

Network Working Group
Request for Comments: 1598
Category: Standards Track

W. Simpson
Daydreamer
March 1994

PPP in X.25

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Abstract

The Point-to-Point Protocol (PPP) [1] provides a standard method for transporting multi-protocol datagrams over point-to-point links. This document describes the use of X.25 for framing PPP encapsulated packets.

This document is the product of the Point-to-Point Protocol Working Group of the Internet Engineering Task Force (IETF). Comments should be submitted to the ietf-ppp@merit.edu mailing list.

Applicability

This specification is intended for those implementations which desire to use facilities which are defined for PPP, such as the Link Control Protocol, Network-layer Control Protocols, authentication, and compression. These capabilities require a point-to-point relationship between peers, and are not designed for multi-point or multi-access environments.

Table of Contents

1.	Introduction	1
2.	Physical Layer Requirements	2
3.	The Data Link Layer	2
3.1	Frame Format	3
3.2	Modification of the Basic Frame	3
4.	Call Setup	4
5.	Configuration Details	5
	SECURITY CONSIDERATIONS	6
	REFERENCES	6
	ACKNOWLEDGEMENTS	6
	CHAIR'S ADDRESS	7
	AUTHOR'S ADDRESS	7

1. Introduction

CCITT recommendation X.25 [2] describes a network layer protocol providing error-free, sequenced, flow controlled, virtual circuits. X.25 includes a data link layer, X.25 LAPB, which uses ISO 3309, 4335 and 6256.

PPP also uses ISO 3309 HDLC as a basis for its framing [3].

When X.25 is configured as a point-to-point circuit, PPP can use X.25 as a framing mechanism, ignoring its other features. This is equivalent to the technique used to carry SNAP headers over X.25 [4].

At one time, it had been hoped that PPP HDLC frames and X.25 frames would co-exist on the same links. Equipment could gradually be converted to PPP. Subsequently, it has been learned that some switches actually remove the X.25 header, transport packets to another switch using a different protocol such as Frame Relay, and reconstruct the X.25 header at the final hop. Co-existence and gradual migration are precluded.

2. Physical Layer Requirements

PPP treats X.25 framing as a bit synchronous link. The link **MUST** be full-duplex, but **MAY** be either dedicated (permanent) or switched.

Interface Format

PPP presents an octet interface to the physical layer. There is no provision for sub-octets to be supplied or accepted.

Transmission Rate

PPP does not impose any restrictions regarding transmission rate, other than that of the particular X.25 interface.

Control Signals

Implementation of X.25 requires the provision of control signals, which indicate when the link has become connected or disconnected. These in turn provide the Up and Down events to the LCP state machine.

Because PPP does not normally require the use of control signals, the failure of such signals **MUST NOT** affect correct operation of PPP. Implications are discussed in [2].

Encoding

The definition of various encodings is the responsibility of the DTE/DCE equipment in use, and is outside the scope of this specification.

While PPP will operate without regard to the underlying representation of the bit stream, X.25 requires NRZ encoding.

3. The Data Link Layer

This specification uses the principles, terminology, and frame structure described in "Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode" [4].

The purpose of this specification is not to document what is already standardized in [4]. Instead, this document attempts to give a concise summary and point out specific options and features used by PPP.

3.1. Frame Format

Since both "PPP in HDLC Framing" [3] and X.25 use ISO 3309 as a basis for framing, the X.25 header is easily substituted for the smaller HDLC header. The fields are transmitted from left to right.

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+
|  Flag (0x7e)  |
+-----+-----+-----+-----+
|   Address     |   Control   |D|Q| SVC# (hi) |   SVC# (lo)   |
+-----+-----+-----+-----+
|p(r) |M|p(s) |0|           PPP Protocol           |
+-----+-----+-----+-----+

```

The PPP Protocol field and the following Information and Padding fields are described in the Point-to-Point Protocol Encapsulation [1].

3.2. Modification of the Basic Frame

The Link Control Protocol can negotiate modifications to the basic frame structure. However, modified frames will always be clearly distinguishable from standard frames.

Address-and-Control-Field-Compression

Because the Address and Control field values are not constant, and are modified as the frame is transported by the network switching fabric, Address-and-Control-Field-Compression MUST NOT be negotiated.

Protocol-Field-Compression

Note that unlike the HDLC framing, the X.25 framing does not align the Information field on a 32-bit boundary. Alignment to a 16-bit boundary occurs when the Protocol field is compressed to a single octet. When this improves throughput, Protocol-Field-Compression SHOULD be negotiated.

4. Call Setup

When the link is configured as a Permanent Virtual Circuit (PVC), support for Switched Virtual Circuit (SVC) call setup and clearing is not required. Calls are Established and Terminated using PPP LCP packets.

When the link is configured as a Switched Virtual Circuit (SVC), the first octet in the Call User Data (CUD) Field (the first data octet in the Call Request packet) is used for protocol demultiplexing, in accordance with the Subsequent Protocol Identifier (SPI) in ISO/IEC TR 9577 [5]. This field contains a one octet Network Layer Protocol Identifier (NLPID), which identifies the encapsulation in use over the X.25 virtual circuit. The CUD field MAY contain more than one octet of information.

The PPP encapsulation MUST be indicated by the PPP NLPID value (CF hex). Any subsequent octet in this CUD is extraneous and MUST be ignored.

Multipoint networks (or multicast groups) MUST refuse calls which indicate the PPP NLPID in the CUD.

The accidental connection of a link to feed a multipoint network (or multicast group) SHOULD result in a misconfiguration indication. This can be detected by multiple responses to the LCP Configure-Request with the same Identifier, coming from different framing addresses. Some implementations might be physically unable to either log or report such information.

Conformance with this specification requires that the PPP NLPID (CF) be supported. In addition, conformance with [4] requires that the IP NLPID (CC) be supported, and does not require that other NLPID values be supported, such as Zero (00), SNAP (80), CLNP (81) or ES-IS (82).

When IP address negotiation and/or VJ header compression are desired, the PPP call setup SHOULD be attempted first. If the PPP call setup fails, the normal IP call setup MUST be used.

The PPP NLPID value SHOULD NOT be used to demultiplex circuits which use the Zero NLPID in call setup, as described in [4]. When such a circuit exists concurrently with PPP encapsulated circuits, only network layer traffic which has not been negotiated by the associated NCP is sent over the Zero NLPID circuit.

Rationale:

Using call setup to determine if PPP is supported should be

inexpensive, when users aren't charged for failed calls.

Using the Zero NLPID call together with PPP could be expensive, when users are charged per packet or for connect time, due to the probing of PPP configuration packets at each call.

PPP configuration provides a direct indication of the availability of service, and on that basis is preferred over the Zero NLPID technique, which can result in "black-holes".

5. Configuration Details

The following Configuration Options are recommended:

- Magic Number
- Protocol Field Compression

The standard LCP configuration defaults apply to X.25 links, except MRU.

To ensure interoperability with existing X.25 implementations, the initial Maximum-Receive-Unit (MRU) is 1600 octets [4]. This only affects the minimum required buffer space available for receiving packets, not the size of packets sent.

The typical network feeding the link is likely to have a MRU of either 1500, or 2048 or greater. To avoid fragmentation, the Maximum-Transmission-Unit (MTU) at the network layer SHOULD NOT exceed 1500, unless a peer MRU of 2048 or greater is specifically negotiated.

The X.25 packet size is not directly related to the MRU. Instead, Protocol Data Units (PDUs) are sent as X.25 "complete packet sequences". That is, PDUs begin on X.25 data packet boundaries and the M bit ("more data") is used to fragment PDUs that are larger than one X.25 data packet in length.

Security Considerations

Implementations MUST NOT consider PPP authentication on call setup for one circuit between two systems to apply to concurrent call setup for other circuits between those same two systems. This results in possible security lapses due to over-reliance on the integrity and security of switching systems and administrations. An insertion attack might be undetected. An attacker which is able to spoof the same calling identity might be able to avoid link authentication.

References

- [1] Simpson, W., Editor, "The Point-to-Point Protocol (PPP)", RFC 1548, December 1993.
- [2] CCITT Recommendation X.25, "Interface Between Data Terminal Equipment (DTE) and Data Circuit Terminating Equipment (DCE) for Terminals Operating in the Packet Mode on Public Data Networks", Vol. VIII, Fascicle VIII.2, Rec. X.25.
- [3] Simpson, W., Editor, "PPP in HDLC Framing", RFC 1549, December 1993.
- [4] Malis, A., Robinson, D., and R. Ullmann, "Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode", RFC 1356, August 1992.
- [5] ISO/IEC TR 9577, "Information technology - Telecommunications and Information exchange between systems - Protocol Identification in the network layer", 1990 (E) 1990-10-15.

Acknowledgments

This design was inspired by the paper "Parameter Negotiation for the Multiprotocol Interconnect", Keith Sklower and Clifford Frost, University of California, Berkeley, 1992, unpublished.

Chair's Address

The working group can be contacted via the current chair:

Fred Baker
Advanced Computer Communications
315 Bollay Drive
Santa Barbara, California 93117

EMail: fbaker@acc.com

Author's Address

Questions about this memo can also be directed to:

William Allen Simpson
Daydreamer
Computer Systems Consulting Services
1384 Fontaine
Madison Heights, Michigan 48071

EMail: Bill.Simpson@um.cc.umich.edu
bsimpson@MorningStar.com

