Ethernet over SONET/SDH;

Extending Ethernet into the Metro

White Paper

By Fredrik Olsson, Jim Shupenis,
Alisdair Gunn

Agilent Technologies
Semiconductor Products Group

Agilent Technologies

Abstract

The overwhelming majority of all data transport starts and ends on an Ethernet LAN card.

Within a campus Ethernet frames flow from source to destination untouched. But when a metro network separates the source and sink, many expensive to implement protocol conversions are often performed. With new standards being adopted, Ethernet protocol conversion ICs targeted at linking the LAN and Metro are now starting to appear.

Ethernet in the LAN

Ethernet has become the uncontested winner in local area networks. Economies of scale and fierce competition have taken Ethernet equipment to an extraordinary price/performance point.

Initially, Ethernet was using a single medium shared among all nodes and the line rate was a modest 10 Mb/s. It has evolved to using a switched medium with dedicated full-duplex bandwidth for each node and the line rate has evolved to 100 Mb/s, 1000 Mb/s and 10 Gb/s. What started out as a quite basic protocol has been extended to allow for network segmentation (VLANs), aggregation of links, automatic topology discovery (spanning tree), flow control etc. Today, Ethernet is a complete solution to all the data transport needs normally found in a corporate environment.

Ethernet in the access

Traditionally, carriers have provided TDM lines like T1s and DS3s to corporations for network access. At the customer premises, this is often the only non-Ethernet connection present. For this connection, special TDM interface equipment is needed like WAN access linecards, FRADs, IADs or DSU/CSUs.

Recently, some carriers have started to offer Ethernet hand-off, i.e. corporations access their carriers through Ethernet lines. Ethernet hand-off to carriers simplifies the management of corporate LANs and lowers cost.

Ethernet hand-off will become increasingly more common as the 802.3 EFM standard specifies a number of new media which Ethernet can run across, specifically for access purposes. Then Ethernet will also be available to run over existing copper wires as well as passive optical networks to further reduce access costs.

Ethernet in the access network is:

- *Scalable* – a simple Gigabit Ethernet interface can offer arbitrary bandwidth between 0 and 1000 Mb/s using flow control.
- *Easy to manage* – Ethernet is well understood by most corporate networking departments.
- *Already present* – No special interface equipment is needed, as Ethernet is already present in the LAN.
- *Cost effective* – The price/performance of Ethernet is excellent.
- *Evolving* – Standardization efforts are adding new media specifically for Ethernet access purposes.

Ethernet in the metro

In a metropolitan area network (MAN), there is a need to transport many different types of traffic;

- Data traffic for transparent LAN services and internet access.
- TDM lines for voice and private line services.
- ATM for mixes of data and voice.

While some new carriers have focused on data transport only, voice and legacy private line services are often a large portion of a carrier's revenue. To maximize profits, the network must support both legacy circuit-oriented as well as newer packet-based traffic.

In a pure Ethernet based network it is difficult to provide legacy TDM line transport reliably. TDM transport requires low latency and no jitter as well as fast restoration times not achievable in Ethernet networks. And as most incumbent carriers have large installed bases of circuit switched oriented SONET equipment, a move to a pure Ethernet network would mean replacing most of the existing equipment. Therefore, a move towards Ethernet in the metro must be a gradual shift without obsolescing the current infrastructure.

But at the same time there is much to gain from a move to Ethernet in metro networks. Data aggregation becomes simple and inexpensive. Access nodes are greatly simplified and no protocol conversions are needed. Network architects can select where to route and where to switch traffic. The metro can be divided into layer 2 “islands”, where routing is only performed at select locations which reduces both cost and complexity.

So while the existing SONET infrastructure must remain, there are many potential benefits from adding Ethernet transport to the metro network.

The solution – Ethernet over SONET

Carriers are handed more and more Ethernet yet have a huge investment in SONT/SDH equipment, and are not about to throw it all away. As revenues from carrying data increase far quicker than the revenues from carrying voice, the carriers must leverage what they can from their legacy SDH equipment. Hence the concept of Ethernet-over-Sonet, hereafter called EoS for their SONET networks.

The solution is quite natural; transport Ethernet over the existing SONET network and thereby preserve the excellent features of SONET that carriers require.

Quality of service and link performance is continuously monitored. Fast sub-50 ms restoration times during equipment failures and fiber breakages. Existing SONET OAM&P network management is still used. Legacy TDM traffic can continue to run over the same links without being impacted. And Ethernet channels can gradually be added to existing links as requirements change.

This year (2001) there has there been a standard defined on how to encapsulate Ethernet into a SONET payload. The legacy way of transporting data across SONET involves protocol conversions at the ingress and then recreation of the Ethernet frame at the egress.

Ethernet over SONET is still Ethernet. The full Ethernet frame is preserved, only a simple encapsulation occurs at the network ingress and egress. All the advanced features of Ethernet are left intactand since Ethernet is a layer 2 protocol and there is no requirement to route traffic at each node. The carrier can choose to route traffic at certain points and just switch traffic at others.

EoS – How it’s done

There are two standardized ways of sending Ethernet over a SONET network. One is approved and one is close to being finalized.

EoS Protocol – LAPS

Ethernet over LAPS, Link Access Procedure – SDH, is defined by ITU-T in X.86. It is a connection-less HDLC-like protocol that goes into the SONET payload.

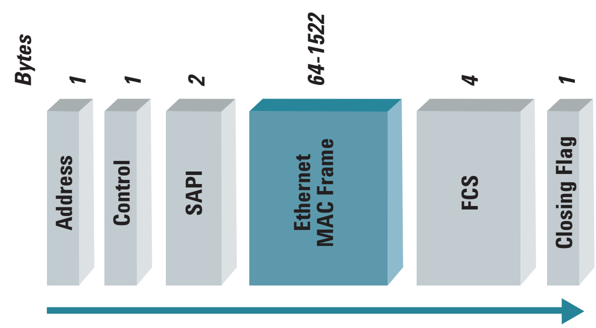

Figure 1. Ethernet in LAPS.

As shown in Figure 1, the first two bytes are the address and control constants (0x04 and 0x03, respectively), for compatibility with HDLC. SAPI, Service Access Point Identifier, indicates the payload type, in this case an Ethernet MAC frame with the constant 0xfe01. Then come the Ethernet MAC frame, a 32-bit LAPS FCS (CRC) and a closing flag. When the link is idle, inter-frame gap bytes, that are the same as the closing flag, are transmitted.

All data is then passed through an $X^{43}+1$ self-synchronous scrambler, before being inserted into the SONET payload. This provides protection against packets from malicious users.

One characteristic of byte-aligned HDLC-like protocols are that some bytes are “escaped” and replaced by a two-byte sequence. Hence, in the unlikely event that some Ethernet frames contain many 0x7d and 0x7e bytes, the link performance is degraded.

EoS Protocol – GFP

GFP, Generic Framing Procedure, is a new ITU standard with its roots in SDL, Simple Data Link. It is a comprehensive protocol with two options on headers and two options on payload mapping. Beside Ethernet, it can also map PPP, Fibre Channel, FICON and ESCON.

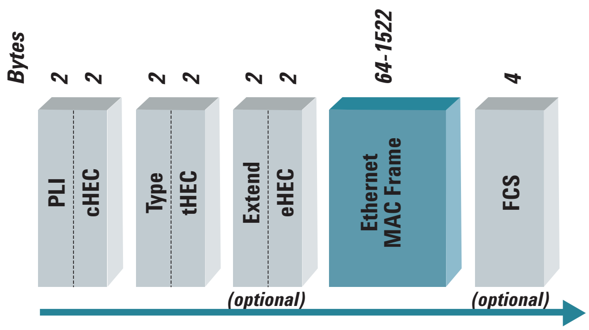

Figure 2. Ethernet in frame-mapped GFP.

A GFP frame, shown in Figure 2, has two or three header fields, payload and optional 32-bit FCS. Each header field is protected by a 16-bit CRC ("HEC").

First is the core header that indicates the length of the current frame in its Payload Length Indicator, PLI. Then comes the type header that contains a code for the type of payload encapsulated. The extended header field is an optional header that deals with logical links within a channel. In simple Ethernet encapsulated point-to-point applications, this header can be omitted altogether. Last come the payload and an optional 32-bit FCS (CRC). All fields except the core header are passed through an $X^{43}+1$ self-synchronous scrambler for protection against packets from malicious users.

Idle frames are transmitted when there is no payload to send. An idle frame is a core header with the length field set to zero. A nice characteristic of GFP is that there are no "escape" characters. The throughput of a link is only dependent on the amount of payload, not which byte values the payload contains. But it also means that the receiver needs to go through a hunt sequence where it searches for valid core headers and when a number of them are found, lock.

Summary

Ethernet over SONET greatly reduces the cost and complexity of carriers' networks by extending Ethernet into a metro network. It allows carriers to add data channels to existing TDM traffic and keep the important features of SONET intact.

Ethernet over SONET eliminates the need for costly protocol conversions when entering or leaving the SONET network, and allow carriers to offer Ethernet hand-off.

www.agilent.com/semiconductors

For product information and a complete list of distributors, please go to our web site.

For technical assistance call:
Americas/Canada: +1 (800) 235-0312 or (408) 654-8675

Europe: +49 (0) 6441 92460

China: 10800 650 0017

Hong Kong: (+65) 271 2451

India, Australia, New Zealand: (+65) 271 2394

Japan: (+81 3) 3335-8152(Domestic/International), or 0120-61-1280(Domestic Only)

Korea: (+65) 271 2194

Malaysia, Singapore: (+65) 271 2054

Taiwan: (+65) 271 2654

Obsoletes 5988-3330EN
December 21, 2001
5988-5108EN

Agilent Technologies