The 2013 Guide to Network Virtualization and SDN

Part 1: Introduction and Network Virtualization

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

Introduction 2

Traditional NV & The NV Use Case 4

Network Overlays via Tunneling: Benefits & Limitations 5

Cloud Orchestration 9

Controller Based NV Solution Architecture 11

Criteria to Evaluate Overlay NV Solutions 12

Tunnel Encapsulation 14

Tunnel Control 15

Comparison of Network Overlay Virtualization Solutions 17

Software Defined NV via Flow Table Segmentation 18

Enterprise Plans for NV Adoption 19

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
EMC^2
Extreme
NEC
Netsocket
Nuage Networks
Packet Design
Pertino
Pica8
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 will be published both in its entirety and in a serial fashion. This is the first of the serial publications. This publication will focus on NV. The three subsequent publications will focus on:

1. SDN
2. The Vendor Ecosystem
3. Planning for NV and SDN

In August and September of 2013 a survey was given to the subscribers of Webtorials. Throughout this document, the IT professionals who responded to the surveys will be referred to as *The Survey Respondents*.

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 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 build 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

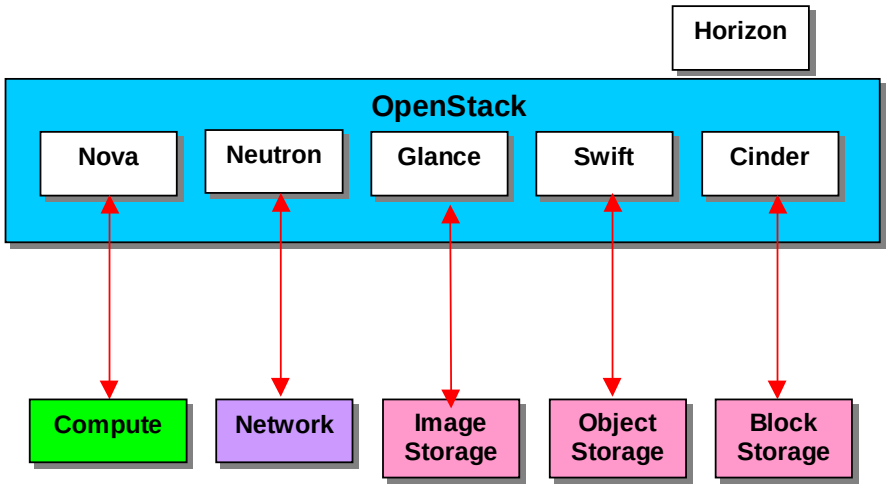

Neutron (formerly called Quantum) allows users to create their own networks, provide connectivity for servers and devices, and control traffic. With appropriate Neutron plug-ins, administrators can take advantage of various NV and SDN solutions to allow for multi-tenancy and scalability. OpenStack networking 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

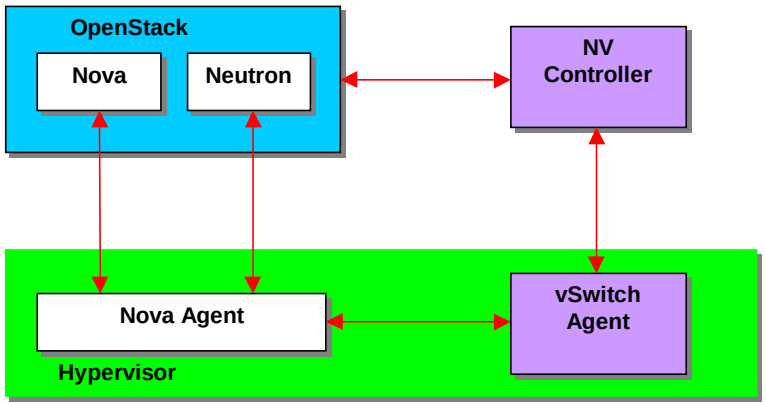

The Nova module in OpenStack instructs the Nova Agent in the hypervisor to create the VM. The Nova agent communicates with the Neutron module in OpenStack to learn the network attributes of the VM. The Nova agent then informs the vSwitch agent to configure the virtual network for the VM and then the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

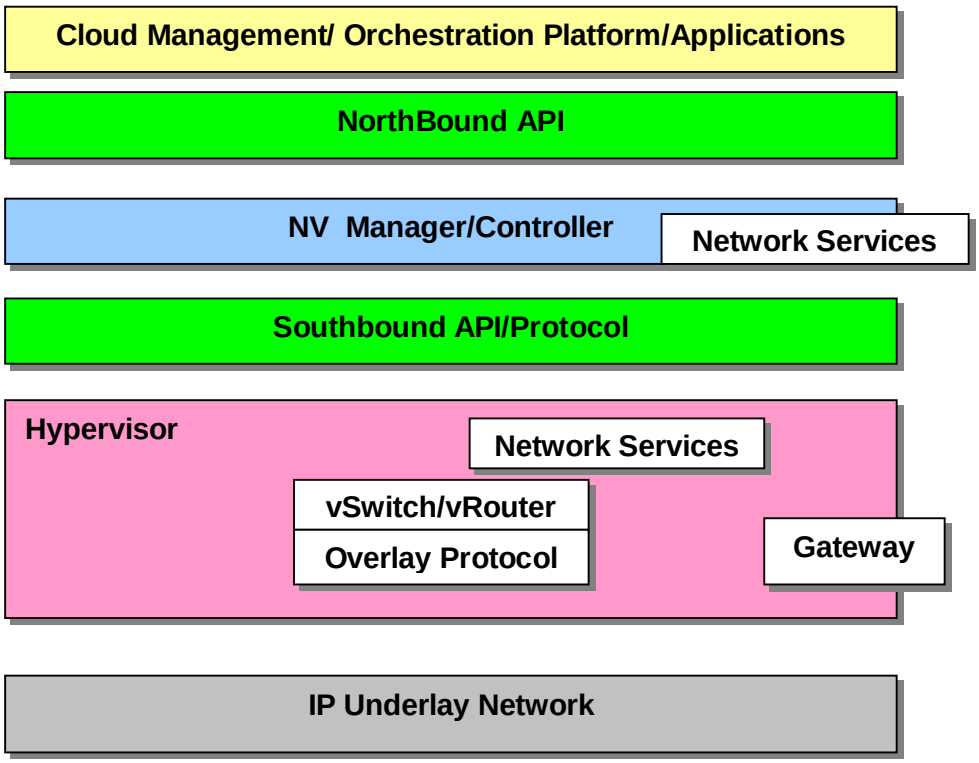

Criteria to Evaluate Overlay NV Solutions

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

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

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

provides added flexibility for the virtual network tunnel control plane. With these features, STT is 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.

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

One controller-based solution for VXLAN control, championed by IBM's Distributed Overlay 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.

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

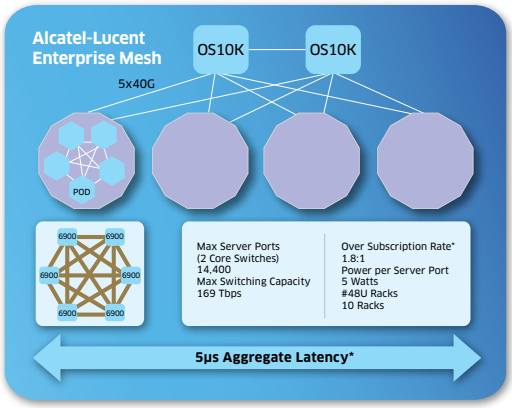

Alcatel-Lucent provides a unique Application Fluent approach to maximize the benefit from virtualization technologies for servers, the desktop, as well as the network. Alcatel-Lucent's application fluent data center fabric can scale from several hundred to over 14,000 server facing ports while keeping 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.

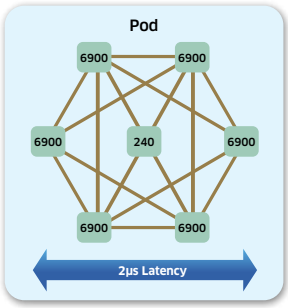

Application fluency in the corporate data center includes its transformation into a multi-site private cloud by extending layer 2 connectivity between data center sites and allowing for seamless delivery of public cloud-based services on the corporate network.

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

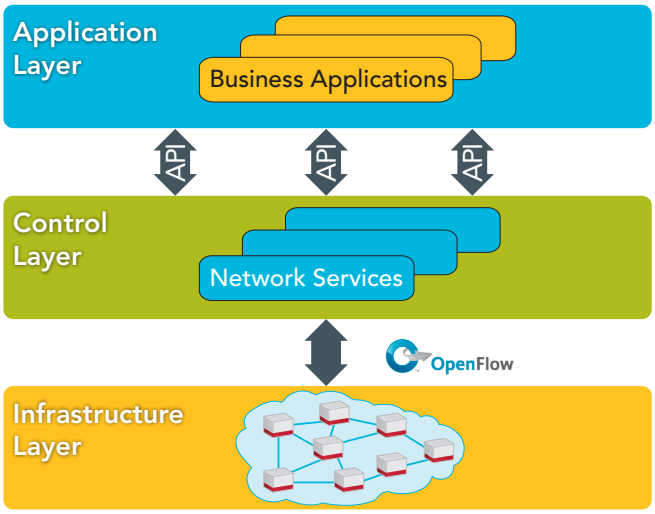

Figure 1. ONF SDN Architecture

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.**

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.

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

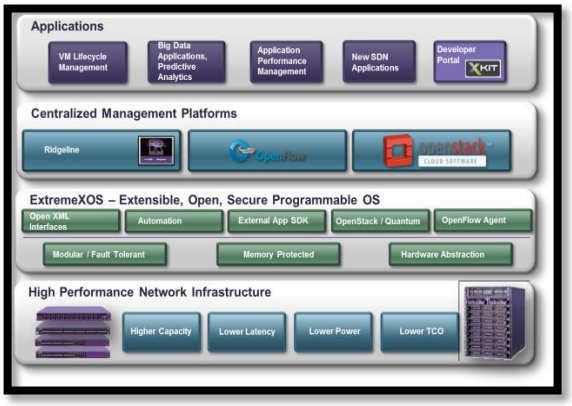

Critically important to the Open Fabric is ExtremeXOS®, the network operating system that delivers the consistent set of features across all platforms while ensuring the security and performance of the Open Fabric. ExtremeXOS is modular, extensible, and has integrated security, while providing a single 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

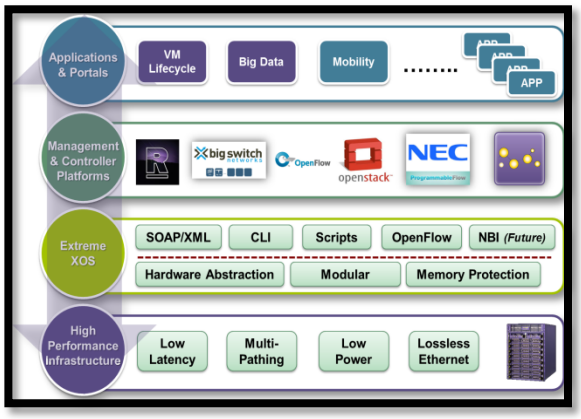

The Extreme Networks Open Fabric SDN strategy therefore extends to include technology partners and systems that leverage the network abstraction capabilities provided by ExtremeXOS.

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

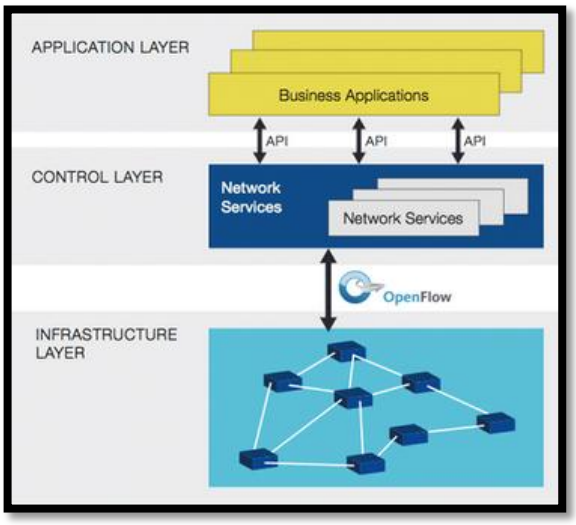

The Open Networking Foundation (ONF) defines OpenFlow: "The physical separation of the network control plane from the forwarding plane, and where a control plane controls several devices."

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

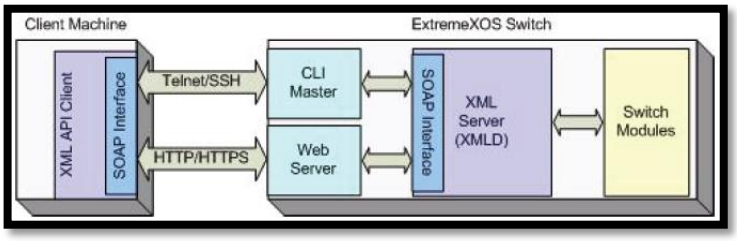

"Open API's enable applications and management systems to directly access the network abstraction layer to manage the control, data and management planes of the infrastructure."

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

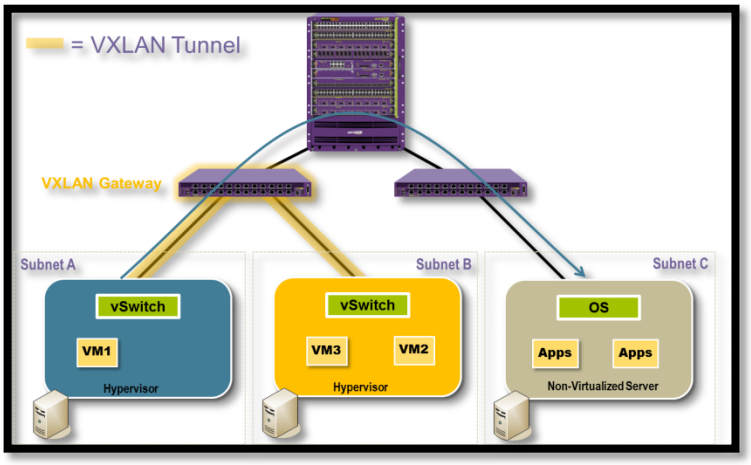

"Network Virtualization Overlays include logical overlays from virtualized server and switching systems that may also include virtualized layer 4-7 services."

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

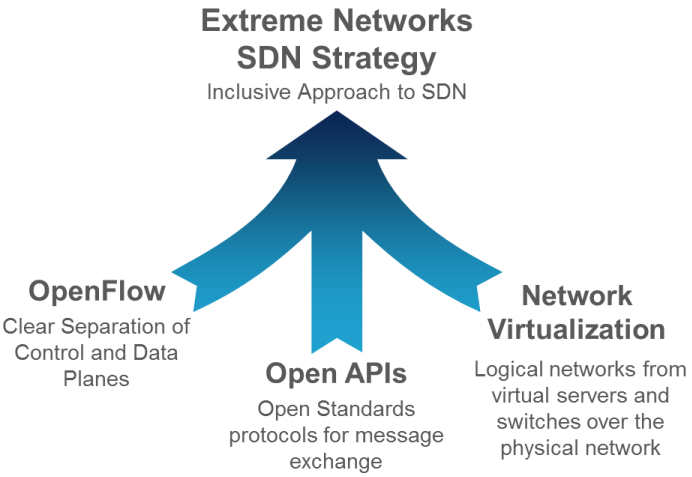

ProgrammableFlow

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

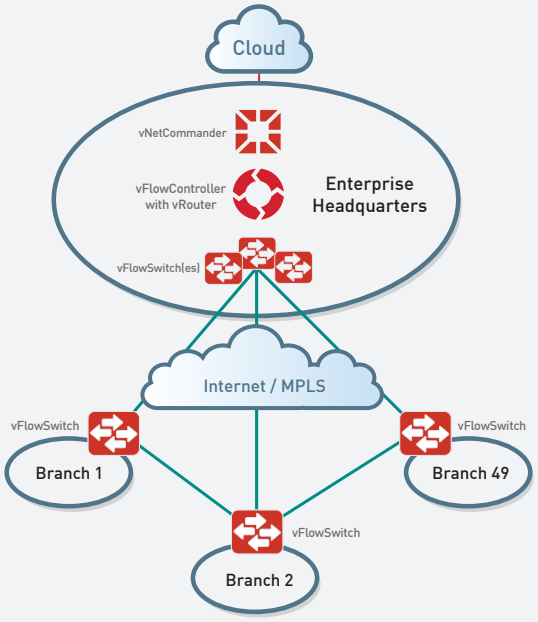

NVN significantly reduces lifecycle CAPEX/OPEX beyond that of traditional site-by-site-managed networking solutions. Immediate benefits include CAPEX savings of 3:1 and OPEX savings of 5:1 over single-purpose, hardware-based legacy networking solutions.

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

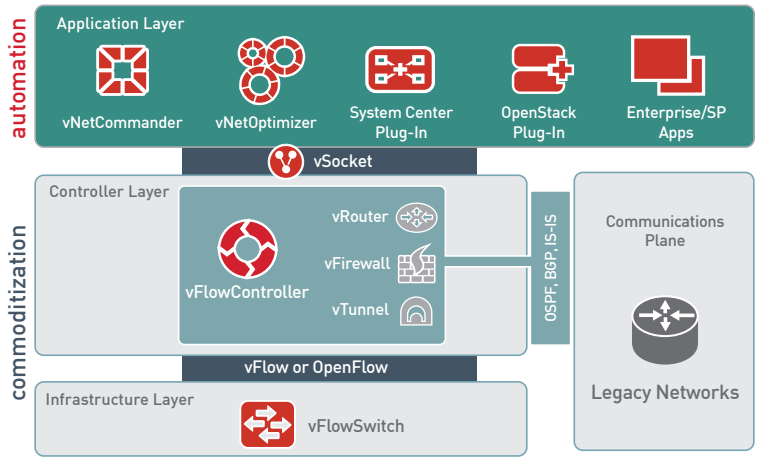

Making The Business Case — Netsocket Virtual Network for Distributed Enterprises

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™

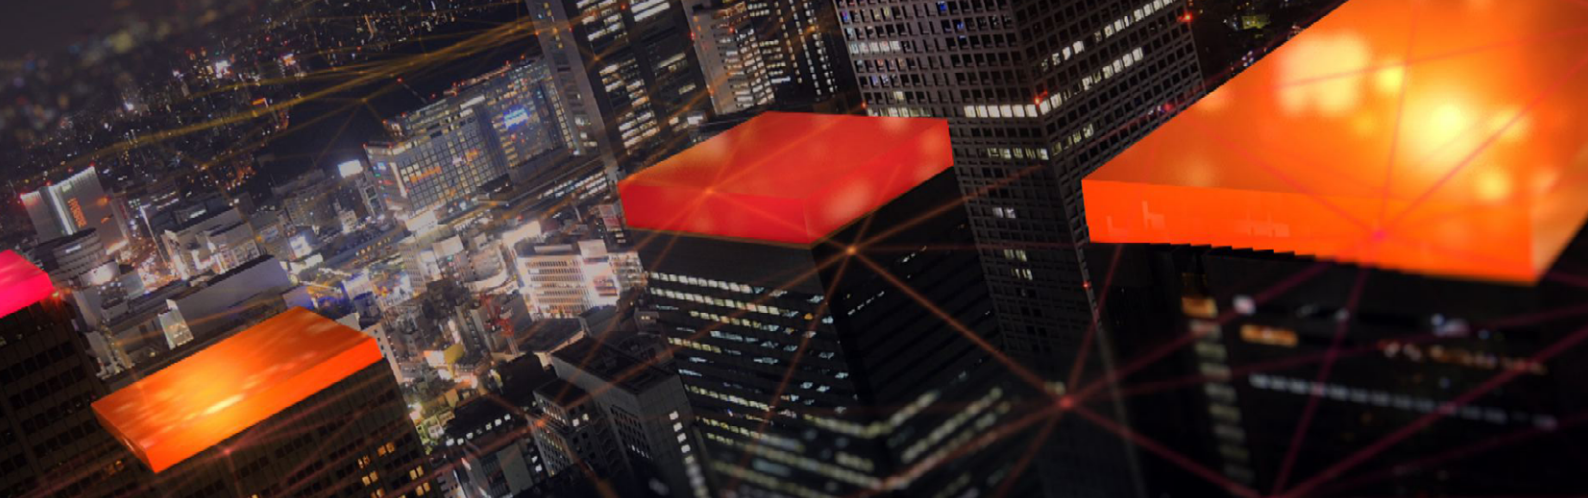

The Consumable Datacenter Network

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

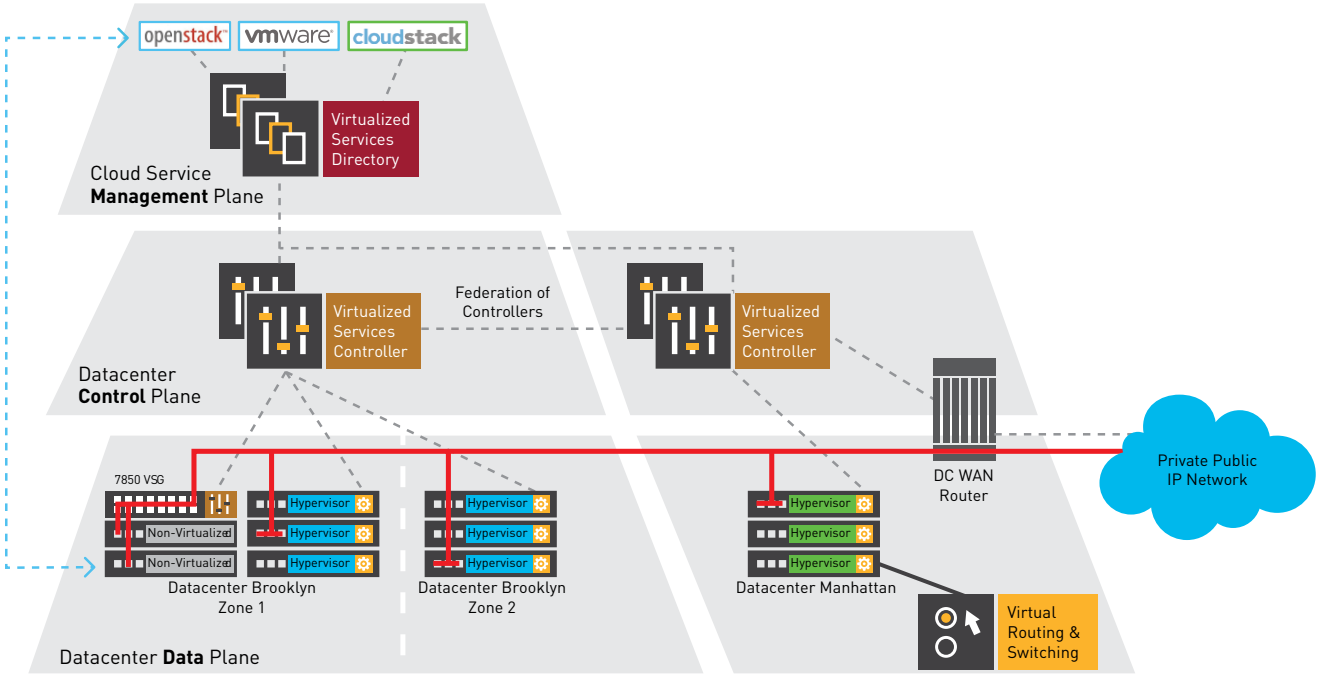

consistently, and in a repeatable manner. The critical need for mobility is also addressed, ensuring network services adjust gracefully and instantly as application endpoints and workloads move from virtual machines within or across datacenters.

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.

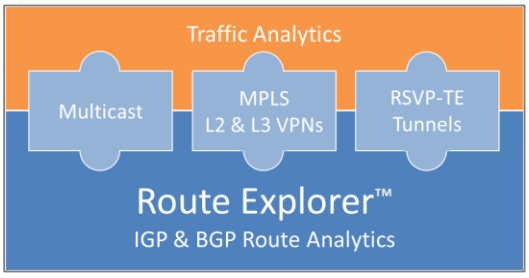

The Route Explorer System

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.

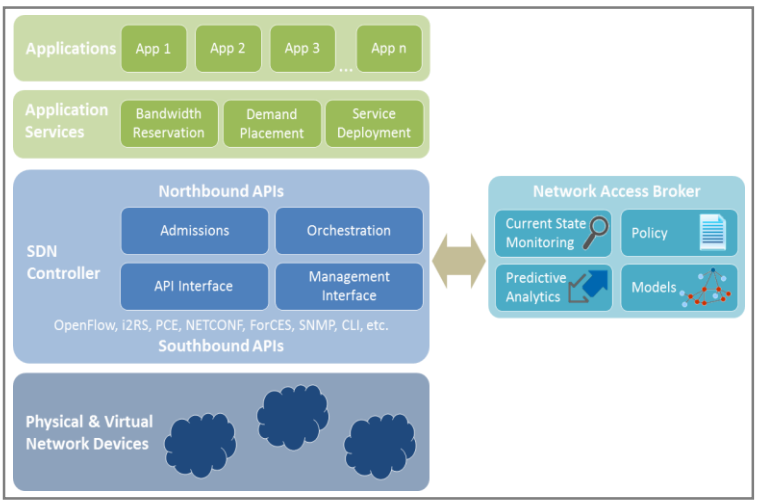

Network Access Broker for SDN

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

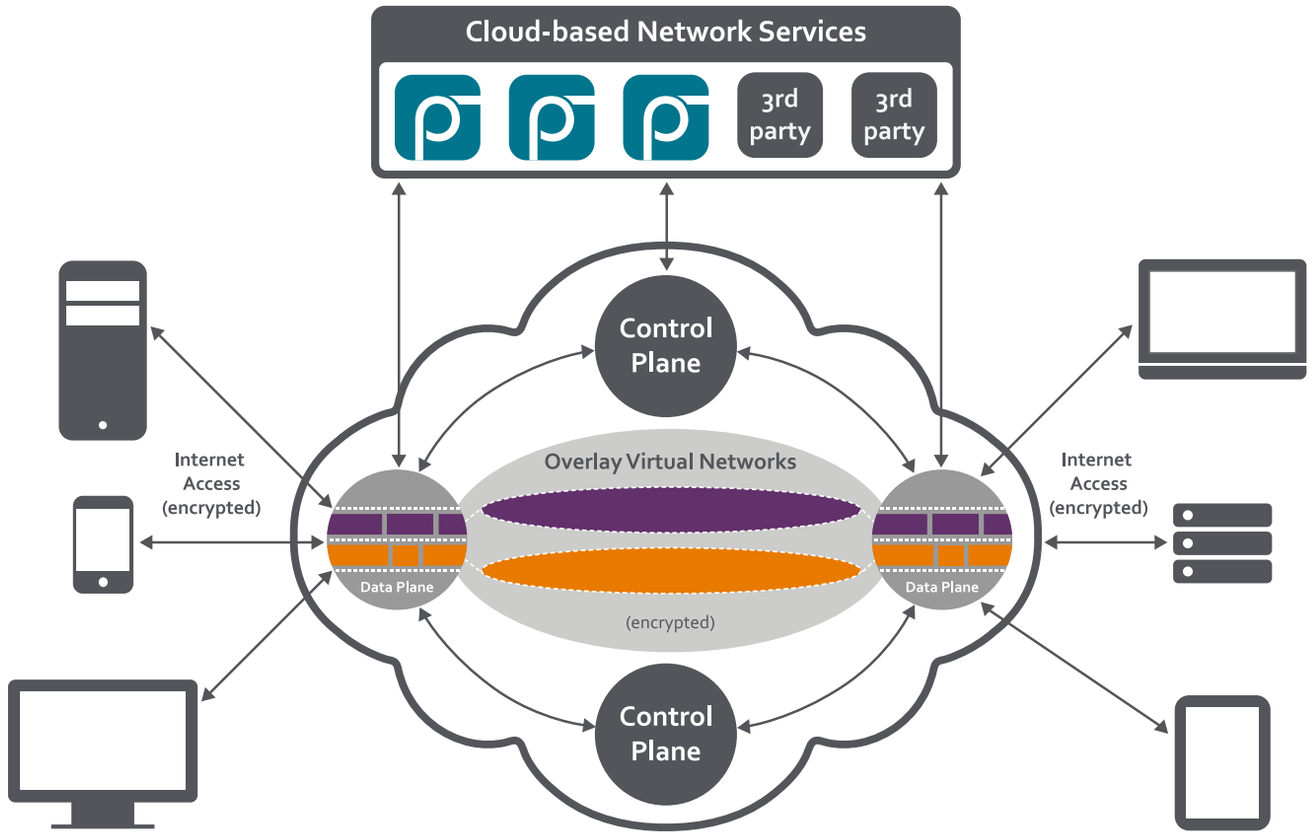

rackspace HOSTING amazon web services Windows Azure linode

pertino.com 408.354.3900 info@pertino.com

PICA8
OPEN NETWORKING

WWW.PICA8.COM

Open Systems for Software Defined Networking (SDN)

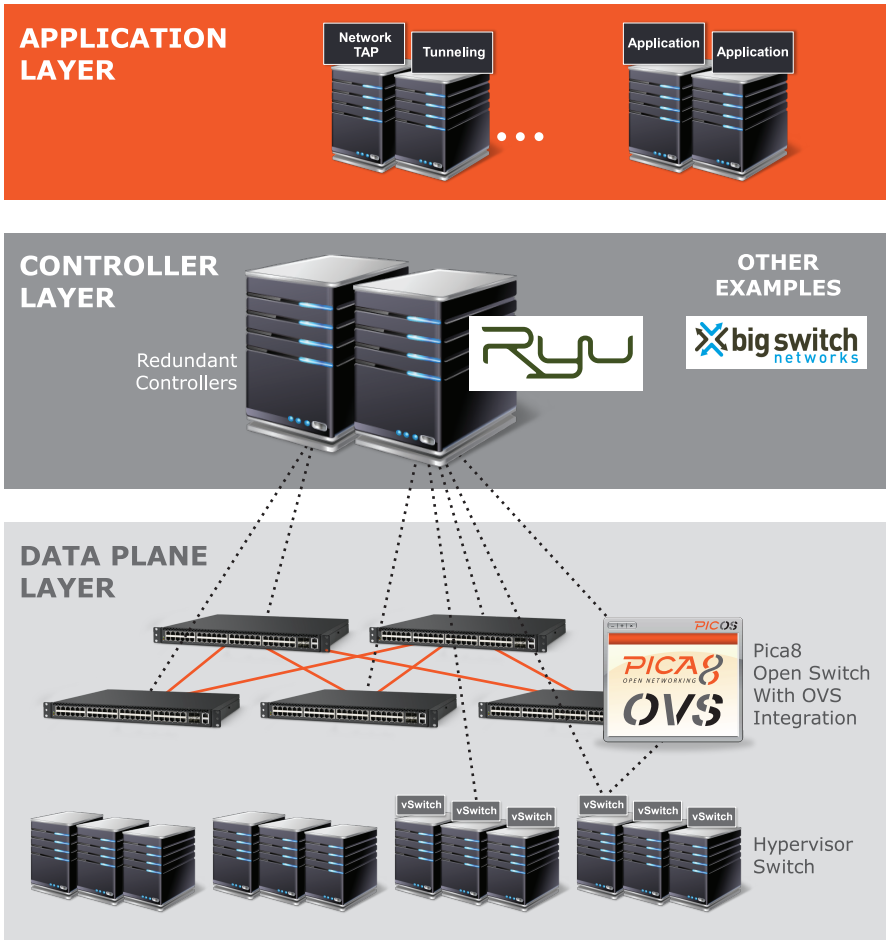

The First Hardware Agnostic, Open Network Operating System

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

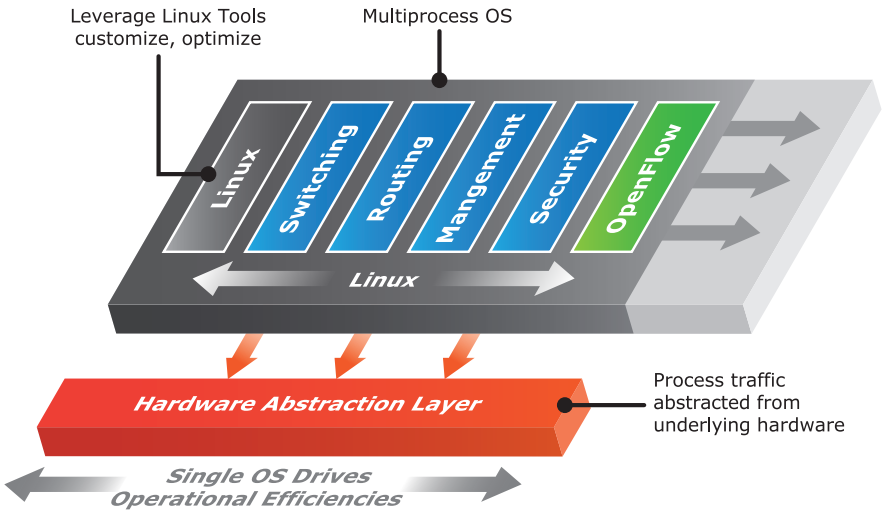

* Only OpenFlow features available in hardware are supported, to ensure optimum performance

Automation for Agile Infrastructure

Corporate Overview

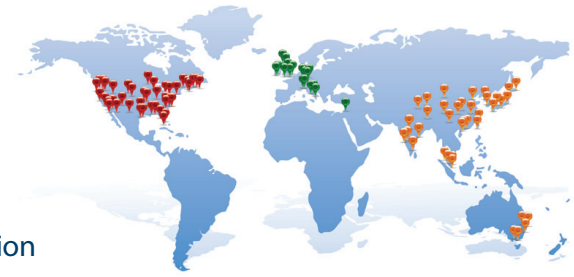

Founded: 2004

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

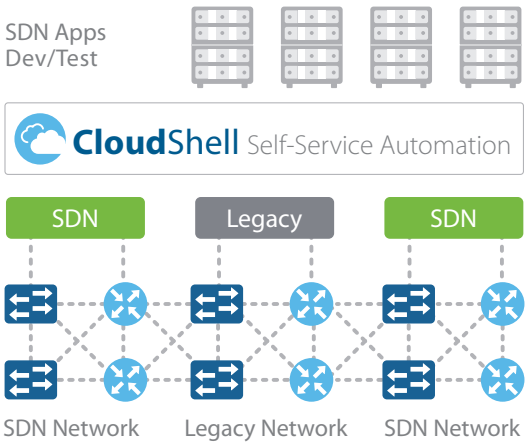

TestShell

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

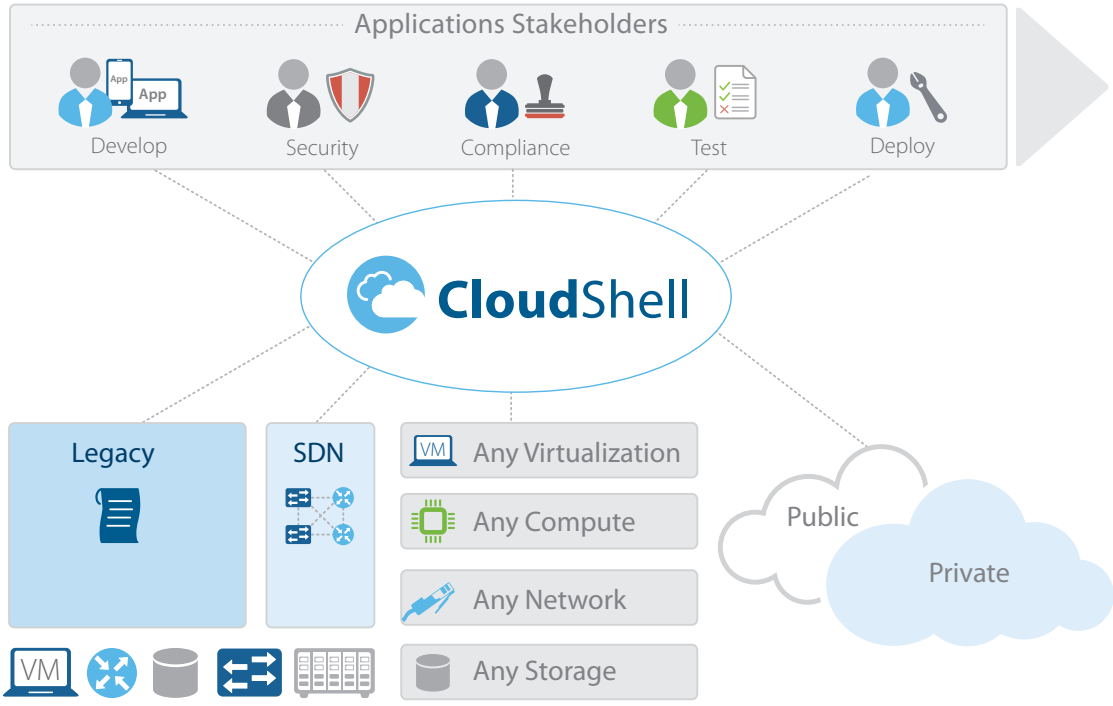

For more information about QualiSystems, visit our website at
www.qualisystems.com

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

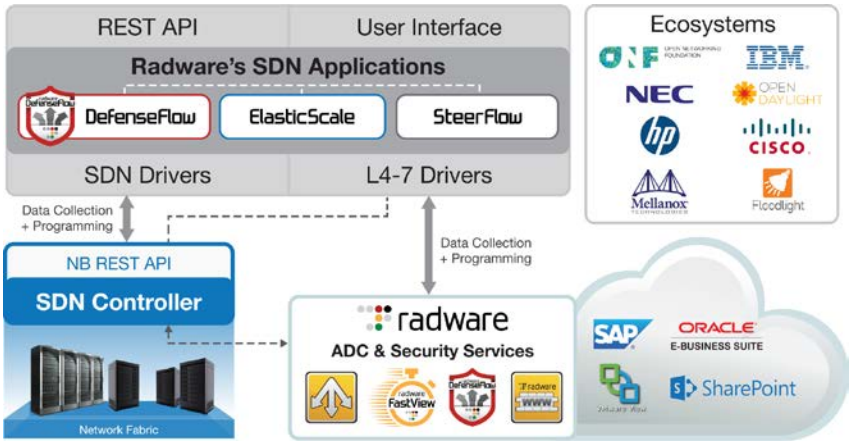

Partnering for Success: Our SDN Ecosystem

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.