

White Paper

1-Gigabit TCP Offload Engine


Achieving greater data center efficiencies by providing Green conscious and cost-effective reductions in power consumption.

June 2009



Background

Broadcom® is a recognized technology leader in providing networking solutions that have consistently benefited end customers in various enterprise networking and storage environments over the years.



Broadcom is a participating member in the Green Grid industry consortium, an organization committed to developing and promoting energy efficiency for data centers and information service delivery.

Introduction

The Converged Network Interface Controller (C-NIC) concept was introduced in the enterprise market in 2006 when Broadcom demonstrated the convergence of network, storage, and management traffic on a single piece of hardware. BCM5708-based LAN on Motherboard (LOM) solutions were introduced shortly afterwards and showed immediate benefits using one of the C-NIC components: the TCP/IP offload engine (TOE) with reduced CPU utilization and increased throughput in a Microsoft® Windows® server environment.

The associated software suite has remained the same for end users since its inception, and has continued to provide the benefits of TOE at speeds such as 1 Gb and 10 Gb without interruption to deployed systems. The 1 Gb TOE has been further refined with hardware and software optimization on the next-generation gigabit controller BCM5709, such as the IPv6 TOE.

TOE Overview

TOE Architecture

TCP Chimney Offload is a networking technology that helps transfer the workload from the CPU to a network adapter during network data transfers. The TCP Chimney Offload creates a direct connection between the top of the TCP/IP protocol stack and the software drivers to enable offloading of the protocol stack while performing processing on the controller.

To improve data-transfer performance over IP networks, the TOE model can relieve much of the overhead of processing TCP/IP from the host CPU. TOE allows the operating system to offload all TCP/IP traffic to specialized hardware on the network adapter while leaving TCP/IP control decisions to the host server.

By relieving the host processor bottleneck, TOE can help deliver the performance benefits that administrators expect from applications running across high-speed network links. TOE is also cost-effective, as it processes the TCP/IP protocol stack on a high speed network device that requires less processing power than a general-purpose high performance CPU (see Figure 1).

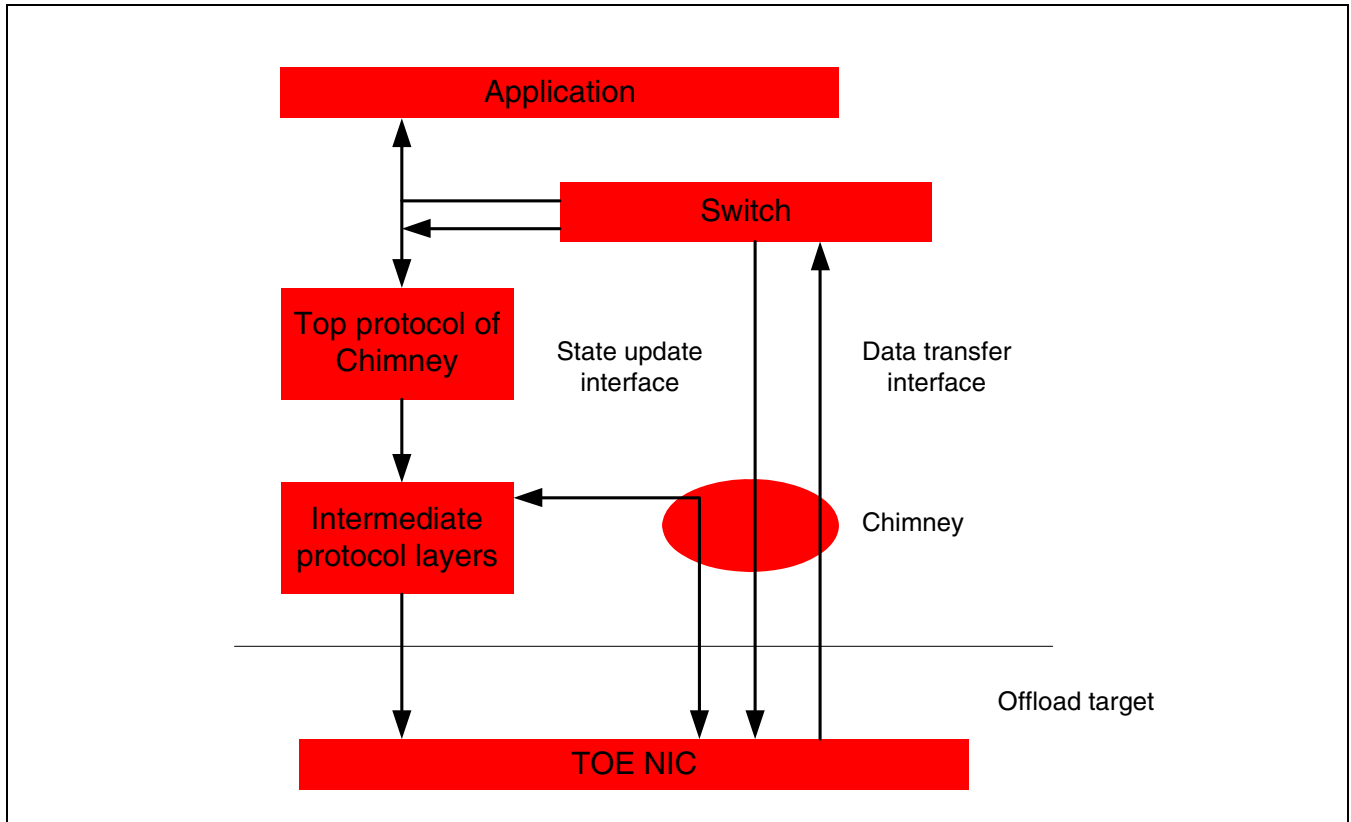


Figure 1: TOE Architecture

Broadcom TOE technology has been covered in detail in various white papers and OEM publications^{1 2}.

1. <http://www.dell.com/downloads/global/power/ps3q06-20060132-Broadcom.pdf>
2. <http://download.microsoft.com/download/3/5/0/3508460c-7bc2-4f7f-b5b9-4aefb981689b/C-NIC-WP102-R.pdf>

System Benefits of TOE

This white paper discusses the system benefits of lower power consumption delivered by using TOE for green data center initiatives. These efficiencies further translate into substantial data center savings for enterprise models such as the consolidation of workloads in virtualization environments.

In recognition of the escalating need to add more LOM controllers due to increased input/output (I/O) activity on the servers, we analyzed the relative power consumption benefits of TOE and compared these to Layer 2 teaming solutions. This analysis produced encouraging results, which are explained in the paper.

To achieve these objectives, Broadcom performance labs used the test setup shown in [Figure 2](#). This included the NetIQ Chariot[®] benchmark. Chariot evaluates the performance of networked applications, performs stress tests on network devices, and predicts networked application performance before deployment.

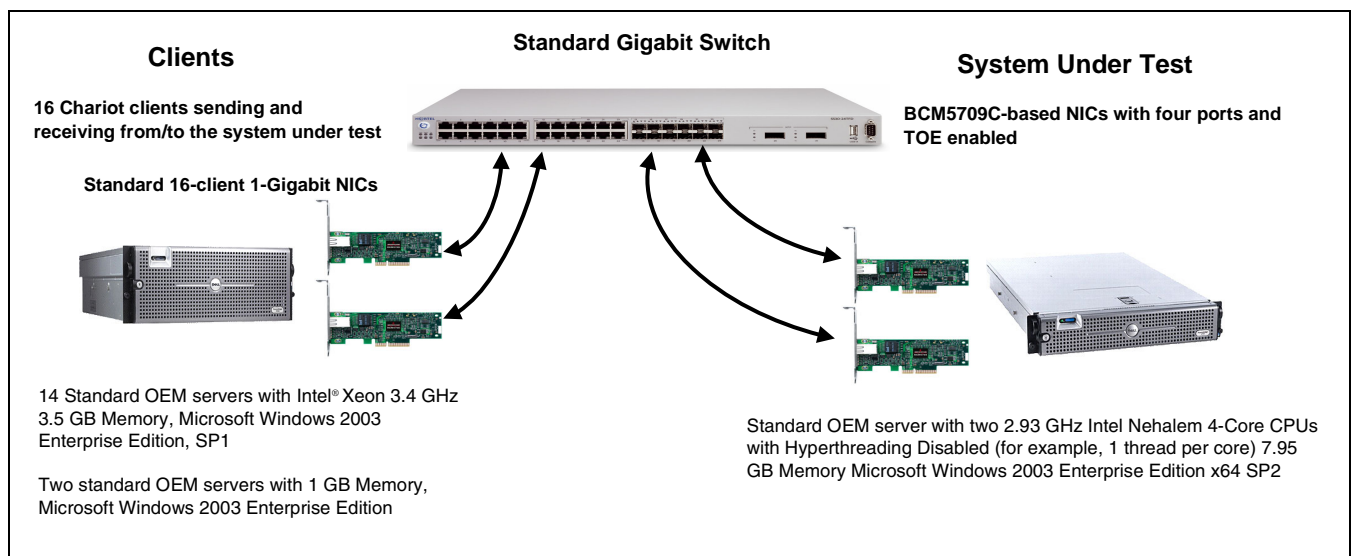


Figure 2: Test Setup

The objective of this test setup was to measure power consumption, CPU utilization, and throughput of systems enabled with TOE. An identical configuration was used with TOE disabled in the Layer 2-only mode to provide a complete comparative analysis.

This setup included two BCM5709 gigabit devices based on dual port NICs in a system under test, with TOE enabled. 16 Windows clients were set up on 16 systems, using single-port gigabit controllers. For current and future data center environments, a majority of server configurations have aggressively moved to 4-gigabit per server port configurations. Therefore, the SUT is based on four GE ports. Chariot clients were set up on each system to send and receive the network traffic to/from the system under test. The setup was created to analyze the maximum benefits that the end user can expect from 100 percent network utilization.

Test setup details, including Chariot script details, can be shared with interested parties so that they can create their own test environment and duplicate Broadcom performance lab results.

Power Benefits

Reduced power consumption is measured as the TOE performs TCP/IP network traffic processing on the BCM5709, rather than passing this task to the host CPU, a process that consumes more power for the same duty. Systems with TOE-enabled networking experience an overall lower power consumption.

Figure 3 shows a throughput and power consumption comparison when TOE is enabled and disabled on the NICs in the system under test, in both teaming and non-teaming modes. Generally, the TOE for 1-gigabit solutions saves overall system power by about 3–8W, while lowering CPU utilization and maintaining the throughput for larger I/O sizes.

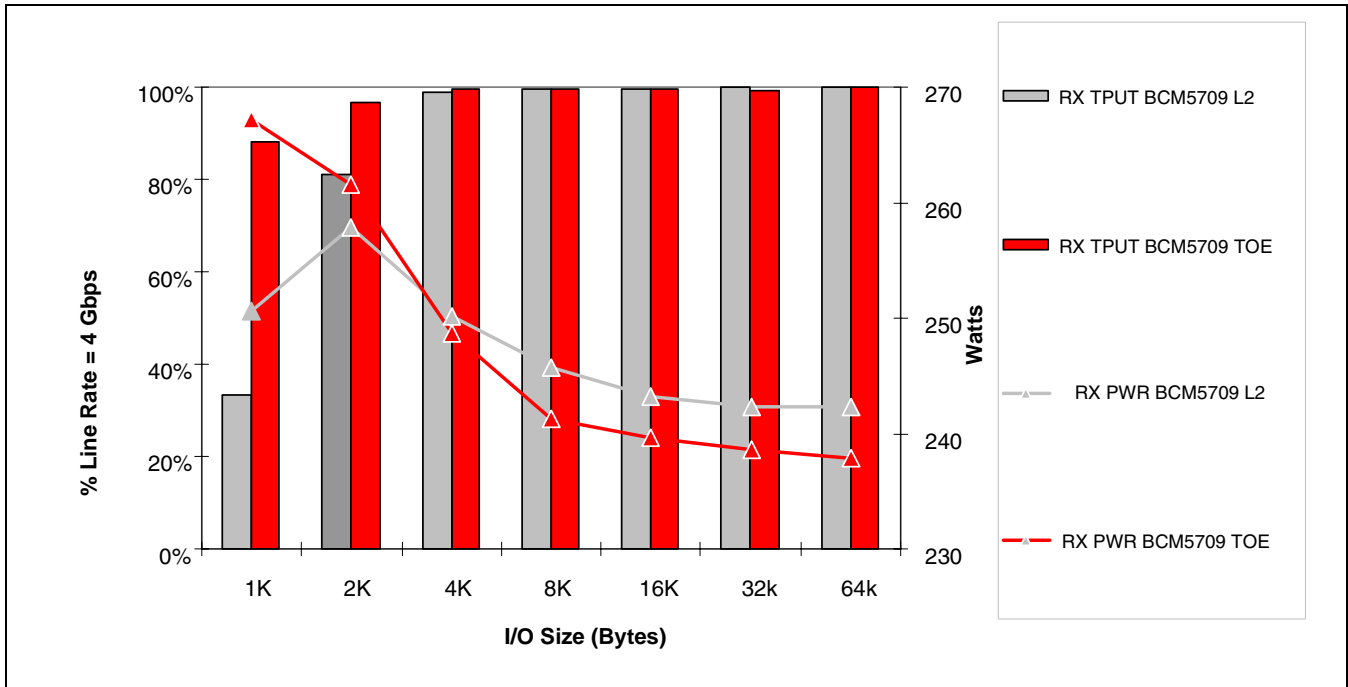


Figure 3: Chariot Four-Port GbE Throughput vs. Platform Power in Non-Teaming Mode

Measuring 100 percent network utilization for each server over a one year period and using the current power usage effectiveness (PUE)^{1 2} value of 1.7 results in a savings of \$10.00 per system. We calculated the savings by multiplying an 8W single-system power reduction by 8.8 cents/kWh, the average electricity commercial use rate reported in the footnoted² energy star report, section 2.2.4.

1. http://svlg.net/campaigns/datacenter/docs/DCEFR_executive_summary.pdf (Page 8)
 2. http://www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf

The savings may, in fact, have increased over the last three years if the commercial rate had trended upwards. When we extrapolate these findings over to a server farm of various capacities as shown in [Figure 4](#), the resulting savings can grow exponentially for a cost-based and "green" data center-focused IT yearly budget. For example, a 2500 server deployment could generate an annual cost savings of \$26,000.

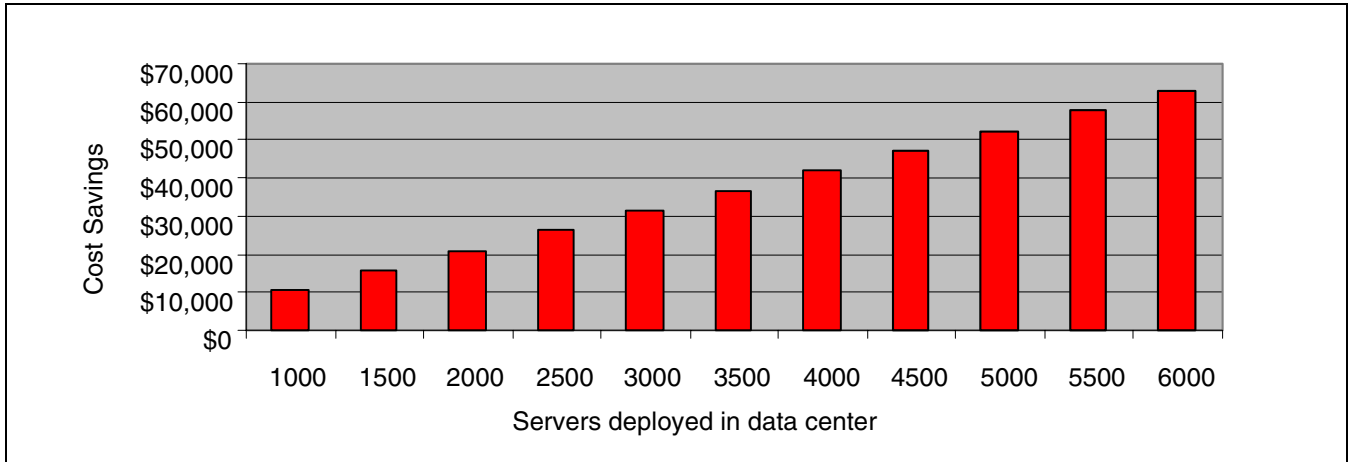


Figure 4: Cost Savings with Lower System Power Consumption

CPU Utilization Benefits

Figure 5 shows the throughput and CPU utilization benefits when TOE is enabled and disabled on the NICs in the system under test, in the non-teaming mode.

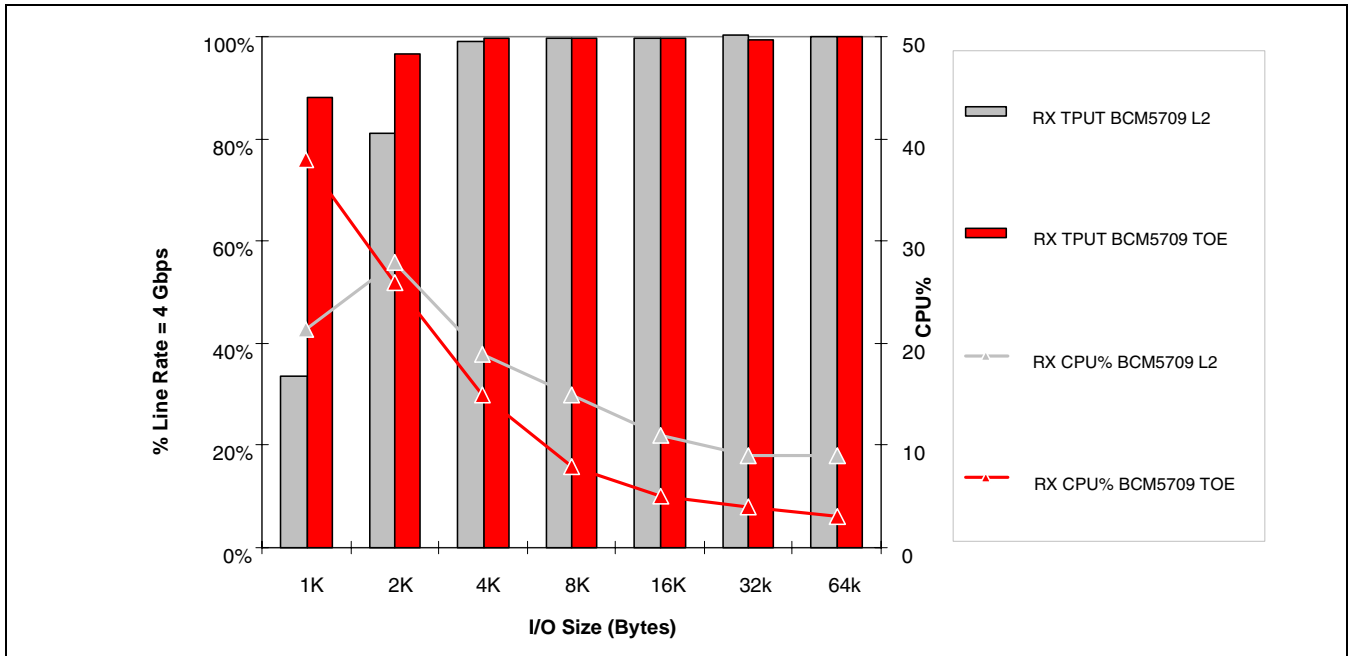


Figure 5: Chariot Four-Port GbE Throughput vs. CPU Utilization in Non-Teaming Mode

The ratio of throughput to CPU utilization is expressed using the performance efficiency (PE) index, which was originally developed by PC Week (now eWeek) in 1995. The PE index is still the most commonly used performance ratio for evaluating adapters. High PE indexes, indicating high throughput with low CPU utilization, suggest favorable overall system performance. Lower CPU utilization and increased PE ratios were illustrated using the same setup shown in [Figure 2 on page 4](#). CPU utilization for larger I/Os is lower by an average of 9 percent in this configuration.

Availability of such CPU computing power in some scenarios can result in substantial infrastructure cost savings for a green data center as shown in [Figure 6](#).

The savings shown in [Figure 6 on page 8](#) were calculated using the values shown in [Table 1](#). This table quantifies servers, switches, and data center resource requirements, and shows ratios and associated costs. These factors represent general estimates which vary according to locations, solution configurations, and requirements. Broadcom performance labs modified the various parameters to assess the impact on the end results. The customer cost savings have remained high and justify a stronger focus on using TOE in existing and future Windows server deployments.

Table 1: CPU Utilization Savings

Description	Metric	Application
Percentage of CPU utilization savings	9.00	For large I/O sizes (2K - 64K)
Cores/server	8	—
Average server cost	\$5,000	—
Server lifetime	3	—
Switch: server ratio	2.5	—
Switch cost	\$1,500	—
Servers per rack	12	2U servers
Rack floor space (sq. ft.)	10	—
Commercial sq. ft. cost	\$1,000.00	—
Servers/IT Administrator	250	—
IT Administrator annual salary	\$60,000	—

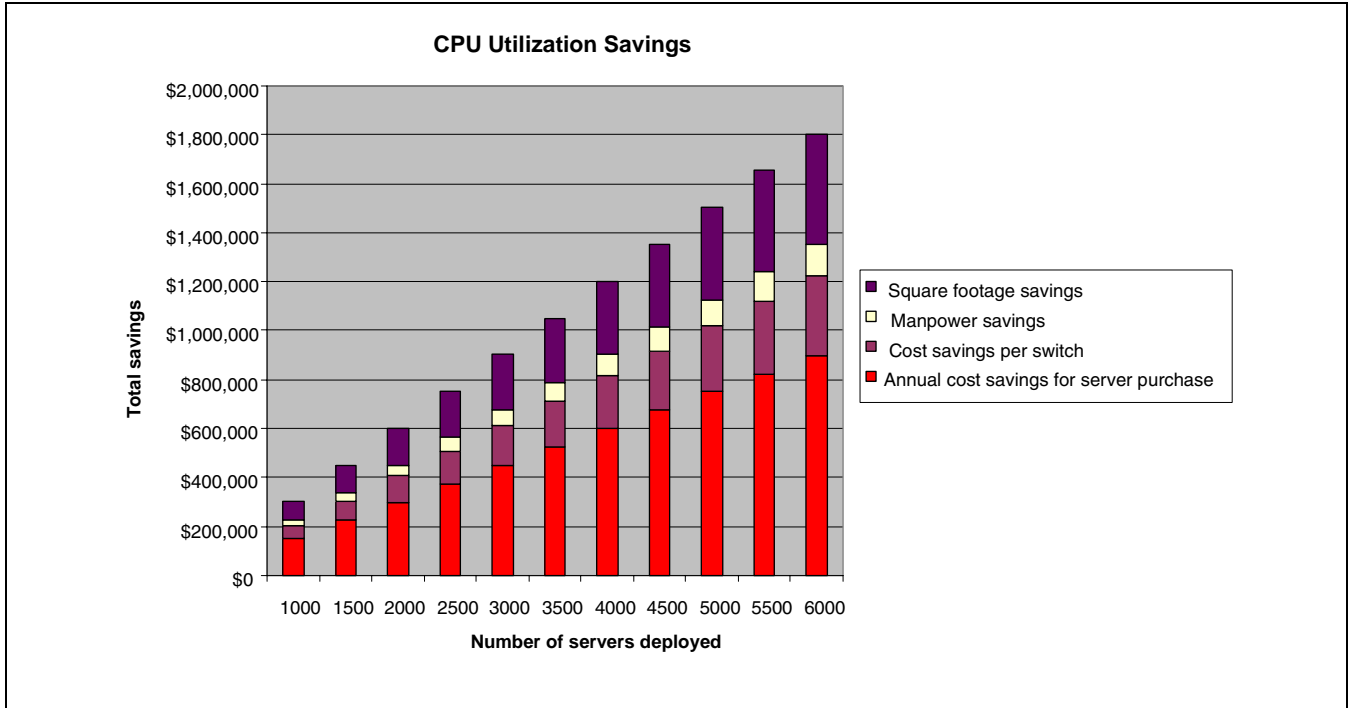


Figure 6: Total Savings for an Average 9 Percent CPU Utilization Reduction

Environments where workloads can be divided into distinct modules, and then distributed over several systems, can benefit substantially. Workload consolidations from one system to another in a virtualized environment will also benefit, as the available CPU resources can be assigned to other processes and activities with greater efficiency.

TOE Teaming Power Consumption and CPU Utilization Benefits

This test setup was used to create a single logical 4-gigabit networking pipe using a Broadcom Advanced Control Suite BACS in L2 and TOE modes. The objective of these tests was to analyze any differences comparing TOE teaming against a standard non-TOE teaming configuration.

TOE-teamed adapters have greater performance efficiency compared to L2-teamed adapters. TOE teaming performance is virtually identical to TOE performance when multiple adapters are run independently. Teaming "overhead" with TOE enabled is negligible.

This serves to keep the realized system power benefits and CPU utilization intact.

Summary

The TOE power consumption and CPU utilization analysis presented herein demonstrates the benefits of TOE deployment into customer data centers. This implementation advances an effective green data center strategy in a corporate climate where businesses are doing more with less, and provides measurable initiatives to contain overall IT budgets.

In recent years, Broadcom has provided these C-NIC solutions using 1 Gb speed TOE. This approach delivers tangible advantages to end customers who need results when managing dense networking traffic environments. As 10 GbE becomes more ubiquitous in the data center, we anticipate that the TOE success story will further extend beyond 1-gigabit links to the more expansive and demanding 10-gigabit links.

Engineering Profiles

Rich Hernandez (rich_hernandez@dell.com) is a senior development engineer with the Server Product Group at Dell™. Rich has been in the computer and data networking industry for over 25 years, and has a Bachelor of Science degree in Electrical Engineering (BSEE) from the University of Houston. He has pursued postgraduate studies at Colorado Technical University.

Dhiraj Sehgal (dhiraj@broadcom.com) is a Senior Product Line Manager in the Enterprise Networking Group at Broadcom Corporation. He has extensive datacenter, enterprise server, and storage experience focusing on various I/O subsystems. Dhiraj has a Master of Science degree in Electrical Engineering (MSEE) from North Carolina State University, Raleigh.

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
© 2009 by BROADCOM CORPORATION. All rights reserved.

Phone: 949-926-5000
Fax: 949-926-5203
E-mail: info@broadcom.com
Web: www.broadcom.com