Lippis Report

A Comprehensive Testing of Cisco Systems Catalyst 6500 Sup2T

The World's Most Widely-Deployed Enterprise Ethernet Campus/Data Center Switch

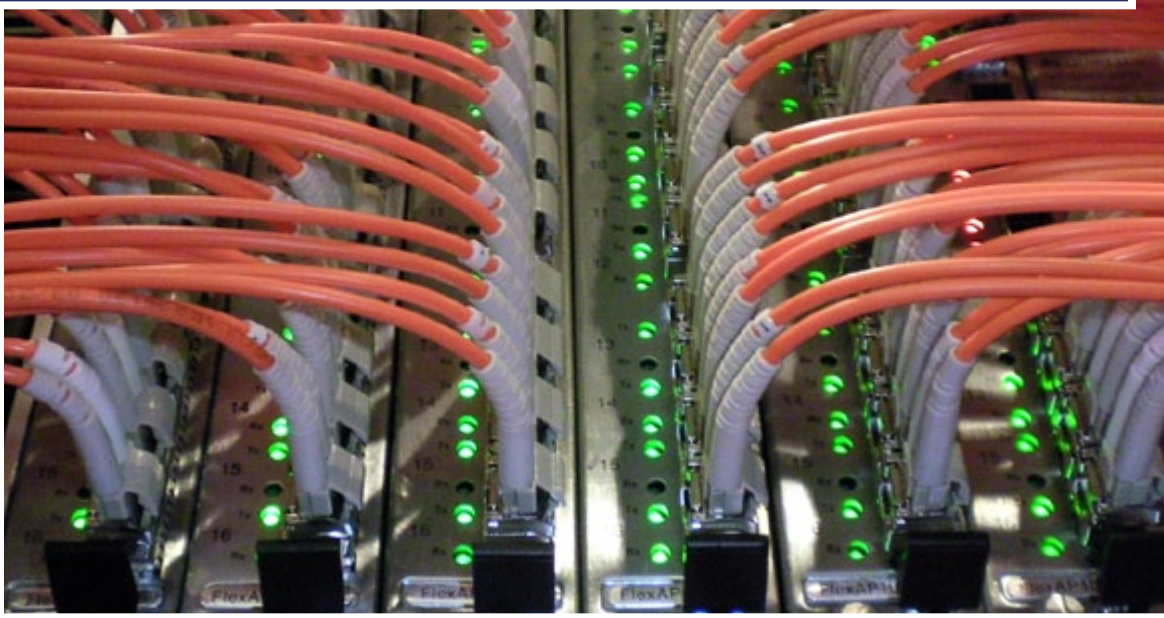

November, 2011

Forward

Packet per Second Calculation Side Bar

For throughput, we measure results in packets per second or PPS and packet loss as a percentage of line rate. Since the Catalyst 6500 Sup2T is a 2- to 4-Terabit switch, PPS is in the millions. PPS is a function of packet size, meaning that smaller packets at the same line rate, 10GbE in the document, will have a higher PPS than larger size packets. The following dimensional analysis demonstrates how to calculate PPS for various packet sizes.

Bits per Second rate /
([Packet Size + Preamble + IPG] * 8) = PPS

(rounded up to whole packets)
IPG= InterPacket Gap

For example: at 10GbE a 64Byte size IPv4 packet would have the following PPS rate.
10,000,000,000 / ([64B + 8B + 12B = 84 Bytes] * 8 = 672 bits) = ~14,880,953 pps or 14.8Mpps

For IPv6 we use a Preamble + IPG of 40Bytes.

A word about packet sizes tested for this report: The Cisco Catalyst 6500 is the most popular, meaning widely deployed, campus core and data center Ethernet switch in the industry. In these parts of the network, small size packets at line rate, i.e., 10GbE, seldom, if ever, occur in real networks. As such, packet sizes ranging between 256 Bytes and jumbo frame of 9260 Bytes were tested as these tend to be most popular, thanks to video and an ever-increasing enterprise application portfolio that relies upon larger data transfer and exchange. To deliver the highest performance possible, most ASIC and Ethernet switch engineers have decided to focus their design efforts on real-world traffic profiles for the targeted market segment of their switches. Therefore, while packet loss at 10GbE line rate for smaller size packets does occur, its impact on real networks is marginal as it's an extremely low probability scenario.

During the week of October 31, 2011, the Lippis Report tested Cisco System's new Catalyst 6500 with Supervisor 2T for performance, upgradability, control and scalability at Ixia's modern iSimCity laboratory in Santa Clara CA. The tests were conducted over a four-day period. Ixia supplied all test equipment needed to conduct the tests, while Cisco provided Catalyst 6500 chassis, modules, management software and engineering resources. Cisco was allocated lab time to run the tests with the assistance of an Ixia engineer. Together a test plan was developed to verify key Catalyst 6500 Supervisor 2T features and stress its performance. Nick Lippis, CEO of the Lippis Report, oversaw testing to assure the test plan was implemented correctly with Ixia and Cisco engineering.

By all counts, Cisco's upgrade of the Catalyst 6500 via its new Supervisor 2T, or Sup2T, is its most ambitious and thoughtful yet for the venerable platform. The Sup2T is a 2-Terabit (Tb) platform that triples the previous Sup720 performance. Thanks to the support of Virtual Switching System (VSS), the platform allows two 2-Tbps switches to combine into a single 4-Tbps virtual switch. The Sup2T is a major upgrade to the most widely-deployed switching platform in campus and data center networking in the industry. But while these performance numbers are impressive, it's the new Catalyst 6500's network services that deliver most of the value, which is partially found in the Sup2T's Policy Feature Card or PFC that increases NetFlow monitoring and a new TCAM design offering improved Access Control List (ACL), Quality of Service design options, encryption security and many other features.

Cisco's Catalyst 6500 is the firm's most successful product with over 700,000 systems and 110 million ports installed, worth some $42 billion in revenue over the years. This product's success increases the stakes for Cisco as it introduces a major upgrade. Cisco had to consider backward and forward customer migration, increased competition and pricing pressure, especially as competitors are starting to offer core switches based upon merchant silicon. In short, Cisco had to eliminate the trade-off of innovation versus investment protection and find a way to deliver both simultaneously. This Lippis Report test document verifies many of Cisco's performance and upgradability claims. While it's impossible to test all of the Catalyst 6500's new 200-plus features with the Sup2T, we rather focus on a select few that will have the widest impact on IT business leaders' product acquisition decision process.

The resources available for this test at Ixia's iSimCity are out of reach for nearly all corporate IT departments with test equipment on the order of $9.5 million, devices under test on the order of $1 million, plus lab costs associated with housing, power and cooling. It's our hope that this report will remove performance, upgradability, scale and management from the purchase decision, allowing IT architects and business leaders to focus on other vendor selection criteria, such as post sales support, platform investment, vision, company financials, etc.

The Lippis test reports are based on independent validation at Ixia's iSimCity and communicate credibility, competence, openness and trust to potential buyers of enterprise and data center switching equipment. The tests are based upon RFC and custom tests that are repeatable. Some tests do not lend themselves to written media, but rather video. Therefore, six video podcasts accompany this Lippis Report test. Video podcast links are embedded where appropriate.

There are three categories of tests to which the Catalyst 6500 was subjected. These include 1) compatibility, upgradeability and investment protection, 2) switching performance, and 3) system networking.

Compatibility, Upgradeability and Investment Protection Test

In this test, we look to measure how smooth the upgrade from Sup720 to Sup2T is. What IT business leaders are looking for are incremental network upgrades with minimal disruption versus major disruption that usually accompanies a significant and, at times, a not so significant network upgrade. Therefore, we will swap out Sup720 for Sup2T and bring up existing service modules and line cards. Remember that line cards represent the largest investment in switching equipment, so we'll demonstrate that older line cards interoperate at high performance when the new Sup2T replaces the Sup720.

Results: We found that upgrading the Catalyst 6500 from Sup720 to Sup2T within the 6513-E chassis was straightforward and compatible with existing line cards and service modules. Those who invested in the E series chassis (i.e. 6503-E to 6513-E) and purchased line cards and service modules will find that this investment is protected and enhanced as new network services such as NetFlow, TCAM architecture improvements, encryption, deeper QoS granularity, Access Control Lists (ACLs), dry-run and atomic commit et al are added during supervisor upgrade from 720 to 2T. Please watch the video to review the methodology used and results gained.

The Supervisor 2T supports existing 6100 series and 6700 series Ethernet line cards, as well as legacy Service Modules. 6700 series line cards equipped with a Distributed Forwarding Card 3 (DFC3) installed require a daughter-card upgrade, but it becomes operational immediately following a simple upgrade.

Much of the new features and performance increases are found within the Policy Feature Card or PFC of the Sup2T. The PFC4 provides hardware accelerated forwarding for packets traversing the switch. This includes forwarding for IPv4 unicast/multicast, IPv6 unicast/multicast, Multi-Protocol Label Switching (MPLS) and Layer 2 packets. Along with forwarding, the PFC4 is also responsible for processing a number of services that are handled during the forwarding lookup process. This includes processing of security ACLs, applying rate limiting policies, quality of service classification and marking, NetFlow flow collection and flow statistics creation, EtherChannel load balancing, packet rewrite lookup, and packet rewrite statistics collection. The Distributed Forwarding Card 4 or DFC4 is a daughter card that is located on selected line cards. The DFC4 contains the same ASIC functional blocks that are found on the PFC4. The main purpose of the DFC4 is to provide local forwarding services for the line card, offloading this function from the PFC4.

Methodology: To assure the Catalyst 6500 Sup720 based modules and chassis are compatible with the Catalyst 6500 Sup2T, we populated a 13-slot enhanced-performance chassis (SKU WS-C6513-E) with the following line cards and service modules:

Slot #	Module Type
1	WS-X6148A-GE-45AF
2	WS-X6148E-GE-45AT
3	WS-X6724-SFP (CFC)
4	WS-X6748-SFP (CFC)
5	WS-SVC-NAM-2-250S
6	WS-SVC-FWM-1
7	VS-S720-10G
8	VS-S720-10G
9	ACE20-MOD-K9
10	WS-SVC-WISM-1-K9
11	WS-X6748-GE-TX (CFC)
12	WS-X6704-10G (DFC3B)
13	WS-X6716-10G (DFC3C)

Notice slots 7 and 8 are occupied by Sup720s while slots 5, 6, 9 and 10 are populated with service modules. The remaining slots are populated with a combination of line cards, some Centralized Forwarding Cards (CFC) and some Distributed Forwarding Cards 3 B and C (DFC 3). A note on Cisco's module nomenclature; please see figure for reference. Line cards 61xx as in slot 1 and 2 are bus based wiring closet facing linecards with POE & POE+ capabilities. 40Gbs bus based line cards include the 67xx-CFC as in slot 3, 4 and 11 and offer no distributed forwarding.

Line card 67xx-DFC3 as in slot 12 and 13 are DFC where packets are forwarded on the line card versus centralized in the supervisor. Line cards 68xx-DFC4 are not included in this test, but they are equipped with the new DFC4 daughter card that deliver new feature and performance benefits as identified above. The Sup2T upgrades module backplane access speed from 40Gbs to 80Gbs. To take advantage of that speed, Cisco introduced the new 6900-DFC4 line cards, which are also equipped with DFC4. Network Analysis, Firewall, ACE and WLAN service modules occupy slots 5, 6, 9 and 10, respectively. Via command line interface, we issue the command to show modules and produce the following snapshot of the Catalyst 6500 Sup720.

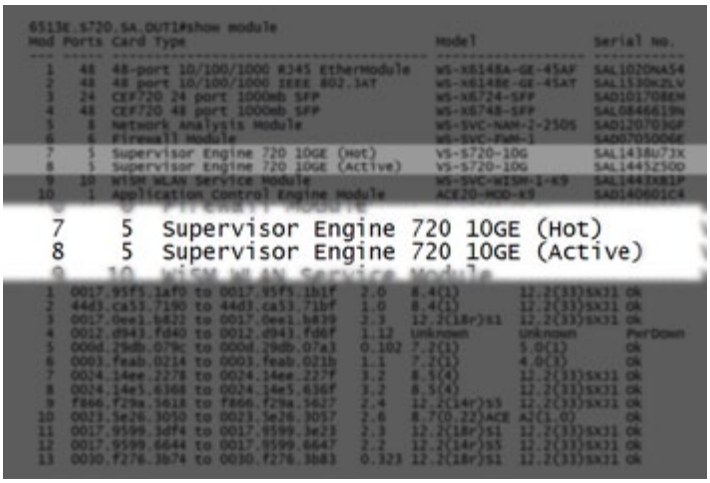

The above confirms the existing system is operational; that is, all modules are online, except for dual-channel line cards, which occupy the upper 6 slots, e.g., slot 4. We copy the running configuration to a Compact Flash drive and then power-down the system. We replace the existing Sup720 (VS-S720-10G) modules, with new Sup2T (VS-SUP2T-10G) modules. We also move the existing SFP and X2 transceivers between modules, and insert the existing Compact Flash disk to the new modules. We then power-up the system. Again we issue the CLI command to show module and produce the following system status.

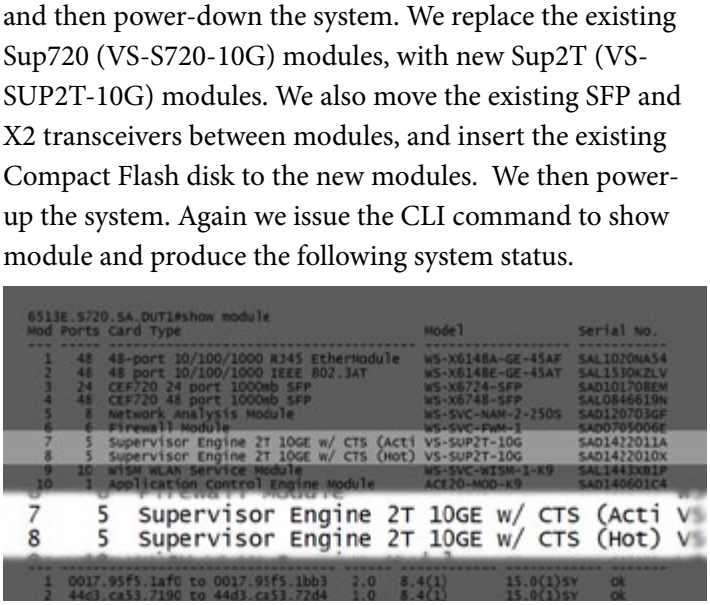

We confirm the new Sup2T modules, as well as service modules, bus-based and CFC-based line cards (including dual-channel modules in upper slots) come online after the Sup2T is installed and the Catalyst 6500 is powered up. We copied the original Sup720 configuration file to the Sup2T running configuration and confirm that it is applied.

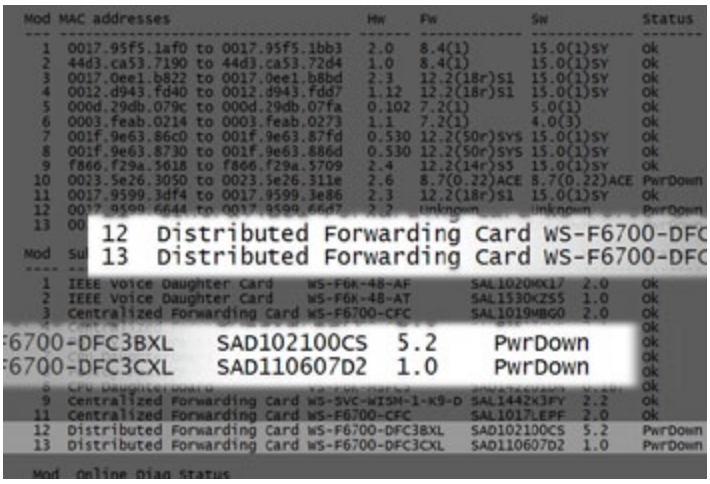

Notice that slots 12 and 13 6704-10G DFC3B and 6716-10G DFC3C are not online after the Sup2T upgrade as DFC3 daughter cards are not compatible with Sup2T PFC4. The DFC3 daughter cards need to be replaced with DFC4. Network operations can perform this swap as it is straightforward, or it can be done by Cisco professional services. Please see the above video for a demonstration. We upgrade the DFC3-based line cards with the new DFC4, reinstall them and confirm that the original configuration was retained and applied. Again we issue the CLI command to show module and produce the following system status.

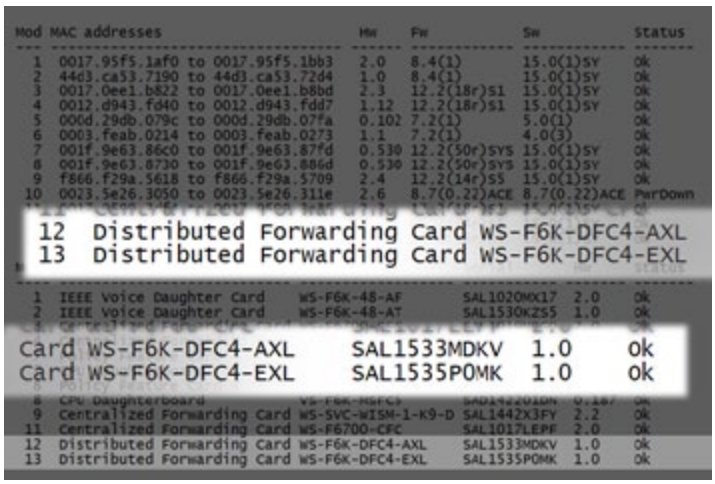

After the DFC4 upgrade, notice that slots 12 and 13 are now online and their module model names change to reflect the new DFC4 daughter cards are installed. Notice that the ACE service module is not online; this was due to a firmware upgrade requirement that was later performed. Therefore, when upgrading from Sup720 to Sup2T with ACE service modules, be sure to be running the latest firmware. Once the firmware was uploaded, the ACE service module came online.

Finally, we test to assure that the Sup2T when combined with the 6513-E chassis enables high-performance dual-fabric line cards to operate in the upper 6 slots. This was also verified when the previously "PwrDown" 6748-SFP (Slot 4) + CFC came online, once Sup2T was booted. To test this, we replace the 6148-GE line cards from slots 1 and 2 with new 6908-10G dual-fabric line cards. Note that the 6908-10G is supported on the Sup2T being a non-blocking 10GbE module with dual-channel support and DFC4 pre-installed, meaning that it supports all the new services such as TrustSec, new ACL enhancements, NetFlow, QoS, etc., features. Again we issue the CLI command to "show module" and produce the following system status. We confirm that the new dual-channel modules come online in the upper slots and that these line cards interoperate with existing modules.

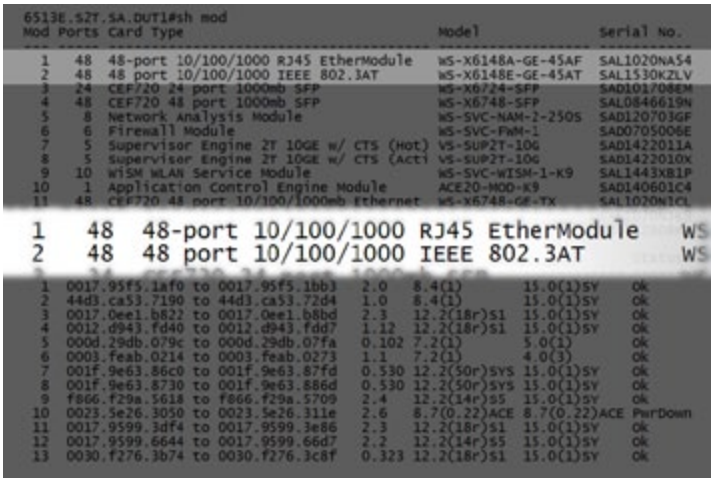

This test verifies backward compatibility of the 6513-E Catalyst 6500 Sup2T with existing service modules, bus-based and CFC-based line cards along with feature and performance benefits afforded by the Sup2T (PFC4). We further verify the upgradability of existing modules which currently employ the DFC3(B and C) daughter card with

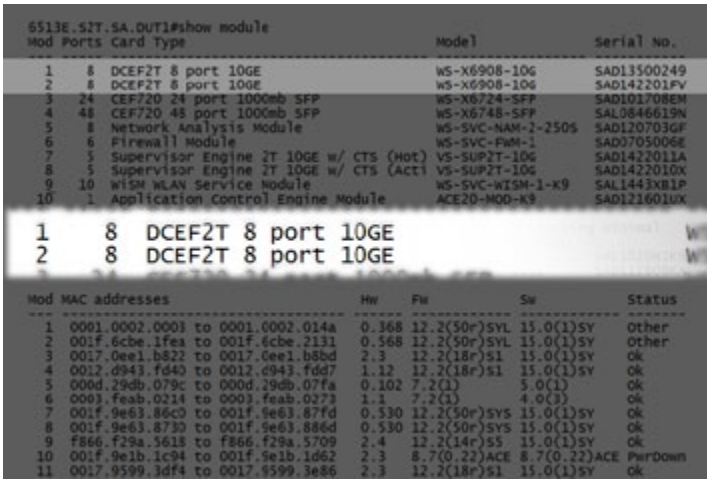

feature and performance benefits afforded by the DFC4 upgrade. We also verify the migration of current IOS configuration (as applicable to existing line cards) as well as the reuse of existing interface transceivers (e.g., SFP & X2). Finally, we verify the Sup2T when combined with the 6513-E chassis enables high-performance (dual-fabric) line cards to operate in the upper 6 slots. The following figures show the front views of the Catalyst 6500 before and after the Sup2T and line card upgrades.

GUI LMS View

GUI LMS View

In this same 6513-E chassis, we replaced the Sup720 for Sup2T, upgraded the line cards in slots 1 and 2 for the new 6908s, upgraded the DFC4 daughter cards in slots 12 and 13 and kept the same service modules. All of this was done while the Catalyst 6500 was operational. The Sup2T triples the performance of Sup720 while adding greater network service features such as Flexible NetFlow monitoring, Mac-Sec of 802.1ae based encryption security, WLAN integration, and firewall protection..

Switching Performance Test

Switching performance in enterprise networks is becoming increasingly important, as IT responsibility has been split between employees and IT departments, thanks to BYOD or Bring Your Own Device and IT consumerization. As a result, the number of devices on the network has increased significantly as employees bring smartphones and other mobile devices into the work force. These devices and their applications are driving unforeseen network requirements in terms of performance and support of both IPv4 and IPv6 as many mobile devices are now set for IPv6 as the default.

For IPv4 and IPv6, dual stack implementations are most popular where desktops and mobile devices run both IPv4 and IPv6, therefore, the network infrastructure needs to support both equally at high performance. IPv6 performance has not been on par with IPv4 until now. To demonstrate how the Catalyst 6500 upgrade with Sup2T has improved IPv6 performance, we measure IPv4 and IPv6 unicast and bidirectional traffic performance via RFC 2544.

To measure IPv4 and IPv6 unicast performance on the Catalyst 6500, we fully populate a 6513-E chassis with 92 10GbE ports via 11 WS-X6908-10G line cards and two VS-SUP2T-10G supervisors running IOS 15.0(1)SY network operating system. Each WS-X6908-10G supports eight 10GbE ports while each SUP2T-10G supports two 10GbE ports. Each Catalyst 10GbE port is connected to Ixia test gear and configured for IPv4 and then IPv6 addressing. The Ixia equipment used was the OPTIXIAXM12, a 12-slot XM form factor high performance chassis populated with six Xcellon FlexAP10G16S 16 port 10GbE load modules.

RFC2544 throughput test is executed first for IPv4 then for IPv6 with results captured.

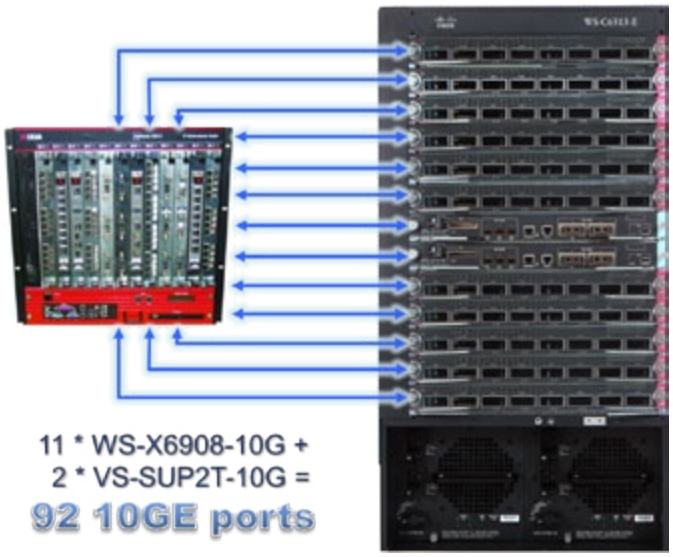

Results: We test the Catalyst 6500 for throughput between popular enterprise network frame sizes ranging from 256 to 9216 byte size packets. We find that each WS-X6908-10G delivers IPv4 and IPv6 throughput at the theoretical maximum possible for packet sizes ranging from 256 to jumbo size 9216 at 10GbE.

Catalyst 6500 Sup2T RFC 2544 IPv4 Throughput Test Unicast

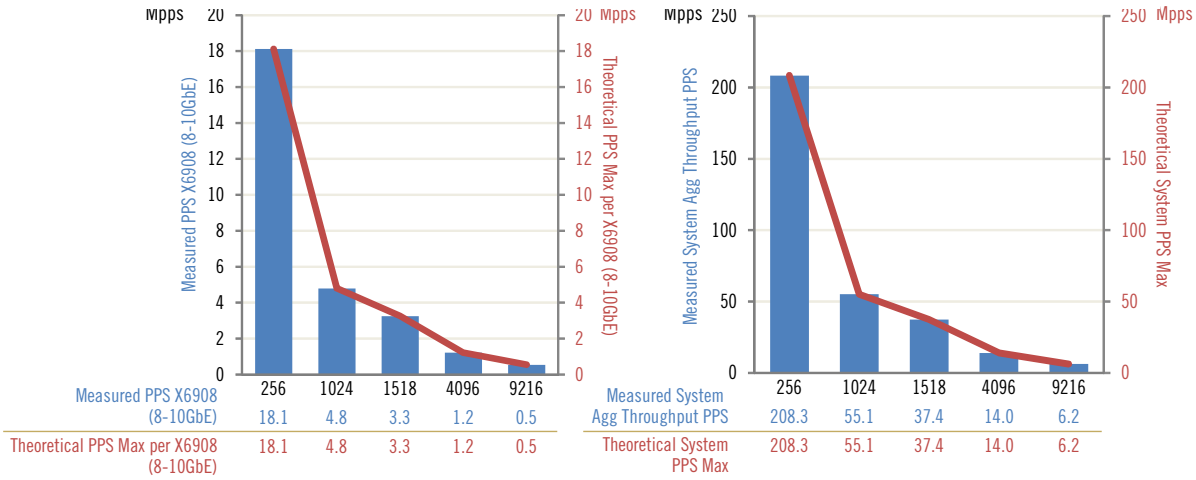

	256	1024	1518	4096	9216
Measured PPS X6908 (8-10GbE)	18.1	4.8	3.3	1.2	0.5
Theoretical PPS Max per X6908 (8-10GbE)	18.1	4.8	3.3	1.2	0.5
Measured System Agg Throughput PPS	208.3	55.1	37.4	14.0	6.2
Theoretical System PPS Max	208.3	55.1	37.4	14.0	6.2

IPv4 Unicast: At 256 byte size packets, the Catalyst 6500 Sup2T configured for IPv4 unicast traffic forwards packets at the theoretical maximum for 92 10GbE ports at 208 Mpps. At 9216 byte size packets, the Catalyst 6500 Sup2T configured for IPv4 unicast traffic forwards packets at the theoretical maximum for 92 10GbE at 6.2 Mpps.

Catalyst 6500 Sup2T RFC 2544 IPv6 Throughput Test Unicast

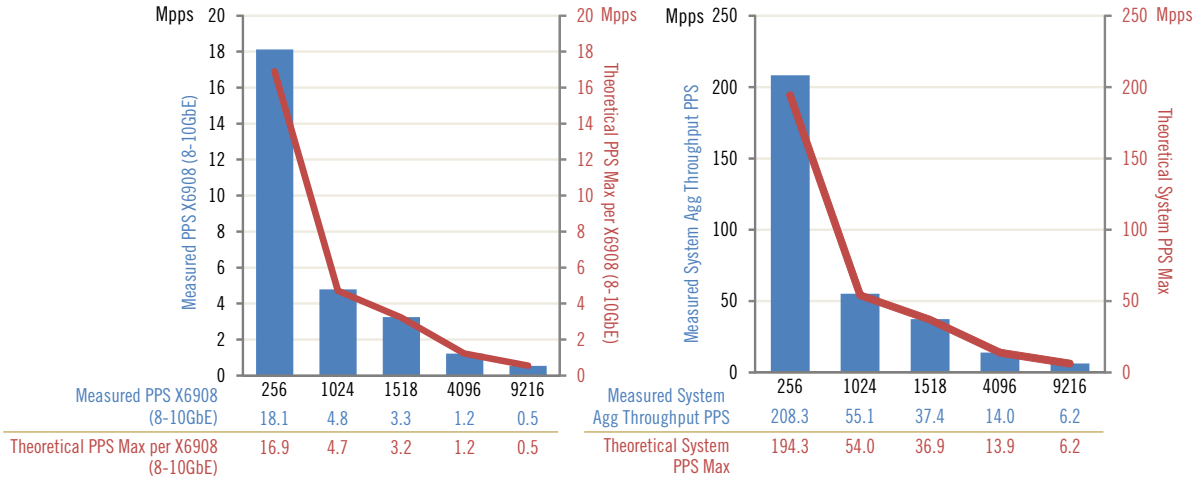

	256	1024	1518	4096	9216
Measured PPS X6908 (8-10GbE)	18.1	4.8	3.3	1.2	0.5
Theoretical PPS Max per X6908 (8-10GbE)	16.9	4.7	3.2	1.2	0.5
Measured System Agg Throughput PPS	208.3	55.1	37.4	14.0	6.2
Theoretical System PPS Max	194.3	54.0	36.9	13.9	6.2

IPv6 Unicast: At 256 byte size packets, the Catalyst 6500 Sup2T configured for IPv6 unicast traffic forwards packets at slightly higher than the theoretical maximum for 92 10GbE ports at 208 Mpps. At 9216 byte size packets, the Catalyst 6500 Sup2T configured for IPv6 unicast traffic forwards packets at the theoretical maximum for 92 10GbE at 6.2 Mpps. Measured performance being slightly higher than theoretical is more than likely due to header calculations of theoretical throughput.

The main observation of the above data is that IPv4 and IPv6 throughput performance is equal during unicast traffic this is especially important observation as IPv6 headers are nearly 4 times larger than IPv4. Therefore, IPv4 and IPv6 endpoints will experience the same high performance. The same configuration is employed for bidirectional traffic. Those results are:

Catalyst 6500 Sup2T RFC 2544 IPv4 Throughput Test Bidirectional

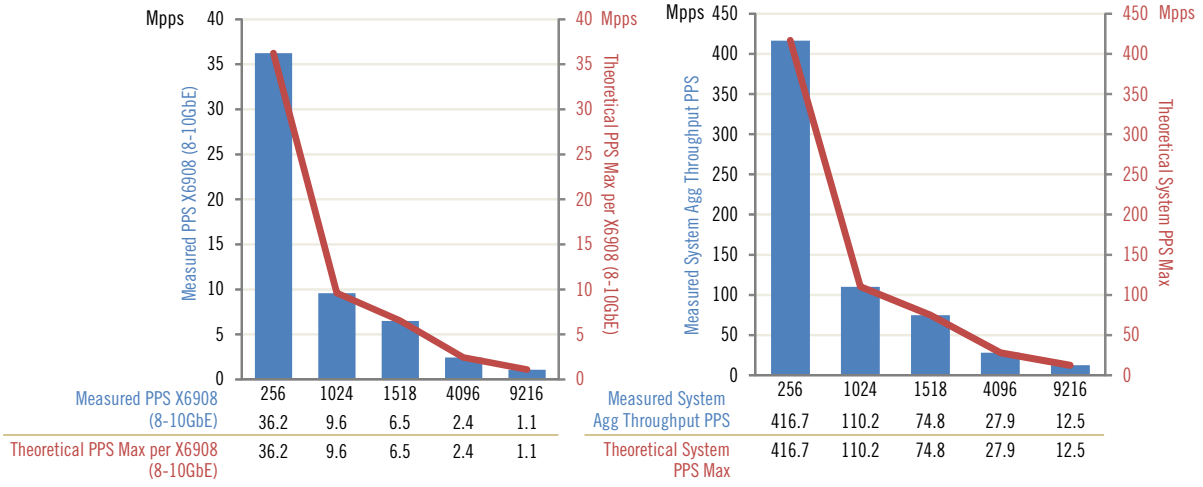

	256	1024	1518	4096	9216
Measured PPS X6908 (8-10GbE)	36.2	9.6	6.5	2.4	1.1
Theoretical PPS Max per X6908 (8-10GbE)	36.2	9.6	6.5	2.4	1.1
Measured System Agg Throughput PPS	416.7	110.2	74.8	27.9	12.5
Theoretical System PPS Max	416.7	110.2	74.8	27.9	12.5

IPv4 Bidirectional: At 256 byte size packets, the Catalyst 6500 Sup2T configured for IPv4 bidirectional traffic forwards packets at the theoretical maximum for 92 10GbE ports at 416 Mpps. At 9216 byte size packets, the Catalyst 6500 Sup2T configured for IPv4 bidirectional traffic forwards packets at the theoretical maximum for 92 10GbE at 12.4 Mpps.

Catalyst 6500 Sup2T RFC 2544 IPv6 Throughput Test Bidirectional

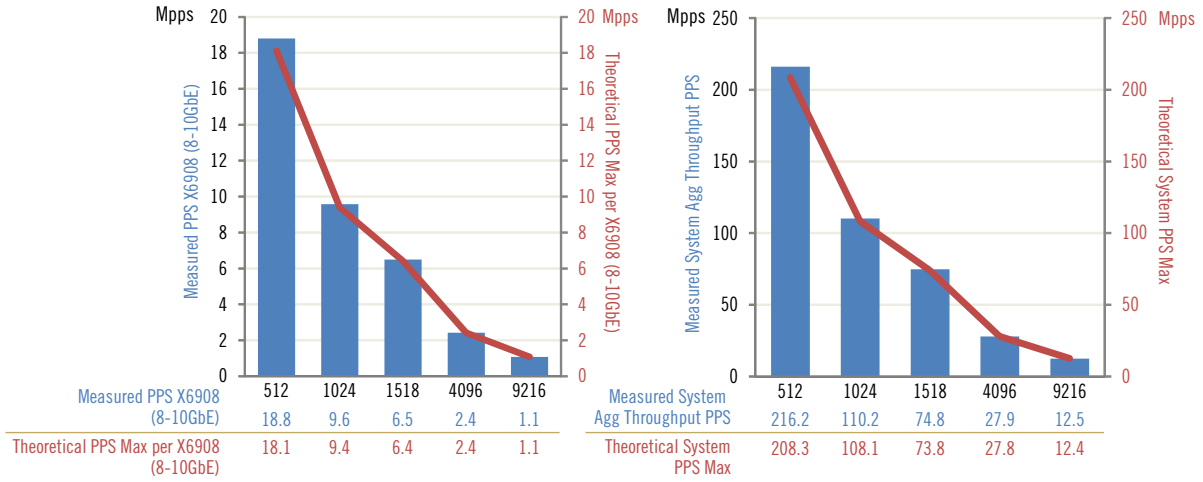

	512	1024	1518	4096	9216
Measured PPS X6908 (8-10GbE)	18.8	9.6	6.5	2.4	1.1
Theoretical PPS Max per X6908 (8-10GbE)	18.1	9.4	6.4	2.4	1.1
Measured System Agg Throughput PPS	216.2	110.2	74.8	27.9	12.5
Theoretical System PPS Max	208.3	108.1	73.8	27.8	12.4

IPv6 Bidirectional: At 512 byte size packets, the Catalyst 6500 Sup2T configured for IPv6 bidirectional traffic forwards packets at slightly higher than the theoretical maximum for 92 10GbE ports at 216 Mpps. At 9216 byte size packets, the Catalyst 6500 Sup2T configured for IPv6 bidirectional traffic forwards packets at the theoretical maximum for 92 10GbE at 12.4 Mpps.

For IPV6 unicast and bidirectional traffic, we measure throughput via RFC2544 and the Catalyst 6500 achieved the performance we expected. We compare IPv4 to IPv6 throughput performance and verify that the Catalyst 6500 delivers IPv6 performance on par with IPv4 noting that IPv6 headers are nearly four times larger than IPv4. Note that the Sup720 delivers IPv6 throughput at approximately half of IPv4. This performance gap has been closed with the new Sup2T, 6908-10G line cards and DFC4 daughter cards.

Catalyst 6500 Sup2T RFC 2544 IPv4/IPv6 Throughput Test - Unicast

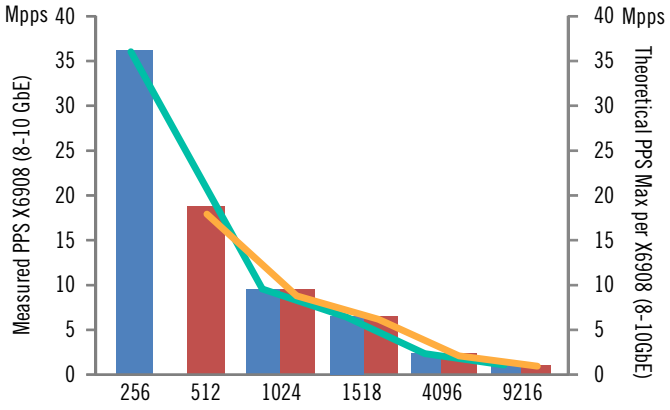

	256	512	1024	1518	4096	9216
IPv4 Theoretical PPS Max per X6908	36.2		9.6	6.5	2.4	1.1
IPv4 Measured PPS X6908	36.2		9.6	6.5	2.4	1.1
IPv6 Theoretical PPS Max per X6908		18.1	9.4	6.4	2.4	1.1
IPv6 Measured PPS X6908		18.8	9.6	6.5	2.4	1.1

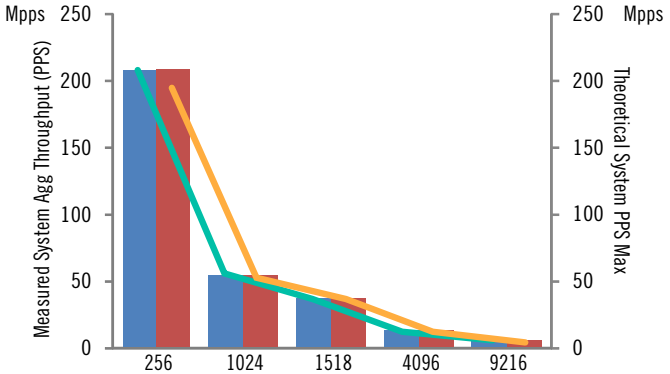

	256	1024	1518	4096	9216
IPv4 Theoretical System PPS Max	208.3	55.1	37.4	14.0	6.2
IPv4 Measured System Agg Throughput PPS	208.3	55.1	37.4	14.0	6.2
IPv6 Theoretical System PPS Max	194.3	54.0	36.9	13.9	6.2
IPv6 Measured System Agg Throughput PPS	208.3	55.1	37.4	14.0	6.2

Catalyst 6500 Sup2T RFC 2544 IPv4/IPv6 Throughput Test - Bidirectional

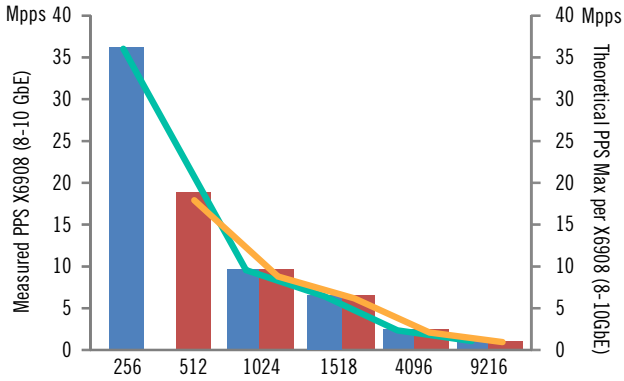

	256	512	1024	1518	4096	9216
IPv4 Theoretical PPS Max per X6908	36.2		9.6	6.5	2.4	1.1
IPv4 Measured PPS X6908	36.2		9.6	6.5	2.4	1.1
IPv6 Theoretical PPS Max per X6908		18.1	9.4	6.4	2.4	1.1
IPv6 Measured PPS X6908		18.8	9.6	6.5	2.4	1.1

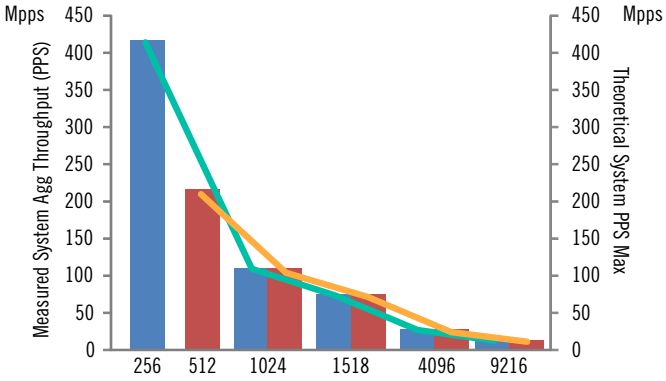

	256	512	1024	1518	4096	9216
IPv4 Theoretical System PPS Max	416.7		110.2	74.8	27.9	12.5
IPv4 Measured System Agg Throughput PPS	416.7		110.2	74.8	27.9	12.5
IPv6 Theoretical System PPS Max		208.3	108.1	73.8	27.8	12.4
IPv6 Measured System Agg Throughput PPS		216.2	110.2	74.8	27.9	12.5

Catalyst 6500 Sup2T RFC 2544 Throughput Test

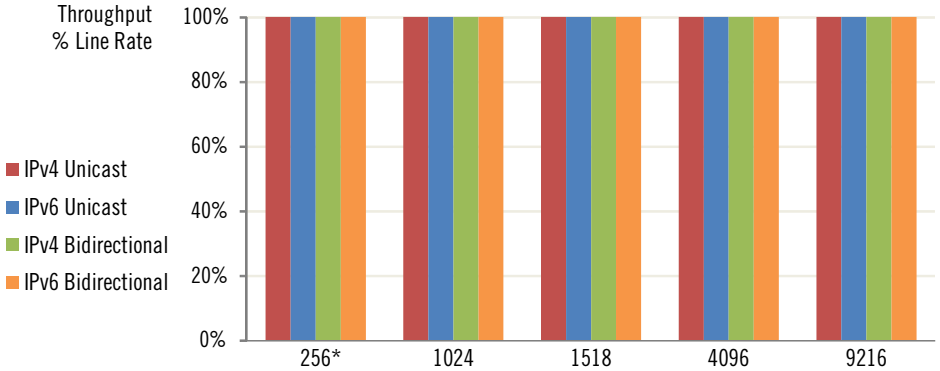

	256*	1024	1518	4096	9216
IPv4 Unicast	100	100	100	100	100
IPv6 Unicast	100	100	100	100	100
IPv4 Bidirectional	100	100	100	100	100
IPv6 Bidirectional	100	100	100	100	100

* The IPV6 RFC2544 Bidirectional test was run between packet sizes 512-9216

IP Multicast Test

IP Multicast traffic has been on the rise, thanks to the increased use of video services within the enterprise. Efficient use of multicast is important to interactive video, video surveillance, video dissemination, etc. Consider 500 to 1000 video surveillance cameras that need to stream their video to five or more locations within the enterprise, for regulation, storage, monitoring, etc. This is a popular requirement in gaming, retail, healthcare, etc. Streaming five streams per camera consumes a lot of bandwidth; therefore, using IP multicast reduces bandwidth consumption making video and other point-multipoint services efficient. Therefore, we test IP Multicast performance on the new catalyst 6500 Sup2T. This test stresses the packet replication ASIC built into the 6908-10G line cards.

The same single device configuration with Ixia test gear running the RFC 3918 IP Multicast test script was employed to test IP multicast traffic. Two IP multicast designs were configured in the Catalyst 6500. The first configuration is point-multipoint where one source is distributed to 91 receivers. The second configuration is a logical mesh topology or multipoint-multipoint where 11 sources distribute/broadcast to 77 receivers at 10GbE each. One 10GbE interface from each of the 11WS-X6908-10G modules broadcast to 77 receivers. Clearly the mesh configuration is overdesigned to stress the replicator ASIC within the line card.

Results: For the point-multipoint configuration, the Catalyst 6500 Sup2T demonstrated zero packet loss or 100% throughput at line rate while a single 10GbE source was broadcast to 92 receivers.

Catalyst 6500 Sup2T RFC 3918 IPv4 Multicast Test Point to Multipoint

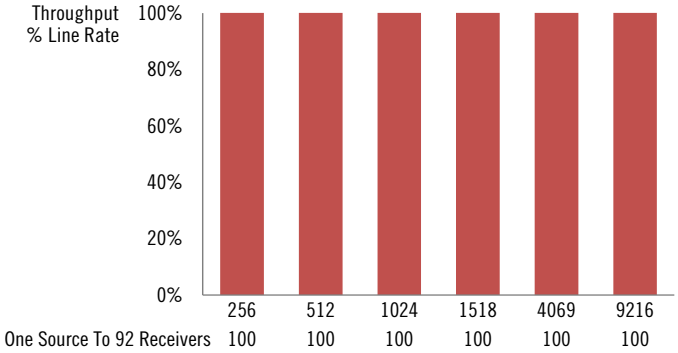

For the mesh multipoint-multipoint configuration, the Catalyst 6500 Sup2T demonstrated throughput performance that ranged from 49.8 Mpps to .53 Mpps for packet sizes that varied between 256 bytes to jumbo size or 9216 bytes. We find that the replication engine that is resident on Catalyst 6500 6908-10G line cards delivers multicast performance scale as there is no performance penalty for point-multipoint and multipoint-multipoint. This is due to the Sup2T having an improved hashing algorithm to support larger IP Multicast flows over the Sup720.

Catalyst 6500 Sup2T RFC 3918 IPv4 Multicast Test Mesh or Multipoint-Mulitpoint

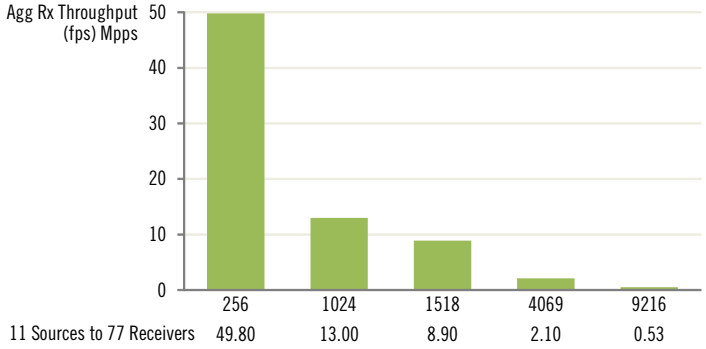

Access Control List Test

Access Control List or ACL are important tools in the configuration and customization of network attributes, especially with the Catalyst 6500. In the Catalyst 6500 upgrade with Sup2T, the TCAM has been both increased and its architecture improved. For ACL, one major concern was the lack of visibility of overflowing the TCAM when new ACL scripts were submitted, which would disrupt network operation. Updating ACLs occur infrequently and over a long period of time. As such multiple network engineers working on the same network may not even be aware of previous ACL updates. Further, an ACL update may drive multiple ACE (ACE = Access Control Entries), which occupy more TCAM resources than anticipated and thus over consume this resource. Therefore, Cisco developed the ACL Dry Run and ACL Atomic Commit to mitigate this scenario from occurring.

To verify the ACL improvements of the Sup2T, we define and configure a large security ACL 50K ACE to evaluate the increase in capacity of the ACL TCAM as well as the

hardware counters associated with each ACE. We use ACL Dry-Run to assure that the TCAM would not overflow, and then implement the changes safely with ACL Atomic-commit. This assures that no network interruption would occur when ACL changes were submitted to the Catalyst 6500.

Results: We verify that this new efficient use of TCAM and ACL safeguards perform as stated and is documented with this short video clip.

System Networking Test

To measure, report and verify throughput performance for a range of popular network designs including MPLS or Multi-Protocol Label Switching, VPLS or Virtual Private LAN Service and VSS or Virtual Switching System, we connect two Catalyst 6500s via eight 10GbE links of EtherChannel. But first, some perspectives on the use of these network designs.

Network virtualization, or the ability to divide a physical network into multiple logical networks with unique attributes, is a design that has grown in popularity as IT business leaders have searched for ways to segment their network with different attributes for different user groups. This is popular in healthcare, education, travel and other industries. Network virtualization can be implemented either in IP and/or MPLS. In addition to network virtualization, connecting the Catalyst 6500 with Sup2T directly to service provider MPLS networks is another popular design; therefore we test for both scenarios here. We test the Catalyst 6500 for throughput performance.

The best practice is to connect data centers via MPLS or VPLS as they provide active-active data center operation, disaster planning and load balancing attributes. VPLS is a layer 2 service, therefore, connected VPLS data centers deliver a LAN-like service over campus and/or wide area networks. Layer 2 connectivity is important as server-server communications expect layer 2 connectivity as most applications have been designed with this assumption. For connecting more than two data centers, VPLS offers mesh connectivity. In essence, data centers connected via VPLS look and act as if they are on the same LAN. Therefore, we test that VPLS throughput performance rates are equally high performance in this scenario as MPLS.

One of the most impressive network design options available on the Catalyst 6500 is the use of VSS. Connecting two Catalyst 6500s equipped with Sup2Ts creates a virtual switch, adding each switch's performance while operating as a single switch thus eliminating spanning tree in favor for active-active links. We configure two Catalyst 6500s via VSS. We measure throughput performance to verify that VSS throughput rates are equally high performance as the MPLS and VPLS scenarios.

System Network Test Configuration

To test MPLS/VPLS and VSS throughput performance, we populate two Catalyst 6500 WS-C6513-Es with eight 10GbE ports each via 6908-10G modules connected directly to Ixia test equipment. The Catalyst 6500s are connected via 8 x 10G Distributed EtherChannels. This configuration created a full end-end 80Gbs path, traversing the Ixia test gear to DUT1 to DUT 2 to Ixia test gear. See graphic. RFC 2544 test script is run in the Ixia test equipment to measure

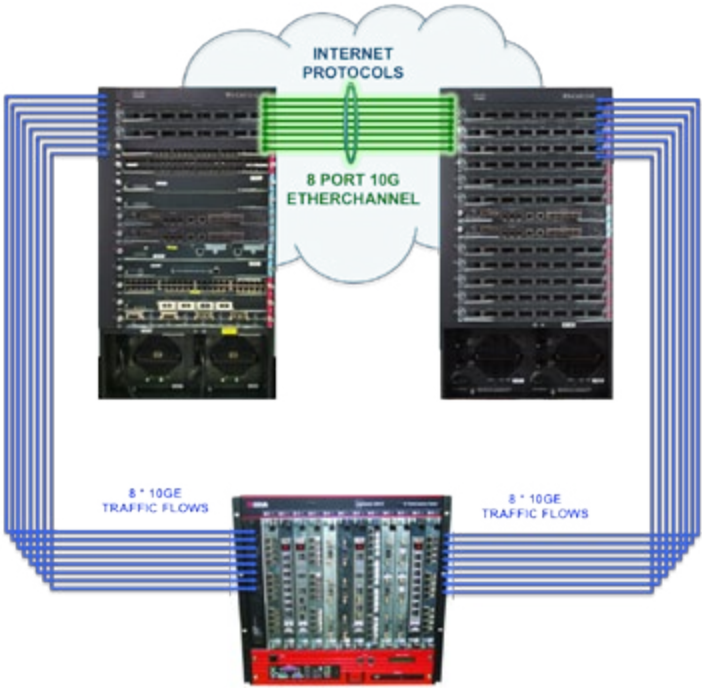

throughput performance in millions of pps. This configuration demonstrated the ability of the Catalyst 6500 Sup2T to channel eight individual 10GbE ports into a single 80Gps path across multiple 6908-10G modules, a first. For the VSS configuration, an eight-port Virtual Switch Link or VSL was configured between the two Catalyst 6500s. See figure.

The traffic profile for these set of MPLS/VPLS/VSS test required a full-mesh combination, where each of the 8 x 10GbE ports on each DUT were assigned different IP endpoints. All endpoints sent data to all other endpoints. This traffic profile is more typical of real-world networks and was necessary to load-balance traffic evenly across the 80G EtherChannel.

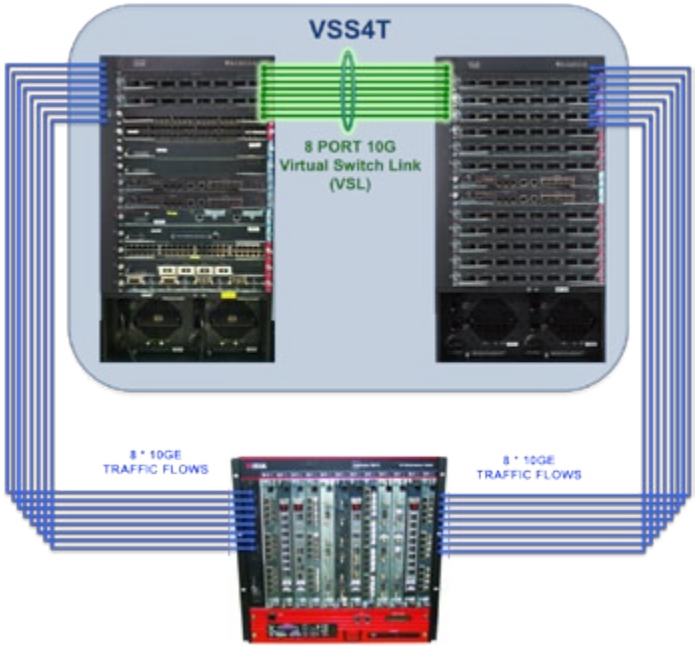

MPLS Results: For MPLS, the Catalyst 6500 Sup2T demonstrated throughput performance that ranged from 32.2 Mpps to 3.8 Mpps for packet sizes that varied between 512 bytes to jumbo size or 4096 bytes.

Catalyst 6500 Sup2T MPLS via RFC2544

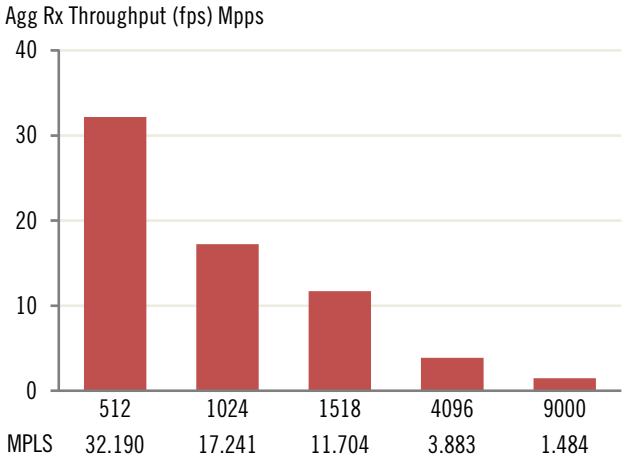

VPLS Results: For VPLS, the Catalyst 6500 Sup2T demonstrated throughput performance that ranged from 33.8 Mpps to 2.4 Mpps for packet sizes that varied between 512 bytes to jumbo size or 4096 bytes.

Catalyst 6500 Sup2T VPLS via RFC2544

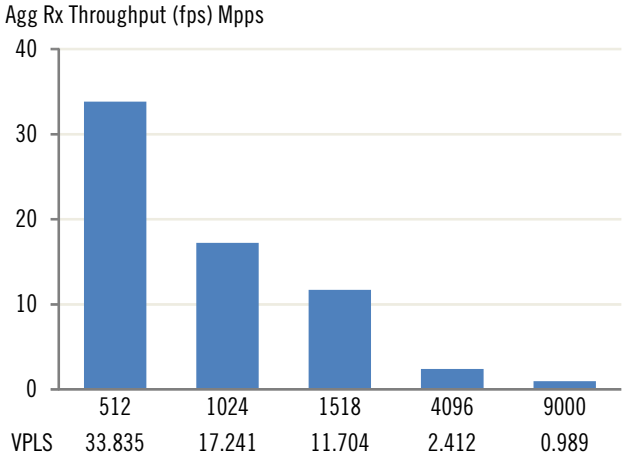

VSS Results: For VSS, the Catalyst 6500 Sup2T demonstrated throughput performance that ranged from 35.3 Mpps to 4.4 Mpps for packet sizes that varied between 512 bytes to jumbo size or 4096 bytes.

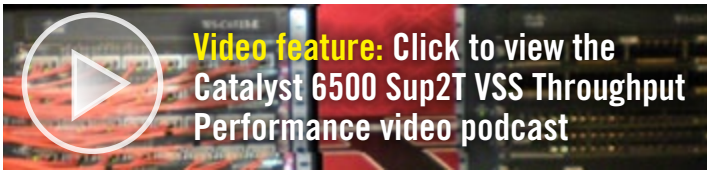

Catalyst 6500 Sup2T VSS via RFC2544

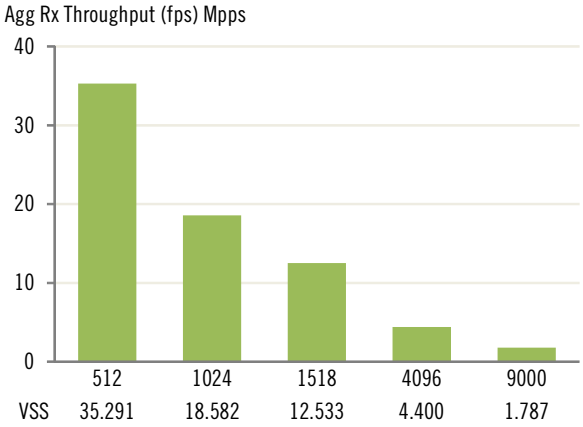

The above data suggest that throughput performance is fairly consistent independent upon protocol that being MPLS, VPLS and VSS. A contributing factor to the differences in throughput is found in different headers associated for each protocol. This result could not occur in the older generation of Catalyst 6500 with Sup720 with its 40Gbs per module backplane access speed.

Catalyst 6500 Sup2T MPLS/VPLS/VSS via RFC2544

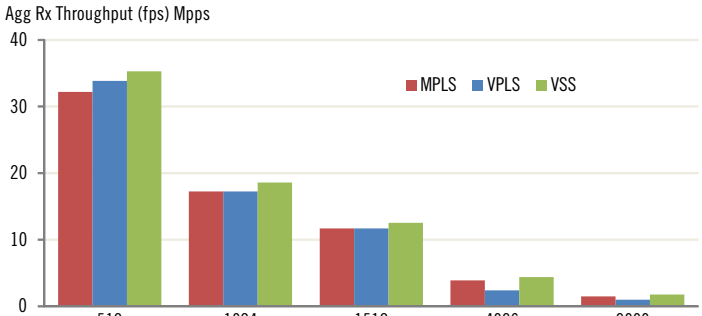

	512	1024	1518	4096	9000
MPLS	32.190	17.241	11.704	3.883	1.484
VPLS	33.835	17.241	11.704	2.412	0.989
VSS	35.291	18.582	12.533	4.400	1.787

Network Encryption with 802.1ae MACSec

With the MPLS/VPLS two Catalyst 6500s configuration, we tested performance for 802.1ae MACSec to verify that there was no throughput performance degradation when encryption was enabled minus the additional 16 Byte overhead of 802.1ae keys. MACSec encryption has become increasingly popular and important to campus network design, but previous switch performance degraded when forwarding encrypted traffic. Here we verify that the Catalyst 6500 does not suffer throughput performance degradation while MACSec traffic is being forwarded.

We tested the Catalyst 6500 via the cPacket Networks cTap Smart 10G passive probe to verify traffic flows were either MACsec encrypted or unencrypted. We found that there is no material difference in throughput performance, other than 802.1ae encryption key overhead, thanks to 16 additional bytes per packet. The cPacket smart passive probe also measured line rate throughput performance.

The cPacket Networks passively tapped a 10GbE link between the Cisco Catalyst 6500s. The cPacket passive probe allows one to visually inspect packets before and after enabling of MACsec to confirm that the Layer 2 payload is encrypted. We configure the Ixia test equipment to run traffic including certain IP address, port number and payload content. MACsec will obfuscate header fields of Layer 3 and above plus payload content, which serves as a marker. By inspecting these header field values and payload content marker, the passive probe determines if a packet is encrypted. In addition and independently of visual inspection, the passive probe sets real time counters, at the tapped link for the specific payload pattern that is used as a marker. Those counters match the Ixia counters when MACsec is disabled and the packets are unencrypted, but will remain zero when MACsec is enabled and the marker is obfuscated. This confirms that all Ixia traffic being forwarding to the Catalyst 6500 has been encrypted when it flows between the two Cisco Catalyst 6500s.

Results: Thanks to the Cisco 6908-10Gline cards supporting in line encryption, there is no performance penalty in the Catalyst 6500s' performance when MACsec is enabled. Therefore, we recommend that IT architects consider encrypting their backbone links as there is no additional dollar and performance cost associated with its use, other than the 16 Byte key overhead.

Conclusion

This test of the Catalyst 6500 is the most extensive to date. We found that upgrading the Catalyst 6500 from Sup720 to Sup2T was straightforward and added significant value in the areas of MACsec encryption, improved ACL capabilities and IPv4/IPv6/MPLS/VPLS/VSS throughput performance. In addition, we found that the Sup2T supported existing service models, such as Network Analysis (NAM), Wireless (WiSM), Application Control Engine (ACE20), Firewall Service Module (FWSM) plus 6148A-GE, 6148E-GE with POE/POE+, 6724-SFP line cards plus 6704 and 6716 line cards after a trivial DFC3 to DFC4 daughter card swap. We found that line cards can be swapped and upgraded while the Sup2T is operational, avoiding off-hour scheduled downtime. In addition, we found that existing interface transceivers SFP and X2 being used in a Sup720 Catalyst 6500 can be reused with the Sup2T. Finally, we found that Sup720 IOS configurations may be copied and migrated to a Sup2T via a flash drive successfully upon boot up.

Much of the throughput performance advantages and scale of network services is due to custom ASICs resident in the Sup2T, 6908-10Gline cards and DFC4 daughter cards. We were particularly impressed with the ease of upgrade, the new ACL dry-run and atomic commit plus MACsec performance. While not visible via this report, the Cisco Prime LAN Management Solution or LMS was used to configure the Catalyst 6500 for each test. We found this management system to be graphically intuitive and complete. For more information, click here.

For existing customers of Cisco's Catalyst 6500 Sup720, we anticipate upgrade experiences similar, if not simpler, than ours as this test was conducted under tight time constraints with limited resources. It's no wonder why the Catalyst 6500 is so popular as it offers a wide variety of network design options such as MPLS/VPLS/VSS. With the new upgrade to Sup2T and supporting line cards we verify that throughput performance doubles over the Sup720 for IPv6, IP Multicast, MPLS/VLPS and VSS.

Terms of Use

This document is provided to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools, the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the vendors that are beyond our control to verify. Among these is that the software/ hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is," and Lippis Enterprises, Inc. (Lippis), gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein.

By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting, directly or indirectly, from any information or material available on it. Lippis Enterprises, Inc., is not responsible for, and you agree to hold Lippis Enterprises, Inc., and its related affiliates, harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Lippis Enterprises, Inc., makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from www.lippisreport.com.

No part of any document may be reproduced, in whole or in part, without the specific written permission of Lippis Enterprises, Inc. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services, which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.

About Nick Lippis

Nicholas J. Lippis III is a world-renowned authority on advanced IP networks, communications and their benefits to business objectives. He is the publisher of the Lippis Report, a resource for network and IT business decision makers to which over 35,000 executive IT business leaders subscribe. Its Lippis Report podcasts have been downloaded over 160,000 times; ITunes reports that listeners also download the Wall Street Journal's Money Matters, Business Week's Climbing the Ladder, The Economist and The Harvard Business Review's IdeaCast. Mr. Lippis is currently working with clients to design their private and public virtualized data center cloud computing network architectures to reap maximum business value and outcome.

He has advised numerous Global 2000 firms on network architecture, design, implementation, vendor selection and budgeting, with clients including Barclays Bank, Eastman Kodak Company, Federal Deposit Insurance Corporation (FDIC), Hughes Aerospace, Liberty Mutual, Schering-Plough, Camp Dresser McKee, the state of Alaska, Microsoft, Kaiser Permanente, Sprint, Worldcom, Cigitel, Cisco Systems, Hewlett Packet, IBM, Avaya and many others. He works exclusively with CIOs and their direct reports. Mr. Lippis possesses a unique perspective of market forces and trends occurring within the computer networking industry derived from his experience with both supply and demand side clients.

Mr. Lippis received the prestigious Boston University College of Engineering Alumni award for advancing the profession. He has been named one of the top 40 most powerful and influential people in the networking industry by Network World. TechTarget, an industry on-line publication, has named him a network design guru while Network Computing Magazine has called him a star IT guru.

Mr. Lippis founded Strategic Networks Consulting, Inc., a well-respected and influential computer networking industry-consulting concern, which was purchased by Softbank/Ziff-Davis in 1996. He is a frequent keynote speaker at industry events and is widely quoted in the business and industry press. He serves on the Dean of Boston University's College of Engineering Board of Advisors as well as many start-up venture firms' advisory boards. He delivered the commencement speech to Boston University College of Engineering graduates in 2007. Mr. Lippis received his Bachelor of Science in Electrical Engineering and his Master of Science in Systems Engineering from Boston University. His Masters' thesis work included selected technical courses and advisors from Massachusetts Institute of Technology on optical communications and computing.