

INTERNET-DRAFT
Category: Informational
Expires: October 29, 2006

B. Berry
H. Holgate
Cisco Systems, Inc.
10 March 2004

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

[draft-bberry-pppoe-credit-05.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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."

This document may not be modified, and derivative works of it may not be created, except to publish it as an RFC and to translate it into languages other than English.

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on October 29, 2006.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

This document extends the Point-to-Point Over Ethernet (PPPoE) Protocol [2] with a credit-based flow control mechanism and Link Quality Metric report. This optional extension should improve the performance of PPPoE over media with variable bandwidth and limited buffering, such as mobile radio links.

Berry, Holgate

Expires October 29, 2006

[Page 1]

Table of Contents

1.	Introduction.....	3
2.	Payload.....	4
3.	Overview of Protocol Extensions.....	4
4.	Discovery Stage.....	4
4.1	PPPoE Active Discovery Request (PADR).....	4
4.2	PPPoE Active Discovery Session-confirmation (PADS).....	5
4.3	PPPoE Active Discovery Session-Grant (PADG).....	5
4.4	PPPoE Active Discovery Session-Credit Response (PADC).....	6
4.5	PPPoE Active Discovery Quality (PADQ).....	7
5.	PPP Session Stage.....	8
6.	Credit Flow Considerations.....	8
7.	Other Considerations.....	9
8.	IANA Considerations.....	9
9.	Security Considerations.....	10
10.	Appendix A: Tag Values.....	10
11.	Appendix B: Example Message Formats.....	12
12.	Acknowledgements.....	19
13.	Normative References.....	19
	Authors' Addresses.....	19
	Intellectual Property and Copyright Statements.....	20

1. Introduction

PPP over Ethernet (PPPoE) [2] is a protocol for establishing and encapsulating sessions between hosts and traffic aggregators (Access Concentrators) for PPP 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, it is insufficient for applications with variable bandwidth and limited buffering, for example, mobile radio links. This document addresses this problem by suggesting an extension to PPPoE to support credit-based session flow control and session-based link metric exchanges.

The diagram below illustrates the problem that this extension is intended to solve, for the case of a radio link. Here PPPoE sessions are used between access concentrators (routers) and radio transmission systems that are shown as radio neighbors. Each radio transmission system establishes point-to-point Radio Link Protocol (RLP) sessions with its neighbors and establishes a corresponding PPPoE session for each neighbor with the transmission system's associated access concentrator (router). The radio logically associates the PPPoE session with the corresponding RLP session.




Figure 1: PPPoE Network

The capabilities of the RF links between RLP neighbors may vary over time due to mobility and environmental conditions. In many instances, the Host Radio has limited buffering capability to handle capacity changes in the RLP sessions. To limit buffering in the Host Radio, the PPPoE credit flow control mechanism provides dynamic buffering feedback to the access concentrator.

In the diagram above, from the access concentrator's perspective, each PPPoE session between it and the Host Radio represent a connection to a remote routable peer. For efficient routing, the local Host Radio uses the link metric mechanism to dynamically update the access concentrator route cost of the associated link.

While the example shows an RF based application, the extensions are applicable to other media.

2. Payload

The Ethernet payload version field retains its value of 0x01. The extensions for credit flow control and link quality metrics are optional and backward compatible.

3. Overview of Protocol Extensions

PPPoE has two distinct stages. There is a Discovery Stage and a PPP Session Stage. During the Discovery Stage, the Host can optionally request a flow controlled PPP Session Stage. Once the Access Concentrator acknowledges the Host flow control request, all PPP Session Stage traffic must be flow controlled.

4. Discovery Stage

The packet exchange of the Discovery Stage is unchanged by this specification. The specifications of the Session Request (PADR) and the Session Confirmation (PADS) packets are extended to include the optional Credit Tag Type-Length-Value (TLV).

In addition, the optional Credit Grant (PADG) packet, the Credit Response (PADC) packet and the Link Quality Metric (PADQ) packets are introduced.

4.1 PPPoE Active Discovery Request (PADR)

The PADR packet is extended to optionally contain a single Credit Tag TLV, indicating that the Host requests credit flow control for this session. The Credit Tag 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 Access Concentrator by the Host. The BCN value is set to 0.

An example packet is shown in [Appendix B](#).

4.2 PPPoE Active Discovery Session-confirmation (PADS)

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

If the PADR contained a Credit Tag, then the Access Concentrator PADS packet indicates support for credit flow control by including a Credit Tag. The PADS Credit Tag FCN represents the number of credits being initially granted to the Host. The Credit Tag BCN is an echo of the number of credits that the Host had granted to the Access Concentrator in the previous PADR packet.

Exchange of the Credit Tag TLV in the PADR and PADS indicates that credit flow control is supported by both the Access Concentrator and the Host 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 can be exchanged.

The Access Concentrator PADS should only carry the Credit Tag in response to a Host PADR with Credits. If the Access Concentrator does not support credit flow, it should not include the Credit Tag in its PADS response. The Host must terminate a credit based session that can not be supported by the Access Concentrator. Credit Tags transmitted outside an established credit based session must be ignored.

An example packet is shown in [Appendix B](#).

4.3 PPPoE Active Discovery Session-Grant (PADG)

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

The peer may then transmit data until the credits are exhausted.

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 (PADC) packet indicating the accumulated credits.

The PADG packet must contain a single Credit Tag TLV, indicating the Forward Credit Notification (FCN) and the Backward Credit

Notification (BCN) of the PPP Session.

Berry, Holgate

Expires October 29, 2006

[Page 5]

The Credit Tag FCN indicates the number of incremental credits being granted to the peer by the node. A value between 1 and 0xffff represent an incremental credit grant. The peer must 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 Credit Tag BCN indicates the remaining absolute credits that have been granted by the peer to the node.

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

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

An example packet is shown in [Appendix B](#).

4.4 PPPoE Active Discovery Session-Credit Response (PADC)

The PPPoE Active Discovery Session-Credit Response (PADC) is a new packet defined in this specification. An Access Concentrator or Host 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 Credit Tag TLV, indicating the Forward Credit Notification (FCN) and the Backward Credit Notification (BCN) of the PPPoE session, and any number of other Tag types.

The Credit Tag FCN represents the absolute credits remaining that have granted to the peer by the node. The Credit Tag BCN represents the remaining absolute credits that have been granted to the node from the peer.

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

An example packet is shown in [Appendix B](#).

Berry, Holgate

Expires October 29, 2006

[Page 6]

4.5 PPPoE Active Discovery Quality (PADQ)

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

When transmitting PPP 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 PADQ must carry a single Metric Tag TYPE, which contains the following fields:

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

Maximum data rate - is the maximum theoretical data rate, in kilobits per second (kbps), that the Host link is capable of providing. When metrics are reported, the maximum data rate must be reported.

Current data rate - is the current data rate, in kilobits per second (kbps), achieved on the Host link. If there is no distinction between maximum data rate and current data rate, current data rate should equal to maximum data rate.

Latency - is the transmission delay that a packet encounters as it is transmitted over the Host link. This is reported in absolute delay, milliseconds. If latency can not be calculated, a value of 0 should be reported.

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

Relative Link Quality (RLQ) - is 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 can not be calculated, a value of 100 should be reported.

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

An example packet is shown in [Appendix B](#).

5. PPP Session Stage

This specification defines the optional use of TLV Tags in the PPP Session Stage. 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 Tag length is subtracted from the overall payload length.

The Credit Tag is the only optional TLV permitted in the PPP Session Stage. The Credit Tag TLV is used to support in-band flow control.

A PPP Session Stage packet with Credits is shown in [Appendix B](#).

6. 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 processing 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 Grant (PADG) and PPPoE Active Discovery Credit (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 unacknowledged. They are however 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 stops transmitting packets to the peer. The values of the PADC are not simply an echo of the PADG. They represent the current internal FCN/BCN values of that node.

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.

The largest possible credit limit is 0xffff. If an incremental credit grant causes the accumulated count to exceed this value, the max value is used.

One unit of credit represents 64-bytes, so a grant of 4 credits translates to 256-bytes.

7. 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 double the waiting period. A PADC response with the associated sequence number will indicate if the previously granted credits were accumulated or not. If not, a PADG with credits, with 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 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.

8. Other Considerations

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

The sending and receiving of PPPoE control 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 should be used by the receiving node to set

the credit values of its peer.

Berry, Holgate

Expires October 29, 2006

[Page 9]



11. Appendix B: Example Message Formats

A PADR packet with OPTIONAL Credit Tag Type 0x0106:

1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																
Access_Concentrator_mac_addr																																															
Access_Concentrator_mac_addr(c)																Host_mac_addr																															
Host_mac_addr (cont)																																															
ETHER_TYPE = 0x8863																v = 1 t = 1 CODE = 0x19																															
SESSION_ID = 0x1234																LENGTH = 0x0C																															
Tag Type = 0x0101																Tag Length=0x00																															
Tag Type = 0x0106																Tag Length=0x04																															
FCN																BCN																															

1										2										3											
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1
Access_Concentrator_mac_addr																															
Access_Concentrator_mac_addr(c)										Host_mac_addr																					
Host_mac_addr (cont)																															
ETHER_TYPE = 0x8863										v = 1 t = 1										CODE = 0x65											
SESSION_ID = 0x1234										LENGTH = 0x0C																					
Tag Type = 0x0101										Tag Length=0x00																					
Tag Type = 0x0106										Tag Length=0x04																					
FCN										BCN																					

A PADG packet with Credit Tag Type 0x0106:

1																2																3															
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																
Destination_mac_addr																																															
Destination_mac_addr(c)																Source_mac_addr																															
Source mac_addr (cont)																																															
ETHER_TYPE = 0x8863																v = 1 t = 1																CODE = 0x0A															
SESSION_ID = 0x1234																LENGTH = 0x0E																															
Tag Type = 0x0108																Tag Length=0x02																															
Sequence Number																Tag Type = 0x0106																															
Tag Length=0x04																FCN																															
BCN																																															

A PADC packet with Credit Tag Type 0x0106:

	1																2																3															
	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1																
	Destination_mac_addr																																															
	Destination_mac_addr(c)																	Source_mac_addr																														
	Source mac_addr (cont)																																															
	ETHER_TYPE = 0x8863																	v = 1		t = 1		CODE = 0x0B																										
	SESSION_ID = 0x1234																	LENGTH = 0x0E																														
	Tag Type = 0x0108																	Tag Length=0x02																														
	Sequence Number																	Tag Type = 0x0106																														
	Tag Length=0x04																	FCN																														
	BCN																																															

A PADQ packet to query for the link metrics: This is indicated by the Metric Tag Length=0.

[illegible]

A PADQ packet with Metric Tag Type 0x0107:

```

          1              2              3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Access_Concentrator_mac_addr      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Access_Concentrator_mac_addr(c)|                                     Host_mac_addr      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Host_mac_addr (cont)                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   ETHER_TYPE = 0x8863          | v = 1 | t = 1 |   CODE = 0x0C   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   SESSION_ID = 0x1234          |   LENGTH = 0x12                    |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Tag Type = 0x0101            |   Tag Length=0x00                  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Tag Type = 0x0107            |   Tag Length=0x0A                  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Reserved                      |R|   RLQ      |   Resource      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Latency (MS)                  |   Current Datarate (kbps)        |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Maximum Datarate (kbps)      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


While the PPP protocol value is shown (0xc021), the PPP payload is left to the reader. This is a packet from the Host to the Access Concentrator.

[illegible]

12. Acknowledgements

The authors would like to acknowledge the influence and contributions from Billy Moon, Fred Baker, Stan Ratliff and Ed Paradise.

13. Normative References

- [1] Simpson, W., Editor, "The Point-to-Point Protocol (PPP)", STD 51, [RFC 1661](#), July 1994
- [2] Mamakos L., et. al., "A Method for Transmitting PPP Over Ethernet (PPPoE)", [RFC 2516](#), February 1999.

Authors' Address

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

Howard Holgate
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA
email: hholgate@cisco.com

Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology under such rights described in this document or the extent to which any license might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Copyright Statement

Copyright (C) The Internet Society (2006). This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.

