The 2016 Guide to SDN and NFV

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

Platinum Sponsors:

A10

CISCO

Hewlett Packard
Enterprise

MASERGY
Performance Beyond Expectations

NETSCOUT™
Guardians of the Connected World

Sonus

Gold Sponsors:

radware

SevOne

Produced by:

Webtorials

TABLE OF CONTENTS

Executive Summary

Introduction

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 2016 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 architectural approaches.

The Guide will be published both in its entirety and in a serial fashion. This document is the eighth of the serial publications and it will present an executive summary of the preceding seven publications. Below is a listing of all of the publications that comprise The Guide:

1. A SDN Status Update
2. The Use Cases and Business Case for SDN
3. The Operational Impediments to Implementing SDN
4. A NFV Status Update
5. Architectural Considerations and Use Cases for NFV
6. The Operational Impediments to Implementing NFV
7. The SDN and NFV Ecosystem
8. An Executive Summary of The Guide

The Guide is based in part on *The 2015 Guide to SDN and NFV* (The 2015 Guide). To limit the size of The Guide, some of the introductory material, such as a description of the basic SDN architecture that was contained in The 2015 Guide has been eliminated. The 2015 Guide, however, is still available online.

This guide contains the results of two surveys. One of the surveys focused on SDN, was administered in October 2015 and was completed by 131 IT professionals. The other survey focused on NFV, was administered in December 2015 and January 2016 and was completed by 144 IT professionals.

A SDN Status Update

The survey data indicated that the majority of organizations are actively involved in analyzing, trialing or using SDN in production. However, that was true a year ago and comparing this year's survey data with last year's data indicates that SDN remains stuck on the edge of crossing the chasm from being used primarily by early adopters to where it is widely used and that it will likely be on the edge for at least another year or two.

One thing that has changed in the last year is the amount of resources that has gone into creating open source solutions. According to Dan Pitt, executive director of the Open Networking Foundation (ONF), the growing interest in SDN-related open source projects will accelerate the adoption of SDN. His belief in the importance of open source based solutions is the reason why in February 2015 the ONF launched an open source community and code repository called OpenSourceSDN.org. The role of this community is to sponsor and develop open SDN solutions in order to provide greater adoption of open SDN.

The interest in open source has led to a situation where there are multiple open source-based SDN controllers including one from ON.Lab and one from the OpenDaylight (ODL) Project. The ON.Lab is a non-profit organization whose mission is to "bring openness and innovation to the Internet and Cloud for the public good". One of the ON.Lab's primary projects is ONOS (Open Network Operating System) – an open source SDN operating system for service providers. In September 2015 ON.Lab released the fourth version of ONOS, code named Drake. According to ON.Lab "Drake adds new security, configuration and application level feature sets with improvements to the northbound and southbound including REST, API and GUI additions and upgrades throughout. In addition to contributing to ONF's Atrium, ONOS has expanded collaboration with other open source communities to develop new distributions including work with the CloudRouter® Project and it will soon be part of the Open Platform for NFV Project (OPNFV)." In October 2015 the ONOS project joined the Linux Foundation.

The ODL Project, which was founded in April 2013, is a collaborative open source project that is also hosted by The Linux Foundation. The goal of the project is to facilitate a community-led, industry-supported open source framework, including code and architecture, to accelerate and advance a common, robust SDN platform and to create a solid foundation for NFV. Towards that end, the ODL project claims that a number of vendors use ODL code as the basis of their SDN products[1] and that its code is also used by the OPNFV platform. In June 2015 the consortium announced the availability of its third software release, called Li7thium.

An important area where progress is being made relative to the evolution of SDN is the North Bound Interface (NBI) that sits between a SDN controller and the business applications and network services that utilize the controller. In 2013 the ONF established a working group to focus on the NBI. Given that traditional standards activities are widely viewed as not being agile enough for the current environment, the goal of the working group was not to develop a standard for the NBI in the traditional sense of the term *standard*. Rather, the goal was to develop a rough consensus and collaboration around community developed NBIs.

Dave Lenrow is the chair of the ONF's NBI working group. Lenrow said that the ONF's NBI initiative is "Essentially doing an experiment in collaborative agile development with open source projects. Instead of spending years trying to prove on paper that our architecture works we throw some experimental API stuff to multiple OSS projects (e.g., ODL, ONOS) and let implementers provide feedback on what works

[1] https://www.opendaylight.org/solutions-provider-directory

and what doesn't with a *fast fail* approach. Our members want ONF as a neutral third party to define the basic software artifacts (e.g., Information model, principles of operation) that get implemented on many vendor's solutions."

The Drivers and Inhibitors of SDN

Unlike the situation last year, there is currently as much interest in either implementing SDN in the WAN or using a SDN-based WAN service as there is in implementing SDN in the data center. Given the breadth of SDN-based use cases described in The Guide, it is becoming difficult to talk about SDN without specifying if that is SDN deployed in the data center, the WAN or the branch/campus.

Table 1 lists the 5 top drivers of implementing SDN in varying segments of the network.

Table 1: Top 5 Drivers

Data Center	WAN	Branch and Campus
Support the dynamic movement, replication and allocation of virtual resources	Ease the administrative burden of configuration and provisioning	Ease the administrative burden of configuration and provisioning
Ease the administrative burden of configuration and provisioning	Better utilize network resources	Better utilize network resources
Better utilize network resources	Perform traffic engineering with an end-to-end view of the network	More easily scale network functionality
Perform traffic engineering with an end-to-end view of the network	More easily scale network functionality	Support the dynamic movement, replication and allocation of virtual resources
More easily scale network functionality	Support the dynamic movement, replication and allocation of virtual resources	Reduce OPEX

Table 2 lists the 5 top inhibitors of implementing SDN in varying segments of the network.

Table 2: Top 5 Inhibitors

Data Center	WAN	Branch and Campus
Concerns about how we would integrate SDN into the rest of our infrastructure	Concerns about how we would integrate SDN into the rest of our infrastructure	Concerns about how we would integrate SDN into the rest of our infrastructure
The immaturity of the enabling technologies	The lack of a compelling business case	The lack of a compelling business case
The confusion and lack of definition in terms of vendors' strategies	The immaturity of the enabling technologies	Possible security vulnerabilities
Other technology and/or business priorities	The immaturity of the current products	The immaturity of the current products
The lack of a compelling business case	Possible security vulnerabilities	Other technology and/or business priorities

Easing the administrative burden of configuration and provisioning tends to be the primary driver of SDN adoption and concerns about how it would be integrated into the rest of the infrastructure is the primary inhibitor. After that, the drivers and inhibitors for SDN vary somewhat based on whether SDN is deployed in the data center, the WAN or the branch/campus.

The Operational Impediments to Implementing SDN

SDN has the potential to make implementing effective security easier and it has the potential to make that harder. One of the ways that SDN can enhance security is by implementing security services on OpenFlow-based access switches that can filter packets as they enter the network. Another such example is role based access that is implemented by deploying a role-based resource allocation application that leverages the control information and capability of the SDN controller.

Some of the security challenges that are associated with SDN include:

- 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 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.

The Guide contains a set of questions that network organizations should ask vendors relative to the security of their SDN solution.

The Guide describes OpenStack and points out that orchestration engines such as OpenStack are important to both SDN and NFV. As explained in The Guide, in conjunction with the orchestration engine, the role of the SDN controller is to translate the abstract model created on the orchestration engine into the appropriate configuration of the virtual and physical resources that will deliver the desired service. For example, the orchestration engine 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;
- Attach Network Services to a VM or chain Network Services between VMs.

The Guide highlights the fact that in SDN environments the challenges associated with end-to-end service performance management are more demanding than they are in traditional network environments. Some of the reasons for that are that in an SDN environment:

- There is a combination of physical and virtual resources that is changing dynamically.
- The SDN controller needs to be instrumented and monitored just as any other application server and the southbound protocol needs to be monitored the same way as any other protocol.
- Network management organizations need tools that enable them to be able to dynamically discover, procure, allocate and reconfigure resources.
- 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 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.

The Guide contains a set of questions that network organizations should ask vendors relative to the security of their SDN solution.

The Guide also positions SDN as being a part of a broader movement to implement all IT functionality in software, referred to as Software Defined Everything (SDE) and points out that 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;
- More focus on setting policy;
- More knowledge of the business;
- More understanding of applications;
- More emphasis on programming.

Almost a third of the survey respondents indicated their belief that over the next two years that the ongoing adoption of software-based IT functionality will have either a significant or a very significant impact on the structure of their IT organization. In addition, over a quarter of the survey respondents indicated their belief that over the next two years that the ongoing adoption of software-based IT functionality will have either a significant or a very significant impact on their jobs. The Guide indicated the types of changes that the survey respondents expect to see.

A NFV Status Update

Roughly three years ago an Industry Specifications Group (ISG) for Network Functions Virtualization (NFV ISG) was formed under the auspices of the European Telecommunications Standards Institute (ETSI NFV ISG). That ETSI group has primarily championed the interest that Communications Service Providers (CSPs) have with NFV. More recently, the Open the Open Networking User Group (ONUG) has emerged to champion the corresponding interest that enterprises have with what the group refers to as Network Services Virtualization (NSV).

There clearly are differences between what ETSI is trying to accomplish with NFV and what ONUG is trying to accomplish with NSV. For example, CSPs hope to virtualize some functionality that few if any enterprise organizations implement and their need for scale far surpasses what is needed by the vast majority of enterprise organizations. In addition, CSPs are notably more likely to have a requirement to link the usage of virtualized network functions to their billing systems than do enterprise organizations. However, if you change at most a few words in how ONUG describes the NSV use case it sounds exactly like what ETSI and others are trying to achieve with NFV. As a result, it makes sense to look at NFV as being applicable to both CSPs and enterprise organizations.

Until recently, many people regarded SDN and NFV as separate initiatives. That is changing. Some of the ways that ETSI believes that NFV and SDN complement each other include:

- The SDN controller fits well into the broader concept of a network controller in an NFV-Infrastructure (NFVI) network domain as defined in ETSI's NFV architectural framework.
- SDN can play a significant role in the orchestration of the NFV Infrastructure resources, both physical and virtual, enabling functionality such as provisioning, configuration of network connectivity, bandwidth allocation, automation of operations, monitoring, security, and policy control.
- SDN can provide the network virtualization required to support multi-tenant NFVIs.

The survey data supported the notion that the perception of the relationship between SDN and NFV is changing. For example, the vast majority of the survey respondents indicated their belief that SDN and NFV are complimentary activities. In addition, only a small percentage of survey respondents indicated that they believe that SDN and NFV are totally independent activities.

The adoption of NFV looks similar to the adoption of SDN. For example, currently only a modest number of IT organizations have implemented NFV in a production network while a somewhat large percentage of IT organizations are currently in varying stages of analyzing NFV. While the state of adoption is similar, the factors driving and inhibiting the adoption of NFV are not very similar to the ones driving and inhibiting the adoption of SDN. For example, by a wide margin, the primary factor driving interest in NFV is the reduction in the time it takes to deploy new services. There isn't a single dominant inhibitor to the adoption of NFV, but a number of inhibitors of roughly equal value. This includes:

- The lack of a compelling business case;
- Concerns about end-to-end service provisioning;
- Concerns about security vulnerabilities;
- The immaturity of the current products;
- The need to reskill our employee base.

NFV: Architectural Considerations and Use Cases

Before an organization adopts NFV they need to address some key considerations relative to how they will architect their data center to support NFV and related initiatives.

The architectural considerations that are discussed in The Guide are:

- Big Bang vs. Piecemeal Approach;
- Software Modularity;
- Technology Considerations;
- Software-Centric Design;
- The Role of Open Source;
- Relationship with SDN;
- Programmatic interfaces;
- A Fresh Approach to High Availability;
- The (potential) end of Moore's Law.

Unfortunately, the survey data indicates that two thirds of IT organizations have made little or no progress towards the development of a NFV architecture. The survey data also indicates that while some organizations are starting to make process towards the development of a NFV architecture, the majority are not.

The Guide discusses nine potential use cases for NFV that have been defined by the ESTI NFV ISG.

Those use cases are:

- NFV Infrastructure as a Service;
- Virtual Network Functions as a Service;
- Virtualization of the Home Environment;
- VNF Forwarding Graph;
- Virtual Network Platform as a Service;
- Virtualization of Mobile Core Network and IP Multimedia Subsystem;
- Virtualization of the Mobile Base Station;
- Virtualization of Content Delivery Networks;
- Virtualization of Fixed Access Network Functions.

NFV: Operational Impediments

The Guide discusses a number of the management challenges that are associated with NFV. Those challenges are the:

- Dynamic relationships between software and hardware components;
- Dynamic changes to physical/virtual device configurations;
- Many-to-Many relationships between network services and the underlying infrastructure:
- Hybrid physical/virtual infrastructures that need to be managed;
- Evolving performance monitoring challenges;
- Fact that network services may span multiple service providers;
- Fact that VNFs will be new types of components in the network;
- Need for tighter IT and Network Operations collaboration;
- Expanding hybrid environments;
- Need for a shared information model;
- Growing need and importance of a policy based architecture.

The survey data showed that there is broad recognition on the part of IT organizations that the adoption of NFV creates new management challenges such as the ones listed above. However, the data also indicated that the vast majority of IT organizations have made little or no progress relative to determining how they will respond to NFV-related management challenges. On a somewhat optimistic note, the data indicated that over the next year the vast majority of IT organizations will spend at least a modest amount of time working on developing an approach to how they will respond to NFV-related management challenges.

Similar to the situation with SDN, the adoption of NFV has the potential to impact the role of network organizations and network professionals. That fact was recognized by the survey respondents, roughly a third of whom indicated their belief that over the next two years that the adoption of NFV is likely to have a significant or very significant impact on the structure of their organization. In addition, roughly 40% of the survey respondents indicated their belief that over the next two years the adoption of NFV is likely to have a significant or very significant impact on the skill base of their organization. The Guide indicated the types of changes the survey respondents expected to see.

The SDN and NFV Ecosystem

The Guide identifies the primary classes of vendors that either currently do, or can be expected to provide either parts or all of a SDN solution. Included in the discussion is the value proposition of this class of vendor as well as a set of representative vendors. The classes of SDN vendors included in The Guide are:

- Merchant Silicon/Chip Vendors;
- HyperScale Data Centers;
- Telecom Service Providers;
- Switch Vendors;
- Network and Service Monitoring, Management and Automation;
- Providers of Network Services;
- Testing Vendors and Services;
- Providers of SDN Controllers;
- Providers of Telcom Service Provider's Infrastructure/ Optical Networking;
- Server Virtualization Vendors.

The Guide also identifies the primary classes of vendors that either currently do, or can be expected to provide either parts or all of a NFV solution. Included in the discussion is the value proposition of this class of vendor as well as a set of representative vendors. The classes of NFV vendors included in The Guide are:

- Telecom Service Providers;
- Merchant Silicon/Chip Vendors;
- Network Systems and Electronic Equipment Vendors;
- Virtualized Network Service and Cloud Service Vendors;
- SDN Controller Software Vendors;
- NFVI Providers;
- Orchestration Software Vendors;
- Network Monitoring, Management and OSS/BSS Vendors;
- Hypervisor Vendors;
- Test Equipment Vendors and Test Services;
- Standards Bodies and Related Communities.

Software Defined Networking (SDN): A Status Update

Status of SDN Adoption

The Survey Respondents were given a set of alternatives and were asked to indicate the alternatives that described their company's current approach to implementing SDN and were allowed to choose all that applied to their company. Their responses are shown in **Table 3.**

Table 3: Approach to Implementing SDN

Approach to Implementing SDN	Percentage
We have not made any analysis of SDN	14%
We will likely analyze SDN sometime in the next year	19%
We are currently actively analyzing the potential value that SDN offers	33%
We expect that within a year that we will be running SDN either in a lab or in a limited trial	17%
We are currently actively analyzing vendors' SDN strategies and offerings	29%
We currently are running SDN either in a lab or in a limited trial	15%
We currently are running SDN somewhere in our production network	9%
We looked at SDN and decided to not do anything with SDN over the next year	9%
We expect that within a year that we will be running SDN somewhere in our production network	12%
Don't know/Other	9%

The data in **Table 3** indicates that the implementation of SDN in production networks remains limited. In addition, comparing the data in **Table 3** to the responses to the same question a year ago yields the conclusion that:

SDN remains stuck on the edge of the chasm and will be there for at least another year or two.

The Open Networking Foundation (ONF)

The Open Networking Foundation (ONF) is the organization that is most closely associated with the development and standardization of SDN. As of September 2015, the ONF had over 140 members.

Most networking professionals associate the ONF with the OpenFlow protocol. That's reasonable because OpenFlow was developed at Stanford, with v1.0 published at the end of 2009 and v1.1 at the beginning of 2011. In March of 2011, the ONF was created and the intellectual property rights of OpenFlow were transitioned to it. Part of the oft-stated vision of the ONF is to make OpenFlow-based SDN the new norm for networks.

While the ONF is bullish on the future of OpenFlow, as described below, there are a number of alternatives to OpenFlow. In a recent blog, Dan Pitt, the executive director of the ONF, addressed the future of OpenFlow. According to Pitt, "OpenFlow is the standard southbound protocol designed for SDN and it is vendor neutral. Nothing else is. It's now appearing in chipsets, white-box switches and branded switches, in addition to the hypervisor switches where it's been pervasive. With forwarding and control separate, OpenFlow-based switches offer amazing price-performance, while separate control software allows operators to tailor the network's behavior to their business priorities. This, of course, is the goal of SDN."

In that blog Pitt went on to discuss some of the large, highly visible current implementations of OpenFlow, including:

- Google's replacement of its worldwide data-center interconnection network with a pure OpenFlow network;
- Google's use of OpenFlow within their data centers;
- AT&T's use of OpenFlow to configure the Open vSwitch (OVS) that it uses in its universal Customer Premise Equipment (uCPE);
- Alibaba's use of OpenFlow within its hybrid SDN cloud network.

In order to accelerate the adoption of OpenFlow, the ONF continues to drive OpenFlow conformance testing and certification. One of the goals of testing and certifying OpenFlow is to get to a state where users can get a controller from one vendor and switches from another. However, in a recent conversation with the author, Pitt said that we haven't reached that state yet and that at least in the short term, that network organizations that implement an OpenFlow-based SDN solution need to buy the controller and the switches from the same company.

In the conversation that the author had with Pitt, Pitt said that he thought that the growing interest in SDN-related open source projects would accelerate the adoption of SDN. His belief in the importance of open source based solutions is the reason why in February 2015 the ONF launched an open source community and code repository called OpenSourceSDN.org. The role of this community is to sponsor and develop open SDN solutions in order to provide greater adoption of open SDN. According to Pitt, the work of OpenSourceSDN.org is complementary and interoperable with work being done by open source organizations such as OpenDaylight (ODL), the Open Networking Lab (ON.Lab) and the Open Platform for NFV (OPNFV).

Open Source projects will likely accelerate the adoption of SDN.

One of the programs that falls under the OpenSourceSDN.org umbrella, referred to as Boulder, is discussed below. Boulder is attempting to develop a consensus and collaboration around a community developed approach to SDN's north bound interface. In June of this year the ONF announced Atrium, another one of the programs that falls under the OpenSourceSDN.org umbrella. One of the issues that Atrium is designed to address is that most open source initiatives are stand-alone activities. Atrium's mission is to integrate existing open source solutions and to possibly add some additional functionality with the goal of responding to user-defined use cases.

Atrium issued its first release, called Atrium 2015/A, in June of this year. Some of the functionality included in that release includes a:

- BGP peering application that runs on ONOS and includes the Quagga BGP stack;
- Collection of OpenFlow v1.3 device drivers, meant for talking to vendor equipment with different hardware pipelines;
- Indigo OpenFlow client along with other support for white-box switches;
- Full testing suite for functionality tests.

The group intends a second release in December of 2015. That release will include porting the first release to run on open source SDN controllers developed by the ODL project.

The Northbound & Southbound Interfaces

Figure 1 contains a graphical representation of the SDN architecture as envisioned by the ONF. Relative to **Figure 1**, the northbound interface is the interface that enables communications between the control layer and the application layer and the southbound interface is the interface that enables communications between the control layer and the infrastructure layer.

Figure 1: ONF's SDN Architecture

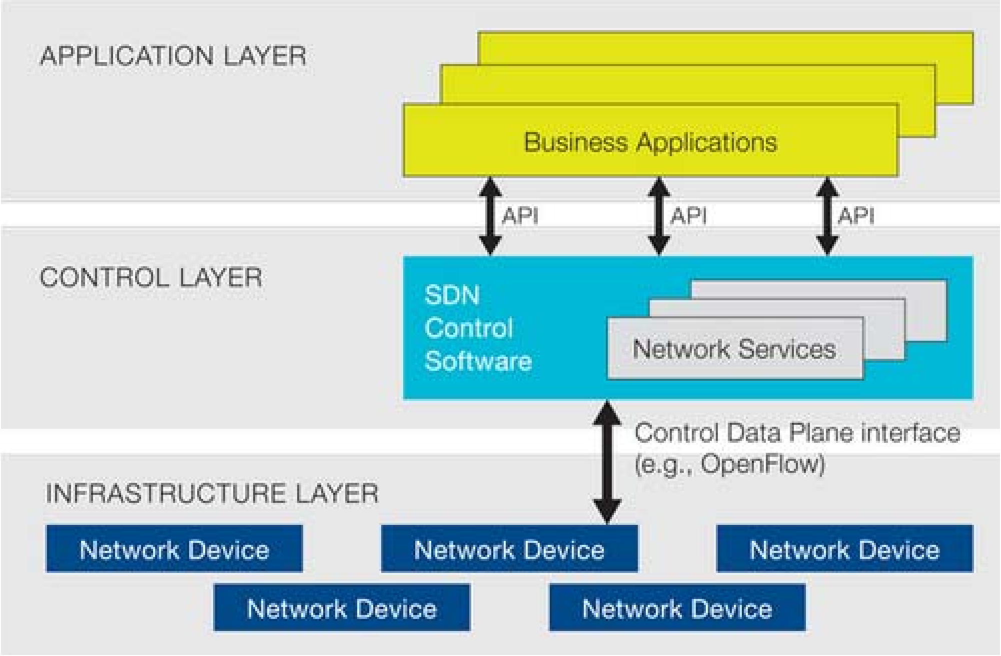

The Northbound Interface

As explained below, there are standards such as OpenFlow that can be used for the southbound interface between the controller and the subtending network elements. However, there isn't a standard for the North Bound Interface (NBI) between the controller and the business applications and network services that utilize the controller. As recently as two years ago there was considerable debate in the SDN community about the viability of either creating such a standard or at least developing a consensus about how the NBI should function. The argument against the approach of developing either a standard or a consensus was that we were so early in the development of SDN that we didn't know what should go into the NBI and hence it made no sense to coalesce around a particular NBI. Part of the argument for such an approach was that it was required in order to avoid vendor lock in. Proponents of the approach also argued that there were numerous controllers on the market, each with their own NBI and none of which had significant market share. According to the proponents, the lack of either a standard or a consensus about how the NBI should function was impeding the development of SDN because without it 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.

In 2013 the ONF established a working group to focus on the NBI. Given that traditional standards activities are widely viewed as not being agile enough for the current environment, the goal of the working group was not to develop a standard for the NBI in the traditional sense of the term *standard*. Rather, the goal was to develop a rough consensus and collaboration around community developed NBIs. The group's complete charter was outlined in a white paper. One of the interesting concepts that that white paper discusses is the need for APIs at different "latitudes". The idea was 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.

Dave Lenrow, a distinguished architect in HP's advanced technology group is the chair of the ONF's NBI working group and he is also on the technical steering committee for ODL and OPNFV. In an interview with the author, Lenrow said that the ONF's NBI initiative is "Essentially doing an experiment in collaborative agile development with open source projects. Instead of spending years trying to prove on paper that our architecture works we throw some experimental API stuff to multiple OSS projects (e.g., ODL, ONOS) and let implementers provide feedback on what works and what doesn't with a *fast fail* approach. Our members want ONF as a neutral third party to define the basic software artifacts (e.g., Information model, principles of operation) that get implemented on many vendor's solutions."

When asked about the progress that the NBI working group was making in general and with regards latitudes in particular, Lenrow said that after publishing the white paper the working group realized that the approach that they established in the white paper could only work if they solved the problem of resource sharing among logic at different latitudes. Currently, every application or service that communicates with the SDN controller acts as if it has total control of all of the subtending network elements and they would "step all over each other". According to Lenrow, the way to solve this problem is to have a single resource arbitrator with a single interface to applications and services.

With that new goal in mind, the NBI working group is focused on developing an intent based interface. In contrast to a prescriptive interface, an intent interface focuses on what the application or service needs and not on the commands to change the network. Some of the other key characteristics and advantages of an intent based interface were explained at a recent *Intent Based Network Summit*.

The intent based interface that the working group develops will be the single interface to applications and services. Subtending that interface will be the various NBIs that are supported by open source and vendor-supplied SDN controllers. The key to making all of this happen is to implement a common information model that enables, via extensibility, every possible use case to be represented by a single NBI. According to Lenrow, there have been a variety of vendor sponsored information models that have not found enough critical mass to create an ecosystem network effect. Lenrow expressed his strong belief that the only way this activity could succeed was to have the interface be developed by a diverse group. He elaborated on this by saying that a "Pay to play approach does not make any sense for this activity and that without paying a fee, people are welcome to join conference calls and to comment on drafts for the ONF's project Boulder, where this work is proceeding." Lenrow said that a document describing the operating principles for an intent based SDN system is "Hopefully in its final draft before being reviewed outside of the Boulder project". He added that at the June 2015 Open Networking Summit in Santa Clara, CA that they were able to demonstrate end-to-end service function chaining. This demo was important because some of the virtual functions were in a domain controlled by the ODL-sponsored open source controller and others were in a domain controlled by the ONOS open source controller. The application that created the end-to-end service used the same interface for each domain.

The ONF NBI initiative has the potential to seamlessly interconnect disparate SDN controllers.

The Southbound Interface

One of the best known protocols used to implement the southbound interface between a SDN controller and the network infrastructure is the OpenFlow protocol. While well-known, OpenFlow isn't the only protocol that can be used to implement the southbound interface. Other options include:

- Border Gateway Protocol (BGP);
- NETCONF;
- Extensible Messaging and Presence Protocol (XMPP);
- Open vSwitch Database Management Protocol (OVSDB);
- MPLS Transport Profile (MPLS-TP).

The Survey Respondents were asked to indicate the likely role that the OpenFlow protocol will play in their company's implementation of SDN. Their responses are shown **Table 4**.

Table 4: Likely Use of OpenFlow

Use of OpenFlow	Percentage
Our implementation of SDN will definitely include OpenFlow	21%
Our implementation of SDN will likely include OpenFlow	25%
Our implementation of SDN might include OpenFlow	22%
Our implementation of SDN will not include OpenFlow	5%
Don't know	25%
Other	2%

One of the conclusions that can be drawn from the data in **Table 4** is that IT organizations have maintained a somewhat favorable view of OpenFlow. In addition:

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

The Overlay and the Underlay Model

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; e.g., supporting the movement of virtual resources or micro-segmentation. 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.

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. 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 5**.

Table 5: The Perceived Value of the Overlay and Underlay-based Models

Response	Percentage
The overlay-based model will provide notably more value	27%
The fabric-based model will provide notably more value	22%
Each model will offer roughly equal value	13%
We don't have an opinion on either model	32%
Other	5%

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

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.

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 correct sequence of VMs.

Open Source SDN Controllers

As previously mentioned, open source is playing a large role in the evolution of SDN and NFV. This section of The Guide will discuss two open source SDN controllers. As a reference, a discussion of the functionality that an OpenFlow-based SDN controller should support is found in *Ten Things to Look for in an SDN Controller*. In addition, a detailed analysis of a number of controllers can be found in *SDN Controllers Report*.

The ON.Lab

The ON.Lab is a non-profit organization founded by people from Stanford University and UC Berkeley. ON.Lab's mission is to "bring openness and innovation to the Internet and Cloud for the public good". ON.Lab believes this mission is best achieved by pursuing the following goals:

- Build tools and platforms that enable and accelerate SDN and make them available through open source;
- Educate the public on the benefits of SDN;
- Provide thought leadership to ensure continued innovation around SDN for the benefit of the public.

One of the ON.Lab's primary projects is ONOS (Open Network Operating System) – an open source SDN operating system. The ONOS partnership goals are to:

- Build an open source SDN operating system for service providers;
- Build open source SDN and NFV solutions;
- Enable vendors to create value with open source and white boxes;
- Create a vibrant and sustainable community.

At least for now, ONOS is targeted at service providers.

The ON.Lab released the first version of ONOS, code named Avocet, in December 2014. ON.Lab sees clear differences between what it is trying to accomplish and what ODL is trying to accomplish. For example, in a November 2014 article in Network World, ON.Lab officials were quoted as saying that ODL was a vendor driven activity that is intended to preserve the incumbency of brand name hardware. Guru Parulkar, ON.Lab's executive director was quoted in that article as saying that "ODL is focused on automation of the command line interface used to configure legacy hardware and does not bring 'SDN value' to service providers, such as lower operation expenditures, speeding service delivery and revenue, and offering white box alternatives."

Current members of the ONOS community include collaborators such as the ONF; vendors such as Cisco, Huawei and NEC; and service providers such as AT&T, NTT Communications and SK Telecom. In September 2015 ON.Lab released the fourth version of ONOS, code named Drake. According to ON.Lab "Drake adds new security, configuration and application level feature sets with improvements to the northbound and southbound including REST, API and GUI additions and upgrades throughout. In addition to contributing to ONF's Atrium, ONOS has expanded collaboration with other open source communities to develop new distributions including work with the CloudRouter® Project and it will soon be part of the Open Platform for NFV Project (OPNFV)."

The ONOS community has expanded to include vendors.

In addition, in October 2015 a partnership was announced between the ONOS project and the Linux Foundation. In a conversation with the author, Parulkar said that it was important initially to have ON.Lab control the development of ONOS, but that now is the time to bring in the broader development community to rapidly expand ONOS's capabilities. As part of that conversation, Jim Zemlin, executive director of the Linux Foundation stated that he saw this partnership as one more proof point that 2016 would be the year of open source in the networking sector. He also stated that the mechanism are in place so that over time the control of ONOS will move away from ON.Lab and to the open source community.

The ONOS project is part of the Linux Foundation.

The OpenDaylight (ODL) Project

The ODL Project, which was founded in April 2013, is a collaborative open source project hosted by The Linux Foundation. The goal of the project is to facilitate a community-led, industry-supported open source framework, including code and architecture, to accelerate and advance a common, robust SDN platform and to create a solid foundation for NFV. Towards that end, the ODL project claims that a number of vendors use ODL code as the basis of their SDN products and that its code is also used by the OPNFV platform, which is described in a subsequent section of The Guide. As of September 2015 the consortium had 50 members: 8 platinum members, 1 gold member and 41 silver members.

In June 2015 the consortium announced the availability of its third software release, called Lithium. Some of the new functionality in Lithium includes:

- Application-Layer Traffic Optimization (ALTO)
  ALTO is an IETF protocol (RFC7285) to provide network information to applications.

- Control And Provisioning of Wireless Access Points (CAPWAP) Protocol
  The protocol, which is described in RFC 5415, enables a central wireless LAN Access Controller to manage a collection of Wireless Access Points.

- Link Aggregation Control Protocol (LACP)
  This capability auto-discovers and aggregates multiple links between an OpenDaylight controlled network and LACP-enabled endpoints or switches.

- Network Intent Composition (NIC)
  This is an interface that allows clients to express a desired state in an implementation-neutral form.

- Opflex Agent
  This is a policy agent that works with OVS to enforce a group-based policy networking model with locally attached virtual machines or containers.

- Reservation
  This capability provides dynamic low level resource reservation so that users can get network as a service, connectivity or a pool of resources for a period of time.

- SNMP
  This is a southbound plugin that allows applications and controller services to interact with devices using SNMP.

- Unified Secure Channel (USC) framework
  This framework provides a central server to coordinate encrypted communications between endpoints.

ODL's Lithium release contains a range of sophisticated functionality.

As mentioned, one of the criticisms of the ODL project is that it is run by vendors who will advocate for proprietary solutions. In an interview with the author, Neela Jacques, the executive director of the ODL project refuted that criticism saying that belonging to an organization such as ODL requires a commitment of resources and so it shouldn't be surprising that the first wave of companies to join ODL were network vendors. Jacques stated that the second wave of companies to join ODL were service providers such as AT&T and Comcast. He pointed to the fact that companies such as NASDAQ and Credit Suisse have joined ODL's board of advisors as proof that a third wave of companies, enterprise organizations, are currently joining ODL.

ODL's membership has expanded to include service providers and enterprises.

When asked about use cases, Jacques said that the use cases of most interest to ODL members are:

- Centralized network management;
- Being able to program the network;
- The ability to implement virtual networks;
- The ability to leverage OpenStack for orchestration and to support NFV use cases such as virtual Customer Premise Equipment (vCPE).

When asked about ONOS, Jacques said that the ON.lab is taking a very different approach to developing a SDN controller than is ODL. According to Jacques, ON.lab is focusing on a few well defined carrier-specific use cases. Jacques stated that ON.lab has already added value and he expressed his belief that they would continue to add value. However, in contrast to ON.lab, Jacques said that part of ODL's goal is to unite the world, not around any one customer, but around a common code base. He added that another part of ODL's goal is to be a place where multiple ideas thrive and incubate and where multiple technologies come together over time.

Why Choose Application Centric Infrastructure (ACI)?

Application Deployment at the Speed of Business

Without ACI

New Application

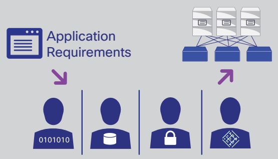

IT administrators work in silos.

Weeks

There is no shared architectural model.

With ACI

New Application

Simpler

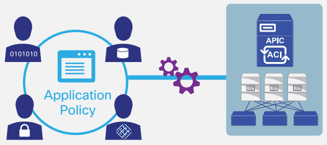

Faster

IT admins work at the application level.

Hours

Efficient

Shared model for policy automation.

ACI cuts deployment time and effort.

Optimal Design

Application Policy

- Connectivity
- Reliability
- Compliance
- Performance
- L4-L7 Services
- Utilization

Faster Deployment

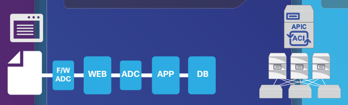

Application Network Profile

- Test
- Certify
- Adapt
- Provision

Simplified Operations

Application Topology

- Monitor
- Troubleshoot
- Optimize
- Scale Out

What does ACI deliver?

- Automation and Visibility
- Performance and Scale
- Security
- Openness

Redefine the Power of IT with ACI

Learn more at www.cisco.com/go/aci

The Use Cases and Business Case for SDN

SDN Use Cases, Drivers and Inhibitors

Focus of SDN Deployment

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 6**.

Table 6: Focus of SDN Deployment

Focus of SDN Deployment	Percentage
Data Center	51%
WAN	31%
Branch and/or Campus	22%
We are likely to implement a service from a WAN service provider that is based on SDN	20%
Don’t know/NA	10%
We are unlikely to implement SDN within the next two years	10%
Other	4%

One observation that can be made from the data in **Table 6** is:

There is currently as much interest in either implementing SDN in the WAN or using a SDN-based WAN service as there is in implementing SDN in the data center.

Below is a discussion of the key use cases for SDN in the data center, the WAN and the campus. In some cases, the distinctions are somewhat arbitrary as some of the use cases that are listed as being appropriate in the data center are also appropriate in the branch and campus and vice versa.

Data Center

Data Center Use Cases

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 a L4 – L7 service such as WAN optimization. The phrase *service chaining* refers to the ability to dynamically steer traffic flows through a sequence of physical or virtual servers that provide 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 is capable of having the switch redirect suspicious traffic flows to higher-layer security controls, such as IDS/IPS systems, application firewalls, and Data Loss Prevention (DLP) devices. Other security applications that run on top of an 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 a 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's 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.

Drivers and Inhibitors of SDN in the Data Center

Table 7 and **Table 8** contain the responses of The Survey Respondents when asked to indicate the two factors that would drive and the two factors that would inhibit their organization's implementation of SDN in a data center.

Table 7: Drivers of Implementing SDN in a Data Center

Challenge or Opportunity	Percentage
Support the dynamic movement, replication and allocation of virtual resources	39%
Ease the administrative burden of configuration and provisioning	36%
Better utilize network resources	22%
Perform traffic engineering with an end-to-end view of the network	18%
More easily scale network functionality	16%
Reduce OPEX	14%
Have network functionality evolve more rapidly based on a software development lifecycle	12%
Reduce CAPEX	11%
Enable applications to dynamically request services from the network	11%
Implement more effective security functionality	9%
Reduce complexity	7%
More easily implement QoS	4%

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

The two primary factors driving SDN deployment in the data center are supporting the dynamic movement, replication and allocation of virtual resources and easing the administrative burden of configuration and provisioning.

Table 8: Inhibitors to the Adoption of SDN in a Data Center

Impediment	Percentage
Concerns about how we would integrate SDN into the rest of our infrastructure	30%
The immaturity of the enabling technologies	28%
The confusion and lack of definition in terms of vendors strategies	23%
Other technology and/or business priorities	20%
The lack of a compelling business case	17%
Possible security vulnerabilities	15%
Concerns about how we would manage SDN	12%
The lack of a critical mass of organizations that have deployed SDN	8%
No inhibitors to implementing SDN	7%
Concerns that the technology will not scale to support enterprise sized networks	6%
Other	4%

Unlike the situation shown in **Table 7** in which there were two clear drivers of SDN deployment in the data center:

There is a wide range of significant inhibitors to the deployment of SDN in the data center.

WAN

WAN Use Cases

As described below, one of the first production implementations of SDN was Google's implementation of their G-Scale WAN. As is also described below, there is currently significant interest in taking a SDN approach to the WAN. This approach is often referred to as a Software-Defined WAN (SD-WAN). SD-WANs are discussed in detail in *The 2015 State-of-the-WAN Report* and *The 2015 Guide to WAN Architecture and Design*.

The Google G-Scale WAN

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 Google's global data centers. G-Scale is a prime example of a production OpenFlow Layer 3 network that is realizing the benefits of FIB centralization. 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%.

SD-WANs

As is the case with any SDN, a SD-WAN centralizes the control function into a SDN controller. The controller abstracts the user's private network services from the underlying IP network and it enables the operations of the user's private network services via centralized policy. The controller also enables the automation of management tasks such as configuration and provisioning.

Leveraging the underlying WAN platforms, which may include physical or virtual routers, the controller sets up virtual overlays that are both transport and technology agnostic. Under the direction of the controller, the WAN platforms implement functionality such as quality of service, path selection, optimization and security, often using dynamic multi-pathing over multiple WAN links.

Drivers and Inhibitors of SDN in the WAN

Table 9 and **Table 10** contain the responses of The Survey Respondents when asked to indicate the two factors that would drive and the two factors that would inhibit their organization's implementation of SDN in the WAN.

Table 9: Drivers of Implementing SDN in a WAN

Challenge or Opportunity	Percentage
Ease the administrative burden of configuration and provisioning	33%
Better utilize network resources	30%
Perform traffic engineering with an end-to-end view of the network	23%
More easily scale network functionality	22%
Support the dynamic movement, replication and allocation of virtual resources	22%
Reduce OPEX	15%
More easily implement QoS	12%
Enable applications to dynamically request services from the network	11%
Reduce CAPEX	10%
Have network functionality evolve more rapidly based on a software development lifecycle	7%
Implement more effective security functionality	6%
Reduce complexity	6%

Some of the conclusions that can be drawn from the data in **Table 9** are:

There are a number of significant drivers of SDN deployment in the WAN.

The two primary factors driving SDN deployment in the WAN are easing the administrative burden of configuration and provisioning and better utilizing network resources.

Table 10: Inhibitors to the Adoption of SDN in the WAN

Impediment	Percentage
Concerns about how we would integrate SDN into the rest of our infrastructure	25%
The lack of a compelling business case	25%
The immaturity of the enabling technologies	25%
The immaturity of the current products	22%
Possible security vulnerabilities	22%
The confusion and lack of definition in terms of vendors strategies	17%
Other technology and/or business priorities	16%
The lack of a critical mass of organizations that have deployed SDN	11%
Concerns about how we would manage SDN	9%
Concerns that the technology will not scale to support enterprise sized networks	9%
No inhibitors to implementing SDN	7%
Other	3%

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.

Three of the major inhibitors to the deployment of SDN in the WAN are concerns about how to integrate SDN into the rest of the infrastructure, the lack of a compelling business case and concerns about security vulnerabilities.

Branch and Campus

Branch and Campus Use Cases

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.

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.

Drivers and Inhibitors of SDN in Branch and Campus Networks

Table 11 and **Table 12** contain the responses of The Survey Respondents when asked to indicate the two factors that would drive and the two factors that would inhibit their organization's implementation of SDN in branch and campus networks.

Table 11: Drivers of Implementing SDN in Branch and Campus Networks

Challenge or Opportunity	Percentage
Ease the administrative burden of configuration and provisioning	37%
Better utilize network resources	25%
More easily scale network functionality	20%
Support the dynamic movement, replication and allocation of virtual resources	18%
Reduce OPEX	17%
Implement more effective security functionality	15%
Perform traffic engineering with an end-to-end view of the network	12%
More easily implement QoS	12%
Enable applications to dynamically request services from the network	11%
Reduce CAPEX	11%
Reduce complexity	11%
Have network functionality evolve more rapidly based on a software development lifecycle	9%
Other	5%

Observations that can be drawn from **Table 11** include:

The primary driver of implementing SDN in branch and campus networks is easing the burden of configuration and provisioning.

While the drivers of implementing SDN in branch and campus networks are similar to the drivers of implementing SDN in the data center, in some cases the relative importance is significantly different.

Table 12: Inhibitors to the Adoption of SDN in Branch and Campus Networks

Impediment	Percentage
Concerns about how we would integrate SDN into the rest of our infrastructure	28%
The lack of a compelling business case	24%
Possible security vulnerabilities	23%
The immaturity of the current products	22%
Other technology and/or business priorities	21%
The immaturity of the enabling technologies	18%
The confusion and lack of definition in terms of vendors strategies	12%
Concerns about how we would manage SDN	11%
The lack of a critical mass of organizations that have deployed SDN	11%
Concerns that the technology will not scale to support enterprise sized networks	10%
No inhibitors to implementing SDN	7%
Other	5%

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

Two of the major inhibitors to the deployment of SDN in branch and campus networks are concerns about how to integrate SDN into the rest of the infrastructure and the lack of a compelling business case.

Taking a holistic view of the factors that are impacting SDN deployment:

Overall, the two primary factors that are driving the implementation of SDN are easing the administrative burden of configuration and provisioning and increasing the utilization of network resources.

However,

Neither reducing complexity nor reducing CAPEX are significant drivers of deploying SDN.

The primary factor inhibiting the adoption of SDN is the concerns that organizations have about how they would integrate SDN into the rest of the infrastructure.

The SDN Business Case

The methodology to develop a business case for an investment in SDN will vary by company. However, it is generally easier to build a business case for an investment in SDN if that investment results in hard dollar savings. For example, there is the potential that the investment it takes to deploy a SD-WAN will result in significant hard dollar savings due to replacing relatively expensive MPLS bandwidth with relatively inexpensive Internet bandwidth.

As described below, one of the potential components of an SDN business case is that implementing SDN makes rolling out new business services easier and faster. In most instances, if a business-related benefit such as that is being used to justify implementing SDN, it would be beneficial to get an appropriate business leader to identify, and where possible quantify the business value of that benefit.

Financial metrics

There are numerous metrics that can be used to measure the financial viability of deploying any kind of technology. One of the most useful metrics is the payback period, which is the amount of time before the resultant savings equals or exceeds the cost of deploying a new technology or service. To demonstrate payback period, assume that a company invests $1,000,000 in SDN equipment in order to implement a SD-WAN and further assume that the SD-WAN saves the company $100,000 a month. In that case, the payback period is ten months.

Another useful financial metric is the internal rate of return (IRR). The IRR of an investment or project is the "annualized effective compounded return rate" that makes the net present value of all cash flows from a particular investment equal to zero.

The Components of a Business Case

WAN Savings
As mentioned, implementing a SD-WAN has the potential to reduce the amount of money that a company spends with communications service providers. Below are two examples of how this savings can be realized:

- Cost reduction
  In this case, a company removes some or all of its MPLS circuits and replaces these circuits with Internet connectivity.

- Cost avoidance
  In this case, a company decides that instead of adding MPLS circuits, that it will add Internet connectivity.

Operational Efficiencies
As highlighted by the preceding survey results, one of the primary advantages of a SDN is that it reduces the cost and time associated with tasks such as configuration and provisioning by centralizing control and allowing network organizations to configure and provision hundreds of devices as if they were one device.

Consolidation of Resources

By virtualizing and pooling compute, storage and network resources, IT organizations can significantly reduce the number and the cost of the required physical resources. However, as described below, a SDN is required both to implement efficient network virtualization and to experience all of the potential benefits that result from compute and storage virtualization.

IT Agility

One of the key characteristics of a SDN is that is supports virtual networks which are decoupled from the physical networks. These virtual networks enable VMs to be dynamically moved between physical servers with no manual intervention. Being able to dynamically move VMs results in considerable operational savings and it makes the IT organization more agile.

Another way that SDN increases the agility of the IT organizations comes from being able to guarantee complete isolation of each user of the SDN. Because of this isolation, an IT organization can allow application developers to run their applications in a production environment without impacting production traffic. This is particularly important for an IT organization that either already has, or soon will embrace DevOps.

Business Agility

In a SDN, network functions such as optimization and security can be coordinated at a policy level with the SDN controller handling all of the details needed to implement those policies across multi-device, multi-platform infrastructure. This enables the IT organization to support new business services notably faster than in a traditional environment in which each device has to be procured and manually configured.

Improved Application Performance

One of the primary characteristics of a SDN is that there are programmatic interfaces into the SDN controller. These interfaces make the control information that has been centralized in the controller available to a potentially unbounded set of SDN applications. These applications are capable of dynamically changing the underlying network to perform tasks such as forwarding packets over the least expensive path or improving application performance by changing the QoS settings based on the available bandwidth or other factors.

Increased Network Availability and Performance

There are a number of ways that a SDN can result in increased availability. For example, one of the many advantages of decoupling the virtual networks from the physical networks is that it enables IT organizations to make changes to the physical network, such as scaling out capacity, without impacting the existing flows or having to take the network out of service.

Another way that a SDN can result in increased availability is relative to how traffic is routed. In a traditional network there is a single data path from origin to destination. If that path becomes unavailable, there is an outage until a new path is determined. A key feature of an SDN controller is its ability to discover multiple paths from the origin of the flow to its destination and to split the traffic for a given flow across multiple links. In normal operating conditions, this capability of SDN increases both the performance and scalability of the solution. In the case of an outage, this capability increases availability because there will still be at least one active path from origin to destination.

Improved Security

There are a number of ways that a SDN can result in improved security. For example, in order to respond to myriad industry and government regulations about data security, IT organizations often need to keep the data generated by one set of users isolated from other users. This can be accomplished by adopting a SDN that provides virtual networks that are fully isolated from one another. In addition, as previously discussed, OpenFlow access switches can filter packets as they enter the network and act as simple firewalls at the edge. OpenFlow switches can also redirect certain suspicious traffic flows to higher-layer security controls, such as IDS/IPS systems, application firewalls, and Data Loss Prevention (DLP) devices.

Enhanced Management and Visibility

As was previously discussed, a SDN dramatically simplifies tasks such as configuration management. A SDN can also help with application performance management. For example, in the majority of instances in which the performance of an application is degrading, the degradation is noticed first by the end user and not by the IT organization. One of the principal reasons why IT organizations are often unaware of degraded application performance is that in the traditional IT environment, IT organizations lack visibility into the end-to-end network flows. One of the key advantages of a SDN is that it enables IT organizations to have end-to-end network flow visibility.

Masergy's Software Defined Platform Delivers the Flexibility Enterprises Require

The enterprise WAN is fast becoming the source of serious innovation. Consider it the central nervous system of corporations and their ability to support even the most advanced and demanding business-critical applications. A modern network should be agile enough to adapt to your rapidly changing business needs.

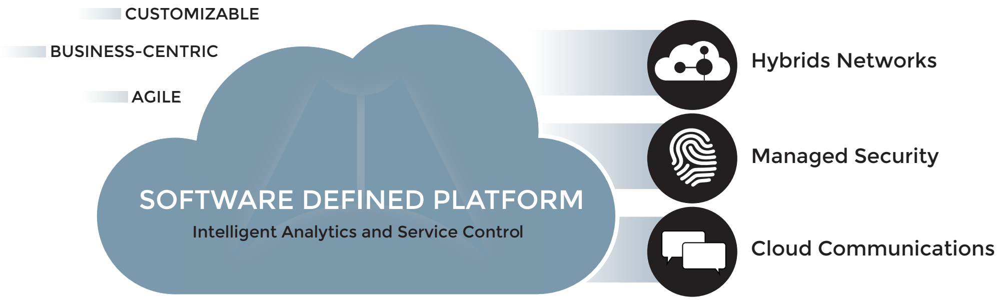

This WAN innovation is being enabled by a Software Defined Networking (SDN) architecture, which enables administrators to rapidly change configurations as performance requirements demand.

Putting theory into practice, Masergy's Software Defined Platform accelerates IT transformation by providing the foundation for an open, automated and programmable network fabric. Our Software Defined Platform is the foundation for our three essential solutions: hybrid networking, managed security and cloud communications.

SDN enables us to build intelligent analytics, automation and service control into all of our solutions.

Here are some of the many benefits of our Software Defined Platform:

- A high-availability, highly resilient hybrid network environment that provides an optimal client and application experience
- A secure, hybrid WAN architecture that permits dynamic traffic engineering across both private and public domains, delivering seamless performance and consistent manageability
- Administrators are afforded full visibility into business-critical applications and the ability to prioritize traffic based on performance, security and business policy needs
- Remote sites and branch offices can be added to the corporate network quickly and with little-to-no on-site administration

SOFTWARE-DEFINED SOLUTIONS THAT ADAPT TO THE DEMANDS OF YOUR BUSINESS

VIRTUALIZATION

Network Function Virtualization (NFV) as a central tenant of our Software Defined Platform. NFV is poised to transform the world of networking as part of a larger shift from rigid, legacy networks where hardware and software are proprietary and tightly integrated, to modern networks that are software-driven and programmable. This gives network architects and administrators a new way to design, deploy and manage network capabilities.

Masergy has implemented NFV in its recently introduced Managed Network Functions f(n). We offer a family of fully managed, distributed network functions that can be delivered in the way that best suits your needs, whether that's on premises, in the cloud, or virtualized via software.

The solution offers Virtual Functions f(n) that lets companies add routing and firewall capabilities in software on their existing Masergy network interface device, eliminating the need for proprietary network appliances and on-site administration.

Our Premise Function f(n), is a complete lifecycle management solution for enterprises, which covers essential on-premises networking functions, including routers, firewalls and session-border controllers. And a third component, Cloud Functions f(n), help companies deliver essential network functions as cloud services over the Masergy network.

> *"Masergy was able to custom design our hybrid network to meet our unique application performance requirements. It's outcome-based approach and ongoing superior support have convinced us we selected the right partner for our needs."*
>
> Patrick Tisdale, CIO – McKenna, Long & Aldridge, LLP

FLEXIBLE BY DESIGN

SDN transforms enterprise networks into modular, scalable assets that can be assembled and rearranged as business needs require. It also reduces IT complexity through automation.

Masergy is helping customers accelerate their IT transformation efforts, providing the foundation for an open, automated and programmable environment. This, in turn, frees up IT staff to focus on strategic, business-driven innovations and less time "keeping the lights on."

MASERGY
Performance Beyond Expecations

For more information, please visit https://www.masergy.com

Compared with legacy approaches, Masergy's NFV offers three compelling advantages:

1. **Extended Flexibility:** Masergy's Managed Functions f(n) gives enterprises complete control over their distributed network resources–plus the ability to scale up new services and decommission outmoded network capabilities as business needs change.
2. **Rapid Deployment:** NFV lets organizations add, remove, configure and modify network services in real time. Rather than ordering and shipping network appliances to branch offices and remote locations, it lets you take advantage of innovative new services and deploy them over your network via software updates.
3. **Lower Costs:** Our managed Network Functions f(n) reduces CAPEX because an organization no longer needs to purchase specialized hardware in many instances. And OPEX drops because of the way NFV lowers the need for dedicated hardware, support personnel and equipment maintenance.

Contact us for a free consultation.

Corporate Headquarters (USA):
2740 North Dallas Parkway, Suite 260
Plano, TX 75093 USA
Phone: +1 (214) 442-5700
Fax: +1 (214) 442-5756

European Headquarters (UK):
29 Finsbury Circus
Salisbury House 5th Floor
London, EC2M 5QQ UK
Phone: +44 (0) 207 173 6900
Fax: +44 (0) 207 173 6899

The Operational Impediments to Implementing SDN

The Operational Implications

One of the implications of adopting SDN is that is increases the need for a DevOps model.

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

Security

Background

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. Another security related use case is to leverage 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 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:

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

SDN creates security opportunities and security challenges.

Vendor Questions

Below are some of the questions that network organizations should ask vendors relative to the security of their SDN solution.

- What functionality does your solution support in order to ensure the security of end-to-end communications?
- How are the components of your solution designed for security? For example, what steps have been taken to harden the SDN controller's operating system?
- What functionality does your solution support in order to ensure the security of communications between the components of your solution?
- What capability does your solution have to detect security breaches?
- How does your solution logically separate traffic?
- What measures are available to deal with possible control flow saturation (controller DDOS) attacks?
- Describe any SDN-based solutions that are available both to detect the communications patterns of spurious traffic (e.g., botnets, spam, and spyware) from internal end systems and to block or quarantine the source.
- How does your solution make implementing security notably less complex than the traditional ways of implementing security?
- What tests have been run to verify the effectiveness of the security measures that have been taken? Is it possible to see those test results?

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 2 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 2: OpenStack

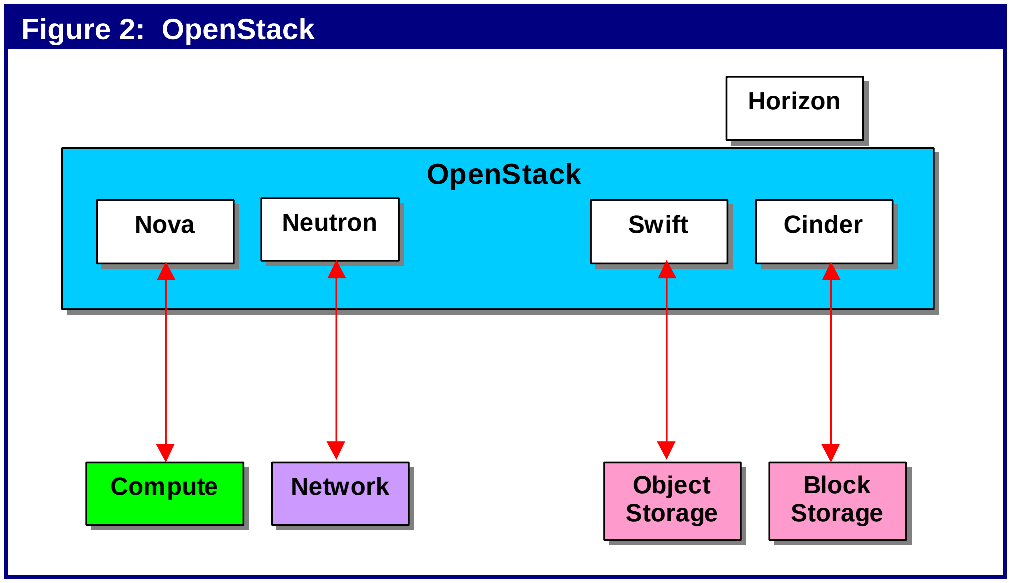

Horizon is the OpenStack Dashboard that provides administrators and users a graphical interface to access, provision and automate cloud-based resources. The dashboard is one of several ways users can interact with OpenStack resources. Developers can automate access or build tools to manage resources 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 (VPNs) to be deployed and managed.

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 3 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 3: VM Creation Workflow with OpenStack

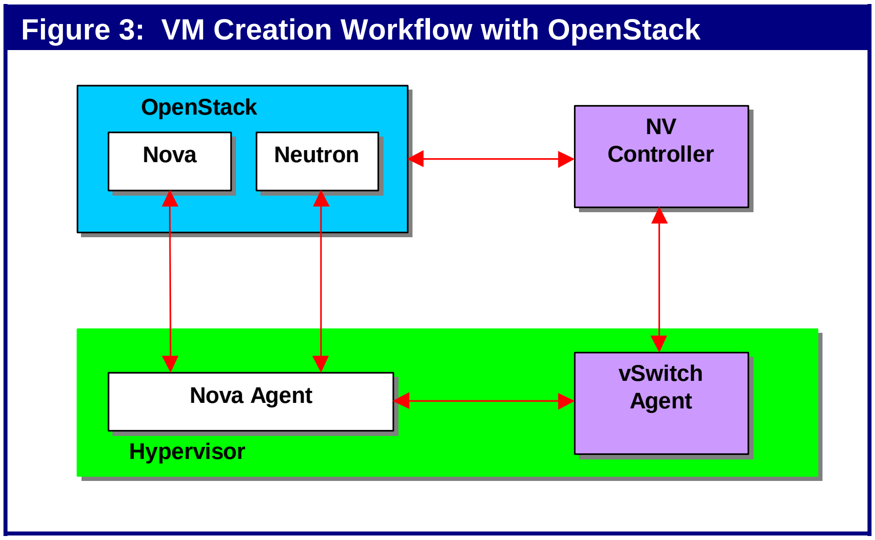

The **Nova** compute module in OpenStack instructs the Nova Agent in the hypervisor to create the VM. The Nova agent communicates with the Neutron module in OpenStack to learn the network attributes of the VM. The Nova agent then informs the vSwitch agent to configure the virtual network for the VM and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

Management

Background

As described in a preceding section of The Guide, one of the two primary factors driving the deployment of SDN is the belief on the part of network organizations that implementing SDN will ease the burden of configuration and provisioning. However, as described below, the adoption of SDN also creates management challenges. This leads to the conclusion that:

SDN creates both management opportunities and management challenges.

A related conclusion 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. 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.

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. A set of management challenges that occurs at the application layer stem from the requirement to ensure acceptable performance. 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 if that is appropriate.

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 network resources.

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 a result:

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.

Vendor Questions

Below are some of the questions that network organizations should ask vendors relative to the management of their SDN solution.

- How extensive is the management functionality that you provide. Include in your answer:
  - The type of management data that is gathered and where it is gathered;
  - Where the management data is stored and where it is processed;
  - How your solution performs monitoring of network and application performance;
  - The level of visibility that your solution provides into network and application performance;
  - The ability of your solution to enable rapid root cause analysis;
  - The level of visibility that your solution provides into the performance of applications and services acquired from a cloud provider;
  - The type and the extent of analytics that are part of your solution.

- What functionality does your solution provide to enable a company to implement and support SLAs for varying types of applications?
- How does your solution provide event correlation and fault management?
- Describe the integration or potential integration that exists between the management tool that you provide to manage your solution and other common business intelligence, security and management tools, whether provided by your company or by a third party.
- Describe the integration or potential integration of your solution with leading orchestration solutions.
- What type of management interface do you provide into your SDC controller? For example, is it based on REST? On something else?
- Describe the ability of your solution to monitor the SDN controller. Include in that description your ability to monitor functionality such as CPU utilization as well as flow throughput and latency. Also describe the statistics you collect on ports, queues, groups and meters; and the error types, codes and descriptors you report on.
- What type of management interface do you provide into your management tool? For example, is it based on REST? On something else?
- Describe the ability of your solution to monitor the network elements in your solution. Include in that description the key performance metrics that you monitor and report on. Also, can the performance data gathered by SDN switches (e.g., counters and meters) be integrated with data from traditional performance management tools based on SFlow and SNMP?
- Describe how your solution monitors the messages that go between the SDN controller and the SDN switches.
- Describe the ability of your solution to provide visualization of traffic flows and service quality.

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*, the adoption of a 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;
- More focus on setting policy;
- More knowledge of the business;
- More understanding of applications;
- More emphasis on programming.

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). The Survey Respondents were 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 13**.

Table 13: Impact of SDN on Organizational Structure

Impact	Percentage
Very Significant Impact	11%
Significant Impact	19%
Moderate Impact	17%
Some Impact	24%
No Impact	8%
Don't Know	21%

Almost a third of the survey respondents believe that over the next two years the ongoing adoption of software-based IT functionality will have either a significant or very significant impact on the structure of their IT organization.

When asked to indicate the type of organizational changes that would likely occur, the responses included that there would likely be:

- An accelerated transition to highly dynamic and flexible cloud architectures;
- Greater investment in logical architectures, systems design thinking and business virtualization;
- An impact on design and purchasing decisions;
- More focus on business processes;
- A redefinition of roles and responsibilities;
- A reorganization based on IT and DevOps skills.

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 14.**

Table 14: Impact of SDN on Jobs

Impact	Percentage
Very Significant Impact	10%
Significant Impact	19%
Moderate Impact	18%
Some Impact	24%
No Impact	11%
Don't Know	18%

Over a quarter of the survey respondents believe that over the next two years the ongoing adoption of software-based IT functionality will have either a significant or very significant impact on their jobs.

When asked to indicate the type of changes that would likely occur to their jobs, the responses included:

- We will spend less time configuring and more time planning;
- Our roles will blend and create some conflicts;
- How we design, deploy and manage networks will change;
- We will need to be re-trained on the skills necessary to support SDE;
- We will need to absorb new skills and evaluate a broader range of vendors;
- The skills needed will change from networking to programming and scripting.

NETSCOUT®

Extending Service Assurance into SDN and NFV Environments

SOLUTION BENEFITS

NETSCOUT's Adaptive Service Intelligence™ (ASI) technology empowers enterprises and service providers to fully realize the benefits of SDN and NFV CapEx and OpEx efficiencies by reducing deployment risk.

- Accelerates migration to virtualized infrastructures with confidence.
- Provides service visibility without compromising user and customer experience.
- Protects and enhances performance of traditional, non-SDN/NFV, deployments.

Solution Core Functionality

NETSCOUT's nGeniusONE™ Service Assurance platform and ASI technology deliver real-time, actionable traffic-based intelligence capabilities.

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

Challenges

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 lower cost. 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 SDN and NFV CapEx and OpEx efficiencies, enterprises and service providers need a comprehensive service delivery monitoring capability which offers end-to-end visibility across physical, virtual, and hybrid environments. To be truly beneficial, the tool needs to offer rapid service triage capabilities to reduce the mean time to resolution (MTTR), 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 IT and Operations teams, and lacks the ability to provide an end-to-end 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 disparate 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 significantly increased mean time to resolution, drastically extended service outages, reduced quality of end-user experience or loss in worker productivity.

Solution Overview

NETSCOUT® offers rapid 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 using one cohesive, consistent set of metadata, based on packet flow data, for service provider and enterprise services. This metadata is generated by the patented Adaptive Service Intelligence technology running on NETSCOUT's physical and virtual 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 points across the service delivery infrastructure using physical and virtual appliances. The packet flow data collection and aggregation is passive and noninstrusive and can scale to collect any required volumes of data across physical, virtual, and hybrid environments.

The nGeniusONE Service Assurance platform aggregates, correlates, and contextually analyzes the metadata gathered from NETSCOUT's physical and virtual Intelligent Data Sources. It then creates real-time holistic views of service performance, establishes performance baselines, and facilitates service-oriented troubleshooting workflows.

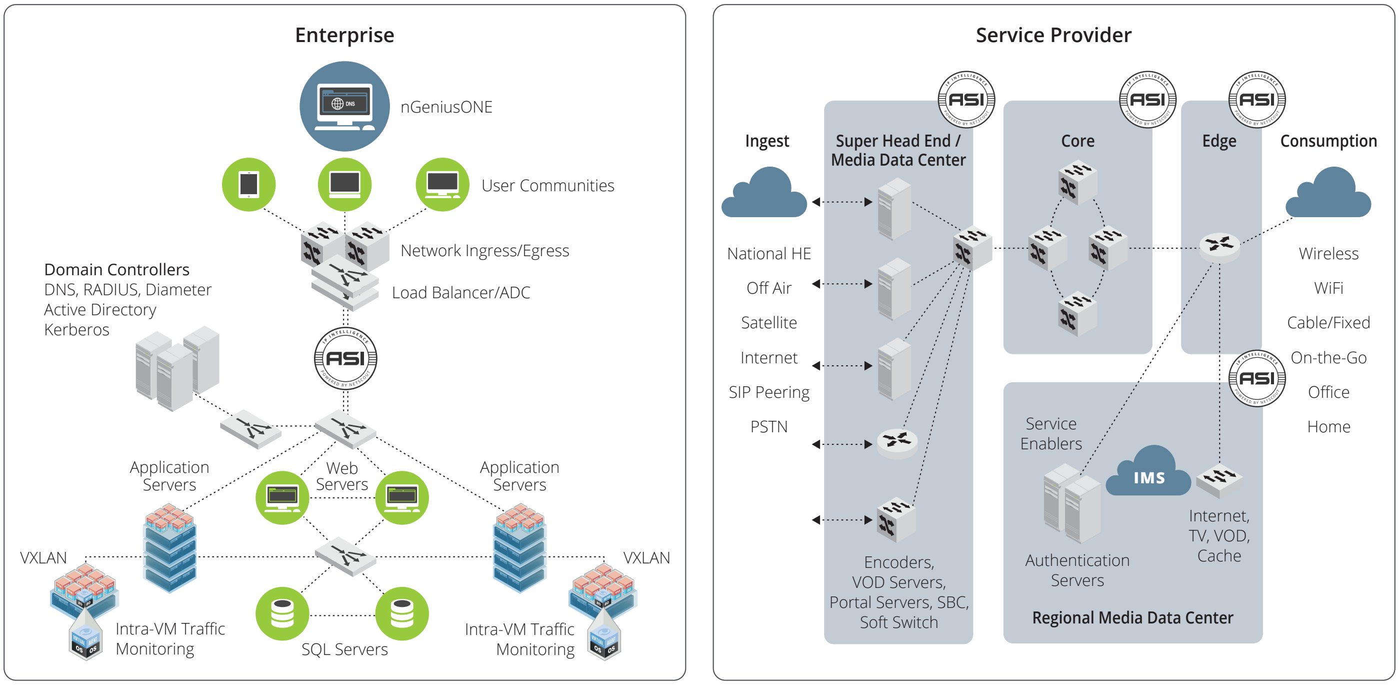

Figure 1: Service Instrumentation in Enterprise and Service Provider Environments.

Core Technologies

NETSCOUT's unique ability to gain a pervasive end-to-end visibility into the service delivery environment, and enable rapid service triage is centered on NETSCOUT's ASI technology, utilizing packet flow data, and providing scalable packet flow access.

Adaptive Service Intelligence (ASI)

Adaptive Service Intelligence is patented technology which uses rich packet-flow data to generate highly scalable metadata that enables a comprehensive real-time and historic view of service, network, application, and server performance. This powerful packet inspection and data mining engine runs on NETSCOUT's Intelligent Data Sources generating metadata based upon actual session traffic in real time as the packets cross physical or virtual links. NETSCOUT's ASI technology is the foundation of a highly scalable service delivery monitoring architecture which seamlessly collects, normalizes, correlates, and contextually analyzes data for all services: voice, data, and video.

Leverage 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 Packet Flow Switches and TAPs provide the foundation for a scalable monitoring architecture needed for service assurance. NETSCOUT's nGenius Packet Flow Switches (PFS) filter, aggregate, and distribute the targeted data to NETSCOUT's Intelligent Data Sources in a transparent, selective, and efficient manner. 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.

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 integrating NETSCOUT's ASI technology into a virtual machine (VM), functioning as a virtual Intelligent Data Source. NETSCOUT's VM either analyzes the intra-VM traffic in a self-contained virtualized mode or redirects the traffic to an external NETSCOUT Intelligent Data Source for analysis.
- **Traffic between VMs that reside in different hypervisors** is monitored by NETSCOUT 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 Comparison

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.
Rapid Service Triage	Reactive and time-consuming triage result in poor user experience, and extended service downtime impacting multiple users.	Rapid service triage helps resolve service degradation in real time before large numbers 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 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 return on investment in performance management by gradually expanding the solution over time.

About NETSCOUT Systems, Inc.

NETSCOUT Systems, Inc. (NASDAQ:NTCT) is a market leader in real-time service assurance and cybersecurity solutions for today's most demanding service provider, enterprise and government networks. NETSCOUT's Adaptive Service Intelligence (ASI) technology continuously monitors the service delivery environment to identify performance issues and provides insight into network-based security threats, helping teams to quickly resolve issues that can cause business disruptions or impact user experience. NETSCOUT delivers unmatched service visibility and protects the digital infrastructure that supports our connected world. To learn more, 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

Network Functions Virtualization (NFV): A Status Update

The Relevance of NFV to Enterprise Organizations

The conventional wisdom has been that Network Functions Virtualization (NFV) is associated exclusively with Communications Service Providers (CSPs). Part of the reason for that is the key role that the European Telecommunications Standards Institute (ETSI) has played in the development of NFV. For example, roughly three years ago an Industry Specifications Group (ISG) for Network Functions Virtualization (NFV ISG) was formed under the auspices of the European Telecommunications Standards Institute (ETSI NFV ISG). While the membership has evolved significantly, the initial members of the ETSI NFV ISG were all CSPs such as AT&T, Deutsche Telekom and NTT.

Table 15 contains examples of functions that the ETSI NFV ISG believes can be virtualized.

Table 15: Potential Functions to be Virtualized

Network Element	Function
Switching elements	Broadband network gateways, carrier grade Network Address Translation (NAT), routers
Mobile network nodes	Home Location Register/Home Subscriber Server, gateway, GPRS support node, radio network controller, various node B functions
Customer premise equipment	Home routers, set-top boxes
Tunneling gateway elements	IPSec/SSL virtual private network gateways
Traffic analysis	Deep packet inspection (DPI), quality of experience measurement
Assurance	Service assurance, service level agreement (SLA) monitoring, testing and diagnostics
Signaling	Session border controllers, IP Multimedia Subsystem components
Control plane/access functions	AAA servers, policy control and charging platforms
Application optimization	Content delivery networks, cache servers, load balancers, accelerators
Security	Firewalls, virus scanners, intrusion detection systems, spam protection

Given the leadership role that ETSI is playing combined with the interest that they have in virtualizing functionality such as broadband network gateways and radio network controllers which has no general applicability in enterprise networks, it is easy to see why some people associate NFV strictly with CSPs.

However, while CSPs typically have a broader range of functionality that they are interested in virtualizing than do enterprises, enterprise IT organizations have been implementing virtualized functionality for several years; e.g., virtualized WAN optimization controllers and virtualized Application Delivery Controllers. While ETSI champions the interest that CSPs have with virtualizing network functions, the Open Networking User Group (ONUG) is one of the organizations that has emerged to champion the corresponding interest that enterprises have. ONUG was founded in 2012 and unlike ETSI its members are primarily enterprise companies such as Fidelity Investments, Citigroup and FedEx. In a white paper entitled *Open Networking Challenges and Opportunities*, the group discussed the cost and complexity of managing a large number of Layer 4 - 7 network appliances from different vendors with different management tools. The appliances they mentioned were:

- Server load balancers and application delivery controllers;
- WAN optimization;
- Firewalls;
- SSL/IPSec VPNs and Intrusion Detection and Prevention Systems.

In that white paper ONUG coined the phrase *Network Services Virtualization* (NSV) to refer to the virtualization of functions such as the ones listed above. The paper also stated that the NSV use case "Seeks to leverage the flexibility and low costs of commodity servers to establish a scale-out pooling of virtual and physical appliances, which can be put to use servicing applications." ONUG went on to say "As each Layer 4 - 7 function is virtualized in software, it provides the following benefits:

- Lower CAPEX costs (approximately 30 percent less);
- Rapid service provisioning;
- Reduced risk through service distribution;
- Eased management and reduced operational costs;
- Consistent policies across different Layer 4-7 services and across data center, campus, and WAN networks;
- Programmatic control and ability to offer network functions as a service to developers.

Other potential benefits of NSV include the ability for IT business leaders to deliver on-demand or self-service IT delivery to business unit managers."

There clearly are differences between what ETSI is trying to accomplish with NFV and what ONUG is trying to accomplish with NSV. As mentioned, CSPs hope to virtualize some functionality that few if any enterprise organizations implement and their need for scale far surpasses what is needed by the vast majority of enterprise organizations. In addition, CSPs are notably more likely to have a requirement to link the usage of virtualized network functions to their billing systems than do enterprise organizations. However, if you change at most a few words in how ONUG describes the NSV use case it sounds exactly like what ETSI and others are trying to achieve with NFV. In addition, if you look at the list of appliances mentioned in the ONUG paper, they are all contained in Table 1.

To test the conventional wisdom about the applicability of NFV, the survey respondents were asked to indicate their view of the relevance of NFV to an enterprise IT architecture. Their responses are shown in **Table 16**.

Table 16: Relevance of NFV to Enterprises

Applicability	Percentage
Very Significant	10%
Significant	42%
Moderate	15%
Some	13%
None	1%
Don’t know/NA	17%
Other	1%

Half of IT professionals believe that NFV has either significant or very significant relevance to enterprise IT architectures.

The bottom line is that conceptually NFV and NSV have far more points of commonality than differences and the perceived relevance of NFV to enterprises is reflected in **Table 16**. As a result of these factors, throughout The Guide, the acronym *NFV* will be used to discuss the virtualization of network functions, whether those functions are used by CSPs or enterprise organizations or both.

The Relationship between SDN and NFV

The conventional wisdom has been that SDN and NFV were separate initiatives which could evolve independently of each other. However, in 2014 the Open Networking Foundation (ONF) and the ETSI NFV ISG announced the signing of a Memorandum of Understanding (MOU). As part of the announcement of 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."

As part of the announcement, the ONF released a document entitled the *OpenFlow-enabled SDN and NFV Solution Brief*. That document discussed how OpenFlow-enabled SDN can meet the need for automated, open, and programmable network connectivity to support some of the ETSI-defined use cases such as *Network Functions Virtualization Infrastructure as a Service* and *Virtual Network Function Forwarding Graph*.

In a white paper ETSI expressed their belief that NFV and SDN are highly complementary efforts. The ETSI view is that both efforts are seeking to leverage virtualization and software-based architectures to make network infrastructures more cost-effective and more agile in their ability to accommodate the dynamic nature of the workflows demanded by applications and end users. While NFV can be implemented using a non-SDN infrastructure, the ETSI vision is that NFV and SDN will increasingly be intertwined into a broad, unified software-based networking paradigm based on the ability to abstract and programmatically control network resources.

Some of the ways that ETSI believes that NFV and SDN complement each other include:

- The SDN controller fits well into the broader concept of a network controller in an NFV-Infrastructure (NFVI) network domain as defined in ETSI's NFV architectural framework.

- SDN can play a significant role in the orchestration of the NFV Infrastructure resources, both physical and virtual, enabling functionality such as provisioning, configuration of network connectivity, bandwidth allocation, automation of operations, monitoring, security, and policy control.

- SDN can provide the network virtualization required to support multi-tenant NFVIs.

- Forwarding Graphs can be implemented using the SDN controller to provide automated provisioning of service chains while ensuring strong and consistent implementation of security and other policies.

- The SDN controller can be run as a virtual network function (VNF), possibly as part of a service chain including other VNFs. For example, applications and services originally developed to run on the SDN controller could also be implemented as separate VNFs.

To test the conventional wisdom, 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 17**.

Table 17: Perceived Relationship between SDN and NFV

Relationship	Percentage
They are totally independent activities	8%
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.	65%
In at least some instances, NFV requires SDN	17%
In at least some instances, SDN requires NFV	12%
Don’t know	18%

Some of the conclusions that can be drawn from the data in **Table 17** are:

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

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

The Adoption of NFV

Extent of NFV Adoption

The Survey Respondents were given a set of alternatives and were asked to indicate the alternatives that described their company's current approach to implementing NFV. Their responses are shown in **Table 18**.

Table 18: Current Approaches to Implementing NFV

Approach to Implementing NFV	Percentage
We are currently actively analyzing the potential value that NFV offers	25%
We will likely analyze NFV sometime in the next year	24%
We are currently actively analyzing vendors' NFV strategies and offerings	23%
We currently are running NFV either in a lab or in a limited trial	18%
We have not made any analysis of NFV	18%
We expect that within a year that we will be running NFV either in a lab or in a limited trial	17%
We currently are running NFV somewhere in our production network	14%
We looked at NFV and decided to not do anything with NFV over the next year	8%
We expect that within a year that we will be running NFV somewhere in our production network	7%
Other	7%

The data in **Table 18** indicates:

While only a modest number of IT organizations have implemented NFV in a production network, a large percentage of IT organizations are currently in varying stages of analyzing NFV.

The Survey Respondents were asked to indicate the primary factor that is driving their company's interest in NFV. Their responses are shown in **Table 19.**

Table 19: Factors Driving NFV

Factor	Percentage
Reduce the time to deploy new services	26%
Greater management flexibility	16%
Better customer experience	14%
Reduce OPEX	11%
Reduce CAPEX	10%
Better network performance	9%
No driver	9%
Other	6%

The data in **Table 19** indicates:

By a wide margin, the primary factor driving interest in NFV is the reduction in the time it takes to deploy new services.

The Survey Respondents were also asked to indicate the three biggest inhibitors to their company broadly adopting NFV sometime in the next two years. Their responses are shown in **Table 20.**

Table 20: Factors Inhibiting NFV

Inhibitor	Percentage
The lack of a compelling business case	29%
Concerns about how we would do end-to-end service provisioning that includes physical and virtual resources and which may cross multiple partners' domains	23%
Concerns about security vulnerabilities	23%
The immaturity of the current products	22%
The need to significantly reskill our employee base	22%
The need for sophisticated orchestration capabilities	16%
The immaturity of the enabling technologies	15%
The need to make significant cultural changes in order to fully realize NFV's promise	13%
The difficulty of doing end-to-end service management	12%
The need to make significant organizational changes in order to fully realize NFV's promise	11%
The need to implement a new generation of agile OSS/BSS	10%
The confusion and lack of definition in terms of vendors' strategies	9%
No inhibitors to implementing NFV	9%
Other technology and/or business priorities	8%
The lack of a critical mass of organizations that have deployed NFV	8%
Concerns about how we would evolve from a POC to broad deployment	8%
The time it will take for standards to be developed and implemented	7%
Other	6%
The reluctance on the part of some of our suppliers to embrace a software model	2%
The requirement to make significant changes to our procurement processes	1%

The data in **Table 20** indicates:

The biggest inhibitors to the broad adoption of NFV are:

- ***The lack of a compelling business case;***
- ***Concerns about end-to-end service provisioning;***
- ***Concerns about security vulnerabilities;***
- ***The immaturity of the current products;***
- ***The need to significantly reskill our employee base.***

The Survey Respondents were also asked to indicate how long it would be before their organization had virtualized 25% of its L4 – L7 functionality such as optimization and security appliances. Their responses are shown in **Table 21.**

Table 21: Time Frame for Deployment	
Time Frame	**Percentage**
Already have	7%
1 – 2 years	30%
3 – 4 years	32%
5 – 6 years	6%
7 or more years	1%
Don’t know/ Not Applicable	24%

The data in **Table 21** indicates that:

Within a few years, the majority of IT organizations are likely to have made a significant deployment of virtualized L4 – L7 functionality.

Industry Organizations Driving the Evolution of SDN and NFV

Although there is some overlap, the organizations driving the development of SDN and NFV fit into three broad classes. One class is industry groups such as the ONF and ETSI. This class of organization develops use cases, best practices, architectures, frameworks, APIs, vocabulary and POCs. When ETSI establishes an Industry Specification Group (ISG) such as the one it established for NFV (ETSI NFV ISG), the ISG has a two year life cycle. After that they either establish a charter for a new phase of the ISG or they go away. This approach tends to make an ISG very action oriented.

Another group driving the evolution of SDN and NFV are Standards Developing Organizations (SDOs) such as the IETF and the Alliance for Telecommunications Industry Solutions (ATIS). Unlike an ETSI ISG, a SDO typically doesn’t have a predetermined life span. As such, they tend to move slowly and focus on technical elegance. There is no doubt that in some situations that technical elegance provides value. There is also no doubt that in many situations the pursuit of technical elegance results in a process that isn’t very agile. The IETF is an example of a SDO that is attempting to become more agile as evidenced of the hackfests that it recently conducted.

The third group driving the evolution of NFV is the open source community including organizations such as OpenDaylight (ODL), ON.Lab and the Open Platform for NFV (OPNFV), all three of which are member of the Linux Foundation. The general charter of this class of organization is captured in the initial announcement that the Linux Foundation made about OPNFV. As part of the announcement the Linux Foundation declared that OPNFV will establish a carrier-grade, integrated, open source reference platform that industry peers will build together to advance the evolution of NFV and ensure consistency, performance and interoperability among multiple open source components. The Foundation also stated that because multiple open source NFV building blocks already exist, OPNFV will work with upstream projects to coordinate continuous integration and testing while filling development gaps. The bottom line being these groups are developing platforms that over time will become quite feature-rich, and many companies are likely to build their offerings based on these platforms.

It is unclear how the relationship between the SDOs and the open source community will develop. One option is that after a group such as the OPNFV has make progress on creating an open source reference platform for NFV, that one or more SDOs will establish working groups to create standards for some of the key tasks that are part of the reference platform. However, since SDO working groups have historically taken years to create new standards, another option is that whatever functionality is part of the reference platform will become defacto standards.

In order to understand the conventional wisdom relative to the value provided by SDO and open source communities, The Survey Respondents were also asked to indicate the type of organization they thought would have greater impact on the evolution of NFV – SDOs such as the IETF or open source organizations such as OPNFV. Their responses are shown in **Table 22**.

Table 22: Influence of SDOs and Open Source Communities

Prime Influencer	Percentage
Open Source Communities	49%
SDOs	27%
Don't know	24%

By almost a 2:1 ratio, IT professionals think that open source communities will have more of an impact on the evolution of NFV than SDOs will.

Key members of the SDN and NFV community

The role that ETSI plays in the evolution of SDN and NFV was described previously in this chapter of The Guide and will be elaborated on in the next chapter (Architectural Considerations and Use Cases for NFV). The role played by OpenDaylight, On.Lab and the ONF was described in Chapter 1 of The Guide (A SDN Status Update). Below is a description of some of the other key organizations driving the evolution of SDN and NFV.

The OpenSwitch Community

In October 2015 the OpenSwitch community was announced. The goal of the community is to develop an open source network operating system (NOS). While there are currently other open source NOSs available, the founders of the OpenSwitch community believe that none of the existing open source NOSs met the requirement for a programmable and scalable NOS that also allows developers to access the source code, rather than just access the NOS through APIs.

Developers and users can download the newly released Linux-based open source NOS, which includes the following functionality and characteristics:

- A fully featured NOS with L2/L3 protocol support;
- An open source cloud database for persistent and ephemeral configuration;
- A system database to support all inter-module communication;
- A universal API approach including CLI, REST, Puppet/Chef, Ansible

The OpenStack Foundation

OpenStack is a cloud computing orchestration project offering free open source Orchestrator 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).

Open Platform for NFV (OPNFV)

As mentioned, the OPNFV mission is to establish a carrier-grade, integrated, open source reference platform that industry peers will build together to advance the evolution of NFV and ensure consistency, performance and interoperability among multiple open source components. In June 2015 OPNFV had their first software release, code named Arno. Arno enables end users to deploy virtual network functions (VNFs) on the platform to test functionality and performance. Arno also reflects OPNFV’s commitment to testing by providing an automated toolchain that allows upstream projects to do automatic builds and verification as they develop independently.

TM Forum

In 2014 the TM Forum announced its Zero-touch Orchestration, Operations and Management (ZOOM) project. According to the Forum, the goal of Zoom is to define a vision of the new virtualized operations environment and a management architecture based on the seamless interaction between physical and virtual components that can easily and dynamically assemble personalized services. In addition, the TM Forum has also been active with a wide range of companies to create what the TM Forum refers to as Catalysts, which are short-term collaborative projects led by members of Forum that address operational and systems challenges.

Internet Engineering Task Force (IETF)

Although their efforts are just getting started, the IETF can be expected to play a significant role in the evolution of standards for SDN and NFV. For example, the IETF Service Function Chaining (SFC) Work Group (WG) currently has a number of Internet drafts on the topic of delivering traffic along predefined logical paths incorporating a number of service functions. As described in one of those Internet drafts, the basic concept of SFC is similar to ETSI NFV ISG’s Virtualized Network Function (VNF)-Forwarding Graphs.

The Alliance for Telecommunications Industry Solutions (ATIS)

ATIS is a standards organization that develops technical and operational standards and solutions for the Information and communications technology (ICT) industry. ATIS has launched an NFV Forum to make contributions to NFV and SDN technologies. Phase I of the NFV Forum work program includes virtual network operator capabilities as well as other high priority use cases. The forum will focus on technical requirements, a catalog of needed capabilities, and the service chaining necessary for a third party service provider or enterprise to integrate NFVs into a business application. This process will result in creation of specifications that are complementary with existing industry work with an emphasis on facilitating inter-provider NFV. The forum will also engage relevant open source activities and agile software methodologies for the implementation of these capabilities.

The 3rd Generation Partnership Project (3GPP)

3GPP is a collaboration between groups of telecommunications associations. While its initial focus was on 3G as well as the completion of the first LTE and the EPC specifications, 3GPP has evolved to become the focal point for mobile systems beyond 3G. 3GPP standardization encompasses Radio, Core Network, and Service architecture. A number of functions defined in the 3GPP architecture are candidates for implementation as NFVs and have been identified as such in ETSI uses case descriptions. As a result, the 3GPP Telecom Management working group will produce a study Item on the management of 3GPP NFVs. 3GPP is also considering how the work in the ETSI NFV ISG might impact 3GPP at the architecture and system level.

The Metro Ethernet Forum (MEF)

The MEF is the defining body for the global market for Carrier Ethernet (CE). MEF's flagship work is CE 2.0, including specifications and related certification programs for services, equipment and professionals. MEF has announced a new Third Network vision that delivers Internet-like agility and ubiquity with CE 2.0-like performance and security. The Third Network vision is based upon the concept of Network as a Service (NaaS) incorporating service orchestration functions, APIs, a protocol independent NaaS information model and service definitions.

AUTOMATE YOUR CLOUD WITH aCLOUD SERVICES ARCHITECTURE

Integrate dynamic services into your Cloud Data Center

www.a10networks.com

Network Functions Virtualization (NFV): Architectural Considerations and Use Cases

Architectural Considerations

Before an organization adopts NFV they need to address some key considerations relative to how they will architect their data center to support NFV and related initiatives. As part of this process, organizations need to thoroughly examine the approach that has been adopted by alternative vendors of Virtual Network Functions (VNFs). It is important to address these issues because they have a major impact on a number of factors, including a solution's:

- Longevity;
- Scalability;
- Interoperability;
- Performance;
- Profitability/Cost[2];
- Risk;
- Availability.

Below are some architecture-related topics that should be addressed prior to adopting NFV.

Big Bang vs. Piecemeal Approach

One of the key architectural questions facing an organization is whether to take a tactical approach (a.k.a., a piecemeal approach) or an architectural approach (a.k.a., a big bang approach) to NFV. Companies that take a piecemeal approach typically focus on one, or at most a small number of use cases for which they see clear business value and for which they participate in Proof of Concept (POC) trials to demonstrate the viability of the possible solutions. For example, an organization could implement just a single use case; e.g., virtual CPE.

In contrast to focusing on rolling out solutions for just a few use cases, when a company takes a big bang approach to NFV they decide on an architecture and some key enabling technologies that the architecture will utilize in order to support any and all NFV use cases. AT&T is an example of a company that is taking a big bang approach to NFV as evidenced by its Domain 2.0 initiative.

Software Modularity

One technique that is associated with maximizing profitability and reusability is modular programming. The phrase *modular programming* refers to a software design technique that emphasizes taking a piece of software-based functionality and decomposing it into independent modules, such that each module contains everything necessary to provide one component of the desired functionality.

[2] Communications Service Providers tend to be more concerned with maximizing profitability and enterprise organizations tend to be more concerned with minimizing cost. Throughout the rest of The Guide, those two goals will be regarded as the same goal.

Organizations should place a preference on acquiring VNFs that were designed in a modular fashion.

Technology Considerations

The adoption of NFV is still in its early stages and these early stages are characterized by rapidly changing technologies. For example, the initial discussion of NFV focused on the use of CloudStack and now OpenStack has largely replaced CloudStack. In addition, most of the discussion of VNFs to date has them running in virtual machines (VMs). However, there is beginning to be discussion about VMs being replaced by containers or Unikernels.

While a VM is certainly one approach to implementing a VNF, there is nothing about a VM that encourages a modular approach. Another disadvantage of a VM-based approach is that because a VM contains a full server hardware stack, instantiating a new VM can be relatively time consuming which negatively impacts both HA and the time it takes to scale functionality to meet peak loads. In addition, while it is challenging to take an application running in a VM and modularize it to run in containers or unikernels, applications that are designed either around containers or unikernals are backwards compatible with VM-based deployments.

Key Virtualization Concepts

Virtual Machines (VMs) are an abstraction of physical hardware in which each VM has a full server hardware stack from virtualized BIOS to virtualized network adapters, storage, and CPU.

Containers don’t virtualize the entire server hardware stack. Instead, the virtualization layer runs as an application within the operating system.

Unikernels are specialized operating system kernels that act as individual software components. A full application or appliance consists of a set of running unikernels working together as a distributed system.

To the degree possible, organizations need to adopt an architecture that can evolve as the enabling technologies change without requiring a major overhaul.

Software-Centric Design

As the industry makes the shift from a hardware centric environment to a virtualized environment, some vendors will offer VNFs that are based on merely porting the code from their hardware-based appliance over to a VM. While that is an expedient approach, it may lead to significant performance problems as that code was originally designed to leverage the underlying specialize hardware which will no longer be present.

In order to achieve maximum performance, organizations should focus their attention on VNFs that were designed to run effectively in a software-centric environment.

The Role of Open Source

As mentioned in Part 4 of The Guide (A NFV Status Update), one of the major groups of players that is driving the evolution of NFV is the open source community. For example, in September 2014 the Linux Foundation announced the founding of the Open Platform for NFV (OPNFV) Project. As discussed in Part 4 of The Guide, in the announcement the Linux Foundation declared that OPNFV will establish a carrier-grade, integrated, open source reference platform that industry peers will build together to advance the evolution of NFV and ensure consistency, performance and interoperability among multiple open source components. The Foundation also stated that because multiple open source NFV building blocks already exist, OPNFV will work with upstream projects to coordinate continuous integration and testing while filling development gaps. One of the goals of working with upstream projects is to ensure that it is easy to load a VNF and have it run correctly regardless of the underlying physical infrastructure.

Organizations need to recognize that solutions that are based on open source solutions will potentially evolve quickly and potentially have a high degree of interoperability.

Relationship with SDN

As mentioned in Part 4 of The Guide (A NFV Status Update), there is a growing relationship between SDN and NFV. To exemplify the potential interaction between SDN and NFV, consider a situation where load balancer services are implemented as VNFs. If demand for load balancing capacity increases, a network orchestration layer can rapidly spin up new load balancing instances and also adjust the network switching infrastructure to accommodate the changed traffic patterns. In turn, the load balancing VNF entity can interact with the SDN controller to assess network performance and capacity and use this additional information to balance traffic better, or even to request provisioning of additional VNF resources.

Organizations should plan for, trial and adopt NFV and SDN in an integrated fashion.

Software Defined Data Center

Using the data center as a model, there is a strong movement away from a static environment in which most functionality is provided on a piece of dedicated hardware and the interface into the hardware is manual and a corresponding movement towards a software defined data center (SDDC). In a SDDC:

- Computing, storage and networking are virtualized and are all pooled resources;
- There are programmatic interfaces into all of the data center resources;
- Automated management delivers a framework for policy-based management of data center application and services.

Few if any organizations will fully implement a SDDC in the near term. SDDCs do, however, represent the general direction that data center design is taking. As such:

Organizations should ensure that whatever NFV related functionality it implements fits with the broader view of a SDDC.

A Fresh Approach to High Availability

Organizations have historically designed their systems for High Availability (HA). Communications Service Providers (CSPs), for example, have typically had the goal of five 9s availability. This means that a system is available 99.999% of the time, or conversely, that a system is unavailable for roughly 5 minutes or less a year. This level of availability is achieved through a variety of techniques. One technique is to acquire equipment that is designed for high availability in part through the use of multiple power supplies and processors. This technique is then supplemented within a central office or data center with generators for longer power outages than batteries can handle, as well as multiple diverse communication lines both within and between facilities.

The implementation of NFV enables organizations to rethink HA.

NFV enables organizations to shift the focus on HA away from redundant systems of highly reliable physical components to a focus on services that are supported by the inherent services and capabilities of the underlying NFV infrastructure (NFVI) layer. With NFV, when there is a failure, the impacted traffic will be re-directed to a new instance or a load-balanced instance of that application either in the same data center or across disparate data centers.

The (potential) end of Moore's Law

One of the key assumptions associated with implementing a growing range of functionality in software is that the general purpose computers on which this software runs will become increasingly powerful. This concept is described in what is usually referred to as Moore's Law in recognition of the fact that in 1965 Gordon Moore, co-founder of the Intel Corporation, described a doubling every year in the number of components per integrated circuit. In 1975, Moore revised the forecast doubling time to two years.

Unfortunately a number of people believe that the doubling of compute power referenced in Moore's Law may be coming to an end. One of the people who believe that is Gordon Moore himself. In a recent article Gordon Moore said that "Moore's Law" is not a law, but an observation and a projection. He also said that the current approach to making integrated circuitry, which is based on continually making things smaller and denser, is coming close to running into some fundamental limits, such as the speed of light. He added that there are other technologies that have been proposed to extend beyond what can currently be done with silicon, but he declined to speculate on how likely it was that they would be successful.

If Moore's law does come to an end, it will limit the range of functions that can successfully migrate from a hardware-centric implementation to a software-centric implementation. There is nothing that an organization can do to impact whether or not the phenomena described by Moore's law is coming to an end. However:

> ***Organization should monitor whether or not Moore's law is coming to an end and if it is, they need to adjust their plans to move from a hardware-centric approach to a software-centric approach.***

Status of NFV-Related Architectures

The Survey Respondents were asked to indicate the progress their organization has made relative to developing an effective architecture for the broad adoption of NFV. Their responses are shown in **Table 23.**

Table 23: Progress Towards a NFV Architecture

Amount of Progress	Percentage
None	26%
A little, with a lot of work ahead of us	41%
A lot, but still some work ahead of us	16%
Already developed an architecture	17%

Table 23 indicates that:

> ***Two thirds of IT organizations have made little or no progress towards the development of a NFV architecture.***

The Survey Respondents were also asked to indicate how much time their organization will spend over the next year developing a NFV-related architecture. Their responses are shown in **Table 24**.

Table 24: Amount of Time to be Spent on a NFV Architecture

Amount of Time	Percentage
None	16%
A modest amount	49%
A significant amount	30%
None, we already have an architecture	6%

The combination of **Table 23** and **Table 24** indicates that:

> ***While some organizations are making significant progress towards the development of a NFV architecture, the majority are not.***

Use Cases and Proof of Concept Trials

The European Telecommunications Standards Institute's (ETSI) Industry Specifications Group (ISG) for Network Functions Virtualization (ETSI NFV ISG) has defined a framework for coordinating and promoting public demonstrations of POC platforms. The PoC Framework outlines:

- The rationale for NFV PoCs;
- The NFV PoC process;
- The format and criteria for NFV PoC proposals;
- The NFV PoC Report format and requirements.

It is ETSI's intention that results from PoCs will guide ongoing standardization work by providing feedback on interoperability and other technical challenges. ETSI POCs are scoped around potential use cases that ETSI identified and which are described below. As of October 2015, ETSI was involved in 11 ongoing POCs.

ETSI NFV Use Cases

Below is a discussion of nine potential use cases for NFV that have been defined by the ESTI NFV ISG. A thorough description of the use cases is available on the ETSI web site.

NFV Infrastructure as a Service (NFVIaaS)

NFVIaaS is analogous to a cloud IaaS that is capable of orchestrating virtual infrastructures that span a range of virtual and physical network, compute, and storage functions. Unlike a traditional IaaS, NFVIaaS would be built on ETSI NFV standard interfaces and would also embrace an information model and network services interfaces that would allow the NFV Infrastructure (NFVI) to span the administrative domains of multiple service providers.

Virtual Network Functions as a Service (VNFaaS)

Many enterprises are deploying numerous network service appliances at their branch offices. Network services commonly installed at the branch can include access routers, WAN optimization controllers, stateful firewalls, intrusion detection systems, and DPI analysis devices. If a number of these functions are implemented on dedicated physical appliance platform, the result can often be a complex, expensive, and difficult-to-manage branch office network.

An alternative solution for enterprise branch office networks is to subscribe to VNFs that are hosted on servers in the network service provider's access network PoP. VNFs delivered as a Service (VNFaaS) are analogous to cloud networking SaaS applications where the subscriber pays only for access to the service and not the infrastructure that hosts the service.

Virtualization of the Home Environment

Virtualization of the Home Environment (VoHE) with NFV is analogous to VNFaaS. In this case the residential gateway (RGW) and the set top box (STB) are virtualized as VNFs residing on servers in the network service provider's PoP. All of the functions of these devices can be supplied as VNFs, including IP routing, NAT, firewall, DHCP, DVR/PVR disk, VoD client, etc. One of the primary benefits of VoHE is that it greatly simplifies the electronics environment of the home, reducing end user and operator CAPEX. In the ultimate scenario, all that is required in the home is a WiFi-enabled Layer 2 switch. Another benefit is that servicing RWGs and STBs is greatly simplified, reducing operator OPEX. However, accessing VNFs remotely would require significantly increased network access bandwidth. Another impediment is that hosting the large numbers of VNFs required in densely populated residential areas would require massive processing power as well as the development of a methodology where multiple VNFs could share a single virtual machine.

VNF Forwarding Graph (FG)

Network Service Providers offering infrastructure-based cloud services (e.g., IaaS) need to be able to orchestrate and manage traffic flows between virtualized service platforms (e.g., VNFs) and physical devices in order to deliver a complete service to the end user.

As noted elsewhere in The Guide, an SDN controller can be programmed to create the desired traffic flow. The VNF Forwarding Graph (VNF FG) is a service that provides flow mapping (a.k.a., service stacking or chaining) from a management and orchestration system that may or may not be part of an SDN infrastructure.

The VNF FG is based on an information model that describes the VNFs and physical entities to the appropriate management and/or orchestration systems used by the service provider. The model describes the characteristics of the entities including the NFV infrastructure requirements of each VNF and all the required connections among VNFs and between VNFs and the physical network included in the IaaS service. In order to ensure the required performance and resiliency of the end-to-end service, the information model must be able to specify the capacity, performance and resiliency requirements of each VNF in the graph. In order to meet SLAs, the management and orchestration system will need to monitor the nodes and linkages included in the service graph. In theory, the VNFs FG are able to span the facilities of multiple network service providers.

Virtual Network Platform as a Service (VNPaaS)

VNPaaS is similar to an NFVIaaS that includes VNFs as components of the virtual network infrastructure. The primary differences are the programmability and development tools of the VNPaaS that allow the subscriber to create and configure custom ETSI NFV-compliant VNFs to augment the catalog of VNFs offered by the service provider. This allows all the 3rd party and custom VNFs to be orchestrated via the VNF FG.

Virtualization of Mobile Core Network and IP Multimedia Subsystem

ETSI has published a document that defines the terminology and acronyms associated with digital cellular communications. That document is helpful when reading any discussion of digital cellular communications, including the discussion below. Some of the acronyms included below are:

EPC	Evolved Packet Core
MME	Mobile Management Entity
S/P GW	Serving gateway/public data network gateway
IMS	IP Multimedia Subsystem
P-CSCF	Proxy - Call Session Control Function
S-CSCF	Serving - Call Session Control Function
PCRF	Policy and Charging Rules Function
HSS	Home Subscriber Server
RLC:	Radio Link Control
RRC:	Radio Resource Control
PDCP:	Packet Data Convergence Protocol
MAC:	Message authentication code
FFT:	Fast Fourier Transformation
RAN	Radio Access Network
EPS	Evolved Packet System
CoMP	Coordinated Multi Point transmission/reception

The 3GPP is the standards organization that defines the network architecture and specifications for Network Functions (NFs) in mobile and converged networks. Each NF typically is run on a dedicated appliance in the mobile network PoP. Running the NFs as VNFs on virtualized industry standard servers is expected to bring a number of benefits in terms of CAPEX, OPEX, as well as flexibility and dynamic scaling of the network to meet spikes in demand.

The latest architecture for the core of cellular systems is the EPC. In this architecture, the NFs specified include the MME and the S/P GW. In the IMS NFs include: the P-CSCF and the S-CSCF, HSS, and the PCRF. HSS and PCRF are NFs that work on conjunction with core and IMS NFs to provide an end-to-end service. One possibility is to virtualize all the NFs in a NFVI PoP or to virtualize only selected NFs.

Virtualization of the Mobile Base Station

3GPP LTE provides the RAN for the EPS. There is the possibility that a number of RAN functions can be virtualized as VNFs running on industry standard infrastructure.
For traditional RAN nodes such as eNodeB, Home eNodeB, and Femto-Picocell, the target virtualization functions are Baseband radio Processing unit (including FFT decoding/encoding), MAC, RLC, PDCP, RRC, control, and CoMP. While this ETSI use case focuses on LTE, it would be possible to virtualize the functions of other RAN types, such as 2G, 3G, and WiMAX.

Virtualization of Content Delivery Networks (CDNs)

Some ISPs are deploying proprietary CDN cache nodes in their networks to improve delivery of video and other high bandwidth services to their customers. Cache nodes typically run on dedicated appliances running on custom or industry standard server platforms. Both CDN cache nodes and CDN control nodes can potentially be virtualized. The benefits of CDN virtualization are similar to those gained in other NFV use cases, such as VNFaaS.

Virtualization of Fixed Access Network Functions

NFV offers the potential to virtualize remote functions in the hybrid fiber/copper access network as well as PON fiber to the home and hybrid fiber/wireless access networks. Advanced versions of DSL (i.e., VDSL2 and G.fast) can deliver between 100 Mbps and 1 Gbps access speeds by leveraging fiber optics from the headend to the neighborhood cabinet or drop point and using legacy twisted pair to reach the final end user premises. In a DSL access network some of the functions that can potentially be virtualized include the DSLAM and Message Display Unit (MDU) forwarding functions, while control functions remain centralized at the central office.

Securing End-User Quality of Experience from the Cloud

There are many applications an enterprise will consider as mission-critical to their business. While it varies by enterprise, some example applications include customized CRM, web-based retail, accounting, and billing. Sonus has a long and successful history of securing and optimizing the end-user experience for real-time communications. As applications continue to migrate to the Cloud, Sonus is applying that knowledge to optimize mission-critical applications that are sensitive to IP transport.

As both service providers and enterprises look to embrace cloud-based environments, there remain some significant challenges that need to be addressed with respect to security, interoperability, portability, quality of service, quality of experience, and performance guarantees.

In a cloud-based environment, service providers must offer connectivity to an enterprise customer that is extremely resilient in order for mission-critical applications to be trusted and operated. The traditional way of provisioning, managing, and selling their network assets is no longer dynamic enough to keep up with the new demands of data center connectivity for the enterprise. Service providers have some challenges. They have to deal with competitive pressures that drive down pricing, yet also be responsive with the delivery of network resources and bandwidth connectivity that mission-critical applications require.

As enterprises will choose between various competing service providers, an additional important differentiator that needs to be addressed is perceived service quality. A service provider should be able to transparently monitor and react quickly to any service quality problems before an enterprise is aware. An optimal Quality of Experience (QoE), when end-users judge the usability of an application based on their own experience, must be achieved, while constraining the application to behave as efficiently as possible to minimize operational costs.

For today's enterprises/CIOs, they need connectivity solutions that allow them to manage their networks more intelligently and dynamically, defining end-to-end policies that align transport with mission-critical applications to deliver a high QoE within their tight operational budgets. The ability to understand and manage QoE for end-users provides a great opportunity to set themselves apart.

What is necessary to meet these needs for both service providers and enterprises requires new approaches that guarantee adherence to concerns on security, as well as to industry requirements for lifecycle management of the services and network resources.

Combining the Intelligence of the Session and Network Control Layers

Sonus provides a solution that combines session layer intelligence with software-defined networking intelligence at the network layer.

Sonus' Session Border Controller (SBC) SWe, integrated with VellOS, Sonus' virtualized network control platform, enables the sharing of security, and policy management information between the session layer and the network control layer. The application-specific intelligence from the Sonus SBC SWe, combined with VellOS' knowledge of traffic flows at the network control layer, gives service providers the ability to offer much higher levels of quality of service than ever before, guaranteeing bandwidth for specific mission-critical applications.

The Sonus solution provides a holistic, systems approach to security—providing a security perimeter in real-time at the network edge. As a result, enterprises can make informed choices and dynamically compose and personalize services in a secure way through transparent interaction with the IP session and transport layer.

The Sonus solution enables delivery of mission-critical applications with an assurance of service level agreements (SLAs) without over-burdening the network. With this holistic view of how one network can intelligently optimize packet flows based on application prioritization, a service provider or enterprise will have a solution that monitors service parameters (like throughput) and automatically proactively react if network conditions may result in QoE degradation.

The combination of the Sonus SBC SWe and VellOS enables a guarantee of SLAs for specified bandwidth in real time for mission-critical applications. By integrating session layer intelligence with network control intelligence, data center network connectivity and cloud-based service delivery are optimized.

About Sonus Networks

Sonus enables and secures real-time communications so the world's leading service providers and enterprises can embrace the next generation of SIP and 4G/LTE solutions, including VoIP, video, instant messaging, and online collaboration. With customers in more than 50 countries and nearly two decades of experience, Sonus offers a complete portfolio of hardware-based and virtualized Session Border Controllers (SBCs), Diameter Signaling Controllers (DSCs), Cloud Exchange Networking Platform, policy/routing servers, and media and signaling gateways. For more information, visit www.sonus.net or call 1-855-GO-SONUS. Sonus is a registered trademark of Sonus Networks, Inc. All other company and product names may be trademarks of the respective companies with which they are associated.

Sonus Networks North American Headquarters

4 Technology Park Drive
Westford, MA 01886
U.S.A.
Tel: +1-855-GO-SONUS

Sonus Networks APAC Headquarters

1 Fullerton Road #02-01
One Fullerton
Singapore 049213
Singapore
Tel: +65-68325589

Sonus Networks EMEA Headquarters

Edison House
Edison Road
Dorcan, Swindon
Wiltshire
SN3 5JX
Tel: +44-14-0378-8114

Sonus Networks CALA Headquarters

Homero No. 1933-902
Col. Los Morales, C.P. 11510
Mexico City, Mexico
Distrito Federal
Mexico Tel: +52-55-1950-3036
Int'l Tel: +1-978-614-8741

To learn more, call Sonus at 855-GO-SONUS
or visit us online at www.sonus.net

Microsoft Partner
Gold Communications

Voice
Unified Communications
Business Productivity Solutions
Midmarket Solution Provider

The content in this document is for informational purposes only and is subject to change by Sonus Networks without notice. While reasonable efforts have been made in the preparation of this publication to assure its accuracy, Sonus Networks assumes no liability resulting from technical or editorial errors or omissions, or for any damages resulting from the use of this information. Unless specifically included in a written agreement with Sonus Networks, Sonus Networks has no obligation to develop or deliver any future release or upgrade, or any feature, enhancement, or function.

Sonus™
Cloud communications made smarter

Network Functions Virtualization (NFV): Operational Impediments

Performance Limitations

In order to obtain the potential cost and agility benefits of adopting NFV, it must be possible to achieve roughly the same performance in a software-based environment as is possible in a traditional hardware-based environment. However, in many cases that isn't possible without an enabling software architecture because of the bottlenecks that are associated with the hypervisors, virtual switches and virtual machines that are the foundation of the emerging software-based approach to IT. In response to the performance bottlenecks that are associated with NFV, ETSI has authored a document entitled "NFV Performance & Portability Best Practices".

Figure 4: Representative Bottlenecks in a Physical Environment

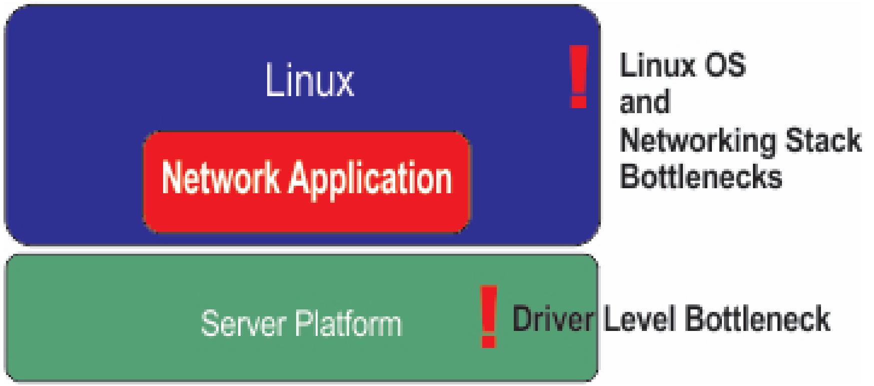

Performance bottlenecks are not unique to virtualized environments. For example, some of the bottlenecks that occur in a physical environment are shown in **Figure 4**.

Unfortunately, as shown in **Figure 5**, as IT organizations adopt a virtualized environment the performance bottlenecks multiply. **Figure 5** demonstrates some, but not all of the bottlenecks that can occur in a virtualized environment. For example, while not explicitly shown in **Figure 5**, VM to VM communications can also result in bottlenecks.

Figure 5: Performance Bottlenecks in a Virtualized Environment

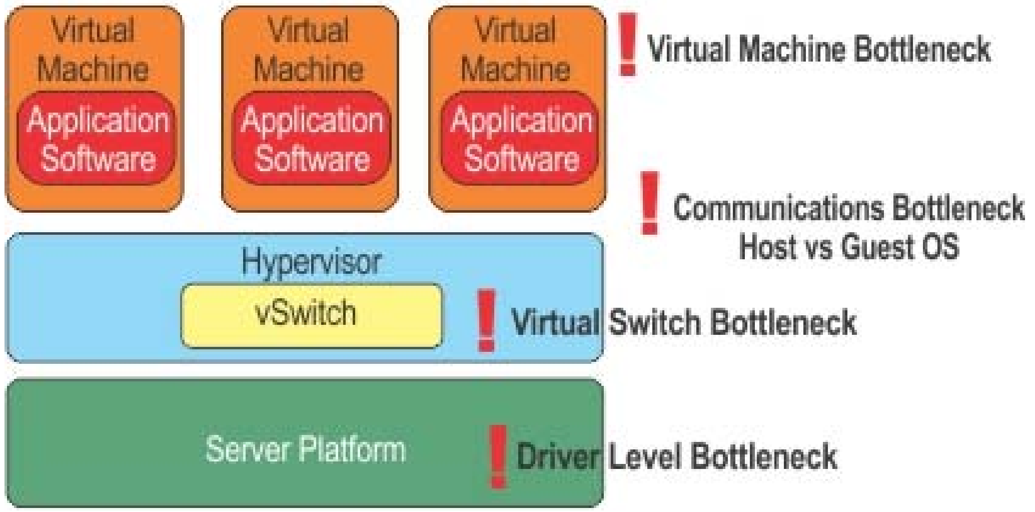

Acquiring solutions that have effective packet processing software that can bypass bottlenecks is one of the primary ways to avoid experiencing unacceptable performance in a virtualized environment. As shown in **Figure 6**, when evaluating the enabling packet processing software, IT organizations should check for the following criteria in order to ensure a cost effective value proposition, and smooth transition to future requirements:

- Equal performance in both physical and virtual environments;
- Transparency: No change should be required to the operating system, the hypervisor, the virtual switch or to the management tools;
- Availability: The solution must work across multi-vendor processors, NICs and hardware platforms.

Figure 6: Breadth of Environments

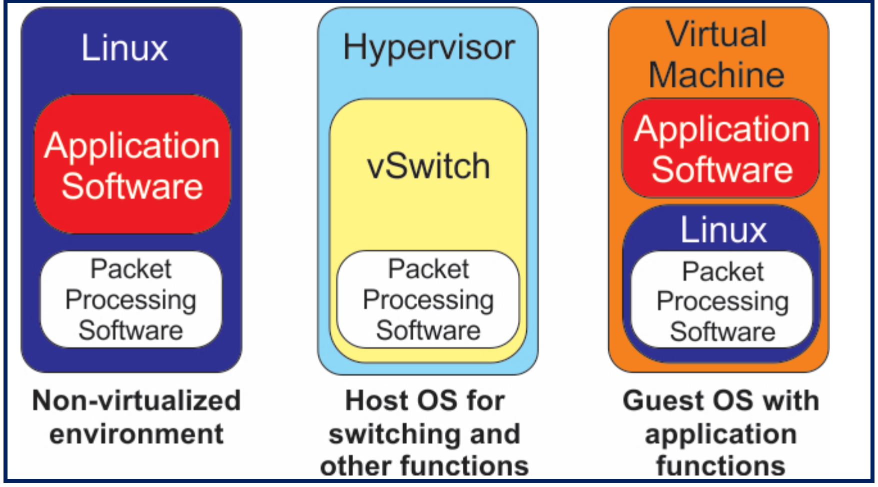
Linux
Application Software
Packet Processing Software
Non-virtualized environment
Hypervisor
vSwitch
Packet Processing Software
Host OS for switching and other functions
Virtual Machine
Application Software
Linux
Packet Processing Software
Guest OS with application functions

The evaluation criteria listed above are intended to ensure that the packet processing software can be easily and universally implemented on any version of Linux or on any hypervisor without requiring changes to existing environments.

End-to-End Management

Status of Management

The Survey Respondents were asked to indicate to what extent NFV creates fundamentally new management challenges. Their responses are shown in **Table 25**.

Table 25: Extent of New NFV Management Challenges

Extent of New Challenges	Percentage
No new challenges	8%
A few new challenges	41%
A broad range of new challenges	51%

Table 25 indicates that:

There is broad recognition on the part of IT organizations that the adoption of NFV creates new management challenges.

The Survey Respondents were asked to indicate how much progress their organization has already made relative to determining how they will respond to NFV's new management challenges. Their responses are shown in **Table 26**.

Table 26: Progress Towards Managing NFV

Amount of Progress	Percentage
None	27%
A little	48%
A lot	14%
We have a well-defined strategy	11%

Table 26 indicates that:

The vast majority of IT organizations have made little or no progress relative to determining how they will respond to NFV-related management challenges.

The Survey Respondents were asked to indicate how much time their organization will spend over the next year working on developing an approach to how they will respond to NFV-related management challenges. Their responses are shown in **Table 27**.

Table 27: Amount of Time to be Spent on NFV Management

Amount of Time	Percentage
None	13%
A modest amount	53%
A significant amount	32%
None – already done	1%

Table 27 indicates that:

Over the next year the vast majority of IT organizations will spend at least a modest amount of time working on developing an approach to how they will respond to NFV-related management challenges.

As discussed in Chapter 3 of The Guide (The Operational Impediments to Implementing SDN), 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. As a result, there is a natural affinity between Orchestration and NFV Management.

The Survey Respondents were asked to indicate the approach that their company is taking to orchestration. Their responses are shown in **Table 28**.

Table 28: Approach to Orchestration

Approach	Percentage
Developing a strategy but concerned that existing solutions are immature	35%
Don't have a strategy and unlikely to develop one in the near term	16%
Have a well thought out strategy and have begun to execute	15%
Developing a strategy and optimistic it will be completed quickly	14%
Don't know/NA	9%
Have a well thought out strategy but not yet begun to execute	6%
Other	5%

Table 28 indicates that:

There is significant interest in orchestration, but only a very small minority of IT organizations are using an orchestration platform in production.

Management Challenges

Throughout this chapter of The Guide, the phrase *service provider* will refer to both Communications Service Providers and to enterprise network organizations.

As is widely recognized, the adoption of NFV poses a number of significant challenges that must be overcome in order to ensure that IT organizations will be able to implement effective end-to-end management. These challenges include:

- **Dynamic relationships between software and hardware components.** In traditional networks, application software and network function software generally run on dedicated hardware that is statically provisioned by manual processes. With virtualization, software running on virtual machines (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. Due to the mobility of VMs, topology changes can occur in a matter of seconds or minutes rather than the days or weeks required for changing software/hardware relationships in traditional networks. In order to accommodate and leverage virtualization technologies, end-to-end management systems will need to be re-architected to be capable of implementing automated processes for virtual resource procurement, allocation, and reconfiguration in accordance with a set of highly granular policies designed to ensure the quality of experience for the user of the network services. Effective operations management also requires tools that give operators clear visibility into the relationships between the virtual and physical networks and their component devices. In particular, when performance or availability problems occur, both root cause analysis and impact analysis require bilateral mapping between the physical and virtual infrastructures.

- **Dynamic changes to physical/virtual device configurations.** To accommodate the dynamic nature of virtualized networks, end-to-end management systems will need to be able to adjust the configuration of devices to react to changing conditions in the network. For example, consider the traffic of an important application flow that has a medium priority class. If the network becomes congested, it may be necessary to change the traffic classification to be high in order to continue to meet an established SLA.

- **Many-to-Many relationships between network services and the underlying infrastructure**. In a typical traditional network infrastructure there is 1-to-1 relationship between a network service and a set of dedicated physical resources. In a virtualized infrastructure a network service can be supported by a number of Virtual Network Functions (VNFs) which may be running on one or several VMs. A single VNF may also support a number of distinct network services. In addition, the group of VNFs supporting a single network service could possibly be running on a number of distinct physical servers. As a result, end-to-end management systems need to support a three-tiered network model based on many-to-many relationships among network services, virtualization infrastructure, and physical infrastructure.

- **Hybrid physical/virtual infrastructures.** As virtualization is gradually adopted, service providers will need to be able to integrate virtual environments into their existing end-to-end traditional/legacy monitoring infrastructures. Therefore, end-to-end management systems developed for the virtual infrastructure will need to be compatible with legacy infrastructure.

- **Performance Monitoring.** Because of the inherent complexity and dynamic nature of NFV, a performance monitoring strategy and methodology must be developed early and applied consistently throughout the service design and development process. This will allow seamless integration of new VNFs into the existing end-to-end monitoring platform and it will also provide development and operations teams with a consistent methodology for service monitoring regardless of what combination of physical and/or virtual functions are used in the delivery of a service. The key will be the ability to consistently and reliably monitor the performance of a service not just the performance of VNFs.

- **Network services spanning multiple service providers.** Some of the VNFs comprising a virtualized network service may be hosted in the clouds of multiple collaborating providers. One major challenge in a multi-cloud environment is managing end-to-end service levels and SLA compliance. Since visibility into portions of the end-to-end path that are external to a service provider will always be limited, some form of aggregated external SLA data will have to be developed and imported from partner providers and the Internet. This requires a flexible and extensible end-to-end management architecture that provides consistent data collection and management interfaces across all on-net and off-net resources and technologies. Multi-cloud environments also require new approaches in managing end-to-end security.

- **IT and Network Operations collaboration.** These organizations will need to cooperate effectively to establish new operational processes that meet the demands of end-to-end management of hybrid physical/virtual infrastructures. This will require an effective DevOps organizational model for the development of network services based on NFV. One of the challenges will be to share the responsibilities for the various tasks involved in rolling out a new service. A key aspect of this cooperation will involve the selection and management of component VNFs, as well as testing and deploying the end-to-end management capability for the network service in question.

- **Hybrid environments**. For the foreseeable future, some services will be based on existing physical network functions while others will be based on VNFs and some others will be based on a hybrid environment made up of both. In a hybrid environment both types of function must have management interfaces built on a common information model (see below) in order to support agile DevOps-style service creation as well as the dynamic management and orchestration. In a hybrid environment it's crucial that management is policy-based and uses control loops to ensure quality of service.

- **Shared information model**. Where dynamic network service configurations are required, the management interfaces presented by both virtual and physical infrastructure elements need to lend themselves to automated plug and play integration. Information models drive consistency in the design of data payloads in automated interfaces by capturing behavior, defining standard interface communications patterns and specifying information representations; e.g., metrics representation and semantics for reporting SLA and QoS performance.

- **Policy based architecture**. Taking full advantage of the dynamic mature of virtualization requires an end-to-end management system that can perform as an autonomic system to support real time operational processes. A policy management architecture is the basis for automated management and orchestration. Policies can be based on hierarchical system of rules designed to deal with the complexities of a hybrid environment and to manage the relationships among users, services, SLAs, and device level performance metrics. For example, if the CPU utilization of a physical server hosting a VNF becomes excessive, the VNF may be moved to a server with lower utilization if that is in accordance with the SLA.

Management Direction

ETSI is working to drive how NFV will be managed. Towards that end, ETSI has established a management and orchestration framework for NFV entitled Network Function Virtualization Management and Orchestration. Some of the key concepts contained in that framework were summarized in another ETSI document. According to that document:

> *"In addition to traditional Fault, Configuration, Accounting, Performance, and Security (FCAPS) Management, the NFV Management and Orchestration framework introduces a new set of management functions associated with the lifecycle management of a VNF. The NFV ISG has focused on detailing these new sets of management functions, which include, but are not limited to: on-board a VNF, instantiate a VNF, scale a VNF, update a VNF, and terminate a VNF. A difference also worth highlighting relates to fault and performance management - in a virtualized environment this is the responsibility of different functional blocks at different layers. As a result, the correlation of faults, alarms and other monitored data such as performance metrics and resource usage, and the consequent fault resolution needed to operate the service in a reliable manner, will typically be distributed.*
>
> *Network Service Orchestration functions are responsible for coordinating the lifecycle of VNFs that jointly realize a Network Service. Network Service orchestration functions include on-boarding a Network Service, management of resources used by the Network Service, managing dependencies between different VNFs composing the Network Service, and managing the forwarding graphs between the VNFs. During the Network Service lifecycle, the Network Service orchestration functions may monitor Key Performance Indicators (KPIs) of a Network Service, and may report this information to support an explicit request for such operations from other functions.*

Expanding on the functional blocks and reference points identified by the NFV Architectural Framework, the NFV Management and Orchestration framework defines requirements and operations on the interfaces exposed and consumed by functional blocks associated with the different management functions; e.g. VNF lifecycle management, virtualized resource management. The objective of such an approach is to expose the appropriate level of abstraction via the interfaces without limiting implementation choices of the functional blocks. The document provides an extensive description of interfaces, which is the basis for future work on standardization and identification of gaps in existing systems and platforms."

Impact on Organizations and Jobs

The Survey Respondents were also asked how much of an impact they thought that NFV will have on the structure of their company's IT organization over the next two years. Their answers are shown in **Table 29**.

Table 29: Impact of NFV on Organizational Structure

Impact	Percentage
Very Significant Impact	9%
Significant Impact	27%
Moderate Impact	19%
Some Impact	19%
No Impact	8%
Don't Know	18%

The data in **Table 29** indicates:

> ***Roughly a third of IT organizations believe that over the next two years that the adoption of NFV is likely to have a significant or very significant impact on the structure of their organization.***

When The Survey Respondents were asked what the impact would be, their answers included that NFV will:

- Reduce the time it takes to deploy new offerings;
- Force realignment between departments;
- Drive changes to security models;
- Result in a reduction in the amount of labor that is required;
- Result in the assignment of new roles once the efficiencies are realized;
- Drive the need for cross domain management;
- Cause a change in their approach to application development.

The Survey Respondents were also asked how much of an impact they thought that NFV will have on the required skill base of their company's employees. Their answers are shown in **Table 30**.

Table 30: Impact of NFV on Employee Skills	
Impact	**Percentage**
Very Significant Impact	7%
Significant Impact	34%
Moderate Impact	26%
Some Impact	19%
No Impact	5%
Don’t Know/Other	10%

The data in **Table 30** indicates:

Over the next two years the adoption of NFV is likely to have a significant or very significant impact on the skill base of roughly 40% of all it professionals.

When The Survey Respondents were asked what the impact would be, their answers included that NFV will:

- Drive the need for employees to develop enhanced skills;
- Increase the time it takes to train new employees;
- Create the need for employees to have a knowledge of software tools;
- Drive the need for cross training of team members;
- Drive the need for combining IT and networking skills;
- Cause companies to replace employees who don’t adapt;
- Increase the need to transition from traditional network skills to cloud and systems skills;
- Result in DevOps skills replacing traditional networking skills.

DevOps

One of the implications of the ongoing virtualization of all forms of IT functionality is the adoption of a DevOps model. The point of adopting DevOps is to establish tight collaboration between a number of the phases of the application development lifecycle, including application development, testing, implementation and ongoing operations. With that goal in mind, some of the key characteristics that are usually associated with DevOps are that the applications development team continuously writes primarily small incremental pieces of code that are tested on an architecture that reflects the production architecture.

Those key principles that characterize DevOps are:

- **Collaboration**
  A key aspect of DevOps is to create a culture of collaboration among all the groups that have a stake in delivery of new software.

- **Continuous integration and delivery**
  With continuous integration, software changes are added to a large code base immediately after development so that new capabilities can be continuously delivered to the entire release chain for testing and monitoring in production-style environments.

- **Continuous testing and monitoring**
  With DevOps, testing is performed continuously at all stage of the release process and not just by the QA organization. Developers do testing and provide test data and procedures that can be used by collaborating groups downstream in the process. The operations group is also typically involved in the test and monitoring processes. Part of their value add is that operations groups can specify load patterns to make testing by other groups more in line with actual usage conditions.

  In addition, operations groups perform continuous monitoring to identify problems with the services being delivered so that they can be fixed in near real-time. Monitoring relies on an appropriate set of tools. The same tools that monitor the production environment can also be employed in development to identify performance problems prior to production deployment.

- **Automation**
  With DevOps all stages of software delivery are highly dependent on automated tools. Automation is essential because it enhances agility and provides the productivity required to support the continuous nature of integration, delivery, testing, and monitoring of many small increments to the code base.

- **API centric automated management interfaces**
  Software Defined Environments (SDEs) are an emerging core capability of DevOps that allow organizations to manage the scale and the speed with which environments need to be provisioned and configured to enable continuous delivery. SDEs use technologies such as API-centric automated management interfaces that define entire systems made up of multiple components. These interfaces are based on information models that define the characteristics, configurations, roles, relationships, workloads, and work- load policies, for all the entities that comprise the system.

All of the basic principles of DevOps are applicable in a network operations (NetOps) setting. However DevOps is generally applied to discreet services that are frequently delivered over the web on a best effort basis. The network environment is different than that and as a result virtualized network services development creates challenges that are not addressed by DevOps. One such challenge is that since Virtual Network Functions (VNFs) such as optimization and security are chained together to create an end-to-end service this creates strong dependencies between the VNFs. For example, if a service provider updates an optimization VNF they need to ensure that it is fully compatible with the security VNF(s). As a result much stronger version control and compatibility testing is needed than would be typical for enterprise applications.

Other challenges created by network services development that must be addressed by NetOps that were not addressed by DevOps include:

- Since for the foreseeable future the vast majority of environments will be a combination of hardware-based and software-based functionality, the NetOps methodology must accommodate services that depend on network functions running on dedicated hardware platforms as well as VNFs.

- Virtualized services will often be created by integrating services from multiple suppliers. This will require NetOps methodologies and best practices to support concurrent synchronized development and integration across the domains of multiple partners.

- Unlike what happens when delivering an application over the Web, NetOps will need to support dynamic and automated management of service performance and SLAs. This can only be achieved by a policy model that supports end-to-end SLA targets.

- Again in contrast to what often happens when delivering an application over the Web, NFV services are often mission critical. This creates a need for high levels of resilience and rapid fallback capabilities.

- Virtualized services will cover a very wide range of network functions and technologies. As a result, consistent frameworks and interfaces are needed in order to achieve the goal of minimizing or eliminating the need for manual intervention of any sort when incorporating VNFs into a network service.

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

Radware SDN applications improve application security, performance, and availability by integrating ADC and security intelligence with SDN to optimally analyze and forward traffic enhancing network services. Radware SDN applications integrate with the SDN application control plane and interact with the SDN controller to work with the Radware technologies throughout the application infrastructure.

Radware SDN-enabled ADC and security services transform applications from device-based solutions to become network wide services that intelligently divert traffic to service engines. Radware enhances SDN functions by leveraging our Virtual Application Delivery Infrastructure (VADI). This enables an EveryWare network service paradigm where applications are available anywhere and everywhere.

Key benefits of the Radware SDN network service infrastructure include:

- **Intelligent application delivery and security –** Optimal application service delivery
- **Easy implementation -** Improved operational efficiency of network management
- **Lower overall network service solution costs** – Deploy network services as needed
- **Greater scalability** – Scale network services throughout the network
- **Easier operational control** – Streamline network operations

DDoS Protection as a Native SDN Application

DefenseFlow is an SDN application features an adaptive, behavioral-based DoS attack detection engine and a traffic steering mechanism taking advantage of the software defined network elements for attack mitigation. DefenseFlow delivers a necessary security control plane for SDN-based networks.

SDN & NFV for a Scalable Application Delivery Network

Radware offers Alteon VA for NFV – the industry's first and highest performing ADC designed from the ground up to run in NFV environments. Alteon NFV provides a unique value proposition consisting of CAPEX/OPEX reduction, vendor agnostic technologies, high performance, enhanced scalability, orchestrated elasticity, and improved network service agility.

Radware's ElasticScale is an SDN application that wraps existing network service virtual appliances, including Alteon NFV to consistently deliver network services in an elastic network environment. ElasticScale can be utilized to help providers adopt network functions virtualization paradigms.

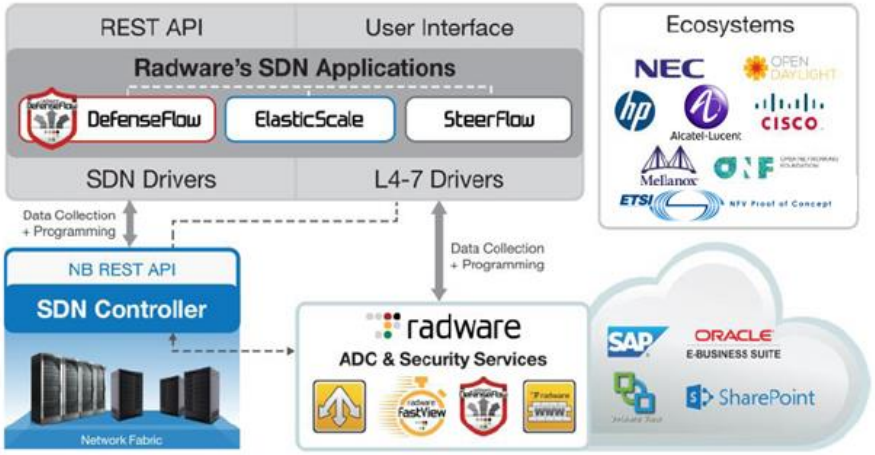

Partnering for Success: Our SDN and NFV Ecosystem

The SDN and NFV eco-systems are a critical focus for Radware. Through partnerships with the industry's leading SDN and NFV consortiums and vendors, Radware ensures customers that our application delivery and security solutions integrate successfully into target architectures.

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 SDN and NFV Ecosystem

The SDN Ecosystem

One measure of the extent of the SDN ecosystem is that there are approximately 130 full time members of the Open Networking Foundation (ONF) along with 32 startup members. This subsection of The Guide identifies the major categories of organizations that are part of the SDN ecosystem and briefly discusses the value proposition of each of the categories. This subsection of The Guide also identifies representative members of each category of organization that are part of the SDN ecosystem. The representative members that are identified either currently provide the indicated functionality or can be expected to provide the indicated functionality in the near term.

Merchant Silicon/Chip Vendors

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

Representative Members:

- Broadcom
- Intel
- Marvell
- Mellanox

HyperScale Data Centers

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

Representative Members:

- Yahoo
- Google
- Facebook

Telecom Service Providers

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

A preceding chapter of The Guide (The Use Cases and Business Case for SDN) discussed the interest that IT organizations have in either using SDN in the WAN or in acquiring a service from a WAN service provider that is based on SDN. Responding to that interest, vendors like Pertino (recently acquired by Cradlepoint) are currently using SDN and Network Function Virtualization (NFV) to enable them to offer a new generation of WAN services. AT&T has announced its interest in using both SDN and NFV to change how it offers services to its customers and Masergy has implemented a number of SDN-based services.

Representative Members:

- Pertino
- Deutsche Telekom
- NTT Communications
- Verizon
- AT&T
- Masergy

Switch Vendors

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

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

Representative Members:

- Alcatel-Lucent
- Cisco
- Dell
- HP
- NEC
- PICA8

Network and Service Monitoring, Management and Automation

Value Proposition: Most, if not all of the providers of SDN solutions will provide at least some ability for the consumers of those solutions to manage the solutions that they provide. The members of this category of the ecosystem don't provide SDN solutions themselves. The vendors listed below either currently provide, or soon will provide management functionality that isn't offered by the providers of SDN solutions and/or they integrate the management of these solutions into a broader management structure.

Representative Members:

- NetScout
- QualiSystems
- EMC
- CA

Providers of Network Services

Value Proposition: The members of this category provide network services such as security and optimization that are part of the overall SDN/NFV solution.

Representative Members:

- A10
- Radware
- Sonus
- HP
- Riverbed
- Citrix
- Cisco

Testing

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

Representative Members:

- QualiSystems
- InCNTRE
- Ixia
- Spirent

Providers of SDN Controllers

Value Proposition: These vendors provide the controllers that are part of any SDN solution.

Representative Members:

- NEC
- Nuage Networks
- HP
- Cisco
- Sonus
- Open Daylight Consortium
- ON.Lab
- VMware/Nicira

Providers of Telcom Service Provider’s Infrastructure/ Optical Networking

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

Representative Members:

- ADVA Optical Networking
- Ciena
- Cyan (Recently acquired by Ciena)
- Infinera
- ZTE Corporation

Server Virtualization Vendors

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

Representative Members:

- Citrix
- Microsoft
- VMware

The NFV Ecosystem

This subsection of The Guide identifies the major categories of organizations that are members of the NFV ecosystem and briefly discusses the value proposition of each of the categories. This subsection of The Guide also identifies representative members of each category of organizations that are part of the NFV ecosystem. The representative members that are identified either currently provide the indicated functionality or can be expected to provide the indicated functionality in the near term.

As a point of reference, an extensive list of NFV-related acronyms can be found in *Network Functions Virtualization (NFV): Use Cases*.

Telecom Service Providers

Value Proposition: Service providers are interested in NFV as a means of improving their ability to deliver services to their customers in a timely, cost-effective, and reliable manner. NFV, possibly in conjunction with SDN, has the potential to enable a new generation of services spanning a wide range of Virtual Network Functions (VNFs) that can generate new revenues from other service providers, enterprises, and residential customers.

Representative Members:

- AT&T
- Cablelabs (representing the cable industry)
- France Telecom S.A.
- Telefonica S.A.
- Masergy
- NTT Corporation

Virtualized Network Service and Cloud Service Vendors

Value Proposition: The members of this category provide VNFs that can be hosted on either the customer's server platforms or provided in the form of a Virtual Network Function as a Service (VNFaaS). Most of these organizations are focused on the communications service providers either as end users or as providers of services to enterprise and residential end users.

Representative Members:

- Sonus
- Allot Communications Systems Ltd
- Mavenir Systems UK Ltd
- NetNumber Inc.
- Virtela Technology Services Inc.

SDN Controller Software Vendors

SDN can be employed by service providers as a means of implementing a Network Functions Virtualization Infrastructure (NFVI) for cloud IaaS services and as a NFVI within their access and core networks. Some SDN implementations provide flow mapping functions that steer traffic flows to VNFs in the proper sequence.

Representative Members:

- Adara Networks Inc
- ConteXtream Inc.
- NEC

NFVI Providers

Value Proposition: The members of this category provide the virtual networking infrastructure including Virtual Switching (Open vSwitch, Linux Bridge), Virtual Networking (IP Forwarding, Virtual Routing, Filtering, NAT, Link Aggregation, etc.), and Overlays such as VXLAN, VLAN, GRE, etc. for multi-tenancy. The NFVI also includes physical NIC poll mode drivers for outside communication and virtual NIC host drivers (such as Virtio) for communication with VMs.

Representative Members:

- 6Wind
- BTI Systems
- Wind River

Orchestration Software Vendors

Orchestration generally involves the assembly of various software components (e.g., VNFs) and hardware components of the end-to-end infrastructure to deliver and manage a defined service. Orchestrators often employ layers of abstraction that facilitate the automation of provisioning, configuration, optimization, and other repetitive operational tasks. Orchestration is another potential solution for mapping flows through VNFs and can be deployed either in conjunction with SDN or independently of SDN.

Representative Members:

- Anuta Networks Inc.
- Cadzow Communications
- CENX Inc.

Network Monitoring, Management and OSS/BSS Vendors

Value Proposition: The members of this category of the ecosystem will provide management functionality that extends to virtualized infrastructures and VNFs and integrates that functionality into a broader management structure.

Representative Members:

- NetScout
- Amdocs Software Systems Ltd
- Comptel Corporation
- Comverse Network Systems Europe B.V.
- EMC
- MetraTech Corp

Hypervisor Vendors

Value Proposition: These vendors provide the VMs, vSwitches, and the hypervisor vSwitch APIs for third party vSwitches that are a key components of SDN and NFV infrastructure solutions.

Representative Members:

- Citrix Systems Inc
- Oracle
- Virtual Open Systems

Test Equipment Vendors and Test Services

Value Proposition: The members of this category either provide products that enable equipment manufacturers and others to test NFV solutions, or they provide the testing as a service.

Representative Members:

- QualiSystems
- European Advanced Networking Test Center
- JDSU Deutschland GmbH
- Spirent Communications
- Tektronix GmbH Co KG
- Yokogawa Europe B.V.

Open Source Communities

Value Proposition: These organizations create working prototypes of key SDN and NFV functionality. Part of the value proposition of these communities is that the prototypes that they develop help to better define the underlying technological challenges. Another part of their value proposition is that the prototypes they create are often used as the basis of commercial products which because they are based on open source solutions can potentially be brought to market more quickly and more cost effectively.

As referred to earlier, the open source community is also very active in the development of SDN.

Representative Members:

- OpenSwitch
- OPNFV
- OpenStack
- OpenDaylight
- ON.Lab

Standards Bodies and Related Communities

Value Proposition: Some of the members of this category develop use cases, architectures and drive POCs. Other members of this category create standards for protocols such as OpenFlow or VXLAN. These standards form the basis for enabling products from disparate vendors to interoperate.

As was also referred to earlier, many members of this category are also very active in the development of SDN.

Representative Members:

- ETSI
- 3GPP
- MEF
- ATIS
- IETF
- TM Forum

Key Vendors

Below is a profile of the sponsoring vendors that focuses on where they fit in the ecosystem, the value add that they provide and the proof points of that value add.

NetScout

Where do you fit in the SDN and/or NFV ecosystem?

NETSCOUT fits into the SDN and NFV ecosystem as a leading real-time service assurance solution provider. According to NETSCOUT, while many of the traditional application and network performance management vendors claim to offer solutions for virtualized and hybrid environments, the reality is that most have just repackaged their existing solutions without the ability to monitor traffic within complex virtual environments. As a result, they have many blind spots when it comes to understanding the root cause of service performance issues. Monitoring traffic into and out of a virtual environment is “table stakes,” but falls well short of meeting the needs of the IT department or service provider. NETSCOUT believes that in order realize the full potential of SDN/NFV CapEx and OpEx efficiencies, organizations need a comprehensive service delivery monitoring solution that expands visibility within virtualized infrastructures and between virtual machines.

Truly managing and understanding the user experience in physical, virtual and hybrid infrastructures requires the ability to have an end-to-end view of the network and services. NETSCOUT stated their belief that using traditional monitoring tools will not suffice as these tools do not provide a common situational awareness. Consequently, users pay the penalty for this as service availability and performance are compromised.

NETSCOUT believes that by extending service assurance to SDN/NFV ecosystems, enterprises and service providers can accelerate digital transformation initiatives. The best way to do that is by proactively collecting, organizing and contextually analyzing traffic data in real time. By reducing service downside risk through continuous monitoring of traffic-based data and real-time analysis, it’s possible for organizations to compete and innovate with confidence.

What is your value add?

NETSCOUT believes that its’ value add in virtualized and SDN/NFV environments is exactly the same as in physical environments, magnified by the extent and breath of the new challenges.

That value add includes:

- Adaptive Service Intelligence^{TM} (ASI) technology to understand the interrelationships and dependencies of the physical and virtual service delivery environment
- Reduced MTTR with proactive service triage
- Enhanced IT efficiencies through a common operational view
- Scalable service assurance architecture

NETSCOUT stated that their solution unlocks the power of traffic-based data to gain real-time insight and to deliver service assurance for the most demanding physical, virtual and hybrid networks. NETSCOUT justifies that statement by pointing out that today ASI technology runs on NETSCOUT’s physical and virtual Intelligent Data Sources. This enables NETSCOUT to extend the monitoring of both enterprise and carrier-scale service delivery infrastructure into both virtual and hybrid environments. NETSCOUT added that their solution provides a holistic view of the entire data center – including VMware NSX resources. These environments include physical and virtual application workloads that exchange greater volumes of traffic between themselves and that also experience a higher risk of service degradations. In both a physical and virtual environment, service degradation often results in a lower quality of end-user experience and may result in increased churn for service providers.

In a virtual environment, there is an additional challenge of collecting management data while having minimal impact on compute and networking resources. This puts significant pressure on service assurance solutions to be as efficient as possible with consuming both compute and networking resources while exchanging monitoring information. NETSCOUT stated that they excel in this area with their ASI technology and the ability to capture, process, and create highly scalable metadata in real time as IP traffic traverses physical or virtual links.

What are the proof points?

NETSCOUT stated that their extensive customer base combined with its integration with VMware's NSX environment means that the company is uniquely positioned as a service assurance leader in the virtual and hybrid environments. A 2015 survey of NETSCOUT customers conducted by TechValidate, a leading "voice of the customer" researcher, revealed the following:

- Four out of five customers cut Mean Time to Knowledge (MTTK) by 80% or more and reduced operational expenses
- 91% of customers get real-time, actionable traffic-based intelligence with ASI technology
- 100% of customers surveyed improved the identification of network issues

In the service provider space, NETSCOUT's nGeniusONE™ Service Assurance platform provides wireless, cable, and wireline network operators with end-to-end network and service performance management. The nGeniusONE platform provides both enterprises and operators a single monitoring infrastructure for hybrid environments of today and the all virtual environments of tomorrow. In addition, NETSCOUT has been at the forefront of working with operators as they test and trial virtual infrastructure.

Cisco

Where do you fit in the SDN and/or NFV ecosystem?

Cisco believes it is a leader in the SDN and NFV ecosystem. Cisco markets an array of SDN solutions, addressing the requirements and use cases of a broad spectrum of customers across a wide range of markets. Cisco has stated its extensible environment includes a broad and growing number of ecosystem technology partners including compliant network, security and services devices, monitoring, analytics and DevOps solutions, as well as cloud automation platforms.

What is your value add?

According to Cisco, their SDN solutions offer a complete portfolio providing choice in automation and programmability for customers. Cisco stated that their solutions are based on open APIs, standards and a broad ecosystem for three approaches: programmable networks, programmable fabrics and a turnkey approach with Cisco Application Centric Infrastructure (ACI). Cisco believes that this approach enables customers to choose the implementation option that best meets their IT and business goals by extending the benefits of programmability and automation across the entire Cisco Nexus switching portfolio.

Cisco ACI: Based on an application centric policy model, ACI provides automated, integrated provisioning of both underlay and overlay networks, L4-7 services provisioning across a broad set of ecosystem partners, and extensive telemetry for application level health monitoring.

Programmable Fabric: Cisco stated that it is providing scale and simplicity to VXLAN Overlays with a standards based approach, based on a Multipoint BGP EVPN Control Plane, on Nexus switches to scale out VXLANs, and simplified provisioning and management of these switches via an overlay management and provisioning system called Virtual Topology System (VTS).

Programmable Network: Cisco is offering programmability and accessibility to the Nexus switches to enable them to automate provisioning and configuration, as well as integrate with orchestration tools.

What are the proof points?

Customer Case studies

- Bowling Green University http://www.cisco.com/web/about/success-stories/docs/bowling-green.html
- Du: http://www.cisco.com/web/about/success-stories/docs/du.html
- KPIT Technologies: http://www.cisco.com/web/about/success-stories/docs/kpit.html
- Qatar University: http://www.cisco.com/web/about/success-stories/docs/qatar-university.html
- More case studies here: http://www.cisco.com/c/en/us/solutions/data-center-virtualization/application-centric-infrastructure/customer-case-study-listing.html

ROI Case Studies / Analyst reports

- IDC: Pulsant Delivers Agile and Cost-Effective Hybrid Cloud Services with Cisco ACI
- IDC: Symantec Delivering on Its Strategic Vision with Next Generation secure Data Center powered by ACI
- Forrester Report: The Total Economic Impact TM Of Cisco Application-Centric Infrastructure (ACI)

Industry Awards

- 2015 Best of Interop | Category SDN - Cisco APIC: http://www.networkcomputing.com/interop/prizewinning-it-products-best-of-interop-2015-/d/d-id/1320139?image_number=8 http://www.interop.com/lasvegas/special-events/best-of-interop-awards.php
- 2015 Winner Excellence in SDN – Cisco APIC http://www.nfvzone.com/topics/newsfeed/articles/400677-winners-the-2015-sdn-excellence-award-announced.htm
- 2014 Winner infoTECH Spotlight Data Center Excellence Award infoTECHhttp://news.tmcnet.com/news/2014/12/18/8128268.htm
- 2014 Best of Interop 014 | Category: Data Center - Cisco Nexus 9516 http://www.interop.com/lasvegas/special-events/best-of-interop-awards.php?itc=we_ilv_le_ilv_drp_text
- 2014 Best of Interop Tokyo – Cisco ACI and Nexus 9000 Best of Interop Tokyo 2014 – Cisco ACI and Nexus 9000

Tech Review/ Articles Customer Business Outcomes

- InfoWorld Tech Review: “Cisco ACI shakes up SDN” – Nov 2015 http://www.infoworld.com/article/3003903/sdn/review-cisco-aci-application-centric-infrastructure-shakes-up-sdn.html?nsdr=true
- Networkworld Article – Oct 2015: Cisco outpacing VMware 2:1 in SDN http://www.networkworld.com/article/2989946/cisco-subnet/cisco-outpacing-vmware-2-1-in-sdn.html
- Forbes Insight Article – Sep 2015 http://www.forbes.com/forbesinsights/cisco_aci/index.html

Supporting information

- Blog Dec 2015 : http://blogs.cisco.com/news/executing-on-our-vision-and-strategy-the-future-of-networking-for-an-application-centric-world
- Blog Nov 2015 : http://blogs.cisco.com/news/co-innovating-with-ecosystem-partners-to-deliver-truly-integrated-solutions-for-our-customers
- http://blogs.cisco.com/news/danske-bank-our-1000th-cisco-aci-customer
- Blog Sep 2015: http://blogs.cisco.com/datacenter/dc-sdn-strategy
- Cisco Ustream channel Oct 2015: Cisco Global Editors Conference on October 6, 2015 http://www.ustream.tv/recorded/74933464

Websites

- Cisco ACI www.cisco.com/go/aci
- Cisco VTS www.cisco.com/go/vts
- Cisco Developers / Open NXOS https://developer.cisco.com/data-center

Masergy

Where do you fit in the SDN and/or NFV ecosystem?

Masergy is a network service provider offering cloud services, virtualized network functions, network monitoring and management capabilities

What is your value add?

Masergy is a 15-year-old company that designs, implements and manages private wide area networks for global enterprises. Masergy has evolved its offerings to include secure Internet and cloud connectivity services as well.

According to Masergy, they have been applying the principals behind software defined networking for a decade - - even before SDN became a widely accepted architectural concept and much-used acronym. Masergy stated that they designed their network fabric to be programmable and adaptable so that it can deliver a customized solution to each customer.

All of Masergy's offerings are built upon its Software Defined Platform. Their goal in making their network programmable is to reduce complexity for their customers relative to the management of their global networks.

Masergy stated that they provide software controls that lets customers make changes to their networks in real time. One such capability is bandwidth on demand, which lets companies increase or throttle back network capacity as application processing requirements change.
According to Masergy, this is becoming a critically important issue as enterprises introduce new, bandwidth-intensive applications into their environments.

An example of the use of bandwidth on demand comes from one of Masergy's customers - PRGX. The company is a leading provider of accounts payable recovery audit services to more than three quarters of the top 20 retailers. PRGX is a global company that also works with enterprises in the oil, gas, pharmaceutical, manufacturing and construction industries. Much of its business involves applying big data analytics on its 3,000 terabytes data to help customer recover unrealized revenue.

Network scalability and flexibility is essential for the company's 1,400 employees operating in 30 countries. A high performance network enables PRGX to apply bandwidth at locations that are doing some serious number crunching. When these tasks are complete, bandwidth can be dialed down to handle routine application processing requirements. Network administrations can make these changes on-the-fly from their Masergy-provided web portal and mobile app using Masergy's Intelligent Service Control. The availability of such network flexibility enabled PRGX to roll out a new analytics service to customers that never would have been possible with its previous network and its inherent performance limitations.

Masergy has implemented NFV in its managed network services as a way to deliver routing, firewall and session border control as cloud services, on-premises in their network interface device (Masergy Intelligent Bridge) and in software that can be pushed down to the MIB without any technician involvement. According to Masergy, this is a boon for remote offices that often lack on-site IT personnel. New routing tables can be pushed down to the local device rather than having to ship a piece of equipment and a technician to install it at the site.

What are the proof points?

Masergy's customers also include Amgen, Brocade, Cornell University, Dolby, E*Trade, Panavision, Pepsico, Tesla Motors, and W.R. Grace and other global brands.
Masergy worked with its hardware provider Overture Networks (now part of ADVA Optical Networking) on an NFV proof of concept project in 2015. The effort resulted in the Global Telecoms Business 2015 Innovation Award.

Other awards include:

- Masergy's Cloud f(n) Router Takes Gold In The 7th Annual 2015 Golden Bridge Awards
- Masergy's Cloud f(n) Router Receives 2015 Internet Telephony TMC Labs Innovation Award
- Masergy Named 2015 Light Reading Leading Lights Finalist For Most Innovative SDN Deployment Strategy
- Masergy's Cloud-Based Router Selected as NFV Pioneer Award Winner
- Masergy's Network Sensor Honored With SDN Award

Sonus

Where do you fit in the SDN and/or NFV ecosystem?

Sonus' primary focus is on migrating real-time communications into the NFV ecosystem. From past experience, one might conclude that Sonus is a Session Border Controller (SBC) company and therefore their primary focus would be on delivering a SBC VNF, but Sonus views that as being insufficient. The Sonus vision encompasses multiple dimensions including virtual session border control, signaling and policy functions, global licensing models, cloud-based toolchains, lifecycle management of VNFs, microservices architectures for NFV applications, and integration with software-defined network intelligence.

For Sonus, NFV represents a new architecture that unlocks powerful new capabilities. Their strategy for NFV includes:

- Ensuring the ability to dynamically instantiate applications based on business triggers, requiring VNFs to be delivered using fully automated, elastic scaling models backed by licensing based on network capacity versus one tied to instances.
- Recognition that the NFVI will provide new capabilities that can be taken advantage of for increased VNF efficiency such as resilient, scalable cloud-DBs and storage models, dynamic VM scaling capabilities and health monitoring.
- Recognition that as applications move from a static hosting model to a dynamic model the toolchains supporting the VNF become just as important as the VNFs themselves as customers make vendor decisions.

According to Sonus, with their solutions it will be possible to assess VNF performance, troubleshoot, and monitor application SLAs in a cloud environment without increased effort compared to the classic method of attaching to a static node. This is critical because if this is not done right, there is a risk of having innovative VNFs that are too difficult to deploy.

Another key part of the Sonus NFV transformation strategy is recognition that SDN and more specifically programmable IP transport goes hand-in-hand with this transformation. According to Sonus, this creates a unique opportunity for them to drive collaboration between the real-time session control layer and the IP transport layer via SDN. They envision the industry moving from a model of independent application, session and transport to a model where the session layer can now broker the wants and needs (i.e. the SLA) of the application. By integrating session layer intelligence with network intelligence, data center network connectivity and cloud-based service delivery will be optimized. As an example, integration between Sonus, SBC VNF and Sonus SDN enable an end-to-end delivery strategy for mission critical real-time applications which are coincidentally most sensitive to IP transport performance.

Sonus has chosen to partner with market-leading orchestration and NFVI vendors to facilitate VNF implementations. To date, this includes SBC NFV validation testing with Juniper, Overture, Alcatel-Lucent, HPE, and Dorado. Sonus also participated in the recent Light Reading/EANTC interop evaluation event, the world's first independent interoperability evaluation of NFV infrastructure, focused (in Phase 1) on multivendor NFV infrastructure-to-virtual network function (VNF) interoperability.

What is your value add?

Sonus has a long history in the management and delivery of one of the most demanding segments of network applications - real-time communications. From there it has been a very logical step for them to broaden the scope of their service delivery architecture to include the IP transport and non-voice applications. This is what Sonus believes that they have architected by including SDN as part of their NFV strategy.

While voice, and real-time communications in general, are obvious beneficiaries of an SDN-managed solution, there are certainly multiple applications that an enterprise may consider as mission-critical to their business success. While it will vary by enterprise, some examples are Salesforce or customized CRM, customer portals and web-based retail applications, license servers, or accounting and billing applications.

Sonus' strategies to optimize the user experience for real-time communications are equally applicable to these mission-critical applications. For Sonus this means that they take a holistic view of how their technology allows "one network" – rather than having a voice, video, and data networks, where that one network intelligently optimizes packet flows based on application prioritization.

What are the proof points?

The following are examples of customer deployments:

- Tier 1 service provider using SBC SWe to offer a virtual CPE solution to their Enterprise customers
- Cloud-based Unified Communications provider deploying SBC SWe in the Amazon cloud, enabling extremely rapid, low-cost instantiation without the need for on premise deployments
- A Fortune 500 financial corporation using SBC SWe for SIP trunking
- Telstra International (SE Asia regional service provider) has deployed VellOS (Sonus software-defined cloud exchange network solution) to optimize their data center interconnections and leverage that infrastructure into a revenue generating Network-as-a-Service.
- Internet Solutions (South Africa) is using VellOS to optimize their data center interconnections and simplify the management of data center connectivity
- State Street Bank, a global financial services company, is using VellOS to optimize data center connectivity and ensure business services continuity and security compliance

Awards in 2015:

- Information Week Elite 100 List of Top Technology Innovators across US
- Computer Technology Review Most Valuable Networking Product for VellOS
- Internet Telephony Unified Communications Product of the Year for SBC SWe

Radware

Where do you fit in the SDN and/or NFV ecosystem?

Radware SDN applications improve application security, performance and availability by integrating ADC and security intelligence with the SDN to collect data and optimally forward traffic to enhance 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 integrate with the SDN application control plane and interact with the SDN controller using dedicated SDN drivers to work together with the Radware technologies by using the Radware API to use specific data collection drivers throughout the application infrastructure.

According to Radware, with their SDN solution, applications such as ADCs 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. The company states that network services can scale to support larger networks at a lower capital and operational cost. Radware also stated that by building SDN applications that continuously interact with the SDN control plane and program the network by leveraging the Radware Virtual Application Delivery Infrastructure (VADI) architecture – which enables pooling of disperse resources to operate uniformly, Radware enables an EveryWare, applications available anywhere and everywhere, network service paradigm.

What is your value add?

According to Radware, key benefits of their 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 under all network conditions is entitled to advanced services.
- **Simpler implementation** of network services allows improved operational efficiency of network management improving application agility. 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 that are protected or load balanced. Offer uniform services throughout the SDN to enable an elastic application-centric infrastructure.
- **Easier operational control** – changing and managing security and ADC functionality becomes simpler through centralized operational deployment models. Not only does SDN streamline network operations, but Radware SDN applications streamline network service operations. Open Radware APIs allow orchestration systems to improve the overall control and automation of network services.

DDoS Protection as a Native SDN Application

Radware's 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 steering mechanism that utilizes the programmable characteristics of the software defined network elements for attack mitigation. Designed as part of the Radware SDN application framework, DefenseFlow delivers a security control plane and operates in traditional network environments while enabling customers to migrate to the customer's future, SDN-based networks.

According to Radware, legacy DDoS protection solutions that make use of scrubbing centers are costly: they need hardware detectors in every network location, BGP for traffic diversion, and GRE tunnels to forward the traffic to its designated network service destination. With SDN, a DDoS protection solution turns into a software application that adds intelligence to the network and does not require additional hardware, BGP, or GRE operations.

Radware stated that 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 an 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 appropriate mitigation device.
- Centralized security control plane including control as part of Radware's Attack Mitigation Network (AMN)

SDN & NFV for a Scalable Application Delivery Network

The NFV initiative was formed in order to enable the standardization of network services by leveraging commercial off-the-shelf (COTS) hardware and running advanced network functions software on them. Towards that end, Radware offers Alteon VA for NFV, which according to Radware is the industry's first and highest performing ADC designed from the ground up to run in NFV environments. Targeted mainly at carriers but also towards large enterprises looking to leverage the NFV architecture, Alteon NFV provides a unique value proposition consisting of CAPEX/OPEX reduction, vendor agnostic technologies, high performance, enhanced scalability, orchestrated elasticity, and improved 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 network environment. ElasticScale can be utilized for service provider internal services, managed services to end customers, and can help providers adopt network functions virtualization paradigms.

According to Radware, 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 ADC product line
- Support for leading hypervisors (Xen/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 solutions

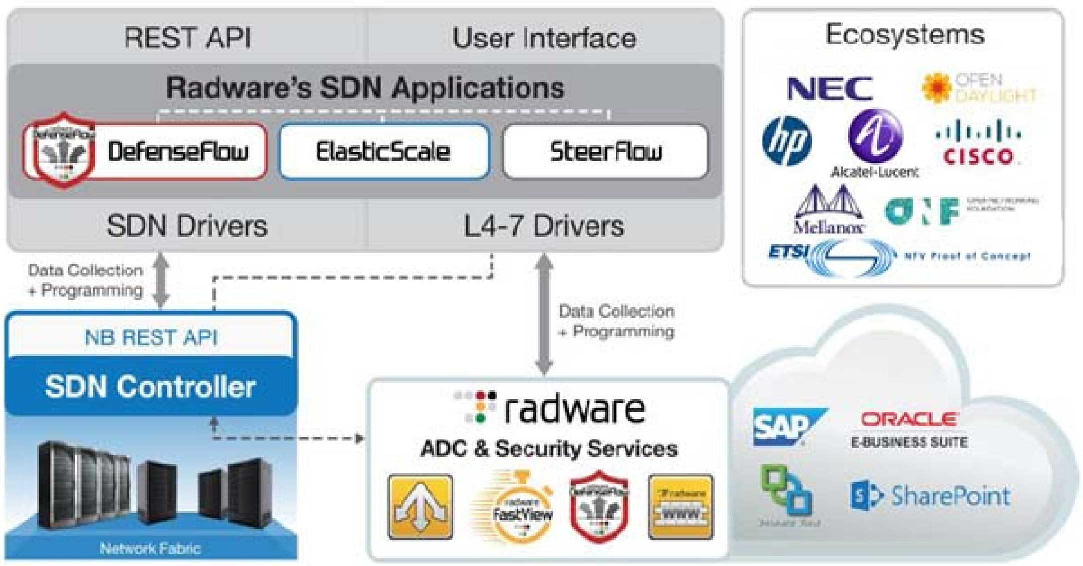

<u>What are the proof points?</u>

The SDN eco-system is a critical focus for Radware. Through partnerships with the industry's leading SDN consortiums and vendors, Radware ensures 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).

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.

Conclusions

The following is a summary of the conclusions that were reached in the preceding sections of **The Guide**.

- SDN remains stuck on the edge of the chasm and will be there for at least another year or two.
- Open Source projects will likely accelerate the adoption of SDN.
- The ONF NBI initiative has the potential to seamlessly interconnect disparate SDN controllers.
- Very few IT organizations have ruled out the use of OpenFlow.
- By a small margin, IT organizations perceive the overlay-based SDN model will provide more value over the next two years than will the fabric-based model. However, many IT organizations are yet to form an opinion.
- At least for now, ONOS is targeted at service providers.
- The ONOS community has expanded to include vendors.
- The ONOS project is part of the Linux Foundation.
- ODL's Lithium release contains a range of sophisticated functionality.
- ODL's membership has expanded to include service providers and enterprises.
- There is currently as much interest in either implementing SDN in the WAN or using a SDN-based WAN service as there is in implementing SDN in the data center.
- The two primary factors driving SDN deployment in the data center are supporting the dynamic movement, replication and allocation of virtual resources and easing the administrative burden of configuration and provisioning.
- There is a wide range of significant inhibitors to the deployment of SDN in the data center.
- There are a number of significant drivers of SDN deployment in the WAN.
- The two primary factors driving SDN deployment in the WAN are easing the administrative burden of configuration and provisioning and better utilizing network resources.
- Three of the major inhibitors to the deployment of SDN in the WAN are concerns about how to integrate SDN into the rest of the infrastructure, the lack of a compelling business case and concerns about security vulnerabilities.
- Two of the major inhibitors to the deployment of SDN in branch and campus networks are concerns about how to integrate SDN into the rest of the infrastructure and the lack of a compelling business case.

- Overall, the two primary factors that are driving the implementation of SDN are easing the administrative burden of configuration and provisioning and increasing the utilization of network resources.

- Neither reducing complexity nor reducing CAPEX are significant drivers of deploying SDN.

- The primary factor inhibiting the adoption of SDN is the concerns that organizations have about how they would integrate SDN into the rest of the infrastructure.

- One of the implications of adopting SDN is that is increases the need for a DevOps model.

- SDN creates security opportunities and security challenges.

- SDN creates both management opportunities and management challenges.

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

- Network management organizations need tools that enable them to be able to dynamically discover, procure, allocate and reconfigure network resources.

- 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.

- 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.

- Almost a third of the survey respondents believe that over the next two years the ongoing adoption of software-based IT functionality will have either a significant or very significant impact on the structure of their IT organization.

- Over a quarter of the survey respondents believe that over the next two years the ongoing adoption of software-based IT functionality will have either a significant or very significant impact on their jobs.

- Half of IT professionals believe that NFV has either significant or very significant relevance to enterprise IT architectures.

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

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

- While only a modest number of IT organizations have implemented NFV in a production network, a large percentage of IT organizations are currently in varying stages of analyzing NFV.

- By a wide margin, the primary factor driving interest in NFV is the reduction in the time it takes to deploy new services.

- The biggest inhibitors to the broad adoption of NFV are:

- The lack of a compelling business case;

- Concerns about end-to-end service provisioning;

- Concerns about security vulnerabilities;

- The immaturity of the current products;

- The need to significantly reskill our employee base.

- Within a few years, the majority of IT organizations are likely to have made a significant deployment of virtualized L4 – L7 functionality.

- By almost a 2:1 ratio, IT professionals think that open source communities will have more of an impact on the evolution of NFV than SDOs will.

- Organizations should place a preference on acquiring VNFs that were designed in a modular fashion.

- To the degree possible, organizations need to adopt an architecture that can evolve as the enabling technologies change without requiring a major overhaul.

- In order to achieve maximum performance, organizations should focus their attention on VNFs that were designed to run effectively in a software-centric environment.

- Organizations need to recognize that solutions that are based on open source solutions will potentially evolve quickly and potentially have a high degree of interoperability.

- Organizations should plan for, trial and adopt NFV and SDN in an integrated fashion.

- Organizations should ensure that whatever NFV related functionality it implements fits with the broader view of a SDDC.

- The implementation of NFV enables organizations to rethink HA.

- Organization should monitor whether or not Moore's law is coming to an end and if it is, they need to adjust their plans to move from a hardware-centric approach to a software-centric approach.

- Two thirds of IT organizations have made little or no progress towards the development of a NFV architecture.

- While some organizations are making significant progress towards the development of a NFV architecture, the majority are not.

- There is broad recognition on the part of IT organizations that the adoption of NFV creates new management challenges.

- The vast majority of IT organizations have made little or no progress relative to determining how they will respond to NFV-related management challenges.

- Over the next year the vast majority of IT organizations will spend at least a modest amount of time working on developing an approach to how they will respond to NFV-related management challenges.

- There is significant interest in orchestration, but only a very small minority of IT organizations are using an orchestration platform in production.

- Roughly a third of IT organizations believe that over the next two years that the adoption of NFV is likely to have a significant or very significant impact on the structure of their organization.

- Over the next two years the adoption of NFV is likely to have a significant or very significant impact on the skill base of roughly 40% of all it professionals.

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.

For editorial and sponsorship information, contact Jim Metzler or Steven Taylor. The Webtorials Editorial/Analyst Division is an analyst and consulting joint venture of Steven Taylor and Jim Metzler.