MPLS 2010 similar to its predecessors will offer its delegates an exclusive opportunity to witness the state of the novel networking technologies in an independent setting. Isocore once again built a comprehensive test bed validating the interoperability of leading vendors, and the co-existence of multiple technologies across a common network infrastructure. MPLS2010 offered a perfect public platform for the delegates to witness the results of a first-ever multi-vendor standards-based MPLS Transport Profile (MPLS-TP) interoperability testing. MPLS-TP testing will showcase statically provisioned label switched paths (LSP) with protection switching, Ethernet service delivery over static pseudowires (PWs), MPLS-TP OAM including BFD connectivity check (CC) and LSP ping for on-demand connection verification (CV), PW status notification and interworking with IP/MPLS. Additionally, Isocore showcased the results of the multicast-LDP, multicast VPNs and BFD and LDP over dynamically signaled RSVP-TE tunnels interoperability. The other objective of MPLS2010 demonstration was to showcase the state of implementations supporting OAM across AS boundaries, through virtual circuit connection verification (VCCV) across Inter-AS Multi-Segment PWs (MS-PW).

The testing referenced a compilation of individual tests extracted from Isocore's library of test plans which have continuously evolved through constant input from its members. Its members comprises of companies from service provider, vendor community and test equipment manufacturers. Isocore primarily focuses on technologies that are standardized by various standard development organizations, such as IEEE, IETF, ITU-T and others.

For the fall leading edge code (LEC)/MPLS2010 staging, a week-long test event was scheduled at Isocore's headquarters in Washington metro area during the week of October 4, 2010. Figure 1 illustrates various technology areas that were included within the scope of the fall LEC event, and results obtained were presented at the public demo.

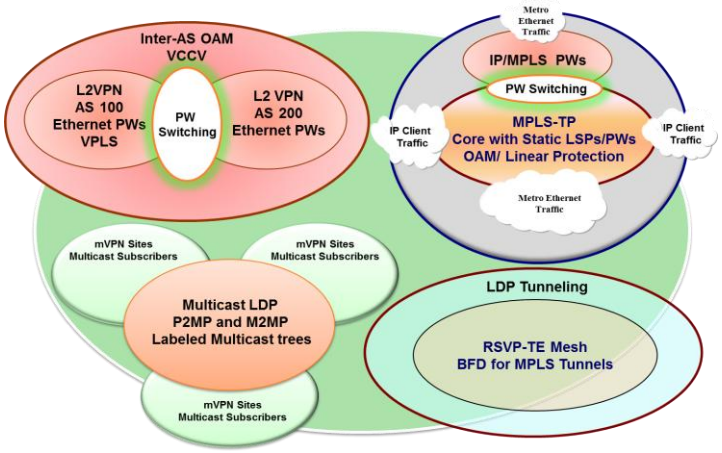

Figure 1: The technologies considered for MPLS2010 Demo

The fall LEC testing saw participation from 9 entities. Representatives from network equipment manufacturers, test equipment vendors worked towards a goal of achieving interoperability in various technology areas and presenting an integrated stable multi-vendor network to the conference delegates. External support from service providers was also present.

This white paper presents a high-level overview of what was tested. For some test areas, results from Isocore spring LEC event are also presented to demonstrate the evolving implementations as standards become stable, and MPLS-TP being one of the classic examples, where in the spring LEC event, we had participation of only two vendors compared to five in the fall LEC event. Figure 2 shows the comprehensive setup highlighting the roles played by all participating nodes and logical representation of the network physical topology.

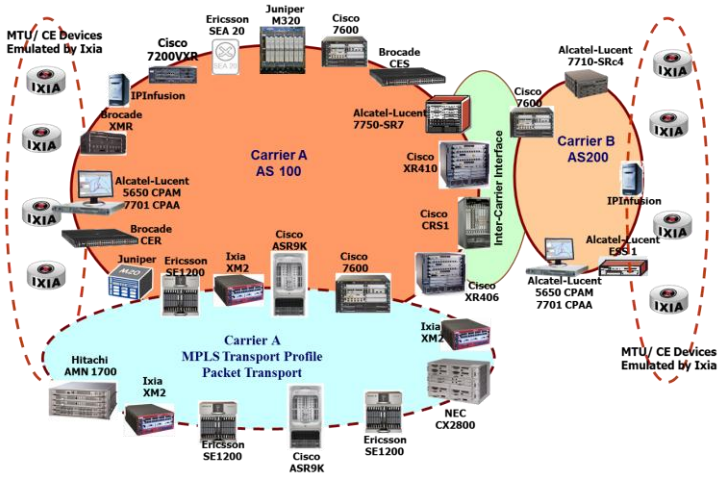

Figure 2: Logical Representation of MPLS2010 Demo Network

During the initial stages of planning for the MPLS2010 demonstration, several technologies were proposed through the feedback received from participants. Upon ranking these topics in the order of priority, lead to the short list of the following areas, which formed the core of the MPLS2010 Interoperability demonstration. These topics also align with the MPLS2010 International conference theme for this year:

1. *MPLS Transport Profile*
   a. *Statically provisioned co-routed LSPs*
   b. *Linear Protection*
   c. *MPLS-TP OAM - including BFD connectivity Check (CC) and LSP Ping using ACH*
   d. *PW Status notification and Interworking with IP/MPLS*
2. *mLDP in both global and mVPN context*
3. *BFD over RSVP-TE Tunnels*
4. *LDP Over RSVP*
5. *VCCV for Inter-AS Multi-Segment PWs*

In addition, for the staging event, Alcatel-Lucent provided their control plane assurance manager (5650 CPAM) which simplified monitoring of the control plane changes happening in the test network throughput the event.

The Isocore IP/MPLS test network started with a flat network with one autonomous system to give vendors an opportunity to test against each other. The network was later split into two AS for testing the Inter-AS aspects of the testing scope. One of the AS, AS100 was built on top of the underlying MPLS transport network, which was used primarily for all MPLS-TP related tests. Figure 2 also illustrates the final integrated testbed at the conclusion of the fall LEC/ MPLS2010 staging.

Testing observations and Results

Similar to earlier events, fall 2010 LEC event offered a perfect staging platform for MPLS2010 public interop demo. The following sections describe the test cases executed and the results observed during the event. Majority of the tests needed more than the time allocated for the LEC event, but what was produced at conclusion of a 4-day testing event is commendable.

1. MPLS Transport Profile

MPLS-TP technology facilitates convergence of carriers' next generation networks onto a single transport technology. The MPLS-TP OAM is a subset of functions within the transport profile used for network performance monitoring, fault management and protection switching. It is a major building block within the MPLS-TP framework with capability to deliver carrier grade OA&M functions including sub-50ms traffic resiliency. In a nutshell, MPLS-TP enables MPLS to support packet transport services with a similar degree of predictability, reliability and OAM to that found in existing transport networks.

This section describes an overview of the MPLS-TP interoperability tests performed in the Fall LEC event. This event focused on OAM and resiliency testing using IP identifiers. The executed tests are categorized in to four areas, Static bidirectional co-routed LSP set up, Linear protection, MPLS-TP OAM - including BFD connectivity Check (CC) and LSP Ping using ACH, Switching of static and dynamic PWs, and IP/MPLS interoperability for verifying end-to-end services.

***Static bidirectional co-routed LSP set up*:** During the test, many bidirectional co-routed LSPs were setup between two LERs (Label Edge Routers) from multiple vendors in a one-hop (or back-to-back) configuration or with an LSR (Label Switch Router) along the path of the LSP between the two LERs. MAC addresses were either statically configured for each LSP or dynamic ARP were used to retrieve the remote MAC addresses. During the testing, agreements on the label range to be used for MPLS-TP label assignments were agreed between the participating vendors to overcome the interoperability of different vendors supporting different label ranges. In a real-world deployment, this could become an interop issue if in a multi-vendor network; these things are not agreed prior to deployment or testing phase.

Figure 3 illustrates the mesh point-to-point MPLS-TP LSPs that were created during the testing and indicate what vendor products participated in the test. The figure differentiate between the working and protecting LSPs by use of different colors, and associate the working and protecting paths by use of special symbols on the LSPs. A comprehensive mesh as seen in the network indicates the stability and the readiness of the implementations, and adherence to the proposed IETF standards.

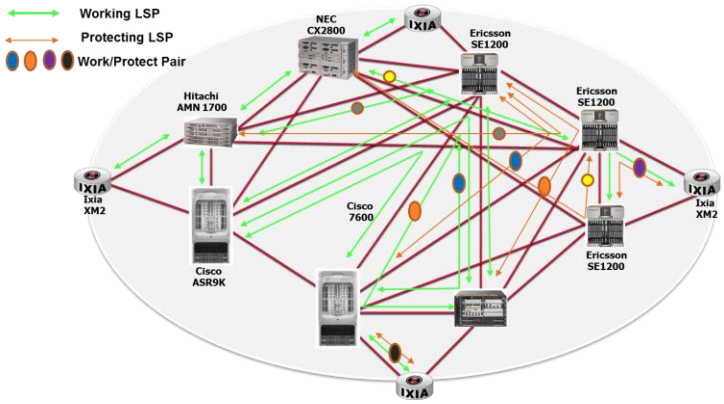

Figure 3: MPLS-TP Demo Topology – Physical

Linear protection: Figure 3 shows the protected LSPs that were created and associated with working LSPs in a 1:1 linear protection configuration. Test cases with either Protection Switching Control (PSC) enabled or disabled were tested. Not all implementations supported the PSC messaging, which helps the LSRs/LER to select the working or recovery path, and to transmit different protocol messages. In tests scenarios, where PSC functionality was disabled, BFD continuity check (BFD CC) was used for detection of loss of continuity to trigger the protection failover. Both revertive and non-revertive configurations were tested during the event. Only Ixia and Ericsson participated in this part of the test.

MPLS-TP OAM - including BFD connectivity Check (CC) and LSP Ping using ACH: Once the LSPs were set up with matching labels, BFD CC (Continuity Check) was enabled to monitor the continuity of the LSPs. BFD CC provides a rapid detection mechanism for LSP LOC (Loss of Continuity), in particular when lower layer may not be able to detect LOC failure at the LSP layer. BFD slow start was not enabled during the test, however, BFD slow start is interoperable with equipment that doesn't support BFD slow start.

LSP Ping using ACH (Associated Channel Header) was tested on each end of an LSP. Each LER supporting the functionality initiated LSP ping to the peering LER; in either a back-to-back configuration or through an LSP in the path, depending on the setup under test.

In the above tests, BFD CC sessions were running concurrently on both primary and backup LSP. When a BFD CC failure was introduced into the primary path, traffic successfully switched to the backup path. In addition, after the BFD CC failure was repaired, the traffic successfully reverted back from the backup LSP to the primary LSP

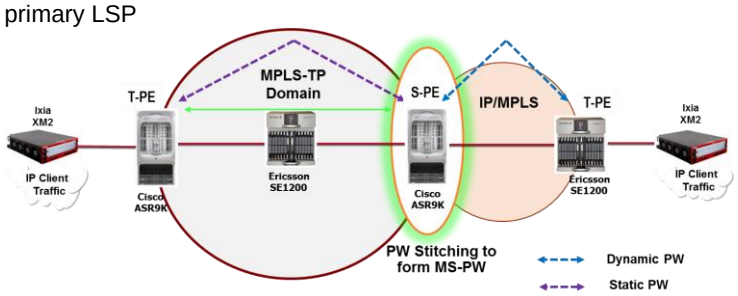

Figure 4: PW switching between dynamic and static segments

Switching of static and dynamic PWs, and MPLS/IP interoperability: Figure 4 illustrates the setup that was attempted to perform MPLS-TP and IP/MPLS interoperability using PW switching between static and dynamic PWs and verifying the status of end-to-end Ethernet services. The testing involved setting up of static PWs between T-PE and S-PE as shown in the figure 4. Following this, a dynamic PW was created between the S-PE and T-PE in the IP/MPLS domain, with S-PE performing the stitching operation connecting the dynamic and static PW, forming a multi-segment PWs crossing from one domain to the other. The end-to-end verification was performed by flapping the attachment circuits, or the transport MPLS-TP LSP.

The MPLS-TP testing included equipment from Ericsson, Cisco, Hitachi, NEC and Ixia. In addition to being a LER for the MPLS-TP, Ixia also provided the client traffic the verification of the MPLS-TP data plane.

2. mLDP within Global and mVPN context

Three vendors participated in multicast Label Distribution Protocol (mLDP) testing at the fall Interop test event. These companies were Alcatel, Cisco, and Ixia. mLDP is a new protocol used to create p2mp and mp2mp labeled multicast trees. The main applications for mLDP are multicast VPN and VPLS. mLDP is typically used in the core and interacts with multicast protocols like PIM or BGP on the edge. These protocols pass the multicast tree request into LDP at

the receiver side of the network and the tree is built back towards the multicast source or Rendezvous Point (RP). When a different branch for the same multicast stream is built, the labels are merged once the paths intersect at the merge point closest to receivers. The testing was based on the IETF draft - draft-ietf-mpls-ldp-p2mp-0x.txt.

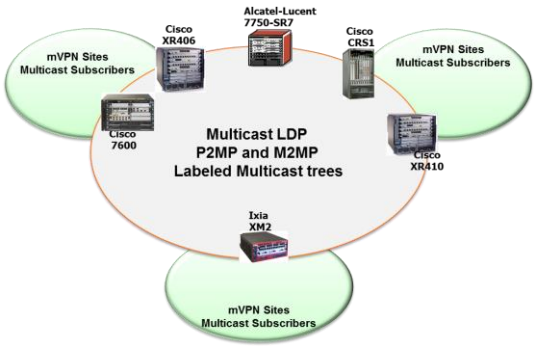

Figure 5: mLDP and mVPN setup – Logical

Two different mLDP scenarios were tested during the event. The first was mLDP in the global context where all multicast streams were part of the same global table and not separated by VPN's. All participating implementations were able to do mLDP signaling during this testing and Cisco and Ixia were able to send multicast traffic into the P2MP LSP's at the source PE router and receive it at the destination PE.

mLDP uses a new protocol element called the p2mp FEC (Forward Equivalency Class) which identifies the root node address and an opaque field used to create a unique p2mp or mp2mp LSP. The opaque field can carry information used to dynamically map multicast traffic to p2mp LSP's on the ingress and egress PE's and is not used by the transit Provider (P) routers. The vendors testing have differences in how they currently used this opaque field which reduced the number of test cases that could be executed.

Multicast VPN architecture allows Service Provider to delivers multicast traffic over existing MPLS/BGP VPN infrastructure. A multicast distribution tree (MDT) is built in the Service Provider network to deliver customer traffic across a common infrastructure, while keeping multicast traffic for each VPN customer separate. There are multiple technologies proposed for implementing mVPN architecture.

In order to leverage MPLS fast forwarding and traffic engineering capability, MPLS technology for multicast delivery has been discussed actively in standard body (IETF). Several drafts have been proposed, but none yet standardized. mVPN with mLDP transport was also tested. mVPN is a method to create per-VPN multicast route tables on the edge PE routers and setup per-VPN paths in the core to carry the aggregated or individual customer traffic. This model currently uses GRE encapsulation but can use RSVP P2MP-TE or mLDP to set up labeled paths as well.

Cisco and Ixia tested mVPN with mLDP. Ixia supports p2mp LSP's for default MDT's which is similar to the existing GRE model. Cisco supports mp2mp LSP's which allows using a single LSP per VPN instead of many p2mp LSP's per VPN. Ixia and Cisco were able to signal LSP for the individual VPN but differences in the LSP type did not allow passing multicast routes or traffic between PEs. Cisco was able to test mp2mp and p2mp LSPs and send per VPN route updates and traffic for both default and data MDTs with Ixia emulating the CE routers.

Due to the plethora of mVPN proposals, different vendor interpretations, and lack of agreement on supporting common references we intend to continue to focus on mVPN in the next round of interoperability testing.

3. BFD over RSVP-TE Tunnels

Bidirectional Forwarding Detection (BFD) for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) is to detect MPLS LSP's data plane failures.

LSP Ping is an existing mechanism for detecting MPLS data plane failures and for verifying the MPLS LSP data plane against the control plane. BFD can be used for the former, but not for the latter. However, the control plane processing required for BFD Control packets is relatively smaller than the processing required for LSP Ping messages. A combination of LSP Ping and BFD can be used to provide faster data plane failure detection and/or to make it possible to provide such detection on a greater number of LSPs. LSP ping is used to carry BFD session parameters.

To use BFD for MPLS LSP fault detection, a BFD session must be established for that particular MPLS LSP. BFD Control packets MUST be sent along the same data path as the LSP being verified and are processed by the BFD processing module of the egress LSR. If the LSP is associated with multiple FECs, a BFD session should be established for each FEC. The default action item after failure detection will be LSP teardown. Sometimes we may need different action item, for example in case of FRR action item should be switch to detour/bypass. So for each BFD session an action item may be specified. Network elements from Juniper Networks, IPInfusion, Cisco, and Brocade were used in this verification.

4. LDP Over RSVP-TE

Traffic engineering in MPLS focuses on optimizing the performance and efficiency of the network service. The MPLS TE is typically building one or multiple fully-meshed LSPs in the network to schedule the network traffic and utilize the bandwidth resource of the network device more efficiently. However, it is not very practical to deploy the RSVP TE throughout the entire carrier network due to the scalability concern. Therefore, the carrier can implement a core RSVP area with core P routers and RSVP TE is deployed in this area. Between PE and P routers, the Carrier can still run LDP signaling protocol which is relatively easier to configure as well as widely deployed and ensure the RSVP TE can tunnel the LDP from PE to PE which is also known as LDP over RSVP-TE tunneling.

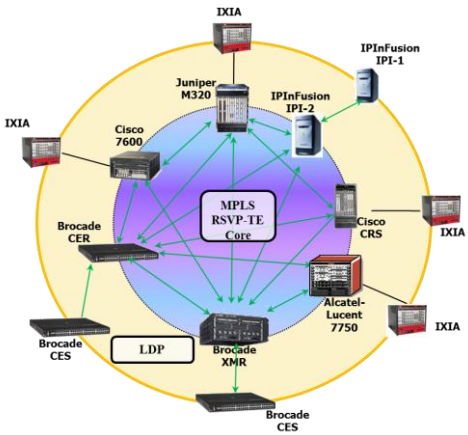

Figure 6: LDP over RSVP-TE Setup

In the Fall LEC event, the LDP over RSVP-TE tunnels was tested by having multiple core P routers from different vendors and between

each other they have fully meshed RSVP-TE LSP configured. The PE routers are running LDP to eliminate the distribution of external routes in the core. The LSPs established by LDP are tunneled through the LSPs established by RSVP. Figure 6 shows the tested logical topology. LDP effectively treats the traffic-engineered LSPs as single hops. When you configure the router to run LDP across RSVP-established LSPs, LDP automatically establishes sessions with the router at the other end of the LSP. LDP control packets are routed hop-by-hop, rather than carried through the LSP. This routing allows one to use simplex (one-way) traffic-engineered LSPs. Traffic in the opposite direction flows through LDP-established LSPs that follow unicast routing rather than through traffic-engineered tunnels. When LDP over RSVP LSPs are configured, carriers can still provision multiple OSPF areas and IS-IS levels in the traffic engineered core and in the surrounding LDP cloud.

Brocade (XMR and CER), Cisco (CRS1 and 7600), Alcatel-Lucent (7750 SR-7), Juniper M320 and IPInfusion participated as Core P routers having RSVP-TE tunnels. The Brocade CES, IPInfusion and IXIA XM2 participated in as PE routers running LDP with its P router peers respectively. LDP over RSVP tunneling was enabled on each P router so that LDP FECs were advertised through the RSVP-TE tunnel and installed over a targeted LDP session for LDP tunneling purposes.

5. VCCV for Inter-AS Multi-Segment PWs

The section focuses on interoperability of switched pseudowires (PW) across an autonomous system (AS) boundary and support of OAM using virtual circuit connectivity verification ping and trace. The participants included Alcatel-Lucent, Cisco, IP Infusion, and Brocade. The diagram shows the setup and vendors which supported this test activity. Not all participants supported the roll of the switching PE at the AS boundary.

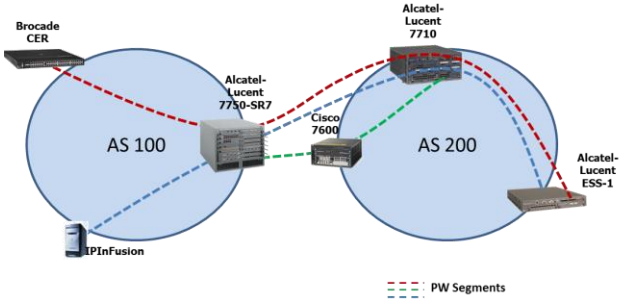

Figure 7: VCCV across PW-Switching

Two separate AS networks were created and PW switching was configured across the AS boundary. The PW within each AS was extended to a remote PE. After the control plane was validate by inspection, bi-directional traffic was forward across the switched PW to verify the end-to-end connectivity was fully established.

Once the PW setup was validated with traffic flow OAM using VCCV ping and trace was tested. VCCV pings were sent from the end node across the PW to verify the proper response on the switched PW. VCCV trace was used to trace the hop by connectivity from the setup. Figure 7 illustrates the test setup and MS-PWs created.

6. IP/MPLS Route & Path Analysis

For this MPLS2010 Fall LEC event Alcatel-Lucent provided Isocore with their 5650 Control Plane Assurance Manager (CPAM) which allowed us to visualize, track and evaluate the control plane of the multi-vendor IP/MPLS LEC network. The 5650 CPAM application is embedded in the 5620 SAM, Alcatel-Lucent's network manager for IP/MPLS networks. Two Alcatel-Lucent 7701 CPAA probes were placed in the network that passively monitored the IGP LSAs and BGP updates originating from all routers. Introducing IP/MPLS route & path analytics to this LEC event proved to be highly effective for control plane assurance and troubleshooting in the multi-vendor IP/MPLS test network. For the first time, we were able to have a detailed view of the dynamic multi-vendor LEC network topology. The tool provided us the visibility into the route changes encountered during testing, allowing us to monitor historical changes to the IGP network topology. Assisted in graphing BGP/IP-VPN route statistics, and reporting on next-hops and number of routes for BGP and IP-VPN route targets. While the probe-based approach provided a multi-vendor route and IP path analysis capability, 5650 CPAM has bundled additional multi-vendor features using the 5620 SAM's SNMP infrastructure.

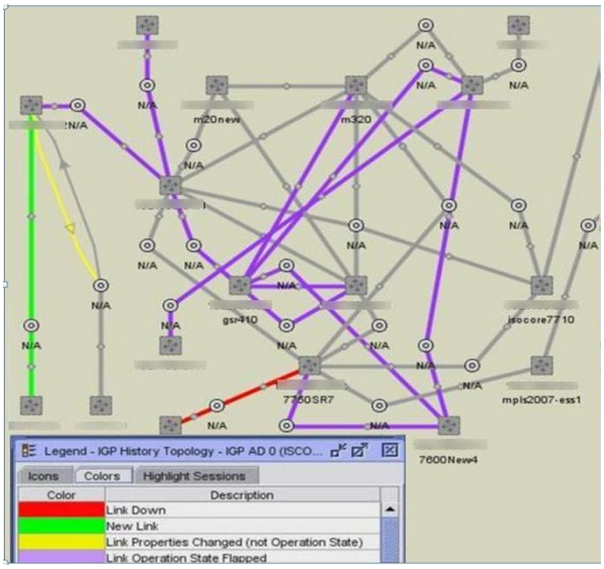

Figure 8: Discovered Test Topology

The MPLS LEC network configuration included multiple AS networks and OSPF Admin Domains in a continuous state of transition as various tests are performed. The 5650 CPAM provided a unique view of this very dynamic network. Figure 8 shows the view of the network as captured during the staging event. The figure illustrates IGP history: added routes (in green), removed routes (in red), flapping routes (in purple) and attribute changes (in yellow).

7. Products and Isocore Members Participated

Alcatel-Lucent	7750-SR7, 7750-SR1, 5650 CPAM, 7710, 7450-ESS
CISCO	ASR 9000, CRS1, GSR XR12410, GSR XR12406, 7606
BROCADE	NetIron XMR, NetIron CES2000 and CER2000
ERICSSON	SE1200, SEA20
HITACHI Inspire the Next	AMN 1710
IXIA	XM2, IxNetwork
ip infusion An ACCESS Company	Protocol Stack
JUNIPER NETWORKS	M320
NEC	CX2800