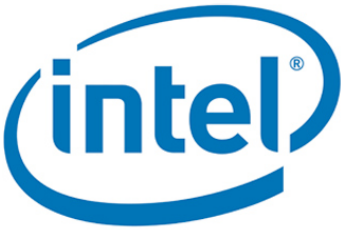

29West Messaging Performance on 10-Gigabit Ethernet

October 10, 2008

Version 1.0

29West, Inc. 29 W 110 Butterfield Rd. Suite 306

Summary

This paper highlights the performance of 29West Messaging on 10-gigabit Ethernet. Commodity Dell workstations running Linux were used along with the new Arista 7124S switch from Arista and NE020 10GbE adapters from Intel/NetEffect. We used 10GBASE-SR PHY and had no interoperability problems.

At this writing, 10-gigabit Ethernet is emerging as a compelling technology for server access to the network. Over 1 million 10-gigabit Ethernet ports shipped in 2007 with most being used for interconnecting network equipment. Many 29West customers are planning deployments of 10-gigabit Ethernet for server access so we have run a series of benchmarks to help our customers estimate the performance they may be able to achieve. Our key findings from these benchmarks include:

- **Payload delivery rates** as high as **9.8 gbps** for large messages
- **Message delivery rates** more than **2.3 million messages per second** for small messages
- **Average ping-pong latency** of **36 microseconds**
- **Scalable performance** as receivers are added when using UDP multicast
- **Message delivery latency** under **50 microseconds** for rates up to 110,000 messages per second

We ran tests using TCP and UDP multicast transport protocols. The benefits of stateless TCP offload features in NICs and accompanying kernel support for these features can be seen in the TCP throughput results. Vendors are now developing kernel bypass libraries that will offer better performance for both UDP and TCP. 29West is planning a follow-up report when such libraries are available and reliable.

10-gigabit Ethernet, once reserved for network-to-network connections, is now a viable alternative for server access to high-speed networks. Latency is low, as is latency jitter. Throughput can reach wire speed with TCP and 3,200-byte messages while UDP rates seem to be limited to about 40% of that until reliable kernel bypass libraries are available. UDP multicast provides the advantage of wire-speed copying for all ports of modern switches, delivering linear growth in receive rates as receivers are added.

The benchmark numbers given here show what we measured using equipment in our testing labs. The numbers that really matter are those run on your production equipment. We invite you to try our messaging software on your production equipment through our free evaluation program.

Test Setup

The system consists of two Dell Precision T3400 machines with Q6600 CPUs networked with 10-gigabit Ethernet running Centos 5 Linux. The measurements were performed using version 3.3.6 of 29West, Inc. Latency Busters® Messaging.

About 29West

29West is the leader in high-performance, low-latency messaging solutions for financial institutions. With its initial release in November 2004, 29West's Latency Busters® Messaging (LBM) set a new standard in performance for financial market messaging and has been deployed in more than 100 firms worldwide. With the introduction of Ultra Messaging® for the Enterprise (UME) at the end of 2006, 29West brought the unique Parallel Persistence® design to guaranteed messaging. Where other solutions send first to a store and then to the end receiver, 29West UME solutions send to the end receiver in parallel with delivery to the store, resulting in dramatic increases in throughput and drops in latency. With offices in Chicago, New York, London and Tokyo, 29West supports the financial markets worldwide. For more information, visit http://www.29west.com