

# White Paper


---

## Doubling the I/O Performance of VMware® vSphere™ 4.1 with Broadcom® 10GbE iSCSI HBA Technology

This document describes the doubling of the I/O performance of vSphere 4.1 by using Broadcom 10GbE full iSCSI HBA technology. By implementing this solution, which has been developed jointly by VMware and Broadcom, data center architects and administrators can increase the number of virtual machines per physical server, lower power consumption, reduce floor space, and reduce the total cost of ownership.

# 3G<sup>C-NIC</sup>

August 2010



## Introduction

There are a number of converging trends in data centers and cloud computing. One is the migration to 10GbE LANs. Another is the growth in the deployment of iSCSI SANs, and a third is the consolidation of operations through server virtualization.

**10GbE**—The migration to 10GbE LANs in virtualized environments is underway. In a recent survey of data center managers, 90 percent of IT professionals reported that they were in some stage of planning for or were implementing 10GbE technology (see [Figure 1](#)). Based on this data, the number of server ports configured for 10GbE would more than double year-over-year for the next several years.

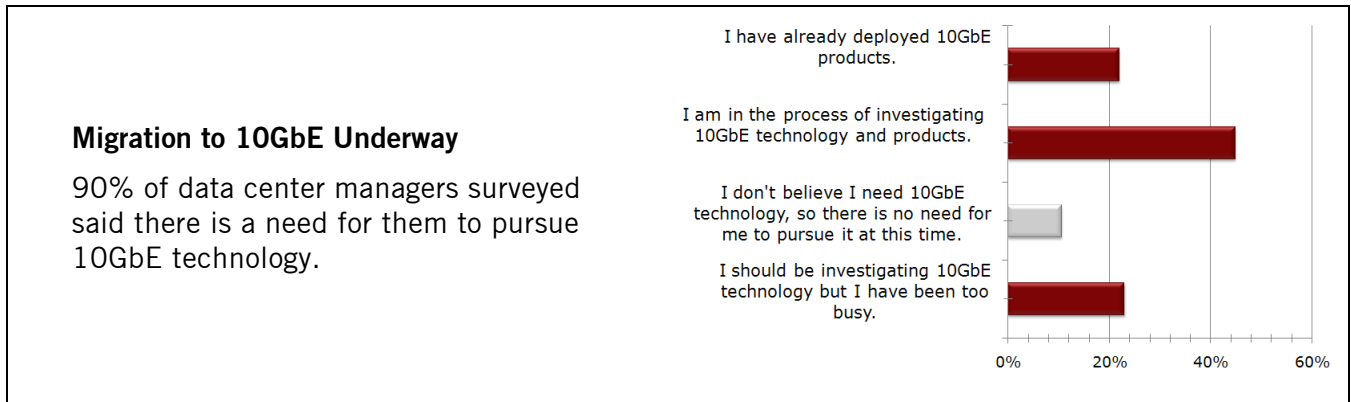


Figure 1: Survey Results

**iSCSI Storage**—iSCSI SANs are high-performance storage networks widely used in virtualization deployments. 1GbE iSCSI technology already dominates SANs in medium-sized businesses, and powerful new 10GbE arrays are making headway as the technology of choice in large data centers (see [Figure 2](#)).

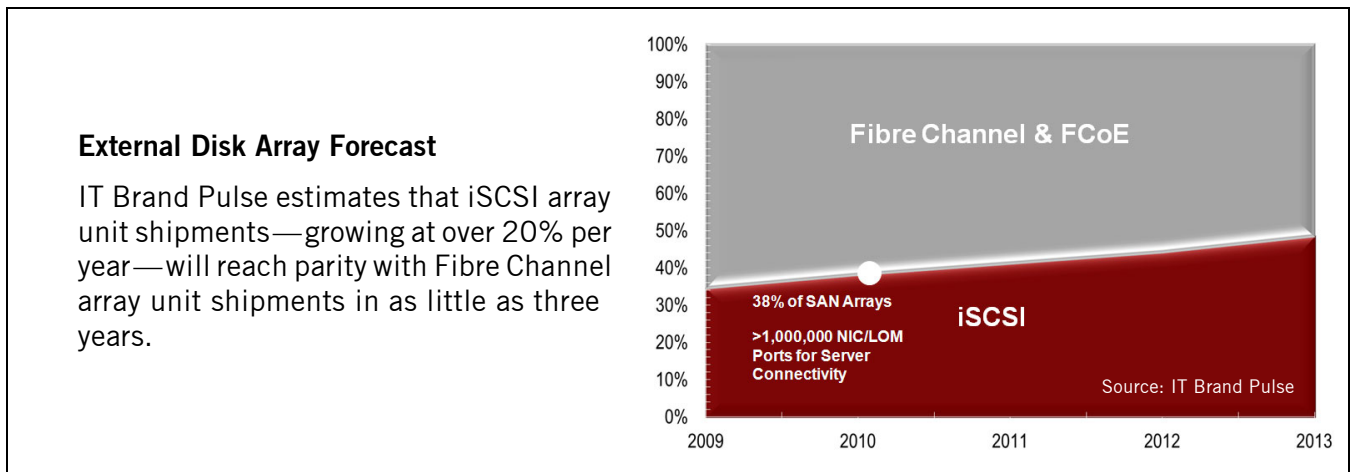
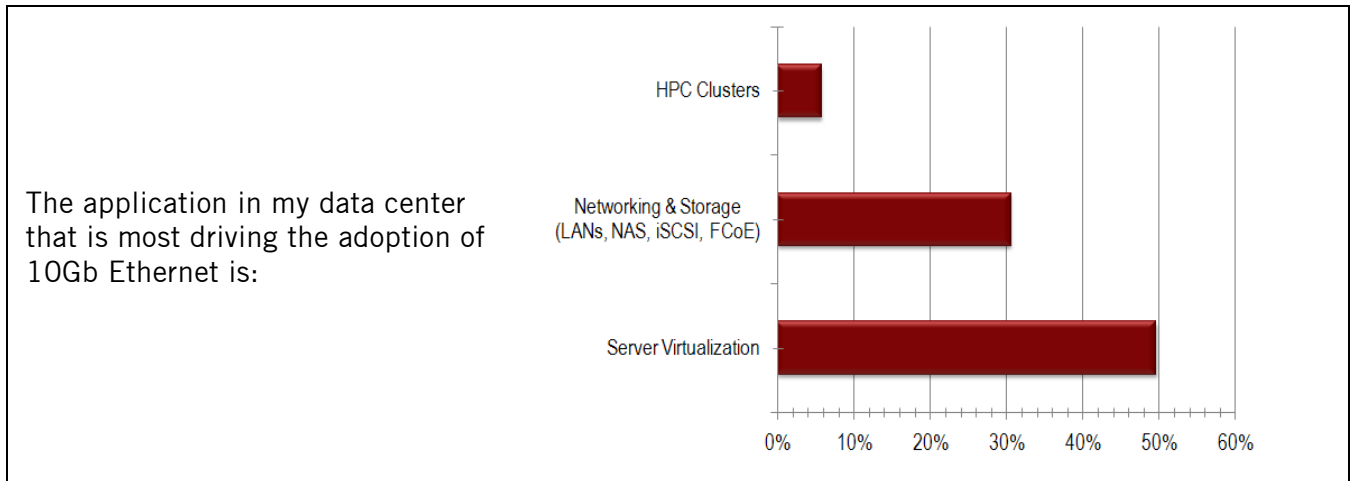


Figure 2: Estimated Growth in iSCSI Array Unit Shipments vs. FC Array Unit Shipments

**Server Virtualization**—The data center initiative most tightly coupled with 10GbE technology is server virtualization. Almost half of the IT professionals surveyed reported that *server virtualization* is the application that is driving the adoption of 10GbE in their data center the most. Clearly, IT professionals realize that new high-powered servers loaded with more virtual machines create an aggregation of I/O that requires more network bandwidth between servers and storage.



**Figure 3: Data Center Applications Driving the Adoption of 10GbE**

## VMware vSphere 4.1 Enables 10GbE Full iSCSI HBA

On July 13, 2010, VMware announced the introduction of VMware vSphere 4.1 to "advance the foundation for cloud computing." Within the category of enhancements called "increased performance through open integration with storage environments," the announcement mentions that VMware vSphere 4.1 now enables Broadcom's 10GbE full iSCSI Host Bus Adapter (HBA) solution.

Before VMware supported 10GbE iSCSI HBA, data center managers connected their VMware servers to 10GbE iSCSI SANs using a combination of disparate NIC or converged network adapter (CNA) hardware and fat iSCSI device drivers, a configuration in which most of the iSCSI protocol processing is done by the VMware ESXi™ server. Now with VMware vSphere 4.1, VMware supports the latest Broadcom converged NIC (C-NIC) that has thin iSCSI drivers and a suite of advanced hardware-based engines. The hardware-based engines within the C-NIC provide most of the iSCSI protocol processing, dramatically increasing platform efficiency and performance. Broadcom's advanced on-chip iSCSI processing is fully integrated with the VMware vSphere 4.1 stack, offering a tightly integrated and optimized data path for iSCSI.

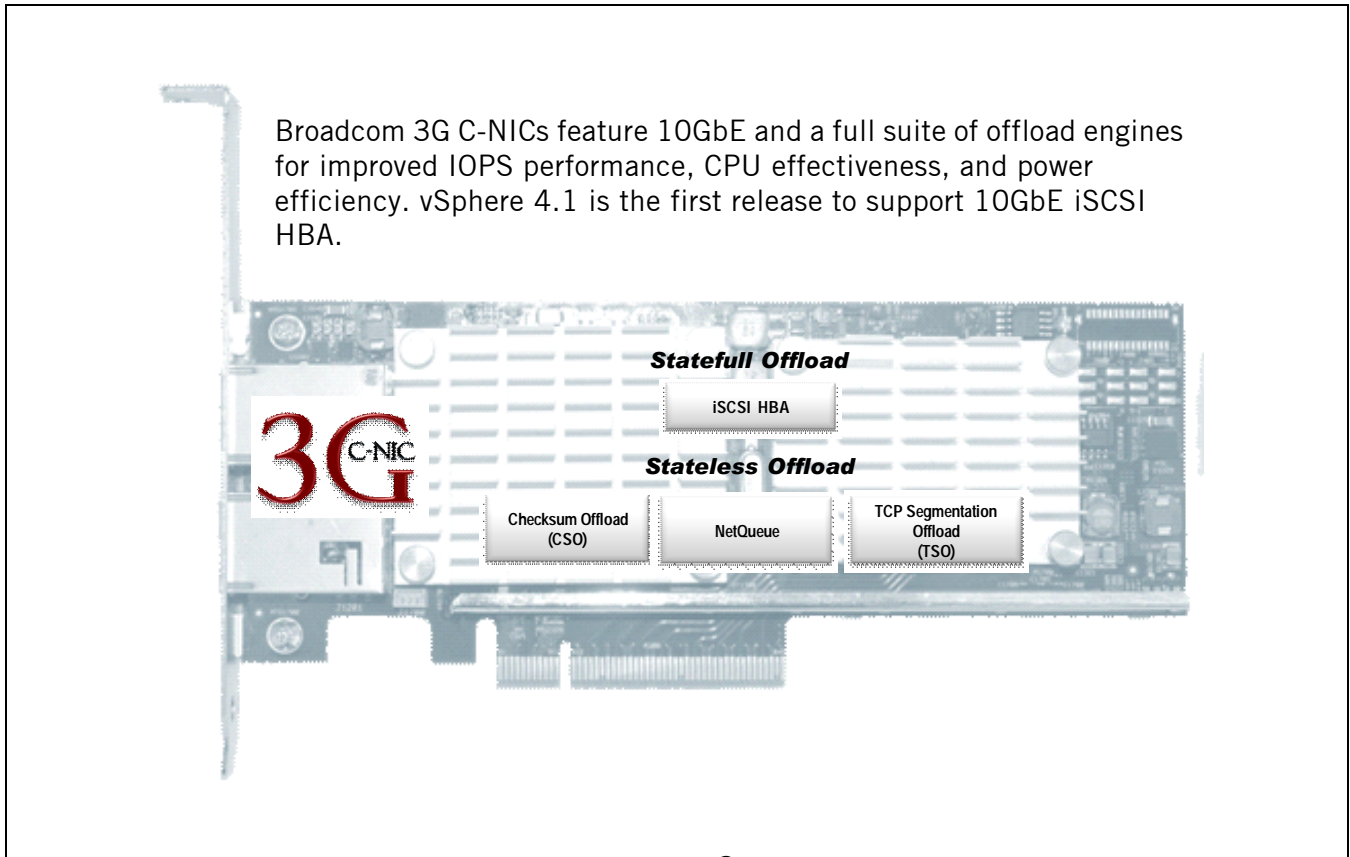


Figure 4: Broadcom NetXtreme II® C-NIC Advanced Offloads

NetXtreme II® C-NICs from Broadcom feature high performance 10GbE technology as well as fine-tuned offload engines implementing stateful and stateless offload service, as illustrated in [Figure 4](#).



**Note:** Today, Broadcom 3G C-NICs are the first and only products with 10GbE iSCSI HBA technology that have been qualified by VMware.

## Doubling IOPS Performance

Testing by Broadcom shows that iSCSI HBA offload doubles the performance of VMware vSphere I/O in the application performance sweet spots of 4 KB and 8 KB block sizes (see [Figure 5](#)). At the 4 KB block size commonly used with Microsoft Exchange deployments, the use of iSCSI HBA almost triples input/output operations per second (IOPS) performance. At the 8 KB block size commonly used with Oracle online transaction processing (OLTP) deployments, the use of iSCSI HBA more than doubles IOPS performance.

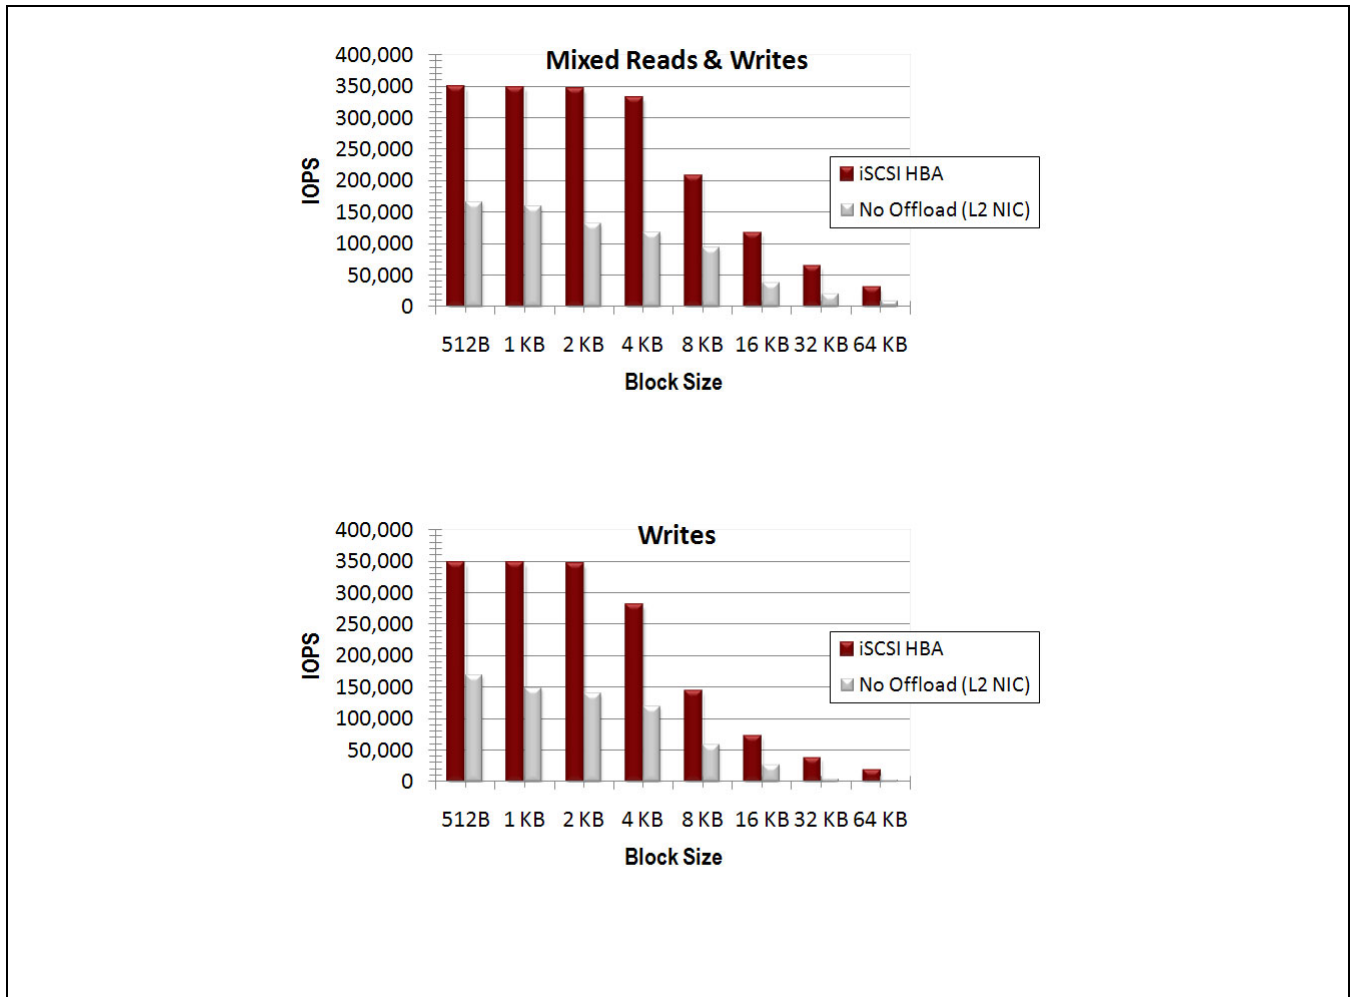


Figure 5: I/O Performance Across Block Size

**Performance Sweet Spot**—Most block sizes used with Microsoft Exchange are 4 KB and 8 KB. Oracle recommends that the database block size match or be multiples of the operating system block size. For an OLTP environment that creates many small transactions, a small block size of approximately 8 KB is needed.

## Doubling CPU Effectiveness

Testing by Broadcom also shows that iSCSI HBA offload technology doubles the performance of VMware vSphere I/O in terms of IOPS per percent of CPU utilization or CPU effectiveness (see [Figure 6](#)). At the 4 KB block size commonly used with Microsoft Exchange deployments, the use of iSCSI HBA almost triples CPU effectiveness. At the 8 KB block size commonly used with Oracle OLTP deployments, the use of iSCSI HBA again almost triples CPU effectiveness.

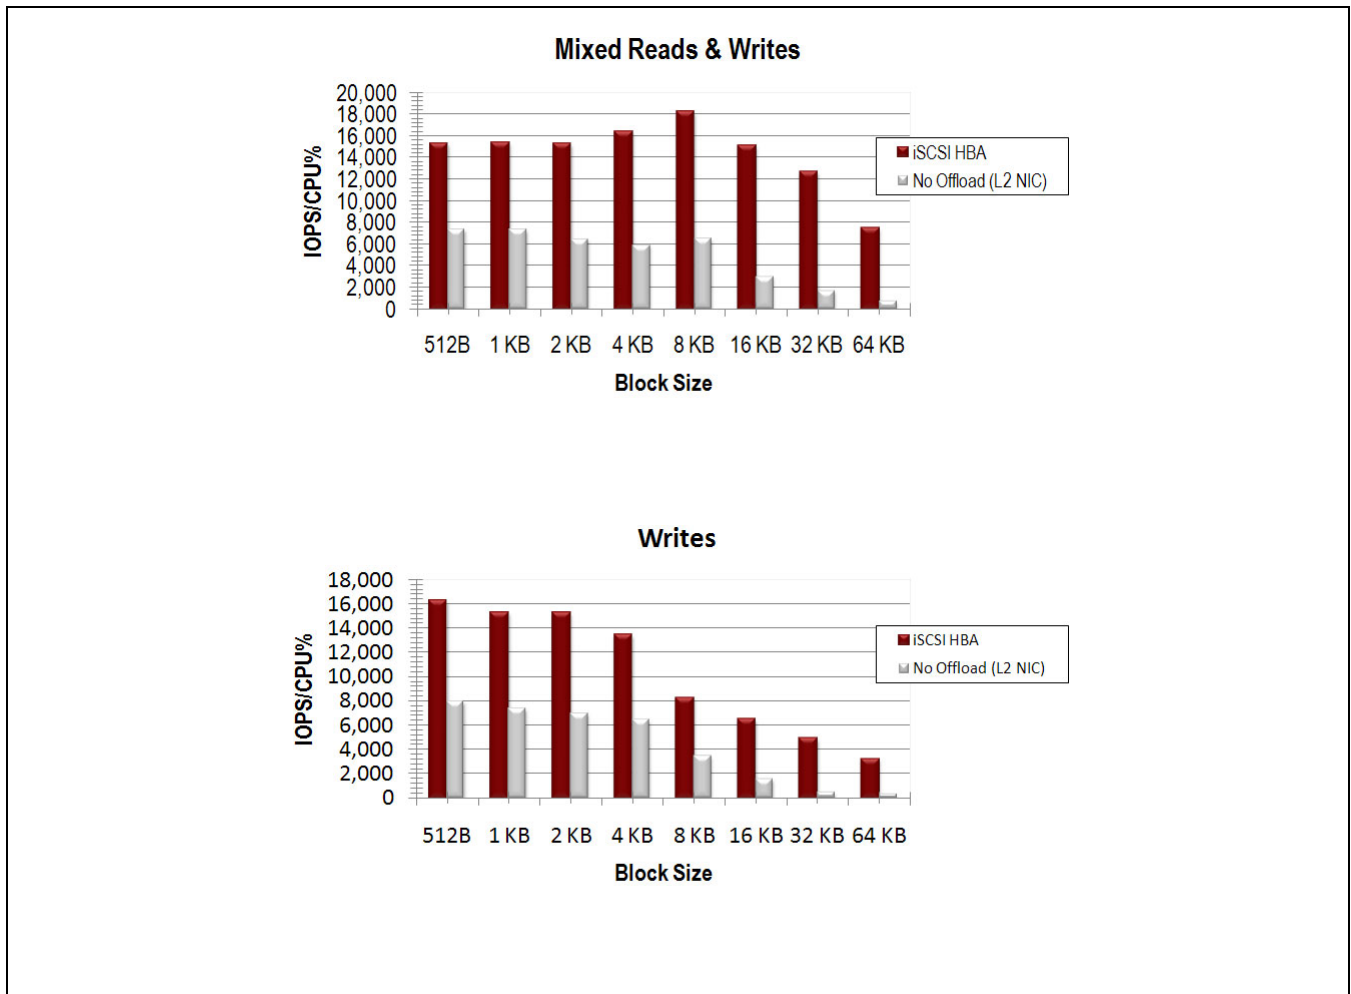


Figure 6: IOPS per Percent of CPU Utilization Across Block Size

**Greater CPU Effectiveness Means Higher VM Density**—The benefit of supporting more IOPS per percent of CPU is that more VMs per percent of CPU can be supported. iSCSI HBA doubles CPU effectiveness, which means doubling the number of VMs per server.

## Summary

If you are a server, storage, or network architect/administrator responsible for virtual server connectivity to a LAN or iSCSI SAN, iSCSI HBA simply is a "must have" feature in your SAN adapter. To learn more about VMware vSphere 4.1, 10GbE adapters, or 10GbE iSCSI storage, see [“Related News and Articles” on page 8](#).

## Checklist for High-Performance Server Virtualization

The following checklist identifies some of the key elements needed for high performance I/O in a virtualized data center:

- ✓ **VMware vSphere 4.1 Software Platform**

The first release to support 10GbE iSCSI HBA, VMware vSphere 4.1 brings many new capabilities to further extend the benefits of vSphere 4.0. The new features and enhancements to core capabilities in vSphere provide more performance optimization and easier provisioning, monitoring, and troubleshooting capabilities.
- ✓ **Converged 1GbE and 10GbE NIC (C-NIC) Architecture for iSCSI and Ethernet LAN**

Adoption of 10GbE in the data center is underway and expected to accelerate rapidly in the coming years. Together, with the move toward convergence, converged NICs offer a combination of no compromise, hardened LAN, and iSCSI SAN connectivity. Broadcom C-NICs are the first and only products with iSCSI hardware offload HBA that have been qualified by VMware.
- ✓ **10GbE iSCSI Storage**

Cost-effective and easy-to-use 1GbE iSCSI storage has been enhanced with powerful 10GbE iSCSI storage. Enterprise-class 10GbE iSCSI storage arrays are available from Dell, EMC, HP, IBM, NetApp, and many others.

## Tested Configuration

### Initiators

- Hardware
  - Dell R710 (two sockets)
  - Two Intel x5570 processors (eight cores at 2.93 GHz)
  - 24 GB RAM
  - Broadcom BCM957711 adapter
- Software
  - vSphere ESXi 4.1
  - “In box” Broadcom drivers (bnx2x, bnx2i, cnic)

### Targets

- Up to 12 1GbE iSCSI target devices

### Network

- HP Procurve 2900-24G (running T.13.63 software)

## More About vSphere 4.1 from VMware

For more information about VMware vSphere 4.1, visit:

<http://www.vmware.com/products/vsphere/mid-size-and-enterprise-business/resources.html>

## More About 10GbE from Broadcom

For more information about Broadcom's 10GbE Ethernet controllers and adapters, visit:

<http://www.broadcom.com/products/Ethernet-Controllers/Enterprise-Server>

## Related News and Articles

*VMware Advances Foundation for Cloud Computing With VMware vSphere 4.1 and Expanded Virtualization Management Portfolio*

<http://www.vmware.com/company/news/releases/vsphere-4-1.html>

*What's New in VMware® vSphere™ 4.1 — Storage*

<http://www.vmware.com/files/pdf/techpaper/VMW-Whats-New-vSphere41-Storage.pdf>

*A unified networking approach: iSCSI storage with Broadcom Controllers*

[http://www.dell.com/content/topics/global.aspx/power/en/unified\\_networking?c=us&cs=555&l=en&s=biz](http://www.dell.com/content/topics/global.aspx/power/en/unified_networking?c=us&cs=555&l=en&s=biz)

*iSCSI HBA Offload and 10GBASE-T*

[http://www.infostor.com/index/blogs\\_new/Frank\\_Berry/blogs/infostor/frank-berry\\_s\\_blog/post987\\_855043080666472366.html](http://www.infostor.com/index/blogs_new/Frank_Berry/blogs/infostor/frank-berry_s_blog/post987_855043080666472366.html)

<http://www.networkcomputing.com/next-gen-network/broadcom-unveils-worlds-1st-converged-nic-with-iscsi-offload-on-10gbase-t.php>

*FCoE I/O convergence and virtualization*

[http://www.infostor.com/index/articles/display/4348346471/articles/infostor/san/fibre-channel/2010/april-2010/fcoe-i\\_o\\_convergence.html](http://www.infostor.com/index/articles/display/4348346471/articles/infostor/san/fibre-channel/2010/april-2010/fcoe-i_o_convergence.html)

*3G C-NICs Address Mass Migration To 10GbE*

<http://www.networkcomputing.com/data-networking-management/3g-c-nics-address-mass-migration-to-10gbe.php>

Broadcom®, the pulse logo, Connecting everything®, and the Connecting everything logo are among the trademarks of Broadcom Corporation and/or its affiliates in the United States, certain other countries and/or the EU. Any other trademarks or trade names mentioned are the property of their respective owners.

Connecting  
everything®



**BROADCOM CORPORATION**

5300 California Avenue  
Irvine, CA 92617

© 2010 by BROADCOM CORPORATION. All rights reserved.

Phone: 949-926-5000

Fax: 949-926-5203

E-mail: [info@broadcom.com](mailto:info@broadcom.com)

Web: [www.broadcom.com](http://www.broadcom.com)