



MPLS and Carrier Ethernet: Service – Connect – Transport

Public Multi-Vendor Interoperability Test

Paris, February 2009

MPLS®
ETHERNET

World
congress
2009

■ EANTC ■

EDITOR'S NOTE



Carsten Rossenhoevel
Managing Director

This year, our hot staging test for the MPLS & Ethernet World Congress interoperability event started in Berlin on a freezing January morning – reminding us of the financial crisis partly affecting our industry. Still, 15 vendors carrying 78 devices arrived at our lab for two weeks of state of the art in MPLS and

Carrier Ethernet testing.

Typically, our winter interop event focusing on MPLS is quite technology-savvy — no different this time. 15 test areas were covered in more detail than before. Many basics had been evaluated in the past already, so vendors were able to focus on truly advanced, new tests.

These days, a big concern for service providers is inter-provider connectivity. Many bilateral agreements are being undertaken in practice, but standardization is lagging behind in the Carrier Ethernet area. We tested the current alternatives in MPLS interconnectivity noting the options and the trade-offs between trust and operational effort.

IPv6 has been dormant for a while. With the advent of next-generation mobile networks and due to new warnings about imminent exhaustion of IPv4 addresses, IPv6 has made it back into service provider RfPs. We successfully tested multi-vendor interop for layer 3 VPNs carrying IPv6 traffic – a reassuring deployment message.

In mobile backhaul, vendors were primarily interested in clock synchronization and ATM and TDM transport over packet switched networks this time. The main issue here is the number of alternative solutions limiting interoperability: We tested five different ways to create TDM tunnels only! We also verified two Ethernet based clocking solutions: IEEE 1588-2008 (Precision Time Protocol) and Synchronous Ethernet; both with positive results for a small number of participating vendors.

This white paper describes all the test efforts of the participating vendors and EANTC team in detail. Enjoy the read.

INTRODUCTION

In recent years market forces in the telecommunications industry have been pulling from several directions. There is an emerging need for packet-oriented aggregation infrastructures complementing the traditional SONET/SDH transport solutions with at least equivalent operational characteristics. Since 2007, vendors have kept us busy testing the interoperability of transport solutions like PBB-TE and T-MPLS in addition to MPLS. Service providers selecting between these need to make an almost political decision. At EANTC, we almost got ourselves in trouble at times because we stayed neutral to these options. Fortunately these times are over, and the industry is about to align on standards, converge on technologies and moves on to progress more advanced solutions.

As an independent test lab organizing interoperability events we bind ourselves by three factors:

- The progress and availability of (draft) standards and their implementations
- The interest of network operators in solutions to their particular problems and their interest in interoperable solutions.
- The goals of the vendors participating in our events.

These rules guided the test plan development which took a holistic view of modern networks – networks are converging and more services are being offered in integrated backbones. As such, we tested MPLS services like IPv6 Virtual Private Networks (VPNs), E-Tree service based on MPLS Virtual Private LAN Services (VPLS) and inter-carrier solutions as well as pre-standard implementations of the MPLS Transport Profile (MPLS-TP).

We tested more methods for mobile backhaul transport than in previous years' events. Not only did we see several implementations of transporting TDM traffic directly over Ethernet, but we also tested TDM transport over MPLS, IP, and Ethernet encapsulations. In the mobile backhaul space we also created a comprehensive clock distribution topology testing five different Precision Time Protocol clients.

Maintaining Service Level Agreements (SLAs) is imperative for retaining a customer base of sensitive business and mobile backhaul services. This area was deemed very important by our service provider panel and was of interest for several participating vendors. We used the ITU-T's Ethernet based performance monitoring recommendations (Y.1731) to verify that vendor devices terminating Ethernet based services are able to monitor the service performance for parameters such as delay and loss.

TABLE OF CONTENTS

Participants and Devices	3
Network Design.....	3
Ethernet Services.....	3
Mobile Backhaul.....	5
TDM Circuit Emulation	5
Synchronous Ethernet	10
IEEE 1588-2008.....	10
Performance Monitoring	11
Pre-Standard MPLS-TP	12
T-MPLS.....	12
MPLS Fast Reroute.....	13
IPv6 MPLS/BGP L3 VPNs.....	14
Acronyms.....	16
References	16

PARTICIPANTS AND DEVICES

Vendor	Devices
Alcatel-Lucent	1850 TSS-320 7750 SR-7 7710 C12 7210 SAS 7705 SAR-8 7701 CPAA 5650 CPAM
Brocade Communications	NetIron XMR 8000
Calnex Solutions	Paragon Sync
Celtro	DMT1000 DMT4000
Corrigent Systems	CM4206 CM4140
Ericsson	Marconi OMS 2430 SPT 2770
Huawei Technologies	NE40E CX600
Ixia	IxNetwork
MRV	OS9024-210 OS9024M OS910M OS904 OS906
NEC Corporation	CX2600/220 PASOLINK NEO iP
RAD Data Communications	ACE-3220 ACE-3105 ACE-3400 Egate-100 ETX-202A ETX-202 ETX-102+MiRiCi-E1T1 Gmux-2000 IPmux-216/24 LA-210 RICi-16E1
Redback Networks — an Ericsson Company	SmartEdge 400
Spirent Communications	Spirent TestCenter AX/4000 GEM / ANUE

Vendor	Devices
Telco Systems — a BATM Company	T5C-XG T5C-24F T5C-24G T5C-48T T-Marc-250P T-Marc-254P T-Marc-340 T-Marc-380 T-Metro-200 T-Metro-XG
UTStarcom	TN705 TN725

NETWORK DESIGN

EANTC and the participating vendors designed the test topology to mimic, in as much as possible, a typical deployment scenario as seen in service providers' networks. Two distinct administrative domains were integrated into the test topology in order to enable inter-provider solution testing – a test area which our service provider panel strongly advised to highlight.

The two Autonomous Systems (AS) created in the two MPLS core areas were therefore used to testing inter-provider connectivity. In addition, the MPLS core networks supported Layer 3 services, such as IPv6 VPNs and extended end-to-end services from one access area to the next.

Two distinct aggregation areas were attached to the MPLS cores: one based on MPLS and one on T-MPLS, where pre-standard MPLS-TP implementations were also tested. Both aggregation areas supported Layer 2 services only. The aggregation areas transported, as is often seen in provider networks, various access solutions to the core of the 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.

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.

ETHERNET SERVICES

In previous interoperability test events a great deal of time was spent to set up a large number of end-to-end services across the network. The large amount of different vendor devices and operating systems inevitably forces the participants to manually

configure all services. While there is an obvious market opening here for a provisioning system that can set up such end-to-end services over this amount of devices, we decided to take a different approach in this event.

Two end-to-end services were configured in the network:

- E-Tree – A point-to-multipoint service defined by the MEF which can be used to transport such services as IPTV (from video headend to customers for example) or DSL customers to a central DSLAM.
- E-Line – A point-to-point service type defined by the MEF that allows to connect two sites over a single virtual wire.

Two connection types were configured in the network:

- Intra-domain – connections that originate and terminate in the same aggregation domain
- Inter-domain – connections that originate on a device attached to one aggregation domain, traverse the MPLS core networks, via either 802.1ad S-Tags or inter-domain pseudowires, and terminate on the other aggregation domain.

Figure 1: MPLS Based Ethernet Services depicts the successful intra-domain services within the MPLS aggregation area. The participating vendors constructed MPLS based Ethernet services using RSVP-TE for Label Switched Paths (LSPs) and LDP for pseudowire (PW) signaling.

The diagram below depicts how the E-Tree was configured in terms of the Root UNI, Leaf UNIs, and Logical Connections. A Logical Connection equates to an Ethernet pseudowire within the MPLS domains, and to a T-MPLS TMC in the MPLS Transport domain.

MPLS nodes with multiple pseudowires created a bridging instance to which the pseudowires were associated, and all nodes with multiple downstream connections required additional configuration to impose the rules of an E-Tree service: the Root UNI must have communication with all Leaf UNIs, however, Leaf UNIs may not communicate with each other directly. Spirent's TestCenter was used at each end node to verify this behavior.




Figure 1: MPLS Based Ethernet Services

A similar configuration was used in the MPLS Transport domain to facilitate interoperability

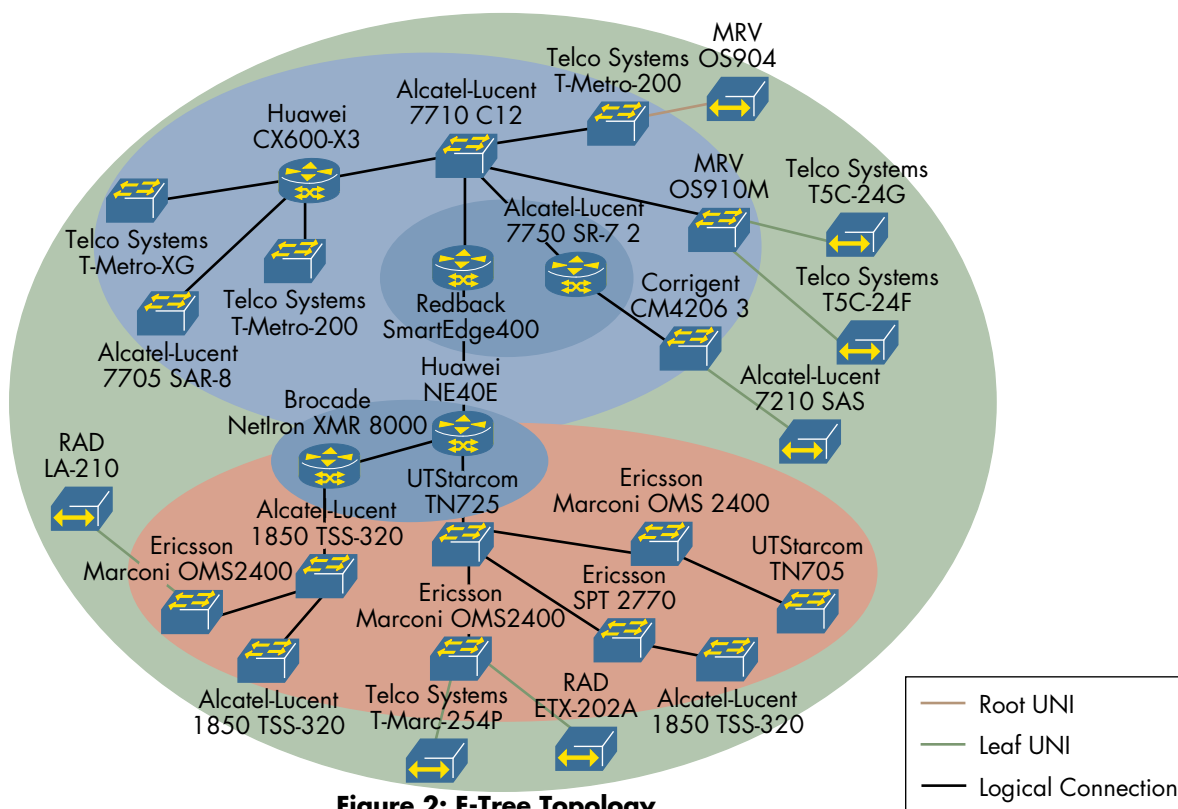


Figure 2: E-Tree Topology

between the vendors participating in that area. The following figure displays the interoperable transport services in that cloud:




Figure 3: T-MPLS and Pre-Standard MPLS-TP Based Ethernet Services

Additionally, services traversed both the T-MPLS and MPLS domains through either S-VLAN tags between the bordering devices or statically configured MPLS labels. In the latter case the statically configured labels translate to a T-MPLS Channel in one direction and an MPLS signaled pseudowire in the other. This was tested between the Alcatel-Lucent 1850 TSS-320 and Alcatel-Lucent 7750 SR-7, Alcatel-Lucent 1850 TSS-320 and Huawei NE40E, and also Alcatel-Lucent 7750 SR-7 and Ericsson Marconi OMS2430.

MOBILE BACKHAUL

The rapid growth of high-speed mobile data services in terms of coverage and bandwidth results in an increasing demand for packet-based mobile backhaul. Fierce competition has driven mobile operators to investigate lowering the backhaul cost. The migration towards packet based mobile backhaul is a logical consequence and has been the topic of various industry forum initiatives, specifically in the MEF and IP/MPLS Forum.

This time, our tests in this area primarily verified the interoperability of ATM and TDM transport mechanisms over packet based networks.

We tested TDM Circuit Emulation Service (CES) over IP, MPLS, and Ethernet, as well as ATM over MPLS. These mobile backhaul transport solutions are key components to the migration path from TDM/ATM to converged packet based services and are essential to the continuing support of legacy base stations which may or may not be upgraded.

All tests were performed by connecting the devices under test back-to-back via their packet interfaces in order to facilitate a maximum number of test pair results. In some cases an impairment generator was used in order to emulate real network conditions between the DUTs.

TDM Circuit Emulation

There are two basic modes for TDM circuit emulation pseudowires: structure-agnostic and structure-aware. While structure-agnostic mode transparently transports the whole E1 frames, structure-aware interprets the E1 timeslot zero, which is used for OAM, and may be configured to transport any subset of E1 timeslots over a packet switched network.

Standards for TDM pseudowires and circuit emulation exist from both the IETF and the MEF. Therefore in this interoperability test event we tested a palette of solutions to transporting TDM traffic over PSN according to the following specifications:

- Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP) – RFC 4553
- Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks – MEF 8
- Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN) – RFC 5086

The frame carrying the emulated TDM traffic may be transported over any kind of PSN. In this event we tested the following: IP, dynamically signalled MPLS LSPs, statically configured MPLS LSPs, and Ethernet.

The devices implementing TDM Circuit Emulation have to exchange clock information in order to synchronize the TDM circuits. We used adaptive clock synchronization for all TDM CES tests. For each device combination, a test run was performed twice in order to test slave and master functionality of the adaptive clock synchronization method on each device. During each test run, an E1 traffic generator was connected to both devices under test. The analyzer generated an E1 signal on one end of the test pair and verified that this signal was received by the far end of the connection. After ten minutes of an error free test run, the PSN connection between the two devices under test was interrupted and restored (by pulling/plugging back in the PSN link between the devices) in order to test that the synchronization function recovers from a network failure and is able to quickly bring back both devices to a synchronous state.

Eight E1 frames were packetized in a single PDU, which resulted in a 256 byte payload for the structure-agnostic mode and between 8 to 248 byte payload for structure-aware mode. The payload size was dependant upon the number of E1 timeslots transported by the service. The packetization of eight E1 frames is a reasonable trade-off between overhead and additional packetization delay and jitter.

For the structure-aware tests, vendors were requested to configure two ranges of E1 timeslots for circuit emulation. The E1 tester verified that only the configured timeslots were forwarded.

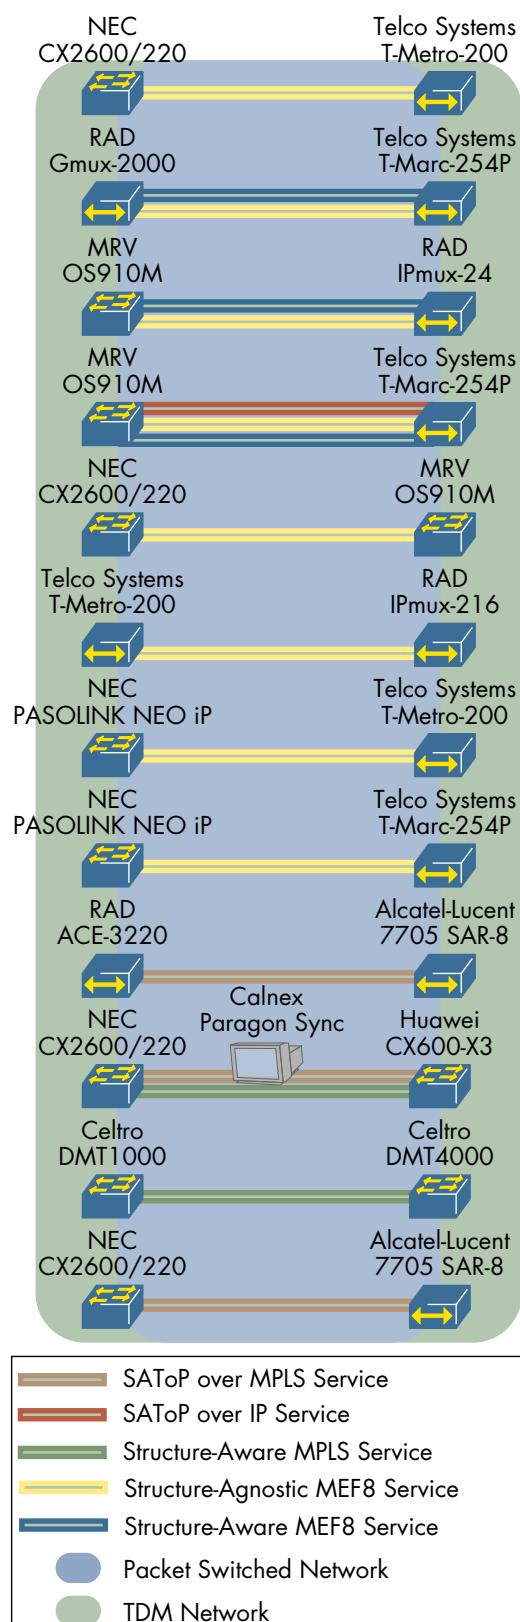


Figure 4: TDM Circuit Emulation Services

The following devices tested structure-agnostic TDM Circuit Emulation according to the MEF Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks (MEF8): Alcatel-Lucent 7705 SAR-8, MRV OS910M, NEC CX2600/220, NEC PASOLINK NEO iP, RAD IPmux-216, RAD IPmux-24, RAD Gmux-2000, Telco Systems T-Metro-200, and Telco Systems T-Marc-254P.

In addition, MRV OS910M, RAD Gmux-2000, RAD IPmux-24, and Telco Systems T-Marc-254P tested structure-aware TDM Circuit Emulation according to MEF 8.

The MRV OS910M and Telco Systems T-Marc-254P tested structure-agnostic TDM pseudowire over IP according to RFC 4553.

The Alcatel-Lucent 7705 SAR-8, Huawei CX600-X3, NEC CX2600/220, and RAD ACE-3220 tested structure-agnostic TDM pseudowire over MPLS according to the RFC 4553.

Huawei CX600-X3 and NEC CX2600/220 tested the interoperability of MPLS based structure-aware CES according to the RFC 5086. In addition, Celtro demonstrated their structure-aware CES between their DMT1000 and DMT4000.

We used a Calnex Paragon Sync impairment generator for the structure-agnostic TDM over MPLS test between Huawei CX600-X3 and NEC CX2600/220. The impairment generator emulated real network conditions of a PSN with 80% average load by adding jitter, packet duplication, corrupting and deleting some packets.

No issues were found while performing the structure-agnostic tests. In the case of the structure-aware tests, the CLI of one vendor created confusion and led to some temporary failed results. While all vendors count the E1 timeslots from 0 to 31, this one vendor counted the timeslots from 1 to 32. This caused a mismatch between the configured slots which was resolved once the issue was understood.

In another case a vendor could not configure multiple E1 timeslot ranges to be mapped into a single virtual connection of a circuit emulation service.

ATM over MPLS Transport

Encapsulation for ATM over MPLS is defined in the IETF RFC 4717. This RFC specifies two modes of ATM virtual circuit or virtual path mapping to an MPLS pseudowire: One-to-one and N-to-one. In order to decrease overhead, the standard specifies "cell concatenation" which allows the encapsulation of multiple cells into a single MPLS PDU.

Initially vendors were requested to configure the N-to-one VCC mapping, with at least two VCCs mapped to a single MPLS pseudowire. However, since some vendors support N=1 only, the tests were performed with N-to-one VCC mapping with N=1. The cell concatenation was set either to 10 cells, or in the case that a vendor did not support it, it was disabled (meaning each cell was transported by a separate MPLS PDU). The concatenation time out was set to 10 ms.

The tests were performed with E1 and STM-1 physical interfaces. The ATM traffic generator sent VBR cells to two ATM VCCs: one that was configured on the devices and another which was not configured. A test passed if all cells of the VCC configured on DUTs were forwarded, but cells sent to the non-configured VCC were dropped.

RFC 4717 does not require usage of an MPLS signaling protocol to setup an ATM pseudowire. We tested statically configured ATM pseudowires as well as dynamically signalled pseudowire using LDP as specified in RFC 4447.

Alcatel-Lucent 7705 SAR-8, Celtro DMT1000, Celtro DMT4000, Huawei CX-600-X3, NEC CX 2600/220, and RAD ACE-3220 successfully tested dynamically signalled ATM pseudowires over MPLS as shown in the figure below.

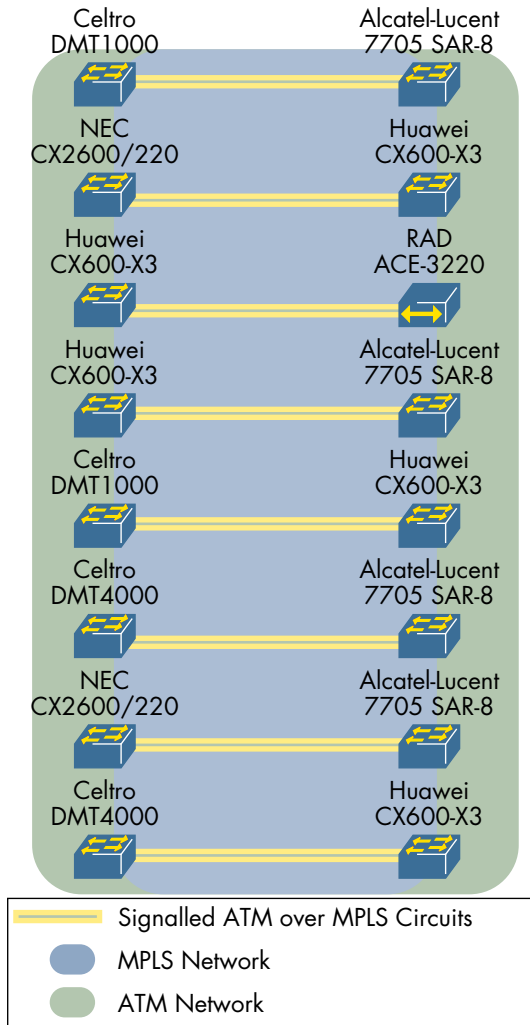


Figure 5: Signalled ATM PW Results

Alcatel-Lucent 7705 SAR-8, RAD ACE-3105, NEC CX 2600/220, NEC PASOLINK NEO iP, and UTStarcom TN725 successfully tested statically configured ATM pseudowires over MPLS as shown in Figure 6: Statically Configured ATM PW Results.

We detected several issues which were resolved during the hot staging event. We also note that different options from the RFCs were chosen amongst the implementations. For example support of cell concatenation, support of N-to-one VCC Service encapsulation with $N > 1$, and support of MPLS penultimate hop popping (PHP) were all implemented in some but not all devices.

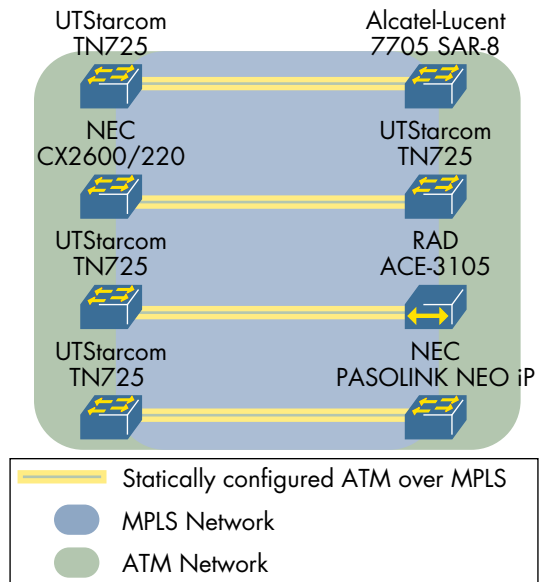


Figure 6: Statically Configured ATM PW Results


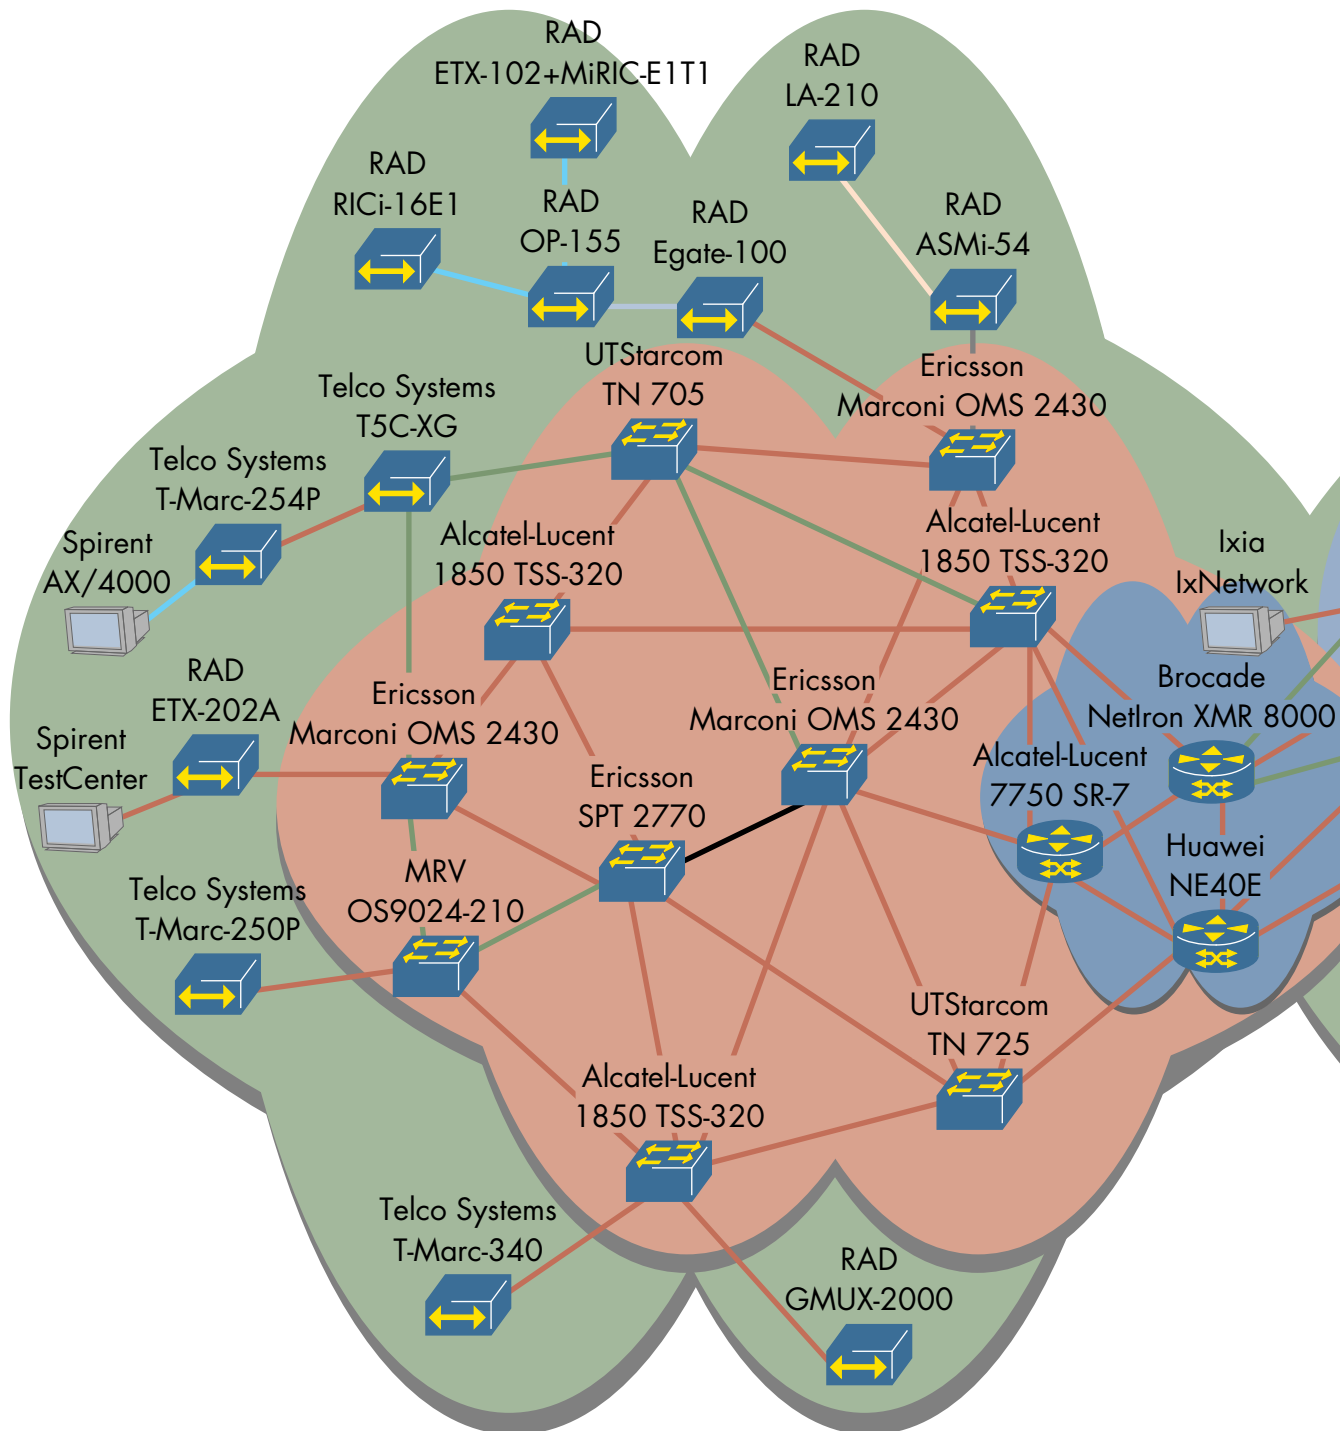
MOBILE BACKHAUL SYNCHRONIZATION REQUIREMENTS

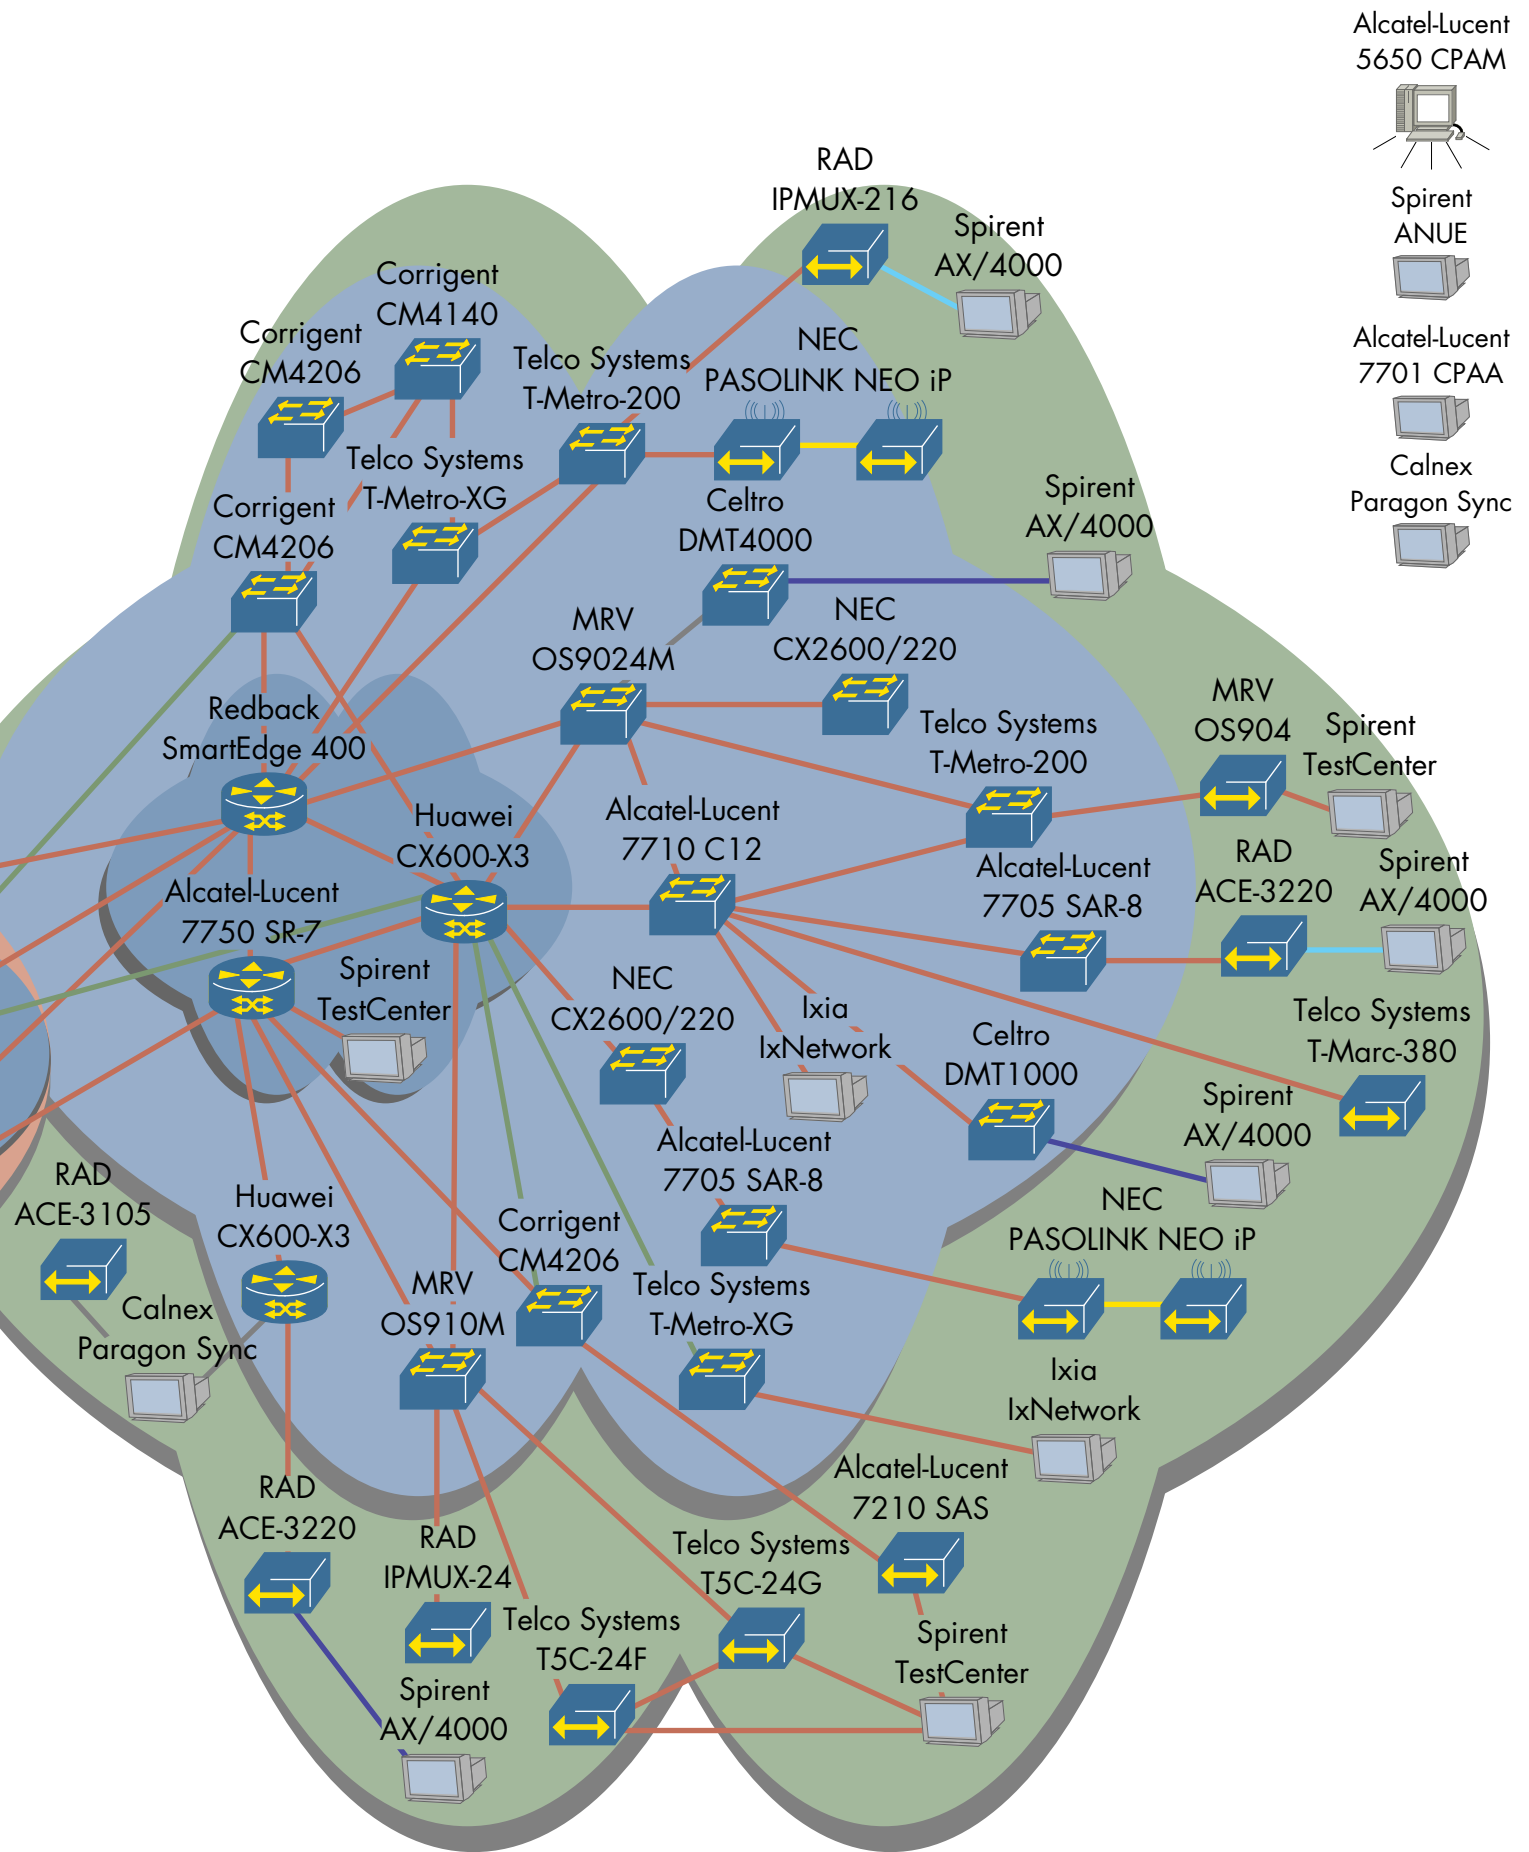
In the "TDM Circuit Emulation" test area described above, we focused on the Interworking Functions (IWFs) that convert data between a TDM circuit and a packet based interface. Another major functionality of such an IWF is to synchronize the local TDM interface with the remote TDM interface served by the remote IWF. This was tested in "TDM Circuit Emulation" for G.823 E1 traffic interface types implicitly since the E1 analyzer was able to receive E1 data from DUTs.

The 2G base stations and 3G NodeBs have higher requirements for their radios' interface synchronization than the requirements for E1 traffic interface quality. The mobile backhaul networks connecting their base stations with their base station controllers, or NodeBs with their radio network controller, have to provide precise frequency synchronization mechanisms regardless of the transport used. Two such methods are the IEEE 1588-2008 and Synchronous Ethernet, both of which were tested in our event for meeting the high synchronization requirements of mobile operators.

The exact requirements on the synchronization depend on the mobile equipment used by the mobile operator and the mobile services. However, the E1 synchronization interface specification of SDH Equipment Clock (SEC) meets most mobile operator requirements. Therefore, the measured Maximum Time Interval Error (MTIE) during the current event was compared with the ITU-T recommendation G.823 SEC mask.

PHYSICAL NETWORK TOPOLOGY





Synchronous Ethernet

Synchronous Ethernet is an upcoming technology for delivering synchronization in a packet based network. The idea of synchronous Ethernet is to provide synchronous capability on the PHY layer similar to SONET/SDH or PDH technologies. An advantage, in comparison to the packet based synchronization methods (like IEEE 1588 or NTP), is that the accuracy of synchronization does not depend on the network load conditions. A trade-off is that all Ethernet network nodes between the slave and the master clock must support synchronous Ethernet in order to deliver the clock information. Therefore, a possible deployment scenario will use Synchronous Ethernet devices in the access networks potentially combined with other synchronization methods transporting clock information across wide area networks.




Figure 7: Synchronous Ethernet Test Results

In order to measure clock accuracy, we connected the devices under test (DUTs) with each other via Synchronous Ethernet a Calnex Paragon Sync tester. The Calnex Paragon device forwarded the clock information unaltered between the two devices and measured the clock accuracy in reference to a common reference clock.

The first DUT (the Synchronous Ethernet master clock) received the signal from the reference clock, locked its internal clock to it, and forwarded this signal over Synchronous Ethernet to the second DUT. The second DUT (the Synchronous Ethernet slave) derived its clock from the received information sent by the master. The second DUT fed a wander analyzer using an E1 interface that was synchronized to the clock signal derived from the Synchronous Ethernet. Both the Calnex Paragon Sync and the E1 wander analyzer received directly the same reference common clock as the clock master device.

The Calnex Paragon Sync measured the clock accuracy signal received from master clock DUT on Synchronous Ethernet in reference to the common clock. In addition, the wander analyzer measured clock accuracy on E1 interface received from the slave DUT that was derived from Synchronous Ethernet. The measurements were compared to the G.823 SEC and G.813 masks which are identical to the Synchronous Ethernet equipment clock mask defined by the ITU-T G.8262 recommendations.




Figure 8: MTIE Results Example for Synchronous Ethernet

IEEE 1588-2008

The IEEE 1588-2008 standard specifies version 2 of the Precision Time Protocol (PTP) for precise synchronization between a clock master and clock slave devices.

Similarly to the Synchronous Ethernet tests, we measured clock accuracy, derived from the IEEE 1588-2008 implementations, by measuring the signal deviation from its reference on the E1 receiving end (MTIE wander measurements). The PTP slave DUT fed a wander analyzer connected over an E1 interface with the signal carried over the PSN while at the same time the wander analyzer received the grandmaster clock signal and compared between the two. The measurements were compared to the G.823 SEC and G.813 masks.

Ultimately we were able to successfully verify interoperability between Huawei CX600-X3 and RAD ACE-3105 with a Calnex Paragon Sync impairment generator between the two. Huawei CX600-X3 was acting as the grandmaster clock receiving its clock signal from a rubidium clock source. The RAD ACE-3105 was acting as PTP client that received the PTP packets encapsulated in IP/UDP from the Huawei CX600-X3 router. Calnex Solutions' Paragon Sync was used to monitor the connection between the master and slave clocks, and to emulate a network between both DUTs. The impairment generator emulated real network conditions of a PSN network with 40% average load by adding jitter specifically for the PTP packets. In addition, two Celtro DMT1000 switches demonstrated IEEE 1588-2008 without an impairment generator.

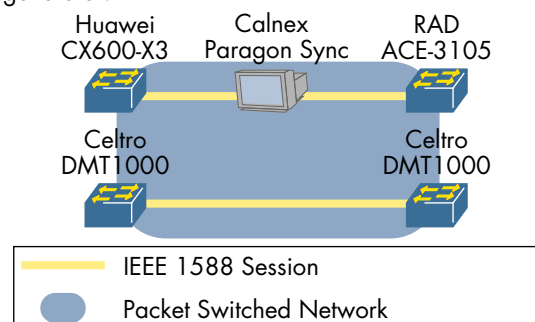


Figure 9: IEEE 1588-2008 Interoperability Results

During the preparation for the testing we found an array of options specified in IEEE 1588-2008 which were supported only by some of the interested vendors. These options included the transport frame format (unicast or multicast), grandmaster functionality, 1 Step or 2 Step and Boundary Clock functionality support. Out of the five vendors interested in this test area several restrictions were discovered allowing only certain test pairs to interoperate.

IEEE 1588-2008 Functions	Supported by All?
Multicast Destination MAC Address	Yes
Unicast Destination MAC Address	No
IP UDP	Yes
1 Step, 2 Step	No
Grandmaster	No
Boundary Clock	No

Another general interoperability issue that we discovered was incompatibilities between different implementations' SYNC messages rate. While all vendors supported configurable SYNC message rates, the actual ranges differed so that some implementations could not synchronize with each other.

PERFORMANCE MONITORING

Ethernet performance monitoring is defined in the ITU-T recommendation Y.1731. The recommendation specifies an array of performance monitoring messages of which only delay and delay variation measurements were tested.

These tests were performed using Calnex Solutions' Paragon Sync and Spirent's ANUE impairment devices. First, messages were exchanged without any impairment to show baseline interoperability. Then, 30 ms delay impairments were generated using the Calnex Paragon Sync. The test pairs using Spirent's ANUE were also required to react to an increase in jitter by the ANUE emulating an Internet traffic jitter distribution.



Figure 10: Jitter Impairments

The following devices tested two-way frame delay and delay variations: MRV OS906, Ixia IxNetwork, NEC CX2600/220, NEC PASOLINK NEO iP, RAD

ETX-202A, RAD ETX-102+MiRIC-E1T1, RAD LA-210, RAD RICi-16E1, and Telco Systems T-Marc-340. Telco Systems T-Marc-380 and RAD ETX-102+MiRIC-E1T1 tested frame delay and frame delay variations across backbone, and MRV demonstrated frame delay measurements across the MPLS Transport area using the OS9024-210 and OS906.

The standard Ethernet performance monitoring was also demonstrated as being agnostic to the access links. The RAD RICi-16E1 provided this functionality over eight bonded E1 access lines, the RAD ETX-102 (equipped with a miniature Ethernet-to-E1 bridge in a standard SFP enclosure, MiRIC-E1) - over a single E1 line, and the RAD LA-210 - over four bonded SHDSL EFM links.

In most cases Delay Measurement Messages (DMMs) were replied to with Delay Measurement Responses (DMRs), and the respective calculations showed a high degree of accuracy between the different devices.

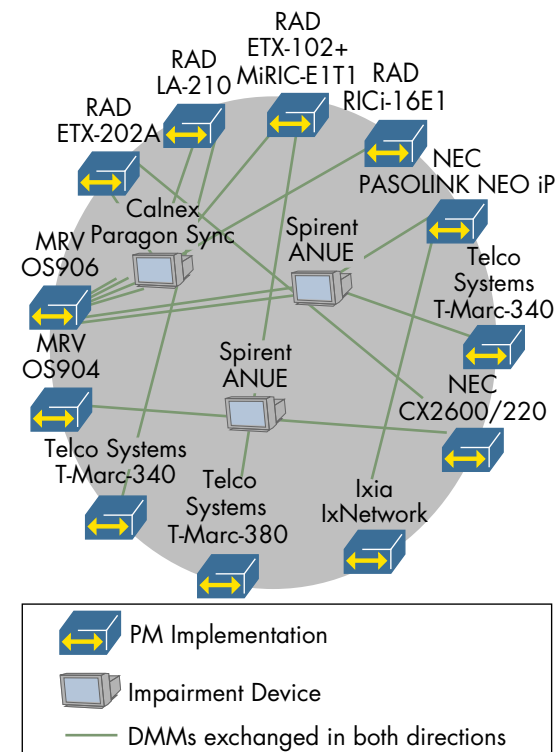


Figure 11: Performance Monitoring Tests

We discovered an issue between two vendors participating in this test. The messages exchanged include a field called End TLV which, according to the standard, should be an all-zero octet. This all-zero octet must be copied into the response message. One vendor relied on the test partner to blindly copy the End TLV format which was not set to zero when answering the measurement request. However, the answer was returned in the form of an all-zero end TLV, as per the standard. The vendor initiating the measurements could therefore not read the received messages and the test did not pass.

Additionally, NEC Corporation demonstrated frame loss measurement and two-way frame delay and frame delay variation measurements between NEC CX2600/220 and PASOLINK NEO iP.

TRANSPORT AND CONNECTIVITY

In this test area, we focus assessing the ability of a technology to provide transport for services, connectivity verification and resiliency. Two distinct network domains existed, MPLS and MPLS Transport (where pre-standard MPLS-TP and T-MPLS were tested). We carried out tests for inter-domain connectivity and resiliency within both network domains.

The Road to MPLS-TP

In 2008 the IETF and the ITU-T agreed to work jointly to bring transport requirements into the IETF and extend IETF MPLS forwarding, OAM, survivability, network management and control plane protocols to meet those requirements through the IETF Standards Process. The two standard bodies agreed to refer to the technology as the Transport Profile for MPLS. They have since been working within the IETF process framework on the definition of the MPLS Transport Profile (MPLS-TP). As the output becomes stable, the ITU-T plans to use the IETF documents as a basis to revise the currently in force T-MPLS Recommendations and to update them with recommendations aligned with MPLS-TP. The future versions of the ITU-T Recommendations will use the name MPLS-TP to reflect the full alignment with the IETF.

A separate domain was established for vendors participating in MPLS transport testing, within which all related testing was performed including forwarding, OAM, and protection.

Pre-Standard MPLS-TP

While much work is still to be done in defining MPLS-TP, some vendors who had implemented T-MPLS have a head start in testing pre-standard MPLS-TP implementations. The IETF, at the time this document was written, is working on several working group drafts including one that is called »MPLS Generic Associated Channel Header«. The Internet-Draft defines a Generic Associated Channel Header (G-ACH) which may be used to transport OAM messages. Further, the Internet-Draft defines a Generic Associated channel header Label (GAL) over which the G-ACH may run. Within the G-ACH, a Channel Type field exists to specify which control or maintenance protocol lies under the G-ACH.

While specific code-point/protocol pairs are still to be defined, the Internet-Draft defines two code-point ranges: One for protocols still to be defined by the IETF, and one for experimental use. Since specific codes are not chosen, the vendors participating in this MPLS-TP test configured the code-point values within the experimental range, with the intention to change them if and when they will be defined by the IETF.

Alcatel-Lucent 1850 TSS-320 and Ericsson SPT 2770 tested this function as an alternative way to transport APS, and also for the transport of a pre-supposed MPLS-TP OAM protocol — Bidirectional Forwarding Detection (BFD). While BFD is not a new protocol, the use of BFD over non-IP-enabled interfaces has not been defined by the IETF yet. The devices under test employed BFD in place of CV as a mechanism to detect failure on a given path. When BFD detected a failure, the APS protocol signaled the 1:1 switchover similarly to that defined in G.8131 as described below.

T-MPLS

The ITU-T originally defined Transport MPLS (T-MPLS) as a solution for packet transport networks. Four vendors were involved in testing their T-MPLS implementations, showing their dedication to using MPLS as transport and the evolution to MPLS-TP. These vendors were Alcatel-Lucent, Ericsson, MRV, and UTStarcom. For further information on the status and process of the standardization, please see the blue box below.

The T-MPLS forwarding plane uses the same packet format defined by the IETF for the MPLS forwarding plane. The following devices successfully established T-MPLS based TMPs (paths) and TMCs (channels) for the transport of Ethernet traffic: Alcatel-Lucent 1850 TSS-320, Ericsson Marconi OMS2430, Ericsson SPT 2770, MRV OS9024-210M, UTStarcom TN705 and UTStarcom TN725.




Figure 12: Test Setup for T-MPLS and MPLS-TP Protection Tests

To support sub-50 ms resiliency mechanisms in T-MPLS, the ITU-T has defined protection mechanisms using a packet based version of the SONET/SDH Automatic Protection Switching (APS) protocol in the G.8131 specification. The mechanisms define both a 1:1 active/backup path scheme and 1+1 protection, in which data flows through both paths simultaneously. The implementations of G.8131 tested currently rely upon either Y.1711 or vendor-specific OAM mechanisms to detect whether or not a given path is functional. APS is then used to signal the switchover. The OAM mechanisms are detailed

below. The Alcatel-Lucent 1850 TSS-320, Ericsson Marconi OMS2430, Ericsson SPT 2770, and UTStarcom TN705 tested T-MPLS protection, which is more specifically shown in Figure 12: Test Setup for T-MPLS and MPLS-TP Protection Tests.

APS administrative commands including manual, force, clear, and lookout were tested between Alcatel-Lucent 1850 TSS-320 and Ericsson Marconi SPT 2770 as well.

One OAM mechanism is Connectivity Verification (CV), which can detect failures along a T-MPLS path. Initially we encountered an interoperability issue between the CV message frame formats used by the participants. Vendors implemented two versions of the protocol, each with its own frame format: CVv0 defined by ITU-T Y.1711, and CVv1 defined in an ITU-T »frozen« draft, trying to describe it the most unbiased way. Two vendors implemented CVv1 while one implemented CVv0. The two message types each had a different frame format which made them incompatible. During the course of the test, the vendor implementing CVv0 updated its code and was able to successfully test CVv1 with the others. Alcatel-Lucent 1850 TSS-320 and Ericsson STP 2770 devices were successfully able to detect misconfigurations and misconnections created within the network, using CVv1.

At EANTC, we had a difficult time choosing appropriate OAM mechanisms for the interop test because the standardization effort in the ITU-T has been frozen in an unpublished draft version when the joint working team was created. There is a huge discussion going on about the status of OAM protocols for T-MPLS in the industry between vendors who focus investment protection for their existing implementations and others who primarily want to provide operators with a clean, clear vision for the future of MPLS-TP. Interestingly, operators have not voiced strong opinions in public yet. Time will tell what exactly the future OAM standard for MPLS-TP will look like.

We found ourselves amidst the industry forces — much alike the PBB-TE testing two years ago. In lack of fully cooked standards, we chose to give vendors a playground for experimental multi-vendor testing of everything available; understanding that there is no such thing as a perfect, unanimously agreed to solution in this area today.

MPLS

Two areas in the network used MPLS: the network core and one of the metro/aggregation areas. Several tests, described over the next few pages, were conducted in these areas.

MPLS Fast Reroute

MPLS Fast Reroute has been repeatedly tested at EANTC's interoperability events. While at some point we hope to reach a stage at which we could safely say that all existing MPLS Fast Reroute imple-

mentations have been tested, at each test event we receive new implementations and with them the demand to test Fast Reroute.

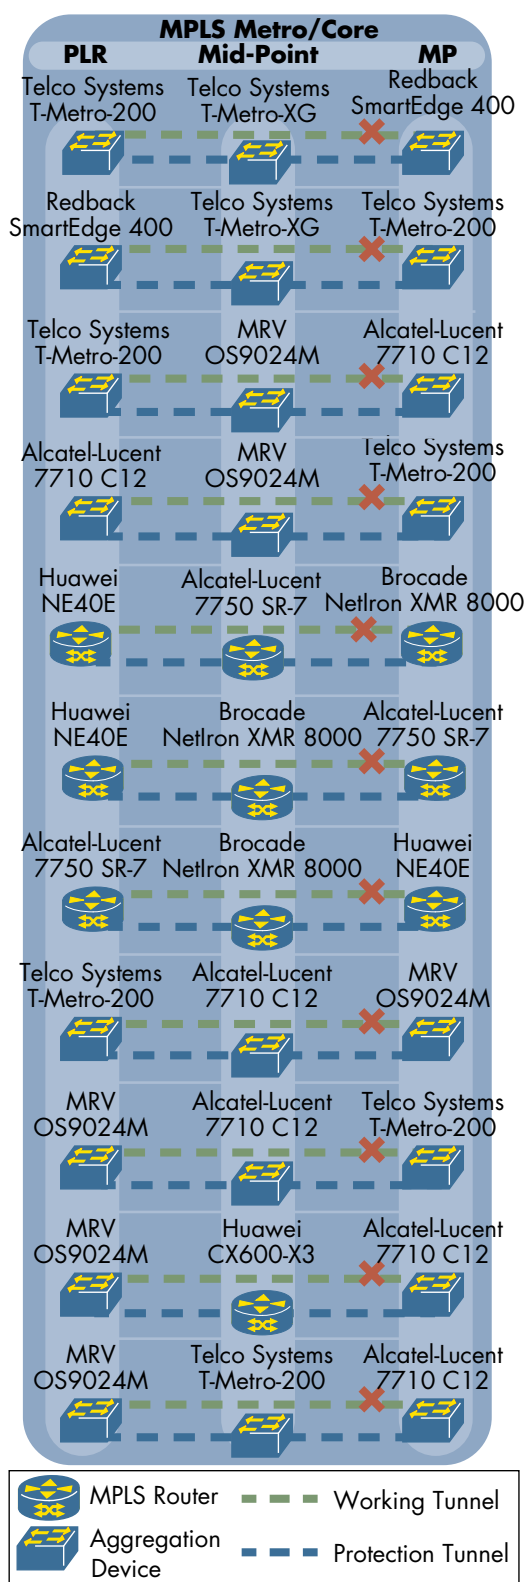


Figure 13: MPLS Fast Reroute Tests

We tested MPLS Fast Reroute by artificially triggering a failover on the primary path between the Point of Local Repair (PLR) router and the Merge Point router. These two routers, respectively, served as the Label Switched Path (LSP) ingress router and egress router. We generated a Loss of Signal (LOS)

by pulling the active link from the PLR. This phase of the test was the simulated failure scenario and, once successful, was followed by a recovery test in which the broken link was reconnected.

Six devices participated in the testing. The following devices participated in the role of Point of Local Repair (PLR): Alcatel-Lucent 7710 C 12, Huawei NE40E, MRV OS9024M, Redback SE400 and Telco Systems T-Metro-200. The Merge Point (MP) which is the router that merges the backup and primary segments of an MPLS tunnel, were: Alcatel-Lucent 7710 C12, Alcatel-Lucent 7750 SR-7, Brocade NetIron XMR 8000, Huawei NE40E, MRV OS9024M, Redback SE400 and Telco Systems T-Metro-200. In six test combinations we measured failover times below 50 ms, two tests were under 200 ms, and in one test case we measured over 700 ms. The importance of the test, however, was the protocol interoperability between the two routers; therefore we concentrated on verifying that both PLR and MPs are able to correctly exchange RSVP-TE extension messages with each other as opposed to spending the time in tweaking configurations for faster convergence times.

In one test case, the DUT became unstable after the recovery phase of the test. The engineers collected debugging information in order to further analyze the issue later. Further interoperability issues were not found, however, we recommended that further analysis would be performed by one vendor who recorded a high out-of-service time.

The Corrigent CM4206 successfully established an LSP with another CM4206 traversing an Alcatel-Lucent 7750 SR-7 and Redback SmartEdge 400. This LSP was protected between the Alcatel-Lucent and Redback devices with Fast Reroute.




Figure 14: IPv6 VPNs

IPv6 MPLS/BGP L3 VPNs

The time and need for service provider backbone migration to IPv6 has been debated quite a bit, however it has not been a huge concern for service providers. More importantly, some enterprise customers start requiring IPv6 VPN service readiness from their service providers.

The IETF has introduced extensions for BGP/MPLS

VPNs such as an IPv6 address family to accommodate for customers using IPv6. Interoperable solutions for this were successfully tested by the following devices: Alcatel-Lucent 7750 SR-7, Huawei CX600-X3, Huawei NE40E, Redback SmartEdge 400, and Spirent TestCenter. Additionally, these scenarios were configured and tested over the different ASes as shown in Figure 16: Inter-Carrier L3VPN.

Management

Alcatel-Lucent brought transparency to the network configuration via their 5650 Control Plane Assurance Manager (CPAM). The 5650 CPAM is a vendor agnostic route analytics solution, which provides real-time configuration insight into the IP/MPLS network. The figure below shows a screen shot representing IGP metric information — highlighting mismatches to be resolved.

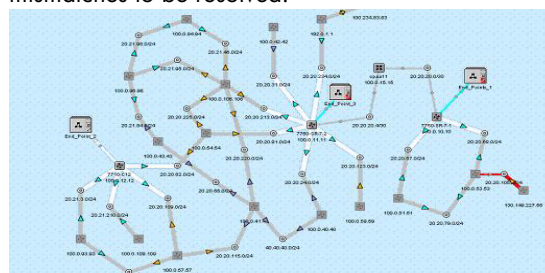


Figure 15: CPAM 5650 Interface

Inter-Carrier MPLS/BGP VPNs

We tested Layer 3 VPN services across the two MPLS domains according to the three options A, B and C described in RFC 4364. In Option A, the adjacent PE router is treated as a CE device and thus uses no MPLS between the domains. With Option B, the AS Border Routers (ASBR) use EBGP to exchange labeled routes from their respective ASes, enabling MPLS between the ASBRs. Option C allows ASBRs to forego their responsibility of maintaining VPN routes. They instead use EBGP to distribute loopback IP address from the PEs in one AS to the PEs in the opposite AS, allowing for end-to-end LSPs. The following vendor pairs were successfully tested:




Figure 16: Inter-Carrier L3VPN

ACRONYMS

Term	Definition
ATM	Asynchronous Transfer Mode
AS	Autonomous System
ASBR	Autonomous System Border Router
BFD	Bidirectional Fault Detection
BGP	Border Gateway Protocol
BSC	Base Station Controller
CES	Circuit Emulation Service
CPE	Customer Premise Equipment
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
DSL	Digital Subscriber Line
E-Line	Point-to-Point Ethernet Service
E-NNI	External Network-to-Network Interface
FRR	Fast ReRoute
HSPA	High-Speed Packet Access
LDP	Label Distribution Protocol
LOS	Loss Of Signal
LSP	Label Switched Path
MPLS	Multi-Protocol Label Switching
MPLS-TP	MPLS Transport Profile
MSC	Mobile Switching Center
MTU-s	Multi Tenant Unit Switch
OAM	Operations, Administration and Maintenance
OSPF	Open Shortest Path First
PE	Provider Edge
PLR	Point of Local Repair
PSN	Packet Switched Network
PTP	Precision Time Protocol
PW	PseudoWire
RFC	Request For Comments
RNC	Radio Network Controller
RSVP-TE	Resource reSerVation Protocol Traffic Engineering
S-Tag	Service Tag
SAToP	Structure-Agnostic Time Division Multiplexing (TDM) over Packet
SLA	Service Level Agreement
T-MPLS	Transport MPLS
TMC	T-MPLS Channel
TMP	T-MPLS Path
UMTS	Universal Mobile Telecommunications System
UNI	User-Network Interface
VLAN	Virtual Local Area Network
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network

REFERENCES

Terms and Definitions for Transport MPLS (G.8101)

Architecture of Transport MPLS (T-MPLS) Layer Network (G.8110.1)

Interfaces for the Transport MPLS (T-MPLS) Hierarchy (G.8112)

Characteristics of multi protocol label switched (MPLS) equipment functional blocks (G.8121)

Provider Bridges (IEEE 802.1ad-2005)

RSVP-TE: Extensions to RSVP for LSP Tunnels (RFC 3209)

IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities (RFC 5073)

An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge (draft-ietf-pwe3-ms-pw-arch-05)

BGP/MPLS IP Virtual Private Networks (VPNs) (RFC 4364)

Provider Bridges (IEEE 802.1ad-2005)

Fast Reroute Extensions to RSVP-TE for LSP Tunnels (RFC 4090)

BFD For MPLS LSPs (draft-ietf-bfd-mpls-05.txt)

Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH) Circuit Emulation over Packet (CEP) (RFC 4842)

Requirements for Edge-to-Edge Emulation of Time Division Multiplexed (TDM) Circuits over Packet Switching Networks (RFC 4197)

Ethernet Service Definitions - Phase 2 (MEF 6.1)

BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN (RFC 4659)

Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP) (RFC 4553)

Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN) (RFC 5086)

Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks (MEF 8)

Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks (RFC 4717)

Precision Time Protocol (PTP) (IEEE 1588-2008)

Timing and Synchronization Aspects in Packet Networks (ITU-T G.8261/Y.1361)

The Control of Jitter and Wander within Digital Networks which are based on the 2048 kbit/s Hierarchy (ITU-T G.823)

The Control of Jitter and Wander within Digital Networks which are based on the 1544 kbit/s Hierarchy (ITU-T G.824)

OAM Functions and Mechanisms for Ethernet Based Networks (ITU-T Y.1731)



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 © 2009 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.

20090204 v1.0