

# TCP Segmentation Offload (TSO) Performance On HP rp4440 Servers Using PCI-X 2-Port Gigabit Ethernet Cards



August 2004

## Table of Contents

Introduction .....	3
rp4440 Performance with GigE and TSO .....	5
Service Demand: Less = More .....	6
Throughput: Faster = Better .....	7
Performance Highlights .....	8
Suggested Use .....	8
Achieving High Availability with 2-Port Cards .....	8
Technology Combinations to Use or to Be Aware Of .....	9
Test Details .....	10
Products Used in Testing .....	10
Test Configuration .....	11
For More Information .....	11
Taking the Next Step .....	11
Legal Notices .....	12



## Introduction

This article highlights performance improvements available when you use HP's TCP Segmentation Offload (TSO) capabilities in conjunction with the new A7011A and A7012A 2-port Gigabit Ethernet cards (Figure 1) in a powerful server such as the HP rp4440.

TSO works especially well in networks sending large amounts of data, such as in: web serving applications, data backups, or file transfer applications including NFS. TSO reduces CPU load thereby improving overall system response. Greater than 50% reduction in CPU utilization has been observed on some FTP workloads.<sup>1</sup>

**Figure 1** HP PCI-X 2-Port Gigabit Ethernet Card (A7012A shown)



1. Not all applications benefit from TSO. Only data-intensive applications that transmit large data buffers using TCP over IPv4 are improved. Other types of applications will not significantly benefit from TSO. Performance improvements vary depending upon the hardware and software used. For details on recommended configurations, please refer to the section in this paper called “Technology Combinations to Use or to Be Aware Of.”

## Introduction

TCP Segmentation Offload (TSO), also known as “large send” enables a system’s protocol stack to offload portions of outbound TCP processing to a network interface card thereby reducing system CPU utilization and enhancing performance. Instead of processing many small maximum transmission unit (MTU)-sized frames during transmit, the system can send fewer larger virtual maximum transmission unit (VMTU)-sized frames. On a typical Ethernet card, the link MTU is 1500 bytes while a VMTU can be much bigger. Currently, the VMTU size is fixed at 32160 bytes.

HP’s TSO is supported on Gigabit Ethernet cards including embedded or “core” LAN cards that are served by the iether or igelan software drivers running on HP-UX 11i v 1.0 and 2.0 or the September 2004 OE release of HP-UX 11i v 2. See “Taking the Next Step” in this document, for how to obtain a free copy. *TSO requires no software configuration*: once all the required components are installed, it’s enabled automatically. TSO is supported on virtual LANs (VLANs). It’s supported at all speed settings on the link and all MTU values including Jumbo frames. Both Ethernet and SNAP encapsulations are supported. TSO does not work over link aggregates (HP Auto Port Aggregation), and if APA is present, TSO is automatically disabled.

The Gigabit Ethernet cards tested with TSO for this paper are the PCI-X 2-port A7011A card (with fiber-based connectors) and the A7012A card (with copper-based connectors). The 2-port cards provide the following benefits:

- Higher port density for I/O slot–constrained systems.
- Higher levels of failover protection, because two cards have no single point of failure. The same level of failover protection would require four single-port I/O cards that would consume twice the number of I/O slots.

The cards also fully support HP Serviceguard and PCI-X/PCI online addition and replacement (OLAR).

- PCI-X I/O cards that are fully backward compatible with PCI 2.2 I/O slots.
- Flexibility for implementing Virtual Partition (vPars) configurations.
- Increased server CPU efficiency and performance through TCP, UDP, and IP checksum protocol off loading.
- Increased network flexibility through support for virtual LANs (VLANs).
- A choice of either a 1000 Base-SX version using multi-mode fiber supporting distances of up to 550m or a 1000Base-T version using CAT5 or better UTP cable supporting distances of up to 100m.

HP rp4440 servers like the one used in this demonstration are highly dependable, adaptable, and efficient servers for your enterprise. The results in this paper show that the A7011A and A7012A cards installed in a server like the rp4440 and combined with features like TCP Segmentation Offload (TSO) can increase your server efficiency while lowering your total cost of ownership.

## rp4440 Performance with GigE and TSO

Efficiency of the HP rp4440 server with TSO over PCI-X 2-port Gigabit Ethernet A7011A and A7012A cards is best demonstrated through tests that show:

- **Service Demand.** Service demand is the amount of time (in microseconds) it takes one CPU to handle one kilobyte of data. It's a straightforward measurement because it eliminates disparities if comparisons are made with different quantities, types, or frequencies of CPUs. Service Demand is an important capacity planning & performance metric that is sometimes considered a better metric than CPU Utilization when comparing different server models.
- **Throughput.** Throughput is the measure of the available bandwidth; in this article, it is shown for one-way transmission alone as well as the combination of transmit and receive throughput, also called bidirectional. It's an important metric of networking performance and server capacity sizing for optimal application performance. These measurements can help determine how well one or a series of batch programs will run with a certain workload or how many user requests can be handled.

The performance gains in service demand and throughput shown in this article were measured with the `netperf` benchmarking software. The following charts show the overall performance of one, two, and four ports on up to 2 PCI-X 2-Port Gigabit Ethernet cards in a 4-way rp4440. TSO was measured with the following traffic types and options:

- Service Demand and Throughput were measured with:  
Transmit (xmit) or bidirectional (bd) with TSO off  
compared to  
transmit or bidirectional with TSO on.
- Transmit and bidirectional tests were run using `netperf` with a socket size of 128K bytes and a message size of 32K bytes.

The results of TSO testing on single port Gigabit Ethernet cards A6825A (Base-T)/A6847A (Base-SX) are also described later in this article though not shown in graphs. For further discussion of the results obtainable with GigE cards other than the 2-port GigE cards, please refer to the section in this paper called "Technology Combinations to Use or to Be Aware Of."

## Service Demand: Less = More

**Figure 2** rp4440 Service Demand with GigE and TSO




Figure 2 shows **Service Demand** results on the rp4440 over GigE both without and with TSO. Service Demand is the amount of time it takes one CPU to handle one kilobyte of data. In this test, the *lower* the number, the better!

- For **transmit only** traffic, TSO *reduces* Service Demand an average of 47%. The actual Service Demand savings are:  
for 1 port with TSO on, 47% *less* Service Demand,  
for 2-ports with TSO on, 47% *less* Service Demand, and  
for 4 ports, 46% *less* Service Demand than when TSO is off!
- For **bidirectional traffic**, TSO *reduces* Service Demand an average of 27%. The actual Service Demand savings are:  
for 1 port with TSO on, 24% *less* Service Demand,  
for 2-ports with TSO on, 30% *less* Service Demand, and  
for 4 ports, 26% *less* Service Demand than when TSO is off!

TSO reduces the amount of load or service demand on the CPU thereby allowing it to work more efficiently. This makes more CPU available for either the same or other kinds of work.

## Throughput: Faster = Better


**Figure 3** rp4440 Throughput with GigE and TSO

Figure 3 shows transmit and bidirectional **throughput** results on the rp4440 over GigE both without and with TSO. In the enclosed graph, the Gigabit Ethernet throughput numbers are shown in *Megabits per second*. In this test, the *higher* the throughput number, the better!

- For **bidirectional traffic**, TSO increases throughput much more with multiple ports. The actual throughput *increases* are:  
for 1 port with TSO on, 3.95 Mbit/s *more* throughput ,  
for 2-ports with TSO on, 178.63 Mbit/s *more* throughput , and  
for 4 ports, 577.18 Mbit/s *more* throughput than when TSO is off!
- TSO does not significantly affect **transmit only** throughput because the ports are already at link rate.

TSO increases the throughput of bidirectional traffic on iether-based GigE cards especially over multiple ports. TSO does not improve throughput on systems that are already nearing link rate. But, because TSO does free up CPU, it can provide improved throughput for systems that are CPU-bound or that are not already at or near link rate on all network ports.

## Performance Highlights

*During bidirectional traffic over multiple ports of the A7011A/A7012A GigE cards, TSO reduces service demand by up to 30% while improving bidirectional throughput by up to 577 Mbit/s. When installed in a 4-way HP rp4440 server, the A7011A and A7012A cards provide exceptional networking throughput, connectivity, and reliability. Performance is excellent even with both ports operating concurrently as well.*

The GigE cards and TSO can be configured and used to best meet your business needs as described in the “Suggested Use” section that follows.

## Suggested Use

Using two A7011A or A7012A cards provides maximum failover. Where performance and high availability are the highest priority, HP recommends that you use 2 cards in a hot-failover configuration. See the next section for more.

### Achieving High Availability with 2-Port Cards

To get high availability using the A7011A/A7012A, HP recommends the following:

- When you use two 2-port cards together, you achieve excellent performance and high availability. In the event of a system slot failure or other network failure, an alternate data port can take over – maximizing system up time!

---

**NOTE** The following recommendations assume that your systems are also configured with failover software such as HP Serviceguard and Logical Volume Manager. Consult <http://docs.hp.com> for full details on configuring your products for high availability.

---

- Connect each GigE port through a different network switch and ensure that the switches are bridged to the same IP subnet. Failover between the two network paths is managed by Serviceguard’s local network failover capability.
- Connect the network ports on multiple rp4440s to a private subnet and use them for a dedicated heartbeat for Serviceguard.

This configuration protects you from a failure of a single GigE port on either card, a failure of either GigE switch, or the failure of any single network cable.

- Use PCI-X mode as opposed to PCI mode. The primary PCI-X bus running at 133MHz PCI-X with 64-bit bus width will yield the best results and the performance comparable to that shown in this paper. HP tested the A7011A and A7012A cards in PCI mode, but because PCI performance is lower than PCI-X performance, HP recommends using these cards in PCI-X mode rather than PCI. In the rp4440 system under test, the PCI-X high-performance or “dual rope” slots are slots 7 and 8. See “Test Configuration” for more details on the PCI-X slots of the rp4440. Please see your systems’ documentation for the location of the PCI-X slots in your systems.

## Technology Combinations to Use or to Be Aware Of

TSO performs differently depending on a LAN card's design and its software driver. In general, all of the cards supported by TSO benefit from a reduction in CPU service demand -- those served by the iether driver and those served by the igelan driver.

Depending on your technical computing or enterprise requirements, the GigE product you use with TSO provides the following choices:

- If you want **optimum system service demand**: both drivers and the cards they support will reduce service demand.
- If you want **maximum network throughput**: please be aware that systems using igelan-based cards running near network saturation will show a significant improvement in service demand but a decrease in per card throughput. The possible reduction in throughput on igelan based cards ranges from 2% on smaller systems to 11% on larger systems. In this situation, systems with hardware partitioning would experience the greatest throughput decrease.

The 2-port GigE cards supported by the iether driver include the A7011A and A7012A 2-port PCI-X GigE cards and the AB352A 2-port GigE core I/O card on the rx4640 server.

The igelan driver supports A6825A copper and A6847A fiber GigE cards, the A6794A GigE/SCSI combination card for rp7410, rp7420, and rx7620, the A7109A GigE card for rp8420/rx8620 (either with or without IO expansion cabinet), and the GigE ports on the A9782A/A9784A combination GigE and Fibre Channel cards.

TSO is not supported on GigE cards that use the gelan driver: the A4926A 1000Base-SX and A4929A 1000Base-T cards.

## Test Details




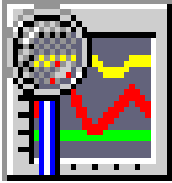
# Test Details

The following subsections describe the hardware and software used in testing and the methods of testing.

## Products Used in Testing

The following products were used for the performance measurement tests:

**Table 1 Products Used in the Performance Measurement Tests**

<b>Server Tested</b>		rp4440 Server. 800 MHz CPU 4-way 16 GB RAM Operating System - HP-UX 11i version 1 (B.11.11.0312). Prerequisite software - The TOUR 2.0 (transport software) needs to be installed in addition to either the IEtherEnh-00 or GigEtherEnh-01 drivers to enable TSO functionality.
<b>Card Tested</b>		A7012A/A7011A 2-port Gigabit Ethernet card PCI-X (64-bit, 133MHz) LAN Driver version – IEtherEnh-00 B.11.11.11
<b>Clients generating the test load for Gigabit Ethernet</b>		Eight rx2600 servers Two, 1.5 GHz Intel Itanium2 CPUs each Operating System - HP-UX 11i version 2 (B.11.23.0303.4) One A6825A PCI 1000Base-T card per rx2600 LAN Driver version – GigEther-00 B.11.23.01
<b>Benchmark software for Gigabit Ethernet tests</b>		Netperf is the benchmarking software suite that generated LAN traffic for the Gigabit Ethernet performance tests. For more information about netperf or to get a free copy, go to <a href="http://www.netperf.org">http://www.netperf.org</a>

## Test Configuration

The test configuration consisted of a 4-way (4-CPU), 800 MHz rp4440 cabinet with 16 Gigabytes of system memory and two A7012A PCI-X 2-port Gigabit Ethernet cards. The card and TSO performance were then measured in a single partition system. The A7011A /A7012A were in slots 7 and 8, the dual-rope or higher-performance slots.

The Gigabit Ethernet test load was generated by connecting the Gigabit Ethernet ports on the cards to one HP 9304M Procurve Routing Switch. Eight 2-way rx2600s each containing a single PCI 1000Base-SX card were connected to the switch (2 rx2600s per port).

### NOTES:

- The networking and I/O performance of the A7011A is identical to the A7012A.
- The observed performance results are consistent across all of the same type of I/O slots of the system.
- The core I/O card in the rp4440 had only minimal site LAN traffic during performance tests.

## For More Information

For more information about the products described in this paper such as a current list of tested HP products or supported systems, please go to:

<http://www.hp.com/products1/unixserverconnectivity>.

This paper is the latest in a series of white papers detailing the performance of HP's link and server products. For a complete list of white papers on HP's networking and I/O products including Gigabit Ethernet and TSO solutions, go to:

<http://docs.hp.com>.

## Taking the Next Step

If you would like to try the TSO product, you can obtain an HP-UX 11i v 1.0-based copy for either the iether or igelan drivers at: <http://software.hp.com> under "enhancement releases and patch bundles."

TSO for both the iether and igelan drivers will be available on the September 2004 OE release of HP-UX 11i v 2.0. On HP-UX 11i v 2.0 versions prior to the September 2004 OE, the iether driver supporting TSO is available as patch PHNE\_30279. Please note that you would need to install PHNE\_30773 and TOUR 2.0 (<http://software.hp.com/portal/swdepot/displayProductInfo.do?productNumber=TOUR>) in addition to PHNE\_30279 in order to enable this functionality. Also note that this functionality is disabled by default on HP-UX 11i v2.0 and can be enabled/disabled on a per interface basis.

For further assistance including a detailed analysis of your specific requirements and needs, please contact your local HP Sales Representative.

## Legal Notices

### Legal Notices

© 2004 Hewlett-Packard Company, L.P. The information contained herein is subject to change without notice.

The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

HP-UX®, Serviceguard®, and Superdome® are registered trademarks of the Hewlett-Packard Corporation. PCI-X is a registered trademark of the PCI SIG.

All other trademarks and registered trademarks are the property of the respective corporations.

