



■ EANTC ■

MEF

Carrier Ethernet World
Congress 2009
Multi-Vendor Interoperability Event
White Paper



**Global
Ubiquitous
Manageable**

Berlin, September 21–25, 2009

EDITOR'S NOTE



Carsten Rossenhövel
Managing Director

Where do Carrier Ethernet implementations stand today? After six months of preparation and a two-week hot staging with 60 engineers from 24 vendors, here is the latest news for the Carrier Ethernet World Congress '09:

Despite the global economic crisis, the majority of Carrier Ethernet vendors (listed on the next page) continue to develop new products with consideration for multi-vendor interoperability.

Our tests show that the technology becomes *broader* — opening up new markets such as in E-NNI and LTE backhaul. It becomes also *deeper* — enabling advanced end-to-end, multi-vendor fault management. And it certainly gets more *mature* — for example, Synchronous Ethernet interoperability.

On the other hand, more deployments in more areas also results in growing challenges for the technology and continuous work for the standards bodies: There is plenty of work to standardize MPLS-TP; E-NNI Phase 1 is in its final stages of standardization in the MEF but not quite fully implemented yet, and good quality synchronization with IEEE-1588:2008 is possible but cannot be taken for granted.

"Global Interconnect", the latest MEF project, is a topic of high visibility these days. Service providers have been connecting their Carrier Ethernet networks for a while — we believe the future will belong to scalable and resilient interconnects, successfully tested this time with MPLS and Ethernet. ENNI failure recovery times were on the long side - understandable given all the protocol translations from one network to the other, but we hope for improvements.

Our Ethernet OAM tests advanced further for both Link and Service OAM. The vendors accepted and mastered the challenge. In the end most issues were configurational: Vendors did not agree how to carry CFM across the network (S-VLAN tagged, untagged, with/without MPLS label) — clarification by standards bodies seems to be required.

On the transport side, MPLS-TP pre-standard, sometimes even pre-IETF-draft testing was interesting. There are still a lot of strategic questions to be decided but technical progress is moving fast. Furthermore, it was a surprise to see PBB-TE back on multiple vendors' agenda. Finally, we tested native Ethernet ring resiliency (ERPS / G.8032).

Please review the detailed results below and visit the live demos in Berlin or the virtual booth after the event!

INTRODUCTION

From a technical perspective, Carrier Ethernet is a rich, broad topic which potentially covers and takes advantage of many different technologies in its implementation. In our efforts to craft the technical scope of our interoperability events in order to give the industry a focused message, we decided to test three major topics. Each topic is correlated to a focus area of the Metro Ethernet Forum (MEF) - the defining body for Carrier Ethernet. Our choices were also influenced by service providers' need to quickly develop new applications and services to generate new revenues. As such, we focused on the following areas:

- Global Interconnect - the MEF coined term for Carrier Ethernet Phase 3 - the worldwide deployment of Carrier Ethernet services. The realization of Global Interconnect will require such tools as External Network to Network Interface (ENNI) which is nearing completion of the first phase specification. Aware that much is to be defined in this phase, we take a look at the various solutions and technologies that are the current state of the art.

- Mobile Backhaul - With the MEF 22 technical specification completed, carriers are ready and waiting to take advantage the Packet Switched Network (PSN) for backhauling their mobile traffic. The tests conducted focus on packet based synchronization mechanisms and legacy mobile transport.

- Managed Ethernet Services - Both from our discussions with service providers and our feedback from a survey given to providers at last year's CEWC, the message was clear: providers are most interested in seeing interoperable solutions for managing Ethernet services. We extended our Ethernet OAM tests (as a response to providers' feedback) and added an extensive set of tests for management solutions.

Last, but not least, we present a detailed view into the current state of transport mechanisms necessary to provide Carrier Ethernet services to customers, including tests for transport technology features from MPLS, MPLS-TP, and PBB-TE.

Within each of these areas, in a tightly scheduled two week hot staging at our Berlin lab, over 60 engineers configured network equipment spanning fifteen racks in order to provide the results we detail in the following report. We hope you enjoy the read.

TABLE OF CONTENTS

Participants and Devices.....	3
Network Design	4
Mobile Backhaul	4
Synchronization	4
GSM and UMTS Transport.....	6
Global Interconnect	8
Managed Ethernet Services	12
Performance Monitoring.....	12
Ethernet in the First Mile (EFM) ...	12
Continuity Fault Management.....	13
Carrier Ethernet Transport.....	15
References.....	18
Acronyms	19

PARTICIPANTS AND DEVICES

Vendor	Participating Devices
Agilent Technologies	N2X
Albis Technologies	ULAF+ MCU-CES ULAF+ Acceed
Alcatel-Lucent	1850 TSS-320 1850 TSS-160
Calnex Solutions	Paragon
Ceragon Networks	FibeAir IP-10
Ciena	CN 5305 CN 3960 CN 3940 CN 3920
Cisco Systems	ME 3400EG-12CS ME 3400EG-2CS Catalyst 3750ME 7606 7604 ASR 1002 Active Network Abstraction (ANA) ^a
ECI Telecom	SR9705 BG9305
Ericsson	Demonstrator SM480
Ethos Networks	E-80 Domain Management System (DMS) ^a
Harris Stratex Networks	Eclipse Packet Node
Huawei Technologies	NE40E-X8 CX600-X3 PTN910 PTN950 PTN1900 PTN3900
Ixia	IxNetwork
JDSU	QT-600 HST-3000 MTS-8000 MTS-6000A EtherNID SmartClass E1
MRV Communications	OS910M OS9124-410G OS904-DSL OS904-MBH
NEC Corporation	Pasolink NEO iP Pasolink NEO HP CX2600

Vendor	Participating Devices
Nokia Siemens Networks	FlexiPacket Microwave Flexi WCDMA BTS
RAD Data Communications	ACE-3220 ACE-3105 ETX-204A ETX-202 IPmux-24
SIAE Microelettronica	ALplus2
Siklu Communication	EtherHaul-250
Spirent Communications	xGEM Spirent TestCenter
Symmetricom	TimeProvider 5000 TimeProvider 500 SSU 2000 TimeWatch Cesium
Tejas Networks	TJ2050 TJ2030
Telco Systems - a BATM Company	T-Metro-XG T-Metro-200 T5C-XG T5C-48T T-Marc-380 T-Marc-280 T-Marc-254P

a. Management system

INTEROPERABILITY TEST RESULTS

The description and results of the technical areas tested are documented in their own respective sections. Each section includes the technical background of the technology, the test procedure used, the successful results, and in some cases some additional information. The document generally follows the structure of the test plan. Also included for some tests is a logical diagram which depicts the results and in some cases the test setup. Some successful tests which involved several vendors were incorporated to the demonstration network, described in the Network Design section below.

It is important to note that the multi-vendor combinations documented detail the successful results completed within the 9 business days available for the hot staging (two weeks minus 1/2 day for setup, 1/2 day for tear-down). Given the time constraint, a full mesh of vendor combinations for some of the technologies tested is impossible to test. If the reader does not find a combination which includes his vendors of choice we encourage him/her to contact the vendor and make sure that the test combination is performed at our next event.

Testers. Throughout the test event the generation of traffic, as well as the measurement of out of service times, was required; both for debugging and in the verification process for many of our tests. We would

like to thank Agilent Technologies, Ixia, Spirent Communications, and JDSU not only for providing the test equipment, but also for the time and effort spent helping debugging and performing tests.

Terminology. For the purposes of readability (and consistency with our previous white papers) we will 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 the same vendor, and also when functionality was used or demonstrated and not thoroughly tested.

NETWORK DESIGN

Based on the interoperability test results, EANTC and the participating vendors constructed a single multi-vendor network topology. The network included all participating devices and a number of Carrier Ethernet demonstrations. The showcase network aimed to mirror modern converged networks including mobile radio access elements, IP/MPLS routers, and both aggregation and access switches. The resulting services spanned equipment from multiple vendors as well as different technologies.

The network showed end-to-end Mobile Backhaul and business services demonstrating the ability of the underlying technologies and devices not only to successfully provide these services, but also to manage them. The detailed results are described in the Mobile Backhaul and Managed Ethernet Services sections below.

The network consisted of two distinct provider domains which were connected using results from the Global Interconnect tests enabling end-to-end network services and applications. Each of the provider domains was created from the conjunction of two transport technologies - one with IP/MPLS and MPLS-TP technologies, and one with IP/MPLS and PBB-TE technologies. The detailed transport technology and administrative domain interconnection test results are described in the Carrier Ethernet Transport and Global Interconnect sections.

MOBILE BACKHAUL

In January 2008 we performed our first Mobile Backhaul interoperability event. We quickly realized that we were a too early for the industry and that not enough Mobile Backhaul solutions existed for interoperability testing. This is clearly no longer the case. We tested deployment and test solutions from 20 of the 24 vendors attending the hot staging from this time around - from solutions to support legacy GSM services all the way to forward looking implementations of Long Term Evolution (LTE).

Mobile backhaul readiness is measured by two yard sticks: The ability to provide transport between base stations and their controllers and the solution's ability to provide clock synchronization to the base stations.

SYNCHRONIZATION IN THE PACKET SWITCHED NETWORK

One of the biggest challenges facing carriers planning to use packet based networks for their mobile backhaul is synchronisation of their mobile base stations. Traditionally, synchronization in mobile networks is provided by the synchronous physical layer (Layer 1) of the network, however, this synchronisation capability is not inherent to Ethernet as it is with TDM/SDH based technologies. Therefore, with the move to Ethernet based transport, carriers must solve this problem independently from the transport network.

Depending on the mobile technology used on the air interface, and base station type, the mobile backhaul plan must consider different aspects and levels of synchronization. It is nearly impossible to define a single set of synchronisation requirements for all mobile backhaul networks. In our event we measured the clock quality of multi-vendor systems typically consisting of a master and slave clock device and provided for each a statement regarding which clock quality level these systems achieved. Whether this clocking quality is sufficient or not lies in the decision of each mobile operator.

Several solutions, from a variety of Standard Development Organizations (SDOs), currently exist to provide synchronization over packet based networks. In our test event we evaluated three mechanisms which can be used to synchronize the packet based network: adaptive clocking, IEEE 1588-2008 protocol, and Synchronous Ethernet - a technology that changes the asynchronous nature of the Ethernet PHY to be synchronous.

For all tests in which we verified the clock accuracy we measured time interval error (TIE) of the slave's clock output (E1 or Synchronous Ethernet interface type) using a jitter and wander analyzer. The TIE was measured against a high precision external atomic clock whose quality was better than the quality of a Primary Reference Clock (PRC) as specified by ITU-T G.811. We then converted the TIE measurements to Maximum Time Interval Error (MTIE) and Time DEVIation (TDEV) curves and compared to the synchronization accuracy requirements expressed as MTIE and TDEV masks defined by ITU-T.

We have seen an increase in participation in these test areas and a significantly higher number of implementations. This facilitated extensive testing which expanded on the procedures used in previous test events.

IEEE 1588-2008 (PTPv2)

In order to verify the synchronization functions of Precision Time Protocol version 2 (PTPv2) defined in IEEE 1588-2008, we measured the time needed for the slaves to synchronize with their master clocks. Following this we measured synchronization quality on the slaves' clock output interface by an external jitter and wander analyzer. Wander was measured

for at least 1 hour after initial synchronization of the slave clock to the master.

During all tests we used the Calnex Paragon was connected between the IEEE 1588-2008 clock master and slave devices. The Paragon emulated a Packet Switched Network (PSN). The impairment generator was configured with a network profile according to ITU-T G.8261 Test Case 12.

The following devices served as PTP masters: Huawei CX600-X3, MRV OS904-MBH, Symmetricom TimeProvider 5000, and Telco Systems T-Marc-254P. The Huawei CX600-X3, Huawei PTN910, RAD ACE-3220 and Telco Systems T-Marc-254P successfully tested PTP slave functionality. All successful test combinations passed the ITU-T G.823 wander traffic mask, and are found in Figure 1. For wander measurements we used the Symmetricom TimeWatch and JDSU wander analyzer.

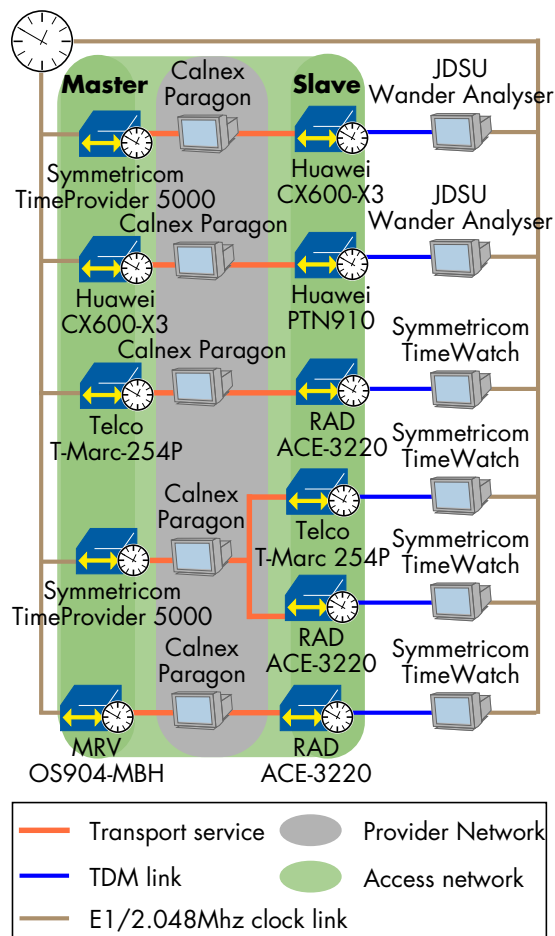


Figure 1: Precision Time Protocol version 2

In this event we housed more PTPv2 implementations in our lab than in any of our previous test events. Much time was spent figuring out the common supported configuration settings and fixing configuration issues. For all test combinations we tested 1-step clock, and used UDP/IP unicast encapsulation. The configured SYNC messages rate was between 16 and 128 per second.

Our results show that the industry is in a progressing stage regarding IEEE 1588-2008 implementations. For our next events we expect to further test advanced IEEE 1588-2008 features like 2-step clock, phase and time synchronization, peer delay,

delay/response protocol mechanisms, boundary and transparent clocks, and other IEEE 1588-2008 encapsulation types.

Synchronous Ethernet and ESMC

Synchronous Ethernet (SyncE) provides mechanisms for high-precision frequency distribution in Ethernet networks similar to functionality provided by SDH. Unlike traditional Ethernet, where the receiving node only performs synchronization on a per-packet basis, internal clocks of the SyncE enabled nodes are constantly synchronized.

A node operating in synchronous mode recovers the clock frequency from an external high precision clock signal which is received either from an Ethernet link or from a clock link such as Building Integrated Timing Supply (BITS). As a result of this operation the node "locks" its internal clock to the external clock signal. This node can then synchronize Ethernet peers in the transmit direction with its internal clock allowing the directly attached SyncE nodes to synchronize their clocks to the node's clock. With this mechanism, synchronization can be provided throughout an Ethernet network.

One of our test goals was to verify the ability of an Ethernet device operating in synchronous mode to synchronize its clock frequency with another SyncE node. We verified the synchronization quality by measuring the jitter and wander of the node's clock output signal. The wander analyzer was connected to the same external clock source as the master SyncE node in order to measure the difference. The results were compared against the synchronization accuracy requirements defined by the ITU-T.

Tests were successfully completed between the devices listed below. The first device in the pair indicates the master, and the second the slave role. All combinations passed at least the ITU-T G.813 SEC Option 1 MTIE and TDEV masks. For wander measurements we used the Calnex Paragon and JDSU MTS-8000 devices: RAD ETX-204A and Alcatel-Lucent 1850 TSS-320, SIAE Microelettronica Alplus2 and Alcatel-Lucent 1850 TSS-320, Ericsson Demonstrator and Alcatel-Lucent 1850 TSS-320, Alcatel-Lucent 1850 TSS-320 and Albis ULAF+ Acceed, RAD ETX-204A and Albis ULAF+ Acceed, SIAE Microelettronica Alplus2 and Albis ULAF+ Acceed, Huawei PTN3900 and Cisco 7604, ECI BG9305 and Albis ULAF+ Acceed, ECI BG9305 and Cisco 7604, ECI BG9305 and RAD ETX-204A, Albis ULAF+ Acceed and RAD ETX-204A, Nokia Siemens Networks FlexiPacket Microwave and Albis ULAF+ Acceed, Nokia Siemens Networks FlexiPacket Microwave and MRV OS904-MBH, ECI BG9305 and Huawei PTN3900, SIAE Microelettronica Alplus2 and MRV OS904-MBH, Ericsson Demonstrator and Nokia Siemens Networks FlexiPacket Microwave, SIAE Microelettronica Alplus2 and Nokia Siemens Networks FlexiPacket Microwave, Alcatel-Lucent 1850 TSS-320 and Albis ULAF+ Acceed, SIAE Microelettronica Alplus2 and Albis ULAF+ Acceed,

Ericsson Demonstrator and MRV OS904-MBH, Huawei PTN910 and MRV OS904-MBH.

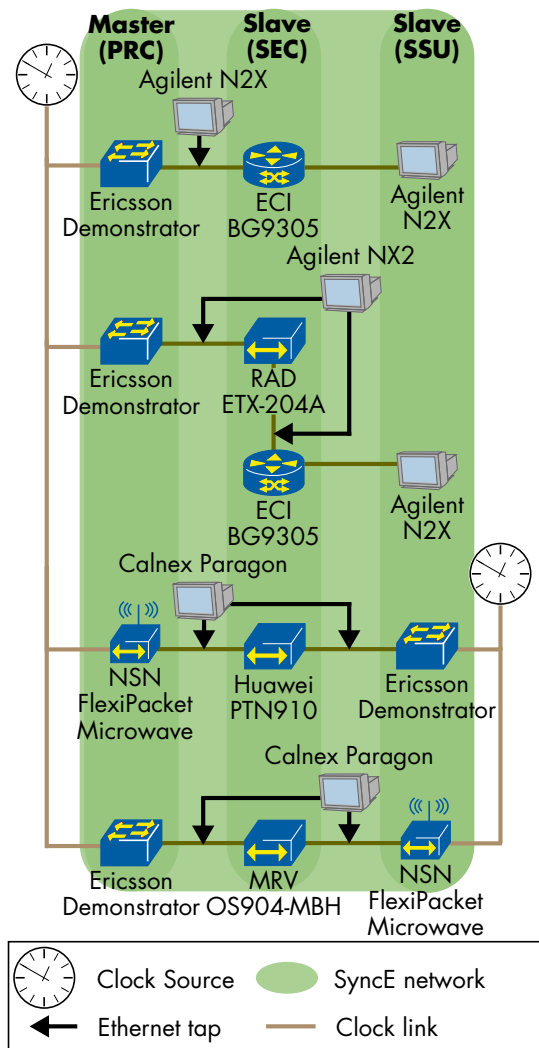


Figure 2: ESMC

In the case of multiple clock sources the network elements should be able to select the best one, and in failure situations to switch from the failed clock source to the working source. This would allow the Service Providers to build resiliency into the clocking infrastructure. This mechanism is called Synchronization Status Messaging (SSM) and its usage for SyncE is defined in G.8264. The origin of the SSM lies in TDM and SDH.

This brings us to the second goal of the SyncE tests - verification of the SSM mechanism. During these tests we first verified the establishment of an Ethernet Synchronization Messaging Channel (ESMC) between the SyncE nodes and then analyzed that the nodes properly generated and interpreted the ESMC status messages. ECI BG9305, Ericsson Demonstrator, Huawei PTN910, MRV OS904-MBH, Nokia Siemens Networks FlexiPacket Microwave, and RAD ETX-204A all successfully passed this test. For ESMC packet analysis we used Agilent N2X and Calnex Paragon. In addition, we used Agilent N2X for ESMC message generation for some tests. All successful test combinations are shown in Figure 2.

Similar to PTPv2 tests, we housed more SyncE implementations in this event than any previous event. The

results showed a high level of interoperability for synchronization between the tested equipment. We also tested Ethernet Synchronization Messaging Channel (ESMC) for the first time - we hope to test still more implementations in the future.

GSM AND UMTS TRANSPORT

Legacy networks such as GSM (commonly referred to as 2G) and UMTS (referred to as 3G) are not going away. Carriers tend to add equipment belonging to the next generation radio technology to a cell site, but rarely do we see legacy technologies decommissioned. As part of the move to packet based backhaul carriers are looking to offload their TDM and ATM base station to controller links to packet networks. This sets a new requirement on mobile backhaul networks: support native and transparent TDM and ATM connectivity between base station and controller.

Both the MEF and the IETF have developed standards to address these requirements before the requirement for mobile backhaul arose. However, with the new drive from mobile carriers to decommission their TDM and ATM networks and to converge all services into one packet base network, these standards are experiencing renewed interests. We verified the interoperability of the devices providing base station to controller connectivity in establishing and maintaining emulated circuits.

TDM and ATM Emulation

Most of today's Mobile Backhaul deployments are based on TDM technology (e.g. GSM) or ATM technology (e.g. UMTS). On the migration path of the existing Mobile Backhuls to Carrier Ethernet, the legacy TDM and ATM equipment can be connected to the new Carrier Ethernet Mobile Backhaul by using translation technologies allowing to transport TDM and ATM data over Carrier Ethernet. During our hot staging we verified ATM transport over MPLS implementations according to the RFC 4717 and TDM structure agnostic and structure aware transport in the IETF version (RFC 4553 and RFC 5086) and MEF version (MEF 8). Additionally to the ability of transport data over a Packet Switched Network, we also verified the quality of adaptive clocking recovery from the Packet Switched Network implemented at the devices. All devices were required to configure a structure aware network service that transport the PCM30 or PCM31 E1 slots 9, 20, 21, and 22.

The following devices were able to successfully transport TDM data over a structure aware network service (CESoPSN): Cisco 7606, Ericsson Demonstrator, Huawei PTN910, Huawei CX600-X3, MRV OS910M, NEC Pasolink NEO TE, and RAD ACE-3220. The Albis ULAF+ MCU-CES, Cisco 7606, ECI BG9305, Ericsson Demonstrator, Huawei CX600-X3, Huawei PTN910, MRV OS910M, MRV OS9124-410G, NEC Pasolink NEO TE, Telco T-Metro-200, RAD ACE-3220, RAD IPmux-24 success-

fully tested TDM structure agnostic CES (SAToP). All tests performed with Ericsson Demonstrator devices were done for STM 1 and STM 4 interfaces towards the circuit network. Also Telco Systems used in one test STM 1 interface. The E1 data structure was encapsulated into STM. For TDM traffic generation and analysis we used the following JDSU tester devices: MTS-8000, MTS-6000A, and SmartClass E1.

We tested Ethernet CES (MEF8) as well as MPLS CES (RFC 4553 and RFC 5086) encapsulation. The following devices tested TDM Circuit Emulation with Ethernet encapsulation: Albis ULAF+ MCU-CES, ECI BG9305, Huawei CX600-X3, Huawei PTN910, Ericsson Demonstrator, MRV OS910M, MRV OS9124-410G, NEC Pasolink NEO TE, RAD IPmux-24, Telco T-Metro-200.

The following devices tested TDM Circuit Emulation with MPLS encapsulation: Cisco 7606, Ericsson Demonstrator, Huawei CX600-X3, RAD ACE-3220, and NEC Pasolink NEO TE.

For both structure agnostic and structure aware TDM Circuit Emulation tests the slave devices were configured to perform adaptive clock recovery from the Packet Switched Network. We used Calnex Paragon and Spirent xGEM impairment devices to emulate a real network between master and slave devices according to G.8261 Test Case 1. The impairment was introduced from the beginning of the test as the slave clock was in Free Running state. As soon as clock changed in synchronous state we measured time interval error from the reference clock for at least one hour, using the JDSU jitter/wander testers.

The following device combinations passed the ITU-T G.823 traffic MTIE and TDEV masks, the first of which indicates master role, the second a slave: RAD IPmux-24 and Albis ULAF+ MCU-CES, Albis ULAF+ MCU-CES and Telco Systems T-Metro-200, Telco Systems T-Metro-200 and Albis ULAF+ MCU-CES, Telco Systems T-Metro-200 and ECI BG9305. These pairs performed TDM Structure Agnostic CES.

The Cisco 7606, Huawei CX600-X3, Huawei NE40E-X8, Huawei PTN910, NEC CX2600, and NEC Pasolink NEO TE tested successfully ATM transport over MPLS. All tests were performed for ATM VPI n-to-1 encapsulation, with n set to 1. For the test between Cisco 7606 and Huawei CX600-X3 and the demonstration between Huawei NE40E-X8 and Huawei PTN910 we configured cell concatenation of 10 cells per PDU and concatenation timeout of 10 ms. In all other tests we tested without cell concatenation, which means the devices encapsulated a single ATM cell per PSN PDU.

Our tests show that TDM structure agnostic and ATM transport over MPLS have achieved a mature implementation status - we did not observe any interoperability issues. In regards to CESoPSN testing, we observed that not all implementations support both PCM30 and PCM31 E1 and also some implementations could not configure different E1 slots (9, 20, 21, and 22) as required by our tests. Also in most cases the implementations either did not fully support

E1 alarm propagation but rather implemented it in a different way. For this reason we see a reason to continue testing CESoPSN interoperability in further events. We also observed increased success for adaptive clocking tests.

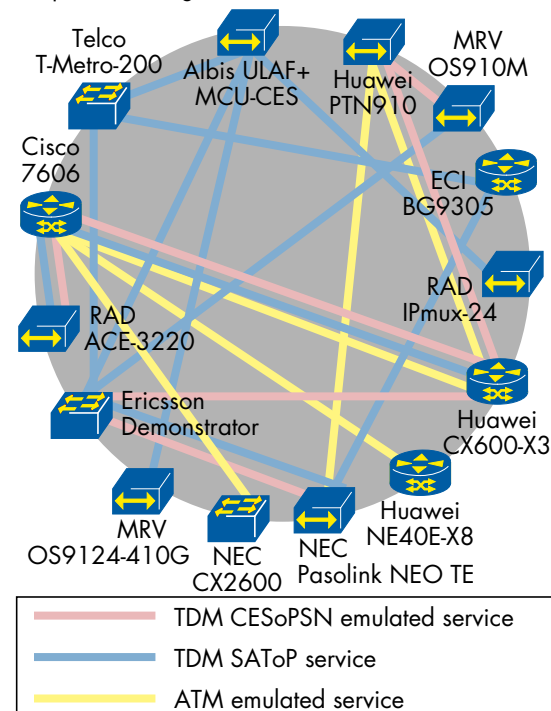


Figure 3: TDM and ATM Emulation

Microwave Transport

We have seen a great deal of development in the products offered by Microwave vendors since we started testing these devices at our interoperability events. As the Microwave transport market became more and more commoditized vendors started implementing more and more intelligence in their devices moving away from pure transport into advanced switching functionality with Microwave uplinks.

The tests focusing on Microwave devices demonstrated these advanced capabilities - from reacting to degraded Microwave link conditions to propagation link state information to the remote end. The former test, more specifically referred to as Adaptive Coding and Modulation (ACM), involved transmitting two classes of Ethernet traffic each distinguished by the Priority Bits in the customer VLAN tag. An attenuator was then used to degrade the signal, at which point the system was expected to automatically lower the Quadrature Amplitude Modulation (QAM). Since lowering the QAM also reduces the bandwidth of the air interface, the system was then expected to only drop traffic of a lower priority. This test was successfully conducted with the Ceragon FibeAir IP-10, Harris Stratex Eclipse Packet Node, Huawei PTN 950 and 910, NEC Pasolink NEO HP, Nokia Siemens Networks FlexiPacket Microwave and the SIAE Microelettronica ALplus2. Siklu performed a similar test with their EtherHaul-250 device where the change in modulation was manually configured as the attenuation was increased.

In order for fault detection to work even over the microwave air interface, some systems have implemented the capability to propagate the loss of signal on the air interface to the Ethernet interfaces, as well as the propagation of one Ethernet link going down to the associated Ethernet link on the opposite side of the air interface. Both functionalities were demonstrated by the Ceragon FibeAir IP-10, Harris Stratex Eclipse Packet Node, NEC Pasolink NEO HP, Nokia Siemens Networks FlexiPacket Microwave and the SIAE Microelettronica ALplus2.

Some vendors demonstrated additional features. Ceragon, for example, showed a resilient ring of six FibeAir IP-10 switches. Failure tests were run where the Ethernet link was disrupted as well as the radio link. In all cases the out of service times ranged from 12 to 50 milliseconds. Harris Stratex demonstrated a Gigabit Ethernet service aggregated over multiple microwave paths. To demonstrate that jumbo frames can be transmitted over microwave links, Harris Stratex, NEC, SIAE Microelettronica, and Siklu all passed traffic consisting of 7600 byte frames.

The following microwave systems demonstrated the transport of Precision Time Protocol version 2 (IEEE 1588-2008): Ceragon FibeAir IP-10, Harris Stratex Eclipse Packet Node, NEC Pasolink NEO HP, Nokia Siemens Networks FlexiPacket Microwave, SIAE Microelettronica ALplus2 and Siklu EtherHaul-250. The master PTPv2 devices were the Symmetricom TimeProvider 5000 and the Telco Systems T-Marc-254P, with the Symmetricom TimeProvider 500 and RAD ACE-3220 as slaves. The demonstration was prepared for the Carrier Ethernet World Congress in Berlin, and was run over a VPLS service provided by the network. We observed the wander with the previously mentioned wander analyzers after the demonstration was setup and recorded that it did not exceed the expected masks.

GLOBAL INTERCONNECT

How are service providers going to offer complete, end-to-end Ethernet services that span multiple administrative domains? Global interconnect aims to answer this question and to expand Carrier Ethernet services from one provider to the next. The end results of such service is true global-encompassing Carrier Ethernet services.

The Global Interconnect term stands also for all interconnections of a network that consist of multiple administrative domains. A single interconnection between two administrative domains is called an ENNI (External Network Network Interface). Devices implementing ENNI must map network services and their attributes between the administrative domain and ENNI, in particular:

- Service delimiting mapping (e.g. MPLS pseudowire labels into 802.1ad S-Tags),
- Service monitoring (e.g. between Ethernet OAM on ENNI and Pseudowire OAM)
- Service attributes (e.g. Availability, QoS)

We verified service delimiting mapping between

MPLS pseudowire (PW) intra-domain segments to ENNI PW segment or ENNI S-Tag, and we were able to create resilient multi-vendor ENNI scenarios, increasing the availability of the network services. The results of these tests are described below.

MPLS and Ethernet Interconnect

An end-to-end Carrier Ethernet network service passing multiple provider domains is realized via segmented network service, a segment per provider domain and a third at the ENNI. In case of Provider Bridging based ENNI, the 802.1ad S-Tags are used for service delimiting and treated as virtual Access Circuits (AC) at the devices implementing ENNI. The switching between the intra-domain PW segments and ACs at ENNI is called "Pseudowire Switching using AC" and described in IETF draft draft-ietf-pwe3-segmented-pw.

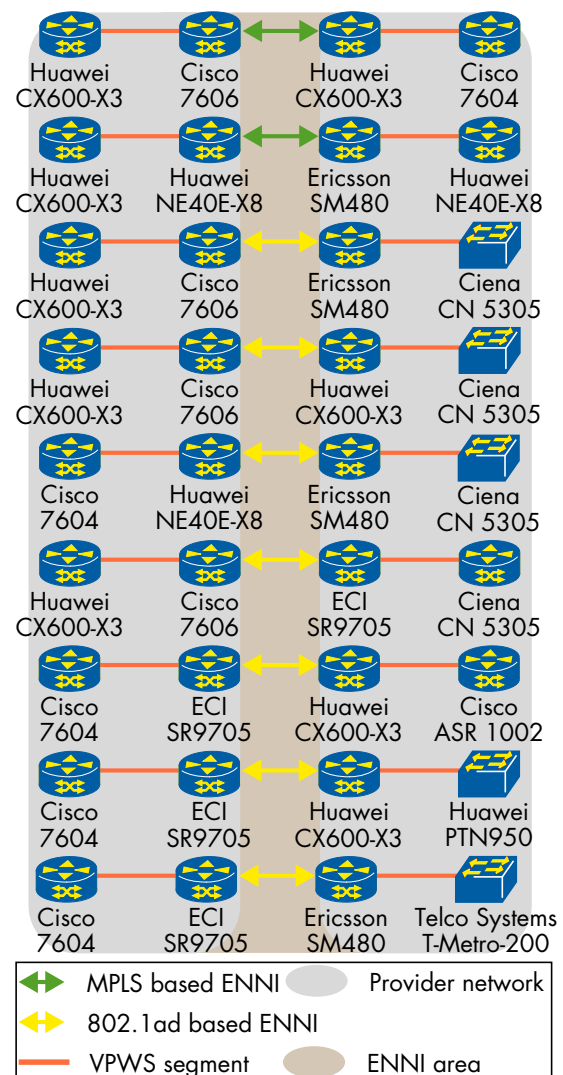


Figure 4: MPLS and Ethernet Interconnect

With MPLS based ENNI, the devices implementing ENNI maintain two PW segments per network service. One PW segment is traversing the ENNI, and the other towards the native domain. The ENNI devices implement switching between the intra-domain and ENNI PW segments. This mode of operation is called "PW Control Plane Switching" in IETF draft draft-ietf-pwe3-segmented-pw.

We tested MPLS based Global Interconnect on the following devices: Cisco 7606, Ericsson SM480, Huawei CX600-X3, and Huawei NE40E-X8. The devices performed MPLS PW switching between provider and ENNI PW segments. eBGP with MPLS label extension was used on ENNI in order to signal the MPLS transport LSP ENNI segment.

We tested Ethernet based Global Interconnect on the following devices: Cisco 7606, ECI SR9705, Ericsson SM480, Huawei CX600-X3, and Huawei NE40E-X8 - which switched between a provider segment PW and an S-VLAN tag on the ENNI.

The following devices served as UNI Provider Edge devices: Ciena 5305, Cisco 7604, Cisco ASR 1002, Huawei CX600-X3, Huawei NE40E-X8, and Telco Systems T-Metro-200. These devices established inter-provider VPWS over the Ethernet and MPLS based ENNIs mentioned above.

The tests were not conducted as quickly as one might think, requiring manual configuration and extensive troubleshooting of Ethernet based ENNIs. The vendor implementations differed regarding the handling of the PW VC types ("Raw" vs "Tagged" mode) and VLAN tag manipulation. For each device combination we discovered a configuration that worked, determining where the S-VLAN tag had to be pushed and popped (at UNI or at ENNI), and which type of PW to use (Ethernet or VLAN type). Figure 4 shows the successfully tested combinations.

ENNI Failure Scenarios

The forwarding of various OAM notifications and failure triggers between the segments of a network service is often required in order to synchronize the status of the network service across all domains and devices. We configured two inter-domain point-to-point network services between two access devices - one service served as primary, and other as backup. We expected the simulated ENNI failure to be signaled across the service from UNI to UNI.

There are different triggers and status notification functions that can be used for such a scenario. During our testing two multi-vendor scenarios were successfully tested as shown in Figure 5.

In one scenario the primary service at the ENNI was realized by Huawei CX600-X3 and ECI SR9705 with a backup through Ericsson SM480 and Cisco 7606. CFM sessions monitored the liveness of the primary service between the Cisco 7604 and ECI SR9705 in one provider domain, and between the Cisco 7604 and Cisco ME3400EG-2CS in the access area. Pseudowire protection was configured in the second provider domain on the Ciena CN 5305. The ENNI failure was triggered by a simulated power failure on the ECI SR9705. Upon the Loss of Signal (LOS) towards the ECI device the Huawei CX600-X3 signaled the PW down via an LDP Withdrawal Message towards the Ciena CN 5305, causing the Ciena CN 5305 to switch to the backup PW. In parallel, the Cisco 7604 received a Continuity Check (CC) error for the CFM session towards the ECI SR9705 and reacted by trans-

mitting a CFM Alarm Indication Signal (AIS), defined in ITU-T Y.1731, towards the Cisco ME3400EG-2CS which then switched over to the backup service based on the CFM AIS.

In the second scenario the failure was triggered by a LOS between the Cisco 7606 and Huawei CX600-X3. The Cisco 7606 reacted by transmitting a Pseudowire Status Notification (AC down) to the Huawei CX600-X3 at the UNI, and the Huawei CX600-X3 at the ENNI reacted to the LOS by sending an LDP Withdraw Message to the Ciena CN 5305. After receiving of the Pseudowire Status Notification the Huawei CX600-X3 device at the UNI shut down the local PW Access Circuit (AC). Upon LOS on the primary AC, the RAD ETX-204A switched over to the backup AC. Agilent N2X was used in both scenarios as an Access Device emulator connected to the Ciena CN 5305.

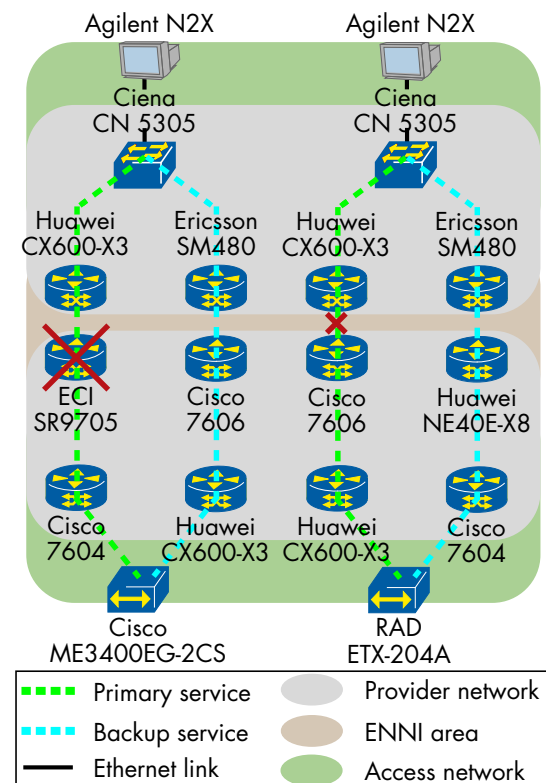
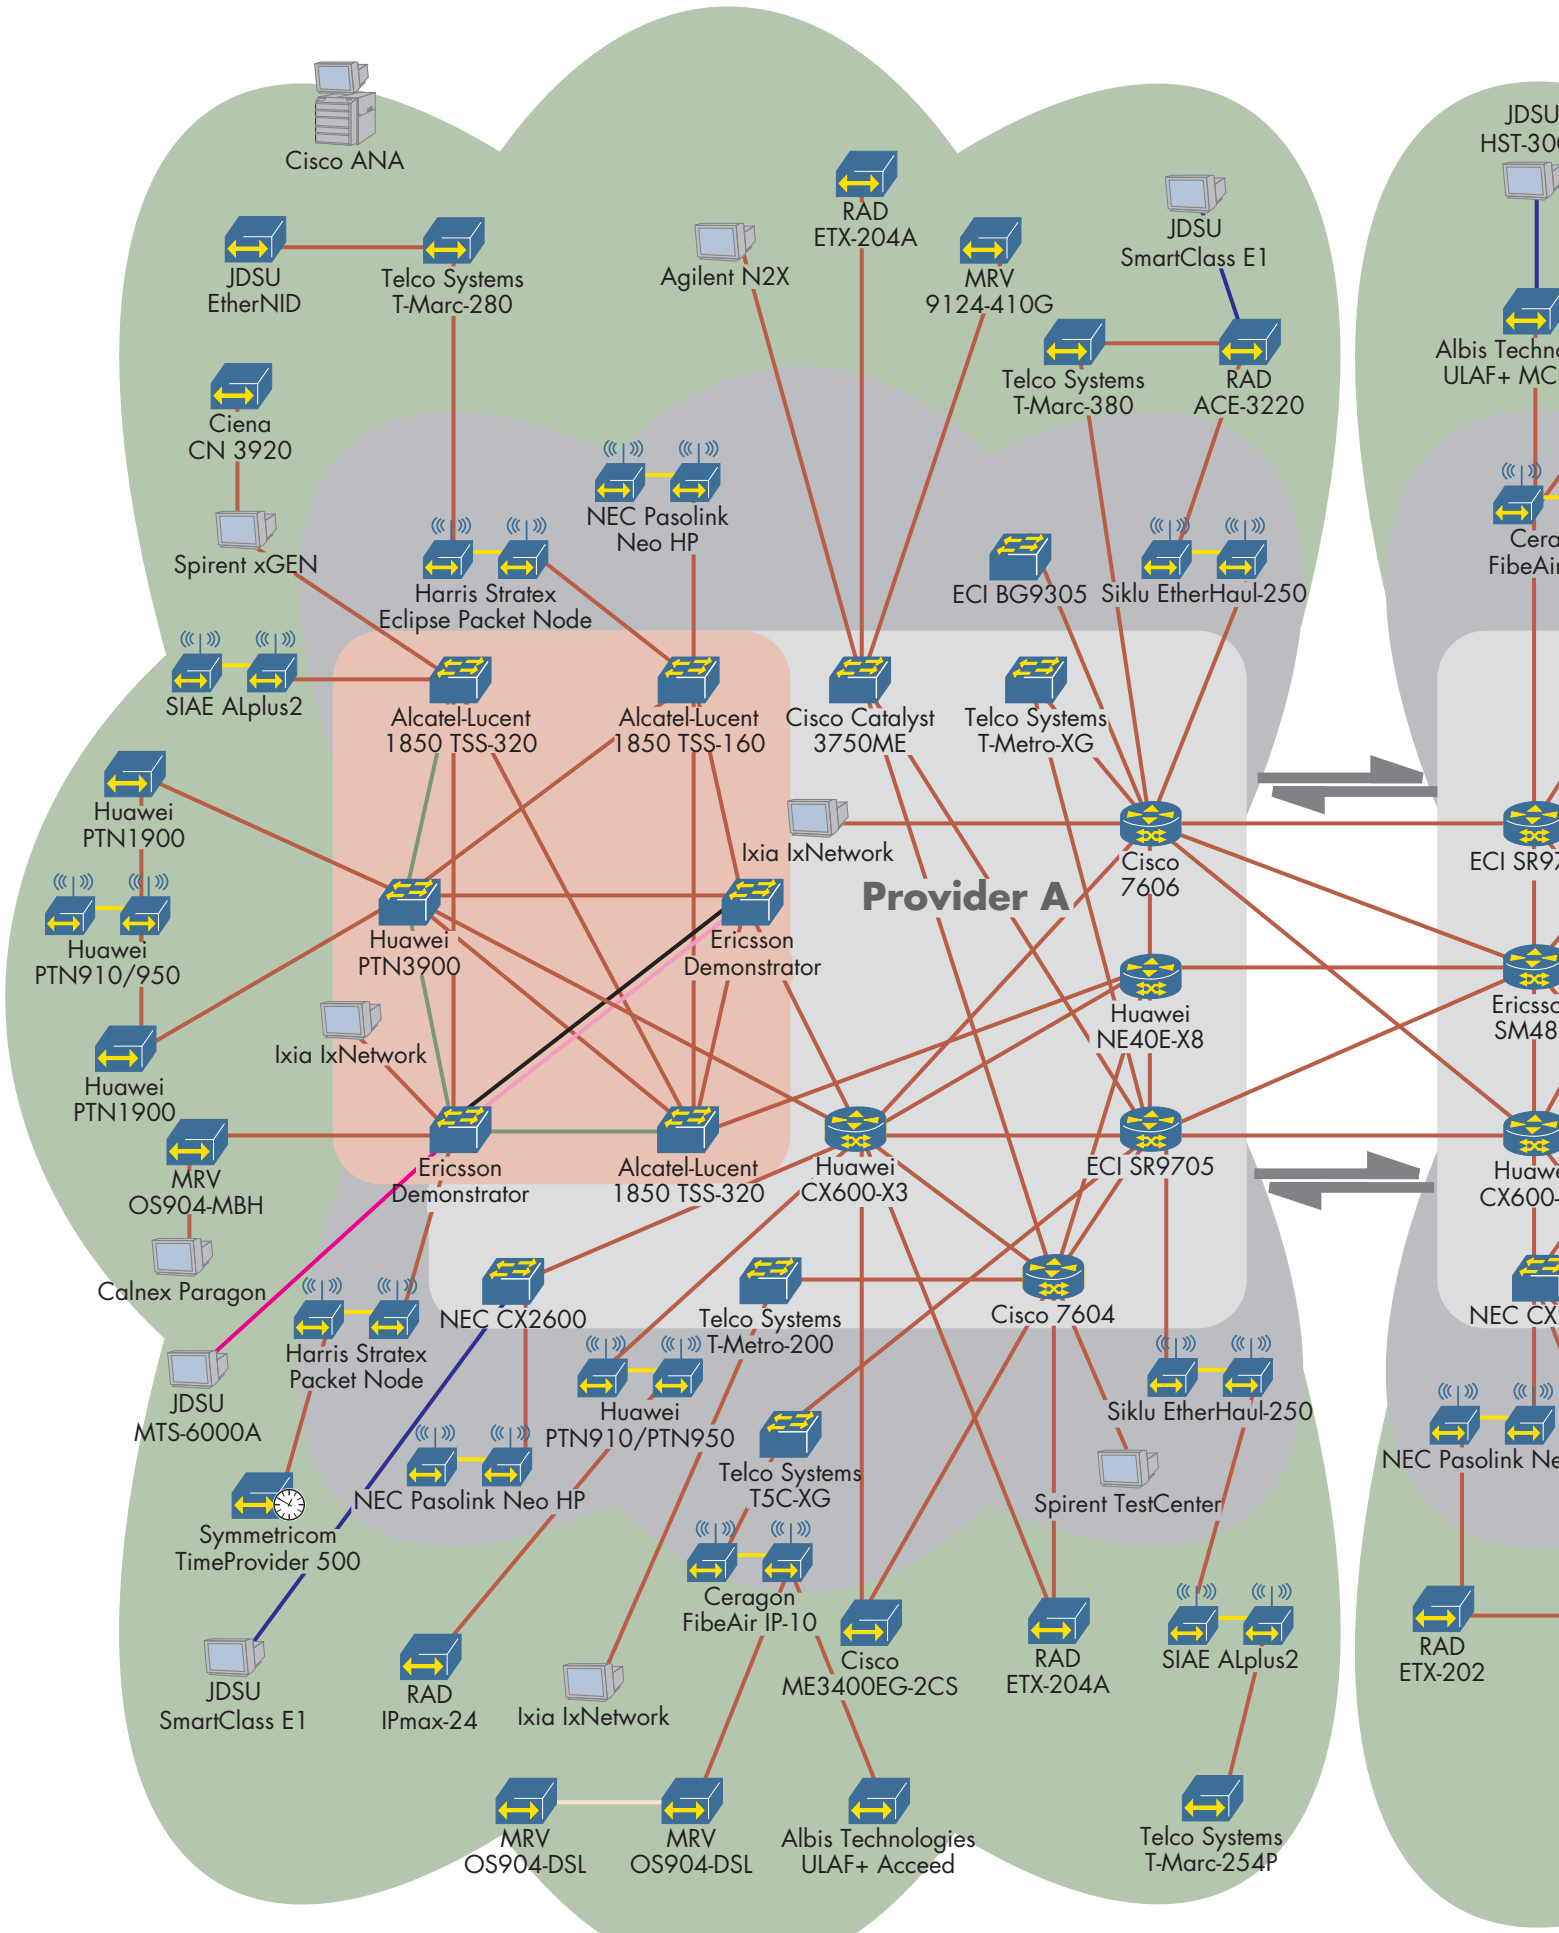
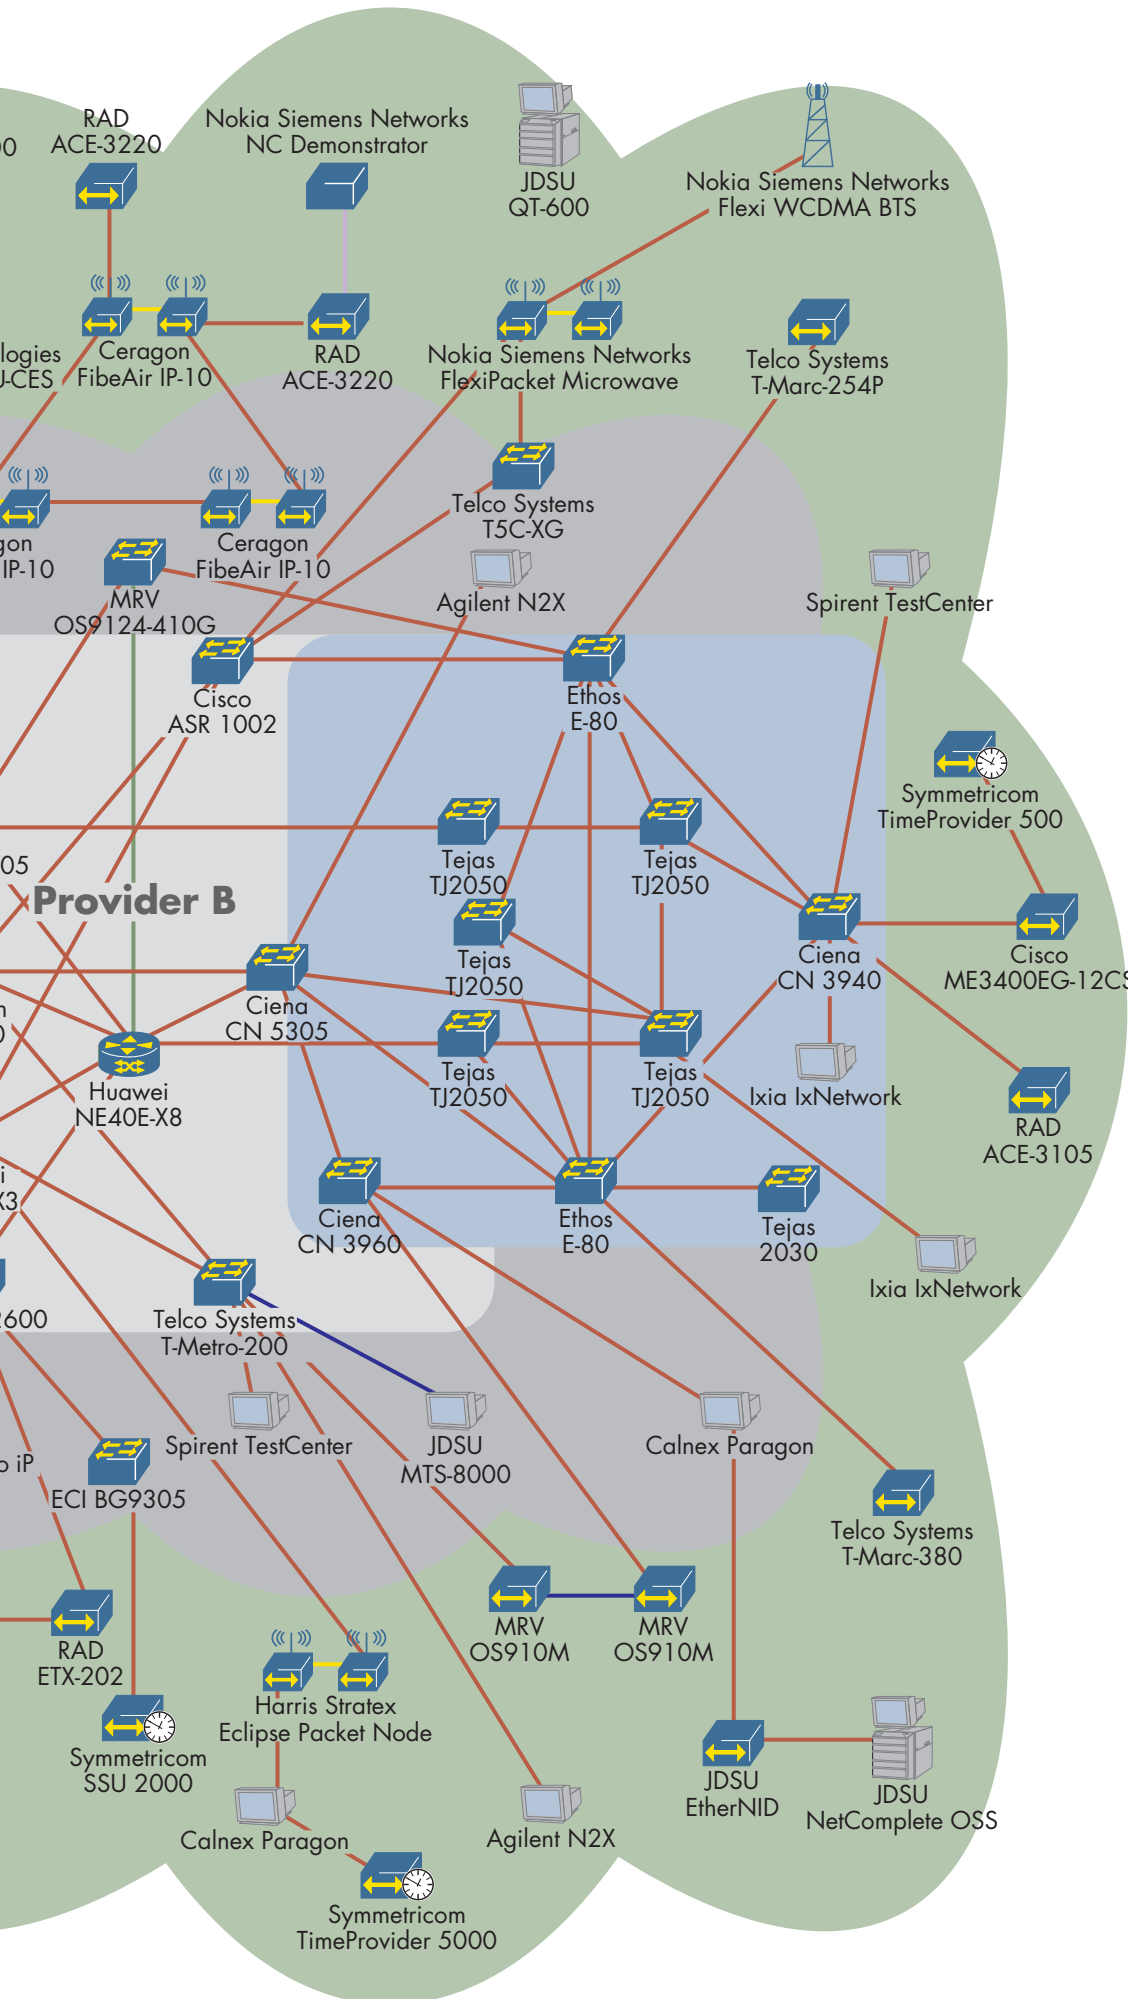











Figure 5: ENNI Failure Scenarios

In both scenarios we measured failover times below 400 milliseconds (ms). The switch-over from the backup to the primary service upon recovery of the primary service was performed manually. Those working in transport networks might find failure times greater than 50 ms to be high, however, these out of service times were expected given the number of different time-outs and protocol translations. Our previous events showed failover mechanisms within a single technology to converge under 50 ms across different implementations, however here the value in testing such multi-domain scenarios was underlined.




















Device Types

-  Metro/Core Network Device
-  Aggregation Device
-  Radio Access Network Controller
-  Access Device
-  1588 Clock Master/Slave Access Device
-  Tester
-  Management System
-  Microwave Device
-  Mobile Base Station

Connection Types

-  10 Gigabit Ethernet
-  Gigabit Ethernet
-  Fast Ethernet
-  TDM link
-  OTN link
-  SDH link
-  SHDSL/ADSL/VDSL link
-  ATM link
-  Radio Link

Network Areas

-  Access network
-  Provider network
-  PBB-TE network
-  MPLS-TP network
-  MPLS network
-  Global interconnect

MANAGED ETHERNET SERVICES

As Carrier Ethernet standards are becoming more mature, service providers are feeling more and more comfortable deploying Ethernet Services. To an extent this is thanks to the evolution from Ethernet Services – an Ethernet based service similar to leased lines in which the customer receives Ethernet network access, to Managed Ethernet Services – a high margin service in which the provider is responsible for the health of the service (usually regulated through strict SLAs) and has a demarcation device at the customer site.

We have been testing Ethernet Services at EANTC's interoperability events since 2005. Now that many of the standards that compromise operations, administration and maintenance (OAM) as well as provisioning are more mature we incorporated all aspects into one test area.

Performance Monitoring

Ethernet service OAM is specified in two complementing standards – IEEE 802.1ag and ITU-T recommendation Y.1731. The ITU-T recommendation specifies the performance monitoring function, including an array of performance monitoring messages. Of these messages, we focused on delay measurement and reply messages (DMMs and DMRs). In addition, the Loopback Messages (LBM) and Loopback Replies (LBR) specified in 802.1ag were used to demonstrate frame delay and delay variation measurement based strictly on the IEEE standard.

All tests were performed using either the Calnex Solutions Paragon or Spirent xGEM impairment equipment. First, messages were exchanged between the DUTs without any impairment to show baseline interoperability. Then, 20 ms delay and 5 ms delay variation impairments were generated using the impairment devices. The DUTs were expected to recognize and report this increase in delay, and we expected the reports to match.

The following devices successfully participated in tests for both two-way frame delay, and two-way frame delay variation: Ciena CN 5305, Ciena CN 3960, Cisco ME 3400EG-12CS, IXIA IxNetwork, JDSU EtherNID, MRV OS904-DSL, NEC CX2600, NEC Pasolink NEO TE, RAD ETX-204A, Telco Systems T-Marc-380 and Telco Systems T-Marc-280. Additionally, the Tejas TJ2050 device tested two-way frame delay only.

During the tests we found that some devices supported more precise delay measurements by using two optional time stamps in the DMMs and DMRs to exclude the processing of delay, reflecting only the transmission delay. This allows a provider to monitor the delay induced by the network, omitting delay from OAM processing by the endpoints of the Ethernet OAM. These tests revealed successful results without issues.

In addition the Cisco ME 3400EG-12CS demonstrated performance monitoring using LBM and LBR

messages towards the Telco Systems T-Marc-280. The ME 3400EG-12CS transmitted LBMs to measure round trip time from which the frame delay and delay variation were calculated. Between both devices the Spirent xGEM was used as an impairment generator. The demonstration showed an accurate measurement, which included OAM processing by the far end nodes.




Figure 6: Performance Monitoring Y.1731

Ethernet in the First Mile (EFM)

Ethernet in the first mile (EFM) is specified in the IEEE standard 802.3ah-2004. The standard defines a loopback mode which intends to assist carriers in performing fault localization as well as testing link performance. Additionally, different types of events are specified in the standard, which can be used for the immediate reporting of local status to a remote device. Operators can use these indications arriving from an Access Device to localize failures and determine the reasons for an inactive customer connection.

We verified the ability of the Access Devices and Provider Edge to generate and properly interpret the link events that signal the number of frame errors, and the number of errored frame seconds within a specified time period. Each device was tested using three thresholds configured as following: 10 errored frames in 10 seconds, 2 errored frame periods per smallest window size supported by the implementation, and 5 errored frame seconds in one minute. First, a traffic generator sent IMIX traffic, filling the window of the peered device. Then, an impairment device introduced CRC errors into the traffic for twenty seconds at one CRC error per second, adding errors to the frames. The device receiving traffic was then expected to send at least one Errored Frame Event, one Errored Frame Period Event and one Errored Frame Second Event.

Four vendors successfully tested these link events as shown in the diagram, using the following devices: Alcatel-Lucent 1850 TSS-320, Ciena CN 3920, Ciena CN 3940, Cisco ME 3400EG-12CS, and Ixia IxNetwork. In tests of Ixia IxNetwork, IxNetwork performed impairment and EFM emulation at the

same time. Otherwise, impairment was provided by the Calnex Paragon or Spirent xGEM as shown in the diagram. Additionally, the Telco Systems T-Marc-380 successfully tested Errored Frame Events with vendor pairs shown in the diagram.

We also verified that each device was able to bring its peer into loopback mode. The device in loopback mode was expected to discard all incoming frames destined towards its peer and loop back all frames coming from the peer. On receipt of these loopback frames the peer was expected to drop them.

The following devices successfully tested loopback mode: Agilent N2X, Alcatel-Lucent 1850 TSS-320, Ceragon FibeAir IP-10, Ixia IxNetwork, MRV OS904-DSL and Telco Systems T-Marc-380.

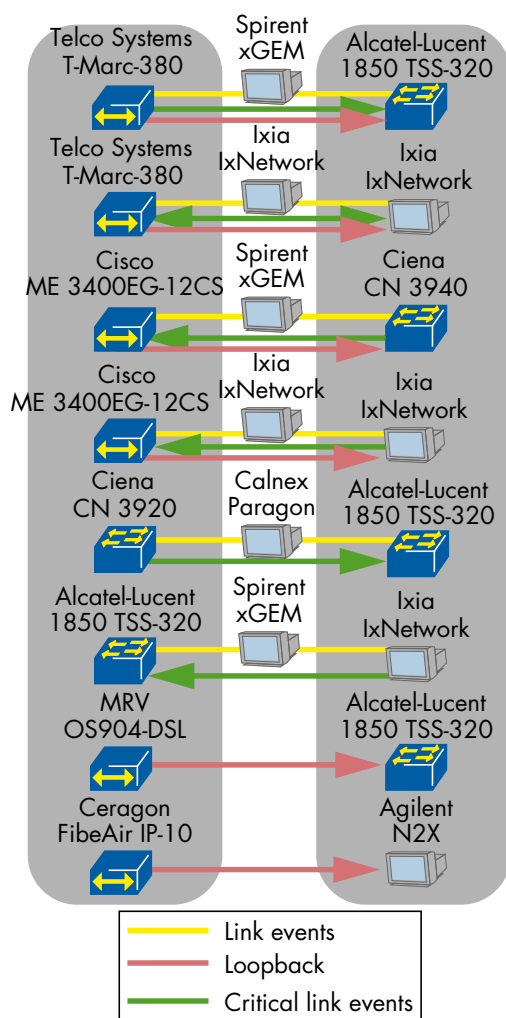


Figure 7: EFM on the UNI

In some cases we also verified the signalling of critical link events such as Dying Gasp when the Access Device was disconnected from power. Dying Gasp is a message defined in the EFM standard which is used by a device when it is being shut down to signal the event to its peer gracefully.

Four vendors also successfully tested the signaling of critical link events: Alcatel-Lucent 1850 TSS-320, Ciena CN 3920, Ciena CN 3940, Cisco ME 3400EG-12CS, Ixia IxNetwork and Telco Systems T-Marc-380. Two such devices: CN 3920 and T-Marc-380 successfully generated Dying Gasp messages and the Alcatel-Lucent 1850 TSS-320,

Cisco ME 3400EG-12CS and Telco Systems T-Marc-380 successfully received the messages. Ixia's IxNetwork was used for emulation and reception of Dying Gasp, Link Fault and Critical Event.

In addition the Ceragon FibeAir IP-10 demonstrated EFM link discovery with the following devices: Albis Technologies ULAF+ Acceed, Cisco ME 3400EG-12CS, RAD ETX 204A, Ixia IxNetwork and MRV OS910M.

Hierarchical Continuity Fault Management

The IEEE 802.1ag Connectivity Fault Management (CFM) standard defines end-to-end Ethernet based OAM mechanisms. Its major use for service providers and enterprises is to verify connectivity across different domains managed by independent entities, such as the Customer, the Service Provider and the Operator(s). The MIPs (Maintenance Association Intermediate Points) present in a certain Maintenance Domain can be used to perform the linktrace procedure which is similar to the Traceroute tool known from IP.

We configured four E-Lines in the network using CFM capable devices to enable the testing. To verify the linktrace function we instantiated CFM Maintenance Associations (MAs) across the E-line services using three Maintenance Domains (MD): Customer MD, Service Provider MD and Operator MD. The MA at the Customer MD Level was built between two Down-MEPs, which were residing at each UNI-C. A Down-MEP is a MEP communicating through a physical port, where as an Up-MEP communicates through the internal switch function of the device. The Service Provider MD Level was configured between two Up-MEPs at the two UNI-N, transmitting LTMs towards the provider network. Finally, the Operator MD Level was built between the provider edges (two providers were simulated, correlating with our Physical Topology in the centerfold of this document). MIPs were then configured on devices which supported MIP functionality, and had a MEP configured at a lower MD Level.

In each MD Level of each scenario 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 participated the test, building six multi-vendor scenarios, and a total of 18 different devices all together: Agilent N2X, Alcatel-Lucent 1850 TSS-320, Alcatel-Lucent 1850 TSS-160, Ceragon FibeAir IP-10, Ciena CN 5305, Cisco Catalyst 3750ME, Cisco 7604, Cisco ME 3400EG-12CS, ECI SR9705, Ixia IxNetwork, JDSU EtherNID, RAD ETX-204A, Ericsson SM480, Spirent TestCenter, Tejas TJ2050 and Telco Systems T5C-XG, Telco Systems T-Metro 200 and Telco Systems T-Marc-280. All MEP to MEP LTMs possible in the six scenarios were successfully returned. Additionally MRV OS904-DSL was able to response LTMs sent by

Agilent N2X.

During the testing we encountered several constraints due to conflicting configuration requirements and support of the DUTs. In one case a device could not create a MEP and MIP using the same VLAN-ID on the same port. Another device would not respond to LTM's which were 802.1ad tagged. One device required VLAN tagged (802.1Q or 802.1ad) CFM frames in order to respond. Finally, one device in the provider domain did not support MIP functionality at all. Due to these constraints, some devices shown in the diagram did not configure a MIP where shown, and we therefore refer you to the following text.

The following MIP devices properly responded to LTM's destined to the pair MEP: Alcatel-Lucent 1850 TSS-320, Alcatel-Lucent 1850 TSS-160, Ceragon

FibeAir IP-10, Cisco Catalyst 3750ME, Ciena CN 3505, Cisco 7604, Ericsson SM480, Tejas TJ2050, Telco Systems T5C-XG, Telco Systems T-Marc-280, and T-Metro 200.

The following MIP devices properly responded to LTM's destined to them as a MIP: Ceragon FibeAir IP-10, Cisco Catalyst 3750ME, Cisco 7604, Ericsson SM480, Tejas TJ2050, Telco Systems T5C-XG, and Telco Systems T-Marc-280.

After our achieved success, and many discussions of configuration requirements from different devices, we are confident that all parties left with a more global understanding of CFM. Additionally we look forward to the specifications from the different standards bodies aligning in regards to CFM and how it is realized over different encapsulations and transport technologies.

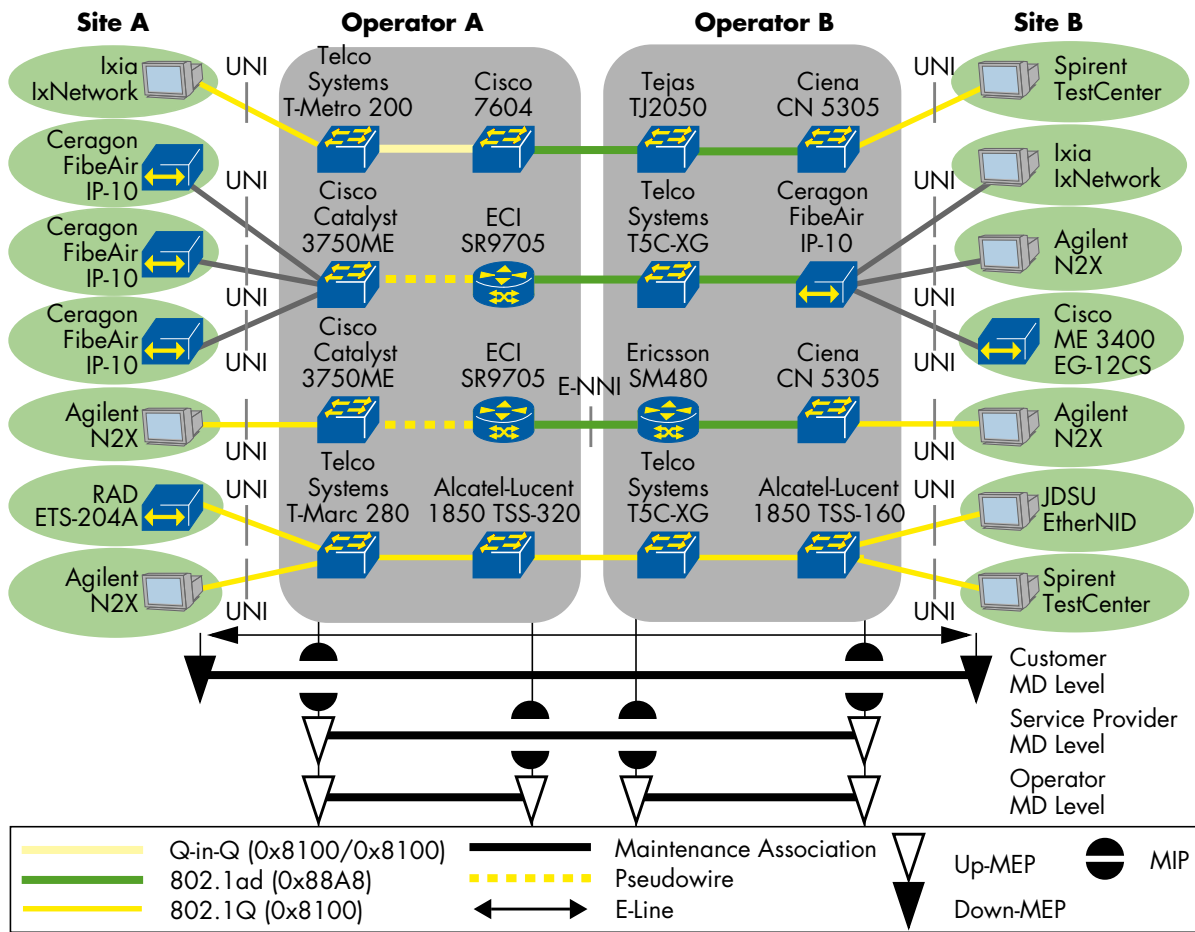


Figure 8: Connectivity Fault Management IEEE 802.1ag

E-LMI and IP/MPLS Interworking

The Ethernet Local Management Interface (E-LMI) specification defines a protocol and procedures that provide the means for Ethernet Virtual Connection (EVC) status notification of the UNI-N to the access device implementing UNI-C. In particular, the E-LMI protocol includes the following procedures:

- Notification to the UNI-C of an added EVC
- Notification to the UNI-C a deleted EVC
- Notification to the UNI-C of the availability state of an EVC (Active, Not Active, Partially Active)

- Communication of UNI and EVC attributes to the UNI-C.

We were able to test also the interworking of MPLS LDP protocol status notification messages with E-LMI. Huawei NE40E-X8 and Cisco 7604 configured an MPLS Virtual Private Wire Service (VPWS). The Cisco 7604 then configured a UNI with the Cisco Catalyst 3750ME. As soon as Huawei NE40E-X8 signalled VPWS Attachment Circuit (AC) status "down", the Cisco 7604 mapped this notification to an asynchronous E-LMI EVC report message towards the Cisco Catalyst 3750ME signaling an EVC "Not Active" status. Upon reception of this E-LMI event,

the Cisco Catalyst 3750ME successfully shut down the EVC and did not forward any traffic to it. Once the Huawei NE40E-X8 signaled VPWS status "up" via an LDP notification message, the Cisco 7604 mapped this notification to asynchronous E-LMI EVC report message and sent it to the Cisco Catalyst 3750ME signalling the EVC "Active" status, causing the device to enable the EVC and forward traffic to it. Finally when we administratively deleted the EVC on the Cisco 7604, the EVC deletion was successfully signalled via E-LMI to Cisco Catalyst 3750ME and Cisco Catalyst 3750ME deleted the EVC and stopped forwarding data to it.

Monitoring and Provisioning

While multi-vendor networks are more common today than they were five years ago, multi-vendor provisioning systems are still a bit rare. Nevertheless, they are of great use to providers. Management systems which display the network topology or the health of certain links and protocols are much more interoperable thanks to the standardization of network protocols, SNMP, and protocol Management Information Bases (MIB). Several of the participating vendors showed their management systems, some of which were able to display multi-vendor information.

Using their Domain Management System (DMS), Ethos Networks successfully provisioned PBB-TE trunks across the Ethos E-80 and Tejas TJ2030 devices and services across their respective access devices - Telco Systems T-Marc-254P and Telco Systems T-Marc-380.

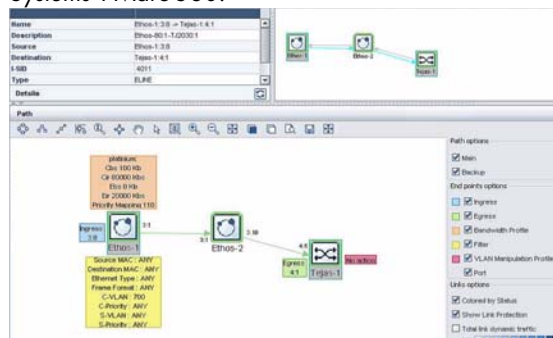


Figure 9: PBB-TE Provisioning on Ethos DMS

Cisco ANA management system demonstrated interoperability across Ceragon, Cisco, Ciena, Ericsson, SIAE Microelettronica, and Telco Systems equipment, automatically displaying inventory and topology between the various devices. Using either the built in interface or by cross launching other vendors' web interfaces, the ANA was able to initiate CFM linktrace and loopback messages during CFM LTM testing. The graphical interface, as shown in Figure 10, aided in visualizing the CFM testing.

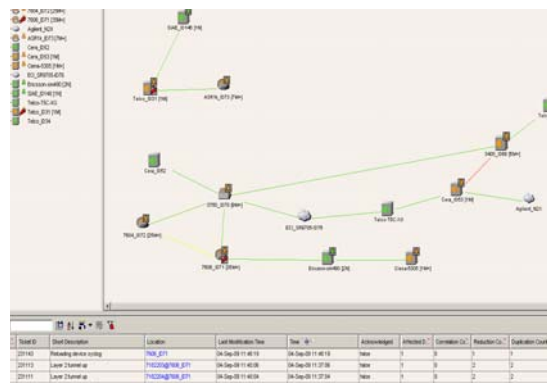


Figure 10: CFM Test Topology on Cisco ANA

CARRIER ETHERNET TRANSPORT

Ethernet transport embodies the basis of a Carrier Ethernet service - the ability of a network to deterministically transport Ethernet frames throughout a service provider network. Answering to the transport call were the following technologies: MPLS, MPLS-TP, and PBB-TE.

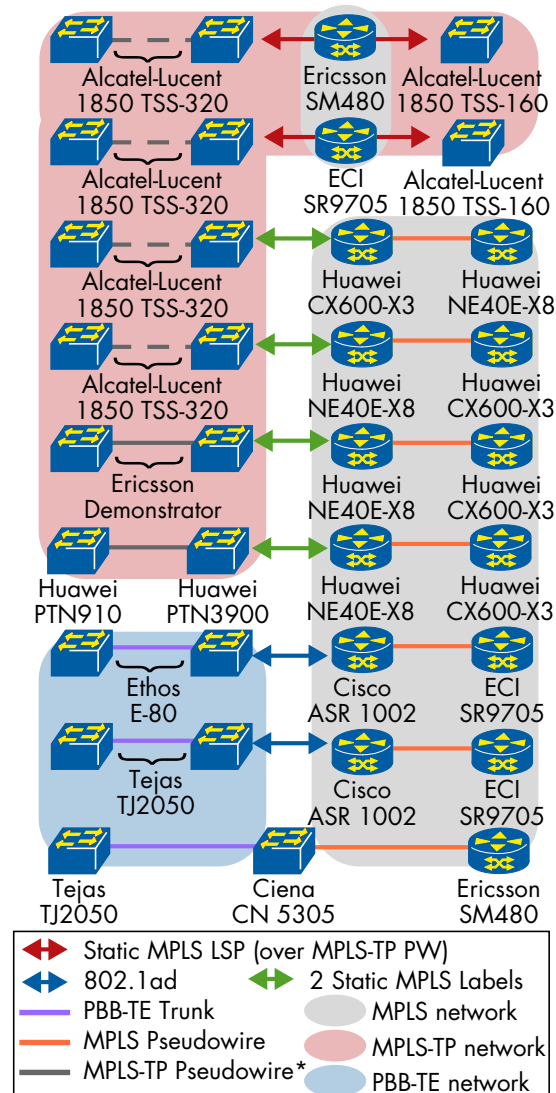


Figure 11: MPLS, MPLS-TP, PBB-TE Interop

* Dashed line - No tunnel label used, leaving one (1) label in MPLS stack.

It is still important to ensure that Carrier Ethernet

services may be realized not only by choosing one of these technologies, but also by a provider who is interested in leveraging multiple. In Figure 11 you can observe the different ways in which the MPLS-TP and PBB-TE devices were configured to interoperate with MPLS equipment.

In the PBB-TE network both Ethos and Tejas equipment terminated PBB-TE trunk into 802.1ad interfaces towards the Cisco ASR 1002 which passed the payload into a pseudowire. Ciena on the other hand, with an implementation for both PBB-TE and MPLS, was able to stitch PBB-TE trunks to MPLS pseudowires, passing the payload from one to the other.

A triple segment pseudowire model was used in the MPLS-TP network. Both Ericsson and Huawei equipment handed off to the MPLS network by means of statically configured Label Switched Path (LSP) and Pseudowire (PW) labels. In these cases the MPLS enabled Huawei NE40E-X8 performed the stitching from the statically configured pseudowire segment to the signaled pseudowire segment.

Alcatel-Lucent also configured the above three-segment approach using MPLS-TP pseudowires without an LSP label in the MPLS-TP part of the network. Additionally, such single-label pseudowires were stitched to statically configured LSPs towards an MPLS router. The MPLS router swapped the outer static label without needing to peer beneath into the pseudowire label. The MPLS-TP nodes on either end of the MPLS router were able to pass OAM over this connection.

MPLS VPLS - BGP Auto-Discovery

With IP/MPLS Provider Edge Equipment (PEs) growing in functionality, the configuration work required is increasing as well. To help reduce configuration efforts when activating new customer connections, some standardization has been made in the IETF's L2VPN working group. The idea is to take advantage of the peering already established with BGP which many PEs will also have configured and leveraging those BGP sessions when setting up VPLS connections. The feature, Auto Discovery (AD), was defined in the VPLS specification which uses BGP for signaling (RFC 4761), but an IETF draft has also specified the use of this feature in LDP signaled VPLS scenarios (RFC 4762) - draft-ietf-l2vpn-signaling, which is currently in RFC Editor Queue, last edited in 2006).

The Cisco 7606, Cisco 7604, and ECI SR9705 successfully discovered each other as VPLS PEs through BGP, where LDP was used to signal the pseudowires. Using BGP signaled pseudowires, the Huawei CX600-X3, Huawei NE40E-X8, and Ixia IxNetwork emulating a VPLS PE also successfully discovered each other as a VPLS PE. Following successful discovery, both scenarios were verified by passing Ethernet traffic.




Figure 12: MPLS BGP AD

MPLS Transport Profile Unfolding

Currently the IETF and ITU-T are working together in a ground breaking liaison to define an MPLS Transport Profile (MPLS-TP). The joint effort has already passed their first RFC document, RFC 5586, which specifies an IP-less mechanism for identifying MPLS and pseudowire (PW) maintenance and management packets. The working group strives to define the framework and requirements for MPLS-TP, standardizing the technology to fulfill the needs of current transmission networks without changing the forwarding rules of MPLS.

MPLS-TP Channels

The latest framework draft proposes multiple services: Virtual Private Wire Service (VPWS), Virtual Private LAN Service (VPLS), Virtual Private Multicast Service (VPMS), and IP-only LAN-like Service (IPLS) that can be provided using the PW client. The following devices successfully forwarded point-to-point Ethernet traffic over MPLS-TP channels: Alcatel-Lucent 1850 TSS-160, Alcatel-Lucent 1850 TSS-320, Ericsson Demonstrator, and the Huawei PTN 3900. The channels were verified to use the encapsulation defined by RFC 5586 for their respective OAM mechanisms. The RFC generalizes the applicability of the Associated Channel Header (ACH) defined in RFC 4385, enabling the realization of a Generic Associated Channel Header (G-ACH) under which to run different OAM mechanisms. The Generic Associated Channel Label (GAL) is used to recognize the presence of the ACH.

Proposed OAM and Protection Solutions (individual drafts)

OAM. Recent discussion in the IETF MPLS working group has been focused on creating a consensus around a standard OAM solution for MPLS-TP. Different solutions have been proposed in individual author drafts such as draft-fulignoli-mpls-tp-bfd-cv-proactive-and-rdi and draft-bhh-mpls-tp-oam-y1731, however the working group has not yet accepted these documents as IETF working group drafts. (The status is easy to distinguish; a working group draft

carries “draft-ietf” in its name.)

Therefore, the solutions tested and demonstrated here are based on individual author drafts not yet accepted into the MPLS-TP framework by the working group — some of them may never be accepted, as they are contradictory to some extent. Implementations based on these drafts are all proprietary today; some of them may become blessed by an RFC in the future, some not.

The purpose of our effort is to evaluate the current status of the industry and to examine progress made since our last test in February. The two OAM proposals tested have received substantial vendor and service provider backing, so we were able to stage multi-vendor interoperability tests.

The Alcatel-Lucent 1850 TSS-160, Alcatel-Lucent 1850 TSS-320, and Huawei PTN 3900 successfully tested an interoperable OAM solution based on draft-bhh-mpls-tp-oam-y1731 which proposes modifications to Ethernet OAM tools as defined in ITU-T Y.1731 for use with MPLS-TP OAM. Ericsson demonstrated an OAM solution consisting of the IETF-defined Bidirectional Fault Detection (BFD) protocol found in draft-ietf-bfd-mpls with proposed enhancements based on draft-fulignoli-mpls-tp-bfd-cv-proactive-and-rdi and draft-boutros-mpls-tp-cc-cv. As of IETF 75 at the end of July 2009 it was planned to converge these two enhancement drafts into draft-asm-mpls-tp-bfd-cc-cv, and submit the converged draft for adoption as an IETF working group draft. The enhanced BFD protocol was used to check the continuity (liveliness) of the MPLS-TP LSP established between the Ericsson devices. Both OAM solutions were verified for the use of GAL and G-ACh.

Protection. The Alcatel-Lucent 1850 TSS-320 and 1850 TSS-160 both with the Huawei PTN 3900 tested an interoperable protection solution, referencing the individual draft draft-weingarten-mpls-tp-linear-protection. Ericsson also demonstrated this protection between Ericsson equipment. The draft specifies the mechanisms needed to coordinate the protection switching on top of the previously mentioned OAM solutions. Out of service times amongst the link failure tests ranged from 7 to 24 milliseconds upon link failure and from less than 1 up to 3 milliseconds upon link recovery. All three vendors also demonstrated such switching via manual commands as opposed to an automatic action as a reaction to a link failure, which is also specified in the author draft. The Alcatel-Lucent 1850 TSS-320 Huawei PTN 3900 additionally demonstrated the ability to switch to a backup path when the far end OAM source was considered invalid based on a mismatch of OAM configurations.

Provider Backbone Bridging Traffic Engineering (PBB-TE)

We have been testing PBB-TE here at our interop events since before the standard received its official IEEE name (back then it was called Provider Backbone Transport (PBT)). Now a completed

standard (IEEE 802.1Qay), three vendors have participated in this year’s interoperability testing with solutions.

The devices involved in PBB-TE testing included the Ciena CN 5305, Ciena CN 3960, Ciena 3940, Ethos E-80, Tejas TJ2050, Tejas TJ2030, and Ixia IxNetwork where Ixia emulated a PBB-TE node transmitting PBB-TE encapsulated frames. All equipment was involved in the establishment of multi-vendor PBB-TE trunks in order to forward Ethernet frames. CFM CCMs as specified by IEEE 802.1ag were used by all devices to monitor the liveliness of a trunk. The standard Ethertypes were also used by all vendors for Ethernet frame encapsulation as well as CFM encapsulation. Additionally, all vendors participated in tests where connectivity was established with the MPLS network, as seen in Figure 11.

Protection. Equipment from all three vendors successfully tested PBB-TE protection across primary and backup tunnels consisting of multiple vendors. For each test, a primary and backup PBB-TE trunk was configured for a particular service. A link not physically connected to both end nodes was then disconnected, and each end node was expected to recognize the failure through a loss of three CCMs. A test was considered a success when the out of service time was under 50 milliseconds, which was the case for all tests. Failover times observed spanned from 15 to 42.5 milliseconds, and recovery tests were hitless.

Ethernet Ring Protection Switching

Carrier Ethernet access networks are often constructed in ring topologies in order to provide resilient path protection while using less cabling resources than other network topologies. The ITU-T has defined Ethernet Ring Protection Switching (ERPS - ITU-T G.8032) in order to provide resiliency mechanisms for such topologies using the defined Automatic Protection Switching (APS) protocol.

We tested a multi-vendor six node ring consisting of the NEC Pasolink NEO TE, RAD ETX-202, and Tejas TJ2050. When the tests were first configured there were inconsistent expectations of how the Maintenance Entity group Level (MEL) values were set, however this issue was quickly resolved through configuration.

We performed several link failure test runs where both NEC and Tejas functioned as the RPL owner. We measured out of service times between 14 and 26 milliseconds in case of link failure and 3 to 14 milliseconds for recovery.

In addition, we successfully tested a protected Ethernet ring between NEC and Tejas using an NEC Pasolink microwave link. A failure of the air interface was recognized by NEC equipment when reception of Continuity Check Message (CCM) frames ceased. This then triggered the ERPS fail-over procedure. It took 20 milliseconds for the ring to converge after the air interface failed and 3 milliseconds to switch back.

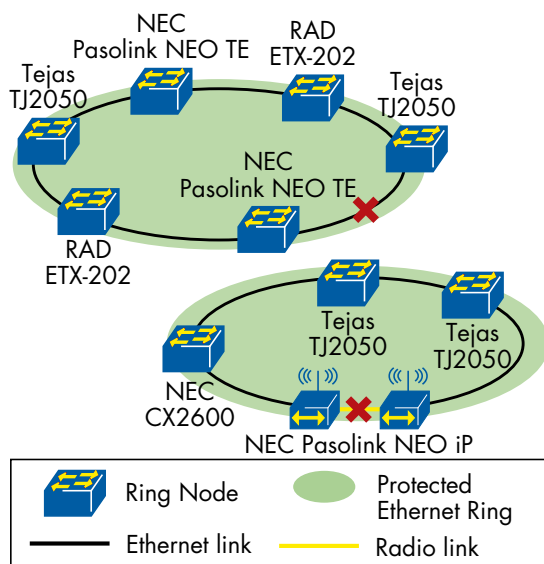


Figure 13: ERPS

Tejas also demonstrated an open ring topology planned for standardization proposal. The ring, consisting of four Provider Backbone Bridging (PBB) nodes, is loosely based on ITU-T G.8032 ERPS however removing the requirement of a closed ring. The open ring created a dual homed connection to a VPLS service in the MPLS network via the Huawei NE40E-X8 and ECI SR9705. Failure times ranged from 4 to 8.6 milliseconds, and recovery times peaked at 4 milliseconds.

Acknowledgements

EANTC would like to acknowledge and thank Kari Asheim and Terje Kroghdahl from nLogic, Oslo, as well as Stephen Murphey and Thomas Peterson from the University of New Hampshire InterOperability Lab (UNH-IOL) for their excellent support during the hot staging.

Thanks to three T5C-48T switches provided by Telco Systems, management connectivity was extended throughout the 15 racks filled by the hot staging equipment in the lab and the work spaces provided.

Editors. This document has been edited by Jambi Ganbar, Sergej Kälberer, Jonathan Morin, Carsten Rossenhövel, Thomas Sladek, and Xiao Tai Yu.

REFERENCES

"Ethernet Services Attributes Phase 2", MEF 10.1
 "Ethernet Service Definitions", MEF 6
 "Implementation Agreement for the Emulation of PDH Circuits over Metro Ethernet Networks", MEF 8
 "Ethernet Local Management Interface", MEF 16
 "User Network Interface (UNI) Type 2 Implementation Agreement", MEF 20
 "External Network Network Interface (ENNI) - Phase 1", Draft 7.1
 "Provider Bridges", IEEE 802.1ad
 "Provider Backbone Bridges", IEEE 802.1ah
 "Provider Backbone Bridges Traffic Engineering", IEEE 802.1Qay

"Media Access Control (MAC) Bridges", IEEE 802.1D
 "Media Access Control (MAC) Bridges Amendment 2: Rapid Reconfiguration", IEEE 802.1W
 "Virtual Bridged Local Area Networks", IEEE 802.1Q
 "Link Aggregation Control Protocol", IEEE 802.3ad
 "Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications. Amendment: Media Access Control Parameters, Physical Layers, and Management", IEEE Std 802.3ah
 "Virtual Bridged Local Area Networks - Amendment 5: Connectivity Fault Management", IEEE 802.1ag
 "OAM Functions and Mechanisms for Ethernet Based Networks", ITU-T Y.1731
 "Ethernet ring protection switching", ITU-T G.8032
 "Segmented Pseudo Wire", draft-ietf-pwe3-segmented-pw
 "Provisioning, Autodiscovery, and Signaling in L2VPNs", draft-ietf-l2vpn-signaling
 "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761
 "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762
 "Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks", RFC 4717
 "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
 "The Control of Jitter and Wander within Digital Networks which are Based on the 2048 kbit/s Hierarchy", ITU-T G.823
 "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE 1588-2008
 "Timing and synchronization aspects in packet networks", ITU-T G.8261/Y.1361
 "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
 "Timing characteristics of primary reference clocks", ITU-T G.811
 "Timing requirements of slave clocks suitable for use as node clocks in synchronization networks", ITU-T G.812
 "Timing characteristics of SDH equipment slave clocks (SEC)", ITU-T G.813
 "MPLS-TP Linear Protection", draft-weingarten-mpls-tp-linear-protection
 "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447
 "BFD For MPLS LSPs", draft-ietf-bfd-mpls
 "MPLS-TP BFD for Proactive CC-CV and RDI", draft-fulgnoli-mpls-tp-bfd-cv-proactive-and-rdi
 "Connection Verification and Continuity Check for MPLS Transport Profile Label Switched Path", draft-boutros-mpls-tp-cc-cv
 "MPLS-TP OAM based on Y.1731", draft-bhh-mpls-tp-oam-y1731

ACRONYMS

Term	Definition
AC	Access Circuit
ACH	Associated Channel Header
AIS	Alarm Indication Signal
APS	Automatic Protection Switching
ATM	Asynchronous Transfer Mode
BITS	Building Integrated Timing Supply
BFD	Bidirectional Fault Detection
BGP	Border Gateway Protocol
CCM	Continuity Check Message
CE	Customer Edge
CE-VLAN	Customer Edge Virtual LAN
CES	Circuit Emulation Service
CFM	Connectivity Fault Management
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
DSL	Digital Subscriber Line
E-LAN	A multipoint-to-multipoint Ethernet service. A LAN extended over a wide area
E-Line	Point-to-Point Ethernet Service similar to a leased line ATM PVC or Frame Relay DLCI
E-LMI	Ethernet Local Management Interface
ENNI	External Network Network Interface
EFM	Ethernet in the First Mile
ERPS	Ethernet Ring Protection Switching
EVC	Ethernet Virtual Connection
ESMC	Ethernet Synchronization Messaging Channel
G-ACh	Generic Associated Channel Header
GAL	Generic Associated Channel Label
GSM	Global System for Mobile
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ITU-T	International Telecommunication Union Telecommunication Standardization Sector
LDP	Label Distribution Protocol
LMM	Loss Measurement Message
LOS	Loss Of Signal
LSP	Label Switched Path
LTE	Long-Term Evolution
MAC	Media Access Control
MBMS	Multimedia Broadcast Multicast Services, integral part of LTE
MEG	Maintenance Entity Group
MEL	Maintenance Entity group Level

Term	Definition
MEP	Maintenance Association End Point
MEF	Metro Ethernet Forum
MIP	Maintenance Association Intermediate Point
MIMO	Multiple-Input and Multiple-Output
MPLS	Multi-Protocol Label Switching
MPLS-TP	MPLS Transport Profile
MTIE	Maximum Time Interval Error
OAM	Operations, Administration and Maintenance
OPEX	OPERating EXpenditure
OSPF	Open Shortest Path First
PBB	Provider Backbone Bridge
PBB-TE	Provider Backbone Bridge Traffic Engineering
PCM30	Pulse-Code-Modulation, 30 of 32 E1 slots used for data
PCM31	Pulse-Code-Modulation, 31 of 32 E1 slots used for data
PE	Provider Edge
PDU	Protocol Data Unit
PHY	PHYsical layer
ppb	parts-per-billion
ppm	parts-per-million
PSN	Packet Switched Network
PTP	Precision Time Protocol
PW	PseudoWire
QAM	Quadrature Amplitude Modulation
RDI	Remote Defect Indication
RFC	Request For Comments
RSVP-TE	Resource reSerVation Protocol Traffic Engineering
RTP	Real Time Protocol
SEC	SDH Equipment slave Clocks
SDH	Synchronous Digital Hierarchy
S-Tag	Service Tag
SAToP	Structure-Agnostic Time Division Multiplexing (TDM) over Packet
SLA	Service Level Agreement
SyncE	Synchronous Ethernet
TDEV	Time DEVIation
TDM	Time Division Multiplexing
TIE	Time Interval Error
UMTS	Universal Mobile Telecommunications System
UNI	User Network Interface
VLAN	Virtual Local Area Network
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
VPWS	Virtual Private Wire Service



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



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 © 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.

v1.1