

Mobile Ad hoc Networks Working
Group
Internet-Draft
Intended status: Standards Track
Expires: August 14, 2014

S. Ratliff
B. Berry
G. Harrison
S. Jury
Cisco Systems
D. Satterwhite
Broadcom
February 10, 2014

Dynamic Link Exchange Protocol (DLEP)
draft-ietf-manet-dlep-05

Abstract

When routing devices rely on modems to effect communications over wireless links, they need timely and accurate knowledge of the characteristics of the link (speed, state, etc.) in order to make forwarding decisions. In mobile or other environments where these characteristics change frequently, manual configurations or the inference of state through routing or transport protocols does not allow the router to make the best decisions. A bidirectional, event-driven communication channel between the router and the modem is necessary.

Status of this Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on August 14, 2014.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	4
1.1 Requirements	8
2. Assumptions	8
3. Credits	9
4. Metrics	9
5. Extensions to DLEP	10
6. Normal Session Flow	11
6.1 DLEP Modem session flow - Discovery case	11
6.2 DLEP Modem session flow - Configured case	11
6.3 DLEP Router session flow	12
6.4 Common Session Flow	12
7. Mandatory Signals and Data Items	13
8. Generic DLEP Message Definition	14
9. DLEP Data Items	15
9.1 DLEP Version	16
9.2 DLEP Port	17
9.3 Peer Type	17
9.4 MAC Address	18
9.5 IPv4 Address	18
9.6 IPv6 Address	19
9.7 Maximum Data Rate (Receive)	20
9.8 Maximum Data Rate (Transmit)	20
9.9 Current Data Rate (Receive)	21
9.10 Current Data Rate (Transmit)	22
9.11 Expected Forwarding Time	23
9.12 Latency	23
9.13 Resources (Receive)	24
9.14 Resources (Transmit)	25
9.15 Relative Link Quality (Receive)	25
9.16 Relative Link Quality (Transmit)	26
9.17 Status	26

9.18	Heartbeat Interval	27
9.19	Link Characteristics ACK Timer	28
9.20	Credit Window Status	28
9.21	Credit Grant Request	29
9.22	Credit Request	30
10.	DLEP Protocol Messages	31
10.1	Signal TLV Values	31
10.2	Peer Discovery Message	32
10.3	Peer Offer Message	32
10.4	Peer Initialization Message	33
10.5	Peer Initialization ACK Message	33
10.6	Peer Update Message	34
10.7	Peer Update ACK Message	35
10.8	Peer Termination Message	35
10.9	Peer Termination ACK Message	36
10.10	Destination Up Message	36
10.11	Destination Up ACK Message	37
10.12	Destination Down Message	37
10.13	Destination Down ACK Message	37
10.14	Destination Update Message	38
10.15	Heartbeat Message	38
10.16	Link Characteristics Request Message	39
10.17	Link Characteristics ACK Message	40
11.	Security Considerations	40
12.	IANA Considerations	40
12.1	Registrations	41
12.2	Expert Review: Evaluation Guidelines	41
12.3	Message (Signal) TLV Type Registration	41
12.4	DLEP Data Item Registrations	42
12.5	DLEP Well-known Port	42
12.6	DLEP Multicast Address	42
13.	Appendix A	42
13.1	Peer Level Message Flows	42
13.1.1	Modem Device Restarts Discovery	43
13.1.2	Modem Device Detects Peer Offer Timeout	43
13.1.3	Router Peer Offer Lost	44
13.1.4	Discovery Success	44
29.1.5	Router Detects a Heartbeat timeout	45
29.1.6	Modem Detects a Heartbeat timeout	45
29.1.7	Peer Terminate (from Modem) Lost	46
29.1.8	Peer Terminate (from Router) Lost	46
29.2	Destination Specific Message Flows	46
29.2.1	Modem Destination Up Lost	47
29.2.2	Router Detects Duplicate Destination Ups	47
29.2.3	Destination Up, No Layer 3 Addresses	48
29.2.4	Destination Up with IPv4, No IPv6	48
29.2.5	Destination Up with IPv4 and IPv6	48
29.2.6	Destination Session Success	49

Acknowledgements	49
Normative References	50
Informative References	50
Author's Addresses	50

[1. Introduction](#)

There exist today a collection of modem devices that control links of variable datarate and quality. Examples of these types of links include line-of-sight (LOS) radios, satellite terminals, and cable/DSL modems. Fluctuations in speed and quality of these links can occur due to configuration (in the case of cable/DSL modems), or on a moment-to-moment basis, due to physical phenomena like multipath interference, obstructions, rain fade, etc. It is also quite possible that link quality and datarate varies with respect to individual destinations on a link, and with the type of traffic being sent. As an example, consider the case of an 802.11g access point, serving 2 associated laptop computers. In this environment, the answer to the question "What is the datarate on the 802.11g link?" is "It depends on which associated laptop we're talking about, and on what kind of traffic is being sent." While the first laptop, being physically close to the access point, may have a datarate of 54Mbps for unicast traffic, the other laptop, being relatively far away, or obstructed by some object, can simultaneously have a datarate of only 32Mbps for unicast. However, for multicast traffic sent from the access point, all traffic is sent at the base transmission rate (which is configurable, but depending on the model of the access point, is usually 24Mbps or less).

In addition to utilizing variable datarate links, mobile networks are challenged by the notion that link connectivity will come and go over time. Effectively utilizing a relatively short-lived connection is problematic in IP routed networks, as routing protocols tend to rely on independent timers at OSI Layer 3 to maintain network convergence (e.g. HELLO messages and/or recognition of DEAD routing adjacencies). These short-lived connections can be better utilized with an event-driven paradigm, where acquisition of a new neighbor (or loss of an existing one) is signaled, as opposed to a timer-driven paradigm.

Another complicating factor for mobile networks are the different methods of physically connecting the modem devices to the router. Modems can be deployed as an interface card in a router's chassis, or as a standalone device connected to the router via Ethernet, USB, or even a serial link. In the case of Ethernet or serial attachment, with existing protocols and techniques, routing software cannot be aware of convergence events occurring on the radio link (e.g. acquisition or loss of a potential routing neighbor), nor can the router be aware of the actual capacity of the link. This lack of

awareness, along with the variability in datarate, leads to a situation where quality of service (QoS) profiles are extremely difficult to establish and properly maintain. This is especially true of demand-based access schemes such as Demand Assigned Multiple Access (DAMA) implementations used on some satellite systems. With a DAMA-based system, additional datarate may be available, but will not be used unless the network devices emit traffic at rate higher than the currently established rate. Increasing the traffic rate does not guarantee additional datarate will be allocated; rather, it may result in data loss and additional retransmissions on the link.

Addressing the challenges listed above, the authors have developed the Data Link Exchange Protocol, or DLEP. The DLEP protocol runs between a router and its attached modem devices, allowing the modem to communicate link characteristics as they change, and convergence events (acquisition and loss of potential routing destinations). The following diagrams are used to illustrate the scope of DLEP packets.




Figure 1: DLEP Network

In Figure 1, when the local modem detects the presence of a remote node, it (the local modem) sends a signal to its router via the DLEP protocol. Upon receipt of the signal, the local router may take whatever action it deems appropriate, such as initiating discovery protocols, and/or issuing HELLO messages to converge the network. On a continuing, as-needed basis, the modem devices utilize DLEP to report any characteristics of the link (datarate, latency, etc) that have changed. DLEP is independent of the link type and topology supported by the modem. Note that the DLEP protocol is specified to run only on the local link between router and modem. Some over the air signaling may be necessary between the local and remote modem in order to provide some parameters in DLEP messages between the local modem and local router, but DLEP does not specify how such over the air signaling is carried out. Over the air signaling is purely a matter for the modem implementer.

Figure 2 shows how DLEP can support a configuration where routers are

connected with different link types. In this example, Modem A implements a point-to-point link, and Modem B is connected via a shared medium. In both cases, the DLEP protocol is used to report the characteristics of the link (data rate, latency, etc.) to routers. The modem is also able to use the DLEP session to notify the router when the remote node is lost, shortening the time required to re-converge the network.

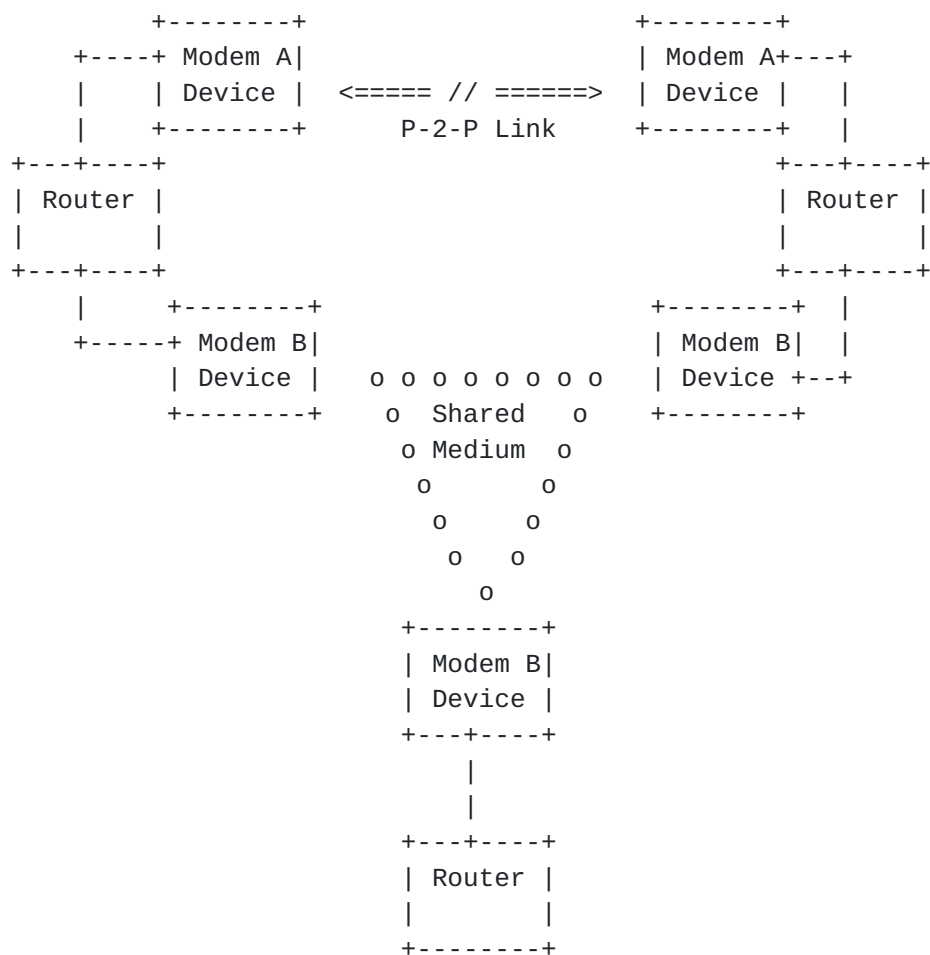


Figure 2: DLEP Network with Multiple Modem Devices

DLEP defines a set of messages used by modems and their attached routers. The messages are used to communicate events that occur on the physical link(s) managed by the modem: for example, a remote node

entering or leaving the network, or that the link has changed. Associated with these messages are a set of data items - information that describes the remote node (e.g., address information), and/or the characteristics of the link to the remote node.

The protocol is defined as a collection of type-length-value (TLV) based messages, specifying the signals that are exchanged between a router and a modem, and the data items associated with the signal. This document specifies transport of DLEP signals and data items via the TCP transport, with a UDP-based discovery mechanism. Other transports for the protocol are possible, but are outside the scope of this document.

DLEP signals are further defined as mandatory or optional. Signals will additionally have mandatory and optional data items. Implementations **MUST** support all mandatory messages and their mandatory data items to be considered compliant. Implementations **MAY** also support some, or all, of the optional messages and data items.

DLEP uses a session-oriented paradigm between the modem device and its associated router. If multiple modem devices are attached to a router (as in Figure 2), a separate DLEP session **MUST** exist for each modem. If a modem device supports multiple connections to a router (via multiple logical or physical interfaces), or supports connections to multiple routers, a separate DLEP session **MUST** exist for each connection. This router/modem session provides a carrier for information exchange concerning "destinations" that are available via the modem device. A "destination" can be either physical (as in the case of a specific far-end router), or a logical destination (as in a Multicast group). As such, all of the destination-level exchanges in DLEP can be envisioned as building an information base concerning the remote nodes, and the link characteristics to those nodes.

Multicast traffic destined for the variable-quality network (the network accessed via the DLEP modem) is handled in IP networks by deriving a Layer 2 MAC address based on the Layer 3 address. Leveraging on this scheme, Multicast traffic is supported in DLEP simply by treating the derived MAC address as any other "destination" (albeit a logical one) in the network. To support these logical destinations, one of the DLEP participants (typically, the router) informs the other as to the existence of the logical neighbor. The modem, once it is aware of the existence of this logical neighbor, reports link characteristics just as it would for any other destination in the network. The specific algorithms a modem would use to report metrics on multicast (or logical) destinations is outside the scope of this specification, and is left to specific implementations to decide.

1.1 Requirements

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

2. Assumptions

Routers and modems that exist as part of the same node (e.g., that are locally connected) can utilize a discovery technique to locate each other, thus avoiding a-priori configuration. The modem is responsible for initialing the discovery process, using the Peer Discovery message.

DLEP utilizes a session-oriented paradigm. A router and modem form a session by completing the discovery process. This router-modem session persists unless or until it either (1) times out, based on the timeout values supplied, or (2) is explicitly torn down by one of the participants. Note that while use of timers in DLEP is OPTIONAL, it is strongly recommended that implementations choose to run with timers enabled.

DLEP assumes that participating modems, and their physical links, act as a transparent bridge. Specifically, the assumption is that the destination MAC address for data traffic in any frame emitted by the router should be the MAC address of a device in the remote node. DLEP also assumes that MAC addresses are unique within the context of the router-modem session.

This document refers to a remote node as a "Destination". Destinations can be identified by either the router or the modem, and represent a specific destination (e.g., an address) that exists on the link(s) managed by the modem. A destination MUST contain a MAC address, it MAY optionally include a Layer 3 address (or addresses). Destinations MAY refer either to physical devices in the network, or to logical destinations, as in a multicast group. As "destinations" are discovered, DLEP routers and modems build an information base on destinations accessible via the modem. Changes in link characteristics MAY then be reported as being "modem-wide" (effecting ALL destinations accessed via the modem) or MAY be neighbor (destination) specific.

The DLEP messages concerning destinations thus become the way for routers and modems to maintain, and notify each other about, an information base representing the physical and logical (e.g., multicast) destinations accessible via the modem device. The information base would contain addressing information (e.g., MAC

address, and OPTIONALLY, Layer 3 addresses), link characteristics (metrics), and OPTIONALLY, flow control information (credits).

DLEP assumes that security on the session (e.g. authentication of session partners, encryption of traffic, or both) is dealt with by the underlying transport mechanism (e.g., by using a transport such as TLS [[TLS](#)]).

This document specifies an implementation of the DLEP messages and data items running over the TCP transport, utilizing a well-known TCP Port number. It is assumed that DLEP running over other transport mechanisms would be documented separately.

3. Credits

DLEP includes an OPTIONAL credit-windowing scheme analogous to the one documented in [[RFC5578](#)]. In this scheme, traffic between the router and modem is treated as two unidirectional windows. This document identifies these windows as the "Modem Receive Window", or MRW, and the "Router Receive Window", or RRW.

If credits are used, they MUST be granted by the receiver on a given window - that is, on the "Modem Receive Window" (MRW), the modem is responsible for granting credits to the router, allowing it (the router) to send data to the modem. Likewise, the router is responsible for granting credits on the RRW, which allows the modem to send data to the router.

DLEP expresses all credit data in number of octets. The total number of credits on a window, and the increment to add to a grant, are always expressed as a 64-bit unsigned quantity.

If used, credits are managed on a neighbor-specific basis; that is, separate credit counts are maintained for each neighbor requiring the service. Credits do not apply to the DLEP session that exists between routers and modems.

4. Metrics

DLEP includes the ability for the router and modem to communicate metrics that reflect the characteristics (e.g. datarate, latency) of the variable-quality link in use. DLEP does NOT specify how a given metric value is to be calculated, rather, the protocol assumes that metrics have been calculated with a "best effort", incorporating all pertinent data that is available to the modem device.

As mentioned in the introduction section of this document, metrics

have to be used within a context - for example, metrics to a unicast address in the network. DLEP allows for metrics to be sent within two contexts - metrics for a specific destination within the network (e.g., a specific router), and "modem-wide" (those that apply to all destinations accessed via the modem). Metrics are further subdivided into transmit and receive metrics. Metrics supplied on DLEP Peer signals are, by definition, modem-wide; metrics supplied on Destination messages signals are, by definition, used for the specific neighbor only.

DLEP modem implementations MUST announce all supported metric items, and provide default values for those metrics, in the Peer Initialization message. In order to introduce a new metric type, DLEP modem implementations MUST terminate the session with the router (via the Peer Terminate message), and re-establish the session.

It is left to implementations to choose sensible default values based on their specific characteristics. Modems having static (non-changing) link metric characteristics MAY report metrics only once for a given neighbor (or once on a modem-wide basis, if all connections via the modem are of this static nature).

The approach of allowing for different contexts for metric data increases both the flexibility and the complexity of using metric data. This document details the mechanism whereby the data is transmitted, however, the specific algorithms (precedence, etc) for utilizing the dual-context metrics is out of scope and not addressed by this document.

5. Extensions to DLEP

While this draft represents the best efforts of the co-authors, and the working group, to be functionally complete, it is recognized that extensions to DLEP will in all likelihood be necessary as more link types are utilized. To allow for future innovation, the draft allocates numbering space for experimental implementations of both signals and data items.

DLEP implementations MUST be capable of parsing and acting on the mandatory messages and data items as documented in this specification. DLEP messages/data items that are optional, or are in the experimental numbering range SHOULD be silently dropped by an implementation if they are not understood.

The intent of the optional messages and data items, as well as the experimental numbering space, is to allow for further development of DLEP protocol features and function. Having experimental space reserved for both signals and data items gives maximum flexibility

for extending the protocol as conditions warrant. For example, experimental data items could be used by implementations to send additional metrics. A combination of experimental messages, and associated data items, could be used to implement new flow control schemes. If subsequent research and development define new features and function, then it should be standardized either as an update to this document, or as an additional stand-alone specification.

6. Normal Session Flow

Normal session flow is slightly different, depending on whether the implementation represents a modem or a router, and whether discovery techniques are used. The normal flow by DLEP partner type is:

6.1 DLEP Modem session flow - Discovery case

If the DLEP modem implementation is utilizing the optional discovery mechanism, then the implementation will initialize a UDP socket, binding it to an arbitrary port. This UDP socket is used to send the Peer Discovery message to the DLEP link-local multicast address and port (TBD). The implementation then waits on receipt of a Peer Offer message, which **MUST** contain the unicast address and port for TCP-based communication with a DLEP router. The Peer Offer message **MAY** contain multiple address/port combinations. If more than one address/port combination is in the Peer Offer, the DLEP modem implementation **SHOULD** consider the list to be in priority sequence, with the "most desired" address/port combination listed first. However, modem implementations **MAY** use their own heuristics to determine the best address/port combination. At this point, the modem implementation **MAY** either destroy the UDP socket, or continue to issue Peer Discovery messages to the link-local address/port combination. In either case, the TCP session initialization occurs as in the configured case.

6.2 DLEP Modem session flow - Configured case

When a DLEP modem implementation has the address and port information for a TCP connection to the router (obtained either via configuration or via the discovery process described above), the modem will initialize and bind a TCP socket. This socket is used to connect to the DLEP router software. After a successful TCP connect, the modem implementation **MUST** issue a Peer Initialization message to the DLEP router. The Peer Initialization message **MUST** contain TLVs for **ALL** supported metrics from this modem (e.g. all **MANDATORY** metrics plus all **OPTIONAL** metrics supported by the implementation), along with the default values of those metrics. After sending the Peer Initialization, the modem implementation should wait for receipt of a

Peer Initialization ACK message from the router. Receipt of the Peer Initialization ACK indicates that the router has received and processed the Peer Initialization, and the session MUST transition to the "in session" state. At this point, messages regarding destinations in the network, and/or Peer Update messages, can flow on the DLEP session between modem and router. The "in session" state is maintained until one of the following conditions occur:

- o The session is explicitly terminated (using Peer Termination), or
- o The session times out, based on supplied timeout values.

6.3 DLEP Router session flow

DLEP router implementations MUST support the discovery mechanism. Therefore, the normal flow is as follows:

The implementation will initialize a UDP socket, binding that socket to the DLEP link-local multicast address (TBD) and the DLEP well-known port number (also TBD). The implementation will then initialize a TCP socket, on a unicast address and port. This socket is used to listen for incoming TCP connection requests.

When the router implementation receives a Peer Discovery message on the UDP socket, it responds by issuing a Peer Offer message to the sender of the Peer Discovery. The Peer Offer message MUST contain the unicast address and port of the TCP listen socket, described above. A DLEP router implementation MAY respond with ALL address/port combinations that have an active TCP listen posted. If multiple address/port combinations are listed, the receiver of the Peer Offer MAY connect on any available address/port pair. Anything other than Peer Discovery messages received on the UDP socket MUST be silently dropped.

When the DLEP router implementation accepts a connection via TCP, it will wait for receipt of a Peer Initialization message. The received Peer Initialization MUST contain metric TLVs for ALL mandatory metrics, and MUST contain metric TLVs for ANY optional metrics supported by the modem. If a new metric is to be introduced, the DLEP session between router and modem MUST be terminated and restarted, and the new metric described in a Peer Initialization message.

6.4 Common Session Flow

In order to maintain the session between router and modem, periodic "Heartbeat" messages SHOULD be exchanged. These messages are intended to keep the session alive, and to verify bidirectional connectivity

between the two participants. DLEP also provides for an OPTIONAL Peer Update message, intended to communicate some change in status (e.g., a change of layer 3 address parameters, or a modem-wide link change).

In addition to the messages above, the participants will transmit DLEP messages concerning destinations in the network. These messages trigger creation/maintenance/deletion of destinations in the information base of the recipient. For example, a modem will inform its attached router of the presence of a new destination via the "Destination Up" signal. Receipt of a Destination Up causes the router to allocate the necessary resources, creating an entry in the information base with the specifics (e.g., MAC Address, Latency, Data Rate, etc) of the neighbor. The loss of a destination is communicated via the "Neighbor Down" signal, and changes in status to the destination (e.g. varying link quality, or addressing changes) are communicated via the "Neighbor Update" signal. The information on a given neighbor will persist in the router's information base until (1) a "Neighbor Down" is received, indicating that the modem has lost contact with the remote node, or (2) the router/modem session terminates, indicating that the router has lost contact with its own local modem.

Again, metrics can be expressed within the context of a specific neighbor via the Neighbor Update message, or on a modem-wide basis via the Peer Update message. In cases where metrics are provided on the router/modem session, the receiver MUST propagate the metrics to all destinations in its information base that are accessed via the originator. A DLEP participant MAY send metrics both in a router/modem session context (via the Peer Update message) and a specific neighbor context (via Neighbor Update) at any time. The heuristics for applying received metrics is left to implementations.

In addition to receiving metrics about the link, DLEP provides an OPTIONAL signal allowing a router to request a different datarate, or latency, from the modem. This signal is referred to as the Link Characteristics Message, and gives the router the ability to deal with requisite increases (or decreases) of allocated datarate/latency in demand-based schemes in a more deterministic manner.

7. Mandatory Signals and Data Items

The following DLEP signals are considered core to the specification; implementations MUST support these signals, and the associated data items, in order to be considered compliant:

Signal	Data Items
=====	=====
Peer Discovery (Modem Only)	DLEP Version

Peer Offer (Router Only)	DLEP Version IPv4 or IPv6 address(es) DLEP Port
Peer Initialization	DLEP Version Maximum Data Rate (Receive) Maximum Data Rate (Transmit) Current Data Rate (Receive) Current Data Rate (Transmit) Latency Relative Link Quality (Receive) Relative Link Quality (Transmit)
Peer Initialization ACK	Status
Peer Termination	None
Peer Termination ACK	Status
Destination Up	MAC Address Maximum Data Rate (Receive) Maximum Data Rate (Transmit) Current Data Rate (Receive) Current Data Rate (Transmit) Latency Relative Link Quality (Receive) Relative Link Quality (Transmit)
Destination Update	MAC Address Maximum Data Rate (Receive) Maximum Data Rate (Transmit) Current Data Rate (Receive) Current Data Rate (Transmit) Latency Relative Link Quality (Receive) Relative Link Quality (Transmit)
Destination Down	MAC Address

All other DLEP signals and data items are OPTIONAL. Implementations MAY choose to provide them. Implementations that do not support optional signals and data items SHOULD parse, and silently drop, all unsupported messages and/or data items.

8. Generic DLEP Message Definition

The Generic DLEP Message consists of a sequence of TLVs. The first TLV represents the signal being communicated (e.g., a "Destination


```

|-----|type string len|-----|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - Length of peer type string.

Peer Type String - Non-Null terminated string, using UTF-8 encoding.
For example, a satellite modem might set this variable to 'Satellite terminal'.

9.4 MAC Address

The MAC address TLV MUST appear in all destination-oriented signals (e.g. Destination Up, Destination Up ACK, Destination Down, Destination Down ACK, Destination Update, Link Characteristics Request, and Link Characteristics ACK). The MAC Address TLV contains the address of the destination on the remote node. The MAC address MAY be either a physical or a virtual destination. Examples of a virtual destination would be a multicast MAC address, or the broadcast MAC (0xFFFFFFFFFFFF).

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD |Length = 6   |               MAC Address               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               MAC Address               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

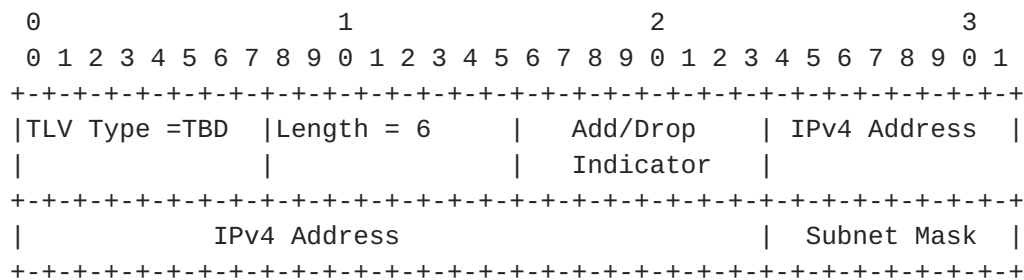
Length - 6

MAC Address - MAC Address of the destination (either physical or virtual).

9.5 IPv4 Address

The IPv4 Address TLV is an OPTIONAL TLV. If supported, it MAY appear in Destination Up, Destination Update, and Peer Update messages. When included in Destination messages, the IPv4 Address TLV contains the IPv4 address of the destination, as well as a subnet mask value. In the Peer Update message, it contains the IPv4 address of the originator of the message. In either case, the TLV also contains an indication of whether this is a new or existing address, or is a deletion of a previously known address.

The IPv4 Address TLV contains the following fields:



TLV Type - TBD

Length - 6

Add/Drop - Value indicating whether this is a new or existing address (0x01), or a withdrawal of an address (0x02).

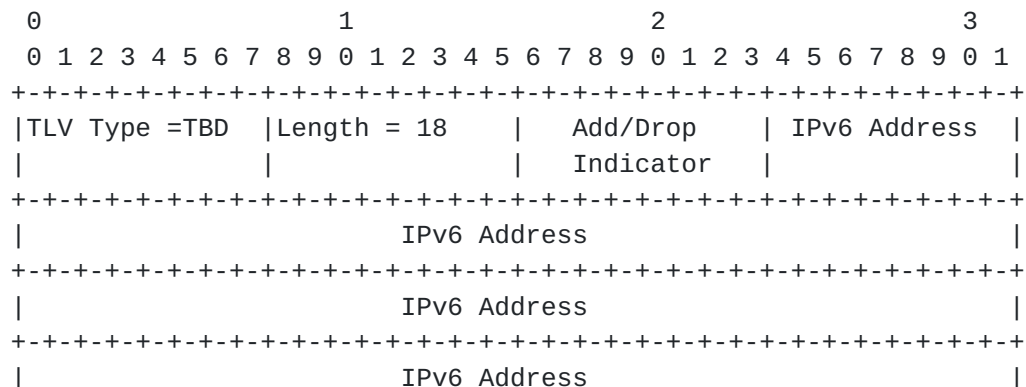
IPv4 Address - The IPv4 address of the destination or peer.

Subnet Mask - A subnet mask (0-32) to be applied to the IPv4 address.

9.6 IPv6 Address

The IPv6 Address TLV is an OPTIONAL TLV. If supported, it MAY be used in the Destination Up, Destination Update, Peer Discovery, and Peer Update Messages. When included in Destination messages, this data item contains the IPv6 address of the destination. In the Peer Discovery and Peer Update, it contains the IPv6 address of the originating peer. In either case, the data item also contains an indication of whether this is a new or existing address, or is a deletion of a previously known address, as well as a subnet mask.

The IPv6 Address TLV contains the following fields:




```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv6 Address                               | Subnet Mask |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 18

Add/Drop - Value indicating whether this is a new or existing address (0x01), or a withdrawal of an address (0x02).

IPv6 Address - IPv6 Address of the destination or peer.

Subnet Mask - A subnet mask value (0-128) to be applied to the Ipv6 address.

9.7 Maximum Data Rate (Receive)

The Maximum Data Rate Receive (MDRR) TLV is used in Destination Up, Destination Update, Peer Discovery, Peer Update, and Link Characteristics ACK Messages to indicate the maximum theoretical data rate, in bits per second, that can be achieved while receiving data on the link. When metrics are reported via the messages listed above, the maximum data rate receive MUST be reported.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| TLV Type =TBD | Length = 8 | MDRR (bps) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| MDRR (bps) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| MDRR (bps) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 8

Maximum Data Rate Receive - A 64-bit unsigned number, representing the maximum theoretical data rate, in bits per second (bps), that can be achieved while receiving on the link.

9.8 Maximum Data Rate (Transmit)

[illegible]

Length - 8

9.9 Current Data Rate (Receive)

The Current Data Rate Receive TLV contains the following fields:

[illegible]


```

|                                     CDRR (bps)                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 8

Current Data Rate Receive - A 64-bit unsigned number, representing the current data rate, in bits per second, that is currently be achieved while receiving traffic on the link. When used in the Link Characteristics Request, CDRR represents the desired receive rate, in bits per second, on the link. If there is no distinction between current and maximum receive data rates, current data rate receive SHOULD be set equal to the maximum data rate receive.

[9.10](#) Current Data Rate (Transmit)

The Current Data Rate Receive (CDRT) TLV is used in Destination Up, Destination Update, Peer Discovery, Peer Update, Link Characteristics Request, and Link Characteristics ACK messages to indicate the rate at which the link is currently operating for transmitting traffic. Current Data Rate Transmit is a MANDATORY data item. In the case of the Link Characteristics Request, CDRT represents the desired transmit data rate for the link. When metrics are reported via the messages above (e.g. Destination Update), the current data rate transmit MUST be reported.

The Current Data Rate Transmit TLV contains the following fields:

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD |TLV Flags=0x10 |Length = 8      |CDRT (bps)      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     CDRT (bps)                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     CDRT (bps)                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 8

Current Data Rate Transmit - A 64-bit unsigned number, representing

the current data rate, in bits per second, that is currently be achieved while transmitting traffic on the link. When used in the Link Characteristics Request, CDRT represents the desired transmit rate, in bits per second, on the link. If there is no distinction between current and maximum transmit data rates, current data rate transmit MUST be set equal to the maximum data rate transmit.

9.11 Expected Forwarding Time

The Expected Forwarding Time (EFT) TLV is is an OPTIONAL data item. If supported, it MAY be used in Destination Up, Destination Update, Peer Discovery, and Peer Update messages to indicate the typical latency between the arrival of a given packet at the transmitting device and the reception of the packet at the other end of the link. EFT combines transmission time, idle time, waiting time, freezing time, and queuing time to the degree that those values are meaningful to a given transmission medium.

The Expected Forwarding Time TLV contains the following fields:

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|TLV Type =TBD  |Length = 4      |              EFT (ms)              |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|              EFT (ms)              |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

TLV Type - TBD

Length - 4

EFT - A 32-bit unsigned number, representing the expected forwarding time, in milliseconds, on the link.

9.12 Latency

The Latency TLV is an MANDATORY data item. It is used in Peer Initialization, Destination Up, Destination Update, Peer Discovery, Peer Update, Link Characteristics Request, and Link Characteristics ACK messages to indicate the amount of latency on the link, or in the case of the Link Characteristics Request, to indicate the maximum latency required on the link.


```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|TLV Type =TBD |Length = 2      |          Latency (ms)          |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

TLV Type - TBD

Length - 2

Latency - A 16-bit unsigned value, representing the transmission delay that a packet encounters as it is transmitted over the link. In Destination Up, Destination Update, and Link Characteristics ACK, this value is reported as delay, in milliseconds. The calculation of latency is implementation dependent. For example, the latency may be a running average calculated from the internal queuing. If a device cannot calculate latency, this TLV SHOULD NOT be issued. In the Link Characteristics Request Message, this value represents the maximum delay, in milliseconds, expected on the link.

9.13 Resources (Receive)

The Receive Resources TLV is an OPTIONAL data item. If supported, it is used in Destination Up, Destination Update, Peer Discovery, Peer Update, and Link Characteristics ACK messages to indicate a percentage (0-100) amount of resources (e.g. battery power), committed to receiving data, remaining on the originating peer.

The Resources TLV contains the following fields:

```

      0               1               2
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|TLV Type =TBD |Length = 1      | Rcv Resources|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

TLV Type - TBD

Length - 1

Receive Resources - A percentage, 0-100, representing the amount of remaining resources, such as battery power, allocated to receiving data. If a device cannot calculate receive resources, this TLV SHOULD NOT be issued.

9.14 Resources (Transmit)

The Transmit Resources TLV is an OPTIONAL data item. If supported, it is used in Destination Up, Destination Update, Peer Discovery, Peer Update, and Link Characteristics ACK messages to indicate a percentage (0-100) amount of resources (e.g. battery power), committed to transmitting data, remaining on the originating peer.

The Resources TLV contains the following fields:

```

      0                   1                   2
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD  |Length = 1      | Xmt Resources|
+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 1

Transmit Resources - A percentage, 0-100, representing the amount of remaining resources, such as battery power, allocated to transmitting data. If the transmit resources cannot be calculated, then the TLV SHOULD NOT be issued.

9.15 Relative Link Quality (Receive)

The Relative Link Quality Receive (RLQR) TLV is a MANDATORY data item. If supported, it is used in Peer Initialization, Destination Up, Destination Update, Peer Discovery, Peer Update, and Link Characteristics ACK messages to indicate the quality of the link for receiving data as calculated by the originating peer.

The Relative Link Quality (Receive) TLV contains the following fields:

```

      0                   1                   2
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD  |Length = 1      |RCV Rel. Link  |
|                |                |Quality (RLQR) |
+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD


```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD |Length = 1      |      Code      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 1

Termination Code - 0 = Success, Non-zero = Failure. Specific values of a non-zero termination code depend on the operation requested (e.g. Destination Up, Destination Down, etc).

9.18 Heartbeat Interval

The Heartbeat Interval TLV is a MANDATORY TLV. It MUST be sent during Peer Initialization to indicate the desired Heartbeat timeout window. The router MUST either accept the timeout interval supplied by the modem, or reject the Peer Initialization, and close the socket. Implementations MUST implement heuristics such that DLEP signals sent/received reset the timer interval.

The Interval is used to specify a period (in seconds) for Heartbeat Messages (See [Section 23](#)). By specifying an Interval value of 0, implementations MAY indicate the desire to disable Heartbeat messages entirely (e.g., the Interval is set to an infinite value), however, it is strongly recommended that implementations use non 0 timer values.

A DLEP session will be considered inactive, and MUST be torn down, by an implementation detecting that two (2) Heartbeat intervals have transpired without receipt of any DLEP messages.

The Heartbeat Interval TLV contains the following fields:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD |Length = 2      |      Interval      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 2

Interval - 0 = Do NOT use heartbeats on this peer-to-peer session. Non-zero = Interval, in seconds, for

heartbeat messages.

9.19 Link Characteristics ACK Timer

The Link Characteristics ACK Timer TLV is an OPTIONAL TLV. If supported, it MAY be sent during Peer Initialization to indicate the desired number of seconds to wait for a response to a Link Characteristics Request. If a router receives this TLV from a modem during Peer Discovery, the router MUST either accept the timeout value, or reject the Peer Discovery. If this TLV is omitted, implementations supporting the Link Characteristics Request SHOULD choose a default value.

The Link Characteristics ACK Timer TLV contains the following fields:

```

      0                   1                   2
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD  |Length = 1      | Interval          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 1

Interval - 0 = Do NOT use timeouts for Link Characteristics requests on this router/modem session. Non-zero = Interval, in seconds, to wait before considering a Link Characteristics Request has been lost.

9.20 Credit Window Status

The Credit Window Status TLV is an OPTIONAL TLV. If credits are supported by the DLEP participants (both the router and the modem), the Credit Window Status TLV MUST be sent by the participant receiving a Credit Grant Request for a given destination.

The Credit Window Status TLV contains the following fields:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD  |Length = 16     | Modem Receive Window Value      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                |                | Modem Receive Window Value      |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Modem Receive Window Value | Router Receive Window Value |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```



```

|                               Router Receive Window Value                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Router Receive Window Value |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 16

Modem Receive Window Value - A 64-bit unsigned number, indicating the current (or initial) number of credits available on the Modem Receive Window.

Router Receive Window Value - A 64-bit unsigned number, indicating the current (or initial) number of credits available on the Router Receive Window.

[9.21](#) Credit Grant Request

The Credit Grant Request TLV is an OPTIONAL TLV. If credits are supported, the Credit Grant Request TLV is sent from a DLEP participant to grant an increment to credits on a window. The Credit Grant TLV is sent as a data item in either the Destination Up or Destination Update messages. The value in a Credit Grant TLV represents an increment to be added to any existing credits available on the window. Upon successful receipt and processing of a Credit Grant TLV, the receiver MUST respond with a message containing a Credit Window Status TLV to report the updated aggregate values for synchronization purposes.

In the Destination Up message, when credits are desired, the originating peer MUST set the initial credit value of the window it controls (e.g. the Modem Receive Window, or Router Receive Window) to an initial, non-zero value. If the receiver of a Destination Up message with a Credit Grant Request TLV supports credits, the receiver MUST either reject the use of credits, via a Destination Up ACK response with the correct Status TLV, or set the initial value from the data contained in the Credit Window Status TLV. If the initialization completes successfully, the receiver MUST respond to the Destination Up message with a Destination Up ACK message that contains a Credit Window Status TLV, initializing its receive window.

The Credit Grant TLV contains the following fields:

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD |Length = 8      |          Credit Increment          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Credit Increment          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|          Credit Increment          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

Length - 8

Reserved - A 64-bit unsigned number representing the additional credits to be assigned to the credit window. Since credits can only be granted by the receiver on a window, the applicable credit window (either the MRW or the RRW) is derived from the sender of the grant. The Credit Increment MUST NOT cause the window to overflow; if this condition occurs, implementations MUST set the credit window to the maximum value contained in a 64-bit quantity.

9.22 Credit Request

The Credit Request TLV is an OPTIONAL TLV. If credits are supported, the Credit Request TLV MAY be sent from either DLEP participant, via a Destination Update signal, to indicate the desire for the partner to grant additional credits in order for data transfer to proceed on the session. If the corresponding Destination Up message for this session did NOT contain a Credit Window Status TLV, indicating that credits are to be used on the session, then the Credit Request TLV MUST be rejected by the receiver via a Destination Update ACK message.

The Credit Request TLV contains the following fields:

```

      0                      1                      2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|TLV Type =TBD |Length = 0      | Reserved, MUST|
|          |          | be set to 0  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

TLV Type - TBD

TBD	Peer Termination
TBD	Peer Termination ACK
TBD	Destination Up
TBD	Destination Up ACK
TBD	Destination Down
TBD	Destination Down ACK
TBD	Destination Update
TBD	Heartbeat
TBD	Link Characteristics Request
TBD	Link Characteristics ACK

10.2 Peer Discovery Message

The Peer Discovery Message is sent by a modem to discover DLEP routers in the network. The Peer Offer message is required to complete the discovery process. Implementations MAY implement their own retry heuristics in cases where it is determined the Peer Discovery Message has timed out.

Given the packet format described in [section 11](#), the initial TLV Type value is set to DLEP_PEER_DISCOVERY (value TBD).

There are NO Data Item TLVs associated with the Peer Discovery message.

10.3 Peer Offer Message

The Peer Offer Message is sent by a DLEP router in response to a Peer Discovery Message. Upon receipt, and processing, of a Peer Offer message, the modem responds by issuing a TCP connect to the address/port combination specified in the received Peer Offer.

The Peer Offer message MUST be sent to the unicast address of the originator of Peer Discovery.

To construct a Peer Offer message, the initial TLV type value is set to DLEP_PEER_OFFER (value TBD). The signal TLV is then followed by all MANDATORY Data Item TLVs, then by any OPTIONAL Data Item TLVs the implementation supports:

Mandatory Data Item TLVs:

- DLEP Version
- Heartbeat Interval
- At least one (1) IPv4 or IPv6 Address TLV
- DLEP Port

Optional Data Item TLVs:

- Peer Type

- Status

10.4 Peer Initialization Message

The Peer Initialization message is sent by a modem to start the DLEP TCP session. It is sent by the modem after a TCP connect to the address/port combination in a received Peer Offer, or to an address/port obtained from a-priori configuration.

All supported metric data items MUST be included in the Peer Initialization message, with default values to be used on a "modem-wide" basis. This can be viewed as the modem "declaring" all supported metrics at DLEP session initialization. Receipt of any DLEP message containing a metric data item NOT included in Peer Initialization MUST be treated as an error, resulting in termination of the DLEP session between router and modem.

To construct a Peer Initialization message, the initial TLV type value is set to DLEP_PEER_INIT (value TBD). The signal TLV is then followed by the required data items:

Mandatory Data Item TLVs:

- DLEP Version
- Heartbeat Interval
- Maximum Data Rate Receive
- Maximum Data Rate Transmit
- Current Data Rate Receive
- Current Data Rate Transmit
- Latency
- Relative Link Quality Receive
- Relative Link Quality Transmit

Optional Data Item TLVs:

- Peer Type

Note that metric data items MUST be supplied with the Peer Initialization, in order to populate default metric values. If, at any time, metrics are reported that were NOT in Peer Initialization, the receiving DLEP peer MUST treat this as a fatal error requiring termination of the DLEP session.

10.5 Peer Initialization ACK Message

The Peer Initialization ACK message is a MANDATORY message, sent in response to a received Peer Initialization message. The Peer Initialization ACK message completes the TCP-level DLEP session establishment; the sender of the message should transition to an "in-session" state when the message is sent, and the receiver should transition to the "in-session" state upon receipt (and successful parsing) of Peer Initialization ACK.

To construct a Peer Initialization ACK message, the initial TLV type value is set to DLEP_PEER_INIT_ACK (value TBD). The signal TLV is then followed by the required data items:

Mandatory Data Item TLVs:

- Status

Optional Data Item TLVs:

- Peer Type

10.6 Peer Update Message

The Peer Update message is an OPTIONAL message, sent by a DLEP peer to indicate local Layer 3 address changes, or for metric changes on a modem-wide basis. For example, addition of an IPv4 address to the router MAY prompt a Peer Update message to its attached DLEP modems. Also, a modem that changes its Maximum Data Rate for all destinations MAY reflect that change via a Peer Update Message to its attached router(s).

Concerning Layer 3 addresses, if the modem is capable of understanding and forwarding this information (via proprietary mechanisms), the address update would prompt any remote DLEP modems (DLEP-enabled modems in a remote node) to issue a "Destination Update" message to their local routers with the new (or deleted) addresses. Modems that do not track Layer 3 addresses SHOULD silently parse and ignore the Peer Update Message. Modems that track Layer 3 addresses MUST acknowledge the Peer Update with a Peer Update ACK message. Routers receiving a Peer Update with metric changes MUST apply the new metric to all destinations (remote nodes) accessible via the modem. Supporting implementations are free to employ heuristics to retransmit Peer Update messages. The sending of Peer Update Messages for Layer 3 address changes SHOULD cease when a either participant (router or modem) determines that the other implementation does NOT support Layer 3 address tracking.

If metrics are supplied with the Peer Update message (e.g. Maximum Data Rate), these metrics are considered to be modem-wide, and therefore MUST be applied to all destinations in the information base associated with the router/modem session.

To construct a Peer Update message, the initial TLV type value is set to DLEP_PEER_UPDATE (value TBD). The Signal TLV is followed by any OPTIONAL Data Item TLVs.

Optional Data Item TLVs:

- IPv4 Address
- IPv6 Address
- Maximum Data Rate (Receive)

- Maximum Data Rate (Transmit)
- Current Data Rate (Receive)
- Current Data Rate (Transmit)
- Latency
- Expected Forwarding Time
- Resources (Receive)
- Resources (Transmit)
- Relative Link Quality (Receive)
- Relative Link Quality (Transmit)

10.7 Peer Update ACK Message

The Peer Update ACK message is an OPTIONAL message, and is sent by implementations supporting Layer 3 address tracking and/or modem-wide metrics to indicate whether a Peer Update Message was successfully processed. If the Peer Update ACK is issued, it MUST contain a Status data item, indicating the success or failure of processing the received Peer Update.

To construct a Peer Update ACK message, the initial TLV type value is set to DLEP_PEER_UPDATE_ACK (value TBD). The Status data item TLV is placed in the packet next, completing the Peer Update ACK.

Mandatory Data Item TLVs:

- Status

Note that there are NO OPTIONAL data item TLVs specified for this message.

10.8 Peer Termination Message

The Peer Termination Message is sent by a DLEP participant when the router/modem session needs to be terminated. Implementations receiving a Peer Termination message MUST send a Peer Termination ACK message to confirm the termination process. The sender of a Peer Termination message is free to define its heuristics in event of a timeout. The receiver of a Peer Termination Message MUST release all resources allocated for the router/modem session, and MUST eliminate all destinations in the information base accessible via the router/modem pair represented by the session. Router and modem state machines are returned to the "discovery" state. No Destination Down messages are sent.

To construct a Peer Termination message, the initial TLV type value is set to DLEP_PEER_TERMINATION (value TBD). The Signal TLV is followed by any OPTIONAL Data Item TLVs the implementation supports:

Optional Data Item TLVs:

- Status

10.9 Peer Termination ACK Message

The Peer Termination Message ACK is sent by a DLEP peer in response to a received Peer Termination order. Receipt of a Peer Termination ACK message completes the teardown of the router/modem session.

To construct a Peer Termination ACK message, the initial TLV type value is set to DLEP_PEER_TERMINATION_ACK (value TBD). The Identification data item TLV is placed in the packet next, followed by any OPTIONAL TLVs the implementation supports:

Optional Data Item TLVs:

- Status

10.10 Destination Up Message

A DLEP participant sends the Destination Up message to report that a new destination has been detected. A Destination Up ACK Message is required to confirm a received Destination Up. A Destination Up message can be sent either by the modem, to indicate that a new remote node has been detected, or by the router, to indicate the presence of a new logical destination (e.g., a Multicast group) exists in the network.

The sender of the Destination Up Message is free to define its retry heuristics in event of a timeout. When a Destination Up message is received and successfully parsed, the receiver should add knowledge of the new destination to its information base, indicating that the destination is accessible via the modem/router pair.

To construct a Destination Up message, the initial TLV type value is set to DLEP_Destination_UP (value TBD). The MAC Address data item TLV is placed in the packet next, followed by any supported OPTIONAL Data Item TLVs into the packet:

Optional Data Item TLVs:

- IPv4 Address
- IPv6 Address
- Maximum Data Rate (Receive)
- Maximum Data Rate (Transmit)
- Current Data Rate (Receive)
- Current Data Rate (Transmit)

- Latency
- Expected Forwarding Time
- Resources (Receive)
- Resources (Transmit)
- Relative Link Factor (Receive)
- Relative Link Factor (Transmit)
- Credit Window Status

10.11 Destination Up ACK Message

A DLEP participant sends the Destination Up ACK Message to indicate whether a Destination Up Message was successfully processed.

To construct a Destination Up ACK message, the initial TLV type value is set to DLEP_Destination_UP_ACK (value TBD). The MAC Address data item TLV is placed in the packet next, containing the MAC address of the DLEP destination. The implementation would then place any supported OPTIONAL Data Item TLVs into the packet:

Optional Data Item TLVs:

- Credit Window Status

10.12 Destination Down Message

A DLEP peer sends the Destination Down message to report when a destination (a remote node or a multicast group) is no longer reachable. The Destination Down message MUST contain the MAC Address data item TLV. Other TLVs as listed are OPTIONAL, and MAY be present if an implementation supports them. A Destination Down ACK Message MUST be sent by the recipient of a Destination Down message to confirm that the relevant data has been removed from the information base. The sender of the Destination Down message is free to define its retry heuristics in event of a timeout.

To construct a Destination Down message, the initial TLV type value is set to DLEP_Destination_DOWN (value TBD). The signal TLV is followed by the mandatory MAC Address data item TLV.

Note that there are NO OPTIONAL data item TLVs for this message.

10.13 Destination Down ACK Message

A DLEP participant sends the Destination Down ACK Message to indicate whether a received Destination Down Message was successfully processed. If successfully processed, the sender of the ACK MUST have removed all entries in the information base that pertain to the referenced destination. As with the Destination Down message, there are NO OPTIONAL Data Item TLVs defined for the Destination Down ACK

message.

To construct a Destination Down message, the initial TLV type value is set to DLEP_Destination_DOWN_ACK (value TBD). The mandatory data item TLVs follow:

- MAC Address Data item
- Status data item

10.14 Destination Update Message

A DLEP participant sends the Destination Update message when it detects some change in the information base for a given destination (remote node or multicast group). Some examples of changes that would prompt a Destination Update message are:

- Change in link metrics (e.g., Data Rates)
- Layer 3 addressing change (for implementations that support it)

To construct a Destination Update message, the initial TLV type value is set to DLEP_Destination_UPDATE (value TBD). Following the signal TLV are the mandatory Data Item TLVs:

MAC Address data item TLV

After placing the mandatory data item TLV into the packet, the implementation would place any supported OPTIONAL data item TLVs. Possible OPTIONAL data item TLVs are:

- IPv4 Address
- IPv6 Address
- Maximum Data Rate (Receive)
- Maximum Data Rate (Transmit)
- Current Data Rate (Receive)
- Current Data Rate (Transmit)
- Latency
- Resources (Receive)
- Resources (Transmit)
- Relative Link Quality (Receive)
- Relative Link Quality (Transmit)
- Credit Window Status
- Credit Grant
- Credit Request

10.15 Heartbeat Message

A Heartbeat Message is sent by a DLEP participant every N seconds, where N is defined in the "Heartbeat Interval" field of the discovery message. Note that implementations setting the Heartbeat Interval to 0 effectively set the interval to an infinite value, therefore, in those cases, this message would NOT be sent.

The message is used by participants to detect when a DLEP session partner (either the modem or the router) is no longer communicating. Participants SHOULD allow two (2) heartbeat intervals to expire with no traffic on the router/modem session before initiating DLEP session termination procedures.

To construct a Heartbeat message, the initial TLV type value is set to DLEP_PEER_HEARTBEAT (value TBD). The signal TLV is followed by the mandatory Heartbeat Interval/Threshold data item.

Note that there are NO OPTIONAL data item TLVs for this message.

10.16 Link Characteristics Request Message

The Link Characteristics Request Message is an OPTIONAL message, and is sent by the router to request that the modem initiate changes for specific characteristics of the link. The request can reference either a real (e.g., a remote node), or a logical (e.g., a multicast group) destination within the network.

The Link Characteristics Request message contains either a Current Data Rate (CDRR or CDRT) TLV to request a different datarate than what is currently allocated, a Latency TLV to request that traffic delay on the link not exceed the specified value, or both. A Link Characteristics ACK Message is required to complete the request. Implementations are free to define their retry heuristics in event of a timeout. Issuing a Link Characteristics Request with ONLY the MAC Address TLV is a mechanism a peer MAY use to request metrics (via the Link Characteristics ACK) from its partner.

To construct a Link Characteristics Request message, the initial TLV type value is set to DLEP_Destination_LINK_CHAR_REQ (value TBD). Following the signal TLV is the mandatory Data Item TLV:

MAC Address data item TLV

After placing the mandatory data item TLV into the packet, the implementation would place any supported OPTIONAL data item TLVs. Possible OPTIONAL data item TLVs are:

Current Data Rate - If present, this value represents the NEW (or

unchanged, if the request is denied) Current Data Rate in bits per second (bps).

- Latency
- If present, this value represents the maximum desired latency (e.g., it is a not-to-exceed value) in milliseconds on the link.

10.17 Link Characteristics ACK Message

The Link Characteristics ACK message is an OPTIONAL message, and is sent by modems supporting it to the router letting the router know the success or failure of a requested change in link characteristics.

The Link Characteristics ACK message SHOULD contain a complete set of metric data item TLVs. It MUST contain the same TLV types as the request. The values in the metric data item TLVs in the Link Characteristics ACK message MUST reflect the link characteristics after the request has been processed.

To construct a Link Characteristics Request ACK message, the initial TLV type value is set to DLEP_Destination_LINK_CHAR_ACK (value TBD). Following the signal TLV is the mandatory Data Item TLV:

MAC Address data item TLV

After placing the mandatory data item TLV into the packet, the implementation would place any supported OPTIONAL data item TLVs. Possible OPTIONAL data item TLVs are:

- Current Data Rate
- If present, this value represents the requested data rate in bits per second (bps).
- Latency
- If present, this value represents the NEW maximum latency (or unchanged, if the request is denied), expressed in milliseconds, on the link.

11. Security Considerations

The protocol does not contain any mechanisms for security (e.g. authentication or encryption). The protocol assumes that any security would be implemented in the underlying transport (for example, by use of DTLS or some other mechanism), and is therefore outside the scope of this document.

12. IANA Considerations

This section specifies requests to IANA.

12.1 Registrations

This specification defines:

- o A new repository for DLEP signals, with fifteen values currently assigned.
- o Reservation of numbering space for Experimental DLEP signals.
- o A new repository for DLEP Data Items, with twenty-one values currently assigned.
- o Reservation of numbering space in the Data Items repository for experimental data items.
- o A request for allocation of a well-known port for DLEP communication.
- o A request for allocation of a multicast address for DLEP discovery.

12.2 Expert Review: Evaluation Guidelines

No additional guidelines for expert review are anticipated.

12.3 Message (Signal) TLV Type Registration

A new repository must be created with the values of the DLEP messages. Valid signals are:

- o Peer Discovery
- o Peer Offer
- o Peer Initialization
- o Peer Initialization ACK
- o Peer Update
- o Peer Update ACK
- o Peer Termination
- o Peer Termination ACK
- o Destination Up
- o Destination Up ACK
- o Destination Down
- o Destination Down ACK
- o Destination Update
- o Heartbeat
- o Link Characteristics Request
- o Link Characteristics ACK

It is also requested that the repository contain space for experimental signal types.

12.4 DLEP Data Item Registrations

A new repository for DLEP Data Items must be created. Valid Data Items are:

- o DLEP Version
- o Peer Type
- o MAC Address
- o IPv4 Address
- o IPv6 Address
- o Maximum Data Rate (Receive)
- o Maximum Data Rate (Transmit)
- o Current Data Rate (Receive)
- o Current Data Rate (Transmit)
- o Latency
- o Expected Forwarding Time
- o Resources (Receive)
- o Resources (Transmit)
- o Relative Link Quality (Receive)
- o Relative Link Quality (Transmit)
- o Status
- o Heartbeat Interval/Threshold
- o Link Characteristics ACK Timer
- o Credit Window Status
- o Credit Grant
- o Credit Request

It is also requested that the registry allocation contain space for experimental data items.

12.5 DLEP Well-known Port

It is requested that IANA allocate a well-known port number for DLEP communication.

12.6 DLEP Multicast Address

It is requested that IANA allocate a multicast address for DLEP discovery messages.

13. Appendix A.

13.1 Peer Level Message Flows

13.1.1 Modem Device Restarts Discovery

Router	Modem	Message Description
=====		
<-----Peer Discovery-----		Modem initiates discovery
-----Peer Offer-----> w/ Non-zero Status TLV		Router detects a problem, sends Peer Offer w/Status TLV indicating the error. Modem accepts failure, restarts discovery process.
<-----Peer Discovery-----		Modem initiates discovery
-----Peer Offer-----> w/ Zero Status TLV		Router accepts, sends Peer Offer w/ Status TLV indicating success. Discovery completed.

13.1.2 Modem Device Detects Peer Offer Timeout

Router	Modem	Message Description
=====		
<-----Peer Discovery-----		Modem initiates discovery, starts a guard timer. Modem guard timer expires. Modem restarts discovery process.
<-----Peer Discovery-----		Modem initiates discovery, starts a guard timer.
-----Peer Offer-----> w/ Zero Status TLV		Router accepts, sends Peer Offer w/ Status TLV indicating success. Discovery completed.

13.1.3 Router Peer Offer Lost

Router	Modem	Message Description
=====		
<-----Peer Discovery-----		Modem initiates discovery, starts a guard timer.
-----Peer Offer-----		Router offers availability
		Modem times out on Peer Offer, restarts discovery process.
<-----Peer Discovery-----		Modem initiates discovery
-----Peer Offer----->		Router detects subsequent discovery, internally terminates the previous, accepts the new association, sends Peer Offer w/Status TLV indicating success.
		Discovery completed.

13.1.4 Discovery Success

Router	Modem	Message Description
=====		
<-----Peer Discovery-----		Modem initiates discovery
-----Peer Offer----->		Router offers availability
<-----Peer Initialization-----		Modem Connects on TCP Port
<-----Peer Heartbeat-----		
-----Peer Heartbeat----->		
<=====		Message flow about destinations (i.e. Destination Up, Destination Down, Destination update)
<-----Peer Heartbeat-----		
-----Peer Heartbeat----->		

-----Peer Term Req-----> Terminate Request

<-----Peer Term Res----- Terminate Response

29.1.5 Router Detects a Heartbeat timeout

Router	Modem	Message Description
=====		
<-----Peer Heartbeat-----		
-----Peer Heartbeat----->		
---Peer Heartbeat-----		
~ ~ ~ ~ ~ ~ ~		
-----Peer Heartbeat----->		
---Peer Heartbeat-----		Router Heartbeat Timer expires, detects missing heartbeats. Router takes down all destination sessions and terminates the Peer association.
-----Peer Terminate ----->		Peer Terminate Request
		Modem takes down all destination sessions, then acknowledges the Peer Terminate
<----Peer Terminate ACK-----		Peer Terminate ACK

29.1.6 Modem Detects a Heartbeat timeout

Router	Modem	Message Description
=====		
<-----Peer Heartbeat-----		
-----Peer Heartbeat-----		
<-----Peer Heartbeat-----		
~ ~ ~ ~ ~ ~ ~		
-----Peer Heartbeat-----		


```

<-----Peer Heartbeat-----
                                Modem Heartbeat Timer expires,
                                detects missing heartbeats. Modem
                                takes down all destination sessions

<-----Peer Terminate-----  Peer Terminate Request

                                Router takes down all destination
                                sessions, then acknowledges the
                                Peer Terminate

-----Peer Terminate ACK-----> Peer Terminate ACK

```

[29.1.7](#) Peer Terminate (from Modem) Lost

Router	Modem	Message Description
=====		
-----Peer Terminate-----		Modem Peer Terminate Request
		Router Heartbeat times out, terminates association.
-----Peer Terminate----->		Router Peer Terminate
<-----Peer Terminate ACK-----		Modem sends Peer Terminate ACK

[29.1.8](#) Peer Terminate (from Router) Lost

Router	Modem	Message Description
=====		
-----Peer Terminate----->		Router Peer Terminate Request
		Modem HB times out, terminates association.
<-----Peer Terminate-----		Modem Peer Terminate
-----Peer Terminate ACK----->		Peer Terminate ACK

[29.2](#) Destination Specific Message Flows

29.2.1 Modem Destination Up Lost

Router	Modem	Message Description
=====		
-----Destination Up -----		Modem sends Destination Up
		Modem timesout on ACK
<-----Destination Up -----		Modem sends Destination Up
-----Destination Up ACK----->		Router accepts the destination session
<-----Destination Update-----		Modem Destination Metrics
.		
<-----Destination Update-----		Modem Destination Metrics

29.2.2 Router Detects Duplicate Destination Ups

Router	Modem	Message Description
=====		
<-----Destination Up -----		Modem sends Destination Up
-----Destination Up ACK-----		Router accepts the destination session
		Modem timesout on ACK
<-----Destination Up -----		Modem resends Destination Up
		Router detects duplicate Destination, takes down the previous, accepts the new Destination.
-----Destination Up ACK----->		Router accepts the destination session
<-----Destination Update-----		Modem Destination Metrics
.		
<-----Destination Update-----		Modem Destination Metrics

29.2.3 Destination Up, No Layer 3 Addresses

Router	Modem	Message Description
=====		
<-----Destination Up -----		Modem sends Destination Up
-----Destination Up ACK----->		Router accepts the destination session
		Router ARPs for IPv4 if defined. Router drives ND for IPv6 if defined.
<-----Destination Update-----		Modem Destination Metrics
.		
<-----Destination Update-----		Modem Destination Metrics

29.2.4 Destination Up with IPv4, No IPv6

Router	Modem	Message Description
=====		
<-----Destination Up -----		Modem sends Destination Up with the IPv4 TLV
-----Destination Up ACK----->		Router accepts the destination session
		Router drives ND for IPv6 if defined.
<-----Destination Update-----		Modem Destination Metrics
.		
<-----Destination Update-----		Modem Destination Metrics

29.2.5 Destination Up with IPv4 and IPv6

Router	Modem	Message Description
=====		
<-----Destination Up -----		Modem sends Destination Up with the IPv4 and IPv6 TLVs
-----Destination Up ACK----->		Router accepts the destination session


```
<-----Destination Update-----      Modem Destination Metrics
      . . . . .
```

[29.2.6](#) Destination Session Success

Router	Modem	Message Description
=====		
		-----Peer Offer-----> Router offers availability
		-----Peer Heartbeat----->
		<-----Destination Up ----- Modem
		-----Destination Up ACK-----> Router
		<-----Destination Update----- Modem
	
		<-----Destination Update----- Modem
		Modem initiates the terminate
		<-----Destination Down ----- Modem
		-----Destination Down ACK-----> Router
		or
		Router initiates the terminate
		-----Destination Down -----> Router
		<-----Destination Down ACK----- Modem

Acknowledgements

The authors would like to acknowledge and thank the members of the DLEP design team, who have provided invaluable insight. The members of the design team are: Teco Boot, Bow-Nan Cheng, John Dowdell, Henning Rogge, and Rick Taylor.

The authors would also like to acknowledge the influence and contributions of Chris Olsen, Martin Duke, Subir Das, Jaewon Kang, Vikram Kaul, and Nelson Powell.

Normative References

[RFC5578] Berry, B., Ed., "PPPoE with Credit Flow and Metrics",
[RFC 5578](#), February 2010.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate

Informative References

[TLS] Dierks, T. and Rescorla, E. "The Transport Layer Security
(TLS) Protocol", [RFC 5246](#), August 2008.

Author's Addresses

Stan Ratliff
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA
EMail: sratliff@cisco.com

Bo Berry
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA
EMail: boberry@cisco.com

Greg Harrison
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA
EMail: greharri@cisco.com

Shawn Jury
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA
Email: sjury@cisco.com

Darryl Satterwhite
Broadcom
USA
Email: dsatterw@broadcom.com

