

Internet Draft
Communications
Corporation

Kory Hamzeh
Ascend
Gurdeep Singh Pall
Microsoft
William Verthein
U.S. Robotics
Jeff Taarud
3COM Corporation
W. Andrew Little
ECI Telematics
June 1996
Expire in six months

Point-to-Point Tunneling Protocol--PPTP
draft-ietf-ppext-pptp-00.txt

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

To learn the current status of any Internet-Draft, please check the "ltd-abstracts.txt" listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Abstract

This document specifies a protocol which allows the Point to Point Protocol (PPP) to be tunneled through an IP network. PPTP does not specify any changes to the PPP protocol but rather describes a new vehicle for carrying PPP. A client-server architecture is defined in order to decouple functions which exist in current Network Access Servers (NAS) and support Virtual Private Networks (VPNs). The PPTP Network Server (PNS) is envisioned to run on a general purpose operating system while the client, referred to as a PPTP Access Concentrator (PAC) operates on a dial access platform. PPTP specifies a call-control and

management protocol which allows the server to control access for dial-in circuit switched calls originating from a PSTN or ISDN or to initiate outbound circuit-switched connections. PPTP uses a GRE-like (Generic Routing Encapsulation) mechanism to provide a flow- and congestion-controlled encapsulated datagram service for carrying PPP packets.

1. Introduction

PPTP allows existing Network Access Server (NAS) functions to be separated using a client-server architecture. Traditionally, the following functions are implemented by a NAS:

- 1) Physical native interfacing to PSTN or ISDN and control of external modems or terminal adapters.

A NAS may interface directly to a telco analog or digital circuit or attach via an external modem or terminal adapter. Control of a circuit-switched connection is accomplished with either modem control or DSS1 ISDN call control protocols.

The NAS, in conjunction with the modem or terminal adapters, may perform rate adaption, analog to digital conversion, sync to async conversion or a number of other alterations of data streams.

- 2) Logical termination of a Point-to-Point-Protocol (PPP) Link Control Protocol (LCP) session.
- 3) Participation in PPP authentication protocols [3].
- 4) Channel aggregation and bundle management for PPP Multilink Protocol.
- 5) Logical termination of various PPP network control protocols (NCP).
- 6) Multiprotocol routing and bridging between NAS interfaces.

PPTP divides these functions between the PAC and PNS. The PAC is responsible for functions 1, 2, and possibly 3. The PNS may be responsible for function 3 and is responsible for functions 4, 5, and 6. The protocol used to carry PPP protocol data units (PDUs) between the PAC and PNS, as well as call control and management is addressed by PPTP.

The decoupling of NAS functions offers these benefits:

Flexible IP address management. Dial-in users may maintain a single IP address as they dial into different PACs as long as they are served from a common PNS. If an enterprise network uses unregistered addresses, a PNS associated with the enterprise assigns addresses meaningful to the private network.

Support of non-IP protocols for dial networks behind IP networks. This allows Appletalk and IPX, for example to be tunneled through an IP-only provider. The PAC need not be capable of processing these protocols.

A solution to the "multilink hunt-group splitting" problem. Multilink PPP, typically used to aggregate ISDN B channels, requires that all of the channels composing a multilink bundle be grouped at a single NAS. Since a multilink PPP bundle can be handled by a single PNS, the channels comprising the bundle may be spread across multiple PACs.

1.1 Protocol Goals and Assumptions

The PPTP protocol is implemented only by the PAC and PNS. No other systems need to be aware of PPTP. Dial networks may be connected to a PAC without being aware of PPTP. Standard PPP client software should continue to operate on tunneled PPP links.

It is envisioned that there will be a many-to-many relationship between PACs and PNSs. A PAC may provide service to many PNSs. For example, an Internet service provider may choose to support PPTP for a number of private network clients and create VPNs for them. Each private network may operate one or more PNSs. A single PNS may associate with many PACs to concentrate traffic from a large number of geographically diverse sites.

While PPTP is not envisioned to be implemented on dial user software, this implementation is not precluded by the protocol.

PPTP uses a GRE-like discipline to carry user PPP packets. These enhancements allow for low-level congestion and flow control to be provided on the tunnels used to carry user data between PAC and PNS. This mechanism allows for efficient use of the bandwidth available for the tunnels and avoids unnecessary retransmissions and buffer overruns. PPTP does not dictate the particular algorithms to be used for this low level control but it does define the parameters that must be communicated in order to allow such algorithms to work. Suggested algorithms are included in Appendix A.

1.2 Terminology

Analog Channel

A circuit-switched communication path which is intended to carry 3.1 Khz audio in each direction.

Digital Channel

A circuit-switched communication path which is intended to carry digital information in each direction.

Call

A connection or attempted connection between two terminal endpoints on a PSTN or ISDN--for example, a telephone call between two modems.

Control Connection

A control connection is created for each PAC, PNS pair and operates over TCP [4]. The control connection governs aspects of the tunnel and of sessions assigned to the tunnel.

Dial User

An end-system or router attached to an on-demand PSTN or ISDN which is either the initiator or recipient of a call.

Network Access Server (NAS)

A device providing temporary, on-demand network access to users. This access is point-to-point using PSTN or ISDN lines.

PPTP Access Concentrator (PAC)

A device attached to one or more PSTN or ISDN lines capable of PPP operation and of handling the PPTP protocol. The PAC need only implement TCP/IP to pass traffic to one or more PNSs. It may also tunnel non-IP protocols.

PPTP Network Server (PNS)

A PNS is envisioned to operate on general-purpose computing/server platforms. The PNS handles the server side of the PPTP protocol. Since PPTP relies completely on TCP/IP and is independent of the interface hardware, the PNS may use any combination of IP interface hardware including LAN and WAN devices.

Session

PPTP is connection-oriented. The PNS and PAC maintain state for each user that is attached to a PAC. A session is created when end-to-end PPP connection is attempted between a dial user and the PNS. The datagrams related to a session are sent over the tunnel between the PAC and PNS.

Tunnel

A tunnel is defined by a PNS-PAC pair. The tunnel protocol is defined by a modified version of GRE [1,2]. The tunnel carries PPP datagrams between the PAC and the PNS. Many sessions are multiplexed on a single tunnel. A control connection operating over TCP controls the establishment, release, and maintenance of sessions and of the tunnel itself.

1.3 Protocol Overview

There are two parallel components of PPTP: 1) a Control Connection between each PAC-PNS pair operating over TCP and 2) an IP tunnel operating between the same PAC-PNS pair which is used to transport GRE encapsulated PPP packets for user sessions between the pair.

1.3.1 Control Connection Overview

Before PPP tunneling can occur between a PAC and PNS, a control connection must be established between them. The control connection is a standard TCP session over which PPTP call control and management information is passed. The control session is logically associated with, but separate from, the sessions being tunneled through a PPTP tunnel. For each PAC-PNS pair both a tunnel and a control connection exist. The control connection is responsible for establishment, management, and release of sessions carried through the tunnel. It is the means by which a PNS is notified of an incoming call at an associated PAC, as well as the means by which a PAC is instructed to place an outgoing dial call.

A control connection can be established by either the PNS or the PAC. Following the establishment of the required TCP connection, the PNS and PAC establish the control connection using the Start-Control-Connection-Request and -Reply messages. These messages are also used to exchange information about basic operating capabilities of the PAC and PNS. Once the control connection is established, the PAC or PNS may initiate sessions by requesting outbound calls or responding to inbound requests. The control connection may communicate changes in operating characteristics of an individual user session with a Set-Link-Info message. Individual sessions may be released by either the

PAC or PNS, also through Control Connection messages.

The control connection itself is maintained by keep-alive echo messages. This ensures that a connectivity failure between the PNS and the PAC can be detected in a timely manner. Other failures can be reported via the Wan-Error-Notify message, also on the control connection.

It is intended that the control connection will also carry management related messages in the future, such as a message allowing the PNS to request the status of a given PAC; these message types have not yet been defined.

1.3.2 Tunnel Protocol Overview

PPTP requires the establishment of a tunnel for each communicating PNS-PAC pair. This tunnel is used to carry all user session PPP packets for sessions involving a given PNS-PAC pair. A key which is present in the GRE header indicates which session a particular PPP packet belongs to. In this manner, PPP packets are multiplexed and demultiplexed over a single tunnel between a given PNS-PAC pair. The value to use in the key field is established by the call establishment procedure which takes place on the control connection.

The GRE header also contains acknowledgment and sequencing information that is used to perform some level of congestion-control and error detection over the tunnel. Again the control connection is used to determine rate and buffering parameters that are used to regulate the flow of PPP packets for a particular session over the tunnel. PPTP does not specify the particular algorithms to use for congestion-control and flow-control. Suggested algorithms for the determination of adaptive time-outs to recover from dropped data or acknowledgments on the tunnel are included in Appendix A of this document.

1.4 Specification Language

In this document, several words are used to signify the requirements of the specification. These words are often capitalized.

MUST	This word, or the adjective "required", means that the definition is an absolute requirement of the specification.
MUST NOT	This phrase means that the definition is an absolute prohibition of the specification.
SHOULD	This word, or the adjective "recommended",

means that, in some circumstances, valid reasons may exist to ignore this item, but the full implications must be understood and carefully weighed before choosing a different course. Unexpected results may result otherwise.

MAY This word, or the adjective "optional", means that this item is one of an allowed set of alternatives. An implementation which does not include this option MUST be prepared to interoperate with another implementation which does include the option.

silently discard The implementation discards the datagram without further processing, and without indicating an error to the sender. The implementation SHOULD provide the capability of logging the error, including the contents of the discarded datagram, and SHOULD record the event in a statistics counter.

1.5 Message Format and Protocol Extensibility

PPTP defines a set of messages sent as TCP data on the control connection between a PNS and a given PAC. The TCP session for the control connection is established by initiating a TCP connection to port 5678. The source port is assigned to any unused port number.

Each PPTP Control Connection message begins with an 8 octet fixed header portion. This fixed header contains the following: the total length of the message, the PPTP Message Type indicator, and a "Magic Cookie".

Two Control Connection message types are indicated by the PPTP Message Type field:

- 1 - Control Message
- 2 - Management Message

Management messages are currently not defined.

The Magic Cookie is always sent as the constant 0x1A2B3C4D. Its basic purpose is to allow the receiver to ensure that it is properly synchronized with the TCP data stream. It should not be used as a means for resynchronizing the TCP data stream in the event that a transmitter issues an improperly formatted message. Loss of

synchronization must result in immediate closing of the control connection's TCP session.

For clarity, all Control Connection message templates in the next section include the entire PPTP Control Connection message header. Numbers preceded by 0x are hexadecimal values.

The currently defined Control Messages, grouped by function, are:

Control Message	Message Code
(Control Connection Management)	
Start-Control-Connection-Request	1
Start-Control-Connection-Reply	2
Stop-Control-Connection-Request	3
Stop-Control-Connection-Reply	4
Echo-Request	5
Echo-Reply	6
(Call Management)	
Outgoing-Call-Request	7
Outgoing-Call-Reply	8
Incoming-Call-Request	9
Incoming-Call-Reply	10
Incoming-Call-Connected	11
Call-Clear-Request	12
Call-Disconnect-Notify	13
(Error Reporting)	
WAN-Error-Notify	14
(PPP Session Control)	
Set-Link-Info	15

The Start-Control-Connection-Request and -Reply messages determine which version of the Control Connection protocol will be used. The version number field carried in these messages consists of a version number in the high octet and a revision number in the low octet. Version handling is described in Section 3. The current value of the version number field is 0x0100 for version 1, revision 0.

The use of the GRE-like header for the encapsulation of PPP user packets is specified in Section 4.

The MTU for the user data packets encapsulated in GRE is 1532 octets, not including the IP and GRE headers.

2.0 Control Connection Protocol Specification

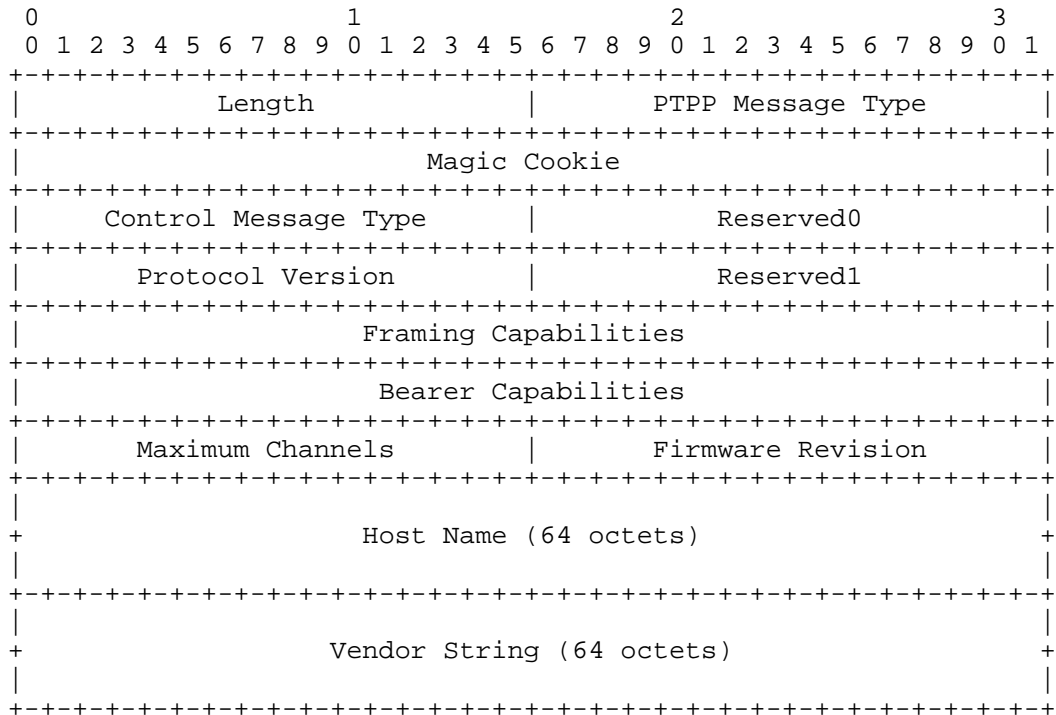
Control Connection messages are used to establish and clear user sessions. The first set of Control Connection messages are used to maintain the control connection itself. The control connection is initiated by either the PNS or PAC after they establish the underlying TCP connection. The procedure and configuration information required to determine which TCP connections are established is not covered by this protocol.

The following Control Connection messages are all sent as user data on the established TCP connection between a given PNS-PAC pair. Note that care has been taken to ensure that all word (2 octet) and longword (4 octet) values begin on appropriate boundaries. All data is sent in network order (high order octets first). Any "reserved" fields MUST be sent as 0 values to allow for protocol extensibility.

The TCP header is followed by the PPTP fields shown in the following:

2.1 Start-Control-Connection-Request

The Start-Control-Connection-Request is a PPTP control message used to establish the control connection between a PNS and a PAC. Each PNS-PAC pair requires a dedicated control connection to be established. A control connection must be established before any other PPTP messages can be issued. The establishment of the control connection can be initiated by either the PNS or PAC. A procedure which handles the occurrence of a collision between PNS and PAC Start-Control-Connection-Requests is described in Section 3.



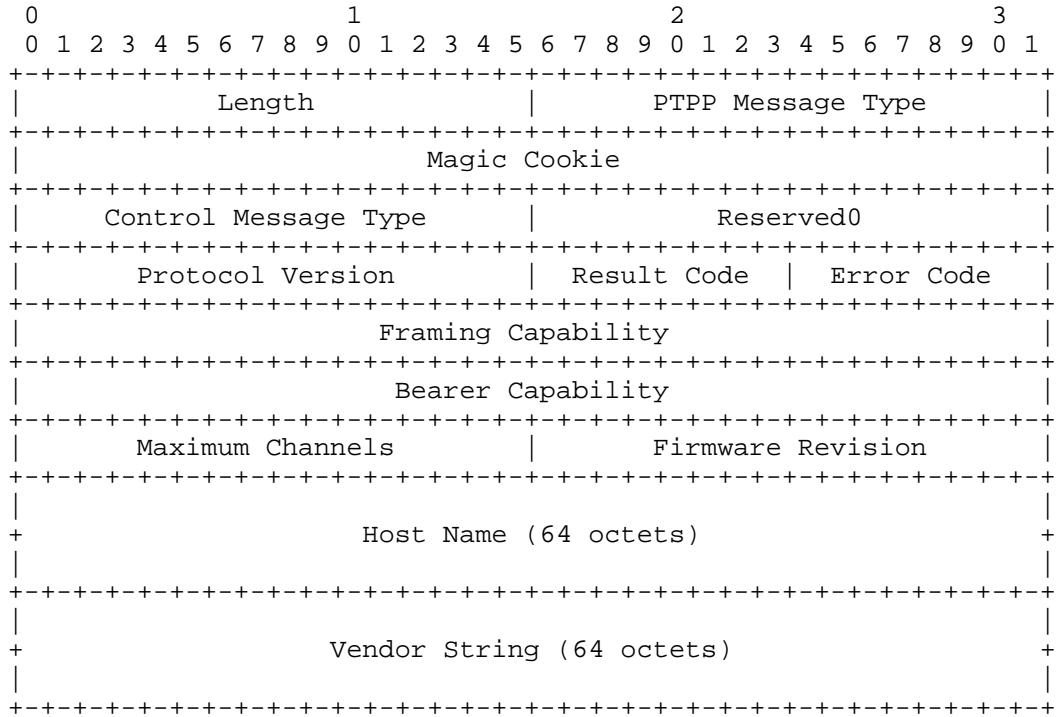
Length	Total length in octets of this PPTP message, including the entire PPTP header.
PTPP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D. This constant value is used as a sanity check on received messages (see Section 1.5).

Control Message Type	1 for Start-Control-Connection-Request.
Reserved0	This field MUST be 0.
Protocol Version	The version of the PPTP protocol that the sender wishes to use.
Reserved1	This field MUST be 0.
Framing Capabilities	A set of bits indicating the type of framing that the sender of this message can provide. The currently defined bit settings are: 1 - Asynchronous Framing supported 2 - Synchronous Framing supported
Bearer Capabilities	A set of bits indicating the bearer capabilities that the sender of this message can provide. The currently defined bit settings are: 1 - Analog access supported 2 - Digital access supported
Maximum Channels	The total number of individual PPP sessions this PAC can support. In Start-Control-Connection-Requests issued by the PNS, this value SHOULD be set to 0. It MUST be ignored by the PAC.
Firmware Revision	This field contains the firmware revision number of the issuing PAC, when issued by the PAC, or the version of the PNS PPTP driver if issued by the PNS.
Host Name	A 64 octet field containing the DNS name of the issuing PAC or PNS. If less than 64 octets in length, the remainder of this field SHOULD be filled with octets of value 0.
Vendor Name	A 64 octet field containing a vendor specific string describing the type of PAC being used, or the type of PNS software being used if this request is

issued by the PNS. If less than 64 octets in length, the remainder of this field SHOULD be filled with octets of value 0.

2.2 Start-Control-Connection-Reply

The Start-Control-Connection-Reply is a PPTP control message sent in reply to a received Start-Control-Connection-Request message. This message contains a result code indicating the result of the control connection establishment attempt.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PTPP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	2 for Start-Control-Connection-Reply.
Reserved0	This field MUST be 0.
Protocol Version	The version of the PPTP protocol that the

sender wishes to use.

Result Code

Indicates the result of the command channel establishment attempt. Current valid Result Code values are:

- 1 - Successful channel establishment
- 2 - General error -- Error Code indicates the problem
- 3 - Command channel already exists;
- 4 - Requester is not authorized to establish a command channel
- 5 - The protocol version of the requester is not supported

Error Code

This field is set to 0 unless a "General Error" exists, in which case Result Code is set to 2 and this field is set to the value corresponding to the general error condition as specified in Section 2.16.

Framing Capabilities

A set of bits indicating the type of framing that the sender of this message can provide. The currently defined bit settings are:

- 1 - Asynchronous Framing supported
- 2 - Synchronous Framing supported.

Bearer Capabilities

A set of bits indicating the bearer capabilities that the sender of this message can provide. The currently defined bit settings are:

- 1 - Analog access supported
- 2 - Digital access supported

Maximum Channels

The total number of individual PPP sessions this PAC can support. In a Start-Control-Connection-Reply issued by the PNS, this value SHOULD be set to 0 and it must be ignored by the PAC. The

PNS MUST NOT use this value to try to track the remaining number of PPP sessions that the PAC will allow.

Firmware Revision

This field contains the firmware revision number of the issuing PAC, or the version of the PNS PPTP driver if issued by the PNS.

Host Name


A 64 octet field containing the DNS name of the issuing PAC or PNS. If less than 64 octets in length, the remainder of this field SHOULD be filled with octets of value 0.

Vendor Name

A 64 octet field containing a vendor specific string describing the type of PAC being used, or the type of PNS software being used if this request is issued by the PNS. If less than 64 octets in length, the remainder of this field SHOULD be filled with octets of value 0.

2.3 Stop-Control-Connection-Request

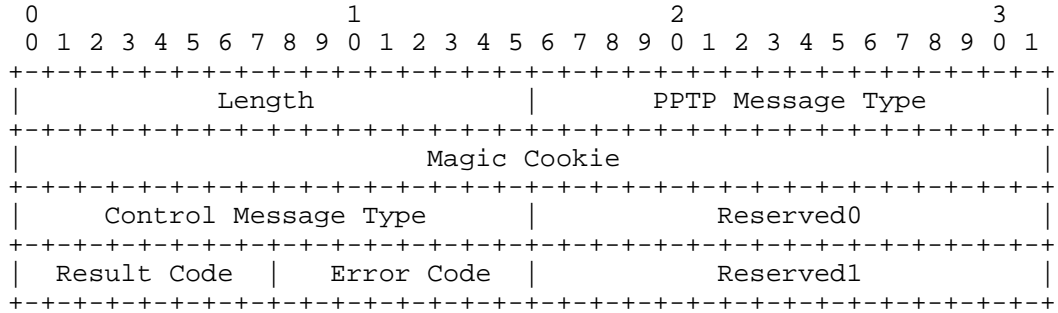
The Stop-Control-Connection-Request is a PPTP control message sent by one peer of a PAC-PNS control connection to inform the other peer that the control connection should be closed. In addition to closing the control connection, all active user calls are implicitly cleared. The reason for issuing this request is indicated in the Reason field.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PTPP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	3 for Stop-Control-Connection-Request.
Reserved0	This field MUST be 0.
Reason	Indicates the reason for the control connection being closed. Current valid Reason values are: <ul style="list-style-type: none"> 1 (None) - General request to clear control connection 2 (Stop-Protocol) - Can't support peer's version of the protocol 3 (Stop-Local-Shutdown) - Requester is being shut down
Reserved1, Reserved2	These fields MUST be 0.

2.4 Stop-Control-Connection-Reply

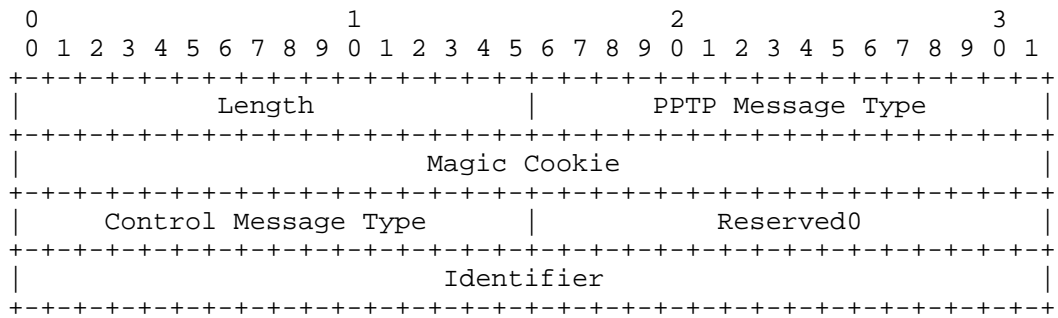
The Stop-Control-Connection-Reply is a PPTP control message sent by one peer of a PAC-PNS control connection upon receipt of a Stop-Control-Connection-Request from the other peer.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PTPP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	4 for Stop-Control-Connection-Reply.
Reserved0	This field MUST be 0.
Result Code	Indicates the result of the attempt to close the control connection. Current valid Result Code values are: <ul style="list-style-type: none"> 1 (OK) - Control connection closed 2 (General Error) - Control connection not closed for reason indicated in Error Code
Error Code	This field is set to 0 unless a "General Error" exists, in which case Result Code is set to 2 and this field is set to the value corresponding to the general error condition as specified in Section 2.16.
Reserved1	This field MUST be 0.

2.5 Echo-Request

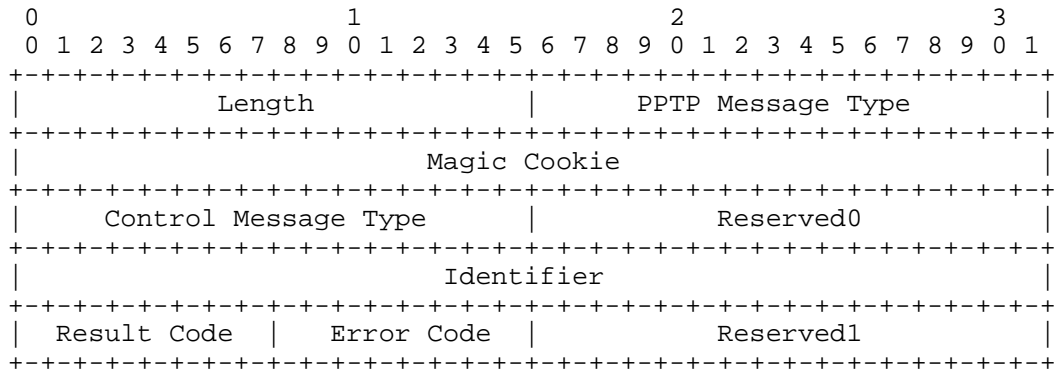
The Echo-Request is a PPTP control message sent by either peer of a PAC-PNS control connection. This control message is used as a "keep-Alive" for the control connection. The receiving peer issues an Echo-Reply to each Echo-Request received. As specified in Section 3, if the sender does not receive an Echo Reply in response to an Echo-Request, it will eventually clear the control connection.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PPTP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	5 for Echo-Request.
Reserved0	This field MUST be 0.
Identifier	A value set by the sender of the Echo-Request that is used to match the reply with the corresponding request.

2.6 Echo-Reply

The Echo-Reply is a POTP control message sent by either peer of a PAC-PNS control connection in response to the receipt of an Echo-Request.



Length	Total length in octets of this POTP message, including the entire POTP header.
POTP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	6 for Echo-Reply.
Reserved0	This field MUST be 0.
Identifier	The contents of the identify field from the received Echo-Request is copied to this field.
Result Code	Indicates the result of the receipt of the Echo-Request. Current valid Result Code values are: <ul style="list-style-type: none"> 1 (OK) - The Echo-Reply is valid 2 (General Error) - Echo-Request not accepted for the reason indicated in Error Code
Error Code	This field is set to 0 unless a "General

Error" condition exists, in which case Result Code is set to 2 and this field is set to the value corresponding to the general error condition as specified in Section 2.16.

Reserved1

This field MUST be 0.

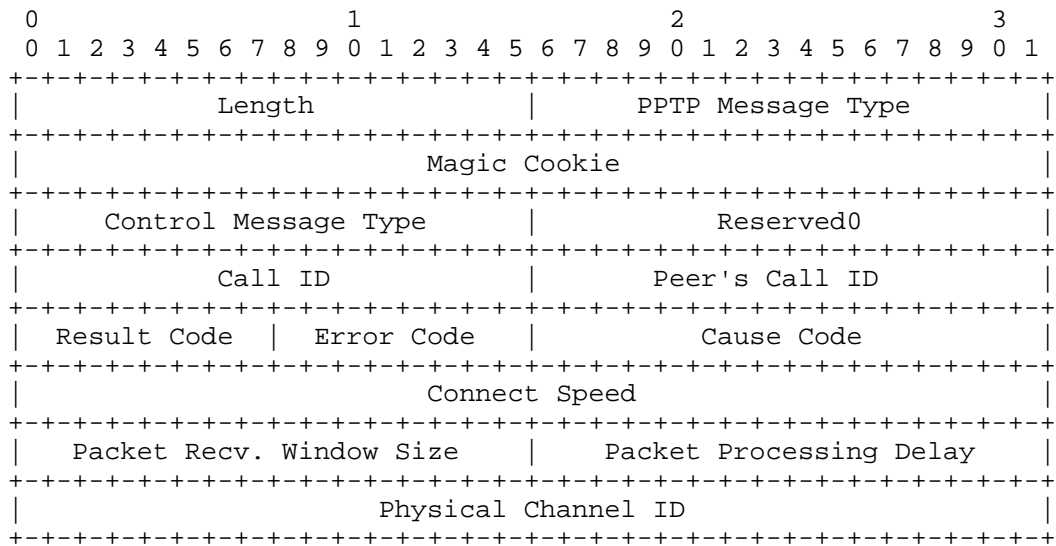
Magic Cookie	0x1A2B3C4D.
Control Message Type	7 for Outgoing-Call-Request.
Reserved0	This field MUST be 0.
Call ID	A unique identifier, unique to a particular PAC-PNS pair assigned by the PNS to this session. It is used to multiplex and demultiplex data sent over the tunnel between the PNS and PAC involved in this session.
Call Serial Number	An identifier assigned by the PNS to this session for the purpose of identifying this particular session in logged session information. Unlike the Call ID, both the PNS and PAC associate the same Call Serial Number with a given session. The combination of IP address and call serial number SHOULD be unique.
Minimum BPS	The lowest acceptable line speed (in bits/second) for this session.
Maximum BPS	The highest acceptable line speed (in bits/second) for this session.
Bearer Type	A value indicating the bearer capability required for this outgoing call. The currently defined values are: 1 - Call to be placed on an analog channel 2 - Call to be placed on a digital channel 3 - Call can be placed on any type of channel
Framing Type	A value indicating the type of PPP framing to be used for this outgoing call. 1 - Call to use Asynchronous framing 2 - Call to use Synchronous framing

3 - Call can use either type of framing

Packet Recv. Window Size	The number of received data packets the PNS will buffer for this session.
Packet Processing Delay	A measure of the packet processing delay that might be imposed on data sent to the PNS from the PAC. This value is specified in units of 1/10 seconds. For the PNS this number should be very small. See appendix A for a description of how this value is determined and used.
Phone Number Length	The actual number of valid digits in the Phone Number field.
Reserved1	This field MUST be 0.
Phone Number	The number to be dialed to establish the outgoing session. For ISDN and analog calls this field is an ASCII string. If the Phone Number is less than 64 octets in length, the remainder of this field is filled with octets of value 0.
Subaddress	A 64 octet field used to specify additional dialing information. If the subaddress is less than 64 octets long, the remainder of this field is filled with octets of value 0.

2.8 Outgoing-Call-Reply

The Outgoing-Call-Reply is a PPTP control message sent by the PAC to the PNS in response to a received Outgoing-Call-Request message. The reply indicates the result of the outgoing call attempt. It also provides information to the PNS about particular parameters used for the call. It provides information to allow the PNS to regulate the transmission of data to the PAC for this session.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PPTP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	8 for Outgoing-Call-Reply.
Reserved0	This field MUST be 0.
Call ID	A unique identifier for the tunnel, assigned by the PAC to this session. It is used to multiplex and demultiplex data sent over the tunnel between the PNS and PAC involved in this session.

Peer's Call ID	This field is set to the value received in the Call ID field of the corresponding Outgoing-Call-Request message. It is used by the PNS to match the Outgoing-Call-Reply with the Outgoing-Call-Request it issued. It also is used as the value sent in the GRE header for mux/demuxing.
Result Code	<p>This value indicates the result of the Outgoing-Call-Request attempt. Currently valid values are:</p> <ol style="list-style-type: none">1 (Connected) - Call established with no errors2 (General Error) - Outgoing Call not established for the reason indicated in Error Code3 (No Carrier) - Outgoing Call failed due to no carrier detected4 (Busy) - Outgoing Call failed due to detection of a busy signal5 (No Dial Tone) - Outgoing Call failed due to lack of a dial tone6 (Time-out) - Outgoing Call was not established within time allotted by PAC7 (Do Not Accept) - Outgoing Call administratively prohibited
Error Code	This field is set to 0 unless a "General Error" condition exists, in which case Result Code is set to 2 and this field is set to the value corresponding to the general error condition as specified in Section 2.16.
Cause Code	This field gives additional failure information. Its value can vary depending upon the type of call attempted. For ISDN call attempts it is the Q.931 cause code.

Connect Speed	The actual connection speed used, in bits/second.
Packet Recv. Window Size	The number of received data packets the PAC will buffer for this session.
Packet Processing Delay	A measure of the packet processing delay that might be imposed on data sent to the PAC from the PNS. This value is specified in units of 1/10 seconds. For the PAC, this number is related to the size of the buffer used to hold packets to be sent to the client and to the speed of the link to the client. This value should be set to the maximum delay that can normally occur between the time a packet arrives at the PAC and is delivered to the client. See Appendix A for an example of how this value is determined and used.
Physical Channel ID	This field is set by the PAC in a vendor-specific manner to the physical channel number used to place this call. It is used for logging purposes only.

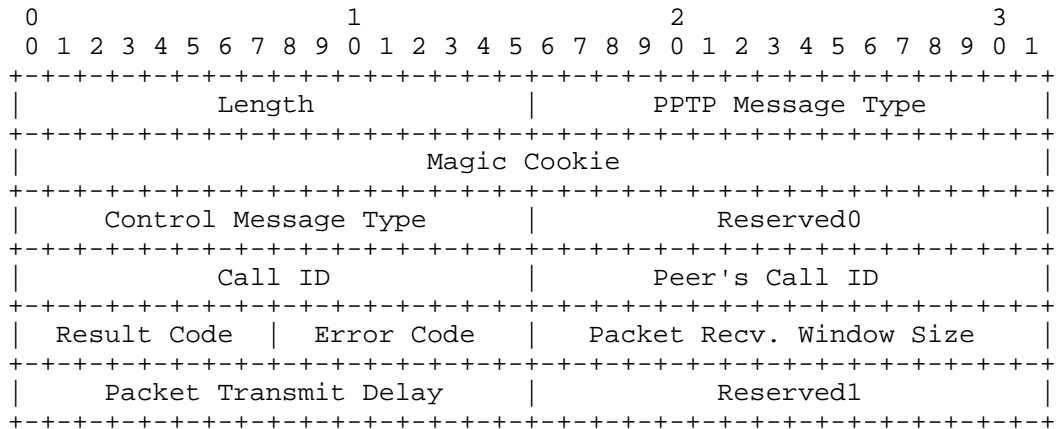
PTPP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	9 for Incoming-Call-Request.
Reserved0	This field MUST be 0.
Call ID	A unique identifier for this tunnel, assigned by the PAC to this session. It is used to multiplex and demultiplex data sent over the tunnel between the PNS and PAC involved in this session.
Call Serial Number	An identifier assigned by the PAC to this session for the purpose of identifying this particular session in logged session information. Unlike the Call ID, both the PNS and PAC associate the same Call Serial Number to a given session. The combination of IP address and call serial number should be unique.
Bearer Type	A value indicating the bearer capability used for this incoming call. Currently defined values are: 1 - Call is on an analog channel 2 - Call is on a digital channel
Physical Channel ID	This field is set by the PAC in a vendor-specific manner to the number of the physical channel this call arrived on.
Dialed Number Length	The actual number of valid digits in the Dialed Number field.
Dialing Number Length	The actual number of valid digits in the Dialing Number field.
Dialed Number	The number that was dialed by the caller. For ISDN and analog calls this field is an ASCII string. If the Dialed Number is less than 64 octets in length, the remainder of this field is filled with octets of value 0.

Dialing Number	The number from which the call was placed. For ISDN and analog calls this field is an ASCII string. If the Dialing Number is less than 64 octets in length, the remainder of this field is filled with octets of value 0.
Subaddress	A 64 octet field used to specify additional dialing information. If the subaddress is less than 64 octets long, the remainder of this field is filled with octets of value 0.

2.10 Incoming-Call-Reply

The Incoming-Call-Reply is a PPTP control message sent by the PNS to the PAC in response to a received Incoming-Call-Request message. The reply indicates the result of the incoming call attempt. It also provides information to allow the PAC to regulate the transmission of data to the PNS for this session.

This message is the second in the three-way handshake used by PPTP for establishing incoming calls. It indicates to the PAC whether the call should be answered or not.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PPTP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	10 for Incoming-Call-Reply.
Reserved0	This field MUST be 0.
Call ID	A unique identifier for this tunnel assigned by the PAC to this session. It is used to multiplex and demultiplex data sent over the tunnel between the PNS and PAC involved in this session.
Peer's Call ID	This field is set to the value received

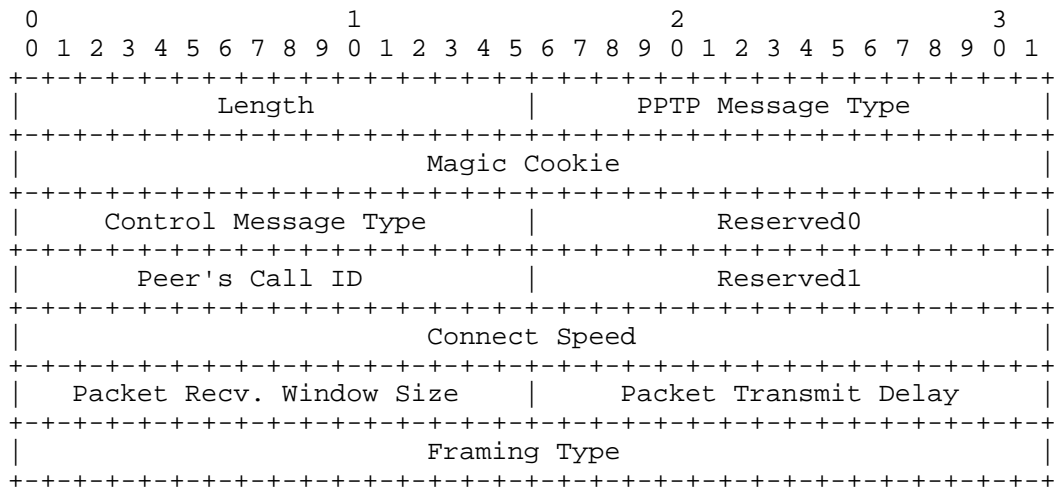
in the Call ID field of the corresponding Incoming-Call-Request message. It is used by the PAC to match the Incoming-Call-Reply with the Incoming-Call-Request it issued. This value is included in the GRE header of transmitted data packets for this session.

Result Code	<p>This value indicates the result of the Incoming-Call-Request attempt. Current valid Result Code values are:</p> <ol style="list-style-type: none">1 (Connect) - The PAC should answer the incoming call2 (General Error) - The Incoming Call should not be established due to the reason indicated in Error Code3 (Do Not Accept) - The PAC should not accept the incoming call. It should hang up or issue a busy indication
Error Code	<p>This field is set to 0 unless a "General Error" condition exists, in which case Result Code is set to 2 and this field is set to the value corresponding to the general error condition as specified in Section 2.16.</p>
Packet Recv. Window Size	<p>The number of received data packets the PAC will buffer for this session.</p>
Packet Transmit Delay	<p>A measure of the packet processing delay that might be imposed on data sent to the PAC from the PNS. This value is specified in units of 1/10 seconds. See Appendix A for a description of how this value is determined and used.</p>
Reserved1	<p>This field MUST be 0.</p>

2.11 Incoming-Call-Connected

The Incoming-Call-Connected message is a PPTP control message sent by the PAC to the PNS in response to a received Incoming-Call-Reply. It provides information to the PNS about particular parameters used for the call. It also provides information to allow the PNS to regulate the transmission of data to the PAC for this session.

This message is the third in the three-way handshake used by PPTP for establishing incoming calls. It provides a mechanism for providing the PNS with additional information about the call that cannot, in general, be obtained at the time the Incoming-Call-Request is issued by the PAC.




Length	Total length in octets of this PPTP message, including the entire PPTP header.
PPTP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	11 for Incoming-Call-Connected.
Reserved0	This field MUST be 0.
Peer's Call ID	This field is set to the value received in the Call ID field of the corresponding Incoming-Call-Reply message. It is used

by the PNS to match the Incoming-Call-Connected with the Incoming-Call-Reply it issued.

Connect Speed	The actual connection speed used, in bits/second.
Packet Recv. Window Size	The number of received data packets the PAC will buffer for this session.
Packet Transmit Delay	A measure of the packet processing delay that might be imposed on data sent to the PAC from the PNS. This value is specified in units of 1/10 seconds. See Appendix A for a description of how this value is determined and used.
Framing Type	A value indicating the type of PPP framing being used by this incoming call. 1 - Call uses asynchronous framing 2 - Call uses synchronous framing

2.12 Call-Clear-Request

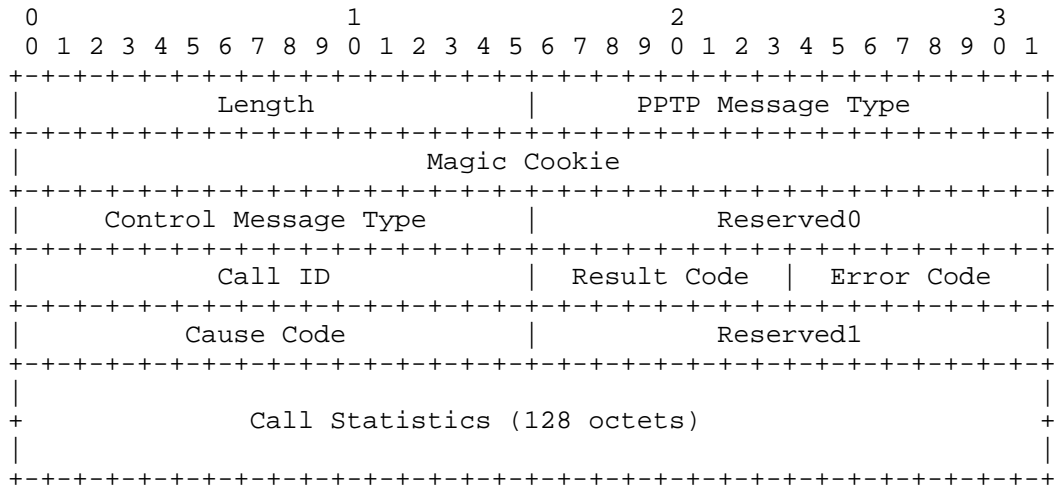
The Call-Clear-Request is a PPTP control message sent by the PNS to the PAC indicating that a particular call is to be disconnected. The call being cleared can be either an incoming or outgoing call, in any state. The PAC responds to this message with a Call-Disconnect-Notify message.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PTPP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	12 for Call-Clear-Request.
Reserved0	This field MUST be 0.
Call ID	The Call ID assigned by the PNS to this call. This value is used instead of the Peer's Call ID because the latter may not be known to the PNS if the call must be aborted during call establishment.
Reserved1	This field MUST be 0.

2.13 Call-Disconnect-Notify

The Call-Disconnect-Notify message is a PPTP control message sent by the PAC to the PNS. It is issued whenever a call is disconnected, due to the receipt by the PAC of a Call-Clear-Request or for any other reason. Its purpose is to inform the PNS of both the disconnection and the reason for it.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PTTP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	13 for Call-Disconnect-Notify.
Reserved0	This field MUST be 0.
Call ID	The value of the Call ID assigned by the PAC to this call. This value is used instead of the Peer's Call ID because the latter may not be known to the PNS if the call must be aborted during call establishment.
Result Code	This value indicates the reason for the disconnect. Current valid Result Code

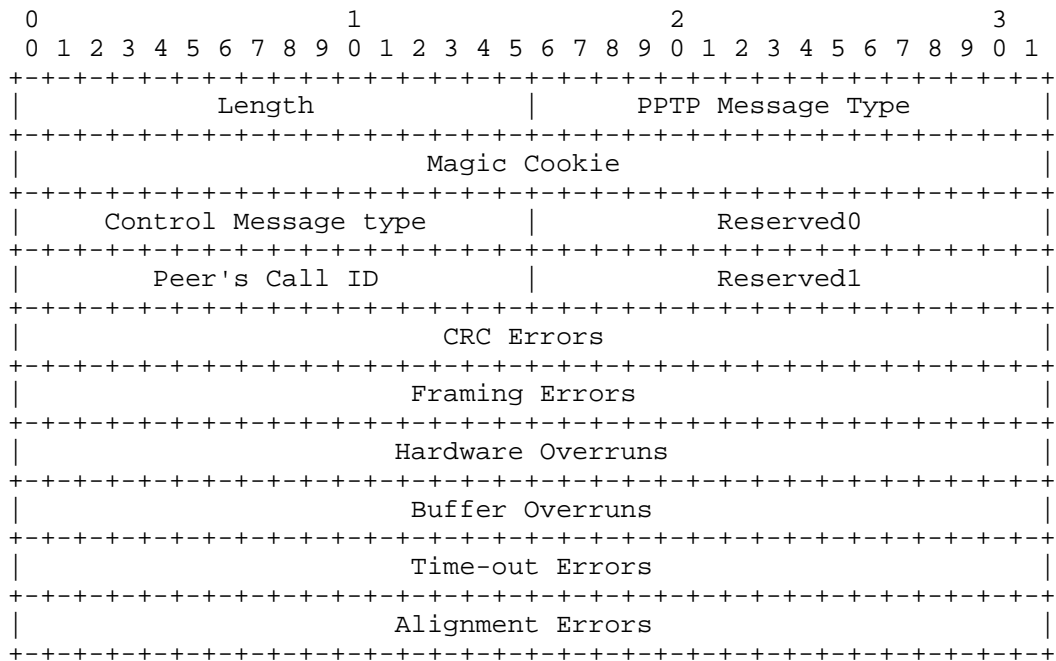
values are:

- 1 (Lost Carrier) - Call disconnected due to loss of carrier
- 2 (General Error) - Call disconnected for the reason indicated in Error Code
- 3 (Admin Shutdown) - Call disconnected for administrative reasons
- 4 (Request) - Call disconnected due to received Call-Clear-Request

Error Code	This field is set to 0 unless a "General Error" condition exists, in which case the Result Code is set to 2 and this field is set to the value corresponding to the general error condition as specified in Section 2.16.
Cause Code	This field gives additional disconnect information. Its value varies depending on the type of call being disconnected. For ISDN calls it is the Q.931 cause code.
Call Statistics	This field is an ASCII string containing vendor-specific call statistics that can be logged for diagnostic purposes. If the length of the string is less than 128, the remainder of the field is filled with octets of value 0.

2.14 WAN-Error-Notify

The WAN-Error-Notify message is a PPTP control message sent by the PAC to the PNS to indicate WAN error conditions (conditions that occur on the interface supporting PPP). The counters in this message are cumulative. This message should only be sent when an error occurs, and not more than once every 60 seconds. The counters are reset when a new call is established.

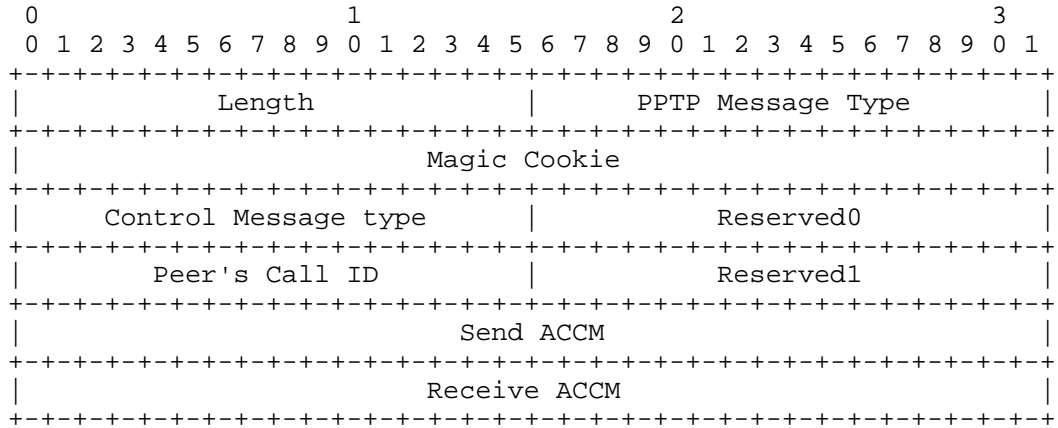


Length	Total length in octets of this PPTP message, including the entire PPTP header.
PTPP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	14 for WAN-Error-Notify.
Reserved0	This field MUST be 0.
Peer's Call ID	Th Call ID assigned by the PNS to this call.

CRC Errors	Number of PPP frames received with CRC errors since session was established.
Framing Errors	Number of improperly framed PPP packets received.
Hardware Overruns	Number of receive buffer over-runs since session was established.
Buffer Overruns	Number of buffer over-runs detected since session was established.
Time-out Errors	Number of time-outs since call was established.
Alignment Errors	Number of alignment errors since call was established.

2.15 Set-Link-Info

The Set-Link-Info message is a PPTP control message sent by the PNS to the PAC to set PPP-negotiated options. Because these options can change at any time during the life of the call, the PAC must be able to update its internal call information dynamically and perform PPP negotiation on an active PPP session.



Length	Total length in octets of this PPTP message, including the entire PPTP header.
PPTP Message Type	1 for Control Message.
Magic Cookie	0x1A2B3C4D.
Control Message Type	15 for Set-Link-Info.
Reserved0	This field MUST be 0.
Peer's Call ID	The value of the Call ID assigned by the PAC to this call.
Reserved1	This field MUST be 0.
Send ACCM	The send ACCM value the client should use to process outgoing PPP packets. The default value used by the client until this message is received is 0xFFFFFFFF. See [7].

Receive ACCM

The receive ACCM value the client should use to process incoming PPP packets. The default value used by the client until this message is received is 0xFFFFFFFF. See [7].

2.16 General Error Codes

General error codes pertain to types of errors which are not specific to any particular PPTP request, but rather to protocol or message format errors. If a PPTP reply indicates in its Result Code that a general error occurred, the General Error value should be examined to determine what the error was. The currently defined General Error codes and their meanings are:

- 0 (None) - No general error
- 1 (Not-Connected) - No control connection exists yet for this PAC-PNS pair
- 2 (Bad-Format) - Length is wrong or Magic Cookie value is incorrect
- 3 (Bad-Value) - One of the field values was out of range or reserved field was non-zero
- 4 (No-Resource) - Insufficient resources to handle this command now
- 5 (Bad-Call ID) - The Call ID is invalid in this context
- 6 (PAC-Error) - A generic vendor-specific error occurred in the PAC

3.0 Control Connection Protocol Operation

This section describes the operation of various PPTP control connection functions and the Control Connection messages which are used to support them. The protocol operation of the control connection is simplified because TCP is used to provide a reliable transport mechanism.

Ordering and retransmission of messages is not a concern at this level. The TCP connection itself, however, can close at any time and an appropriate error recovery mechanism must be provided to handle this case.

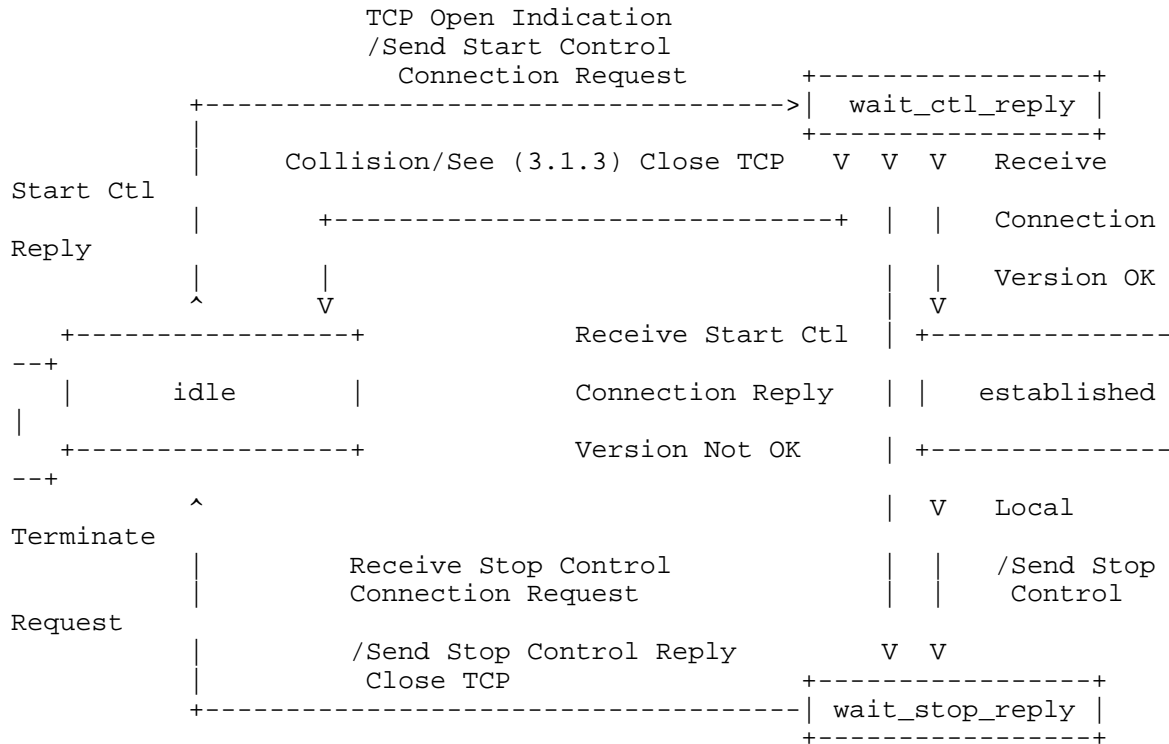
Some error recovery procedures are common to all states of the control connection. If an expected reply does not arrive within 60 seconds, the control connection is closed, unless otherwise specified. Appropriate logging should be implemented for easy determination of the problems and the reasons for closing the control connection.

Receipt of an invalid or malformed Control Connection message should be logged appropriately, and the control connection should be closed and restarted to ensure recovery into a known state.

3.1 Control Connection States

The control connection relies on a standard TCP connection for its service. The PPTP control connection protocol is not distinguishable between the PNS and PAC, but is distinguishable between the originator and receiver. The originating peer is the one which first attempts the TCP open. Since either PAC or PNS may originate a connection, it is possible for a TCP collision to occur. See Section 3.1.3 for a description of this situation.

3.1.1 Control Connection Originator (may be PAC or PNS)



idle

The control connection originator attempts to open a TCP connection to the peer during idle state. When the TCP connection is open, the originator transmits a send Start-Control-Connection-Request and then enters the wait_ctl_reply state.

wait_ctl_reply

The originator checks to see if another TCP connection has been requested from the same peer, and if so, handles the collision situation described in Section 3.1.3.

When a Start-Control-Connection-Reply is received, it is examined for a compatible version. If the version of the reply is lower than the version sent in the request, the older (lower) version should be used provided it is supported. If the version in the reply is earlier and supported, the originator moves to the established state. If the version is earlier and not supported, a Stop-Control-Connection-Request SHOULD be sent to the peer and the originator moves into the wait_stop_reply state.

established

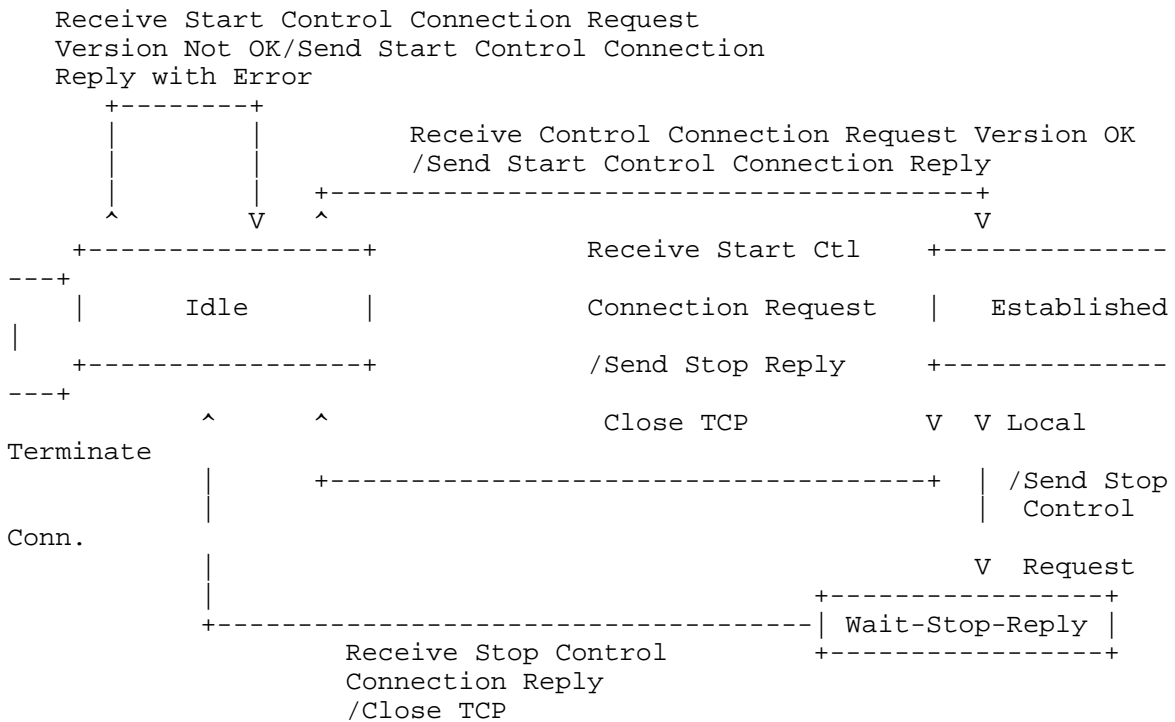
An established connection may be terminated by either a local condition or the receipt of a Stop-Control-Connection-Request. In the event of a local termination, the originator MUST send a Stop-Control-Connection-Request and enter the wait_stop_reply state.

If the originator receives a Stop-Control-Connection-Request it SHOULD send a Stop-Control-Connection-Reply and close the TCP connection making sure that the final TCP information has been "pushed" properly.

wait_stop_reply

If a Stop-Control-Connection-Reply is received, the TCP connection SHOULD be closed and the control connection becomes idle.

3.1.2 Control connection Receiver (may be PAC or PNS)



idle

The control connection receiver waits for a TCP open attempt on port 5678. When notified of an open TCP connection, it should prepare to receive PPTP messages. When a Start-Control-Connection-Request is received its version field should be examined. If the version is earlier than the receiver's version and the earlier version can be

supported by the receiver, the receiver SHOULD send a Start-Control-Connection-Reply. If the version is earlier than the receiver's version and the version cannot be supported, the receiver SHOULD send a Start-Connection-Reply message, close the TCP connection and remain in the idle state. If the receiver's version is the same as earlier than the peer's, the receiver SHOULD send a Start-Control-Connection-Reply with the receiver's version and enter the established state.

established

An established connection may be terminated by either a local condition or the reception of a Stop-Control-Connection-Request. In the event of a local termination, the originator MUST send a Stop-Control-Connection-Request and enter the wait_stop_reply state.

If the originator receives a Stop-Control-Connection-Request it SHOULD send a Stop-Control-Connection-Reply and close the TCP connection, making sure that the final TCP information has been "pushed" properly.

wait_stop_reply

If a Stop-Control-Connection-Reply is received, the TCP connection SHOULD be closed and the control connection becomes idle.

3.1.3 Start Control Connection Initiation Request Collision

A PAC and PNS must have only one control connection between them. It is possible, however, for a PNS and a PAC to simultaneously attempt to establish a control connection to each other. When a Start-Control-Connection-Request is received on one TCP connection and another Start-Control-Connection-Request has already been sent on another TCP connection to the same peer, a collision has occurred.

The "winner" of the initiation race is the peer with the higher IP address (compared as 32 bit unsigned values, network number more significant). For example, if the peers 192.33.45.17 and 192.33.45.89 collide, the latter will be declared the winner.

The loser will immediately close the TCP connection it initiated, without sending any further PPTP control messages on it and will respond to the winner's request with a Start-Control-Connection-Reply message. The winner will wait for the Start-Control-Connection-Reply on the connection it initiated and also wait for a TCP termination indication on the connection the loser opened. The winner MUST NOT send any messages on the connection the loser initiated.

3.1.3 Keep Alives and Timers

A control connection SHOULD be closed by closing the underlying TCP connection under the following circumstances:

1. If a control connection is not in the established state (i.e., Start-Control-Connection-Request and Start-Control-Connection-Reply have not been exchanged), a control connection SHOULD be closed after 60 seconds by a peer waiting for a Start-Control-Connection-Request or Start-Control-Connection-Reply message.
2. If a peer's control connection is in the established state and has not received a control message for 60 seconds, it SHOULD send a Echo-Request message. If an Echo-Reply is not received 60 seconds after the Echo-Request message transmission, the control connection SHOULD be closed.

3.2 Call States

3.2.1 Timing considerations

Because of the real-time nature of telephone signaling, both the PNS and PAC should be implemented with multi-threaded architectures such that messages related to multiple calls are not serialized and blocked. The transit delay between the PAC and PNS should not exceed one second. The call and connection state figures do not specify exceptions caused by timers. The implicit assumption is that since the TCP-based control connection is being verified with keep-alives, there is less need to maintain strict timers for call control messages.

Establishing outbound international calls, including the modem training and negotiation sequences, may take in excess of 1 minute so the use of short timers is discouraged.

If a state transition does not occur within 1 minute (except for connections in the idle or established states), the integrity of the protocol processing between the peers is suspect and the ENTIRE CONTROL CONNECTION should be closed and restarted. All Call IDs are logically released whenever a control connection is started. This presumably also helps in preventing toll calls from being "lost" and never cleared.

3.2.2 Call ID values

Each peer assigns a Call ID value to each user session it requests or accepts. This Call ID value MUST be unique for the tunnel between the PNS and PAC to which it belongs. Tunnels to other peers can use the

same Call ID number so the receiver of a packet on a tunnel needs to associate a user session with a particular tunnel and Call ID. It is suggested that the number of potential Call ID values for each tunnel be at least twice as large as the maximum number of calls expected on a given tunnel.

A session is defined by the triple (PAC, PNS, Call ID).

3.2.3 Incoming calls

An Incoming-Call-Request message is generated by the PAC when an associated telephone line rings. The PAC selects a Call ID and serial number and indicates the call bearer type. Modems should always indicate analog call type. ISDN calls should indicate digital when unrestricted digital service or rate adaption is used and analog if digital modems are involved. Dialing number, dialed number, and subaddress may be included in the message if they are available from the telephone network.

Once the PAC sends the Incoming-Call-Request, it waits for a response from the PNS but does not answer the call from the telephone network. The PNS may choose not to accept the call if:

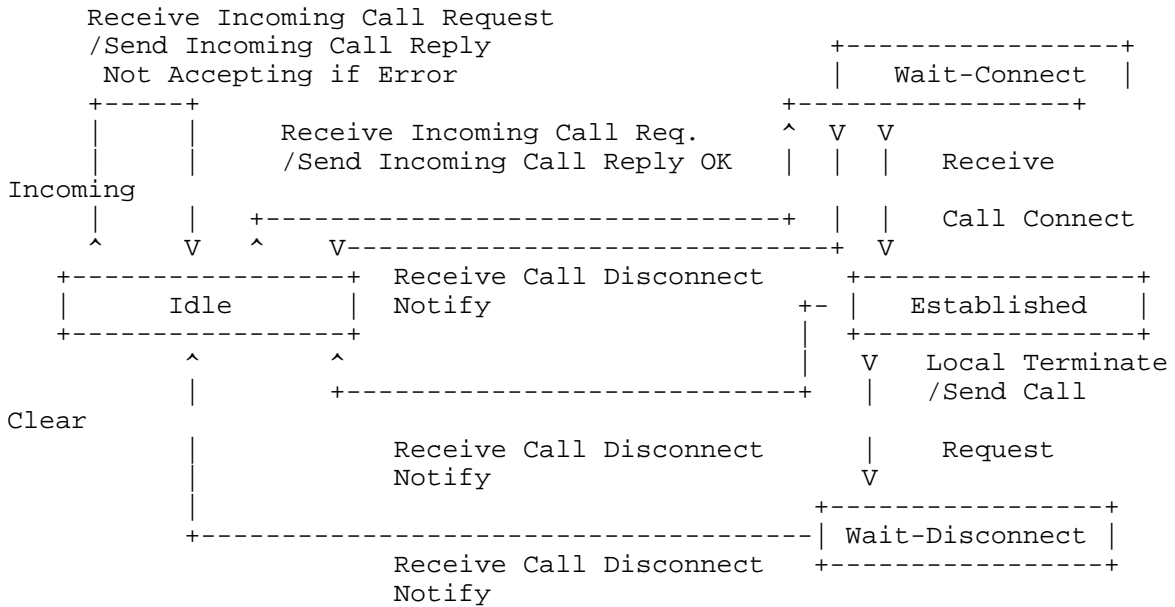
- No resources are available to handle more sessions
- The dialed, dialing, or subaddress fields are not indicative of an authorized user
- The bearer service is not authorized or supported

If the PNS chooses to accept the call, it responds with an Outgoing-Call-Reply which also indicates window sizes (see Appendix B). When the PAC receives the Outgoing-Call-Reply, it attempts to connect the call, assuming the calling party has not hung up. A final call connected message from the PAC to the PNS indicates that the call states for both the PAC and the PNS should enter the established state.

When the dialed-in client hangs up, the call is cleared normally and the PAC sends a Call-Disconnect-Notify message. If the PNS wishes to clear a call, it sends a Call-Clear-Request message and then waits for a Call-Disconnect-Notify.

A local reason. The PAC sends a Call-Disconnect-Notify message.

3.2.3.2 PNS Incoming Call States



The states associated with the PNS for incoming calls are:

idle

An Incoming-Call-Request message is received. If the request is not acceptable, an Incoming-Call-Reply is sent back to the PAC and the PNS remains in the idle state. If the Incoming-Call-Request message is acceptable, an Incoming-Call-Reply is sent indicating accept in the result code. The session moves to the wait_connect state.

wait_connect

If the session is connected on the PAC, the PAC sends an incoming call connect message to the PNS which then moves into established state. The PAC may send a Call-Disconnect-Notify to indicate that the incoming caller could not be connected. This could happen, for example, if a telephone user accidentally places a standard voice call to a PAC resulting in a handshake failure on the called modem.

established

The session is terminated either by receipt of a Call-Disconnect-Notify message from the PAC or by sending a Call-Clear-Request. Once a Call-Clear-Request has been sent, the session enters the wait_disconnect state.

wait_disconnect

Once a Call-Disconnect-Notify is received the session moves back to the idle state.

3.2.4 Outgoing calls

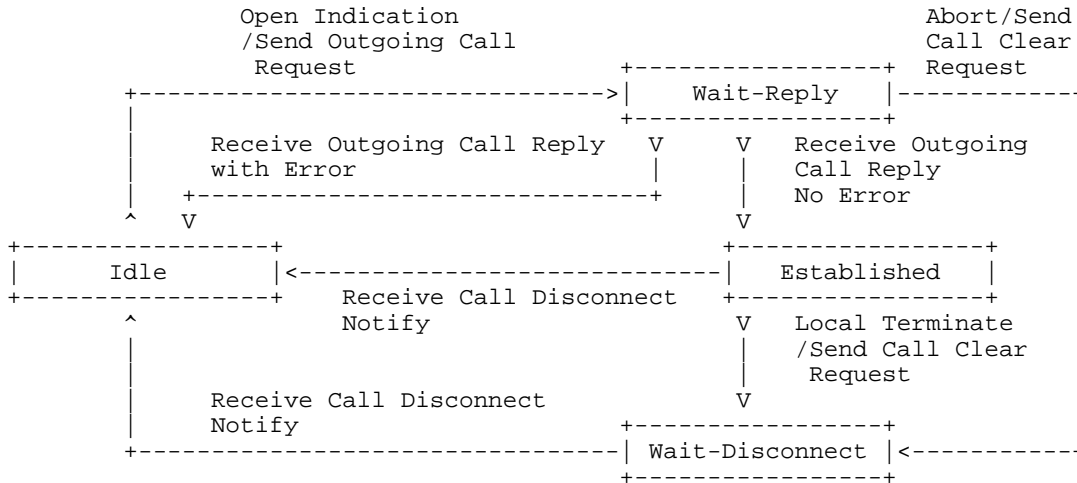
Outgoing messages are initiated by a PNS and instruct a PAC to place a call on a telco interface. There are only two messages for outgoing calls: Outgoing-Call-Request and Outgoing-Call-Reply. The PNS sends an Outgoing-Call-Request specifying the dialed party phone number and subaddress as well as speed and window parameters. The PAC MUST respond to the Outgoing-Call-Request message with an Outgoing-Call-Reply message once the PAC determines that:

The call has been successfully connected

A call failure has occurred for reasons such as: no interfaces are available for dial-out, the called party is busy or does not answer, or no dial tone is detected on the interface chosen for dialing

released via appropriate mechanisms and a Call-Disconnect-Notify message SHOULD BE sent to the PNS. If the call is disconnected by the client or by the telco interface, a Call-Disconnect-Notify message SHOULD be sent to the PNS.

3.2.4.2 PNS Outgoing Call States



The states associated with the PNS for outgoing calls are:

idle

An Outgoing-Call-Request message is sent to the PAC and the session moves into the wait_reply state.

wait_reply

An Outgoing-Call-Reply is received which indicates an error. The session returns to idle state. No telco call is active. If the Outgoing-Call-Reply does not indicate an error, the telco call is connected and the session moves to the established state.

established

If a Call-Disconnect-Notify is received, the telco call has been terminated for the reason indicated in the Result and Cause Codes. The session moves back to the idle state. If the PNS chooses to


terminate the session, it sends a Call-Clear-Request to the PAC and then enters the wait_disconnect state.

wait_disconnect

A session disconnection is waiting to be confirmed by the PAC. Once the PNS receives the Call-Disconnect-Notify message, the session enters idle state.

4.0 Tunnel Protocol Operation

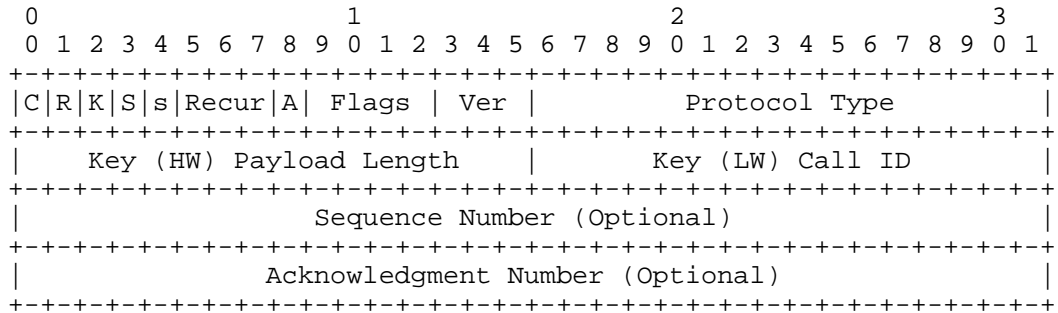
The user data carried by the PPTP protocol are PPP data packets. PPP packets are carried between the PAC and PNS, encapsulated in GRE packets which in turn are carried over IP. The encapsulated PPP packets are essentially PPP data packets less any media specific framing elements. No HDLC flags, bit insertion, control characters, or control character escapes are included. No CRCs are sent through the tunnel. The IP packets transmitted over the tunnels between a PAC and PNS has the following general structure:



4.1 Enhanced GRE header

The GRE header used in PPTP is enhanced slightly from that specified in the current GRE protocol specification [1,2]. The main difference involves the definition of a new Acknowledgment Number field, used to determine if a particular GRE packet or set of packets has arrived at the remote end of the tunnel. This Acknowledgment capability is not used in conjunction with any retransmission of user data packets. It is used instead to determine the rate at which user data packets are to be transmitted over the tunnel for a given user session.

The format of the enhanced GRE header is as follows:



- C (Bit 0) Checksum Present. Set to zero (0).
- R (Bit 1) Routing Present. Set to zero (0).
- K (Bit 2) Key Present. Set to one (1).
- S (Bit 3) Sequence Number Present. Set to one (1) if a payload (data) packet is present. Set to zero (0) if payload is not present (GRE packet is an Acknowledgment only).
- s (Bit 4) Strict source route present. Set to zero (0).
- Recur (Bits 5-7) Recursion control. Set to zero (0).
- A (Bit 8) Acknowledgment sequence number present. Set to one (1) if packet contains Acknowledgment Number to be used for acknowledging previously transmitted data.
- Flags (Bits 9-12) Must be set to zero (0).
- Ver (Bits 13-15) Must contain 1 (enhanced GRE).
- Protocol Type Contains the protocol ID for PPTP [6].
- Key Use of the Key field is up to the

implementation.

PPTP uses it as follows:

Payload Length

(High 2 octets of Key) Size of the payload, not including the GRE header

Call ID

(Low 2 octets) Contains the Peer's Call ID for the session to which this packet belongs.

Sequence Number

Contains the sequence number of the payload. Present if S bit (Bit 3) is one (1).

Acknowledgment Number

Contains the sequence number of the highest numbered GRE packet received by the sending peer for this user session. Present if A bit (Bit 8) is one (1).

The payload section contains a PPP data packet without any media specific framing elements.

The sequence numbers involved are per packet sequence numbers. The sequence number for each user session is set to zero at session startup. Each packet sent for a given user session which contains a payload (and has the S bit (Bit 3) set to one) is assigned the next consecutive sequence number for that session.

This protocol allows acknowledgments to be carried with the data and makes the overall protocol more efficient, which in turn requires less buffering of packets.

4.2 Sliding Window Protocol

The sliding window protocol used on the PPTP data path is used for flow control by each side of the data exchange. The enhanced GRE protocol allows packet acknowledgments to be piggybacked on data packets. Acknowledgments can also be sent separately from data packets. Again, the main purpose of the sliding window protocol is for flow control--retransmissions are not performed by the tunnel peers.

4.3 Multi Packet Acknowledgment

One feature of the PPTP sliding window protocol is that it allows the acknowledgment of multiple packets with a single acknowledgment. All outstanding packets with a sequence number lower or equal to the acknowledgment number are considered acknowledged. Time-out calculations are performed using the time the packet corresponding to the highest sequence number being acknowledged was transmitted.

Adaptive time-out calculations are only performed when an Acknowledgment is received. When multi-packet acknowledgments are used, the overhead of the adaptive time-out algorithm is reduced. The PAC is not required to transmit multi-packet acknowledgments; it can instead acknowledge each packet individually as it is delivered to the PPP client.

4.4 Out-of-Sequence Packets

Occasionally packets lose their sequencing across a complicated internetwork. Say, for example that a PNS sends packets 0 to 5 to a PAC. Because of rerouting in the internetwork, packet 4 arrives at the PAC before packet 3. The PAC acknowledges packet 4, and may assume packet 3 is lost. This acknowledgment grants window credit beyond packet 4.

When the PAC does receive packet 3, it **MUST** not attempt to transmit it to the corresponding PPP client. To do so could cause problems, as proper PPP protocol operation is premised upon receiving packets in sequence. PPP does properly deal with the loss of packets, but not with reordering so out of sequence packets between the PNS and PAC **MUST** be silently discarded, or they may be reordered by the receiver. When packet 5 comes in, it is acknowledged by the PAC since it has a higher sequence number than 4, which was the last highest packet acknowledged by the PAC. Packets with duplicate sequence numbers should never occur since the PAC and PNS never retransmit GRE packets. A robust implementation will silently discard duplicate GRE packets, should it receive any.

5.0 Security Considerations

Security issues are not addressed in this document. End to end security is addressed by PPP. Further security considerations will be addressed by the next version of PPTP.

6.0 Authors' Addresses

Kory Hamzeh
Ascend Communications
1275 Harbor Bay Parkway
Alameda, CA 94502
Email: kory@ascend.com

Gurdeep Singh Pall
Microsoft Corporation
Redmond, WA
Email: gurdeep@microsoft.com

William Verthein
U.S. Robotics

Jeff Taarud
3COM Corporation

W. Andrew Little
ECI Telematics

7.0 References

- [1] Hanks, S. Li, T., Farinacci, D. Traina, P., "Generic Routing Encapsulation (GRE)", RFC 1701, October 1994.
- [2] Hanks, S. Li, T., Farinacci, D. Traina, P., "Generic Routing Encapsulation (GRE) over IPv4 Networks", RFC 1702, October 1994.
- [3] Lloyd, B. Simpson, W., "PPP Authentication Protocols", RFC 1334, October 1992.
- [4] Postel, J. B., "Transmission Control Protocol", RFC 791, September 1981
- [5] Postel, J. B., "User Data Protocol", RFC 768, August 1980.
- [6] Reynolds, J., Postel, J., "Assigned Numbers", RFC 1700, October, 1994.
- [7] Simpson, W., editor, "The Point-to-Point Protocol (PPP)", RFC 1661, July 1994.

Appendix A. Acknowledgment Time-Outs

PPTP uses sliding windows and time-outs to provide both user session flow-control across the internetwork and to perform efficient data buffering to keep the PAC-PNS data channels full without causing receive buffer overflow. PPTP requires that a time-out be used to recover from dropped data or acknowledgment packets. The exact implementation of the time-out is vendor-specific. It is suggested that an adaptive time-out be implemented with backoff for congestion control. The time-out mechanism proposed here has the following properties:

Independent time-outs for each session. A device (PAC or PNS) will have to maintain and calculate time-outs for every active session.

An administrator-adjustable maximum time-out, MaxTimeOut, unique to each device.

An adaptive time-out mechanism that compensates for changing throughput. To reduce packet processing overhead, vendors may choose not to recompute the adaptive time-out for every received acknowledgment. The result of this overhead reduction is that the time-out will not respond as quickly to rapid network changes.

Timer backoff on time-out to reduce congestion. The backed-off timer value is limited by the configurable maximum time-out value. Timer backoff is done every time an acknowledgment time-out occurs.

In general, this mechanism has the desirable behavior of quickly backing off upon a time-out and of slowly decreasing the time-out value as packets are delivered without time-outs.

Some definitions:

Packet Processing Delay (PPD) is the amount of time required for each side to process the maximum amount of data buffered in their receive packet sliding window. The PPD is the value exchanged between the PAC and PNS when a call is established. For the PNS, this number should be small. For a PAC making modem connections, this number could be significant.

Sample is the actual amount of time incurred receiving an acknowledgment for a packet. The Sample is measured, not calculated.

Round-Trip Time (RTT) is the estimated round-trip time for an

Acknowledgment to be received for a given transmitted packet. When the network link is a local network, this delay will be minimal (if not zero). When the network link is the Internet, this delay could be substantial and vary widely. RTT is adaptive: it will adjust to include the PPD and whatever shifting network delays contribute to the time between a packet being transmitted and receiving its acknowledgment.

Adaptive Time-Out (ATO) is the time that must elapse before an acknowledgment is considered lost. After a time-out, the sliding window is partially closed and the ATO is backed off.

Packet Processing Delay (PPD)

The PPD parameter is a 16-bit word exchanged during the Call Control phase that represents tenths of a second (64 means 6.4 seconds). The protocol only specifies that the parameter is exchanged, it does not specify how it is calculated. The way values for PPD are calculated is implementation-dependent and need not be variable (static time-outs are allowed). The PPD must be exchanged in the call connect sequences, even if it remains constant in an implementation. One possible way to calculate the PPD is:

```
PPD' = ((PPP_MAX_DATA_MTU - Header) * WindowSize * 8) / ConnectRate
PPD = PPD' + PACFudge
```

Header is the total size of the IP and GRE headers, which is 36. The MTU is the overall MTU for the internetwork link between the PAC and PNS. WindowSize represents the number of packets in the sliding window, and is implementation-dependent. The latency of the internetwork could be used to pick a window size sufficient to keep the current session's pipe full. The constant 8 converts octets to bits (assuming ConnectRate is in bits per second). If ConnectRate is in bytes per second, omit the 8. PACFudge is not required but can be used to take overall processing overhead of the PAC into account.

The value of PPD is used to seed the adaptive algorithm with the initial RTT[n-1] value.

Sample

Sample is the actual measured time for a returned acknowledgment.

Round-Trip Time (RTT)

The RTT value represents an estimate of the average time it takes for an acknowledgment to be received after a packet is sent to the remote

end of the tunnel.

A.1 Calculating Adaptive Acknowledgment Time-Out

We still must decide how much time to allow for acknowledgments to return. If the time-out is set too high, we may wait an unnecessarily long time for dropped packets. If the time-out is too short, we may time out just before the acknowledgment arrives. The acknowledgment time-out should also be reasonable and responsive to changing network conditions.

The suggested adaptive algorithm detailed below is based on the TCP 1989 implementation and is explained in Richard Steven's book TCP/IP Illustrated, Volume 1 (page 300). 'n' means this iteration of the calculation, and 'n-1' refers to values from the last calculation.

```
DIFF[n] = SAMPLE[n] - RTT[n-1]
DEV[n] = DEV[n-1] + (beta * (|DIFF[n]| - DEV[n-1]))
RTT[n] = RTT[n-1] + (alpha * DIFF[n])
ATO[n] = MIN (RTT[n] + (chi * DEV[n]), MaxTimeOut)
```

DIFF represents the error between the last estimated round-trip time and the measured time. DIFF is calculated on each iteration.

DEV is the estimated mean deviation. This approximates the standard deviation. DEV is calculated on each iteration and stored for use in the next iteration. Initially, it is set to 0.

RTT is the estimated round-trip time of an average packet. RTT is calculated on each iteration and stored for use in the next iteration. Initially, it is set to PPD.

ATO is the adaptive time-out for the next transmitted packet. ATO is calculated on each iteration. Its value is limited, by the MIN function, to be a maximum of the configured MaxTimeOut value.

Alpha is the gain for the average and is typically 1/8 (0.125).

Beta is the gain for the deviation and is typically 1/4 (0.250).

Chi is the gain for the time-out and is typically set to 4.

To eliminate division operations for fractional gain elements, the entire set of equations can be scaled. With the suggested gain constants, they should be scaled by 8 to eliminate all division. To simplify calculations, all gain values are kept to powers of two so that shift operations can be used in place of multiplication or

division.

A.2 Congestion Control: Adjusting for Time-Out

This section describes how the calculation of ATO is modified in the case where a time-out does occur. When a time-out occurs, the time-out value should be adjusted rapidly upward. Although the GRE packets are not retransmitted when a time-out occurs, the time-out should be adjusted up toward a maximum limit. To compensate for shifting internetwork time delays, a strategy must be employed to increase the time-out when it expires. A simple formula called Karn's Algorithm is used in TCP implementations and may be used in implementing the backoff timers for the PNS or the PAC. Notice that in addition to increasing the time-out, we are also shrinking the size of the window as described in the next section.

Karn's timer backoff algorithm, as used in TCP, is:

$$\text{NewTIMEOUT} = \text{delta} * \text{TIMEOUT}$$

Adapted to our time-out calculations, for an interval in which a time-out occurs, the new ATO is calculated as:

$$\begin{aligned} \text{RTT}[n] &= \text{delta} * \text{RTT}[n-1] \\ \text{DEV}[n] &= \text{DEV}[n-1] \\ \text{ATO}[n] &= \text{MIN} (\text{RTT}[n] + (\text{chi} * \text{DEV}[n]), \text{MaxTimeOut}) \end{aligned}$$

In this modified calculation of ATO, only the two values that contribute to ATO and that are stored for the next iteration are calculated. RTT is scaled by chi, and DEV is unmodified. DIFF is not carried forward and is not used in this scenario. A value of 2 for Delta, the time-out gain factor for RTT, is suggested.

Appendix B. Acknowledgment Time-Out and Window Adjustment

B.1 Initial Window Size

Although each side has indicated the maximum size of its receive window, it is recommended that a slow start method be used to begin transmitting data. The initial window size on the transmitter is set to half the maximum size the receiver requested, with a minimum size of one packet. The transmitter stops sending packets when the number of packets awaiting acknowledgment is equal to the current window size. As the receiver successfully digests each window, the window size on the transmitter is bumped up by one packet until the maximum is reached. This method prevents a system from flooding an already

congested network because no history has been established.

B.2 Closing the Window

When a time-out does occur on a packet, the sender adjusts the size of the transmit window down to one half its value when it failed. Fractions are rounded up, and the minimum window size is one.

B.3 Opening the Window

With every successful transmission of a window's worth of packets without a time-out, the transmit window size is increased by one packet until it reaches the maximum window size that was sent by the other side when the call was connected. As stated earlier, no retransmission is done on a time-out. After a time-out, the transmission resumes with the window starting at one half the size of the transmit window when the time-out occurred and adjusting upward by one each time the transmit window is filled with packets that are all acknowledged without time-outs.

B.4 Window Overflow

When a receiver's window overflows with too many incoming packets, excess packets are thrown away. This situation should not arise if the sliding window procedures are being properly followed by the transmitter and receiver. It is assumed that, on the transmit side, packets are buffered for transmission and are no longer accepted from the packet source when the transmit buffer fills.