

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 28 September 2025

B. Cheng
MIT Lincoln Laboratory
D. Wiggins

L. Berger
D. Fedyk, Ed.
LabN Consulting, L.L.C.
27 March 2025

Dynamic Link Exchange Protocol (DLEP) Traffic Classification Data Item
draft-ietf-manet-dlep-traffic-classification-17

Abstract

This document defines a new Data Item for the Dynamic Link Exchange Protocol (DLEP) to support traffic classification. Traffic classification information identifies traffic flows based on frame/packet content such as destination address. The Data Item is defined in an extensible and reusable fashion. Its use will be mandated in other documents defining specific DLEP extensions. This document also introduces DLEP Sub-Data Items, and Sub-Data Items are defined to support DiffServ and Ethernet traffic classification.

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 28 September 2025.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Key Words	3
2. Traffic Classification	3
2.1. Traffic Classification Data Item	4
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. Compatibility	11
4. Security Considerations	11
5. IANA Considerations	11
5.1. Data Item Values	11
5.2. DLEP Traffic Classification Sub-Data Item Registry	12
6. References	13
6.1. Normative References	13
6.2. Informative References	14
Appendix A. Acknowledgments	15
Authors' Addresses	16

1. Introduction

The Dynamic Link Exchange Protocol (DLEP) is defined in [RFC8175]. This protocol 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 DLEP specification does not include any flow identification beyond DLEP endpoints, i.e., flows are identified based on their DLEP endpoint.

This document defines DLEP Data Item formats which provide flow identification on a more granular basis. Specifically, it enables a router to use traffic flow classification information provided by the modem to identify traffic flows based on a combination of information found in a data plane header. (For general background on traffic classification see [RFC2475] Section 2.3.) The Data Item is

structured to allow for use of the defined traffic classification information with applications such as credit window control as specified in [I-D.ietf-manet-dlep-credit-flow-control]. The credit window control document provides an example of combining traffic classification and credit window flow control.

This document defines traffic classification based on a DLEP destination and flows identified by either DiffServ [RFC2475] Differentiated Services Codepoints (DSCPs) or IEEE 802.1Q [IEEE8021Q] Ethernet Priority Code Points (PCPs). The defined mechanism allows for flows to be described in a flexible fashion and when combined with applications such as credit window control, allows credit windows to be shared across traffic sent to multiple DLEP destinations and as part of multiple flows, or used exclusively for traffic sent to a particular destination and/or belonging to a particular 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 [RFC8651] or even 5-tuple IP flows.

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.

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 represents a list of flows that may be used at the same time to provide different service classes 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. Note that when extensions supporting multiple Sub-Data Item types are negotiated, these types MAY be combined in a single Data Item.

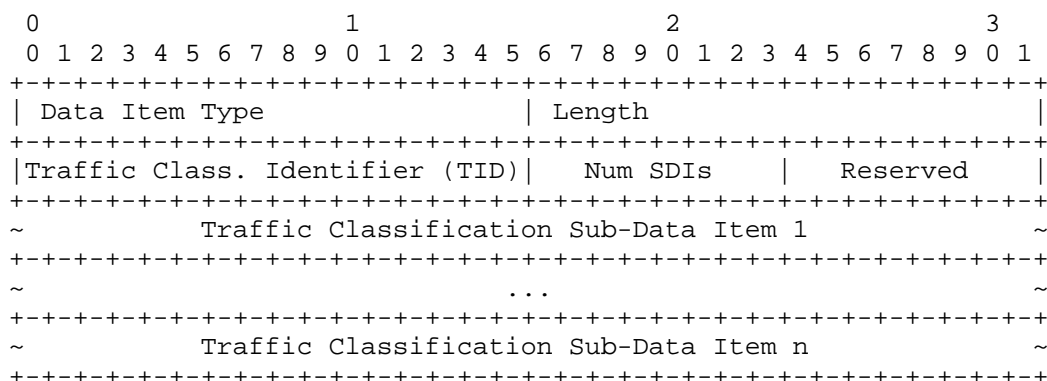
Each list of flows is identified using a "Traffic Classification Identifier" or "TID" and is expected to represent a valid combination of data plane identifiers that may be used at the same time. Each flow is identified via a "Flow Identifier" or "FID". Each FID is defined in a Sub-Data Item which carries 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 [I-D.ietf-manet-dlep-credit-flow-control]. TID and FID values have modem-local scope.

2.1. Traffic Classification Data Item

This section 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 any Data Item, the Data Items, including this Traffic Classification Data Item SHOULD be included by a modem in any Session Initialization Response Message, e.g., see [I-D.ietf-manet-dlep-credit-flow-control]. 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. The length here is limited by the packet data unit (PDU) length supported. For example if the Packet is limited to 1400 bytes then the length MUST NOT exceed this value. If larger packets are supported the maximum MUST be adjusted to be smaller or equal to the maximum PDU. Multiple messages can be used if there is more than fits in a single TLV.

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:

For the Traffic Classification Data Item this reserved field is currently unused. It MUST be set to all zeros for this version of the Data Item and it is currently ignored on reception. This allows for future extensions of the Data Item if needed.

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. If the associated traffic classification information is found, the router MUST replace the corresponding information using the values carried in the Data Item. In both cases, a router MUST also ensure that any data plane state, e.g., [I-D.ietf-manet-dlep-credit-flow-control], 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, see [RFC8175].

The format of the Traffic Classification Sub-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
+-----+-----+-----+-----+-----+-----+-----+-----+
| Sub-Data Item Type                                     | Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
~                                                     Value... ~
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Sub-Data Item Type:

A 16-bit unsigned integer that indicates the type and corresponding format of the Sub-Data Item's Value field. Sub-Data Item Types are scoped within the Data Item in which they are carried, i.e., the Sub-Data Item Type field MUST be used together with the Traffic Classification Data Item Type to identify the format of the Sub-Data Item. Traffic Classification Sub-Data Item Types are managed according to the IANA registry described in Section 5.2.

Length:

Variable


2.2. DiffServ Traffic Classification Sub-Data Item

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 that does not have an explicit match to a FID. A typical use of this is mapping any DSCPs that are not explicitly mapped to a default queue.

DS Field:

Each DS Field is an 8-bit that carries the DSCP field defined in [RFC2474].



DSCP: Differentiated Services Codepoint (RFC 2474).

MBZ: Must Be Zero - set to zero when transmitted.

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 in Section 2.1.1.

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 in Section 2.1.1.

2.3. Ethernet Traffic Classification Sub-Data Item

The Ethernet Traffic Classification Sub-Data Item identifies 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 [IEEE8021Q] 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:

```

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
+-----+-----+-----+-----+-----+-----+-----+-----+
| Sub-Data Item Type (2)          | Length          |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Flow Identifier (FID)           | NumPCPs          | VLAN Identifier (VID) |
+-----+-----+-----+-----+-----+-----+-----+-----+
| Pri. 1| Pri. 2| ..... | ..... | ..... | Pad  |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Sub-Data Item Type:

Sub-Data Item Type with value two (2) identifies the Ethernet Traffic Classification Sub-Data Item Type in the format defined in Section 2.1.1.

Length:

Variable

Length is defined above. For this Sub-Data Item, it is equal to four (4) plus the number of octets needed to accommodate the number of Priority fields indicated by the NumPCPs field. Note that as length is in octets and each Priority field is 4 bits, the additional length is the value carried in the NumPCPs field divided by two and rounded up to the next higher integer quantity. This TLV has maximum length of 4 plus 8 divided by 2 or 16 octets.

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 4-bit unsigned integer indicating the number of Priority fields carried in the Sub-Data Item. A zero (0) indicates a (wildcard) match against any PCP value that does not have an explicit match to a FID. A typical use of a wildcard is mapping any PCPs that are not explicitly mapped to a default queue. The maximum number of PCPs 8.

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. Any explicitly mapped VLANs are match first and then any VLANs that do not have a mapping map to this default mapping.

Priority:

Each Priority Field is 4-bits long and indicates a PCP field defined in [IEEE8021Q]. Note that zero (0) is a valid value for either PCP.

```

0   1   2   3
+---+---+---+---+
|   PCP   |MBZ|
+---+---+---+---+

```

PCP: Priority Code Point (IEEE8021Q)

MBZ: Must Be Zero - set to zero when transmitted.

Pad:

A 4-bit long field included when NumPCPs is an odd number. This field MUST be set to zero by the sender, and MUST be ignored on receipt.

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. Note that validation can include usage specific semantics such as those found in [I-D.ietf-manet-dlep-credit-flow-control]. Any failures MUST be treated as an error as described above in Section 2.1.1.

After successful validation, 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 Items. 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 in Section 2.1.1.

In cases where both Traffic Classification Sub-Data Item types are defined, matching on Ethernet information takes precedence. More specifically, when a packet matches both a DSCP indicated in a DiffServ Traffic Classification Sub-Data Item (Section 2.2) and a VID/PCP identified in an Ethernet Traffic Classification Sub-Data Item (Section 2.3), then the TID associated with the matching VLAN/PCP MUST be used.

3. Compatibility

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

The DLEP specification [RFC8175] defines handling of unexpected appearances of any Data Items, including those defined in this document.

4. Security Considerations

This document introduces finer grained flow identification mechanisms to DLEP. These mechanisms expose vulnerabilities similar to existing DLEP messages. An example of a threat to which traffic classification might be susceptible is a malicious actor masquerading as a DLEP peer could inject an alternate Traffic Classification Data Item, changing the mapping of traffic to queues that in turn causes delay, congestion or loss to one or more service classes. Other possible threats are given in the Security Considerations of [RFC8175] and are also applicable to, but not specific to, traffic classification.

The transport layer security mechanisms documented in [RFC8175], with some updated references to external documents listed below, can be applied to this document. Implementations following the "networked deployment" model described in the "Implementation Scenarios" of [RFC8175] SHOULD refer to [BCP195] for additional details. The Layer 2 security mechanisms documented in [RFC8175] can also, with some updates, be applied to the mechanism defined in this document. Examples of technologies that can be deployed to secure the Layer 2 link include [IEEE-802.1AE] and [IEEE-8802-1X].

5. IANA Considerations

5.1. Data Item Values

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

Type Code	Description
TBA1	Traffic Classification

Table 1: Requested Data Item Values

5.2. DLEP Traffic Classification Sub-Data Item Registry

Upon approval of this document, IANA is requested to create a new DLEP registry, named "Traffic Classification Sub-Data Item Type Values".

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

Type Code	Description	Reference
0	Reserved	
1	DiffServ Traffic Classification	[RFC2474]
2	Ethernet Traffic Classification	[IEEE8021Q]
3-65407	Specification Required	
65408-65534	Private Use	
65535	Reserved	

Table 2: Initial Registry Values

This section provides guidance to the Internet Assigned Numbers Authority (IANA) regarding registration of values related to the Traffic Classification Sub-Data Item Type Values registry for the DLEP protocol, in accordance with BCP 26 and [RFC8126].

This registry encompasses packet traffic classification, where standard packet header identifiers in packets or data frames indicate Quality of Service (QoS) treatment. It includes two specific registries for widely recognized identifiers used in QoS management for IP and Ethernet networks. Reserved values are set aside for similar future identifiers that may emerge to denote QoS treatment. However, requests for new entries are not expected to be frequent.

Allocations within the registry are subject to the following requirements:

1. Documentation of the intended use of the requested value, in compliance with the "Specification Required" policy defined in [RFC8126].
2. Approval by the Designated Expert (DE) appointed by the IESG. The DE must:
 - * Verify that the requested value is clearly documented and its purpose and usage are unambiguous.
 - * Ensure the proposed value does not conflict with existing work or ongoing efforts within the IETF.
 - * Confirm that any specification requesting a code point has undergone review by the MANET working group (or a successor mailing list designated by the IESG).
 - * Validate that external specifications requesting code points are publicly available, permanently archived, and do not conflict with active or published IETF work.
 - * Ensure the review process is conducted in a timely manner, with any disputes resolved through consultation with the appropriate working groups.

To simplify future registrations, it is recommended that this guidance serves as a standard reference for all DLEP-related registries. Future specifications may include a header note pointing to this guidance document.

6. References

6.1. 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>>.

6.2. Informative References

- [BCP195] Best Current Practice 195,
<<https://www.rfc-editor.org/info/bcp195>>.
At the time of writing, this BCP comprises the following:
- Sheffer, Y., Saint-Andre, P., and T. Fossati,
"Recommendations for Secure Use of Transport Layer
Security (TLS) and Datagram Transport Layer Security
(DTLS)", BCP 195, RFC 9325, DOI 10.17487/RFC9325, November
2022, <<https://www.rfc-editor.org/info/rfc9325>>.
- Moriarty, K. and S. Farrell, "Deprecating TLS 1.0 and TLS
1.1", BCP 195, RFC 8996, DOI 10.17487/RFC8996, March 2021,
<<https://www.rfc-editor.org/info/rfc8996>>.
- [I-D.ietf-manet-dlep-credit-flow-control]
Cheng, B., Wiggins, D., Ratliff, S., Berger, L., and E.
Kinzie, "Dynamic Link Exchange Protocol (DLEP) Credit-
Based Flow Control Messages and Data Items", Work in
Progress, Internet-Draft, draft-ietf-manet-dlep-credit-
flow-control-18, 19 March 2025,
<[https://datatracker.ietf.org/api/v1/doc/document/draft-
ietf-manet-dlep-credit-flow-control/](https://datatracker.ietf.org/api/v1/doc/document/draft-ietf-manet-dlep-credit-flow-control/)>.
- [I-D.ietf-manet-dlep-da-credit-extension]
Cheng, B., Wiggins, D., Berger, L., and D. E. Eastlake,
"DLEP DiffServ Aware Credit Window Extension", Work in
Progress, Internet-Draft, draft-ietf-manet-dlep-da-credit-
extension-21, 3 March 2025,
<[https://datatracker.ietf.org/doc/html/draft-ietf-manet-
dlep-da-credit-extension-21](https://datatracker.ietf.org/doc/html/draft-ietf-manet-dlep-da-credit-extension-21)>.
- [IEEE-802.1AE]
"802.1AE-2018 - IEEE Standard for Local and metropolitan
area networks-Media Access Control (MAC) Security",
<<https://ieeexplore.ieee.org/document/8585421>>.

[IEEE-8802-1X]

"8802-1X-2021 - IEEE/ISO/IEC International Standard-Telecommunications and exchange between information technology systems--Requirements for local and metropolitan area networks--Part 1X:Port-based network access control",
<<https://ieeexplore.ieee.org/document/9650828>>.

[IEEE8021Q]

IEEE, "IEEE Standard for Local and Metropolitan Area Networks--Bridges and Bridged Networks",
DOI 10.1109/IEEESTD.2022.10004498, IEEE 802.1Q-2022, July 2022, <<https://ieeexplore.ieee.org/document/8403927>>.

[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>>.

[RFC8651]

Cheng, B., Wiggins, D., and L. Berger, Ed., "Dynamic Link Exchange Protocol (DLEP) Control-Plane-Based Pause Extension", RFC 8651, DOI 10.17487/RFC8651, October 2019,
<<https://www.rfc-editor.org/info/rfc8651>>.

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 IETF 101. Many useful comments were received from contributors to the MANET working group, notably Ronald in't Velt and David Black.

We had the honor of working too briefly with David Wiggins on this and related DLEP work. His contribution to the IETF and publication of the first and definitive open source DLEP implementation have been critical to the acceptance of DLEP. We mourn his passing on November 23, 2023. We wish to recognize his guidance, leadership and professional excellence. We were fortunate to benefit from his leadership and friendship. He shall be missed.

Authors' Addresses

Bow-Nan Cheng
MIT Lincoln Laboratory
Massachusetts Institute of Technology
244 Wood Street
Lexington
Email: bcheng@ll.mit.edu

David Wiggins
Email: david@none.org

Lou Berger
LabN Consulting, L.L.C.
Email: lberger@labn.net

Don Fedyk (editor)
LabN Consulting, L.L.C.
Email: dfedyk@labn.net