


Carrier Ethernet

WORLD CONGRESS

Public Multi-Vendor Carrier Ethernet Interoperability Test

Geneva, September 2007

EANTC

IIR

Editor's Note

Carsten Rossenhoevel
Managing Director

For the third time since 2005, EANTC invited network equipment manufacturers to a comprehensive interoperability test of Carrier Ethernet solutions. This became our largest Carrier Ethernet event so far — a total of 24 vendors participated, representing a large majority of the Carrier Ethernet market. The overwhelming vendor interest clearly shows that Carrier Ethernet services are a hot topic, and that more products are becoming available.

This time, our event focused on new infrastructure solutions in the core, metro and access areas. We are proud to report ground breaking progress at the first public multi-vendor interoperability verification of new metro transport technologies: Carrier Ethernet services were tested successfully over a PBT network (with nine vendors) and a T-MPLS cloud (with four vendors).

Our tests of MPLS (with nine vendors) confirmed its status as a mature and reliable solution for large Carrier Ethernet networks. Multi-vendor interoperability of established access solutions (Ethernet over copper, fiber, radio links) was confirmed at an advanced level. Furthermore, we duly noted that the number of implementations supporting Ethernet OAM has grown substantially since our last event. Last but very certainly not least we are pleased to report that multi-vendor provisioning solutions were tested for the first time at a public EANTC event.

The number and level of achievements has certainly been outstanding given an extremely heterogeneous test environment of 65 different systems. Preparations began in April 2007 and culminated in two weeks of intense hot-stage testing with more than 70 engineers from participating vendors at EANTC in August 2007. We hope that the test results as documented in this white paper can provide guidance to service providers and enterprises in a very progressive and rapidly changing technology.

Introduction

The third EANTC multi-vendor Carrier Ethernet interoperability event addresses four groups of industry stakeholders:

- Service providers: To understand the status of new Carrier Ethernet technologies, the choices and implications for every-day operations.
- Enterprises: To gain a better understanding of the types and service levels of Carrier Ethernet services they can expect from service providers today.
- Participating vendors: To greatly improve the interoperability of their implementations, understand the challenges and opportunities of heterogeneous environments.
- Industry forums and standards bodies (MEF, IEEE): To provide feedback regarding the readiness of standards for real world deployment.

Document Structure

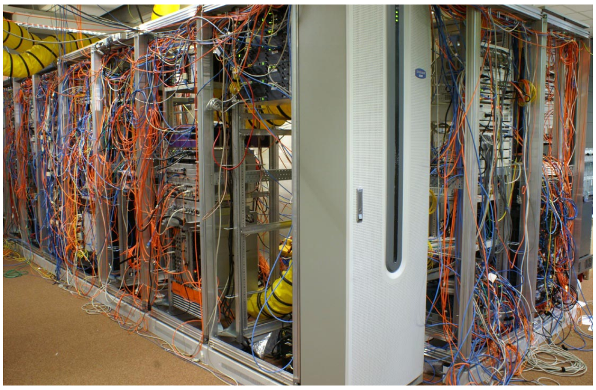

Image 1: Hot-staging at EANTC, Berlin

Based on EANTC's experience in organizing and executing multi-vendor interoperability events a 10 working days, closed doors, hot-staging event was conducted in EANTC's lab in Berlin, Germany. The results are summarized in this white paper.

Participants and Devices

Alcatel-Lucent	1850TSS/320 7450 ESS-1 7450 ESS-7
ANDA Networks	ER1006 ER1002 EE4000
Ceragon Networks	FibeAir IP-MAX^2
Ciena	CN 3102 CN 3106 CN 5060
Cisco Systems	ME 3400G-12CS ME 3400G-2CS ME 3400-24TS ME 3400-24FS ME 6524 Catalyst 3750 Metro Catalyst 6504 (Supervisor 720) 7604
Extreme Networks	Black Diamond 12802
Gridpoint Systems	VMS
Hammerhead Systems	HSX 6000
Harris Stratex Networks	Eclipse (Gigabit) Radio
Huawei Technologies	Quidway CX600 Quidway CX300A Quidway CX300B Quidway CX200A Quidway CX200B
Ixia Communications	XM12/IxNetwork
Juniper Networks	MX960 M7i
MRV Communications	OS304 OS910 OS9024-M

Nokia Siemens Networks	hiD 6670 hiD 6650
Nortel	Metro Ethernet Routing Switch 8600 Metro Ethernet Services Unit 1850
RAD Data Communications	ACE-3200 ACE-3400 ETX-102 ETX-202 ETX-202A RICi-16 LA-210 Egate-100 OP-1551 ASMi-54
Shenick Network Systems	diversifEye 8400
Soapstone Networks	PNC
Spirent Communications	Spirent TestCenter
Telco Systems, a BATM Company	T-Metro-200 T5C-24F T5C-24G T-Marc-250 T-Marc-254 T-Marc-340 T-Marc-380 Edge Gate483
Tellabs	6305 Ethernet Media Converter 6315 Metro Ethernet Node 6325 Edge Node 6345 Switch Node
TPACK	Longmorn Seagram
World Wide Packets	LE-311v
ZTE Corporation	ZXMP S385

Carrier Involvement

T-Systems engineers participated for the duration of the whole EANTC hot-staging event. T-Systems engineers were responsible for performing the Ethernet OAM tests.

Physical Network Topology

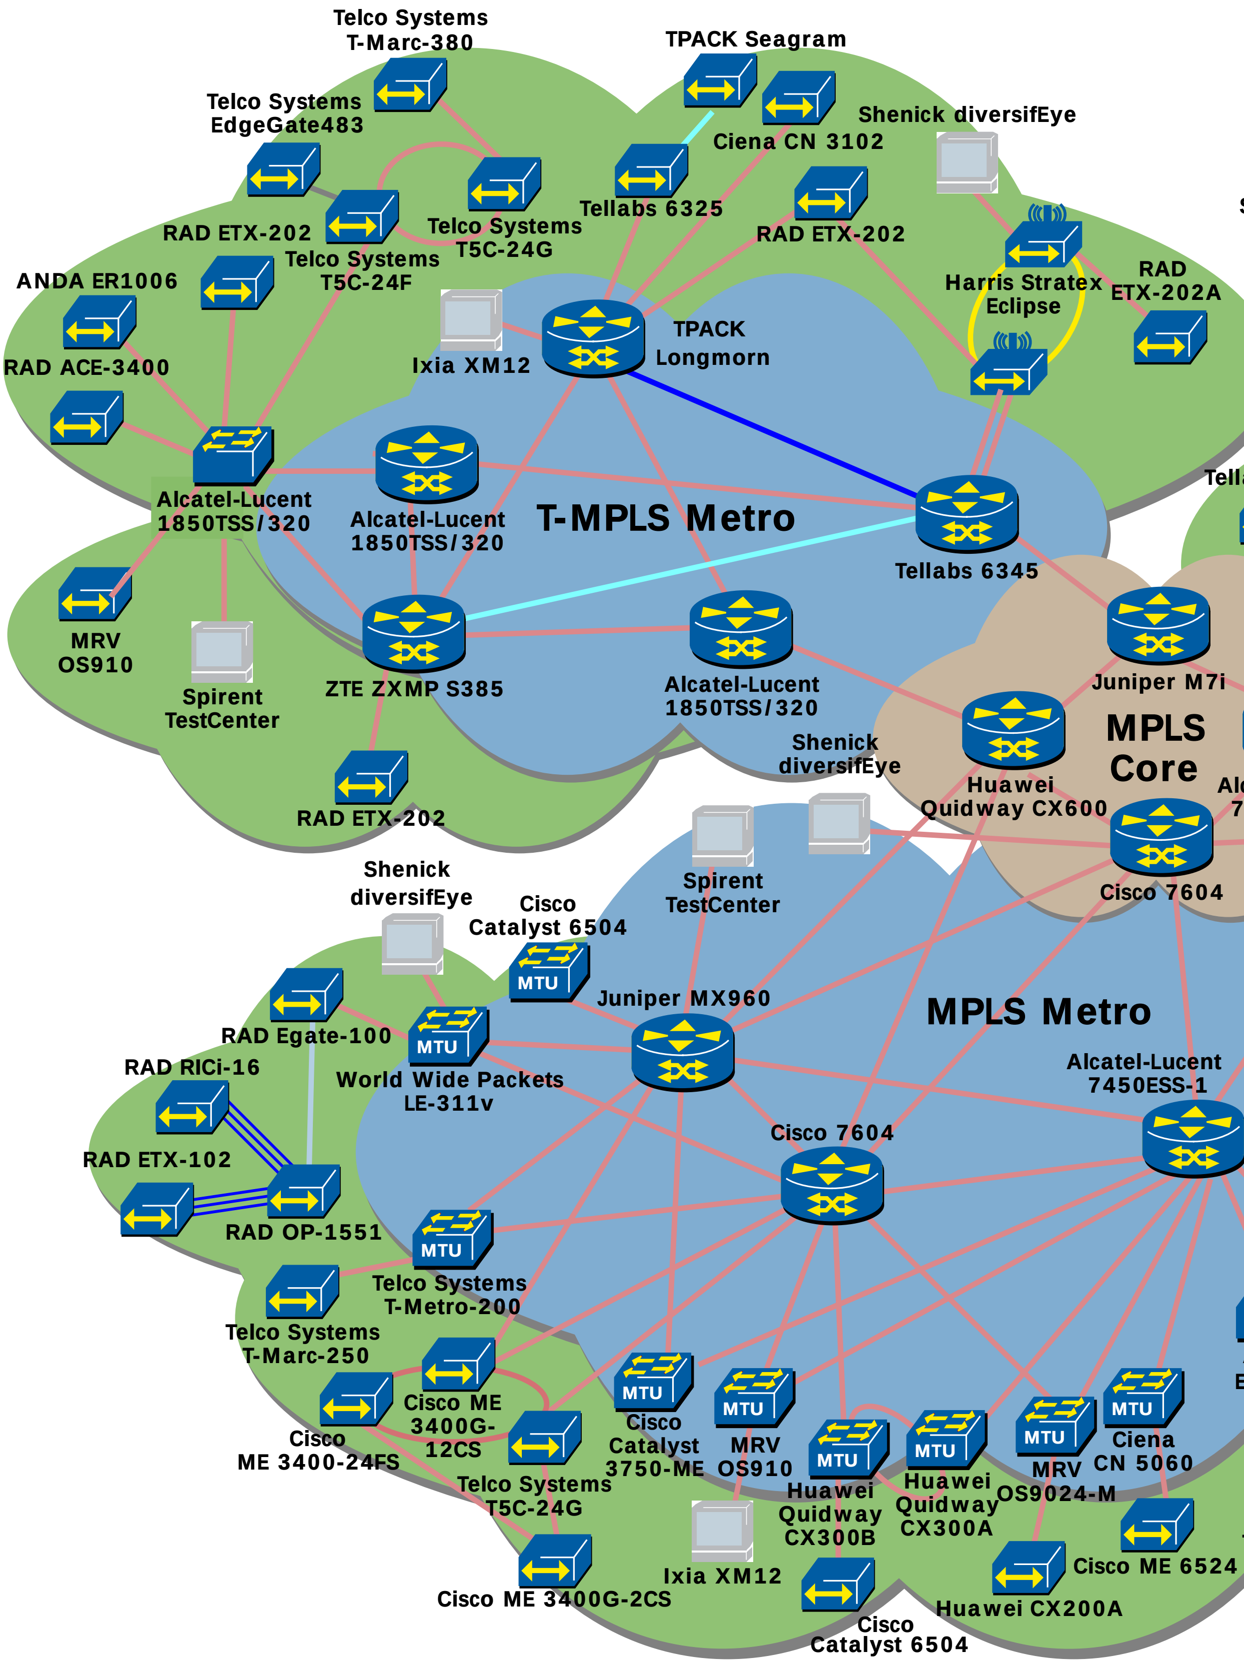

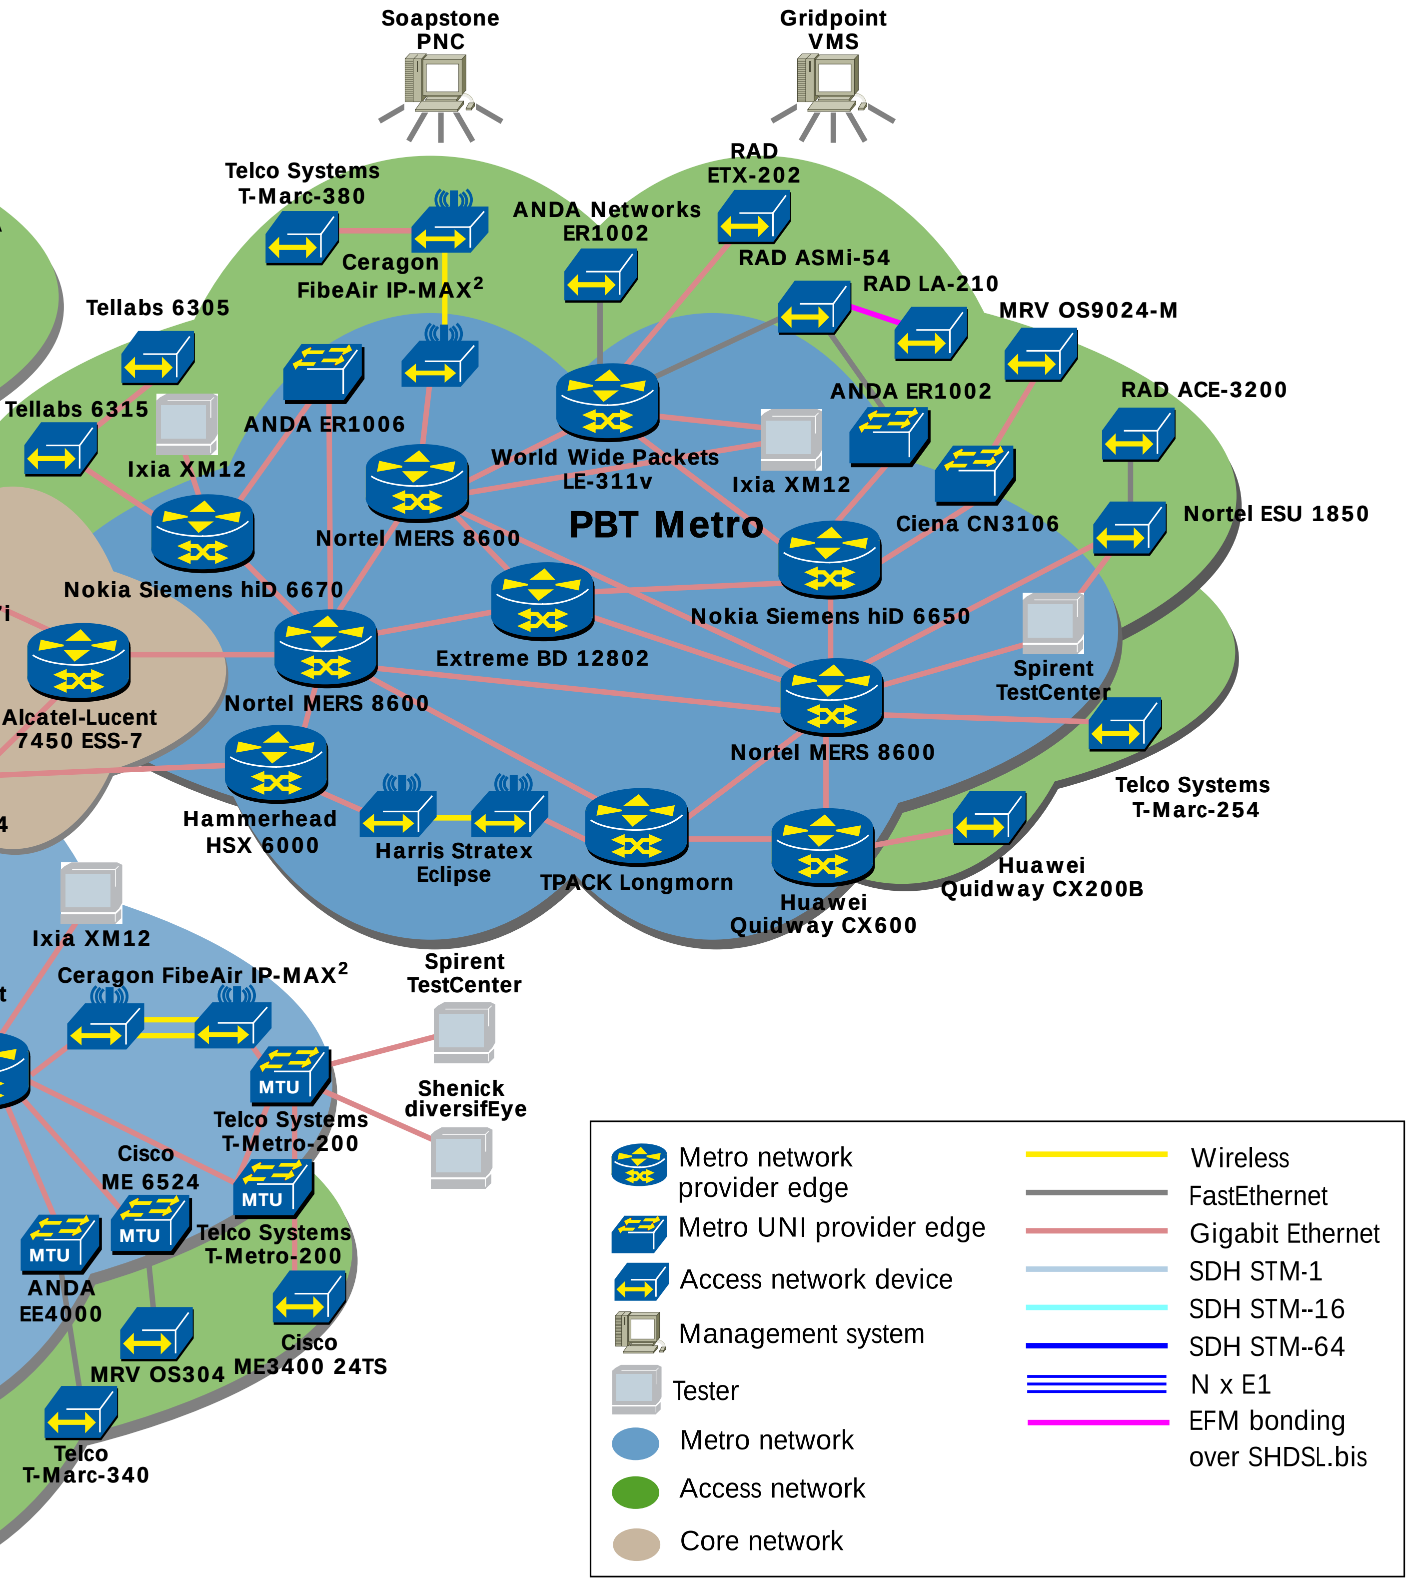
Soapstone PNC
Gridpoint VMS
RAD ETX-202
Telco Systems T-Marc-380
ANDA Networks ER1002
Ceragon FibeAir IP-MAX2
RAD ASMi-54
RAD LA-210
MRV OS9024-M
Tellabs 6305
ANDA ER1002
RAD ACE-3200
Tellabs 6315
ANDA ER1006
World Wide Packets LE-311v
Ixia XM12
Ixia XM12
Ciena CN3106
Nortel ESU 1850
PBT Metro
Nortel MERS 8600
Nokia Siemens hiD 6670
Nokia Siemens hiD 6650
Alcatel-Lucent 7450 ESS-7
Extreme BD 12802
Spirent TestCenter
Nortel MERS 8600
Nortel MERS 8600
Telco Systems T-Marc-254
Hammerhead HSX 6000
Harris Stratex Eclipse
TPACK Longmorn
Huawei Quidway CX600
Huawei Quidway CX200B
Ixia XM12
Ceragon FibeAir IP-MAX2
Spirent TestCenter
Shenick diversifEye
Telco Systems T-Metro-200
Cisco ME 6524
Telco Systems T-Metro-200
ANDA EE4000
MRV OS304
Cisco ME3400 24TS
Telco T-Marc-340
MTU
Metro network provider edge
Metro UNI provider edge
Access network device
Management system
Tester
Metro network
Access network
Core network
Wireless
FastEthernet
Gigabit Ethernet
SDH STM-1
SDH STM--16
SDH STM--64
N x E1
EFM bonding over SHDSL.bis

Network Design

There were two design guidelines for the Carrier Ethernet World Congress Interoperability Network:

- Resemble as much as possible next generation service provider networks for Carrier Ethernet services
- Allow for as many vendors as possible to test interoperability with each other

The participating vendors intended to test three metro/aggregation transport technologies: MPLS (Multi Protocol Label Switching), PBT (Provider Backbone Transport, which is being standardized by the IEEE as PBB-TE: Provider Backbone Bridge – Traffic Engineering) and T-MPLS (Transport-MPLS). We therefore planned for three metro area clouds, one for each transport technology.

In addition a broad range of access solutions and network termination devices were to be tested. Technologies such as Ethernet over PDH and SDH as well as Ethernet over microwave and Ethernet over DSL were verified. In order to realize residential and business services we attached access areas to each metro cloud.

Today, Metro Ethernet Forum (MEF) standards focus on UNIs (User-Network Interface) – the logical interface between the user and the service provider. MEF services are, therefore, defined over that interface. In the test network, the UNI was located between the access and the metro areas. Since the network was to resemble a carrier's network, we facilitated the creation of user services regardless of the type of metro cloud through which they connected, in much the same way that carriers strive to offer all network services independent of the underlying transport solution. Our intent was to verify if participating vendors implemented Carrier Ethernet services with comparable functionality and resilience across all technologies, so that customers would not feel any difference in the service regardless of their point of attachment to the network.

Last, but most certainly not least, was the interconnection of the different metro clouds. We decided, together with participants, to use MPLS in the core of the network as service providers commonly use MPLS in their core networks.

Interoperability Test Results

This time we extended the hot-staging test to full two working weeks in order to accommodate the large number of participants and the growing number of test areas. The substantial amount of test results collected during the hot-staging test are summarized in the following sections.

Access Technology Tests

A number of technologies are available today to carry Ethernet frames over xDSL lines, point-to-point radio links, PDH, SDH/Sonet and obviously fiber. Ethernet's ubiquity is one of the highlighted reasons for its success.

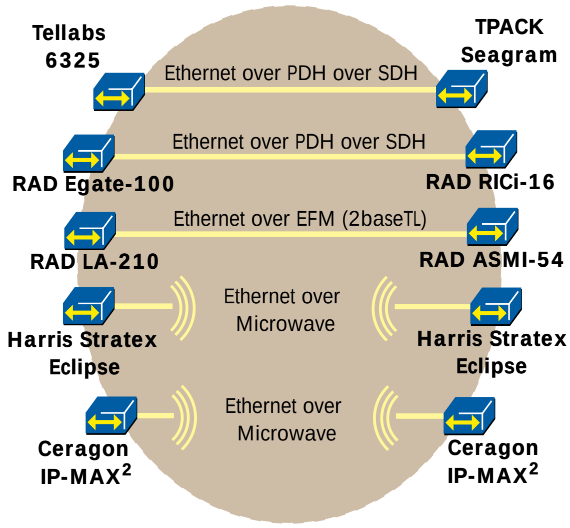

Figure 1: Diverse Access Solutions

The vendors participating in this test presented the following solutions[1]:

- Tellabs 6325 and TPACK Seagram tested the interoperability of Ethernet over PDH over SDH. This solution is typically applied where Ethernet services are provided over bonded PDH access lines, which in turn are backhauled over SDH.

  A total of 16 concatenated E1 containers were configured over an STM-16 interface to transport native Ethernet frames according to G.8040 (GFP Frame mapping into PDH), G.7043 (VCAT, Virtual Concatenation of PDH) and G.7042 (LCAS, Link Capacity Adjustment Scheme). LCAS can, when combined with physical diverse routing of VCG (virtual concatenation group) members, substantially increase the available network capacity when compared to a traditional SDH full protection scheme, which reserves spare bandwidth that is not utilized under normal conditions. LCAS can also change bandwidth during operation, meaning a VCG can be up- or downgraded in a hitless manner.

1. 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.

- RAD demonstrated a similar Ethernet over PDH/SDH service based on GFP, VCAT and LCAS using the RAD Egate-100 aggregator and RICi-16 CPE where each bonded group was mapped to a VLAN into the Metro network.
- Ethernet frames were also carried over point-to-point microwave links. One such solution was presented by Ceragon using their FibeAir IP-MAX2 solution and the second by Harris Stratex using their Eclipse product line. The radios demonstrated their capabilities to drop low priority traffic in the event that the wireless access link suffers degradation (such as seen by bad weather), using Adaptive Coding Modulation based on QoS parameters. The tests showed that both vendors can transport high priority traffic and drop low priority traffic when radio link capacity is reduced. This ability is different from traditional TDM or SDH links which fail under degraded conditions.
- RAD demonstrated Ethernet over EFM based SHDSL (2baseTL) in LA-210 using eight bonded wires employing EFM bonding.

Metro Transport Tests

A large number of service providers worldwide plan to upgrade, extend or exchange their metro aggregation networks in the near future. Increased requirements for scalability and coverage, range of Ethernet services offered and competitive pricing force carriers to search for new solutions for their metro business services, residential triple play offerings and inter-carrier services.

In response to these requirements, standards bodies and industry forums are working to develop new technologies to carry Ethernet services.

All metro aggregation devices were assigned to the three groups based on the transport technology of the vendors' choice: MPLS, Provider Backbone Transport (PBT) and Transport MPLS (T-MPLS). Transport was realized mostly over Gigabit Ethernet links, however, some links used STM-16, STM-64 and point-to-point microwave links.

As an independent test lab we were aiming to test the interoperability of common baseline services for all transport technologies. Since residential and business services were going to be offered we asked the participants in each cloud to support point-to-point and multipoint-to-multipoint services. We required each of the three transport technologies to establish the logical connections to support the end-user services and allowed for resiliency demonstrations within each cloud.

Multipoint Service Support

Both PBT and T-MPLS are considered pure transport technologies over which Metro Ethernet Forum defined services can be supported. Multipoint-to-multipoint services require explicit configuration of point-to-point tunnels or paths between all the service end points, plus a bridging function on those points.

A protocol called VPLS exists for multipoint services over MPLS. There are two incompatible IETF specifications, one based on LDP (supported by most players in the industry) and one based on BGP (supported by Juniper and Huawei in addition).

Provider Backbone Transport

PBT, like many other networking protocols, has started its life as a proprietary vendor solution. We decided to test its interoperability for three reasons:

- There is a tremendous interest from service providers to learn more about PBT's latest status and interoperability.
- PBT is being morphed into a multi-vendor industry specification: It has just begun its way through the IEEE standardization process under the name Provider Backbone Bridge – Traffic Engineering (PBB-TE, 802.1Qay). The standard will be considered an extension to the already established Provider Backbone Bridges (IEEE 802.1ah). The current draft carries the version number »0.0«.

  The work on the PBB-TE standard was officially started in 2007 and as a result most of the vendors implementations are still based on PBT, so we decided to base all tests on the PBT specification. For upcoming test events, we will switch to the PBB-TE standard.
- We were surprised to see that, albeit the protocol is so young, nine vendors signed up their products for the PBT interoperability testing.

The following devices were part of the PBT metro cloud and have successfully tested PBT trunk establishment and data forwarding: ANDA ER 1002 and ER 1006, Ciena CN 3106, Extreme Black Diamond 12802, Hammerhead HSX 6000, Huawei CX600, Nokia Siemens Networks hiD 6650 and hiD 6670, Nortel Metro Ethernet Routing Switch (MERS) 8600, TPACK Longmorn and World Wide Packets LE-311v. In addition, Ixia XM12 and Spirent TestCenter emulated PBT trunk endpoints and generated Ethernet L2 traffic over these ports.

The majority of the PBT trunks were configured manually since the current PBT specifications aims to separate the data plane (implemented in the switch) and the control plane (implemented externally). Several trunks were provisioned by two multi-vendor management systems: Gridpoint VMS and Soapstone PNC (please see the provisioning section for a discussion on the topic).

The external control plane must support all devices in the target network in order to be able to set up end-to-end trunks. In a diverse network environment as was constructed, provisioning and trunk setups was time consuming and performed by hand. In a real world deployment carriers are likely to adopt their provisioning system to their choice of PBT devices to allow a faster configuration and provisioning of trunks.

PBT is based on existing Ethernet standards, and as expected data forwarding worked without any problems. Since some of the vendors required Connectivity Fault Management (CFM) in order to establish the PBT trunks, CFM was used. CFM, however, is not required by all vendors in order to establish the trunks. We also found that one vendor had a proprietary version of CFM which required assistance from Ixia in order to setup the PBT trunks in our multi-vendor environment.

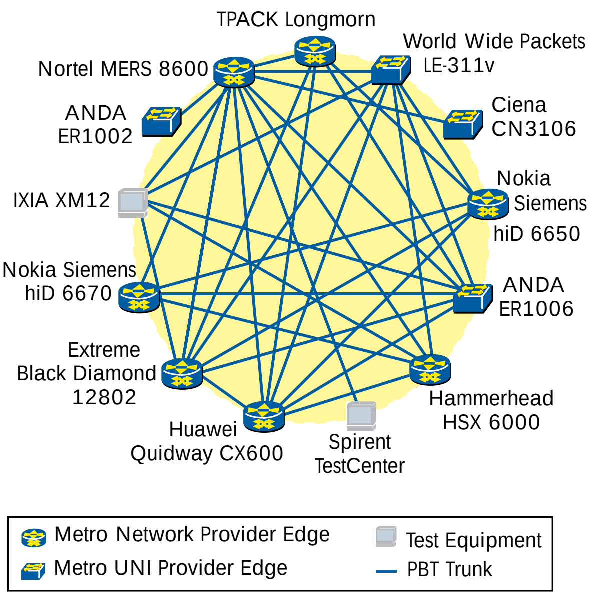

Figure 2: Provider Backbone Transport Trunks

Two further issues arouse before the PBT trunks could be established. In order to exchange CFM messages an Ethertype must be used. All vendors require a common Ethertype and since it is not yet defined in the standard, the best solution was to find an Ethertype that all participants can support. The Ethertype chosen was non-standard, however, it allowed all vendors to exchange CFM messages. The CFM message interval was also a source of discussion as some vendors could not change their CFM interval to the one agreed upon by the group. As a result, some trunks could not initially be established. Vendors were able to receive new code that allowed the CFM intervals to be changed and matched the group's configuration and were able to bring up the trunks.

Since there is no standard implementation of PBT, it comes as no surprise that there were inconsistencies in the implementations of different vendors with regard to EtherType values, protection and OAM. However, once agreement was reached, full interoperability was demonstrated by all the vendors. The conclusion is, therefore, that while there are currently differences in the implementation of PBT, interoperability is achievable.

Transport MPLS (T-MPLS)

T-MPLS is another new transport technology aiming to provide a way for carriers to virtualize a wire and offer the same predictive paths as carriers are used to in the SDH world. T-MPLS is defined in five ITU-T draft documents either approved or under approval (see reference section). T-MPLS and MPLS have a common forwarding plane: from this perspective, T-MPLS paths (TMP) correspond to MPLS Tunnels, while T-MPLS channels (TMC) correspond to MPLS Pseudowires.

The T-MPLS reference standard does not define a specific control plane, but does not preclude the future use of the GMPLS control plane (e.g. RSVP-TE for signaling; OSPF-TE and ISIS-TE for routing). At the time of the hot-staging, only manual configuration was available. Four vendors participated in the T-MPLS interoperability. Given the small size of the network, manual configuration was not a problem and T-MPLS Paths (TMPs) were configured quickly.

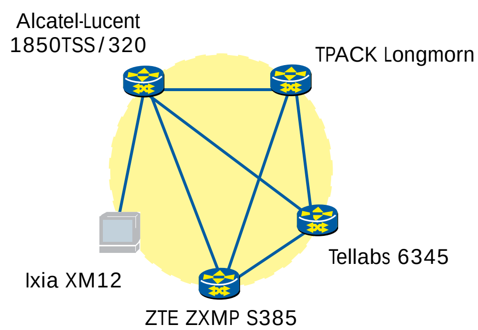

Figure 3: T-MPLS Paths (TMPs)

The following devices were part of the T-MPLS metro cloud and have successfully tested and demonstrated the T-MPLS Path (TMP) and T-MPLS Channel (TMC) establishment, data forwarding over T-MPLS and T-MPLS Multiplexing in which a number of TMCs were transported within a TMP:
Alcatel-Lucent 1850TSS/320, Tellabs 6345, TPACK Longmorn and ZTE ZXMP S385. Between Tellabs 6345

and TPACK Longmorn, T-MPLS over SDH was successfully tested over an STM-64 link. Another SDH link (STM-16) was successfully tested between ZTE ZXMP S385 and Tellabs 6345. The Ixia XM12 emulated T-MPLS endpoints and together and generated Ethernet traffic over the same ports. Spirent TestCenter, also, generated Ethernet traffic to verify data forwarding.

Since T-MPLS borrows the concept of label switching from MPLS some of the participants required IP addresses for their interfaces and used the ARP protocol to resolve Ethernet addresses for next-hop forwarding. Other vendors did not support ARP — hence the Ethernet addresses of some of the participating interfaces had to be entered manually into the partnering T-MPLS switches.

EANTC defined a list of end-to-end Ethernet Virtual Circuits (EVCs). They needed to be established from the access network across the T-MPLS cloud and then handed over to the core network on a single joint physical interface. This required that T-MPLS equipment at the edge of the T-MPLS cloud would be able to map from VLAN tags to TMCs. Some implementations had not reached that level of maturity yet — they supported only port-based mapping where all VLANs at an access port are mapped into a single T-MPLS channel.

Another issue discovered during the event was the range of values for channels (TMCs) and paths (TMPs). By definition, the T-MPLS label range is the same as MPLS. Due to some vendor software issues a subset of the values was initially supported. This was fixed with a software upgrade during the test event.

MPLS for Aggregation

The tests in this area were based on previous experience largely gained in previous Carrier Ethernet World Congress and MPLS World Congress interoperability test events. 16 MPLS vendors took part in the most recent effort in February this year. The technology is considered mature and the standards defined by the IETF are enjoying wide deployment. In addition to using MPLS as a core technology, the participating vendors also tested MPLS interoperability as metro / aggregation technology.

Participants in the MPLS cloud proposed to configure point-to-point services over Ethernet Pseudowires (PWs) and multipoint-to-multipoint services over Virtual Private LAN Services (VPLS).

The MPLS metro / aggregation cloud included 9 vendors and operated independently from the core cloud (where MPLS was used as well). MPLS-enabled Carrier Ethernet platforms participating in this area included Alcatel-Lucent 7450 ESS-1, Cisco 7604 and Juniper MX960. All three routers constructed an LDP-based VPLS domain — this was the first public interoperability test of LDP-VPLS including Juniper.

The above devices served as Provider Edge routers (PEs) and were surrounded by hierarchical VPLS (H-VPLS) multi-tenant units (MTUs). ANDA EE4000, Ciena CN 5060, Cisco Catalyst 6504, Cisco ME 6524, Cisco Catalyst 3750-ME, Huawei Quidway CX300B, Huawei Quidway CX300A, MRV OS910, MRV OS9024-M, Telco Systems T-Metro-200 and World Wide Packets LE-311v all participated as MTUs. The connection between Telco Systems T-Metro-200 and Alcatel-Lucent 7450 ESS-1 was realized over a point-to-point microwave link using Ceragon's FibeAir IP-MAX2.

A few problems were encountered while constructing the H-VPLS domain. In one case, a signaling session could not be established due to a mismatch in the maximum PDU size for the signaling protocol LDP. In another case, a system stripped VLAN headers before forwarding the Ethernet frames, a major cause of confusion on the receiving end of the MPLS tunnel which were expecting a VLAN header.

In one case, a Label Switched Path (LSP) could not be established because an implementation did not accept the implicit null or explicit null label values it received.

These issues were found in rather new implementations participating at an EANTC interoperability event for the first time — an indication that new MPLS implementations still need to be verified for compliance and interoperability carefully. We are certain that participating vendors will fix the issues identified in Berlin so that the interoperability of all participating devices will improve based on the event.

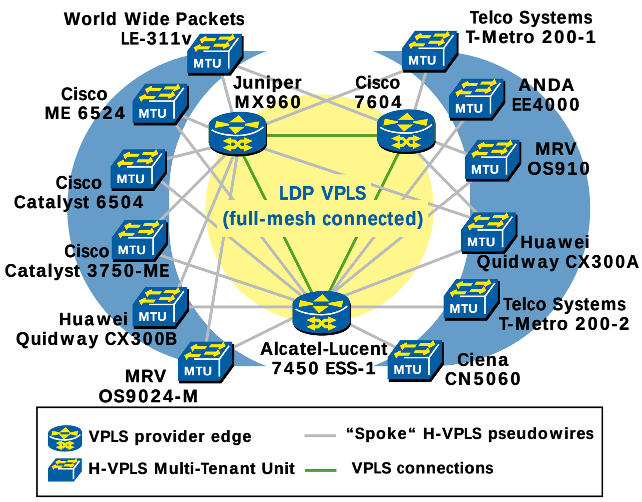

Figure 4: H-VPLS Metro / Aggregation

In conclusion, MPLS was the only metro transport technology at the event that demonstrated native support for both point-to-point and multipoint-to-multipoint services.

As we will describe later in the "MPLS Backbone Tests" section, MPLS also played a key role in enabling multi-point-to-multipoint connectivity for other metro transport technologies.

OSS and Provisioning Tests

To solve the provisioning challenges in building services primarily (but not limited to) across PBT networks, Gridpoint VMS (Value Management Suite) and Soapstone PNC (Provider Network Controller) implemented service activation tools for the PBT network under test. Since the PBT standard does not define a control plane, an external software-based control plane was integrated with the network to verify PBT equipment deployment readiness. Furthermore, management vendors strive to decouple services from the underlying network technology, although the two system types present at CEWC 2007 focused just the PBT control plane for the moment.

The primary tasks were to test the ability:

- To engineer and provision PBT tunnels across multiple vendors' equipment automatically
- To identify the optimum path for PBT tunnels between any two endpoints while meeting its QoS requirements and optimizing business rules
- To create suitable backup PBT trunks across (if possible) disjunct parts of the network infrastructure

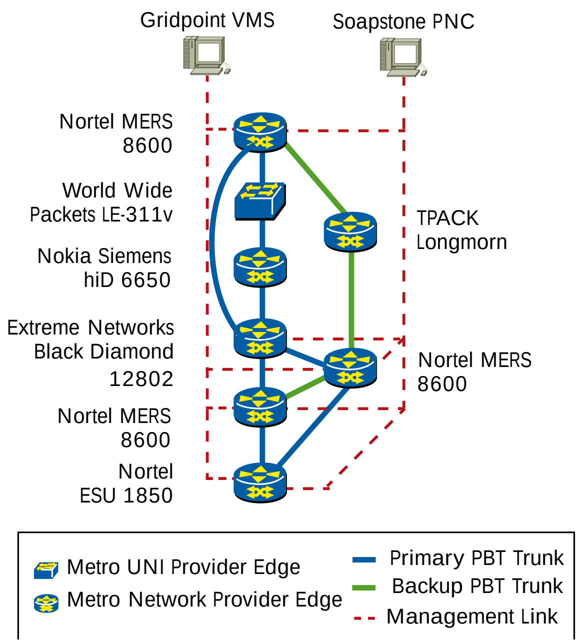

Figure 5: Multi-Vendor PBT Provisioning

Gridpoint VMS and Soapstone PNC configured primary and backup trunks across the PBT cloud.

Gridpoint VMS interfaced with Nortel MERS 8600 and Nortel ESU 1850 equipment. With the PBT trunk group in place, Gridpoint VMS created a service, delivering a rate-limited 5 Mbps service.

Soapstone's PNC participated in the PBT demo where it dynamically configured five services across a multi-node PBT network comprised of Nortel Networks MERS8600 and Extreme Networks BD12802 switches.

Soapstone PNC mediated between multiple resource requests from the services/control layer and OSS. The PNC delivered network modeling and network-to-service fault as well as SLA correlation.

In a separate step, Gridpoint VMS modeled the complete test network with 80 network elements. The modeling was carried out manually, not by reading out switch configurations. The highly interconnected network presented more that 50,000 possible routes across the network to form a service. VMS was used to calculate the most suitable paths across the network, with regards to bandwidth and service SLAs. Backup paths were calculated along disjunct routes where possible. The resulting end-to-end paths could be read off the screen and configured manually by participating vendors' engineers.

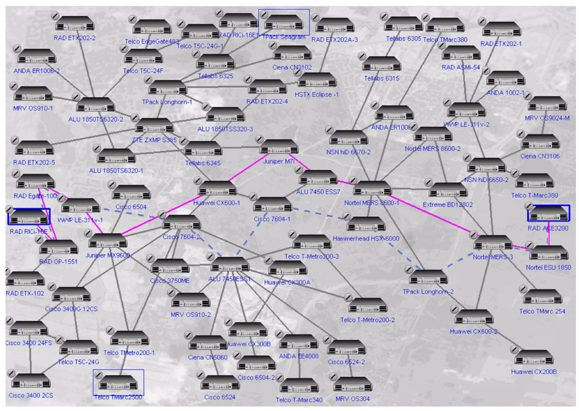

Figure 6: Network Model for Trunk Planning

MPLS Backbone Tests

One major design aspect for service providers will be the interface between the metro/aggregation and the core network. We built an MPLS backbone to interconnect the three aggregation networks using an Alcatel-Lucent 7450 ESS-7, a Cisco 7604, a Huawei Quidway CX600 and a Juniper M7i. The backbone provided transport for both point-to-point and multipoint-to-multipoint Ethernet services. Point-to-point services were implemented with Ethernet pseudowires while multipoint services were implemented using VPLS with LDP signaling between the four vendors.

The connections between the three metro/aggregation clouds are discussed below. We made no assumptions regarding the inter-provider or intra-provider nature of the network-to-network interface.

MPLS Metro to MPLS Core Connectivity. An MPLS interface with multi-segment pseudowires was used for the connection between the MPLS metro cloud and the MPLS core. This MPLS network-to-network interface (NNI) confined all external signaling between the two networks exclusively to the peering devices. This minimized the reachability information shared across the two networks and maintained common data, control and management planes across both networks and their interconnection interface. Participating vendors opted for this approach, instead of a native Ethernet interface, in order to allow testing of multi-segment pseudowire solution.

The participating vendors decided to test the following combinations:

- Cisco 7604, Huawei Quidway CX600 and Juniper MX960 interoperated using eBGP with label (RFC 3107). In this case, the peering devices functioned as Autonomous System Border Routers (ASBRs).
- Alcatel-Lucent 7450 ESS-1 and Cisco 7604 were tested using LDP and a separate OSPF process between the directly attached interfaces.

eBGP and LDP (according to the list above) were used to exchange tunnel label information while three pseudowire segments were stitched together for each end-to-end point-to-point service:

1. Pseudowire segment within the MPLS metro area
2. Pseudowire segment between the metro ASBR and the core ASBR
3. Pseudowire segment inside the core to the egress point form the core network

For the multipoint-to-multipoint service, an H-VPLS spoke (between Alcatel-Lucent 7450 ESS-1 and Cisco 7604) was created to join the VPLS instances in the MPLS metro and MPLS backbone.

PBT Metro to MPLS Core Connectivity. In the case of PBT, where the external control plane solutions provided by the participating vendors did not extend into the core backbone, the data forwarding between PBT and MPLS core clouds was accomplished in the following two ways:

- IEEE 802.1Q interface between Nortel MERS 8600 and Alcatel-Lucent 7450 ESS-7

  This connection was used for point-to-point services only. The Nortel MERS 8600 terminated the PBT trunks and forwarded traffic to Alcatel-Lucent 7450 ESS-7 via Ethernet VLANs. The Alcatel-Lucent router then mapped each of the VLANs into a separate MPLS pseudowire.
- MPLS pseudowire interface between Hammerhead HSX 6000 and Cisco 7604.

This link was used to establish multipoint-to-multipoint services based on the HSX 6000's ability to map PBT trunks directly to MPLS pseudowires. Hammerhead was then connected to Cisco 7604 in the core of the network which treated the Hammerhead as a spoke H-VPLS MTU. As discussed above, PBT supports only point-to-point tunnels so instead the Hammerhead HSX 6000 terminated multiple PBT trunks into multiple EoMPLS PWs (one to one mapping). These PWs were then associated as H-VPLS spokes to a common VPLS instance on the Cisco 7604 and thus bridging traffic among the spokes.

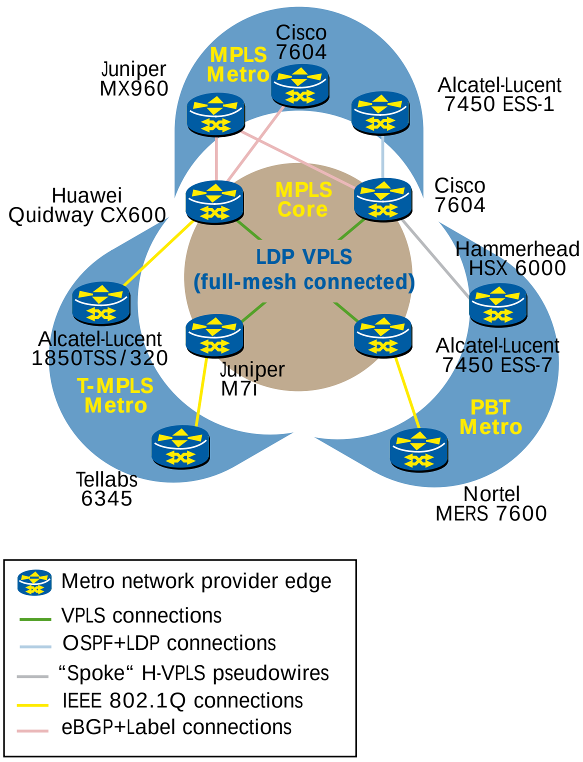

Figure 7: MPLS Core Logical Connections

T-MPLS Metro to MPLS Core Connectivity. Similar to the previous scenario, no control plane was provided on the T-MPLS side to interface with the MPLS core and therefore, a native Ethernet interface was used to connect both networks. The data connections between the MPLS core and the T-MPLS cloud were constructed using the following two solutions:

- VLAN-based (IEEE 802.1Q) interface between Tellabs 6345 and Juniper Networks M7i

  The TMCs terminating on the Tellabs 6345 were mapped to VLAN tagged Ethernet frames and were then forwarded to the Juniper Networks M7i located in the core. The frames coming from the core device were mapped to T-MPLS channels based on their VLAN tags by the receiving Tellabs 6345 device.

- T-MPLS TMCs originating from Alcatel-Lucent 1850 TSS320 and stitched together by Huawei Quidway CX600 to a PW through the core to the MPLS aggregation cloud.

  The multipoint-to-multipoint service was realized building a bridge instance on the Tellabs 6345 device for a set of T-MPLS channels and physical port. As a result, the external multipoint service implementation was designed similar to the PBT cloud.

Ethernet OAM Tests

We tested two Ethernet Operation, Administration and Maintenance (OAM) standards in previous Carrier Ethernet World Congress Interoperability events:

- IEEE 802.1ag – Connectivity Fault Management (CFM)
- IEEE 802.3ah – Ethernet in the First Mile (EFM), specifically section 57.

In 2006, only a few implementations existed so we were pleased to receive a great deal more interest in both test areas this time.

Native Ethernet fault management protocols are essential for the delivery of high quality end-to-end services to business customers and increasingly also consumers (for example for IPTV or VoIP). The service availability is key for many customer applications; it can be managed using Ethernet OAM protocols. The tests described below are a subset of the rich set of OAM tools available today for Ethernet.

Ethernet Link OAM

The Ethernet in the First Mile IEEE standard extends the Ethernet Media Access Control (MAC) layer to additional physical layers such as voice grade copper cable. As part of the standard, a set of OAM tools are defined aiming to help carriers monitor link operations. We tested the following three functions:

- Link OAM Discovery – Detecting the existence and configuration of the OAM sublayer in the neighboring Ethernet equipment
- Link OAM Loopback – Configuring a loopback mode for fault localization and link performance testing
- Link OAM Dying Gasp – Notifying the neighbor about an imminent system shutdown

The OAM discovery test involved 28 unique vendor-pairs for the discovery test. Only four implementations had problems detecting neighbor capabilities at this initial discovery phase of the testing.

A slightly smaller group of 17 devices supported the loopback testing. Seven of these implementations managed to enter loopback mode against their test partner. The problems plaguing this test were wide in range and included systems not able to reach “SEND ANY” state or being stuck in loopback state once the test was done, just to name a few.

We believe OAM interoperability testing is specifically important: It is very likely that different vendors will supply the customer premises equipment (CPE) and the switches terminating the metro/aggregation cloud. Carriers need to rely on remote fault localization and performance verification in order to keep SLAs and offer a

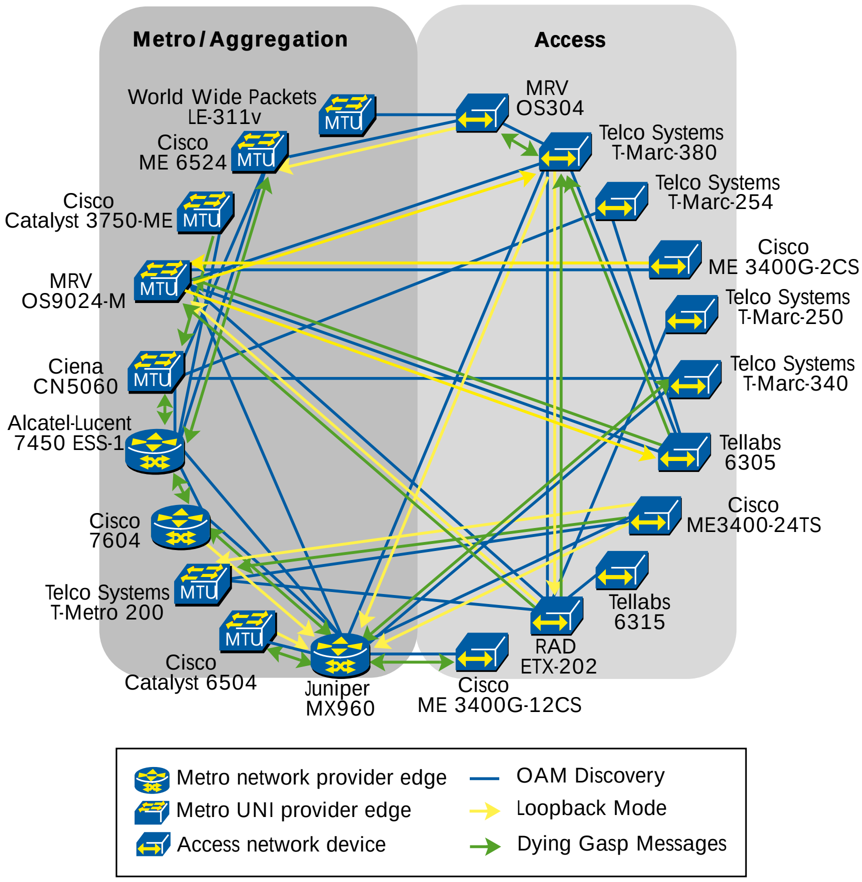

Figure 8: Ethernet Link OAM Tests

profitable service. While it is positive to see a growing number of Ethernet OAM implementations, a lot of work still needs to be done in testing compliance and interoperability.

The Dying Gasp test was performed between 16 pairs of vendors. Participating devices' roles were asymmetric: The customer-side device generates the dying gasp while the network-side device is responsible for receiving and processing the message.

This test was successful between almost all pairs. In one case we monitored an outgoing dying gasp message which was not registered by the receiver and one node failed to generate the dying gasp message as it was being powered off.

End-to-End Ethernet Service OAM

In contrast to link OAM, service OAM (as defined in IEEE 802.1ag) allows up to eight layers of end-to-end service monitoring through different management domains and can be used across logical connections. This standard is often called Connectivity Fault Management (CFM).

The major use of CFM for service providers and enterprises is to verify connectivity across different management domains. A carrier can define a management domain level to be used internally while allowing their customers to verify end-to-end connectivity over the network using a different CFM level.

At EANTC, we already tested service OAM – specifically the Continuity Check functionality – at Carrier Ethernet World Congress 2006. Last year, six devices from four vendors (early adopters of the draft standard) had participated in the test.

This time, 16 vendors with an impressive 34 CPE implementations expressed interest in joining the service OAM tests. We were able to add Link Trace and Loopback testing to the test plan based on the progress of implementations.

Because IEEE 802.1ag is still in draft status — and naturally, different draft versions are not backwards compatible — it was important to check the versions supported by participants. Implementations ranged from draft version 1.0 to the most current draft 8.1. This fact by itself is interesting as several vendors are waiting for the standard to be finalized before making any further CFM code updates.

At the beginning, we noticed a lot of configuration issues including Ethertype mismatches, Connectivity Check (CC) interval times incompatibility problems as well as a range of domain name format issues.

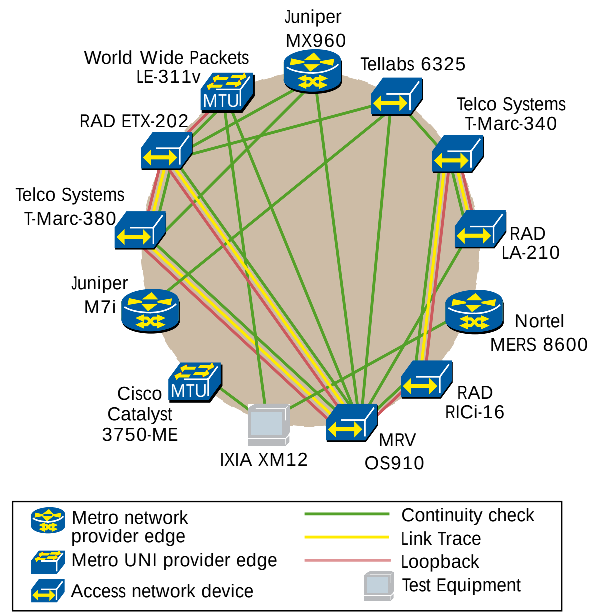

Figure 9: Back to Back Service OAM Tests

Once the configuration issues were resolved, we were able to test a large number of devices successfully. In order to use the testing time effectively, 19 CPE vendors started by testing CFM Continuity Check (CC) messages back-to-back, in directly connected configurations. Six of these implementations tested CFM loopback and two tested CFM link trace as well. Figure 9 shows an overview of the back-to-back tests.

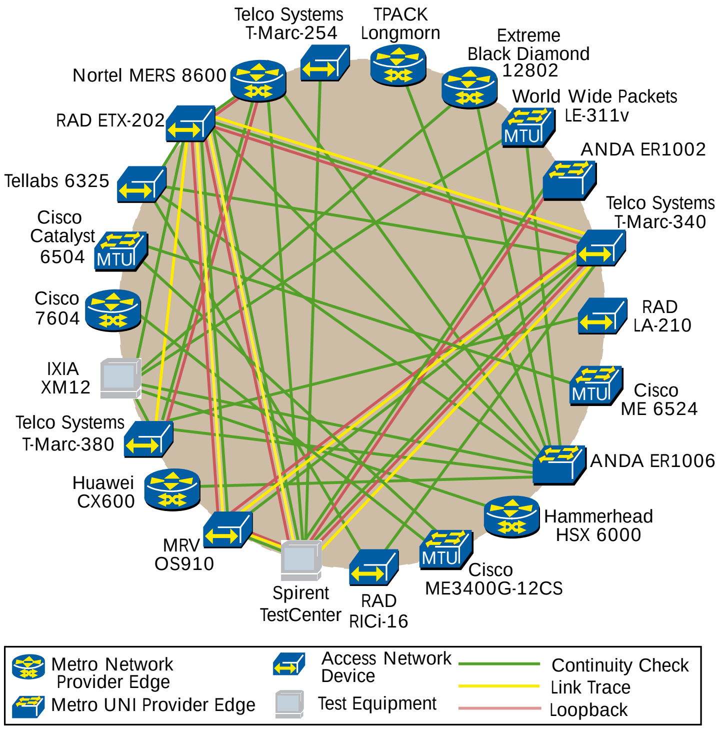

Figure 10: End-to-End Service OAM Tests (CFM)

Further tests were conducted across the complete network once it was set up. Connectivity Check tests were conducted with 36 device pairs; Loopback functionality was successfully evaluated in nine configurations and eight pairs successfully tested CFM's link trace. End-to-end test pairs had been carefully chosen based on information collected from participants in advance in order to facilitate testing between vendors who supported the same IEEE 802.1ag draft versions. We were pleased to see that careful planning helped achieve full interoperability, however, full mesh testing might have had different results.

Ethernet Local Management Interface (E-LMI). The E-LMI protocol is specified by the Metro Ethernet Forum in MEF 16. It enables the customer premises equipment to request and receive status and service attributes information from the Carrier Ethernet Network, including Ethernet Virtual Channel (EVC) and remote UNI status information as well as EVC information (including C-VLAN to EVC map). The information propagated by E-LMI is provided by the other OAM protocols as mentioned above.

Cisco demonstrated E-LMI capabilities using Cisco ME3400G-2CS, ME3400-24FS, ME3400-24T and Catalyst 3750-ME, showing end-to-end propagation of EVC status to the customer equipment. This information allowed the customer equipment to dynamically reroute traffic around failed services/VLANs.

Ethernet OAM Summary. Altogether, the Ethernet OAM test area proved to be extremely popular in comparison to last year's test event. Both native Ethernet transport and Carrier Ethernet services are gaining traction in the market so more and more carriers require Ethernet OAM tools, on par with the tools available in existing networks.

Resiliency Tests

A range of resiliency technologies can be employed to deliver Carrier Ethernet services with high availability. Resiliency is the ability of the network to react to node or link failures and to reconverge quickly. Users expect different levels of resiliency according to service types. TDM services are typically expected to restore service within 50 ms, Voice over IP services in a similar order of magnitude. Sub-second restoration is typically accepted for video and data services. In the test results presented here, Ixia XM12 was used to generate test traffic and calculate the out of service time during failover and restoration.

Resiliency using Link aggregation

Link aggregation is typically used for two reasons: To scale bandwidth and to protect links. This test case focused on link resiliency.

Link aggregation facilitates the grouping of several interfaces into a single logical unit. Once the unit is created the device can use any set of algorithms to divide traffic across the link aggregation group (LAG). Traffic across a LAG with n links must be kept below $(n – 1) / n$ of the total available bandwidth (i.e. 1/2 = 50% for two links). In this case the LAG can carry all traffic during a single link failure. LACP, the Link Aggregation Control Protocol, is responsible for maintaining links and synchronizing neighboring switches.

A single fiber of a duplex link was disconnected during each test run. We measured how fast the unidirectional link failure was relayed to the other side using LACP and how quickly both systems agreed not to use the link anymore and to rearrange for the data streams to take different links.

The following devices participated in the link aggregation tests with LACP: Ciena CN5060, MRV OS9024, Nokia Siemens Networks 6670, Nortel MERS 8600, Telco Systems T-Marc-380 and World Wide Packets LE-311v.

We recorded an out of service time of less than 100 ms in all failover situations – a very positive result for carriers planning on using LACP as a resiliency mechanism. On the other hand we recorded service outage of almost seven seconds during restoration in one case and in another test run the restored link was not used to forward traffic at all again, pointing to a few implementation issues that can certainly be resolved based on the experience the vendors gained in our event.

Static link aggregation groups (without LACP) were configured between Harris Stratex Eclipse and the Tellabs 6345 and a failover time of less than 150 ms was recorded between the two with only 1.4 ms out of service time when the link was restored.

Another demonstration in this area was Alcatel-Lucent's "Multi-chassis LAG" (MCLAG), a solution combining MPLS-based protocols to provide protection against link and node failure by terminating the LAG end points over different routers. The demonstration consisted of Alcatel-Lucent's 7450 ESS7 and 7450 ESS1 which were connected to a VPLS domain. The MCLAG control traffic utilized the VPLS service between the two nodes to control the active and standby LAG ports.

Several resilience test scenarios were explored within each individual metro/aggregation cloud, according to the abilities of each technology. All three metro areas provided native resilience mechanisms. In addition, several innovative alternatives were proposed by some of the participating vendors to further improve coverage and efficiency of resiliency solutions.

PBT Protection

Resiliency within the PBT domain was accomplished by defining two PBT trunks (»working path« and »protection path«) to the same destination. As the PBT ingress device detects that the working path has failed, it switches traffic over to the protected path. The failure detection is done by CFM (Connectivity Fault Management, see the explanation in the box on the right).

A large majority of the PBT vendors present at the event also participated in the PBT resiliency testing. We intentionally failed each primary path and measured frame loss as an indication of the time required to switch to the backup path. Since participants had agreed to a Continuity Check (CC) interval of 10 ms, we expected around 50 ms out of service time. This was achieved in almost all test configurations — a reassuring result given that these were multi-vendor configurations. In very few cases, out of service times up to 350 ms were measured.

The restoration phase of the test showed less consistent results. Two of the trunk pairs performed the restoration phase of the test without losing a single frame while an out of service time of up to 800 ms was recorded between two other pairs.

PBT allows for an optional external system to calculate and configure primary and backup paths. The Gridpoint VMS system provisioned both paths in the test and was able to trigger a switchover manually. No frames were lost during the scheduled rerouting. When switching back to the primary trunk, we measured an out-of-service time of 0.6 ms.

Soapstone PNC provided automated path selection for customizable services, thus instantiating, disabling, tearing down and failing over services across the PBT network.

T-MPLS Protection

T-MPLS path protection and PBT trunk protection are mostly similar in design. Again, we planned to define a primary and a backup path for each trunk and measure the out of service time when the primary path was failed.

An interoperability issue was discovered during preparation of the test. One of the vendors involved in the test supported T-MPLS's Connectivity Verification protocol version 1 (CV_v1), two other supported version 0 (CV_v0) and the fourth vendor supported both CV types.

PBT, T-MPLS and Connectivity Fault Management

Both PBT and T-MPLS use similar mechanisms for connectivity verification, as recommended in IEEE802.1ag and ITU-T G.8114 respectively. Continuity Check messages are exchanged between the end points in a connection (PBT trunk end points and TMP end points). Once a number of messages are no longer received, the devices declare the path dead and switch over to the alternative path.

A carrier using PBT and T-MPLS native failure detection mechanisms must keep in mind the results presented above in the CFM tests. Even within the two technology specific clouds we discovered interoperability issues in such areas as Ethertype, intervals and lack of support for CFM.

Consequently we were able to evaluate only two multi-vendor combinations in this area because each test scenario required three devices (ingress, egress, midpoint): Two devices using CV_v0 and two devices using CV_v1 with one of the other units as the bypass point.

Tellabs 6345 and TPACK Longmorn-1 used Y.1711 for connectivity verification (CV) and Y.1720 for resiliency across their TMPs and over an STM-64 link. The CV interval used was 50 ms which is defined in the standard as Fast Failure Detection (FFD).

Alcatel-Lucent 1850TSS/320 tested resiliency against ZTE ZXMP S385 using CV_v1. The two agreed on a CV interval of 3.3 ms which is in accordance to the ITU-T G.8114 standard. Both units were configured to wait for three lost CV messages before switching to the backup path. The results showed a switchover time of 34 ms and an out-of-service time during restoration of 19 ms.

Path Redundancy Using CFM

RAD demonstrated the use of CFM on customer edge devices (outside the metro areas) to determine path failure in the network and allow the device to switch to an alternative path.

In this demonstration two TMCs were configured between two RAD ETX-202 devices across the T-MPLS cloud. One path was defined as primary while the second configured as a backup. The two RAD ETX-202 units exchanged CFM Loopback messages across the primary path and noted when a failure in the primary path occurred. Upon detection, the units shut down the primary path and switched to the backup path.

We recorded less than 185 ms out of service time during failure. Restoration could not be tested because the implementation is non-revertive.

Resiliency through MPLS

Several resilience mechanisms exist in MPLS and most has been tested in previous MPLS World Congress Interoperability (MPLSWC) events. MPLS Fast Reroute (FRR) is one such mechanism that was tested in MPLSWC 2006 and showed under 50 ms rerouting times (the white paper is available on EANTC's web site). MPLS Fast Reroute has the ability to protect both against node and link failures.

Dual Homed Multi-Tenant Units (MTUs). Dual homed MTUS are discussed in two specifications related to multipoint MPLS services, RFC 4761 and RFC 4762. Since the multipoint service protocol H-VPLS is a realization of a hub and spoke design with the MTU acting as a spoke and the PE router as a hub, the standards authors saw a need to address an inherent drawback of hub and spoke – the spoke could easily be separated from the hub and, therefore, lose connectivity. In order to protect the MTU from such a disaster, an additional backup pseudowires can be defined to a separate provider edge (PE) router. If the primary pseudowire fails, failover to the secondary pseudowire is facilitated by use of native MPLS protocols or other mechanisms.

Four MTUs participated in this test: Cisco Catalyst 3750-ME, MRV OS910, Telco Systems T-Metro-200 and World Wide Packets LE-311v. They were homed to a combination of the three PE routers in the MPLS metro aggregation network: Alcatel-Lucent 7450 ESS-1, Cisco 7604 and Juniper MX960.

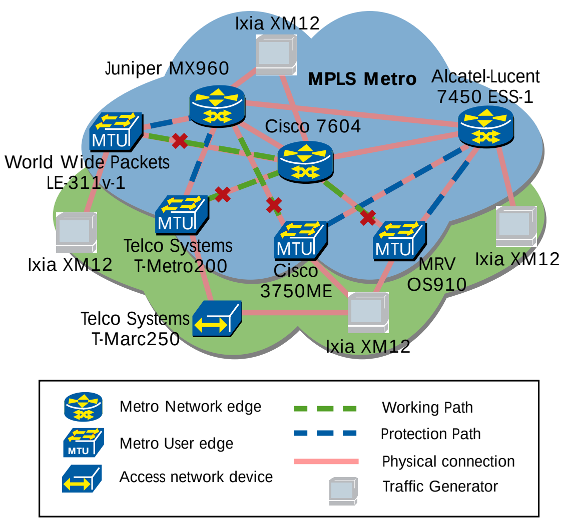

Figure 11: H-VPLS Dual Homed MTUs

The standard describes the concept of MAC address withdrawal to further reduce convergence time by explicitly withdrawing learned MAC addresses from the VPLS domain. Alcatel-Lucent 7450 ESS-1 and MRV OS910 supported MAC address withdrawal.

A number of the vendors succeeded in switching over to the backup pseudowire within less than 100 ms. We would have required more testing time to evaluate the root cause for inconsistent results in other configurations. Vendors will investigate this topic further.

MPLS Fast Reroute. Two Telco Systems T-Metro-200 MTU devices were connected in a ring topology to the Alcatel-Lucent 7450 ESS-1 router. The primary path traversed Ceragon's FibeAir IP-MAX2 microwave link while the secondary path was directly connected to the Alcatel-Lucent router. We simulated a failure of the wireless link by shutting down both wireless transmitters. Fast Reroute mechanisms were evaluated successfully.

Cisco demonstrated a link failure scenario with a traffic outage under 50 ms in a three-node REP ring built with Cisco ME 3400 family switches.

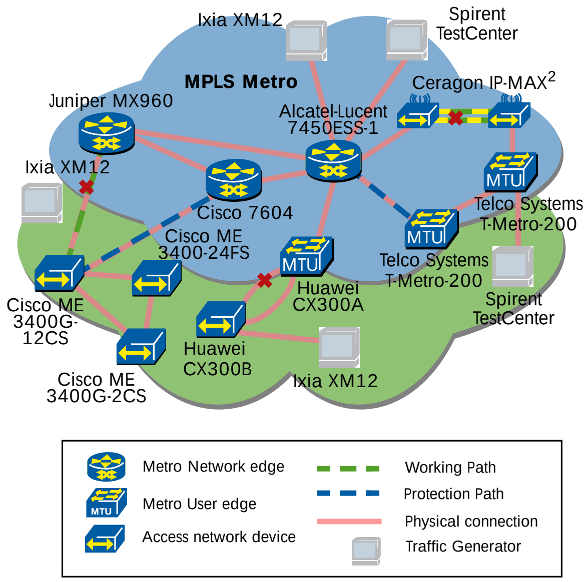

Figure 12: MTU Protection – MPLS Fast Reroute and Ethernet Resiliency

Native Ethernet Resiliency

Several native Ethernet solutions were demonstrated at the event. One such solution was Cisco's Resilient Ethernet Protocol (REP). The protocol provides an alternative to Spanning Tree Protocol (STP) in the access network to control network loops while handling link and node failures. Cisco demonstrated a link failure scenario in a 3-node REP ring with 3 vlans over Cisco ME 3400 family switches.

Cisco also demonstrated its Flex Link feature which allows a switch to make a local protection decision by designating an interface as a backup for another one. Flex Link provides an alternative solution to the Spanning Tree Protocol (STP) in order to provide link and node redundancy in hub-and-spoke topologies. In the demonstration, a Cisco ME 3400G-12CS was dual homed to a Juniper MX960 (configured as primary NPE) and Cisco 7604 (secondary NPE). Upon failure of the primary link, the Cisco ME 3400G-12CS switched to the secondary link in under 50 ms.

Harris Stratex Networks Eclipse wireless transport platform demonstrated resilience of wireless Ethernet networks using a ring-based architecture. They demonstrated their Resilient Wireless Packet Ring (RWPR) protocol which uses Ethernet's Rapid Spanning Tree protocol (RSTP) in combination with a rapid failure detection algorithm to improve convergence times.

Huawei demonstrated their Rapid Ring Protection Protocol (RRPP) between Huawei CX300B and Huawei CX300A. The protocol is meant to provide fast convergence in ring topologies regardless of the number of nodes in the ring. In a non-failed condition, master node will maintain a port in a blocked mode. Once a link failure occurs, the node detecting the link failure will send a "link_down" message immediately to the master node, which then opens the blocking port to forward traffic.

Telco Systems demonstrated their Ring Protection Protocol between Telco Systems T5C-24F and Telco Systems T5C-24G. The protocol is meant to provide fast convergence in an Carrier-Ethernet based ring regardless of the number on nodes in the ring or various types of other vendors sharing the same ring as illustrated in the MPLS Metro area of the network topology for this event.

End-to-End Services

Two types of end-to-end UNI services were implemented in the test network, spanning all metro clouds and the core. Point-to-point Ethernet Private Line services (EPL) were defined from UNI to UNI and provided a high degree of transparency for Service Frames between the UNIs. The goal of delivering identical Service Frames at the destination UNI as were seen in the source UNI was verified using the Ixia XM12 and Spirent TestCenter. As described in MEF 6 "Ethernet Services Definitions - Phase I" the whole UNI was dedicated to the service and no Service Multiplexing was allowed.

The second end-to-end service configured was Ethernet multipoint-to-multipoint (EP-LAN). The service was configured on *all* customer edge devices and was verified using Shenick DiversifEye 8400 emulating Triple Play Clients and Servers. The emulated IP based Multicast video server, as well as unicast HTTP and SIP servers was placed in the MPLS cloud and the triple play clients were emulated on a dedicated port on various customer edge devices.

In both cases provisioning the services across the diverse clouds was extremely time consuming. Engineers were required to manually configure the services on the UNIs and provide mapping configuration from the UNIs to the transport (regardless of the transport technologies).

Emulation of ATM and TDM services over Metro Ethernet

RAD demonstrated the emulation of TDM and ATM over the PBT and T-MPLS metro clouds and across the core. The service emulation used E1 TDM and E1 ATM as is typically seen in cellular base stations. Both E1 TDM and

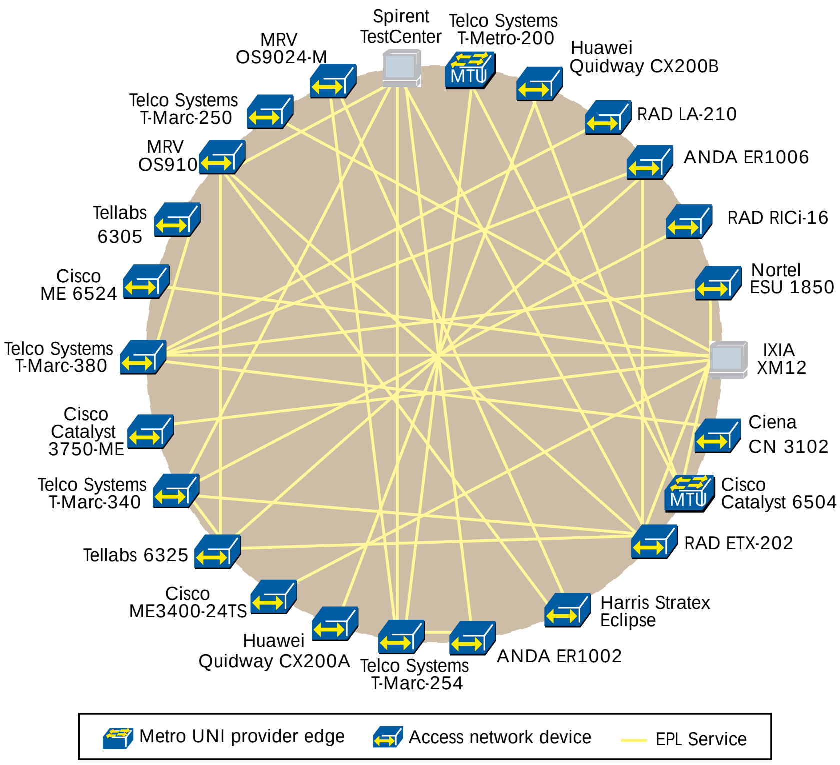

Figure 13: End-to-End Ethernet Private Line (EPL) Services

ATM services were generated into RAD Hub Site Gateway ACE-3200 tunneled over the metro clouds and terminated on RAD's ACE-3400. Timing synchronization was maintained to prevent clocking issues in the end-to-end ATM and TDM services.

Flexible UNI

In addition to the standard UNI services defined by the MEF, Cisco demonstrated the "flexible UNI" service on the 7604 router in the MPLS Metro/Aggregation cloud. "Flexible UNI" facilitates fine granularity Ethernet frame matching, VLAN tag translation and manipulation and flexible service mapping. Cisco showed the “flexible UNI” using one physical port, where several service instances were defined. Each service instance selectively matched single IEEE 802.1Q VLAN tags and double tags based on unique or a range of or list of tags. Additionally, the demonstration included VLAN tag manipulation such as pop, push and translation, and mapping service instances to VPWS, VPLS and L3/VRF.

Quality of Service Tests

Several Quality of Service (QoS) tests were planned for the event. The tests focused on the devices' abilities to enforce the MEF defined service attributes per UNI, EVC or Ethernet's CoS bits. One additional test was defined to verify hierarchical QoS capabilities over the UNI.

Four vendors were interested in performing the tests and the Ixia XM12 tester was used to verify the results, however, due to time constraints, the testing could not be properly completed.

Acknowledgments

We would like to thank Ralf-Peter Braun and Thomas Kessler from T-Systems for their extensive support during the hot-staging event.

We would also like to thank Mark Lum for his support.

Many thanks to Jonathan Morin and Kyle Price from the University of New Hampshire Interoperability Lab (UNH-IOL) for helping to move things forward at the hot-staging test.

The test would not have been possible without the steady flow of coffee and smiles supplied by Melanie Schern.

This document has been edited by Jambi Ganbar, Sergej Kaelberer, Carsten Rossenhoevel and Gabriele Schrenk (EANTC).

Acronyms

Term	Definition
B-MAC	Backbone MAC
B-Tag	Backbone Tag
BSC	Base Station Controller
BNC	Broadband Network Connector
CBS	Committed Burst Size
CE	Customer Edge.
CE-VLAN	Customer Edge Virtual LAN
CFM	Continuous Fault Management
CIR	Committed Information Rate
CM	Color Mode
CoS	Class of Service
DTE	Data Terminal Equipment
EBS	Excess Burst Size
EEPP	End to End Path Protection
EIR	Excess Information Rate
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.
EVC	Ethernet Virtual Connection
FCS	Frame check sequence (4 bytes at the end of Ethernet frame)
LAG	Link aggregation Group
LSP	Label Switched Path
MPLS	MultiProtocol Label Switching
MTU	Multi Tenant Unit
MTU	Maximum Transmission Unit
Metro NPE	Metro Network Provider Edge device
OSPF	Open Shortest Path First
PBB	Provider Backbone Bridge
PBB-TE	Provider Backbone Bridge Traffic Engineering
PE	Provider Edge.
QoS	Quality of Service
RNC	Radio Network Controller
RSVP-TE	Resource reSerVation Protocol Traffic Engineering
SLS	Service Level Specifications
STP	Spanning Tree Protocol
TMP	T-MPLS Path
TMC	T-MPLS channel
T-MPLS	Transport MPLS
UNI	User to Network Interface
Metro UPE	Metro UNI Provider Edge device
VPLS	Virtual Private LAN Service
VPWS	Virtual Private Wired 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 Framework (MEF 4)

Metro Ethernet Services Definitions Phase I (MEF 6)

Abstract Test Suite for Ethernet Services at the UNI (MEF 9)

Abstract Test Suite for Traffic Management Phase 1 (MEF 14)

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 802.3, clause 43)

Media Access Control Parameters, Physical Layers, and Management for Subscriber Access Networks, also referred as "Ethernet in the First Mile"(IEEE 802.3, clause 56 - 67)

Virtual Bridged Local Area Networks - Connectivity Fault Management (IEEE 802.1ag/D6.1-D8.1)

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

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

Provider Backbone Bridges Traffic Engineering (IEEE 802.1Qay)

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

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

Operation & Maintenance mechanisms for T-MPLS layer networks (G.8114)

Characteristics of Transport MPLS equipment functional blocks (G.8121)

Linear protection switching for Transport MPLS (T-MPLS) networks (G.8131)

GFP Frame mapping into PDH (G.8040)

Virtual Concatenation of PDH (G.7043)

Operation & Maintenance mechanism for MPLS networks (Y.1711)

Protection switching for MPLS networks (Y.1720)

Link Capacity Adjustment Scheme (G.7042)

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 Conferences

29 Bressenden Place
London SW1E 5DR, UK
Tel: +44 (0) 20 7017 7483
Fax: +44 (0) 20 7017 7825
registration@iir-conferences.com
www.carrierethernetworld.com

This report is copyright © 2007 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.