

Executive Summary of STAC Report on Juniper/IBM/Solarflare

Summary

The Securities Technology Analysis Center (STAC[®]) is a vendor-neutral specialist in high-performance capital markets technology. STAC facilitates the STAC Benchmark Council[™], which creates a set of STAC Benchmarks[™] to help vendors and customers understand the performance of workloads in the financial industry using different products in combination. Juniper joined the Council in January of 2011.

STAC Benchmarks are defined by end-user IT executives within the financial industry with input from vendors. Once benchmarks are approved, the supporting STAC Test Harness[™] software is made available to vendor and end-user members of the Council. Vendors wishing to publish STAC Benchmark results must have those results audited, whereby STAC runs tests under carefully controlled standard scenarios on the vendor's choice of hardware and software.

Juniper Networks, IBM, and Solarflare Communications have completed an audited STAC-M2[™] Benchmark, which tests the ability of a given "stack" of switches, middleware, NICs and servers to handle real-time market data under a variety of conditions. The tests provide key performance metrics such as latency, throughput, power efficiency, and CPU/memory consumption under several scenarios, including both undisturbed flow and exception conditions like slow consumers.

This set of products has generated the lowest latency and highest messages per second (msg/sec) throughput seen to date in any audited report in which a 10 Gigabit Ethernet (GbE) switch was the network transport. This summary report captures a small snapshot of the tests available in the STAC Report Highlights published by STAC at www.STACresearch.com/juniper. The full STAC Report[™] is available in the STAC Vault[™] to STAC's premium end-user subscribers.

Test Overview

Terminology

STAC-M2 uses a software library created by STAC to simulate flows of market data in a trading environment. Controls are provided in the test suite to vary the data rate as well as simulate a variety of traffic conditions, such as a slow consumer. Key terms:

- **Producer:** The Producer is the software application that uses the STAC Library to initiate the message flows. In different tests there are one or two Producers. The Producer in these tests used IBM[®] WebSphere[®] MQ Low Latency Messaging v2.4.0.2 (LLM) software to send messages via multicast so that multiple consumers could receive a single message flow.
- **Consumer:** The Consumer is the software application that receives messages. Using the STAC Library, the Consumer takes the messages it receives from the Producer and reflects a small proportion of them back to the Producer. It is these reflected messages that are used to ensure messages are being delivered and consumed properly, as well as measure the round-trip latency.

- **Latency:** Latency is measured from the time the STAC Library supplies the message to be sent to the Producer application to the time the reflected message is received back from the Consumer and received by the STAC Library in the Producer. This latency is divided by two to yield what STAC-M2 designates as “Hybrid Latency,” which is an estimate of the one-way latency from Producer to Consumer.
- **Stack Under Test (SUT):** The SUT is the combination of the vendor software and hardware that comprise the environment used for the STAC audit. The STAC tests measure the performance of the complete SUT, not just any single component.

Environment Overview

Juniper Networks

The Juniper Networks QFX3500 was used as the network hardware for this environment. Featuring 48 dual-mode small form-factor pluggable transceiver (SFP+/SFP) ports and four quad small form-factor pluggable plus (QSFP+) ports in a 1 U form factor, the QFX3500 Switch delivers feature rich Layer 2 and Layer 3 connectivity to networked devices such as rack servers, blade servers, storage systems, and other switches in highly demanding, high-performance data center environments. For converged server edge access environments, the QFX3500 is also a standards-based Fibre Channel over Ethernet (FCoE) Transit Switch and FCoE to Fibre Channel (FCoE-FC) Gateway, enabling customers to protect their investments in existing data center aggregation and Fibre Channel storage area network (SAN) infrastructures.

Solarflare Communications

The Solarflare SFN5122F is a low-latency, low-power 10GbE server adapter. Solarflare Server Adapters are designed to provide high performance in the most demanding application environments. Solarflare's OpenOnload application acceleration middleware was used in combination with the SFN5122F to enable full operating system bypass, which Solarflare reports dramatically reduces host processing overheads and enables high transaction rates while substantially reducing application latency with very low jitter. OpenOnload performs network processing at user level and is binary compatible with existing applications that use TCP/UDP with BSD sockets.

IBM

IBM's[®] WebSphere[®] MQ Low Latency Messaging v2.4.0.2 (LLM) is a messaging transport that IBM targets at the very high-volume, low-latency requirements of financial market firms. The product is daemonless and provides peer-to-peer transport for one-to-many and many-to-many multicast messaging, as well as point-to-point unicast messaging.

IBM also supplied sixteen IBM[®] x3550 M2 Model 7946 servers, equipped with Dual Intel Xeon Processor X5570 running Red Hat Linux.

Interpreting Results

The STAC-M2 Test Harness software produces a standardized report with a set of tables summarizing the results. Among other things, these numbers describe how many messages were passed in a second, how fast they were passed, and the deviation (jitter) observed.

Highest supply rate: The maximum number of messages sent per second without causing lost messages or other conditions defined as failures by the benchmark specs.

Mean latency: The arithmetic mean—or average value—of hybrid latency for every message measured during the three-minute test cycle.

Max latency: The highest value of hybrid latency measured during the test cycle, signifying the worst-case message completion at a given load.

Standard deviation of latency: The deviation is a particularly important measure, because it signifies how predictable (or deterministic) the traffic is across the environment. A low standard deviation implies that almost all traffic completed close to the average, or mean, latency value.

Results Overview

The following table highlights some of the high-level results from a single test sequence from the STAC-M2 Audit and compares them against the other STAC-M2 Audits that have been performed. While the stacks under test did not always have the same middleware, NIC, or server, the comparisons are nonetheless interesting.

STAC-M2 “BASELINE” Test Sequence (1 Producer to 5 Consumers)

(See the links below for the STAC Report Highlights for each product stack)

Description	QFX3500/IBM LLM	Cisco 4900M/29West	Cisco Nexus 5010/29 West	Voltaire IB/IBM LLM
Stack Under Test (SUT) ID	LLM110203	LBM100414	LBM100402	LLM091105
Highest Supply Rate (msg/sec) STAC.M2.v1.0.BASELINE.TPUT1	1,500,000	1,300,000	1,300,000	1,000,000
Mean latency (microseconds) STAC.M2.v1.0.BASELINE.LAT1.MEAN	9	15	14	8
Max latency (microseconds) STAC.M2.v1.0.BASELINE.LAT1.MAX	23	30	33	47
Standard deviation of latency (microseconds) STAC.M2.v1.0.BASELINE.LAT1.STDV	0*	1	1	1

* Less than the current STAC-M2 precision of 1 microsecond.

Summary

The STAC-M2 Benchmark tests represent a simulation of market data messaging performance.

The Juniper QFX3500, in combination with IBM server and middleware with Solarflare NICs, delivered the best public performance to date for product combinations involving 10GbE switches.

In nearly every test setup, this product combination delivered more messages faster with lower jitter than any other audited STAC-M2 Benchmark.

The low latency and jitter results of this environment guarantee the highest performance and the most predictable delivery of messages, which are critical for transactions within financial trading environments.

Links

For additional information on the products used in this STAC-M2 Audit, please see the following links:

STAC Highlights Report for Juniper/IBM/Solarflare: www.STACresearch.com/juniper

STAC Highlights Report for Cisco4900M/29West: <http://www.stacresearch.com/node/6839>

STAC Highlights Report for Cisco Nexus 5010/29West: <http://www.stacresearch.com/node/6792>

STAC Highlights Report for Voltaire/IBM/Solarflare: <http://www.stacresearch.com/node/5947>

STAC-M2 Test Harness: http://STACresearch.com/STAC-M2_Overview.pdf

Juniper Networks QFX3500 Switch data sheet: <http://www.juniper.net/us/en/local/pdf/datasheets/1000361-en.pdf>

IBM x3550 Server: <http://www-03.ibm.com/systems/x/hardware/rack/x3550m3/index.html>

IBM WebSphere MQ Low Latency Messaging Software[®] <http://www-01.ibm.com/software/integration/wmq/llm/>

Solarflare SFN5122F Dual-Port 10GbE Adapter[®] http://www.solarflare.com/products/products_adapters.php