THE TOLLY GROUP

No. 208346 OCTOBER 2008

BLADE Network Technologies, Inc.

RackSwitch G8100

Competitive Performance Evaluation versus Cisco Catalyst 4900M Switch

Tolly Up to Spec Certified

Test Summary

Premise: *As bandwidth- and latency-sensitive applications for Web 2.0, high-performance computing (HPC) clusters and virtualized data centers proliferate, enterprise data centers that have consolidated their server resources using blade server systems with 10 Gigabit Ethernet are discovering that top-of-rack 10GbE low-latency switches can provide unified and rack-level virtualization of servers, blade server systems and storage resources. Network performance and cost remain the most important factors when choosing an Enterprise data center solution. However, low power consumption, optimal cooling and floor space are also critical requirements.*

BLADE Network Technologies, Inc. (BLADE) commissioned The Tolly Group to evaluate the Layer 2 performance and power consumption of its RackSwitch G8100 10G Low Latency Switch with 24 10GbE ports — designed for top-of-rack aggregation for the large data center — against a Catalyst 4900M from Cisco Systems, Inc.

Tolly Group engineers measured the Layer 2 throughput, latency and multicast performance of the switches in a variety of common network topologies as well as the energy usage for the switch operations with and without load. Engineers also compared the price/performance of the switches from both companies.

Tests were conducted in September 2008.

Test Highlights

- Demonstrates an average of 12.4X less latency than a comparable Cisco solution for various frame sizes in a port-to-port configuration using 20 10GbE CX4 ports
- Sustains line-rate throughput for all frame sizes tested up to 9,216 bytes in a 24-port full-mesh test, while the Catalyst 4900M achieves only 49.7% of the theoretical maximum
- Consumes 73% less energy than the Catalyst 4900M when handling 50% traffic load with all ports active, which is equivalent to $227 per year, per switch deployed
- Exhibits 6.5X better price/performance compared to the Catalyst 4900M

BLADE RackSwitch G8100 vs. Cisco Catalyst 4900M Layer 2 Latency over 20 10GbE Ports in a Port-to-Port Scenario

as Reported by Spirent TestCenter 2.22

(Lower bars are better)

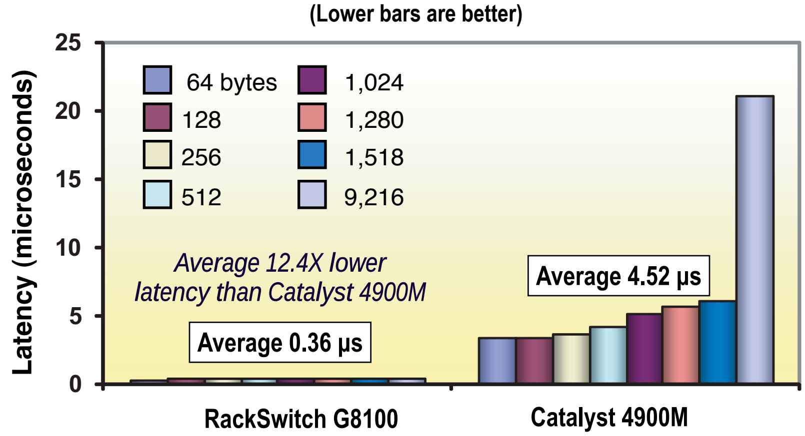

Notes: 20 10GbE CX4 ports were used. Engineers measured cut-through latency for the RackSwitch G8100 and store-and-forward latency for the Cisco 4900M. Latency was measured at 45% of theoretical maximum. Average did not include Jumbo Frame result.

Source: The Tolly Group, September 2008

Figure 1

Executive Summary

BLADE's RackSwitch G8100 outperforms the comparable Cisco solution by demonstrating 2X higher throughput, 12.4X lower latency, 2.7X greater energy efficiency and 6.5X better price/performance than the Cisco Catalyst 4900M.

Application-driven network architectures are very common in large Enterprise data centers. As the application complexity increases, different network design and operations are required. However, non-blocking switch architecture and extremely low latency have been the best weapons to deal with the new demands from ever evolving applications. Recently, energy efficiency has attracted more attention than ever before, becoming a key issue in the Enterprise data center.

The Tolly Group's hands-on evaluation of BLADE's RackSwitch G8100, based on a FocalPoint FM4224 low-latency 10GbE switch chip from Fulcrum Microsystems, demonstrates that the device offers non-blocking 10GbE switching, far lower latency and energy consumption than the Catalyst 4900M. These are very important productivity indicators for the top-of-rack switches deployed in the data center where latency sensitive applications such as HPC clusters and financial applications set strict requirements. As data centers move forward to Converged Enhanced Ethernet (CEE), the high-bandwidth, non-blocking network is required to support loss-less I/O transactions over Ethernet such as FibreChannel over Ethernet (FCoE), iSCSI and IP NAS. (CEE, or Data Center Ethernet, is an enhanced Ethernet that enables convergence of various applications in data centers [LAN, SAN and HPC] onto a single interconnect technology.)

Layer 2 Throughput Comparison of BLADE RackSwitch G8100 vs. Cisco Catalyst 4900M across 24 10GbE Ports in a Full-Mesh Configuration
as Reported by Spirent TestCenter 2.22

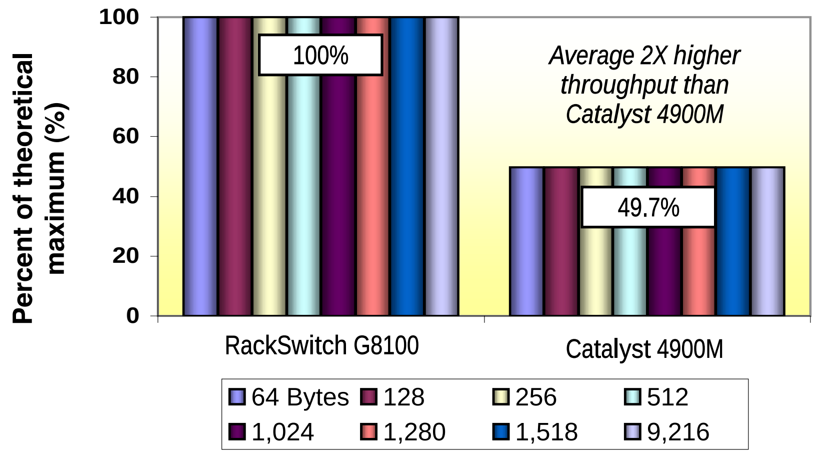

Notes: 20 10GbE CX4 ports and 4 10GbE SR ports were used for the measurement.

Source: The Tolly Group, September 2008

Figure 2

The throughput test proves that the RackSwitch G8100 sustains wire-speed throughput across all frame sizes tested in a 24-port full-mesh configuration due to its non-blocking switch fabric. The Cisco solution handled only half of the traffic that the RackSwitch G8100 processed.

The performance differences between these two solutions become even more apparent in the latency test. While the BLADE RackSwitch G8100 offers under 0.4 microsecond packet switching for all frame sizes tested from 64 bytes to 9,216 bytes, the Catalyst 4900M showed an average of 12.4X higher latency by recording 3~21 microseconds of latency.

The power consumption test proves that data center users can reduce electrical costs by 75% with the RackSwitch G8100 versus Catalyst 4900M under normal conditions.

Another feature that makes the BLADE RackSwitch G8100 more suitable for the data center rack environment is its cooling feature. The RackSwitch G8100 provides a choice of either front-to-rear or rear-to-front cooling to match the airflow of other servers in the rack.

These tests, combined with the price/performance analysis, illustrate that BLADE's RackSwitch G8100 is a viable solution for data center users who look for cost-effective, strong Layer 2 network performance and an energy-efficient top-of-rack solution. For the test, Tolly Group engineers focused on measuring the base system Layer 2 performance in their default configurations.

Results

Layer 2 Throughput

The BLADE RackSwitch G8100 achieved 100% zero-loss throughput for all standard frame sizes, as well as one Jumbo Frame size (64 to 9,216 bytes) tested with 24 10GbE ports in a full-mesh configuration. Even though The Tolly Group's acceptable frame loss is $\leq$0.001%, the BLADE RackSwitch G8100 did not drop any packets while handling the traffic at line rate. On the other hand, a similarly equipped Cisco solution only achieved 49.7% of zero-loss throughput across all the frame sizes tested. This is due, primarily, to the low-capacity switch fabric from the eight-port 10GbE half-cards.

Layer 2 Latency

Tolly Group engineers measured the latency of the BLADE RackSwitch G8100 and the Cisco Catalyst 4900M. BLADE's Layer 2 performance advantage was even more apparent in this test.

The RackSwitch G8100 achieved an average of 0.36 µs of latency by measuring from 0.31 µs to 0.38 µs for 64-byte and 1,518-byte frames, respectively. This represents 10X to 16X less latency than the comparable Cisco solution. The Catalyst 4900M measured from 3.3 µs to 6.1 µs for 64-byte to 1,518-byte frames, respectively. See Figure 1. While 9,216-byte Jumbo Frames were tested, they were not factored into average latency since they represent a small percentage of traffic.

Layer 2 Multicast Forwarding

Multicast data was transmitted from one 10GbE port to a set of 23 10GbE ports. This test simulates a "1 to 23" fan-out with Layer 2 unidirectional multicast traffic with 10-Gbps test traffic of the standard Ethernet frame sizes between 64 bytes to 9,216 bytes.

Similarly, with the throughput test, the RackSwitch G8100 achieved 100% throughput for all frame sizes. However, the Catalyst 4900M was able to handle half the traffic the RackSwitch G8100 processed.

Power Consumption

Tolly Group engineers found that the RackSwitch G8100 used almost 83.4 watts versus 351.9 watts for the Catalyst 4900M in an idle state scenario with all ports active. This equates to US$222.50 savings per year. This power consumption does not include the separate power usage to cool the equipment. This heat dissipation is normally in proportional to the power consumption.

Even when data traffic passed through the switches, the performance gap remained the same.

With 50% of traffic load, the RackSwitch G8100 consumed 101.3 watts (0.84 watts per Gbps), while the Catalyst 4900M consumed 375.3 watts (3.14 watts per Gbps). This proves that the RackSwitch G8100 used 73% less energy than the Catalyst 4900M and, in turn, generates US$227 in savings yearly on a per-switch basis.

With 100% of traffic load, the RackSwitch G8100 consumed 119.7 watts. This represents 0.50 watts per Gbps and a single-year operational cost of just over US$99. The Catalyst 4900M was not included in this test because it could not handle more than 50% of the theoretical maximum throughput. See Figure 4.

Price/Performance Comparison

Engineers calculated the price-performance ratio of the DUTs in terms of US dollars per Gbps (US$/Gbps) of maximum throughput measured by Tolly Group engineers.

For this comparison, engineers configured the DUTs as similarly as possible. Both base systems were outfitted with two power supplies, 20 10GbE CX4 ports and four 10GbE SR ports.

The prices of the RackSwitch G8100 were obtained from BLADE Network Technologies' sales channel, while those for the Cisco Catalyst 4900M were obtained from an invoice of the DUT, omitting the reseller discount to arrive at the manufacturer's list price.

BLADE's RackSwitch G8100 offered a 6.5X better price/performance ratio than the Catalyst 4900M with a cost of $68.54 per Gbps. The Catalyst 4900M as tested cost $443.72 per Gbps. See Figure 4.

Test Setup & Methodology

The test bed consisted of the DUT and a traffic generator. See Figure 6 for the test bed diagram.

BLADE Network Technologies

RackSwitch G8100

Throughput, Latency, Multicast Forwarding & Energy Consumption

TOLLY Up to Spec CERTIFIED

Designed as a blade switch

BLADE's RackSwitch G8100 is not just any other 10-Gigabit Ethernet switch; its design is a direct result of the company's heritage as the leading manufacturer of blade server switches.

In addition to its extremely low latency and line-rate throughput, the switch design was perfected by engineers who are used to designing networking equipment that meets the strict power and space constraints of blade server systems. RackSwitch G8100 was designed from the ground up with the most efficient silicon, power management subsystem and printed circuit board design. BLADE's RackSwitch G8100 also employs airflow efficiencies with the data center environment in mind.

Two RackSwitch G8100 models are available—with either front-to-rear or rear-to front airflow—that match the airflow of the servers in the racks to maintain separation of hot aisles and cool aisles in the data center. BLADE RackSwitch G8100 uses the same operating system as the company's blade server switches for IBM HP and NEC, and can provide rack-level provisioning of servers and storage supporting network virtualization, lower power and cooling costs, and easier management. The G8100 is compatible in Cisco core networks, and can provide rack-level provisioning of servers and storage supporting network virtualization, lower power and cooling costs, and easier management. Learn more about "Rackonomics"—Virtual, Cooler and Easier networking—at www.bladenetwork.net/rackonomics

Source: BLADE Network Technologies Sept. 2008

Avg. Power Consumption of BLADE RackSwitch G8100 vs. Cisco Catalyst 4900M and Projected One-Year Operational Costs with Different Traffic Loads
as Reported by "Watts up? PRO"

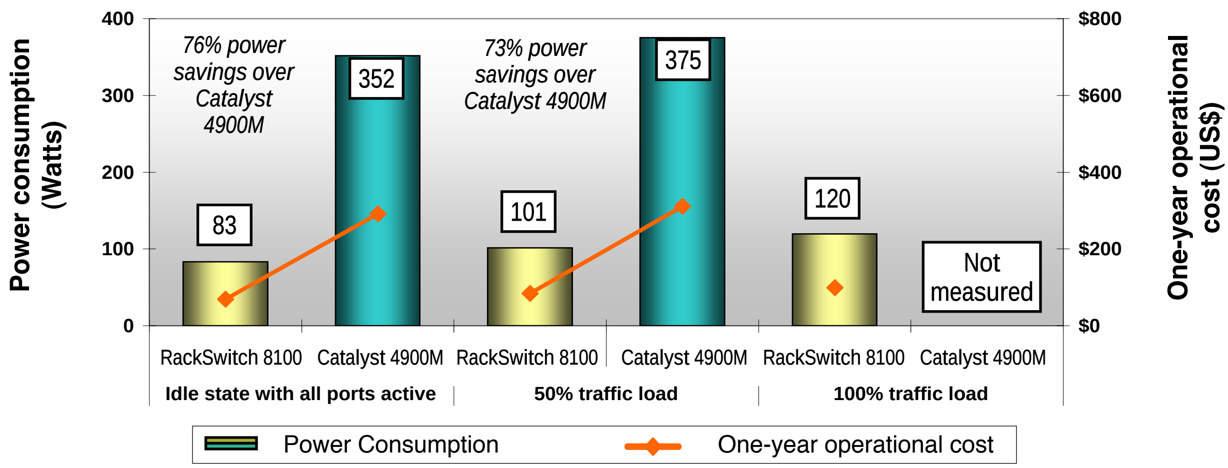

Notes: 20 10GbE CX4 ports and 4 10GbE SR ports were active. Power consumption of the Catalyst 4900M was not measured at 100% load because the switch cannot handle more than 50% of the theoretical maximum throughput. Cost projection is based on the 2006/2007 Average Commercial Electric Price of US$0.0946 per kilowatt hour. The power consumption estimates in this table do not include the energy for cooling so the actual energy consumption and operational costs would be higher than reported here.

Source: The Tolly Group, September 2008 *Figure 3*

The DUTs were BLADE Network Technologies RackSwitch G8100 (Ver. 1.0.1.4) and Cisco Systems Catalyst 4900M top-of-rack switch (IOS version 12.2(40)XO). These switches are specifically designed for data-center networks. In this test, both switches were equipped with 20 10GbE CX4 and four 10GbE SR ports, as well as two AC power supplies. RackSwitch G8100 provides two AC power supplies by default for high-availability.

The RackSwitch G8100 was outfitted with a base system (BN-8100F-BDL) and four SFP+ Transceivers (BN-CKM-SP-SR).

The Catalyst 4900M switch was outfitted with a base system (WS-C4900M), two AC power supplies (PWR-C49M-1000AC), two 8-port 10GbE half cards (WS-X4908-10GE), 20 10GbE CX4 transceivers (X2-10GB-CX4) and four 10GbE SR transceivers (X2-10GB-SR). A Spirent TestCenter SPT-9000A test traffic generator was used with MSA-2001A (two-port 10GbE) test modules, with TestCenter application version 2.22.0606.

Test Tool Accuracy

The Spirent TestCenter 9000A solution, using MSA-2001A modules, has been tested back-to-back to measure the inherent latency of the CX4 transceivers. The inherent latency measured was 85 nanoseconds. Engineers subtracted this from the measured latency to calculate the actual latency by the DUTs. All reported latency values are the actual latency.

Price/Performance Comparison of BLADE and Cisco Switches Tested

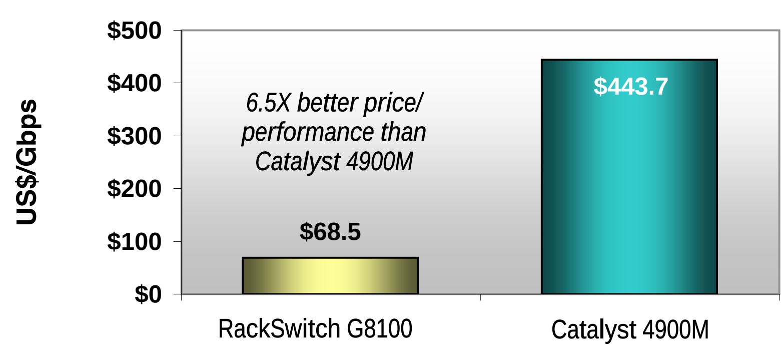

Notes: MSRP price for the RackSwitch G8100 tested was $16,450 while MSRP for the Cisco product tested was $52,980. Also, aggregate throughput measured for the RackSwitch G8100 and the Catalyst 4900M was 240 Gbps and 119.4 Gbps, respectively.

Source: The Tolly Group, September 2008 *Figure 4*

Layer 2 Throughput and Latency Tests

The tests were performed with test traffic consisting of one flow per port for frame sizes of 64, 128, 256, 512, 1,024, 1,280, 1,518 and 9,216 bytes, transmitted at the appropriate rate for each test scenario. A binary search algorithm was used to find the maximum throughput. An acceptable frame loss was set at ≤ 0.001%. Each iteration was run for 60 seconds. The same test was run three times and the results averaged. All ports on the DUT were placed in a single VLAN. The DUTs remained in their default configuration throughout the tests except that the Spanning Tree Protocol was disabled.

For the throughput test, 24 10GbE ports transmitted test traffic in a full-mesh pattern.

For the latency test, 20 10GbE CX4 ports were used and flowed the traffic in port-to-port traffic pattern (e.g. Port1 ⇔ Port 20, Port2 ⇔ Port 19, ...). Test traffic was generated at 45% of the theoretical maximum to accommodate lower throughput of the Catalyst 4900M. Since the RackSwitch G8100 is a cut-through switch, engineers used the LILO method to measure the latency. The Catalyst 4900M is a store-and-forward switch so the LIFO method was used to measure latency.

Layer 2 Multicast Forwarding

This set of tests used the same set of frame sizes of 64, 128, 256, 512, 1,024, 1,280, 1,518 and 9,216 bytes with unidirectional Layer 2 multicast forwarding with IGMP disabled on the DUTs. Engineers configured the test tool to generate the multicast traffic from Port1 to every other ports (e.g. Port1 ⇒ Port2, Port1 ⇒ Port3, ...)

Power Consumption

Engineers measured the power consumption of both products in idle state with all available ports connected, with 50% traffic load and with 100% traffic load. Engineers placed two "Watts up? PRO" power measuring tools between the wall outlets and two AC power supplies in the DUT and started the tool to log the reading (in watts) every second. During the course of testing, engineers used the Spirent TestCenter to generate the traffic loads for 50% and 100% of the theoretical maximum throughput. The test was based upon a 24-port full-mesh configuration with 64-byte frames.

Engineers used the 2006/2007 national average retail price of commercial electricity, from the Energy Information Administration of the Official Energy Statistics from the U.S. Government (US$0.0946 cents per kilowatt hour) to calculate the cost of operating a switch over a year.

Test Results Summary

Packet size (bytes)	Layer 2 throughput (% of line rate)		Average Layer 2 latency (24 10GbE ports in full mesh) in microseconds		Layer 2 multicast throughput (% of theoretical maximum)	
	RackSwitch G8100	Catalyst 4900M	RackSwitch G8100	Catalyst 4900M	RackSwitch G8100	Catalyst 4900M
64	100	49.71	0.311	3.32	100	50
128	100	49.67	0.339	3.44	100	50
256	100	49.73	0.377	3.69	100	50
512	100	49.72	0.374	4.2	100	50
1,024	100	49.72	0.374	5.17	100	50
1,280	100	49.71	0.379	5.69	100	50
1,518	100	49.73	0.382	6.14	100	50
9,216	100	49.73	0.374	21.05	100	50
Avg.	**100**	**49.71**	**0.36**	**6.59**	**100**	**50**

Source: The Tolly Group, September 2008

Figure 5

Test Bed Diagram

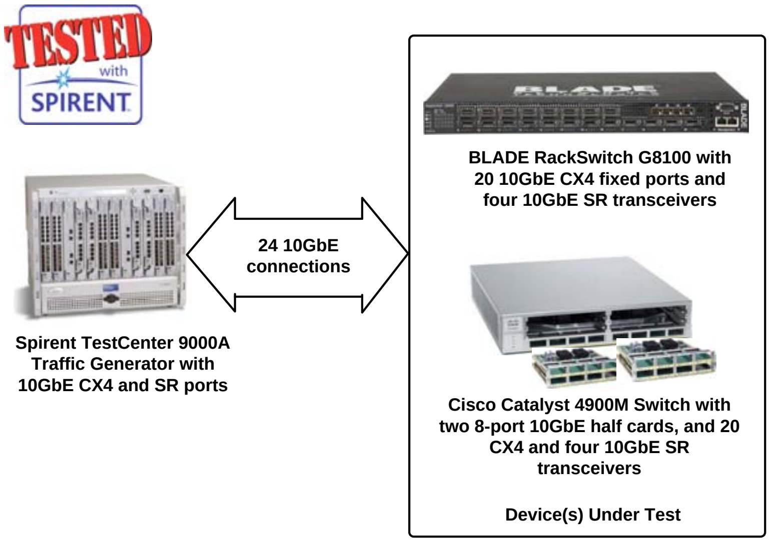

Source: The Tolly Group, September 2008 *Figure 6*

The Tolly Group is a leading global provider of third-party validation services for vendors of IT products, components and services.

The company is based in Boca Raton, FL and can be reached by phone at (561) 391-5610, or via the Internet at:

Web: http://www.tolly.com,
E-mail: sales@tolly.com

Fair Testing Charter™

Interaction with Competitors

The Tolly Group invited representatives from Cisco Systems to participate in the test as per The Tolly Group's Fair Testing Charter. Representatives from Cisco did not respond to the invitation.

BLADE Network Technologies supplied the products under test to The Tolly Group; Cisco products were acquired through normal distribution channels. Base system with default configurations (except that Spanning Tree Protocol was disabled) were used for all tests.

208346-cbdi1-cdb-08OCT08