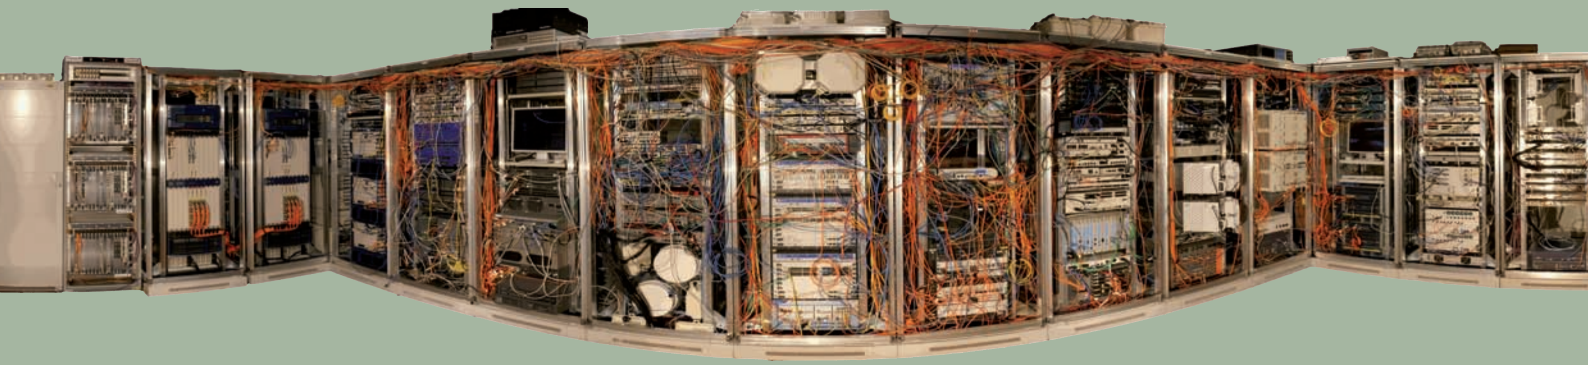

Carrier Ethernet Services - The Future

Public Multi-Vendor Interoperability Test

Singapore, November 2008

cew APAC
Carrier Ethernet World

EANTC

Editor's Note

Carsten Rossenhoevel
Managing Director

Carrier Ethernet service deployments are rapidly taking up speed in many Asia Pacific countries — Carrier Ethernet solutions are well positioned for the large-scale, broadband service provider networks in the region.

Many (if not most) end-to-end Carrier Ethernet networks employ equipment from multiple vendors today. Service providers typically implement a dual-vendor strategy separately for core backbone, metro aggregation and access components. Robust interoperability across the board for all vendors and types of equipment is an important cornerstone of Carrier Ethernet's success story.

For the first time in Singapore, EANTC showcases the results of a large-scale live interoperability test focusing all aspects of end-to-end Carrier Ethernet networks.

More than 30 engineers from ten vendors with over 35 systems participated in the hot staging preparation of this test at EANTC's lab in Berlin, Germany, in August.

The participating vendors verified more than 20 test areas in any-to-any combinations in ten days. Carrier Ethernet implementations support more functions and cover more markets today — ranging from core to microwave to access, E-Lines to E-Tree and triple play to mobile backhaul.

It was a great experience to witness the testing session, a unique get-together of many leading players with one single goal: To improve multi-vendor interoperability of advanced Carrier Ethernet implementations.

Interestingly, market forces are operating at full strength. This year, we once again tested three transport technologies in the test event's metro/aggregation networks: MPLS, PBB-TE and T-MPLS. These three compete to some extent — at our test, they all proved being well suited for the transport of Carrier Ethernet services. We noticed that service OAM support is becoming mandatory for aggregation and CPE devices. In addition we saw substantial advances in Ethernet microwave solutions.

This white paper summarizes in detail the monumental effort that the participating vendors and EANTC team underwent. Enjoy the read.

Table of Contents

Introduction

The interoperability event focused on the Future of Carrier Ethernet Services. EANTC has been conducting such events at the Carrier Ethernet World Congress in Europe since its inception in 2005 and while each previous event concentrated on specific topics such as mobile backhaul or service creation, this event aimed to congregate the knowledge and experience the industry gained in the last four years into a single modern, converged network thus demonstrating the state of the art for all technologies that a tier-one service provider is likely to encounter. We therefore tested:

- Converged residential and business services realized using E-Line, E-LAN and for the first time E-Tree services
- The leading access, transport and aggregation technologies
- Microwave access and transport
- Ethernet OAM: Fault management and performance monitoring
- High availability
- Management and SLA reporting

EANTC's interoperability showcases are driven by three main goals:

Technical – Through participation in the event, vendors have the opportunity to verify the interoperability of their devices and protocol implementations against the majority of the industry's leading vendors.

Marketing – The participants can showcase the interoperability of their latest solutions on a unique, large-scale platform.

Standards – When fundamental issues are found during the hot staging event EANTC reports the discoveries to the standard bodies. These in turn update the standards.

The test plan was created by EANTC based on the test topics suggested by vendors, our service provider panel and experience gained from previous events, and was lined up with recent IEEE, IETF, ITU-T and MEF standards.

Service Provider Test Plan Review

The test plan was reviewed by a panel of global service providers in July this year. Their feedback and comments were reflected in the final version of the test plan. EANTC and the participating vendors would like to thank: COLT, GVT Brazil, PT TELKOM Indonesia, T-Systems and Metanoia Inc.

Participants and Devices

Vendor	Participating Devices
Alcatel-Lucent	1850 TSS-40 5650 CPAM 7450 ESS-6 7705 SAR 7750 SR7 9500 MPR
ECI Telecom	SR9705
Ericsson	Marconi OMS 2400
Harris Stratex Networks	Eclipse (Gigabit) Radio
Ixia	XM2 IxNetwork
Juniper Networks	M10i MX240 MX480
Nortel	Metro Ethernet Routing Switch (MERS) 8600
Spirent Communications	Spirent TestCenter
Tejas Networks	TJ2030
Tellabs	8830 Multiservice Router

Network Design

As in previous events we set off to construct a network that would allow all participating vendors to establish end-to-end Ethernet services with any of the other vendors. One of the central design considerations for the network was to enable any device positioned in the access network to reach any other access network device regardless of the other device's point of attachment to the network.

We also aimed to build a network that would look familiar to service providers. It is perhaps unrealistic to expect that service providers will incorporate all current transport technologies into their network. Nevertheless the familiar network domains are likely to exist: access, aggregation, metro and core, regardless of the chosen transport technology. It is also realistic to expect service providers to use MPLS in the core of their network.

Looking at the network from a customer's perspective, we used the following network areas:

- *Access*: The devices that normally exist at the customer premise were positioned here. We were lucky to see a diverse number of access technologies for transporting Ethernet such as microwave links, copper, and fiber. These devices implemented the UNI-C construct as defined by the MEF.
- *Aggregation:* The aggregation area of a network consisted of a variety of solutions meant to aggregate customer premise devices. This included Provider Bridges and H-VPLS Multi-Tenant Unit Switches (MTU-s). When applicable these devices performed the UNI-N role in the network.
- *Metro*: Three different transport technologies were used in each of the three metro area networks: MPLS, PBB-TE and T-MPLS. This allowed each transport technology to test its own resiliency and Network-to-Network Interface (NNI) solutions.
- *Core*: As stated above, IP/MPLS was used to support connectivity between the different metro area networks in order to realize end-to-end services. In addition, MPLS Layer 3 VPNs as defined in RFC 4364 were tested in the core of the network.

The physical network topology presented here depicts the roles of all the devices and their respective placement in the network. Please note that many tests required logical connectivity between the devices, often at an end-to-end nature, which will be shown, where applicable, using logical topologies in each test section.

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

Please note that we use the term »tested« when reporting on multi-vendor interoperability tests. The term »demonstrated« refers to scenarios where a service or protocol was terminated by equipment from a single vendor on both ends.

Ethernet Service Types

The Metro Ethernet Forum (MEF) has defined three Ethernet service types in order to allow the industry and specifically the customers interested in the services to have a common language to discuss such Ethernet based services. The three service types are defined in terms of the Ethernet Virtual Connection (EVC) construct:

- E-Line – Point-to-point EVC
- E-LAN – Multipoint-to-multipoint EVC
- E-Tree – Rooted-multipoint EVC

While the E-Line service type provides a service to exactly two customer sites, the E-LAN and E-Tree service types allow the connection of more than two customer sites. In contrast to the E-LAN service type which allows an any-to-any connectivity between customer sites, E-Tree introduces two different roles for customer sites: leaf and root. An E-Tree service facilitates communication between leaves and roots, however, leaves can not communicate with each other directly. An E-Tree service implemented by a rooted-multipoint EVC can be used to provide multicast traffic distribution and hub-and-spoke topologies (e.g. DSL customers to BRAS).

In the test network we instantiated three specific definitions of service types: Ethernet Virtual Private Line (EVPL), Ethernet Virtual Private LAN (EVP-LAN), and Ethernet Virtual Private Tree (EVP-Tree). All services were configured manually in the network. Due to the large number of devices present at the hot staging, this process was time consuming and prone to mistakes. A multi-vendor provisioning tool would have been ideal for the testing and is recommended for any service provider planning to deploy Carrier Ethernet services.

The services created in the network were configured in two ways:

- EVCs that remained within the same metro area network
- EVCs that crossed the network core

The sections below describe the services in the network in detail.

E-Tree

For the first time at an EANTC interoperability event, an E-Tree service instantiation was established. One EVP-Tree was configured with one root node within the MPLS metro area and leaves throughout all network areas. The MEF defines an E-Tree service to be a rooted Ethernet service where the roots are able to communicate with all leaves, and all leaves are able to communicate with the roots, but not with each other.

This service utilized each metro technology in a unique way. The MPLS metro and core each used a separate VPLS instance to create this service, using different split horizon groups to ensure that leaf UNIs could only communicate with the root UNI, but could not establish communication between each other. The PBB-TE and T-MPLS metros respectively configured PBB-TE trunks and TMCs to a static rule set in order to propagate the tree service.

Devices included in this test were Alcatel-Lucent 7750 SR7, ECI SR9705, Ericsson Marconi OMS 2400, Harris Stratex Eclipse, Juniper MX480, Nortel MERS 8600, Tejas TJ2030, and Tellabs 8830.

E-LAN

One EVP-LAN was configured in the network with customer ports in all three metro areas. The construction of the EVP-LAN service used different mechanisms in each metro area. These mechanisms are described in details in the diverse transport section.

E-Line

The E-Line service type configured in the network used Virtual LAN (VLAN) IDs to distinguish between the various services. In some cases, much like real world networks, a switch positioned at the customer site would add a Service VLAN tag (S-VLAN) to the Ethernet traffic provided by the customer, therefore, allowing the customer to maintain its private VLAN addressing scheme and separate the customer VLAN space from the provider's.

All vendor devices successfully participated in creation of E-Line services. From the number of combinations tested, we are confident that an any-to-any combination of endpoints is possible.

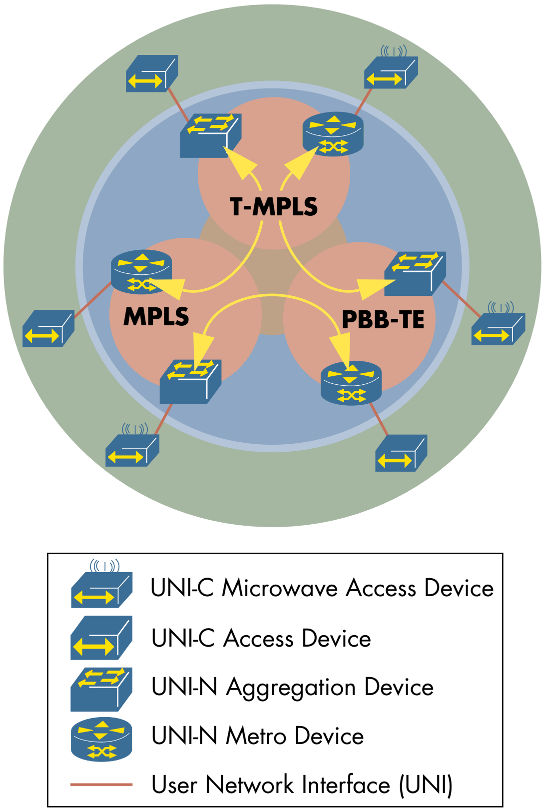

Figure 1: E-Line service creation

Diverse Access Technologies

The different services in the test network used a variety of access technologies to reach the simulated last mile customer access device. Most services used fiber (multi-mode) and copper based Gigabit Ethernet.

Several Ethernet access links comprised of two Ethernet links with a microwave signal in between. These systems are described in more detail below.

Microwave for Access and Transport

In recent years we have seen an increased interest in our interoperability events from vendors offering microwave connectivity and network solutions. Microwave solutions alleviate the need to roll out physical wire infrastructure and are especially prevalent in such areas as cellular backhaul, emerging markets, large corporation networks, hospitals, and mobile-fixed operators. This event enjoyed the participation of the Alcatel-Lucent 9500 MPR and Harris Stratex Eclipse.

Since the radios rely on a signal through the air some weather events such as rain and heavy fog can cause the signal to degrade effectively decreasing the range or capacity of the link. Radio devices can recognize the decrease in air-link capacity and some solutions can distinguish which frames should be prioritized and further transported versus which frames will be dropped. The Alcatel-Lucent 9500 MPR and Harris Statex Eclipse showed this functionality by decreasing the modulation scheme in Quadrature Amplitude Modulation (QAM) which caused traffic loss only to best effort frames and no or minimal loss to prioritized traffic with unaffected latency.

Services relying on microwave equipment will need to be made aware when the microwave signal is too weak to transmit traffic. The link state propagation function disables the Ethernet link state for all ports associated with the microwave link. This functionality was demonstrated by the Harris Stratex Eclipse. This device also showed the capability to propagate an incoming loss of signal on a tributary Ethernet port across the microwave link and switching off the appropriate physical port on the other side of the radio connection.

Diverse Transport

The Carrier Ethernet architecture specified by the MEF is agnostic to the underlying technology used to provide Carrier Ethernet services. The creation and support of such services is, however, an essential component of the interoperability test event. Mainly three technologies compete for Carrier Ethernet Transport: MPLS, PBB-TE and T-MPLS. During this event we had the opportunity to verify all three technologies. The following sections describe test results for each technology in detail.

MPLS

MPLS is defined in a set of protocols standardized by the Internet Engineering Task Force (IETF) and the IP/ MPLS Forum. MPLS is positioned to deliver layer 2 and layer 3 services including Ethernet services as defined by the MEF while being agnostic to the underlying transport technology.

The tests in this area were based on previous experience gained from EANTC's Carrier Ethernet World Congress and MPLS World Congress interoperability test events. The MPLS metro domain operated independently from the MPLS core network.

The MPLS metro network was built solely for the purpose of Carrier Ethernet services. Multipoint-to-multipoint services were facilitated with the creation of a single Virtual Private LAN Services (VPLS) instance utilizing both VPLS PEs and H-VPLS MTU switches established between the following devices: Alcatel-Lucent 7450 ESS-6, ECI SR9705, Ixia XM2 IxNetwork, Juniper MX240 and Juniper MX480, and Tellabs 8830. This VPLS instance used LDP for signaling statically configured peers as described in RFC 4762. These devices also established Ethernet pseudowires using LDP to facilitate point-to-point Ethernet services.

In order to test the interoperability of VPLS implementations which use BGP for signaling as described in RFC 4761, another separate VPLS instance was configured. This was demonstrated with BGP-based Auto-Discovery enabled by the Juniper MX240 and Juniper MX480. The Juniper MX480 performed an interworking function between this BGP signaled VPLS domain and an LDP signaled VPLS domain.

Provider Backbone Bridge Traffic Engineering (PBB-TE)

One of the potential solutions to delivering MEF defined services using Ethernet technologies only is the IEEE defined Provider Backbone Bridge Traffic Engineering (PBB-TE). The technical specification is defined in 802.1Qay which is working its way through the standard process and is in draft version 3.0 at the time of the testing. The standard extends the functionality of the Provider Backbone Bridges (802.1ah) adding a connection-oriented forwarding mode by creating point-to-point trunks. These trunks deliver resiliency mechanisms and a configurable level of performance.

The vendors participating in the PBB-TE transport domain included Ixia XM2 IxNetwork, Nortel MERS 8600, and Tejas TJ2030.

In the PBB-TE Metro network we were able to test the establishment of E-Line, E-LAN, and E-Tree services. The establishment of E-Line services was straightforward as we tested it in several previous events. E-LAN and E-Tree services creation was tested for the first time within the PBB-TE cloud. For the E-LAN service the Tejas TJ2030 switch established a bridging instance per PBB-TE trunk and C-VLAN/S-VLAN IDs. The PBB-TE edge devices established a trunk towards this device for each particular UNI. A few issues related to usage of different Ethertype values in CFM messages, padding, and different interpretation of CCM intervals were discovered in the initial configuration phase of PBB-TE trunks, however, these issues were resolved quickly.

In addition, the Nortel MERS 8600 and the Spirent TestCenter tested one of the latest additions to Ethernet - Provider Link State Bridging (PLSB) – a pre-standard implementation of the IEEE 802.1aq (Shortest Path Bridging) which is in draft version 0.3. The protocol uses the IETF defined IS-IS protocol for distributing Backbone MAC addresses and Service IDs of participating nodes across the network. Once the network topology has been learned, IS-IS is used to establish loop-free multipoint-to-multipoint services. The forwarding plane uses PBB (802.1ah), however since the other devices in the PBB-TE network did not support PLSB the three Nortel MERS 8600 devices were able to use a re-encapsulation of either PBB-TE trunks or VLAN tags to peer within the PBB-TE network. The Nortel MERS 8600 devices and the Spirent TestCenter emulated nodes successfully learned the appropriate B-MAC addresses, and forwarded the respective traffic accordingly.

Tejas Networks demonstrated a logical Ethernet LAN network with an IEEE 802.1ad based Ethernet Ring Protection Switching (ERPS). This ring based control protocol being standardized under ITU-T G.8032 is a protection mechanism which offers carriers a deterministic sub-50 ms network convergence on a fiber failure as opposed to the conventional loop-breaking mechanisms like Rapid Spanning Tree Protocol (RSTP). ERPS convergence time is independent to the number of nodes in the network, thereby vastly enhancing the scalability of a carrier network. We measured failover and restoration of <35 ms for the demonstrated ERPS.

T-MPLS to MPLS-TP Migration

Following the approval of the first version of the ITU recommendations on T-MPLS, the IETF and ITU-T jointly agreed to work together to extend MPLS protocols to meet transport network requirements in order to ensure a smooth convergence of MPLS-based packet transport technology. A Joint Working Group (JWT) was formed between the IETF and the ITU to achieve mutual alignment of requirements and protocols and to analyze the different options for T-MPLS standard progress. On the basis of the JWT activity, it was agreed that the future standardization work will focus on defining a transport profile of MPLS (named MPLS-TP) within IETF and in parallel aligning the existing T-MPLS Recommendations within ITU-T to the MPLS-TP work in IETF.

At their Dublin meeting in July 2008, the IETF has initiated the work on MPLS-TP. Due to the fact that IETF MPLS-TP standard or drafts do not exist yet, we tested the implementations based on the T-MPLS ITU-T Recommendations currently in force and its relevant drafts.

It is our intention also to include the first implementations of MPLS-TP drafts at our next event.

Transport MPLS (T-MPLS)

This test marked the third T-MPLS interoperability testing at EANTC. The following devices successfully participated in the T-MPLS area during the event: Alcatel-Lucent TSS-40, Ericsson Marconi OMS 2400, and Ixia XM2 IxNetwork.

The T-MPLS standards specify the networking layer for packet transport networks based on MPLS data plane and designed for providing SONET/SDH-like OAM and resiliency for packet transport networks.

Alcatel-Lucent and Ericsson successfully tested the creation of E-Line, E-LAN and E-Tree services, the last of which was a first at an EANTC interoperability event. Both participants constructed T-MPLS paths (TMP) which are end-to-end tunnels that aggregate T-MPLS channels (TMC) representing the services. The TMPs and TMCs were transported over different physical layer types including 1 Gbit Ethernet, 10 Gbit Ethernet, ITU-T G.709, and SDH STM-16. The Alcatel-Lucent 7705 SAR was used as a non-T-MPLS switch in the aggregation area, interfacing to the T-MPLS domain by means of statically configured MPLS labels.

The E-LAN and E-Tree services were configured using a multipoint architecture similar to VPLS. On one particular E-Line service, both Alcatel-Lucent 1850 TSS-40 and Ericsson Marconi OMS 2400 were able to successfully test Quality of Service (QoS) by distinguishing between three different classes of service within the same Ethernet service and only drop low priority traffic when interfaces were oversubscribed.

Since the T-MPLS standards do not define a control plane protocol, the T-MPLS connections between vendors were manually configured. Ericsson used two proprietary management tools (ENEA and DiToNe) to setup the T-MPLS network configuration on their devices.

MPLS Core

Since MPLS is used by the majority of service providers as core technology it is only logical that when providers add Carrier Ethernet services to their product offering the MPLS core will be used. We followed this approach and used the MPLS core to connect between three different Ethernet transport metro areas. The core area was constructed using the following edge devices, all of which successfully established multiple VPLS domains and Virtual Private Wire Services (VPWS) using LDP for various Ethernet services: Alcatel-Lucent 7750 SR7, ECI SR9705, Juniper M10i, and Tellabs 8830.

In addition to providing transport for Ethernet services, all edge devices in the core established an IP/MPLS L3VPN service using BGP (based on RFC 4364). The Alcatel-Lucent 7750 SR7 terminated Ethernet pseudowires into this VPN providing the potential to offer layer 3 services to customers which are not reachable otherwise.

External Network to Network Interface (E-NNI)

As we described above three different technologies were used in the metro areas. The problem that every service provider then faces is to connect the metro area with the existing network core. In our test network, much like in most service provider networks, the core used MPLS for transport and services. Therefore, we required mechanisms to allow services originating on one metro area to cross the core and be received on other metro areas.

The following subsections describe the specific Network to Network Interface (E-NNI) solutions used in the network.

MPLS Metro Connectivity to the Core

Several options exist to allow connectivity between two MPLS areas. The preferred options were MPLS based, but one option used IEEE 802.1ad Provider Bridging tags, or simply 802.1Q VLAN tags to transport services between the two areas. The Label Edge Router (LER) in the MPLS core would strip the MPLS header from traffic before it forwarded the Ethernet frames to the LER in the MPLS metro. The S-Tag or VLAN tag would then signal to the MPLS metro device which pseudowire to forward the frames onto. Devices which used this method included the Alcatel-Lucent 7450 ESS-6, ECI SR9705, Juniper M10i and Juniper MX480.

The other option used to connect between administratively separated MPLS core and metros is referred to as pseudowire (PW) stitching. This involves the creation of two pseudowires, one in each domain, and then interconnecting them either within one device, or with a third pseudowire between the two edge devices. Vendors who took this approach chose the latter. In this case two MPLS labels must be signaled: the inner label (PW label) signaled by LDP, and the transport label (PSN label) signaled by either eBGP (IPv4+label) or LDP. To facilitate the transmission of LDP sessions, either a separate OSPF area

Physical Network Topology

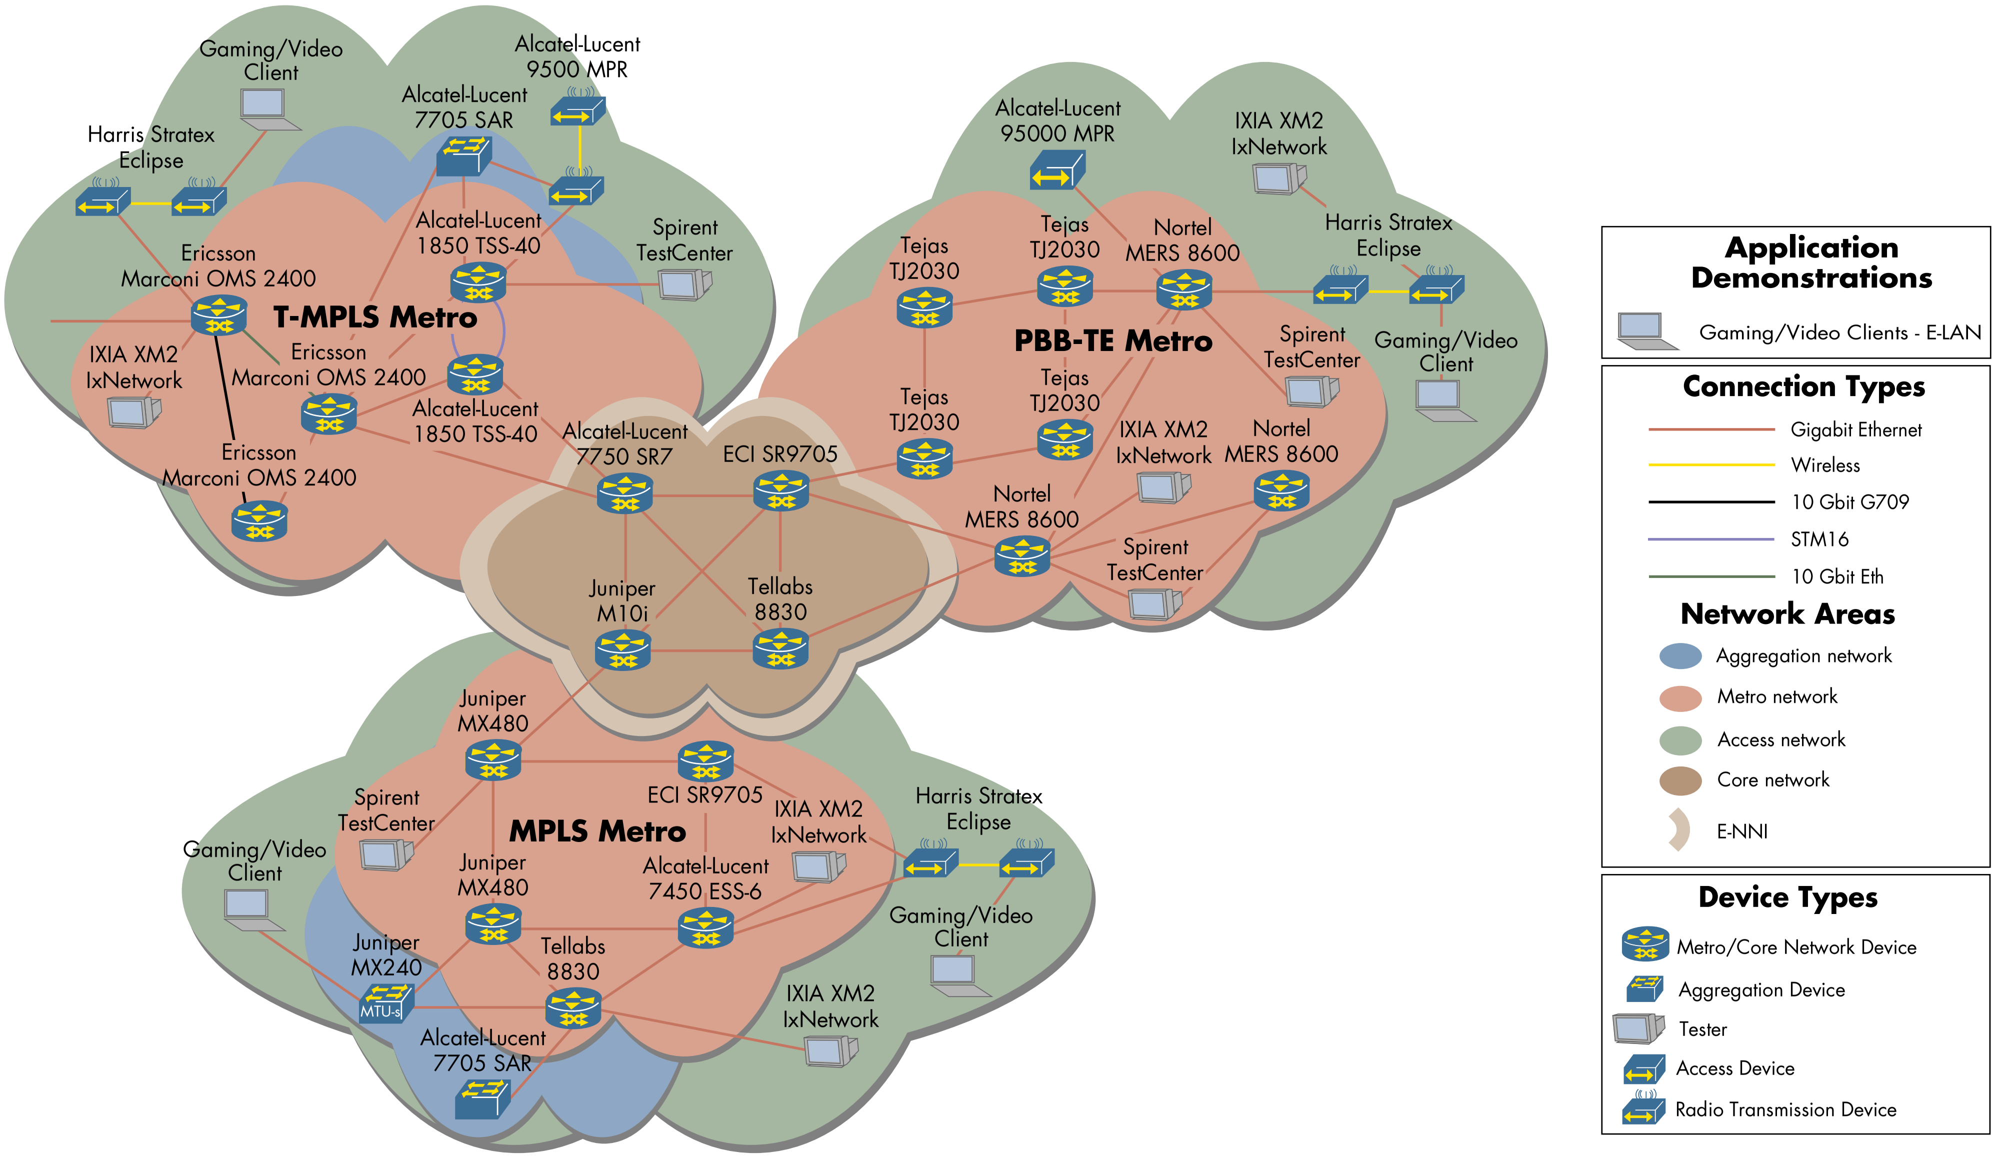

was enabled between the two edge devices or a static route was used. Juniper M10i and Juniper MX480 took the PW stitching approach.

PBB-TE Connectivity to the Core

As PBB-TE and 802.1ad are both part of the IEEE Provider Bridging domain of technologies, it is not surprising that Provider Bridging tags were supported across the board in the PBB-TE metro domain. All services crossing the core into the PBB-TE cloud used S-Tags (Service Tags) to distinguish each service between a core edge router and a PBB-TE switch. These devices included ECI SR9705, Nortel MERS 8600, Tejas TJ2030, and Tellabs 8830. One PBB-TE trunk was configured between Nortel MERS 8600 and Tejas TJ2030 and traversed the MPLS core.

T-MPLS Connectivity to the Core

Two options were used to establish services over an MPLS core into a T-MPLS network. The first was to use 802.1ad S-Tags, similarly to the MPLS metro.

The second option used was pseudowire stitching. The T-MPLS edge device terminated TMCs coming from the edges of the network and stitched them to an MPLS Ethernet pseudowire which was established with the neighboring core edge device. This was done using LDP for both MPLS labels, which ran over a separate OSPF area than the core. This option was tested between Alcatel-Lucent 1850 TSS-40 and the Alcatel-Lucent 7750 SR7, which also had some services configured over statically configured MPLS pseudowires.

Ethernet OAM

EANTC interoperability events have integrated Ethernet Operations, Administration and Management (OAM) testing since 2006. Several different protocols fall under the category of *Ethernet OAM*. In this event we have tested Ethernet in the First Mile (EFM), and Connectivity Fault Management (CFM), standardized by the IEEE under 802.3ah and 802.1ag respectively, Y.1731, standardized by ITU-T, and E-LMI, specified by MEF. Over the past years we have seen a significant increase in support in this area – in the first event four vendors tested their CFM implementations and five vendors tested their EFM code. At our current event almost all participating vendors tested their EFM and CFM implementations.

Our service provider panel placed a high value on both Ethernet OAM test areas and with the inclusion of both EFM and CFM in the new MEF 20 “User Network Interface (UNI) Type 2 Implementation Agreement” technical specifications a clear continuous need for testing has been established

Link OAM

Link OAM is the name used by the Metro Ethernet Forum (MEF) to refer to clause 57 of the IEEE 802.3 standards where OAM is defined for Ethernet in the First Mile. The protocol monitors the health and operations of the UNI’s physical layer. The MEF requires the usage of Link OAM between the UNI-N and UNI-C starting from UNI type 2.2 and recommends the use of the protocol starting from UNI type 2.1.

The idea behind the protocol is to assist operators in localizing last mile physical layer errors. The most logical location in the network where the protocol could be used is between the Provider Edge device (a router or a switch) and the customer premises device. Both devices can identify themselves to each other and communicate problems with each other. The Provider Edge device can use a palette of tools to verify that it can reach the customer premises device therefore saving the provider the extreme cost of sending engineers to the customer site to verify physical link issues.

The following products successfully participated in the Link OAM tests discovering and setting each other in loopback mode: Alcatel-Lucent 7450 ESS-6, Ixia XM2 IxNetwork, Juniper MX240, Spirent TestCenter, and Tellabs 8830.

We performed an additional test within this area verifying that a customer premises device is able to notify the Provider Edge device that the loss of signal is due to the customer premise device being shut down. The message carried in the Link OAM frame is aptly called Dying Gasp. Juniper’s MX480 and MX240 demonstrated their ability to receive Dying Gasp messages and notify the operator using CLI flags that the customer device has been shut down.

Service OAM

The IEEE 802.1ag standard defines end-to-end Ethernet based OAM mechanisms which are referred to by the MEF as Service OAM. The support for Service OAM is mandatory starting from UNI type 2.1. In contrast to link OAM the major use of CFM for service providers and enterprises is to verify connectivity across different management domains. A carrier can define one level for internal use while allowing their customers to verify end-to-end connectivity over the network using a different CFM level.

Since the 802.1ag standard has been published in December 2007 this has been the first interoperability event where we could specify a finished version of the standards for testing which simplified the testing.

For this interoperability event we added a test at the request of the participating vendors to the Service OAM area: Ethernet Remote Defect Indication (ETH-RDI). This message type, integrated into the Continuity Check Messages (CCMs), communicates to a remote Maintenance End Point (MEP) that a defect condition has been encountered. When a defect condition is encountered on a MEP, that MEP will set the RDI bit in CCMs for all Maintenance Entity Group (MEG) levels affected. When the defect condition is repaired or removed, the MEP will reset the RDI bit in the appropriate CCMs. This message type is particularly useful to notify a remote MEP about MAC layer problems, changes to the configurations and such error conditions as sudden unidirectional connectivity.

The following vendors successfully tested Service OAM’s Continuity Check, Linktrace and Loopback features as well as Remote Defect Indication: Alcatel-Lucent 7450 ESS-6, ECI SR9705, Ixia XM2 IxNetwork, Juniper MX240 and Juniper XM480, Nortel MERS 8600, Spirent TestCenter, and Tellabs 8830.

Performance Monitoring

The ITU-T specification Y.1731 defines two message types for calculating loss and delay between two Ethernet OAM endpoints. Loss Measurement Messages (LMMs) are sent which should be replied to on arrival with Loss Measurement Replies (LMRs). The initiating device can make a calculation of average loss by comparing the number of frames sent by the initiating end with those received at the far end. Delay is measured in a similar way, calculating the time it takes for a Delay Measurement Reply (DMR) to come back for each Delay Measurement Message (DMM). Similar to CFM, this tool helps both customers to learn about the service they are receiving, and operators to learn about the service they are providing. Both functionalities were demonstrated by the Nortel MERS 8600.

Resilience and Fault Detection

One of the key aspects of carrier grade transport technologies is SONET/SDH-like resiliency mechanisms and fault detection. Based on historical precedence, SONET/SDH's 50 ms restoration capabilities has been used as the measurement stick for all transport technologies to follow and these days, with Voice over IP (VoIP) and Mobile Backhaul we see a real need for 50 ms restoration time.

The sections to follow depict several failover and restoration tests that were performed during the event. We tested each transport technology's resilience protocols interoperability and augmented these by native fault detection mechanisms (CFM for PBB-TE and T-MPLS and Bidirectional Fault Detection (BFD) for MPLS). We focused our Link Aggregation tests on the UNI therefore providing a complete end-to-end resilience and fault detection story.

Link Aggregation

Several physical Ethernet ports of Ethernet switches can be bundled into a single logical Ethernet interface. This capability is used to increase the amount of bandwidth and/or to provide a link protection mechanism. The capability is specified in IEEE 802.3 clause 43, and commonly known as 802.3ad or just LAG (Link Aggregation Group). In order to create or change the members of a LAG group dynamically, Link Aggregation Control Protocol (LACP) has been standardized and is used between two Ethernet devices providing LAG.

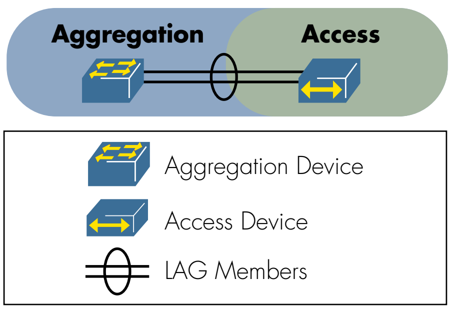

Figure 2: LACP Test Setup

The MEF specified LAG as a protection mechanism on the UNI. Devices supporting MEF UNI type 2.2 must support LAG and LACP which led us to require, as opposed to previous years, that vendors participating in this test will use LACP and not the static configuration of link aggregation groups.

As the test was defined for devices that implement UNI functionality we separated the participants into UNI-C and UNI-N devices. At the UNI-C were Ixia XM2 IxNetwork and Spirent TestCenter, while on the UNI-N Alcatel-Lucent 7450 ESS-6 and ECI SR9705. We first verified that the link aggregation groups were established between the two devices and then measured the failover time by pulling out of the primary link, followed by a measurement of the restoration time by reconnecting the primary link. In almost all cases we measured failover time below 31 ms, and restoration time below 25 ms. In one case we measured failover time of 550 ms and around two seconds restoration time.

Apart from the UNI, Link Aggregation can be used in the core of the network at the External Network to Network Interfaces (E-NNI) or to provide added link capacity. In addition to the LAG tests at the UNI, we performed some LAG tests between directly connected devices within the metro and core area networks with the following: Alcatel-Lucent 7750 SR7, Ericsson Marconi OMS 2400, Juniper MX480, and Spirent TestCenter.

MPLS Protection

Several resilience mechanisms exist in MPLS and most have been tested in previous EANTC interoperability events. MPLS Fast Reroute (FRR) is one such mechanism that was tested in 2006 and 2007 (the white papers are available on EANTC's web site).

MPLS Fast Reroute provides link and node protection for Label Switched Paths (LSPs), therefore, protecting the services running within the LSPs. The following devices participated in the role of the Point of Local Repair (PLR): Alcatel-Lucent 7450 ESS-6, Alcatel-Lucent 7750 SR7, and ECI SR9705. The Merge Point (MP) is the router which merges the backup and primary segments of an MPLS tunnel. Alcatel-Lucent 7450 ESS-6 and Tellabs 8830 both participated as MPs. Failover times reached as low as 19 ms. The importance of the test, however, was the protocol interoperability between the PLRs and the MPs therefore we spent no time on trying to reconfigure the devices for faster failover times.

In two MPLS Fast Reroute test combinations the tunnel failure was recognized by Bidirectional Fault Detection protocol (BFD). BFD allows the routers to

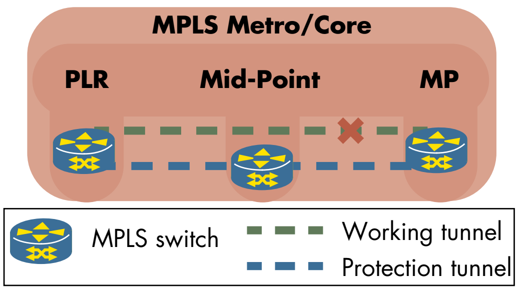

Figure 3: MPLS Fast Reroute Test Setup

detect failures of non-directly attached links and signal the failure to its neighbors quicker than the IGP used in the network. In both tests the BFD transmission interval was configured to 50 ms. The following devices were involved in BFD testing: ECI SR9705, Juniper MX480, and Tellabs 8830.

PBB-TE Protection

Resiliency in the PBB-TE domain was accomplished by defining a protection PBB-TE trunk for a working trunk. In case of the single homed access devices the working and protection trunks had the same endpoints, but different mid-point PBB-TE switches. In case of the dual homing, a single access device was connected to two different PBB-TE switches, so that one endpoint of the working and protection trunks was different.

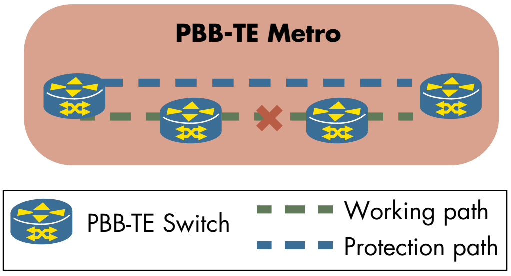

Figure 4: PBB-TE Protection Setup

In both cases, the detection was done by trunk endpoint switches using CFM (CFM is described in the Service OAM section). In almost all cases the Continuity Check Messages (CCM) transmission interval was configured to 10 milliseconds (ms), and in one case to 3.3 ms.

All vendors participating in the PBB-TE cloud were able to show PBB-TE protection interoperability. In most cases the failover time was below 51 ms, and restoration time below 5 ms. In one case the restoration time was 718 ms. The high restoration time was explained by differences between implementations. Some vendors shutdown traffic on a trunk when they revert over to the restored primary trunk and drop any traffic still running on the backup trunk. Other vendors continue to receive traffic on both trunks, but only transmit on one trunk. This ensures that the traffic is still received and reaches the destination.

T-MPLS Protection

Multiple protection mechanisms exist in T-MPLS. Both T-MPLS path (T-MPLS path is equivalent to an MPLS tunnel) and T-MPLS channel (a T-MPLS channel is equivalent to an MPLS pseudowire) can be protected. During this event we demonstrated only the per path protection mechanisms.

The T-MPLS linear 1:1 and 1+1 path protection schemes were demonstrated by three Ericsson Marconi OMS 2400 devices for 10 Gigabit Ethernet, STM-16 and 10 Gigabit G.709 interfaces. The protection schemes required two paths to be build to the same destination. In the 1:1 protection one path is declared as active and the other is set in backup mode. When the active path fails traffic is switched to the backup path. The second scheme, 1+1 protection,

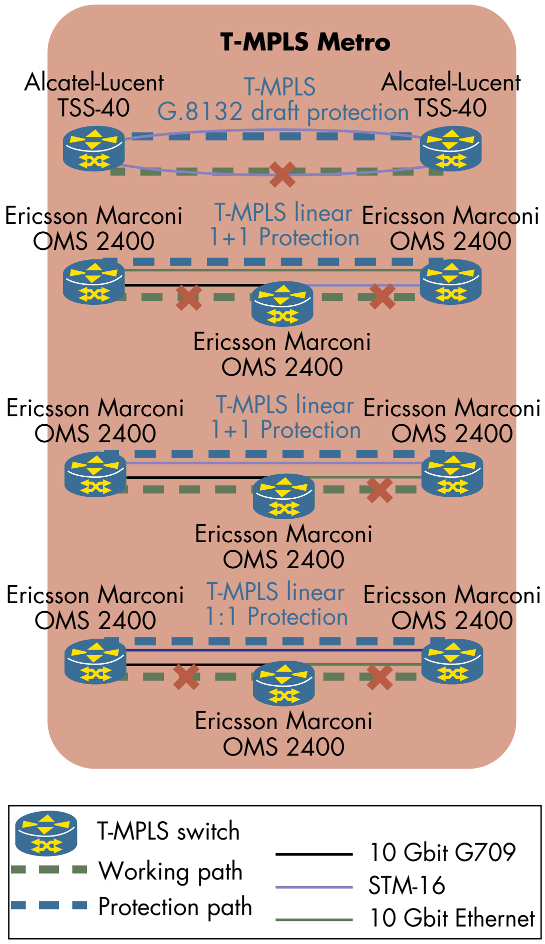

Figure 5: T-MPLS Protection Results

replicates all frames across both active and backup paths such that when the primary path fails, the backup path is already carrying the lost frames. In almost all cases Ericsson demonstrated the failover and restoration times below 35 ms for 1+1 protection and failover times below 30 ms for 1:1 protection. The restoration time for 1:1 protection was under 0.5 ms. In one test run we measured 76 ms failover time over 10 Gigabit Ethernet link for 1+1 protection.

Alcatel-Lucent demonstrated with the TSS-40 devices a pre-standard G.8132 T-MPLS ring protection implementation over STM-16 interfaces. In its demonstration Alcatel-Lucent showed failover times below 30 ms and restoration time below 16 ms. The two vendors already showed T-MPLS protection interoperability during EANTC's Mobile Backhaul event in January 2008 (report available on EANTC's web site).

The T-MPLS demonstration results are summarized in the figure above.

Management

Alcatel-Lucent demonstrated the 5650 Control Plane Assurance Manger (CPAM), a vendor agnostic IP/MPLS control plane management solution, which provided real-time IGP topology maps, control plane configuration, and a look into route advertisements within the layer 3 services. The tool was helpful since many configurations were being made by many different operators, helping to quickly identify configuration issues.

Editors

This document has been edited by Jambi Ganbar, Sergej Kaelberer, Jonathan Morin, Bionda Mueffke, Carsten Rossenhoevel and Gabriele Schrenk.

Acronyms

Term	Definition
AIS	Alarm Indication Signal
B-MAC	Backbone MAC
BFD	Bidirectional Fault Detection
BGP	Border Gateway Protocol
BRAS	Broadband Remote Access Server
CCM	Continuity Check Message
C-VLAN	Customer Edge Virtual LAN
CFM	Connectivity Fault Management
CPE	Customer Premise Equipment
DMM	Delay Measurement Message
DMR	Delay Measurement Reply
DSL	Digital Subscriber Line
E-NNI	External Network-to-Network Interface
EFM	Ethernet in the First Mile
ERPS	Ethernet Ring Protection Switching
EVC	Ethernet Virtual Connection
EVPL	Ethernet Virtual Private Line
FRR	Fast ReRoute
IGP	Internal Gateway Protocol
IS-IS	Intermediate System to Intermediate System
LACP	Link Aggregation Control Protocol
LAG	Link Aggregation Group
LDP	Label Distribution Protocol
LER	Label Edge Router
LMM	Loss Measurement Message
LOS	Loss Of Signal
LSP	Label Switched Path
MAC	Media Access Control
MEG	Maintenance Entity Group
MEP	Maintenance End Point
MPLS	Multi-Protocol Label Switching
MPLS-TP	MPLS Transport Profile
MTU-s	Multi Tenant Unit Switch
OAM	Operations, Administration and Maintenance
OSPF	Open Shortest Path First
PBB	Provider Backbone Bridge
PBB-TE	Provider Backbone Bridge Traffic Engineering
PE	Provider Edge
PLR	Point of Local Repair
PLSB	Provider Link State Bridging
PSN	Packet Switched Network
PW	PseudoWire
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
RDI	Remote Defect Indication
RFC	Request For Comments
RSTP	Rapid Spanning Tree Protocol
RSVP-TE	Resource reSerVation Protocol Traffic Engineering
S-Tag	Service Tag
SLA	Service Level Agreement
T-MPLS	Transport MPLS
TMC	T-MPLS Channel
TMP	T-MPLS Path
UNI	User-Network Interface
VLAN	Virtual Local Area Network
VoIP	Voice over IP
VPLS	Virtual Private LAN Service
VPN	Virtual Private Network
VPWS	Virtual Private Wire Service

References

Ethernet Services Attributes Phase 2 (MEF 10.1)
Requirements and Framework for Ethernet Service Protection in Metro Ethernet Networks (MEF 2)
Metro Ethernet Network Architecture Framework - Part 1: Generic Frame-work (MEF 4)
Ethernet Service Definitions (MEF 6)
User Network Interface (UNI) Requirements and Framework (MEF 11)
User Network Interface (UNI) Type 1 Implementation Agreement (MEF 13)
Ethernet Local Management Interface (E-LMI) (MEF 16)
User Network Interface (UNI) Type 2 Implementation Agreement (MEF Ballot)
Provider Bridges (IEEE 802.1ad-2005)
Provider Backbone Bridges (IEEE 802.1ah, draft 4.2)
Provider Backbone Bridges Traffic Engineering (IEEE 802.1qay, draft 3)
Shortest Path Bridging (IEEE 802.1aq)
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-2004)
Virtual Bridged Local Area Networks - Amendment 5: Connectivity Fault Management (IEEE 802.1ag-2007)
Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) (RFC 4447)
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)
Fast Reroute Extensions to RSVP-TE for LSP Tunnels (RFC 4090)
Architecture of Transport MPLS (T-MPLS) Layer Network (ITU-T G.8110.1)
Interfaces for the Transport MPLS (T-MPLS) Hierarchy (ITU-T G.8112)
Characteristics of Multi Protocol Label Switched (MPLS) Equipment Func-tional Blocks (ITU-T G.8121)
BFD For MPLS LSPs (draft-ietf-bfd-mpls-06.txt)
OAM Functions and Mechanisms for Ethernet Based Networks (ITU-T Y.1731)

EANTC AG
European Advanced Networking Test Center

Einsteinufer 17
10587 Berlin, Germany
Tel: +49 30 3180595-0
Fax: +49 30 3180595-10
info@eantc.de
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

v1.0