Generic Framing Procedure (GFP): The Catalyst for Efficient Data over Transport

Paul Bonenfant and Antonio Rodriguez-Moral, Photuris, Inc.

ABSTRACT

We discuss the newly defined Generic Framing Procedure (GFP) in the context of emerging nontraditional data over transport applications. Coupled with complementary efforts to define virtual concatenation, automatic link capacity adjustment schemes, and distributed control planes for transport networks, we contend that GFP serves as the catalyst for *efficient* and *standard* data over transport service offerings.

INTRODUCTION

While admittedly cliché, rumors of the impending death of synchronous optical network/synchronous digital hierarchy (SONET/SDH) and transport networking in general have been greatly exaggerated. A great many carriers worldwide count SONET/SDH as their transport infrastructure of choice, and have accumulated tremendous valuable experience operating, maintaining, and deriving revenue from these networks. A few recent developments are poised to further adapt SONET/SDH networks to the changing times.

Specifically, a new set of enhancements will make the transport network better suited to carrying data signals, driving its evolution toward increased efficiency and flexibility in supporting new data over transport services. These include the Generic Framing Procedure (GFP), virtual concatenation, and the Link Capacity Adjustment Scheme (LCAS). Virtual concatenation, GFP, and LCAS enable generalized mappings of variable-length multiprotocol packets into SONET/SDH, and allow for flexible and elastic data transport over SONET/SDH networks. When implemented in a combined fashion, these developments provide an efficient and standard means of carrying data signals over existing transport networks, offering the transport equivalent (roughly speaking) of statistical multiplexing while leveraging embedded networks and network management systems, and exploiting the transport carriers' comfort level with SONET/SDH bandwidth management. And, while initially geared toward existing transport networks, these developments can also be applied to emerging optical transport networks (OTNs).

In this article we briefly discuss both the current, as well as GFP and LCAS enabled approaches to data over transport. Complementing the GFP and LCAS transport data plane developments are the efforts to define a distributed control plane for transport networks — generalized multiprotocol label switching (GMPLS) in the Internet Engineering Task Force (IETF), and automatic switched transport networks (ASTNs) in the Interntational Telecommunication Union — Telecommunication Standardization Sector (ITU-T). We describe how these complementary developments render the transport network capable of increasingly efficient data over transport applications, breathing new life into carriers' existing SONET/SDH and emerging OTN infrastructures. We further outline the networking issues and considerations facing network planners charged with implementing this palette of transport network enhancements.

LIFE BEFORE GFP: PROPRIETARY MAPPINGS FOR DATA OVER TRANSPORT

A variety of approaches exist for mapping data signals over transport networks [1], and their use varies widely, depending on carrier business models, customer service offerings, and corresponding service requirements. With the possible exception of IP — so-called packet over SONET/SDH (POS) and IP over asynchronous transfer mode (ATM) — proprietary data over SONET/SDH mappings rule the day. This includes mappings for Ethernet and Gigabit Ethernet (GbE), Enterprise System Connect (ESCON), Fiber Connection (FICON), Fibre

 0163-6804/02/$17.00

Channel (FC), and storage area networking (SAN) applications in general. Proprietary mappings offer limited opportunity for interworking between equipment from different vendors and between carriers, and preclude the economies of scale in application-specific integrated circuit (ASIC) manufacturing inherent with standard techniques, since every proprietary mapping is a custom development effort.

Take GbE, for example. The tremendous acceptance of GbE in LAN and campus environments has created pressure on carriers to offer native Ethernet services at gigabit rates in the MAN/access environment. Applications such as transparent LAN extension interfacing at 1-Gb/s with Ethernet switches require a transport infrastructure that extends the reach of GbE signals, while maintaining the low cost of ownership expected from Ethernet network applications.

Multiple methods are currently used for transporting Ethernet signals over a MAN/access infrastructure. The IEEE 802.3 standard [2] ensures that GbE switches can be interconnected by dark fiber over distances up to 5 km using single-mode fiber (1000BASE-LX). Note that the IEEE distance limitations for GbE are usually very conservative. Many 1000BASE-LX vendors guarantee their products for much longer distances (10 km is a typical distance). Moreover, some vendors have developed fiber extenders allowing for distances up to 80 km.

The simplest method of deploying a transport infrastructure for distances that cannot be covered directly by GbE interfaces consists of using bit/byte interleaving or SONET/SDH framing to encapsulate Ethernet packets onto a wavelength or fiber. With this approach, any SONET/SDH or DWDM system with SONET/SDH transponders can transport Ethernet signals over metropolitan and regional distances. For example, a simple bit or byte interleaving device can take two GbE signals (whose line rate is actually 1.25 Gb/s due to the 8B/10B encoding) and multiplex them into a 2.5 Gb/s signal for native transport on a 2.5 Gb/s wavelength. The same scheme applies to eight GbE signals interleaved to a 10.0 Gb/s signal.

Another consists of using SONET/SDH framers at OC-48/STM-16 and OC-192/STM-64 rates, and performing some rate adaptation if statistical multiplexing of Ethernet frames is desired (e.g., by packing more than two GbE signals into an OC-48/STM-16 signal or more than eight GbE signals into an OC-192/STM-64 signal). Today, many proprietary implementations exist, and multivendor interoperability is not guaranteed. A mapping that is both standardized, such as GFP, and widely deployed is required to achieve such interoperability.

Efficient Data over Transport: GFP + Virtual Concatenation + LCAS

To combat the spread of proprietary solutions for mapping "data" signals into SONET/SDH frames that have emerged in the absence of a standard, the ITU-T and American National Standards Institute (ANSI) chartered a work effort in 1999 on data over SONET/SDH to promote vendor equipment and carrier interworking. This effort culminated in the development and specification of GFP [3]. The attractiveness of GFP lies in its combination with the co-developed virtual concatenation [4, 5] and link capacity adjustment scheme [6]. We provide a brief overview of each below.

The Generic Framing Procedure

GFP provides — for the first time — a standard means of mapping, in a very efficient way, a wide variety of data signals into SONET/SDH frames, enabling compliant equipment from different manufacturers to transport both traditional and nontraditional data signals over a SONET/SDH infrastructure.

GFP provides a generic mechanism to adapt traffic from higher-layer client signals over SONET/SDH, or even OTNs. Client signals may be protocol data unit (PDU)-oriented (e.g., IP/PPP or Ethernet MAC) or block-code-oriented constant-bit-rate streams such as ESCON, FICON, or FC. GFP consists of both common and client-specific aspects. Common aspects of GFP apply to all GFP adapted traffic, and include the definition of the basic signal structure for GFP frames, the types of GFP frames (client data frames and client management frames), and the frame-level processes common to all payloads that are mapped via GFP: frame delineation, frame multiplexing, client signal fail indication generation and propagation, and defect handling.

Currently, two modes of client signal adaptation are defined for GFP: a PDU-oriented adaptation mode, referred to as frame-mapped GFP (or GFP-F), and a block-code-oriented adaptation mode, referred to as transparent GFP (or GFP-T).

GFP-F is a type of GFP mapping in which a client signal frame is received and mapped in its entirety into one GFP frame. In this adaptation mode the client/GFP adaptation function may operate at the data link layer (or higher layer) of the client signal. Client PDU visibility is required. GFP-F mappings are currently defined for Ethernet MAC payloads and IP/PPP payloads.

GFP-T mapping provides a block-code-oriented adaptation mode in which the client/GFP adaptation function operates on the coded character stream rather than the incoming client PDUs. Transparent GFP provides a way for a number of client data characters to be mapped into efficient block codes for transport within a GFP frame. With this type of mapping block-coded client characters are decoded and then mapped into a fixed-length GFP frame, and may be transmitted immediately without waiting for the reception of an entire client data frame. This allows for some network applications — LAN/SAN extension applications — wherein client equipment running protocols that require very low latency, such as FC, ESCON, and FICON, may be interconnected via GFP in their native mode. GFP-T mappings are currently defined for FC, ESCON, FICON, and GbE.

The mapping of the framed payloads into an SDH path is specified in ITU-T Recommenda-

GFP provides a generic mechanism to adapt traffic from higher-layer client signals over SONET/SDH, or even OTN transport networks. Client signals may be PDU-oriented, or block-code oriented constant bit rate streams such ESCON, FICON, or Fiber Channel.

Virtual concatenation breaks the integral payload into individual channels, transports each channel separately, and then recombines them into a contiguous bandwidth at the endpoint of the transmission — in essence a form of inverse multiplexing.

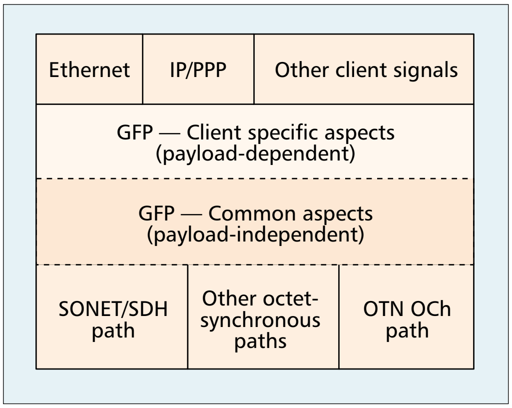

■ **Figure 1.** *GFP's relationship to payloads and SONET paths [3].*

tion G.707 [5]. The mapping of the framed payloads into an OTN optical channel (OCh) path is specified in ITU-T Recommendation G.709 [7].

Figure 1 illustrates the relationship between the higher-layer payloads, GFP, and SONET/SDH or OTN paths.

VIRTUAL CONCATENATION

The GFP effort for SONET/SDH leveraged a parallel activity to standardize virtual concatenation of SONET/SDH (and later OTN) paths. Virtual concatenation allows for relaxation of the "rigidity" of SONET/SDH payload bit rates, originally designed based on the digital hierarchy defined for the telephone (voice) network. Hence, in combination with virtual concatenation, GFP will allow the efficient mapping of a wide variety of data signals over SONET/SDH. Table 1 and Fig. 2 provide a sample listing of the target virtually concatenated SONET/SDH path sizes for various client signal protocols.

In the context of SONET/SDH, virtual concatenation is intended to support the transport of payloads that do not fit efficiently into the standard set of SONET/SDH payloads. Virtual concatenation breaks the integral payload into individual channels, transports each channel separately, and then recombines them into a contiguous bandwidth at the endpoint of the transmission — in essence a form of inverse multiplexing. This type of concatenation requires concatenation functionality only at the path termination equipment.

Using SONET payload types and terminology [4] for the sake of example, 10 Mb/s Ethernet could be carried across a VT1.5-7v link instead of using up a full STS-1 link. Similarly, a (near line rate) 100 Mb/s Ethernet link could be carried across an STS-1-2v link instead of an STS-3c link, or a VT1.5-64v, which provides a payload of 102.4 Mb/s, could be used. See Fig. 2 for the increased bandwidth efficiency enabled by virtual concatenation for 10, 100, and 1000 Mb/s Ethernet signals.

As a result, an OC-48 (2.5 Gb/s) link can be more efficiently used to simultaneously transport data and voice traffic: 2 GbE signals can fit into 2 STS-1-21v virtually concatenated groups, which leaves 6 STS-1s (290 Mb/s) for other applications, including voice traffic.

GFP can employ virtual concatenation to enable efficient mapping of client signals, substantially improving the bandwidth efficiency of SONET/SDH infrastructures for transporting data signals to levels close to 100 percent, as discussed above.

The flexibility and bandwidth efficiency provided by a combination of GFP and virtual concatenation can be exploited in so-called Multi-Service Provisioning Platform (MSPP) systems at the edge of the network, and in SONET/SDH and DWDM/OTN aggregation and switching equipment in regional network segments.

LINK CAPACITY ADJUSTMENT SCHEME

LCAS further enhances virtual concatenation by enabling increase or decrease of capacity of virtually concatenated links without interrupting the traffic flow. In essence, LCAS endows SONET/SDH and OTN with the ability to automatically "tune" the bandwidth of virtually concatenated signals.

LCAS Applications — This fine management of the bandwidth of virtually concatenated signals is particularly attractive for efficient transport of data services that are inherently of variable bit rates. For example, consider the transport of a partially filled GbE signal. Although its nominal bandwidth rate is 1 Gb/s, the instantaneous rate can typically be only 200–300 Mb/s. Allocating 1 Gb/s of continuous bandwidth to this GbE signal (as is done in pure transport applications) wastes, on the average, 70 percent of network bandwidth, whereas the use of virtual concatenation can increase bandwidth efficiency. The amount of bandwidth to assign to the virtually concatenated signal can be determined by balancing average and peak bandwidth. If average bandwidth is used, the network elements on both ends of the GbE path must provide enough buffering for flow control. If peak bandwidth is used, the virtually concatenated signal will require more bandwidth and its average utilization will be lower. For example, again using SONET terminology, an STS-1-7v that provides a bandwidth of 338.688 Mb/s could be used if the average bandwidth is considered. The network elements on both ends

Client payload (unencoded bandwidth/line rate)	SONET/SDH path (bandwidth)
ESCON (160/200 Mb/s)	STS-1-4v/VC-3-4v (196 Mb/s)
Fibre Channel — FC100 (850/1062.5 Mb/s)	STS-3c-6v/VC-4-6v (900 Mb/s)
FICON (850/1062.5 Mb/s)	STS-3c-6v/VC-4-6v (900 Mb/s)
Gigabit Ethernet (1000/1250 Mb/s)	STS-3c-7v/VC-4-7v (1050 Mb/s)

■ **Table 1.** *Virtually-concatenated path size for 8B/10B client signals.*

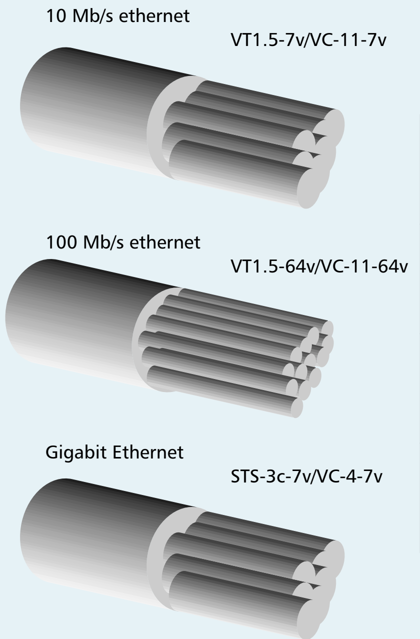

Data signal	SONET/SDH payload mapping and bandwidth efficiency	SONET/SDH with virtual concatenation payload mapping and bandwidth efficiency
Ethernet (10 Mb/s)	STS-1/VC-3 — 21%	VT1.5-7v/VC-11-7v — 89%
Fast Ethernet (100 Mb/s)	STS-3c/VC-4 — 67%	VT1.5-64v/VC-11-64v — 98%
Gigabit Ethernet (1000 Mb/s)	STS-48c/VC-4-16c — 42%	STS-3c-7v/VC-4-7v — 95% STS-1-21v/VC-3-21v — 98%

Figure 2. *Virtual concatenation for increased bandwidth efficiency.*

A proper combination of these LCAS applications could help provide carriers employing a SONET/SDH infrastructure to turn up bandwidth in increments that closely match the purported abilities of native-Ethernet service providers.

of the GbE path would then have to provide enough buffering and/or local port flow control to handle instantaneous transmission rates over this bandwidth.

LCAS was designed with this type of application in mind. In LCAS operation, an initial group of channels is assigned to a virtual concatenation group (channels are referred to as *members* of the group) for transport of variable rate data signals. When bandwidth must be increased or decreased — this decision is the responsibility of a bandwidth management application, which can monitor the source rate to dynamically adjust the bandwidth the network will provide for transporting the signal — LCAS adds/removes individual channels to/from the virtual concatenation group. In the previous example, the initial virtual concatenated signal could be an STS-1-7v, which, when implemented with LCAS, could in principle expand as large as STS-1-21v. Note that a sound bandwidth management application will avoid excessive LCAS bandwidth add/remove operations, and will probably use rate monitoring and thresholds to trigger such operations.

Note that in the current transparent mapping of FC, ESCON, and FICON into GFP frames, idles are not removed. Consequently, transport of these signals always requires full bandwidth, and LCAS provides little advantage.

A second — and somewhat complementary — application for LCAS relates to service survivability. For some data services, transport layer protection switching of the entire bandwidth may no longer be needed or required on all signals by carriers and service providers. Using LCAS, it is possible to handle failures of individual members of the concatenated signal by simply reducing the capacity and providing a minimum bandwidth. This is possible because the individual members of a virtual concatenation group can be diversely routed.

Still another related application deals with load balancing of data signals. This LCAS application also uses the diverse routing of individual members of a virtual concatenation group to split the traffic load between two points in a network. Some Ethernet applications, for example, use this principle for layer 2 failover restoration techniques.

A proper combination of these LCAS applications could help provide carriers employing a SONET/SDH infrastructure to turn up bandwidth in increments that closely match the purported abilities of native Ethernet service providers.

LCAS Operation — When LCAS is triggered at the source node of a virtual concatenation group link, by either the instance of the distributed control plane running in that node or the provisioning application running in the EMS/NMS, it exchanges signaling messages with the remote end to synchronize the addition/removal of SONET/SDH channels. This synchronization is accomplished by exchanging multiframed LCAS control packets.

To ensure that the capacity adjustments to a virtual concatenation group are hitless, the two ends of the link must agree on precisely when the virtual concatenation group transitions to a new payload in which new group members have been added or previous members removed. This coordination unavoidably requires hardware-level synchronization, that is, bit-interval-accurate indications to the SONET/SDH payload mappers as to when to begin/stop inserting/extracting a payload to/from a virtual concatenation group member. In LCAS, this indication is provided by information in the member path

The use of GFP renders SONET/SDH (or OTNs) both a versatile and flexible transport infrastructure, while the combination of LCAS (and its associated network management application or control plane) and virtual concatenation renders SONET/SDH (or OTNs) an elastic transport infrastructure.

overhead, in particular in the H4 byte for high order virtual concatenation (HOVC), which is used by the payload mappers at the two ends of the link to coordinate sequence numbers and exchange LCAS control packets. Coordinating and synchronizing member additions and deletions involves, in part, control of these sequence numbers. Extending H4 to control member additions and deletions is therefore consistent with the sequence number functionality of the H4 byte. The state machine/protocol specifications for LCAS also support resilient management of a virtual concatenation group. The LCAS extensions to the functionality of H4 support sequence number reassignment through member exclusions and inclusions required to respond to member path failure and recovery.

In summary, the use of GFP renders SONET/SDH (or OTNs) both a *versatile* and *flexible* transport infrastructure, while the combination of LCAS (and its associated network management application or control plane) and virtual concatenation renders SONET/SDH (or OTNs) an *elastic* transport infrastructure.

A Powerful Solution for Data over Transport: GFP/LCAS + GMPLS/ASTN

LCAS has been designed as a natural extension to virtual concatenation that can operate in both "traditional" transport networks (i.e., those in which the setup and release of individual channel connections is performed in the management plane by a centralized EMS/NMS) or next-generation transport networks (i.e., those that incorporate a distributed control plane tasked with, among other things, path setup and teardown operations). In LCAS it is assumed that in cases of capacity initiation, increases ,or decreases, the responsibility for the construction or destruction of the end-to-end path — before (or after) it is assigned to (or deleted from) a virtual concatenation group by LCAS — is outside of the LCAS process itself. LCAS is concerned only with signaling related to virtual concatenation group operations, such as addition or removal of a member, or (logically) renumbering the members in a virtual concatenation group.

With the advent of distributed control planes for transport networks (e.g., GMPLS [8] or ASTN [9]), it is envisaged that the protocols constituting those control planes will be used for topology discovery and signaling the end-to-end setup and teardown of member paths — LCAS would then add (or delete) the new member to (or from) the virtual concatenation group. Therefore, it is important that such control planes — currently being standardized — incorporate the appropriate extensions to operate with LCAS [10]. These extensions must preserve the clear separation of responsibilities between LCAS and the control plane protocols, while at the same time preserving the expected end-to-end behavior of the transport network for these dynamic bandwidth applications.

Figure 3 shows a possible high-level (abstract) architecture for GMPLS and LCAS. The figure also illustrates how LCAS can be deployed in the context of an EMS/NMS-based provisioning model in the absence of a distributed control plane, since the interface to LCAS consists of the same set of operations.

In the case of a GMPLS-based control plane that uses Resource Reservation Protocol (RSVP), the sequence for the interaction of RSVP and LCAS when increasing the bandwidth of an existing virtual concatenation group would be as follows (roughly speaking— this is an oversimplification for the sake of example):

- The existing virtual concatenation group is an STS-1-3v.
- The bandwidth management application (e.g., residing within the source node) decides that a bandwidth increase is needed, and requests that an additional STS-1 (which may or may not exist between the endpoints) be added to the virtual concatenation group — *Increase Bandwidth* (*GID, Delta*) operation.
- The GMPLS control plane would check whether there is some existing label switched path (LSP) in the path database that satisfies the selection criteria. If this is the case, no LSP setup is required, and the result of the invocation would simply be the selected LSP. Otherwise, the operation would result in a computation of the path (done by the RTA algorithm) and further LSP setup by RSVP.
- If LSP setup is required, RSVP will perform the required signaling (*Path* and *Resv* messages) in order to set up the end-to-end STS-1 path, instantiating along the way the required crossconnect points in the network.
- Once the additional STS-1 path is set up, LCAS is triggered at the source node with a request to add a new STS-1 member to the virtual concatenation group.
- LCAS signaling (through LCAS control packets) takes place between the two end nodes; the effect is the synchronization of the payload mappers at both ends, and an expansion of the virtual concatenation group to STS-1-4v.

The combination of GFP, virtual concatenation, and LCAS with a distributed control plane enables the deployment of powerful dynamic bandwidth applications in which the required intelligence for topology and resource discovery, constraint-based path computation, path setup and teardown, and bandwidth usage monitoring and bandwidth adjustment triggering all resides within the network. This opens up the possibility of reducing operating expenses, and can also be a critical element for future internetworking applications in which the dynamic bandwidth adjustment application runs between client devices at the edge of the network.

Network Application Issues

Some issues worth mentioning related to the network application of data over transport enabled by the combination of GFP and LCAS include: interworking with the GMPLS/ASTN protocols, since the combination of LCAS and a control plane must ensure the correct behavior of the network with respect to path setup and

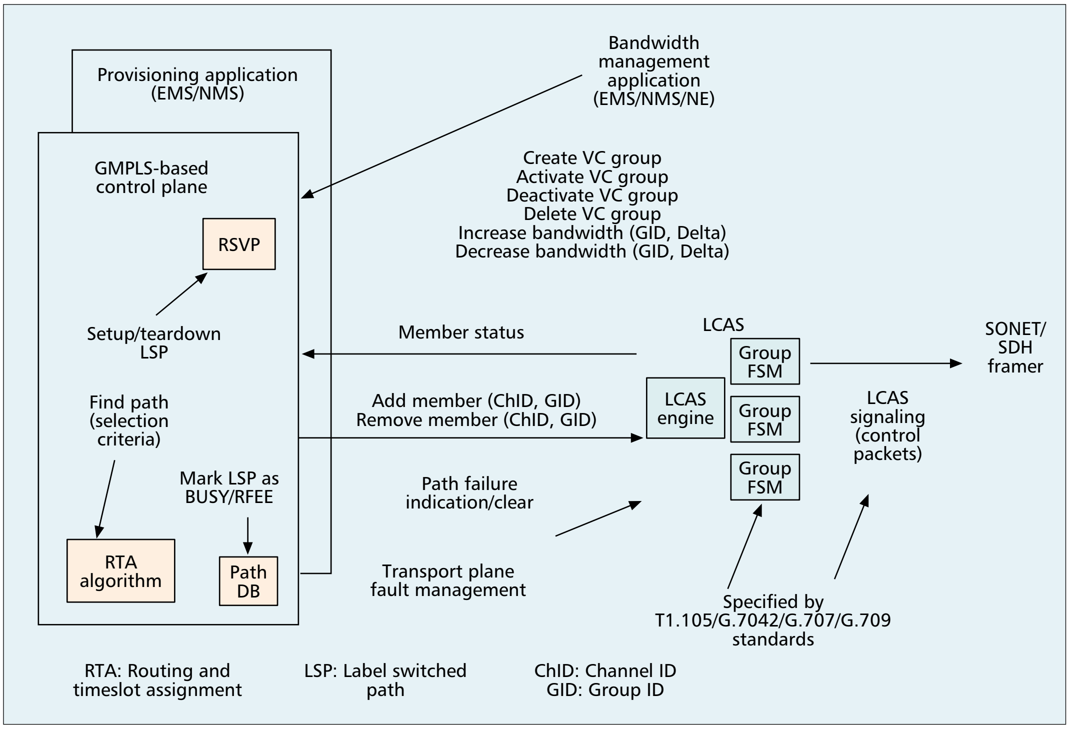

Figure 3. *High-level GMPLS/LCAS architecture.*

Since the bandwidth assigned to a given signal can be increased or decreased by LCAS, the provisioning and management of bandwidth in a given SONET/SDH link, in an efficient way, becomes a non-trivial task.

teardown and associated capacity adjustment operations; protection layer interworking, since transport layer protection may be provided on the signals constituting an LCAS virtual concatenation group; the use of GFP as an alternative mapping for 10 GbE signals over SONET/SDH; and the concern that GFP might evolve into a network layer in its own right.

The Interaction of GFP/LCAS with GMPLS/ASTN Control Planes

The powerful combination of GFP, LCAS and a distributed control plane such as GMPLS or ASTN still has some unresolved challenges that will likely result in some extensions to the GMPLS/ASTN protocols. These include:

- ASTN/GMPLS will have to support connection bandwidth modification between virtual concatenated endpoints that do not both support LCAS.
- ASTN/GMPLS will have to provide mechanisms that enable several LSPs (i.e., circuits or member channels) to be combined to form a larger virtually concatenated group. This allows dealing with the non-co-routing of SONET/SDH paths.
- A GMPLS/ASTN control plane maintains states in end systems (e.g., at the O-UNI) and intermediate systems. Thus, bandwidth modifications introduced by LCAS (e.g., by the autonomous removal of failed members) must be reflected in the control plane to avoid inconsistencies. This requires coordination between an LCAS engine and the control plane.
- Some extensions to the network resource information advertised by a GMPLS/ASTN control plane could be required for GFP and LCAS signals. For example, link reliability and maximum differential delay between members would be necessary when computing a path for adding a new member to an existing virtual concatenation group (e.g., in cases where diverse routing of virtual concatenation group members is used).
- Although the operation of LCAS is unidirectional, a bandwidth management application might request the setup of bidirectional virtual concatenation groups. A bidirectional LSP is set up in such circumstances with source and sink nodes on either side. It might be possible to set up two unidirectional LSPs, but it is rather cumbersome for signaling from both sides. Existing extensions to GMPLS protocols for transport networks now allow setting up bidirectional LSPs. Coordination between a control plane that uses this functionality and LCAS signaling may be required.

A separate issue concerns the design of bandwidth management applications for LCAS-enabled networks. Since the bandwidth assigned to a given signal can be increased or decreased by LCAS, the provisioning and management of bandwidth in a given SONET/SDH link, in an efficient way, becomes a nontrivial task. This issue — while neither unique to LCAS nor new (as it was addressed some time ago for resource management of B-ISDN networks) — promises to be important to dynamic bandwidth management in networks employing both LCAS and GMPLS/ASTN control planes.

We can also consider greenfield network applications where a GMPLS control plane — with appropriate extensions — would be responsible for dynamic bandwidth adjustment, where-

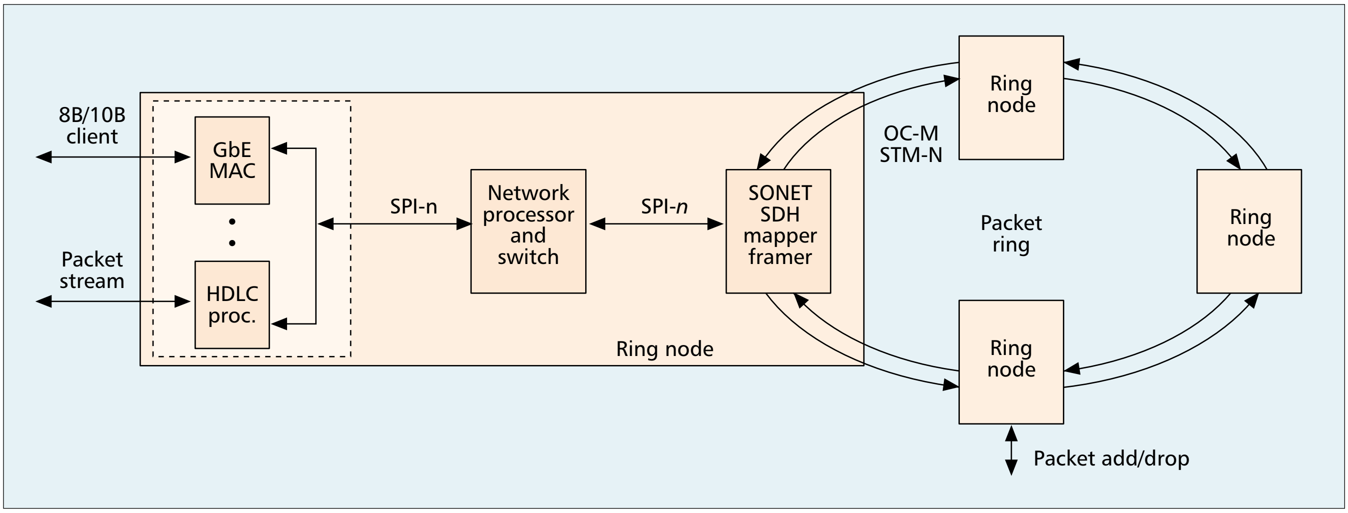

Figure 4. *Candidate architecture for GFP: resilient packet rings [12].*

in the transport network could be deployed without LCAS and still provide some form of dynamic capacity adjustment. For example, this functionality is part of RSVP-TE for packet networks and could be extended to SONET/SDH and OTN networks as well. To achieve the hitless nature of LCAS, however — required for traffic-engineered virtual private line services with strict service level contracts — some type of bit-accurate synchronization of the payload mappers at both ends (which implies a hardware implementation) would be required. Without such capability, an extended GMPLS control plane could be used for some limited applications, such as Internet service provider backbones, where hitless operation may not be necessary.

Transport Layer Protection and Diverse Routing of LCAS Members

Some may consider the diverse routing of members of a virtual concatenation group a "survivability option" in LCAS. Diverse routing of virtual group members can be used to satisfy the requirement that a virtual concatenation group guarantee some minimum bandwidth in the event of a network failure. To achieve this goal, virtual concatenation group members can be diversely routed (e.g., split into two subgroups) such that they traverse disjoint paths.

In principle, a virtual concatenation group (or each of its component group members) could be protected by standard SONET/SDH protection schemes. However, the dynamic characteristics of an LCAS enabled virtual concatenation group, especially when diverse routing is employed, would seem to dictate that SONET/SDH protection be disabled for these signals.

GFP as a Mapping for 10-Gigabit Ethernet over SONET/SDH

GFP, as part of its frame-mapped mode, also defines a mapping for 10GbE signals over SONET/SDH or OTN paths. This implies that network operators will have an alternative to the WAN physical layer signal (PHY) defined in IEEE 802.3ae [2] to deploy 10GbE services over a SONET/SDH infrastructure.

In 10GbE, the WAN PHY differs from the LAN PHY by the inclusion of a simplified SONET/SDH framer in the WAN interface sublayer (WIS) [1]. Since the line rate of a SONET OC-192/SDH STM-64 signal is within a few percent of 10.0 Gb/s, it is feasible to implement a MAC layer able to operate with a LAN PHY at 10.0 Gb/s or a WAN PHY at the SONET/SDH payload rate (9.584640 Gb/s for STS-192c). However, there are some slight differences between a full SONET/SDH layer and 10GbE WAN PHY. For example, SONET/SDH systems use synchronized high-accuracy stratum clocks to form a synchronous clock hierarchy. SONET/SDH regenerators recreate the signals moving from one SONET/SDH segment to the next by using this synchronous clock hierarchy. On the other hand, the WAN PHY operates like any other asynchronous network interface, retaining the asynchronous nature of Ethernet. Each link is separated from the clock domain of the next link by a store-and-forward buffer device implemented in a router or layer 2 switch. Furthermore, the 10GbE WAN PHY does not use the entire SONET/SDH overhead. Certain overhead functions considered unnecessary are simply not used.

These considerations drive the need for an interworking function between the 10GbE WAN PHY and existing SONET/SDH infrastructures. However, this interworking function would not be required for the GFP mapping of 10GbE signals, since GFP does not require any modification to a SONET/SDH infrastructure. For some carriers, this could represent a clear advantage of the GFP mapping.

GFP as a Network Layer

Originally conceived for "simple" multiprotocol mappings in point-to-point applications, GFP has raised some concerns regarding its possible complexity. This is true in particular for ring topologies, either with the use of GFP as the scheme for resilient packet ring (RPR) [11] to SONET/SDH mappings or for native GFP ring applications.

The RPR effort, as it relates to GFP, has also led to some concerns. In brief, RPR presents an alternative to traditional SONET/SDH time-division multiplexing (TDM) rings by combining the resilient nature of ring topologies with statistical multiplexing and QoS capabilities of a packet-optimized MAC protocol. While the protection mechanism is handled at the packet layer by RPR, the individual RPR spans can be composed of SONET/SDH links employing an RPR over SONET/SDH mapping, as shown in Fig. 4.

To support mapping of RPR frames into SONET/SDH via GFP, and of RPR or native GFP ring topologies, the ITU-T has reserved a code point for RPR frames as a client signal for GFP, and an (optional) extension header for ring topologies (for further study in the current GFP standard).

These extensions raised some concerns that as the functionality associated with GFP increased and its use expanded beyond simple point-to-point applications, it could become a network layer in its own right, requiring end-to-end network management and control, similar to the network layers above (e.g., Ethernet) or below (e.g., SONET/SDH, OTN). It has become clear, however, that the true value of GFP lies in its multiprotocol mapping capabilities, and network operators need not be concerned with GFP as a network layer. The networking functionality will continue to reside either at the layer above or below the GFP mapping.

SUMMARY AND CONCLUSIONS

A convergence of developments under the guise of data over SONET/SDH — GFP, virtual concatenation, and LCAS — will render existing transport networks capable of increasingly efficient data over transport, leveraging existing infrastructures to offer new services. This is music to the ears of capital expenditure constrained carriers in a difficult economic environment, since next-generation solutions must lessen the degree to which a "paradigm shift" is required to support new services. In other words, they must allow for the creation of new revenue-generating services that, to the extent possible, leverage the existing network and have minimal impact on operating procedures rather than require a wholesale replacement of a carrier's network infrastructure.

Solutions enabled by GFP and LCAS will be deployed in transport networks built both with and without distributed control planes. The GMPLS/ASTN control planes for next-generation transport networks will incorporate the appropriate extensions to handle GFP signals, virtually concatenated structures, and LCAS operation. Using virtual concatenation, neither GFP nor LCAS require end-to-end upgrades to the embedded base of network equipment; they can and will be deployed only at the ingress and egress of a carrier's transport network. This is a critical factor for carrier and service provider acceptance of GFP-based solutions, since they enable new service offerings while leveraging the existing network infrastructure.

Combined, GFP and LCAS offer an attractive option for carrying data protocols over transport networks, and may present a compelling alternative to the use of ATM and MPLS for transport-oriented statistical multiplexing gain. As the enabler for standard (and, we envision, widely deployed) mappings for data over SONET/SDH and OTNs, GFP will indeed serve as the catalyst for efficient data over transport.

Combined, GFP and LCAS offer an attractive option for carrying data protocols over transport networks, and may present a compelling alternative to the use of ATM and MPLS for transport-oriented statistical multiplexing gain.

REFERENCES

[1] P. Bonenfant and A. Rodriguez-Moral, "Framing Techniques for IP over Fiber," *IEEE Network*, Jul./Aug. 2001, pp. 12–18.

[2] IEEE 802.3 CSMA/CD (Ethernet): http://grouper.ieee.org/groups/802/3

[3] ITU-T Draft Rec. G.7041, "Generic Framing Procedure (GFP)," 2001.

[4] ANSI T1.105, "Synchronous Optical Network (SONET) — Basic Description Including Multiplex Structure, Rates and Formats," 2001.

[5] ITU-T Rec. G.707, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)," 2000.

[6] ITU-T Draft Rec. G.7042, "Link Capacity Adjustment Scheme (LCAS) for Virtual Concatenated Signals," 2001.

[7] ITU-T Rec. G.709, "Network Node Interface for the Optical Transport Network (OTN)," 2001.

[8] E. Mannie *et al.*, Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," IETF draft, <draft-ietf-ccamp-gmpls-architecture>, Mar. 2002, work in progress.

[9] ITU-T Draft Rec. G.8070, "Requirements for Automatic Switched Transport Networks (ASTN)," 2001.

[10] A. Rodriguez Moral *et al.*, "Potential Issues and Questions of Clarification on LCAS," ANSI T1X1.5/2001–207, Sept. 2001.

[11] IEEE 802.17 Resilient Packet Ring Working Group: http://www.ieee802.org/17/

[12] M. Scholten, "Overview of Generic Framing Procedure (GFP)," OIF 2001-172, Apr. 2001.

NOTES

ANSI T1X1.5 documents are available from: http://www.t1.org/t1x1/_x1-grid.htm

IETF drafts and RFCs are available from: http://www.ietf.org/

ITU-T documents are available from: http://www.itu.int/ITU-T/index.html

OIF documents are available from: http://www.oiforum.com

BIOGRAPHIES

PAUL BONENFANT [M'89] (paul@photuris.com) serves as chief architect at Photuris, Inc. His experience spans SONET/SDH, WDM, and optical networking transport architecture, product evolution planning, network survivability, and associated global standards development. Before joining Photuris he served as a business development manager for Mergers and Acquisitions in Lucent's Optical Networking Group, and led a group responsible for optical network architecture evolution at Bell Laboratories. Prior to joining Lucent, he led requirements and standards development for SONET/SDH self-healing rings and dense WDM systems at Bell Communications Research (Bellcore, now Telcordia). He holds a B.S. degree in engineering and applied science, and an M.S. degree in electrical engineering, both from the California Institute of Technology, Pasadena.

ANTONIO RODRIGUEZ-MORAL [M'91] (arodmor@photuris.com) leads the definition of network architectures and next-generation systems at Photuris, Inc. Before joining Photuris, he was at Bell Laboratories, Lucent Technologies, where he was involved in the definition and analysis of network architectures, product evolution planning, and strategy for next-generation DWDM networks, and in research and development of network management systems and high-speed IP networks. Prior to joining Bell Laboratories he was with AT&T Network Systems in Europe, where he led several research and development projects for SDH and passive optical networks. He received his M.S. degree in electrical engineering from the Technical University of Madrid, Spain.