TEST REPORT

Tolly.

#209115
May 2009

Commissioned by BLADE Network Technologies, Inc.

BLADE Network Technologies RackSwitch G8100 & G8124: Competitive Performance Evaluation versus Cisco Catalyst 4900M Switch

EXECUTIVE SUMMARY

BLADE's RackSwitch G8100 and RackSwitch G8124 outperform the comparable Cisco solution by demonstrating 2X higher throughput, 8.2X lower latency, 2.9X greater energy efficiency and 5.8X better price/performance on average than the Cisco Catalyst 4900M.

...<continued on next page>

BLADE RackSwitch G8100 & G8124 vs. Cisco Catalyst 4900M Layer 2 Latency over 20 10GbE Ports in a Port-to-Port Scenario
as reported by Spirent TestCenter 2.32

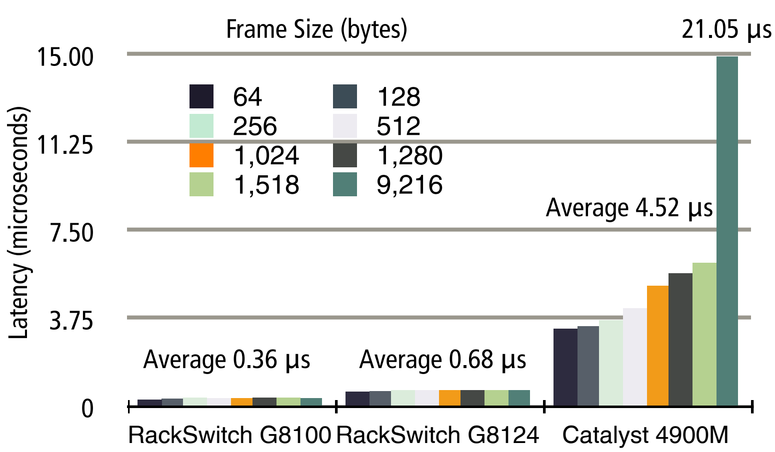

Note: 20 10GbE CX4 ports were used for Catalyst 4900M and RackSwitch G8100, while 24 10GbE SR ports were used for RackSwitch G8124. Latency was measured at 45% of theoretical maximum, due to Cisco's throughput limitations. Jumbo Frames not included in average.

Source: Tolly, March 2009

Figure 1

THE BOTTOM LINE

1. Demonstrates an average of 8.2X less latency than a comparable Cisco solution for various frame sizes in a port-to-port configuration
2. 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
3. Consumes an overall 65% less energy than the Catalyst 4900M, equivalent to a average savings of $198.37 per year, per switch deployed
4. Exhibits an average of 5.35X better price/performance when compared to 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 tools 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.

Tolly's hands-on evaluation of BLADE's RackSwitch G8100 and G8124, both based on a FocalPoint FM4224 low-latency 10GbE switch chip from Fulcrum Microsystems, demonstrates that the devices offer non-blocking 10GbE switching, far lower latency and energy consumption than the Catalyst 4900M. These are very important measurements 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 lossless I/O transactions over Ethernet such as Fibre Channel 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.)

The throughput test proves that both the RackSwitch G8100 and G8124 sustain 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 was only able to handle half the amount of traffic that the RackSwitch G8100 and G8124 processed.

The performance differences between these solutions become even more apparent in the latency test. While the BLADE RackSwitch G8100 and G8124 offer under 0.4 and 0.7 microsecond (μs) packet switching, depending on the frame size, tested from 64 bytes to 9,216 bytes, respectively. The Catalyst 4900M showed an average of 8.2X higher latency, ranging from 3 to 21 microseconds.

The power consumption test proves that data center users can reduce electrical costs by up to 75% when compared to the Catalyst 4900M under normal conditions.

Another feature that makes the BLADE RackSwitches more suitable for the data center rack environment is its cooling feature. The RackSwitch 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 is a viable solution for data center users who are looking 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.

BLADE Network Technologies

RackSwitch G8100 & G8124

Throughput, Latency, Multicast Forwarding & Energy Consumption

Tolly. Certified

Tested March 2009

Results

Layer 2 Throughput

Both the BLADE RackSwitch G8100 and G8124 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. The BLADE RackSwitches 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, apparently, to the low-capacity switch fabric from the eight-port 10GbE half-cards.

Layer 2 Latency

Tolly engineers measured the latency of the BLADE RackSwitches and the Cisco Catalyst 4900M. BLADE's Layer 2 performance advantage was even more apparent in these tests.

The RackSwitch G8100 achieved an average 0.36 µs of latency, ranging from 0.31 µs to 0.38 µs for 64-byte and 1,518-byte frames, respectively. This represents a 10X to 16X lower latency than the comparable Cisco solution, while the RackSwitch G8124's latency varied from 0.63 µs to 0.7 µs across the same set of frame sizes, a 5X to 8.7X lower latency than the Cisco solution.

The Catalyst 4900M measured from 3.3 µs to 6.1 µs for 64-byte to 1,518-byte frames, respectively. While 9,216-byte Jumbo Frames were tested, they were not factored into average latency since they represent a small percentage of traffic in a production network.

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 and 9,216 bytes.

Layer 2 Throughput Comparison of BLADE RackSwitch G8100 &G8124 vs. Cisco Catalyst 4900M across 24 10GbE Ports in a Full-Mesh Configuration

as reported by Spirent TestCenter 2.32

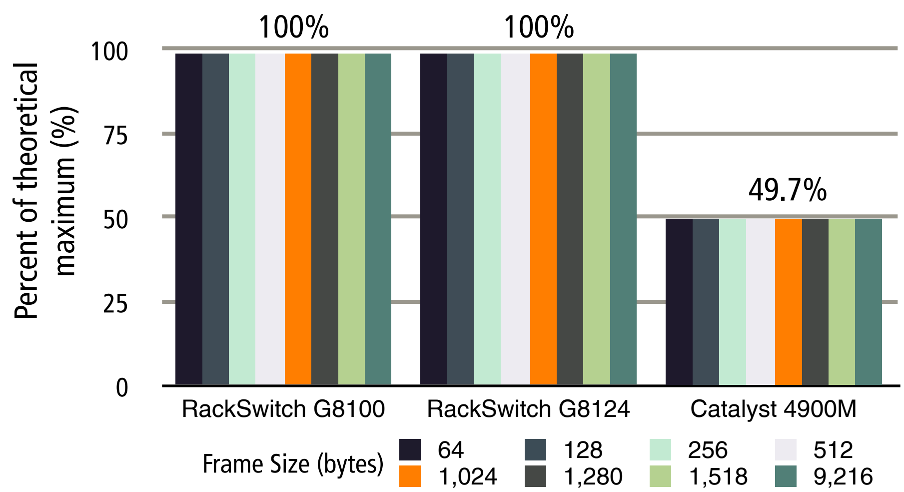

Note: 20 10GbE CX4 ports were used for Catalyst 4900M and RackSwitch G8100, while 24 10GbE SR ports were used for RackSwitch G8124.

Source: Tolly, March 2009

Figure 2

Similarly, with the throughput test, both BLADE RackSwitches achieved 100% throughput for all frame sizes However, the Catalyst 4900M was only able to handle half the traffic processed by BLADE.

Power Consumption

Tolly engineers found that the RackSwitch G8100 and G8124 used 83.4 and 147 watts, respectively, versus 351.9 watts for the Catalyst 4900M in an idle state scenario with all ports active. This equates to US $222.50 savings for the G8100 and a $169.72 savings for the G8124 on a yearly basis. This power consumption does not include the separate power usage to cool the equipment, though this heat dissipation is normally proportional to the power consumption.

Even when test traffic passed through the switches, the efficiency gap remained the same.

With a 50% 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.

For the RackSwitch G8124, the power consumption while operating at 50% of the maximum theoretical throughput was 165.8 watts (1.38

Avg. Power Consumption of BLADE RackSwitch G8100 & G8124 vs. Cisco Catalyst 4900M and Projected One-Year Operational Costs with Different Traffic Loads

as reported by "Watts up? PRO"

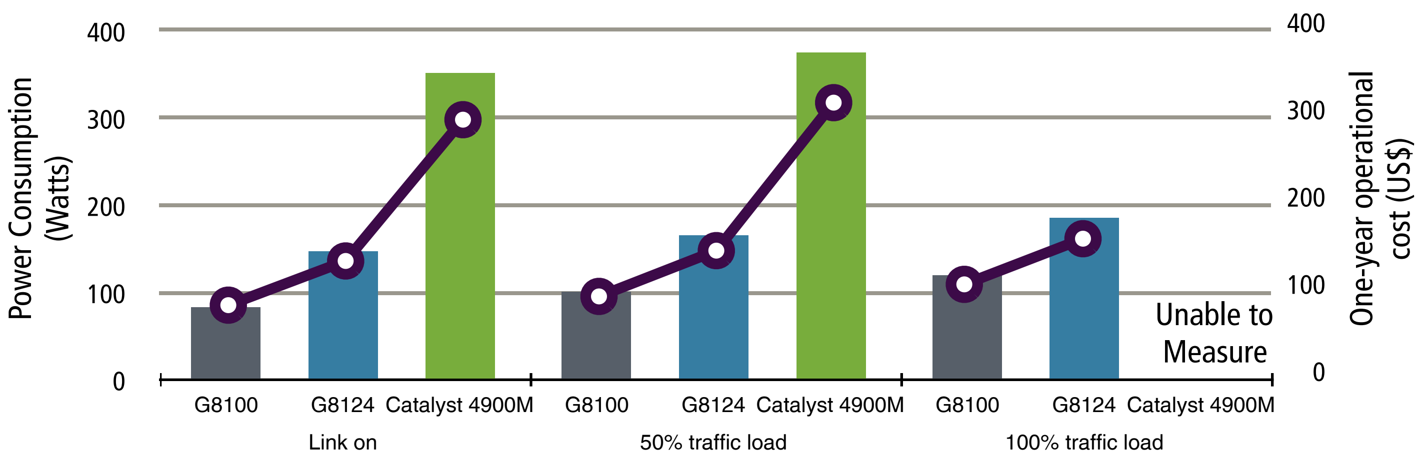

Note: Bar graph denotes power consumption in Watts, line graph denoted one-year operational cost.
20 10GbE CX4 ports were used for Catalyst 4900M and RackSwitch G8100, while 24 10GbE SR ports were used for RackSwitch G8124.
Power consumption of Catalyst 4900M was not measured at 100% because the switch cannot support more than 50% of the theoretical maximum throughput. Cost projection based on the 2006/2007 Average Commercial Electric Price of $0.0946 per kilowatt-hour. Power consumption estimates in this table do not include energy expended by cooling systems.
Test traffic consisted of 64-byte frames.

Source: Tolly, March 2009 Figure 3

watts per Gbps), 55.8% less energy than the Catalyst 4900M, generating $173.61 in savings over the course of a year.

With 100% traffic load, the RackSwitch G8100 consumed 119.7 watts. This represents .50 watts per Gbps and a single-year operational cost of just over US$99, while the RackSwitch G8124 topped out at 185.6 watts (.77 watts per Gbps), translating to an operational cost of $153.81 a year. The Catalyst 4900M was not included in this test because it could not handle more than 50% of the theoretical maximum throughput.

Price/Performance Comparison of BLADE RackSwitch G8100, RackSwitch G8124, and Cisco Catalyst 4900M

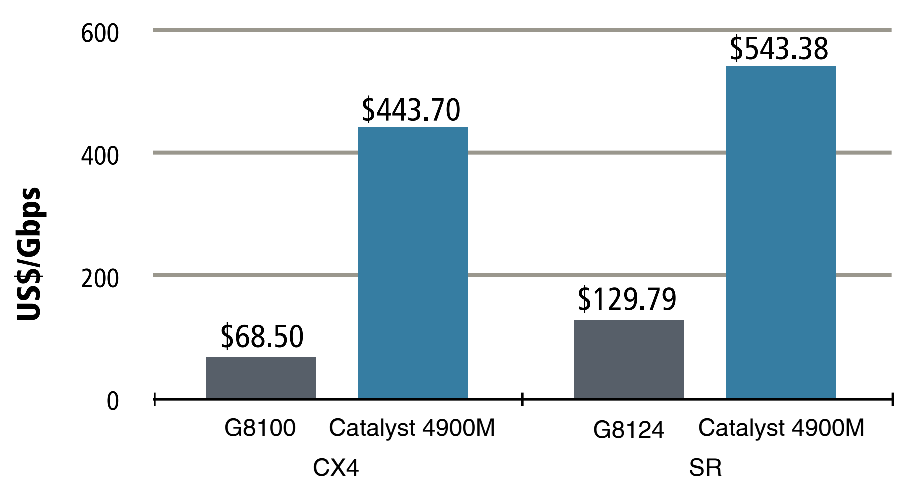

Note: 20 10GbE CX4 ports were used for Catalyst 4900M and RackSwitch G8100 comparison, 24 10GbE SR ports were used for Catalyst 4900M and RackSwitch G8124 comparison.
Test traffic consisted of 64-byte frames.

Source: Tolly, March 2009 Figure 4

Price/Performance Comparison

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

For this comparison, engineers configured the DUTs as similarly as possible. All base systems were outfitted with two power supplies. 20 10GbE CX4 and four 10GbE SR ports were used on the RackSwitch G8100 and Catalyst 4900M for one price comparison, while 24 10GbE SR ports were used on both DUTs when comparing the RackSwitch G8124 and Catalyst 4900M.

The prices of the RackSwitch G8100 and G8124 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

Tolly Common RFP™

The test methodology used for this report relies upon test procedures, metrics and documentation practices as defined by Tolly Common RFP, #1080 LAN Switch Power Consumption. V1.0.

To learn more about Tolly Common RFPs, go to:

http://CommonRFP.com

Test Results Summary

Packet size (bytes)	Layer 2 throughput (% of line rate)			Average Layer 2 port -to-port latency in microseconds (24 10GbE ports)			Layer 2 multicast throughput (% of theoretical maximum)		
	RackSwitch G8100	RackSwitch G8124	Catalyst 4900M	RackSwitch G8100	RackSwitch G8124	Catalyst 4900M	RackSwitch G8100	RackSwitch G8124	Catalyst 4900M
64	100	100	49.71	0.311	0.63	3.32	100	100	50
128	100	100	49.67	0.339	0.65	3.44	100	100	50
256	100	100	49.73	0.377	0.7	3.69	100	100	50
512	100	100	49.72	0.374	0.7	4.2	100	100	50
1024	100	100	49.72	0.374	0.69	5.17	100	100	50
1280	100	100	49.71	0.379	0.7	5.69	100	100	50
1518	100	100	49.73	0.382	0.7	6.14	100	100	50
9216	100	100	49.73	0.374	0.7	21.05	100	100	50
Avg.	**100**	**100**	**49.71**	**0.36**	**0.68**	**4.52**	**100**	**100**	**50**

Source: Tolly, March 2009

Figure 5

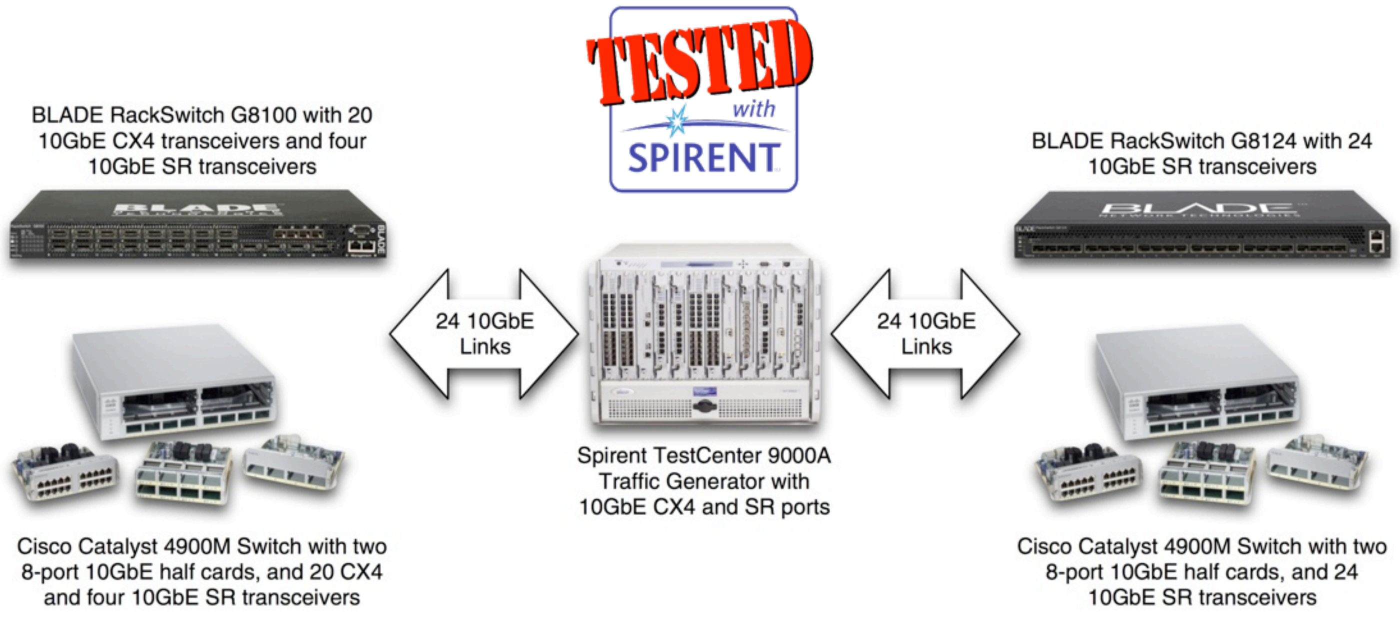

Note: Two test beds were used to compare both BLADE switches to the Cisco Catalyst 4900M. The RackSwitch G8100 has 20 built-in CX4 ports, while the RackSwitch G8124 is equipped with 24 SFP+ ports. Cisco Catalyst 4900M was custom configured to compare to the BLADE DUT.

Source: Tolly, March 2009

Figure 6

of $68.54 per Gbps. The Catalyst 4900M as tested cost $443.72 per Gbps. The RackSwitch G8124 had a cost per Gbps of $129.79, 4.2X better than the Catalyst 4900M, which cost $543.38 per Gbps when configured with 10GbE SR ports.

Test Setup & Methodology

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

The DUTs were BLADE Network Technologies RackSwitch G8100 and G8124 (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. For this testing, the switches were equipped with either 20 10GbE CX4 and four 10GbE SR ports (G1800 vs 4900M), or 24 10GbE SR transceivers (G8124 vs 4900M), as well as two AC power supplies. BLADE RackSwitches provide 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 RackSwitch G8124 was outfitted with a base system (BN-8124F-BDL) with 24 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 2.32.

The Spirent TestCenter 9000A solution, using MSA-2001A modules, has been tested back-to-back to measure the inherited latency of the both the CX4 and SFP+ transceivers. The inherent latency was 85 nanoseconds for the CX4 transceivers, and 60 nanoseconds for the SFP+ transceivers. Engineers subtracted this from all the measured latency to calculate the actual latency of all DUTs.

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. Each of the three iterations was run for sixty seconds 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 and flow control were disabled.

For the throughput test, 24 10GbE ports transmitted test traffic in a full-mesh pattern. Test traffic was generated at 45% of the theoretical maximum to accommodate the lower throughput of the Catalyst 4900M.

For the latency test, engineers connected the 24 ports on the DUT to the test tool in a port-to-port configuration. Since the RackSwitch G8100 and G8124 are cut-through switches, 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 with unidirectional Layer 2 multicast forwarding with IGMP disabled on the DUTs. Engineers configured the test tool to generate the multicast traffic from a single port to all other ports.

Power Consumption

Tolly engineers measured the power consumption of both products in an idle state, with all available ports connected, with 50% traffic load, and with 100% traffic load, consisting entirely of 64-byte frames. 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 readings 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. This 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 per kilowatthour) to calculate the cost of operating a switch over a year.

Designed Like a Blade Switch

BLADE's RackSwitch G8100 and G8124 are not just any other 10-Gigabit Ethernet switches; their design is a direct result of the company's heritage as the leading manufacturer of blade server switches.

In addition to their extremely low latency and line-rate throughput, the switches' design was perfected by engineers who are accustomed to designing networking equipment that meets the strict power and space constraints of blade server systems. RackSwitch G8100 and G8124 were designed from the ground up with the most efficient silicon, power management subsystem and printed circuit board design. BLADE's RackSwitch G8100 and G8124 also employ airflow efficiencies with the data center environment in mind.

Both G8100 and G8124 can be ordered with either a front-to-rear or rear-to-front airflow, so customers can match the airflow of the servers in their racks to maintain separation of hot aisles and cool aisles in their data center. BLADE's RackSwitch G8100 and G8124 use the same operating system as the company's blade server switches for IBM, HP and NEC. The G8100 and G8124 are 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
March 2009

About Tolly…

The Tolly Group companies have been delivering world-class IT services for 20 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

You can reach the company via E-mail at sales@tolly.com, or via telephone at +1 561.391.5610.

Visit Tolly on the Internet at:
http://www.tolly.com

Interaction with Competitors

Tolly personnel invited representatives from Cisco Systems, Inc. to participate in the test as per the Tolly Fair Testing Charter. Representatives from Cisco did not respond to the invitation.

For more information on the Tolly Fair Testing Charter, visit:
http://www.tolly.com/FTC.aspx

209115-cb3-jft-11MAY09-verJ