

Point-to-Point Protocol

The Point-to-Point Protocol (PPP) is designed for simple links which transport packets between two peers. These links provide full-duplex simultaneous operation, and are assumed to deliver packets in order. PPP provides a common solution for easy connection of a wide variety of hosts, bridges and routers.

PPP is a complete specification for transmitting datagrams between data communications equipment from different manufacturers over dial-up and dedicated serial point-to-point links. It is a recommended standard of the Internet Advisory Board (IAB) and is contained in a number of request for comments (RFCs) produced by the PPP Protocol Working Group.

Traditionally, interoperability across serial links was restricted to equipment supplied by the same manufacturer. Now, PPP allows for multi-vendor interoperability.

PPP was first proposed as a standard in 1990 to replace an older de facto standard known as Serial Line Internet Protocol (SLIP), which requires links to be established and torn down manually. However, unlike SLIP which only supports IP, PPP is not limited in protocol support. PPP provides the flexibility to add support for other protocols through software upgrades. PPP can also simultaneously transmit multiple protocols across a single serial link, limiting the need to set up a separate link for each protocol. PPP is also ideal for interconnecting dissimilar devices such as hosts, bridges, and routers over serial links. For example, a standalone TCP/IP host can communicate with a router across a serial PPP link.

PPP Components

The three main components of PPP are the:

- Encapsulation scheme
- Link Control Protocol (LCP)
- Network Control Protocol (NCP)

These components are responsible for creating the frame, controlling the link, and managing the network layer protocol respectively.

Main Components of PPP

- Encapsulation Scheme**
- Link Control Protocol**
- Network Control Protocols**

Encapsulation

Standard encapsulation schemes exist for the transmission of datagrams over most Local Area Networks (LANs) such as Ethernet, Token Ring, ARCnet, and FDDI. In the past, the only Wide Area Network (WAN) encapsulation scheme that provided a standard for the transmission of datagrams was X.25. The introduction of new WAN schemes, such as Frame Relay, expanded the variety of encapsulation schemes available. However, the majority of LAN-to-LAN traffic is still carried over dedicated leased lines. The introduction of PPP allows these existing proprietary leased lines the opportunity to convert to a new encapsulation scheme that gives the user the true interoperability that traditionally could only be found on LANs.

PPP is a full-duplex, bit-oriented protocol that can run over synchronous or asynchronous links. PPP uses a variant of HDLC as the basis for encapsulation. Links may be dedicated or circuit-switched, and PPP can work over copper, fiber optic, microwave, or satellite leased lines. PPP provides data error detection while higher layer protocols are responsible for error recovery. There are three parts to PPP encapsulation and these are:

- Protocol field
- Information field
- Padding

Protocol Field

The protocol field is one or two octets, and its value identifies the datagram encapsulated in the information field of the packet.

Information


Field The information field is zero or more octets. The information field contains the datagram for the protocol specified in the protocol field.

The maximum length of the information field, including padding is termed the maximum receive unit (MRU), which defaults to 1500 octets. By negotiation, consenting PPP implementations may use other values for the MRU.

Padding

On transmission, the information field may be padded with an arbitrary number of octets up to the MRU. It is the responsibility of each protocol to distinguish padding octets from real information.

PPP Encapsulation



LCP

The second component of PPP is the LCP protocol, which operates at the data link layer to manage communication functions.

In order to establish communications over a point-to-point link, each end of the PPP link must first send LCP packets. LCP, using a four phase process, establishes the link between two PPP peers and negotiates configuration options. Only phases one (link establishment and negotiation) and four (link ready) are necessary to establish communications. Phases two (authentication) and three (link quality determination) are optional and completely dependent on the PPP implementations at both ends of the link.

LCP Management Functions

The LCP management functions do the following:

- Determine encapsulation format options
- Negotiate optimal packet size
- Terminate the link
- Authenticate the identity of the peer on the link (optional)
- Negotiate PPP multilink data compression (optional)
- Link quality monitoring (optional)

In addition to providing procedures for establishing, configuring, testing, and terminating data link connections, LCP also negotiates other non-default LCP options such as the MRU.

NCP

NCPs are a series of independently-defined protocols that encapsulate network layer protocols such as TCP/IP, DECnet, AppleTalk, IPX, XNS, and OSI. Each NCP has individual requirements for addressing and advertising connectivity for its network layer protocol. Each NCP is defined in a separate RFC. Future protocols can be supported by defining new NCPs.

After the link has been established (LCP packets have been exchanged), PPP must send NCP packets to choose and configure one or more network-layer protocol. Once each of the chosen network-layer protocols has been configured, datagrams from each can be sent.

The link will remain configured for communications until explicit LCP or NCP packets close the link down, or until some external event occurs (an inactivity timer expiry or network administrator intervention).

Link Operation

In the process of configuring, maintaining and terminating the point-to-point link, the PPP link goes through several distinct phases:

- Link dead (physical layer not ready)
- Link establishment
- Authentication
- Network-layer protocol phase
- Link termination

Link Dead When an external event, such as carrier detection or network administrator configuration, indicates that the physical layer is ready to be used, PPP proceeds to the link establishment phase.

Link Establishment Phase

The LCP is used to establish the connection through an exchange of configuration packets:

- Send Configure Request (SCR)
- Send Configure Ack (SCA)
- Send Configure Nak (SCN)

This exchange is complete, and the LCP opened state entered, once an SCA packet has been both sent and received.

All configuration options are assumed to be at default values unless altered by the configuration exchange.

Authentication Phase

On some links it may be desirable to require a peer to authenticate itself before allowing network-layer protocol packets to be exchanged.

By default authentication is not mandatory. If an implementation desires that the peer authenticate with some specific authentication protocol, then it must request the use of the authentication protocol during the link establishment phase.

Network Layer Protocol Phase

Once PPP has finished the previous phases, each network-layer protocol, such as IP, IPX or AppleTalk, must be separately configured by the appropriate NCP.

After an NCP has reached the opened state, PPP carries the corresponding network layer packets. During this phase, link traffic consists of any possible combination of LCP, NCP and network-layer packets.

Link Termination Phase

PPP can terminate the link at any time. This might happen because of loss of carrier, authentication failure or link-quality determination. LCP is used to close the link through an exchange of terminate packets.

PAP

Packet-level Procedure (PAP) is a protocol for the transfer of packets between an X.25 DTE and an X.25 DCE. PAP is a full-duplex protocol that supports data sequencing, flow control, accountability, and error detection and recovery.

CHAP

Challenge Handshake Application Protocol (CHAP) is an authentication method that can be used, for example when connecting to an Internet service provider (ISP). CHAP allows you to log in to your provider automatically, without the need for a terminal screen. It is more secure than PAP.

CHAP is used to periodically verify the identity of the peer using a three-way handshake. This is done upon initial link establishment, and may be repeated anytime after the link has been established.

PPP Over ISDN

Since the ISDN B-channel is by definition a point-to-point circuit link, PPP is well suited for use over these links.

PPP treats ISDN channels as bit- or octet-oriented synchronous links. These links must be full-duplex, but may be either dedicated or circuit-switched. PPP presents an octet interface to the physical layer. The octet stream is applied primarily at the R or T reference points.

Transmission Rate

PPP does not impose any restrictions regarding transmission rate.

PPP Configuration

With a multitude of encapsulation, LCP, and NCP options available for PPP, configuration of the link could be a major concern. However, configuration can be simplified in bridges or routers by having a default configuration which automatically allows peers across a link to negotiate each option. Having such a default configuration makes PPP one of the easiest data link protocols to configure. PPP eliminates manual, time-consuming configuration. It is virtually plug and play.

PPP operability Testing

Inter-Setting up a PPP connection between two peers involves four phases:

- LCP negotiation
- LCP steady state
- NCP negotiation
- Network layer protocol data flow

When interconnecting equipment from different manufacturers, problems can occur at any of these four stages. In general, the closer two peers get towards the fourth stage, the more compatible they are. However, even reaching the fourth stage doesn't guarantee proper operation since basic differences in functionality can prevent them from communicating.