# GAINING FULL CONTROL OF YOUR NETWORK WITH SERVICE PROVIDER SDN

STRATEGIC WHITE PAPER

While the focus of software-defined networking (SDN) so far has been the automation of data center (DC) networks, service providers (SPs) are beginning to define a new role for SDN in helping them gain more control of their metro and wide area networks (WAN). To fulfill their requirements, the networking industry is working on a service provider version of SDN (SP SDN) that goes beyond the capabilities of traditional network management systems. The primary focus of SP SDN will be to:

- Simplify how the operations support system (OSS) interfaces with the network to accelerate the introduction of new network elements and features, and the automation of service provisioning in multilayer, multivendor networks
- Provide greater visibility and control of multilayer, multivendor networks so they can run more efficiently and support new on-demand services
- Support dynamic integration of the DC with metro and wide area networks to simplify cloud inter-connection and cloud access

This paper will share Alcatel Lucent's insight into SP requirements for SDN in metro and wide area networks (WANs), and provide an overview of the open, extensible SP SDN framework required to address them.

# TABLE OF CONTENTS

# KEY DRIVERS FOR SERVICE PROVIDER SDN

Today's networks are optimized for the relatively static and predictable task of interconnecting remote sites/users to centralized offices. Network connectivity services are provisioned by complex IT/OSS systems that rely on low-level, vendor-specific APIs. Everything is designed with long-term continuity in mind – provisioning can take days or weeks, connections are expected to stay up indefinitely, and service innovations can take years to roll out. Network programming (i.e. traffic engineering) is typically implemented in windows several months apart to accommodate the complex analysis and planning required.

The advent of distributed cloud architectures is placing considerable pressure on this model. Enterprises want SPs to deliver connectivity services as quickly as cloud vendors deliver cloud services. They only want to pay for the latency, jitter, bandwidth, etc. they need, when and where they need it. To meet their demands, SPs will have to offer innovative new connectivity services such as bandwidth-on-demand and scheduled bandwidth. These services will have to be developed quickly and enhanced frequently to retain a competitive edge. And as SPs add IT and advanced network services to the mix (for example, Software as a Service (SaaS), firewall), their customers will expect seamless integration with the underlying WAN service that delivers them. All this will be difficult to achieve — if not impossible — in an environment where application development and service provisioning is a complex and time-consuming task.

Another key driver for SP SDN is the desire for greater network efficiency. It is normal today for SPs to reserve large amounts of bandwidth to accommodate fail-overs and unexpected surges in demand. While this approach typically ensures the highest possible service levels, it is not what customers want — or are willing to pay for — all the time. It is also an increasingly expensive way for operators to make use of their network assets. As the on-demand nature of cloud consumption changes network consumption patterns by making them less static and predictable, SPs will have to allocate even more bandwidth to ensure their services don't run out of capacity, or start drifting from pre-defined quality of service attributes.

So where does SP SDN fit in? It empowers SPs to solve these problems and gain full control of their networks without incurring huge costs. SP SDN complements existing network management systems, which continue to provide comprehensive service and infrastructure management of vendor-specific gear, by adding the following new capabilities:

- Allows SPs to quickly introduce new network elements and new network vendors
- Allows SPs to quickly develop and deploy new applications that increase network efficiency and enable dynamic pay-as-you go services
- Simplifies OSS applications and reduces IT/network integration complexity
- Uses network resources more efficiently to reduce costs and increase operating flexibility

## WHAT'S REQUIRED

To achieve the network automation and optimizing benefits described above, SP SDN relies on four key capabilities: network abstraction, global network visibility, centralized network control, and policy.

The network abstraction capabilities of SP SDN simplify how the network (and the services it supports) appears to applications so they can be developed and enhanced quicker than ever before. This is accomplished by presenting only the subset of network topologies and service capabilities that is relevant to a specific application, greatly reducing the complexity it is exposed to. Abstractions will often leverage industry-standard data models so new applications can work transparently across network gear from multiple layers and vendors. Operators will also be provided with the ability to extend standard models with advanced capabilities not covered by standards to provide differentiation or to satisfy customer demand.

SP SDN allows new management applications and services to make optimal use of network assets by providing real-time, global network visibility that spans multiple layers, domains and vendors. SP SDN can assess the state of the entire network — not just a portion of it — before mapping service requests to underlying network resources. This ensures requirements for latency, bandwidth, etc. are met, and that optimal use is made of IP/optical network assets. Armed with a real-time view of the whole network, path computation engines can operate with the assurance that they have all the information necessary to make optimal choices.

Automated service provisioning and dynamic network re-programming is achieved via a centralized network control capability that spans multiple layers, domains and vendors. Application requests for connectivity can trigger automated provisioning of service instances across multiple network elements, layers and vendor boundaries — without the bottlenecks associated with manual, vendor or technology-specific provisioning approaches. SPs can directly control how and when MPLS paths and topologies are modified or they can dynamically create short-cut routes via the introduction of segment routing. New paths can be calculated whenever existing paths are not available, when they no longer support the requested service attributes, or when they need to satisfy pre-defined policies for network optimization.

Policy provides SPs with a quick and flexible mechanism to specify how and when network assets are to be consumed and who consumes them. Network policies are used to create network abstractions, to specify how service instances are mapped to transport resources, and to define triggers that push changes to network paths and topologies. Service policies are used to simplify the definition of services and the business rules that govern how, when and who uses them.

# SERVICE PROVIDER SDN FRAMEWORK

Figure 1 illustrates how these capabilities are applied in the context of an SP SDN framework.

**Figure 1. Service provider SDN framework**

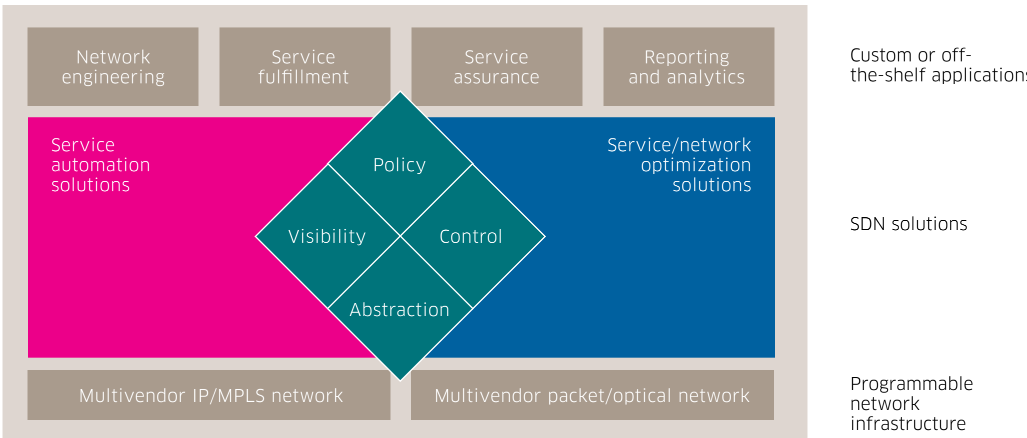


At the top of the framework are the applications that consume, manage or optimize networks under SDN control. In the WAN, these applications are off-the-shelf or custom-developed OSS tools that undergo rapid simplification as high-level REST APIs eliminate the need for complex, low-level OSS/network integration code. Key capabilities, like network and service assurance, are significantly improved since a centralized SP SDN architecture allows them to collect and correlate alarm and status data across a broader section of the network.

SDN solutions leverage the capabilities defined in the previous section to fulfill two distinct but complementary roles: network optimization and service automation. Both product types support extended visibility of L0-L3 network topology, statistics, status, analytics, connectivity services and flows. Optimization products contain IP, optical and hierarchical PCEs, along with policies that specify how and when the PCEs should be used to optimize the network. Service automation products use normalized data models to automate the provisioning of network services and the network infrastructure that supports them. Multivendor mediation is achieved directly through plug-ins that support industry-standard protocols such as Netconf/YANG, BGP LS, SNMP, PCEP, CLI, and OpenFlow. Mediation is also achieved by leveraging the Open Daylight (ODL) framework and the many vendor plug-ins it supports. Like Nuage Networks VSP, all SP SDN products delivered by Alcatel-Lucent will interoperate with external controllers and architectures that comply with SDN standards, such as ODL.

A software defined network is not complete without a programmable network infrastructure that is just as open, flexible and easy to use as the SDN solution software that controls it. Openness requires full support of traditional and new SDN network protocols (see above) so network infrastructure can be controlled with off-the-shelf SDN solutions, by in-house SDN controllers, or by a combination of the two. Flexibility requires fine-grain control of network assets — from flexible ODU sizes to configurable flow steering — so that network resources can be used in the most efficient manner possible. Easy-to-use requires new functionality, like IP/optical integration, to simplify the control of multilayer networks while making them more flexible and efficient.

While the above principles enable the transformation of an SP's core assets — traditional metro and wide area networks — into software-defined networks, they can also be applied to SP data center and CPE devices to extend SDN benefits all the way from the virtual machine (VM) to its consumer at the branch (see Figure 2).

**Figure 2. Service Provider SDN solutions**

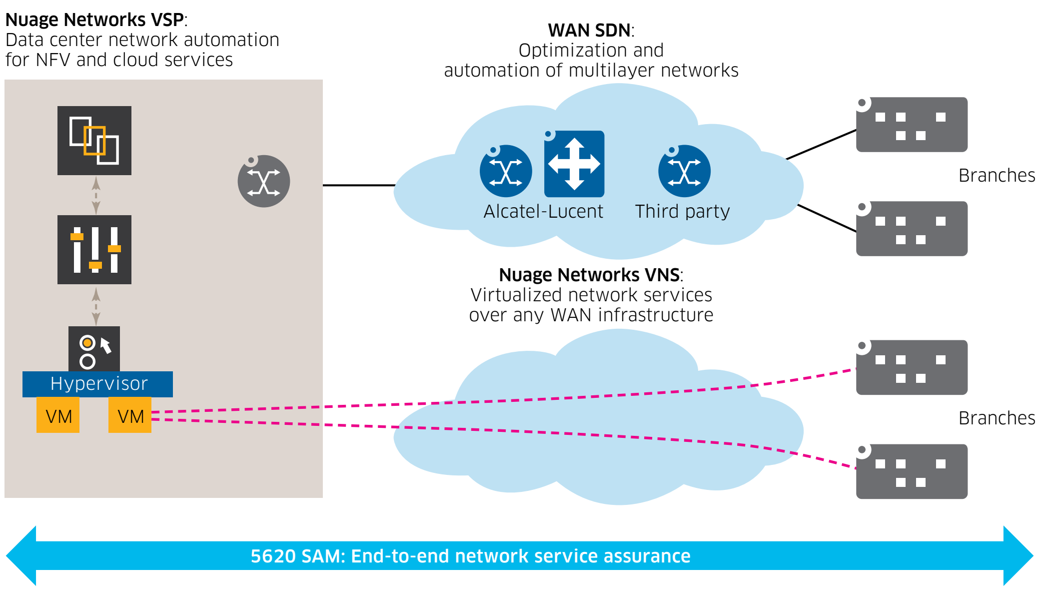


In the DC, SDN solutions such as Nuage Network Virtual Services Platform (VSP) enable full automation of local networks in support of virtualized network functions (vNFs) and IT applications/services. DC SDN solutions communicate with WAN SDN solutions to dynamically connect applications to remote users and other DCs in a manner that takes full advantage of underlying metro and wide area network services. For SPs with both application and network infrastructure, this approach provides the greatest flexibility to quickly deliver exactly what customers want while making optimal use of network assets.

In contrast to the "underlay" approach discussed thus far, cloud and business services groups within SPs, as well as providers with many disparate networks, may prefer the "overlay" approach of bringing SDN to the WAN. This involves extending the dynamic creation of tunnels (i.e. via DC SDN solutions) to encompass remote DCs and branches across any network. The advantage is network transparency: cloud providers can instantiate applications and IT infrastructure services that reach any customer, anywhere — without regard to who owns the underlying network infrastructure. SPs that lack network reach can rapidly extend their enterprise network services to reach customers wherever they may be. An example of an overlay SDN solution for the WAN is Nuage Networks Virtualized Network Services (VNS).

The other major components at play here are traditional network and service management systems, such as the Alcatel-Lucent 5620 Service Aware Manager (SAM). These products will continue to provide comprehensive service provisioning and comprehensive infrastructure management for a given vendor's network gear. More importantly, from an SDN context, they contain the service assurance functions — such as fault and performance management — that are required to achieve end-to-end network service assurance for software defined networks. The rich set of data they collect and analyze, from link utilization status to predictive analytics, are critical to guiding the SDN solutions that assign network resources, trigger network changes and govern network behavior. Support for end-to-end assurance — from the DC, to the WAN to the branch — is critical to making SDN deployments manageable, and hence achievable, in the near future.

# USE CASES

The following customer-derived uses case illustrate how SPs are planning to use SP SDN to optimize their networks and to eliminate roadblocks to dynamic operation and cloud service adoption.

## On-demand, inter-DC connectivity services

Problem: Customers want their network and bandwidth services to dynamically follow data center VM movements. They want to support changes in traffic patterns such as nightly backups and bursts of customer activity, while adhering to existing network and security policies.

Solution: Empower customers to dynamically manage connectivity between DCs. Create new services that allow them to re-allocate purchased bandwidth between different DCs, or provide additional bandwidth to match changing traffic patterns. Allow customers to specify their needs, such as bandwidth, time of day and the length of time the connection is required.

## On-demand service chain extension into WAN

Problem: Cloud services – including IT services and network services such as firewalls – and the users that consume them are difficult to stitch together in service chains that incorporate metro and wide area networks. Provisioning often entails manual creation of DC/WAN connections and manual injection of custom routes or ACLs into the network.

Solution: Accelerate the adoption of cloud services by leveraging DC SDN solutions to create service chains, and SP SDN to dynamically extend them to the WAN. SP SDN dynamically sets up connections between the DC and WAN, and dynamically injects ACLs into the network to steer service chain subscribers to the appropriate DC.

## Dynamic path selection and traffic steering

Problem: Customers that leverage IP/optical integration to provide multiple transport choices with varying degrees of redundancy require a way to intelligently steer flows onto these new tunnels.

Solution: SP SDN allows operators to create policies that take customer preferences (i.e. the service level they are prepared to pay for) or network optimization rules (i.e. overflow traffic steered onto unused links during peak hours) into account when mapping services instances and flows to transport links.

### Real-time network re-engineering

Problem: Network engineering windows are spread too far apart to ensure service availability and quality as traffic patterns become less static and predictable. The process takes time; operators typically monitor each network layer in isolation, correlate the results offline and run an impact analysis to make changes.

Solution: SP SDN provides a policy-driven way to automate network engineering changes at layers 1-3. Operator-defined policies leverage rich data collection, alarm correlation and predictive analytics to make changes to the network to avoid imminent outages, accommodate surges in network demand, and to ensure SLAs.

### Network as a Service

Problem: Operators need a simpler, more effective way to virtualize WAN infrastructure so that resources can be more efficiently allocated in “slices” across a growing and increasingly diverse number of retail, wholesale and infrastructure services.

Solution: Allow operators to create and manage multilayer, traffic-engineered service planes that meet specific SLA requirements for retail, wholesale and infrastructure services. Service planes can be used to drive new revenue; for instance, new enterprise storage and database replication services can be mapped to a service plane that monetizes unused bulk bandwidth. Optical forward error correction (FEC) can be turned off in one service plane to improve bandwidth efficiency and minimize latency, and turned on in another service plane to support applications that can tolerate higher latency but no packet loss.

## CONCLUSION

The need to better align the network with the cloud, coupled with the drive for better operational and bandwidth efficiency is driving SDN into metro and wide area networks. A new generation of SDN solutions will fulfill these needs by:

- Simplifying OSS/network integration to accelerate service automation and innovation in multivendor, multilayer networks
- Providing greater network visibility and control so networks can run more efficiently and support new on-demand services
- Supporting dynamic integration between the DC and WAN to simplify cloud interconnection and cloud access

Integration with existing management systems to leverage assurance functions and data is key to SDN transformation, both in terms of feeding SDN solution engines to make the right mapping and optimization choices, and in enabling the end-to-end assurance capability required to deploy SDN-based services in real-world environments.

# ACRONYMS

| | |
|---|---|
| ACL | access control list |
| API | application programming interface |
| BGP LS | Border Gateway Protocol Link-State |
| CLI | Command Line Interface (or Instruction) |
| CPE | customer premises equipment |
| DC | data center |
| FEC | forward error correction |
| MPLS | Multiprotocol Label Switching |
| ODL | Open Daylight |
| OSS | operations support system |
| PCE | path computation element |
| PCEP | Path Computation Element Protocol |
| SDN | software defined networking |
| SLA | service level agreement |
| SNMP | Simple Network Management Protocol |
| SP | service provider |
| VM | virtual machine |
| vNF | virtualized network function |
| VNS | Virtualized Network Services |
| VSP | Virtualized Services Platform |
| WAN | wide area network |



Alcatel·Lucent