MPLS: Resilient & Scalable

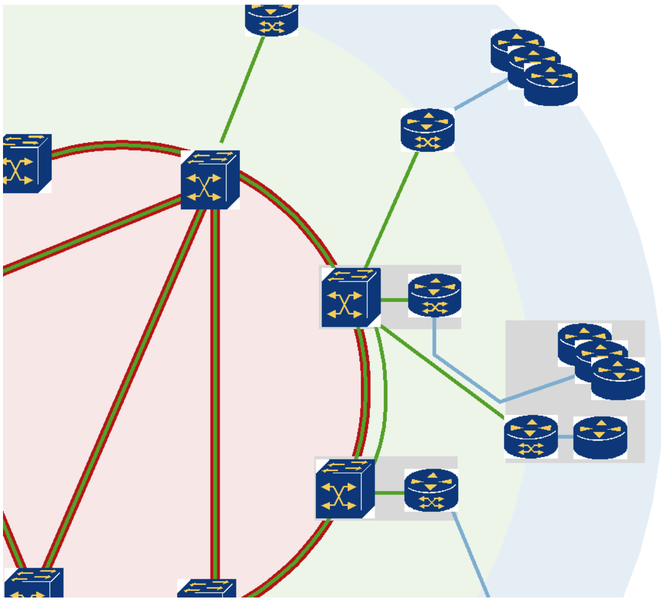

Public Interoperability Event

Test Plan and Results

Organized by:

In collaboration with:

Hosted by

www.upperside.fr

Introduction

This interoperability event is organized by the MPLS Forum and the European Advanced Networking Test Center (EANTC), in cooperation with the ETSI Interoperability Service, 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 5-7, 2003. The demo was hot-staged at EANTC labs in Berlin, Germany, in January 2003.

The test scenarios were designed upon the experiences of previous interoperability test events and also covered new MPLS capabilities which have not been shown before. The focus was on demonstrating scalability of multi-vendor VPNs (layer 2 and layer 3) and interoperability of MPLS Fast Rerouting (FRR).

The demonstration consisted of a network running the two major MPLS signalling protocols in a hierarchical manner: RSVP-TE was used in the core and LDP at the edge of the network allowing VPN services to be established over the hierarchical network.

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 Forum, including Agilent Technologies, Alcatel, Avici Systems, EANTC, Nortel Networks 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	7770 OBX
Avici Systems	QSR SSR
Cisco Systems	7206VXR 12404
Data Connection	DC-MPLS, DC-BGP, DC-OSPF, DC-VPN Manager
IXIA	1600T

NetPlane	Powercode OPTIRoute LTCS
Nortel Networks	Shasta 5000 BSN
Quallaby	Proviso
RAD Data Communications	IPmux-1
Redback	SmartEdge 800
Riverstone Networks	RS8000
Spirent	AX4000

Test Areas and Test Plan

After speaking with a number of European service providers, we developed an idea of what service scalability needs they have this year and the next. Realistic numbers were chosen so as to make the results useful to the industry in the immediate future.

Although the main focus of these tests was to show the scalability of MPLS networks, it is an ongoing important effort to further improve interoperability. For this reason we conducted test cases to cover both interoperability and scalability test areas.

	L2 VPNs	BGP/MPLS VPNs	Fast Rerouting	Management Solutions	TDM over IP
Agilent Tech.	*	*			
Alcatel		*	*		
Avici Systems			*		
Cisco Systems	*	*	*		
Data Connection		*	*		
Ixia	*	*			
NetPlane		*	*		
Nortel Networks		*			
Quallaby				*	
RAD					*
Redback	*	*			
Riverstone	*	*			
Spirent	*	*			

Vendors participating in the different test areas

Layer 3 VPNs

Interoperability. This test area was aimed at determining the level of interoperability that can be achieved between RFC2547bis implementations of the various vendors. First we started to test VPN establishment between different pairs of PE devices. 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.

Scalability. This test area aimed to define the scalability characteristics of the PE routers under test with regard to the number of routes that can be held in the network across multiple VRF's and the maximum number of VRFs that can be maintained in the network. The term *maximum* here refers to the required scalability for a typical Service Provider network.

The scalability performance of an BGP/MPLS VPN network is an important consideration for service providers wanting to deploy IP VPN services based upon RFC2547bis standards. Any network solution must scale linearly as the number of customers and the size of the customer networks increase.

Based on the input from service provides, we established the following scalability goals for the test:

- 255 VPNs configured per provider edge router (virtual router functions)
- 1000 routes per VPN in ten VPNs, plus 10 routes per VPN in 245 VPNs
- All VPNs are configured on all *n* provider edge routers available in the demo network

These figures represent two scenarios based upon a service provider with 255 VPN customers.

The first scenario assuming 10 large customers with 1000 routes from the headquarters. The second scenario assuming the remaining customers would have *n* locations with only 10 routes each.

Layer 2 VPNs

Interoperability. The L2 VPN tests according to the Martini draft covered:

- VC Label Binding Distribution via targeted LDP sessions between the provider edge routers
- Data Encapsulation of Ethernet VLAN frames
- Point-to-point transmission of Ethernet packet data over MPLS

Scalability. Each vendor combination was tested with up to 200 parallel »Martini« Ethernet VLAN tunnels (pseudo-wires). The tunnels were setup between the provider edges; load generators injected and analyzed bidirectional traffic via Ethernet VLAN (IEEE 802.1q) interfaces attached to both provider edge routers.

Fast Rerouting

Service guarantees using Fast Rerouting:

The aim of the fast reroute testing was to prove provisioning of guaranteed VPN services within an MPLS network. An extension to traffic engineering known as Fast Reroute is a newly developed MPLS application (work in progress, still in draft form). Fast Reroute is a link and node protection mechanism to minimize the packet loss during a label-switched path (LSP) failure by rerouting traffic onto a backup LSP. The primary and backup LSPs are signaled using an optional extension to the Resource Reservation Protocol with traffic engineering extensions (RSVP-TE).

There are two Fast Reroute modes:

- Facility (Bypass) and
- Detour.

Both offer distinct advantages and disadvantages, but previous fast reroute related drafts have not accounted for interoperability between the two. The most recent IETF draft for fast reroute, draft-ietf-mpls-rsvp-lsp-fastreroute-01(work in progress) provides guidelines to develop the ability for either draft to interoperate while maintaining the advantages of both modes.

Routers in a network supporting MPLS fast reroute take different roles:

- PLR — Point of Local Repair: Head-end label switch router of a backup tunnel (LSP) that is also a label switch router along the primary tunnel (LSP); responsible for switching to the backup path in the case of a failure or to restore the primary path.
- MP — Merge Point: Tail-end label switch router of a backup tunnel (LSP) where the backup path rejoins the protected primary tunnel (LSP); responsible for merging traffic arriving on the backup and primary path.
- Backup Mid-Point: Mid-Point label switch router along the backup path; takes a passive role.

Interoperability. The following areas were covered during the tests:

- Rerouting of traffic onto the backup LSP during a failure of the primary LSP
- Link Failure between the point of local repair (PLR) and the merge point (MP) nodes
- Failures of the mid point node

The routers were expected to switch-over to backup LSPs with a reasonable low packet loss (equivalent to small rerouting time), in the order of magnitude of 50 ms. To verify this time, a voice circuit application, TDM over IP gateway was used to show that the rerouting took place without service interruption. In addition, IP traffic generators were used to verify the packet loss when the MPLS paths were sent over the backup tunnels.

Final Integrated Test Network

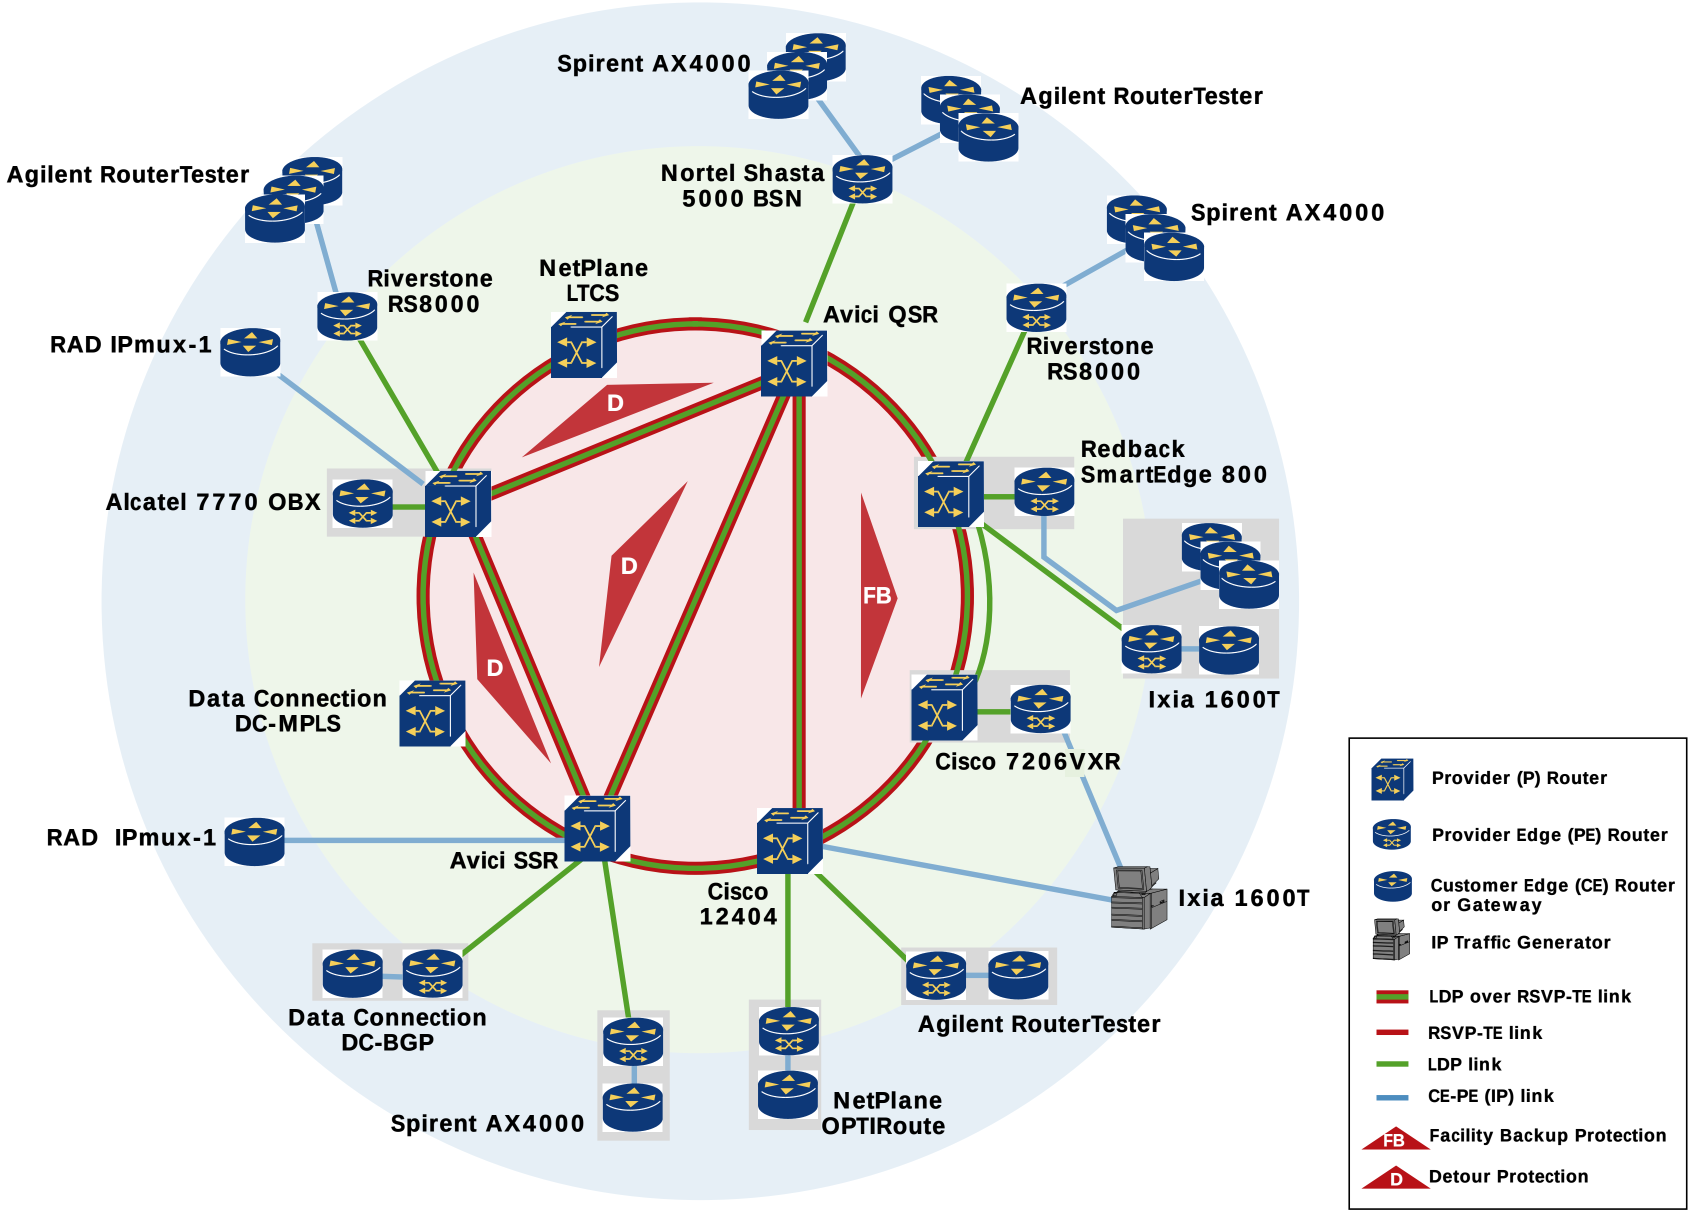

Evolution To The Joint Multi-Vendor Network

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 tests were done in pairs of two (VPN) or three (fast reroute) devices, verifying the interoperability first, followed by one-to-one scalability evaluation. Pairs who passed a couple of these tests were integrated into a larger group to verify scalability at the next level. Later on, all devices of each scenario were connected to a multi-vendor network. This integration went forward smoothly.

As soon as full-mesh interoperability and scalability had been verified, the VPN Layer 2 and VPN Layer 3 networks were attached to the edges of the Fast Reroute network and finally integrated into one large common infrastructure. Further interop testing was necessary at this stage because the roles of the routers changed and because the LDP signalling of the VPN groups was tunneled through the RSVP-TE signalling of the core.

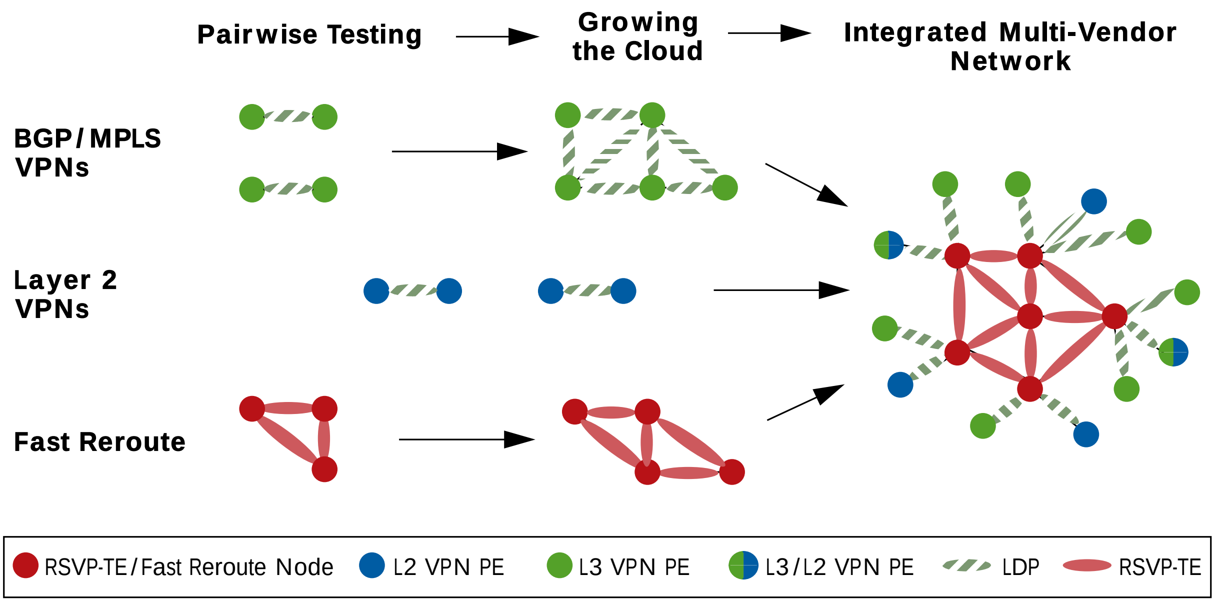

Evolution Of Test Stages

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 MPLS networks knowing that the network will meet the immediate scalability requirements.

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: Layer 2 VPN Tests

Ethernet over MPLS was tested according to the Martini draft. In the hot-staging, all tested point-to-point connections interoperated as expected and achieved or exceeded the scalability goal of 200 Ethernet VLAN tunnels per provider edge router. This test included devices from Cisco, Redback and Riverstone, plus analyzers from Agilent Technologies, IXIA and Spirent.

Results: MPLS/BGP VPN Tests

Basic interoperability of all the vendors' MPLS/BGP VPN implementations was achieved in the vast majority of combinations as expected; these VPNs have already been demonstrated at several interop events. The different provider edge routers (PEs) interworked almost without any problems; a total of 45 pairs with ten vendors were evaluated in three days.

The testing included routers and protocol stack implementations from Alcatel, Cisco Systems, Data Connection, NetPlane, Nortel Networks, Redback and Riverstone Networks. Provider edge and customer edge routers were emulated by test equipment supplied by Agilent Technologies, IXIA and Spirent.

The scalability goals set in the test plan were reached by more than 2/3 of the vendors:

- The final network consisted of 255 VPNs
- Each with ten provider edge routers and
- 10–1000 routes injected by each provider edge into each VPN.

So the network was scaled in two dimensions: The number of VPNs and the number of routes per virtual route function. All participating devices scaled easily for VPN signalling and routing in a multi-vendor environment. This is especially impressive because it was the first scalability interop test for MPLS/BGP VPNs ever.

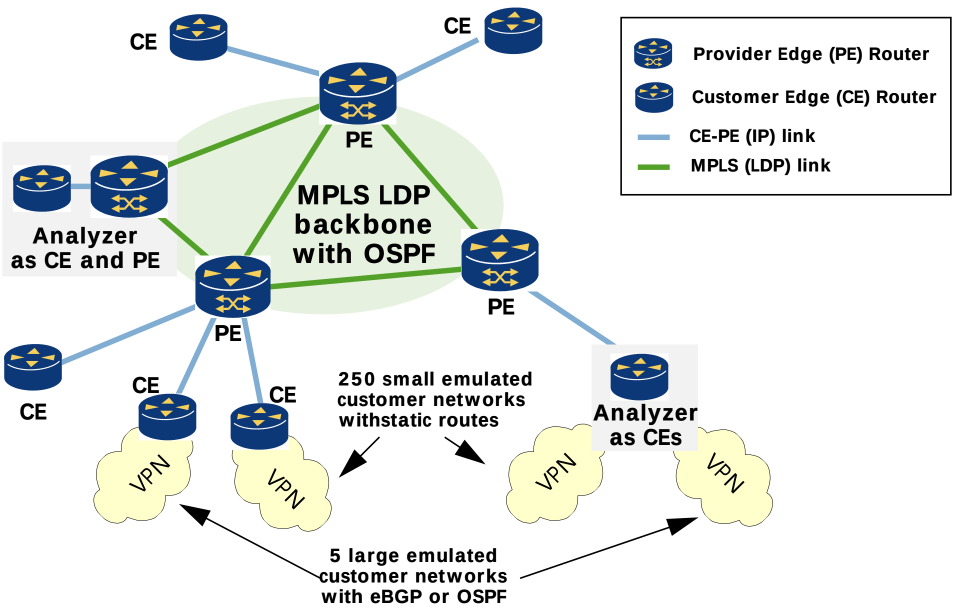

Topology for MPLS/BGP VPN Tests

Data path basic connectivity was evaluated in several cases as soon as the backbone came up, but traffic performance measurements were not the focus of this test. Obviously, there was not enough time during the hot-staging to test if the large-scale network can be maintained for an extended time.

It is safe to say that customers can rely on robust operation and scalability of all BGP/MPLS VPN routing and signalling implementations tested.

Results: Fast Rerouting Tests

For the first time ever, this interop demo showed both alternatives for MPLS RSVP-TE fast rerouting — and proved that both can reroute within the order of magnitude of 50 ms, and scale for carrier requirements:

- Facility Backup (one common hierarchical tunnel protects all label switched paths (LSPs) of a relation)
- Detour Backup (each LSP is protected by its own backup LSP)

The scenario included Alcatel (detour backup), Avici (facility and detour backup), Cisco (facility backup), Data Connection and NetPlane (both supporting detour).

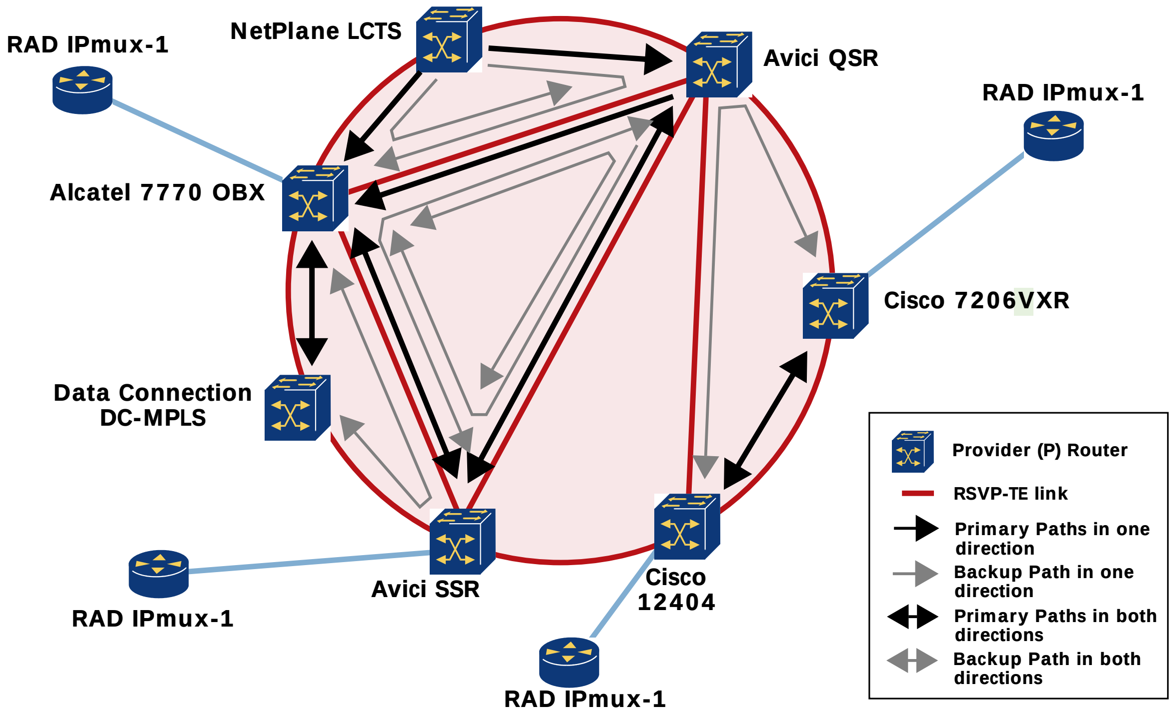

Topology for Fast Rerouting Tests

In both scenarios, it was possible to set up protected paths. The switch-over times were measured from the users point of view by the use of RAD's IPmux-1, a TDM over IP gateway carrying E1 circuits and attached to the network as customer edge (CE) router. In all the cases that could be verified, the gateway showed rerouting times between 14—34 ms; in one test case, there was no packet loss noted at all (< 1 ms / all packets buffered during rerouting).

Given that this was the first public interop test for detour and the second for facility backup, the test results were promising. They helped to identify interoperability issues, mainly caused by not well-defined draft standards.

Results: LDP and RSVP-TE Tests

It was very positive that RSVP-TE implementations were interoperable in all cases — in contrast to what had been seen at the last MPLS Forum organized interop event at Supercomm in June 2002. None of the issues noticed there were noticed again. The only problems in this area were related to interworking between signalling and routing (OSPF-TE).

A few basic LDP interoperability issues were still seen, mainly with new implementations that have not been tested for interoperability often before.

As Supercomm 2002 paved the way for RSVP-TE interoperability, this event was a milestone for LDP over RSVP-TE tunneling:

In the final network, the VPN provider edge routers used the Label Distribution Protocol (LDP) for MPLS signalling, while the backbone required the Resource Reservation Protocol (RSVP-TE) to enable fast rerouting. For end-to-end VPN connectivity, the LDP paths were tunneled through the RSVP-TE core — resulting in a two-tier hierarchical network.

The hierarchical LDP over RSVP-TE tunnels added a level of complexity that was underestimated. It is important to understand exactly how the penultimate and egress routers will negotiate and process the stacked labels. Interoperability issues were mainly caused by misconfigurations that could not be resolved in time. In general, it seems advisable to perform more multi-vendor testing in this area to understand implementation specifics and configurations.

Results Summary

Key Features Tested		Results
L2 VPNs	Interoperability Ethernet VLANs	OK
	Scalability 200 Ethernet VLANs	OK
	Data Transfer	OK
L3 VPNs	Interoperability LDP	Most combinations interoperable
	Interoperability MP-BGP	OK
	Scalability 255 VPNs	All implementations reached 255 VPNs
	Scalability 10–1000 routes per provider edge	All implementations reached 10 VPNs x 1000 routes (BGP/OSPF routing) plus 245 VPNs x 10 routes (static routing)
	Data Transfer through VPNs	A few combinations tested; no issues found
Fast Reroute Interoperability Detour and Facility Backup		Backup tunnels established in most cases, switch-over verified (< 50 ms, SDH grade resiliency)
LDP over RSVP-TE Tunnel Interoperability		Unresolved configuration issues in a couple of cases

Problems

The format of this table is based upon the similar problem table presented within the MPLS Forum's SUPER DEMO report from Supercomm 2002. This allows the reader to see how the industry is progressing.

Problem Area	Description of the Problem	Temporary Resolution, if any	Recommendation
RSVP-TE Signalling	Tunnels were not established	Many issues were related to routing and the routing tables were corrected to work.	Investigate how to provide more cohesion between RSVP-TE and OSPF-TE
LDP Signalling	Sessions were not established; LDP options used in an incompatible way	Arrange topology to keep incompatible LDP implementations away from each other	Work to reach an implementation agreement regarding the use use of label ranges over Ethernet.
OSPF-TE support	A few vendors did not support OSPF-TE.	—	Provide support for OSPF-TE.
	Issues with some OSPF-TE implementations.	Different vendors modified bandwidth values and priorities to work around some of the issues.	Work to clarify acceptable default TE parameters and enhance error checking capabilities in the draft-katz-yeoung-ospf-traffic specification
Fast Reroute (Detour)	Class types in fast reroute object not backward / forward compatible between implementations of different versions of the fast-reroute draft	Use only the latest version of the draft	Start a discussion at IETF to make the standard more clear about recommended / required behavior.
Fast Reroute (Facility)	LSR didn't know that it is a merge point until a failure occurs; so it didn't merge paths	—	An LSR which is on the path of a protected LSP should always assume that it is a merge point
LDP over RSVP-TE Tunnels	Data Transfer through hierarchical tunnel was not always possible.	—	Perform more multi-vendor testing to identify implementation specifics and configurations. Enhance implementations if required.

Conclusion

This event was a significant achievement and another confident step forward for the advancement of MPLS technology and deployment. Here, for the first time, scalability of MPLS implementations has been demonstrated in a public event. The Service Providers have been quite unsure about the scalability claims being made, and this event was to prove that the current needs for Service Providers can be met by the implementations. This was also the first time both »facility backup« and »detour backup« techniques of fast reroute have been tested in a public event.

In conclusion, the testing verified the fast reroute techniques can support the typical services a Service Provider will have to provide to its customers over their MPLS network. This is indeed reassuring for the Service Providers to see the equipment manufacturers coming together and proving that their implementations can meet the Service Provider needs in a scalable and robust fashion.

References

All IETF drafts mentioned here are work in progress.

RFC3036 — LDP Specification

RFC2205 — Resource ReSerVation Protocol (RSVP)

RFC3209 — RSVP-TE: Extensions to RSVP for LSP Tunnels

draft-martini-l2circuit-trans-mpls-10 — Transport of Layer 2 Frames Over MPLS

draft-martini-ethernet-encap-mpls-01 — Encapsulation Methods for Transport of Ethernet Frames Over IP and MPLS Networks

draft-ietf-ppvpn-rfc2547bis-03 — BGP/MPLS VPNs

draft-ietf-mpls-rsvp-lsp-fastreroute-01 — Fast Reroute Extensions to RSVP-TE for LSP Tunnels

draft-katz-yeung-ospf-traffic-09 — Traffic Engineering Extensions to OSPF Version 2

MPLS Forum: Super Demo 2002, Test Plan & Results, June 2002

Acknowledgements

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.

MPLS Forum
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

ETSI
650, Route des Lucioles
06921 Sophia-Antipolis
Cedex, France

Tel: +33.492.94.42.00
Fax: +33.493.65.47.16

www.etsi.org / plugtests
infocentre@etsi.fr

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 © 2003 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.