

# The 2015 Guide to SDN and NFV

## Part 2: Network Functions Virtualization (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

<b>EXECUTIVE SUMMARY .....</b>	<b>1</b>
<b>INTRODUCTION .....</b>	<b>2</b>
BACKGROUND .....	2
ETSI .....	2
TM FORUM .....	4
INTERNET ENGINEERING TASK FORCE (IETF) .....	5
OPEN PLATFORM FOR NFV (OPNFV) .....	5
RELATIONSHIP BETWEEN SDN AND NFV .....	6
STATUS OF NFV ADOPTION .....	8
<b>USE CASES AND PROOF OF CONCEPT .....</b>	<b>17</b>
<b>ETSI NFV USE CASES .....</b>	<b>17</b>
<i>NFV Infrastructure as a Service (NFVlaaS)</i> .....	17
<i>Virtual Network Functions as a Service (VNFaaS)</i> .....	17
<i>Virtualization of the Home Environment</i> .....	17
<i>VNF Forwarding Graph (FG)</i> .....	18
<i>Virtual Network Platform as a Service (VNPaaS)</i> .....	18
<i>Virtualization of Mobile Core Network and IP Multimedia Subsystem</i> .....	18
<i>Virtualization of the Mobile Base Station</i> .....	19
<i>Virtualization of Content Delivery Networks (CDNs)</i> .....	19
<i>Virtualization of Fixed Access Network Functions</i> .....	19
<b>TM FORUM CATALYST POCs .....</b>	<b>20</b>
<i>Closing the Loop: Data-driven network performance optimization for NFV &amp; SON</i> .....	20
<i>CloudNFVTM: Dynamic, data-driven management and operations Catalyst</i> .....	21
<i>Orchestrating Software-Defined Networking (SDN) and NFV while Enforcing Service Level Agreements (SLAs) over Wide Area Networks (WANs)</i> .....	21
<i>Service bundling in a B2B2X marketplace</i> .....	21
<b>PRIVATE POCs .....</b>	<b>22</b>
<i>Virtualized S/Gi-LAN</i> .....	22
<b>THE OPERATIONAL IMPLICATIONS .....</b>	<b>32</b>
<b>PERFORMANCE LIMITATIONS .....</b>	<b>32</b>
<b>END-TO-END MANAGEMENT .....</b>	<b>34</b>
<i>Management Challenges</i> .....	34
<i>Management Direction</i> .....	36
<b>THE ORGANIZATIONAL IMPLICATIONS .....</b>	<b>37</b>
<i>Impact on Organizations and Jobs</i> .....	37
<i>DevOps</i> .....	39

# Executive Summary

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

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

This document has three chapters:

## 1. Introduction

Provides background on topics such as the role of both the European Telecommunications Standards Institute (ETSI) and the TM Forum in the development of NFV; the current status of NFV deployment; and the perceived relationship between SDN and NFV. This section also analyzes the primary factors that are driving and inhibiting adoption.

## 2. Use Cases

Describes the nine NFV use cases identified by ETSI and describes some of the key NFV Proof of Concept trials (POCs) that are currently underway.

## 3. Operational Considerations

Analyzes the performance, management and organizational issues relative to implementing NFV and identifies how impactful these issues are likely to be in the near term. This section also discusses the value that DevOps brings to IT and identifies how a DevOps methodology would have to be enhanced to be applicable to network operations.

This section 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.***

# Introduction

## 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:

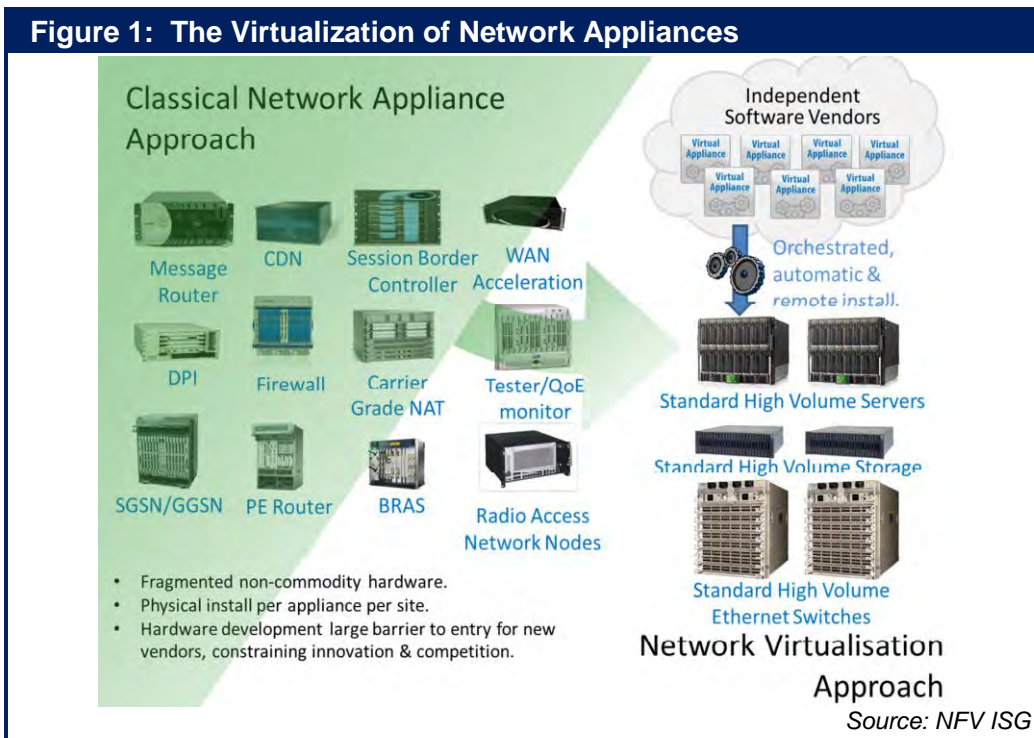
- Industry standard servers running Linux or Windows;
- Virtual appliances running over hypervisors on industry standard hardware servers;
- COTS-compute blade integrated in network elements;
- Proprietary hardware appliances.

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 1**.



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 1** contains examples of functions that could be virtualized.

**Table 1: Potential Functions to be Virtualized**

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

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 uses 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 2**.

<b>Table 2: Perceived Relationship between SDN and NFV</b>	
<b>Relationship</b>	<b>% of Respondents</b>
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 2** 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 3**.

<b>Applicability</b>	<b>% of Respondents</b>
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 4**.

<b>Approach to Implementing NFV</b>	<b>% of Respondents</b>
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 4** 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 5**.

<b>Table 5: Factors Driving NFV</b>	
<b>Factor</b>	<b>% of Respondents</b>
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 5** 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 6**.

<b>Table 6: Factors Inhibiting NFV</b>	
<b>Inhibitor</b>	<b>% of Respondents</b>
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 6** 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 7**.

<b>Table 7: Time Frame for Deployment</b>	
<b>Time Frame</b>	<b>% of Respondents</b>
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 7** plus the data in **Table 4** 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.***

~ Continued on page 17 ~



# **AUTOMATE YOUR CLOUD WITH** **aCLOUD SERVICES ARCHITECTURE**

Integrate dynamic  
services into your  
Cloud Data Center

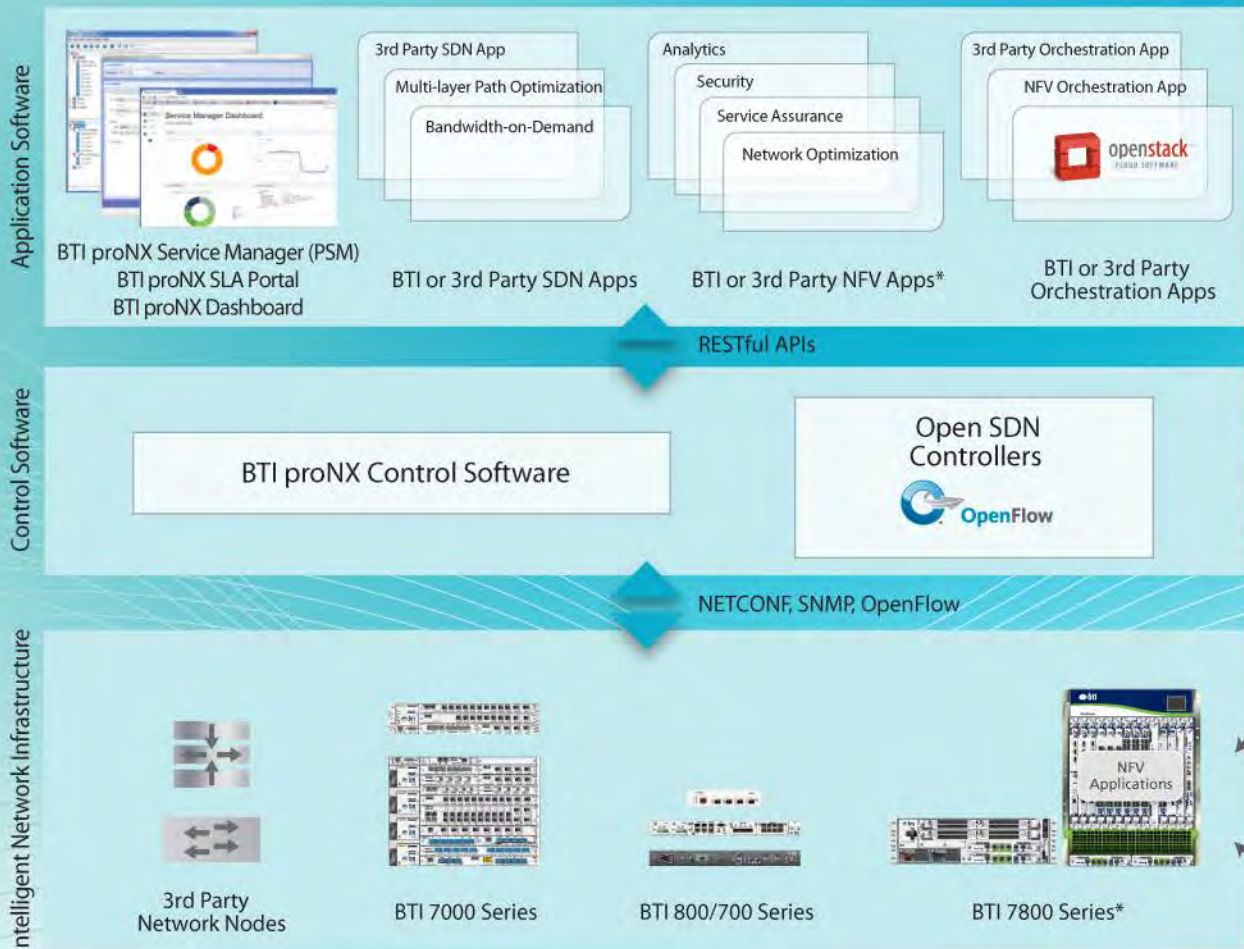
---

[www.a10networks.com](http://www.a10networks.com)






# 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](http://btisystems.com)

# 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:

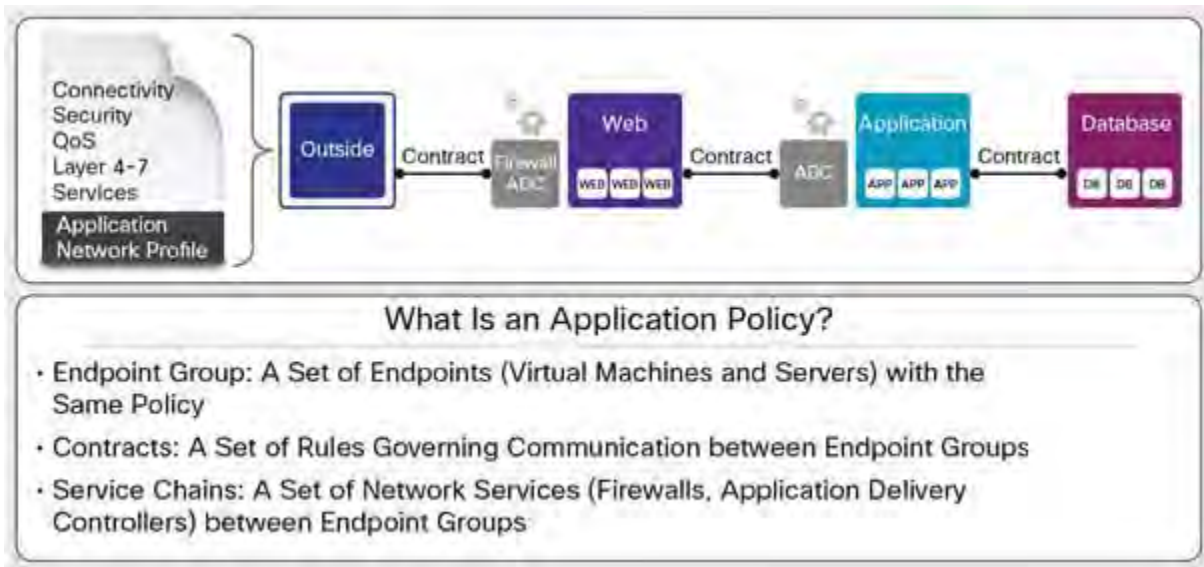
- A centralized policy store and infrastructure controller: In SDN and Cisco ACI, this feature is generally known as the controller (Cisco [Application Policy Infrastructure Controller \[APIC\]](#) for Cisco ACI).
- Programmable, or automated, network devices: All infrastructure devices, such as switches, 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](http://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](http://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)

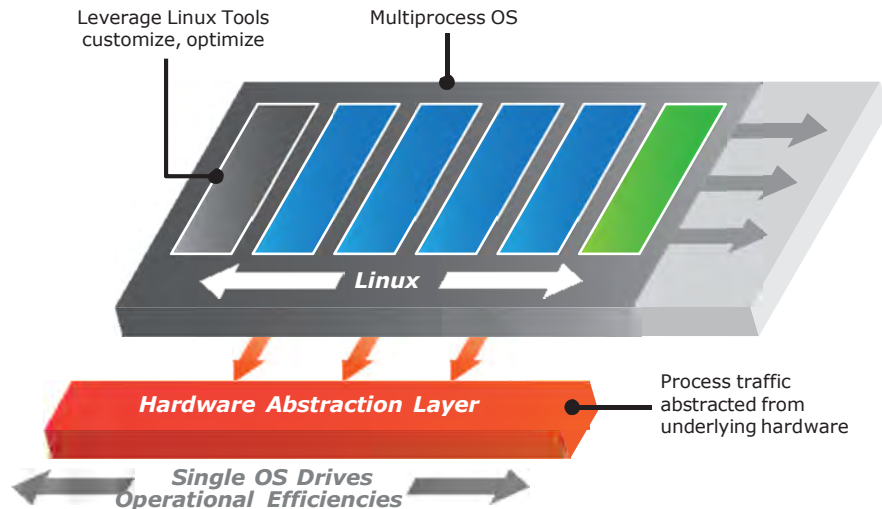


## PicOS Overview

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

## 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 3<sup>rd</sup> 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 8**.

<b>Table 8: Interest in ETSI Use Cases</b>	
<b>Use Case</b>	<b>% of Respondents</b>
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 8** 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 3<sup>rd</sup> 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.

*~ Continued on page 32 ~*

# 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](http://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.

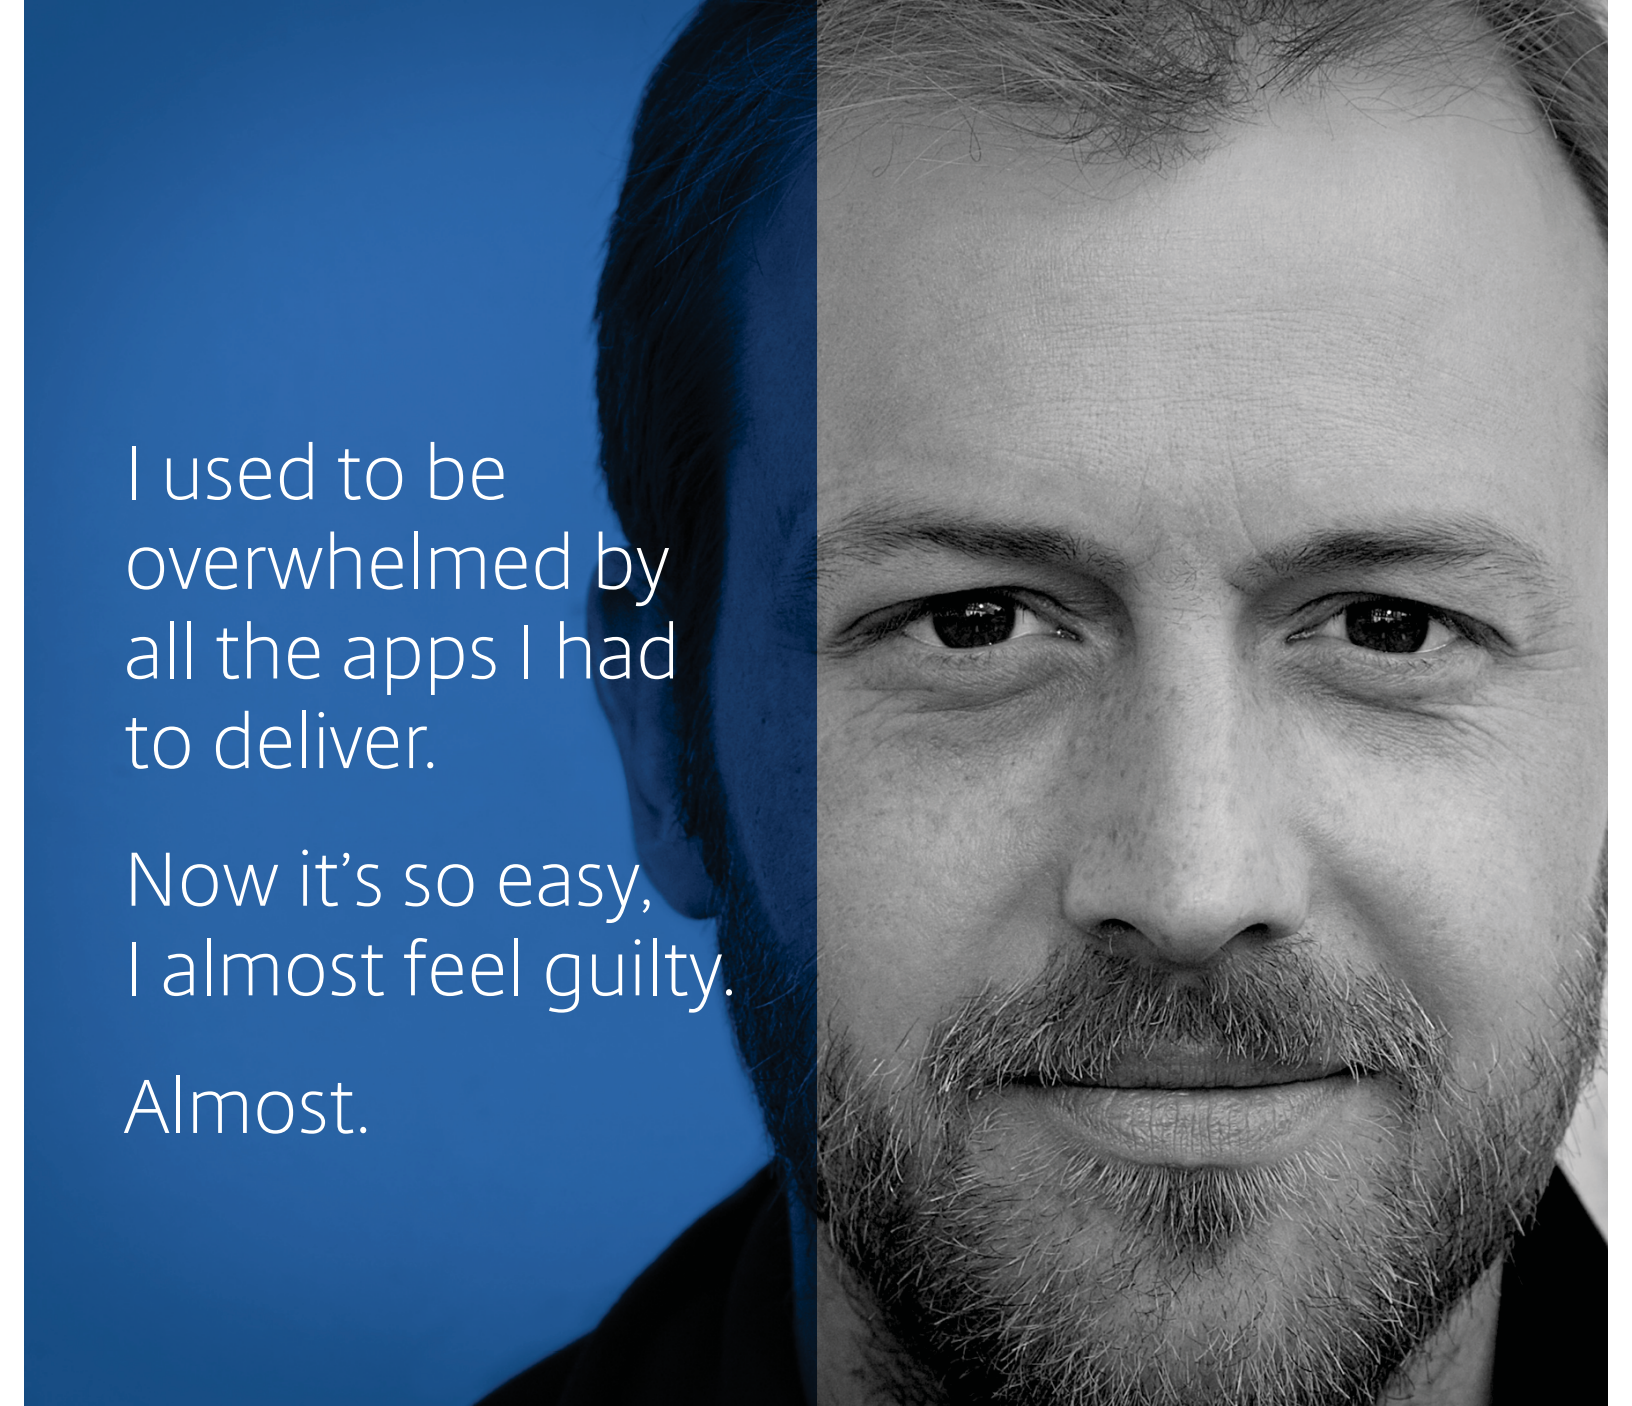
## Nuage Networks SDN

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.





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.



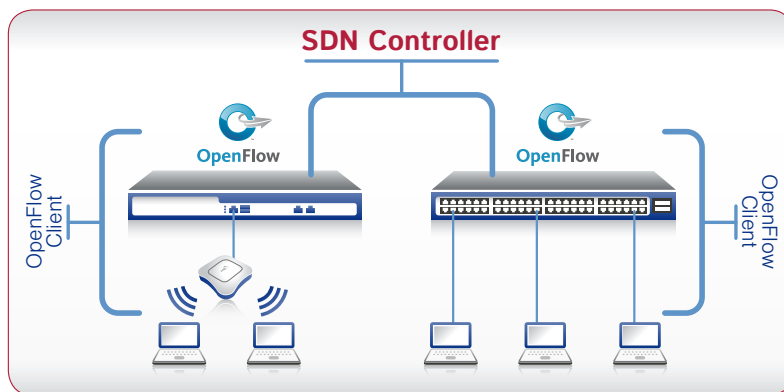
**NetScaler with TriScale**  
SOFTWARE SMART. HARDWARE STRONG.

**CITRIX®**

[www.citrix.com/netscaler](http://www.citrix.com/netscaler)



## 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](http://www.merunetworks.com) or email your questions to: [meruinfo@merunetworks.com](mailto: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](mailto:meruinfo@merunetworks.com)



# 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"> <li>Define network service requirements in clear, IT-friendly language</li> <li>Bring services up using automated, policy-based instantiation of network connectivity</li> <li>Dramatically reduce time to service and limit potential for errors</li> </ul>	<ul style="list-style-type: none"> <li>Adapt datacenters and private networks dynamically</li> <li>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</li> </ul>	<ul style="list-style-type: none"> <li>Benefit from distributed, policy-based approach that allows multiple virtualization platforms to interoperate over a single network</li> <li>Optimize the datacenter network and private network by separating service definition from service instantiation</li> </ul>

## 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"

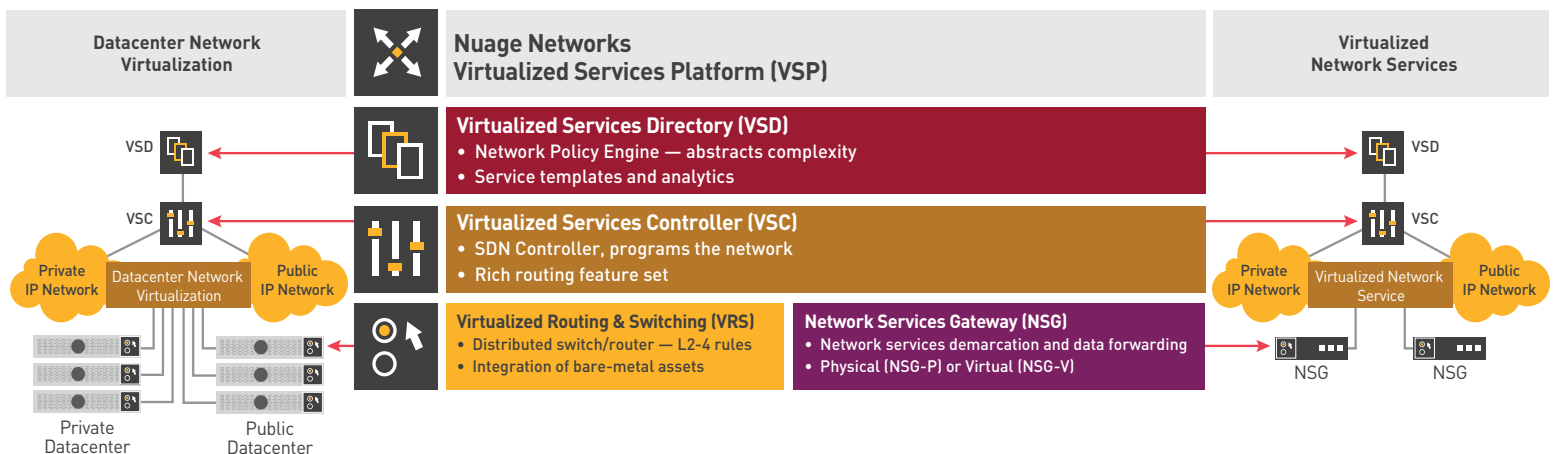
The cloud can be more than what it is. In fact, it needs to be. When we founded Nuage Networks, it was with the idea that it's time for the cloud to come of age. From the beginning we recognized the unique challenges that cloud service providers and large enterprises face delivering and managing large, 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)**

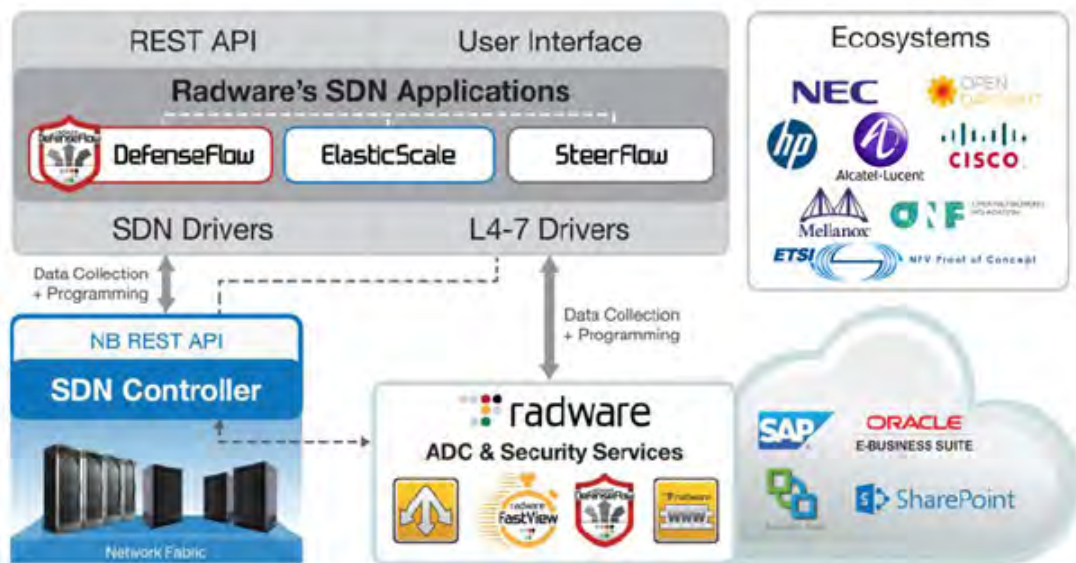
### SDN & NFV for a Scalable Application Delivery Network

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](mailto:info@radware.com) or go to [www.radware.com](http://www.radware.com).

## 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 9**.

<b>Table 9: Effort to go from POC to Production</b>	
<b>Amount of Effort</b>	<b>% of Respondents</b>
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%

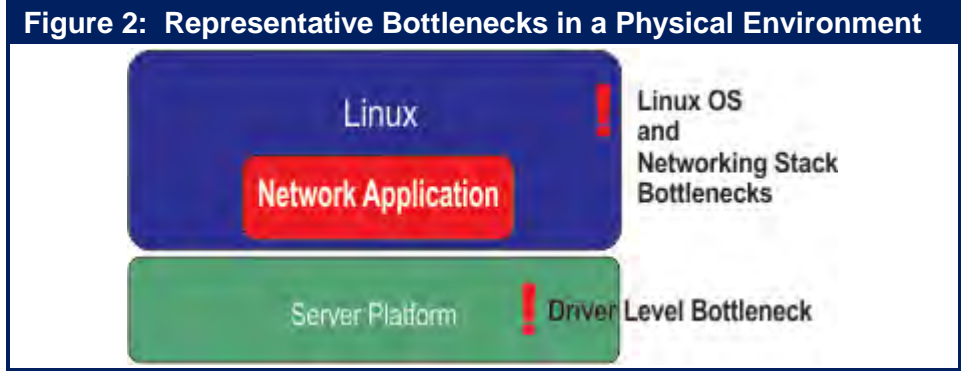
The data in **Table 9** indicates:

***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.***

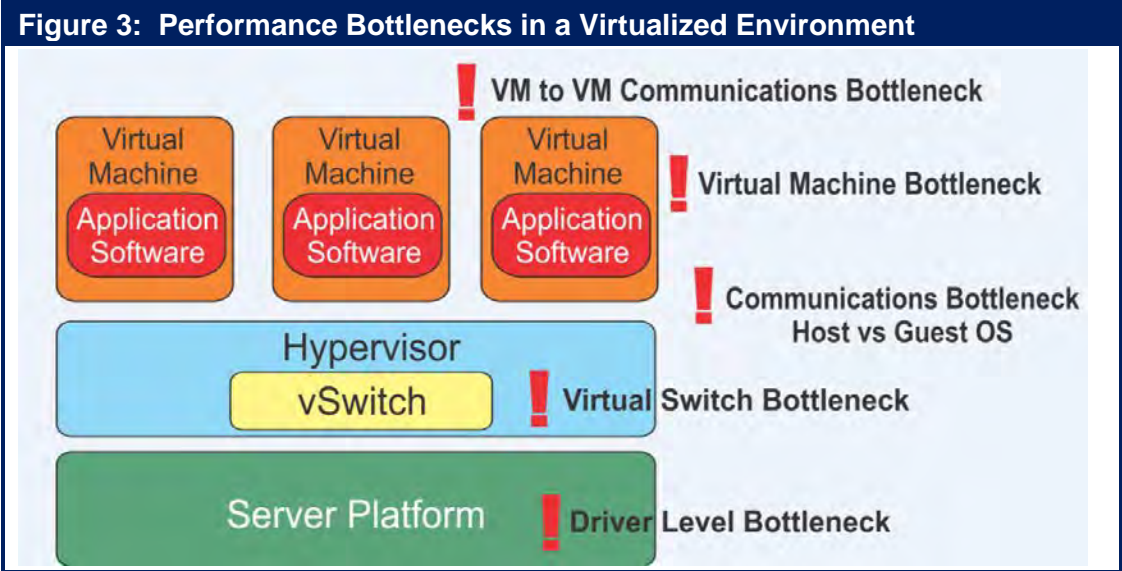
### Performance Limitations

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 2**.

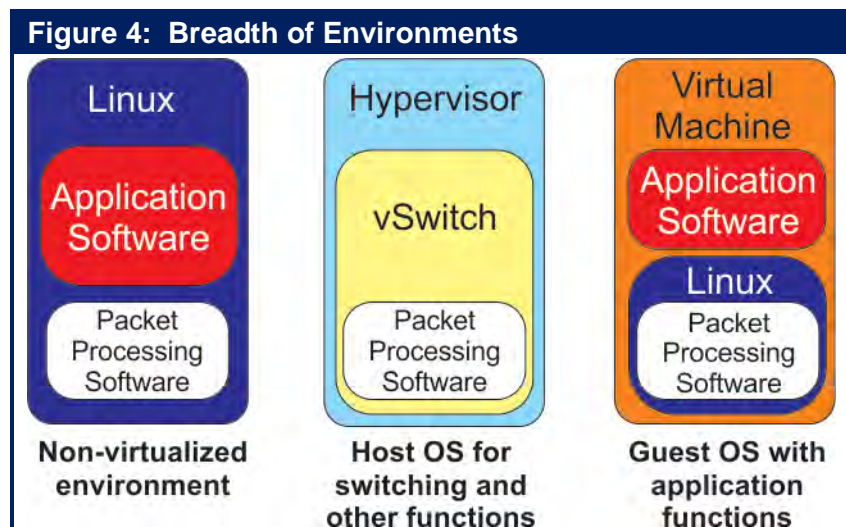


Unfortunately, as shown in **Figure 3**, as IT organizations adopt a virtualized environment, the performance bottlenecks multiply.



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 4**, 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 be 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 10**.

<b>Table 10: Impact of NFV on Organizational Structure</b>	
<b>Impact</b>	<b>Percentage of Responses</b>
Very Significant Impact	6%
Significant Impact	28%
Moderate Impact	24%
Some Impact	19%
No Impact	12%
Don’t Know	9%

The data in **Table 10** 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 11**.

<b>Table 11: Impact of NFV on Employee Skills</b>	
<b>Impact</b>	<b>% of Responses</b>
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 11** 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.

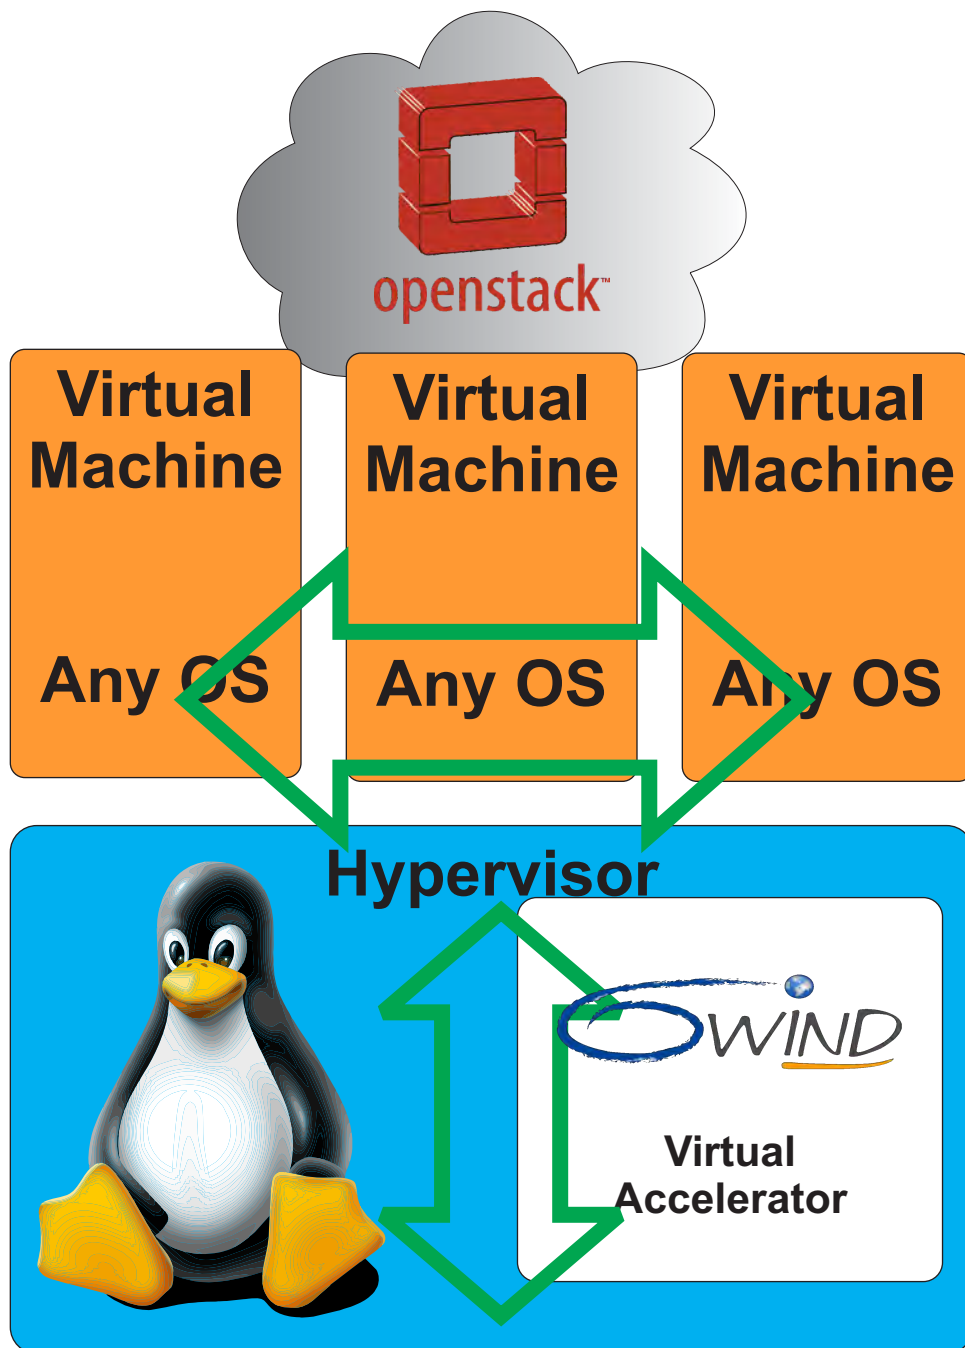
All of the basic principles of DevOps are applicable in a network operations (NetOps) setting. However DevOps is generally applied to discreet services that are frequently delivered over the web on a best effort basis. The network environment is different than that and as a result virtualized network 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.

# 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](http://www.6WIND.com)



## 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)

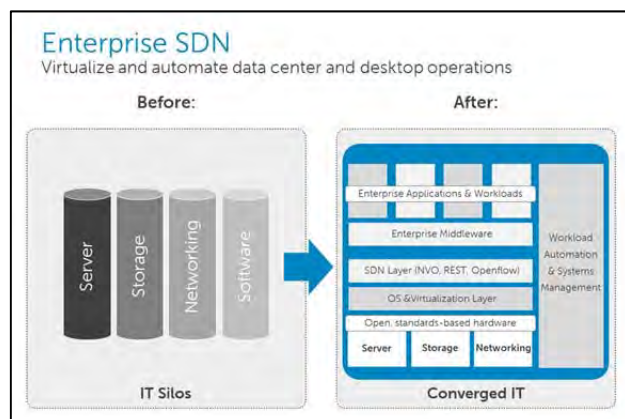
Other IT teams prefer to use programmatic or scripting languages, such as Puppet or Chef, to interface with their infrastructure and automate operations. Rather than a controller that speaks to multiple devices, they implement a programmatic language to define and implement policies across the infrastructure.

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

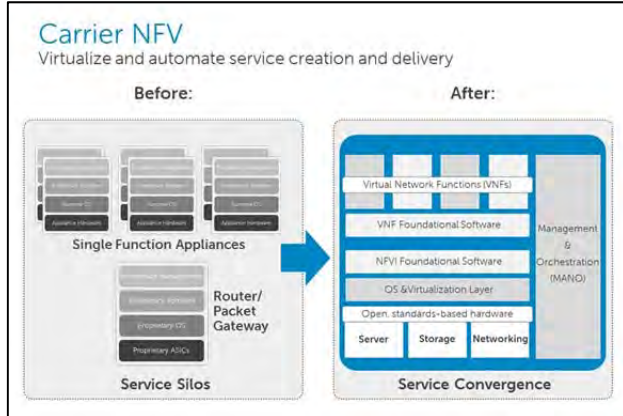
Traditionally, when a carrier delivers IP services, data packets are sent from customer site or device to a carrier's router or switch, and then daisy-chained through a set of boxes performing additional service-related functions.

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.

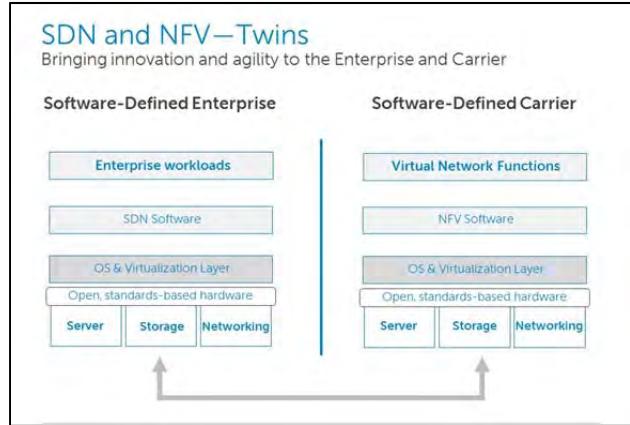
There are also fundamental differences in the vocabulary surrounding each. Instead of *workloads*, carriers are concerned with *services*, and instead of *business continuity*, carriers are interested in *carrier-grade 5-9's and 6-9's technology*.



**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](http://Dellnetworking.com)**

# Extending Service Performance Management into SDN and NFV Environments

## Solution Benefits

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

## Problem Overview

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

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

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

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

## Solution Overview

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

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

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




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

### Core Technologies

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

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

#### Utilize Packet Flow Data

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

#### Provide Scalable Packet Flow Access

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

#### Adaptive Session Intelligence (ASI)

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



## Service Delivery Monitoring in SDN Environments

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

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

## Solution Benefits

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

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

## 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](http://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](http://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.



# 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](http://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](mailto:asd@emc.com) or call 866-438-3622.

**EMC<sup>2</sup>**

## 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](http://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 © 2014 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.