



By Dhiraj Sehgal

INTRODUCING BROADCOM iSCSI OFFLOAD ENGINE TECHNOLOGY FOR DELL SERVERS

Internet SCSI (iSCSI) is rapidly becoming a convergent data center technology for disparate types of networking. The Broadcom® iSCSI Offload Engine technology available in 9th- and 10th-generation Dell™ PowerEdge™ servers can free server resources to help increase performance and throughput, enabling organizations to maximize the benefits of iSCSI in their environments.

Increasing storage and networking demands have made Internet SCSI (iSCSI) technology a key component of enterprise IT infrastructures. Built on the familiar TCP/IP protocol, iSCSI helps increase storage flexibility by allowing access to content on a server through an Ethernet fabric. As a result, servers can either converge data applications and storage on the same network to help lower total cost of ownership, or dedicate one network for data applications and another for storage. In addition, by enabling an OS to be initialized across a storage area network and accessed over an IP network, multiple client systems can access the same available storage space over the network while allowing individual clients to access a different storage space over the same network.¹

By enabling organizations to use the same standard Ethernet equipment for multiple purposes—including high-speed networking, storage, clustering, and remote management—iSCSI can help lower power consumption, enhance performance, localize patch management, and conserve data center space. However, a standard server equipped with Layer 2 Ethernet controllers cannot efficiently run network, storage, and cluster traffic simultaneously over

Ethernet at the full line rate without consuming a significant amount of processing power. This heavy network traffic can then consume vital system resources required to process critical applications.

Broadcom converged network interface controllers (C-NICs) with iSCSI Offload Engine (iSOE) technology, available in 9th- and 10th-generation Dell PowerEdge servers, are designed to overcome this drawback. By handling traffic from disparate network functions in a unified Ethernet fabric and offloading iSCSI processing from host processors to LAN on Motherboards (LOMs), network interface cards (NICs), or host bus adapters (HBAs), these controllers can help increase both performance and throughput while helping optimize server processor utilization.

INTRODUCING BROADCOM iSOE TECHNOLOGY

Broadcom BCM5708C, BCM5708S, BCM5709C, and BCM5709S C-NICs with iSOE technology are featured in Dell PowerEdge server models 840, 860, 900, 1950, 2900, 2950, 2970, 6950, R200, R300, R805, R900, R905, SC1435, T105, T300, and T605. By offloading iSCSI protocol processing overhead from host processors to a BCM5708 or BCM5709 C-NIC,

Related Categories:

Broadcom

Internet SCSI (iSCSI)

Network interface card (NIC)

Visit DELL.COM/PowerSolutions for the complete category index.

¹ For more information on the iSCSI protocol, visit www.rfc-archive.org/getrfc.php?rfc=3720. For more information on the advantages of iSCSI, see the three-part series "iSCSI: Changing the Economics of Storage": "Part 1—Understanding iSCSI in Enterprise Environments," by Travis Vigil, in *Dell Power Solutions*, May 2007, DELL.COM/Downloads/Global/Power/ps2q07-20070335-Vigil.pdf; "Part 2—Deploying iSCSI in Virtualized Data Centers," by Matt Baker and Travis Vigil, in *Dell Power Solutions*, August 2007, DELL.COM/Downloads/Global/Power/ps3q07-20070401-Baker.pdf; and "Part 3—Using iSCSI in Small and Medium Businesses," by Travis Vigil, in *Dell Power Solutions*, November 2007, DELL.COM/Downloads/Global/Power/ps4q07-20070402-Vigil.pdf.

LOMs and NICs with iSOE technology can help free processor cores and memory resources to help increase I/Os per second (IOPS) and reduce processor utilization at line rate for various I/O sizes. This offloading helps increase performance for file-oriented storage, block-oriented storage, backups, database transactions, and tightly coupled distributed applications such as high-performance computing.

The minimum system requirements for iSOE functionality are as follows:

- Broadcom BCM5708 or BCM5709 add-in iSCSI HBA or integrated LOM
- Microsoft® Windows Server® 2003, Microsoft Windows Server 2008, Red Hat® Enterprise Linux® 5, or Novell® SUSE® Linux Enterprise Server 10 OS, which can be factory-installed on supported Dell PowerEdge servers for an iSOE-enabled LOM or installed as part of an add-in iSCSI HBA kit
- iSOE hardware license key installed on LOMs before server boot (this license key is built into select add-in NICs and is an optional feature on Dell PowerEdge servers equipped with Broadcom BCM5708C-, BCM5708S-, BCM5709C-, and BCM5709S-based LOMs)
- Activation of embedded iSOE feature set (a Broadcom LOM is integrated into the base Dell PowerEdge server configuration at no additional cost, but organizations must elect to activate iSOE as a purchase option at the point of sale; licenses are not optional on all Broadcom controller-based NICs, but certain Broadcom controller-based NICs contain the license)
- Broadcom NetXtreme® II drivers for specific operating systems (pre-installed or provided for installation)

NICs operate in non-iSOE mode if any of the preceding components are missing or if the system is running an unsupported OS. After organizations have activated iSOE as purchase option, no additional

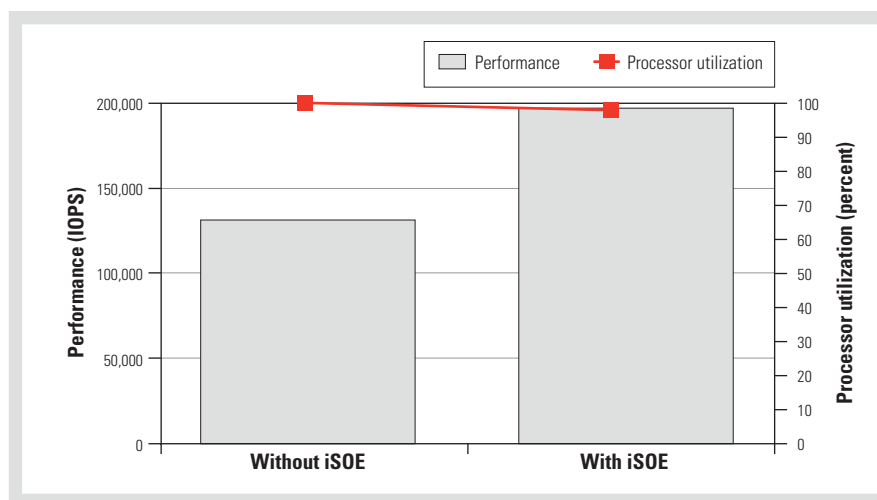


Figure 1. Performance and processor utilization with and without Broadcom iSOE technology

work is typically required to enable the hardware feature and the iSCSI software components; for iSOE-capable NICs, no additional hardware key is required.

Organizations that want to use iSOE functionality in their LOMs can purchase iSOE keys as an upgrade to the initial basic configuration. To enable iSOE, they must install the appropriate iSOE key in the TCP/IP Offload Engine (TOE) key connector on the server system board and update the LOM firmware and drivers to the latest versions.²

EVALUATING iSOE PERFORMANCE AND THROUGHPUT

Increased performance for key metrics—such as IOPS and processor utilization—was a key focus during iSOE development. The result is that for applications with small I/O sizes, such as Web servers and file database servers, Broadcom iSOE technology can provide higher bandwidth and lower processor utilization than a software solution such as the Microsoft iSCSI Software Initiator. For applications with I/O sizes greater than or equal to 4 KB, such as remote storage data backup, iSOE can provide tremendously increased bandwidth and correspondingly reduced processor utilization. In addition, iSOE technology is designed to address both I/O reads and writes, enhancing

performance for both types of operations as well as freeing processor cycles for other critical applications such as databases, media streaming, and file sharing.

To demonstrate the performance advantages of Broadcom C-NICs with iSOE technology, in October 2007 Broadcom engineers configured two Dell PowerEdge 2950 servers with dual-core Intel® Xeon® 5160 processors at 2.66 GHz, 8 GB of double data rate 2 (DDR2) RAM at 2.66 GHz, a frontside bus at 1,066 MHz, and Microsoft Windows Server 2003 with Service Pack 1. Each server had two on-board Broadcom BCM5708C C-NICs, with two software-based iSCSI targets connected through switches to each server; one server had iSOE enabled on the C-NICs, while the other had iSOE disabled and used the Microsoft iSCSI Software Initiator. The benchmark used in the tests included the Softpedia Iometer 2006.07.27 network performance analyzer (1 worker per target and 32 outstanding I/Os).

Figure 1 shows the performance results for iSCSI read operations using a 512-byte I/O size along with processor utilization. In these tests, the server with iSOE technology handled 50.2 percent more IOPS than the server without iSOE technology, at comparable levels of processor utilization.

² Dell uses four part numbers for iSOE keys: YR232 for single-port TOE and iSOE technology, WY733 for dual-port TOE and iSOE technology, CR774 for quad-port TOE and iSOE technology in Dell PowerEdge R805 and PowerEdge R905 servers, and C402D for quad-port TOE and iSOE technology in other supported PowerEdge servers. For more information on installing iSOE software components, refer to the user's guide included with the network driver package.


Figures 2 and 3 show the throughput results for sequential iSCSI read and write operations, respectively, along with processor utilization. For sequential read operations, the iSOE-enabled server provided up to 28 percent higher throughput and up to 85 percent lower processor

utilization than the iSOE-disabled server using the Microsoft iSCSI Software Initiator; at large I/O sizes, in fact, the iSOE-enabled server had a processor utilization of less than 10 percent. For sequential write operations, the iSOE-enabled server provided up to 38 percent

higher throughput with up to 85 percent lower processor utilization than the iSOE-disabled server using the Microsoft iSCSI Software Initiator, with the iSOE-enabled server again having processor utilization of less than 10 percent at large I/O sizes.

These results demonstrate the performance and throughput increases possible when using Broadcom C-NICs with iSOE technology compared with a software-based iSCSI solution such as the Microsoft iSCSI Software Initiator. Organizations should keep in mind that real-world performance will vary based on the specific configuration running in a given environment.

OPTIMIZING iSCSI ENVIRONMENTS

iSCSI technology can offer significant benefits in enterprise IT environments, including reduced power consumption, increased performance, and reduced total cost of ownership. Broadcom C-NICs with iSOE technology are designed to enhance these benefits even further, enabling organizations to unify disparate network functions over Ethernet while increasing performance compared with software-based iSCSI initiators, minimizing the burden on host processors, and enabling the efficient use of enterprise resources. 

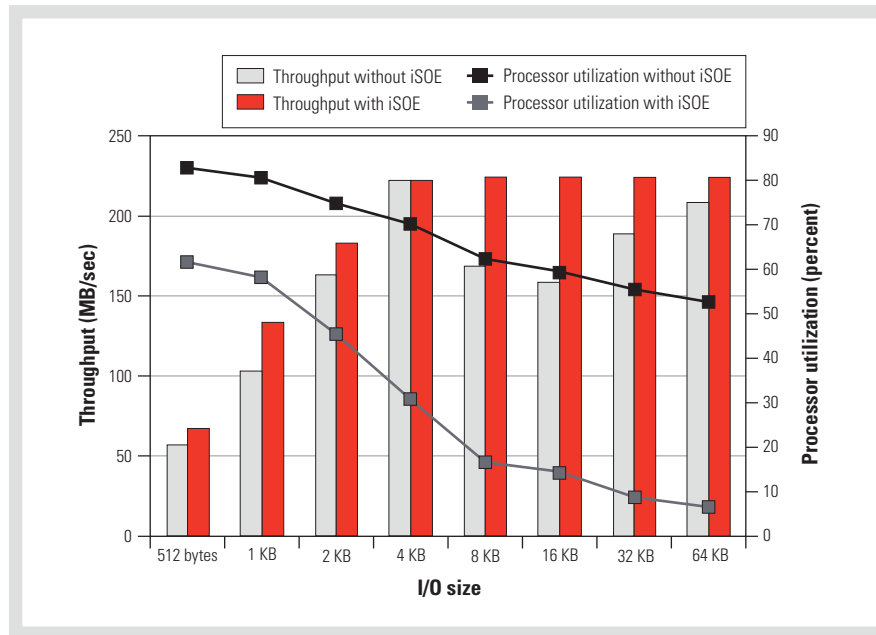


Figure 2. Throughput and processor utilization for sequential iSCSI read operations with and without Broadcom iSOE technology

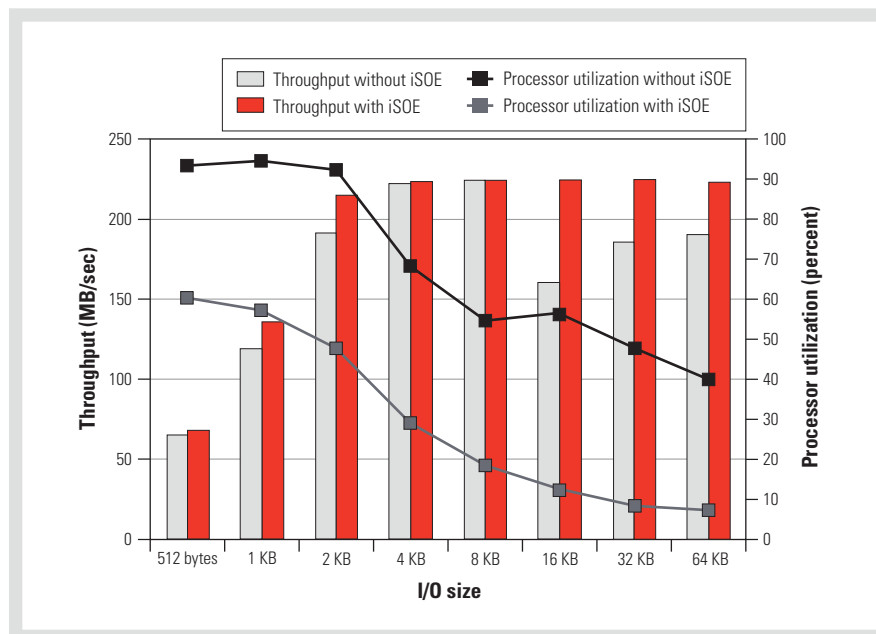


Figure 3. Throughput and processor utilization for sequential iSCSI write operations with and without Broadcom iSOE technology

Dhiraj Sehgal is a senior product line manager for Ethernet controllers at Broadcom. Dhiraj has an M.S.E.E. from North Carolina State University, Raleigh.

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Broadcom controllers:
www.broadcom.com/products/Enterprise-Networking/Gigabit-Ethernet-Controllers

Dell PowerEdge servers:
DELL.COM/PowerEdge