

Catalyst 6500 with Supervisor720 — 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 720Gbps 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-X6748-GE-TX, 48-Port 10/100/1000 Ethernet interface card, priced at just \$22,500 US List Price, drive the price/performance of the Catalyst 6500 to over 400Mpps whilst protecting customers' investment in existing Catalyst interface cards by ensuring complete interoperability with next-generation cards installed in the same chassis.

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

Gigabit tests used the Catalyst 6513 chassis equipped with 410 ports, making this test the industry's largest Gigabit Ethernet test to date, confirming Cisco's claim to provide the highest GigE port density in a single-chassis solution.


The test configuration comprised 7 x WS-X6724-SFP Gigabit Ethernet Cards, plus 5 x WS-X6748-GE-TX 10/100/1000 Ethernet Cards, and all cards were equipped with optional WS-F6700-DFC3a Distributed Cisco Express Forwarding (dCEF) daughter cards.

Venue & Test Equipment

The tests were conducted at Spirent's independent "SmartLab" in Calabasas, California and employed Spirent's new 4-port dual-media LAN-3325A Gigabit Ethernet cards, facilitating a mix of copper 10/100/1000 and fiber Gigabit interfaces within the test.

Overall Test Highlights
Cisco demonstrated full investment protection for existing customers, including support of 1st, 2nd & 3rd generation interface cards in the same chassis.
Catalyst demonstrated up to 407 million packets per second (Mpps) IPv4 aggregate performance.
No performance degradation with 10,000 line access control lists.
Catalyst demonstrated up to 230 Mpps IPv6 aggregate performance.
Support for the equivalent of 425,000 S,G mroutes when using Bidirectional PIM multicast routing.

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



EANTC extensively tested Cisco's Catalyst 6500 next generation line cards. They showed the expected performance values, confirming Cisco's claims for switching speed, scalability and investment protection.

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-RJ45 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. This prolongs the investment in Catalyst interface cards purchased as far back as 1999.

Test Methodology

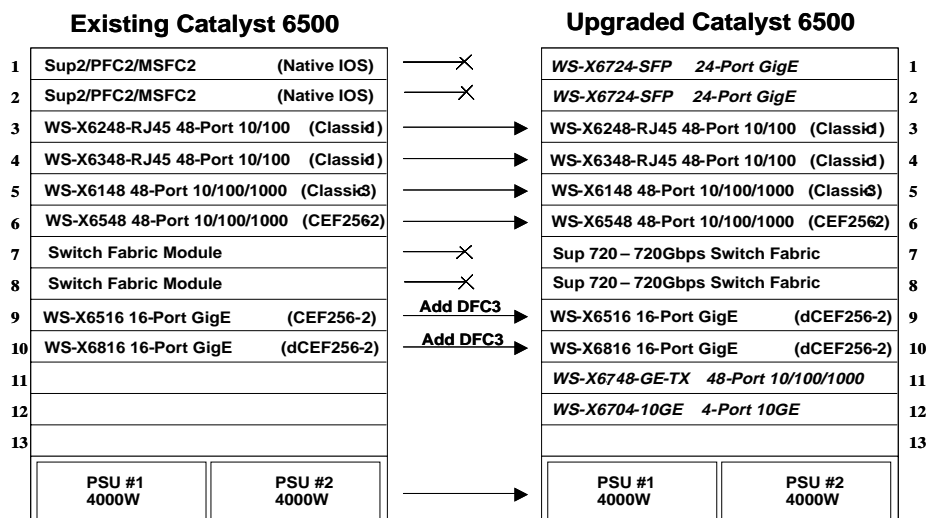
The test was split into four stages:

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 old 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 cards 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

- Each card slightly exceeded Cisco's performance claims.

Test Objectives

This test was designed to confirm the performance claims made by Cisco for each of the new third-generation interface cards introduced alongside the Supervisor720 module.

Test Methodology

All ports on the card under test were connected to the SmartBits and tested with IPv4 data streams using SmartFlow for 60 seconds. Traffic was sent at 100 % load using 64-byte frames and the aggregate forwarding rate of the card was recorded.

Test Results

Product Code	Description	Performance @ 64-byte packet size	
		Claimed	Actual
WS-X6724-SFP	24-Port Gigabit Fiber Ethernet Card	24 Mpps	26.04 Mpps
WS-X6748-GE-TX	48-Port 10/100/1000 Copper 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.

10/100/1000 Cable Fault Detection Test

Test Highlights

- TDR built into the 10/100/1000 48-port line card accurately detected wire breaks, open circuit pair and short circuit pair errors at various points along a 97 meter Category 5 cable.

Test Objectives

Cisco claims to have integrated Time Domain Reflectometry (TDR) into the ports of its latest generation of copper 10/100/1000 interface cards. At the time of testing, only the WS-X6548-GE-TX and WS-X6148-GE-TX 48-Port 10/100/1000 interface cards support this feature. However, Cisco confirmed other line cards will be enhanced with this feature in the future.

Time Domain Reflectometry (TDR) is a test process able to detect a number of common faults on structured cabling installations. TDR can determine the cause of the fault and report approximately how far from the switch port the fault is located. This allows rapid fault location and repair without the need for expensive test equipment.

This test was designed to confirm the accuracy of the TDR feature in detecting and reporting a number of common cable faults found in structured wiring systems.

At the time of writing, EANTC could only identify one other vendor that offered integrated TDR functionality in its switches.

Test Methodology

A special 97-meter cable was constructed with the ability to insert various fault conditions at 1-meter, 50-meter and 95-meter lengths from the switch port under test. The cable was attached to the Catalyst port at one end and a normal PC at the other.

Broken wires, open and short circuited twisted pairs were simulated and the results recorded.

Test Results

Fault Type	Distance from switch port	Fault as reported by TDR feature
Single broken wire in a twisted pair	1, 50 & 95 meters	Detected and accurately located
Open circuit twisted pair	1, 50 & 95 meters	Detected and accurately located
Short circuited twisted pair	1, 50 & 95 meters	Detected and accurately located
Detection of whether cable was terminated or not	N/A	TDR was able to detect whether the PC was attached or not and accurately recorded the overall cable length.

EANTC Test Analysis

We used a PC to terminate the 97-meter cable, as we wanted to assess whether the TDR function might damage end-user equipment connected to the line when the test is run. No damage occurred.

We believe this capability will be of major benefit to any customer with a structured cabling system. The TDR function proved highly accurate and its reporting format was easy to understand. This integrated feature can reduce the time to locate a cable fault and can be used without having to visit the wiring closet. This feature is included at no additional cost and could save an investment in cable test equipment.

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 40 Gbps switch fabric channel capacity per card slot.

Test Objectives

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

In the Gigabit Ethernet tests, only the WS-X6748-GE-TX, 48-Port GigE Card has sufficient ports to generate 40 Gbps across the backplane/switch-fabric. Each group of 24 ports on the card is assigned to one of two 20 Gbps (40 Gbps full-duplex) switch fabric channels on the backplane of the switch.

With only 24 ports and a single 20 Gbps fabric channel on the WS-X6724-SFP, this module cannot generate sufficient traffic to stress an 40 Gbps backplane connection. This module was therefore excluded from this test.

Test Methodology

Two WS-X6748-GE-TX cards were used in the test. Each card had 40 Gigabit Ethernet 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.

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.

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	1,070,453,302	39.678 Gbps
1504	1536	381,671,654	39.389 Gbps
6016	6048	35,762,224	38.688 Gbps
8992	9024	64,034,636	38.611 Gbps

a. Calculated as follows: (net packet size + 32 byte tag + 12 byte inter-frame gap) * total packets received * 8 bits / 60 seconds / 2 directions

EANTC Test Analysis

This test demonstrates that the new WS-X6748-GE-TX, 48-Port GigE cards utilized 79.356 Gbps of the 80 Gbps full-duplex fabric channel capacity per slot. Each of the two fabric channels serving the slot contributed 39.678 Gbps to this aggregate figure.

Additional tests with more cards saw no degradation in a slot's fabric channel capacity.

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 and high latency can adversely affect end to end performance.

Test Methodology

The tests were run using 64-byte packets sent at 90 % load over 20 ports for 120 seconds.

Test Results

Product Code		WS-X6724-SFP	WS-X6748-GE-TX
Description		24-Port GigE Card	48-Port 10/100/1000 Ethernet Card
IPv4 Latency	Average	12.234 μs	12.349 μs
	Max	16.400 μs	16.600 μs
IPv6 Latency	Average	12.449 μs	12.571 μs
	Max	16.600 μs	16.700 μs

EANTC Test Analysis

The Catalyst demonstrated remarkably consistent latency under all tests. 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 410-Port Gigabit Ethernet Tests (Please read first)

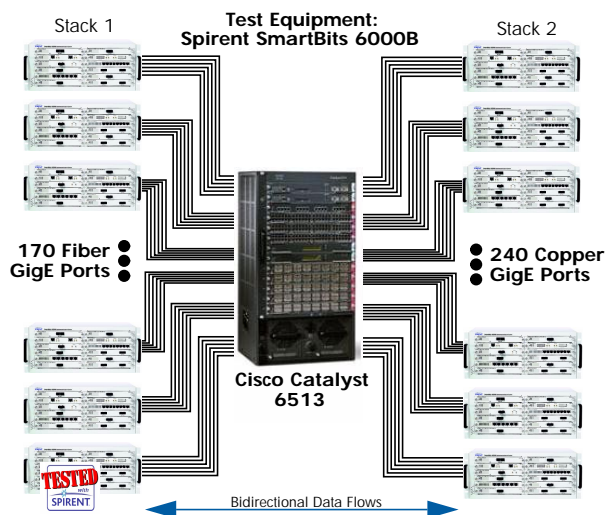
To understand the 410-port Gigabit Ethernet (GigE) tests it is first necessary to explain some of the techniques and terminology used in the test report.

The 410-port GigE tests in this report represent the industry's largest-ever Gigabit Ethernet test. No other vendor has ever been tested with such a large number of interfaces before.

Mixed media test bed

Cisco's 410-port test used a Catalyst 6513; thirteen slot chassis, equipped with 7 x WS-X6724-SFP 24-port fiber GigE interface cards, plus 5 x WS-X6748-GE-TX 48-port 10/100/1000 copper-based interface cards.

With the addition of the two fiber GigE ports on the Supervisor720 module, a total of 170 fiber ports were in the test, the remaining 240 ports being copper.



Fortunately, Spirent Communications had just released their new LAN 3325A 4-port mixed media card for the Smartbits at the same time and we were able to employ these in the test, reducing the number of SmartBits chassis needed and making the whole test a lot easier to run.

Two separate SmartBits stacks used

Due to the large number of ports involved in the test, two separate SmartBits stacks had to be used to perform the test. The test traffic from both SmartBits stacks was synchronized to start at exactly the same time using a SmartBits Automation Script developed by Spirent Communications, thereby ensuring that the switch under test had to handle the full traffic load for the full test duration.

Catalyst card types with slightly different performance characteristics

As the two Catalyst card types exhibit slightly different performance characteristics, SmartBits Stack#1 was connected exclusively to the WS-X6724-SFP cards in the top seven card slots, while SmartBits Stack #2 was connected to the WS-X6748-GE-TX cards in the lower five slots.

This allowed us to independently vary the traffic rate for each card type to achieve the maximum aggregate performance from the Catalyst. The results from each SmartBits stack were then combined to find the overall aggregate forwarding rate and packet loss for the switch as a whole.

Traffic patterns used in the tests

Other industry tests have used a large single full-mesh traffic pattern to test core switches like the Catalyst 6500.

Using such a traffic pattern in this test would mean each of the 410 GigE ports would send and receive traffic from every other GigE port in the test.

Firstly, it is not physically possible to test a single full-mesh comprising 410 ports, as the two SmartBits stacks discussed earlier cannot send traffic between each other.

Secondly, in EANTC's testing experience such enormous full mesh traffic patterns are not representative of the traffic flows found in typical enterprise networks, and we wanted to test the Catalyst under as realistic conditions as possible.

EANTC's experience of testing and working with large enterprise networks shows that application traffic follows the client-server model. Traffic tends to flow between groups of users (clients) and a subset of the servers in the server farm. For instance, the sales force is one such user-group and would tend to share information within the group and mainly use a specific subset of the servers in the data center. It is unlikely that the sales force would regularly access the Finance or R&D servers, but would make much use of the servers containing product information, pricing, forecasts etc.

Traffic tends to remain within these "communities of interest" and results in a series of Full Mesh Traffic Groups (FMTGs) within the network.

Catalyst backplane switch fabric overview

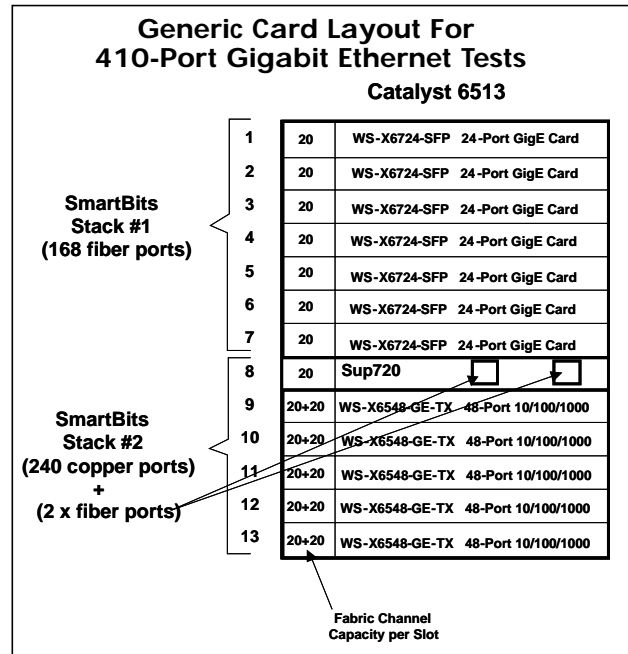
Traffic within each FMTG is a full-mesh between a total of 48 ports (50 ports in the case of FMTG #9). Each FMTG is spread across two interface cards, 24 ports on each card.

In the case of the Catalyst 6500, each group of 24 ports on a card is connected to a 20Gbps (40Gbps full-duplex) backplane switch fabric channel for inter-card traffic flows. If all 24 ports on one card tried to forward traffic across the backplane to the other card, the 24 ports would generate 24Gbps of traffic, which would have to contend for just 20 Gbps backplane capacity. This would result in some packet loss.

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 with all the ports on the other interface card. During this brief period all 24Gbps of traffic must contend for the 20Gbps switch fabric channel's limited capacity, resulting in packet loss.

As the 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 condition for the Catalyst is where all the traffic stays local to the ports served by a particular switching engine; 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 oversubscription rate on the switch fabric is just 1.2:1, so a small amount of traffic remaining local to the card would relieve the congestion on the backplane switch fabric channel.

Worst-Case and Best-Cast test scenarios

To provide the reader with the most accurate performance information possible, we ran each GigE 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. Only when >20 Gbps of traffic per 24-GigE port group attempts to flow across the backplane will the maximum-forwarding-rate-without-loss be lower than the best-case result.

410-Port GigE — (L2) Maximum Forwarding Rate Test

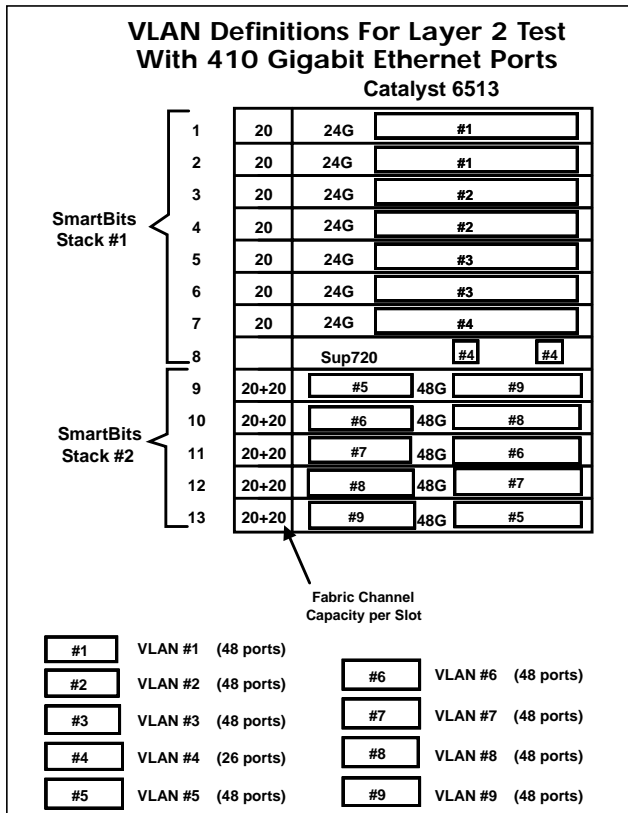
Test Highlights

- Catalyst demonstrated a maximum Layer 2 forwarding-rate with zero-loss of up to 407 Mpps at 64 byte packet size, confirming Cisco's performance claims.

Test Objectives

Test the maximum forwarding rate of the Catalyst 6500 equipped with 410 Gigabit Ethernet ports when Layer 2 switching within multiple VLANs. Measure the worst-case and best-case traffic scenarios.

Test Methodology



The diagram above shows how multiple Full-Mesh Traffic Groups (FMTGs) were configured and mapped to Ethernet VLANs. One Ethernet VLAN was configured for each

FMTG to simulate a typical VLAN deployment at an enterprise site.

Within each of the 48-port VLANs, traffic was sent in a full mesh traffic pattern to determine the aggregate forwarding rate of the switch as a whole.

Test Parameter Settings	
Test Application	SmartBits Automation Script, 100 hosts per port, frame-loss test
L2 Addresses	41,000
Traffic Pattern	Full Mesh within FMTG, Bidirect. 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 2 Forwarding Rate	
	Best-Case Traffic Pattern	Worst-Case Traffic Pattern
64 Byte	407.19 Mpps	355.97 Mpps
Loss	0.000 %	0.000 %

EANTC Test Analysis

In real networks, the Catalyst will reliably perform at maximum forwarding rates between the worst-case and best-case figures shown in the results table above.

In most cases the exact figure 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.

Be aware that many other vendors only show their best-case performance figures achieved with all flows locally switched and we applaud Cisco's decision to share both best-case and worst-case performance figures with the reader.

In best-case conditions, the Catalyst 6500 exceeded Cisco's performance claims of 400 Mpps for L2 forwarding and in all tests the forwarding rates achieved were done so with no packet loss.

IPv4 Baseline (L3) Maximum Forwarding Rate Test

Test Highlights

- Catalyst demonstrated a maximum IPv4 L3 forwarding-rate of up to 407Mpps @ 64-byte packet size, confirming Cisco's performance claims for the switch.
- Catalyst performance was very consistent, posting almost exactly the same results as achieved when L2 switching.
- The test routed packets between remote subnets advertised using OSPF and emulated 200 hosts per advertised subnet. Traffic was sent in a full host-mesh with a total of 16.8 million active L3 flows.

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.

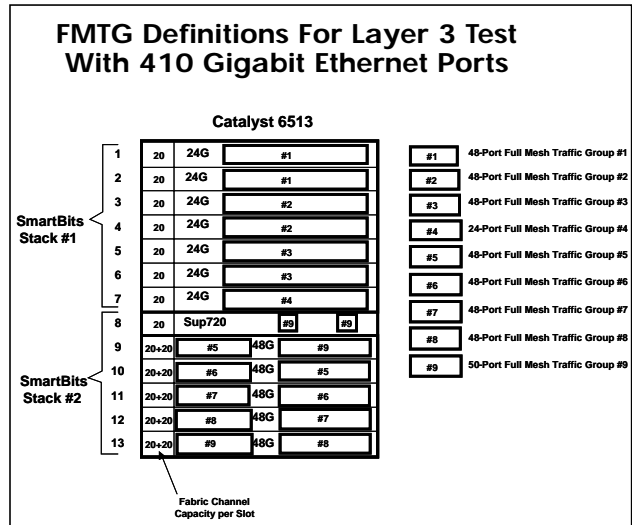
Test Methodology

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

Spirent's TeraRouting API was configured to emulate a next-hop router attached to each of the 410 interfaces and used OSPF to advertise a single remote subnet to the Catalyst. This required the Catalyst to support 410 OSPF neighbors, a notable achievement in its own right.

Next, a specially developed SmartBits Automation Script sent traffic that varied the source and destination IP address in each packet transmitted. The script cycled through 200 source host addresses per advertised subnet, with each source host sending to 200 hosts on each of 47 destination subnets within the FMTG. When multiplied by the 9 full-mesh traffic groups of 48 ports in the test, this generated 16.835 million source/destination L3 flows that the Catalyst had to handle.

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 applies most stress to the switch under test, and provides an accurate assessment of the switch's performance under the most stressful traffic conditions.



The diagram on this page shows the nine groups of 48-port Full Mesh Traffic Groups (FMTGs) used in this test and illustrates that traffic was switched both locally and across the fabric between cards. As with the previous L2 performance tests, we adopted two test configurations, the worst-case configuration shown in the diagram above and a best-case configuration where all traffic stayed local to each card.

In each test scenario, transmit rates were adjusted until we achieved maximum forwarding rates with zero or negligible packet loss. The worst-case performance result is artificially limited by the fabric channel capacity as explained in the "Understanding the 410-Port GigE Tests" section at the beginning of this test report.

Test Parameter Settings	
Test Application	TeraRouting API + Smartbits Automation Script. Use OSPF to advertise 1 subnet per port, emulate 200 hosts per subnet
Number of L3 flows	16.835 million
Traffic Pattern	Full Mesh within FMTG, 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	407.17 Mpps	355.97 Mpps
Loss	0.01 %	0.00 %

EANTC Test Analysis

The Catalyst posted virtually 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.

In the best-case test scenario, the Catalyst exceeded Cisco's performance claims of 400Mpps.

The Catalyst also handled a staggering 16.84 million source/destination L3 flows with ease.

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 Layer 3 forwarding rate of the Catalyst 6500 when multiple services are also configured on each port. Record both worst-case and best-case performance characteristics.

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 410 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 410 interfaces. This was active at the same time as the Security and QoS ACLs.

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.

2. To confirm ACLs are active, send traffic from one FMTG that matches the 10,000th ACL DENY statement. All traffic within this FMTG should be dropped.
3. Use the switch command line interface to confirm that QoS and Netflow statistics gathering are active during the tests.

Test Parameter Settings	
Test Application	TeraRouting API + Smartbits Automation Script. Use OSPF to advertise 1 subnet per port, emulate 200 hosts per subnet
Number of L3 flows	16.835 million
Traffic Pattern	Full Mesh within FMTG 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

Test Stage	Maximum Layer 3 With Services Forwarding Rate	
	Best-Case Traffic Pattern	Worst-Case Traffic Pattern
IPv4 max-fwd with no services from previous test	407.17 Mpps 0.01 % loss	355.97 Mpps 0.00% loss
#1 – No Traffic Matches ACL	407.17 Mpps 0.00 % Loss	355.65 Mpps 0.005 % Loss
#2 – FMTG1 Traffic Matches 10,000th ACL DENY statement	0.00 Mpps 100 % Loss	0.00 Mpps 100 % Loss
#3 – Confirm QoS and Netflow statistics are active during test	Confirmed	Confirmed

EANTC Test Analysis

The Catalyst produced identical 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.

Bidirectional PIM Multicast Scalability and Maximum Forwarding Rate Test

Test Highlights

- Catalyst 6500 demonstrated support for the equivalent of 425,000 PIM-SM s,g mroutes.
- Forwards >216 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). In validated tests conducted last year, the Catalyst 6500 with Supervisor2 and 256Gbps Switch Fabric demonstrated PIM-SM 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.

Test Methodology

Test Parameter Settings	
Test Application	Smartbits Automation Script (Multicast), maximum forwarding rate with zero or negligible loss
Traffic Pattern	Multicast, unidirectional. 17 x Multicast Transmission Groups comprising 1 sender and 23 receivers. Each sender emulates 50 multicast sources and transmits to 50 unique multicast groups. Equivalent of 425,000 s,g mroutes.
Packet Sizes	64 byte
Traffic Rate	Maximum rate to achieve zero or negligible loss.
Test duration	60 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.

The test uses a Smartbits Automation Script developed by Spirent Communications to generate very large numbers of s,g mroutes.

All 410 ports were used in the test, with 17 ports being designated multicast sender ports and the rest designated multicast receivers. For each sender port, we allocated 23 receiver ports on different cards that would request to join all the multicast groups associated with the sender. A total of 17 such Multicast Transmission Groups were configured as shown in the diagram on this page.

The Catalyst was tested using the most stressful minimum-size 64-byte packet size. This applies most stress to the switch under test, and provides an accurate assessment of the switch's PIM BiDir performance under the most stressful traffic conditions.

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

memory. The test was scaled down and generated traffic simulating a total of 425,000 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	Equivalent number of PIM-SM s,g mroutes	Number of OIL entries	Maximum Forwarding Rate
64 Byte	425,000	195,500	216,38 Mpps 0.00 % Loss

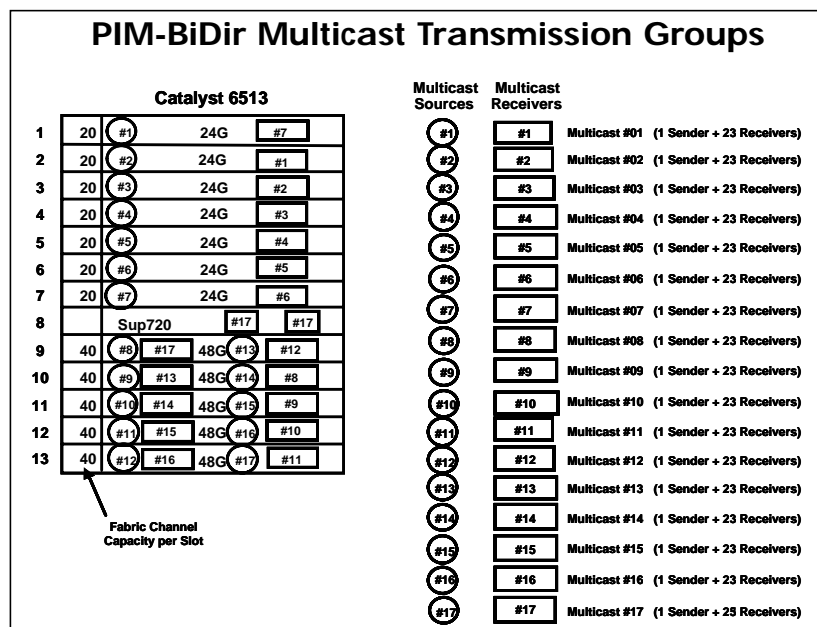
EANTC Test Analysis

Cisco was keen to demonstrate the Catalyst 6500's PIM-BiDir scalability limits, however the test applications simply couldn't generate the 2.125 Million s,g mroutes in the test plan.

The test had to be scaled down, but we believe the Catalyst could have scaled even further.

Over 8,500 multicast groups were generated in the test, however, Cisco claim the Catalyst can easily support up to 30,000 multicast groups and this would indeed have seen the Catalyst support the equivalent of >2 million s,g mroutes.

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 even 425,000 mroutes.



IPv6 Baseline (L3) Maximum Forwarding Rate Test

Test Highlights

- Catalyst demonstrated a maximum forwarding rate of >216 Mpps with just 0.015% loss.

Test Objectives

Verify Cisco's claim that the Catalyst 6500 with Supervisor720 can forward IPv6 traffic in hardware up to 200 Mpps.

Test Methodology

The same test configuration as in the IPv4 Baseline Maximum Forwarding-Rate Test was used. However, it was found that the SmartBits Automation Script was unable to emulate multiple hosts per port when running IPv6 traffic. In addition, due to the larger IPv6 packet header and the insertion of SmartBits' special tag into the data area of the packet, the minimum IPv6 packet size that can be produced is 76 bytes.

We ran out of time to run best-case and worst-case test scenarios. It was therefore decided only to proceed with the Worst-Case traffic scenario, recording the minimum performance customers can expect from the Catalyst.

Test Parameter Settings	
Test Application	SmartBits Automation Script, emulate single host per port (SmartBits limitation)
Traffic Pattern	Full Mesh within FMTG, Bidirect. Traffic
Packet Size	76 bytes (Minimum IPv6 Pkt-Size)
Traffic Rate	Maximum rate to achieve maximum forwarding-rate with zero or negligible packet loss.
Test duration	60 seconds

Test Results

Packet Size	Maximum Layer 3 Forwarding Rate	
	Best-Case Traffic Pattern	Worst-Case Traffic Pattern
76 Byte	- not tested -	229.69 Mpps
Loss		0.015 %

EANTC Test Analysis

Even in the worst-case traffic scenario as discussed in the "Understanding the GigE Performance Tests" section at the front of this document, the Catalyst put in a respectable 229.69 Mpps performance with 76-byte packets and exceeded Cisco's conservative 200 Mpps performance claim.

IPv6 Maximum Forwarding Rate with Services Test

Test Highlights

- Catalyst demonstrated no performance degradation when services were enabled.
- Catalyst exhibited problems when attempting to DENY IPv6 traffic matching the 10,000th IPv6 ACL statement. Traffic on some flows were still forwarded.

Test Objectives

Test Cisco's claim that the addition of significant levels of added-value services will not affect the performance of the Catalyst's 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 410 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 IPv6 per-flow statistics on each of the 410 interfaces. This was active at the same time as the Security and QoS ACLs.

Test Parameter Settings	
Test Application	SmartBits Automation Script, emulate single host per port (SmartBits limitation)
Traffic Pattern	Full Mesh within FMTG, Bidirectional Traffic
Packet Size	76 (Minimum IPv6 Pkt-Size) bytes
Traffic Rate	Maximum rate to achieve maximum forwarding-rate with 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 from one FMTG that matches the 10,000th ACL DENY statement. All traffic within this FMTG should be dropped, confirm other FMTG's traffic is unaffected.
3. Use the switch command line interface to confirm that QoS and Netflow statistics gathering are active during the tests.

Test Results

Test Stage	Maximum Layer 3 IPv6 With Services Forwarding Rate	
	Best-Case Traffic Pattern	Worst-Case Traffic Pattern
IPv6 max-fwd with no services from previous test	not tested	229.69 Mpps 0.015 % loss
#1 – No Traffic Matches ACL	not tested	229.70 Mpps 0.011 % Loss
#2 – FMTG1 Traffic Matches 10,000th ACL DENY statement	not tested	0.00 Mpps 100 % Loss
#3 – Confirm QoS and Netflow statistics are active during test	not tested	Confirmed

EANTC Test Analysis

In Stage #1 of the test the Catalyst displayed no degradation in performance, 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.

In Stage #2, the SmartBits was reconfigured to send traffic that matched the 10,000th deny statement in the ACL. Unfortunately, the Catalyst was unable to prevent all of this traffic from getting through, with some flows making it to the other side.

Numerous runs were conducted and on each run different flows managed to evade the ACL and get through. Cisco has raised a bug report and Cisco engineering is currently investigating the problem.

It should be mentioned that the tests were conducted using early engineering prototype cards and that this problem may not be present in later production versions of the card.

Conclusion

The Catalyst 6500 equipped with the new Supervisor720 and populated with 3rd generation 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/>