

CCAMP
Internet Draft
Document: <[draft-bms-optical-sdhsonet-mpls-control-frmrk-01.txt](#)>
Category: Informational

G. Bernstein
Ciena
E. Mannie
Ebony
V. Sharma
Metanoia
July 2001

Expires January 2002

Framework for GMPLS-based Control of SDH/SONET Networks
<[draft-bms-optical-sdhsonet-mpls-control-frmrk-01.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#) [1].

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at

<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at

<http://www.ietf.org/shadow.html>.

1. Abstract

The suite of protocols that defines Multi-Protocol Label Switching (MPLS) is in the process of enhancement to generalize its applicability to the control of non-packet based switching, that is, optical switching. One area of prime consideration is to use these generalized MPLS (GMPLS) protocols in upgrading the control plane of optical transport networks. This document illustrates this process by describing those extensions to MPLS protocols that are directed towards controlling SONET/SDH networks. SONET/SDH networks make very good examples of this process since they possess a rich multiplex structure, a variety of protection/restoration options, are well defined, and are widely deployed. The document discusses extensions to MPLS routing protocols to disseminate information needed in transport path computation and network operations together with the extensions to MPLS label distribution protocols needed for the provisioning of transport circuits. New capabilities that an MPLS control plane would bring to SONET/SDH networks, such as new

restoration methods and multi-layer circuit establishment, are also discussed.

Bernstein, Mannie, Sharma Informational - January 2002

1

GMPLS based Control of SDH/SONET

July 2001

[2.](#) Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [2].

[3.](#) Introduction

A few years ago, the Internet Engineering Task Force (IETF) began work on the specification of a new connection-oriented transport technology called Multi-Protocol Label Switching (MPLS). The MPLS forwarding plane was inspired mainly by concepts from virtual circuit switching in ATM, while its control plane was inspired mainly by the routing protocols found in IP. As work on defining the components of MPLS progressed, it soon became apparent that the principles upon which MPLS technology was based were generic, and were applicable to multiple layers of the transport network. As such, MPLS-based control of other network layers, such as the time division multiplexing (TDM) and optical layers was also possible. The motivation behind introducing such control was to provide new services, such as dynamic establishment of TDM and optical circuits, which were hitherto not possible in transport networks. With MPLS-based control, transport operators or service providers would be able to offer on-demand services to their customers, due to the reduction in provisioning time of their circuits, thus adding considerable flexibility in their service portfolios.

The CCAMP Working Group of the IETF is currently working on extending MPLS protocols to support multiple network layers and these new services. This extended MPLS, which was initially known as Multi-Protocol Lambda Switching, is now better referred to as Generalized MPLS (or GMPLS). The authors of this work are among the co-authors of the GMPLS specifications, and - focus mainly on those aspects of GMPLS that relate to the control of SDH/SONET networks.

The GMPLS effort is, in effect, extending IP technology to control and manage lower layers. Using the same framework and similar signaling and routing protocols to control multiple layers can not only reduce the overall complexity of designing, deploying and

maintaining networks, but can also make it possible to operate two contiguous layers by using either an overlay model, a peer model, or an integrated model. The benefits of using a peer or an overlay model between the IP layer and its underlying layer(s) will have to be clarified and evaluated in the future. In the mean time, GMPLS could be used for controlling each layer independently.

The goal of this work is to highlight how GMPLS could be used to dynamically establish, maintain and tear down SDH/SONET circuits. The objective of using these extended MPLS protocols is to provide at least the same kinds of SDH/SONET services as are provided today, but using signaling instead of provisioning via centralized

Bernstein, Mannie, Sharma Informational - January 2002

2

GMPLS based Control of SDH/SONET

July 2001

management to establish those services. This will allow operators to propose new services, and will allow clients to create SONET/SDH paths on-demand, in real-time, through the provider network. We first review the essential properties of SDH/SONET networks and their operations, and we show how the label concept in MPLS can be extended to the SONET/SDH case. We then look at important information to be disseminated by a link state routing protocol and look at the important signal attributes that need to be conveyed by a label distribution protocol. Finally, we look at some outstanding issues and future possibilities. [3], [4], [5], [6], [7], [8], [9], [10], [11], [12].

3.1. MPLS Overview

A major advantage of the MPLS architecture for use as a general network control plane is its clear separation between the forwarding plane (or data plane) the signaling (or connection control) plane, and the routing (or topology discovery/resource status) plane. This allows the work on MPLS extensions to focus on the forwarding and signaling planes, while allowing well-known IP routing protocols to be reused in the routing plane. This clear separation also allows for MPLS to be used to control networks that do not have a packet-based forwarding plane.

An MPLS network consists of MPLS nodes called Label Switch Routers (LSRs) connected via circuits called Label Switched Paths (LSPs). An LSP is unidirectional and could be of several different types such as point-to-point, point-to-multipoint, and multipoint-to-point. Border LSRs in an MPLS network, act either as ingress or egress LSRs depending on the direction of the traffic being forwarded.

Each LSP is associated with a Forwarding Equivalence Class (FEC), which may be thought of as a set of packets that receive identical

forwarding treatment at an LSR. The simplest example of an FEC might be the set of destination addresses lying in a given address range. All packets that have a destination address lying within this address range are forwarded identically at each LSR configured with that FEC.

To establish an LSP, a signaling protocol (or label distribution protocol) such as LDP/CR-LDP or RSVP-TE is required. Between two adjacent LSRs, an LSP is locally identified by a short, fixed length identifier called a label, which is only significant between these two LSRs. The signaling protocol is responsible for the inter-node communication that assigns and maintains these labels.

When a packet enters an MPLS-based packet network, it is classified according to its FEC and, possibly, additional rules, which together determine the LSP along which the packet must be sent. For this purpose, the ingress LSR attaches an appropriate label to the packet, and forwards the packet to the next hop. The label may be attached to a packet in different ways. For example, it may be in the form of a header encapsulating the packet (the "shim" header) or

Bernstein, Mannie, Sharma Informational - January 2002

3

GMPLS based Control of SDH/SONET

July 2001

it may be written in the VPI/VCI field (or DLCI field) of the layer 2 encapsulation of the packet. In case of SDH/SONET networks, we will see that a label is simply associated with a segment of a circuit, and is mainly used in the signaling plane to identify this segment (e.g. a time-slot) between two adjacent nodes.

When a packet reaches a core packet LSR, this LSR uses the label as an index into a forwarding table to determine the next hop and the corresponding outgoing label (and, possibly, the QoS treatment to be given to the packet), writes the new label into the packet, and forwards the packet to the next hop. When the packet reaches the egress LSR, the label is removed and the packet is forwarded using appropriate forwarding, such as normal IP forwarding. We will see that for a SONET/SDH network these operations do not occur in quite the same way.

3.2. SDH/SONET Overview


There are currently two different multiplexing technologies in use in optical networks: wavelength division multiplexing (WDM) and time division multiplexing (TDM). This work focuses on TDM technology. SDH and SONET are two TDM standards widely used by operators to transport and multiplex different tributary signals over optical links, thus creating a multiplexing structure, which we call the SDH/SONET multiplex. SDH, which was developed by the ETSI and later

standardized by the ITU-T, is now used worldwide, while SONET, which was standardized by the ANSI, is mainly used in the US. However, these two standards have several similarities, and to some extent SONET can be viewed as a subset of SDH. Internetworking between the two is possible using gateways.

The fundamental signal in SDH is the STM-1 that operates at a rate of about 155 Mbps, while the fundamental signal in SONET is the STS-1 that operates at a rate of about 51 Mbps. These two signals are made of contiguous frames that consist of transport overhead (header) and payload. To solve synchronization issues, the actual data is not transported directly in the payload but rather in another internal frame that is allowed to float over two successive SDH/SONET payloads. This internal frame is named a Virtual Container (VC) in SDH and a Synchronous Payload Envelope (SPE) in SONET.

The SDH/SONET architecture identifies three different layers, each of which corresponds to one level of communication between SDH/SONET equipment. These are, starting with the lowest, the regenerator section/section layer, the multiplex section/line layer, and (at the top) the path layer. Each of these layers in turn has its own overhead (header). The transport overhead of a SDH/SONET frame is mainly sub-divided in two parts that contain the regenerator section/section overhead and the multiplex section/line overhead. In addition, a pointer (in the form of the H1, H2 and H3 bytes) indicates the beginning of the VC/SPE in the payload of the overall STM/SDH frame.

The VC/SPE itself is made up of a header (the path overhead) and a payload. This payload can be further subdivided into sub-elements (signals) in a fairly complex way. In the case of SDH, the STM-1 frame may contain either one VC-4 or three multiplexed VC-3s. The SONET multiplex is a pure tree, while the SDH multiplex is not a pure tree, since it contains a node that can be attached to two parent nodes. The structure of the SONET/SDH multiplex is shown in Figure 1. In addition, we show reference points in this figure that are explained in later sections.



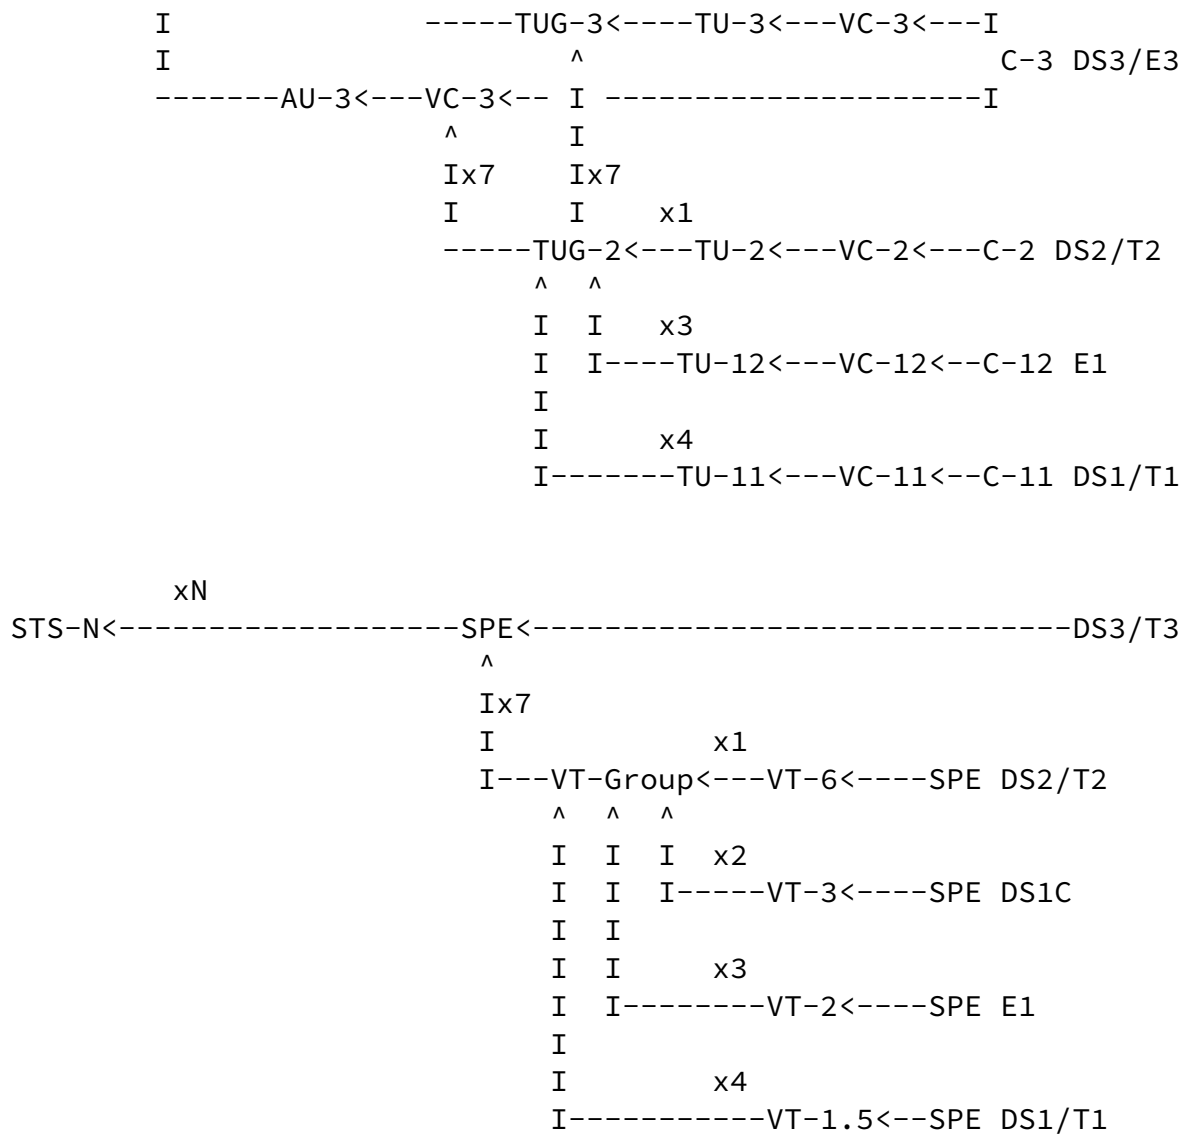


Figure 1. SDH and SONET multiplexing structure and typical PDH payload signals.

The leaves of these multiplex structures are time slots (positions) of different sizes that can contain tributary signals. These tributary signals (e.g. E1, E3, etc) are mapped into the leaves using standardized mapping rules. In general, a tributary signal does not fill a time slot completely, and the mapping rules define precisely how to fill it.

What is important for the MPLS-based control of SDH/SONET circuits is to identify the elements that can be switched from an input multiplex on one interface to an output multiplex on another interface. The only elements that can be switched are those that can be re-aligned via a pointer, i.e. a VC-x in the case of SDH and a SPE in the case of SONET.

An STM-N/STS-N signal is formed from $N \times$ STM-1/STS-1 signals via byte interleaving. The VCs/SPEs in the N interleaved frames are independent and float according to their own clocking. To transport tributary signals in excess of the basic STM-1/STS-1 signal rates, the VCs/SPEs can be concatenated, i.e., glued together. In this case their relationship with respect to each other is fixed in time and hence this relieves, when possible, an end system of any inverse multiplexing bonding processes. Different types of concatenations are defined in SDH/SONET.

For example, standard SONET concatenation allows the concatenation of $M \times$ STS-1 signals within an STS- N signal with $M \leq N$, and $M = 3, 12, 48, 192, \dots$). The SPEs of these $M \times$ STS-1s can be concatenated to form an STS- M_c . The STS- M_c notation is short hand for describing an STS- M signal whose SPEs have been concatenated.

3.3. The Current State of Circuit Establishment in SDH/SONET Networks

Today, June 2001, SDH and SONET networks are statically configured. When a client of an operator requests a point-to-point circuit, the request sets in motion a process that can last for several weeks or more. This process is composed of a chain of shorter administrative and technical tasks, some of which can be fully automated, resulting in significant improvements in provisioning time and in operational savings. In the best case, the entire process can be fully automated allowing, for example, customer premise equipment (CPE) to contact a SDH/SONET switch to request a circuit. Currently, the provisioning process involves the following tasks.

3.3.1. Administrative Tasks

The administrative tasks represent a significant part of the provisioning time. Most of them can be automated using IT applications, e.g., a client still has to fill a form to request a circuit. This form can be filled via a Web-based application and can be automatically processed by the operator. A further enhancement is

to allow the client's equipment to coordinate with the operator's

network directly and request the desired circuit. This could be achieved through a signaling protocol at the interface between the client equipment and an operator switch, i.e., at the UNI interface, where GMPLS signaling can be used.

3.3.2. Manual Operations

Another significant part of the time may be consumed by manual operations that involve installing the right interface in the CPE and installing the right cable or fiber between the CPE and the operator switch. This time can be especially significant when a client is in a different time zone than the operator's main office. This first-time connection time is frequently accounted for in the overall establishment time.

3.3.3. Planning Tool Operation

Another portion of the time is consumed by planning tools that run simulations using heuristic algorithms to find an optimized placement for the required circuits. These planning tools can require a significant running time, sometimes on the order of days. These simulations are, in general, executed for a set of demands for circuits, i.e., a batch mode, to improve the optimality of network resource usage and other parameters. Today, we do not really have a means to reduce this simulation time. On the contrary, to support fast, on-line, circuit establishment, this phase may be invoked more frequently, i.e., we will not "batch up" as many connection requests before we plan out the corresponding circuits. This means that the network may need to be re-optimized periodically, implying that the signaling should support re-optimization with minimum impact to existing services.

3.3.4. Circuit Provisioning

Once the first three steps have been completed, the circuits must be provisioned by the operator using the outputs of the planning process. The time required for provisioning varies greatly. It can be fairly short, on the order of a few minutes, if the operators already have tools that help them to do the provisioning over heterogeneous equipment. Otherwise, the process can take days. Developing these tools for each new piece of equipment and each vendor is a significant burden on the service provider. A standardized interface for provisioning, such as GMPLS signaling, could significantly reduce or eliminate this development burden. In general, provisioning is a batched activity, i.e., a few times per week an operator provisions a set of circuits. GMPLS will reduce this provisioning time from a few minutes to a few seconds and could help to transform this periodic process into a real-time process.

When a circuit is provisioned, it is not delivered directly to a client. Rather, the operator first tests its performance and

behavior and if successful, delivers the circuit to the client. This

testing phase lasts, in general, for up to 24 hours. The operator installs test equipment at each end and uses pre-defined test streams to verify performance. If successful, the circuit is officially accepted by the client. To speed up the verification (sometimes known as "proving") process, it would be necessary to support some form of automated performance testing.

3.4. Centralized Approach versus Distributed Approach

Whether a centralized approach or a distributed approach will be used to control SDH/SONET networks is an open question, since each approach has advantages and disadvantages. The application of GMPLS to SONET/SDH networks does not preclude either model although MPLS is itself a distributed technology.

The basic tradeoff between the centralized and distributed approaches is that of complexity of the network elements versus that of the network management system (NMS). Since adding functionality to existing

SDH network elements may not be possible, a centralized approach may be needed in some cases. The main issue facing centralized control via an NMS is one of scalability. For instance, this approach may be limited in the number of network elements that can be managed (e.g. one thousand). It is, therefore, quite common for operators to deploy several NMSs in parallel at the Network Management Layer, each managing a different zone. In that case, however, a Service Management layer must be built on the top of several individual NMSs to take care of end-to-end on-demand services. On the other hand, in a complex and/or dense network, restoration could be faster with a distributed approach than with a centralized approach.

Let's now look at how the major control plane functional components are handled via the centralized and distributed approaches:

3.4.1. Topology Discovery and Resource Dissemination

Currently NMS's maintain a consistent view of all the networking layers under their purview. This can include the physical topology, such as information about fibers and ducts. Since most of this information is entered manually, it remains error prone.

A link state GMPLS routing protocol, on the other hand, could perform automatic topology discovery and dissemination the topology as well as resource status. This information would be available to all nodes in the network, and hence also the NMS. Hence one can

look at a continuum of functionality between manually provisioned topology information (of which there will always be some) and fully automated discovery and dissemination as in a link state protocol. Note that, unlike the IP datagram case, a link state routing protocol applied to the SDH/SONET network does not have any service impacting implications.

Bernstein, Mannie, Sharma Informational - January 2002

8

GMPLS based Control of SDH/SONET

July 2001

3.4.2. Path Computation (Route Determination)

In the SDH/SONET case, unlike the IP datagram case, there is no need for network elements to all perform the same path calculation. In addition, path determination is an area for vendors to provide a potentially significant value addition in terms of network efficiency, reliability, and service differentiation. In this sense, a centralized approach to path computation is easier to operate and upgrade. For example, new features such as new types of path diversity or new optimization algorithms can be introduced with a simple NMS software upgrade. On the other hand, updating switches with new path computation software is a more complicated task. In addition, many of the algorithms are quite computationally intensive and may be completely unsuitable for the embedded processing environment available on most switches. In restoration scenarios the ability to perform a reasonably sophisticated level of path computation on the network element can be particularly useful for restoring traffic during major network faults.

3.4.3. Connection Establishment (provisioning)

The actual setting up of circuits, i.e., a coupled collection of cross connects across a network, can be done either via the NMS setting up individual cross connects or via a "soft permanent LSP" (SPLSP) type approach. In the SPLSP approach, the NMS may just kick off the connection at the "ingress" switch with GMPLS signaling setting up the connection from that point onward. Connection establishment is the trickiest part to distribute, however, since errors in the connection setup/tear down process are service impacting.

Distributed approach

Centralized approach

Control plane like MPLS or PNNI Do we really need it? Being added/specified by several standardization bodies	Management plane like TMN or SNMP Always needed! Already there, proven and understood.
High survivability (e.g. in case of partition)	Potential single point(s) of failure
Distributed load	Bottleneck: #requests and actions to/from NMS
Individual local routing decision	Centralized routing decision, can be done per block of requests
Routing scalable as for the	Assumes a few big

Bernstein, Mannie, Sharma Informational - January 2002

9

GMPLS based Control of SDH/SONET

July 2001

Internet	administrative domains
Complex to change routing protocol/algorithm	Very easy local upgrade (non-intrusive)
Requires enhanced routing protocol (traffic engineering)	Better consistency
Ideal for inter-domain	Not inter-domain friendly
Suitable for very dynamic demands	For less dynamic demands (longer lived)
Probably faster to restore, but more difficult to have reliable restoration.	Probably slower to restore, but could effect reliable restoration.
High scalability (hierarchical)	Limited scalability: #nodes, links, circuits, messages
Planning (optimization) harder to achieve	Planning is a background centralized activity
Easier future integration with other control plane layers	

Table 1. Qualitative comparison between centralized and distributed

approaches.

3.5. Why SDH/SONET will not Disappear Tomorrow

As IP traffic becomes the dominant traffic transported over the transport infrastructure, it is useful to compare the statistical multiplexing of IP with the time division multiplexing of SDH and SONET.

Consider a scenario where IP over WDM is used everywhere and lambdas are optically switched. In such a case, a carrier's carrier would sell dynamically controlled lambdas with each customer building his/her own IP backbone over these lambdas.

This simple model implies that a carrier would sell lambdas instead of bandwidth. The carrier's goal will be to maximize the number of wavelengths/lambdas per fiber, with each customer having to fully support the cost for each end-to-end lambda whether or not the wavelength is fully utilized. Although, in the near future, we may have technology to support up to several hundred lambdas per fiber, a world where lambdas are so cheap and abundant that every individual customer buys them, from one point to any other point, appears an unlikely scenario today.

Bernstein, Mannie, Sharma Informational - January 2002

10

GMPLS based Control of SDH/SONET

July 2001

More realistically, there is stillroom for a multiplexing technology that provides circuits with a lower granularity than a wavelength. (Not everyone needs a minimum of 10 Gbps or 40 Gbps per circuit, and IP does not yet support all telecom applications in bulk efficiently.)

SDH and SONET possess a rich multiplexing hierarchy that permits fairly fine granularity and that provides a very cheap and simple physical separation of the transported traffic between circuits, i.e., QoS.

Moreover, even IP datagrams are not transported directly over a wavelength. A framing or encapsulation is always required to delimit IP datagrams. The Total Length field of an IP header cannot be trusted to find the start of a new datagram, since it could be corrupted and would result in a loss of synchronization. The typical framing used today for IP over DWDM is defined in [RFC1619](#)/RFC2615 and known as POS (Packet Over SONET/SDH), i.e., IP over PPP (in HDLC-like format) over SDH/SONET. SDH and SONET are actually efficient encapsulations for IP. For instance, with an average IP datagram length of 350 octets, an IP over GBE encapsulation using an

8B/10B encoding results in 28% overhead, an IP/ATM/SDH encapsulation results in 22% overhead and an IP/PPP/SDH encapsulation results in only 6% overhead. (New simplified encapsulations could reduce this overhead to as low as 3%, but the gain is not huge compared to SDH and SONET -, which have other benefits as well.)

Any encapsulation of IP over WDM should at least provide error monitoring capabilities (to detect signal degradation), error correction capabilities, such as FEC (Forward Error Correction) that are particularly needed for ultra long haul transmission, sufficient timing information, to allow robust synchronization (that is, to detect the beginning of a packet), and capacity to transport signaling, routing and management messages, in order to control the optical switches. SDH and SONET cover all these aspects natively, except FEC, which tends to be supported in a proprietary way.

Since IP encapsulated in SDH/SONET is efficient and widely used, the only real difference between an IP over WDM network and an IP over SDH over WDM network is the layers at which the switching or forwarding can take place. In the first case, it can take place at the IP and optical layers. In the second case, it can take place at the IP, SDH/SONET, and optical layers.

Almost all transmission networks today are based on SDH or SONET. A client is connected either directly through an SDH or SONET interface or through a PDH interface, the PDH signal being transported between the ingress and the egress interfaces over SDH or SONET. What we are arguing here is that it makes sense to do switching or forwarding at all these layers.

[4. GMPLS Applied to SDH/SONET](#)

Bernstein, Mannie, Sharma Informational - January 2002

11

GMPLS based Control of SDH/SONET

July 2001

4.1. Controlling the SDH/SONET Multiplex

Controlling the SDH/SONET multiplex implies deciding which of the different components of the SDH/SONET multiplex that can be switched do we wish to control using GMPLS. Essentially, every SDH/SONET element that is referenced by a pointer can be switched. These component signals are the VC-4, VC-3, VC-2, VC-12 and VC-11 in the SDH case; and the VT and STS SPEs in the SONET case. The SONET case is a bit difficult to explain since, unlike in SDH, SPEs in SONET do not have individual names. We will refer to them by identifying the structure that contains them, namely the STS-1, VT-6, VT-3, VT-2 and VT-1.5.

The STS-1 SPE corresponds to a VC-3, a VT-6 SPE corresponds to a VC-2, a VT-2 SPE corresponds to a VC-12, and a VT-1.5 SPE corresponds to a VC-11. The SONET VT-3 SPE has no correspondence in SDH, however SDH's VC-4 corresponds to SONET's STS-3c SPE.

In addition, it is possible to concatenate some of the structures that contain these elements to build larger elements. For instance, SDH allows the concatenation of X contiguous AU-4s to build a VC-4- Xc and of m contiguous TU-2s to build a VC-2- mc . In that case, a VC-4- Xc or a VC-2- mc can be switched and controlled by MPLS. Note that SDH also defines virtual (non-contiguous) concatenation of TU-2s, but in that case each constituent VC-2 is switched individually.

4.2. SDH/SONET LSR and LSP Terminology

Let a SDH or SONET Terminal Multiplexer (TM), Add-Drop Multiplexer (ADM) or cross-connect (i.e. a switch) be called an SDH/SONET LSR. A SDH/SONET path or circuit between two SDH/SONET LSRs now becomes a GMPLS LSP. An SDH/SONET LSP is a logical connection between the point at which a tributary signal (client layer) is adapted into its virtual container, and the point at which it is extracted from its virtual container.

To establish such an LSP, a signaling protocol is required to configure the input interface, switch fabric, and output interface of each SDH/SONET LSR along the path. An SDH/SONET LSP can be point-to-point or point-to-multipoint, but not multipoint-to-point, since no merging is possible with SDH/SONET signals.

To facilitate the signaling and setup of SDH/SONET circuits, an SDH/SONET LSR must, therefore, identify each possible signal individually per interface, since each signal corresponds to a potential LSP that can be established through the SDH/SONET LSR. It turns out, however, that not all SDH signals correspond to an LSP and therefore not all of them need be identified. In fact, only those signals that can be switched need identification.

[5. Decomposition of the MPLS Circuit-Switching Problem Space](#)

Although those familiar with MPLS may be familiar with its application in a variety of application areas, e.g., ATM, Frame Relay, and so on, here we quickly review its decomposition when applied to the optical switching problem space.

(i) Information needed to compute paths must be made globally available throughout the network. Since this is done via the link state routing protocol, any information of this nature must either be in the existing link state advertisements (LSAs) or the LSAs must be supplemented to convey this information. For example, if it is desirable to offer different levels of service in a network based on whether a circuit is routed over SDH/SONET lines that are ring protected versus being routed over those that are not ring protected (differentiation based on reliability), the type of protection on a SDH/SONET line would be an important topological parameter that would have to be distributed via the link state routing protocol.

(ii) Information that is only needed between two "adjacent" switches for the purposes of connection establishment is appropriate for distribution via one of the label distribution protocols. In fact, this information can be thought of as the "virtual" label. For example, in SONET networks, when distributing information to switches concerning an end-to-end STS-1 path traversing a network, it is critical that adjacent switches agree on the multiplex entry used by this STS-1 (but this information is only of local significance between those two switches). Hence, the multiplex entry number in this case can be used as a virtual label. Note that the label is virtual in that it is not appended to the payload in any way, but it is still a label in the sense that it uniquely identifies the signal locally on the link between the two switches.

(iii) Information that all switches in the path need to know about a circuit will also be distributed via the label distribution protocol. Examples of such information include bandwidth, priority, and preemption for instance.

(iv) Information intended only for end systems of the connection. Some of the payload type information in may fall into this category. [8],[10].

6. MPLS Routing for SDH/SONET

Modern transport networks based on SONET/SDH excel at interoperability in the performance monitoring (PM) and fault management (FM) areas., They do not, however, inter-operate in the areas of topology discovery or resource status. Although link state routing protocols, such as IS-IS and OSPF, have been used for some time in the IP world to compute destination-based next hops for routes (without routing loops), their value in providing timely topology and network status information in a distributed manner, i.e., at any network node, is immense. If resource utilization information is disseminated along with the link status (as was done in ATM's PNNI routing protocol) then a very complete picture of

network status is available to a network operator for use in planning, provisioning and operations.

The information needed to compute the path a connection will take through a network is important to distribute via the routing protocol. In the optical TDM case, this information includes, but is not limited to: the available capacity of the network links, the switching and termination capabilities of the nodes and interfaces, and the protection properties of the link.

When applying routing to circuit switched networks it is useful to compare and contrast this situation with the datagram routing case. In the case of routing datagrams all routes on all nodes must be calculated exactly the same to avoid loops and "black holes". In circuit switching, this is not the case since routes are established per circuit and are fixed for that circuit. Hence, unlike the datagram case, routing is not service impacting in the circuit switched case. This is helpful, because to accommodate the optical layer routing protocols need to be supplemented with new information, much more than the datagram case. This information is also likely to be used in different ways for implementing different user services. Due to the increase in information transferred in the routing protocol, it is important to separate the relatively static parameters concerning a link from those that may be subject to frequent changes. This is particularly important in the case of available capacity advertisements.

6.1. Switching Capabilities

The main switching capabilities that characterize a SONET/SDH end system and thus need to be advertised via the link state routing protocol are: the switching granularity, supported forms of concatenation, and the level of transparency.

6.1.1. Switching Granularity

From references [3], [4] and the overview section on SONET/SDH we see that there are a number of different signals that compose the SONET/SDH hierarchies. Those signals that are referenced via a pointer, i.e., the VCs in SDH and the SPEs in SONET are those that will actually be switched within a SONET/SDH network. These signals are subdivided into lower order signals and higher order signals as shown in Table 2.

Table 2. SDH/SONET switched signal groupings.

Signal Type	SDH	SONET
Lower Order	VC-11, VC-12, VC-2	VT-1.5 SPE, VT-2 SPE, VT-3 SPE, VT-6 SPE
Higher	VC-3, VC-4	STS-1 SPE

Bernstein, Mannie, Sharma Informational - January 2002

14

GMPLS based Control of SDH/SONET

July 2001

Order

Manufacturers today differ in the types of switching capabilities their systems support. Many manufacturers today switch signals starting at VC-4 for SDH or STS-1 for SONET (i.e. the basic frame) and above (see concatenation section), but they do not switch lower order signals. Some of them only allow the switching of entire aggregates (concatenated or not) of signals such as 16 VC-4s, i.e. a complete STM-16, and nothing finer. Some go down to the VC-3 level for SDH. Finally, some offer highly integrated switches that switch at the VC-3/STS-1 level down to lower order signals such as VC-12s. In order to cover the needs of all manufacturers and operators, GMPLS must consider both higher order and lower order signals.

6.1.2. Signal Concatenation Capabilities

As stated in the SONET/SDH overview, to transport tributary signals with rates in excess of the basic STM-1/STS-1 signal, the VCs/SPEs can be concatenated, i.e., glued together. Different types of concatenations are defined: contiguous standard concatenation, arbitrary concatenation, and virtual concatenation with different rules concerning their size, placement, and binding.

Standard SONET concatenation allows the concatenation of $M \times \text{STS-1}$ signals within an STS- N signal with $M \leq N$, and $M = 3, 12, 48, 192, \dots$). The SPEs of these $M \times \text{STS-1}$ s can be concatenated to form an STS- Mc . The STS- Mc notation is short hand for describing an STS- M signal whose SPEs have been concatenated. The multiplexing procedures for SONET and SDH are given in references [3] and [4]. Constraints are imposed on the size of STS- Mc signals, i.e., they must be a multiple of 3, and on their starting location and interleaving.

This has the following advantages: (a) restriction to multiples of 3 helps with SDH compatibility (there is no STS-1 equivalent signal in SDH); (b) the restriction to multiples of 3 reduces the number of connection types; (c) the restriction on the placement and interleaving could allow more compact representation of the "label";

The major disadvantages of these restrictions are:

(a) Limited flexibility in bandwidth assignment (somewhat inhibits finer grained traffic engineering). (b) The lack of flexibility in starting time slots for STS-Mc signals and in their interleaving (where the rest of the signal gets put in terms of STS-1 slot numbers) leads to the requirement for re-grooming (due to bandwidth fragmentation).

Due to these disadvantages some SONET framer manufacturers now support "flexible" or arbitrary concatenation, i.e., no restrictions on the size of an STS-Mc (as long as $M \leq N$) and no constraints on the STS-1 timeslots used to convey it, i.e., the signals can use any combination of available time slots.

Bernstein, Mannie, Sharma Informational - January 2002

15

GMPLS based Control of SDH/SONET

July 2001

Standard and flexible concatenations are network services, while virtual concatenation is a SONET/SDH end-system service recently approved by the committee T1 of ANSI and ITU-T. The essence of this service is to have SONET/SDH end systems "glue" together the VCs or SPEs of separate signals rather than requiring that the signals be carried through the network as a single unit. In one example of virtual concatenation, two end systems supporting this feature could essentially "inverse multiplex" two STS-1s into a virtual STS-2c for the efficient transport of 100Mbps Ethernet traffic. Note that this inverse multiplexing process can be significantly easier to implement with SONET/SDH signals rather than packets. Since virtual concatenation is provided by end systems, it is compatible with existing SONET/SDH networks. Virtual concatenation is defined for both higher order signals and low order signals. Table 3 shows the nomenclature and capacity for several lower-order virtually concatenated signals contained within different higher-order signals.

Table 3 Capacity of Virtually Concatenated VTn-Xv (9/G.707)

	Carried In	X	Capacity	In steps of
VT1.5/ VC-11-Xv	STS-1/VC-3	1 to 28	1600kbit/s to 44800kbit/s	1600kbit/s
VT2/ VC-12-Xv	STS-1/VC-3	1 to 21	2176kbit/s to 45696kbit/s	2176kbit/s
VT1.5/ VC-11-Xv	STS-3c/VC-4	1 to 64	1600kbit/s to 102400kbit/s	1600kbit/s

VT2/ VC-12-Xv	STS-3c/VC-4	1 to 63	2176kbit/s to 2176kbit/s 137088kbit/s
------------------	-------------	---------	--

6.1.3. SDH/SONET Transparency

The purposed of SONET/SDH is to carry its payload signals in a transparent manner. This can include some of the layers of SONET itself. For example, situations where the path overhead can never be touched, since it actually belongs to the client. This was another reason for not coding an explicit label in SDH/SONET path overhead. It may be useful to transport, multiplex and/or switch lower layers of the SONET signal transparently.

As mentioned in the introduction, SONET overhead is broken into three layers: Section, Line and Path. Each of these layers is concerned with fault and performance monitoring. The Section overhead is primarily concerned with framing, while the Line overhead is primarily concerned with multiplexing and protection. To perform multiplexing, a SONET network element should be line terminating. However, not all SONET multiplexers/switches perform

Bernstein, Mannie, Sharma Informational - January 2002

16

GMPLS based Control of SDH/SONET

July 2001

SONET pointer adjustments on all the STS-1s contained within a higher order SONET signal passing through them. Alternatively, if they perform pointer adjustments, they do not terminate the line overhead. For example, a multiplexer may take four SONET STS-48 signals and multiplex them onto an STS-192 without performing standard line pointer adjustments on the individual STS-1s. This can be looked at as a service since it may be desirable to pass SONET signals, like an STS-12 or STS-48, with some level of transparency through a network and still take advantage of TDM technology. Transparent multiplexing and switching can also be viewed as a constraint, since some multiplexers and switches may not switch with as fine a granularity as others. Table 4 summarizes the levels of SONET/SDH transparency.

Table 4. SONET/SDH transparency types and their properties.

Transparency Type	Comments
Path Layer (or Line Terminating)	Standard higher order SONET path switching. Line overhead is terminated or modified.
Line Level (or Section	Preserves line overhead and switches

Terminating)	the entire line multiplex as a whole. Section overhead is terminated or modified.
Section layer	Preserves all section overhead, Basically does not touch any of the SONET/SDH bits.

6.2. Protection

SONET and SDH networks offer a variety of protection options at both the SONET line (SDH multiplex section) and SONET/SDH path level[5][6]. Standardized SONET line level protection techniques include Linear 1+1 and Linear 1:N automatic protection switching (APS) and both two-fiber and four-fiber bi-directional line switched rings (BLSRs). At the path layer, SONET offers uni-directional path switched ring protection. Both ring and 1:N line protection also allow for "extra traffic" to be carried over the protection line when that line is not being used, i.e., when it is not carrying traffic for a failed working line. These protection methods are summarized in Table 5. It should be noted that these protection methods are completely separate from any MPLS layer protection or restoration mechanisms.

Table 5. Common SONET/SDH protection mechanisms.

Protection Type	Extra Traffic Optionally	Comments
-----------------	--------------------------	----------

Bernstein, Mannie, Sharma Informational - January 2002

17

GMPLS based Control of SDH/SONET

July 2001

	Supported	
1+1 Unidirectional	No	Requires no coordination between the two ends of the circuit. Dedicated protection line.
1+1 Bi-directional	No	Coordination via K byte protocol. Lines must be consistently configured. Dedicated protection line.
1:1	Yes	Dedicated protection.
1:N	Yes	One Protection line shared

by N working lines.

4F-BLSR (4 fiber bi- directional line switched ring)	Yes	Dedicated protection, with alternative ring path.
2F-BLSR (2 fiber bi- directional line switched ring)	Yes	Dedicated protection, with alternative ring path
UPSR (uni- directional path switched ring)	No	Dedicated protection via alternative ring path. Typically used in access networks.

It may be desirable to route some connections over lines that support protection of a given type, while others may be routed over unprotected lines, or as "extra traffic" over protection lines. Also, to assist in the configuration of these various protection methods it can be extremely valuable to advertise the link protection attributes in the routing protocol. For example suppose that a 1:N protection group is being configured via two nodes. One must make sure that the lines are "numbered the same" with respect to both ends of the connection or else the APS (K1/K2 byte) protocol will not correctly operate.

Table 6. Parameters defining protection mechanisms.

Protection Related Link	Comments
----------------------------	----------

Bernstein, Mannie, Sharma Informational - January 2002

18

GMPLS based Control of SDH/SONET

July 2001

Information

Protection Type	Indicates which of the protection types delineated in Table 5.
-----------------	--

Protection Group Id	Indicates which of several protection groups (linear or ring) that a node belongs to. Must be unique for all groups that a node participates in
Working line number	Important in 1:N case and to differentiate between working and protection lines
Protection line number	Used to indicate if the line is a protection line.
Extra Traffic Supported	Yes or No
Layer	If this protection parameter is specific to SONET then this parameter is unneeded, otherwise it would indicate the signal layer that the protection is applied.

An open issue concerning protection is the extent of information regarding protection that must be disseminated. The contents of Table 6 represent one extreme while a simple enumerated list of: Extra-Traffic/Protection line, Unprotected, Shared (1:N)/Working line, Dedicated (1:1, 1+1)/Working Line, Enhanced (Ring) /Working Line, represents the other.

There is also a potential implication for link bundling, that is, for each link, the routing protocol could advertise whether that link is a working or protection link and possibly some parameters from Table 6. A possible drawback of this scheme is that the routing protocol would be burdened with advertising properties even for those protection links in the network that could not, in fact, be used for routing working traffic, e.g., dedicated protection links. An alternative method, would be to bundle the working and protection links together, and advertise the bundle instead. Now, for each bundled link, the protocol would have to advertise the amount of bandwidth available on its working links, as well as the amount of bandwidth available on those protection links within the bundle that were capable of carrying "extra traffic." This would reduce the amount of information to be advertised. An issue here would be to decide which types of working and protection links to bundle together. For instance, it might be preferable to bundle working

links (and their corresponding protection links) that are "shared" protected separately from working links that are "dedicated" protected.

6.3. Available Capacity Advertisement

Each SDH/SONET LSR must maintain an internal table per interface that indicates each signal in the multiplex structure that is allocated at that interface. This internal table is the most complete and accurate view of the link usage and available capacity.

For use in path computation, this information needs to be advertised in some way to all others SONET/SDH LSRs in the same domain . There is a trade off to be reached concerning: the amount of detail in the available capacity information to be reported via a link state routing protocol, the frequency or conditions under which this information is updated, the percentage of connection establishments that are unsuccessful on their first attempt due to the granularity of the advertised information, and the extent to which network resources can be optimized. There are different levels of summarization that are being considered today for the available capacity information. At one extreme, all signals that are allocated on an interface could be advertised, while at the other extreme, a single aggregated value of the available bandwidth per link could be advertised.

Consider first the relatively simple structure of SONET and its most common current and planned usage. DS1s and DS3s are the signals most often carried within a SONET STS-1. Either a single DS3 occupies the STS-1 or up to 28 DS1s (4 each within the 7 VT groups) are carried within the STS-1. With a reasonable VT1.5 placement algorithm within each node it may be possible to just report on aggregate bandwidth usage in terms of number of whole STS-1s (dedicated to DS3s) used and the number of STS-1s dedicated to carrying DS1s allocated for this purpose. This way a network optimization program could try to determine the optimal placement of DS3s and DS1s to minimize wasted bandwidth due to half-empty STS-1s at various places within the transport network. Similarly consider the set of super rate SONET signals (STS-Nc). If the links between the two switches support flexible concatenation then the reporting is particularly straightforward since any of the STS-1s within an STS-M can be used to comprise the transported STS- Nc. However, if only standard concatenation is supported then reporting gets trickier since there are constraints on where the STS-1s can be placed. SDH has still more options and constraints, hence it is not yet clear which is the best way to advertise bandwidth resource availability/usage in SONET/SDH. However, due to the multiplexed nature of the signals reporting of bandwidth particular to signal types rather than as a single aggregate bit rate is highly

desirable.

6.4. Path Computation

Although a link state routing protocol can be used to obtain network topology and resource information, this does not imply the use of an "open shortest path first" route. The path must be open in the sense that the links must be capable of supporting the desired signal type and that capacity must be available to carry the signal. Other constraints may include hop count, total delay (mostly propagation), and underlying protection. In addition, it may be desirable to route traffic in order to optimize overall network capacity, or reliability, or some combination of the two. Dijkstra's algorithm computes the shortest path with respect to link weights for a single connection at a time. This can be much different than the paths that would be selected in response to a request to set up a batch of connections between a set of endpoints in order to optimize network link utilization. One can think of this along the lines of global or local optimization of the network in time.

Due to the complexity of some of the connection routing algorithms (high dimensionality, non-linear integer programming problems) and various criteria by which one may optimize a network, it may not be possible or desirable to run these algorithms on network nodes. However, it may still be desirable to have some basic path computation ability running on the network nodes, particularly for use during restoration situations. Such an approach is in line with the use of MPLS for traffic engineering, but is much different than typical OSPF or IS-IS usage where all nodes must run the same routing algorithm.

7. LSP Provisioning/Signaling for SDH/SONET

Traditionally, end-to-end circuit connections in SDH/SONET networks have been set up via network management systems (NMSs), which issue commands (usually under the control of a human operator) to the various network elements involved in the circuit, via an equipment vendor's element management system (EMS). Very little multi-vendor interoperability has been achieved via management systems. Hence, end-to-end circuits in a multi-vendor environment typically require the use of multiple management systems and the infamous configuration via "yellow sticky notes". As discussed in [Section 2](#), a common signaling protocol, such as RSVP with TE extensions or CR-

LDP appropriately extended for circuit switching applications, could therefore help to solve these interoperability problems. In this section, we examine the various components involved in the automated provisioning of SONET/SDH LSPs.

7.1.1. What do we Label in SDH/SONET? Frames or Circuits?

MPLS was initially introduced to control asynchronous technologies like IP, where a label was attached to each individual block of data, such as an IP packet or a Frame Relay frame. SONET and SDH,

Bernstein, Mannie, Sharma Informational - January 2002

21

GMPLS based Control of SDH/SONET

July 2001

however, are synchronous technologies that define a multiplexing structure (see [Section 3.2](#)), which we referred to as the SDH (or SONET) multiplex. This multiplex involves a hierarchy of signals, lower order signals embedded within successive higher order ones (see Fig. 1). Thus, depending on its level in the hierarchy, each signal consists of frames that repeat periodically, with a certain number of byte time slots per frame.

The question then arises: is it these frames that we label in GMPLS? It will be seen in what follows that each SONET or SDH "frame" need not have its own label, nor is it necessary to switch frames individually. Rather, the unit that is switched is a "flow" comprised of a continuous sequence of time slots that appear at a given position in a frame. That is, we switch an individual SONET or SDH signal, and a label associated with each given signal.

For instance, the payload of an SDH STM-1 frame does not fully contain a complete unit of user data. In fact, the user data is contained in a virtual container (VC) that is allowed to float over two contiguous frames for synchronization purposes. A pointer in the Section Overhead (SOH) indicates the beginning of the VC in the payload. Thus, frames are now inter-related, since each consecutive pair may share a common virtual container. From the point of view of GMPLS, therefore, it is not the successive frames that are treated independently or labeled, but rather the entire user signal. An identical argument applies to SONET.

Observe also that the GMPLS signaling used to control the SDH/SONET multiplex must honor its hierarchy. In other words, the SDH/SONET layer should not be viewed as homogeneous and flat, because this would limit the scope of the services that SDH/SONET can provide. Instead, GMPLS tunnels should be used to dynamically and hierarchically control the SDH/SONET multiplex. For example, one unstructured VC-4 LSP may be established between two nodes, and

later lower order LSPs (e.g. VC-12) may be created within that higher order LSP. This VC-4 LSP can, in fact, be established between two non-adjacent internal nodes in an SDH network, and later advertised by a routing protocol as a new (virtual) link called a Forwarding Adjacency (FA).

A SONET/SDH-LSR will have to identify each possible signal individually per interface to fulfill the GMPLS operations. In order to stay transparent the LSR obviously should not touch the SONET/SDH overheads; this is why an explicit label is not encoded in the SDH/SONET overheads. Rather, a label is associated with each individual signal. This approach is similar to the one considered for lambda switching, except that it is more complex, since SONET and SDH define a richer multiplexing structure. Therefore a label is associated with each signal, and is locally unique for each signal at each interface. This signal could, and will most probably, occupy different time-slots at different interfaces.

7.2. Label Structure in SDH/SONET

The signaling protocol used to establish an SDH/SONET LSP must have specific information elements in it to map a label to the particular signal type that it represents, and to the position of that signal in the SONET/SDH multiplex. As we will see shortly, with a carefully chosen label structure, the label itself can be made to function as this information element.

In general, there are two ways to assign labels for signals between neighboring SDH/SONET LSRs. One way is for the labels to be allocated completely independently of any SDH/SONET semantics; e.g. labels could just be unstructured 16 or 32 bit numbers. In that case, in the absence of appropriate binding information, a label gives no visible information about the flow that it represents. From a management and debugging point of view, therefore, it becomes difficult to match a label with the corresponding signal, since, as we saw in [Section 7.1.1](#), the label is not coded in the SDH/SONET overhead of the signal.

Another way is to use the welldefined and finite structure of the SDH/SONET multiplexing tree to devise a signal numbering scheme that makes use of the multiplex as a naming tree, and assigns each multiplex entry a unique associated value. This allows the unique identification of each multiplex entry (signal) in terms of its type and position in the multiplex tree. By using this multiplex entry

value itself as the label, we automatically add SDH/SONET semantics to the label! Thus, simply by examining the label, one can now directly deduce the signal that it represents, as well as its position in the SDH/SONET multiplex. We refer to this as multiplex-based labeling. This is the idea that was incorporated in the GMPLS signaling specifications [7].

7.3. Signaling Elements

In the preceding sections, we defined the meaning of a SDH/SONET label and specified its structure. A question that arises naturally at this point is the following. In an LSP or connection setup request, how do we specify the signal for which we want to establish a path (and for which we desire a label)?

Clearly, information that is required to completely specify the desired signal and its characteristics must be transferred via the label distribution protocol, so that the switches along the path can be configured to correctly handle and switch the signal. This information is specified in three parts, each of which refers to a different network layer.

The first specifies the nature/type of the LSP or the desired SDH/SONET channel, in terms of the particular signal (or collection of signals) within the SDH/SONET multiplex that the LSP represents, and is used by all the nodes along the path of the LSP.

Bernstein, Mannie, Sharma Informational – January 2002

23

GMPLS based Control of SDH/SONET

July 2001

The second specifies the payload carried by the LSP or SDH/SONET channel, in terms of the termination and adaptation functions required at the end points, and is used by the source and destination nodes of the LSP.

The third specifies certain link selection constraints, which control, at each hop, the selection of the underlying link that is used to transport this LSP.

[8.](#) Summary and Conclusions

We provided a detailed account of the issues involved in applying MPLS-based control to TDM networks.

We began with a brief overview of MPLS and SDH/SONET networks, discussing current circuit establishment in TDM networks, and

arguing why SDH/SONET technologies will not be "outdated" in the foreseeable future. We then looked at MPLS applied to SDH/SONET networks, where we considered why such an application makes sense, and reviewed some MPLS terminology as applied to TDM networks. We then considered the two main areas of application of MPLS methods to TDM networks, namely routing and signaling. We considered in detail the switching capabilities of TDM equipment, and the requirement to learn about the protection capabilities of underlying links, and how these influence the available capacity advertisement in TDM networks. We focused briefly on path computation methods, pointing out that these were not subject to standardization. We then examined optical path provisioning or signaling, considering the issue of what constitutes an appropriate label for TDM circuits, how this label should be structured, and we focused on the importance of hierarchical label allocation in a TDM network. We then reviewed the signaling elements involved when setting up an optical TDM circuit, focusing on the nature of the LSP, the type of payload it carries, and the characteristics of the links that the LSP wishes to use at each hop along its path, for achieving a certain reliability.

9. Security Considerations

This draft raises no new security issues in the MPLS specifications.

10. References

- [1] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997

Bernstein, Mannie, Sharma Informational - January 2002

24

GMPLS based Control of SDH/SONET

July 2001

- [3] Synchronous Optical Network (SONET) Basic Description including Multiplex Structure, Rates, and Formats, ANSI T1.105-1995.
- [4] G.707, Network Node Interface for the Synchronous Digital Hierarchy (SDH), International Telecommunication Union, 03/96.
- [5] ANSI T1.105.01-1995, Synchronous Optical Network (SONET) Automatic Protection Switching, American National Standards institute.

- [6] G.841, Types and Characteristics of SDH Network Protection Architectures, ITU-T, 07/95.
- [7] Peter Ashwood-Smith and Lou Berger, Editors, "Generalized MPLS: Signaling Functional Description," Internet Draft, [draft-ietf-mpls-generalized-signaling-04.txt](#), Work in Progress, May 2001.
- [8] E. Mannie, Editor, "GMPLS Extensions for SONET and SDH Control", Internet Draft, [draft-ietf-ccamp-gmpls-sonet-sdh-01.txt](#), Work in Progress, June 2001.
- [9] E. Mannie, Greg Bernstein "Extensions to OSPF and IS-IS in support of MPLS for SDH/SONET Control", Internet Draft, Work in Progress, [draft-mannie-mpls-sdh-ospf-isis-00.txt](#), July 2000.

11. Acknowledgments

12. Author's Addresses

Greg Bernstein
Ciena Corporation
10480 Ridgeview Court
Cupertino, CA 94014
Phone: +1 510 573-2237
E-mail: greg@ciena.com

Eric Mannie
EBONE
Terhulpesteenweg 6A
1560 Hoeilaart - Belgium
Phone: +32 2 658 56 52
Mobile: +32 496 58 56 52
Fax: +32 2 658 51 18
E-mail: eric.mannie@ebone.com

Bernstein, Mannie, Sharma Informational - January 2002

25

GMPLS based Control of SDH/SONET

July 2001

Vishal Sharma

Metanoia, Inc.
335 Elan Village Lane, Unit 203
San Jose, CA 95134
Phone: +1 408 943 1794
Email: v.sharma@ieee.org

Full Copyright Statement

"Copyright (C) The Internet Society (date). All Rights Reserved.
This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into

Bernstein, Mannie, Sharma Informational - January 2002

27

GMPLS based Control of SDH/SONET

July 2001

