

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 30, 2017

B. Cheng
D. Wiggins
Lincoln Laboratory
L. Berger
LabN Consulting, L.L.C.
October 27, 2016

DLEP DiffServ Aware Credit Windowing Extension
draft-cheng-manet-dlep-da-credit-extension-00

Abstract

This document defines an extension to the DLEP protocol that enables a DiffServ aware credit-windowing scheme for destination-specific flow control.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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 Internet-Draft will expire on April 30, 2017.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Key Words	3
2.	Extension Overview	3
3.	Extension Usage and Identification	4
4.	Extension Messages	4
4.1.	Credit Control Message	4
4.2.	Credit Control Response Message	5
5.	Extension Data Items	5
5.1.	Queue Parameters	5
5.2.	DiffServ Aware (DA) Credit Grant	8
5.3.	DiffServ Aware (DA) Credit Window Status	10
5.4.	DiffServ Aware (DA) Credit Request	11
6.	Compatibility	12
7.	Security Considerations	13
8.	IANA Considerations	13
8.1.	Extension Type Value	13
8.2.	Message Values	13
8.3.	Data Item Values	14
9.	References	14
9.1.	Normative References	14
9.2.	Informative References	14
	Authors' Addresses	15

[1.](#) Introduction

The Dynamic Link Event Protocol (DLEP) is defined in [\[I-D.ietf-manet-dlep\]](#). It provides the exchange of link related control information between DLEP peers. DLEP peers are comprised of a modem and a router. DLEP defines a base set of mechanisms as well as support for possible extensions. This document defines one such extension.

The base DLEP specification does not include any flow control capability. There are various flow control theoretically possible with DLEP. For example, a credit-windowing scheme for destination-specific flow control which provides aggregate flow control for both modem and routers has been proposed in [\[I-D.ietf-manet-credit-window\]](#).

This document defines a DLEP extension which provides flow control for DiffServ [\[RFC2475\]](#) traffic sent from a router to a modem. Flow control is provided for multiple DSCPs (differentiated services codepoint), which are grouped into logical sets of logical queues. The extension defined in this document is referred to as "DiffServ Aware Credit Windowing" or, more simply, the "DA Credit" extension.

This document defines a new DLEP Extension Type Value in [Section 3](#) which is used to indicate the use of the extension. Two new messages are defined in [Section 4](#), and four new DLEP Data Items in [Section 5](#).

1.1. Key Words

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#), [RFC 2119](#) [[RFC2119](#)].

2. Extension Overview

The DA Credit extension can be used to support credit based flow control of traffic sent from a router to a modem. The extension can be used to support DiffServ and non-DiffServ based flow control. Both types of DLEP endpoints, i.e., a router and a modem, negotiate the use of extension during session initialization, see [Section 5.2](#) [[I-D.ietf-manet-dlep](#)].

When the extension is to be used, a modem passes to a router its known DSCPs, if any. DSCPs are grouped by logical queues, each of which are given a logical queue index. The queue index zero (0) is special and is used for any DSCP value, including 0, which is not otherwise identified by the modem. The modem also provides a size, in bytes, of the logical queue for informative and, potential, future uses. Currently, only DSCP to logical queue index value mapping is used in flow control operation.

The DA Credit extension supports credit based flow control on a per MAC address destination, per queue index basis. Modems provide the initial size of the associated "Credit Window", i.e., the amount in octets (bytes) that may be sent by the router to the modem, when a MAC destination becomes reachable and, optionally, when its rate information changes. "Credit Increments", i.e., increases in a Credit Window size, are provided using new "Credit Control" messages. A router provides its view of the Credit Window, which is known as "status", in Destination Up Response and the new "Credit Control Response" messages. Routers can also request credits using the new "Credit Control" message.

Credit information, both grants and status, is provided in new credit grant related data items. Each data item contains a single credit value that applies to one or more queue indexes. Different (grant and status) values for different queue indexes can be provided in a single message by including multiple grant data items. The values indicate a number of octets (bytes), including MAC headers on the router to modem link, that may be sent.

Note that credit information related to different destination MAC addresses is always passed in different DLEP messages.

3. Extension Usage and Identification

The use of the extension defined in this document SHOULD be configurable. To indicate that the DiffServ Aware Credit Windowing Extension is to be used, an implementation MUST include the DiffServ Aware Credit Windowing Type Value in the Extensions Supported Data Item. The Extensions Supported Data Item is sent and processed according to [[I-D.ietf-manet-dlep](#)].

The DiffServ Aware Credit Windowing Extension Type Value is TBA1, see [Section 8](#).

4. Extension Messages

Two new messages are defined by this extension: the Credit Control and the Credit Control Response message. Sending and receiving both message types MUST be supported by any implementation that advertises use of this extension per [Section 3](#).

4.1. Credit Control Message

Credit Control Messages are sent by modems to provide Credit Increments. For messages sent by modems, only one message per MAC address can be outstanding at one time. That is, a modem MUST NOT send a second (or any subsequent) message containing the same MAC Address until a Credit Control Response message is received from its peer router with that MAC address.

Credit Control Messages MAY be sent by routers, e.g., to request credits or provide window status. No specific response message is required from a message transaction perspective.

[TBD: Should anything be said about sending, or limiting, multiple credit requests?]

The Message Type value in the DLEP Message Header is set to TBA2.

The message MUST contain a DLEP MAC Address Data Item.

A message sent by a modem, MUST contain one or more DiffServ Aware Credit Grant data items as defined below in [Section 5.2](#). A router receiving this message MUST respond with a Credit Control Response Message.

A message sent by a router, MUST contain the DiffServ Aware Credit Request data item defined below in [Section 5.4](#). A modem receiving this message MUST provide a Credit Increment for the indicated MAC address and queue indexes via a new Credit Control Message.

Specific processing associated with each Credit data item is provided below.

[4.2.](#) Credit Control Response Message

Credit Control Response Messages are sent by routers to report the current Credit Window for a destination.

The Message Type value in the DLEP Message Header is set to TBA3.

The message MUST contain a DLEP MAC Address Data Item.

A message sent by a router, MUST contain one or more DiffServ Aware Credit Window Status data items as defined below in [Section 5.3](#).

Specific processing associated with the DA Credit Window Status data item is provided below.

[5.](#) Extension Data Items

Four data items are defined by this extension. The Queue Parameters Data Item is used by a modem to provide information on the DSCPs it uses in forwarding. The DA Credit Grant is used by a modem to provide credits to a router. The DA Credit Request is used by a router to request additional credits. The DA Credit Window Status is used to advertise the sender's view of number of available credits for synchronization purposes.

The defined data items and operations are similar to those found in [\[I-D.ietf-manet-credit-window\]](#). One notable difference from this extension is that in this document credits are never provided by the router to the modem.

[5.1.](#) Queue Parameters

The Queue Parameters Data Item is used by a modem to indicate DSCP values that may be independently controlled. This data item MUST be included in a Session Initialization Response Message that also contains the DiffServ Aware Credit Windowing Extension Type Value in the Extensions Supported Data Item. Updates to these parameters MAY be sent by a modem by including the data item in Session Update Messages.

The Queue Parameters Data Item identifies DSCPs based on groups of logical queues. The number of logical queues is variable as is the number of DSCPs associated with each queue. A queue size (in bytes) is provided for informational purposes. An implementation that does not support DSCPs would indicate 1 queue with 0 DSCPs, and the number of bytes that may be in its associated link transmit queue.

The format of the Queue Parameters Data Item is:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Data Item Type                               | Length                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   Num Queues   | Scale |           Reserved           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|Num DSCPs Q0(0)|           Queue Size Q0                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Num DSCPs Q1  |           Queue Size Q1                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Num DSCPs Q2  |           Queue Size Q2                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:                                     ...                                     :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Num DSCPs Qn  |           Queue Size Qn                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| DS Field Q1   | DS Field Q1   | DS Field Q1   | DS Field Q2   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:                                     ...                                     | DS Field Qn |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Data Item Type: TBA4

Length: Variable

Per [[I-D.ietf-manet-dlep](#)] Length is the number of octets in the data item, excluding the Type and Length fields.

Num Queues:

An 8-bit unsigned integer indicating the number of queues represented in the data item. This field **MUST** contain a value of at least one (1). Note that this number is one larger than the largest queue index value included in the data item.

Scale:

An 4-bit unsigned integer indicating the scale used in the Queue Size fields. The valid values are:

Value	Scale

0	B - Bytes (Octets)
1	KB - Kilobytes (B/1024)
2	MB - Megabytes (KB/1024)
3	GB - Gigabytes (MB/1024)

Reserved:

MUST be set to zero by the sender (a modem) and ignored by the receiver (a router).

Num DSCPs Qn:

An 8-bit unsigned integer indicating the number of DSCPs associated with the indexed queue. Other than the special case covered in the next paragraph, this field MUST contain a value of at least one (1). Queue indexes start at zero (0) and the maximum queue index "Qn" is one less than the value carried in the Num Queues field. Queue indexes are implicit in the position in the data item.

Queue index zero "Q0" is a special case. It is used for any traffic that does not carry a DSCP value represented in the data item. Therefore the value of the Queue index zero field, "Num DSCPs Q0", field MUST be zero (0).

Queue Size Qn:

A 24-bit unsigned integer representing the size, in the octet scale indicated by the Scale field, of the queue supporting traffic with the DSCPs associated with the queue index.

DS Field Qn:

The data item contains a sequence of 8 bit DS Fields. The position in the sequence identifies the associated queue index. The number of DS Fields present should equal the sum of all Num DSCPs field values.

The DS Field structure is the same as [[RFC2474](#)].


```

      0   1   2   3   4   5   6   7
+---+---+---+---+---+---+---+---+
|           DSCP           | CU |
+---+---+---+---+---+---+---+---+

```

DSCP: differentiated services codepoint

CU: currently unused, MUST be zero

5.2. DiffServ Aware (DA) Credit Grant

The DiffServ Aware, or DA, Credit Grant data item is used by a modem to provide credits to a router. One or more DA Credit Grant data items MAY be carried in the DLEP Destination Up, Destination Announce Response, Destination Update, and Credit Control messages. Multiple DA Credit Grant data items in a single message are used to indicated different credit values for different logical queues.

In Destination type messages, this data item provides the total number of octets available in the Credit Window to the destination indicated in the message for the specified logical queues. In the Credit Control message, this data item provides an additional number of octets to be added to the Credit Window to the destination indicated in the message for the specified logical queues.

Logical queues are identified using a Queue Index as defined above in [Section 5.1](#). Multiple Queue Indexes MAY be present to allow for the case where same credit information applies to multiple queues. As mentioned above, multiple DA Credit Grant Data Items MAY be present to provide different queue-specific credit information in one message. The special Queue Index value of 255 is used to indicate that the credit information applies to all queues.

The format of the DA Credit Grant Data Item is:

```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Data Item Type           | Length           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                           Credit Value       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                           Credit value      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Queue Index | ... |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
: ... | Queue Index |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


Data Item Type: TBA5

Length: Variable

Per [[I-D.ietf-manet-dlep](#)], Length is the number of octets in the data item, excluding the Type and Length fields. It will equal 8 plus the number of Queue Index fields carried in the data item and MUST be at least 9.

Credit Value:

A 64-bit unsigned integer representing the credits, in octets, to be applied to the Credit Window. This value includes MAC headers as seen on the link between the modem and router.

Queue Index:

One or more 8-bit fields used to indicate a queue index defined by a Queue Parameters Data Item. The special value of 255 indicates that the information in the data item applies to all queue indexes.

Receive processing of this data item is based on the message in which it is carried. When this data item is received in a Destination type message, the Credit Window of the indicated destination MAC address and indicated queue indexes MUST be set to the value contained in the Credit Value field.

When this data item is received in a Credit Control message, the Credit Window of the indicated destination MAC address and indicated queue indexes MUST be increased by the value contained in the Credit Value field. If the increase results in a window overflow, i.e., the Credit Window resulting after the increase is smaller than the original Credit Window, the Credit Window must be set to its maximum (0xFFFFFFFFFFFFFFFF).

Independent of the received message, the receiving router MUST send a DA Credit Window Status data item or items reflecting the resulting Credit Window value of each modified queue index. When the Credit Grant data item is received in a Destination Up message, the DA Credit Window Status data item(s) MUST be sent in the corresponding Destination Up Response message. In all other cases, the a Credit Control message MUST be sent.

5.3. DiffServ Aware (DA) Credit Window Status

The DiffServ Aware, or DA, Credit Window Status data item is used by a router to report the current Credit Window to its peer modem. One or more DA Credit Window Status data items MAY be carried in a Destination Up Response message or a Credit Control Response message. As discussed above, the Destination Up Response message is used when the data item is sent in response to a Destination Up message, and the Credit Control Response message is sent in response to a Credit Control message. Multiple DA Credit Window Status data items in a single message are used to indicate different credit window values for different logical queues.

Similar to the DA Credit Grant, logical queues are identified using a Queue Index as defined above in [Section 5.1](#). Multiple Queue Indexes MAY be present to allow for the case where same credit information applies to multiple queues. Multiple DA Credit Window Status Data Items are used to provide different queue-specific credit window information in one message. The special Queue Index value of 255 is used to indicate that the Credit Window information applies to all queues.

The format of the DA Credit Window Status Data Item is:

```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Data Item Type                               | Length                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Credit Window                               :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:                               Credit Window                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Queue Index |                               ...                               :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:                               ...                               | Queue Index |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Data Item Type: TBA6

Length: Variable

Per [[I-D.ietf-manet-dlep](#)] Length is the number of octets in the data item, excluding the Type and Length fields. It will equal 8 plus the number of Queue Index fields carried in the data item and MUST be at least 9.

Credit Window:

A 64-bit unsigned integer, indicating the current number of credits, in octets, available for the router to send to the modem. This is referred to as the Modem Receive Window in [\[I-D.ietf-manet-credit-window\]](#).

Queue Index:

One or more 8-bit fields used to indicate a queue index defined by a Queue Parameters Data Item. The special value of 255 indicates that the information in the data item applies to all queue indexes.

The receiving modem SHOULD check the received Credit Window value against the outstanding credits available at the time the last Credit Increment associated with the indicated MAC address and Queue Indexes were sent. If the values significantly differ, i.e., greater than can be accounted for based on observed data frames, then the modem SHOULD send a Destination Update message carrying a DA Credit Grant data item to reset the associated Credit Window(s) to the data item indicated value. Multiple values and queue indexes SHOULD be combined into a single Destination Update Control message when possible. Alternatively, and also in cases where there are small differences, the modem MAY adjust the values sent in DA Credit Grant data items to account for the reported Credit Window.

5.4. DiffServ Aware (DA) Credit Request

The DiffServ Aware, or DA, Credit Request Data Item data item is used by a router to request additional credits for a specific destination and Queue Index associated Credit window. DA Credit Request data items are carried in Credit Control messages, and only one DA Credit Request data item SHOULD be present in a message.

Logical queues are identified using a Queue Index as defined above in [Section 5.1](#). Multiple Queue Indexes MAY be present to allow for the case where the credit request applies to multiple queues. The special Queue Index value of 255 is used to indicate that a credit request is being made across all queues.

The format of the DA Credit Request Data Item is:


```

      0              1              2              3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Data Item Type                               | Length                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Queue Index |                               ...                        :
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
:                               ...                        | Queue Index |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

[LB note: a list of Queue Indexes is now supported as is special value 255.]

Data Item Type: TBA7

Length: Variable

Per [[I-D.ietf-manet-dlep](#)] Length is the number of octets in the data item, excluding the Type and Length fields. It will equal the number of Queue Index fields carried in the data item and MUST be at least 1.

Queue Index:

One or more 8-bit fields used to indicate a queue index defined by a Queue Parameters Data Item. The special value of 255 indicates that the request applies to all queue indexes.

A modem receiving this data item MUST provide a Credit Increment for the indicated MAC address and queue indexes via a DA Credit Grant carried in a new Credit Control Message. Multiple values and queue indexes SHOULD be combined into a single Credit Control message when possible.

6. Compatibility

Sessions established with both peers identified as supporting the DiffServ Aware Credit Windowing Extension Type, see [Section 3](#), SHOULD NOT use the [[I-D.ietf-manet-credit-window](#)] defined Credit data items. If a node supporting the extension defined in this document, receives a [[I-D.ietf-manet-credit-window](#)] defined data item, the recipient MUST treat the received credit information as applying to Queue Index zero (0).

7. Security Considerations

The extension introduces a DiffServ awareness to the mechanisms defined in [[I-D.ietf-manet-credit-window](#)]. The extension does not inherently introduce any additional threats above those documented in [[I-D.ietf-manet-dlep](#)]. The approach taken to Security in that document and [[I-D.ietf-manet-credit-window](#)] apply equally when running the extension defined in this document.

8. IANA Considerations

This document requests the assignment of 5 values by IANA. All assignments are to registries defined by [[I-D.ietf-manet-dlep](#)].

8.1. Extension Type Value

This document requests 1 new assignment to the DLEP Extensions Registry named "Extension Tyoe Values" in the range with the "Specification Required" policy. The requested value is as follows:

Code	Description
TBA1	DiffServ Aware Credit Windowing

Table 1: Requested Extension Type Value

8.2. Message Values

This document requests 2 new assignments to the DLEP Message Registry named "Message Values" in the range with the "Specification Required" policy. The requested values are as follows:

Type Code	Description
TBA2	Credit Control
TBA3	Credit Control Response

Table 2: Requested Message Values

8.3. Data Item Values

This document requests 4 new assignments to the DLEP Data Item Registry named "Data Item Values" in the range with the "Specification Required" policy. The requested values are as follows:

Type Code	Description
TBA4	Queue Parameters
TBA5	DA Credit Grant
TBA6	DA Credit Window Status
TBA7	DA Credit Request

Table 3: Requested Data Item Values

9. References

9.1. Normative References

- [I-D.ietf-manet-dlep]
 Ratliff, S., Jury, S., Satterwhite, D., Taylor, R., and (. (Unknown), "Dynamic Link Exchange Protocol (DLEP)", [draft-ietf-manet-dlep-24](#) (work in progress), July 2016.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.

9.2. Informative References

- [I-D.ietf-manet-credit-window]
 Ratliff, S., "Credit Windowing extension for DLEP", [draft-ietf-manet-credit-window-04](#) (work in progress), April 2016.
- [RFC2474] Nichols, K., Blake, S., Baker, F., and D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", [RFC 2474](#), DOI 10.17487/RFC2474, December 1998, <<http://www.rfc-editor.org/info/rfc2474>>.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", [RFC 2475](https://www.rfc-editor.org/info/rfc2475), DOI 10.17487/RFC2475, December 1998, <<http://www.rfc-editor.org/info/rfc2475>>.

Authors' Addresses

Bow-Nan Cheng
Lincoln Laboratory
Massachusetts Institute of Technology
244 Wood Street
Lexington, MA 02420-9108

Email: bcheng@ll.mit.edu

David Wiggins
Lincoln Laboratory
Massachusetts Institute of Technology
244 Wood Street
Lexington, MA 02420-9108

Email: David.Wiggins@ll.mit.edu

Lou Berger
LabN Consulting, L.L.C.

Email: lberger@labn.net

