SERVER VIRTUALIZATION

X86 Virtual Servers and Network Challenges

This technical brief describes the dramatic impact of virtualization on server utilization, rack space, and power and cooling costs. It discusses the software used for moving virtual machines (VMs) across physical servers, which adds another layer of switching to the network access layer and removes control and configuration from the network administrator.

BROCADE

CONTENTS

introduction **3**

Meeting the Virtualization Challenge **5**

Virtual Machine Identification 5

Persistence of Network Polices 5

Scalable Networks for Virtual Machines 6

Virtualization Management 6

Product Features: 7

INTRODUCTION

Virtualization has proven itself an essential technology for lowering cost in the data center. Mainframe computing introduced this capability, which was later extended to UNIX systems. Today, x86 platforms are using virtualization software from a variety of vendors to consolidate multiple applications onto a single x86 server. The impact on server utilization, rack space, and power and cooling costs has been dramatic. Bandwidth has seen explosive growth as well, and is expected to continue at approximately 30% year after year.

However, there are unintended consequences from virtualization that directly affect the network, one of which is the use of an embedded software Ethernet bridge, or softswitch, integrated with the hypervisor. This softswitch is used for activities such as moving VMs across physical servers, to manage internal hypervisor and management traffic. This adds another layer of switching to the network access layer and removes control and configuration of that layer from the network administrator. This, in turn, impacts how network policy and security profiles are managed, complicates network administration, and increases the vulnerability of the access layer to security violations. Combine this with the need for compliance with standards such as the Sarbanes-Oxley Act (SOX), the Peripheral Component Interconnect (PCI), and the Health Insurance Portability and Accountability Act (HIPAA), and a simple procedure to discover why a host in rack A cannot talk to a host in rack B could result in opening tickets with three different groups.

Softswitches have limited support for Ethernet protocols that are typically provided in access layer switches, and implementations often use multiple softswitches per physical server. This dramatically increased the number of "moving parts" in the network access layer, with a corresponding increase in operations effort and complexity.

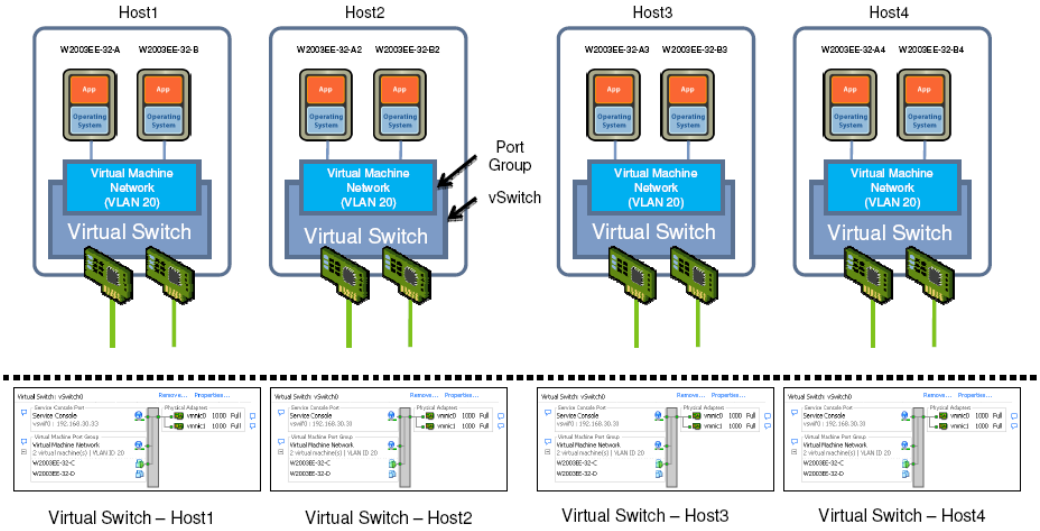

Figure 1. Softswitch (virtual switch) per physical server

One important capability of virtualization is the ability to move a VM and the application running inside it from one physical server to another. This provides 24x7 application availability when the server needs maintenance or upgrades. VM mobility can be used to scale physical resources up or down as the application workload varies. To provide VM mobility without disrupting client connections, the VM IP address must stay the same. This means that any server that the VM moves to must be on the same IP subnet and Ethernet VLAN.

Further, the softswitches in different servers must have the same policies, VLAN, and security settings when a VM migrates—requiring close attention to softswitch configuration. Failure to do so results in a failed migration. This increases complexity and administration time for softswitches. It also impacts the network

administrator, who has to coordinate with the server administrator to ensure that the physical access layer switch configuration matches the configuration of multiple softswitches.

Recently a distributed softswitch that logically extends across all physical servers in the cluster has become available. However, the distributed softswitch is still administered by the sever administrator, not the network administrator. Also, switching VM traffic still requires the server CPU to perform the switching—with unpredictable impact on application performance as network traffic ebbs and flows. This is a particular concern when VMs are migrated from one physical server to another, since migration traffic can monopolize softswitch resources, slowing business applications running on the cluster. Also, there is still the need to address the fragmentation of management responsibility between the server administration and network administrator.

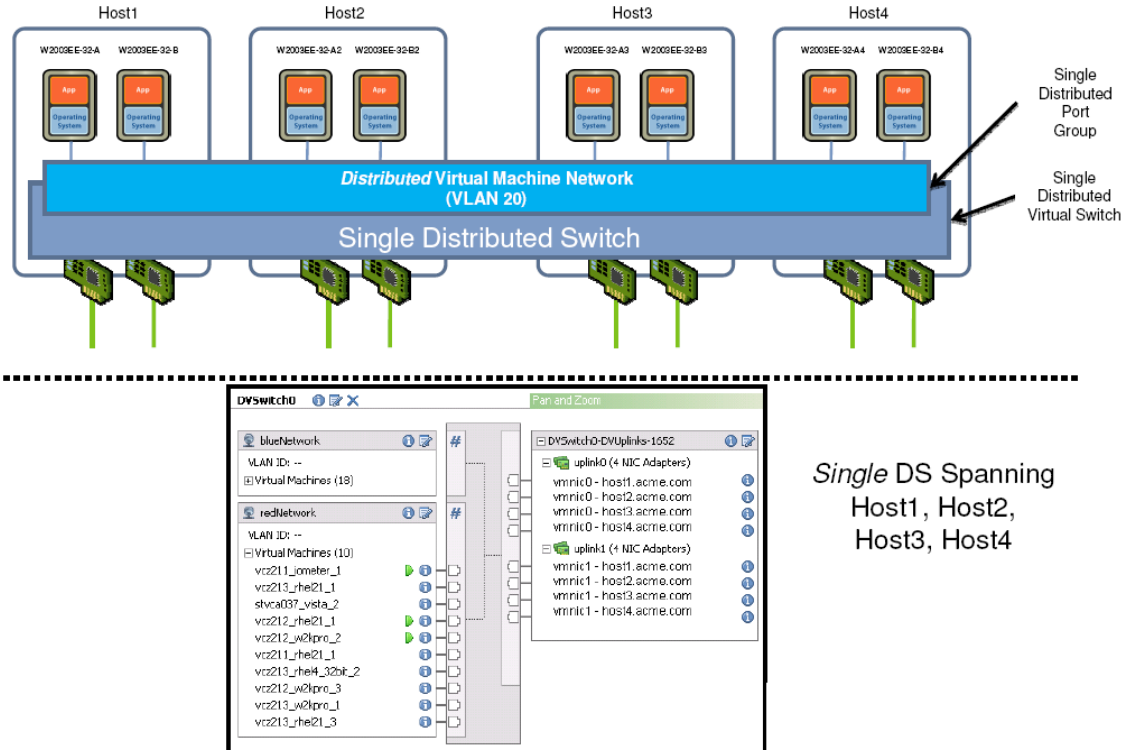

Figure 2. Distributed softswitch (DS) shared by all servers in a cluster

At a deeper level, the introduction of server virtualization has disrupted the accepted network architecture strategy of requiring more switching stages, as well as the proliferation of Network Interface Cards (NICs) in each server for traffic isolation. It has also raised concerns about how to scale the network as softswitches proliferate. These concerns focus on several issues: the negative impacts on network management and operations (due to limited or no visibility to access layer traffic flows), the segregation of responsibility for network security and policy administration, and increased complexity in identifying the root cause of network problems.

To summarize, x86 server virtualization creates these challenges for the network:

- The softswitch is hidden from network management.
- The softswitch adds another hop in the access layer, but it prevents the uniform application of network policies.
- The softswitch stretches the already overcommitted three-tier network model, which adds latency, increases oversubscription, and complicates management.

- Migration of VMs is not transparent to the network, thus limiting its application.
- End-to-end network management—which is already complex, error prone, and time consuming—becomes more so when softswitches are added.

Meeting the Virtualization Challenge

The benefits of server virtualization, which dramatically reduces both initial and operating costs, are important and will drive increased adoption as IT continues to have to do "more with less." Therefore, the network has to evolve, particularly at the access layer, to remove the negative impacts of server virtualization on the network. The new network requirements include the following:

- Creating a data center-wide access layer for easy extension of virtualization across the data center
- Simplifying virtual server mobility by adding a virtual network server in the physical network for consistent network policy management
- Aggregating all "network" functions back into the network for network administrator access
- Expanding the scope of the Layer 2 network outside the data center to include remote data centers
- Preserving existing server and network management practices and separation of responsibilities
- Uniformly managing data center-wide network policy management based on applications and VMs, instead of physical switch ports

Brocade® is working on enhancements to its product portfolio to meet these requirements. Our strategy is best captured by the phrase, "revolution through evolution." Since virtual server adoption tends to occur one project and one server rack at a time, changes in the network should be limited to a single rack and should not require changes all the way back to the network core. This provides the lowest cost solution, minimizes risk when adopting new products and technologies, and keeps customers in control at all times.

Specific network enhancements supporting improved network manageability are described in the following sections.

Virtual Machine Identification

Virtual machines are logical computing entities that are created, moved, and deleted as required. Therefore, they are no longer identified by the physical MAC address of a server NIC, and network policies can no longer be statically applied to physical switch ports. VMs need a persistent identifier so that the network can ensure that policies and security profiles are automatically and consistently applied whenever the VM moves from one physical server and access layer switch port to another. In network switches, port configuration settings have to be virtualized and portable. The settings are no longer restricted to physical switch ports but belong to logical "port profiles." Introducing a logical port profile that abstracts the policy, VLAN, and other settings from the physical port provides the portability required for network configuration, as well as policy mobility that matches the mobility of the VM. VM logical identifiers can be mapped to the correct port profile, thus ensuring that policies are always applied consistently to application traffic.

The ability to consistently identify network traffic by VM, even when VMs move across physical servers, is important to ensuring that application service-level agreements (SLAs) are met. To do this, network monitoring and traffic management has to be at the VM level, not just at the physical NIC level.

Persistence of Network Polices

Port profiles need to persist and also be available across all the physical switches that are connected to virtual servers. Port profile management is under the control of the network administrator, so network management functions in virtual server environments are moved back under the control of the network administrator. All network management platforms can access the port profile database and have a consistent view of policies to ensure consistent policy enforcement. Port profiles also ensure that inter-VM traffic is subject to consistent policy and security profiles, a task that is not easy today, given the use of softswitches.

Scalable Networks for Virtual Machines

Scalable networks require simpler management and operations: VM-aware network automation and Automated Migration of Port Profiles (AMPP) automate the discovery of VMs and port profile creation, drastically eliminating the network configurations required for supporting virtual server deployments.

The projected increase in the number of VMs over the next five years indicates that the access layer of the network will have to scale up significantly. For example, the oversubscription ratios that are commonly seen with one application per server will see 10 to 20 times the traffic when applications are consolidated onto a single physical server at ratios of 10-, 20-, or even 40-to-one. Access layer switching will require more performance and more new services, such as port profiles, to meet the increased demands of network traffic.

Newer switches—with the support necessary for port profiles, higher performance, and reduced power and cooling requirements—can be installed in new racks supporting virtualization projects. Or, the newer switches can replace existing switches in racks with servers being repurposed to host virtualization software. Existing access layer, aggregation, and core switches need to continue to work with the newer virtualization-enabled switches, without needing to be replaced prematurely.

The Brocade "revolution through evolution" strategy supports and builds on open standards from IEEE, IETF, and ANSI. Management interfaces such as sFlow are open and used by leading network management platforms. Brocade delivers new products and enhancements to current products with interoperability and backward compatibility, so that scaling the network can be cost-effective and non-disruptive.

The Brocade VDX Data Center Switches are available in 24-, 60-, and 48-cu and 60 port rack-mount form factors and can be deployed in 16-, 24-, 40-, 48-, 50-, or 60-port configurations with Brocade exclusive Ports-on-Demand (POD) licensing. The 24-port and 60-port optical versions now also have the option of adding 8 and 16 ports of Fibre Channel. The Brocade VDX™ family interoperates with existing Ethernet switching products, providing an evolutionary path from legacy Ethernet networks to Ethernet fabrics. Existing legacy switches do not have to be replaced until the end of their useful life; Brocade VDX Data Center Switches can be added at the top of rack or the middle or end of row, forming a scalable, simple-to-manage Ethernet fabric.

For example, several Brocade VDX Data Center Switches deployed in a Top of Rack (ToR) configuration create a single logical chassis with a single distributed control plane across multiple racks of servers, delivering compelling reductions in capital and operating costs, while simplifying VM migration. AMPP is hypervisor-agnostic, which is an important feature, as most data centers deploy different hypervisor stacks based on application and server requirements.

Also, since the Brocade VDX Data Center Switch family supports and builds on the emerging Ethernet Data Center Bridging (DCB), Transparent Interconnection of Lots of Links (TRILL), and Fiber Channel over Ethernet (FCoE) standards, the Ethernet fabric is lossless, low latency, and convergence-ready.

Virtualization Management

Brocade VCS™ technology allows IT organizations to create efficient data center networks that "just work." Ethernet fabric architectures built on Brocade VCS technology share information across nodes and can be managed as a single logical chassis, greatly simplifying management and reducing operational overhead. Brocade VCS technology offers unmatched VM awareness and automation versus traditional architectures and competitive fabric solutions. It also supports storage over a unified fabric when IT organizations are ready. AMPP and VM-Aware Network is a part of the Brocade virtualization management portfolio, which enables organizations to fully align virtual server and network infrastructure resources.

Product Features:

- **VM-aware network automation:** Provides dynamic discovery, zero-touch configuration, and secure communication to VMware vCenter
- **Automated Migration of Port Profiles (AMPP):** Simplifies management by dynamically migrating network policies as VMs move throughout the Brocade Ethernet fabric.

The Brocade VCS solution drastically reduces the operational challenge in virtualized data centers by automating the discovery of VMs and port profile creation. It relives the burden of managing the server and network domain through separate administrative domains. The Brocade VCS solution makes the network totally intelligent, dynamically applying all the port profiles, such as VLANs, access control lists (ACLs), and quality of service (QoS) attributes, significantly reducing the time required for network configurations to support VMs in the network. The solution also enhances the end user experience by seamlessly managing VM migration anywhere in the cluster. Brocade VCS technology allows IT organizations to create efficient data center networks that "just work." Ethernet fabric architectures built on Brocade VCS technology share information across nodes and can be managed as a single logical chassis, greatly simplifying management and reducing operational overhead. Brocade VCS technology offers unmatched VM awareness and automation versus traditional architectures and competitive fabric solutions, and it supports storage over a unified fabric when IT organizations are ready.

Figure 3. Brocade VDX 6720-24 (left) and Brocade VDX 6720-60 (right) for Ethernet fabrics

© 2011 Brocade Communications Systems, Inc. All Rights Reserved. 08/11 GA-TB-387-01-02

Brocade, the B-wing symbol, DCX, Fabric OS, and SAN Health are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, CloudPlex, MLX, VCS, VDX, and When the Mission Is Critical, the Network Is Brocade are trademarks of Brocade Communications Systems, Inc., in the United States and/or in other countries. Other brands, products, or service names mentioned are or may be trademarks or service marks of their respective owners.

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.