Second OPTICAL SIGNALING, ROUTING AND MANAGEMENT Test Event

September 20 – 26, 2004

Whitepaper

OSRM Test Event
InterOperability Laboratory
Research Computing Center
University of New Hampshire

121 Technology Drive Suite 2
Durham, NH 03824
(603) 862-0090
http://www.iol.unh.edu

Executive summary

Service providers continue to look to developments in Optical Signaling, Routing and Management (OSRM) technologies to provide increased network automation that will allow them to offer new, diversified services while concurrently reducing operational costs. OSRM enables the dynamic creation, maintenance and and tear down of connections within and between optical networks.

Ongoing OSRM test events at the University of New Hampshire's InterOperability Laboratory (UNH-IOL) are bringing together service providers and telecom equipment vendors in this emerging market to demonstrate and facilitate standards based interoperability. As part of this effort, the UNH-IOL recently conducted a Generalized Multi-Protocol Label Switching (GMPLS) test event. GMPLS is one of several OSRM technologies that Service providers can utilize to automate the delivery and management of bandwidth how and deliver to their customers.

The results of the independent GMPLS testing conducted at UNH-IOL are detailed in this white paper, and suggest that significant progress continues to be made in the areas of OSRM standards development and vendor implementation; this progress, in turn, will enable service providers to sell optical services to customers with greater efficiency, differentiation, and profitability.

Introduction

GMPLS consists of a suite of OSRM protocols that enable dynamic end-to-end provisioning, maintenance and tear down of connections across the electrical and optical transport domains. In effect, it merges IP-based routing, signaling, and management with the optical realm. Based on the standards efforts of the Internet Engineer Task Force (IETF's) Common Control and Management Plane Working Group (CCAMP WG), GMPLS also provides a foundation for multi-vendor interoperability.

The most recent UNH-IOL OSRM test event, held September 20-26, 2004, was designed to validate and improve GMPLS functionality in a multi-vendor network. The event provided a vendor-neutral setting that gave participants the opportunity to assess interoperability and provided valuable feedback that could help refine their implementations.

The UNH-IOL test suites were built in cooperation with Nippon Telegraph and Telephone Corp. (NTT). This collaboration generated an exceptionally realistic and demanding test suite in line with service providers' operational demands, rather than simple conformance or interoperability scenarios.

Participants included NTT, as well as telecom equipment providers Agilent Technologies, Juniper Networks, Lambda Optical Systems, Navtel Communications, Sycamore Networks, and Tellabs Inc.

Test Methodology

The OSRM test event GMPLS enabled network elements from leading telecom equipment vendors in a diversified, heterogeneous network designed to carrier-class specifications. Test cases probed the functionality of various aspects of GMPLS, including hierarchical LSP Setup and Teardown, Traffic Engineering capabilities using OSPF-TE and Constrained Shortest Path First (CSPF) calculations, and LSP re-optimization. Test cases supported both IPv4 and IPv6 traffic, the next-generation Internet protocol. The test event culminated with the creation of a GMPLS LSP provisioned by Layer 1 VPN to the public Internet which participants used to accesse-mail and Web based applications. Building upon previous UNH-IOL tests that verified critical OSRM features, the October 2004 UNH-IOL test plan pushed interoperability and functionality to territories not previously explored. These cases are detailed in the following sub-sections.

TE link Configuration

The GMPLS architecture offers OSRM functionality over a variety of data-plane resources, called TE links. TE links can be configured to support many attributes, including numbered links, unnumbered links, bundled links, Forwarding Adjacency (FA)-LSPs, protection types, etc. Support for all of these attributes allows service providers the maximum flexibility in establishing a G-LSP end-to-end. However, not all features are suitable for all types of devices. For example, implementing numbered links is a likely requirement for a Packet Switch Capable (PSC) device, but unnumbered links are more applicable to Time Division Multiplexed (TDM), Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC) devices. As TE links become more complex, proper encoding and decoding of the sub-TLVs in OSPF-TE LSAs is vital to interoperability among multi-layer devices.

Test Case #1. TE links Advertisement
Properly interconnected devices were configured to exchange TE links via OSPF-TE LSAs. The type of TE links tested were numbered and unnumbered links, and the feature tested was FA-LSP. For the most part, the various configurable parameters characterizing the numbered and unnumbered TE links were properly exchanged, and

LSPs were allowed to setup via these TE links. The Unreserved Bandwidth sub-TLV was properly accounted for as LSPs were set up and torn down. FA-LSP was established in both the control-plane and data-plane.

Bidirectional LSPs

A bidirectional LSP allows for rapid LSP establishment, symmetrical paths, and equal resource utilization in each direction. A bidirectional LSP is established by including an Upstream label object in the corresponding Path message generated by the initiator. The terminator may begin forwarding data via the LSP towards the initiator as soon as it accepts the Upstream Label, and the initiator may begin forwarding data via the LSP towards the terminator as soon as it accepts the Downstream Label provided by the corresponding Resv message. Including a Suggested Label in the Path message can reduce the time between LSP signaling and actual data forwarding. It allows the initiator and internal nodes to concurrently configure its database and hardware (e.g. aligning the mirrors) for data forwarding while the Path message is traveling downstreams. Interoperability of bidirectional LSP setup is a very basic test, however it is a function that is extremely important to verify, as it allows for more complex features to be built over G-LSPs. The establishment of bidirectional LSPs across the data-plane was verified in three separate scenarios.

Test Case #2. Bidirectional LSP Setup and Teardown (graceful) – Dynamic
In this case, the devices automatically generated an ERO based on CSPF. The LSP was observed to take the most optimal path available.

Test Case #3. Bidirectional LSP Setup and Teardown (graceful) – Strict ERO
The second scenario was LSP setup with strict ERO configuration. In this case, the route was explicitly (manually)configured for each hop, and the path chosen was a less optimal path. The LSP was observed to take the less optimal path as specified by the strict ERO.

Test Case #4. Bidirectional LSP Setup and Teardown (graceful) – Loose ERO
In this case, the route was explicitly configured for certain hops, and the devices calculated the most optimal path to the hop(s) specified by the loose ERO configuration. The LSP was observed to take the most optimal path available to the hop(s) specified by configuration.

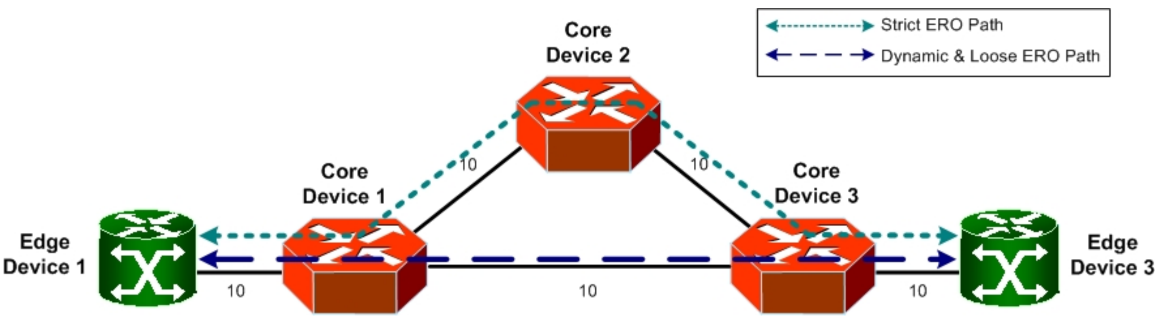

Figure 1: Bidirectional LSP topology

Hierarchical LSPs

In the GMPLS architecture, a higher layer LSP may be nested within a lower layer LSP. Fundamentally, hierarchical LSPs are FA-LSPs – the higher layer LSP traverses nodes that appear to be another TE link, but is in fact a lower layer LSP. An obvious advantage to the deployment of hierarchical LSPs is that multiple LSPs can be grouped into a very high bandwidth pipe (i.e. the lower layer LSP). From the network level perspective, scalability is a merit that hierarchical LSPs offer.

Test Case #5. Hierarchical LSP Setup and Teardown (graceful)
A comparatively lower layer LSP was first established within the core network and advertised as an FA-LSP (essentially a TE link) via OSPF-TE LSAs. After the routing tables synchronized, a comparatively higher layer LSP was then signaled, with an ERO specifying the FA-LSP as a hop, between the edge end points. The hierarchical LSP was established in both the control-plane and data-plane.

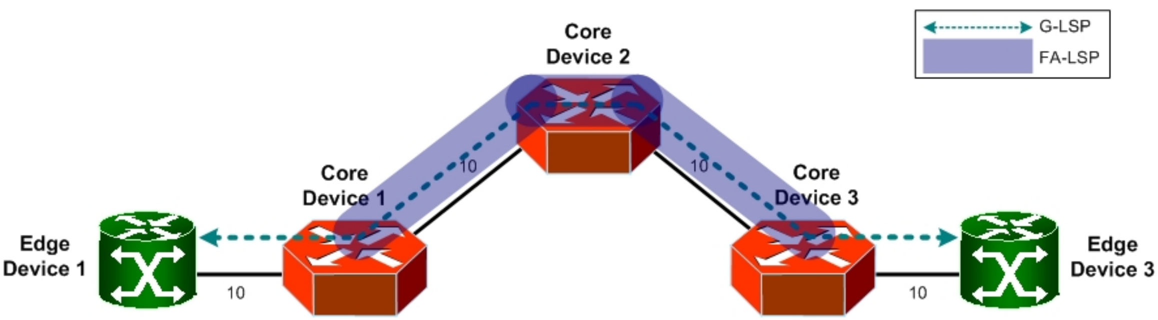

Figure 2: Hierarchical LSP topology

LSP Re-optimization

A large network periodically changes; these changes can be categorized as soft changes – peak and off-peak hours, administrative costs, etc., and hard changes – link failure, new equipment, new fiber, etc. In each case, it is ideal that selected LSPs are able to recalculate the most optimal path (within constraints) based on the latest network topology.

Test Case #6. Single Layer Reoptimization
An LSP was established between the edge devices via the most optimal path based on OSPF costs. Once the LSP was verified to be stable, the OSPF costs were reconfigured such that the original path was no longer the most optimal. The LSP was observed to reroute to the new optimal path. The data-plane was not widely available for this test during the test event, and testing involved only a single layer. In addition to the reconfiguration of OSPF costs which caused the LSP rerouting, the edge devices were observed to be capable of rerouting each LSP to the most optimal path based the volume of data traffic and resource considerations.

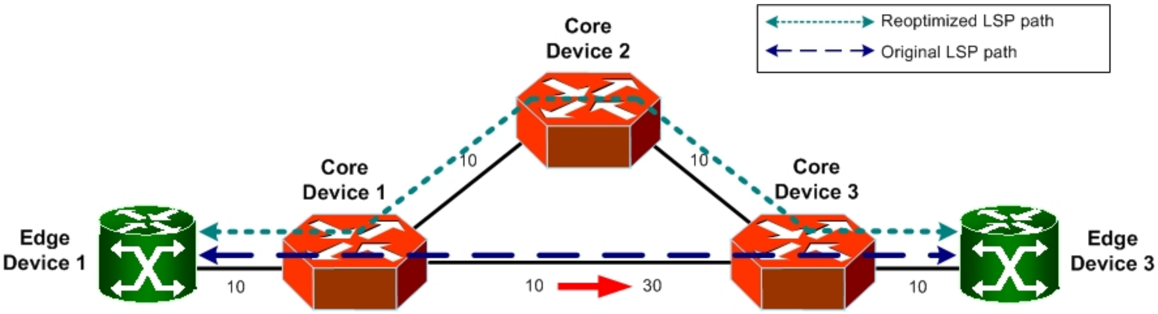

Figure 3: Reoptimization topology

CSPF – Constraint: Protection Type

Link level and path level protection types are proposed in the current GMPLS standards. From a performance perspective, link-level protection offers much quicker restoration than path-level protection. GMPLS protection options are requested during the signaling process, and an LSP can only be setup if the requested protection (or better protection than requested) can be provided. At the time of writing, type 0x02 Unprotected was supported by the majority of vendors, and it was tested to allow for a better understanding of the CSPF interoperability by Protection Type.

Test Case #7. LSP Setup with Protection Object
TE-links were configured and advertised for 0x02 Unprotected type, and the other link was not configured for any link protection type. An LSP was signaled with a Protection Object requesting 0x02 Unprotected type. The LSP was observed to establish via the 0x02 Unprotected link, and data traffic was transported across via the LSP.

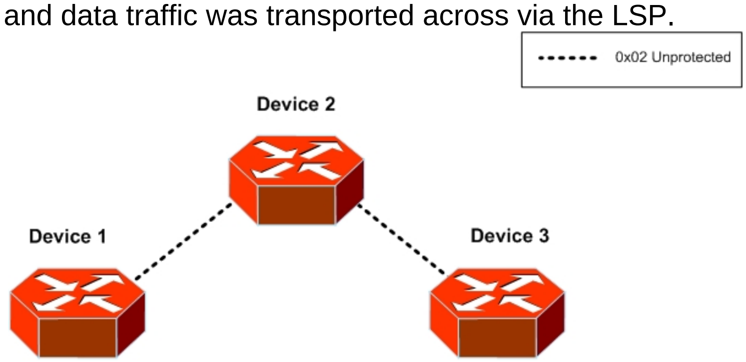

Figure 4: CSPF topology

GMPLS Traffic Engineering

The volume of data traffic supported over service provider networks continues to grow as demand for new data services increases. IP/optical networks, essentially multi-layer service networks, are required to accommodate both existing and future multi-service needs. A profitable IP/optical network must be bandwidth efficient for all traffic types, including deterministic bandwidth utilization as well as random traffic surges. Deterministic usage is easier to plan for and manage, as careful design is generally sufficient. Random usage, on the other hand, is more difficult to efficiently accommodate, as accurate predictions are not possible. Therefore, it is desired, and necessary, to have dynamic multi-layer resource allocation functions in an IP/optical network that can provide resources on demand. In all scenarios, service interruption must be kept to a minimum.

Test Case #8. Bandwidth on Demand

LSP1 with bandwidth X was established between the edge devices across all the core devices. LSP2 and LSP3 were then established, each reserving a bandwidth of 0.2X and 0.8X respectively. In the absence of data-plane, an increase in bandwidth utilization was simulated. When bandwidth utilization exceeded the configured threshold, the edge devices calculated a new path and established a new LSP via the path ‘Edge Device 1’ – ‘Core Device 1’ – ‘Core Device 3’ – ‘Edge Device 2’.

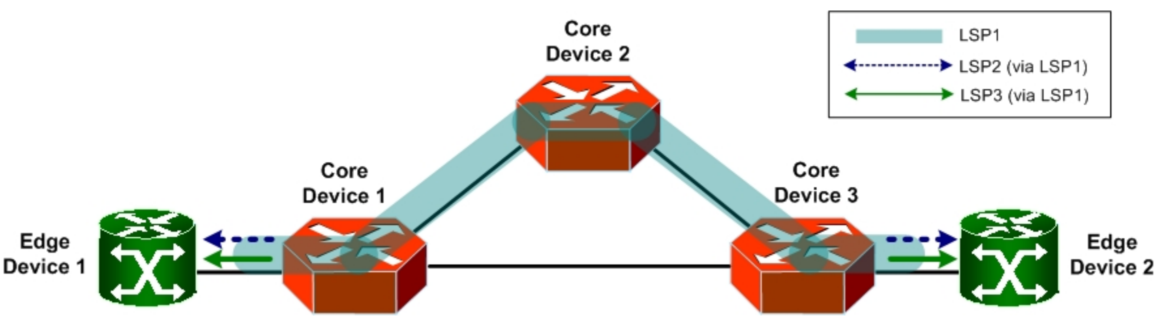

Figure 5: Bandwidth on Demand topology

Test Case #9. Cut Through

LSP1 was established between edge device 1 and edge device 2, and LSP2 was established between edge device 2 and edge device 3 as shown. LSP3 (PSC-LSP) was created via the path ‘Edge Device 1’ – ‘Edge Device 2’ – ‘Edge Device 3’. Each LSP reserved a bandwidth of 0.8X. An increase in bandwidth utilization was simulated. When bandwidth utilization exceeded the configured threshold, the edge devices calculated a new path and established LSP4 via the path ‘Edge Device 1’ – ‘Core Device 1’ – ‘Core Device 3’ – ‘Edge Device 3’, and LSP5 (PSC-LSP) via the LSP4.

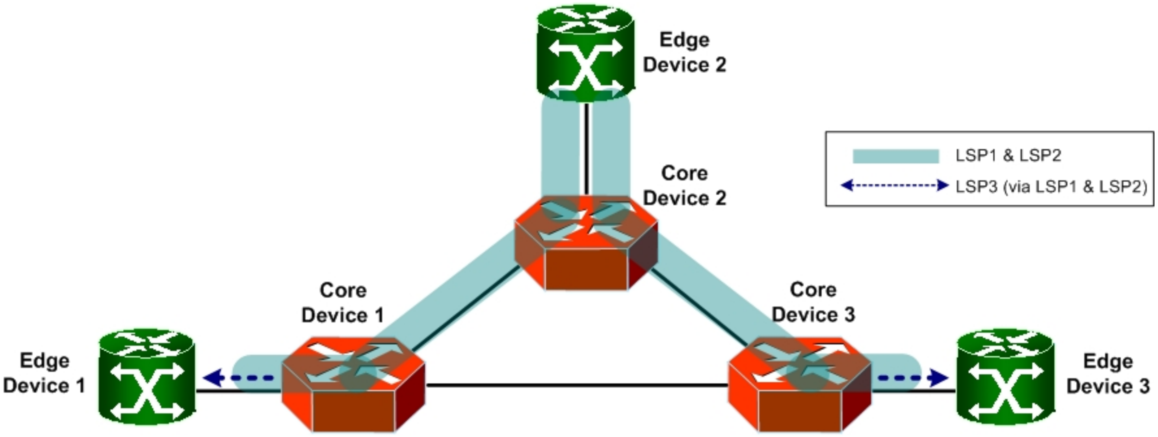

Figure 6: Cut Through topology

Layer 1 VPN

Currently, three Layer 1 VPN models are described by IETF Internet Draft draft-takeda-l1vpn-framework-01. Each model differs in the customer interface type (i.e. management plane vs. control plane) and in the type of information exchanged over the customer interface (e.g. signaling vs. signaling & routing). The Management-based Service Model allows for total control over the SPCs by the optical network provider. In this model, provider management systems provision the Layer 1 connection across the optical network by communicating with the appropriate edge routers or switches to establish and teardown a G-LSP. Once the LSP is setup, the customer end-points may begin communicating with each other across the optical network via the LSP.

Test Case #10. L1VPN Control-Plane Setup

Provider management systems, namely L1VPN Servers, were connected to the PE switches in the control-plane in this test event. The PE switches were configured to advertise the TE links, and the OSPF tables were synchronized, and the L1VPN Servers configured the parameters for the TE links between the appropriate PEs and CEs. The L1VPN Servers then calculated an appropriate path between the CEs and requested a layer 1 connection between the CEs. The Layer 1 connection (G-LSP) was established by the switches as indicated by the ERO, which specifies the hops as well as the outgoing interface of the downstream PE – thus established the SPC.

Test Case #11. Applications over GMPLS

With the SPC created in Test Case #10, both IPv4 and IPv6 packets were transmitted into the optical network and transported across via the G-LSP. Then, a pair of CE routers connected to the PE switches were able to establish an MPLS-LSP across the optical network via the SPC. Video streaming provided by web-enabled Panasonic cameras was tested using both IPv4 and IPv6 addressing.

Stability Testing

To verify the stability of a network containing a variety of GMPLS elements, the GMPLS test network was used as the uplink to the Internet for the test participants. An onsite switch providing Internet access to the participants was connected to one of the two

MPLS-LSP endpoints, and the other endpoint was plugged into the onsite uplink to the Internet (through a local server). The participants were able to acquire DHCP addresses, access web browsers, secure email servers, and other Internet applications through the provisioned SPC. No loss of packets or interruption of services was observed during this test. When the L1VPN Servers removed the SPC, all online services were lost.

Discoveries for Further Investigation

Advertising Transport Links

The IETF Drafts draft-ietf-ccamp-gmpls-routing-09 and draft-ietf-ccamp-ospf-gmpls-extensions-12 describe ways of advertising resource information in terms of bandwidth. While this advertisement scheme is adequate for PSC devices, the scheme does not work for LSC devices, because an LSC device makes reservations at the wavelength level, rather than at the bandwidth level. If an LSC LSP is setup with a particular wavelength, availability of bandwidth alone does not guarantee successful LSP setup because bandwidth may be available on some other wavelength, but not on the requested wavelength. To perform a successful LSC LSP path computation satisfying wavelength continuity constraint, information about available wavelengths on the TE links should be advertised.

The concept is better illustrated with the following example:

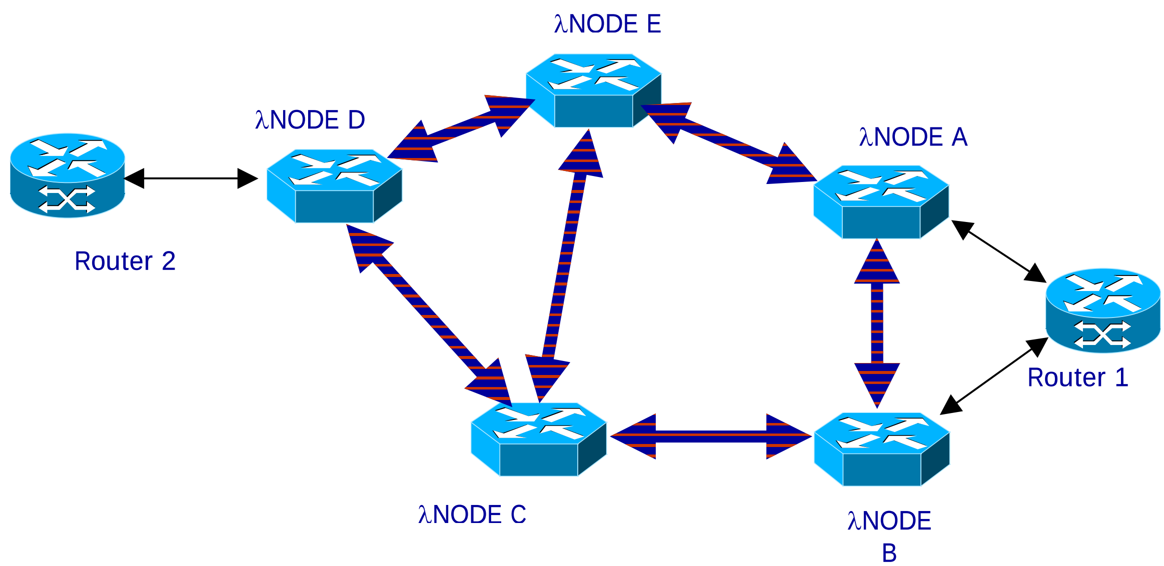

Figure 7: Abstract topology

In this scenario, there are two 2-hop paths from Router 1 to Router 2 through the λNodes – A-E-D or through B-C-D. Suppose Router 1 choose path A-E-D, but the wavelength constraints cannot be satisfied on the link A-E, then node A may choose a longer path A-B-C-D or may perform wavelength conversion or may force crank back to Router 1 to try another path. If Router 1 can make use of wavelength as a constraint when computing the LSC LSP, it can choose the best path through the LSC network.

Thus when computing LSPs, different resource constraints must be employed by the nodes computing the path for the LSP depending on the type of LSP being signaled. Following table summarizes the type of resource constraints that should be used for each LSP type:

LSP Type	**Resource Constraint**
PSC	Bandwidth
LSC	Wavelength
TDM	Timeslot
FSC	Port availability

When advertising an LSC Capable TE link, a node should include information about wavelengths available on that TE link. Each available wavelength should be individually listed in the TE link in an available Wavelength sub-TLV. The following structure for this TLV is proposed:

0	31
Type	Length = 4N (N = number of wavelengths available)
Wavelength 1 in floating point format	
Wavelength 2 in floating point format	
….	
Wavelength N in floating point format	

List of Issues Encountered

Problem Area	Problem in General	Problem in Test Event	Solution in Test Event	Proposed Solution
IP	TTL value for multicast OSPF must be “1”. However, TTL>1 is required in the case that a control-plane adjacency is multiple hops away. This is typical in GMPLS as the control-plane topology is not necessarily identical to the data-plane topology.	Some vendors implement an additional “virtual” hop, so packets with TTL = 1 was dropped.	TTL was configured to a value greater than 2.	When a multicast address is used, the TTL value should be kept at 1 as specified by RFC2328 since it is already widely implemented and deployed. In this case, tunneling (e.g. GRE) can be used to ensure the multicast OSPF packet reaches its destination. When a unicast address is used, the TTL value should be configurable. This scenario applies to all types of OSPF that uses unicast address – point-to-point, point-to-multipoint, etc.
Control Plane	Dissimilar control-plane configurations that do not interoperate are an impediment to forming control-plane adjacencies.	Some devices supported tunneling and others supported plain IP. Within tunneling, some supported GRE and others supported IP-	Some vendors implemented GRE, IP-in-IP, or some forms of plain IP during the test event.	Plain IP is simple, scalable, and allows for dynamic discovery of adjacencies. Tunneling has its purposes and applications, and prevents the intermediate hops that are irrelevant to GMPLS from interacting with the routing tables of the OSPF adjacencies in the

		in-IP.		GMPLS domain. Both tunneling and plain IP are likely to be used in a control-plane network. However, to ensure interoperability, there need to be at least one common implementation among all control-plane adjacencies. It is recommended that all three configurations (i.e. GRE, IP-in-IP, and plain IP) SHOULD be supported to allow for flexibility, and two of the three configurations MUST be supported.
OSPF	Dissimilar OSPF types that do not interoperate (e.g. point-to-point and multicasting) becomes an impediment to forming control-plane adjacencies.	Devices supported a combination of point-to-point, point-to-multipoint, and multicasting. Some devices do not have a common OSPF implementation.	Some vendors implemented other OSPF types during the test event. In some cases, routes were indirectly learned through a third neighbor.	Since each of the OSPF types has its purposes and applications, and the type of OSPF used depends on the design and requirements of the network itself. A recommendation to select a default is not possible or practical. However, to allow for maximum flexibility, all OSPF types SHOULD be implemented.
RSVP-TE	Some devices accept ERO subobjects that specify incoming interface addresses, and some devices only accept those that specify outgoing interface addresses.	Same as description in “Problem in General”.	Some vendors made changes to implementation so that both interface addressing are accepted, and some made changes so that both types can be generated in a Path message.	For an ERO, the incoming interface should always be used as the default. Identifying a hop by the incoming interface is desirable because it allows for no ambiguity in the intended explicit path. This is especially true in the case of strict ERO, in which multiple paths may exist to the next strict hop if the hops are identified by the outgoing interfaces. For loose ERO, the use of incoming interfaces is also recommended as the default. With these rules, the RRO in a Resv message can always be assumed to be a set of incoming interfaces from the perspective of the direction of the Path message – that is, the sub-objects in an RRO in a Resv message are the same as those in an ERO in a Path message.
GMPLS Architecture	Some devices support numbered interfaces only, and some devices	Same as description in “Problem in General”	Some vendors implemented numbered or unnumbered	Unnumbered interfaces are often implemented and preferred by lower layer switches that do not switch

	support unnumbered interfaces only.		interfaces during the test event so they can be data-plane adjacency.	data by the IP header because 1) numbered interfaces are usually required for routers, and 2) numbered interfaces require more resources as the identifiers (i.e. IP addresses) are meaningful. PSCs typically support numbered interfaces, and that they are most likely to be connected to both IP routers and lower layer switches. Therefore it is recommended that PSCs support both numbered and unnumbered links.
OSPF	Values for OSPF HelloInterval and RouterDeadInterval are not specified for GMPLS, and there are no default values recommended.	Vendors implemented different interval values by default, resulting in loss of OSPF adjacency periodically.	Vendors implemented and configured common values.	The OSPF timers MUST be configurable.
OSPF-TE	The amount of information present in TE-link advertisements is insufficient for TDM and LSC links. For TDM switches, the number of timeslots available for an associated TE-Link is important information. For LSC switches, how many wavelengths are available for an associated TE-Link is important information.	Some TDM platforms were unable to synchronize their database because there was difference in how the TE-link information is advertised and parsed, and a reference RFC or draft is not available to address this.	No solution identified. The tested devices tried to setup the LSP as usual. When the LSP setup was successful, it was implied that timeslots or wavelengths were available. When the LSP setup failed, it was indicated (by messages) that timeslots or wavelengths were not available.	For TDM TE-links, available timeslots should be advertised. For LSC TE-links, available wavelengths should be advertised. Bandwidth information should also be advertised for each. Encoding and decoding of these information must be precisely defined to avoid discrepancies in interpretation.

Conclusion

To remain competitive, service providers must control the cost of network operation and provisioning while at the same time providing more efficient transport and new value-added services. Service providers are evaluating a wide variety of OSRM technologies to accomplish this. Both standard-based and vendor proprietary OSRM solutions are available, however, it is not clear how multiple vendor proprietary implementations will interoperate in operational networks. GMPLS and other standard-based OSRM approaches (ITU-T, OIF) promise to overcome interoperability issues, but to do so, it crucial that multi-vendorOSRM implementations are tested for standards-compliance.

Organizations such as the IETF's CCAMP WG are addressing the standards for a unified optical control plane, validation of these efforts is crucial to the ongoing process of protocol standardization and commercial adoption. Assurance of interoperability based on these standards is indispensable to service providers that deploy network elements from multiple vendors. Unbiased, cooperative interoperability tests provide both service providers and equipment vendors with a neutral forum to evaluate standards based interoperability. The results of these events help vendors refine and validate their OSRM implementations, which ultimately promotes the future acceptance of their products and services into service provider networks. Assurance of interoperability also reduces the risk and complexity of deploying multi-layer networks while concurrently reducing the operational expenses associated with provisioning, service activation and re-allocation.

The UNH-IOL provides an aggressive operative test scenario tailored to stringent service providers' demands. A comprehensive series of MPLS and GMPLS tests has been completed, including initial investigations into Layer 1 VPN functionality. Layer 1 VPN evaluation takes OSRM validation beyond the basic LSP failure/recovery features essential for commercial adoption, and moves it firmly into the area of new services and associated revenue generation – a strong step in the direction of realizing the benefits of GMPLS from a business perspective. Commercial adoption, driven by demonstrated standards compliance and interoperability, will require additional validation. The UNH-IOL will continue to foster this validation process in answer to industry demand.

Recommendations for Further Investigation

As a result of the OSRM test methodologies and findings described above, several facets of OSRM technology emerged as compelling candidates for further testing. Among these are the following:

- LSP formation with multiple switching capabilities: Testing with multiple nodes that have different switching capabilities.
- Hierarchical LSPs, with multiple switching capabilities
- Additional Layer 1 VPN network models
- Scalability Testing, involving increased number of nodes, TE links, LSPs and setup/teardown patterns.
- Additional protection and restoration scenarios
- GMPLS and UNI and NNI interworking
- Link Management Protocol (LMP)

The UNH-IOL looks forward to working with service providers and all participants in the OSRM test events to further investigate these and other aspects that are equally important to realizing validation in complex operational networks.

Glossary

Abbreviation	**Definition**
FSC	**F**iber **S**witch **C**apable device. An optical cross connect that switches the contents of a whole fiber to another fiber.
GMPLS	**G**eneralized **M**ulti**P**rotocol **L**abel **S**witching. An architecture that extends the MPLS architecture defined by RFC3031 to support transport links.
L1VPN	**L**ayer **1** **V**irtual **P**rivate **N**etwork. A service that provides an optical path end-to-end that is virtually isolated from other users.
LSC	**L**ambda **S**witch **C**apable device. An optical cross connect that switches incoming data based on wavelengths.
LSP	**L**abel **S**witched **P**ath. A virtual tunnel across an MPLS domain that switches data by labels in the Shim header.
MPLS	**M**ulti**P**rotocol **L**abel **S**witching. An architecture that introduces the concept of forwarding packets across a network by labels instead of routing.
NNI	**N**etwork-to-**N**etwork **I**nterface. The interface between two network nodes.
OSPF	**O**pen **S**hortest **P**ath **F**irst. A routing protocol that calculates the least OSPF-cost path to all network points within a routing area.
OSPF-TE	**T**raffic **E**ngineering Extensions to **O**pen **S**hortest **P**ath **F**irst.
PSC	**P**acket **S**witch **C**apable device. A device that switches incoming data based on the packet header.
RSVP-TE	Resource **R**e**S**er**V**ation **P**rotocol – **T**raffic **E**xtensions.
SPC	**S**oft **P**ermanent **C**onnection. An virtual dedicated path that provides a layer 1 connection between users across an optical network.
TDM	**T**ime **D**ivision **M**ultiplexed device. A device that switches incoming data based on a specific slot in time.
TLV	**T**ype/**L**ength/**V**alue. A message format in which OSPF elements are specified.
UNI	**U**ser-to-**N**etwork **I**nterface. The interface between an edge node and a network node.

References

Request for Comments 2205 – Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification

Request for Comments 2328 – Open Shortest Path First (OSPF) Version 2

Request for Comments 3031 – Multiprotocol Label Switching Architecture

Request for Comments 3032 – MPLS Label Stack Encoding

Request for Comments 3209 – RSVP-TE: Extensions to RSVP for LSP Tunnels

Request for Comments 3471 – Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description

Request for Comments 3473 – Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions

Request for Comments 3477 – Signalling Unnumbered Links in Resource ReSerVation Protocol – Traffic Engineering (RSVP-TE)

Request for Comments 3945 – Generalized Multi-Protocol Label Switching (GMPLS) Architecture

Internet Draft draft-ietf-ccamp-gmpls-recovery-e2e-signaling-01.txt – RSVP-TE Extensions in support of End-to-End GMPLS-based Recovery

Internet Draft draft-ietf-ccamp-gmpls-routing-09.txt – Routing Extensions in Support of Generalized MPLS

Internet Draft draft-ietf-ccamp-gmpls-recovery-terminology-03.txt – Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)

Internet Draft draft-ietf-ccamp-ospf-gmpls-extensions-12.txt – OSPF Extensions in Support of Generalized Multi-Protocol Label Switching

Internet Draft draft-ietf-mpls-lsp-hierarchy-08.txt – LSP Hierarchy with Generalized MPLS TE

Contributors

The whitepaper content is an accumulation of agreements, input, and comments from all participants. Takumi Ohba, Kaori Shimizu and Yumiko Kawashima from NTT Network Service Systems Laboratories and Henry He from UNH-IOL MPLS Services Consortium wrote the initial contents. Special thanks to Tom DiMicelli from Sycamore Networks and Chris Volpe from UNH-IOL for their work especially on the executive summary, introduction, and conclusion sections. Much appreciation to Rajeev Veettil and David Walters from Lambda Optical Systems for their documentation and proposal on wavelength advertisement for LSC TE links.