

The 2015 Guide to SDN and NFV

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

Platinum Sponsors:



Gold Sponsors:



Produced by:



TABLE OF CONTENTS

EXECUTIVE SUMMARY.....	1
SOFTWARE DEFINED NETWORKING (SDN)	1
<i>Introduction</i>	1
<i>SDN Use Cases</i>	2
<i>The Operational Implications</i>	2
NETWORK FUNCTIONS VIRTUALIZATION (NFV).....	5
<i>Introduction</i>	5
<i>Use Cases and Proof of Concept (POC)</i>	6
<i>The Operational Implications</i>	8
THE SDN AND NFV ECOSYSTEM	10
SOFTWARE DEFINED NETWORKING (SDN).....	11
INTRODUCTION	11
<i>Definition of SDN</i>	11
<i>Context for SDN</i>	12
<i>Status of SDN Adoption</i>	12
<i>The SDN Architecture</i>	13
<i>The Northbound Interface</i>	14
<i>Architectural Distinctions between Approaches</i>	15
<i>The Overlay and the Underlay Model</i>	16
<i>Service Chaining</i>	17
<i>The OpenDaylight Consortium</i>	17
<i>The Relationship Between SDN and NFV</i>	19
SDN USE CASES	20
<i>Drivers and Inhibitors</i>	20
<i>SDN Deployment Plans</i>	21
<i>Data Center</i>	22
<i>WAN</i>	25
<i>Campus</i>	25
THE OPERATIONAL IMPLICATIONS.....	28
<i>Security</i>	28
<i>Cloud Orchestration</i>	29
<i>Management</i>	31
<i>Organizational Impact</i>	34
NETWORK FUNCTIONS VIRTUALIZATION (NFV).....	50
INTRODUCTION	50
<i>Background</i>	50
<i>ETSI</i>	50
<i>TM Forum</i>	52
<i>Internet Engineering Task Force (IETF)</i>	53
<i>Open Platform for NFV (OPNFV)</i>	53
<i>Relationship between SDN and NFV</i>	54
<i>Status of NFV Adoption</i>	56
USE CASES AND PROOF OF CONCEPT.....	60

<i>ETSI NFV Use Cases</i>	60
<i>TM Forum Catalyst POCs</i>	63
<i>Private POCs</i>	65
THE OPERATIONAL IMPLICATIONS	66
<i>Performance Limitations</i>	66
<i>End-to-End Management</i>	68
<i>The Organizational Implications</i>	71
THE SDN AND NFV ECOSYSTEM	82
THE SDN ECOSYSTEM	82
<i>Merchant Silicon/Chip Vendors</i>	82
<i>HyperScale Data Centers</i>	82
<i>Telecom Service Providers</i>	82
<i>Switch Vendors</i>	83
<i>Network and Service Monitoring, Management and Automation</i>	83
<i>Providers of Network Services</i>	84
<i>Testing</i>	84
<i>Standards Bodies and Related Communities</i>	84
<i>Providers of SDN Controllers</i>	85
<i>Providers of Telcom Service Provider's Infrastructure/ Optical Networking</i>	85
<i>Server Virtualization Vendors</i>	85
THE NFV ECOSYSTEM	86
<i>Telecom Service Providers</i>	86
<i>Network Systems and Electronic Equipment Vendors</i>	86
<i>Merchant Silicon/Chip Vendors</i>	87
<i>Virtualized Network Service and Cloud Service Vendors</i>	87
<i>SDN Controller Software Vendors</i>	87
<i>NFVI Providers</i>	87
<i>Orchestration Software Vendors</i>	88
<i>Network Monitoring, Management and OSS/BSS Vendors</i>	88
<i>Hypervisor Vendors</i>	88
<i>Test Equipment Vendors and Test Services</i>	89
<i>Standards Bodies and Related Communities</i>	89
KEY VENDORS	90
<i>NetScout</i>	90
<i>Cisco</i>	92
<i>A10</i>	93
<i>Alcatel-Lucent</i>	94
<i>Meru Networks</i>	96
<i>6WIND</i>	97
<i>Dell</i>	99
<i>EMC</i>	100
<i>Citrix</i>	101
<i>Nuage Networks</i>	103
<i>Pica8</i>	105
CONCLUSIONS	107

Executive Summary

Software Defined Networking (SDN)

Introduction

This e-book is based in part of two surveys that were administered in September and October of 2014. One of the surveys focused on SDN and the other on NFV. Throughout this executive summary, the respondents to those surveys will be referred to respectively as The SDN Survey Respondents and The NFV Survey Respondents.

The responses to the SDN survey indicated that the general familiarity with SDN has increased significantly over the last year and that while the percentage of IT organizations that have implemented SDN in production is still small, it has increased somewhat significantly over the last year. The SDN Survey Respondents also indicated that the percentage of IT organizations who have SDN in production will likely increase somewhat over the next year, but the percentage will remain small.

The e-book identified a number of changes that have occurred with SDN over the last year. One thing that has changed is that most of the discussion around whether or not an overlay network virtualization solution is indeed SDN has gone away. Today, most IT professionals regard an overlay solution as being a form of SDN. The e-book discusses the pros and cons of the overlay and the underlay SDN models and presents market research that indicates that by a small margin that The SDN Survey Respondents believe that the underlay model will provide more value over the next two years.

Another change that has occurred in the SDN landscape within the last year is that the Open Networking Foundation (ONF) established the Northbound Interface (NBI) working group with the goal of eventually standardizing SDN's northbound interface. Sarwar Raza, the chairman of the working group, is quoted as saying that standardization was not a short term goal of the group and that "Our goal in the next year is to formalize the framework along with the information and data models and then iterate some with code before we even start a standards discussion." The NBI working group intends to work with one or more open source initiatives to develop working code for the NBIs and the group aims to work on standardization at an appropriate time in the future.

Another change in the SDN landscape that is discussed in the e-book is that in February 2014 the OpenDaylight community issued its first software release, called Hydrogen and in September 2014 issued its second software release called Helium. A number of vendors have announced their intention to use the OpenDaylight solution as the basis of their SDN controller. This creates the potential for SDN solutions based on OpenDaylight solutions to reach critical mass in the near term and hence accelerate the adoption of SDN.

The majority of The SDN Survey Respondents indicated that they thought that SDN and NFV are complimentary activities and a quarter of the respondents indicated that they thought that in at least some instances that NFV requires SDN. That second school of thought is in line with the ONF who in March of 2014 published a white paper that included uses cases that the ONF believes demonstrate how OpenFlow-enabled SDN can meet the need for automated, open, and programmable network connectivity to support NFV.

SDN Use Cases

The SDN Survey Respondents indicated that a wide range of factors were driving their interest in SDN including the desire to better utilize network resources and to perform traffic engineering with an end-to-end view of the network. However, very few of the respondents indicated that they thought that SDN would help them reduce CAPEX or reduce complexity. The SDN Survey Respondents also indicated that a wide range of factors were inhibiting their interest in SDN. 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 of the key inhibitors, such as the lack of a compelling business case, need to be addressed or they will continue to impede SDN adoption.

The SDN Survey Respondents indicated that over the next two years that the primary focus of their SDN deployment will likely to be in the data center. However, they expressed considerable interest in deploying SDN in the WAN as well as in branch and campus networks. In addition, when asked to look forward three years, The SDN Survey Respondents indicated that three years from now that they will have deployed SDN pervasively in their data centers and that they will also have made significant SDN deployment both in their WAN and in their campus networks.

The e-book discussed a number of SDN use cases. The WAN use case that was discussed was how Google has deployed SDN to connect its data centers and as a result of that deployment, has driven its network utilization to 95%. The campus use cases that were discussed were:

- Dynamic QoS and traffic engineering;
- Unified wired and wireless networks;
- QoS management for Microsoft Lync across wired and wireless networks;
- Personal Bonjour;
- Roll based access.

The data center use cases that were discussed were:

- Virtual machine migration;
- Service chaining;
- Security services;
- Load balancer services;
- Software defined clouds;
- Cloud hosting.

The Operational Implications

Thirty five percent of The SDN Survey Respondents indicated that SDN will enable them to implement more effective security functionality and 12% of The SDN Survey Respondents indicated that concerns about possible security vulnerabilities is a significant inhibitor to SDN deployment. As the e-book discusses, 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.

The e-book discusses some of the security challenges including:

- 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 e-book describes OpenStack and points out that orchestration engines such as OpenStack are important to both SDN and NFV. As explained in the e-book, 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 or a VN;
- Attach Network Services to a VM or chain Network Services between VMs.

In spite of the importance of orchestration, only a small minority of The SDN Survey Respondents indicated that their organization had a well thought out strategy for how they would do orchestration.

Similar to the situation with security, the e-book shows how management is a double edged sword. Fifty three percent of network organizations believe that SDN will ease the administrative burden of management tasks such as configuration and provisioning while 13% of network organizations believe that concerns about how to manage SDN is a significant inhibitor to SDN deployment.

The e-book 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:

- The combination of physical and virtual infrastructure and dynamically changing resources requires a more holistic approach to instrumentation, consolidation of individual datasets, and analysis of the consolidated dataset in a service contextual fashion.
- 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 e-book 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.

The Survey Respondents were asked how they thought that the SDE movement would likely impact their organization. Their answers included:

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

The Survey Respondents were also asked how they thought told that the SDE movement would likely impact their jobs. Their answers included:

- The way to design, implement and troubleshoot networks will change a lot;
- The job will require new skill sets in general and more programming knowledge in particular;
- There will be new security requirements;
- New architectures will need to be developed;
- There will be a lot of re-training and re-trenching.

Network Functions Virtualization (NFV)

Introduction

NFV is being driven by a number of different types of players who are described in the e-book. This includes industry organizations such as the TM Forum and ETSI, open source communities such as OPNFV and traditional standards development organizations such as IETF.

As described in the e-book, early 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. As of November 2014, the ZOOM team has delivered an assessment of how virtualization impacts SLAs and is currently working on information and policy models, NFV preparedness, and a set of operational support system (OSS) design principles needed for NFV adoption to become widespread.

The ETSI NFV ISG has identified nine NFV use cases and is currently driving 25 POCs. The ETSI NFV ISG was established with a two year life span that expires in January 2015. In late July and early August 2014 the NFV ISG met in Santa Clara, CA. At that meeting the [primary objectives of NFV Phase 2](#) were identified. Whereas ETSI characterizes Phase 1 as being the Requirements Phase, ETSI characterizes Phase 2 as being the Implementation Phase. The objectives of Phase 2 include building on the achievements that were made in the first two years of the ISG and consist of an enhanced focus on interoperability, formal testing, as well as working closer with projects developing open source NFV implementations. In addition, the NFV ISG also released nine draft NFV documents for [industry comments](#) and published a publically available document that summarizes the key concepts that are contained in those [documents](#).

In September 2014 the Linux Foundation announced the founding of the [Open Platform for NFV Project](#) (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.

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 over forty active 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.

In spite of the fact that the vast majority of The NFV Survey Respondents believe that NFV is applicable in both an enterprise and a service provider environment, only a modest number of IT organizations have implemented NFV in a production network. However, driven primarily by the belief that NFV will enable them to reduce the amount of time it takes to deploy new services, a large percentage of IT organizations are currently in varying stages of analyzing NFV.

The NFV Survey Respondents indicated that the primary impediments that would keep their organization from broadly implementing NFV are:

- Concerns about end-to-end provisioning;
- The lack of a compelling business case;
- The immaturity of the current products.

Use Cases and Proof of Concept (POC)

The e-book discusses some of the use cases and POCs being sponsored by ETSI and by the TM Forum. The ETSI use cases are:

- NFV Infrastructure as a Service (NFVlaaS)
NFVlaaS 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.
- Virtual Network Functions as a Service (VNFaaS)
Many enterprises are deploying numerous network service appliances at their branch offices; e.g., access routers, WAN optimization controllers, stateful firewalls and intrusion detection systems. Virtual Network Functions delivered as a Service (VNFaaS) is an alternative solution for enterprise branch office networks whereby VNFs are hosted on servers in the network service provider's access network PoP.
- Virtualization of the Home Environment (VoHE)
Virtualization of the Home Environment 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.
- VNF Forwarding Graph (FG)
IT organizations 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.

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.

- Virtual Network Platform as a Service (VNPaaS)
VNPaaS is similar to an NFVlaaS 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.
- Virtualization of Mobile Core Network and IP Multimedia Subsystem
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.

- Virtualization of the Mobile Base Station
3GPP LTE provides the Radio Access Network (RAN) for the Evolved Packet System (EPS). There is the possibility that a number of RAN functions can be virtualized as VNFs running on industry standard infrastructure.
- 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.
- 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. 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.

The POCs that are being driven by the TM Forum that are discussed in this e-book are:

- Closing the Loop: Data-driven network performance optimization for NFV & SON
In this context closing the loop means collecting and analyzing data to identify how the network can be optimized and then implement those changes. This POC showed how network operators can use Self-Organizing Networks (SON) and Network Functions Virtualization (NFV) in tandem to automate closing the loop and improve performance for customers.
- CloudNFV: Dynamic, data-driven management and operations Catalyst
This POC builds on TM Forum's Information Framework to create a meta-data model using *active virtualization*, a term coined by the CloudNFV™ consortium. The specific challenge this POC is addressing is that without these connections, services like dynamic quality of service likely won't work at scale.
- Orchestrating Software-Defined Networking (SDN) and NFV while Enforcing Service Level Agreements (SLAs) over Wide Area Networks (WANs)
One set of challenges that this POC addressed are the challenges that service providers face when offering private clouds to enterprises and managing SLAs in a virtualized environment. Another set of challenges are the challenges that geographically diversified enterprises encounter when integrating data centers.
- Service bundling in a B2B2X marketplace
This POC showed how a buyer can bundle a collection of services sourced from different suppliers and deliver them seamlessly to a customer in a business-to-business or business-to-business-to-consumer arrangement. These components could include traditional network access products, as well as NFV and infrastructure-as-a-service products.

The Operational Implications

The majority of The NFV Survey Respondents indicated that they believe that even if a NFV-related POC is successful, it will take between a significant and a tremendous amount of effort to broadly implement that solution in production. One of the operational challenges that can make it difficult to move from POC to production is performance. As discussed in the e-book, in order to move VNFs into production, it must be possible to achieve the same or greater performance in a software-based environment as is possible in a traditional hardware-based environment. However, 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.

The adoption of NFV poses a number of other significant challenges that must be overcome in order to ensure the ability to continue to implement effective end-to-end management. These challenges include:

- **Dynamic relationships between software and hardware components.** With NFV, 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.
- **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.
- **Many-to-Many relationships between network services and the underlying infrastructure.** In a virtualized infrastructure a network service can be supported by a number of 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.
- **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.
- **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.
- **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.

Roughly a third of IT The NFV Survey Respondents 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 asked what type of changes they expected, a number of The NFV Survey Respondents commented that it would require them to change how they implemented SLAs, how they developed a business case and it would cause them to rethink their business models. Other comments included:

- We will need to adopt a new approach to service provisioning and management;

- It will cause us to consolidate our physical platforms;
- It will change how we do network planning;
- We will need to determine how we are going to orchestrate end-to-end systems.

Almost half of The NFV Survey Respondents indicated that over the next two years that the adoption of NFV will likely have a significant or very significant impact on the skill base of IT professionals. When asked to indicate the type of impact, the answers included:

- We will need to know multiple technologies;
- We will need to think in software and end-to-end terms rather than in component terms;
- It will require the skills to drive the integration between legacy equipment and management systems and NFV management systems;
- We will need to modify our change management, incident and problem management processes.

An additional hurdle that has to be overcome before the full benefits of NFV can be realized is that IT organizations must take a DevOps-like approach to network operations. The e-book describes the key principles that characterize DevOps and also describes how a DevOps approach has to be modified in order to be applied to network operations.

The SDN and NFV Ecosystem

The e-book 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 e-book 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;
- Standards Bodies and Related Communities;
- Providers of SDN Controllers;
- Providers of Telcom Service Provider's Infrastructure/ Optical Networking;
- Server Virtualization Vendors.

The e-book 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 e-book 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)

Introduction

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

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

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

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

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

Definition of SDN

The Open Networking Foundation (ONF) is the group that is most associated with the development and standardization of SDN. According to the [ONF](#), "Software-Defined Networking (SDN) is an emerging architecture that is dynamic, manageable, cost-effective, and adaptable, making it ideal for the high-bandwidth, dynamic nature of today's applications. This architecture decouples the network control and forwarding functions enabling the network control to become directly programmable and the underlying infrastructure to be abstracted for applications and network services. The OpenFlow™ protocol is a foundational element for building SDN solutions." Many vendors have announced support for OpenFlow V1.3. An overview of that protocol and the use cases it enables can be found in [An Overview of OpenFlow V1.3](#).

The ONF states that the SDN architecture is:

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

Context for SDN

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

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

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

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

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

Status of SDN Adoption

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

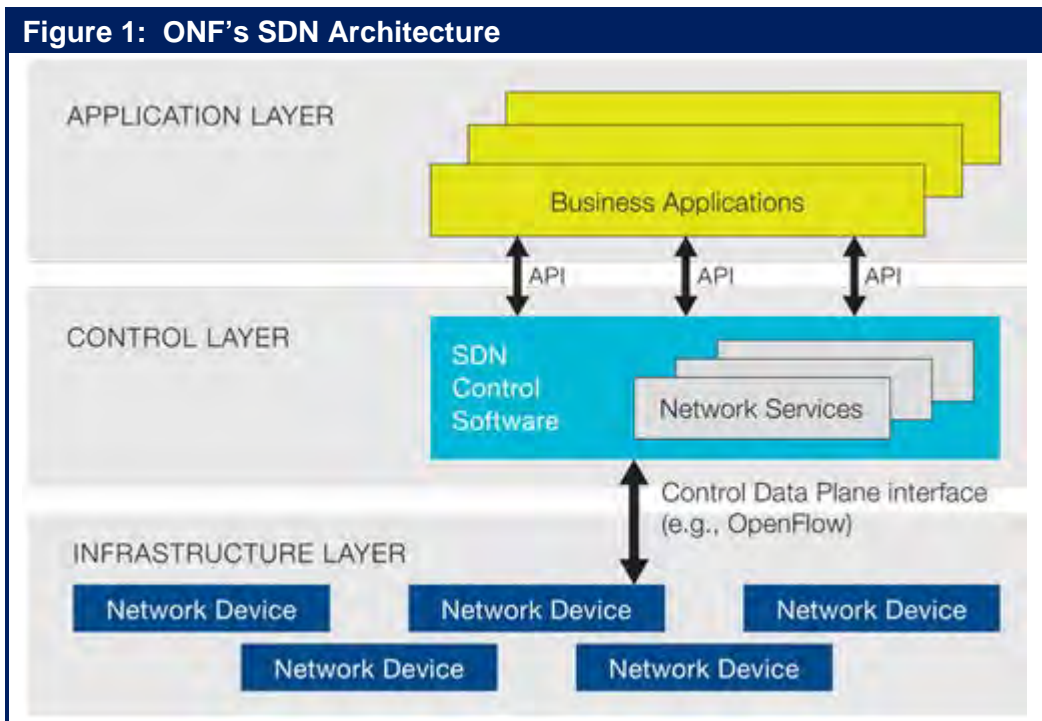
Table 1: SDN Utilization		
Approach to Implementing SDN	Responses to The 2014 Survey	Responses to The 2013 Survey
We have not made any analysis of SDN	16%	19%
We will likely analyze SDN sometime in the next year	22%	26%
We are currently actively analyzing the potential value that SDN offers	32%	36%
We expect that within a year that we will be running SDN either in a lab or in a limited trial	22%	19%
We are currently actively analyzing vendors' SDN strategies and offerings	25%	20%
We currently are running SDN either in a lab or in a limited trial	18%	13%
We currently are running SDN somewhere in our production network	11%	6%
We looked at SDN and decided to not do anything with SDN over the next year	10%	5%
We expect that within a year that we will be running SDN somewhere in our production network	18%	10%
Don't know	2%	4%

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

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

The SDN Architecture

Figure 1 contains a graphical representation of the SDN architecture as envisioned by the ONF. One key component of a complete SDN solution that is missing from **Figure 1** is cloud orchestration platforms such as OpenStack. The role that these platforms play in both SDN and NFV is described later in this document.



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

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

The Northbound Interface

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

After over a year of discussion, in late 2013 the ONF created the NBI working group and outlined the group's charter in a [white paper](#). As part of their charter, the NBI working group intends to work with one or more open source initiatives to develop working code for the NBIs that the group standardizes.

According to [Sarwar Raza](#) the chair of the NBI working group, the working group has a good relationship with both the OpenStack and the OpenDaylight initiatives but that when dealing with open source initiatives “there is no magic handshake”. Raza elaborated by saying that none of the open source initiatives are going to agree in advance to produce code for NBIs that are under development. He expects that what will happen is that after the standards have been developed the NBI working group will have detailed technical discussions with multiple open source communities and will see if there is a consensus about developing code.

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

Architectural Distinctions between Approaches

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

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

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

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

As indicated above, there is a divergence of opinion relative to the role of dedicated hardware. One example of that divergence of opinion is that some vendors believe it is possible to fully support network virtualization in the data center without using dedicated hardware and some vendors believe that dedicated hardware is needed at least some times. The Survey Respondents were asked to indicate if they believed that with the current technologies and products it’s possible to broadly support network virtualization in the data center without using any dedicated hardware? The *no* responses outnumbered the *yes* responses by almost a 2:1 ratio.

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

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

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

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

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

The Overlay and the Underlay Model

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

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

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

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

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

other than network virtualization, but that the underlay solution also uses an overlay-based approach to implement network virtualization.

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

Table 3: The Perceived Value of the Overlay and Underlay-based Models	
Response	Percentage of Respondents
The overlay-based model will provide notably more value	22%
The fabric-based model will provide notably more value	28%
Each model will offer roughly equal value	12%
We don't have an opinion on either model	31%
Other	7%

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

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

Service Chaining

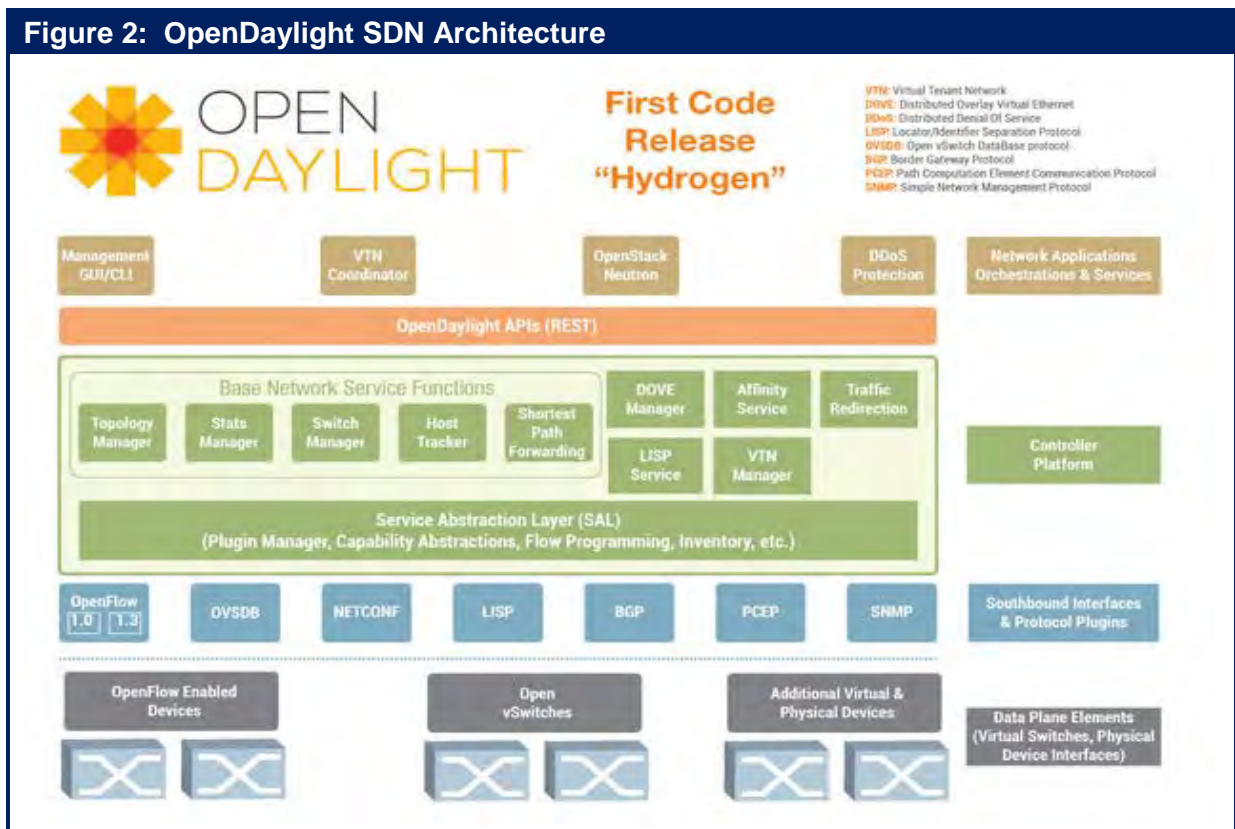
The phrase *service chaining* refers to the ability to steer virtual machine (VM)-VM traffic flows through a sequence of physical and/or virtual servers that provide network services, such as firewalls, IPS/IDS, DPI, or load balancers. In an underlay-based solution, the controller configures the forwarding plane switches to direct the flows along the desired paths. In an overlay-based solution, the controller adjust the Forwarding Information Bases (FIBs) of the vSwitches/vRouters to force the traffic through the right sequence of VMs. The next section of The Guide focuses on Network Functions Virtualization (NFV). That section will discuss what the European Telecommunications Standards Institute (ETSI) refers to as VNF forwarding graphs, which are similar in concept to service chains.

The OpenDaylight Consortium

The [OpenDaylight Consortium](#) was founded in April 2013. The consortium's stated mission is to facilitate a community-led, industry-supported open source framework, including code and architecture,

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

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



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

The Relationship Between SDN and NFV

Until recently, the conventional wisdom in the IT industry in general, and on the part of the ONF and the ETSI NFV ISG¹ in particular, was that SDN and NFV were separate topics and didn't need to be formally coordinated. That conventional wisdom officially changed in March 2014 when the ONF and the ETSI NFV ISG announced the signing of a Memorandum of Understanding (MOU).

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

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

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

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

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

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

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

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

SDN Use Cases

Drivers and Inhibitors

The Survey Respondents were shown a number of challenges and opportunities and were asked to indicate which of them they thought that SDN could help them to respond to and they were allowed to check all that applied. Their responses are shown in **Table 5**.

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

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

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

The Survey Respondents were also shown a set of impediments and were asked to indicate the two impediments that would be the biggest inhibitors to their company adopting SDN sometime in the next two years. Their responses are shown in **Table 6**.

Table 6: Inhibitors to the Adoption of SDN

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

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

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

SDN Deployment Plans

While the use of SDN in data centers receives the majority of attention, it is also possible to implement SDN in branch and campus networks as well as in wide area networks (WANs). In order to understand where SDN will likely be implemented, The Survey Respondents were asked "If your organization is likely to implement SDN sometime over the next two years, where are you likely to implement it?" Their responses are summarized in **Table 7**.

Table 7: Focus of SDN Deployment

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

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

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

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

Table 8: Planned SDN Deployment			
	Campus Networks	WAN	Data Center Networks
Exclusively based on SDN	1%	2%	6%
Mostly SDN	10%	6%	20%
Hybrid, with SDN and traditional coexisting about equally	34%	36%	50%
Mostly traditional	29%	31%	10%
Exclusively traditional	13%	13%	4%
Don't know	12%	12%	10%

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

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

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

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

Data Center

Virtual Machine Migration

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

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

Service Chaining

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

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

Security Services

By virtue of Layer 2-4 flow matching capability, OpenFlow access switches can perform filtering of packets as they enter the network, acting as simple firewalls at the edge. With OpenFlow switches that support modification of packet headers, an OpenFlow-enabled controller will also be able to have the switch redirect certain suspicious traffic flows to higher-layer security controls, such as IDS/IPS systems, application firewalls, and Data Loss Prevention (DLP) devices. Other security applications built on OpenFlow controller can match suspicious flows to databases of malware signatures or divert DDoS attacks.

Load Balancer Services

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

Indiana University (IU) has developed an OpenFlow-based, load-balancing application called FlowScale. According to the [University](#), “FlowScale provides complex, distributed load balancing of network traffic using an OpenFlow-capable Top of Rack (ToR) switch. IU deployed the application into its Intrusion Detection System (IDS) to distribute traffic evenly to sensors. FlowScale is currently being deployed as part of the Intrusion Detection Systems operated by the Indiana University Information Security Office.

New Software Defined Cloud model

A national Cloud Leader is creating a new architecture that will allow all network and value-added services to be software defined, based on SDN. They are using OpenStack both for overall orchestration and also for its end user-friendly Horizon dashboard. While customers interact with the system via the dashboard, their administrators leverage REST APIs to interact both with OpenStack and with the SDN. The SDN provides a virtual network overlay for a consistent, unified fabric over the entire network and all datacenters (planned expansion is 10 datacenters by 2018).

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

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

New Distributed Cloud Hosting model

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

New capabilities include:

- **Precision SLAs for each Customer:** Adding precise network controls to server virtualization and OpenStack orchestration, the TSP's key customer demand of precise, end-to-end SLAs can be reliably delivered even on shared vPODs.
- **Consistent Performance Across Datacenters:** Similar to server virtualization, network virtualization provides consistent and predictable performance that is independent of each datacenter's build-out, hardware configurations, and network architectures. This capability enables a range of new customer-centric capabilities such as load balancing workloads across datacenters.
- **Fluid Disaster Recovery:** Having consistent performance independent of datacenter build-out changes how disaster recovery is performed. Instead of having idle resources standing by in a dedicated datacenter, a customers' implementation can be stretch-clustered across datacenters for truly fluid disaster recovery. In this fashion, loss of one or even multiple datacenters can be accommodated without disruption to operations.
- **Effortless Datacenter Scalability:** With this architecture, the TSP can scale-out to accommodate the needs of each customer just by adding vPODs or by adding racks to a dedicated vPOD. They also can easily scale out to 100 times their initial rack count and up to millions of managed endpoints without having to change the architecture or the configuration.
- **Fast and Non-disruptive Provisioning:** Outside of physical racking and cabling, new vPODs can be added in an automated and non-disruptive manner – the entire installation or de-installation process takes only a few hours. New servers can be allocated to a vPOD nearly instantaneously via automation.

WAN

The Google G-Scale WAN

As is discussed in [An Overview of OpenFlow V1.3](#), one of the primary benefits of OpenFlow is the centralized nature of the Forwarding Information Base (FIB). Centralization allows optimum routes to be calculated deterministically for each flow by leveraging a complete model of the end-to-end topology of the network. Based on an understanding of the service levels required for each type of flow, the centralized OpenFlow controller can apply traffic engineering principles to ensure each flow is properly serviced. Bandwidth allocations can be controlled dynamically to provide bandwidth on demand with changing traffic patterns. The result can be much better utilization of the network without sacrificing service quality. Centralized route processing also allows the pre-computation of a set of fail-over routes for each possible link or node failure.

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

Campus

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

Dynamic QoS & Traffic Engineering

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

In an SDN environment, as the Lync call controller is sending the IP address to the Lync client running on User A's laptop, the Lync controller can be configured to also send it to an SDN application, whose function is to communicate with an SDN controller and have the priority set to specified values for specific IP pairs in a network. A Lync call, for instance, could be set to a high priority. The SDN application communicates to the SDN controller that the priority level for traffic between a specific pair of IP addresses needs to be set to high and that this traffic should run over non-congested links. The SDN controller takes this information and determines the optimal path for the packets to flow through the network from User A to User B. This flow matching information, along with the required actions, are pushed out to each of the OpenFlow-enabled switches.

Unified Wired and Wireless Networks

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

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

QoS Management for Microsoft Lync across wired and wireless networks

This use case can be viewed as a combination of the preceding two use cases. As previously noted, enterprises are rapidly adopting Microsoft's Lync as their unified communications solution of choice, but until recently a unifying Lync wired and wireless solution wasn't available in the market. That is important because wireless has become the edge of the network and mobile users have a growing dependency on wireless services for performing critical job tasks. This situation creates a challenge relative to how wireless users can effectively and reliably access Lync services.

Recently an OpenFlow-based application that bridges wired and wireless networks to ensure a user the highest quality of experience with Microsoft's Lync has entered the market. The solution can detect quality of service issues, identify resolutions and prioritize traffic across any OpenFlow-enabled network. The solution also enables the wireless and wired network to dynamically change in response to application traffic requirements.

Personal Bonjour

When Apple announced Bonjour, its zero-configuration application, it filled a void in the market. Users could simply access a network attached television or printer, as long as the device was on the same sub-net. Businesses quickly saw value in this class of application and commercial network centric solutions opened Bonjour up to more expansive networks. However, this created a management challenge relative to user and device associations with larger populations of network users and devices. Recently an OpenFlow-based application has been introduced to the market that implements the highest granularity policy management for Bonjour service access available. This application has functionality that enables IT organizations to ensure that individual users may be allowed to only access selected devices, in selected locations, at selected times of day. One way that this can be used is that a dormitory full of students, each with their own printer or TV, can be isolated from all other users without the network congestion often encountered with a standard Bonjour implementation.

Role Based Access

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

The Operational Implications

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

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

Security

SDN creates security opportunities and security challenges.

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

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

Some of the security challenges related to SDN are described in [SDN Security Considerations in the Data Center](#). As pointed out in that document:

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

Other security-related considerations include that IT organizations should:

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

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

Cloud Orchestration

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

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

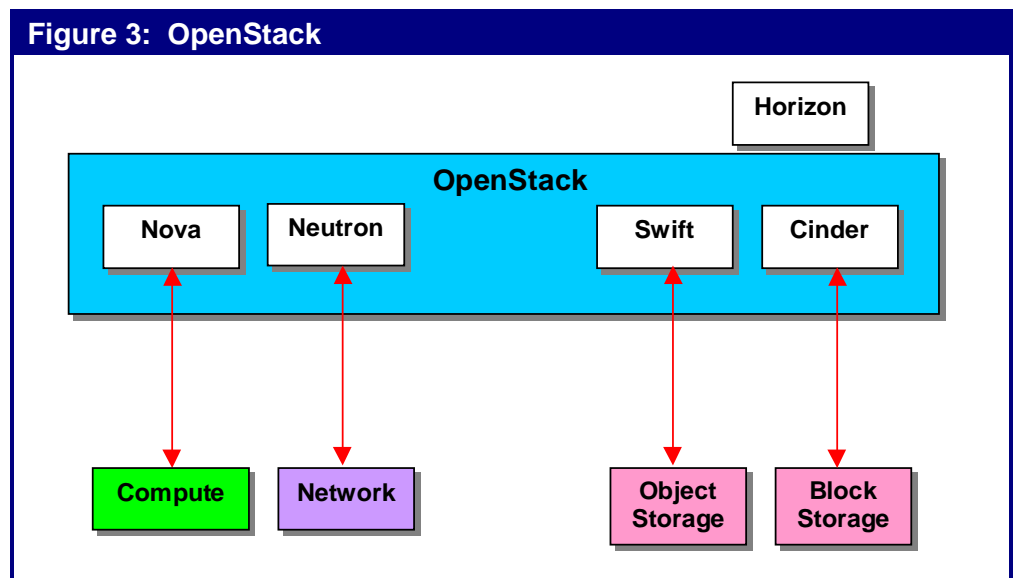
In addition, there are a number of proprietary orchestrators that offer open APIs to allow integration across vendor boundaries. These include VMware's vCloud Director and IBM's SmartCloud Orchestrator.

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

Horizon is the OpenStack Dashboard that provides administrators and users a graphical interface to access, provision and automate cloud-based resources. The dashboard is one of

several ways users can interact with OpenStack resources. Developers can automate access or build tools to manage resources using the native OpenStack API or the EC2 compatibility API. The dashboard also provides a self-service portal for users to provision their own resources within set limits.

Neutron (formerly called Quantum) allows users to create their own networks, provide connectivity for servers and devices, and control traffic. With appropriate Neutron plug-ins, administrators can take advantage of various SDN solutions to allow for multi-tenancy and scalability. A number of drivers/plugins are included with the OpenStack source code. OpenStack networking also has an extension framework allowing additional network services, such as intrusion detection systems (IDS), load balancing, firewalls and virtual private networks (VPN) to be deployed and managed. One example of the extension service is the Load Balancer as a Service (LBaaS) driver for Neutron available starting with the October 2013 Havana release. The driver enables ADC vendors to offer simple LBaaS plugins



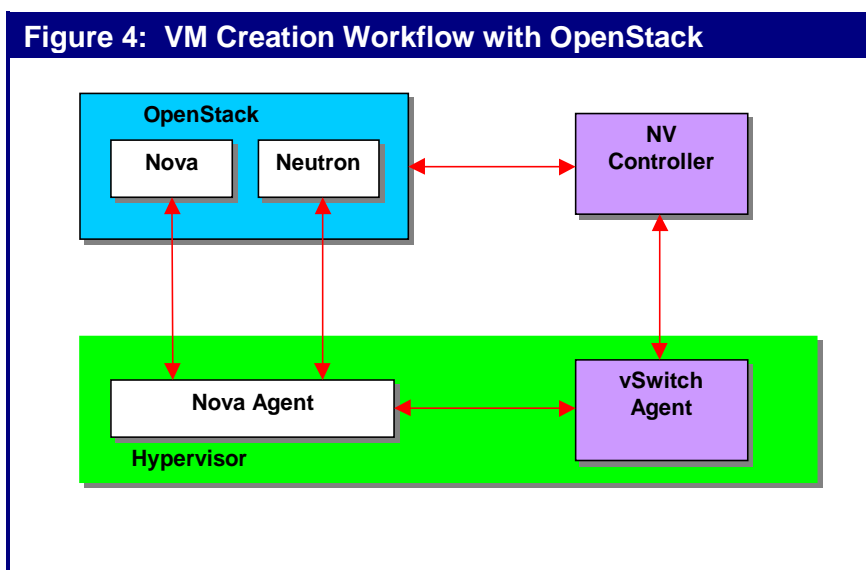
for Neutron, allowing their ADCs to be directly provisioned by OpenStack. Vendor-specific driver plugins that are contributed to the project are included in the OpenStack source code.

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

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

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

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



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

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

Table 9: Approaches to Orchestration	
Approach	Percentage of Respondents
We have a well thought out strategy and we have begun to execute against that strategy	16%
We have a well thought out strategy but we have not yet begun to execute against that strategy	6%
We are in the process of developing a strategy and are optimistic that it will come together relatively quickly	21%
We are in the process of developing a strategy but have some concerns that the existing solutions are immature	30%
Don't know/NA	22%
Other	5%

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

Management

SDN creates management opportunities and security challenges.

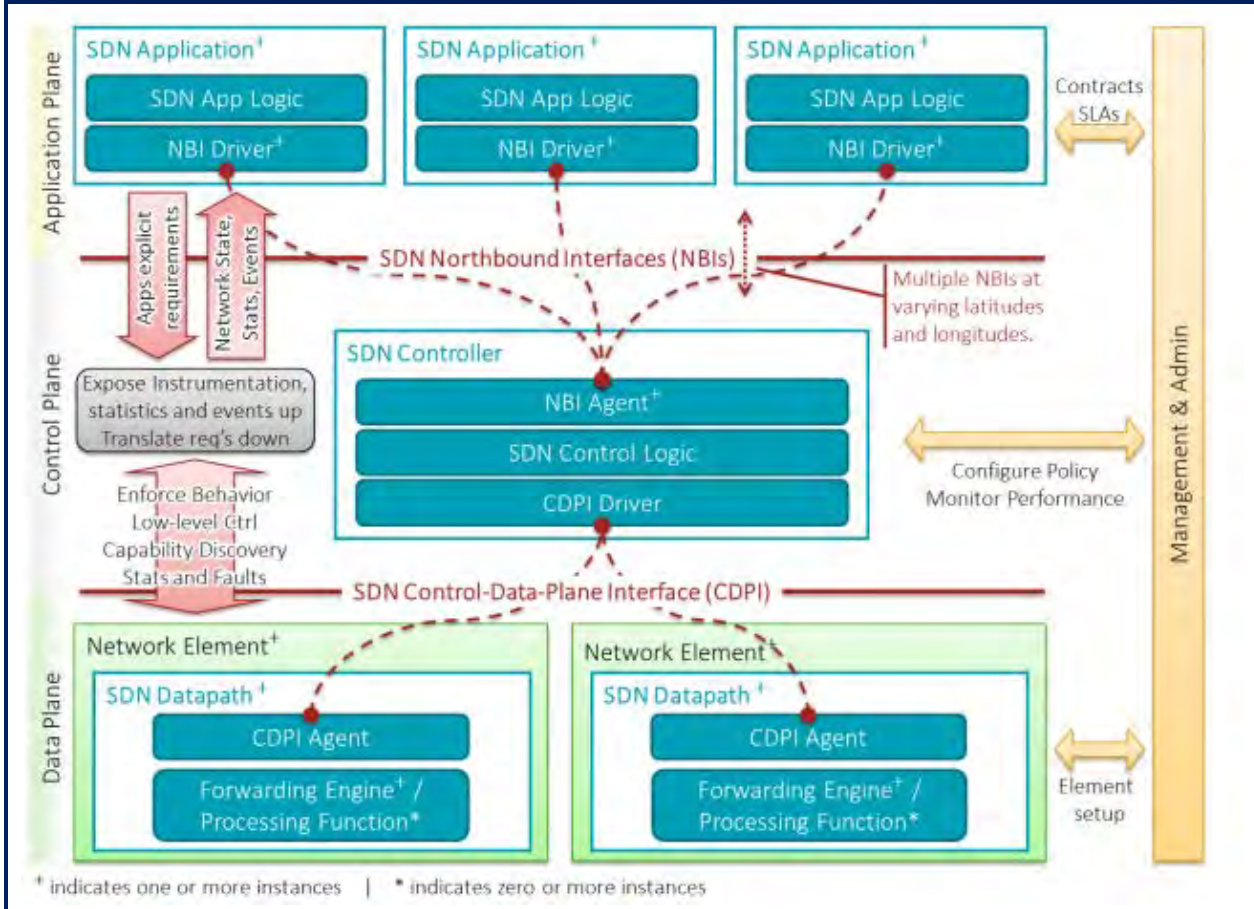
The fact that SDN poses both management opportunities and management challenges was demonstrated by **Table 5** and **Table 6**. **Table 5** shows that 53% of network organizations believe that SDN will ease the administrative burden of management tasks such as configuration and provisioning. **Table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

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

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

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

Figure 5: SDN Management Challenges



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

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

At the infrastructure tier, one of the primary challenges is to perform element management potentially of both virtual and physical network elements. One of the management challenges at the control layer results from the fact that the controller is in the data path for new flows entering the network. During periods when many new flows are being created, the controller can potentially become a performance bottleneck adding significant latency for flow initiation. Performance management systems need visibility not only into application performance but also controller performance in processing flows.

As described in the preceding discussion of the North Bound Interface (NBI), one of the management challenges that occurs at the application tier is that based on the type of application (e.g., business

application vs. a firewall) the service or application needs varying levels of visibility into the underlying network. Another set of management challenges that occurs at the application layer stem from the requirement to ensure acceptable performance. As described below, one thing this means is that network management organizations must have visibility into the SLA requirements of the application so that resources can be dynamically allocated to meet those requirements.

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

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

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

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

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

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

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

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

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

Organizational Impact

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

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

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

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

- **More focus on setting policy**

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

- **More knowledge of the business**

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

- **More understanding of applications**

While client server and n-tier applications are still common, as pointed out in [The 2013 Application and Service Delivery Handbook](#), many applications are now based on a wide range of architectures; e.g., a Services Oriented Architecture (SOA). In addition, complex applications, such as Customer Relationship Management (CRM), are actually comprised of several modules, with a range of network requirements. IT infrastructure and network professionals in particular need to better understand these new architectures and complex applications in order to ensure that the emerging set of technologies are designed and architected appropriately.

- **More emphasis on programming**

While it is not true that all networking and data center professionals will become programmers, it is true that many senior level IT professionals will need an understanding of programming in order to better interact with the company's software development organization. It is also true

that some network organizations will want to leverage the API functionality that the emerging technologies provide by having network professionals write programs that utilize those APIs.

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

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

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

Table 10: Impact of SDN on Organizational Structure	
Impact	Percentage of Responses
Very Significant Impact	4%
Significant Impact	16%
Moderate Impact	14%
Some Impact	18%
No Impact	23%
Don't Know	25%

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

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

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

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

- A shift from siloed specialists to service aligned generalists;
- A likely re-org around application development and network operations;
- An increase in cross functional teams and projects;

- Moving from a tower based organization to a DevOps model;
- An increased focus on software engineering;
- Team work will involve an enhanced mix of skills including programming, networking, virtualization and DevOps;

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

Table 11: Impact of SDN on Jobs	
Impact	Percentage of Responses
Very Significant Impact	6%
Significant Impact	19%
Moderate Impact	16%
Some Impact	23%
No Impact	19%
Don't Know	18%

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

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

- The product development life cycle will change;
- The job will require new skills in general and more knowledge of software in particular;
- The customer demands are unknown;
- Product development needs to be able to provide tools to manage and monitor the environment;
- There will be new business models, new product offerings that must be supported.

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

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

~ Continued on page 50 ~

Cisco ACI: An Application Centric Approach to SDN

IT Trends and the Advent of Software Defined Networking

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

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

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

A Comparison of ACI to Traditional SDN Architectures

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

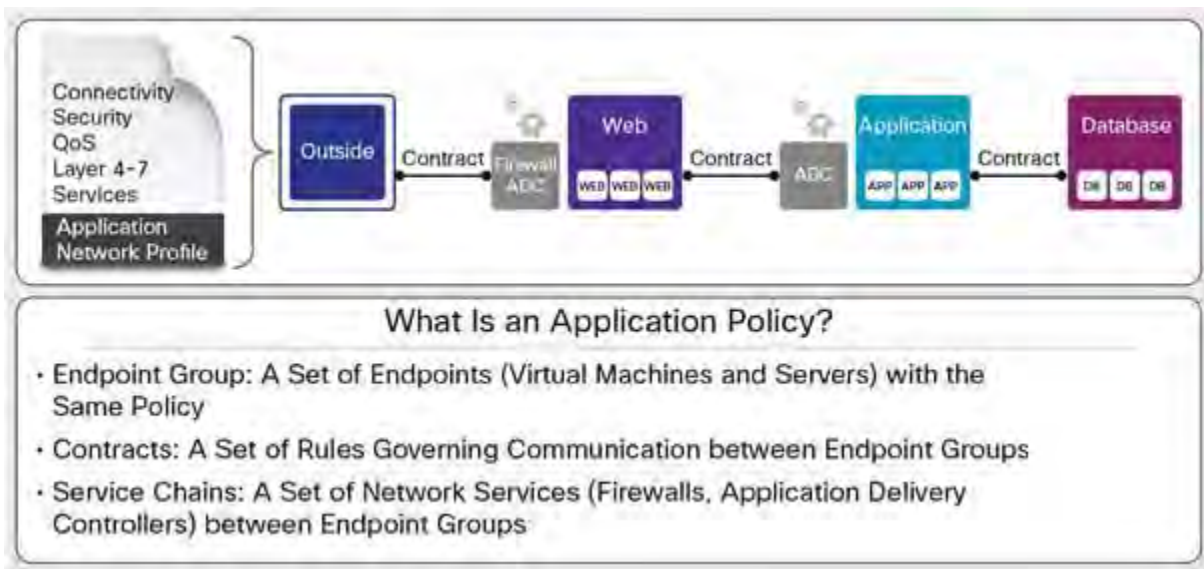
- A centralized policy store and infrastructure controller: In SDN and Cisco ACI, this feature is generally known as the controller (Cisco [Application Policy Infrastructure Controller \[APIC\]](#) for Cisco ACI).
- Programmable, or automated, network devices: All infrastructure devices, such as switches,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

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

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

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

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



For More Information

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



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

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

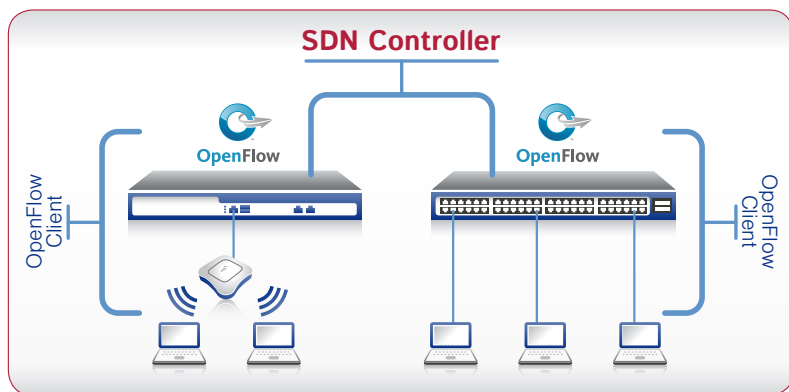
Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

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

Recycling symbol: Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)



Delivering the promise of SDN across wired/wireless networks



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

The rapid acceptance of SDN and this new approach to design, build and manage data centers addresses the top challenges experienced by organization related to networks: namely, too many manual processes, and

difficulties changing configurations. SDN tackles these challenges in the data center, but SDN can equally address the same issues for the enterprise campus. Without bringing SDN to the edge of the network, its true promise is lost.

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

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



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



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



Meru Collaborator

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



Meru Personal Bonjour

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

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



Making SDN a Reality for Wi-Fi

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

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

[Click for more information](#)



Corporate Headquarters
894 Ross Drive
Sunnyvale, CA 94089

T +1 (408) 215-5300

F +1 (408) 215-5301

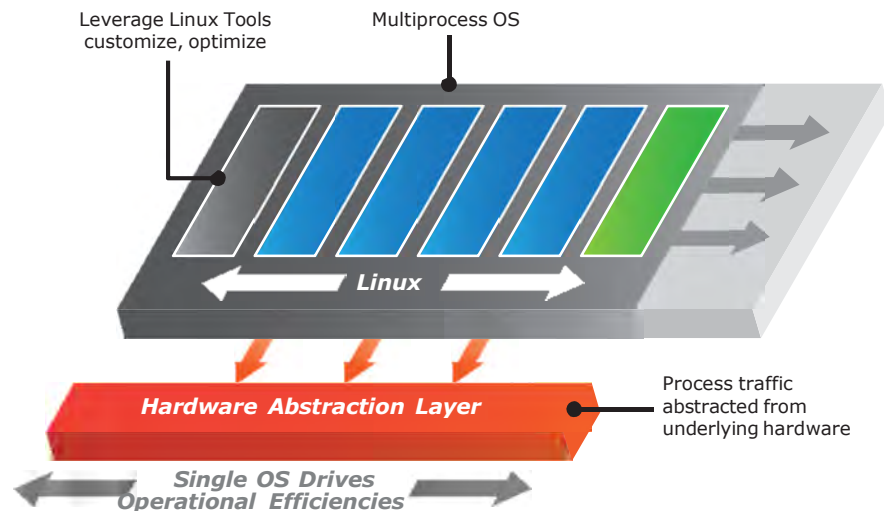
E meruinfo@merunetworks.com

PicOS Overview

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

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

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



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

PicOS - Three Editions to Leverage

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

Features Included	Required PicOS Editions		
	Linux Switching OS	Routing	OpenFlow
<ul style="list-style-type: none"> • Network operation system using user space standard Debian Linux environment • Leverage vast array of standard Linux tools as a common management and operations framework • Zero Touch Provisioning (ZTP) functionality coupled with ONIE delivers a true bare metal to application environment • Rich Layer-2 protocol stack with MLAG, seamlessly integrating into existing architectures • Full Layer-2 & Layer-3 ACL support • IPv4 & IPv6 Static Routing 	✓		
<ul style="list-style-type: none"> • Rich OSPF and BGP protocol stacks integrating into existing spine / leaf architectures • IPv6 routing protocol support (OSPFv3, MBGP) • Multicast PIM support • NAT (depends on ASIC support) • VXLAN network virtualization (depends on ASIC support) 	✓	✓	
<ul style="list-style-type: none"> • Leading OpenFlow 1.4 support through OVS 2.0 • Deliver true seamless migration to SDN through CrossFlow mode (Layer-2 / Layer-3 and OpenFlow simultaneously) • Leveraging OpenFlow to control MPLS, GRE, NVGRE or VXLAN tunnels, delivering on the promise of open programmability • Support for all major OpenFlow controllers (for example: OpenStack Neutron ML2, OpenDaylight, Ryu) 	✓		✓
PICOS Bundle	✓	✓	✓

Enterprise SDN and Carrier NFV: Distant Cousins or Twins?

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

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

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

OPERATIONAL NEEDS DRIVE ENTERPRISE SDN ADOPTION

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

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

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

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

1. Network virtualization overlay (NVO)

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

2. Controller-based solutions (ex: Openflow)

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

3. Programmable solution (ex: REST)


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

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

SERVICE DELIVERY GOALS DRIVE CARRIER NFV ADOPTION


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

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



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

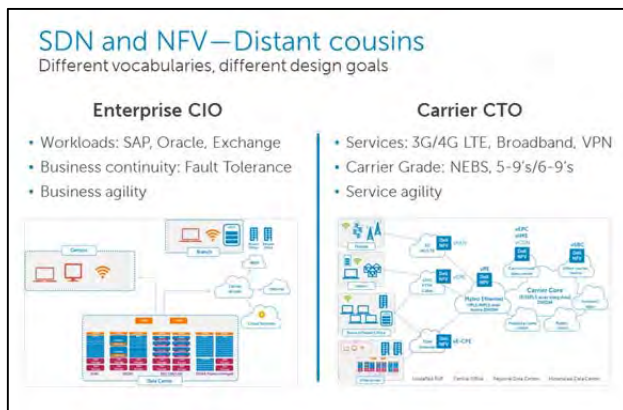
Carriers globally are now turning to virtualization and in particular NFV as a way to simplify and automate service delivery infrastructure, while also introducing greater agility for new service creation and delivery. For carriers, then, the drivers for virtualization are to improve both CAPEX and OPEX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



THE LANGUAGE BARRIER

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

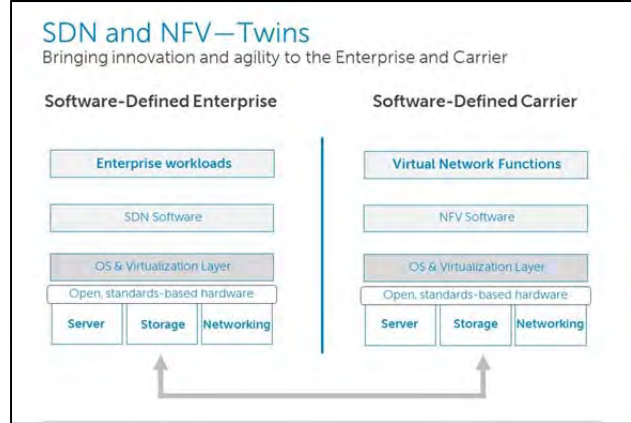
There are also fundamental differences in the vocabulary surrounding each. Instead of *workloads*, carriers are concerned with *services*, and instead of *business continuity*, carriers are interested in *carrier-grade 5-9's and 6-9's technology*.



DISTANT COUSINS OR TWINS?

Considering this laundry list of differences, you might wonder how we can propose that enterprise SDN and carrier NFV are actually twins. It's not until

you look at their technological DNA that you start to see the remarkable similarities.



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

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

COMMON TRAITS FOR THE FUTURE

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

1. Carriers who are new to network virtualization can learn best practices from Web 2.0 and large enterprises who have already made significant strides in that area and apply in context..
2. Organizations that operate both production and provisioned infrastructure—enterprise-style for their own operations and carrier-style to provide services—can cross-pollenate, leveraging common technology assets, best-practices, and purchasing power.

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

Dell is one of the world's leading providers of SDN and NFV, and the only provider of truly open networking with software/hardware disaggregation. Learn more at Dellnetworking.com

Dynamic Cloud, Dynamic Services

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

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

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

The NFV Journey

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

1. Virtualization

The abstraction of the Telecom functions software from dedicated hardware to run on open commercial-off-the-shelf (COTS) hardware, as well as the need to balance performance and cost reductions, will force service providers to make critical roadmap decisions. Some

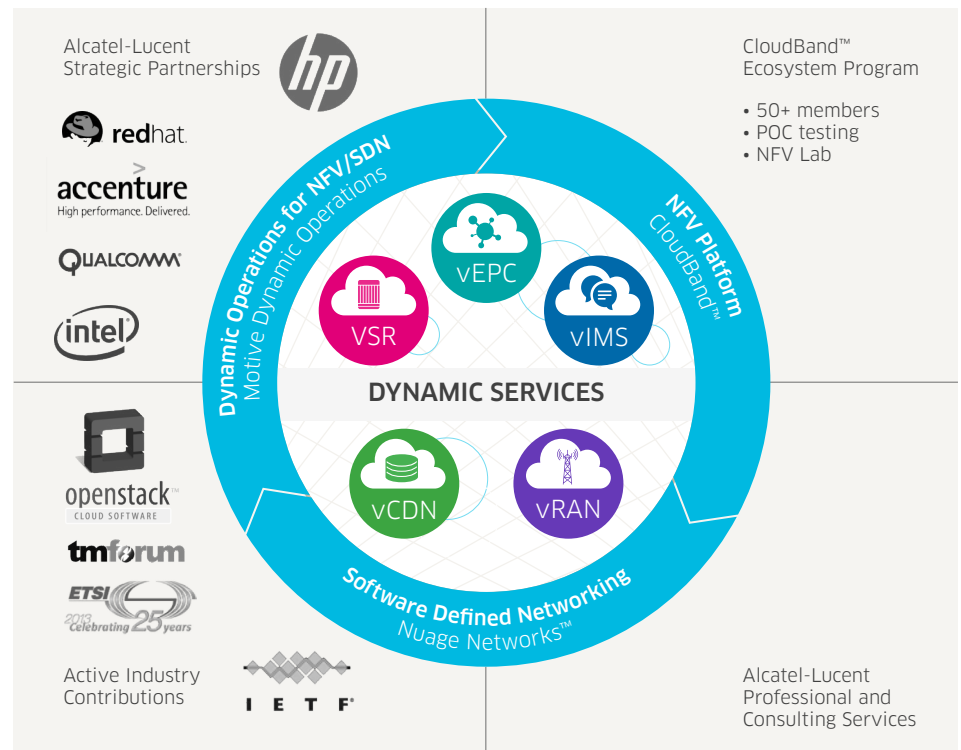


functions will achieve significant advantages of scale and flexibility from COTS hardware, while other functions, or even the same function, may benefit from the performance advantages of dedicated hardware. This duality is likely to exist for a while, as we pass through a transition phase, but this should not complicate the operational model, provided the same management entities exist. While exact feature parity may not be critical, function performance and robustness cannot be compromised. Service providers should consider the many years

of feature development put into the existing functions, and consider carefully how this work will be carried forward into the new mode.

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

Alcatel-Lucent's uniquely open approach and ecosystem



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

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

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

2. Orchestration

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

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

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

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

3. Automation

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

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

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

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

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

CloudBand

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

Virtualized Service Routing

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

Virtualized IMS

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

Virtualized IP Mobile Core

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

Nuage Networks SDN

Nuage Networks is a leader in SDN. It focuses on modern datacenter requirements for multi-tenancy, full-featured routing and security at scale. The Nuage Networks platform transforms the physical network into a simple to manage, rack-once and wire-once, vendor-independent IP backplane. As a result, network resources within and across datacenters can be treated as an elastic resource pool of capacity that can be consumed and re-purposed on demand.

Motive Dynamic Operations

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

Software-Defined Networking

Are your management tools prepared?



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

Are you prepared for this evolution?

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



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

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

EMC²

Extending Service Performance Management into SDN and NFV Environments

Solution Benefits

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

Problem Overview

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

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

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

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

Solution Overview

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

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

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

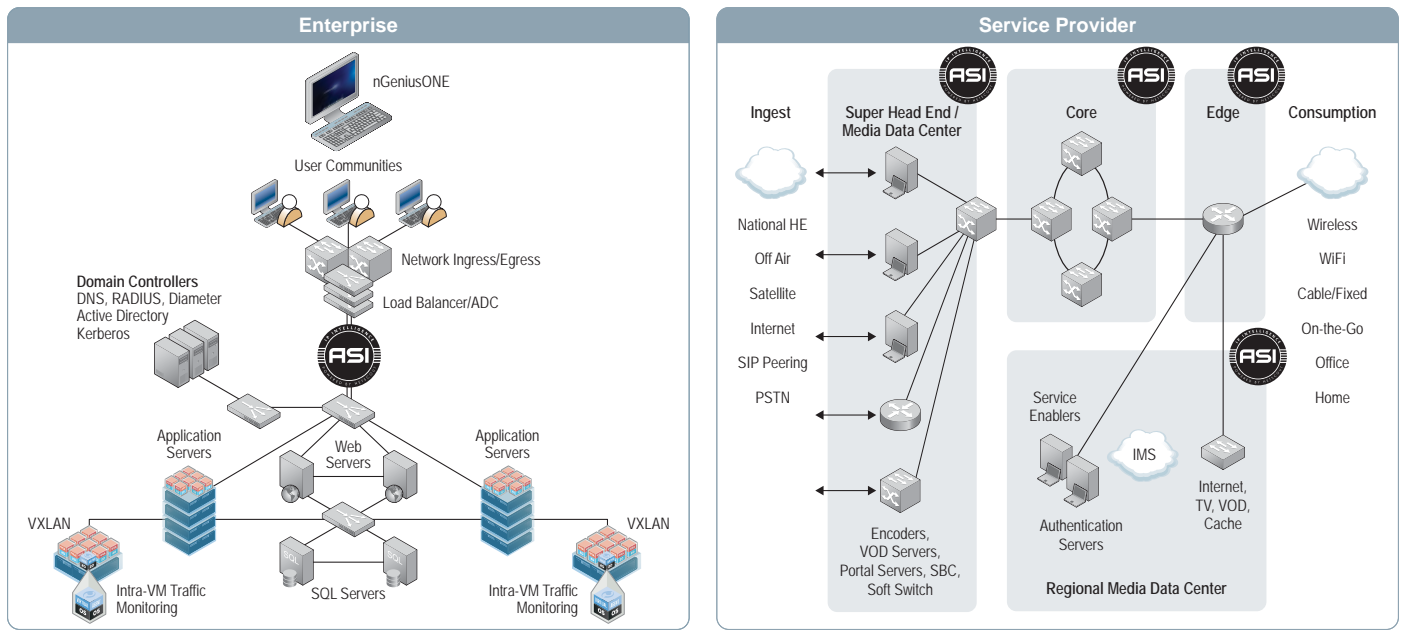


Figure 1: Service Instrumentation in Enterprise and Service Provider Environments.

Core Technologies

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

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

Utilize Packet Flow Data

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

Provide Scalable Packet Flow Access

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

Adaptive Session Intelligence (ASI)

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

Service Delivery Monitoring in SDN Environments

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

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

Solution Benefits

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

Attribute	Bottom-Up Triage Problems	NetScout’s Solution	IT Benefits
End-to-End Visibility	<ul style="list-style-type: none"> • 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. 	<ul style="list-style-type: none"> • 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. 	<ul style="list-style-type: none"> • Optimize experience of user communities and customers. • Comprehensive solution from a single vendor. • Full visibility into services running in physical, virtual, and hybrid environments.
Effective Service Triage	<ul style="list-style-type: none"> • Reactive and time consuming triage results in poor user experience, and extended service downtime impacting multiple users. 	<ul style="list-style-type: none"> • Proactive service triage helps resolve service degradation in real time, before a large number of users are impacted. 	<ul style="list-style-type: none"> • Increase service uptime and end-user productivity. • Support more services with existing IT resources. • Reduce time wasted in war rooms.
Scalability	<ul style="list-style-type: none"> • Lacks scalability required to assure delivery of modern business services for service providers and enterprises. 	<ul style="list-style-type: none"> • Scales to assure service delivery across any size of service provider and enterprise infrastructure. 	<ul style="list-style-type: none"> • Optimize your investment in performance management by gradually expanding the solution over time.

About NetScout Systems, Inc.

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



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

Copyright © 2014 NetScout Systems, Inc. All rights reserved. NetScout, nGenius and InfiniStream are registered trademarks, nGeniusONE and Adaptive Session Intelligence are trademarks and MasterCare is a service mark of NetScout Systems, Inc. and/or its affiliates in the United States and/or other countries. All other brands and product names, and registered and unregistered trademarks are the sole property of their respective owners. NetScout reserves the right, at its sole discretion, to make changes at any time in its technical information, specifications, and service and support programs.

Network Functions Virtualization (NFV)

Introduction

This section of the e-book contains the results of a survey that was distributed in October 2014. Throughout The Guide, the 135 professionals who completed the survey will be referred to as **The Survey Respondents**. Of the 135 hundred IT professionals who completed the survey, only 2 indicated that they were extremely familiar with NFV.

The general awareness of NFV is low in general and it is lower than the general awareness of SDN.

Background

The acronym **NFV** is often associated with telecommunications service providers. Their interest in NFV stems from the fact that, in the current environment, telecommunications and networking software is being run on four types of platforms:

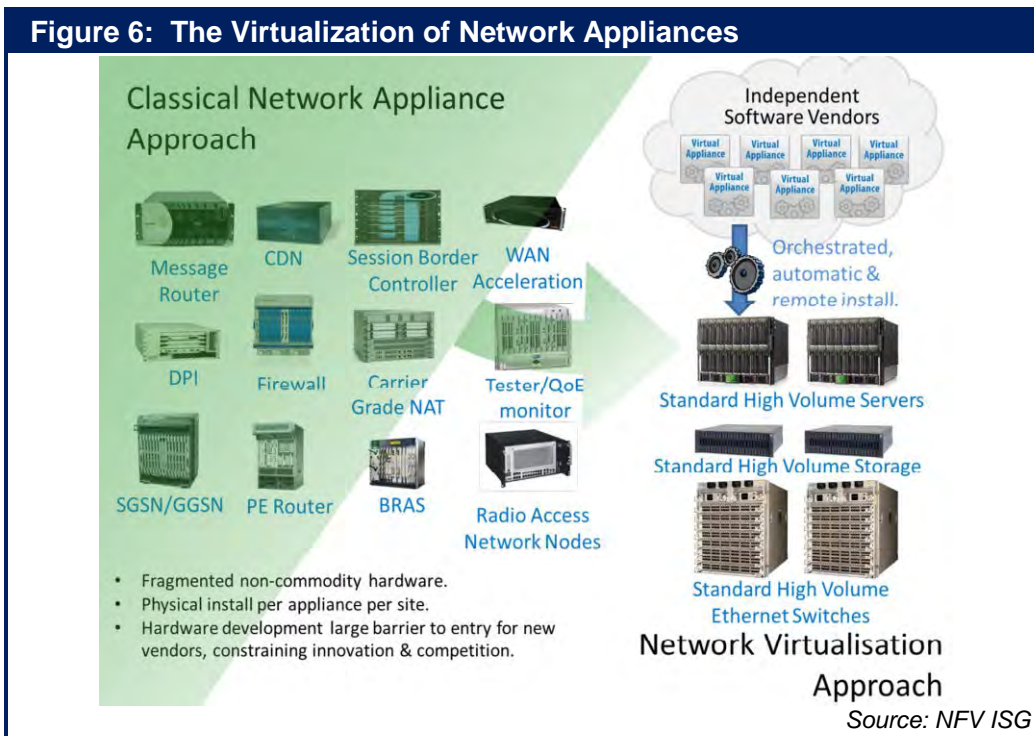
- Industry standard servers running Linux or Windows;
- Virtual appliances running over hypervisors on industry standard hardware servers;
- COTS-compute blade integrated in network elements;
- Proprietary hardware appliances.

A large part of the initial motivation to develop NFV came from the fact that telecommunications service providers felt that they can greatly simplify their operations and reduce cost if all network functions were available as virtual appliances that can be easily provisioned and integrated regardless of the vendor who provided the appliance or the hypervisor(s) on which it runs. However, while service providers 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. As such, NFV can be regarded as an important topic for both service providers and for enterprises.

The subsequent section of this document contains recent market research that indicates the relative importance of a variety of the criteria that are driving the development and implementation of NFV.

ETSI

In order to bring the vision of NFV to fruition, an Industry Specifications Group (ISG) for Network Functions Virtualization (NFV ISG) has been formed under the auspices of the European Telecommunications Standards Institute ([ETSI NFV ISG](#)). Their vision for the transition from hardware appliances of today to a fully virtualized appliance environment is depicted in **Figure 6**.



The approach that the ETSI NFV ISG is taking is that the virtualization of network functionality is applicable to any data plane packet processing and control plane function in both fixed and mobile networks. **Table 12** contains examples of functions that could 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

The initial members of the NFV ISG were service providers such as AT&T, Deutsche Telekom and NTT. Its [membership](#) has since grown and, as of October 2014, there were more than 90 organizations that are full members of the ETSI NFV ISG, with approximately another 140 organizations listed as participants.

The first meeting of the group was held in January 2013 and a number of smaller working groups were created in April 2013. In October 2013, ETSI published the first five specifications relative to [NFV](#). According to [ETSI](#), “The five published documents include four ETSI Group Specifications (GSs) designed to align understanding about NFV across the industry. They cover NFV use cases, requirements and the architectural framework. The fifth GS defines a framework for coordinating and promoting public demonstrations of Proof of Concept (PoC) platforms illustrating key aspects of NFV. Its objective is to encourage the development of an open ecosystem by integrating components from different players.” One of the documents that ETSI published identified [NFV-related terminology](#) and it is a useful reference when reading any NFV-related document, including this document. As of October 2014, ETSI is sponsoring twenty-five [POCs](#).

One of the interesting aspects of the ETSI NFV ISG is that it has a two year life span that expires in January 2015. As a result, there is work underway to identify what happens after that. For example, in late July and early August 2014 the NFV ISG met in Santa Clara, CA. At that meeting the primary objectives of NFV Phase 2 were [identified](#). Whereas ETSI characterizes Phase 1 as being the Requirements Phase, ETSI characterizes Phase 2 as being the Implementation Phase. The objectives of Phase 2 include building on the achievements that were made in the first two years of the ISG and consist of an enhanced focus on interoperability, formal testing, as well as working closer with projects developing open source NFV implementations. In addition, the NFV ISG also released nine draft NFV documents for industry [comments](#) and published a document that summarizes the key concepts that are contained in those [documents](#). Those nine documents describe an infrastructure overview, the virtualized network functions architecture and the compute, hypervisor and infrastructure network domains. They also cover management and orchestration, resiliency, interfaces and abstractions, and security.

TM Forum

Another industry group that is closely associated with the development of NFV is the [TM Forum](#). The TM Forum came into existence as the OSI/Network Management Forum in 1988 with the goal of solving the systems and operational management issues that were associated with the OSI protocols. The name was changed to TeleManagement Forum in 1998 and to TM Forum in 2013.

The Forum has over 1,000 member companies, including more than 250 communications service providers. One of the ways that the TM Forum delivers value is by bringing together working groups to rapidly address specific business issues by defining standards-based tools and best practices. Early 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, ZOOM aims to identify and define new security approaches to protect infrastructure, functions and services across all layers of software and hardware. It is also a goal of ZOOM to compliment the ongoing work of ETSI and other industry leaders. As of October 2014, the ZOOM team has delivered an assessment of how virtualization impacts SLAs and is currently working on information and policy

models, NFV preparedness, and a set of operational support system (OSS) design principles needed for NFV adoption to become widespread.

The TM Forum has also been active with companies such as Microsoft to create Catalysts, which are short-term collaborative projects led by members of Forum that address operational and systems challenges. Like POCs, Catalysts are a way to quickly test new approaches and best practices. In June 2014 at the TM Forum Live! event in Nice, France there was a demonstration of fifteen Catalyst POCs including five that focused on virtualization. Four additional virtualization centric Catalysts will be demonstrated at TM Forum's Digital Disruption conference in San Jose, CA in December 2014.

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 SDN, the IETF can develop standards that complement the efforts of the Open Networking Foundation ([ONF](#)) and other relevant Standard Defining Organizations (SDOs). In the case of NFV, the IETF can possibly play a more central role in creating standards that fit into the overall architectural frameworks defined by the ETSI NFV ISG because ETSI's work is focused on frameworks and broad specifications rather than standards per se.

The IETF Service Function Chaining (SFC) Work Group (WG) currently has over forty active Internet drafts on the topic of delivering traffic along predefined logical paths incorporating a number of service functions. The basic concept of SFC is similar to ETSI NFV ISG's Virtualized Network Function (VNF)-Forwarding Graphs. Service chaining is also an important capability of SDN networks. It is likely that the IETF's work on SFC will apply to both SDN and non-SDN environments. Some of the topics being investigated by the SFC WG include:

- Service function instances discovery;
- Service function resource management;
- Service chain creation;
- Traffic flow steering rules on a router to define network forwarding paths;
- Service chain monitoring and adaptability for reliability and optimized performance;
- Information and data models for SFC and NFV.

Another area of IETF activity related to SDN and NFV is the work the IETF has done on a security architecture that is based on horizontal (a.k.a., east/west) APIs in addition to the northbound and southbound [APIs](#). One IETF SDN-specific activity focuses on centralized security services (i.e., firewalls and DDOS mitigation systems) designed specifically for SDN [environments](#). Another SDN-specific Internet draft addresses the possible application of DevOps principles to service provider software defined telecom [networks](#).

Open Platform for NFV (OPNFV)

In September 2014 the Linux Foundation, announced the founding of the Open Platform for NFV Project ([OPNFV](#)). As part of the announcement the Linus 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 initial project objectives are to:

- Develop an integrated and tested open source platform that can be used to investigate and demonstrate core NFV functionality;
- Include proactive participation of leading end users to validate that OPNFV meets the needs of the end user community;
- Contribute to and participate in relevant open source projects that will be leveraged in the OPNFV reference platform;
- Establish an open ecosystem for NFV solutions based on open standards and open source software; and
- Promote OPNFV as the preferred open reference platform.

Relationship between SDN and NFV

The majority of the material in this section comes from the corresponding section in the first chapter in The Guide. That material is being included in this section so that this chapter is self-contained. The only material in this section that isn't included in the first chapter of The Guide is the survey question about the applicability of NFV to both enterprises and service providers.

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

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

Marc Cohn functions as a liaison between the ONF and the ETSI NFV ISG. In a recent [blog](#), Cohn was quoted as saying that SDN and NFV are inextricably linked. He backed that statement up by saying that half of the use cases that the ISG has defined are cloud based and require the type of dynamic network functionality that SDN provides, but which is not provided by a traditional network [architecture](#). Two of those use cases ("Network Functions Virtualization Infrastructure as a Service" and "Virtual Network Function Forwarding Graph") are described in detail in *OpenFlow-enabled SDN and NFV Solution Brief*.

In a recent [white paper](#), ETSI made the following comments about the relationship between SDN and NFV:

"NFV creates a very dynamic network environment, driven by customers needing on-demand services and operators needing to manage utilization and performance of services. Tenant networks will come and go, and VNFs and their connectivity will change frequently to balance load across the infrastructure. The capability to programmatically control network resources (through a centralized or distributed controller) is important in an era of continuous change. Complex network connectivity topologies may be readily built to support automated provisioning

of service chains as a realization of NFV ISG Forwarding Graphs while ensuring strong and consistent implementation of security and other policies. The SDN controller maps to the overall concept of network controller identified in the NFV architectural framework, as a component of the NFVI network domain. As such, an SDN controller can efficiently work with orchestration systems and control both physical and virtual switching, as well as provide the necessary comprehensive network monitoring. However, special attention is needed to ensure that when SDN is applied to telecommunications networks, the separation of control plane and data plane does not cause additional traffic overhead, latency, jitter, etc., as well as redevelopment of existing protocols especially for switching, routing and high availability.”

The first chapter of The Guide included market research that was based on a survey that was distributed in September 2014. The respondents to that survey 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 13**.

Table 13: Perceived Relationship between SDN and NFV	
Relationship	% of Respondents
They are complementary activities in that each can proceed without the other but the value of each activity may be enhanced by the other activity.	61%
In at least some instances, NFV requires SDN	25%
Don't know	16%
In at least some instances, SDN requires NFV	10%
They are totally independent activities	6%

Some of the conclusions that can be drawn from the data in **Table 13** are:

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

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

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

As was previously mentioned, this chapter of The Guide includes market research that is based on a survey that was distributed in October 2014. The conventional wisdom is that NFV is applicable only to service providers. To test that conventional wisdom the respondents to the October 2014 survey were asked about their view of the applicability of NFV in both the enterprise and service provider environments. Their responses are shown in **Table 14**.

Table 14: Applicability of NFV	
Applicability	% of Respondents
NFV is applicable equally in a service provider and an enterprise environment	42%
NFV is applicable primarily in a service provider environment but it provides some value in an enterprise environment	40%
NFV is applicable primarily in an enterprise environment but it provides some value in a service provider environment	6%
NFV is applicable only in a service provider environment	5%
Don't know	4%
NFV is applicable only in an enterprise environment	1%
Other	1%

Only a very small percentage of IT professionals think that NFV is only applicable in a service provider environment.

Almost half of IT professionals think that NFV is equally applicable in a service provider environment and an enterprise environment.

Status 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 15**.

Table 15: Current Approaches to Implementing NFV	
Approach to Implementing NFV	% of Respondents
We are currently actively analyzing the potential value that NFV offers	39%
We are currently actively analyzing vendors' NFV strategies and offerings	24%
We currently are running NFV either in a lab or in a limited trial	21%
We will likely analyze NFV sometime in the next year	16%
We expect that within a year that we will be running NFV either in a lab or in a limited trial	14%
We currently are running NFV somewhere in our production network	13%
Other	9%
We have not made any analysis of NFV	8%
We looked at NFV and decided to not do anything with NFV over the next year	6%
We expect that within a year that we will be running NFV somewhere in our production network	6%

The data in **Table 15** 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 16**.

Table 16: Factors Driving NFV	
Factor	% of Respondents
Reduce the time to deploy new services	33%
Reduce OPEX	14%
Greater management flexibility	13%
Better network performance	12%
Reduce CAPEX	11%
Better customer experience	9%
Other	7%
No driver	2%

The data in **Table 16** 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 17**.

Table 17: Factors Inhibiting NFV	
Inhibitor	% of Respondents
Concerns about how we would do end-to-end service provisioning that includes physical and virtual resources and which may cross multiple partners' domains	30%
The lack of a compelling business case	28%
The immaturity of the current products	24%
The need to significantly reskill our employee base	19%
The need to make significant organizational changes in order to fully realize NFV's promise	17%
Concerns about security vulnerabilities	17%
The need to implement a new generation of agile OSS/BSS	17%
The need for sophisticated orchestration capabilities	15%
The immaturity of the enabling technologies	14%
Concerns about how we would evolve from a POC to broad deployment	13%
The difficulty of doing end-to-end service management	12%
The time it will take for standards to be developed and implemented	12%
The lack of a critical mass of organizations that have deployed NFV	11%
Other technology and/or business priorities	10%
The confusion and lack of definition in terms of vendors' strategies	9%
The need to make significant cultural changes in order to fully realize NFV's promise	6%
The reluctance on the part of some of our suppliers to embrace a software model	6%
No inhibitors to implementing NFV	5%
The requirement to make significant changes to our procurement processes	4%
Other	4%

The data in **Table 17** indicates:

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

- ***Concerns about end-to-end provisioning;***
- ***The lack of a compelling business case;***
- ***The immaturity of the current products.***

Of the three primary inhibitors listed above, the impact of the immaturity of the current products will diminish over time due to the natural evolution of products. The TM Forum is working to ease the challenges associated with end-to-end provisioning. It is unclear if any industry-wide source will create

a business case for NFV. It is also worth noting that as pointed out in the preceding chapter of The Guide, the lack of a compelling business case is also a major inhibitor to the adoption of SDN.

The Survey Respondents were also asked to indicate how long it would be before their organization has made a significant deployment of virtualized IT and/or network functionality. Their responses are shown in **Table 18**.

Table 18: Time Frame for Deployment	
Time Frame	% of Respondents
Already have	21%
1 – 2 years	30%
3 – 4 years	32%
5 – 6 years	4%
7 or more years	0%
Don't know/ Not Applicable	13%

The combination of the data in **Table 18** plus the data in **Table 15** that highlighted the significant commitment that IT organizations have made in analyzing NFV indicates:

Within a few years, the majority of IT organizations are likely to have made a significant deployment of NFV.

Use Cases and Proof of Concept

As mentioned, the 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.

As mentioned, as of October 2014, 25 POCs has been defined. 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 the nine potential use cases that ETSI identified and which are described below. Also described below are some NFV-related POCs. The uses cases are generic. However, given the nature of a POC, the POCs involve vendors and/or service providers.

ETSI NFV Use Cases

The ESTI NFV ISG has identified nine potential use cases for NFV. This section of The Guide provides an overview of these possible use cases. A thorough description of the use cases is available on the [ETSI web site](#).

NFV Infrastructure as a Service (NFVlaaS)

NFVlaaS 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, NFVlaaS 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 NFV/IaaS 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.

The Survey Respondents were given a listing without description of the nine use cases that ETSI has defined and they were asked to indicate the two use cases that they think will gain the most traction in the market over the next two years. Their responses are shown in **Table 19**.

Table 19: Interest in ETSI Use Cases	
Use Case	% of Respondents
Network Functions Virtualization Infrastructure as a Service	51%
Virtual Network Function as a Service (VNFaaS)	37%
Virtualization of Mobile Core Networks and IMS	32%
Virtual Network Platform as a Service (VNPaaS)	22%
Fixed Access Network Functions Virtualization	13%
Virtualization of CDNs (vCDN)	12%
Virtualization of Mobile base station	11%
Don't know	6%
Virtualization of the Home Environment	4%
VNF Forwarding Graphs	1%
Other (Please specify)	1%

The data in **Table 19** indicates:

While IT organizations have interest in a number of the ETSI-defined use cases, by a wide margin they are most interested in the Network Functions Virtualization Infrastructure as a Service use case.

TM Forum Catalyst POCs

In June 2014 at the TM Forum Live! event in Nice, France there was a demonstration of 15 Catalyst POCs including the four POCs discussed below.

Closing the Loop: Data-driven network performance optimization for NFV & SON

In this context *closing the loop* means collecting and analyzing data to identify how the network can be optimized and then implement those changes. This POC showed how network operators can use Self-Organizing Networks (SON) and Network Functions Virtualization (NFV) in tandem to automate closing the loop and improve performance for customers.

Participants in the project included Mycom, TEOCO and Wipro, while Telecom Italia and Reliance Communications were the champions of the project. The POC demonstrated how to build a closed loop using key performance indicators, including network performance, customer experience and service quality data, to enable network changes, optimization and self-healing. TM Forum's Performance Management Interface was used along with 3rd Generation Partnership Project (3GPP) interfaces to link operational support systems (OSS) with network elements, both physical and virtual. As part of the demonstration in Nice, configuration and performance data was collected from a mobile network and then the data was analyzed to identify where problems exist or where there is potential for improvement.

CloudNFV™: Dynamic, data-driven management and operations Catalyst

This POC builds on TM Forum's Information Framework to create a meta-data model using **active virtualization**, a term coined by the CloudNFV™ [consortium](#). That consortium is a group of NFV technology suppliers working together to develop solutions aimed at solving the problem of how to link orchestration systems in a virtual network with the other business and operational support systems that control network policy. The specific challenge this POC is addressing is that without these connections, services like dynamic quality of service likely won't work at scale. Participants in the CloudNFV™ Catalyst include EnterpriseWeb, Huawei and Qosmos, plus several other companies supplying hardware and software components. Champions of the project include AT&T, BT, Orange and Sprint.

Orchestrating Software-Defined Networking (SDN) and NFV while Enforcing Service Level Agreements (SLAs) over Wide Area Networks (WANs)

One set of challenges that this Catalyst addressed are the challenges that service providers face when offering private clouds to enterprises and managing SLAs in a virtualized environment.

Another set of challenges are the challenges that geographically diversified enterprises encounter when integrating data centers.

This Catalyst used OpenFlow version 1.3 to demonstrate full OpenFlow 1.3 interoperability with OpenFlow-enabled controllers and it implemented a gateway for public to private data center connectivity. A number of NFV's Virtualized Network Functions (VNF) were run and a cloud management system monitored and adjusted parameters on a virtual machine and on the OpenFlow controller to the desired performance levels. This illustrates how performance levels in an enterprise data center or network can be changed and it also demonstrates how SLAs can be adjusted quickly.

As part of creating this Catalyst the team developed a cloud reference architecture that can be used to help firms design and operate data centers and to help service providers offer private clouds and digital services. This reference architecture can be connected to TM Forum APIs for SLA management and billing. In addition, the TM Forum Application Framework (TAM) can be used to specify the approach for creating the infrastructure design and implementation.

The project has worked with several groups that are establishing best practices and recommendations for offering cloud services. These include the Open Networking Foundation, the Open Data Center Alliance (ODCA), the Open Mobile Alliance, and ETSI.

Service bundling in a B2B2X marketplace

This Catalyst showed how a buyer can bundle a collection of services sourced from different suppliers and deliver them seamlessly to a customer in a business-to-business or business-to-business-to-consumer arrangement. These components could include traditional network access products, as well as NFV and infrastructure-as-a-service products. Catalyst participants included Cisco Systems, DGIT and Liberated Cloud, and the champions of the project were AT&T, NBN Co, Uecomm, Ultrafast Fibre and Vodafone New Zealand.

The B2B2X Catalyst combined a cloud service, a software-defined network mocked up by Cisco and a fiber access service provided by NBN Co of Australia. The three components were bundled into a

product that combines high-speed Internet, firewall service and virtual servers as a bundled service for small to mid-sized businesses.

In order for the process to work, products were defined with characteristics that are orderable attributes of the product, and those characteristics were encoded in product definitions. So for a business service, for example, orderable attributes can include things like the service level agreement and throughput speed. The Catalyst showed the product definitions and how they are built and expressed by using the dynamic extensibility of the Information Framework and by creating templates for dynamic data, the specifications for which are shared between the two provider organizations.

Private POCs

In addition to the POCs being driven by organizations such as ETSI and the TM Forum, a number of vendors are conducting private POCs with one or more service providers.

Virtualized S/Gi-LAN

These trials enable the operator to develop expertise necessary to conduct full life-cycle management of the virtualized applications that reside between the mobile packet gateway (PGW) and the Internet—a domain commonly referred to as either the Gi-LAN (3G) or the SGi-LAN (LTE). As the predominant application in the Gi-LAN and SGi-LAN, the Citrix ByteMobile Adaptive Traffic Manager (ATM) is part of these network virtualization trials.

Citrix is partnering with operators to develop a solution that: a) is readily integrated with an operator's chosen NFV management and operations (MANO) framework; and b) meets NFV requirements such as rapid service provisioning. The Citrix ByteMobile ATM function must scale in parallel with broadband data traffic growth and an NFV implementation will enable the automated scaling of this function within the S/Gi-LAN domain. To achieve this end, Citrix offers a complete virtualized application stack that includes the virtual Adaptive Traffic Manager and the Citrix NetScaler VPX virtual application delivery controller. In preparation for expected operator demand, Citrix has conducted lab demonstrations of this application stack using both XenServer/CloudPlatform and KVM/OpenStack as hypervisor /virtual infrastructure manager.

The Operational Implications

It is very positive for the development and deployment of NFV that organizations such as ETSI and the TM Forum are currently conducting a wide range of POCs. However, even if a POC is successful it can be very challenging to deploy that solution into a production environment. To quantify that challenge, The Survey Respondents were told to assume that one of the NFV POCs has been a technical success. They were then asked to indicate how much of an effort they thought it would require in order to take the solution that formed the basis of the POC and implement it broadly in production inside of their company. Their responses are shown in **Table 20**.

Table 20: Effort to go from POC to Production	
Amount of Effort	% of Respondents
A tremendous amount	7%
A very significant amount	23%
A significant amount	35%
No more of an effort than is required to implement any new technology or architecture; i.e., virtual servers	17%
Less than the typical amount of effort	2%
Don't know	14%
Other (Please specify)	1%

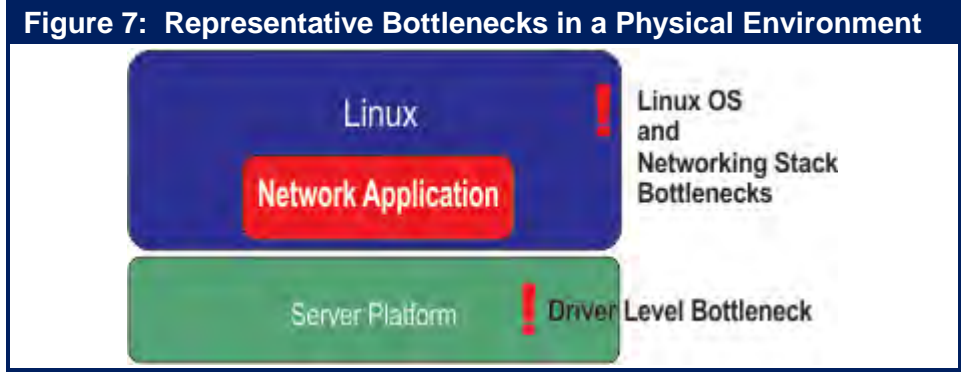
The data in **Table 20** indicates:

The majority of IT organizations believe that even if a NFV-related POC is successful, it will take between a significant and a tremendous amount of effort to broadly implement that solution in production.

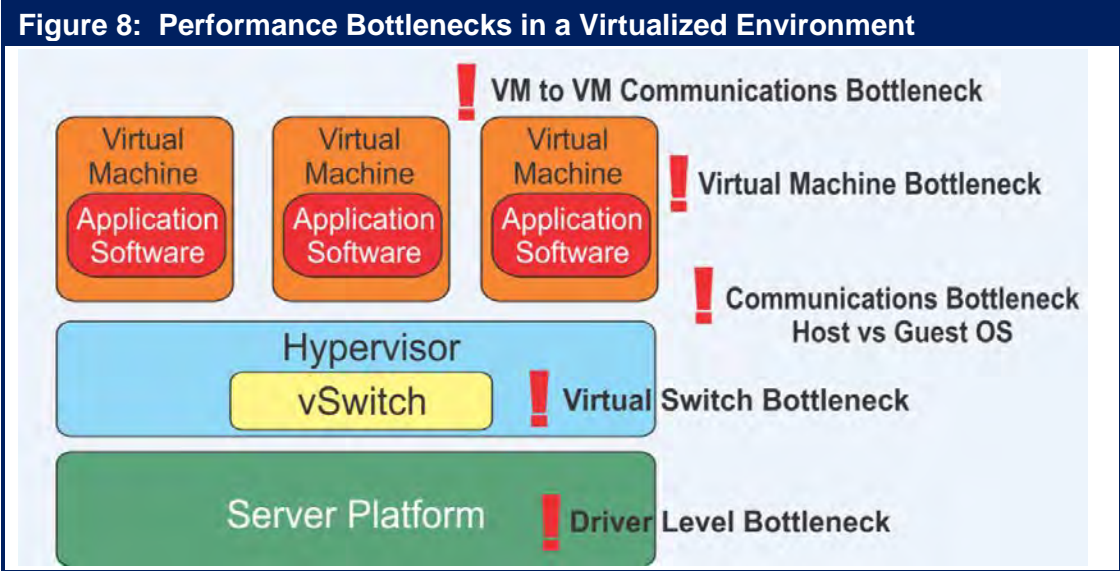
Performance Limitations

In order to obtain the potential cost and agility benefits of a software-based approach to providing IT functionality, it must be possible to achieve the same or greater performance in a software-based environment as is possible in a traditional hardware-based environment. However, 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](#)".

Performance bottlenecks are not unique to virtualized environments. For example, some of the bottlenecks that occur in a physical environment are shown in **Figure 7**.

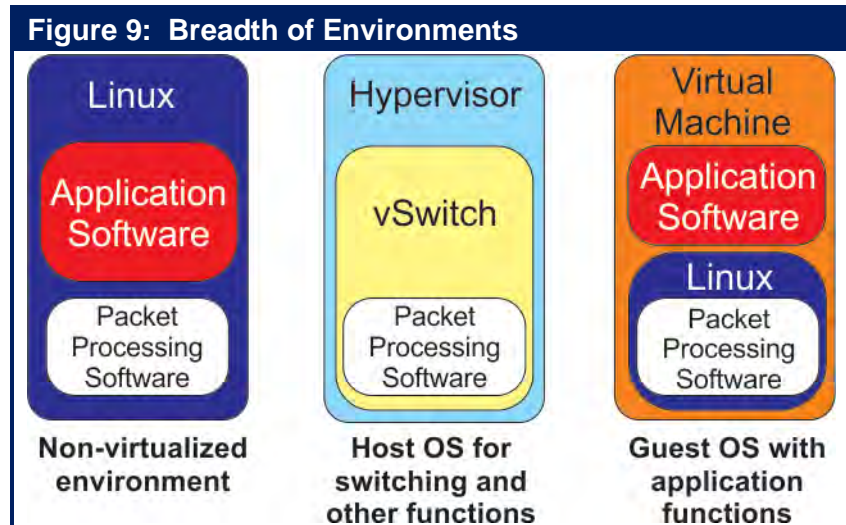


Unfortunately, as shown in **Figure 8**, as IT organizations adopt a virtualized environment, the performance bottlenecks multiply.



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 9**, 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;
- Portability: Live migration of VNFs over disparate hardware platforms and from one server to another.



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.

The types of performance improvements that are possible are significant. For example, it is possible to leverage packet processing software to accelerate the performance of a virtual switch, such as Open vSwitch, by a factor of 10 or more. Some examples of high performance Virtual Network Functions (VNFs) designed with effective packet processing software include:

- An accelerated TCP/UDP stack that enables the building of products such as stateful firewalls, DPI engines, cloud servers and web servers that support millions of concurrent sessions and also support session setup rates above one million sessions per second.
- A high performance IPsec stack that can sustain more than 190 Gbps of encrypted traffic on a single server.
- High performance and capacity for encapsulation protocols such as GRE, GTP, PPP, L2TP. An example of this is a vBRAS server that can handle 256,000 PPPoE tunnels with 70 Gbps throughput.

End-to-End Management

Management Challenges

The adoption of NFV poses a number of significant challenges that must be overcome in order to ensure the ability to continue 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 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.

- **VNFs will be new types of components in the network.** In order for a service provider to be able to mix and match VNFs from a variety of network equipment vendors It will necessary for the industry to establish some standards for the functionality of VNFs, the hypervisors that are supported, and the management interfaces they present to end-to-end management systems. For their part, end-to-end management systems will need to support these standards as they evolve.
- **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.

Management Direction

As mentioned, the TM Forum is working 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 intends to identify and define new security approaches to protect infrastructure, functions and services across all layers of software and hardware.

ETSI is also 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, virtualised 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 standardisation and identification of gaps in existing systems and platforms.”

The Organizational Implications

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 21**.

Table 21: Impact of NFV on Organizational Structure	
Impact	Percentage of Responses
Very Significant Impact	6%
Significant Impact	28%
Moderate Impact	24%
Some Impact	19%
No Impact	12%
Don't Know	9%

The data in **Table 21** 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.

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

- It will change the way out networks are operated and managed;
- It will require us to have a mature and more streamlined end to end service management function with better understanding of what will benefit our client's and the value we can provide to them;
- We will need to overhaul of our networking architecture;
- It will change how we provision and deliver service to our clients;
- It will require a reorganization of the groups that plan and operate the network;
- We will need to productize and update our provisioning processes;

- It will impact us by being another step along the way to our company being a Service Provider more than just a Telecom Provider.

In addition to the changes listed above, one respondent expressed concern that his company would suffer lost productivity during the transition to NFV.

When asked the same question, a number of enterprise respondents commented that it would require them to change how they implemented SLAs, how they developed a business case and it would cause them to rethink their business models. One respondent mentioned that it would also require their IT organization to change its culture. Other comments from the enterprise respondents include:

- It will reduce the time it takes us to deploy new services;
- It will give us greater management flexibility;
- We will need to adopt a new approach to service provisioning and management;
- It will cause us to consolidate our physical platforms;
- It will change how we do network planning;
- We will need to determine how we are going to orchestrate end-to-end systems.

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 22**.

Table 22: Impact of NFV on Employee Skills	
Impact	% of Responses
Very Significant Impact	8%
Significant Impact	35%
Moderate Impact	22%
Some Impact	19%
No Impact	6%
Don't Know/Other	11%

The data in **Table 22** 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 nearly half of all IT professionals.

Some of the answers from service provider respondents when asked to indicate the type of impact that NFV will have on the skill base of their company’s employees include:

- The sales and marketing people are going to have to learn a whole new way of thinking;
- We need a plan for the acquisition, evolution and retention of the required talent and skills;
- We have to transition to more software-based skills from the current set of hardware-based skills;
- We need to transition to where we have more computer science skills in our organization;
- We need to develop a new training curriculum.

One of the survey respondents expressed their concern about how much of a transition has to be made by commenting that “Our network staff is IT illiterate.”

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

- We will need to know multiple technologies;
- It will make virtualization know how the most significant skill;
- We will need to think in software and end-to-end terms rather than in component terms;
- This will create the requirement to become more of a programmer than was required in a traditional network role;
- It will require the skills to drive the integration between legacy equipment and management systems and NFV management systems;
- We will need to modify our change management, incident and problem management processes;
- This is a paradigm shift for network engineers to retool and relearn new methods.

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. According to a recent [Information Week Report](#), eighty-two percent of the IT organizations that implemented DevOps saw at least some improvement in infrastructure stability and eighty three percent saw at least some improvement in the speed of application development.

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, behaviors, 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 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.

~ Continued on page 83 ~



The Cloud Network Unbound

Virtualized and automated networking across datacenters and branch offices

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

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

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

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

Our solutions close the gap between the network and cloud-based consumption models, creating an infrastructure in which network resources are as readily consumable as compute and storage resources. Our approach enables enterprises to transform the way they build and use their networks, which has a profound effect inside and across multiple datacenters and across the wide area network.

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

Take advantage of a fully virtualized services platform

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

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

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

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

Responsive datacenter networking

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

Virtual private networking on your terms

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

Nuage Networks SDN solutions are specifically designed to:

Simplify operations for rapid service instantiation	Address changing business requirements with flexible, adaptable services	Support massive scalability and hybrid models with secure, open infrastructure
<ul style="list-style-type: none"> Define network service requirements in clear, IT-friendly language Bring services up using automated, policy-based instantiation of network connectivity Dramatically reduce time to service and limit potential for errors 	<ul style="list-style-type: none"> Adapt datacenters and private networks dynamically Detect newly created and updated virtual machines within the datacenter and respond automatically by adapting network services according to established policies, instantly making available new applications to all users regardless of location 	<ul style="list-style-type: none"> Benefit from distributed, policy-based approach that allows multiple virtualization platforms to interoperate over a single network Optimize the datacenter network and private network by separating service definition from service instantiation

Nuage Networks SDN solution components

Nuage Networks VSP is the first network virtualization platform to address modern datacenter requirements for multi-tenancy, full-featured routing and security at scale. It is a software solution that transforms the physical network into a simple to manage, rack-once and wire-once, vendor-independent IP backplane. As a result, network resources within and across datacenters can be treated as an elastic resource pool of capacity that can be consumed and repurposed on demand.

Nuage Networks VSP integrates seamlessly with wide area business VPN services. It is also particularly effective when deployed with Nuage Networks VNS for a cloud-optimize network that spans the datacenter right out to your remote locations.

NU•ÂHJ: FROM FRENCH, MEANING "CLOUD"

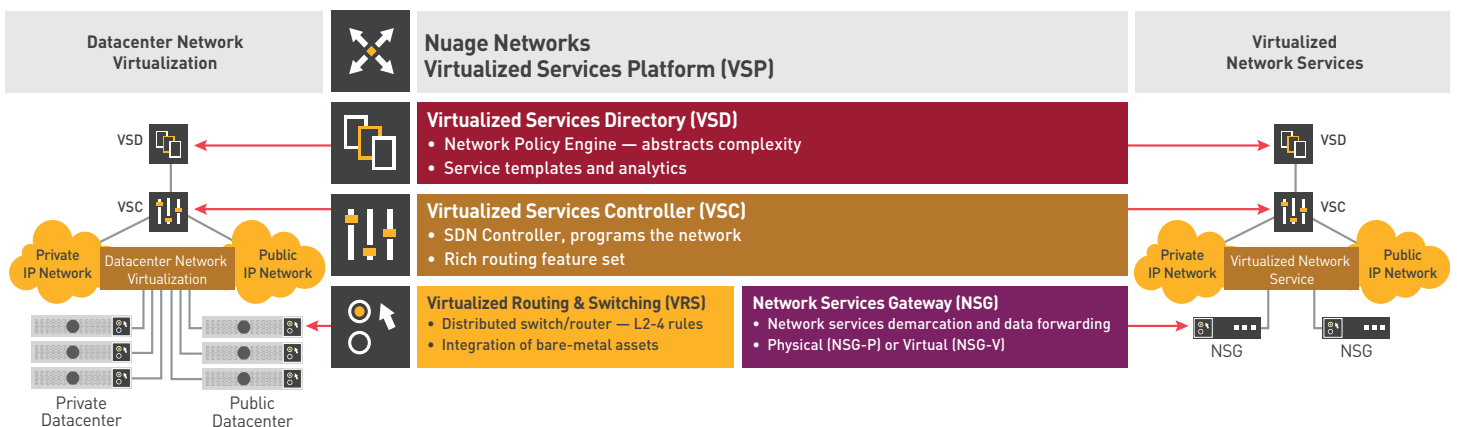
The cloud can be more than what it is. In fact, it needs to be. When we founded Nuage Networks, it was with the idea that it's time for the cloud to come of age. From the beginning we recognized the unique challenges that cloud service providers and large enterprises face delivering and managing large,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

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

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

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

Nuage Networks SDN Portfolio



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

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

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

Key benefits from the Radware SDN network service infrastructure include:

- **More intelligent application delivery and security decisions** throughout the network break existing network barriers when developing business applications. Every application everywhere is entitled for advanced services.
- **Simpler implementation** of network services allows improved operational efficiency of network management alongside application changes. Not every project needs to become a networking project.
- **Lower overall network service solution costs** – as network service delivery is partially offloaded to the SDN, there is no need to invest in excess network service appliances and capacity. Deploy network services as needed, and use by many tenants and applications throughout the datacenter.
- **Greater scalability** – scale your network services throughout the network. No more limited areas are protected or load balanced. Offer uniform services throughout the SDN.
- **Easier operation** – changing and managing security and ADC functionality becomes simpler as the deployment operates as if it is centralized. Not only does SDN streamline network operations, but Radware SDN applications streamline network service operations. In addition, API to various orchestration systems enables to improve the overall control and automation of network services.

DDoS Protection as a Native SDN Application

[DefenseFlow](#) is an SDN application that enables network operators to program the network to provide DDoS protection as a native network service. DefenseFlow features an adaptive behavioral-based DoS attack detection engine and a traffic diversion mechanism that utilizes the programmable characteristics of the software defined network elements for attack cleansing. Designed as part of the Radware SDN application framework, DefenseFlow delivers a security control plane and operates in traditional network environments while enabling to migrate to customer's future, SDN-based networks.

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

DefenseFlow equips network operators with the following key advantages:

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

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


SDN & NFV for a Scalable Application Delivery Network

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

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

ElasticScale offers network operators the following key features and benefits:

- Ultra scalable traffic steering solution (80Gbps-1Tbps and beyond)
- Based on industry leading, carrier grade Alteon load balancing product line
- Support for leading hypervisors (oXen/KVM/Hyper-V/ESXi)
- Compatible with leading SDN controllers; OpenDaylight, Cisco XNC, NEC pFlow & HP Flare
- Seamless integration with OpenStack and vCloud Director
- Runs over any physical SDN network equipment



Partnering for Success: Our SDN Ecosystem

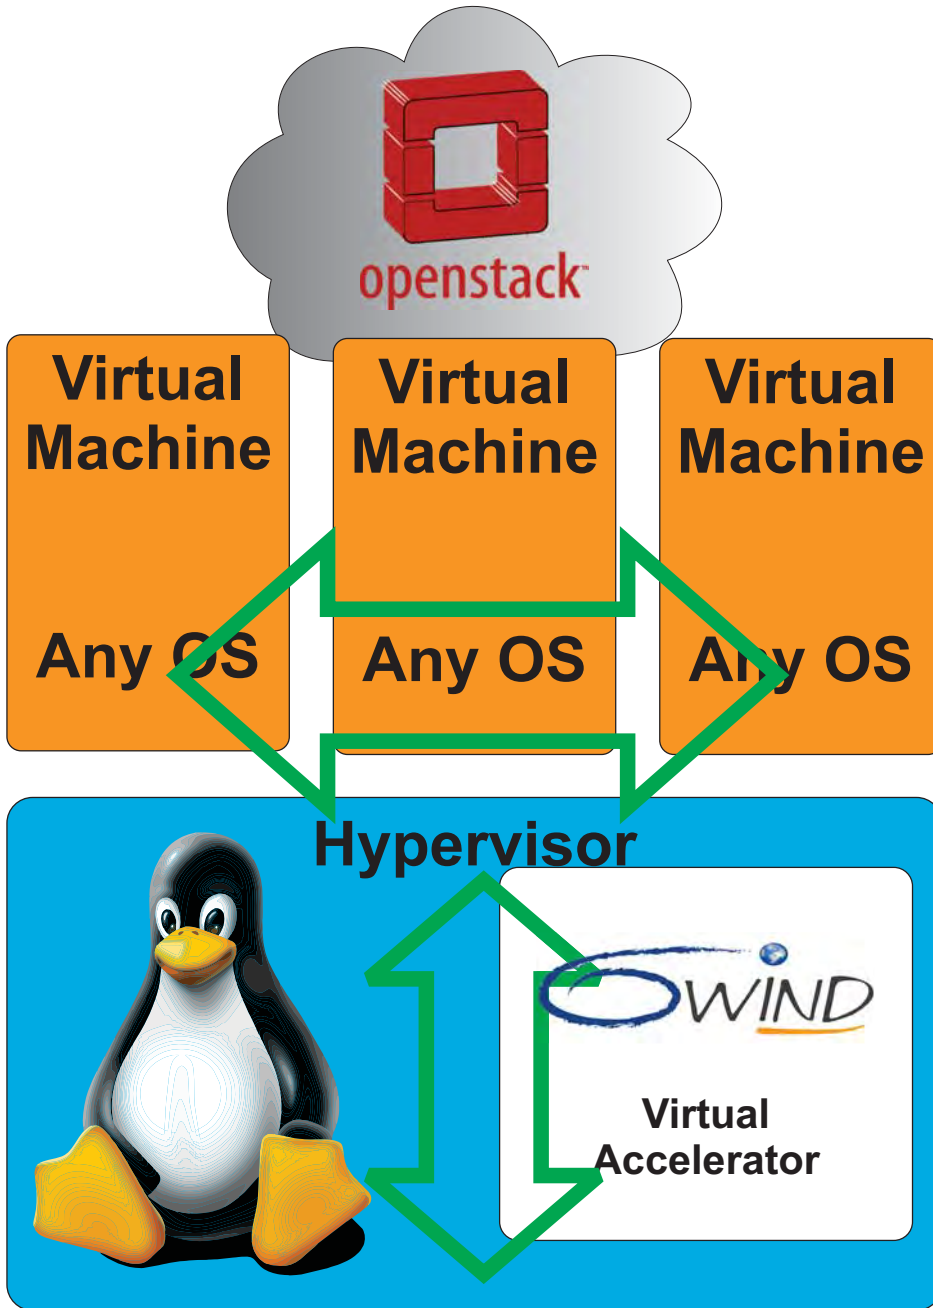
The SDN eco-system is a critical focus for Radware. Through partnerships with the industry's leading SDN forums and vendors, Radware can ensure customers that our application delivery and security solutions integrate successfully into target architectures. Radware is an active contributor in the following industry and vendor SDN initiatives: Cisco Application Centric Infrastructure (ACI), HP Virtual Application Networks, NEC, Mellanox, Alcatel Lucent, ETSI, Open Daylight Project, and the Open Networking Forum (ONF). Radware is also a member of VMware's NSX partner ecosystem for network functions virtualization (NFV).

Learn More

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

6WIND Virtual Accelerator

Enable NFV And Virtual Networking



Transparent OpenStack orchestration support

High bandwidth for VM performance, density and communications

Complete virtual networking infrastructure

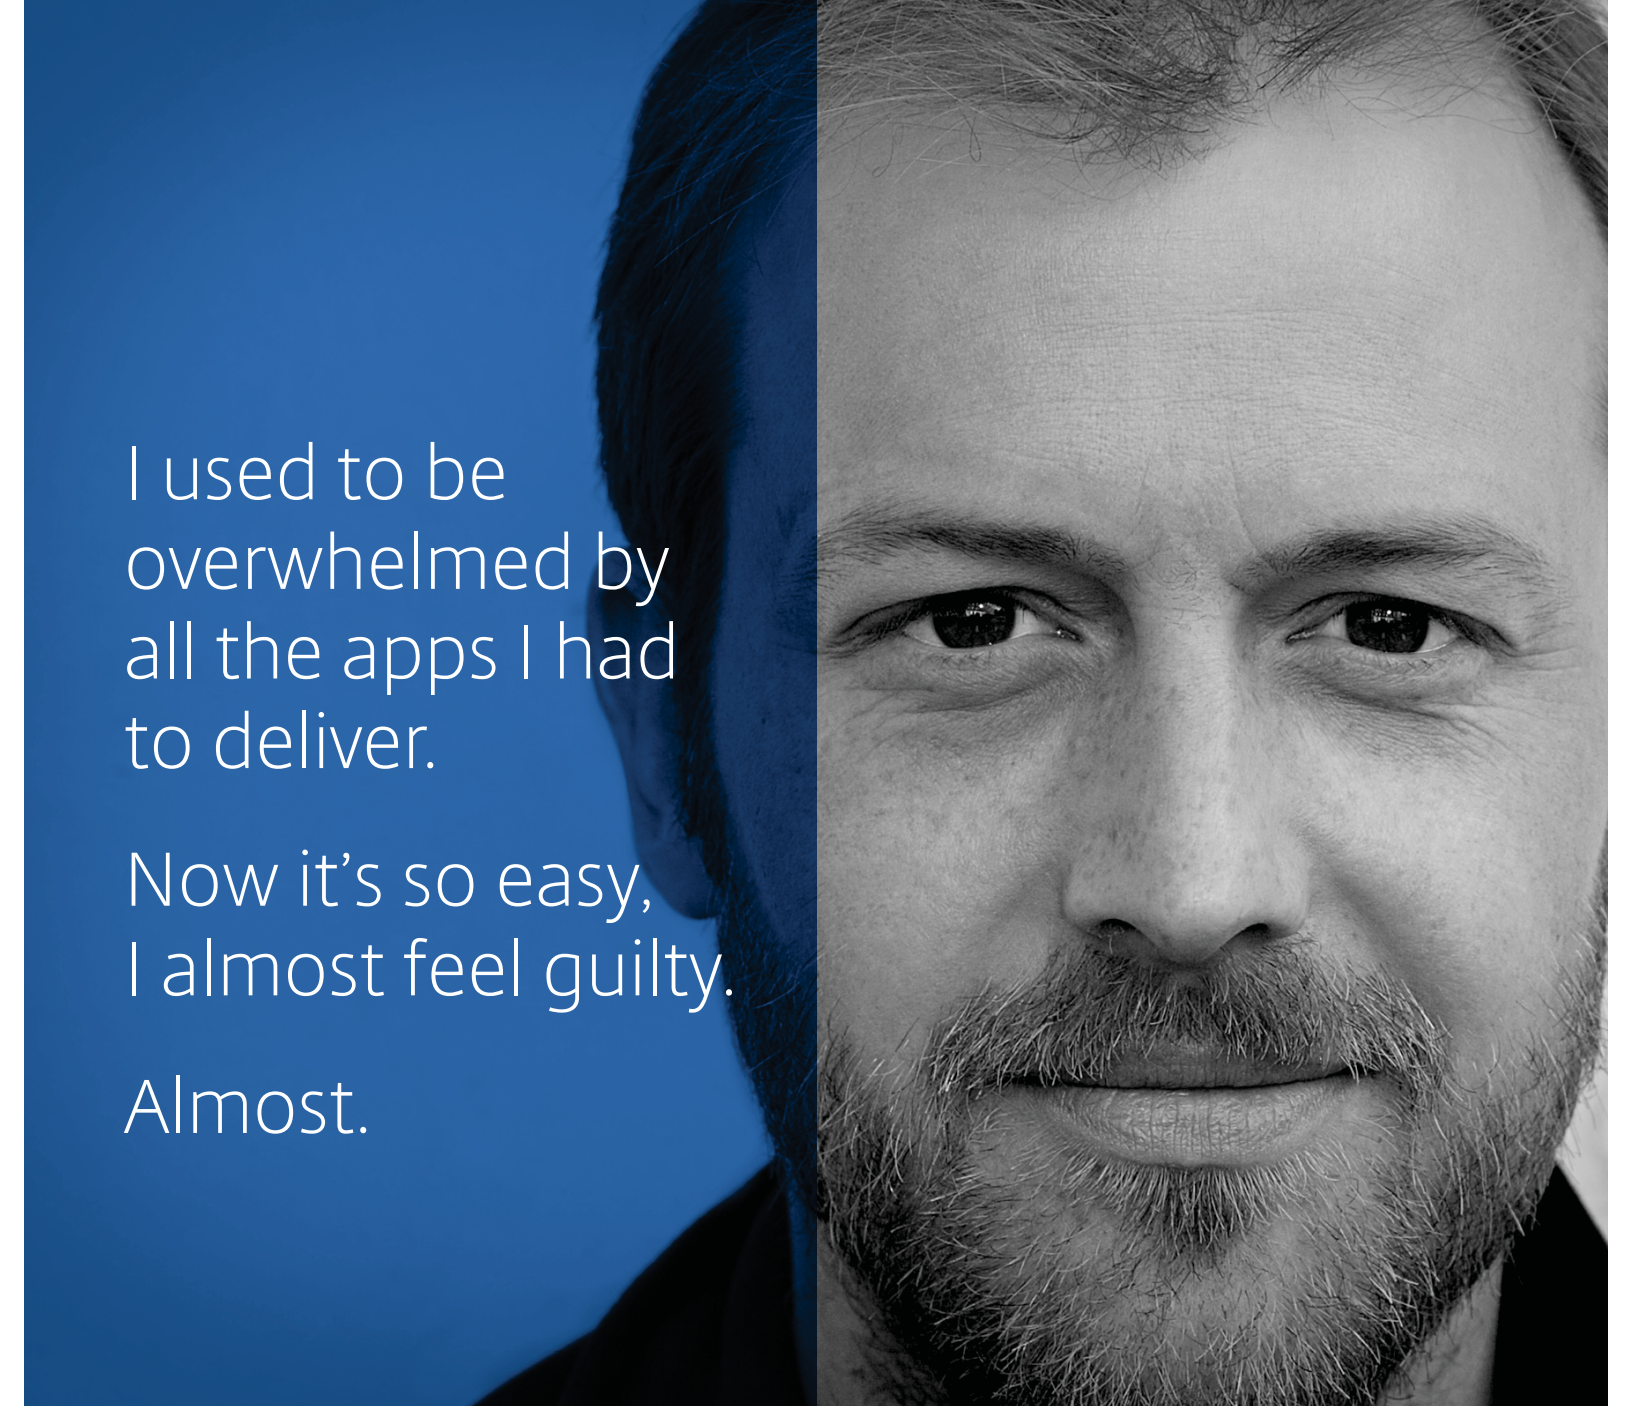
Support for Open vSwitch and Linux Bridge with no modifications

Network hardware independence for seamless hardware upgrades



Wire Speed Virtual Switching
From Common Hardware

www.6WIND.com



I used to be
overwhelmed by
all the apps I had
to deliver.

Now it's so easy,
I almost feel guilty.

Almost.

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



NetScaler with TriScale
SOFTWARE SMART. HARDWARE STRONG.

CITRIX[®]

www.citrix.com/netscaler



AUTOMATE YOUR CLOUD WITH **aCLOUD SERVICES ARCHITECTURE**

Integrate dynamic
services into your
Cloud Data Center

www.a10networks.com



The SDN and NFV Ecosystem

The SDN Ecosystem

One measure of the extent of the SDN ecosystem is that there are currently more than 100 members of the Open Networking Foundation ([ONF](#)). 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 organizations 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. As is explained below, in some instances there can be a very wide range in terms of the functionality provided by the members of a given category.

Merchant Silicon/Chip Vendors

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

Representative Members:

- Broadcom
- Intel
- Marvell
- Mellanox

HyperScale Data Centers

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

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 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](#) are currently using SDN and Network Function Virtualization (NFV)² to enable them to offer a new generation of WAN services and [Verizon](#) has announced a trial based on using SDN to enable a new generation of data center to data center WAN services. AT&T has announced its interest in using both SDN and NFV to change how it offers services to its [customers](#).

Representative Members:

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

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
- Extreme Networks
- HP
- Meru Networks
- 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.

² NFV was explained in the preceding chapter of The Guide

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 solution. There is the possibility that over time that a large number of independent software vendors (ISVs) will also provide these services.

Representative Members:

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

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

Standards Bodies and Related Communities

Value Proposition: Some of the members of this category develop use cases, architectures and drive POCs. In some cases, the work of these members helps to clarify the problems that need to be solved and the standards that need to be developed. 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.

Representative Members:

- ONF³

³ The ONF is active developing a standards based protocol (OpenFlow) for communicating between a SDN controller and a network element. Its scope of work, however, is broader than just developing OpenFlow.

- IEEE
- IETF
- MEF
- OpenStack
- OpenDaylight

Providers of SDN Controllers

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

Representative Members:

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

Providers of Telecom Service Provider's Infrastructure/ Optical Networking

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

Representative Members:

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

Server Virtualization Vendors

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

Representative Members:

- Citrix
- Microsoft
- VMware

The NFV Ecosystem

One measure of the extent of the NFV ecosystem is that there are currently more than 90 organizations that are full members of the ETSI NFV Industry Specification Group (ISG), with approximately another 140 organizations listed as participants. 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 is explained below, in some instances there can be a very wide range in terms of the functionality provided by the members of a given category.

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.
- NTT Corporation

Network Systems and Electronic Equipment Vendors

Value Proposition: This category includes a very wide variety of the components of service provider network infrastructures, and, in some cases, enterprise network infrastructure. In order to accommodate NFV and SDN these vendors will need to take at least some of the control functionality that has typically resided in their products and now rely on that functionality being provided by an SDN controller or NFV management system or orchestrator. These vendors need to implement protocols in their products to support communication with central control entities. In a number of cases, vendors will be called upon to migrate the functionality of their products from dedicated hardware platforms to virtual appliances that can run on industry standard servers.

Representative Members:

- ADTRAN Europe Ltd
- Cisco Systems
- Ericsson
- IBM Europe

- Huawei Technologies (UK) Co. Ltd
- Spidercloud Wireless Inc.

Merchant Silicon/Chip Vendors

Value Proposition: These vendors are in a position to provide hardware support for protocols that support SDN and NFV in switching chips and other semiconductors. This will have the effect of increasing the speed and scalability of infrastructures that support NFV as well as the platforms that support VFNs.

Representative Members:

- Broadcom
- Freescale Semiconductor
- Intel
- Marvell

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:

- 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.

Standards Bodies and Related Communities

Value Proposition: Some of the members of this category develop use cases, architectures and drive POCs. In some cases, the work of these members helps to clarify the problems that need to be solved and the standards that need to be developed. 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.

Representative Members:

- ETSI
- 3GPP
- MEF
- ATIS
- IETF
- OPNFV
- OpenStack
- OpenDaylight
- 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 ecosystem as one of the leading Service Performance Management (SPM) providers. The SPM vendor landscape in virtual SDN and hybrid environments is very similar to what it is in physical environments. The Application Performance Management players are the same in physical and virtual environments. The Network Performance Management vendors include the traditional players in physical environments as well as the leading SDN / virtualization vendors such as VMware, Microsoft and HP. The bottoms up management approach in virtual/SDN and hybrid environments is conceptually the same as it is in physical environments. The deficiencies of this approach are magnified due to the increase in the overall SPM related big data that needs to be collected, normalized, contextually analyzed and visualized across virtual/SDN and hybrid environments.

For the next few years, the industry will operate in a hybrid environment. Some services will migrate to an NFV/SDN environment while others remain on traditional purpose-built hardware. Even for those services that do migrate to a virtual environment, given the nature of a carrier network, users will still need to traverse functions and services that reside on purpose-built hardware and virtual environments.

For an operator to truly manage and understand the user experience, it will require the ability to have an end-to-end view of the network and services. A view of just the virtual environment or just the traditional environment will not suffice in providing the level of service and experience demanded to truly leverage the agility provided by a virtual infrastructure.

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 breadth of the new challenges.

That value-add includes:

- Pervasive end-to-end visibility into service delivery;
- Reduced MTTR with Proactive Service Triage;
- Enhanced IT Efficiencies through Common Operational View;
- Scalable service delivery management architecture.

NetScout believes that it is uniquely positioned within the industry to be the market leader in monitoring both the hybrid environment as well as the future all-virtual environments. NetScout justifies that statement by pointing out that today it provides a virtual Adaptive Session Intelligence (ASI) probe for VMware NSX environments. This enables NetScout to extend the monitoring of both enterprise and carrier-scale service delivery infrastructure into both virtual and hybrid environments. This capability is necessary to address the need to monitor services deployed in more complex environments. These environments include physical and virtual application workloads that exchange greater volumes of traffic between themselves and that also experience higher risk of service degradations. In both a

physical and a virtual environment, service degradation often results in a lower quality of end-user experience and may result in increased churn for the 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 monitoring companies to be as efficient as possible with consuming both compute and networking resources in passing monitoring information for performance management. This is another example of where NetScout believes that its ASI technology and the ability to capture, process, and create metadata on packet flow data will be a critical success factor.

What are the proof points?

NetScout claims to have 20,000 of the world's largest enterprises, government agencies, and more than 165 service providers as customers. NetScout believes that the breadth of its customer base combined with their integration with VMware's NSX environment means that NetScout is uniquely positioned to become the leader of SPM in the virtual and hybrid environments.

In the service provider space, NetScout's nGeniuONE platform provides wireless, cable, and wireline network operators with end-to-end network and service performance management. The nGeniuONE platform provides both enterprises and operators a single, monitoring infrastructure for monitoring the hybrid environments of today and the all virtual environments of tomorrow.

NetScout believes that additional proof points of the value-add of its SPM solution in virtual and SDN/NFV environments are based on:

- Shortcoming of existing management tools
 - Existing management tools, such as VMware's vCenter, don't offer end-to-end SPM in hybrid environments.
- Same operational best practices as in the physical environment
 - VMware's vRealize Suite used for cloud operations offers IT the following service troubleshooting tools
 - Syslog – Log analytics using vCenter Log Insight to analyze large volumes of data.
 - NetFlow – class of service, congestion, flows, # of flows (UDP, TCP...)
 - "Deep Troubleshooting" with DPI tools such as Wireshark
 - Manually pinging hypervisors (App VMs, vFW...), check vFW rules (centralized UI – basic ACL using NSX attributes)
 - Configuration management, orchestration, dashboard, capacity mgmt., application awareness (discovery & dependency mapping)
 - These tools are very similar in their scope and functionality to traditional mgmt. tools used in physical environments. Hence SPM solutions have a similar, but significantly augmented, value proposition in virtual and hybrid environments
- Unique agility requirements

Cisco

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

Cisco believes that its Application Centric Infrastructure (ACI) is an entire open SDN ecosystem unto itself. Cisco markets the SDN/ACI-capable network devices and the SDN controller. Cisco has also defined a policy model and created the communication protocols and interfaces between devices, controller and orchestration platforms. Cisco has stated that its open, extensible environment includes over 60 ecosystem partners, including ACI-compliant network, security and services devices, monitoring, analytics and DevOps solutions, as well as cloud automation platforms.

What is your value add?

Cisco's application-centric approach to SDN:

- Extends beyond network devices to include L4-7 services, security, and eventually servers and storage;
- Includes a policy model which is defined in terms of application requirements and which reflects business activity and requirements, making it easier to align IT with business strategy;
- Applies application policies equally across physical and virtual environments, so the solution is not just an overlay network that has to be managed separately from the physical infrastructure.

What are the proof points?

The proof points include:

- A number of case study [videos](#)
- Other case [studies](#)
- [Awards](#)

A10

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

IT organizations are evolving their IT strategy by adopting various cloud computing models and SDN architectures for their internal private data centers in order to achieve automation, business agility, and dramatically reduce operational costs. These organizations need an equally automated and agile L4 - L7 network services architecture to ensure that application networking and security policies are fully integrated within these emerging cloud data center architectures, and to deliver equal automation and cost of ownership benefits.

As part of the industry's SDN ecosystem, A10 Networks delivers a portfolio of products and solutions that enable seamless integration with cloud orchestration platforms and SDN network fabrics through API calls to dynamically provision application and security policies per tenant.

As part of the industry's NFV ecosystem, A10 Networks delivers virtualized network functions on its vThunder virtual appliances. Automation through OpenStack and integration with on-demand licensing makes it possible to turn up new services for customers as they are needed, and tear them down once they're no longer needed. In addition to flexibility, the A10 appliances allow customers to optimize performance so they can maximize their investment in resources

What is your value add?

Integration with leading SDN networks ensures that network and security policies are applied on any of A10's appliances for automated L4 - L7 services provisioning. Overlay and SDN fabric integration ensures automated provisioning of network segmentation and security policies on a per tenant basis. In addition, A10's NFV solutions allow virtualization of L4 - L7 services so they can be chained together to create customized communications services quickly and as needed.

What are the proof points?

There are several case studies that demonstrate A10's value proposition. This includes:

- <http://www.a10networks.com/resources/files/A10-CS-80103-EN.pdf>
- http://www.a10networks.com/resources/files/A10-CS_Micron21.pdf
- Additional proof points will be available in Q1 2015

Alcatel-Lucent

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

Alcatel-Lucent (including Alcatel-Lucent's venture Nuage Networks) offers a comprehensive SDN and NFV solution architecture which is comprised of the Cloudband™ management system, Nuage Virtualized Services platform, Motive Dynamic operations and solution specific virtualized software. Alcatel-Lucent's design goal is to enable scaling of networks with virtualized networking and communications solutions including LTE packet core, VoLTE IMS architectures, virtualized CDN, virtualized RAN, virtualized routing as well as tuneable and scalable packet optical and routing solutions.

What is your value add?

Within the SDN and NFV portfolio, Alcatel-Lucent's focus is on relevant virtualization of its existing networking and communications portfolio and providing the core networking infrastructure to support virtual functions, with its Nuage Networks division for SDN control, Cloudband, a platform for NFV service orchestration, and Motive for dynamic operations of virtualized operations.

Alcatel-Lucent stated that all of its products are developed within an open standards based philosophy and are sold as best of breed solutions on their own or combined within the broader Alcatel-lucent framework. The key relevant products within the Alcatel-Lucent NFV portfolio are:

- vRAN with NFV offering virtual network functions for control, performance and delivery optimization at the RAN level. Alcatel-Lucent's first vRAN platform is already in commercial service with its vRNC solution, which facilitates advanced RNC requirements for geo-redundancy, hitless software upgrades, load balancing, and dynamic reconfiguration.
- Nuage Networks VSP interconnects multi-tenant infrastructures and hybrid clouds with an enterprise's existing Ethernet Layer 2 or IP Layer 3 VPN. A distributed, policy-based approach separates the evolution of compute and networking technologies. This separation allows multiple virtualization platforms to interoperate over a single network. For large scale and high traffic volume environments, the Nuage Networks 7850 VSG provides gateway functionality with native support for 1GE, 10GE and 40GE connections.
- For NFV, the CloudBand Management System orchestrates, automates, and optimizes virtual network functions across the service provider's distributed network and data centers. The CloudBand Node is a turn-key, all-in-one compute and storage node system. It includes hardware and software designed for efficient remote operation of distributed clouds.
- For next generation OSS/BSS solutions, the Motive Dynamic Operations provides Service & unified resource engine (SURE). SURE allows service providers to make their operation systems as agile as their virtualized network and data center, providing a unified view of the network and cloud infrastructure.

What are the proof points?

Alcatel-Lucent states that it has the capacity to scale its business quickly, and that it has an end-to-end offering from an SDN Controller, NFV orchestration platform with routing, optics, and virtualized appliances such as the LTE Packet Core, and VoLTE IMS solutions. The company claims that it can leverage its size and its expertise in broadband and wireless access solutions.

Alcatel-Lucent is highly active in promoting the notion of SDN and NFV, including:

- An ecosystem of partners as integral parts of its architecture, with 50 members including 6Wind, HP, F5 Networks, Intel, RedHat, VMWare, Contextream and others.
- Deployments in various countries (including Verizon, AT&T, NTT, DT and Telefonica) and more than 30 trials.

Alcatel-Lucent says that it is differentiated by taking a holistic end-to-end approach by combining its SDN and NFV solutions with operational support systems, a broad range of network solutions, strategic partnerships and professional transformation services.

Meru Networks

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

Meru Networks stated that it is taking a leadership role in developing and deploying best-of-breed wireless LAN solutions that are SDN enabled. According to Meru, its solutions can integrate with any wired vendor that also supports the OpenFlow solution. This capability allows customers to manage and control their wired/wireless network as a single unified network.

What is your value add?

The Meru SDN solutions provide:

- ONF certified OpenFlow wireless network solutions;
- End-to-end application QoS enabling enforceable service-level agreements (SLAs);
- Single-pane-of glass management of the unified wired and wireless network, with policy automation;
- Support for multi-vendor solutions through the ability to mix-and-match best-of-breed solutions.

What are the proof points?

- Meru Networks is the first WLAN vendor to receive the Certificate of Conformance through the ONF OpenFlow™ Conformance Testing Program (June, 2014);
- Meru Networks is the first vendor to complete qualification with Microsoft Lync® for 802.11ac wireless networking solutions (Aug 2014);
- Meru Networks is the winner of the October [SearchNetworking Network Innovation Award](#) for its achievements in the wireless software-defined-networking (SDN) space. (November 2014).

6WIND

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

6WIND stated that it enables NFV by accelerating Linux based networking environments to provide over 10X network performance improvements compared to standard Linux software architectures. As a result, service providers benefit from bare metal performance in their virtual environments.

The two products that 6WIND delivers for NFV are:

Solution 1: Data Plane Acceleration using 6WINDGate packet processing software

Description: Data plane performance enhancements that enable OEMs to build accelerated applications in bare metal and virtual environments. By leveraging a fast path architecture outside of the Linux kernel, 6WINDGate is deployed transparently with no change to OpenStack, the OS, hypervisor or virtual switch. 6WINDGate delivers the following features:

- High performance Layer 2-4 packet processing software for generic servers with a choice of multicore processors including Broadcom, Cavium, Intel and EZchip/Tilera;
- Cryptographic acceleration (software and hardware acceleration for built-in or external crypto engines);
- Fast path-based data plane solution on Intel leveraging DPDK and extensions (multi-vendor 10G and 40G NICs, smart NICs and more);
- Accelerated IPsec and IKE stack supporting over 190 Gbps over tens of thousands of tunnels on Intel servers;
- Accelerated TCP/UDP stack supporting over 100 million concurrent sessions and session setup rates of 5 million sessions per second;
- High capacity firewall and NAT;
- Wide tunneling support: GTP, PPP, L2TP, GRE, MPLS, VXLAN, etc.

Virtual Network Functions (VNFs) that can be built with 6WINDGate packet processing software include: routers, firewalls, Carrier Grade NAT, IPsec Gateways, EPC, HTTP-based applications and more.

Solution 2: NFV Infrastructure using 6WIND's Virtual Accelerator

Description: The 6WIND Virtual Accelerator runs within the hypervisor domain with a hardware-independent architecture that allows new and existing VMs to be integrated quickly onto x86-based servers. As a transparent virtual infrastructure acceleration solution, 6WIND Virtual Accelerator is provided as a simple software package so that customers do not have to replace or modify existing software such as Open vSwitch (OVS), Linux, Hypervisors and OpenStack.

Features include:

- Network hardware independence for seamless hardware updates, including 10G to 40G to 100G ports;
- Wire speed performance required to enable high density, compute intensive VMs on a single server;
- Flexible virtual switching support for Open vSwitch and Linux Bridge with no modifications;
- Complete virtual networking infrastructure with VLAN, VXLAN, Virtual Routing, IP Forwarding, Filtering and NAT;

- Native Virtio support for VMs based on different OSs;
- High bandwidth for VM to VM communications required for Service Chaining;
- No modification to OpenStack for orchestration.

What is your value add?

Scalability: With 6WIND's data plane acceleration and NFV Infrastructure, performance scales linearly with the number of processor cores. This means that fewer processor cores can be used for networking tasks so that more cores can be saved for the actual VNFs. At Dell World 6WIND recently demonstrated 240 Gbps aggregate bandwidth running an IP Forwarding VM. Another demonstration showed over 200 Gbps throughput with 80% of the processing cores left to run [Virtual Network Functions](#).

Hardware Independence: As an independent software vendor, 6WIND supports multi-vendor NICs from vendors such as Intel, Mellanox and Emulex so that there is not vendor hardware lock-in. This network hardware independence enables seamless hardware upgrades including 10G to 40G to 100G ports.

Performance: NFV is cost effective if performance can be achieved in virtual environments at least similar to physical environments. 6WIND enables 200 Gbps of virtual switching, 190 Gbps of IPsec and over 100 million concurrent TCP connections, to give a couple examples of performance that is not sacrificed with virtualization.

What are the proof points?

High Performance Virtualized SBC with Metaswitch and 6WIND

In November 2014 Metaswitch Networks announced a test with Perimeta, its virtualized session border controller (SBC), with 6WIND Virtual Accelerator for NFV Infrastructure. The background for this test is that conventional cloud environments built for IT workloads are not designed to provide the high rates of packet throughput needed to support virtualized network functions such as session border controllers. This has the effect of limiting the number of concurrent media sessions that can be supported by a virtualized SBC on a given amount of hardware.

For example, Perimeta SBC software running directly on current generation Intel architecture servers can relay about 60,000 concurrent bi-directional audio sessions using 6 CPU cores. However, the same software, running in a virtual machine in a conventional cloud configuration, using Open vSwitch (without specialized tuning) and using the same number of CPU cores can only manage 700 sessions.

By substituting 6WIND Virtual Accelerator, the results are much better. Virtual Accelerator helps increase the media capacity of a virtualized Perimeta SBC, running on 6-CPU cores, to 36,000 sessions, a more than 50 times improvement over Open vSwitch. In this configuration, 2 CPU cores are dedicated to running 6WIND Virtual Accelerator while 4 CPU cores are forwarding media. At 9,000 sessions per CPU core, the 6WIND Virtual Accelerator solution delivers 90 percent of the capacity per core of Perimeta SBC running on bare metal.

Dell

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

Dell stated that it enables the Open Networking Ecosystem. As Dell points out, the world has changed, the cutting edge is no longer found in proprietary solutions and closed ecosystems. Yet while open infrastructure technologies have already delivered exponential cost efficiencies and paradigm shifting innovation in some of the world's largest data centers, these technologies remain difficult to access for service providers and effectively impossible to access for most enterprises. Dell stated that it stands alone in its bold embrace of open technologies and that Dell's focus is on making the latest innovations from open ecosystems available to all consumers without vendor lock-in. Dell believes that this approach fundamentally changes the economics and accessibility of open, web-scale technologies.

What is your value add?

Dell stated that its solutions leverage open technologies with open interfaces throughout all layers and that this enables them to deliver the simplicity of vertically integrated solutions with the openness, flexibility and economics of web-scale technologies. For carriers investing in NFV, Dell offers pre-engineered NFV infrastructure bundles with validated reference architectures for leading VNF offerings, management & orchestration solutions and full support for popular Linux and OpenStack distributions without proprietary hardware or software requirements. For enterprises investing in private cloud solutions, Dell supports Microsoft, VMware and OpenStack environments equally without forcing customers into a vertically integrated & closed solutions. Dell Networking offers integrations with Microsoft, VMware, Openstack, Cloudstack, Puppet, Chef and other ecosystem solutions without additional licensing fees.

All advanced software interfaces including Perl/Python/Puppet/Chef/Shell/REST/Openflow and other API's are included in the base license for Dell Data Center switches with full support and no additional fees for use with either Dell or 3rd party management software, controllers or applications. Dell's embrace of open hardware and software ecosystems allows the company to offer the broadest and most flexible array of technology solutions with industry-leading performance, economics and efficiency.

What are the proof points?

Dell Networking has added over 3,000 new customers and continues to outpace the market in growth. Dell stated that its extensive list of available case studies with marquis customers and cutting edge use cases provides a testimony of the efficacy and performance of their solutions. Dell Networking products have received numerous awards from leading technology publications, events and analysts and have received superior ratings in performance reviews from leading independent testing firms including Miercom, The Lippis Report and others.

EMC

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

EMC provides data center management software that delivers comprehensive monitoring, diagnostics and Service Assurance across software-defined networks, integrating natively between virtual network infrastructure and the physical hardware, providing detailed real-time topology mapping from physical port through the tenant. These capabilities are the foundation for EMC's NFV management strategy, extending the functionality across physical and virtual network boundaries providing enterprise and service providers the tools they need to effectively manage these emerging virtual network functions across heterogeneous network infrastructure.

What is your value add?

EMC Service Assurance Suite provides complete operational visibility across storage, compute and networks providing detailed fault correlation, root cause and impact analysis across large, complex data center infrastructure. Service Assurance increases availability by providing detailed performance analysis and event correlation to address issues before service is impacted, enabling network operations teams to determine the specific root cause of network issues, minimizing any potential downtime.

What are the proof points?

In April 2014, EMC Service Assurance Suite was selected by Enterprise Management Associates as the network management product offering the best scalability in their annual [Enterprise Network Availability and Monitoring System radar report](#).

As an additional proof point, Compucom saved over \$550K in the first 12 months of operation because of 80% faster root cause analysis of system and network problems, cutting their discovery time for their 15,000 node network from 2 weeks to 30 minutes with EMC Service Assurance Suite. They also saw a 4x reduction in trouble ticket reduction in the first year after deploying the [EMC solution](#).

Citrix

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

Citrix stated that it has always believed in software-based networking and offering its customers complete freedom of choice between platforms and features. Citrix claimed that its NetScaler VPX product line is the leader in the virtual ADC space and has the exact same binary as their hardware appliances. Citrix added that its virtualized products are the fastest growing products within their NetScaler product line, both in terms of revenue and net new customers and that they are positioning the NetScaler SDX as an open and elastic platform that consolidates network services into a unified service delivery layer accessible as a whole by an application through open APIs.

Citrix believes that the network needs an application control layer. SDN enables a programmable networking model that allows Citrix to disseminate deep and broad application intelligence into the network, making the network a unified Layer 2– Layer 7 intelligent application fabric. Based on that belief, Citrix stated its goal of seamlessly integrating its technology into SDN environments as an always-on, elastic service that can be consumed on demand. Towards that end, Citrix is building the NetScaler Control Center, which is a common multi-tenant platform that orchestrates NetScaler services across both physical and virtual appliances. The Control Center will allow customers to use all NetScaler appliances as an aggregate pool of capacity.

Citrix acknowledged that most vendors of L4 - L7 services have already made their solutions available as virtual appliances. They also stated their belief that advanced L4 - L7 services play a vital role in lending application intelligence to NFV environments through intelligent traffic steering between various virtualized services and enabling seamless availability, scalability, and performance of those services. These advanced L4 – L7 services can also provide the ability to integrate into the application orchestration environment as well as open APIs to drive configuration programmatically.

Citrix's view is that the core value of its products and technology should remain the same across both physical and virtual form factors. The choice of a physical appliance is primarily for performance and scalability reasons, which is generally addressed through a scale-out architecture in NFV environments which Citrix supports through its TriScale clustering technology.

What is your value add?

Citrix believes that NetScaler is very well positioned to play a critical role in the SDN value chain. Citrix stated that value for customers lies in networks having a deep understanding of applications and Citrix believes that's where NetScaler's application intelligence becomes an indispensable asset. Citrix extracts application information and they disseminate that information through the network using the programmable interfaces that SDN offers. The company believes that by tightly integrating with SDN environments they become a core part of the fabric and that they can interact with the switching layer to augment network intelligence with functionality such as application visibility, application-based QoS, advanced security and application-aware routing.

In NFV environments, NetScaler's value add goes beyond just large-scale load balancing of an operator's infrastructure. Citrix's stated that its orchestration capabilities, open APIs, and TriScale technology form the key enablers for the agility and scalability needed in environments like the Evolved Packet Core. NetScaler's native intelligence of various signaling protocols such as SIP and Diameter

allows for optimization of virtualized voice and AAA services in both mobile and fixed line operator networks. Purpose built functions such as CG-NAT and NAT64 enable a seamless transition to IPV6, while NetScaler's content and front-end optimization capabilities allow providers to offer a rich end-user experience for their mobile customers. Finally, NetScaler's layer 7 intelligence and traffic steering capabilities enables intelligent chaining of virtualized services that can be customized per subscriber

Citrix is an active participant in the OpenDaylight community and is working to shape the direction of SDN by working closely with Cisco and many other industry leaders to forge innovation in the areas of Group Policy (i.e., an advanced policy abstraction model to describe all networking), OpFlex (i.e., a declarative policy protocol that enables highly scalable solutions), and Network service header (NSH) for intelligent traffic steering and service chaining.

What are the proof points?

There are a number of NFV-related Proof of Concept (POC) trials being sponsored by organizations such as the European Telecommunications Standards Institute (ETSI). In addition, there are a number of additional private trails underway. Some of these trials are focused on enabling the operator to develop expertise necessary to conduct full life-cycle management of the virtualized applications that reside between the mobile packet gateway (PGW) and the Internet—a domain commonly referred to as either the Gi-LAN (3G) or the SGi-LAN (LTE). As the predominant application in the Gi-LAN and SGi-LAN, the Citrix ByteMobile Adaptive Traffic Manager (ATM) is part of these network virtualization trials.

Citrix is partnering with operators to develop a solution that: a) is readily integrated with an operator's chosen NFV management and operations (MANO) framework; and b) meets NFV requirements such as rapid service provisioning. The Citrix ByteMobile ATM function must scale in parallel with broadband data traffic growth and an NFV implementation will enable the automated scaling of this function within the S/Gi-LAN domain. To achieve this end, Citrix offers a complete virtualized application stack that includes the virtual Adaptive Traffic Manager and the Citrix NetScaler VPX virtual application delivery controller. In preparation for expected operator demand, Citrix has conducted lab demonstrations of this application stack using both XenServer/CloudPlatform and KVM/OpenStack as hypervisor /virtual infrastructure manager.

Nuage Networks

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

Nuage Networks Virtualized Services Platform (VSP) is a Software Defined Networking architecture that acts as a non-disruptive overlay for all existing virtualized and non-virtualized network resources. VSP works across Cloud Management Software packages (such as OpenStack, CloudStack, and VMware), across hypervisors, across vSwitch architectures (embedded into a hypervisor or standalone as part of an Open Source platform such as OpenFlow), across legacy network gear, across converged network hardware / software approaches, across network software platforms, and across datacenters and WANs.

In contrast to controller-based SDN approaches, Nuage Networks VSP is completely hardware-independent. Platform components run in virtual machines, in hypervisors, or in Docker servers. Further, in contrast to server virtualization-based SDN approaches, Nuage Networks VSP is completely independent from hypervisor versions.

A fundamental component of Nuage Networks VSP is network virtualization to eliminate the need for applications and Virtual Machines (VMs) to deal with complexities of the physical network layer, including both datacenter and branch office WANs via the Nuage Networks Virtualized Network Services (VNS) use case. Building on the virtualization layer, replacing hardcoded addresses and relationships with intelligent network policies enables applications and VMs to move fluidly and rapidly as needed. Lastly, NFV capabilities such as Service Chaining of virtual functions enables the automation of complex configuration tasks such as defining cascading firewalls for a multi-tier web application or eliminating network appliances at the branch location.

In summary, Nuage Networks VSP occupies a unique place in the SDN market. Nuage Networks' VSP is able to seamlessly operate as an overlay on the incumbent network vendors and technologies and also virtualization and programming constructs such as Docker. By seamlessly removing all constraints in the datacenter network through the WAN with a comprehensive network policy framework, Nuage Networks VSP makes the network completely programmable and instantly available to the business applications and end users.

What is your value add

What differentiates Nuage Networks:

- Non-disruptive network virtualization solution that's agnostic to the existing network equipment and topology, choice of compute management platform or hypervisor deployed.
- Unified Fabric – from the Datacenter to the WAN: With the Nuage Networks Virtualized Network Services (VNS) the solution efficiently virtualizes and unifies the entire network fabric across datacenters and out to the very edge of the WAN within a single policy domain.
- Intelligent Automation: By providing intelligent policies that are interpreted where needed at each network end point – such as within the hypervisor, the rack switch, or the WAN edge, the Nuage Networks solution provides unique efficiencies in CapEx and OpEx as well as provide a consistent set of network and security policies across the data center and WAN environments.

- **Service Provider Scale with full control:** Leveraging the operating system that today runs roughly a quarter of the Internet, the Nuage Networks solution provides the scale that F500 Enterprises, major Service Providers, and Governments demand.
- **Open Approach:** The Nuage Networks solution is unique in simultaneously supporting any Cloud Management System (e.g. OpenStack, CloudStack, and VMware), any Hypervisor (e.g. Docker, VMware, KVM and Xen), and any networking vendor (e.g. Alcatel-Lucent, Cisco, and Juniper).
- **Leading-edge Capabilities:** In addition to supporting new programming constructs such as Docker, the Nuage Networks solution enables the self-service capabilities needed to both enter new markets and to compete effectively with public cloud giants.

What are the proof points

In around 18 months of sales activity, Nuage Networks has been implemented at scale within the networks of leading global companies, including financial services and healthcare companies, cloud providers, telecommunications providers, and infrastructure providers across the three key worldwide theatres (NA, EMEA and APAC).

The best proof points for any solution are when customers leverage it to create entirely new ways to do business. Select examples include:

- **New Software Defined Cloud model:** Numergy, the national cloud provider for France, has created a new architecture that will allow all network and value-added services to be software defined, based largely on Nuage Network's capabilities. See <http://www.nuagenetworks.net/press-releases/numergy-selects-nuage-networks-software-defined-networking-solution-new-cloud-infrastructure/>
- **New Cloud Distributed Hosting model:** A service provider in EMEA (EVONET) has created a virtual Platform Optimized Design (vPOD) architecture that provides Cloud efficiencies along with the flexibility of offering either shared or dedicated resources distributed among datacenters. Nuage Networks VSP provides the interconnection within and among all vPODs and among all datacenters. See <http://vimeo.com/98173520>.
- **New Cloud Service model:** OVH, the number one hosting provider in Europe, has introduced a game-changing OpenStack-as-a-Service offering with Nuage Networks VSP. See <http://www.nuagenetworks.net/press-releases/ovh-sdn-nuage-networks/>
- **Unified Healthcare model:** UPMC, a \$10 billion integrated global health enterprise, has unified their 450 sites and 2 datacenters to provide services to 62,000 employees with Nuage Networks. See <http://www.nuagenetworks.net/press-releases/upmc-selects-nuage-networks/>

Pica8

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

Pica8 is one of the pioneers of open networking in part based on [PicOS™](#), the first Linux-based network operating system that enables customers to easily integrate conventional networking with software-defined networking (SDN) using commodity bare metal switches. Pica8 is challenging 20 years of the traditional proprietary approach to networking by providing an open switching system, which users can personalize to meet the needs of their application environment.

As traffic demands skyrocket and data centers must scale, traditional networking has proven increasingly inefficient. Current architectures, comprised of disparate network devices, are difficult to provision and to customize. Pica8 believes that SDN resolves these challenges because it abstracts the network data plane from the control plane, enabling users to easily address all network services through myriad external programming interfaces. The result is that a company's data center network can now be customized and personalized.

PicOS provides extensive support for traditional switching and routing protocols and Linux. PicOS offers a comprehensive and flexible configuration management environment from either a Linux shell or feature rich command line interface (CLI). PicOS runs as an application on an un-modified Linux kernel. This enables familiar system administration and automation tools such as CFEngine, Chef or Puppet to run natively and seamlessly.

PicOS delivers SDN solutions through Pica8's adoption of Open-vSwitch (OVS). Pica8 can provide PicOS, switching hardware through a growing list of [Hardware Ecosystem Partners](#), or both in a fully integrated package leveraging [ONIE](#) as part of an end-to-end data center SDN solution. Pica8 also provides a [Starter Kit](#) that Pica8 stated enables deployment in hours rather than months.

PicOS interoperates with leading network virtualization products such as VMware's NSX (Nicira), Midokura and PLUMgrid.

What is your value add?

As customers progress down the path of SDN, one thing that is clear in the early rollouts is that there are no cookie-cutter deployments. Everyone has a different application, a different use case, and a different rate of adoption. Pica8 believes that the network should be flexible and easy enough to help any customer on their journey.

To that end, Pica8 offers [CrossFlow Networking](#), where customers get the flexibility and granularity of OpenFlow-based network policy on the same device that's running traditional Layer-2 / Layer-3 protocols for efficient packet forwarding. CrossFlow Networking users can use OpenFlow to fine tune the switching or router tables in a switch. This allows users to inject OpenFlow rules for specific applications and policies, while also preserving the networking topology that's built on tried and true Layer-2 and Layer-3 protocols.

Benefits of CrossFlow Networking:

- Save money by reducing CapEx. With CrossFlow, users don't have to purchase a different set of networking equipment for their OpenFlow implementation.

- Save time by reducing the complexity of different modes of operation between Layer-2, Layer-3 and OpenFlow.
- Deliver business logic and policy into the network by simplifying the integration of SDN into the existing network.

More and more customers are deploying SDN in new use cases and in real networks. According to Pica8, these customers are discovering that some SDN implementations have significant limitations relative to the scalability of the solution. Pica8 enables large-scale OpenFlow based networks through software innovation with PicOS™, their operating system for open switching.

Pica8 accomplishes this in part by combining the switching and routing memory in the switch. Another technique that Pica8 has implemented that enables scalability is optimizing the memory allocation across these tables to drastically increase the number of flow entries that their switches can support. Pica8 claims that they can support over 200,000 flows — enough for even the largest of networks.

Pica8 enables seamless SDN integration with CrossFlow Networking, and supports networks of any scale with PicOS. If you want to deploy OpenFlow at scale in your production network, take a look at Pica8.

What are the proof points?

Pica8 has over 350 customers worldwide, including web services companies, global carriers, and leading research labs. Their corporate headquarters is located in Palo Alto, California, and their R&D facility is in Beijing, China, with sales and support offices worldwide.


Conclusions

The following is a summary of the conclusions that were drawn throughout the e-book

- Over the last year, the familiarity with SDN has increased significantly.
- The use of SDN in production networks should increase somewhat significantly in the next year.
- IT organizations are highly skeptical that they can implement network virtualization in the data center without using at least some dedicated hardware.
- Very few IT organizations have ruled out the use of OpenFlow.
- By a small margin, IT organizations perceive the fabric-based SDN model will provide more value over the next two years than will the overlay model. However, many IT organizations are yet to form an opinion.
- 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.
- Relatively few IT organizations believe that SDN will help them reduce CAPEX or reduce complexity.
- Two of the major inhibitors to SDN adoption are concerns about how to integrate SDN into the rest of the infrastructure and the lack of a compelling business case.
- Over the next two years, the primary focus of SDN deployment is likely to be in the data center. However, there is considerable interest in deploying SDN in the WAN as well as in branch and campus networks.
- Network organizations are very optimistic that over the next three years that there will be a significant increase in SDN deployment.
- Network organizations believe that three years from now that SDN deployment in the WAN and branch and campus networks will be almost as common as SDN deployment in data centers.
- SDN creates security opportunities and security challenges.
- The vast majority of IT organizations don't have a well thought out strategy for how they will implement orchestration.
- SDN creates management opportunities and security challenges.
- SDN holds the potential to enable IT organizations to dynamically change the environment in order to meet SLAs.
- 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.
- Over the next two years the ongoing adoption of software-based IT functionality is likely to have an impact on the structure of IT organizations.
- Over the next two years the ongoing adoption of software-based IT functionality is likely to have an impact on the jobs of IT professionals.

- The general awareness of NFV is low in general and it is lower than the general awareness of SDN.
- The vast majority of IT organizations believe that SDN and NFV are complimentary activities.
- A significant percentage of IT organizations believe that in at least some instances NFV requires SDN.
- Only a small percentage of IT organizations believe that SDN and NFV are totally independent activities.
- Only a very small percentage of IT professionals think that NFV is only applicable in a service provider environment.
- Almost half of IT professionals think that NFV is equally applicable in a service provider environment and an enterprise environment.
- 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 three biggest inhibitors to the broad adoption of NFV are:
 - Concerns about end-to-end provisioning;
 - The lack of a compelling business case;
 - The immaturity of the current products.
- Within a few years, the majority of IT organizations are likely to have made a significant deployment of NFV.
- While IT organizations have interest in a number of the ETSI-defined use cases, by a wide margin they are most interested in the Network Functions Virtualization Infrastructure as a Service use case.
- The majority of IT organizations believe that even if a NFV-related POC is successful, it will take between a significant and a tremendous amount of effort to broadly implement that solution 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 nearly half of all it professionals.

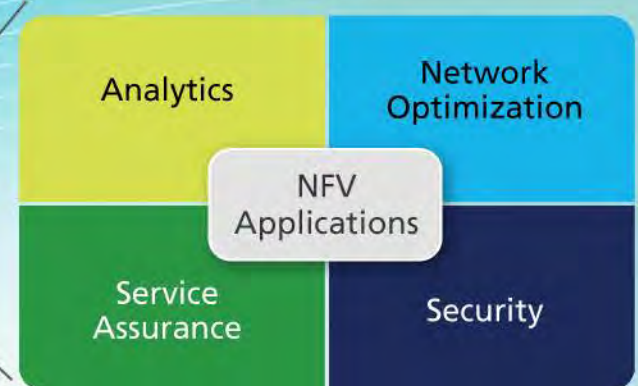
BTI's SDN & NFV Solutions



BENEFITS:

- Open, standards-based
- End-to-end visibility & control
- Improved efficiencies & utilization
- Increased scale & performance
- Rapid service innovation
- Reduced opex

Component	BTI Applications Blade Options
CPU	x.86 options based on application specification



*Note: NFV Applications scan operate on the BTI 7800 Series [x.86/Linux] Applications Blade



BTI Systems, Inc.

Corporate Headquarters
1000 Innovation Drive, Suite 200
Ottawa, Ontario K2K 3E7 Canada

US Headquarters
One Monarch Drive, Suite 105
Littleton, MA 01460 USA

btisystems.com

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.

Copyright © 2015 Webtorials

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.