

Cisco Catalyst 6500 with Supervisor720 — 10 Gigabit Ethernet Performance Test

Introduction

In March 2003, Cisco Systems announced the next phase in the evolution of its market-leading Catalyst 6500 switching system, providing improved price/performance and switching capacity, plus new features such as high density 10 Gigabit Ethernet and IPv6 forwarding in hardware.

These enhancements are based around a new 720Gbps Supervisor Module, which combines a 720 Gbps switch fabric and management processors on a single card; enabling up to 40 Gbps (80 Gbps full-duplex) backplane capacity per card slot. New high density, high performance, interface cards such as the WS-X6704-10GE 4-Port 10 Gigabit Ethernet card, priced at just \$27,500 (includes WS-F6500-DFC3a @ \$7,500) US List Price, drive the price/performance of the Catalyst 6500 to over 400 Mpps (million packets per second) whilst protecting customers' investment in existing Catalyst interface cards by ensuring complete interoperability with next-generation cards installed in the same chassis.

In July 2003, Cisco Systems commissioned European Advanced Network Test Center (EANTC) to independently validate the performance of the upgraded Catalyst 6500 in a series of Gigabit and 10 Gigabit Ethernet tests designed to confirm Cisco's performance claims for the enhanced system.



Test Setup

The 10 Gigabit Ethernet (10GE) tests used a Catalyst 6509 chassis with eight WS-X6704-10GE 4-port 10 Gigabit Ethernet cards to provide a 32-Port 10GE test bed.

Venue & Test Equipment

The tests were conducted at Spirent's independent "SmartLab" in Calabasas, California; they employed Spirent's XLW-3721A 10-Gigabit Ethernet cards for the Smartbits traffic generator.

The test configuration comprised 8 x WS-X6704-10GE 10 Gigabit Ethernet Cards, all cards were equipped with optional WS-F6700-DFC3a Distributed Cisco Express Forwarding (dCEF) daughter cards.

Overall Test Highlights
Cisco demonstrated full investment protection for existing customers, including support of 1st, 2nd & 3rd generation interface cards in the same chassis as wire-rate 10GE.
Catalyst demonstrated up to 351 million packets per second (Mpps) IPv4 aggregate performance.
No performance degradation with 10,000-line access control lists.
Catalyst demonstrated >187 Mpps IPv6 aggregate performance.

This test report details the findings of the 10 Gigabit Ethernet tests. A separate report »Catalyst 6500 with Supervisor720 — Gigabit Ethernet Performance Test« contains the results for Gigabit Ethernet tests conducted at the same time.



Supervisor720 Upgrade & Investment Protection Test

Test Highlights

- Upgrading to Supervisor720 increases the number of available card slots due to its integrated switch fabric design. Two card slots previously occupied by Supervisor2 modules are recovered in a fully redundant switch.
- Catalyst demonstrated IPv4 any-to-any connectivity amongst 1st, 2nd and 3rd generation interface cards in the same chassis. Confirming all interface cards can interoperate and that the Catalyst chassis backplane can automatically adjust to the different card types.
- Cisco demonstrated that all generations of line cards supported IPv6 forwarding in hardware, including a 1st generation WS-X6248-RJ-45 card, first shipped in 1999.
- Cisco demonstrated that the presence of 1st and 2nd generation interface cards do not affect the performance of third-generation line cards such as the WS-X6704-10GE 4-Port 10GE card.

- 1st & 2nd generation interface cards can benefit from new "third generation" features such as IPv6 in hardware, MPLS etc. This prolongs the investment in Catalyst interface cards purchased as far back as 1999.

Test Methodology


The test was split into four stages:

1. Confirm any-to-any IPv4 (Layer 3) connectivity between existing interface cards, in a chassis configured with current-generation Supervisor2 and 256Gbps Switch Fabric Modules. Send full-mesh traffic at 100Mbps on each port. (4 ports per card in test)
2. Replace the Supervisor2 and 256Gbps Switch Fabric Modules with the new Supervisor720 and add third generation interface cards into the test. Demonstrate any-to-any IPv4 connectivity between cards from all generations as in Stage #1.
3. Run full-mesh traffic test on 4 x 10GE ports to confirm Cisco's claim that 3rd generation interface card performance is not degraded by the presence of older interface cards in the same chassis.
4. Test IPv6 any-to-any Layer 3 forwarding between cards from all generations to assess Cisco's claim regarding new Supervisor720 features being available on 1st and 2nd generation interface cards. Full mesh IPv6 traffic was sent at 100Mbps on each port. (4 ports per card in the test)

Test Objectives

To simulate the upgrade of an existing Catalyst 6500 with Supervisor2 and 256Gbps Switch Fabric Modules to the new 720Gbps Supervisor720 and assess Cisco's claims regarding investment protection for its existing customers:

- Cisco is the only major switch vendor to support high-density 10GE in the same chassis as existing interface cards. (Other vendors either require all interface cards to be replaced or do not support existing cards in the same chassis as high-density 10GE cards).
- 1st & 2nd generation interface cards do not degrade the performance of new generation line cards installed in the same chassis.



Classic-1 First-generation, 32Gbps bus-based card

Classic-3 Third-generation, 32Gbps bus-based card for wiring closet applications

CEF256-2 Second-generation "Fabric Enabled" card without distributed forwarding card (DFC)

dCEF256-2 Second-generation "Fabric Enabled" card with distributed forwarding card (DFC)

Bold italic: Third-generation, 720Gbps fabric enabled cards with distributed forwarding card (DFC3)

Test Results

Test Stage	Result
Stage #1	Zero packet loss confirming IPv4 any-to-any connectivity between 1st and 2nd generation cards.
Stage #2	Zero packet loss confirming IPv4 any-to-any connectivity between 1st, 2nd and 3rd generations of interface cards.
Stage #3	Confirmed no performance degradation. 10GE performance was in line with results obtained in later tests, even though 1st and 2nd generation interface cards were present and passing data in the same chassis.
Stage #4	Confirmed that all generations of interface cards were able to forward IPv6 traffic.

EANTC Test Analysis

This test confirms Cisco's investment protection claims, offering a path for Cisco's customers to enhance the performance and functionality of their Catalyst 6500's, while maintaining full interoperability with their existing 1st and 2nd generation interface cards.

We were very impressed by the ability of first generation WS-X6248-RJ45 interface cards, (no longer sold by Cisco) to suddenly support IPv6 Layer 3 forwarding in hardware. IPv6 wasn't even a standard when these cards first shipped, yet by upgrading to the Supervisor720, Cisco demonstrated how these legacy cards could be given a new lease on life.

This test highlights the benefits of Cisco's unique hybrid Centralized/Distributed forwarding architecture and wide variety of interface card types in protecting customer's investment in Catalyst 6500 products.

To our knowledge, none of the other major switch vendors have provided an upgrade path to high-density 10GE and IPv6 whilst allowing older generations of cards in the same chassis.

Single Card Maximum Forwarding Rate Test

Test Highlights

- *The WS-X6704-10GE card exceeded Cisco's performance claims at 64 byte packet size, forwarding 48.28 million packets per second (Mpps).*

Test Objective

This test is designed to confirm the performance claims made by Cisco for the third-generation WS-X6704-10GE, four port 10GE interface card introduced alongside the Supervisor720 module.

Test Methodology

All ports on the card under test were connected to the SmartBits and IPv4 traffic was sent at 100% load using 64-byte frames for 60 seconds. The aggregate forwarding rate of the card was recorded and compared to Cisco's performance claims to verify these were accurate.

Test Results

Product Code	Description	Performance at 64-byte packet size	
		Claimed	Actual
WS-X6704-10GE	4-Port 10 Gigabit Fiber Ethernet Card	48 Mpps	48.28 Mpps

EANTC Test Analysis

Cisco's performance claims were found to be accurate and the card slightly exceeded Cisco's claimed performance number.

40 Gbps (80 Gbps full-duplex) Card Slot Capacity Test

Test Highlights

- *When internal tagging is taken into consideration, the Catalyst demonstrated it could make full use of the 40Gbps switch fabric channel capacity per card slot when populated with the WS-X6704-10GE, four port 10 Gigabit Ethernet card.*
- *Catalyst demonstrated support for "Jumbo Frames" over 10 Gigabit Ethernet which is important for data center applications.*

Test Objective

Assess Cisco's claim that the upgraded Catalyst 6500 with Supervisor720 can support up to 40 Gbps (80 Gbps full-duplex) backplane/switch-fabric bandwidth per card slot.

Test Methodology

Two WS-X6704-10GE cards were used in this test. Each card had 4 x 10GE ports connected to the

SmartBits, thereby generating a maximum of 40 Gbps of traffic in each direction. SmartFlow was used to send traffic at 100% load using 512, 1504, 6016 and 8992 byte packet sizes in a partial mesh between the two cards. All traffic flowed across the backplane/switch-fabric.

Test Parameter Settings	
Test Application	SmartFlow, 100 hosts per port
Traffic Pattern	Partial Mesh, Bidirectional Traffic
Packet Size	512, 1504, 6016, 8992 bytes (Add 32-byte internal tag to each packet size)
Traffic Rate	100%
Test duration	60 seconds

After compensating for the effect of a 32-byte internal tag that is applied to each packet before it is forwarded across the switch fabric, we confirmed the effective Gbps capacity of the fabric channels serving each card slot using the calculations below.

Test Results & Calculations

Packet Size		Total Packets Received in 60 seconds	Total Fabric Channel per Slot ^a
net	w/Tag		
512	544	1,070,470,144	39.679 Gbps
1518	1550	383,034,715	39.887 Gbps
6016	6048	98,938,820	39.971 Gbps
8992 ^b	9024	66,369,641	39.981 Gbps

- Calculated as follows: (net packet size + 32 byte tag + 12 byte inter-frame gap) * total packets received * 8 bits / 60 seconds / 2 directions
- The Catalyst's maximum MTU is 9,216 bytes. The test was run at 8,992 bytes to remain within the 9,016 byte maximum MTU size of the SmartBits test equipment.

EANTC Test Analysis

This test demonstrates that the new WS-X6704-10GE, 4-Port Ten Gigabit Ethernet cards utilized 79.962 Gbps of the 80 Gbps full-duplex fabric channel capacity per slot. Each of the two fabric channels serving the slot contributed 39.981 Gbps to this aggregate figure.

Additional tests with the maximum Catalyst 6509 chassis capacity saw no degradation in a card's ability to use the full fabric channel capacity per slot, demonstrating that this card slot capacity is available on all slots simultaneously in the Catalyst 6509.

WS-X6704-10GE 2-Port Maximum Forwarding Rate Test

Test Highlights

- The Catalyst provided totally non-blocking, zero-loss performance at all packet sizes tested when subjected to 100% traffic load using two ports of each 4-port 10GE card.

Test Objective

Cisco Systems asked to include this test after requests from prospective customers who were interested in using just two of the four interfaces on the card to provide a totally non-blocking 10GE environment.

Test Methodology

Eight WS-X6704-10GE cards are used in the test, each with just two of the available four ports on each card in use. A total of 16 x 10GE ports were used.

Traffic at 100% load is passed between the two cards using various packet sizes, including Jumbo Frames.

Test Parameter Settings	
Test Application	SmartFlow, 100 hosts per port
Traffic Pattern	Partial Mesh, Bidirectional Traffic
Packet Size	64, 512, 1504, 6016, 8992 bytes
Traffic Rate	100%
Test duration	60 seconds

Test Results

Packet Size	Packets Transmitted (100% Load)	Maximum Forwarding Rate (pps)	Packet Loss
64	14,285,713,920	238,095,232	0%
512	2,255,639,040	37,593,984	0%
1518	780,234,048	13,003,901	0%
6016	198,807,152	3,313,453	0%
8992 ^a	133,155,792	2,219,263	0%

- Note: The Catalysts maximum MTU is 9,216 bytes. The test was run at 8,992 bytes to remain within the 9,016 byte maximum MTU size of the SmartBits test equipment.

EANTC Test Analysis

Using two ports of each WS-X6704-10GE card, the Catalyst demonstrated zero-loss, wire-rate performance at all packet sizes tested.

L3 Switching Latency Tests

Test Objectives

This test is designed to measure the average and maximum latency for IPv4 and IPv6 traffic being Layer 3 switched by the Catalyst 6500.

Low latency is important for time-sensitive applications such as real-time video conferencing and IP telephony. High latency can adversely affect end-to-end performance.

Test Methodology

The tests were run using 64-byte packets sent at 90% load for 120 seconds.

Test Results

Product Code	WS-X6704-10GE
Description	4-Port 10 Gigabit Ethernet Card
IPv4 Average Latency	10.743 μ s
IPv4 Max Latency	11.400 μ s
IPv6 Average Latency	11.076 μ s
IPv6 Max Latency	12.300 μ s

EANTC Test Analysis

The Catalyst demonstrated remarkably consistent latency under all tests, with very little difference between IPv4 and IPv6 latencies. This is remarkable given the more difficult packet parsing required by IPv6 forwarding decisions. These results show the Catalyst is highly suited to time-sensitive applications.

Understanding the 32-Port Ten Gigabit Ethernet Tests (Please read first)


Catalyst backplane switch fabric overview

The Catalyst 6509 was equipped with eight WS-X6704-10GE, 4-port 10 Gigabit Ethernet cards, resulting in a total of 32 interfaces. They were connected with a Spirent SmartBits stack of 32 Ten Gigabit Ethernet test ports.

In the case of the 4-port WS-X6704-10GE, each pair of 10 Gigabit interfaces is connected to one of two 20Gbps (40 Gbps full-duplex) backplane switch fabric channels serving the card slot. Before a packet is forwarded across the backplane, an internal 32 byte tag is added to aid internal switching. If all four ports on a card tried to forward traffic across the backplane switch fabric simultaneously, the four ports

would generate more than 40 Gbps of traffic, because each frame has the additional 32 byte internal tag associated with it.

Under such conditions some packets loss would occur.



Characteristics of full-mesh traffic patterns and their effect on this test

As a full-mesh traffic pattern cycles through the many thousands of source and destination port combinations (Layer 3 IP flows), there will inevitably come a moment in time (a particular combination of source and destination ports) when all the ports on one interface card will try to talk to all the ports on the other interface card. During this brief period all traffic must contend for the switch fabric channel's limited capacity, resulting in packet loss.

As the 10 GigE performance tests called for us to measure the maximum forwarding rate of the switch with zero or negligible packet loss, we had to lower the SmartBits transmit-rate to a point where no loss occurred across the fabric channel.

This represents the worst-case test scenario for the Catalyst and artificially lowers the aggregate zero-loss performance to a rate that doesn't oversubscribe the switch fabric channels. The switch fabric channel becomes the sole factor that determines the result.

The best-case conditions for the Catalyst is where all the traffic stays local to the ports served by a particular switching engine and no traffic at all passes over the switch fabric; thereby avoiding the oversubscription problem outlined above. Under these conditions, the Catalyst will post its best results and these results will represent the maximum forwarding rate of the switching engines.

The overhead of the internal tag and the effect of interframe gap means the oversubscription rate on the switch fabric is just 1.08:1, so a small amount of traffic remaining local to the card would relieve the congestion on the backplane switch fabric channel and deliver the Catalyst's best-case performance figure.

Worst-Case and Best-Case test scenarios

To provide the reader with the most accurate performance information possible, we ran each 10 Gigabit Ethernet test in two test scenarios. One representing the worst case scenario, where all traffic contends for the backplane capacity, then again in the best-case scenario where all traffic is switched between ports on the same card, exercising the distributed switching engines to their maximum capacity.

The Catalyst's maximum forwarding rate in a real network will fall somewhere between these two figures. If sufficient traffic remains local to the card, so that the backplane is not oversubscribed, the aggregate forwarding rate will be the same as the best-case results.

32-Port Ten Gigabit Ethernet — (L2) Maximum Forwarding Rate Test

Test Highlights

- Catalyst demonstrated a maximum L2 forwarding rate with zero-loss of up to 351 Mpps @ 64-byte packet size, confirming Cisco's performance claims for the switch.

Test Objectives

Test the maximum forwarding rate of the Catalyst 6500 equipped with 32 Ten Gigabit Ethernet ports when Layer 2 (Ethernet) switching within a single 32-port VLAN.

Test Methodology

Traffic was sent in a full mesh traffic pattern to determine the aggregate forwarding rate of the switch as a whole. A second test was then run where traffic remained local to each card, testing the aggregate switching capacity of the distributed switching engines on each card.

Test Parameter Settings	
Test Application	SmartFlow, 100 hosts per port, frame-loss test
Number of L2 Addresses	3,200
Traffic Pattern	Full Mesh, Bidirectional Traffic
Packet Sizes	64 bytes
Traffic Rate	Maximum-rate per card type to achieve zero or negligible loss
Test duration	60 seconds

Test Results

Packet Size	Maximum Layer 2 Forwarding Rate	
	Best-Case Traffic Pattern	Worst-Case Traffic Pattern
64 Byte	351.26 Mpps	316.67 Mpps
Loss	0.00%	0.00%

EANTC Test Analysis

In real networks, where a portion of the traffic will stay local to the card, the maximum performance figure for 64 byte packets will typically be nearer the best-case result, as only a small amount of traffic must be switched locally between ports on the same card to prevent oversubscription of the backplane switch fabric channel.

Of course, if customers wanted a totally non-blocking environment, they could employ just two of the four ports available on the card. In this configuration the Catalyst delivered totally non-blocking wire-rate performance at 64 byte packet size.

IPv4 Baseline (L3) Maximum Forwarding Rate Test

Test Highlights

- Catalyst demonstrated a maximum IPv4 Layer 3 forwarding rate of up to 351 Mpps @ 64-byte packet size.
- Catalyst performance was very consistent, posting exactly the same results as achieved when Layer 2 switching.

Test Objectives

Assess the maximum L3 forwarding rate with zero or negligible loss for IPv4 traffic. Record both worst-case and best-case performance characteristics for the switch under test.

Test Methodology

Each port on the Catalyst was configured as a unique subnet and the test traffic simulated a remote routing scenario where the Catalyst had to forward traffic between remote subnets advertised via OSPF.

SmartBits TeraRouting was configured to emulate a next-hop OSPF router attached to each of the 32 interfaces.

Due to limited time and a large number of tests in the test plan, the Catalyst was tested using only the most stressful minimum-size 64-byte packet size. This provides an accurate assessment of the switches performance under the most stressful traffic conditions.

In each test scenario, transmit rates were adjusted until we achieved maximum forwarding rates with negligible or zero packet loss. The worst-case performance result is artificially limited by the fabric channel capacity as explained in the “Understanding the 32-Port 10 Gigabit Ethernet Tests” section above.

Test Parameter Settings	
Test Application	TeraRouting. Use OSPF to advertise 1 subnet per port, emulate 200 hosts per subnet
Traffic Pattern	Full Mesh, Bidirectional Traffic
Packet Size	64 bytes
Traffic Rate	Maximum-rate per card type to achieve zero or negligible loss.
Test duration	60 seconds

Test Results

Packet Size	Maximum Layer 3 Forwarding Rate	
	Best-Case Traffic Pattern	Worst-Case Traffic Pattern
64 Byte	351.26 Mpps	316.67 Mpps
Loss	0.00%	0.00%

EANTC Test Analysis

The Catalyst posted absolutely identical results to the previous L2 switching test. This demonstrates the Catalyst maintains a very consistent baseline forwarding performance whether forwarding is based on L2 or L3 addresses.

IPv4 Maximum Forwarding Rate with Services Test

Test Highlights

- Catalyst performance unaffected by the addition of value-added services such as ACLs, QoS traffic classification and NetFlow Statistics gathering.

Test Objectives

Assess the impact on the maximum L3 forwarding-rate of the Catalyst 6500 when multiple services are also configured on each port.

Test Methodology

The Catalyst was configured with the following services prior to the start of the test:

- 10,000-entry access-control list denying a non-repeating, non-sequential set of TCP port numbers, applied as both an incoming and outgoing list on each of the 32 Ten Gigabit Ethernet interfaces. This causes the Catalyst to do a double-lookup for each and every packet traversing the switch, once on ingress and once on egress.
- 500-entry QoS traffic classification filter applied to all interfaces and in operation at the same time as the 10,000-entry security ACL.
- NetFlow statistics gathering per-flow statistics on each of the 32 Ten Gigabit interfaces. This was active at the same time as the Security and QoS ACLs.

Test Parameter Settings	
Test Application	TeraRouting. Use OSPF to advertise 1 subnet per port, emulate 200 hosts per subnet.
Traffic Pattern	Full Mesh, Bidirectional Traffic
Packet Size	64 bytes
Traffic Rate	Maximum-rate per card type to achieve zero or negligible loss.
Test duration	60 seconds

The test was run in a set of stages designed to confirm all services were simultaneously active during the tests.

1. Send traffic that doesn't match any of the DENY statements in the 10,000-entry ACL. This forces the switch to compare each packet against the full ACL list, finally matching a PERMIT-ALL entry in the 10,001st ACL entry.

- To confirm ACLs are active, send traffic that will match the 10,000th ACL DENY statement from 25% of the flows. This should result in 25% of the overall traffic being dropped.
- Use the switch command line interface to confirm that QoS and Netflow statistics gathering are active during the tests.

Due to lack of time, only the worst-case traffic scenario was tested.

Test Results

Test Stage	Maximum Layer 3 With Services Forwarding Rate
	Worst-Case Traffic Pattern
IPv4 max-fwd with no services from previous test	316.67 Mpps 0.00% loss
#1 – No Traffic Matches ACL	315.10 Mpps
	0.00% Loss
#2 – 25% of Traffic Matches 10,000th ACL DENY statement	0.00 Mbps 25% Loss
#3 – Confirm QoS and Netflow statistics are active during test	Confirmed

EANTC Test Analysis

This test was only run using the worst-case traffic pattern due to time constraints. The test demonstrated the Catalyst 6500's performance is not affected even when very large numbers of services are configured on the switch. The Catalyst's performance only reduced by 1.5 Mpps when compared to the worst-case results from the previous 10GE maximum forwarding rate test without services.

Overall the Catalyst produced comparable forwarding-rates to the previous baseline L2 and L3 forwarding rate tests, even though substantial numbers of ACLs, QoS and NetFlow Statistics were active on the switch.

PIM-SM Multicast Scalability and Maximum Forwarding-Rate Test

Test Highlights

- Catalyst 6500 demonstrates support for 25,000 PIM-SM s,g mroutes.
- Catalyst 6500 forwards 195.83 Mpps of multicast traffic with zero loss.

Test Objectives

PIM Sparse Mode (PIM-SM) is the most widely deployed multicast routing protocol in the industry. To provide optimal multicast routing, PIM-SM employs source,group multicast routes, commonly referred to as s,g mroutes.

PIM-SM allows the recipients of the multicast group to find the shortest path back to the multicast source, thereby reducing network latency to a minimum. If a multicast group has multiple sources sending traffic to it (quite common in financial trading rooms etc.), it must keep a separate mroute for each source,group pair.

The result is that the amount of multicast mroutes the switch must support can quickly run into the thousands.


In independently validated tests conducted last year, the Catalyst 6500 with Supervisor2 and 256Gbps Switch Fabric demonstrated support for a maximum of 12,000 s,g mroutes using GigE interfaces.

The objective of this test is to show that the Catalyst 6500 running PIM-SM can support up to 30,000 s,g mroutes as claimed by Cisco.

Test Methodology

The test used Spirent SmartFlow 3.0 to generate very large numbers of s,g mroutes.

For each of the two sender ports, we allocated 14 receiver ports on different cards that would request to join all the multicast groups associated with the sender. A total of 2 such Multicast Transmission Groups were configured as shown below.



This topology doubles the overall rate at which multicast traffic must be replicated and retains a relatively large number of receiver ports to increase the number of packet replications needed.

Test Parameter Settings	
Test Application	SmartFlow 3.0
Traffic Pattern	Multicast, unidirectional. 2 x Multicast Transmission Groups comprising 1 sender and 14 receivers. Each sender emulates 50 multicast sources and transmits to 250 unique multicast groups. This test generates 25,000 s,g mroutes.
Packet Size	64, 1024 bytes
Traffic Rate	Maximum rate to achieve zero or negligible loss.
Test duration	100 seconds
Record	Number of s,g mroutes in the test, the total number of Outgoing Interface List (OIL) entries needing to be maintained and the maximum forwarding rate achieved.

Test Results

Packet Size	64 bytes	1024 bytes
Number of PIM-SM s,g mroutes in the test	25,000	25,000
Number of OIL entries	375,000	375,000
Maximum Forwarding Rate (Mpps)	195,83 Mpps	31,85 Mpps
% Loss	0	0

EANTC Test Analysis

Cisco demonstrated large-scale multicast scalability in this test, doubling the number of PIM-SM mroutes supported when compared to similar tests run on the Supervisor2 and 256Gbps Switch Fabric models.

The maximum number of mroutes that could be supported by the SmartBits limited the test to just 25,000 s,g mroutes. Cisco claim the Catalyst could have gone on to support up to 30,000 s,g mroutes although this claim could not be tested.

Support for 25,000 s,g mroutes is a very impressive result and the Catalyst did not appear to be stressed by this test.

Bidirectional PIM Multicast Scalability and Maximum Forwarding Rate Test

Test Highlights

- Catalyst 6500 demonstrates support for the equivalent of 200,000 PIM-SM s,g mroutes
- Catalyst 6500 forwards 195.83 Mpps of multicast traffic with zero loss.

Test Objectives

Bidirectional PIM (PIM-BiDir) is a new multicast routing protocol, currently an IETF internet draft (work in progress). The protocol has been developed to overcome scalability issues with PIM Sparse Mode (PIM-SM), resulting from each mroute comprising a source,group address pair. In validated tests conducted last year, the Catalyst 6500 with Supervisor2 and 256Gbps Switch Fabric demonstrated support for a maximum of 12,000 s,g mroutes.

The objective of this test is to show that the Catalyst 6500 running PIM-BiDir can support the equivalent of hundreds of thousands of PIM-SM s,g mroutes using this new highly scalable multicast routing protocol.

Test Methodology

The test uses SmartFlow 3.0 to generate very large numbers of s,g mroutes.

For each sender port, we allocated 14 receiver ports on different cards that would request to join all the multicast groups associated with the sender. A total of two such Multicast Transmission Groups were configured, the same as in the previous PIM-SM test.

Test Parameter Settings	
Test Application	SmartFlow 3.0
Traffic Pattern	Multicast, unidirectional. 2 x Multicast Transmission Groups comprising 1 sender and 14 receivers. Each sender emulates 200 multicast sources and transmits to 1,000 unique multicast groups. Equivalent of 200,000 s,g mroutes.
Packet Size	64, 1024 bytes
Traffic Rate	Maximum rate to achieve zero or negligible loss.
Test duration	100 seconds
Record	Equivalent number of s,g mroutes simulated in the test, the total number of Outgoing Interface List (OIL) entries needing to be maintained and the maximum forwarding rate achieved.

Each physical sender port was originally configured to emulate 250 multicast sources, however, this proved impossible, with the PCs driving the test applications running out of memory.

The test was scaled down and eventually generated traffic simulating a total of 200,000 s,g mroutes.

EANTC used the Catalyst's command line interface to confirm that the Catalyst had the correct amount of multicast state during the tests.

Test Results

Packet Size	64 bytes	1024 bytes
Equivalent Number of PIM-SM s,g mroutes in the test	200,000	200,000
Maximum Forwarding Rate (Mpps)	195,83 Mpps	31,85 Mpps
% Loss	0	0

EANTC Test Analysis

The test used 200 sources x 1,000 groups to generate the equivalent of 200,000 s,g mroutes.

As the number of multicast groups rather than the combination of source and group address pairs stresses PIM-BiDir scalability, we did a further test with 2,000 groups and a single source per group.

This was the maximum number of groups we could generate with the SmartBits in this configuration.

In this latter test the Catalyst posted exactly the same results as in the test results table above. Cisco claims the Catalyst will support up to 30,000 multicast groups, although we were unable to test this on this occasion.

In the final analysis, the Catalyst can easily provide massive multicast scalability when using PIM-BiDir and we doubt there are many applications that would require the equivalent of even 200,000 s,g mroutes.

IPv6 Baseline (L3) Maximum Forwarding Rate Test

Test Highlights

- Catalyst demonstrates a maximum forwarding rate of 187.5 Mpps using 76-byte packets in a worst-case traffic pattern.

Test Objectives

Test Cisco's claim that the Catalyst 6500 with Supervisor720 can forward IPv6 traffic in hardware up to 200 Mpps at 64 byte packet size.

Due to lack of time, it was only possible to test the Catalysts IPv6 maximum forwarding rate in the worst-case traffic pattern. The SmartBits minimum packet size with IPv6 is 76 bytes, so we were unable to confirm Cisco's performance claim of 200 Mpps at 64 byte packet size.

Test Methodology

The test used the same methodology as the IPv4 Baseline Maximum Forwarding Rate Test with the exception that due to the larger IPv6 packet header and the insertion of SmartBits special tag into the data area of each packet, the minimum IPv6 packet size that can be produced is 76 bytes.

Test Parameter Settings	
Test Application	SmartFlow3.0, emulate single host per port (SmartBits limitation)
Traffic Pattern	Full Mesh, Bidirectional Traffic
Packet Size	76 bytes (Minimum IPv6 Packet Size)
Traffic Rate	Maximum rate to achieve zero or negligible loss
Test duration	60 seconds

Test Results

Packet Size	Maximum IPv6 Forwarding Rate	
	Best-Case Traffic Pattern	Worst-Case Traffic Pattern
76 Byte	not tested	187.50 Mpps
Loss		0.00%

EANTC Test Analysis

Even in the worst-case traffic scenario, the Catalyst put in a respectable 187.50 Mpps performance with 76-byte packets.

It is unlikely that customers will exceed the IPv6 forwarding capabilities of the Catalyst 6500 for many years, as the amount of native IPv6 traffic in most networks today is very limited.

IPv6 Maximum Forwarding Rate with Services Test

Test Highlights

- Catalyst demonstrated no IPv6 performance degradation when services were enabled.
- Catalyst exhibited difficulties reloading the 10,000-entry ACL after reboot when running the 2nd stage of the test.

Test Objectives

Test Cisco's claim that the addition of significant levels of value-added services will not affect the performance of the Catalysts IPv6 maximum forwarding rate.

Test Methodology

The Catalyst was configured with the following services prior to the start of the test:

- 10,000-entry access-control list denying a non-repeating, non-sequential set of TCP port numbers, applied as both an incoming and outgoing list on each of the 32 Ten Gigabit Ethernet interfaces. This causes the Catalyst to do a double-lookup for each and every packet traversing the switch, once on ingress and once on egress.
- Netflow statistics, gathering IPv6 per-flow statistics on each of the 32 Ten Gigabit interfaces. This was active at the same time as the Security ACLs.

Test Parameter Settings	
Test Application	SmartFlow3.0, emulate single host per port (SmartBits limitation)
Traffic Pattern	Full Mesh, Bidirectional Traffic
Packet Size	76 bytes (Minimum IPv6 Packet Size)
Traffic Rate	Maximum rate to achieve zero or negligible packet loss
Test duration	60 seconds

The test was run in a set of stages designed to confirm all services were simultaneously active during the tests.

1. Send traffic that doesn't match any of the DENY statements in the 10,000-entry IPv6 ACL. This forces the switch to compare each packet against the full ACL list, finally matching a PERMIT-ALL entry in the 10,001st ACL entry.

2. To confirm ACLs are active, send traffic that matches the 10,000th ACL DENY statement from 25% of the flows. 25% of the overall traffic should be dropped, confirm other 75% of the traffic is unaffected.
3. Use the switch command line interface to confirm that Security ACLs and NetFlow statistics gathering are active during the tests.

Test Results

Test Stage	Maximum IPv6 L3 With Services Forwarding Rate
	Worst-Case Traffic Pattern
IPv6 max-fwd with no services from baseline test	187.50 Mpps 0.00% loss
#1 – No Traffic Matches ACL	187.50 Mpps
	0.005% Loss
#2 – FMTG ^{a1} Traffic Matches 10,000th ACL DENY statement	Test failed, Catalyst was unable to reload the 10,000 ACL entries after reboot
#3 – Confirm security ACLs and Netflow statistics are both active during test	Confirmed

a. Fully Meshed Traffic Group

EANTC Test Analysis

In Stage #1 of the test the Catalyst displayed no degradation in performance when compared to the previous maximum IPv6 forwarding-rate test without services. Even though a 10,000-entry IPv6 access control list was configured to check each packet on both ingress and egress to/from the switch performance didn't drop by even one packet.

In Stage #2, the SmartBits was re-configured to send traffic that matched the 10,000th deny statement in the ACL. Unfortunately, after we rebooted the switch, we had difficulty reloading the 10,000-entry ACL and had to abandon the test due to lack of time to investigate further. According to Cisco this was due to a bug with one of the pre-production cards we were using.

After the tests Cisco investigated the problem further at Cisco and have fixed it in the production cards.

Conclusion

The Catalyst 6500 equipped with the new Supervisor720 module and populated with third generation Ten Gigabit Ethernet cards achieved or in many cases exceeded Cisco's performance claims for the switch.

Performance was unaffected by very significant quantities of value added features and the Catalyst demonstrated massive IP Multicast scalability.

These tests re-affirm that the Catalyst 6500 remains one of the most potent switches on the market and we believe that based on the results of these tests it will remain a popular choice for both enterprise and service provider customers.

About Cisco Systems

Cisco Systems, Inc. (Nasdaq:CSCO) is the worldwide leader in networking for the Internet. News and information are available at <http://www.cisco.com>.

Cisco, Cisco Systems and the Cisco Systems logo are registered trademarks of Cisco Systems, Inc. in the US and certain other countries. All other trademarks mentioned in this document are the property of their respective owners.

For further information, please e-mail sup720-test-questions@external.cisco.com.

About Spirent Communications



Spirent Communications is a worldwide provider of integrated performance analysis and service assurance systems for next-generation network technologies. Spirent's solutions enable customers to develop and deploy network equipment and services more economically and efficiently by emulating real-world conditions in the lab and assuring end-to-end performance of large-scale networks.

Spirent Communications is a wholly owned business group of Spirent plc, an international network technology company. Spirent, Spirent Communications and the Spirent logo are trademarks of Spirent plc.

<http://www.spirentcom.com/>

About EANTC



The European Advanced Networking Test Center (EANTC) offers vendor neutral network test services for manufacturers, service providers and enterprise customers. Primary business areas include interoperability, conformance and performance testing for IP, ATM, MPLS, and broadband voice related network technologies and applications.

EANTC AG
Einsteinufer 17, 10587 Berlin, Germany
Gabriele Schrenk, Managing Director
schrenk@eantc.com
<http://www.eantc.com/>