

PPP Reliable Transmission

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 defines a method for negotiating and using Numbered-Mode, as defined by ISO 7776 [2], to provide a reliable serial link.

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@ucdavis.edu mailing list.

Table of Contents

1.	Introduction	1
2.	Physical Layer Requirements	2
3.	The Data Link Layer	2
3.1	Frame Format	2
4.	Configuration Option Format	4
5.	Numbered-Mode Operation	5
5.1	Single Link	6
5.2	Inverse Multiplexing	6
5.3	Using Multi-Link Procedure... ..	7
5.4	LAPB Parameter defaults	8
	SECURITY CONSIDERATIONS	9
	REFERENCES	9
	ACKNOWLEDGEMENTS	9
	CHAIR'S ADDRESS	10
	AUTHOR'S ADDRESS	10

1. Introduction

By default, PPP packets over HDLC framed links consist of "connectionless" datagrams. If reliable transmission over the HDLC link is desired, the implementation MUST specify the Numbered-Mode Configuration Option during Link Establishment phase.

Generally, serial link reliability is not a major issue. The architecture of protocols used in datagram networking presume best-effort non-sequential delivery. When errors are detected, datagrams are discarded.

However, in certain circumstances, it is advisable to provide a reliable link, at least for a subset of the messages. The most obvious case is when the link is compressed. Since the dictionary is recovered from the compressed data stream, and a lost datagram corrupts the dictionary, datagrams must not be lost. Not all compression types will require a reliable data stream, since the cost to detect and reset a corrupt dictionary is small.

The ISO 7776 LAPB can be used guarantee delivery. This is referred to in this document as "Numbered Mode" to distinguish it from the use of "Unnumbered Information", which is standard PPP framing practice.

Where multiple parallel links are used to emulate a single link of higher speed, Bridged traffic, Source Routed traffic, and traffic subjected to Van Jacobsen TCP/IP header compression must be delivered to the higher layer in a certain sequence. However, the fact of the links being relatively asynchronous makes traffic ordering uncertain.

The ISO 7776 Multi-Link Procedure MAY be used to restore order. Implementation of the ISO Multi-Link Procedure is deprecated. It is recommended that the PPP multilink procedure [4] be used instead.

2. Physical Layer Requirements

PPP Reliable Transmission imposes the same requirements that are described in "PPP in HDLC Framing" [3], with the following exceptions.

Control Signals

While PPP does not normally require the use of control signals, implementation of Numbered-Mode LAPB or LAPD 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.

3. The Data Link Layer

Numbered-Mode affects only the Address and Control fields. The remainder of the frame conforms to the framing in use for PPP.

The Address Field of the frame MUST take the value announced in the Numbered-Mode Configuration Option, and the Control Field MAY take any value valid in ISO 7776.

Once the link enters Numbered-Mode, Numbered-Mode MUST be used on all frames, as some implementations do not support the use of the Unnumbered-Information control field or the use of the All-Stations address intermixed with Numbered-Mode frames.

3.1. Frame Format

The following frame format is valid under Numbered-Mode. The fields are transmitted from left to right.

Numbered Mode

+-----+-----+-----+		
Flag	Address	Control
01111110	1-2 octets	1-2 octets
+-----+-----+-----+		
+-----+-----+-----+		
Protocol	Information	Padding
1-2 octets	*	*
+-----+-----+-----+		
+-----+-----+-----+		
FCS	Flag	Inter-frame Fill
16 bits	01111110	or next Address
+-----+-----+-----+		

The Protocol, Information and Padding fields are described in the Point-to-Point Protocol Encapsulation [1]. The FCS and Flag Sequence fields are described in "PPP in HDLC Framing" [3].

4. Configuration Option Format

Description

The LCP Numbered-Mode Configuration Option negotiates the use of Numbered-Mode on the link. By default or ultimate disagreement, Unnumbered-Mode is used.

A summary of the Numbered-Mode Configuration Option format is shown below. 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
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Type      |      Length      |      Window      |      Address...
+-----+-----+-----+-----+-----+-----+-----+

```

Type

11

Length

>= 4

Window

A value between 1 and 127. This indicates the number of frames the receiver will buffer, which is the maximum number that the sender should send without receiving an acknowledgement. If window < 8, then modulo 8 sequencing is used on the link. Otherwise, modulo 128 sequencing is used.

It is conceivable and legal that differing window values might be announced. However, it is not permitted for one system to use modulo 8 sequencing and the other to use modulo 128. Therefore, the rule is: a Configure-Nak may reduce the window but may not increase it.

Address

An HDLC Address as specified in ISO 3309. ISO 7776 specifies four of the possible values: 1 and 3 for single link operation, 7 and 15 for the Multi-Link Procedure. Other values consistent with ISO 3309 are considered legal.

Implementation of the Multi-Link Procedure is optional; A

Configure-Nak may therefore force a change from MLP to single link mode, but not the reverse.

Should the address be zero upon receipt, the receiver MUST Configure-Nak with an appropriate address. If both peers send address zero, the system advertising the numerically smaller window will select the smaller address. If both windows are the same size, a random choice MUST be made; when good sources of randomness are used, the link will converge in a reasonable time.

If magic numbers have been negotiated on the link, the system with the numerically smaller magic number SHOULD specify the smaller address.

5. Numbered-Mode Operation

When using the Numbered-Mode, each link is established in the usual manner for the type of link. The Numbered-Mode Configuration Option is negotiated, the Magic-Number Configuration Option MUST also be negotiated, and the Address-and-Control-Field-Compression Configuration Option MUST NOT be negotiated.

Following the successful negotiation of the Numbered-Mode Configuration Option during LCP Link Establishment phase, the system with the numerically smaller Magic-Number will send a SABM or SABM(E), and the other will respond with a UA. In the event that either the SABM or UA is lost, this exchange may be repeated according to the same parameters as the configuration exchange itself, using the Restart Timer and counter values. Authentication, Link Quality Determination, and NCP Configuration follow this step.

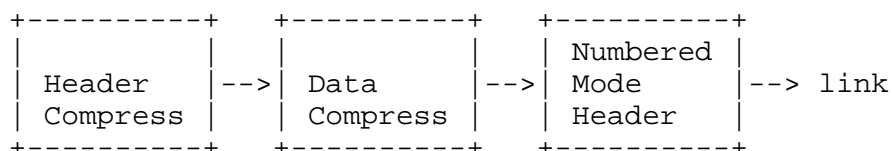
Once the link has been established with Numbered-Mode, when re-negotiation of link configuration occurs, the entire re-negotiation MUST be conducted in Numbered-Mode. If the Numbered-Mode Configuration Option is not successfully re-negotiated, the link reverts to Unnumbered-Information operation prior to Authentication, Link Quality Determination, and NCP Configuration.

When an implementation which is capable of Numbered-Mode, and is not currently configured for Numbered-Mode operation, detects a frame which has a correct FCS but does not have a UI Control octet, the implementation MUST send a DM message, immediately followed by a LCP Configure-Request.

When an implementation which is currently configured for Numbered-Mode operation receives a DM message, it MUST revert to Unnumbered-Information operation, and immediately send a LCP Configure-Request.

5.1. Single Link

When Network-Layer packets are sent over a single link, the packets are encapsulated in the following order:

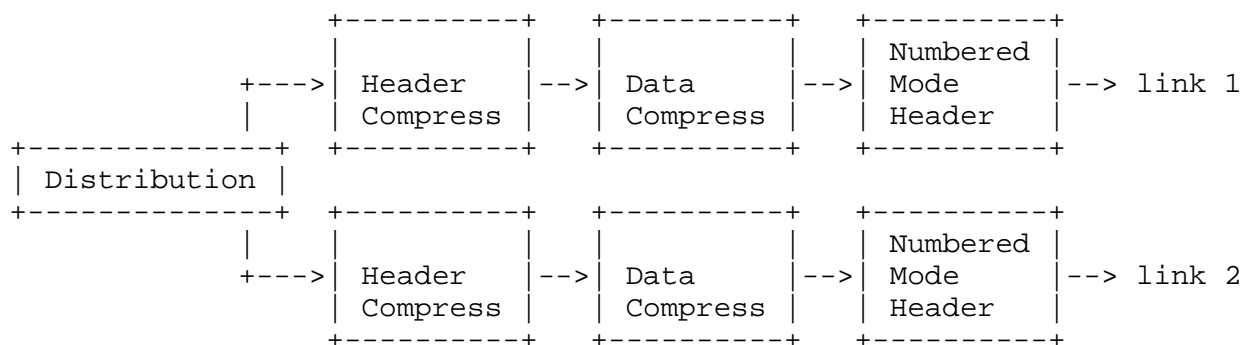


5.2. Inverse Multiplexing

Since sending several connections over a single link is often called "multiplexing", sending packets from a single connection over multiple parallel links is sometimes called "inverse-multiplexing". By default, PPP performs no special processing for such links. Each link is established and terminated independently, negotiates its own configuration options, and may have different combinations of such options as ACCM, Protocol Field Compression and IP-Address. This facilitates using the links simultaneously over dissimilar media, such as 56K sync with async backup.

Every link in a single machine **MUST** have different Magic Numbers, and each end of every link between two peers **SHOULD** have Magic Numbers which are unique to those peers. This protects against patch-panel errors in addition to looped-back links.

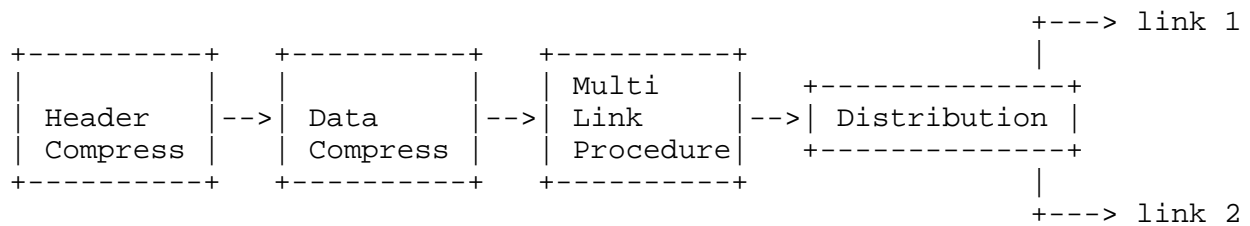
The distribution to each link is controlled by higher level routing mechanisms. When Network-Layer specific compression techniques (such as Van Jacobsen Compression) rely on sequential delivery, without Multi-Link Procedure support such compression **MUST** be applied on a link by link basis.



5.3. Using Multi-Link Procedure

This document does not offer a standard for ISO Multi-Link, but does offer a method for agreeing on the addressing scheme usable with Multi-Link. A sample implementation is shown below. Implementation of Multi-Link is not required.

When using the ISO 7776 Multi-Link Procedure, each link is established as described above. In addition, the Numbered-Mode Configuration Option is negotiated with appropriate addresses for the Multi-Link Procedure. The distribution to each link is controlled by the Multi-Link Procedure, as is the recovery of sequence in the receiving system.



5.4. LAPB Parameter defaults

The following guidelines specify the default values of LAPB configurable parameters.

Timer T1

Timer T1 is the maximum time permitted before a retransmission is started, as a result of no response to a transmitted I frame. This value must be greater than the time required for a maximum sized frame to be received by the other side of the link, and for a response to be generated for the frame. This SHOULD be determined dynamically, based on the measured round trip time delay of the link at the LAPB level. In the event that the system cannot determine the round trip time of the link, this value SHOULD be set to twice the bit rate of the link, divided by the maximum number of bits per frame, plus 100 milliseconds processing time. For example, on a 14,400 bps link, with a maximum frame size of 8000 bits (1000 octets), the T1 value would be set to 3.7 seconds.

Timer T3

Timer T3 gives an indication of the idle state of the link. Its value must be greater than the T1 value.

Maximum number of attempts to complete a transmission, N2

Parameter N2 gives the maximum number of retransmission attempts for a given frame. If this value is exceeded, the link SHOULD be terminated. The default value for parameter N2 SHOULD be 3.

Security Considerations

Security issues are not discussed in this memo.

References

- [1] Simpson, W., Editor, "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, Daydreamer, July 1994.
- [2] ISO 7776, Information Processing Systems - Data Communication - High Level Data Link Control Procedures - Description of the X.25 LAPB-Compatible DTE Data Link Procedures
- [3] Simpson, W., Editor, "PPP in HDLC Framing", STD 51, RFC 1662, Daydreamer, July 1994.
- [4] Sklower, K., "PPP MultiLink Procedure", Work in Progress.

Acknowledgments

Fred Baker was the original author of this document.

Bill Simpson contributed materially to the document.

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:

Dave Rand
2180 Fortune Drive
San Jose, CA 95131

Phone: +1 408 321-1259
EMail: dave_rand@novell.com

