



A unified networking approach to iSCSI storage with **Broadcom controllers**

By Dhiraj Sehgal, Abhijit Aswath, and Srinivas Thodati

In environments based on Internet SCSI (iSCSI) and 10 Gigabit Ethernet, deploying Dell™ PowerEdge™ servers with Broadcom® NetXtreme II® converged network controllers can help IT departments to eliminate network bottlenecks, reduce power consumption, and maximize data center efficiency.

Use of the Internet SCSI (iSCSI) protocol has been growing significantly in enterprise data centers, largely because it offers a variety of advantages over traditional Fibre Channel storage technologies. Because iSCSI provides a simple method for transporting SCSI commands, data, and status messages over standard TCP/IP networks, it lets organizations take advantage of existing infrastructure and knowledge bases while using cost-effective, familiar components. It is also standards based, which facilitates industry adoption and helps ensure interoperability. And its performance can scale up along a common technology path—Ethernet.

But the widespread use of iSCSI has also created challenges as network usage has grown and evolved. For example, many organizations are running an increasing number of rich-content and data-intensive applications or virtualized environments on their servers, which can significantly increase the iSCSI-based network traffic in a data center. To help keep up with this growth, many IT departments have begun moving from Gigabit Ethernet (GbE) network controllers

to 10 Gigabit Ethernet (10GbE) controllers. Under heavy workloads, however, traditional controllers can consume significant amounts of server processing power, reducing the amount available for critical applications. In addition, the traditional approach typically requires using separate network controllers to handle different types of tasks (such as high-speed networking, storage, and clustering), which leads to a burgeoning number of devices, which drives up cost and complexity—and limits the ability of IT departments to take advantage of 10GbE controllers.

The Broadcom NetXtreme II family of 10GbE converged network interface controllers (C-NICs) is designed to address these problems. Available in Dell PowerEdge servers as LAN on Motherboards (LOMs), mezzanine cards, and stand-up network interface cards (NICs), these controllers provide the speed and efficiency to support heavy network traffic workloads—enabling IT departments to simplify their systems and provide network, storage, and clustering capabilities over existing TCP/IP and Ethernet infrastructures, while also helping reduce power consumption and enabling highly efficient use of processing resources.



“As converged controllers, Broadcom NetXtreme II C-NICs can handle multiple types of networking tasks simultaneously.”

Easing the network processing load

As converged controllers, Broadcom NetXtreme II C-NICs can handle multiple types of networking tasks simultaneously. A traditional server is typically equipped with four GbE controllers, along with other storage adapters in some cases. Simultaneously running network, storage, and clustering traffic at high rates typically requires multiple adapters and a large number of CPU cycles. The Broadcom C-NIC approach, in contrast, enables administrators to use a single 10GbE port for multiple traffic types. Alternatively, if administrators choose to run just one traffic type, they can do so without having to deploy custom hardware and software, which helps simplify IT deployment and management. In addition, NetXtreme II C-NICs available in Dell PowerEdge servers can support Microsoft® Windows®, Linux®, VMware®, and other platforms, helping maximize deployment flexibility.

To help reduce the burden of packet header processing on the host server, NetXtreme II C-NICs provide iSCSI host bus adapter (HBA) functionality with iSCSI Offload Engine (iSOE) technology (see Figure 1). By offloading iSCSI header processing from host processors to HBAs, these controllers can help optimize server processor utilization while helping increase both performance and throughput for file-oriented storage, block-oriented storage, backups, database transactions, and tightly coupled distributed applications such as high-performance computing workloads. The iSOE technology is designed to free up host processor cores and memory resources and increase I/Os per second (IOPS)—including, in the Broadcom test environment detailed in the next section of this article, enabling up to 400,000 IOPS at 10GbE line rates over a single Ethernet port while substantially reducing processor utilization.

iSOE enables NetXtreme II C-NICs to effectively handle both TCP/IP and iSCSI processing. By offloading the TCP/IP and iSCSI stacks, the controller does not need to compete with upper-layer software such as e-mail or Web applications for CPU cycles: iSCSI performance is unaffected by application workload. The convergence of block storage and network processing over a standard TCP infrastructure helps eliminate the need for a separate storage adapter and additional cabling, while providing performance and reliability comparable to Fibre Channel at a significantly reduced cost.

Table with 8 columns: Broadcom controller model, Ports, Broadcom part number, Dell part number, Card type, I/O bus, Physical interface, iSCSI features. It lists various Broadcom NetXtreme II C-NIC models and their specifications.

Figure 1. iSOE-enabled Broadcom NetXtreme II C-NICs available in Dell PowerEdge servers

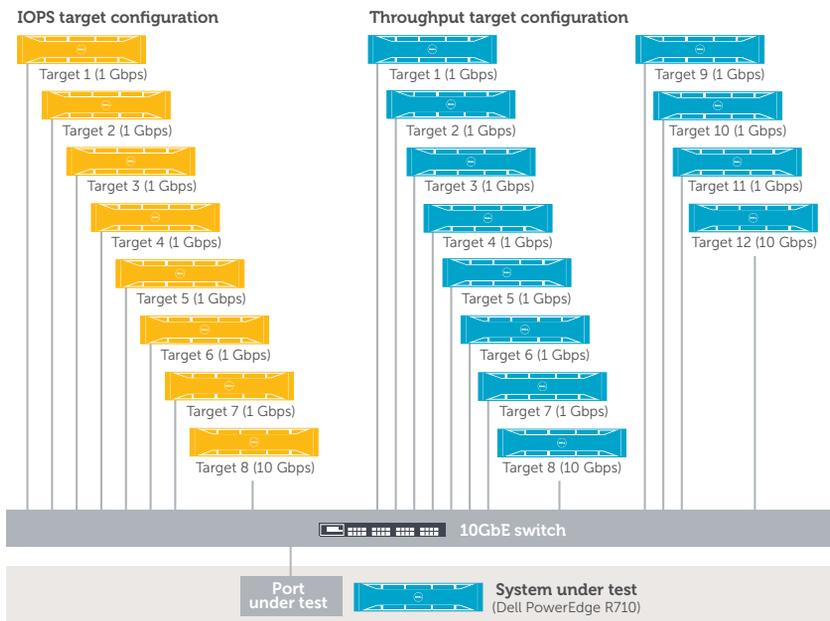


Figure 2. High-level configuration used in the test environment

Easing management and reducing costs are ongoing goals in the data center—and to that end, the NetXtreme II family helps to simplify administration of controllers across the environment through the Broadcom Advanced Control Suite (BACS) 3 management application, which provides a single platform for network and iSCSI HBA I/O management. BACS 3 includes at-a-glance status reports of LAN adapters and controllers in a system, network testing to confirm connectivity to a remote station, and detailed performance statistics on adapters and controllers.

Evaluating the converged approach

To help IT departments understand the enhancements and benefits of the converged approach, in October 2009 Broadcom's performance laboratories ran a series of tests to analyze the power consumption, processor utilization, throughput, and processor effectiveness (IOPS per CPU cycle) of a Broadcom NetXtreme II C-NIC and a third-party NIC with an iSCSI software initiator.

Figure 2 shows the test environment, which was designed to evaluate the maximum performance value organizations could expect in this type of deployment. The system under test was a Dell PowerEdge R710 server with two quad-core

Intel® Xeon® X5570 processors at 2.93 GHz, 12 GB of RAM, and the Microsoft Windows Server® 2008 Enterprise Edition OS. The setup included one NetXtreme II BCM57711 dual-port 10GbE C-NIC with iSCSI HBA functionality enabled, along with a third-party NIC using an iSCSI software initiator. The tests used 8 targets for IOPS measurement and 12 targets for throughput measurement; the target systems were PowerEdge R710 servers with two quad-core Intel Xeon X5570 processors at 2.93 GHz, 2 GB of RAM, the Red Hat® Enterprise Linux 5.2 OS, and iSCSI Enterprise Target (IET) 0.4.16 software.

The tests utilized the Iometer I/O subsystem measurement and characterization tool, which is designed to evaluate the performance of storage and networked applications, perform stress tests on storage and network devices, and predict storage and networked application performance. The test environment used version 2006.07.27 of this tool, configured for 128 outstanding I/Os. Power consumption was measured as the average of three one-second samples taken at the power inlet.

As these tests showed, using NetXtreme II C-NICs as iSCSI HBAs can provide a variety of advantages in data center environments, including helping to reduce power consumption and increase processing efficiency. In the test environment, the NetXtreme II C-NIC in iSCSI HBA

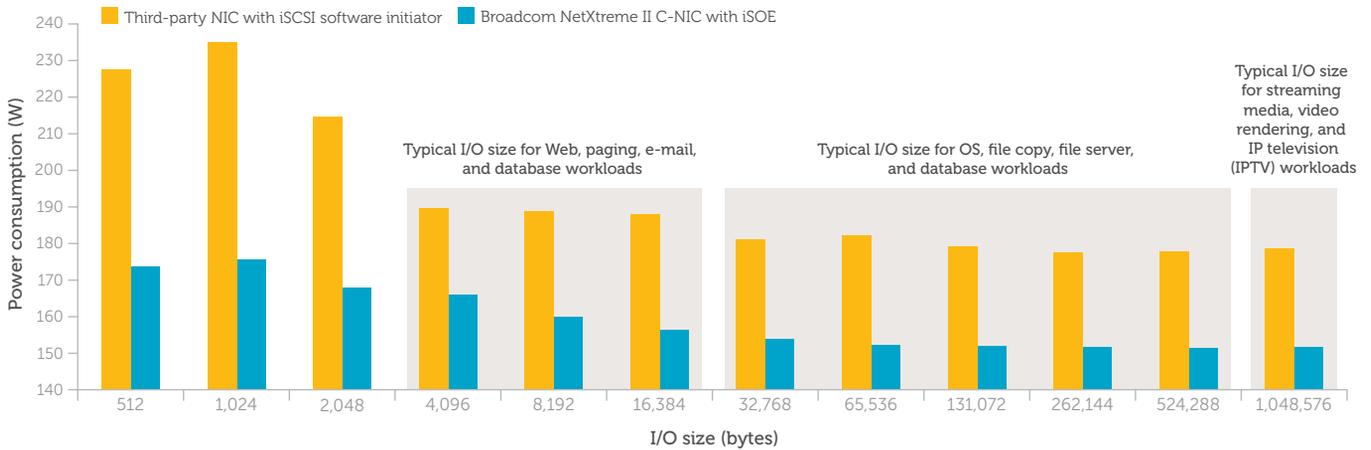


Figure 3. Power consumption at different I/O sizes for the third-party NIC and the Broadcom NetXtreme II C-NIC

mode provided a system-level savings of approximately 60 W per port compared with the third-party NIC (see Figure 3)—a difference that can lead to dramatic cost savings. For example, consider a data center with 2,500 servers containing four 10GbE ports each. Assuming the same load characteristics over the same period as the test environment, using this NetXtreme II C-NIC

in place of the third-party NIC could reduce power consumption by 600 kW at the port level alone; when factoring in an example Power Usage Effectiveness (PUE) value of 1.8 and assuming matching savings in power consumption across the supporting infrastructure, the total reduction would reach 1,080 kW. At an energy cost of US\$0.088/kWh, then, the cost savings over the course of a year

would amount to US\$832,550. IT departments could also take advantage of these enormous power savings to help them rightsize the data center's supporting infrastructure, helping to reduce total cost of ownership and improve PUE efficiency.

In addition to reducing power consumption, the NetXtreme II C-NIC reduced processor utilization by

approximately 30 percent at large I/O sizes (2–64 KB) and delivered three times the IOPS per CPU cycle of the third-party NIC, enabling highly efficient processing of storage workloads (see Figure 4). This efficient operation in turn can provide increased IOPS performance to end users for storage workloads while also freeing up processing power to support additional applications

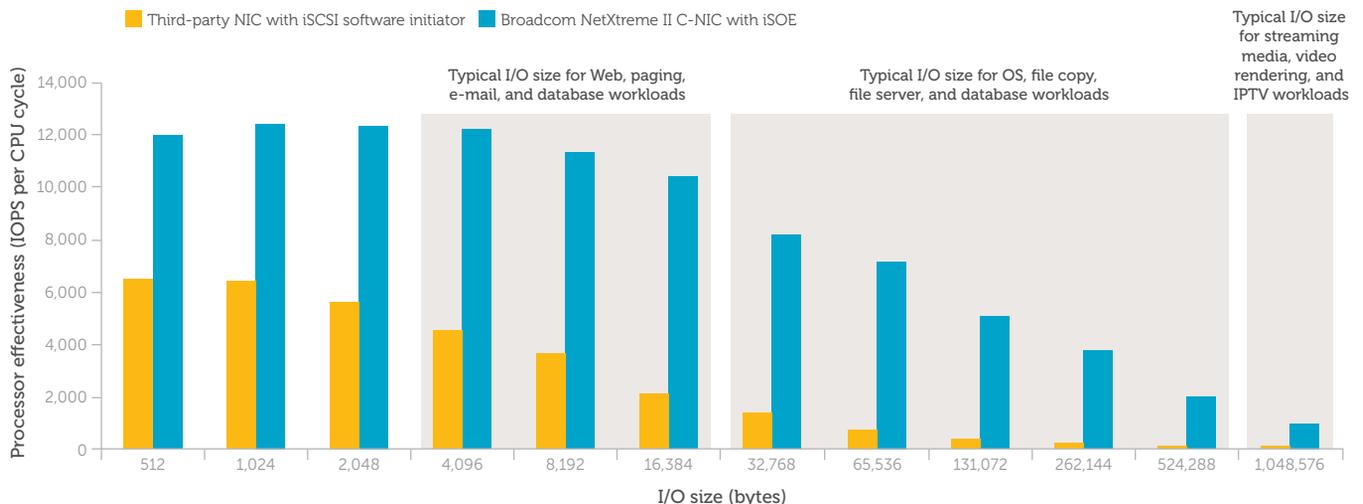


Figure 4. Processor effectiveness at different I/O sizes for the third-party NIC and the Broadcom NetXtreme II C-NIC

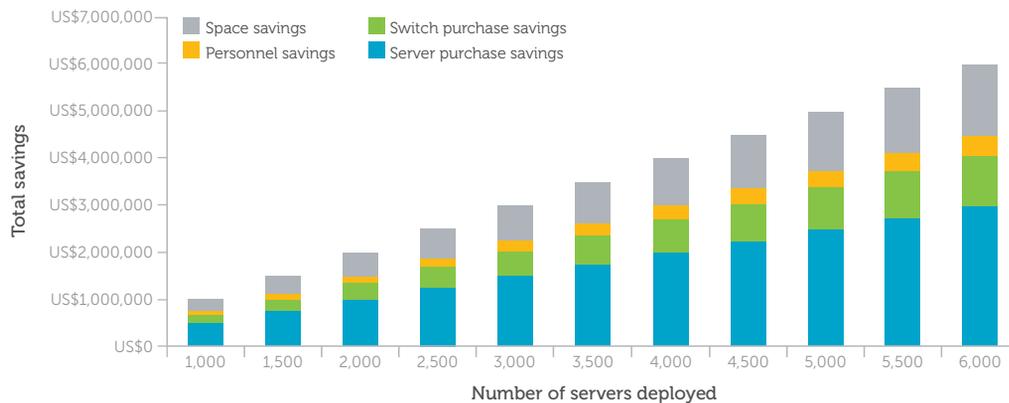


Figure 5. Example onetime savings possible from proper planning and efficient processor utilization with Broadcom NetXtreme II C-NICs

or increased usage of high-performance applications—all without increasing the total cost of operations.

This efficient operation can also help IT departments during installation, allowing them to plan for increased load capacity. In a data center with 5,000 servers, for example, the onetime savings from proper planning and improved utilization could reach US\$5 million (see Figure 5).¹

Reducing costs through converged networking

Converging multiple functions over one wire using Dell PowerEdge servers with Broadcom NetXtreme II C-NICs can help simplify networking, increase processor effectiveness, and reduce total cost of ownership. In addition, by enabling IT departments to take advantage of existing Ethernet infrastructure and avoid the need for stand-alone HBA cards, these converged controllers help lower the cost of iSCSI acquisition, deployment, and management—avoiding the need to maintain a separate storage infrastructure and the need for specialized training.

The Broadcom NetXtreme II family can also enable organizations to take advantage of 10GbE iSCSI networking while providing the flexibility to either converge data and storage

traffic onto a single network or use a dedicated network for each. This approach holds the promise of accelerating iSCSI networking adoption and helping IT departments pursue a variety of critical initiatives, from environmentally friendly computing to virtualization and consolidation—and the constant need to keep costs down while maximizing data center efficiency. **PS**



Dhiraj Sehgal is a senior product line manager for Ethernet controllers at Broadcom.

Abhijit Aswath is a senior product line manager for Ethernet controller software at Broadcom.

Srinivas Thodati is a senior product marketing manager for PowerEdge M-Series servers at Dell.

Learn more



Broadcom Ethernet controllers:
broadcom.com/products/Ethernet-Controllers



Dell PowerEdge servers:
dell.com/poweredge

¹Based on a 30 percent reduction in server count from an average 30 percent reduction in processor utilization. Calculations assume an average server cost of US\$5,000, a server lifetime of 3 years, 1 switch for every 2.5 servers, an average switch cost of US\$1,500, 12 servers per rack, 10 square feet of rack floor space at a commercial cost of US\$1,000 per square foot, 1 administrator for every 250 servers, and an average administrator salary of US\$60,000.