Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: November 2, 2018

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

May 1, 2018

DLEP Credit-Based Flow Control Messages and Data Items
draft-ietf-manet-dlep-credit-flow-control-00

Abstract

This document defines new DLEP protocol Data Items that are used to support credit-based flow control. The Data Items enable separate but related functions: traffic classification and credit window control. Traffic classification information is used to identify traffic flows based on frame/packet content such as destination address. Credit window control is used to regulate when data may be sent to an associated virtual or physical queue. The Data Items are defined in an extensible and reusable fashion. Their use will be mandated in other documents defining specific DLEP extensions. This document also introduces DLEP sub-data items.

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 https://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 November 2, 2018.

Copyright Notice

Copyright (c) 2018 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

(https://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 4
2. Traffic Classification 4
2.1. Traffic Classification Data Item 5
2.1.1. Traffic Classification Sub Data Item 6
2.2. DiffServ Traffic Classification Sub Data Item 7
2.2.1. Router Receive Processing 8
2.3. Ethernet Traffic Classification Sub Data Item 8
2.3.1. Router Receive Processing 10
3. Credit Window Control 10
3.1. Data Plane Considerations 12
3.2. Credit Window Messages 12
3.2.1. Credit Control Message 12
3.2.2. Credit Control Response Message 13
3.3. Credit Window Control Data Items 13
3.3.1. Credit Window Initialization 14
3.3.2. Credit Window Associate 16
3.3.3. Credit Window Grant 17
3.3.4. Credit Window Status 18
3.3.5. Credit Window Request 20
3.4. Management Considerations 21
4. Compatibility 21
5. Security Considerations 21
6. IANA Considerations 21
6.1. Message Values 21
6.2. Data Item Values 22
6.3. DLEP Sub Data Item Registry 22
7. References 23
7.1. Normative References 23
7.2. Informative References 23
Appendix A. Acknowledgments 24
Authors' Addresses 24

1. Introduction

The Dynamic Link Exchange Protocol (DLEP) is defined in [RFC8175]. 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. DLEP defines Data Items which are sets of information that can be reused in DLEP messaging. The base DLEP specification does not include any flow identification beyond DLEP endpoints or flow control capability. There are various flow control techniques theoretically possible with DLEP. For example, a credit-window 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], and a control plane pause based mechanism is defined in [I-D.ietf-manet-dlep-pause-extension].

This document defines DLEP Data Items and Messages which provide flow identification, and a flow control mechanism for traffic sent from a router to a modem. Flow control is provided using one or more logical "Credit Windows", each of which will typically be supported by an associated virtual or physical queue. Traffic sent by a router will use traffic flow classification information provided by the modem to identify which traffic is associated with each credit window. In this case, a flow is identified based on information found in a data plane header and one or more matches are associated with a single flow. (For general background on traffic classification see [RFC2475] Section 2.3.) Credit windows may be shared or dedicated on a per flow basis. The Data Items are structured to allow for reuse of the defined traffic classification information with non-credit window applications as well as reuse of the credit window based flow control with different traffic classification techniques.

This document defines traffic classification based on a DLEP destination and flows identified by either DiffServ [RFC2475] DSCPs (differentiated services codepoints) or IEEE 802.1Q [IEEE.802.1Q_2014] Ethernet Priority Code Points (PCP). The defined mechanism allows for credit windows to be shared across traffic sent to multiple DLEP destinations and flows, or used exclusively for traffic sent to a particular destination and/or flow. The extension also supports the "wildcard" matching of any flow (DSCP or PCP). Traffic classification information is provided such that it can be readily extended to support other traffic classification techniques, or be used by non-credit window related extensions, such as [I-D.ietf-manet-dlep-pause-extension] or even 5-tuple IP flows.

Note that this document defines common Messages, Data Items and mechanisms that are reusable. They are expected to be required by DLEP extensions defined in other documents such as found in [I-D.ietf-manet-dlep-da-credit-extension].

This document defines support for traffic classification using a single new Data Item in Section 2.1 for general support and two new

sub Data Items are defined to support identification of flows based on DSCPs and PCPs. The document supports credit window control by introducing two new DLEP messages in Section 3.2, and five new DLEP Data Items in Section 3.3.

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 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Traffic Classification

The Traffic Classification Data Item is used to represent a list of flows that may be used at the same time for traffic sent from a router to a modem. The data plane information used to identify each flow is represented in a separate sub Data Item. The Data Item and Sub Data Item structure is intended to be independent of any specific usage of the flow identification, e.g., flow control. The Sub Data Item structure is also intended to allow for future traffic classification types, e.g., 5-tuple flows. While the structure of the Data Items is extensible, actual flow information is expected to be used in an extension dependent manner. Support for DSCP and PCP-based flows are defined via individual sub Data Items below. Other types of flow identification, e.g., based on IP protocol and ports, may be defined in the future via new sub Data Items.

The list of flows contained in the Data Item can be used per sender or shared across multiple senders. Each list of flows is identified using a "Traffic Classification Identifier" or "TID" and is expected to represent a valid combination of data plane identifiers that may be used at the same time. Each flow is identified via a "Flow Identifier" or "FID". Each FID is defined in a sub Data Item which carries the data plane identifier or identifiers used to associate traffic with the flow. A DLEP destination address is also needed to complete traffic classification information used in extensions such as flow control. This information is expected to be provided in an extension specific manner. For example, this address can be provided by a modem when it identifies the traffic classification set in a Destination Up Message using the Credit Window Associate Data Item defined in Section 3.3.2 .

2.1. Traffic Classification Data Item

This sections defines the Traffic Classification Data Item. This Data Item is used by a modem to provide a router with traffic classification information. When an extension requires use of this Data Item the Traffic Classification Data Item SHOULD be included by a modem in any Session Initialization Response Message that also indicates support for an extension that requires support for the credit window control mechanisms defined in this document, e.g., see [I-D.ietf-manet-dlep-da-credit-extension]. Updates to previously provided traffic classifications or new traffic classifications MAY be sent by a modem by including the Data Item in Session Update Messages. More than one Data Item MAY be included in a message to provide information on multiple traffic classifiers.

The set of traffic classification information provided in the data item is identified using a Traffic Classification Identifier, or TID. The actual data plane related information used in traffic classification is provided in a variable list of Traffic Classification Sub Data Items.

The format of the Traffic Classification Data Item is:

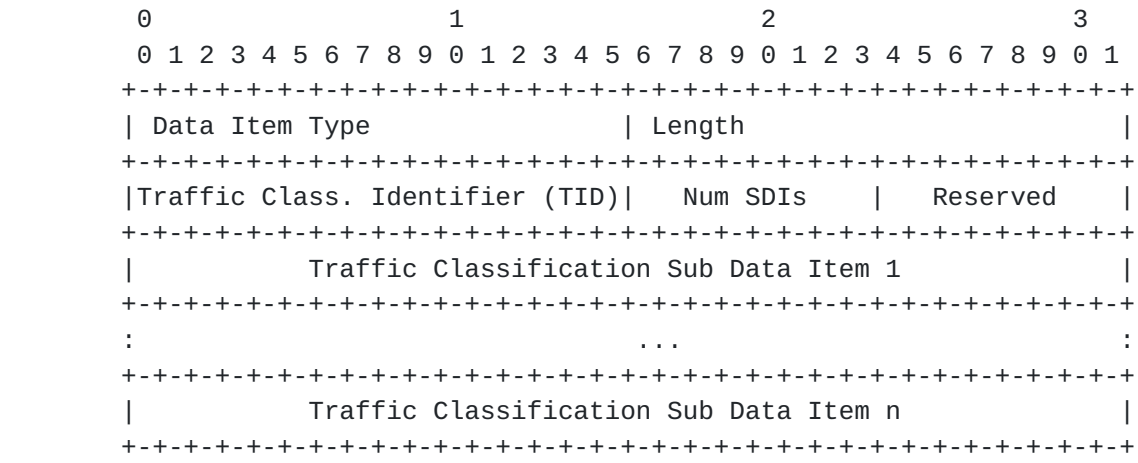

Data Item Type: TBA1

Length: Variable

Per [RFC8175] Length is the number of octets in the Data Item, excluding the Type and Length fields.

Traffic Classification Identifier (TID):

A 16-bit unsigned integer identifying a traffic classification set. There is no restriction on values used by a modem, and there is no requirement for sequential or ordered values.

Num SDIs:

An 8-bit unsigned integer indicating the number of Traffic Classification Sub Data Items included in the Data Item. A value of zero (0) is allowed and indicates that no traffic should be matched against this TID.

Reserved:

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

Traffic Classification Sub Data Item:

Zero or more Traffic Classification Sub Data Items of the format defined below MAY be included. The number MUST match the value carried in the Num SDIs field.

A router receiving the Traffic Classification Data Item MUST locate the traffic classification information that is associated with the TID indicated in each received Data Item. If no associated traffic classification information is found, the router MUST initialize a new information set using the values carried in the Data Item. When associated traffic classification information is found, the router MUST update the information using the values carried in the Data Item. In both cases, a router MUST also ensure that any data plane state, e.g., see Section 3.1, that is associated with the TID is updated as needed.

2.1.1. Traffic Classification Sub Data Item

All Traffic Classification Sub Data Items share a common format that is patterned after the standard DLEP Data Item format, see [RFC8175] Section 11.3. There is no requirement on, or meaning to sub Data Item ordering. Any errors or inconsistencies encountered in parsing sub Data Items are handled in the same fashion as any other Data Item parsing error encountered in DLEP.

The format of the Traffic Classification Sub Data Item is:

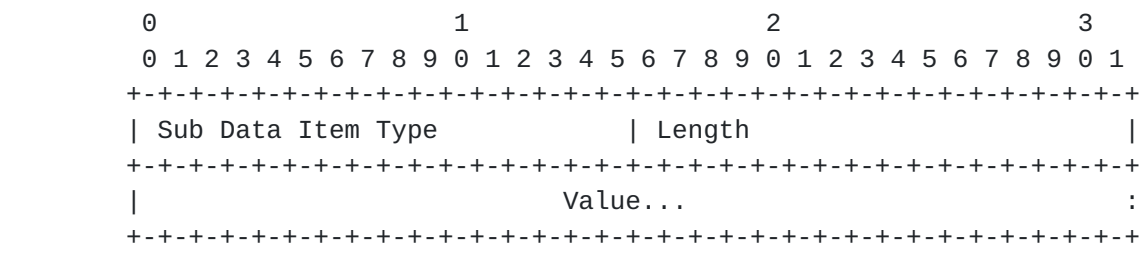

Sub Data Item Type:

A 16-bit unsigned integer that indicates the type and corresponding format of the Data Item's Value field. Sub Data Item Types are managed according to the IANA registry described in Section 6.3.

Length: Variable

Copying [RFC8175], Length is a 16-bit unsigned integer that is the number of octets in the sub Data Item, excluding the Type and Length fields.

2.2. DiffServ Traffic Classification Sub Data Item

The DiffServ Traffic Classification Sub Data Item is used to identify the set of DSCPs that should be treated as a single flow, i.e., receive the same traffic treatment. DSCPs are identified in a list of DiffServ fields. An implementation that does not support DSCPs and wants the same traffic treatment for all traffic to a destination or destinations would indicate 0 DSCPs.

The format of the DiffServ Traffic Classification Sub Data Item is:

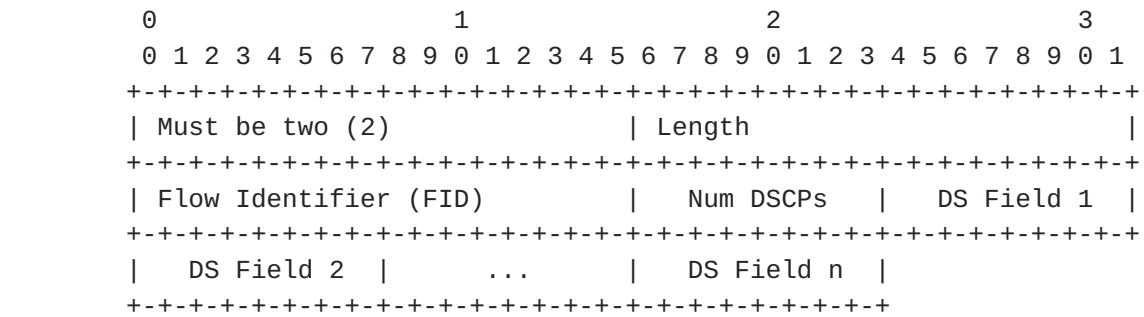

Length: Variable

Length is defined above. For this Sub Data Item, it is equal to three (3) plus the value of the Num DSCPs field.

Flow Identifier (FID):

A 16-bit unsigned integer representing the data plane information carried in the sub Data Item that is to be used in identifying a flow. The value of 0xFFFF is reserved and MUST NOT be used in this field.

Num DSCPs:

An 8-bit unsigned integer indicating the number of DSCPs carried in the sub Data Item. A zero (0) indicates a (wildcard) match against any DSCP value.

DS Field:

Each DS Field is an 8-bit whose definition is the same as [RFC2474].

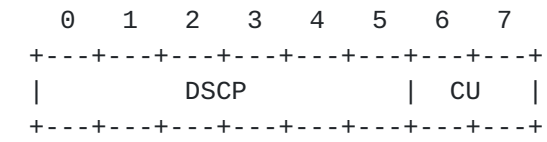

DSCP: differentiated services codepoint
CU: currently unused, MUST be zero

2.2.1. Router Receive Processing

A router receiving the Traffic Classification Sub Data Item MUST validate the information on receipt, prior to using the carried information, including potentially updating the data behavior as determined by the extension requiring the use of the Sub Data Item. Validation failures MUST be treated as an error as described above.

Once validated, the receiver MUST ensure that each DS Field value is listed only once across the whole Traffic Classification Data Item. Note, this check is across the Data Item and not the individual sub Data Item. If the same DS Field value is listed more than once within the same Traffic Classification Data Item, the Data Item MUST be treated as an error as described above.

2.3. Ethernet Traffic Classification Sub Data Item

The Ethernet Traffic Classification Sub Data Item is used to identify the VLAN and PCPs that should be treated as a single flow, i.e., receive the same traffic treatment. Ethernet Priority Code Point support is defined as part of the IEEE 802.1Q [IEEE.802.1Q_2014] tag format and includes a 3 bit "PCP" field. The tag format also includes a 12 bit VLAN identifier (VID) field. PCPs are identified in a list of priority fields. An implementation that does not

support PCPs and wants the same traffic treatment for all traffic to a destination or destinations would indicate 0 PCPs. Such an implementation could identify a VLAN to use per destination.

The format of the Ethernet Traffic Classification Sub Data Item is:

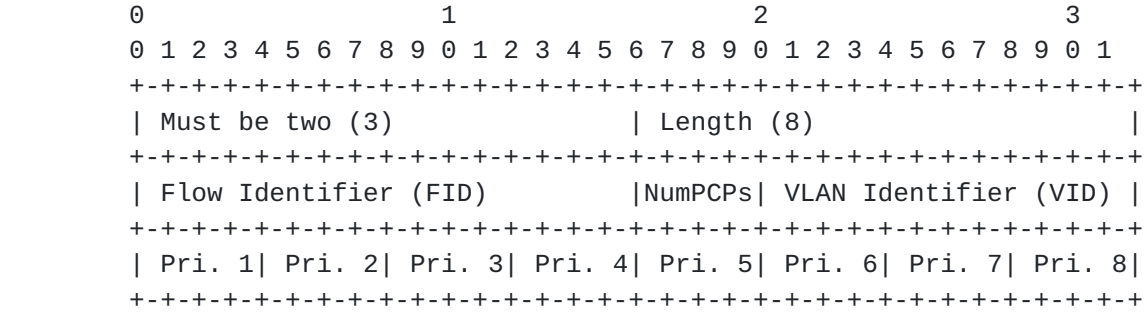

Length: Variable

Length is defined above. For this Sub Data Item, it is equal to eight (8).

Flow Identifier (FID):

A 16-bit unsigned integer representing the data plane information carried in the sub Data Item that is to be used in identifying a flow. The value of 0xFFFF is reserved and MUST NOT be used in this field.

Num PCPs:

A 16-bit unsigned integer representing the data plane information carried in the sub Data Item that is to be used in identifying a flow.

VLAN identifier (VID):

A 12-bit unsigned integer field indicating the VLAN to be used in traffic classification. A value of zero (0) indicates that the VID is to be ignored and any VID is to be accepted during traffic classification.

Priority:

Each Priority Field is an 4-bit whose definition includes the PCP field defined in [IEEE.802.1Q_2014]

```
0   1   2   3
+---+---+---+---+
|    PCP    |DEI|
+---+---+---+---+

PCP: Priority code point
DEI: currently unused, MUST be zero
```

2.3.1. Router Receive Processing

A router receiving the Traffic Classification Sub Data Item MUST validate the information on receipt, prior to the using the carried information, including potentially updating the data behavior as determined by the extension requiring the use of the Sub Data Item. Validation failures MUST be treated as an error as described above.

Once validated, the receiver MUST ensure that each Priority Field value is listed only once across the whole Traffic Classification Data Item. Note, this check is across the Data Item and not the individual sub Data Item. If the same Priority Field value is listed more than once within the same Traffic Classification Data Item, the Data Item MUST be treated as an error as described above.

3. Credit Window Control

This section defines additions to DLEP used in credit based flow control. In addition to the Traffic Classification Data Item, two new messages and five Data Items are defined to support credit window control. The use of credit window control impacts the data plane.

The credit window control mechanisms defined in this document support credit based flow control of traffic sent from a router to a modem. The mapping of specific flows of traffic to a particular credit window is based on the Traffic Classification Data Item defined in Section 2.1. Both types of DLEP endpoints, i.e., a router and a modem, negotiate the use of extension during session initialization, e.g., see [I-D.ietf-manet-dlep-da-credit-extension]. When using credit windows, data traffic is only allowed to be sent by the router to the modem when there are credits available.

Credits are managed on a per logical "Credit Windows" basis. Each credit window can be thought of as corresponding to a queue within a modem. Credit windows may be shared across, or dedicated to, destinations and data plane identifiers, e.g., DSCPs, at a granularity that is appropriate for a modem's implementation and its attached transmission technology. As defined below, there is a direct one-for-one mapping of credit windows to flows as identified by FIDs carried within the Traffic Classification Data Item. Modems

pass to the router information on their credit windows and FIDs prior to a router being able to send data when an extension requiring the use of credit window control is used. In addition to the traffic classification information associated with an FID, routers provide an initial credit window size, as well as the maximum size of the logical queue associated with each credit window. The maximum size is included for informative and potential future uses.

Modems provide an initial credit window size at the time of "Credit Window Initialization". Such initialization can take place during session initiation or any point thereafter. It can also take place when rate information changes. Additional "Credit Grants", i.e., increments to Credit Window size, are provided using a Destination Up or the new "Credit Control" Message. 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.

When modems provide credits to a router, they will need to take into account any overhead of their attached transmission technology and map it into the credit semantics defined in this document. In particular, the credit window is defined below to include per frame (packet) MAC headers, and this may not match the actual overhead of the modem attached transmission technology. In that case a direct mapping, or an approximation will need to be made by the modem to provide appropriate credit values.

Actual flows of traffic are mapped to credit windows based on flow identification information provided by modems in the Traffic Classification Data item defined in Section 2. This data item supports traffic classification on a per destination or more fine grain level. Routers use the combination of the DLEP identified destination and flow information associated with a credit window in order to match traffic they send to specific credit windows.

When a destination becomes reachable, a modem "Associates" (identifies) the appropriate traffic classification information via the TID to be used for traffic sent by the router to that destination. As defined, each credit window has a corresponding FID. This means that the use of FIDs, TIDs and the association of a TID to a DLEP destination enables a modem to share or dedicate resources as needed to match the specifics of its implementation and its attached transmission technology.

The defined credit window control has similar objectives as the control found in [I-D.ietf-manet-credit-window]. One notable difference from that credit control is that in this document, credits are never provided by the router to the modem.

3.1. Data Plane Considerations

When credit windowing is used, a router MUST NOT send data traffic to a modem for forwarding when there are no credits available in the associated Credit Window. This document defines credit windows in octets. A credit window value MUST be larger than the number of octets contained in a packet, including any MAC headers used between the router and the modem, in order for the router to send the packet to a modem for forwarding. The credit window is decremented by the number of sent octets.

A router MUST identify the credit window associated with traffic sent to a modem based on the traffic classification information provided in the Data Items defined in this document. Note that routers will typically view a DLEP destination as the next hop MAC address.

3.2. Credit Window Messages

Two new messages are defined in support for credit window control: the Credit Control and the Credit Control Response Message. Sending and receiving both message types is REQUIRED to support the credit window control defined in this document.

3.2.1. Credit Control Message

Credit Control Messages are sent by modems and routers. Each sender is only permitted to have one message outstanding at one time. That is, a sender (i.e., modem or router) MUST NOT send a second or any subsequent Credit Control Message until a Credit Control Response Message is received from its peer (i.e., router or modem).

Credit Control Messages are sent by modems to provide credit window increases. Modems send credit increases when there is transmission or local queue availability that exceeds the credit window value previous provided to the router. Modems will need to balance the load generated by sending and processing frequent credit window increases against a router having data traffic available to send, but no credits available.

Credit Control Messages MAY be sent by routers to request credits and provide window status. Routers will need to balance the load generated by sending and processing frequent credit window requests against a having data traffic available to send, but no credits available.

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

A message sent by a modem, MUST contain one or more Credit Window Grant Data Items as defined below in Section 3.3.3. A router receiving this message MUST respond with a Credit Control Response Message.

A message sent by a router, MUST contain one or more Credit Window Request Data Items defined below in Section 3.3.5 and SHOULD contain a Credit Window Status Data Item, defined in Section 3.3.4, corresponding to each credit window request. A modem receiving this message MUST respond with a Credit Control Response Message based on the received message and Data Item and the processing defined below, which will typically result in credit window increments being provided.

Specific processing associated with each Credit Data Item is provided below.

3.2.2. Credit Control Response Message

Credit Control Response Messages are sent by routers to report the current Credit Window for a destination. A message sent by a router, MUST contain one or more Credit Window Status Data Items as defined below in Section 3.3.4. Specific receive processing associated with the Credit Window Status Data Item is provided below.

Credit Control Response Messages sent by modems MUST contain one or more Credit Window Grant Data Items. A Data Item for every Credit Window Request Data Item contained in the corresponding Credit Control Response Message received by the modem MUST be included. Each Credit Grant Data Item MAY provide zero or more additional credits based on the modem's transmission or local queue availability. Specific receive processing associated with each Grant Data Item is provided below.

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

3.3. Credit Window Control Data Items

Five new Data Items are defined to support credit window control. The Credit Window Initialization Data Item is used by a modem to identify a credit window and set its size. The Credit Window Association Data Item is used by a modem to identify which traffic classification identifiers (flows) should be used when sending traffic to a particular DLEP identified destination. The Credit Window Grant is used by a modem to provide additional credits to a router. The Credit Request is used by a router to request additional credits. The Credit Window Status is used to advertise the sender's

view of number of available credits for state synchronization purposes.

Any errors or inconsistencies encountered in parsing Data Items are handled in the same fashion as any other data item parsing error encountered in DLEP, see [RFC8175]. In particular, the node parsing the Data Item MUST terminate the session with a Status Data Item indicating Invalid Data.

3.3.1. Credit Window Initialization

The Credit Window Initialization Data Item is used by a modem to identify a credit window and set its size. This Data Item SHOULD be included in any Session Initialization Response Message that also indicates support for an extension that requires support for the credit window control mechanisms defined in this document, e.g., see [I-D.ietf-manet-dlep-da-credit-extension]. Updates to previously identified credit windows or new credit windows MAY be sent by a modem by including the Data Item in Session Update Messages. More than one data item MAY be included in a message to provide information on multiple credit windows.

The Credit Window Initialization Data Item identifies a credit window using a Flow Identifier, or FID. It also provides the size of the identified credit window. Finally, a queue size (in bytes) is provided for informational purposes. Note that to be used, a FID must be defined within a Traffic Classification Data Item and the associated TID must be provided via a Credit Window Association Data Item.

The format of the Credit Window Initialization Data Item is:

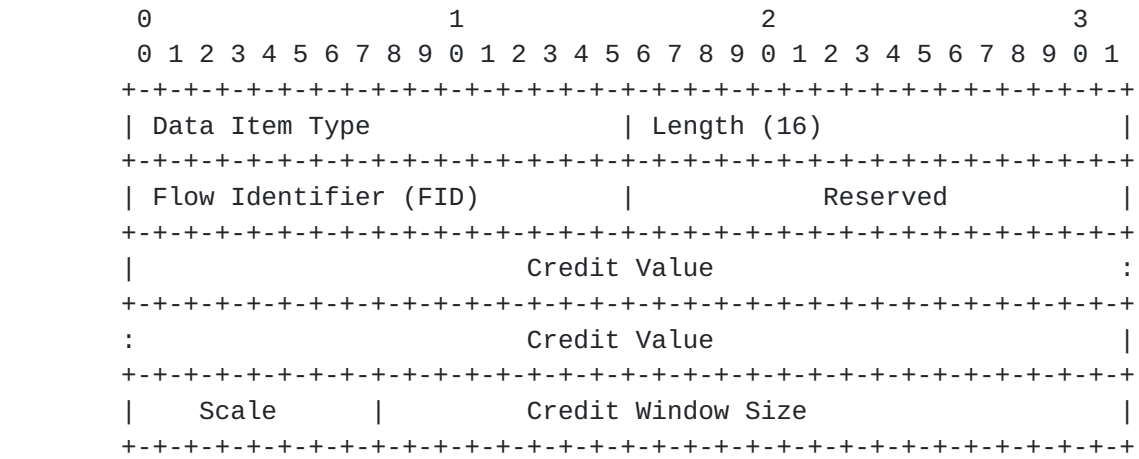

Data Item Type: TBA4

Length: 16

Per [RFC8175] Length is the number of octets in the Data Item. It MUST be equal to sixteen (16).

Flow Identifier (FID):

A flow identifier as defined by the Traffic Classification Data Item. The FID also uniquely identifies a credit window.

Reserved:

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

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.

Scale:

An 8-bit unsigned integer indicating the scale used in the Credit Window 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)

Credit Window Size:

A 24-bit unsigned integer representing the maximum size, in the octet scale indicated by the Scale field, of the associated credit window.

A router that receives a Credit Window Initialization Data Item MUST ensure that the FID field value has been provided by the modem in a Traffic Classification Data Item carried in either the current or previous message. If the FID cannot be found the router SHOULD report or log this information. Note that no traffic will be associated with the credit window in this case. After FID validation, the router MUST locate the credit window that is associated with the FID indicated in each received Data Item. If no associated credit window is found, the router MUST initialize a new

credit window using the values carried in the Data Item. When an associated credit window is found, the router MUST update the credit window and associated data plane state using the values carried in the Data Item. It is worth noting, that such updates can result in a credit window size being reduced, for example, due to a transmission rate change on the modem.

3.3.2. Credit Window Associate

The Credit Window Associate Data Item is used by a modem to associate traffic classification information with a destination. The traffic classification information is identified using a TID value that has previously been sent by the modem or is listed in a Traffic Classification Data Item carried in the same message as the Data Item.

A single Credit Window Associate Data Item MUST be included in all Destination Up and Destination Update Messages sent by a modem when the credit window control defined in this document is used. Note that a TID will not be used unless it is listed in a Credit Window Associate Data Item.

The format of the Credit Window Associate Data Item is:

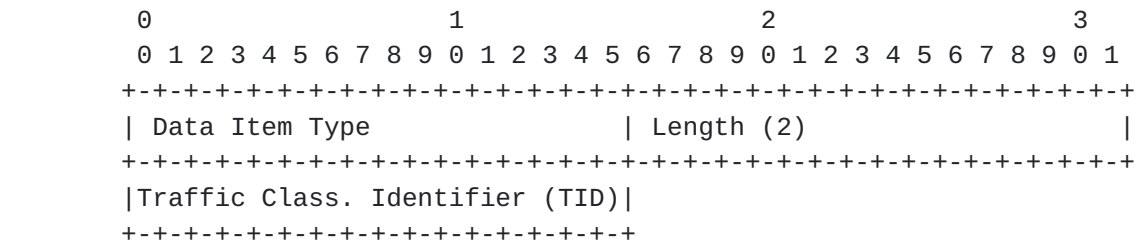

Data Item Type: TBA5

Length: 2

Per [RFC8175] Length is the number of octets in the Data Item. It MUST be equal to two (2).

Traffic Classification Identifier (TID):

A 16-bit unsigned integer identifying a traffic classification set that has been identified in a Traffic Classification Data Item, see Section 2.1.

A router that receives the Credit Window Associate Data Item MUST locate the traffic classification information indicated by the received TID. If no corresponding information can be located, the

Data Item MUST be treated as an error as described above. Once the traffic classification information is located, it MUST be associated with the destination and the router MUST ensure that any data plane state, see Section 3.1, that is associated with the TID and its corresponding FIDs is updated as needed.

3.3.3. Credit Window Grant

The Credit Window Grant Data Item is used by a modem to provide credits to a router. One or more Credit Window Grant Data Items MAY be carried in the DLEP Destination Up, Destination Announce Response, Destination Update, Credit Control Messages, and Credit Control Response Messages. Multiple Credit Window Grant Data Items in a single message are used to indicate different credit values for different credit windows. In all message types, this Data Item provides an additional number of octets to be added to the indicated credit window. Credit windows are identified using FID values that have been previously been sent by the modem or are listed in a Credit Window Initialization Data Item carried in the same messages as the Data Item.

The format of the Credit Window Grant Data Item is:

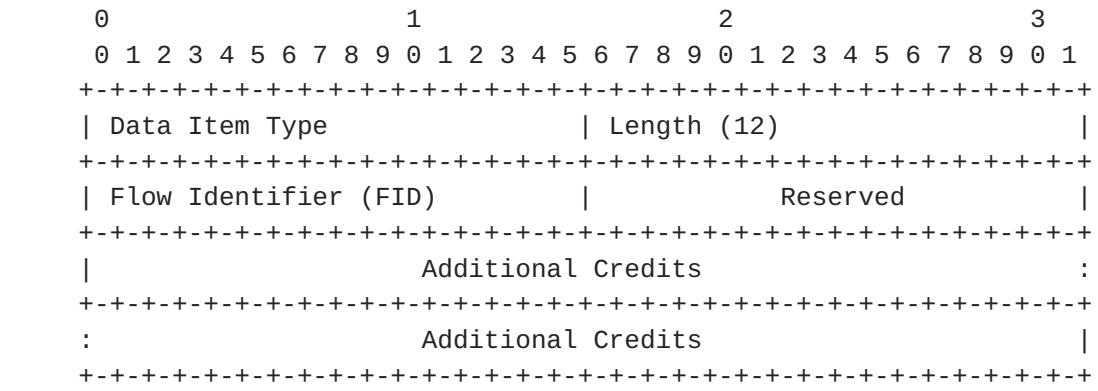

Data Item Type: TBA6

Length: 12

Per [RFC8175], Length is the number of octets in the Data Item. It MUST be equal to twelve (12).

Flow Identifier (FID):

A flow identifier as defined by the Traffic Classification Data Item. The FID also uniquely identifies a credit window.

Additional Credit:

A 64-bit unsigned integer representing the credits, in octets, to be added to the Credit Window. This value includes MAC headers as seen on the link between the modem and router. A value of zero indicates that no additional credits are being provided.

When receiving this Data Item, a router MUST identify the credit window indicated by the FID. If the FID is not known to the router, it SHOULD report or log this information and discard the Data Item. It is important to note that while this Data Item can be received in a destination specific message, credit windows are managed independently from the destination identified in the message carrying this Data Item, and the indicated FID MAY even be disjoint from the identified destination.

Once the credit window is identified, the credit window size MUST be increased by the value contained in the Additional Credits field. If the increase results in a window overflow, i.e., the size of the credit window after the increase is smaller than the original credit window size, the Credit Window must be set to its maximum (0xFFFFFFFFFFFFFFFF).

No response is sent by the router to a modem after processing a Credit Window Grant Data Item received in a Credit Control Response Message. In other cases, the receiving router MUST send a Credit Window Status Data Item or items reflecting the resulting Credit Window value of the updated credit window. When the Credit Grant Data Item is received in a Destination Up Message, the Credit Window Status Data Item(s) MUST be sent in the corresponding Destination Up Response Message. Otherwise, a Credit Control Message MUST be sent.

3.3.4. Credit Window Status

The Credit Window Status Data Item is used by a router to report the current credit window size to its peer modem. One or more 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 Credit Window Status Data Items in a single message are used to indicate different sizes of different credit windows. Similar to the Credit Window Grant, credit windows are identified using FID values that have been previously been sent by the modem.

The format of the Credit Window Status Data Item is:

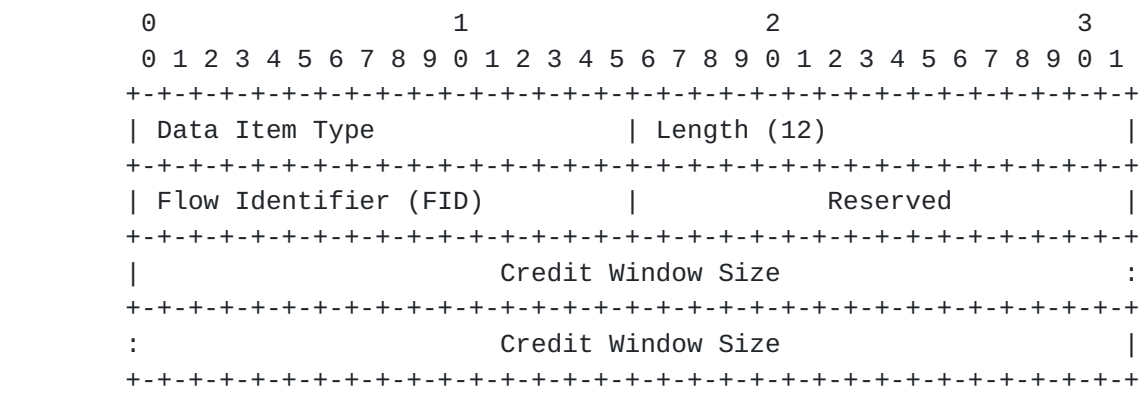

Data Item Type: TBA7

Length: 12

Per [RFC8175] Length is the number of octets in the Data Item. It MUST be equal to twelve (12).

Flow Identifier (FID):

A flow identifier as defined by the Traffic Classification Data Item. The FID also uniquely identifies a credit window.

Credit Window Size:

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].

When receiving this Data Item, a modem MUST identify the credit window indicated by the FID. If the FID is not known to the modem, it SHOULD report or log this information and discard the Data Item. As with the Credit Window Grant Data Item, the FID MAY be unrelated to the Destination indicated in the message carrying the Data Item.

Once the credit window is identified, the modem SHOULD check the received Credit Window Size field value against the outstanding credit window's available credits at the time the most Credit Window Initialization or Grant Data Item associated with the indicated FID was 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 Credit Window Initialization Data Item to reset the associated credit window size to the modem's current view of the available credits. As defined above, Credit Window Initialization Data Items are sent in Session Update Messages. When multiple Data Items need to be sent, they SHOULD be combined into a single message when

possible. Alternatively, and also in cases where there are small differences, the modem MAY adjust the values sent in Credit Window Grant Data Items to account for the reported Credit Window.

3.3.5. Credit Window Request

The Credit Window Request Data Item is used by a router to request additional credits for particular credit windows. Credit Window Request Data Items are carried in Credit Control Messages, and one or more Credit Window Request Data Items MAY be present in a message.

Credit windows identified using a FID as defined above in Section 3.3.1. Multiple FIDs MAY be present to allow for the case where the router identifies that credits are needed in multiple credit windows. A special FID value, as defined below, is used to indicate that a credit request is being made across all queues.

The format of the Credit Window Request Data Item is:

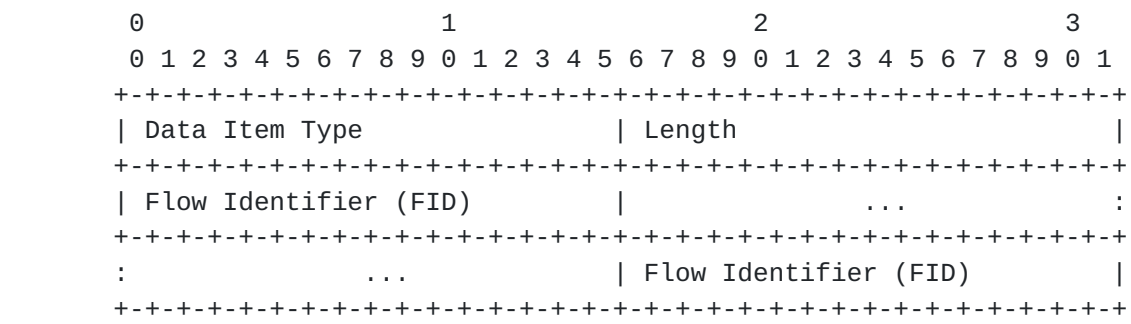

Data Item Type: TBA8

Length: Variable

Per [RFC8175] Length is the number of octets in the Data Item, excluding the Type and Length fields. It will equal the number of FID fields carried in the Data Item times 2 and MUST be at least 2.

Flow Identifier (FID):

A flow identifier as defined by the Traffic Classification Data Item. The FID also uniquely identifies a credit window. The special value of 0xFFFF indicates that the request applies to all FIDs.

A modem receiving this Data Item MUST provide a Credit Increment for the indicated credit windows via Credit Window Grant Data Items carried in a new Credit Control Message. Multiple values and queue

indexes SHOULD be combined into a single Credit Control Message when possible. Unknown FID values SHOULD be reported or logged and then ignored by the modem.

3.4. Management Considerations

This section provides several network management guidelines to implementations supporting the credit window mechanisms defined in this document.

Modems MAY support the configuration of the number of credit windows (queues) to advertise to a router.

Routers may have limits on the number of queues that they can support and, perhaps, even limits in supported credit window combinations, e.g., if per destination queues can even be supported at all. When modem-provided credit window information exceeds the capabilities of a router, the router MAY use a subset of the provided credit windows. Alternatively, a router MAY reset the session and indicate that the extension is not supported. In either case, the mismatch of capabilities SHOULD be reported to the user via normal network management mechanisms, e.g., user interface or error logging.

4. Compatibility

The data items defined in this document will only be used when extensions require their use.

5. Security Considerations

This document introduces credit window control and flow mechanisms to DLEP. These mechanisms do not inherently introduce any additional threats above those documented in [RFC8175]. The approach taken to Security in that document applies equally to the mechanism defined in this document.

6. IANA Considerations

This document requests the assignment of several values by IANA. All assignments are to registries defined by [RFC8175].

6.1. 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 1: Requested Message Values

6.2. Data Item Values

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

Type Code	Description
TBA1	Traffic Classification
TBA4	Credit Window Initialization
TBA5	Credit Window Association
TBA6	Credit Window Grant
TBA7	Credit Window Status
TBA8	Credit Window Request

Table 2: Requested Data Item Values

6.3. DLEP Sub Data Item Registry

Upon approval of this document, IANA is requested to create a new DLEP registry, named "Sub Data Item Type Values". The registry shall identify the type code value, the Data Item which may use the value, and a description of the value. While the same value may be reused in different Data Items, this is not recommended at this time.

The following table provides initial registry values and the [RFC8126] defined policies that should apply to the registry:

Type Code	Valid Data Items	Description
0	N/A	Reserved
1	N/A	Reserved (for pause sub-DI)
2	DiffServ Traffic Classification	DiffServ Traffic Classification
3	Ethernet Traffic Classification	Ethernet Traffic Classification
4-65407		Specification Required
65408-65534		Private Use
65535		Reserved

Table 3: Initial Registry Values

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8175] Ratliff, S., Jury, S., Satterwhite, D., Taylor, R., and B. Berry, "Dynamic Link Exchange Protocol (DLEP)", RFC 8175, DOI 10.17487/RFC8175, June 2017, <https://www.rfc-editor.org/info/rfc8175>.

7.2. Informative References

[I-D.ietf-manet-credit-window] Ratliff, S., "Credit Windowing extension for DLEP", draft-ietf-manet-credit-window-07 (work in progress), November 2016.

[I-D.ietf-manet-dlep-da-credit-extension]
Cheng, B., Wiggins, D., and L. Berger, "DLEP DiffServ Aware Credit Window Extension", draft-ietf-manet-dlep-da-credit-extension-04 (work in progress), March 2018.

[I-D.ietf-manet-dlep-pause-extension]
Cheng, B., Wiggins, D., and L. Berger, "DLEP Control Plane Based Pause Extension", draft-ietf-manet-dlep-pause-extension-03 (work in progress), March 2018.

[IEEE.802.1Q_2014]
IEEE, "IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks", IEEE 802.1Q-2014, DOI 10.1109/ieeestd.2014.6991462, December 2014, <http://ieeexplore.ieee.org/servlet/opac?punumber=6991460>.

[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, <https://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, DOI 10.17487/RFC2475, December 1998, <https://www.rfc-editor.org/info/rfc2475>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

Appendix A. Acknowledgments

The sub Data Item format was inspired by Rick Taylor's "Data Item Containers". He also proposed the separation of credit windows from traffic classification at IETF98. Many useful comments were received from contributors to the MANET working group. This document was derived from [I-D.ietf-manet-dlep-da-credit-extension] as a result of discussions at IETF101.

Authors' Addresses

Bow-Nan Cheng
Lincoln Laboratory
Massachusetts Institute of Technology
244 Wood Street
Lexington, MA 02421-6426

Email: bcheng@ll.mit.edu

David Wiggins
Lincoln Laboratory
Massachusetts Institute of Technology
244 Wood Street
Lexington, MA 02421-6426

Email: David.Wiggins@ll.mit.edu

Lou Berger
LabN Consulting, L.L.C.

Email: lberger@labn.net