The 2013 Guide to Network Virtualization and SDN

Part 2: The What, Why and How of SDN

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

Platinum Sponsors:


Gold Sponsors:

Produced by:

Table of Contents

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

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 second of the serial publications. The first publication[1] focused on NV and this publication will focus on SDN. The two subsequent publications will focus on:

1. The Vendor Ecosystem
2. 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*.

1 webtorials.com/Metzler

Background

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

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

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

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

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

Potential SDN Use Cases

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

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

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

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

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

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

The left hand column of **Table 1** contains some of the primary challenges & opportunities facing the typical IT organization. The Survey Respondents were shown those challenges & opportunities and were asked to indicate which of them they thought that SDN could help them to respond to and they were allowed to check all that applied. Each row of the right hand column of **Table 1** contains the percentage of The Survey Respondents that indicated that they thought that SDN could help them to respond to the challenge or opportunity in the corresponding left hand column.

Table 1: Opportunities & Challenges that SDN Can Address	
Challenge or Opportunity	**Percentage**
Better utilize network resources	51%
Ease the administrative burden of configuration and provisioning QoS and Security	47%
Perform traffic engineering with an end-to-end view of the network	44%
More easily scale network functionality	39%
Support the dynamic movement, replication and allocation of virtual resources	38%
Establish virtual Ethernet networks without the limitations and configuration burden of VLANs	35%
Reduce Complexity	34%
Enable applications to dynamically request services from the network	32%
Reduce OPEX	30%
Have network functionality evolve more rapidly based on a software development lifecycle	27%
More easily implement QoS	27%
Implement more effective security functionality	26%
Reduce CAPEX	25%
We don't see any challenges or opportunities that SDN can help us with	3%
Don't know	3%
Other	3%

One observation that can be drawn from the data in **Table 1** is that there is a wide range of challenges and opportunities that The Survey Respondents believe that SDN can help with and conversely very few IT organizations believe that SDN won't be beneficial. Having a wide range of potential challenges and opportunities to respond to bodes well for the long-term adoption of SDN. However, having so many challenges and opportunities to respond to can create confusion in the short term and can possibly delay SDN adoption.

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

The situation is quite different if the opportunity that the hypothetical IT organization is trying to respond to is the need to make it easier to implement QoS or the need to enable applications to dynamically request services from the network. The hypothetical IT organization can potentially respond to both of these challenges by implementing an SDN solution whereas that organization couldn't respond to those challenges by just implementing one of the controller based NV solutions that were discussed in the

preceding chapter. As will be pointed out in the following discussion of a federated overlay/underlay model, it would potentially be possible for the hypothetical IT organization to respond to those challenges using a federation of NV overlay solutions and SDN solutions.

The challenges and opportunities that are identified in **Table 1** aren't dependent on any particular technology. For example, there are a number of technologies that can be implemented in order to ease the burden of configuration management. That said, a subsequent sub-section of this document identifies some of the specific use cases and benefits that are associated with the OpenFlow protocol.

While the use of SDN in data centers receives the majority of attention, it is also possible to implement SDN in branch and campus networks as well as in wide area networks (WANs). In order to understand where SDN will likely be implemented, The Survey Respondents were asked "If your organization is likely to implement SDN sometime over the next two years, where are you likely to implement it?" Their responses are summarized in **Table 2**.

Table 2: Focus of SDN Deployment

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

One observation that can be made from the data in **Table 2** is that while the primary interest in deploying SDN is focused on the data center, there is strong interest in deploying SDN broadly throughout an organization's entire network.

A Working Definition of SDN

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

The ONF is the group that is most associated with the development and standardization of SDN. According to the ONF[2], "Software-Defined Networking (SDN) is an emerging architecture that is dynamic, manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. This architecture decouples the network control and forwarding functions enabling the network control to become directly programmable and the underlying infrastructure to be abstracted for applications and network services. The OpenFlow™ protocol is a foundational element for building SDN solutions."

According to the ONF, the SDN architecture is:

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

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

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

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

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

The Survey Respondents were given a set of characteristics that are often associated with SDN and were asked to indicate which two characteristics would provide the most value to their company's network. Their responses are shown in **Table 3**.

Table 3: Value of SDN Characteristics

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

One observation that can be drawn from **Table 3** is that the characteristic of SDN that offers the most value to The Survey Respondents is tactical: The centralization of configuration and policy management. However, the second most important characteristic, the programmability of network elements, is strategic. That characteristic is strategic because the programmability of network elements is a key component of the overall functionality that is required in order to enable applications to dynamically request the network services they need.

Another observation that can be drawn from **Table 3** is that in spite of all of the discussion in the industry about open networking, The Survey Respondents were not very enthusiastic about the value that open protocols would bring to their networks.

The SDN Solution Architecture

Figure 1 contains a graphical representation of the SDN architecture as envisioned by the ONF. One key component of a complete SDN solution that is missing from **Figure 1** is cloud orchestration platforms such as OpenStack. The role that these platforms play in both NV and SDN solutions was described in the preceding section of The Guide.

Figure 1: ONF's SDN Architecture

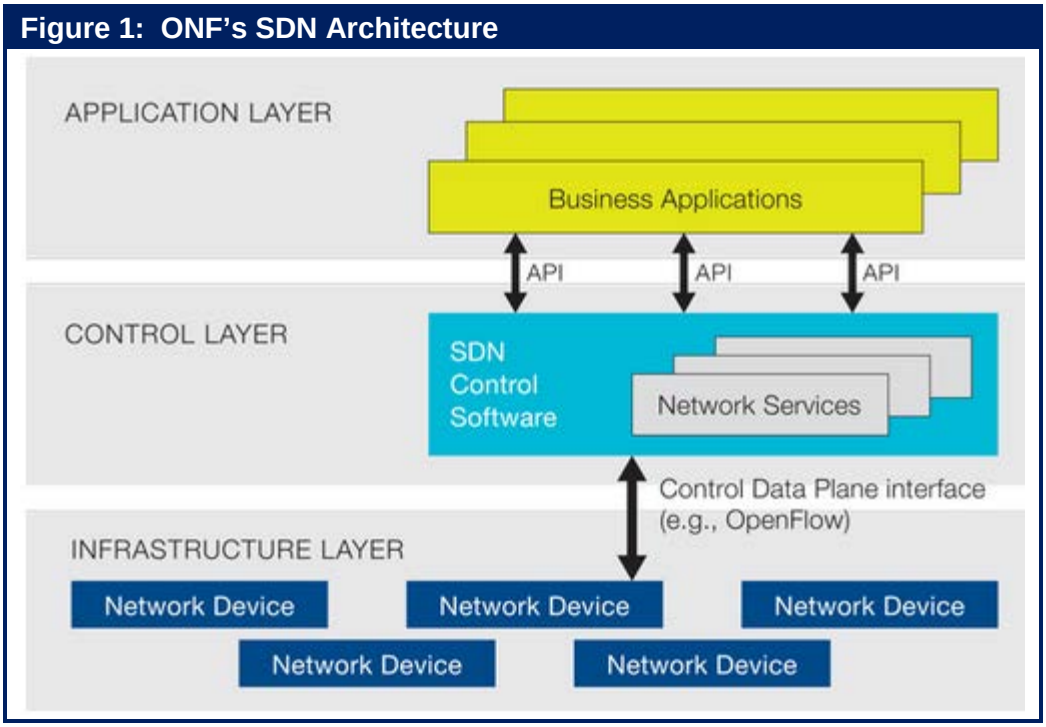

Below are definitions of some terms that are commonly associated with SDN, some of which appear in **Figure 1**.

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

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

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

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

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

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

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

- **Southbound API**
  Relative to **Figure 1**, the southbound API is the API that enables communications between the control layer and the infrastructure layer.

- **Service Chaining**[3]
  Service chaining is the ability to steer VM-VM traffic flows through a sequence of physical or virtual servers that provide network services, such as firewalls, IPS/IDS, DPI, or load balancers.

Figure 1 shows an API between the SDN control layer and the business applications. This API is commonly referred to as *the Northbound API*. The role of the northbound API is to enable communications between the control layer and the application layer. Currently there isn't a standard for the Northbound API, although the ONF has recently begun a process that could lead to a standards based API. While it isn't possible to state how the development of the northbound API will evolve, it is likely that there won't be a single northbound API, but multiple northbound APIs. One viable alternative is that there will be a northbound API between the SDN control software and each of the following entities:

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

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

Criteria to Evaluate SDN Solution Architectures

Below is a set of 7 questions that IT organizations should ask vendors who provide all or the majority of the SDN solution architecture that is shown in **Figure 1**. These questions focus on key criteria that IT organizations should use relative to evaluating alternative SDN solutions. A more complete set of criteria can be found in *A Mock RFI for SDN Solutions*[4].

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

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

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

2. Describe the SDN solution that you are proposing and include in that description how the SDN architecture for the solution you are proposing is similar to the architecture shown in **Figure 1** and also describe how it is different. In your answer, identify the southbound protocols that you support and provide the rationale for supporting those protocols.

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

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

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

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

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

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

4 Will be published at: webtorials.com/Metzler

The Inhibitors to SDN Adoption

The left hand column of **Table 4** contains some of the primary impediments to the adoption of SDN. The Survey Respondents were shown these impediments and were asked to indicate the two impediments that would be the biggest inhibitors to their company adopting SDN sometime in the next two years. Each row of the right hand column of **Table 4** contains the percentage of The Survey Respondents that indicated that the impediment in the corresponding left column was one of the two primary inhibitors.

Table 4: Inhibitors to the Adoption of SDN

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

One clear observation that can be drawn from **Table 4** is that immaturity, broadly defined, is the primary inhibitor to the adoption of SDN. That includes the immaturity of the current products, the immaturity of the enabling technologies and the confusion and lack of definition in terms of vendor strategies.

The role that a compelling business case plays relative to driving and inhibiting the adoption of SDN is somewhat subtle. As shown in **Table 4**, only 7% of *The Survey Respondents* indicated that the lack of a compelling value proposition was an inhibitor to their adoption of SDN. It would be easy to conclude from that metric that business cases that demonstrate the compelling value of SDN exist and that these business cases are widely understood. Drawing the conclusion would be a mistake.

Arguing against that conclusion is the fact that 24% of *The Survey Respondents* indicated that "other technology and/or business priorities" was an inhibitor and that 21% of *The Survey Respondents* indicated that "the lack of resources to evaluate SDN" was an inhibitor. If indeed, there were compelling, well-understood SDN business cases, these organizations would rearrange their priorities and find the resources to evaluate SDN solutions.

The Overlay/Underlay Model

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

- Virtual and physical networks are separate entities, possibly with separate service assurance solutions, policy management, provisioning, and control points.

- As the virtual networks grow and evolve, the physical network does not automatically adapt to the changes. As a result, overlay NV requires a lightly oversubscribed or non-oversubscribed physical underlay network.

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

An emerging approach to overcome the limitations of the overlay model is referred to as an overlay/underlay model. The cornerstone of this approach is a federation between the overlay network virtualization controller and the underlay SDN controller. In August 2013, HP and VMware announced their intention to work together to create an overlay/underlay solution[5]. As part of that announcement, HP stated their intention to develop a new application called ConvergedControl that will enable HP's Intelligent Management Center (IMC) to share information about the network with both the HP and the VMware controllers. As part of the announced solution, VMware's NSX controller will continue to provision the virtual network overlay and HP's SDN controller will continue to provision physical network flows on its switches. The solution is intended to enable the two controllers to work together to ensure that the virtual network gets the physical flows it needs. The solution is also intended to provide visibility across the virtual and physical environment so that, for example, if there is congestion or a failure on the physical network, the virtual environment is aware of the issue and can respond accordingly.

Network Function Virtualization

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

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

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

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

In order to bring this vision to fruition, an Industry Specifications Group for Network Functions Virtualization (NFV ISG) has been formed under the auspices of the European Telecommunications Standards Institute (ETSI). Their vision for the transition from hardware appliances of today to a fully virtualized appliance environment is depicted in **Figure 2**.

Figure 2: The Virtualization of Network Appliances

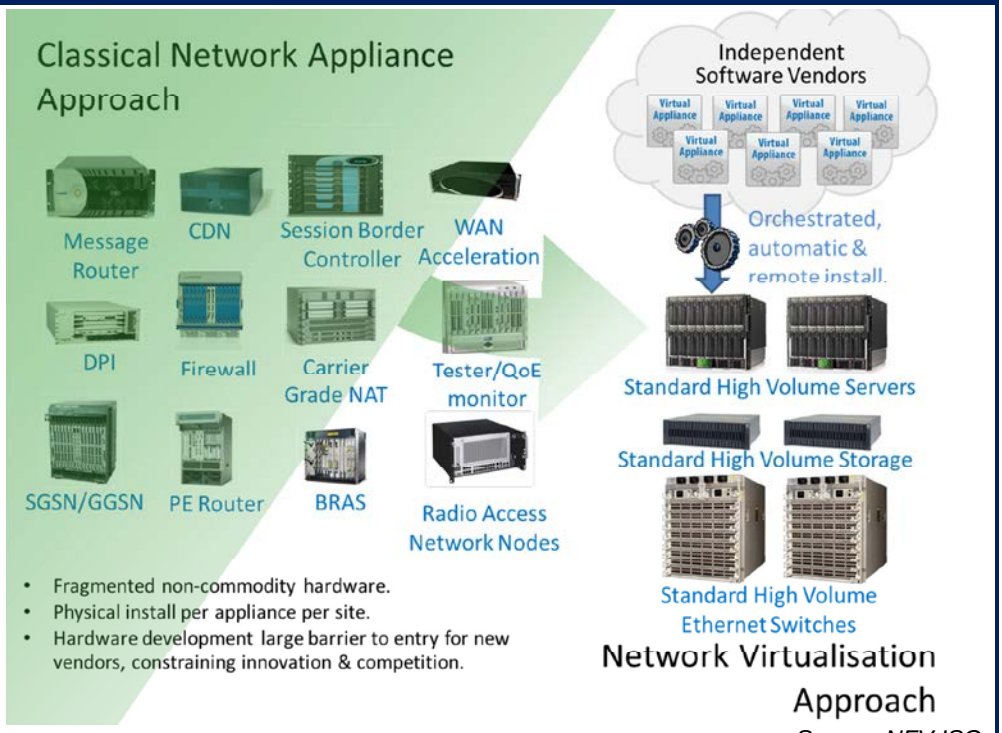

Source: NFV ISG

The approach that the NFV ISG is taking is that the virtualization of network functionality is applicable to any data plane packet processing and control plane function in both fixed and mobile networks. As shown in **Figure 2**, examples of these functions include:

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

The initial members of the NFV ISG were service providers such as AT&T, Deutsche Telekom and NTT. Its membership[6] has since grown and now includes a number of equipment vendors, but currently relatively few of the top vendors of virtual appliances are members.

6 http://portal.etsi.org/NFV/NFV_List_members.asp

The first meeting of the group was held in January 2013 and a number of smaller working groups were created in April 2013. In October 2013, ETSI published the first five specifications relative to NFV[7]. According to ETSI[8], "The five published documents include four ETSI Group Specifications (GSs) designed to align understanding about NFV across the industry. They cover NFV use cases, requirements, the architectural framework, and terminology. The fifth GS defines a framework for co-ordinating and promoting public demonstrations of Proof of Concept (PoC) platforms illustrating key aspects of NFV. Its objective is to encourage the development of an open ecosystem by integrating components from different players."

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

7 http://www.etsi.org/technologies-clusters/technologies/nfv

8 http://www.etsi.org/index.php/news-events/news/700-2013-10-etsi-publishes-first-nfv-specifications

The Open Networking Foundation and OpenFlow

The Open Networking Foundation

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

Figure 3 shows the ONF organizational model. More information on the ONF working and study groups as well as the activities that the ONF is sponsoring can be found at the ONF web site[9].

Figure 3: The ONF Organizational Model

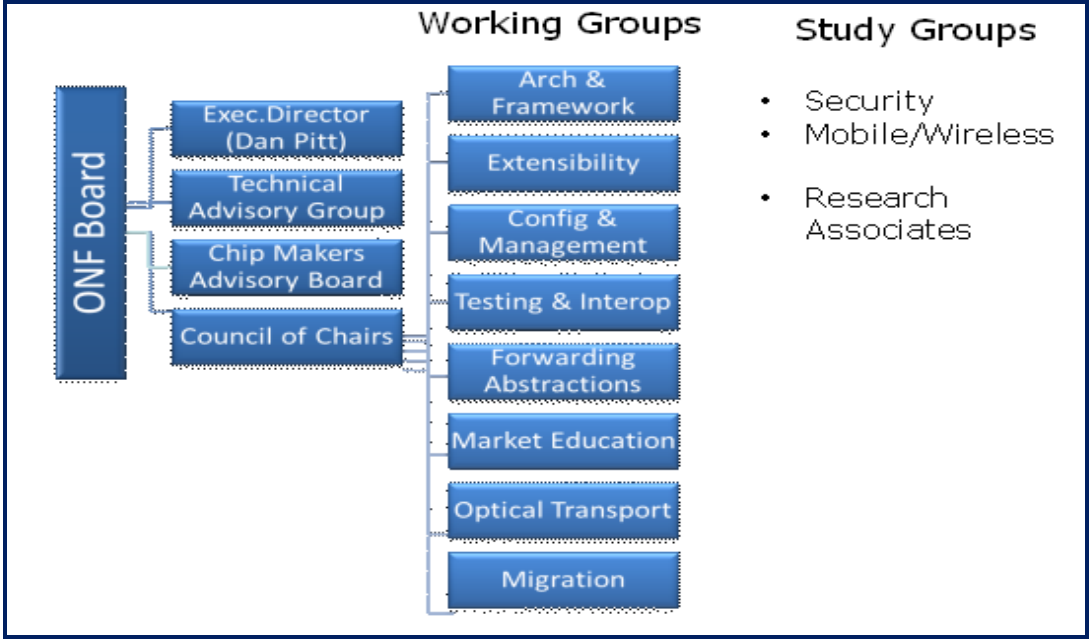

The OpenFlow Protocol

Referring back to **Figure 1** (ONF's SDN Architecture), OpenFlow is a standards-based protocol that enables an SDN controller to program the behavior of an OpenFlow-enabled switch. OpenFlow V1.0 was developed by Stanford University and was published in December 2009. The basic elements of an OpenFlow V1.0 network are shown on the left hand side of **Figure 4**. The central controller communicates with the switch's OpenFlow agent over a secure TLS (Transport Layer Security) channel. This channel could be either in-band or out-of-band. The OpenFlow agent on the switch populates the flow table as directed by the controller. Note that within **Figure 4**, OpenFlow is referred to as OF.

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

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

Figure 4: The Major Components of an OpenFlow Switch V1.0-V1.3

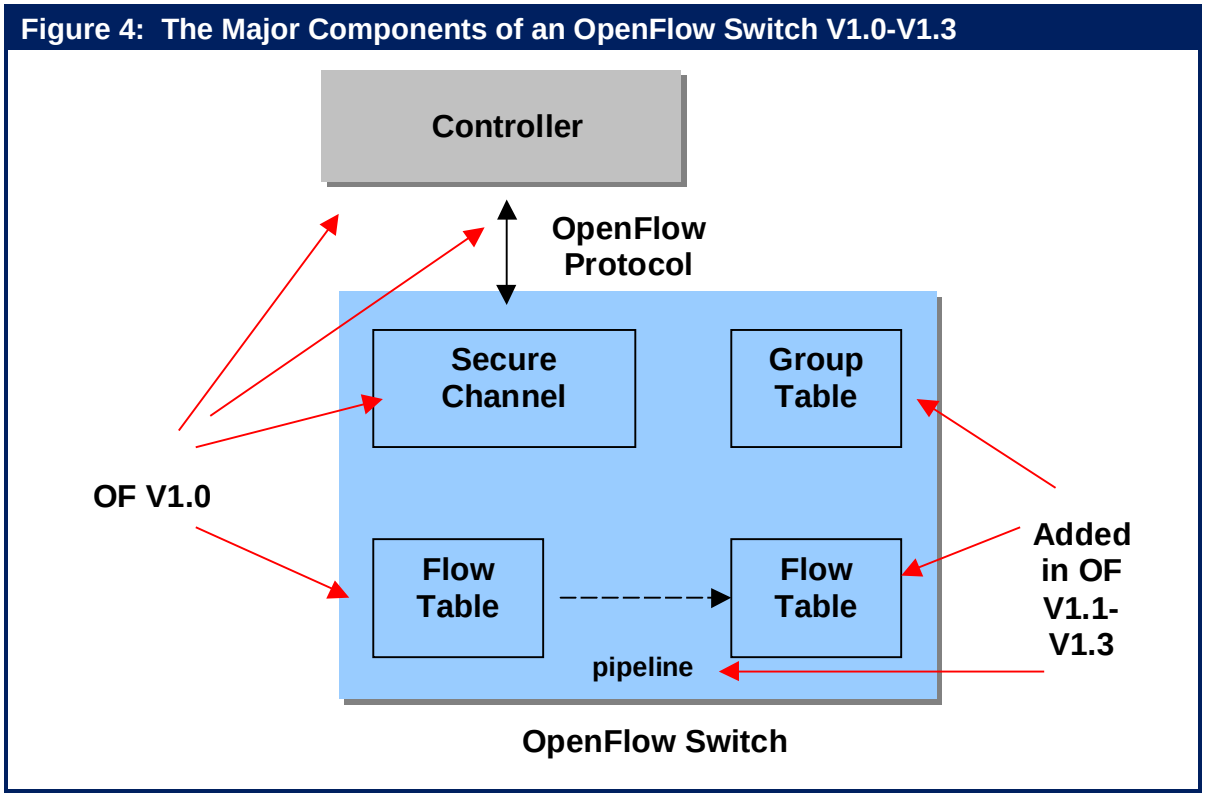

Throughout most of 2012, SDN and OpenFlow were tightly linked in the trade press as if they were either the same thing, or as if OpenFlow was required in order to implement an SDN. Neither statement is true. OpenFlow is one possible protocol that can be used to implement an SDN. In order to understand how IT organizations currently view OpenFlow, The Survey Respondents were given a set of options and were asked to indicate which option best describes the role that the OpenFlow protocol will play in their company's implementation of SDN. Their possible options and the percentage of the respondents who indicated that option are shown in **Table 5**.

Table 5: Planned Use of OpenFlow

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

The data in **Table 5** indicates that there is strong interest in using the OpenFlow protocol as part of implementing an SDN. The data also shows, however, that there is still a high level of uncertainty and whether or not OpenFlow will be used.

Potential Use Cases and Benefits of OpenFlow

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

Centralized FIB/Traffic Engineering
One of the primary benefits of OpenFlow is the centralized nature of the Forwarding Information Base (FIB). Centralization allows optimum routes to be calculated deterministically for each flow by leveraging a complete model of the end-to-end topology of the network. This model can be build using a discovery protocol, such as the Link Layer Discovery Protocol (LLDP). Based on an understanding of the service levels required for each type of flow, the centralized OpenFlow controller can apply traffic engineering principles to ensure each flow is properly serviced. Bandwidth allocations can be controlled dynamically to provide bandwidth on demand with changing traffic patterns. The result can be much better utilization of the network without sacrificing service quality. Centralized route processing also allows the pre-computation of a set of fail-over routes for each possible link or node failure. Centralized processing also can take advantage of the virtually unlimited processing power of multi-core processors and cluster computing for calculating routes and processing new flows. As shown in **Table 1**, being able to do end-to-end traffic engineering is one of the top three opportunities that The Survey Respondents associate with SDN.

The Google G-Scale WAN backbone links its various global data centers. G-Scale is a prime example of a production OpenFlow Layer 3 network that is realizing the benefits of FIB centralization. The G-Scale control plane is based on BGP and IS-to-IS and the OpenFlow-only switches are very simple 128 port 10 GbE switches built by Google using merchant silicon (when Google built these switches, 128 port 10 GbE switches had not yet been introduced in the commercial market). Google has identified a number of benefits that are associated with its G-Scale WAN backbone including that Google can run the network at utilization levels up to 95%[10]. As shown in **Table 1**, being able to increase resource utilization is the primary opportunity that The Survey Respondents associate with SDN.

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

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

OpenFlow-Based Virtual Networking
With OpenFlow V1.3 virtual ports, an OpenFlow switch can be programmed to perform tunnel encapsulation and de-capsulation. Therefore, an OpenFlow switch can be programmed to be a overlay NV VTEP/NVE or gateway, as described in the section on overlay NV. As also described in that section, OpenFlow can provide another type of network virtualization for isolating network traffic based on flows segregation or segmentation. Flows are separated based on a subset of the match fields listed earlier in the section.

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

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

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

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

Indiana University has developed an OpenFlow-based, load-balancing application called FlowScale. According to the University[11], "FlowScale provides complex, distributed load balancing of network traffic using an OpenFlow-capable Top of Rack (ToR) switch. IU deployed the application into its Intrusion Detection System (IDS) to distribute traffic evenly to sensors. When fully deployed, the system will span the IU Bloomington and IUPUI networks and have the capability to distribute traffic at rates exceeding 500Gb/s."

Network Taps
With OpenFlow virtual ports, the functionality of a network tap can be programmed into the OpenFlow switch, allowing selected traffic to be monitored without deploying physical taps. Traffic can also be replicated and redirected to any monitoring device in the network. Big Switch networks has announced such a network monitoring application referred to as Big Tap[12].

Service Insertion/Chaining
OpenFlow's ability to dynamically reroute flows allows network services provided by physical or virtual appliances (e.g., firewalls, NATs, load balancers, and WOCs) to be inserted in the path of the flow.

[11] http://incntre.iu.edu/research/flowscale
[12] http://www.bigswitch.com/blog/2013/07/26/network-monitoring-with-big-tap-your-first-sdn-application

Redirecting the flow to the next service can be based on encapsulation or rewrite of the destination MAC address.

Circuit Provisioning

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

The OpenDaylight Consortium

The OpenDaylight Consortium[13] was founded in April 2013. The consortium's stated mission is to facilitate a community-led, industry-supported open source framework, including code and architecture, to accelerate and advance a common, robust Software-Defined Networking platform.

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

Figure 5: Members of the OpenDaylight Consortium

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

The OpenDaylight Consortium has announced its intention for the first release of code. That code release is called *Hydrogen* and is expected to occur in December 2014. **Figure 6** depicts the

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

OpenDaylight SDN Architecture and indicates some of the functionality that will be included in the first code release.

Figure 6: OpenDaylight SDN Architecture

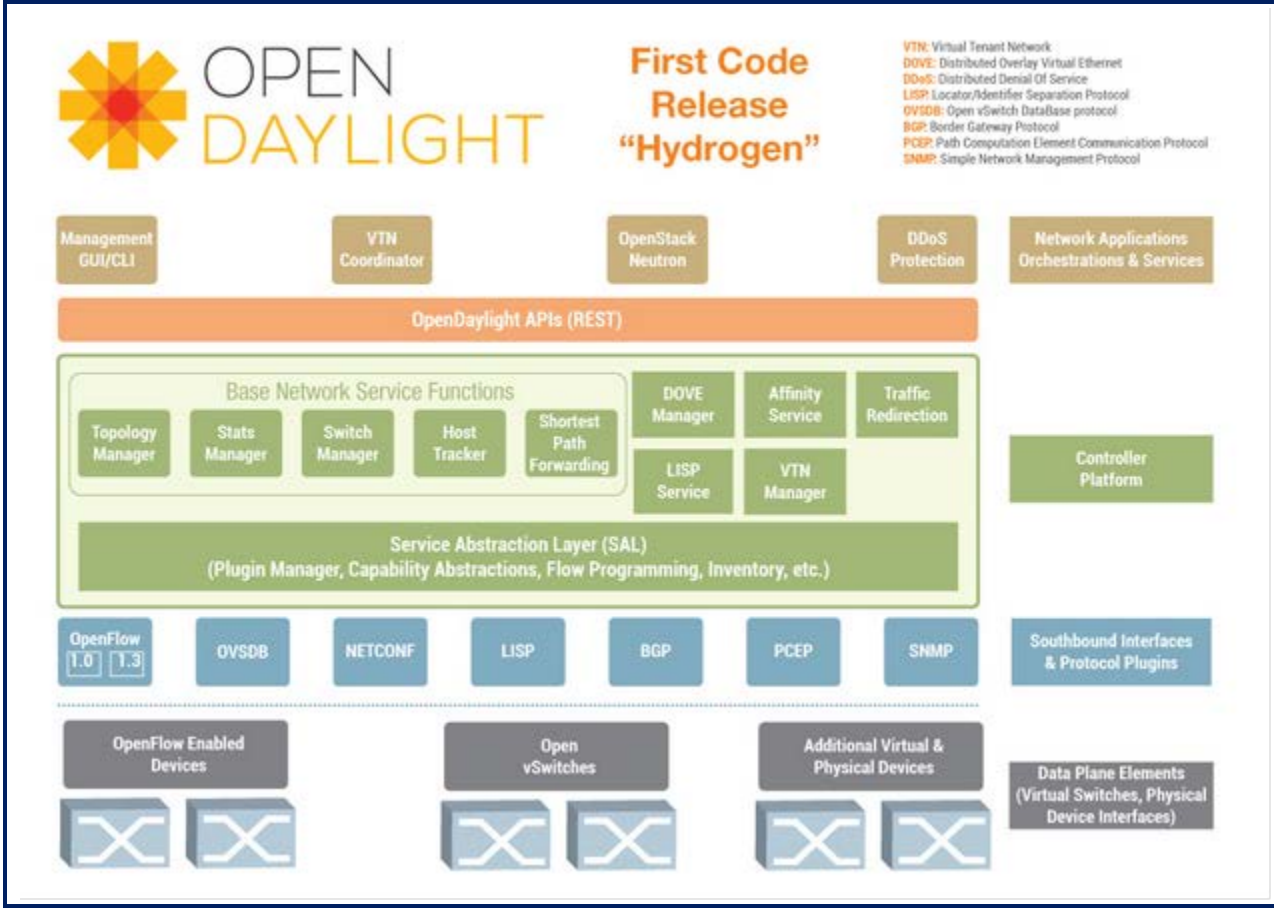

Some vendors, such as Cisco, have announced that they will use the code produced by the OpenDaylight Consortium as the basis for their SDN controller. Other vendors are taking a wait and see attitude.

Security

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

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

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

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

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

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

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

As noted, in addition to creating security challenges, SDN also presents opportunities to improve security by implementing security related applications that leverage the control information that has been centralized in the SDN controller. One example of such an application is DefenseFlow that was recently announced by Radware[14]. Relative to the terminology of **Figure 1**, DefenseFlow is a network service that provides DDoS protection. Another such example is HP's Sentinel application[15] that was designed to combat the security challenges that are associated with BYOD by leveraging the HP TippingPoint Repudiation Digital Vaccine data base.

To quantify the concern that IT organization have relative to security, The Survey Respondents were given the following question. "Some in the industry suggest that the implementation of SDN will make organizations less secure because if the SDN controller is hacked, the hacker has access to all of the subtending switches. Others argue that new security-oriented applications will be developed that take advantage of the SDN controller and make organizations more secure. What is the overall impact that you believe that SDN will have on network security? (Choose only one.)" Their responses are shown in **Table 6**.

[14] http://www.radware.com/Products/DefenseFlow/
[15] http://h20195.www2.hp.com/v2/GetPDF.aspx/4AA4-7496ENW.pdf

Table 6: Perceived Impact of SDN on Security	
Impact	**Percentage**
Networks will be much more secure	7%
Networks will be somewhat more secure	31%
It will have no impact on network security	20%
Networks will be somewhat less secure	20%
Networks will be much less secure	3%
Don't know	19%

One observation that can be drawn from the data in **Table 6** is that overall The Survey Respondents believe that SDN will have a positive impact on security.

Management

As is the case with security, SDN presents both management opportunities and management challenges. One of the primary opportunities was highlighted in **Table 3**. That table showed the characteristic of SDN that The Survey Respondents stated would provide the most value to their company's network was the centralization of configuration and policy management. In addition, as previously described, new network management applications, such as network taps, that leverage SDN functionality are now coming to market.

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

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

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

As was previously mentioned, one of the characteristics of NV and SDN is that network functions such as load balancing and firewalls are increasingly implemented in software as network services that can integrated with virtual networks or SDN flows under programmatic control; a.k.a., service chaining. Implementing these functions in software both increases the delay associated with performing these functions and it also increases the variability of that delay. The result is an increased need for insight into the performance of each component of the overall SDN solution.

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

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

2. Describe the ability of your solution to monitor the SDN controller. Include in that description your ability to monitor functionality such as CPU utilization as well as flow throughput and latency. Also describe the statistics you collect on ports, queues, groups and meters; and the

error types, codes and descriptors you report on. Also, does your solution monitor the number of flow set-ups being performed by the SDN controller?

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

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

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

The Survey Respondents were asked to indicate how much of an impact they thought that SDN will have on network management. Their responses are shown in **Table 7**.

Table 7: Perceived Impact of SDN on Management

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

One observation that can be drawn from the data in **Table 7** is that the vast majority of The Survey Respondents believe that SDN will have a positive impact on management.

Appendix

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

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

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

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

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

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

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

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

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

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

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

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

Match Fields	Counters	Instructions/Actions	Priority

- Flexible pipeline processing through multiple flow tables, as shown in the right hand side of **Figure 4**. As a packet is processed through the pipeline, it is associated with a set of accumulating actions and metadata. The action set is resolved and applied at the end of the pipeline. The metadata allows a limited amount of state to be passed down the pipeline.

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

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

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

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

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

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

- Support for multiple controllers to improve reliability

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

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

Professional Opinions Disclaimer
All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

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

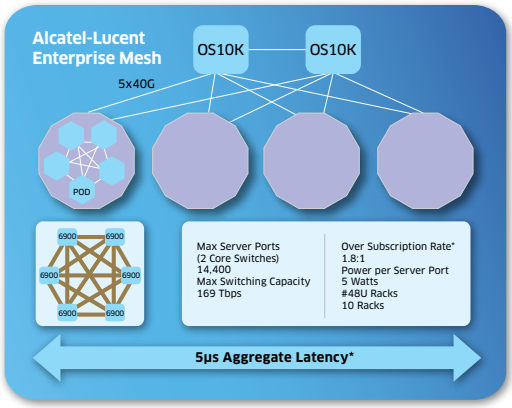

Alcatel-Lucent provides a unique Application Fluent approach to maximize the benefit from virtualization technologies for servers, the desktop, as well as the network. Alcatel-Lucent's application fluent data center fabric can scale from several hundred to over 14,000 server facing ports while keeping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

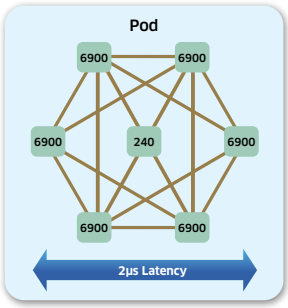

Application fluency in the corporate data center includes its transformation into a multi-site private cloud by extending layer 2 connectivity between data center sites and allowing for seamless delivery of public cloud-based services on the corporate network.

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

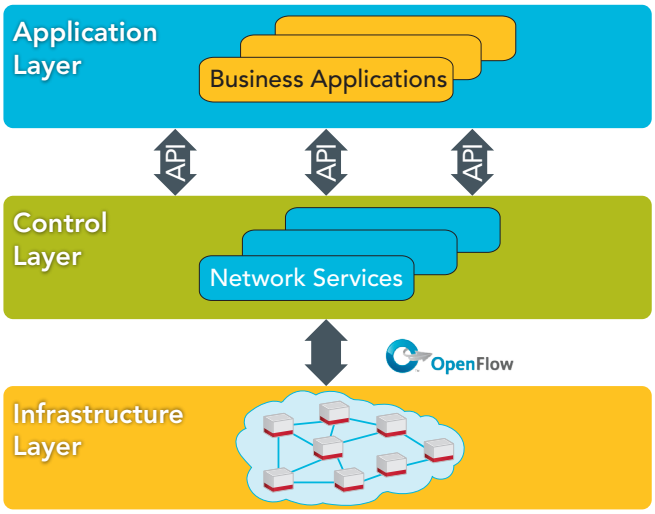

Figure 1. ONF SDN Architecture

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

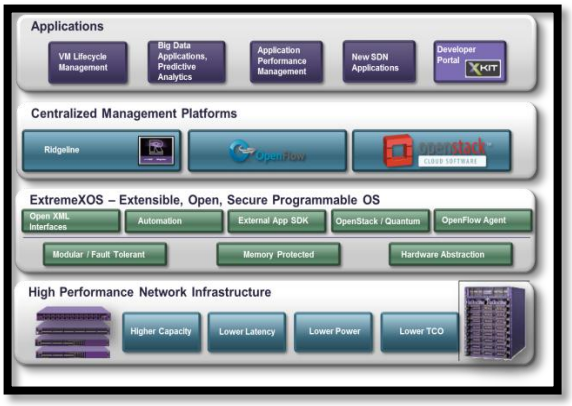

Critically important to the Open Fabric is ExtremeXOS®, the network operating system that delivers the consistent set of features across all platforms while ensuring the security and performance of the Open Fabric. ExtremeXOS is modular, extensible, and has integrated security, while providing a single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

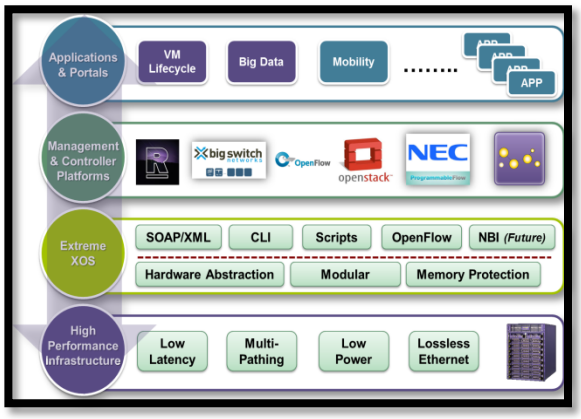

The Extreme Networks Open Fabric SDN strategy therefore extends to include technology partners and systems that leverage the network abstraction capabilities provided by ExtremeXOS.

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

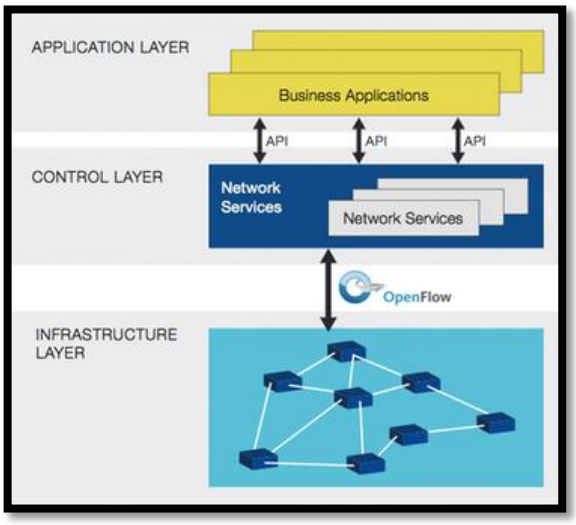

The Open Networking Foundation (ONF) defines OpenFlow: "The physical separation of the network control plane from the forwarding plane, and where a control plane controls several devices."

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

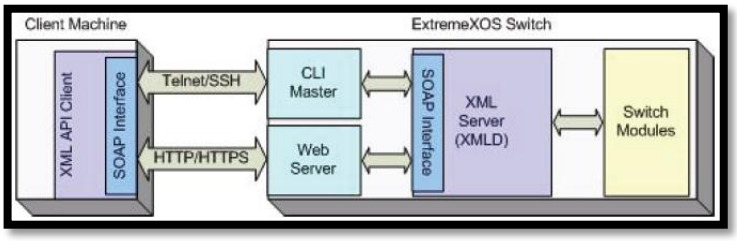

"Open API's enable applications and management systems to directly access the network abstraction layer to manage the control, data and management planes of the infrastructure."

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

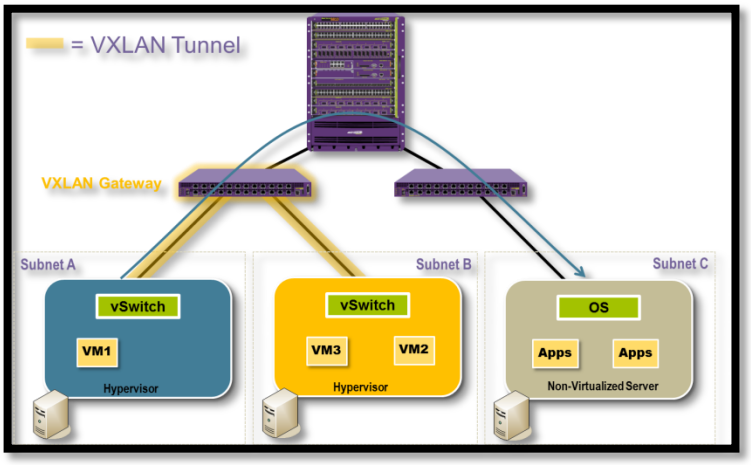

"Network Virtualization Overlays include logical overlays from virtualized server and switching systems that may also include virtualized layer 4-7 services."

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

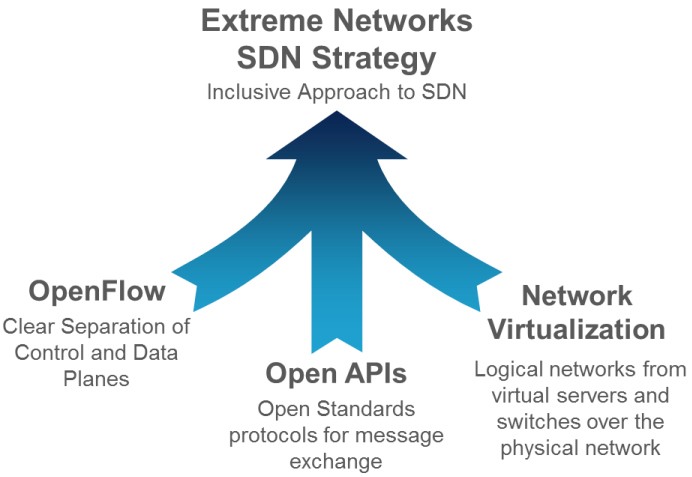

ProgrammableFlow

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

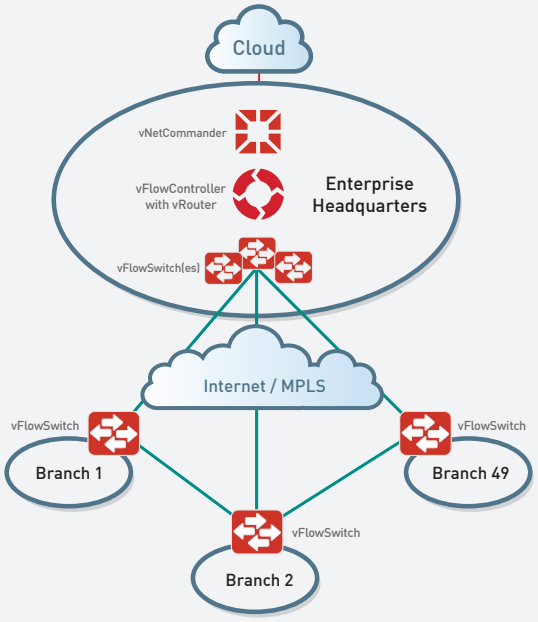

NVN significantly reduces lifecycle CAPEX/OPEX beyond that of traditional site-by-site-managed networking solutions. Immediate benefits include CAPEX savings of 3:1 and OPEX savings of 5:1 over single-purpose, hardware-based legacy networking solutions.

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

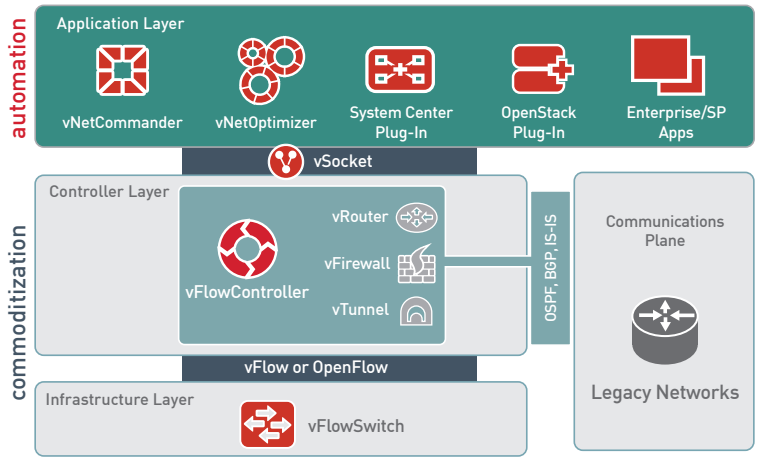

Making The Business Case — Netsocket Virtual Network for Distributed Enterprises

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™

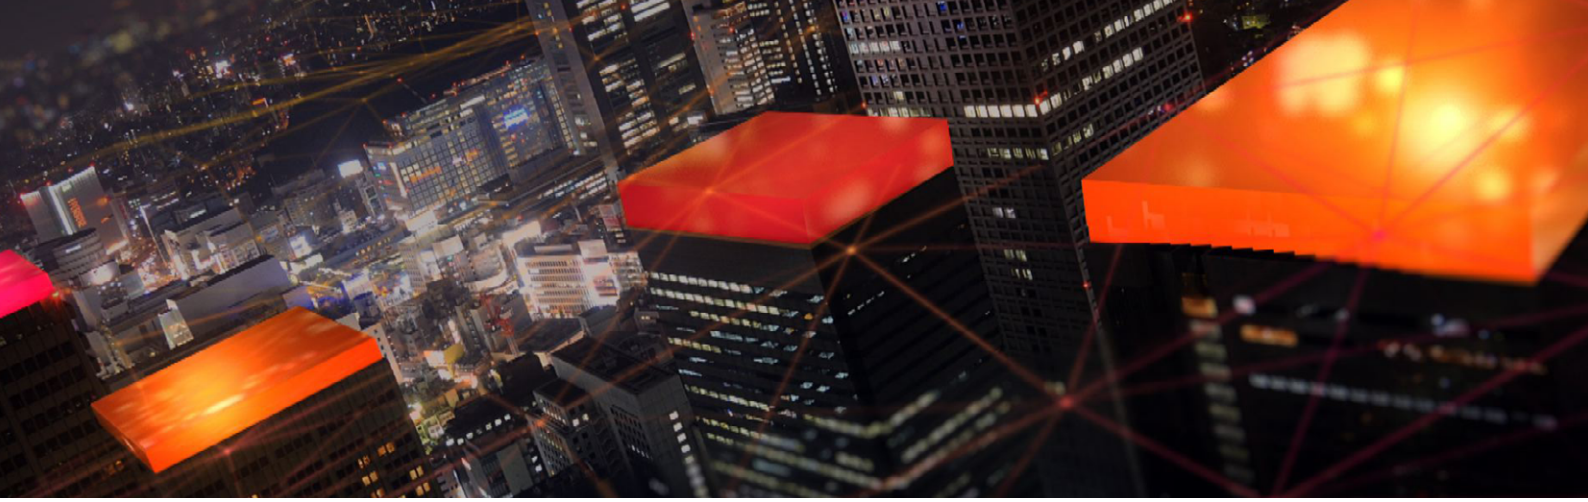

The Consumable Datacenter Network

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

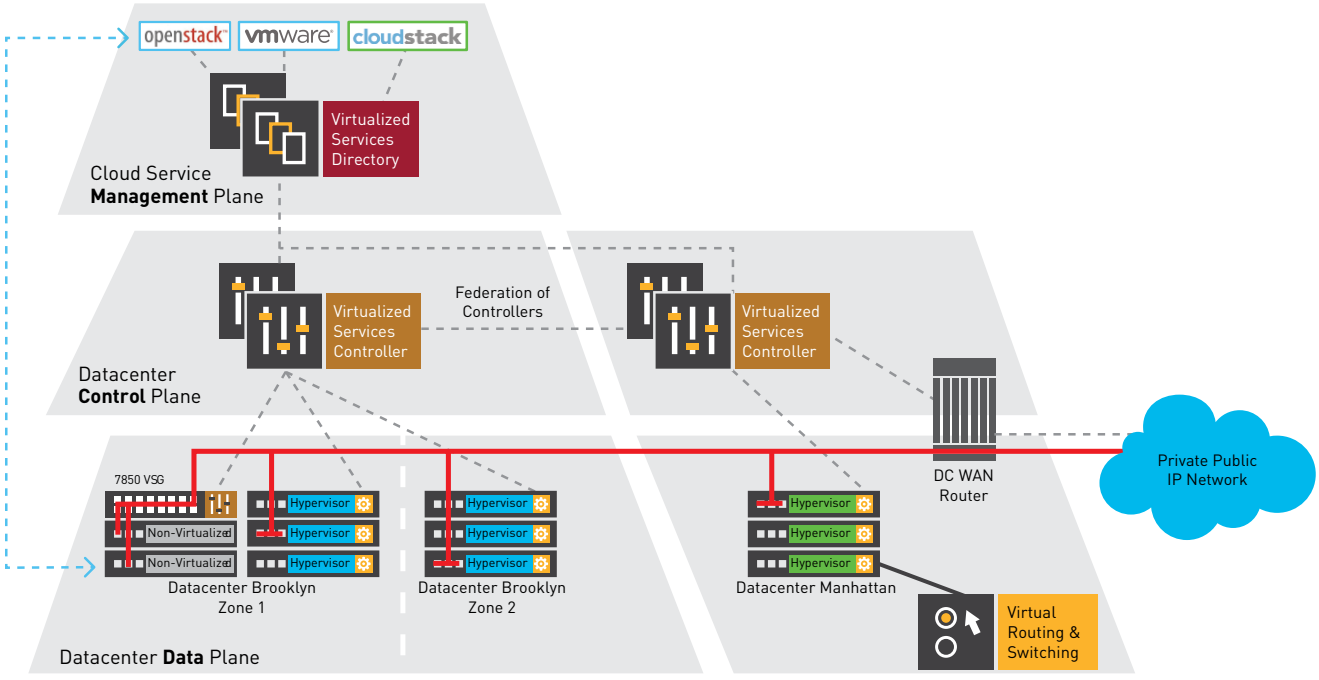

consistently, and in a repeatable manner. The critical need for mobility is also addressed, ensuring network services adjust gracefully and instantly as application endpoints and workloads move from virtual machines within or across datacenters.

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.

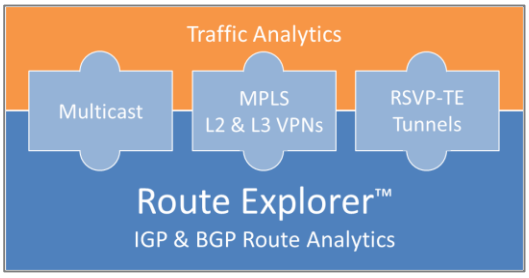

The Route Explorer System

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.

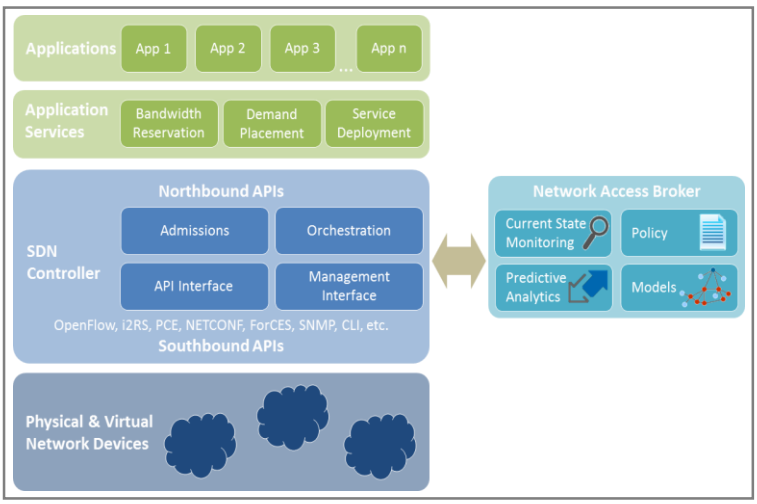

Network Access Broker for SDN

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

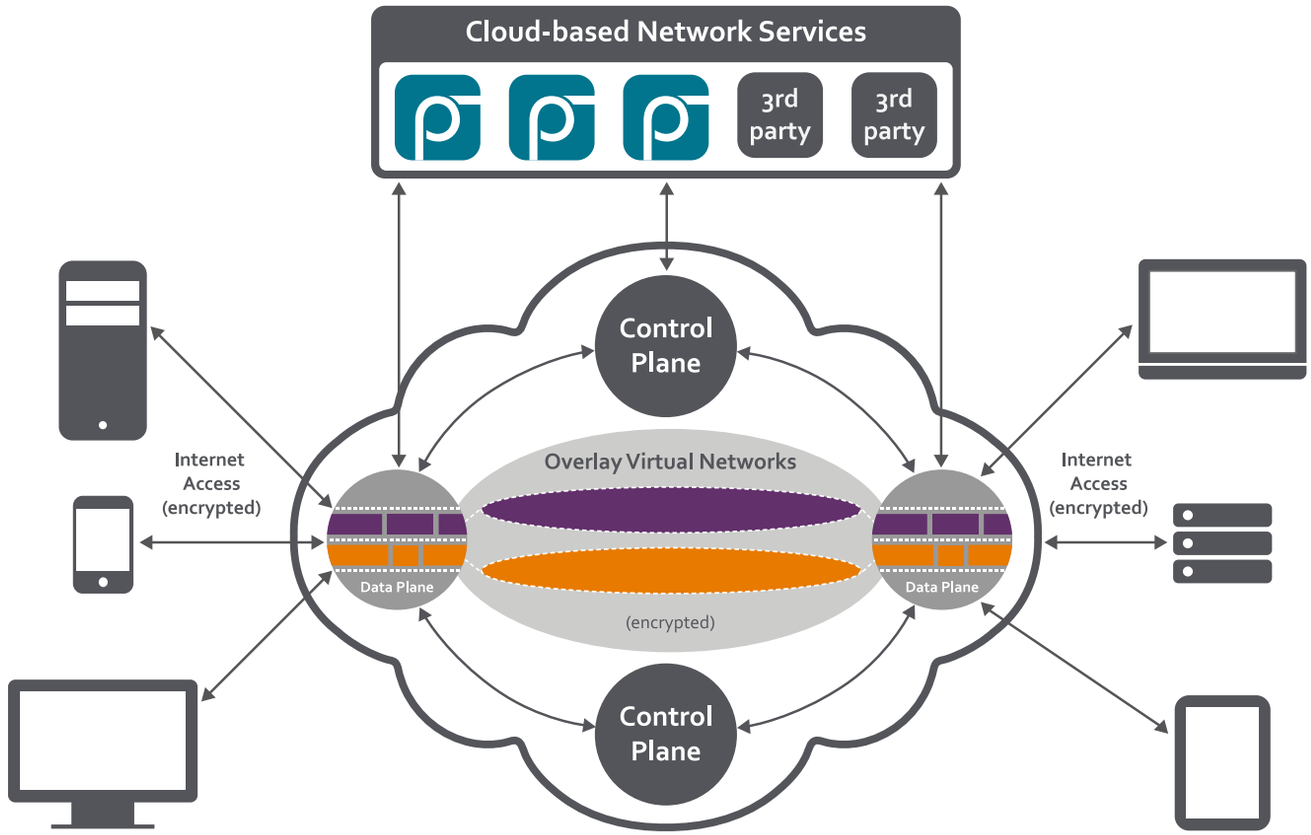

rackspace HOSTING amazon web services Windows Azure linode

pertino.com 408.354.3900 info@pertino.com

PICA8
OPEN NETWORKING

WWW.PICA8.COM

Open Systems for Software Defined Networking (SDN)

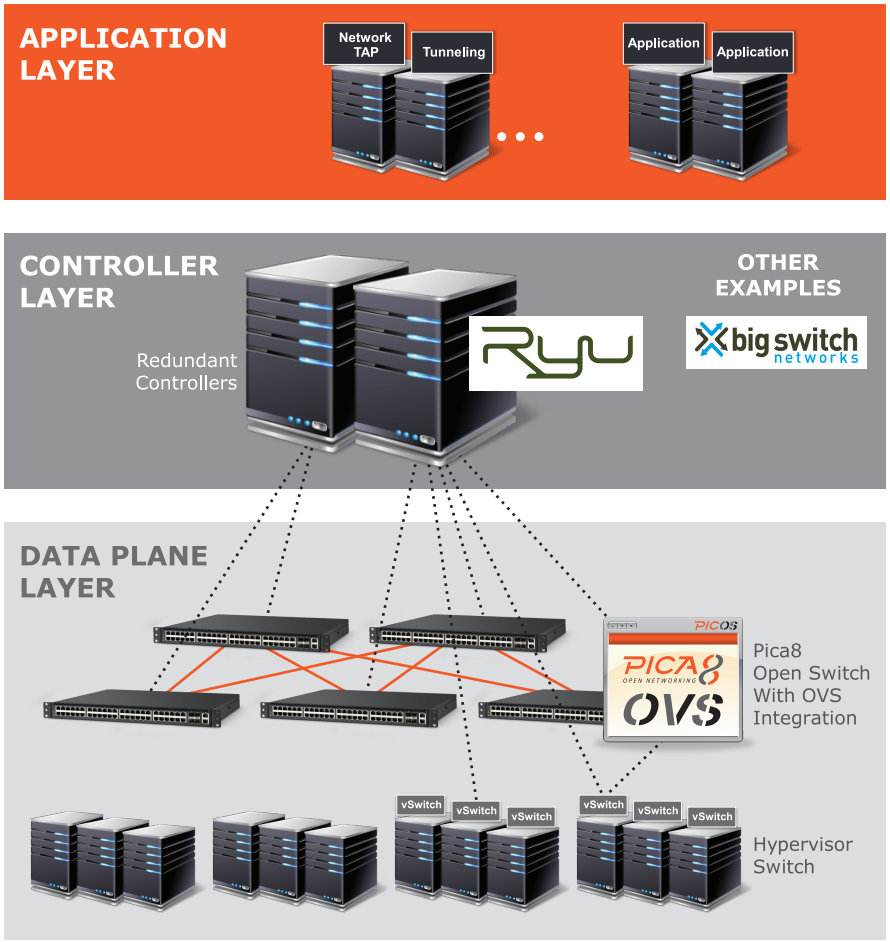

The First Hardware Agnostic, Open Network Operating System

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

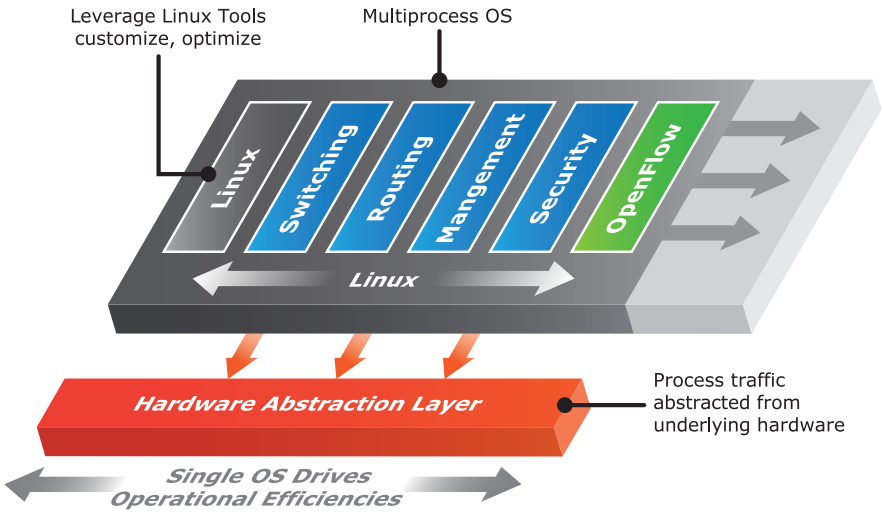

* Only OpenFlow features available in hardware are supported, to ensure optimum performance

Automation for Agile Infrastructure

Corporate Overview

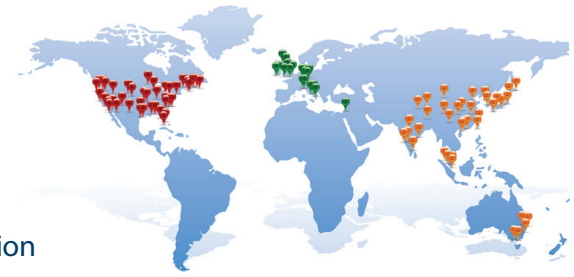

Founded: 2004

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

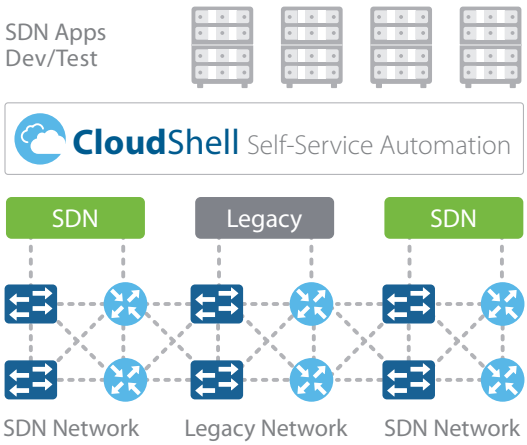

TestShell

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

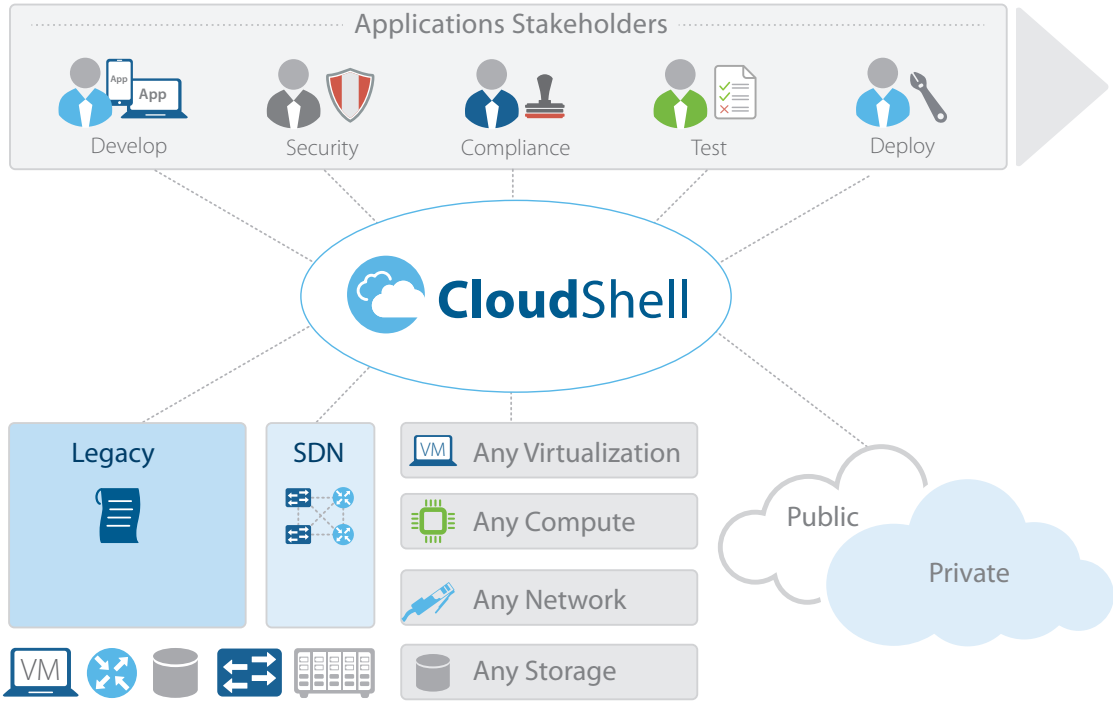

For more information about QualiSystems, visit our website at
www.qualisystems.com

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

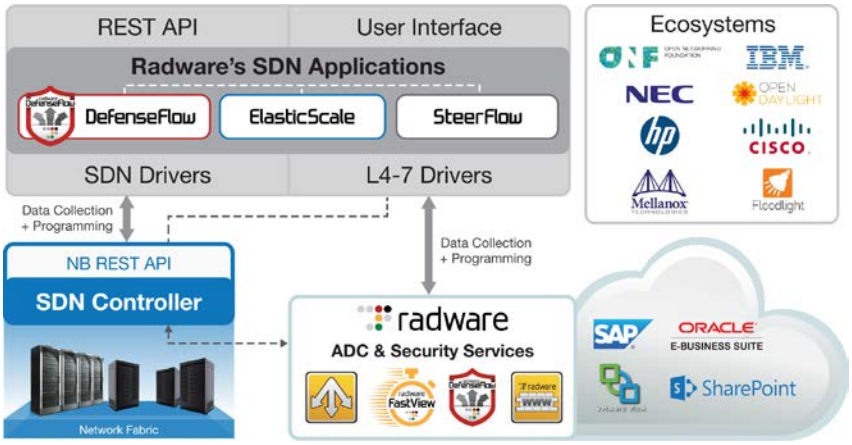

Partnering for Success: Our SDN Ecosystem

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.