

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 5, 2019

B. Cheng
D. Wiggins
MIT Lincoln Laboratory
L. Berger, Ed.
LabN Consulting, L.L.C.
April 3, 2019

DLEP Control Plane Based Pause Extension
draft-ietf-manet-dlep-pause-extension-06

Abstract

This document defines an extension to the Dynamic Link Exchange Protocol (DLEP) that enables a modem to use DLEP messages to pause and resume data traffic coming from its peer router.

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 October 5, 2019.

Copyright Notice

Copyright (c) 2019 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	3
2.	Extension Usage and Identification	3
3.	Extension Data Items	3
3.1.	Queue Parameters	3
3.1.1.	Queue Parameter Sub Data Item	5
3.2.	Pause	6
3.3.	Restart	8
4.	Security Considerations	9
5.	IANA Considerations	9
5.1.	Extension Type Value	9
5.2.	Data Item Values	9
5.3.	Queue Parameters Sub Data Item Values	10
6.	References	10
6.1.	Normative References	10
6.2.	Informative References	11
Appendix A.	Acknowledgments	11
	Authors' Addresses	11

[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. This document defines one such extension.

The base DLEP specification does not include any data plane flow control capability. Various flow control methods are possible, e.g., see [\[I-D.ietf-manet-dlep-da-credit-extension\]](#). The extension defined in this document supports flow control of data traffic based on explicit messages sent via DLEP by a modem to indicate when a router should hold off sending traffic, and when it should resume. The extension also optionally supports DSCP (differentiated services codepoint) aware flow control. For general background on Differentiated Services see [\[RFC2475\]](#). The extension defined in this document is referred to as "Control Plane Based Pause". Note that this mechanism only applies to traffic that is to be transmitted on the modem's attached data channel and not to DLEP control messages themselves.

This document defines a new DLEP Extension Type Value in [Section 2](#) which is used to indicate the use of the extension, and three new DLEP Data Items in [Section 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. Extension Usage and Identification

The use of the Control Plane Based Pause Extension SHOULD be configurable. To indicate that the implementation supports use of the Control Plane Based Pause Extension, an implementation MUST include the Control Plane Based Pause Extension Type Value in the Extensions Supported Data Item. The Extensions Supported Data Item is sent and processed according to [[RFC8175](#)].

The Control Plane Based Pause Extension Type Value is TBA1, see [Section 5](#).

3. Extension Data Items

Three data items are defined by this extension. The Queue Parameters Data Item is used by a modem to provide information about the DSCPs it uses in forwarding. The Pause Data Item is used by a modem to indicate when a router should cease sending packets and the Restart Data Item is used by a modem to indicate when a router can resume sending packets.

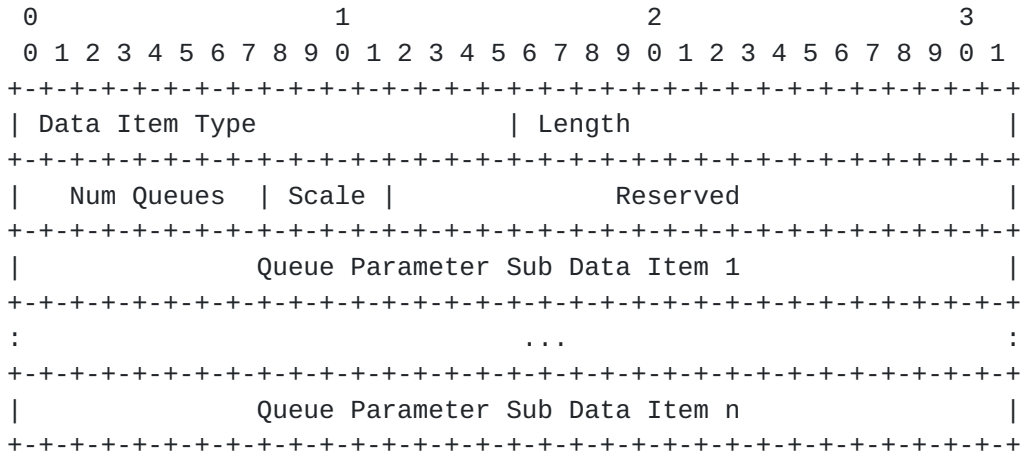
3.1. Queue Parameters

The Queue Parameters Data Item is used by a modem to indicate DSCP values that may be independently paused. This data item MUST be included in a Session Initialization Response Message that also contains the Control Plane Based Pause 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 groups DSCPs into logical queues, each of which is identified by a "Queue Index". The number of logical queues, or queue indexes, is variable as is the number of DSCPs associated with each queue. A queue size (in bytes) is provided for informational purposes. Queue Indexes are numbered sequentially from zero, where queue index zero is a special case covering DSCPs which are not otherwise associated with Queue Index.

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. Additional logical queues are represented in a variable series of Queue Parameter sub data items.

The format of the Queue Parameters Data Item is:



Data Item Type: TBA2

Length: Variable

Per [RFC8175] 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), and is equal to one greater than the number of included Queue Parameter Sub Data Items.

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	KiB - Kilobytes (1024 B)
2	MiB - Megabytes (1024 KiB)
3	GiB - Gigabytes (1024 MiB)


Reserved:

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

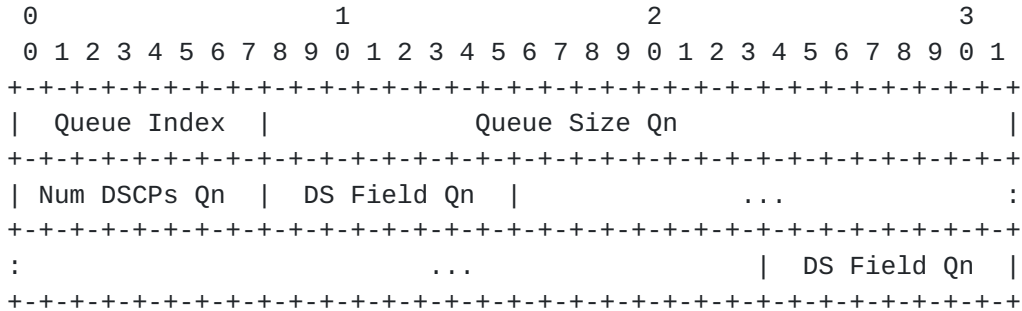
3.1.1. Queue Parameter Sub Data Item

Queue Parameter Sub Data Items are an unordered list composed of sub data items with a common format. The format of the Queue Parameter Sub Data Item is patterned after the standard DLEP data item format, see [RFC8175] Section 11.3. 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 Queue Parameter Sub Data Item is:



and Value has the format:



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 Queue Parameters Data Item Type to identify the format of the Sub Data Item. This field MUST be set to one (1) for the Queue Parameter Sub Data Item.

Length: Variable

Length is the number of octets in the sub data item, excluding the Type and Length fields.

Queue Index:

An 8-bit field indicating the queue index of the queue parameter represented in the sub data item. Only the first instance of a particular Queue Index value is meaningful. Subsequent sub data items containing the same Queue Index values, if present, MAY be logged via a management interface and MUST otherwise be ignored. Note that the value 255 is reserved and MUST NOT be used in this field.

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.


Num DSCPs Qn:

An 8-bit unsigned integer indicating the number of DSCPs associated with the queue index associated with the sub data item. This field MUST contain a value of at least one (1).

DS Field Qn:

The data item contains a sequence of 8 bit DS Fields. The number of DS Fields present MUST equal the sum of all Num DSCPs field values.

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



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

3.2. Pause

The Pause Data Item is used by a modem to indicate to its peer that traffic is to be suppressed. The motivating use case is for this data item is when a modem's internal queue length exceeds a particular threshold. Other use cases are possible, e.g., when there a non queue related congestion points within a modem, but such are not explicitly described in this document.


A modem can indicate that traffic is to be suppressed on a device-wide or destination-specific basis. An example of when a modem might use device wide indications is when output queues are shared across all destinations, and destination specific might be used when per destination queuing is used. To indicate that suppression applies to all destinations, a modem MUST send the Pause Data Item in a Session Update Message. To indicate that suppression applies to a particular destination a modem MUST send the Pause Data Item in a Destination Update Message.

Each Pause Data Item identifies the traffic to be suppressed by the Queue Index defined by [Section 3.1](#), which in turn indicates a set of traffic identified by DSCPs. The special value of 255 is used to indicate that all traffic is to be suppressed.

While there is no restriction on the number of Messages containing Pause Data Item that may be sent by a modem, a modem SHOULD include multiple queue indexes in the same message when possible.

A router which receives the Pause Data Item MUST cease sending the identified traffic to the modem. This may of course translate into the router's queues exceeding their own thresholds. If a received Pause Data Item contains a Queue Index value other than 255, or a queue index established by a Session Initialization or Session Update Message, the router MUST terminate the session with a Status Data Item indicating Invalid Data.

The format of the Pause Data Item is:



Data Item Type: TBA3

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 Queue Index fields carried in the data item.

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 all traffic is to be suppressed to the modem, when the data item is carried in a Session Update Message, or is to be suppressed to a destination, when the data item is carried in Destination Update Message.


3.3. Restart

The Restart Data Item is used by a modem to indicate to its peer that transmission of previously suppressed traffic may be resumed. An example of when a modem might send this data item is when an internal queue length drops below a particular threshold.

The sending of this data item parallels the Pause Data Item, see the previous section, and follows the same rules. As above, to indicate that transmission can resume to all destinations, a modem MUST send the Restart Data Item in a Session Update Message. It also includes that to indicate that transmission can resume to a particular destination a modem MUST send the Pause Restart Item in a Destination Update Message. Finally, queue indexes are interpreted in the same way as in the Pause Data Item..

A router which receives the Restart Data Item SHOULD resume transmission of the identified traffic to the modem.

The format of the Restart Data Item matches the Pause Data Item and is:



Data Item Type: TBA4

Length: See [Section 3.2](#).

Queue Index: See [Section 3.2](#).

4. Security Considerations

The extension introduces a new mechanism for flow control between a router and modem using DLEP. The extension does not inherently introduce any additional vulnerabilities above those documented in [RFC8175]. The approach taken to Security in that document applies equally when running the extension defined in this document.

Note that this extension does allow a compromised or impersonating modem to suppress transmission by the router, but this is not a substantively different attack by such a compromised modem simply dropping all traffic destined to, or sent by a router. [RFC8175] defines the use of TLS to protect against the impersonating attacker.

5. IANA Considerations

This document requests the assignment of 4 values by IANA. All assignments are to registries defined by [RFC8175].

5.1. Extension Type Value

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

```

+-----+-----+
| Code | Description          |
+-----+-----+
| TBA1 | Control Plane Based Pause |
+-----+-----+

```

Table 1: Requested Extension Type Value

5.2. Data Item Values

This document requests 3 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
TBA2	Queue Parameters
TBA3	Pause
TBA4	Restart

Table 2: Requested Data Item Values

5.3. Queue Parameters Sub Data Item Values

Upon approval of this document, IANA is requested to create a new DLEP registry, named "Queue Parameters 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/Policy
0	Reserved
1	Queue Parameter
2-65407	Specification Required
65408-65534	Private Use
65535	Reserved

Table 3

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

- [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-07](#) (work in progress), March 2019.
- [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" idea.

Authors' Addresses

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

Email: bcheng@ll.mit.edu

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

Email: David.Wiggins@ll.mit.edu

Lou Berger (editor)
LabN Consulting, L.L.C.

Email: lberger@labn.net