Whitepaper

DECODING SOFTWARE DEFINED NETWORKING

DECODING SDN

For the past year, software-defined networking (SDN) has been the buzz of the networking world. But in many ways, networking has always been defined by software. Software is pervasive within all of the technology that impacts our lives and networking is no different. However, networks have been constrained by the way software has been configured, delivered and managed—literally within a box, updated monolithically, managed through command lines that are throw-back to the days of mini-computers and DOS in the 1980's.

THE CHALLENGES WITH NETWORKING SOFTWARE

Networking software has been a drag on innovation across our industry. Because each network device must be configured individually—usually manually; literally from a keyboard—networks can't keep pace with the on-the-fly-changes required by modern cloud systems. Internet companies like Amazon or Google that dedicate hundreds of engineers to their cloud systems have built their own solution to network configuration but this is not a reasonable approach for most companies to build their private cloud. As virtualization and the cloud has revolutionized computing and storage, the network has lagged behind.

In the service provider world, carriers struggle to configure and manage their networks. Like Google, they too have built operational support systems to configure their networks but these systems are often 20+ years old and they are crumbling from the burden placed upon them by networking software. For a service provider, the network is their business, so they must look to networking vendors to introduce new capabilities in order to enable new business opportunities. Here again, networking software is failing the industry—it is developed as a monolithic, embedded system and there is no concept of an application. Every new capability requires an update of the entire software stack. Imagine needing to update the OS on your Smartphone every time you load a new application. Yet that is what the networking industry imposes on its customers. What's worse is that each update often comes with many other changes—and these changes sometimes introduce new problems. So service providers must carefully and exhaustively test each and every update before they introduce it into their networks.

WHAT IS SDN?

Enterprise and service providers are seeking solutions to their networking challenges. They want their networks to adjust and respond dynamically, based on their business policy. They want those policies to be automated so that they can reduce the manual work and personnel cost of running their networks. They want to quickly deploy and run new applications within and on top of their networks so that they can deliver business results. And they want to do this in a way that allows them to introduce these new capabilities without disrupting their business. This is a tall order but SDN has the promise to deliver solutions to these challenges. How can SDN do this? To decode and understand SDN, we must look inside networking software. From this understanding, we can derive the principles for fixing the problems. This is what SDN is all about.

Here are six principles of SDN with corresponding customer benefits:

1. **Cleanly separate** networking software into four layers (planes): Management, Services, Control, and Forwarding—providing the architectural underpinning to optimize each plane within the network.

2. **Centralize** the appropriate aspects of the Management, Services and Control planes to simplify network design and lower operating costs.

3. **Use the Cloud** for elastic scale and flexible deployment, enabling usage-based pricing to reduce time to service and correlate cost based on value.

4. **Create a platform** for network applications, services, and integration into management systems, enabling new business solutions.

5. **Standardize protocols** for interoperable, heterogeneous support across vendors, providing choice and lowering cost.

6. **Broadly apply SDN principles** to all networking and network services including security—from the data center and enterprise campus to the mobile and wireline networks used by service providers.

THE FOUR PLANES OF NETWORKING

Inside every networking and security device–every switch, router, and firewall—you can separate the software into four layers or planes. As we move to SDN, these planes need to be clearly understood and cleanly separated. This is absolutely essential in order to build the next generation, highly scalable network.

NETWORK PLANES

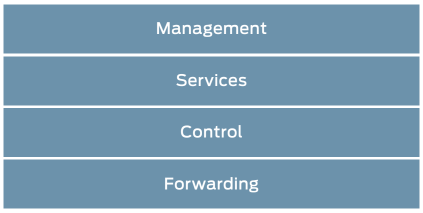

Forwarding. The bottom plane, Forwarding, does the heavy lifting of sending the network packets on their way. It is optimized to move data as fast as it can. The Forwarding plane can be implemented in software but it is typically built using application-specific integrated circuits (ASIC's) that are designed for that purpose. Third party vendors supply ASIC's for some parts of the switching, routing, and firewall markets. For high performance and high scale systems, the Forwarding ASIC's tend to be specialized and each vendor provides their own, differentiated implementation. Some have speculated that SDN will commoditize switching, routing, and firewall hardware. However, the seemingly insatiable demand for network capacity generated by thousands of new consumer and business applications creates significant opportunity for differentiation in Forwarding hardware and networking systems. In fact by unlocking innovation, SDN will allow further differentiation from the vendors who build these systems.

Control. If the Forwarding plane is the brawn of the network, Control is the brains. The Control plane understands the network topology and makes the decisions on where the flow of network traffic should go. The Control plane is the traffic cop that understands and decodes the alphabet soup of networking protocols and ensures that the traffic flows smoothly. Very importantly, the Control plane learns everything it needs to know about the network by talking to its peer in other devices. This is the magic that makes the Internet resilient to failures, keeping traffic flowing even when a major storm like Sandy brings down thousands of networking devices.

Services. Sometimes network traffic requires more processing and for this, the Services plane does the job. Not all networking devices have a Services plane—you won't find this plane in a simple switch. But for many routers and all firewalls, the Services plane does the deep thinking, performing the complex operations on networking data that cannot be accomplished by the Forwarding hardware. Services are the place where firewalls stop the bad guys and parental controls are enforced. They enable your Smartphone to browse the web or stream a video, all the while ensuring you're properly billed for the privilege. The Services plane is ripe for innovation.

Management. Like all computers, network devices need to be configured, or managed. The Management plane provides the basic instructions of how the network device should interact with the rest of the network. Where the Control plane can learn everything it needs from the network itself, the Management plane must be told what to do. Today's networking devices are often configured individually. Frequently, they are manually configured using an esoteric command line interface (CLI), understood by a small number of network specialists. Because the configuration is manual, mistakes are frequent and these mistakes sometimes have serious consequences—cutting off traffic to an entire data center or stopping traffic on a cross-country networking highway. Service providers worry about backhoes cutting fiber optic cables but more frequently, their engineers cut the cable in a virtual way by making a simple mistake in the complex CLI used to configure their network routers or security firewalls.

TODAY'S NETWORK DEVICES

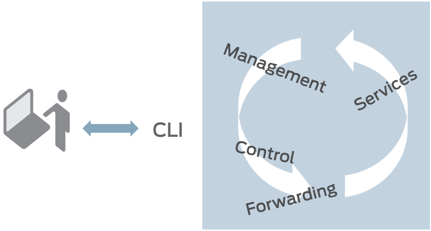

While the Forwarding plane uses special purpose hardware to get its job done, the Control, Services, and Management planes run on one or more general purpose computers. These vary in sophistication and type, from very inexpensive processors within consumer devices to what is effectively a high-end server in larger, carrier-class systems. But in all cases today, these general purpose computers use special purpose software that is fixed in function and dedicated to the task at hand. That inflexibility is the root of the issue that has sparked the interest in SDN.

If you crawled through the software inside a router or firewall today, you'd find all four of the networking planes. But with today's software, that networking code is built monolithically without cleanly defined interfaces between the planes. What you have today are individual networking devices, with monolithic software, that must be manually configured. This makes everything harder than it needs to be.

CENTRALIZATION

So if today's networking software is the root of the problem, better software is the solution and that's where SDN comes in. How do we go from today's networking software to a modern architecture? We start by looking at the way cloud providers build their software. Amazon, Google, and Facebook use racks of industry-standard, x86 servers running software that is designed to scale-out by adding more servers as the need for capacity increases. The use of industry standard, x86 hardware combined with scale-out software is how modern, highly available systems are built.

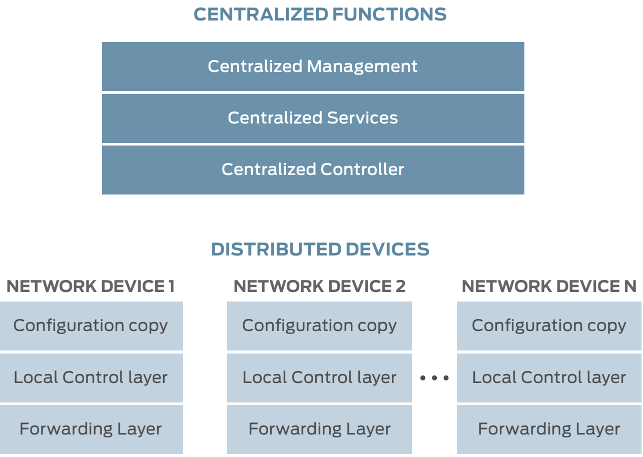

Unlike most cloud applications, networks are inherently decentralized. That's really what networks are all about—moving data from one place to another. So while Facebook can run in a small number of huge data centers, networks are distributed —throughout a data center, over a campus, within a city, or in the case of the Internet, across the entire planet. That's why networks have always been built as a collection of separate, self-contained, individually managed devices. But centralization is powerful; it is a key principle for SDN and it's very appropriate to apply centralization to networking software. However, you can't take this too far. Centralization only makes sense within a highly-connected, contained geographic area—for example, within a data center, throughout a campus, or in the case of a service provider, across a city. Even with this centralization, network devices themselves will remain distributed and they must have local intelligence.

When you add the concept of centralization to networking software, the four planes move around a bit. Regardless of the number of distributed devices, you'd like to manage the network as a system and Centralized Management does that job. When you centralize management, it becomes the configuration master; all of the devices keep just a copy. This is very similar to the way publications work with our Smartphones and tablets. If you run the New York Times app on your iPad, it pulls down today's edition. During the day, it keeps checking for updates and downloads them when they appear. This is analogous to how Centralized Management works; the full truth lives in the center and only a copy of the configuration data is stored on the networking devices.

Services have historically been implemented within each networking and security device but with SDN, Services can move to the center and are performed on behalf of all devices. However, this only makes sense in a highly-connected, contained geographic area. If you're accessing the Internet from your Smartphone, you want to get onto the Internet highway from the city you're in, not someplace half-way across the country.

When SDN enters the picture and some things are centralized, the changes to the Control plane are the most complex. The control plane is the cop that directs the traffic. The way the Control plane works is each networking device talks to the networking devices they directly connect with. They tell each other what they know about the network. Think about it as an electronic version of smoke signals. Each device passes information about the network on to the next device. This works incredibly well in the highly connected, networking world. Many years of work across the entire networking industry ensure that networks continue to do their job even when things go wrong. When a major router goes offline, there is a buzz of chatter between the networking devices as they scurry to restructure their view of the network—and keep you connected.

But sometimes having a central, birds-eye view of traffic also makes sense. That's where the Centralized Controller comes in. The Centralized Controller has a broad view of the network and can connect things together in a way that optimizes the overall traffic.

Forwarding is one plane that always stays decentralized in an SDN world. This makes sense because Forwarding actually moves the data—and this is by-definition decentralized.

GETTING FROM HERE TO THERE

So how do we go from today's fully decentralized networks to a new world where some things are centralized with SDN? You can't start with a clean sheet of paper because networks are actively running and must continue to function as SDN is introduced. SDN is like a remodel; you need to do it one step at a time. Like most remodels, there is more than one way to get to the SDN result, but here is a reasonable set of steps to reach the goal:

Step 1: Management is the best place to start as this provides the biggest bang for the buck. The key is to centralize network management, analytics, and configuration functionality to provide a single master that configures all networking devices. This lowers operating cost and allows customers to gain business insight from their networks.

STEP 1: CENTRALIZE MANAGEMENT

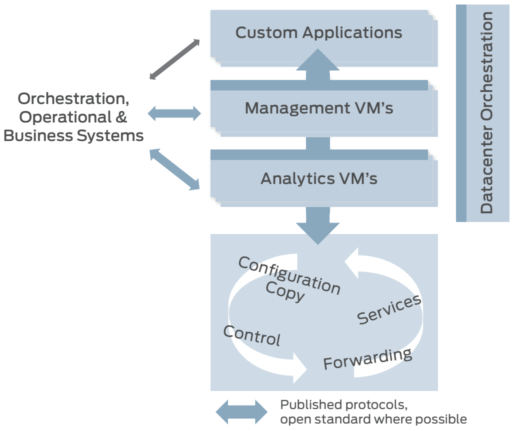

Centralizing Management does several things, each of which provides significant value. You start by creating a Centralized Management system. Similar to cloud applications, this centralized management system is packaged in x86 virtual machines (VM's) running on industry standard servers. Those VM's are orchestrated using one of the commonly available orchestration systems such as VMware's vCloud Director, Microsoft System Center, or OpenStack.

In the case of the service provider, their operational and business systems connect to the centralized management VM's which configure the network. Similarly within a data center, that same data center orchestration system (VMware vCloud Director, OpenStack, etc.) can now directly manage the network.

Configuration is performed through published API's and protocols; where possible these protocols are industry-standard. As SDN is still nascent, industry standard protocols are still emerging but it is very important that moving forward these standards get created.

Networking and security devices generate huge amounts of data about what is happening across the network. Much can be learned by analyzing this data and like other aspects of business, "Big Data" analytics techniques applied to networking and security data can transform our understanding of business.

Pulling management from the network device into a centralized service provides the first step to creating an application platform. Of greatest urgency is simplifying the connection to the operational systems used by enterprises and service providers. But as this platform takes shape, new applications will emerge. The analytics provides insight into what's happening within the network, enabling better business decisions and new applications which will dynamically modify the network based on business policy. Centralized management enables changes to be performed quickly—enabling service providers to try out new applications, packages and plans, quickly expanding those that work and dropping those that don't. In fact, like other new platforms we've seen over the years, the possibilities are endless and the most interesting applications will only emerge once that platform is in place.

Step 2: Extracting Services from network and security devices by creating service VM's is a great next step because Services are an area that is terribly underserved by networking. This enables network and security services to independently scale using industry-standard, x86 hardware based on the needs of the solution.

STEP 2: EXTRACT SERVICES

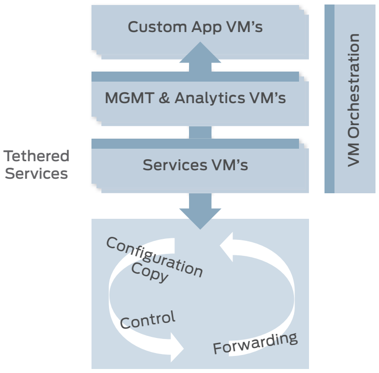

Creating a platform that enables services to be built using modern, x86 VM's opens up a whole new world of possibility. For example, the capacity of a security firewall today is completely limited by the amount of general-purpose processing power you put into a single networking device—the forwarding plane is faster by an order of magnitude or more. So if you can pull the security services out of the device and then run them on a bank of inexpensive x86 servers, you dramatically increase capacity and agility.

As a first step, you can tether, or connect these services back to a single networking device. You can put the x86 servers in a rack next to the networking device or they can be implemented as server blades within the same networking device. Either way, this step opens up the possibilities for a whole new set of network applications.

Step 3: Creating a Centralized Controller is a big step forward. The Centralized Controller enables multiple network and security services to connect in series across devices within the network. This is called “SDN Service Chaining”—using software to virtually insert services into the flow of network traffic. Service chaining functionality is physically accomplished today using separate network and security devices. Today’s physical approach to service chaining is quite crude; separate devices are physically connected by Ethernet cables; each device must be individually configured to establish the service chain. With SDN Service Chaining, networks can be reconfigured on the fly, allowing them to dynamically respond to the needs of the business. SDN Service Chaining will dramatically reduce the time, cost and risk for customers to design, test and deliver new network and security services.

STEP 3: CENTRALIZE CONTROLLER

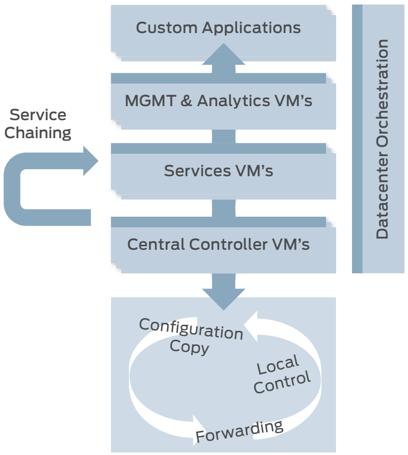

Here are several examples of SDN Service Chaining. The first example *(fig.1)* is a cloud data center connection between the Internet and a web server. In this example, the Stateful Firewall service protects the application and the Application Delivery Controller provides load balancing of network traffic across multiple instances of the web server. SDN Service Chaining allows each service within the chain to elastically scale based on need; the SDN Service Chain dynamically adjust the links within the chain as instances of the services come and go.

The second example *(fig.2)* is between two components of a cloud application; in this case between the web server and the middle-tier application VM’s. The traffic between these application components must be isolated from other traffic within the cloud data center and the load needs to be balanced across application instances with an Application Delivery Controller service. With SDN Service Chaining, all of this is done in software—the chain forms a virtual network where the endpoints are the virtual switches within the hypervisors of the servers that run the application VM’s. The SDN Service Chain dynamically adjusts the links in the chain when the data center orchestration system moves a VM from one physical server to another. Of course, there is still a physical network underneath the SDN Service Chain but it does not need to be reconfigured when changes are made within the SDN Service Chain.

fig.1

SDN SERVICE CHAIN EXAMPLE 1

Datacenter cloud application Internet to Web Server

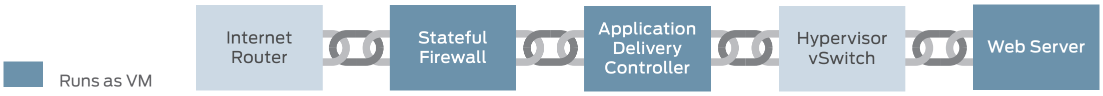

fig.2

SDN SERVICE CHAIN EXAMPLE 2

Datacenter cloud Web Server VM to Mid-Tier VM

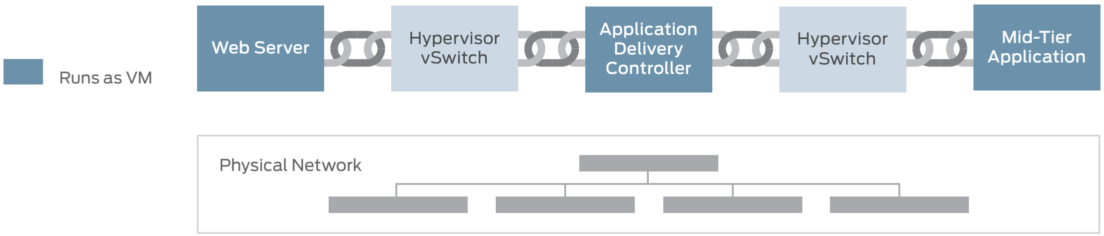

fig.3

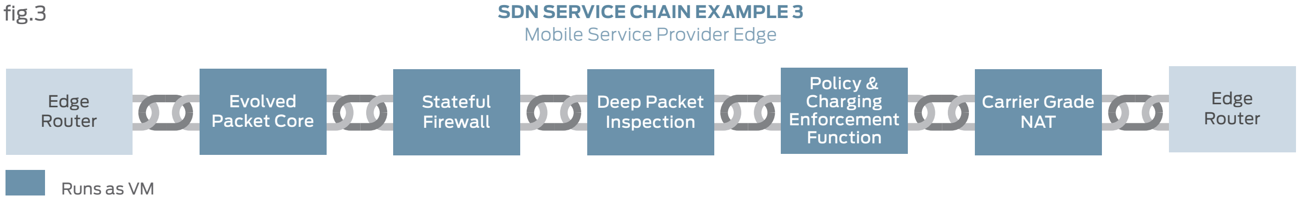

While the first two SDN Service Chain examples apply to the cloud data center, the third example is in a completely different domain—the mobile service provider edge. In this case, the network traffic is coming from a cell phone tower; it moves through an edge router and then a set of processing steps are performed in series. The Evolved Packet Core extracts the Internet Protocol (IP) sessions from the network tunnels connected to the cell tower base stations. Immediately, this traffic is analyzed and protected by a Stateful Firewall. Deep Packet Inspection is used to determine traffic patterns and generate analytics information. The Policy Charging and Enforcement Function applies subscriber polices such as enhancing the quality of service for premium subscribers. Finally, as the traffic heads out to the Internet, Carrier Grade Network Address Translation (NAT) provides the traffic with an IP address.

In the third example ***(fig.3)***, both end-points of the SDN Service Chain are edge routers. While the specific application in the mobile service provider edge is very different from the data center, the SDN Service Chaining architecture is exactly the same.

SDN Service Chaining dramatically increases the flexibility of service deployment. Most significantly, it allows network and security devices to be managed and upgraded independently from the services within the SDN Service Chain. SDN Service Chaining enables services to be treated like applications on your Smartphone—the network can still operate when new services are installed. This is a huge advance over the current situation where these upgrades are highly disruptive, thus requiring immense care and planning.

SDN Service Chaining is a new innovation and thus extensions to existing protocols and new protocols will need to be defined. As these emerge, it is important that they are established as industry standards to enable multi-vendor interoperability.

Step 4: The final step of optimizing network and security hardware can proceed in parallel with the other three. As services are disaggregated from devices and SDN Service Chains are established, network and security hardware can be used to optimize performance based on the needs of the solution. Network and security hardware will continue to deliver 10x or better Forwarding performance then can be accomplished in software alone. The combination of optimized hardware together with SDN Service Chaining allows customers to build the best possible networks.

STEP 4: OPTIMIZE THE HARDWARE

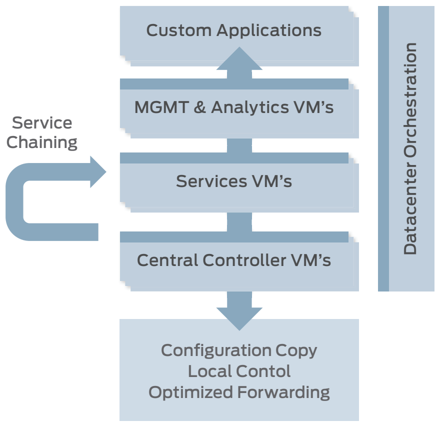

The separation of the four planes helps to identify functionality that is a candidate for optimization within the Forwarding hardware. This unlocks significant potential for innovation within the ASIC's and system design of networking and security devices. While an x86 is general purpose, the ASIC's within networking devices are optimized to forward network traffic at extreme speeds. This hardware will evolve to become more capable –every time you move something from software into an ASIC, you can achieve a 10x performance improvement or more. This requires close coordination between ASIC design, hardware systems, and the software itself. As SDN becomes pervasive, the ability to optimize the hardware will create lots of opportunity for networking and security system vendors.

SUMMARY

SDN is a major shift in the networking and security industries. Its impact will extend far beyond the data center and is thus actually much broader than many predict today. SDN will create new winners and losers. We will see new companies successfully emerge and we'll watch as some incumbents unsuccessfully struggle to transition. But like any major industry trend, the customer benefit is real and we've now reached a tipping point where the technology shift is inevitable.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC and EMEA Headquarters
Juniper Networks International B.V.
Boeing Avenue 240
1119 PZ Schiphol-Rijk
Amsterdam, The Netherlands
Phone: 31.0.207.125.700
Fax: 31.0.207.125.701

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

2000510-001-EN Feb 2013

♻ Printed on recycled paper