First European Multi-Vendor Carrier Ethernet Interoperability Test

Carrier Ethernet World Congress
Berlin, September 2005

Introduction

(See page 10 for an executive summary)

This report describes the results of one of the first public multi-vendor Carrier Ethernet interoperability tests in Europe. As opposed to previous such events, we constructed a fully functional, end-to-end Carrier Ethernet network with business and residential applications. 12 different vendors participated in the event with devices ranging from customer premise equipment through wireless aggregation units to high density/high capacity core routers.

Overall, the results were very promising — some aspects worked immediately in the multi-vendor environment. We were successful in reaching the goal of creating a scalable end-to-end multi-vendor Carrier Ethernet network carrying business and residential triple play services. In other areas of the test, the experiences made were helpful to identify potential improvements.

The interoperability event at the Carrier Ethernet World Congress 2005 has been organized by the European Advanced Networking Test Center (EANTC) and sponsored by the Metro Ethernet Forum (MEF). T-Systems supported the test.

Hot-staging at EANTC (Berlin, Germany)

The interoperability tests detailed in this document were conducted using Carrier Ethernet switches, MPLS routers, emulators, as well as customer premises equipment from various vendors, during a *hot-stage* event. Through several rounds of testing and methodology refining, a final network of interoperable devices was successfully constructed. This network and the test results were demonstrated at Carrier Ethernet World Congress 2005 in Berlin, September 13–15, 2005.

The test scenarios designed specifically for this showcase by EANTC were based upon the MEF technical specifications and reviewed by participating vendors and T-Systems.

To ensure the event's success, a one week hot-staging event with all the participating vendors was conducted before Carrier Ethernet World Congress. The hot-stage took place at EANTC's labs (European Advanced Networking Test Center) in Berlin, Germany.

Participants and Devices

The following companies and devices demonstrated their interoperability in the test event:

Actelis Networks	MetaLIGHT 1300 MetaLIGHT 50
Agilent Technologies	N2X
Alcatel	7450 ESS-7 7450 ESS-1
Cisco Systems	7606 3750M
IXIA	1600T 400T
RAD Data Communications	Egate-100 ETX-102 ETX-202 Gmux-2000 IPmux-14 RICi-8E1
Riverstone Networks	15008 RS3100
Shenick	DiversifEye
Siemens	Surpass HiD 6650
Spirent Communications	SmartBits 6000C Abacus 5000
Stratex Networks	Eclipse
Tpack	6 Gbps Carrier Class Ethernet Reference System (Millburn)

Test Areas and Test Plan

The following section describes the test plan in detail. Results are documented on page 4.

EANTC and T-Systems created the test plan between May and July 2005. The test plan was reviewed and critiqued by participating vendors.

The network topology was designed to be representative for a typical service provider Metro Ethernet network. The design included all devices under test:

- An MPLS backbone ring with three nodes and redundant links
- A »city pop« Ethernet aggregation network with four links, organized as a ring with uplinks to the MPLS cloud
- Extensive CPE infrastructure attached to the city pop via Ethernet fiber and via SDH/PDH access networks: business services, voice over IP and analog phones, TDM applications, and video clients
- A second Ethernet access network attached to the MPLS backbone emulated a »service area«, attaching the video and voice over IP servers as well as emulated business customers

All tests were run over this network. Primarily, test cases were constructed around MEF 2 (Requirements and Framework for Ethernet Service Protection).

- **E-Line and E-LAN support**. The MEF defines two fundamental transport services, a point-to-point E-Line service and a multipoint-to-multipoint E-LAN service. In this test, E-LANs and E-Lines were to be created end to end throughout the backbone and access network.
- **Multicast**. For *Triple Play* video broadcasting services, multicast support in the Ethernet aggregation network is mandatory. IGMP (Internet Group Management Protocol) snooping ensures multicast streams are distributed only to subscribing ports in the Ethernet domain. The tests included transmitting multicast video signal across *E-LANs* in order to verify that only set top boxes that subscribe to the signal are receiving the desired content.
- **Bandwidth Profiles at the UNI**. Service providers often use traffic classification and bandwidth limitation enforcement (Committed and Excess Information Rate) for different applications at the user-network interface (UNI) facing the customer. These functions were verified in the test. The requirements for Bandwidth Profiles at the UNI are clearly defined by the MEF and are required in order to comply with MEF specifications.
- **End-to-end Quality of Service (QoS)**. To properly cope with congestion in the carrier network and to maintain service-level agreements, mechanisms to prioritize traffic classes are required. This test aimed to define four traffic classes for voice, video, management, and high-speed Internet traffic, and to emulate congestion situations in the core of the network.
- **MPLS Protection**. Various mechanisms to implement a backbone for a layer 2 Carrier Ethernet network exist, one of them being MPLS. The Virtual Private LAN Service (VPLS) standard enables multipoint-to-multipoint Ethernet services, and Ethernet pseudo-wires provide point-to-point services.

  Node and link protection in the backbone are important topics for service providers. We verified pre-allocated MPLS backup tunnels (no signaling protocol required). Backup tunnels and fast reroute allow for the realization of the MEF goal, parallel with other transport mechanism and industry standards, of sub-50 ms line and node protection.
- **Ethernet Protection**. In native Ethernet based aggregation networks, protocols like Rapid Spanning Tree (RSTP) are applied. We verified RSTP interoperability and restoration times. The standardization work in different forums will in the coming future focus on how to combine e.g. RSTP and MPLS protection switching to optimize protection times even more to come down in the 50 ms space.
- **End-to-End Performance**. The MEF defines delay, delay variation and packet loss performance metrics for *E-Lines*. We intended to verify these end-to-end metrics by attaching traffic generators to the network edge.
- **TDM Support**. Legacy TDM application support is important in a next-generation packet network. We verified whether circuit emulation solutions worked through the backbone.
- **Versatile Access Technologies**. We verified Carrier Ethernet services over different access technologies

Protocol Support

Items that are not applicable to analyzers and CPE devices have been grayed out.

	Jumbo Frames (< 1522 Bytes)	Bandwidth Profiles[a]	Multicast (IGMP Snooping)	Q-in-Q VLAN Stacking	Link Aggregation	Rapid Spanning Tree
Actelis MetaLIGHT	•	•	•	•	•	•
Alcatel 7450	•	•	•	•	•	•
Cisco 7606	•	•	•	•	•	•
Cisco 3750M	•	•	•	•	•	•
RAD ETX-102/202	•	•		•	•	
RAD Egate-100				•		
RAD Gmux-2000	•				•	
RAD IPmux-14	•	•		•		
RAD RICi-8E1	•	•		•		
Riverstone 15008	•	•	•	•	•	•
Riverstone RS3100	•	•	•	•	•	•
Siemens Surpass HiD6650	•	•	•	•	•	•
Stratex Networks Eclipse	•				•	•
Tpack Millburn	•	•	•	•	•	•
Agilent N2X	•	•	•	•		
IXIA 1600T	•	•	•	•		•
Shenick diversifEye	•	•	•	•		
Spirent SmartBits 6000	•	•	•	•		•

a. Analyzers do not implement rate enforcement, but test it

Interoperability Test Results

The hot-staging event had two targets: First, as in most interoperability test events, the test event sought to confirm and improve the interworking of vendors' implementations; second, the test served to verify clarity and completeness of the current standards. The findings will be communicated to relevant industry forums and standards organizations.

The following sections summarize results of all test areas.

Results: MPLS Backbone Services

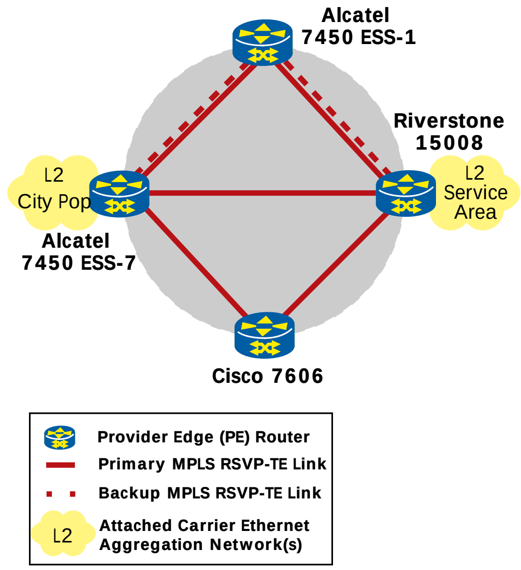

After all vendor shipments had arrived and devices had been powered up, the MPLS backbone was quickly established. RSVP-TE was used for transport signaling. Ethernet pseudowires and non-hierarchical Virtual Private LAN Services (VPLS) were shown to be interoperable between Alcatel 7450, Riverstone 15008, and Cisco 7606 equipment. Backup paths were established between the left side (emulated city pop) and right side (emulated server area) of the backbone. Vendors decided to employ backup paths instead of MPLS fast rerouting since complex protocol extensions (RSVP-TE) were not require for the realization of backup paths. While this is an elegant solution for such a short test event each backup tunnel needs to be created manually.

We measured restoration time to be within industry standards of sub-50 ms. This value represents switch-over time for one tunnel. It may be different for many simultaneously active tunnels.

Results: Business Services (E-LAN and E-Line)

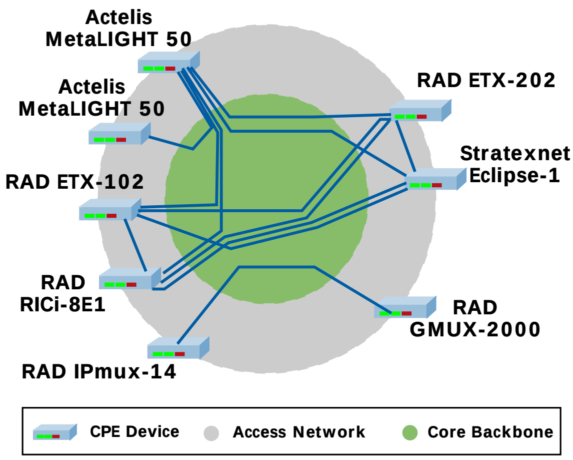

E-Line Services

As in a real Carrier Ethernet network, point-to-point and multipoint-to-multipoint services were created in the test environment. We configured an E-LAN service with multiple network access locations. The service employed VLAN stacking, enabling the customer to transport an arbitrary number of customer VLANs without additional provisioning by the service provider. Basic E-LAN interoperability was reached between all CPE vendors including across different access technologies such as direct Ethernet fiber, SDH/PDH networks, and wireless point-to-point links.

There were no technical interoperability issues with E-Lines; however, manual provisioning of each individual E-Line on every node is a time consuming process. A common edge-to-edge provisioning platform based on standardized configuration mechanisms would have sped up testing in this multi-vendor network.

Results: Residential Services (Triple Play)

Next, we configured VLANs for video, voice, and high-speed Internet access for emulated residential customers. In contrast to E-LANs for business services, no VLAN stacking was used; each of the three applications were carried through the backbone as a flat VLAN. (For the purpose of the test, we ignored security and billing mechanisms that are typically required in a real carrier network.)

Again, multi-vendor interoperability for unicast data transfer was quickly reached. We verified interoperability using the following applications:

- The Shenick diversifEye emulator supplied a video server and video clients to test real video streams. In addition, PC-based video clients were attached to the network receiving the data stream generated by Shenick.
- The Spirent Abacus 5000 emulated a voice over IP server and voice over IP clients attached to RAD ETX-102 and RAD ETX-202, proving that at least 8 calls/second could be forwarded across the whole network without issues. The most impressive voice path was: RAD Ethernet CPE, Stratex long-distance wireless Gigabit Ethernet link, VPLS backbone using Alcatel and Riverstone devices, Ethernet aggregation network with a Siemens switch, RAD Ethernet-over-8xE1 bundle CPE, RAD Ethernet CPE.

Results: Ethernet Protection

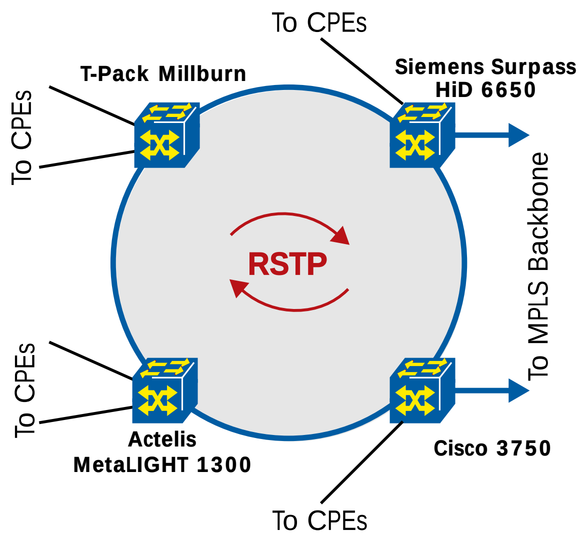

We verified Rapid Spanning Tree (RSTP) interoperability and switch-over times in the Ethernet aggregation network. Four devices participated in the test: Actelis MetaLIGHT 1300, Cisco 3750, Siemens Surpass HiD 6650, and T-Pack Millburn.

A few misunderstandings occurred because Rapid Spanning Tree (RSTP) relied on Multiple Spanning

Tree (MSTP) in one vendor's implementation. MSTP runs a separate spanning tree protocol on each VLAN — which was not our goal because of the large number of VLANs transported through the aggregation network. After each device had been configured correctly, interoperability was reached and RSTP converged with Siemens Surpass HiD 6650 being the root bridge. Typically, Spanning Tree blocked the link between Actelis MetaLIGHT 1300 and the Tpack Millburn.

Restoration times were measured with the Agilent N2X attached to the Actelis MetaLIGHT 1300 and the Cisco 3750:

Action	# Addrs	Restoration Time	
Break Link	1	250 ms	+/- 50 ms
Restore Link	1	350 ms	
Break Link	8,000	3,000 ms	+/- 500 ms
Restore Link	8,000	3,000 ms	

We experienced a number of broadcast storms during the test session. The reason was that the VPLS service over MPLS extended the Ethernet cloud across the core network, creating a loop. Initially MPLS was connected to the aggregation network with two redundant links via the Siemens Surpass HiD 6650 and the Cisco 3750. However, RSTP was not configured over the VPLS domain — in general, it is common practice not to run spanning tree protocols in wide-area networks. The mechanisms to transport Ethernet BPDUs across the MPLS cloud exist in all core vendors, however, due to the short duration of the event the decision was made to remove one of the aggregation to core links. Once the link was removed the Ethernet loop problem was resolved.

Troubleshooting Ethernet broadcast storms proved to be a difficult task. The only viable solution was to temporarily disable each link. A standardized method to improve Ethernet network topology debugging would be welcomed by engineers.

Fortunately, all devices under test survived the service VLAN broadcast storms easily. On one instance, one of the devices became stalled, when the management network experienced a storm.

Results: Multicast

Ethernet switches do not actively participate in IP multicast protocols, but they *snoop* (listen to) Internet Group Management Protocol (IGMP) and maintain subscriber status for each multicast group.

The only alternative to IGMP snooping is to broadcast all packets to each destination port in the VLAN. We started with this basic configuration and found it working throughout the MPLS backbone and Ethernet aggregation network.

Next, IGMP snooping was enabled on all devices including the MPLS provider edge routers. (VPLS-aware edge routers can implement IGMP snooping and other IGMP features.) Shenick verified that nine multicast video streams could be transported over the core.

Unfortunately, multicast transport through the backbone became unstable at some point when RSTP tests were carried out in parallel by coincidence. It was not possible to forward multicast traffic end-to-end with IGMP snooping enabled. Even though end stations tried to subscribe to multicast groups, traffic was not forwarded. Time did not permit a full investigation for the root cause of the problem. In the demo network, unicast video streams will be transmitted. IGMP tests need to be repeated in the future.

Results: UNI Bandwidth Profiles

Bit	Use	CIR	EIR	CBS
7	Mgmt	1 Mbps	1 Mbps	None
5	Voice	1 Mbps	10 Mbps	10 kb (1ms)
3	Video	100 Mbps	200 Mbps	200 kb (1 ms)
0	Best Effort	—	Full rate	full rate (1 s)

We configured bandwidth profiles at the UNI compliant to the Metro Ethernet Forum standard MEF10. Rate limiting was then verified using Agilent N2X, Ixia 1600T and Spirent SmartBits 6000 load generators.

Two devices were tested for UNI bandwidth profiles, the Cisco 3750 and Tpack Millburn. They implemented UNI bandwidth profiles correctly. In another case, bandwidth units were misconfigured

(10 Gbit/s instead of 10 Mbit/s). We recommend to standardize the bandwidth units used in command line interfaces.

Results: End-to-End Performance

We verified end-to-end forwarding delay, jitter, and packet loss through the backbone over selected E-Lines. The Agilent N2X measured end-to-end maximum Ethernet propagation delay of only 0.5 ms under normal operation and 1.4 ms during an RSTP switch-over. No packet loss was observed.

The network performance was superb. All devices including the edge systems showed a very small forwarding delay, virtually no jitter and lossless transport. These figures were determined when the backbone network was not congested.

Results: TDM Support

RAD verified functionality of TDM over Ethernet E-Lines The IPmux-14 and Gmux-2000 equipment set up tunnels for TDM traffic over an E-Line. The Ethernet backbone network showed sufficient performance for TDM transport and voice quality was as high as expected from TDM systems.

Results: Wireless Support

Stratex Networks provided an Eclipse 23GHz wireless link with Switched Ethernet Gigabit interfaces. The wireless link was included in end-to-end interoperability tests, as a CPE from RAD was connected behind it. Throughput was measured to be 308 Mbit/s and the latency was below 500us with no packet loss in a reference configuration. Antennas were bypassed in this indoor test event. Queuing according to the assigned class of services was also verified.

Summary of Test Results		
MPLS Backbone	RSVP-TE	OK
	VPLS (non-hierarchical) and Pseudowires	OK
Carrier Ethernet Basics	MAC Frame Size Support up to 1522 bytes (service frames)	OK
	EVC Leakage	OK
	CE-VLAN ID Preservation across a point-to-Point EVC	OK in all but one cases
Class of Service	CE-VLAN CoS Preservation across an EVC.	OK
	Bandwidth Profile per CoS Service attribute	OK, tested for two devices only
	E-LAN Unicast bandwidth preservation	OK
Multicast	E-LAN broadcast and multicast propagation	OK
	IGMP snooping and RSTP combined	Failed for unknown reasons
Protection	Restoration time over MPLS backup tunnels (no fast reroute)	OK, within 50 ms industry standard
	ALNP Path Protection	OK
	E-LAN protection using Rapid Spanning Tree Protocol (RSTP)	Reached 250-3000 ms depending on the number of MAC addresses
Versatile Access	Ethernet service interoperability over different access options (Ethernet fiber, SDH/PDH, wireless links)	OK
TDM Support	Emulated E1 TDM over Ethernet	OK

Final Integrated Test Network

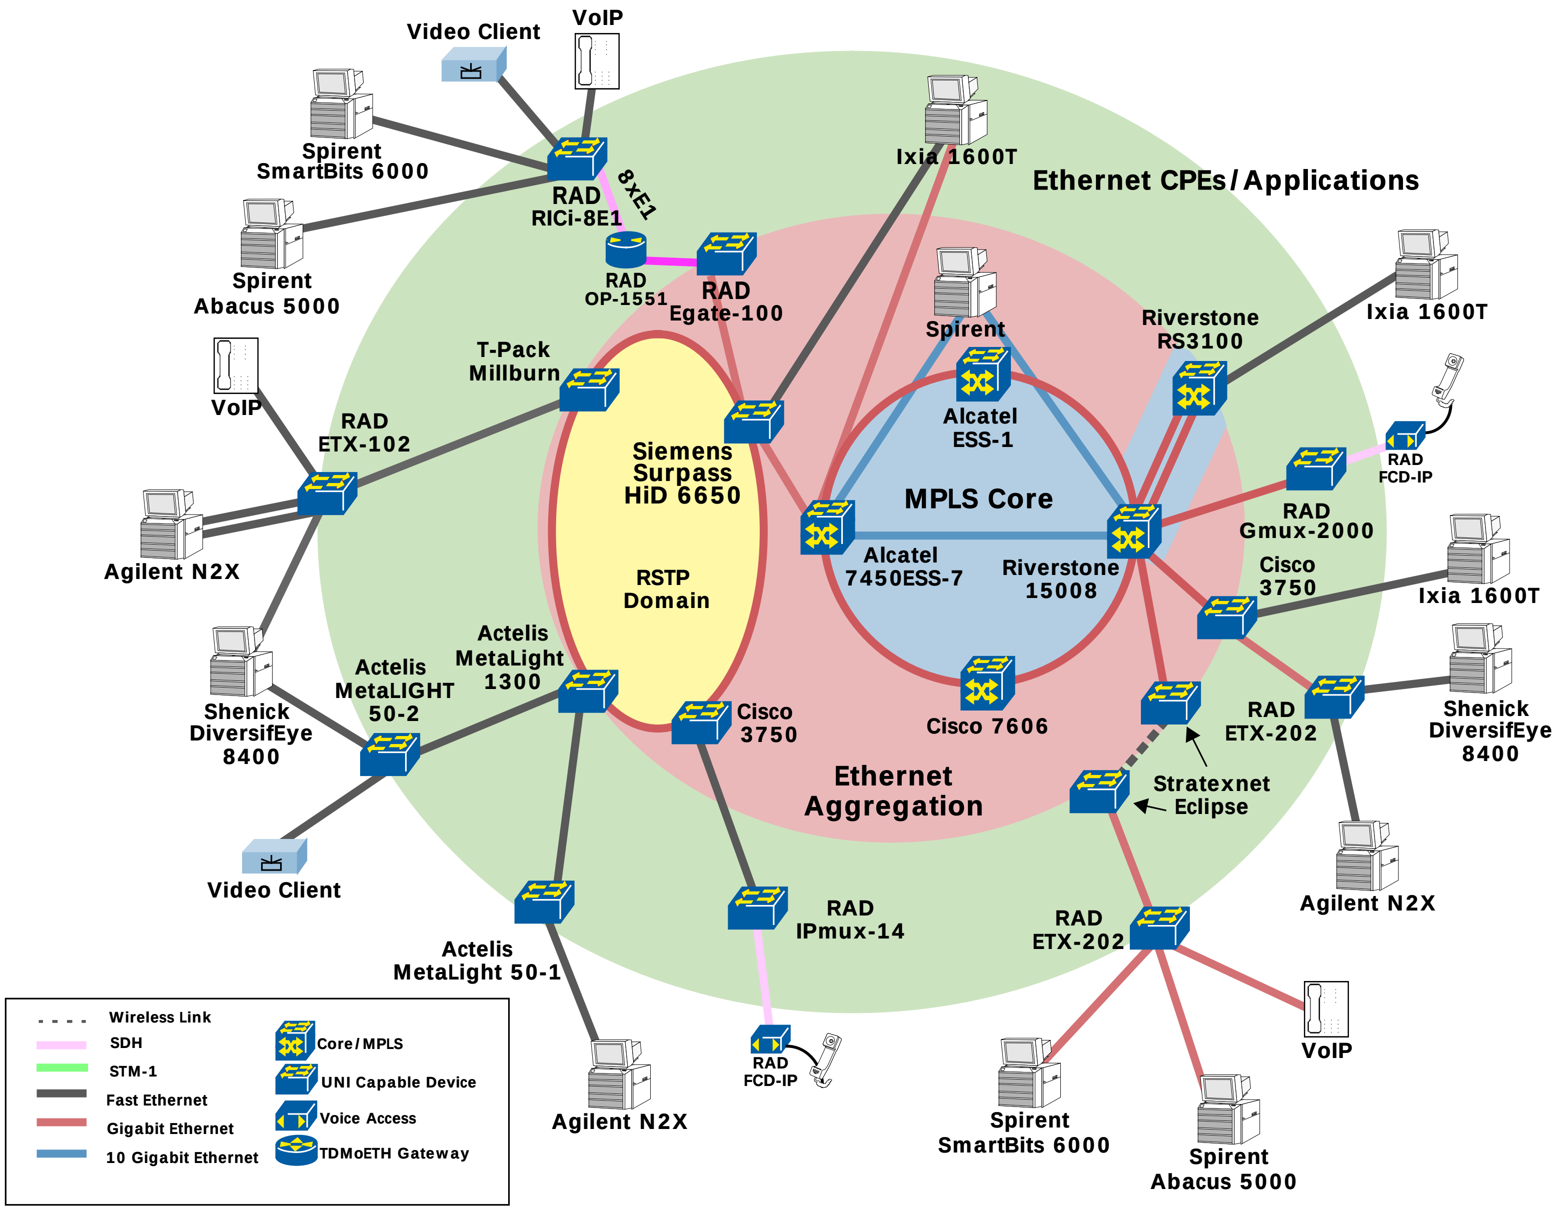

Problem Summary

Problem Area	Description	Temporary Solution, if any	Recommendation
RSTP	No implementation guidelines by MEF Easy to misconfigure, confusion with MSTP running in parallel to RSTP	None	Define guidelines for use of MSTP/RSTP for protection purposes
IGMP Snooping	PDUs not forwarded	Further analyze causes in future interop events	—
Management	Broadcast Storm Debugging	Disable links	Define more advanced methods to locate loops in Ethernet backbones faster and more reliably
Class of Service preservation	In some cases, the CoS bits are not preserved across a device under test when pushing a VLAN tag	Define a single priority at the DUT for all traffic	Become MEF2 compliant (test case groups 12.1 / 12.2).
QoS in the core	Not defined by MEF	Use IEEE 802.1q to prioritize and queue packets	Define guidelines for CoS in an Ethernet backbone

Acknowledgments

We would like to thank Ralf-Peter Braun from T-Systems for supporting the event, guiding test plan development and network design throughout all phases of the event.

Many thanks to Curtis Simonson and Henry He from the University of New Hampshire InterOperability Lab (UNH-IOL) for support during test plan development and hot-staging.

This document has been edited by Carsten Rossenhoevel and Jambi I. Ganbar, EANTC.

Conclusions / Executive Summary

Interoperability based on Metro Ethernet Forum (MEF) specifications has been demonstrated at SUPERCOMM 2005 and at the Carrier Ethernet World Congress. Each highlighted different aspects of the MEF services and Infrastructure.

The Carrier Ethernet World Congress event is unique in that it demonstrated an integrated, end-to-end multi-vendor Carrier Ethernet network. Twelve vendors participated, contributing core backbone switches, aggregation network equipment, and CPE devices. The test plan was developed with help of T-Systems, ensuring a realistic network design and test cases relevant for service providers.

The majority of the test areas could be completed successfully. The interoperability at the Ethernet User-Network Interface (UNI) was outstanding, thanks to the MEF 10 standard detailing the exact service attributes to be used at this interface.

We were able to run triple play residential services and point-to-point (E-Line) as well as multipoint-to-multipoint (E-LAN) services over the Carrier Ethernet network without any issues. Voice over IP, Video on Demand and high-speed Internet worked instantaneously as expected.

The protection tests showed that sub-50 ms restoration times can be reached in an MPLS core while native Ethernet networks achieve typically between 250–3,500 ms. More work is also required to improve multicast snooping interoperability and broadcast storm debugging. In addition, an NNI service attribute standard would be helpful to provide guidance for QoS configuration in the backbone.

The event showed that test equipment vendors are keeping pace with the unique demands of ensuring Carrier Ethernet interoperability. Such demands include critical interoperability measurements which ensure no EVC leakage, verify restoration times and confirm adherence to bandwidth profiles.

In total, the interoperability event proved that Carrier Ethernet is already a multi-vendor capable technology in many aspects. The series of interoperability tests will be continued in the future to address remaining issues.

References

- MEF 2 — Requirements and Framework for Ethernet Service Protection in Metro Ethernet Networks.
- MEF 10 — Ethernet Services Attributes Phase 1
- MEF 11 — User Network Interface (UNI) Requirements and Framework
- IEEE 802.1D-2004 — Media Access Control (MAC) Bridges (revision incorporates 802.11c, 802.1p, 802.12e)
- IEEE 802.1Q — Virtual Bridged Local Area Networks, July 1998
- IEEE 802.3-2000 — Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer definitions
- RFC 3916 — Requirements for MPLS Pseudo-Wire Emulation Edge-To-Edge. Xi-Peng Xiao, Danny McPherson, Prayson Pate. September 2004
- draft-ietf-pwe3-control-protocol-14 — Transport of Layer 2 Frames over MPLS. Luca Martini et.al. December 2004 (work in progress) Expires: June 2005
- draft-ietf-pwe3-ethernet-encap-08 — Encapsulation Methods for Transport of Ethernet Frames over IP/MPLS Networks. Luca Martini et.al. September 2003 (work in progress)
- draft-ietf-l2vpn-vpls-ldp-05 — Virtual Private LAN Services over MPLS. Marc Lasserre, Vach Kompella, et.al. September 2004. (work in progress)

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

Metro Ethernet Forum

19900 MacArthur Blvd. Suite 810
Irvine, CA 92612, USA
Tel: +1.949.250-7188

manager@metroethernetforum.org
www.metroethernetforum.org