How to Succeed With SD-WAN Using Virtualized Service Assurance

An Operational Practice prepared for SCTE/ISBE by

Etienne Martel
Solution Manager
Accedian
2351 Blvd Alfred-Nobel, Suite N-410
Saint-Laurent (Montreal), Quebec, H4S 2A9, Canada
514-331-6181
emartel@accedian.com

Gregory Spear
Solution Manager
Accedian
2351 Blvd Alfred-Nobel, Suite N-410
Saint-Laurent (Montreal), Quebec, H4S 2A9, Canada
514-331-6181
gspear@accedian.com

Table of Contents

Title **Page Number**

List of Figures

Title **Page Number**

Introduction

1. Executive Summary

To succeed in the enterprise market, cable multiple system operators (MSOs) must serve nationwide or global corporations that expect their providers to reach all their sites, so they do not have to assemble their own network. This means crossing multiple service areas, including outside the MSOs' footprint.

Software-defined wide area networks (SD-WAN) can help MSOs reach on and off-net sites, over any access media, uniformly—allowing them to leverage the scale and reach of their extensive DOCSIS and Carrier Ethernet services, augmented where required with third party access.

The pitfall: SD-WAN appliances do not offer standards-based test, turn up and monitoring functions required to offer service level agreement (SLA)-grade services. SD-WAN solutions use proprietary monitoring and reporting methods, which do not interoperate with existing network equipment. Because SD-WAN may only be required in certain customer locations, any implementation has to interact seamlessly with traditional service delivery methods.

This is not optional. All MSOs in North America offer SLA-backed services over fiber, and the majority over DOCSIS too, according to a 2017 Heavy Reading study[1]; best-effort only services will not satisfy the enterprise market requirements for uniform services and stringent SLAs.

Virtualized test probes and test reflectors cost-efficiently replicate network interface device (NID) functionality, bringing standards based turn-up testing, monitoring and operations & maintenance (OAM) functions to SD-WAN endpoints. Virtualized instrumentation uplifts SD-WAN with carrier-grade functionality, making it interoperate with existing network infrastructure, operations procedures, and support systems.

In today's SD-WAN market, the two main deployment architectures are centralized or distributed. They will both benefit from service assurance solutions that can be deployed as software and optionally enhanced with NFV-powered hardware modules.

When selecting a service assurance solution, it is important to choose an industry-proven method for extending standards-based test, measurement, and OAM to virtualized environments for SD-WAN and x86 infrastructure to ensure satisfactory coverage and unified visibility. The interoperability provided by a standards-based solution enables centralizing performance monitoring data into existing reporting and fault-management systems to achieve the operational success of the technology. It also opens the possibility to optimizing the SD-WAN performance by leveraging the highly granular and micro-second accurate end-to-end monitoring data.

[1] *Breznick, A. Heavy Reading*, January 2017. How cable can conquer the enterprise market. Retrieved from https://accedian.com/wp-content/uploads/2017/01/HR_Accedian_Cable_Enterprise_WP_1-24-17.pdf

Content

2. SD-WAN adoption — Enablers and Drivers

The ever-growing use of cloud-based applications by enterprises is making SD-WAN more relevant every day. Software-defined networking (SDN), intially reserved for data center applications—along with a number of other technology enablers—have set the table for SD-WAN to disrupt traditional WAN architectural models prevalent within most enterprises.

Technology enablers for enterprise SD-WAN include:

- Widespread availability of scalable and elastic cloud computing for the enterprise.
- Widespread availability of faster, more diverse, and more reliable wired and wireless broadband internet access technologies.
- SDN concepts: control plane and data plane separation.
- Application-aware routing with deep-packet inspection (DPI) instead of traditional IP routing.
- Availability of affordable x86 network appliances (commercial off-the-shelf hardware/COTS)

From hybrid models to full scale implementations, the advantages offered by SD-WAN cannot be ignored. SD-WAN offers increased network agility, better overall performance, lower cost per megabit, and a new WAN architecture that complies with the software as-a-service (SaaS) and cloud computing business model.

All enterprises are now considering SD-WAN solutions for their next WAN refresh; none plan to fully retain the traditional single WAN, hub and spoke model with centralized internet access at the hub. Enterprises are compelled to consider SD-WAN, if only for the ability to have branch locations locally break-out to the internet in a secure and controlled way instead of backhauling all internet traffic across the WAN to the hub.

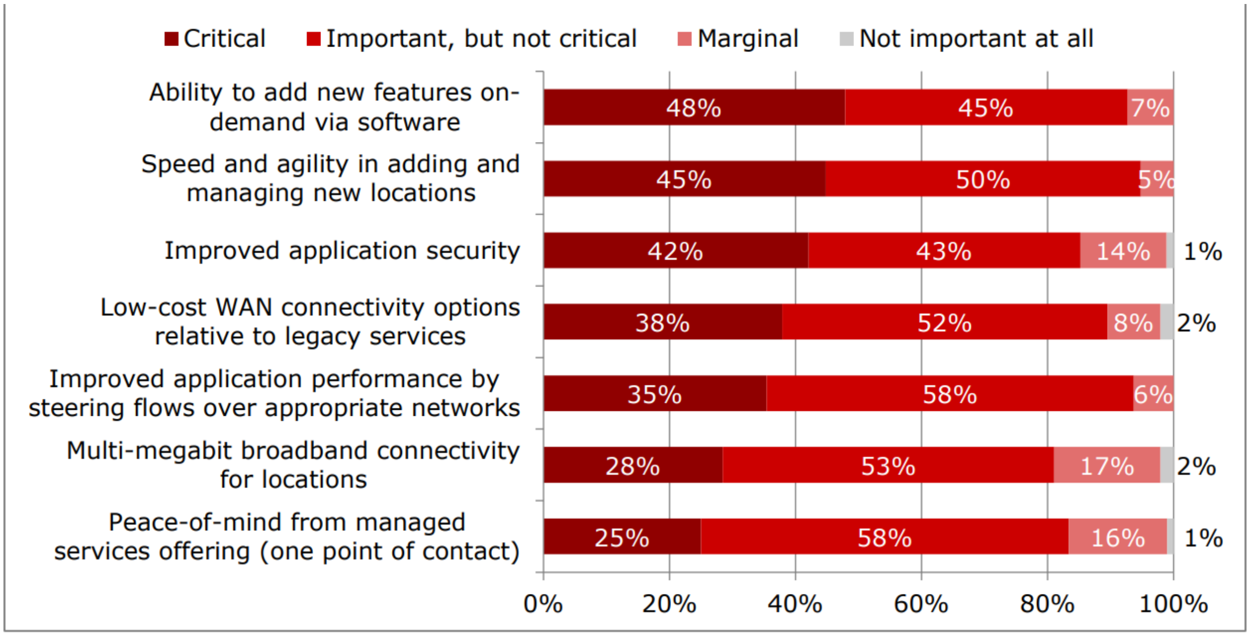

N=96
Source: Heavy Reading December 2016 Operator Views on Emerging SD-WANs Survey, Sponsored by ADVA

Figure 1 - Most Important Expected Benefits for Operators' Customer[2]

From the operators point-of-view, SD-WAN is appealing for similar reasons, primarily to offer a low-cost bandwidth enhancement to offered services—but also because it unlocks the ability to turn up new features on-demand, enhancing agility and speed for service delivery and the initial service turn-up.

As shown in Figure 1, operators are deploying SD-WAN in their network to gain agility and flexibility first and foremost. The software automation at the heart of the SD-WAN solution will allow for the creation of fully dynamic networks and give end-users a control into the nature and level of services they require on an ongoing basis. From the MSO point-of-view, having the ability to both deploy and maintain features and services via software deployment is crucial. Running software on COTS servers dramatically lowers both risks and costs when compared with the traditional dedicated hardware appliance solutions that required extensive trials to approve and the trained personnel, space, power and cooling to run.

As it stands, SD-WAN deployments are still in the early adoption phase and do not offer the same service-level expectations as traditional business service WAN offerings. Now, as widespread adoption continues, operators find that SD-WAN managed services must deliver the same quality levels as traditional WAN offerings. Operators therefore need tools that offer the visibility and reporting capabilities to manage network performance and SLAs.

[2] *Sterling, P. Heavy Reading.* February 2017. Operator Success in the New Age of the Software-Defined WAN, Retrieved from https://resources.ext.nokia.com/asset/201132

3. Business Services Over SD-WAN Lifecycle Overview

Specific operational practices pertain to each phase of the business services over SD-WAN service lifecycle as illustrated below in Figure 2:

1. Provisioning and Turn Up: Deployment and service activation testing (SAT)
2. Performance Management: Performance monitoring and SLA reporting; collecting and presenting key performance metrics
3. Fault Management: techniques to identify, isolate, and troubleshoot service issues

These three phases are consistent with the Metro Ethernet Forum (MEF) definition of the Carrier Ethernet service lifecycle[3], which serves as an established model for commercial connectivity.

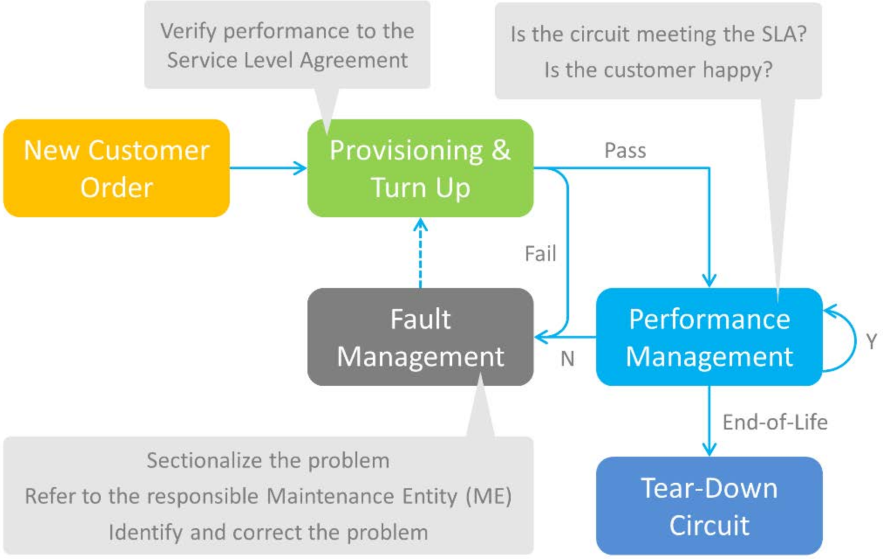

Figure 2 - Metro Ethernet Forum Service Lifecycle[4]

This paper is structured to address the operational practices associated with each stage of the service lifecycle, as it applies to business services over SD-WAN.

[3] *MEF Forum.* April 2012. Introducing the Specifications of the MEF. MEF 38: Service OAM Fault Management YANG Modules Technical Specification. Retrieved from: http://slideplayer.com/slide/5687304/

[4] *MEF Forum.* March 2016. Service Operations Specification MEF 55: Lifecycle Service Orchestration (LSO): Reference Architecture and Framework. Retrieved from: http://dev.mef.net/Assets/Technical_Specifications/PDF/MEF_55.pdf

4. Deployment Options Over Common SD-WAN Architectures and Hybrid Services

Any performance monitoring solution requires four key elements: 1) test session control; 2) a test packet generator; 3) a test packet reflector or receiver; and 4) precision timestamping. In a traditional WAN network, these elements were often customer premises equipment (CPE)-based network interface devices (NIDs) and centralized test suites, all of which required time to install and configure—something that virtualized network services have eliminated.

The operational practices can be implemented using a network-embedded architecture that employs small footprint, programmable service assurance hardware modules (vCPE modules) augmented by virtualized service assurance functions hosted on a centralized, virtualized performance assurance controller (vPAC). A lightweight, stand-alone orchestratable software agent is another viable way to instrument the network; this architecture can offer a complete software-only solution. However, it cannot rival the precision and the feature-set offered by vCPE modules. Section 4 introduces these architectures, as well as operational considerations that facilitate integration of these approaches with existing operational support systems (OSS), network management systems (NMS), and virtual network function (VNF) orchestrators.

4.1. Architecture Overview: Using Virtualized Performance Assurance Controller VNFs and Lightweight Stand-Alone Orchestratable Software Agent VNFs.

For deployments where service assurance using standard-based protocols is needed, but the added-benefits offered by the NFV-enabled modules are not required, stand-alone orchestratable software agents can be used to offer the reflection capabilities needed to complement a centralized performance monitoring approach using vPAC VNFs as probe generators.

The main benefit of this software-only architecture is the deployment speed and agility offered by being able to remotely and centrally deploy, configure, and run everything needed to instrument an existing network, on-demand and with minimal expense. Standards-based monitoring methods integrate the network itself into a ubiquitous instrumentation layer. With this visibility centralized in data centers shared with SDN control and big data analytics, providers have an integrated foundation to deliver a new level of customer experience.

The vPAC assumes all session setup, control, and sequencing functions, as well as results analysis and reporting to file servers. As a virtual network function (VNF), vPAC instances can be deployed and orchestrated seamlessly with the network service descriptors, allowing fully-automated setup and assurance of virtual service chains.

The lightweight software agent VNF offers reflection capabilities required to instrument the network with any orchestrator and can easily run un-privileged on any Linux based operating system.

The lightweight software agent VNF also enables bi-directional measurements, unrivaled metrics set, measurement granularity, and third party interoperability—features that are unavailable when using built-in standard open-source tools (such as ICMP ping) or even proprietary measurement methods offered by SD-WAN vendors.

4.2. Architecture Overview: Using Controller VNFs and NFV-Powered Hardware Modules.

In the context of this document, the term *vCPE* will refer to the strategy of virtualizing as many customer-located networking functions as possible, while retaining the minimum hardware necessary for service delivery, consistent with performance, reliability, and quality of experience (QoE) expectations. An example of a vCPE strategy—where onsite hardware appliances performing firewall, PBX, and routing functions have been virtualized—is illustrated in Figure 3, below. Virtualization is accomplished by transferring local networking functionality to software-based VNFs, which can be hosted on low-cost COTS servers or cloud infrastructure.

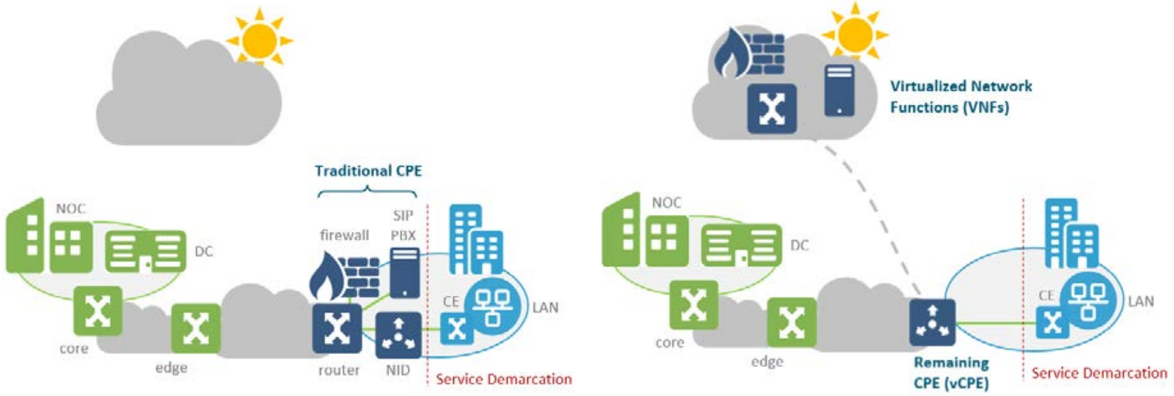

Figure 3 - vCPE: Traditional vs. Virtualized Customer Premises Equipment Example

In the context of SD-WAN, this approach can be used to introduce customer premises-located performance monitoring, turn-up test, service OAM (SOAM) and troubleshooting functionality, which—in the case of fiber business services—is normally provided using a NID. Reducing hardware appliances required at the branch site is a key benefit of SD-WAN; installing a standard NID along with the SD-WAN appliance is not normally a feasible CPE option.

NFV-powered hardware modules can offer the same level of performance monitoring precision, as well as loopback and full line-rate turn-up test capabilities at a fraction of the cost of a NID, making this approach an economically viable fit when deploying SLA-grade business services over SD-WAN. The solution delivers complete quality of service (QoS) and QoE insight without compromise. In addition to supplementing the SD-WAN appliance (or COTS server) with service assurance features, this approach has a number of other benefits:

1. Truck-rolls are reduced over the service lifecycle when compared to handheld test sets, as a single vCPE module can remotely perform turn-up testing, continuous monitoring, and on-demand troubleshooting.
2. Compatibility with existing hand-held Ethernet test sets and third-party centralized monitoring probes allows straightforward integration into existing operational practices and infrastructure.
3. By employing NFV, new functionality can be added to the vCPE module remotely, without impacting the service. This allows MSOs to introduce new, performance-assured commercial services without requiring new equipment on-site.

An example of how NID functionality can be virtualized using NFV is shown in Figure 4 below. On the left, you can see a traditional CPE NID, it it contains both a control plane and a data plane. On the right, we have disaggregated the functionality into a layer of software to deliver the control plane in the form of a service assurance VNF and a layer of hardware in the form of vCPE Modules which contain just enough hardware to deliver the required data plane features at the site while leveraging the VNFs for any compute intensive job.

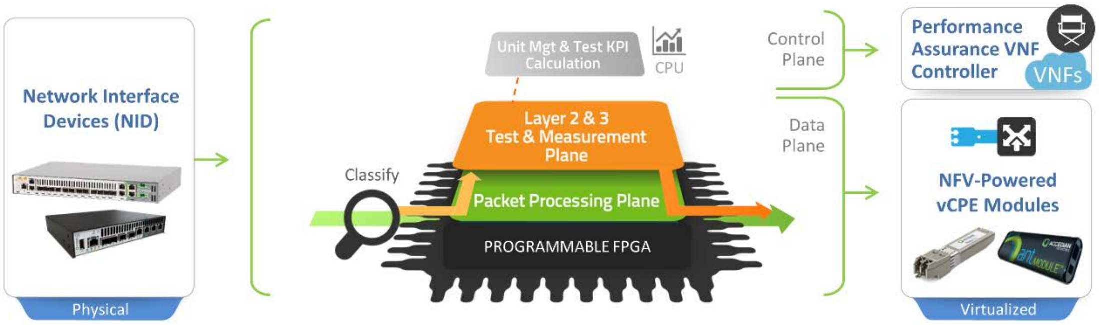
Network Interface Devices (NID)
Physical
Unit Mgt & Test KPI Calculation
CPU
Layer 2 & 3 Test & Measurement Plane
Classify
Packet Processing Plane
PROGRAMMABLE FPGA
Control Plane
Data Plane
Performance Assurance VNF Controller
VNFs
NFV-Powered vCPE Modules
Virtualized

Figure 4 - Virtualization of NID Architecture Using NFV

The connection between the service assurance VNF controller and each module needs to be reliable, secure, and lossless (e.g. transmission connection protocol (TCP) based) to ensure the vCPE module can assume the same level of functionality as a traditional NID. As shown in Figure 5 below, this management 'tunnel' is critical to support service assurance VNFs, as raw data is returned to the controller for test results calculation, performance montoring, and fault reporting, in addition to performance monitoring session control, module management, synchronization information, etc. In an NFV-based vCPE architecture, the 'lossless' control sessions allow each remote module to virtually become a remote 'port' of the controller, which is analogous to a virtualized NID that can support many remote endpoints.

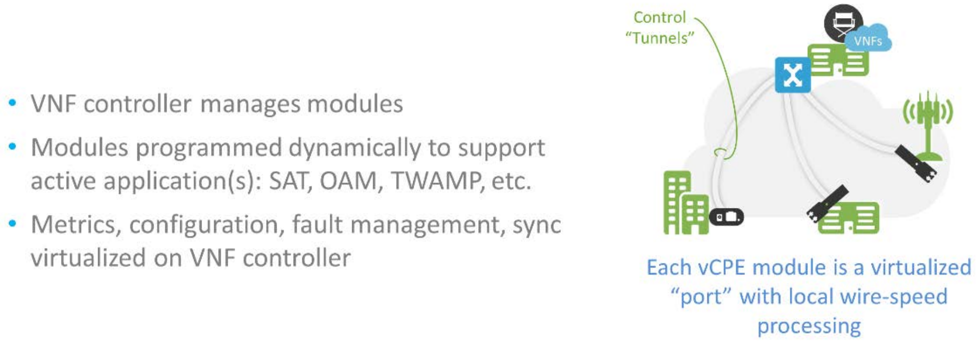
- VNF controller manages modules
- Modules programmed dynamically to support active application(s): SAT, OAM, TWAMP, etc.
- Metrics, configuration, fault management, sync virtualized on VNF controller

Control "Tunnels"
VNFs
Each vCPE module is a virtualized "port" with local wire-speed processing

Figure 5 - NFV-based vCPE Control Tunnel & Function Locations

SD-WAN architectures virtualize some or all customer premises functions with a simple COTS server at the customer site. As part of their standard feature-set, SD-WAN solutions implement path monitoring and measurement. However, these measurements are typically lacking for managed business services over SD-WAN deployments because those service assurance functions implemented purely in software:

1. lack sufficient time stamping precision and packet transmission scheduling control to meet the requirements of:
   a. Full line-rate test traffic generation and loopback for SAT and troubleshooting.
   b. Precise traffic generation sequencing required by common turn-up test standards (where inter-packet delay needs to be controlled for burst testing, for example).
   c. Microsecond-level latency measurement precision required to monitor and report on commerical services SLAs.
2. are subject to the resource-sharing of the x86 system. This causes additional uncertainty in the results by bundling the performance of the x86 system with the perfomance of the network itself.

In addition, SD-WAN solutions use proprietary monitoring and reporting methods which do not interoperate with existing network equipment (or other SD-WAN vendors). Because SD-WAN may only be required in certain locations, any service assurance implementation has to interact seamlessly with the traditional service delivery methods.

Relying on built-in SD-WAN monitoring also has the effect of creating a potential blind spot. This is especially true when considering service activation testing (SAT) such as RFC2544[5] or ITU-T Y.1564[6] which have no support from the SD-WAN vendors. Moreover, the SD-WAN built-in performance monitoring functions can also only provide a top-down view of performance—the over-the-top (OTT) path. This view presents no insight into why a specific path is operating badly, just that it is not performing. Complementing this top-down view with a bottom-up perspective provided by hop-by-hop or layer 2-3 path monitoring can add the missing pieces to more efficiently run an assured SD-WAN services, enabling detailed troubleshooting and measurable quality improvements.

Running a unified service assurance solution across both the incumbent part of the network and the SD-WAN part of the network also has the benefit of offering a unified level of precision and reporting intervals. As such, pin-pointing events and segmenting the network will ease troubleshooting and accelerate mean time to resolution (MTTR) when issues arise.

Aside from enabling these capabilties, vCPE modules also offer a number of other advantages over traditional test set and centralized probe solutions:

1. Modules can monitor and test between themselves: for site-to-site monitoring, end-to-end turn-up testing, and troubleshooting between customer service endpoints. Most probe-based solutions are limited to loopbacks or monitoring tests from a central location to a service endpoint. This ‘hub-and-spoke’ topology does not test the actual service path between customer locations, or between a customer and a remotely hosted data center, for example.

[5] *Bradner S., McQuaid J.*, March 1999. RFC2544. Benchmarking Methodology for Network Interconnect Devices. Retrieved from https://www.ietf.org/rfc/rfc2544.txt
[6] *ITU-T*. February 2016. Y.1564: Ethernet Service Activation Test Methodology. Retrieved from: https://www.itu.int/rec/T-REC-Y.1564/en

2. Test sets require trained technician dispatch to each service endpoint requiring service activation testing or troubleshooting, which is much less responsive and much more costly than a remotely initiated test using vCPE modules that are initially installed during service provisioning.

4.3. Service Assurance Functions

NFV-based vCPE solutions must be capable of all service assurance functions required to support the business services lifecycle, as described in Section 3 and shown in the Figure 6 displayed below. These include, but are not limited to:

1. Standards-based SAT supporting commonly employed IEEE RFC-2544 and ITU-T Y.1564 turn-up testing approaches.
2. Ethernet connectivity fault management (CFM), as defined by IEEE 802.1ag[7] to ensure service availability meets SLA definitions, and to measure continuity and latency using CCM and DMM/DMR messages, respectively.
3. Standards-based performance monitoring for Layer 2 (Ethernet) and Layer 3 (IP) services, typically implemented using ITU-T Y.1731[8]/ IEEE 802.3ah[9] Ethernet SOAM and RFC-5357 Two-Way Active Measurement Protocol [10] (TWAMP), respectively.
4. Bandwidth utilization monitoring, per port and per service flow (as defined by the MSO: VLAN, class of service/CoS, source or destination MAC or IP address, etc.) for usage-based billing, trending, and troubleshooting.

[7] *IEEE*. December 2007. 802.1ag - Connectivity Fault Management. Retrieved from: http://www.ieee802.org/1/pages/802.1ag.html

[8] *ITU-T*. August 2015. G.8013/Y.1731: OAM functions and mechanisms for Ethernet-based networks. Retrieved from: https://www.itu.int/rec/T-REC-Y.1731

[9] *IEEE*. September 2015. IEEE Standard for Ethernet. Retrieved from: http://standards.ieee.org/about/get/802/802.3.html

[10] *Yum, K. IETF*. October 2008. RFC-5357: A Two-Way Active Measurement Protocol (TWAMP) Retrieved from: https://tools.ietf.org/html/rfc5357

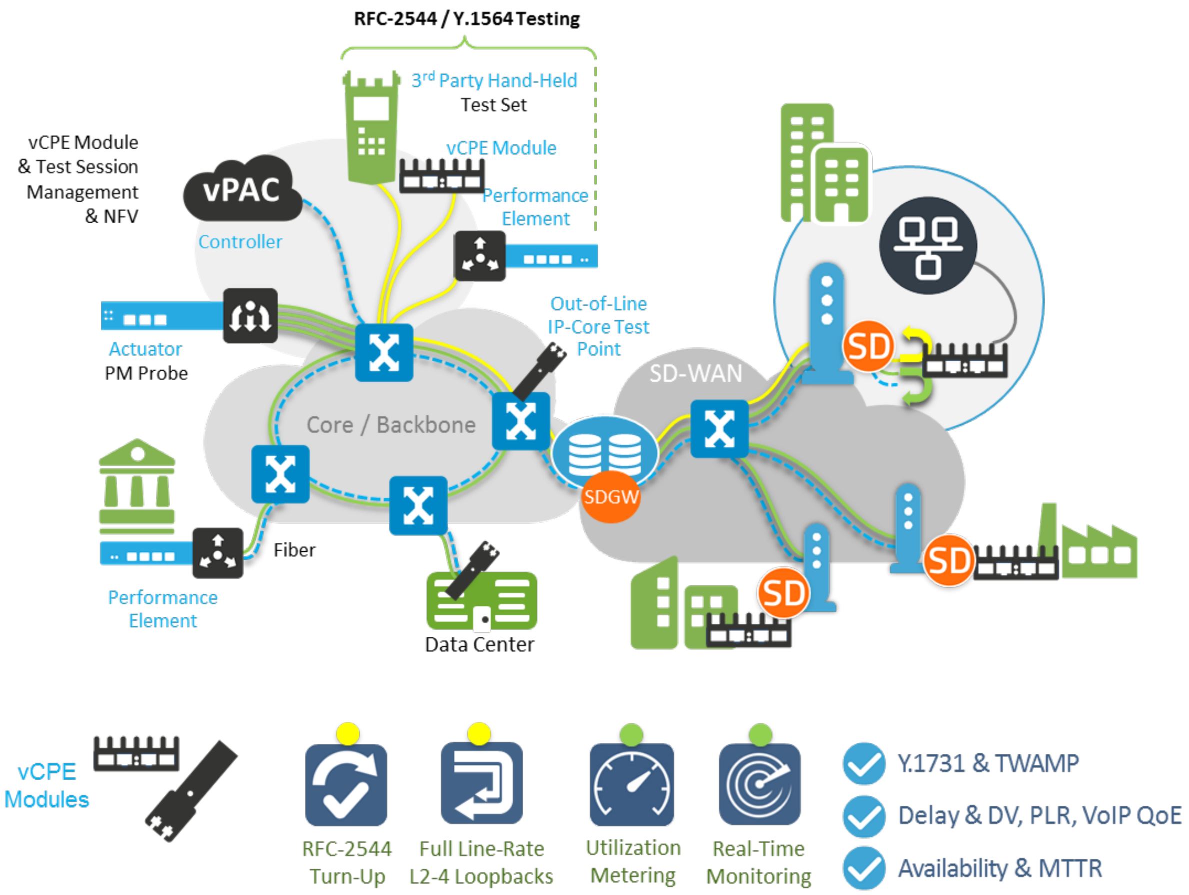

Figure 6 - Typical SD-WAN Service Assurance Functionality & Implementation

The full gamut of the Service Assurance solution can be seen in Figure 6. The provisioning and turn-up stage of the service lifecycle is delivered with the ability to generate and reflect Turn-up tests (yellow line) using traditional NIDs, Performance Elements, or hand-help test sets up to full line rate at any Ethernet supported packet size. The Performance Management stage of the service lifecycle is delivered through performance and SLA monitoring (green and blue lines) with micro-second accurate continuity and latency measurements and one second granularity bandwidth monitoring

4.4. Operations Integration Considerations

Implementation of an NFV-based solution should interwork with existing OSS to permit integration with existing management practices and procedures, and to make deployment of vCPE modules—as well as the monitoring and maintenance of the services they support—as operationally efficient as possible. Main areas to consider include:

- Deployment and management of the solution itself, to facilitate and automate element management of remote vCPE modules and SD-WAN service provisioning.
- Integration with SLA reporting platforms and fault management systems to harmonize monitoring, and reporting and within existing tools.

As introduction and general guidelines, the following opeational aspects should be considered during solution selection and deployment.

4.4.1. Low-Touch vCPE Module Provisioning

Ideally, vCPE modules should be ready to install in 'factory default' configuration, without requiring pre-staging by the MSO. To make that possible, the modules must be discoverable by an inventory system that can attribute the module to a particular customer site. This may be accomplished by relating the module to the MAC or IP address of the customer's SD-WAN appliance, for example.

To come under management control without requiring pre-staging or on-site configuration by a trained technician, the units require a method to 'discover' the management environment, have their managmeent IP address defined, have the desired configuration provisioned on the unit, have the module registered in the inventory and potentially have the service turn-up testing start automatically followed by service performance monitoring.

One commonly employed method to bring devices under managmeent involves using Dynamic Host Configuration Protocol (DHCP) with options 60 & 43 (refer to Figure 7 below):

1. When a vCPE module is connected to the network, it announces itself using a DHCP request with option 60, which communicates the module's identifier (device type) to the DHCP server.
2. When properly configured, the server assigns a dynamic IP address to the unit, and responds with option 43, providing the module with address of its "inventory node," responsible for managing the module.
3. Once under management control, a static IP can be assigned (if desired). A Fully Qualified Domain Name (FQDN) can also be assigned to the module. The FQDN must remain in sync with any link-state change (per RFC-2131), typically realized using automated DNS queries by the module inventory node.
4. Once the module is under management control, automation may be used to trigger an immediate or scheduled turn-up test to validate the service, then provision customer/SLA-specific monitoring sessions, etc. to allow customer-level self-install.

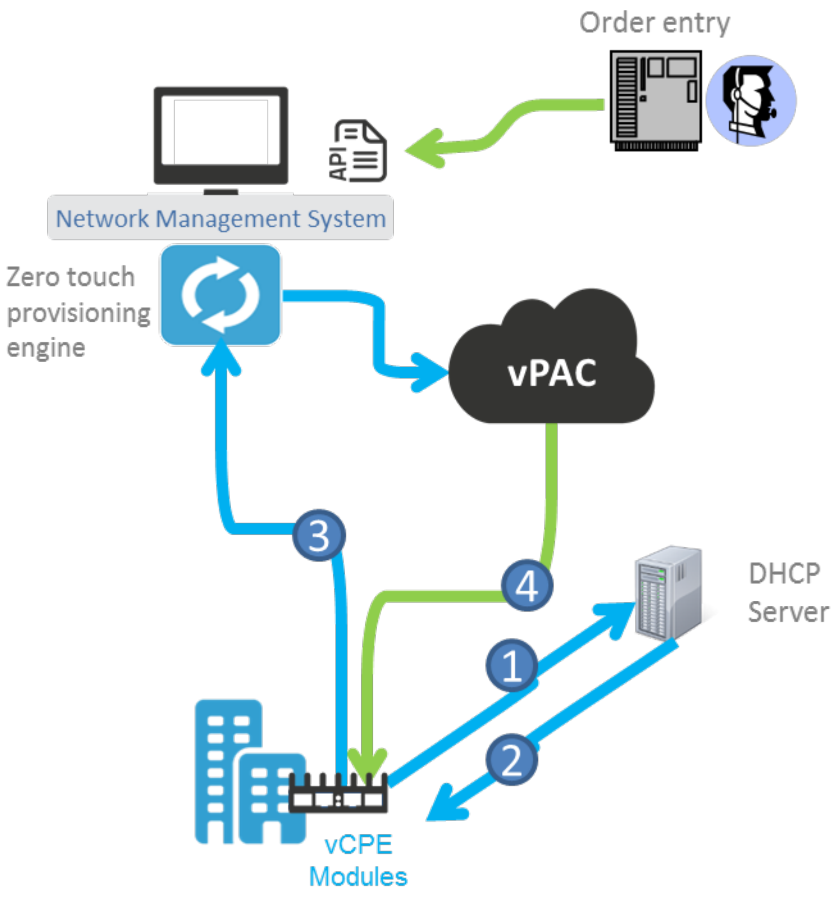

Figure 7 - Low-Touch vCPE Module Provisioning

4.4.2. Stand-alone Software Agents Provisioning

Software agent deployment is very simple and essentially consists of transferring a single executable binary file to the target system.

Software agents should be service-chained in such a way as to be accessible from the WAN link(s). Their configuration is typically very simple and straigtforward requiring only a limited number of options to be passed as arguments on the command-line during the agent instantiation. The agents can be configured to reflect performance monitoring probes only on specific system interfaces if desired.

The most basic type of software agent requires no separate management communications channel as it doesn't communicate management information, but merely act as responder to performance monitoring probes sent from a vPAC-type node.

4.5. vCPE Module Management Communication Options

The vCPE module should be capable of adapting to an MSO's particular management and device addressing methodology. This requires the unit to distinguish customer/test traffic from management communication. A variety of methods are commonly used, all implying that the vCPE module must offer support for each of these schemes:

- Layer 2 addressing: separate management and customer MAC addresses. In this case, the vCPE module must support two MAC addresses, one for management traffic, and another for customer, test, CFM, SOAM, and active performance monitoring traffic.
- Layer 2 addressing: separate management and customer VLANs. A single MAC address is used in combination with Q-in-Q VLAN support (C/S tagging, IEEE 802.1ad).
- Layer 3 addressing: similar to the Layer 2 scheme described above; separate management and customer traffic IP addresses may need to be supported by the vCPE module, depending on the method used by the operator. VLAN support, including Q-in-Q (S-VLAN), may also be required to support this scenario.
- Layer 3, transparent IP addressing: an operator may elect not to assign new IP address(s) to the vCPE module, instead using the address of a device located 'behind' the module. This is practical when the operator has a known device at the customer premises (e.g. an SD-WAN appliance, a set-top box, Wi-Fi controller, security gateway, etc.) In this case the vCPE module detects management traffic using a combination of this other device's IP address in combination with other identifiers (such as a management VLAN tag).

5. Industry Proven Methods to Extend Standards-Based Test, Measurement and OAM to Virtualized Environments for SD-WAN and x86 Infrastructure

5.1. SD-WAN Traffic Routing and Performance Monitoring

vCPE Modules are capable of full line-rate test traffic generation, able to create and analyze up to four Layer 2 or Layer 3 unique flows, and run a fully-fledged RFC-2544 or 8-flow Y.1564 SAT suite toward other vCPE modules or third-party endpoints.

Each vPAC is capable of generating up to 4000 performance monitoring flows toward other vPACs, other vCPE modules, or third party-endpoints.

When paired together, the vCPE modules and vPAC will allow service providers to test the multiple service paths at turn-up and re-validate on-demand the capacity of a specific path or service during maintenance windows for troubleshooting.

To direct performance monitoring flows, the SD-WAN controller pushes policies throughout the network, stating that network traffic complying with defined profiles is kept on chosen paths. Then the vCPE modules or the vPAC generate performance monitoring flows that comply with the traffic profile, giving the operator the ability to proactively monitor all SD-WAN paths concurrently and quickly take action whenever a fault is detected.

5.2. SD-WAN Deployment Models and Service Assurance Instrumentation

From a network instrumentation point-of-view, providers can either use a blanket approach to simplify and unify their deployments or use a right-size approach to customize and adapt the instrumentation to the site and its significance in the overall network. The selected service assurance solution should have interoperable solutions that scale from supporting built-in integrated third-party reflectors to fully featured dedicated devices.

As shown in Figure 8, the ability to use a diversity of standards-based service assurance endpoints unlocks the ability to deploy the best tool for the job at each location. Some locations, like aggregations and cores sites, will require the enhanced features provided by a traditional Performance Element NID. Many Enterprise and SMB customers will benefit from the essential features and ease of use offered by vCPE modules, and for those locations where an x86 platform is readily available, a low-touch light-weight software agent can be deployed.

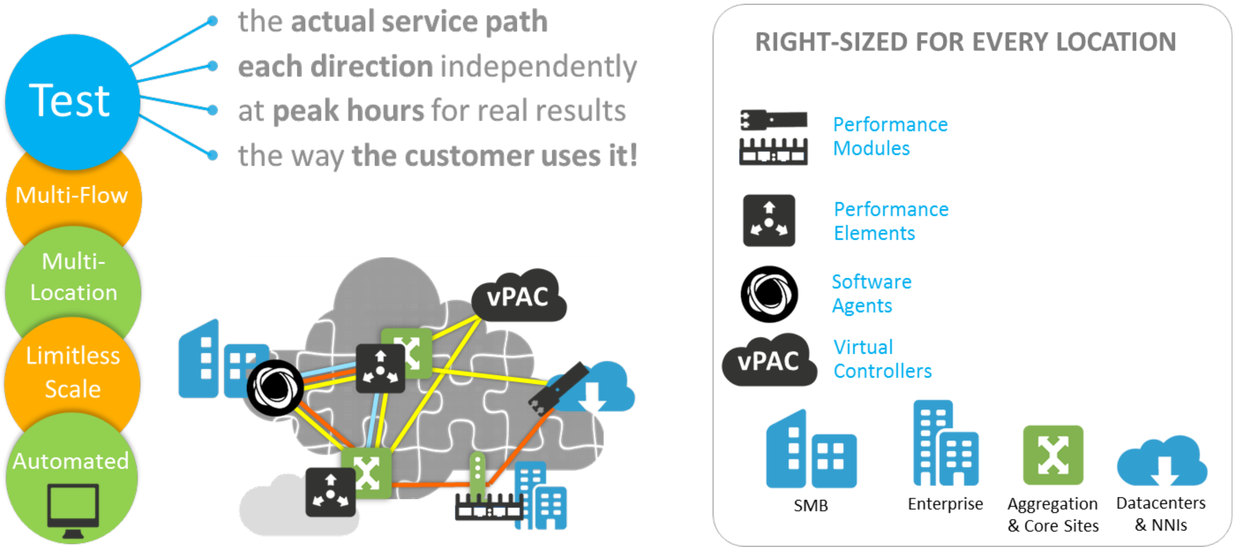

Figure 8 - Performance Monitoring Through Multiple Paths to a Diversity of Endpoints

5.2.1. Centralized Gateway SD-WAN Models

Many SD-WAN vendors offer an architecture based on centralized gateways to act as the virtual hub for any number of remote locations (spokes) as displayed in Figure 9. The connected sites (Branch, Head Office, HQ) need little hardware and a number of network transports (internet links or traditional WAN links) to establish the overlay network needed for the SD-WAN to operate illustrated using the gray lines to the SD-WAN Gateway. The overlay network is built by having each remote site establish encrypted tunnels to the SD-WAN gateway over each provisioned path.

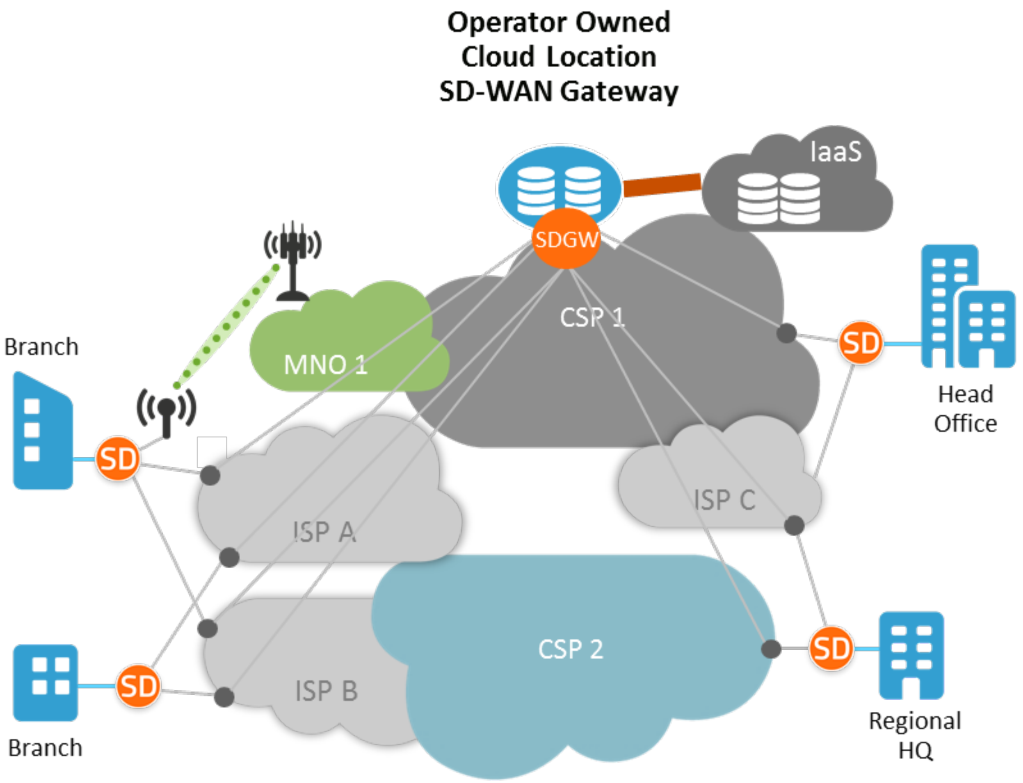

Figure 9 - Centralized Gateway SD-WAN Model

5.2.1.1. With vCPE Modules

In such a model, it is recommended to co-locate a vPAC along with the SD-WAN gateway(s) as compute resources are typically plentiful in the cloud.

When deployed in this manner, the vPAC can be used in large scale hub-spoke and full-mesh topologies to perform active, micro-second accurate, standards-based performance monitoring towards thousands of endpoints continuously.

To ensure the most flexible and featureful performance monitoring solution, the vPAC is supplemented by a vCPE module at each connected site. The remote vCPE modules effectively become remote ports of the centralized vPAC and therefore this deployment model ensures that the full performance monitoring feature-suite is available for the operator. This also guarantees that the solution will evolve along with the network as the functionality is delivered through NFV capabilities to each remote endpoint.

Further, as displayed in Figure 10, having vCPE modules at each site enables NFV performance monitoring (NFV-PM) and remote SAT generation. Gaining the ability to source the performance monitoring traffic flows from any location and targeting any location ensure that the results will match the user experience 1-to-1.

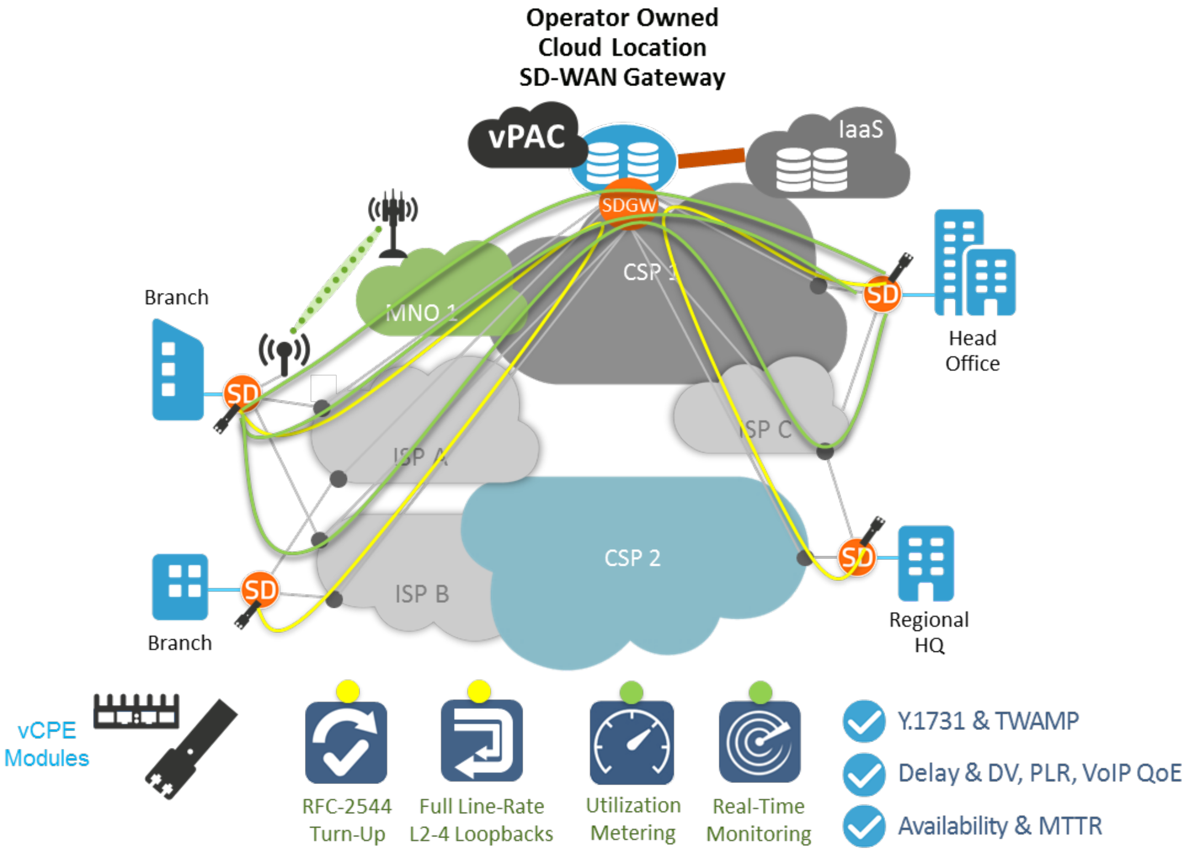

Figure 10 - Centralized Gateway SD-WAN Model —SAT and Performance Monitoring with vCPE Modules

5.2.1.1. With Software Agents

For locations where the added benefits of vCPE modules, NFV-PM, and remote SAT generation/reflection are not required, a lightweight software agent can be embedded directly into the SD-WAN software stack on-site. The agent can run on bare-metal, in unprivileged mode on a Linux OS, or inside a dedicated container within the NFV infrastructure (NFVi) supporting the SD-WAN deployment.

The benefits of this approach are two-fold: 1) the simplicity of deployment is second-to-none as only a simple software program need to be transferred and started; 2) the site running the software agent can now be the target of highly precise and granular probes, supporting 1-way metrics and return standard-based metrics to the vPAC that will process the results, analyze and monitor the network quality, and report the network state northbound to NMS, analytics, and big data systems.

When software agents are installed on x86 CPE, the overall performance strategy is slightly changed. As shown in Figure 11, the performance monitoring flows are no longer endpoint-to-endpoint (spoke to spoke), but instead generated at the central cloud location (hub) and reflected at the spoke. This falls in line with the visibility provided by the SD-WAN gateway where each path is measured independently. While the model covers the complete network, it does not completely reflect the end-user experience as accurately as the vCPE modules strategy explained previously can.

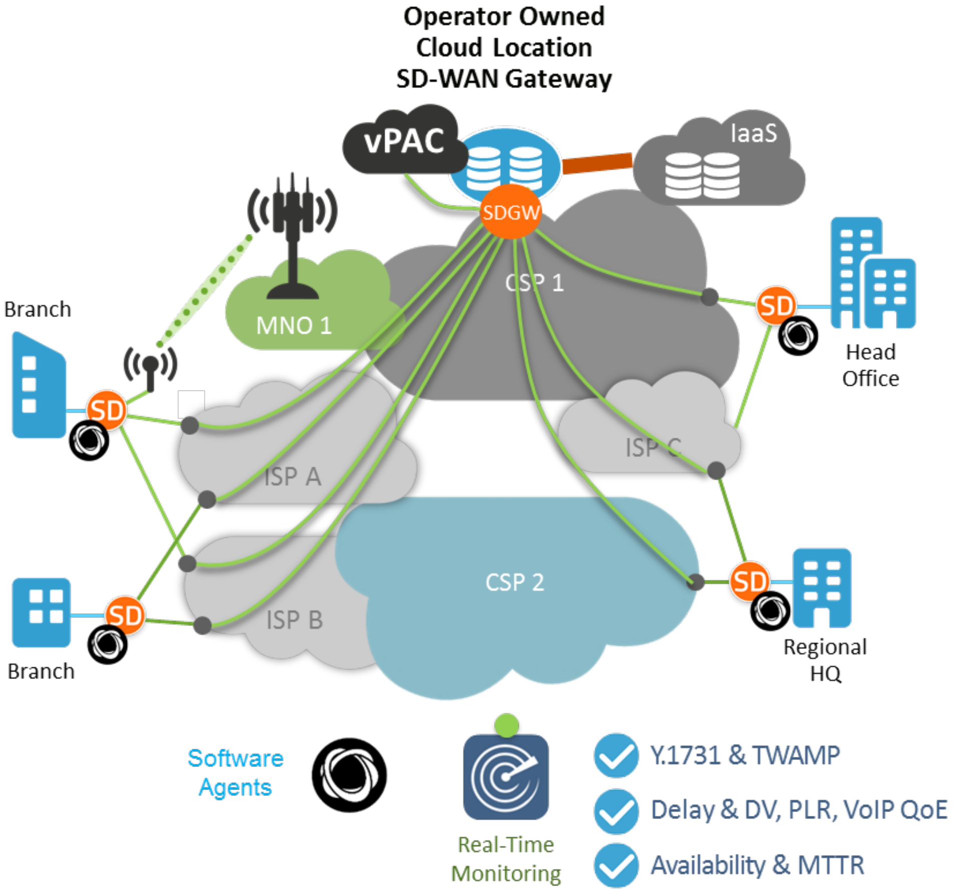

Figure 11 - Hub-to-Site Performance Monitoring from vPAC to Software Agents

5.2.2. Distributed SD-WAN Models

For SD-WAN architectures where each site can directly connect to each other site dynamically and on demand, the need for a centralized SD-WAN gateway is removed. This approach has slightly different pros and cons because the SD-WAN intelligence is distributed into the end-points and the provider can use separate controller (or director) software to centrally configure and control each of the endpoints. Removing the centralized gateway can improve overall resiliency and reduce the required cloud compute resources, but it also removes a centralized location where additional network services (such as the vPAC) could easily be hosted. As shown in Figure 12, in this SD-WAN model, each site is able to establish overlay tunnels over each potential path to any other site dynamically and as a response to underlying network conditions.

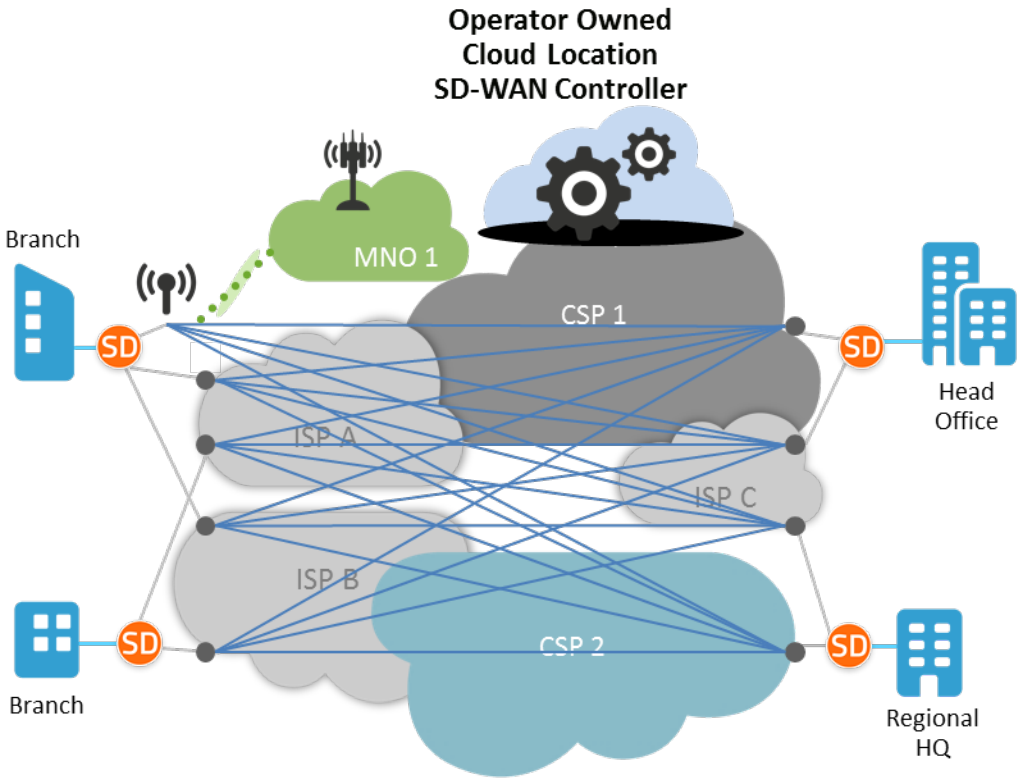

Figure 12 - Distributed SD-WAN Model

5.2.2.1. With vCPE Modules

When operating in a distributed SD-WAN architecture, the vPAC can be hosted in a head office or data center location as it requires limited compute resources. In the rare case where no compute resources are available inside the existing network, the vPAC can be hosted in a data center location and then connected to the SD-WAN as a synthetic site.

Distributed SD-WAN architectures are best served when using controller VNFs and NFV-powered hardware modules because this deployment model makes it possible to easily test and monitor each path directly and individually as shown in Figure 13. Having the ability to retrieve all this performance data enables the operations team to easily pinpoint issues and bottlenecks in situations where traffic from one location might hamper the performance when communicating to another site. Given the exponential number of potential paths and the dynamic nature of the traffic, being able to monitor the whole network with a single solution, and easily export the performance data to an open data repository, proves to be instrumental to supporting this new type of environment.

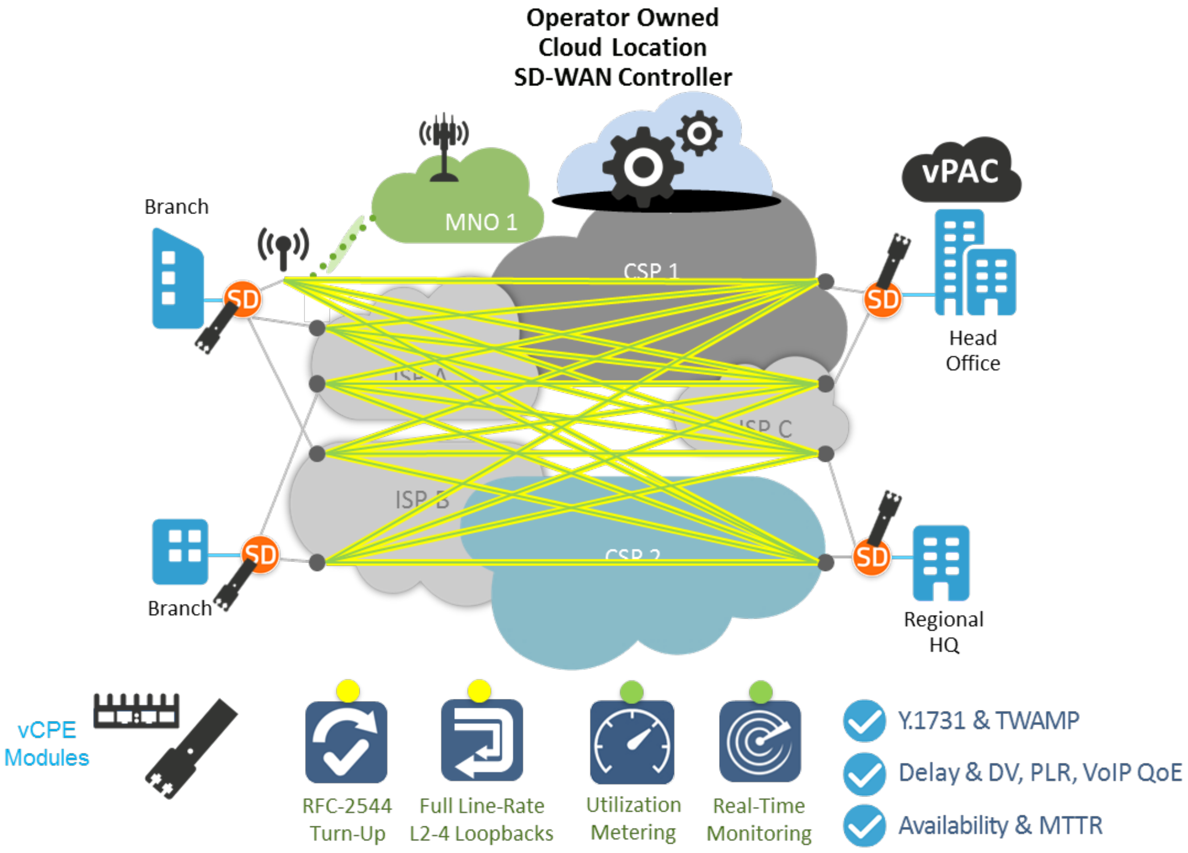

Figure 13 - Distributed SD-WAN Model — SAT and Performance Monitoring with vCPE Modules

5.2.2.2. With Software Agents

When operating in a distributed SD-WAN architecture and using software agents or third-party reflectors, the vPAC should be hosted in locations of greater interest as it will be used to generate performance monitoring traffic flows. When many sites have significant operational importance, each one of these sites should run the vPAC VNF in order to be used as a probe generator.

Because the vPAC is used as the flow generator and the software agent is a reflector, the result of combining a distributed SD-WAN architecture with a 100% software performance monitoring solution results in star-shaped measurements to represent a mesh-like network. Depending on the end-user's use-case and typical network usage, the limited monitoring could be less than ideal and therefore it is recommended to supplement this model with cost effective vCPE modules in key locations to offer the coverage needed for best results.

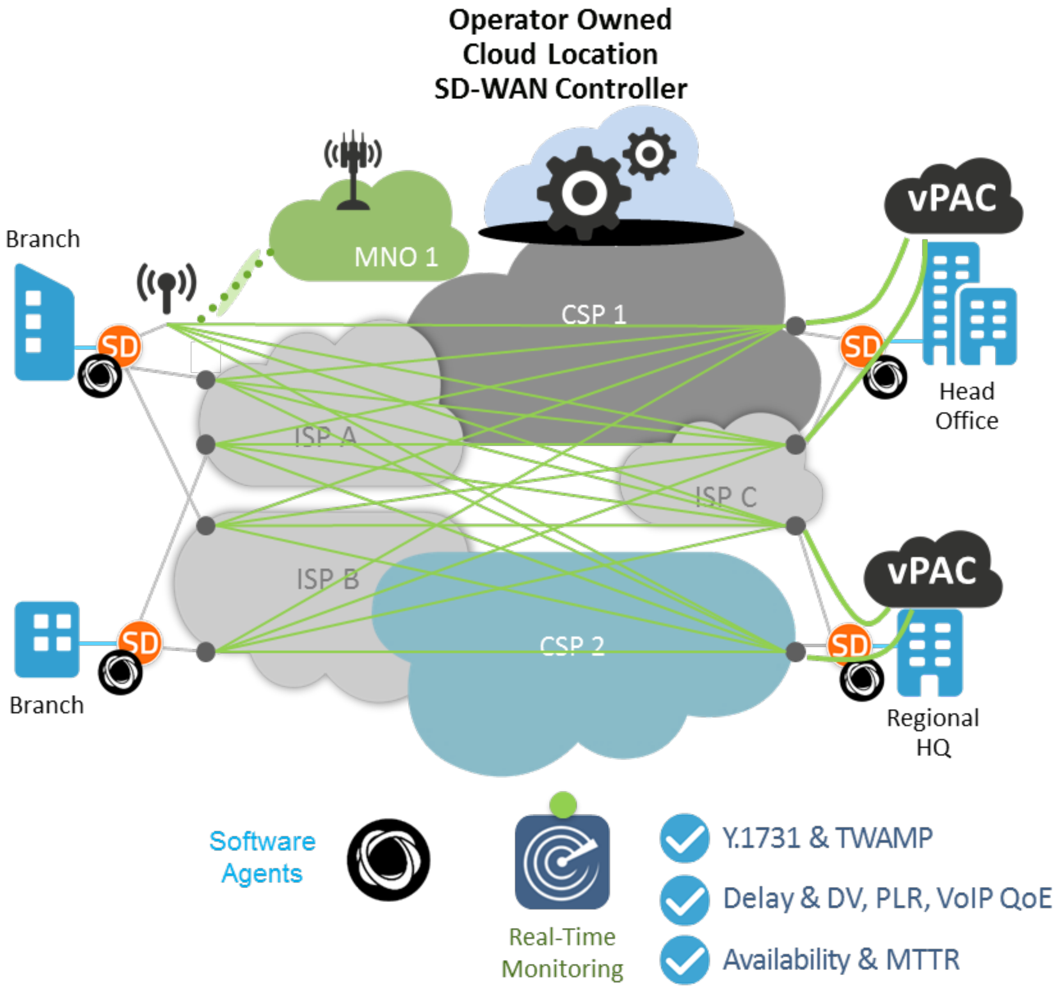

Figure 14 - vPAC to Software Agent Performance Monitoring

6. Methods to Centralize Performance Monitoring Data into Existing Reporting and Fault-Management Systems.

The virtualized NFV-based active and passive monitoring solutions typically provide both fault management (FM) and performance management (PM) metrics and data into an operator’s existing FM, PM, and OSS systems. FM data is traditionally provided as threshold-based simple network management protocol (SNMP) traps or extensible markup language (XML) type events, whereas PM data often uses comma separated values (CSV) files or XML/JSON encoded data streams or file transfer.

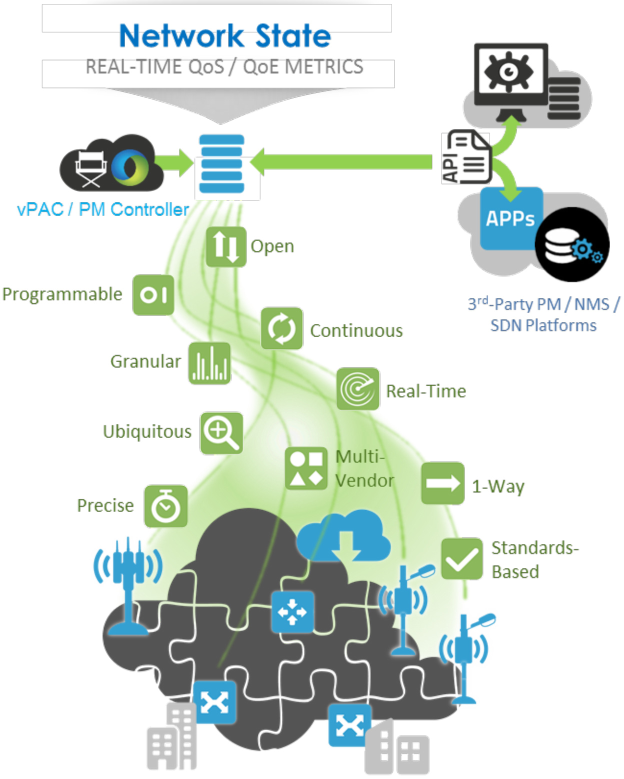

The metrics data provided to the 3rd party system should fulfil as many as possible of the 10 metrics fundamentals:

1. Open – well documented
2. Continuous 24/7
3. 1-way – path separation
4. Granular – seconds not minutes
5. Programmable – configurable
6. Real-Time – immediately available
7. Ubiquitous – as many sources as possible
8. Multi-vendor – no lock-in
9. Standards based
10. Precise – microsecond where applicable

Figure 15 - Export and Centralize Network State to Northbound Systems

The metrics reported northbound should be tangible, near real-time, and granular. For a modern SD-WAN type deployment the historically de-facto standard of using SNMP-polling every 5 or 15 minutes is far from sufficient to capture the fast transients that occur in today’s networks. Active (TWAMP / Y.1731 / UDP) type tests should be able to report at least every 30 seconds, possibly down to every 5 seconds, and use a sampling rate of at least 10 probe packets per second to properly capture any short-lived events. The highly granular data may be aggregated in retrospect, to reduce the amount of metrics stored for long-term reports. It is important, though, that such aggregation does not result in loss of the measured extremes. Use of percentiles (95th, 98th, 99th, etc.) helps preserve a view of the distribution of delay or jitter values, while metrics such as loss burstiness aid in determining length of service interruptions, down to millisecond level with a sufficiently high sample rate.

Ideally, the northbound metrics transport should be streaming based (XML / JSON) for larger installations due to the efficiency of such machine-to-machine (M2M) methods. For smaller, enterprise-type applications, a CSV-based file transfer may be sufficient. The monitoring solution providing the metrics must have methods to filter out any type of monitoring data not desired, and be able to hold the real-time granular data for a limited period of time to enable detailed review of the measurement results for troubleshooting or incident reports. The northbound consumer of the data should keep records for at least one year, possibly rolled up to daily aggregates to reduce data storage requirements.

Reports and dashboards for the monitoring data should be created to suit the SD-WAN applications offered. This may include management-level geographical overview charts, as well as detailed regional views for network operations center (NOC) troubleshooting audiences.

7. Approaches to Optimizing SD-WAN Performance Using End-to-End Monitoring Data

For SDN solutions in general, and SD-WAN in particular, using active and passive end-to-end or hop-by-hop monitoring metrics in a feedback loop towards the SDN/SD-WAN controller provides a means to automatically assist the traffic and service control functions with external quality metrics.

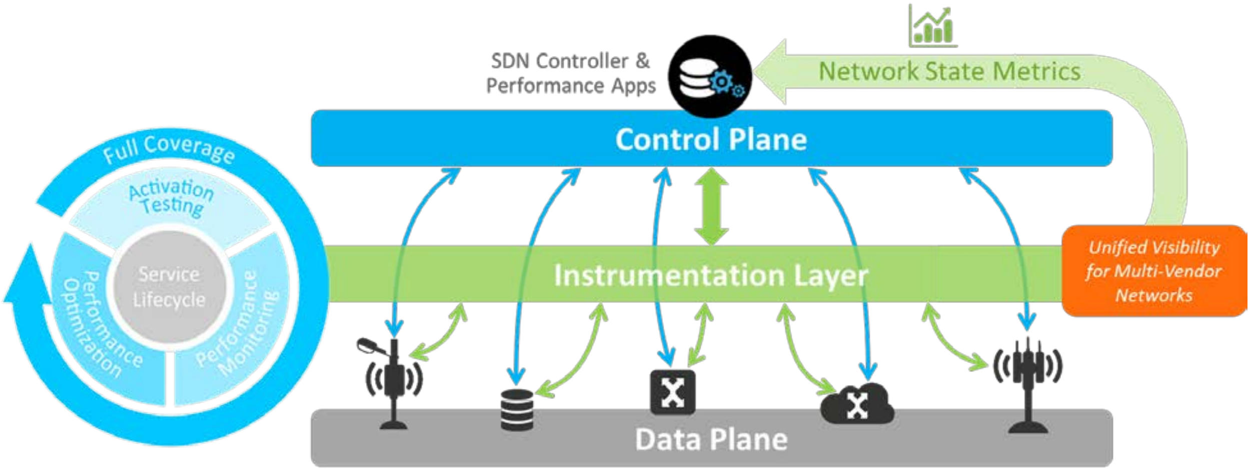

Figure 16 - Unified and Complete Instrumentation Layer Unlocks Value for the Control Plane

Providing the SDN controller with direct network quality metrics allows for immediate feedback on the impact of reroute decisions or path changes, and ultimately allows the controller to assess—via machine learning methodologies—the impact of each decision before it is taken. Automation of the full service lifecycle is thus possible.

This monitoring is not limited to active Layer 2-3 performance data (delay, jitter, loss) but could also be augmented with software- based agent type monitoring of a specific service type or path as well as hardware-based, highly granular utilization type metrics that can alert on links or paths experiencing high frequency of micro-bursts or micro outages.

Conclusion

SD-WAN for enterprises is becoming more relevant all the time, driven by growing use of cloud-based applications. The advantages of SD-WAN—regardless of whether it is deployed using a hybrid model or a full-scale implementation—cannot be ignored. And, as widespread adoption of this technology continues, operators must be able to deliver the same level of quality with their SD-WAN managed services as they do with traditional WAN offerings.

With that goal in mind, all SD-WAN lifecycle phases can benefit from a flexible, NFV-based performance monitoring solution that scales beyond the footprint of the SD-WAN cloud and can send performance flows from any starting location to any destination in the network infrastructure. Such a solution can be used to:

- Cover large scale hub-spoke and full-mesh topologies with active, micro-second accurate, standards-based performance monitoring towards thousands of endpoints continuously.
- Bring standards-based turn-up testing, monitoring, and OAM functions to all SD-WAN endpoints, by adding NFV-enabled vCPE modules or orchestratable lightweight software agents. Since the solution is standards-based, standard networking devices can also act as responders to performance monitoring flows.
- Monitor micro-outages, one-way delay and variation, and SLA compliance by delivering precise and granular metrics.
- Centralize test control and automation, integrated with existing OSS, by pairing vPACs with NMS solutions.
- Deliver a new level of PM workflow automation with results centrally stored for comparison to predefined QoS templates or SLA levels. Tests—conducted one-way or bi-directionally, in an end-to-end or segmented manner—can be scheduled on demand or triggered by service endpoint installation.
- Provide open access to turn-up data and results—including customer-ready reports reflecting their specific SLAs—using the API.

Abbreviations

API	application programming interface
CCM	continuity check message
CoS	class of service
COTS	commercial off-the-shelf
CPE	customer premises equipment
CSV	comma separated values
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data Over Cable Service Interface Specification
DMM/DMR	delay measurement message / delay measurement response
DPI	deep packet inspection
FM	fault management
FQDN	fully qualified domain name
ICMP	Internet Control Message Protocol
JSON	JavaScript object notation
MEF	metro ethernet forum
M2M	machine-to-machine
MSO	multiple systems operator
NFV	network functions virtualization
NFV-PM	network functions virtualization performance monitoring
NFVi	network functions virtualization infrastructure
NID	network interface device
NMS	network management system
NOC	network operations center
OAM	operations and maintenance
OTT	over the top
OSS	operational support systems
PM	performance management
QoE	quality of experience
QoS	quality of service
SAT	service activation test
SaaS	software as-a-service
SD-WAN	software-defined WAN
SDN	software-defined networking
SLA	service level agreement
SNMP	Simple Network Management Protocol
SOAM	service OAM
TCP	transmission connection protocol
UDP	user datagram protocol
vCPE	virtualized customer premises equipment
VLAN	virtual local area network
VNF	virtual network function
vPAC	virtualized performance assurance controller

WAN	wide area network
XML	extensible markup language

Bibliography & References

MEF Forum. July 2014. CE 2.0 Service Management Life Cycle White Paper. Retrieved from: http://mef.net/Assets/White_Papers/CE_2_0-Service_Life_Cycle_White_Paper.pdf

ITU-T. August 2015. G.8013/Y.1731: OAM functions and mechanisms for Ethernet-based networks. Retrieved from: https://www.itu.int/rec/T-REC-Y.1731

MEF Forum. January 2012. Implementation Agreement MEF 23.1: Carrier Ethernet Class of Service – Phase 2. Retrieved from: http://dev.mef.net/Assets/Technical_Specifications/PDF/MEF_23.1.pdf

ITU-T. February 2016. Y.1564: Ethernet Service Activation Test Methodology. Retrieved from: https://www.itu.int/rec/T-REC-Y.1564/en

Yum, K. IETF. October 2008. RFC-5357: A Two-Way Active Measurement Protocol (TWAMP) Retrieved from: https://tools.ietf.org/html/rfc5357

Breznick, A. Heavy Reading, January 2017. How cable can conquer the enterprise market. Retrieved from https://accedian.com/wp-content/uploads/2017/01/HR_Accedian_Cable_Enterprise_WP_1-24-17.pdf

Sterling, P. Heavy Reading. February 2017. Operator Success in the New Age of the Software-Defined WAN, Retrieved from https://resources.ext.nokia.com/asset/201132

MEF Forum. April 2012. Introducing the Specifications of the MEF. MEF 38: Service OAM Fault Management YANG Modules Technical Specification. Retrieved from: http://slideplayer.com/slide/5687304/

MEF Forum. March 2016. Service Operations Specification MEF 55: Lifecycle Service Orchestration (LSO): Reference Architecture and Framework. Retrieved from: http://dev.mef.net/Assets/Technical_Specifications/PDF/MEF_55.pdf

Bradner S., McQuaid J., March 1999. RFC2544. Benchmarking Methodology for Network Interconnect Devices. Retrieved from https://www.ietf.org/rfc/rfc2544.txt

IEEE. December 2007. 802.1ag - Connectivity Fault Management. Retrieved from: http://www.ieee802.org/1/pages/802.1ag.html