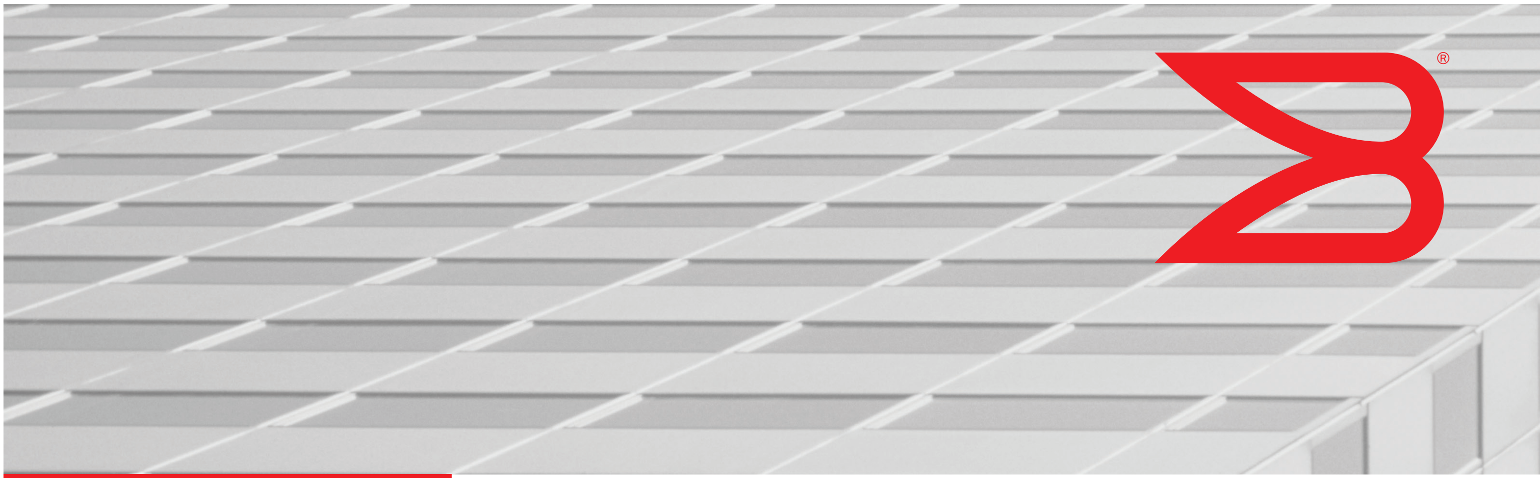

NETWORK FUNCTIONS VIRTUALIZATION

Segmenting Virtual Network with Virtual Routers

INTRODUCTION

For the past 20 years, network architects have used segmentation strategies to make their networks more manageable and secure. Deploying firewalls between servers with different purposes or trust levels has long been a "must have" for any production network—especially those intended to rise to the level of PCI compliance.

The rise of virtualization has caused some network designers to rethink the need for network segmentation. Virtual environments seem to naturally lend themselves to the use of big flat networks. vSwitch, the basic virtual switch provided by VMware, doesn't even support Layer 3 functionality—so absent other technology, virtual machines within a hypervisor are not isolated or segmented. Some engineers have gone so far as to declare that it is time to do away with 3-tiered networks altogether. This paper will look at the question of network segmentation in highly virtualized environments.

IS THE WORLD FLAT?

A flat network is one where the hosts have IP addresses on the same subnet—they are all in the same broadcast domain. Because the hosts are within a shared subnet, routing using a Layer 3 network device isn't required for traffic remaining inside the network.

Flat networks have the advantages of being both simple and, provided there aren't too many devices on it, fast. Flat networks are also supportive of virtual machine migration, an important consideration in today's virtualized world.

Frank Ohlhorst made the case for flat networks when he wrote, "Flat network design came into being because an alternative was needed to interconnect systems relying on massive amounts of connections, caused by heavy virtualization and the convergence of networking technologies. Flat networks eschew the need for Layer 3 routing, which effectively removes traditional security technologies, such as firewalls, filters and other security appliances from the subnet[1]."

The seeming simplicity of large, flat networks comes at a cost—flat networks are limited in the number of devices they can support, troubleshooting and isolating network faults on large flat networks can be a challenge and unsegmented networks allow machines of different trust levels to share traffic—essentially lowering the trust level of all the network hosts to the lowest common denominator. As Ivan Pepelnjak has pointed out[2], Layer 2 networks are a single failure domain. That is to say, when all servers are on the same broadcast domain, and a network loop occurs, all networking to those servers are affected. Ivan wrote, "If you're serious about the claims that you have mission-critical applications that require high availability (and everyone claims they have them), then you simply have to create multiple availability zones in your network, and spread multiple copies of the same application across them." It is worth noting that Layer 2 networks are prone to broadcast and multicast storms and additional mechanisms (configurations) have to be put into place to prevent these storms from hogging bandwidth.

The nail in the flat network coffin is security—and its evil twin—compliance. Yes, there are ways to partially segment flat networks. However, where achieving PCI compliance is an issue, the absence of true network segmentation means that the scope of assessment will be, well, everything; all the devices on the network will need to be assessed for compliance.

Though compliance regimes such as PCI and HIPPA are vague on the specifics of network design, compliance mandates that security best practices be followed. For example, PCI states that credit card processing and user data need be walled off from the rest of the network - placing the rest of that system outside the scope of the assessment. The point is that user and credit card data be isolated and encrypted and that systems with different trust

[1] Frank J. Ohlhorst, Network Computing, March 22, 2012 - http://www.networkcomputing.com/next-gen-network-tech-center/232700055

[2] "LAYER-2 NETWORK IS A SINGLE FAILURE DOMAIN," http://blog.ioshints.info/2012/05/layer-2-network-is-single-failure.html

levels be firewalled to limit potential breaches. In the case of PCI, the specific requirements state, “At a high level, adequate network segmentation isolates systems that store, process or transmit cardholder data from those that do not[3].”

The need to tier network services has left network architects of virtualized data centers with a conundrum; it is not easy to build properly segmented networks within the hypervisor. It takes planning and effort to replicate physical network security policies in virtualized environments.

The challenges of properly networking and securing virtualized environments will only grow larger. Next-generation processors from Intel and others are leading to ever-greater VM densities. As the number of VMs grows, network demands increase. Each server added to a hypervisor increases the network traffic entering, leaving and traversing the host.

PATHS TO NIRVANA—STRATEGIES FOR BUILDING TIERED, VIRTUALIZED NETWORKS

FIRST PATH: ZONE ISOLATION

One approach to segmenting traffic within hypervisors is to only put servers of the same trust zone within any one hypervisor. With this approach, traditional physical routers and firewalls are placed between virtual hosts—creating an air-gap between trust zones. The only difference between this configuration and a traditional, physical data center is that the servers within the trust zone are virtualized.

Organizing virtual servers into common trust zones has a number of advantages, including:

- Simplicity
- Clarity of responsibilities
- Ease of configuration
- Limits the scope of PCI assessment

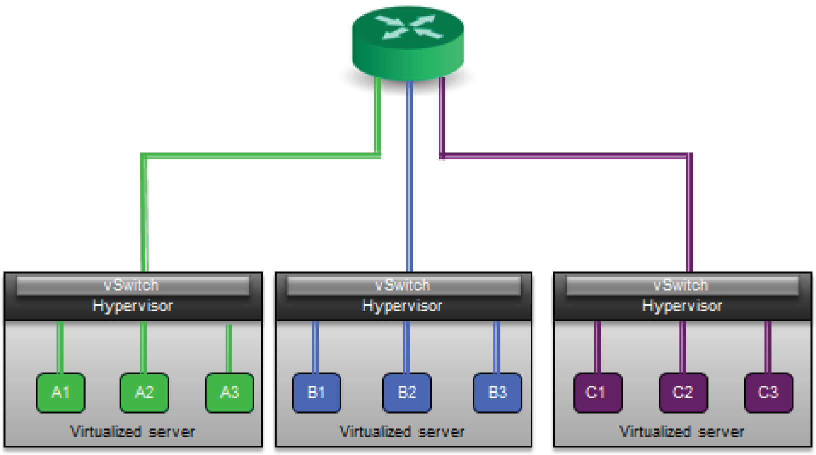

Figure 1.
Hypervisors organized by trust zones lack flexibility and limit server densities.

However, while organizing your virtual systems by trust zone sounds good, implementation can be difficult. IT professionals are under pressure to maximize and balance compute resources. This approach is notably rigid and may prevent organizations from maximizing server densities. The result of organizing your virtual data center by trust zones will probably be greater resource requirements and the loss of operational efficiencies. This approach also assumes near perfect foresight as to how the network will need to evolve—an expectation that few achieve.

SECOND PATH: THROW GEAR AT IT

A second approach to segmenting virtual data centers is to solve the problem with hardware. Network traffic within and between hypervisors can be routed out of the hypervisor and through physical firewalls and routers. This hybrid approach (virtual data center/physical router-firewall) is probably the most common solution to the segmentation problem. The hybrid approach has many advantages, including: that using hardware to segment networks second nature to network professionals, using the same firewall for both virtual and physical segments eases learning and management and buying a few more firewalls from an approved vendor is often a relatively easy purchase.

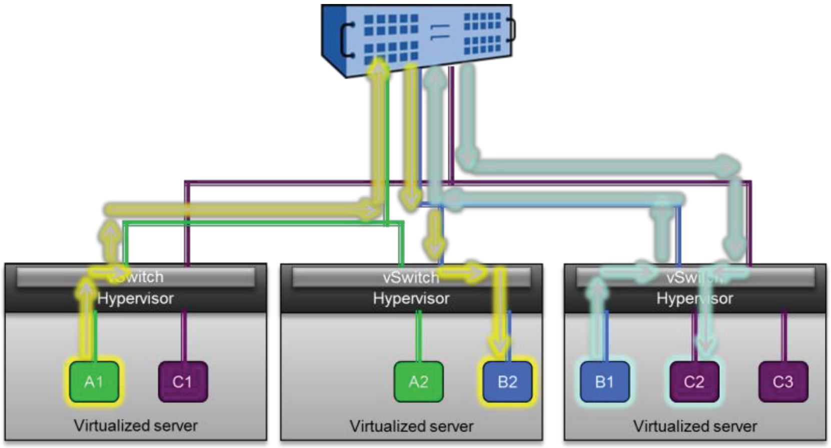

Figure 2.
Inter- and intra-hypervisor traffic can be routed through physical network gear.

However, the hybrid method creates as many issues as it solves. Some network architects are concerned about the “hairpin” effect where traffic intended to go from one virtual machine to another has to exit the hypervisor, go through one or more layers of physical network gear and then return to the virtual environment. While the resulting traffic flow looks inelegant and potentially adds latency and even bottlenecks, these network issues are probably not that serious. In most cases, a few microseconds of added latency won’t be noticed. A larger drawback is the cost and loss of flexibility of relying on extra hardware to help build your virtual data center. Whether your virtual machines are on a local hypervisor or in a public cloud, you probably don’t want to be installing new boxes every time your network needs to grow or change. Not only are proprietary routers expensive, they also require space, power, cooling, spares, etc. The point of the cloud is to reduce reliance on hardware, not add to it.

THIRD PATH: VIRTUAL NETWORKING

Many cloud architects are opting for a third approach to solving the segmentation issue—one more in keeping with the vision of cloud computing. Virtual networking technology can move Layer 3 network functions such as routing, VPN and firewall into the hypervisor. The use of virtual routers and virtual firewalls can solve the conundrum of how to maximize compute resources and agility without sacrificing the network segmentation and machine isolation of physical networks.

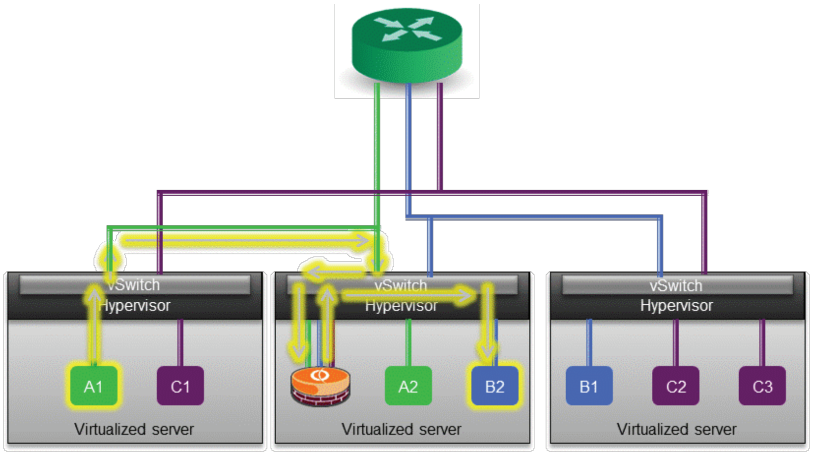

Figure 3.
A virtual router can increase agility while decreasing costs and latency.

Virtual networking relies on software networking - which should not be confused with Software Defined Networking (SDN). SDN is the idea that "network traffic flow can be made programmable at scale, thus enabling new dynamic models for traffic management[4]." Software networking, on the other hand, is the delivery of network services in software, able to run on either standard x86 servers or as virtual machines. In virtualized environments, software networking allows a virtual machine to provide networking services within or between hypervisors.

Virtual networking offers some significant advantages, notably:

- Agility—new networking VMs can be spun up when and where you need them
- Scalability—additional resources can be assigned to the network VM as traffic grows
- Utility Pricing—costs are incurred only has new services are added

Virtual networking is a useful approach, but it has drawbacks. Organizations committed to one brand of router or firewall may not find a suitable virtual edition—requiring training, as well as support of products from multiple vendors. In some cases, central management may be an issue. Additionally, dedicated hardware devices may have performance advantages, especially where deep packet inspection or extensive firewall rule sets are required. Some virtual networking products have significant performance issues, as *Network World Magazine* found when they reviewed Cisco's new Cloud Services Router 1000v.[5]

[4] Wikibon - http://wikibon.org/wiki/v/SDN,_OpenFlow_and_OpenStack_Quantum

[5] Joel Snyder, Network World, Cisco virtual router targets the cloud, 2/25/13:

THE BROCADE SOLUTION

The Brocade Vyatta vRouter is a single, virtualization-optimized solution that includes powerful routing functionality along with stateful firewall, traffic management, IPSec VPN, SSL-based OpenVPN and more. Brocade Vyatta vRouter virtual machines can be employed as virtual gateways on a per server basis to provide hypervisor and application security by establishing zone or rule-based firewalling, detailed traffic inspection and secure remote access.

Complex N-Tier Security

The enterprise-class routing, firewall and VPN capabilities enable tenants to define advanced multi-tier networks, preserving the security and compliance policies enforced within physical networks.

Combat VLAN Sprawl

Deploying Brocade Vyatta vRouter on a per-customer basis provides application isolation and security policy compliance while minimizing reliance on VLANs.

It also eliminates unnecessary latency by reducing multi-trip packet flows between the hypervisor and external physical devices.

PCI Compliance

Using Brocade to build a properly segmented virtual network will ease the path to PCI compliance both by limiting access to critical assets, such as credit card information, and by limiting the scope of compliance assessment efforts.

Auto Provisioning & Remote Management

The Brocade Vyatta Remote Access API and advanced configuration scripting options enable simplified management, orchestration and provisioning through 3rd-party tools. The result is simple button-click deployment and user-defined, template-based configuration of network connectivity and security.

CONCLUSION

Despite the flat network hype, it is clear that the requirement for tiered networks based on networking segmentation hasn't gone away. Now that server virtualization has left the lab and become a common means of delivering production services, the need for network solutions that match the agility and ROI of server virtualization has become critical.

Software-based networking solutions optimized for virtual environments promise a solution for network architects looking to build sophisticated, multi-tiered networks within and between their virtual environments.

Learn more about the Brocade virtual networking solution at www.brocade.com.

Corporate Headquarters
San Jose, CA USA
T: +1-408-333-8000
info@brocade.com

European Headquarters
Geneva, Switzerland
T: +41-22-799-56-40
emea-info@brocade.co

Asia Pacific Headquarters
Singapore
T: +65-6538-4700
apac-info@brocade.com

Notice: This document is for informational purposes only and does not set forth any warranty, expressed or implied, concerning any equipment, equipment feature, or service offered or to be offered by Brocade. Brocade reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. Contact a Brocade sales office for information on feature and product availability. Export of technical data contained in this document may require an export license from the United States government.

BROCADE