OPEN NETWORKING
FOUNDATION

Operator Network Monetization Through OpenFlow™-Enabled SDN

ONF Solution Brief
April 3, 2013

Table of Contents

Executive Summary

In a broad industry effort spearheaded by the Open Networking Foundation (ONF), Software-Defined Networking (SDN) is transforming outmoded WAN designs by decoupling the control and data planes, centralizing network intelligence, and abstracting applications from the underlying network infrastructure using the OpenFlow standard.

As a result of this effort, data center and telecom operators will gain unprecedented network programmability, automation, and control, enabling highly scalable and flexible WAN solutions that readily adapt to changing business needs.

The emergence of viable SDN solutions is timely to address the challenges that operators face in today's business environment. Smart devices, video content, and cloud services are all generating double-digit growth in network traffic. Operators, however, are still locked into cost-avoidance strategies, bearing most of the costs and risks of maintaining and upgrading the network but with little ability to monetize this new network traffic. Declining average revenue per user (ARPU), market saturation, and a volume-based subscriber acquisition model combined with intense competition from over-the-top (OTT) services are all leading operators to the inescapable conclusion that the future lies in new services and experiences. In order to extract the full value of their core asset, the network, operators must drive a customer-centric market shift, moving from commoditized connectivity services to leveraging subscriber and network intelligence for service innovation and new business models.

But legacy networks are inflexible, static, and closed to innovation. This severely limits the operator's ability to cost-effectively respond to the scale, performance, and user experience requirements of today's dynamic environments, or to roll out differentiated services. Little application visibility into real-time network state, and the lack of network control and application interaction, result in compromised application and user experience and inefficient use of network resources.

Connectivity service features currently are tightly coupled with the network infrastructure, and operators are required to configure devices individually—and, often, manually—leading to capital and operational inefficiencies. The lack of demand visibility coupled with the lack of a feedback loop between the network and applications forces operators either to over-provision the network by 60–100% or more to minimize the likelihood of service degradation, or to be resigned to best-effort service delivery. This impacts their ability to offer and get value-added revenue from customized services and advanced SLAs based on application type, subscriber profile, customer end point device and location, and other characteristics.

OpenFlow-based SDN Enables Network Monetization

SDN with OpenFlow provides the framework and tools to enable symbiotic linkage between the operator's network and the applications that use the network to deliver end-customer services. As shown in Figure 1, this linkage is mediated by a logically centralized network control layer that maintains a global inventory of all network resources and completely controls resource allocation in response to evolving application-specified traffic demands. This linkage is further enabled through open APIs between the applications and the control layer and by the OpenFlow interface and protocol between the control layer and the infrastructure layer. OpenFlow enables the control layer to program flows and their treatment into the data-plane forwarding tables of the network devices.

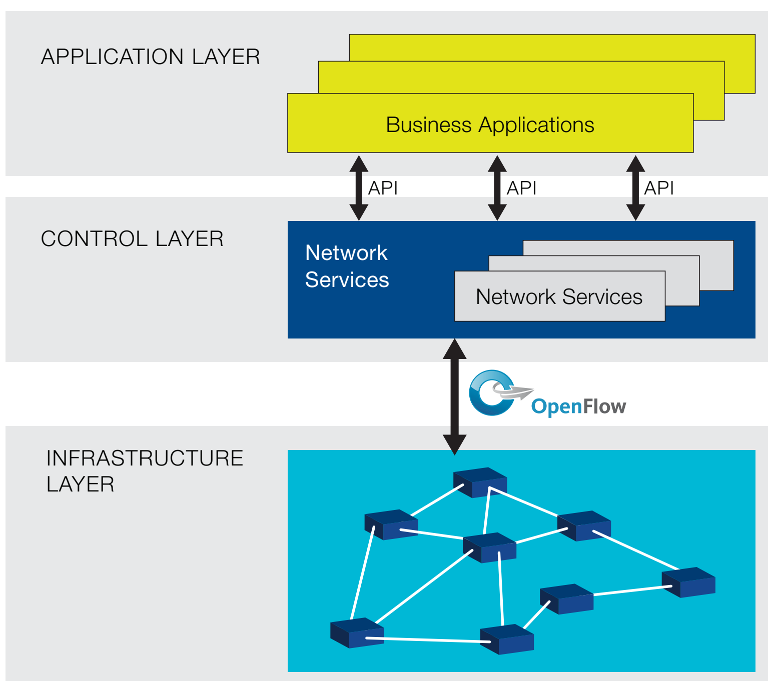

FIGURE 1
SDN framework

This explicit control of and linkage with the application layer enables operators to monetize their network investments through new services by allowing and automating direct association between individual or logical groupings of clients and applications and the network resources necessary to deliver differentiated and measurable quality of service (QoS). The SDN framework may also enhance monetization by markedly improving existing services—e.g., through faster turn-up, more competitive cost, broader coverage, mass customization—so that their revenue generation is significantly increased.

This brief describes several examples of network monetization, mainly of the new service variety, including Bandwidth on Demand, Bandwidth Exchange, Pay for QoS, and Network Features for Pay.

Bandwidth on Demand

Big-data file sets, cloud networking, and ad hoc inter-enterprise collaboration projects are all on the rise. And all of these trends result in WAN bandwidth demand peaks that are 10 to 20 times[1] greater than their mean, with the peaks lasting anywhere from less than an hour to several weeks or more. Contracting for a private line or committed information rate (CIR) sized to address the peak is at best wasteful—and for some, prohibitively expensive. Bandwidth on Demand (BWoD) services enable enterprises or cloud providers to dynamically establish or resize connectivity from the fixed or wireless access network through the core as necessary, so they pay only for what they consume. These connections can be established between subscriber locations, from the subscriber to a service gateway (e.g., a cloud data center), or from the subscriber to a third-party interconnect point.

[1]Based on modeling performed by Ciena on behalf of their customers

BWoD services have been offered by a limited number of operators, but the current model faces several challenges.

- The lack of automation capabilities makes it difficult to roll out self-provisioned services and respond to time-sensitive changes in bandwidth requirements. While customers are given some control and are able to invoke the services through a portal (reducing the need for network operators to be involved in the service order provisioning process), this is very limited in scope.
- Frequent changes in a distributed control environment sometimes lead to transient overloads caused by the network having to manage traffic from multiple sources that share the same network link, resulting in congestion and instability.
- The lack of a standard interface means operators today must interface their OSS/BSS systems to a vendor-specific network infrastructure. This requires them to either redesign their control applications for each vendor or to limit their services to a single vendor.

Deploying BWoD from an OpenFlow-based SDN architecture with a programmatic northbound API allows operators to have centralized, granular control over the networking infrastructure. As depicted in Figure 2, it also enables customers to automatically request dynamic changes to bandwidth allocation and other Quality of Service parameters at the packet and/or optical layers, either immediately or scheduled in the future. The SDN control layer can leverage topology-aware path computation to cost-effectively enable bandwidth on demand. SDN provides a real-time topological view of the network, enables network virtualization, and allows network bandwidth reservation to provide guaranteed performance on a per-connection or flow basis to meet SLA requirements.

As an example, BWoD offered as Network as a Service (NaaS) enables enterprise or cloud providers' applications to interact with the network control layer to query network performance and resource availability. This allows those organizations to better coordinate end device capabilities, data center workload placement, and NaaS provisioning to meet their user application bandwidth, class of service, priority, security, and cost-sensitivity needs.

FIGURE 2
Bandwidth on demand

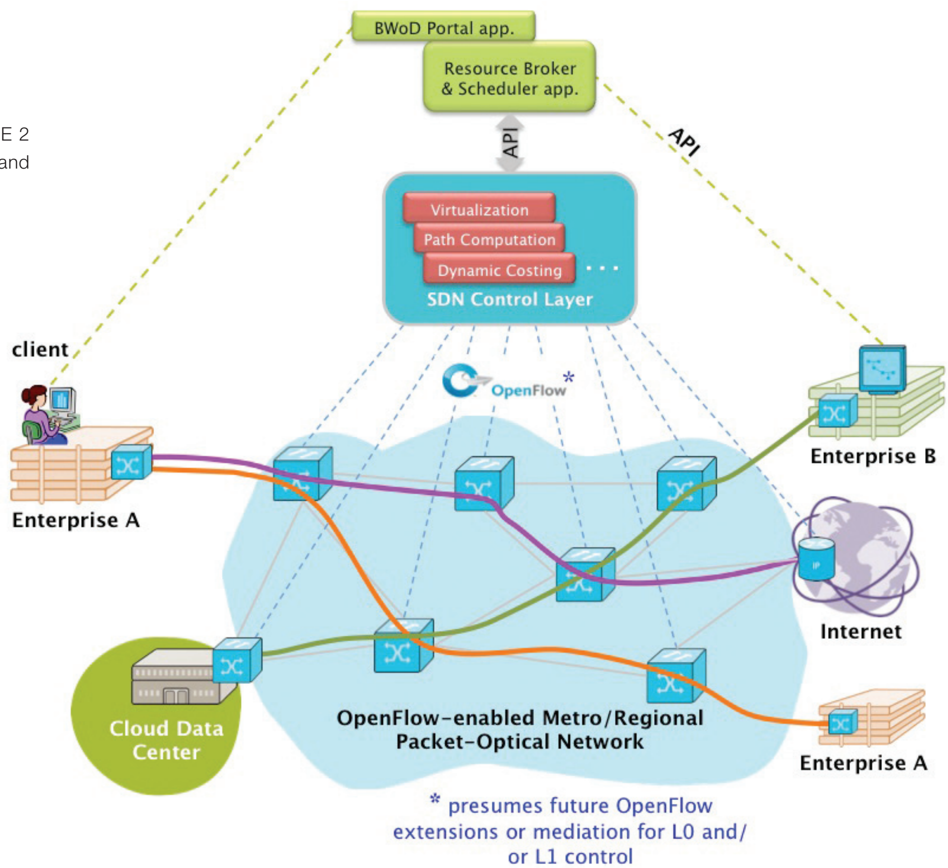

SDN's centralized global view of network resource supply and customer service demands paves the way for more intelligent and dynamic BWoD pricing. An analytics engine could evaluate current supply and demand as well as historical temporal demand peaks and supply shortages to raise network cost database entries during shortages and lower them during gluts. Through continual learning of the price elasticity of demand, these adjustments can become more refined, enabling the analytics engine to maximize network revenue per bit. Network virtualization allows operators to leverage the same networking and operational infrastructure on which they deliver traditional services to create and offer flexible BWoD services and new billing models. Multiple, diverse pay-as-you-go service pricing options increase the addressable market such that even smaller enterprises can afford short-term, high-capacity connectivity. These capabilities significantly increase profitability and drive new revenue streams.

Bandwidth Exchange

The sourcing of raw bandwidth to deliver networking services presents a constant challenge for both operators and virtual network operators (VNOs). Both must acquire sufficient bandwidth to supply target markets and service anticipated demand growth, while avoiding excessive idle inventory. They often rely on predictive business models to evaluate bandwidth requirements on existing and future extensions of their networks. For VNOs, the purpose of these models is to determine when and where to lease wholesale capacity. For operators, it is to drive their "buy/build versus lease" decisions. While buy decisions involve capital and lifecycle operational expenses, they tend to provide operators with the greatest control over their network infrastructure. Conversely, bandwidth lease decisions may make better sense where capacity does not warrant dedicated builds or is not expected to grow substantially over time.

However, planning and forecasting are by nature imperfect. The amount, timing, and location of bandwidth demands are difficult to predict, leading to network segments or links with either excessive or insufficient capacity. Lengthy lead times from equipment suppliers or traditional wholesale bandwidth providers can exacerbate the issue, particularly in areas of the network where responsiveness to new demand opportunities is a critical differentiator. Whereas idle bandwidth that is not generating revenue is undesirable, the inability to competitively respond to new opportunities is particularly detrimental for network operators.

Bandwidth exchange marketplaces provide an alternative means for network operators to address capacity planning challenges. These markets enable bandwidth suppliers and buyers to easily trade bandwidth like a commodity (with options and futures), and provide a means to transfer the control of real bandwidth resources between parties. Bandwidth exchange markets facilitate a common inventory of bandwidth contracts that can be priced and exchanged in an automated fashion. In addition to providing a vehicle for operators and pure bandwidth providers to monetize excess capacity, or for operators and VNOs to obtain spot capacity, bandwidth exchange markets provide other benefits, including:

- Quicker ROI from projects requiring capacity expansion
- Reduced operational costs from the automation of bandwidth acquisition and provisioning
- Optimized network expansion costs associated with "leased" bandwidth, taking advantage of general market pricing efficiency and elasticity

Bandwidth exchange markets create opportunities to make bandwidth trading readily available to global markets, but several challenges need to be addressed. From a networking perspective, these include:

- Registration, tracking, and management of available time-based bandwidth inventory across multiple network domains
- Secure and automated orchestration, scheduling, coordination, and provisioning of resources between multiple supply and demand entities
- Integration of network operators' management systems into the bandwidth exchange BSS/OSS
- Policy management and administration of bandwidth resources according to parameters such as time, duration, volume, and location
- Monitoring and enforcement of standardized service-level agreements across the market

Figure 3 illustrates the concept of bandwidth exchange markets and a reference implementation architecture. Suppliers and buyers meet at common, neutral exchange "points" where bandwidth transactions can occur, such as for carrier Ethernet or optical transport bandwidth. Enterprise companies, for example, may utilize the bandwidth exchange to acquire the most cost-effective leased bandwidth service to Cloud Provider A for a transaction-oriented cloud operation. Similarly, Operator A may use the bandwidth exchange to locate cost-effective bandwidth to interconnect two of its disjoint networks. Diverse bandwidth requirements could be met by leasing bandwidth from different exchange points over different suppliers' networks (Operator B and Bandwidth Provider C).

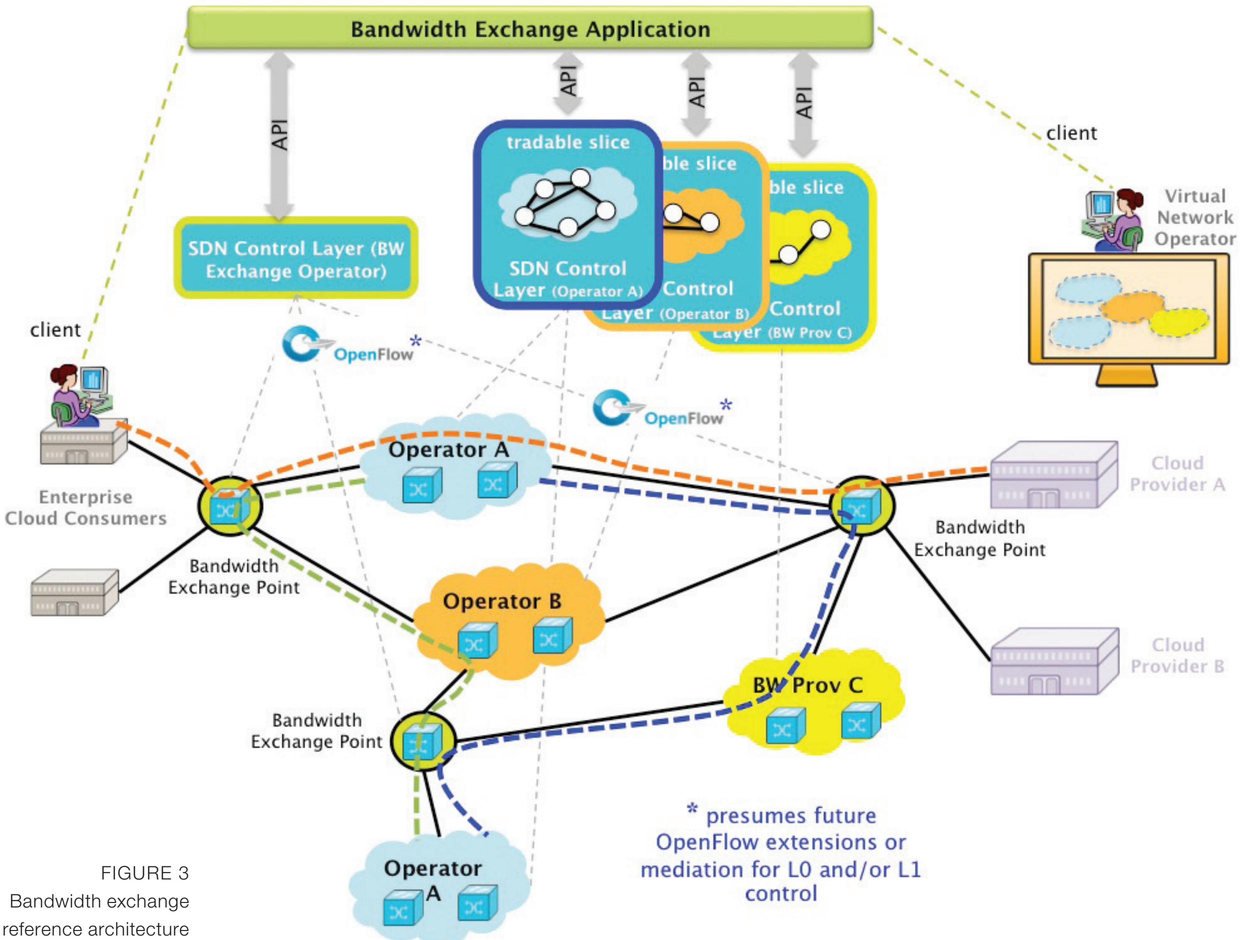

FIGURE 3
Bandwidth exchange reference architecture

SDN with OpenFlow provides essential tools for enabling automation of the processes necessary to achieve real-time bandwidth trading. It can play a key role in supplier networks as well as in the bandwidth exchange market. With bandwidth resources spanning multiple networks, canonical abstractions of available network services and bandwidth resources across multivendor environments are essential. SDN can play an important role in directly controlling elements of the network architecture, such as provisioning the bandwidth exchange point switches interconnecting these networks in this implementation model, or slicing and dynamically allocating bandwidth in the supplier's network.

Standardized representation and configurability of bandwidth flows are essential for enabling a bandwidth exchange marketplace that can uniformly manage bandwidth inventory and provision flows in a multitenant environment. With an SDN approach, suppliers in the bandwidth exchange ecosystem can easily partition their networks into "public/tradable" and "private/non-tradable" network slices.

They can then provide secure access and control of their designated tradable bandwidth to the bandwidth exchange application, where their bandwidth can be pooled with tradable resources from other suppliers. The matching of demand with supply, based on duration and temporal availability, can be achieved through the logically centralized SDN control layer's global perspective. The open northbound API to the bandwidth exchange application and the programmatic OpenFlow southbound network API facilitate the rapid, on-demand provisioning of resources and the automation of workflow processes that make bandwidth trading possible and create an efficient market that can help mitigate both undesirable bandwidth shortage and over-supply situations.

Pay for Quality of Service

Operators providing Internet access services, whether wireline or wireless, are being challenged by the volume and unpredictability of traffic they are facing. With the popularity of smart phones and tablets and the rise of connected machines such as sensors, traffic signals, and other "Internet of things" endpoints, the number of connected devices has surged past the number of people in the world. When the impact of those devices is combined with the growth in video content and communications, the volume of bandwidth is projected to grow two orders of magnitude over the next decade. Most of these devices are mobile, and services are migrating to a cloud delivery model in which the virtual machines and data stores processing and serving this traffic may be flexibly placed or moved among data centers. The resulting unpredictability of device and server locations only exacerbates the traffic volume dilemma.

Much of that traffic is coming from over-the-top (OTT) content or cloud Software as a Service (SaaS) providers. Unfortunately, operators are not generally sharing in the OTT revenue. Today, operators can either add capacity continually to deal with this growth and over-provision to deal with the uncertainty, or they can hold the line on investment and allow service to degrade. Neither of these choices is attractive. So as not to lose customers, most operators have opted for the former. But with customer willingness to pay for basic monthly access having plateaued and their cost per bit not falling fast enough to outpace traffic growth, operators need to find new ways to monetize their investment.

If the eventual alternative were degraded service, then customers who want a quality experience and/or the OTT and cloud providers that rely on the operators' networks to ensure that experience may be willing to pay for the value they receive. This requires that some traffic receive better than best-effort treatment within the operator's network. This of course means that some traffic will get worse treatment than other traffic even if it is of the same nature (application type, source, destination). Such an approach may contravene net-neutrality rules, yet we are seeing those circumvented in wireless networks where resource constraints in shared radio access networks are particularly acute. As bandwidth demand swells, wireline operators in some countries are convincing regulators to take a harder look at modifying those rules.

SDN provides tools that enable network operators to transform their networks from an undifferentiated bit-transport commodity to a value-based resource. The crux of the solution, shown in Figure 4, is to identify and distinguish flows and determine the treatment each deserves. It is then possible to allocate resources appropriately by having the SDN control layer program the forwarding table on the edge OpenFlow switches (facing the traffic source), the OpenFlow-enabled broadband access node or wireless base station (facing the customer), and potentially all OpenFlow switches along the path. Alternatively, end-to-end tunnels of different capacities could be pre-established between these edge nodes, with tables programmed to forward flows into the appropriate tunnel.

FIGURE 4
SDN-enabled QoS-differentiated and monetizable services

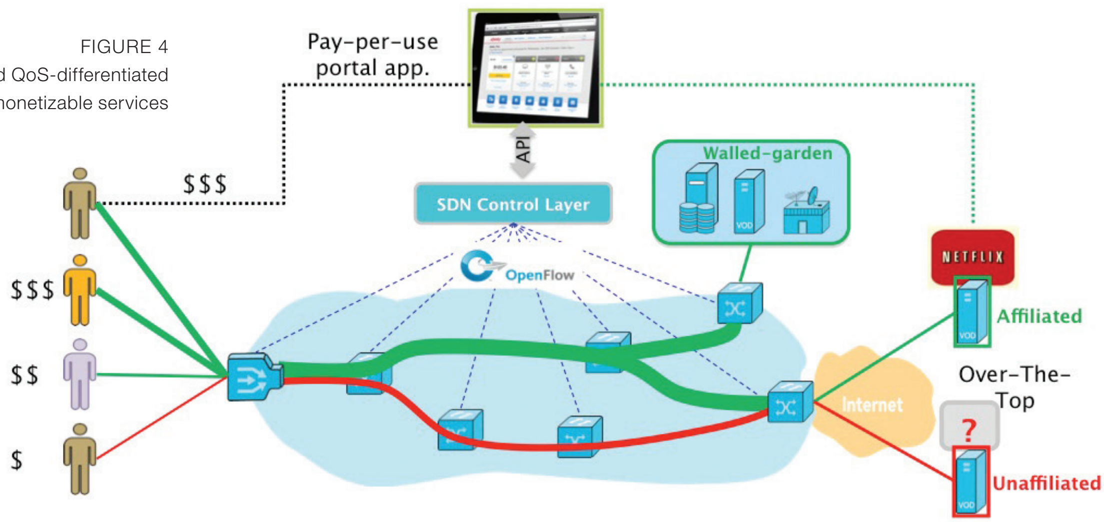

There are several ways to identify and distinguish flows. One method is to distinguish based on the content or application provider as determined by the source IP address or other data extracted from the packets. Flows from the operator's internal "walled-garden" service platform get premium treatment and are placed on a high-performance, traffic-engineered path. Flows from affiliated (partner) OTT providers (who pay some monthly fee) also get premium treatment. Generic, unidentified flows (e.g., from unaffiliated OTT sources) can be forwarded onto a best-effort path.

Alternatively, flows can be distinguished based on subscriber—for example, as identified by destination IP address. Flows from Gold subscribers (who pay the highest monthly fee) get the highest bit rate and carriage on the high-performance path. Flows from Silver subscribers get a medium bit rate (or perhaps a high bit rate up to a monthly maximum volume) as well as carriage on the high-performance path. Flows from Bronze subscribers (who pay the lowest monthly fee) get either a lower bit rate or carriage on the best-effort path.

Another approach is to distinguish flows associated with a particular session identified by Layer 4–7 information. In this scenario, any user can go to the pay-per-use portal and select a title (for a price) from an affiliated OTT provider, even if they do not have a monthly subscription. After an initial session handshake between the OTT server and subscriber device through the portal path to identify the unique flow identifier (e.g., IP Source Address + IP Destination Address + TCP Port Number), the portal can request the SDN control layer to program the flow to be redirected over the high-performance path at high bit rate. This would allow even Bronze subscribers to pay for "turbo" experience when they so desire.

The open API above the SDN control layer enables operators to create their own methods for distinguishing flows and algorithms for determining treatment, such as integration with Policy and Charging Rules Function (PCRF). The OpenFlow API below the control layer enables operators to implement Pay for Quality of Service consistently across different vendors and access technologies. And the logically centralized global view of network resources and service demands allows operators to ensure that their customers or partners get what they pay for and pay for what they get. Ultimately, this allows operators to garner the revenues necessary to support all these traffic demands more profitably.

Network Features for Pay

Due to the downward price pressure that operators are facing in their connectivity services, they have begun to offer network features and managed services as a way to drive new revenue and move up the value chain for their enterprise and small-to-medium-sized business (SMB) customers. These value-added services include managed router, managed security, and network performance assurance, among others.

Today's managed services are built on proprietary systems, often with closed and proprietary interfaces, requiring significant expertise to integrate and operationalize. Many of these managed services—such as routing and security—are based on network features built into sophisticated customer-located equipment (CLE). The sheer volume and distributed physical locations of equipment create operational challenges for the operator throughout the lifecycle of the service. This makes it difficult to develop and roll out new services, taking 12–18 months and many millions of dollars.

While some of these services have shown to be strong revenue generators, the net effect of today's model is a lengthened product development cycle that stifles innovation. An accelerated innovation cycle would allow operators to better monetize their networks by quickly rolling out new services, testing new business models, and creating a differentiated competitive advantage.

SDN provides the framework of tools to accelerate, automate, and simplify the creation, monetization, and operation of these network features-based services. For operators already offering these services, SDN can provide substantial improvements that increase revenue generation. For those operators considering the services, SDN can offer a more financially and operationally viable way to get started:

- The OpenFlow API between the SDN control software and OpenFlow-enabled aggregation switches separates the data plane from the control plane functions, enabling the operator to program CLE forwarding tables to "route" traffic based on flexible flow definitions.
- Open northbound APIs into the application and orchestration layers allow the operator to create dynamic, on-demand, and customized services by easily integrating them into back office operations and billing systems.

The following use-case example illustrates the capabilities and the benefits of using an SDN approach for these services.

VIRTUALIZED MANAGED ROUTER SERVICES

Managed router services are a big source of revenue for many operators. They provide customers with clear benefits such as Internet access, firewall protection, and L2 and L3 VPN services without an upfront capital outlay, and without requiring in-house expertise with routers or WAN protocols.

In the traditional managed router approach, the operator installs a small router on the customer premises. Typical customer premises installations also include a Layer 2 demarcation device (NID) to terminate and manage the metro Ethernet access connection to the central office. The CLE router performs control-plane functions such as route discovery and network address translation as well as data plane forwarding. The connection back in the central office terminates on an Ethernet port on an aggregation router.

The traditional approach to providing managed router services (see Figure 5) has significant costs for the operator:

- On-premises router deployment and manual provisioning
- Complex configuration, management, and upgrades of the customer-located router
- An operational burden of supporting different router versions to cost-effectively support the specific subset of service features for each customer, or else always deploying the same router even when a more basic version would suffice for some customers
- Expensive port on central office aggregation router
- Router change-out as technology churns or customer's needs change

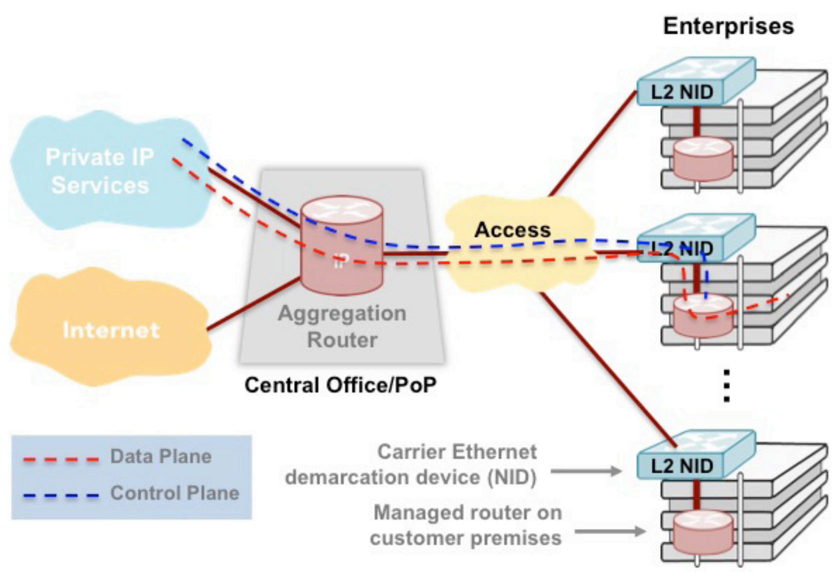

FIGURE 5
Traditional Managed Router Service requires expensive and difficult-to-manage customer-located routers and an expensive central office router port for each customer endpoint.

A better approach is to deploy only a simpler OpenFlow-enabled NID at the customer site, obviating the need to install and manage a separate router. Similarly, in the central office the expensive aggregation router can be replaced by a more cost-effective Layer 2+ OpenFlow-enabled switch (see Figure 6).

Both are controlled by SDN control software running in the central office or data center. Rather than running an embedded, vendor-specific L3 control plane extended out to the customer premises, the complex routing functions are centralized and instantiated as an SDN application on a virtual machine. New L3 flows are automatically directed to the SDN control layer and to the virtual router application that manages route discovery, network address translation, and other features of the managed router service.

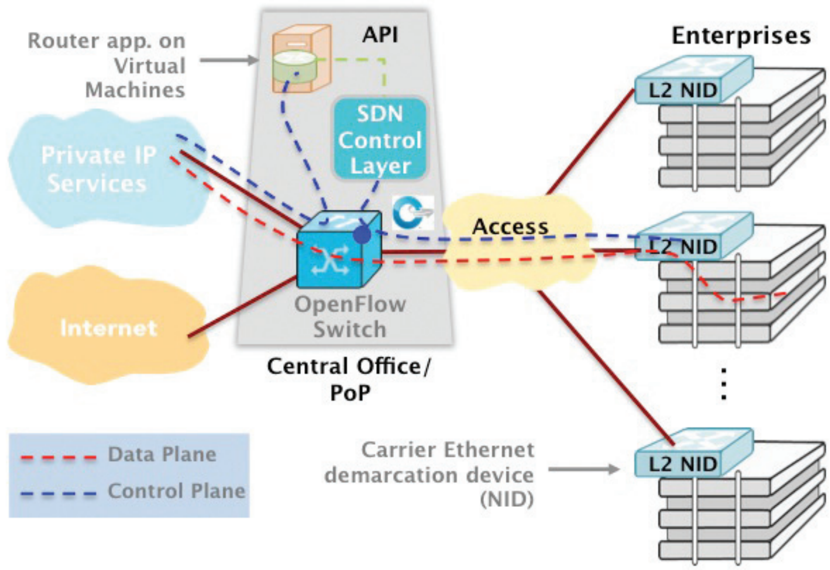

FIGURE 6
SDN-enabled virtual managed router service removes the customer located router and central office aggregation router, replacing it with a lower-cost OpenFlow switch and virtual router application running on a low-cost server.

Service delivery and management operations become much more scalable. Revenue-generating feature enhancements are installed centrally and become easily available to any customer. Compared to traditional approaches, this virtual model also:

- Eliminates the management burden and expense of the customer located routers.
- Reduces technology obsolescence at the customer premises.
- Eliminates the need for a CO router, dramatically reducing the cost of aggregation.
- Supports a pay-as-you grow CAPEX model, requiring a modest upfront investment in virtual machines to support router functions that scale with service success.
- Facilitates simplified integration into automated commissioning systems.
- Enables a simple service upsell to managed firewall and managed security services on demand and without truck roll.

Key Benefits

SDN with OpenFlow provides several benefits that facilitate network monetization.

- The ability to view and control allocation of network resources globally makes it easier to dynamically broker network services.
- An open northbound API enables customers and applications to automatically invoke network services that match their needs, making for a more fluid and mass-customized marketplace.
- Global views of both supply and demand enable analytics-based dynamic pricing. Such pricing can maximize revenues and profitability by increasing demand when resources are idle and increasing revenue per bit when resources are scarce.
- A modular, open control framework enables operators to develop their own differentiated network services and network-aware applications.

Conclusion

In this era of traffic growth outpacing cost reductions, competitive pressure reducing traditional wireline revenues and profitability, and the threat of unaffiliated OTT providers garnering a larger wallet share while network operators foot the bill, operators need new ways to monetize their network assets. SDN provides tools that enable operators to transform their networks from an undifferentiated bit-transport commodity to a value-based resource through direct linkage with—and therefore greater relevance to—revenue-generating consumer and enterprise applications. And it does so with both lower CAPEX and lower OPEX, allowing more competitive pricing. The use cases illustrated above are only examples of what is possible. With Software-Defined Networking, network monetization potential may be as expansive as the operator's imagination.

Contributors

Mitch Auster, Editor
Smita Deshpande
Mark Durrett
Chris Liou