

MPLS
WORLD


CONGRESS '07 CONGRESS '07

Test Report

Multi-Vendor MPLS Interoperability Event

Paris, February 2007

■ **EANTC** ■



www.upperside.fr

Editor's Note

Multi-Protocol Label Switching (MPLS) has "crossed the chasm" between early adoption and significant deployment. Service providers today are using MPLS to deliver point-to-point services over interconnected backbone networks, for example, while placing increasingly complex demands and ever-larger economies of scale upon these networks. Carrier-grade network operators now rely on MPLS to facilitate inter-carrier connectivity, mobile backhaul strategies, and multicast VPNs, to name only three common uses. And yet, misconfigurations and points of needed clarification persist. Clearly not all aspects of interoperability have been tested, and not all ambiguities in the standard have been resolved.

EANTC has followed the development of the MPLS protocol family for its entire duration. Over the course of the last seven years and through many interoperability events, we have seen a lot of MPLS services mature, perform and interoperate flawlessly. A few examples include signaling, routing, traffic engineering and VPN services. Buzzwords come and go, but MPLS has proved itself the most robust, interoperable and scalable carrier-class technology available for packet networks today.

This year's MPLS World Congress interoperability event focused on multiple areas of recent MPLS advances. The results further demonstrated that multi-vendor MPLS networks actually do work. Both returning vendors and first-time participating companies brought mature implementations to the test. There were a few issues as always (see the problem section for details), but there were no serious showstoppers.

To summarize our findings, inter-carrier connectivity functioned well between all participating vendors — at least one of the three options works for any combination. Basic multicast VPN (»Rosen draft«) implementations interworked seamlessly in most cases (although not across inter-area boundaries, there is no standard yet). End-to-end interoperability was verified using new mobile backhaul (SAToP) and encrypted pseudowire services.

The hot-stage event at EANTC resulted in a substantial amount of technical data, summarized in this white paper. We hope you'll find it useful.

Carsten Rossenhoewel

Introduction

EANTC's interoperability tests share two targets. The first target is device and equipment manufacturers. Our events allow network device and equipment vendors a rare opportunity to test implementations against those of other vendors and to verify implementations of new technologies in accordance with standards. Vendors gather information about the nature of any problems and in many cases receive new code updates from developers during the event. Interoperable vendors use the results to expand their footprints in new markets and display their level of technological readiness.

Our second target is network operators. The EANTC interoperability events provide carriers and service providers with a realistic picture of available technical solutions, potential deployment issues, and best practices network design. Carriers' involvement in the interoperability events has grown steadily in recent years.

A resulting effect of the interoperability event is feedback to technical committees and standards bodies, including opportunities to improve the standards' readiness for real world deployment.

EANTC and UNH-IOL (University of New Hampshire Interoperability Lab) developed the test plan between October and December 2006. The test areas were finalized after extensive review with participating vendors and service providers.

Based on EANTC's experience in organizing and executing multi-vendor interoperability events an eight days, closed doors, hot-staging event was conducted in EANTC's lab in Berlin, Germany. The results are presented in this White Paper.



Figure 1: Hot-staging at EANTC, Berlin

Document Structure

Participants	→ Page 3
Network Design	→ Page 3
Test Areas	→ Page 4
Interoperability Results	→ Page 5
Topology Diagram	→ Page 10
Problem Summary	→ Page 11
References	→ Page 11

Participants and Devices

Alcatel-Lucent	7710 SRc-12 7250 SAS 7750 SR1 7750 SR7 5620 SAM 1850 TSS-40
Cisco Systems	CRS-1 12000 Series
Foundry Networks	Netiron XMR 16000
Huawei Technologies	CX600 ME60 NE40 NE40E
IXIA	Optixia X16
MRV Communications	OS9024M-210Gx OSM207
RAD Data Communications	ACE 3100 ETX-202 Gmux-2000
Redback Networks	SmartEdge 400
Rohde & Schwarz SIT	SITLink R&S SITLine ATM
Siemens	SURPASS HiD 6650
Spirent	TestCenter
Telco Systems, a BATM Company	T-Metro-100 T-Metro-200 T-Marc-250 T-Marc-254
ZTE Corporation	ZXR10 T128

Carrier Involvement

Engineers from the German service providers T-Systems and Versatel provided detailed comments to the test plan and participated for the duration of the whole hot-staging event. T-Systems engineers were responsible for carrying out the service verification tests; Versatel coordinated the inter-area tests.

Network Design

The interoperability test infrastructure resembled two next-generation carrier networks designed to support end-to-end residential IPTV, Triple Play, business services and cellular backhaul.

Given a total of 13 participating vendors, we created a multi-vendor network with many interoperability test opportunities. This amount of vendors would be quite unusual in a real service provider network; carriers often use heterogeneous MPLS networks from a few vendors — typically two vendors for the core, aggregation and CPEs.

Figure 2 illustrates the five logical parts of the network that were the focus of the tests:

- 1. Inter-carrier domain.** Extending existing VPN technologies to cross the traditional carrier domain can provide more service offerings to the customer, and institute new business models into the service provider's portfolio. Several inter-AS VPNs were demonstrated, supporting both Layer 2 and Layer 3 applications across the boundary. L2VPNs included both single-segment and multi-segment pseudowires end-to-end, as well as a full mesh VPLS domain. L3VPNs featured the three options (known separately as Option A, Option B, and Option C) described in IETF RFC 4364.
- 2. Service provider domains.** Within each of the two carrier networks, VPLS and MPLS/BGP VPN were created among the participants. In addition, Fast Reroute provided each backbone with carrier grade Resilience and Protection.
- 3. Aggregation.** Several distinct solutions to implement aggregation and access networks were verified during the test. In particular, H-VPLS MTUs, and TDM and ATM pseudowire access switches were evaluated. The Resilience and Protection for the aggregation part of the network was covered by Dual-Homed MTU and MPLS BFD.
- 4. Access.** In this area encrypted end-to-end ATM and TDM connections as well as end-to-end Ethernet path redundancy were tested.
- 5. Network Services Area.** Triple play applications such as VoIP and IPTV (IP-Multicast based) servers were positioned in this area.

Test Areas and Test Plan

The variety of test cases in our interoperability events continues to grow annually as MPLS becomes applicable to more types of products and network environments. The following test areas have been defined in cooperation with the participants who attended the hotstaging event:

- **MPLS based services.** The basic premise of the test network relied on an underlying MPLS infrastructure. Participants were only allowed to partake in testing if their device supported at least one of the aforementioned MPLS services. Specific interoperability problems inherent to MPLS have been addressed in previous events. Hence, in this year's event, the MPLS services were simply used to facilitate testing of advanced functionalities.
- **Resilience and Protection.** To date, disruption prevention and downtime minimization have been a top priority for service providers in every generation of telecom networks. As networks become departmentalized, failure restoration capabilities must also be installed towards the edge (e.g. aggregation), in addition to the service provider's backbone where sub-50ms restoration traditionally has been. In this event, aggregation and access protection were verified via Dual Homed MTUs and VPLS based Ethernet Path Protection.




Figure 2: Schematic Network Design

- **Fault Monitoring.** Persistent fault monitoring is essential to the delivery of high quality services to both residential and business customers. During the hotstaging event, various fault monitoring tools were tested among the participants, including MPLS ping, MPLS traceroute and MPLS BFD.
- **Inter-carrier L3VPN.** Several solutions exist to enable an L3VPN to cross different administrative domains. In particular, the IETF defined three options where MPLS/BGP VPNs can be supported across an AS boundary using readily available protocols. To the customers, implementation of the three options are transparent. To the service providers, each option requires a different level of trust relationship at the boundary. Option A is the simplest of the three, as each service provider views the other as a customer. Option B overcomes some drawbacks of Option A as a result of its simplicity, but a higher level of trust must exist between the two providers. Option C is the most integrated method, in which the two provider networks must function as one, forming one single large VPN domain end-to-end. Each option has its pros and cons, as well as its value proposition. All three options were tested.
- **Inter-carrier L2VPN.** Similar to inter-carrier L3VPN, a number of options to extend L2VPNs over AS borders exist. The options available to L2VPN use the IETF defined multi-segment pseudowire (MS-PW). MS-PW allows for the termination of a pseudowire in one AS, and continuation of the same service using another pseudowire in the next AS. The available options differ in the way the pseudowires are stitched together. Popular procedures include the use of VLANs, manual stitching of pseudowires, LSP tunnels stitching and dynamic extensions of pseudowires. All four options were tested.
- **ATM and TDM Support.** ATM and TDM application support is still an important issue in next-generation packet network. These applications are time tested and provide a steady source of revenue for many service provide. TDM and encrypted ATM traffic was carried over Ethernet Pseudowires in addition to a demonstration from ATM circuit encryption devices.

Interoperability Test Results

After eight intense days of hot-stage testing we collected a substantial amount of results. The results provide a detailed insight into the current state of MPLS solutions. The following sections summarize our findings.

We evaluated VPLS/H-VPLS and L3VPN services in the MPLS backbone. VPLS was based on LDP signalling using FEC 128 while the L3VPN services were based on RFC 4364 (2547bis). Additionally, the ability to provide a point-to-point service for TDM and ATM native traffic was tested.

As a first step, two MPLS backbone networks were configured resembling a typical service provider infrastructures. Both backbones included access and CPE components. They were interconnected via redundant connections as shown in Figure 2.

Results: Inter-Carrier Connectivity

The inter-connection between the two mock carrier networks was tested using all variations defined in RFC 4364. This specification provides three options for inter-AS connectivity for different network scenarios (see the Test Plan section for a detailed discussion.)

Autonomous System Border Routers (ASBRs) were connected in as many combinations as possible — many more than are shown in the final network diagram. Four VPN service types were verified across the inter-area links: IP-based VPNs, Ethernet-based multipoint VPNs (VPLS), single-segment pseudowires and multi-segment («stitched») pseudowires.

Signaling and Routing. The original test plan had called for RSVP-TE signaling to create transport tunnels. Instead, we decided to use LDP at the hot-staging event. LDP (in downstream unsolicited mode) simplifies the configuration of the border routers, because the signaling protocol automatically creates a label for each route — RSVP-TE would have required manual tunnel setup. This decision was taken purely for convenience; RSVP-TE would have been a working option for all participating vendors.

For the exchange of pseudowire VPN labels, targeted LDP sessions were established between the provider edge routers as usual (required by the specification).

Inside the MPLS domain of each autonomous system, OSPFv2 was used as the interior gateway protocol (IGP). Exterior BGP (eBGP) was configured on the inter-area links. In addition, vendors used the interior BGP (iBGP) protocol to exchange VPN labels for the IP-based VPNs as specified in the standard (RFC 4364).

iBGP was *not* used to exchange transport labels inside the autonomous systems.

Inter-Area Connections were configured on the border routers Alcatel-Lucent 7750 SR7, Cisco 12000, Cisco CRS-1, Foundry Netiron XMR 16000, Huawei NE40E, Huawei CX600, Redback SE400 and ZTE ZXR T128. See figure 3 for a diagram of all successfully evaluated combinations. Due to limited time, not all possible connections were verified; for the same reason, routers were not relocated to the other autonomous system (AS) to allow for more combinations.

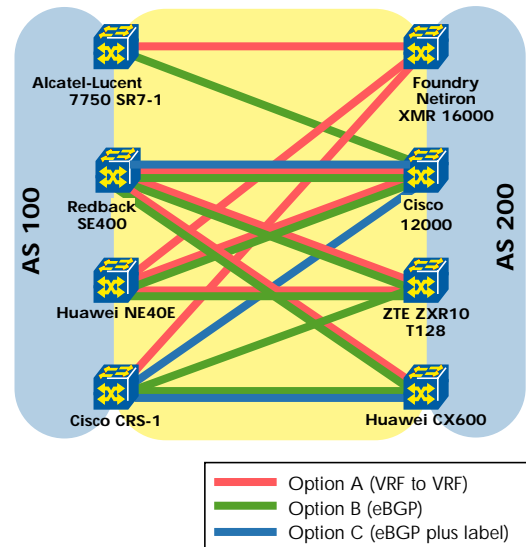


Figure 3: Inter-Area L3VPN Connections

The **Option A** links (marked in red) were quickly established and did not show any issues between participating vendors. Option A is similar to normal PE-CE connections. The inter-area link is simply a routed unicast IPv4 link where multiple VRFs are differentiated on layer 2, in our case by VLAN IDs.

The major part of the testing focused on **Option B** which was supported by Alcatel-Lucent 7750 SR7, Cisco 12000, Cisco CRS-1, Huawei NE40E, Huawei CX600, Redback SE400 and ZTE ZXR10 T128. Eight connections in total were tested successfully.

During the Option B tests, we encountered eBGP configuration issues which caused loops when multiple Option B links were active simultaneously for resilience. In some cases, the redistribution of eBGP inter-area routes into an interior gateway protocol (OSPF) in both autonomous systems led to loops. These were purely configuration issues which were resolved during the test by proper filtering.

Finally, several vendors worked on **Option C** links. Option C was supported by Cisco 12000, Cisco CRS-1, Huawei CX600 and Redback SE400.

No multi-vendor interoperability was achieved due to a mapping issue: Once the border routers have exchanged labeled eBGP routes on the inter-area link, the labels needed to be mapped inside each of the attached networks.

End-to-end label-switched paths can only be created if labeled paths towards the inter-AS links exist at the provider edge routers. Some implementations supported mapping the eBGP routes to iBGP only — not to LDP. While these implementations did implement the inter-area link properly, we were unable to verify label exchange inside the area (label switched paths between the Option C link and the provider edge routers) because the implementations did not support it.

To our knowledge, mapping procedures are not specified by the IETF, because they are internal to a router. It would be good to provide some guidance to vendors, for example by means of an application note.

As a backup (because Option C links were required for other test areas), a single-vendor Option C connection was configured.

BGP/MPLS VPNs. Based on the inter-area links created above, IP VPN services were established between the provider edge routers (PEs) Alcatel-Lucent 7710 and 7750 routers, Cisco 12000 and CRS-1, Foundry Netiron XMR 16000, Huawei CX600 / NE40 / NE40E / ME60, Redback SE400 and ZTE ZXR10 T128. Notice that the border routers were also configured as PEs in the hot-stage lab environment.

No route reflectors were used in the test. Inside each carrier network, full-mesh peers were established to increase the number of test combinations.

The inter-carrier connections were established over two Option B links in parallel. No load-sharing methods were employed; failover was enabled by use of the regular BGP routing update mechanisms.

For this and all further test areas, one physical port at each router was configured for a PE-CE link and connected to the Ixia X16 and Spirent TestCenter emulators. Using these test links for traffic generation and analysis, we verified full-mesh connectivity between all provider edge routers for the MPLS/BGP VPN service.

IP VPN intra-area and inter-area connectivity was achieved successfully between all participants without any issues:

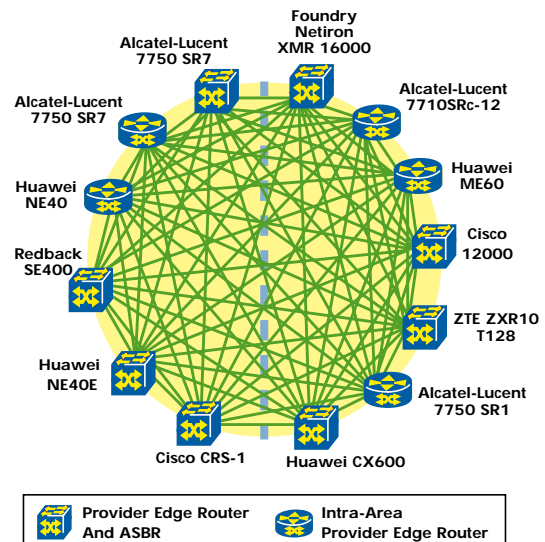


Figure 4: MPLS/BGP VPN Logical Diagram

VPLS Multipoint Ethernet Service. Virtual Private LAN Services (VPLS) were established between provider edge routers inside and across the two carrier networks.

All MPLS-enabled devices participated in this test.

In the hierarchical VPLS domain, the Alcatel-Lucent 1850-TSS40, Alcatel-Lucent 7250SAS, MRV OS9024M-210Gx and Telco Systems T-Metro were configured as multi-tenant units (MTUs) with one or two redundant uplinks to a PE-RS router. All other systems were configured as PE-RS routers, connecting to each other in a near full-mesh logical configuration as shown in Figure 5.

We configured one common VPLS domain on all participating devices so that the Ixia X16 and Spirent TestCenter systems were able to verify full-mesh connectivity.

As far as vendors could spare the time, full-mesh combinations were configured and successfully verified to forward end-to-end traffic. Only two interoperability issues occurred of which one could be resolved during the event.

Quite a few configuration issues delayed the creation of the VPLS domain. We have seen similar delays in previous EANTC multi-vendor interoperability tests:

- Choice of pseudowire type. Both the »raw Ethernet« and the »VLAN« pseudowire types can be chosen for VPLS links — independent of whether the Ethernet frames are to be delivered on VLANs at the edge or not. Some vendors were able to configure both alternatives; some could even choose different encapsulation types for different links within the

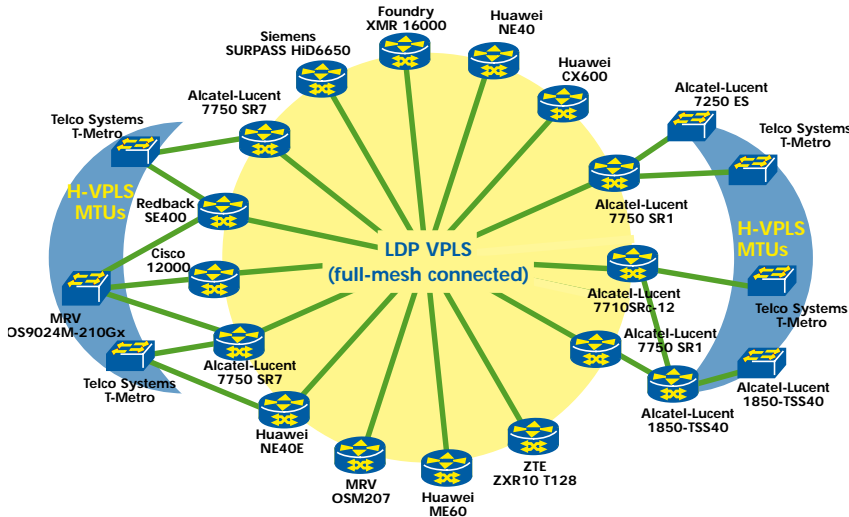


Figure 5: VPLS Connections

same VPLS domain; two vendors supported only the »raw Ethernet« option.

- Maximum Transmission Unit (MTU) is a common source of misconfiguration within Ethernet-based MPLS networks. The MTU has to be increased on MPLS transport links so that labeled Ethernet frames with 1500 bytes service payload can be carried.
- LDP interoperability was delayed due to address configuration issues and label ranges. Newcomers were still confused by the question whether interface or loopback addresses should be used for LDP sessions. In one case, the large label range sent by a peer router confused the implementation of the receiver.

These issues should be taken seriously. Experience shows that they do not disappear over time. Why do multiple options of limited use still exist in MPLS standards? Two pseudowire types for the same purpose and MTU sizes that are too small to work in any reasonable MPLS configuration are potential sources of misconfiguration. We hope the IETF will consider clarifying their use, reducing the amount of options and avoid similar situations in future specifications (similar trouble is ahead in the Ethernet multicast over MPLS area).

End-to-End Pseudowires. In addition to multipoint services, point-to-point tunnels were created as in a real MPLS service provider network.

We did not aim to create a fully meshed pseudowire end-to-end scenario because this task was completed as part of the VPLS tests, so we were certain all participants would interoperate.

Instead we focused on inter-area pseudowires as well as interoperability issues. There are two options for inter-area pseudowires to be established. From the point of view of the provider edge router, creating end-to-end LSPs is easiest, where the border router maps labels appropriately over the inter-carrier links. An end-to-end pseudowire can be established just like an intra-area pseudowire. No special configuration or protocol support is required at the provider edge router.

TDM pseudowires were established dynamically (e.g. using LDP) with the following encapsulation options:

- SAToP (RFC4553)
- TDM over Ethernet (MEF 8)

In addition, end-to-end ATM pseudowires (according to RFC 4717) were established between RAD ACE-3100 and Alcatel-Lucent 7750.

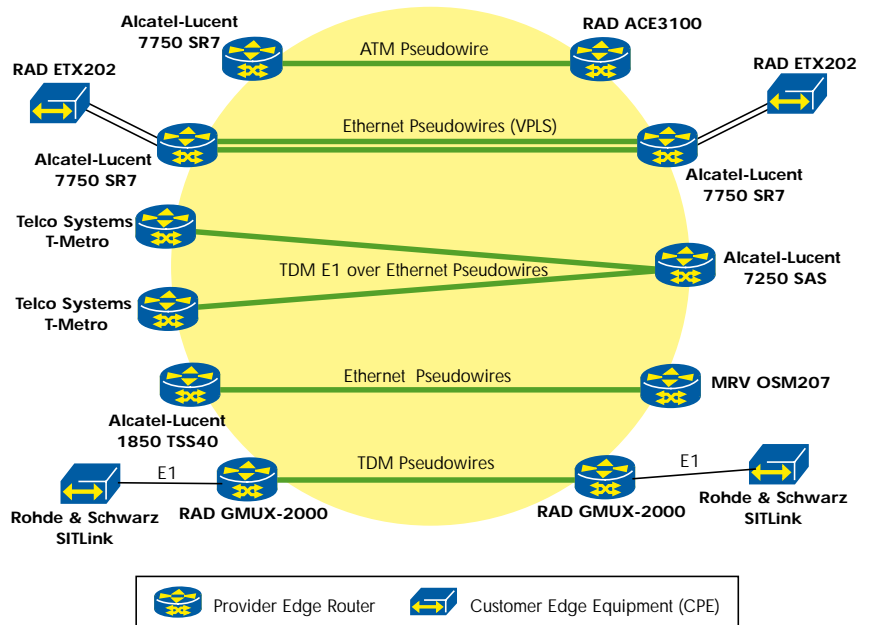


Figure 6: End-to-End Pseudowires

Please see the application section for more details.

All pseudowires were established without any issues after the inter-area connections were established.

Multi-Segment Pseudowires. Multi-segment Pseudowires (MS-PW) are an important solution to enable inter-area point-to-point layer 2 services. In the test event, MS-PWs were used solely for Ethernet Pseudowires.

Multi-segment Pseudowires are created of multiple pseudowires that are »stitched« together at intermediate MPLS routers. Each segment can be statically defined by using LDP on each segment. Pseudowire stitching was performed by Cisco 12000, Huawei NE40E and ZTE ZXR10 T128.

In addition, VLAN stitching of pseudowires was tested between the ZTE ZXR10 T128 and Huawei NE40E/NE40/CX600. In this case label-switched paths from one carrier network were terminated by a border router and reconnected to label-switched paths in the second carrier network by using VLAN tags. This is a static solution that is simple, but requires manual configuration.

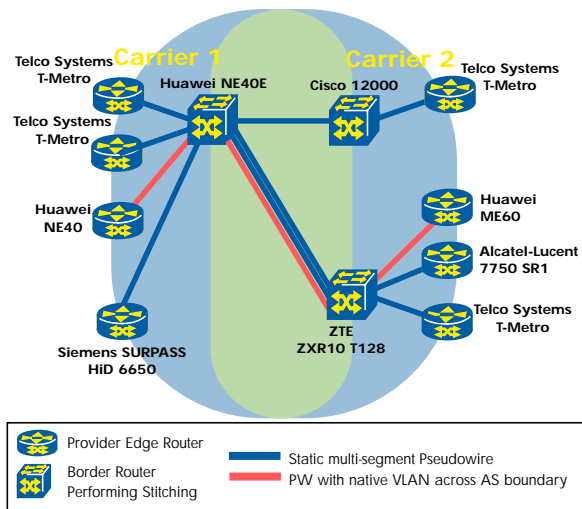


Figure 7: Multi-Segment Pseudowires

Results: MPLS Protection

The following vendors were scheduled to participate in **Fast Reroute** testing: Alcatel-Lucent 7710 and 7750 routers, Cisco 12000 and CRS-1, Foundry Netiron XMR 16000, Huawei CX600/NE40/NE40E/ME60, Redback SE400, Siemens SURPASS HiD 6650, Telco Systems T-Metro and ZTE ZXR10 T128.

Given the limited time, we prioritized the inter-carrier related test areas because they were tested for the first time. Fast Rerouting has been evaluated for multi-vendor interoperability in several EANTC events in the past.

Fast Rerouting could not be evaluated across the inter-carrier connections. There is an IETF draft describing a potential solution, but it is in its early stages. Parti-

pants were not yet ready for multi-vendor interop testing.

There was only one case of Fast Rerouting tests between Telco Systems T-Metro and Alcatel-Lucent 7750 SR7 and 7710. The rerouting was verified to work in under 50 ms.

Dual Homing MTUs. Dual homing MTUs to two PEs provides resiliency for the MTU-PE connection and hence significantly increases the service availability to the customer. In order to verify Dual Homing interoperability three MTUs were connected to two upstream PEs.

One Telco Systems T-Metro MTU was connected to the Redback SE400 and the Alcatel-Lucent 7750 SR7 utilizing local link failure detection in order to switch from the primary link to the secondary when the primary link was disconnected. A second link Telco Systems T-Metro MTU was connected to Huawei NE40E and Alcatel-Lucent 7750 SR7. On the primary connection BFD was configured to detect OSPF adjacency loss and automatically switch to the secondary connection. The MRV OS9024M was the third MTU used in this test. It was connected to the Redback SE400 and the Alcatel 7750 SR7. This combination also showed no interoperability problems in using a secondary connection when the primary was down.

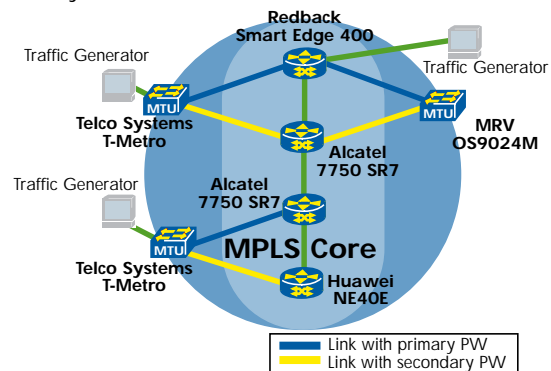


Figure 8: Dual Homing MTU Tests

Results: Fault Monitoring and OAM

LSP ping tests were carried out in the MPLS backbone with a total of 15 implementations. We did not invest much time in configuring and troubleshooting the LSP ping protocol; it was expected that it should just work, being a troubleshooting protocol itself.

We discovered a total of 10 inconclusive results; this calculates to a 57 % success rate for the LSP ping fault detection. It seems LSP ping is still not a commodity service that works reliably in multi-vendor environments.

Furthermore, MPLS traceroute was tested in several combinations. Most problems were related to the imple-

mentation of different protocol versions. Some vendors implemented draft versions while others already supported the final RFC 4379 — which is incompatible with its predecessor drafts. In general, interop problems were more visible in MPLS traceroute than in LSP ping since the former uses two additional TLVs (“Downstream-Mapping” and “Interface and Label Stack”) that were changed during the evolution of the drafts.

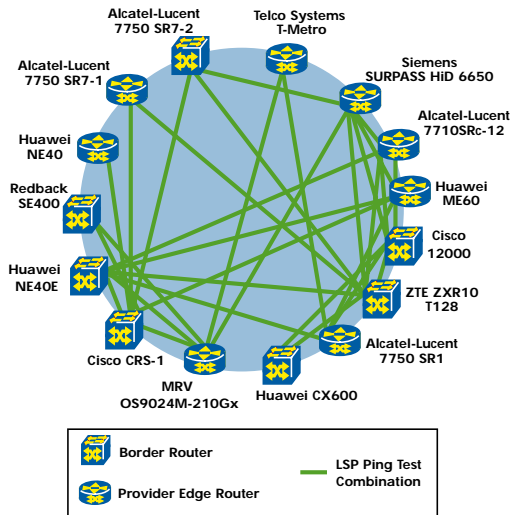


Figure 9: LSP Ping Tests

Results: End-to-end Ethernet Path Protection

Two RAD ETX-202 Ethernet demarcation devices, located at both ends of the Ethernet path across the VPLS core, exchanged Ethernet loopback OAM to control end-to-end path integrity. In case of a path break in the core, the ETX-202 units stopped receiving OAM, indicating a failure, and automatically switched the user traffic (emulated by IXIA) to a backup VLAN. The backup VLAN was recognized by PE (Alcatel 7750) that re-established the backbone Ethernet pseudowire connection by switching over to the backup VPLS path.

Results: Multicast VPN Service

Multicast in Layer 3 VPNs was tested in accordance to the L3VPN IETF working group Internet draft known as the »Rosen draft«.

In each of the two carrier networks a separate multicast topology was defined. In the respective network backbones, PIM-SM (protocol independent multicast) was used to create multicast distribution trees. All participating provider edge routers joined a single group. Generic routing encapsulation (GRE) tunnels were used in accordance with the »Rosen draft« specification to transport multicast traffic within the network.

We created a customer VPN dedicated to multicast inside each of the two carrier backbones. Spirent's Test-Center was used to emulate customer edge routers against every participating provider edge router. In order to facilitate multicast traffic distribution in the customer domain, PIM-SM was configured on all links towards customer edge routers (in addition to OSPF routing). A multicast group was sourced behind one emulated customer edge router and joined by receivers attached to all other endpoints within the multicast VRF.

In one of the carrier networks, Alcatel 7750 SR7, Redback SE400 and Huawei NE40E constructed a multicast domain within L3VPN. All devices were able to construct data distribution trees and deliver multicast to receivers.

In the second autonomous system, Huawei ME60, Cisco 12000, Alcatel 7710SRc-12 and Alcatel 7750 SR1 constructed the multi-vendor multicast topology. After having fixed configuration issues all vendors showed no problems interoperating with each other.

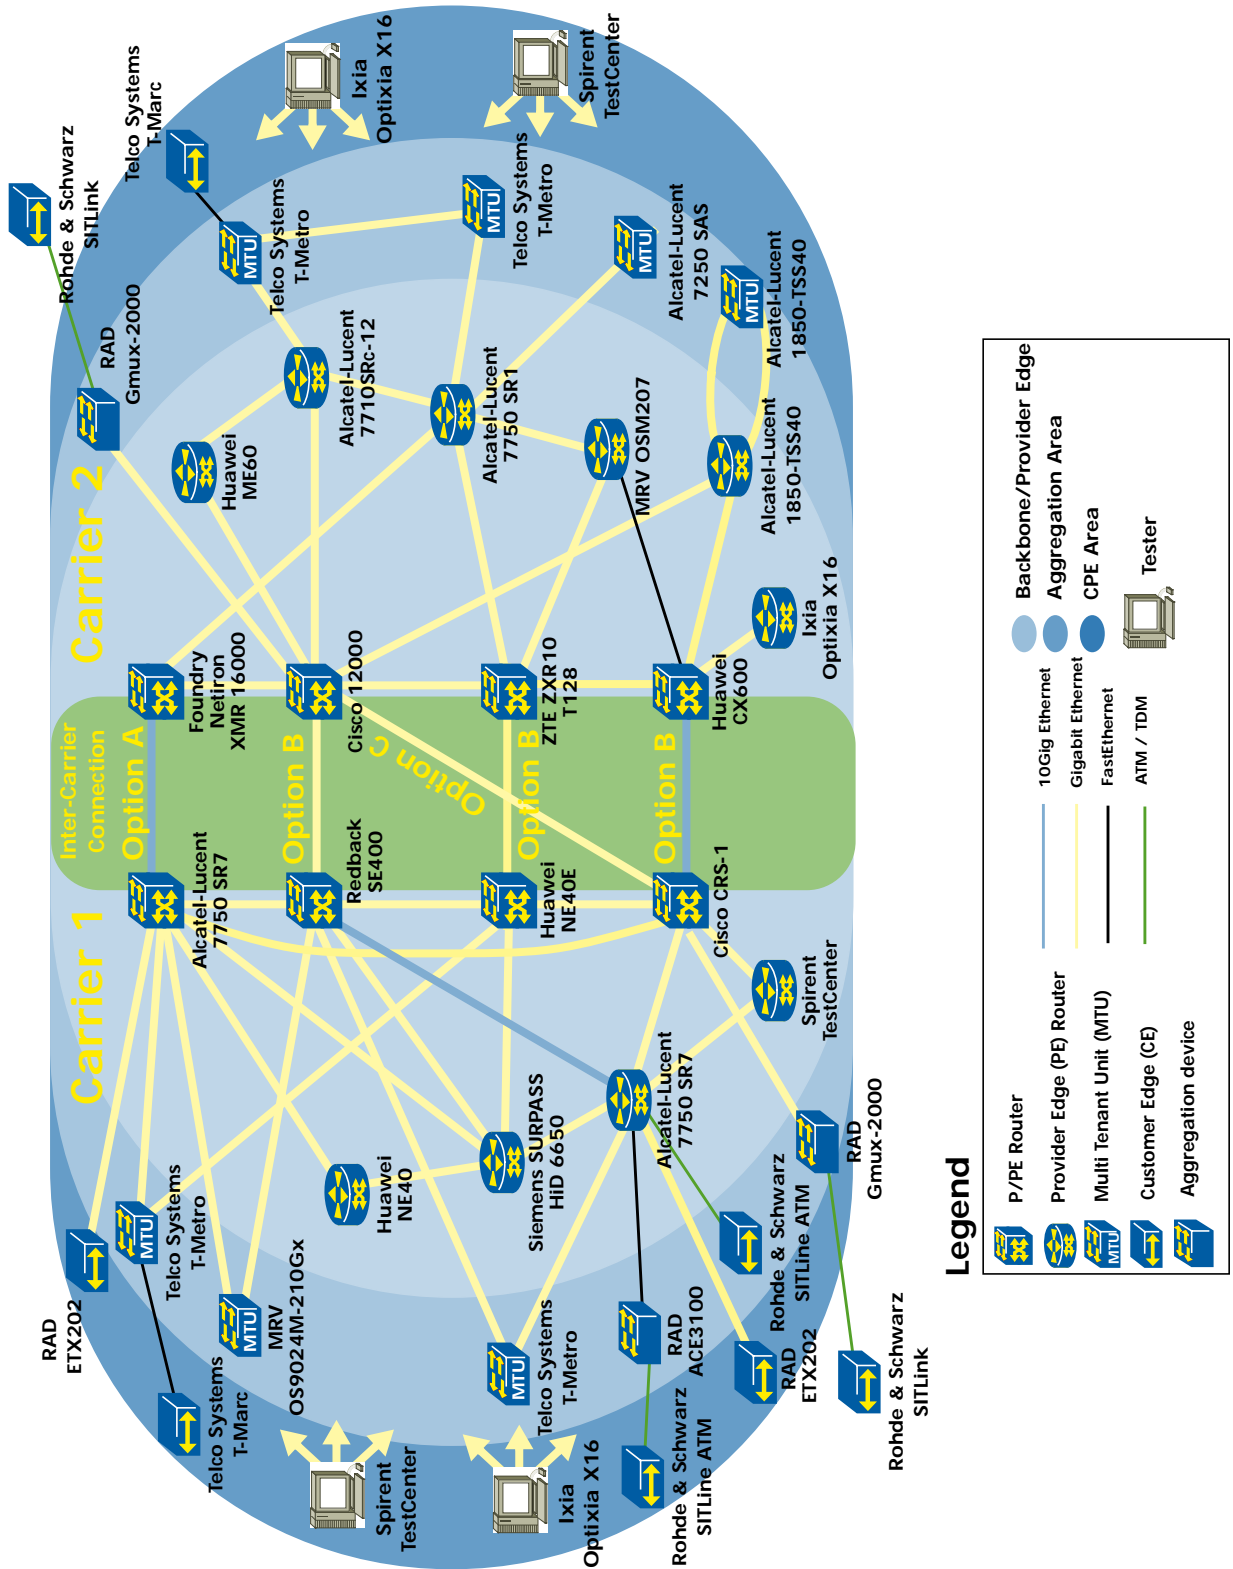
Participating vendors discussed the potential of interconnecting the two multicast domains across the inter-area connections. There is no specific standard yet. The only option would have been to interconnect the domains using *Option A* (see above), but for this purpose multiple virtual multicast routing instances would have been required on each of the border routers — one instance per customer domain. Participants were not ready for multi-vendor interoperability testing in this area yet.

Results: TDM and ATM Pseudowires

RAD Gmux-2000 gateways were used to dynamically establish single-segment TDM pseudowires across the backbone using LDP signaling. The encrypted E1 TDM traffic was encapsulated according to RFC4553 (SAToP). To demonstrate the transport of the TDM traffic, two applications were run over E1 TDM interfaces: voice and transparent E1 TDM traffic from the Rohde & Schwarz SITLink encryption system.

A RAD ACE-3100 access gateway and Alcatel-Lucent 7750 PE were used to establish single-segment ATM pseudowires across the backbone. The ATM traffic was encapsulated according to RFC 4717, N-to-one mode without control word. To demonstrate the transport of the encrypted ATM traffic, Rohde & Schwarz R&S SITLine ATM encryption systems were connected via STM-1 ATM UNI interfaces.

Final Integrated Test Network



Problem Summary

Problem Area	Description	Temporary Solution, if any	Recommendation
L3VPN Inter-Area with Option C	No label allocation via LDP for routes that were imported from eBGP into OSPF	Use a different Inter-area option for L3VPN	None
	No label allocation via eBGP if eBGP+label was enabled		None
LDP Session Establishment	LDP Hello messages were sent to a unicast address sourced on the interface address	The LDP session was established through an intermediate hop	Correct the implementation to transmit LDP hellos to multicast address with the correct source IP
MPLS Traceroute	Traceroute stops at intermediate hops	None	Implementations must be updated in accordance to the final RFC
Static LSP Establishment	LSPs could not be established due to label range incompatibilities	None	The maximum label range values that routers must accept in incoming messages should be specified in the RFC
Multi-Segment Pseudowires	Some vendors implement the stitching based on the LDP label, some on the VCID label	None	Increase flexibility of implementations or specify more clearly in the RFC which label should be used for stitching

Acknowledgments

We would like to thank Ralf-Peter Braun, Manuel Paul, and Sabine Szuppa from T-Systems and Reiner Rommel from Versatel for the extensive support during the hot-staging event, guidance during test plan development and network design. This document has been edited by Carsten Rossenhoevel, Jambi Ganbar, Sergej Kaelberer (EANTC); Henry He, Kyle Price, Chris Volpe (UNH-IOL).

References

- LDP Specification (RFC 3036)
- Fast Reroute Extensions to RSVP-TE for LSP Tunnels (RFC 4090)
- OSPF Version 2 (RFC 2328)
- Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP) (RFC 4447)
- Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures (RFC 4379)
- BFD For MPLS LSPs (draft-ietf-bfd-mpls-03.txt)
- Pseudo Wire Virtual Circuit Connectivity Verification (VCCV) (draft-ietf-pwe3-vccv-12.txt)
- Bidirectional Forwarding Detection (BFD) (draft-ietf-bfd-base-05.txt)
- BGP/MPLS IP Virtual Private Networks (VPNs) (RFC 4364)
- Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling (RFC 4762)
- Encapsulation Methods for Transport of Ethernet Frames Over IP/MPLS Networks (RFC 4448)
- Segmented Pseudo Wire (draft-ietf-pwe3-segmented-pw-03.txt)
- Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised) (RFC 4601)
- Multicast in MPLS/BGP IP VPNs, Work in progress (draft-ietf-l3vpn-2547bis-mcast-03.txt)
- Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP) (RFC 4553)
- Encapsulation Methods for Transport of Asynchronous Transfer Mode (ATM) over MPLS Networks (RFC 4717)
- Implementation Agreement for the Emulation of PDH circuits over Metro Ethernet Networks (MEF8)




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
qhe@iol.unh.edu (Henry He)
www.iol.unh.edu



www.uppertime.fr

Upper Side

Tel: +33 1 53 46 63 80
Fax: +33 1 53 46 63 85
www.uppertime.fr

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. (v1.1)