EANTC Independent Test Report

Cisco Catalyst 9000

Scalability and Forwarding Performance

January 2019

Introduction

Cisco commissioned EANTC to conduct an independent performance test of its new Catalyst 9000 family switches.

The switches are designed to work in enterprise campus environments. Cisco offers a range of form factors, port combinations and fixed or flexible modular configurations for specific use case scenarios. All of these switch types are part of the same family with regards to hardware and software design.

EANTC evaluated the following specific new Catalyst configurations:

- Catalyst 9300
- Catalyst 9300 Stackwise 480
- Catalyst 9400
- Catalyst 9500-24Q
- Catalyst 9500-40X

Cisco provided the IOS XE software release 16.6 for all tests. The verification was carried out at Cisco's headquarters in San Jose, California, in July and August 2018.

EANTC focused on the evaluation of Ethernet and IP performance using fully deployed configurations with up to one Terabit per second throughput and up to 392 ports. We evaluated unicast and multicast, IPv4 and IPv6, switching and routing scenarios that are typical for campus environments.

Keysight Technologies provided IxNetwork load generators for this project to enable the large-scale GigE, 10GigE and 40GigE test scenarios. The IxNetwork 8.40.1124.8 EA software version was used for all tests.

Test Highlights

- → Ethernet Forwarding: Catalyst 9300, Catalyst 9400, Catalyst 9500-24Q and Catalyst 9500-40X exhibited full line rate for fixed frame sizes and IMIX
- → Ethernet Forwarding: Catalyst 9300 Stackwise 480 exhibited full line rate for fixed frame sizes
- → In IPv4, IPv6 and mixed scenarios, the Catalyst 9300 and Catalyst 9400 exhibited full line rate performance for fixed frame sizes and IMIX
- → In IPv4, IPv6 and mixed scenarios, the Catalyst 9500-24Q and Catalyst 9500-40X exhibited full line rate performance for fixed frame sizes
- → The Catalyst 9300 Stackwise 480 exhibited line rate performance for fixed frame sizes both in IPv4-only and IPv6-only scenarios
- → Catalyst 9300 stackwise 480 successfully installed 8,000 IPv4 routes into the Forwarding Information Base (FIB)
- → Catalyst 9400 and Catalyst 9500-24Q successfully installed 65,000 IPv4 routes into the Forwarding Information Base (FIB)
- → Catalyst 9400 and Catalyst 9500-24Q successfully learned 63,000 MAC addresses
- → Catalyst 9300 stackwise 480 successfully learned 30,000 MAC addresses
- → Multicast Forwarding: Catalyst 9300 stackwise 480, Catalyst 9400- and Catalyst 9500-24Q exhibited full line rate for fixed frame sizes

Executive Summary

The Cisco Catalyst family of campus switches with routing capabilities has more than a decade of legacy in the Enterprise market. It is one of Cisco's most longstanding, largest and most successful product areas.

Enterprise campus networks are facing an imminent need to support ever-increasing bandwidth demand. They need to support the rapid growth of powerful endpoints that can deliver richer content such as HD video and wireless access points that deliver advanced wireless connectivity technologies such as 802.11ax. Switch interconnections are moving from 10GigE to 25GigE/40GigE/100GigE. More importantly, the demand for bandwidth is growing in the enterprises. Cloud services and streaming applications alike result in much higher throughput in the campus. Some industry sectors use multicast extensively. And even IPv6 becomes more widespread used in some regions.

That said, it would not be very smart to simply deploy the same type of switch everywhere. Port density, port type and performance requirements differ a lot across collapsed core, distribution and access switching locations. For this reason Cisco has asked us to evaluate five very different Catalyst configurations.

In all cases, Cisco uses the same hardware design principles, optimized for the respective deployment scenarios. As a result, our tests found that all five new Catalyst configurations we evaluated support full line rate or nearly full line rate in all test scenarios. The service scale of the backbone switches was impressive, as well as their total throughput supported. With the new Catalyst 9000 series solutions, Cisco continues to lead the industry. All standard Ethernet/IP forwarding scenarios are supported very well.

Test Methodology

This section introduces generic test methodology applicable to all Catalyst 9000 switches evaluated.

Our report includes the verification of throughput performance for IP packets (IPv4 and IPv6), Ethernet unicast and multicast frames. In addition, this report also contains scalability tests of the IP Forwarding Information Base (FIB) and Ethernet MAC table.

Both for IP and Ethernet layer forwarding performance tests, we determined the maximum throughput and latency based on RFC 2544 methodology for fixed packet sizes — 64, 128, 256, 512, 1024 and 1518 bytes. Additionally, we tested with Internet traffic mix (IMIX). IMIX is a packet mix that aims to replicate real Internet traffic. IMIX consists of 5 % 64-byte packets, 41 % 100-byte packets, 10 % 373-byte packets, 8 % 570-byte packets, 10 % 1300-byte packets, 25 % 1518-byte and 1 % 9000-byte packets.

For the IP forwarding test, we evaluated scenarios with IPv4-only traffic, IPv6-only traffic and a mix of IPv4 and IPv6 (50:50 ratio). In each case, 100 emulated endpoints per IP protocol and port were configured on the test equipment. We generated the bi-directional IP traffic and measured the IP packet forwarding performance metrics while transmitting the data traffic into the DUTs.

In the Ethernet unicast forwarding test, we configured 10 MAC addresses per port for Catalyst 9300 stackwise 480 and 100 MAC addresses per port for all other devices under test (DUTs) on the IxNetwork traffic generator.

In the Ethernet multicast forwarding test, we configured one source port and multiple receiver ports per DUT. We enabled IGMPv2 on all receiver ports in the emulator. Cisco Catalyst switches have IGMP snooping enabled by default. IGMP snooping constrains the flooding of multicast traffic by dynamically configuring Layer 2 interfaces so that multicast traffic is forwarded to only those interfaces associated with IP multicast devices. In our test, we configured the load generator to subscribe to all multicast groups on each of the ports.

Three Catalyst switch types participated in the multicast tests. Given their specific roles in the network, Cisco asked to configure them with specific numbers of multicast groups in total and per port as shown in the following table.

Once the receiver ports had successfully joined all multicast groups via IGMP, we generated Ethernet multicast traffic on the single source port for a number of fixed packet sizes — 64, 128, 256, 512, 1024, 1518 bytes. There was no IMIX multicast test.

DUT	Source Port	Destination Ports	IGMP Groups per Port
Catalyst 9300 Stack-wise 480	1x40 GigE	288x1GigE	1
		15x40GigE	40
Catalyst 9400	1x10 GigE	384x1GigE	1
		7x10GigE	10
Catalyst 9500-24Q	1x40 GigE	23x40GigE	40

Table 1: Multicast Setup

For the Forwarding Information Base (FIB) scalability test, we created four emulated OSPF routers on the tester. Each emulated router established MD5-authenticated OSPF sessions with the DUTs. We configured the maximum number of unique IPv4 prefixes on the emulated routers (based on Cisco's claims) and sent traffic to each route to verify that the DUTs are able to sustain as many forwarding (FIB) table entries.

To avoid route aggregation on the DUT, we used different prefix lengths and did not use consecutive prefixes. Instead, they were increased by a step of two. Before starting the test, we collected the number of prefixes present in the FIB table before (typically static, interface and management network routes). By subtracting these preexistent entries from the total number of FIB entries, we ensured that all emulated routes were indeed installed in the FIB. We sent test traffic at the rate of 800 Mbit/s via all advertised prefixes and expected no packet loss.

The FIB size test contained an additional step. Once all FIB entries had been installed and verified, we increased the number of /24 prefixes to exceed the FIB table limit. The DUT was expected to continue to operate reliably. It was also expected to raise an alarm indicating the FIB table exhaustion.

The MAC table scalability test was carried out using all physical links of each DUT. Untagged traffic on all ports was configured into the same single VLAN. While using one tester port to monitor if any flooded frames were received, we sent Ethernet frames carrying the maximum number of source MAC addresses from all other ports to determine the number of MAC addresses successfully learned by the DUT. Flooded or dropped frames were not expected during the test. The test was performed following RFC 2889 methodology.

The following sections explain specific results with each of the five devices under test, obtained using the generic methodology described above.

Results: Cisco Catalyst 9300

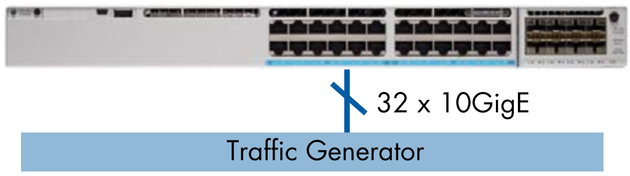

Figure 1: Catalyst 9300

Hardware Type	Software
Catalyst 9300	IOS XE 16.6.3

Table 2: Hardware Catalyst 9300

Hardware Type	Software
Catalyst 9300 Stackwise 480	
3 x C9300-48U	
2 x C9300-24P	IOS XE 16.6.3
1 x C9300-48UXM	
2 x C9300-24U	

Table 3: Hardware Catalyst 9300 Stackwise 480

Ethernet and IP Unicast Forwarding Performance

➔ Catalyst 9300 exhibited line rate performance for fixed frame sizes and IMIX

We tested the Catalyst 9300 in a configuration with 32 10GigE ports, fully deployed. All tests showed the expected throughput for fixed frame sizes and IMIX. No other tests were carried out with this DUT. The results of the Catalyst 9300 stacked configuration (next section) are applicable to the single unit Catalyst 9300 configuration as well.

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	320	5	7	8
128	320	5	7	8
256	320	5	7	8
512	320	5	7	8
1024	320	5	6	8
1518	320	5	6	8
IMIX	320	5	14	21

Table 4: Catalyst 9300 Forwarding Performance Ethernet

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	320	5	7	8
128	320	5	7	8
256	320	5	7	9
512	320	5	7	9
1024	320	5	7	10
1518	320	5	8	12
IMIX	320	5	14	21

Table 5: Catalyst 9300 Forwarding Performance IPv4

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
78	320	5	7	8
128	320	5	7	8
256	320	5	7	8
512	320	5	7	8
1024	320	5	7	8
1518	320	5	7	9
IMIX	320	5	14	21

Table 6: Catalyst 9300 Forwarding Performance IPv6

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
78	320	5	7	8
128	320	5	7	8
256	320	5	7	8
512	320	5	7	8
1024	320	5	6	8
1518	320	5	6	8
IMIX	320	5	14	21

Table 7: Catalyst 9300 Forwarding Performance Mix of IPv4 and IPv6

Results: Cisco Catalyst 9300 Stackwise 480

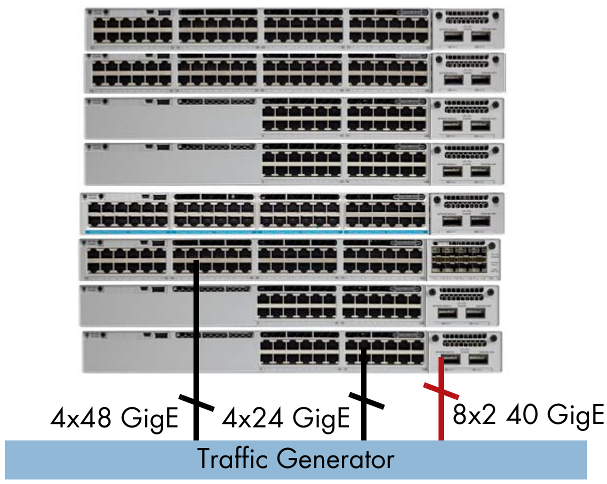

Figure 2: Catalyst 9300 Stackwise 480

We tested the Catalyst 9300 Stackwise 480 in a configuration of multiple systems that were stacked, resulting in a total of 288 GigE ports plus 16 40GigE ports.

Ethernet and IP Unicast Forwarding Performance

→ Catalyst 9300 Stackwise 480 exhibited line rate performance for fixed frame sizes

Stackwise 480 architecture builds high-speed 480Gbps per stack switch member in the stack ring. Together with Cisco, we designed the following traffic streams to simulate the real-world traffic patterns for campus environments when switches are stacked together:

- Traffic between downlink and uplink ports representing North-South traffic (80 Gbit/s bidirectional); this traffic leverages the stackwise-480 bandwidth
- Inter-switch traffic between the stack members representing East-West traffic (40 Gbit/s bidirectional); this traffic leverages the stackwise-480 bandwidth
- Intra-switch traffic within the stack members also representing East-West Traffic (344 Gbit/s bidirectional).

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	928	4	6	19
128	928	4	6	23
256	928	4	7	31
512	928	2	8	46
1024	928	4	9	24
1518	928	4	30	32
IMIX	927.5	3	78	104

Table 8: Catalyst 9300 stackwise 480 - Forwarding Performance Ethernet

Frame Size [Bytes]	928 Gbit/s Throughput, Latency [µs]				
	Min	Avg[a]	Max[a]	Avg[b]	Max[b]
64	4	12	17	41	78
128	4	12	16	66	127
256	4	12	15	114	224
512	2	14	15	211	420
1024	4	19	20	404	811
1518	4	24	25	Inconclusive	

a. 1 GigE downlinks to 40 GigE uplinks/1 GigE downlinks
b. 40 GigE uplinks to 1 GigE downlinks

Table 9: Catalyst 9300 stackwise 480 - Forwarding Performance IPv4

Frame Size [Bytes]	928 Gbit/s Throughput, Latency [µs]				
	Min	Avg[a]	Max[a]	Avg[b]	Max[b]
78	4	12	15	46	85
128	4	12	16	66	127
256	4	12	15	114	224
512	2	14	19	211	420
1024	17	18	19	404	811
1518	19	19	21	inconclusive	

a. 1 GigE downlinks to 40 GigE uplinks/1 GigE downlinks
b. 40 GigE uplinks to 1 GigE downlinks

Table 10: Catalyst 9300 stackwise 480 - Forwarding Performance IPv6

In the IP forwarding test, IMIX frames and mixed traffic types (IPv4 and IPv6) were excluded from this test.

In the Ethernet forwarding performance test for IMIX traffic, the Catalyst 9300 Stackwise 480 was able to forward 927.5 Gbit/s traffic — 0.5 Gbit/s less than expected (equivalent to 0.05% packet loss). The tremendous stack setup involved physical complexity with over 300 ports. Due to time limitations, we performed only frame loss tests (without binary search) with this hardware combination.

Ethernet Multicast Forwarding Performance

➔ Catalyst 9300 Stackwise 480 showed line rate Ethernet multicast forwarding performance for 40 groups with 303 receivers

The Catalyst 9300 Stackwise 480 successfully joined all ports 40 IGMP multicast groups on 303 ports as configured. Full line rate throughput performance was measured as expected.

DUT	Expected Throughput	Measured Throughput
Catalyst 9300 Stackwise 480	888 Gbit/s	888 Gbit/s

Table 11: Catalyst 9300 Stackwise 480 Multicast Forwarding Performance

Forwarding Information Base Scalability

➔ Catalyst 9300 Stackwise 480 successfully installed a total of 8,000 entries

In the FIB scalability test, Cisco asked us to verify that the Catalyst 9300 Stackwise-480 would be able to program 8,000 entries into the FIB table. The system was configured with static IPv4 and IPv6 routes (totaling 2,712 static entries) and we used IXIA to advertise additional 6,148 IPv4 routes via OSPF. The 2,712 static entries consisted of 914 IPv4 routes and 899 IPv6 routes. Cisco explained that each IPv6 route takes two entries in the FIB table.

As expected, the Catalyst 9300 Stackwise 480 learned successfully 8,105 entries.

Network Prefixes Count			
Expected	Initial	Advertised	Installed
8,000	914 IPv4, 899 IPv6[a]	6,148	8,105

a. Each IPv6 entry was expected to occupy double entry size in the FIB table

Table 12: Catalyst 9300 Stackwise 480 FIB Scalability

As expected, the Catalyst 9300 Stackwise 480 learned successfully 6,043 prefixes without any packet loss. The FIB population time was less than one minute.

The DUT remained stable after receiving the increased number of emulated routes exceeding the limit. The CPU utilization and memory utilization was still unchanged (CPU: 2% and memory: 31%). Via CLI, we observed the expected alarm message indicating that the FIB table was exhausted.

MAC Table Scalability

➔ Catalyst 9300 Stackwise 480 successfully learned 30,000 MAC addresses

When injecting 32,000 sequential MAC address (theoretical maximum size), we observed flooded frames. The Cisco engineer explained that the hash algorithm optimizes MAC capacity efficiency. The more random the MAC addresses, the higher the utilization of the table capacity. We moved the configuration that turned all addresses into unique random addresses. We configured the test tool to advertised 30,000 MAC address and the Catalyst 9300 stackwise 480 successfully learned all 30,000 MAC address without any frame flooding and reached 94% size utilization efficiency. However, we didn't have time to increase the number of MAC address to find the true maximum which is between 30,000 and 32,000."

DUT	MAC Addresses Count	
	Theoretical	Learned
Catalyst 9300 Stackwise 480	32,000	30,000

Table 13: Catalyst 9300 Stackwise 480 MAC Table Scalability

Results: Cisco Catalyst 9400

The Catalyst 9400 is a large-scale, modular campus switch featuring ten slots in total.

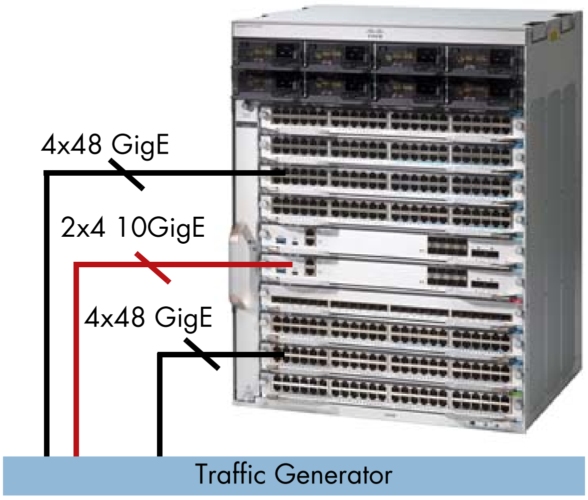

Figure 3: Catalyst 9400

We tested the Catalyst 9400 in a fully loaded configuration with eight 48-port GigE line cards resulting in a total of 384 GigE ports. Two supervisory cards were deployed, each carrying four 10GigE ports, resulting in eight 10GigE ports in total.

Hardware Type	Software
Catalyst 9400	IOS XE 16.6.3
2 x C9400-LC-48U	
6 x C9400-LC-48T	
2 x C9400-SUP-1XL	

Table 14: Hardware Catalyst 9400

Ethernet and IP Unicast Forwarding Performance

➔ Catalyst 9400 exhibited line rate performance for fixed frame sizes and IMIX

Cisco explained that supervisory card has central ASICs (Application-Specific Integrated Circuit) servicing all its port. Together with Cisco, we designed traffic streams aiming to utilized all the ASICs and the backplane as well:

- 80 Gbit/s bidirectional North-South flows in a 1-to-1 port scenario.
- 192 Gbit/s bidirectional East-West flows in a 1-to-1 port scenario.

For all three test scenarios (IPv4 only, IPv6 only and a mix of IPv4 & IPv6) and all the frame sizes including the IMIX frames, the throughput was in full line speed without any packet loss.

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	464	4	11	13
128	464	4	10	12
256	464	4	11	13
512	464	4	13	15
1024	464	4	14	18
1518	464	4	16	22
IMIX	464	4	100	133

Table 15: Catalyst 9400 Forwarding Performance Ethernet

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	464	4	11	14
128	464	4	10	13
256	464	4	11	14
512	464	4	13	16
1024	464	4	14	20
1518	464	4	17	23
IMIX	464	4	100	134

Table 16: Catalyst 9400 Forwarding Performance IPv4

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
78	464	4	11	14
128	464	4	10	13
256	464	4	11	14
512	464	4	13	17
1024	464	4	14	19
1518	464	4	17	23
IMIX	464	4	100	132

Table 17: Catalyst 9400 Forwarding Performance IPv6

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
78	464	4	11	13
128	464	4	10	13
256	464	4	11	15
512	464	4	13	16
1024	464	4	14	20
1518	464	4	17	23
IMIX	464	4	100	134

Table 18: Catalyst 9400 Forwarding Performance Mix of IPv4 and IPv6

Ethernet Multicast Forwarding Performance

→ Catalyst 9400 showed line rate performance for 1 source and 391 receivers over 10 groups with all frame sizes

The Catalyst 9400 showed that 10 IGMP multicast groups were successfully joined as configured and full line rate throughput performance was measured as expected.

DUT	Expected Throughput	Measured Throughput
Catalyst 9400	454 Gbit/s	454 Gbit/s

Table 19: Catalyst 9400 Multicast Forwarding Performance

Forwarding Information Base Scalability

→ Catalyst 9400 successfully installed 65,000 IPv4 routes into the FIB

The Catalyst 9400 successfully learned 65,000 prefixes without any packet loss. The FIB population took less than one minute.

When we increased the number of advertised route beyond the known limit, the switch did not learn the additional routes anymore as expected.

Network Prefixes Count			
Expected	Initial	Advertised	Installed
64,000	258	65,000	65,258

Table 20: Catalyst 9400 FIB Scalability

At 65,240 prefixes advertise, the switch started losing traffic associated with two prefixes as expected. During the whole test, CPU load and memory utilization were stable (CPU: 24 %, memory: 23 %). Via CLI, we observed the expected alarm message indicating that the FIB table was exhausted.

MAC Table Scalability

→ Catalyst 9400 successfully learned 63,000 MAC address

Cisco explained that the theoretical maximum size of the MAC address table would be 64,000 entries. Before reaching this limit, we observed flooded frames. We reduced the number of MAC addresses to 63,000 subsequently. In this configuration, the Catalyst 9400 successfully learned all 63,000 MAC addresses without any frame flooding and reached 98 % size utilization efficiency.

DUT	MAC Addresses Count	
	Theoretical	Learned
Catalyst 9400	64,000	63,000

Table 21: Catalyst 9400 MAC Table Scalability

Results: Cisco Catalyst 9500-40X

The Catalyst 9500-40X configuration was tested with a total of 40 10GigE ports and two 40GigE ports.

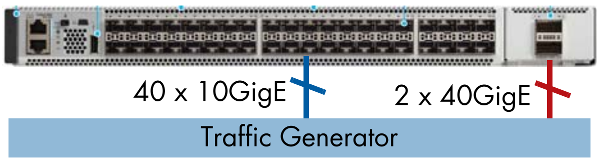

Figure 4: Catalyst 9500-40X

Ethernet and IP Unicast Forwarding Performance

→ Catalyst 9500-40X exhibited line rate performance for fixed frame sizes and IMIX

Cisco explained that the Catalyst 9500-40X features a specific ASIC design for connecting the 10GigE ports with the two 40GigE uplink ports. To fully load the ASICs managing the 40GigE uplinks, we designed a traffic layout consisting of full-load North-South traffic between eight 10GigE and two 40GigE ports respectively, totaling at 80 Gbit/s traffic in each direction. To load the remaining 10GigE ports as well, we applied East-West traffic between pairs of these 24 ports. With Ethernet (Layer 2), the traffic profile was set to one-to-one and with IP (Layer 3), the traffic profile was set many to many.

Stream	North-South	East-West
Bandwidth	160 Gbit/s	320 Gbit/s
10 GigE Ports	8	32
40 GigE Ports	2	0

Table 22: Catalyst 9500-40X Traffic Flow

In the IP packet forwarding performance test, we configured the test with many to many traffic profile as well as mixed speed ports (10 Gbit/s and 40 Gbit/s). The IMIX frames contained Jumbo frame. The combination of mix frame sizes and jumbo frames with test tool's scheduling create bursts and that results in minor traffic loss for the IMIX test. The loss in this test causes the latency measurement to be very high. The loss rate was 0.9% (4Gbit/s loss out of 480Gbit/s of total traffic) when forwarding traffic at line rate with the IMIX frame setup.

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	480	3	4	5
128	480	3	4	4
256	480	3	4	4
512	480	3	4	4
1024	480	3	4	5
1518	480	3	4	5
IMIX	480	3	10	13

Table 23: Catalyst 9500-40X Forwarding Performance Ethernet

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	480	3	4	5
128	480	3	4	5
256	480	3	4	6
512	480	3	4	6
1024	480	3	6	9
1518	480	3	5	8
IMIX	476	3	93	1,552

Table 24: Catalyst 9500-40X Forwarding Performance IPv4

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
78	480	3	4	5
128	480	3	4	5
256	480	3	4	5
512	480	3	5	6
1024	480	3	5	8
1518	480	3	6	8
IMIX	476	3	92	1,166

Table 25: Catalyst 9500-40X Forwarding Performance IPv6

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
78	480	3	4	5
128	480	3	4	6
256	480	3	5	7
512	480	3	5	8
1024	480	3	6	11
1518	480	3	8	13
IMIX	476	3	90	1,222

Table 26: Catalyst 9500-40X Forwarding Performance Mix of IPv4 and IPv6

Results: Cisco Catalyst 9500-24Q

The Catalyst 9500-24Q is a campus core component. It is equipped with 24 ports of 40GigE. Its typical use case is to connect the 40GigE uplinks of many other campus switches at a central location. In this function, the core switch will typically experience high load and nearly full mesh configuration.

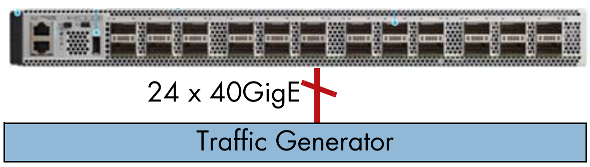

Figure 5: Catalyst 9500-24Q

Ethernet and IP Unicast Forwarding Performance

➔ Catalyst 9500-24Q exhibited line rate performance for fixed frame sizes and IMIX

To fully utilized the forwarding capacity of the core component, the test traffic carried the full load for all 24 x 40 GigE ports. The flows were configured in a full-mesh layout to maximize the load of the ASIC. The goal was to verify whether each port can exchange data with all other ports of the same ASIC at maximum speed. In the IP packet forwarding performance test, we configured full-mesh traffic profile and the IMIX frames contained Jumbo frame. The combination of mix frame sizes and jumbo frames with test tool's scheduling create bursts and that results in minor traffic loss for the IMIX test. We observed 0.7% packet loss (6 Gbit/s loss of 960 Gbit/s of total traffic) when forwarding traffic at line rate. The loss in this test causes the latency measurement to be very high.

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	960	3	4	4
128	960	3	4	4
256	960	3	4	5
512	960	3	4	5
1024	960	3	4	5
1518	960	3	4	5
IMIX	960	3	6	7

Table 27: Catalyst 9500-24Q Forwarding Performance Ethernet

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
64	960	3	4	4
128	960	3	4	4
256	960	3	4	4
512	960	3	4	5
1024	960	3	4	6
1518	960	3	4	6
IMIX	954	3	29	412

Table 28: Catalyst 9500-24Q Forwarding Performance IPv4

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [µs]		
		Min	Avg	Max
78	960	3	4	4
128	960	3	4	4
256	960	3	4	5
512	960	3	4	5
1024	960	3	4	5
1518	960	3	4	5
IMIX	954	3	29	363

Table 29: Catalyst 9500-24Q Forwarding Performance IPv6

Frame Size [Bytes]	Throughput [Gbit/s]	Latency [μs]		
		Min	Avg	Max
78	960	3	4	4
128	960	3	4	5
256	960	3	4	5
512	960	3	4	6
1024	960	3	4	6
1518	960	3	4	6
IMIX	954	3	29	395

Table 30: Catalyst 9500-24Q Forwarding Performance Mix of IPv4 and IPv6

Ethernet Multicast Forwarding Performance

➔ Catalyst 9500-24Q showed line rate performance for 23 receivers over 40 groups

The Catalyst 9500-24Q showed that 40 IGMP multicast groups were successfully joined as configured and full line rate throughput performance was measured as expected.

DUT	Expected Throughput	Measured Throughput
Catalyst 9500-24Q	920 Gbit/s	920 Gbit/s

Table 31: Catalyst 9500-24Q Multicast Forwarding Performance

Forwarding Information Base Scalability

➔ Catalyst 9500-24Q successfully installed 65,000 IPv4 routes into the FIB

The Catalyst 9500-24Q successfully learned 65,000 prefixes without any packet loss. The FIB population took less than one minute.

Network Prefixes Count			
Expected	Initial	Advertised	Installed
64,000	173	65,000	65,173

Table 32: Catalyst 9500-24Q FIB Scalability

After exceeding 65,250 prefixes, we observed packet loss in two prefixes as expected and CPU and memory utilization was stable. Via CLI, we observed the expected alarm message indicating that the FIB table was exhausted.

MAC Table Scalability

➔ Catalyst 9500-24Q successfully learned 63,150 MAC address

Cisco explained the theoretical maximum size of the MAC address table would be 64,000 entries. Before reaching this limit, we observed flooded frames. We reduced the MAC address to 63,150 subsequently. The Catalyst 9500-24Q successfully learned all 63,150 MAC addresses without any frame flooding and reached 97% size utilization efficiency.

DUT	MAC Addresses Count	
	Theoretical	Learned
Catalyst 9500-24Q	64,000	63,150

Table 33: Catalyst 9500-24Q MAC Table Scalability

About EANTC

EANTC (European Advanced Networking Test Center) is internationally recognized as one of the world's leading independent test centers for telecommunication technologies.

Based in Berlin, the company offers vendor-neutral consultancy and realistic, reproducible high-quality testing services since 1991. Customers include leading network equipment manufacturers, tier 1 service providers, large enterprises and governments worldwide. EANTC's Proof of Concept, acceptance tests and network audits cover established and next-generation fixed and mobile network technologies.

EANTC AG
Salzufer 14, 10587 Berlin, Germany
info@eantc.com, http://www.eantc.com/
[v1.0 20190124]