

MPLS & ETHERNET World congress 2010

Public Multi-Vendor
Interoperability Test Event
MPLS & Carrier Ethernet

■ EANTC ■

Paris, February 2010

EDITOR'S NOTE



Carsten Rossenhövel
Managing Director
EANTC AG

Most corporations, when they are just as successful as they always wanted to be, become restless. They start messing around with their logo, change their name (usually to something starting with »A«) and generally believe they should modify something to increase business.

Some of the more successful ones, though, often come out of this exercise strengthened by just changing — nothing at all.

This applies to MPLS as well. Two or three years ago, vendors got restless and proposed a *new world* of packet transport networks under various names. There were major discussions here at the MPLS & Ethernet World Congress in Paris; some people thought it was time for MPLS to go.

Today, we are back with another MPLS and Carrier Ethernet interoperability event. Certainly some of the new ideas helped the industry to evolve (around protection, for example), some helped to make the technology more attractive for new customer groups (such as MPLS-TP), but in general we have seen a stronger focus on advanced MPLS testing this time than in previous years.

The technology and its Carrier Ethernet applications are approaching the next level. Let's be honest — many new technologies lack the depth, breadth and robustness required for many applications. The fact that MPLS and Carrier Ethernet are well established today means that the industry can rely on a rock-solid, widely supported technology platform with solutions for almost every market.

The topics we have tested this year evolve around maturing mobile backhaul — a very demanding application when it comes to synchronization. We focused the integration of packet and optical networks, an area with a number of technology options and with ongoing debates about the optimal solution. We re-evaluated IPv6 after a long time, now that demand is growing, and found that all the bits and pieces for VPN transport are there (v6-to-v4 NAT yet to come as a test topic). We tested a number of pre-standard draft MPLS-TP implementations and are waiting for the standards bodies to agree on additional specifications to be evaluated.

Would I have imagined in 2003, when we presented the results of our first interoperability event with four racks in Paris, to be back in 2010 with the largest ever MPLS and Carrier Ethernet interoperability event, in terms of rack space (16), power (55 kW) and noise (>79 dB)? I don't think so. We are definitely looking forward to green technology and higher port densities in the future!

INTRODUCTION

Over the past ten years, Multi-Protocol Label Switching (MPLS) has become a technology of vital importance in the networking industry. The Internet Engineering Task Force (IETF) has not been sitting on its laurels, but rather continues to expand the protocol families surrounding MPLS, to extend the technology uses, and to embed more features. We see the results of this continuing development in the diverse set of topics chosen by the participants from a set of options proposed by us (EANTC). This year's event topics allowed us to explore the broad use cases that exist for MPLS today, and to test interoperability for the features that vendors felt are of importance today. By December 2009 when we summarized our test plan discussions, the test topics were:

- **Multicast MPLS Services** - A topic which is still under development process in many labs. We test the tools vendors have available today to provide multicast, or point-to-multipoint, connectivity in an MPLS network.
- **Integration of Packet and Optical Networks** - There is a lot of talk about different solutions for packet optical networks, so without specifying any technology, we asked vendors to propose a solution for integrating packet based MPLS networks, with Dense Wave Division Multiplexing (DWDM) networks.
- **Mobile Backhaul** - As expected, the mobile space is increasing in relevance and importance, and service providers are starting to deploy the tools necessary to scale their mobile data in an efficient way. How do they do that? By moving to packet based backhaul networks.
- **MPLS and Ethernet Based Transport and Access** - Under this test section we tested new MPLS features applicable to several use cases - IPv6 VPN Services, MPLS-TP, Ethernet Ring Protection Switching (ERPS) and more.
- **Management** - As in every event, we encourage vendors to test their systems that can manage, provision, and monitor services running on equipment from multiple suppliers; interoperable management solutions.

The variety of tests in the agenda proved to be a challenge to complete in the two weeks of hot staging. If, as the reader, you are searching for a particular test combination (Vendor A with Vendor B) but do not find it in the following report, it may not necessarily be the case that a test failed or a vendor did not have the proper support, it could have been simply a factor of time. In any case, what the team was able to accomplish was not trivial and allows for some interesting reporting as follows.

TABLE OF CONTENTS

Participants and Devices	3
Demonstration Network Design.....	3
Interoperability Test Results.....	3
Multicast Services.....	4
MPLS Integration with Optical Transmission...	5
Synchronization in the Backhaul	6
Legacy Mobile Backhaul Transport.....	9
MPLS and Ethernet Transport and Access....	12
Management.....	15

PARTICIPANTS AND DEVICES

Vendor	Devices
Alcatel-Lucent	1850 TSS-160 1850 TSS-320 1850 TSS-5 9500 MPR
Brocade	MLX-8 NetIron CER 2024C
Calnex	Paragon
Cisco	7604 7606 ASR 1000 ASR 9000 CRS-1
Ericsson	OMS 1410 SmartEdge1200 SmartEdge100
Huawei	CX600-X3 GPON&DSLAM MA5600T NE40E-X8 NE40E-X16 NE5000E OSN1500 OSN3500 OSN6800 PTN910 PTN950 PTN3900 RTN910 RTN950 S9300 U2000 U2520
Ixia	IxN2X IxNetwork
MRV	OS904 OS904MBH OS910M OS9124-410G
Orckit-Corrigent	CM-11 CM-4140 CM-4206
RAD Data Communications	ACE-3220 ETX-204A IPmux-1.55L IPmux-24

Vendor	Devices
Spirent Communication	Spirent TestCenter Spirent GEM Spirent xGEM
Telco Systems - a BATM Company	T-Marc-280 T-Marc-380 T-Metro-XG T-Metro-200 T5C-XG T5C-XG-ES T-Marc-254H T-Marc-254P
ZTE	ZXCTN 6100 ZXCTN 6300 ZXCTN 9008 ZXR10 5928E ZXR10 8905 ZXR10 M6000 ZXR10 T1200

DEMONSTRATION NETWORK DESIGN

While some test scenarios had to be tested back to back in an isolated scenario, we tried as much as possible to perform the tests in a single multi-vendor network, shown as the Physical Topology in the centerfold of this document. An MPLS core was built which facilitated all Layer 3 VPN related tests. Devices for the Layer 2 VPN tests were included in the MPLS aggregation network, which was an extension of the MPLS core. Additionally there was an MPLS-TP network and a special area where optical transport tests were demonstrated. Finally nodes connecting at the edge and all nodes considered as cell site, customer site, or last mile equipment are classified as "Access Devices" and fit into the edge access network.

INTEROPERABILITY TEST RESULTS

In the following sections of the white paper we describe the test areas and results of the interoperability event. The document generally follows the structure of the test plan.

Terminology. We use the term "tested" when reporting on multi-vendor interoperability tests. The term "demonstrated" refers to scenarios where a service or protocol was terminated by equipment from a single vendor on both ends.

Test Equipment. In order to run our tests we required the ability to generate, measure, impair, and analyze Ethernet traffic as well as analyze synchronization. We are indebted to Calnex Solutions, Ixia, and Spirent Communications for their equipment and support throughout the hot staging.

MULTICAST SERVICES

Broadcast TV has been transported over IP packets for several years under the service name IPTV. That signal that used to come from antennas, satellites or cable is now offered as an integral part of triple play service offerings. A report posted by the French telecom regulator ARCEP has stated that the penetration of IPTV services has reached 34.6% of the ADSL lines in the country. It stated that there were 5.643 million IPTV subscribers and 16.299 million ADSL subscribers in the third quarter of 2008.

Multicast, or point-to-multipoint services, allows service providers to optimize network resources by sending only one copy of the data through the network until the point closest to the subscriber. Then the network replicates the data to the viewers that explicitly requested it. Multicast lends itself perfectly to the distribution of broadcast TV as well as other single source to multiple receivers services.

IP multicast has been around for over a decade, however, interest in multicast architectures in MPLS has become a topic of more recent interest. Given the numbers presented above it is also easy to understand why: service providers have been using MPLS in their networks for some 10 years and wish to continue using these network to provide TV services. We take a look at where vendors are, and what is currently interoperable in this regard.

Layer 3 Multicast Virtual Private Networks

Virtual Private Networks, built using MPLS and BGP, are currently deployed by many service providers to offer layer 3 services to their enterprise customers. Nowadays, providers are looking into applications to enrich their services and increase revenues. The increase of bandwidth from point-to-multipoint applications such as video means efficient use of the network is required including in enterprise scenarios.

We successfully tested multicast VPNs based on GRE for the following devices: Cisco 7604, Cisco CRS-1, Ericsson SE1200, Huawei NE40E-X16, Huawei NE5000E, ZTE ZXR10 M6000, ZTE ZXR10 T1200. For multicast signaling PIM-SM was mostly used, as shown in the diagram. In addition we tested multicast VPN based on GRE and PIM-SSM signaling for the following devices: Cisco 7604, Cisco ASR 1000, Cisco ASR 9000, Cisco CRS-1, Ixia IxNetwork, Spirent TestCenter.

We verified the multicast connectivity by sending IGMP joins and leaves into the PE from the customer link and verifying that only the joined ports received the traffic.

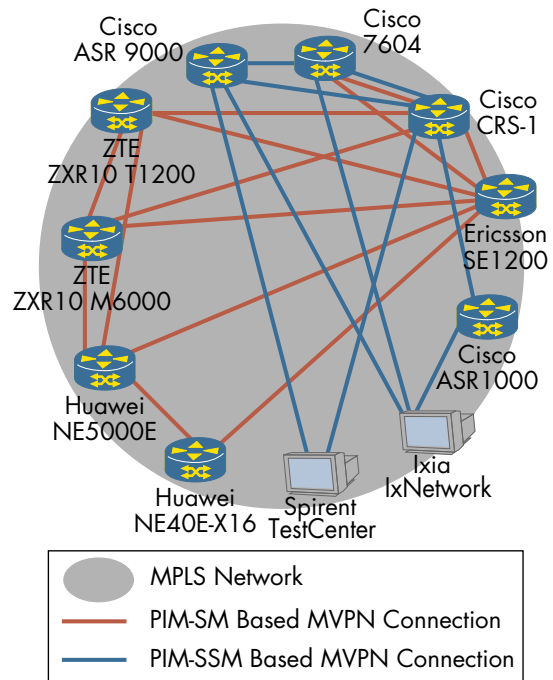


Figure 1: Multicast VPN

Multicast in VPLS Networks

A few methodologies have been defined for a proper multicast behavior in VPLS networks. In the following test the devices did not go so far as to build complex multicast trees, however we did test interoperability for LDP based VPLS signaling in conjunction with IGMP snooping to constrain traffic to PEs without multicast receivers. As the diagram shows, several devices were tested in several combinations. The arrows depict where traffic was sourced and egressed, amongst all scenarios tested (thus it is the inverse of a diagram depicting where IGMP joins were sent).

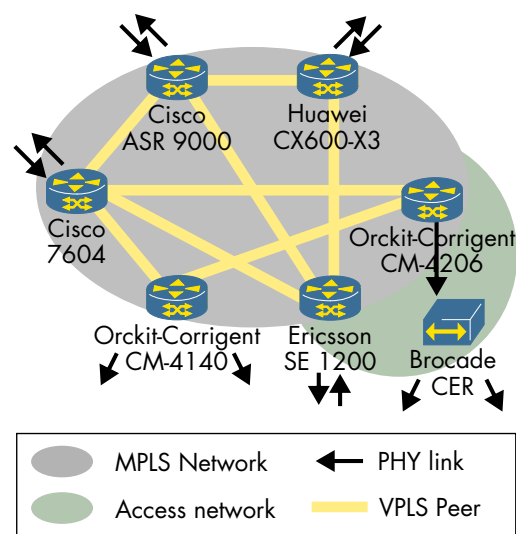


Figure 2: Multicast over VPLS

Point to Multipoint RSVP-TE

A familiar protocol to those engineering their MPLS networks, ReSource reservation Protocol with Traffic Engineering can be used to provision and signal a point-to-point MPLS tunnel through the network. In response to a growing need for point-to-multipoint (p2mp) services, the IETF has defined extensions to RSVP-TE to solve this. Implementations from core and aggregation router vendors are quite new and therefore a topic of particular interest right now.

The test was setup such that an ingress MPLS Label Switched Router (LSR) was directly connected to a branch LSR (point where tunnel splits and traffic is replicated), and the branch LSR was connected directly to two egress LSRs. The devices under test were Cisco 7604, Cisco 7606, Cisco CRS-1, and Ixia IxNetwork which emulated an MPLS node.

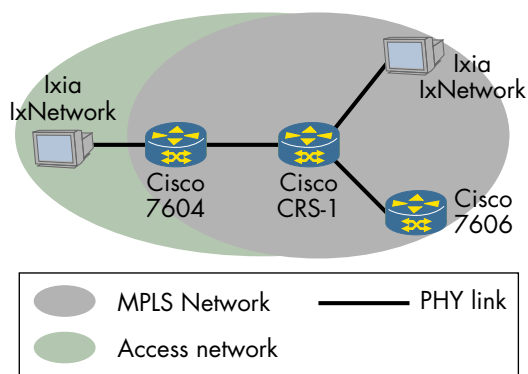


Figure 3: Example P2MP RSVP-TE Scenario

Three scenarios were tested allowing both the 7604 and CRS-1 to each be tested as both ingress LSR and branch LSR, IxNetwork to test as emulated ingress LSR and emulated egress LSR, and the 7606 to be tested as an egress LSR. One of the three scenarios is shown in the Figure. Ixia IxNetwork was used to emulate traffic coming from behind the emulated routers, unless no routers were emulated in which case the Ixia transmitted IP packets into the Cisco devices. We verified that the devices under test could establish a p2mp LSP, replicate traffic at the branch node, and add/remove a source-to-leaf (S2L) sub-LSP on the existing p2mp LSP. In all scenarios each of the three verifications were successful.

MPLS INTEGRATION WITH OPTICAL TRANSMISSION

When we started discussion for this test our requirements were flexible. We were hoping that the vendors who have been stating support and solutions for integrating optical infrastructures with MPLS networks would bring their solutions to test with one another. The IETF has done much work to define Generalized MPLS (GMPLS), which implements ITU-T ASON control plane architecture. One important aspect is that different data plane technologies may be managed with this solution. The data planes used by participating vendors included

Optical Transport Network (OTN) (Huawei) and pre-standard MPLS-TP (Alcatel-Lucent).

Generalized Multiprotocol Label Switching (IETF) defines both routing and signaling protocols for the creation of Label Switched Paths in networks built using various transport technologies. There are in total three models for deploying GMPLS: The first is called the peer model, where the IP network and optical networks are considered an integrated network with a unified control plane. Then there is the overlay model, where the IP network and optical networks operate completely independently from each other (IP networks and optical network run their own set of routing and signaling protocols, respectively). A third and less common model, the augmented model, is where both IP and optical networks have their own routing instances, but reachability information of distributed IP networks is passed by the optical networks.

The overlay model describes a User-to-Network Interface (UNI), which is used for the client (for example, a router) to communicate with the server layer (the optical domain). Two standards bodies have worked to define this interface - the Optical Internetworking Forum UNI (OIF) (OIF UNI) and the Internet Engineering Task Force (IETF) (GMPLS UNI). The following demonstrations focused on the latter.

GMPLS over OTN. In their demonstration, Huawei leveraged the overlay model. In Huawei's network, the server layer network was built using three OSN6800 Optical Cross Connect (OXC) devices physically connected with 10 Gigabit OTU2 Interfaces. Across the three OSN6800 devices Automatic Switching of Optical Networks (ASON) protocols were configured, including OSPF-TE, RSVP-TE, and Link Management Protocol (LMP). The Client layer consisted of two Huawei NE5000E routers, one on each side of the optical domain, each connected via 10 Gigabit Ethernet links. Separately, Gigabit Ethernet links were run between all devices, including at the GMPLS UNI, for out-of-band signaling.

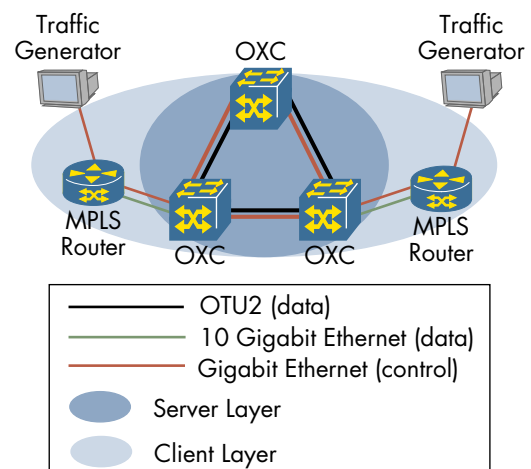


Figure 4: Packet+Optical Arch. (overlay)

In this model, requests for Label Switched Path (LSP) establishment are initiated at the client layer

(Huawei NE5000E). RSVP-TE PATH messages are sent at the GMPLS UNI through the server layer, destined to the far-end client layer node. Since out-of-band signaling was used, static LMP was configured to correlate physical links or "TE links" for the LSP (the data plane).

In addition to verifying that the client network could establish an LSP using GMPLS UNI interfaces, we also verified that it could successfully delete LSPs. LSPs were deleted by both client endpoints by using a graceful method (where the configuration remains but the LSP is disable), and in addition, the initiating client node also deleted an LSP using a more forceful method (configuration is removed completely throughout the network). In one setup, two Packet over SONET (POS) interfaces (one 10 Gig and one 40 Gig, each of which served a separate LSP) were used between the NE5000E routers and their corresponding OSN6800s in order to demonstrate an asymmetrical load balancing. Finally, Huawei demonstrated an offloading methodology which is can relieve stress from a core node by signaling a new path in the optical network. In the demonstration we also verified that only a load consisting of high priority traffic would cause a new service to be signaled.

Protection. Automatically Switched Optical Network (ASON) architecture is standardized by the ITU-T in recommendation G.8080. ASON aims to provide OTN with an intelligent control plane for dynamic network provisioning, along with mechanisms for protection and restoration. The standard specifies the architecture and requirements for the automatic switching of the transport network. It also defines a connection management function as well as a protection and restoration technique, thus providing network survivability of IP networks across optical networks.

For their demonstration Huawei configured a 1+1 protection scheme within the server (OTN) network. We verified this by first sending bi-directional Ethernet traffic between the two traffic generators as shown in the Figure. We then physically removed a fiber link between the two OSN 6800 switches which served the primary LSP, and measured frame loss. Given the number of lost frames, we were able to calculate the time it took to switch to the backup LSP (using the third OSN 6800) in milliseconds. Huawei was able to demonstrate sub-50 millisecond failure in the scenario.

GMPLS over MPLS-TP. The second vendor which participated in this test section, Alcatel-Lucent, also demonstrated their GMPLS implementation over their MPLS-TP based transport network equipment - one 1850 TSS-160 and two 1850 TSS-320s. In the demonstration, Alcatel-Lucent was able to leverage out of band GMPLS control plane, using the peer-model with static LMP on data links, to signal MPLS-TP LSPs across their equipment using OSPF-TE and RSVP-TE. To demonstrate the flexibility, control channels were not only out-of-fiber out-of-band with respect to data channels, but also data plane topology and control plane topology were not

aligned. Then a bidirectional LSP was setup and a static Pseudowire configured on its end nodes.

Additionally, Alcatel-Lucent used a section of this out-of-band scenario to successfully interoperate with MPLS based equipment by incorporating in-band signaling, referred to as a border node scenario. The Alcatel-Lucent 1850 TSS-160 and 1850 TSS-320 signaled bi-directional LSP using GMPLS (out-of-band). Then, the 1850 TSS-320 established two unidirectional MPLS signaled LSP (with a static pseudowire label) with a Cisco 7604. Finally, the 1850 TSS-320 stitched the GMPLS-based and MPLS-based LSPs together to create an end-to-end service.

SYNCHRONIZATION IN THE BACKHAUL

Base stations, according to their mobile technology, require at least frequency synchronization and at most phase, frequency and time of day synchronization. Traditionally synchronization was provided by the physical transport layer like PDH and SDH. Migrating the traditional TDM-based backhaul networks to packet-based backhaul, as a way to save OPEX and increase the capacity, demands a different approach as packet networks do not, traditionally, support synchronization. Currently, the two prominent synchronization solutions are IEEE 1588-2008 and Synchronous Ethernet. Both were included in our interoperability test.

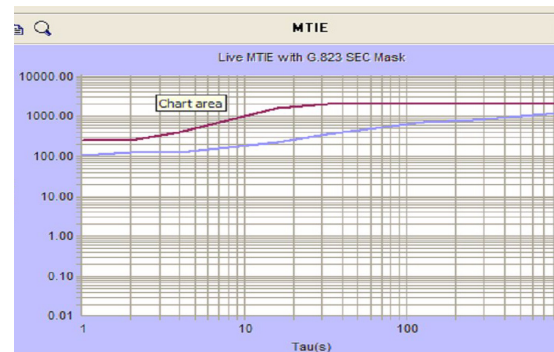


Figure 5: MTIE for End-to-end Scenario

IEEE 1588-2008

The IEEE 1588 standard is known as "Precision Clock Synchronization Protocol for Networked Measurement and Control Systems" or "Precision Time Protocol" (PTP). PTP is intended to be used to synchronize the real-time clocks of network nodes through a packet based network. Initially the standard was developed for highly accurate measurement tasks such that are used in industrial Ethernet. In 2008 the standard has been updated and its use in mobile backhaul is being explored by providers. IEEE 1588-2008 is able to provide base stations (used as a general term here to include NodeBs and eNodeBs) with three requirements for advanced mobile services: phase, frequency and Time of Day (ToD) synchronization.

In our tests we used a high precision reference clock that was synchronized with a GPS signal. The reference clock was then attached to the clock master nodes. The slaves used IEEE 1588-2008 in order to synchronize their clocks while we applied impairments using the Calnex Paragon between the master slave nodes. To make the results comparable and reproducible we configured impairment parameters matching the network profile as described in G.8261, section VI.5.2.2, test case 12.

We measured the Maximum Time Interval Error (MTIE) in order to evaluate the accuracy of the clock as its received on the 1588-2008 clients. The signal was compared to the reference clock. This process allowed us to evaluate if the vendor's implementation met the synchronization accuracy requirements defined by ITU-T G.823 or G.824 (depending on interface type).

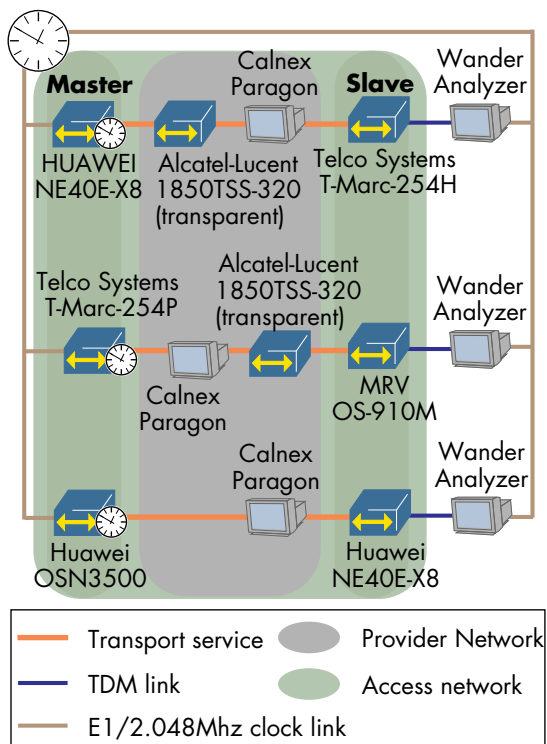


Figure 6: IEEE 1588

One of the first results in this test area was the realization that vendors' implementations differed slightly from each other which lead to the case that not all vendors could test interoperability. We have seen this in previous events and were not very surprised. Some vendors only supported PTP message exchange via Ethernet where other vendors only supported UDP or plain IP as the PTP transport mechanism. Not all vendors were able to send 1588-2008 messages via multicast and some vendors only supported unicast. Additional factors were encountered when vendors configured their subdomain ranges. Some were only able to use a 1588 subdomain fixed set to 10, where other vendors could only use subdomains "0" and "1".

The last interoperability issue we found related to message transmission intervals which are specified

in section 7.7.2.1 of the IEEE 1588-2008 specification. Some vendors sent messages with a slightly different interval than described in the above mentioned specification. The critical point we observed here was that the client announced that it has been locked to the sync signal where as the Wander measurement showed a huge drift. This issue was resolved during the hot staging through new code patches in the system.

IEEE 1588-2008 Functions	Supported by Some?	Supported by All?
Multicast Destination MAC Address	Yes	No
Unicast Destination MAC Address	Yes	Yes
UDP over IP	Yes	No
Plain IP	Yes	No
Ethernet	Yes	No

Despite the issues catalogued above we successfully tested interworking of the IEEE 1588-2008 clock synchronization protocol between Huawei NE40E-X8 (Grand Master) and the Telco Systems T-Marc-254H (Slave) with the Alcatel-Lucent 1850 TSS-320 acting as a transparent clock. The second combination that we could verify was between Telco Systems T-Marc-254P (Grand Master) and MRV OS-910M (Slave). In addition Huawei demonstrated 1588-2008 based clock synchronization between the Huawei OSN35000 (Grand Master) and the Huawei NE40E-X8 (Slave). We measured the wander and verified that the measured Maximum Time Interval Error (MTIE) met the synchronization accuracy requirements defined by ITU-T G.823 respectively G.824. The figure shows all successful results for IEEE 1588 tests.

Clock Source

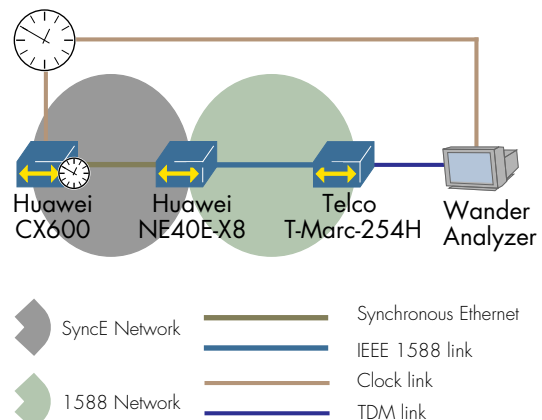


Figure 7: SyncE and IEEE 1588 Combined

Some networks are attempting to use a combination of 1588-2008 in the core and Synchronous Ethernet in the aggregation to the Base Stations. This could

simplify the delivery of frequency information to the Base Stations. We successfully tested a combination of SyncE and IEEE 1588-2008 between the Huawei CX600-X3, Huawei NE40E-X8, and Telco Systems T-Marc-254H. The NE40E-X8 received its clock signal from the GPS-based reference clock and distributed the clock through SyncE to the CX600-X3. The T-Marc-254H synchronized its clock through IEEE 1588 from the CX600-X3. We verified that the measured Maximum Time Interval Error (MTIE) meets the synchronization accuracy requirements defined by ITU-T G.823 respectively G.824.

Synchronous Ethernet

Synchronous Ethernet enables clock synchronization in packet based networks. While IEEE 1588-2008 performs clock synchronization using packets or frames, exchanging specific messages with the clock source, Synchronous Ethernet (SyncE) uses the physical layer to recover clock frequency from an external high precision clock signal.

Even though the clock synchronization mechanism differs in Synchronous Ethernet from the mechanisms used in IEEE 1588-2008, the procedure to verifying the correct and interoperable implementation between two nodes is the same. We again measured the accuracy of the clock frequency at one Synchronous Ethernet link connected to the slave node using the Calnex Paragon as wander analyzer.

The following figure shows all successfully tested combination of vendor nodes which were able to synchronize their clock through a Synchronous Ethernet link where the measured Maximum Time Interval Error (MTIE) met the synchronization accuracy requirements defined by ITU-T G.823.

ESMC Interoperability

Synchronous Ethernet networks topologies could be complex. In order to assist slave nodes in selecting the clock source, and to provide clock source redundancy, Synchronization Status Messages (SSM) could be used to exchange clocking status information. SSM messages are not new, they were already used in SDH networks, and represent the quality level of the system clocks located in various network elements.

We verified, using several failover scenarios, that the nodes were sending in their SSM messages the correct quality level. In order to evaluate the correct implementation we used 3 scenarios, one with reference clock connected to the Primary Reference Clock (PRC) node, one with reference clock connected to both the PRC and the SSU and another scenario in which we connected the reference clock only to the SSU. All successfully tested devices were able to signal their quality level to the network and changed it accordingly if the clock source has been changed.

To verify the quality level within the SSM messages

we placed two different in-line monitoring/capturing tools between the nodes. In some scenarios we used the Calnex Paragon's ESMC in-line monitor and in other scenarios we used Spirent GEM. Both were used to verify the SSM messages quality level as a real-time diagram. The following devices successfully passed all our ESMC test procedures: Alcatel-Lucent 1850 TSS-320, Alcatel-Lucent 1850 TSS-5, Cisco ASR 9000, Huawei OSN 3500, Ixia IxN2X, MRV 904-MBH, Orckit-Corrigent CM-4140, Orckit-Corrigent CM-4206, RAD ACE-3220, Telco Systems T5C-XG-ES, ZTE ZXCTN 6300, ZTE ZXCTN 6100.

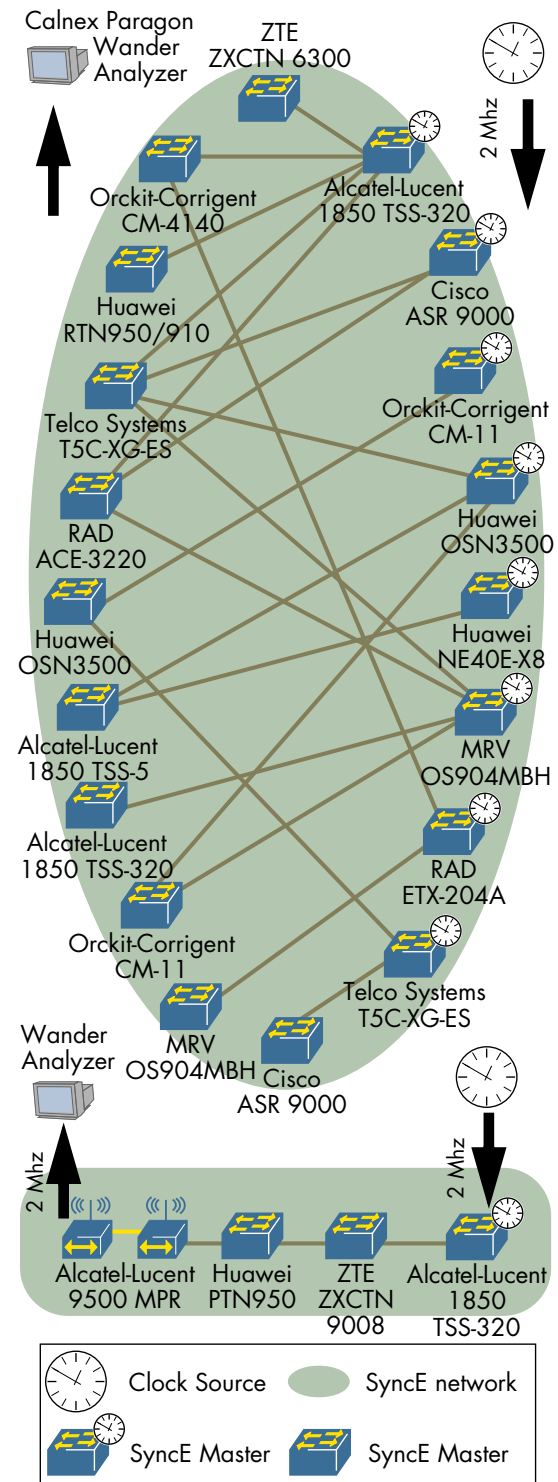


Figure 8: Synchronous Ethernet

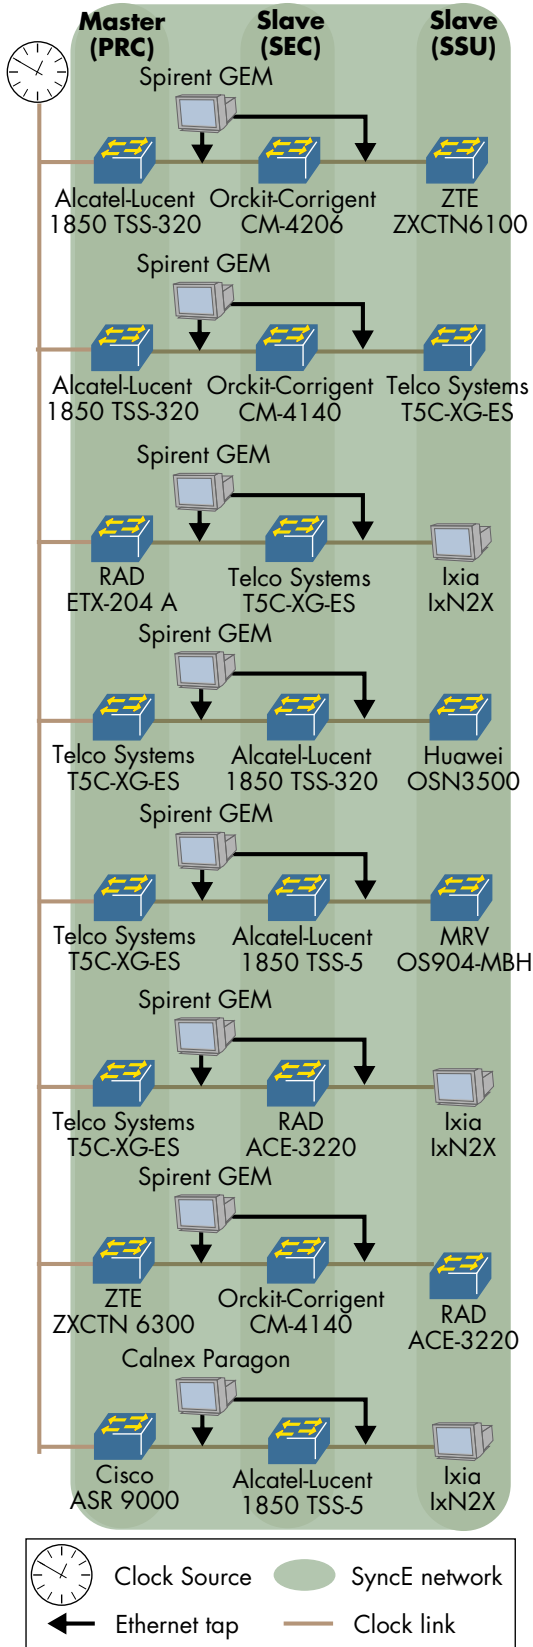


Figure 9: Ethernet Synchronization Messaging Channel

LEGACY MOBILE BACKHAUL TRANSPORT

Today most of the existing Mobile Backhaul infrastructure is based on TDM or ATM technology.

While migrating their infrastructure to next generation packet based networks mobile operators need to maintain their existing legacy infrastructure and need reliable solutions to transport TDM and ATM services through their packet based mobile backhaul network. In order to test this, we tested interoperability for structure agnostic pseudowire emulation (RFC 4353 (SAToP) and MEF8), structure aware pseudowire emulation (RFC 5086), and ATM pseudowire transport over MPLS (RFC 4717).

In order to verify running TDM services we sent a bidirectional E1 traffic pattern through the pseudowire and observed if any bit error occurred. The following devices successfully passed our SAToP test: Alcatel-Lucent 9500 MPR, Cisco 7606, Orckit-Corrigent CM-11, Orckit-Corrigent CM-4140, Huawei DSLAM 5600T, Huawei NE40E-X8, Huawei RTN950/910, MRV OS-910M, RAD ACE-3220 (through the RAD ASMi-54), Telco Systems T-Metro-200. Additionally, we verified the signaling of faults detected at the pseudowire egress to the pseudowire ingress for the following devices: Alcatel-Lucent 9500 MPR, Cisco 7606, Huawei RTN950/910, MRV OS-910M, Orckit-Corrigent CM-11, Orckit-Corrigent CM-4140. Some vendors also supported structure aware TDM pseudowire emulation (CESoPSN), which was successfully tested by the Cisco 7606, Huawei NE40E-X8, and RAD ACE-3220.

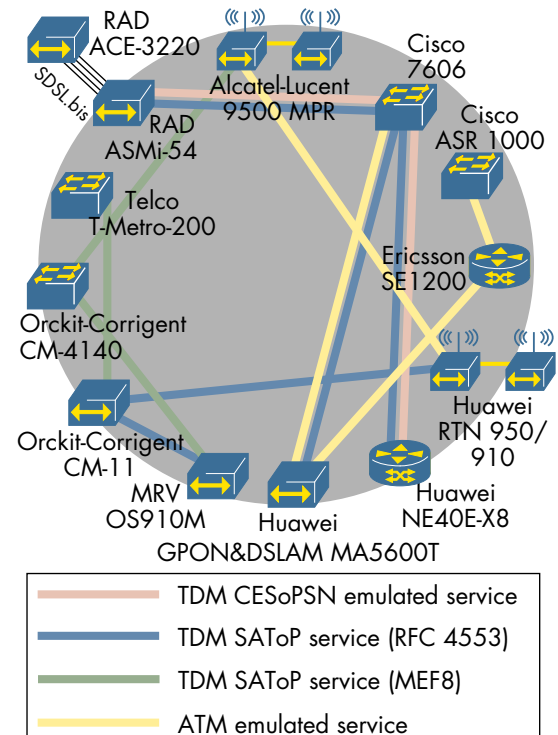
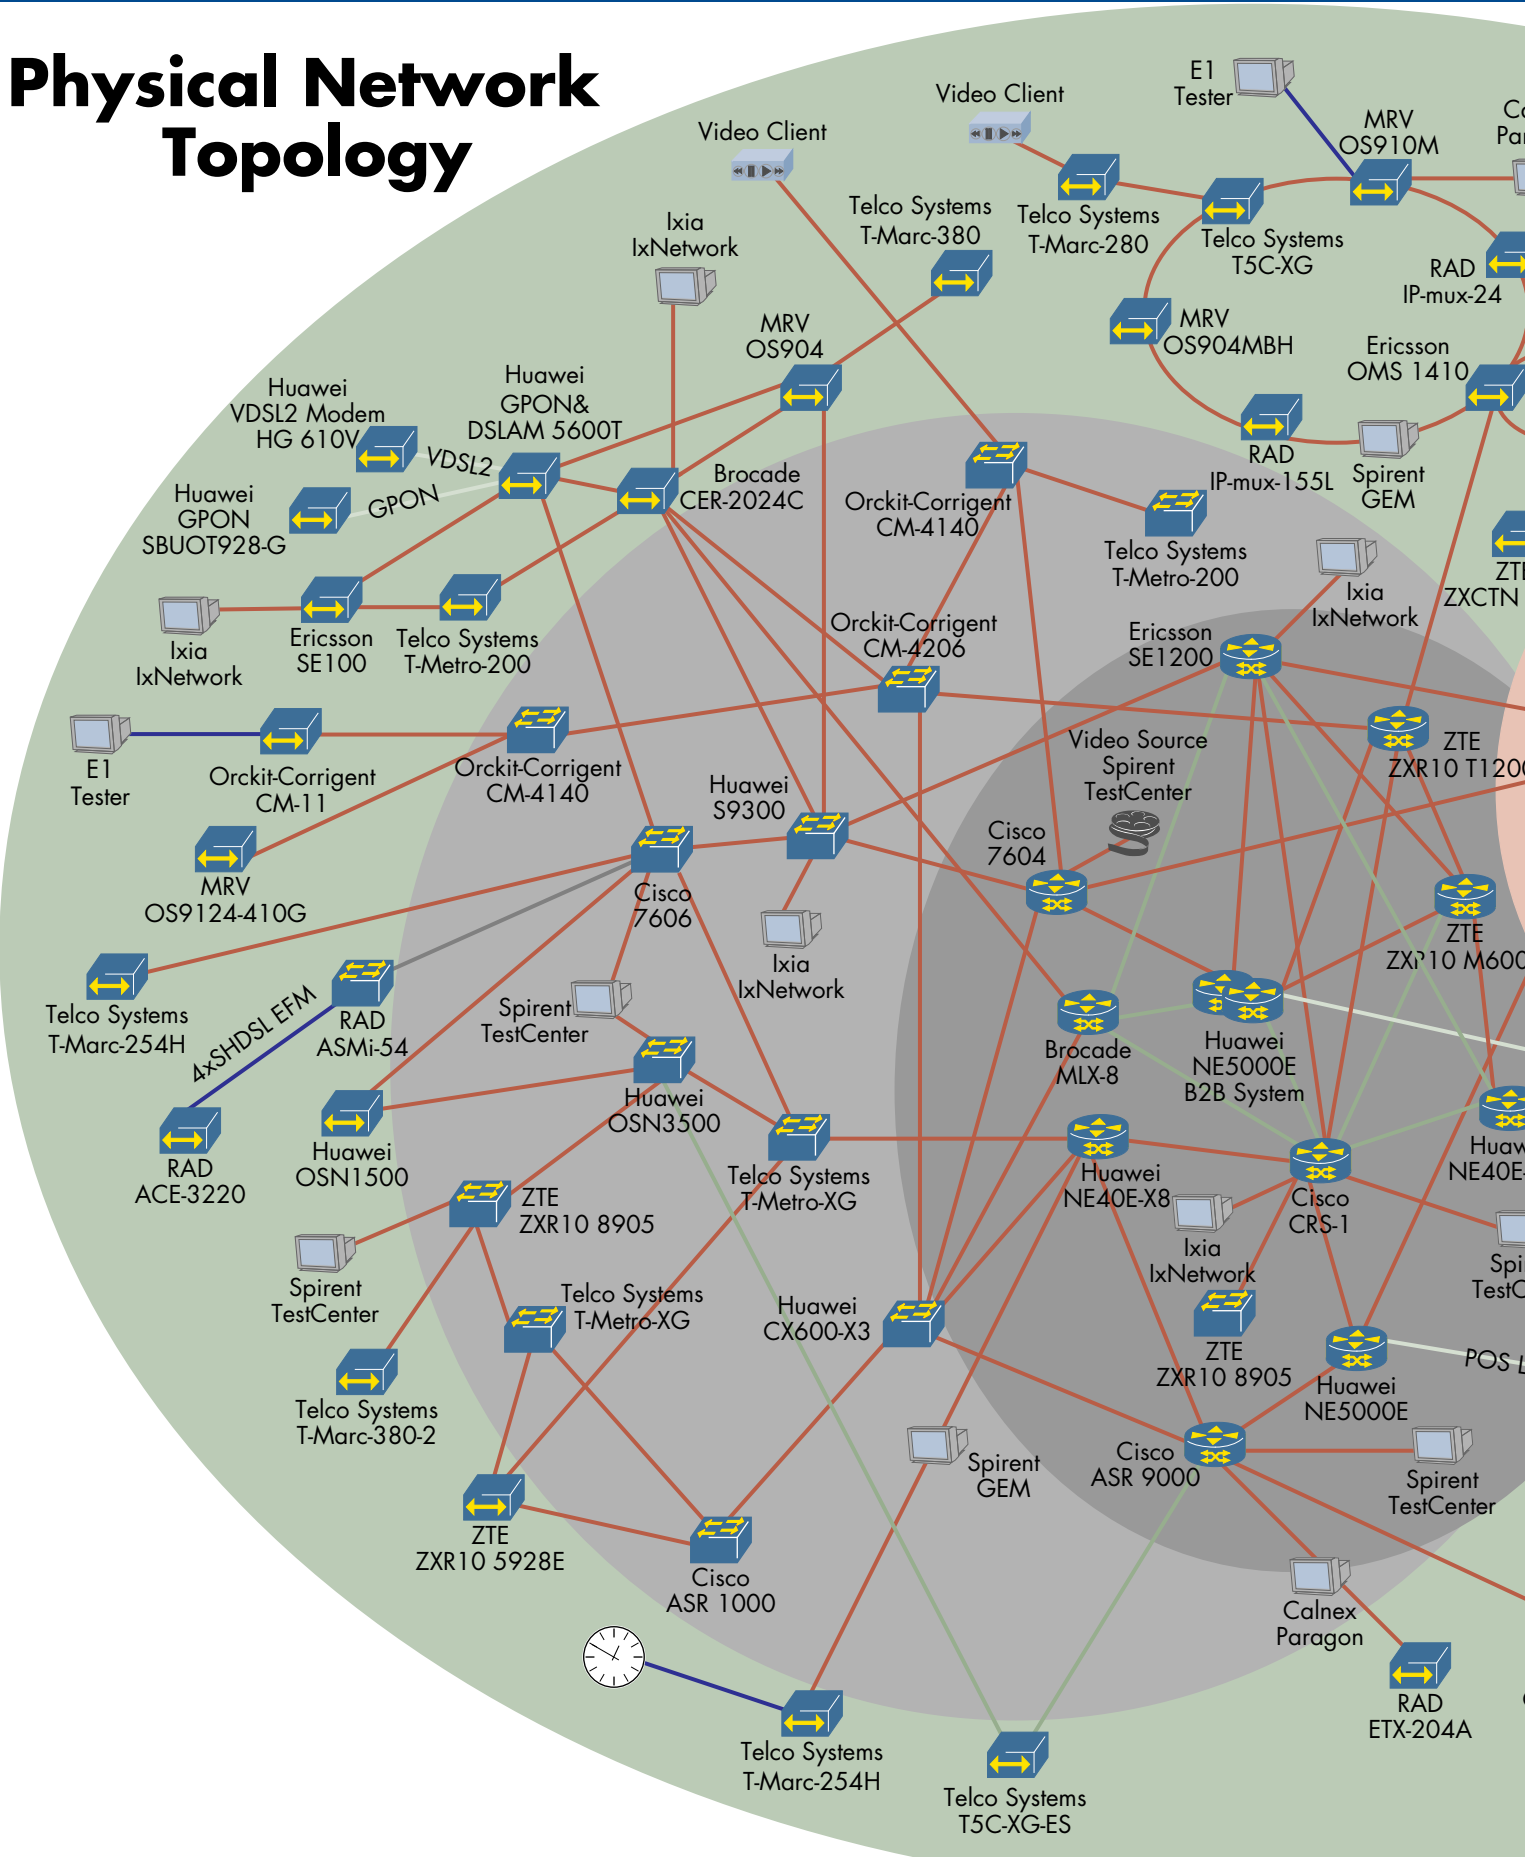



Figure 10: TDM and ATM Transport

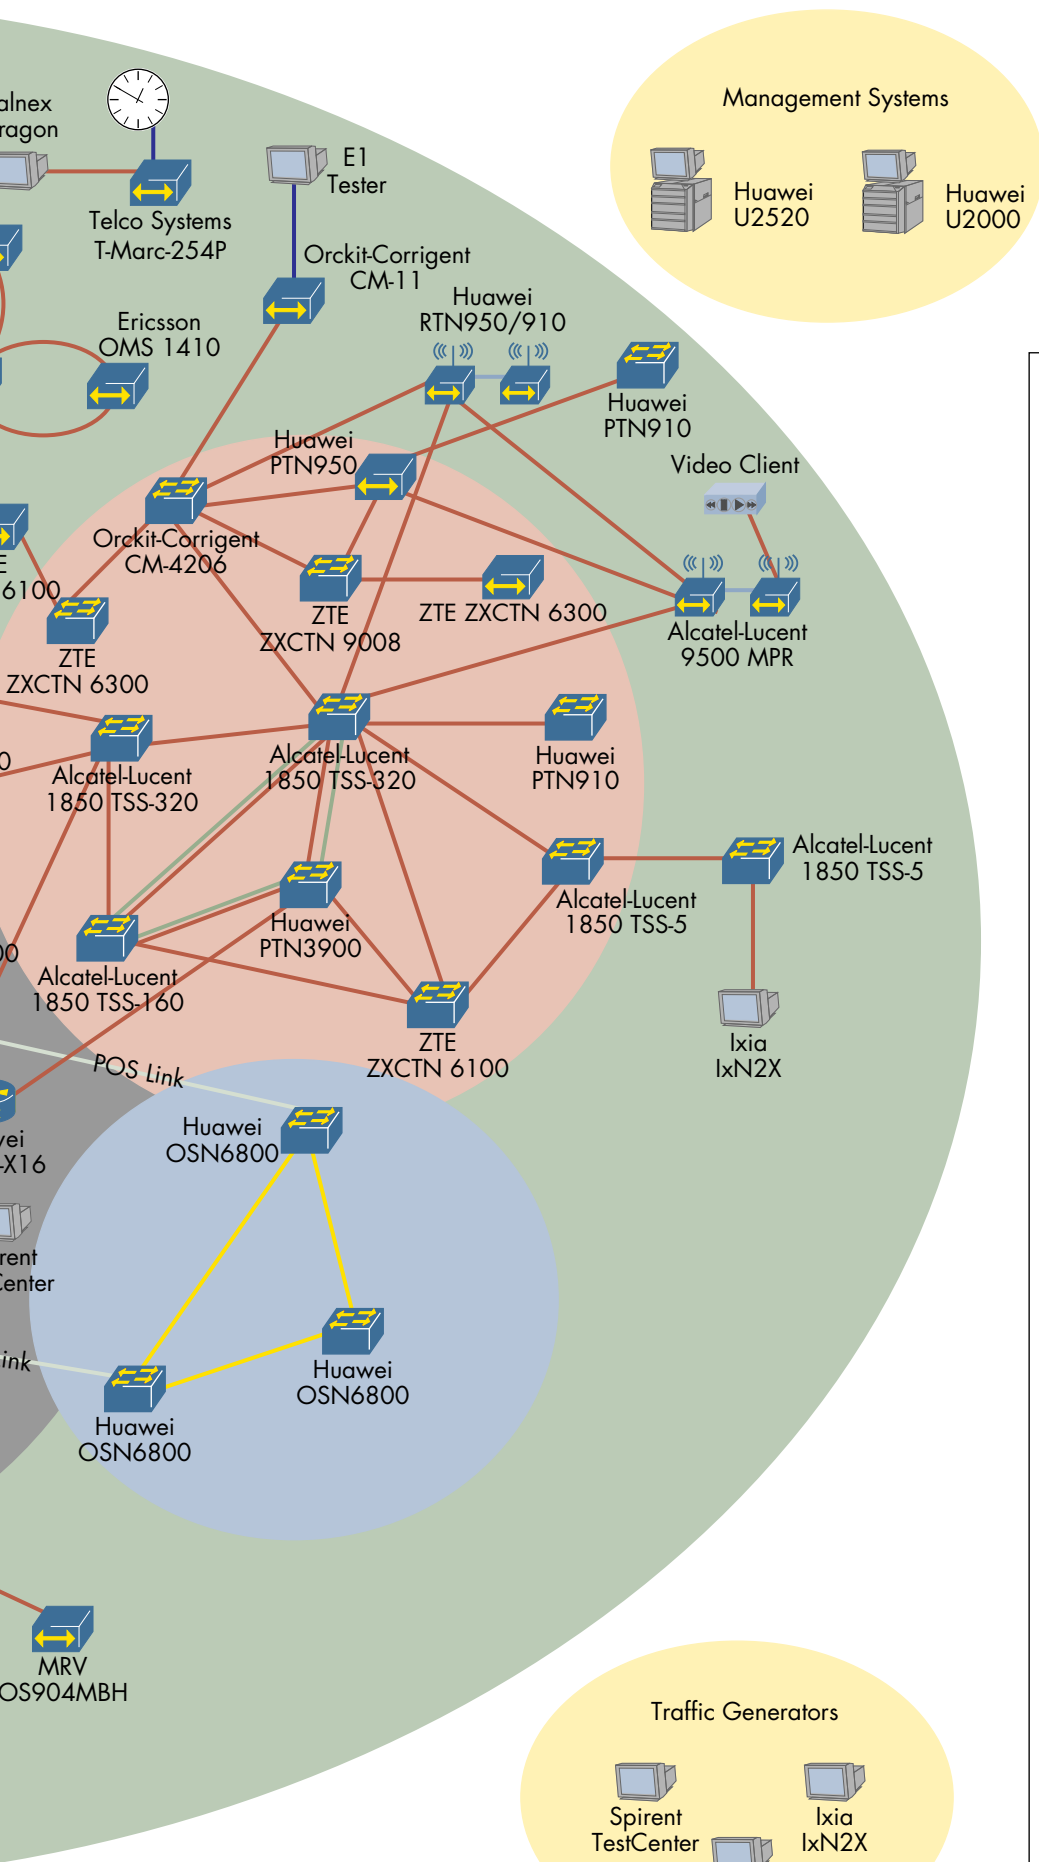
In addition to the PDH tests, we also tested ATM traffic through MPLS based networks, which is particularly relevant for 3G mobile networks. We successfully tested ATM pseudowire emulation based on IETF RFC 4717 for the following devices: Alcatel-Lucent 9500 MPR, Cisco 7606, Cisco ASR 1000, Ericsson SE1200, Huawei GPON&DSLAM MA5600T, Huawei RTN 950/910.

Physical Network Topology



10





Device Types

- Metro/Core Network Device
- Aggregation Device
- Access Device
- Clock Source
- Tester
- Management System
- Microwave Device

Connection Types

- TDM Link
- Clock Link
- ATM Link
- OTU 2 Link (10 Gbit/s)
- 10 Gigabit Ethernet Link
- Gigabit Ethernet Link
- Fast Ethernet Link
- Radio Link

Network Areas

- Access Network
- MPLS Aggregation
- MPLS Core Network
- Optical Transport Network
- MPLS-TP Network

Application Demonstrations

- Video Client
- Video Source



MPLS AND ETHERNET TRANSPORT AND ACCESS

With the growth and further development of MPLS and carrier-grade Ethernet, it is apparent that there is still much to be tested, and much to still be matured. In this section we look into the more recent and more relevant test topics based on the opinion of our service provider panel, the participating vendors, and ourselves.

IPv6 BGP/MPLS VPNs

As mentioned earlier in the multicast VPN section, MPLS-based Layer 3 VPNs allow providers to offer connectivity and routing services to enterprise customers. As many enterprises started to roll their network infrastructures to IPv6 from IPv4, the IETF defined an extension to BGP/MPLS VPNs (mainly a new address family type) which let providers to continue the same VPN services to IPv6 based customers while keeping the IPv4 carrier infrastructure already in place.

In each test, we emulated two distinct enterprise customers by configuring IPv6 interfaces on test equipment, which was then connected to the Provider Edge (PE) routers under test which thus had to establish two separate VPN services. PEs either peered with the emulated customer using eBGP (shown in diagram) or simply redistributed the directly connected subnet into the customer's routing table.

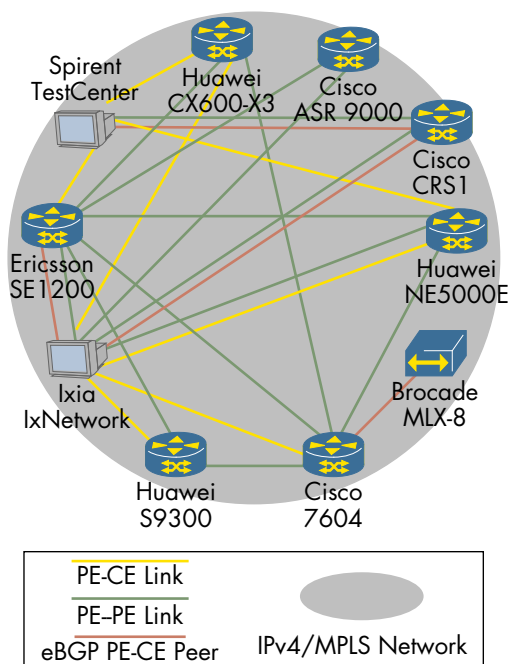


Figure 11: IPv6 BGP/MPLS VPNs

We verified on each PE device that the IPv6 routes were learned from the neighboring PE device, and were present in the Virtual Routing and Forwarding instance (VRF) associated with that customer. Bidirectional

traffic streams were transmitted on each VPN by either a Spirent TestCenter or Ixia LxNetwork, and we observed that all packets traversed the IPv4/MPLS network and arrived only on the designated VPN (no dropped frames, no cross-talk).

The following devices were configured to serve as PE routers: Cisco 7604, Cisco ASR 9000, Cisco CRS-1, Ericsson SE1200, Huawei CX600-X3, Huawei NE5000E, Huawei S9300, Ixia LxNetwork, and Spirent TestCenter. The Brocade MLX-8 functioned as a CE device in one of the tests, and all other CE devices were emulated by either an Ixia LxNetwork or Spirent TestCenter.

Ethernet Ring Protection Switching (ERPS - G.8032)

End-to-end availability is one of the most vital yet challenging aspects of offering a service. Providers are quite assured by the resiliency of their SDH/SONET networks, and now also by their Fast Reroute enabled MPLS networks, but what about their native Ethernet networks? Rapid Spanning Tree has improved upon Spanning Tree, but this still does not fit the requirements of many networks. For this reason, the ITU-T has defined Ethernet Ring Protection Switching (ERPS) in G.8032.

Two methodologies were used to break the ring's connectivity - a physical link break, and an explicit impairment of Continuity Check Messages (CCMs) in a single direction.

Several multi-vendor rings were tested. One ring was comprised of an Ericsson OMS1410, an MRV OS910M, and a RAD IPmux-155L. For the second ring, we substituted a RAD IPmux-24 where the IPmux-155L had been. Unidirectional traffic flows were transmitted between the Ring Protection Link (RPL) owner and the non-RPL node by a Spirent TestCenter or Ixia LxNetwork. Physical link down caused the failover in both cases. Failover times ranged from 20 to 138 milliseconds, and revertive switchover times from 4 to 105 milliseconds.

The third ring was made up of an MRV OS910M and OS904-MBH, as well as a Telco Systems T5C-XG. Unidirectional traffic flows were transmitted between the RPL (non-owner) node and the non-RPL node by the test equipment. The CCM rate was set at 1 per second, and the CCMs were blocked between the RPL (non-owner) node and the non-RPL node in one direction per test by either a Spirent GEM or Calnex Paragon. Both the MRV OS910M and the Telco T5C-XG were the RPL owner in varying tests. Failover times ranged from 5 to 55 milliseconds, and revertive switchover times from 12 to 64 milliseconds.

In addition, Ericsson demonstrated an ERPS version 2 ring between two OMS1410 nodes. Manual switchover times were measured (per flow) at 51 and 79 milliseconds, forced block switchover at 68 and 85 milliseconds, and link failure switchover at 74 and 86 milliseconds. We verified that enabling non-revertive failover caused the RPL owner to not switch traffic back to the original direction until non-

revertive mode was cleared, at which point switchover times were measured at 49 and 61 ms.




Figure 12: Logical Diagram of ERPS Results

Finally, a four-node ring was built, consisting of an Ericsson OMS1410, an MRV OS910M and OS904-MBH, and a Telco Systems T5C-XG. The section between MRV devices was based on CCM exchange. The Ericsson OMS1410 was the RPL owner. Traffic was supplied to the Ericsson and Telco Systems devices and the non-RPL link was administratively disabled on the Ericsson. Failover times were measured at 70 and 44 milliseconds, and revertive switchover times at 54 and 17 milliseconds.

VPLS Enhancements

It need not be debated the usefulness of a service which provides a service where customers miles away are seemingly connected via a single Layer 2 switch - thus is Virtual Private LAN Service (VPLS). Nevertheless, there has been some room for improvement, namely in two specific areas - ease of configuration, and scalability.

Auto-Discovery. If an MPLS network already has

iBGP running between peers, why not take advantage of this for multiple services? This was inherent to establishing Layer 3 VPNs as BGP is required, however not every VPLS network requires BGP (two types of VPLS have been defined - one using LDP for signaling, and one using BGP). Nevertheless, since BGP is common in many networks providing VPLS services, an extension was defined to use these BGP sessions to automatically discover PEs in your VPLS domain, thus removing the requirement to explicitly configure each peering. This feature was successfully tested and verified between the Cisco 7604, Cisco ASR 9000, and Ixia IxNetwork.

PBB Interworking. One aspect of a VPLS service is to manage (learn) the MAC addresses in each customer site, so the network knows which site to forward any given incoming Ethernet frame to. The problem with this is the sheer load of MAC addresses within a given site, all of which must be stored in the memory of the attached PE router on the provider side. One way to manage this is to leverage the use of Provider Backbone Bridges, which abstract the MAC addresses learned, so long as the PE supports the PBB interface (802.1ah) and its respective extension to VPLS.

Cisco demonstrated this feature by first establishing a VPLS domain between a 7606 and 7604 router. The 7606 terminated the VPLS into an 802.1ah interface towards an Ixia IxNetwork traffic generator, as well as an 802.1ah interface towards a Cisco ASR 9000. The ASR 9000 played the role of the Backbone Edge Bridge (BEB), accepting 802.1ad frames from a second Ixia interface, and sending PBB frames towards the 7606. On the side of the 7604, a PBB instance was created within the router, thus terminating the PBB area and configuring a 802.1ad interface towards a third Ixia interface. Ixia then exchanged traffic amongst the three interfaces, using their respective encapsulation, and witnessed zero frame loss.

MPLS-TP Resiliency via OAM

As the IETF continues to discuss and define the characteristics of the MPLS Transport Profile (MPLS-TP), we continue to do our part by holding a magnifying glass over the current implementations and solutions from the vendors that bring them to our interop events. Being a technology intended for transport networks, we have continued to test their more vital aspects, namely OAM and resiliency. While the IETF are discussing several options for MPLS-TP OAM tools (ITU-T Recommendation Y.1731, LSP Ping, MPLS BFD, PW VCCV) we tested the approach supported by Alcatel-Lucent, Huawei, and ZTE based on Y.1731.

The Alcatel-Lucent 1850 TSS-320, Huawei PTN 3900, and ZTE ZXCTN 6100 all successfully tested interoperability for their MPLS-TP OAM implementations, based on ITU-T Y.1731 which defines a protocol for Ethernet OAM. We tested this by first verifying that the protocol was used by all vendors to

establish connectivity on the respective MPLS-TP transport path, and their ability to switch over to a backup transport path upon loss of such connectivity. Between those devices was a LAN segment, to make sure that the trigger was the loss of CCM frames (not the LOS).




Figure 13: MPLS-TP Protection

APS administrative commands including “Manual”, “Force”, “Clear” and “Lockout” were also tested for 1:1 protection. Three scenarios were successfully tested, resulting in a full mesh of test pairs between the Alcatel-Lucent 1850 TSS-320, Huawei PTN 3900, and ZTE ZXCTN 6100. In each test traffic switched to the backup, and reverted, respectively. Finally, the Alcatel-Lucent 1850 TSS-320, Huawei PTN 3900 and ZTE ZXCTN 6100 tested that in case of a mis-connection, emulated by purposefully mis-configuring OAM fields, caused traffic to switch over to the backup path.

Additionally, the following devices participated in providing services in the MPLS-TP domain by statically assigning MPLS labels to establish a path: Alcatel-Lucent 1850 TSS-5, Alcatel-Lucent 1850 TSS-160, Alcatel-Lucent 1850 TSS-320, Huawei PTN910/PTN950, Huawei PTN3900, Orckit-Corrigent CM-4206, ZTE ZXCTN 6100, ZTE ZXCTN 6300, and ZTE ZXCTN 9008.

MPLS to the Access

The baseline for providing Ethernet services over an MPLS network is the establishment of Ethernet pseudowires (PW) or Pseudo Wire Emulation Edge to Edge (PWE3). A majority of the participating equipment supports this, and a majority of this equipment has been tested for this feature at our interop events dating back several years. Nevertheless, extending MPLS to the access to provide such services is becoming a topic of interest, and also potentially requires new implementations to be tested.

This test is intended for equipment which implements MPLS but is classified as an Access Device. We archived a group of test pairs which supported diverse E-Line capabilities such like Ethernet and Ethernet tagged mode defined in RFC4446 and 4448, and based on the Label Switched Paths (LSP) created by both RSVP-TE and LDP signaling protocols, where the latter one also used for the label distribution of the pseudowires.

The following vendors and devices successfully participated the tests: Brocade CER-2024C, Cisco ASR 1000, Ericsson SE100, Huawei GPON&DSLAM MA5600T, Ixia IxNetwork, MRV OS904, Spirent TestCenter, Telco Systems T-Metro-200, ZTE ZXR10 5928E, ZTE ZXR10 8905 and ZTE ZXR10 M6000.

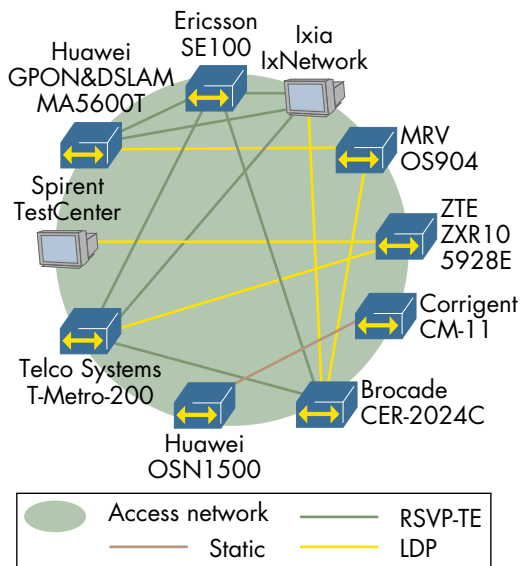


Figure 14: MPLS to Access Equipment

We found that while some implementations may be new, LDP and RSVP-TE are still mature protocols and issues were minimal. Most devices supported both protocols for tunneling, still, vendors demonstrated that they could run this test using either as shown in the diagram. It is important to note that many of the core and aggregation routers and switches could also perform this test however since Ethernet pseudowires is now a mature and common feature in this equipment, and since we have tested it in interop events years, we attempted to limit this test to access devices only.

PWE3 Termination to L3VPN

Layer 3 VPNs, which have been previously introduced in this report, continue to offer what enterprise customers need. However, service providers run into an issue when their customer has an office location outside of the reachability of their network. In these cases, it is often possible, and a feasible solution, to establish a Layer 2 service from a third party to gain this reachability.

In this test a PWE3 configured on Device Under Test (DUT) is used to backhaul L3VPN services to the core network, where Layer 3 services were offered. The following figure depicts the test setup in detail, where the devices making the termination are posited in the middle.

Five vendors successfully participated the test using the following devices: Cisco 7604, Cisco CRS-1, Ericsson SE1200, Huawei NE40E-X8, Huawei S9300, ZTE ZXR10 M6000, ZTE ZXR10 T1200. We were happy to find minimal issues in this test,

proving this to be a reasonable architecture for providers who may need to deploy it.

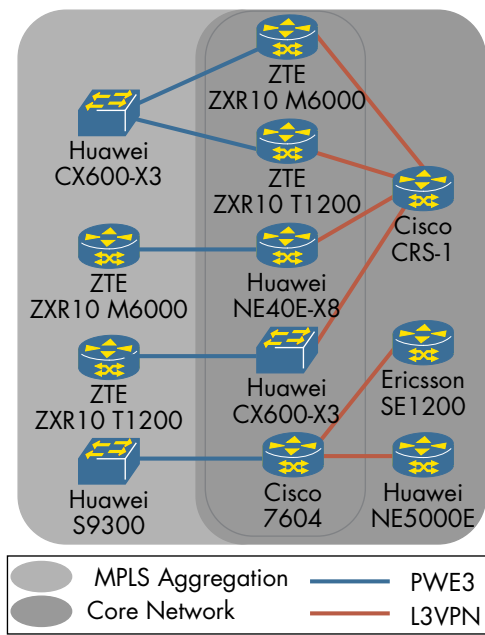


Figure 15: PWE3 Termination to L3VPN

Pseudowire Resiliency

While Fast Reroute is in many cases a great solution for adding resiliency to the transport layer of an MPLS network, service-aware resiliency is still to be desired. Vendors showed that they were already prepared to start testing redundant pseudowires (PWs) based on the “Muley” drafts (draft-ietf-pwe3-redundancy-02, and draft-ietf-pwe3-redundancy-bit-02), which define a mechanism to signal the standby status of a redundant PWE3 between their terminating nodes, using the Preferential Forwarding Status Bit.

In each test an upstream Provider Edge (PE) node had two pseudowires configured, each terminating on a separate downstream PE, and the PEs agree which of the two PWs is primary. When the access circuit connected to the active downstream PE went down (when Ethernet link to the customer went down) the downstream PE was expected to signal to the upstream PE to switchover to the secondary PW. Since this meant there were two access circuits connected to the downstream customer site, a Telco Systems T5C-XG (a layer 2 switch) was used to converge the two, as shown in the diagram. A tester sent traffic between the upstream PE and the T5C-XG to measure fail-over times.

The following devices supported the preferential forwarding status bit, and performed the test in multiple configurations (as shown in the diagram): Cisco 7604, Cisco ASR 1000, Ericsson SE 1200, Huawei CX600-X3, Huawei S9300, and Telco Systems T-Metro-XG. We found that failover times were not consistent amongst the vendors, ranging from 3 ms to 525 ms. We did however expect some length and inconsistency to the numbers to begin with, given the complexity of the MPLS environment

the devices were testing in, as well as the additional factor of the layer 2 switch. Additionally, we encountered some time consuming issues regarding OSPF-TE configuration and interoperability, but these were eventually solved for the most part. Ultimately we were happy to see that interoperability amongst the many implementations that were tested given the youth of this feature.

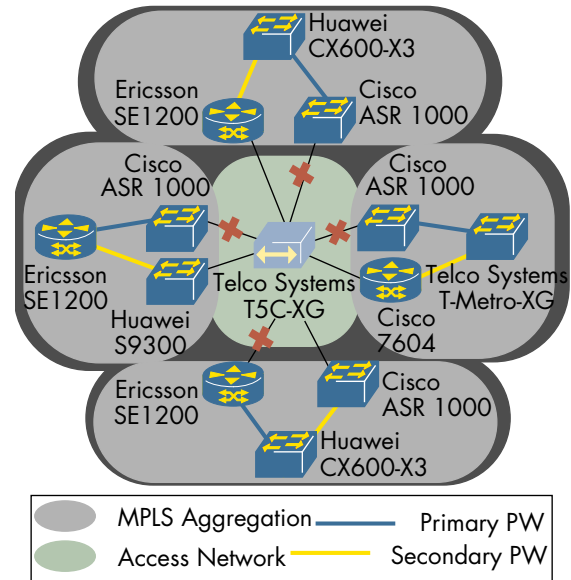


Figure 16: Pseudowire Resiliency

On a separate note, Cisco demonstrated OAM over multi-segmented pseudowires with their 7604, 7606, and ASR 1000 routers. The 7606 established a pseudowire with each of the other two routers, both of which had Virtual Circuit Connectivity Verification (VCCV) enabled. Through support of Control Channel (CC) type 3, each end router was able to perform LSP pings and traceroutes across the entire service, including each of the two segments which were created.

Management

Each interoperability event we host, we call out to the industry for interoperable management solutions. The request is relatively flexible, potentially including anything from systems which support SNMP, to systems which have built in plug-ins for multiple vendors. We look for features such as SLA monitoring, provisioning, and network monitoring.

Three participating vendors demonstrated standardized IETF MIBs defined in RFCs 3812, 3815, 4382 with the following devices: Brocade MLX-8, Cisco CRS-1, and Huawei NE5000E. The MIBs were sent to Huawei’s cross domain Unified NMS iManager U2000. We observed that the management system was able to report alarm notifications either when an LSP or a Layer 3 VPN VRF interface went down.

Performance Monitoring (Y.1731)

We have polled our service provider panel to see which test areas are most interesting to them, and to see where interoperability is most important. A recurring result of this question, is Performance Monitoring. We continue these tests by seeing how vendors can measure loss and delay in an interoperable way - namely, by use of frames defined in ITU-T recommendation Y.1731. The standard extends IEEE 802.1ag Ethernet OAM to define specific frames for performance monitoring including Delay Measurement Messages (DMM) and Loss Measurement Messages (1DMs and LMMs). In one case, where these were not supported, Loopback Messages (LBM) were used to measure the round trip delay.

In order to verify that the devices under test (DUTs) could measure loss and delay, we had to inject loss and delay into the test setup. For this, we used either the Calnex Solutions Paragon or Spirent xGEM impairment equipment. First, we verified OAM connectivity between the DUTs without any impairment to get baseline interoperability, and a baseline measurement. Then, depending on the test, we introduced a specific type of impairment with the impairment tool.

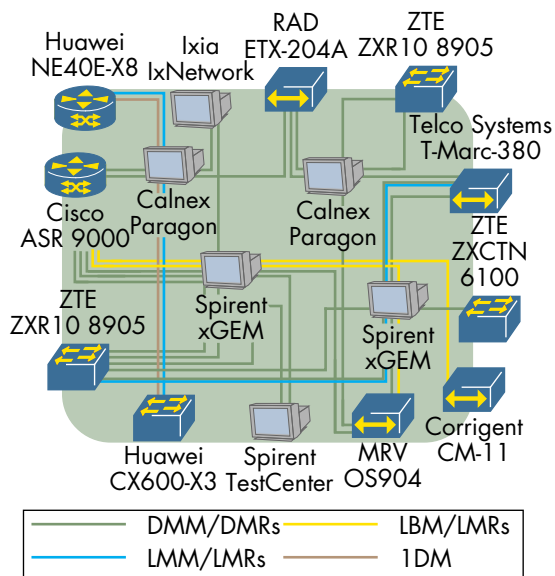


Figure 17: Performance Monitoring

For delay measurement we added a constant delay of 20 milliseconds (ms), and expected the measurement reported from the DUTs to increase by this amount, compared to the baseline. For delay variation, we configured the impairment tool to vary the delay it introduced from 15 ms to 25 ms, on a per-packet basis. The following devices participated in the tests: Cisco ASR 9000, Ixia IxNetwork, MRV OS904, RAD ETX-204A, Telco Systems T-Marc-380, Spirent TestCenter and ZTE ZXR10 8905. The majority of tests did not have issues, and reported relatively accurate values. In the case of the test between Cisco ASR 9000 and Corrigent CM-11 LBMs and LBRs were used for delay measurements.

While interoperability issues were quite minimal, only a rough majority of the devices under test matched each others measurements wive-2—l of accuracy. Initially, the impairment tool was configured to alter all packets, regardless of if they were customer frames or OAM measurement frames. After witnessing some inconclusive results between a particular pair of implementations, we added a filter on the impairment tool to only delay OAM measurement frames (DMMs and LMMs). This helped to focus on the discrepancies in the reported values.

Finally, for frame loss, we configured a constant 10% loss. The following vendors accurately reported on this increase of loss: Telco Systems T-Marc-380 and ZTE ZXR10 8905. Additionally, ZTE demonstrated the frame loss measurement using both ZTE ZXR10 8905 and ZXCTN 6100 devices.

Using the NetNumen management system ZTE demonstrated frame delay measurement between both ZTE ZXR10 8905 and ZXCTN 6100, where the latter device responded to NetNumen with the measured delay, and NetNumen displayed the values via its graphical user interface.

Using the CX600-X3 and NE40E-X8 Huawei demonstrated dual and single frame loss, one way frame delay over a PW, demonstrating that a VC Label could carry the LMM/LMR frames. The Calnex Solutions Paragon impairment tool was used to introduce frame loss impairments. Huawei demonstrated the reporting of the frame loss measurements to their iManager U2000 Unified NMS, which first provisioned the PW, then displayed the frame loss measurement taken by the two routers.

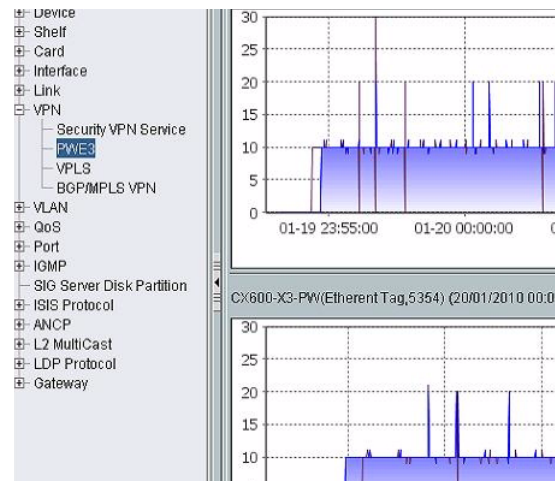


Figure 18: Huawei Unified NMS iManager U2000 - Frame loss measurement

Video Quality Monitoring

Services these days are becoming increasingly application-rich in hopes to not only increase revenue for providers, but also to cut out the number of parties a customer must interface with in order to enjoy certain products. One clear application which belongs in this category, is video, supplied one way or another (IPTV, web based video, video conferencing, etc).

This brings new challenges to providers. If they would like to offer video, it is no longer sufficient to monitor the network providing it, but it should be also possible to monitor the video itself. We asked vendors to test their equipment (be it a management system or a router itself) that can do such video monitoring.

Two vendors signed up for the test. Cisco tested with the ASR 9000 which supports in-line video monitoring. Huawei tested their NE40E-X8 which supports in-line monitoring by iVSE solution and the NEU100, an external probe, both monitored by a single management system - the U2520. The IETF has defined a proposed Media Delivery Index (MDI) in RFC 4445 which defines metrics to monitor media such as video. Both vendors implementations were based on the RFC, Cisco's monitoring at an IP/UDP level, and Huawei at MPEG level.

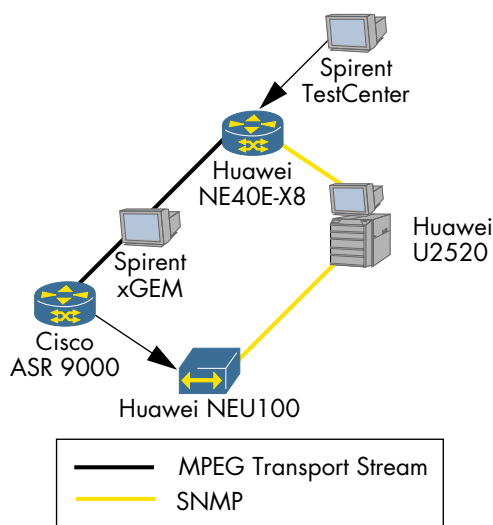


Figure 19: Video Quality Monitoring

We ran a single test which included all participating equipment, as shown in the diagram. The Spirent TestCenter was used to host an MPEG Transport Stream which was enabled for CBR. Each implementation includes a metric which measures loss (Media Loss Ratio or MLR for MDI) and delay (Delay Factor or DF for MDI) so the test had two parts. In each part we injected impairment (either delay or loss) using the Spirent xGEM impairment tool, and expected only the associated metric to be affected on the Cisco ASR 9000 and the Huawei NEU100 equipment. In each test the Huawei NE40E-X8 was expected to continue reporting the same values as it did without impairment, since it was placed before the impairment tool. The test proved successful, as both downstream devices updated their respective values as delay and loss was imposed on the video stream. On the Cisco ASR 9000, the values were shown directly on the Command Line Interface (CLI) while Huawei displayed values coming from both the NE40E-X8 and NEU100 on their management system (U2520) GUI.

Acknowledgements

We would like to thank Stephen Murphey from the University of New Hampshire InterOperability Lab (UNH-IOL) and Robert Jones for their hard work and efforts during the two week hot staging.

Editors. This report has been edited by Jambi Ganbar, Robert Jones, Jonathan Morin, Stephen Murphy, Ronsard Pene, Carsten Rossenhövel, Thomas Sladek, and Xiao Tai Yu.

ACRONYMS

Term	Definition
AIS	Alarm Indication Signal
APS	Automatic Protection Switching
ASON	Automatic Switching of Optical Networks
ATM	Asynchronous Transfer Mode
BEB	Backbone Edge Bridge
BGP	Border Gateway Protocol
CCM	Continuity Check Message
CESoPSN	Circuit Emulation Service over Packet Switched Network
CFM	Connectivity Fault Management
DF	Delay Factor
DMM	Delay Measurement Message
DSL	Digital Subscriber Line
DUT	Device Under Test
E-Line	Point-to-Point Ethernet Service similar to a leased line ATM PVC or Frame Relay DLCI
EFM	Ethernet in the First Mile
ERPS	Ethernet Ring Protection Switching
ESMC	Ethernet Synchronization Messaging Channel
GMPLS	Generalized Multiprotocol Label Switching
GPS	Global Positioning System
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IP	Internet Protocol
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
LBM	Loopback Message
LDP	Label Distribution Protocol
LMM	Loss Measurement Message
LMP	Link Management Protocol
LOS	Loss Of Signal
LSP	Label Switched Path
LSR	Label Switched Router
MAC	Media Access Control
MDI	Media Delivery Index
MLR	Media Loss Rate
MPLS	Multi-Protocol Label Switching
MPLS-TP	MPLS Transport Profile
MTIE	Maximum Time Interval Error
MVPN	Multicast VPN

Term	Definition
OAM	Operations, Administration and Maintenance
OPEX	OPerating EXpenditure
OSPF	Open Shortest Path First
OTN	Optical Transport Network
OTU 2	Optical Transport Unit 2 (10 Gbit/s)
P2MP	Point-to-Multipoint
PBB	Provider Backbone Bridge
PBB-TE	Provider Backbone Bridge Traffic Engineering
PE	Provider Edge
PHY	PHYsical layer
PIM-SM	Protocol Independent Multicast - Sparse Mode
PIM-SSM	Protocol Independent Multicast - Source Specific Multicast
POS	Packet over SONET
PRC	Primary Reference Clock
PTP	Precision Time Protocol
PW	PseudoWire
PWE3	Pseudowire Emulation Edge to Edge
RFC	Request For Comments
RPL	Ring Protection Link
RSVP-TE	Resource reSerVation Protocol Traffic Engineering
S2L sub-LSP	Source-to-leaf sub-LSP
SAToP	Structure-Agnostic Time Division Multiplexing (TDM) over Packet
SDH	Synchronous Digital Hierarchy
SEC	SDH Equipment slave Clocks
SLA	Service Level Agreement
SSM	Synchronization Status Message
SSU	Synchronization Supply Unit
SyncE	Synchronous Ethernet
TDM	Time Division Multiplexing
TE	Traffic Engineering
ToD	Time of Day
UDP	User Datagram Protocol
UNI	User-to-Network Interface
UNI	User Network Interface
VCCV	Virtual Circuit Connectivity Verification
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network

REFERENCES

- "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007
- "Signalling Requirements for Point-to-Multipoint Traffic-Engineering MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006
- "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", draft-ietf-l3vpn-2547bis-mcast-bgp-08, October 2009
- "Multicast in MPLS/BGP IP VPNs", draft-ietf-l3vpn-2547bis-mcast-08, March 2009
- "Multicast in MPLS/BGP IP VPNs", draft-rosen-vpn-mcast-12, August 2009
- "OAM Functions and Mechanisms for Ethernet Based Networks", ITU-T Y.1731, February 2008
- "Interfaces for the Optical Transport Network (OTN)", G.709/Y.1331, March 2003
- "Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extensions for Automatically Switched Optical Network (ASON)", RFC4139, July 2005
- "Architecture for the Automatic Switched Optical Network (ASON)", ITU-T G.8080, June 2006
- "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional description", RFC 3471, January 2003
- "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource Reservation Protocol - Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003
- "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4202, October 2005
- "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005
- "Precision Time Protocol (PTP)", IEEE 1588-2008
- "The Control of Jitter and Wander within Digital Networks which are Based on the 2048 kbit/s Hierarchy", ITU-T G.823, March 2000
- "The Control of Jitter and Wander within Digital Networks which are Based on the 1544 kbit/s Hierarchy", ITU-T G.824, March 2000
- "Timing and Synchronization Aspects in Packet Networks", ITU-T G.8261/Y.1361, April 2008
- "Timing characteristics of synchronous Ethernet equipment slave clock (EEC)", ITU-T G.8262/Y.1362, August 2007
- "Distribution of timing through packet networks", ITU-T G.8264/Y.1364, October 2008
- "Synchronization layer functions", ITU-T G.781, September 2008
- "Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks", RFC 4717, December 2006
- "Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)", RFC 4553, June 2006
- "Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)", RFC 5086, December 2007
- "Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks", MEF 8, October 2004
- "MPLS Generic Associated Channel", RFC 5586, June 2009
- "MPLS-TP OAM Analysis", draft-ietf-mpls-tp-oam-analysis, November 2009
- "MPLS-TP Linear Protection", draft-weingarten-mpls-tp-linear-protection, December 2009
- "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006
- "MPLS-TP OAM based on Y.1731", draft-bhh-mpls-tp-oam-y1731, July 2009
- "Multicast in VPLS", draft-ietf-l2vpn-vpls-mcast-05, July 2009
- "Protocol Independent Multicast - Sparse Mode (PIM-SM)", RFC 4601, August 2006
- "Internet Group Management Protocol, Version 3", RFC 3376, October 2002
- "Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)", RFC 3916, September 2004
- "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005
- "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006
- "Ethernet ring protection switching", ITU-T G.8032, June 2008
- "Preferential Forwarding Status bit definition", draft-ietf-pwe3-redundancy-bit-02.txt, October, 2009
- "Pseudowire (PW) Redundancy", draft-ietf-pwe3-redundancy-02, October 2009
- "Provisioning, Autodiscovery, and Signaling in L2VPNs," draft-ietf-l2vpn-signaling-08, May 2006
- "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007
- "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007
- "BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN", RFC 4659, September 2006
- "Extensions to VPLS PE model for Provider Backbone Bridging", draft-ietf-l2vpn-pbb-vpls-pe-model, January 2010
- "VPLS Interoperability with Provider Backbone Bridges", draft-ietf-l2vpn-pbb-vpls-interop, January 2010
- "A Proposed Media Delivery Index (MDI)", RFC 4445, April 2006



EANTC AG
European Advanced Networking Test Center

Einsteinufer 17
10587 Berlin, Germany
Tel: +49 30 3180595-0
Fax: +49 30 3180595-10
info@eantc.de
<http://www.eantc.com>



Upperside Conferences

54 rue du Faubourg Saint Antoine
75012 Paris - France
Tel: +33 1 53 46 63 80
Fax: + 33 1 53 46 63 85
info@upperside.fr
<http://www.upperside.fr>

This report is copyright © 2010 EANTC AG. While every reasonable effort has been made to ensure accuracy and completeness of this publication, the authors assume no responsibility for the use of any information contained herein.

All brand names and logos mentioned here are registered trademarks of their respective companies in the United States and other countries.

20100309 v1.2

