




SERVER VIRTUALIZATION

Addressing Server Virtualization Challenges with Brocade Converged Network Adapters (CNAs) and Brocade Network Advisor

According to analysts, the key challenges faced by customers in virtualized data centers are *application performance under load* and *manageability*. This technical brief discusses how Brocade CNAs and Brocade Network Advisor provide unique benefits in virtualized IT environments. Brocade CNAs address scalability and performance for virtualized workloads with Virtual Machine Optimized Ports (VMOPs). Brocade Network Advisor provides virtual machine visibility and end-to-end flow monitoring to help address manageability challenges.



Third party information provided to you courtesy of Dell.

OVERVIEW

Data centers continue to grow as digital assets increase and more applications are deployed. Organizations expect agile application deployment—in minutes, not months—to keep their competitive edge as markets and competitors become global in scale. At the same time, data center resources such as rack space, power, and cooling are growing more scarce and costly. For these reasons, server virtualization is now aggressively deployed in data centers for consolidation of applications hosted on x86 servers.

In addition, server Input/Output (I/O) consolidation, in the form of stand-up and mezzanine 10 Gigabit Ethernet (GbE) CNAs and top-of-rack and embedded Data Center Bridging (DCB)/Fibre Channel over Ethernet (FCoE) switches, is starting to be actively deployed to help organizations take the benefits of virtualization to the next level by further lowering physical server capital and operational expenses and drastically reducing the number of adapters and access layer switches and the cabling required to deploy them.

In physical environments, the access layer of the network was the physical edge switch, and each server and application was statically associated with a physical access layer switch port. With server virtualization, the access layer has moved into the server via software-based Ethernet switches inside the virtualization Virtual Machine Manager (VMM) or hypervisor. The migration of the access layer into the server has created multiple challenges.

According to industry analysts, the two greatest challenges for IT administrators in virtual environments are application performance under load and manageability. In this paper, you will learn how Brocade® stand-up and mezzanine CNAs—together with Brocade Network Advisor—help overcome these two challenges.

Server Virtualization Concerns

Which of these is your greatest concern regarding server virtualization for this year? Next year?

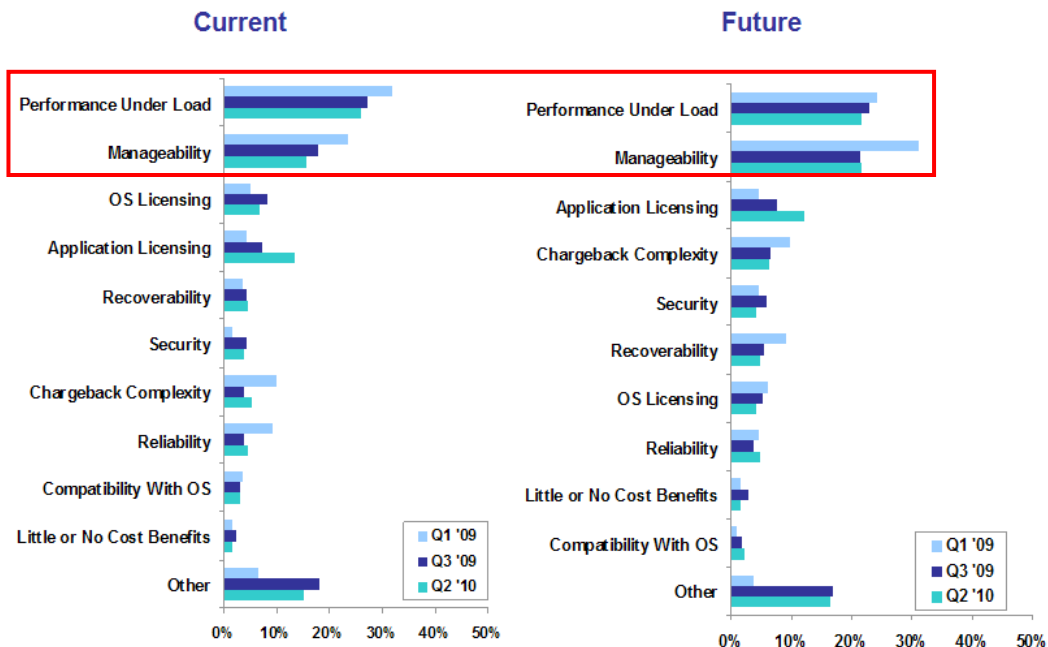


Figure 1. TheInfoPro Server Wave 8 (November 2009)

Brocade offers both stand-up and mezzanine CNAs for Dell rack and blade servers, including the following models:

- **Brocade 1020.** Dual-port 10 GbE CNA for rack-based servers, supported across all major industry-standard operating systems and hypervisors
- **Brocade BR1741M-k 10 GbE CNA for Dell PowerEdge M-Series Blade Servers.** Dual-port 10 GbE CNA specifically designed for Dell M-Series blade servers, providing the same features as the Brocade 1020 in a blade-optimized mezzanine form factor

ADDRESSING THE PERFORMANCE CHALLENGE WITH VM-OPTIMIZED PORTS

As the number of Virtual Machines (VMs) running on a single physical server increases, so do the switching tasks that the hypervisor needs to perform in software. As each incoming network packet arrives on the Network Interface Card (NIC), the hypervisor needs to examine that packet and determine—primarily based on the destination MAC address and VLAN tagging—the destination VM. The packet is then forwarded to the processor core that has an affinity with that VM and ultimately to the appropriate VM. These packet classification and sorting tasks create significant I/O overhead and they can take up a large portion of CPU processing, severely impacting network performance. As a result, the hypervisor becomes a bottleneck for network I/O. Even at 10 Gigabit per second (Gbps) speeds, the resulting performance can be in the range of 2 to 4 Gbps.




Figure 2. Packet classification/sorting at the hypervisor

Brocade addresses virtual switching technology by essentially integrating an open, standards-based virtual switch into our CNAs. And this host-based hardware virtual switching solution will be extended in future to provide the flexibility to switch packets between VMs either at the server or outside the server at the access layer switch by leveraging the Ethernet Virtual Bridging (EVB) standards such as Virtual Ethernet Port Aggregator (VEPA) or Virtual Ethernet Bridge (VEB).

Today, the Brocade CNAs leverage hypervisor multi-queue network technologies such as VMware NetQueue and Microsoft VMQ through a feature called Virtual Machine Optimized Ports (VMOPs), which provide a highly optimized networking infrastructure to efficiently route packets, offload the CPU, and reduce latency. Brocade VMOPs implement up to 128 optimized “ports” or queues (64 queues per port) in hardware, enabling significant VM scalability, while guaranteeing high performance and CPU efficiency.

VMOP core hardware technology is a hardware-based Layer 2 classifier, which sorts incoming packets based on the L2 header—primarily MAC address and/or VLAN tags. VMOP multiplexes and demultiplexes workloads based on this distribution logic. Moreover, each VMOP queue is assigned an MSI-X interrupt vector to distribute the workload across multiple processors. VMOP requires no user configuration and is enabled on the data path once VMware NetQueue or Microsoft VMQ is enabled.

The data packets arriving on the network adapter port are passed through the hardware based L2 classifier. Based on the MAC address and/or VLAN tags, the packet is placed on the appropriate receive queue associated with the processor core and/or destination VM. Since a unique MSI-X interrupt vector is associated with the receive queue, the adapter can interrupt the affinity processor for the work item and deliver the packet to the processor and core associated with that particular VM.

From a data transmission perspective, the data packets transmitted from the VMs are placed into the corresponding transmit queues by the hypervisor¹. The adapter implements advanced queue servicing algorithms to optimize the downside link utilization and prevent head-of-line blocking. The queue servicing also guarantees no starvation and fairness is maintained across the ports.

Finally the CNAs provide L2/L3 offloads such as TCP Segmentation Offload (TSO), IPv4/IPv6/TCP/UDP checksum offload, Receive Side Scaling (RSS), Large Receive Offload (LRO) and Header Data Split (HDS). These stateless offloads provide significant performance benefit and remove the processing burden from the operating system or hypervisor CPU.

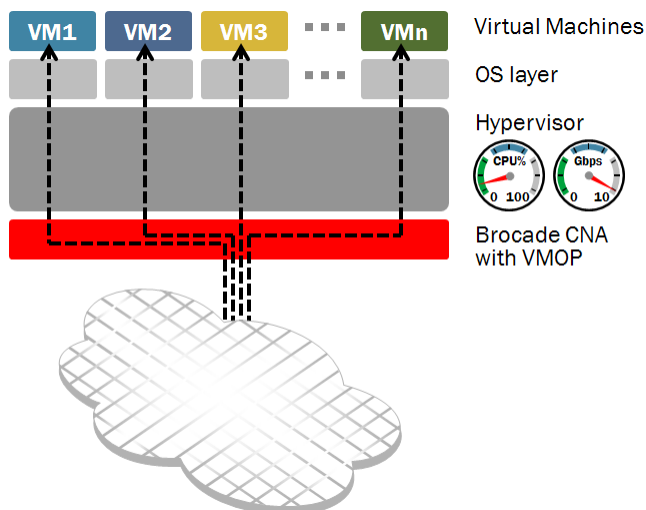


Figure 3. Packet classification/sorting offloaded to the CNA with VMOP

¹ Currently available for VMware ESX. Server; support for Microsoft VMQ transmit queues will be added in a future driver release

VM VISIBILITY AND CONTROL WITH BROCADE NETWORK ADVISOR

The virtualized data center presents a number of manageability challenges. Virtualization creates an abstraction layer that results in multiple applications sharing the same physical adapter and network switch ports. From the network point of view, there is no way to differentiate each of the applications that access the network through the same adapter and port, so application visibility is essentially lost. Monitoring application performance or troubleshooting and diagnosing problems such as congestion in these conditions is nearly impossible.

Server I/O consolidation is the first phase of network convergence that is starting to be deployed at the server edge. Combining Ethernet NICs and Fibre Channel (FC) HBAs in a single adapter, as well as separate top-of-rack Ethernet and Fibre Channel switches in a single access layer switch, means that administrators need to consider the Ethernet and storage networking aspects of their server deployments. Providing a single “pane of glass” to manage both environments becomes critical to achieve the management efficiencies that are required in very dynamic virtualized IT environments.

For data center administrators looking to simplify their network management and to consolidate tools for their virtual, cloud, or physical infrastructure, Brocade Network Advisor delivers the industry’s only comprehensive unified easy-to-use management application with end-to-end network visibility from a single application, supporting Ethernet, DCB, FCoE and FC networks, and associated server connectivity.

One of the key features of Brocade Network Advisor is the ability to discover Brocade server connectivity resources across a data center environment to monitor and manage groups of adapters (HBAs and CNAs) from a central location—together with Brocade switches, directors, and backbone platforms. By communicating with the adapter agent running on the servers, Network Advisor can discover adapter properties, providing end-to-end visibility from the adapter all the way to the storage LUN. It provides host-based topology views (shown in Figure 4) from initiator to target, automatic server-to-adapter mapping, real-time performance monitoring for adapters, end-to-end switch and adapter configuration operations (such as FCSP), and the ability to receive adapter events and collect adapter diagnostic and support data from multiple servers simultaneously. It also integrates with Brocade Host Connectivity Manager (HCM), the element manager for Brocade adapters, to provide a fully integrated management stack.

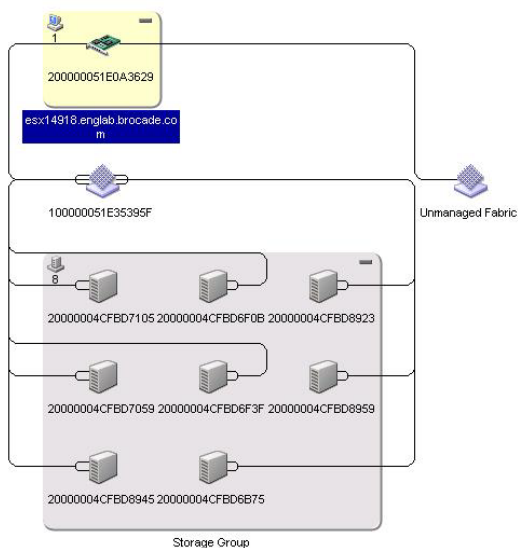


Figure 4. Host-based topology view

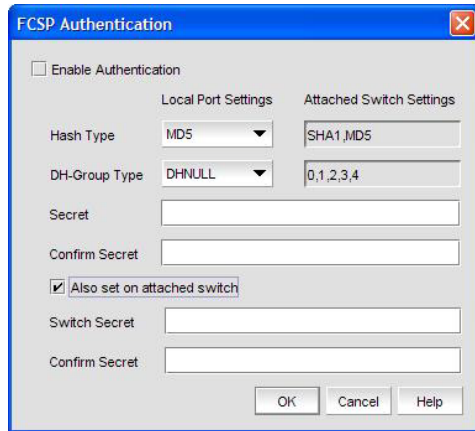


Figure 5. End-to-end management operations

In virtualized environments, in addition to the discovery of the adapters on the servers, Brocade Network Advisor uses the VMware ESX Server² APIs to communicate with the hypervisor and discover all VM-related information from the server, providing unmatched levels of visibility into the virtual server environment for network administrators. Through this integration, Network Advisor can provide detailed information about all the VMs running in a physical server, including CPU and memory allocations, ESX Server and guest OS version information, VM datastore information, and end-to-end VM-to-LUN visibility, as shown in Figure 6.

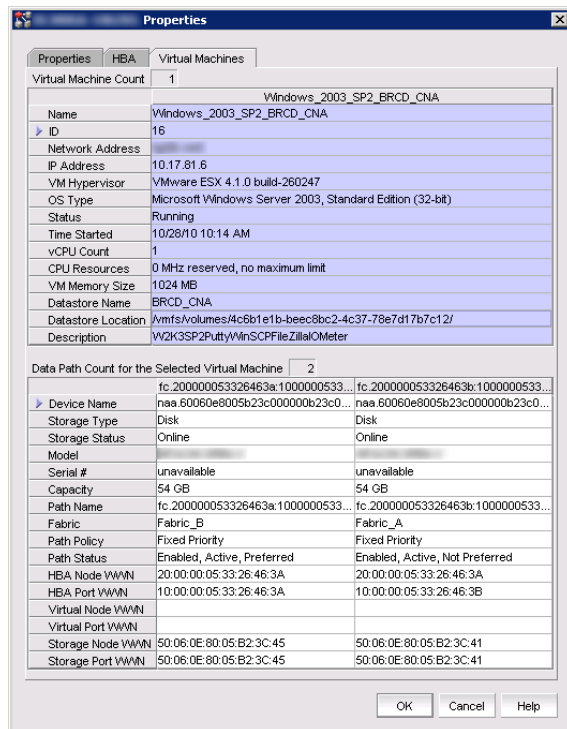


Figure 6. VM information displayed in the Properties dialog box

² Microsoft Hyper-V support will be added in a future release

In addition, if N_Port ID Virtualization (NPIV) is being used at the hypervisor, Brocade Network Advisor can identify each VM in the fabric and correlate the information discovered from the fabric with the information obtained from ESX Server to provide unprecedented levels of detail. In this case, Network Advisor can also monitor performance—in real time (shown in Figure 7) and on a historical basis—and assign network policies such as zoning, security, or Quality of Service (QoS) for individual VMs.

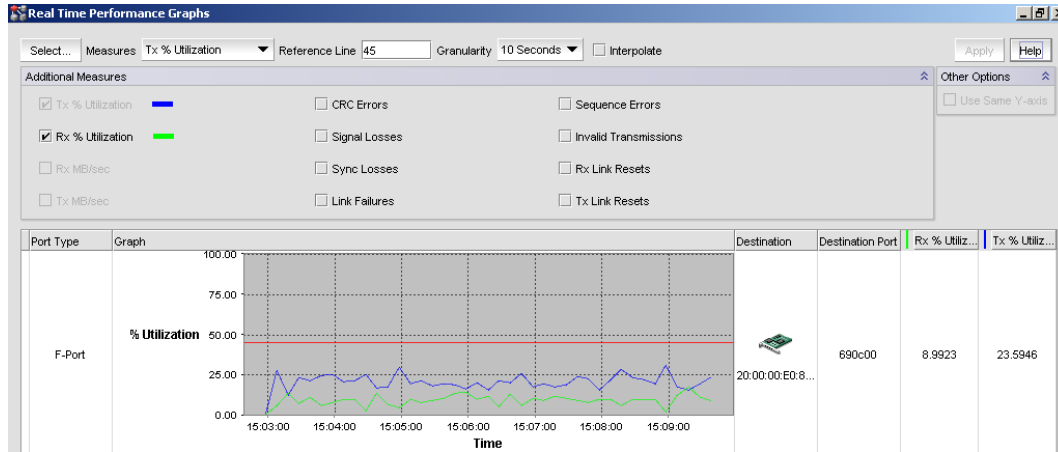


Figure 7. Real-time performance monitoring

This end-to-end performance monitoring capability, from VM to storage device, provides the visibility and pulse of application behavior, helping identify application receive/transmit flows and assisting in the diagnosis and troubleshooting of application performance issues.

In addition to providing VM visibility for SAN administrators, Brocade Network Advisor also offers network visibility for VM administrators by providing integration with industry-leading third party management applications such as VMware vCenter or Microsoft System Center Operations Manager (SCOM) via management packs, or plug-ins, that integrate with Network Advisor.

The Brocade Management Plug-in for VMware vCenter provides VMware administrators critical insight into network connectivity of their VMs and VM health to enable proactive response via policy-based actions for network issues. Brocade Network Advisor populates VMware vCenter with SAN connectivity information as it pertains to the VMs—fabric details, switch details, performance statistics, event messages, and intelligent recommendations that IT administrators can then implement, as shown in Figure 8. This provides VM administrators with a single pane of glass view across VMs, their connections through the SAN, and their connectivity to storage LUNs.

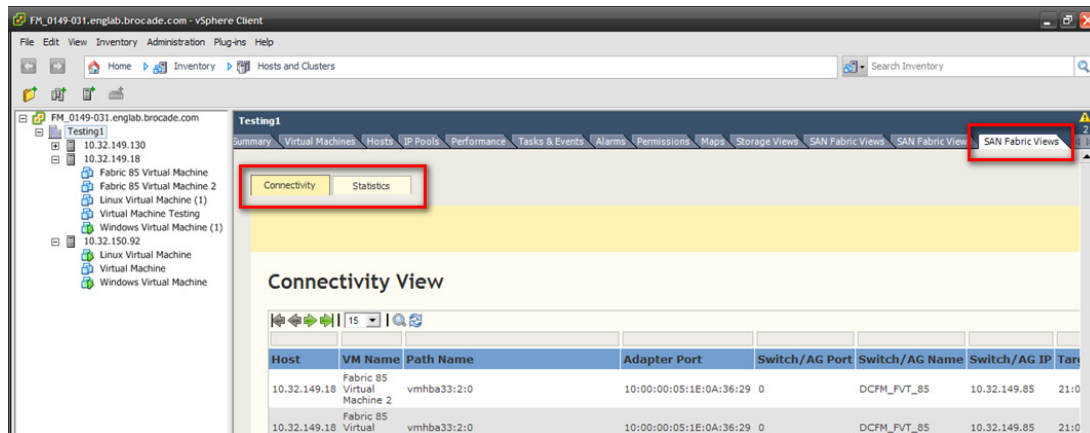


Figure 8. The vCenter plug-in provides end-to-end fabric visibility and performance statistics

The Brocade Management Pack for Microsoft System Center Operations Manager uses open, standards-based interfaces to seamlessly integrate with Microsoft SCOM and to help bridge operational gaps across server, network, and storage administrators. Available directly within SCOM, it provides switch and fabric details, Call Home events, and end-to-end traffic statistics. With this level of information, Microsoft server administrators can quickly identify network bottlenecks in the environment, including the ability to view performance statistics or CRC errors that could be causing poor application performance.

SUMMARY

Brocade CNAs and Brocade Network Advisor are uniquely positioned to help customers address their main challenges when it comes to virtualizing their data center, helping them transition smoothly in their journey to the private cloud.

This whitepaper was created by Brocade and is provided to you as a courtesy. Dell makes no warranties or representations regarding the accuracy of any information in this paper. Any questions or comments regarding this paper should be addressed to Brocade.

© 2010 Brocade Communications Systems, Inc. All Rights Reserved. 11/10 GA-TB-358-00-Dell

Brocade, the B-wing symbol, BigIron, DCFM, DCX, Fabric OS, FastIron, IronView, NetIron, SAN Health, ServerIron, Turbolron, and Wingspan are registered trademarks, and Brocade Assurance, Brocade NET Health, Brocade One, Extraordinary Networks, MyBrocade, VCS, and VDX 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.