



Carrier Ethernet World Congress 2010
Multi-Vendor Interoperability Event
White Paper

Evolving Universal Services

Warsaw, September 20-23, 2010
Hong Kong, December 1-3, 2010



EDITOR'S NOTE



Carsten Rossenhövel
Managing Director

As I am writing this note, our team prepares to ship five tons of equipment — more than 100 devices under test from 24 participating vendors — to Warsaw in Poland and, later, some of it to Hong Kong.

The new Eastern European location of the Carrier Ethernet World Congress and the flourishing Carrier Ethernet

World APAC conference in the Asia Pacific region are symbolic for Carrier Ethernet expansion into new markets and applications. Analysts across the board see two thirds of the world's Carrier Ethernet market volume in Europe and APAC, with double-digit cumulated annual growth rates¹. More carriers than ever before substitute legacy private lines with Carrier Ethernet offerings for businesses, consumers and mobile backhaul.

In our two-week hot staging test, we noticed that mobile backhaul solutions are driving a lot of innovation in the industry. It is great to see that the industry keeps pace with challenging operator requirements, even in multi-vendor scenarios. With the advent of Long Term Evolution (LTE), six vendors interoperated successfully in the world's first multi-vendor end-to-end phase clock synchronization test.

Ethernet-based microwave and millimeter-wave solutions — helping to overcome the fiber gap in the last mile, specifically to cell sites — are another important area of innovation. Five otherwise competitive

vendors joined our test for a collaborative evaluation of their latest advances side by side, such as adaptive modulation, integration in Ethernet OAM, and clock synchronization support. The speed of technological advances in this sector is amazing as well, and certainly fueled by the strong competition.

A range of high availability technologies were a recurring topic this year. Enterprises and mobile operators recognize the cost of a network outage — with IPTV deployments, highly available aggregation networks are even a need for consumer services. Draft and pre-draft standard MPLS-TP protection, Ethernet ring protection (ERPS) and pseudowire redundancy in the access were three main focus areas. We saw most advances with eight interoperable solutions in ERPS (up from four last time).

Management topics, specifically OAM for fault management and performance monitoring, are on the agenda for many service provider Requests for Proposal (RFPs) now. At EANTC, these tests are a staple — we have tested Ethernet OAM since 2005 and Y.1731 performance monitoring since 2008. It

is great to witness how the industry matures. Delay performance monitoring went really well with 12 implementations; there is still more work upcoming for loss performance monitoring.

This year's EANTC Carrier Ethernet interoperability test shows the industry in an intermediate state. Basic services were taken up across 23 vendors in minutes rather than days this time; key technologies such as Ethernet OAM, MPLS signaling, Synchronous Ethernet worked like a charm; now advanced solutions for challenging markets are being targeted. Carrier Ethernet is on the way to the next level of enlightenment.

INTRODUCTION

There are several metrics with which the scale of our event can be recorded. In terms of two which we find to be the most noteworthy, this was indeed our largest event yet — effort, and results. Over 80 engineers worked together in the two week test event in our lab, bringing what we feel is more test results than we have ever had in the past.

Two years since our first tests of IEEE 1588, this is perhaps the first event with some (not all!) virtually seamless test runs. The same could be said about Ethernet Service OAM (802.1ag) where so many link trace tests were conducted, we could barely fit them in a single-page diagram. These tests were much more straightforward than in the past.

This is evidence of two factors. First, the wide-spread support across vendor equipment and device types has increased. Second the implementation have grown in maturity. On the other hand, we welcomed five first-time participants this year.

This is not to say that all tests were seamless. In fact, we experienced as many hiccups and interesting issues as we have in the past if not more. Most, if not all, vendors fixed issues based on the findings of the testing, in some cases re-running the test for a successful result. In this white paper, we point out issues we think are most relevant for those who standardize or deploy the technology.

The test plan was divided into four main umbrella test areas - Synchronization, Transport, Resiliency, and Management. We are open to additional topics for future events and encourage you to send us your suggestions and comments.

TABLE OF CONTENTS

Participants and Devices.....	3
Interoperability Test Results	4
Synchronization	4
Converged Transport	8
Managed Ethernet Services	9
Carrier Ethernet Resiliency	15
References	19
Acronyms	19

1. Source: Vertical Systems Group presentation to the MEF, January 2010

PARTICIPANTS AND DEVICES

Vendor	Participating Devices
Actus Networks	Ganesh
Alcatel-Lucent	1850 TSS-160 1850 TSS-320 7750 SR7 7210 SAS-M 7705 SAR8 7750 SR-c12 9500 MPR 9500 MPR-e
Albis Technologies	ACCEED 2202 ACCEED 1416
Aviat Networks	Eclipse Packet Node
Calnex Solutions	Paragon
Chronos Technology	SyncWatch 200 SyncWatch 300
Ciena	CN3920 CN3940 CN3960 CN5305
Cisco	ANA ASR 9010 ME3600X ME3800X MWR2941
Ericsson	MINI-LINK CN 500 MINI-LINK CN 1010 MINI-LINK TN OMS 1410 SEA 10 SEA 20
Hitachi	AMN1710
Huawei	CX600-X1 CX600-X2 CX600-X3 PTN910 PTN950 PTN1900 PTN3900
Ixia	IxNetwork (software) IxN2X
MRV	OS940 OS904-DSL4 OS904-MBH OS904-MBH-4 OS906

Vendor	Participating Devices
NEC	iPASOLINK 200
Omnitron Systems	iConverter GM3
Orckit-Corrigent	CM11 CM4140 CM4206 CM4314 CMview
RAD Data Communications	ETX-204A IPmux-216 ETX-203A ACE-3220 ACE-3105 MITOP-E1T1/GE
Raisecom	iPN201
Siklu	EtherHaul1200
Spirent	Spirent TestCenter XGEM/GEM
Symmetricom	CsIII SSU 2000e TimeMonitor (software) TimeProvider 5000 TimeProvider 500
Telco Systems	T-Metro-7224 T-Metro-7124S T-Metro-200 T5C-XG T-Marc-380 T-Marc-254H
Vitesse	VTSS Caracal CE10
ZTE	ZXR10 T8000 ZXR10 M6000 ZXR10 8905E ZXR10 5928E ZXR10 5128E ZXR10 2928E ZXCTN 9008 ZXCTN 9004 ZXCTN 6100 ZXCTN 6200 ZXCTN 6300

INTEROPERABILITY TEST RESULTS

Each section of this document covering individual test areas includes the technical background of the technology, the test procedures employed, a logical topology diagram and a detailed review of the results achieved. Some successful tests which involved several vendors were incorporated to the demonstration network, described in the Network Design section below.

The multi-vendor combinations documented detail the successful results completed within the two-week hot staging time. Given the time constraint, we were able to reach fully meshed vendor combinations only for some of the technologies tested. For any missing test combinations, we encourage the reader to contact the respective vendors to ask that the test is performed at our next event.

Testers. What is a test without test equipment? Without the use of some key tools, a majority — or even all — of what was accomplished here would not have been possible. We would like to thank Ixia and Spirent for their traffic generation (user emulation), protocol emulation, and analyzing tools made available by Ixia's IxNetwork, IxOS, and IxN2X, and the Spirent TestCenter. Throughout all test areas you will find that the use of impairment tools, the Calnex Paragon and Spirent GEM, allowed us to emulate certain network conditions. Finally, our synchronization tests were only made possible by Symmetricom's Cesium atomic clock CslII, the Chronos SyncWatch 200 with TimeMonitor analyzer software and Calnex Paragon.

Terminology. For consistency, we use the term "tested" to describe interoperability tests between equipment from multiple vendors and performed according to the test plan. The term "demonstrated" will be used both when a test was performed with equipment from a single vendor and in cases when functionality was not thoroughly tested.

NETWORK DESIGN

In the first week of the hot staging, all vendor engineers and the EANTC support team focused achieving test results in appropriate, dedicated small test configurations and topologies.

In the second week of testing, we asked vendors to freeze their configurations and connections on the way to building a fully integrated end-to-end network. The end product is a snapshot of core, aggregation, and access network areas that facilitate realistic demonstration scenarios to be showcased at the Carrier Ethernet World Congress.

SYNCHRONIZATION FOR LTE MOBILE BACKHAUL

One of the 3GPP-Long Term Evolution (LTE) requirements for the LTE Backhaul is support of clock synchronisation. There are multiple ways to implement synchronisation services in a network: Packet-based and network-based methods. While network-based synchronisation relies on the

synchronous signal of the physical network layer, packet-based synchronisation is based on specific control protocols.

We tested IEEE 1588-2008, a packet-based clock synchronisation protocol, and Synchronous Ethernet which represents the network-based method. While Synchronous Ethernet provides frequency synchronisation only, IEEE 1588-2008 can provide phase or time of day synchronisation as well. The time of day synchronisation is required by LTE Multimedia Broadcast Multicast Service (MBMS), for example.

At this event, we successfully tested multi-vendor interoperability of IEEE 1588-2008 for phase synchronisation in an end-to-end network for the first time publicly worldwide.

Frequency and Time of Day Synchronisation over IEEE 1588-2008 (PTPv2)

We identified three important evaluation scenarios:

1. Precision Time Protocol (PTP) Interoperability between a PTP grandmaster and PTP slave clocks
2. PTP Transparent Clock interoperability, where a transparent clock was inserted between the grandmaster and the slave clock to improve latency awareness and slave clock stability
3. PTP and Synchronous Ethernet interoperability combined between the PTP grandmaster, a PTP slave clock and a Synchronous Ethernet slave clock. The PTP slave clock was connected to the PTP grandmaster over the packet network as before; in addition, it served as a Synchronous Ethernet master and provided clock recovered via IEEE 1588 to the Synchronous Ethernet slave.

Figure 1 shows the logical test topology. We used the Symmetricom TimeProvider 5000 as PTP grandmaster and either Spirent GEM or Calnex Paragon the as impairment tools for each of the tests. The impairment tools emulated reproducible conditions of a packet-switched network, using ITU-T G.8261 Test Case 12 impairment profiles.

Given the number of variables still undefined by the G.8261 standard in regards to how impairment profiles are recorded, it is known that not all implementations of a given G.8261 profile will be equivalent. Since this would have implications in regards to which vendors would pass or fail a test depending on what impairment was used, we validated Test Case 12 profile implementations of both Calnex and Spirent prior to the event. We are glad to report that both profile implementations resulted in comparable impairment.

PTP is a rich protocol with a range of options. In the preparation phase, we reviewed key protocol options with participating vendors, resulting in the creation of three PTP configuration profiles shown in table 1 below.

We measured frequency accuracy using the Chronos SyncWatch 200 frequency analyzer on the slaves via either 2048KHz or E1 clock output interfaces. In our third scenario with Synchronous Ethernet (SyncE), we connected the tester to the SyncE slaves. Otherwise the analyzer was connected to the PTP slaves' clock outputs.

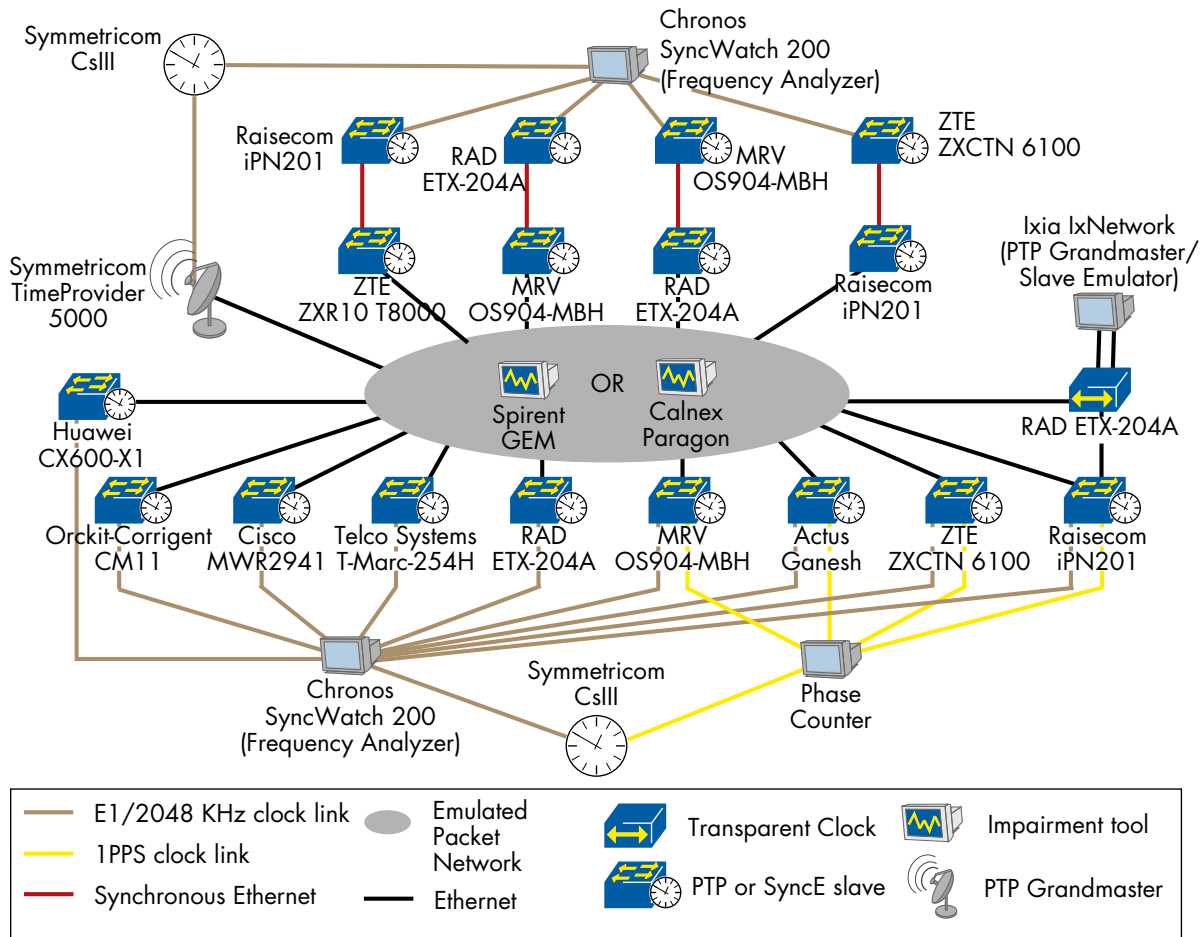


Figure 1: PTP Test Results

In the two first scenarios we also measured phase (time of day / TOD) synchronization. The support of phase synchronisation was optional for the basic PTP Interoperability (scenario 1). It was mandatory for the Transparent Clock interoperability tests (scenario 2). We used frequency counters to take phase accuracy measurements, which were connected to the slaves via 1 PPS (Pulse Per Second) interfaces.

Measurements for all three scenarios were performed for at least four hours. We started tests with the PTP slaves in a “free-running” state. We considered a PTP test passed as soon as the MTIE measurement for frequency accuracy was below the ITU-T G.823 SEC (SDH equipment slave clocks) mask and the absolute phase deviation, if tested, from the reference clock was below 10 μ s.

Figure 1 depicts the PTP slaves that successfully tested frequency synchronization (bottom half of the diagram): Actus Ganesh, Cisco MWR2941, Huawei CX600-X1, MRV OS904-MBH, Orckit-Corrigent CM11, RAD ETX-204A, Raisecom iPN201, Telco Systems T-Marc-254H, ZTE ZXCTN 6100.

The following devices were successfully tested as PTP slaves for phase (time of day) synchronization: Actus Ganesh, MRV OS904-MBH, Raisecom iPN201, and ZTE ZXCTN 6100. The tests are indicated by yellow 1PPS clock links between the PTP slave and the Phase Analyzer in the figure.

TABLE 1. PTP Profiles

Parameter	Profile1	Profile2	Profile3
Address Type	unicast	unicast	multicast
Encapsulation	IP/UDP	IP/UDP	IP/UDP
Clock Mode ^a	one-way/ two-way	one-way	two-way
Master Clock Type	1-Step or 2- Step	1-Step or 2- Step	1-Step or 2- Step
Sync/Delay_ RequestDelay_ Response rate per second	64	32	64
Announce message rate/s	0.5	0.5	0.5
Unicast Request Mechanism ^b	No	No	No
Pdelay Request/ Pdelay Response	No	No	No
Correction Field	Yes	Yes	Yes
Phase Output Interface ^c	1PPS	N/A	1PPS
Frequency Output Interface ^d	2048 KHz	E1	E1
PTP Domain	1	1	1
VLAN ID	10	20	30

- one-way clock type can be used for frequency synchronisation only
- Clause 16.1 of IEEE 1588-2008
- Slave's clock output interface type for phase measurement
- Slave's clock output interface type for frequency measurement

For the "PTP Transparent Clock Interoperability" scenario we successfully verified the RAD ETX-204A which served as the Transparent Clock and Raisecom iPN201 as the PTP slave. For this test we used an Ixia IxNetwork which emulated a PTP Grandmaster and PTP Slave and measured the Transparent Clock correction under normal conditions and conditions with unidirectional traffic load.

In the "PTP and Synchronous Ethernet" scenario we tested successfully MRV OS904-MBH, RAD ETX-204A, Raisecom iPN201, and ZTE ZXR10 T8000 as PTP slave/Synchronous Ethernet master and the MRV OS904-MBH, RAD ETX-204A, Raisecom iPN201, and ZTE ZXR10 6100 as Synchronous Ethernet clients.

In one test, an implementation could not receive Delay Response messages from the grandmaster because it used a wrong sub-domain ID in its Delay Request messages (10 instead of 1). In another situation a different implementation mixed a unicast destination IP address and a multicast Ethernet address in its PTP Delay Request messages. Apart from these hiccups, protocol interoperability was great due to the alignment of protocol implementation options in advance of the test.

Synchronous Ethernet and ESMC

Clearly, Synchronous Ethernet as a network-based synchronization mechanism is independent from packet network utilisation and packet forwarding performance of network elements. The disadvantages in comparison to the packet-based technique are that SyncE works in pure Ethernet environments only, every active network element in a clocking chain must support Synchronous Ethernet, and it provides frequency synchronisation only.

We verified SyncE interoperability by connecting a Synchronous Ethernet master device directly to a very precise reference clock provided by Symmetricom and measuring the clock output signal quality on the Synchronous Ethernet slave via a frequency analyzer. We used the Chronos SyncWatch 200 as a frequency analyzer connected to the slaves via either 2048KHz or E1 clock output interfaces; alternatively, we used the Calnex Paragon connected to slaves via a SyncE interface.

Since synchronisation quality does not depend on the network utilisation, it was not necessary to emulate a network between SyncE master and slaves. A measurement interval of at least 1 hour was deemed sufficient. We set the requirements to pass a test higher than for PTP tests: A test was considered as passed if the slave clock quality stayed below the ITU-T G.823 SSU masks.

The following systems were successfully tested as Synchronous Ethernet master: Albis Aceeed, Cisco ASR 9010, Cisco ME3600X, Huawei CX600-X2, MRV OS904-MBH, NEC iPasolink 200, Orckit-Corrigent CM4206, Orckit-Corrigent CM4314, RAD ETX-204A, Raisecom iPN201, Telco Systems T-Metro-7124S, ZTE ZXR10 5928E, ZTE ZXCTN 6300, ZTE ZXCTN 9008, ZTE ZXR10 M6000. During the test run between ZTE ZXR10 M6000 and Actus Ganesh, we observed two unexpected jumps in TIE values. We assume they were caused by

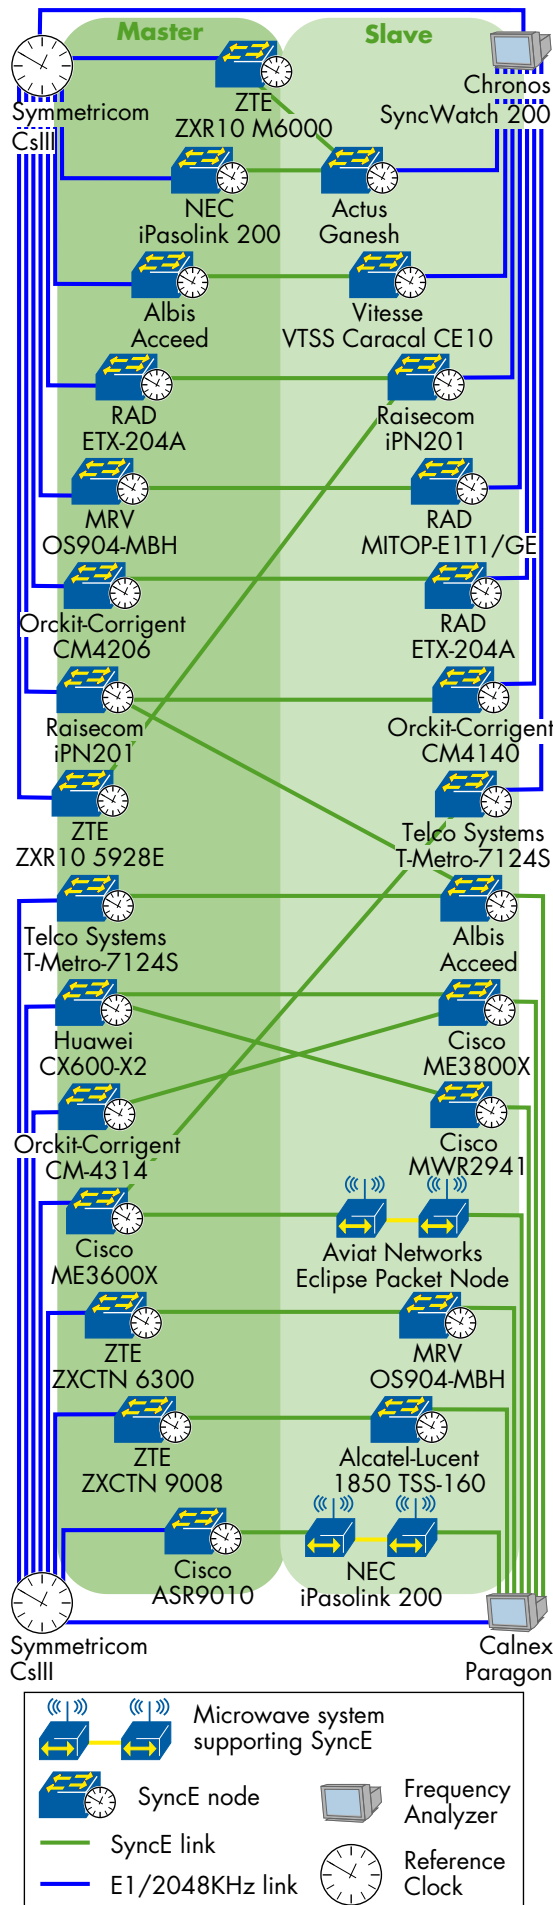


Figure 2: SyncE Test Results

electric issues in the lab, however we had no time to either retest or verify the assumption in detail.

The following systems were successfully tested as Synchronous Ethernet slave clocks: Actus Ganesh, Albis Aceeed, Alcatel-Lucent 1850 TSS-160, Aviat Eclipse Packet Node, Cisco ME3800X, Cisco MWR2941, MRV OS904-MBH, NEC iPasolink 200, Orckit-Corrigent CM4140, RAD ETX-204A, RAD MITOP-E1T1/GE, Raisecom iPN201, Telco Systems T-Metro-7124S, Vitesse VTSS Caracal CE10. Please refer to Figure 2 for results.

Ethernet Synchronization Messaging Channel (ESMC)

ESMC is a protocol defined in ITU-T G.8264 for usage in Synchronous Ethernet networks with more than one clock source. The SyncE nodes exchange Synchronization Status Messages (SSM) that distribute clock quality level information to adjacent nodes. Each SSM message contains a single QL-TLV with a QL (Quality Level) value binary-compatible with the SDH network definition (ITU-T G.781). The QL information allows the receiving SyncE node to select the best clock source available at any time.

In each of our ESMC tests we verified interoperability between three SyncE nodes. Two nodes were directly connected to different external clocks and advertised their internal quality level to the network which was related to the quality level of the external clock — one of them at primary reference clock (PRC) quality level and another at SSU quality level. A third node was connected to the first two nodes over SyncE and had no external clocks directly attached to it. This node advertised the quality level of its internal clock to the network, which was often SEC, so we call the device SEC for simplicity.

During the ESMC tests we connected, disconnected, and reconnected the external clocks on PRC and SSU devices and observed the ESMC protocol exchange. The following devices successfully participated in the ESMC tests: Actus Ganesh, Albis Aceeed, Alcatel-Lucent 1850 TSS-160, Cisco ASR 9010, Huawei CX600-X3, Ixia IxNetwork and IxN2X, MRV OS904-MBH, Orckit-Corrigent CM11 and CM4140, RAD ETX-204A, Raisecom iPN201, Telco Systems T-Metro-7124S, Vitesse VTSS Caracal CE10, ZTE ZXCTN 6300 and ZXR10 M6000. Please refer to figure 3 for documentation of each individual test combination. The tests were supported by Calnex Paragon and Ixia IxN2X as ESMC packet analyzers.

We discovered an issue with one implementation during a switchover from PRC to SSU clock: The SEC slave initially sent clock code EEC1 as expected followed by an unexpected EEC2. Later on, the vendor resolved the issue and the device sent SSU as expected. In another combination, a device was incompatible as it sent periodic ESMC messages every 2–4 seconds instead of every second as expected. In yet another case, a device under test could not lock on its external PRC clock after it had been reconnected. The problem could be solved by a software shutdown/restart of the interface to the external clock. In general, interoperability was very reassuring - SyncE ESMC support is growing.

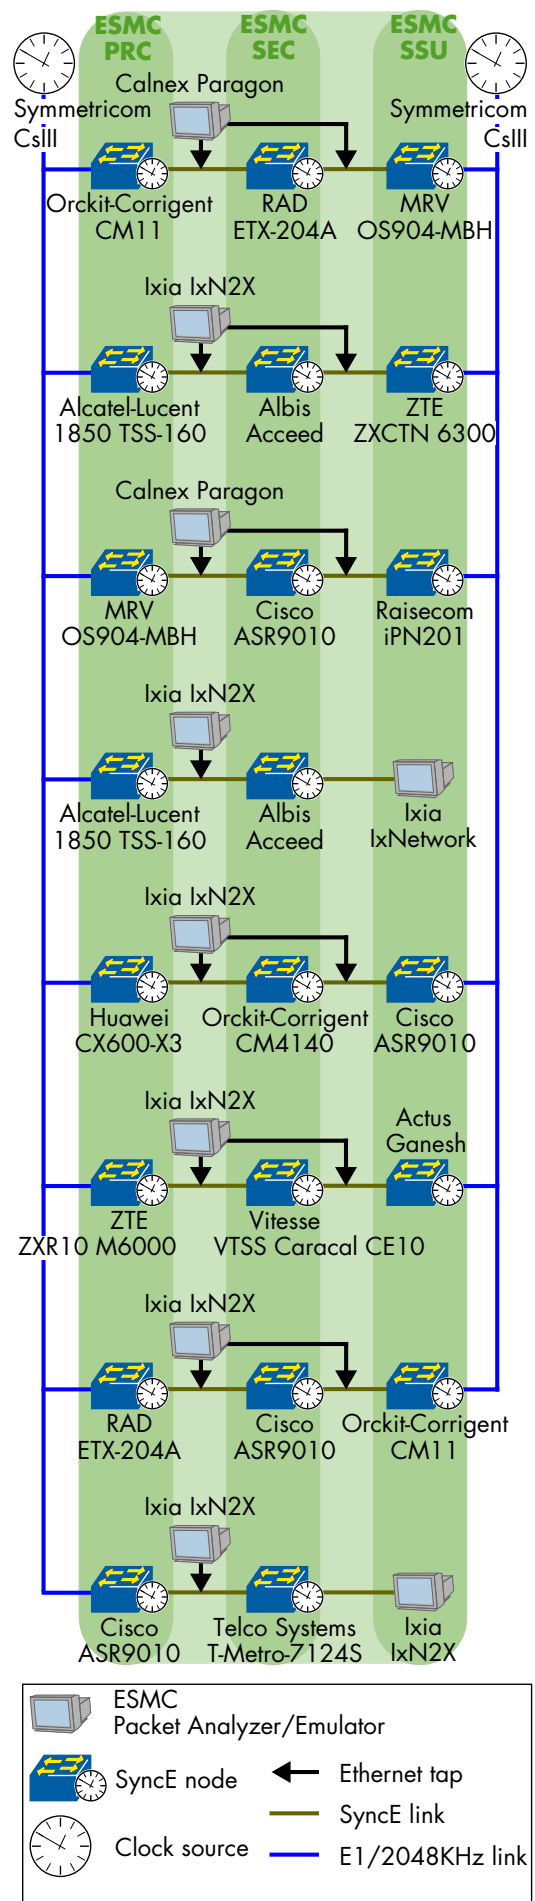


Figure 3: ESMC Results

CONVERGED TRANSPORT

Data Center Interconnection. With the streamlining of data center services gaining quick traction this year, and cloud computing, we attempted to add this focus on our transport tests. We have tested several packet based transport technologies quite a bit in the past, and the idea was to add some context. For such a concept, we thought both multipoint connectivity - for site-to-site interconnection - and jumbo frames - for performance optimization - should be required given the requirements from such data centers and their load balancing.

Interworking Between Transport Domains

The premise of this test was to ensure that services could interwork across the boundary of IP/MPLS and MPLS-TP domains. Unicast, broadcast, and multicast frames (in IMIX and 9000 Byte jumbo frame patterns) were transmitted into the VPLS clouds. Figure 4 depicts multipoint and point-to-point services were configured. In all cases a static MPLS label was used to tunnel the service between the domains.




Figure 4: IP/MPLS and MPLS-TP Interconnect

PBB-VPLS Interworking

In a Virtual Private LAN Service, the provider edge (PE) nodes need to learn customer MAC addresses in order to forward customer Ethernet traffic appro-

priately. This may often require the PE to maintain thousands of MAC addresses or more. Scaling to provider requirements was one of the motivations when the IEEE designed Provider Backbone Bridging (PBB): To create an additional abstraction layer for bridges.

We tested interworking between provider backbone bridges and VPLS domains in two cases: In the first test combination, a Cisco ASR9010 acted as a native PBB switch, connected to the Alcatel-Lucent 7750 SR7 for encapsulation and forwarding into a VPLS domain. The Alcatel-Lucent SRc12 performed both the PBB and VPLS PE functionality required for this test case. Some full-duplex traffic was passed with no loss.

In a second test scenario, an Alcatel-Lucent 7750 SR7 was configured as a VPLS PE, and an Ixia IxNetwork emulating a PBB network was attached. Ixia IxNetwork emulated VPLS/PBB devices. All traffic was passed back and forth with no loss.

Ethernet Radio Transport

Five microwave and millimetric wave systems joined this year's EANTC event to interoperate with other Carrier Ethernet systems and to run through our feature benchmarking. Such systems are important for operators to offer connectivity to places difficult or expensive to reach with wire-line services — this is particularly helpful for base station deployment.

All participating microwave devices demonstrated their adaptive modulation capabilities. Aviat started at 256-QAM with their Eclipse Packet Node, stepping down to 64-QAM with no loss in prioritized traffic. NEC showed their iPasolink 200, also going from 256-QAM to 128-QAM with no loss in prioritized traffic. Alcatel-Lucent demonstrated interoperability between their split-mount 9500 MPR and the stand-alone 9500 MPR-e, stepping from 16-QAM down to 4-QAM with no loss in prioritized traffic. Ericsson started at 512-QAM, stepping down to 128-QAM with no loss in prioritized traffic on both the Mini-Link TN and Mini-Link CN500. Siklu's EtherHaul1200 stepped from QPSK-500MHz to QPSK-250MHz with no drop in prioritized traffic.

Link State Propagation is a feature enabling communication about link state between communicating endpoints of the radio link. When the wireline Ethernet link goes down on one side of the radio, the radio tells the other side to also disable its Ethernet link. The following systems demonstrated this feature: Aviat Eclipse Packet Node, Ericsson Mini-Link TN, Ericsson Mini-Link CN500, NEC iPasolink 200, and Siklu EtherHaul1200. Alcatel-Lucent presented a different strategy, pointing out that this feature also disables links by default to the provider, limiting their ability to localize the failure. It has therefore been discussed at the ITU-T to define a "Client-Signal-Fail" within Y.1731 to allow for a specific alarm for this case.

In addition to the Synchronous Ethernet (SyncE) tests described in the next section, several SyncE chains with microwave systems included were constructed. For the following test combinations, we quickly observed that the running clock quality stayed well below the mask at a glance: The first chain involved

the Alcatel-Lucent 9500 MPR and 9500 MPR-e, Aviat Packet Eclipse Radio, Cisco ME3600X, and MRV OS904-MBH; and a second chain involved the Aviat Packet Eclipse Radio, Cisco ME3600X, Cisco MWR2941, and MRV OS904-MBH.

RAD demonstrated a bonding feature - traffic was distributed across two Ethernet links. On those two links sat two Siklu EtherHaul 1200 systems.

Additionally, Ericsson demonstrated the radio link utilization of their Mini-Link TN radio using stateful TCP traffic.

ATM Pseudowire Emulation

A test was performed between Huawei PTN 3900 and Alcatel-Lucent 1850TSS-320. ATM cells were successfully transported over MPLS-TP by using encapsulation as described in RFC 4717. We configured cell transport type 0x000A, packetization delay of 10ms, and the Maximum Number of Cells Packed (MNCP) of 5.

MANAGED ETHERNET SERVICES

Performance Monitoring

Network operators require toolsets that can monitor performance of Ethernet Virtual Connections (EVC) during operation to verify Service Level Agreements (SLAs). The ITU-T specification Y.1731 provides such a toolset by introducing techniques to measure frame loss, frame delay, and frame delay variation on point-to-point Ethernet services. On the IP layer, the IETF has defined IP-based performance measurements in RFC 5357. We evaluated both.

As a baseline reference test, we first validated performance monitoring implementations without introducing any impairment. All further tests were performed using either a Calnex Paragon or a Spirent GEM impairment tool that introduced artificial static delay, delay variations, or packet loss. For each specific type of impairment we verified the measurement results provided by the implementations under test against the emulator configuration. We expected the delta between the impaired configuration and the reference test to be equivalent to the impairment tool settings.

For delay measurement we added a constant delay of 20 milliseconds (ms). In order to test delay variation, we then configured the impairment tool to introduce packet delay of 15 ms to every second packet and 25 ms to the remaining packets — the average remained at 20 ms but with a delay variation of 10 ms.

The following devices successfully participated in the Y.1731 delay/delay variation tests: Alcatel-Lucent 7750 SR7 and 1850 TSS-320 and 7705 SAR8, Ciena CN5305 and CN3960, Cisco ASR9010, Huawei CX600-X3 and PTN3900, Ixia IxNetwork, MRV OS940, Omnicon iConverter GM3, Orckit-Corrigent CM11, RAD ETX-203A, Raisecom iPN201, Siklu EtherHaul1200, Spirent TestCenter, Telco Systems T-Marc-380 and ZTE ZXR10 8905E and ZXCTN 6200.

Cisco ASR9010 and Huawei CX600-X3 measured delay and delay variation over a Virtual Private Wire Service (VPWS) across an MPLS network. Using the Calnex Paragon we introduced delay and delay variation impairment on the MPLS encapsulated DMM and DMR messages.

Based on Y.1731 over MPLS-TP, as defined in the draft-bhh-mpls-tp-oam-y1731, the Alcatel-Lucent 1850 TSS-320 performed delay and delay variation measurement with Huawei PTN3900 and ZTE ZXCTN6200.

During the measurement one implementation exhibited an incorrect delay value when the remote peer was configured with a CCM interval of 100 ms. Oddly enough, after the remote peer changed the interval to 1 s, the implementation showed an accurate delay measurement.

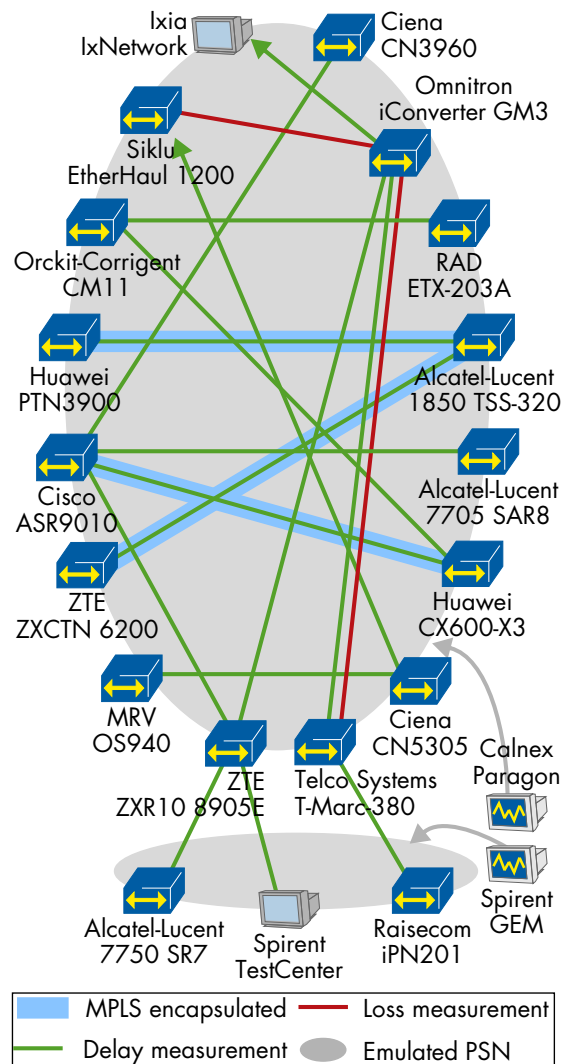


Figure 5: Performance Monitoring Y.1731

We also tested one way frame loss measurements based on Y.1731 Loss Measurement Messages and Replies (LMMs and LMRs). Using a traffic generator, we sent bidirectional Ethernet traffic over the network service and introduced 10% frame loss in one direction with the impairment tool. We verified whether the far-end and the near-end frame loss displayed on the device under test showed the same loss values. Three participants successfully tested the frame loss measurement: Omnicon iConverter GM3, Siklu EtherHaul1200, and Telco Systems T-Marc-380.

Physical Network Topology



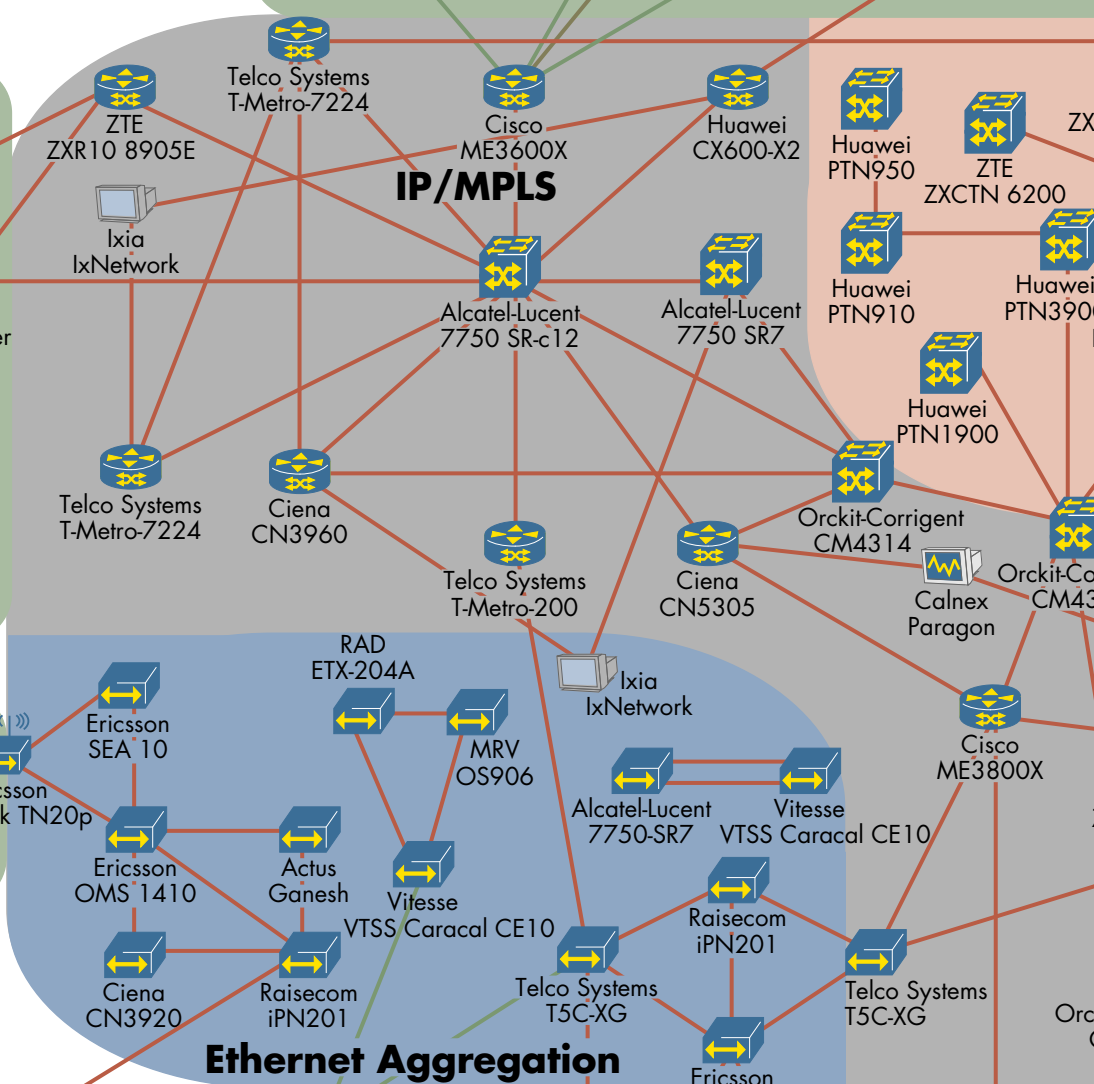
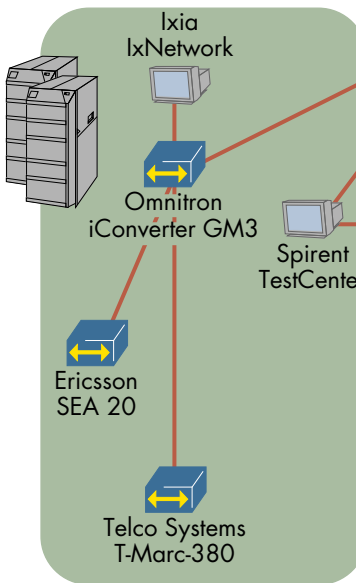
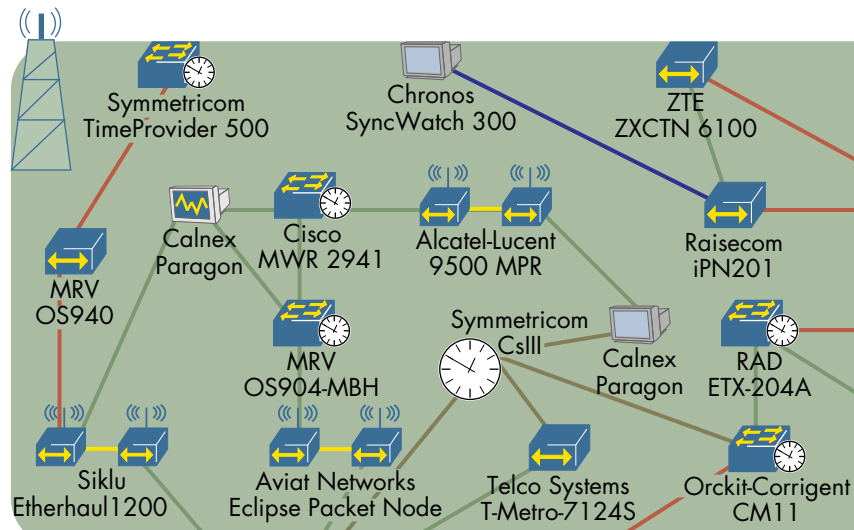
Management Systems

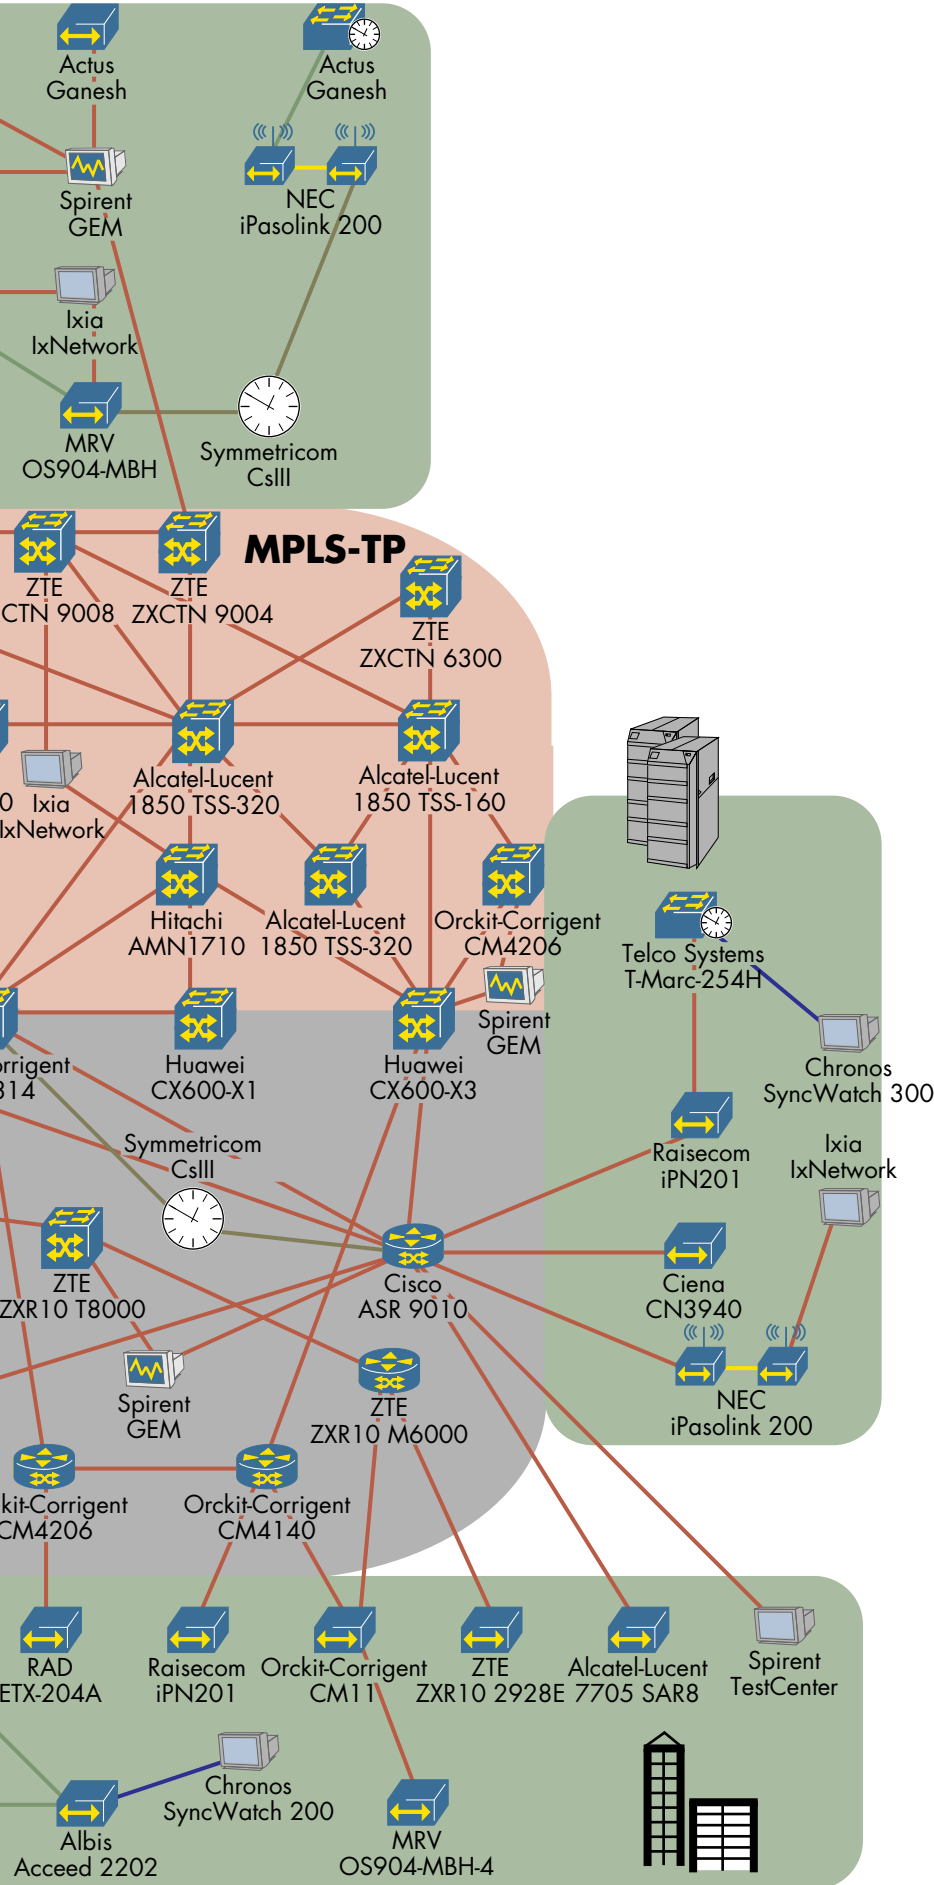


Cisco ANA













Orckit-Corrigent CMview
















Device Types

-  Provider Edge device
-  Provider router
-  Access device
-  IEEE1588-2008 client
-  Analyzer/Emulator
-  Impairment tool
-  Reference Clock
-  Management system
-  IEEE1588-2008 Grandmaster
-  Microwave radio system

Connection Types

-  Ethernet link
-  Synchronous Ethernet
-  Microwave Radio link
-  TDM link
-  Clock link

Network Areas

-  IP/MPLS network
-  MPLS-TP network
-  Ethernet Aggregation network
-  Radio access
-  Data center
-  Business access

In addition the Ericsson SEA 20 and the Ericsson SEA 10 demonstrated frame loss, frame delay and frame delay variation measurement.

Tests of the IP-based TWAMP protocol were performed between Ixia IxNetwork and Ciena CN3960 using Calnex Paragon as impairment device. Since both DUTs supported the SERVWAIT defined by RFC 5357 to last 900 seconds, the test was performed over a few hours with at least 15 minutes spent on each step.

We verified that both DUTs were able to show the baseline delay. The DUTs tested delay measurement in the direction from the IxNetwork to the CN3960. The test in the other direction had started but unfortunately not completed simply due to time constraints. Based on the configuration using the Ixia IxNetwork as the TWAMP Server and the Ciena CN3960 as the Session-Receiver, we also tested delay measurement based on Y.1731.

Link OAM

Link OAM is a term that the MEF uses to refer to implementations of OAM specified for Ethernet in the First Mile (EFM), as defined in IEEE 802.3ah-2004. Link OAM describes multiple features, our tests verified the following Link OAM features:

- Loopback operation mode, which is meant to assist carriers in performing fault localization and testing link performance.
- Link events, that signal symbol and frame errors in terms of either the number of errors or a seconds - a specified time period with errors.
- Dying Gasp, which is a critical link event, that should be generated, when a station is about to reboot or shutdown. Operators can use this indication arriving from an Access Device to localize the reason a customer connection is currently inactive.

During the loopback mode tests we verified that a DUT port operating in active mode was able to bring a remote port into loopback mode. The port operating in the loopback mode was expected to discard all incoming frames destined towards the active mode port and loopback all frames coming from the active port. The active port was expected to drop frames which were received from the remote port operating in a loopback mode, after it has updated its port statistics. The following devices participated in the loopback test: Albis Acceed 2202, Albis Acceed 1416, Alcatel-Lucent 7750 SR7, Alcatel-Lucent 7750 SR-c12, Cisco ME3600X, Ericsson SEA 20, Ericsson OMS1410, Huawei CX600-X2, Ixia IxN2X, Ixia IxNetwork, Omnitron iConverter GM3, Orckit CM11, RAD ETX-204A, Raisecom iPN201, Vitesse VTSS Caracal CE10, ZTE ZXR10 5128E, ZTE ZXR10 8905E, and ZTE ZXCTN 9008.

During the test we observed the following issues: one vendor device did not discard the loopback traffic coming from the peer device port that was working in the loopback mode; two other devices (each in a different test pair) when operating in loopback mode

did not discard traffic arriving on the other interfaces (destined toward the active port / remote device, from the device in loopback mode). The traffic was unexpectedly forwarded together with the loopback traffic to the remote DUT.

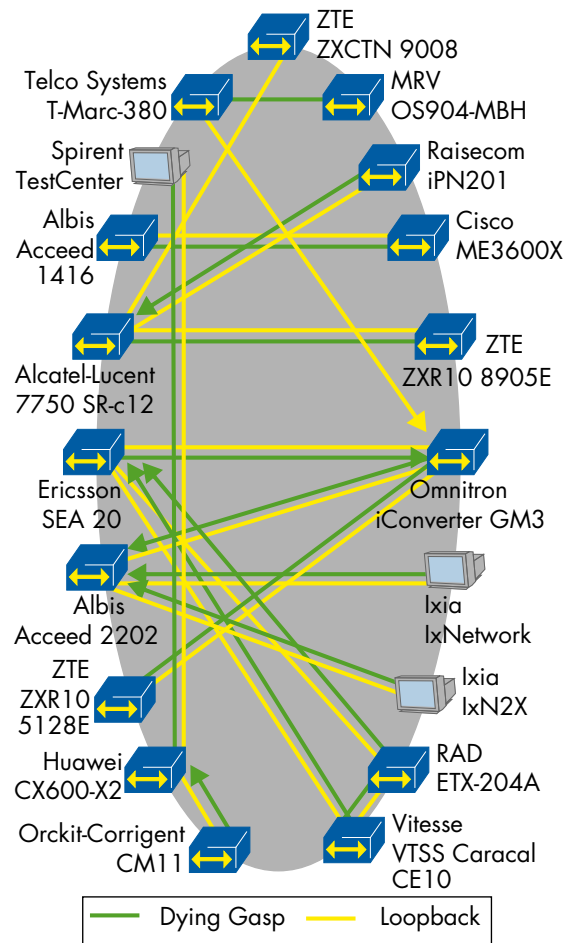


Figure 5: Link OAM - Dying Gasp and Loopback

We verified the ability of the DUTs to generate the following link events: the number of frame errors, and the number of errored frame seconds within a specified time period. We verified the DUT behavior by using CLI or a management system. Each device was tested using three thresholds: 10 errored frames in 10 seconds, 10 errored frames 64 Byte frames at 1 Gbit/s, and 5 errored frame seconds in one minute. In order to test link events we introduced errors via either Calnex Paragon or Spirent GEM impairment tools on the link between the DUTs. First, a traffic generator sent traffic filling the window of the peered device. Then the impairment tool introduced CRC errors into the traffic for five seconds at two CRC errors per second. The device receiving traffic was then expected to send one Errored Frame Event, one Errored Frame Event, one Errored Frame Period Event and one Errored Frame Second Event with a TLV value of 10 Errored Frames and 5 Errored Frame Seconds. As shown in the diagram, seven vendors passed the test: Ciena CN5305, Ixia IxNetwork, Huawei CX600-X3, Omnitron iConverter GM3, Spirent TestCenter, Telco Systems T-Marc-380, Vitesse VTSS Caracal-1 CE10 and ZTE ZXR10 5128E.

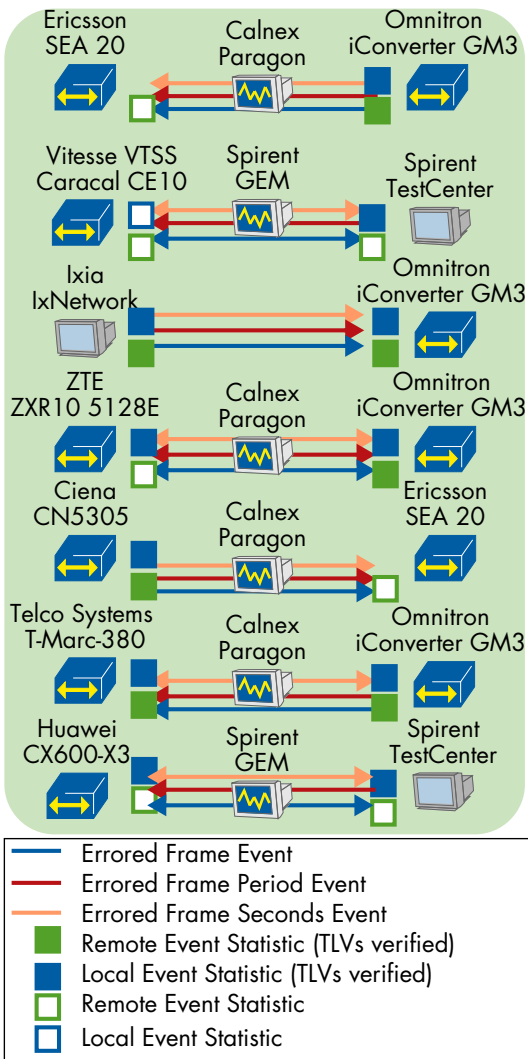


Figure 6: Link OAM - Events

The following devices successfully tested the signaling of critical link events (Dying Gasp): Albis Aceeed 1416, Alcatel-Lucent 7750 SR-c12, Cisco ME3600X, Huawei CX600-X2, MRV OS904-MBH, Omnitron iConverter GM3, RAD ETX-204A, Telco Systems T-Marc-380, Vitesse VTSS Caracal CE10, ZTE ZXR10 5128E and ZTE ZXR10 8905E. The majority of tests were performed by powering down of the device. The Alcatel-Lucent 7750 SR-c12 verified the test by disabling the EFM session. The Cisco ME3600X and the Huawei CX600-X2 performed the test by rebooting the device via CLI.

Connectivity Fault Management

Carrier Ethernet services are often built over multiple networks belonging to different administrative entities. As part of the hierarchical model provided by Connectivity Fault Management (CFM), defined in IEEE 802.1ag, the Maintenance Domain (MD) and MD levels are key components that allow monitoring of a single service within different administrative domains. The entity of a higher MD level generally has no access to the connectivity information of the lower level, in order to provide clean separation of responsibilities and prevent customers from learning the internal structure of operator networks without consent of the operator.

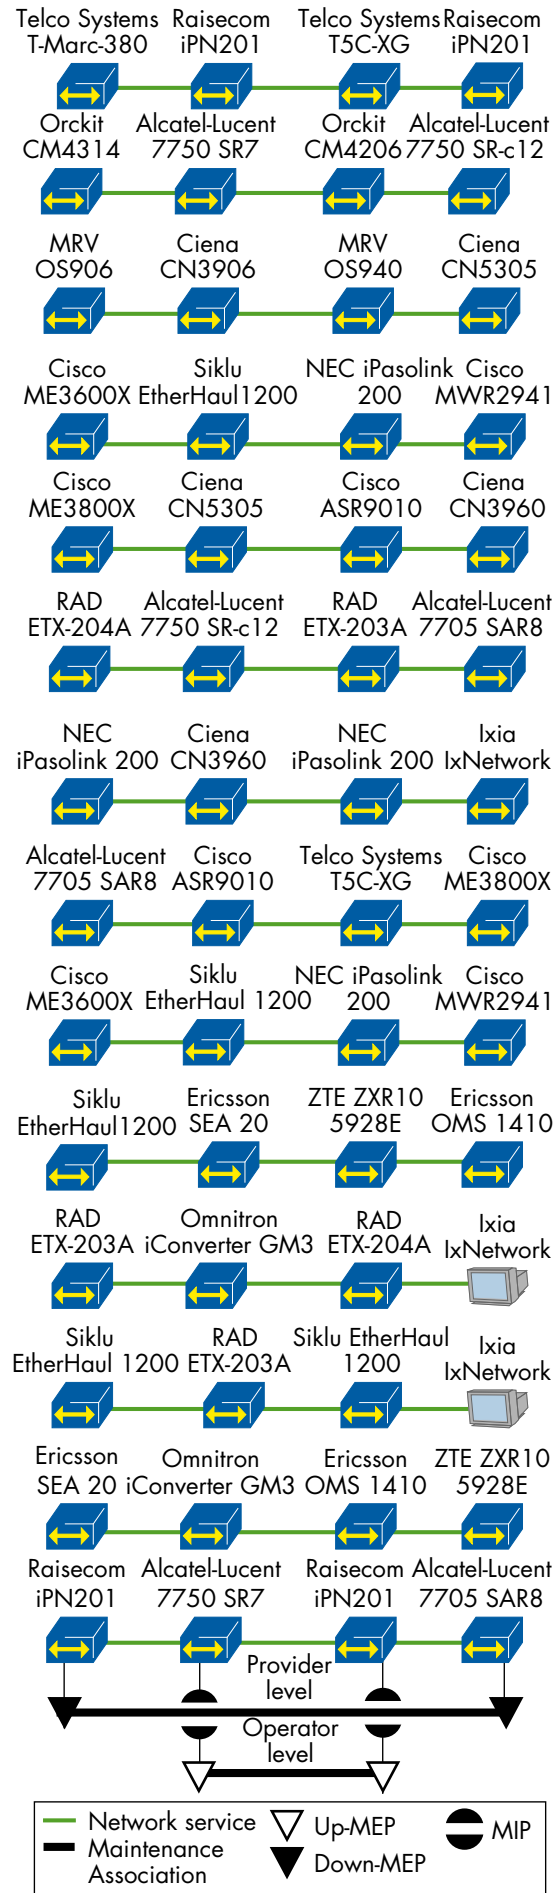


Figure 7: Connectivity Fault Management IEEE 802.1ag

In many cases, especially when performing path discovery and fault isolation, it is however useful to provide more fine-grained continuity information to the higher entity. This information is provided by configuring a Maintenance Association Intermediate Point (MIP) on the higher level and at the position of lower level's MEP; loss of connectivity between a pair of MIPs corresponds to a MEP connectivity failure at a lower MD Level. The MIPs present on a certain MD level can be used to perform the Linktrace procedure which is similar to the Traceroute tool known from IP.

In this test we verified the CFM Linktrace function of the MIPs, configured to be accessible for administrative entities situated on the Provider MD Level. The lower Operator MD level was configured between two MEPs situated at the position of the Provider MD Level's MIPs. In each MD Level the Linktrace Messages (LTMs) were transmitted by every MEP in each direction both to the pair MEP at that MD level, and also to all MIPs configured at that MD Level.

Thirteen vendors with a total of 25 devices participated successfully in the test: Alcatel-Lucent 7750 SR7, Alcatel-Lucent 7750 SR-c12, Alcatel-Lucent 7705 SAR8, Alcatel-Lucent 7750 SR-c12, Ciena CN3960, Ciena CN5305, Cisco ME3600X, Cisco ME3800X, Cisco ASR9010, Cisco MWR2941, Ericsson SEA 20, Ericsson OMS1410, Ixia IxNetwork, MRV OS906, MRV OS940, NEC iPasalink 200, Omnitron iConverter GM3, Orckit CM4314, Orckit CM4206, RAD ETX-203A, RAD ETX-204A, Raisecom iPN201, Siklu EtherHaul1200, Telco Systems T5C-XG and ZTE ZXR10 5928E.

At one point we recognized that an LTM destined to one vendors MIP was not terminated, when a device further upstream responded to the LTM.

Cisco demonstrated the use of their ANA management tool to run through the CFM and Y.1731 procedure. Some network nodes were automatically discovered (Alcatel-Lucent 7705 SR-c12, Ciena CN5305). Using the tool, they initiated link trace messages from their ASR9010, as well as Siklu's EtherHaul1200. Additionally, Y.1731 delay measurements were taken from the ASR9010 and the Ciena CN5305, where Ciena CFM error message traps were used for the reporting.

Service Failure

There are a number of heartbeat-like verification tools defined for different technologies. Such tools have the feature in common where they verify the liveliness of network services, which can be used by the service provider or the network operator to recognize failures on the service path. As soon as a heartbeat-like tool recognizes a failure, ping-like tools are used in order to localize failures. Support of such tools is a typical requirement for the Provider Edge devices.

The participating vendors intended to configure two types of network services: the intra- and the inter-domain network services. The liveliness of these network services were verified by the "heartbeat"-like OAM session configured on the DUTs.

We focused the test based on the inter-domain services which were transported over the virtual

tunnels, such as the MPLS tunnels or the Q-in-Q (inner Ethertype: 0x8100/ other Ethertype: 0x88a8 or 0x8100 for both Ethernets) encapsulation. In case of intra-domain service the network service was transported via a single C-VLAN (Ethertype: 0x8100) configured between two DUTs.

We verified three "heartbeat"-like OAM protocols during the test: the CFM Continuity Check (CC CFM), the LSP BFD and the VCCV BFD. We used the impairment devices Calnex Paragon and the Spirent GEM for all tests.

As a first step we performed a baseline test without any impairment to verify that the proper configuration of the OAM protocols. In the next step using the impairment tool we introduced service failure by dropping all frames carrying the C-VLAN-ID associated with the network service in one direction, and verified that the failure was reported in the direction in which the impairment was introduced, the Remote Defect Indication (RDI) was propagated in the opposite direction. In addition, we verified that the ping-like OAM protocol requests were not responded by the remote peer indicating the service failure between both DUTs.

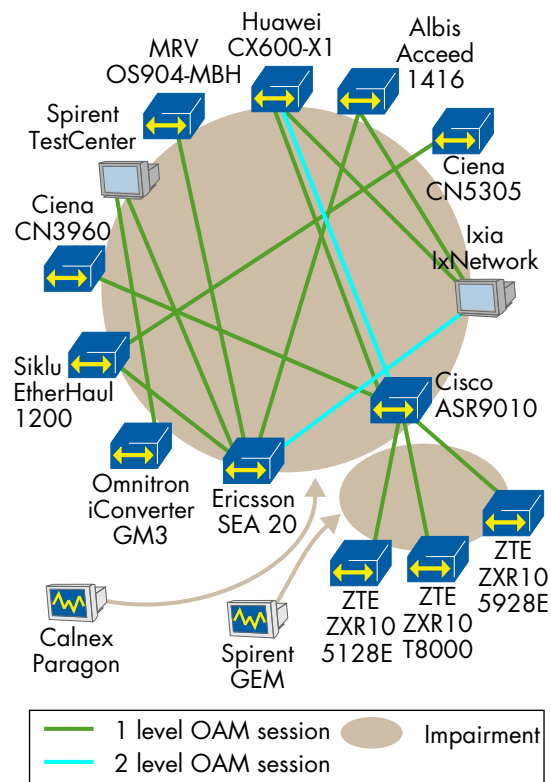


Figure 8: Service Failure (Inter-Domain Network Services)

In case of the inter-domain service we verified the liveliness of the tunnel when the DUT supported the second "heartbeat"-like OAM session for the tunnel. When performing the service failure we verified that the tunnel was up and the OAM ping went through over the tunnel. Then, we introduced tunnel failure by dropping all frames associated with a tunnel ID such like S-VLAN-ID or the MPLS label, and verified that the failure was reported at the tunnel level and at the service level. In this case the OAM ping was not responded at both levels by the remote peer.

The following devices successfully tested the service level failure using inter-domain network services:

Albis Aceeed 1416, Albis Aceeed 1416, Ciena CN3960 and CN5305, Cisco ASR9010, Ericsson SEA 20, Huawei CX600-X1, Ixia XM12, Ixia IxNetwork, MRV OS940, Omnitron iConverter GM3, Siklu EtherHaul1200, Spirent TestCenter, ZTE ZXR10 5928E, ZXR10 T8000 and ZXR10 5128E.

The Cisco ASR9010 and Huawei CX600-X1 tested the two level service failure using CFM over an MPLS Pseudowire and MPLS OAM (LSP Ping/Traceroute). The Ixia IxNetwork and the Ericsson SEA 20 tested the two level service failure using network service with Q-in-Q (0x8100/0x88a8) encapsulation.

In addition, the following devices successfully tested the service level failure using intra-domain services: Albis Aceeed 1416, Alcatel-Lucent 7750 SR7, Ciena CN5305, Cisco ASR9010, Ericsson SEA 20, Huawei CX600-X1, Ixia IxNetwork, MRV OS904-MBH, Omnitron iConverter GM3, Orckit-Corrigent CM4314, Orckit-Corrigent CM11, Spirent TestCenter, Telco Systems T-Marc-380, ZTE ZXR10 T8000, ZTE ZXR10 8905E and ZTE ZXR10 M6000.

We discovered an interoperability issue: Most vendors supported MD names as string, two vendors implemented MD name using No Maintenance Domain Name present. As a consequence, the OAM discovery failed in some cases.

Failure Propagation

Finally we verified two scenarios of OAM Failure Propagation. In one scenario (CFM level interworking) we verified propagation of failure notifications between different CFM levels. In the other scenario (OAM protocol interworking) we verified OAM Failure Propagation between MPLS and CFM OAM protocols.

The Alcatel-Lucent 1850 TSS-320, Alcatel-Lucent 1850 TSS-160, ZTE ZXCTN9004 and ZTE ZXCTN9008 successfully tested CFM level interworking. The failure propagation was performed over MPLS-TP based on the draft-bhh-mpls-tp-oam-y1731 using AIS messages. Huawei CX600-X2, Huawei CX600-X1 and Cisco ASR9010 successfully tested OAM protocol interworking. Between Huawei CX600-X2 and Huawei CX600-X1 MPLS VCCV BFD was configured. Huawei CX600-X1 and Cisco ASR9010 configured CC CFM at MD level 5. Huawei CX600-X1 performed failure propagation between CC CFM and MPLS VCCV BFD.

LLDP

Link Layer Discovery Protocol (LLDP) is used to discover information about neighbors on a network node. The node can obtain information such as Time To Live (TTL), Port Identification (Port ID), and Chassis Identification (Chassis ID) about its neighbors by implementing LLDP.

We identified several different implementations of LLDP. Specifically the Port ID and Chassis ID fields had a range of outputs, be that a string, Media Access Control (MAC) address, or a hexadecimal number. Also in several scenarios vendors did not display their neighbor's TTL field.

There were 9 vendors who participated and a total of 18 different devices. The following devices successfully participated in tests for LLDPs excluding the Chassis ID change: Ciena CN3960 with MRV OS906, Orckit-Corrigent CM4314; Cisco ME3800X with Alcatel-Lucent 7750 SR7, Ciena CN5305, MRV OS906, and Telco Systems T5C-XG; Ciena CN 5305 with Spirent TestCenter; Telco Systems T-Marc-380 with MRV OS904-DSL4, and Raisecom iPN201 with Ciena 3920, ZTE ZXR10 5128E, ZTE ZXR10 5928E, ZTE ZXCTN9004 and ZTE ZXR10 T8000. Spirent TestCenter tested the changing of their Chassis ID.

CARRIER ETHERNET RESILIENCY

Ethernet Ring Protection Switching (ERPS)

As network operators deploy Carrier Ethernet services, they are falling a bit short on resiliency mechanisms for their access networks. Spanning Tree is the old go-to, but many operators are weary of using this enterprise protocol for their services. Ethernet Ring protection Switching (ERPS), defined in ITU-T Recommendation G.8032, provides a means of achieving such resiliency. The original version of the standard (commonly referred to as "version 1") supports resilient ring topologies. The more recent version, updated this year and referred to as "version 2" introduces a range of additional features like ring interconnection, and administrative Automatic Protection Switching (APS) commands like Force Switch and Manual Switch for manual control of the network.

After the "version 1" rings were properly configured with the R-APS VLAN and service VLAN in all scenarios, bidirectional traffic flows were sent between the two ring nodes. We successfully verified in all scenarios that the RPL owner unblocked his RPL port upon physical link failure between two ring nodes. The observed failover times ranged from 2 – 55 ms, and the restoration time ranged from 6 – 21 ms. There were no major issues. Three multi-vendors rings based on version 1 were tested: Telco Systems T5C-XG, Raisecom iPN201 and Ericsson OMS 1410; RAD ETX-204A, Vitesse Caracal CE10 and MRV OS906; Vitesse VTSS Caracal CE10, MVR OS906 and Ciena CN3920. The RPL owner of each ring was TelcoSystems T5C-XG, MRV OS906 and Ciena CN3920 respectively. The Ericsson OMS 1410, implementing "version 2" of the standard demonstrated interoperability between the first "version 1" ring — a new feature of the standard, currently under development by ITU-T from SG15 Q9.

The following systems took part in tests of "version 2": Telco Systems T5C-XG, Raisecom iPN201, Ericsson OMS 1410, Actus Ganesh and Ciena CN3920. Figure 9 depicts the two topologies which were tested, and reflects how in each case there were two interconnecting nodes between the sub-rings. Failover times did not exceed 35.5 ms and the restoration time from the Telco Systems ring was 4.5 ms. In each scenario, we also verified that a failure in the one sub-ring didn't affect the operation of the

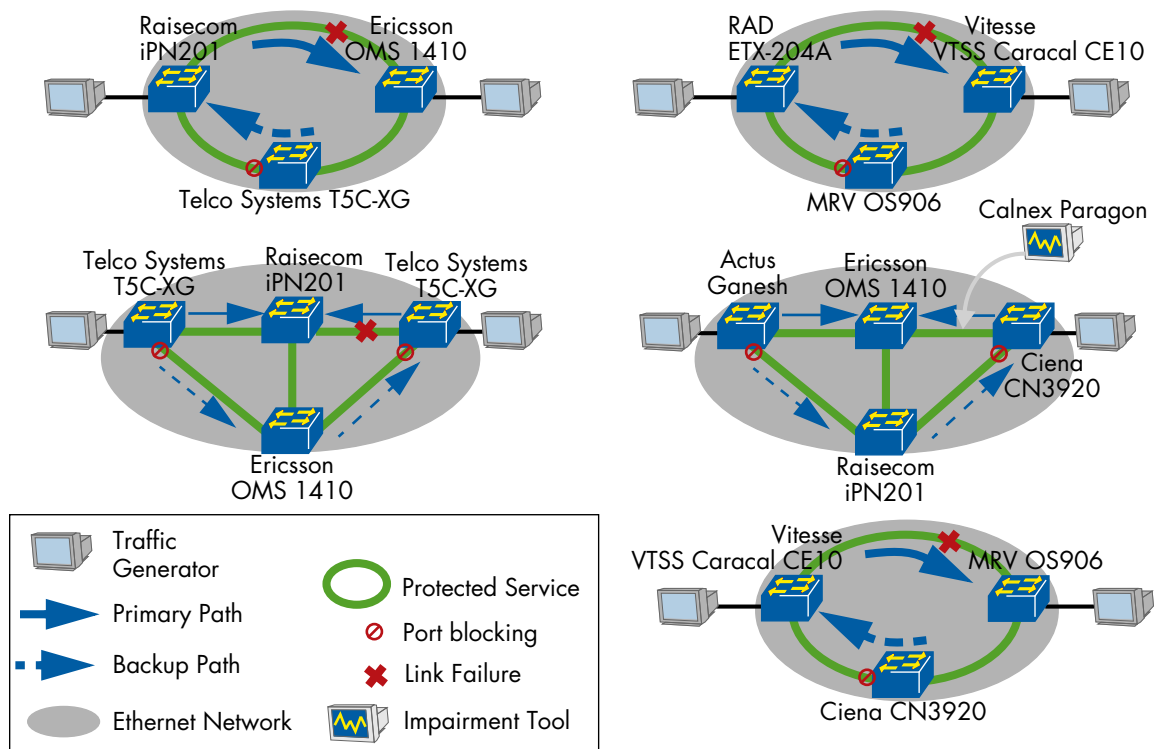


Figure 9: ERPS Test Results

other sub-ring.

We also successfully tested administrative commands in both rings — manual and force switch — to move the port block from Ring Protection Link (RPL) port to a port on a different ring link. The measured force switch times were 6.5 ms and 2 ms and the measured manual switch times were 2.5 ms and 6.5 ms. Manual and force switch command was also successfully removed by a Clear command in all scenarios.

When moving to the “version 2” tests, we initially encountered a difficulty of finding which multicast MAC to be used in the sub-ring for sending R-APS because of some ambiguity in the standard about interpreting the last byte of the MAC addresses used for G.8032 R-APS messages communication — the standard mandates the use 01-19-A7-00-00-[Ring ID] on one hand and to use 01-19-A7-00-00-01 on the other hand. Naturally, different vendors interpreted this differently. Vendors have taken two different approaches in implementing the ETHDi/ETH_A function that extracts and generates the R-APS messages. Some require a MEP to be configured and running CCM, others do not. This can lead to interworking problems based on CCM behavior: In particular we observed that some vendors were not able to configure their CCM interval below 1 s (as soon as the CCM was configured below this value the CCM session failed). Moreover, some vendors could not function as the interconnection device, though they claimed support for other features in “version 2”.

Currently ITU-T SG15 Q9 is developing version 3 of G.8032, where the topic of connecting non-ERP nodes into an ERP network topology is being addressed. To demonstrate this, Ericsson connected their Mini-Link TN into a ring built using an Actus Ganesh, Ericsson OMS 1410, Ericsson SEA 10. Failover times were measured as we removed the

link between two ring nodes (Ericsson OMS 1410 and Ericsson SEA 10).

Ethernet Linear Protection Switching

Ethernet Linear Protection Switching (ELPS) allows carrier networks to quickly recover from failure to ensure high reliability and network survivability.

ITU-T G.8031 defines the specification of ELPS to protect point-to-point ethernet service paths by using disjoint protection paths. Automatic Protection Switching (APS) is also defined in the same standard to allow both head-end and tail-end of the protected domain to co-ordinate and switchover the protection path in the failure event.

We accomplished this test in two scenarios using 1:1 bidirectional architecture and revertive mechanism. Both head-ends and tail-ends were using APS to coordinate the switchover in the case of a link failure event. We tested Ethernet Linear Protection among Alcatel-Lucent 7750-SR7, Vitesse VTSS Caracal, RAD ETX-204A and Raisingcom iPN201 in a 1:1 architecture with revertive mode enabled.

The first test scenario was performed between Alcatel-Lucent 7750-SR7 and Vitesse VTSS Caracal using Loss Of Signal as trigger; the second scenario was between RAD ETX-204A and Raisingcom iPN201 using impaired CCMs as a trigger. CCMs were sent at 100 ms intervals to check the liveness of the link. For both tests primary and secondary paths were configured. Each path was configured with different VLAN IDs for control traffic. After breaking down the link on the primary, traffic was redirected to the secondary path. Neither failover time nor restoration exceeded 50 milliseconds.

We also successfully tested the APS command including Lockout of protection, force switch normal traffic-to-protection and Manual switch normal traffic-to-protection. After sending Lockout of protection and breaking down the link on the primary path 100% loss was observed in both scenarios. By sending Force switch normal traffic-to-protection and Manual switch normal traffic-to-protection from both DUT and breaking down the link on the secondary path in both scenarios the expected behavior was observed, the recorded switchover time ranged between 3 ms and 300 ms.

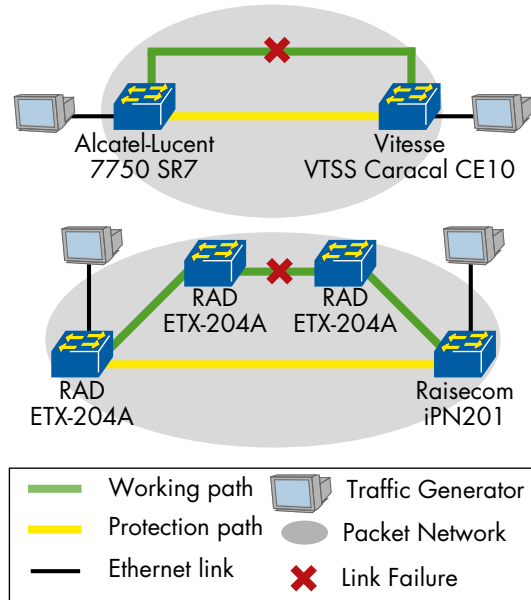


Figure 10: ELPS Results

In another vendor pair, we discovered a section of the Y.1731 specification which was interpreted in multiple ways. Section 9.1.1 talks about the MEG ID Byte set and references to appendix A. In appendix A it is stated the UMC ID SHALL consist of 7-12 characters with trailing NULLs. This was interpreted in two ways: a) There should only be 12 characters and if there was less use NULL as padding in the last character; b) There is an additional character after the UMC ID.

Connectivity Redundancy using Continuity Check

In a relatively basic test, we verified that IEEE 802.1ag defined CFM could be used to detect failure on a specific service on a given link, while a service on the same link was still healthy.

Two services were established between DUT1 and DUT2 using VLANs — two services, each with a different VLAN ID. CFM with a period of 3.33 ms was configured for both services to monitor their liveness. We first sent traffic on working path and verify that the traffic was indeed using the working path by capturing frames and checking the VLAN ID. We then used the impairment tools to drop all CCM packets on the working path for a single service. Traffic was then expected to move from the primary to the secondary path and we measured the failover time. After restoring the CCM connectivity, the traffic was returned back to the working path. This demonstration was conducted between two

Raisecom iPN201 nodes. The Alcatel-Lucent 7750-SR7 and Cisco ME3800X performed this test using the CFM CC protocol to drive interface state and trigger IP Routing reconvergence in the event of failure of the most preferred route configured with the best administrative distance.

MPLS-TP 1:1 LSP Protection

The telecommunications industry has watched the ITU-T and IETF discuss and define the MPLS Transport Profile (MPLS-TP) with great interest, and seems ready to come to some conclusions. While there may not be complete consensus regarding the direction of the technology, most agree that MPLS-TP is a technology that operators can compare to their transmission networks — currently based on SDH/SONET lines. Resiliency is surely a key aspect, as is OAM, where the past two years of discussions have been heavily focused. Most recently, one of the author drafts for a Bidirectional Forwarding Detection (BFD) based OAM titled «Proactive Connection Verification, Continuity Check and Remote Defect Indication for MPLS Transport Profile» has been accepted by the IETF MPLS working group. In parallel, a series of vendors registered to the interop event ready to test their OAM solutions based on ITU-T Recommendation Y.1731.

Members of our service provider review panel confirmed their interest of a test of all available draft or pre-draft MPLS-TP OAM options to explore in how far current vendor solutions can satisfy operators' requirements for superior protection, fault and performance management in the MPLS-TP space already. The carriers confirmed they are actively watching the MPLS-TP standardization, hoping for a fast and comprehensive implementation of latest drafts to progress the industry. While it would be desirable to converge to a single protocol option in the end, the viability of all options on the table should be explored at this point.

The channel over which OAM should be transmitted has been standardized for some time now by RFC 5586, "MPLS Generic Associated Channel". The channel specifies how OAM frames can be detected based on MPLS labels, as opposed to MAC or IP addresses. The following devices established Label Switched Paths (LSPs) and exchanged OAM messages over this channel: Alcatel-Lucent 1850 TSS-160, Alcatel-Lucent 1850 TSS-320, Huawei CX600-X3, Hitachi AMN1710, Huawei PTN910, Huawei PTN950, Huawei PTN3900, Ixia IxNetwork, Orckit-Corrigent CM4314, Orckit-Corrigent CM4206, and ZTE ZXCTN6300.

The next step was to verify that the loss of OAM frames would cause the network to switch service frames over to a backup LSP. Two methodologies were used to break connectivity — a physical link break — performed between intermediate nodes so that the end device would not experience a link failure — and an impairment of continuity in a single direction. The following devices successfully participated in the protection test: Alcatel-Lucent 1850TSS-320, Huawei PTN3900, Huawei CX600-X3, Orckit-Corrigent CM4206, Orckit-Corrigent CM4314, and

ZTE ZXCTN6300.

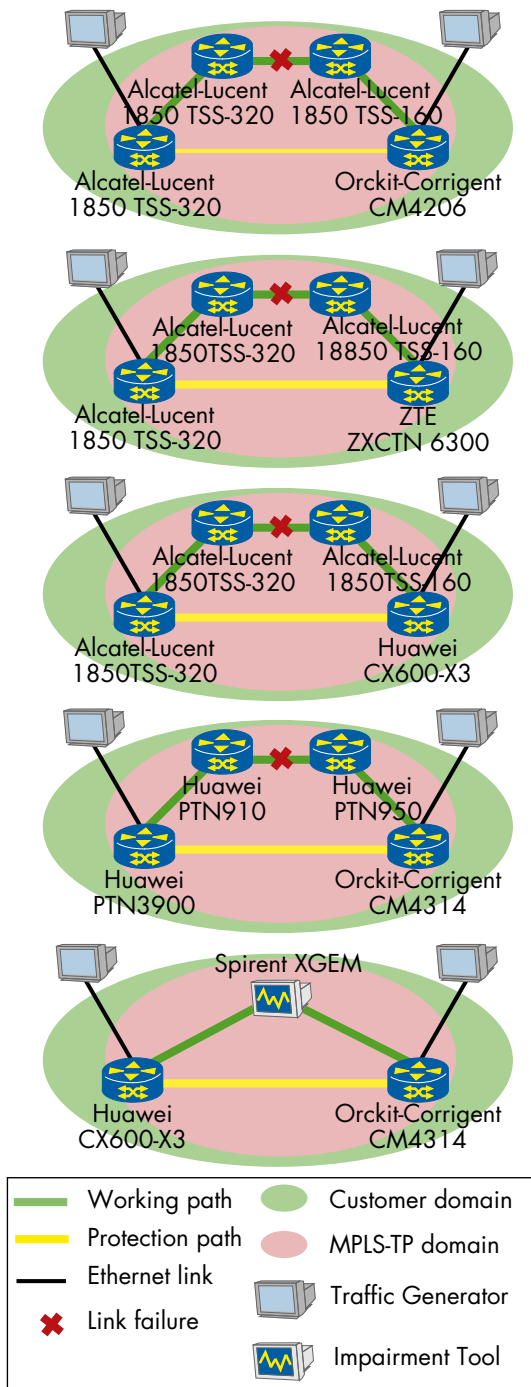


Figure 11: MPLS-TP 1:1 LSP Protection

In order to perform this test several multi-vendors pairs were built (see diagram). The observed failover times ranged between 13 to 28 ms for link failure. The OAM tools from these vendors was based on draft-bhh-mpls-tp-oam-y1731 and the linear protection was tested according to draft-zulr-mpls-tp-linear-protection-switching-01. Besides the protection switching due to link failure, APS commands "force switch normal traffic signal-to-protection" and "manual switch normal signal-to-protection" were successfully tested between all the above mentioned vendor's devices. However in one scenario when the "Lockout of protection" command was tested, we observed 100% loss in one direction and 24% in other direction between Huawei CX600-X3 and Orckit-Corrigent CM4206, 100% loss was observed

between other vendors pair. Finally, we also successfully tested failure based on test parameter mismatch including MEP_IG and MEG_ID and observed that traffic was moved to protection path. Hitachi AMN1710 and Ixia IxNetwork, supporting BFD based OAM according to draft-ietf-mpls-tp-cc-cv-rdi, also interoperated at an OAM level, bringing up an MPLS-TP pseudowire and BFD session between the two.

MPLS Pseudowire Redundancy and H-VPLS Dual Homing

MPLS pseudowire redundancy and multi-homed Multi Tenant Unit switches (MTU-s) in H-VPLS can each be used to protect different parts of an MPLS network - an attachment circuit failure, and a spoke-PW failure, respectively.

Pseudowire Redundancy is being standardized by the IETF in two drafts: draft-ietf-pwe3-redundancy which defines a framework and architecture of pseudowire redundancy, and draft-ietf-pwe3-redundancy-bit, which describes a mechanism for standby signaling of a redundant pseudowire.

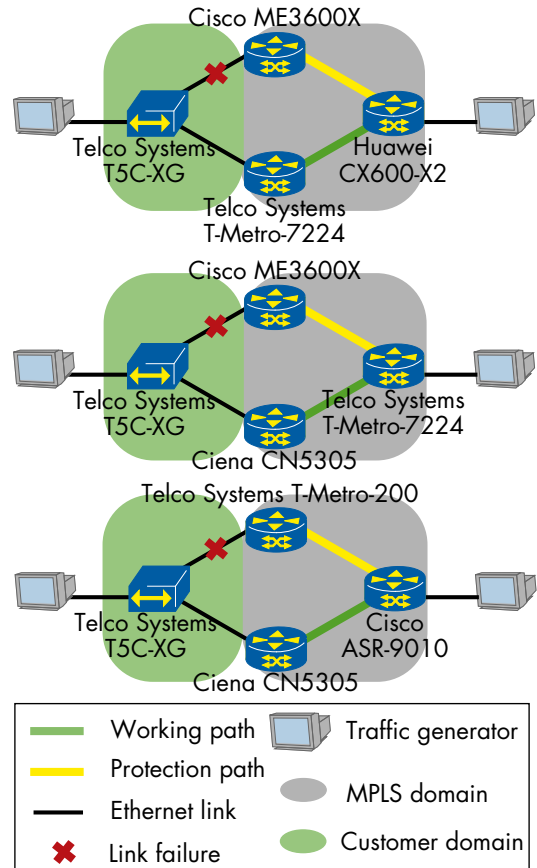


Figure 12: Pseudowire Redundancy

In all test scenarios a Telco Systems T5C-XG was dual-homed using two attachment circuits each to a separate PE, and each configured with a pseudowire terminated at a common "switchover PE". Both pseudowires were operationally up but the preferential forwarding status of one of the pseudowire was active. In all test scenarios the failure of the attachment circuit triggered pseudowire switchover on the "switchover PE" and traffic was redirected to the secondary pseudowire as

expected. The failover times was ranged everywhere from 3 ms to 3400 ms depending on which vendor's equipment was the "switchover PE" doing the switchover. The Telco Systems T-Metro-7724, Huawei CX600-2 and Cisco ASR9010 were each used as the «switchover PE». The Cisco ME3600X and Ciena 5305 were used as the near end PEs, dual homing the attachment circuit and signaling the PW status to the far end PE.

H-VPLS Dual Homing can be used by service providers to protect the otherwise isolated customer (hence "spoke"). Since PEs can be a single point of failure, MTU-s can peer with two hub PEs using two active/standby spoke PWs. Since both spoke PWs are grouped under the MTU-s only one of them can be active at any point of time and forward traffic. Upon failure of the active spoke PW MTU-s causes a switchover from active to standby spoke-PW. Combined with the LDP feature of MAC Address Withdraw, H-VPLS dual homing can be used to avoid a possible traffic black hole.

Three vendors participated in this test: Ciena CN5305, Telco Systems T-Metro-7224 and Cisco ASR9010. The Ciena CN5305 was configured as a MTU-s, dual-homed to the Cisco ASR9010 and Telco T-Metro-7724 PEs. After the active spoke PW was torn down, the Ciena CN5305 successfully switched the bidirectional traffic over the secondary PW. The Cisco ASR9010 was able to send LDP MAC Address Withdraw after the recovery of the primary spoke PW and failure of secondary spoke PW where the Cisco ASR9010 was connected. Since automatic restoration was not supported, the PWs were reverted back to active manually.

Acknowledgements

Editors. Jonathan Morin, Sergej Kälberer, Ronsard Pene, Xiao Tai Yu and Carsten Rossenhövel (all EANTC) co-authored this document.

In addition, we would like to thank Joe Miller and Stephen Murphy from the University of New Hampshire Interoperability Lab (UNH-IOL) for supporting the test and contributing to the white paper under the EANTC/UNH-IOL collaboration program.

REFERENCES

- "Precision Time Protocol (PTP)", IEEE 1588-2008
- "Mobile Backhaul Implementation Agreement Phase 2", MEF technical specification, work in progress
- "Timing and Synchronization Aspects in Packet Networks", ITU-T G.8261/Y.1361
- "Precision Time Protocol Telecom Profile for frequency synchronization", ITU-T G.8264.1, work in progress
- "Timing characteristics of synchronous Ethernet equipment slave clock (EEC)", ITU-T G.8262/Y.1362
- "Distribution of timing through packet networks", ITU-T G.8264/Y.1364
- "Synchronization layer functions", ITU-T G.781
- "Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks", RFC 4717
- "Ethernet ring protection switching", ITU-T G.8032/Y.1344, March 2010
- "Ethernet linear protection switching", ITU-T G.8031/Y.1342, November 2009
- "Pseudowire (PW) Redundancy" draft-ietf-pwe3-redundancy, May 2010
- "Pseudowire Preferential Forwarding Status Bit", draft-ietf-pwe3-

redundancy-bit, May 2010

"Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447

"MPLS transport Profile Data Plane Architecture", draft-ietf-mpls-tp-data-plane

"A Framework for MPLS in Transport Networks", draft-ietf-mpls-tp-framework

"MPLS-TP OAM based on Y.1731", draft-bhh-mpls-tp-oam-y1731

"Linear Protection Switching in MPLS-TP", draft-zulr-mpls-tp-linear-protection-switching

"Proactive Connection Verification, Continuity Check and Remote Defect Indication for MPLS Transport Profile", draft-ietf-mpls-tp-cc-cv-rdi

"Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling" RFC 4762

"Virtual Bridged Local Area Networks - Connectivity Fault Management", IEEE 802.1ag

"OAM functions and mechanisms for Ethernet based networks", ITU-T Y.1731

ACRONYMS

Term	Definition
APS	Automatic Protection Switching
BFD	Bidirectional Forwarding Detection
CCM	Continuity Check Message
CLI	Command Line Interface
CFM	Command Line Interface
EEC	Synchronous Ethernet Equipment clock
ELPS	Ethernet Linear Protection Switching
ERPS	Ethernet Ring Protection Switching
ESMC	Ethernet Synchronization Messaging Channel
H-VPLS	Hierarchical VPLS
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
LDP	Label Distribution Protocol
LOS	Loss of Signal
LSP	Label Switching Path
LTE	3GPP Long Term Evolution
MBMS	Multimedia Broadcast Multicast Service
MEG	Maintenance Entity Group
MEP	Maintenance Entity Point
MIP	Maintenance Domain Intermediate Point
MNCP	Maximum Number of Cells Packed
MPLS	Multi-Protocol Label Switching
MPLS-TP	MPLS Transport Profile
MTIE	Maximum Time Interval Error
MTU-s	Multi Tenant Unit - switch
OAM	Operation, Administration and Maintenance
PE	Provider Edge
PRC	Primary Reference Clock
PTP	Precision Time Protocol
PW	Pseudowire
QL	Quality Level
RPL	Ring Protection Link
SEC	SDH equipment slave clocks
SSM	Synchronization Status Messages
SSU	Synchronization Supply Unit
SyncE	Synchronous Ethernet
TOD	Time of Day
VLAN	Virtual Local Area Network



Authored by:
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
www.eantc.com



Hosted by:
IIR Telecoms

29 Bressenden Place
London SW1E 5DR, UK
Tel: +44 20 7017 7483
Fax: +44 20 7017 7825
registration@iir-conferences.com
www.carrierethernetworld.com

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.
v1.0



Test Setup at EANTC Lab, August 2010