Ethernet Over SONET

Technology White Paper

Mimi Dannhardt
Technical Advisor

Issue 1: February, 2002

PMC-2020296

Abstract

This White Paper describes the use of the new ITU G.7041 GFP and ITU G.7042 LCAS standards to provide Ethernet leased line service over SONET Virtual Concatenation.

About the Author

Mimi Dannhardt is a Technical Advisor in the Product Research Group and oversees ICs for SONET, optical transmission and access systems.

Contents

Table Of Figures

Table Of Tables

1 Introduction

The communications landscape is dominated by two differing technologies: Ethernet in the LAN for internal business communications; and SONET/SDH in the Telco/PTT WAN. When businesses need to communicate with each other, or a business wants to connect its head office and branch offices to the same LAN, there is an interconnection problem. To interface the LAN to the WAN provided by the Telco/PTT, historically, has required an inter-working protocol, as Ethernet is not directly supported over the SONET/SDH network.

Figure 1 Traditional Private Line Network

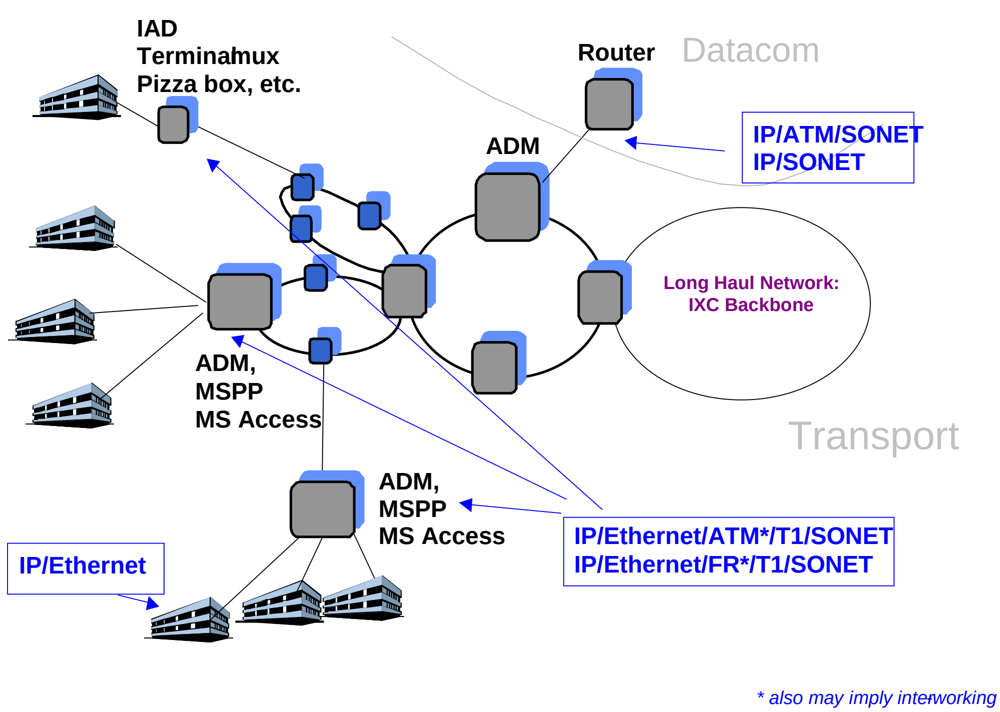

Ironically, in today's networks, in order to transport the customer's traffic over the Telco network in a "standard" manner, different technologies have been employed. These technologies (frame relay, ATM, Packet over SONET, ML-PPP to name a few) each require inter-working the native Ethernet traffic to the transport protocol prior to transmission. Often, the inter-working function must terminate the Ethernet and map the underlying IP traffic into a new Layer 2 (L2). Alternatively, the Ethernet is encapsulated within another L2 technology. Both of these techniques introduce additional complexity and cost into the WAN interface.

The various technologies also influence the demarcation point between the customer's network and the carrier's network. In some service models, the customer is required to inter-work their network traffic prior to handing it off to the public network. In alternative models, the carrier takes complete responsibility for the inter-working function. In both cases, network management issues exist as well as the need for additional equipment for inter-working. Furthermore, problems can arise when businesses are required to hand-off non-Ethernet traffic to the WAN provider. Although the local Enterprise IT staff are experts with Ethernet, they do not have the necessary expertise in the various WAN protocols to deal with issues that may occur.

Given the complexities and expense of the current inter-working solutions, why hasn't Ethernet been carried directly over SONET/SDH? Why has all the protocol inter-working evolved in the first place? In a nutshell, Ethernet rates do not match SONET/SDH rates and encapsulation methods have not been efficient.

As illustrated in Table 1, Ethernet rates are typically 10 Mbit/s, 100 Mbit/s or 1 Gbit/s, always increasing in factors of 10. On the other hand, SONET rates are optimized for telecommunications (or voice traffic) and do not match the optimal rates for transporting the typical Ethernet data stream. These rate mismatches make carrying a single Ethernet connection over a SONET pipe bandwidth inefficient.

Table 1 Typical Ethernet Rates vs. SONET rates

	SONET		
Data Bit Rate	**SONET Rate**	**Effective Payload Rate**	**Bandwidth Efficiency**
10 Mbit/s Ethernet	STS-1	~48.4 Mbit/s	21%
100 Mbit/s Fast Ethernet	STS-3c	~150 Mbit/s	67%
1Gbit/s Ethernet	STS-48c	~2.4 Gbit/s	42%

To address the inefficiencies of carrying a single Ethernet stream over a SONET pipe, a multitude of WAN technologies evolved to share the large transport pipes among multiple users' data streams while still providing Quality of Service and Bandwidth guarantees.

Furthermore, WAN technologies evolved that needed to address the reality of the user bandwidth being distributed across sub-rates as a result of limitations in the access bandwidth. Often, WAN bandwidth was severely constrained to fractional T1 rates due to cost. As the costs have dropped and demand for bandwidth has increased, the WAN access solutions (ATM or FR) provided via discrete T1 lines are no longer sufficient. One method of increasing bandwidth is by jumping to a T3 access solution (a larger pipe) and incurring the higher costs associated with it. A more cost-effective method currently being employed for acquiring additional bandwidth from a Telco-based WAN provider is through multilink services, which combines the T1s into a virtual fatter pipe either through FR multilink or IMA for ATM. These multilink technologies add inefficiencies due to the encapsulation that counteracts the benefits of bandwidth efficiencies through aggregation.

In spite of these issues, network equipment vendors have developed solutions that transport Ethernet over SONET using a variety of proprietary approaches. Unfortunately, the obvious problem with these proprietary approaches is that inter-operability between various vendors' equipment is difficult, if not impossible.

To help optimize the transport of Ethernet over SONET/SDH links, two new technologies have been standardized. The first, Virtual Concatenation allows for non-standard SONET/SDH multiplexing in order to address the bandwidth mismatch problem. The second, Generic Framing Procedure (GFP) provides deterministic encapsulation efficiency and eliminates inter-working.

2 Virtual Concatenation

Virtual Concatenation is a technique that allows SONET channels to be multiplexed together in arbitrary arrangements. This permits custom-sized SONET pipes to be created that are any multiple of the basic rates. Virtual concatenation is valid for STS-1 rates as well as for Virtual Tributary (VT) rates. All the intelligence to handle virtual concatenation is located at the end-points of the connections, so each SONET channel may be routed independently through the network without it requiring any knowledge of the virtual concatenation. In this manner, virtually concatenated channels may be deployed on the existing SONET/SDH network with a simple end-point upgrade. All the equipment currently in the center of the network need not be aware of the virtual concatenation.

In contrast, arbitrary contiguous concatenation also allows for custom sized pipes to be created but requires that the concatenated pipe be treated as single entity through the network. This requirement makes deployment of this service virtually impossible over the legacy SONET networks.

Using virtual concatenation, the SONET/SDH transport pipes may be “right-sized” for Ethernet transport. In effect, the SONET pipe size may be any multiple of 50Mbit/s for high-order virtual concatenation (STS-1), or 1.6 Mbit/s (VT1.5)/2.176 Mbit/s (VT2) for low-order virtual concatenation. Virtual concatenation rates are designated by STS-*m*-*n*v for high-order concatenation, where the *n*v indicates a multiple *n* of the STS-*m* base rate. Similarly, low-order virtual concatenation is designated by VT-*m*-*n*v.

Table 2 Typical Ethernet Rates vs SONET/SDH rates using Virtual Concatenation

	SONET		
Data Bit Rate	**SONET Rate**	**Effective Payload Rate**	**Bandwidth Efficiency**
10 Mbit/s Ethernet	VT-1.5-7v	~11.2 Mbit/s	89%
10 Mbit/s Ethernet	VT-2.0-5v	~10.88 Mbit/s	92%
100Mbit/s Fast Ethernet	STS-1-2v	~96.77 Mbit/s	103%
1Gbit/s Ethernet	STS-1-21v	~1.02 Gbit/s	98%
1Gbit/s Ethernet	STS-3c-7v	~1.05 Gbit/s	95%

Virtual concatenation provides flexibility in choosing the transport size to better match the desired bandwidth requirements. In addition to sizing the transport paths to handle the peak bandwidth expected, virtual concatenation may be used to create an arbitrary sized transport pipe. The pipe may be sized for the average bandwidth consumed for a single connection, or may be sized in order to provide a statistically multiplexed transport pipe.

In virtual concatenation, data is striped over the multiple channels in the Virtual Concatenation Group. This is illustrated in Figure 2. Control packets, which contain the necessary information for reassembling the original data stream, are inserted in some of the currently unused SONET overhead bytes. This information contains the sequence order of the channels and a frame number, which is used as a time stamp.

Figure 2 Virtual Concatenation

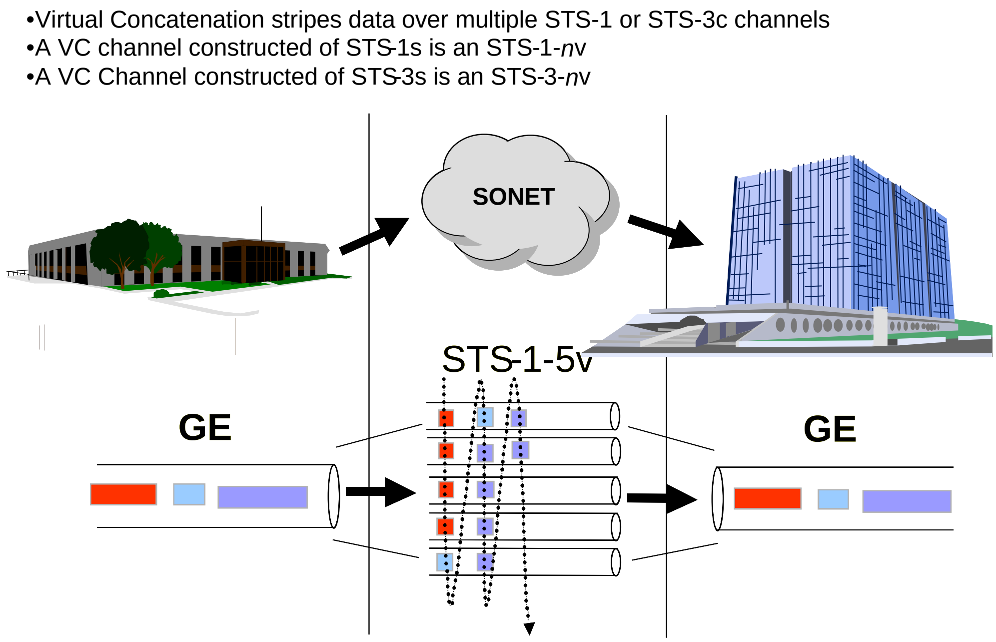

The receiving end-point is then responsible for reassembling the original byte stream. This includes compensating for differential delay that may have occurred by different routings or paths that the channels took through the network (Figure 3).

Figure 3 Differential Delay

•Individual STS-1's or STS-3c's sub-channels can take different paths through the SONET network. This can introduce *differential delay*.
•Buffering at the far end is required to align the sub- channels and extract the original frames.

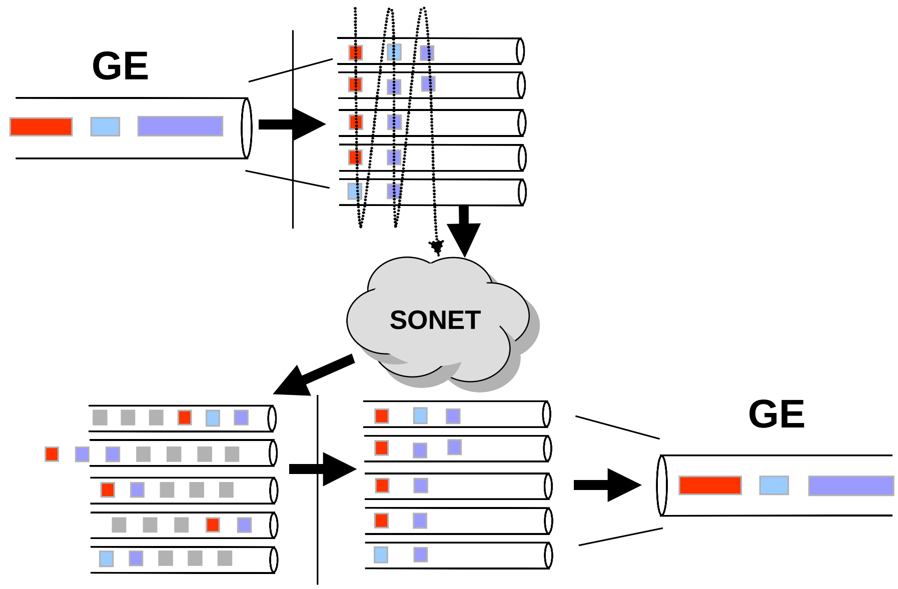

2.1 Dynamic Bandwidth Allocation

Along with virtual concatenation, the capability to dynamically change the amount a bandwidth used for a virtual concatenated channel is being developed. This capability is commonly referred to as Link Capacity Adjustment Scheme (LCAS). Signaling messages are exchanged within the SONET overhead in order to change the number of tributaries being used by a Virtually Concatenated Group (VCG). The number of tributaries may be either reduced or increased, and the resulting bandwidth change may be applied without loss of data in the absence of network errors.

The ability to change the amount of bandwidth allows for further engineering of the data network and providing new services. Bandwidth can be adjusted based on time-of-day demands and seasonal fluctuations. For example, businesses can subscribe to higher bandwidth connections (for backup etc.) when the demand for bandwidth is low and hence the cost is lower.

LCAS can further provide "tuning" of the allocated bandwidth. If the initial bandwidth allocation is only for the average amount of traffic rather than the full peak bandwidth, and the average bandwidth usage changes over time, the allocation can be modified to reflect this change. This tunability can then be used to provide (and charge for) only as much bandwidth as the customer requires.

LCAS is also useful for fault tolerance and protection since the protocol has the ability to remove failed links from the VCG. As the data stream is octet-striped across the tributaries in the virtual concatenation group, without such a mechanism if one of the tributaries has errors, the entire data stream has errors for the duration of the error within the tributary. The LCAS protocol provides a mechanism to detect the tributary in error and automatically remove it from the group. The VCG ends up operating at a reduced bandwidth, but the VCG still continues to carry data that is error-free.

Figure 4 Dynamic bandwidth adjustment

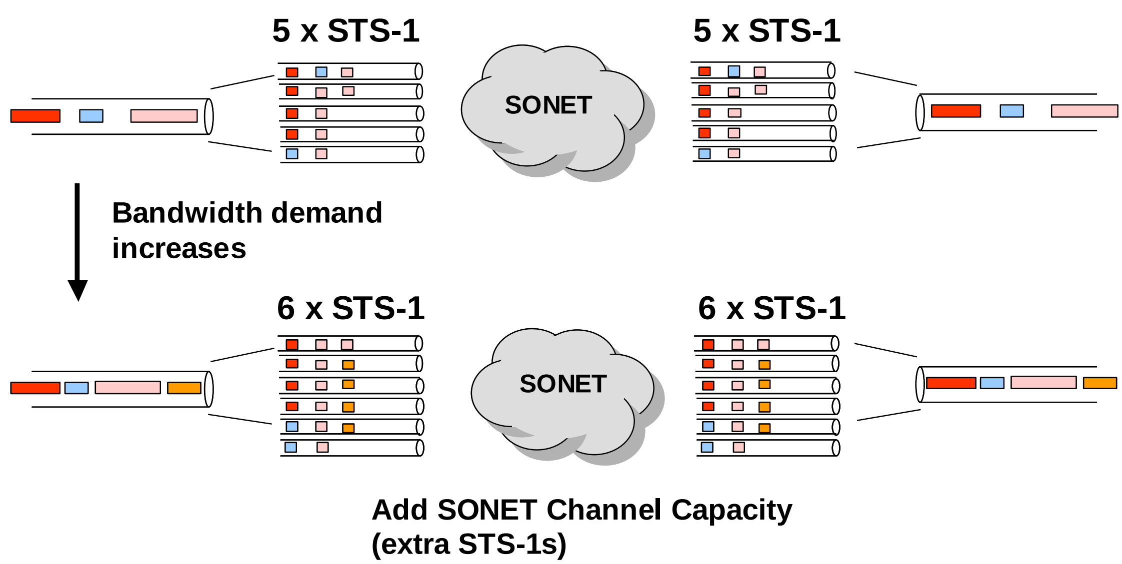

3 GFP

Generic Framing Procedure (GFP) is protocol for mapping packet data into an octet-synchronous transport such as SONET. Unlike HDLC-based protocols, GFP does not use any special characters for frame delineation. Instead, it has adapted the cell delineation protocol used by ATM to encapsulate variable length packets. A fixed amount of overhead is required by the GFP encapsulation that is independent of the contents of the packets. In contrast to HDLC whose overhead is data dependent, the fixed amount of overhead per packet allows deterministic matching of bandwidth between the Ethernet stream and the virtually concatenated SONET stream.

Figure 5 GFP Encapsulation Format

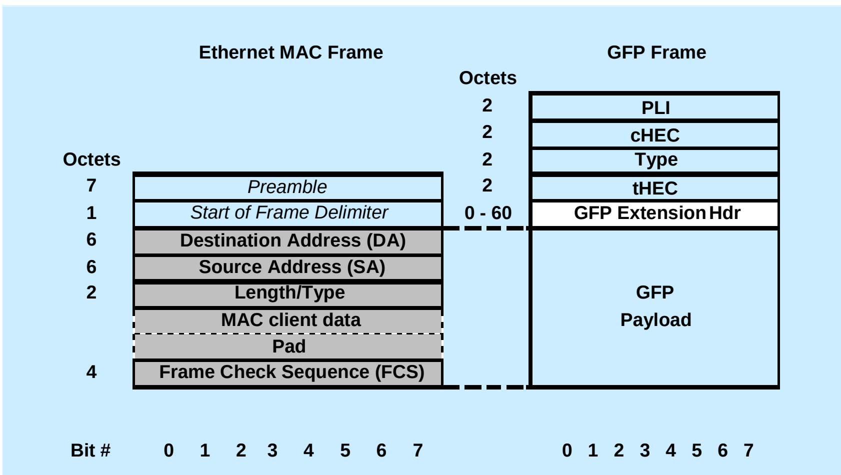

The GFP overhead can consist of up to 3 headers:

- a Core header containing the packet length and a CRC which is used for packet delineation;
- a Type header identifying the payload type;
- an Extension header, which is optional.

Frame delineation is performed on the core header. The core header contains the two byte packet length and a CRC. The receiver would hunt for a correct CRC and then use the received packet length to predict the location of the start of the next packet.

Within GFP, there are two different mapping modes defined: frame based mapping and transparent mapping. Each mode is optimized for providing different services.

3.1 Frame based GFP

Frame based GFP is used for connections where efficiency and flexibility are key. In order to support the frame delineation mode utilized within GFP, the frame length must be known and prepended to the head of the packet. In many protocols, this forces a store-and-forward encapsulation architecture in order to buffer the entire frame and determine its length. This buffering may add undesirable latency. Frame based GFP is good for sub-rate services and statistically multiplexed services as the entire overhead associated with the line coding and interpacket gap (IPG) are discarded and not transported.

3.2 Transparent GFP

Transparent GFP is useful for applications that are sensitive to latency or for unknown physical layers. In this encapsulation, all code words from the physical interface are transmitted. Currently, only physical layers that use 8B/10B encoding are supported. In order to increase efficiency, the 8B/10B line code are transcoded into a 64B/65B block code and then the block codes are encapsulated into fixed sized GFP packets.

This coding method is primarily targeted at Storage Area Networks (SANs) where latency is very important and the delays associated with frame based GFP cannot be tolerated.

4 What new services are enabled?

The advent of virtual concatenation, LCAS and GFP encapsulation will act as enabling technologies for deployment of some new Ethernet based services from the carriers. These types of services are available today but require that the traffic be inter-worked to a different Layer 2 technology prior to transport and switching. This inter-working increases the complexity and cost of these services.

4.1 Private Leased line

Private leased line services are used to interconnect various business locations. They are widespread today and are typically provided via ATM, Frame relay or multi-link Frame-relay.

An Ethernet based leased line service could be carried through the currently deployed SONET network using the GFP encapsulation and virtual concatenation technologies. Ethernet private lines may be provisioned at various service rates from 50 Mbit/s to 1Gbit/s utilizing STS-1 concatenation and from 1.6Mbit/s to 100 Mbit/s utilizing VT1.5 concatenation.

Ethernet private lines deployed over SONET offer the reliability and broad service area coverage associated with the carrier infrastructure. As it is a private line, data rate guarantees and security are key offerings, as well as upgradable bandwidth utilizing the LCAS protocol to adjust the bandwidth supplied.

4.2 Virtual Leased Line

Virtual Leased Lines or Virtual Private Networks (VPNs) would be deployed similarly to the Private Leased line. The main difference is that a Virtual Leased line is a shared service where many customers share the same transport bandwidth. This leads to more efficient use of the transport bandwidth via statistical multiplexing and thus lower costs. Since the transport bandwidth is shared, this service is generally a more economical service than a Private Leased line but generally does not have the QoS provided by the Private Leased Line. Instead, service parameters are control with Service Level Agreements (SLAs).

5 What products use these techniques?

The discussion of mapping Ethernet into SONET/SDH has thus far focused on the technology and the advantages of the techniques. The next question that arises is "what products would use these techniques?”

Before answering this question it is necessary to look at a simplified network, and discuss the typical products in the network.

Figure 6 Simplified Network Definition

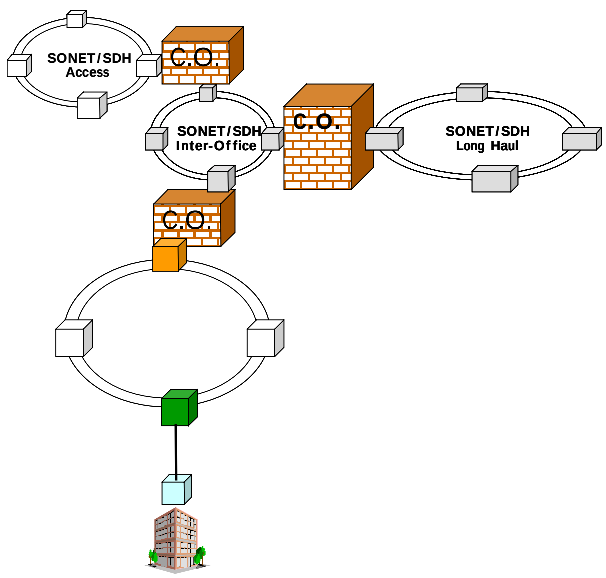

The Telco/PTTs networks are based on SONET/SDH rings. Multiple Rings are connected together to provide complete connectivity around a metropolitan area (city) and from city to city. There are four basic building blocks, or types of equipment, used to provide this connectivity. These are defined in Figure 7.

Figure 7 Definition of Connectivity Equipment

DCS, (Digital Cross-connect Switch) Provides grooming and connectivity between Rings.
Has no drops, only serves to provide connection between rings.

ADM, (Add Drop Multiplexer) Provides grooming and connectivity for access equipment.
Has interfaces to PDH and/or optical services. Typically layer 1 only, and closer to network edge than DCS

TM, (Terminal Multiplexer) Provides aggregation of optical and PDH services for transport on a typically unprotected Optical uplink to a higher order ADM

MSPP, (Multi-Service Provisioning Platform)
Adds layer 2 and/or Layer 3 functionality to the ADM.
Usually resides at the Metro Edge (might take the form of a POS card etc.)

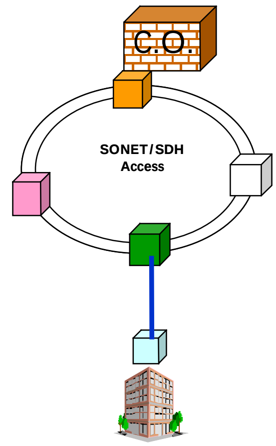

Of these four building blocks, three have customer-facing interfaces. These are Add Drop Multiplexers, Terminal Multiplexers and Multi-Service Provisioning Platforms.

5.1 Add Drop Multiplexer and Multi-Service Provisioning Platform

Some forms of MSPP are ADMs with data interfaces, some are Routers with SONET interfaces and switching. For the purposes of this discussion MSPP are assumed to be ADMs with Ethernet interfaces.

ADM's have traditionally been used to provide PDH (T1/E1/T3, etc.) and SONET/SDH drops to connect to specific customers. These platforms provide a good place for Ethernet over SONET/SDH customer interfaces.

Figure 8 ADM Functional Block Diagram

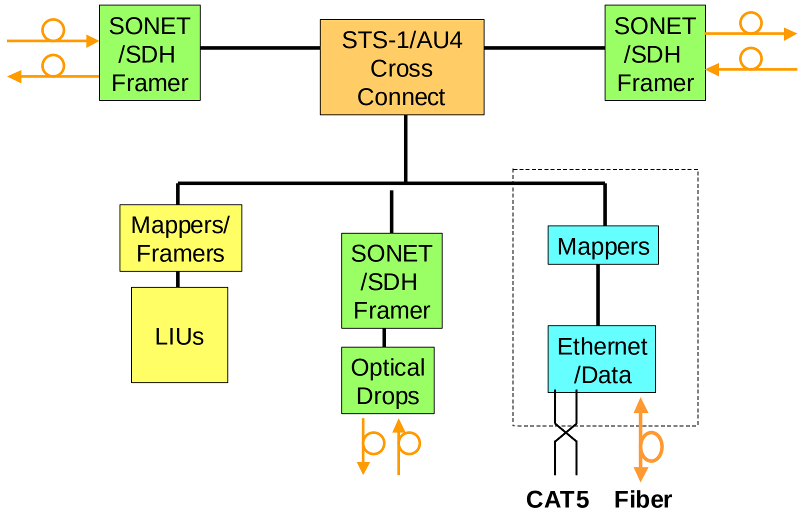

ADMs may be placed in large office building or in a more geographically central location at a CO. When providing Ethernet over SONET/SDH from the CO the interface is most likely to be a Fiber connection (i.e. Ethernet/Fiber) due to the potentially long paths between the CO and the customer.

When placed in the telephone closet of a multi-tenant office building the ADM could provide CAT5 cable interfaces to the customer.

5.2 Terminal Mulitplexer

TM's are similar to an ADM except that a TM terminates the SONET/SDH path (there is no through path). TMs are traditionally used to provide PDH (T1/E1/T3, etc.) and SONET/SDH drops to specific customers. These platforms also provide a good place for Ethernet over SONET/SDH customer interfaces.

TMs are often placed in multi-tenant office buildings. These products are likely to be found in the telephone closet of a multi-tenant office building with CAT5 cable interfaces to the customer.

Figure 9 Terminating Multiplexer Functional Block Diagram

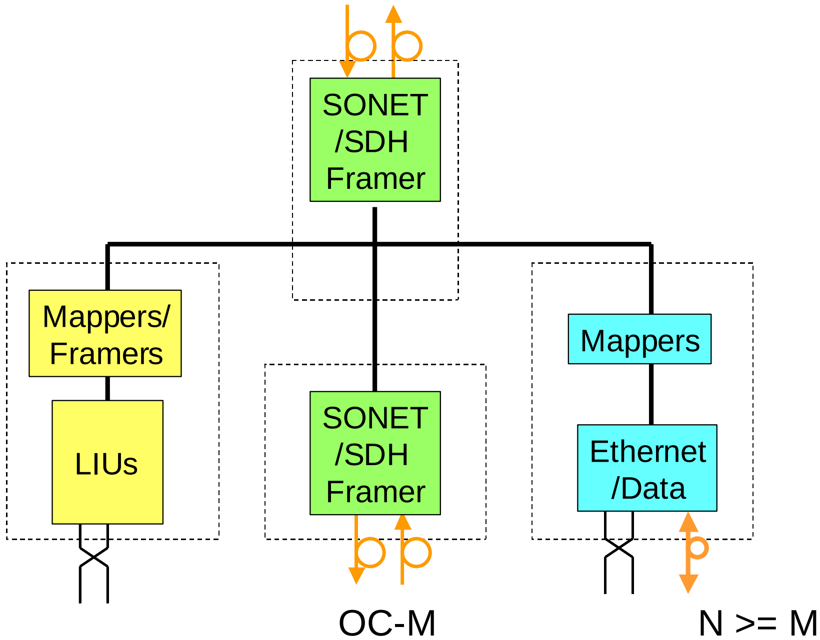

Head Office:
PMC-Sierra, Inc.
#105 - 8555 Baxter Place
Burnaby, B.C. V5A 4V7
Canada
Tel: 604.415.6000
Fax: 604.415.6200

To order documentation,
send email to:
document@pmc-sierra.com
or contact the head office,
Attn: Document Coordinator

All product documentation is
available on our web site at:
http://www.pmc-sierra.com
For corporate information,
send email to:
info@pmc-sierra.com

PMC-2020296 (1)