Moving SDN Forward:

Deploy with Confidence

SDN Introduction

Software Defined Networking (SDN) is an architectural revolution that enables application controlled programming and management of network resources in a dynamic and scalable manner. According to 2015 forecasts by *Infonetics Research, SDN is being evaluated and trialed by Cloud Providers (CSPs) and Service Providers (SPs) today with significant deployments expected to start in 2016 and growing to over $18.8B addressable market by the end of 2019. The strong motivation for deploying SDN is driven by many factors:

- The compelling economics of building programmable networks using commodity hardware.
- The agility to optimize network performance potentially offers new services and revenue opportunities.
- Simplify network management and complexity to facilitate the transition toward 'cloudification' of IT.

Cloud and Service Providers face a challenging road ahead of them before achieving the desired end goal—a programmable network that runs on commodity hardware. Today's provider networks are a collection of proprietary, purpose-built routers and switches that are expensive and are at different stages in their depreciation cycle. These traditional elements are all-in-a-box devices that do everything: switching, routing, access control or firewalling, policy enforcement, deep packet inspections (DPI), traffic engineering, service deployment, flow tracking and management. The complexity and pervasiveness of these devices make it very challenging to manage or upgrade existing network infrastructure. In turn, this leads to major delays in the introduction of new revenue-generating services.

Promise of SDN

SDN helps address lack of programmability and vendor lock-in by introducing an intuitive 3-tier architecture. The architecture enables providers to source from different hardware and software vendors and the pace of innovation in one domain is not slowed by the product release schedule in another domain.

SDN Promises	
Programmability	SDN defines standardized (API protocol) to communicate with infrastructure
Improved agility	Flexible, elastic networks to react quickly to change
Network visibility	End-to-end visibility for optimal traffic engineering
Innovative services	Enables offerings of innovative services, dynamically, based on user demands
Vendor neutrality	Promotes multi-vendor and open source solutions based on user collaboration

** Infonetics Research: Total SDN Hardware and Software Worldwide and Regional Market Size and Forecast.*

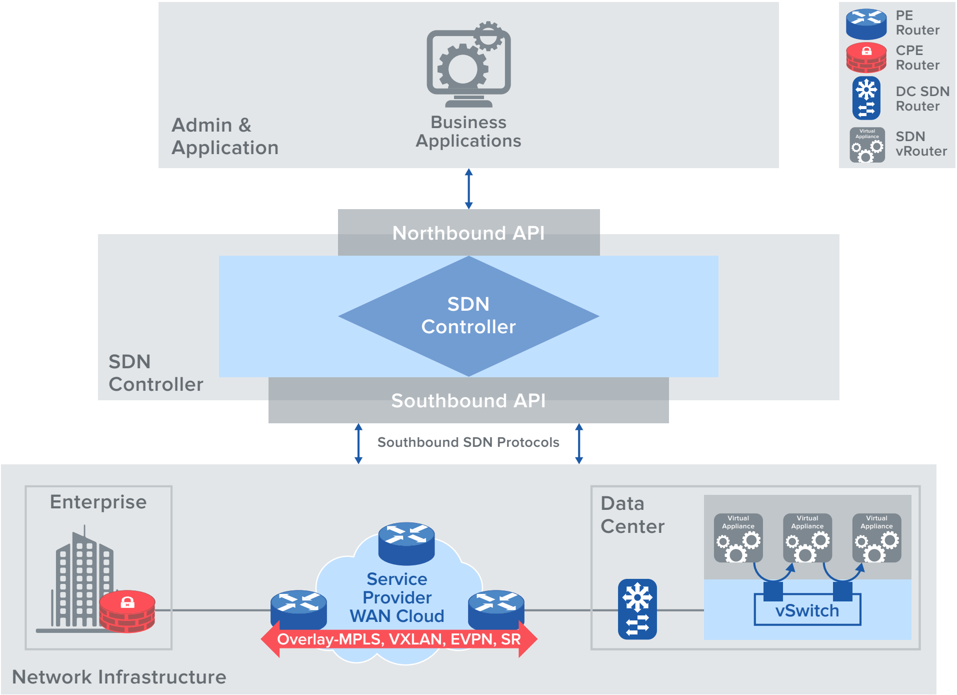

Figure 1: The 3-tier SDN architecture

At a high level, this 3-tier architecture offers:

- Programmatic network control through the use of SDN controllers
- A forwarding infrastructure that is separated from the control functions and leverages commodity hardware, including SDN-compatible routers, switches, and other devices
- Administrative or business applications become the service consumers

The following sub-sections describe the SDN architectural components.

SDN Controller

SDN controller is the most significant component of the SDN architecture, providing centralized control and management. It communicates with network infrastructure nodes using southbound protocols, with application layer using northbound API, and with peer SDN controllers using standardized control protocols. SDN controllers collaborate across domains to attain a global network view that helps make optimal routing and traffic steering decisions. Much of the benefits of SDN are derived from the ability of the controller to communicate with other components in an open and standardized manner.

SDN controllers can perform many network functions, including switching, routing, traffic engineering and path computation, network automation, and policy enforcement. On the southbound interface, controllers communicate with the infrastructure using SDN protocols such as OpenFlow, path computation element protocol (PCE-P), BGP-Flowspec, BGP LinkState (BGP LS), and NetConf.

SDN Infrastructure

SDN infrastructure includes all the programmable physical and virtual forwarding nodes that participate in packet forwarding associated with SDN flows. Examples of nodes in SDN infrastructure are routers, vRouters, switches, vSwitches, and transport nodes.

SDN infrastructure can deployed using one of two models:

- overlay networking
- hop-by-hop SDN forwarding

Operators deploy SDN overlay networks using tunneled encapsulations such as VXLAN, MPLS, GRE, segment routing or EVPN, which allow them to deploy service oriented networks on the fly. In SDN overlay network, only the edge devices (virtual or physical) need to provide support for SDN protocols and APIs. Such network architecture enables operators to expand and contract on demand with minimal disruption to underlying non-SDN nodes. With the hop-by-hop SDN forwarding model, every node in the forwarding path is SDN-aware and programmable. Building the hop-by-hop model is more expensive and disruptive.

The figure shown below expands on the 3-tier architecture model and highlights various southbound and overlay protocols that are used in SDN deployments.

SDN Application Layer

SDN simplifies network complexity and management by providing centralized end-to-end visibility of the network. Business and operator applications can push new service requests and new policies on the fly to SDN controller using northbound APIs (RESTful, Thrift, etc.). Bandwidth calendaring applications can auto-program the network capacity for high bandwidth or low-latency applications at certain times of the day. Capacity planning applications can query the centralized traffic engineering and network management controllers to gauge overall network utilization.

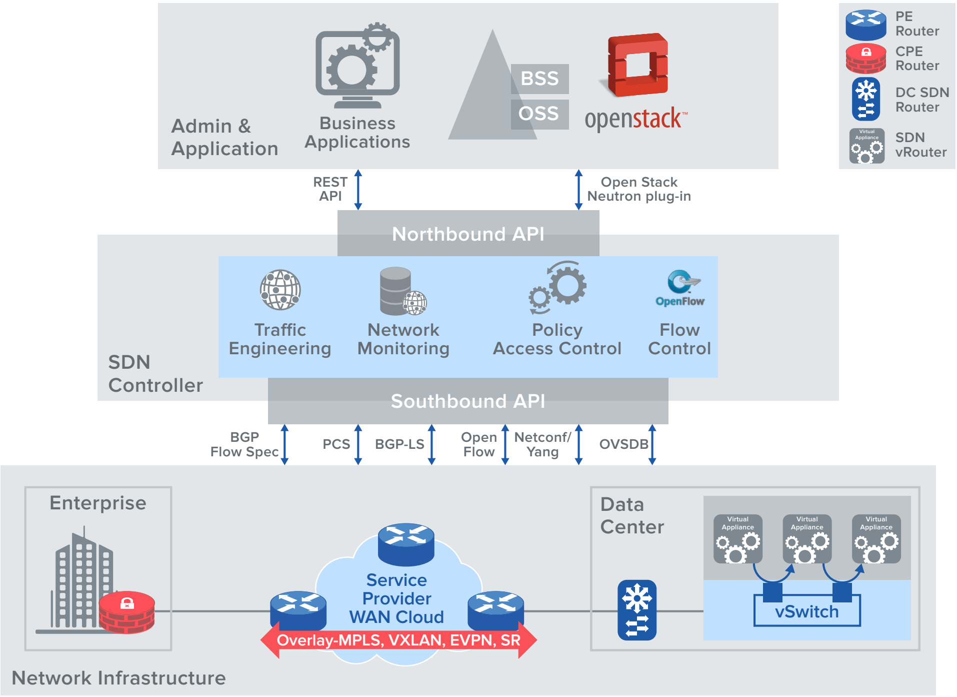

Figure 2: SDN southbound and overlay protocols

Testing Software Defined Networks

Network administrators are drawn to SDN's programmability, agility and manageability. Concerns about network scalability, reliability, and security, however, are slowing down real-world deployments. The key to addressing these concerns is thorough pre-deployment validation of SDN hardware and software. SDN testing can be broken down into testing each of its three main aspects: SDN controller, SDN infrastructure and overlay networks, and end-to-end services.

Testing the SDN controller

SDN decomposes the conventional router into multiple reduced function applications and distributes them across multiple nodes. New SDN protocols have been introduced for communication between SDN controllers and underlying infrastructure.

Compute-intensive tasks such as routing, end-to-end path computation, and flow management are moved from the forwarding devices to the SDN controller. It is imperative to verify the performance, reliability, and stability of the controller and relevant infrastructure under high load conditions. Spirent recommends that SDN controllers be validated along multiple dimensions as described below:

Functional Tests

Verify:

- The interface and API for each SDN protocol, including OpenFlow, PCE-P, BGP Flowspec, and BGP link state.
- Interoperability in multi-vendor environments with branded vendor equipment, open source controllers such as ODL and ONOS white box switches, and open source controllers such as ODL and ONOS.

High Availability Tests

Verify:

- The SDN controller is a single point of failure, so controller redundancy is a must, therefore, verify how long it takes for a back-up controller to take over when SDN controller becomes unavailable.
- State and database synchronization between primary and backup controllers.
- Whether the underlying infrastructure can load share between multiple controllers and synchronize state between controllers.

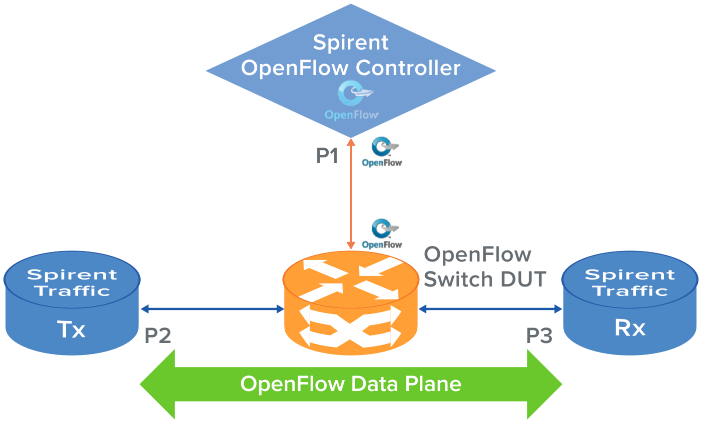

FLOW TABLE Match	Action
In Port 2, dst mac=00:00:01	Out Port 3
In Port 2, dmac=00:00:01, dIP=0.0.0.1	Out Port 3
In Port 2, dmac=00:00:01, vlan=1, dIP=0.0.0.1	Out Port 3
In Port 2, dmac=00:00:01, vlan=1, dIP=0.0.0.1, udp=1	Out Port 3

Figure 3: OpenFlow switch forwarding

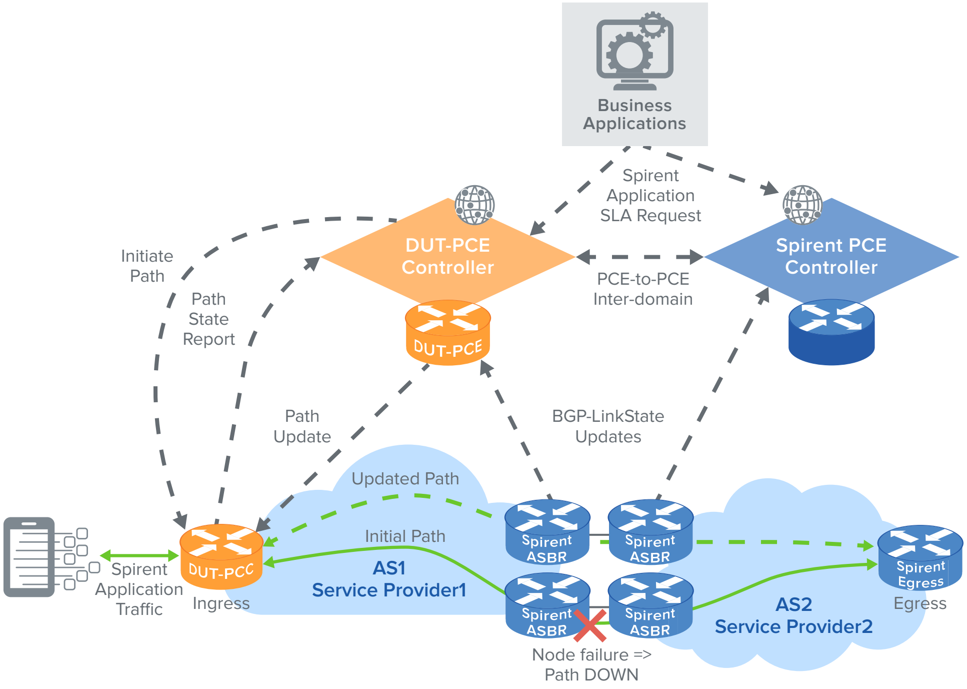

Figure 4: Validating the PCE controller and client

Performance Tests

Measure:

- Response time of the controller when new unknown flows or traffic engineering (TE) paths are requested.
- How long it takes for the client router or switch to program the flow or set up the TE paths.
- For Openflow, the flow table setup rate from controller to switch with data forwarding for a large number of flows.
- For traffic engineering use cases, the end-to-end TE path computation rate for thousands of applications.
- Path setup rate with traffic for thousands of applications on a PCC client.
- When TE paths span multiple routing domains, the PCE controllers in each domain compute paths for their respective domain.
- With ingress PCE, the controller requests path information from next egress PCE controller.
- For inter-domain flows, verify end-to-end path computation and setup time.

Scalability Tests

Determine:

- Centralized controllers must be able to initiate, update, and synchronize millions of paths from hundreds of clients both within the domain and inter-domain.
- Underlying infrastructure devices (switches or ingress routers) must be able to setup, report state, and synchronize state for thousands of paths and flows.

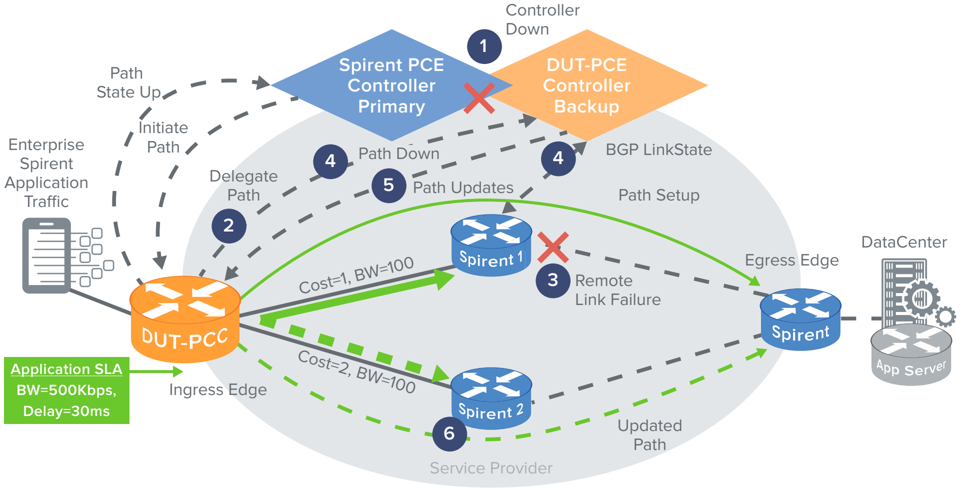

PCC LSP Scale
- LSP Setup Rate
- Time taken from Update/Initiate to LSP UP and traffic being forwarded on LSP

PCE High Availability
- Primary PCE rejects some LSPs, PCC forwards to backup PCE
- Primary PCE does not respond, PCC delegates control to Backup PCE
- Primary PCE recovers, PCC

Figure 5: Validating redundancy and failover of PCE controllers

Reliability Tests

Verify:

What happens if a link goes down before it affects thousands of paths since network failures are inevitable. In order to maintain carrier and enterprise SLAs, the network must be resilient and recover quickly after failures.

How long does it will take to re-calculate an optimized path, once the SDN controller is notified that a path is down or temporarily recovered.

How long it will take for the network to re-converge on that new path.

Security Tests

Validate:

If the SDN controller is compromised, control could be lost for the entire network. This is far worse than the effect of a single router being compromised. All communication to the SDN controller from the underlying infrastructure (southbound) and from application side (northbound) must be secured via authentication and encryption.

Controllers must be able to handle distributed denial of service (DDoS) attacks, where there is a flood of TCP connection requests from malicious clients or applications. DDoS attacks could slow down the controller's response time for legitimate requests, resulting in a sluggish network that is no longer agile and adaptable on the fly.

Testing SDN Infrastructure & Overlay Networks

Enterprise workloads change dynamically, requiring data centers to move workloads based on resource availability. SDN demands flexible and agile networking, making overlay networks the perfect fit. Overlay networks provide flexible, logical networks that can easily scaled as needed. Overlay network components need to be validated for their performance, scalability, and reliability.

Performance Tests – How long does it take to set up and deploy new overlay tunnels on the fly?

Scalability Tests – How many overlay tunnels can be setup?

Reliability Tests – Can the underlying infrastructure handle faults such as link failure, node failure, and congestion? Can they re-route the affected overlay traffic with minimal disruption?

Two widely-used SDN overlay technologies and the metrics that must be validated:

- EVPN with VXLAN for data center interconnect
- Segment routing for traffic engineering in WAN service provider domain

EVPN with VXLAN

Virtual eXtensible LAN, or VXLAN, provides highly scalable layer two logical networks on top of the layer three IP network. VXLAN increases traditional VLAN limits from 4,000 to 16 million logical networks. It uses IP multicast flooding to forward packets to unknown destinations.

EVPN is a layer two VPN (L2VPN) solution that enables integrated layer two and layer three services over Ethernet, with multi-homing. For datacenter interconnect (DCI), top of rack (ToR) access switches use EVPN BGP auto-discovery for learning VXLAN tunnel end-points (VTEPs) and BGP signaling for remote VM MAC learning. EVPN with BGP auto discovery and MAC learning eliminates the need for VXLAN multicast flooding between data centers. EVPN provides a mechanism to relearn MAC addresses efficiently after VM migrations with MAC mobility feature support.

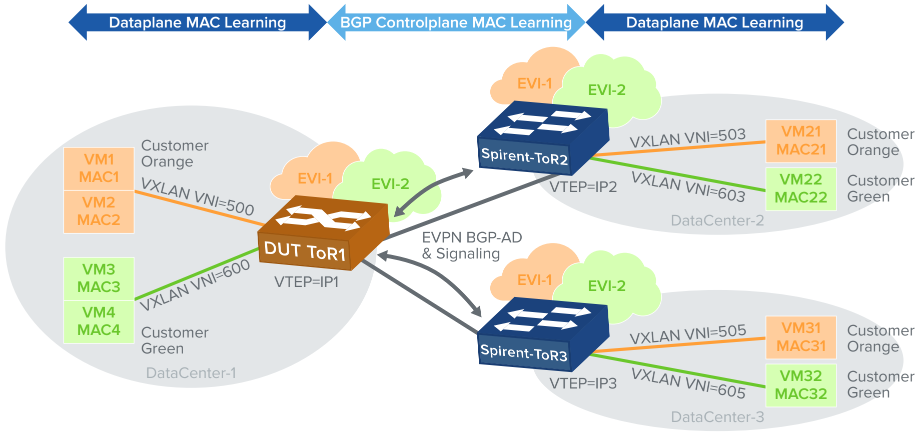

Customer	VM MAC	VTEP IP	VXLAN ID
Orange EVI-1	VM3-MAC1 VM4-MAC2	IP1	500
	VM21-MAC21	IP2	503
	VM31-MAC31	IP3	505

Customer	VM MAC	VTEP IP	VXLAN ID
Green EVI-2	VM1-MAC1 VM2-MAC2	IP1	600
	VM22-MAC22	IP2	603
	VM32-MAC32	IP3	605

Figure 6: Validating EVPN VXLAN for data center interconnect

Performance tests

Measure:

- Measure the time taken to setup thousands of VXLAN tunnels with end-to-end traffic. Measure the time taken for EVPN ToR switches to converge and learn millions of MAC addresses and their corresponding VXLAN tunnel.

Scalability tests

Determine:

- Maximum number of end-to-end VXLAN tunnels that can be setup on a top of rack (ToR) switch.
- Maximum number of remote VTEP connections supported.
- Maximum number of MACs that can be advertised and learned with EVPN BGP.

Reliability tests

Verify:

- How long it takes for network to converge after VM migration using MAC mobility.
- For multi-homed sites, measure the time to re-converge on a standby path when an Ethernet segment goes down on the active path.

Segment Routing

Segment routing (SR) provides highly scalable overlay networks for WAN service provider (SP) networks. In WAN SDN, SR operates with PCE controllers to setup up millions of traffic engineering paths on the fly. The data plane uses multi-protocol label switching (MPLS), which is well established in SP networks. This means no change to existing forwarding infrastructure and support for all L2 and L3 VPN services over SR. SR provides reduced state traffic engineering since the entire state of the label switched path (LSP) is carried in the data packet in the form of an MPLS label stack. Each label in the MPLS stack refers to a segment in the SR network that the packet must traverse. Segment routing inherently supports high reliability with equal cost load sharing and fast convergence. SR runs on traditional internal gateway protocols (IGP) IS-IS and OSPF that support equal cost multi-pathing (ECMP) for load sharing and fast re-route using loop-free alternate (LFA) mechanisms.

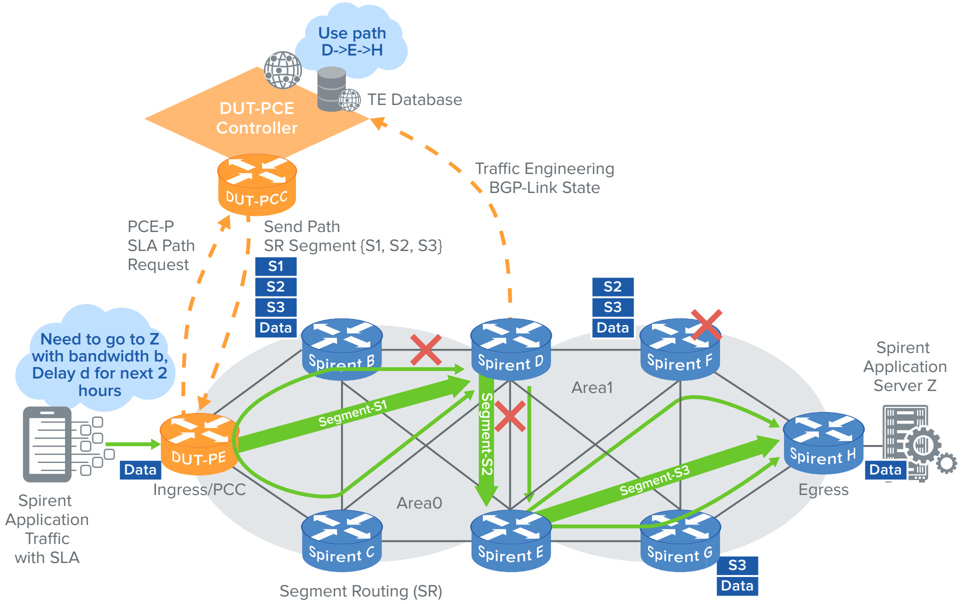

Figure 7: Traffic Engineering with Segment Routing

Performance Tests

Determine:

- Amount of time taken for device under test (DUT) PCE to compute a path and determine SR TE segments for each domain or area.
- Time taken on the ingress DUT node to set up traffic engineering paths and forward traffic in a segment routing network.
- TE path setup rate with different tunnel depths, where the tunnel depth defines the number of segments to traverse in the SR network.

Scalability Tests

Measure:

- Maximum number of end-to-end SR TE tunnels that can be setup on the DUT ingress PE. This determines the number of applications with traffic engineering that can be supported by the ingress PE.
- Maximum tunnel depth or maximum number of SR segments supported for a TE path by the underlying infrastructure. This is useful when the TE path must pass through multiple domains or areas.
- Then, verify the maximum number of SR TE tunnels that can be managed by a DUT PCE controller within a domain and across domains.

Reliability Tests

Verify:

- How long it takes for the network to converge after link failure or node failure. For SR, the DUT at the point of local repair (PLR) will use a built-in IGP fast convergence mechanism to recover the critical time sensitive applications to back up paths. The DUT ingress PE will report to the controller that the path status is down and temporarily recovered. Then, controller will re-compute the optimized SR path and update the end-to-end TE path.
- The time it takes to go from failure to re-convergence on optimized TE path.

Summary

Software Defined Networks (SDN) offer the promise of increased agility, programmability, and revenue services opportunities. This comes with a cost, however, in the form of additional routing and traffic engineering mechanisms – with associated protocols and APIs. Many of SDN's capabilities result from the centralization of control, engendering concerns over controller reliability, scalability and performance. Hence, it's critical to test, measure, and verify the functionality and performance of all SDN components. Thorough validation of new technologies in the lab helps reduce Capex by preventing expensive truck rolls and reducing customer churn. This white paper has provided the configurations and tests that must be performed in order to guarantee a trouble-free SDN deployment.

Spirent's SDN and NFV Testing solutions enable our customers to accelerate the deployments of SDN networks and help them ensure carrier-grade performance, scalability, availability and security of their SDN offerings prior to field deployments. Find out more about how to validate your SDN controllers and SDN infrastructures and learn about Spirent's SDN testing solutions by visiting us at www.spirent.com. Check out our SDN and NFV Testing solutions.

Additional resources:

- OpenFlow White Paper
- NFV eBook
- PCEP data sheet
- Segment Routing data sheet
- OpenFlow data sheet
- VXLAN data sheet
- EVPN data sheet

About Spirent Communications

Spirent Communications (LSE: SPT) is a global leader with deep expertise and decades of experience in testing, assurance, analytics and security, serving developers, service providers, and enterprise networks.

We help bring clarity to increasingly complex technological and business challenges.

Spirent's customers have made a promise to their customers to deliver superior performance. Spirent assures that those promises are fulfilled.

For more information, visit: www.spirent.com

Contact Us

For more information, call your Spirent sales representative or visit us on the web at www.spirent.com/ContactSpirent.

www.spirent.com

Americas 1-800-SPIRENT
+1-800-774-7368 | sales@spirent.com

US Government & Defense
info@spirentfederal.com | spirentfederal.com

Europe and the Middle East
+44 (0) 1293 767979 | emeainfo@spirent.com

Asia and the Pacific
+86-10-8518-2539 | salesasia@spirent.com