THE CLOUD-OPTIMIZED MAN AND WAN

LEVERAGING A MULTI-LAYER SDN FRAMEWORK TO DELIVER SCALABLE AND AGILE CLOUD SERVICES

STRATEGIC WHITE PAPER

Over the next few years, service providers will adopt Software-Defined Networking (SDN) to prepare their MANs and WANs for the delivery of cloud applications and services. A multi-layer SDN framework will be required to deliver the cloud-optimized routing and transport capabilities necessary for this transition. This document summarizes the need for multi-layer SDN and illustrates how it can be applied to increase the agility, scale and efficiency of carrier MANs and WANs as they enter the cloud era.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

By allowing applications to consume IT resources on demand, cloud technologies have fundamentally changed the way data centers (DCs) are being designed and operated. Software-Defined Networking (SDN) extends this cloud model to the network by allowing applications to consume network connections as quickly and as easily as they consume virtualized compute and storage. In an SDN-based DC network, connections are not permanent entities; they are initiated by virtual machines (VMs) in seconds, and are dismantled when no longer required. As the cloud expands to encompass multiple DCs and users across metro and wide area networks (MANs and WANs), applications will expect the carrier network connections that underpin this new, distributed cloud fabric to leverage SDN and be just as easy to set up and consume.

Achieving SDN in the MAN/WAN is no trivial task. While DCs are simple, homogeneous and can count on essentially limitless bandwidth, carrier networks are complex, multi-vendor, and subject to many technology and bandwidth constraints. Any attempt to introduce a faster and more automated way of provisioning the MAN/WAN must work in concert with provisioning systems for existing services. The dynamic consumption nature of distributed clouds and their users will put significant pressure on operational models designed for static and predictable traffic patterns. Service providers will need to augment existing traffic engineering processes with a dynamic resource management capability that can ensure network efficiency, resiliency and availability in the face of rapidly changing cloud connectivity needs.

The industry has already begun the process of evolving SDN for carrier networks with proposals to abstract and open the proprietary packet/optical layer to external control using OpenFlow™. While this is a good starting point, little value can be realized without a broader SDN framework that can provision and correlate topology, resource and physical constraint information across multiple layers — including packet/optical, Ethernet, IP and IP VPNs. Only a multi-layer SDN framework has the global visibility and universal control necessary to deliver on the SDN promise of increased network agility, efficiency and scale.

This document details requirements and use cases for cloud-optimized MAN/WAN services and the cloud-optimized routing and transport that delivers them. It is directed at service providers and large private enterprises that want to better understand how SDN can help them prepare their MAN/WAN infrastructure for the cloud.

KEY DRIVERS FOR SDN IN THE MAN AND WAN

Traditional MAN and WAN services are optimized for a relatively static and predictable task: interconnecting remote enterprise sites and users with centralized offices that serve as hubs for enterprise business processes. Network service provisioning is handled by complex IT/OSS systems that use low-level vendor-specific APIs, by manual provisioning processes based on command-line interfaces, or by script files. Everything is designed with long-term continuity in mind — service provisioning can take days or weeks, connections are expected to stay up indefinitely, and service innovations can take years to roll out. Major adjustments to the network (such as traffic engineering and resource re-allocation) are typically implemented in windows several months apart to accommodate for the complex analysis and planning required. The advent of distributed cloud architectures is placing considerable pressure on this model.

The need for real-time service delivery and rapid service innovation

- Once a physical network link is in place, enterprise customers are demanding that the MAN/WAN service that underpins a cloud service be delivered just as quickly and as efficiently as the DC portion — not in days, weeks or months, but in seconds. They want service providers to ensure that their changing service requirements (including acceptable delay, jitter and bandwidth) can be fulfilled quickly and on demand.
- As VMs move from DC to DC based on the availability of virtualized compute and storage, and as network connectivity needs shift from web/transactions in the daytime to backups at night, enterprise customers are looking for more flexible and dynamic ways to consume bandwidth. New usage-based services and service attributes, such as bandwidth-on-demand or scheduled bandwidth, are required to satisfy their changing connectivity needs.
- Service providers need a rapid and cost-effective way of delivering advanced network service bundles to customers. They need an approach that eliminates the need for manual, time-consuming manipulation of routing tables to steer traffic to individual application servers, across multiple DCs.
- New service capabilities need to be brought to market much quicker to ensure competitiveness.

Achieving the above requires a shift away from the low-level, vendor-specific APIs that are used to create and provision network services today. To speed up service innovation, service providers need standards-based, high-level APIs that “abstract” (that is, provide simplified views) of network capabilities. In contrast, low-level commands force them to deal with the full complexities of each network layer and vendor.

The transformation of the MAN/WAN into an easily-consumable resource for applications creates a fertile environment for the rapid evolution of new, usage-based IT and network services. To speed up service provisioning, service providers need to eliminate the errors and delays associated with manual coordination of multiple provisioning systems (per layer and/or per vendor), or with wholly-manual provisioning processes.

The need for multi-layer visibility and control

As traffic patterns start to change and become less static, service providers will have to manage the increased complexity this introduces into their networks. Network operations will no longer be able to ensure service availability and quality with planning cycles spread months apart. To ensure services don’t run out of bandwidth or drift from their pre-defined QoS attributes, network operators will have to revisit the network allocation and engineering decisions they make on a more frequent basis.

The process of mapping network service requests to IP/optical transport resources will require a real-time analysis component. The service provider must be able to pick the path that both satisfies a customer’s SLA requirements and ensures the IP/optical infrastructure is used in the most efficient manner. New paths must be calculated whenever existing paths are not available, or when they no longer support the requested service attributes due to changing network conditions.

To facilitate real-time path selection and computation, service providers require instant, global visibility of all network resources and topologies — from IP/MPLS to the optical layers — along with the capability to provision these resources on demand. This is not possible today given the complexity involved in gathering, correlating and analyzing information, and the difficulty in coordinating provisioning across many layer- and/or vendor-specific systems. Instead, operators are assigning ever increasing amounts of bandwidth to accommodate the growing peak-to-mean gap. Global visibility and control will allow service providers to run their networks “hotter”. They can assign every retail service and selected wholesale customers their own tightly controlled slice of the MAN/WAN that has been engineered to satisfy their unique connectivity requirements.

CLOUD-OPTIMIZED ROUTING AND TRANSPORT

Preparing the network for the cloud requires a multi-layer SDN framework that partitions carrier networks into two major components: network virtualization and automation, and programmable IP/optical transport (see Figure 1). The capabilities of each are briefly described below, while the next section provides use cases that articulate how they can be applied to remove roadblocks to cloud service adoption and revenue.

Figure 1. Cloud-optimized routing and transport based on multi-layer SDN

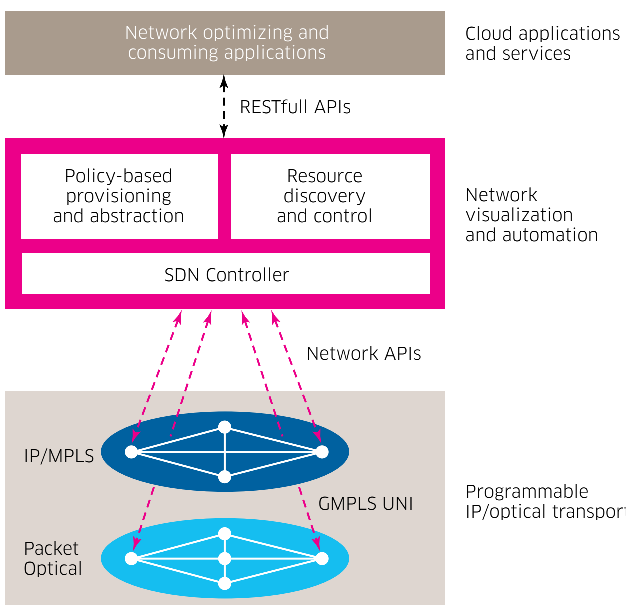

The primary role of the **programmable IP/optical transport** component is to provide highly reliable, high performance packet/optical, Ethernet, IP and VPN transport to enable retail, wholesale and infrastructure services. Key capabilities include:

- Specialized NPUs and optical hardware deliver Layers 0 to 3 grooming, switching and forwarding capabilities at an optimal price for performance.
- Proven, distributed protocols are embedded within network hardware to ensure network scale, stability and resiliency.
- Open APIs (for example, OpenFlow, NetConf, SNMP, RADIUS and DIAMETER) allow standards-based applications and SDN controllers to monitor and control network resources at any IP/optical layer, and from any standards-compliant vendor.
- New capabilities, such as segment routing, can be introduced to enable offline path computation for new cloud-based services and capabilities (see use case example in following section).

The **network virtualization and automation** component provides a high-level, abstracted view of the network to applications, and has real-time visibility and control of all network resources — from Layers 0 to 3. It resides in generic CPU platforms to achieve the highest scale for the lowest cost. Key capabilities include:

- SDN controllers are used to provision all elements in one or more network layers.
- Policy-based, need-driven service provisioning allows network tasks and services to be defined as policies so that they can be instantiated faster, and on a mass scale. A policy framework allows enterprise services to be delivered in the same way that wireless services are delivered to cellular subscribers today. Once a customer application or portal requests connectivity, the desired service, along with all the attributes provisioned for the customer, is automatically provisioned across all layers — including IP, Ethernet and packet/optical. The service can be turned up or down based on need.
- Policies can be used to combine many lower-level network tasks into a higher level function that shields applications from the unnecessary complexity of vendor-specific, low-level provisioning. The abstraction of low-level network functions into a standards-based, high-level "business language" allows service providers to innovate faster and compete better.
- **Resource discovery and control** (see Figure 2) consists of the **Resource Manager (RM)**, and the **Path Computation Element (PCE)**. These features provide the following:
  - ¬ The RM is responsible for dynamic discovery of network topologies, resources and constraints across all IP/optical layers. SDN controllers leverage this global view to dynamically map services, applications and user flows to IP/optical resources. The RM is aware of constraints across all layers — from jitter and delay at the IP layer, to modulation techniques, latency and FEC at the optical layer. This allows for optimal service selection, path selection and path computation in a changing network environment.
  - ¬ Analytics are available via ALTO+. If multiple network paths and services are available to satisfy an online connection request (such as a DC-to-DC connection), ALTO+ can analyze the global topology and state information and provide a ranked list of options. ALTO+ can also enable optimization after service turn-up.
  - ¬ To provide admission control, the RM performs bandwidth reservation to ensure key transport links do not become oversubscribed.
  - ¬ The PCE allows operators to calculate custom network paths (routes) for selected services and applications, removing this burden from IP/optical routers and switches whenever it makes more sense to implement it externally.

Figure 2. Resource discovery and control

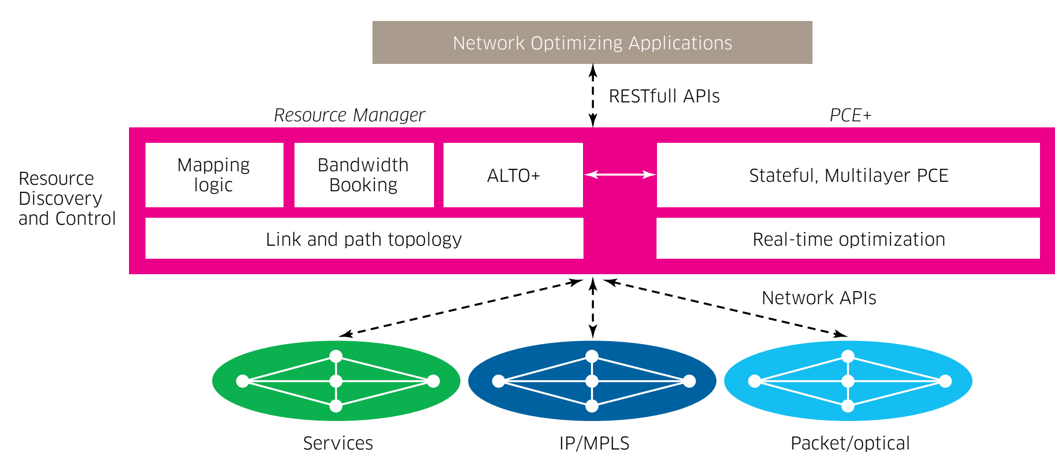

CLOUD-OPTIMIZED MANS AND WANS – SELECTED USE CASES

Like any new technology, the introduction of SDN into carrier networks will be driven entirely by economic considerations. The ability to create new revenue streams, increase efficiency and agility, and reduce OPEX and CAPEX will be required to validate any investment decision. As the use cases below highlight, SDN can eliminate roadblocks to cloud service adoption and enable new retail and wholesale revenue. It does this without impacting operational models and processes that are well-tuned to meeting the needs of existing network services.

On-demand, inter-DC connectivity services

Problem: Customers want their network and bandwidth services to dynamically follow DC 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 on demand to match changing traffic patterns. Allow customers to identify specific needs, such as bandwidth, time of day and the length of time the connection is required.

SDN role: The new service can be quickly defined using the policy-based abstraction framework. User requirements are mapped to the service policy, which is executed by some form of trigger, such as when the VM that starts the backup process is initiated. The policy-based provisioning framework automatically provisions all network resources. It can increase the shaping rate of an IP port or the ODU size of an optical connection as required, and brings the service down when no longer needed, or based on a pre-arranged contract. The RM manages bandwidth to ensure it is not oversubscribed.

On-demand service chain creation over the MAN/WAN

Problem: Cloud services — including IT services and network services such firewalls — are often distributed across multiple DCs, making it difficult to stitch them together into custom packages for remote users. Provisioning can take weeks, involves manual changes to route tables, and results in static, inflexible routes through the network.

Solution: Accelerate adoption of advanced services and minimize operational costs by dynamically creating service "chains" for each desired service combination.

SDN role: PCE is used to calculate optimal paths between virtualized service appliances, eliminating the manual, time-consuming task of editing route tables to stitch together connectivity. Policy is used to dynamically map each customer to their own optimized service chain as they enter the MAN/WAN. Service chain topology is kept in the RM so that it does not tax routers and force unnecessary network hardware upgrades due to the large amount of CPU required.

Dynamic path selection and traffic steering

Problem: Service providers and cloud providers must take network conditions into account when determining how and where to steer traffic, or where to place an IT resource. This information is required to determine which video cache is best placed to serve the needs of a specific set of users, which DC should host a VM, or which public cache or DC is best suited to provide additional capacity when local resources are oversubscribed.

Solution: Allow customers to state connectivity preferences between two endpoints that include latency, jitter and bandwidth. Provide a means to rank the multiple paths and services that satisfy these criteria, in real time.

SDN role: Multi-layer analysis allows ALTO to arrive at the best choice. For instance, a direct optical path between two points may seem the best option from an optical point of view. However, visibility at the IP layer also reveals an IP path with one hop that has acceptable latency and provides greater bandwidth. Peering point congestion and cost data is used to determine which peering point is best suited to reach a public DC or cache. The necessary mappings can be automatically provisioned at the IP and optical layers, or alternatively, the ranking information can be passed to a CDN or cloud orchestration layer.

Dynamic, multi-layer network optimization

Problem: Network engineering windows are spread too far apart to ensure service availability and quality as traffic patterns become less static and predictable. Service providers must monitor each network layer in isolation, correlate the results offline and run an impact analysis for even the smallest change.

Solution: Provide a policy-based way to track SLAs and automate traffic and network engineering at Layers 1 to 3. Transport connections, metrics and bandwidth can be dynamically adjusted in operational windows that are separated by minutes and weeks instead of months.

SDN role: External applications are used to access and analyze the global topology and state information maintained by the RM on a continual basis. The output is used to dynamically trigger polices that re-provision network attributes such as MPLS-TE parameters or ODU container sizes, and service mappings such as IP flows into optical lambdas. This allows service providers to avoid out-of-resource conditions, improve QoE, or make optimal use of bandwidth. Network resilience can be improved via the dynamic creation of short cuts and backup paths that use segment routes to avoid disrupting MPLS topology.

Traffic-engineered service planes

Problem: Operators need a simpler, more effective way to virtualize the MAN/WAN infrastructure so resources can be more efficiently allocated in slices across a growing and increasingly diverse number of retail, wholesale and infrastructure services (Figure 3).

Solution: Allow operators to create and manage multi-layer, traffic-engineered service planes that meet the SLA requirements of a targeted set of applications. 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 FEC can be turned off in one service plane to improve bandwidth efficiency and minimize latency. It can be turned on in another service plane to support applications that can tolerate higher latency but no packet loss.

SDN role: Service planes can be created manually, or with network optimizing applications, by accessing L0 to L3 topology and constraint information in the RM and grouping together paths with common service attributes. Policies can be created to dynamically map traffic identified by a wholesale customer, service or application to the appropriate service plane.

Figure 3. Traffic engineered service planes

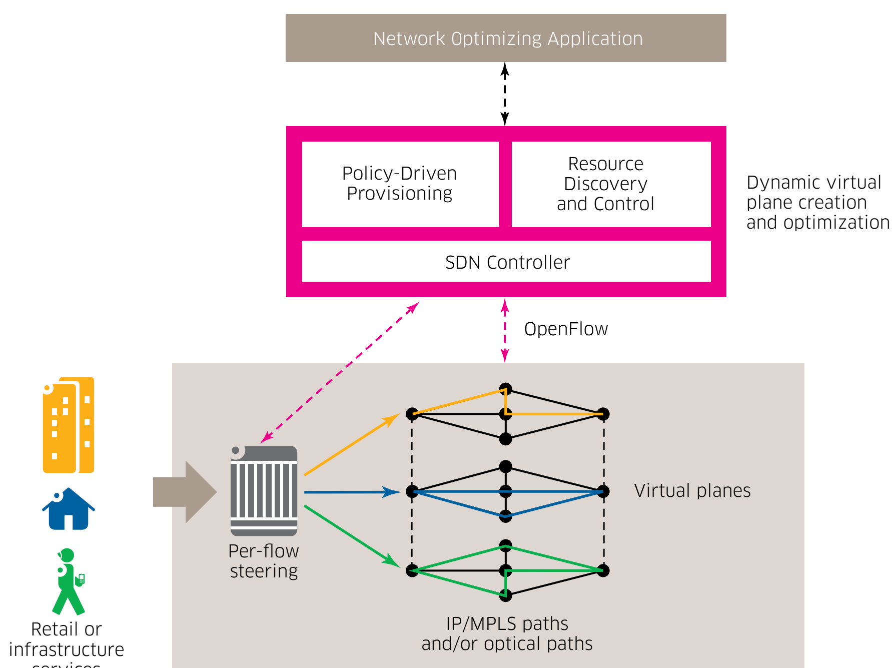

CONCLUSION

The advent of cloud computing and cloud-based IT services is changing the way service providers — and their enterprise customers — are building and using their networks within the DC, between DCs and across the MAN/WAN. Initially positioned as a means of virtualizing and automating DC networks, SDN is now being applied to the MAN/WAN to deliver cloud-optimized routing and transport. While some work has begun on abstracting the optical layer within OpenFlow, this is only a partial solution. Only a multi-layer SDN framework — from L0 to L3 — has the global visibility and universal control necessary to deliver the network agility, efficiency and scale required by distributed cloud architectures.

ACRONYMS

API	Application programming interface
CDN	Content Delivery Network
DC	Data center
FEC	Forward Error Correction
MPLS-TE	Multiprotocol Label Switching-Traffic Engineering
NPU	Network processing unit
ODU	Optical Data Unit
OSS	Operations support system
PCE	Path Computation Element
QoE	Quality of Experience
QoS	Quality of Service
RM	Resource Manager
SDN	Software-Defined Networking
VM	Virtual machine
VPN	Virtual Private Network

Alcatel·Lucent