

10, 11, 12 & 13 February 2004
Sofitel Forum Paris Rive Gauche - France

MPLS WORLD CONGRESS 2004

New Revenue Streams With MPLS Service Differentiation

PUBLIC INTEROPERABILITY EVENT
TEST PLAN AND RESULTS



Hosted by



www.upperside.fr

Introduction

This interoperability event has been organized by the MPLS & Frame Relay Alliance and the European Advanced Networking Test Center (EANTC), and hosted by Upperside.

The interoperability test scenarios were performed using a multi-vendor network of real MPLS routers, complemented by emulators. Test methodologies were checked and improved throughout the testing. The end result was that the network was successfully constructed. This achievement, along with the advantages and capabilities of this technology, will be demonstrated at the MPLS World Congress in Paris, February 11–13, 2004. The demo was hot-staged at EANTC labs in Berlin, Germany, in January 2004.



**Hotstaging at EANTC lab
(Berlin, Germany)**

The test scenarios were designed specifically for this showcase, based upon the experiences of previous interoperability test events. They covered new MPLS capabilities which have not been shown before. The focus was on demonstrating differentiated services over MPLS and MPLS traffic engineering. Multi-vendor MPLS/BGP VPNs and Layer 2 Ethernet VPNs (Martini and VPLS) were configured to prove that services were traffic engineering-enabled and could process differentiated services.

To ensure the success of the event, a one week hot-staging event with all the participating vendors was conducted before the MPLS World Congress. It took place at the EANTC (European Advanced Networking Test Center) in Berlin, Germany. The test plan was defined by the Interoperability Working Group of the MPLS & Frame Relay Alliance, including EANTC and UNH IOL (University of New Hampshire Interoperability Lab).

Participants and Devices

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

Agilent Technologies	RouterTester
Alcatel	7670 RSP
Avici Systems	QSR
Cisco Systems	12404 12406
Ixia	1600T
Marconi	BXR 48000
Alcatel/Native Networks	ISA PR & 1662SMC
Navtel Communications	InterWatch 95000
Nortel Networks	Shasta 5000 BSN
RAD Data Communications	IPmux-1 IPmux-1000
Riverstone Networks	RS8000
Tellabs	8820

Test Areas and Test Plan

This time, the interoperability evaluation focused the transport of differentiated services a multi-vendor MPLS backbone supporting MPLS traffic engineering.

Based on the guaranteed backbone transport, multi-vendor MPLS virtual private networks (VPNs) were established to demonstrate that VPNs can benefit from the enhanced backbone services.

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

	MPLS DiffServ	Constraint-Based Routing	VPLS	L2 VPNs	BGP/MPLS VPNs	TDM over MPLS
Agilent Tech.	●		●	●	●	
Alcatel	●	●		●	●	
Avici Systems	●	●				
Cisco Systems	●	●			●	
Ixia			●	●	●	
Marconi	●	●			●	
Alcatel/ Native Networks			●	●		
Navtel Comm.	●	●		●	●	
Nortel Networks	●				●	
RAD						●
Riverstone			●	●	●	
Tellabs		●	●	●	●	

MPLS Traffic Engineering And Differentiated Services

The MPLS & Frame Relay Alliance test plan *mpls2003.149.03* was used for this section.

The tests asked to enable OSPF-TE routing and to ensure full-mesh exchange of dynamic routes in the MPLS backbone. Based on a network-wide agreed per-hop behavior configured on each device, multi-vendor constraint-based MPLS tunnels should be established using the signalling protocol RSVP-TE and the relevant extensions for MPLS Diff-Serv.

Next, transport of emulated application traffic in the Diff-Serv classes Assured Forwarding (AF), Expedited Forwarding (EF) and Best Effort classes should be configured at the label edge routers. A mapping between per-hop behaviors and experimental bit settings (PHB-EXP mapping) was required; the correct prioritization of the per-hop behavior at the MPLS layer during congestion at ingress, transit, and egress MPLS routers should be ensured by sending traffic and creating congestion points at ingress and transit MPLS routers.

There are two different options for the creation of constraint-based label switched paths: The E-LSPs where the per-hop behavior is inferred from the experimental bit setting in each labeled packet, and the L-LSPs where each per-hop behavior group of traffic uses its own label switched path. The test plan focused both E-LSPs and L-LSPs, depending on vendor support.

The last test cases dealt with real constraint-based routing. Since OSPF-TE continuously updates information about the reserved bandwidth in the MPLS backbone network, it is possible to calculate best routes at the ingress label edge router based on the current network usage. The ability of MPLS routers to explicitly select a route based on traffic class and reserved bandwidth, and to establish a label-switched path using this route, was verified in this section.

Virtual Private LAN Service (VPLS) and Ethernet Point-to-Point VPNs

The MPLS & Frame Relay Alliance test plans *mpls2003.091.00* and *mpls2003.092.00* were used for the tests in this section. They covered:

- Label Binding Distribution via targeted LDP sessions between the provider edge routers
- Data encapsulation of Ethernet and tagged Ethernet frames
- Data encapsulation of ATM and Frame Relay frames

- VPLS service establishment by label exchange between provider edge routers
- Data forwarding to unknown and known Ethernet addresses
- Path tear down and withdraw between provider edge routers

Since VPLS is basically a multipoint extension of point-to-point Ethernet pseudowire links, the tests for point-to-point evaluation were used as a prerequisite for the VPLS tests.

The test plan requested that the transport tunnels were established using the signalling protocol RSVP-TE in order to benefit from the traffic engineering and Diff-Serv processing in the backbone.

In addition to regular Ethernet traffic, the Ethernet pseudowire tunnels also effectively transported TDM emulated traffic (data and voice) using TDMoMPLS technology.

BGP/MPLS VPNs

The MPLS & Frame Relay Alliance test plans *mpls2002.049.01* was used for this section.

This test area was aimed at determining the level of interoperability that can be achieved between RFC2547bis implementations of the various vendors. First, VPN establishment between PE devices was tested. The Layer 3 VPN tests, based on the RFC2547bis draft standard covered:

- Full-mesh Multi Protocol BGP (MP-BGP) peering
- MPLS signalled tunnels between provider edge (PE) routers, using the Label Distribution Protocol (LDP)

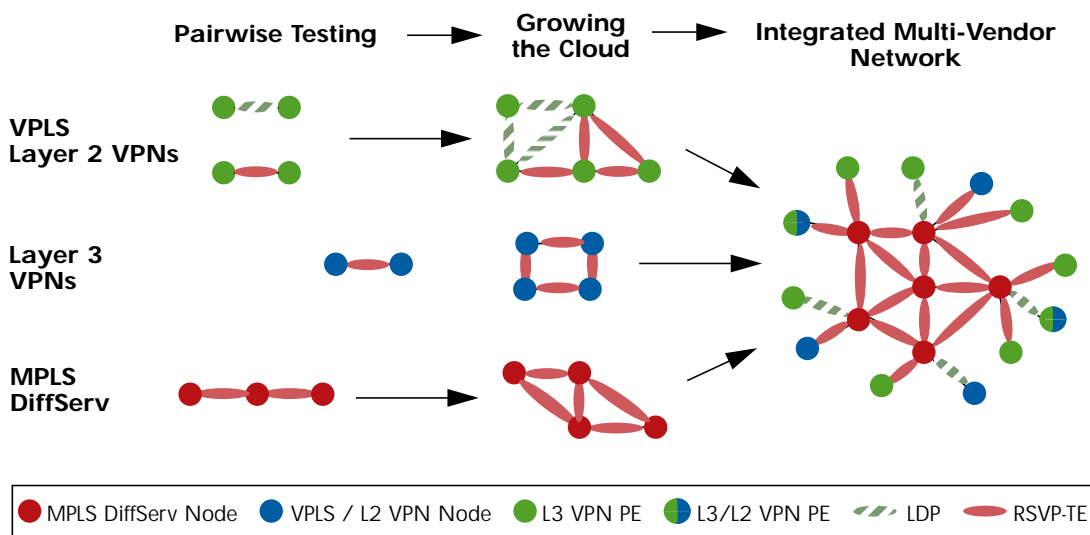
- Dynamic backbone routing with OSPF including traffic engineering extensions (OSPF-TE)
- Dynamic route propagation using BGP or OSPF between customer edge routers (CE) and provider routers (PE) and also between the PE routers themselves.

Similar to the Ethernet VPN area, the test plan requested that the transport tunnels were established using the signalling protocol RSVP-TE in order to benefit from traffic engineering and Diff-Serv processing in the backbone.

Evolution To The Joint Multi-Vendor Network

Three different test areas were tested in parallel during the hot-staging. Small groups were used to check interoperability and scalability. Successful groups in this stage were then connected to form larger groups to build the multi-vendor network. The integration of the groups went forward smoothly and at the end of the hot-staging, all participants formed a piece of the demonstration network

Initial interoperability tests were done in groups of three devices connected to each other. Later on, all devices of each scenario were connected to separate multi-vendor networks for VPNs and the MPLS-DiffServ backbone. Finally, the VPN networks were attached to the edges of the MPLS-DiffServ network and thus integrated into one large common infrastructure. Further interop testing and traffic load generation was done at this stage to verify that VPN traffic benefited from traffic differentiation.



Evolution of Test Stages

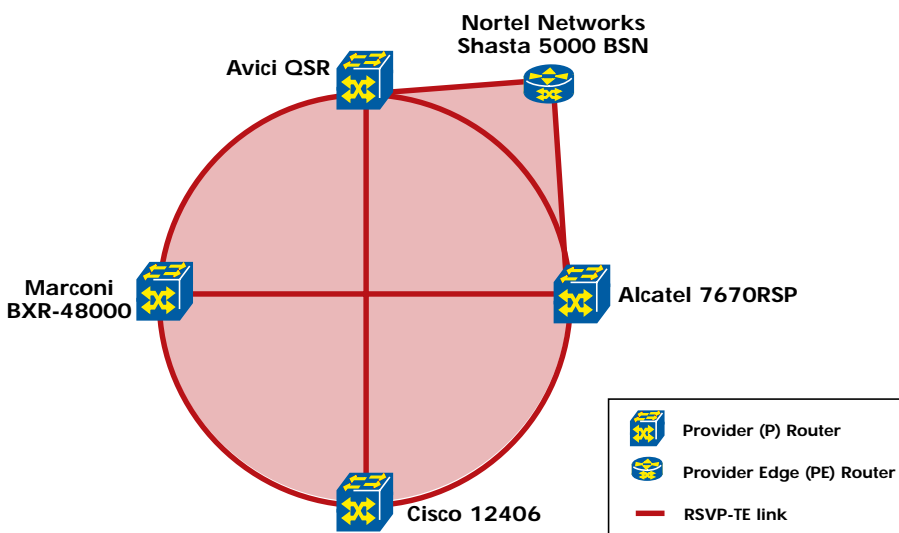
Interoperability Test Results

The goal of this event was two-fold — first, as for typical interoperability test events, to verify and improve the interworking of vendors' implementations, and secondly, to prove that the Service Providers will be able to deploy VPN services over MPLS networks knowing that the network will provide the differentiated services requested by customers for their business-critical applications.

Today, this means more than just to find bugs and correct them to advance the standards compliance. In many cases, implementations rely on draft standards — vendors need to adapt their features to customers' requirements so fast they cannot wait until the final standard is adopted. Thus, the test event was another effort to verify clarity of the current standards.

Results: MPLS Diff-Serv Tests

Basic interoperability of RSVP-TE and OSPF-TE was achieved between all vendors without any signaling or routing issues. Problems that had been noticed in previous events have obviously been corrected. It was very positive to see that the signaling and routing implementations of all participating



MPLS Diff-Serv/Traffic Engineering Backbone

manufacturers were easily interoperable.

The test went forward without major issues when all vendors defined the mapping between experimental bits and per-hop behaviors for Assured Forwarding (AF), Expedited Forwarding (EF) and Best Effort traffic. Initially, the test plan was unclear

about the drop precedence so we discussed the expected handling for different AF types in the same AF class. The IETF standard defines a relative drop precedence relationship instead of absolute values. Some devices were flexible to modify the relative drop precedence while others had fixed relationships. For example, drop precedence is smaller for AF11 than for AF12. The probability to drop AF11 traffic could be varied from zero up to the drop precedence of AF12 in some devices; it was fixed in other devices.

In this multi-vendor test it was of course impossible to use any management application to configure the per-hop behavior policies on all switches. The vendors' engineers had to configure everything manually instead. Naturally, some misconfiguration happened, leading to incorrect prioritization. We recommend to use management systems to control the distribution of policies within a Diff-Serv/traffic engineering enabled network.

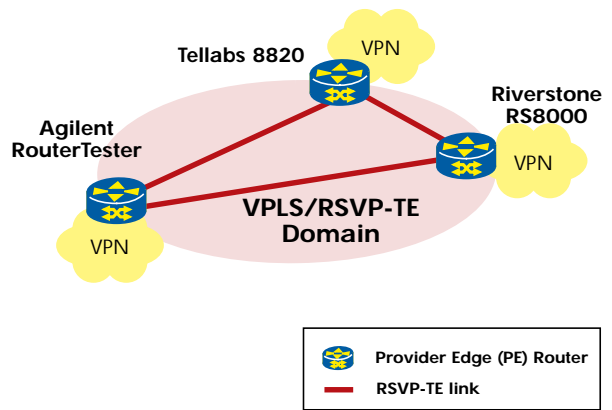
All vendors supported E-LSPs (no matter what physical link type was chosen) and were able to use the full range of experimental bits. We tested with three queues which is realistic compared to today's network requirements. A few vendors also supported L-LSPs over ATM or PoS interfaces.

Only one vendor implemented the DiffServ object as per RFC3270 — all other implementations were designed to accept locally configured Diff-Serv mappings only.

When setting up label-switched paths with traffic engineering and DiffServ, a few interoperability problems were observed by the vendors. Most notably, it is important to recognize an RSVP-TE path as DiffServ-enabled no matter whether the Diff-Serv object is present or not. Therefore, RFC3270 defines a backwards compatibility mode. Other

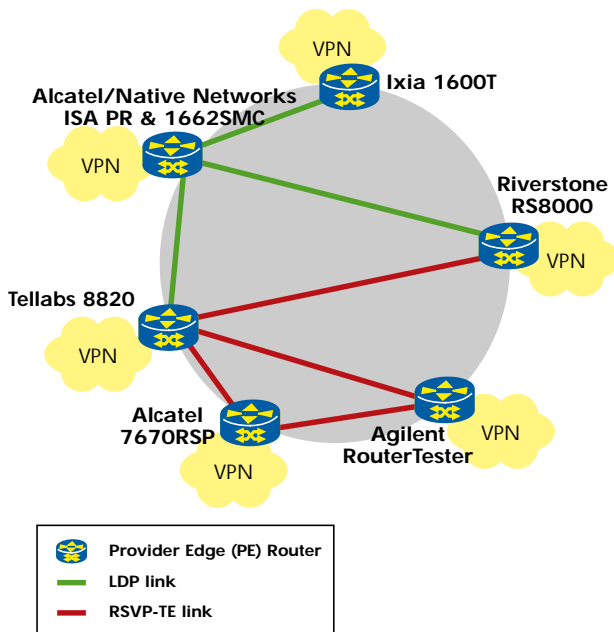
interop issues with DiffServ-enabled path messages were also noted, leading to path rejection. The participants informed us that these issues have either already been resolved during the event, or will be fixed in the following weeks.

Due to limited time, it was not possible to evaluate constraint-based routing in detail. However, five vendors (see table on page 3) confirmed that they used OSPF-TE constraint-based routing during the whole session. Based on current OSPF-TE reserved bandwidth information, an ingress router calculates a suitable end-to-end path that satisfies the tunnel resource requirements — not just based on static path cost. One vendor demonstrated briefly that the constraint-based routing process works as expected.



VPLS Test Network

Results VPLS / Ethernet VPN Tests



Ethernet Point-to-Point Tunnels

Point-to-point Ethernet over MPLS tunnels («pseudowires») were tested according to the Martini draft. In the hot-staging, almost all tested point-to-point connections interoperated as expected. A few issues were only observed on the transport layer — paths could not be established between the provider edge routers.

While it was common for all vendors to use LDP as the transport label distribution protocol last year, this has changed since. Since the use of RSVP-TE has grown and because it is the only protocol with traffic engineering support, RSVP-TE is now used by the majority of vendors for VPN transport labels. Only two participants did not support RSVP-TE for VPLS / Ethernet-VPNs; they mentioned they will implement RSVP-TE support soon.

Two Ethernet point-to-point VPN domains were created, one with LDP and one with RSVP-TE trans-

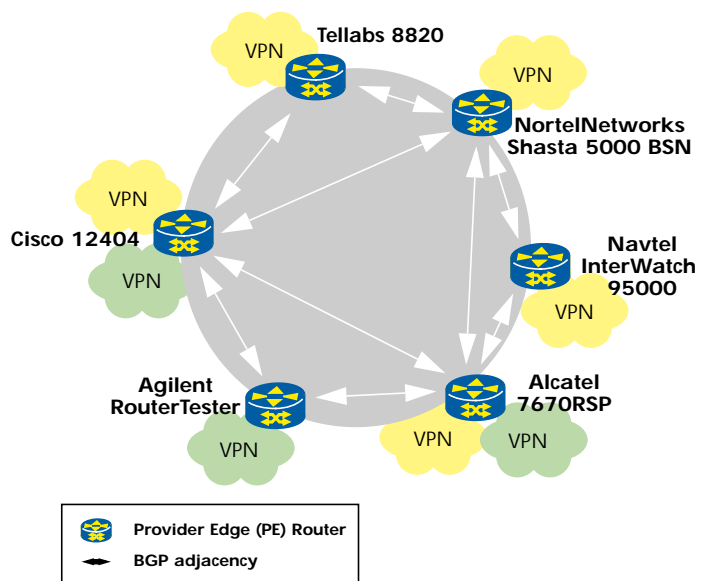
port. Full-mesh VPLS was established in the RSVP-TE domain.

E1 traffic (both data and voice) was emulated over Ethernet pseudowire using TDMoMPLS gateway equipment.

Results: BGP/MPLS VPN Tests

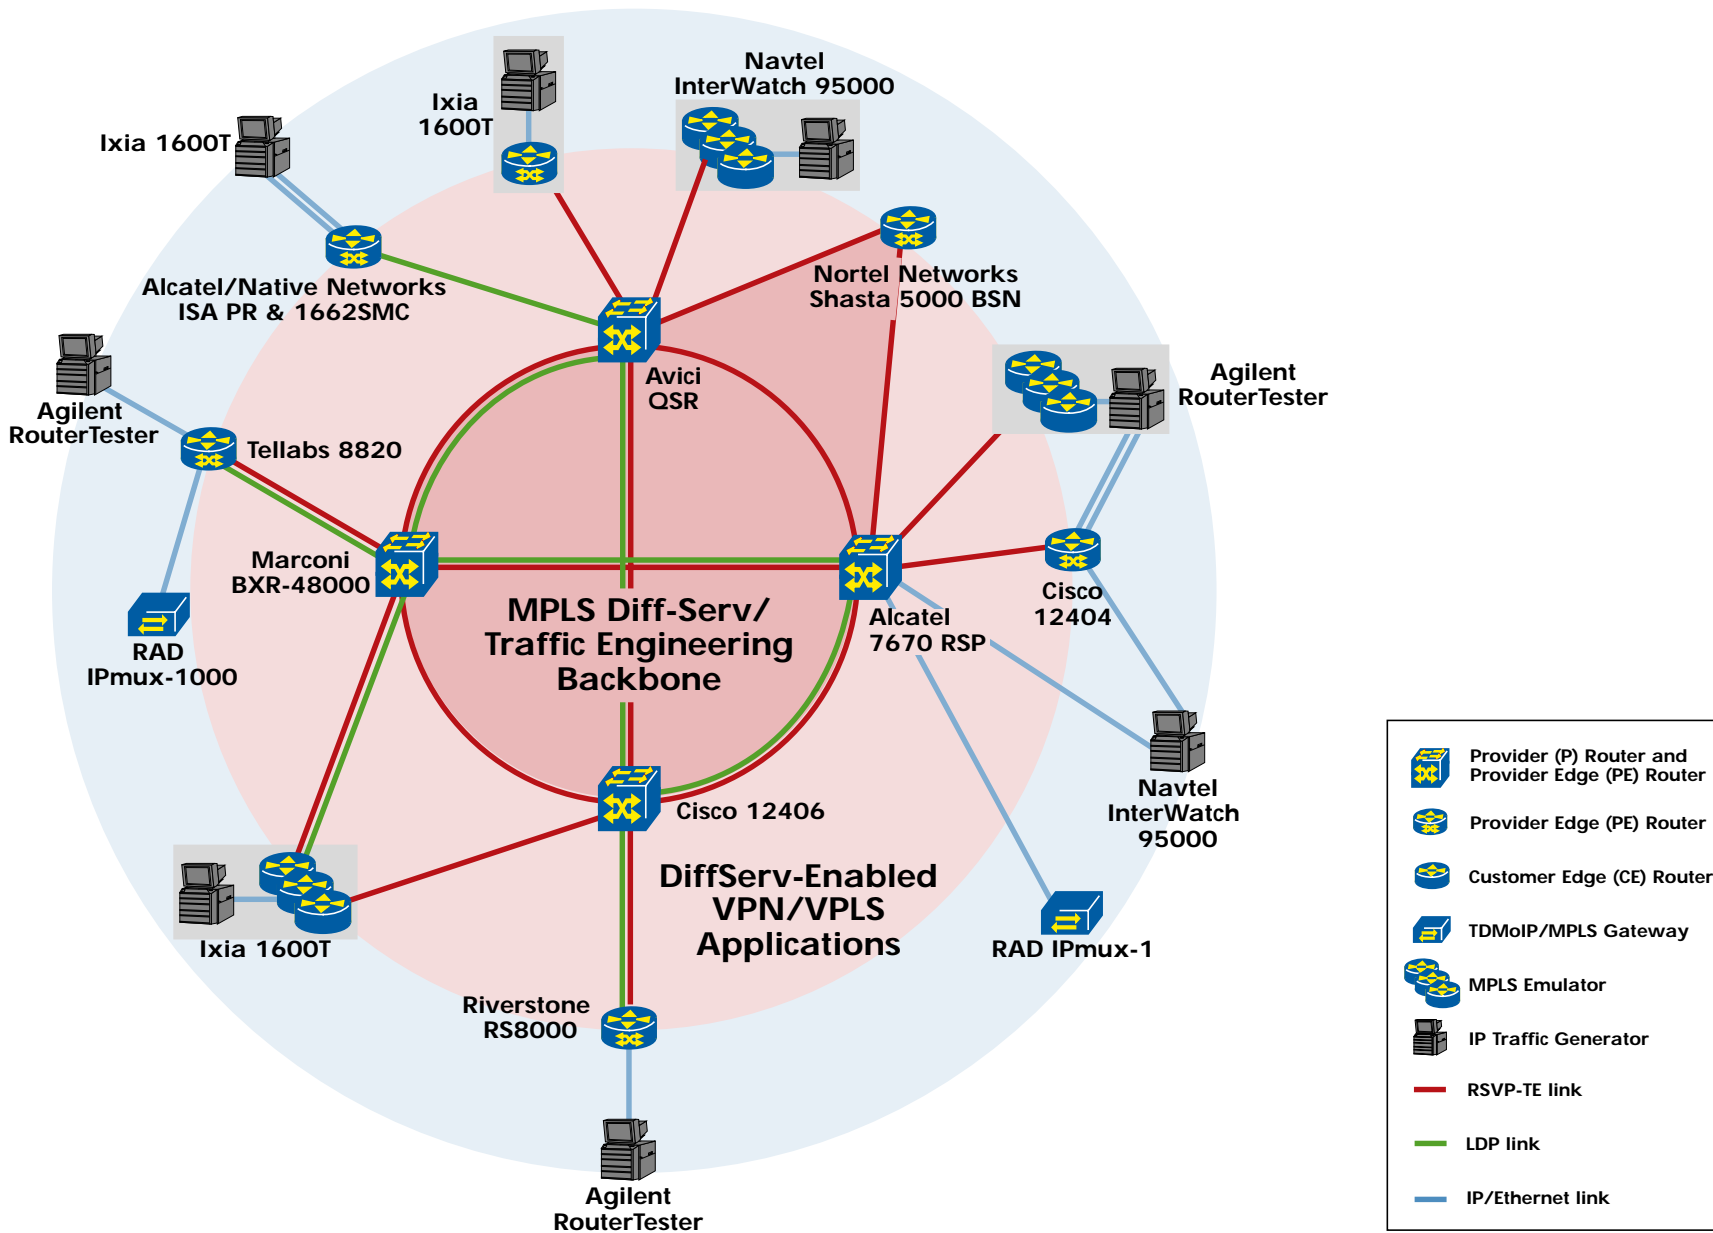
BGP/MPLS VPNs have been used in the industry for a few years already. The test session did not focus testing this area in detail again; they were merely used to demonstrate Diff-Serv transport over VPNs. Connectivity was achieved between two groups of participating vendors; only in one case we observed issues with a BGP implementation.

After establishing BGP sessions, EF and AF traffic was generated from two sites destined to one, and a physical link was oversubscribed to test queuing. Correct traffic prioritization was observed.



BGP/MPLS VPN Test Network

Final Integrated Test Network



Results Summary

Key Features Tested		Results
DiffServ MPLS	Basic Signaling	Majority of combinations interworked
	Per-Hop Behavior	OK
	Packet Exchange During Congestion	All implementations mastered congestion as expected; configuration issues solved later
	Diff-Serv specific RSVP-TE signaling	Only one implementation supported the Diff-Serv object already
L2 VPNs	Interoperability LDP, RSVP-TE	Majority of combinations interworked
	Data Transfer	OK
	Frame Relay / ATM tunnels	Not tested
	Traffic Transfer Over RSVP-TE Tunnels	OK, prioritization observed as expected
	E1 (Data and Voice) Emulated Traffic Transfer	OK
VPLS	Full-Mesh Establishment	OK
	Traffic Transfer Over RSVP-TE Tunnels	OK, prioritization observed as expected
	MAC Address Withdraw	Not tested
L3 VPNs	Interoperability iBGP-MP	Almost all implementations interworked seamlessly
	Data Transfer	OK
	Traffic Transfer Over RSVP-TE Tunnels	OK, prioritization observed as expected

Problems

Problem Area	Description of the Problem	Temporary Resolution, if any	Recommendation
RSVP-TE DiffServ Object	Tunnels were not established	Always define a static EXP-PHB mapping in addition unless you are absolutely sure the network consists only of routers with DiffServ object support	Ensure that the backward compatibility definition in RFC3270 is implemented
DiffServ Provisioning	Traffic was not prioritized as expected	Verify correct configuration using load generators	Use centralized management applications to control network wide per-hop behavior
LDP Application Data Transport	LDP does not support traffic engineering	Configure LDP-over-RSVP-TE hierarchical tunnels to carry LDP data	Substitute LDP transport by RSVP-TE (does not relate to VC labels!)

Conclusion

Today, generally service providers have not implemented a full range of differentiated services within their Virtual Private Networks (VPNs) yet. New revenue to replace decreasing profit from voice and TDM legacy services has to date not been realized. However, competition will grow. There are two main reasons for carriers to implement differentiated services in MPLS: First, the integration of multiple networks into one common backbone for voice, leased lines and Internet-style data. Second, the creation of new differentiated product offers which require accurate, application-specific backbone service levels.

This event has proven that the participating vendors can now support differentiated services over MPLS and provide a common mapping of services and queues. It gives carriers the confidence that converging different networks can still provide a common service guaranteeing/SLA network.

As the deployment of new VPN services and MPLS applications grows, congestion will occur. At this point new bandwidth must be purchased. Traffic engineering across multiple networks allows for better utilization of available bandwidth — therefore absorbing the growth and delaying the time before you have to invest more.

Dynamic constraint-based routing opens a path to more flexible and efficient network usage for operators. Specifically in conjunction with a differentiated services offer, congestion can be avoided more efficiently. The test event started multi-vendor evaluation of dynamic constraint-based routing and signalling.

The test event also verified that edge services are now able to make advantage of traffic engineering and differentiation in the network across multiple vendors.

This was verified for the full range of typical applications: Virtual Private LAN Service (VPLS), Point-to-Point Ethernet tunnels, and BGP/MPLS VPNs carrying IP traffic. It was shown that all these can be easily enabled to carry differentiated services data.

In the past few years, MPLS has grown from a VPN service enabler and a backbone traffic engineering tool to a technology capable of integrating legacy products and new service offers, serving them all according to their specific requirements. This has been made possible by the powerful and versatile network features defined in Multi-Protocol Label Switching.

The variety of standard-conforming products available from many vendors is a particular strength of MPLS. The MPLS & Frame Relay Alliance and the supporting test labs, UNH and EANTC, are proud that the series of interoperability test events conducted since 2001 have been able to improve interoperability dramatically.

References

All IETF drafts mentioned here are work in progress.

- Diff-Serv / Traffic Engineering
 - mpls2003.149.03 — MPLS & Frame Relay Alliance MPLS DiffServ and IGP-TE Interoperability Test Suite
 - RFC3270 — MPLS Support of Differentiated Services
 - RFC2205 — Resource ReSerVation Protocol (RSVP)
 - RFC3209 — RSVP-TE: Extensions to RSVP for LSP Tunnels
 - RFC2597 — Assured Forwarding PHB Group
 - RFC3246 — An Expedited Forwarding PHB
 - RFC2474 — Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers
 - RFC2598 — An expedited forwarding PHB group
 - RFC3630 — Traffic Engineering (TE) Extensions to OSPF Version 2
 - RFC3140 — Per Hop Behavior Identification Codes
 - RFC3272 — Overview and Principles of Internet Traffic Engineering
 - RFC2702 — Requirements for Traffic Engineering over MPLS
- Virtual Private LAN Services
 - mpls2003.092.00 — MPLS Forum VPLS Interoperability Test Suite
 - draft-ietf-l2vpn-vpls-ldp-01 — Virtual Private LAN Services over MPLS
- MPLS Point-to-Point Ethernet VPNs (Pseudo-Wires)
 - mpls2003.091.00 — MPLS Forum Layer 2 VPN Interoperability Test Suite
 - draft-ietf-pwe3-control-protocol-05 — Pseudowire Setup and Maintenance using LDP
 - draft-ietf-pwe3-ethernet-encap-05 — Encapsulation Methods for Transport of Ethernet Frames over IP/MPLS Networks
 - IEEE 802.1D — Media Access Control (MAC) Bridges
 - IEEE 802.1Q — Virtual Bridged Local Area Networks
- BGP/MPLS VPNs
 - mpls2002.094.01 — MPLS Forum BGP/MPLS VPN (RFC-2547bis) Interoperability Test Suite
 - RFC2547 — BGP/MPLS VPNs
 - draft-ietf-ppvnp-rfc2547bis-03 — BGP/MPLS VPNs
- 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

Acknowledgments

We would like to thank Ben Schultz and Ankur V. Chadda from the University of New Hampshire Interoperability Lab (UNH IOL) for supporting this event.



www.upperside.fr

MPLS & Frame Relay Alliance

39355 California St.
Suite 307
Fremont, CA 94538, USA

Tel: +1.510.608.5910
Fax: +1.510.608.5917

www.mplsforum.org

EANTC AG

Einsteinufer 17
10587 Berlin
Germany

Tel: +49 30 3180595-0
Fax: +49 30 3180595-10

www.eantc.com
info@eantc.com

Upperside

54, rue du Fbg
Saint-Antoine
75012 Paris, France

Tel: +33.153.46.63.80
Fax: +33.153.46.63.85

www.upperside.fr
info@upperside.fr

This report is copyright © 2004 EANTC AG. While every reasonable effort has been made to ensure accuracy and completeness of this publication, EANTC assumes no responsibility for the use of any information contained herein. All trademarks are property of their respective owners.