



APPLICATION NOTE 3849

Ethernet-over-PDH Technology Overview

Abstract: Ethernet-over-PDH (EoPDH) is a collection of technologies and standards that allow the transport of native Ethernet frames over the well-established PDH telecommunications infrastructure. This enables carriers to make use of the extensive network of legacy PDH and SDH equipment to provide new Ethernet-centric services. In addition, EoPDH paves the pathway for interoperability and the gradual migration of carriers to Ethernet networks. This article covers the technologies used in EoPDH including frame encapsulation in GFP as defined in G.7041, mapping Ethernet-over-PDH framing, link aggregation through G.7043, link-capacity adjustment through G.7042, management messaging as defined in Y.1731 and Y.1730, VLAN tagging, QoS prioritization, and higher layer applications such as DHCP servers and HTML user interfaces.

Ethernet transport over non-Ethernet networks has existed for decades. A myriad of technologies, protocols, and equipment have been created to accomplish one seemingly simple task: connect network node A with network node B over distance X. The set of solutions to that simple equation has thus far been unbounded. From the first computer gateway with a 300-baud acoustic FSK modem, to today's advanced Ethernet-over-SONET/SDH systems, the task has remained essentially the same. However, varying forces over the years have caused the evolution of technological implementations used to solve this task and shaped them to the needs of the day. Some of the evolutionary "branches" have been miserable failures. Others have seen extensive global deployments, such as DSL. How does one identify an emerging evolutionary branch that will endure? Using hindsight as a guide, the enduring technologies struck an optimal balance of service quality, dependability, available bandwidth, scalability, interoperability, ease of use, equipment cost, and cost of operation. Technologies that perform poorly in any of these areas are not selected for widespread usage, and eventually disappear or are relegated to niche environments. With this perspective in mind, the emerging Ethernet-over-PDH (EoPDH) technology can be evaluated.

In a broad stroke, EoPDH is the transport of native Ethernet frames over the existing telecommunications copper infrastructure by leveraging the well-established Plesiochronous Digital Hierarchy (PDH) transport technology. EoPDH is actually a collection of technologies and new standards that allow carriers to make use of their extensive networks of legacy PDH and SDH (Synchronous Digital Hierarchy) equipment to provide new Ethernet-centric services. In addition, the collection of EoPDH standards paves the pathway for interoperability and the gradual migration of carriers to Ethernet networks. The standardized technologies used in EoPDH (in simple terminology) include frame encapsulation, mapping, link aggregation, link capacity adjustment, and management messaging. Common practices in EoPDH equipment also include the tagging of traffic for separation into virtual networks, prioritization of user traffic, and a broad range of higher layer applications such as DHCP servers and HTML user interfaces.

Frame encapsulation is the process by which Ethernet frames are placed as payload inside an auxiliary format for transmission on a non-Ethernet network. The primary purpose of encapsulation is the identification of the beginning and ending bytes of the frame. This is known as frame delineation. In actual Ethernet networks, a start of frame delimiter and length field perform the frame delineation function. A secondary role of encapsulation is to map the sporadic ("bursty") Ethernet transmissions into a smooth, continuous data stream. In some technologies, encapsulation also performs error detection by appending a Frame Check Sequence (FCS) to each frame. Many encapsulation technologies exist, including High-Level Data Link Control (HDLC), Link Access Procedure for SDH (LAPS/X.86), and Generic Framing Procedure (GFP). Although any encapsulation technology could theoretically be used for EoPDH applications, GFP has significant advantages and has emerged as the preferred encapsulation method. Most EoPDH equipment also supports HDLC and X.86 encapsulation for interoperability with legacy systems.

GFP is defined in ITU-T G.7041, and makes use of Header with Error Control (HEC) frame delineation. In some other encapsulation protocols that use start/stop flags, such as HDLC, bandwidth expansion occurs when start/

stop flags occur in the user data and must be replaced with longer escape sequences. By making use of HEC frame delineation, there is no need for GFP to perform flag substitution in the data stream. This gives GFP the significant advantage of having a consistent and predictable payload throughput. This is critical for carriers needing to provide customers with a guaranteed throughput. In **Figure 1**, the frame format of Frame-mapped GFP (GFP-F) is shown, along with HDLC for comparison. Note that the octet count is the same for native Ethernet and GFP-F encapsulated Ethernet. This small detail simplifies rate adaptation. Once the Ethernet frames are encapsulated in a higher level protocol that performs frame delineation, they are then ready to be mapped for transport.

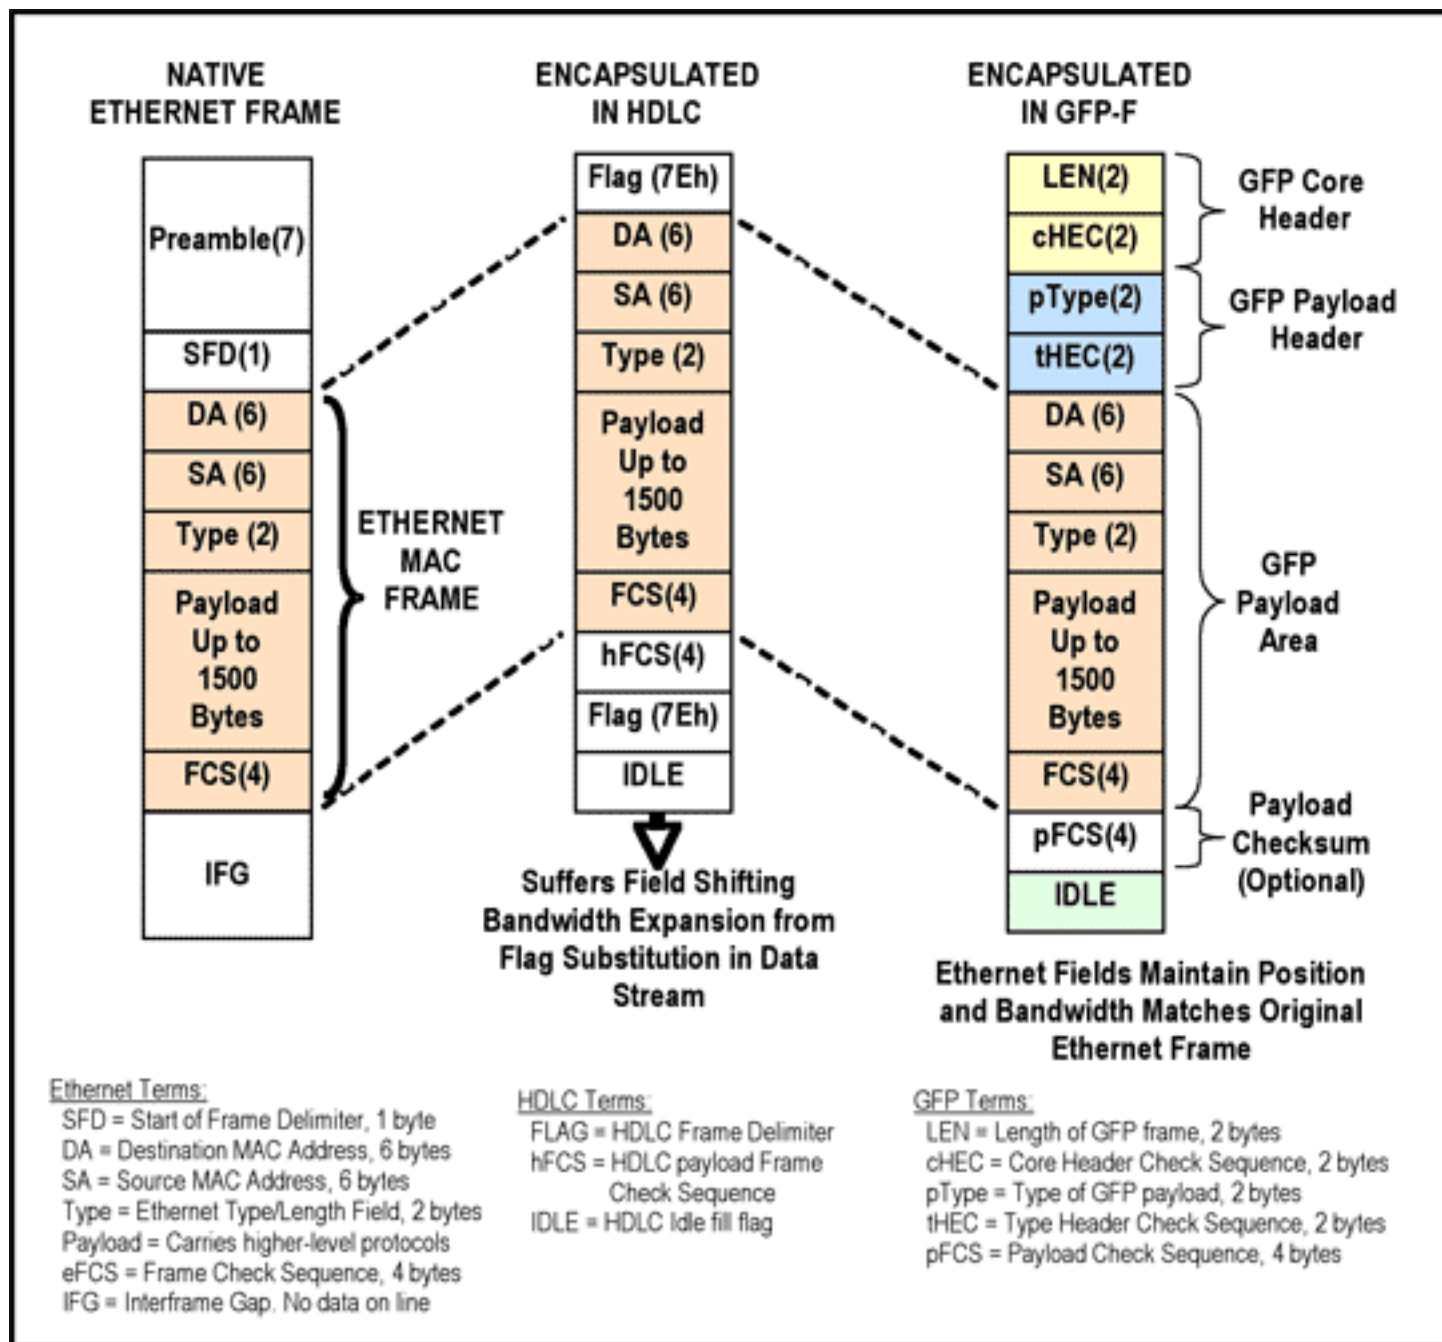


Figure 1. Comparison of HDLC and GFP frame encapsulation.

Mapping is the process by which the encapsulated Ethernet frames are placed in a "container" for transport across a link. These containers have various names across technologies. To generalize the term "container" for purposes of discussion, a container's primary purpose is to provide alignment of information. Some containers also provide management/signaling paths and link quality monitoring. Containers normally have rigid formatting with predefined locations for overhead and management traffic. Some examples of containers in the SDH include C-11, C-12, and C-3. The terms "trunk" and "tributary" are commonly used to refer to PDH containers. Some examples in the PDH include the DS1, E1, DS3, and E3 framing structures. In most cases, one or more lower data-rate containers can be placed inside ("mapped") into a higher data-rate container. In SONET/SDH

networks, virtual channels (VCs) and tributary units have also been defined, which work around some of the rigid requirements of the basic containers to provide greater flexibility.

Frame formats for the basic DS1 and E1 tributaries are shown in **Figure 2**. Note that each frame has a reserved location for framing information. The purpose of the framing bit (or byte) is to provide alignment information to the receiving node. The structured frame format is repeated every 125ms. A group of 24 DS1 frames is an Extended Super Frame (ESF). A group of 16 E1 frames is an E1 multiframe. By using the framing information, the receiving node can separate the incoming bits into individual time slots or channels. In traditional telephony, each time slot (or channel) carries the digitized information for a single telephone call. When transporting packetized data, all time slots can be collectively utilized as a single container.

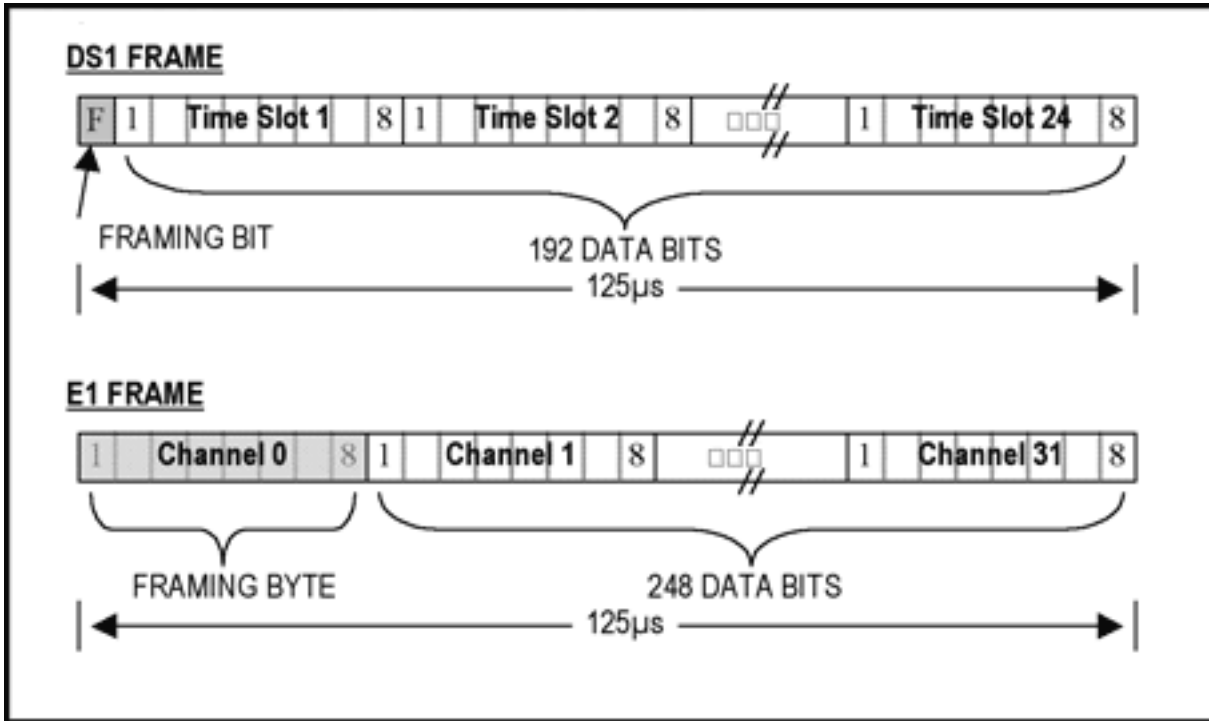


Figure 2. Examples of PDH frame formats.

When encapsulated Ethernet frames are transported over PDH, the time between Ethernet frames is filled with an idle pattern. When GFP encapsulation is being carried over DS1 or E1, the information is byte-aligned. Alignment is slightly more complicated when a DS3 link is used. Nibble alignment is specified for DS3 links in ITU-T G.8040. **Figure 3** shows an example of GFP-encapsulated Ethernet over a DS1 link. Note that the positioning of the encapsulated Ethernet frame is independent of the DS1 framing pattern bits ("F") and is byte-aligned. Though not shown in the diagram, the payload information has an X43+1 scrambling algorithm applied prior to transmission. Similar mapping and scrambling techniques are utilized for SDH transport containers. The complete specification for mapping Ethernet frames directly over SDH can be found in ITU-T G.707.

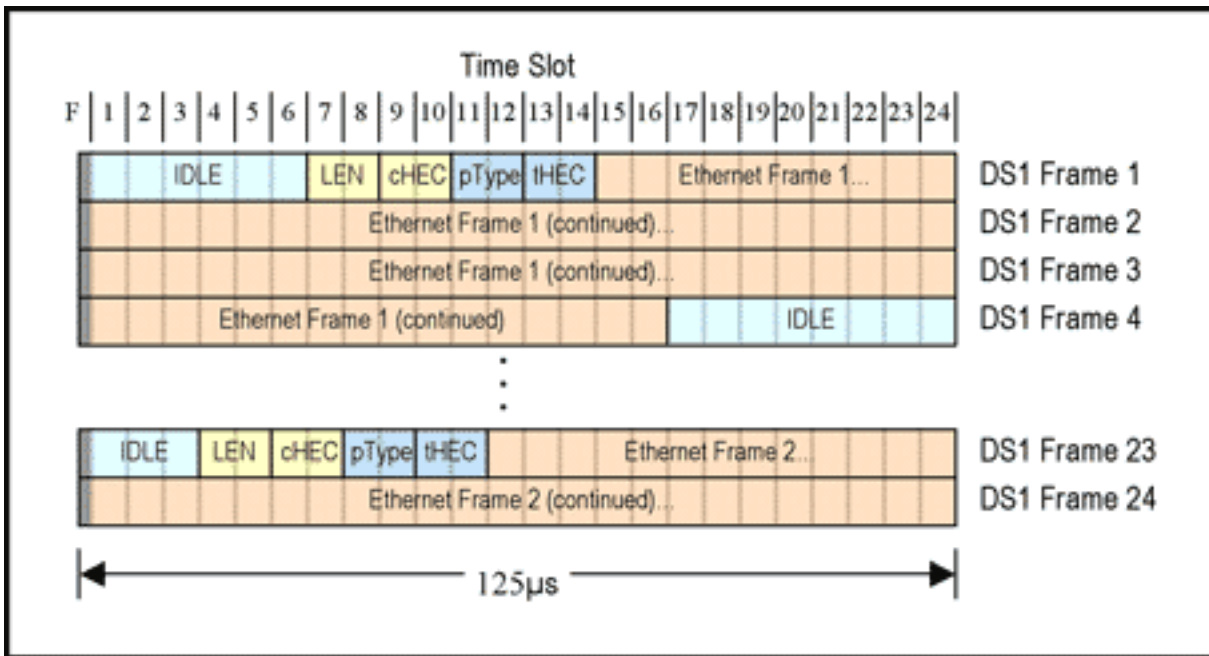


Figure 3. GFP Encapsulated ethernet frames mapped into a DS1 extended super frame (ESF).

Link aggregation is functionally the combination of two or more physical connections into a single, virtual connection. Link aggregation is actually a structured methodology for distributing data across multiple signal paths, aligning information received from paths with dissimilar latencies, and recompiling the data correctly for a transparent handoff with higher level protocols. Link aggregation is also not new. Multi-Link Frame Relay (MLFR), Multi-Link PPP (MLPPP), Multi-Link Procedure (X.25/X.75 MLP), and Inverse Multiplexing over ATM (IMA) are just a few examples of link-aggregation technologies. Of these, IMA and MLFR are the most widely deployed.




Figure 4. Link aggregation application example.

Link aggregation has typically been used to increase bandwidth between two network nodes, as shown in **Figure 4**, allowing deferral of a migration to a higher throughput PDH or SDH tributary. One form of link aggregation, Ethernet in the First Mile (EFM, defined in IEEE 802.3ah) bonds multiple DSL lines together to either increase throughput at a given distance or, often more importantly, to effectively increase the distance able to be served at a given throughput.

The primary link aggregation technology used in SONET/SDH networks today is called Virtual Concatenation (VCAT) and is defined in ITU-T G.707. The standard makes use of existing overhead paths for VCAT overhead. However, when VCAT the concept was extended to PDH networks, the existing management paths were insufficient and a new field was assigned for the VCAT overhead path. **Figure 5** shows the location of the VCAT overhead for a DS1 connection. The overhead byte occupies the first time slot of the ESF on each of the concatenated DS1s.




Figure 5. Virtual concatenation (VCAT) overhead for DS1.

The management channel created by the VCAT overhead byte is used to convey information about each link. With each transmitted DS1 ESF or E1 multiframe, one byte of VCAT overhead is placed on the link. Thus, 1/576th of the available DS1 bandwidth is used for VCAT overhead.

The VCAT overhead definition is shown in **Figure 6**. The 16 bytes shown in the figure are transmitted one byte at a time over 16 consecutive DS1 ESFs. Every 48ms, the bytes are repeated.

The lower nibble of the VCAT overhead byte contains a Multi-Frame Indicator (MFI), used to align the frames from links with varying transmission delays. The high nibble contains a uniquely defined control word for each of the 16 values of MFI. This upper nibble is called the VLI, for Virtual Concatenation and Link Capacity Adjustment Scheme (LCAS) information.

| ESF/MF# | VLI CONTROL | | | | MF1 SEQUENCE | | | |
|---------|-----------------|----------------|----------------|----------------|--------------|---|---|-----|
| | MSB | | | | | | | LSB |
| 1 | MST (7-4) | | | | 1 | 0 | 0 | 0 |
| 2 | MST (3-0) | | | | 1 | 0 | 0 | 1 |
| 3 | 0 | 0 | 0 | RS-ACK | 1 | 0 | 1 | 0 |
| 4 | RESERVED (0000) | | | | 1 | 0 | 1 | 1 |
| 5 | RESERVED (0000) | | | | 1 | 1 | 0 | 0 |
| 6 | RESERVED (0000) | | | | 1 | 1 | 0 | 1 |
| 7 | RESERVED (0000) | | | | 1 | 1 | 1 | 0 |
| 8 | SQ Bits 7-4 | | | | 1 | 1 | 1 | 1 |
| 9 | MF12 MSBs (7-4) | | | | 0 | 0 | 0 | 0 |
| 10 | MF12 LSBs (3-0) | | | | 0 | 0 | 0 | 1 |
| 11 | CTRL / LCAS | | | | 0 | 0 | 1 | 0 |
| 12 | 0 | 0 | 0 | GID | 0 | 0 | 1 | 1 |
| 13 | RESERVED (0000) | | | | 0 | 1 | 0 | 0 |
| 14 | RESERVED (0000) | | | | 0 | 1 | 0 | 1 |
| 15 | C ₁ | C ₂ | C ₃ | C ₄ | 0 | 1 | 1 | 0 |
| 16 | C ₅ | C ₆ | C ₇ | C ₈ | 0 | 1 | 1 | 1 |

* Source: ITU-T G.7043

Figure 6. VCAT overhead byte definition for DS1/E1.

Collectively, the concatenated links are referred to as a Virtually Concatenated Group (VCG). All members of a VCG have their own VCAT overhead path, as shown in **Figure 7**. Figure 7 also diagrams the placement of data on the members of the VCG. The complete EoPDH link bonding specification can be found in ITU-T G.7043.

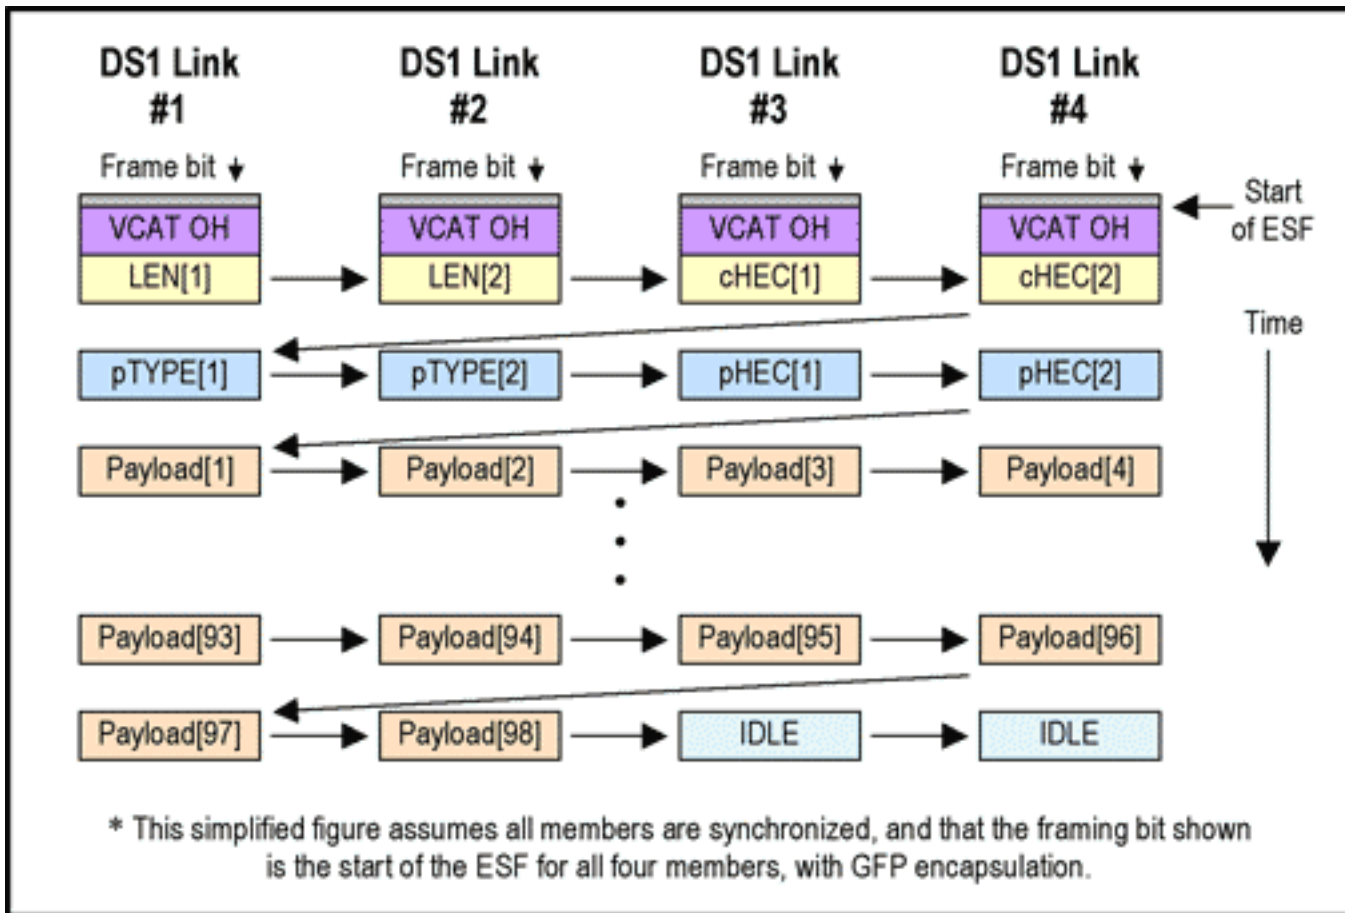


Figure 7. Distribution of data on a four-member DS1 VCG.

Link capacity adjustment is used to change the aggregate throughput by the addition or removal of logical connections between two nodes. When members of a VCG are added or need to be removed, the two end nodes negotiate the transaction using the LCAS. LCAS makes use of the VCAT overhead path to perform the negotiation. By using LCAS, bandwidth can be added to the VCG without interrupting the flow of data. In addition, failed links are automatically removed with minimal impact on traffic. The complete standard for LCAS can be found in ITU-T G.7042/Y.1305.

Management messaging is primarily used to communicate status, report failures, and test connectivity between network nodes. In carrier Ethernet networks, this is typically referred to as Operation, Administration, and Maintenance (OAM). OAM is important because it eases network operation, verifies network performance, and reduces operational costs. OAM contributes greatly to the level of service received by the subscriber by automatically detecting network degradation or failure, automatically implementing recovery operations when possible, and ensuring that the length of downtime is recorded.

The messages exchanged are known as OAM Protocol Data Units (OAMPDUs). More than 16 OAMPDUs have been defined for many purposes: monitoring status, checking connectivity, detecting failures, reporting failures, localizing errors, looping back data, and preventing security breaches. The International Telecommunication Union (ITU) has defined layers of management domains, allowing a user's network-management traffic to pass through the network while the carrier's OAM manages each point-to-point link. The ITU has also defined interaction between management entities, allowing multiple carriers to seamlessly manage end-to-end flows. The format and usage of OAMPDUs has been jointly defined by the IEEE, ITU, and the Metro Ethernet Forum (MEF). Applicable standards are IEEE 802.3ah and 802.3ag, as well as ITU-T Y.1731 and Y.1730.

Tagging allows the carrier to uniquely identify a customer's data traffic at any location in the carrier's network. Several techniques are used for this purpose: VLAN tagging, MPLS, and GMPLS. All of these techniques insert several bytes of identification into each Ethernet frame at the ingress point (when traffic first enters the network), and remove the information when the frame leaves the network. Each of these techniques also provides functions other than just tagging. For example, VLAN tags also include fields for prioritizing traffic, and MPLS/GMPLS was designed to be used to "switch" traffic (i.e., to determine a frame's destination, and forward it only to the applicable part of the network).

Prioritization can be used when Ethernet frames are buffered at any point in the network. While the frames are waiting in a buffer, the highest priority traffic can be scheduled to be transmitted first. One could visualize prioritization as the rearranging of cars at a stop light. Buffering must occur when the output data rate from a node is less than the input data rate. Usually, these conditions are transients due to network congestion and exist only for very brief periods. If the long-term output rate of a node is less than the input data rate, flow control must be used to exert "backpressure" to slow data from the data source. The latter condition is common at nodes where Local Area Network (LAN) traffic enters a Wide Area Network (WAN) connection, due to the relatively higher cost of bandwidth over long distances. This node is usually called the "Access Node" and plays the most important role in prioritizing traffic. These two concepts, prioritization and flow control, are the cornerstone of what is commonly called Quality of Service (QoS). Many people have the misconception that using prioritization provides a guaranteed "clear pipe" for high priority traffic. In actuality, prioritization and scheduling merely allows "more important" traffic to be delayed the least at buffered nodes. There are several other dimensions that must be considered for properly implemented QoS.

Higher level applications performed by a network node can cover a wide range of purposes. Layer-2 (Data Link Layer) and Layer-3 (Network Layer) applications are most common. Layer-2 applications include protocols that impact node-to-node communications. These include protocols such as Address Resolution Protocols (ARP/RARP/SLARP/GARP), Point-to-Point Protocols (PPP/EAP/SDCP), and Bridging Protocols (BPDU/VLAN). Layer-3 applications include protocols for inter-host communication. These include protocols such as the Bootstrap Protocol (BOOTP), Dynamic Host Configuration Protocol (DHCP), Internet Group Management Protocol (IGMP), and Resource Reservation Protocol (RSVP). Layer-4 (Transport) protocols are occasionally implemented, but normally only to service higher-level applications.

Layer-7 (Application Layer) protocols are occasionally utilized in EoPDH equipment. These include the Hyper-Text Transfer Protocol (HTTP) for serving a HTML user-interface web page, and the Simple Network Management Protocol (SNMP) for providing automated equipment monitoring by the subscriber's network management tools.

EoPDH technology provides a method of transporting native Ethernet frames over the existing telecommunications infrastructure by leveraging the well-established PDH transport technology. The long-term outlook for EoPDH can be evaluated on several metrics:

Service Quality and Dependability Advanced Ethernet OAM increases service quality above that of the underlying DS1/E1 or DS3/E3 transport. Links are monitored, performance degradation and link failures are automatically reported, and recovery operations can be automated. Because the underlying transport is PDH, existing PDH management tools can also be utilized. Over time, the PDH and Ethernet management tools may merge, providing transparency and a single management interface.

Bandwidth Needs and Scalability EoPDH link aggregation allows scaling of bandwidth utilized for transport in increments as small as 1.5Mbps, from 1.5Mbps to 360Mbps. This range covers all access applications on the near horizon, including bandwidth-hungry applications such as IPTV. The use of Committed Information Rate (CIR) circuits at the ingress point allows even finer granularity for bandwidth served to the end customer.

Interoperability and Ease of Use Because EoPDH leverages existing PDH technology, a large infrastructure of knowledge and equipment already exists for PDH tributaries. Trained craftspeople already understand provisioning and maintenance of PDH tributaries, and PDH test equipment is readily available. Legacy equipment can be used to transport, switch, and monitor the PDH tributaries. Because of interoperability, significant cost advantages exist when EoPDH is used in conjunction with legacy SONET/SDH networks. The combination of these technologies is called Ethernet-over-PDH-over-SONET/SDH, or EoPoS. EoPoS reduces cost by allowing reuse of legacy TDM-over-SONET/SDH equipment. Rather than replacing existing SONET/SDH nodes with "next-generation" Ethernet-over-SONET/SDH (EoS) boxes, PDH tributaries can be dropped from legacy ADMs to lower cost CPE or demarcation devices that perform EoPDH VCAT/LCAS link aggregation.

Equipment Cost and Cost of Operation Because existing equipment can be used in the transport network, only access nodes need to be EoPDH enabled. Often, enabling an access node for EoPDH requires adding only a small DSU (modem/media converter). Advanced Ethernet OAM reduces operational cost through link monitoring and rapid fault location. Future equipment may make use of Ethernet-based protocols for self-configuration, greatly simplifying installation. EoPDH not only saves the carrier money, the subscriber service fee for multiple (aggregated) DS1 or E1 connections is usually much less than the service fee for a higher speed connection, such as DS3, saving the carrier's customer money as well.

Applications for EoPDH technology span the realm of telecom equipment. All equipment that places Ethernet frames on a PDH, TDM, or other serial links is able to benefit from the advantages of EoPDH technology.

Example equipment types include remote DSLAMs, cellular backhaul, WAN routers, Ethernet access, multi-tenant access units, and EFM equipment. For more information on how Dallas Semiconductor enables EoPDH technology, visit: [Ethernet Access Products](#).

Application Note 3849: www.maxim-ic.com/an3849

More Information

For technical support: www.maxim-ic.com/support

For samples: www.maxim-ic.com/samples

Other questions and comments: www.maxim-ic.com/contact

Automatic Updates

Would you like to be automatically notified when new application notes are published in your areas of interest?

[Sign up for EE-Mail™](#).

Related Parts

DS33R11: [QuickView](#) -- [Full \(PDF\) Data Sheet](#) -- [Free Samples](#)

DS33R41: [QuickView](#) -- [Full \(PDF\) Data Sheet](#) -- [Free Samples](#)

DS33Z11: [QuickView](#) -- [Full \(PDF\) Data Sheet](#) -- [Free Samples](#)

DS33Z41: [QuickView](#) -- [Full \(PDF\) Data Sheet](#) -- [Free Samples](#)

DS33Z44: [QuickView](#) -- [Full \(PDF\) Data Sheet](#) -- [Free Samples](#)

AN3849, AN 3849, APP3849, Appnote3849, Appnote 3849

Copyright © by Maxim Integrated Products

Additional legal notices: www.maxim-ic.com/legal