

29
WEST

Now a part of
INFORMATICA®
The Data Integration Company™

29West Ultra Messaging® IPC Performance
with the
Cisco Unified Computing System

by Greg Lorence

May 2010

Contents

- Executive Summary** 1
- Results** 2
 - Latency Testing 3
 - Thoughtput Tests 4
 - Full Results 5
 - Latency Tests 5
 - Throughput Tests 5
- Methodology of Measurements** 6
 - Throughput 6
 - Latency 6
- Conclusion** 7
- Detailed Test Configuration** 8
 - Hardware Configuration 8
 - BIOS Configuration 8
 - Operating System Configuration 9
 - Software Configuration 9
- About 29 West** 10

Executive Summary

Experts agree that in the future CPU scaling will go out, not up. Clock cycles seem to have plateaued around 3 Gigahertz, and while five years ago people got excited about dual-core machines, today the major commodity CPU manufacturers deliver 4, 6 and 8-core processors and system vendors link these CPUs together in multi-socket servers, as Cisco does in their UCS line.

Financial players like 29West's exchange and investment banking customers are always interested in gaining whatever performance advantage they can in the ultra competitive world of high-frequency electronic trading. In order to take advantage of this highly horizontal movement in CPU architecture, 29West, now a part of Informatica, has integrated a new mode of transporting data from application to application via shared memory. Called IPC (Inter-Process Communication), this transport mechanism routes data messages among applications all running on the same physical machine.

The following report outlines a recent testing effort by Cisco and 29West to highlight the kind of next-generation performance that is achievable when using world-class enterprise software and hardware. As a result of these tests, we are able to show:

- Sub-microsecond average latencies, as low as **660 nanoseconds**
- Maximum fan-out throughput of **21.1 million messages per second**
- Nearly perfect **linear horizontal scaling** to multiple CPU cores

The full test results are provided here. We also explain the testing methodology, along with full descriptions of the technology stack under test, hardware and software.

Results

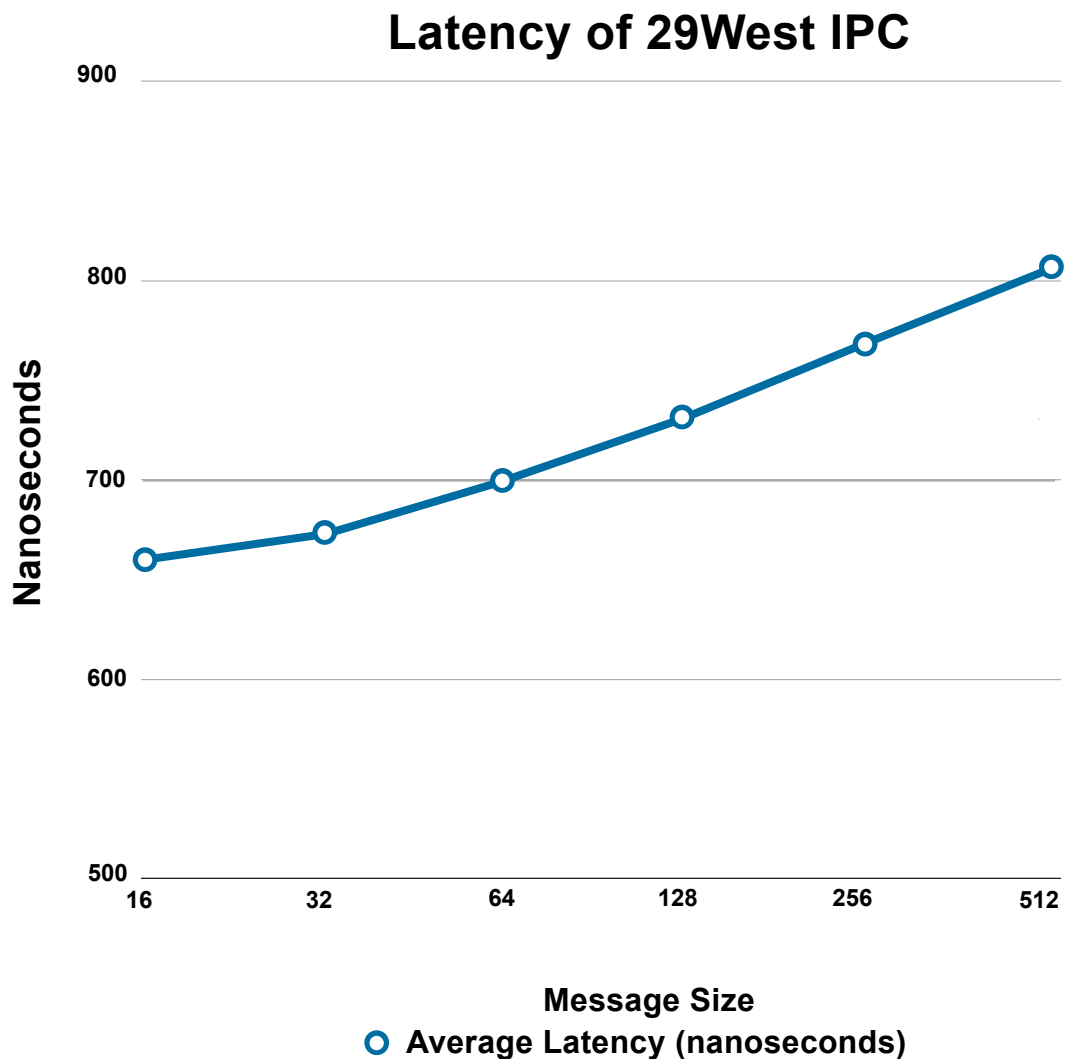
Testing was performed to determine the maximum performance of the technology stack for two simple benchmarks: highest throughput and lowest latency.

Throughput tests were done using both a 1:1 model (one sending process and one receiving process, each running on an individual CPU core) and a 1:11 fanout model (1 sending process to 11 receiving processes, using all 12 available CPU cores) to test not only how fast one application would be, but to show how well the multi-core architecture of the new Intel CPUs scales horizontally.

Latency testing was done using a 1:1 model.

Latency Testing

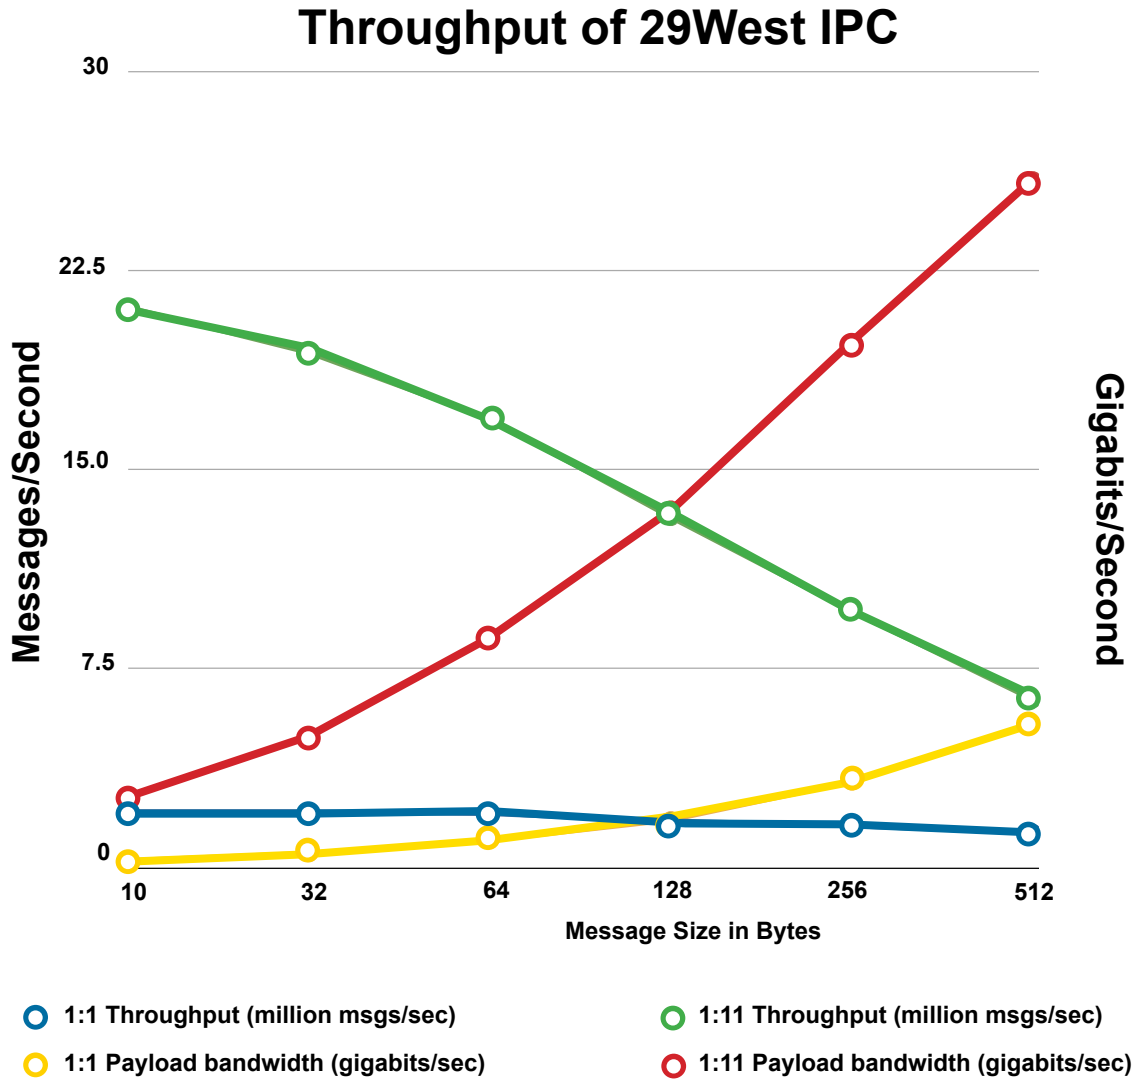
The goal of the latency tests is to show the remarkable speed of the LBM IPC transport when run on Cisco's world-class enterprise UCS servers. The results are impressive: an order of magnitude faster than the fastest network-based transports we've tested.



Observed latencies were significantly under 1 microsecond for all message sizes through 512 bytes. This shows an increase in latency as message size increases, similar to all other transports 29West has tested.

Throughput Testing

The graph below visually represents the result of our throughput testing:



These results show extremely good throughput in both the 1:1 case and the 1:11 fanout case. Over 2 million messages per second per thread is significantly more throughput than any network-based benchmark 29West has ever seen. Not to mention the excellent horizontal scalability shown in the fanout test, which shows a mere 9% drop from perfectly linear scaling.

Full results of our throughput testing appear below.

Full Results

Latency Tests

Message Size (bytes)	Average Latency (nanoseconds)
16	660
32	672
64	701
128	732
256	763
512	803

Throughput Tests

Message Size (bytes)	1:1 Throughput (millions of messages/sec)	1:11 Fan-out Throughput (millions of messages/sec)
16	2.11	21.1
32	2.11	19.4
64	2.04	16.9
128	1.76	13.2
256	1.58	9.66
512	1.32	6.34

Methodology of Measurement

Throughput

Throughput measurements were performed using the 29West *lbmsrc* and *lbmrcv* applications. *lbmsrc* generates a series of messages, as fast as it can, and publishes them to an IPC shared memory segment. *lbmrcv* meanwhile, reads messages from the shared memory segment simultaneously, also going as fast as it can.

In the 1:11 fanout tests, the same is true, except that there are 11 identical *lbmrcv* applications running, and they are all reading from the same IPC shared memory segment. Total throughput numbers in the fanout tests are therefore the aggregate of all 11 receiving applications, which are all copies of the originally published set of messages.

Latency

Latency measurements were performed using the 29West *lbmping* application. The *lbmping* app has two sides: the 'ping' side and the 'pong' side. On the ping side, messages are generated and sent to the 'pong' side application. The pong side simply reflects the messages it receives back to the sending application, using the same settings as the pong side, to ensure that both sides of the round-trip will be equal.

When measuring extremely fast transports, like the 29West IPC transport, because of the amount of CPU cycles taken by the timestamp system call, taking even one timestamp can materially skew the results. For this reason, the *lbmping* applications were configured to take only two timestamps for the entire test run. One timestamp is taken at the very beginning of the test, and one is taken at the very end. The difference in time is then divided by the total number of messages sent. The result of this calculation is the average round-trip latency for each message in the batch. This number is then divided by 2 to get the one-way latency. To ensure an accurate average, 500,000 samples were taken for each test run.

Conclusion

29West Ultra Messaging IPC Transport can send and receive information inside of a fully publish/subscribe, powerful messaging API at significantly sub-microsecond latency. The more than *2 million messages per second per thread* is a level of performance is unmatched in the industry, and shows why Ultra Messaging IPC transport is a compelling solution for high-frequency trading and other high-performance application systems, providing nanosecond data delivery to the businesses that deploy 29West Ultra Messaging on Cisco's UCS servers.

Detailed Test Configuration

Hardware Configuration

The test system consists of one machine, with 2 sockets, 6 cores per socket.

Hardware Configuration Table

Processors	2 Intel Xeon X5670(six-core) @2.93GHz
RAM	96 Gigabyte
Cache	8192 kB (each)

BIOS Configuration

The following changes were made in the BIOS of both sending and receiving machines:

BIOS Configuration Table

C3-Report	Disable
C6-Report	Disable
Hyperthreading	Disable
Turbo Mode	Enable
Intel SpeedStep	Enable

Operating System Configuration

The operating system used was an unmodified version of 64 bit RHEL5, kernel version 2.6.32.9-67.fc12.x86_64

Software Configuration

The measurements were performed using version 4.0EA2 of 29West LBM

Software Configuration Table

Transport	lbtipc (29West Inter-Process Communication)
Advertisement Minimum Interval	0
Advertisement Sustain Interval	0
Query Initial Interval	0
Query Threshold	1
Receiver Thread Behavior	busy_wait
Session Message Interval	2000
Source Behavior	receiver_paced
Transport Activity Timeout	3000

About

29West

29West is a wholly-owned subsidiary of Informatica Corporation and the leader in high-performance, low-latency messaging solutions for financial institutions. The Ultra Messaging® family of products provides up to a 10X latency reduction and 100X bandwidth improvement over traditional messaging designs across a wide range of use cases, including streaming, persistence, caching, queuing and desktop distribution. 29West provides customers with the highest performing, most tunable, most controllable, most powerful messaging API in the industry, and allows full end-to-end monitoring capabilities. With more than 140 production deployments since 2004, 29West provides enterprise strength, enterprise-wide messaging for next-generation trading platforms. For more information, call +1 630 836 2990 or send an email to info@2west.com.

Informatica

Informatica Corporation (NASDAQ: INFA) is the world's number one independent provider of data integration software. Organizations around the world gain a competitive advantage in today's global information economy with timely, relevant and trustworthy data for their top business imperatives. More than 4,000 enterprises worldwide rely on Informatica to access, integrate and trust their information assets held in the traditional enterprise, off premise and in the Cloud. For more information, call +1 650-385-5000 (1-800-653-3871 in the U.S.), or visit www.informatica.com.

Cisco

Cisco, (NASDAQ: CSCO), the worldwide leader in networking that transforms how people connect, communicate and collaborate, this year celebrates 25 years of technology innovation, operational excellence and corporate social responsibility. Information about Cisco can be found at www.cisco.com. For ongoing news, please go to <http://newsroom.cisco.com>. Cisco products are supplied in Taiwan by Cisco Systems International, BV, a wholly owned subsidiary of Cisco Systems, Inc.

Contact Information

We look forward to the opportunity to answer any questions you may have about [29West Ultra Messaging®](#).

We have offices in Chicago, New York, London and Tokyo and would welcome a chance to discuss your needs and see how we can help.

For more information on [29West Ultra Messaging®](#) or to request a *Free Software Evaluation*, please call or e-mail us at an office listed below. If you would like to be added to our mailing list or receive copies of our networking and messaging white papers, please e-mail us at sales@29West.com.

If you would like to learn more about 29West, now a part of Informatica, please visit www.29West.com or contact us via e-mail at info@29West.com.

Chicago

29W110 Butterfield Road
Suite 306 Warrenville, IL 60555
Phone: 630-836-2990
Fax: 630-836-7508
E-mail: info@29west.com

New York

100 Church Street, 8th Floor
New York, NY 10007
Phone: 646-429-0800
E-mail: kclancey@29west.com

London

1st Floor
10 Arthur Street
London, EC4R 9AY
Phone: +44 (0) 207 122 0310
E-mail: annalisa@29west.com

Tokyo

Burex Kojimachi
3-5-2 Kojimachi, Chiyoda-ku
Tokyo, 102-0083 Japan
Phone: +81 (0)3 6268 9803
E-mail: swatanabe@29west.com