White Paper

Carrier Ethernet Implementation:

How to Profit without the Pain

Agilent N2X

Agilent Technologies

Background and Motivation

This white paper logically follows the flow of the Agilent-LightReading webinar held on June 7, 2007. This can be reviewed at the following URL until June 2008.

http://www.lightreading.com/webinar_archive.asp?doc_id=28280

Introduction

Developing and deploying Carrier Ethernet devices and services offers new revenue generation opportunities while reducing infrastructure and operational costs. However, the implementation path is fraught with scalability, robustness and interoperability risks:

- Will network infrastructure scale to meet customer growth and new services?
- How well can evolving fault-management technologies such as CFM, BFD, LACP, RSVP FRR and MSTP cooperate to reduce network outage?
- Can switched technologies, such as PBB/PBT, and VPLS technologies from different vendors transparently co-exist, providing end-to-end service?

Through anticipating and solving the challenges via carefully planned testing throughout all stages of development and deployment, both vendors and operators can realize the gain without feeling the pain.

This white paper:

- Briefly examines the market and technology drivers;
- Assesses the key development and deployment risks on several Carrier Ethernet technologies;
- Answers the questions above by revealing tangible test scenarios and showing real test results from testing commercially available devices using the N2X Multiservices Test System.

Carrier Ethernet Market Drivers

Carriers are deploying Ethernet for two major reasons: To improve their "top line" by creating new revenue from end-to-end Ethernet services, and to improve their "bottom line" by reducing capital and operational network expenditure.

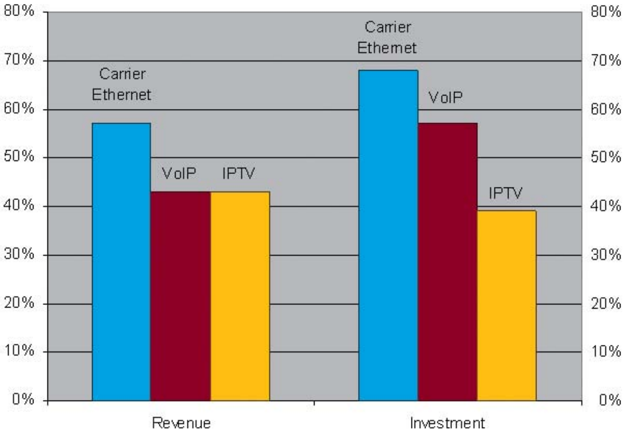

Figure 1: Ethernet is driving carrier investment and revenue expectations even more than VoIP or IPTV

The graphs in figure 1 are derived from the Infonetics report "Service Provider Plans for IP/MPLS: North America, Europe, and Asia Pacific 2006". In a survey of major services providers worldwide, respondents rated the drivers for investment in data networks and new architectures in the next 12 months on a scale of 1 to 7, where 1 is not a driver and 7 is a driver. As can be seen, the rollout of Carrier Ethernet services is up sharply from last year's survey, in which it rated sixth on the list, at 38%.

Surprisingly, Carrier Ethernet is driving service provider investment even more than IPTV and VoIP. In addition, Carrier Ethernet service revenue is expected to grow at a faster rate than IPTV, VoIP and VPNs.

Development & Deployment Challenges

Revenue and cost drivers are spurring equipment developers and standards bodies to invent new protocols and technologies to make Ethernet – traditionally a best-effort Enterprise LAN technology – truly carrier-class. These technologies improve network robustness, scalability and manageability, enabling end-to-end Ethernet services across different network infrastructure technologies such as IP/MPLS and SONET/SDH.

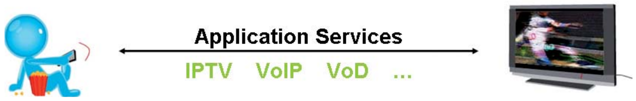

Figure 2A: Application Services

Carrier Ethernet Services
E-LAN & E-Line - MEF 9 & 14 UNI

Figure 2B: Carrier Services

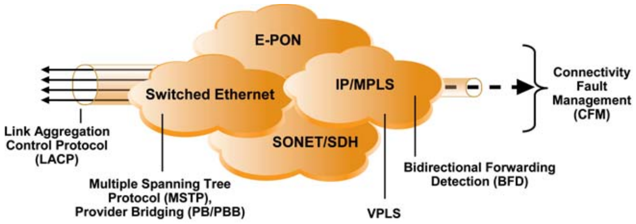
E-PON
Switched Ethernet
IP/MPLS
SONET/SDH
Connectivity Fault Management (CFM)
Link Aggregation Control Protocol (LACP)
Multiple Spanning Tree Protocol (MSTP), Provider Bridging (PB/PBB)
VPLS
Bidirectional Forwarding Detection (BFD)

Figure 2C: Evolving Ethernet infrastructure and service technologies create new implementation challenges

Figures 2 conceptually shows how Ethernet infrastructure and services technologies coexist today:

1. Figure 2A shows the Application Services, such as IPTV, VoIP and Video on Demand. These may operate over Carrier Ethernet services or directly over the infrastructure.
2. Figure 2B shows the Carrier Ethernet Services (E-LAN and E-Line), defined by the Metro Ethernet Forum (MEF), which operate over the different infrastructure technologies.
3. Figure 2C shows the network infrastructure. This includes switched Ethernet, as well as other technologies such as IP/MPLS and LACP. We will introduce these technologies and describe associated implementation challenges.

We now define three types of testing and provide an example to illustrate the need for each:

- Functional and negative testing help prevent functional problems. For example, an unexpected or malformed frame can bring down a service.
- Conformance and interoperability testing help prevent service and device interworking problems. For example, a software upgrade can cause the failure of a provider's network to peer with an upstream service provider.
- Performance and Scalability testing help ensure Quality of Service (QoS) and network robustness, and the ability for services to scale to cope with future customer demands. For example, system loading can cause service delays, incurring Service Level Agreement (SLA) penalties.

Technologies, New Challenges, and Real Test Results

In the next section, we discuss several emerging Carrier Ethernet technologies: CFM, MSTP, BFD, and LACP. We then discuss Carrier Ethernet services, focusing on QoS. We then conclude with one example of an application service, IPTV, which must be delivered to the end user with a high Quality-of-Experience (QoE).

For each, we will review the technology; walk through a typical test scenario; and then present actual test results in which N2X was used to qualify a real device or system.

We will also show the results of two audience polls taken during Agilent's online Carrier Ethernet webinar.

Ethernet OAM and CFM Technology

Ethernet Operations, Administration and Maintenance (E-OAM) is a group of network management functions that provide network fault indication, performance information, and data and diagnosis functions in the context of Ethernet. The IEEE and ITU-T have co-developed two standards for E-OAM: IEEE 802.1ag "Connectivity Fault Management" (CFM) and ITU-T Y.1731 "OAM functions and mechanisms for Ethernet based networks."

Connectivity Fault Management detects, verifies and isolates Ethernet faults from end-to-end. Y.1731 includes similar fault management mechanisms but also defines additional diagnostic and performance management functions.

These standards break down the last barrier to Ethernet adoption as a Carrier-class technology in that they provide OAM mechanisms to ensure network manageability and robustness.

As shown in figure 3, the standards define four message types to enable fault isolation in the horizontal plane, from end to end. These 4 messages are:

- Layer-2 Loopback (effectively an Ethernet MAC address ping)
- Layer-2 Link Trace (similar to IP traceroute)
- Layer-2 Continuity Check (a periodic heartbeat)
- Alarm Indication Signal (AIS)

This concept of domains allows fault isolation in the vertical plane from operator through provider to customer. Figure 3 illustrates domain separation, a key concept defined in the standards.

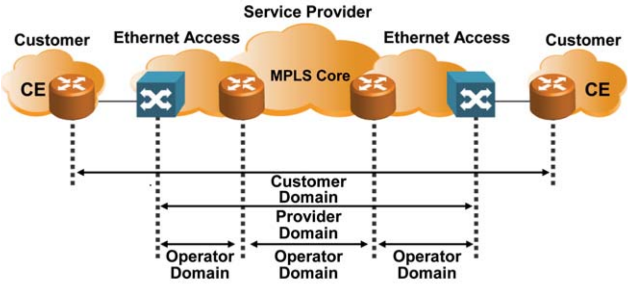

Figure 3: Ethernet fault management provides fault isolation in the operator, provider and customer domains

Given that Ethernet fault management (FM) is a very new technology, recently consented in the ITU-T and currently (at the time of writing) at draft 8.1 in the IEEE, there are plenty of challenges for those who wish to implement this Carrier Ethernet enabler. For example,

- Incompatible FM frames from another vendor – will they cause interoperability issues?
- What is the optimal Continuity Check message timer value, in the presence of other line-card processing?
- Scalability – How many Maintenance Entities (end points and intermediate points) can be supported?
- Are AIS notifications propagated to higher levels? We will now discuss this particular test challenge.

Ethernet OAM and CFM Test Challenges

To avoid confusion, in this section and in figure 4, we use IEEE terminology to describe an Ethernet CFM test challenge.

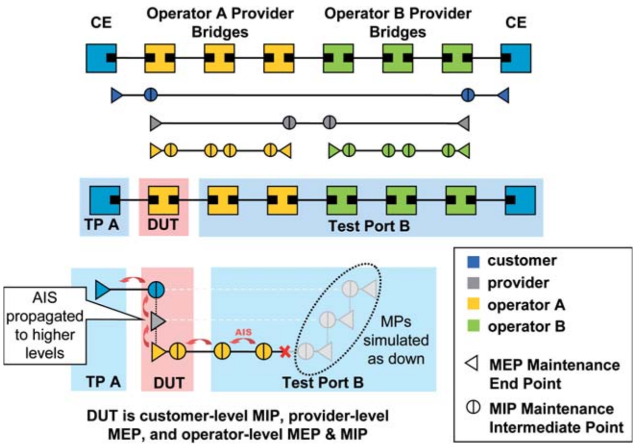
CE
Operator A Provider Bridges
Operator B Provider Bridges
CE
TP A
DUT
Test Port B
AIS propagated to higher levels
AIS
MPs simulated as down
TP A
DUT
Test Port B
DUT is customer-level MIP, provider-level MEP, and operator-level MEP & MIP
customer
provider
operator A
operator B
MEP Maintenance End Point
MIP Maintenance Intermediate Point

Figure 4: Verification of AIS propagation to higher levels

Figure 4 illustrates Maintenance End Points (MEPs) and Maintenance Intermediate Points (MIPs) at different domain levels. For example, at the blue customer level, the end CE devices are MEPs, and the connected ports on the operator devices are MIPs. At the grey provider level, the end Operator bridges (depicted as grey triangles) are MEPs, while the inter-operator bridges (depicted as grey circles) are MIPs. And so on.

The central part of figure 4 conceptually shows the test topology for AIS propagation. The pink-delineated Device-Under-Test (DUT) is surrounded by emulated MIPs and MEPs behind the blue-delineated test ports. To the DUT, the emulated devices behave exactly as real network elements. The value of emulation is that the test engineer can replace a huge and costly testbed of real devices with a single compact test system. Another benefit is that the tester can measure key functional, performance and scalability parameters that would otherwise not be available on a real device operating system.

This conceptual representation is expanded in the lower part of figure 4 to show the 3 vertical domains, using IEEE symbology. A key point to note here is that the DUT is a customer-level MIP, a provider-level MEP, and a operator-level MEP plus MIP.

The objective of this test is to verify that an AIS notification generated at Test Port B propagates up through the 3 DUT levels and is detected at Test Port A. The test steps are listed below.

1. Simulate network fault by suppressing Continuity Check and generating AIS
2. Check AIS propagation to upper levels
3. Verify that upper layers suppress alarms – e.g. Loss of Continuity (LOC)

With a technology as new as CFM, it’s also imperative to verify conformance to the standard. Conformance testing is a prerequisite to ensuring device interoperability. The next section shows some real conformance test results.

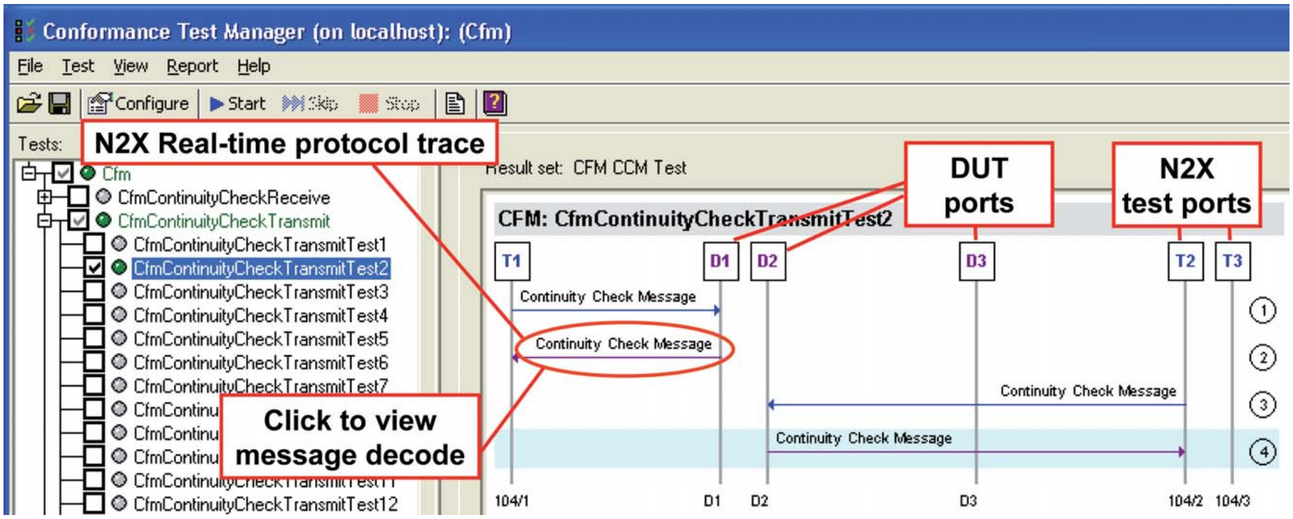

Conformance Test Manager (on localhost): (Cfm)

File Test View Report Help

Configure | Start | Skip | Stop

Tests:

- Cfm
  - CfmContinuityCheckReceive
  - CfmContinuityCheckTransmit
    - □ CfmContinuityCheckTransmitTest1
    - ☑ CfmContinuityCheckTransmitTest2
    - □ CfmContinuityCheckTransmitTest3
    - □ CfmContinuityCheckTransmitTest4
    - □ CfmContinuityCheckTransmitTest5
    - □ CfmContinuityCheckTransmitTest6
    - □ CfmContinuityCheckTransmitTest7
    - □ CfmContinu
    - □ CfmContinu
    - □ CfmContinu
    - □ CfmContinuityCheckTransmitTest11
    - □ CfmContinuityCheckTransmitTest12

Result set: CFM CCM Test

CFM: CfmContinuityCheckTransmitTest2

T1 | D1 | D2 | D3 | T2 | T3

1. Continuity Check Message
2. Continuity Check Message
3. Continuity Check Message
4. Continuity Check Message

104/1 | D1 | D2 | D3 | 104/2 | 104/3

N2X Real-time protocol trace

DUT ports

N2X test ports

Click to view message decode

Figure 5A: N2X verification of one CFM conformance test case for Continuity Check messages

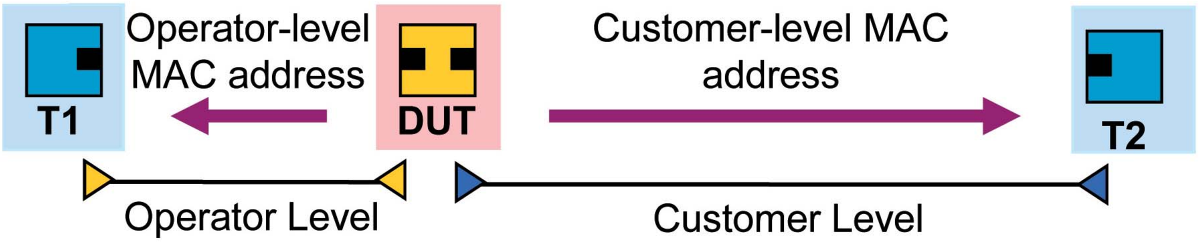

Figure 5B: Protocol trace to the physical test topology.

CFM Conformance Test Results on N2X

Figure 5A shows a screen capture of a conformance test protocol trace from the Agilent N2X test platform.

This particular CFM conformance test verifies that a MEP's Continuity Check (CC) messages are sent to the group MAC address corresponding to the Maintenance Domain level.

In figure 5A, on the left of the N2X Conformance Test Manager, we can see some of the possible test cases, including the CC Transmit Test Case being run. On the right, we can see the real-time protocol traces between the DUT and N2X. As time elapses, successive protocol exchanges are appended on the screen and numbered (1 through 4), as can be seen on the far right.

Figure 5B relates the protocol trace to the physical test topology.

Subsequent releases of device operating software need to be retested for conformance – a process known as regression testing. Failure to do so can cause interoperability issues such as the inability to peer with an upstream provider.

MSTP Technology

The Multiple Spanning Tree Protocol (MSTP) calculates a loop-free LAN topology on multiple VLAN sets in a scalable manner, enabling load balancing and manageability. Figure 6 shows three MSTP regions interconnected by the black lines, delineating the Common and Internal Spanning Tree (CIST).

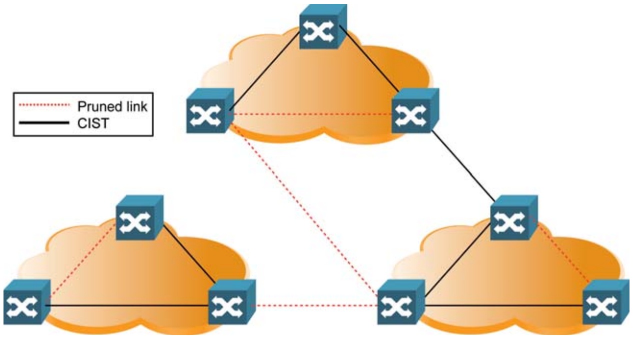

Figure 6: MSTP calculates a loop-free LAN topology on multiple VLAN sets in a scalable manner

Within each region, multiple spanning tree instances (MSTIs) can operate, with each instance carrying an independent set of VLANs. The red lines in figure 6 are links that have been pruned by the protocol to remove loops for a particular MSTI.

Creation of regions assists the network administrator to manage and scale the network. MSTP presents the multiple instances within each region to the 'outside' as a single bridge device.

From a development and deployment perspective, there are several potential problems:

- Scalability – Can the full range of 4,094 VLAN IDs map to the 64 MSTIs?
- MSTP interaction with CFM – do/should blocked ports pass CFM messages?
- Slow MSTP reconvergence – how much packet loss and what impact on service QoE; or reconvergence to lower-bandwidth links – is high-priority traffic given precedence? We will discuss this latter test challenge in the next section.

MSTP Test Challenge

Figure 7 shows a system under test (SUT) in blue and 3 test ports in gold. The test ports emulate bridges, including the root bridge.

The table in figure 7 shows the mapping of two different sets of VLANs onto MSTI 10 and MSTI 20.

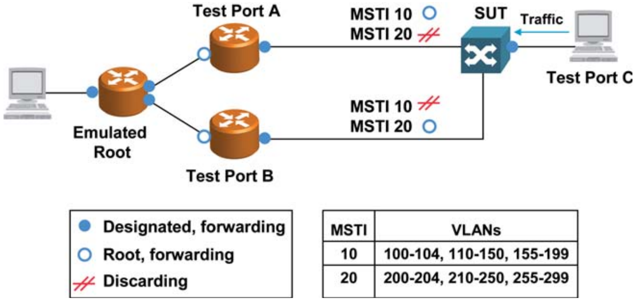

MSTI	VLANs
10	100-104, 110-150, 155-199
20	200-204, 210-250, 255-299

Figure 7: Verifying MST load-balancing and measuring the impact of path cost change

The test objective is to verify MST load-balancing and measure the impact of a path cost change. In this test, we:

- Configure two MSTIs for two sets of VLANs to balance traffic between Test Ports A and B
- Send traffic from Test Port C to the emulated hosts behind the root
- Trigger MST recalculation by modifying the path cost on MSTI 10 from Port A to root, while the test is running
- Verify that MSTI 10 traffic on Test Port A switches to Test Port B
- Measure the packet loss and time to recalculate the spanning tree and restore the services.

We actually ran this test against a real device. The results follow.

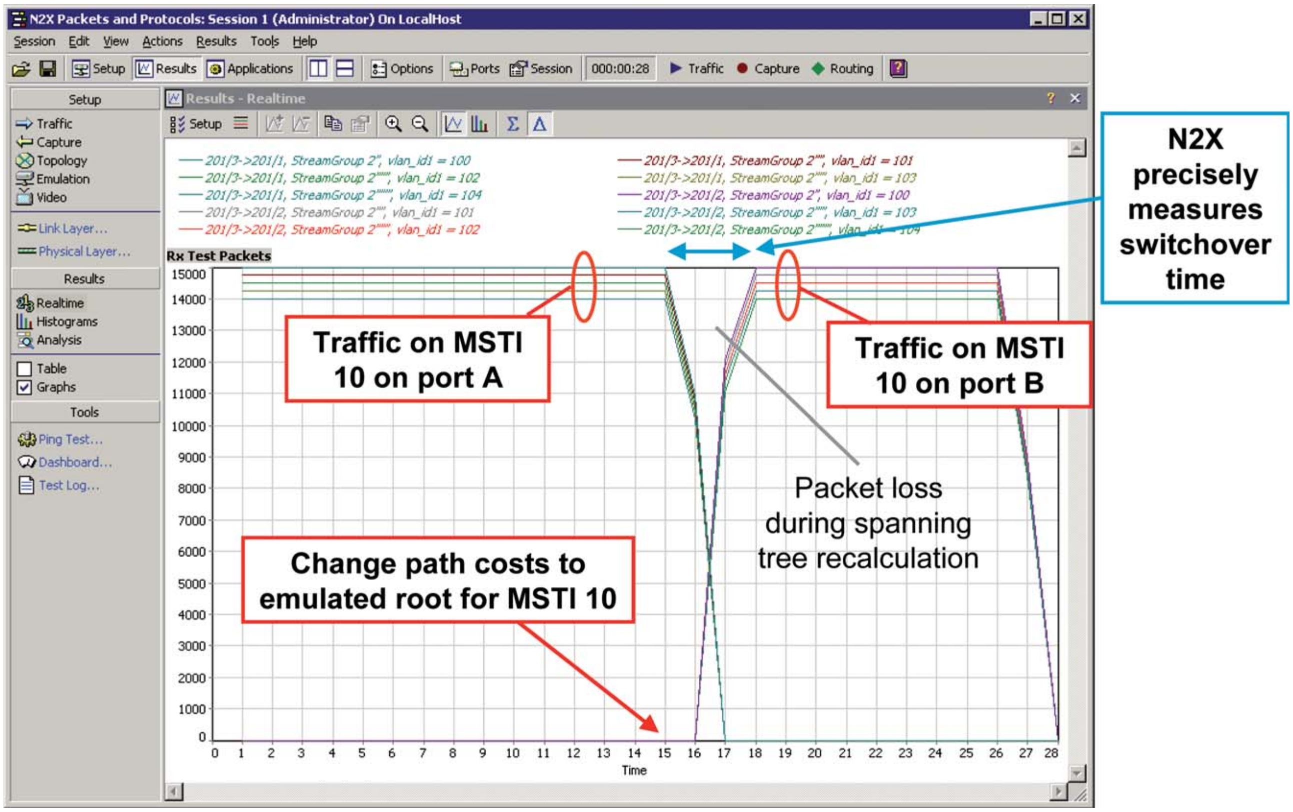
N2X Packets and Protocols: Session 1 (Administrator) On LocalHost
Session Edit View Actions Results Tools Help
Setup Results Applications Options Ports Session 000:00:28 Traffic Capture Routing
Setup
Traffic
Capture
Topology
Emulation
Video
Link Layer...
Physical Layer...
Results
Realtime
Histograms
Analysis
Table
Graphs
Tools
Ping Test...
Dashboard...
Test Log...
Results - Realtime
Setup
201/3->201/1, StreamGroup 2", vlan_id1 = 100
201/3->201/1, StreamGroup 2""""", vlan_id1 = 102
201/3->201/1, StreamGroup 2"""""", vlan_id1 = 104
201/3->201/2, StreamGroup 2"""", vlan_id1 = 101
201/3->201/2, StreamGroup 2""""", vlan_id1 = 102
201/3->201/1, StreamGroup 2"""", vlan_id1 = 101
201/3->201/1, StreamGroup 2"""", vlan_id1 = 103
201/3->201/2, StreamGroup 2", vlan_id1 = 100
201/3->201/2, StreamGroup 2"""", vlan_id1 = 103
201/3->201/2, StreamGroup 2"""""", vlan_id1 = 104
Rx Test Packets
15000
14000
13000
12000
11000
10000
9000
8000
7000
6000
5000
4000
3000
2000
1000
0
0 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28
Time
N2X precisely measures switchover time
Traffic on MSTI 10 on port A
Traffic on MSTI 10 on port B
Packet loss during spanning tree recalculation
Change path costs to emulated root for MSTI 10

Figure 8: N2X measures MSTP switchover time and packet loss

MSTP Test Results on N2X

Figure 8 shows a screen capture of the MSTI load-balancing test from the Agilent N2X test platform.

The graph in figure 8 plots received test packets over time. We are plotting the PDU (Ethernet frames) received on just five of the VLANs on MSTI 10. These are shown in different colors. We have set up these streams at slightly different rates (between 14,000 and 15,000 packets per 1-second interval) so that they are easy to distinguish in the graph.

On the left of figure 8, we see MSTI 10 traffic on port A from the SUT. At time 15 seconds, we change the path costs to the emulated root for MSTI 10. This change causes the SUT to recalculate the spanning tree for MSTI 10, whereupon the traffic to Port A ramps down to zero.

At the same time, the MSTI 10 traffic starts to ramp up on test port B. This switchover takes about 3 seconds, as indicated by the blue arrow. The packet loss that occurs during switchover is shown by the grey shaded area.

At Gigabit data rates, 2 or 3 seconds constitutes significant service outage – far more than the oft-quoted 50 ms restoration time! The switchover time can vary between vendor implementations and would be something significant to quantify.

While MSTP is used in switched networks, MPLS can also be used to transport Carrier Ethernet services. We will now discuss BFD, a technology used in IP/MPLS deployments.

BFD Technology

The Bidirectional Forwarding Detection (BFD) protocol detects Ethernet forwarding plane connectivity via a keep-alive mechanism, which speeds up fault detection and accelerates service restoration. BFD is important for Ethernet, which has no native fault notification mechanism.

Figure 9A shows BFD used in an IP-routed network to trigger fast routing reconvergence, rather than relying on slower timeout approaches inherent in OSPF, IS-IS and IGPs.

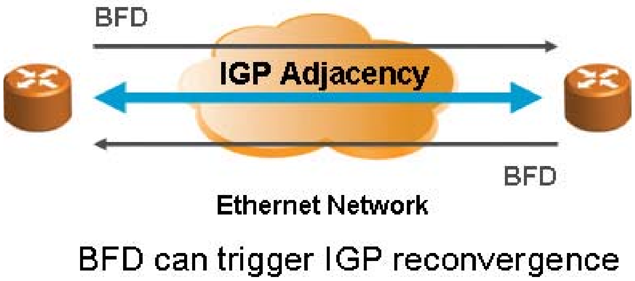

Figure 9A: BFD is used to trigger routing reconvergence

Figure 9B illustrates how BFD can be used to accelerate the triggering of RSVP fast re-route.

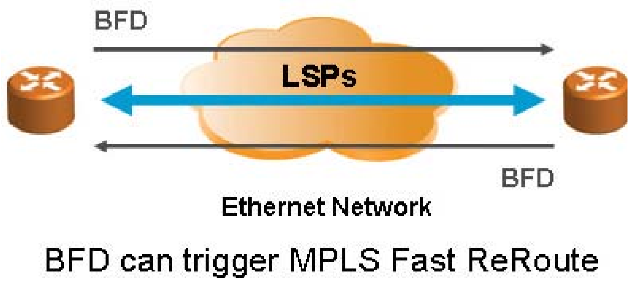

Figure 9B: BFD is used to trigger MPLS Fast Reroute

BFD protection is particularly important when a device's forwarding engine fails but the physical link is still up colloquially speaking, the light is on, but no-one is home and it is the loss of light that would normally trigger fast re-route.

For vendors and service providers, there are several potential BFD pitfalls:

- Scalability – What is the smallest BFD timer value possible for the number of sessions required?
- Incompatible BFD packets from another vendor – will they cause interoperability issues?
- Will different interpretations of evolving draft specifications lead to interoperability issues?
- Can adequate recovery time be achieved to minimize service disruption? We will now discuss this particular test challenge.

BFD Test Challenge – BFD Triggering RSVP Fast Reroute

Figure 10 shows the SUT in blue surrounded by three test ports in gold.

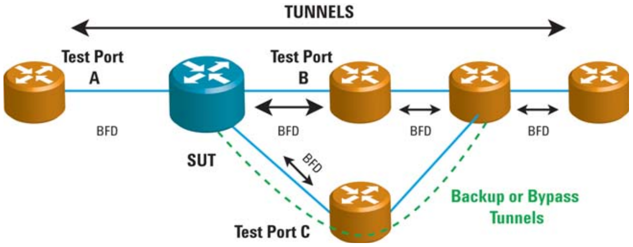

Figure 10: Measuring LSP reroute time after BFD goes down

The objective of this test is to measure LSP reroute time after BFD goes down. In this test:

- A primary Label Switched Path (LSP) is emulated from test port A to test port B, through the DUT, and on to the emulated Label Switched Routers (LSRs) further down the chain.
- A backup or bypass LSP is emulated from test port A to test port C, through the DUT, and on to the same emulated LSRs.
- BFD is used to protect all the LSPs.
- A BFD down event is initiated at test port B and a timestamp T1 is recorded.
- The first MPLS packets received at Port C are timestamped as T2.
- Finally, the MPLS Fast Reroute time is calculated as the difference between T2 and T1.

This test can be scaled to multiple LSPs to find the scaling and performance limits of the SUT. This scaling is depicted graphically in the test results shown in the next section.

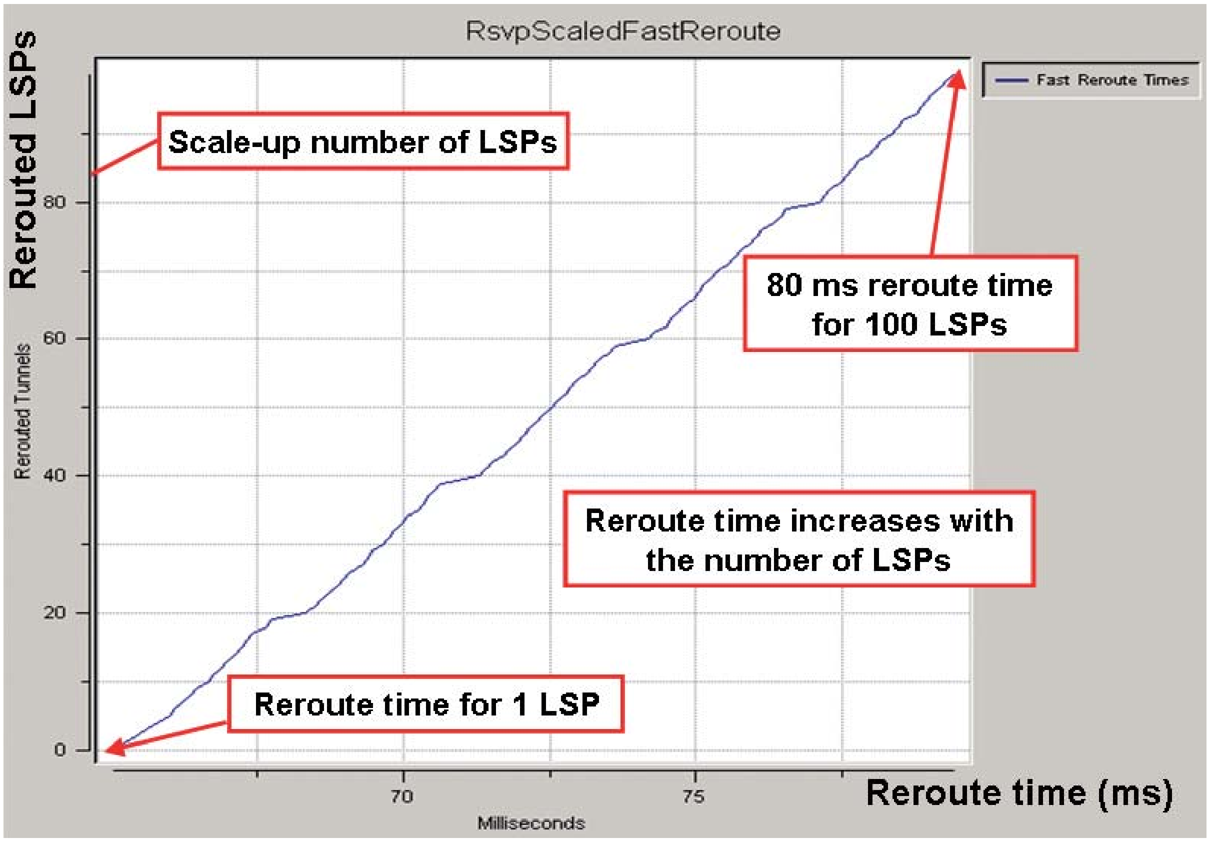

Figure 11: N2X graphs RSVP Reroute time against the number of LSPs

BFD Test Results on N2X

The graph in figure 11 is relatively self-explanatory, and shows that reroute time increases as we scale the number of LSPs.

The reroute time increases linearly and at some point will exceed the acceptable recovery time. In this example, we reached an 80 ms reroute time after only 100 LSPs.

Many people perceive BFD as just a simple keep-alive protocol. However, it is currently defined in 4 different IETF documents that are still evolving, with changes as recent as March 2007.

Highlighting the complexity of this protocol, our BFD conformance test suite contains over 140 test cases. Another BFD complexity is the dependency between scaling and performance; smaller timer values stress the SUT, which then impacts scalability.

Given the importance of this emerging protocol, we were interested in the audience's perception of BFD during our Carrier Ethernet webinar on June 7, 2007, and so, we conducted a live audience poll.

We asked the audience to indicate their timeframe for testing their BFD implementations. The pie chart in figure 12 shows the poll results. The pie chart shows the responses for those participants for whom the poll was applicable.

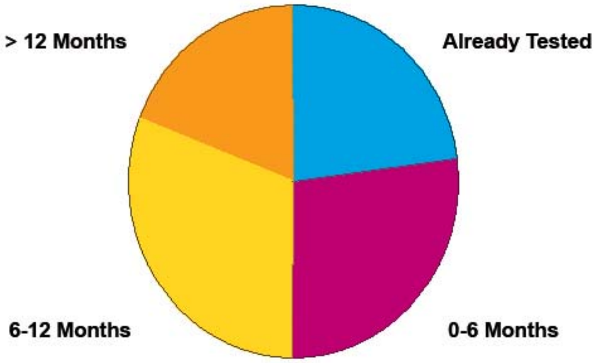

Figure 12: Poll results for BFD testing timeframe

This poll clearly shows that BFD is still in the early stages of adoption and, like all new technologies, requires testing to avoid the pain and realize the profits of Carrier Ethernet.

LACP Technology

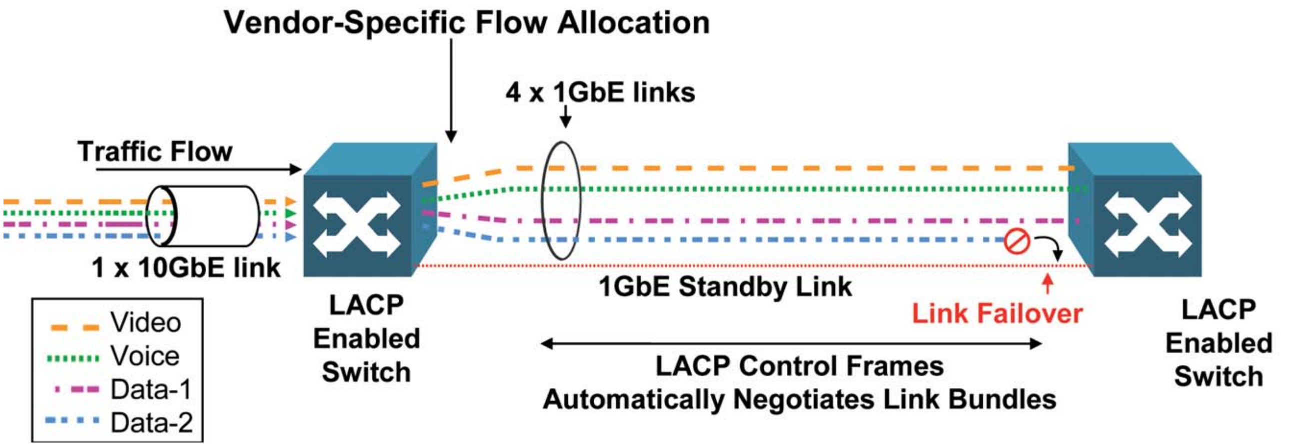

Figure 13: LACP bundles physical links into a single logical link for bandwidth flexibility and high reliability

Link Aggregation is a technology that bundles several physical Ethernet links into a single, logical link. This allows a service provider to provision links at speeds other than 10/100, GbE and 10 GbE, and it makes better use of the available bandwidth.

For example, a provider can set up a link bundle of four GbE links, as illustrated in figure 13. Link Aggregation Control Protocol (LACP) is the IEEE-defined control protocol that is used to negotiate and setup a link bundle, also known as a Link Aggregation Group (LAG).

Each traffic flow from the upstream 10GbE link (on the left of figure 13) is allocated by the switch to one of the physical links in the bundle on the right. Flow allocation is vendor-specific – it is not specified in the IEEE standard – so that is an important area to test because a poor implementation can split a flow between links and cause frame mis-sequencing.

The other benefit of LACP is High Availability. If a link goes down, the flows on that link are reallocated onto the other active links. Alternatively, the flows can be reallocated onto a hot standby link, as shown in red in figure 14.

The greatest risk for service providers is in regard to honoring SLAs: How quickly are the flows reallocated onto the other links? How many packets are lost? Is there a spike in latency? How much slower is reallocation if the traffic is scaled?

In the next section, we will discuss the measurement of link failover.

LACP Test Challenge – Simulate link failure and measure reconvergence

Figure 14 is the same as figure 13 except the system under test is connected not to a second switch, but rather to four test ports that emulate a second switch.

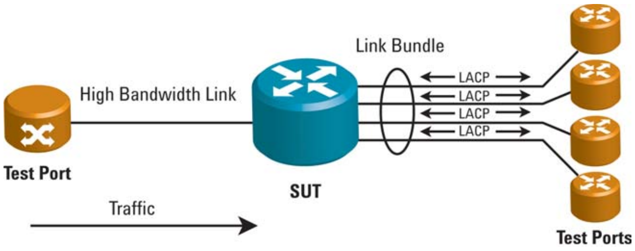

Figure 14: Measuring the time taken to reallocate traffic flows, when a link fails

The objective of this test is to measure the time taken to reallocate traffic flows when a link fails. In this test:

- A link bundle is created and traffic is sent through the DUT with multiple flows, evenly distributed by the SUT
- The laser is turned off to simulate link failure, and the ‘down’ event is time-stamped
- The time taken to reallocate traffic onto alternative active or standby links is then measured
- Packet loss incurred during reallocation is also measured

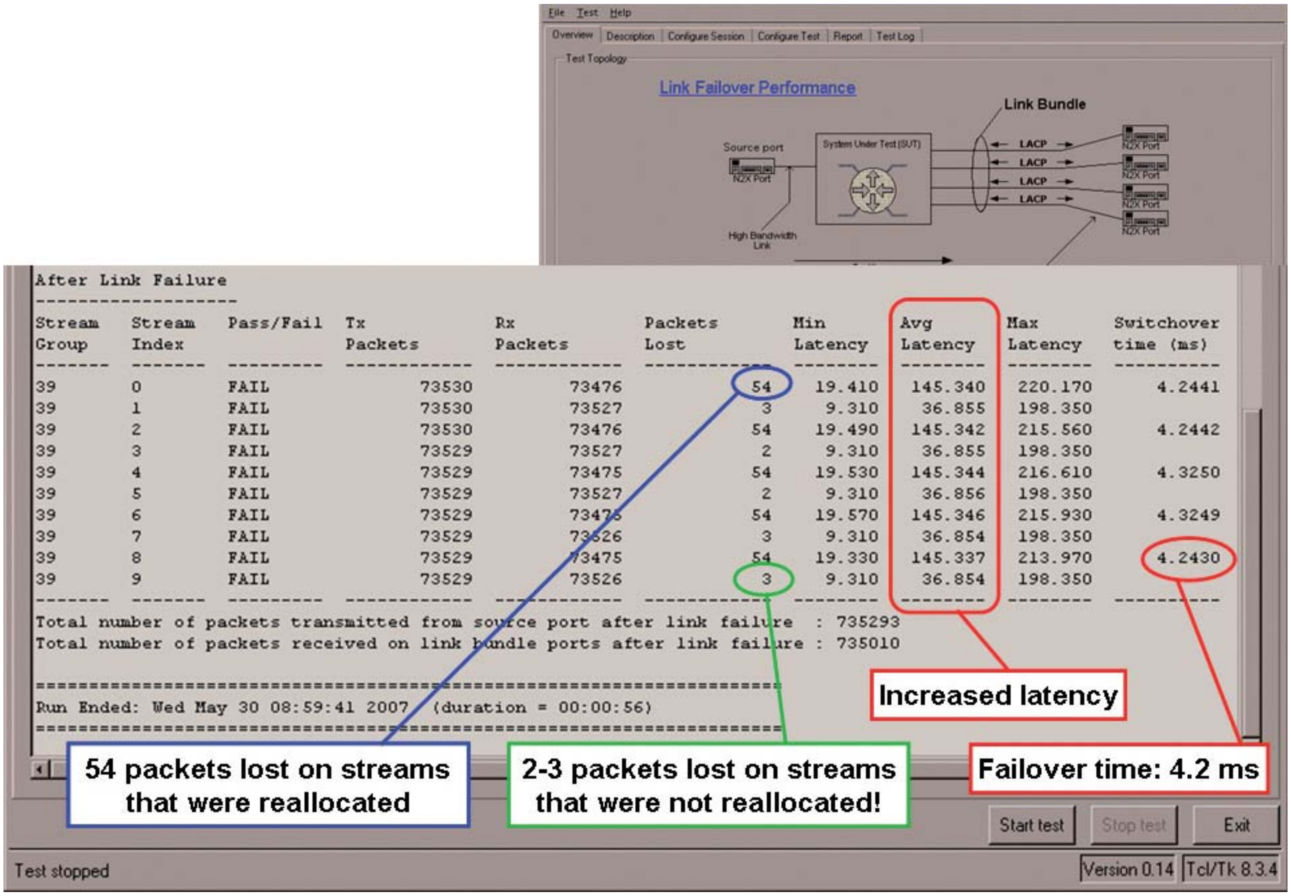

```
After Link Failure
--------------------
Stream   Stream   Pass/Fail  Tx            Rx            Packets       Min       Avg       Max       Switchover
Group    Index               Packets       Packets       Lost          Latency   Latency   Latency   time (ms)
-------  -------  ---------  ------------  ------------  ------------  --------  --------  --------  ----------
39       0        FAIL              73530         73476            54    19.410   145.340   220.170      4.2441
39       1        FAIL              73530         73527             3     9.310    36.855   198.350
39       2        FAIL              73530         73476            54    19.490   145.342   215.560      4.2442
39       3        FAIL              73529         73527             2     9.310    36.855   198.350
39       4        FAIL              73529         73475            54    19.530   145.344   216.610      4.3250
39       5        FAIL              73529         73527             2     9.310    36.856   198.350
39       6        FAIL              73529         73475            54    19.570   145.346   215.930      4.3249
39       7        FAIL              73529         73526             3     9.310    36.854   198.350
39       8        FAIL              73529         73475            54    19.330   145.337   213.970      4.2430
39       9        FAIL              73529         73526             3     9.310    36.854   198.350
-------  -------  ---------  ------------  ------------  ------------  --------  --------  --------  ----------
Total number of packets transmitted from source port after link failure  : 735293
Total number of packets received on link bundle ports after link failure : 735010

=======================================================================
Run Ended: Wed May 30 08:59:41 2007  (duration = 00:00:56)
=======================================================================
```

Figure 15: N2X LACP QuickTest performed against a well-known, commercially available switch

It is possible that following a failure, the remaining links of a link bundle become over-subscribed. Therefore, an important variation of this test is to over-subscribe the traffic after the failure, which forces the SUT to discard low-priority flows. This is important for meeting SLAs for revenue-bearing traffic or mission-critical applications.

In the next section, we highlight some actual results from performing this test.

LACP Test Results

The results of the LACP test are shown in figure 15. This is not smoke-and-mirrors; these are actual results from testing a commercially available carrier-grade device. This is disturbing because it uncovers unexpected behavior.

To obtain these results, we used one of the automated applications called 'QuickTests' that run on the Agilent N2X test platform, as shown in figure 15. This QuickTest makes test configuration, execution and results reporting fast and easy.

Notice the following four points in figure 15:

- In red text box on the far right, the failover time was 4.2 ms. This may be adequate for voice, but probably not for video.
- In the other red text box, we can see that the average frame latency of the re-allocated flows increased from 10 microseconds to 145 microseconds, an order-of-magnitude increase:
- In the blue text box at the left, we can see that 54 packets were lost on each of the reallocated flows.
- Surprisingly, in the green text box, we can see that even the other flows that were not reallocated suffered some packet loss, even though the link bundle was not over-subscribed. So if this particular device was chosen for a carrier network, it is very possible that all customer traffic on these LAGs will be affected.

So here we've seen frame loss in an automated test. It's also powerful to quantify and visualize frame performance in real-time, as we'll see in the next section.

Carrier Ethernet Services QoS Testing on N2X

To facilitate SLAs, carriers always prioritize their customers' traffic, as conceptually shown by the gold, silver and bronze pipes on the right of figure 16. Consequently, it is important to ensure that network devices can meet SLA requirements – especially for the gold (premium) traffic class

Ideally, the gold traffic will have the lowest frame delay variation – in other words, jitter – as shown in the small histogram representation at the top of figure 16.

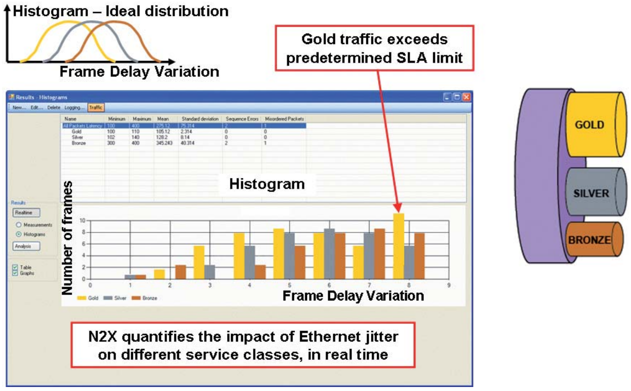
Histogram – Ideal distribution
Frame Delay Variation
Gold traffic exceeds predetermined SLA limit
Results - Histograms
New... Edit... Delete Logging... Traffic
Histogram
Number of frames
Frame Delay Variation
Results
Realtime
Measurements
Histograms
Analysis
Table
Graphs
Gold Silver Bronze
GOLD
SILVER
BRONZE
N2X quantifies the impact of Ethernet jitter on different service classes, in real time

Name	Minimum	Maximum	Mean	Standard deviation	Sequence Errors	Misordered Packets
All Packets Latency	100	400	275.12	75.314	2	1
Gold	100	110	105.12	2.314	0	0
Silver	102	140	128.2	8.14	0	0
Bronze	300	400	345.243	40.314	2	1

Figure 16: N2X real-time jitter histogram shows whether E-LAN & E-Line Services SLAs can be met

However, it is possible that under certain conditions, (such as the LACP failover event shown previously), a network device will violate an SLA.

By displaying a continuously updating histogram of the jitter of all three traffic classes, it's easy to visualize disturbances like this, in real time.

Poll Result – Carrier Ethernet Technologies

We polled our customers during a recent Carrier Ethernet webinar, asking them which Carrier Ethernet technology will cause the most implementation pain, if they don't adequately test. The pie chart below shows the results.

As shown by the equal distributions of poll results, all of these technologies are important – depending on the stage of device or network implementation. However, the key thing to note is that, as shown earlier in figure 2, it is critical to test all these technologies together to ensure that they all coexist.

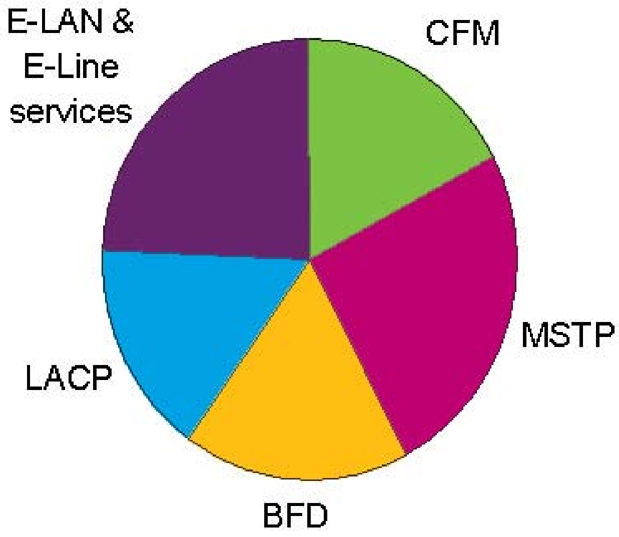

Figure 17: Poll results showing the Carrier Ethernet technologies that cause the most implementation pain

IPTV Services over Ethernet – Technology

It is essential to test the impact of frame impairments occurring at the infrastructure layer on the user's experience at the application layer. This is called Quality of Experience or QoE.

The IETF RFC 4445 defines a method for quantifying QoE and an associated metric called Media Delivery Index (MDI) that has two components:

- Media Loss Rate (MLR), which is related to frame loss.
- Delay Factor (DF), which is related to jitter and the maximum required buffer needed to smooth that jitter (and avoid buffer overflow and underflow)

In practical terms, this means that if a video stream suffers a maximum Delay Factor of 50 ms, then the customer Set Top Box must have at least a 50 ms buffer to smooth the jitter and see a perfect picture.

MDI has two components:
- Delay Factor (DF)
- Media Loss Rate (MLR)

Service	Maximum acceptable Delay Factor (DF)
All	9 – 50 ms

Service	Maximum acceptable Media Loss Rate (MLR)
SDTV	0.004 packets/sec
VOD	0.004 packets/sec
HDTV	0.0005 packets/sec

Typical MLR target: 1 packet dropped per 30 minutes for SDTV programming
Ideal: Zero packet loss.

Possible Dis-eases

Ethernet frame loss, latency and jitter will contribute to MDI

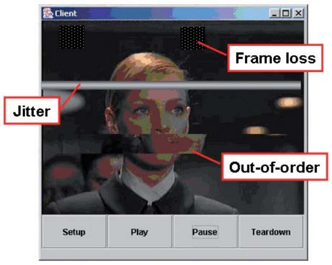

Figure 18: MDI measures the impact of network impairments on QoE – the end user's perception of service

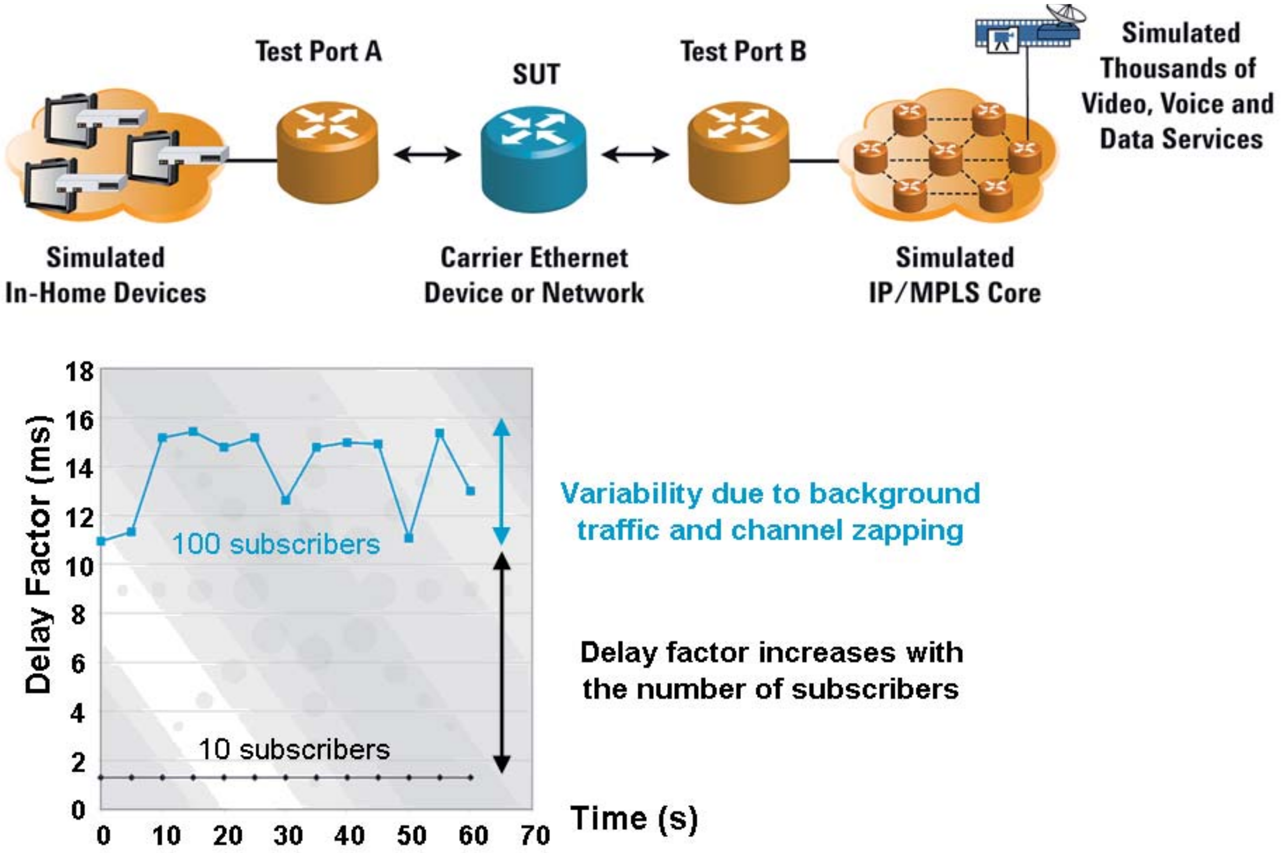

Figure 19: Measuring MDI Delay Factor on the Agilent N2X

The tables on the left of figure 18 on the previous page show typical MDI targets for various video services, such as Standard Definition TV, Video on Demand, and High Definition TV.

As shown in the impaired video sample in figure 18, Ethernet frame loss, latency and jitter will contribute to MDI and reduce the Quality of Experience.

In the next section, we measure the Delay Factor performance of a real device.

IPTV Quality of Experience Test Results on N2X

In this test, we simulate thousands of services and multiple subscribers.On the right of figure 19, behind Test Port B, we simulate thousands of voice, video and data services.

On the left, behind Test Port A, we simulate multiple subscribers (with in-home devices) who are receiving multiple IPTV services and rapidly changing (also known as zapping) channels.

With only 10 subscribers, as we can see from the black line on the graph, the Delay Factor is less than 2 ms. However, with 100 simulated subscribers, the Delay Factor increases to a maximum of 15 ms.

If several devices in the IPTV delivery chain introduced a similar delay, the compounded Delay Factor could easily exceed the recommended 50 ms threshold for high-end Set Top Boxes.

N2X Carrier Ethernet Test Solution

In the previous sections, we presented a technology overview, highlighted key test challenges, and described typical test scenarios, concluding with real test results, for the six technologies CFM, MSTP, BFD, LACP, QoS and IPTV.

In order to perform these tests and profit without the pain, two things are needed: Test plans and Carrier Ethernet test equipment.

As shown in figure 20, Agilent's Journal of Internet Test Methodologies is a book, compiled over several years from real customer experiences, of over 130 test cases.

This document is regularly updated with additional test scenarios. Test engineers can use this as the building blocks to develop their own test plans.

The Journal can be downloaded for free by registering at www.agilent.com/find/thejournal. Additionally, the Agilent INSIGHT eMagazine is available for free subscription, containing the latest Journal Test Cases and technology papers.

The Journal test cases can be performed on the Agilent N2X Multiservices test system, shown in figure 21. More information on N2X can be obtained from www.agilent.com/find/n2x.

Many of the Journal test cases are implemented as N2X QuickTests, as we've seen in the LACP test results. These QuickTests automate common test scenarios, making test configuration, execution and results reporting fast and easy.

The table in figure 20 shows that N2X provides emulation and conformance capabilities for all of the Carrier Ethernet technologies discussed in this paper.

N2X also allows you to test device performance, scalability, interoperability and robustness for other layer-2 to 7 technologies such as VPLS, MPLS, routing and access protocols.

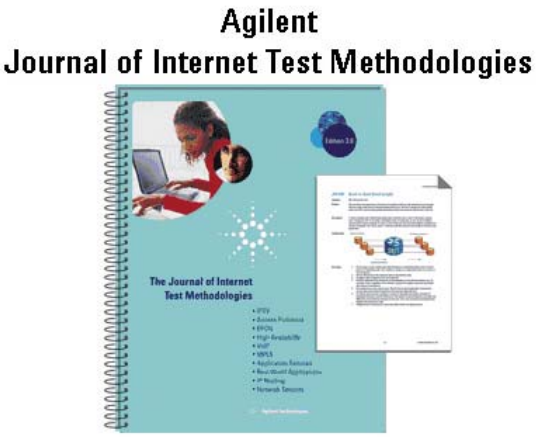
Agilent
Journal of Internet Test Methodologies
The Journal of Internet Test Methodologies

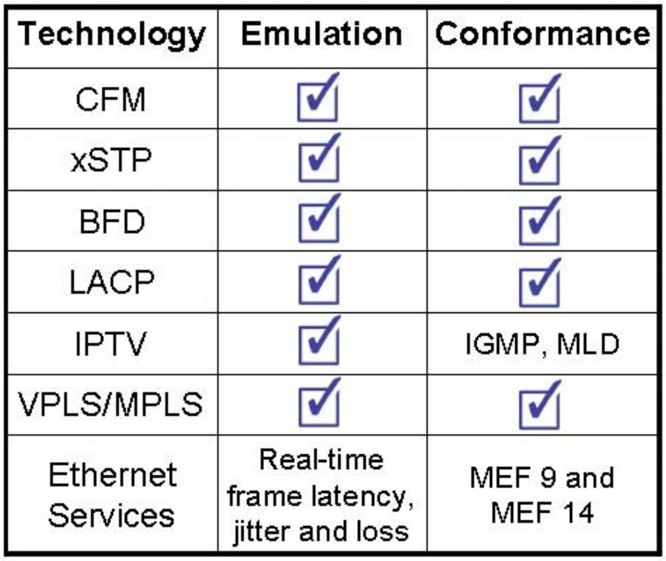

Technology	Emulation	Conformance
CFM	☑	☑
xSTP	☑	☑
BFD	☑	☑
LACP	☑	☑
IPTV	☑	IGMP, MLD
VPLS/MPLS	☑	☑
Ethernet Services	Real-time frame latency, jitter and loss	MEF 9 and MEF 14

Figure 20: The Agilent Journal of Test Cases and N2X Carrier Ethernet test capabilities (July 2007)

Figure 21: N2X Multiservices Test System

Conclusion

Ethernet is increasingly playing a strategic role as service providers open their networks to the revenue promises and cost reductions of Ethernet-based services and infrastructures.

There are many co-existing technologies that are evolving to make Ethernet carrier-class. The technologies present unique challenges that must be tested in isolation and concurrently.

To avoid Carrier Ethernet implementation pain and realize the profits, register for Agilent's Journal of Internet Test Methodologies and INSIGHT eMagazine, and integrate the Agilent N2X Multiservices Test Solution into your testbed.

www.agilent.com/find/n2x

Sales, Service and Support

United States:
Agilent Technologies
Test and Measurement Call Center
P.O. Box 4026
Englewood, CO 80155-4026
1-800-452-4844

Canada:
Agilent Technologies Canada Inc.
2660 Matheson Blvd. E
Mississauga, Ontario
L4W 5M2
1-877-894-4414

Europe:
Agilent Technologies
European Marketing Organisation
P.O. Box 999
1180 AZ Amstelveen
The Netherlands
(31 20) 547-2323

United Kingdom
07004 666666

Japan:
Agilent Technologies Japan Ltd.
Measurement Assistance Center
9-1, Takakura-Cho, Hachioji-Shi,
Tokyo 192-8510, Japan
Tel: (81) 426-56-7832
Fax: (81) 426-56-7840

Latin America:
Agilent Technologies
Latin American Region Headquarters
5200 Blue Lagoon Drive, Suite #950
Miami, Florida 33126
U.S.A.
Tel: (305) 269-7500
Fax: (305) 267-4286

Asia Pacific:
Agilent Technologies
19/F, Cityplaza One, 1111 King's Road,
Taikoo Shing, Hong Kong, SAR
Tel: (852) 3197-7777
Fax: (852) 2506-9233

Australia/New Zealand:
Agilent Technologies Australia Pty Ltd
347 Burwood Highway
Forest Hill, Victoria 3131
Tel: 1-800-629-485 (Australia)
Fax: (61-3) 9272-0749
Tel: 0-800-738-378 (New Zealand)
Fax: (64-4) 802-6881

This information is subject to change without notice.
Printed on recycled paper

Printed in USA June 25, 2007
5989-7048EN

Agilent Technologies