

Large-Scale Multi-Vendor Layer 2 VPNs with MPLS

**PARIS 2005
PUBLIC INTEROPERABILITY EVENT**

Introduction

The MPLS World Congress 2005 interoperability event has been organized by the MPLS & Frame Relay Alliance, and facilitated by the European Advanced Networking Test Center (EANTC) and the University of New Hampshire (INH-IOL).

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 stage* event. 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 2005 in Paris, February 16–18, 2005.



**Hot-staging at EANTC
(Berlin, Germany)**

The test scenarios designed specifically for this showcase were based upon the experiences of previous interoperability test events:

- MPLS protocols for the creation of virtual private networks (VPNs) have been verified in interoperability tests before. However, scalability to real-life network sizes in multi-vendor environments has always been an important question for service providers.
- Multi-vendor layer 3 (IP) VPNs can be considered mature by now. Except of carrier-carrier interworking protocols and new extensions to the RFCs (multicast), they are proven to be interoperable regarding functionality and scalability, as shown, for example, during the MPLS World Congress 2004 interoperability event.
- A lot of carriers offer, or are considering the possibility of offering, large-scale multi-point

services using Hierarchical Virtual Public LAN Services (H-VPLS). The hierarchical part of the protocol enables service providers to scale the number of customers and endpoints per customer offered using VPLS, without stressing the backbone network.

Consequently, this test event primarily focused on the scalability of hierarchical VPLS provider edge routers (PE-RS) and multi-tenant units (MTUs). Industry support is growing fast; we tested a total of four PE-RS and eight MTU implementations.

- Furthermore, we verified the scalability of Layer 2 Ethernet pseudowires over MPLS beyond previous limits. Almost every participating device supported them.
- A few vendors tested ATM and TDM pseudowires.
- Finally, we continued to investigate MPLS label switch path (LSP) ping and traceroute interoperability.

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-stage took place at the EANTC (European Advanced Networking Test Center) in Berlin, Germany.

The Interoperability Working Group of the MPLS & Frame Relay Alliance, 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
Ciena	DN 7100
Cisco Systems	12406
IXIA	1600T
MRV	OSM 207 OSM 800

Native Networks	EMX 3700
RAD Data Communications	ACE-3200 ETX-510 IPmux-14 IPmux-11 Gmux-2000
Resolute Networks	Arranto 100
Riverstone	15008 RS8600

Test Areas and Test Plan

The interoperability evaluation focused on Ethernet, ATM and TDM pseudowires (layer 2 point-to-point service), MPLS hierarchical Virtual Private LAN Service (H-VPLS) and enhanced management support for MPLS networks via LSP ping and trace-route facilities.

MPLS Protocol Support	RSVP-TE Signalling	L2 Ethernet Pseudowires	L2 ATM Pseudowires	L2 Frame Relay Pseudowires	L2 TDM Pseudowires	Non-Hierarchical VPLS	Hierarchical VPLS	LSP Ping	LSP Traceroute
Agilent N2X	•	•	•	•		•	•	•	
Alcatel 1662 PRS	•	•					•		
Alcatel 7670 RSP	•	•							
Alcatel 7750 SR	•	•				•	•	•	•
CIENA DN 7100	•	•	•	•	•			•	•
Cisco 12406	•	•	•	•				•	•
IXIA 1600T	•	•	•			•	•		

MPLS Protocol Support	RSVP-TE Signalling	L2 Ethernet Pseudowires	L2 ATM Pseudowires	L2 Frame Relay Pseudowires	L2 TDM Pseudowires	Non-Hierarchical VPLS	Hierarchical VPLS	LSP Ping	LSP Traceroute
MRV OSM	•	•				•	•		
Native EMX3706	•	•					•		
RAD ACE-3200			•						
RAD ETX-510		•					•		
RAD IPmux-14/ Gmux-2000					•				
Riverstone 15008	•	•				•	•	•	•
Riverstone RS8600	•	•				•	•	•	•

The following section describes the test plan in detail. Results are documented on page 5.

MPLS Signalling and Routing

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

Ethernet Point-to-Point VPNs (Pseudowires)

Point-to-point Ethernet VPN Services over MPLS were tested using the MPLS & Frame Relay Alliance 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
- Transfer of Ethernet traffic from many emulated end stations (up to 8,000) and via many tunnels running in parallel inside one targeted LDP session (up to 2,000)
- Circuit emulation over Ethernet pseudowires was also tested.

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-12vpn-vpls-ldp-05*, using the MFA test plan *mpls2003.092.03* (recently extended for H-VPLS scalability).

- 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
- Scalability tests with up to 50 VPLS instances in parallel on each PE-RS and MTU
- Scalability tests with up to 25 PE-RS (most of them emulated) connected full-mesh for all the instances mentioned above
- Scalability tests with up to 4,000 MAC addresses in total

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

LSP Ping and Traceroute

LSP ping was tested according to *draft-ietf-mpls-lsp-ping-07.txt* and MFA test plan *mpls2004.091.00*. Tests involved sending MPLS echo request packets from one label edge router to another, and receiving MPLS echo reply packets from the remote end. This verified connectivity as well as congruence between the data plane and the control plane.

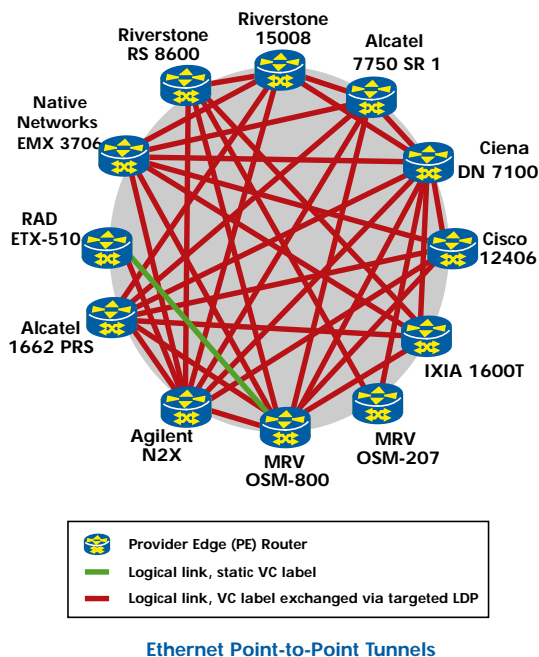
LSP traceroute was verified according to *draft-ietf-mpls-lsp-ping-07*, involving each device as an active initiator and a responder of traceroute commands.

Interoperability Test Results

The goal of this event was two-fold. First, as in most interoperability test events, the test event sought to verify and improve the interworking of vendors' implementations, and second, it sought to prove that service providers may confidently deploy large-scale point-to-point and multipoint Ethernet services over MPLS networks knowing that the network will carry these services with the required levels of resilience and availability.

Today, this means more than just finding bugs and correcting them to advance standards compliance. In many cases, implementations rely on draft standards — vendors need to adapt their features to customers' requirements so quickly that they cannot wait until the final standard is adopted. Thus, the test served secondarily to verify clarity of the current standards.

Results: Ethernet Point-to-Point Pseudowire Tests



Point-to-point Ethernet over MPLS tunnels («pseudowires») were tested according to the IETF PWE3 specifications. During the hot-staging event, all tested point-to-point connections interoperated as expected. Ethernet pseudowires were successfully tested among Agilent N2X, Alcatel 1662 PRS,

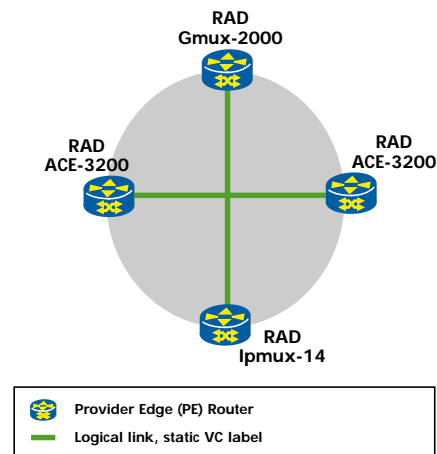
Alcatel 7750 SR 1, Ciena DN 7100, Cisco 12406, Ixia 1600T, MRV OSM-800, Native Networks EMX 3706, RAD ETX-510, Riverstone 15008 and Riverstone RS8600.

Scalability tests were conducted with up to 2,000 pseudowires established within one transport tunnel between two devices. Agilent, Alcatel, Ciena, Cisco, Ixia, RAD and Riverstone participated in this test. Some vendors established fewer pseudowires because the configuration time would have been excessive otherwise (using CLIs because there were no provisioning systems installed during the test).

The large majority of vendors used RSVP-TE signaling for VPN transport labels. Only one participating device supported LDP signaling exclusively.

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

RAD verified functionality of TDM over MPLS pseudowires according to MPLS & Frame Relay Alliance implementation agreement 4.0 «TDM Transport over MPLS using AAL1». The IPmux-14, and Gmux-2000 equipment set up label-switched paths for TDM traffic using static labels without signaling over Cisco and MRV MPLS routers. Ciena claimed support of static labels but did not test them due to time limitations. The other vendors did not support static labels.



Due to limited time, only a few vendors focused the creation of ATM pseudowires. Multi-vendor interoperability in this area has been proven in previous tests, for example SUPERCOMM 2004. The

RAD ACE-3200 devices successfully established an ATM pseudowire tunnel over the backbone using static label assignment.

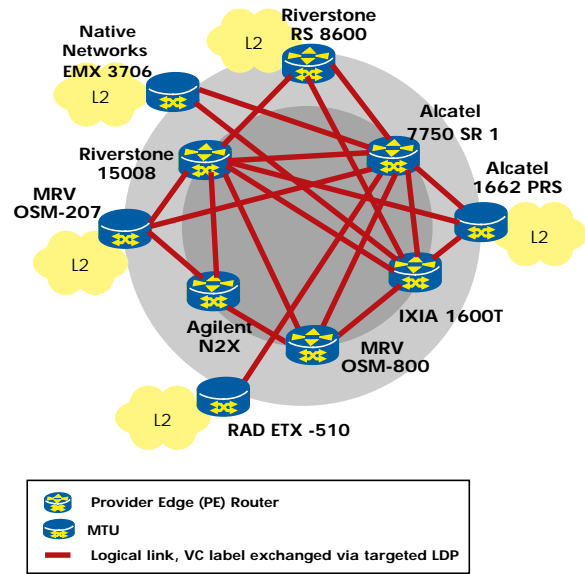
Two vendors also verified transport of E1 over IP over MPLS with their customer premise products (Resolute Networks Arranto 100, RAD IPmux-11). E1 over IP tunnels were based on pseudowires from Native Networks, Alcatel, Riverstone and MRV operating through the MPLS core.

Results: VPLS and H-VPLS Tests

During the hot-staging event, the hierarchical VPLS interoperability between Provider Edge (PE-RS) implementations as well as PE-RS and Multi-Tenant Unit (MTU) systems went very well. As illustrated in the diagram below, all PE-RS implementations were interoperable without any issues: The two PE-RS routers (Alcatel 7750 SR 1, Riverstone 15008), two emulators (Agilent N2X, Ixia 1600T), and a non-hierarchical PE implementation (MRV OSM-800) were able to establish tunnels and to exchange data.

A number of multi-tenant units (Alcatel 1662 PRS, MRV OSM-207, Native Networks EMX 3706, RAD ETX-510, and Riverstone RS-8600) were connected to different PE-RS systems for different VPLS instances in parallel, for example attaching half of the configured instances of the Native Networks EMX 3706 to the Riverstone 15008, and the other half to the Alcatel 7750 SR 1.

MTUs worked with most participating PE-RS systems — all issues were identified, and vendors stated that they will fix the remaining bugs soon.




VPLS & H-VPLS


Furthermore, the test network was loaded with an additional 23 emulated provider edge (PE-RS) routers by the Agilent and Ixia emulators, resulting in a total of 25 PE-RS systems. Each of these nodes had 50 VPLS instances (different customer domains) configured, and there was traffic sent from 4,000 emulated MAC addresses in total.

All systems were interconnected fine on the VPLS layer. A few LDP signaling and OSPF routing/static routing issues created interoperability problems (details see problem section below), but these were the only source of problems.

Results: LSP Ping/Traceroute Tests



LSP Ping and Traceroute



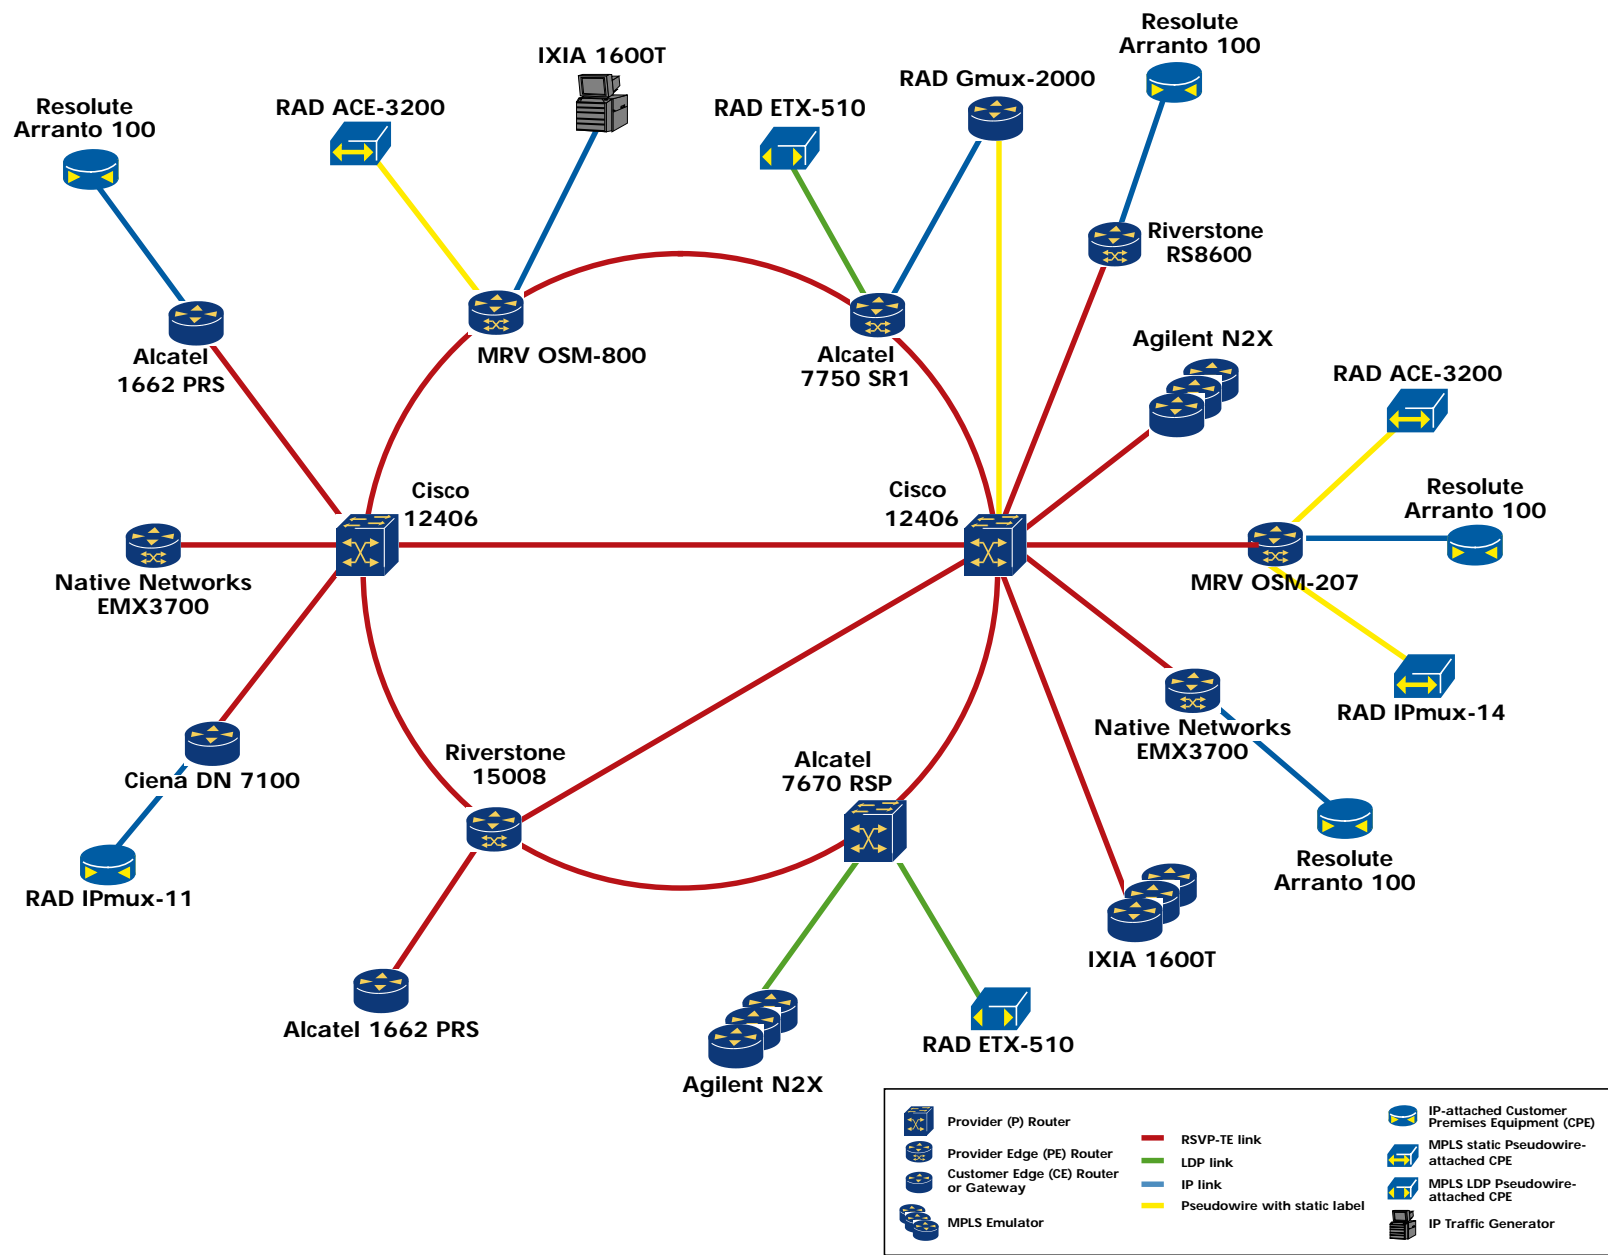
LSP ping and traceroute were tested between four

devices, Alcatel 7750 SR 1, Ciena DN 7100, Cisco 12406, and Riverstone 15008. Unfortunately, only a few combinations worked successfully. The major hindrance seemed to be that vendors implemented different versions of the IETF internet draft — versions 2, 3, 6, and 7 were present. Vendors believed that incompatible changes of the IETF between different versions of the draft prevented interoperability. Using passive

protocol monitors, we verified that LSP ping and traceroute requests or replies were generated properly by all systems, but in some cases they were not detected on receipt. In the end, only the combinations shown in the diagram above worked.

Results Summary

Key Features Tested		Results
L2 Pseudo-wires	Interoperability RSVP-TE	OK
	Interoperability LDP	OK in most cases
	Static LSPs without signaling	OK, supported by two vendors
	Data Transfer	OK
	Ethernet tunnels	OK
	Traffic Transfer Over RSVP-TE and LDP Tunnels	OK
	ATM Pseudowires	OK, tested with one vendor
	Frame Relay Pseudowires	not tested
	TDM Pseudowires	OK, tested with one vendor
VPLS	Full-Mesh Tunnel Establishment between PE-RS systems	OK
	Traffic Transfer Over RSVP-TE Transport Tunnels	OK
	MAC Address Withdraw	not tested
	Hierarchical VPLS PE-RS functionality	OK
	Hierarchical VPLS MTU functionality	OK
LSP Ping and Traceroute	Generating LSP Echo Request	OK
	Generating LSP Echo Reply	OK
	Receiving LSP Echo Requests	Issues
	Receiving LSP Echo Replies	Issues
	Generating LSP Traceroute Request	OK
	Receiving LSP Traceroute Messages	Issues



Final Integrated MPLS Test Network

Problem Summary

Problem Area	Description	Temporary Solution, if any	Recommendation
LDP	Vendor-specific MTU size calculations do not always match	None	If the advertised MTU size is derived from the access interface MTU, the exact value of signalled MTU should be known and configured at both sides.
	LDP Hello message with Hold time set to 0 (default) was not accepted as correct	Set Hold time to 45 sec explicitly	Implementations shall accept the value 0 to match RFC 3036 3.5.2.
	The negotiation of timer values (Hello, Keep Alive) did not work properly which caused LDP sessions to disconnect.	Set timers to identical values	Fix implementations that do not negotiate timers correctly
	Some vendors do not support pseudowires on LDP tunnels	None	To be most flexible for multi-vendor networks, it is recommended to support both RSVP-TE and LDP signaling protocols for tunnel transport
RSVP-TE	Some vendors do not support pseudowires on RSVP-TE tunnels	None	
OSPF	Some vendors do not support OSPF	The Router ID had to be configured statically, TE tunnels had to be established without CSPF.	OSPF-TE support is recommended to improve interoperability and to simplify provisioning.
	Some vendors always expect OSPF-TE to be enabled in order to establish MPLS tunnels, even without Traffic Engineering parameters. It took a lot of time to configure MPLS properly.		Simplify MPLS configuration for the case that the peer does not run a routing protocol.
MPLS OAM	MPLS LSP ping/traceroute packets were sent without an „IP router alert option“. The intermediate router did not recognize this packet as an MPLS OAM packet.	None	This is a common bug with MPLS implementations; vendors should keep it in mind.
	Interoperability couldn't be achieved in some cases because vendors implemented different LSP ping/traceroute draft versions.	Upgrade implementations to the latest draft.	The IETF should verify if incompatible draft updates can be avoided.
Configuration	Traffic Engineering had to be enabled separately in order to establish RSVP-TE tunnels. It took a lot of time to configure MPLS properly.		Simplify the configuration of MPLS.

Conclusion

Since 2002, the MPLS and Frame Relay Alliance has tested and publicly demonstrated different aspects of MPLS interoperability. In a total of six large multi-vendor test events, the participating vendors verified many different MPLS protocols for multi-vendor interoperability — from basic signalling to different flavors of Layer 2 and Layer 3 VPN services as well as DiffServ traffic engineering.

The interoperability event at MPLS World Congress 2005 reassures us that MPLS Layer 2 Ethernet-based VPN implementations are ready for large-scale deployments with many customers in a similar way to Layer 3 VPNs. Of course, each of these two solutions has its particular applications, inherent strengths and weaknesses which are well-known. In both cases, the IETF still continues to develop recommendations (see RFC2547bis, and the battle between BGP-based and LDP-based VPLS).

The more MPLS becomes a complete set of protocols suitable for a multitude of applications, the more we recognize how many work areas still remain to be addressed. As an example, it was quite surprising to notice that there were still issues with the interpretation of the basic LDP standard defined four years ago. At one of the first interoperability events of the MFA in 2003, an urgent need for implementation agreements was noted in order to reduce the amount of protocol options. This statement still holds, and implementation agreements defining use cases are more important than ever.

MPLS LSP ping and traceroute are a different topic. The corresponding internet draft still undergoes regular changes — which is the nature of an internet draft —, and vendors implement different incompatible versions. We contacted the draft authors; they hope that the next draft (number 8) will be fairly mature and might go to last call in March. There is hope that LSP ping and traceroute implementations will work more reliably in heterogeneous environments at the next MFA interoperability test event.

Despite these small issues, Multi-Protocol Label Switching has grown to support a full set of standardized and interoperable VPN types — making MPLS way more flexible than network technologies of the past. A vast number of vendors implement MPLS by now, and the majority of carriers worldwide use MPLS as the foundation for their IP and layer 2 service backbones.

The MPLS & Frame Relay Alliance and the supporting test labs, UNH-IOL and EANTC, are proud that the series of interoperability test events conducted since 2001 have been able to improve interoperability dramatically.

Acknowledgements

The white paper was authored by Jonathan Morin, Michael Pergament and Carsten Rossenhoevel.

References

- MPLS Ethernet and ATM Point-to-Point Pseudo-Wires
 - RFC 3916 — Requirements for Pseudo-Wire Emulation Edge-To-Edge. Xi-Peng Xiao, Danny McPherson, Prayson Pate. September 2004
 - mpls2003.091.03 — MPLS Forum Layer 2 VPN Interoperability Test Suite
 - 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) Expires: March 2005
 - draft-ietf-pwe3-atm-encap-07.txt — Encapsulation Methods for Transport of ATM Cells / Frame over IP and MPLS Networks. Luca Martini et.al. October 2004. (work in progress) Expires: April 2005
 - draft-ietf-pwe-frame-relay-03.txt. Frame Relay over Pseudo-Wires. Luca Martini et.al. August 2004 (work in progress) Expires: January 2005
 - 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
- Virtual Private LAN Services
 - mpls2003.092.03 — MPLS & Frame Relay Alliance H-VPLS Interoperability Test Suite
 - draft-ietf-l2vpn-vpls-ldp-05 — Virtual Private LAN Services over MPLS. Marc Lasserre, Vach Kompella, et.al. September 2004. (work in progress) Expires: February 2005
- LSP Ping
 - draft-ietf-mpls-lsp-ping-07.txt — Detecting MPLS Data Plane Failures
- Generic
 - MPLS & Frame Relay Alliance: Super Demo 2002, Test Plan & Results, June 2002
 - MPLS & Frame Relay Alliance: Resilient & Scalable — MPLS World Congress 2003 Interoperability Demonstration, Test Plan & Results, February 2003
 - MPLS & Frame Relay Alliance: New Revenue Streams With MPLS Service Differentiation — MPLS World Congress 2004 Interoperability Demonstration, Test Plan & Results, February 2004




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



MPLS & Frame Relay Alliance

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

The MPLS & Frame Relay Alliance is an international industry organization that is advancing the recognition and acceptance of MPLS and Frame Relay technologies in the global telecom industry. The Alliance is driving worldwide deployment of multi-vendor MPLS and Frame Relay networks, applications and services, through interoperability initiatives, implementation agreements, and educational and marketing resources and programs. The Alliance currently has more than 50 members.

This report is copyright © 2005 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.