The 2016 Guide to SDN and NFV

Part 2: The Use Cases and Business Case for SDN

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 second of the serial publications. It will focus on describing the most popular SDN use cases. This document will also provide insight into how to create a business case for 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.

SDN Use Cases, Drivers and Inhibitors

Focus of SDN Deployment

While the use of SDN in data centers receives the majority of attention, it is also possible to implement SDN in branch and campus networks as well as in wide area networks (WANs). In order to understand where SDN will likely be implemented, The Survey Respondents were asked "If your organization is likely to implement SDN sometime over the next two years, where are you likely to implement it?" Their responses are summarized in **Table 1**.

Table 1: Focus of SDN Deployment

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

One observation that can be made from the data in **Table 1** is:

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

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

Data Center

Data Center Use Cases

Virtual Machine Migration

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

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

Service Chaining

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

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

Security Services

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

Load Balancer Services

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

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

Drivers and Inhibitors of SDN in the Data Center

Table 2 and **Table 3** contain the responses of The Survey Respondents when asked to indicate the two factors that would drive and the two factors that would inhibit their organization's implementation of SDN in a data center.

Table 2: Drivers of Implementing SDN in a Data Center

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

One observation that can be drawn from the data in **Table 2** is that:

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

Table 3: Inhibitors to the Adoption of SDN in a Data Center	
Impediment	**Percentage**
Concerns about how we would integrate SDN into the rest of our infrastructure	30%
The immaturity of the enabling technologies	28%
The confusion and lack of definition in terms of vendors strategies	23%
Other technology and/or business priorities	20%
The lack of a compelling business case	17%
Possible security vulnerabilities	15%
Concerns about how we would manage SDN	12%
The lack of a critical mass of organizations that have deployed SDN	8%
No inhibitors to implementing SDN	7%
Concerns that the technology will not scale to support enterprise sized networks	6%
Other	4%

Unlike the situation shown in **Table 2** in which there were two clear drivers of SDN deployment in the data center:

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

WAN

WAN Use Cases

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

The Google G-Scale WAN

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

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

identified a number of benefits that are associated with its G-Scale WAN backbone including that Google can run the network at utilization levels up to 95%.

SD-WANs

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

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

Drivers and Inhibitors of SDN in the WAN

Table 4 and **Table 5** contain the responses of The Survey Respondents when asked to indicate the two factors that would drive and the two factors that would inhibit their organization's implementation of SDN in the WAN.

Table 4: Drivers of Implementing SDN in a WAN

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

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

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

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

Table 5: Inhibitors to the Adoption of SDN in the WAN	
Impediment	**Percentage**
Concerns about how we would integrate SDN into the rest of our infrastructure	25%
The lack of a compelling business case	25%
The immaturity of the enabling technologies	25%
The immaturity of the current products	22%
Possible security vulnerabilities	22%
The confusion and lack of definition in terms of vendors strategies	17%
Other technology and/or business priorities	16%
The lack of a critical mass of organizations that have deployed SDN	11%
Concerns about how we would manage SDN	9%
Concerns that the technology will not scale to support enterprise sized networks	9%
No inhibitors to implementing SDN	7%
Other	3%

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

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

Branch and Campus

Branch and Campus Use Cases

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

Dynamic QoS & Traffic Engineering

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

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

Unified Wired and Wireless Networks

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

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

Role Based Access

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

Drivers and Inhibitors of SDN in Branch and Campus Networks

Table 6 and **Table 7** contain the responses of The Survey Respondents when asked to indicate the two factors that would drive and the two factors that would inhibit their organization's implementation of SDN in branch and campus networks.

Table 6: Drivers of Implementing SDN in Branch and Campus Networks

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

Observations that can be drawn from **Table 6** include:

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

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

Table 7: Inhibitors to the Adoption of SDN in Branch and Campus Networks

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

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

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

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

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

However,

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

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

The SDN Business Case

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

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

Financial metrics

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

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

The Components of a Business Case

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

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

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

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

Consolidation of Resources

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

IT Agility

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

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

Business Agility

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

Improved Application Performance

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

Increased Network Availability and Performance

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

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

Improved Security

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

Enhanced Management and Visibility

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

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

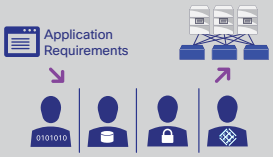

IT administrators work in silos.

Weeks

There is no shared architectural model.

With ACI

New Application

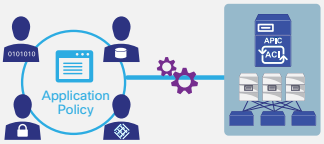

Simpler

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

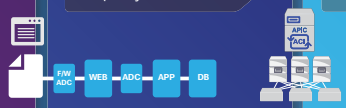

Application Network Profile

- 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.

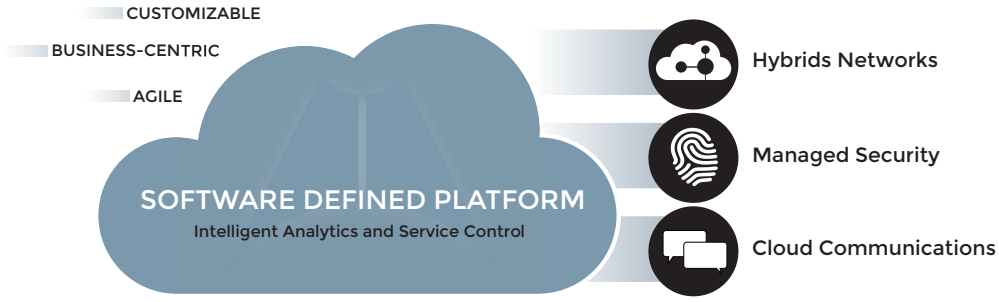

This WAN innovation is being enabled by a Software Defined Networking (SDN) architecture, which enables administrators to rapidly change configurations as performance requirements demand.

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.

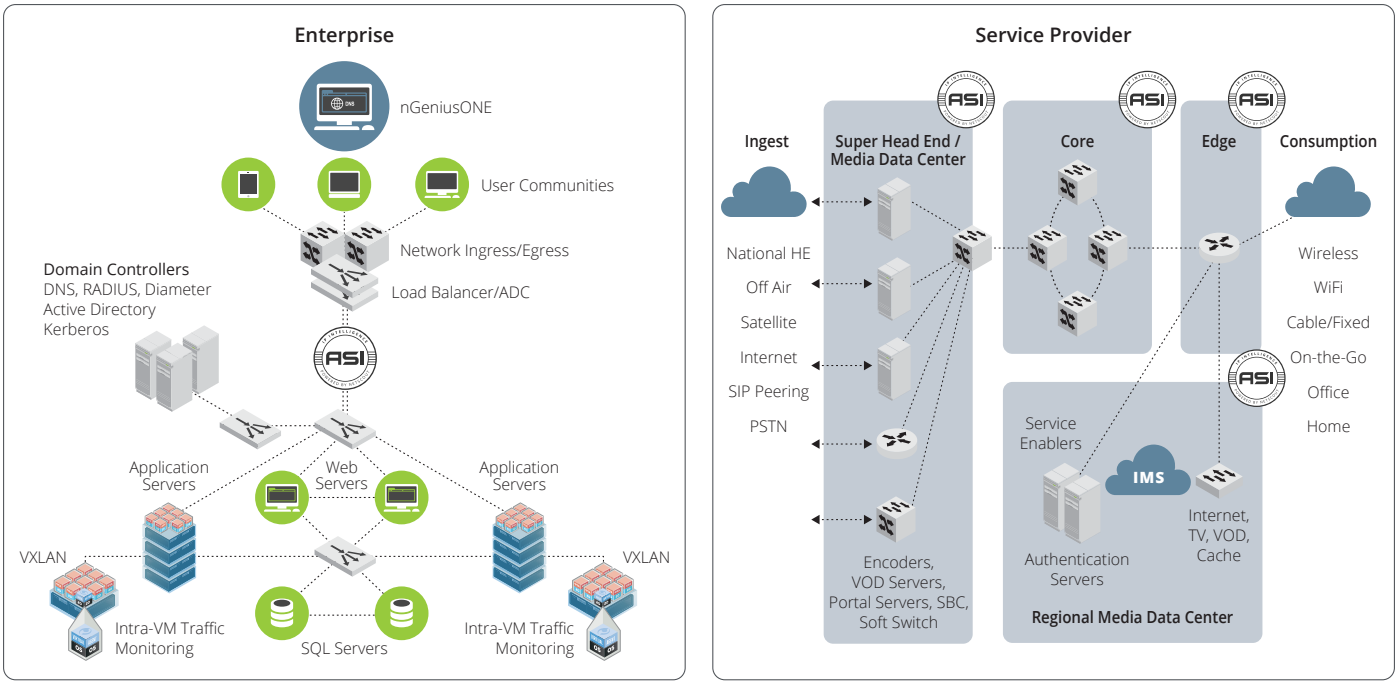

Figure 1: Service Instrumentation in Enterprise and Service Provider Environments.

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.

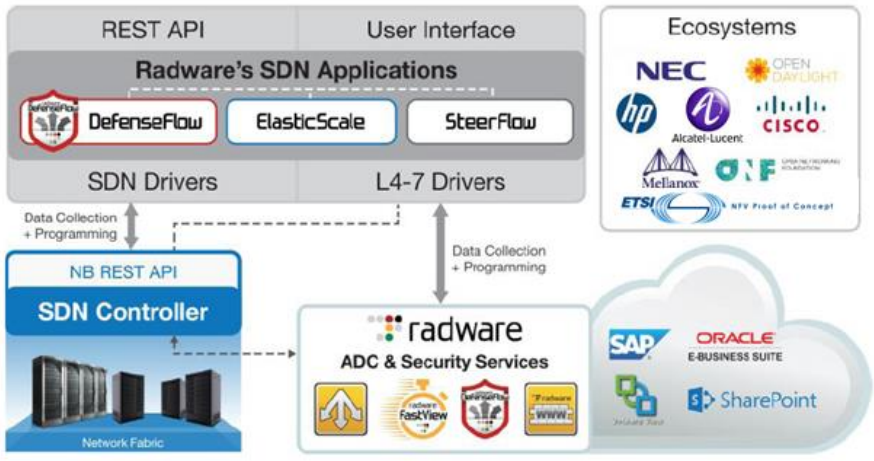

Partnering for Success: Our SDN and NFV Ecosystem

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.

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.

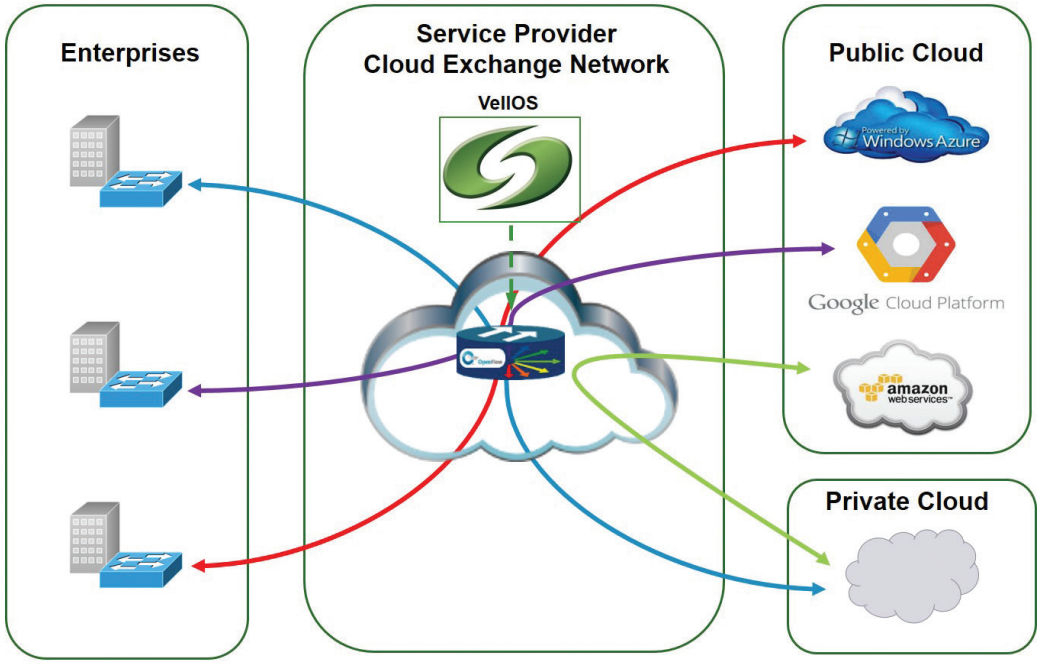

VellOS leverages SDN to transform the Cloud Exchange Network

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.

Sonus™
Cloud communications made smarter