The Rayno Report
Innovation and Investment Intelligence

TECH STRATEGIC BRIEF
February 2014

Analytics and the Software-Defined Data Center

Introduction

The modern data center, where business services and applications are deployed through the cloud, has been the subject of an enormous amount of hype. Cloud is everything, now. Beyond the noise, however, many businesses and service providers are saddled with the challenge of managing a cloud-computing infrastructure.

The first benefits of cloud technology – to automate the provisioning of applications – has been achieved to a large degree today, though there is still more work to do. The second phase is just as important: the automation of the management, maintenance, and security of the infrastructure. Yes, eventually, we want the machine to watch over itself. It's convenient enough that humans don't have to crawl into the back of a wiring closet to provision a new customer – but then what? After the customer is provisioned, it would be nice to automate as much of the management as possible.

As virtualization – the enabling technology of cloud computing – is being extended deeper in the network to include Network Functions Virtualization (NFV) and Software Defined Networking (SDN) in data centers and telecom networks, the pieces are slowly being put into place to build a more adaptive and responsive system.

How do we get there? The path to a fully automated and adaptive computing infrastructure is pervasive software control, coupled with powerful measurement and analytics tools to provide instantaneous monitoring and feedback.

Measurement and analytics tools are the next step to making this happen. If the system can't watch itself, how does it know what to do? The real key to the value creation in network virtualization and automation will be the measurement and analytics tools put into place to monitor the system and build a self-adaptive infrastructure.

The result, in the future, will be a fully automated and adaptive data center – one that can measure performance activity and respond in real time to problems and security risks.

The Search for Visibility & Flexibility

One great analogy for a new world of an adaptive, automated computing and networking system is aviation technology. We talk about self-driving vehicles as if they are some futuristic fantasy, yet airplanes already feature sophisticated avionics and self-piloting capabilities today. Planes are capable of landing themselves, though pilots only don't always utilize the technology.

The autopilot capabilities of a modern airplane require important data inputs and measurement platforms to work together as a dynamic system. The automated technology platform of the airplane is constantly measuring functions of speed, thrust, vertical and horizontal navigation, temperature, fuel levels, electrical systems, and many other factors.

This is where data centers and networks need to go: integrated data measurement and analysis. In computing and networking, a new quiver of management tools and technologies will be needed to help manage and grow the data center and put it on "autopilot." Measurement and analytics tools will play a key role in gaining visibility inside these rapidly growing and changing computing and networking environments.

The problem to this point has been that many systems – whether they are computing systems, storage systems, or networking systems – often exist in islands, without sharing information with one another. This is where networks and cloud computing platforms may not be as integrated as airplanes – are all of the different parts of the airplane talking to one another?

One of the core elements of the next big waves of virtualization – SDN, NFV, and Software Defined Data Centers (SDDC) – is framework for more open communications between various layers of the network.

Let's take a look at how this is described visually. To get another view, take a look at the same concept as described in one of many new open-source virtualization platforms – the OpenDaylight Project. OpenDaylight is defining how a new generation of software "Controllers" – or pilots – can communicate with other parts of the network.

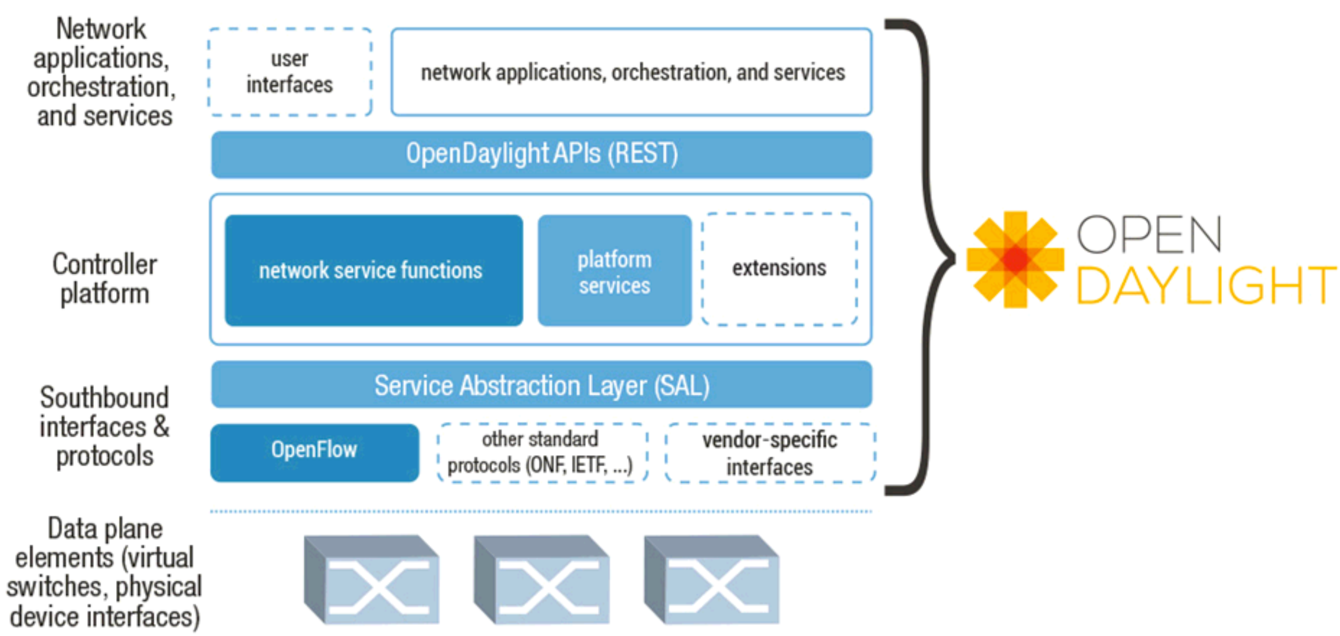

Source: OpenDaylight

Applications Programming Interfaces (APIs) are crucial to this vision, enabling different pieces of software to talk to one another. Think of how APIs fueled the world of mobile services, enabling the explosion of smartphone apps as thousands of developers were able to produce software for a specific mobile phone platform. The same kind of explosion will occur in data-center and service-provider applications, once these open application frameworks are put into place.

The way the framework is being built with NFV and SDN technology, with these APIs, the controller can connect to applications, "northbound," or back into the computing infrastructure, or the switching and routing machines at the core of the data center, known as "southbound." Let's look at another visualization of the APIs and how they enable communication between the layers of the infrastructure. Here is yet another way to visualize it: The diagram below describes this new integrated system of communication between the layers.

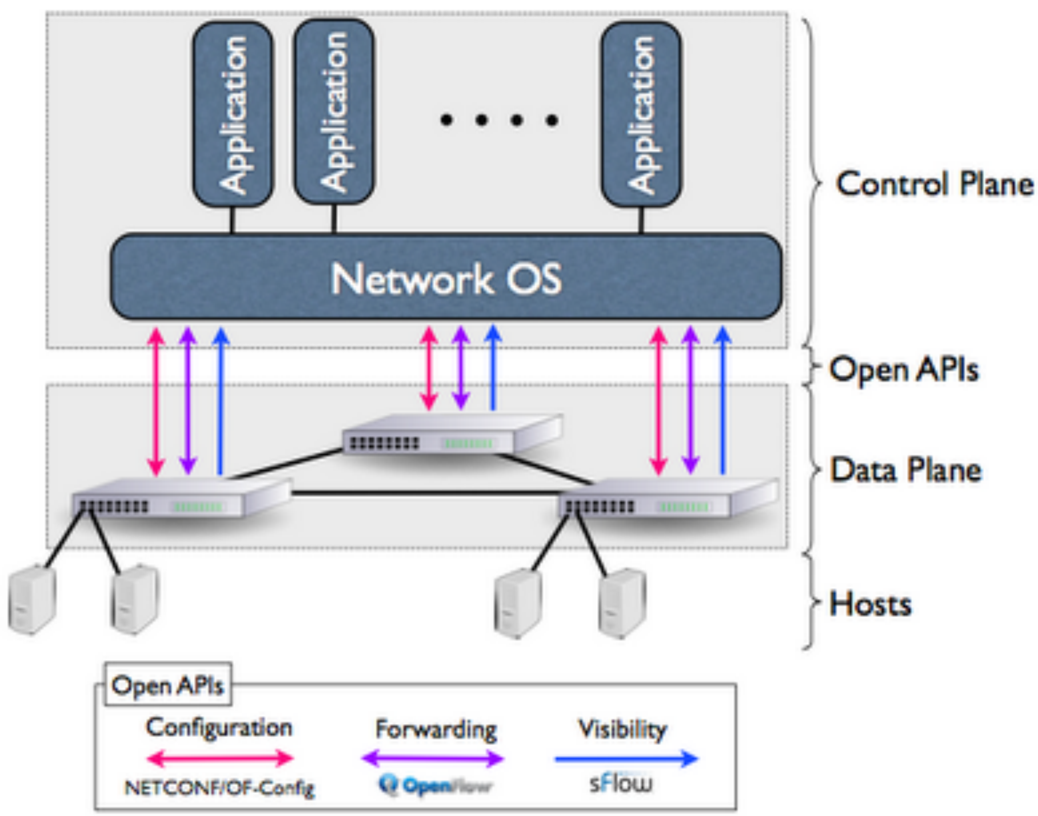

Source: InMon Corp.

So what's the big deal? The point is that as these APIs open up channels of information between the network layers, it will increase visibility extracted by measurement and analysis tools. This will produce a new breed of analytics and management applications, enabling real-time measurement and analysis of what's going through the entire computing and networking system.

Conclusion: Measurement & Control Yields Real-World Apps

So are we ready to move forward? As more open measurement and analytics tools will lay the foundation for increasingly advanced applications to build the self-aware network, how exactly will it unfold? What are the next real-world applications that will be produced by the arrival of more open, integrated data centers and networks?

Let's get back to airplane analogy. We want to move deeper into autopilot mode. There are plenty of ways to look at data centers and networks like air travel. Key functions are speed, topology (navigation), and network traffic and congestion (weather and/or air traffic control). The more inputs an avionics system has, the smoother the flight. Navigation systems can be programmed to navigate a plane around storms, or avoid other planes. But first, the information is needed.

The diagram below demonstrates the concept, as shown by InMon Corp., which provides measurement and monitoring tools. By aggregating data from all layers of the network, more comprehensive analytics tools can be built to manage the system.

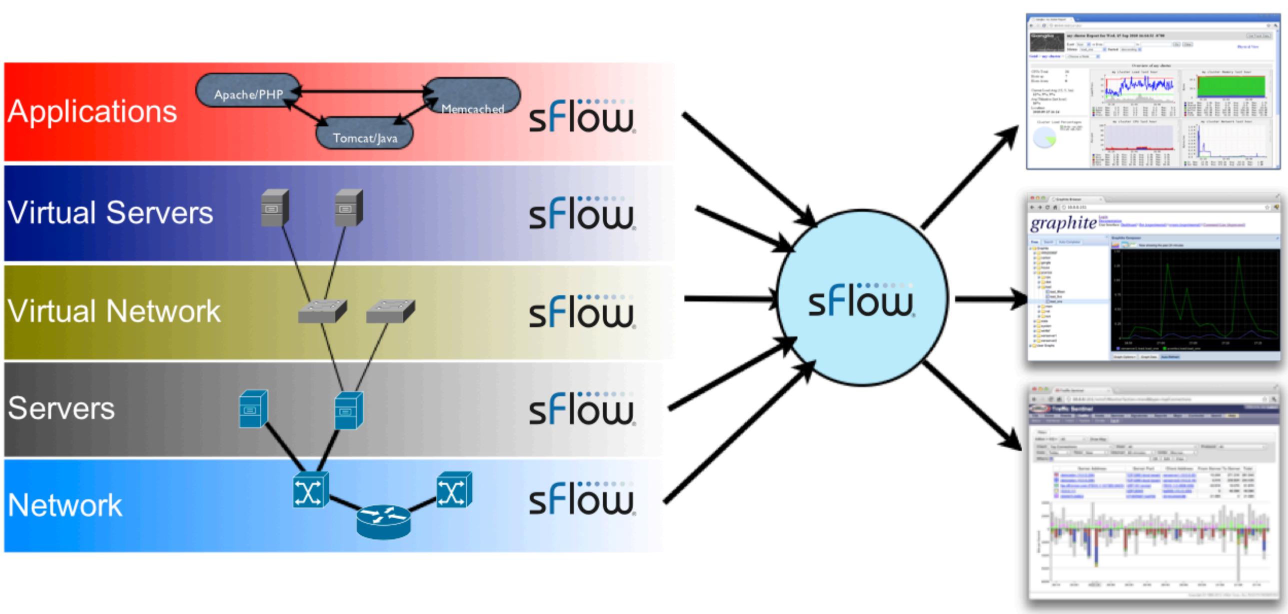

Source: InMon Corp.

Such a system could be thought of as measurement-based SDN and SDDC, in which the system is monitoring itself – across many layers and systems – and how the different applications, data, and hardware elements are interacting with each other. The key characteristics of a dynamically monitored network are this:

- Pervasive monitoring of all points in the network
- A dynamic network that's able to make its own adjustments
- Open exchange of data through APIs
- Real-time, dynamic reaction.

Once you grasp the amount of information that can be collected about the system, the possibilities open up.

Recently when in a meeting with a technology executive, I was asked, "But what are the applications of SDN?" As it turns out, they are probably infinite. It's like asking, when Apple launched its App Store, "But what kind of apps will people come up with?"

In fact, large commercial data centers are already finding many applications for open SDN and NFV-based networking.

Here are a few examples of real SDN applications, leveraging measurement and analytics, which are already in place:

Security & Denial of Service Mitigation: A well-known, global e-commerce operator data experienced a serious of vicious hacker attacks that shut down its systems. Programmers were able to define the behavior common to these attacks and place measurement devices in the system to discern the activity and program automated protective responses to the attack whenever this behavior surfaced.

Traffic Engineering: The SDN community speculates about the future potential for traffic engineering, as if it's some sort of black magic, but the reality is that it's already possible. According to its own technology paper, Google says it has implemented SDN to get a "unified view of the network" to simplify configuration, management, and provisioning. How did Google do this? It built its own routers and switches based on OpenFlow technology. All of these pieces communicate with each other across the Wide Area Network (WAN). According to Google:

> *The service collects real-time utilization metrics and topology data from the underlying network and bandwidth demand from applications/services. With this data, it computes the path assignments for traffic flows and then programs the paths into the switches using OpenFlow. In the case of changing demand or network events, the service re-computes path assignments and reprograms the switches.*

Service Chaining: The concept of service chaining is to replicate the work done in building applications for a specific service so that it does not have to be done over and over again for each new client. When tied to physical hardware, service chaining is a difficult, manual process. The advent of SDN and NFV technologies will allow this activity to be moved into software, where it takes place in the Controller.

Large service providers such as Verizon say this is a top application for SDN, but isn't Amazon already doing it? You can access nearly any physical product you need from Amazon, though it doesn't actually originate on Amazon. I like to think of Amazon as a virtual service chain for the entire retail industry.

Appendix: Selected Measurement & Analytics Companies

A number of companies are developing next-generation solutions for managing the network. Many of these are based on open standards designed to leverage data across the network and provide measurement and analytics capabilities.

This is not a comprehensive list, but rather, a selection of several examples of companies using an open-data framework to produce holistic next-generation data analysis and network management platforms.

Boundary is an operations management platform delivered as a software-as-a-service (SaaS) offering. It enables customers to monitor the entire IT infrastructure through a centralized management console that combines events, alerts, and notifications from any source in a single view. This includes the capability to provide a real-time snapshot of the logical application topology.

(www.boundary.com)

Cloud Physics produces a SaaS data analytics platform that enables managers to gain operational insights into virtualized data centers. In a unique approach, Cloud Physics allows the community to share data in a common place and supplies APIs so that users can use a collection of data from the community to build their own applications for data analytics. For example, users can build applications for simulation and resource-management.

(www.cloudphysics.com)

InMon is the inventor of sFlow, an open measurement agent that can be adopted in networking, server, and storage gear so that managers can collect information about what's happening on all levels of the network. sFlow has been adopted by most of the major networking players, including Cisco, Dell, Juniper, and HP, as well as silicon vendors such as Broadcom, Intel, and Marvell. In addition to providing sFlow-based network monitoring tools, InMon markets software that provides an integrated picture of network, storage, server, and communications performance, including Traffic Sentinel and sFlow RT. It is compatible with all the major virtualization platforms including VWware, KVM, Xen, Linux, and HyperV.

(www.inmon.com)

GigaMon produces a series of network-monitoring devices, known as GigaVUE, that help manage, analyze, and secure the network. The hardware and software solution provides traffic visibility, switching, and passive monitoring instrumentation to enable security, forensics review, application performance, VoIP QoS, uptime, and other network management tasks.

(www.gigamon.com)

Sideband's management platform provides single point of visibility and control to the networking team. The company uses real-time management technology, including its eXtensible Response Engine (XRE), to process real-time network traffic and allow for responses to network performance issues and threats. Packaged as a "virtual device," the platform operates at Layer 2-7 and provides networking monitoring up to 40Gbps with millisecond reaction to live traffic.

Splunk (SPLK) was a 2013 IPO. Its flagship product, Splunk Enterprise, is a data analysis platform. Splunk monitors the data from IT systems and technology infrastructure and analyzes it, giving managers insight into potential problems and solutions. Splunk's goal is to analyze the machine data from end-to-end infrastructure and avoid service degradation or outages.

(www.splunk.com)

Editor & Publisher: R. Scott Raynovich

(406) 582-5886

scott@rayno.com