack and Cisco's own UCS Director.

In June, Cisco augmented its virtual networking and automation capabilities with a new vision for the data center: an Application Centric Infrastructure (ACI). ACI is a cloud and data center fabric designed around application policies that will further simplify and automate the provisioning and deployment of applications, as well as configuring and optimizing the network and network services for application-specific requirements.

The resulting ACI capabilities will further reduce IT costs by automating nearly all application and network provisioning tasks, while allowing IT to be dramatically more responsive to changing business needs by accelerating application deployment, policy changes and fundamentally improving resource allocation and efficiency. The ACI Fabric will be ideally designed for both physical and virtual applications, and also removes obstacles to scale and network visibility that competitive virtual overlay solutions introduce. Nexus 1000V technology and key components of the Cisco virtual network architecture will be part of the ACI fabric.

For More Information

Learn more about the Cisco Nexus 1000V virtual networking portfolio: http://cisco.com/go/1000v

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**

Making **Sense** of **SDN**

and putting it to work for you

The Cloud is here, and things are changing – fast. As the engine of the hyper-connected world, cloud computing has changed the landscape of business , bringing with it a new wave of untapped opportunities. Modern consumers are increasingly drawn to brands that can offer a more savvy and immersive high-tech experience. New technologies can not only offer a deeper understanding of customers than ever before possible, but also promise to refine business operations with an analytical precision that can redefine operational efficiency. While the hyper-connected world can offer tremendous impact, in this new era the opportunities belong to those with the greatest mastery of technology, those who can execute with agility and maximize the impact of the latest innovations.

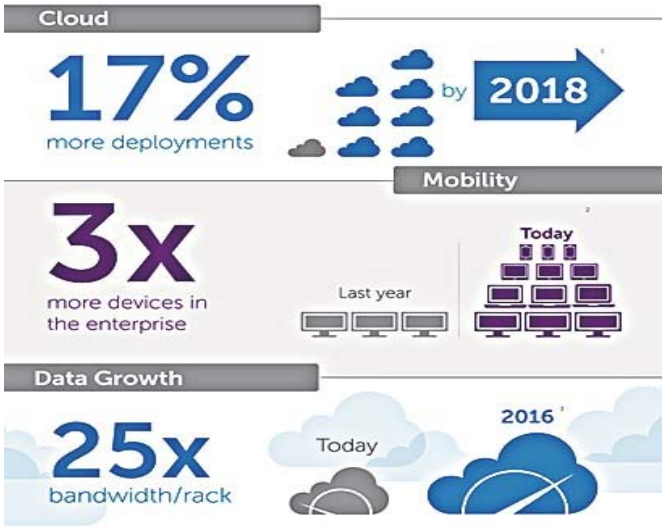

Software Defined Networking is one of the most significant new technologies as it holds the key to the next wave of automation and dynamic integration, igniting business and operational agility while empowering a deeper end-user experience.

At Dell, SDN is not a confusing choice, it is baked into every data center platform we sell. Dell Active Fabric offers a single high-performance architecture that improves the performance of legacy applications while fully preparing enterprises for the rigorous & intelligent demands of next-generation applications. Our robust software suite empowers IT staff to immediately take advantage of SDN, offering out of the box wizard-based design, fully automated deployment and single-pane of glass operations - designed from the ground up to provide a new lifecycle-based approach to highly-automated operations that can redefine enterprise IT.

As the world's largest startup, Dell is embracing the latest innovations and invite you to join in – our staff of experts is waiting to show you how to put SDN to work for you and realize its benefits, today.

Visit **DellNetworking.com** to learn more about Dell SDN solutions or to contact a Dell Representative

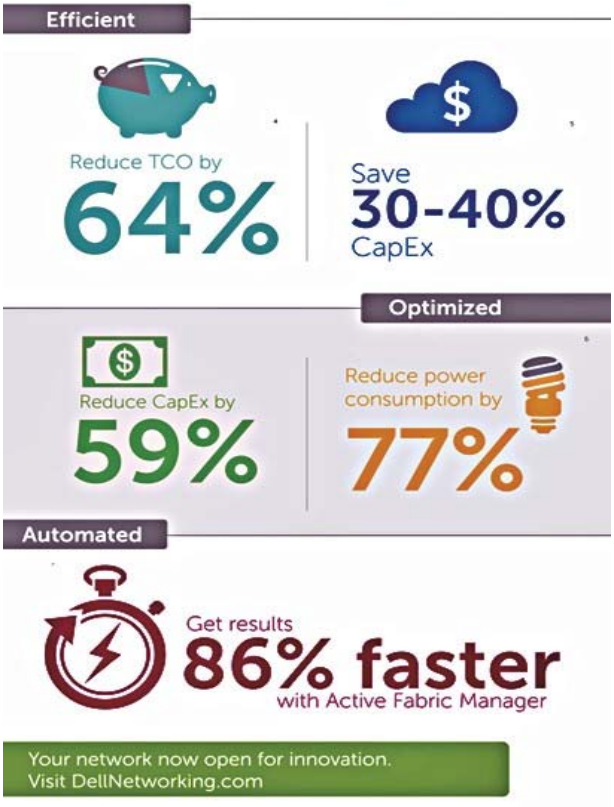

Software-Defined Networking

Are your management tools prepared?

Software-Defined Networking (SDN) and Network Virtualization (NV) are quickly becoming priorities because of the promise to dynamically manage traffic loads while lowering costs in response to changing business requirements…

Are you prepared for this evolution?

EMC understands these challenges. Designed to manage physical, virtual and cloud environments, the EMC Service Assurance Suite helps IT operations teams manage infrastructure across each phase of this evolution.

Empower your IT operations team to visualize, analyze, and optimize your service-delivery infrastructure. Learn more at www.emc.com/sa.

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, email us at asd@emc.com or call 866-438-3622.

EMC²

Extreme Networks Open Fabric as the Foundation for SDN

The Extreme Networks ***Open Fabric*** framework includes the key attributes of the data center network, such as high speed, low latency switching, lossless connectivity, multiple paths for resiliency, low power use, automation capabilities, and open standards that are also important to the campus, enterprise and other mission critical networks that require high performance, high scale and resiliency.

Figure 1 Extreme Networks Open Fabric

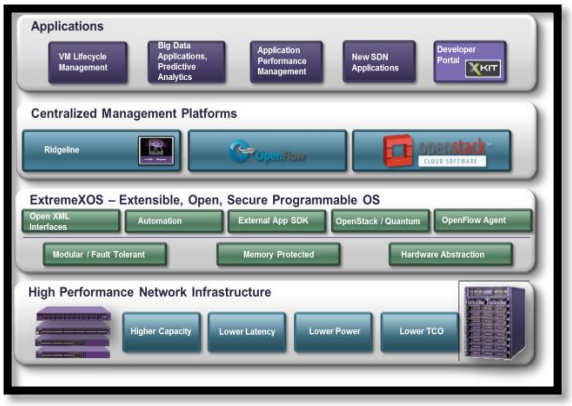

Critically important to the Open Fabric is ExtremeXOS®, the network operating system that delivers the consistent set of features across all platforms while ensuring the security and performance of the Open Fabric. ExtremeXOS is modular, extensible, and has integrated security, while providing a single linux-based OS from the core of your network all the way down to the edge. In essence, ExtremeXOS is the system wide ***network abstraction*** layer that allows both seamless introduction of new hardware while opening up the network to management platforms and applications.

The Open Fabric and Extremes are the foundation of the Open Fabric SDN framework. The Open Fabric provides the attributes for the high performing infrastructure while ExtremeXOS abstracts the intelligence of the network, uniquely bonding together to create the Open Fabric SDN framework. The ***network abstraction*** of the Open Fabric SDN approach is found at the ExtremeXOS layer and includes SOAP/XML open APIs, the OpenFlow protocol, CLI and scripting, and the operating system itself. Again, note that network abstraction is available on all Extreme Networks platforms, from edge to core, from 1GE to 100GE. The multitude of network abstraction components allows many different methods for applications and management platforms to access network intelligence, including OpenFlow controllers from NEC and Big Switch Networks, and the OpenStack cloud orchestration system for provisioning storage, compute and network elements.

Figure 2 Extreme Networks Open Fabric SDN

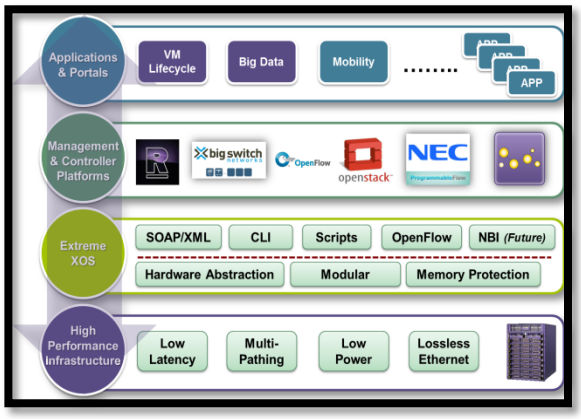

The Extreme Networks Open Fabric SDN strategy therefore extends to include technology partners and systems that leverage the network abstraction capabilities provided by ExtremeXOS.

Open Fabric SDN – Inclusive Approach to SDN

From a pure networking standpoint, The Extreme Networks Open Fabric SDN approach includes OpenFlow, Open API's and Network Virtualization as 3 main technology areas inclusive of a broad definition of SDN.

OpenFlow

The OpenFlow protocol is one of the leading new technologies driving the SDN market. OpenFlow is an open standards-based specification led by the Open Networking Foundation.

Figure 3 OpenFlow Protocol

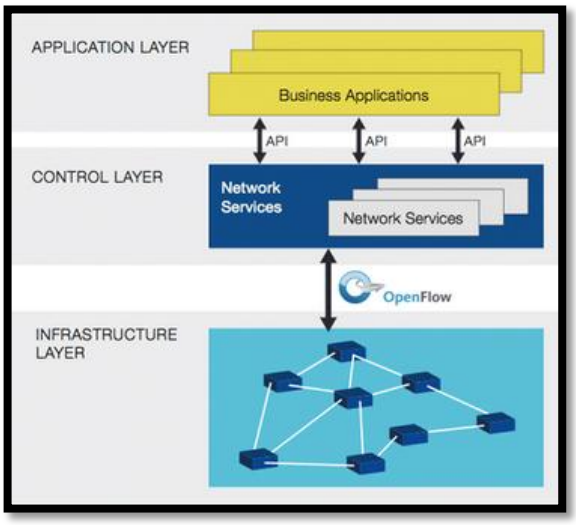

The Open Networking Foundation (ONF) defines OpenFlow: "The physical separation of the network control plane from the forwarding plane, and where a control plane controls several devices."

Open APIs

Using industry standard messaging protocols allow client and server systems to exchange configuration, statistics and state information. OpenStack is a cloud management and orchestration system that uses API's to provision and manage storage, compute and network resources. Extreme Networks has created a software plugin that allows the OpenStack platform to access the network abstraction layer using open API's (SOAP/XML).

As an example, the XML server (XMLD) shown in Figure 4 is responsible for providing a gateway between the external interface and the switch modules. It enforces security; wraps, unwraps, and validates messages; and performs the mechanical translations of results from the modules to the client machine. The XML APIs use the SOAP protocol over telnet/SSH or HTTP/HTTPS to exchange XML configuration messages between the client machine and the ExtremeXOS switch modules.

Figure 4 Extreme Networks Open APIs

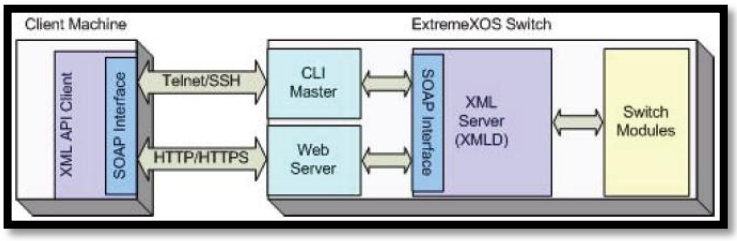

"Open API's enable applications and management systems to directly access the network abstraction layer to manage the control, data and management planes of the infrastructure."

Network Virtualization

Network Virtualization Overlays, commonly called Network Virtualization (NV) or just Overlays, includes a virtual logical network construct over a physical topology. Overlays still require a high performing, robust physical infrastructure and can be leveraged at various networking layers, including:

- Network Virtualization at Layer 2 with VLANs and MPLS
- Network Virtualization at Layer 3 with MPLS VRF's and Virtual Routers (VR) as well as VXLAN and NVGRE for the transport of Layer 2 protocols.

Also, using Open API's and OpenFlow can enable custom applications to create an overlay as well.

Figure 5 Network Virtualization Overlay with VXLAN

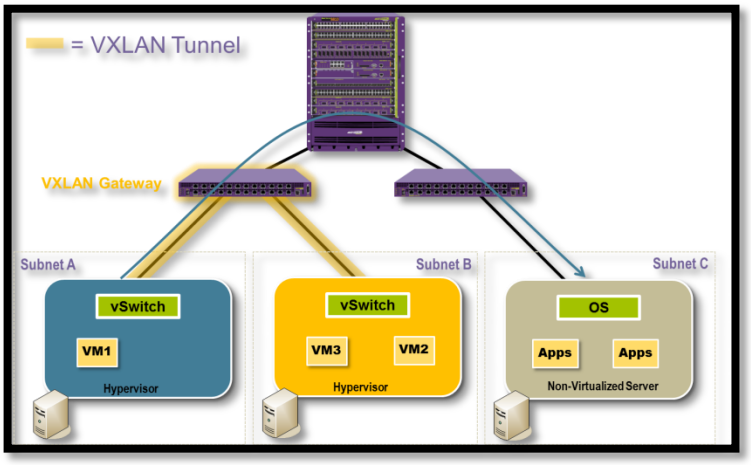

"Network Virtualization Overlays include logical overlays from virtualized server and switching systems that may also include virtualized layer 4-7 services."

Extreme Networks: The Inclusive Approach to SDN - Summary

This inclusive approach to SDN allows a complementary mix of industry and customer perspectives, enabling multiple different SDN strategies. From OpenFlow to Open APIs to Network Virtualization, the Extreme Networks Open Fabric SDN framework enables an inclusive approach to SDN that leverages the ExtremeXOS network abstraction capabilities of a single binary OS ubiquitous from edge to core.

Figure 6 Inclusive SDN Approach

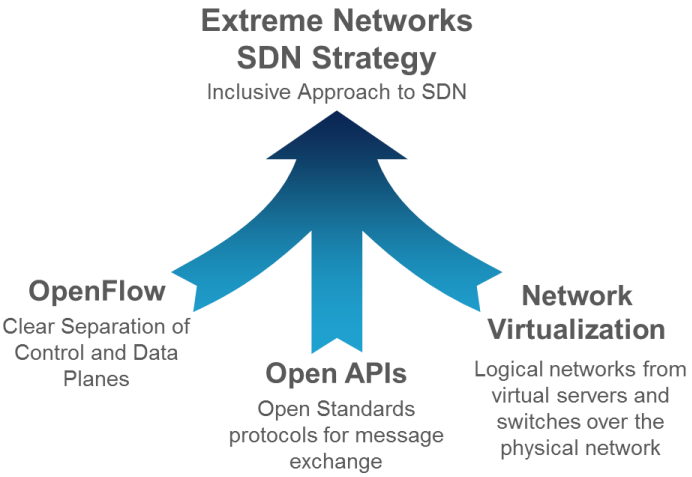

HP is leading Software-defined Networking

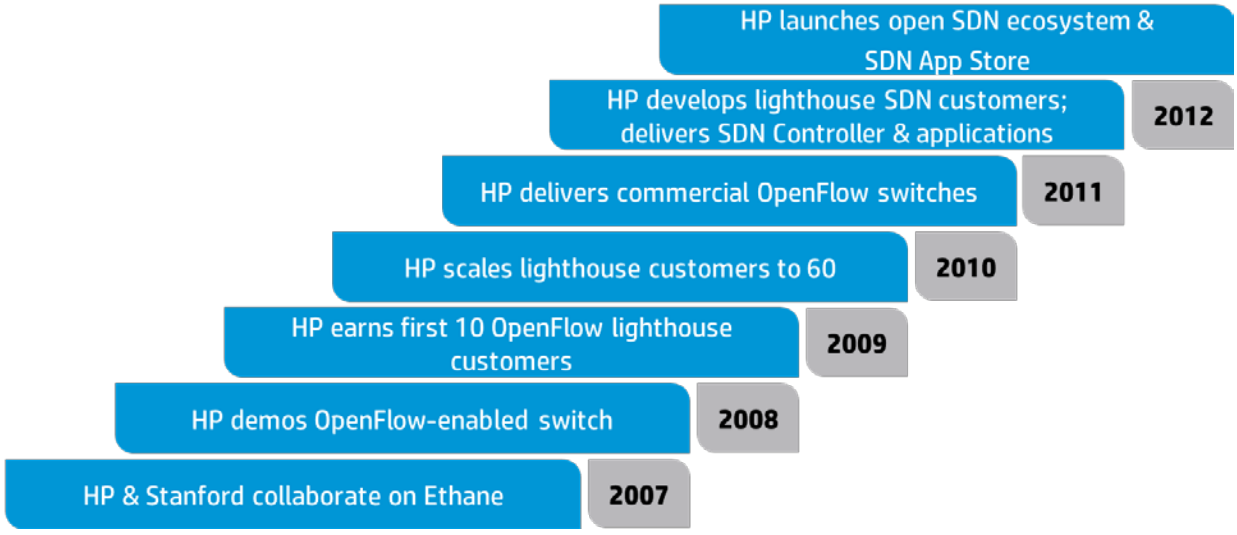

HP Open SDN portfolio, SDK, App Store enables Ecosystem

Programmable network aligned to business application delivers agility

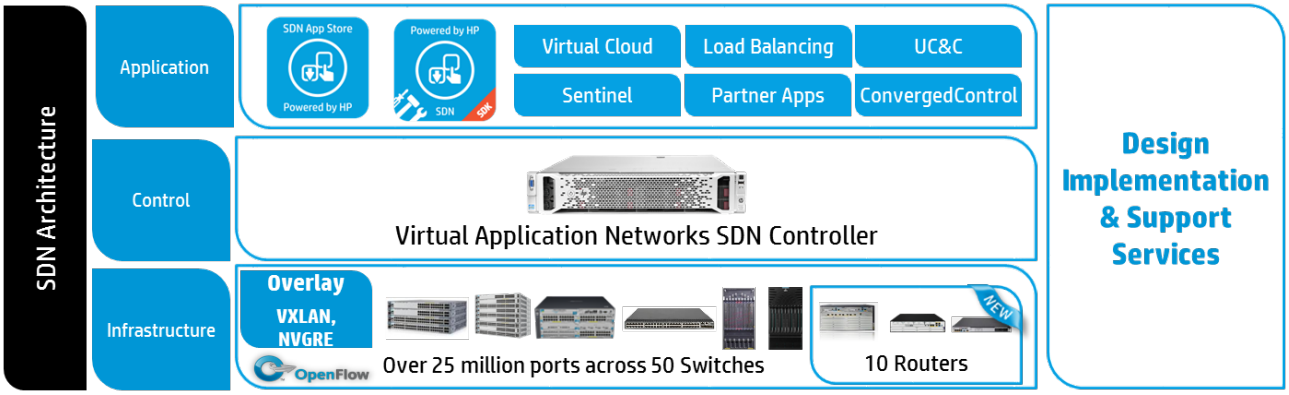

ProgrammableFlow

Empowered by Innovation NEC

Kanazawa University Hospital

Reaping the benefits of a successful production SDN deployment

Keisuke Nagase, M.D., Ph.D., is a professor of medicine, healthcare administration, and medical informatics at Kanazawa University in Kanazawa, Japan. He also serves as vice-director in charge of budget/management and director, department of corporate planning for Kanazawa University Hospital. The hospital recently began overhauling its cumbersome network infrastructure by deploying NEC's ProgrammableFlow® solutions, a network solution based on OpenFlow technology. With 839 beds and 33 clinical departments, Kanazawa University Hospital is one of the largest and oldest teaching and research hospitals in the country.

金沢大学 KANAZAWA UNIVERSITY

Interview excerpt from www.SDNcentral.com

What kind of IT environment do you have at the hospital? What is the rough annual IT spend?

Our network is essential to the day-to-day business of providing patient care. From electronic medical records to medical equipment, IT is critical for everything in the hospital. The patient management system and billing system are the largest in scale in terms of IT, but everything is connected – ICU, operating rooms, medical equipment. We spend roughly $600M Yen ($6M-$7M USD) per year on IT.

What are the major IT problems you have had to solve at the hospital?

As an educational hospital, we are large and armed with innovative new healthcare technologies. The problem is, many computer networks have been deployed independently because each medical equipment manufacturer and vendor wanted to simplify the environment around their equipment.

When I moved to this hospital from a previous position, I faced a chaotic situation. Information technology is not our core business, patient care is. As a result, human resources for information system management were limited for a long time. The existing network was high risk and high cost, and poor control over the network led to many unfavorable incidents and accidents. For example, packet storms caused by large-scale loops would interrupt daily jobs for four hours.

Even daily operations were challenging. Technologies evolve rapidly in the medical field, and doctors often try new equipment. Connecting this equipment to the network involved changing settings and verifying connections, and sometimes even rewiring, putting a considerable strain on the hospital's budget. A network that requires setting changes and rewiring every time a new piece of equipment is connected cannot be called stable. The other issue is slow reconfiguration of the network due to the processes in place, adding a new piece of equipment could take 3 months including time to initiate the contract for the add/move/change.

Why did you decide to use OpenFlow technology to address these problems?

We were looking for a more agile solution that had the same or lower risk as our existing network, at the same or lower cost. That was OpenFlow. We did not select SDN as a result of passion for a new technology. Our business is not IT -- our system is directly related to the life or death of our patients. Education, research and healthcare are our business.

There was no breakthrough or epoch-making technologies in SDN, we believe, but rather an innovation of philosophy. We wanted to be free from any specific manufacturer. We selected OpenFlow because we need it. We consider OpenFlow switches and controllers to be stable.

"We did not select SDN as a result of passion for a new technology. Our business is not IT—our system is directly related to the life or death of our patients."

"Now we are enjoying rapid recovery time and flexibility in a network with reduced maintenance and operational costs. The time for recovery was reduced to seconds rather than minutes."

As you know, many manufacturers are modifying their existing products to be OpenFlow enabled. With such consideration, we felt the stability of OpenFlow switches and controllers to be the same or better than conventional switches, even at their worst. Because the software is simple, it is essentially more stable than our legacy technology. The only exception is if an incompetent person codes the applications running on the controller.

How did you introduce OpenFlow to the existing system?

We added a new general research building to our campus more than one year ago. Each clinical department and its corresponding university department moved to the new building. In the new building, four independent networks were requested to be deployed, and the existing network also needed to be deployed to the new building. We introduced SDN/OpenFlow in the new building to eliminate complexity of network.

We thought the deployment of SDN to the new building was quite a good opportunity to evaluate SDN. Multiple in-house LANs are required to implement SDN, making the situation a good test case for network slicing with SDN. By adopting SDN in the new building, we also decided it would be a good test for migration from our legacy network to SDN.

Even if the SDN network failed somehow, the effect would be limited because the new building is connected to the old hospital building and legacy network via a corridorwe ran a parallel network initially that the staff could still access in different rooms but only a short walk away. We concluded adopting SDN/OpenFlow in the new building would at worst be the same risk, same cost.

We integrated the existing independent network using SDN/OpenFlow in the new research building. With OpenFlow, the network within the building was kept simple, and our new virtual tenant networks are merged with the existing hospital network using link aggregation.

"…the operational expenses and maintenance cost has reduced markedly. I estimate a savings of 80% on my operational expenses."

Why did you choose NEC ProgrammableFlow switches and controllers?

An NEC network Systems Engineer (SE) understood the deeply unstable situation of our network, and he suggested we use OpenFlow. NEC was the only supplier of production quality OpenFlow switches at the time of our contract, and they have been our partner for many years. The NEC SE built a good relationship with the assistant professor in charge of the hospital information system.

NEC installed two ProgrammableFlow controllers and 16 switches in our new building. It allowed us to install devices one floor at a time and expand gradually and safely. We could manage each department's LAN without impacting our existing network.

With NEC's ProgrammableFlow solution, the entire network is managed like a large virtual switch, making an independent virtual network. Our OpenFlow switch was implemented as edge (floor) switches. We have full mesh wiring between switches. In the center, the OpenFlow network is connected under the existing L3 switch (core switch) using link aggregation, so as to be configured as single L2 switch network from L3 switch.

For redundancy, we have two sets of OpenFlow controllers. For OpenFlow switches, we have two sets in center side, two sets in the new building side, and two sets on each floor, for a total of 16 sets. We also have two sets of secure channel switches—in the system operation center and the new building. NEC required only one month to get the new network up and running.

How does the SDN network compare in cost and price?

The acquisition cost of the hardware was almost the same as the legacy network. However, the operational expenses and maintenance cost has reduced markedly. I estimate a savings of 80% on my operational expenses, including reduction in staff hours required to manage the network. We also expect that the price of OpenFlow switches and OpenFlow controllers will be reduced further as a result of competition in the market. Furthermore, with the flexible configurability of OpenFlow, a full mesh configuration is not required, and our next phase will be in realized in less cost per switch.

"I can now provision the network after new equipment installations or equipment moves in minutes instead of the 3 months it used to take."

What benefits have you seen from deploying SDN?

As I've mentioned, I've seen significantly lower maintenance costs, allowing me to make much better use of my human resources at the hospital. More importantly, I now have the ability to perform moves, adds and changes to my network much faster than before. I can now provision the network after new equipment installations or equipment moves in minutes instead of the 3 months it used to take. This is achieved via ProgrammableFlow, leveraging the OpenFlow protocol, which will automatically connect the equipment to the right network instantly.

So, what's your final evaluation of SDN and NEC's ProgrammableFlow solution?

I would say that the network has been successfully delivering critical patient health records as well as MRI and CT scan data, reliably and efficiently. With this experience we decided to expand our ProgrammableFlow OpenFlow network to the entire hospital network over the next two years. We also expect to refresh and clean up our IP address space from a chaotic situation utilizing flexibility we gained from our SDN network.

In summary, I would declare our SDN deployment highly successful and would recommend other medical centers take a serious look at deploying SDN and reaping the significant benefits today.

"I would declare our SDN deployment highly successful and would recommend other medical centers take a serious look at deploying SDN and reaping the significant benefits today."

Key Features of the NEC ProgrammableFlow Networking Suite:

- **Drag and drop network design**: The GUI interface to the ProgrammableFlow Controller includes the familiar CLI found on most routers and switches today, so with minimal training a network admin can easily point and click to design an entire network from the single pane provided by the ProgrammableFlow Controller. This can radically reduce network programming and design time and errors caused previously by human intervention.
- **VM mobility**: With the ability to readily direct traffic throughout the data center—or throughout multiple data centers, it is possible to better manage all of the resources in a data center. For example, in NEC's own data centers in Japan, where they have recently implemented the ProgrammableFlow Fabric, it has enabled them to spread traffic between East and West Japan, offloading servers in East Japan that were nearing capacity, and postponing purchase of new servers, for a substantial saving. VM Mobility also enabled Nippon Express to complete a data center consolidation move that normally would have taken 2 months down to 10 days.
- **Bandwidth monitoring and traffic flow visualization**: This feature of the ProgrammableFlow Controller provides performance monitoring of network flows and centralized management of network traffic, reducing bottlenecks and enabling smooth, streamlined network operations with substantially improved network admin productivity.
- **Secure, multi-tenant networks**: Secure, multi-tenant networks from the ProgrammableFlow Controller enables customers like Genesis Hosting to expand their service offering with new sources of revenue potential. Genesis also reports software engineering investments were reduced by 100 hours each month with the advancements provided by ProgrammableFlow multi-tenancy.
- **Automation and administration of business policy to network management**: With network services aligned with business policy, automation such as prioritizing classes of applications or specific applications over other enterprise activity during peak loads is now possible with the ProgrammableFlow Network Suite, with multiple paths provided automatically. These capabilities offer significant value, particularly to enterprises engaged in heavy transaction loads.
- **Load balancing**: Traditional networking protocols often lead to performance-reducing bottlenecks. ProgrammableFlow uses path selection algorithms to analyze traffic flow across the network, check all available paths, and customize traffic flows to maintain performance and fully utilize network capacity. This increases the utilization of the network and improves application performance.

A fully optimized, automated, cost-effective networking solution, Netsocket Virtual Network provides end-to-end virtual networking, unified network management, real-time network service analytics with intelligent network remediation as well as superior interoperability with legacy routed networks.

Virtual network distributed on commodity x86 hardware

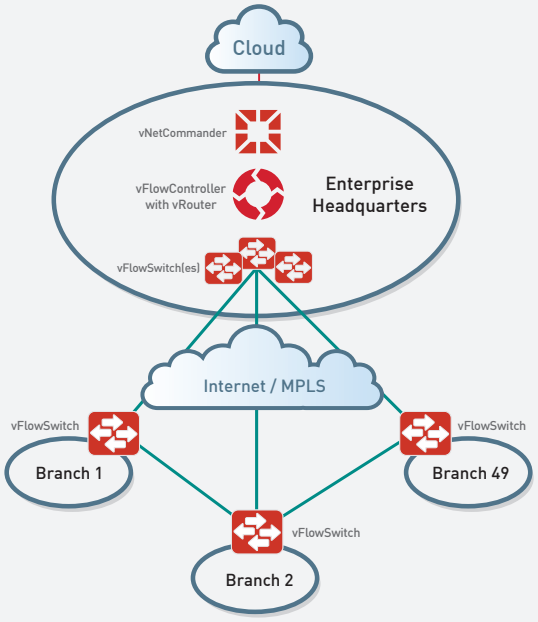

NVN significantly reduces lifecycle CAPEX/OPEX beyond that of traditional site-by-site-managed networking solutions. Immediate benefits include CAPEX savings of 3:1 and OPEX savings of 5:1 over single-purpose, hardware-based legacy networking solutions.

Go "Virtual" Networking Today

Software-defined Networks (SDN) offer a vision of networks evolving to a virtualized world where the networks of yesterday can live harmoniously with the software-based network elements of tomorrow. This virtualized world of SDN offers service providers and enterprises the promise of doing this in a way that allows users to introduce new features and functionality without disrupting their business along the way. Coupled with the pledge of automating fast deployment of new applications that can be integrated into and layered on top of networks, virtual networks hold the potential to deliver optimum business results and an increased bottom line.

So, how do network innovators bridge the gap between rigidly inflexible and costly 'stone-age' networks and the seemingly futuristic network nirvana that SDN promises?

Netsocket Virtual Network (NVN) delivers on the promise of SDN with a network solution that can address the needs of today's dynamic business applications with a virtualized infrastructure that provides end-to-end visibility and centralized remediation for the entire network, transforming it into an asset that is responsive to the needs of the business.

Netsocket Virtual Network (NVN)

Software Defined Networking (SDN) Three-tier Architecture

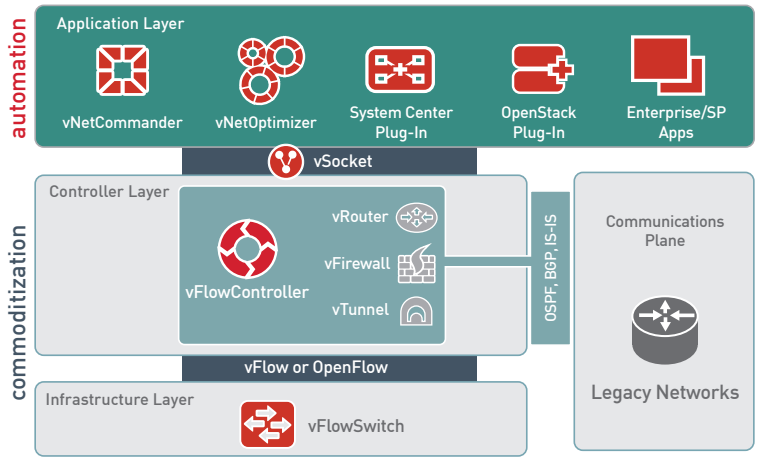

Making The Business Case — Netsocket Virtual Network for Distributed Enterprises

Today's data center centric SDN solutions simply do not address the underserved distributed enterprise use case requirements. They lack necessary functionality such as flexible logical addressing, inter-site quality of service and diverse off-net access per site. Netsocket fills this void with the Netsocket Virtual Network (NVN) delivering a flexible, low-cost, centrally managed virtual network optimized for the enterprise LAN and WAN edge network deployments. Deployed on commodity x86 servers, the Netsocket Virtual Network interconnects enterprise branches in just a few minutes, with no networking expertise required at the site. Its switching and routing components are automatically deployed and provisioned to each branch office using the centralized, intuitive network management application vNetCommander. Utilizing its robust, web-based GUI, the vNetCommander is designed to handle automated deployment, installation, configuration and orchestration of virtualized networks—all from a centralized console.

Netsocket Virtual Network delivers on the promise of SDN through a dramatic reduction in lifecycle costs, impressive network flexibility and deployment response time, and exceptional scalability. NVN provides for legacy network interoperability as well as the ability to easily and cost-effectively incorporate new software or make network changes and updates based on future business needs.

Explore how Netsocket can virtualize your world, visit www.netsocket.com.

Experience your own virtual network today, download the complimentary NVN Early Experience version at www.virtualnetwork.com.

netsocket™

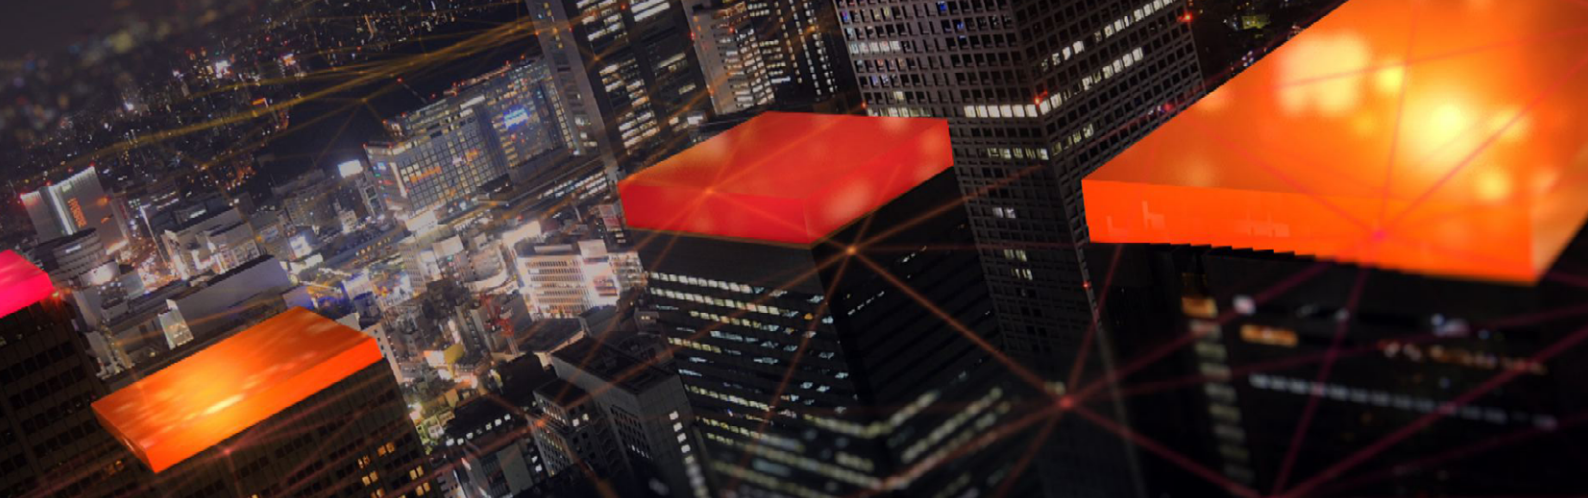

The Consumable Datacenter Network

Taking cloud computing to the next level

The move to cloud computing and storage has changed the way Enterprise users access and consume data. Unfortunately, today's data communications networks aren't keeping pace with this dynamic business environment, and they're struggling to deliver consistent, on-demand connectivity.

That's where we come in. Nuage Networks™ closes the gap between the network and the cloud-based consumption model, creating an infrastructure in which network resources are as readily consumable as compute and storage resources. Our approach enables enterprises to transform the way they build and use their networks, which has a profound effect inside and across multiple datacenters. The transformation is also felt at the critical remote working environment, through a seamless connection to the Enterprise's Wide Area Network.

WOULDN'T IT BE NICE IF...

- Datacenter infrastructures were so simple and standards-based that you could break the vendor lock and work with whichever suppliers offered you the best solutions for your business?
- The network could expand and evolve transparently with the needs of applications, bypassing the datacenter's arbitrary boundaries?
- The datacenter network team could set up controlled, secure templates that application teams could use to deploy applications on the network for and by themselves — without manual transactions or unnecessary project overhead?

Before the move to the cloud, enterprises had to purchase large compute systems to meet the peak processing needs of a limited set of specific events, such as financial milestones (month end or year end), or annual retail events (holiday shopping). Outside of the specific events, the systems were underutilized. This approach was therefore expensive, both in terms of CAPEX and OPEX, requiring significant outlay for power, space and air-conditioning.

Cloud-based datacenters have unshackled the IT environment, making it possible for applications to request additional compute and storage on an as-needed basis. Peak demands can be provisioned "just in time", which lowers operational costs and provides the ability to share compute resources across applications.

The term "cloud" means many things to many people. We focus on two key benefits that cloud computing delivers to Enterprises:

Abstraction of the application from the infrastructure. Cloud computing separates the application from the physical compute and storage infrastructure. This allows workloads to be consistently configured remotely, and templated for mass deployment. End users don't need to worry about the location and specifications of individual hosts. Virtualization and cloud management tools abstract those details to make the infrastructure more readily consumable.

Customer self-fulfillment. Cloud Management Systems (CMS) like Alcatel-Lucent CloudBand™ and the abstraction layer enabled by server virtualization allow IT departments to minimize the tedious and cumbersome processing of application-to-network transactions. For example, IT can provision end customer access policies in the CMS to govern who is authorized to create virtual machine instances, in which location, how many are allowed, and who is the funding department. Users and work groups get instant application deployment, which in turn, makes the business more agile and responsive — critical

attributes in today's enterprise environment. At the same time, operational expenses associated with the handling of work orders is greatly reduced.

As a result of these innovations, Enterprises enjoy a powerful new IT environment in which applications can consume compute resources easily. However as the dynamic nature of cloud computing becomes mainstream, the underlying datacenter network is struggling to match the flexibility of the applications. In fact, most often the network is the weak link, inhibiting the enterprise's ability to profit from the benefits that moving to the cloud should provide.

While virtual compute resources can be instantiated in seconds, it often takes days for network connectivity to be configured and established. Furthermore, the static configurations used by today's networks do not provide the efficiencies and flexibility needed to drive maximum server utilization and application availability.

Consuming the Network

Nuage Networks ensures your network elements are as efficient and flexible as your cloud computing. The result is a choreographed datacenter environment where the compute resources and network work seamlessly.

Imagine the possibilities when network resources are easily consumable. A Nuage Networks datacenter network is as dynamic, automated and virtualized as the server infrastructure, and supports the needs of applications with instantaneous network connectivity.

Nuage Networks eliminates the constraints that have been limiting the datacenter network as it scales out to meet growing demand. With Nuage Networks, you can:

- Define the network service design per application
- Optimize your workload placement across datacenter zones or even across geo-diverse datacenters
- Maximize efficiency of your compute and storage resources

Nuage Networks paves the way for datacenters of the future to be the heartbeat of a powerful cloud infrastructure. Enterprises and user groups could conceive and consume their own secure slices of a robust multi-tenant infrastructure, with appropriate operational visibility and control.

Nuage Networks Virtualized Services Platform

Nuage Networks Virtualized Services Platform (VSP) is the first network virtualization platform that addresses modern datacenter requirements for multi-tenancy, full-featured routing and security at scale. It also integrates seamlessly with wide area business VPN services. It is a software solution that transforms the physical network into a simple to manage, rack-once and wire-once, vendor-independent IP backplane. As a result, network resources within and across datacenters can be treated as an elastic resource pool of capacity that can be consumed and repurposed on demand. Nuage Networks enables unconstrained datacenter networks for the cloud era.

Nuage Networks delivers virtualization and automation of business networks through the three key elements in the Nuage Networks VSP:

Virtualized Services Directory (VSD). Configuration of networks is complex. To eliminate unnecessary complexity while leaving full control and visibility of applications with the IT administrator, the VSD abstracts networking constructs down to their base primitives in four categories: Connectivity Domains, Security, Quality of Service, and Analytics. This allows the requirements for network services to be expressed simply,

FIGURE 1. Nuage Networks Virtualized Services Platform

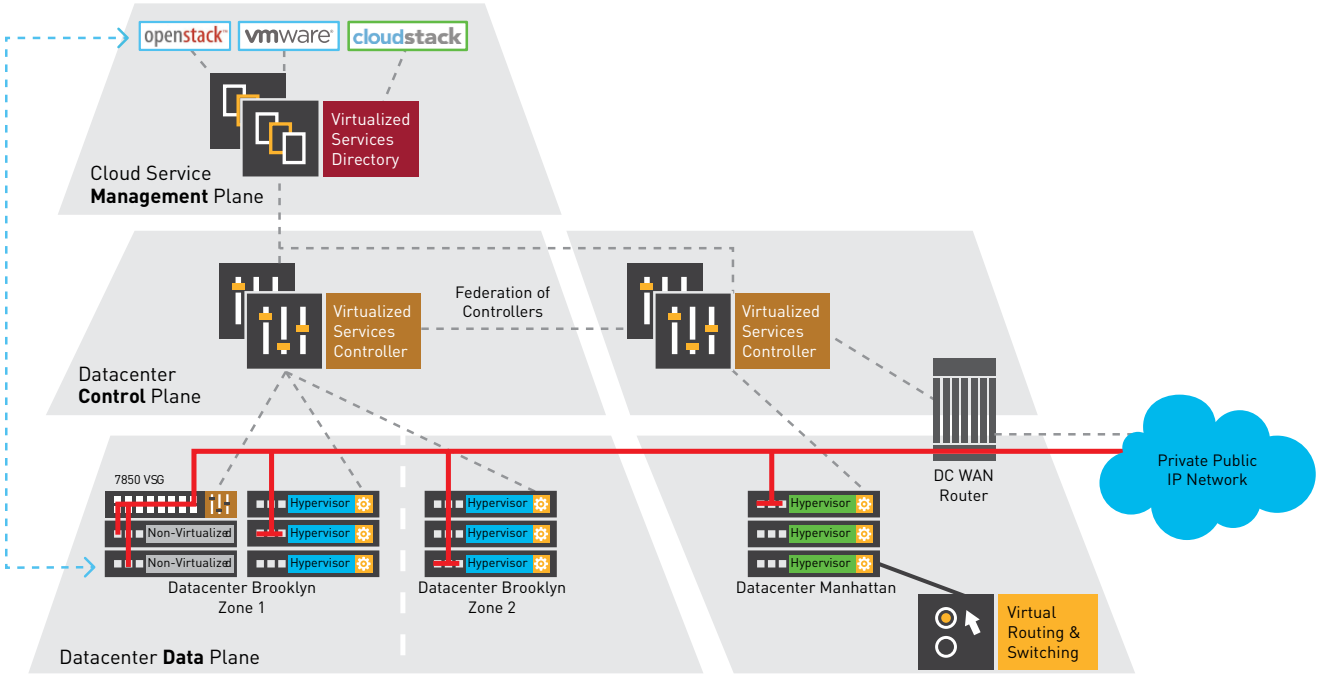

consistently, and in a repeatable manner. The critical need for mobility is also addressed, ensuring network services adjust gracefully and instantly as application endpoints and workloads move from virtual machines within or across datacenters.

The VSD also provides a rich permission-based multi-tenant interface to enable end user provisioning by application owners. Through its role-based hierarchy of permissions, the VSD eliminates operational delays and minimizes transactions between organizations while providing visibility and control of the network "slices" that each group is given in support of their application requirements.

Virtualized Services Controller (VSC). The VSC is an advanced SDN controller that manages the provisioning of virtual network services by programming the edges of the network using OpenFlow™. The VSC ensures that the network follows the application instantaneously. Parting with cumbersome and error-prone device-by-device manual provisioning, Nuage Networks introduces an event-triggered and pull-based configuration model. Once application events such as moves, adds or changes are detected, appropriate policy-based configurations are instantaneously applied. Leveraging Alcatel-Lucent's proven Service Router Operating System, which has been deployed in over 400 service provider networks worldwide for over a decade, the VSC runs a full and robust IP routing stack that allows it to communicate and seamlessly integrate into existing networks.

Virtual Routing and Switching (VRS) is a true hypervisor for the network. The first of its kind in the industry, the VRS fully virtualizes network offerings ranging from distributed virtual Layer 2, Layer 3 forwarding and Layer 4 security. These virtual network services leverage the existing network infrastructure and are offered in a standards-based manner compliant with IETF NVO3. Operators can use whatever servers, hypervisors, and cloud management systems they choose; the Nuage Networks solution abstracts and automates the cloud-networking infrastructure.

In many real-world installations, datacenter environments are a mix of virtualized and non-virtualized assets. To help all datacenters benefit from automation and network virtualization, Nuage Networks supports the full range of options. Software gateways such as the Nuage VRS-G are ideal for environments with relatively low density of bare metal servers and appliances, just as hardware VTEPs from our ecosystem partners provide a viable alternative for certain use cases and environments. For environments with significant investment in bare metal servers and appliances, a new breed of high performance gateway is needed.

The **Nuage Networks 7850 Virtualized Services Gateway (VSG)** is a high-performance gateway that extends Nuage Networks SDN 2.0 functionality seamlessly between virtualized and non-virtualized assets in the datacenter. Working in concert with the Nuage Networks VSP, policies devised for applications automatically extend across virtualized and non-virtualized assets for a fully automated network infrastructure.

FIGURE 2. Nuage Networks datacenter network benefits

	Status Quo	NUAGE NETWORKS DELIVERS What is Needed
Virtualization of network services	LAYER 2 VIRTUALIZATION	FULL NETWORK VIRTUALIZATION, L2 THROUGH L4
Breadth of application models	SIMPLE SCENARIOS	HYBRID CLOUD SERVICES, SEAMLESS VPN CONNECTIVITY
Availability & scale	FRAGILE, NOT MULTI-TENANT	ROBUST, THOUSANDS OF TENANTS
Reach & mobility of network resources	ISLANDS, WITHIN RACKS OR CLUSTERS	SEAMLESS VIRTUALIZED FABRIC, THROUGHOUT & ACROSS DATACENTERS
Network service turn-up time	SLOW, MANUAL, CONFIGURATION DRIVEN	INSTANTANEOUS, AUTOMATED POLICY-DRIVEN
Openness	SPECIFIC TO VENDOR IMPLEMENTATIONS	INDEPENDENCE FROM HARDWARE CHOICES
Breadth of assets automated	VIRTUALIZED ASSETS, LIMITED OPTIONS FOR NON-VIRTUALIZED	ALL DATACENTER ASSETS, VIRTUALIZED & NON-VIRTUALIZED

NU•ÂHJ: FROM FRENCH, MEANING "CLOUD"

The cloud can be more than what it is. In fact, it needs to be. When we founded Nuage Networks, it was with the idea that it's time for the cloud to come of age. From the beginning we recognized the unique challenges that cloud service providers and large enterprises face delivering and managing large, multi-tenant clouds. While the virtualization of compute and storage has evolved quickly, the network simply has not kept up. The result is that today your cloud is being held back. And so is your business.

When we started Nuage Networks, it was with the mission that we could empower our customers to finally deliver on the true promise of the cloud. We envision a world in which IT and IP are no longer in conflict, but rather work in concert to propel your business and elevate the cloud for every one of your customers. We see a world where innovation isn't hampered by infrastructure, and network resources are as effortlessly consumable as compute and storage.

To make this vision a reality, Nuage Networks brings a unique combination of groundbreaking technologies and unmatched networking expertise. This enables us to create solutions that do more than provide incremental improvement. It allows us to introduce radically new thinking and pick up where others have left off, delivering a massively scalable SDN solution that makes the datacenter network able to respond instantly to demand and boundary-less.

Visit us on the web

Our mission is to help you harness the full value of the cloud.

nuagenetworks

www.nuagenetworks.net

While much of the current industry focus on software defined networking (SDN) is in the context of the software-defined data center, Packet Design is enabling SDN in the routed wide area network (WAN) where network programmability and automation demand best practices and tools for management visibility and policy-based control. Always-current network models and traffic load profiles are required for real-time network provisioning by the SDN controller as well as for the successful monitoring and management of SDN applications, such as bandwidth calendaring and workload placement, as well as virtualized network functions and overlay networks.

Packet Design's Route Explorer™ system, available today, maintains a 100% accurate model of the network topology in real time, including IGP areas, BGP autonomous systems, RSVP-TE tunnels, and Layer 2 and Layer 3 VPNs. This is augmented by the recording and analysis of traffic flows to create traffic load profiles. These network models and traffic matrices are available for a variety of network deployment models, including networks with or without RSVP-TE tunnels. Whether the network is programmed or configured (or a combination), network performance can degrade under a variety of conditions, including link or node failures. Route Explorer compares and contrasts network state to a baseline and identifies the root cause of problems quickly. Its monitoring, diagnostics, modeling and reporting capabilities are directly applicable to SDN deployments, providing real-time monitoring, back-in-time forensic analysis, and network event and demand modeling.

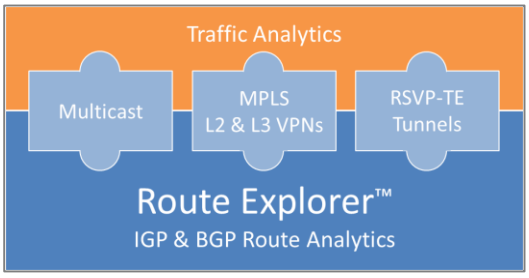

The Route Explorer System

The Packet Design Network Access Broker (NAB), currently in development, uses topology models, traffic profiles and business policies to determine in real time whether or not application requests for network resources can be satisfied. It calculates the impact that requested changes will have on other services by determining the resulting network topology and traffic behavior. The NAB also examines historical traffic profiles to determine if network load is likely to change significantly after the application request is satisfied (for example, the predictable increase in market data and trading traffic that occurs when stock markets open). With Packet Design's unique real-time network models, traffic profiles and analytics, the NAB, which may be integrated in the SDN Controller or exist as an independent software function, provides the intelligence required for mainstream viability of software defined networking in the WAN.

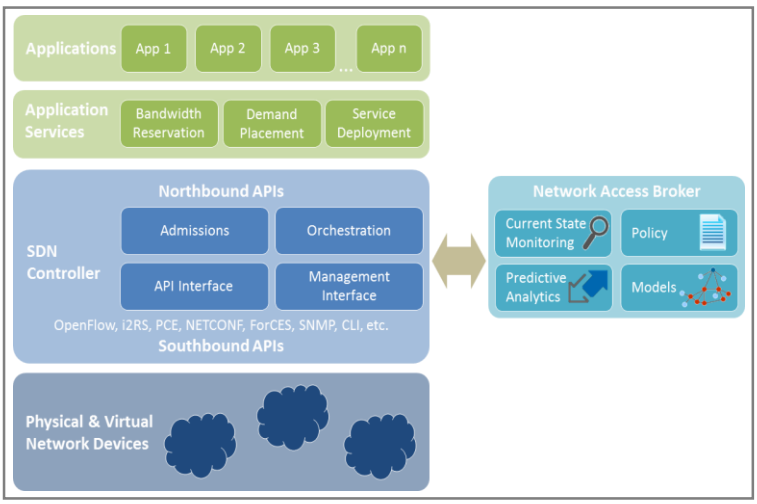

Network Access Broker for SDN

pertino

Cloud Network Engine

Create secure, optimized cloud networks in minutes, add people and devices instantly, and deploy network services on demand.

- Multi-cloud overlay
- Distributed control panel
- L3 switching data plane
- Network service virtualization
- Real-time orchestration
- App store

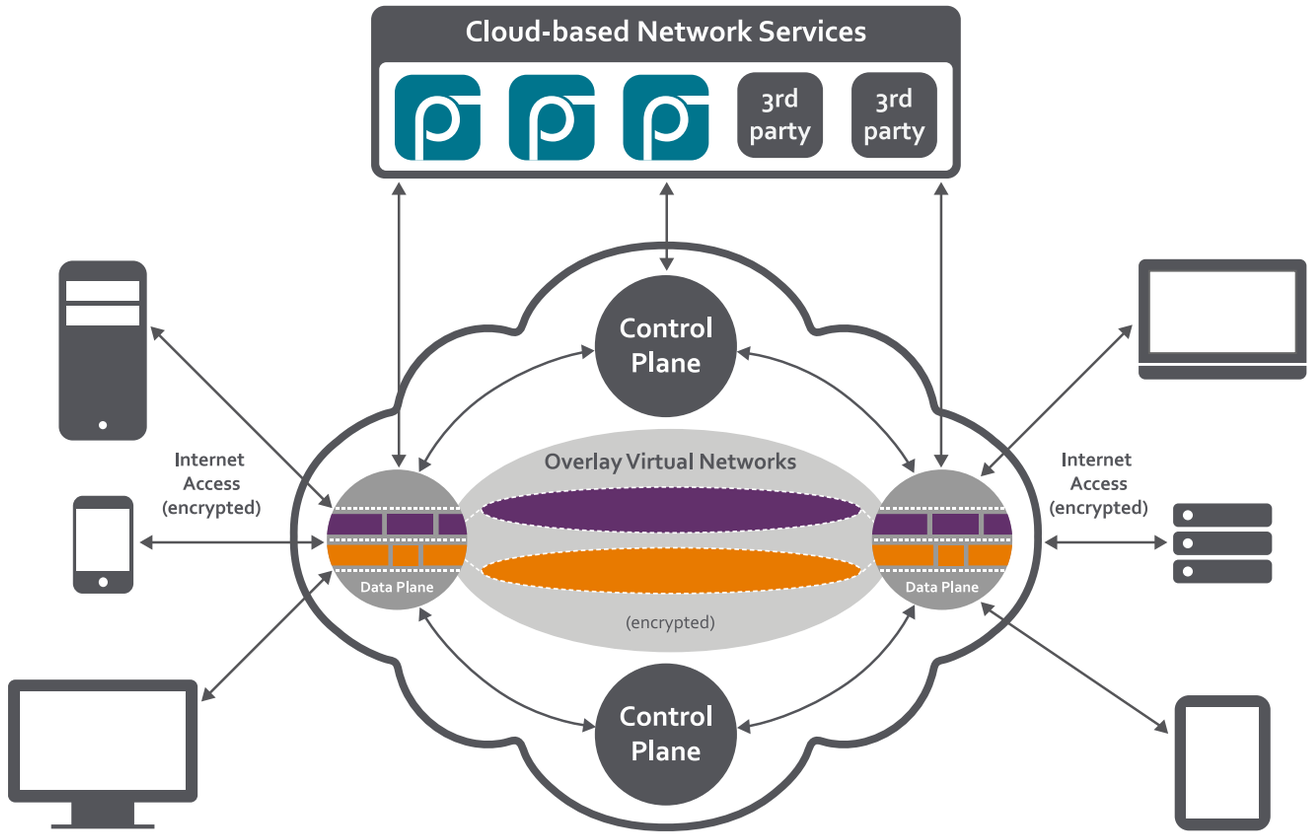

rackspace HOSTING amazon web services Windows Azure linode

pertino.com 408.354.3900 info@pertino.com

PICA8
OPEN NETWORKING

WWW.PICA8.COM

Open Systems for Software Defined Networking (SDN)

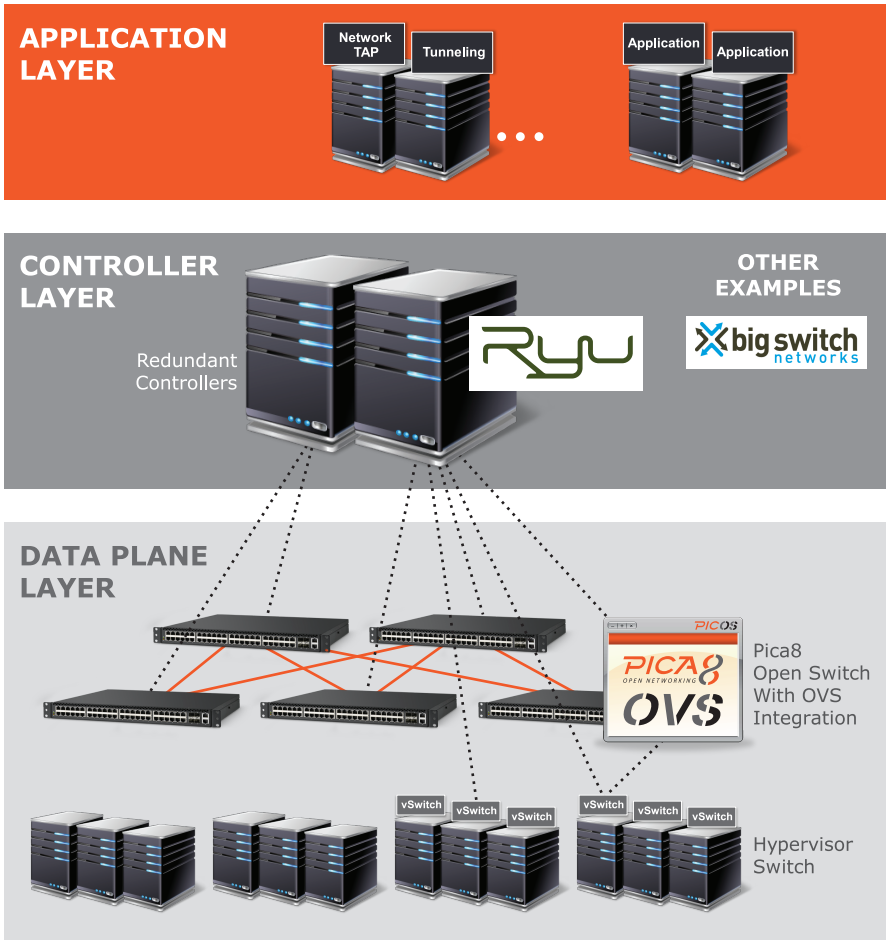

The First Hardware Agnostic, Open Network Operating System

Pica8™ is the first in the world to offer hardware-agnostic open switches. A pioneer in software-defined networking (SDN), we pair high-performance, white box switch hardware with PicOS: our hardware-agnostic, open network operating system that supports standards-based Layer 2 / Layer 3 protocols and Industry-leading OpenFlow* 1.3. In one complete package, Pica8 provides the physical switch, comprehensive switching and routing features, and the fulfilled promise of open networking.

What makes PicOS open?

- **PicOS is hardware agnostic**: because of PicOS's hardware abstraction layer, the operating system is not tightly coupled to any switching ASIC, CPU or memory hardware. We continue to expand our ODM partners, offering a portfolio of pre-qualified white box, bare metal switches to select from
- **Debian Linux is exposed**, so you can use your existing tools (such as Puppet, Chef or CFEngine) for hands-free provisioning and myriad APIs through the Debian-Linux environment, helping you personalize Pica8 switches to support your open network
- **PicOS supports OpenFlow 1.3**, through Open vSwitch (OVS) v1.9 integration: OVS runs as a process within PicOS, providing the OpenFlow interface for external programmability

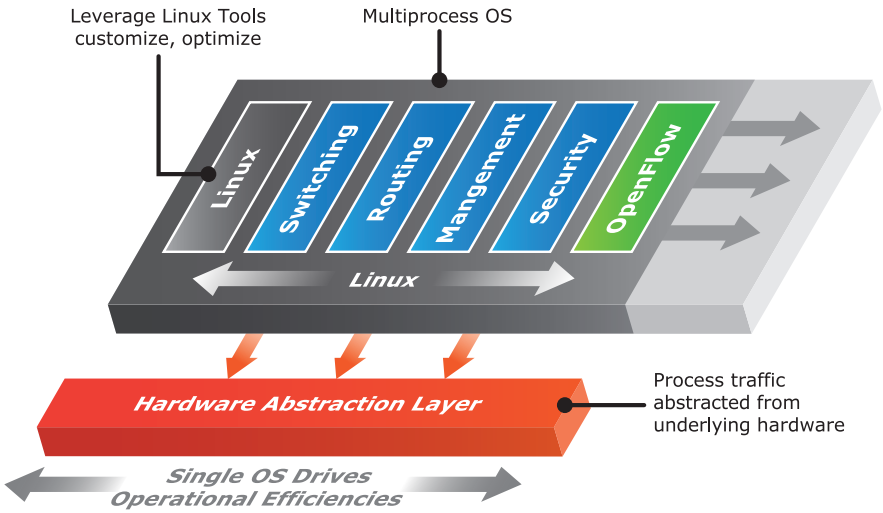

* Only OpenFlow features available in hardware are supported, to ensure optimum performance

Automation for Agile Infrastructure

Corporate Overview

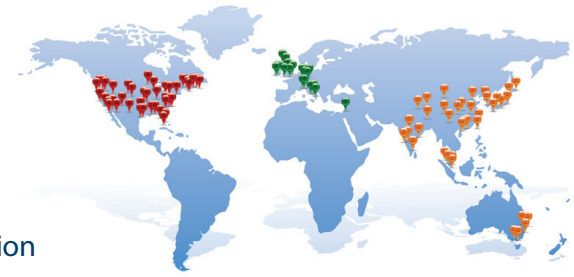

Founded: 2004

North America HQ: Santa Clara, CA

Market-leading supplier of automation solutions for:

- Network test and test lab efficiency, productivity and savings
- IT infrastructure self-service for DevOPS agility and cloud evolution

Mature, proven technology:

- Hundreds of customer deployments
- Millions of infrastructure elements managed
- $Billions in infrastructure managed

Automation Platform

Comprehensive Automation Framework

- Resource management
- Heterogeneous environment design + workflow authoring
- Reporting and business intelligence
- Self service portal

Object library-based architecture

- Supports & enforces best practices
- Optimizes programming staff skills
- Achieves high ROI through ease of maintenance and scalability

Any-Stack Integration

- Key API integration libraries + open driver creation
- Freedom from vendor roadmaps, allows integration with legacy, home-grown components
- Overcomes interface silos

User-friendly GUI-based automation design

- Break open expertise bottlenecks
- Systematize knowledge, increase reusability
- Maximize total team productivity

SDN Self-Service Automation

- SDNs offer northbound API's for applications to drive network behavior
- Yet SDN adopters will need to manage heterogeneous network environments with both legacy and SDN elements
- CloudShell provides the means to automate the delivery of SDN/legacy network environments for DevOPS network application development, testing and deployment

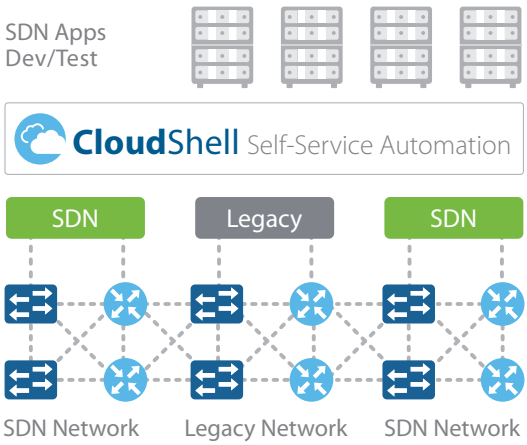

TestShell

TestShell is an object-oriented test and lab automation platform. It delivers powerful lab infrastructure management, and test automation solutions for network, data center, tech support, and demo/PoC lab environments. TestShell is deployed by leading service providers, technology manufacturers, enterprise and government IT departments around the world.

TestShell's object-oriented architecture revolutionizes network, data center and cloud infrastructure testing by:

- Dramatically increasing the efficiency and ROI of test infrastructure through improved resource sharing
- Simplifying the creation, maintenance and re-use of automated device control interfaces, provisioning actions and testing tasks through a shared object library
- Empowering non-programmers to create, save, share, integrate and reuse complex test topologies and automation workflows
- Enabling seamless hand-offs of topologies and automation workflows between developers, architects, QA teams, pre-production, technical support, field operations and customer engineers

CloudShell

CloudShell is a self-service automation platform for heterogeneous, multi-generational IT infrastructures and networks. It helps infrastructure and networking teams to deliver agile, end-to-end infrastructure to application delivery stakeholders including developers, testers, compliance and security engineers, and deployers.

Self-service automation of heterogeneous, multi-generational IT infrastructure

- Legacy systems and stack
- Traditional datacenter and network environments
- Industry-specific IT components
- Software-Defined Networking
- Private and public clouds

Helps IT infrastructure and network teams achieve DevOPS agility

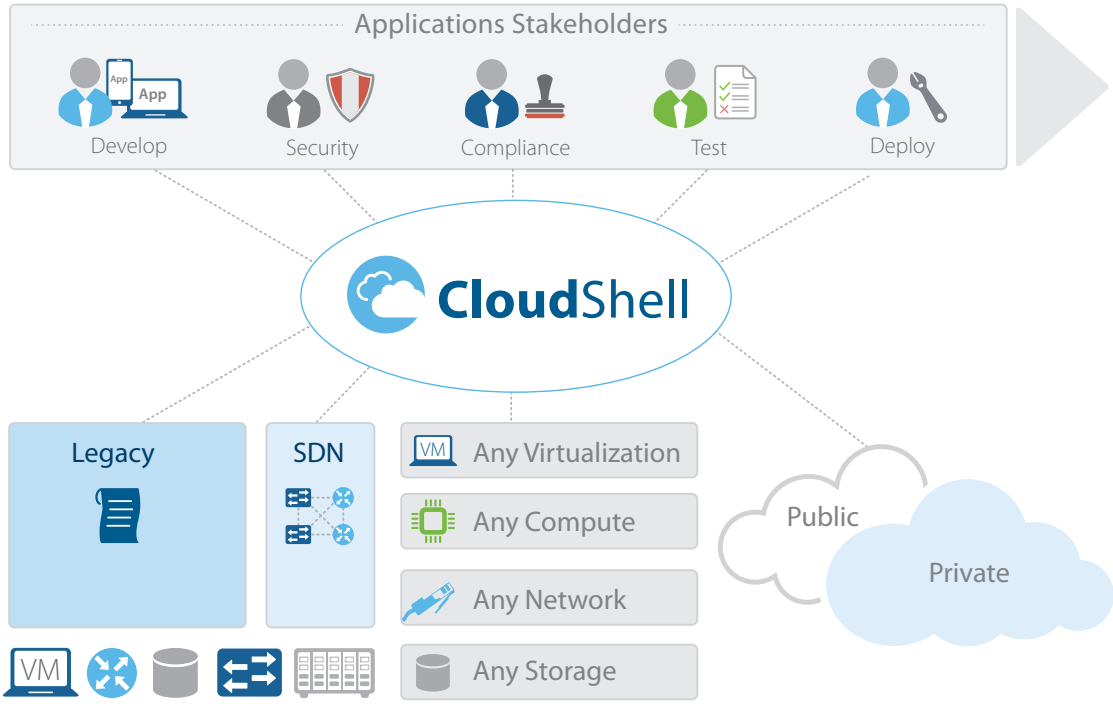

For more information about QualiSystems, visit our website at
www.qualisystems.com

Software Defined Networking Solutions
Enable Network Wide Services via SDN Applications

Radware SDN applications improve application security, performance and availability by programming the SDN to collect data and optimally forward traffic to deliver network services. The native component of the new network stack introduced by SDN includes the data plane networking devices and the control plane SDN controllers. The Radware SDN applications constructing the SDN application control plane, interact with the SDN controller using dedicated SDN drivers and work together with the Radware systems' using the Radware API to collect data throughout the application infrastructure using specific data collection drivers.

With Radware SDN applications, ADC and security services transform from device-based solutions requiring a static traffic forwarding configuration, to network wide services that intelligently divert traffic to service engines. Network services can scale to support larger networks at lower capital and operational cost. By building SDN applications that continuously interact with the SDN control plane and program the network (and by leveraging the Radware Virtual Application Delivery Infrastructure (VADI) architecture – which enables pooling of disperse resources to operate uniformly) Radware enables an anywhere and everywhere network service paradigm.

Key benefits from the Radware SDN network service infrastructure include:

- **More intelligent application delivery and security decisions** throughout the network break existing network barriers when developing business applications. Every application everywhere is entitled for advanced services.

- **Simpler implementation** of network services allows improved operational efficiency of network management alongside application changes. Not every project needs to become a networking project.

- **Lower overall network service solution costs** – as network service delivery is partially offloaded to the SDN, there is no need to invest in excess network service appliances and capacity. Deploy network services as needed, and use by many tenants and applications throughout the datacenter.

- **Greater scalability** – scale your network services throughout the network. No more limited areas are protected or load balanced. Offer uniform services throughout the SDN.

- **Easier operation** – changing and managing security and ADC functionality becomes simpler as the deployment operates as if it is centralized. Not only does SDN streamline network operations, but Radware SDN applications streamline network service operations.

DDoS Protection as a Native SDN Application

DefenseFlow is an SDN application that enables network operators to program the network to provide DDoS protection as a native network service. DefenseFlow features an adaptive behavioral-based DoS attack detection engine and a traffic diversion mechanism that utilizes the programmable characteristics of the software defined network elements for attack cleansing. Designed as part of the Radware SDN application framework, DefenseFlow operates in any SDN enabled network infrastructure.

Legacy DDoS protection solutions that make use of scrubbing centers are costly: need hardware detectors in every network location; BGP for traffic diversion; and GRE tunnels to forward the traffic to its designated network object. With SDN, a DDoS protection solution turns into a software application that adds intelligence to the network – no need for additional hardware, BGP or GRE operations.

DefenseFlow equips network operators with the following key advantages:

- **Unprecedented coverage against** all type of network DDoS attacks
- **Best design for attack mitigation**
    - Attack detection is always performed out of path (OOP)
    - During attack only suspicious traffic is diverted through the mitigation device
- **Most scalable mitigation solution** – DefensePro mitigation devices can be placed in any location, DefenseFlow diverts the traffic to the nearest mitigation device.

SDN for a Scalable Application Delivery Network

Radware's ElasticScale is an SDN application that wraps existing network service virtual appliances and provides provisioning and traffic distribution logic to consistently deliver network services in an elastic demand environment. ElasticScale can be utilized for service provider internal services, managed services to end customers and can providers adopt network function virtualization paradigms.

ElasticScale offers network operators the following key features and benefits:

- Ultra scalable traffic steering solution (100's of Gbps)
- Ultra scalable load balancing solution
- Based on industry leading, carrier grade Alteon load balancing product line
- Support for leading hypervisors (oXen/KVM/Hyper-V/ESXi)
- Compatible with leading SDN controllers; OpenDaylight, Cisco XNC, NEC pFlow & HP Flare
- Seamless integration with OpenStack and vCloud Director
- Runs over any physical SDN network equipment

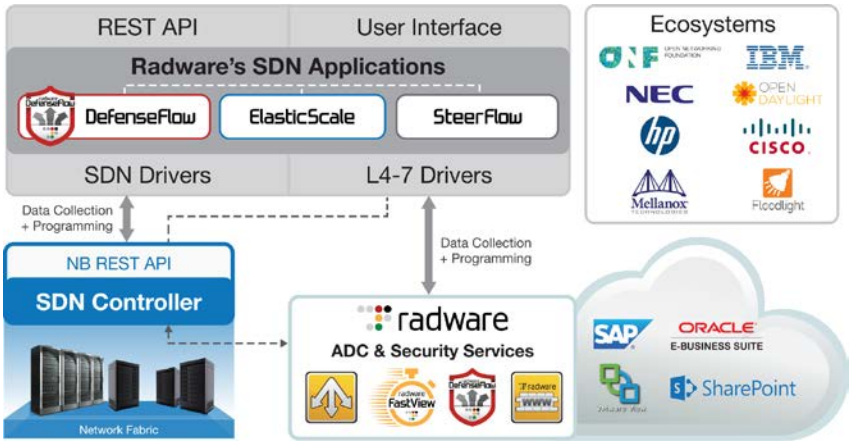

Partnering for Success: Our SDN Ecosystem

The SDN eco-system is a critical focus for Radware. Through partnerships with the industry's leading SDN forums and vendors, Radware can ensure customers that our application delivery and security solutions integrate successfully into target architectures.

Radware is an active contributor in the following industry and vendor SDN initiatives: Big Switch Networks, Cisco Open Network Environment (ONE), Floodlight, HP Virtual Application Networks, IBM Distributed Overlay Virtual Ethernet (DOVE), NEC, Mellanox, Open Daylight Project, and the Open Networking Forum (ONF). Radware is also a member of VMware's NSX partner ecosystem for network functions virtualization (NFV).

Learn More

To learn more about how Radware's SDN solutions can enable you to get the most of your business and IT investments, email us at info@radware.com or go to www.radware.com.