

10GBE, SERVERS, STORAGE AND VIRTUALIZATION – INTEROPERABILITY REVIEW AND HIGHLIGHTS

A Dell™ White Paper

June 2009



Information in this document is subject to change without notice.

© Copyright 2009 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the written permission of Dell Inc. is strictly forbidden.

This white paper is for informational purposes only and may contain typographical errors or technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

Dell, the *DELL* Logo, *EqualLogic*, *PowerEdge*, *FlexAddress*, and *OpenManage* are trademarks of Dell Inc. *Microsoft* is a registered trademark of Microsoft Corporation.; *VMware*, *vCenter*, and *VMotion* are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. *Intel* and *Xeon* are registered trademarks of Intel Corporation in the US and other countries. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims proprietary interest in the marks and names of others.

Table of Contents

Introduction.....	4
A Sampling of 10GbE Capable Solution Builders from Dell	4
Dell PowerEdge Blade Server.....	4
PowerConnect Blade Switch	6
Broadcom® NetXtreme™ II 57710 Ethernet Card.....	6
VMware Infrastructure 3	7
Dell blades + 10GbE + VMware = A Robust Infrastructure	7
Benefits of Dell Blades, 10GbE and Virtualization	8
Summary	12

Introduction

IT departments are consolidating servers and implementing virtualization to simplify their environments and lower the total cost of ownership. Some organizations are achieving consolidation ratios of as high as 100:1, and 20:1 consolidation is commonplace, using VMware® and multi-core server technology. But success in reducing the number of servers, and the concomitant reduction in management and environmental costs, has a price: each reduction in physical server presence also subtracts a number of NIC ports. A proliferation of virtual machines on a single server competes for that server's resources, often resulting in I/O bottlenecks.

iSCSI over 10GbE provides the bandwidth needed to handle the increased I/O burden of a consolidated and virtualized server environment—along with the scalability and flexibility to support new servers and storage devices in the future. Organizations moving to 10GbE iSCSI can lower equipment and power costs, consolidate and manage their servers more easily, and increase throughput in their virtual server environment.

For reasons such as these, analysts are predicting a rapid adoption of 10GbE in the next two years. Forrester recommends targeted deployments now, expecting widespread implementation in 2010.¹

By implementing 10GbE solutions, organizations can enjoy the economies they hoped to obtain by consolidating and virtualizing their server environment, along with increased performance, reliability, security, and ease of use. Dell can assist customers implement 10GbE solutions that are right for their environment using products similar to the ones described in this paper.

A Sampling of 10GbE Capable Solution Builders from Dell

Dell offers complete 10GbE iSCSI solutions today, comprising products such as:

- Dell PowerEdge™ M1000e blade enclosures
- PowerEdge M610 blade servers
- Dell PowerConnect™ M8024 blade switches
- Broadcom® NetXtreme™ II 57710 Ethernet cards
- VMware Infrastructure 3

Dell PowerEdge Blade Server

Dell PowerEdge blade servers offer high density for consolidation, plus abundant I/Os for virtual machines.

¹ Forrester Research, Inc., May 2008, *10GbE: Its Time Is Coming*

Physical server consolidation benefits from using blades that demand less rack space, and lower environmental costs including power costs. Blades also require less management than traditional rack-mounted systems.

Blade Modular Enclosure: The Dell PowerEdge M1000e is a high-density, energy-efficient blade chassis that supports up to sixteen half-height blade servers, or eight full-height blade servers, and six I/O modules in one enclosure. A passive high-speed midplane connects the servers to the I/O modules' management and power in the rear of the enclosure. The front side of the enclosure includes a flip out LCD screen for local systems management configuration. The enclosure includes six hot pluggable and redundant power supplies, and nine hot pluggable, N+1, redundant fan modules. Longevity is key for a blade enclosure because it needs to support multiple generations of blade servers. The PowerEdge M1000e blade enclosure was designed with lifecycle management in mind, and is targeted to support new blades through 2013, with warranty support through 2018.

Blade servers: Blade servers are becoming increasingly valuable, with robust memory and CPU capabilities. The PowerEdge M1000e supports the recently announced 11th generation PowerEdge M610 and M710 blade servers that are based on the new Intel® Xeon® processors. The PowerEdge M610 supports 12 DIMM slots, and the M710 supports 18 DIMM slots. These blade servers use Intel Quickpath technology to provide high-speed links to the memory modules. These 11th generation blade servers come with the next generation system management tool: Unified Server Configurator (USC) to help reduce operating costs by simplifying deployment and management. USC supports diagnostics, firmware updates, and hardware configuration. The enclosure also supports AMD-based PowerEdge M605, M805 and M905 blade servers, as well as the older Intel-based PowerEdge M600 blade server.

I/O connectivity: The enclosure provides three redundant fabrics, using six hot-swappable I/O modules. Ethernet options for these modules are PowerConnect blade switches, Cisco Catalyst blade switches, and pass through modules. FibreChannel through 8G/s and Infiniband Dual Data Rate (DDR) and Quad Data Rate (QDR) switch modules are also supported. Every fabric is available to each server, and all server I/O connections cards have dual ports with each port mapping to a different switch for redundancy, for a total of 6 ports on half height blades. Full height blade servers have complete redundancy with each I/O card replicated to the motherboard, for a total of 12 total ports. With the full height blade, you can have 4x1G/s NIC ports to fabric A (2x dual port physical connections) and 8x10Gb/s ports total to fabrics B & C (4x dual port physical connections).

Management: The Dell PowerEdge M1000e has integrated management features, including redundant Chassis Management Controller (CMC) modules for enclosure management and an integrated keyboard, video, and mouse (iKVM) module. Through the CMC, the enclosure supports FlexAddress™ technology. FlexAddress technology allows any M-Series blade enclosure to lock the World Wide Name (WWN) of the FibreChannel controller and Media Access Control (MAC) of the Ethernet and iSCSI controller into a blade slot. This enables seamless swapping or upgrade of blades servers within the chassis for

Ethernet, iSCSI, and FibreChannel controllers without impacting the LAN/SAN mapping or zoning.

PowerConnect Blade Switch

The Dell PowerConnect M8024 is a member of a broad family of Dell PowerConnect products. Each PowerConnect M8024 switch provides twenty four 10GbE ports to the Dell™ PowerEdge M1000e blade chassis; 16 internal (i.e., to the blades) ports and up to 8 external ports, for a 2:1 aggregation. This valuable capability delivers 10GbE to each of the 16 server blades as needed. The enhanced bandwidth and performance satisfies the increasingly important requirement of servers embracing technologies like virtualization.

Extreme Flexibility

In addition to the 16 internal 10GbE ports, the M8024 delivers flexible external I/O connectivity choices, utilizing two modular CX-4 and/or SFP+ uplink modules as needed. Each SFP+ module provides four 10GbE ports. When using two SFP+ modules, a total of eight external 10GbE ports are provided. Each CX-4 module provides three 10GbE ports. Two CX-4 modules yield a total of six external 10GbE ports. For even more flexibility, the M8024 supports mixed environments using both the SFP+ and the CX-4 modules.

High Density

The PowerConnect M8024 is designed specifically for the Dell PowerEdge M1000e enclosure, and provides several benefits including optimized data center space utilization, reduced cable infrastructure and complexity, shared power and cooling infrastructure, and lower total cost of ownership.

The PowerConnect M8024 provides valuable features, functions, and benefits available in other Dell switches installed in the data center. This not only eases integration, but also simplifies on-going management. IT personnel will find familiar products recognized for their valuable capabilities. Additionally, IT server administrators will find deployment of the M8024 to be simplified, and will enjoy the benefits delivered by the simple switch firmware feature included with the M8024. These capabilities make deploying a 10GB Ethernet utilizing M8024 switches a quick and effective solution for growing data centers.

Broadcom® NetXtreme™ II 57710 Ethernet Card

The Broadcom® NetXtreme™ II 57710 Ethernet card introduces a new level of cost effective networking performance with increased bandwidth, reduced latency, and lower CPU utilization at 10GB speeds with Dell PowerEdge blade servers.

Based on Dell blade server adapter designs, the 57710 card has two ports for redundancy and connects to networks at 10GB using the Dell PowerConnect M8024 10GB Ethernet switches. Together, these products provide the necessary infrastructure, performance, and bandwidth critical to I/O intensive applications such as virtualization.

VMware Infrastructure 3

VMware on Dell blades offers excellent virtualization capabilities. VMware Infrastructure 3 includes VMware ESX/ESXi and vCenter Server, with the following components:

- **VMware VMotion:** VMware VMotion technology provides real-time migration of running virtual machines from one server to another with no disruption or downtime. VMotion can be used to achieve manual load balancing and zero downtime maintenance.
- **VMware Distributed Resource Scheduler (DRS):** VMware DRS technology enables VMotion to automatically achieve load balancing according to resource requirements.
- **VMware High Availability (HA):** Upon server failure, VMware HA automatically restarts virtual machines on other physical servers. VMware HA provides high availability at the virtual machine level.
- **VMware Update Manager:** VMware Update Manager automates patch management, enforcing compliance to patch standards for VMware ESX hosts as well as Microsoft or Linux virtual machines. Secure patching of offline virtual machines, as well as automated snapshots and rollback capabilities, reduce risk in the environment. Integration with DRS allows for non-disruptive VMware ESX patching with little-to-no administrative overhead.
- **VMware Storage VMotion:** VMware Storage VMotion enables real-time migration of running virtual machine disks from one storage array to another with no disruption or downtime. It minimizes service disruptions due to planned storage downtime previously incurred for rebalancing or retiring storage arrays.

For more information on VMware Infrastructure, refer to <http://www.vmware.com/products/vi/>.

Dell blades + 10GbE + VMware = A Robust Infrastructure

VMware Infrastructure on Dell blade servers and a 10GbE Dell EqualLogic™ SAN constitute a robust infrastructure that meets the demands of today's datacenters. They offer rapid deployment, optimized resource utilization, simplified management, seamless scalability, and energy efficiency with no compromise in performance.

Rapid Deployment

The solution offers several tools that help customers go from an out-of-the-box environment to a production environment quickly and seamlessly. Dell's USC provides a single interface to configure the server BIOS, network controller, and storage controller. USC also supports VMware ESX installations. Should customers opt for VMware ESXi, Dell ships the PowerEdge M610, M710, M805 and M905 servers with integrated ESXi on an embedded SD card. Dell EqualLogic SAN provides a highly simplified deployment process and can go from an out-of-the-box environment to a production deployment in under an

hour, with an intuitive, wizard-based GUI to integrate easily into the Ethernet backbone of the virtual infrastructure. Using virtualization capabilities, new VMs can be created or existing VMs can be cloned or deployed from templates with just a few steps.

Optimized Resource Utilization

In addition to increased resource utilization through server consolidation, VMware's VMotion and DRS technologies can dynamically distribute workloads to all servers in the cluster, ensuring that all server computing resources are optimized. Dell EqualLogic SAN complements this by optimizing the storage devices. The virtual machine data is intelligently distributed between all the members (storage devices) in the storage pool, ensuring that all the storage controllers and network interfaces are most appropriately utilized. This distribution is done automatically and is completely transparent to the user.

Simplified Management

One of the key benefits of virtualization is simplified management. VMware ESXi software comes integrated with Dell OpenManage™ management software, providing customers with a simplified systems management interface.

Seamless Scalability

VMware virtualization in combination with Dell servers and storage provides a highly scalable solution to meet the growing IT demands of today's datacenters. New server resources can be easily added to VMware DRS clusters. Virtual machines will automatically migrate to utilize the newly added resources. Similarly, new storage devices can be easily added to the Dell EqualLogic storage pools. The EqualLogic software will automatically migrate the virtual machine data to the newly added storage resources with no downtime.

Energy Efficiency

The Dell PowerEdge M1000e enclosure takes advantage of thermal design efficiencies, such as ultra-efficient power supplies and dynamic power-efficient fans with optimized airflow design, to efficiently cool the enclosure and enable better performance in a lower power envelope.

Consolidating workloads running on older servers on the new energy-efficient Dell blade servers greatly reduces power consumption. This capability, combined with VMware's DPM technology that shuts down underutilized servers, substantially improves the energy efficiency of the solution. Power "not to exceed" thresholds and alerts can be set by the user, and the enclosure will automatically manage resources (per user defined policies) to stay within these defined parameters. Additionally, the enclosure can automatically move power supply units (PSU) on and offline to dynamically respond to power demand. This keeps PSU utilization high and increases energy efficiency even more.

Benefits of Dell Blades, 10GbE and Virtualization

As businesses feel the need to better utilize datacenter resources and lower TCO, virtualized infrastructure deployment will steadily rise. Virtualization improves the availability of common enterprise applications. It helps IT lower expenses by improving

server utilization. Having fewer physical servers can reduce the number of person hours needed to manage the hardware and lower power and space costs, potentially resulting in a significant annual savings in operating expense. Virtualization also results in a more flexible infrastructure in which storage space can be provisioned faster and more easily.

Physical server consolidation benefits from blades include requiring less rack space, power, management, and cooling than standard rack-mount servers. Moving from older, lower performing servers to new, higher performing blades yields a faster, more capable server environment, with better performance and more easily managed hardware. Adding server virtualization results in a higher density of consolidation, extending these benefits even further.

Dell PowerEdge M610 and M710 blade servers and PowerEdge M1000e blade enclosure exemplify the new generation of multicore servers that can support a larger number of virtual machines than ever before. This translates into a greater potential savings for businesses through server virtualization.

However, with increased virtualization comes increased network traffic. As the number of virtual machines per physical server increases, the physical network layer must be expanded to accommodate the increased network traffic or risk creating a bottleneck. To meet the requirements of the increased network traffic, an organization can add more 1GbE network adapters, but this would mean more cabling and more management. Alternatively, migrating to 10GbE NICs would provide the required bandwidth without adding complexity and increasing the management burden as is shown in Figure 1.

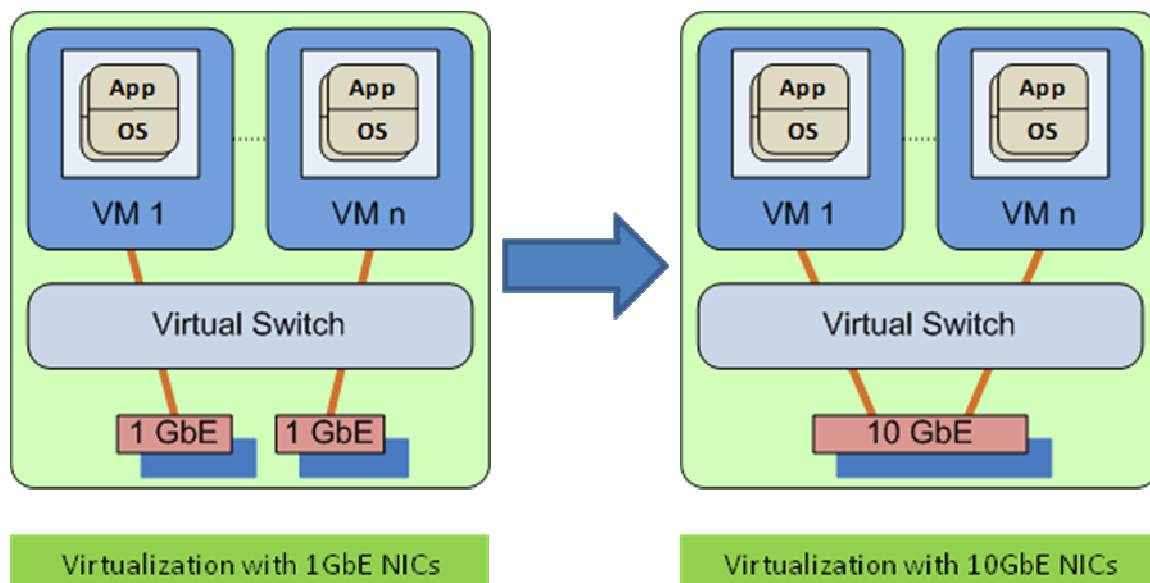


Figure 1: 1GbE vs. 10GbE network adapter configurations

With 10GbE, an organization can enjoy the cost savings and agility of virtualization while avoiding I/O bottlenecks for high performance servers as the massive bandwidth

eliminates network congestion. Additionally, 10GbE enhances high availability, because NIC teaming and multipath I/O can be implemented to provide redundancy.

10GbE adds value to Dell virtualization solutions, because the increased bandwidth and high performance let NIC virtualization flourish. Virtual NICs mean less physical hardware to manage, and the higher throughput of virtual NICs means better application response time.

Dell blades offer excellent virtualization capabilities, and are optionally available with VMware preinstalled. A typical deployment would look similar to the one shown in Figure 2 below. In this illustration, a Dell PowerEdge M1000e enclosure contains PowerEdge M610 blade servers running VMware ESX™ hypervisor.

Each modular blade is equipped with multiple 1GbE and 10GbE NICs to provide for high availability and additional bandwidth. VM management traffic is segregated and managed using the 1GbE network. iSCSI SAN, Vmotion, and VM general traffic pass separately through the 10GbE network that has been divided using VLANs.

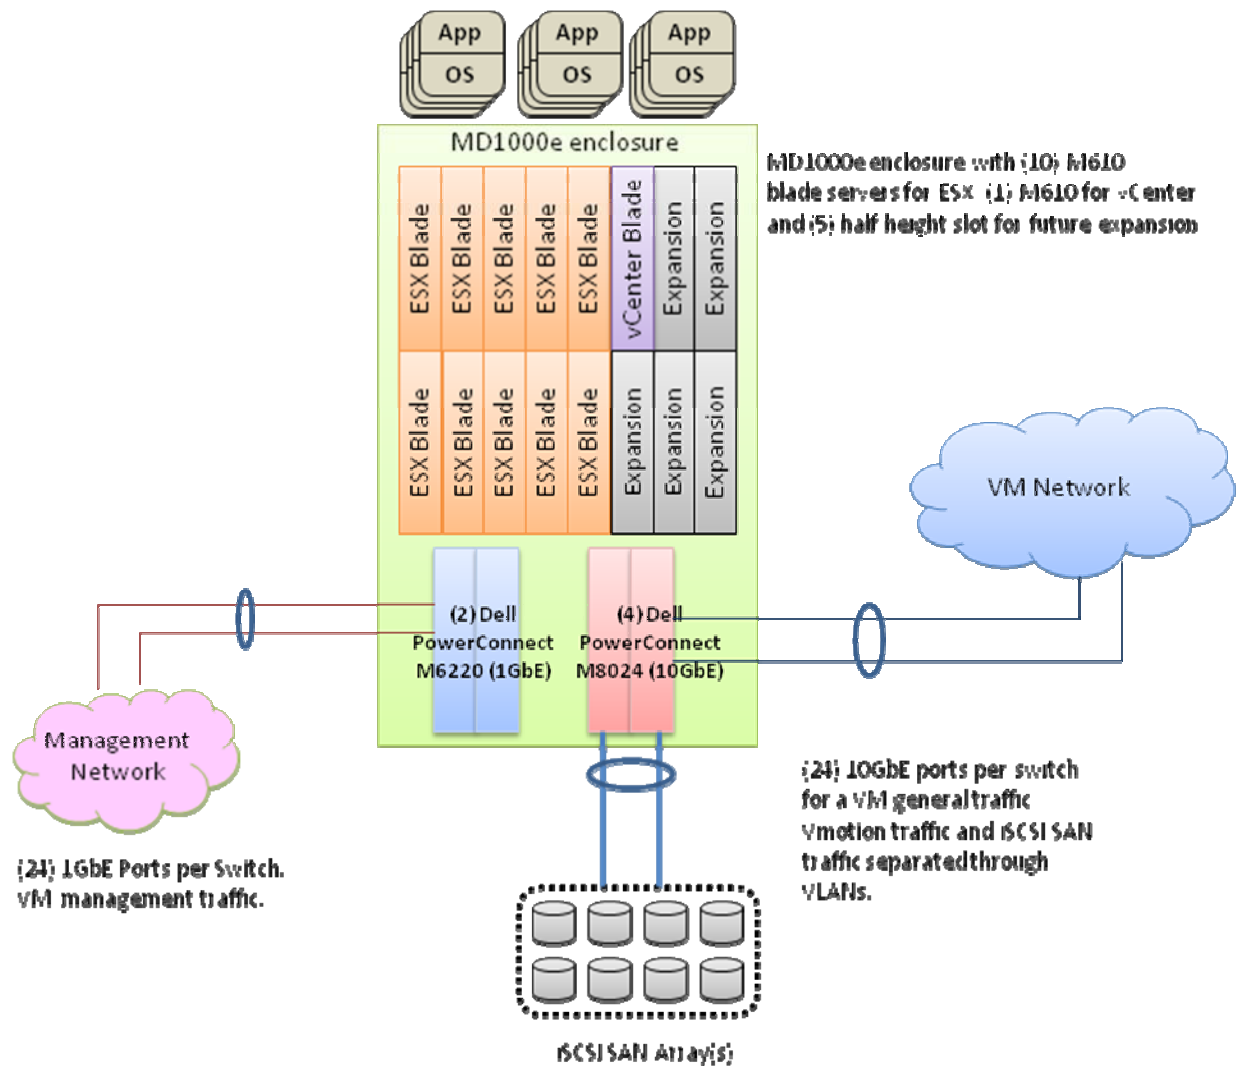


Figure 2: Typical virtualization deployment

Summary

Virtualizing a server environment has the potential to lower operating expense dramatically because it increases server utilization, resulting in needing fewer physical servers. Having fewer servers means lower management time and costs, less floor space used, and less power consumed. Today's multicore blade servers provide organizations with the ability to massively virtualize their server environment, with resulting savings; however, adding virtual machines to a server increases the network traffic and the demand for network bandwidth. The potential downside of a virtualized server environment is network bottlenecks. 10GbE NICs address this issue better than adding 1GbE NICs, since they provide ample bandwidth without the additional cabling and management complexity required by 1GbE NICs.

Dell offers complete 10GbE iSCSI solutions, including PowerEdge blade servers, blade enclosures designed for multiple upgrades, and PowerEdge blade switches. PowerEdge blade servers are available with optional VMware Infrastructure 3. Broadcom Ethernet cards based on Dell blade server adapter designs offer 10GbE speeds for PowerEdge blade servers.

Organizations can benefit from the benefits of 10GbE, consolidation on Dell PowerEdge blade servers, and virtualization with VMware. Dell has all the components, along with the expertise and services, to help organizations create solutions that improve performance in their data centers.