

Point-to-Point Protocol

Background

In the late 1980s, the Internet (a large international network connecting many research institutions, government agencies, universities, and private businesses) began to experience explosive growth in the number of hosts supporting the Internet Protocol (IP). The vast majority of these hosts were connected to local-area networks (LANs) of various types, Ethernet being the most common. Most of the other hosts were connected through wide-area networks (WANs) such as X.25-style *public data networks* (PDNs). Relatively few of these hosts were connected with simple point-to-point (that is, serial) links. Yet point-to-point links are among the oldest methods of data communications and almost every host supports point-to-point connections. For example, asynchronous EIA/TIA-232-C (formerly RS-232-C) interfaces are essentially ubiquitous.

One reason for the small number of point-to-point IP links was the lack of a standard Internet encapsulation protocol. The Point-to-Point Protocol (PPP) was designed to solve this problem. In addition to solving the problem of standardized Internet encapsulation of IP over point-to-point links, PPP was also designed to address other issues, including assignment and management of IP addresses, asynchronous (start/stop) and bit-oriented synchronous encapsulation, network protocol multiplexing, link configuration, link quality testing, error detection, and option negotiation for such capabilities as network-layer address negotiation and data compression negotiation. PPP addresses these issues by providing an extensible *Link Control Protocol* (LCP) and a family of *Network Control Protocols* (NCPs) to negotiate optional configuration parameters and facilities. Today, PPP supports other protocols besides IP, including Internetwork Packet Exchange (IPX) and DECnet.

PPP Components

PPP provides a method for transmitting datagrams over serial point-to-point links. It has three main components:

- A method for encapsulating datagrams over serial links—PPP uses the *High-Level Data Link Control* (HDLC) protocol as a basis for encapsulating datagrams over point-to-point links. See Chapter 11, “Synchronous Data Link Control and Derivatives,” for more information on HDLC.
- An extensible LCP to establish, configure, and test the data-link connection.
- A family of NCPs for establishing and configuring different network-layer protocols—PPP is designed to allow the simultaneous use of multiple network-layer protocols.

General Operation

In order to establish communications over a point-to-point link, the originating PPP first sends LCP frames to configure and (optionally) test the data link. After the link has been established and optional facilities have been negotiated as needed by the LCP, the originating PPP sends NCP frames to choose and configure one or more network-layer protocols. When each of the chosen network-layer protocols has been configured, packets from each network-layer protocol can be sent over the link. The link will remain configured for communications until explicit LCP or NCP frames close the link, or until some external event occurs (for example, an inactivity timer expires or a user intervenes).

Physical-Layer Requirements


PPP is capable of operating across any DTE/DCE interface. Examples include EIA/TIA-232-C (formerly RS-232-C), EIA/TIA-422 (formerly RS-422), EIA/TIA-423 (formerly RS-423) and International Telecommunication Union Telecommunication Standardization Sector (ITU-T) (formerly CCITT) V.35. The only absolute requirement imposed by PPP is the provision of a duplex circuit, either dedicated or switched, that can operate in either an asynchronous or synchronous bit-serial mode, transparent to PPP link-layer frames. PPP does not impose any restrictions regarding transmission rate, other than those imposed by the particular DTE/DCE interface in use.

PPP Link Layer

PPP uses the principles, terminology, and frame structure of the International Organization for Standardization (ISO) HDLC procedures (ISO 3309-1979), as modified by ISO 3309:1984/PDAD1 “Addendum 1: Start/stop transmission.” ISO 3309-1979 specifies the HDLC frame structure for use in synchronous environments. ISO 3309:1984/PDAD1 specifies proposed modifications to ISO 3309-1979 to allow its use in asynchronous environments. The PPP control procedures use the definitions and control field encodings standardized in ISO 4335-1979 and ISO 4335-1979/Addendum 1-1979.

The PPP frame format appears in Figure 9-1.

Figure 9-1 PPP Frame Format



The fields of a PPP frame are as follows:

- *Flag*—A single byte that indicates the beginning or end of a frame. The flag field consists of the binary sequence 01111110.
- *Address*—A single byte that contains the binary sequence 11111111, the standard broadcast address. PPP does not assign individual station addresses.

- *Control*—A single byte that contains the binary sequence 00000011, which calls for transmission of user data in an unsequenced frame. A connectionless link service similar to that of Logical Link Control (LLC) Type 1 is provided. For more information about LLC types and frame types, refer to Chapter 11, “Synchronous Data Link Control and Derivatives.”
- *Protocol*—2 bytes that identify the protocol encapsulated in the information field of the frame. The most up-to-date values of the protocol field are specified in the most recent *Assigned Numbers Request for Comments* (RFC).
- *Data*—Zero or more bytes that contain the datagram for the protocol specified in the protocol field. The end of the information field is found by locating the closing flag sequence and allowing 2 bytes for the FCS field. The default maximum length of the information field is 1,500 bytes. By prior agreement, consenting PPP implementations can use other values for the maximum information field length.
- *Frame check sequence* (FCS)—Normally 16 bits (2 bytes). By prior agreement, consenting PPP implementations can use a 32-bit (4-byte) FCS for improved error detection.

The LCP can negotiate modifications to the standard PPP frame structure. However, modified frames will always be clearly distinguishable from standard frames.

PPP Link Control Protocol

The PPP LCP provides a method of establishing, configuring, maintaining, and terminating the point-to-point connection. LCP goes through four distinct phases:

- Link establishment and configuration negotiation—Before any network-layer datagrams (for example, IP) can be exchanged, LCP must first open the connection and negotiate configuration parameters. This phase is complete when a configuration acknowledgment frame has been both sent and received.
- Link quality determination—LCP allows an optional link quality determination phase following the link establishment and configuration negotiation phase. In this phase, the link is tested to determine if the link quality is sufficient to bring up network-layer protocols. This phase is optional. LCP can delay transmission of network-layer protocol information until this phase is completed.
- Network-layer protocol configuration negotiation—After LCP has finished the link quality determination phase, network-layer protocols can be separately configured by the appropriate NCP and can be brought up and taken down at any time. If LCP closes the link, it informs the network-layer protocols so that they may take appropriate action.
- Link termination—LCP can terminate the link at any time. This will usually be done at the request of a user, but can happen because of a physical event such as the loss of carrier or the expiration of an idle-period timer.

There are three classes of LCP frames:

- Link establishment frames—Used to establish and configure a link.
- Link termination frames—Used to terminate a link.
- Link maintenance frames—Used to manage and debug a link.

These frames are used to accomplish the work of each of the LCP phases.

