PICA8®
WHITE BOX SDN

WHITEPAPER

SDN Services at the Customer Edge

White Box SDN – Foundation for Strategic Advantage

SDN Services at the Customer Edge

White Box SDN – Foundation for Strategic Advantage

Businesses and consumers are always clamoring for more services, more bandwidth, and lower prices, and service providers large and small must find ways to deliver. Subscribers are used to dialing up new capabilities on their smartphones with the click of a button, and they want Internet services to have the same ease of use. For their part, service providers want to enable customer self-provisioning, new services, and automated provisioning not only as competitive differentiators but also as a way to reduce costs and speed time-to-revenue or time-to-payback for new services. Software-defined networking (SDN) is a key to service agility and self-service provisioning for service providers and their customers.

How Does SDN Help?

Simply put, SDN helps service providers monetize their networks. By implementing policy-based control over the network, SDN enables very rapid implementation of new services in ways that would be impossible under today's Layer 2/Layer 3 networking models.

Network administrators interact with software on an SDN controller to set policies and spin up services, and the controller implements those policies and services on network switches. By disaggregating hardware and software, SDN creates new levels of service agility. It also enables the use of white box switches, which save CAPEX when compared with traditional edge switching and routing solutions.

Meter, match, and act are the three steps SDN undertakes to execute tasks in a policy-driven network. SDN enables the metering of traffic conditions, application and user behavior to match those conditions against a set of pre-defined criteria and then to act on the match according to a policy. Part of a policy framework is to pre-set conditions that are metered against.

This embedded network intelligence can enable customers to drive the network themselves. One application might be a customer dashboard as depicted in Figure 1, where customers can select the amount of bandwidth, the duration, and the latency level. For example, this would enable a business customer to throttle up bandwidth for three hours for a backup to the off-site data center or cloud, and then throttle the bandwidth back down to conserve spending. With SDN, these kinds of automated operations give service providers tremendous agility once set up, without any effort on their part to scale and enable this self-service / utility mindset.

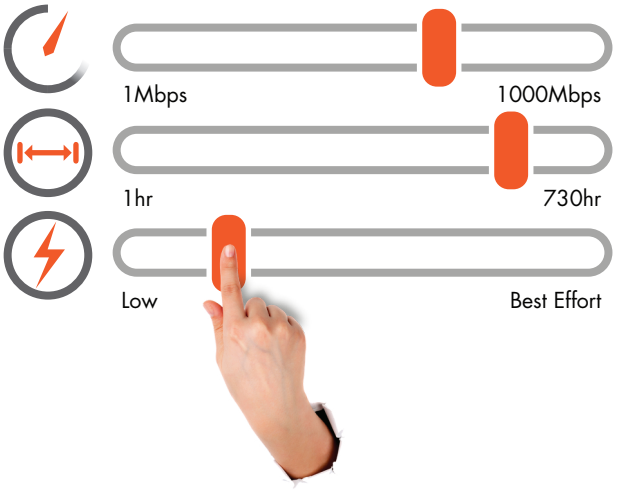

Figure 1. Customer-driven service portal

Disaggregation is the Key

Disaggregation is the key to network abstraction and flexibility. By freeing themselves from legacy models of tightly coupled hardware and software, service providers can speed provisioning and automate processes, thereby saving money and improving service agility for their customers.

Servers, PCs, and most smartphones have disaggregated hardware and software, and it's time for this trend to overtake the network. With a disaggregated network, the network OS becomes the abstraction layer, and white-box switching hardware reduces CAPEX and OPEX. Here's the difference:

Suppose a customer wants a new service, and your switches were implemented as traditional, fully integrated stacks. If the new service required a specific feature, the provider would first need to upgrade the network, and this could take weeks (or potentially months) to procure, install, and configure the new equipment. With a disaggregated hardware/software model, implementing the new service can be done in less than 5 minutes – it's a simple matter of deploying new software. More and more service providers are seeing this flexibility as critical to their business success.

Service Injection at the Edge

There are three service injection points, or service edges, in a service provider network: the business customer's hosted data center, the business customer premises, and the last mile to the consumer. At these service edges, we can drop in a SDN-enabled white box switch to implement SDN and enable self-provisioning networks to better support new applications for these customers.

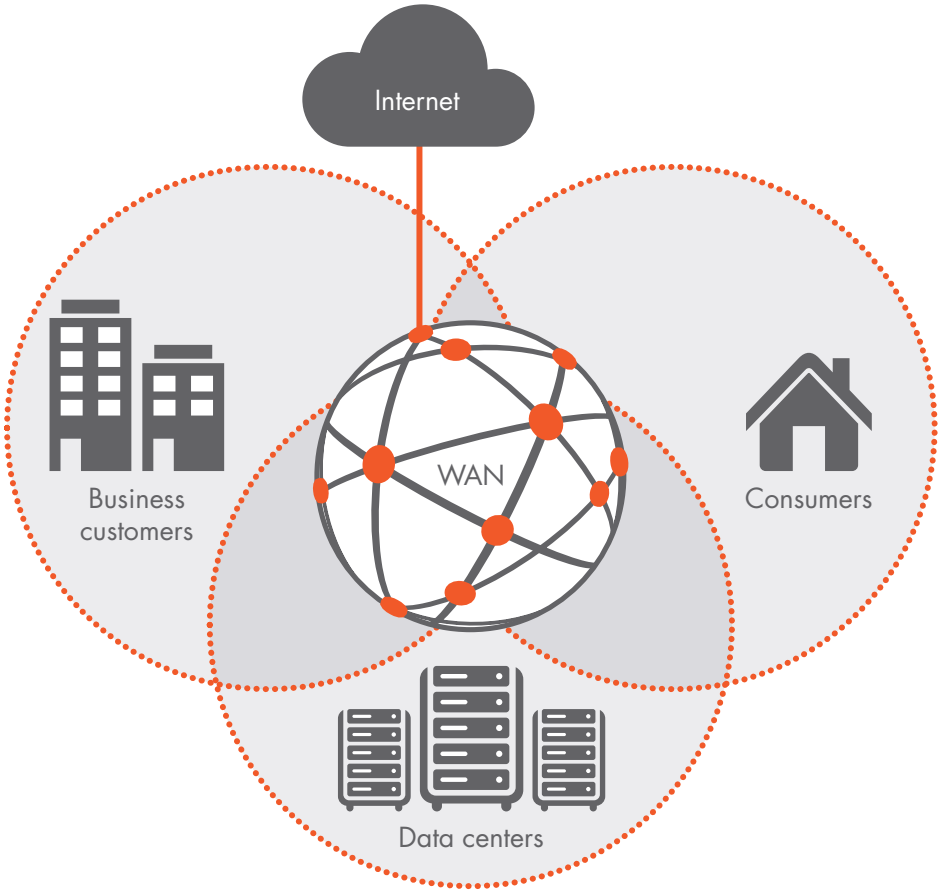

Figure 2. Services in a white box SDN-driven network

The prime examples of these SDN-driven services are called out in Table 1.

Business Need	Service	SDN Features
Connectivity – Smart pipes	• Liquid bandwidth • Metered access • QoS on demand	• OpenFlow metering and action matched rules • OpenFlow control of MPLS headers • Interoperable with leading controllers such as ONOS, ODL, HP VAN, and Ryu
Dynamically provisioned VPNs	• On demand disaster recovery (DR) • Setting up multi-tenant buildings • Hybrid Cloud services	• VXLAN, NVRE, L2/L3oGRE • Network virtualization and orchestration integration • OpenFlow control of BGP
IT operations and configuration management	• Self provisioning • Automation / DevOps • API Integration	• RESTful, Linux and OpenFlow APIs • Network automation: Chef, Puppet, Salt

Table 1. SDN-driven services

The applications vary based on location and end customer. For business customers, the typical desired applications are VPNs for simplified disaster recovery (DR), liquid bandwidth, and self-service bandwidth provisioning. These are implemented in the network devices at the business data center and business edge.

For consumers, we can look at telecommuters, home-based businesses, and power users who access the Internet via fiber to the home (FTTH). All of these want self-provisioned bandwidth and easy-to-implement VPNs for connectivity to suppliers, customers, and the corporate office. These are implemented in the network devices at the customer premises.

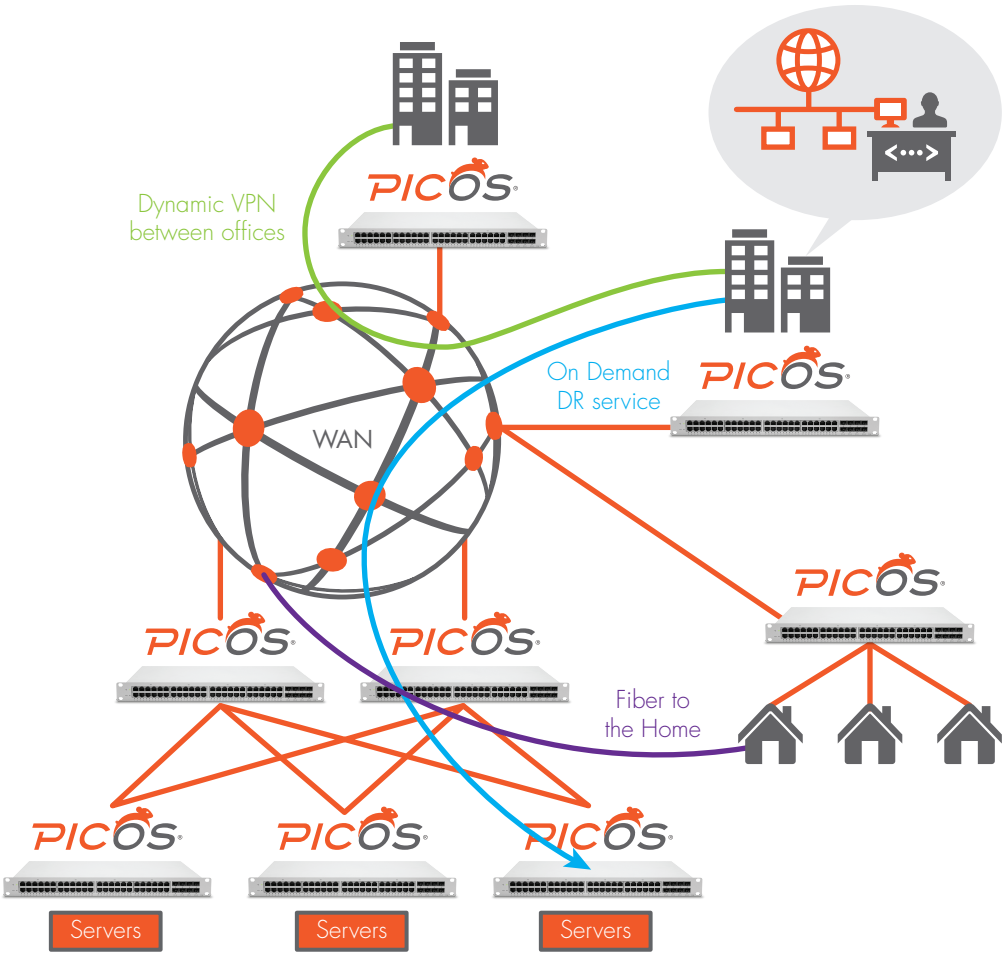

Figure 3. Illustration of three SDN-driven services

Consider an example of compute bursting. This involves liberating new compute resources with the click of a button, setting up a VPN and grabbing some idle CPUs in the data center with pay-for-what-you-use billing.

There is also an example for on-demand DR. The data center terminates the DR tunnel the customer set up – it ends up in the ISP's storage cloud. This enables customers to do routine backups without the cost of leased lines – the customer can use a portal to create a VPN to the service provider's data center, do the backup, and then turn off the VPN when finished. Services are based on how much bandwidth and storage you used, rather than on a fixed monthly leased line cost.

In each of these service injection points, you need a network that can 1) support lots of diverse applications (voice, video, wireless, big data, IoT , etc.), 2) support change management very quickly (DevOps and automation), and 3) be flexible enough to support new technologies and protocols (disaggregation and SDN).

White box networking is the key to supporting the requirements for these new ISPs. And since these services go all the way out to the consumer in many cases, you need white box and SDN in all parts of the network…in your data center, at your Internet edge, and all the way to the customer edge.

Enabling SDN Technologies

To enable liquid bandwidth, QoS on demand, self-configured VPNs and other services for customers, the service provider implements SDN to build this idea of meter, match and act into the network.

For enterprises, they want to be able to leverage their IT efficiently (read: elastic and self service) across a spectrum of on-premise and in-the-cloud services. In this hybrid cloud model, they want to be able to create logical networks, share data and information easily and securely across geographies, and get access to differentiated services when they need them (e.g. traffic engineering, application acceleration, monitoring, and security).

For service providers, they want all of the same things that the enterprises do, with the ability to monetize services, and without any additional burden on their existing IT operations, staff, and budgets.

To do this, customers should consider three areas of SDN to make their decisions: controllers, overlays, and the ability to scale.

CONTROLLERS

In any SDN-driven implementation, the controller plays the role as the 'brains'. It is responsible for the policies, and communicating with the network infrastructure to implement those policies. It also has the programmatic interfaces that allow the operators to customize and provision the applications to manipulate the network.

There are different SDN controllers with different communication protocols on the market today. OpenDaylight, HP VAN, Ryu, ONOS, OpenFlow, VMware NSX, Midokura MidoNet are just a handful of examples.

The network infrastructure needs to be flexible and programmatic to support as many of these approaches as possible.

OVERLAYS

An overlay is a logical network that enables you to create paths and connections on top of (and in many cases, regardless of) the physical connections between the end points. More importantly, overlays are a critical construct because they enable network operators to create more virtual subnets. Subnets in turn support multi-tenancy connections; the overlays are virtual pipes, delivering not just bits but the ability to have virtual machines and workloads move between locations.

One overlay approach is VXLAN. A big reason for this is the laundry list of vendors that have backed it – Cisco, Arista, Broadcom, and of course, VMware (based on the capabilities of their NSX controller) just to name a few. One of the reasons VXLAN was introduced was to address the problem of limited logical scale and to create Layer 2 adjacencies across different IP networks. It all sounds great – particularly if you have infrastructure that understands VXLAN and can behave as a VXLAN Tunnel End Point (VTEP).

The other key overlay approach is MPLS via Labeled BGP. VXLAN is an option for providers, but the downside is that it's a relatively new protocol. It might require new equipment to support VTEP functionality, and it will definitely require education and training on how to build networks with VXLAN.

Labeled BGP is a perfectly viable solution. Service providers have extensive experience and tools to solve these problems across the WAN. They can use MPLS to establish tunnels within and between datacenters, and use Labeled BGP as the signaling mechanism to exchange the MPLS labels between BGP peers.

Providers building MPLS tunnels with protocols like Labeled BGP can use this protocol as a familiar tool set to deploy new and differentiated services for their clients

Regardless of the approach, the network infrastructure should have a support a wide range of overlay protocols and techniques to maximize flexibility.

SCALING SDN

One of the key questions that arises when service providers consider implementing SDN is, "Will it scale?" If we examine OpenFlow as an example, old-line switch vendors will warn that OpenFlow only supports up to 2000 flows on the Trident II chipset, which is the most widely used chip in network switches today. Clearly, 2000 flows are not enough to scale a network.

SDN vendors like Pica8 have overcome this limitation by making the Ternary Content Addressable Memory (TCAM) much more efficient, and by combining the TCAM with the Forwarding Information Base (FIB). This architecture supports over 200,000 flows – enough to scale any network.

Conclusion

Service providers need to compete for customers, and those that can offer differentiated services quickly are going to win the race. SDN is critical for service agility and the ability to monetize services, and white box networking delivers the hardware/software disaggregation while lowering CAPEX. By using a switching OS and an array of white box switching hardware options, service providers can deploy infrastructure at the service edges of their networks that will drive rapid rollouts of differentiated services and higher revenues.

Pica8 Software Inc.
Corporate Headquarters

1032 Elwell Court, Suite 105
Palo Alto, California 94303 USA
650-614-5838 | www.pica8.com