Data Transport Applications Using GFP

Mike Scholten and Zhenyu Zhu, AMCC
Enrique Hernandez-Valencia, Lucent
John Hawkins, Nortel Networks

ABSTRACT

An emerging generic framing procedure (GFP) standard defines a new data encapsulation protocol designed to accept and transport multiple protocols over metropolitan, storage, and wide area networks. Why is another data encapsulation method needed for transporting data traffic over WANs? How does GFP encapsulate data for transport? This article provides an introduction to data transport applications, reviews existing data transport options (e.g., ATM and HDLC), introduces GFP as a new data transport technology, and compares GFP transport to existing popular alternatives. Companion articles provide more in-depth understanding of the actual GFP encapsulation method, based on the current GFP standard (ITU-T G.7041).

INTRODUCTION

It has been widely noted that data traffic has overtaken voice traffic in terms of WAN network bandwidth utilization. Just as WAN line rates have continued to increase (from 622 Mb/s to 2.5 and 10 Gb/s, with 40 Gb/s emerging), datacom data rates are also increasing. Gigabit Ethernet (GbE) is being widely deployed, 10 10GbE products are being introduced, Fibre Channel is already in service at 1.0625 Gb/s and expected at 2.125 Gb/s, and 10 Gb/s Fibre Channel (10GFC) standards are in development. Not only is data traffic volume increasing, but traffic rates are also increasing as high-rate services are offered, such as digital video broadcast and storage area networking (SAN). Demand for access to these high-speed services across wide areas is rapidly growing as companies become increasingly geographically diverse, e-commerce data exchange proliferates, and direct access to offsite data storage is required.

Consequently, not only are new high-speed data services being deployed, but they are being extended across MANs and WANs. GFP provides a generic mechanism for transporting data protocols across these networks.

DATA OVER SONET/SDH

Data protocols are widely transported today across synchronous optical network/synchronous digital hierarchy (SONET/SDH) telecommunications networks using packet over SONET (POS), frame relay, and asynchronous transfer mode (ATM). Efficient transport of Ethernet over SONET/SDH (EoS) is proposed using either X.86 EoS/LAPS [1] or International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) G.7041 generic framing procedure (GFP) [2]. IP is transported using either Internet Engineering Task Force (IETF) RFC 2615 POS [3] or X.85 IP over SDH [4]. ATM may be used to transport a variety of protocols, and provides routing and prioritization through cell headers. GFP uniquely offers a mechanism for transparent transport of various data protocols based on 8B/10B coding. Figure 1 illustrates various protocol stacks for transporting data traffic over the public network infrastructure.

The following discussion focuses on applications in which layer 2 or 3 protocols (Ethernet frames or IP/PPP packets) are encapsulated for transport over a SONET/SDH physical layer network, primarily to implement router- or switch-based networks. In these applications, packets or frames are extracted from the payload of the received physical layer signal, routed or switched based on addressing information contained in the packets or frames, and delivered to output line interfaces for reencapsulation into the payload of outgoing physical layer signals.

ATM TRANSPORT

As the most widely deployed broadband transport technology for data transport, ATM is a natural candidate as the unifying transport solution for various types of data traffic over the public transport infrastructure. ATM adaptation creates a fixed size cell stream that relies on implicit information about the cell size (53 bytes) and explicit header error control (HEC) cyclic redundancy check (CRC) in the ATM cell head-

0163-6804/02/$17.00 © 2002 IEEE

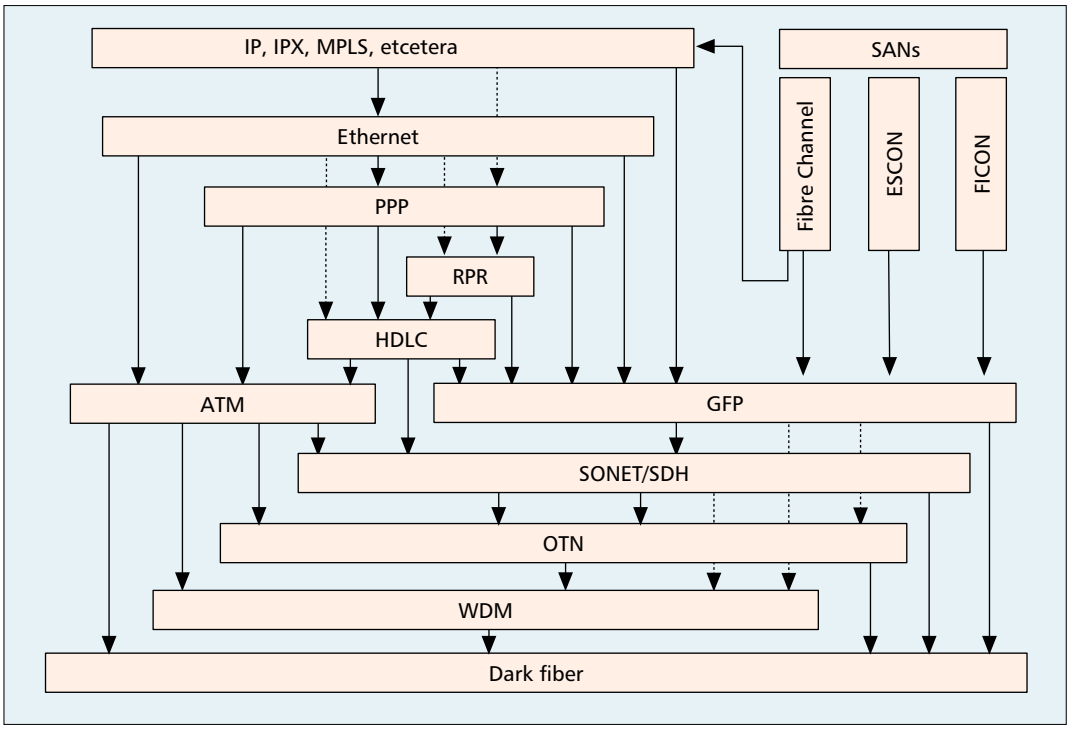

Figure 1. *Protocol stacks for various data transport applications*

ATM can be overkill for simple point-to-point traffic transport among broadband switches. As the volume of IP traffic increases, the point-to-point traffic among core IP routers becomes high enough to make the direct use of the entire SONET/SDH line without further bandwidth partitioning more efficient and desirable.

ers for the purpose of cell delineation. A header CRC hunting mechanism is employed by the receiver to extract the ATM cells from the bit/byte synchronous stream. Starting from the assumed cell boundary, the ATM receiver compares its computed HEC value for the assumed ATM cell header against the HEC value indicated by the assumed HEC field. Cell stream delineation is declared after positive validations of the incoming HEC fields of a few consecutive ATM cells. The fixed ATM cell size significantly simplifies delineation of cells within a byte stream by providing a strict bound on the duration of the hunting procedure. Protocol data units (PDUs) themselves are typically adapted to ATM via ATM adaptation layer 5 (AAL5) via a simple packet segmentation procedure. Figure 2a shows the resulting byte stream of PDUs adapted to the public transport infrastructure using ATM/AAL5.

ATM is more than just a client adaptation mechanism; it provides a full complement of switching, multiplexing, and networking functions. ATM supports sophisticated traffic engineering, flexible quality of service (QoS)-aware routing that provides granular partitioning of SONET/SDH bandwidth, and multiservice integration. In addition, emulation capabilities are specified to support a wide variety of transport services. For example, both routed and bridged Ethernet transport services [5, 6] are already defined, and deployed, making it uniquely suitable as a multiservice (voice, data, video) transport platform. Commercial component interconnect interfaces such as UTOPIA (ATM Forum) and SPI-3/SPI-4 (OIF) are readily available to guarantee a large pool of inexpensive system components and ease of integration. Given its sophisticated traffic management mechanisms and technological maturity, ATM has been the technology of choice for fine-grained traffic engineering and QoS support of data traffic over a SONET/SDH infrastructure, particularly at the edge of the core transport network.

However, ATM can be overkill for simple point-to-point traffic transport among broadband switches. As the volume of IP traffic increases, the point-to-point traffic among core IP routers becomes high enough to make the direct use of the entire SONET/SDH line without further bandwidth partitioning more efficient and desirable. In such scenarios, the transport overhead generated by ATM/AAL-5 adaptation procedure and by the partial fill of ATM cells make the transport of best effort IP traffic over ATM less bandwidth-efficient than IP directly over SONET/SDH or over an optical channel. Better traffic engineering and strict/differential QoS capabilities are being defined for IP traffic under the multiprotocol label switching framework [7]. Switches may further decrease the need for an intervening ATM layer.

IP over SONET/SDH

IP traffic is one of the most dominant forms of data communications today. IP traffic is frequently transported across WAN networks through either frame relay (FRF.14 [8]), PPP/HDLC [9], PoS (IETF RFC 2615 [3]), or IP over SDH using LAPS (ITU-T X.85 [4]).

Figure 2b illustrates the format of HDLC encapsulated signals.

With frame relay and POS in widespread use, why transition to IP/PPP over GFP?

Any transport based on byte-aligned HDLC relies on byte-stuffed transparency processing, described in IETF RFC 1662. In these HDLC applications, any data bytes within each packet matching either HDLC flag or escape characters

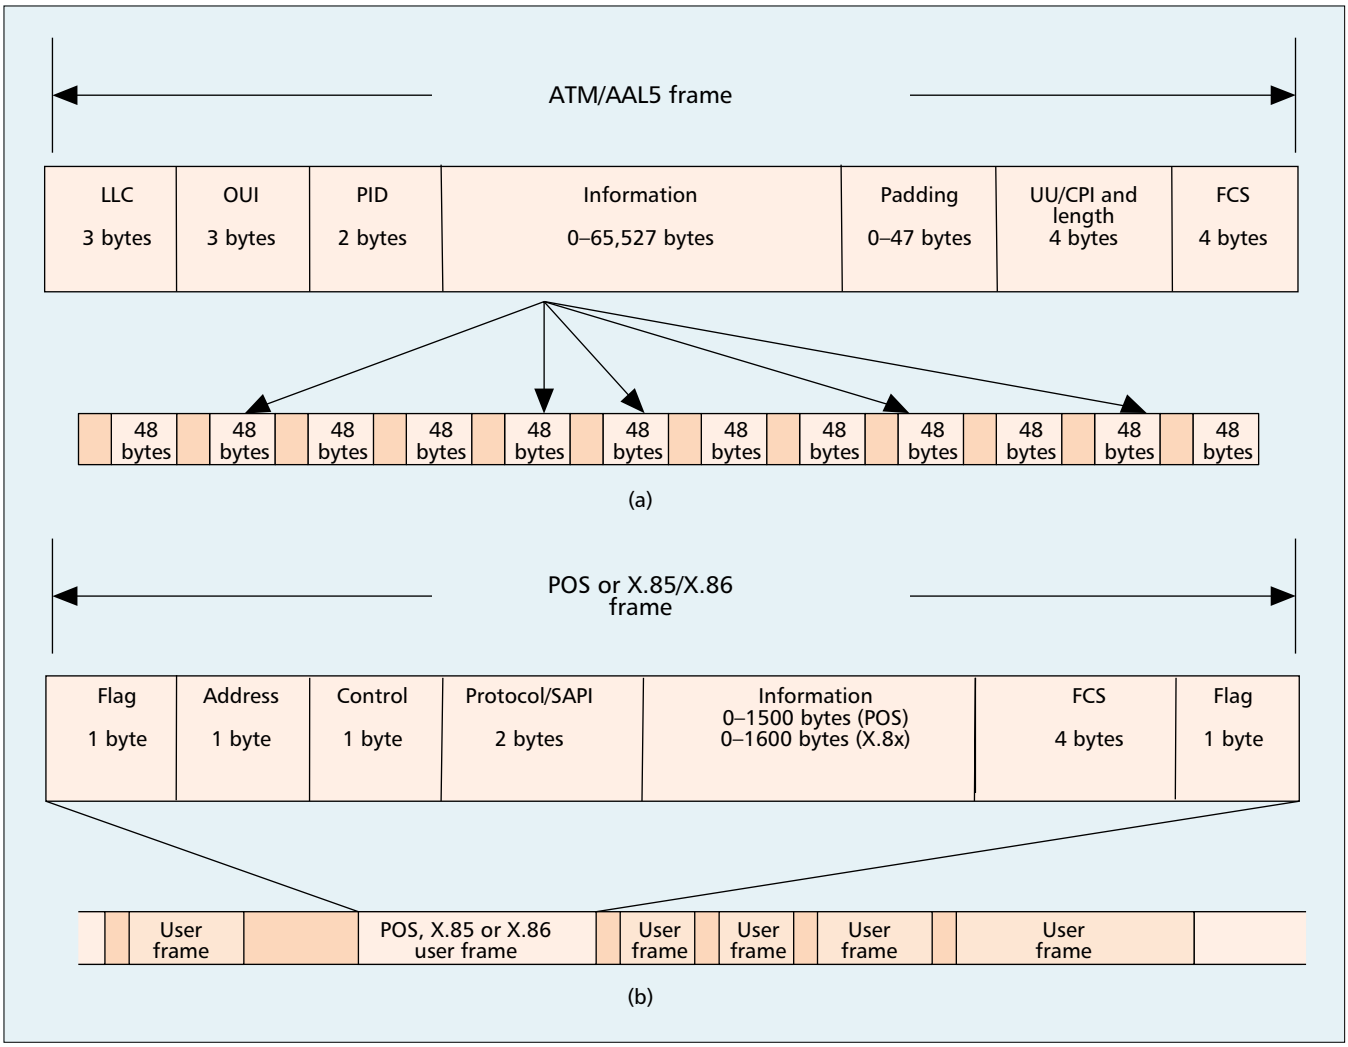

Figure 2. *Frame formats for a) ATM/AAL5; b) HDLC frames.*

(0x7e and 0x7d, respectively) must be replaced with an escape character followed by the original character exclusive-ORed with 0x20. As a result, the signal to be transported grows by 1 byte for each byte requiring transparency processing. While rare, it is possible in some applications, including digital video, for flag- and escape-emulating characters to occur frequently, leading to excessive and unpredictable bandwidth expansion.

These transparency byte stuffing operations randomly expand the effective PDU size in bytes, hence increasing the PDU size presented to the physical layer in a nondeterministic fashion. The average payload expansion may vary between 1 and 20 percent for random data (depending on the control code character set supported by the transmitter), but short-term variations over a few hundred characters can be rather large. This uncontrolled data expansion interferes with traffic engineering and QoS management mechanisms implemented by the layer 2 (Ethernet MAC) or layer 3 (IP) transport protocols in several ways:

- Creates random variable-length transmission overhead for different PDUs of the same size, which increases storage requirement on the data link receiver for PDU extraction.
- Increases the effective data rate of the user flow on the transmitter side, which increases the effective offered load to the transmission system. The traffic increase may exceed the allocated bandwidth on the outbound port for a QoS-enabled user flow (e.g., as implemented by a rate-based QoS scheduler or accounted for by a priority-based QoS scheduler), which may lead to internal buffer overflows and increased forwarding delay.
- Increases the bandwidth consumed by user flows, which decreases the real free bandwidth available on a network. This decreased bandwidth may not be visible to external operations, administration, maintainance, and provisioning (OAM&P) systems, particularly those responsible for capacity planning, traffic engineering, and QoS management, which may lead to improper route selection, imprecise usage forecast, and hence overall lack of control of QoS resources.

This payload expansion also provides malicious users with a trivial mechanism to significantly inflate a data flow's bandwidth requirements, by simply inserting a sequence of flag patterns within a user PDU.

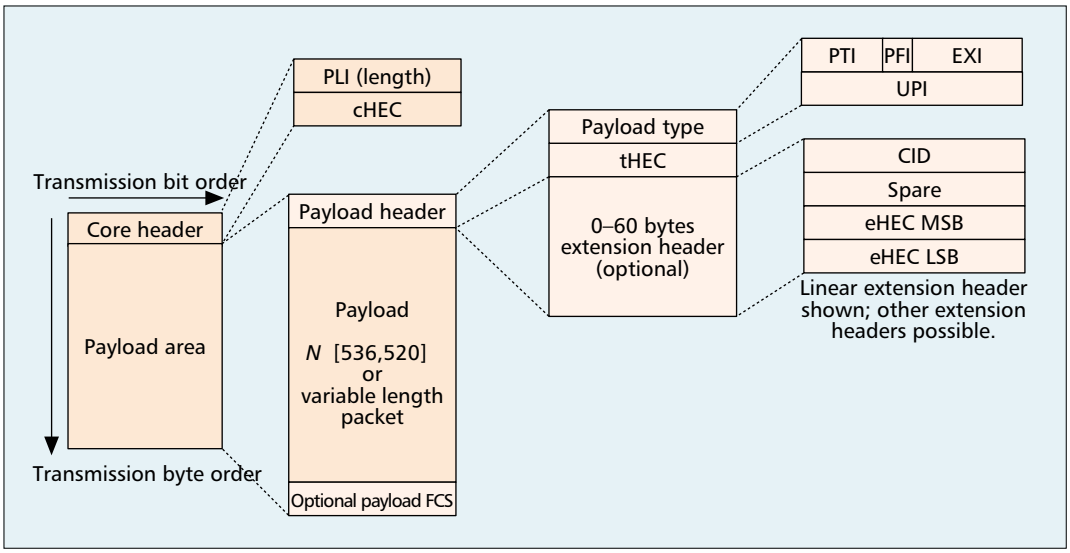

Figure 3. *GFP frame structure.*

Flag-based delineation mechanisms can result in error multiplication. When a non-flag data byte within a frame is corrupted and emulates a flag to the receiver, the frame is prematurely terminated at the receiver. The reception of an incomplete frame leads to an error burst that may not be detected by the frame CRC. Such error events may be caused by even a single bit error. For this reason, HDLC-like data link receivers always discard all errored frames.

HDLC frame delineation is based on the recognition of flag characters (0x7E). In the ingress direction (client packets being encapsulated for transport over SONET/SDH WAN), cut-through packet processing is possible. That is, as packet data is received, it can be transparency processed, scrambled per RFC 2615, and mapped into SONET/SDH paths. There is no need to buffer entire packets before mapping them into SONET/SDH using HDLC. In the egress direction, the need for reverse transparency processing, in which received escape characters are discarded, as well as the lack of knowledge of the length of the incoming packet imply store-and-forward packet handling to ensure that mid-packet underflow does not occur when delivering packets back to the client.

GFP solves the bandwidth expansion problem by encapsulating packets with a length/HEC header providing robust frame delineation while eliminating the need to replace any "reserved" characters with escape sequences. A companion article describes GFP length/HEC headers and frame delineation in greater detail.

Length/HEC encapsulation requires knowledge of each packet's length before the first bytes of the encapsulated packet — the length/HEC core header — can be transmitted. Consequently, GFP requires store-and-forward processing on ingress. On egress, however, cut-through operation is possible. With GFP, no reverse transparency processing is required, and the length of the following packet is known as soon as the core header is received.

Figure 3 illustrates the structure of GFP frames. Details of GFP frame fields and construction are provided in a companion article.

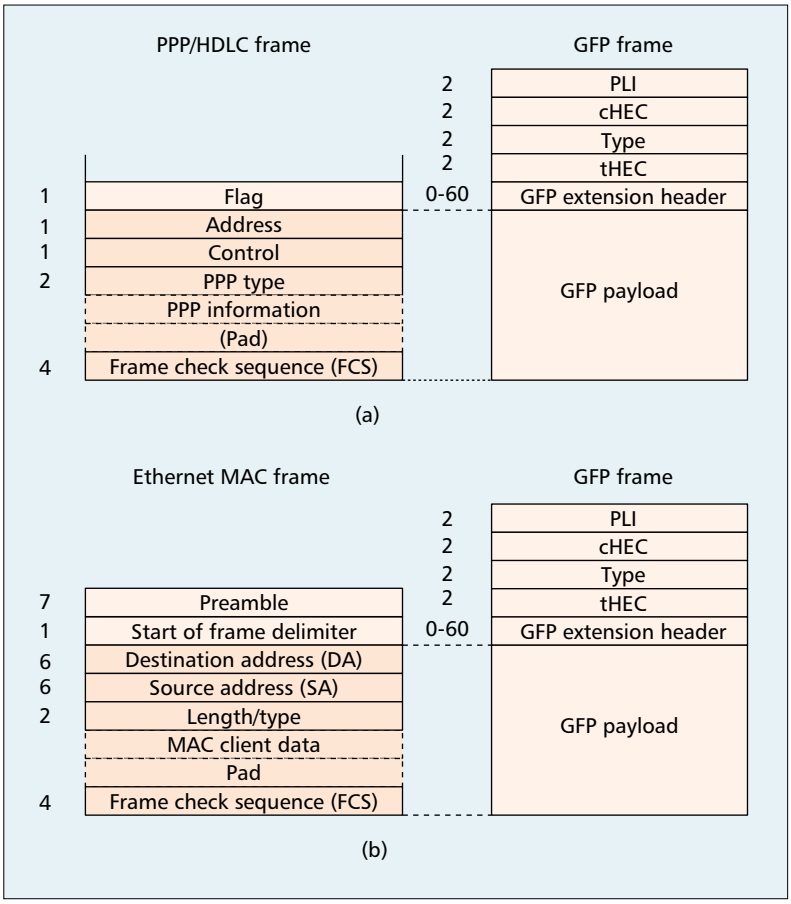

Figure 4. *GFP frame encapsulation of a) PPP/HDLC3z and b) GbE frames.*

Figure 4a illustrates the mapping of IP/PPP packets into the payload area of GFP frames.

Ethernet over SONET/SDH

Widespread acceptance of Ethernet and the emergence of GbE and 10GbE have generated interest in transporting Ethernet frames across

When fiber is not available, a dedicated wavelength per client is deemed too expensive, or when other benefits of in-place WAN networks are desired, efficient transport of client network protocols over SONET/SDH is desired.

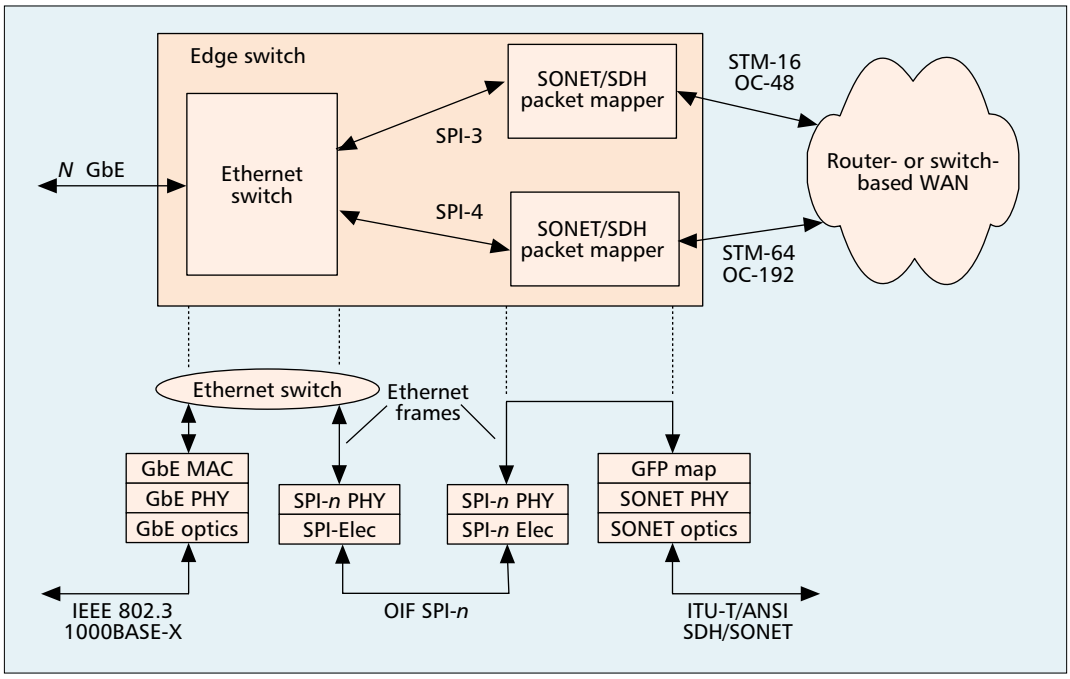

Figure 5. *Ethernet and GFP frame relationships.*

SONET/SDH networks. One such approach, ITU-T X.86 Ethernet over SONET/SDH (EoS), relies on familiar HDLC technology. Rather than encapsulate IP/PPP packets, EoS encapsulates complete Ethernet frames in HDLC packets. A unique protocol ID (service access point identifier, SAPI, in ITU-T terminology), prepended to the encapsulated frame, differentiates HDLC-encapsulated Ethernet frames from HDLC-encapsulated IP/PPP packets. As with IP over SONET/SDH, EoS suffers from the same nondeterministic bandwidth expansion resulting from byte-stuffed transparency processing.

GFP encapsulates Ethernet frames with the same length/HEC core header used for IP encapsulation, again eliminating the need for transparency processing. A payload header immediately following the length/HEC core header uniquely identifies the type of encapsulated traffic. Consequently, Ethernet over GFP is distinguishable from IP/PPP over GFP or other traffic encapsulated in GFP frames simply by examining the payload header. Again, more details of this may be found in the companion article detailing GFP frame format.

Figure 4b illustrates the mapping of Ethernet frames into GFP frames. Figure 5 illustrates an Ethernet "frame-mapped" application. A similar mapping supports IP packet transport by encapsulating IP packets, rather than Ethernet frames, within GFP frames.

Transparent GbE, FC, FICON, and ESCON Transport

In some applications, it may be desirable to simply transport a native network signal to a peer device without extracting packets or frames, re-encapsulating them with HDLC or GFP, and mapping encapsulated packets or frames into SONET/SDH paths. With the advent of wavelength-division multiplexing (WDM), a number of manufacturers developed products to simply optically multiplex together client network optical signals and transport them to peer devices. This approach works as long as fiber is in place where needed, and standards-based WAN network protection or OAM&P is not required.

When fiber is not available, a dedicated wavelength per client is deemed too expensive, or when other benefits of in-place WAN networks are desired, efficient transport of client network protocols over SONET/SDH is desired. A common attribute of many widely used high-speed datacom protocols is their reliance on 8B/10B physical layer coding. 8B/10B coding converts 8-bit bytes into 10-bit codewords, ensuring high transition density and DC balance for clock recovery circuits on the receiving end. The additional codespace created by 10-bit codewords provides for transmission error detection (but not correction) as well as unique control words without placing restrictions on data codewords. Control codewords are utilized for client link configuration, frame/packet delineation, and interpacket gap filler.

In frame-based IP and Ethernet over SONET/SDH approaches, a physical media adapter (PMA), physical coding sublayer (PCS), and media access controller (MAC) are required to synchronize on received 10B codewords, extract control and data characters, delineate frames, and extract packets or frames. Extracted packets or frames can then be reencapsulated inside HDLC or GFP frames, mapped into SONET/SDH paths, and transported over SONET/SDH lines. Not only is this a lot of work to simply transport the signal, but the link configuration information isn't transported, preventing client network nodes on either side of the transport link from "seeing" each other.

Users may wish to transparently transport native network signals over SONET/SDH MANs or WANs, using the MAN/WAN to simply stretch the distance between network nodes. Since SONET payloads are byte-aligned, directly mapping 10B codewords into SONET payloads is awkward. In addition, such an approach requires 25 percent more bandwidth than needed to transport just the information bits. Since SONET/SDH provides its own data scrambling and error detection, no transition density or error detection benefit is provided by preserving the 10-bit coding.

HDLC and ATM provide no generalized and efficient mechanism for transparently transporting 8B/10B client protocols over MANs and WANs. GFP solves this problem by recoding 8B/10B protocols using 64B/65B coding. The recoding scheme is more bandwidth-efficient, but retains and transports all of the data as well as special control characters, ensuring that data, link configuration, and frame delineation information is exchanged between client devices at either end of the transport link. A companion article provides details of 64B/65B encoding and encapsulation into GFP frames.

The current GFP standard supports transparent transport of GbE, Fibre Channel, ESCON, and FICON. Additional proposals are being considered to extend transparent transport support to digital video broadcast (DVB) asynchronous serial interface (ASI) and Fast Ethernet.

Transparent vs. Frame-Based Transport

Why does GFP provide two schemes for handling data protocols? The simple answer is that each offers a unique set of advantages over the other, depending on the application. While GbE, Fibre Channel, and ESCON share common 8B/10B coding, they do not share common frame delineation and packet formats. As a result, a protocol-specific PCS and MAC is required to extract and transport each protocol via frame-based GFP encapsulation. Transparent transport requires only minimal protocol awareness, allowing a single hardware design to transparently transport multiple protocols based on 8B/10B coding.

However, transparent transport delivers the complete client signal regardless of the information content of that signal. That is, even when a significant amount of the bandwidth in the signal received from the client network contains client idles, transparent-mapped GFP transports the idles as well as the information, consuming WAN bandwidth to deliver meaningless idles. Frame-mapped GFP only delivers client data frames, allowing for more efficient WAN bandwidth utilization to transport lightly loaded clients.

Recognizing the benefits of each approach, the current GFP standard supports both frame- and transparent-mapped client data transport.

Efficient Data Transport Using Virtual Concatenation

Unfortunately, the line and data rates of GbE, Fiber Channel, ESCON and FICON are not well matched to standard SONET/SDH contiguously concatenated payload rates (e.g., 1.0 Gb/s data rate of GbE does not fit into STS-12c/VC-4-4c at 622 Mb/s, and wastes bandwidth if mapped singly into a 2.5 Gb/s STS-48c/VC-4-16c). The recent addition of a virtual concatenation standard to SONET/SDH solves this problem by allowing any number of SONET/SDH paths to be virtually concatenated (in multiples of STS-1/VC-3 or STS-3c/VC-4, in the case of high-order virtual concatenation). Consequently, virtual paths can be created in multiples of 50 or 150 Mb/s, and multiple LAN/SAN data signals can be transported over single OC-48/STM-16 or OC-192/STM-64 lines with much better bandwidth utilization. Furthermore, each virtually concatenated group of signals can be independently routed through existing SONET/SDH networks without requiring any upgrades to deploy SONET/SDH network switches.

For example, one GbE signal can be transparently GFP-mapped into an STS-3c-7v/VC-4-7v (virtually concatenated group assembled from 7 STS-3c/VC-4 paths). The remaining 9 STS-3c/VC-4 paths in the originating OC-48 / STM-16 can be filled with other TDM voice signals, transparent- or frame-mapped GFP data signals, or even POS or ATM signals mapped into either contiguously concatenated or virtually concatenated paths.

Figure 6 illustrates the transparent mapping and independent routing of multiple GbE or FC client signals through a SONET/SDH WAN.

Data over Optical Transport Networks

SONET/SDH is the dominant MAN/WAN transport technology in use today for line rates between 51.84 Mb/s (STS-1) and 9953.28 Mb/s (STS-192/STM-64). At high bit rates (typically 10 Gb/s and above), forward error correction (FEC) can be required to achieve desired transmission distances without signal regenerators. The deployment of WDM systems transmitting multiple wavelengths on single fibers and the emergence of all-optical add/drop multiplexing and switching equipment have created additional layers to manage in today's long-haul networks. A new digital wrapper standard, ITU-T G.709, has been developed to provide sufficient overhead to both manage these additional layers in a nonproprietary manner and support stronger FEC.

In many applications, G.709 digital wrappers may be used to encapsulate 2.5, 10, or even 40 Gb/s SONET/SDH signals. However, G.709 also defines data mapping into the payload of G.709 signals using GFP. No other data mapping into G.709 is defined (e.g., mapping of HDLC-encapsulated signals into G.709 is not included in the standard). Since G.709 replicates and improves on the OAM&P capabilities of SONET/SDH, mapping GFP into SONET/SDH into G.709 consumes bandwidth for possibly redundant functionality. It is envisioned that future networks may eliminate redundant overhead by mapping data directly into G.709 using GFP. Both frame-mapped packet transport and transparent-mapped client signal transport applications may exist and can be supported.

Note that G.709 defines standard line rates

Even when a significant amount of the bandwidth in the signal received from the client network contains client idles, transparent-mapped GFP transports the idles as well as the information, consuming WAN bandwidth to deliver meaningless idles.

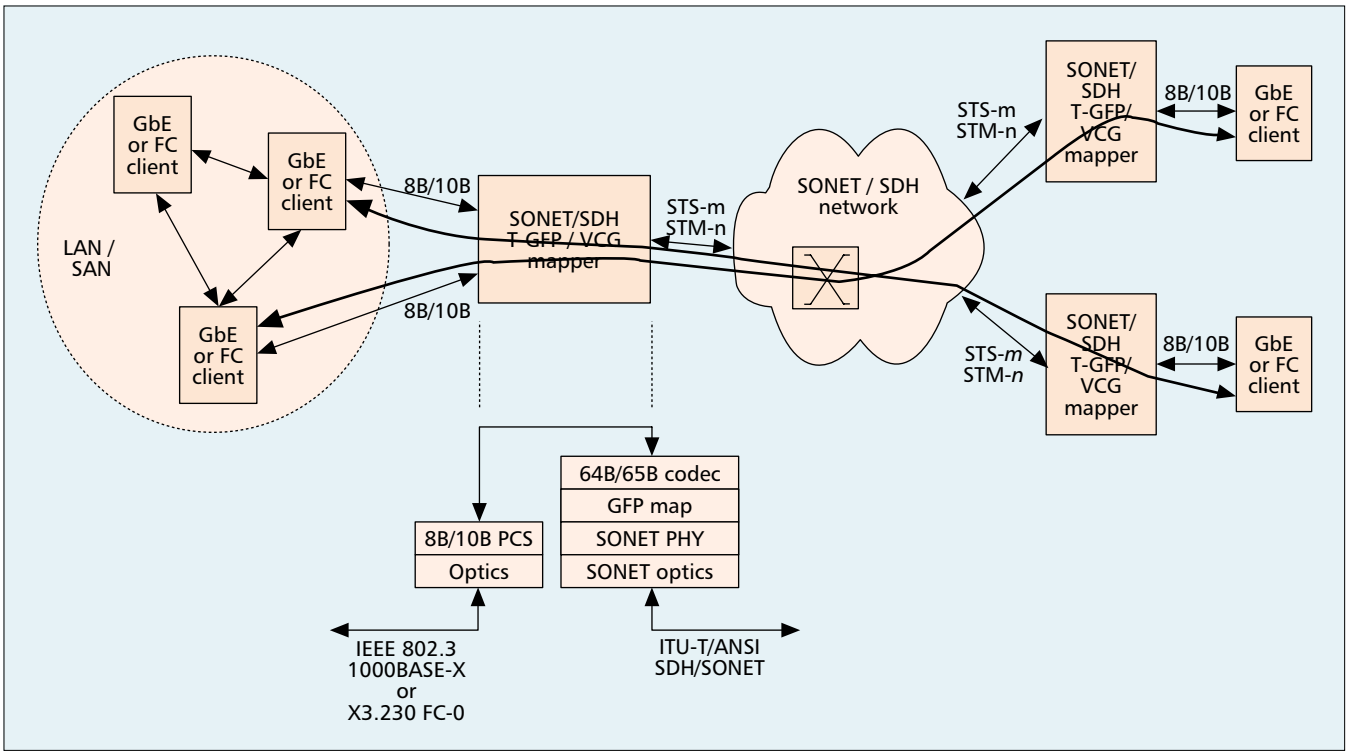

Figure 6. *Transparent-mapped GbE or FC independently routed through a WAN using virtual concatenated groups.*

slightly above 2.5, 10, and 40 Gb/s to support efficient transport of SONET/SDH standard rates. G.709 does not support channelization below the 2.666 Gb/s OTU-1 rate. Consequently, without multiplexing together frames from multiple clients, G.709 transport of client signals with line rates below 2.5 Gb/s is inefficient. However, as described below, GFP extension headers support frame-multiplexed port aggregation.

Port Aggregation in Point-to-Point Applications

SONET/SDH standards provide for robust point-to-point transport across short-, medium-, long-, and even ultra-long-haul optical networks. These same standards support channelization of the transport bandwidth in increments of 64 kb/s up to STS-1/STM-0 51.84 Mb/s SONET/SDH line rate, and in multiples of 51.84 or 155.52 Mb/s up to nearly 40 Gb/s. Using contiguous or virtual concatenation, SONET/SDH channel bandwidth can be tailored to fit the bandwidth required for individual data communication channels.

The current G.709 OTN standard does not support channelization below 2.5 Gb/s, but does support channelization in multiples of 2.5 Gb/s up to 40 Gb/s. Virtual concatenation of OTN channels has recently been included in the ITU-T G.709 OTN standard.

With optional extension headers, GFP provides another means of aggregating physical channels and transporting them across a single physical fiber or wavelength. Using linear extension headers, GFP allows each transported GFP frame to be associated with a specific channel, and allows frames from multiple channels to be multiplexed frame by frame over a single large concatenated payload. In applications where the instantaneous bandwidth utilization of any one channel may vary widely, but the average bandwidth across all channels is or can be constrained to the payload capacity of a single high-speed SONET/SDH payload (e.g., STS-48c/VC-4-16c or STS-192c/VC-4-64c), such port aggregation supports more efficient transport bandwidth utilization. In the case of OTNs, port aggregation provides a mechanism for supporting multiple data channels within a single G.709 physical link without requiring SONET/SDH to provide such channelization.

While GFP linear extension headers provide a mechanism for such port aggregation, the GFP standard does not attempt to suggest or standardize fairness algorithms to ensure that ports competing for shared point-to-point bandwidth obtain either prioritized or equal access to that bandwidth.

Resilient Packet Ring Applications

Today's SONET/SDH networks are deployed in a variety of network topologies, including point-to-point, mesh, and ring. Automatic protection switching schemes provide rapid and robust protection for data transported over SONET/SDH networks. However, these protection schemes require dedicated protection bandwidth, and do not prioritize traffic (other than to preemptively drop traffic being carried as extra unprotected traffic on protection channels when those channels are required for a protection switch). Resilient packet ring (RPR) technology attempts to better utilize ring network bandwidth by providing protection switching at the data rather than physical layer.

Members of the IEEE 802.17 Working Group (RPR) have indicated interest in using GFP for adaptation of client data payloads into SONET/SDH and OTN. While their work on an RPR standard is in its early stages, the group has requested that unique code points be reserved for an RPR user payload identifier as well as an RPR extension header identifier.

GFP provides several key advantages for RPR mapping including:

- Uniform mapping across all path types maximizing silicon and equipment commonality
- Robust delineation with deterministic bandwidth expansion
- Rate adaptation from RPR frame rate to the underlying SONET/SDH or OTN rate via idle frame insertion

The 802.17 group has indicated its intention to support the use of GFP's null extension header to implement a logical point-to-point pipe between stations on an RPR. The RPR MAC itself would be used to multiplex packet traffic from higher-layer clients onto these pipes as necessary. In the presence of a GFP-compliant PHY, the RPR MAC, if necessary, would populate the payload length indicator field in the core header in the case of ingress traffic and pass the PLI information seamlessly through its transit buffer in the case of pass-through traffic.

Conclusion

Data transport continues to drive demand for increased bandwidth in WANs. While a number of data transport mechanisms exist and are widely deployed, each has its own limitations. GFP provides a common data transport mechanism for a variety of traffic types, includes mechanisms to overcome the drawbacks of existing ATM and HDLC-based transport, and extends data transport capabilities to include transparent transport of LAN, SAN, and DVB signals. Coupled with virtual concatenation, GFP provides for both economical and efficient transport of multiple protocols, including RPR, over existing SONET/SDH networks as well as emerging OTN networks.

References

[1] ITU-T Draft X.86, "Ethernet over LAPS," Feb. 2001.
[2] ITU-T Draft G.7041, "Generic Framing Procedure (GFP)," Oct. 2001.
[3] A. Malis and W. Simpson, "PPP over SONET/SDH," IETF RFC 2615, June 1999.
[4] ITU-T Draft X.85/Y.1321, "IP over SDH Using LAPS," Feb. 2001.
[5] D. Grossman and J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5," IETF RFC 2684, Sept. 1999.
[6] IEEE 802.1D, "IEEE Standard for Information Technology —Telecommunications and Information Exchange Between Systems — IEEE Standard for Local and Metropolitan Area Networks —Common Specifications — Media Access Control (MAC) Bridges," 1998.
[7] E. Rosen, A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," IETF RFC 3031, Jan. 2001.
[8] FRF.14, "Physical Layer Interface Implementation Agreement," Dec. 1998.
[9] W. Simpson, "PPP in HDLC-like Framing," IETF RFC 1662, July 1994.

Additional Reading

[1] IEEE 802.3, "Information Technology—Telecommunications and Information Exchange between Systems—Local and Metropolitan Area Networks — Specific Requirements — Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications," 2000

Biographies

MICHAEL SCHOLTEN (mscholte@amcc.com) is a senior principal engineer at Applied Micro Circuits (AMCC) functioning as a senior systems engineer. He is responsible for specifying new products within AMCC's Digital Product Division. He represents AMCC in T1X1 standards committees and consults with AMCC's representatives to ITU-T, OIF, IEEE, and other standards organizations. Previously, he has worked as a software engineer, software development manager, product development manager, program manager, and systems engineer in various optical telecommunications companies. He is a graduate of Michigan Technological University.

ZHENYU ZHU (zzhu@amcc.com) is a senior staff engineer at Applied Micro Circuits (AMCC) functioning as a senior systems engineer. He is responsible for specifying new products within AMCC's Digital Product Division. Previously he has worked as a system engineer, ASIC design engineer, and network engineer in Lucent Technologies, Siemens Mobile Communications, and other communication companies. He holds degrees from Lehigh University, Chinese Academy of Science, and Shanghai Jiao Tong University. His interests include optical communications, data networks, wireless communications, and communication/networking ASIC implementation.

ENRIQUE HERNANDEZ-VALENCIA [M] (enrique@lucent.com) is a distinguished member of technical staff at Lucent Technologies Bell Laboratories. He received his B.Sc. degree in electrical engineering from Universidad Simon Bolivar, Caracas, Venezuela, and his M.Sc. and Ph.D. degrees in electrical engineering from the California Institute of Technology, Pasadena. Since joining Bell Laboratories in 1987, he has worked in the research and development of high-speed data communications protocols and systems. He is a member of ACM and Sigma Xi.

JOHN HAWKINS (jhawkins@nortelnetworks.com) is a senior marketing manager at Nortel Networks responsible for OPTera Packet Edge technology (Nortel's pre-standard implementation of RPR). He is a member of the IEEE 802.17 working group and serves as Chairman of the Board of the RPR Alliance. Previously he has worked as IC designer, development manager, and product manager. He holds degrees from North Carolina State, Southern Methodist, and Duke Universities.

Data transport continues to drive demand for increased bandwidth in wide area networks. While a number of data transport mechanisms exist and are widely deployed, each has its own limitations.