

Internet Engineering Task Force (IETF)
Request for Comments: 8651
Category: Standards Track
ISSN: 2070-1721

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

Dynamic Link Exchange Protocol (DLEP)
Control-Plane-Based Pause Extension

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 is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8651>.

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
 - 1.1. Key Words
 2. Extension Usage and Identification
 3. Extension Data Items
 - 3.1. Queue Parameters
 - 3.1.1. Queue Parameter Sub-Data Item
 - 3.2. Pause
 - 3.3. Restart
 4. Security Considerations
 5. IANA Considerations
 - 5.1. Extension Type Value
 - 5.2. Data Item Values
 - 5.3. Queue Parameter Sub-Data Item Values
 6. References
 - 6.1. Normative References
 - 6.2. Informative References
- Acknowledgments
Authors' Addresses

1. Introduction

The Dynamic Link Exchange Protocol (DLEP) is defined in [RFC8175]. It provides the exchange of link-related control information between 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. 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. This functionality parallels the flow-control mechanism found in PPP over Ethernet (PPPoE) per [RFC5578]. The extension also optionally supports DSCP-aware flow control ("DSCP" stands for "Differentiated Services Code Point") for use by Diffserv-aware modems. (For general background on Differentiated Services, see [RFC2475].) This functionality is very similar to that provided by Ethernet priority-based flow control; see [IEEE.802.1Q_2014]. The extension defined in this document is referred to as "Control-Plane-Based Pause". Other flow-control methods are possible with DLEP; for example, see [DLEP-DIFFSERV] and [DLEP-CREDIT].

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. Furthermore, it applies only to the single subnetwork that is used to connect a modem and a router, and for traffic sent from a router to a modem.

This document defines a new DLEP Extension Type Value that is used to indicate the use of the extension; see Section 2. Three new DLEP Data Items are defined 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 the 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 2; 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

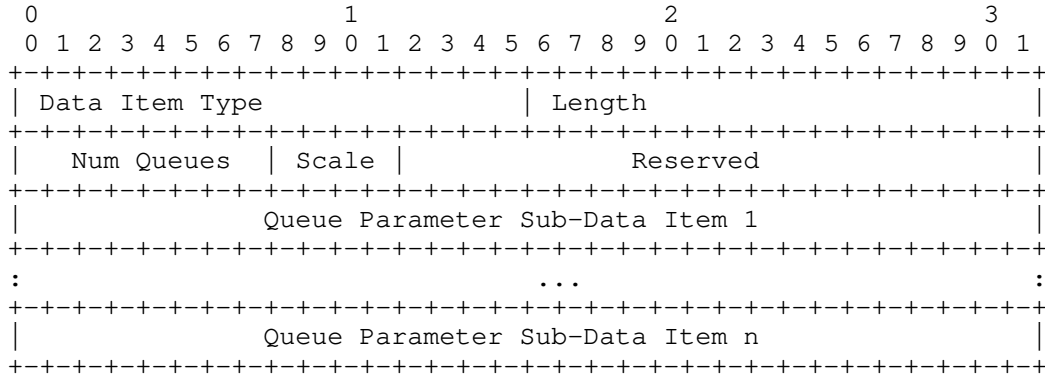
The Queue Parameters Data Item is sent by a modem to a router to indicate DSCP values that may be independently paused. This Data Item MUST be included in a Session Initialization Response Message that also contains the Control-Plane-Based Pause Extension Type Value in the Extensions Supported Data Item. 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" field. 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. Queue Index fields are numbered sequentially from zero, where queue index zero is a special case covering DSCPs that are not otherwise associated with a Queue Index field.

An implementation that does not support DSCPs would indicate one queue with zero 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:
23

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 Queue Parameter Sub-Data Items that follow. This field MUST contain a value of at least one (1).

Scale:
A 4-bit unsigned integer indicating the scale used in the Queue Size Qn field. The valid values are:

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

Table 1: Queue Size Qn Field Values

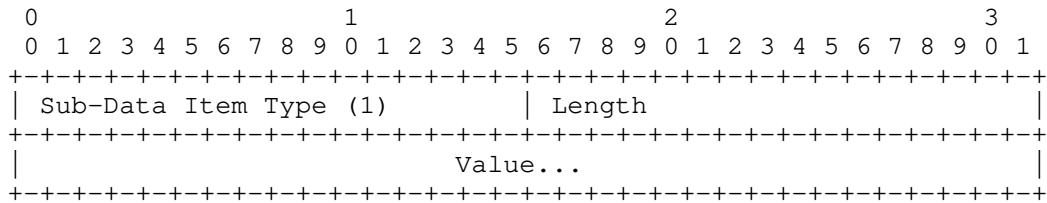
Reserved: A 20-bit field that MUST be set to zero (0) by the sender (a modem) and ignored by the receiver (a router).

3.1.1.1. Queue Parameter Sub-Data Item

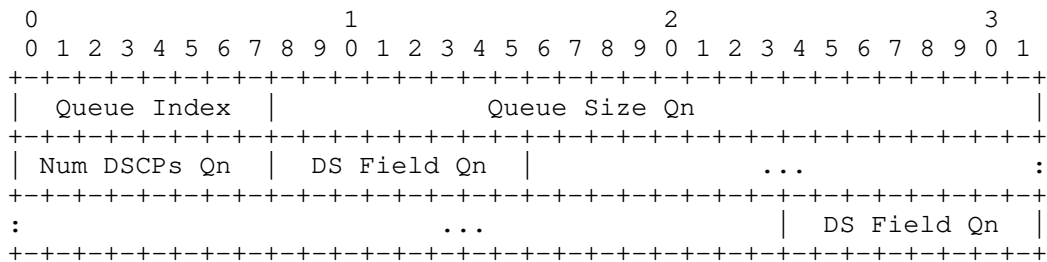
Queue Parameter Sub-Data Items are an unordered list composed of Sub-Data Items with a common format. The format of the Queue Parameter Sub-Data Item is patterned after the standard format for

the DLEP Data Item; 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. In particular, the receiving implementation MUST issue a Session Termination Message containing a Status Data Item with status code set to 130 ("Invalid Data") and transition to the Session Termination state.

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 that supports the traffic with the DSCPs associated with the queue index.

Num DSCPs Qn:

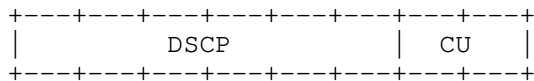
An 8-bit unsigned integer indicating the number of DSCPs associated with the queue index associated with the Sub-Data Item.

DS Field Qn:

The Data Item contains a sequence of 8-bit DS fields. The number of DS fields present MUST equal the Num DSCPs Qn field value.

The DS field structure is the same as the structure shown in [RFC2474].





DSCP: Differentiated Services Code Point

CU: Currently Unused; MUST be zero

3.2. Pause

The Pause Data Item is sent by a modem to a router to indicate to its peer that traffic is to be suppressed, i.e., paused. The motivating use case 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 are non-queue-related congestion points within a modem. Such cases are not explicitly described in this document.

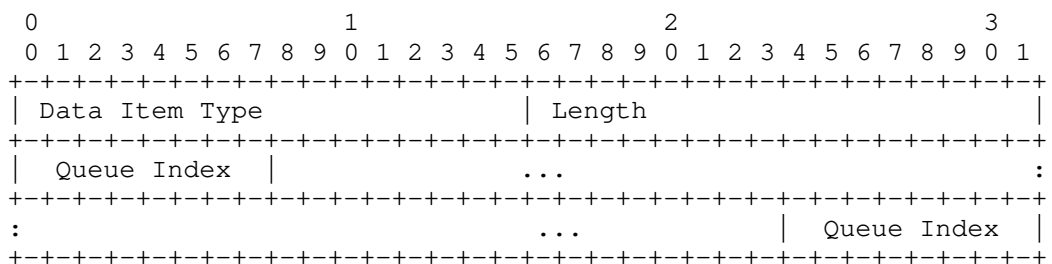
A modem can indicate that traffic is to be suppressed on a device-wide or destination-specific basis. An example of when a modem might use device-wide suppression is when output queues are shared across all destinations. Destination-specific suppression might be used when per-destination queuing is used. 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 field (Section 3.1), which in turn indicates traffic identified by one or more 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 Items that may be sent by a modem, a modem SHOULD include multiple queue indexes in the same message when possible.

A router that 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:
24

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 that (1) all traffic to the modem is to be suppressed when the Data Item is carried in a Session Update Message or (2) all

traffic to a particular destination is to be suppressed when the Data Item is carried in a Destination Update Message.

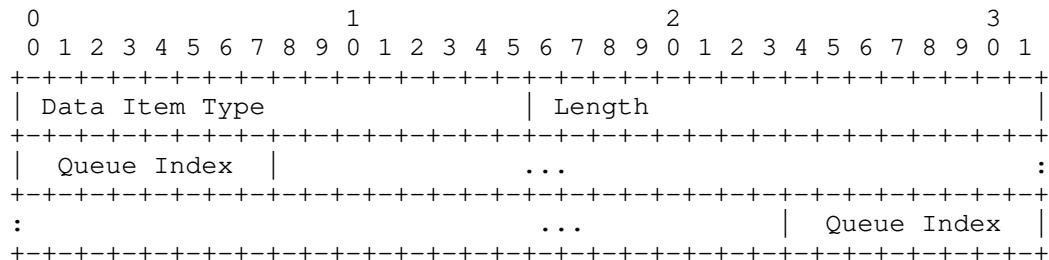
3.3. Restart

The Restart Data Item is sent by a modem to a router 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 Section 3.2) and follows the same rules. To indicate that transmission can resume to all destinations, a modem MUST send the Restart Data Item in a Session Update Message. To indicate that transmission can resume to a particular destination, a modem MUST send the Restart Data Item in a Destination Update Message. Finally, the same rules apply to queue indexes.

A router that 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: 25

Length: See Section 3.2.

Queue Index: See Section 3.2.

4. Security Considerations

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

Implementations of the extension defined in this document MUST support the configuration and use of TLS, as described in [RFC8175], in order to protect configurations where injection attacks are possible, i.e., when the link between a modem and router is not otherwise protected.

Note that this extension does allow a compromised or impersonating modem to suppress transmission by the router or a switch that interconnects the modem and router. Similar attacks are generally possible with base DLEP -- for example, an impersonating modem may cause a session reset, or a compromised modem can simply drop all traffic destined for or sent by a router. [RFC8175] defines the use of TLS to protect against such impersonating attackers.

5. IANA Considerations

This document assigns four new values and creates a new subregistry in the "Dynamic Link Exchange Protocol (DLEP) Parameters" registry.

5.1. Extension Type Value

This document adds a new assignment to the DLEP extensions registry

named "Extension Type Values" [RFC8175], per the "Specification Required" policy [RFC8126]. IANA has assigned the following value:

Code	Description
2	Control-Plane-Based Pause

Table 2: Extension Type Value

5.2. Data Item Values

This document adds three new assignments to the DLEP Data Item registry named "Data Item Type Values" [RFC8175], per the "Specification Required" policy [RFC8126]. IANA has assigned the following values:

Type Code	Description
23	Queue Parameters
24	Pause
25	Restart

Table 3: Data Item Values

5.3. Queue Parameter Sub-Data Item Values

IANA has created a new DLEP registry named "Queue Parameter Sub-Data Item Type Values".

Table 4 provides initial registry values and the registration policies [RFC8126] that apply:

Type Code	Description/Policy
0	Reserved
1	Queue Parameter
2-65407	Specification Required
65408-65534	Private Use
65535	Reserved

Table 4: Initial Registry Values

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

[DLEP-CREDIT]

Cheng, B., Wiggins, D., Berger, L., and S. Ratliff, "DLEP Credit-Based Flow Control Messages and Data Items", Work in Progress, Internet-Draft, draft-ietf-manet-dlep-credit-flow-control-04, 6 March 2019, <<https://tools.ietf.org/html/draft-ietf-manet-dlep-credit-flow-control-04>>.

[DLEP-DIFFSERV]

Cheng, B., Wiggins, D., and L. Berger, "DLEP DiffServ Aware Credit Window Extension", Work in Progress, Internet-Draft, draft-ietf-manet-dlep-da-credit-extension-07, 6 March 2019, <<https://tools.ietf.org/html/draft-ietf-manet-dlep-da-credit-extension-07>>.

[IEEE.802.1Q_2014]

IEEE, "IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks", IEEE 802.1Q-2014, <<https://ieeexplore.ieee.org/document/6991462>>.

[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>>.

[RFC5578] Berry, B., Ed., Ratliff, S., Paradise, E., Kaiser, T., and M. Adams, "PPP over Ethernet (PPPoE) Extensions for Credit Flow and Link Metrics", RFC 5578, DOI 10.17487/RFC5578, February 2010, <<https://www.rfc-editor.org/info/rfc5578>>.

[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>>.

Acknowledgments

The format for the Sub-Data Item 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
United States of America

Email: bcheng@ll.mit.edu

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

United States of America

Email: David.Wiggins@ll.mit.edu

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

Email: lberger@labn.net