The 2016 Guide to SDN and NFV

Part 1: Software Defined Networking (SDN): A Status Update

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

Platinum Sponsors:

A10

CISCO

hp

MASERGY
Performance Beyond Expectations

NETSCOUT™
Guardians of the Connected World

Sonus®

Gold Sponsors:

radware

SevOne

Produced by:

Webtorials

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 first of the serial publications and it will focus on providing insight into the current status of SDN. 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. The Use Cases and Business Considerations 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 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 October 2015. Throughout The Guide the 131 network professionals who completed the survey will be referred to as The Survey Respondents.

Status of SDN Adoption

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

Table 1: Approach to Implementing SDN

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

The data in **Table 1** indicates that the implementation of SDN in production networks remains limited. In addition, comparing the data in **Table 1** to the responses to the same question a year ago yields the conclusion that:

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

The Open Networking Foundation (ONF)

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

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

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

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

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

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

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

Open Source projects will likely accelerate the adoption of SDN.

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

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

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

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

The Northbound & Southbound Interfaces

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

Figure 1: ONF's SDN Architecture

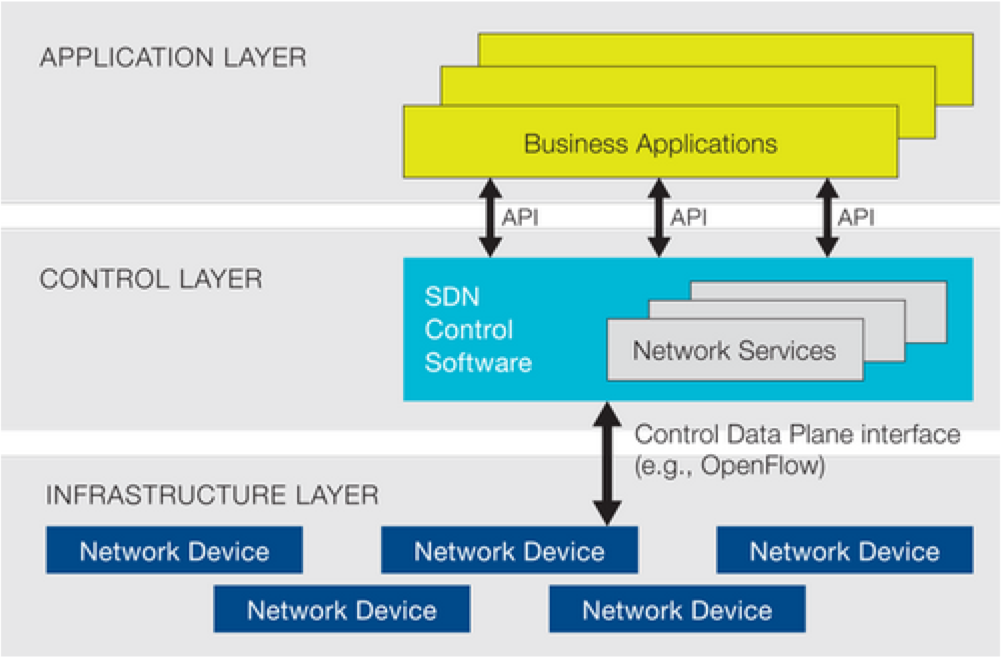

The Northbound Interface

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

In 2013 the ONF established a working group to focus on the NBI. Given that traditional standards activities are widely viewed as not being agile enough for the current environment, the goal of the working group was not to develop a standard for the NBI in the traditional sense of the term *standard*. Rather, the goal was to develop a rough consensus and collaboration around community developed NBIs. The group's complete charter was outlined in a white paper. One of the interesting concepts that that white paper discusses is the need for APIs at different "latitudes". The idea was that a business application that uses the NBI should not require much detailed information about the underlying network. Hence, applications like this would require a high degree of abstraction. In contrast, network services such as load balancing or firewalls would require far more granular network information from the controller and hence, not need the same level of abstraction.

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

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

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

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

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

The Southbound Interface

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

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

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

Table 2: Likely Use of OpenFlow

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

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

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

The Overlay and the Underlay Model

There are two primary approaches that vendors are taking to implement the architecture depicted in **Figure 1**. These two approaches are the:

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

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

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

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

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

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

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

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

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

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

Open Source SDN Controllers

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

The ON.Lab

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

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

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

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

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

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

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

The ONOS community has expanded to include vendors.

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

The ONOS project is part of the Linux Foundation.

The OpenDaylight (ODL) Project

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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.

Division Cofounders:

Jim Metzler

Steven Taylor

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.

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

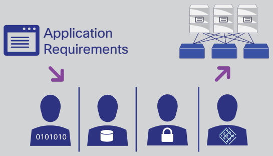

IT administrators work in silos.

Weeks

There is no shared architectural model.

With ACI

New Application

Simpler

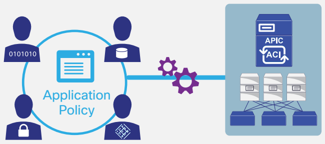

Faster

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

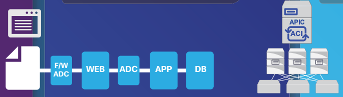

Application Network Profile

- 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.

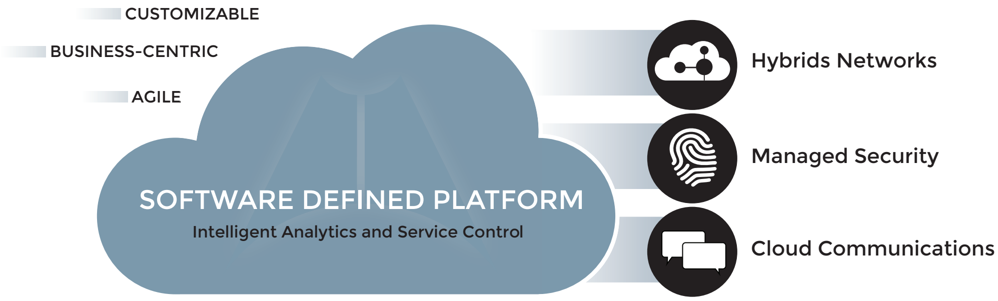

This WAN innovation is being enabled by a Software Defined Networking (SDN) architecture, which enables administrators to rapidly change configurations as performance requirements demand.

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

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

VIRTUALIZATION

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

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

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

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

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

FLEXIBLE BY DESIGN

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

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

MASERGY
Performance Beyond Expecations

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

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

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

Contact us for a free consultation.

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

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

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.

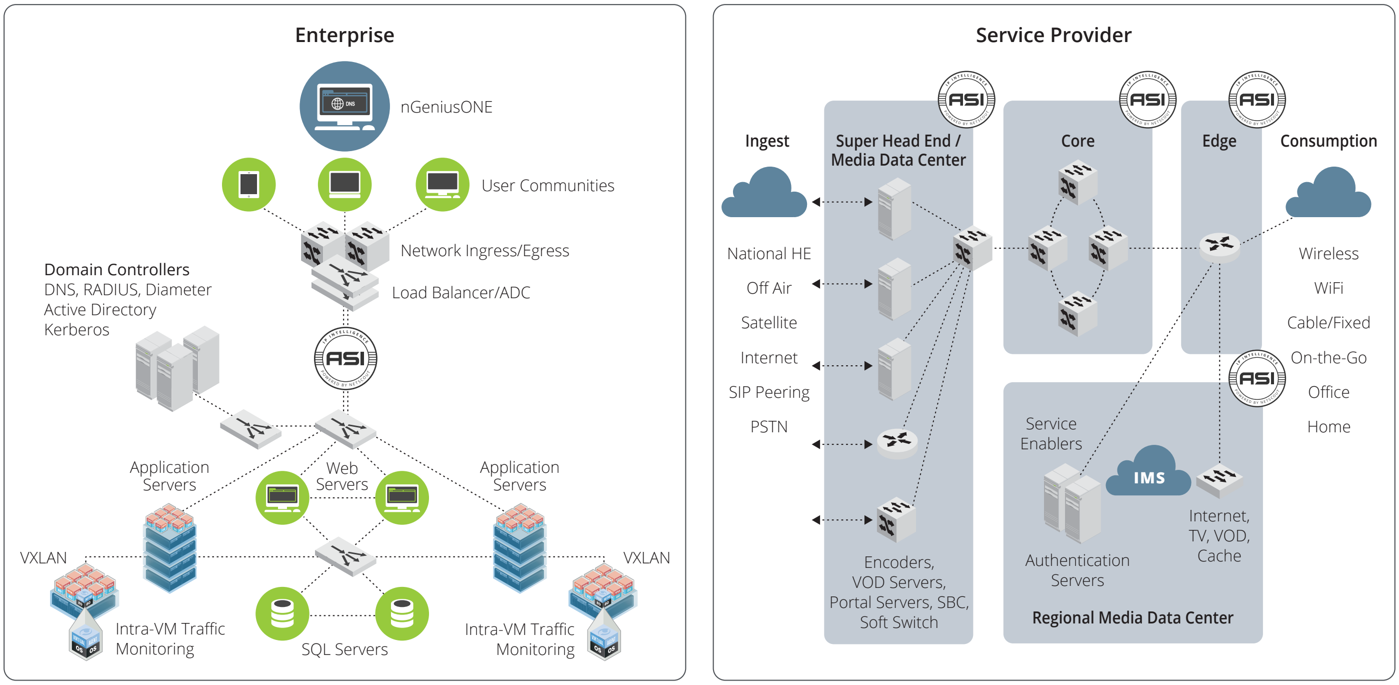

Figure 1: Service Instrumentation in Enterprise and Service Provider Environments.

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.

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

radware

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.

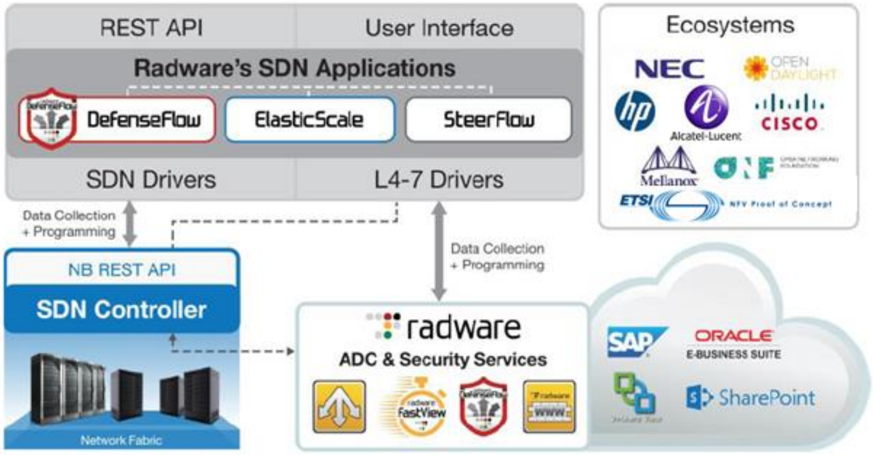

Partnering for Success: Our SDN and NFV Ecosystem

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.

VellOS® Cloud Exchange Networking Platform

New Cloud Services for Service Providers; Efficient Cloud Networking for Enterprises

As enterprises and service providers discover the value of cloud networking architectures, they seek technology that transcends traditional networking products and wide area networks (WANs). Traditional networking infrastructure is inadequate to address the requirements of migrating applications and data to Cloud environments. With limited capability to cost-effectively handle rapidly changing traffic dynamics, the infrastructure cannot optimize network resources, and is cumbersome to perform network-wide policy and configuration management.

VellOS is a market-leading, virtualized software-defined platform in a new class of Cloud exchange networking products that integrates private, hybrid, and public clouds into a seamless IT environment. VellOS unifies the control of the optical, Ethernet, and virtual networking planes under common open software, enabling real-time network visibility, control, and optimization.

With the application-aware network intelligence provided by VellOS, enterprises and service providers can orchestrate and dynamically deliver just-in-time data center connectivity, improve business-critical application performance, ensure business continuity, and transition the WAN into a strategic asset. Via its policy-driven architecture and logically centralized resource control, VellOS simplifies network and application roll-out.

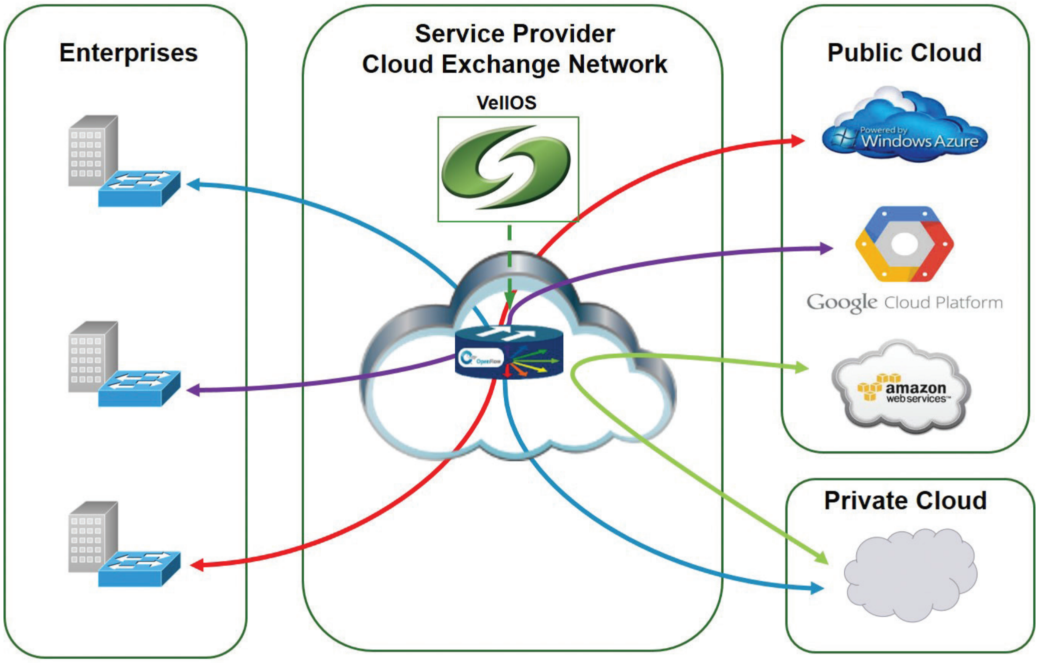

VellOS leverages SDN to transform the Cloud Exchange Network

Dynamic Data Center Interconnect

As enterprises place increasing emphasis on the Internet and Cloud-based services to support internal and external communications and transactions, Cloud deployments must have optimal network connectivity. There are three critical aspects the cloud exchange network must address. First, connectivity should be managed in real-time, based on business priorities, including those defined in business continuity plans. Second, traffic between Data Centers should be prioritized based on business priorities for each different application. Third, bandwidth utilization should be maximized, thereby reducing the cost of Data Center connectivity.

With VellOS, business policies are automatically translated into simplified provisioning and network infrastructure configuration, on-demand connectivity based on application requirements is orchestrated, and infrastructure changes are dynamically implemented.

VellOS makes it possible to achieve significant cost savings by maximizing bandwidth utilization, as well as ensuring the availability and continuity of mission-critical business applications. As a result, service providers and Cloud hosting providers can turn their Cloud exchange network into a revenue-generating opportunity by providing Data Center Interconnect as a "Network-as-a-Service".

Real-Time Service Quality for Unified Communications

Applications that deliver real-time unified communications (UC) have specific requirements for low latency, jitter, and packet loss. When provided from the Cloud and delivered across a cloud exchange network, these applications require policies that are application-aware to ensure bandwidth allocation and packet prioritization on a per-session basis. VellOS provides API interworking with various vendors' UC implementation, so that UC traffic will get priority over other IP application traffic across the cloud exchange network. This traffic prioritization is dynamic and can be performed at the granularity of each session, each user, or each UC application.

> "SDN is one of the most transformative business and technology trends the telecomm industry has seen in decades. Faster time-to-market for new services, faster provisioning, delivery, and upgrade times for existing services, and reduced human error and lower OPEX costs from software automation are the primary reasons for adopting SDN."
>
> – Sterling Perrin,
> Senior Analyst, Heavy Reading

Security

As applications and data migrate to the Cloud, solutions previously used to secure traffic across the corporate LAN or across private lines may no longer be viable. In situations where an application relies on the cloud exchange network for security, VellOS will block unknown traffic from accessing the cloud exchange network. Only traffic from known users or devices is allowed. Data will not flow across the cloud exchange network by default. Only provisioned IP addresses that are allowed to exchange data can ARP each other. Bandwidth allocation enforcement prevents traffic flooding, as well as ping, ARP, or LLDP floods. For ease of implementation, VellOS can integrate with an existing authorization/authentication solution to validate users.

Business Continuity and Disaster Recovery

As enterprises place increasing emphasis on Internet and Cloud-based services to support their business, they require a different business continuity paradigm for their data, applications, and transactions. These enterprises must ensure they have access to Cloud-based applications and Data Centers, even during disasters, for transparent and continuous operation. The availability of application and data must flow quickly to minimize any negative impact on business operations. With VellOS, business continuity/disaster recovery policies are automatically translated into simplified provisioning and network infrastructure configuration. This ensures cloud exchange network changes are rapidly implemented in the event of unforeseen or unplanned changes in network topology or network behavior.

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.

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