Performance & Interoperability
Dell Networking
7000 and 8100 Switch Series

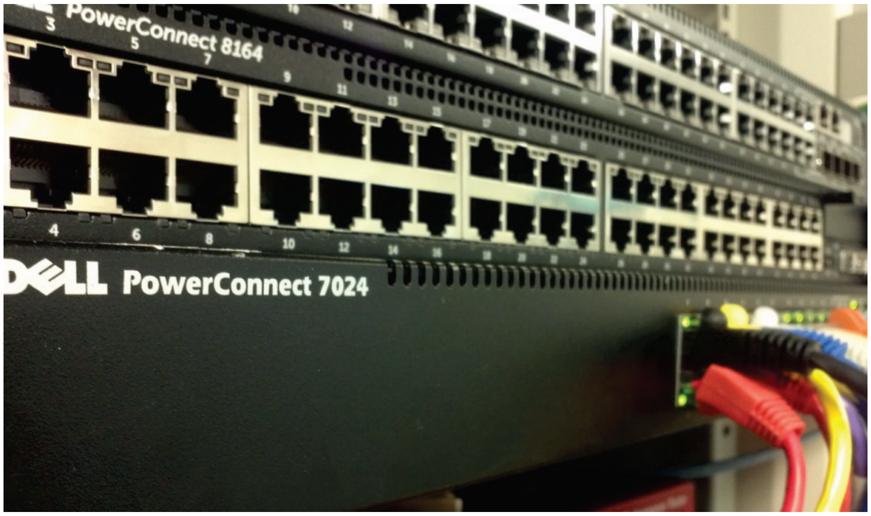

Miercom Lab Testing Report

April 2013

Report 130301

Contents

1.0 Executive Summary

Dell engaged Miercom to verify the performance and interoperability of the Dell Networking 7000 and 8100 switches. Power efficiency, switch features and resiliency testing was also included in this evaluation. The test results are discussed with detail in this report. Overall, we were impressed with the Dell Networking 7000 and 8100 series switches. These switches are designed for use in an enterprise campus network, either in a standalone or stackable configuration.

The Dell Networking switches are operational out-of-the-box and do not require any additional configuration. Switch configuration for testing was quick and straight-forward when using the CLI. Administrators that are familiar with the Cisco CLI, would find that the Dell switches support a similar CLI design. There is also a Web interface available, which is intuitive and allows for easy access to configuration parameters and settings.

The Dell Networking 7024 and 8164 both achieved 100% line rate throughput with zero loss, without compromising latency. Both the 7024 and 8164 switches saved power when Energy Efficient Ethernet (EEE) was enabled. There was a reduction of 89 watts for the 8164 and 8.1 watts for the 7024 when EEE was enabled, when compared to power used without EEE.

Resiliency was examined by disconnecting cables when in a stacking environment. The switches failed over and the network remained operational. Spanning Tree was tested by removing a forwarding STP link. There was sub-zero convergence with traffic being forwarded at 100% line rate.

The security assessment revealed that the Dell Networking switches are capable of effectively blocking DoS attacks; ensuring legitimate traffic is not dropped and the GUI remains operational.

Additional tests performed showed that the Dell Networking switches had excellent performance, scalability, resiliency and interoperability in campus network environments.

Key Findings

- Saved 89 watts when Energy Efficient Ethernet (EEE) was enabled
- Forwarded 100% line rate throughput, while maintaining low latency values
- Dell Networking switches had a better overall five year TCO when specifically compared to HP 8206 zl, HP 3800G PoE+, and Cisco Catalyst WS-C4500X/WS-C3750X switches
- Stacking proved to be resilient to network failures
- DoS attacks were dropped by the switch, protecting endpoint devices
- Proved to be interoperable with Cisco Catalyst 3750x and 4500x switches
- Web interface is intuitive

Overall, the Dell Networking 7024 and 8164 switches are easy to configure, can achieve 100% line rate throughput with zero loss and low latency, save power consumption by properly utilizing EEE, and are easily interoperable with Cisco Catalyst 3750x and 4500x switches.

Robert Smithers
CEO
Miercom

2.0 Methodology

Testing was conducted on the Dell Networking 7024 and 8164 switches, running firmware version 5.0.1.3. Each switch was tested as a standalone unit, as well as in a stacking configuration. Standalone tests included RFC 2544, 2889 and 3918 throughput and latency on Layers 2 and 3 with IPv4 and IPv6 modes with bidirectional traffic. Various frame sizes ranging from 64-bytes to 9612-bytes were used.

Latency was tested using the Ixia's built in Store and Forward parameter. RFC 2889 was used to verify the cross processor throughput, to measure CPU utilization, and to verify the packet buffer size of the switches. Maximum multicast throughput without any frame loss was verified by using RFC 3918.

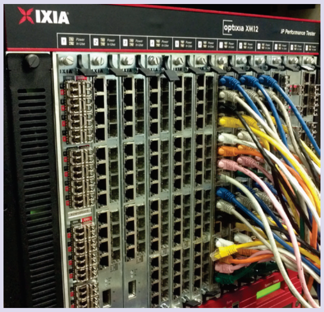

Ixia XM12 Chassis

Ixia (www.ixiacom.com) is an industry leader in performance testing of networking equipment. Ixia's unique approach utilizes coordination of energy measurements with network traffic load – allowing energy consumption to be graphed against network traffic volume. Real-world traffic is generated by Ixia's test platform and test applications, principally IxAutomate for Layer 2-3 switching and routing traffic.

Stacking tests included resiliency tests and failover. For failover, sending traffic through the stack links and removing a link provided the convergence time of the stack architecture. Resiliency was tested by verifying a stack slave can instantly become the stack master, once the master switch is removed from the stack.

Additionally, we tested route capacity, MAC address cache size, power efficiency, link aggregation throughput, QoS, multiple STP modes, and interoperability with Cisco Catalyst 3750X and Cisco Catalyst 4500X switches in various network configurations.

The Ixia XM12 chassis was used with the Ixia IxNetwork application to drive network traffic through the switches using various test methodologies. Ixia (www.ixiacom.com) is an industry leader in performance testing of networking equipment. Ixia's exclusive approach employs coordination of energy measurements with network traffic load, allowing energy consumption to be charted against network traffic volume. Real-world traffic is generated by Ixia's test platform and test applications, principally IxAutomate for Layer 2-3 switching and routing traffic.

The security assessment was conducted using Spirent Studio Security (www.spirent.com). Spirent Studio Security provides a complete service assurance solution for determining the reliability, availability and security of IP-based applications and services. The Spirent solution is highly automated, with lights-out fault isolation. To aid in the remediation of software flaws, Spirent Studio Security provides actionable reports and complete data on any faults. Spirent-based testing is managed via a variety of interfaces, including a highly visual Web-based graphical user interface. Testing can also be remotely controlled using REST- or XML-based APIs for integration into laboratory automation frameworks such as HPQC or STAF.

2.1 Test Bed Diagram

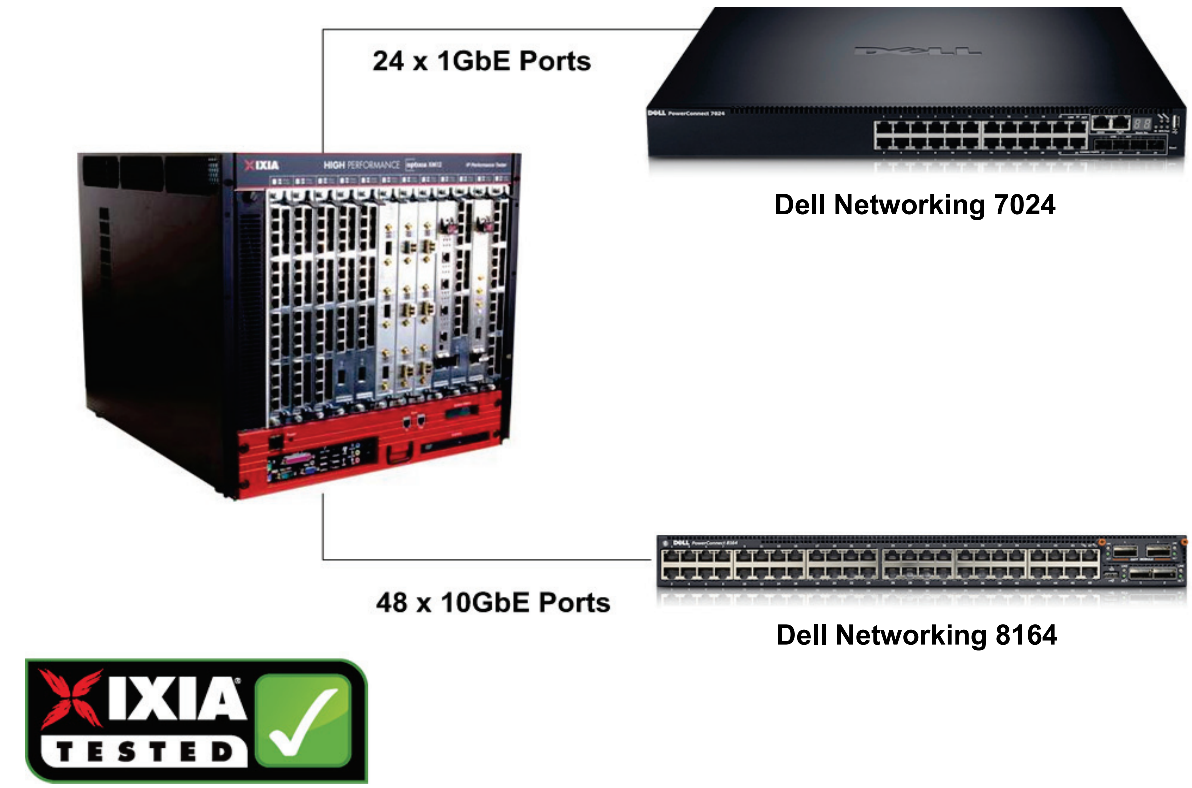

Source: Miercom, April 2013

2.2 RFC Methodologies

RFC 2544 Throughput and Latency

The RFC 2544 throughput test determines the maximum rate at which the DUT receives and forwards frames without frame loss. Frames are initially sent at a specified rate and a binary search algorithm is used to obtain a rate at which the DUT does not lose frames. Frames can be MAC only, IPv4, IPv6 (with or without Extension Headers) or an IPv4/IPv6 mixture. Additionally, latency is calculated by subtracting the transmit timestamp from the receive timestamp. Results include: throughput rates in frames per second obtained for each frame size and latencies for each frame size, and the average latencies for all trials, with high granularity and precision.

Methodology

The test load generator will forward and receive traffic to and from each directly connected port on the switch. Frames are initially sent at a user-specified rate, generally the maximum theoretical rate based on the speed of the port. A binary search algorithm is used to obtain a rate at which the DUT does not lose frames. This test is configured with a one-to-one traffic mapping.

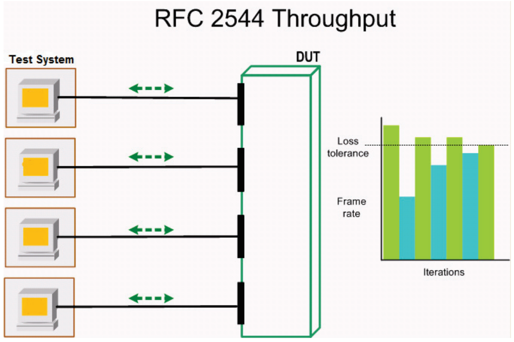

Results

The results will show the maximum throughput the switch is able to achieve without any frame loss. Additionally, latency values will be captured for each frame size tested.

RFC 2889 Fully Meshed

Objective

This test determines the maximum rate at which the DUT can handle frames while receiving frames on all its ports.

Methodology

Simultaneously, all ports transmit and receive traffic at a specified transmission rate. Each port on the DUT transmits frames to all other ports in an evenly distributed round-robin manner, and receives frames from all the ports. The test uses a binary, step, or combo search algorithm to determine the maximum no drop rate.

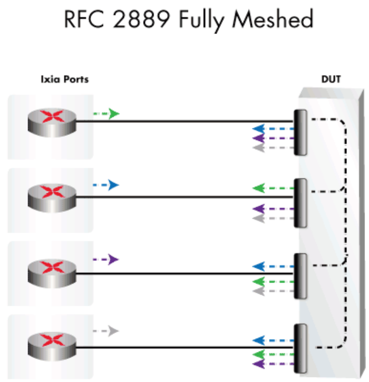

Results

Total number of frames transmitted from all ports, total number of frames received on all ports, percentage of lost frames obtained for each frame size, latency, jitter, sequence errors and data integrity errors.

RFC 3918 Throughput No Drop Rate

Objective

This test determines the maximum throughput the DUT can support while receiving and transmitting multicast traffic.

Methodology

This test calculates the maximum DUT throughput for IP Multicast traffic using either a binary or a linear search, and to collect latency and data integrity statistics. The test is patterned after the ATSS Throughput test; using multicast traffic. A one-to-many traffic mapping is used, with a minimum of two ports required. The IGMPv2 protocol was used in this test.

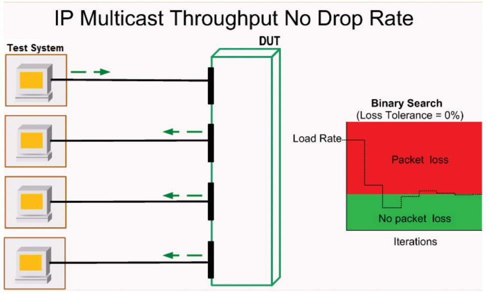

Results

Results will show the maximum throughput per port and frame loss per multicast group.

3.0 Performance and Scalability

Performance and scalability tests focused on throughput, latency, microbursts, fully meshed throughput, stacking, spanning tree, route capacity, and MAC address cache size. Testing was based on RFC 2544, 2889 and 3918 benchmark tests.

Overall, we found that configuring the switches for our performance and scalability tests was straightforward. Layer 2 traffic passed through the switch without any configurations made. To transmit Layer 3 traffic, VLANs were configured with an IP address associated with each VLAN. This is typically done to perform Layer 3 RFC 2544 throughput tests. Configuring VLANs and IP addresses was simple and did not take much time. Dell's stacking deployment was as simple as connecting stacking cables between the switches, and forwarding traffic between the stack in just a few minutes. The Dell's CLI and GUI were both intuitive and allowed for easy configuration of the switches, reducing the amount of time to deploy the switches in a network.

3.1 Throughput

The Dell 7024 switch was tested with 24 ports of 1GbE. The Dell 8164 switch was tested using 2 configurations. The first configuration was done with 48 10GbE ports and a four port 10GbE module. The second configuration was 48 1GbE ports and a four port 10GbE module.

Figure 1: Dell 7024 Switch RFC 2544 Throughput

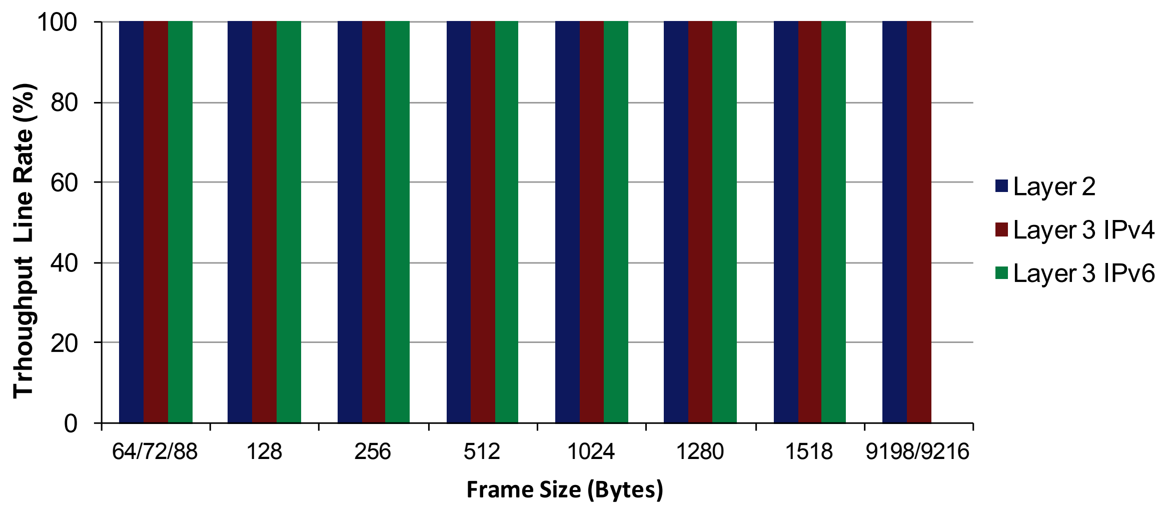

Source: Miercom, April 2013

Dell 7024 switch achieved 100% line rate for Layer 2 and Layer 3 (IPv4 and IPv6) with zero frame loss.

Figure 2: Dell Networking 8164 RFC 2544 Throughput

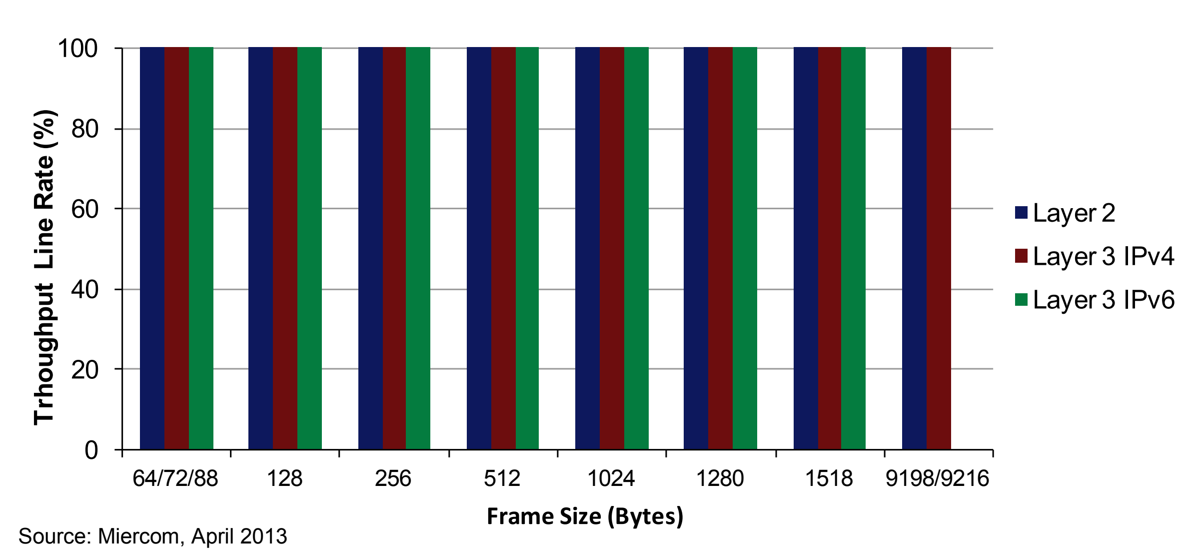

Dell 8164 switch achieved 100% line rate for Layer 2 and Layer 3 (IPv4 and IPv6) with zero frame loss.

The smallest frame size tested on the both the 7024 and 8164 switches were:
Layer 2 test was 64-bytes,
Layer 3 IPv4 was 72-bytes,
and Layer 3 IPv6 was 88-bytes.

The maximum frame size:
Layer 2 was 9216-bytes
9198-bytes for Layer 3 IPv4
Layer 3 IPv6 is 1518-bytes.

The Dell Networking 7024 and 8164 switches performed at 100% line rate without dropping a single packet for both Layer 2 and Layer 3 (IPv4 and IPv6).

3.2 Latency

Figure 3: Dell 7024 Switch RFC 2544 Latency

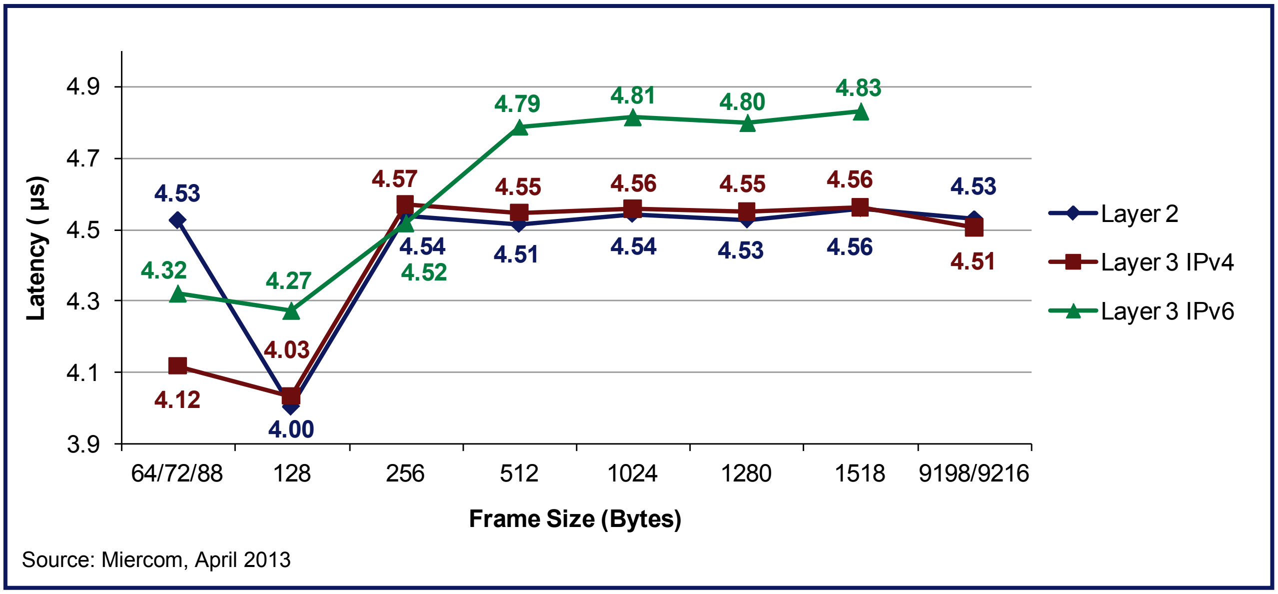

The Dell Networking 7024 maintained low latency values while forwarding 100% line rate traffic. The measured latency average ranged was:

4 µs to 4.5 µs for Layer 2 traffic,

4.0 µs to 4.5 µs for Layer 3 IPv4 traffic, and

4.3 µs to 4.8 µs for Layer 3 IPv6 traffic.

Figure 4: Dell Networking 8164 RFC 2544 Throughput

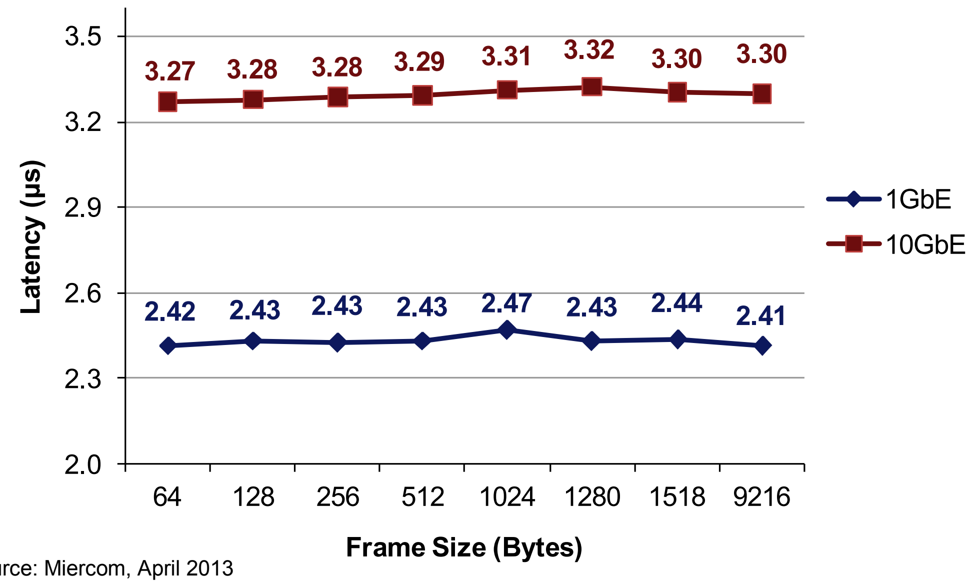

Source: Miercom, April 2013

With latency tested in store and forward mode, we found that low latency values were maintained for both 1GbE and 10GbE traffic with Layer 2 traffic on the Dell Networking 8164. IPv4 and IPv6 latency was also low, and close to the Layer 2 latency values.

3.3 Microbursts

Microbursts are sub-second bursts of traffic that can occur in networks. Microbursts could cause traffic loss if the switch is unable to buffer excess packets. It is important for switches to process bursty traffic efficiently, especially in networks where multiple devices are transmitting to the same receiver. We verified microburst sustainability using several port configurations and tested microbursts using a port count ranging from two to thirty. Our results were obtained from testing the Dell 8164 and 7024 switches.

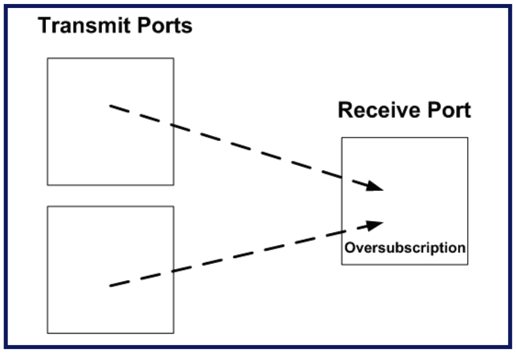

Oversubscription occurs when multiple ports send to one receiving port.

Figure 5: Packet Buffer Size

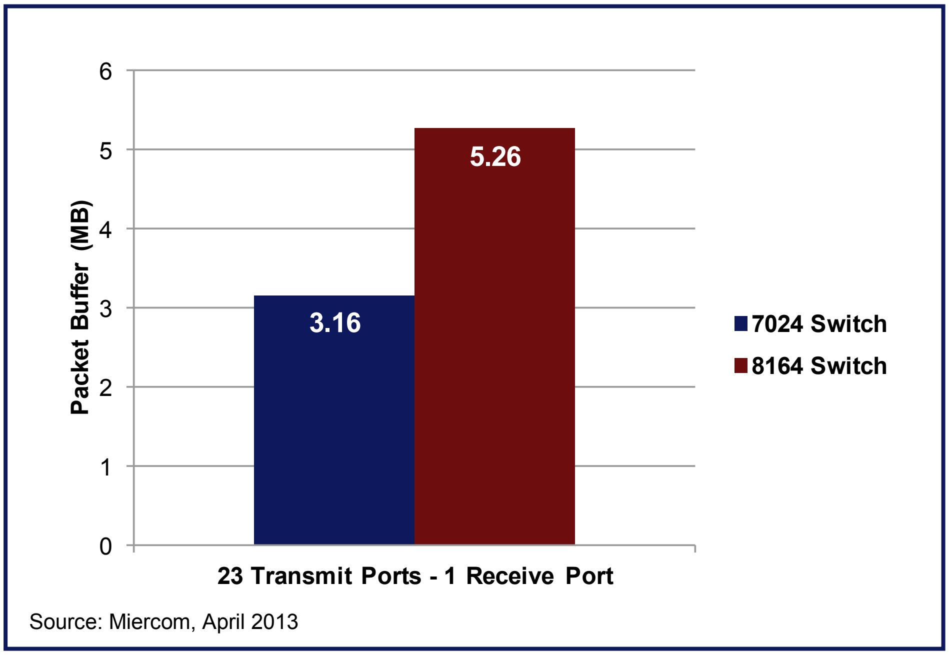

Maximum packet buffer size using 23 transmitting ports and 1 receive port.

From our testing, we found that the 7024 switch was capable of buffering 3.16MB on the receive port, and the 8164 switch was able to buffer 5.26MB on the receive port. If the packet buffer was exceeded, there was frame loss present.

3.4 Fully Meshed Throughput

The fully meshed throughput test is a many-to-many port mapping which validates the switches cross processor performance. This test stresses the switch more than the standard RFC 2544 one-to-one port mapping.

To perform the test, we used RFC 2889 Fully Meshed with all ports fully loaded. In addition to verifying throughput without frame loss, CPU utilization was monitored during the test.

Figure 6: Dell 7024 Switch RFC 2889 Throughput and CPU Utilization

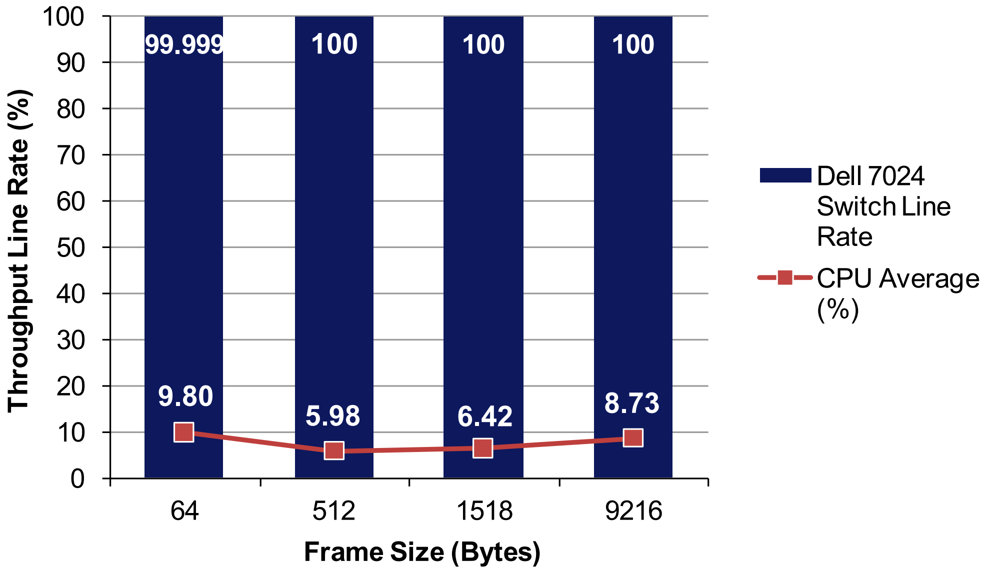

Source: Miercom, April 2013

Dell Networking 7024 switch maintained a low CPU utilization for all frame sizes.

The Dell Networking 7024 switch had a minimal frame loss of .001 when tested with 64-byte frames while maintaining an average CPU utilization of 9.80%. At the 64-byte frame size, the aggregate frame per second was 35,714,062.77. This equates to 29,233 frames lost out of the 6,428,571,432 frames that were transmitted.

All of the remaining frame sizes showed zero loss and maintained an average CPU utilization ranging from 5.98% to 8.73%.

Figure 7: Dell 8164 Switch RFC 2889 Throughput and CPU Utilization

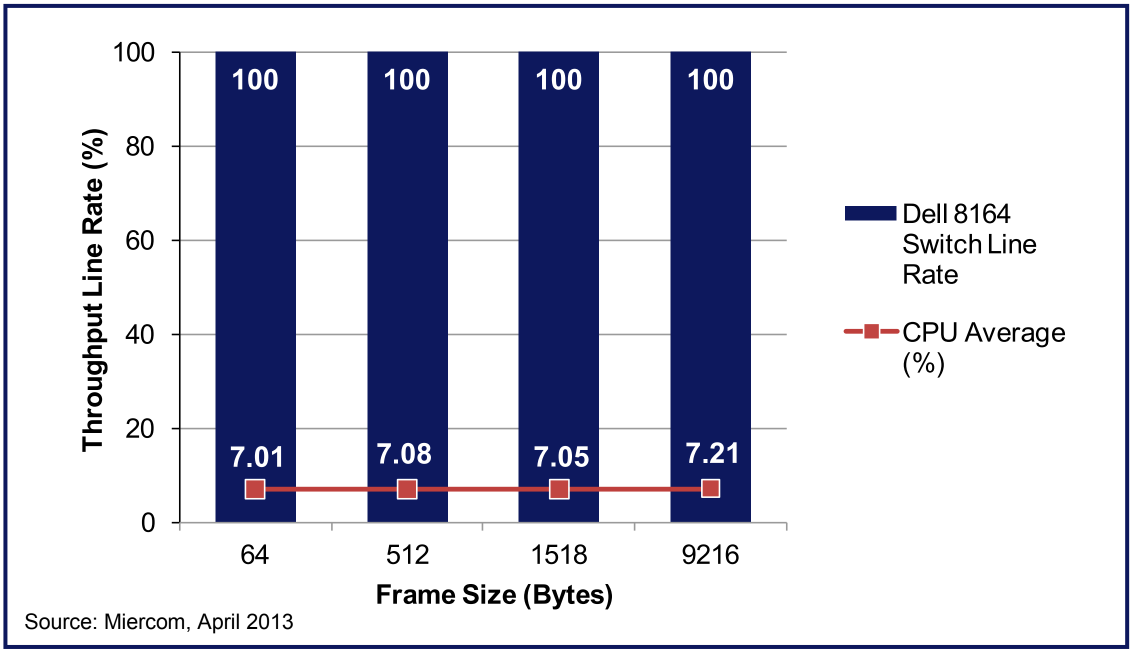

Dell Networking 8164 switch maintained low CPU utilization while forwarding line rate traffic with zero loss.

The 8164 switch achieved 100% line rate fully meshed traffic, with zero loss. During each frame size tested, the CPU utilization ranged from 7.01% and 7.21%.

3.5 Stacking

The Dell Networking 7000 series supports stack LED indicators on the front panel.

Stackable switches provide many benefits in networks today. This architecture allows for the management of two or more switches in the stack using one single IP address. Additionally, stacking provides port density, high-availability and resiliency. In our testing we focused on these key areas. When the switches are stacked, each unit supports up to 184 Gbps in switch capacity, with up to 2 Tbps of capacity in a single stack.

The Dell Networking 7000 series switches support stack LED indicators to help identify master and stack switch members, *see image to the right*. The CLI is accessible from only the master switch. If a console session is attempted on a stack member, the user is prompted to initiate the connection with the master switch. However, users can telnet to the master switch through a stack member to access the CLI. The web GUI can also be accessed through either the master or any stack member for management. Managing the stack is a simple task, with a single point of management from the stack master. Administrators can configure VLANs, ports, security and many other Layer 2 and Layer 3 features on stack members from one interface or CLI session. Within the GUI, there is an option to switch between the different switches in the stack for configuration using the drop down menu.

There is also an option within the GUI to configure the standby unit in stacking: To configure the standby unit via the Web interface, view the Stacking > Unit Configuration page and select Standby for the Unit Type.

Each bay can contain a CX-4, SFP+, or a 10GBase-T module. Stacking is supported only on CX-4 modules in either or both bays. CX-4 stacking modules (with two ports per module) can be configured in 16 Gbps stacking mode or 10 Gbps Ethernet uplink mode. CX-4 stacking modules default to stacking mode, where the maximum cable length is 3 meters.

Dell Networking 7024 Device View

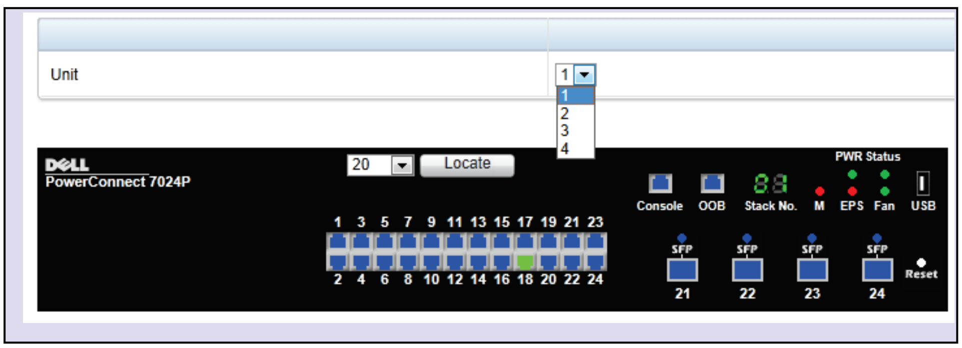

Device view shows the front panel status on unit one in the stack.

On the following page, the master switch stack summary feature shows details about each switch in the stack. For example, the master switch is displayed as the management switch (Mgmt Sw), while the standby switch is labeled as the configuration standby (Cfg Stby). Additionally, the unit number, model number, switch status and firmware versions are displayed.

Stack Summary: Detail

Unit	Management Status	Standby Status	Preconfigured Model Identifier	Plugged-in Model Identifier	Switch Status	Firmware Version	NSF Unit Support	SFS Status	SFS Last Attempt Status
1	Stack Mbr		PCT7024P	PCT7024P	OK	5.0.1.3	Enable	No Action	None
2	Stack Mbr		PCT7048	PCT7048	OK	5.0.1.3	Enable	No Action	None
3	Stack Mbr	Cfg Stby	PCT7024	PCT7024	OK	5.0.1.3	Enable	No Action	None
4	Mgmt Sw		PCT7024P	PCT7024P	OK	5.0.1.3	Enable	No Action	None

The Dell 7000 series switch stack is created automatically when the stack links are connected between each switch. The switch with the highest MAC address is assigned as the stack master. The master of the stack can also be manually configured. Installing the latest firmware version is only required on the stack master, which then automatically pushes the firmware to all other switches or a new switch when added.

We verified the stacks' resiliency by removing the master from the stack. The standby switch took over the role as the master, and maintained all configurations from the previous stack master.

We determined the convergence time when removing a stack link. To test convergence, traffic was sent through the stack at 1 Gbps from one switch to another, and then we removed a stack link. We performed three iterations to ensure our results were repeatable. There was zero frame loss when the blocked stack link was removed.

When the active stack link was removed, we saw an average convergence time of 1,956.66ms. There was no loss when plugging the stack link back in.

We also wanted to see how much loss there was when forcing the master offline, which results in the automatic standby to master switchover. There was a convergence time of 53.54ms.

Dell Networking Switch Stack Test Bed

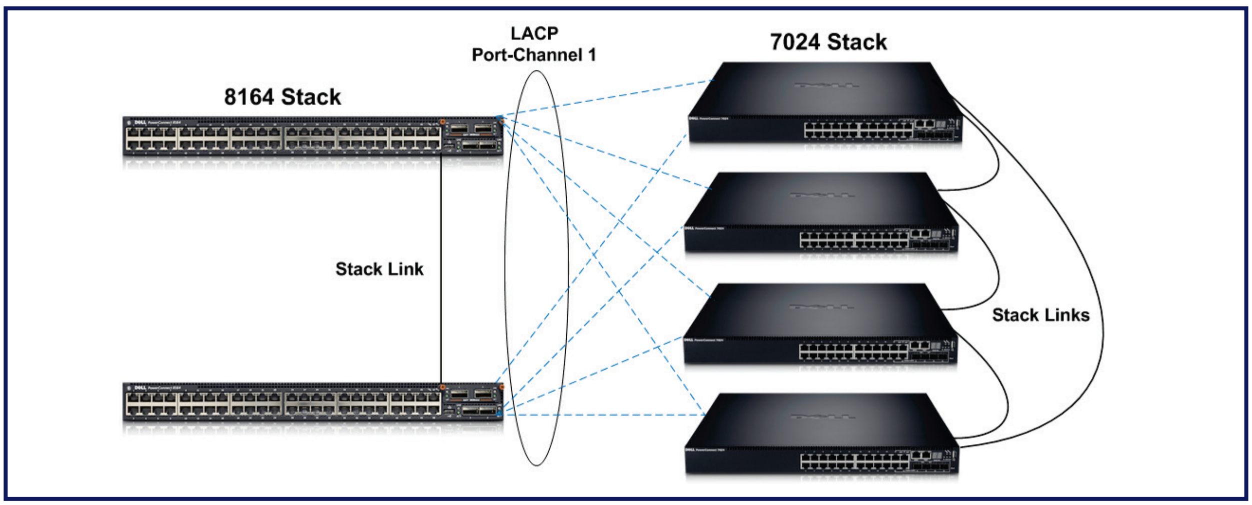

Test bed configuration used for testing stacking capabilities. The test bed included two 8100 series switches and four 7000 series switches.

Stacking capabilities were also verified on the Dell Networking 8164 switches. The 8164 switches were added to the 7000 series switch stack using eight ports in one port channel LAG. When sending traffic through the stack, the switches successfully forwarded traffic at 100% line rate with zero loss. Additionally, we verified 100% line rate traffic with zero loss, with four cables removed from the port channel.

3.6 Spanning Tree

The original Spanning Tree Protocol does not support VLAN-based blocking; however, newer STP protocols have been designed that support VLANs. RSTP, MSTP and other per VLAN Spanning Tree Protocols can set different block links for each VLAN. RSTP and MSTP can passively contribute to load balancing of network utilization because of the different data paths for each VLAN. Original Spanning Tree, 802.1d has a slower convergence time, whereas RSTP and MSTP have much faster convergence times.

The default spanning tree protocol for the Dell Networking 8100 and 7000 switches is RSTP 802.1w. RSTP is backwards compatible with 802.1d neighbors, allowing devices that only support 802.1d to interoperate.

Original Spanning Tree Protocol Test Bed

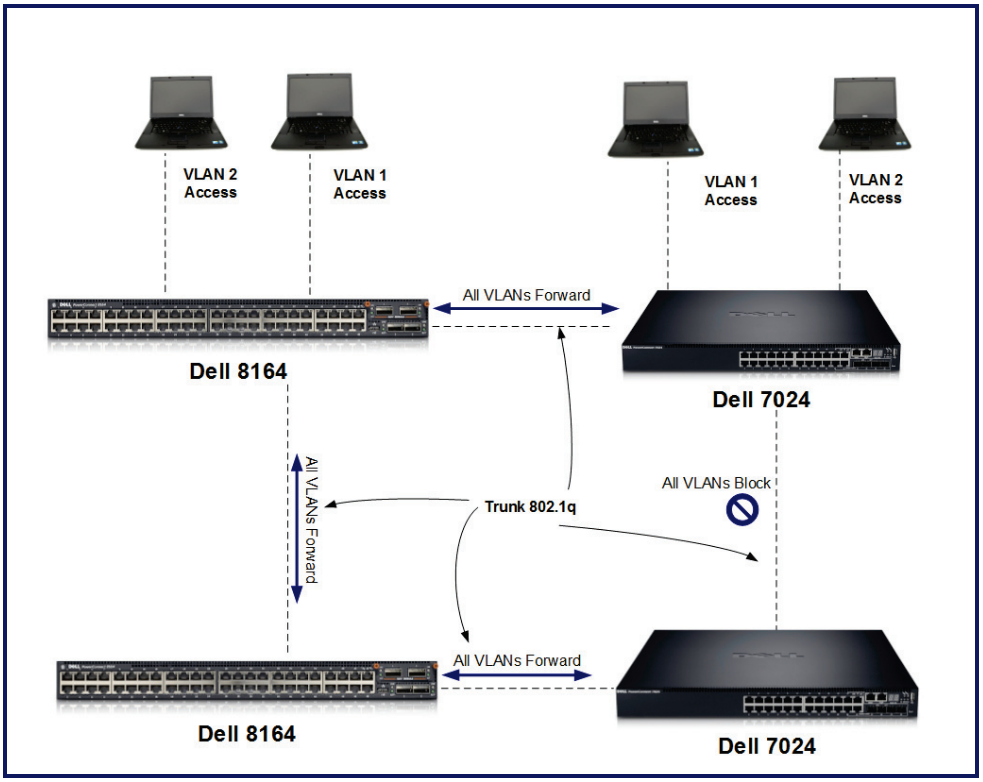

Our testing focused on validating proper STP function and convergence time. Convergence was tested by transmitting traffic between the switches, and disconnecting a link in the forwarding state. This caused a link in the blocked state to activate, requiring all network traffic to reroute to the new link in the forwarding state.

RSTP and MSTP Test Bed

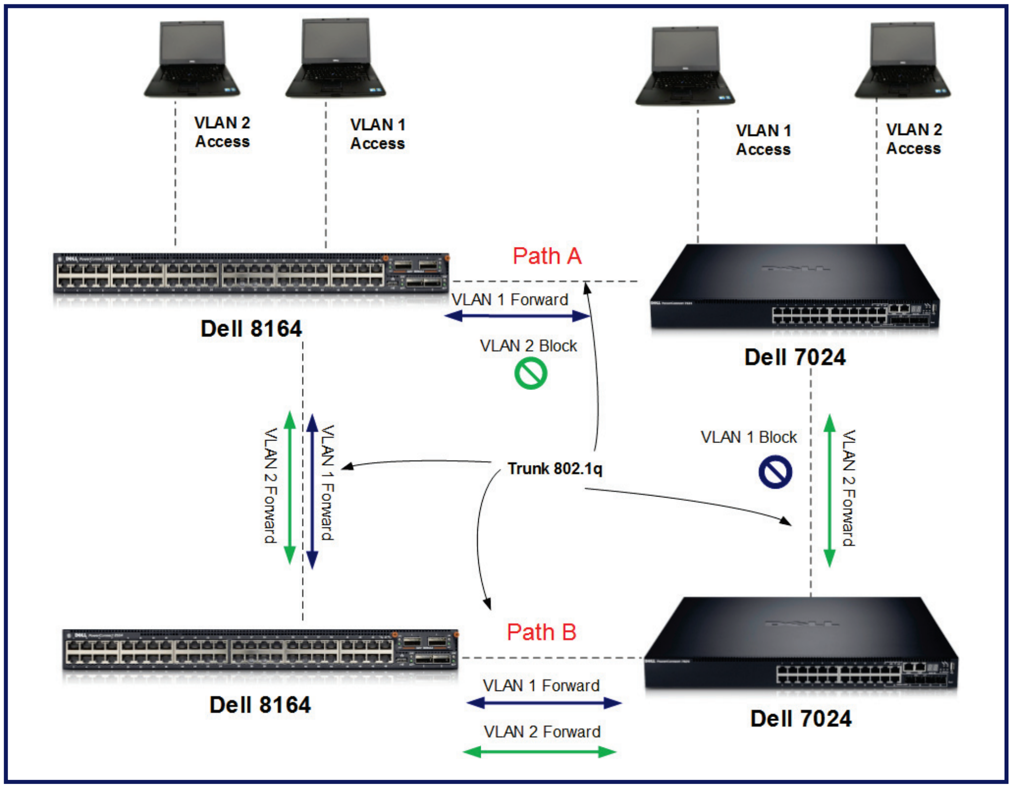

Results

Convergence Time for Spanning Tree Protocol

Spanning Tree Protocol	Average Convergence Time (Sec)
Original STP, 802.1d	32.53
RSTP, 802.1w	1.237
MSTP, 802.1s - without per VLAN STP	1.102
MSTP, 802.1s - with per VLAN STP (Average of all VLANs)	0.932

As expected, the original STP takes more time to converge traffic. Since RSTP can interact and exchange messages with neighbor switches, the average convergence time is faster than the original spanning tree protocol.

Link Utilization of STP and MSTP

	VLAN1 Forwarding Path	VLAN 2 Forwarding Path	Throughput Mbps	Network Utilization
Original STP 802.1d	Path A		1970	50%
MSTP	Path B	Path A	3908	100%
MSTP	Path A		1970	50%
MSTP	Path A	Path B	3908	100%
MSTP	Path B		1970	50%

Additionally through testing we confirmed that MSTP can utilize 100% of the redundant network without creating a loop or broadcast radiation. This is because MSTP can support multiple VLANs with different STP instances. Therefore, VLANs can create different traffic forwarding paths and blocking paths. MSTP can utilize 100% line rate traffic by load balancing traffic through different VLANs.

3.7 Link Aggregation

Link aggregation (LAG) involved linking several Dell Networking switches together and testing the performance using RFC 2544 to verify throughput. Two 8164 Dell Networking switches were linked together using two CAT6A cables to allow 10 Gbps traffic to transmit on each link. Each switch was configured with two 10GbE fiber uplinks to pass traffic between the two switches. With bidirectional traffic flowing through the LAG, we verified that the switches were capable of forwarding 100% line rate traffic with zero loss using a 1518-byte frame size. With the 1518-byte frame size, the switch performed at full line rate, with zero loss. The aggregate throughput calculated was 38.55 Gbps through the two LAG ports bidirectionally.

3.8 QoS

Our testing focused on VLAN-based and protocol-based QoS. For VLAN based QoS, VLAN 2 was configured with the highest priority, VLAN 1 with mid priority, and VLAN 3 with the lowest priority. Each of the VLANs forwards traffic from the Ixia to a single receive port, which in reality would oversubscribe the port causing frame loss. However, when VLAN based QoS features were implemented, the switch properly prioritized the traffic to allow all of the higher priority traffic to be delivered to the receive port. The middle and lower priority VLAN traffic was properly limited, allowing the receive port to receive all of the higher priority VLAN traffic.

QoS VLAN-based				
	VLAN1 (Mid Priority)	VLAN2 (Highest Priority)	VLAN3 (Low Priority)	Total % Line Rate
Tx Line Rate [%]	75	75	75	225
Rx Throughput per VLAN [Mbps]	198.00	740.25	50.06	988.31
Rx Line Rate [%]	20.06	75.00	5.07	100.13
Frame Loss [%]	73.23	0.00	93.23	124.84

QoS Protocol-based				
	TCP 8 (Mid Priority)	UDP 7 (Highest Priority)	TCP 9 (Low Priority)	Total % Line Rate
Tx Line Rate [%]	75	75	75	225
Rx Throughput per Protocol [Mbps]	198.00	740.25	50.06	988.31
Rx Line Rate [%]	20.06	75.00	5.07	100.13
Frame Loss [%]	73.23	0.00	93.23	124.84

Protocol-based QoS functions similarly to the VLAN-based QoS, however, traffic is prioritized using a protocol rather than a VLAN ID. Transmit traffic was tagged with three different protocols (UDP 7, TCP 8, and TCP 9), and each transmitted at 75% of line rate to one receive port. As with the VLAN-based QoS test, the receive port would essentially be oversubscribed if the QoS function is not enabled. Protocol-based QoS properly prioritized traffic for each protocol type. All traffic transmitted with UDP 7 traffic was received. The excess traffic with TCP 8 and TCP 9 packets were dropped. The Dell Networking switches use VLAN- and protocol-based QoS correctly.

3.9 Multicast

The RFC 3918 throughput no drop rate methodology was used to verify the performance of both the Dell Networking 7024 and 8164 switches. Traffic was transmitted from one port and received on the remaining ports on the switch.

Figure 8: Dell 7024 and 8164 Switches RFC 3918 Throughput

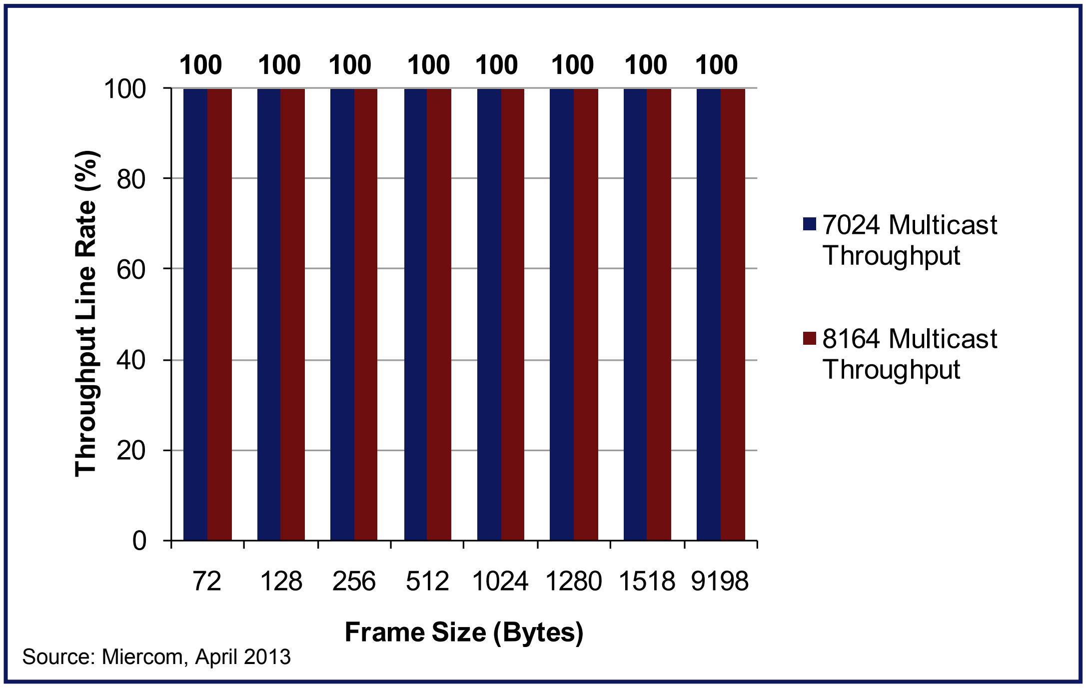

Dell Networking 7024 and 8164 switches achieved 100% line rate multicast traffic with zero loss.

Both switches were capable of snooping the multicast groups and then properly transmitted multicast traffic at 100% line rate with zero loss to each multicast group member.

4.0 Scalability

As networks grow, switches need to support more devices, and will be required to support higher capacities. To test the Dell Networking 7024 and 8164 switches for scalability, we focused on route capacity and MAC address cache size.

4.1 Route Capacity

To verify the route capacity of each switch, we used the Ixia IxAutomate OSPF route capacity test methodology. The test finds the maximum number of OSPF routes the switch can sustain. We verified that each Dell Networking switch achieved 12,288 routes, as stated in the Dell Networking datasheet.

Route Capacity

Dell Networking 7024		Dell Networking 8164	
Datasheet Value	Tested Value	Datasheet Value	Test Value
12,288 Routes	12,288 Routes	12,288 Routes	12,288 Routes

Dell Networking 7024 and 8164 each achieved 12,288 routes.

By default, each switch reserves 8,192 routes for IPv4 and 4,096 for IPv6. In the switch configuration, we disabled the IPv6 routing table memory allocation to test the maximum amount of IPv4 routes.

4.2 MAC Address Cache Size

We tested to verify the maximum number of supported MAC addresses. Three ports were used in the port mapping. One port was configured as the transmit port, while another was configured as a receive port. The third port was used for monitoring flooding. The transmit port was configured to transmit random MAC addresses to the receive port, until the MAC table was filled.

MAC Address Cache Size

Dell Networking 7024		Dell Networking 8164	
Datasheet Value	Tested Value	Datasheet Value	Test Value
32,768 MACs	32,768 MACs	131,072 MACs	131,072 MACs

Dell Networking 7024 and 8164 each achieved their datasheet values.

Both Dell Networking switches achieved their stated MAC address cache. During the learning process, the average CPU utilization for the 7024 was 13.3%, and 7.21% for the 8164 switch.

5.0 Power Efficiency

Networks are expanding rapidly, which can result in more power consumption by devices. With technology becoming much more advanced, switch vendors are designing their switches more efficiently. Switches are now designed to consume less power and are built with energy efficient standards, such as IEEE 802.3az Energy Efficient Ethernet (EEE) and various other green features. The Dell Networking 7024 and 8164 switches specifically support EEE and Energy Detect. EEE mode allows the switch to reduce the overall power consumption during periods of little or no link utilization. Energy Detect mode shuts down a port when it is not in use. The switches are also built with 80 PLUS certified power supplies and support multi-speed fan operation.

5.1 Energy Efficient Ethernet (EEE)

To take advantage of the EEE energy saving feature, an EEE-capable device needs to be connected to the switch. To test the overall power consumption of the switch with EEE, Miercom engineers configured each switch in a snake configuration. Essentially this configuration allows traffic to forward from port one on the switch to the last port on the switch, *see image below*.

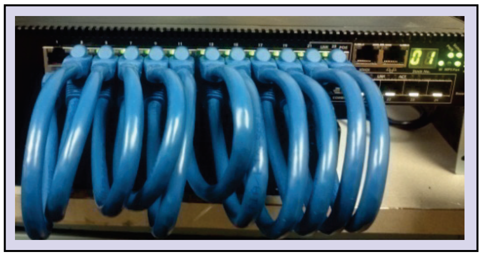

Ixia traffic is injected into port one, which is then forwarded to port two internally via VLAN 2. Port two is then directly connected with a CAT5 Ethernet cable to port three. Traffic received on port three is then forwarded to port four via VLAN 3, and so on until it reaches the last port and returns to the Ixia.

The traffic pattern must be unique for EEE to activate. Sending traffic at 100% line rate will not show any power savings. There must be an Inter-Burst Gap (IBG) between each packet. Power consumption is dependent on the length of the gap. With a tiny gap of 96ns, there will not be any power savings. However, with an IBG of 100ms, there should be a noticeable reduction in energy consumption. Traffic is sent with an IBG ranging from 96ns to 100ms, and a continuous burst of 1,000 packets. It is during these gaps where the PHY goes into Low Power Idle (LPI) mode.

While the 7024 switch was idle with link, there was a savings of 8.2 watts. When sending traffic with an IBG of 100ms, there was a savings of 8.1 watts, or 13%. Power savings ranged from 4.6 watts to 8.2 watts (7% to 13%) when EEE was enabled.

Figure 9: Dell Networking 7024 Energy Efficient Ethernet

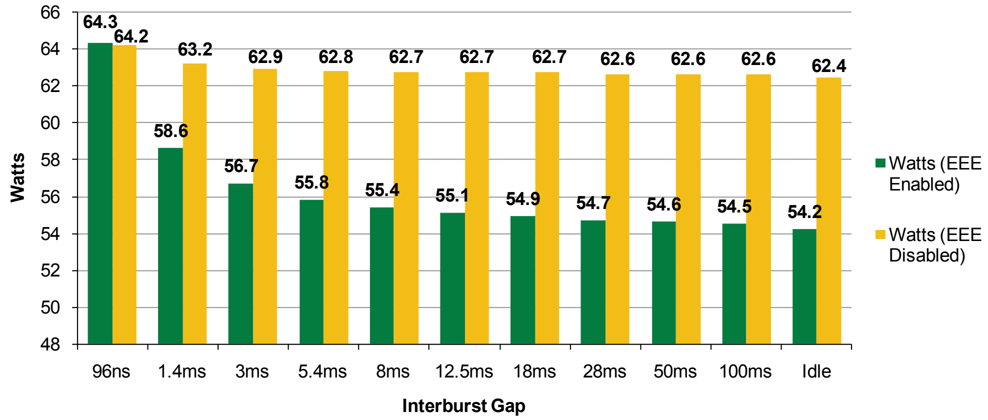

Source: Miercom, April 2013

Figure 10: Dell Networking 8164 Energy Efficient Ethernet

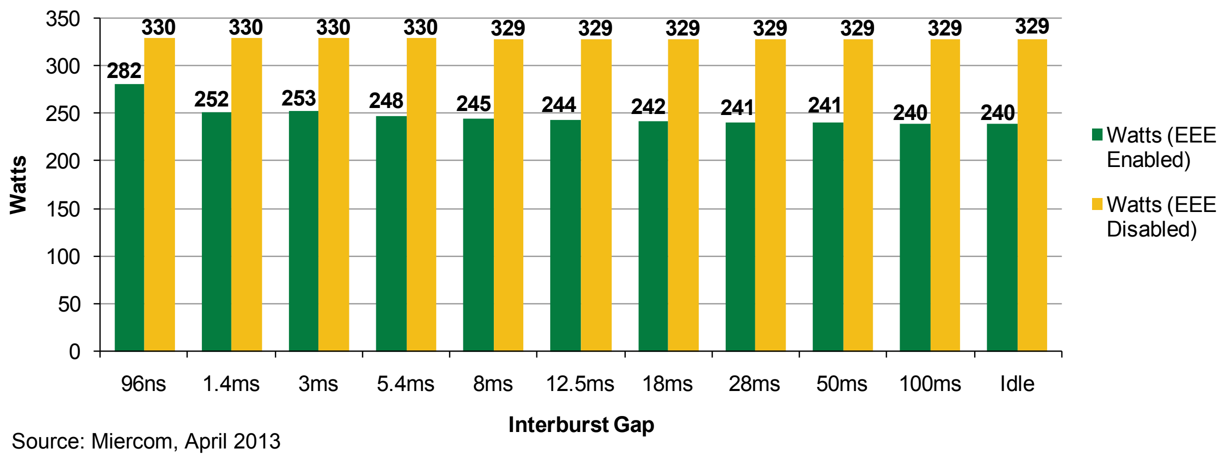

Source: Miercom, April 2013

The Dell Networking 8164 was set up with the snake configuration, interconnected using CAT6A cables to allow 10GbE traffic to pass. The 8164 consumed more power than the 7024 switch; however, the 8164 saved more power. We validated a total of 89 watts saved when using a 100ms IBG, which resulted in a 27% power savings. The Dell datasheet for the 8164 switch specifies that EEE is only compatible with 10GbE ports. We attempted to test EEE with 1GbE ports, and did not see any power savings.

5.2 Energy Consumption at Full Load

Measuring the power consumption of the switch at 100% line rate at various frame sizes is an important test. It shows how efficient the switch is and verifies the maximum power consumed by the switch. Typically, smaller packets require more processing power, causing more power to be consumed. Our testing focused on Layer 2 and Layer 3 traffic at 100% line rate.

Figure 11: Dell Networking 7024 Layer 2 Power Consumption

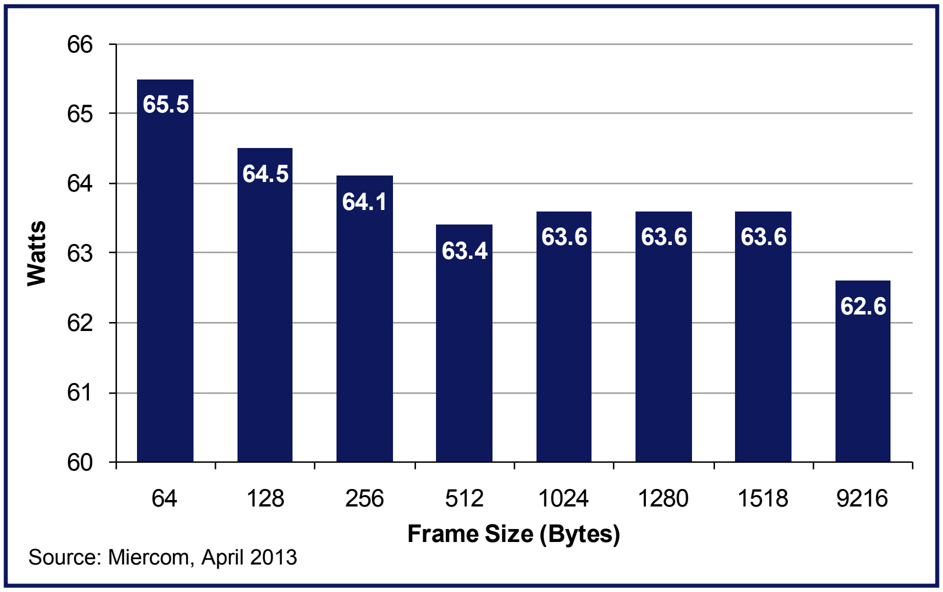

Testing was performed with EEE and Energy Detect enabled.

The difference between Layer 2 and Layer 3 power consumption was ~.4 watts.

Power consumption at the 64-byte frame size was 65.5 watts, while the 9216-byte jumbo frame consumed 62.6 watts.

Figure 12: Dell Networking 8164 Layer 2 Power Consumption

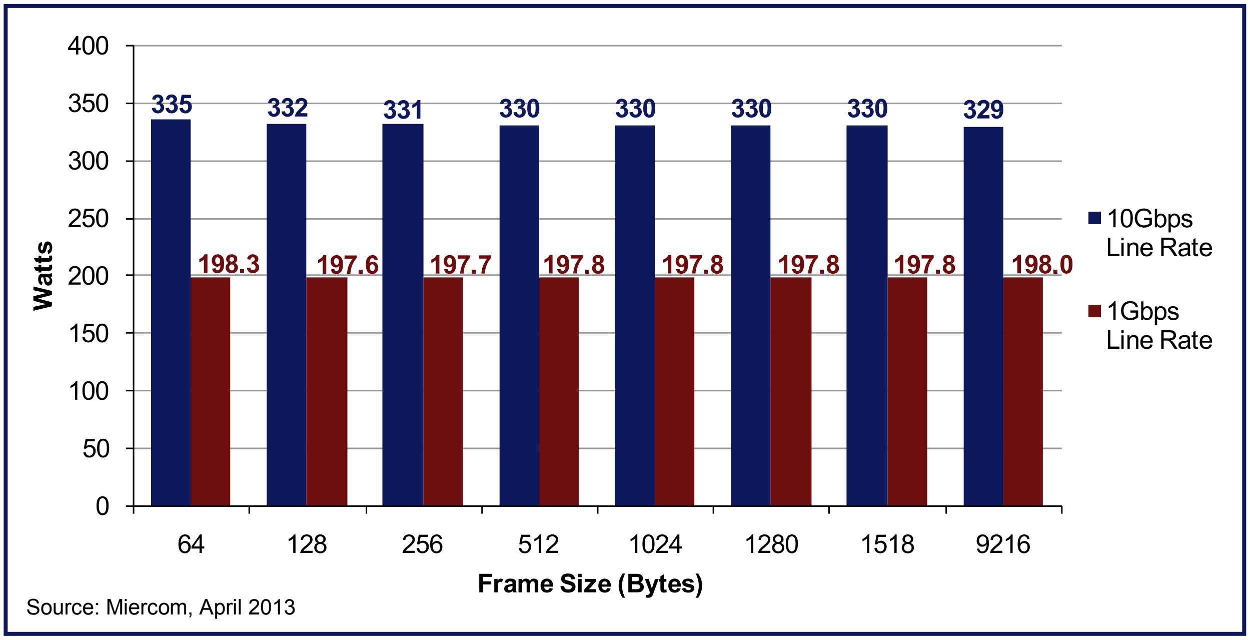

With 1GbE ports fully loaded, the 8164 consumed a maximum of 198.3 watts. When fully loaded with 10GbE ports, the maximum usage was 335 watts.

6.0 Security

We performed a brief security assessment to see how the Dell Networking 7024 and 8164 switches behaved during an IP DoS attack. Spirent Studio Security was used to send the DoS attack at a rate of 100,000 packets per second.

Several different attack iterations were performed:

- DoS against management IP address, with DoS protection disabled
- DoS against management IP address, with DoS protection enabled
- DoS to a laptop connected to the switch, with DoS protection disabled
- DoS to a laptop connected to the switch, with DoS protection enabled

During each attack, we monitored the Web GUI for any sluggishness, and the CPU utilization through a console CLI connection.

When issuing the attack to the management IP address with DoS protection disabled on the switch, the GUI stopped responding for the duration of the attack. The CPU utilization reached 52%. Once the attack was stopped, the GUI was functional once again. For the duration of the attack, the CLI remained functional.

When the IP DoS attack was issued against the management IP address with DoS protection enabled, the Web GUI remained functional. The CPU utilization averaged 6%.

When the attack was sent through the switch to a laptop on a different VLAN with DoS protection disabled, the laptop eventually became unresponsive. With DoS protection enabled, the attempted attack to the laptop was blocked.

It is important to note that the Dell Networking switches employ effective measures to block DoS attacks; however, DoS prevention is not enabled by default as inspection rules and thresholds must be planned to account for devices connected to the switch ports to enable the highest level of protection.

7.0 Interoperability with Cisco Catalyst 3750X and 4500X Switches

Ensuring interoperability with other switch vendors is an important step when designing and developing a switch. Each switch vendor is different. Some vendors develop their own protocols, while others use standard protocols for accomplishing switch functions.

For example, Cisco supports two different Link Aggregation Group (LAG) protocols. One is Port Aggregation Protocol (PagP), which is developed by Cisco and exclusive to only their devices. Cisco also supports the standard IEEE 802.3ad LACP link aggregation protocol, which is interoperable with most switches in the market today.

To verify LAG interoperability between the Cisco Catalyst 4500X, 3750X and Dell Networking 7024 and 8164 switches, we used LACP. All other interoperability tests used standard protocols supported by both Dell and Cisco.

7.1 Management VLANs and Trunking

Management VLANs allow administrators to manage and configure switches either through a telnet session or Web GUI. Our test bed included the Dell 8164, 7024, and Cisco 4500X and 3750X. The test bed diagram below accurately depicts the switch configurations and connections. Essentially this configuration allows access devices part of VLAN 50 or VLAN 60 to manage all devices.

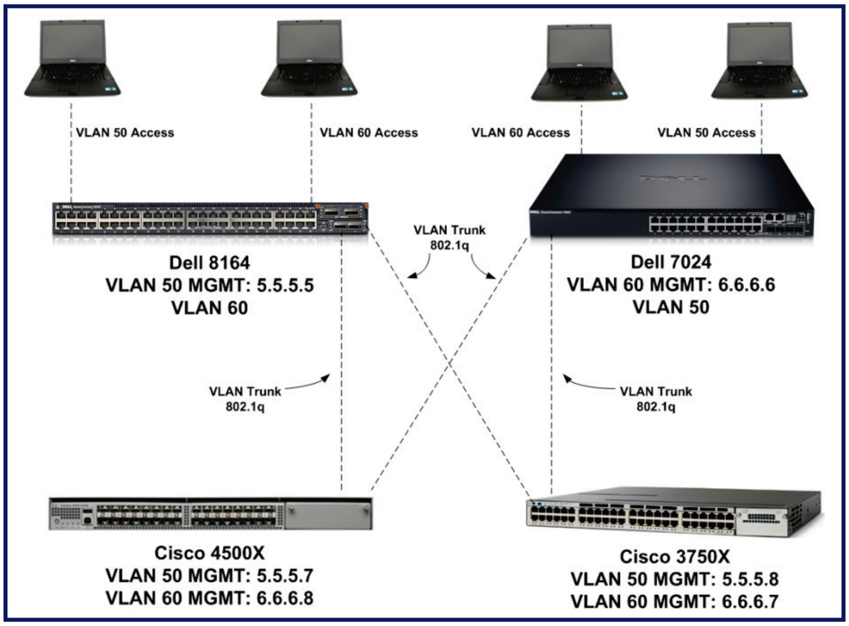

When connected to the Dell 8164 via VLAN 50 and VLAN 60, we successfully managed the directly attached switch, the Cisco Catalyst 4500X and 3750X. When we connected to the Dell 7024 via VLAN 50 and VLAN 60, we successfully managed the directly attached switch, the Cisco 3750X and 4500X. The Dell Networking 7024 and 8164 switches are fully interoperable with the Cisco 4500X and 3750X switches for management purposes when using trunks between the switches. The Dell Networking switches correctly forwarded management traffic to the Cisco switches.

7.2 Link Aggregation

To test the interoperability for link aggregation between the Dell and Cisco switches, we used LACP as the protocol to link the switches together. A port channel group is created on each switch pair, and they are then connected with each other using two Ethernet cables. The port channel was then configured as a trunk, to allow multiple VLANs. Once this configuration is set, the “show interfaces port-channel 1” command was issued to check the port channel status. The command shows the port channel name, assigned ports and status. Overall, we found that if we used the Dell and Cisco configuration guide, it was pretty straightforward to configure them both for LACP.

To send traffic through the port channel, two VLANs were created with different access ports assigned to them. Traffic was transmitted through these access ports. Traffic was then forwarded through the port channel to the destination switch. We successfully verified that the Dell and Cisco switches were interoperable using LACP, with 100% of traffic received.

7.3 Spanning Tree Protocol

The Spanning Tree Protocol (STP) prevents loops in networks that could cause broadcast storms. There are several STP variations and our interoperability testing focused on STP, Rapid Spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP). STP (802.1d) prevents loops; however, convergence time can take longer than 30 seconds. RSTP (802.1w) and MSTP (802.1s) also prevent loops, but convergence time is quicker than STP. For each STP protocol tested, we configured the test bed with fully redundant links. Each switch was plugged into each other to create a loop. STP detects the loop and configures ports into a blocking state.

The main focus of testing was to verify that Dell and Cisco switches could be configured with STP, and whether the protocol was utilized appropriately.

Traffic was transmitted through the switches to determine if failover worked efficiently. Links in the forwarding state were removed to force links in the blocking state to activate. With each STP variation, the switches seamlessly converged traffic and provided a loop-free network. The Dell Networking 7024 and 8164 and Cisco Catalyst 3750x and 4500x are interoperable when configured with different STP versions.

8.0 Total Cost of Ownership

Total cost of ownership (TCO) is an important factor when making purchasing decisions. Two main areas of focus for TCO are energy expense and the purchase price. We calculated the five year cost of ownership for a network consisting of 3,136 ports based upon energy consumption results.

Figure 13: Five Year Total Cost of Ownership

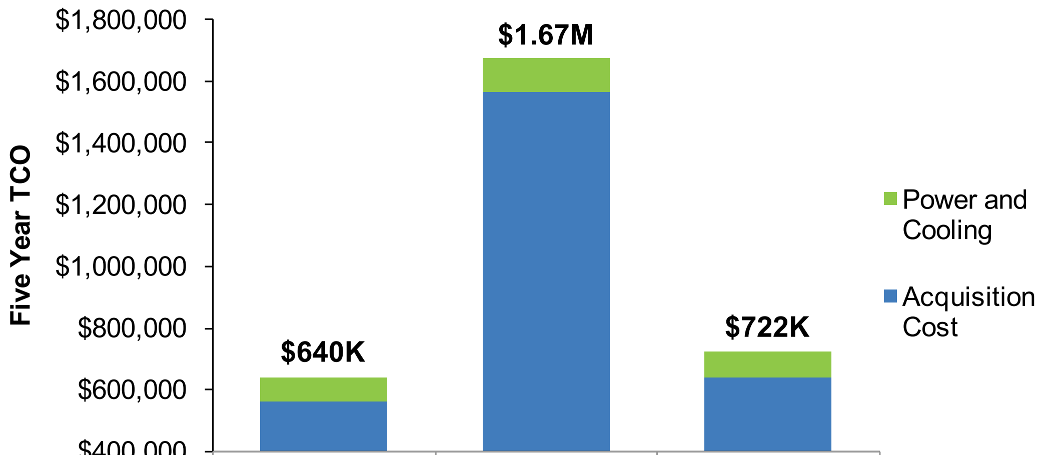

	Dell	Cisco	HP
Power and Cooling	$81,514.32	$106,221.63	$83,782.57
Acquisition Cost	$558,776.00	$1,566,400.00	$638,527.00

Source: Miercom, April 2013

Estimated pricing includes product and energy costs over five years for a 3,136-port network of these switches.

As a percentage of a network, the edge switch represents the vast majority of network devices and a significant percentage of network cost. For this analysis, the measurable components of TCO are: energy expense and purchase price. We calculated the five year cost of ownership for a network of 3,136 ports based upon energy consumption results. The Dell network contained two Dell Networking 8164 and one hundred twenty eight Dell Networking 7024P switches. The Cisco network contained two Catalyst WS-C4500X-32SFP+ and one hundred twenty eight Catalyst WS-C3750X-24P-E switches. The HP network was made up of two 8206 zl and one hundred twenty eight 3800-24G-PoE+ switches.

Figure 13 shows an estimated Five Year Total Cost of Ownership for the Dell, Cisco and HP switches. The Five Year Total Cost of Ownership was calculated by looking at current switch price costs using sites such as www.cdw.com and www.costcentral.com. Power and cooling costs were estimated using power consumption values at various loads from our testing. Power consumption values for the two HP switches and Cisco Catalyst WS-C4500X-32SFP+ were obtained from published datasheets.

Based on our analysis, we found that the Dell Networking switches had the lowest overall TCO when comparing the total acquisition costs and power and cooling costs for a five year period.

Calculation used to determine Power and Cooling costs:

1. The average power consumption, multiplied by the average duty cycle (21%). The idle power consumption of the switch multiplied by the idle duty cycle (79%).
2. Cooling is estimated by multiplying the average duty cycle by 33% of the power consumption.
3. Total of power and cooling is multiplied by the national average of cost per kWh (12.9 cents), then by 24 hours, then by 365 days, and then 5 years.
4. This number is then multiplied by the number of switches in the network.

These product prices and support costs are based on open source research and manufacturer published list prices. For more information, contact reviews@miercom.com for a specific TCO analysis on your products and applications.

About Miercom

Miercom has hundreds of product comparison analyses published in leading network trade periodicals including *Network World, Business Communications Review - NoJitter, Communications News, xchange, Internet Telephony* and other leading publications. Miercom's reputation as the leading, independent product test center is unquestioned.

Miercom's private test services include competitive product analyses, as well as individual product evaluations. Miercom features comprehensive certification and test programs including: Certified Interoperable, Certified Reliable, Certified Secure and Certified Green. Products may also be evaluated under the NetWORKS As Advertised program, the industry's most thorough and trusted assessment for product usability and performance.

Other Notes and Comments

Product names or services mentioned in this report are registered trademarks of their respective owners. Miercom makes every effort to ensure that information contained within our reports is accurate and complete, but is not liable for any errors, inaccuracies or omissions. Miercom is not liable for damages arising out of or related to the information contained within this report. Consult with professional services such as Miercom Consulting for specific customer needs analysis.

The tests in this report are intended to be reproducible for customers who wish to recreate them with the appropriate test and measurement equipment. Current or prospective customers interested in repeating these results may contact reviews@miercom.com for details on the configurations applied to the Device Under Test and test tools used in this evaluation. Miercom recommends customers conduct their own needs analysis study and test specifically for the expected environment for product deployment before making a product selection.