Pipeline
Technology for Service Providers.
www.pipelinepub.com Volume 14, Issue 2

The Criticality of SDN Analytics Data

By: Cengiz Alaettinoglu

SDN technologies promise dynamic, flexible networks that can be reconfigured quickly to accommodate different business requirements. However, automation alone will not achieve this. Current SDN infrastructures lack the management intelligence to run autonomous networks effectively. SDN presents many management challenges, including loss of visibility into and control of changes occurring in the network, and the need to capture engineering know-how in SDN applications.

SDN analytics are what is required for real-time orchestration and enhanced service visibility across both legacy and SDN network infrastructures. Another layer of management is required to obtain these analytics and create truly adaptive networks.

Here's a look at the issues facing service providers, the case for SDN analytics, including use cases, and where analytics should reside. This article focuses on WAN-SDN (aka Carrier SDN), which applies the software-defined concept to wide area networks. For service providers, this includes the core, metro, and backhaul networks.

First, here’s a look at the unique challenges pushing service providers to SDN.

Supporting Unique Service Requirements and Accelerating Service Activation

Today’s service provider networks are complex, because they must support many applications: Internet access, streaming video, voice-over-IP, Layer 2 and Layer 3 VPNs, 3GPP mobile backhaul and core transport, cloud services, and more. Unlike networks of the past, many applications now run as a service on top of converged IP/MPLS packet-switching networks. These are much more efficient, scalable, and fault tolerant. However, their performance is less predictable and requires closer monitoring of service paths.

Running multiple applications on a converged network presents management challenges, because each application has unique performance requirements, growth rates, and fault-tolerance characteristics. A service provider may need to optimize the network for corporate connectivity services during the day (e.g., a financial services enterprise may be willing to pay premium prices for very short delay paths to support its time-sensitive trading application), and for over-the-top content delivery at night (e.g., Netflix and YouTube for residential customers).

Another challenge is an increase in the rate of service activation and deactivation requests. Customers are also expecting faster service provisioning times – from weeks to hours, and even seconds. For example, many service providers offer self-service portals that allow customers to request more bandwidth.

Satisfying Customer Demands Manually

Preparing the network for any service is difficult and requires significant time investment by

expensive engineers, especially as focus has shifted into the IP/MPLS network. In the past, IP/MPLS was used in the backbone of the network, and there may have been 500 routers. The growth in traffic, especially by mobile users, has forced service providers to extend IP/MPLS to the access and aggregation networks.

This has dramatically increased the number of routers that service providers must manage, often upwards of 20,000. It is extremely difficult to operate an IP/MPLS network of that size. For example, in a medium-sized service provider, five percent of the tunnels used for traffic engineering may be down at any one time. This could be more than 1,000 tunnels.

It is not possible for engineers to manually determine why these tunnels are down in a timely manner – it would take hours or even days. Making matters worse, by the time they complete the analysis, the data is no longer valid because the network has changed.

As these challenges demonstrate, optimizing an effective multi-services network is probably not achievable without automation.

Why Traditional SDN Architectures Are Insufficient

SDN can help address these challenges and streamline network provisioning. Figure 1 illustrates a simple two-tier architecture in which SDN applications control network behavior. Network devices (both physical and virtual) are not configured manually. Rather, they are programmed via southbound APIs by one or more SDN controllers or service orchestrators, which perform higher-level orchestration functions across domains and sometimes across the IP/MPLS and optical layers. The controllers provide access to applications via northbound APIs, enabling the applications to modify network behavior to meet their needs.

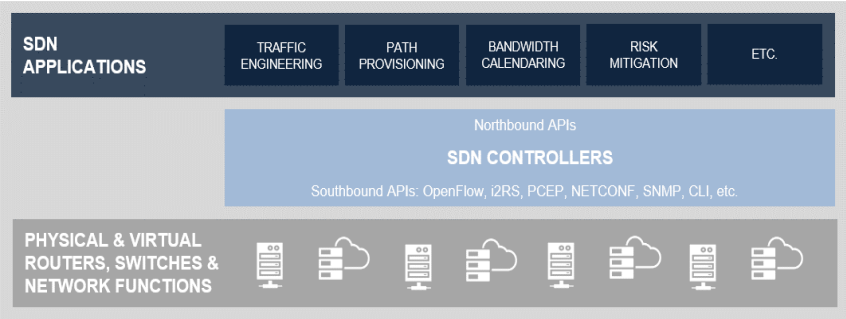

Figure 1 — Traditional Two-tier WAN-SDN Architecture

Although SDN controllers provide the means to change network configurations through software, they lack management intelligence. For example, they cannot answer basic questions such as:

- If applications and services are being rolled out without operator intervention and adequate visibility, how can network operators plan for them?
- Who or what determines if these programmatic changes should be made?
- Will the changes negatively impact existing applications and services?
- When changes are problematic, how does the operator diagnose the problem or find the root cause?

To answer these questions in dynamic networks requires analytics.

Analytics Defined

SDN management requires much data about what is happening in the network, such as IGP topology, BGP routes, traffic demands, jitter, performance, delay, and interface utilizations. The industry talks a lot about this telemetry, but it is simply the collection of this data. It represents the beginning of effective SDN management. Many big data projects overload engineers with data, but

then don't tell them what to do with it.

Analytics are the actionable conclusions drawn from the data. SDN analytics provide the visibility and management intelligence service providers need to run their automated networks effectively. They help answer in seconds the planning and troubleshooting questions that engineers may spend hours or days on using manual methods.

Two Functions for SDN Analytics

The first function of SDN analytics is to maintain management visibility into the network as programmatic changes are being made. SDN analytics should provide visibility into the devices and controllers by recording real-time telemetry from the network's control and data planes, including the routing topology, performance metrics, and traffic flow data. Recorded data helps with back-in-time forensics to identify the root cause of issues.

The second and more important function of SDN analytics is to provide management intelligence. Analytics software replicates the expertise of network planning groups, assessing the network's readiness and capacity for making significant changes, acquiring a new enterprise customer, or turning on a new service.

Once the SDN analytics software comes up with a solution based on telemetry data, the SDN controller or orchestrator can provision it in the network.

SDN Use Cases

The importance of analytics becomes even more apparent as service providers continuously create new uses for SDN. Here are a few of them.

Rapid service provisioning: To speed up service creation and activation/deactivation times from weeks to minutes, SDN analytics are a must. For example, if a customer requests more bandwidth via a self-service portal, using path computation technology will automatically generate optimized network configuration recommendations, based on the supplied constraints, to the SDN controller.

Data sovereignty: Many organizations cannot have their data leave their country's physical borders. This requires service providers to create policies that specify the devices and links that traffic can traverse. They must know which paths should be used and have recovery options. This is normally a very labor-intensive process. SDN makes automating data sovereignty protection possible, if service providers have the intelligence needed to drive automated path provisioning.

Run networks hotter: To reduce costs, SDN can make it possible to run service provider networks at around 70 percent or greater link utilizations. However, keeping utilizations high requires real-time and predictive analytics that drive automated network configuration to accommodate changing demands.

Hybrid Cloud Use Cases: Bandwidth on demand and calendaring capabilities are required to efficiently use WAN resources and support new services such as cloud backup and data-center disaster recovery needs. These use cases require the ability to record and baseline network routing, traffic, and performance data to feed machine-learning algorithms that calculate optimum network configurations.

Where Should Analytics Reside?

For analytics to be viable, they must reside in the right place in SDN architectures. A few years ago the industry thought that analytics would be in the controller. This did not happen. First, the controllers became a commodity. Vendors viewed analytics as a value-add and kept them in their offerings.

Second, the controller is a control plane device. Doing big data analytics in the controller is not advisable. Analytics could be put in applications (such as traffic engineering), but not all applications may have access to the same telemetry. One application does not know what is happening with another one, and so forth.

This is why a new layer is being introduced into the SDN architecture: An analytics and automation layer. Given the importance of SDN analytics, therefore, the traditional two-tier SDN architecture needs to be expanded to include an analytics-based orchestration layer, as shown in Figure 2. This new layer delivers both management visibility and intelligence to SDN applications.

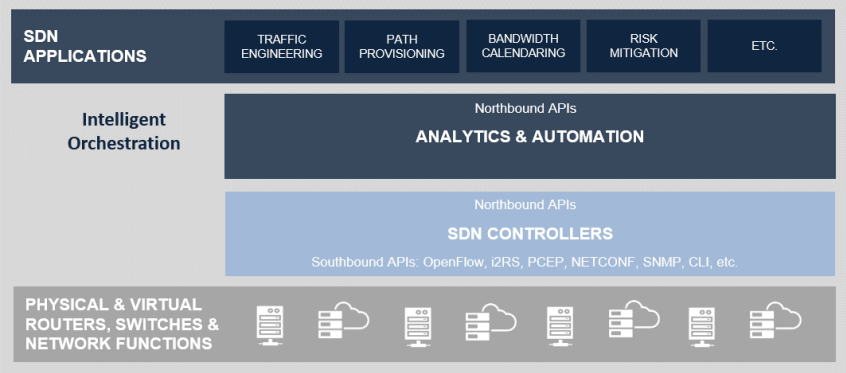

Figure 2 — Three-tier WAN-SDN architecture

The Bottom Line

When software is making decisions that engineers traditionally have made, that software needs to be powered by the same expertise and deliver transparency to human operators. SDN analytics, fed by real-time and historical telemetry, projections, and algorithms, make it possible to manage a much larger sized network and satisfy customer demands quickly via automation. These analytics should be delivered via a management layer between the SDN controller and SDN applications. Only then can service providers effectively manage their multi-service networks to increase business agility, optimally use their capital investments, and add revenue opportunities.