



Converged Network Services Using MPLS

**PARIS 2006
PUBLIC INTEROPERABILITY EVENT**

Editor's Note



Gabriele Schrenk
Managing Director

EANTC has been organizing the interoperability showcase at the MPLS World Congress for the last four years, while independently testing MPLS for eight years. As a technology, we consider MPLS mature. Since many topics have been tested for interoperability in the past, the decision to conduct another interop event was not easy. We decided that this year's event would unite all the previ-

ous experiences and test areas while adding new MPLS capabilities testing.

MPLS supports many technologies and services such as Triple Play and Metro Ethernet and allows service providers to converge a large number of networks and services into a single unified backbone. With this ideology we broadened the scope of the testing and set the motto for this year's event to be all inclusive - a single network could be demonstrated to support as wide a range of services as possible.

The widened scope and the success of the past events attracted more interested vendors than ever. A quarter of the participants were newcomers ranging from access to core devices. In total we had 15 participants with over 30 devices. With the success of MPLS and the push of the technology to the network edges, we expect even more implementations to be available soon.

As the scope of the testing and the number of devices increased, so did the findings. Implementation issues specifically related to traffic engineering (RSVP-TE, OSPF-TE) and high availability (Fast Reroute) slowed down the progress of the testing but were mostly overcome by the end of the test event. We observed again that vendors continue to overcome the challenges associated with implementing network services using a common subset of protocol options.

As an independent test lab, we see an urgent need for standards committees and industry forums to reduce the number of protocol options and to clarify implementation options. We believe that a clear definition would help to improve interoperability substantially, and further the deployment of multi-vendor MPLS networks.

Introduction

The MPLS World Congress 2006 interoperability event has been organized and facilitated by the European Advanced Networking Test Center (EANTC) and the University of New Hampshire InterOperability Laboratory (UNH-IOL) and endorsed by the MFA Forum.

The interoperability tests detailed in this document were conducted using MPLS routers and switches, emulators, as well as customer premises equipment from various vendors, during a *hot-staging* event in January 2006. Through several rounds of testing and refining the methodology, a final network of interoperable devices was successfully constructed. This network and the test results were demonstrated at MPLS World Congress 2006 in Paris, February 7-10, 2006.



Hot-staging at EANTC
(Berlin, Germany)

Several new test scenarios were designed specifically for this showcase. In addition, previously used test plans were employed for regression testing since we intended to test *converged* network services:

- »99.999%« carrier-grade high availability is one of the cornerstones of MPLS benefits. The Fast Reroute mechanism uses automatic pre-established backup paths to realize fast (sub-50 milliseconds) switchover in case of link or node failure. We had run Fast Reroute tests with a small number of vendors in 2004, and wanted to expand on these tests.
- Differentiated services have been available over MPLS in a simple IP quality-like fashion for a while. Now, new standards are on their way to enhance the integration of traffic engineering and application-specific differentiation. We intended verify the multi-vendor readiness of implementations.
- Multi-vendor layer 3 (IP) VPNs can be considered mature by now. Except for carrier-carrier interworking protocols, multicast and IPv6 traffic

forwarding, they are proven to be interoperable regarding functionality and scalability, as shown, for example, during the MPLS World Congress 2004 and 2005 interoperability events.

- The standards for Ethernet and ATM pseudowires have existed for a long time. Previous tests have shown that there are a lot of mature and stable implementations. We see a growing number of vendors implementing pseudowires and out of the 15 participants in the event most vendors were interested in verifying interoperability of their pseudowire implementations.
- Multipoint Ethernet services (Virtual Private LAN Service, VPLS) are offered by a growing number of carriers. The hierarchical part of the protocol (H-VPLS) enables service providers to scale the number of customers and endpoints per customer offered using VPLS, without stressing the backbone network.

Our regression test verified that previous years' results were still valid. We tested scalability of hierarchical VPLS provider edge routers (PE-RS) and multi-tenant units (MTUs) in 2005. This time a total of six PE-RS and two MTU implementations were checked.

- A major topic of the test program was supposed to be multicast traffic forwarding in Ethernet and IP VPNs. There were quite a few hurdles in testing multicast in the context of MPLS — see test results section.
- With the proliferation of Ethernet access we intended to evaluate the relevant access solutions for IP/MPLS core networks, specifically carrier-class Ethernet access solutions and pseudowire access to support Layer 1 (TDM) and Layer 2 (ATM) services.

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

The Interoperability Working Group of the MFA Forum, including EANTC and UNH-IOL, defined the MPLS test plans.

Participants and Devices

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

Agilent Technologies	N2X
Alcatel	1662 PRS 7670 RSP 7750 SR1 / SR7
Ciena	DN 7100
Cisco Systems	12406 CRS-1
Huawei	NE40E
IXIA	1600T
Lucent	CBX 3500
MRV	OSM 207 OS 9024
Nortel	MPE 9500
RAD Data Communications	ACE-3100 ACE-3402 ETX-202 IPmux-14 Gmux-2000 FCD-IP
Riverstone Networks	15008 15101
Spirent Communications	Test Center SPT-5000A
Telco Systems (BATM)	T-Metro
Tellabs	8840
Tpack	Millburn

Test Areas and Test Plan

The following table displays the different areas of testing and the roles and interests the various vendors had for the technology in this event. The table represents all the devices available at the hot-staging event in Berlin and includes test/traffic generators, customer premise equipment (CPE), Provider Edge (PE) and Provider core (P) routers and Multi-tenant units (MTU).

MPLS Protocol Support	RSVP-TE Signaling	OSPF-TE v2 Routing	DiffServ-Traffic Engineering	MPLS Fast Reroute	BGP/MPLS IP VPNs	Ethernet Pseudowires	TDM Pseudowires	ATM Pseudowires	Flat VPLS	Hierarchical VPLS	Multi-Segment Pseudowires
Agilent N2X	•	•	•	•	•	•		•	•	•	
Alcatel 1662 PRS	•	•				•				•	
Alcatel 7670 RSP	•	•	•	•	•	•		•			
Alcatel 7750 SRx	•	•		•	•	•		•	•	•	
Ciena DN 7100	•	•		•	•	•	•	•	•		
Cisco 12406	•	•	•	•	•	•			•		• a
Cisco CRS-1	•	•	•	•	•	•					• a
IXIA 1600T	•	•	•	•	•	•			•	•	
Huawei NE40E	•	•	•	•	•	•			•		
Lucent CBX 3500	•	•						•			
MRV OSM 207	•	•	•			•			•	•	
MRV OS 9024	•	•	•			•				•	
Nortel MPE 9500	•	•			•	•		•		•	• b
RAD ACE-3100/3402								•			
RAD IPmux-14/Gmux-2000						•					
Riverstone 15008	•	•	•	•		•			•	•	
Riverstone 15101	•	•	•			•			•	•	
Spirent AX4000	•	•	•	•	•	•			•		

MPLS Protocol Support	RSVP-TE Signaling	OSPF-TE v2 Routing	DiffServ-Traffic Engineering	MPLS Fast Reroute	BGP/MPLS IP VPNs	Ethernet Pseudowires	TDM Pseudowires	ATM Pseudowires	Flat VPLS	Hierarchical VPLS	Multi-Segment Pseudowires
Spirent Test-Center	•	•	•	•	•	•			•		
Telco Systems (BATM) T-Metro	•			•		•	•		•	•	•
Tellabs 8840	•	•	•	•	•	•	•	•	•	•	
Tpack Millburn	•	•				•			•		

- a. Static Multi-Segment Pseudowires
- b. Static and Dynamic MS-PW

The following section describes the test plan in detail. Results are documented on page 6.

MPLS Signaling and Routing

Test engineers first constructed the backbone network. All test cases required RSVP-TE signaling for MPLS transport and dynamic routing in the backbone using OSPF with traffic engineering extensions.

The OSPF link state databases and link costs were configured carefully to prepare for the DiffServ-Traffic Engineering tests (see below).

Fast Reroute

RFC4090 defines extensions to RSVP-TE to allow for the redirection of traffic to backup LSPs in less than a second. An interoperability test plan has been defined by the MFA Forum and is in final straw ballot under *mpls2005.129.00*. The test plan aims to verify the functionality of Fast Reroute's two topologies (link and/or node protection), the correct handling of RSVP-TE objects defined for Fast Reroute and measure scalability with a realistic number of tunnels. The vendors participating in this test can be positioned as P or PE nodes.

DiffServ – Traffic Engineering

DiffServ-TE followed the MFA Forum interoperability test suite defined in *mpls2004.149.03*. The tests focus on the logical path packets take through an MPLS network and the actions LSRs have to take in order to accommodate differentiated classification for packets forwarded. Specifically the test plan defines the following areas:

- Verify that an LSR can preempt an LSP when bandwidth is insufficient for all LSPs
- Ensure appropriate constraint-based routing behavior (CBR) using OSPF-TE
- Validate TE path calculation
- Verify the correct behavior of the three bandwidth constraint models (MAM, RDM, MAR)

Multicast

The various solutions for transporting multicast traffic over layer 2 and layer 3 MPLS based VPNs have been a subject of a heated debate in the respective IETF working groups recently. We tested L3 VPN multicast functionality according to the IETF working group draft *draft-ietf-l3vpn-2547bis-mcast-01.txt*

For L2 VPN multicast, we investigated testing according to the solution discussed in the IETF »l2vpn« working group in January, that mandated PIM and IGMP snooping in VPLS (*draft-hemige-serbest-l2vpn-vpls-pim-snooping-00.txt*).

However, we found that there are not enough implementations yet that could be tested. We had to adhere to the traditional method of forwarding multicast traffic, using the broadcast mechanism built into the VPLS protocol.

Ethernet Point-to-Point VPNs (Pseudowires)

Point-to-point Ethernet VPN Services over MPLS were tested using the MFA Forum test methodology defined in the test plan *mpls2003.091.03*. The tests covered:

- Label binding and distribution for Ethernet pseudowires via targeted LDP sessions between the provider edge routers
- Data encapsulation of Ethernet and tagged Ethernet frames

Hierarchical VPLS (H-VPLS)

Since VPLS is basically a multipoint extension of point-to-point Ethernet pseudowire links, point-to-point evaluation tests provided a prerequisite for the VPLS tests. They were carried out in accordance to *draft-ietf-l2vpn-vpls-ldp-05*, using the MFA test plan *mpls2003.092.03*.

- VPLS service establishment by label exchange between provider edge routers
- Hierarchical VPLS service establishment for provider edge (PE-RS) VPLS switches
- Hierarchical VPLS configuration for multi-tenant unit (MTU) VPLS switches

Vendors supporting H-VPLS could be either the provider edge (PE) or the multi-tenant unit (MTU) device while participating in the test.

Access Pseudowires

The extension of pseudowires into the access was evaluated with dedicated customer premises equipment and access gateways. Native TDM and ATM services are transported towards the provider edge (PE) using pseudowires and in turn are further forwarded using Multi-Segment pseudowires. TDM pseudowire are implemented in accordance with MFA 4.0 Implementation Agreement, and ATM pseudowires with *draft-ietf-pwe3-atm-encap-10* (one-to-one mode).


Multi-Segment Pseudowires

Multi-Segment pseudowires represent a set of two or more contiguous pseudowire segments that behave and function as a single point-to-point pseudowire. This architecture provides control plane scalability when looking at extending PWs into the metro/access network and allows for inter-domain/inter-provider pseudowire set-up. Due to time constraints this technology was not tested during the hot-staging; however, it will be demonstrated in Paris at the MPLS World Congress event. The demonstration will be conducted in accordance with *draft-ietf-pwe3-segmented-pw* to show manual configuration of multi-segment pseudowires and in accordance to *draft-ietf-pwe3-dynamic-ms-pw-00.txt* (previously known as *draft-balus-bocci-martini-dyn-ms-pwe3-00.txt*) to show dynamic placement of multi-segment pseudowires.

Interoperability Test Results

This section summarizes all the results obtained during the hot-staging week, sorted by test sessions.

Results: Ethernet Point-to-Point Pseudowire Tests



Ethernet Point-to-Point Tunnels

Point-to-point Ethernet over MPLS tunnels («pseudowires») were tested according to the IETF PWVE3 specifications. During the hot-staging event, all tested point-to-point connections interoperated as expected. Ethernet pseudowires were successfully tested among Alcatel 1662 PRS, Alcatel 7750 SR1/SR7, Alcatel 7670 RSP, Telco Systems (BATM) T-Metro, Ciena DN 7100, Cisco 12406, MRV OS 9024 and OSM 207, Nortel MPE 9500, Riverstone 15101 and 15008, Spirent TestCenter (acting as a PE) and Tellabs 8840.

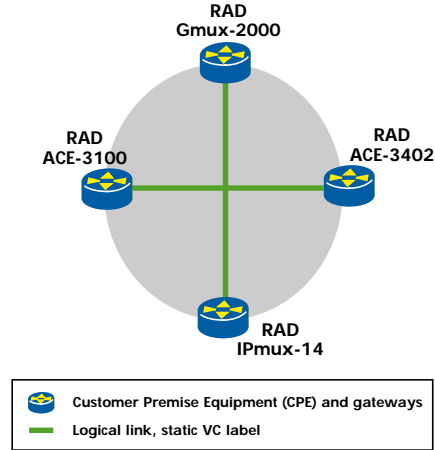
Since scalability had been tested in previous years with up to 2,000 pseudowires established within one transport tunnel between two devices we did not repeat the test.

Unlike previous test events in which some vendors supported only LDP, all vendors supported RSVP-TE signaling for VPN transport labels this time.

Results: TDM and ATM Point-to-Point Pseudowire Tests

The access pseudowire solutions were demonstrated with RAD IPmux-14 and Gmux (TDM pseudowire) and ACE-3100/3402 (ATM pseudowire) access gateways. RAD verified functionality of TDM over MPLS pseudowires

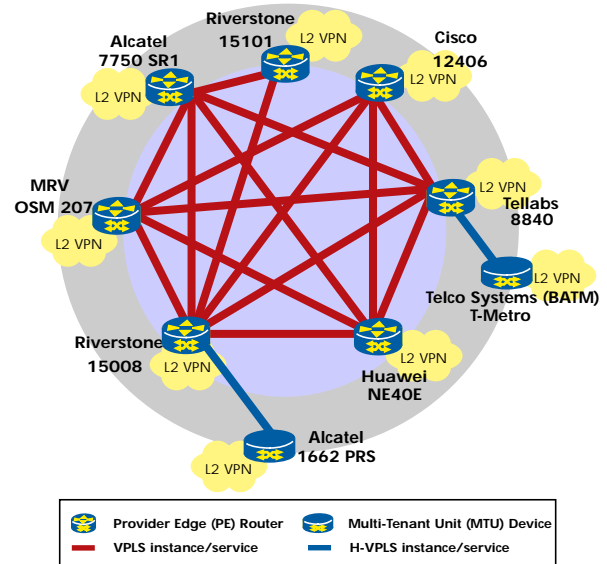
according to MFA Forum implementation agreement 4.0 »TDM Transport over MPLS using AAL1«. The IPmux-14, and Gmux equipment set up label-switched paths for TDM traffic using static labels over Cisco and MRV MPLS routers. Ciena participated in the static label exchange configuration with RAD and MRV, however, due to time constraints the configuration was not fully verified.



ATM and TDM Point-to-Point Tunnels

Due to limited time, only a few vendors focused on the creation of ATM pseudowires. The RAD ACE-3000 devices successfully established an ATM pseudowire tunnel over the backbone using static label assignment.

Results: VPLS and H-VPLS Tests



VPLS / H-VPLS Multipoint Ethernet Services

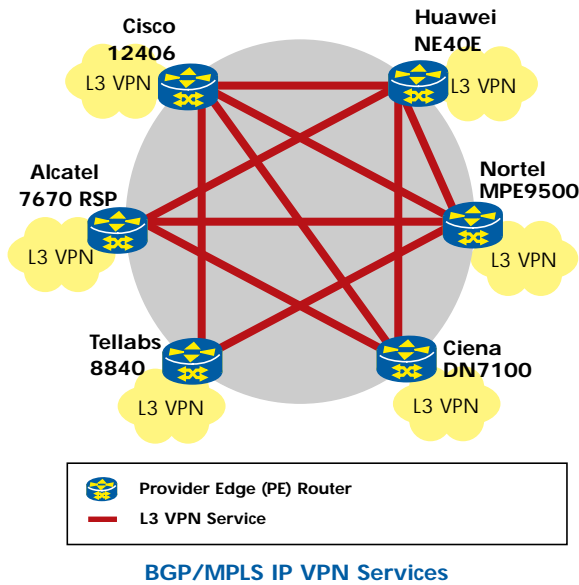
During the hot-staging event, the hierarchical VPLS interoperability tests between Provider Edge (PE-RS) implementations as well as PE-RS and Multi-Tenant Unit (MTU) systems were very successful. As illustrated in the

diagram above, most PE-RS implementations were interoperable without any issues: The six PE-RS routers (Alcatel 7750 SR1, Cisco 12406, Huawei NE40E, MRV OSM207, Riverstone 15008, Tellabs 8840), two emulators (Agilent N2X, Ixia 1600T), and two multi-tenant units (Alcatel 1662 PRS, Telco Systems T-Metro) were able to establish tunnels and exchanged data.

All systems were able to interconnect on the VPLS layer. A few LDP signaling issues created interoperability problems occasionally (details see problem section below), but these were the only source of problems.

Multicast over VPLS. The tests distributed multicast traffic through the VPLS network as broadcast and unknown traffic. At the moment, the IETF has only a preliminary draft on the subject making testing of more advanced solutions impossible.

Results: RFC 2547bis, L3 VPN Tests



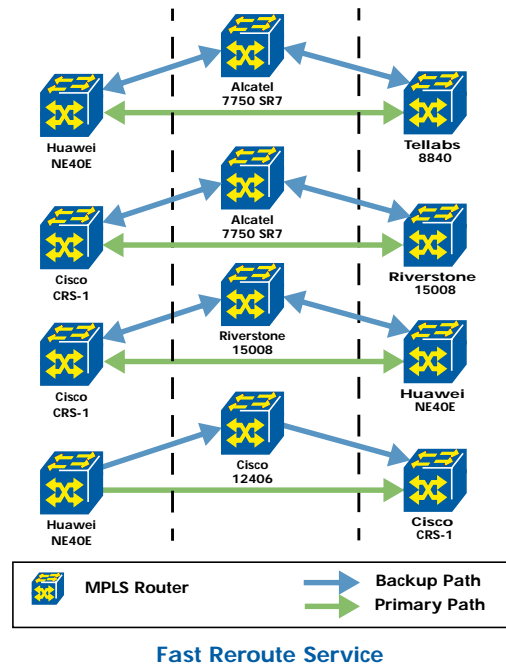
IP Virtual Private Networks were constructed easily. They are one of the oldest applications for MPLS networks so we did not expect any issues. In fact, there were none. The Alcatel 7670 RSP, Ciena DN7100, Cisco 12406, Huawei NE40E, Nortel MPE9500, and Tellabs 8840 routers participated in the test. More participating devices supported BGP/MPLS IP VPNs, but these vendors focused on other areas.

Multicast over MPLS/BGP VPNs. Multicast support over MPLS/BGP VPNs requires substantial protocol addition as mentioned in the test plan section above. We had a first glance at IETF draft implementations with Cisco 12406, Huawei NE40E and Ixia 1600T.

Results: Fast Reroute

The Fast Reroute interoperability tests showed that the support for Fast Reroute is growing. We evaluated six implementations from Alcatel, Cisco (2x), Huawei, Riverstone, and Tellabs. All of the systems were able to setup protected tunnels that were signaled by the Agilent N2X tester (acting as a PE router). We did note that a working implementation does not automatically translate to interoperability — surprisingly, we saw issues similar to our test two years ago!

Most problems were caused by misunderstandings of the IETF Fast Reroute RFC and by RSVP Objects being incorrectly understood. Luckily we were able to overcome all problems and set up the combinations shown in the diagram below. Rerouting times were always below 50 milliseconds, which is an improvement from previous tests. The rerouting times relate to only one tunnel, though; performance tests with many tunnels may show different results.



Ethernet Access to MPLS Core

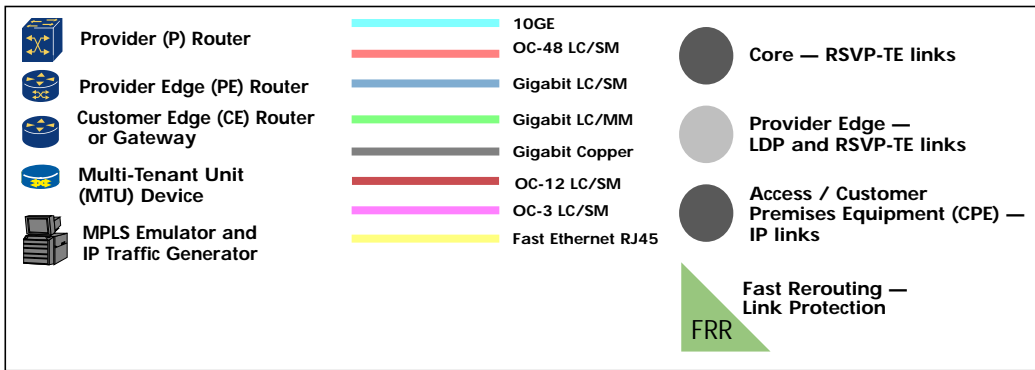
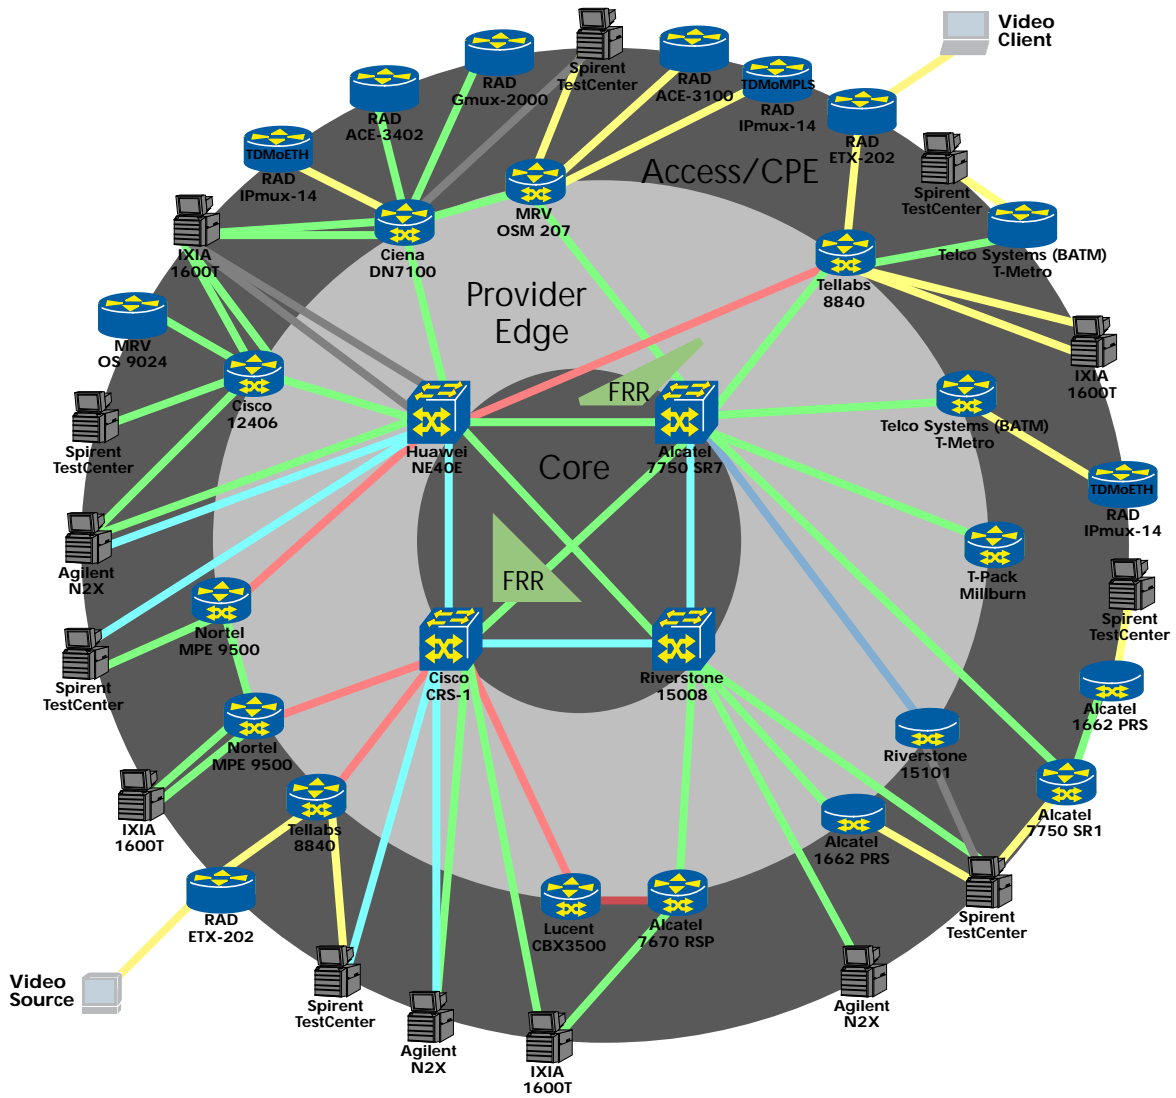
RAD demonstrated pre-standard implementation of IEEE 802.1ag/ITU-T Y.17ethoam Ethernet OAM with ETX-202, an Ethernet NTU. Ethernet OAM Loopback was used for end-to-end path protection by switching over to a backup core Ethernet pseudowire.

Another carrier-class feature of Ethernet access, which was verified at the demo, was support of TDM traffic according to MEF 8 (SAToP encapsulation). The feature was demonstrated with RAD IPmux-14.

Results Summary

Key Features Tested		Results
Pseudowires	Interoperability RSVP-TE	OK
	Data Transfer	OK
	Ethernet tunnels	OK
	Traffic Transfer Over RSVP-TE and LDP Tunnels	OK
	ATM Pseudowires	OK, tested with one vendor
	TDM Pseudowires	OK, tested with one vendor
VPLS	Basic LSP Establishment between PE routers	OK
	Label Exchange Between PE routers	OK
	Forwarding to Unknown MAC addresses	OK, mostly
	Traffic Forwarding	OK
	Tunnel Teardown and Withdraw	OK, tested with 3 vendors
	Hierarchical VPLS PE-RS functionality	OK
	Hierarchical VPLS MTU functionality	OK,
BGP/MPLS IP VPNs	VPN Establishment	OK
	Basic PE Data forwarding	OK
	Backbone Data Forwarding	OK
	Two VPNs with Overlapping Address Space	OK
	VPN Route Uniqueness	OK
	Extranet access with Route Targets	OK
	Customer Control of Routes using target attributes	OK
	Internet access from VPN	OK
MPLS Fast Reroute	Facility Backup LSP Signaling and Creation	OK
	Link Protection	OK
Multicast	Multicast/Broadcast Transport over VPLS	Tests in progress at time of printing
	Multicast over BGP/MPLS VPNs	

Final Integrated MPLS Test Network



Problem Summary

Problem Area	Description	Temporary Solution, if any	Recommendation
LDP	TLV TE setting of experimental bit caused message to be dropped. The Traffic Engineering Data Base was not built.		Resolved by loading different code
	Some vendors established tunnels to Host FECs as mentioned in RFC3036, some supported the new draft ietf-mpls-rfc3036bis which prohibits the usage of Host FEC.	Configure LDP stack so that prefix FEC is being sent instead of a host FEC.	The IETF should verify if incompatible draft updates can be avoided.
	Some vendors had problems to establish stable targeted LDP sessions.	New release installed to suppress topology LDP session.	Fix bug completely
RSVP-TE	Some vendors had problems to decode and encode the RSVP-TE Object label recording (RRO). Tunnels were not established.	When the option was turned off the tunnel came up. Unfortunately Fast Reroute does not work without.	More detailed interoperability tests are needed.
	Illegal bandwidth value in RESV message. No RSVP-TE tunnels could be established.	None	Implementation should be corrected.
OSPF and OSPF-TE	OSPF-TE databases were inconsistent sometimes.	None	OSPF-TE interoperability needs to be improved.
	Some LSA transmitted by a vendor was not supported by others. The OSPF adjacencies couldn't be established.	A new software release ignored this type of LSAs.	The IETF should verify if incompatible draft updates can be avoided.
L2 VPNs	When a vendor was setting up Ethernet PW in VLAN mode they were stripping the VLAN so the other end just dropped the packets.		New software release solved the issue.
	One vendor couldn't run ATM pseudowires over Ethernet as transport layer.	ATM PW could only run over PoS links.	This will be fixed in future software release.
	VLAN labelled traffic sent into VPLS via the core comes out at the edge as 16x as much packets and a second VLAN is appended.	None	Additional tests are necessary.

Conclusion

EANTC has been organising interoperability events at the MPLS World Congress since 2001. Every year the event presented new technological challenges, attracting more participants and allowing us, as an independent test lab, to evaluate what progress was made since the previous year in the world of MPLS protocols and services.

The results from this year's event are reassuring to MPLS vendors and to Service Providers relying on the technology. A larger group of vendors than ever before demonstrated interoperability in all areas pertaining to services. Network services that depend on MPLS for success, for example Carrier Ethernet and IP-based VPN services, are now free to choose from an ever-expanding list of vendors providing MPLS support. Another important achievement is the interoperability of a key backbone component, MPLS Fast Reroute, between a larger than before set of vendors.

As MPLS advances and evolves both in the network core and to the edge, new protocols are being discussed and defined by the IETF. MPLS can only advance into new services and wider reach when suitable protocols become available. We hope that by next year's event, some of the areas, such as Multicast over VPLS and over BGP/MPLS VPNs, Multi-Segment Pseudowires and Carrier's carrier protocols, will be mature enough to enable vendors to implement them and demonstrate interoperability.

MPLS is moving towards a ubiquitous role as a networking technology. As such, more devices that traditionally did not play a role in MPLS topologies also begin to include MPLS stacks and signaling capabilities along with increased protocol support. It is clear to us that the MPLS interop-erability efforts are far from finished.

Acknowledgements

The white paper was authored by Jambi Ganbar, Jonathan Morin, Carsten Rossenhoefel and Gabriele Schrenk.

References

All tests were conducted in accordance to MFA Forum Interoperability Test Suites detailed in the Test Plan section of this publication.



EANTC AG
European Advanced Networking Test Center

Tel: +49 30 3180595-0
Fax: +49 30 3180595-10
info@eantc.de
www.eantc.com



University of New Hampshire
InterOperability Laboratory

Tel: +1.603.862.4212
Fax: +1.603.862.0898
morin@iol.unh.edu (Jonathan Morin)
www.iol.unh.edu



MFA Forum

Tel: +1.510.608.5910
Fax: +1.510.608.5917
amorris@mfaforum.org
www.mfaforum.org

The MFA Forum is an international, industry-wide, nonprofit association of telecommunications, networking, and other companies focused on advancing the deployment of multi-vendor, multi-service packet-based networks, associated applications, and interworking solutions. The Forum currently has more than 50 members.

This report is copyright © 2006 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 trademarks are property of their respective owners.
