The 2016 Guide to SDN and NFV

Part 5: Network Functions Virtualization (NFV): Architectural Considerations and Use Cases

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

Introduction

Over the last couple of years the hottest topics in networking have been Software Defined Networking (SDN) and Network Functions Virtualization (NFV). While both enterprises and service providers have shown great interest in these topics, the vast majority of organizations are still either on the sidelines or in the analysis stage of adoption. The primary goals of **The 2016 Guide to SDN & NFV** (The Guide) are to eliminate the confusion that surrounds SDN and NFV and to accelerate the analysis and potential adoption of these new architectural approaches.

The Guide will be published both in its entirety and in a serial fashion. This document is the fifth of the serial publications. It will focus on some of the key NFV architectural considerations as well as some of the key NFV use cases. Below is a listing of all of the publications that comprise The Guide:

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

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

The Guide contains the results of a survey that was distributed in December 2015 and January 2016. Throughout The Guide the 144 network professionals who completed the survey will be referred to as The Survey Respondents.

Architectural Considerations

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

- Longevity;
- Scalability;
- Interoperability;
- Performance;
- Profitability/Cost[1];
- Risk;
- Availability.

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

Big Bang vs. Piecemeal Approach

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

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

Software Modularity

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

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

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

Technology Considerations

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

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

Key Virtualization Concepts

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

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

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

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

Software-Centric Design

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

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

The Role of Open Source

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

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

Relationship with SDN

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

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

Software Defined Data Center

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

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

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

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

A Fresh Approach to High Availability

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

The implementation of NFV enables organizations to rethink HA.

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

The (potential) end of Moore's Law

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

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

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

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

Status of NFV-Related Architectures

The Survey Respondents were asked to indicate the progress their organization has made relative to developing an effective architecture for the broad adoption of NFV. Their responses are shown in **Table 1.**

Table 1: Progress Towards a NFV Architecture	
Amount of Progress	**% of Respondents**
None	26%
A little, with a lot of work ahead of us	41%
A lot, but still some work ahead of us	16%
Already developed an architecture	17%

Table 1 indicates that:

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

The Survey Respondents were also asked to indicate how much time their organization will spend over the next year developing a NFV-related architecture. Their responses are shown in **Table 2**.

Table 2: Amount of Time to be Spent on a NFV Architecture	
Amount of Time	**% of Respondents**
None	16%
A modest amount	49%
A significant amount	30%
None, we already have an architecture	6%

The combination of **Table 1** and **Table 2** indicates that:

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

Use Cases and Proof of Concept Trials

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

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

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

ETSI NFV Use Cases

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

NFV Infrastructure as a Service (NFVIaaS)

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

Virtual Network Functions as a Service (VNFaaS)

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

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

Virtualization of the Home Environment

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

VNF Forwarding Graph (FG)

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

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

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

Virtual Network Platform as a Service (VNPaaS)

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

Virtualization of Mobile Core Network and IP Multimedia Subsystem

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

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

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

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

Virtualization of the Mobile Base Station

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

Virtualization of Content Delivery Networks (CDNs)

Some ISPs are deploying proprietary CDN cache nodes in their networks to improve delivery of video and other high bandwidth services to their customers. Cache nodes typically run on dedicated appliances running on custom or industry standard server platforms. Both CDN cache nodes and CDN

control nodes can potentially be virtualized. The benefits of CDN virtualization are similar to those gained in other NFV use cases, such as VNFaaS.

Virtualization of Fixed Access Network Functions

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

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

Jim Metzler has a broad background in the IT industry. This includes being a software engineer, an engineering manager for high-speed data services for a major network service provider, a product manager for network hardware, a network manager at two Fortune 500 companies, and the principal of a consulting organization. In addition, he has created software tools for designing customer networks for a major network service provider and directed and performed market research at a major industry analyst firm. Jim's current interests include cloud networking and application delivery.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler or Steven Taylor.

Published by Webtorials Editorial/Analyst Division

www.Webtorials.com

Division Cofounders:

Jim Metzler

Steven Taylor

Professional Opinions Disclaimer

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

AUTOMATE YOUR CLOUD WITH aCLOUD SERVICES ARCHITECTURE

Integrate dynamic services into your Cloud Data Center

www.a10networks.com

Why Choose Application Centric Infrastructure (ACI)?

Application Deployment at the Speed of Business

Without ACI

New Application

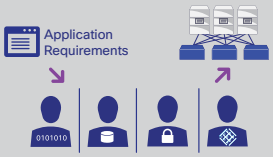

IT administrators work in silos.

Weeks

There is no shared architectural model.

With ACI

New Application

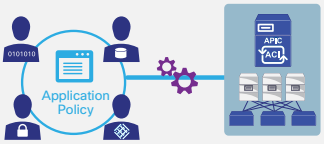

Simpler

Faster

IT admins work at the application level.

Hours

Efficient

Shared model for policy automation.

ACI cuts deployment time and effort.

Optimal Design

Application Policy

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

Faster Deployment

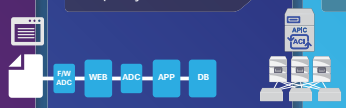

Application Network Profile

- Test
- Certify
- Adapt
- Provision

Simplified Operations

Application Topology

- Monitor
- Troubleshoot
- Optimize
- Scale Out

What does ACI deliver?

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

cisco

Redefine the Power of IT with ACI

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

Masergy's Software Defined Platform Delivers the Flexibility Enterprises Require

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

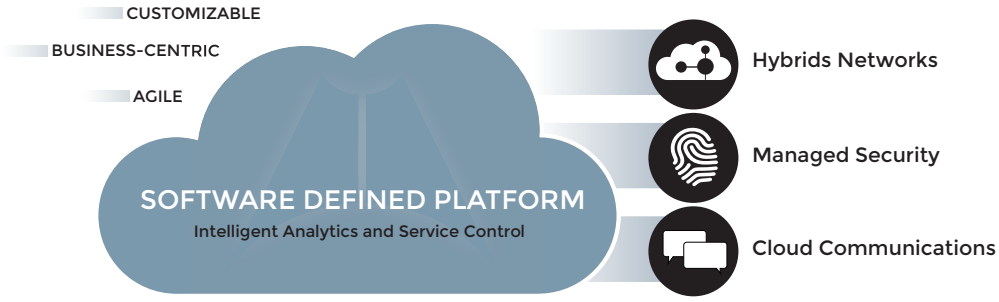

This WAN innovation is being enabled by a Software Defined Networking (SDN) architecture, which enables administrators to rapidly change configurations as performance requirements demand.

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

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

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

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

VIRTUALIZATION

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

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

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

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

"Masergy was able to custom design our hybrid network to meet our unique application performance requirements. It's outcome-based approach and ongoing superior support have convinced us we selected the right partner for our needs."

Patrick Tisdale, CIO – McKenna, Long & Aldridge, LLP

FLEXIBLE BY DESIGN

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

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

MASERGY
Performance Beyond Expecations

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

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

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

Contact us for a free consultation.

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

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

NETSCOUT

Extending Service Assurance into SDN and NFV Environments

SOLUTION BENEFITS

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

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

Solution Core Functionality

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

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

Challenges

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

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

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

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

Solution Overview

NETSCOUT® offers rapid service triage based on pervasive end-to-end visibility across physical, virtual, and hybrid service delivery environments. The triage is performed proactively by detecting service degradations in real time using one cohesive, consistent set of metadata, based on packet flow data, for service provider and enterprise services. This metadata is generated by the patented Adaptive Service Intelligence technology running on NETSCOUT's physical and virtual Intelligent Data Sources, and offers meaningful and contextual view of all interrelationships and dependencies across all service delivery components in physical, virtual, and hybrid environments.

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

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

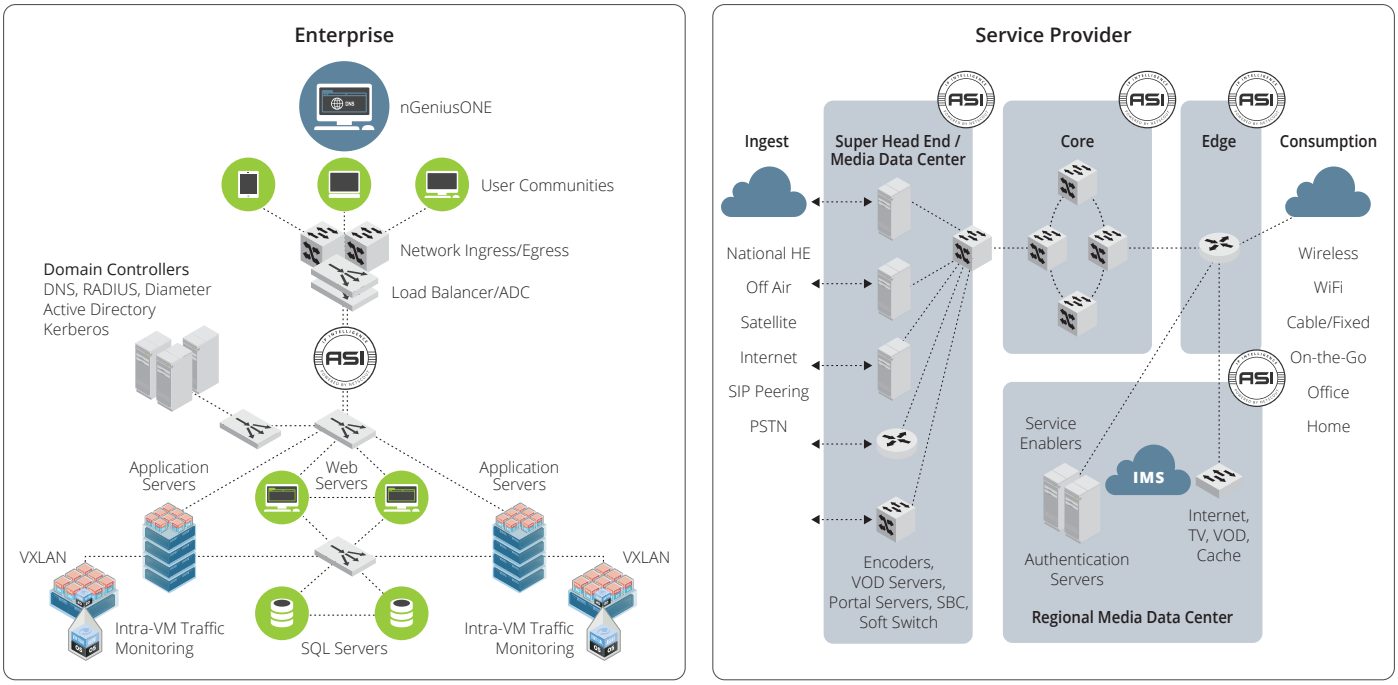

Figure 1: Service Instrumentation in Enterprise and Service Provider Environments.

Core Technologies

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

Adaptive Service Intelligence (ASI)

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

Leverage Packet Flow Data

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

Provide Scalable Packet Flow Access

NETSCOUT Packet Flow Switches and TAPs provide the foundation for a scalable monitoring architecture needed for service assurance. NETSCOUT's nGenius Packet Flow Switches (PFS) filter, aggregate, and distribute the targeted data to NETSCOUT's Intelligent Data Sources in a transparent, selective, and efficient manner. NETSCOUT physical and virtual TAP network monitoring devices provide comprehensive and reliable access to packet flow data and establish strategic service visibility points across the entire service delivery infrastructure.

Service Delivery Monitoring in SDN Environments

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

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

Solution Comparison

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

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

About NETSCOUT Systems, Inc.

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

NETSCOUT

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

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

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

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

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

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

ADV_001-1500 11/2015

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

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

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

Key benefits of the Radware SDN network service infrastructure include:

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

DDoS Protection as a Native SDN Application

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

SDN & NFV for a Scalable Application Delivery Network

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

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

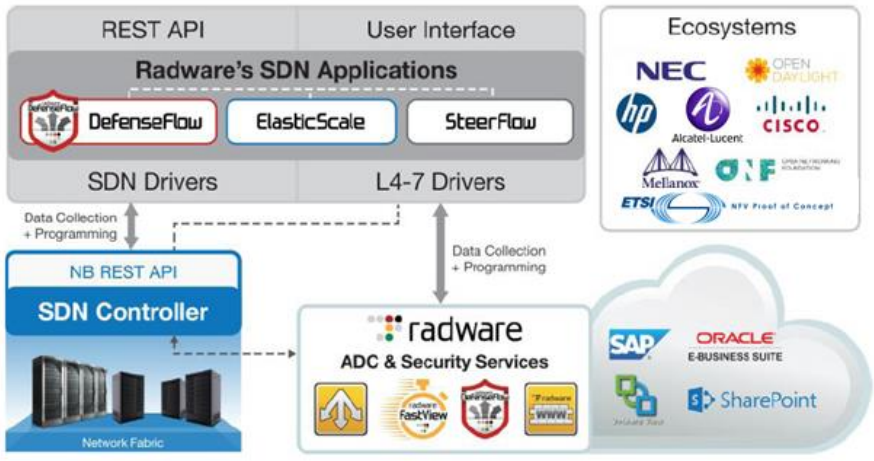

Partnering for Success: Our SDN and NFV Ecosystem

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

Learn More

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

Sonus SBC SWe—Virtual Session Border Controller

Service Providers and Enterprises Leverage NFV from Customer Premises to the Cloud

In the migration to private and public Cloud environments, Network Function Virtualization (NFV) is being adopted as one of the most disruptive changes in telecommunications since the transition to all-IP networks. NFV focuses on new methods for deployment and delivery of telecom services over a software-based network infrastructure. Through NFV, applications that were previously coupled to proprietary hardware can now be instantiated on generic commercial off-the-shelf (COTS) computing hardware.

Designed to operate in virtualized public and private cloud environments, the Sonus SBC SWe is the industry's only software-based SBC that delivers unmatched scalability using the same code base, resiliency, session management, media processing, transcoding, and security technology of a hardware-based SBC.

Sonus provides a migration path to private and public Clouds that can begin at multiple points within a service provider or enterprise network. Because there are many possible paths—from the traditional vertically integrated, custom-hardware-centric way of building networks, to the more flexible, software-defined, and highly elastic future way of building networks—Sonus is committed to working with each service provider or enterprise to achieve the best possible deployment model for their specific network.

Deploying SBC SWe as Virtual CPE (vCPE)

When deployed as a virtual CPE (vCPE), the Sonus SBC SWe enables service providers and enterprises to take advantage of shorter and more flexible deployment cycles for new services and cost savings from virtualization. A perfect application for a virtualized SBC, the emerging vCPE market is based on leveraging the value of NFV and orchestration for network edge services at the customer premises. Deployment models could include a service provider that manages the CPE, a Cloud-based service provider, or an enterprise that wants to retain ownership and control.

Service Providers Offer SBC-as-a-Service (SBCaaS)

With NFV, service providers can deliver 'SBC as a Service' to enterprise customers by moving to a more efficient business model leveraging their virtualization and cloud infrastructure. Hosting SBCaaS eliminates the need to install, deploy, and maintain SBCs at the customer premises. Because NFV enables the rapid instantiation of an SBC, SBCaaS makes it far easier to serve enterprises that have seasonal business or high variability in traffic level, by adding or reducing capacity with a pay-as-you-grow model. Enterprises benefit because they eliminate SBC capital expenses and associated operational expenses. Additionally, the need to carry and manage spare physical inventory, to deal with space, power, and equipment installation issues, and possible CPE obsolescence are also eliminated.

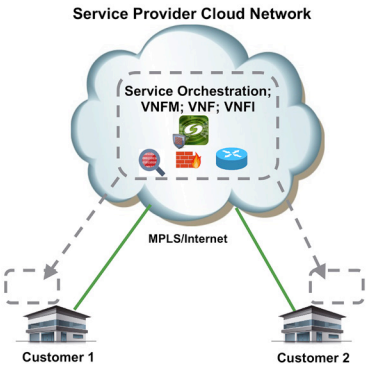

Service Provider SBC-as-a-Service

Virtual SBCs Deployed as Service Provider-owned Managed CPE

Service providers can choose to include a virtual SBC as part of a 'Managed CPE' service, in which low-cost COTS servers on the customer premises are used to host Virtual Network Functions (VNFs), and only the management and service orchestration of the virtualized SBC is conducted from their cloud environment. Orchestration and automation eliminate costly and manual turn-up/turn-down processes, and seamlessly and dynamically scale resources. In this mixed model, new services can be deployed quickly, and since the SBC is a collection of virtual functions, a service provider's upfront CPE costs are contained using COTS servers.

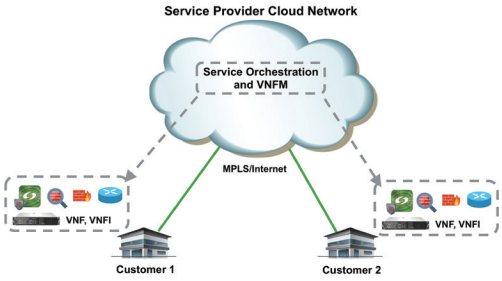

Service Provider-Owned Managed CPE

Enterprise-owned Virtual SBCs Deployed on Customer Premises

When a premises-based SBC is a necessity for an enterprise, the Sonus SBC SWe can be virtualized on common hardware platforms to contain operational costs. The portability and flexibility of the SBC SWe solution make it the ideal choice for enterprises that have invested in virtualization technology, require prepackaged solutions in a box, require remote network deployments, or are supporting opportunities such as entering a new geographic market where deploying a hardware solution would be impractical or cost prohibitive.

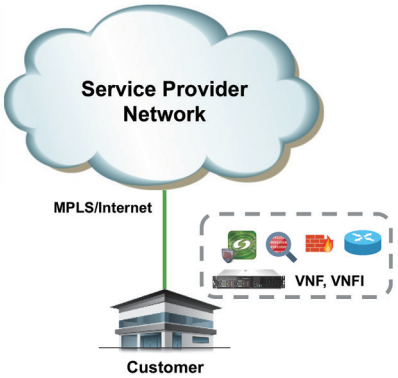

Enterprise-Owned Virtual SBC on Premise

Deploying SBC SWe as a Virtual Interconnect SBC

Interconnection peering points between service providers can be characterized as having dynamic traffic demands—which applies to mobile operators, fixed network operators, and cable providers, and especially to hub/wholesale service providers. Since each service provider is trying to optimize the cost of their interconnection points, traffic patterns are dynamic because they frequently make routing changes on the fly based on least-cost routes and quality-of-service conditions.

Because of this dynamic traffic demand, it's difficult to optimize the SBC capacity required at these interconnection points with a fixed hardware-based solution. This typically leads to overprovisioning of capacity, increasing both capital expense and operational expense. With a virtualized interconnect SBC, it is possible to overcome this limitation.

Deploying virtualized interconnect SBCs in a Cloud environment provides elasticity—the ability to have on-demand instantiation and/or reconfiguration of SBC VNFs to match dynamic traffic demand. A significant advantage of the Cloud environment is the ease and speed at which a new "logical" instantiation of an SBC can be deployed. Given the ability to perform scaling on-demand, service providers can deploy SBC VNFs in their network and scale a single instance or multiple instances independently from very low to very high session counts. With NFV orchestration, this on-demand scaling will be automated and touchless.

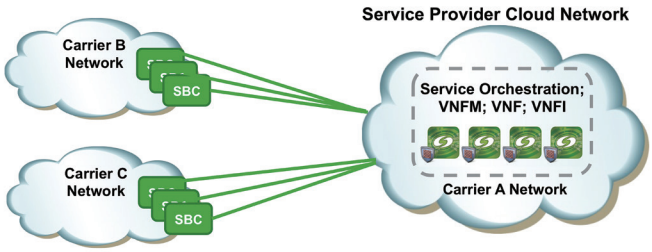

Virtualized Interconnect SBC

Load balancing and high availability are also key to addressing dynamic traffic demands of interconnect SBCs. Load balancing enables efficient scaling by allocating work load across a resource pool. In a Cloud environment, with only virtual resources, this is essential for the scalable deployment of SBCs. With a well-designed load balancing strategy, virtual resources are optimized to fine-tune the overall status of the application processing.

Real-time applications in a virtualized, cloud-based environment have the same high availability requirements for service, subscriber, and call resiliency as they do in traditional network environments. High availability requires an architecture where critical state information in an SBC is backed up in another node, ready to handle the traffic in the event of failure. Any change in the network is transparent to all peers, and no action is required to achieve this seamless transition.

About Sonus Networks

Sonus brings the next generation of Cloud-based SIP and 4G/VoLTE solutions to its customers by enabling and securing mission-critical traffic for VoIP, video, IM, and online collaboration. With Sonus, enterprises can intelligently secure and prioritize real-time communications, while service providers can deliver reliable, secure real-time services for mobile, UC, and social applications. Sonus offers an award-winning portfolio of hardware-based and virtualized Session Border Controllers (SBCs), Diameter Signaling Controllers (DSCs), Cloud Exchange Networking Platform, Policy/Routing servers, and media/signaling gateways. Visit www.sonus.net or call 1-855-GO-SONUS.

Sonus™
Cloud communications made smarter