



Light Reading
OC-192 and OC-768 Router
Test
(640 Gbit/s in total)

Test Plan

Version 2.2

EANTC AG
European Advanced Networking Test Center

Copyright (C) 2004

EANTC European Advanced Networking Test Center Aktiengesellschaft

This test plan document is copyrighted by EANTC AG. It may not, in whole or in part, be reproduced, transmitted by any means or stored in any web site or electronic retrieval system without the prior written permission of EANTC AG. EANTC AG grants the receiving party of this test plan a non-transferable right to use this document for internal purposes with regards to projects with EANTC.

All copies must retain and reproduce this copyright notice and all other copyright notices contained within the original material.

Einsteinufer 17
D-10587 Berlin
Germany

Tel. +49. (0)30. 318 05 95-0
Fax +49. (0)30. 318 05 95-10
E-Mail info@eantc.de
WWW <http://www.eantc.de/>

This document was created for Light Reading.

Table of Contents

SCHEDULE AND LOCATION	4
ANALYSIS EQUIPMENT	4
SOFTWARE AND HARDWARE VERSIONS	6
IPV4 IMIX-BASED FORWARDING	18
IPV4/IPV6 IMIX-BASED FORWARDING	20
IPV4/IPV6 IMIX-BASED FORWARDING WITH SERVICES	22
BGP PEERING AND ROUTE RESOLUTION TEST	24
COMBINED IPV4 UNICAST/MULTICAST FORWARDING TEST	28
CLASS OF SERVICE TEST	30
SOFTWARE MAINTENANCE	34
MPLS SCALABILITY TEST	35

Introduction

Schedule and Location

Each test group will take roughly one day.

Day	Test Session
Monday	IPv4 and IPv4/IPv6 forwarding
Tuesday	BGP peering test
Wednesday	Multicast forwarding
Thursday	Class of Service test
Friday	Software maintenance and MPLS scalability test

The vendor, Agilent and EANTC will cooperatively execute the tests. The vendor may decide on the order of tests, however Test Case 1 has to run first because it is prerequisite to the others. Tests will typically be carried out at a vendor lab.

Analysis Equipment

A cluster of Agilent N2X systems (software release 6.3) with 40 OC-192 interfaces will be used for the test, as well as a second cluster of Agilent N2X systems (software release 5.1.1) with 16 OC-192 interfaces and two OC-768 interfaces. Agilent analysis equipment will be used.

Before starting the test, EANTC will run a reference loopback test (Test Case 1 in a port-to-port configuration) with the analysis equipment, to ensure that there are no physical or software problems that might affect the test results. The vendor may access the results of the reference test to confirm that the analysis equipment is working correctly. During the test, the vendor may request additional loopback reference tests if there is any concern that test results are not valid.

All test cases will be run at least twice. If the results are not exactly identical, the test case will be run a third time. EANTC

will calculate an error margin (confidence interval) for test results.

Software and Hardware Versions

Versions of the analysis equipment:

Model	SW Version
Agilent N2X	6.3 6.4 (MPLS test only)
Agilent RouterTester	5.1.1

Versions of the systems under test:

Model	HW Version	SW Version
Chassis		
Interface OC-192		
Interface OC-768		
...		

Only true SONET OC-768, a.k.a. SDH STM-256, interface modules will be allowed to participate in this test. 4x OC-192 WDM ports are not considered true OC-768 interfaces.

All cards of the same type (line cards, switching processors, etc.) are required to be the same production revision. Only production release or public beta (purchasable) hardware is allowed for the test.

All tests will be carried out with the emulators connected to the same physical ports. If the device under test is modular, modules will be used in a way that both the back plane and the modules' internal processing units are stressed: Multiple modules are used, and at least one module of each type is fully populated.

The vendor chooses one single software release for all tests. During the test session, software must not be exchanged (except for the software upgrade test). Only production releases or public beta versions are allowed. At the beginning of the test, EANTC will download the software for installation from the vendor's web server

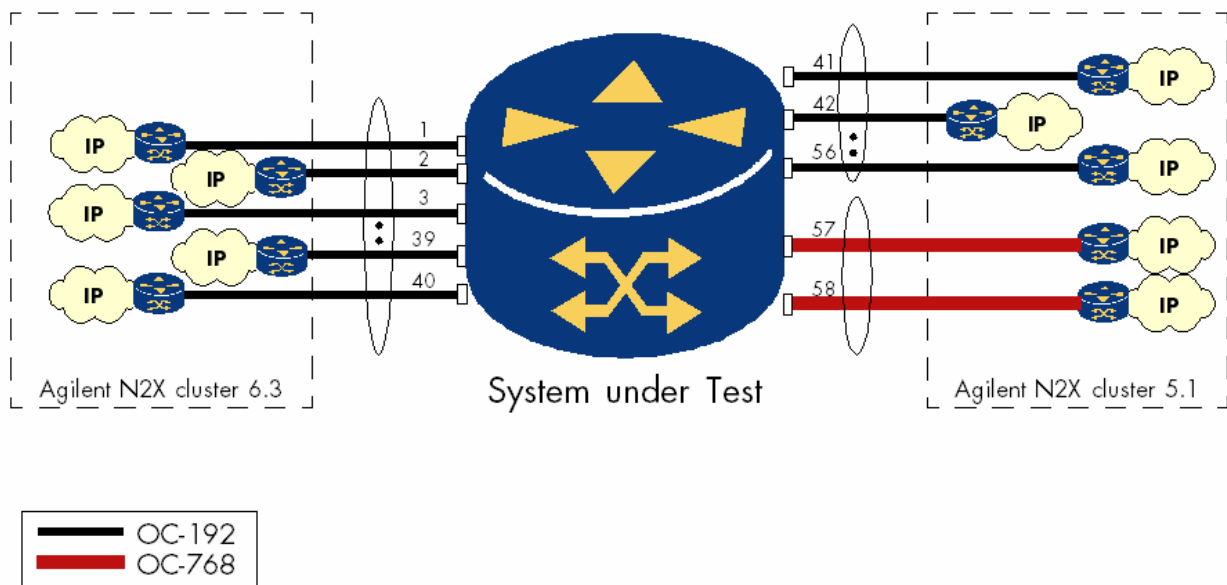
Configuration changes are allowed under EANTC supervision. System uptime and configuration activities will be monitored by EANTC. Reboots are not expected to happen during the test session. The occurrence of a frozen or any interrupted module will also be noted,

regardless of whether the vendor documentation notices that it becomes unusable. This type of event will be mentioned in the article, and will be considered in the calculations for overall system availability.

General Test Setup

The system under test will be connected to two Agilent N2X clusters. (The clusters are separate because of limited interoperability between different software/hardware versions.)

All test streams will be separated into two groups because each cluster is controlled by a separate graphical user interface.



All tests will use all the interfaces mentioned above:

56 x OC-192

2 x OC-768

All tests will use the following IPv4 and IPv6 addresses:

A /30 subnet mask will be used for IPv4 transit networks; the router under test always uses the host address .1, the emulator host address .2. The attached customer networks use various subnet masks. These and the BGP peers are detailed in the BGP test case. For all tests without routing protocols, the customer networks will be used flat (/8 subnet mask).

All IP streams of one port originate from the same source IP address; only the destination IP addresses differ.

The interfaces corresponding to each test cluster are interleaved to distribute the load across the switch fabric.

IPv4 Addressing:

Router Port	Transit Network	Cluster	Agilent Port	Bandwidth	BGP AS	Attached Networks
1	10.1.1.0/30	CLUSTER #1		OC-192	1	128.0.0.0/x
2	10.1.2.0/30	CLUSTER #12		OC-192	2	129.0.0.0/x
3	10.1.3.0/30	CLUSTER #1		OC-192	3	130.0.0.0/x
4	10.1.4.0/30	CLUSTER #1		OC-192	4	131.0.0.0/x
5	10.1.5.0/30	CLUSTER #1		OC-192	5	132.0.0.0/x
6	10.1.6.0/30	CLUSTER #1		OC-192	6	133.0.0.0/x
7	10.1.7.0/30	CLUSTER #1		OC-192	7	134.0.0.0/x
8	10.1.8.0/30	CLUSTER #1		OC-192	8	135.0.0.0/x
9	10.1.9.0/30	CLUSTER #1		OC-192	9	136.0.0.0/x
10	10.1.10.0/30	CLUSTER #1		OC-192	10	137.0.0.0/x
11	10.1.11.0/30	CLUSTER #1		OC-192	11	138.0.0.0/x
12	10.1.12.0/30	CLUSTER #1		OC-192	12	139.0.0.0/x
13	10.1.13.0/30	CLUSTER #1		OC-192	13	140.0.0.0/x
14	10.1.14.0/30	CLUSTER #1		OC-192	14	141.0.0.0/x
15	10.1.15.0/30	CLUSTER #1		OC-192	15	142.0.0.0/x
16	10.1.16.0/30	CLUSTER #1		OC-192	16	143.0.0.0/x
17	10.1.17.0/30	CLUSTER #1		OC-192	17	144.0.0.0/x
18	10.1.18.0/30	CLUSTER #1		OC-192	18	145.0.0.0/x
19	10.1.19.0/30	CLUSTER #1		OC-192	19	146.0.0.0/x
20	10.1.20.0/30	CLUSTER #1		OC-192	20	147.0.0.0/x
21	10.1.21.0/30	CLUSTER #1		OC-192	21	148.0.0.0/x
22	10.1.22.0/30	CLUSTER #1		OC-192	22	149.0.0.0/x
23	10.1.23.0/30	CLUSTER #1		OC-192	23	150.0.0.0/x
24	10.1.24.0/30	CLUSTER #1		OC-192	24	151.0.0.0/x
25	10.1.25.0/30	CLUSTER #1		OC-192	25	152.0.0.0/x
26	10.1.26.0/30	CLUSTER #1		OC-192	26	153.0.0.0/x
27	10.1.27.0/30	CLUSTER #1		OC-192	27	154.0.0.0/x
28	10.1.28.0/30	CLUSTER #1		OC-192	28	155.0.0.0/x
29	10.1.29.0/30	CLUSTER #1		OC-192	29	156.0.0.0/x
30	10.1.30.0/30	CLUSTER #1		OC-192	30	157.0.0.0/x
31	10.1.31.0/30	CLUSTER #1		OC-192	31	158.0.0.0/x
32	10.1.32.0/30	CLUSTER #1		OC-192	32	159.0.0.0/x
33	10.1.33.0/30	CLUSTER #1		OC-192	33	160.0.0.0/x
34	10.1.34.0/30	CLUSTER #1		OC-192	34	161.0.0.0/x

35	10.1.35.0/30	CLUSTER #1	OC-192	35	162.0.0.0/x
36	10.1.36.0/30	CLUSTER #1	OC-192	36	163.0.0.0/x
37	10.1.37.0/30	CLUSTER #1	OC-192	37	164.0.0.0/x
38	10.1.38.0/30	CLUSTER #1	OC-192	38	165.0.0.0/x
39	10.1.39.0/30	CLUSTER #1	OC-192	39	166.0.0.0/x
40	10.1.40.0/30	CLUSTER #1	OC-192	40	167.0.0.0/x
41	10.1.41.0/30	CLUSTER #2	OC-192	41	168.0.0.0/x
42	10.1.42.0/30	CLUSTER #2	OC-192	42	169.0.0.0/x
43	10.1.43.0/30	CLUSTER #2	OC-192	43	170.0.0.0/x
44	10.1.44.0/30	CLUSTER #2	OC-192	44	171.0.0.0/x
45	10.1.45.0/30	CLUSTER #2	OC-192	45	172.0.0.0/x
46	10.1.46.0/30	CLUSTER #2	OC-192	46	173.0.0.0/x
47	10.1.47.0/30	CLUSTER #2	OC-192	47	174.0.0.0/x
48	10.1.48.0/30	CLUSTER #2	OC-192	48	175.0.0.0/x
49	10.1.49.0/30	CLUSTER #2	OC-192	49	176.0.0.0/x
50	10.1.50.0/30	CLUSTER #2	OC-192	50	177.0.0.0/x
51	10.1.51.0/30	CLUSTER #2	OC-192	51	178.0.0.0/x
52	10.1.52.0/30	CLUSTER #2	OC-192	52	179.0.0.0/x
53	10.1.53.0/30	CLUSTER #2	OC-192	53	180.0.0.0/x
54	10.1.54.0/30	CLUSTER #2	OC-192	54	181.0.0.0/x
55	10.1.55.0/30	CLUSTER #2	OC-192	55	182.0.0.0/x
56	10.1.56.0/30	CLUSTER #2	OC-192	56	183.0.0.0/x
57	10.1.57.0/30	CLUSTER #2	OC-768	57	184.0.0.0/x
58	10.1.58.0/30	CLUSTER #2	OC-768	58	185.0.0.0/x

IPv6 Addressing:

Router Port	Transit Network	Agilent Cluster	Agilent Port	Bandwidth	BGP AS	Attached Networks
1	5:1::0/124	CLUSTER #1		OC-192	1	6:1::0/x
2	5:2::0/124	CLUSTER #1		OC-192	2	6:2::0/x
3	5:3::0/124	CLUSTER #1		OC-192	3	6:3::0/x
4	5:4::0/124	CLUSTER #1		OC-192	4	6:4::0/x
5	5:5::0/124	CLUSTER #1		OC-192	5	6:5::0/x
6	5:6::0/124	CLUSTER #1		OC-192	6	6:6::0/x
7	5:7::0/124	CLUSTER #1		OC-192	7	6:7::0/x
8	5:8::0/124	CLUSTER #1		OC-192	8	6:8::0/x
9	5:9::0/124	CLUSTER #1		OC-192	9	6:9::0/x
10	5:A::0/124	CLUSTER #1		OC-192	10	6:A::0/x
11	5:B::0/124	CLUSTER #1		OC-192	11	6:B::0/x
12	5:C::0/124	CLUSTER #1		OC-192	12	6:C::0/x
13	5:D::0/124	CLUSTER #1		OC-192	13	6:D::0/x
14	5:E::0/124	CLUSTER #1		OC-192	14	6:E::0/x
15	5:F::0/124	CLUSTER #1		OC-192	15	6:F::0/x
16	5:10::0/124	CLUSTER #1		OC-192	16	6:10::0/x

17	5:11::0/124	CLUSTER #1	OC-192	17	6:11::0/x
18	5:12::0/124	CLUSTER #1	OC-192	18	6:12::0/x
19	5:13::0/124	CLUSTER #1	OC-192	19	6:13::0/x
20	5:14::0/124	CLUSTER #1	OC-192	20	6:14::0/x
21	5:15::0/124	CLUSTER #1	OC-192	21	6:15::0/x
22	5:16::0/124	CLUSTER #1	OC-192	22	6:16::0/x
23	5:17::0/124	CLUSTER #1	OC-192	23	6:17::0/x
24	5:18::0/124	CLUSTER #1	OC-192	24	6:18::0/x
25	5:19::0/124	CLUSTER #1	OC-192	25	6:19::0/x
26	5:1A::0/124	CLUSTER #1	OC-192	26	6:1A::0/x
27	5:1B::0/124	CLUSTER #1	OC-192	27	6:1B::0/x
28	5:1C::0/124	CLUSTER #1	OC-192	28	6:1C::0/x
29	5:1D::0/124	CLUSTER #1	OC-192	29	6:1D::0/x
30	5:1E::0/124	CLUSTER #1	OC-192	30	6:1E::0/x
31	5:1F::0/124	CLUSTER #1	OC-192	31	6:1F::0/x
32	5:20::0/124	CLUSTER #1	OC-192	32	6:20::0/x
33	5:21::0/124	CLUSTER #1	OC-192	33	6:21::0/x
34	5:22::0/124	CLUSTER #1	OC-192	34	6:22::0/x
35	5:23::0/124	CLUSTER #1	OC-192	35	6:23::0/x
36	5:24::0/124	CLUSTER #1	OC-192	36	6:24::0/x
37	5:25::0/124	CLUSTER #1	OC-192	37	6:25::0/x
38	5:26::0/124	CLUSTER #1	OC-192	38	6:26::0/x
39	5:27::0/124	CLUSTER #1	OC-192	39	6:27::0/x
40	5:28::0/124	CLUSTER #1	OC-192	40	6:28::0/x
41	5:29::0/124	CLUSTER #2	OC-192	41	6:29::0/x
42	5:2A::0/124	CLUSTER #2	OC-192	42	6:2A::0/x
43	5:2B::0/124	CLUSTER #2	OC-192	43	6:2B::0/x
44	5:2C::0/124	CLUSTER #2	OC-192	44	6:2C::0/x
45	5:2D::0/124	CLUSTER #2	OC-192	45	6:2D::0/x
46	5:2E::0/124	CLUSTER #2	OC-192	46	6:2E::0/x
47	5:2F::0/124	CLUSTER #2	OC-192	47	6:2F::0/x
48	5:30::0/124	CLUSTER #2	OC-192	48	6:30::0/x
49	5:31::0/124	CLUSTER #2	OC-192	49	6:31::0/x
50	5:32::0/124	CLUSTER #2	OC-192	50	6:32::0/x
51	5:33::0/124	CLUSTER #2	OC-192	51	6:33::0/x
52	5:34::0/124	CLUSTER #2	OC-192	52	6:34::0/x
53	5:35::0/124	CLUSTER #2	OC-192	53	6:35::0/x
54	5:36::0/124	CLUSTER #2	OC-192	54	6:36::0/x
55	5:37::0/124	CLUSTER #2	OC-192	55	6:37::0/x
56	5:38::0/124	CLUSTER #2	OC-192	56	6:38::0/x
57	5:39::0/124	CLUSTER #2	OC-768	57	6:39::0/x
58	5:3A::0/124	CLUSTER #2	OC-768	58	6:3A::0/x

Physical port assignment by router port:

Router Port	Slot	Interface
1	6	0
2	6	1
3	6	2
4	6	3
5	7	0
6	7	1
7	7	2
8	7	3
9	8	0
10	8	1
11	8	2
12	0	1
13	9	0
14	9	1
15	9	2
16	0	3
17	10	0
18	10	1
19	10	2
20	1	1
21	11	0
22	11	1
23	11	2
24	1	3
25	12	0
26	12	1
27	12	2
28	2	1
29	13	0
30	13	1
31	13	2
32	2	3
33	14	0
34	14	1
35	14	2
36	3	1
37	15	0

38	15	1
39	15	2
40	3	3
41	0	0
42	8	3
43	0	2
44	9	3
45	1	0
46	10	3
47	1	2
48	11	3
49	2	0
50	12	3
51	2	2
52	13	3
53	3	0
54	14	3
55	3	2
56	15	3
57	4	0
58	5	0

Slot by router port and cluster:

Slot	Interface	Router Port	Cluster
0	0	41	CLUSTER #2
0	1	12	CLUSTER #1
0	2	43	CLUSTER #2
0	3	16	CLUSTER #1
1	0	45	CLUSTER #2
1	1	20	CLUSTER #1
1	2	47	CLUSTER #2
1	3	24	CLUSTER #1
2	0	49	CLUSTER #2
2	1	28	CLUSTER #1
2	2	51	CLUSTER #2
2	3	32	CLUSTER #1
3	0	53	CLUSTER #2
3	1	36	CLUSTER #1

3	2	55	CLUSTER #2
3	3	40	CLUSTER #1
4	0	57	CLUSTER #2
5	0	58	CLUSTER #2
6	0	1	CLUSTER #1
6	1	2	CLUSTER #1
6	2	3	CLUSTER #1
6	3	4	CLUSTER #1
7	0	5	CLUSTER #1
7	1	6	CLUSTER #1
7	2	7	CLUSTER #1
7	3	8	CLUSTER #1
8	0	9	CLUSTER #1
8	1	10	CLUSTER #1
8	2	11	CLUSTER #1
8	3	42	CLUSTER #2
9	0	13	CLUSTER #1
9	1	14	CLUSTER #1
9	2	15	CLUSTER #1
9	3	44	CLUSTER #2
10	0	17	CLUSTER #1
10	1	18	CLUSTER #1
10	2	19	CLUSTER #1
10	3	46	CLUSTER #2
11	0	21	CLUSTER #1
11	1	22	CLUSTER #1
11	2	23	CLUSTER #1
11	3	48	CLUSTER #2
12	0	25	CLUSTER #1
12	1	26	CLUSTER #1
12	2	27	CLUSTER #1
12	3	50	CLUSTER #2
13	0	29	CLUSTER #1
13	1	30	CLUSTER #1
13	2	31	CLUSTER #1
13	3	52	CLUSTER #2
14	0	33	CLUSTER #1
14	1	34	CLUSTER #1
14	2	35	CLUSTER #1
14	3	54	CLUSTER #2

15	0	37	CLUSTER #1
15	1	38	CLUSTER #1
15	2	39	CLUSTER #1
15	3	56	CLUSTER #2

The router ID and IPv6 loopback address is defined as 100.100.100.100/32 (IPv4) and ::100.100.100.100/128 (IPv6)

eBGP Sessions

eBGP4+ will be used on all ports if not mentioned otherwise. The emulator will advertise the customer network prefixes dynamically. The BGP session transport uses IPv6 on native IPv6 interfaces, IPv4 on IPv4 interfaces. IPv4 and IPv6 routes are exchanged over all BGP connections.

The emulator will send data traffic on one IP address per emulated network.

L2 Definitions

The Packet over SONET (PoS) parameters will be chosen as follows:

- PPP encapsulation (IPCP, IPv6CP, MPLSCP protocols)
- Standard SDH and PPP timers
- 32-bit checksum

IP Packet Stream Definitions

All load tests will use a mix of packet sizes, representing Internet mix traffic (IMIX). This is a deterministic way of simulating real network traffic according to the frame size usage. Some studies indicate that Internet traffic consists of fixed percentages of different frame sizes. IMIX traffic contains a mixture of frame sizes in a ratio to each other that approximates the overall makeup of frame sizes observed in real Internet traffic. A detailed definition of IMIX packet sizes and distributions can be found above.

Using IMIX traffic allows us to test the SUT under realistic conditions, as compared to single packet sizes tested sequentially.

The number of packets per second is related to full load ("wire-speed") on OC-192/OC-768. During the binary search of RFC2544, absolute packet numbers are modified but the packet size distribution (percentages) remain unchanged.

IPv4 packet streams

IP Packet Size (Bytes)	Bandwidth %
40	57 %
552	7 %
576	16 %
1500	20 %

IPv6 packet streams

IP Packet Size (Bytes)	Bandwidth %
64	38 %
174	23 %
750	16 %
1500	23 %

An exception to this rule is the loss-less software upgrade. It uses only 40-byte packets (IPv4) or 60-byte packets (IPv6) because the interrupt time can be measured more exactly with small packets.

Also, the Forwarding With Services test uses 48-byte UDP packets instead of 40-byte packets.

NOTE: The N2X traffic generator will calculate the correct wire-speed packet rate, taking PoS byte stuffing on the transmit interface to the SUT into account. The router will modify the IP header (TTL) and the checksums, potentially generating byte stuffing at different locations. The extra bandwidth required by byte stuffing at incoming and outgoing interfaces should be statistically identical.

However, the maximum throughput will be reduced for the latency measurements to ensure that byte stuffing does not exceed the line rate on the outgoing interfaces.

RFC2544 Parameters

All RFC2544 tests are carried out with the following parameters:

- Precision 1%
- Lost packets accepted: 0
- Test duration per run: 60 seconds
- Pure IPv4 or IPv6 packets, no UDP or TCP header used
- Packet payload is checked for in-sequence delivery and transmission errors with test blocks and CRC data integrity checksums.

Test Cases

IPv4 IMIX-Based Forwarding

Test Objectives

Determine the maximum IPv4 forwarding performance of the system under test under realistic conditions.

Description

The test will use the IMIX IPv4 packet mix as defined above. We calculate 10 flows per IP route. Based on the estimated worst-case maximum number of BGP routes (1,500,000; see below), there will be 15 million flows in total (per router). This amount of flows exceeds today's requirements and, according to carrier comments, is a good number.

However, this test case does not use BGP, because routing protocols are tested later. Instead, the calculated number of 15 million flows will be generated as follows:

One route per port (/8, Class A) will be statically configured as defined in the address plan section.

From this network, the emulator will use multiple (see below) non-consecutive destination IP addresses and one source IP address per port.

The emulator will be configured to run "partial-mesh" tests (from each IP address of the port to each other IP address on all other ports, excluding local loopback flows returning to the same port), resulting in

$32 \times 31 \times 28,925 / 2 = 14,346,800$ flows for Cluster #1

$20 \times 19 \times 3,440 / 2 = 653,600$ flows for Cluster #2

The number of flows per cluster is defined by the Agilent N2X flow scalability.

Flows (specifically in cluster #2) are distributed in a balanced way so that, in theory, no port is congested.

Procedure

1. Connect all 56 ports with the emulators.
2. Configure all ports for IPv4, configure one route per port to the customer network defined in the address plan, using a class A (255.0.0.0) network mask.

3. Send full mesh IPv4 IMIX traffic within both Agilent N2X clusters for all ports with n IP addresses per port as defined above.
4. Determine the maximum throughput according to RFC 2544.
5. Execute one additional test run at the maximum throughput for an extended time of 300 seconds to uncover potential long-term robustness issues.
6. Repeat the test at 99.5% load (or at the maximum throughput as measured in step 4 if it was less than 99.5% load) to find the forwarding latency. Set the test duration to 120 seconds.
7. The forwarding and latency tests (steps 4-6) may be repeated with 40-byte IP packets only, to find the IPv4 forwarding performance limits.

Expected Results

With IMIX traffic, we expect wire-speed forwarding, no packet loss, no out-of-sequence packets, no stray (misrouted) frames, low latency at 99.5% wire-speed.

With 40-byte IP packets, we expect at least the same forwarding performance as measured in the IMIX test (counting packets per second), no out-of-sequence packets.

IPV4/IPV6 IMIX-Based Forwarding

Test Objectives

Determine the maximum forwarding performance for mixed IPv4 and IPv6 traffic of the system under test.

Description

In 2004, the percentage of IPv6 traffic in the Japanese Internet backbone reached 1% of all traffic. A worst-case guess is that IPv6 traffic might reach roughly 25% of all Internet traffic until 2012, which is what we are going to use in this test. We estimate that the current generation of core routers will have a maximum lifetime of eight years in carrier networks.

The emulator will be configured to run "partial-mesh" tests (from each IP address of the port to each other IP address on all other ports, excluding local loopback flows returning to the same port), resulting in

$32 \times 31 \times 30,077 / 2 = 14,918,192$ flows for Cluster #1

$20 \times 19 \times 431 / 2 = 81,890$ flows for Cluster #2

The number of flows per cluster is defined by the Agilent N2X flow scalability.

Flows (specifically in cluster #2) are distributed in a balanced way so that, in theory, no port is congested.

Procedure

1. Connect all 56 ports to the emulator.
2. Configure all ports for both IPv4 and IPv6 as detailed in the address plan. Add one static IPv4 route (/8) and one static IPv6 route (/32) to emulated customer networks per port.
3. Send partial-mesh IPv4- and IPv6-IMIX traffic within both Agilent N2X clusters for all ports with n IPv4 addresses and m IPv6 addresses per port.
4. Determine the maximum throughput using the binary search algorithm according to RFC 2544.
5. Execute one additional test run at the maximum throughput for an extended time of 300 seconds to uncover potential long-term robustness issues.
6. Repeat the test at 99.5% load (or at the maximum throughput as measured in step 4 if it was less than 99.5% load) to find the forwarding latency. Set the test duration to 120 seconds.

Expected Results

We expect wire-speed forwarding for both IPv4 and IPv6 flows, no packet loss, no out-of-sequence packets, no stray (misrouted) frames, low latency at wire-speed.

IPV4/IPV6 IMIX-Based Forwarding With Services

Test Objectives

Determine the maximum forwarding performance for mixed IPv4 and IPv6 traffic of the system under test while access lists and ACL logging are enabled.

Description

Most IP switches in carrier and enterprise environments need to observe basic access control rules today, sometimes with basic intrusion detection mechanisms (ACL logging).

This test case verifies that the system under test can forward IPv4 and IPv6 packets at wire-speed even when access lists and ACL logging are configured.

Procedure

1. Configure a 5,001 entry ACL for IPv4, where 5,000 entries are DENY and the last entry is a PERMIT-ALL. (ACL DENY criteria are pseudo random and not sequential, thereby preventing sequential ranges being converted into a single ACL entry. Also, 50% of the DENY entries are configured matching the IP addresses of the data streams in step 5, but not matching the UDP ports of the data traffic: DENY UDP dst port = 1, UDP dst port = 2, ... UDP dst port = 2500). Apply the ACL as both an ingress and egress filter on each port. Configure an identical group of access lists for IPv6.
2. Configure ACL Logging on the 5,000th ACL entry (both IPv4 / IPv6), but do not enable it yet.
3. Configure a 500-entry QoS traffic classification filter on each interface (both IPv4 / IPv6), but do not enable it yet.
4. Connect all ports to the emulator.
5. Configure all ports for both IPv4 and IPv6 as detailed in the address plan. Add one static IPv4 route (/8) and one static IPv6 route (/32) to emulated customer networks per port.
6. Send full mesh IPv4- and IPv6-IMIX traffic within both Agilent N2X clusters for all ports with 16,500 IPv4 addresses and 5,500 IPv6 addresses per port. In contrast to other test cases, send UDP packets instead of IP packets without a specific protocol - consequently, the smallest packet size will be 48 bytes instead of 40 bytes. Choose permutating UDP ports 5000... 7,500. Make sure that all traffic only matches the final "PERMIT-ALL" ACL entry, thereby forcing the switch to compare the traffic to all 5,001 ACL entries per port in both the ingress and egress directions (double lookups).

7. Determine the maximum throughput using the binary search algorithm and forwarding latency according to RFC 2544.
8. To confirm the ACLs are actively protecting the network, re-run the test and send a mix of traffic. One of the IP addresses per port matches the 5,000th DENY entry in the ACL, the other IP addresses match none of the DENY entries. Ensure that 100% of the traffic matching the DENY entry is dropped and that the remaining traffic is forwarded without loss.
9. Enable ACL logging. Send the same traffic as in step 9. Confirm that the ACL logging feature does not affect performance or latency and that the details recorded help identify the individual flows that are attempting to break through the ACL barrier.
10.U
se the CLI to confirm QoS is enabled and that QoS Traffic Classification Lists are active at the same time as the security ACLs.
11.E
xecute one additional test run at the maximum throughput for an extended time of 300 seconds to uncover potential long-term robustness issues.

Expected Results

Compare the results of this test to those recorded in the IPv4/IPv6 throughput tests without services configured. The addition of significant levels of services should not affect the router performance.

Check that the ACL Logging accurately records details of the flows that have matched the DENY entries in the 10,001-entry ACL. There should be sufficient information to determine the source of the attack.


BGP Peering and Route Resolution Test

Test Objectives

Determine the performance and the route convergence time of the BGP implementation using both IPv4 and IPv6 routes. The first test determines eBGP convergence for changes in the eBGP table. The second test determines the convergence for the case that changes in the IGP cause changes in the iBGP route recursion.

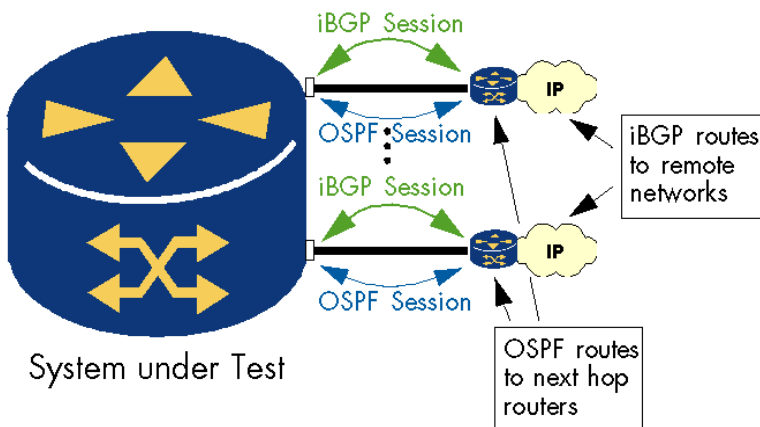
Test Setup

eBGP Test:



System under Test
(does not show all ports/clusters)

iBGP/OSPF Test:



(does not show all ports/clusters)

Description

This test emulates a realistic environment with a large number of routes. We will send traffic over all routes using the same IMIX packet mix as before.

Today, the full Internet routing table has about 160,000 entries, with each entry representing one network (or group of networks) attached to the public Internet. The number of entries is growing fast. The goal of this test is to verify how SUT would handle future growth; we will use an eBGP routing table with 1,500,000 unique routes.

While customer network routes will be distributed via eBGP4/eBGP4+, backbone (internal) routes will be advertised with OSPFv2/OSPFv3. All BGP sessions in Cluster #1 will be configured over IPv6; BGP sessions in Cluster #2 will be created over IPv4.

The following eBGP prefixes will be used ("x" represents the IPv4 customer network prefix reserved for a given port; "y" represents the IPv6 customer network prefix reserved for that port):

Prefix address group	Prefix length	Fixed routes per port	Changing routes per port	Total number of routes per port
x.0.0.0 - x.191.255.255	/19	536	1000	1536
x.192.0.0 - x.240.255.255	/24	8544	4000	12,544
x.241.0.0 - x.241.240.255	/29	5952	2000	7952
y::	/48	500	500	1000
y::	/60	1250	750	2000
y::	/96	3750	1250	5000

Procedure eBGP Test

1. Remove all static routes from the system under test.
Configure eBGP4 over IPv4 peers for Cluster #2, eBGP4+ over IPv6 peers for Cluster #1.
Configure one BGP peer per port.
2. Configure all the fixed and changing IPv4 and IPv6 routes advertised to the system under test as outlined in the table above. These routes are fixed and will not be flapped.

3. Add a second route to each of the "changing" destinations as outlined above to each of the port groups 1/17, 2/18... 16/32; 33/41, 34/41... 40/48; 49/51, 50/52. These routes are a subset of the total routes listed above, but carry a better AS-path so they would be preferred.
4. Send mixed IPv4- and IPv6-IMIX traffic over all routes. The load of IPv4 traffic should be 40% of the maximum load; the load for IPv6 traffic should be 10% of the maximum load. Use one source IP address per route, sending full mesh traffic within each of the two clusters.
5. Keep sending traffic; of the routes configured in step 3, flap all in cluster #1, and flap the largest route pool in cluster #2.
6. Measure the convergence time until all traffic is forwarded to the new routes without loss.
7. Re-advertise the routes configured in step 3 back to their original next-hop destination; measure convergence time again.
8. Steps 2 and 4 may be repeated with gradually increasing numbers of eBGP routes until the system shows degraded forwarding performance on any route (or otherwise unexpected behavior) to find the maximum supported number of routes.

Procedure Route Resolution Test

9. Remove all static routes from the system under test.
Configure one iBGP session per port group (port groups as above, use lower port range).
Configure one OSPF session per port.
10.C
Configure the total BGP routes advertised to the system under test as outlined above, however make the next-hop address different from the analyzer port address (e.g., 10.99.1.1, 10.99.1.2...)
11.C
Configure the OSPF sessions on each port pair so that the corresponding next-hop address is advertised towards the system under test on both ports with different metrics, so that the router picks one of the two ports as the egress port for the set of BGP routes. **Do not redistribute OSPF routes to BGP or vice versa.**
12.S
Send mixed IPv4- and IPv6-IMIX traffic over all routes. The load of IPv4 traffic should be 40% of the maximum load; the load for IPv6 traffic should be 10% of the maximum load. Use one source IP address per route, sending full mesh traffic within each of the two clusters.
13.K
Keep sending traffic; withdraw the routes of the "better" OSPF session.

- 14.....M
 easure the convergence time until all traffic is forwarded to the
 new routes without loss.
- 15.....R
 e-advertise the OSPF routes back to their original next-hop
 destination; measure convergence time again.

Expected Results

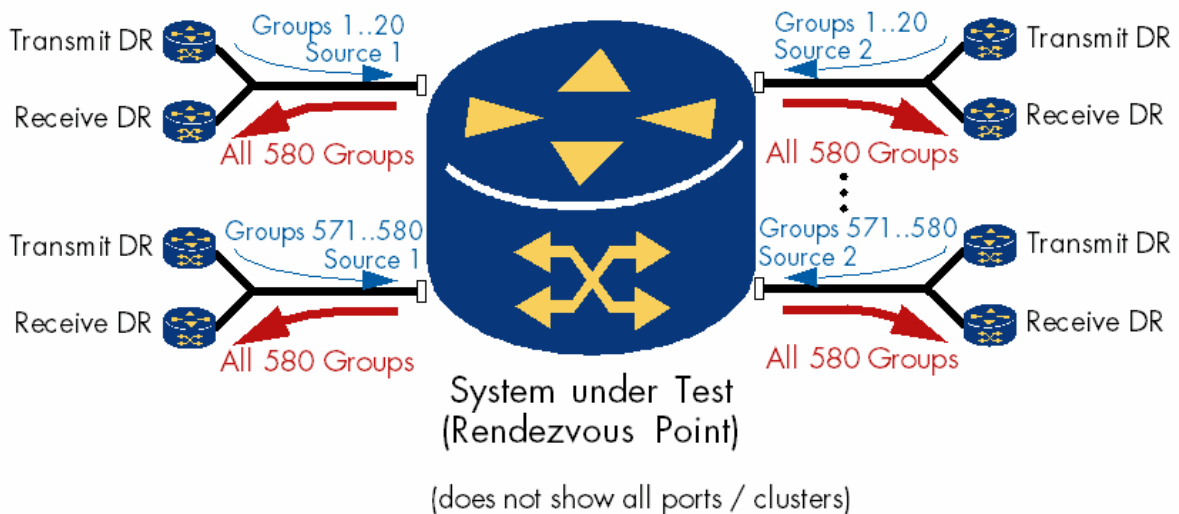
The system under test is expected to process all BGP routes and updates successfully, resulting in zero loss throughput for all routes during normal operation. The system should converge quickly after route flapping (resulting in only a small amount of lost packets), with zero loss throughput at some short time after the flap event.

Combined IPv4 Unicast/Multicast Forwarding Test

Test Objectives

This test evaluates the multicast forwarding capability of the SUT under load (multicast and unicast). For multicast routing we will use PIM-SM, because this is the dominating protocol in carrier networks today. The test is divided into two parts. In part A we will emulate a multicast routing environment, where the SUT is not configured as a rendezvous point (RP). In Part B we will configure the SUT as an RP.

Test Setup



Part A - SUT not configured as RP

Procedure

1. The SUT is connected to the emulator on all ports and employs PIM Sparse Mode. Configure one of the analyzer clusters to emulate the RP.
2. Set up full-mesh IMIX IPv4 unicast traffic within each of the two clusters at 30% load per port. Use a fixed packet size of 125 Bytes at a packet rate of 3 million packets per second (Mpps) to simplify statistics evaluation.
3. Configure an emulated multicast routing environment with 58 emulated designated routers (DRs): Each port sources 20 groups, resulting in a total of 580 multicast groups (thus each group has two different sources on different ports). At the same time, each DR receives

multicast traffic, subscribing to all multicast (S,G) pairs except the ones that are sourced on that port (i.e. each port joins $57 \times 20 = 1140$ (S,G) source-specific groups).

4. The emulator generates multicast traffic with a fixed packet size of 125 Bytes to all groups for 60 seconds. The load is 5000 packets per second per source/group, adding up to 5.7 Mpps (= roughly 5.7 Gbit/s) multicast Rx traffic per port (57 ports \times 20 sources \times 5 Mbit/s --). Naturally there will be packet loss observed on one cluster, because the emulated RP does not forward traffic back to the router under test.

Part B - SUT is configured as RP

Procedure

1. The SUT is connected to the emulator on all ports and employs PIM Sparse Mode.
2. The SUT is configured as an RP.
3. Setup the same traffic streams as in part A and transmit data for 60 seconds.

Expected Results

In both parts, the router under test should process JOIN requests, construct a multicast routing table with source-specific routes, and forward multicast packets accordingly, depending on the amount of traffic. Each port should receive the expected unicast and multicast traffic without packet loss and without out-of-order packets. (With Agilent N2X, the packet loss criteria can be verified only within the cluster of the transmitter port. We will assume that the other cluster does not show packet loss if all packets are received on all ports of the first cluster). In part B, the rendezvous point should calculate source-specific trees and use optimized the multicast paths.

Class of Service Test

Test Objectives

The test case verifies correct implementation of IPv4 Differentiated Services (DiffServ).

Description

The test is divided into two parts:

- Part A: The links will be oversubscribed by 10% using constant data rate traffic flows. The queue management of the system under test has to handle the resulting congestion.
- Part B: Live IP network traffic is usually bursty. The second part of the test will evaluate queue management under burst conditions. We will determine whether bursty traffic has an influence on EF-class traffic.

The test will use five traffic classes. All traffic will be pure IP packets. The following table shows the used traffic classes and the respective IP packet sizes:

DS-Class	DS-Name	Class Name	DS Code Point (decimal)	Packet Size IPv4
EF	Expedited Forwarding	Premium	46	60 Byte
AF11	Assured Forwarding class 1	Gold	10	40 Byte
AF21	Assured Forwarding class 2	Silver	18	552 Byte
AF31	Assured Forwarding class 3	Bronze	26	576 Byte
BE	Best-effort	Best-effort	0	1500 Byte

The traffic streams will be set up in a full mesh as before, however without ports 1-6. Instead, ports 1, 3, and 5 will send their streams

to ports 7 through 47, and ports 2, 4, and 6 will send their streams to ports 8 through 56, resulting in roughly 10% oversubscription. (Ports 1-6 will not receive any data)

IPv6 does not use Differentiated Services in the same way; this test deals with IPv4 only.

Part A - Congestion

- This part verifies the behavior of the SUT under congestion. It will be tested whether the SUT drops frames of higher priority, even if lower priority traffic causes link congestion.

Parameters:

DS-Class	Class Name	DS Code Point	Bandwidth Limitation on Router (for OC-192; multiply by four for OC-768)	Bandwidth transmitted from traffic generator (for OC-192)
EF	Premium	46	0.5 Gbit/s	0.5 Gbit/s
AF11	Gold	10	35 % of line speed without EF	3 Gbit/s
AF21	Silver	18	25 % of line speed without EF	2 Gbit/s
AF31	Bronze	26	25 % of line speed without EF	2 Gbit/s (3 Gbit/s in a second step)
BE	Best-effort	0	(rest = 15 %)	3.5 Gbit/s (1 Gbit/s coming from the extra ports) (2.5 Gbit/s in the second step where again 1 Gbit/s is coming from the extra ports)

If traffic in one DiffServ class exceeds the limit, it shall not be remarked by the router (this is typically an edge router function) but will be forwarded according to queue configuration. The purpose of this test is to verify congestion handling and queueing, not remarking.

Procedure

1. Set up five traffic classes for the incoming traffic on all interfaces of the SUT as described above.
2. Start sending IMIX IPv4 traffic marked with the appropriate DiffServ code points on all interfaces as described above.
3. Observe packet loss and latency of all traffic classes.

Part B - Bursty Traffic

In this part we will observe packet loss and latency of EF traffic while sending bursty traffic on the other traffic classes.

Parameters

For OC-192 ports; multiply by 4 for OC-768 ports

DS-Class	Average bandwidth	Burst Load	Burst Length (packets)	Repeat Count
EF	0.5 Gbit/s	-	-	-
AF11	1 Gbit/s	1.5 Gbit/s	5	1
AF21	1.5 Gbit/s	2.0 Gbit/s	10	1
AF31	2 Gbit/s	4 Gbit/s	25	1
BE	3 Gbit/s	6 Gbit/s	50	1

.....

Procedure

1. Set up five traffic classes for the incoming traffic on all interfaces of the SUT as described.
2. Start sending IMIX IPv4 traffic marked with the appropriate DiffServ code points on all interfaces as described above.
3. Observe packet loss and latency of all traffic classes.

Expected Results

The router should prioritize EF traffic in all conditions. No packet should be lost in the EF traffic class, and latency should not be increased even when the link is oversubscribed. Traffic in all other classes should be forwarded as defined, dropped if necessary.

The CPU (or any other software-based part of the system running on a processor) should not show significant utilization at wire-speed forwarding - QoS classification should be carried out in hardware completely.

Bursty traffic flows should be forwarded in an optimized way, using any kind of weighted selection protocol implemented by the vendor.

Software Maintenance

Test Objectives

The aim of this test is to observe the behavior of the SUT during Software upgrade and hardware module exchange.

Part A Software Upgrade/ Downgrade

Procedure

1. Connect all ports to the traffic generator.
2. Rerun the eBGP peering test, part A (no flapping).
3. Upgrade/downgrade the software of several modules.
4. Analyze the latency, throughput and frame loss.

Part B Hardware Module Exchange Procedure

Procedure

1. Connect all ports to the traffic generator.
2. Rerun the eBGP Peering test.
3. Remove and reinsert a redundant switch processor module.
4. Analyze the latency, throughput and frame loss.
5. Rerun steps 2-4 for the primary switch processor module and an interface module of each type. If possible, also exchange a module within each line card.

Expected Results

We expect that a software upgrade/downgrade influences forwarding and routing only to the extent claimed by the vendor, prior to testing. Exchanging an interface module should not affect other modules at all.

Should go without saying, but really clarify what the vendor says will happen. Hitless should mean hitless.

MPLS Scalability Test

Test Objectives

This test verifies the functionality and scalability of the router under test as an MPLS label switch router (LSR) using RSVP-TE.

Discussion

The system under test is used in the core backbone, not at the edge. Regarding MPLS, it works as an LSR (Label Switch Router) and is not involved in VPN services. Its main function is to negotiate transport MPLS labels, and to maintain traffic-engineering tunnels.

Test Setup

Bidirectional RSVP-TE tunnels are created among the following groups of three ports:

C 1	27	C 11	32	C 21	37	C 45	53
B 2	27	B 12	32	B 22	37	B 46	53
C 3	28	C 13	33	C 23	38	C 47	54
B 4	28	B 14	33	B 24	38	B 48	54
C 5	29	C 15	34	C 25	39	C 49	55
B 6	29	B 16	34	B 26	39	B 50	55
C 7	30	C 17	35	C 41	51	C 56	58
B 8	30	B 18	35	B 42	51	B 57	58
C 9	31	C 19	36	C 43	52		
B 10	31	B 20	36	B 44	52		

Each group consists of

A port originating best-effort tunnels (B)

A port originating constraint-based tunnels (C)

A third port terminating the best-effort and constraint-based tunnels

Procedure

1. Connect all ports to the emulator. (Port 40 is not in use in this test.)
2. Configure all ports with IPv4 addresses as in the IPv4 forwarding test. To simplify the setup, use a fixed size of 64 bytes for all packets. Enable OSPF over IPv4 to distribute the core backbone topology. Configure two emulated OSPF routers behind each port to ensure that the SUT interfaces are not the penultimate hop, and non-NULL labels will be issued. IPv6 is not used in this test.

3. Enable MPLS and Traffic Engineering on all ports, connecting one RSVP-TE peer per port.

Establish 750 bi-directional LSPs (1,500 unidirectional LSPs) between each "C" port and each terminating port.

In addition, establish 750 bi-directional LSPs between each "B" port and each terminating port.

NOTE: In effect, each terminating port will carry 1,500 bi-directional LSPs (3,000 unidirectional LSPs).

Configure an EXP-to-DSCP mapping at the device under test with two classes: Traffic marked with one arbitrary EXP bit setting shall be mapped to the EF class, traffic marked with another arbitrary EXP bit setting shall be mapped to a best-effort scheduling class.

Please note that bi-directional tunnels are created, although only one direction can be oversubscribed later.

4. Generate bi-directional traffic for each port group. Set the appropriate EXP bits for each port. In order to overload each termination port, generate 0.2 x line rate / 750 traffic on all tunnels originating from each best-effort port. Generate 90 % load on ports 56 and 57 to over-subscribe port 58.
5. Verify that all LSPs can be established and that there is no packet loss on TE-LSPs.

Expected Results

The router under test should be able to establish all RSVP-TE tunnels. Tunnel creation delay should be minimal. Throughput should reach wire-speed for all RSVP-TE tunnels. The EF traffic should not observe any packet loss, latency should be minimal.

Frequently Asked Questions

Q: Did vendors have a chance to modify the test plan?

A: No. Vendors who committed to participate in the test had a chance to view the test plan in advance to pre-stage the testing. In addition, participating vendors were allowed to point out technical flaws in the test plan. EANTC carefully examined these suggestions, rejecting modifications that would have improved the test results for a specific vendor.

Q: Why are all cards of each type required to be the same production revision?

A: Because we are testing the latest generation of terabit core routers, and there is no history of different hardware revisions yet.

Q: Why does the BGP peering test use only one IP address per route, while the forwarding tests use many IP addresses per route?

A: The BGP peering test makes use of many more routes per port (22,000 IPv4 and 7000 IPv6 routes). One IP address per route results in 29,000 IP addresses per port, which generates the same amount of flows as the forwarding tests. The goal is to stress the router with a similar number of routes, verifying that the use of dynamic routing does not change the forwarding performance.

Q: Why are there no MPLS fast re-route or LSP re-convergence tests?

A: This test plan addresses only one router under test. We would need multiple routers to achieve representative fail-over/re-convergence times.

Q: Why are almost all RFC2544 tests run with 60-second intervals?

A: It takes a lot of time to run binary search RFC2544 test cases with long intervals (each test is usually run around 5-7 times before the analyzer finds the correct optimum bandwidth). At the ultra broadband speeds of the interfaces under test, throughput granularity is minimal at 60-second test intervals. To accurately measure the latency, we run one additional test with the maximum speed for an extended interval of time.

Q: Why does the test plan not contain any MPLS churn test, measuring the RSVP-TE tunnel establishment delay?

A: In its current software version, the analysis equipment would be the limiting factor for the establishment latency.

Q: Why does the test plan require only 750 MPLS tunnels per port?

A: This is the limit of the Agilent N2X analyzer running 6.4 beta code with standard refresh timers. Other software versions (5.1.1) show larger numbers, especially if the refresh timers are increased, but we needed to use 6.4 because of other added features.