

The PPP OSI Network Layer Control Protocol (OSINLCP)

Status of this Memo

This RFC specifies an IAB standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "IAB Official Protocol Standards" 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 of encapsulating Network Layer protocol information over point-to-point links. PPP also defines an extensible Link Control Protocol, and proposes a family of Network Control Protocols (NCPs) for establishing and configuring different network-layer protocols.

This document defines the NCP for establishing and configuring OSI Network Layer Protocols.

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

Table of Contents

1.	Introduction	2
1.1	OSI Network Layer Protocols over PPP	2
2.	A PPP Network Control Protocol (NCP) for OSI	5
2.1	Sending OSI NPDUs	6
2.2	NPDU Alignment	6
2.3	Network Layer Addressing Information	6
3.	OSINLCP Configuration Options	7
3.1	Align-NPDU	7
	REFERENCES	9
	ACKNOWLEDGEMENTS	9
	SECURITY CONSIDERATIONS	10
	CHAIR'S ADDRESS	10
	AUTHOR'S ADDRESS	10

1. Introduction

PPP has three main components:

1. A method for encapsulating datagrams over serial links.
2. A Link Control Protocol (LCP) for establishing, configuring, and testing the data-link connection.
3. A family of Network Control Protocols (NCPs) for establishing and configuring different network-layer protocols.

In order to establish communications over a point-to-point link, each end of the PPP link must first send LCP packets to configure and test the data link. After the link has been established and optional facilities have been negotiated as needed by the LCP, PPP must send NCP packets to choose and configure one or more network-layer protocols. Once each of the chosen network-layer protocols has been configured, datagrams from each network-layer protocol can be sent over the link.

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 expires or network administrator intervention).

1.1. OSI Network Layer Protocols over PPP

A number of protocols have been defined for the Network Layer of OSI, including the Connectionless Network Layer Protocol (CLNP, ISO 8473) [3], the End System to Intermediate System routing protocol (ES-IS, ISO 9542) [4], the Intermediate System to Intermediate System routing protocol (IS-IS, ISO 10589) [5], and the Inter-Domain Routing Protocol (IDRP, CD 10747) [6]. Generally, these protocols were designed to run over non-reliable data link protocols such as PPP.

Network Layer Protocol Identifier (NLPID)

OSI Network Layer protocols can be discriminated according to the first octet in each Network Protocol Data Unit (NPDU, that is, packet), known as the Network Layer Protocol Identifier (NLPID), which is defined in ISO/TR 9577 [7]. This allows the various protocols to be run over a common data link without any discriminator below the network layer.

Inactive Network Layer Protocol

ISO/TR 9577 reserves a NLPID value of zero to represent the "Inactive Network Layer Protocol", as defined in ISO 8473. The inactive network layer protocol MUST NOT be used over PPP. This assures that whichever OSI network layer protocol is used will have a non-zero NLPID value.

Connection-Oriented Network Protocol

The OSI Connection-Oriented Network Protocol (ISO 8208) [8], effectively the Packet Layer of CCITT X.25, is intended to be run over a reliable data link, such as IEEE 802.2 type II or LAPB. Therefore, the unreliable data link service provided by PPP is not appropriate for use with ISO 8208.

ConnectionLess Network Protocol (CLNP)

The ConnectionLess Network Protocol offers a simple non-reliable datagram service very similar to IP, and is designed to run over a non-reliable data link service, such as provided by PPP.

End-System to Intermediate-System Protocol (ES-IS)

ES Hellos and IS Hellos are retransmitted on a periodic timer-driven basis (based on expiration of the "Configuration Timer"). The resulting ES and IS configuration information is invalidated on a timer driven basis, based on expiration of the "Holding Timer" for each piece of information. The value of a Holding Timer is set by the source of the information, and transmitted in the Holding Time field of the appropriate ES-IS packet. ISO 9542 recommends that the holding time field is set to approximately twice the Configuration Timer parameter, such that even if every other Hello packet is lost the configuration information will be retained (implying that the Holding Timer is actually set to slightly more than twice the Configuration Timer).

Generally, the recommendation in ISO 9542 is sufficient for PPP links. For very unreliable links, it may be necessary to set the Holding Timer to be slightly more than three times the Configuration Timer to ensure that loss of configuration information is an unusual event.

Redirect information is not transmitted on point-to-point links, but may be transmitted on general topology subnetworks, which in turn may make use of PPP. Redirect information is sent on an event-driven basis (based on a CLNP packet being forwarded by a router out the incoming interface), but redirect information is

invalidated on a timer-driven basis. Loss of a single redirect may result in a subsequent data packet being sent to the same incorrect router, which will re-issue the redirect. This operates in the same manner as ICMP redirects for IP packets, and does not pose any problem for operation over PPP links.

Intermediate-System to Intermediate-System Protocol (IS-IS)

IS-IS allows for broadcast links (typically LANs), point-to-point links (such as PPP), and general topology links (such as X.25 networks) which are modelled as a collection of point-to-point links.

There are four types of IS-IS packets: IS-IS Hello Packets, Link State Packets (LSPs), Complete Sequence Number Packets (CSNPs), and Partial Sequence Number Packets (PSNPs).

IS-IS Hello messages are transmitted periodically on point-to-point links (based on expiration of the "ISISHello" timer). Routers expect to receive IS-IS Hello packets periodically. Specifically, the IS-IS Hello packet specifies a "Holding Time". If no subsequent IS-IS Hello is received over the corresponding link for the specified time period, then the neighboring router is assumed to have been disconnected or to be down. It is highly undesirable for links to "flap" up and down unnecessarily, which implies that the holding time needs to be large enough that a link is very unlikely to be declared down due to a failure to receive an IS-IS Hello. This implies that running IS-IS over unreliable data links requires the Holding time to be greater than "k" times the ISISHello timer, where k is chosen such that the loss of k consecutive IS-IS Hello's is rare. If the quality of the link is poor, then the Holding Time will need to be increased or the "ISISHello" time decreased.

LSPs are acknowledged by the IS-IS protocol (via use of partial sequence number packets). A lost LSP will be recovered from with no problem provided that PPP links are treated the same way as other point-to-point links. On those rare occasions where a partial sequence number packet is lost, this might result in the retransmission of a link state packet over a single link, but will not impact the correct operation of the routing algorithm.

CSNPs are sent upon link startup on a point-to-point link. This does not need to be changed for PPP. If a CSNP fragment is lost upon startup it is merely loss of an optimization -- LSPs that did not need to be transmitted over the link will be transmitted. If a periodic CSNP fragment is lost it merely means that detection of low probability database corruption will take longer.

PSNPs function as ACKs. Loss of a PSNP may result in an unnecessary retransmission of an LSP, but does not prevent correct operation of the routing protocol.

Inter-Domain Routing Protocol (IDRP)

IDRP expects to run over datagram links, but requires reliable exchange of IDRP information. For this reason, IDRP contains built-in reliability mechanisms which ensure that packets will be received correctly.

2. A PPP Network Control Protocol (NCP) for OSI

The OSI Network Layer Control Protocol (OSINLCP) is responsible for configuring, enabling, and disabling the OSI protocol modules on both ends of the point-to-point link. OSINLCP uses the same packet exchange mechanism as the Link Control Protocol (LCP). OSINLCP packets may not be exchanged until PPP has reached the Network-Layer Protocol phase. OSINLCP packets received before this phase is reached should be silently discarded.

The OSI Network Layer Control Protocol is exactly the same as the Link Control Protocol [1] with the following exceptions:

Frame Modifications

The packet may utilize any modifications to the basic frame format which have been negotiated during the Link Establishment phase.

Data Link Layer Protocol Field

Exactly one OSINLCP packet is encapsulated in the Information field of a PPP Data Link Layer frame where the Protocol field indicates type hex 8023 (OSI Network Layer Control Protocol).

Code field

Only Codes 1 through 7 (Configure-Request, Configure-Ack, Configure-Nak, Configure-Reject, Terminate-Request, Terminate-Ack and Code-Reject) are used. Other Codes should be treated as unrecognized and should result in Code-Rejects.

Timeouts

OSINLCP packets may not be exchanged until PPP has reached the Network-Layer Protocol phase. An implementation should be prepared to wait for Authentication and Link Quality Determination to finish before timing out waiting for a Configure-Ack or other

response. It is suggested that an implementation give up only after user intervention or a configurable amount of time.

Configuration Option Types

OSINLCP has one Configuration Option, which is defined below.

2.1. Sending OSI NPDUs

Before any Network Protocol Data Units (NPDUs) may be communicated, PPP must reach the Network-Layer Protocol phase, and the OSI Network Layer Control Protocol must reach the Opened state.

Exactly one OSI NPDU is encapsulated in the Information field of a PPP Data Link Layer frame where the Protocol field indicates type hex 0023 (OSI Network Layer).

The maximum length of an OSI NPDU transmitted over a PPP link is the same as the maximum length of the Information field of a PPP data link layer frame. Larger NPDUs must be segmented as necessary. If a system wishes to avoid segmentation and reassembly, it should use transport layer mechanisms to discourage others from sending large PDUs.

2.2. NPDU Alignment

OSI protocols have peculiar alignment problems due to the fact that they are often encapsulated in data link protocols with odd-length headers, while PPP defaults to even-length headers. A router switching an OSI packet may find that the beginning of the packet falls on an inconvenient memory boundary when the hardware used to transmit the packet to its next hop requires a particular alignment. This situation can be addressed by the use of leading zero padding.

When sending, an implementation MAY insert one to three octets of zero between the PPP header and the OSI NPDU. These zero octets correspondingly reduce the maximum length of the NPDU that may be transmitted.

On reception, any such leading zero octets (if present) MUST be removed. Regardless of whether leading zero padding is used, an implementation MUST also be able to receive a PPP packet with any arbitrary alignment of the NPDU.

2.3. Network Layer Addressing Information

OSINLCP does not define a separate configuration option for the exchange of OSI Network Layer address information. Instead, the ES-

IS protocol, ISO 9542, should be used. This protocol provides a mechanism for determining the Network Layer address(es) of the neighbor on the link, as well as determining if the neighbor is an End System or an Intermediate System.

A draft addendum to ES-IS [9] is being defined in ISO to add support for dynamic address assignment. This addendum has currently passed the formal "Committee Draft" (CD) letter ballot.

3. OSINLCP Configuration Options

OSINLCP Configuration Options allow negotiation of desirable Internet Protocol parameters. OSINLCP uses the same Configuration Option format defined for LCP [1], with a separate set of Options.

The most up-to-date values of the OSINLCP Option Type field are specified in the most recent "Assigned Numbers" RFC [2]. Current values are assigned as follows:

1 Align-NPDU

3.1. Align-NPDU

Description

This Configuration Option provides a way for the receiver to negotiate a particular alignment of the OSI NPDU. Empirical evidence suggests that the greatest time deficit for re-alignment exists at the receiver.

The alignment is accomplished through combination of PPP header compression with leading zero padding (see above). It is recommended that alignment be entirely through header compression combinations whenever possible. For example, an alignment of 3 could be achieved by combining uncompressed PPP Address and Control fields (2 octets) with a compressed PPP Protocol field (1 octet).

This option is negotiated separately in each direction. A receiver which does not need alignment MUST NOT request the option. A sender which desires alignment prior to sending SHOULD Configure-Nak with an appropriate value.


Implementation Note: In a complex environment, there might be several conflicting needs for alignment. It is recommended that the receiver request alignment based on the needs of the highest speed next hop link. Also, greater efficiency might be obtained by negotiating upstream the values requested by

downstream PPP links, since those packets will not need a change in alignment on transit.

The alignment request is advisory, and failure to agree on an alignment MUST NOT prevent the OSINLCP from reaching the Opened state. By default, the alignment is done according to the needs of the sender, and all receivers MUST be capable of accepting packets with any alignment.

Vernacular: If you don't like this option, you can refuse to negotiate it, and you can send whatever alignment you want. However, if you accept the peer's alignment option, then you MUST transmit packets with the agreed alignment.

A summary of the Align-NPDU Configuration Option format is shown below. The fields are transmitted from left to right.



Type

1

Length

3

Alignment

This field specifies the offset of the beginning of the OSI NPDU relative to the beginning of the PPP packet header (not including any leading Flag Sequences).

A value of 1 through 4 requires an offset of that specific length, modulo 4. For example, a value of 1 would require no padding when the PPP Address, Control, and Protocol fields are compressed. One octet of leading zero padding would be necessary when the PPP header is full sized.

A value of 255 requests an offset of an odd length (1 or 3). A value of 254 requests an offset of an even length (2 or 4). If the sender is not capable of dynamically varying the amount of padding, it MUST NAK with one of the two specific values.

References

- [1] Simpson, W., "The Point-to-Point Protocol (PPP)", RFC 1331, Daydreamer, May 1992.
- [2] Reynolds, J., and J. Postel, "Assigned Numbers", STD 2, RFC 1340, USC/Information Sciences Institute, July 1992.
- [3] ISO, "Information processing systems -- Data communications -- Protocol for providing the connectionless-mode network service", ISO 8473, 1988.
- [4] ISO, "Information processing systems -- Telecommunications and information exchange between systems -- End system to Intermediate system Routeing exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO 9542, 1988.
- [5] ISO, "Information processing systems -- Telecommunications and information exchange between systems -- Intermediate system to Intermediate system Intra-Domain routeing exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO 10589, 1990.
- [6] ISO, "Protocol for Exchange of Inter-domain Routeing Information among Intermediate Systems to Support Forwarding of ISO 8473 PDUs", ISO CD 10747, 1991.
- [7] ISO, "Information technology -- Telecommunications and information exchange between systems -- Protocol identification in the network layer", ISO/IEC TR9577:1990.
- [8] ISO, "Information processing systems -- Data communications -- X.25 packet level protocol for Data terminal equipment", ISO 8208, 1984.
- [9] Taylor, E., "Addendum to ISO 9542 (PDAM 1 - Dynamic Discovery of OSI NSAP Addresses by End Systems)", SC6/N7248.

Acknowledgments

Some of the text in this document is taken from previous documents produced by the Point-to-Point Protocol Working Group of the Internet Engineering Task Force (IETF).

Special thanks to Ross Callon (DEC), and Cyndi Jung (3Com), for contributions of text and design suggestions based on implementation

experience.

Thanks also to Bill Simpson for his editing and formatting efforts, both for this document and for PPP in general.

Security Considerations

Security issues are not discussed in this memo.

Chair's Address

The working group can be contacted via the current chair:

Brian Lloyd
Lloyd & Associates
3420 Sudbury Road
Cameron Park, California 95682

Phone: (916) 676-1147
EMail: brian@lloyd.com

Author's Address

Questions about this memo can also be directed to:

Dave Katz
cisco Systems, Inc.
1525 O'Brien Dr.
Menlo Park, CA 94025

Phone: (415) 688-8284
EMail: dkatz@cisco.com