Internet Engineering Task Force (IETF)
Request for Comments: 8175
Category: Standards Track
ISSN: 2070-1721

S. Ratliff
VT iDirect
S. Jury
Cisco Systems
D. Satterwhite
Broadcom
R. Taylor
Airbus Defence & Space
B. Berry
June 2017

Dynamic Link Exchange Protocol (DLEP)

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 routing 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. This document introduces a new protocol called the Dynamic Link Exchange Protocol (DLEP), which provides a bidirectional, event-driven communication channel between the router and the modem to facilitate communication of changing link characteristics.

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 http://www.rfc-editor.org/info/rfc8175.

Table of Contents

1. Introduction

There exist today a collection of modem devices that control links of variable data rate and quality. Examples of these types of links include line-of-sight (LOS) terrestrial radios, satellite terminals, and broadband modems. Fluctuations in speed and quality of these links can occur due to configuration, 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 data rate vary with respect to individual destinations on a link and with the type of traffic being sent. As an example, consider the case of an IEEE 802.11 access point serving two associated laptop computers. In this environment, the answer to the question "What is the data rate on the 802.11 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 data rate of 54 Mbps for unicast traffic, the other laptop, being relatively far away or obstructed by some object, can simultaneously have a data rate of only 32 Mbps 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 24 Mbps or less).

In addition to utilizing links that have variable data rates, mobile networks are challenged by the notion that link connectivity will come and go over time, without an effect on a router's interface state (Up or Down). Effectively utilizing a relatively short-lived connection is problematic in IP routed networks, as IP routing protocols tend to rely on interface state and independent timers to maintain network convergence (e.g., HELLO messages and/or recognition of DEAD routing adjacencies). These dynamic 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 paradigm driven by timers and/or interface state. DLEP not only implements such an event-driven paradigm but does so over a local (1-hop) TCP session, which guarantees delivery of the event messages.

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 or serial link. In the case of Ethernet 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 data rate, leads to a situation where finding the (current) best route through the network to a given node is 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 data rate may be available but will not be used unless the network devices emit traffic at a rate higher than the currently established rate. Increasing the traffic rate does not guarantee that additional data rate will be allocated; rather, it may result in data loss and additional retransmissions on the link.

Addressing the challenges listed above, the Dynamic Link Exchange Protocol, or DLEP, has been developed. DLEP runs between a router and its attached modem devices, allowing the modem devices to communicate (1) link characteristics as they change and (2) convergence events (acquisition and loss of potential routing next hops). Figures 1 and 2 illustrate the scope of DLEP packets.

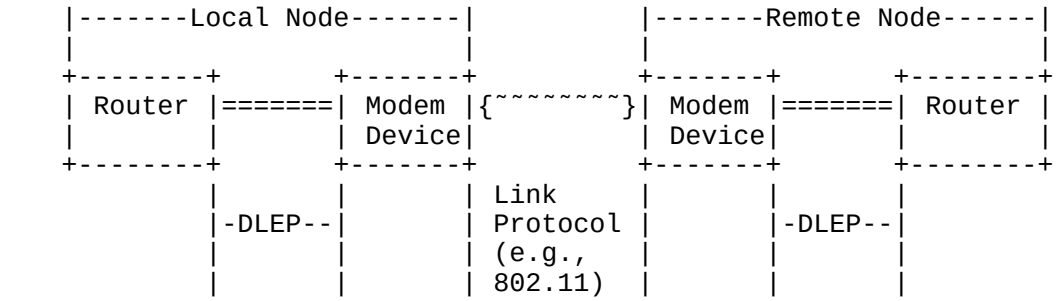

Figure 1: DLEP Network

In Figure 1, when the local modem detects the presence of a remote node, it (the local modem) sends a message to its router via DLEP. The message consists of an indication of what change has occurred on the link (e.g., the presence of a remote node detected), along with a collection of DLEP-defined Data Items that further describe the change. Upon receipt of the message, 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 use DLEP to report any characteristics of the link (data rate, latency, etc.) that have changed. DLEP is independent of the link type and topology supported by the modem. Note that DLEP 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 Device Type A implements a point-to-point link, and Modem Device Type B is connected via a shared medium. In both cases, DLEP 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 reconverge the network.

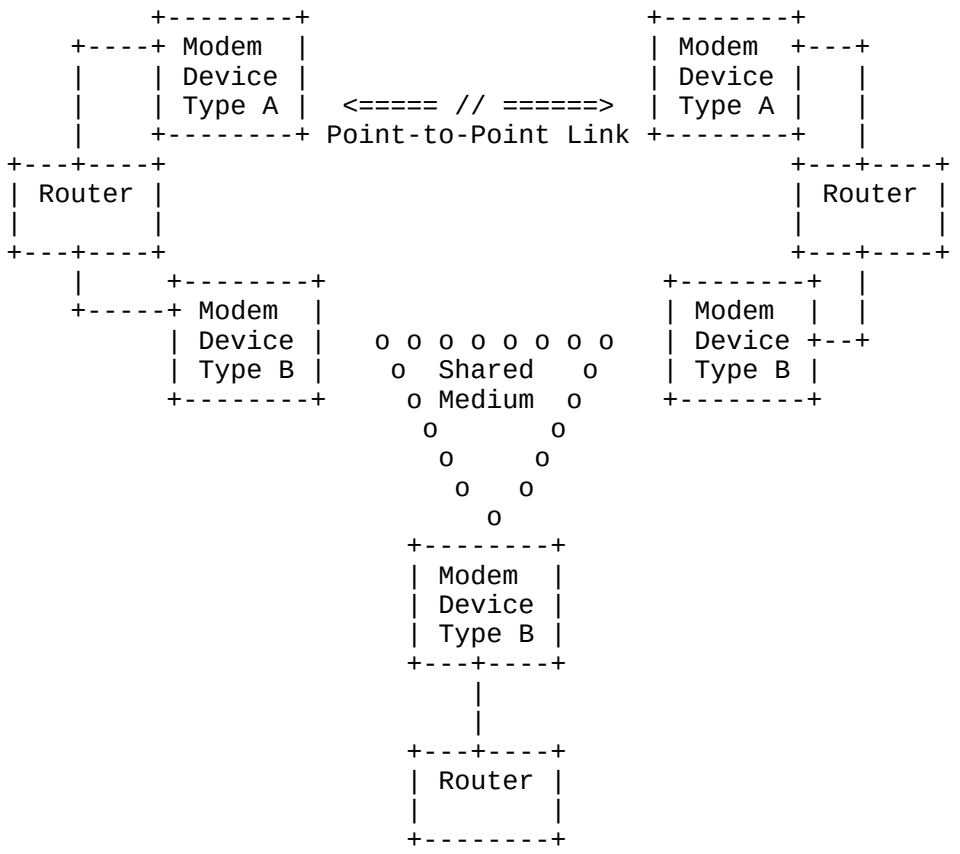

Figure 2: DLEP Network with Multiple Modem Devices

2. Protocol Overview

DLEP defines a set of Messages used by modems and their attached routers 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. Throughout this document, we refer to modems/routers participating in a DLEP session as "DLEP Participants", unless a specific distinction (e.g., modem or router) is required.

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) or the modem supports multiple connections

(via multiple logical or physical interfaces), then separate DLEP sessions exist for each modem or connection. A router and modem form a session by completing the discovery and initialization process. This router-modem session persists unless or until it either (1) times out, based on the absence of DLEP traffic (including heartbeats) or (2) is explicitly torn down by one of the DLEP participants.

While this document represents the best efforts of 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 used. Such extensions are defined as additional Messages, Data Items, and/or status codes, and associated rules of behavior, that are not defined in this document. DLEP contains a standard mechanism for router and modem implementations to negotiate the available extensions to use on a per-session basis.

2.1. Destinations

The router-modem session provides a carrier for information exchange concerning "destinations" that are available via the modem device. Destinations can be identified by either the router or the modem and represent a specific, addressable location that can be reached via the link(s) managed by the modem.

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 destinations accessible via the modem device, as well as the link characteristics to those destinations.

A destination can be either physical or logical. The example of a physical destination would be that of a remote, far-end router attached via the variable-quality network. It should be noted that for physical destinations the Media Access Control (MAC) address is the address of the far-end router, not the modem.

The example of a logical destination is Multicast. Multicast traffic destined for the variable-quality network (the network accessed via the 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 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 destination. The modem, once it is aware of the existence of this logical destination, reports link characteristics just as it

would for any other destination in the network. The specific algorithms a modem would use to derive metrics on logical destinations are outside the scope of this specification; these algorithms are left to specific implementations to decide.

In all cases, when this specification uses the term "destination", it refers to the addressable locations, either logical or physical, that are accessible by the radio link(s).

2.2. Conventions and Terminology

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.

3. Requirements

DLEP MUST be implemented on a single Layer 2 domain. The protocol identifies next-hop destinations by using the MAC address for delivering data traffic. No manipulation or substitution is performed; the MAC address supplied in all DLEP Messages is used as the Destination MAC address for frames emitted by the participating router. MAC addresses MUST be unique within the context of the router-modem session.

To enforce the single Layer 2 domain, implementations MUST support the Generalized TTL Security Mechanism [RFC5082], and implementations MUST adhere to this specification for all DLEP Messages.

DLEP specifies UDP multicast for single-hop discovery signaling and TCP for transport of the Messages. Modems and routers participating in DLEP sessions MUST have topologically consistent IP addresses assigned. It is RECOMMENDED that DLEP implementations utilize IPv6 link-local addresses to reduce the administrative burden of address assignment.

DLEP relies on the guaranteed delivery of its Messages between router and modem, once the 1-hop discovery process is complete -- hence, the specification of TCP to carry the Messages. Other reliable transports for the protocol are possible but are outside the scope of this document.

4. Implementation Scenarios

During development of this specification, two types of deployments were discussed.

The first can be viewed as a "dedicated deployment". In this mode, DLEP routers and modems are either directly connected (e.g., using crossover cables to connect interfaces) or connected to a dedicated switch. An example of this type of deployment would be a router with a line-of-sight radio connected into one interface, with a satellite modem connected into another interface. In mobile environments, the router and the connected modem (or modems) are placed into a mobile platform (e.g., a vehicle, boat, or airplane). In this mode, when a switch is used, it is possible that a small number of ancillary devices (e.g., a laptop) are also plugged into the switch. But in either event, the resulting network segment is constrained to a small number of devices and is not generally accessible from anywhere else in the network.

The other type of deployment envisioned can be viewed as a "networked deployment". In this type of scenario, the DLEP router and modem (or modems) are placed on a segment that is accessible from other points in the network. In this scenario, not only are the DLEP router and modem(s) accessible from other points in the network; the router and a given modem could be multiple physical hops away from each other. This scenario necessitates the use of Layer 2 tunneling technology to enforce the single-hop requirement of DLEP.

5. Assumptions

DLEP assumes that a signaling protocol exists between modems participating in a network. This specification does not define the character or behavior of this over-the-air signaling but does expect some information to be carried (or derived) by the signaling, such as the arrival and departure of modems from this network, and the variation of the link characteristics between modems. This information is then assumed to be used by the modem to implement DLEP.

This specification assumes that the link between router and modem is static with respect to data rate and latency and that this link is not likely to be the cause of a performance bottleneck. In deployments where the router and modem are physically separated by multiple network hops, served by Layer 2 tunneling technology, DLEP statistics on the RF links could be insufficient for routing protocols to make appropriate routing decisions. This would

especially become an issue in cases where the Layer 2 tunnel between router and modem is itself served in part (or in total) with a wireless backhaul link.

6. Metrics

DLEP includes the ability for the router and modem to communicate metrics that reflect the characteristics (e.g., data rate, 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 by a "best effort", incorporating all pertinent data that is available to the modem device. Metrics based on large-enough sample sizes will preclude short traffic bursts from adversely skewing reported values.

DLEP allows for metrics to be sent within two contexts -- metrics for a specific destination within the network (e.g., a specific router), and "per session" (those that apply to all destinations accessed via the modem). Most metrics can be further subdivided into transmit and receive metrics. In cases where metrics are provided at the session level, the router propagates the metrics to all entries in its information base for destinations that are accessed via the modem.

DLEP modems announce all metric Data Items that will be reported during the session, and provide default values for those metrics, in the Session Initialization Response Message (Section 12.6). In order to use a metric type that was not included in the Session Initialization Response Message, modem implementations terminate the session with the router (via the Session Termination Message (Section 12.9)) and establish a new session.

A DLEP modem can send metrics in both (1) a session context, via the Session Update Message (Section 12.7) and (2) a specific destination context, via the Destination Update Message (Section 12.17), at any time. The most recently received metric value takes precedence over any earlier value, regardless of context -- that is:

1. If the router receives metrics in a specific destination context (via the Destination Update Message), then the specific destination is updated with the new metric.

2. If the router receives metrics in a session-wide context (via the Session Update Message), then the metrics for all destinations accessed via the modem are updated with the new metric.

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

In addition to communicating existing metrics about the link, DLEP provides a Message allowing a router to request a different data rate or latency from the modem. This Message is the Link Characteristics Request Message (Section 12.18); it gives the router the ability to deal with requisite increases (or decreases) of allocated data rate/latency in demand-based schemes in a more deterministic manner.

7. DLEP Session Flow

All DLEP participants of a session transition through a number of distinct states during the lifetime of a DLEP session:

- Peer Discovery
- Session Initialization
- In-Session
- Session Termination
- Session Reset

Modems, and routers supporting DLEP discovery, transition through all five of the above states. Routers that rely on preconfigured TCP address/port information start in the Session Initialization state.

Modems MUST support the Peer Discovery state.

7.1. Peer Discovery State

Modems MUST support DLEP Peer Discovery; routers MAY support the discovery signals or rely on a priori configuration to locate modems. If a router chooses to support DLEP discovery, all signals MUST be supported.

In the Peer Discovery state, routers that support DLEP discovery MUST send Peer Discovery Signals (Section 12.3) to initiate modem discovery.

The router implementation then waits for a Peer Offer Signal (Section 12.4) response from a potential DLEP modem. While in the Peer Discovery state, Peer Discovery Signals MUST be sent repeatedly by a DLEP router, at regular intervals. It is RECOMMENDED that this interval be set to 60 seconds. The interval MUST be a minimum of 1 second; it SHOULD be a configurable parameter. Note that this operation (sending Peer Discovery and waiting for Peer Offer) is outside the DLEP transaction model (Section 8), as the transaction model only describes Messages on a TCP session.

Routers receiving a Peer Offer Signal MUST use one of the modem address/port combinations from the Peer Offer Signal to establish a TCP connection to the modem, even if a priori configuration exists. If multiple Connection Point Data Items exist in the received Peer Offer Signal, routers SHOULD prioritize IPv6 connection points over IPv4 connection points. If multiple connection points exist with the same transport (e.g., IPv6 or IPv4), implementations MAY use their own heuristics to determine the order in which they are tried. If a TCP connection cannot be achieved using any of the address/port combinations and the Discovery mechanism is in use, then the router SHOULD resume issuing Peer Discovery Signals. If no Connection Point Data Items are included in the Peer Offer Signal, the router MUST use the source address of the UDP packet containing the Peer Offer Signal as the IP address, and the DLEP well-known port number.

In the Peer Discovery state, the modem implementation MUST listen for incoming Peer Discovery Signals on the DLEP well-known IPv6 and/or IPv4 link-local multicast address and port. On receipt of a valid Peer Discovery Signal, it MUST reply with a Peer Offer Signal.

Modems MUST be prepared to accept a TCP connection from a router that is not using the Discovery mechanism, i.e., a connection attempt that occurs without a preceding Peer Discovery Signal.

Implementations of DLEP SHOULD implement, and use, Transport Layer Security (TLS) [RFC5246] to protect the TCP session. The "dedicated deployments" discussed in "Implementation Scenarios" (Section 4) MAY consider the use of DLEP without TLS. For all "networked deployments" (again, discussed in "Implementation Scenarios"), the implementation and use of TLS are STRONGLY RECOMMENDED. If TLS is to be used, then the TLS session MUST be established before any Messages are passed between peers. Routers supporting TLS MUST prioritize connection points using TLS over those that do not.

Upon establishment of a TCP connection, and the establishment of a TLS session if TLS is in use, both modem and router enter the Session Initialization state. It is up to the router implementation if Peer Discovery Signals continue to be sent after the device has

transitioned to the Session Initialization state. Modem implementations MUST silently ignore Peer Discovery Signals from a router with which a given implementation already has a TCP connection.

7.2. Session Initialization State

On entering the Session Initialization state, the router MUST send a Session Initialization Message (Section 12.5) to the modem. The router MUST then wait for receipt of a Session Initialization Response Message (Section 12.6) from the modem. Receipt of the Session Initialization Response Message containing a Status Data Item (Section 13.1) with status code set to 0 'Success' (see Table 2 in Section 13.1) indicates that the modem has received and processed the Session Initialization Message, and the router MUST transition to the In-Session state.

On entering the Session Initialization state, the modem MUST wait for receipt of a Session Initialization Message from the router. Upon receipt of a Session Initialization Message, the modem MUST send a Session Initialization Response Message, and the session MUST transition to the In-Session state. If the modem receives any Message other than Session Initialization or it fails to parse the received Message, it MUST NOT send any Message, and it MUST terminate the TCP connection and transition to the Session Reset state.

DLEP provides an extension negotiation capability to be used in the Session Initialization state; see Section 9. Extensions supported by an implementation MUST be declared to potential DLEP participants using the Extensions Supported Data Item (Section 13.6). Once both DLEP participants have exchanged initialization Messages, an implementation MUST NOT emit any Message, Signal, Data Item, or status code associated with an extension that was not specified in the received initialization Message from its peer.

7.3. In-Session State

In the In-Session state, Messages can flow in both directions between DLEP participants, indicating changes to the session state, the arrival or departure of reachable destinations, or changes of the state of the links to the destinations.

The In-Session state is maintained until one of the following conditions occurs:

- The implementation terminates the session by sending a Session Termination Message (Section 12.9), or
- Its peer terminates the session, indicated by receiving a Session Termination Message.

The implementation MUST then transition to the Session Termination state.

7.3.1. Heartbeats

In order to maintain the In-Session state, periodic Heartbeat Messages (Section 12.20) MUST be exchanged between router and modem. These Messages are intended to keep the session alive and to verify bidirectional connectivity between the two DLEP participants. It is RECOMMENDED that the interval timer between Heartbeat Messages be set to 60 seconds. The interval MUST be a minimum of 1 second; it SHOULD be a configurable parameter.

Each DLEP participant is responsible for the creation of Heartbeat Messages.

Receipt of any valid DLEP Message MUST reset the heartbeat interval timer (i.e., valid DLEP Messages take the place of, and obviate the need for, additional Heartbeat Messages).

An implementation MUST allow a minimum of 2 heartbeat intervals to expire with no Messages from its peer before terminating the session. When terminating the session, a Session Termination Message containing a Status Data Item (Section 13.1) with status code set to 132 'Timed Out' (see Table 2) MUST be sent, and then the implementation MUST transition to the Session Termination state.

7.4. Session Termination State

When an implementation enters the Session Termination state after sending a Session Termination Message (Section 12.9) as the result of an invalid Message or error, it MUST wait for a Session Termination Response Message (Section 12.10) from its peer. A sender SHOULD allow 4 heartbeat intervals to expire before assuming that its peer is unresponsive and before continuing with session termination. Any other Message received while waiting MUST be silently ignored.

When the sender of the Session Termination Message receives a Session Termination Response Message from its peer or times out, it MUST transition to the Session Reset state.

When an implementation receives a Session Termination Message from its peer, it enters the Session Termination state, and then it MUST immediately send a Session Termination Response and transition to the Session Reset state.

7.5. Session Reset State

In the Session Reset state, the implementation MUST perform the following actions:

- Release all resources allocated for the session.
- Eliminate all destinations in the information base represented by the session. Destination Down Messages (Section 12.15) MUST NOT be sent.
- Terminate the TCP connection.

Having completed these actions, the implementation SHOULD return to the relevant initial state:

- For modems: Peer Discovery.
- For routers: either Peer Discovery or Session Initialization, depending on configuration.

7.5.1. Unexpected TCP Connection Termination

If the TCP connection between DLEP participants is terminated when an implementation is not in the Session Reset state, the implementation MUST immediately transition to the Session Reset state.

8. Transaction Model

DLEP defines a simple Message transaction model: only one request per destination may be in progress at a time per session. A Message transaction is considered complete when a response matching a previously issued request is received. If a DLEP participant receives a request for a destination for which there is already an outstanding request, the implementation MUST terminate the session by issuing a Session Termination Message (Section 12.9) containing a Status Data Item (Section 13.1) with status code set to 129 'Unexpected Message' (see Table 2) and transition to the Session

Termination state. There is no restriction on the total number of Message transactions in progress at a time, as long as each transaction refers to a different destination.

It should be noted that some requests may take a considerable amount of time for some DLEP participants to complete; for example, a modem handling a multicast Destination Up request may have to perform a complex network reconfiguration. A sending implementation MUST be able to handle such long-running transactions gracefully.

Additionally, only one session request, e.g., a Session Initialization Message (Section 12.5), may be in progress at a time per session. As noted above for Message transactions, a session transaction is considered complete when a response matching a previously issued request is received. If a DLEP participant receives a session request while there is already a session request in progress, it MUST terminate the session by issuing a Session Termination Message containing a Status Data Item with status code set to 129 'Unexpected Message' and transition to the Session Termination state. Only the Session Termination Message may be issued when a session transaction is in progress. Heartbeat Messages (Section 12.20) MUST NOT be considered part of a session transaction.

DLEP transactions do not time out and are not cancellable, except for transactions in flight when the DLEP session is reset. If the session is terminated, canceling transactions in progress MUST be performed as part of resetting the state machine. An implementation can detect if its peer has failed in some way by use of the session heartbeat mechanism during the In-Session state; see Section 7.3.

9. Extensions

Extensions MUST be negotiated on a per-session basis during session initialization via the Extensions Supported mechanism. Implementations are not required to support any extensions in order to be considered DLEP compliant.

If interoperable protocol extensions are required, they will need to be standardized as either (1) an update to this document or (2) an additional standalone specification. The IANA registries defined in Section 15 of this document contain sufficient unassigned space for DLEP Signals, Messages, Data Items, and status codes to accommodate future extensions to the protocol.

As multiple protocol extensions MAY be announced during session initialization, authors of protocol extensions need to consider the interaction of their extensions with other published extensions and specify any incompatibilities.

9.1. Experiments

This document registers Private Use [RFC5226] numbering space in the DLEP Signal, Message, Data Item, and status code registries for experimental extensions. The intent is to allow for experimentation with new Signals, Messages, Data Items, and/or status codes while still retaining the documented DLEP behavior.

During session initialization, the use of the Private Use Signals, Messages, Data Items, status codes, or behaviors MUST be announced as DLEP extensions, using extension identifiers from the Private Use space in the "Extension Type Values" registry (Table 3), with a value agreed upon (a priori) between the participants. DLEP extensions using the Private Use numbering space are commonly referred to as "experiments".

Multiple experiments MAY be announced in the Session Initialization Messages. However, the use of multiple experiments in a single session could lead to interoperability issues or unexpected results (e.g., clashes of experimental Signals, Messages, Data Items, and/or status code types) and is therefore discouraged. It is left to implementations to determine the correct processing path (e.g., a decision on whether to terminate the session or establish a precedence of the conflicting definitions) if such conflicts arise.

10. Scalability

The protocol is intended to support thousands of destinations on a given modem/router pair. On a large scale, an implementation should consider employing techniques to prevent flooding its peer with a large number of Messages in a short time. For example, a dampening algorithm could be employed to prevent a flapping device from generating a large number of Destination Up / Destination Down Messages.

Also, the use of techniques such as a hysteresis can lessen the impact of rapid, minor fluctuations in link quality. The specific algorithms for handling flapping destinations and minor changes in link quality are outside the scope of this specification.

11. DLEP Signal and Message Structure

DLEP defines two protocol units used in two different ways: Signals and Messages. Signals are only used in the Discovery mechanism and are carried in UDP datagrams. Messages are used bidirectionally over a TCP connection between the participants, in the Session Initialization, In-Session, and Session Termination states.

Both Signals and Messages consist of a Header followed by an unordered list of Data Items. Headers consist of Type and Length information, while Data Items are encoded as TLV (Type-Length-Value) structures. In this document, the Data Items following a Signal or Message Header are described as being "contained in" the Signal or Message.

There is no restriction on the order of Data Items following a Header, and the acceptability of duplicate Data Items is defined by the definition of the Signal or Message declared by the type in the Header.

All integers in Header fields and values MUST be in network byte order.

11.1. DLEP Signal Header

The DLEP Signal Header contains the following fields:

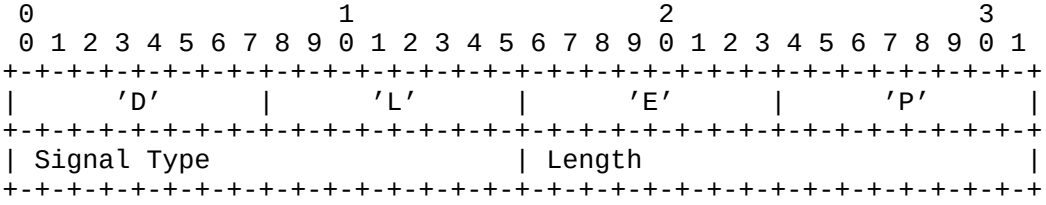

Figure 3: DLEP Signal Header

"DLEP": Every Signal MUST start with the following characters: U+0044, U+004C, U+0045, U+0050.

Signal Type: A 16-bit unsigned integer containing one of the DLEP Signal Type values defined in this document.

Length: The length in octets, expressed as a 16-bit unsigned integer, of all of the DLEP Data Items contained in this Signal. This length MUST NOT include the length of the Signal Header itself.

The DLEP Signal Header is immediately followed by zero or more DLEP Data Items, encoded in TLVs, as defined in this document.

11.2. DLEP Message Header

The DLEP Message Header contains the following fields:

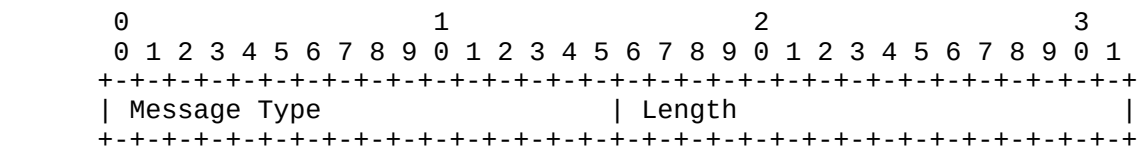

Figure 4: DLEP Message Header

Message Type: A 16-bit unsigned integer containing one of the DLEP Message Type values defined in this document.

Length: The length in octets, expressed as a 16-bit unsigned integer, of all of the DLEP Data Items contained in this Message. This length MUST NOT include the length of the Message Header itself.

The DLEP Message Header is immediately followed by zero or more DLEP Data Items, encoded in TLVs, as defined in this document.

11.3. DLEP Generic Data Item

All DLEP Data Items contain the following fields:

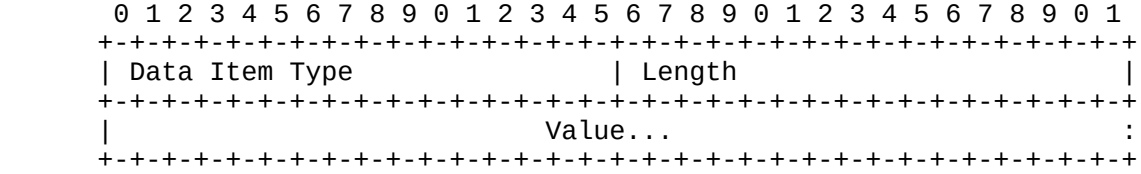

Figure 5: DLEP Generic Data Item

Data Item Type: A 16-bit unsigned integer field specifying the type of Data Item being sent.

Length: The length in octets, expressed as a 16-bit unsigned integer, of the Value field of the Data Item. This length MUST NOT include the length of the Data Item Type and Length fields.

Value: A field of <Length> octets that contains data specific to a particular Data Item.

12. DLEP Signals and Messages

12.1. General Processing Rules

If an unrecognized or unexpected Signal is received or if a received Signal contains unrecognized, invalid, or disallowed duplicate Data Items, the receiving implementation MUST ignore the Signal.

If a Signal is received with a TTL value that is NOT equal to 255, the receiving implementation MUST ignore the Signal.

If an unrecognized Message is received, the receiving implementation MUST issue a Session Termination Message (Section 12.9) containing a Status Data Item (Section 13.1) with status code set to 128 'Unknown Message' (see Table 2) and transition to the Session Termination state.

If an unexpected Message is received, the receiving implementation MUST issue a Session Termination Message containing a Status Data Item with status code set to 129 'Unexpected Message' and transition to the Session Termination state.

If a received Message contains unrecognized, invalid, or disallowed duplicate Data Items, 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.

If a packet in the TCP stream is received with a TTL value other than 255, the receiving implementation MUST immediately transition to the Session Reset state.

Prior to the exchange of Destination Up (Section 12.11) and Destination Up Response (Section 12.12) Messages, or Destination Announce (Section 12.13) and Destination Announce Response (Section 12.14) Messages, no Messages concerning a destination may be sent. An implementation receiving any Message with such an unannounced destination MUST terminate the session by issuing a Session Termination Message containing a Status Data Item with status code set to 131 'Invalid Destination' and transition to the Session Termination state.

After exchanging Destination Down (Section 12.15) and Destination Down Response (Section 12.16) Messages, no Messages concerning a destination may be sent until a new Destination Up or Destination Announce Message is sent. An implementation receiving a Message about a destination previously announced as 'down' MUST terminate the

session by issuing a Session Termination Message containing a Status Data Item with status code set to 131 'Invalid Destination' and transition to the Session Termination state.

12.2. Status Code Processing

The behavior of a DLEP participant receiving a Message containing a Status Data Item (Section 13.1) is defined by the failure mode associated with the value of the status code field; see Table 2. All status code values less than 100 have a failure mode of 'Continue'; all other status codes have a failure mode of 'Terminate'.

A DLEP participant receiving any Message apart from a Session Termination Message (Section 12.9) containing a Status Data Item with a status code value with failure mode 'Terminate' MUST immediately issue a Session Termination Message echoing the received Status Data Item and then transition to the Session Termination state.

A DLEP participant receiving a Message containing a Status Data Item with a status code value with failure mode 'Continue' can continue normal operation of the session.

12.3. Peer Discovery Signal

A Peer Discovery Signal SHOULD be sent by a DLEP router to discover DLEP modems in the network; see Section 7.1.

A Peer Discovery Signal MUST be encoded within a UDP packet. The destination MUST be set to the DLEP well-known address and port number. For routers supporting both IPv4 and IPv6 DLEP operation, it is RECOMMENDED that IPv6 be selected as the transport. The source IP address MUST be set to the router IP address associated with the DLEP interface. There is no DLEP-specific restriction on source port.

To construct a Peer Discovery Signal, the Signal Type value in the Signal Header is set to 1 (see "Signal Type Registration" (Section 15.2)).

The Peer Discovery Signal MAY contain a Peer Type Data Item (Section 13.4).

12.4. Peer Offer Signal

A Peer Offer Signal MUST be sent by a DLEP modem in response to a properly formatted and addressed Peer Discovery Signal (Section 12.3).

A Peer Offer Signal MUST be encoded within a UDP packet. The IP source and destination fields in the packet MUST be set by swapping the values received in the Peer Discovery Signal. The Peer Offer Signal completes the discovery process; see Section 7.1.

To construct a Peer Offer Signal, the Signal Type value in the Signal Header is set to 2 (see "Signal Type Registration" (Section 15.2)).

The Peer Offer Signal MAY contain a Peer Type Data Item (Section 13.4).

The Peer Offer Signal MAY contain one or more of any of the following Data Items, with different values:

- IPv4 Connection Point (Section 13.2)
- IPv6 Connection Point (Section 13.3)

The IPv4 and IPv6 Connection Point Data Items indicate the unicast address the router MUST use when connecting the DLEP TCP session.

12.5. Session Initialization Message

A Session Initialization Message MUST be sent by a DLEP router as the first Message of the DLEP TCP session. It is sent by the router after a TCP connect to an address/port combination that was obtained either via receipt of a Peer Offer or from a priori configuration.

To construct a Session Initialization Message, the Message Type value in the Message Header is set to 1 (see "Message Type Registration" (Section 15.3)).

The Session Initialization Message MUST contain one of each of the following Data Items:

- Heartbeat Interval (Section 13.5)
- Peer Type (Section 13.4)

If DLEP extensions are supported, the Session Initialization Message MUST contain an Extensions Supported Data Item (Section 13.6).

The Session Initialization Message MAY contain one or more of each of the following Data Items, with different values and with the Add/Drop (A) flag (Section 13) set to 1:

- IPv4 Address (Section 13.8)
- IPv6 Address (Section 13.9)
- IPv4 Attached Subnet (Section 13.10)
- IPv6 Attached Subnet (Section 13.11)

If any optional extensions are supported by the implementation, they MUST be enumerated in the Extensions Supported Data Item. If an Extensions Supported Data Item does not exist in a Session Initialization Message, the modem MUST conclude that there is no support for extensions in the router.

DLEP Heartbeats are not started until receipt of the Session Initialization Response Message (Section 12.6), and therefore implementations MUST use their own timeout heuristics for this Message.

As an exception to the general rule governing an implementation receiving an unrecognized Data Item in a Message (see Section 12.1), if a Session Initialization Message contains one or more Extensions Supported Data Items announcing support for extensions that the implementation does not recognize, then the implementation MAY ignore Data Items it does not recognize.

12.6. Session Initialization Response Message

A Session Initialization Response Message MUST be sent by a DLEP modem in response to a received Session Initialization Message (Section 12.5).

To construct a Session Initialization Response Message, the Message Type value in the Message Header is set to 2 (see "Message Type Registration" (Section 15.3)).

The Session Initialization Response Message MUST contain one of each of the following Data Items:

- Status (Section 13.1)
- Peer Type (Section 13.4)
- Heartbeat Interval (Section 13.5)

- Maximum Data Rate (Receive) (Section 13.12)
- Maximum Data Rate (Transmit) (Section 13.13)
- Current Data Rate (Receive) (Section 13.14)
- Current Data Rate (Transmit) (Section 13.15)
- Latency (Section 13.16)

The Session Initialization Response Message MUST contain one of each of the following Data Items, if the Data Item will be used during the lifetime of the session:

- Resources (Section 13.17)
- Relative Link Quality (Receive) (Section 13.18)
- Relative Link Quality (Transmit) (Section 13.19)
- Maximum Transmission Unit (MTU) (Section 13.20)

If DLEP extensions are supported, the Session Initialization Response Message MUST contain an Extensions Supported Data Item (Section 13.6).

The Session Initialization Response Message MAY contain one or more of each of the following Data Items, with different values and with the Add/Drop (A) flag (Section 13) set to 1:

- IPv4 Address (Section 13.8)
- IPv6 Address (Section 13.9)
- IPv4 Attached Subnet (Section 13.10)
- IPv6 Attached Subnet (Section 13.11)

The Session Initialization Response Message completes the DLEP session establishment; the modem should transition to the In-Session state when the Message is sent, and the router should transition to the In-Session state upon receipt of an acceptable Session Initialization Response Message.

All supported metric Data Items MUST be included in the Session Initialization Response Message, with default values to be used on a session-wide basis. This can be viewed as the modem "declaring" all supported metrics at DLEP session initialization. Receipt of any

further DLEP Message containing a metric Data Item not included in the Session Initialization Response Message MUST be treated as an error, resulting in the termination of the DLEP session between router and modem.

If any optional extensions are supported by the modem, they MUST be enumerated in the Extensions Supported Data Item. If an Extensions Supported Data Item does not exist in a Session Initialization Response Message, the router MUST conclude that there is no support for extensions in the modem.

After the Session Initialization / Session Initialization Response Messages have been successfully exchanged, implementations MUST only use extensions that are supported by both DLEP participants; see Section 7.2.

12.7. Session Update Message

A Session Update Message MAY be sent by a DLEP participant, on a session-wide basis, to indicate local Layer 3 address changes and/or metric changes.

To construct a Session Update Message, the Message Type value in the Message Header is set to 3 (see "Message Type Registration" (Section 15.3)).

The Session Update Message MAY contain one or more of each of the following Data Items, with different values:

- IPv4 Address (Section 13.8)
- IPv6 Address (Section 13.9)
- IPv4 Attached Subnet (Section 13.10)
- IPv6 Attached Subnet (Section 13.11)

When sent by a modem, the Session Update Message MAY contain one of each of the following Data Items:

- Maximum Data Rate (Receive) (Section 13.12)
- Maximum Data Rate (Transmit) (Section 13.13)
- Current Data Rate (Receive) (Section 13.14)

- Current Data Rate (Transmit) (Section 13.15)
- Latency (Section 13.16)

When sent by a modem, the Session Update Message MAY contain one of each of the following Data Items, if the Data Item is in use by the session:

- Resources (Section 13.17)
- Relative Link Quality (Receive) (Section 13.18)
- Relative Link Quality (Transmit) (Section 13.19)
- Maximum Transmission Unit (MTU) (Section 13.20)

If metrics are supplied with the Session Update Message (e.g., Maximum Data Rate), these metrics are considered to be session-wide and therefore MUST be applied to all destinations in the information base associated with the DLEP session. This includes destinations for which metrics may have been stored based on received Destination Update messages.

It should be noted that Session Update Messages can be sent by both routers and modems. For example, the addition of an IPv4 address on the router MAY prompt a Session Update Message to its attached modems. Also, for example, a modem that changes its Maximum Data Rate (Receive) for all destinations MAY reflect that change via a Session Update Message to its attached router(s).

Concerning Layer 3 addresses and subnets: if the modem is capable of understanding and forwarding this information (via mechanisms not defined by DLEP), the update would prompt any remote DLEP-enabled modems to issue a Destination Update Message (Section 12.17) to their local routers with the new (or deleted) addresses and subnets.

12.8. Session Update Response Message

A Session Update Response Message MUST be sent by a DLEP participant when a Session Update Message (Section 12.7) is received.

To construct a Session Update Response Message, the Message Type value in the Message Header is set to 4 (see "Message Type Registration" (Section 15.3)).

The Session Update Response Message MUST contain a Status Data Item (Section 13.1).

12.9. Session Termination Message

When a DLEP participant determines that the DLEP session needs to be terminated, the participant MUST send (or attempt to send) a Session Termination Message.

To construct a Session Termination Message, the Message Type value in the Message Header is set to 5 (see "Message Type Registration" (Section 15.3)).

The Session Termination Message MUST contain a Status Data Item (Section 13.1).

It should be noted that Session Termination Messages can be sent by both routers and modems.

12.10. Session Termination Response Message

A Session Termination Response Message MUST be sent by a DLEP participant when a Session Termination Message (Section 12.9) is received.

To construct a Session Termination Response Message, the Message Type value in the Message Header is set to 6 (see "Message Type Registration" (Section 15.3)).

There are no valid Data Items for the Session Termination Response Message.

Receipt of a Session Termination Response Message completes the teardown of the DLEP session; see Section 7.4.

12.11. Destination Up Message

Destination Up Messages MAY be sent by a modem to inform its attached router of the presence of a new reachable destination.

To construct a Destination Up Message, the Message Type value in the Message Header is set to 7 (see "Message Type Registration" (Section 15.3)).

The Destination Up Message MUST contain a MAC Address Data Item (Section 13.7).

The Destination Up Message SHOULD contain one or more of each of the following Data Items, with different values:

- IPv4 Address (Section 13.8)
- IPv6 Address (Section 13.9)

The Destination Up Message MAY contain one of each of the following Data Items:

- Maximum Data Rate (Receive) (Section 13.12)
- Maximum Data Rate (Transmit) (Section 13.13)
- Current Data Rate (Receive) (Section 13.14)
- Current Data Rate (Transmit) (Section 13.15)
- Latency (Section 13.16)

The Destination Up Message MAY contain one of each of the following Data Items, if the Data Item is in use by the session:

- Resources (Section 13.17)
- Relative Link Quality (Receive) (Section 13.18)
- Relative Link Quality (Transmit) (Section 13.19)
- Maximum Transmission Unit (MTU) (Section 13.20)

The Destination Up Message MAY contain one or more of each of the following Data Items, with different values:

- IPv4 Attached Subnet (Section 13.10)
- IPv6 Attached Subnet (Section 13.11)

A router receiving a Destination Up Message allocates the necessary resources, creating an entry in the information base with the specifics (MAC Address, Latency, Data Rate, etc.) of the destination. The information about this destination will persist in the router’s information base until a Destination Down Message (Section 12.15) is received, indicating that the modem has lost contact with the remote node or that the implementation transitions to the Session Termination state.

12.12. Destination Up Response Message

A router MUST send a Destination Up Response Message when a Destination Up Message (Section 12.11) is received.

To construct a Destination Up Response Message, the Message Type value in the Message Header is set to 8 (see "Message Type Registration" (Section 15.3)).

The Destination Up Response Message MUST contain one of each of the following Data Items:

- MAC Address (Section 13.7)
- Status (Section 13.1)

A router that wishes to receive further information concerning the destination identified in the corresponding Destination Up Message MUST set the status code of the included Status Data Item to 0 'Success'; see Table 2.

If the router has no interest in the destination identified in the corresponding Destination Up Message, then it MAY set the status code of the included Status Data Item to 1 'Not Interested'.

A modem receiving a Destination Up Response Message containing a Status Data Item with a status code of any value other than 0 'Success' MUST NOT send further Destination Messages about the destination, e.g., Destination Down (Section 12.15) or Destination Update (Section 12.17) with the same MAC address.

12.13. Destination Announce Message

Usually, a modem will discover the presence of one or more remote router/modem pairs and announce each destination's arrival by sending a corresponding Destination Up Message (Section 12.11) to the router. However, there may be times when a router wishes to express an interest in a destination that has yet to be announced, typically a multicast destination. Destination Announce Messages MAY be sent by a router to announce such an interest.

A Destination Announce Message MAY also be sent by a router to request information concerning a destination (1) in which the router has previously declined interest, via the 1 'Not Interested' status code in a Destination Up Response Message (Section 12.12) (see Table 2) or (2) that was previously declared as 'down', via the Destination Down Message (Section 12.15).

To construct a Destination Announce Message, the Message Type value in the Message Header is set to 9 (see "Message Type Registration" (Section 15.3)).

The Destination Announce Message MUST contain a MAC Address Data Item (Section 13.7).

The Destination Announce Message MAY contain zero or more of the following Data Items, with different values:

- IPv4 Address (Section 13.8)
- IPv6 Address (Section 13.9)

One of the advantages of implementing DLEP is to leverage the modem's knowledge of the links between remote destinations, allowing routers to avoid using probed neighbor discovery techniques; therefore, modem implementations SHOULD announce available destinations via the Destination Up Message, rather than relying on Destination Announce Messages.

12.14. Destination Announce Response Message

A modem MUST send a Destination Announce Response Message when a Destination Announce Message (Section 12.13) is received.

To construct a Destination Announce Response Message, the Message Type value in the Message Header is set to 10 (see "Message Type Registration" (Section 15.3)).

The Destination Announce Response Message MUST contain one of each of the following Data Items:

- MAC Address (Section 13.7)
- Status (Section 13.1)

The Destination Announce Response Message MAY contain one or more of each of the following Data Items, with different values:

- IPv4 Address (Section 13.8)
- IPv6 Address (Section 13.9)
- IPv4 Attached Subnet (Section 13.10)
- IPv6 Attached Subnet (Section 13.11)

The Destination Announce Response Message MAY contain one of each of the following Data Items:

- Maximum Data Rate (Receive) (Section 13.12)
- Maximum Data Rate (Transmit) (Section 13.13)
- Current Data Rate (Receive) (Section 13.14)
- Current Data Rate (Transmit) (Section 13.15)
- Latency (Section 13.16)

The Destination Announce Response Message MAY contain one of each of the following Data Items, if the Data Item is in use by the session:

- Resources (Section 13.17)
- Relative Link Quality (Receive) (Section 13.18)
- Relative Link Quality (Transmit) (Section 13.19)
- Maximum Transmission Unit (MTU) (Section 13.20)

If a modem is unable to report information immediately about the requested information -- for example, if the destination is not currently reachable -- the status code in the Status Data Item MUST be set to 2 'Request Denied'; see Table 2.

After sending a Destination Announce Response Message containing a Status Data Item with a status code of 0 'Success', a modem then announces changes to the link to the destination via Destination Update Messages.

When a successful Destination Announce Response Message is received, the router should add knowledge of the available destination to its information base.

12.15. Destination Down Message

A modem MUST send a Destination Down Message to report when a destination (a remote node or a multicast group) is no longer reachable.

A router MAY send a Destination Down Message to report when it no longer requires information concerning a destination.

To construct a Destination Down Message, the Message Type value in the Message Header is set to 11 (see "Message Type Registration" (Section 15.3)).

The Destination Down Message MUST contain a MAC Address Data Item (Section 13.7).

It should be noted that both modem and router may send a Destination Down Message to their peer, regardless of which participant initially indicated the destination to be 'up'.

12.16. Destination Down Response Message

A Destination Down Response Message MUST be sent by the recipient of a Destination Down Message (Section 12.15) to confirm that the relevant data concerning the destination has been removed from the information base.

To construct a Destination Down Response Message, the Message Type value in the Message Header is set to 12 (see "Message Type Registration" (Section 15.3)).

The Destination Down Response Message MUST contain one of each of the following Data Items:

- MAC Address (Section 13.7)
- Status (Section 13.1)

12.17. Destination Update Message

A modem SHOULD send a 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 as follows:

- Change in link metrics (e.g., data rates)
- Layer 3 addressing change

To construct a Destination Update Message, the Message Type value in the Message Header is set to 13 (see "Message Type Registration" (Section 15.3)).

The Destination Update Message MUST contain a MAC Address Data Item (Section 13.7).

The Destination Update Message MAY contain one of each of the following Data Items:

- Maximum Data Rate (Receive) (Section 13.12)
- Maximum Data Rate (Transmit) (Section 13.13)
- Current Data Rate (Receive) (Section 13.14)
- Current Data Rate (Transmit) (Section 13.15)
- Latency (Section 13.16)

The Destination Update Message MAY contain one of each of the following Data Items, if the Data Item is in use by the session:

- Resources (Section 13.17)
- Relative Link Quality (Receive) (Section 13.18)
- Relative Link Quality (Transmit) (Section 13.19)
- Maximum Transmission Unit (MTU) (Section 13.20)

The Destination Update Message MAY contain one or more of each of the following Data Items, with different values:

- IPv4 Address (Section 13.8)
- IPv6 Address (Section 13.9)
- IPv4 Attached Subnet (Section 13.10)
- IPv6 Attached Subnet (Section 13.11)

Metrics supplied in this Message overwrite metrics provided in a previously received Session Message, Destination Message, or Link Characteristics Message (e.g., Session Initialization, Destination Up, Link Characteristics Response).

It should be noted that this Message has no corresponding response.

12.18. Link Characteristics Request Message

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

To construct a Link Characteristics Request Message, the Message Type value in the Message Header is set to 14 (see "Message Type Registration" (Section 15.3)).

The Link Characteristics Request Message MUST contain a MAC Address Data Item (Section 13.7).

The Link Characteristics Request Message MUST also contain at least one of each of the following Data Items:

- Current Data Rate (Receive) (Section 13.14)
- Current Data Rate (Transmit) (Section 13.15)
- Latency (Section 13.16)

The Link Characteristics Request Message MAY contain either a Current Data Rate (Receive) (CDRR) or Current Data Rate (Transmit) (CDRT) Data Item to request a different data rate than is currently allocated, a Latency Data Item to request that traffic delay on the link not exceed the specified value, or both.

The router sending a Link Characteristics Request Message should be aware that a request may take an extended period of time to complete.

12.19. Link Characteristics Response Message

A modem MUST send a Link Characteristics Response Message when a Link Characteristics Request Message (Section 12.18) is received.

To construct a Link Characteristics Response Message, the Message Type value in the Message Header is set to 15 (see "Message Type Registration" (Section 15.3)).

The Link Characteristics Response Message MUST contain one of each of the following Data Items:

- MAC Address (Section 13.7)
- Status (Section 13.1)

The Link Characteristics Response Message SHOULD contain one of each of the following Data Items:

- Maximum Data Rate (Receive) (Section 13.12)
- Maximum Data Rate (Transmit) (Section 13.13)
- Current Data Rate (Receive) (Section 13.14)
- Current Data Rate (Transmit) (Section 13.15)
- Latency (Section 13.16)

The Link Characteristics Response Message MAY contain one of each of the following Data Items, if the Data Item is in use by the session:

- Resources (Section 13.17)
- Relative Link Quality (Receive) (Section 13.18)
- Relative Link Quality (Transmit) (Section 13.19)
- Maximum Transmission Unit (MTU) (Section 13.20)

The Link Characteristics Response Message MUST contain a complete set of metric Data Items, referencing all metrics declared in the Session Initialization Response Message (Section 12.6). The values in the metric Data Items in the Link Characteristics Response Message MUST reflect the link characteristics after the request has been processed.

If an implementation is not able to alter the characteristics of the link in the manner requested, then the status code of the Status Data Item MUST be set to 2 ’Request Denied’; see Table 2.

12.20. Heartbeat Message

A Heartbeat Message MUST be sent by a DLEP participant every N milliseconds, where N is defined in the Heartbeat Interval Data Item (Section 13.5) of the Session Initialization Message (Section 12.5) or Session Initialization Response Message (Section 12.6).

To construct a Heartbeat Message, the Message Type value in the Message Header is set to 16 (see "Message Type Registration" (Section 15.3)).

There are no valid Data Items for the Heartbeat Message.

The Heartbeat Message is used by DLEP participants to detect when a DLEP session peer (either the modem or the router) is no longer communicating; see Section 7.3.1.

13. DLEP Data Items

The core DLEP Data Items are as follows:

Type Code	Description
0	Reserved
1	Status (Section 13.1)
2	IPv4 Connection Point (Section 13.2)
3	IPv6 Connection Point (Section 13.3)
4	Peer Type (Section 13.4)
5	Heartbeat Interval (Section 13.5)
6	Extensions Supported (Section 13.6)
7	MAC Address (Section 13.7)
8	IPv4 Address (Section 13.8)
9	IPv6 Address (Section 13.9)
10	IPv4 Attached Subnet (Section 13.10)
11	IPv6 Attached Subnet (Section 13.11)
12	Maximum Data Rate (Receive) (MDRR) (Section 13.12)
13	Maximum Data Rate (Transmit) (MDRT) (Section 13.13)
14	Current Data Rate (Receive) (CDRR) (Section 13.14)
15	Current Data Rate (Transmit) (CDRT) (Section 13.15)
16	Latency (Section 13.16)
17	Resources (RES) (Section 13.17)

18	Relative Link Quality (Receive) (RLQR) (Section 13.18)
19	Relative Link Quality (Transmit) (RLQT) (Section 13.19)
20	Maximum Transmission Unit (MTU) (Section 13.20)
21-65407	Unassigned (available for future extensions)
65408-65534	Reserved for Private Use (available for experiments)
65535	Reserved

Table 1: DLEP Data Item Types

13.1. Status

For the Session Termination Message (Section 12.9), the Status Data Item indicates a reason for the termination. For all response messages, the Status Data Item is used to indicate the success or failure of the previously received Message.

The Status Data Item includes an optional Text field that can be used to provide a textual description of the status. The use of the Text field is entirely up to the receiving implementation, e.g., it could be output to a log file or discarded. If no Text field is supplied with the Status Data Item, the Length field MUST be set to 1.

The Status Data Item contains the following fields:

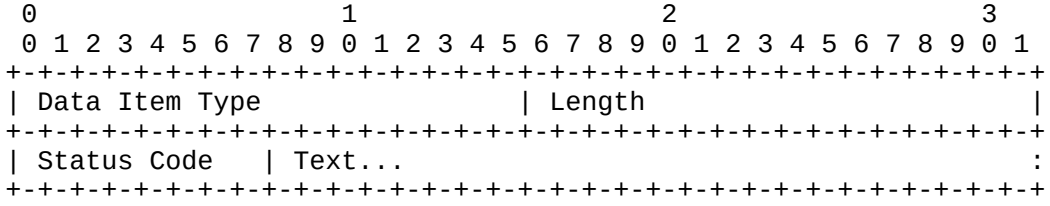

Data Item Type: 1

Length: 1 + Length of Text, in octets.

Status Code: One of the status codes defined in Table 2 below.

Text: UTF-8 encoded string of Unicode [RFC3629] characters, describing the cause, used for implementation-defined purposes. Since this field is used for description purposes, implementations SHOULD limit characters in this field to printable characters.

An implementation MUST NOT assume that the Text field is a NUL-terminated string of printable characters.

Status Code	Failure Mode	Description	Reason
0	Continue	Success	The Message was processed successfully.
1	Continue	Not Interested	The receiver is not interested in this Message subject, e.g., in a Destination Up Response Message (Section 12.12) to indicate no further Messages about the destination.
2	Continue	Request Denied	The receiver refuses to complete the request.
3	Continue	Inconsistent Data	One or more Data Items in the Message describe a logically inconsistent state in the network -- for example, in the Destination Up Message (Section 12.11) when an announced subnet clashes with an existing destination subnet.

4-111	Continue	<Unassigned>	Available for future extensions.
112-127	Continue	<Reserved for Private Use>	Available for experiments.
128	Terminate	Unknown Message	The Message was not recognized by the implementation.
129	Terminate	Unexpected Message	The Message was not expected while the device was in the current state, e.g., a Session Initialization Message (Section 12.5) in the In-Session state.
130	Terminate	Invalid Data	One or more Data Items in the Message are invalid, unexpected, or incorrectly duplicated.
131	Terminate	Invalid Destination	The destination included in the Message does not match a previously announced destination -- for example, in the Link Characteristics Response Message (Section 12.19).
132	Terminate	Timed Out	The session has timed out.
133-239	Terminate	<Unassigned>	Available for future extensions.

240-254	Terminate	<Reserved for Private Use>	Available for experiments.
255	Terminate	Shutting Down	The peer is terminating the session, as it is shutting down.

Table 2: DLEP Status Codes

13.2. IPv4 Connection Point

The IPv4 Connection Point Data Item indicates the IPv4 address and, optionally, the TCP port number on the modem available for connections. If provided, the router MUST use this information to initiate the TCP connection to the modem.

The IPv4 Connection Point Data Item contains the following fields:

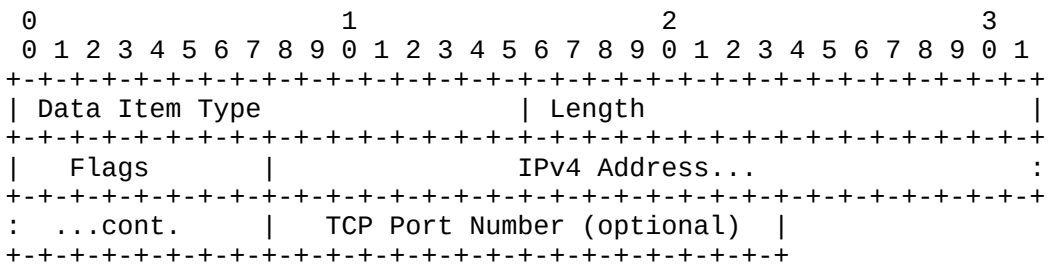

Data Item Type: 2

Length: 5 (or 7 if TCP Port Number included).

Flags: Flags field, defined below.

IPv4 Address: The IPv4 address listening on the modem.

TCP Port Number: TCP port number on the modem.

If the Length field is 7, the port number specified MUST be used to establish the TCP session. If the TCP Port Number is omitted, i.e., the Length field is 5, the router MUST use the DLEP well-known port number (Section 15.14) to establish the TCP connection.

The Flags field is defined as:

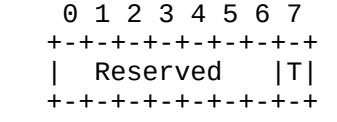

T: Use TLS flag, indicating whether the TCP connection to the given address and port requires the use of TLS [RFC5246] (1) or not (0).

Reserved: MUST be zero. Left for future assignment.

13.3. IPv6 Connection Point

The IPv6 Connection Point Data Item indicates the IPv6 address and, optionally, the TCP port number on the modem available for connections. If provided, the router MUST use this information to initiate the TCP connection to the modem.

The IPv6 Connection Point Data Item contains the following fields:

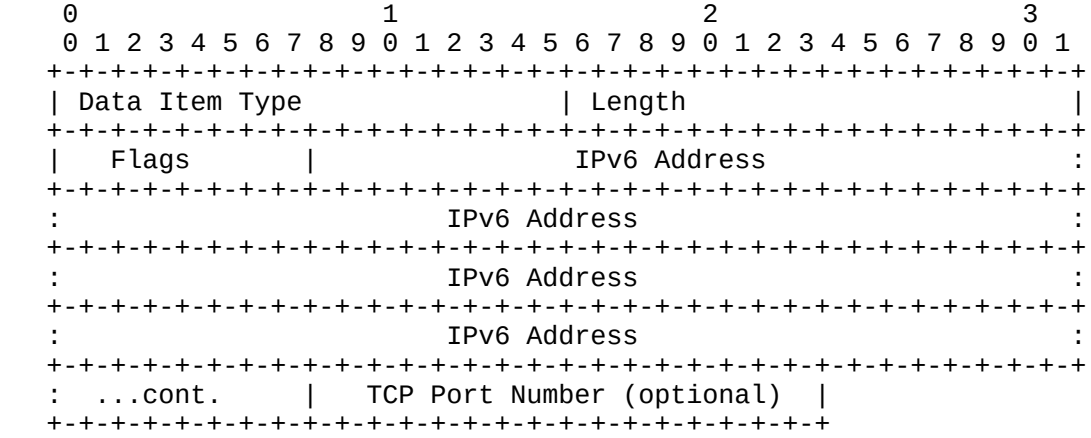

Data Item Type: 3

Length: 17 (or 19 if TCP Port Number included).

Flags: Flags field, defined below.

IPv6 Address: The IPv6 address listening on the modem.

TCP Port Number: TCP port number on the modem.

If the Length field is 19, the port number specified MUST be used to establish the TCP session. If the TCP Port Number is omitted, i.e., the Length field is 17, the router MUST use the DLEP well-known port number (Section 15.14) to establish the TCP connection.

The Flags field is defined as:

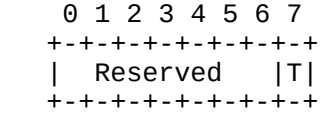

T: Use TLS flag, indicating whether the TCP connection to the given address and port requires the use of TLS [RFC5246] (1) or not (0).

Reserved: MUST be zero. Left for future assignment.

13.4. Peer Type

The Peer Type Data Item is used by the router and modem to give additional information as to its type and the properties of the over-the-air control plane.

With some devices, access to the shared RF medium is strongly controlled. One example of this would be satellite modems -- where protocols, proprietary in nature, have been developed to ensure that a given modem has authorization to connect to the shared medium. Another example of this class of modems is governmental/military devices, where elaborate mechanisms have been developed to ensure that only authorized devices can connect to the shared medium. Contrasting with the above, there are modems where no such access control is used. An example of this class of modem would be one that supports the 802.11 ad hoc mode of operation. The Secured Medium (S) flag is used to indicate if access control is in place.

The Peer Type Data Item includes a textual description of the peer; it is envisioned that the text will be used for informational purposes (e.g., as output in a display command).

The Peer Type Data Item contains the following fields:

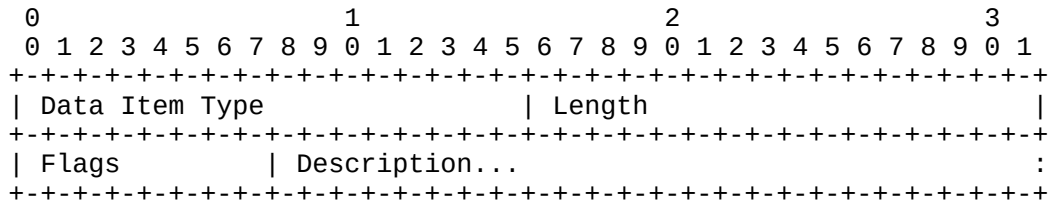

Data Item Type: 4

Length: 1 + Length of Description, in octets.

Flags: Flags field, defined below.

Description: UTF-8 encoded string of Unicode [RFC3629] characters. For example, a satellite modem might set this variable to "Satellite terminal". Since this Data Item is intended to provide additional information for display commands, sending implementations SHOULD limit the data to printable characters.

An implementation MUST NOT assume that the Description field is a NUL-terminated string of printable characters.

The Flags field is defined as:

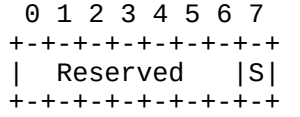

S: Secured Medium flag, used by a modem to indicate whether the shared RF medium implements access control (1) or not (0). The Secured Medium flag only has meaning in Signals and Messages sent by a modem.

Reserved: MUST be zero. Left for future assignment.

13.5. Heartbeat Interval

The Heartbeat Interval Data Item is used to specify a period in milliseconds for Heartbeat Messages (Section 12.20).

The Heartbeat Interval Data Item contains the following fields:

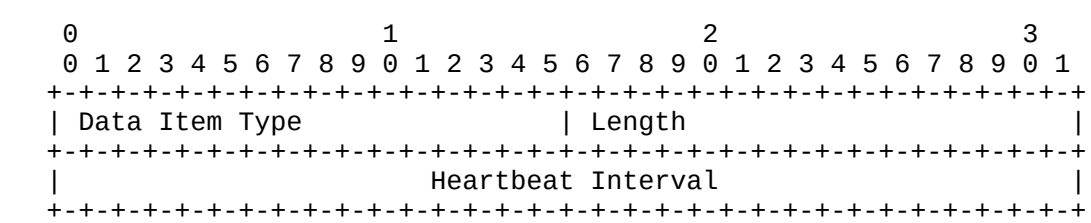

Data Item Type: 5

Length: 4

Heartbeat Interval: The interval in milliseconds, expressed as a 32-bit unsigned integer, for Heartbeat Messages. This value MUST NOT be 0.

As mentioned before, receipt of any valid DLEP Message MUST reset the heartbeat interval timer (i.e., valid DLEP Messages take the place of, and obviate the need for, additional Heartbeat Messages).

13.6. Extensions Supported

The Extensions Supported Data Item is used by the router and modem to negotiate additional optional functionality they are willing to support. The Extensions List is a concatenation of the types of each supported extension, found in the IANA registry titled "Extension Type Values". Each Extension Type definition includes which additional Signals and Data Items are supported.

The Extensions Supported Data Item contains the following fields:

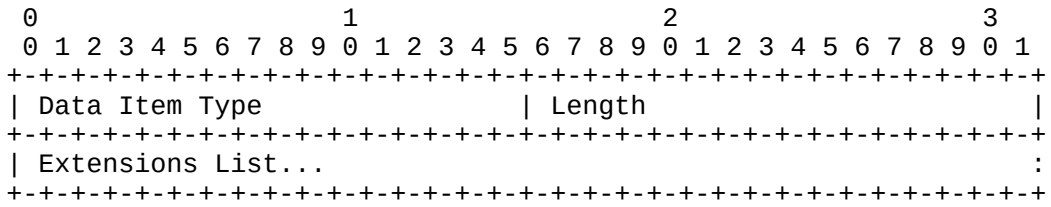

Data Item Type: 6

Length: Length of the Extensions List in octets. This is twice (2x) the number of extensions.

Extensions List: A list of extensions supported, identified by their 2-octet values as listed in the "Extension Type Values" registry.

13.7. MAC Address

The MAC Address Data Item contains the address of the destination on the remote node.

DLEP can support MAC addresses in either EUI-48 or EUI-64 format ("EUI" stands for "Extended Unique Identifier"), with the restriction that all MAC addresses for a given DLEP session MUST be in the same format and MUST be consistent with the MAC address format of the connected modem (e.g., if the modem is connected to the router with an EUI-48 MAC, all destination addresses via that modem MUST be expressed in EUI-48 format).

Examples of a virtual destination would be (1) a multicast MAC address or (2) the broadcast MAC address (FF:FF:FF:FF:FF:FF).

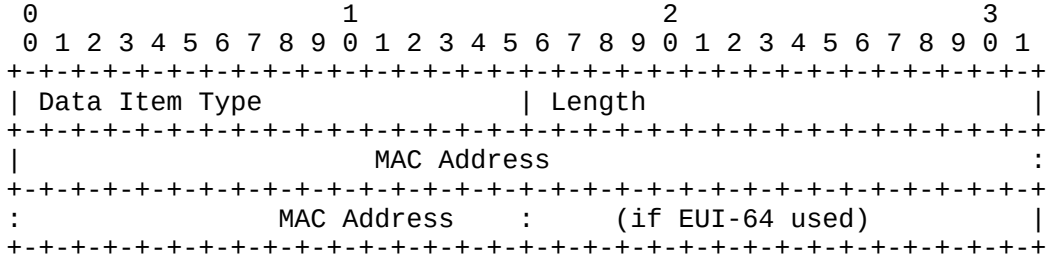

Data Item Type: 7

Length: 6 for EUI-48 format or 8 for EUI-64 format.

MAC Address: MAC address of the destination.

13.8. IPv4 Address

When included in the Session Update Message, this Data Item contains the IPv4 address of the peer. When included in Destination Messages, this Data Item contains the IPv4 address of the destination. In either case, the Data Item also contains an indication of whether this is (1) a new or existing address or (2) a deletion of a previously known address.

The IPv4 Address Data Item contains the following fields:

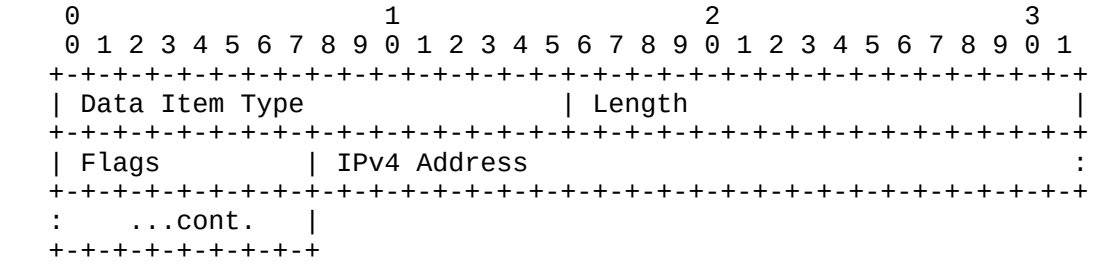

Data Item Type: 8

Length: 5

Flags: Flags field, defined below.

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

The Flags field is defined as:

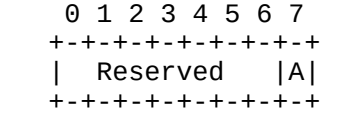

A: Add/Drop flag, indicating whether this is a new or existing address (1) or a withdrawal of an address (0).

Reserved: MUST be zero. Reserved for future use.

13.8.1. IPv4 Address Processing

Processing of the IPv4 Address Data Item MUST be done within the context of the DLEP peer session on which it is presented.

The handling of erroneous or logically inconsistent conditions depends upon the type of the message that contains the Data Item, as follows:

If the containing message is a Session Message, e.g., a Session Initialization Message (Section 12.5) or Session Update Message (Section 12.7), the receiver of inconsistent information MUST issue a Session Termination Message (Section 12.9) containing a Status Data Item (Section 13.1) with status code set to 130 'Invalid Data' and transition to the Session Termination state. Examples of such conditions are:

- An address Drop operation referencing an address that is not associated with the peer in the current session.
- An address Add operation referencing an address that has already been added to the peer in the current session.

If the containing message is a Destination Message, e.g., a Destination Up Message (Section 12.11) or Destination Update Message (Section 12.17), the receiver of inconsistent information MAY issue the appropriate response message containing a Status Data Item with status code set to 3 'Inconsistent Data' but MUST continue with session processing. Examples of such conditions are:

- An address Add operation referencing an address that has already been added to the destination in the current session.
- An address Add operation referencing an address that is associated with a different destination or the peer in the current session.
- An address Add operation referencing an address that makes no sense -- for example, defined as not forwardable in [RFC6890].
- An address Drop operation referencing an address that is not associated with the destination in the current session.

If no response message is appropriate -- for example, the Destination Update Message -- then the implementation MUST continue with session processing.

Modems that do not track IPv4 addresses MUST silently ignore IPv4 Address Data Items.

13.9. IPv6 Address

When included in the Session Update Message, this Data Item contains the IPv6 address of the peer. When included in Destination Messages, this Data Item contains the IPv6 address of the destination. In either case, the Data Item also contains an indication of whether this is (1) a new or existing address or (2) a deletion of a previously known address.

The IPv6 Address Data Item contains the following fields:

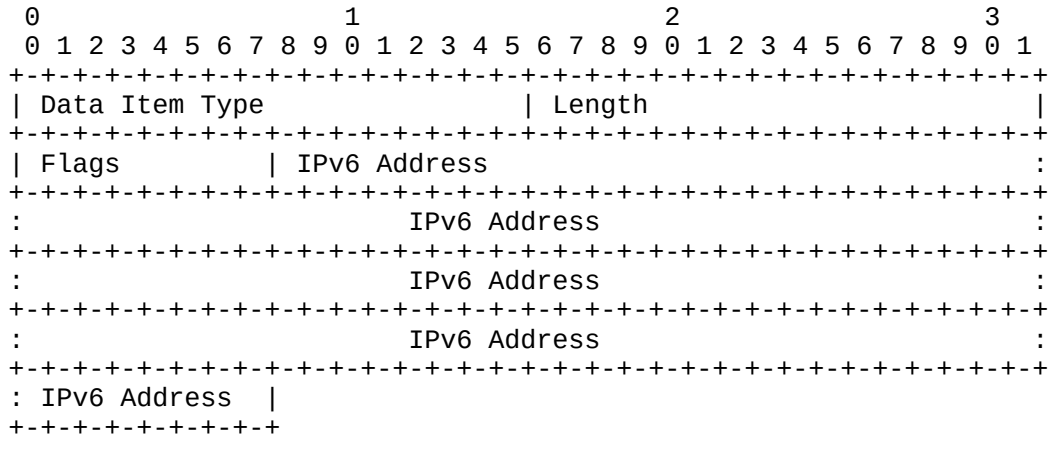

Data Item Type: 9

Length: 17

Flags: Flags field, defined below.

IPv6 Address: The IPv6 address of the destination or peer.

The Flags field is defined as:

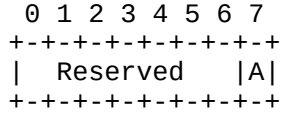

A: Add/Drop flag, indicating whether this is a new or existing address (1) or a withdrawal of an address (0).

Reserved: MUST be zero. Reserved for future use.

13.9.1. IPv6 Address Processing

Processing of the IPv6 Address Data Item MUST be done within the context of the DLEP peer session on which it is presented.

The handling of erroneous or logically inconsistent conditions depends upon the type of the message that contains the Data Item, as follows:

If the containing message is a Session Message, e.g., a Session Initialization Message (Section 12.5) or Session Update Message (Section 12.7), the receiver of inconsistent information MUST issue a Session Termination Message (Section 12.9) containing a Status Data Item (Section 13.1) with status code set to 130 'Invalid Data' and transition to the Session Termination state. Examples of such conditions are:

- An address Drop operation referencing an address that is not associated with the peer in the current session.
- An address Add operation referencing an address that has already been added to the peer in the current session.

If the containing message is a Destination Message, e.g., a Destination Up Message (Section 12.11) or Destination Update Message (Section 12.17), the receiver of inconsistent information MAY issue the appropriate response message containing a Status Data Item with status code set to 3 'Inconsistent Data' but MUST continue with session processing. Examples of such conditions are:

- An address Add operation referencing an address that has already been added to the destination in the current session.
- An address Add operation referencing an address that is associated with a different destination or the peer in the current session.
- An address Add operation referencing an address that makes no sense -- for example, defined as not forwardable in [RFC6890].
- An address Drop operation referencing an address that is not associated with the destination in the current session.

If no response message is appropriate -- for example, the Destination Update Message -- then the implementation MUST continue with session processing.

Modems that do not track IPv6 addresses MUST silently ignore IPv6 Address Data Items.

13.10. IPv4 Attached Subnet

The DLEP IPv4 Attached Subnet Data Item allows a device to declare that it has an IPv4 subnet (e.g., a stub network) attached, that it has become aware of an IPv4 subnet being present at a remote destination, or that it has become aware of the loss of a subnet at the remote destination.

The DLEP IPv4 Attached Subnet Data Item contains the following fields:

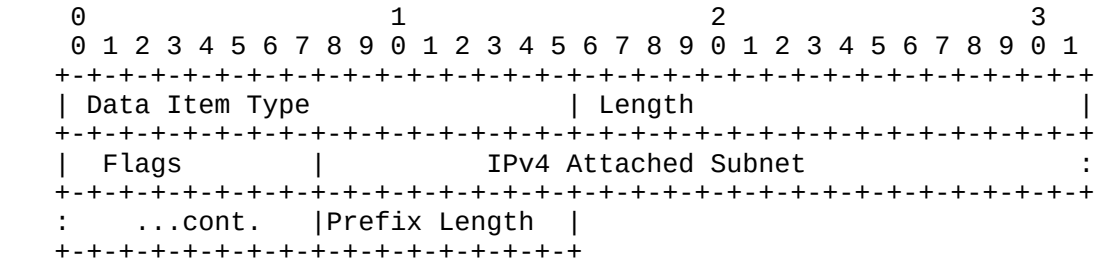

Data Item Type: 10

Length: 6

Flags: Flags field, defined below.

IPv4 Attached Subnet: The IPv4 subnet reachable at the destination.

Prefix Length: Length of the prefix (0-32) for the IPv4 subnet. A prefix length outside the specified range MUST be considered as invalid.

The Flags field is defined as:

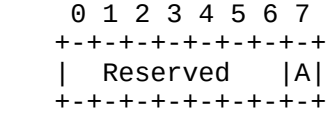

A: Add/Drop flag, indicating whether this is a new or existing subnet address (1) or a withdrawal of a subnet address (0).

Reserved: MUST be zero. Reserved for future use.

13.10.1. IPv4 Attached Subnet Processing

Processing of the IPv4 Attached Subnet Data Item MUST be done within the context of the DLEP peer session on which it is presented.

If the containing message is a Session Message, e.g., a Session Initialization Message (Section 12.5) or Session Update Message (Section 12.7), the receiver of inconsistent information MUST issue a Session Termination Message (Section 12.9) containing a Status Data Item (Section 13.1) with status code set to 130 'Invalid Data' and transition to the Session Termination state. Examples of such conditions are:

- A subnet Drop operation referencing a subnet that is not associated with the peer in the current session.
- A subnet Add operation referencing a subnet that has already been added to the peer in the current session.

If the containing message is a Destination Message, e.g., a Destination Up Message (Section 12.11) or Destination Update Message (Section 12.17), the receiver of inconsistent information MAY issue the appropriate response message containing a Status Data Item with status code set to 3 'Inconsistent Data' but MUST continue with session processing. Examples of such conditions are:

- A subnet Add operation referencing a subnet that has already been added to the destination in the current session.
- A subnet Add operation referencing a subnet that is associated with a different destination in the current session.
- A subnet Add operation referencing a subnet that makes no sense -- for example, defined as not forwardable in [RFC6890].
- A subnet Drop operation referencing a subnet that is not associated with the destination in the current session.

If no response message is appropriate -- for example, the Destination Update Message -- then the implementation MUST continue with session processing.

Modems that do not track IPv4 subnets MUST silently ignore IPv4 Attached Subnet Data Items.

13.11. IPv6 Attached Subnet

The DLEP IPv6 Attached Subnet Data Item allows a device to declare that it has an IPv6 subnet (e.g., a stub network) attached, that it has become aware of an IPv6 subnet being present at a remote destination, or that it has become aware of the loss of a subnet at the remote destination.

The DLEP IPv6 Attached Subnet Data Item contains the following fields:

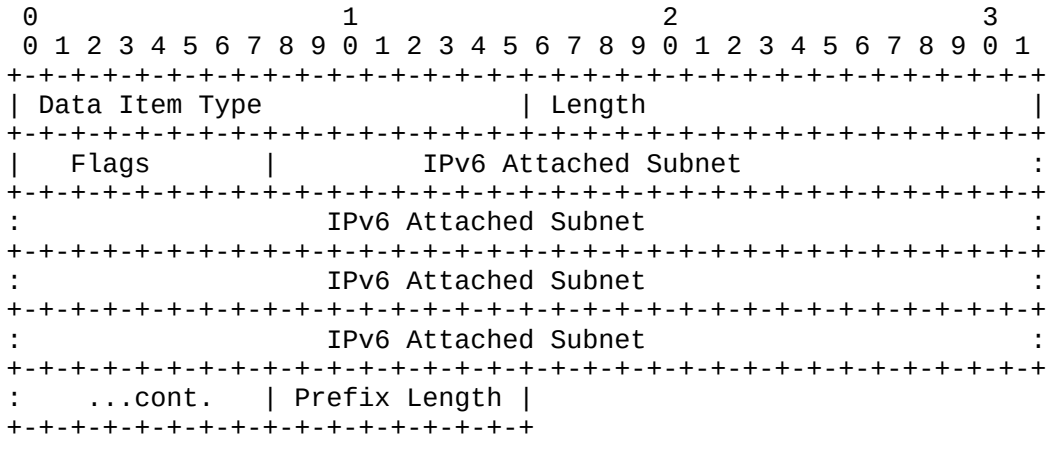

Data Item Type: 11

Length: 18

Flags: Flags field, defined below.

IPv6 Attached Subnet: The IPv6 subnet reachable at the destination.

Prefix Length: Length of the prefix (0-128) for the IPv6 subnet. A prefix length outside the specified range MUST be considered as invalid.

The Flags field is defined as:

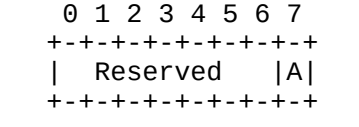

A: Add/Drop flag, indicating whether this is a new or existing subnet address (1) or a withdrawal of a subnet address (0).

Reserved: MUST be zero. Reserved for future use.

13.11.1. IPv6 Attached Subnet Processing

Processing of the IPv6 Attached Subnet Data Item MUST be done within the context of the DLEP peer session on which it is presented.

If the containing message is a Session Message, e.g., a Session Initialization Message (Section 12.5) or Session Update Message (Section 12.7), the receiver of inconsistent information MUST issue a Session Termination Message (Section 12.9) containing a Status Data Item (Section 13.1) with status code set to 130 'Invalid Data' and transition to the Session Termination state. Examples of such conditions are:

- A subnet Drop operation referencing a subnet that is not associated with the peer in the current session.
- A subnet Add operation referencing a subnet that has already been added to the peer in the current session.

If the containing message is a Destination Message, e.g., a Destination Up Message (Section 12.11) or Destination Update Message (Section 12.17), the receiver of inconsistent information MAY issue the appropriate response message containing a Status Data Item with status code set to 3 'Inconsistent Data' but MUST continue with session processing. Examples of such conditions are:

- A subnet Add operation referencing a subnet that has already been added to the destination in the current session.
- A subnet Add operation referencing a subnet that is associated with a different destination in the current session.

- A subnet Add operation referencing a subnet that makes no sense -- for example, defined as not forwardable in [RFC6890].

- A subnet Drop operation referencing a subnet that is not associated with the destination in the current session.

If no response message is appropriate -- for example, the Destination Update Message -- then the implementation MUST continue with session processing.

Modems that do not track IPv6 subnets MUST silently ignore IPv6 Attached Subnet Data Items.

13.12. Maximum Data Rate (Receive)

The Maximum Data Rate (Receive) (MDRR) Data Item is used to indicate the maximum theoretical data rate, in bits per second (bps), that can be achieved while receiving data on the link.

The Maximum Data Rate (Receive) Data Item contains the following fields:

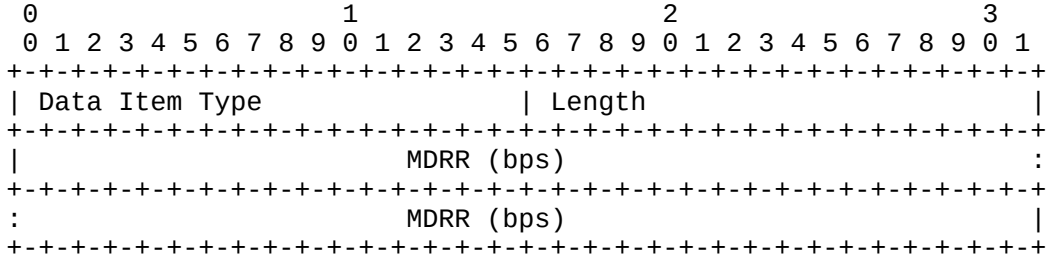

Data Item Type: 12

Length: 8

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

13.13. Maximum Data Rate (Transmit)

The Maximum Data Rate (Transmit) (MDRT) Data Item is used to indicate the maximum theoretical data rate, in bits per second, that can be achieved while transmitting data on the link.

The Maximum Data Rate (Transmit) Data Item contains the following fields:

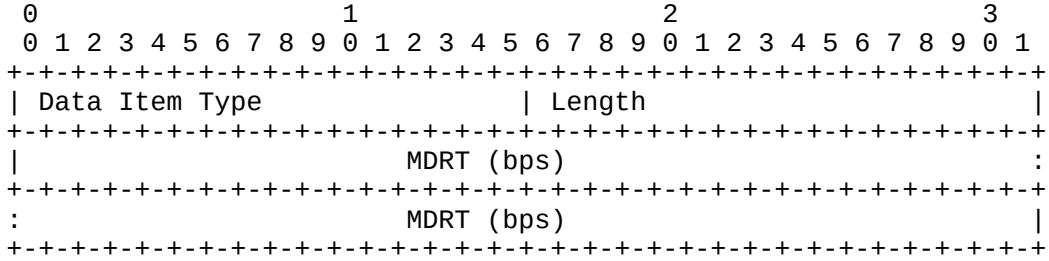

Data Item Type: 13

Length: 8

Maximum Data Rate (Transmit): A 64-bit unsigned integer, representing the maximum theoretical data rate, in bits per second, that can be achieved while transmitting on the link.

13.14. Current Data Rate (Receive)

The Current Data Rate (Receive) (CDRR) Data Item is used to indicate the rate at which the link is currently operating for receiving traffic.

When used in the Link Characteristics Request Message (Section 12.18), Current Data Rate (Receive) represents the desired receive rate, in bits per second, on the link.

The Current Data Rate (Receive) Data Item contains the following fields:

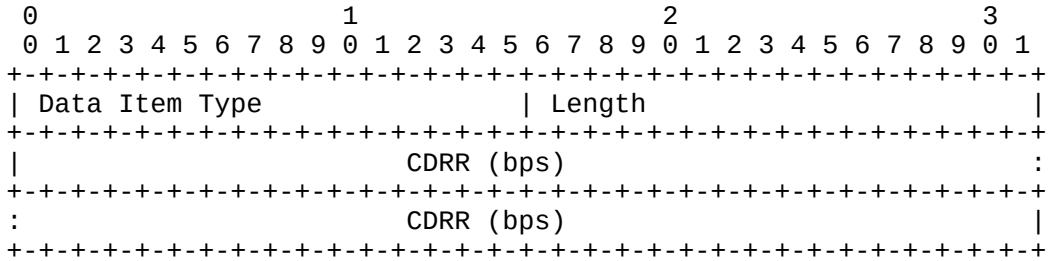

Data Item Type: 14

Length: 8

Current Data Rate (Receive): A 64-bit unsigned integer, representing the current data rate, in bits per second, that can currently be achieved while receiving traffic on the link.

If there is no distinction between Current Data Rate (Receive) and Maximum Data Rate (Receive) (Section 13.12), Current Data Rate (Receive) MUST be set equal to Maximum Data Rate (Receive). Current Data Rate (Receive) MUST NOT exceed Maximum Data Rate (Receive).

13.15. Current Data Rate (Transmit)

The Current Data Rate (Transmit) (CDRT) Data Item is used to indicate the rate at which the link is currently operating for transmitting traffic.

When used in the Link Characteristics Request Message (Section 12.18), Current Data Rate (Transmit) represents the desired transmit rate, in bits per second, on the link.

The Current Data Rate (Transmit) Data Item contains the following fields:

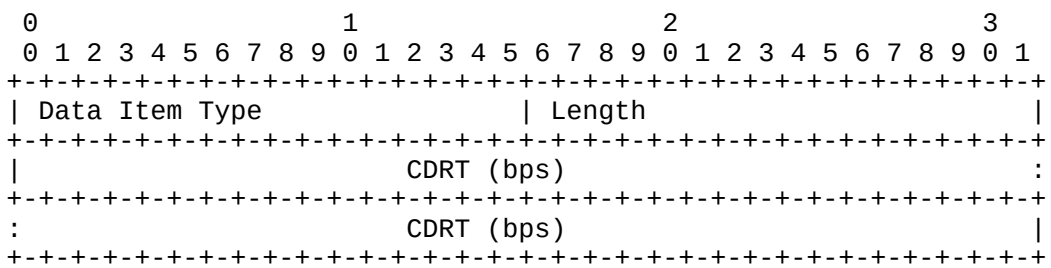

Data Item Type: 15

Length: 8

Current Data Rate (Transmit): A 64-bit unsigned integer, representing the current data rate, in bits per second, that can currently be achieved while transmitting traffic on the link.

If there is no distinction between Current Data Rate (Transmit) and Maximum Data Rate (Transmit) (Section 13.13), Current Data Rate (Transmit) MUST be set equal to Maximum Data Rate (Transmit). Current Data Rate (Transmit) MUST NOT exceed Maximum Data Rate (Transmit).

13.16. Latency

The Latency Data Item is used to indicate the amount of latency, in microseconds, on the link.

The Latency value is reported as transmission delay. The calculation of latency is implementation dependent. For example, the latency may be a running average calculated from the internal queuing.

The Latency Data Item contains the following fields:

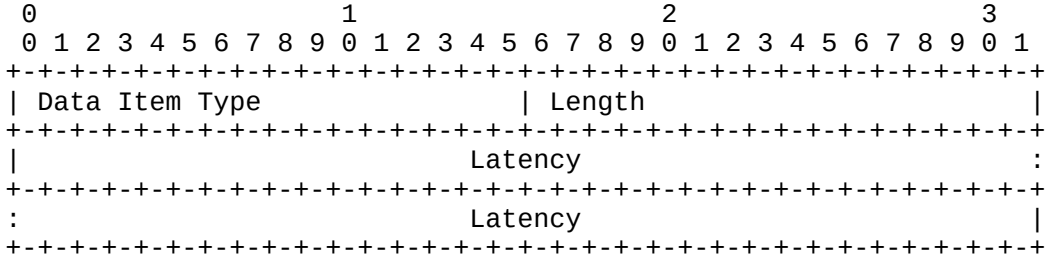

Data Item Type: 16

Length: 8

Latency: A 64-bit unsigned integer, representing the transmission delay, in microseconds, that a packet encounters as it is transmitted over the link.

13.17. Resources

The Resources (RES) Data Item is used to indicate the amount of finite resources available for data transmission and reception at the destination as a percentage, with 0 meaning 'no resources remaining' and 100 meaning 'a full supply', assuming that when Resources reaches 0 data transmission and/or reception will cease.

An example of such resources is battery life, but this could also include resources such as available memory for queuing, or CPU idle percentage. The specific criteria to be used for this metric is out of scope for this specification and is implementation specific.

This Data Item is designed to be used as an indication of some capability of the modem and/or router at the destination.

The Resources Data Item contains the following fields:

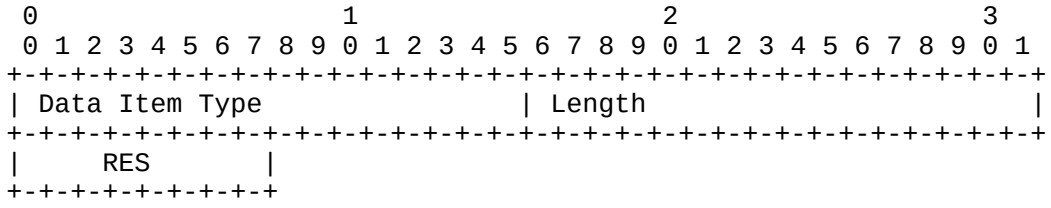

Data Item Type: 17

Length: 1

Resources: An 8-bit unsigned integer percentage, 0-100, representing the amount of resources available. Any value greater than 100 MUST be considered as invalid.

If a device cannot calculate Resources, this Data Item MUST NOT be issued.

13.18. Relative Link Quality (Receive)

The Relative Link Quality (Receive) (RLQR) Data Item is used to indicate the quality of the link to a destination for receiving traffic, with 0 meaning 'worst quality' and 100 meaning 'best quality'.

Quality in this context is defined as an indication of the stability of a link for reception; a destination with high Relative Link Quality (Receive) is expected to have generally stable DLEP metrics, and the metrics of a destination with low Relative Link Quality (Receive) can be expected to rapidly fluctuate over a wide range.

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

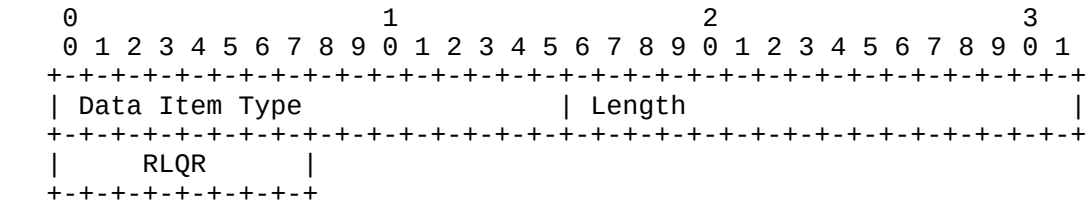

Data Item Type: 18

Length: 1

Relative Link Quality (Receive): A non-dimensional unsigned 8-bit integer, 0-100, representing relative quality of the link for receiving traffic. Any value greater than 100 MUST be considered as invalid.

If a device cannot calculate Relative Link Quality (Receive), this Data Item MUST NOT be issued.

13.19. Relative Link Quality (Transmit)

The Relative Link Quality (Transmit) (RLQT) Data Item is used to indicate the quality of the link to a destination for transmitting traffic, with 0 meaning 'worst quality' and 100 meaning 'best quality'.

Quality in this context is defined as an indication of the stability of a link for transmission; a destination with high Relative Link Quality (Transmit) is expected to have generally stable DLEP metrics, and the metrics of a destination with low Relative Link Quality (Transmit) can be expected to rapidly fluctuate over a wide range.

The Relative Link Quality (Transmit) Data Item contains the following fields:

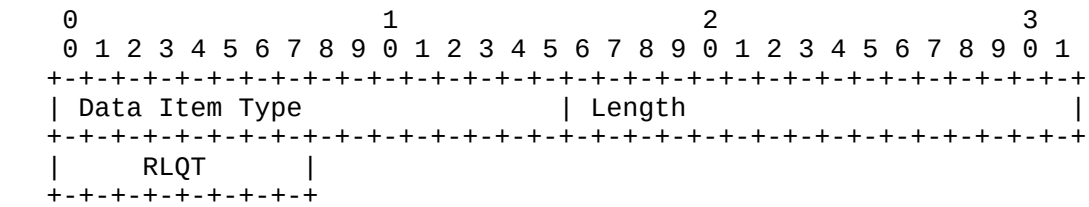

Data Item Type: 19

Length: 1

Relative Link Quality (Transmit): A non-dimensional unsigned 8-bit integer, 0-100, representing relative quality of the link for transmitting traffic. Any value greater than 100 MUST be considered as invalid.

If a device cannot calculate Relative Link Quality (Transmit), this Data Item MUST NOT be issued.

13.20. Maximum Transmission Unit (MTU)

The Maximum Transmission Unit (MTU) Data Item is used to indicate the maximum size, in octets, of an IP packet that can be transmitted without fragmentation, including headers, but excluding any lower-layer headers.

The Maximum Transmission Unit Data Item contains the following fields:

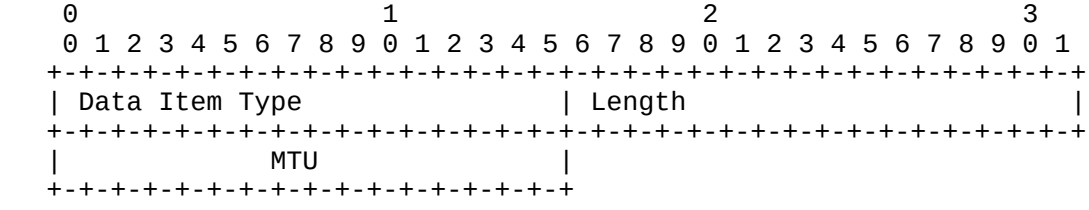

Data Item Type: 20

Length: 2

Maximum Transmission Unit: The maximum size, in octets, of an IP packet that can be transmitted without fragmentation, including headers, but excluding any lower-layer headers.

If a device cannot calculate Maximum Transmission Unit, this Data Item MUST NOT be issued.

14. Security Considerations

The potential security concerns when using DLEP are as follows:

1. An attacker might pretend to be a DLEP participant, either at DLEP session initialization or by injection of DLEP Messages once a session has been established.

2. DLEP Data Items could be altered by an attacker, causing the receiving implementation to inappropriately alter its information base concerning network status.

3. An attacker could join an unsecured radio network and inject over-the-air signals that maliciously influence the information reported by a DLEP modem, causing a router to forward traffic to an inappropriate destination.

The implications of attacks on DLEP peers are directly proportional to the extent to which DLEP data is used within the control plane. While the use of DLEP data in other control-plane components is out of scope for this document, as an example, if DLEP statistics are incorporated into route cost calculations, adversaries masquerading as a DLEP peer and injecting malicious data via DLEP could cause suboptimal route selection, adversely impacting network performance. Similar issues can arise if DLEP data is used as an input to policing algorithms -- injection of malicious data via DLEP can cause those policing algorithms to make incorrect decisions, degrading network throughput.

For these reasons, security of the DLEP transport must be considered at both the transport layer and Layer 2.

At the transport layer, when TLS is in use, each peer SHOULD check the validity of credentials presented by the other peer during TLS session establishment. Implementations following the "dedicated deployments" model attempting to use TLS MAY (1) need to consider the use of pre-shared keys for credentials, (2) provide specialized techniques for peer identity validation, and (3) refer to [RFC5487] for additional details. Implementations following the "networked deployment" model described in "Implementation Scenarios" (Section 4) SHOULD refer to [RFC7525] for additional details.

At Layer 2, since DLEP is restricted to operation over a single (possibly logical) hop, implementations SHOULD also secure the Layer 2 link. Examples of technologies that can be deployed to secure the Layer 2 link include [IEEE-802.1AE] and [IEEE-802.1X].

By examining the Secured Medium flag in the Peer Type Data Item (Section 13.4), a router can decide if it is able to trust the information supplied via a DLEP modem. If this is not the case, then the router SHOULD consider restricting the size of attached subnets, announced in IPv4 Attached Subnet Data Items (Section 13.10) and/or IPv6 Attached Subnet Data Items (Section 13.11), that are considered for route selection.

To avoid potential denial-of-service attacks, it is RECOMMENDED that implementations using the Peer Discovery mechanism (1) maintain an information base of hosts that persistently fail Session Initialization, even though those hosts have provided an acceptable Peer Discovery Signal and (2) ignore any subsequent Peer Discovery Signals from such hosts.

This specification does not address security of the data plane, as it (the data plane) is not affected, and standard security procedures can be employed.

15. IANA Considerations

15.1. Registrations

IANA has created a new protocol registry for the Dynamic Link Exchange Protocol (DLEP). The remainder of this section details the new DLEP-specific registries.

15.2. Signal Type Registrations

IANA has created a new DLEP registry, named "Signal Type Values".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Type Code	Description/Policy
0	Reserved
1	Peer Discovery Signal
2	Peer Offer Signal
3-65519	Unassigned / Specification Required
65520-65534	Reserved for Private Use
65535	Reserved

15.3. Message Type Registrations

IANA has created a new DLEP registry, named "Message Type Values".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Type Code	Description/Policy
0	Reserved
1	Session Initialization Message
2	Session Initialization Response Message
3	Session Update Message
4	Session Update Response Message
5	Session Termination Message
6	Session Termination Response Message
7	Destination Up Message
8	Destination Up Response Message
9	Destination Announce Message
10	Destination Announce Response Message
11	Destination Down Message

12	Destination Down Response Message
13	Destination Update Message
14	Link Characteristics Request Message
15	Link Characteristics Response Message
16	Heartbeat Message
17-65519	Unassigned / Specification Required
65520-65534	Reserved for Private Use
65535	Reserved

15.4. DLEP Data Item Registrations

IANA has created a new DLEP registry, named "Data Item Type Values".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Type Code	Description/Policy
0	Reserved
1	Status
2	IPv4 Connection Point
3	IPv6 Connection Point
4	Peer Type
5	Heartbeat Interval
6	Extensions Supported
7	MAC Address
8	IPv4 Address
9	IPv6 Address
10	IPv4 Attached Subnet

11	IPv6 Attached Subnet
12	Maximum Data Rate (Receive) (MDRR)
13	Maximum Data Rate (Transmit) (MDRT)
14	Current Data Rate (Receive) (CDRR)
15	Current Data Rate (Transmit) (CDRT)
16	Latency
17	Resources (RES)
18	Relative Link Quality (Receive) (RLQR)
19	Relative Link Quality (Transmit) (RLQT)
20	Maximum Transmission Unit (MTU)
21-65407	Unassigned / Specification Required
65408-65534	Reserved for Private Use
65535	Reserved

15.5. DLEP Status Code Registrations

IANA has created a new DLEP registry, named "Status Code Values".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Status Code	Failure Mode	Description/Policy
0	Continue	Success
1	Continue	Not Interested
2	Continue	Request Denied
3	Continue	Inconsistent Data
4-111	Continue	Unassigned / Specification Required

112-127	Continue	Private Use
128	Terminate	Unknown Message
129	Terminate	Unexpected Message
130	Terminate	Invalid Data
131	Terminate	Invalid Destination
132	Terminate	Timed Out
133-239	Terminate	Unassigned / Specification Required
240-254	Terminate	Reserved for Private Use
255	Terminate	Shutting Down

15.6. DLEP Extension Registrations

IANA has created a new DLEP registry, named "Extension Type Values".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Code	Description/Policy
0	Reserved
1-65519	Unassigned / Specification Required
65520-65534	Reserved for Private Use
65535	Reserved

Table 3: DLEP Extension Types

15.7. DLEP IPv4 Connection Point Flags

IANA has created a new DLEP registry, named "IPv4 Connection Point Flags".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Bit	Description/Policy
0-6	Unassigned / Specification Required
7	Use TLS [RFC5246] indicator

15.8. DLEP IPv6 Connection Point Flags

IANA has created a new DLEP registry, named "IPv6 Connection Point Flags".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Bit	Description/Policy
0-6	Unassigned / Specification Required
7	Use TLS [RFC5246] indicator

15.9. DLEP Peer Type Flags

IANA has created a new DLEP registry, named "Peer Type Flags".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Bit	Description/Policy
0-6	Unassigned / Specification Required
7	Secured Medium indicator

15.10. DLEP IPv4 Address Flags

IANA has created a new DLEP registry, named "IPv4 Address Flags".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Bit	Description/Policy
0-6	Unassigned / Specification Required
7	Add/Drop indicator

15.11. DLEP IPv6 Address Flags

IANA has created a new DLEP registry, named "IPv6 Address Flags".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Bit	Description/Policy
0-6	Unassigned / Specification Required
7	Add/Drop indicator

15.12. DLEP IPv4 Attached Subnet Flags

IANA has created a new DLEP registry, named "IPv4 Attached Subnet Flags".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Bit	Description/Policy
0-6	Unassigned / Specification Required
7	Add/Drop indicator

15.13. DLEP IPv6 Attached Subnet Flags

IANA has created a new DLEP registry, named "IPv6 Attached Subnet Flags".

The following table provides initial registry values and the policies, as defined by [RFC5226], that apply to the registry:

Bit	Description/Policy
0-6	Unassigned / Specification Required
7	Add/Drop indicator

15.14. DLEP Well-Known Port

IANA has assigned the value 854 in the "Service Name and Transport Protocol Port Number Registry" found at <http://www.iana.org/assignments/service-names-port-numbers/> for use by "DLEP", as defined in this document. This assignment is valid for TCP and UDP.

15.15. DLEP IPv4 Link-Local Multicast Address

IANA has assigned the IPv4 multicast address 224.0.0.117 in the registry found at <http://www.iana.org/assignments/multicast-addresses> for use as "DLEP Discovery".

15.16. DLEP IPv6 Link-Local Multicast Address

IANA has assigned the IPv6 multicast address FF02:0:0:0:0:0:1:7 in the registry found at <http://www.iana.org/assignments/ipv6-multicast-addresses> for use as "DLEP Discovery".

16. References

16.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

[RFC3629] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, DOI 10.17487/RFC3629, November 2003, <http://www.rfc-editor.org/info/rfc3629>.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, DOI 10.17487/RFC5082, October 2007, <http://www.rfc-editor.org/info/rfc5082>.

[RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, DOI 10.17487/RFC5246, August 2008, <http://www.rfc-editor.org/info/rfc5246>.

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

16.2. Informative References

[IEEE-802.1AE] "IEEE Standards for Local and Metropolitan Area Networks: Media Access Control (MAC) Security", DOI 10.1109/IEEESTD.2006.245590, <http://ieeexplore.ieee.org/document/1678345/>.

[IEEE-802.1X] "IEEE Standards for Local and metropolitan area networks-- Port-Based Network Access Control", DOI 10.1109/IEEESTD.2010.5409813, <http://ieeexplore.ieee.org/document/5409813/>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, DOI 10.17487/RFC5226, May 2008, <http://www.rfc-editor.org/info/rfc5226>.

[RFC5487] Badra, M., "Pre-Shared Key Cipher Suites for TLS with SHA-256/384 and AES Galois Counter Mode", RFC 5487, DOI 10.17487/RFC5487, March 2009, <http://www.rfc-editor.org/info/rfc5487>.

[RFC6890] Cotton, M., Vegoda, L., Bonica, R., Ed., and B. Haberman, "Special-Purpose IP Address Registries", BCP 153, RFC 6890, DOI 10.17487/RFC6890, April 2013, <http://www.rfc-editor.org/info/rfc6890>.

[RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <http://www.rfc-editor.org/info/rfc7525>.

Appendix A. Discovery Signal Flows

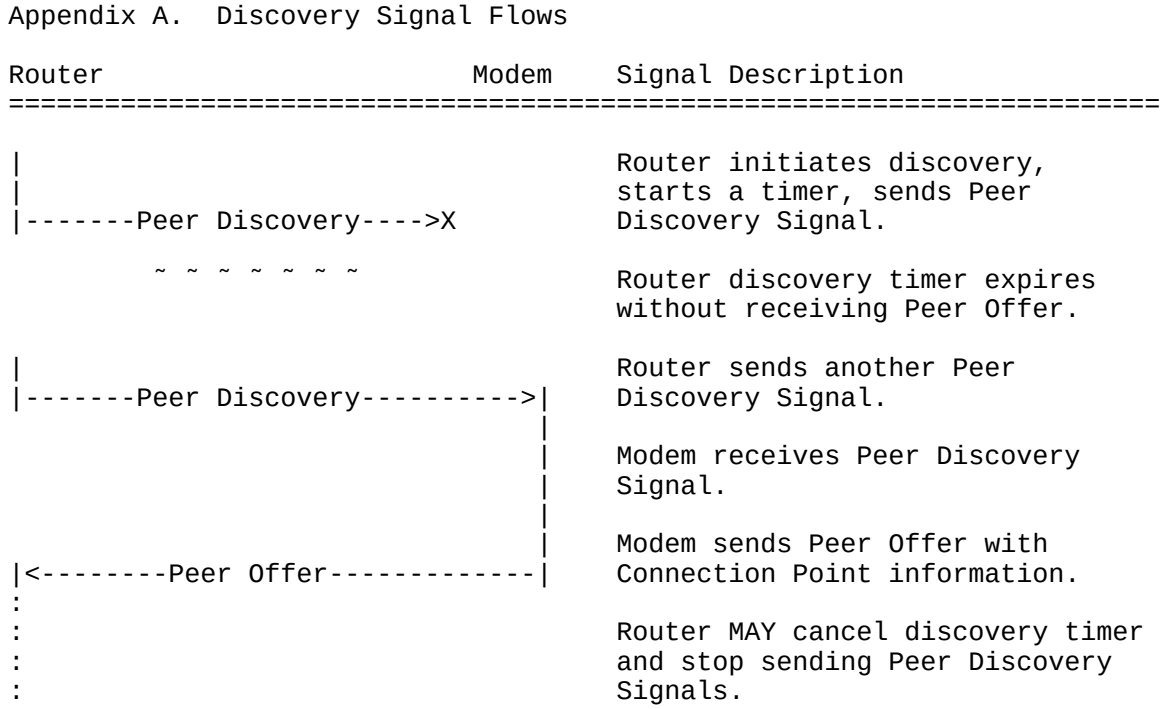

Appendix B. Peer-Level Message Flows

B.1. Session Initialization

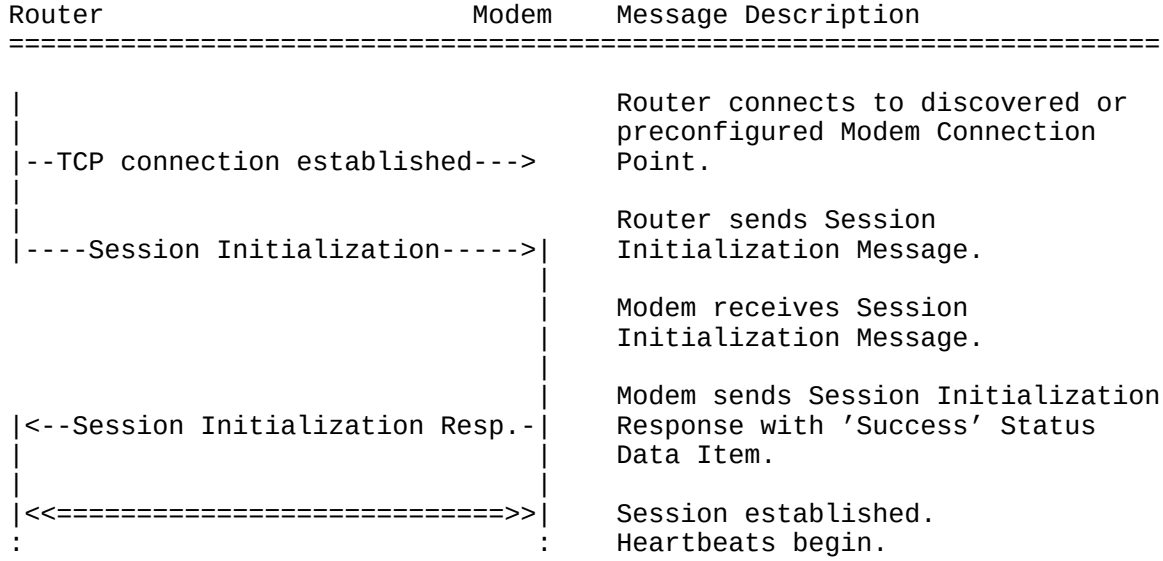

B.2. Session Initialization - Refused

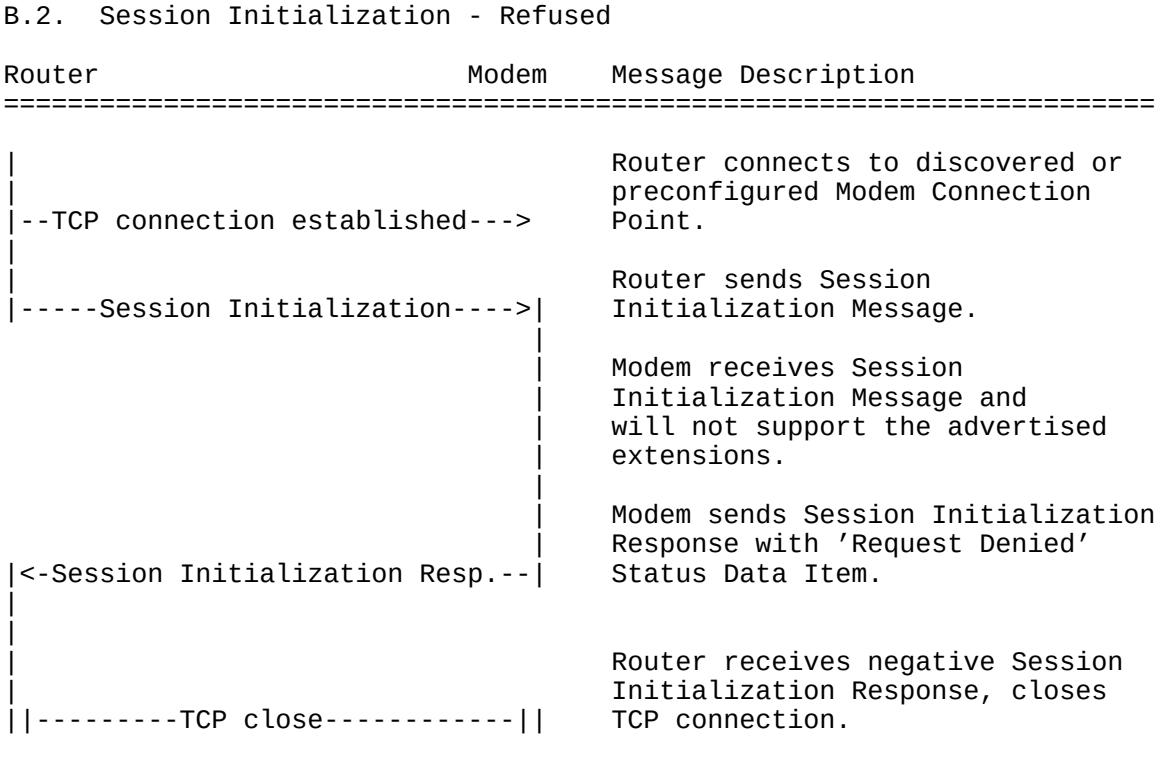

B.3. Router Changes IP Addresses

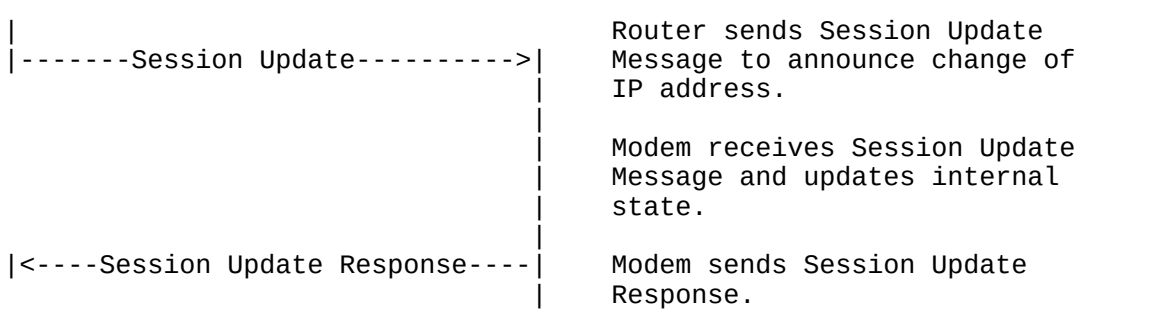

B.4. Modem Changes Session-Wide Metrics

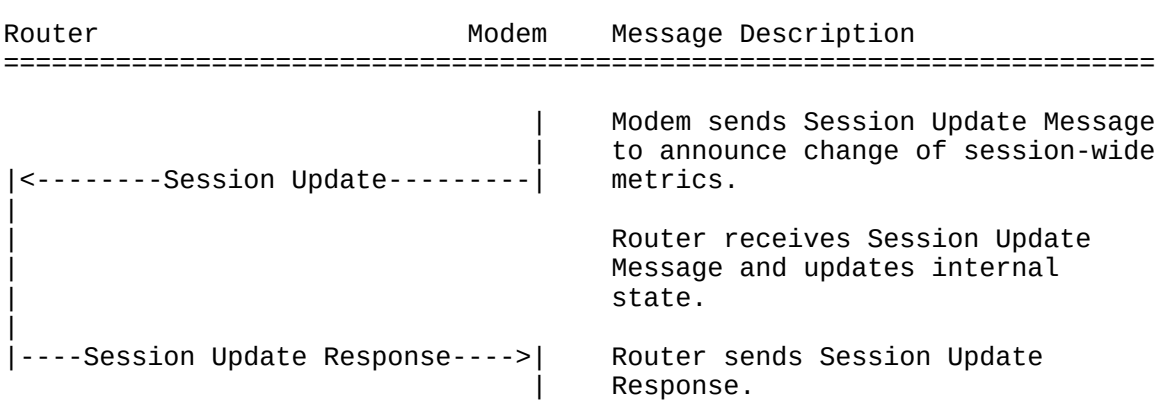

B.5. Router Terminates Session

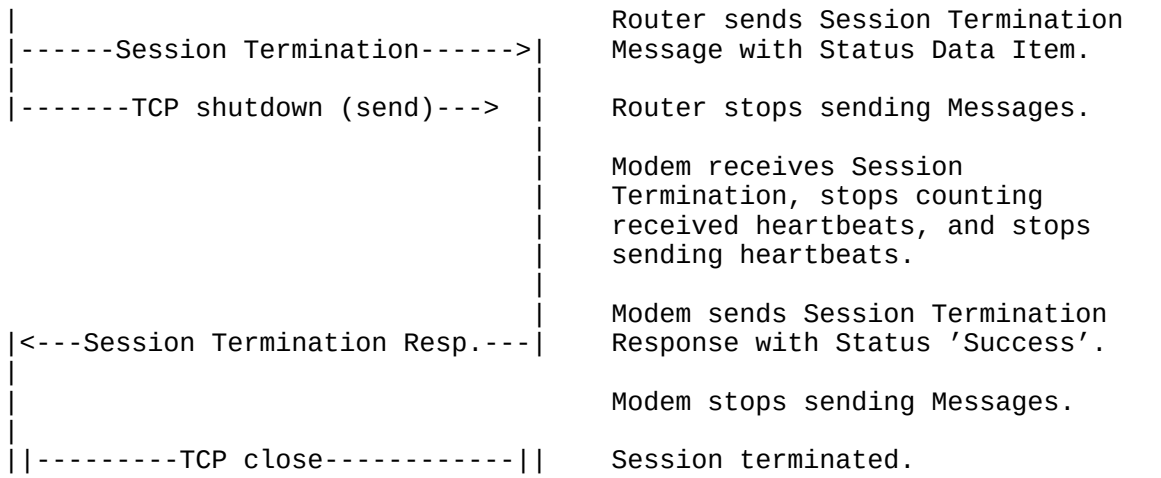

B.6. Modem Terminates Session

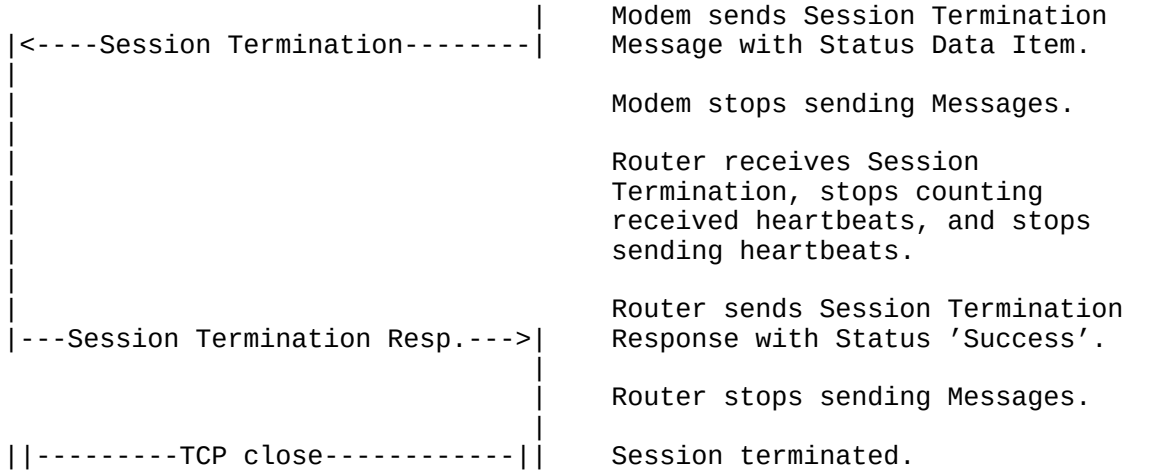

B.7. Session Heartbeats

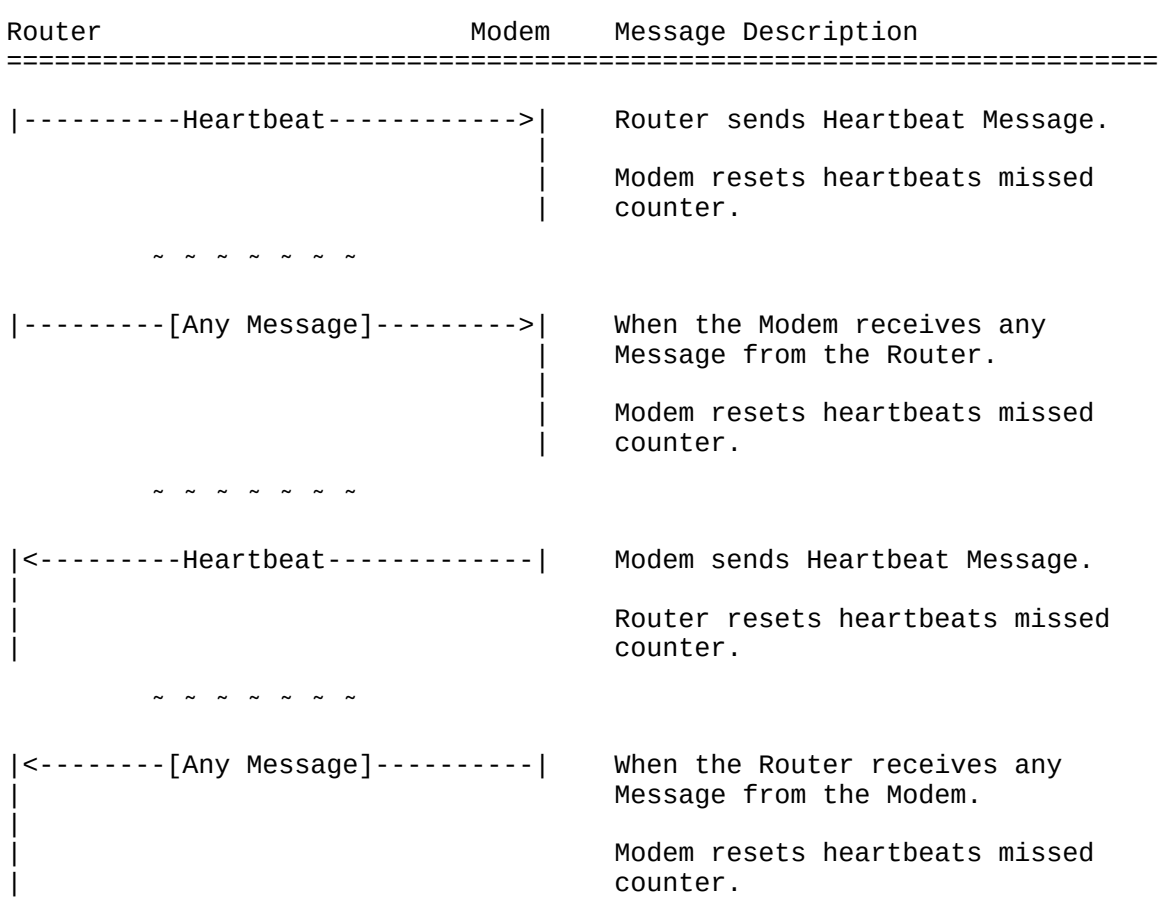

B.8. Router Detects a Heartbeat Timeout

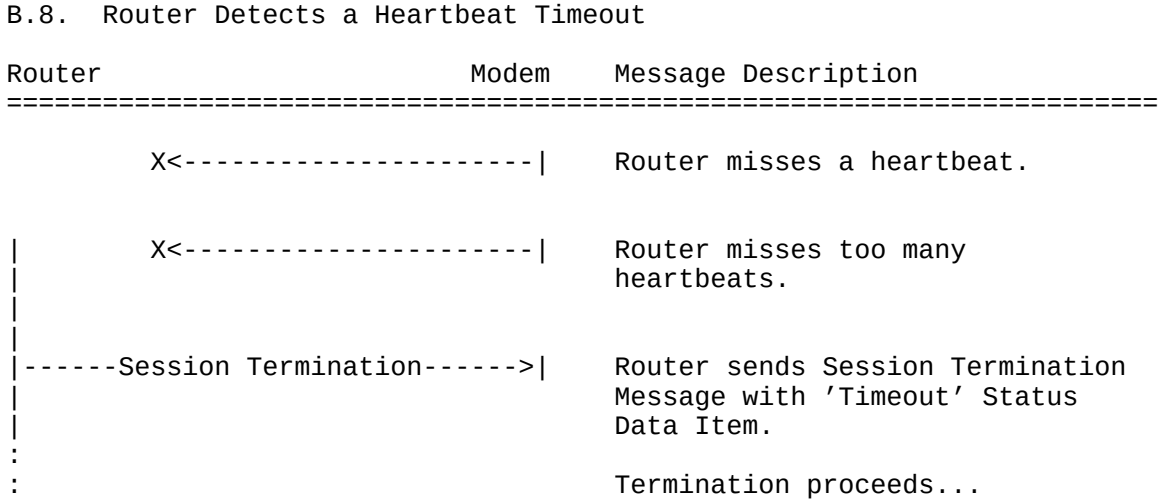

B.9. Modem Detects a Heartbeat Timeout

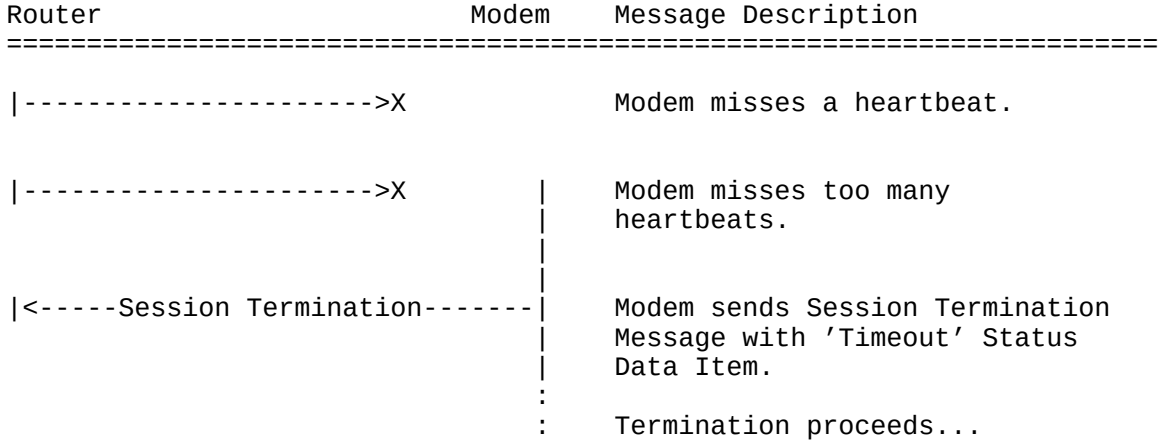

Appendix C. Destination-Specific Message Flows

C.1. Common Destination Notification

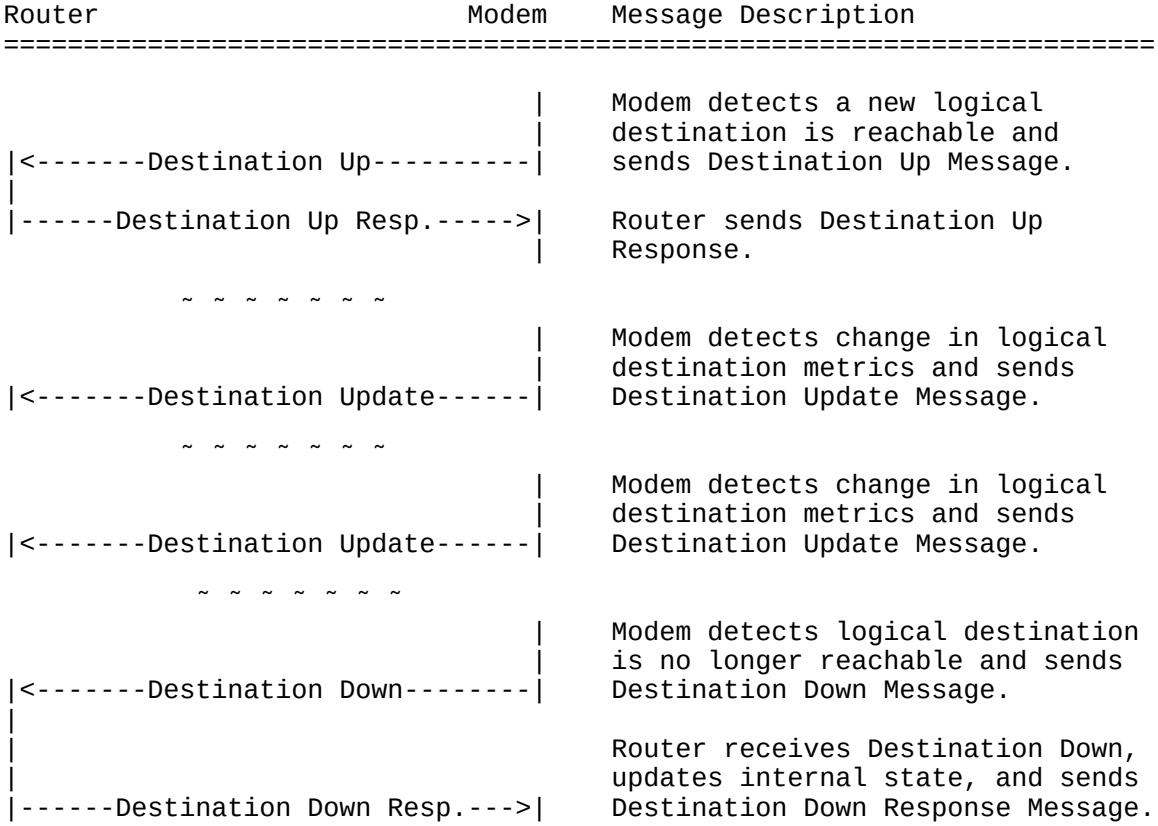

C.2. Multicast Destination Notification

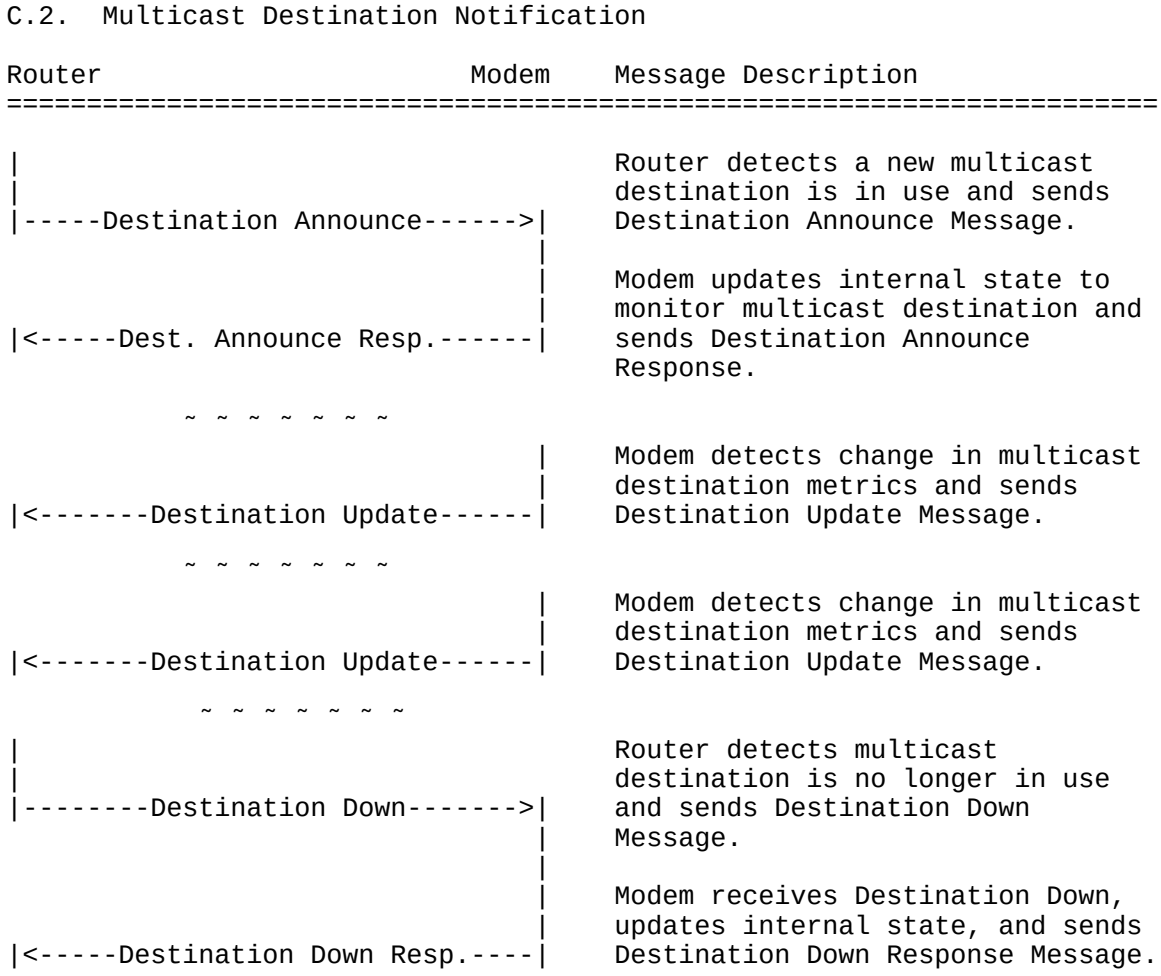

C.3. Link Characteristics Request

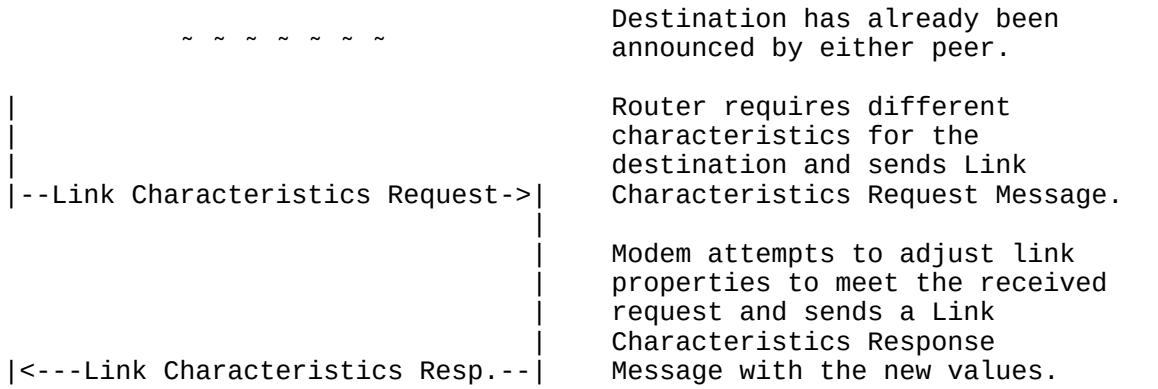

Acknowledgments

We 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, and Henning Rogge.

We would also like to acknowledge the influence and contributions of Greg Harrison, Chris Olsen, Martin Duke, Subir Das, Jaewon Kang, Vikram Kaul, Nelson Powell, Lou Berger, and Victoria Pritchard.

Authors' Addresses

Stan Ratliff
VT iDirect
13861 Sunrise Valley Drive, Suite 300
Herndon, VA 20171
United States of America
Email: sratliff@idirect.net

Shawn Jury
Cisco Systems
170 West Tasman Drive
San Jose, CA 95134
United States of America
Email: sjury@cisco.com

Darryl Satterwhite
Broadcom
Email: dsatterw@broadcom.com

Rick Taylor
Airbus Defence & Space
Quadrant House
Celtic Springs
Coedkernew
Newport NP10 8FZ
United Kingdom
Email: rick.taylor@airbus.com

Bo Berry