The 2015 Guide to SDN and NFV

Part 1: Software Defined Networking (SDN)

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

Platinum Sponsors:

Gold Sponsors:

Produced by:

Table of Contents

Executive Summary

Over the last couple of years, the hottest topics in networking have been Software Defined Networking (SDN) and Network Functions Virtualization (NFV). While both enterprises and service providers have shown great interest in these topics, the vast majority of organizations are still either on the sidelines or in the analysis stage of adoption. The primary goals of **The 2015 Guide to SDN & NFV** (The Guide) are to eliminate the confusion that surrounds SDN and NFV and to accelerate the analysis and potential adoption of these new design and architectural approaches.

The Guide will be published both in its entirety and in a serial fashion. This document is the first of the serial publications and it will focus on SDN. The second publication will focus on NFV and the third on the SDN and NFV ecosystem. The fourth publication will include a complete copy of The Guide as well as an executive summary of the complete guide.

This document has three chapters:

1. **Introduction**
   Provides background on topics such as the current status of SDN deployment and the interest in both the fabric-based and the overlay based models.

2. **Use Cases**
   Discusses SDN use cases in the data center, WAN and campus and analyzes the factors that are currently limiting deployment.

3. **Operational Considerations**
   Identifies the security, management and organization issues relative to implementing SDN and analyzes how impactful these issues are likely to be in the near term.

Introduction

This chapter of The Guide is based in part on *The 2013 Guide to Network Virtualization and SDN* (The 2013 Guide). To limit the size of this chapter, some of the introductory SDN material that was contained in The 2013 Guide has been eliminated. That document, however, is still available online. Also with the goal of limiting the size of this chapter, detailed analyses of a number of topics are avoided and URLs are provided that point to relevant material. That material includes:

- An analysis of OpenFlow V1.3 and the use cases it enables;
- Criteria to evaluate a vendor's overall SDN solution as well as specific criteria to evaluate a SDN controller and the subtending network devices;
- A framework to plan for SDN;
- An analysis of the advantages and disadvantages of the overlay-based SDN model;
- Criteria to evaluate overlay-based SDN solutions.

This section contains the results of a survey that was distributed in September 2014 (The 2014 Survey). Throughout The Guide, the 176 network professionals who completed the survey will be referred to as The Survey Respondents. Where appropriate, the results of The 2014 Survey will be compared to the results of a similar survey given in 2013 (The 2013 Survey).

Thirty-two percent of The Survey Respondents indicated that they were either very familiar or extremely familiar with SDN. In response to The 2013 Survey, only twenty-one percent of the respondents indicated that they were either very familiar or extremely familiar with SDN.

Over the last year, the familiarity with SDN has increased significantly.

Definition of SDN

The Open Networking Foundation (ONF) is the group that is most associated with the development and standardization of SDN. According to the ONF, "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." Many vendors have announced support for OpenFlow V1.3. An overview of that protocol and the use cases it enables can be found in *An Overview of OpenFlow V1.3*.

The ONF states that the SDN architecture is:

- Directly programmable;
- Agile;
- Centrally managed;
- Programmatically configured;
- Open standards-based and vendor-neutral.

Context for SDN

As described later in this document, one of the key SDN use cases is the engineering of data traffic. A high level metaphor that both explains the value that SDN brings to engineering data traffic and that also provides insight into SDN's overall value proposition stems from the world of vehicular traffic.

In terms of the engineering of vehicular traffic, things are largely the way they were upon the introduction of the traffic signal. A traffic signal is a good thing in that it helps vehicles to avoid a collision and it gives priority to higher volume roads. With the exception of HOV lanes and toll roads, until recently there has been very little else to assist in improving traffic flow. That situation changed several years ago with the introduction of GPS and real time maps which together provide a connected driver with the information that enables that driver to take a less congested route, which presumably results in a shorter travel time.

The metaphor is that in a computer network a packet is similar to a car in part because it has an origin and a destination and in part because the switches and routers along the end to end path from origin to destination play a role somewhat similar to traffic signals and road signs. The computer network also now has the equivalent of an HOV lane for important traffic and it is getting better at routing that traffic in ways that reduce travel times.

As mentioned, one of the key SDN use cases is traffic engineering. In addition, a defining characteristic of SDN is that it separates the control of the network from the process of forwarding the packets. Staying with the metaphor, one way to think about how SDN concepts could be applied to vehicular traffic involves thinking not of a traditional car, but of a Google-inspired, driverless car. Before it starts to move, the driverless car connects to a central control point that has a deep understanding of conditions that impact travel. The car informs the control point of its starting point and its destination, its status, (i.e. number of passengers, mission etc.) and in return, the control point sends the car a route. The route is based on factors such as the roads that are available and the other vehicles which are using those roads. The car merges onto a road, travels both at high speed and at a distance of only a few inches from other driverless cars - both front and back and side to side. Because the central control point has a deep level of understanding of the roads and the cars, openings are made for exiting and merging traffic and accidents are eliminated.

Centralized control points, driverless cars and cars traveling within a few inches of other cars may sound far-fetched. However, a subsequent section of this document details how Google applied SDN to its Wide Area Network and is now able to increase the utilization of that network to be 95%. A traditional WAN typically runs at a utilization rate between 60% and 65%. Increasing the utilization to 95% should cut the monthly cost of the WAN in half.

Status of SDN Adoption

The Survey Respondents in both 2013 and 2014 were given a set of alternatives and were asked to indicate the alternatives that described their company's current approach to implementing SDN and were allowed to choose all that applied to their company. The responses of the two survey groups are shown in **Table 1.**

Table 1: SDN Utilization

Approach to Implementing SDN	Responses to The 2014 Survey	Responses to The 2013 Survey
We have not made any analysis of SDN	16%	19%
We will likely analyze SDN sometime in the next year	22%	26%
We are currently actively analyzing the potential value that SDN offers	32%	36%
We expect that within a year that we will be running SDN either in a lab or in a limited trial	22%	19%
We are currently actively analyzing vendors' SDN strategies and offerings	25%	20%
We currently are running SDN either in a lab or in a limited trial	18%	13%
We currently are running SDN somewhere in our production network	11%	6%
We looked at SDN and decided to not do anything with SDN over the next year	10%	5%
We expect that within a year that we will be running SDN somewhere in our production network	18%	10%
Don't know	2%	4%

The data in **Table 1** indicates that while the utilization of SDN in production networks remains limited, it has increased somewhat significantly in the last year. In addition:

The use of SDN in production networks should increase somewhat significantly in the next year.

The SDN 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 SDN and NFV is described later in this document.

Figure 1: ONF's SDN Architecture

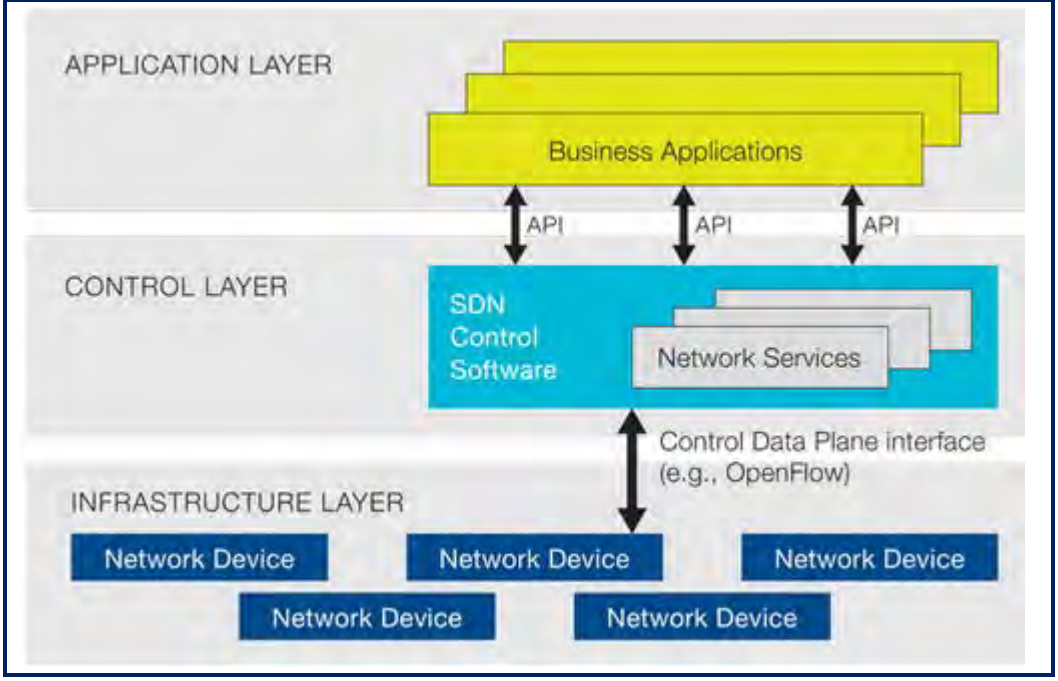

As is discussed below, in some implementations of the architecture depicted in **Figure 1**, the infrastructure layer is just the virtual switch in a hypervisor. This will be referred to as the overlay-based model. In other implementations, the infrastructure layer is a combination of virtual and physical network devices. This will be referred to as either the underlay-based model or the fabric-based model.

The white paper entitled *How to Plan for SDN* discusses criteria to evaluate a vendor's overall SDN solution as well as specific criteria to evaluate a SDN controller and the subtending network devices. That white paper also contains a framework for how network organizations can plan for the adoption of SDN.

The Northbound Interface

The 2013 Guide contains definitions of the key terms and concepts that are embodied in **Figure 1**. One of those concepts is the North Bound Interface (NBI), which is the interface between the control layer and the application layer. When The 2013 Guide was published there were not any standards associated with the NBI and there was an ongoing debate in the industry about the viability of creating such standards. Proponents of standardizing the NBI argued that there were numerous controllers on the market, each with their own NBI and none of which had significant market share. Their argument was that the lack of standardization impeded the development of SDN because without standardization application developers wouldn't be very motivated to develop applications for a controller with small market share knowing that they will likely have to modify their application to work on other controllers. The argument against standardization was that given where the industry was relative to the development of SDN it wasn't possible to really know what should go into the NBI and hence it made no sense to standardize it.

After over a year of discussion, in late 2013 the ONF created the NBI working group and outlined the group's charter in a white paper. As part of their charter, the NBI working group intends to work with one or more open source initiatives to develop working code for the NBIs that the group standardizes. According to Sarwar Raza the chair of the NBI working group, the working group has a good relationship with both the OpenStack and the OpenDaylight initiatives but that when dealing with open source initiatives "there is no magic handshake". Raza elaborated by saying that none of the open source initiatives are going to agree in advance to produce code for NBIs that are under development. He expects that what will happen is that after the standards have been developed the NBI working group will have detailed technical discussions with multiple open source communities and will see if there is a consensus about developing code.

The NBI working group has introduced the need for APIs at different *latitudes*. The idea is that a business application that uses the NBI should not require much detailed information about the underlying network. Hence, applications like this would require a high degree of abstraction. In contrast, network services such as load balancing or firewalls would require far more granular network information from the controller and hence, not need the same level of abstraction. One conclusion to be drawn from this approach is that the NBI working group won't come out with one NBI that works for every type of application. It is also highly likely that there will be further segmentation of NBIs based on industry sector. For example, there may be different NBIs for enterprises than there are for service providers.

Architectural Distinctions between Approaches

Network virtualization isn't a new topic. IT organizations have implemented various forms of network virtualization for years; i.e., VLANs, VPNs, VRF. However, in the context of SDN the phrase *network virtualization* refers to the creation of logical, virtual networks that are decoupled from the underlying network hardware to ensure the network can better integrate with and support increasingly virtual environments.

As previously noted, the predecessor to The Guide was entitled The 2013 Guide to Network Virtualization and SDN. The genesis of that title was that in 2013 there was disagreement in the industry about whether or not SDN and network virtualization were the same thing. Today most of that disagreement has gone away and there is general agreement that network virtualization is a critical SDN application and as described below, there are multiple ways to implement network virtualization.

In addition to having multiple ways of implementing network virtualization, other key architectural distinctions between the varying ways that vendors are implementing SDN include the:

- Role of dedicated hardware;
- Amount of control functionality that is centralized;
- Use of protocols such as OpenFlow.

As indicated above, there is a divergence of opinion relative to the role of dedicated hardware. One example of that divergence of opinion is that some vendors believe it is possible to fully support network virtualization in the data center without using dedicated hardware and some vendors believe that dedicated hardware is needed at least some times. The Survey Respondents were asked to indicate if they believed that with the current technologies and

products it's possible to broadly support network virtualization in the data center without using any dedicated hardware? The *no* responses outnumbered the *yes* responses by almost a 2:1 ratio.

IT organizations are highly skeptical that they can implement network virtualization in the data center without using at least some dedicated hardware.

The Survey Respondents were also asked to indicate the likely role that the OpenFlow protocol will play in their company's implementation of SDN. Their responses are shown in **Table 2**.

Table 2: Likely Use of OpenFlow

Use of OpenFlow	Percentage of Responses
Our implementation of SDN will definitely include OpenFlow	18%
Our implementation of SDN will likely include OpenFlow	24%
Our implementation of SDN might include OpenFlow	24%
Our implementation of SDN will not include OpenFlow	4%
Don't know	29%
Other	2%

One of the conclusions that can be drawn from the data in **Table 2** is that IT organizations have a favorable view of OpenFlow. In addition:

Very few IT organizations have ruled out the use of OpenFlow.

The Overlay and the Underlay Model

As mentioned, there are two primary approaches that vendors are taking to implement the architecture depicted in **Figure 1**. These two approaches are the:

- Overlay-based model;
- Fabric-based or underlay model.

The overlay-based model focuses on the hypervisor and it uses tunneling and encapsulation. Since the overlay-based model focuses on the hypervisor, its use cases tend to be focused on responding to challenges and opportunities that are associated with virtualized servers. A discussion of the pros and cons of the overlay-based model is found in *The Advantages and Disadvantages of the Overlay-Based SDN Model*. A detailed set of criteria that IT organizations can use to evaluate some of the specific characteristics of the overlay-based model is found in *Architectural Criteria to Evaluate Overlay-Based SDN Solutions*.

Whereas the overlay-based model focuses on the hypervisor and uses tunneling and encapsulation, the underlay-based model focuses on a range of virtual and physical network elements and relies on the SDN controller manipulating flow tables in the network elements. In addition, whereas the use cases for the overlay-based model are focused on responding to challenges and opportunities that are associated with virtualized servers, the use cases that are

associated with the underlay-based model are broader in scope; i.e., ease the burden of configuring and provisioning both physical and virtual network elements.

One way that network virtualization can be implemented within an underlay solution is by having virtual networks be defined by policies that map flows to the appropriate virtual network based on the L1-L4 portions of the header. In line with the general philosophy of an underlay-based model, the SDN controller implements these virtual networks by configuring the forwarding tables in OpenFlow-based physical and virtual switches. However, another option is that an underlay solution manipulates the flow tables in OpenFlow-based physical and virtual switches in order to provide a range of functionality other than network virtualization, but that the underlay solution also uses an overlay-based approach to implement network virtualization.

The Survey Respondents were asked to indicate how their company sees the value that the overlay- and the underlay-based models will provide over the next two years. Their responses are shown in **Table 3**.

Table 3: The Perceived Value of the Overlay and Underlay-based Models

Response	Percentage of Respondents
The overlay-based model will provide notably more value	22%
The fabric-based model will provide notably more value	28%
Each model will offer roughly equal value	12%
We don't have an opinion on either model	31%
Other	7%

By a small margin, IT organizations perceive the fabric-based SDN model will provide more value over the next two years than will the overlay model. However, many IT organizations are yet to form an opinion.

Another step in the evolution of SDN is that a year ago the discussion of the overlay-based and underlay-based models was typically phrased as the overlay-based model vs. the underlay-based model. While that is still an interesting discussion, some providers of overlay-based solutions either have already started to ship products or have announced their intention to ship products based on federating their controllers with those of one or more providers of underlay-based solutions; a.k.a., an overlay/underlay solution. A large part of the motivation to deliver federated overlay/underlay solutions is that effective operations management requires that IT organizations have tools that give them clear visibility into the relationships between the virtual networks that are set up by the overlay solution and the physical networks and their component devices that are controlled and managed by the underlay solution. That is required because when performance or availability problems occur, both root cause analysis and impact analysis require bilateral mapping between the physical and virtual infrastructures.

Service Chaining

The phrase *service chaining* refers to the ability to steer virtual machine (VM)-VM traffic flows through a sequence of physical and/or virtual servers that provide network services, such as firewalls, IPS/IDS, DPI, or load balancers. In an underlay-based solution, the controller configures the forwarding plane switches to direct the flows along the desired paths. In an

overlay-based solution, the controller adjust the Forwarding Information Bases (FIBs) of the vSwitches/vRouters to force the traffic through the right sequence of VMs. The next section of The Guide focuses on Network Functions Virtualization (NFV). That section will discuss what the European Telecommunications Standards Institute (ETSI) refers to as VNF forwarding graphs, which are similar in concept to service chains.

The OpenDaylight Consortium

The OpenDaylight Consortium 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 of September 2014 the consortium had 41 members: 9 platinum members, 2 gold members and 30 silver members. Platinum members commit to dues of $500,000 a year for two years and to also provide at least ten developers a year. The financial commitment for Gold and Silver members is determined by a sliding scale based on the company's revenues. Gold members pay annual dues that range between $50,000 and $250,000 and provide at least three developers while Silver members pay annual dues that range between $5,000 and $20,000 and provide at least one developer.

In February 2014 the consortium issued its first software release, called Hydrogen (**Figure 2**). A number of vendors have announced their intention to use Hydrogen as the basis of their SDN controller. A discussion of the functionality that Hydrogen provides can be found at the consortium's Web site.

Figure 2: OpenDaylight SDN Architecture

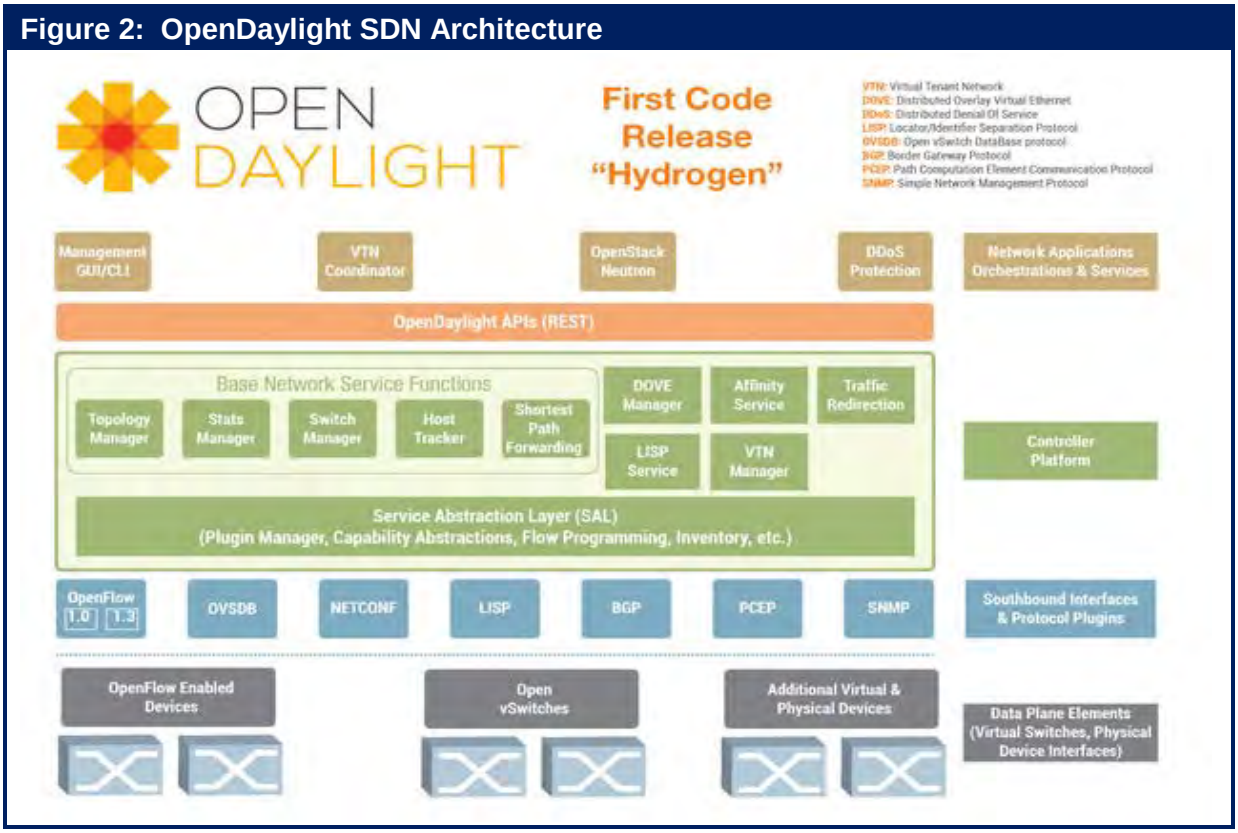

According to Neela Jacques, the Executive Director of OpenDaylight, the consortium's next release of software, code named Helium, will likely be released in late 2014. He stated that some of the new functionality that may be included in Helium includes service chaining, the federation of SDN controllers, additional network virtualization options as well as more L4 – L7 functionality.

The Relationship Between SDN and NFV

Until recently, the conventional wisdom in the IT industry in general, and on the part of the ONF and the ETSI NFV ISG[1] in particular, was that was that SDN and NFV were separate topics and didn't need to be formally coordinated. That conventional wisdom officially changed in March 2014 when the ONF and the ETSI NFV ISG announced the signing of a Memorandum of Understanding (MOU).

As part of the announcing the MOU, the ONF and ETSI stated that "Together the organizations will explore the application of SDN configuration and control protocols as the base for the network infrastructure supporting NFV, and conversely the possibilities that NFV opens for virtualizing the forwarding plane functions." Also as part of the announcement, the ONF released a document entitled the *OpenFlow-enabled SDN and NFV Solution Brief* . The solution brief showcases how operators are combining NFV and SDN to achieve the common goals of both technologies to achieve greater agility of the networks. The brief discusses the network challenges that operators will need to overcome to implement NFV, and it presents use cases that demonstrate how OpenFlow-enabled SDN can meet the need for automated, open, and programmable network connectivity to support NFV.

The Survey Respondents were asked to indicate the relationship that their company sees between SDN and NFV and they were allowed to check all that applied. Their answers are shown in **Table 4**.

Table 4: Perceived Relationship between SDN and NFV

Relationship	Percentage of Respondents
They are totally independent activities	6%
They are complementary activities in that each can proceed without the other but the value of each activity may be enhanced by the other activity.	61%
In at least some instances, NFV requires SDN	25%
In at least some instances, SDN requires NFV	10%
Don't know	16%

Some of the conclusions that can be drawn from the data in **Table 4** are:

The vast majority of IT organizations believe that SDN and NFV are complimentary activities.

A significant percentage of IT organizations believe that in at least some instances NFV requires SDN.

[1] The role that this group plays in the development of NFV is explained in the next chapter of The Guide.

Only a small percentage of IT organizations believe that SDN and NFV are totally independent activities

Advertorial

Cisco ACI: An Application Centric Approach to SDN

IT Trends and the Advent of Software Defined Networking

IT departments and lines of business are looking at cloud automation tools and software-defined networking (SDN) architectures to accelerate application delivery, reduce operating costs, and increase business agility. The success of an IT or cloud automation solution depends largely on the business policies that can be carried out by the infrastructure through the SDN architecture.

The emergence of SDN promised a new era of centrally managed, software-based automation tools that could accelerate network management, optimization, and remediation. Gartner has defined SDN as "a new approach to designing, building and operating networks that focuses on delivering business agility while lowering capital and operational costs." (*Source: "Ending the Confusion About Software-Defined Networking: A Taxonomy",* Gartner, March 2013)

The Cisco Application Centric Infrastructure (ACI) architecture, Cisco's expanded vision of SDN that encompasses the entire data center infrastructure, supports a more business-relevant application policy language than alternative software overlay solutions or traditional SDN designs. What makes the Cisco SDN policy model application-centric? And what are the benefits? First we need a comparison of ACI to traditional SDN designs.

A Comparison of ACI to Traditional SDN Architectures

Although traditional SDN and Cisco ACI have important differences, both have essentially the same architectural components and concepts for policy-based IT infrastructure automation:

- A centralized policy store and infrastructure controller: In SDN and Cisco ACI, this feature is generally known as the controller (Cisco Application Policy Infrastructure Controller [APIC] for Cisco ACI).
- Programmable, or automated, network devices: All infrastructure devices, such as switches, application delivery controllers and firewalls, must be able to respond to and implement policies according to commands from the controller. This feature may involve agents running on the device, APIs in the devices themselves, or management hooks to the devices that are implemented in the controller.
- A controller southbound protocol to communicate with the managed or controlled devices and to communicate policy information: Initially, the OpenFlow protocol was used in SDN architecture, and vendors released OpenFlow-compliant switches. In Cisco ACI, OpFlex is the primary protocol used, although other mechanisms for integrating devices into the Cisco ACI policy model are supported.
- Northbound controller interfaces for integrating higher-level automation solutions on top of the policy and controller framework, including workflow automation tools and analytics: Modern SDN controllers, as does Cisco APIC, include northbound APIs allowing for the integration of OpenStack or other vendor-specific cloud automation tools (e.g., Cisco UCS Director).

What's unique about ACI is that the policy language (the rules that tell your cloud infrastructure what to do) is not modeled on arcane networking concepts like VLAN's and IP addresses, but on application requirements, and especially how application workloads can and can't communicate, and what kind of services they are entitled to. Policies are applied to classes of applications or workloads (e.g., the web tier of an application), also called endpoint groups (EPG), which can be either physical or virtual workloads (or containers).

An application policy will consist of the EPG's that make up the application, and the contracts and services between the EPG's. This is fundamentally all we need to automate the deployment, provisioning and optimization of our application network anywhere, on any cloud resources we want.

The result is an SDN-automated infrastructure that extends beyond just network devices, to include layer 4-7 application services like load balancers, as well as security devices and policies for IPS and firewall components. Because applications are the best reflection of business activity, an application-centric policy is ideal to align IT with business policies, and to automate policies that reflect real business and application requirements.

Figure – Cisco ACI provisions the entire network infrastructure through application polices managed in a centralized SDN controller, the APIC.

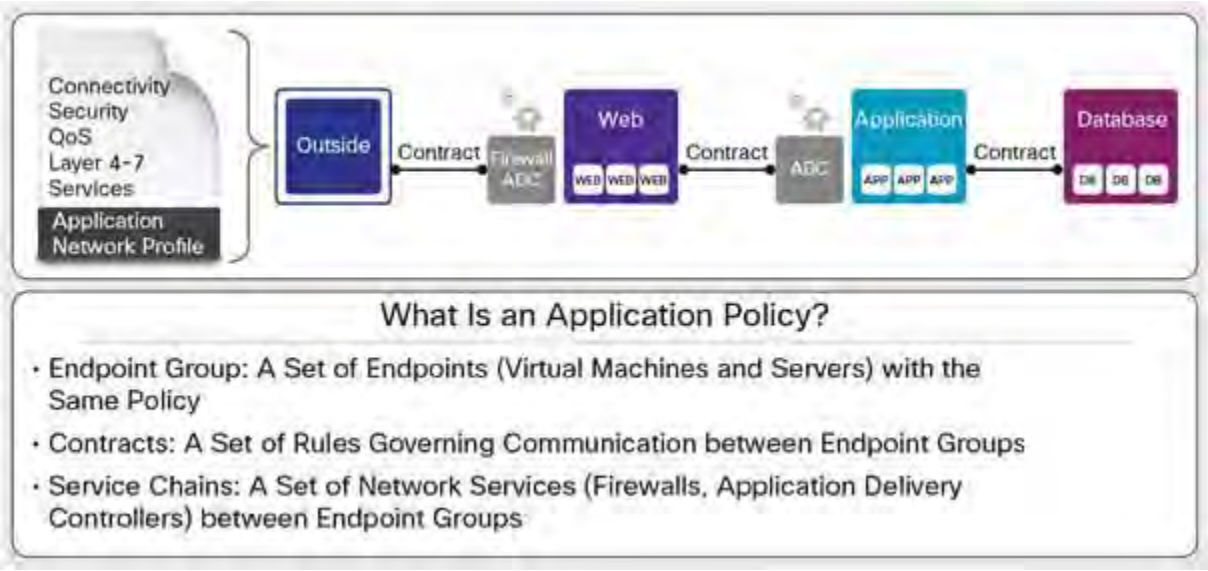

For More Information

For more information, please visit http://cisco.com/go/aci.

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.

Printed in USA C11-728045-00 05/13

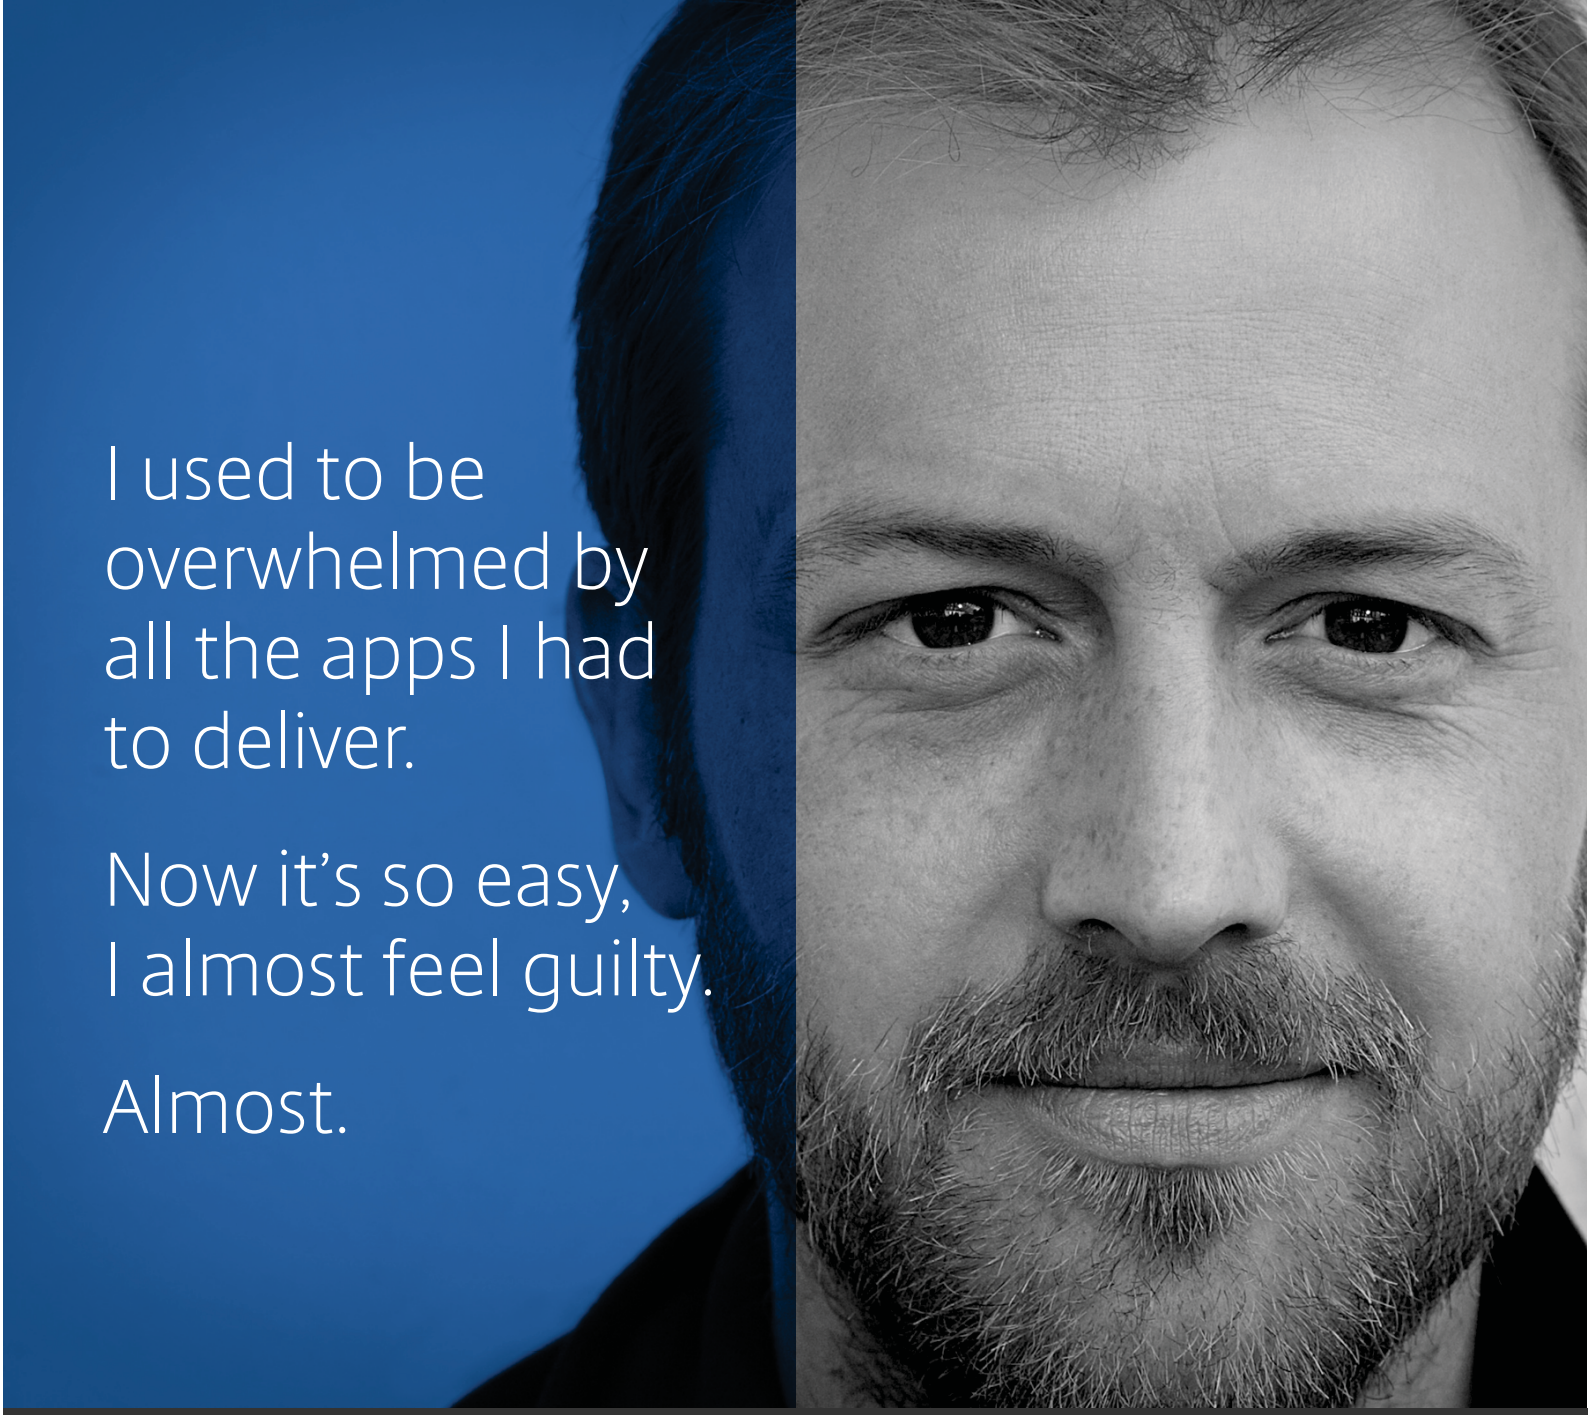

NetScaler with TriScale harnesses the power of software so you can effortlessly customize your app delivery for any business need.

NetScaler with TriScale
SOFTWARE SMART. HARDWARE STRONG.

CITRIX®

www.citrix.com/netscaler

Enterprise SDN and Carrier NFV: Distant Cousins or Twins?

It's not unusual to hear SDN and NFV mentioned together as part of a broad, conceptual discussion about network virtualization. But in practical use, the two represent entirely different worlds—SDN having been born out of the needs of large enterprises principally focused on data center virtualization, and NFV being embraced by telcos and communication service providers for virtualizing service delivery.

One major reason enterprise and carrier technology, including SDN and NFV, are treated as fundamentally different is that the IT goals driving enterprises and carriers are noticeably dissimilar. Add to that the historic differences in vocabulary, infrastructure, and scale between the two camps, and it's not surprising most IT professionals still think of these worlds as wholly unrelated.

But as IT evolves toward virtualization and convergence, the fact is that undeniable similarities have started to emerge. In fact, there are common infrastructural elements striking enough to raise the question: should we think of enterprise SDN and carrier NFV as distant cousins, or are they actually more like twins?

OPERATIONAL NEEDS DRIVE ENTERPRISE SDN ADOPTION

Data centers have been virtualizing server and storage functions using software and hypervisors from VMware, Microsoft Hyper-V, and Red Hat/OpenStack for years now. Virtual machines (VMs) give enterprise data centers the flexibility and agility they need to scale and operate efficiently on a day-to-day basis while also reducing the amount of physical infrastructure required.

Naturally, IT teams have begun applying the same philosophy to networking, seeking the greatest level of virtualization, automation, and programmability possible to simplify their back-end operations.

In many cases, this involves the deployment of virtual switching technology (aka vSwitches) and networking them along with physical switches to create more efficient workflows for applications and workloads.

Right now, there are three common approaches to virtualizing networking infrastructure and and introducing greater levels of programmability and automation.

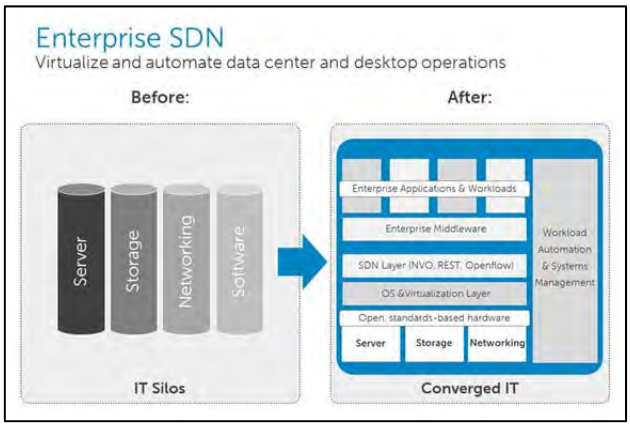

1. Network virtualization overlay (NVO)

NVO stitches the data center's vSwitches together by building tunnels (VXLAN, NV-GRE, etc.) through the physical switch infrastructure, requiring no additional effort at the physical switch level.

2. Controller-based solutions (ex: Openflow)

Controller-based solutions change what takes place in the physical switch by establishing a protocol among the deployed physical switches and a controller. The controller can then be used to program all the switches in any way desired for policy control.

3. Programmable solution (ex: REST)

Other IT teams prefer to use programmatic or scripting languages, such as Puppet or Chef, to interface with their infrastructure and automate operations. Rather than a controller that speaks to multiple devices, they implement a programmatic language to define and implement policies across the infrastructure. .

Each of these approaches has arisen from challenges that are inherent in operating large-scale data centers. Meanwhile, carriers have their own reasons for virtualizing their infrastructure.

SERVICE DELIVERY GOALS DRIVE CARRIER NFV ADOPTION

Traditionally, when a carrier delivers IP services, data packets are sent from customer site or device to a carrier's router or switch, and then daisy-chained through a set of boxes performing additional service-related functions.

Just as it sounds, this process of service creation and delivery has been very physical in nature, involving many pieces of equipment, cables, and

moving parts and requiring similarly large number of staff for rollout and support.

Carriers globally are now turning to virtualization and in particular NFV as a way to simplify and automate service delivery infrastructure, while also introducing greater agility for new service creation and delivery. For carriers, then, the drivers for virtualization are to improve both CAPEX and OPEX structures, making existing service delivery more cost effective, and enabling new, high-margin, services quickly.

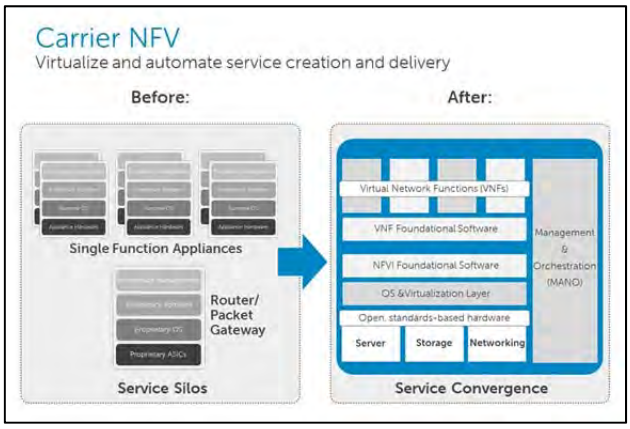

THE LANGUAGE BARRIER

To further compound these differences, enterprise SDN and carrier NFV generally fall under the purview of different executive roles—typically the CIO at enterprises and the CTO for carriers.

There are also fundamental differences in the vocabulary surrounding each. Instead of *workloads*, carriers are concerned with *services*, and instead of *business continuity*, carriers are interested in *carrier-grade* 5-9's and 6-9's *technology*.

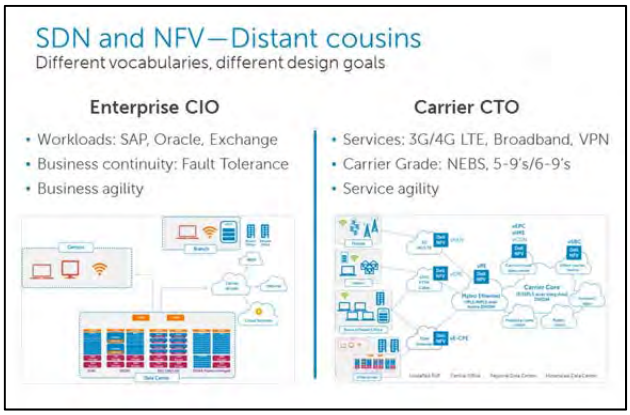

DISTANT COUSINS OR TWINS?

Considering this laundry list of differences, you might wonder how we can propose that enterprise SDN and carrier NFV are actually twins. It's not until you look at their technological DNA that you start to see the remarkable similarities.

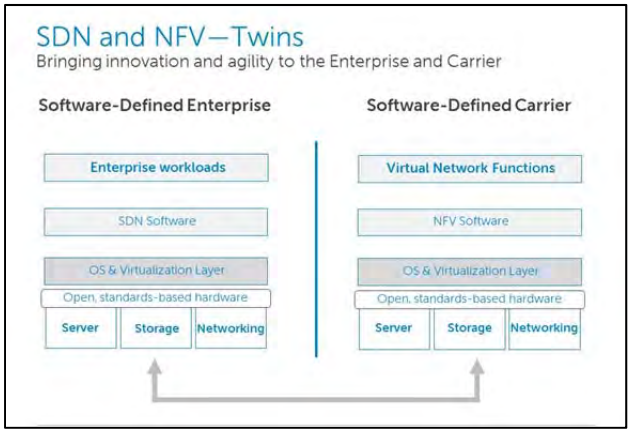

As the above image shows, beneath the disparate business goals and terminology, the infrastructures that support enterprise SDN and carrier NFV are practically identical.

At its core, in both enterprise SDN and carrier NFV, exists x86 server-centric DNA that forms the foundation of the converged infrastructure for compute, storage and networking. Yet, just as twins share the same DNA but can have very different personalities based on environmental factors, enterprise SDN and carrier NFV are really only distinguishable at the application level (e.g. enterprise application vs. carrier VNF)

COMMON TRAITS FOR THE FUTURE

The full implications of this shift in perspective remain to be discovered, but a couple of opportunities immediately arise when we recognize the structural similarities of enterprise SDN and carrier NFV:

1. Carriers who are new to network virtualization can learn best practices from Web 2.0 and large enterprises who have already made significant strides in that area and apply in context..

2. Organizations that operate both production and provisioned infrastructure—enterprise-style for their own operations and carrier-style to provide services—can cross-pollenate, leveraging common technology assets, best-practices, and purchasing power.

While the vocabulary and topologies may never fully converge, the thinking can, having the potential to open new doors for positive collaboration and greater operational efficiencies. Recognizing the common traits behind enterprise SDN and carrier NFV is the first step.

PICA8 OPEN NETWORKING

WWW.PICA8.COM

PicOS Overview

PicOS™ is the first bare metal compatible network operating system that:

- Enables customers to seamlessly and easily integrate conventional networking and SDN.
- Provides extensive support for traditional switching and routing protocols that is extendable by SDN and OpenFlow capacity through Pica8's hardware accelerated Open-vSwitch (OVS).
- Offers a unique, comprehensive and flexible configuration management environment from either a Linux shell, a feature-rich command line interface (CLI) or a comprehensive set of APIs (JSON RPC and OpenFlow).

PicOS runs as an application in user space in an un-modified Linux kernel, thereby leveraging kernel thread protection, and compatible with DevOps tools such as Chef and Puppet that are popular with server and system administrators.

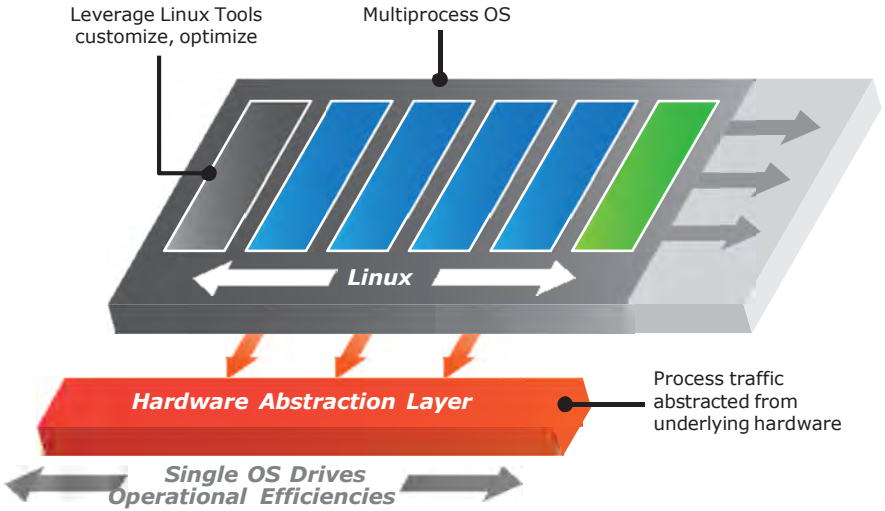

* Only OpenFlow features available in hardware are supported, to ensure optimum performance

PicOS - Three Editions to Leverage

A base configuration starts with the Linux Switching OS package. For additional functionality, select either the Routing or OpenFlow Editions, or the PicOS Bundle depending on your use case.

Features Included	Required PicOS Editions		
	Linux Switching OS	Routing	OpenFlow
• Network operation system using user space standard Debian Linux environment • Leverage vast array of standard Linux tools as a common management and operations framework • Zero Touch Provisioning (ZTP) functionality coupled with ONIE delivers a true bare metal to application environment • Rich Layer-2 protocol stack with MLAG, seamlessly integrating into existing architectures • Full Layer-2 & Layer-3 ACL support • IPv4 & IPv6 Static Routing	✓		
• Rich OSPF and BGP protocol stacks integrating into existing spine / leaf architectures • IPv6 routing protocol support (OSPFv3, MBGP) • Multicast PIM support • NAT (depends on ASIC support) • VXLAN network virtualization (depends on ASIC support)	✓	✓	
• Leading OpenFlow 1.4 support through OVS 2.0 • Deliver true seamless migration to SDN through CrossFlow mode (Layer-2 / Layer-3 and OpenFlow simultaneously) • Leveraging OpenFlow to control MPLS, GRE, NVGRE or VXLAN tunnels, delivering on the promise of open programmability • Support for all major OpenFlow controllers (for example: OpenStack Neutron ML2, OpenDaylight, Ryu)	✓		✓
PICOS Bundle	✓	✓	✓

SDN Use Cases

Drivers and Inhibitors

The Survey Respondents were shown a number of challenges and 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. Their responses are shown in **Table 5.**

Table 5: Opportunities & Challenges that SDN Can Address

Challenge or Opportunity	Percentage
Better utilize network resources	55%
Perform traffic engineering with an end-to-end view of the network	54%
Ease the administrative burden of configuration and provisioning	53%
Support the dynamic movement, replication and allocation of virtual resources	52%
More easily scale network functionality	45%
Enable applications to dynamically request services from the network	45%
Have network functionality evolve more rapidly based on a software development lifecycle	41%
Reduce OPEX	40%
Implement more effective security functionality	35%
More easily implement QoS	33%
Reduce CAPEX	29%
Reduce complexity	24%
Other	5%

One observation that can be drawn from the data in **Table 5** is that IT organizations are optimistic that SDN can help them respond to a wide range of opportunities and challenges. However:

***Relatively few IT organizations believe that SDN will help them reduce CAPEX* or reduce complexity.**

The Survey Respondents were also shown a set of 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. Their responses are shown in **Table 6.**

Table 6: Inhibitors to the Adoption of SDN	
Impediment	**Percentage**
The immaturity of the current products	29%
Concerns about how we would integrate SDN into the rest of our infrastructure	23%
The immaturity of the enabling technologies	23%
The lack of a compelling business case	21%
The confusion and lack of definition in terms of vendors strategies	16%
Other technology and/or business priorities	14%
Concerns about how we would manage SDN	13%
Possible security vulnerabilities	12%
The lack of a critical mass of organizations that have deployed SDN	9%
No inhibitors to implementing SDN	7%
Concerns that the technology will not scale to support enterprise sized networks	6%
Other	5%

Some of the inhibitors to SDN adoption, such as the immaturity of current products and the immaturity of enabling technologies, will naturally dissipate over time. However some on the key inhibitors won't just naturally dissipate over time. These inhibitors need to be aggressively addressed by vendors and network organizations.

Two of the major inhibitors to SDN adoption are concerns about how to integrate SDN into the rest of the infrastructure and the lack of a compelling business case.

SDN Deployment Plans

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 7**.

Table 7: Focus of SDN Deployment	
Focus of SDN Deployment	**Percentage**
Data Center	64%
WAN	26%
Branch and/or Campus	25%
We are unlikely to implement SDN within the next two years	12%
Don't know/NA	10%
We are likely to implement a service from a WAN service provider that is based on SDN	8%
Other	6%

One observation that can be made from the data in **Table 7** is:

Over the next two years, the primary focus of SDN deployment is likely to be in the data center. However, there is considerable interest in deploying SDN in the WAN as well as in branch and campus networks.

The Survey Respondents were also asked to indicate how broadly they expected their campus, WAN and data centers networks would be based on SDN three years from now. Their responses are summarized in **Table 8**.

Table 8: Planned SDN Deployment

	Campus Networks	WAN	Data Center Networks
Exclusively based on SDN	1%	2%	6%
Mostly SDN	10%	6%	20%
Hybrid, with SDN and traditional coexisting about equally	34%	36%	50%
Mostly traditional	29%	31%	10%
Exclusively traditional	13%	13%	4%
Don't know	12%	12%	10%

Given the relatively low penetration of SDN currently, the data in **Table 8** shows that:

Network organizations are very optimistic that over the next three years that there will be a significant increase in SDN deployment.

Network organizations believe that three years from now that SDN deployment in data centers will be highly pervasive and that there will also be significant SDN deployment both in the WAN and in campus networks.

The sections below describe possible SDN use cases in the data center, the WAN and the campus. In some instances the use cases are generic and in some instances the use cases reflect actual implementations. In many cases the placement of the use case is somewhat arbitrary. For example, most of the use cases that are included in the data center section could also be included in the campus networks section.

Data Center

Virtual Machine Migration

One of the advantages of server virtualization is that it enables moving VMs between physical servers. However, when a VM is moved between servers, the VM needs to be on the same VLAN after it was moved as it was on prior to the migration. Extending VLANs across a data center in order to support workload mobility adds to the operational cost and complexity and it adds time to the process because it requires that each switch in the end-to-end path be manually reconfigured.

Network virtualization resolves that challenge because with network virtualization when a VM changes location, even to a new subnet in the physical network, the switches at the edge of the overlay automatically update their mapping tables to reflect the new physical location of the VM. One of the advantages of network virtualization is that since the necessary changes are performed only at the network edge, nothing has to be done to the remainder of the network.

Service Chaining

In a traditional data center implementing L4 – L7 services such as firewalls and WAN optimization is cumbersome and time consuming as it requires acquiring the requisite network appliances and cabling them together in the correct order. Since each appliance has its own unique interface, configuring these appliances is a time consuming, error-prone task.

SDN overcomes the challenges of implementing L4 – L7 services by implementing two closely related techniques: service insertion and service chaining. The phrase *service insertion* refers to the ability to dynamically steer traffic flows to a physical or virtual server that provides one of the L4 – L7 services that were listed above. The phrase *service chaining* refers to the ability to dynamically steer traffic flows through a sequence of physical or virtual servers that provide the same type of L4 – L7 services.

Security 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, an OpenFlow-enabled 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 OpenFlow controller can match suspicious flows to databases of malware signatures or divert DDoS attacks.

Load Balancer Services

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 (IU) has developed an OpenFlow-based, load-balancing application called FlowScale. According to the University, “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. FlowScale is currently being deployed as part of the Intrusion Detection Systems operated by the Indiana University Information Security Office.

New Software Defined Cloud model

A national Cloud Leader is creating a new architecture that will allow all network and value-added services to be software defined, based on SDN. They are using OpenStack both for overall orchestration and also for its end user-friendly Horizon dashboard. While customers

interact with the system via the dashboard, their administrators leverage REST APIs to interact both with OpenStack and with the SDN. The SDN provides a virtual network overlay for a consistent, unified fabric over the entire network and all datacenters (planned expansion is 10 datacenters by 2018).

This use of SDN helps the Cloud Leader change how data center services are offered. New capabilities include:

- Full Datacenter Capabilities: Most public clouds offer compute and storage but do not systematically address networking. Their approach provides a complete datacenter approach that spans compute, network, and storage.
- Full UI-driven Self-service: Customers can control every aspect of their virtualized environment using their user interface. This capability both increases customer control and enables the cloud leader to handle huge volumes of customers and VMs – projected to be 1 million VMs at the end of the first year of operation.
- Full Network Programmability: SDN provides a coherent cloud network fabric that enables programmability from the datacenter endpoint all the way through the network. The fabric enables a number of new capabilities including consistent network service independent of underlying hardware, full workload portability among datacenters, and full programmability for future services.
- High Security within the Datacenter: Legacy security approaches focus on external threats rather than threats within the datacenter. The SDN's built-in security, including a default "Zero Trust" model, operates at the virtual machine level. These capabilities provide security and isolation within the rack, within each customer's operations, and within the datacenter.

New Distributed Cloud Hosting model

A telecommunications service provider (TSP) in EMEA has created a virtual Platform Optimized Design (vPOD) architecture that provides Cloud efficiencies along with the flexibility of offering either shared or dedicated resources distributed among datacenters. SDN provides the interconnection within and among all vPODs and among all datacenters. Having a cohesive, unified cloud across datacenters enables consistent performance critical for SLAs, robust disaster recovery, and other value-added services.

New capabilities include:

- Precision SLAs for each Customer: Adding precise network controls to server virtualization and OpenStack orchestration, the TSP's key customer demand of precise, end-to-end SLAs can be reliably delivered even on shared vPODs.
- Consistent Performance Across Datacenters: Similar to server virtualization, network virtualization provides consistent and predictable performance that is independent of each datacenter's build-out, hardware configurations, and network architectures. This capability enables a range of new customer-centric capabilities such as load balancing workloads across datacenters.
- Fluid Disaster Recovery: Having consistent performance independent of datacenter build-out changes how disaster recovery is performed. Instead of having idle resources standing by in a dedicated datacenter, a customers' implementation can be stretch-clustered across datacenters for truly fluid disaster recovery. In this fashion, loss of one or even multiple datacenters can be accommodated without disruption to operations.
- Effortless Datacenter Scalability: With this architecture, the TSP can scale-out to accommodate the needs of each customer just by adding vPODs or by adding racks to

a dedicated vPOD. They also can easily scale out to 100 times their initial rack count and up to millions of managed endpoints without having to change the architecture or the configuration.

- Fast and Non-disruptive Provisioning: Outside of physical racking and cabling, new vPODs can be added in an automated and non-disruptive manner – the entire installation or de-installation process takes only a few hours. New servers can be allocated to a vPOD nearly instantaneously via automation.

WAN

The Google G-Scale WAN

As is discussed in *An Overview of OpenFlow V1.3*, 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. 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.

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%.

Campus

Below are some popular use cases associated with deploying SDN in branch and campus networks.

Dynamic QoS & Traffic Engineering

The hop-by-hop routing and queuing techniques currently used in branch and campus networks yield a best effort network that results in poor quality for applications such as unified communications (UC). For the sake of example, consider the case of two users, User A and User B, of a popular UC application: Microsoft Lync. When User A asks Lync to make a call to User B, the Lync call controller converts User B's contact information to an IP address. The Lync call controller sends this IP address to the Lync client running on User A's laptop. A call is then started between the two users, but there is nothing in the call setup to indicate that the traffic for this call should have higher priority than other traffic.

In an SDN environment, as the Lync call controller is sending the IP address to the Lync client running on User A's laptop, the Lync controller can be configured to also send it to an SDN application, whose function is to communicate with an SDN controller and have the priority set

to specified values for specific IP pairs in a network. A Lync call, for instance, could be set to a high priority. The SDN application communicates to the SDN controller that the priority level for traffic between a specific pair of IP addresses needs to be set to high and that this traffic should run over non-congested links. The SDN controller takes this information and determines the optimal path for the packets to flow through the network from User A to User B. This flow matching information, along with the required actions, are pushed out to each of the OpenFlow-enabled switches.

Unified Wired and Wireless Networks

Typically, wireless networks have been built as overlays to a wired network. As a result, in the vast majority of cases the wired and wireless networks in a campus operate as separate entities. This situation has a negative impact on users because it means that users will likely have different experiences based on whether they are using a wired or a wireless access device. This situation also negatively impacts IT organizations because maintenance and troubleshooting are unduly complex due to the fact there are two separate management systems, two separate sets of policies and two separate authentication processes.

One of the advantages of integrating the wired and wireless networks in a campus is that it results in a single-pane-of-glass management of the unified wired and wireless network. Using SDN technologies for this integration will make network provisioning more dynamic. For example, as wireless devices roam from AP (access point) to AP the policy associated with the user moves as well. Another advantage of the SDN architecture and related technologies is that they enable enforcing policy at a very granular level. This means, for example, that it is possible to set quality of service policies on a per user or per device basis. Another example of a granular policy option that is enabled by SDN is that if the IT organization trusts traffic from a specific SSID, it can decide to let that traffic bypass the firewall and hence not consume firewall resources needlessly.

QoS Management for Microsoft Lync across wired and wireless networks

This use case can be viewed as a combination of the preceding two use cases. As previously noted, enterprises are rapidly adopting Microsoft's Lync as their unified communications solution of choice, but until recently a unifying Lync wired and wireless solution wasn't available in the market. That is important because wireless has become the edge of the network and mobile users have a growing dependency on wireless services for performing critical job tasks. This situation creates a challenge relative to how wireless users can effectively and reliably access Lync services.
Recently an OpenFlow-based application that bridges wired and wireless networks to ensure a user the highest quality of experience with Microsoft's Lync has entered the market. The solution can detect quality of service issues, identify resolutions and prioritize traffic across any OpenFlow-enabled network. The solution also enables the wireless and wired network to dynamically change in response to application traffic requirements.

Personal Bonjour

When Apple announced Bonjour, its zero-configuration application, it filled a void in the market. Users could simply access a network attached television or printer, as long as the device was on the same sub-net. Businesses quickly saw value in this class of application and commercial network centric solutions opened Bonjour up to more expansive networks. However, this

created a management challenge relative to user and device associations with larger populations of network users and devices.
Recently an OpenFlow-based application has been introduced to the market that implements the highest granularity policy management for Bonjour service access available. This application has functionality that enables IT organizations to ensure that individual users may be allowed to only access selected devices, in selected locations, at selected times of day. One way that this can be used is that a dormitory full of students, each with their own printer or TV, can be isolated from all other users without the network congestion often encountered with a standard Bonjour implementation.

Role Based Access

It is often useful to control what users can and cannot do on a network based on the role they play within the organization. One of the strengths of the SDN architecture and the OpenFlow protocol is that they offer a hardware- and software-independent abstraction model to access and manipulate resources. One way that the abstraction model can be leveraged to implement role-based resource allocation is by leveraging the authentication functionality that exists between the user and the NAC (Network Access Control) application in such a way that when the authentication process is complete, a message is sent to a role-based resource allocation SDN application. The message contains the MAC address of the user, the port of entry in the network, and the role of the user. The application then finds the user in a previously configured capabilities list. This list contains information such as which devices and other users this new user can communicate with; which VLAN the user should be assigned to; how much bandwidth the user can have assigned to its traffic; and what IP addresses are off limits. These capabilities are converted to a network resource message that is sent to the SDN controller. The SDN controller then communicates with the appropriate network device and configures the OpenFlow tables on that device to ensure the appropriate priority setting for the user's traffic, the appropriate bandwidth as well as instructions to drop flows to restricted addresses.

AUTOMATE YOUR CLOUD WITH aCLOUD SERVICES ARCHITECTURE

Integrate dynamic services into your Cloud Data Center

www.a10networks.com

Delivering the promise of SDN across wired/wireless networks

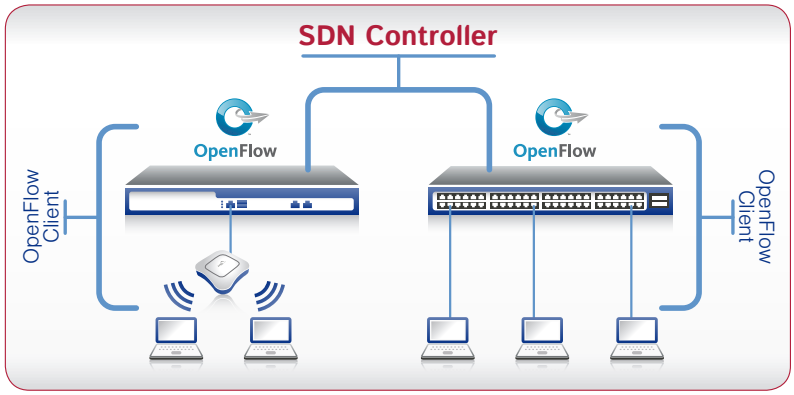

As the enterprise network edge transitions to an all wireless network, software-defined networking (SDN) and OpenFlow are emerging as a way to bring new levels of agility to organizations beyond the data center where SDN first gained traction.

The rapid acceptance of SDN and this new approach to design, build and manage data centers addresses the top challenges experienced by organization related to networks: namely, too many manual processes, and difficulties changing configurations. SDN tackles these challenges in the data center, but SDN can equally address the same issues for the enterprise campus. Without bringing SDN to the edge of the network, its true promise is lost.

Meru is leading the way being the first wireless vendor to receive a Certificate of Conformance through the ONF OpenFlow™ Conformance Testing Program within our wireless LAN controllers to enable third-party control all the way down to the access point. This provides customers with confidence in the products that they adopt will provide multi-vendor support.

Meru is also collaborating with IT giants such as NEC to enable seamless interoperability between the NEC ProgrammableFlow® Networking Suite and Meru 802.11ac intelligent Wi-Fi solutions. NEC and Meru are the world's first vendors to receive OpenFlow Conformance Certification respectively as a wired and wireless vendor - a natural pairing.

Meru has introduced Meru Center, a network application management platform, unifying network applications under a single platform and permits easy activation of pre-installed network tools.

With Meru Center, new SDN applications are delivered via the Meru App Store. This library function hosts a growing set of qualified applications that may be selected and installed on a user's network. Initial Meru SDN applications available will include:

Meru Collaborator
An SDN application that integrates with Microsoft's Lync unified communication solution with the ability to detect QoS (quality of service) issues on a heterogeneous wired/wireless network, deliver prescriptive resolution options and prioritize traffic across multi-vendor wired and wireless networks.

Meru Personal Bonjour
An application that minimizes Bonjour broadcast storms of Apple related devices across unified networks and advertises services only to the correct users according to established policies.

Making SDN a Reality for Wi-Fi

The promise of SDN is that networks will no longer be closed, proprietary, and difficult to manage. Meru is taking a leadership position in the emerging wireless market for SDN, and is committed to delivering the most robust SDN Wi-Fi solution in the market while providing a best-of-breed wireless solution.

With innovative solutions from Meru and a robust SDN ecosystem, organizations can meet the unprecedented demand for Wi-Fi with ease.

Click for more information

Meru delivers an all-wireless network that fully supports the enterprise, delivering a consistent, interactive experience for all users. No matter what applications they are running. No matter how many other users are on the network. For more information, visit **www.merunetworks.com** or email your questions to: meruinfo@merunetworks.com.

Meru Networks®
Corporate Headquarters
894 Ross Drive
Sunnyvale, CA 94089
T +1 (408) 215-5300
F +1 (408) 215-5301
E meruinfo@merunetworks.com

Radware's Software Defined Networking Solutions: Enable Network Wide Services via SDN Applications and NFV

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. In addition, API to various orchestration systems enables to improve the overall control and automation of network services.

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 delivers a security control plane and operates in traditional network environments while enabling to migrate to customer's future, SDN-based networks.

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.**

- **Centralized security control plane including control part of Radware's Attack Mitigation Network (AMN)**

SDN & NFV for a Scalable Application Delivery Network

The Network Functions Virtualization (NFV) initiative was formed in order to enable the standardization of network equipment by leveraging commercially off-the-shelf (COTS) hardware and running advanced network function software on them. Radware is proudly introducing Alteon VA for NFV – the industry's first and only ADC designed from the ground up to run in NFV environments. Targeted mainly at carriers but also at high-end online businesses, Alteon NFV provides unique value proposition including CAPEX/OPEX reduction, eliminate "vendor lock", high performance, high-end scalability and greater network service agility.

Radware's ElasticScale is an SDN application that wraps existing network service virtual appliances, including Alteon NFV, 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 help providers adopt network function virtualization paradigms.

ElasticScale offers network operators the following key features and benefits:

- Ultra scalable traffic steering solution (80Gbps-1Tbps and beyond)
- 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

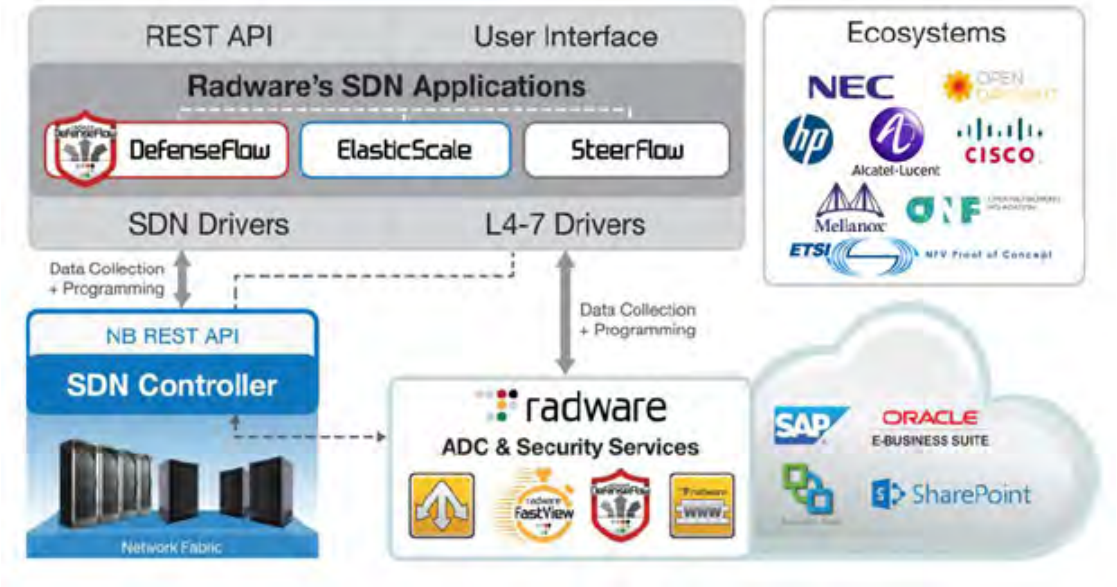

Partnering for Success: Our SDN Ecosystem

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: Cisco Application Centric Infrastructure (ACI), HP Virtual Application Networks, NEC, Mellanox, Alcatel Lucent, ETSI, 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.

The Operational Implications

One of the operational implications of adopting SDN is the movement to a DevOps operational model.

A detailed discussion of DevOps is contained in the subsequent chapter of The Guide.

Security

SDN creates security opportunities and security challenges.

The fact that SDN poses both security opportunities and security challenges was demonstrated by **Table 5** and **Table 6**. **Table 5** shows that 35% of network organizations believe that SDN will enable them to implement more effective security functionality. **Table 6** shows that 12% of network organizations believe that concerns about how possible security vulnerabilities is a significant inhibitor to SDN deployment.

Two examples of how SDN can enhance security were already discussed. In one of those examples, security services were implemented based on OpenFlow-based access switches filtering packets as they enter the network. In the second example, role based access is implemented by deploying a role-based resource allocation application that leverages the control information and capability of the SDN controller. Other security related use cases include leveraging the control information and capability of the SDN controller to provide DDoS protection.

Some of the security challenges related to SDN are described in *SDN Security Considerations in the Data Center*. As pointed out in that document:

- The centralized controller emerges as a potential single point of attack and failure that must be protected from threats.
- The southbound interface between the controller and underlying networking devices (that is, OpenFlow), is vulnerable to threats that could degrade the availability, performance, and integrity of the network.
- The underlying network infrastructure must be capable of enduring occasional periods where the SDN controller is unavailable, yet ensure that any new flows will be synchronized once the devices resume communications with the controller.

Other security-related considerations include that IT organizations should:

1. Implement measures to deal with possible control flow saturation (controller DDOS) attacks;
2. Harden the SDN controller's operating system to ensure availability of the controller function;
3. Implement effective authentication and authorization procedures that govern operator access to the controller.

Chapter 2 of The 2013 Guide to Network Virtualization and SDN contains a set of 5 key questions that network organizations can ask vendors about the security of their SDN solutions.

Cloud Orchestration

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

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

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

Figure 3 shows a block diagram of the OpenStack system, including the OpenStack modules that are used to control resource pools in the data center, including Horizon and Neutron.

Figure 3: OpenStack

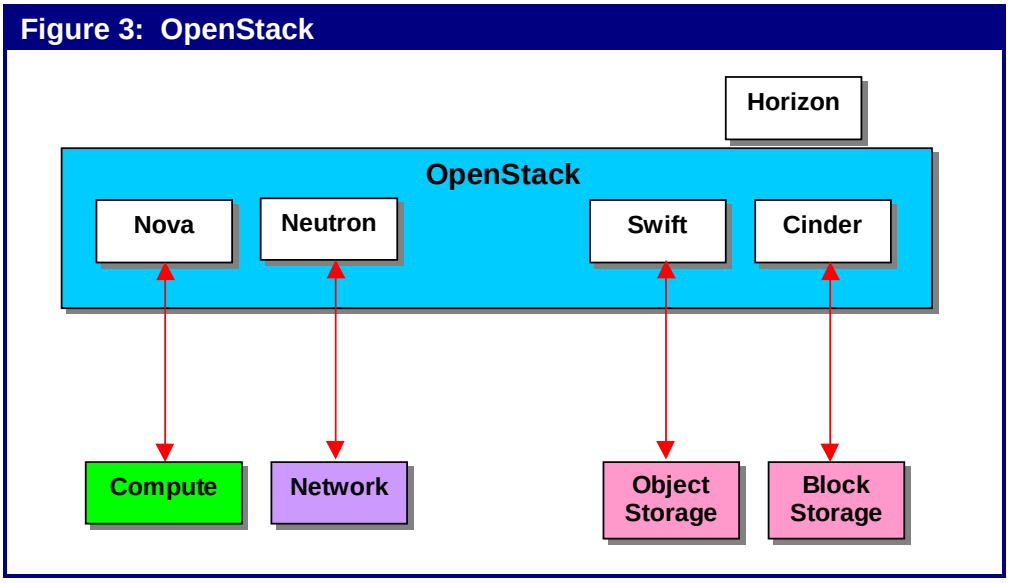

Horizon is the OpenStack Dashboard that provides administrators and users a graphical interface to access, provision and automate cloud-based resources. The dashboard is one of several ways users can interact with OpenStack resources. Developers can automate access or build tools to manage resources using the native OpenStack API or the EC2 compatibility API. The dashboard also provides a self-service portal for users to provision their own resources within set limits.

Neutron (formerly called Quantum) allows users to create their own networks, provide connectivity for servers and devices, and control traffic. With appropriate Neutron plug-ins, administrators can take advantage of various SDN solutions to allow for multi-tenancy and scalability. A number of drivers/plugins are included with the OpenStack source code. OpenStack networking also has an extension framework allowing additional network services,

such as intrusion detection systems (IDS), load balancing, firewalls and virtual private networks (VPN) to be deployed and managed. One example of the extension service is the Load Balancer as a Service (LBaaS) driver for Neutron available starting with the October 2013 Havana release. The driver enables ADC vendors to offer simple LBaaS plugins for Neutron, allowing their ADCs to be directly provisioned by OpenStack. Vendor-specific driver plug-ins that are contributed to the project are included in the OpenStack source code.

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

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

Figure 4 provides a high level depiction of how an orchestrator (OpenStack) and an overlay-based SDN controller might interact to place a VM into service within a VN.

Figure 4: VM Creation Workflow with OpenStack

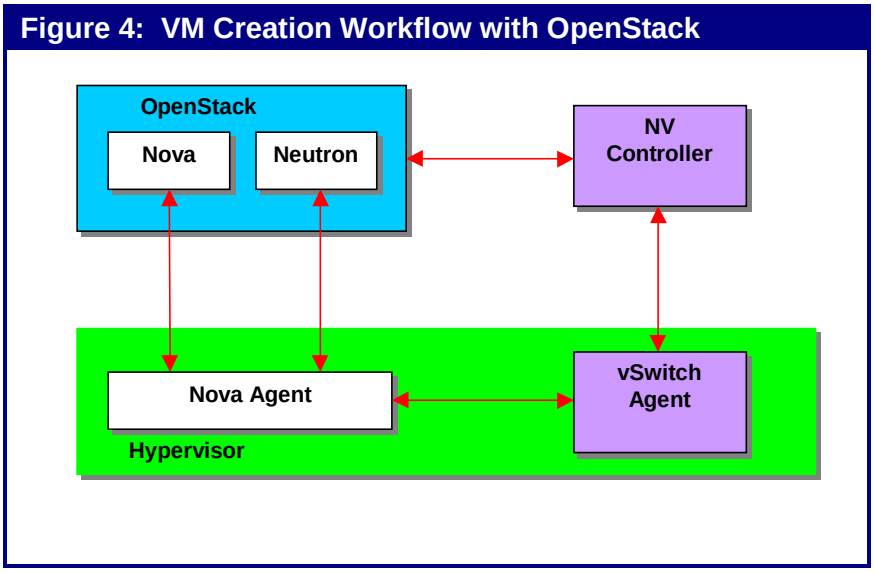

The **Nova** compute module in OpenStack instructs the Nova Agent in the hypervisor to create the VM. The Nova agent communicates with the Neutron module in OpenStack to learn the network attributes of the VM. The Nova agent then informs the vSwitch agent to configure the virtual network for the VM and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

With the April 2014 Icehouse release of OpenStack the **Heat** Orchestration Service has been added. Heat is a template-driven engine that allows application developers to describe and automate the deployment of infrastructure through both an OpenStack-native REST API and a CloudFormation-compatible Query API. The flexible template language can specify compute, storage, and networking configurations to automate the full provisioning of infrastructure as well as services and applications. Through integration with the **Celiometer** Telemetry service, the Heat engine can also perform auto-scaling of certain infrastructure elements. Celiometer aggregates usage and performance data across the services deployed in an OpenStack cloud. This capability provides visibility and insight into the usage of the cloud across multiple data points and allows cloud operators to view service level metrics globally or by individual deployed resources. Usage data can be used for billing and charge back purposes.

The Survey Respondents were asked to indicate the approach that their company is taking relative to orchestration. Their responses are shown in **Table 9**.

Table 9: Approaches to Orchestration

Approach	Percentage of Respondents
We have a well thought out strategy and we have begun to execute against that strategy	16%
We have a well thought out strategy but we have not yet begun to execute against that strategy	6%
We are in the process of developing a strategy and are optimistic that it will come together relatively quickly	21%
We are in the process of developing a strategy but have some concerns that the existing solutions are immature	30%
Don't know/NA	22%
Other	5%

The vast majority of IT organizations don't have a well thought out strategy for how they will implement orchestration.

Management

SDN creates management opportunities and security challenges.

The fact that SDN poses both management opportunities and management challenges was demonstrated by **Table 5** and **Table 6**. **Table 5** shows that 53% of network organizations believe that SDN will ease the administrative burden of management tasks such as configuration and provisioning. **Table 6** shows that 13% of network organizations believe that concerns about how to manage SDN is a significant inhibitor to SDN deployment.

An architectural view of the key management challenges at each tier of the SDN architecture is depicted in **Figure 5** which was published in the ONF document entitled *SDN Architecture Overview*. One of the conclusions that can be drawn from **Figure 5** is that:

> ***In SDN environments the challenges associated with end-to-end service performance management are more demanding than they are in traditional network environments.***

This follows because in a SDN environment there is a need to monitor additional components, such as SDN controllers, in an environment that is a combination of physical and virtual resources and which is changing dynamically. From a service performance management perspective, the SDN controller can be viewed as a *service enabler* that needs to be instrumented and monitored just as any other application server. Whether it is OpenFlow or some other protocol that enables communications between the SDN controller and the network elements that protocol needs to be monitored the same way as any other protocol. In similar fashion, the combination of virtual and physical network elements need to be instrumented end-to-end and monitored across the entire infrastructure.

Figure 5: SDN Management Challenges

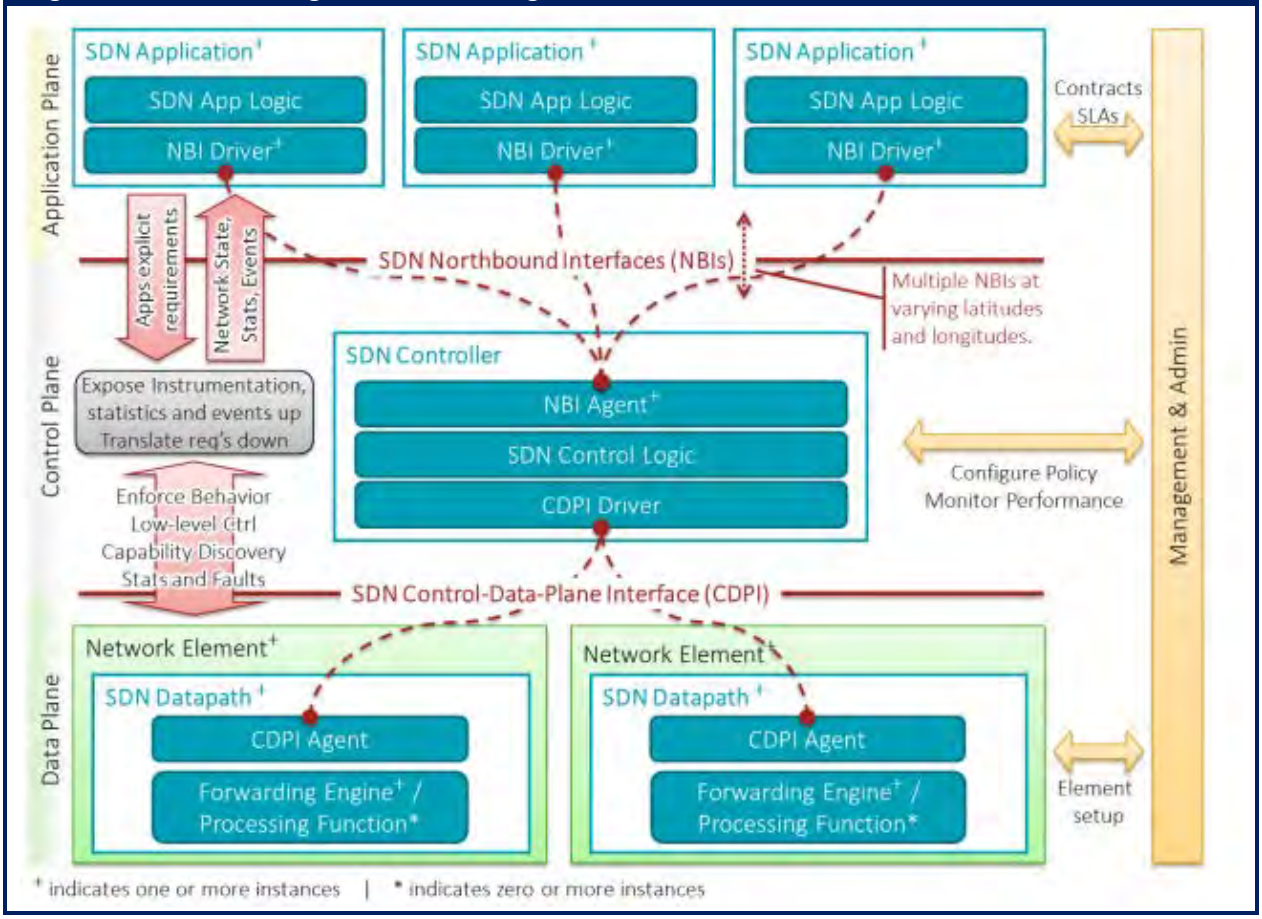

At the bottom of **Figure 5**, the data plane is comprised of network elements, whose *SDN Datapath*s expose their capabilities through the *Control-Data-Plane Interface (CDPI) Agent*. At the top of **Figure 5**, *SDN Applications* communicate their requirements via *NBI Drivers*. In the middle of the figure, the *SDN Controller* translates these requirements and exerts low-level control over the SDN Datapaths, while providing relevant information up to the SDN Applications.

One of the management challenges that applies across multiple tiers of the SDN architecture is the requirement to manage the messaging that goes between tiers; e.g., between the application tier and the control tier as well as between the control tier and the infrastructure tier. Another challenge that goes across tiers is the requirement to assign the SDN Datapaths to their SDN Controller and to configure policies that define the scope of control given to the SDN Controller or SDN Application.

At the infrastructure tier, one of the primary challenges is to perform element management potentially of both virtual and physical network elements. One of the management challenges at the control layer results from the fact that 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.

As described in the preceding discussion of the North Bound Interface (NBI), one of the management challenges that occurs at the application tier is that based on the type of application (e.g., business application vs. a firewall) the service or application needs varying levels of visibility into the underlying network. Another set of management challenges that occurs at the application layer stem from the requirement to ensure acceptable performance. As described below, one thing this means is that network management organizations must have visibility into the SLA requirements of the application so that resources can be dynamically allocated to meet those requirements.

Looking at network virtualization as an application of SDN, another one of the performance management challenges stems from the fact that one of the primary benefits of overlay-based SDN solutions is the ability to support multiple virtual networks that run on top of a physical network. As previously mentioned, in order to perform management functions such as root cause analysis and impact analysis, network management organizations need the ability to see the bilateral mapping between the virtual networks and the physical network that supports them.

While understanding the mapping between the virtual networks and the physical infrastructure is necessary, it is not sufficient. For example, with the virtualization of L4 – L7 functions, software running on VMs can readily be moved among physical servers or replicated to run on newly created VMs in order to dynamically maintain availability, expand/shrink capacity, or balance the load across physical resources. Many of these changes in the infrastructure can be automated and programmatically activated to conform to configured policies under specific sets of circumstances. For example, consider the traffic of an important IP application flow that has a medium priority class. If congestion in the network results in excessive packet loss, it may be necessary to change the traffic classification to be high in order to continue to meet an established SLA.

SDN holds the potential to enable IT organizations to dynamically change the environment in order to meet SLAs.

However, due to the mobility of VMs or the need to change QoS settings, topology changes can occur in a matter of seconds rather than the days or weeks required for changing software/hardware relationships in traditional networks. In order to accommodate and leverage the virtualization technologies, network management organizations need tools that enable them to be able to dynamically discover, procure, allocate and reconfigure resources. In addition:

Network management organizations need to be able to perform a two-way mapping between an application or service and all of the virtual services that support it and they must also be able to perform a two-way mapping between the virtual services that support a given service or application and the physical infrastructure that supports them.

Given the challenges described above as well as the requirement to integrate the traditional legacy environment with the emerging software-centric environment:

Applications and services need to be instrumented end-to-end.

The physical and virtual environments should be instrumented independently and network management organizations should have the ability to contextually correlate and consolidate the two management datasets into one consistent and cohesive dataset which offers operational insight into the end-to-end service delivery.

Chapter 2 of The 2013 Guide to Network Virtualization and SDN contains a set of 5 key questions that network organizations can ask vendors about the management of their SDN solutions.

Organizational Impact

SDN can be viewed as being a part of a broader movement to implement all IT functionality in software, referred to as Software Defined Everything (SDE). The primary drivers of the SDE movement are the need to support a more agile IT operational model as well as increasingly more agile business processes.

As described in *The Changing Role of the IT & Network Professional,* because of the growing adoption of an SDE approach many organizations are implementing DevOps. DevOps is described in the next chapter of this e-book. As is also described in *The Changing Role of the IT & Network Professional* the adoption of an SDE approach is causing the role of network and IT infrastructure professionals to change. Some of the key characteristics of the emerging roles are:

- **An increased knowledge of other IT disciplines**

  In a recent blog, GE Capital's CTO Eric Reed explained the need for all IT professionals to expand their area of expertise. According to Reed, "Our experience [GE Capital's] on this journey to date has been that the small, self-directed teams required in a DevOps world require an amalgamation of skills spanning everything from IT security to database design and application architecture, plus everything in between. While each individual on the team has a particular strength (say, application design and coding), each one also needs to have working knowledge in other areas (maybe UX or network design)."

- **More focus on setting policy**

  Emerging technologies and architectures (e.g., Software Defined Networking, Network Functions Virtualization) enable IT organizations to implement a policy driven infrastructure in a more dynamic and granular fashion than was previously possible. It will take some time to adjust to these new capabilities, but the vast majority of IT organizations will adjust and will place more emphasis on setting policy.

- **More knowledge of the business**

  The need for more knowledge of the business is driven in part by the need for IT and network professionals to implement a policy driven infrastructure that is based on the specific requirements of the business. In addition, the ability of the IT organization to justify an investment in IT is increasingly tied to the ability of the organization to concretely demonstrate the business value of that investment.

- **More understanding of applications**

  While client server and n-tier applications are still common, as pointed out in The 2013 Application and Service Delivery Handbook, many applications are now based on a wide range of architectures; e.g., a Services Oriented Architecture (SOA). In addition, complex applications, such as Customer Relationship Management (CRM), are actually

comprised of several modules, with a range of network requirements. IT infrastructure and network professionals in particular need to better understand these new architectures and complex applications in order to ensure that the emerging set of technologies are designed and architected appropriately.

- **More emphasis on programming**

  While it is not true that all networking and data center professionals will become programmers, it is true that many senior level IT professionals will need an understanding of programming in order to better interact with the company's software development organization. It is also true that some network organizations will want to leverage the API functionality that the emerging technologies provide by having network professionals write programs that utilize those APIs.

The Survey Respondents were told that SDN is part of a broader movement to implement all IT functionality in software, referred to as Software Defined Everything (SDE) and they were asked a number of questions about how the SDE movement has and would likely impact their organization as well as how it would likely impact their jobs. There were 122 responses from people who are involved with enterprise communications networks and 19 responses from people who work for a service provider that offers WAN services. Given that there are only 19 responses from service providers that is not a large enough sample size to be statistically significant. It is, however, large enough to provide insight into the organizational impact that the ongoing adoption of software based functionality is having on WAN service providers.

For example, The Survey Respondents were asked if within the last year the SDE movement had prompted their IT organization to do a re-org. Nine percent of enterprise respondents said yes and 32% of service provider respondents said yes. These responses make it appear as if service providers are further along relative to reorganizing the company to leverage software-based IT functionality.

The Survey Respondents were also asked how much of an impact they thought that the SDE movement will have on the structure of their company's IT organization over the next two years? Their answers are shown in **Table 10**.

Table 10: Impact of SDN on Organizational Structure

Impact	Percentage of Responses
Very Significant Impact	4%
Significant Impact	16%
Moderate Impact	14%
Some Impact	18%
No Impact	23%
Don't Know	25%

Over the next two years the ongoing adoption of software-based IT functionality is likely to have an impact on the structure of IT organizations.

Some of the answers from service provider respondents when asked to indicate the type of organizational changes that had either already occurred or that they expected would occur include:

- The operations group is likely to be restructured;
- We now need to gather management from virtual devices;
- The company's technical experts have been consolidated into a single group;
- The company has set up a subsidiary and are in the process of moving IT employees to that subsidiary;
- The organization's OSS/BSSs need to be revamped.

When asked the same question, the answers from the enterprise respondents included:

- A shift from siloed specialists to service aligned generalists;
- A likely re-org around application development and network operations;
- An increase in cross functional teams and projects;
- Moving from a tower based organization to a DevOps model;
- An increased focus on software engineering;
- Team work will involve an enhanced mix of skills including programming, networking, virtualization and DevOps;

In addition, the Survey Respondents were asked how much of an impact they thought that the SDE movement will have on the nature of their jobs over the next two years? Their answers are shown in **Table 11**.

Table 11: Impact of SDN on Jobs

Impact	Percentage of Responses
Very Significant Impact	6%
Significant Impact	19%
Moderate Impact	16%
Some Impact	23%
No Impact	19%
Don't Know	18%

Over the next two years the ongoing adoption of software-based IT functionality is likely to have an impact on the jobs of IT professionals.

Some of the answers from service provider respondents when asked to indicate the type of impact on their jobs that had either already occurred due to the ongoing adoption of software-based IT functionality or that they expected would occur include:

- The product development life cycle will change;
- The job will require new skills in general and more knowledge of software in particular;
- The customer demands are unknown;

- Product development needs to be able to provide tools to manage and monitor the environment;
- There will be new business models, new product offerings that must be supported.

When asked the same question, the answers from the enterprise respondents included:

- The way to design, implement and troubleshoot networks will change a lot;
- The job will require new skill sets in general and more programming knowledge in particular;
- There will be new security requirements;
- As we adopt DevOps, broad based skills are required;
- There will be less emphasis on technology silos;
- New architectures will need to be developed;
- There will be a lot of re-training and re-trenching.

Dynamic Cloud, Dynamic Services

Service providers are on a journey to the cloud. Network function virtualization (NFV) and software-defined networking (SDN), when fully implemented, will create highly dynamic networks with an unprecedented level of scale, resiliency and programmability.

The result will be new dynamic services, where the network adapts to users' demands, rather than limits what the user can do. These new services promise to be more flexible and offer a better user experience. However, for service providers to remain viable businesses, it is critical that the migration to this new architecture does not disrupt existing services, and the new services do not cost more to deliver than users are willing to pay.

Alcatel-Lucent and Bell Labs have been with you on this journey from the beginning. From the first telephone, to the invention of the transistor, from the earliest digital telephone systems and cellular networks to today's advanced IP/optical and LTE networks, we have been the industry's leading pioneers. We are also an early leader in adapting cloud technologies to the telecom world, and we have the key solutions to get you started on the next stage of your journey.

The NFV Journey

NFV is the start of a multi-year journey; a journey that is being made possible as a result of many technical advances coming together simultaneously. The journey to a fully operational NFV network requires the coordination of three interlinked but separate development paths: virtualization, orchestration and automation. Balancing the investments a service provider allocates to each path has much to do with where they start and their strategy. No path should be considered in isolation.

1. Virtualization

The abstraction of the Telecom functions software from dedicated hardware to run on open commercial-off-the-shelf (COTS) hardware, as well as the need to balance performance and cost reductions, will force service providers to make critical roadmap decisions. Some functions will achieve significant advantages of scale and flexibility from COTS hardware, while other functions, or even the same function, may benefit from the performance advantages of dedicated hardware. This duality is likely to exist for a while, as we pass through a transition phase, but this should not complicate the operational model, provided the same management entities exist. While exact feature parity may not be critical, function performance and robustness cannot be compromised. Service providers should consider the many years of feature development put into the existing functions, and consider carefully how this work will be carried forward into the new mode.

While virtualization of the function is one activity on the path to full NFV, consideration must also be given to how the function will scale. Initially scaling may happen manually, but ultimately, it should be fully automated. Scale and distribution will drive a need for tight inter-virtual machine (VM) communication, and this must be achieved without performance impact.

Alcatel-Lucent's uniquely open approach and ecosystem

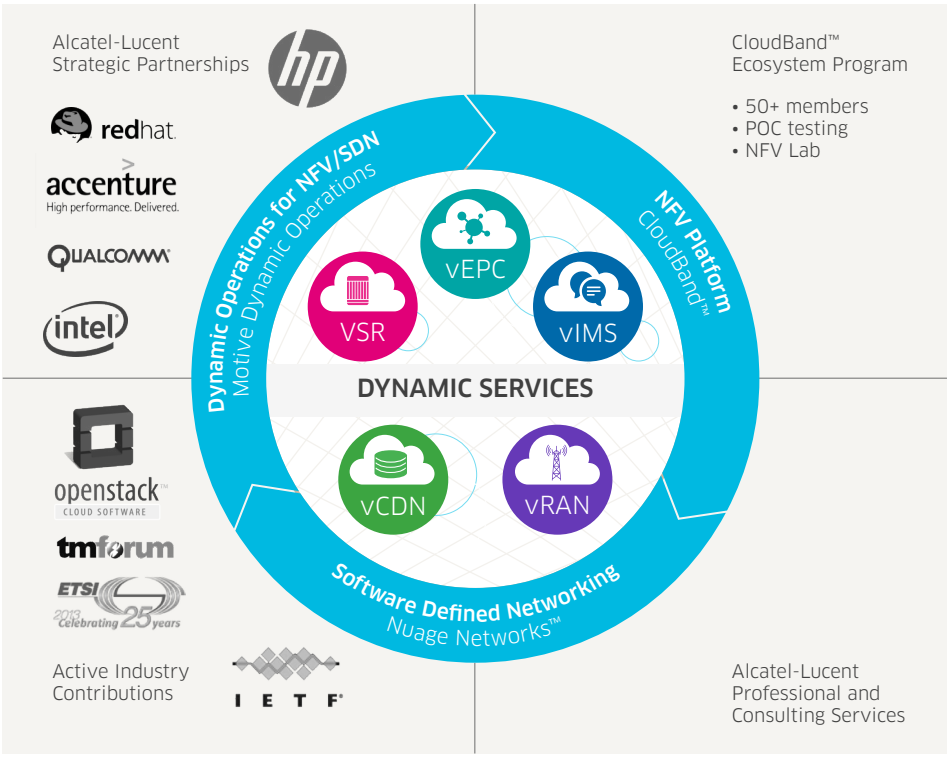

The path to full NFV may follow a number of steps as systems evolve:

1. Virtualized software running in a static mode on a defined COTS hardware and software build
2. Virtualized software functions on any COTS or other specialized virtual servers with manually triggered scaling
3. Full cloud implementation with auto scaling, resiliency and open APIs that enable dynamic service activation by third parties, including control of core network functions

When making a decision on which step to take first, the end game should be in sight or it may delay other decisions later on.

2. Orchestration

The orchestration and management of virtual machines needs to be done differently in a telecom network than a typical IT data center. Whether the service provider is offering a mobile app or real-time voice within a Web app (WebRTC) there will be many software routines all interconnected and sharing data across internal and external APIs. Each software module is uploaded onto a virtual machine image within a server. As a result, the telecom domain requires many thousands of virtual machines, which for reasons of resiliency and SLA integrity may be widely distributed. Managing the distribution to assure service performance requires a higher degree of orchestration.

The orchestrator automates the process of preparing and tracking virtual machines within the service provider's network. Each telecom function requires a different virtual machine setup and configuration. Through templates and recipes the orchestrator knows the configuration required to support each application. When a new function and/or more capability is required, an available virtual machine will be located and made available with the correct configuration.

The orchestrator is responsible for the lifecycle management of the virtual machine and its hosted function, including the creation of VM profiles and a wide variety of other functions. A horizontally scalable VNF management function enables the NFV platform to be set up as a Carrier Platform as a Service (CPaaS). The industry still needs to converge on a common scripting tool to create the VNF profiles. The Topology and Orchestration Specification for Cloud Applications (TOSCA) is considered a front-runner.

Quality of service metrics must also be standardized to ensure that when application performance is measured and monitored the performance is considered against a consistent metric and appropriate actions are taken to improve the metric.

3. Automation

As NFV scales, the operator must simultaneously manage the underlying network infrastructure. To do this cost effectively, it is necessary to automate the network to ensure it is in step with application demand. This is the role of SDN.

SDN is currently deployed in data centers where an overlay control layer is proving critical to meet the networking demands of the rapidly rising number of virtual machines. In these deployments, SDN ensures that network connections can be made as fast as the virtual machines within a server are created. The adoption of cloud computing within telecom networks additionally brings much shorter service lifecycles combined with increased application mobility. For typical telecom services, the location of the host for a service can move very rapidly. Thus the wide area network (WAN) environment is more dynamic than in data-center applications.

Adoption of SDN within the WAN will improve the resource and capacity utilization of the network by automating adjustments based on real-time usage. A fully dynamic network will be achieved by implementing NFV and SDN on top of a converged and programmable IP/optical network fabric to scale and automate application and service performance when and where it's needed.

Alcatel-Lucent has already developed the pieces, partners and ecosystem that operators will need to start down these three interconnected paths. We offer best of breed solutions for the different layers of NFV, using industry-supported open platforms and standards that avoid vendor lock-in. Our professional services organization operates a fully featured test bed environment where our partners, ecosystems of developers and service provider customers can ensure the continuity and resilience that real world deployments will demand.

Find out how we can help you on your journey to virtualization: **www.alcatel-lucent.com/solutions/cloud**

CloudBand

The industry reference NFV platform, CloudBand is a management and orchestration platform for open and massive distribution of virtualized telecom functions. With more than 30 customer trials, including most Tier 1 operators, CloudBand also has over 50 ecosystem members who share experiences, as well as implement and test services.

Virtualized Service Routing

The Alcatel-Lucent Virtualized Service Router (VSR) is a highly flexible, virtualized IP edge router optimized for x86 server environments. The VSR delivers a broad and rich set of virtualized IP edge applications and services. It is built to deliver high performance and elastic scalability, and enables rapid service innovation, extends service reach, opens new markets, and accelerates time to market while lowering operating costs with a homogenized physical infrastructure.

Virtualized IMS

The full portfolio of Alcatel-Lucent IMS solutions is now virtualized and commercially available. It has complete feature parity with native solutions, including the same committed SLAs, OpenStack with HEAT support today, migrating to TOSCA. New service innovations beyond VoLTE are enabled by our IMS APIs and WebRTC in partnership with leading application developers.

Virtualized IP Mobile Core

Alcatel-Lucent has virtualized the IP Mobile Core, including gateways, management, policy and charging, subscriber management and element and network management. It is a proven solution, widely deployed and fully supportive of 2G, 3G and LTE Mobile Core features. Deployed and tested in many NFV trials in conjunction with IMS, it has demonstrated tangible benefits for VoLTE.

Nuage Networks SDN

Nuage Networks is a leader in SDN. It focuses on modern datacenter requirements for multi-tenancy, full-featured routing and security at scale. The Nuage Networks platform 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 re-purposed on demand.

Motive Dynamic Operations

The new OSS for SDN and NFV, the Motive Dynamic Operations suite brings Motive's rich history with customer experience solutions to the management of SDN automation and NFV abstraction, as well as analytics and professional services – all designed to address different, critical touch points in the relationship between communications service providers and their customers.

Alcatel·Lucent

6WIND Virtual Accelerator

Enable NFV And Virtual Networking

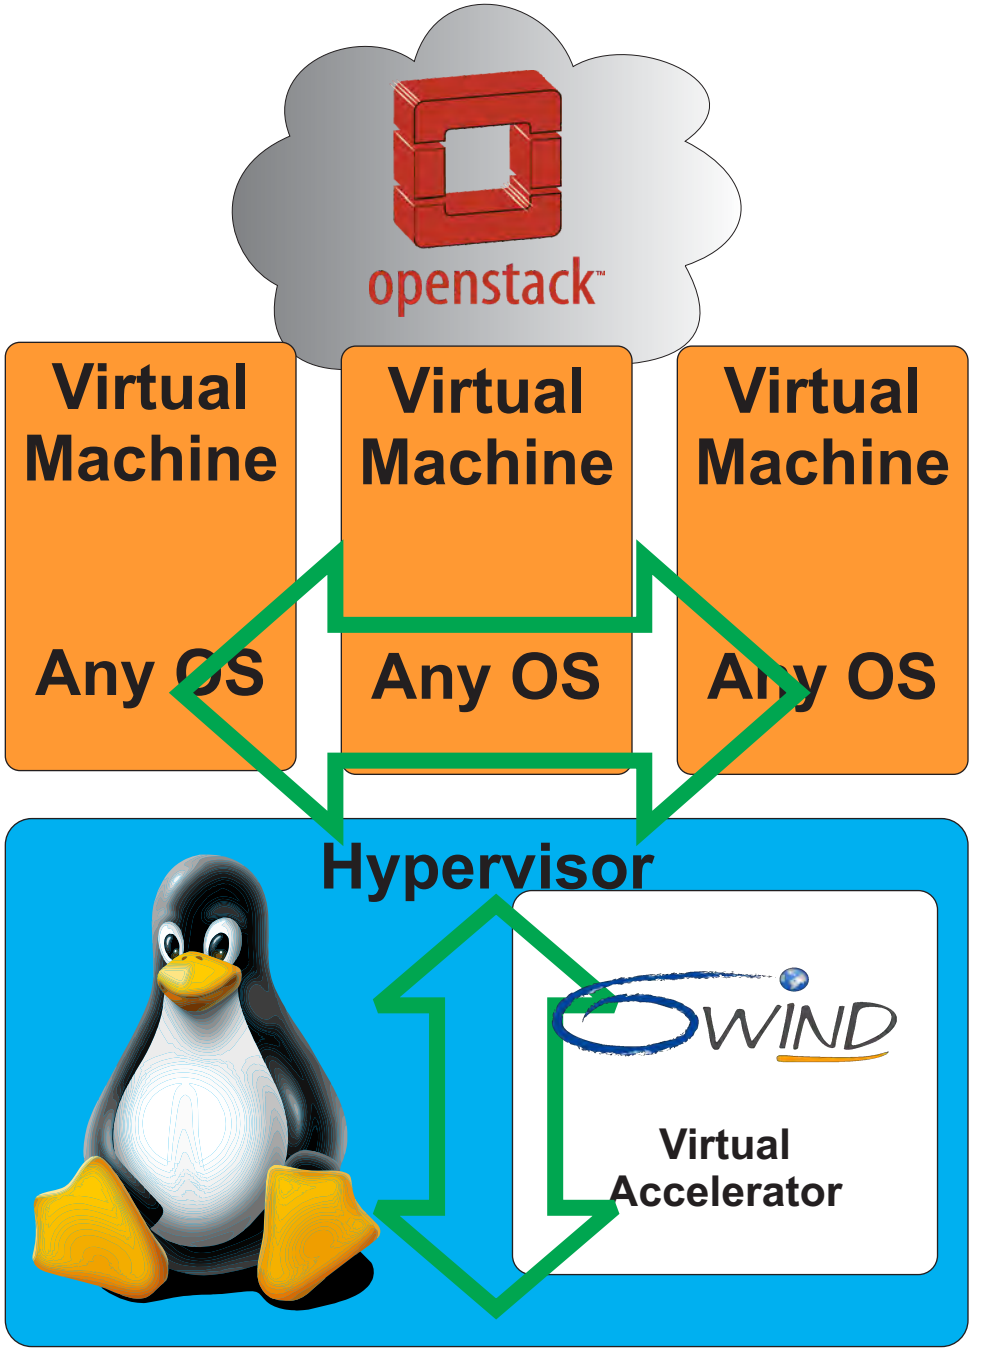

Transparent OpenStack orchestration support

High bandwidth for VM performance, density and communications

Complete virtual networking infrastructure

Support for Open vSwitch and Linux Bridge with no modifications

Network hardware independence for seamless hardware upgrades

6WIND

Wire Speed Virtual Switching
From Common Hardware

www.6WIND.com

Extending Service Performance Management into SDN and NFV Environments

Solution Benefits

- Holistic end-to-end visibility into physical, virtual, and hybrid service delivery infrastructure
- Proactive service triage helps resolve problems in real time and assures a positive customer/user experience
- Comprehensive service performance management platform across voice, data, and video services and applications
- Ultra high scalability assures service delivery across any size of service provider and enterprise infrastructure

Problem Overview

While the strategic importance of delivering IP-based services is constantly increasing, enterprises and service providers are being pressured to find ways to deliver these services faster, with higher quality, and more cost effectively. To achieve these goals, enterprises and service providers are gradually migrating their data center workloads onto a virtual infrastructure.

To realize the full potential of virtualization CapEx and OpEx efficiencies, enterprises and service providers need a comprehensive service delivery monitoring tool which offers end-to-end visibility across physical, virtual, and hybrid environments. To be truly beneficial, the tool needs to offer proactive service triage capabilities to reduce the mean-time-to-resolution, by identifying the root cause of service degradations and outages in real time.

Unfortunately, the traditional bottom-up triage methodology based on multi-vendor silo-specific Network Performance Management (NPM) and Application Performance Management (APM) tools is ineffective. It does not offer service level triage capabilities to key organizations, and lacks the ability to provide a view of the overall service.

The bottom-up triage methodology relies on disparate sets of data collected from multiple silo-specific tools, which makes it virtually impossible to gain an end-to-end holistic view of the service performance. Furthermore, these datasets lack the insight on the interrelationships and dependencies between service delivery components and therefore inhibit service triage activities. The overall result of relying on the bottom-up triage methodology is drastically increased service unavailability, reduced quality of end-user experience and loss in worker productivity.

Solution Overview

NetScout offers efficient service triage based on pervasive end-to-end visibility across physical, virtual, and hybrid service delivery environments. The triage is performed proactively by detecting service degradations in real time and is based on one cohesive, consistent set of metadata, for service provider and enterprise services. This metadata is generated by the patented Adaptive Session Intelligence™ (ASI) technology running in both virtual environments as well as nGenius® Intelligent Data Sources, and offers meaningful and contextual view of all interrelationships and dependencies across all service delivery components in physical, virtual, and hybrid environments.

NetScout's pervasive and scalable data collection is established by instrumenting strategic access points across the service delivery infrastructure using physical and virtual appliances. The packet flow data collection and aggregation is passive and non-intrusive and can scale to collect any required volumes of data across physical, virtual, and hybrid environments.

The nGeniusONE™ Performance Management platform aggregates, correlates, and contextually analyzes the metadata gathered from the nGenius Intelligent Data Sources in both physical and virtual environments. It then creates real-time holistic views of service performance, establishes performance baselines, and facilitates service-oriented troubleshooting workflows.

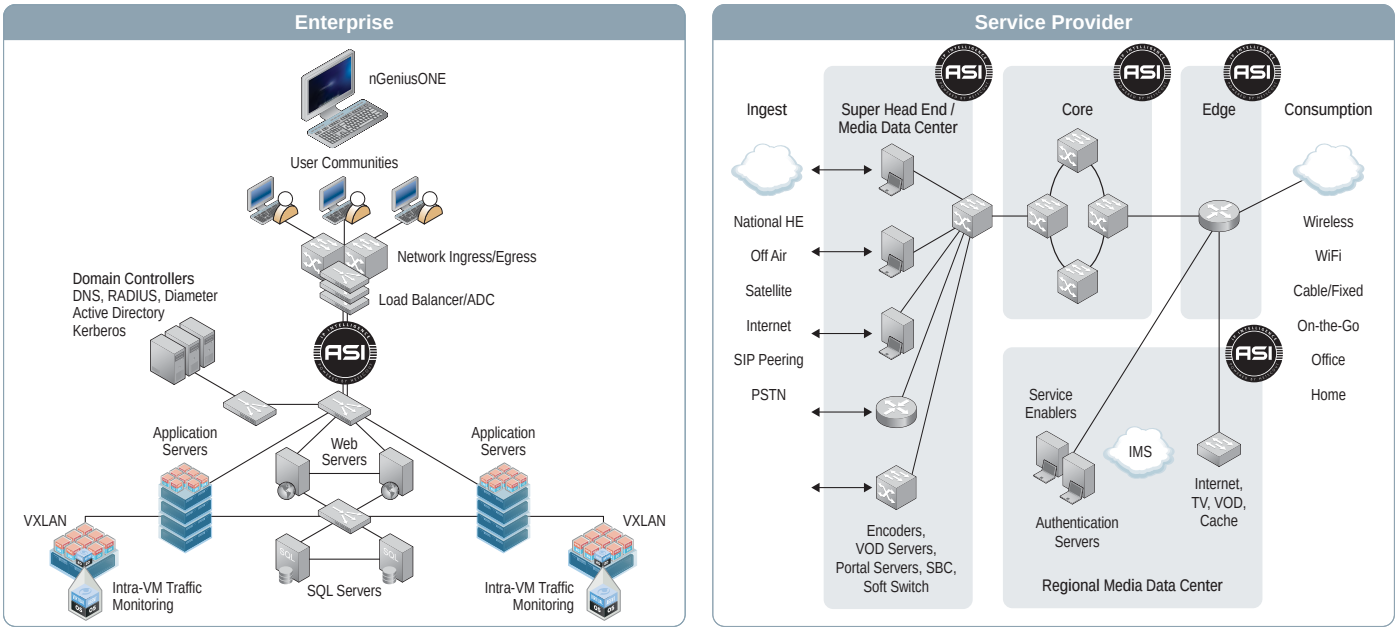

Figure 1: Service Instrumentation in Enterprise and Service Provider Environments.

Core Technologies

NetScout's unique ability to gain a pervasive end-to-end visibility into the service delivery environment, and offer proactive service triage is attributed to the following architectural principles and technologies:

- Utilize Packet Flow Data
- Provide Scalable Packet Flow Access
- Adaptive Session Intelligence (ASI)

Utilize Packet Flow Data

NetScout uses packet flow data as the foundation for generating highly scalable metadata that enables a comprehensive real time and historic view of all service components including physical and virtual networks, n-tier applications, workloads, protocols, servers, databases, users, and devices.

Provide Scalable Packet Flow Access

NetScout physical and virtual TAP network monitoring devices provide comprehensive and reliable access to packet flow data and establish strategic service visibility points across the entire service delivery infrastructure. The nGenius Packet Flow Switches (PFS) filter, aggregate, and distribute the targeted data to the nGenius Intelligent Data Sources in a transparent, selective, and efficient manner.

Adaptive Session Intelligence (ASI)

ASI is patented technology which uses a rich packet-flow data Deep Packet Inspection (DPI) engine to generate highly scalable metadata that enables a comprehensive real time and historic view of service, network, application, and server performance. This powerful deep packet inspection and data mining engine runs on nGenius Intelligent Data Sources, generating metadata based upon actual session traffic in real time as the packets cross physical or virtual links. The generated metadata provides important metrics such as application traffic volumes, application server response times, server throughputs, aggregate error counts, error codes specific to application servers and domain, as well as other data related to network and application performance. The ASI technology is the foundation of a highly scalable service delivery monitoring architecture which seamlessly collects, normalizes, correlates, and contextually analyzes data for all business services.

Service Delivery Monitoring in SDN Environments

NetScout has partnered with VMware, the global leader in virtualization and cloud infrastructure, to provide service delivery monitoring solutions in VMware NSX™ environments. These solutions enable NetScout to gain full visibility into applications traversing NSX environments in the following use cases:

- **Traffic between the VMs on the same hypervisor** is monitored by embedding NetScout's ASI patented technology into a virtual machine (VM) probe, which resides on the same hypervisor as the monitored VMs. NetScout's VM either analyzes the intra-VM traffic in a self-contained virtualized probe mode or redirects the traffic to an external nGenius Intelligent Data Source for analysis.
- **Traffic between VMs that reside in different hypervisors** is monitored by the nGenius Intelligent Data Sources that decode the VXLAN encapsulation and access the original packet flow data between the VMs.
- **Multi-tier East-West and North-South Data Center traffic** is monitored by collecting data from a combination of multi-tier physical and virtual service delivery environments, correlating, and contextually analyzing all the interrelationships and dependencies across all monitored service delivery components. These include n-tier applications, workloads, protocols, servers, databases, users, and devices.

Solution Benefits

NetScout's ability to provide end-to-end visibility into multi-tier physical, virtual, and hybrid service delivery environments combined with proactive service triage, helps address the key problems associated with silo-specific, component-based, bottom-up performance management approaches.

Attribute	Bottom-Up Triage Problems	NetScout's Solution	IT Benefits
End-to-End Visibility	• Point visibility into individual service delivery components from a variety of multi-vendor silo-specific tools. Lacks the necessary insight into interrelationships of service delivery components.	• Holistic end-to-end visibility into service delivery infrastructure using one cohesive, consistent set of data, for service provider and enterprise services delivered in physical and virtual environments.	• Optimize experience of user communities and customers. • Comprehensive solution from a single vendor. • Full visibility into services running in physical, virtual, and hybrid environments.
Effective Service Triage	• Reactive and time consuming triage results in poor user experience, and extended service downtime impacting multiple users.	• Proactive service triage helps resolve service degradation in real time, before a large number of users are impacted.	• Increase service uptime and end-user productivity. • Support more services with existing IT resources. • Reduce time wasted in war rooms.
Scalability	• Lacks scalability required to assure delivery of modern business services for service providers and enterprises.	• Scales to assure service delivery across any size of service provider and enterprise infrastructure.	• Optimize your investment in performance management by gradually expanding the solution over time.

About NetScout Systems, Inc.

NetScout Systems, Inc. (NASDAQ:NTCT) is the market leader in application and network performance management solutions that enable enterprise and service provider organizations to assure the quality of the user experience for business and mobile services. Used by 92 percent of Fortune 100 organizations and more than 165 service providers worldwide, NetScout's technology helps these organizations proactively manage service delivery and identify emerging performance problems, helping to quickly resolve issues that cause business disruptions or negatively impact users of information technology. For more information about NetScout, visit www.netscout.com.

NETSCOUT

Americas East
310 Littleton Road
Westford, MA 01886-4105
Phone: 978-614-4000
Toll Free: 800-357-7666

Americas West
178 E. Tasman Drive
San Jose, CA 95134
Phone: 408-571-5000

Asia Pacific
17F/B
No. 167 Tun Hwa N. Road
Taipei 105, Taiwan
Phone: +886 2 2717 1999

Europe
One Canada Square
29th floor, Canary Wharf
London E14 5DY, United Kingdom
Phone: +44 207 712 1672

NetScout offers sales, support, and services in over 32 countries.

For more information, please visit www.netscout.com or contact NetScout at 800-309-4804 or +1 978-614-4000

nuagenetworks

@nuagenetworks

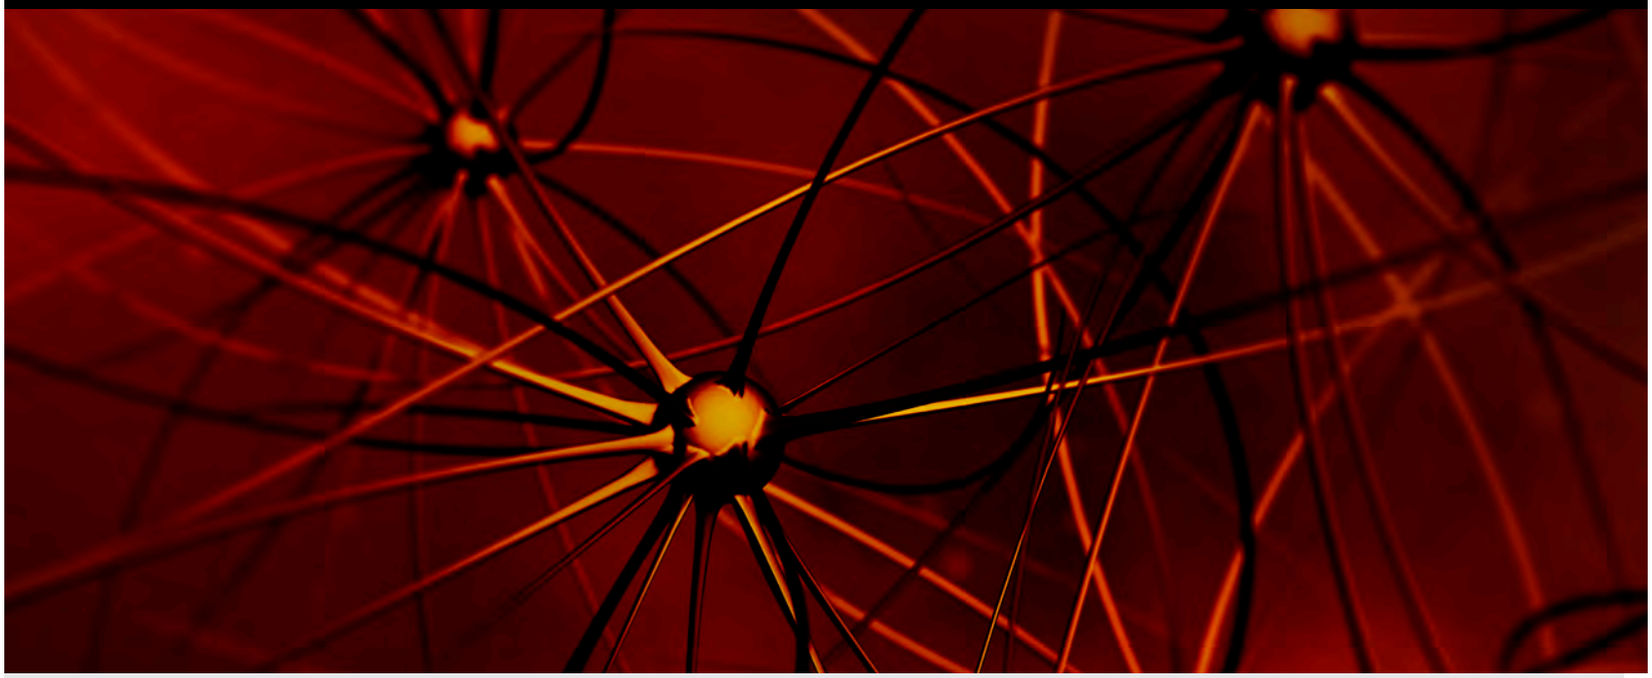

The Cloud Network Unbound

Virtualized and automated networking across datacenters and branch offices

Cloud computing is changing the way enterprises access and consume data. To remain competitive, businesses know they must be able to react quickly to market changes. The cloud addresses their need for speed, agility and responsiveness. Unfortunately, today's data communications networks aren't keeping pace. In fact, they're struggling to deliver consistent, on-demand connectivity and things are only going to get more challenging. Fortunately, Nuage Networks has a solution.

Nuage Networks leverages Software Defined Networking (SDN) to unleash the power of the cloud, giving enterprises the freedom and flexibility to:

- Connect sites, workgroups and applications faster, more securely and more cost effectively
- React to change easily
- Respond to growth seamlessly

Nuage Networks makes the network as responsive as your business needs it to be — from the datacenter to remote locations.

Our solutions close the gap between the network and cloud-based consumption models, 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 and across the wide area network.

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.

Take advantage of a fully virtualized services platform

Cloud-based datacenters have unshackled the IT environment, making it possible for applications to request additional compute and storage on an as-needed basis. Extending the reach of virtualized network services from the datacenter to remote locations further enhances the enterprise's ability to respond to business imperatives at cloud speed. Peak demands can be provisioned "just in time", which lowers operational costs and makes it possible to share compute resources across applications. Geography is taken out of the equation.

Nuage Networks SDN solutions enable you to react to changes in your datacenter or at branch locations with speed, agility, and flexibility. Our solutions seamlessly connect your datacenters and the wide area network, so networking across the whole environment is fluid and responsive to changing business conditions.

By improving efficiency, resiliency and security, our products enable networks to be built and operated at any scale – from a single rack to Fortune 500 scale.

Our SDN solutions work closely together and deployment is flexible, so you can focus on the area most in need of help.

Responsive datacenter networking

Build robust and highly scalable networking infrastructures with the **Nuage Networks Virtualized Services Platform (VSP)**. These new infrastructures will let you instantaneously deliver compute, storage and networking resources securely to thousands of user groups.

Virtual private networking on your terms

The **Nuage Networks Virtualized Network Services (VNS)** enables you to respond faster and with greater agility to changes in your wide are network environment. A self-serve portal allows enterprise end users to self-manage moves, adds and changes, significantly reducing the time and effort required to manage the wide area network.

Nuage Networks SDN solutions are specifically designed to:

Simplify operations for rapid service instantiation	Address changing business requirements with flexible, adaptable services	Support massive scalability and hybrid models with secure, open infrastructure
■ Define network service requirements in clear, IT-friendly language ■ Bring services up using automated, policy-based instantiation of network connectivity ■ Dramatically reduce time to service and limit potential for errors	■ Adapt datacenters and private networks dynamically ■ Detect newly created and updated virtual machines within the datacenter and respond automatically by adapting network services according to established policies, instantly making available new applications to all users regardless of location	■ Benefit from distributed, policy-based approach that allows multiple virtualization platforms to interoperate over a single network ■ Optimize the datacenter network and private network by separating service definition from service instantiation

Nuage Networks SDN solution components

Nuage Networks VSP is the first network virtualization platform to address modern datacenter requirements for multi-tenancy, full-featured routing and security at scale. 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 VSP integrates seamlessly with wide area business VPN services. It is also particularly effective when deployed with Nuage Networks VNS for a cloud-optimize network that spans the datacenter right out to your remote locations.

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 ground breaking 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 ensures the datacenter and wide area network are able to respond instantly to demand and are boundary-less.

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

Nuage Networks SDN Portfolio

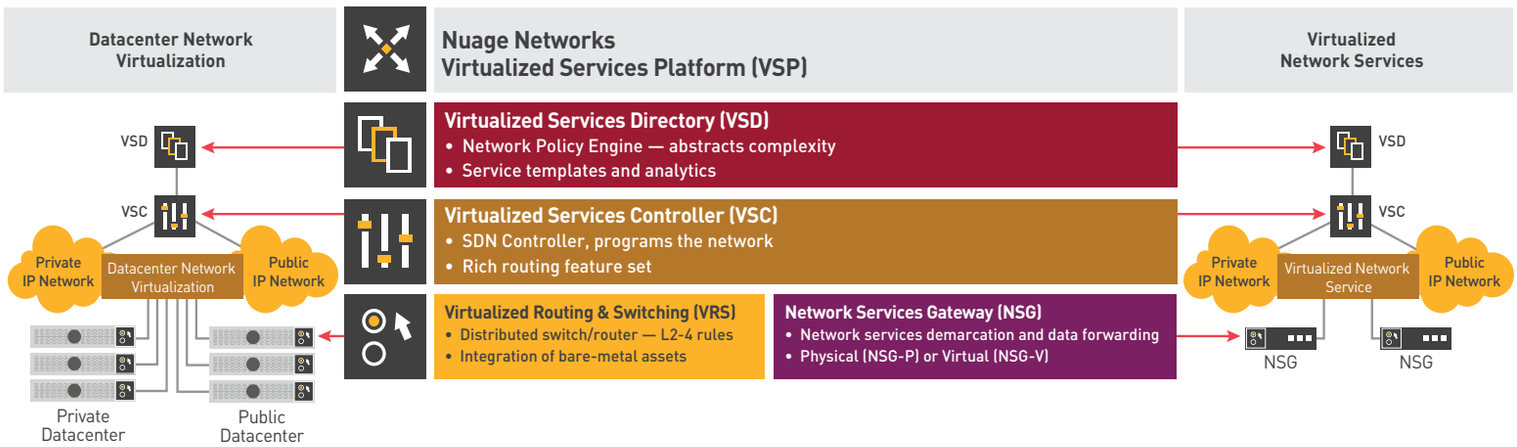

nuagenetworks

Software-Defined Networking

Are your management tools prepared?

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

Are you prepared for this evolution?

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

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

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

EMC²

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 or Steven Taylor.

Published by Webtorials Editorial/Analyst Division

www.Webtorials.com

Division Cofounders:

Jim Metzler

Steven Taylor

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.