



**Improving Windows<sup>®</sup>  
Performance with TCP/IP  
Offload Engine (TOE) on  
IBM<sup>®</sup> System x<sup>®</sup> Servers**

*Steve Worley  
IBM Corporation*

## **Abstract**

This paper discusses the performance of Microsoft® Windows® when TCP/IP Offload Engine (TOE) is used on IBM System x® servers. TOE provides offloaded processing of the TCP/IP stack from the main system processor to the network controller. This offloading can provide performance advantages for certain configurations and network loads. The results presented will show which configurations can benefit the most from using TOE on IBM System x servers.

## **Introduction**

TCP/IP (Transmission Control Protocol / Internet Protocol) has become the standard protocol used by computers throughout the world. IP contains the information needed to connect two computers and routes data from one computer to another. TCP segments data for transmission and reassembles the data at the receiving end, making sure the data is error-free. TCP requires a great deal of processing to perform these tasks. In the past, attempts have been made to shorten path lengths of the software code making up the TCP stack. An alternative is to offload the TCP stack to the network device itself, thus allowing more processing cycles for the main system CPU(s). TCP/IP Offload Engine (TOE) is the technology that addresses this issue.

## ***Overview of TCP/IP Offload Engine (TOE)***

Because TCP is a connection-oriented protocol, there is a great deal of processing required by the protocol. This processing must establish a connection, acknowledge packets, provide sequence number and checksum calculation, provide congestion control, and terminate the connection. TCP Offload Engine frees the server's CPU(s) for other tasks. There are at least three types of TCP/IP offload: Parallel Stack Full Offload, HBA Full Offload, and TCP Chimney Offload. TCP Chimney Offload is the type of offload discussed in this paper. At this time, TOE is supported on Windows 2003 and 2008. TOE is generally not supported on the various versions of Linux® because The Linux Foundation does not support implementation of TOE.

TOE does not offload every connection. Only older connections are offloaded to the network adapter. The Windows operating system determines which connections will be offloaded. TOE applies only to TCP connections. Other protocols, which are not offloaded, are processed normally by the operating system.

## **Requirements**

The following are required for using TOE with IBM System x servers:

- Microsoft Windows Server™ 2003 Service Pack 1 or later plus the Scalable Networking Pack or Windows 2008
- IBM System x server with Integrated dual Gigabit Ethernet with TCP/IP Offload Engine (TOE) and associated IBM network driver

## **Receive Side Scaling**

IBM System x servers also support Receive Side Scaling (RSS). With RSS, incoming packets are balanced across multiple processors. This is most beneficial for any application with heavy network traffic running on a multi-processor system such file or Web servers.

### Configuration for Analysis

The configuration for the system under test was an IBM System x3650 M2 server with two Intel® Xeon® X5570 (2.93GHz) processors with 12GB (6 x 2GB DIMMs) of memory.

### TOE vs non-TOE Windows 2008 100% Client Writes

Data was gathered for 2- and 4-port TOE and non-TOE configurations for both throughput and CPU utilization. TOE shows lower CPU utilization for most data sizes; however, TOE shows lower throughput for smaller data sizes.

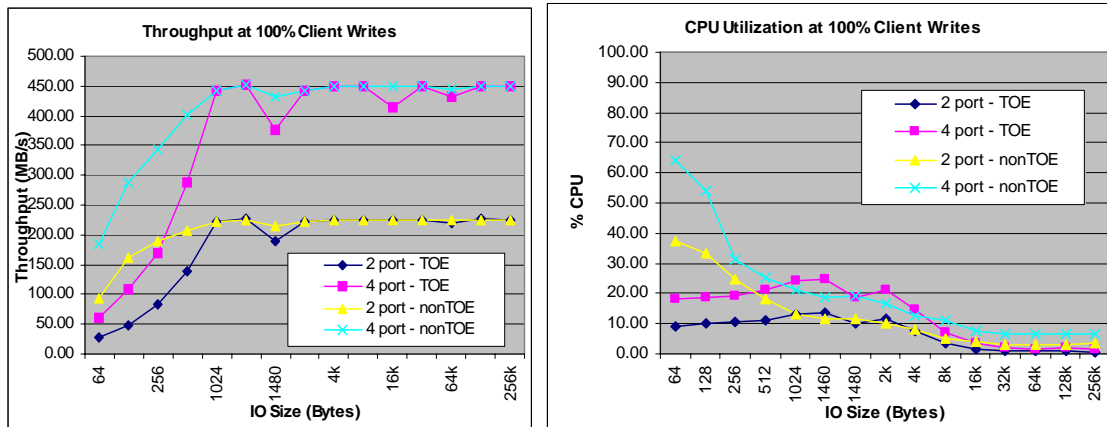


Figure 1: TOE vs non-TOE Windows 2008 100% Client Writes

### TOE vs non-TOE Windows 2008 100% Client Writes Efficiency

Perhaps a better comparison for TOE is the efficiency. Efficiency is the throughput per CPU utilization. TOE shows greater efficiency for large data sizes but very little difference in efficiency for small data sizes.

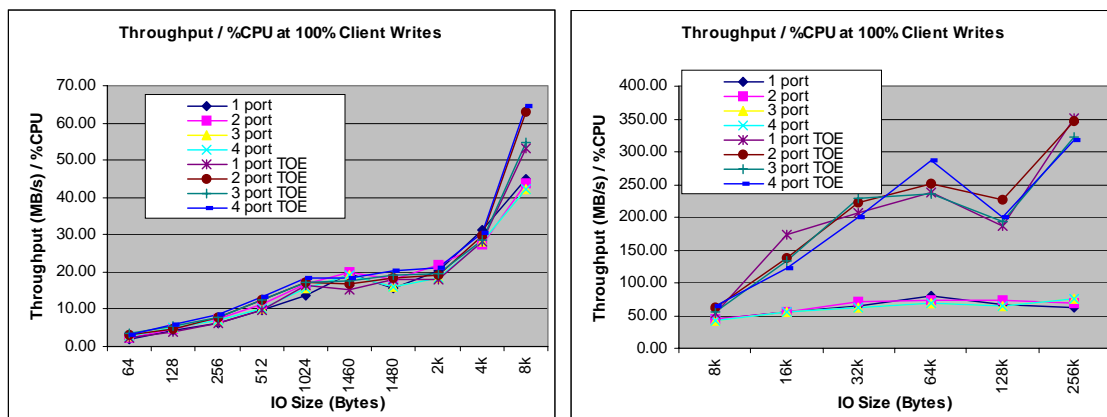


Figure 2: TOE vs non-TOE Windows 2008 100% Client Writes Efficiency

### TOE vs non-TOE Windows 2008 70% Client Reads/30% Client Writes

TOE vs non-TOE shows little difference for 70% client reads/30% client writes.

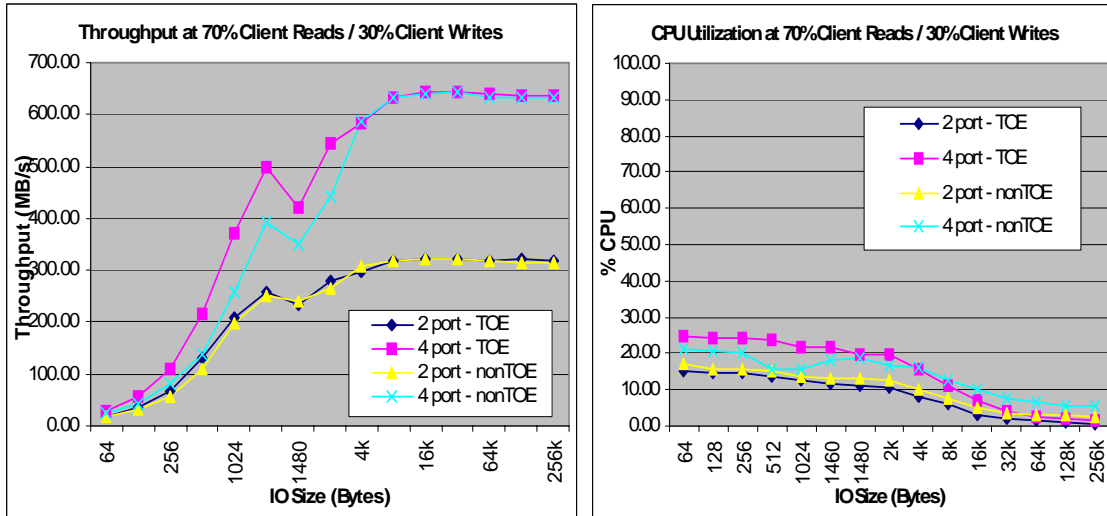


Figure 3: TOE vs non-TOE Windows 2008 70% Client Reads/30% Client Writes

### TOE vs non-TOE Windows 2008 70% Client Reads/30% Client Writes Efficiency

It is more useful to look at efficiency for 70% client reads/30% client write configurations. Just as with 100% client writes, it is more useful to look at efficiency. TOE shows greater efficiency for large data sizes and very little difference in efficiency for small data sizes.

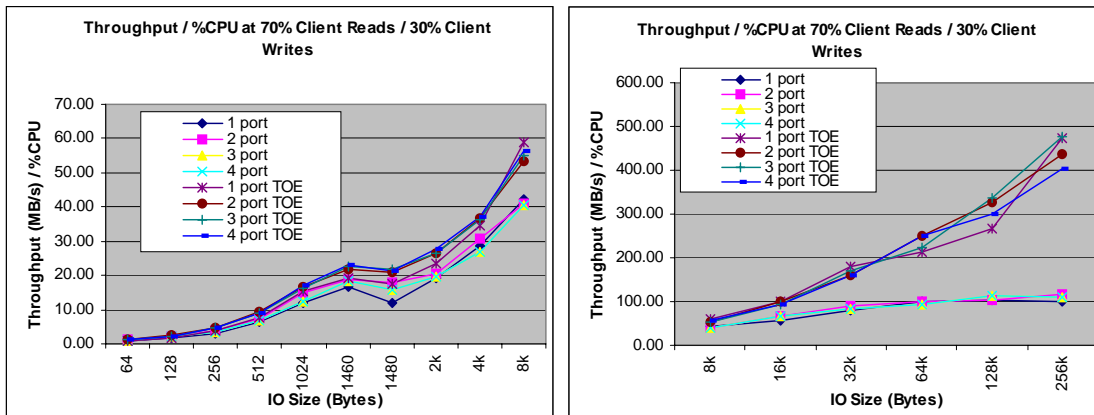


Figure 4: TOE vs non-TOE Windows 2008 70% Client Reads/30% Client Writes Efficiency

### TOE vs non-TOE Windows 2003 100% Client Writes Efficiency

TOE on Windows 2003 shows similar efficiencies to Windows 2008. TOE shows greater efficiency for large data sizes and very little difference in efficiency for small data sizes.

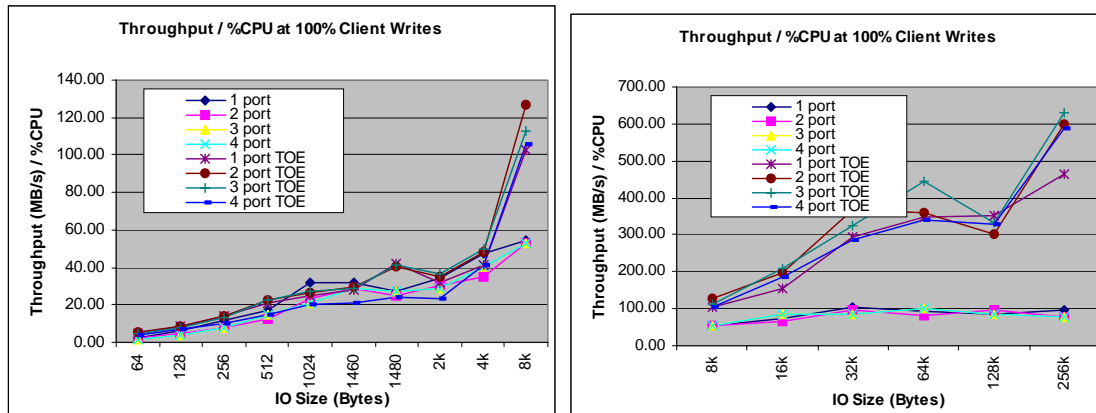


Figure 5: TOE vs non-TOE Windows 2003 100% Client Writes Efficiency

## Conclusion

TCP/IP Offload Engine provides greater efficiency for large data sizes for a mix of traffic of 70% Client Read/30% Client Write. Additionally for 100% client writes, TOE shows lower CPU utilization for most data sizes. TOE also shows greater efficiency for large data sizes. The benefits in performance are similar for both Windows 2003 and 2008.

TCP/IP Offload Engine does not increase wire speeds, but does improve server efficiency. If packets are small and TCP/IP connections do not last long, it does not pay to offload the protocol.

Network performance benefits when the TCP/IP stack is offloaded to an adapter that supports TOE. With IBM System x servers and adapters running TOE, the server's efficiencies can be increased significantly.



© IBM Corporation 2009  
IBM Systems and Technology Group  
3039 Cornwallis Road  
Research Triangle Park, NC 27709

Produced in the USA  
11-09  
All rights reserved

Warranty Information: For a copy of applicable product warranties, write to: Warranty Information, P.O. Box 12195, RTP, NC 27709, Attn: Dept. JDJA/B203. IBM makes no representation or warranty regarding third-party products or services including those designated as ServerProven or ClusterProven.

IBM, the IBM logo, and System x are trademarks or registered trademarks of the International Business Machines Corporation in the United States and/or other countries. For a complete list of IBM Trademarks, see [www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml). Intel and Xeon are registered trademarks of Intel Corporation. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. Microsoft and Windows are registered trademarks of Microsoft Corporation in the United States, other countries, or both. All other products may be trademarks or registered trademarks of their respective companies.

Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Performance is based on measurements using industry standard or IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve performance levels equivalent to those stated here.

IBM reserves the right to change specifications or other product information without notice. References in this publication to IBM products or services do not imply that IBM intends to make them available in all countries in which IBM operates. IBM PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.