Independent Submission
Request for Comments: 5578
Obsoletes: 4938
Category: Informational
ISSN: 2070-1721

B. Berry, Ed.
S. Ratliff
E. Paradise
Cisco
T. Kaiser
Harris Corporation
M. Adams
L3 Communications

February 2010

PPP over Ethernet (PPPoE) Extensions for Credit Flow and Link Metrics

Abstract

This document extends the Point-to-Point Protocol over Ethernet (PPPoE) with an optional credit-based flow control mechanism and an optional Link Quality Metric report. These optional extensions improve the performance of PPPoE over media with variable bandwidth and limited buffering, such as mobile point-to-point radio links.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc5578.

IESG Note

The PPP Extensions Working Group (PPPEXT) has reservations about the desirability of the feature described in this document. In particular, it solves a general problem at an inappropriate layer and it may have unpredictable interactions with higher and lower level protocols. The techniques described in this document are intended for use with a particular deployment technique that uses a PPP termination separated from a radio termination by an Ethernet, and that has radio-side flow control for a slower PPP-only link to remote

nodes. Implementors are better advised to avoid split termination with inter-media protocol translation, and use standard Internet Protocol routing instead.

Table of Contents

1. Introduction

PPP over Ethernet (PPPoE) [2] is a protocol for establishing and encapsulating sessions between Hosts (Clients) and traffic-access aggregators (Servers) for PPP [1] transport over real or emulated Ethernet. PPPoE works well when both session endpoints have similar bandwidth, forwarding, and buffering capabilities that do not vary over time. However, improvements can be made for applications with variable bandwidth and limited buffering. This document addresses these improvements with optional extensions to PPPoE that support credit-based session flow control and session-based link metric quality reports.

These extensions are designed to support radio systems that exhibit point-to-point waveforms. The diagram below is used to illustrate the improvement that these extensions address. When the local Client (Radio) detects the presence of a remote Radio neighbor, it initiates a PPPoE session with its local Server (router). The Radio also establishes a radio link connection with the remote Radio over the point-to-point RF (radio frequency) link (which is beyond the scope of this document). The remote Radio is also establishing a PPPoE session with its local Server (router). The Radios associate the two PPPoE sessions and the point-to-point radio link protocol (RLP), creating a complete data path. Now a PPP session is established via the PPP IP Control Protocol (IPCP) as described in RFC 1661. Included in this IPCP exchange is the router IP address. With the exchange of the IPCP IP addresses, each router inserts the remote IP address into its local routing tables. Note that the radio IP information is not inserted into the routing tables.

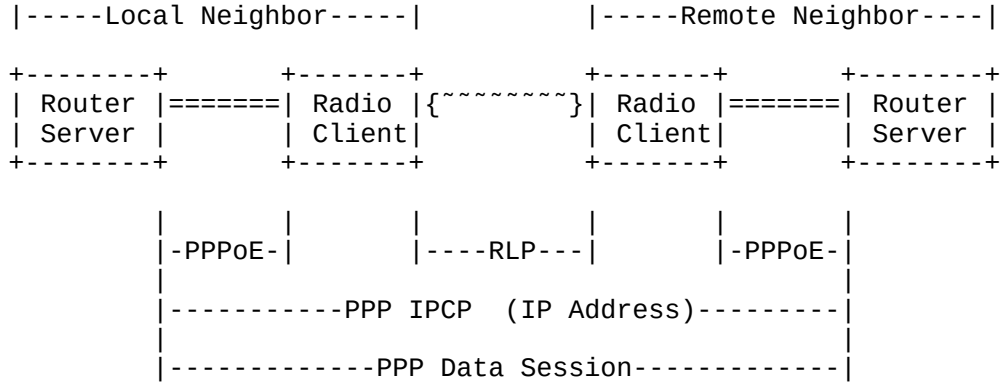

Figure 1: PPPoE Network

The capabilities of the RF links between RLP neighbors will vary over time due to mobility and environmental conditions as well as changes in the RF waveforms and encoding. To reflect these dynamic changes, the Radio may periodically generate Link Quality Metrics to the router. The router uses the link metric to update route costs and influence route selection. The influence upon the routing protocols is beyond the scope of this document.

A PPPoE Client implementation can be found at [3]. It is open source (GNU GPLv2 -- General Public License).

2. Terminology

Access Concentrator
: The RFC 2516 term used to describe the Server. This document uses the terms Router or Server instead.

BCN
: Backward Credit Notification. The BCN represents the number of remaining credits at the local node that were granted by peer node. The local node uses these credits to transmit payload back to the peer node. BCN ranges from 0-65535.

CDR
: The Current Datarate.

FCN
: Forward Credit Notification. The FCN represents the credits that the local node is granting to the peer node. The peer node uses these granted credits to transmit payload back to the local node. FCN ranges from 0-65535.

Gbit/s
: gigabits (1,000,000,000) per second.

Host
: The RFC 2516 term used to describe the Server. This document uses the terms Radio or Client.

kbit/s
: kilobits (1,000) per second.

LCP
: PPP Link Control Protocol, RFC 1661.

Mbit/s
: megabits (1,000,000) per second.

MDR
: The Maximum Datarate.

NCP
: PPP Network Control Protocol, RFC 1661.

RLP
: Radio Link Protocol.

TAG
: The RFC 2516 PPPoE synonym for TLV. This document uses TLV.

Tbit/s terabits (1,000,000,000,000) per second.

TLV Type-Length-Value.

3. Overview of Protocol Extensions

The extensions consist of optional TLVs as well as enhanced and newly defined Discovery packets. The enhancements are applied to the Discovery Stage and the PPP Session Stage.

3.1. TLVs

The new TLVs are listed in the table below:

TLV Value	TLV Description
0x0106	Credits TLV
0x0107	Metrics TLV
0x0108	Sequence Number TLV
0x0109	Credit Scale Factor TLV

3.1.1. Credits TLV

This TLV contains the Forward Credit Notification (FCN) and the Backward Credit Notification (BCN). The Credits TLV is optional with the PADR (PPPoE Active Discovery Request) and PADS (PPPoE Active Discovery Session-Confirmation) and required in the PADG (PPPoE Active Discovery Session Grant) and PADC (PPPoE Active Discovery Session-Credit) Discovery packets (ETHER_TYPE=8863).

The Credits TLV is optionally carried in the PPPoE data payload packet of the PPP Stage (ETHER_TYPE=8864).

The FCN represents the number of credits being granted by the local node to the peer node. The BCN represents the number of credits remaining at the local node that were granted by the peer node.

The Credits TLV is shown below:

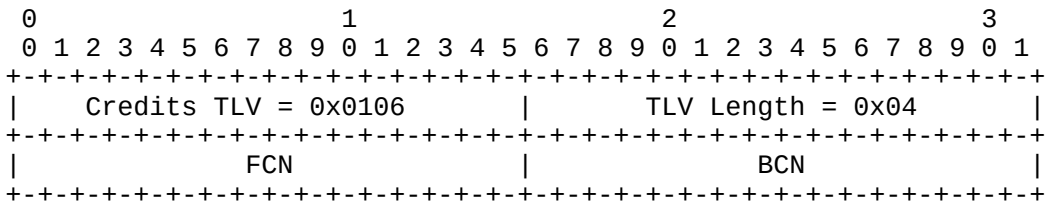

These fields are transmitted in network byte order. The same byte order is used throughout the document in other structures as well.

3.1.2. Metrics TLV

The Metrics TLV is used to report the link-quality parameters. The Metrics TLV is required with the PADQ (PPPoE Discovery Quality) Discovery packet.

The Metrics TLV contains the following fields:

Resources - a percentage, 0-100, representing the amount of remaining or available resources, such as battery power. If resources cannot be calculated, a value of 100 should be reported.

Relative Link Quality (RLQ) - a non-dimensional number, 0-100, representing the relative link quality. A value of 100 represents a link of the highest quality. If the RLQ cannot be calculated, a value of 100 should be reported.

Receive only - a bit that indicates whether the link is bi-directional or receive-only. A value of "1" indicates that the link is receive-only.

Reserved - reserved fields are zeroed unless otherwise specified.

CD - two bits that designate the units of the Current Datarate.

```
CD Scale:  00 == kbit/s (default)
           01 == Mbit/s
           10 == Gbit/s
           11 == Tbit/s
```

MD - two bits that designate the units of the Maximum Datarate.

```
MD Scale:  00 == kbit/s (default)
           01 == Mbit/s
           10 == Gbit/s
           11 == Tbit/s
```

Current Datarate - the Current Datarate, in un-scaled units per second, achieved on the RLP link. If the Radio makes no distinction between Maximum Datarate and Current Datarate, Current Datarate should equal the Maximum Datarate. When metrics are reported, the Current Datarate must be reported. The Current Datarate must be less than or equal to the Maximum Datarate.

Latency - the transmission delay that a packet encounters as it is transmitted over the link. This is reported in absolute delay, milliseconds. If latency cannot be calculated, a value of 0 should be reported. The calculation of latency is Radio dependent. For example, the latency may be a running average calculated from the internal queuing.

Maximum Datarate - the maximum theoretical data rate, in un-scaled units per second, that the RLP link is capable of providing. When metrics are reported, the Maximum Datarate must be reported.

The Metrics TLV is shown below:

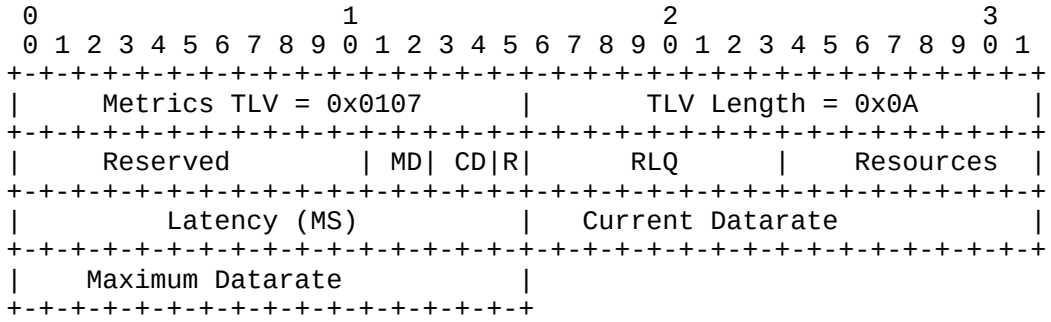

3.1.3. Sequence Number TLV

This TLV is used to carry a unique 16-bit sequence number in order to identify a specific request and the associated response. The sequence number should be initialized to zero and incremented by one for each new request. For retransmitted packets, the same sequence number that was used in the previous packet transmission is repeated. The PADG and PADC packets require the Sequence Number Tag.

The Sequence Number TLV is shown below:

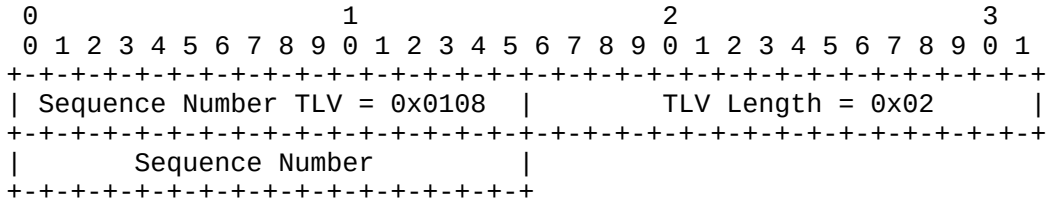

3.1.4. Credit Scale Factor TLV

This TLV contains the scale factor value that is to be applied to the session credit calculations. The Credit Scale Factor TLV is optional with the PADR and PADS packets. Once the session is established with specified scale factors, the scale factors are set for the entire session. The scale factor value represents the units that the local node grants to the remote node. The remote node is responsible for maintaining the credit accounting relative to the data flow back to the local node.

The Credit Scale Factor TLV can be used to change from the default 64-byte credit unit during the PADR-PADS exchange. The credit scale factor value can range from 1 byte to 65535 bytes. A zero value is ignored and the default 64-byte unit remains set. The scale factor is set per each payload flow: peer-to-local and local-to-peer.

The Credit Scale Factor TLV is shown below:

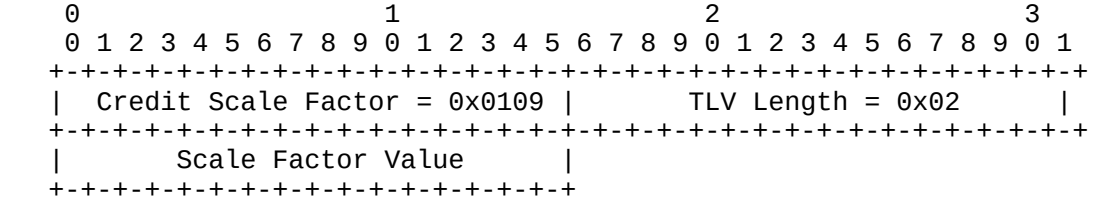

3.2. Discovery Stage Extensions

The specifications of the PPPoE Active Discovery Request (PADR) and the PPPoE Active Discovery Session-confirmation (PADS) packets are extended to include the optional Credits TLV and the Credit Scale Factor TLV. The PPPoE Active Discovery Session-Grant (PADG) packet, the PPPoE Active Discovery Session-Credit Response (PADC), and the Quality packets are newly defined Discovery Stage packets.

Discovery Packet	Status	
PADR	Enhanced	Optionally includes the Credits TLV and the Credit Scale Factor TLV
PADS	Enhanced	Optionally includes the Credits TLV and the Credit Scale Factor TLV
PADG	New	Includes the Credits TLV and the Sequence Number TLV
PADC	New	Includes the Credits TLV and the Sequence Number TLV
PADQ	New	Includes the Metrics TLV

3.2.1. PPPoE Active Discovery Request (PADR)

The PADR packet is extended to optionally contain a single Credits TLV, indicating that the Client requests credit flow control for this session. The Credits TLV contains the Forward Credit Notification (FCN) and the Backward Credit Notification (BCN) to be applied to the PPP Session Stage. The FCN provides the initial credits granted to the Server by the Client. The BCN value is set to 0, as the Client has not yet been granted credits from the Server.

The PADR packet is enhanced to optionally contain a single Credit Scale Factor TLV. The Credit Scale Factor TLV defines the credit scale factor value. If the Credit Scale Factor TLV is omitted, the default 64-byte value is used for the session. When the Client includes the optional Credit Scale Factor TLV in the PADR, this credit scale factor value is applied to all credit grants associated with the Client credits that are granted to the Server.

The Server must echo the Credit Scale Factor TLV in the PADS response to confirm the credit scaling session and to designate the Server credit scaling factor. This PADS Credit Scaling Factor TLV represents the scale factor value that is applied to all credits granted from the Server to the Client.

Once the session is established during the PADR-PADS exchange, the credit scale factor value cannot be changed.

A Discovery PADR packet with the optional Credits TLV is shown below:

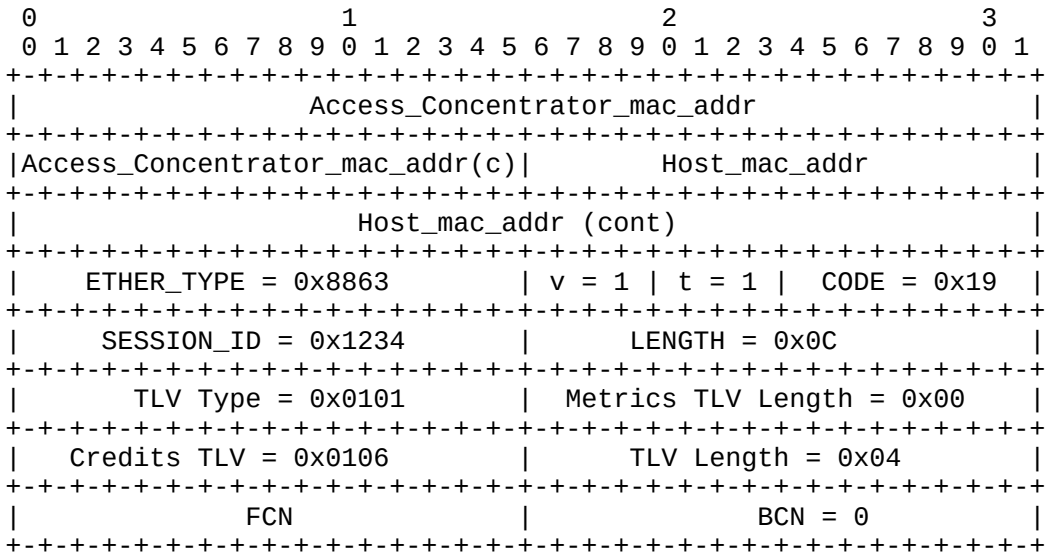

The credit units are expressed in the default 64-byte units.

A Discovery PADR packet with the optional Credits TLV and the optional Credit Scale Factor TLV is shown below:

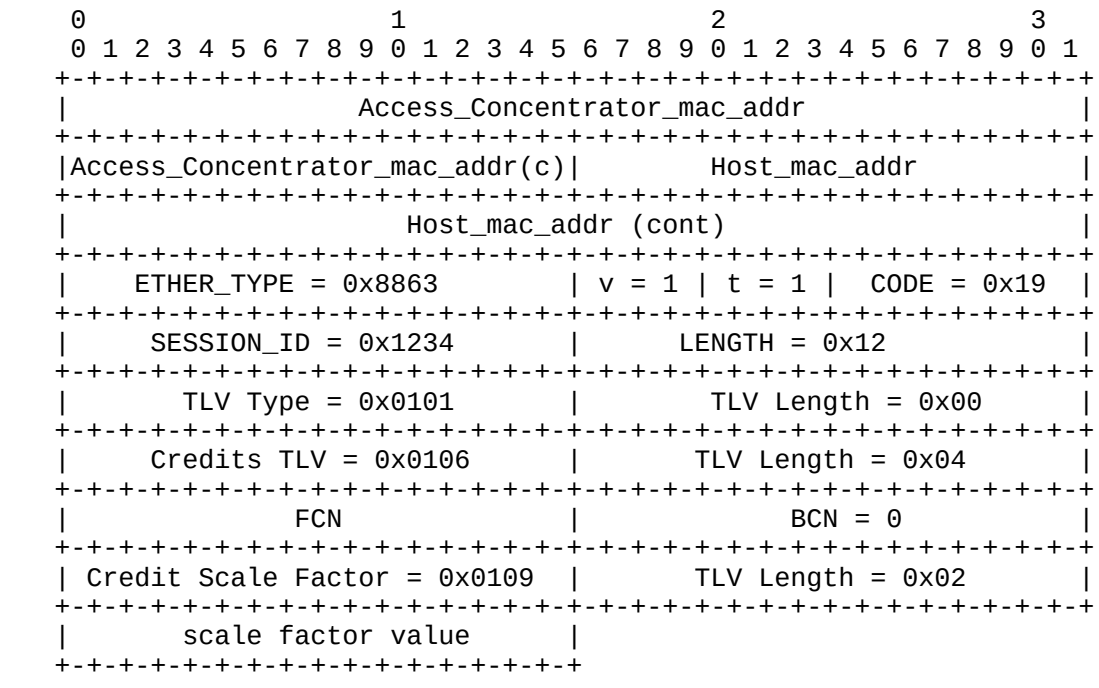

The Credits TLV FCN value is expressed in units of the session's credit scale factor value.

3.2.2. PPPoE Active Discovery Session-Confirmation (PADS)

The Server PADS is extended to optionally contain a single Credits TLV, indicating the Forward Credit Notification (FCN) and the Backward Credit Notification (BCN) of the PPP Session Stage.

If the Client PADR contained a Credits TLV, then the Server PADS must indicate support for credit flow control by including a Credits TLV. The PADS Credits TLV FCN represents the number of credits initially granted to the Client. The Credits TLV BCN is an echo of the number of credits that the Client had granted to the Server in the originating PADR packet.

Exchange of the Credits TLV in the PADR and PADS indicates that credit flow control is supported by both the Server and the Client for the designated PPP Session Stage. This is binding and must be

followed for the entire duration of the PPP Session Stage. A session's credit binding must be established prior to any other credit indications being exchanged.

The Server PADS should only include the Credits TLV in response to a Client PADR that included the Credits TLV. If the Server does not support credit flow, it should not include the Credits TLV in its PADS response. The Client must terminate a credit-based session that cannot be supported by the Server. A Credits TLV transmitted outside an established credit-based session must be ignored.

The Server PADS is enhanced to optionally contain a single Credit Scale Factor TLV. The Credit Scale Factor TLV defines the credit scale unit value. The Credit Scale Factor TLV must be included if it was included in the Client PADR. If the Credits TLV was not included in the originating PADR, it must be omitted, indicating that the 64-byte default is used for the directional flow. This credit scale factor is applied to Server grants to the Client.

A Discovery PADS packet with the optional Credits TLV is shown below:

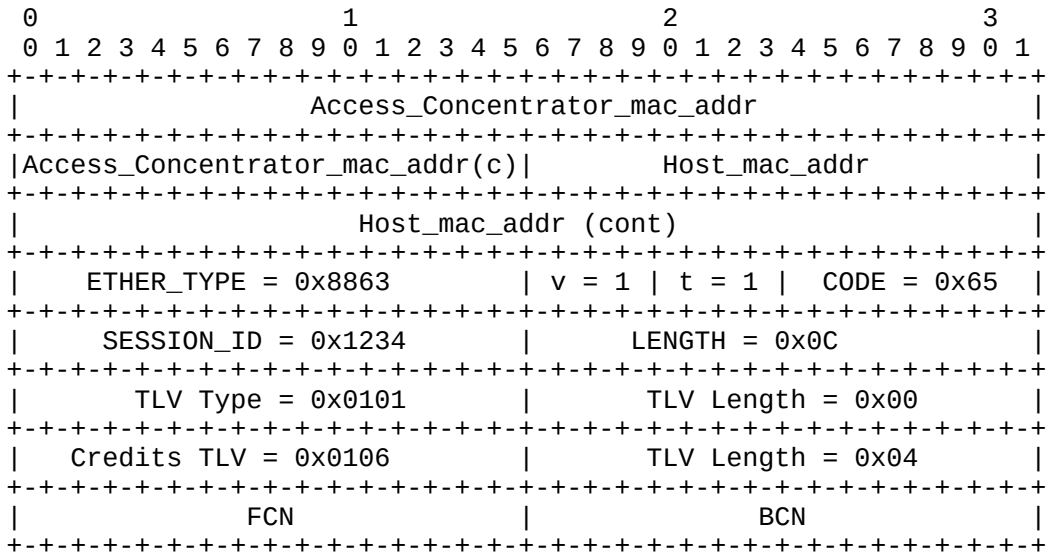

The BCN is expressed in the default 64-byte units.

A Discovery PADS packet with the optional Credits TLV and the optional Credit Scale Factor TLV is shown below:

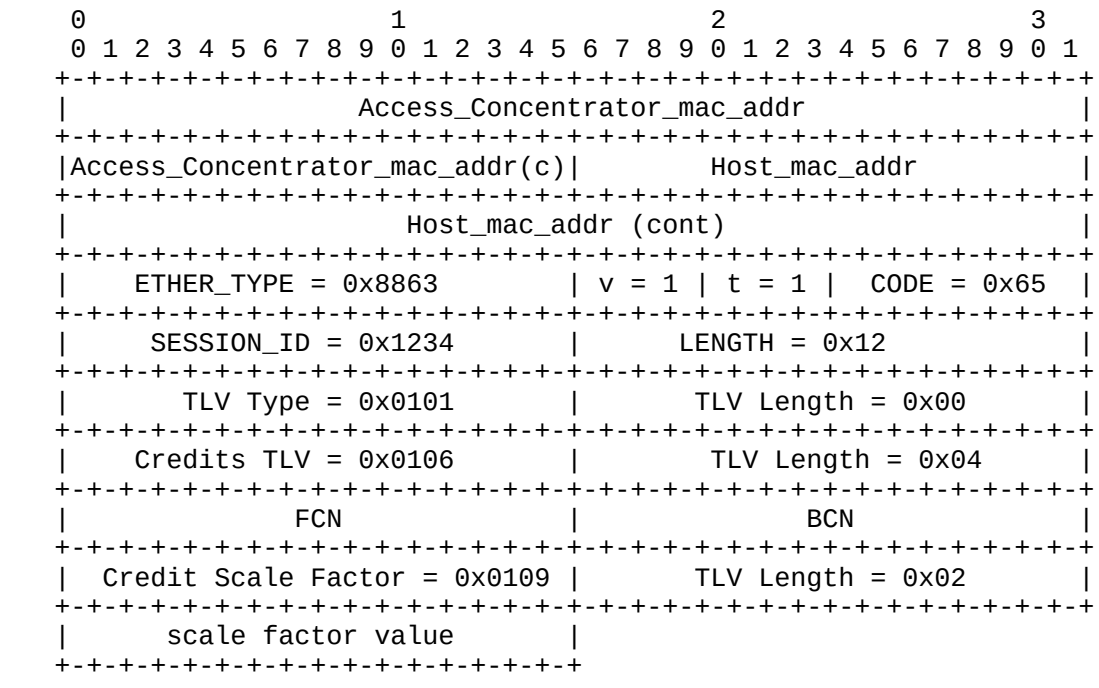

The Credits TLV BCN value is expressed in units of the session scale factor value received in the PADR.

3.2.3. PPPoE Active Discovery Session-Grant (PADG)

The PPPoE Active Discovery Session-Grant (PADG) is a new packet defined in this specification. The local node (Server or Client) may send a PADG at any time after the PADR/PADS exchange to grant incremental flow control credits to a peer. The CODE field is set to 0x0A and the SESSION_ID must be set to the unique value generated for this PPPoE Session.

Each flow control credit corresponds to the amount of PPP payload bytes that can be sent or received. For example, if the default credit scale factor of 64 bytes is used, and 128 bytes of PPP payload data are sent, then 2 credits would be consumed. When calculating credits to consume, all credit calculations must be rounded up. If, in the previous example, 130 bytes of PPP payload data were sent, 3 credits would have been consumed.

When the peer receives a PADG packet, it adds the incremental credits to its working credit count and responds with a PPPoE Active Discovery Session-Credit Response (PADC) packet, indicating the accumulation of the credits. The FCN and BCN values must be scaled by the value established during session establishment in the Credit Scale Factor TLV or by the default 64-byte value prior to processing.

The PADG packet must contain a single Credits TLV, indicating the Forward Credit Notification (FCN) and the Backward Credit Notification (BCN) of the PPPoE Session.

The Credits TLV FCN indicates the number of incremental credits being granted to the peer by the node. A value between 1 and 0xffff represents an incremental credit grant. The peer must multiply the credit units by the credit scale factor and add these credits to its accumulated transmit credit count. A value of 0x0000 represents a NULL grant, meaning that there are no additional credits being granted.

The Credits TLV BCN indicates the remaining absolute credits that have been granted by the peer to the local node. When the local node exhausts the BCN, it must stop transmitting payload packets.

Once a credit has been granted, it must be honored. The largest number of incremental credits at any time is 0xffff.

The PADG packet must contain a single Sequence Number TLV. This TLV is used to carry a unique 16-bit sequence number to uniquely identify each request. The sequence number should be initialized at zero and incremented by one for each new PADG. For retransmitted PADGs, the same sequence number that was used in the previous packet transmission is repeated.

A Discovery PADG packet with the Sequence Number and Credits TLVs is shown below:

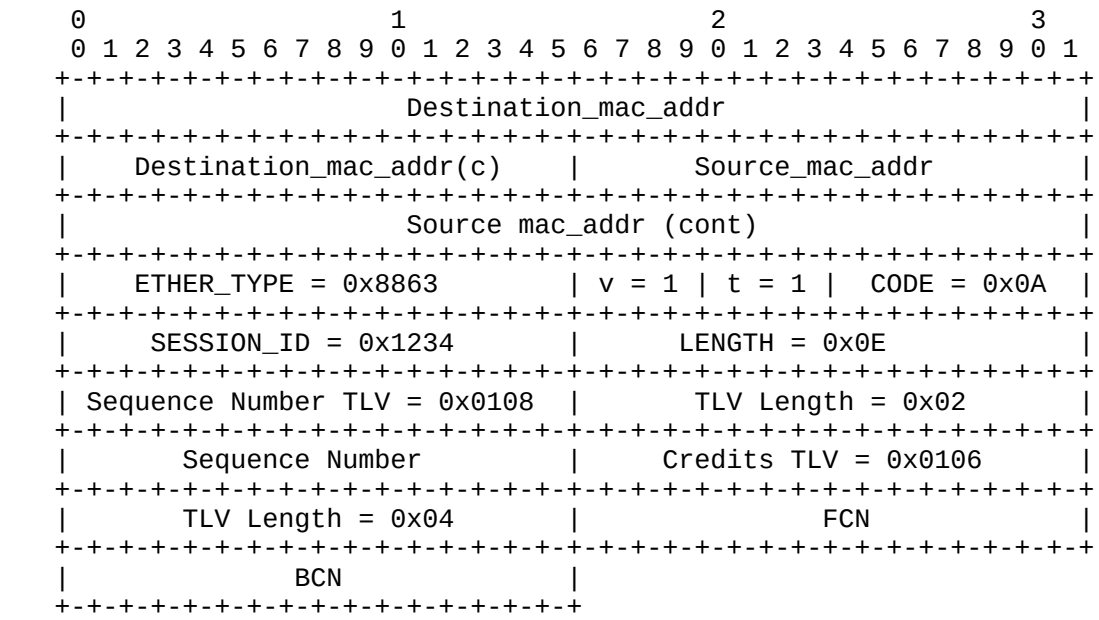

3.2.4. PPPoE Active Discovery Session-Credit Response (PADC)

The PPPoE Active Discovery Session-Credit Response (PADC) is a new packet defined in this specification. A Server or Client must send a PADC in response to a PADG. The CODE field is set to 0x0B, and the SESSION_ID must be set to the unique value generated for this PPPoE session.

The PADC packet must contain a single Credits TLV, indicating the Forward Credit Notification (FCN) and the Backward Credit Notification (BCN) of the PPPoE session.

The Credits TLV FCN represents the absolute credits remaining that have been granted to the peer by the node. The Credits TLV BCN represents the remaining absolute credits that have been granted to the local node from the peer. The FCN and BCN values must be scaled by the value established during session establishment in the Credit Scale Factor TLV or by the default 64-byte value prior to processing.

The PADC packet must contain a single Sequence Number TLV. The sequence number must be the sequence number associated with the PADG.

A Discovery PADC packet with the Sequence Number and Credits TLV is shown below:

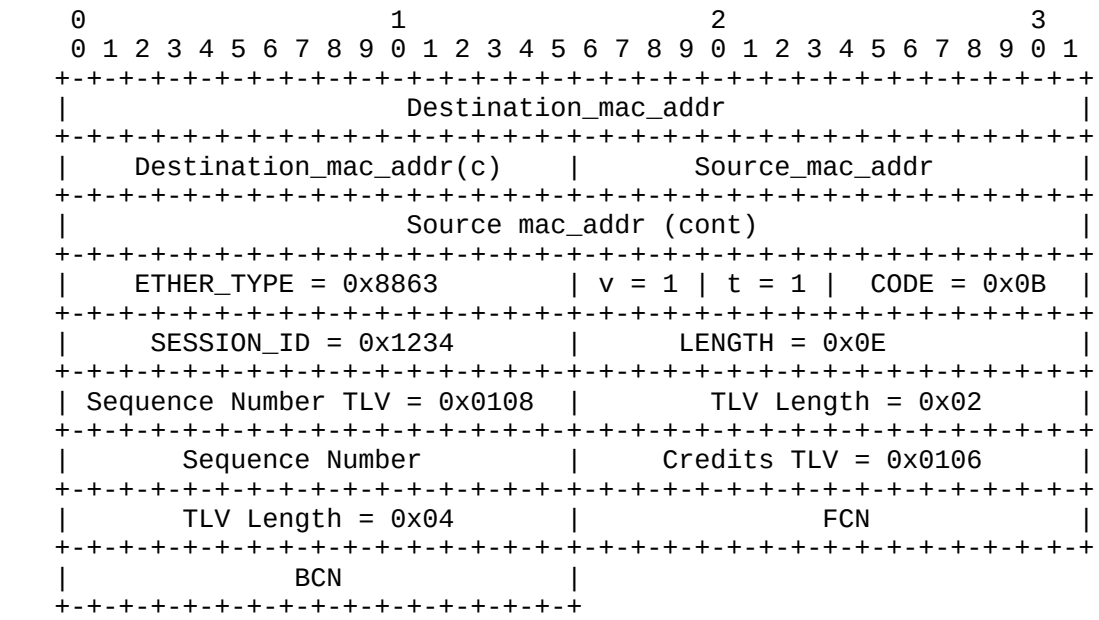

The FCN and BCN values are expressed in the respective units defined by the Credit Scale Factor TLV or the 64-byte default.

3.2.5. PPPoE Active Discovery Quality (PADQ)

The PPPoE Active Discovery Quality (PADQ) is a new packet defined in this specification. An Server or Client may send an optional PADQ at any time to query or report link-quality metrics.

When transmitting PPP [1] streams over wireless links through radio modems, the quality of the RF link directly affects the throughput. The PPPoE Active Discovery Quality (PADQ) packet can be used by the radio modem to report RF link metrics. The CODE field is set to 0x0C, and the SESSION_ID must be set to the unique value generated for this PPPoE session.

The PPPoE Active Discovery Quality (PADQ) packet can be used to query link metrics by setting the PADQ Metrics TLV Length to zero.

The PADQ must carry a single Metrics TLV. When processing the data rates, the values must be converted using the indicated data rate units. This document enhances the Metrics TLV as described below.

A Discovery PADQ packet with the required Metrics TLV is shown below:

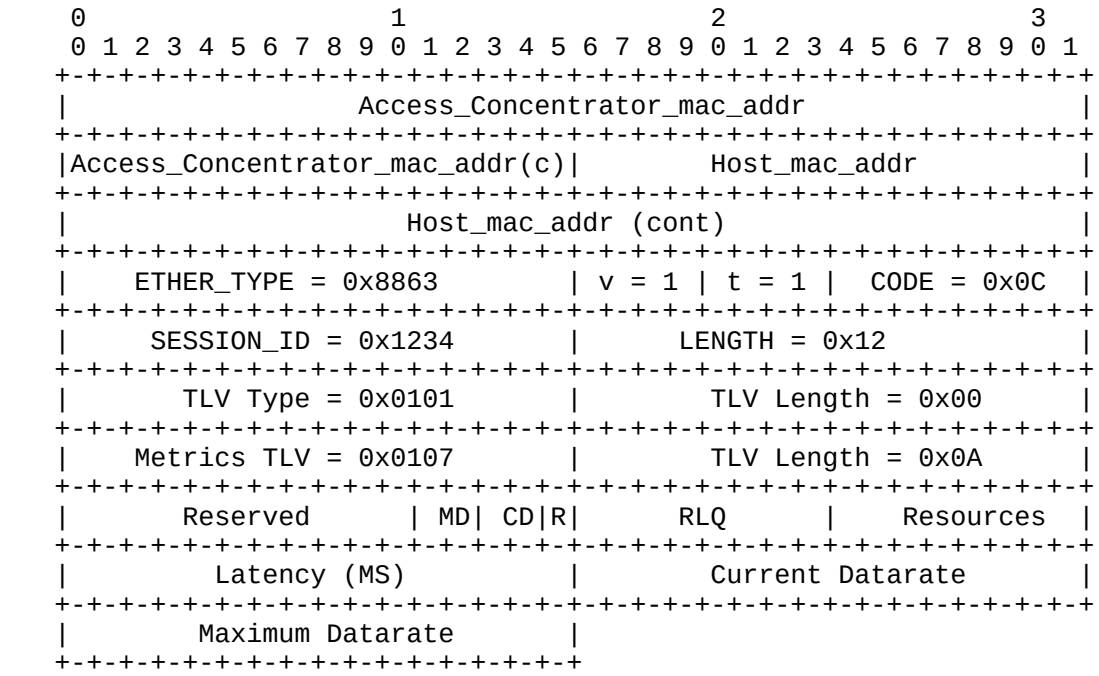

The Maximum Datarate and the Current Datarate are expressed in units determined by the MD and CD bits, respectively.

A Discovery PADQ packet with a Metrics TLV Length=0 to query is shown below:

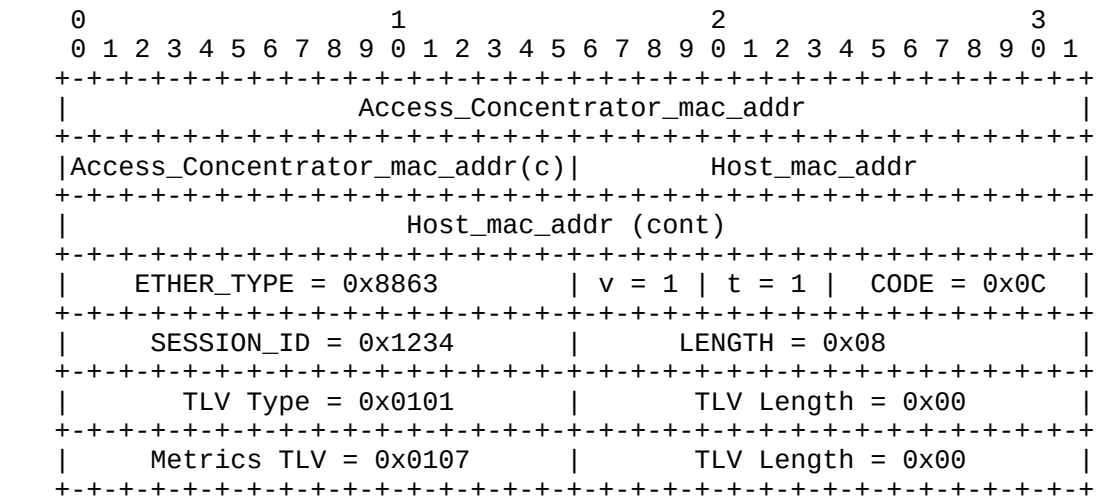

3.3. PPP Session Stage Extensions

The PPP Session Stage Extensions define the optional use of Credits TLV. Use of the Credits TLV in the PPP Session Stage is referred to as an in-band credit grant.

The first field following the PPP Session Stage LENGTH must be checked. If the value is equal to the PPP Protocol identifier (0xc021), then normal packet (payload) processing occurs. When the field following the PPP Session Stage LENGTH is not the PPP Protocol identifier (0xc021), a TLV is assumed. In this case, the TLV length is subtracted from the overall payload length.

A PPP LCP packet with optional Credits TLV is shown below:

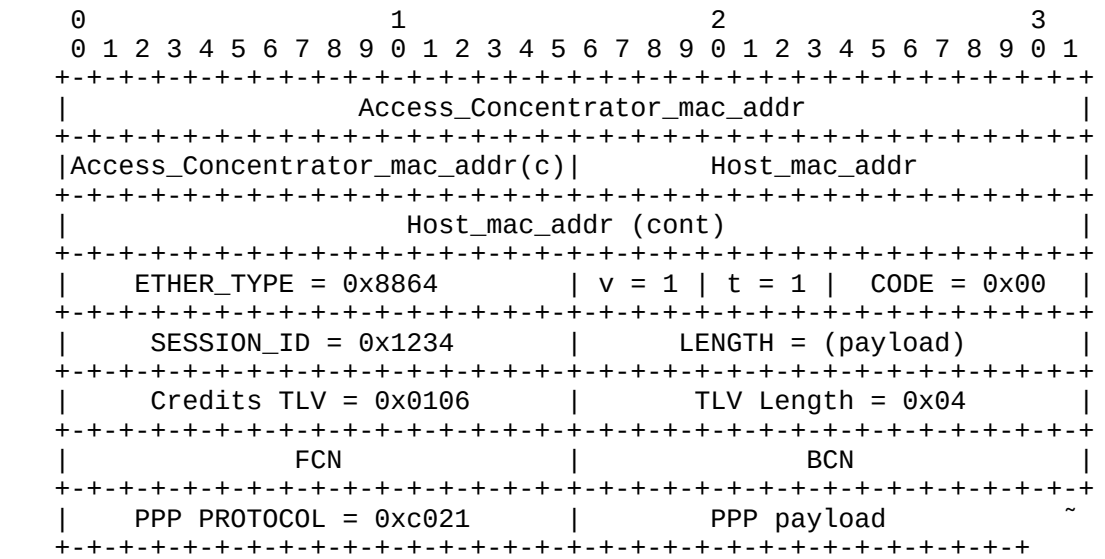

4. Credit Flow Considerations

For a given session, credit grants exchanged in the Discovery Stage, PADG-PADC, are referred to as out-of-band. Credit grants exchanged in the PPP Session Stage are referred to as in-band. Credit accounting is only applied to the packets transmitted in the PPP Session Stage.

Out-of-band credit management is handled by periodic exchange of the PPPoE Active Discovery Session-Grant (PADG) and PPPoE Active Discovery Credit Response (PADC) packets.

In-band credit management allows credits to be incrementally granted with each PPP Session Stage packet. These in-band incremental credit grants are not explicitly acknowledged. However, they are reflected in the in-band credit flow from the peer node. This offers the greatest credit-granting efficiency when traffic rates are high.

Once agreed upon during the Discovery Stage, credit grants are required to transmit packets in the PPP Session Stage. A node must grant credits to its peer before the peer can transmit packets to the granting node.

Credits are granted incrementally in the forward direction. Locally, a node manages the credits that it has granted to a peer, as well as the credits that a peer has granted to it.

Grants received from a peer are added to a local running credit counter. The accumulated credits are decremented with each packet the node transmits to the peer. When the running counter reaches zero, the node must stop transmitting packets to the peer.

To manage the credits that a node has granted, the node maintains a running counter. With each PPP Session Stage packet received from the peer, the running counter is decremented. When the running counter reaches zero, no additional packets are expected. The node incrementally grants more credits to the peer to maintain packet flow. Packets received when granted credits have been exhausted are discarded.

5. PADG and PADC Retransmission

When a node does not receive a PADC packet in response to a PADG within a specified amount of time, it should transmit a new PADG packet with zero credits, using the same sequence number and doubling the waiting period. A PADC response with the associated sequence number will indicate whether or not the previously granted credits were accumulated. If they were not, a PADG with credits and an incremented sequence number should be transmitted. This process should be repeated until granted credits are properly acknowledged or as many times as desired.

When a node does not receive a PADQ metric packet in response to a query within a specified amount of time, it should resend the PADQ query packet and double the waiting period. This can be repeated as many times as desired.

6. Other Considerations

A node may autonomously generate PADQ metric packets. The rate of autonomously generated PADQ metric packets may need to be throttled so as not to overrun the peer.

The sending and receiving of PPPoE Discovery packets are independent of credit counts. For example, a node must always be able to receive a PADG and send a PADC.

During normal operation, nodes may disagree about the number of credits. Operational credit mismatches would occur due to packets in transit on the wire. Much larger credit mismatches can occur if there are transmission errors. To correct these larger errors, the BCN fields of the PADG and PADC packets and in-band credit grants from a peer can be used by the receiving node to reset the credit values of its peer.

7. IANA Considerations

IANA has assigned the following PPPoE TLV Values for which this RFC serves as the reference:

TLV Value	TLV Name	Description	Reference
262 0x0106	Credits	See the reference	[RFC5578]
263 0x0107	Metrics	See the reference	[RFC5578]
264 0x0108	Sequence Number	See the reference	[RFC5578]
265 0x0109	Credit Scale Factor	See the reference	[RFC5578]

IANA has assigned the following PPPoE Code fields for which this RFC serves as the reference:

Code	PPPoE Packet Name	Description	Reference
10 0x0a	PADG, Session-Grant	See the reference	[RFC5578]
11 0x0b	PADC, Session-Credit Response	See the reference	[RFC5578]
12 0x0c	PADQ, Quality	See the reference	[RFC5578]

8. Security Considerations

This memo defines a mechanism for adding flow control to the existing PPP over Ethernet (PPPoE) sessions. These extensions are subsequent to the existing PPPoE security mechanisms as described in RFC 2516 [2]. It is required that the Service TLV and Session ID always be validated prior to processing credits.

9. References

9.1. Normative References

[1] Simpson, W., Ed., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.

[2] Mamakos, L., Lidl, K., Evarts, J., Carrel, D., Simone, D., and R. Wheeler, "A Method for Transmitting PPP Over Ethernet (PPPoE)", RFC 2516, February 1999.

9.2. Informative References

[3] An open source (GPLv2) PPPoE Client implementation of RFC 5578, PPP Over Ethernet (PPPoE) Extensions for Credit Flow and Link Metrics, http://rfc4938.sourceforge.net/.

Appendix A. Examples of Session Credit Flows

Session Credit Flow with the default 64-byte credit unit.

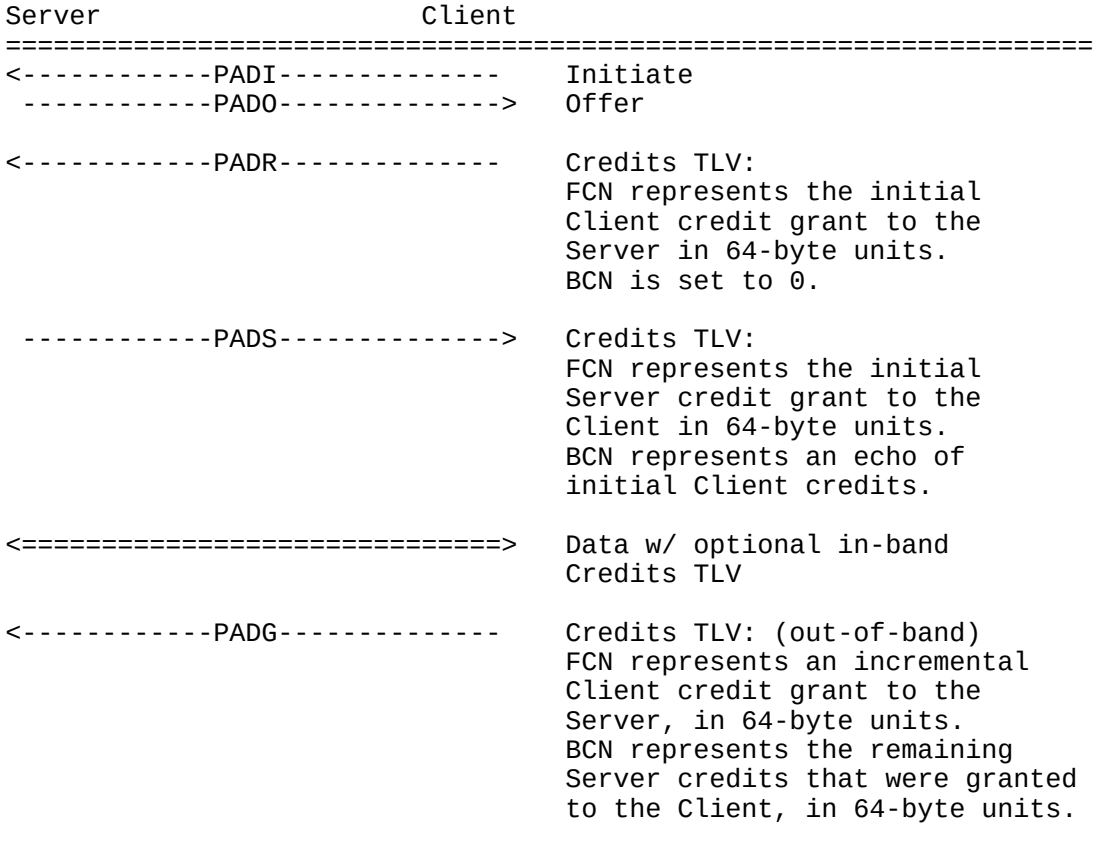

Session Credit Flow with specific credit scale factor units for the Server and the Client.

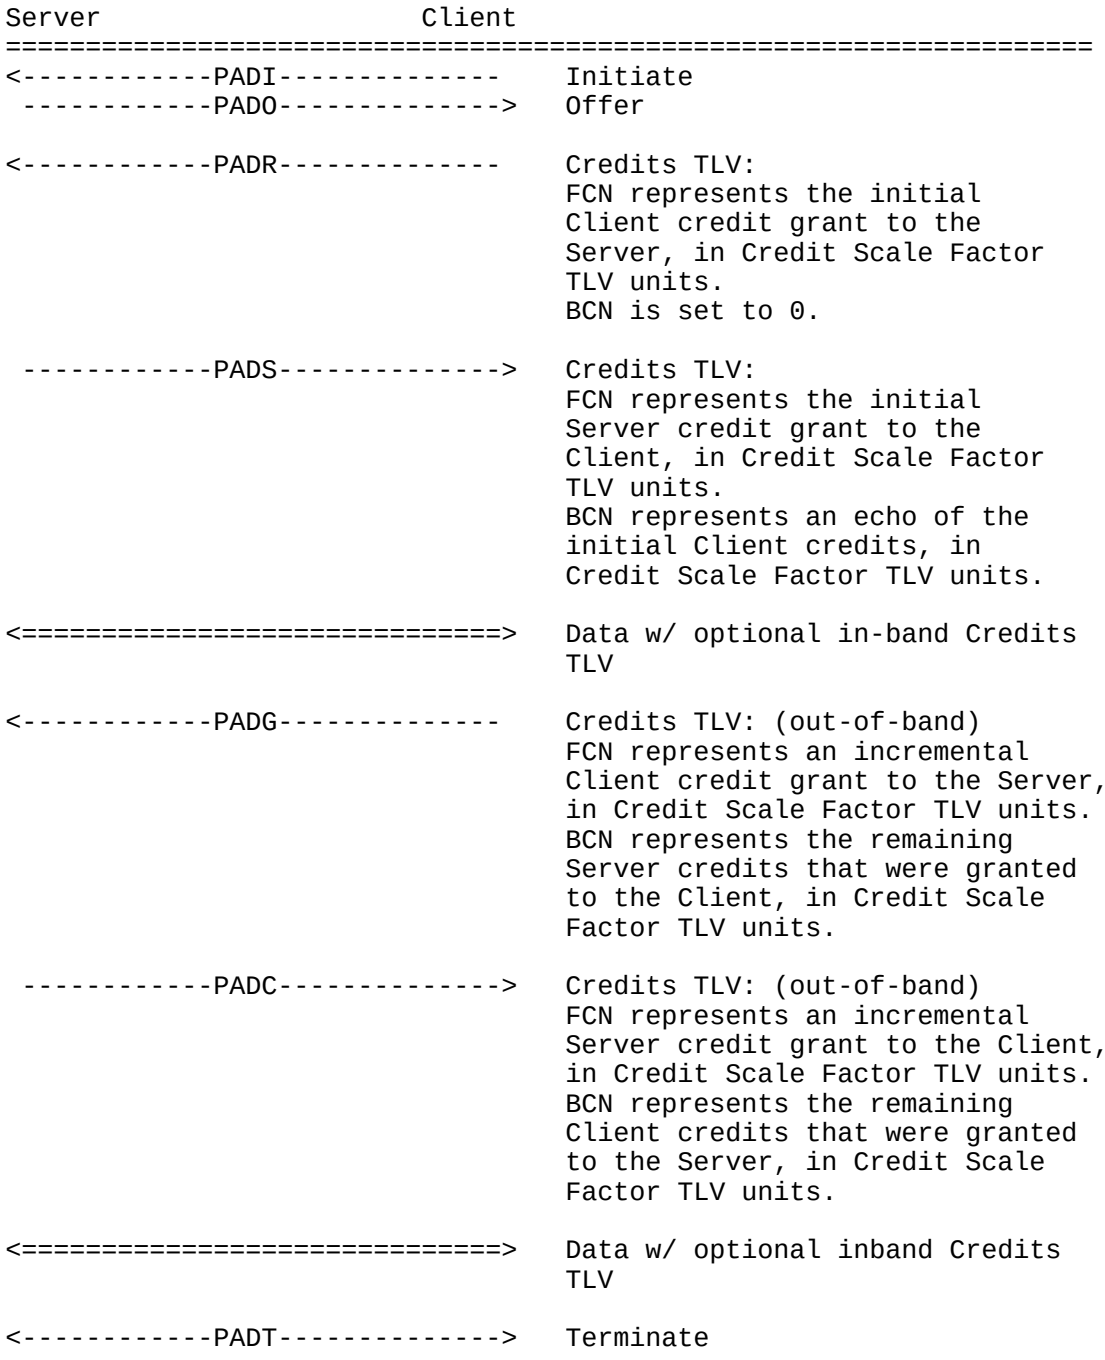
Server
Client
<----------PADI----------
Initiate
----------PADO---------->
Offer
<----------PADR----------
Credits TLV:
FCN represents the initial
Client credit grant to the
Server, in Credit Scale Factor
TLV units.
BCN is set to 0.
----------PADS---------->
Credits TLV:
FCN represents the initial
Server credit grant to the
Client, in Credit Scale Factor
TLV units.
BCN represents an echo of the
initial Client credits, in
Credit Scale Factor TLV units.
<==========>
Data w/ optional in-band Credits
TLV
<----------PADG----------
Credits TLV: (out-of-band)
FCN represents an incremental
Client credit grant to the Server,
in Credit Scale Factor TLV units.
BCN represents the remaining
Server credits that were granted
to the Client, in Credit Scale
Factor TLV units.
----------PADC---------->
Credits TLV: (out-of-band)
FCN represents an incremental
Server credit grant to the Client,
in Credit Scale Factor TLV units.
BCN represents the remaining
Client credits that were granted
to the Server, in Credit Scale
Factor TLV units.
<==========>
Data w/ optional inband Credits
TLV
<----------PADT---------->
Terminate

Authors' Addresses

Bo Berry, Editor
Cisco
170 West Tasman Drive
San Jose, CA 95134
EMail: bberry@cisco.com

Stan Ratliff
Cisco
170 West Tasman Drive
San Jose, CA 95134
EMail: sratliff@cisco.com

Ed Paradise
Cisco
170 West Tasman Drive
San Jose, CA 95134
EMail: pdice@cisco.com

Tim Kaiser
Harris Corporation
Government Communications System Division
Mail Stop 25-11F
P.O. Box 37
Melbourne, FL 32902-0037
EMail: timothy.kaiser@harris.com

Michael D Adams
640 N 2200 W MS F1J12
Salt Lake City, Utah 84116
EMail: Michael.D.Adams@L-3com.com