

Broadcom® FCoE/iSCSI and IP Networking Adapter Evaluation

Evaluation report prepared under contract with Broadcom Corporation

Introduction

Enterprises are moving towards 10 Gigabit networking infrastructures and are examining the uses for multiprotocol adapters, especially as server virtualization continues to be deployed. With an increasing number of operating systems and/or applications installed in today's servers, the need for network and storage adapters that can support multiple protocols becomes more important.

IT administrators are also turning to 10 Gigabit Ethernet (10GbE) technologies for cost-effective and flexible methods of addressing growing network traffic demands. A key component of the 10GbE network is the network adapter with network protocol processing offload capabilities for greater host processor utilization. 10GbE networking adapters by Broadcom allow servers to fully leverage the available 10GbE bandwidth, while reducing processor utilization. As a result, 10GbE end-to-end performance is comparable to more specialized and costlier cloud based data center server to fabric interconnects. With 10GbE, organizations can expand application capabilities, increase scalability and improve responsiveness to address dynamic business environments.

Broadcom® has introduced its BCM957810S dual-port Converged NIC that supports 10 Gbps IP networking with network protocol offload capabilities, Fibre Channel over Ethernet (FCoE) with full hardware offload and iSCSI with full hardware offload. This adapter fits the description of a Converged Network Adapter (CNA). Broadcom commissioned Demartek® to evaluate this adapter and compare its performance to other competitive 10 Gbps adapters.

Evaluation Environment

For the storage protocol offload testing, the evaluation environment consisted of a Windows Server 2008 R2 SP1 platform with an Intel® Xeon® E5-2690 processor and 32GB of RAM using various 10 Gbps adapters communicating with two different DRAM-based storage targets. The FCoE storage target consisted of 32 DRAM LUNs and the iSCSI storage target consisted of 24 DRAM LUNs. The open source I/O load generator IOMeter was used to drive the test workloads. As is customary for storage performance tests, throughput results are measured in MegaBytes per second.

For the IP networking tests, the evaluation environment consisted of a server with the 10GbE adapters installed, connected to a 10GbE/1GbE network environment with 16 clients, and separately, the same server and 10GbE adapters connected to specialized networking hardware. As is customary for network performance tests, throughput results are measured in megabits per second.

The 10 Gbps adapters shown in the table below were tested. The most current drivers and/or firmware were installed for each of the adapters for these tests.

Vendor	Adapter Model	Hardware Offload
Broadcom®	BCM957810S	Yes
Emulex®	OCe11102-F & OCe11102-I	Yes
Intel®	X520	No
QLogic®	QLE8242	Yes

Evaluation Summary

Storage Offload Performance

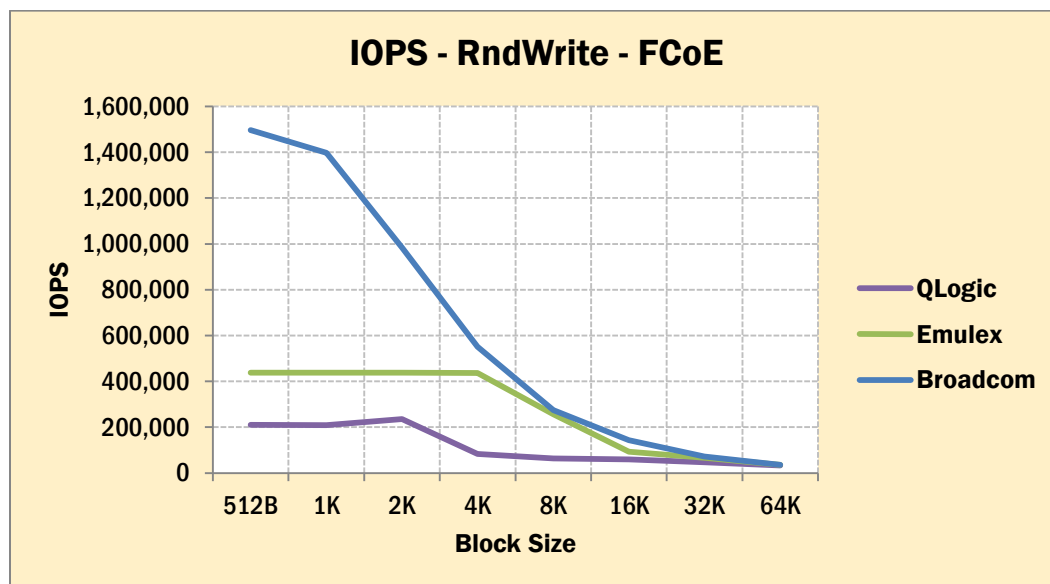
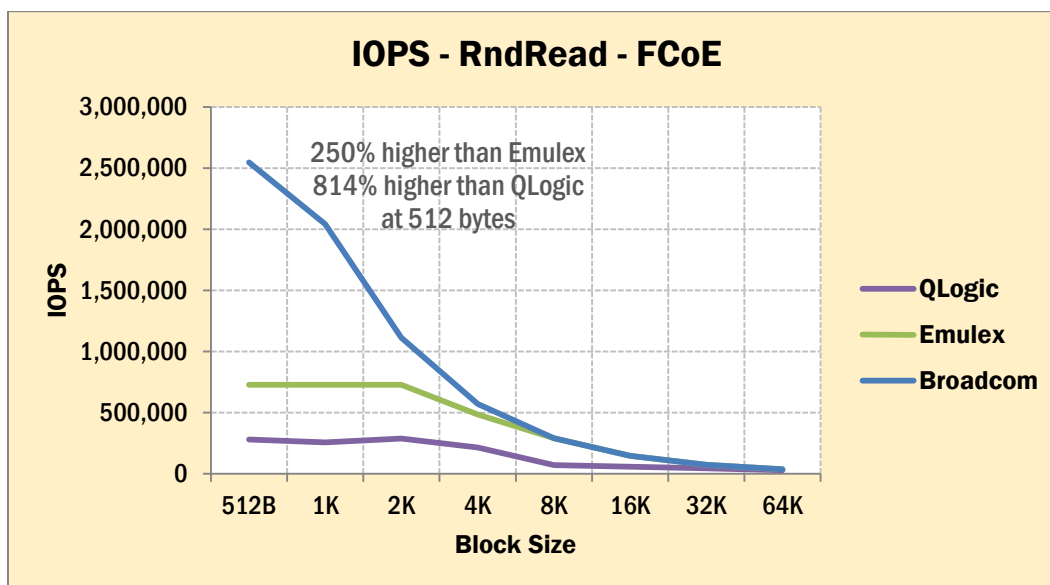
We found that the Broadcom BCM957810S adapter outperformed the competitive adapters. We were able to achieve approximately **2.5 million IOPS** for FCoE random reads and more than 1.5 million IOPS for iSCSI random reads using the Broadcom BCM957810S adapter. Broadcom has demonstrated very strong performance and good CPU efficiency when compared to both full-hardware-offload adapters and non-hardware-offload adapters.

IP Networking Performance

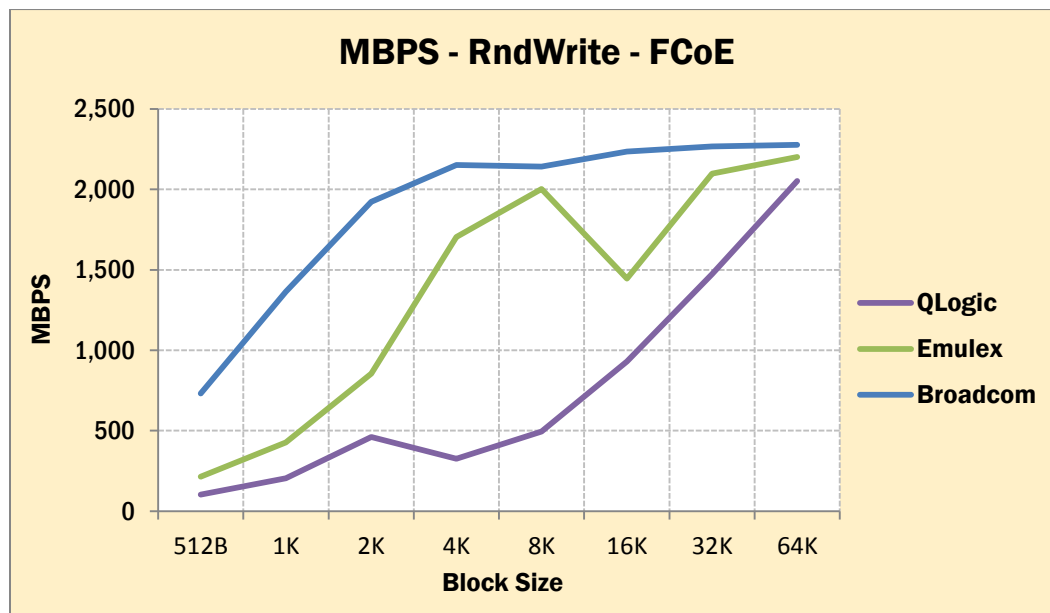
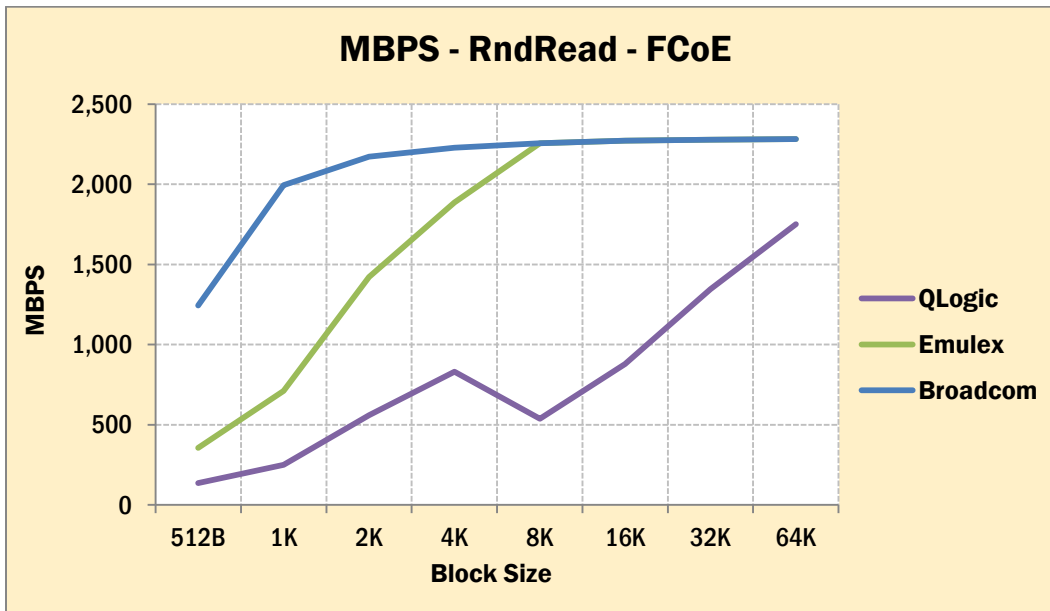
We found that the Broadcom BCM957810 10GbE adapter outperformed the equivalent competitive adapters in our tests, for the transmit and receive tests as well as the small packet tests.

1 – FCoE Performance: Full Hardware Offload Comparisons

As enterprises review their long-term server, networking and storage plans, consideration must be given to converged networking technologies such as Fibre Channel over Ethernet (FCoE). The long-term economics of reduced numbers of adapters, cables and disparate switches can be a compelling reason to explore the use of FCoE. The Broadcom adapter significantly outperformed the competitive adapters in the smaller block sizes in terms of I/Os per second (IOPS).

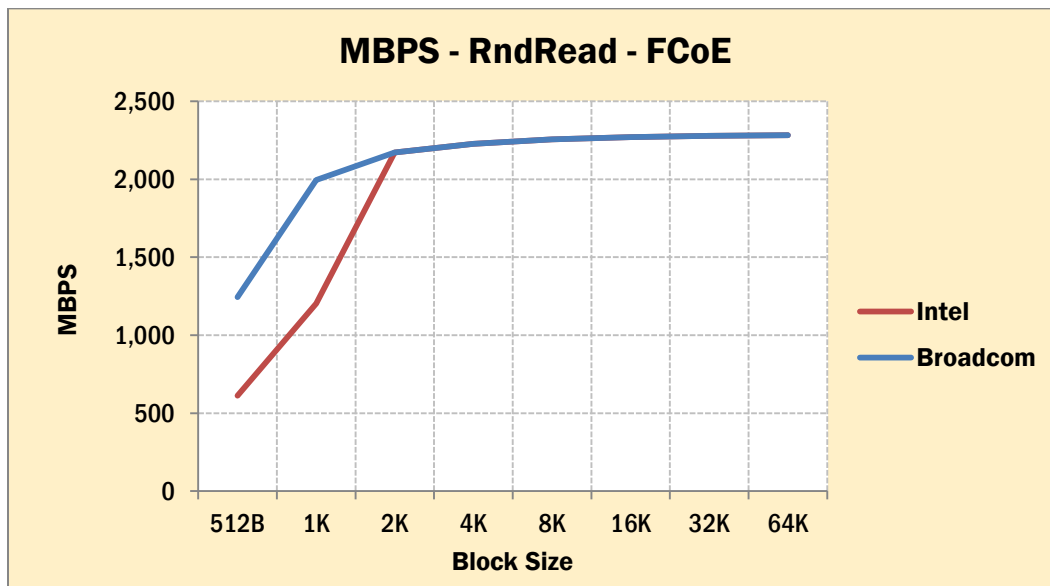
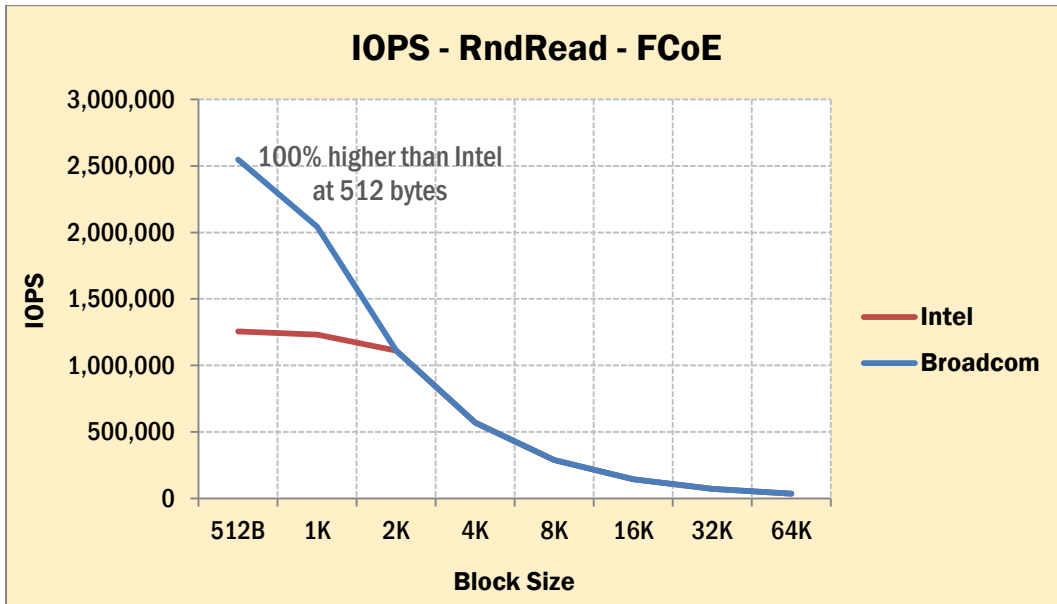


The Broadcom adapter achieved line rate at smaller block sizes than the competitive adapters in terms of bandwidth throughput. At 8K block sizes and above, one of the competitive adapters matched the performance of the Broadcom adapter for FCoE random reads. In the graph, the two data lines overlap each other for those data points.

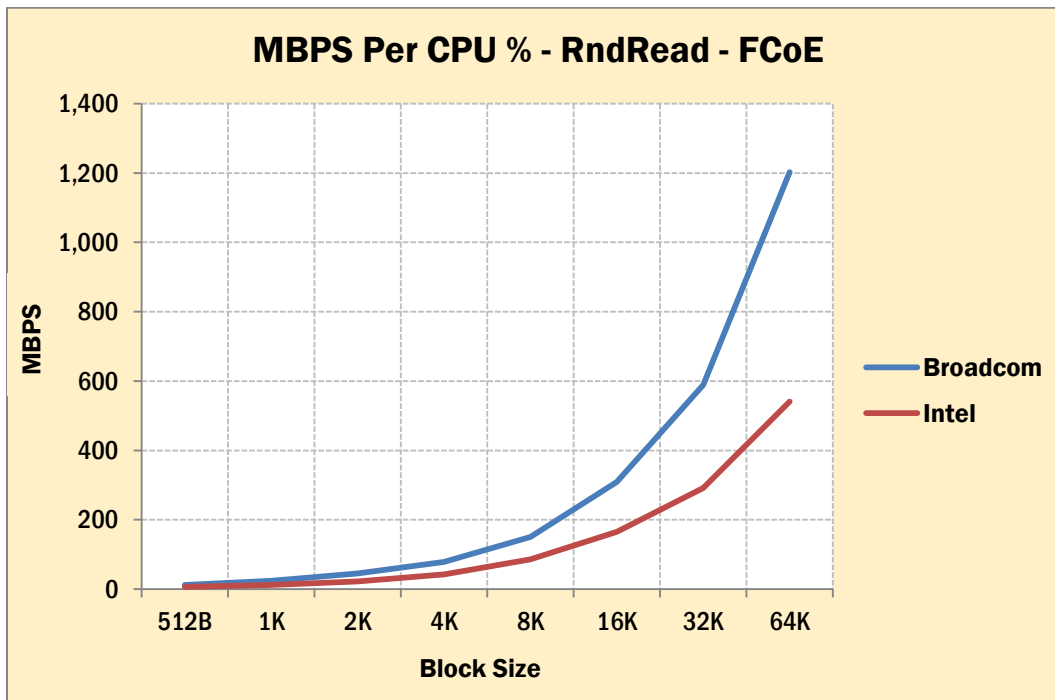
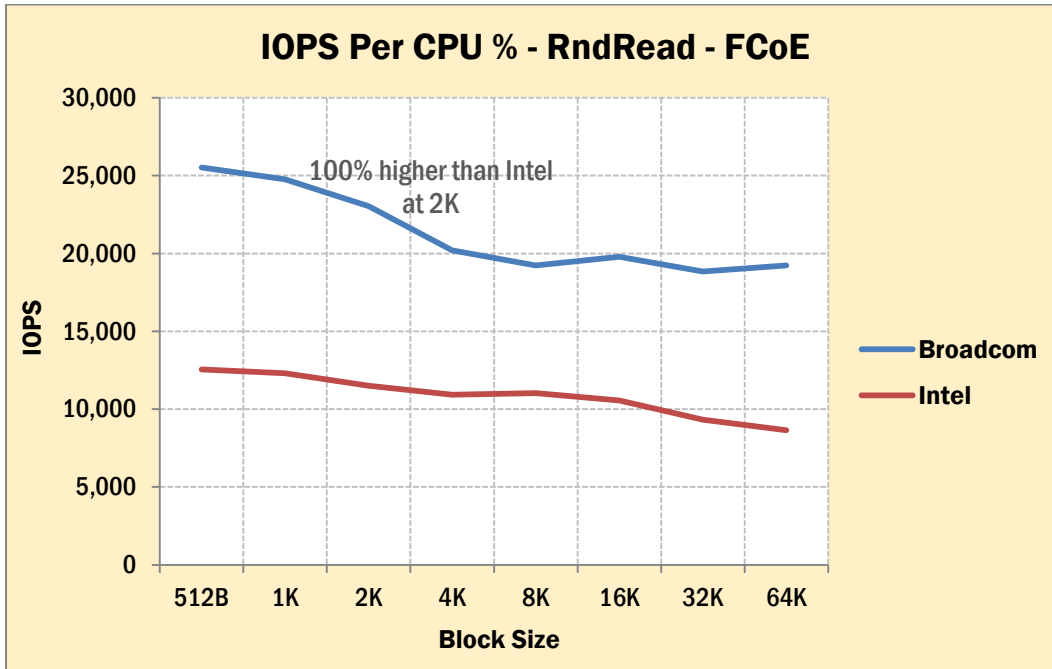


2 – FCoE Performance: Non-Hardware Offload Comparisons

These comparison charts show the Broadcom adapter performance compared to non-hardware offload (Intel) adapter. At the smaller block sizes, the Broadcom adapter significantly outperformed the Intel adapter for FCoE random reads. For the larger block sizes, the two adapters performed approximately the same, so there is some overlap in the lines on the graph.

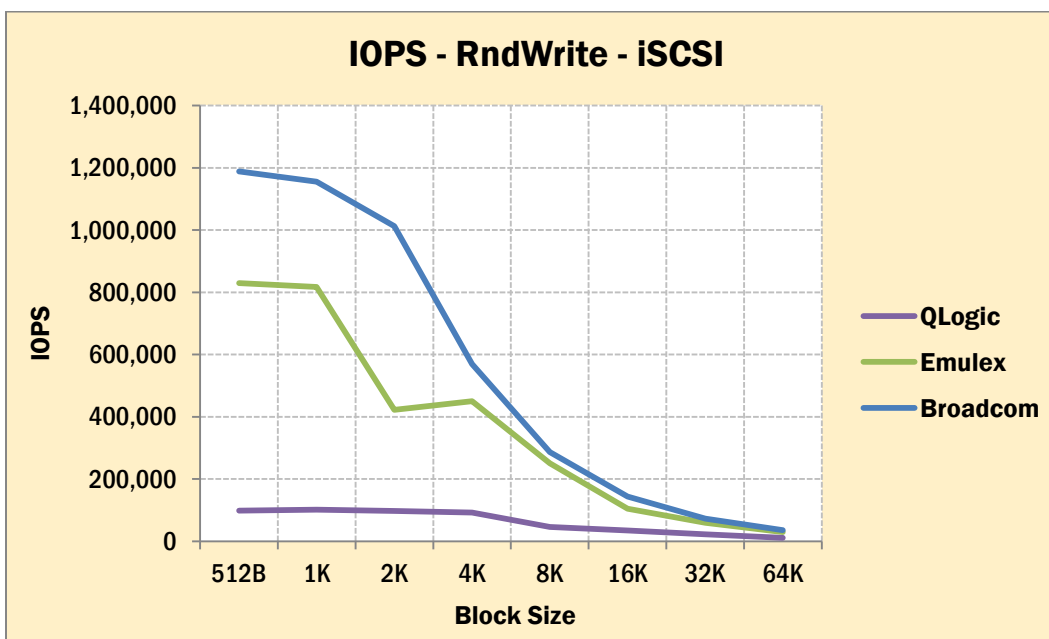
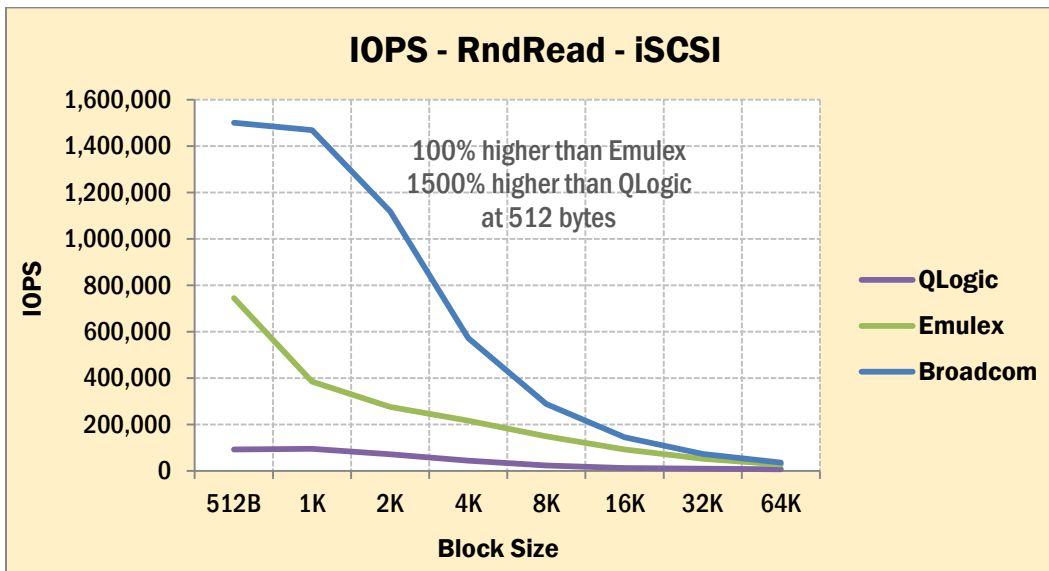


For FCoE random reads, the Broadcom adapter was better in terms of CPU effectiveness for IOPS and MBPS for each percentage of CPU utilization.

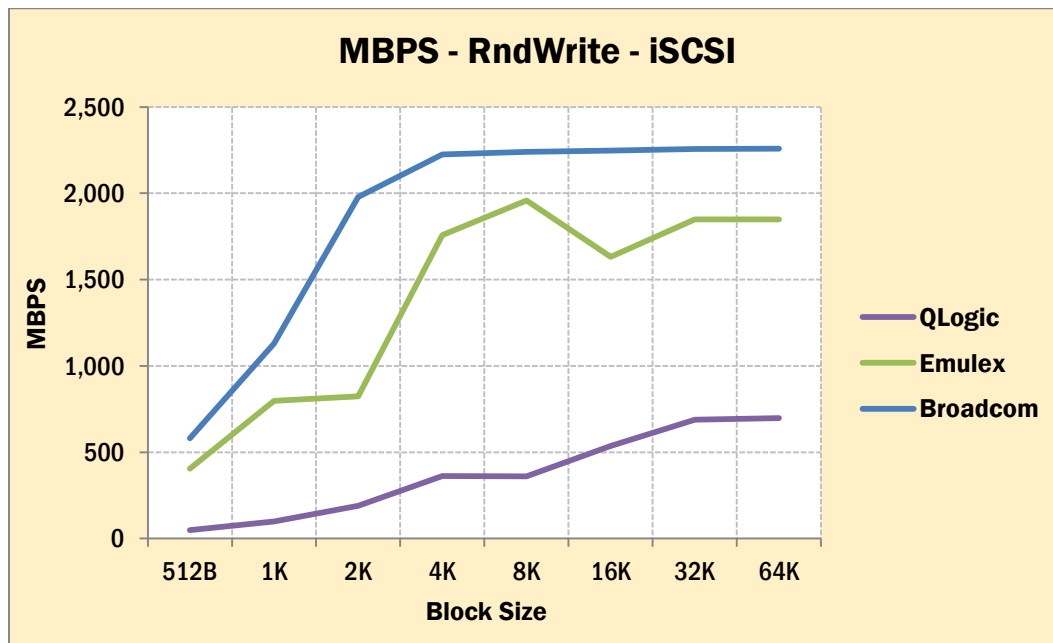
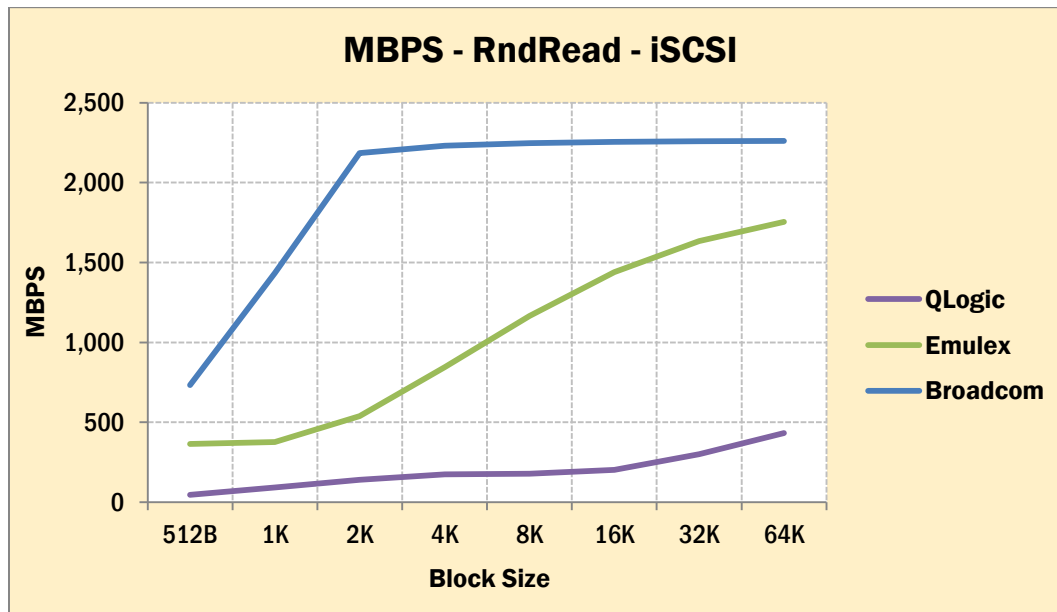


3 – iSCSI Performance: Full Hardware Offload Comparisons

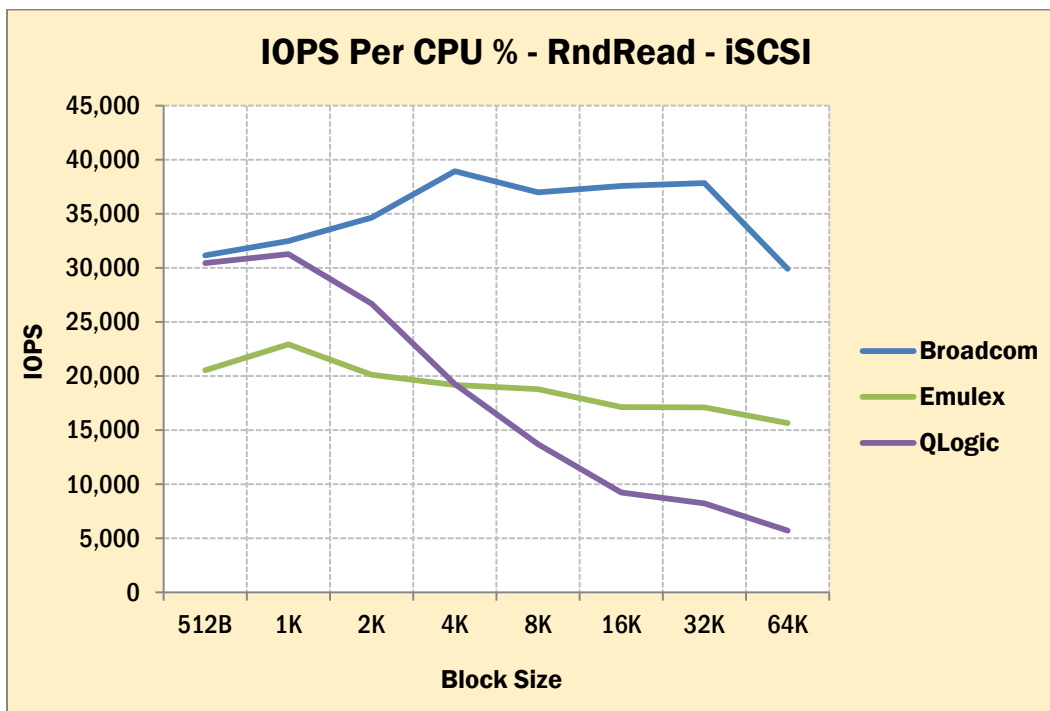
As enterprises consider deploying iSCSI in 10GbE environments, iSCSI adapter performance can be a critical factor in the overall iSCSI solution. Broadcom has demonstrated very strong performance for iSCSI, as shown by the performance results below. These comparison charts show the Broadcom adapter performance compared to iSCSI hardware-offload (Emulex and QLogic) adapters. Again, the Broadcom adapter significantly outperformed the competitive adapters in the smaller block sizes in terms of I/Os per second (IOPS).



The Broadcom adapter achieved line rate at the block sizes shown on the charts for iSCSI offload, while the hardware offload competitive adapters were unable to achieve line rate at these block sizes.

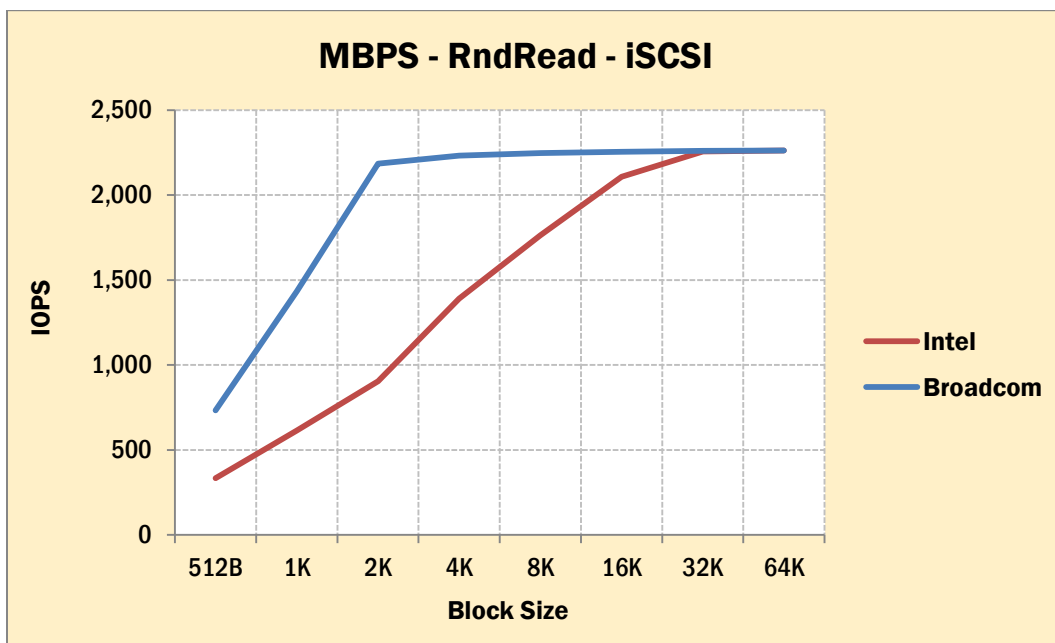
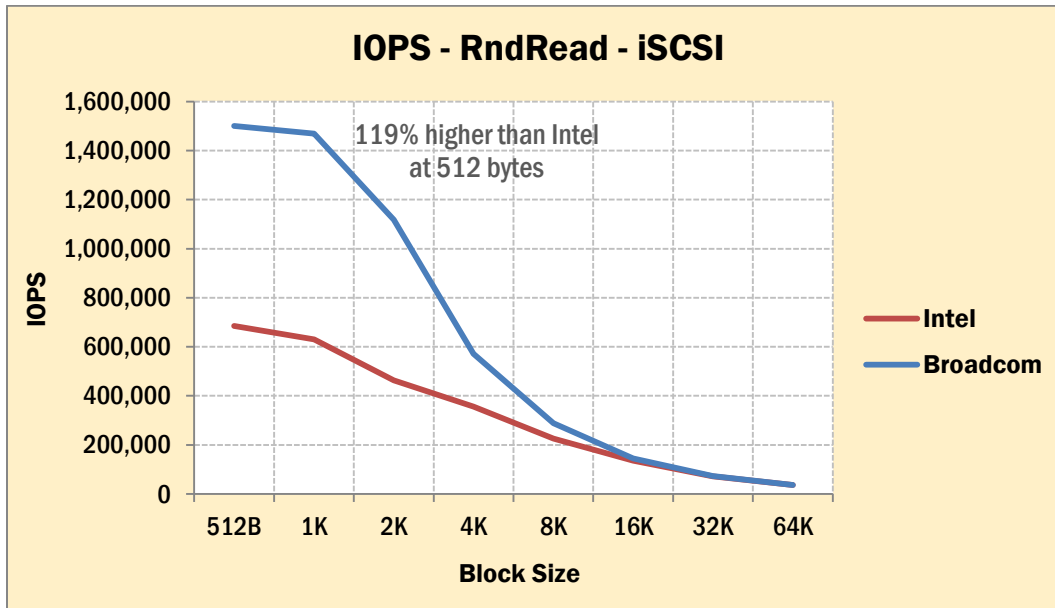


For iSCSI hardware offload random reads, the Broadcom adapter was better in terms of CPU effectiveness for IOPS.

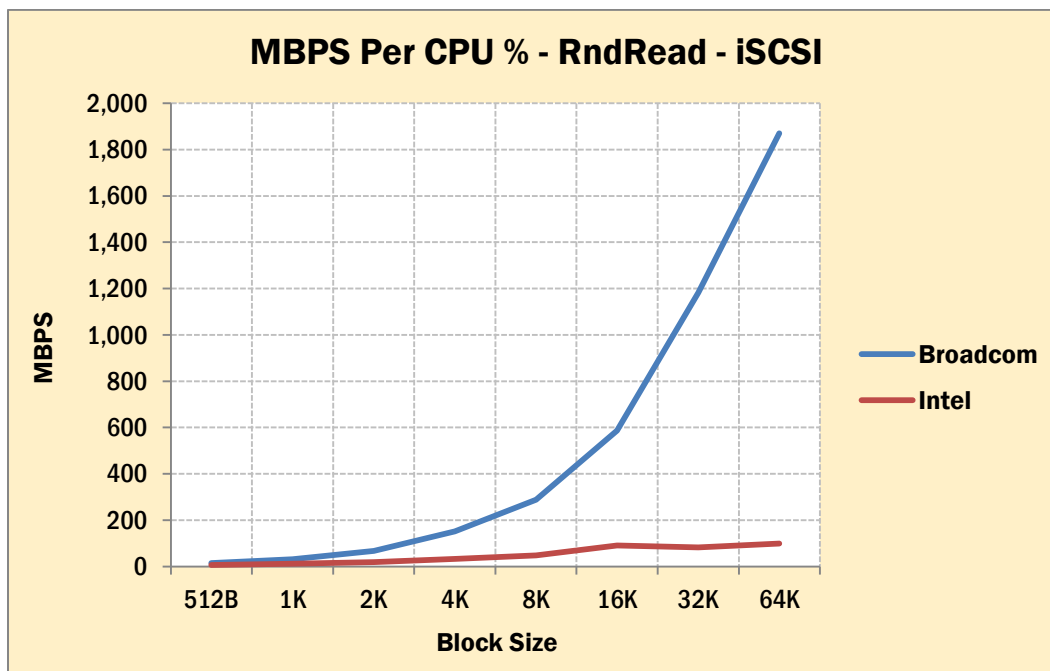
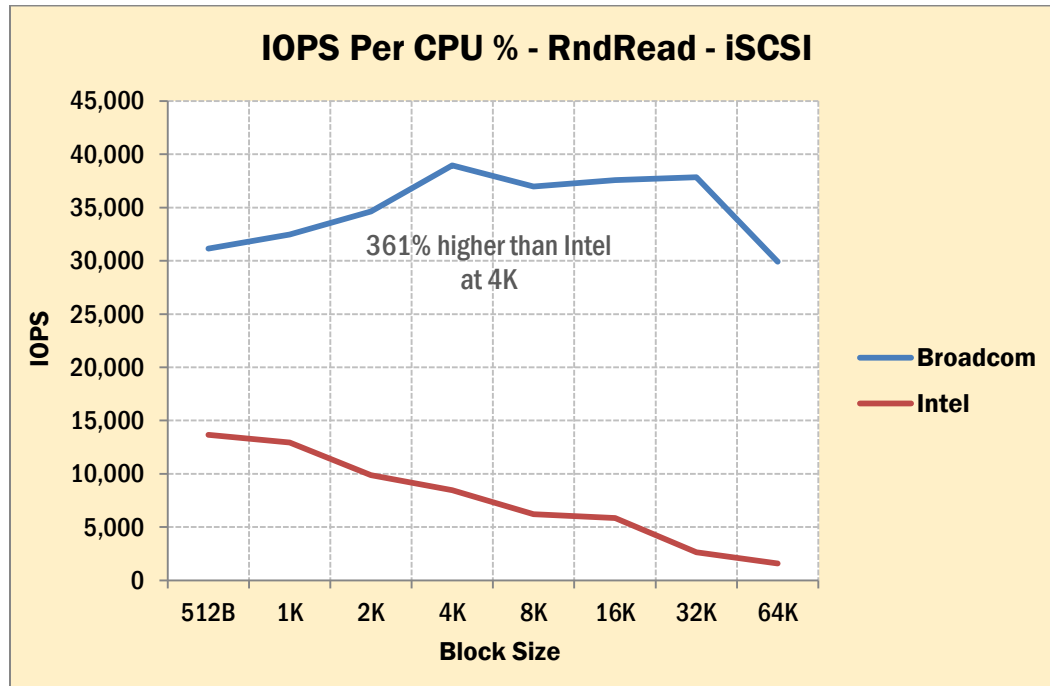


4 – iSCSI Performance: Non-Hardware Offload Comparisons

These comparison charts show the Broadcom adapter performance compared to non-hardware offload iSCSI (Intel) adapters. For iSCSI random read operations, the Broadcom adapter significantly outperformed the competitive adapter, especially for many of the smaller block sizes.



For iSCSI random reads, the Broadcom adapter was better in terms of CPU effectiveness for IOPS and MBPS for each percentage of CPU utilization.



5 – IP Networking Performance: Transmit and Receive Tests - Linux

As public and private cloud providers move to 10GbE infrastructure, the choice of 10GbE adapters in host servers plays an important role. Many cloud environments use Linux hosts due to the low cost of the operating system.

IXIA IxChariot

For these tests, dual-port 10GbE adapters were tested using IXIA IxChariot, a popular network test tool for simulating real-world applications to predict device and system performance under realistic load conditions. Comprised of the IxChariot Console, Performance Endpoints and IxProfile, the IxChariot product family offers thorough network performance assessment and device testing by simulating protocols across the Ethernet network endpoints. IxChariot provides the ability to assess the performance characteristics of any application running on wired and wireless networks.

Three Tests

Three tests were run, measuring throughput performance and CPU utilization. These three tests were:

- Transmit/Receive (TX/RX) Test
- Receive-only (RX) Test
- Transmit-only (TX) Test

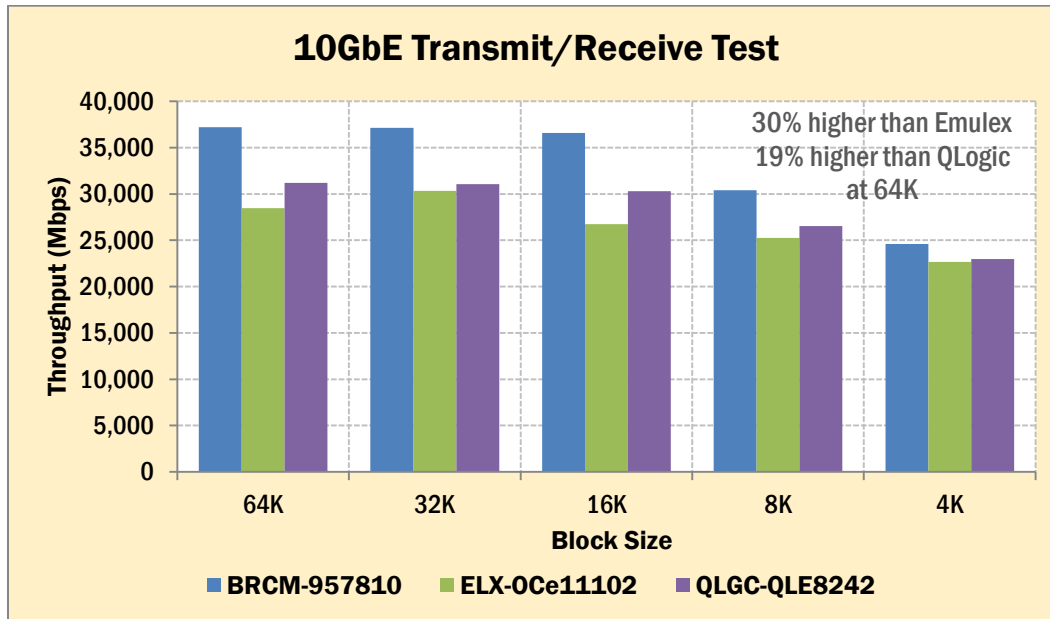
For these tests, the system under test (SUT) was configured as shown in the diagram. Each 10GbE port was connected to a different 10GbE switch. These switches have a combination of 1GbE and 10GbE ports. Network traffic was sent from the server to each of 16 clients. Each client had two 1GbE adapters, one connected to the first 10GbE switch and the other port connected to the second switch. The goal was to make sure that there was enough traffic coming to and from the clients to keep the 10GbE adapter ports on the SUT fully busy.

For these tests, both of the 10Gb ports on the adapters were connected and tested together. For the Transmit/Receive test, the maximum combined theoretical line rate is 40Gb per second because traffic is flowing in both directions on both ports simultaneously. For the receive-only and transmit-only tests, traffic is flowing in one direction only, resulting in a maximum combined line rate of 20Gb per second.

However, traffic must flow through the adapter, its driver, and the full TCP/IP stack on the host server. These tests measure the throughput performance and CPU utilization of the complete solution.

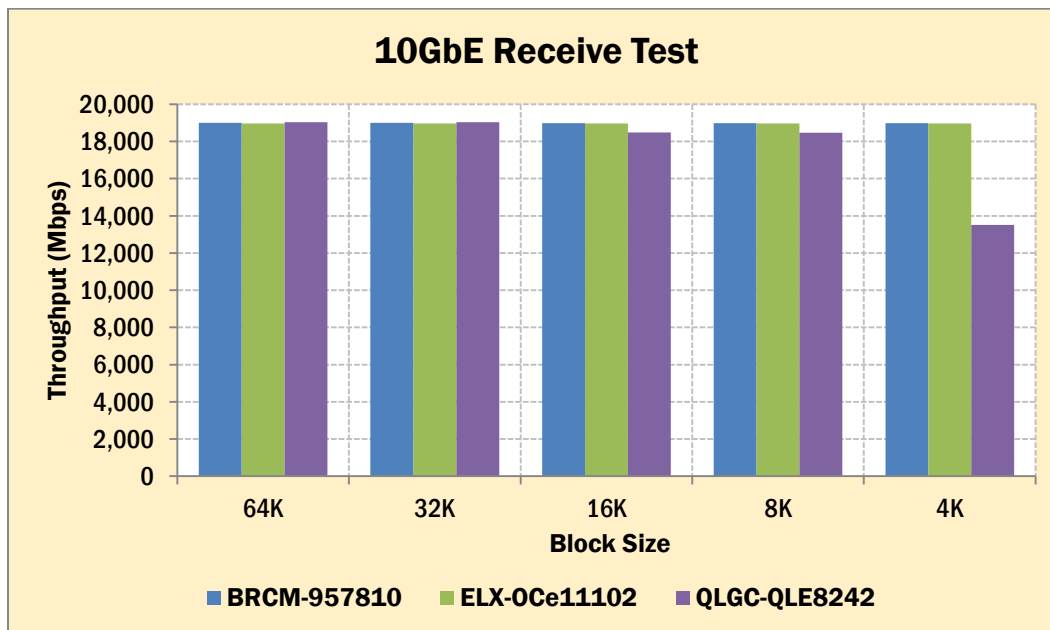
Transmit/Receive (TX/RX) Test

The first set of tests are combined transmit and receive tests. For these tests, traffic is flowing in both directions on both of the 10GbE ports on the adapters and different block sizes were tested. In these combined transmit and receive tests, the Broadcom adapter had higher performance.



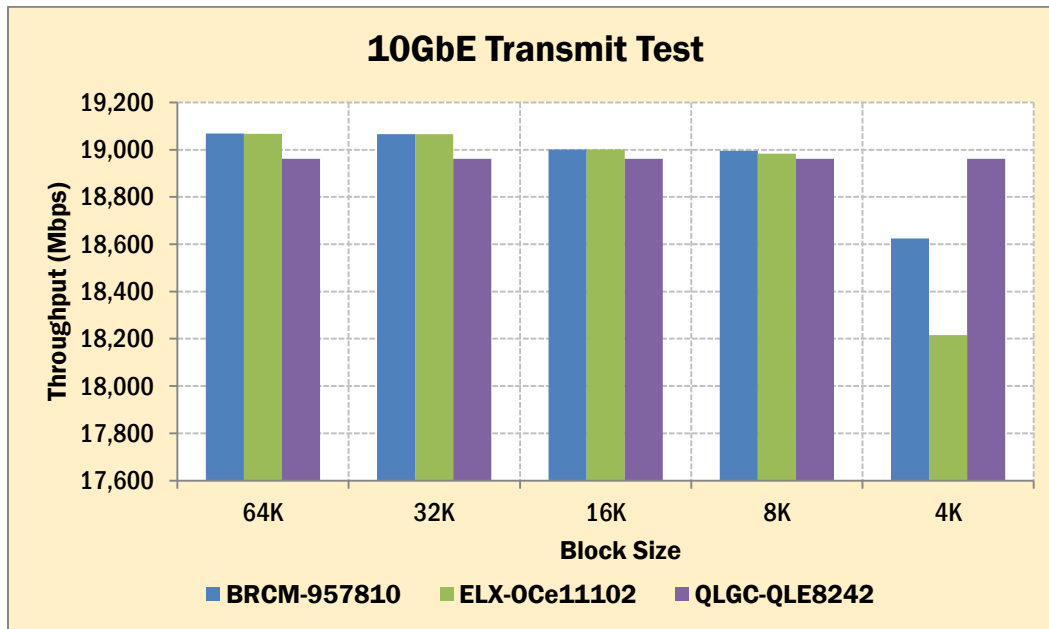
Receive (RX) Test

For the receive-only test, traffic is flowing in one direction only. Because the throughput values were quite close, the scale is adjusted to highlight the differences.



Transmit (TX) Test

For the transmit-only test, traffic is flowing in one direction only. Because the throughput values were very close, the scale is adjusted to highlight the differences. Note the results were almost identical for the larger block sizes, but the Broadcom adapter had a performance advantage at 4K block size.



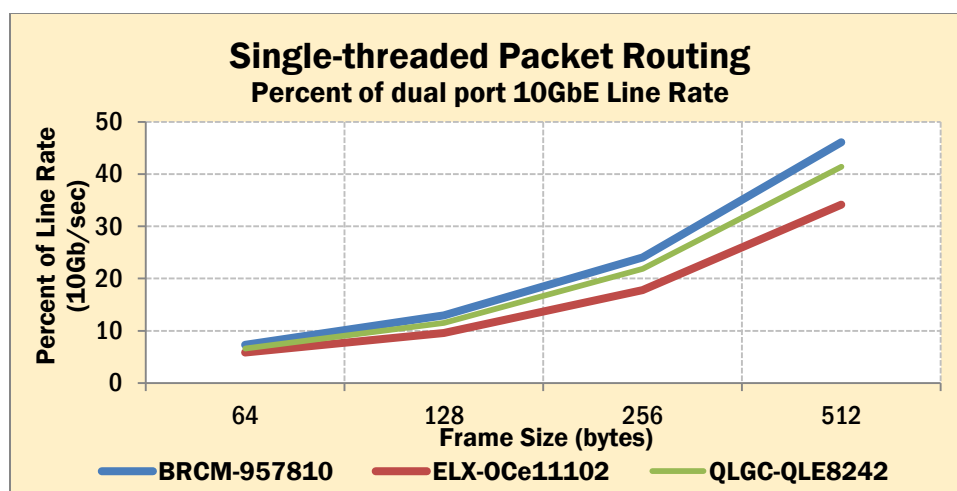

6 – IP Networking Performance: Small Packet Performance

For the small packet performance test, we used the IXIA “IxAutomate” application test tool to configure and automate test scenarios and analyze test results for back to back, one to many, many to one throughput results. Using the capabilities of Ixia test hardware, such as wire-speed traffic generation, filtering, capturing, and statistics collection, we utilized IxAutomate suite of pre-built tests based on industry-standard RFCs requirements.

These tests handle packet routing through the between the test hardware and the adapter installed in the host server. These tests focus on small packets and utilize the adapter, driver and the IP portion of the TCP/IP stack.

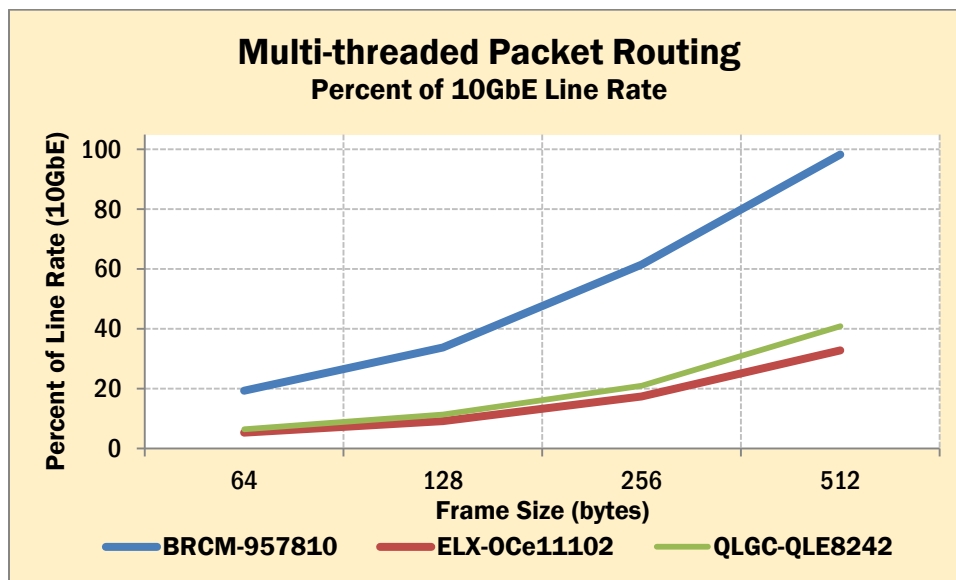
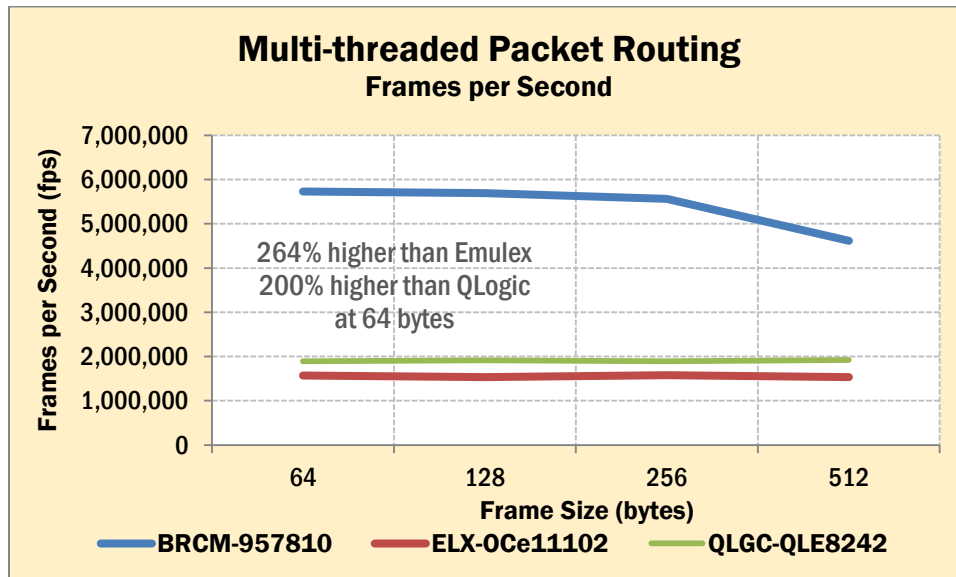
Single-threaded Packet Routing

For these tests, the Broadcom adapter was able to handle more frames per second than the Emulex and QLogic adapters. These results are for both active ports.



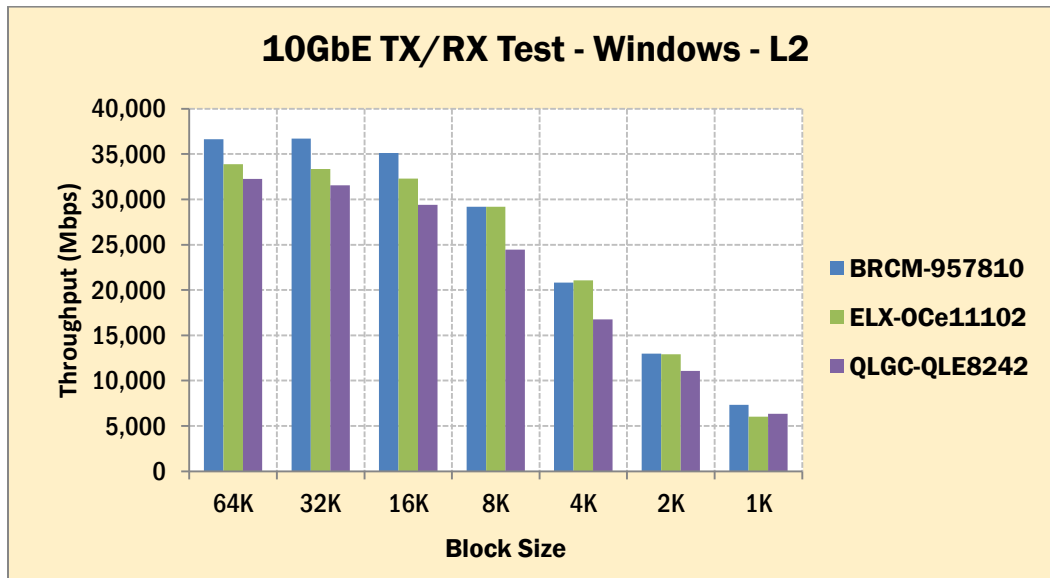
Multi-threaded Packet Routing

For these tests, the Broadcom adapter was able to handle more frames per second than the Emulex and QLogic adapters. These results are for both active ports.

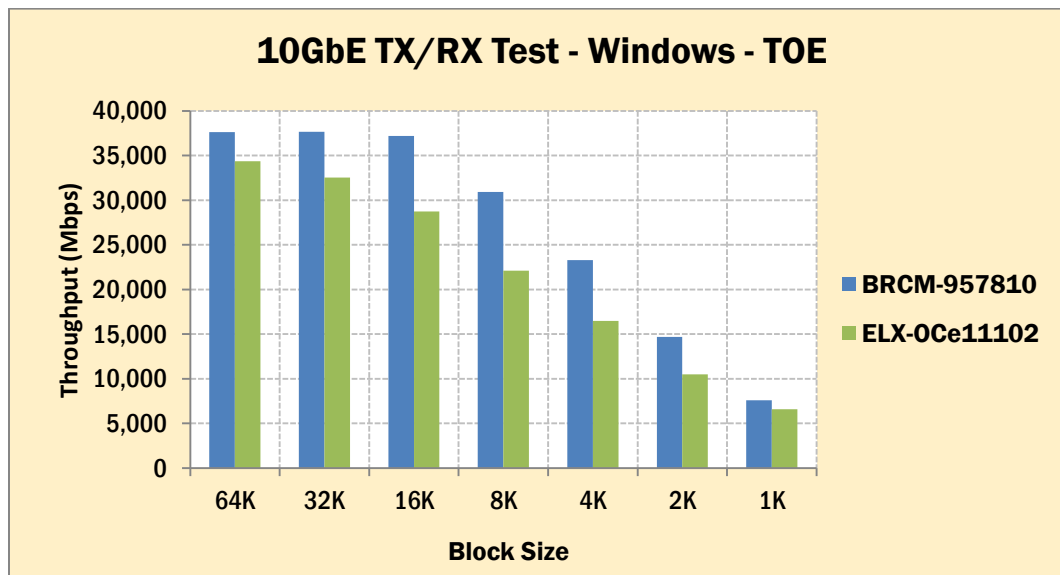


7 – IP Networking Performance: Transmit and Receive Tests - Windows

We also conducted transmit and receive tests in Windows, and found that the Broadcom adapter outperformed the other adapters in almost every case.



TCP/IP Offload Engine (TOE) is a feature available in Microsoft Windows operating systems that offloads the TCP/IP stack. The QLogic card does not support TOE mode, so it is not represented on the chart below.



Summary and Conclusion

For IT managers considering a Converged Network Adapter (CNA) with storage over Ethernet capabilities as well as high-performance IP networking, we recommend evaluating the Broadcom BCM957810S adapter.

In these tests, we found that the Broadcom adapter provided very strong performance across the board for IP networking and storage protocol offloads (FCoE and iSCSI), and would be an excellent choice for 10GbE deployment in enterprise data center and cloud environments.



Appendix – Evaluation Environment

Host Server

- Dell PowerEdge R720
- Intel Xeon E5-2690, 2.90 GHz, 8 cores, 16 logical processors
- 32GB RAM
- Windows Server 2008 R2 SP1, version 6.1.7601

Adapters

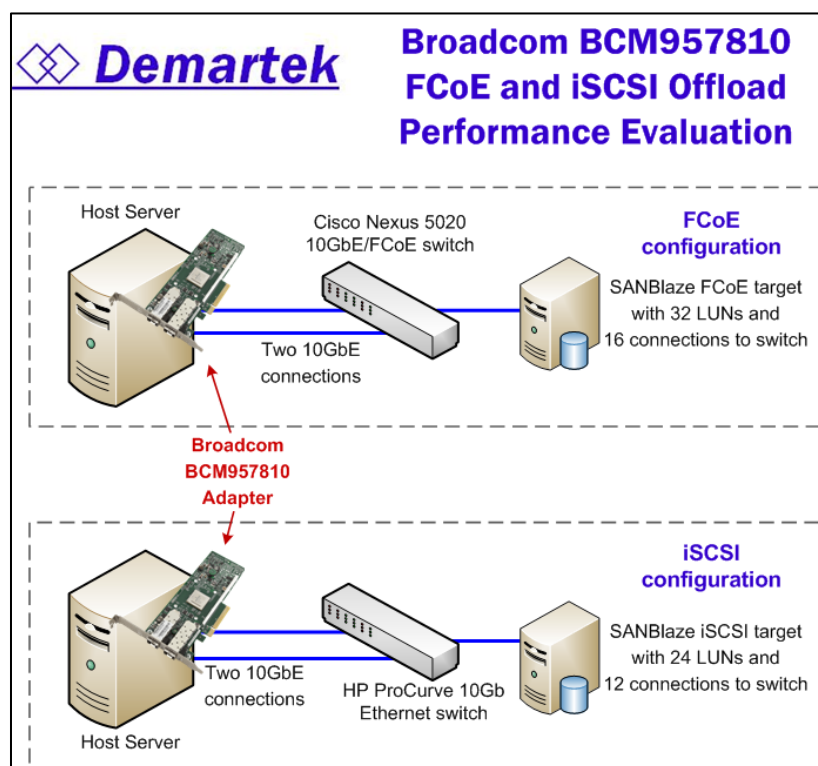
- Broadcom BCM957810S, driver version 7.2.8.0
- Emulex OCE11102, FCoE driver version 7.2.50.7, iSCSI driver version 4.0.317.0
- Intel X520, driver version 2.9.71.0
- QLogic QLE8242, FCoE driver version 9.1.9.39, iSCSI driver version 2.1.5.27

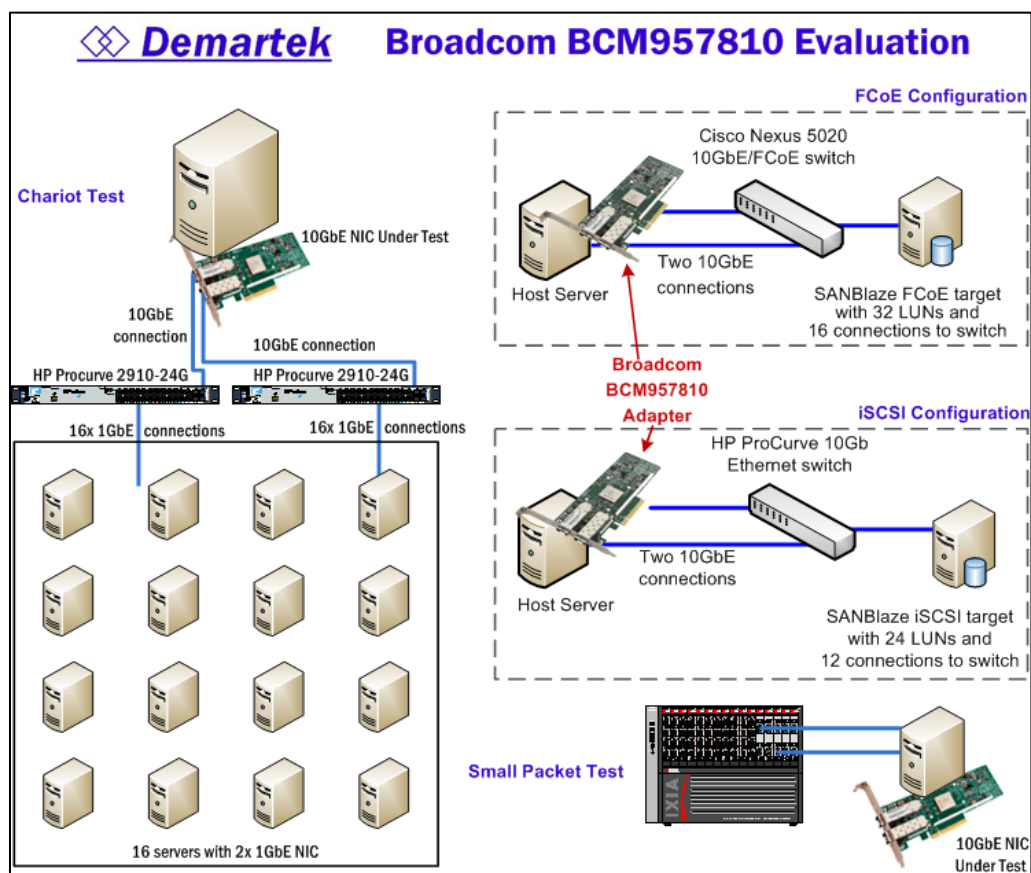
Storage Targets

- FCoE: SANBlaze DRAM targets, 32 total targets connected to Cisco Nexus 5020 switch
- iSCSI: SANBlaze DRAM targets, 24 total targets connected to HP ProCurve 10 Gbps switch

IOmeter Parameters

- FCoE: 16 workers, 2 physical drives per worker, queue depths 1 through 64
- iSCSI: 12 workers, 2 physical drives per worker, queue depths 1 through 64





The original version of this document is available at
http://www.demartek.com/Demartek_Broadcom_BCM957810_FCoE-iSCSI_Adapter_Evaluation_2012-06.html.

Broadcom, Connecting everything logo and NetXtreme are trademarks or registered trademarks of Broadcom Corporation and/or its affiliates in the United States and certain other countries.

Emulex is a registered trademark and OneConnect is a trademark of Emulex Corporation.

Intel and Xeon are registered trademarks of Intel Corporation.

QLogic and the QLogic logo are registered trademarks of QLogic Corporation.

Demartek is a registered trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.