WHITE PAPER

Segmenting Virtual Networks with Virtual Routers

Brocade VCS® Fabric technology helps organizations deliver efficient, highly automated networks for modern data centers. Ethernet fabrics built on Brocade VCS Fabric technology provide unmatched simplicity and manageability compared to traditional network architectures and competitive fabric offerings. In addition, they increase flexibility and IT agility, enabling organizations to transition smoothly to offering elastic, mission-critical services within cloud data centers.

Brocade VCS Logical Chassis functionality allows organizations to manage an entire Brocade VCS fabric as a single switch. This eliminates the need to manually configure and manage each switch, simplifying management and reducing operational costs.

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 Payment Card Industry (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 large, flat networks. vSwitch, the basic virtual switch provided by VMware, does not even support Layer 3 (L3) functionality, so that—absent other technology—Virtual Machines (VMs) within a hypervisor are not isolated or segmented. Some engineers are even calling for an end to the use of 3-tiered networks altogether. This white paper examines the use 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—all in the same broadcast domain. Because the hosts exist within a shared subnet, routing using an L3 network device is not required for traffic that remains inside the network.

Flat networks have the advantage of being both simple and fast, as long as the network does not have to support too many devices. Flat networks also support VM migration, an important consideration in today's virtualized world.

Frank J. 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 [a] massive [number] 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.
- Unsegmented networks allow machines of different trust levels to share traffic, lowering the trust level of all the network hosts to the lowest common denominator.

As Ivan Pepelnjak pointed out, Layer 2 (L2) networks are a single-failure domain.[2] In other words, when all servers are on the same broadcast domain, and a network loop occurs, all networking to those servers is affected. Pepelnjak 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."[3] It is worth noting that L2 networks are prone to broadcast and multicast storms; thus, additional mechanisms (configurations) must be put into place to prevent such storms from hogging bandwidth.

Security and regulatory compliance are perhaps two of the greatest hurdles to overcome and consider when architecting complex multitier networks. 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 is everything; that is, all the devices on the network must be assessed for compliance. Though compliance regimes such as PCI and the Health Insurance Portability and Accountability Act of 1996 (HIPAA) are vague on the specifics of network design, compliance mandates that security best practices are followed. For example, PCI states that credit card processing and user data must be walled off from the rest of the network—placing the rest of the network outside the scope of the assessment. The goals are to isolate and encrypt user and credit card data and to install firewalls on systems with different trust levels in order 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."[4]

[1] Frank J. Ohlhorst, "Flat Network Strength Also A Security Weakness", Network Computing, March 22, 2012, at www.networkcomputing.com/next-gen-network-tech-center/232700055.

[2] Ivan Pepelnjak, "Layer-2 Network is a Single Failure Domain," IP Space blog at blog.ioshints.info/2012/05/layer-2-network-is-single-failure.html.

[3] Pepelnjak, ibid.

[4] PCI Requirements and Security Assessment Procedures, Version 2.0, October 2010, at www.pcisecuritystandards.org/documents/pci_dss_v2.pdf.

The need to tier network services creates a conundrum for network architects of virtualized data centers: 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 in the future. Next-generation processors from Intel and other vendors are leading to ever-greater VM densities. As the number of VMs grows, network demands increase. Each server that is added to a hypervisor increases the network traffic entering, leaving, and traversing the host.

Strategies for Building Tiered, Virtualized Networks

First Path: Zone Isolation

One approach to segmenting traffic within hypervisors is to put only 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 provides a number of advantages, including:

- Simplicity
- Clarity of responsibilities
- Ease of configuration
- Limited scope of PCI assessment

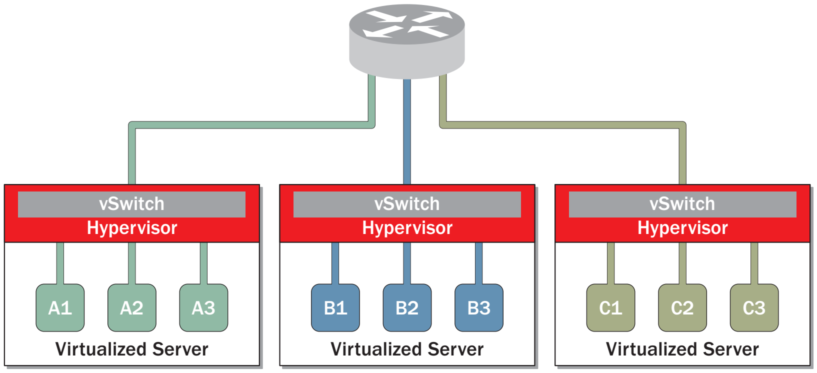

Figure 1: Hypervisors organized by trust zones lack flexibility and limit server densities.

However, although it might seem like a good idea to organize your virtual systems by trust zone, 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. If you organize your virtual data center by trust zones, the result will probably be greater resource requirements and a loss of operational efficiencies. Such an approach also assumes that you can accurately predict how the network might need to evolve in the future—which is nearly impossible to do.

Second Path: Install More Hardware

A second approach to segmenting virtual data centers is to solve the problem by installing more 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 these:

- Using hardware to segment networks is highly familiar to network professionals.
- Using the same firewall for both virtual and physical segments eases learning.
- It is usually easy to purchase and manage a few additional firewalls from an approved vendor.

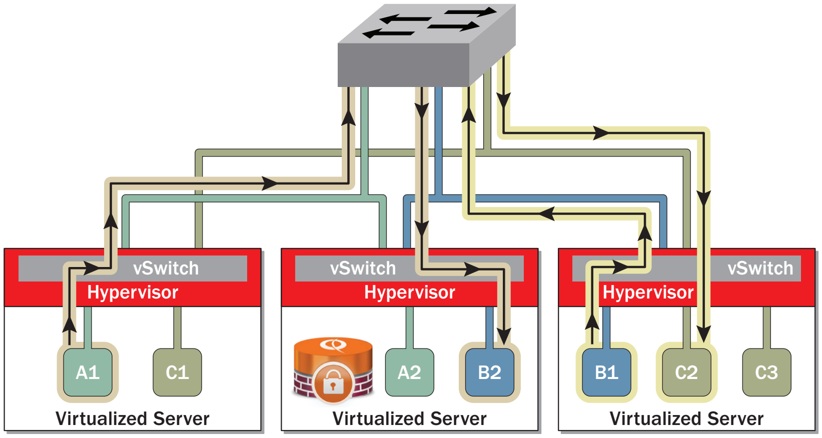

Figure 2: Inter- and intra-hypervisor traffic can be routed through physical network gear.

However, the hybrid method creates as many issues as it solves. Some network architects are concerned about the "hairpin" effect, in which traffic intended to go from one VM to another has to exit the hypervisor, traverse 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 very serious. In most cases, a few microseconds of added latency are not 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 VMs are on a local hypervisor or in a public cloud, it is probably not the best method to install new boxes every time your network needs to grow or change. Not only are proprietary routers expensive, they also require space, power, cooling, spares, and so forth. The usefulness 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—an option that is more compatible with the vision of cloud computing. Virtual networking technology can move L3 network functions such as routing, Virtual Private Networking (VPN), and firewall services into the hypervisor. The use of virtual routers and virtual firewalls can solve the problem of how to maximize compute resources and agility without sacrificing the network segmentation and machine isolation of physical networks.

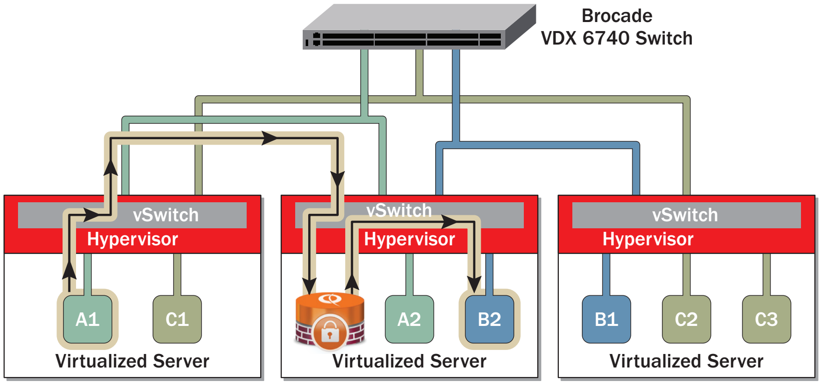

Figure 3: A virtual router can increase agility while decreasing costs and latency.

Virtual networking relies on software networking, which is not the same as Software-Defined Networking (SDN). SDN is based on the idea that "network traffic flow can be made programmable at scale, thus enabling new dynamic models for traffic management."[5] Software networking, on the other hand, is the delivery of network services in software that is capable of running on either standard x86 servers or as VMs. In virtualized environments, software networking allows a VM to provide networking services within or between hypervisors.

Virtual networking offers some significant advantages, notably these:

- **Agility:** You can spin up new networking VMs when and where you need them.
- **Scalability:** You can assign additional resources to the network VM as traffic grows.
- **Utility Pricing:** Costs are incurred only when 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. This requires extra 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 reported in Network World Magazine in a review of the new Cisco Cloud Services Router 1000v.[6]

[5] Stuart Miniman, "SDN, OpenFlow and OpenStack Quantum, Wikibon at www.wikibon.org/wiki/v/SDN,_OpenFlow_and_OpenStack_Quantum.

[6] Joel Snyder, "Cisco virtual router targets the cloud," Network World, 2/25/13, at www.networkworld.com/reviews/2013/022513-cisco-virtual-router-test-266658.html.

The Brocade Solution

Brocade® vRouters are a single virtualization-optimized solution that include powerful routing functionality along with stateful firewall, traffic management, IP security (IPsec) VPN, Secure Sockets Layer (SSL)-based OpenVPN, and more. Brocade vRouter VMs 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 of the Brocade vRouter enable tenants to define advanced multitier networks, preserving the security and compliance policies enforced within physical networks.

Combat VLAN Sprawl

Deploying Brocade 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 multitrip packet flows between the hypervisor and external physical devices.

PCI Compliance

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

Auto Provisioning and Remote Management

The Brocade Remote Access Application Programming Interface (API) and advanced configuration scripting options enable simplified management, orchestration, and provisioning through third-party tools. The result is simple button-click deployment and user-defined, template-based configuration of network connectivity and security.

Conclusion

Despite the hype about flat networks, it is clear that the requirement for tiered networks based on networking segmentation has not disappeared. 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 Return on Investment (ROI) of server virtualization is critical. Software-based networking solutions that are optimized for virtual environments promise a solution for network architects looking to build sophisticated, multitiered 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.com

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 features, 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 information 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.