THE CASE FOR VIRTUALIZING NETWORKS

Contents

Executive Summary

Traditional hardware-bound network infrastructure is brittle and inflexible, but network virtualization technology is available to allow an end-to-end virtualization of key datacenter infrastructure.

The network infrastructure has been historically the most hardware-bound part of the datacenter. A history of single vendor dominance has prevented the level of innovation and squelched competing disruptive forces necessary to separate the intelligence and software from the devices and hardware. IT staff have accommodated the trend by selecting single-vendors for their entire network infrastructure thereby avoiding the risks of inserting incompatible gear into their network. But there is an array of pressures to rethink this hardware-bound approach to networking and new drivers around network virtualization:

1. The proliferation of server virtualization (virtual machines) and storage virtualization into the production datacenter has put pressure to do the same on the network side. Although virtualized workloads can be provisioned quickly, they still need to be assigned network properties, such as IP addresses, virtual LANs (VLANs) etc., to start communicating. In the age of the cloud, setting up the network manually for the virtual machines is just not acceptable. Network admins are increasingly faced with the dilemma of "lead, follow, or get out of the way" with respect to the accelerating the level of network virtualization.

2. The virtualization architects who have risen from the ranks of system administrators have been taking on more and more of a networking role. They are innately drawn to increased levels of virtualization on the network side. Their virtual machines' NICs have been virtualized for over a decade, they have been using virtual switches for years, and they have been deploying distributed logical switching and routing capability in recent years. The appeal of higher order abstraction of network infrastructure and network functions is a great draw for these staff and these organizations.

3. Pressures from public (off-premise) IaaS cloud services are putting pressure on the entire IT team to be more agile and responsive to the needs of the business units. Public IaaS is also setting new precedents on the line-of-businesses ability to control their own resources and have visibility into their own resource consumption. The business unit is expecting the rapid provisioning, reduced costs, easy scaling, tenant control, and isolation that they are provided by public cloud services.

4. The networking industry is seeing the introduction of multiple models that are increasingly decoupling network hardware from the software. The availability of white box network gear and independent software for such purposes is too much cost savings to be ignored and is helping to drive the phenomenon.

5. The availability and acceptance of network virtualization software that is platform independent and "overlays" the existing network infrastructure allows users more choice to build, run, and manage their networks separate from the hardware. A network admin maintains the underlying networking hardware and sets policies for the provisioning of network overlays. Tenants can self-provision networks in the overlay based on the policies set by network administrators. This reduces workload for network administrators and improves the provisioning time in the process. Network virtualization not only removes bottlenecks in the traditional networking technologies but also improves scalability as well as the security of networks by providing required levels of isolation in a multi-tenant cloud.

The drivers and technologies exist to deliver on the promise of network virtualization are readily apparent. But what are the key variables and differentiators the network staff need to consider in their journey to transform the network? What are the key functional considerations for achieving scalability, performance, and resiliency? How does one achieve the simplicity of an overlay approach without being hamstrung by the limitations of the physical underlay? The answer stems from deploying overlay virtualization by distributing control to the edge of the network… away from the core (to minimize traffic) and where the applications are (for better control).

What is Network Virtualization?

Virtualization in IT infrastructure generally implies the abstraction of resources into logical constructs thereby allowing much more flexible resource usage independent of hardware. Each time a level of abstraction has been introduced in the IT industry, the user tends to benefit with greater flexibility. Turning infrastructure into virtual infrastructure allows infrastructure to be provisioned more quickly, allows greater utilization of hardware, and allows optimization of existing resources. These tendencies have driven compute and storage-side virtualization for years.

Network virtualization carries some of these core value propositions associated with compute/storage virtualization. Additionally, network virtualization technology is delivering value propositions previously beyond capability of other virtualization technology. Firstly, network virtualization today can create entire networks and goes beyond the creation of virtual gadgets (virtual switches, virtual routers, virtual machines, virtual LUNs, virtual disks, virtual appliances). Secondly, network virtualization stands to actually deliver performance advantages whereas virtualization overhead may have implied the opposite in the past.

These seemingly contradicting goals (of laying down a virtual overlay and enhancing performance; of achieving increased agility and enhancing control) can be achieved by combining the advantages of overlay virtualization with distributed networking.

Why Virtualize Your Network?

a. Meet Competition demands on Network Responsiveness and Costs

With the increase in the number of virtualized workloads, together with the bottlenecks inherent in traditional networking technologies, both enterprises and service providers are recognizing the need to deploy SDN in their datacenters.

Deploying an SDN solution not only enables enterprise and service providers to reduce time to market for their applications, but also improves the customer experience and allows the enterprise or provider to be more competitive. Figure 2 below shows the results of a survey conducted by InformationWeek that confirms the biggest selling points for an SDN solution.

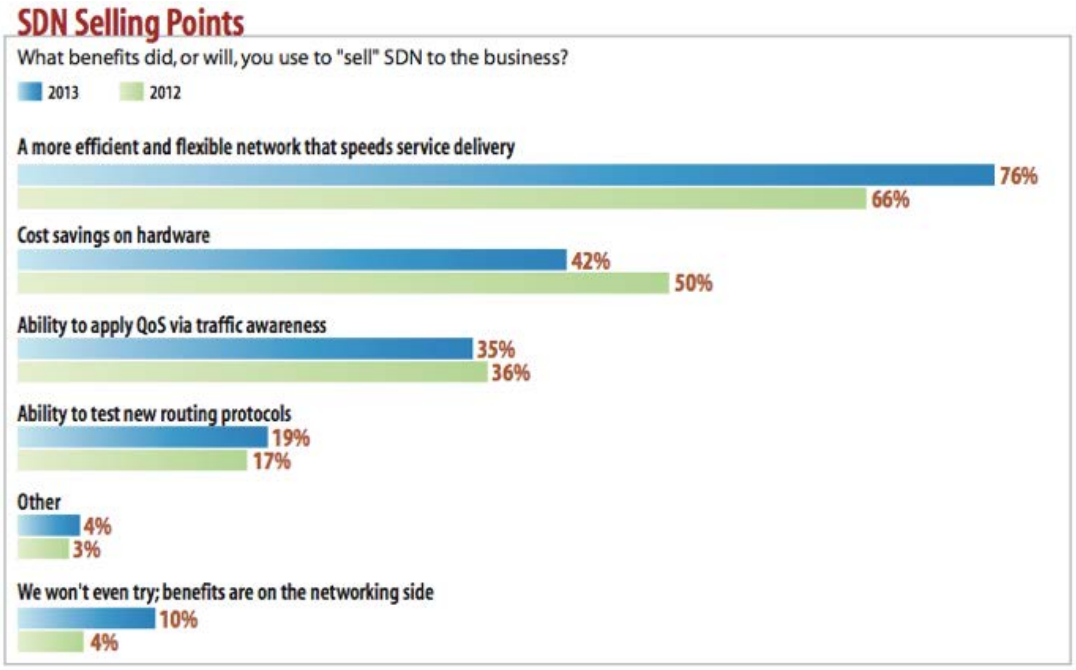

Figure1[1]

1. 116 respondents in July 2013 and 116 in July 2012 at organizations with, or planning to have, SDN in production. Data: InformationWeek – software defined networking survey of business technology professionals.

b. Server Virtualization Forces You to Reconsider Your Network

Network administration and System Administration have become intertwined. In recent years, IT has moved towards a more software-centric environment where the underlying hardware is de-coupled from the operating system. Control of the overall system that once resided in the hardware has moved "up" into the layers of software. Virtual machines, with their agility and flexibility, are mainstream for server operations. In a world where speed equals survival, being able to instantly set up complex and repeatable environments is essential. This movement of control from hardware to software not only has reduced the overall operational complexity but has also brought agility and flexibility to the IT environment.

In the old world, when servers were physical machines, VLANs were good enough to isolate the traffic without having to worry about its limitations (max. of 4096 VLANs). With virtualization, network isolation must be provided to VMs in addition to the physical machines running these VMs. This means network functions extend outside of the physical network infrastructure and into the server boxes that run the virtual machines. In order to provide network isolation to these VMs, network as well as cloud administrators face challenges of greater complexity and abstraction.

A further consequence of virtualization, especially for service providers, is the need to rethink their implementation of network "services", such as routing, firewall and load balancing. These services largely run on hardware appliances that are very costly and hard to scale. These services running on expensive hardware no longer makes sense, especially if they aren't meeting the need to be more agile and scalable at low cost. Hardware appliances cannot be deployed quickly, which means that services cannot keep up with increasing demand. CIOs are looking to provide network services to end users on-demand, drive additional business revenue and cut cost over time.

Hypervisors enable VMs to be moved en masse from one physical machine to another. This allows for machine maintenance and upgrades, machine replacement and general reorganization of the physical layer with no downtime required. However, while the hypervisor may guarantee that the VMs keep running smoothly on the destination machine, IT agility is incomplete unless the network properties are also guaranteed to be exactly the same on the destination machine. Network connectivity, addresses, network sessions etc., everything just has to work. Network overlays provide network connectivity and allow for the migration of the network properties associated with the VMs in motion to achieve the required agility.

c. Challenges in Providing Isolated Networks in a Virtualized Platform

1. VLAN Limitations

As VMs became mainstream in data centers, VLANs became the popular technology for network isolation. Ask most network administrators about how to provide network segmentation to VMs, and an immediate answer would be: "VLANs."

Historically, VLANs were easy to understand due to their presence in the physical network. The way they might be applied to VMs was easy to explain and fairly easy to implement. The argument was that VLANs work with VMs in the same way that they work with physical machines. If a network interface is configured with a VLAN ID, the traffic that belongs to this VLAN can be delivered to the devices on the VLAN. People easily grasped the idea that there was a "switch" inside a physical machine where virtual NICs for VMs get connected and traffic could be isolated in accordance with the VLAN configurations. For years, the concept of VLANs has enabled network engineers to secure and manage their networks through segmentation.

As virtualization keeps evolving, limitations to the VLAN approach become increasingly apparent and even more critical. Service providers withmany users on shared physical infrastructure, as well as big enterprises running private clouds, may hit the limit of 4096 VLANs. This number may sound large; however, it is easy to breach this limit on networks with multiple customers running hundreds of VMs operating in a multi-tenant network. When users start using networks as a service, they feel it is very easy and cheap to create more network segments (VLANs), which naturally dramatically accelerates the required bandwidth of the segments..

Big enterprises and service providers have multi-tenant clouds with regulatory and compliance requirements making network isolation a must-have feature. IT requires multiple isolated networks within their own environment such that the required number of isolated networks exceeds the number of customers and those tenants must have the right level of control of only their network. For service providers to comfortably operate their businesses at scale, they require more network isolation than VLANs alone can provide.

As a protocol of layer 2 networking, VLAN are subject to the constraints of how layer 2 networks functions. In Ethernet, the switches have to learn the MAC addresses to transmit traffic to the right destinations when new nodes are connected. This learning process is very expensive, time consuming and constrains the scalability of the network (in this case, scalability refers to the number of entries and the speed of convergence). In order to sustain this, physical switches have to follow the state of each VM, constantly learning and unlearning new addresses. While there are switches that have evolved to handle large numbers of MAC entries and that are capable of converging the network at sufficient speed, these switches are extremely cost prohibitive.

Using only layer 2 technology also puts resilience at risk. Traditionally the Spanning Tree Protocol has been used to bring redundancy to the network and avoid network loops. However network administrators have to be extra careful to avoid network loops that could crash the whole network. MLAG is a technology that allows link aggregation with multiple switches to overcome the limitations of Spanning Trees. While MLAG has a following among network administrators, implementation is highly proprietary and vendor specific.

2. Limitations of Virtual Appliances

In a traditional network, a network admin deploys physical appliances, such as routers, firewalls and load balancers, to meet the requirements of the system. Appliances are wired together to apply specific services to the traffic that needs them. In this way, appliances provide value-added services to basic connectivity.

With increasing pressure for rapid deployment, that is, “networks with a click,” the challenge becomes how to provide services in such a way that management of these services can easily be delegated to the end user, virtual network by virtual network.

The first generation of solutions, which is still with us, is the “Virtual Appliance.” On the positive side, virtual appliances themselves are VMs that have network functions installed in the image and are relatively easy to deploy in a virtualized environment. As with physical appliances, users may “wire up” the necessary virtual appliances in their system, configuring the network topology as desired. Many procedures may become automated aspects of the management system of the virtualized platform. This brings agility and, to some extent, an “on demand” flavor to the services that the end user may require.

The problem with virtual appliances is that they are not easily scaled. Architectures become complex, single points of failure abound, and many benefits of the virtual environment are lost in the implementation.

When network functions are packaged into virtual appliances, and these appliances are connected to VMs, everything is fine as long as the traffic moves smoothly through the appliances. But what happens if the appliances get deleted or if the physical hardware that runs the virtual appliances goes down? Network and System administrators who wish to build robust systems must avoid such single points of failure.

The centralized architecture of virtual appliances also introduces performance bottlenecks to the environment. When specific traffic is always destined to go through a particular virtual appliance, that appliance forms a potential bottleneck on network bandwidth as a whole. The bottleneck may not be obvious in a small environment, but will become a highly undesirable problem when you expand the system. Once the problem arises, there are no easy solutions. In the virtual world, more than ever, scalability is constrained by the fundamental choices you make at the outset.

The Power of the "Overlay" Network

Traditional networking technologies are coupled tightly to physical boxes and virtual machines. For maximum efficiency, we must minimize this entanglement with physical boxes and traditional networking technologies. To do this, we need to promote networks into a new world of "click to create." This new world of networking is a high-visibility world, with powerful and programmable tools and on-demand services. It is a world of great flexibility and elasticity. This is a world where unlimited virtual networks can be created with a click and then monitored at a very deep level. All we need from hardware is a simple but powerful backplane. Between the ends of the physical network, traffic becomes entirely encapsulated, in a purely virtual layer.

Because traffic is encapsulated as it moves between the ends of the physical network, the overlay network is sometimes referred to as a "tunnel". A header is placed in front of the traffic before it leaves the initial endpoint. The header maintains validity for only as long as the traffic is tunneling through the physical network. Upon reaching the destination endpoint, the header is removed and the receiver sees the contents as it exited from the source. In other words, the traffic in the virtual

network is completely independent from how the physical network is wired. It is also not constrained by any proprietary technology that may be used in the physical layer.

Network isolation through tunneling is not the only benefit provided by overlay networks. They also overcomes VLAN constraints, which allow only for a maximum of 4096 isolated networks, and the need for switches to learn MAC addresses each time a new VM is created. Overlay networking can use tunneling protocols that far exceed the address space of VLANs, allowing for more isolated networks. The traffic in the virtual network does not touch the state of the physical network, because the hardware is not required to interact with the contents of the encapsulated traffic. The virtual network topology is completely decoupled from the hardware underlay. The only MAC addresses that the physical switches need to learn are those of the physical devices between and including the endpoints (for example, the hosts running the hypervisor and cloud-controller software). Using network overlays also gives network administrators the option to use inexpensive commodity hardware for the physical layer.

In an environment where VMs are created and destroyed very frequently, the process for spinning a VM goes through network administrators. This can lead to long provisioning times. With network overlays, network administrators can focus on just maintaining the hardware underlay and network policies for an organization. The tenants or end users themselves can provision networks in the overlay. This shortens the process to provision a VM by an order of magnitude.

a. One time hardware setup – virtual changes anytime

With network functions decoupled from specific boxes, there are significant improvements in network agility and flexibility of networks. Because virtual network are managed separately from building and managing the physical network, problems of scale are made orders of magnitude simpler.

By contrast, when the physical network and the virtual network are tightly coupled, network administrators have to manage complex, entrenched and non-scalable physical configurations as they struggle to satisfy the demands coming from the virtual network. Virtual networks are configured in a matter of a few mouse clicks, while the physical networks have administrators literally walking around, plugging cables and administering boxes.

In the overlay scenario, network administrators can focus on deploying simple, robust and scalable hardware, building a physical backplane that can be managed on its own track. Everything above this backplane is encapsulated traffic that only needs to be delivered end to end.

Provisioning the virtual network takes place in a software environment, on its own track. Network administrators can create virtual networks on demand and provide connectivity and network services to the end user in a matter of minutes.

Overlays are the innovation that unifies the marriage between physical networks and VMs. The fresh challenge that becomes possible thanks to the overlay approach is to offer a truly agile way of providing network services to virtualized devices.

b. OpenFlow or Overlay Networks

In a traditional router or a switch, the forwarding plane and the control plane exist on a single device. OpenFlow, an open standard implemented by various switch vendors, separates these two functions. With the OpenFlow model, high-level routing decisions are moved to a centralized controller. OpenFlow allows the use of any switching or routing device from a vendor that supports OpenFlow and offers a lot of flexibility to customers. However, it has its drawbacks. When an OpenFlow switch receives a packet it has never seen before, and for which it has no matching entries, it sends this packet to the controller. This increases network latency and affects overall system performance. OpenFlow also requires you to manually configure devices each time you add or remove them from a network.

An overlay network allows running software-based networks on top of any hardware network underlay. Controllers for an overlay network are distributed. A distributed architecture enhances scalability and improves network performance. Overlay networks are hardware-vendor agnostic and allows addition and removal of network devices from a network without the need for any manual network configuration.

c. The Distributed Control Plane

Any network virtualization solution provides virtual networks to end users by means of a centralized or a distributed control plane running on top of a powerful overlay. The distributed controllers are distributed at the edges and allow high scalability and high performance of the overlay networks. A centralized controller, on the other hand, also provides high scalability and high performance to the overlay networks, but given the nature of this design, can represent a single point of failure (SPoF).

In distributed environments that employ REST APIs, the control plane also provides access to 3rd party network services and the entire topology becomes programmable from end to end.

In modern network deployments, fault tolerance is more critical than ever. Embracing decentralization more strongly than its competitors even in overlay-based networks, decentralized controllers are designed specifically to avoid the SPoF and bottleneck risks of centralized controllers (such as fabric controllers based on OpenFlow).

Due to its distributed architecture, decentralized controllers demonstrate very high intelligence at the edge of the network. Computational power at the edge is leveraged to the maximum extent possible. Controller agents installed on every hypervisor node process all the network traffic arriving into the virtualized environment. This allows you to leverage your existing processing resources and avoids the need to purchase costly processing power elsewhere in the network environment.

When an agent receives a packet, it first simulates the virtual topology configured by the user. This enables it to modify, forward or drop the packet, before the packet has travelled deeper into the physical network. For example, in the virtual topology, if there you've created a filter that drops specific traffic, the controller agent will drop the packet at the edge, thereby preventing unnecessary traffic from reaching the underlay network. This culling of unnecessary traffic makes the system become dramatically more efficient at scale.

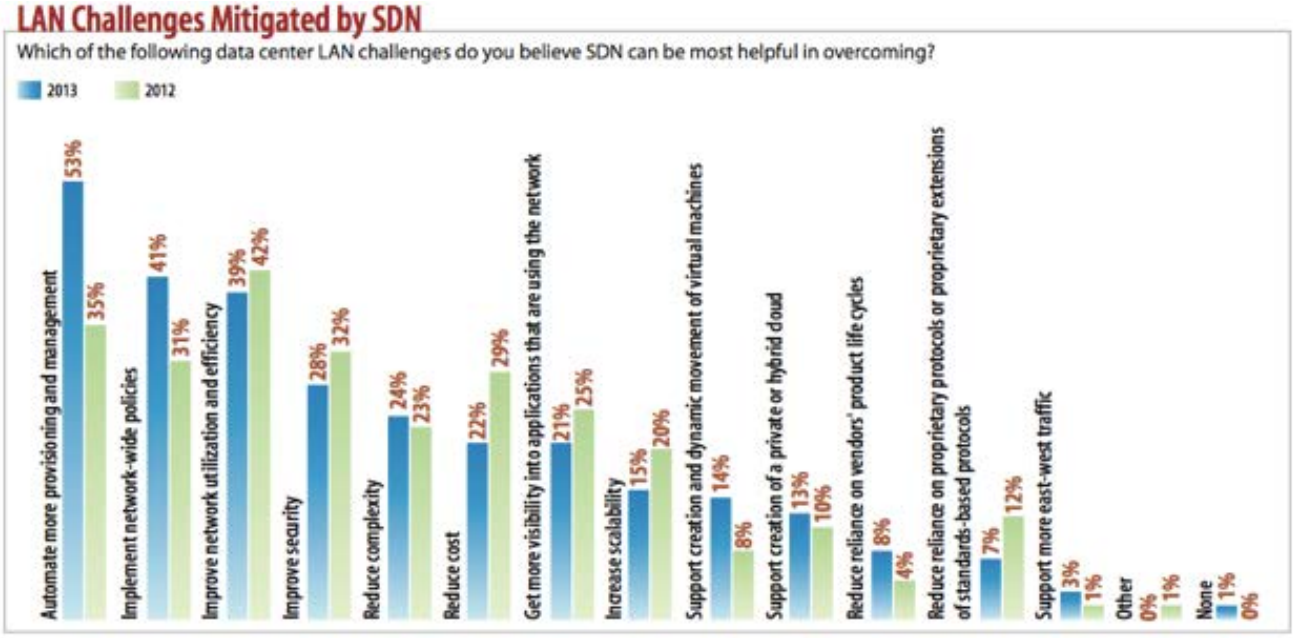

Figure 2[2]

A key value of distributing controller functionality in this way is the traffic on the underlay is minimized. Less "traffic tromboning" suggests performance advantages as your network scales. Additionally, it becomes intuitively clear the advantages of the approach to pushing controller functionality to the edge when you consider the insertion of higher order functionality. Consider all the advantages of distributing network controller and then other network services to the edge with the application. It is not hard to see how network services around policy management, load balancing, security management can all be optimized when they reside at the edge close to the application without cluttering the core network.

2. Data: InformationWeek – software defined networking survey of business technology professionals.

Conclusion

A combination of overlay-based network virtualization with distributed control is a key determinant of a completely automated and secure public/private cloud environment. Network virtualization brings advanced networking capabilities to the cloud and eliminates barriers to cloud adoption by making networks dynamic and agile.. SDN will help service providers and enterprises to:

1. Provision networks for applications in a matter of seconds – thereby reducing time to market.
2. Avoid limitations of physical network topologies and VLANs, thus providing better flexibility and scalability than physical networks.
3. Provides greater resiliency at the overlay and underlay.
4. Minimizes "traffic tromboning" on the underlay and the adverse impacts on performance.
5. Monitor traffic, troubleshoot networks and apply quality of service (QoS) to networks.
6. Implement granular network security using policy-based access without the need for manually configuring access control lists.
7. Reduce CapEx and OpEx – no new hardware required, easier to manage a virtualized network.

About Midokura

Midokura is a global company that offers MidoNet, a network virtualization solution, to enterprise and service providers. Founded in 2010, the team has a pedigree that includes veterans from Amazon and Google, and has spent more than three years building MidoNet, a complete overlay network virtualization solution that integrates with cloud platforms, such as OpenStack. Midokura has offices in San Francisco, Tokyo and Barcelona, and is on the web at http://www.midokura.com. Follow us on twitter: @midokura