GMPLS Operations and Management: Today's Challenges and Solutions for Tomorrow

Thomas D. Nadeau and Hari Rakotoranto, Cisco Systems, Inc.

ABSTRACT

This article examines the technical challenges involved in the delivery of GMPLS, and in particular the challenges in managing these networks. This article focuses first on requirements gathered from customers who are deploying GMPLS in their networks, or will be doing so in the near future. The motivations for these requirements are often based on a service provider's next-generation network designs that are motivated primarily by network convergence, and the promised operational and capital expense reductions and architectural simplifications over existing segregated network designs. This article then considers various solution options to manage GMPLS networks based on these requirements. The key areas examined include provisioning/configuration, performance monitoring and operations and management, and tools that can be used for these purposes. We conclude with a recap of standardization issues and opportunities.

TODAY'S NETWORKS AND MOVING BEYOND THEM

The majority of today's service provider core networks are based on multiprotocol label switching (MPLS). This includes those networks that carry data, video, and voice traffic. Those that do not still carry some portion of services such as leased line service (frame relay or asynchronous transfer mode, ATM) or public switched telephone network (PSTN) services on separate dedicated networks for these purposes. This latter case is driving the convergence of network services onto MPLS and possibly generalized MPLS (GMPLS) networks. The advantages of such a move are clear: single capital expenditures for equipment, as well as single operations and management systems and associated costs (i.e., minimizing the number of network management systems that are maintained and operated). Furthermore, a reduction in complexity of the network implies a move toward simplified operations, maintenance, and management.

In this article we look at and discuss some of issues operators of large networks are facing today, especially in terms of recurring operational expenditures, when considering the convergence of existing networks onto MPLS and GMPLS networks. Since much work has been focused in the past on the operations and management of MPLS networks, we focus on discussing the current challenges and possible solutions for GMPLS networks. To this end, we begin with a taxonomy of the various standard types of GMPLS configurations. Then we discuss what tools and techniques are available to manage these networks today. Next, we discuss why these tools are inadequate and do not lend themselves to a comprehensive solution. Finally, we suggest some tools the various standards fora might tackle in their ongoing work to provide a more comprehensive standards-based management solution for GMPLS networks going forward.

GMPLS ARCHITECTURE

Let us first begin by giving a brief overview of the GMPLS architecture.

The idea is to define a common set of control functions and interconnection mechanisms that allow unified communication, routing, and control across disparate types of underlying transport technologies, such as IP, ATM, synchronous optical network/synchronous digital hierarchy (SONET/SDH), and dense wavelength-division multiplexing (DWDM). Traditionally, each specific technology has its own control protocols, and as a result each set of control protocols do not communicate directly with each other on a peer-to-peer level. Instead, networks are layered one on top of the other, creating overlays at each layer to collectively provide end-user services. Obviously, this process requires knowledge of each technology domain, provisioning of each layer, and separate management of per-domain operations functions.

GMPLS refers to a set of protocols that will provide interoperable end-to-end provisioning of

 0163-6804/05/$20.00

optical networks as well as other devices. The protocols consist of:

- Generalized Resource Reservation Protocol with traffic engineering (RSVP-TE) for signaling
- Internal Gateway Protocol (IGP) with TE extensions for routing (Open Shortest Path First, OSPF, or Intermediate System to Intermediate System, ISIS)
- Link Management Protocol (LMP) and LMP-DWDM for link management and discovery functions

To determine the appropriate control plane architecture, two efforts have been carried out simultaneously by both the Internet Engineering Task Force (IETF) and International Telecommunication Union (ITU). The IETF has embraced extending MPLS for integrating data and optical network technologies. MPLS provides an attractive foundation for the optical control plane architecture because MPLS has a natural separation between its data and control planes. Hence, the IETF has extended the MPLS label-switching concept to include other types of forwarding planes. For example, if we extend the definition of a label, MPLS can be applied to wavelengths (lambdas). The extended MPLS protocols, which are considered a superset of MPLS, are called GMPLS for generalized MPLS.

One of the main architectural enhancements introduced as part of GMPLS is the decoupling or separation of the control plane from the data plane. The advantage of this is that it allows for complete out-of-band control of a GMPLS network. Some of the reasoning behind this approach is that it allows GMPLS LSRs to control optical networks that currently exist without having to retrofit too much of the existing network. While the technology used by the control plane remains based on the same guiding principles of the Internet Protocol (IP), the data plane can now include a much wider variety of traffic such as fiber switch capable (FSC), lambda switch capable (LSC), packet switch capable (PSC), and time-division multiplex (TDM).

GMPLS enables seamless interconnection across different networking technologies as well as the capability to perform end-to-end "one touch" provisioning across a heterogeneous network. Such a network configuration is illustrated in Fig. 1, in whichthe various types of networks are shown together: packet and optical transport. The figure shows how LSPs can be established across an optical transport. The LSRs on the left of the figure establish (i.e., signal) TE tunnel LSPs using RSVP-TE that begin in the packet network, traverse the optical core, and then culminate at the remote LSR situated at the edge of the optical and packet networks. Traffic from the packet networks can then be routed into these tunnels transparently to the packet networks. It should be noted that the main constraint of this configuration is the requirement that a label switched path (LSP) end on a network node of the same type. More precisely, the GMPLS LSPs both begin and end on a GMPLS label switching router (LSR). This does have the advantage, however, that provisioning of the LSP can be done much in the same manner in which it is performed in an MPLS network. Also, the LSP can later be managed using similar tools and techniques used to manage MPLS TE tunnel LSPs. Although errors that occur within the optical transport still need to be handled by systems designed to operate and manage that network, it is possible to integrate these systems into those used to manage the MPLS network, thus providing a seamless management system.

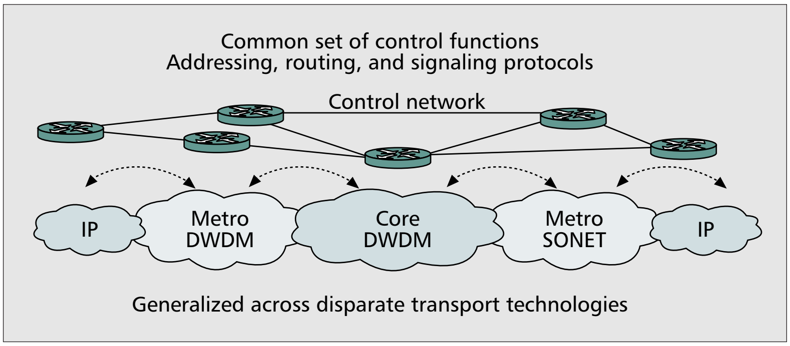

Figure 1. *Generalized/unified control plane.*

At the same time, in response to a demand from ITU members to specify an automated switched transport network, the ITU took on the task to specify a reference architecture for optical networks known as automatic switched optical network (ASON). ASON defines the different components in an optical control plane and the interaction between those components.

The purpose of the ASON control plane (as per ITU-T G.8080) is to:

- Facilitate fast and efficient configuration of connections within a transport layer network to support both switched and soft permanent connections
- Reconfigure or modify connections that support calls previously set up
- Perform a restoration function

ASON provides support for both switched connection (as in setup/teardown initiated by a user request, typically a client) as well as soft permanent connection (initiated via a management/operations support system, OSS). Fault management as well as connection state information are also provided.

ASON also integrates the notion of domain. The control plane might be subdivided in domains that, for example, match the administrative domains of the networks. ASON specifies the possible interaction between different domains in terms of routing, exchange of information, and so on. The points of interaction between different domains are known as *reference points*. The reference point between an administrative domain and an end user is the user–network interface (UNI). The reference point between domains is the external network–network interface (E-NNI). The reference point within a domain between routing areas and, where required, between sets of control components within routing areas is the internal NNI (I-NNI).

As the ITU does not define any protocol specifications, the Optical Internetworking Forum (OIF) has extended several GMPLS components and defined a set of UNI and NNI protocols to address ASON requirements. It should be noted

The network management functions provided within a GMPLS LSR should extend those provided in an MPLS LSR, thereby extending the operational investment in managing MPLS networks.

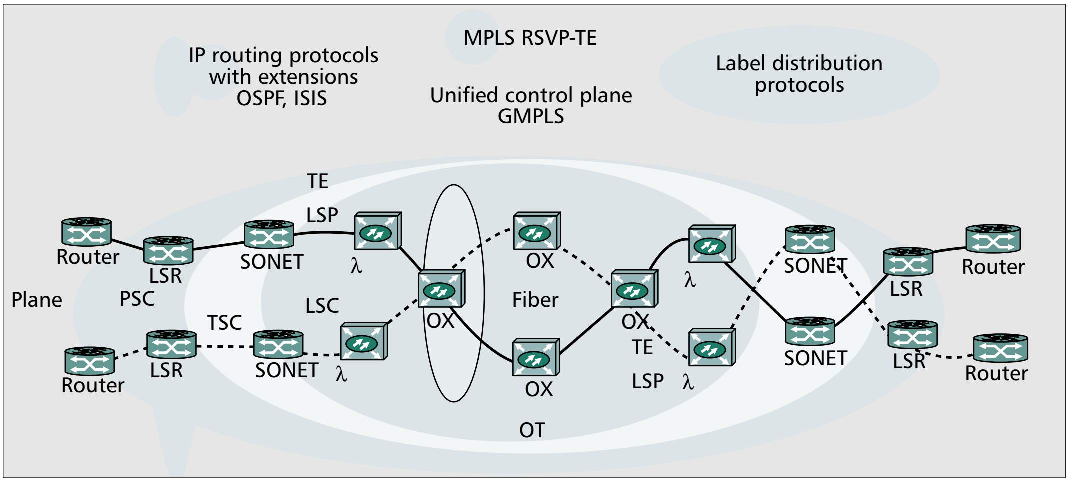

■ **Figure 2.** *GMPLS architecture.*

that these specifications have also been fed back into IETF contributions, including draft-ietf-ccamp-gmpls-ason-reqts-07.txt, draft-ietf-ccamp-gmpls-rsvp-te-ason-03.txt, and RFC 3473.

The protocols are known as optical UNI (O-UNI) and NNI. The client side device runs O-UNI-C protocols, and the network side device runs O-UNI-N protocols (O-UNI provides a user-to-network bidirectional signaling interface between the service requester and service provider control plane entry point and does not share routing information across these domains.). Inter-/intradomain interactions are handled via the I-NNI and E-NNI.

It should also be mentioned for completeness that the approach of extending the MPLS control plane for GMPLS also implies that both may coexist at the same time. That is, a GMPLS LSR can simultaneously support both GMPLS and MPLS LSPs. Furthermore, the network management functions provided within a GMPLS LSR should extend those provided in an MPLS LSR, thereby extending the operational investment in managing MPLS networks. We see that this is indeed the case (Table 1).

GMPLS Network Models

Given the general architecture described above, two network models have been proposed for operation of this architecture. The first is referred to as the *overlay model*, and is illustrated in Fig. 2. This approach assumes that the optical network topology and its various characteristics will not be exposed to the edge devices, or to any devices in different administrative areas or domains. This model also assumes separate routing domains. The reasons for deployment of such networks vary, but one reason for their use is regulatory rules that require their separation, meaning different departments within a single network operator must not allow others to modify (i.e., provision) the devices controlled by that department. This approach allows them to maintain this requirement while still garnering the benefits of GMPLS. Figure 2 illustrates how a GMPLS LSR can compute a path across the optical network while the edge router attached to it has no visibility of the path details other than that a path exists toward the destination advertised (i.e., the far LSR in the

Requirements	Specified in RFC 3474/3476	Proposal
Soft permanent connection	Yes (SPC label)	Yes (RFC 3473)
E2e capability negotiation	No	Yes
Call without connection setup	No	Yes
Call with (single) connection setup	Yes (limited to single-hop sessions)	Yes
Multiple connections per call (add/remove)	No	Yes
Call segments	No	Yes
Restart (CP failures)	Limited	Yes
Crankback signaling	No	Ongoing
Backward capability	No	Yes

■ **Table 1.** *Some of the enhancements required for GMPLS protocols to support ASON.*

IETF GMPLS overlay	ITU-T/OIF UNI
Shared GMPLS address space	Fully decoupled address spaces
If full ERO provided by client, infers visibility into topology and reachability or shared route server	Topology and reachability hidden from clients
Single signaling session across all domains (end-to-end connection)	UNI signaling session independent of interdomain signaling (connection segments supported)

Table 2. *IETF overlay vs. ITU-T/OIF.*

The Peer model is suitable whenever the transfer of full routing information is required. As such it is mainly relevant for intra-domain applications. There is no assumption made in this model as far as policy and security at the network inter-connection boundaries is concerned.

figure) and is reachable via some routing metric value. (An edge node [router] will just send a request for a path toward the destination; the GMPLS LSR [on the edge of the optical network] will determine the best path to the far LSR where the destination router is connected.) ASON architecture relies on the overlay model.

The second model, referred to as the *peer model*, provides an alternative GMPLS deployment configuration. In this model all the devices participating in the GMPLS topology share the same network topology information. The model is an extension of the IETF MPLS TE protocol (routing, RSVP extension to support TDM, wavelength, fiber, etc.). A Link Management Protocol (LMP) for neighbor configuration and discovery is also specified, which allows neighbors to discover each other and their link configurations.

In this network model the optical-electrical-optical LSR takes full part in routing and path selection. The edge node has full control over the path through the optical network.

In a nutshell, the peer model is suitable whenever the transfer of full routing information is required. As such, it is mainly relevant for intradomain applications. There is no assumption made in this model as far as policy and security at the network interconnection boundaries is concerned. The overlay model, on the other hand, is mainly geared toward user-provider scenarios, where policy and security are well defined. In particular, the overlay model is generally used where specific policies are defined as a means to allow a specific domain to not disclose its topology.

It should also be noted for the sake of completeness that the IETF has also proposed an overlay model mainly derived from the peer-to-peer model, where routing information from the optical network to the edge is filtered in order to hide it. Table 2 summarizes some of the differences between both overlay model approaches.

CHALLENGES IN MANAGING GMPLS

Now that we have introduced and discussed the various types of network configurations for GMPLS, we first discuss and describe the tools that are currently available from various standards fora for the management of GMPLS networks. Next, we describe why these tools provide an inadequate solution for managing GMPLS networks. Finally, we propose some solutions that may rectify the situation.

TODAY'S TOOLS AND TECHNIQUES

We now move on to the tools and techniques that exist today and can be applied to GMPLS networks. These tools cannot be applied in every one of the deployment configurations described above. For simplicity, we will note where they do not apply.

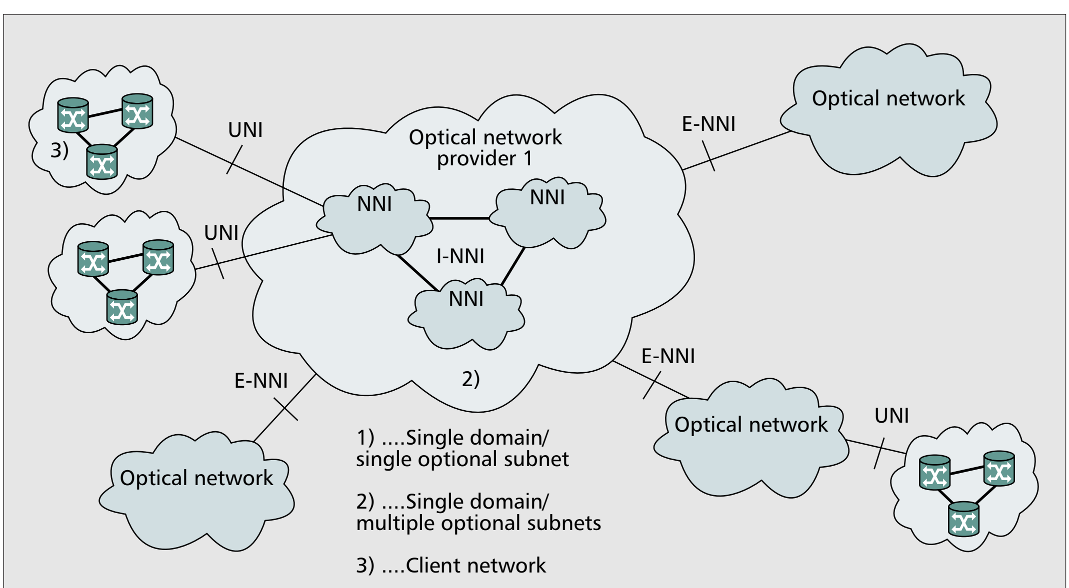

Figure 3. *ASON reference points.*

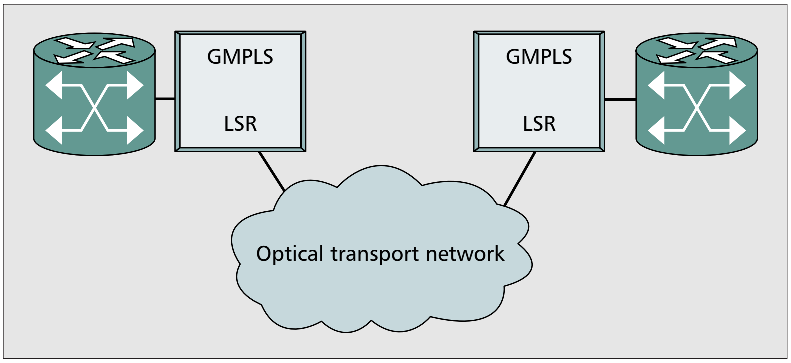

■ **Figure 4.** *Overlay network model.*

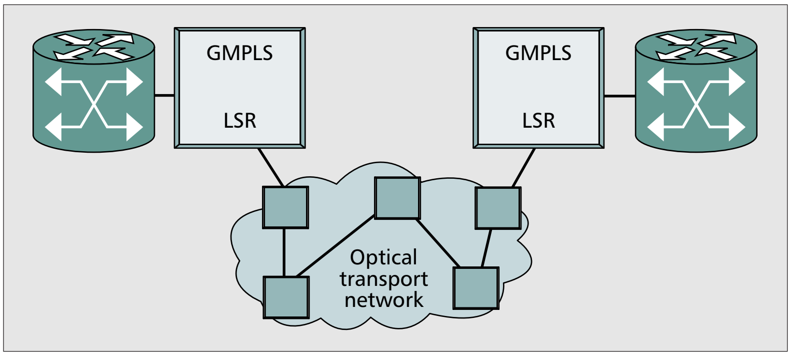

■ **Figure 5.** *Peer-to-peer network model.*

Since GMPLS is built on the foundation of MPLS, we first need to consider what can be applied from the management of MPLS networks. First and foremost, a comprehensive set of Simple Network Management Protocol (SNMP) management information base (MIB) modules has been defined by the IETF for MPLS that are useful for monitoring and configuring the standard elements of an MPLS deployment. These are defined in RFCs 3811, 3812, 3813, and 3815. Since these modules were designed to work together, a network manager is able to utilize pieces from one module in another. RFCs 3811 and 3813 contain modules that allow a manager to view the performance and configuration of basic MPLS label switching, label imposition and disposition, as well as some other basic attributes present in all MPLS LSRs. RFC 3815 provides similar modules for the management of the Label Distribution Protocol (LDP). Finally, RFC 3812 provides a module used to manage the RSVP-TE configuration of an LSR. One interesting thing to note is that although it shows a comprehensive set of configuration details for all TE tunnels on an LSR, it also refers to the signaled LSPs represented in RFC 3814 by providing "pointers" to them. In this way, RFC 3812 leverages the label switching infrastructure provided in RFC 3814 without having to redefine it. RFC 3812 is also particularly interesting for LSRs that support GMPLS, as a set of SNMP MIB modules has been drafted by the IETF's CCAMP Working Group that extend and enhance RFC 3812 to support GMPLS. This includes the CCAMP-GMPLS-LSR-MIB, CCAMP-GMPLS-TC-MIB and CCAMP-GMPLS-TE-MIB. In addition to these, which extend existing MPLS MIB modules, the CCAMP Working Group has also defined the CCAMP-LMP-MIB, which is designed to manage the LMP.

The advantage to this approach is that all of the existing management infrastructure provided by an MPLS LSR can be leveraged for GMPLS deployments. Furthermore, management applications such as Cisco's Internet Solution Center® can be extended incrementally to support GMPLS. Service providers that have deployed this application and trained their staff to operate it will also only have an incremental amount of training required to operate GMPLS when deployed.

A number of other MIB modules can be leveraged to manage GMPLS LSRs, such as the interfaces MIB (RFC 2233) because GMPLS is implemented as an extension to MPLS. However, other MIB modules are required that are unrelated to MPLS, such as those required to manage SONET/SDH.

Beyond the SNMP MIB modules devised for the management of specific MPLS and GMPLS features, other tools are provided for managing the MPLS portion of the network, such as *LSP ping* and *traceroute*, which may be useful for managing the LSRs in a GMPLS network. This tool is designed to test the data plane forwarding of an MPLS LSP by providing an "echo" and a tracing function. In the case of the former, most functionalityfound in a packet-based network should be retained. The only exception would be if a node within the optical core failed. In this case, since GMPLS views the entire optical core as one network hop, the failure point would only be registered as the last GMPLS LSR. The same applies to the tracing function of LSP traceroute. In this case, the trace would not return any hops within the optical core, just the adjacent GMPLS LSRs. Still, this functionality is useful for managing the GMPLS network and can be integrated into the management system, which knows the specifics of the optical core for further levels of troubleshooting and monitoring.

In addition to the tools devised by the IETF that are useful for managing GMPLS, the TeleManagement Forum (TMF) has defined TMF814 that provides an XML data model useful for provisioning and configuration, inventory, provisioning, and service activation. The TMF814 specification is based on a network model that uses the G.805 functional model of transport networks as a premise. This data model is in fact deployed in some networks today.

The ITU-T has defined the ASON Management Framework that can be used as a basis for the definition of new management tools for the overlay model described above, in addition to peer interfaces such as NNIs. The ITU has defined G.7718 that specifies a management framework for ASON control planes as specified in ITU Recommendation G.8080. ITU G.8080 describes functional components of the control plane, including abstract interfaces and primitives, and the interactions between a caller and these components. This includes interactions among these different pieces of the architecture, including how they apply during connection establishment, maintenance, and teardown. The ASON framework describes these interactions in

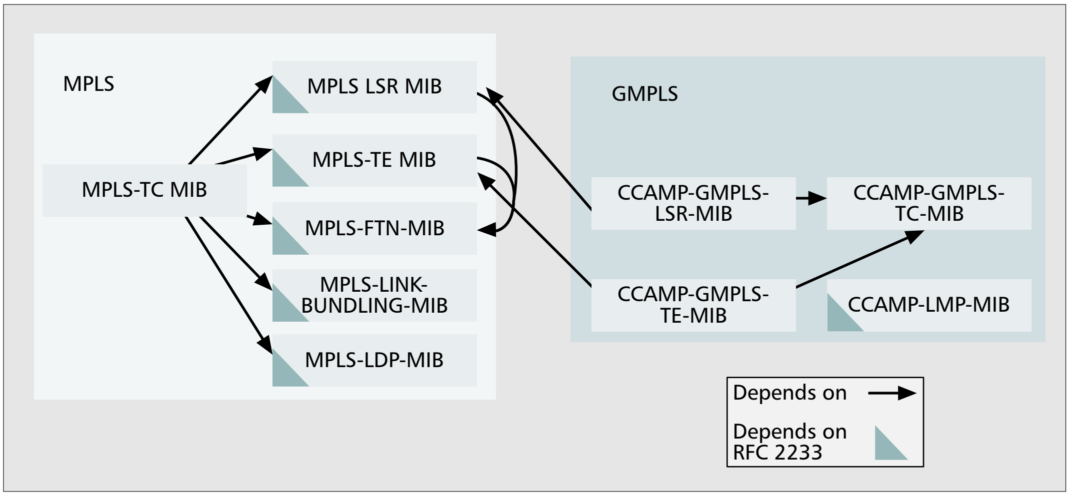

Figure 6. *SNMP MIB modules for managing GMPLS.*

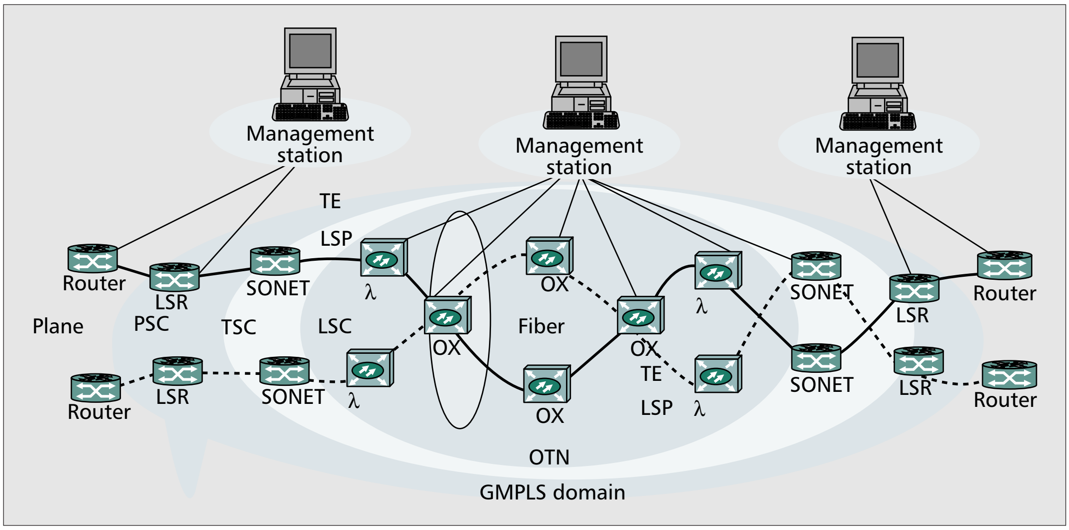

Figure 7. *Managing GMPLS networks with multiple management systems.*

two capacities: first, how they apply to an enterprise messaging server (EMS) that might speak directly to devices (and their control planes) in order to modify or query their management interfaces; and second, how an EMS might contact a number of EMS components in order to facilitate service-level management. The ASON framework defines a number of key components that are designated to perform key network management functions including fault, configuration, authentication, provisioning, and security. Additionally, the framework provides requirements for how each of these components should function both independently and in relation to the other components of the framework.

Finally, the most common interface used to manage optical core devices today is still the outdated TL-1 interface. Although the TL-1 interface may undergo OSMINE certification by Telcordia and thus may be used by third-party OSSs to facilitate multivendor interoperability, the TL-1 interface is still essentially an unstructured, nonstandard interface much like the proprietary command line interface (CLI) provided by most device vendors. Another important point about TL-1 is that despite the possibility of providing a management interface suitable for managing some of one vendor's devices, it does not provide a common interface that is typically supported by all devices in the network.

Although the TL-1 interface may undergo OSMINE certification by Telcordia, the TL-1 interface is still essentially an unstructured, nonstandard interface much like the proprietary command line interface (CLI) provided by most device vendors.

Why Are the Existing Tools Inadequate?

The previous section identified and discussed a variety of standards-based tools available for the management of GMPLS networks. However, the aforementioned tools, neither in whole nor in part, provide a comprehensive solution for managing GMPLS networks. Let us discuss why.

In the section above we discuss the variety of SNMP MIB modules and XML interfaces that are available to manage MPLS and GMPLS networks. While this collection provides a way of managing the GMPLS LSRs, they are typically implemented in a read-only fashion, leaving the provisioning to the proprietary TL-1 interface. These interfaces also cannot provide any insight into the configuration/status of the nodes that lie within the optical core. Even where SNMP MIB modules and standards-based XML models are

Vendors should begin migrating away from antiquated TL-1 interfaces toward relatively modern interfaces such as SNMP or XML. In general, we recommend that GMPLS networks migrate toward ultimately being managed as MPLS networks are managed today.

supported on the optical core devices, access to those nodes is restricted from the GMPLS LSRs and requires an out-of-band management solution comprising a dedicated subnetwork connecting all of the optical nodes to the management station. This network is somehow interconnected with the one connecting the GMPLS LSRs together. This in essence defeats the purpose of having a single unified network. We illustrate this configuration in Fig. 4. This also results in greater difficulty tracking down and diagnosing failures or errors because multiple networks need to be consulted when a failure occurs. Similarly, the standard network configurations discussed also support a variety of dedicated out-of-band control plane signaling networks. These further complicate troubleshooting of the networks as the standard MPLS operations, administration, and maintenance (OAM) troubleshooting tools such as LSP ping and traceroute are not allowed any insight into the topology of the optical core. Further reducing the effectiveness of existing tools is the fact that LSP ping and traceroute have not yet been adapted to carry extended GMPLS label information.

Finally, despite the discussion above regarding standards-based management interfaces, the fact of the matter is that the most common interface used to manage optical core devices today is still the outdated TL-1 interface. Since it provides what is essentially an unstructured nonstandard interface suitable for managing the devices from one vendor, it further complicates matters in a multivendor network.

Recommendations

Given the challenges described above, we would like to recommend a few solutions that, combined with existing tools, may contribute to a more comprehensive management solution for GMPLS networks. First, tools like LSP ping and traceroute might be adapted to carry optical label information. We should also consider providing topology information to these tools for troubleshooting purposes in cases where the transfer of such information is acceptable to all parties involved. Second, although a standards-based northbound interface exists in the form of an XML data model, a similar model needs to be constructed that can be used to control the optical nodes and GMPLS LSRs. This will facilitate management from a common management station. Along these lines, vendors should begin migrating away from antiquated TL-1 interfaces toward relatively modern interfaces such as SNMP or XML. In general, we recommend that GMPLS networks migrate toward ultimately being managed as MPLS networks are managed today.

Conclusions

In this article we first introduce the various network configurations of GMPLS. We then discussed the tools and techniques that were available today for managing these networks. We then discuss why these tools are inadequate in providing operators a comprehensive solution for managing GMPLS networks. Finally, we make an attempt to suggest some solutions to these challenges that, when implemented, will allow for a complete management solution for GMPLS networks.

Additional Reading

[1] IETF RFC 3945, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture."

[2] IETF RFC 3471, "Generalized Multi-Protocol Label Switching (GMPLS) — Signaling Functional Description."

[3] IETF RFC 3473,"Generalized Multi-Protocol Label Switching (GMPLS) — Signaling Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions."

[4] IETF draft-ietf-ccamp-gmpls-routing-09.txt, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching."

[5] IETF draft-ietf-ccamp-ospf-gmpls-extensions-12.txt, "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching."

[6] IETF draft-ietf-ccamp-lmp-10.txt, "Link Management Protocol (LMP)."

[7] IETF draft-ietf-ccamp-gmpls-overlay-05.txt,"GMPLS UNI: RSVP Support for the Overlay Model."

[8] IETF draft-ietf-ccamp-gmpls-te-mib-09.txt,"Generalized Multiprotocol Label Switching (GMPLS) Traffic Engineering Management Information Base."

[9] IETF draft-ietf-ccamp-gmpls-lsr-mib-08.txt, "Generalized Multiprotocol Label Switching (GMPLS) Label Switch Router Management Information Base."

[10] IETF draft-ietf-ccamp-lmp-mib-10.txt, "Link Management Protocol Management Information Base."

[11] IETF draft-ietf-ccamp-gmpls-ason-reqts-07.txtx, "Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extensions for Automatically Switched Optical Network (ASON)."

[12] OIF-UNI-01.0, "User Network Interface(UNI) 1.0 Signaling Specification."

[13] Lang *et. al.*, "Link Management Protocol (LMP)," OIF2001.461.

[14] ITU Rec. G.8080/Y.1304, "Architecture for Automatically Switched Optical Network (ASON)."

[15] ITU Rec. G.807/Y.1302, "Requirements for Automatic Switched Transport Networks (ASTN)."

[16] ITU Rec. G.7713.2/ Y.1704.2, "Distributed Call and Connection Management: Signaling Mechanism Using GMPLS RSVP-TE."

[17] ITU Rec. G7718/Y.1709, "Framework for ASON Management."

Biographies

Thomas D. Nadeau (tnadeau@cisco.com) works at Cisco Systems, where he is a technical leader responsible for leadership of operations and management and network management standards, development, and architecture for MPLS-related components. He is an active participant in the IETF, ITU, and IEEE. He is co-author of many IETF MIBs, protocol and architecture documents in the L2/L3VPN, TE, PWE3, GMPLS, and MPLS areas. He is co-author of RFCs 3564, 3811, 3812, 3813, 3814, 3815, 3916, 3945, and 3985. He was recently co-editor of the October 2004 *IEEE Communications Magazine* feature topic on MPLS operations and management. He has filed a number of patents in the area of networking operations and management. He received his B.S. in computer science from the University of New Hampshire, and an M.Sc. from the University of Massachusetts in Lowell, where he has been an adjunct professor of computer science since 2000. He currently teaches courses on the topic of data communications. He is also on the technical committees of several prominent networking conferences where he provides technical guidance on their content. He is technical editor of *Enabling VPN Aware Networks with MPLS* (Prentice-Hall, 2001), and author of *MPLS Network Management: MIBs, Tools, and Techniques* (Morgan-Kaufman, 2002). He is currently working on a second text covering the topic of MPLS operations and management, which is expected in late 2005.

Hari Rakotoranto is currently product manager for GMPLS in ITD. He is also working very closely with service providers/Interoperability bodies (ISOCORE/EANTC/etc.) in the field of MPLS (VPN, TE, GMPLS, etc.) on technology deployment and overseeing future directions. He has worked as a software engineer, pre-/post sales technical consulting on different technologies: layer2/3 switches, routing protocol, network management and Unix kernel design. He is an active member of ITU-T SG13 focusing mainly on MPLS, MPLS OAM, and MPLS management.