Mobile Ad hoc Networks Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 14, 2011

S. Ratliff
B. Berry
G. Harrison
S. Jury
D. Satterwhite
Cisco Systems

September 14, 2010

Dynamic Link Exchange Protocol (DLEP)
draft-sratliff-dlep-01

Abstract

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

Status of this Memo

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

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

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

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

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

This Internet-Draft will expire on March 14, 2011 .

Table of Contents

1. Introduction

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

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

Another complicating factor for mobile networks are the different methods of physically connecting the modem devices to the router. Modems can be deployed as an interface card in a router's chassis, or as a standalone device connected to the router via Ethernet, USB, or even a serial link. In the case of Ethernet or serial attachment, with existing protocols and techniques, routing software cannot be aware of convergence events occurring on the radio link (e.g. acquisition or loss of a potential routing

neighbor), nor can the router be aware of the actual capacity of the link. This lack of awareness, along with the variability in bandwidth, leads to a situation where quality of service (QoS) profiles are extremely difficult to establish and properly maintain. This is especially true of demand-based access schemes

such as Demand Assigned Multiple Access (DAMA) implementations used on some satellite systems. With a DAMA-based system, additional bandwidth may be available, but will not be used unless the network devices emit traffic at rate higher than the currently established rate. Increasing the traffic rate does not guarantee additional bandwidth will be allocated; rather, it may result in data loss and additional retransmissions on the link.

In attempting to address the challenges listed above, the authors have developed the Data Link Exchange Protocol, or DLEP. The DLEP protocol runs between a router and its attached modem devices, allowing the modem to communicate link characteristics as they change, and convergence events (acquisition and loss of potential routing neighbors). The diagram below is used to illustrate the scope of DLEP sessions. When a local client (Modem device) detects the presence of a remote neighbor, it sends an indication to its local router via the DLEP session. Upon receipt of the indication, the local router would take appropriate action (e.g. initiation of discovery or HELLO protocols) to converge the network. After notification of the new neighbor, the modem device utilizes the DLEP session to report the characteristics of the link (bandwidth, latency, etc) to the router on an as-needed basis. Finally, the Modem is able to use the DLEP session to notify the router when the remote neighbor is lost, shortening the time required to re-converge the network.

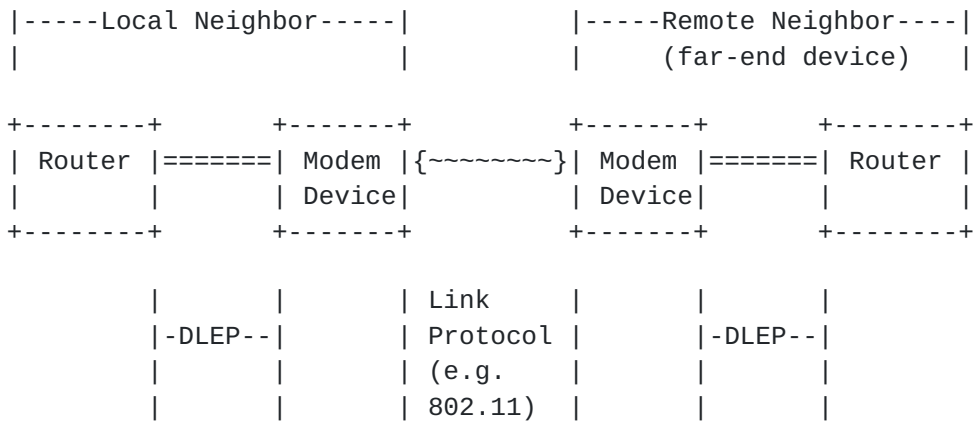

Figure 1: DLEP Network

DLEP exists as a collection of type-length-value (TLV) based messages formatted using RFC 5444. The protocol can be used for both Ethernet-attached modems (utilizing, for example, a UDP socket for transport of the RFC 5444 packets), or in environments where the modem is an interface card in a chassis (via a message passing scheme). DLEP utilizes a session paradigm between the modem device and its

associated router. If multiple modem devices are attached to a router, a separate DLEP session MUST exist for each modem. If a modem device supports multiple connections to a router (via multiple interfaces), or supports connections to multiple routers, a separate DLEP session MUST exist for each connection.

1.1 Requirements

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

2. Assumptions

In order to implement discovery in the DLEP protocol (thereby avoiding some configuration), we have defined a first-speaker and a passive-listener. Specifically, the router is defined as the passive-listener, and the modem device defined as the first-speaker (e.g. the initiator for discovery). Borrowing from existing terminology, this document refers to the first-speaker as the 'client', even though there is no client/server relationship in the classic sense.

DLEP assumes that participating modem devices appear to the router as a transparent bridge - specifically, the assumption is that the destination MAC address for data traffic in any frame emitted by the router should be the MAC address of the next-hop router or end-device, and not the MAC address of any of the intervening modem devices.

DLEP assumes that security on the session (e.g. authentication of session partners, encryption of traffic, or both) is dealt with by the underlying transport mechanism for the RFC 5444 packets (e.g. by using a transport such as DTLS [DTLS]).

The RFC 5444 message header Sequence Number MUST be included in all DLEP packets. Sequence Numbers start at 1 and are incremented by one for each original and retransmitted message. The unsigned 16-bit Sequence Number rolls over at 65535 to 1. A Sequence Number of 0 is not valid. Peer level Sequence Numbers are unique within the context of a DLEP session. Sequence numbers are used in DLEP to correlate a response to a request.

3. Normal Session Flow

A session between a router and a client is established by exchanging the "Peer Discovery" and "Peer Offer" messages described below.

Once that exchange has successfully occurred, the client informs the router of the presence of a new potential routing partner via the "Neighbor Up" message. The loss of a neighbor is communicated via the "Neighbor Down" message, and link quality is communicated via the "Neighbor Update" message. Note that, due to the issue of metrics varying depending on neighbor (discussed above), DLEP link metrics are expressed within the context of a neighbor relationship, instead

of on the link as a whole.

Once the DLEP session has started, the session partners exchange heartbeat messages based on a negotiated time interval. The heartbeat

messages are used to assure the session partners are in an appropriate state, and that bidirectional connectivity still exists.

In addition to receiving metrics about the link, DLEP provides for the ability for the router to request a different amount of bandwidth, or latency, for its client via the Link Characteristics Message. This allows the router to deal with requisite increases "Neighbor Update" message. This allows the router to request an (or decreases) of allocated bandwidth/latency in demand-based schemes in a more deterministic manner.

4. Generic DLEP Packet Definition

The Generic DLEP Packet Definition follows the format for packets defined in RFC 5444.

The Generic DLEP Packet Definition contains the following fields:

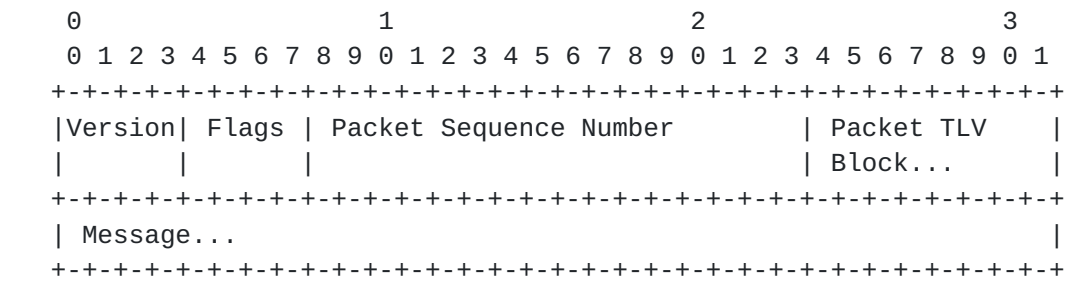

Version - Version of RFC5444 specification on which the packet/messages/TLVs are constructed.

Flags - 4 bit field. Only bit 1 (phastlv) is set/used. All other bits MUST be ignored by DLEP implementations.

Packet Sequence Number - If present, the packet sequence number is parsed and ignored. DLEP does NOT use or generate packet sequence numbers.

Packet TLV block - a TLV block which contains packet level TLV information.

Message - the packet MAY contain zero or more messages.

5. Generic DLEP Message Format

The Generic DLEP Message Format follows the format for MANET messages

defined in RFC 5444. The <msg-seq-num> field, which is OPTIONAL in RFC 5444, MUST exist in all DLEP messages.

The Generic DLEP Message Format contains the following fields:

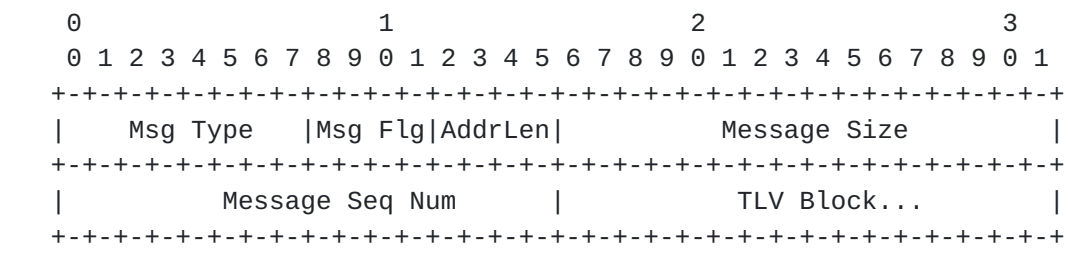

Message Type - an 8-bit field which specifies the type of the message

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - a 4-bit unsigned integer field encoding the length of all addresses included in this message. DLEP implementations do not use this field; contents SHOULD be ignored.

Message Size - a 16-bit unsigned integer field which specifies the number of octets that make up the message including the message header.

Message Sequence Number - a 16-bit unsigned integer field that contains a sequence number, generated by the originator of the message. Sequence numbers range from 1 to 65535. Sequence numbers roll over at 65535 to 1; 0 is invalid.

TLV Block - TLV Block included in the message.

6. Generic DLEP TLV Block Format

The Generic DLEP TLV Block Format follows the format for MANET message TLVs defined in RFC 5444.

The Generic DLEP TLV Block Format contains the following fields:

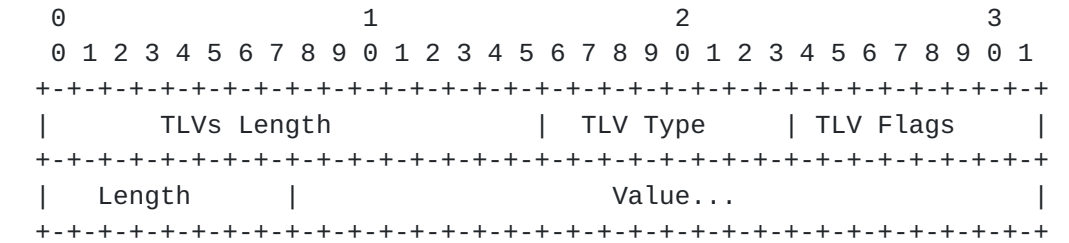

TLVs Length - a 16-bit unsigned integer field that contains the total number of octets in all of the immediately following TLV elements (tlvs-length not included).

TLV Type - an 8-bit unsigned integer field specifying the type of the TLV.

TLV Flags - an 8-bit flags bit field. Only bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - Length of the value field of the TLV

Value - A field of length <Length> which contains data specific to a particular TLV type.

7. DLEP TLVs

TLV Value	TLV Description
TBD	Identification TLV
TBD	DLEP Version TLV
TBD	Peer Type TLV
TBD	MAC Address TLV
TBD	IPv4 Address TLV
TBD	IPv6 Address TLV
TBD	Maximum Data Rate (MDR) TLV
TBD	Current Data Rate (CDR) TLV
TBD	Latency TLV
TBD	Resources TLV
TBD	Relative Link Quality (RLQ) TLV
TBD	Status TLV

7.1 Identification TLV

This TLV MUST be in the Packet Header TLV Block for all DLEP messages. It contains client and router identification information used for all messages contained within the packet.

The Identification TLV contains the following fields:

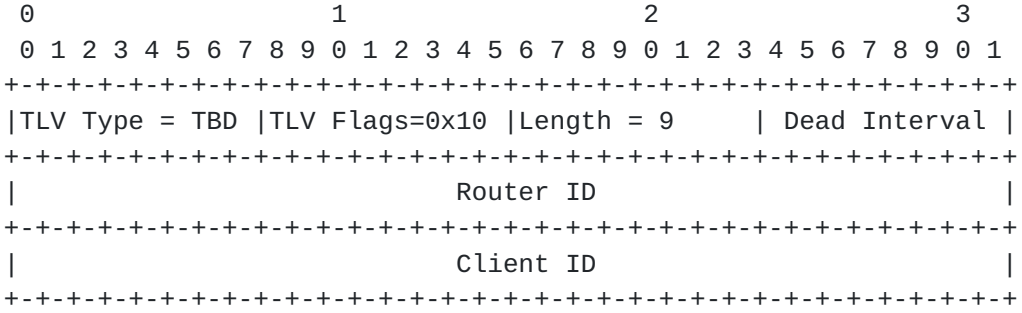

TLV Type - Value TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are unused and MUST be set to '0'.

Length - 9

Dead Interval - An 8-bit, unsigned value containing the maximum number of seconds during which no messages can be received before determining that the session is dead. A value of '0' indicates that the field is ignored. This value is used during the Peer Discovery/Peer Offer exchange. In other packets, the value MUST be ignored.
The Dead timer runs at a peer-to-peer level, that is, it runs between a router and a modem device. If a peer does NOT receive any messages for a complete dead interval, it should initiate DLEP session termination procedures.

Router ID - indicates the router ID of the DLEP session.

Client ID - indicates the client ID of the DLEP session.

When the client initiates discovery (via the Peer Discovery message), it MUST set the Client ID to a 32-bit quantity that will be used to uniquely identify this session from the client-side. The client MUST

set the Router ID to '0'. When responding to the Peer Discovery message, the router MUST echo the Client ID, and MUST supply its own unique 32-bit quantity to identify the session from the router's perspective. After the Peer Discovery/Peer Offer exchange, both the Client ID and the Router ID MUST be set to the values obtained from the Peer DIscovery/Peer Offer sequence.

7.2 DLEP Version TLV

The DLEP Version TLV is OPTIONAL, and is used to indicate the client or router version of the protocol. The client and router MAY use this information to decide if the peer is running at a supported level.

The DLEP Version TLV contains the following fields:

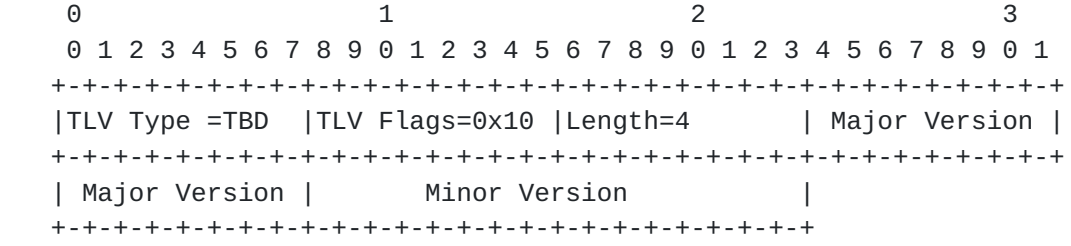

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are

not used and MUST be set to '0'.

Length - Length is 4

Major Version - Major version of the client or router protocol.

Minor Version - Minor version of the client or router protocol.

Support of this draft is indicated by setting the Major Version to '1', and the Minor Version to '0' (e.g. Version 1.0).

7.3 Peer Type TLV

The Peer Type TLV is used by the router and client to give additional information as to its type. It is an OPTIONAL TLV in both the Peer Discovery Message and the Peer Offer message. The peer type is a string and is envisioned to be used for informational purposes (e.g. display command).

The Peer Type TLV contains the following fields:

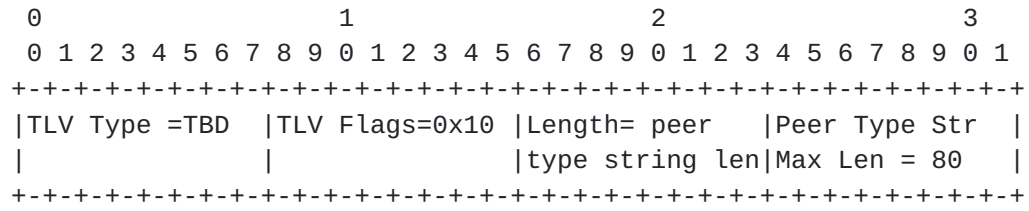

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - length of peer type string (80 bytes maximum)

Peer Type String - Non-Null terminated peer type string, maximum length of 80 bytes. For example, a satellite modem might set this variable to 'Satellite terminal'.

7.4 MAC Address TLV

The MAC address TLV MUST appear in all neighbor-oriented messages (e.g. Neighbor Up, Neighbor Up ACK, Neighbor Down, Neighbor Down ACK, Neighbor Update, Link Characteristics Request, and Link Characteristics ACK). The MAC Address TLV contains the address of the far-end (neighbor) router.

The MAC Address TLV contains the following fields:

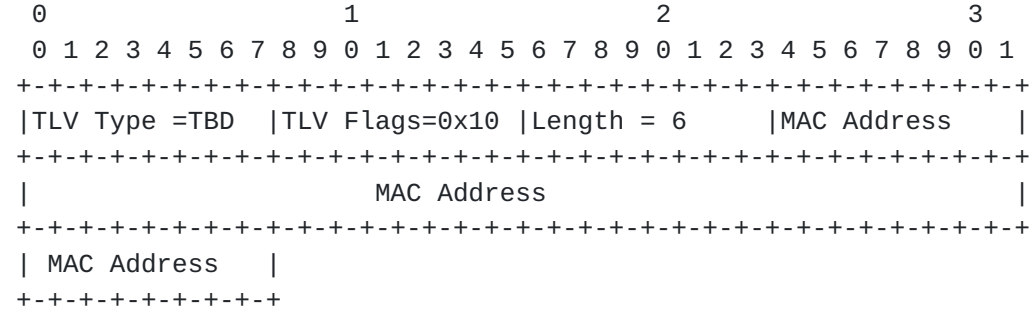

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 6

MAC Address - MAC Address of the far-end router.

7.5 IPv4 Address TLV

The IPv4 Address TLV MAY be used in Neighbor Up, Neighbor Update, and Peer Update Messages, if the client is aware of the Layer 3 address. When included in Neighbor messages, the IPv4 Address TLV contains the IPv4 address of the far-end router (neighbor). In the Peer Update message, it contains the IPv4 address of the local router. In either case, the TLV also contains an indication of whether this is a new or existing address, or is a deletion of a previously known address.

The IPv4 Address TLV contains the following fields:

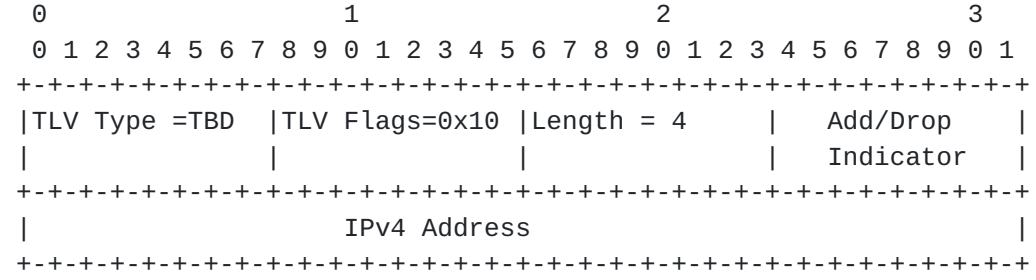

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 5

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

IPv4 Address - IPv4 Address of the far-end router.

7.6 IPv6 Address TLV

The IPv6 Address TLV MAY be used in Neighbor Up, Neighbor Update, and Peer Update Messages, if the client is aware of the Layer 3 address. When included in Neighbor messages, the IPv6 Address TLV contains the IPv6 address of the far-end router (neighbor). In the Peer Update, it contains the IPv6 address of the local router. In either case, the TLV also contains an indication of whether this is a new or existing address, or is a deletion of a previously known address.

The IPv6 Address TLV contains the following fields:

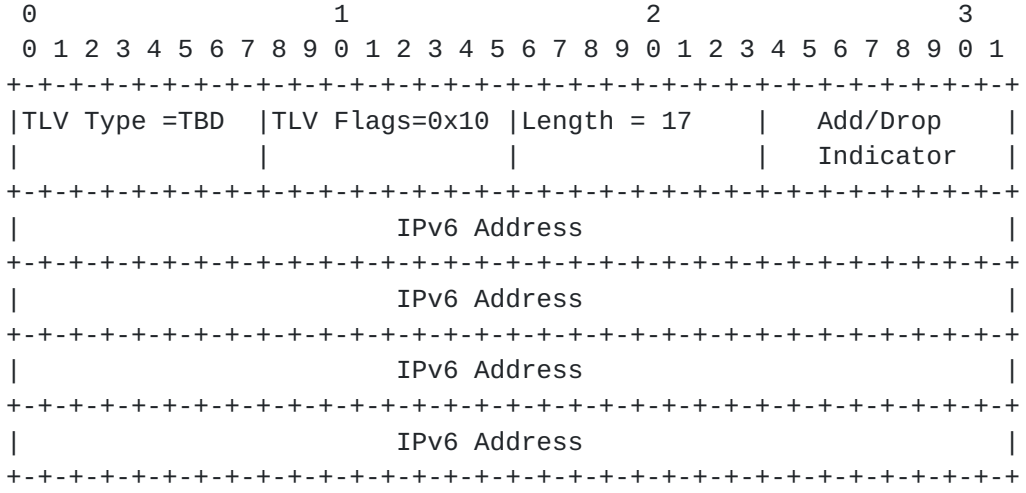

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 17

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

IPv6 Address - IPv6 Address of the far-end router.

7.7 Maximum Data Rate TLV

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

The Maximum Data Rate TLV contains the following fields:

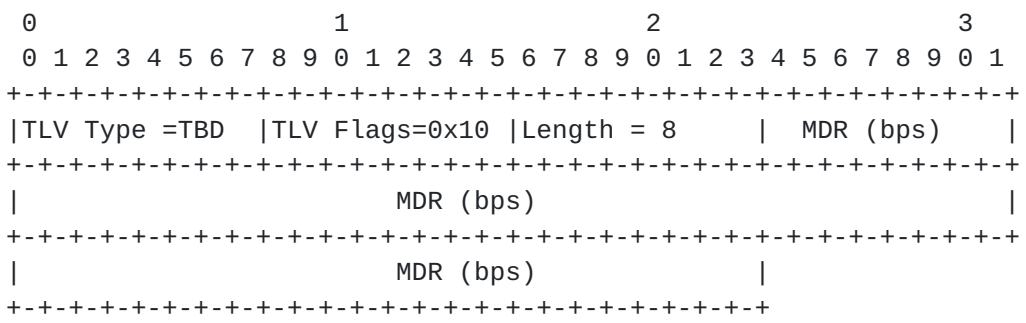

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 8

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

7.8 Current Data Rate TLV

The Current Data Rate (CDR) TLV is used in Neighbor Up, Neighbor Update, Link Characteristics Request, and Link Characteristics ACK messages to indicate the rate at which the link is currently operating, or in the case of the Link Characteristics Request, the desired data rate for the link.

The Current Data Rate TLV contains the following fields:

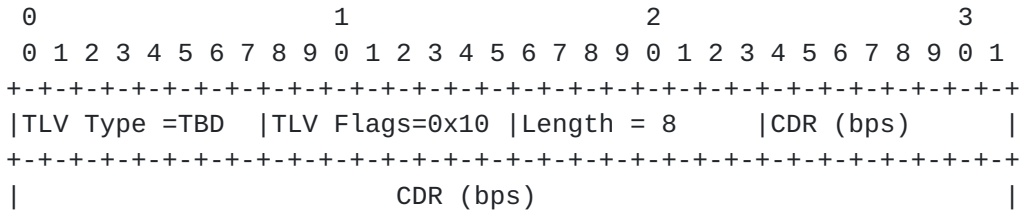

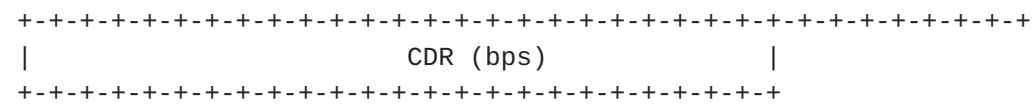
CDR (bps)

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 8

Current Data Rate - A 64-bit unsigned number, representing the current data rate, in bits per second (bps), on the link. When reporting metrics (e.g, in Neighbor Up, Neighbor Down, or Link Characteristics ACK), if there is no distinction between current and maximum data rates, current data rate SHOULD be set equal to the maximum data rate.

7.9 Latency TLV

The Latency TLV is used in Neighbor Up, Neighbor Update, Link Characteristics Request, and Link Characteristics ACK messages to indicate the amount of latency on the link, or in the case of the Link Characteristics Request, to indicate the maximum latency required (e.g. a should-not-exeed value) on the link.

The Latency TLV contains the following fields:

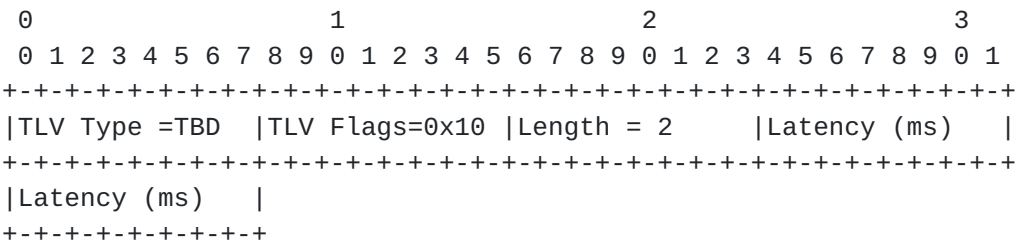

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 2

Latency - the transmission delay that a packet encounters as it is transmitted over the link. In Neighbor Up, Neighbor Update, and Link Characteristics ACK, this value is reported in absolute delay, in milliseconds. The calculation of latency is modem-device dependent. For example, the latency may be a running average calculated from the internal queuing. If

the modem device cannot calculate latency,
it SHOULD be reported as 0.

In the Link Characteristics Request Message, this value represents the maximum delay, in milliseconds, expected on the link.

7.10 Resources TLV

The Resources TLV is used in Neighbor Up, Neighbor Update, and Link Characteristics ACK messages to indicate a percentage (0-100) amount of resources (e.g. battery power) remaining on the modem device.

The Resources TLV contains the following fields:

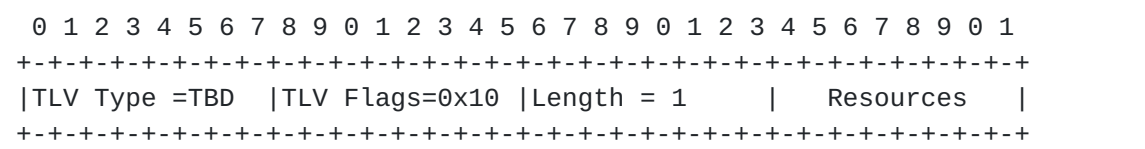

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 1

Resources - a percentage, 0-100, representing the amount of remaining resources, such as battery power. If resources cannot be calculated, a value of 100 SHOULD be reported.

7.11 Relative Link Quality TLV

The Relative Link Quality (RLQ) TLV is used in Neighbor Up, Neighbor Update, and Link Characteristics ACK messages to indicate the quality of the link as calculated by the modem device.

The Relative Link Quality TLV contains the following fields:

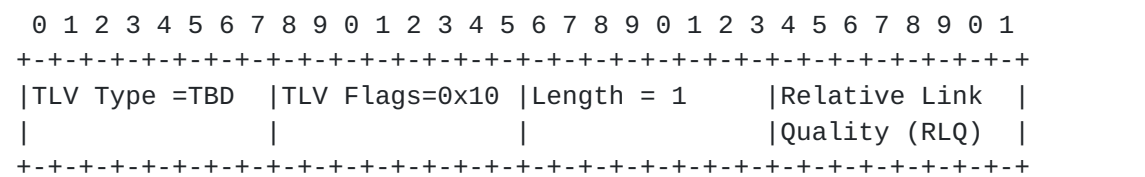

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 1

Relative Link Quality - a non-dimensional number, 0-100, representing the relative link quality. A value of 100 represents a link of the highest quality. If the RLQ cannot be calculated, a value of 100 should be reported.

7.12 Status TLV

The Status TLV is sent from either the client or router to indicate the success or failure of a given request

The Status TLV contains the following fields:

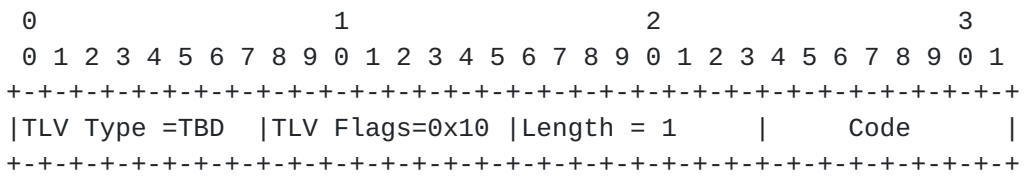

TLV Type - TBD

TLV Flags - 0x10, Bit 3 (thasvalue) is set, all other bits are not used and MUST be set to '0'.

Length - 1

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

8. DLEP Messages

The DLEP Packet, being based on [RFC5444], contains the following fields:

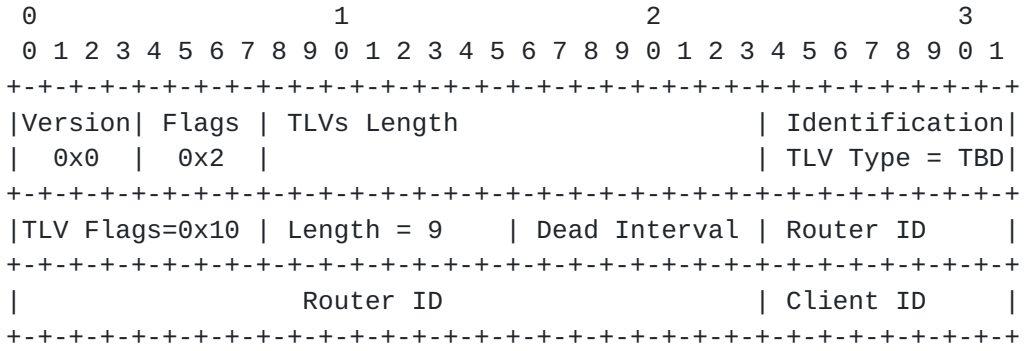

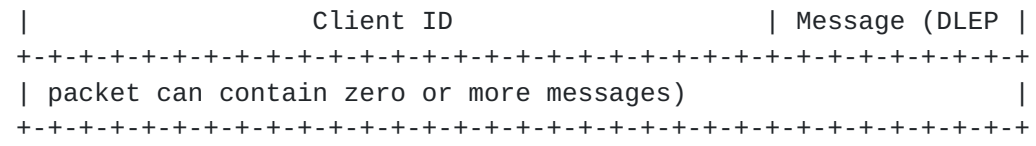
| Client ID | Message (DLEP |
| packet can contain zero or more messages) |

Version - Version of RFC5444 specification on which the packet/ messages/TLVs are constructed.

Flags - 0x2 Only bit 1 (phastlv) is set/used, all other bits are not used and MUST be set to '0'.

Packet Header TLV Block which contains:
Identification TLV

Message - the packet may contain zero or more messages.

8.1 Message TLVs

TLV Value	TLV Description
TBD	Attached Peer Discovery
TBD	Detached Peer Discovery
TBD	Peer Offer
TBD	Peer Update
TBD	Peer Update ACK
TBD	Peer Termination
TBD	Peer Termination ACK
TBD	Neighbor Up
TBD	Neighbor Up ACK
TBD	Neighbor Down
TBD	Neighbor Down ACK
TBD	Neighbor Update
TBD	Heartbeat
TBD	Link Characteristics Request
TBD	Link Characteristics ACK

9. Peer Discovery Messages

There are two different types of Peer Discovery Messages, Attached and Detached. Attached Peer Discovery Messages are sent by the client when it is directly attached to the router (e.g. the client exists as a card in the chassis, or it is connected via Ethernet with no intervening devices). The Detached Peer Discovery message, on the other hand, is sent by a "remote" client -- for example, a client at a satellite hub system might use a Detached Discovery Message in order to act as a proxy for remote ground terminals. To explain in another way, a detached client uses the variable link itself (the radio or satellite link) to establish a DLEP session with a remote router.

9.1 Attached Peer Discovery Message

The Attached Peer Discovery Message is sent by an attached client to a router to begin a new DLEP association. The Peer Offer message

is required to complete the discovery process. The client MAY implement its own retry heuristics in the event it (the client) determines the Attached Peer Discovery Message has timed out.

The Attached Peer Discovery Message contains the following fields:

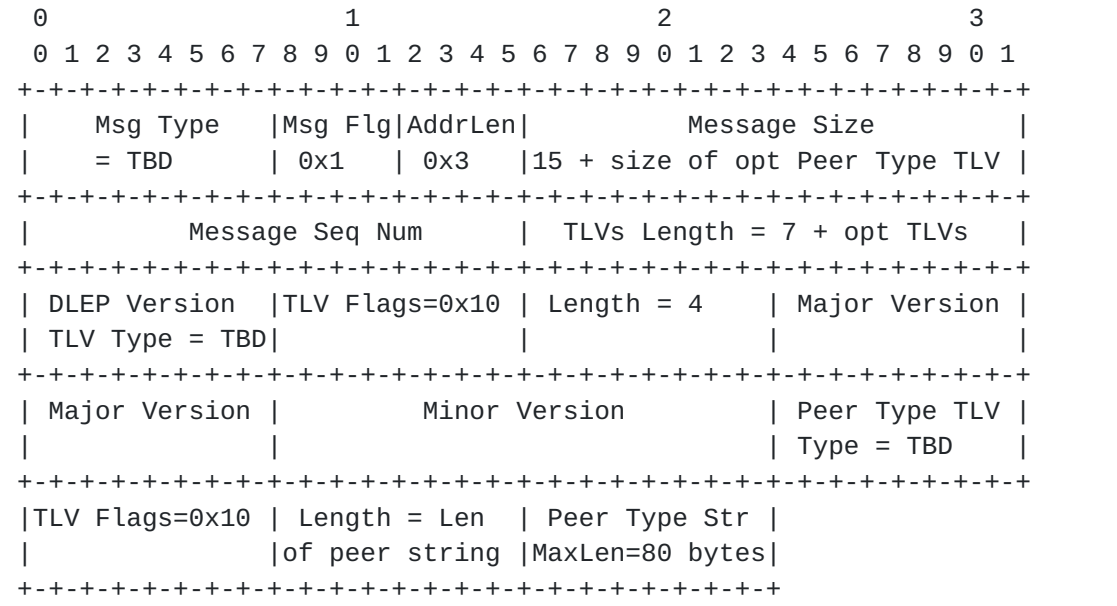

Attached Peer Discovery Message - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). No other bits are used and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 15 + size of optional Peer Type TLV

Message Sequence Number - a 16-bit unsigned integer field containing a sequence number generated by the message originator.

TLV Block - TLVs Length: 7 + size of OPTIONAL Peer Type TLV.
DLEP Version TLV
Peer Type TLV (OPTIONAL)

9.2 Detached Peer Discovery Message

The Detached Peer Discovery Message is sent by a detached client proxy to a router to begin a new DLEP session. The Peer Offer message is required to complete the discovery process. The client MAY implement its own retry heuristics in the event it (the client) determines the Detached Peer Discovery Message has timed out.

The Detached Peer Discovery Message contains the following fields:

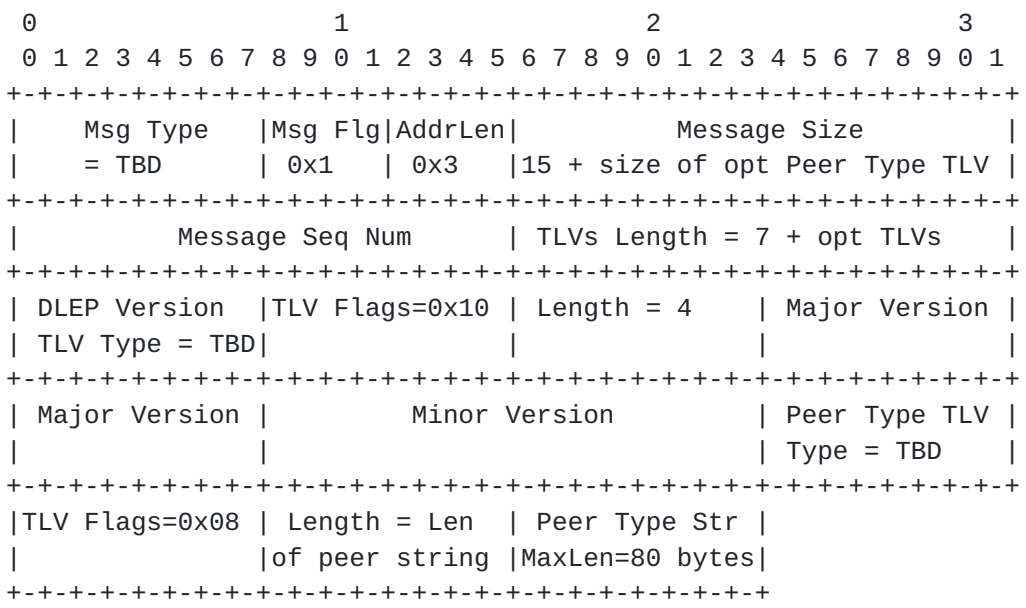

Detached Peer Discovery Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are not used and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 15 + size of optional Peer Type TLV

Message Sequence Number - A 16-bit unsigned integer field containing a sequence number, generated by the message originator.

TLV Block - TLVs Length: 7 + size of OPTIONAL Peer Type TLV.
DLEP Version TLV
Peer Type TLV (optional)

10. Peer Offer Message

The Peer Offer Message is sent by a router to a client or client proxy in response to a Peer Discovery Message. The Peer Offer Message is the response to either of the Peer Discovery messages (either Attached or Detached), and completes the DLEP session establishment.

The Peer Offer Message contains the following fields:

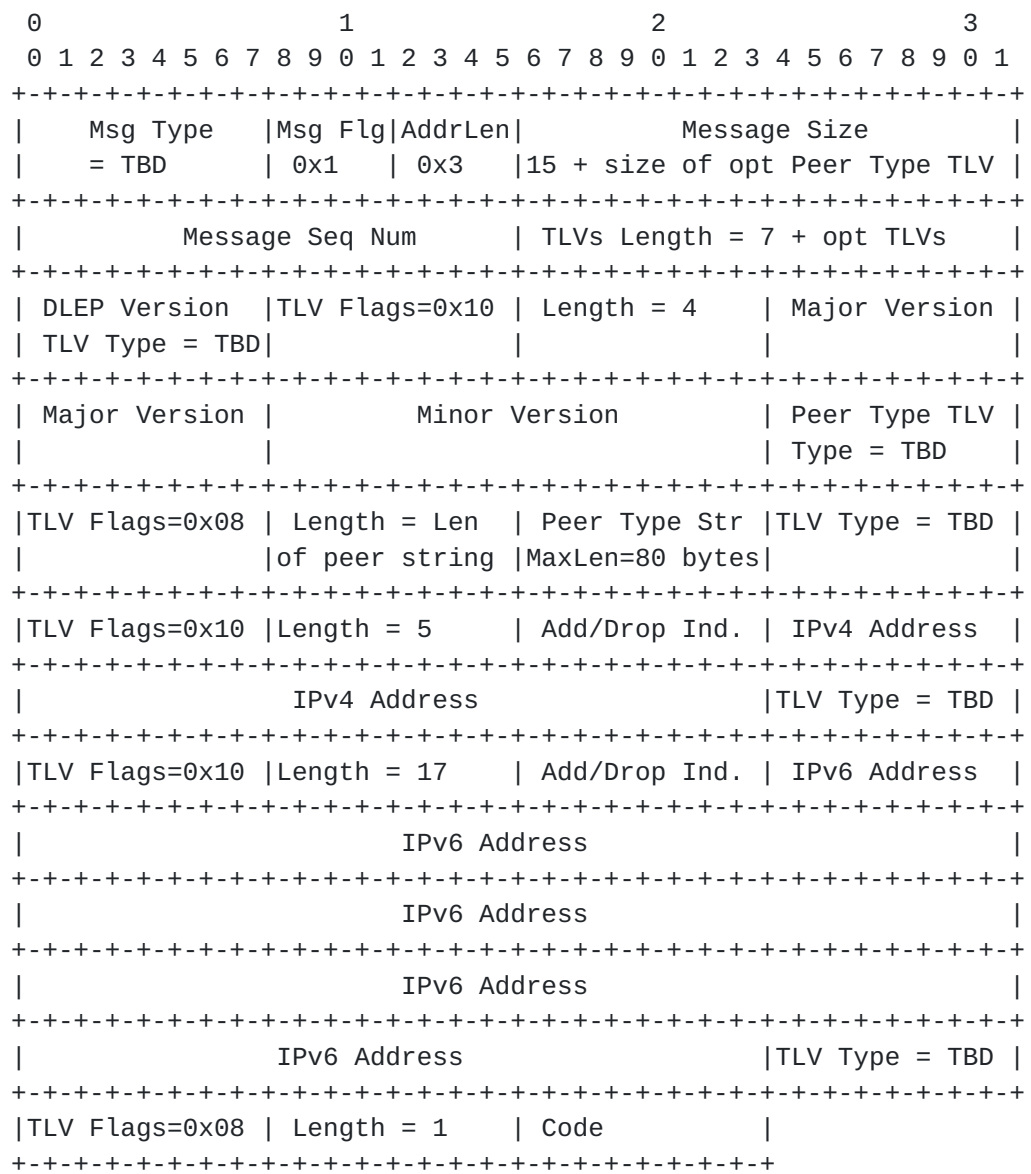

Peer Offer Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and

MUST be set to '0'.

Message Address Length - 0x3

Message Size - 15 + size of optional Peer Type TLV

Message Sequence Number	- A 16-bit unsigned integer field containing a sequence number, generated by the message originator.
TLV Block	- TLV Length: 7 + size of optional Peer Type TLV. DLEP Version TLV Peer Type TLV (OPTIONAL) IPv4 Address TLV (OPTIONAL) IPv6 Address TLV (OPTIONAL) Status TLV (OPTIONAL)

11. Peer Update Message

The Peer Update message is sent by the router to indicate local Layer 3 address changes. For example, addition of an IPv4 address to the router would prompt a Peer Update message to its attached DLEP clients. If the modem device is capable of understanding and forwarding this information, the address update would prompt any remote DLEP clients (DLEP clients that are on the far-end of the variable link) to issue a "Neighbor Update" message to their local routers, with the address change information. Clients that do not track Layer 3 addresses MUST silently ignore the Peer Update Message. Clients that track Layer 3 addresses MUST acknowledge the Peer Update with a Peer Update ACK message. Routers MAY employ heuristics to retransmit Peer Update messages. Sending of Peer Update Messages SHOULD cease when a router implementation determines that a partner modem device does NOT support Layer 3 address tracking.

The Peer Update Message contains the following fields:

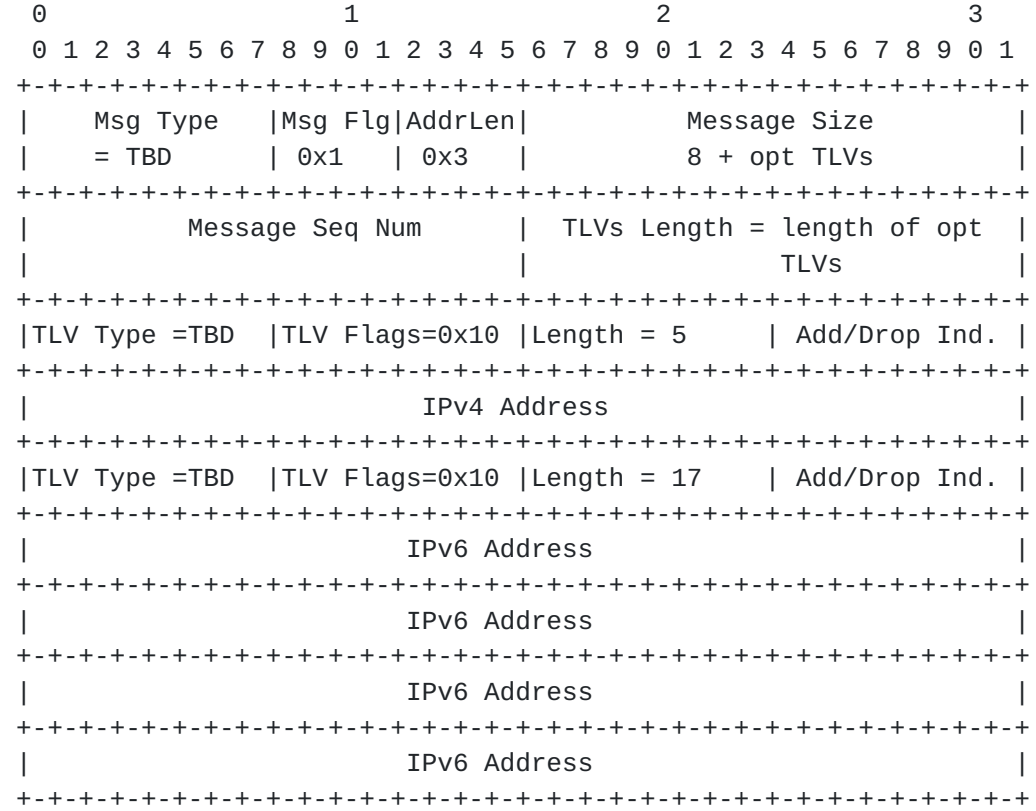

Peer Update Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 8 + optional TLVs

Message Sequence Number - A 16-bit unsigned integer field containing a sequence number generated by the message originator.

TLV Block - TLV Length: length of optional TLVs.
IPv4 Address TLV (OPTIONAL)
IPv6 Address TLV (OPTIONAL)

12. Peer Update ACK Message

The client sends the Peer Update ACK Message to indicate whether a Peer Update Message was successfully processed.

The Peer Update ACK message contains the following fields:

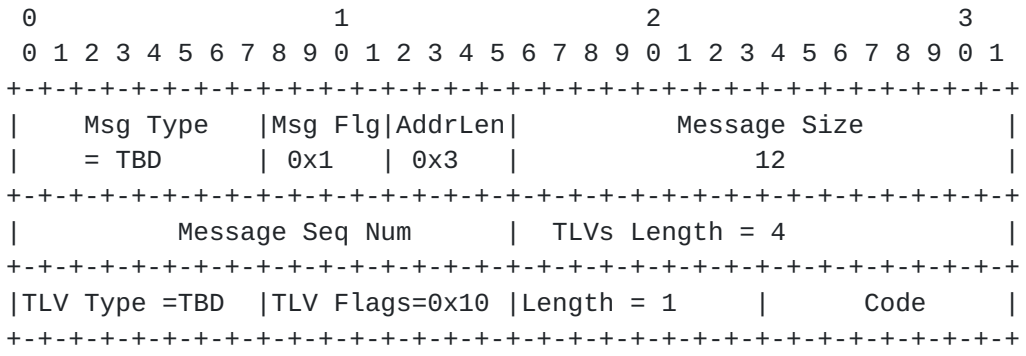

Peer Update ACK

Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 12

Message Sequence Number - A 16-bit unsigned integer field containing the sequence number from the Neighbor Up Message that is being acknowledged.

TLV Block - TLV Length: 4
Status TLV

13. Peer Termination Message

The Peer Termination Message is sent by either the client or the router when a session needs to be terminated. Transmission of a Peer Termination ACK message is required to confirm the termination process. The sender of the Peer Termination message is free to define its heuristics in event of a timeout. The receiver of a Peer Termination Message MUST terminate all neighbor relationships and release associated resources. No Neighbor Down messages are sent.

The Peer Termination Message contains the following fields:

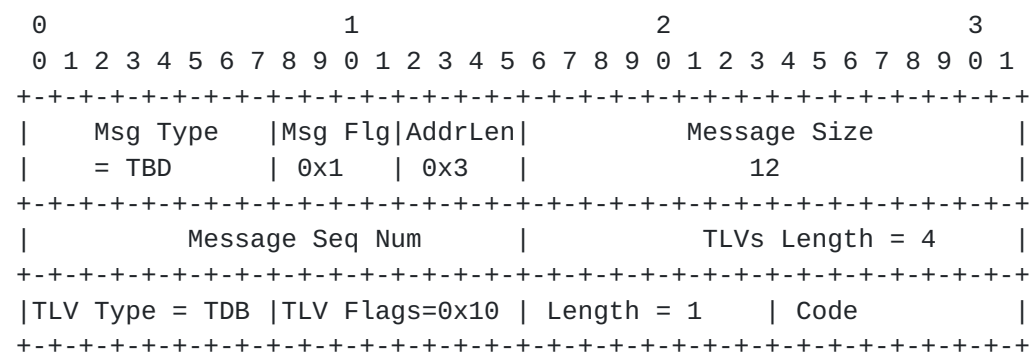

Peer Termination Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 12

Message Sequence Number - A 16-bit unsigned integer field containing a sequence number generated by the message originator.

TLV Block - TLV Length = 4.
Status TLV

14. Peer Termination ACK Message

The Peer Termination Message ACK is sent by either the client or the router when a session needs to be terminated.

The Peer Termination ACK Message contains the following fields:

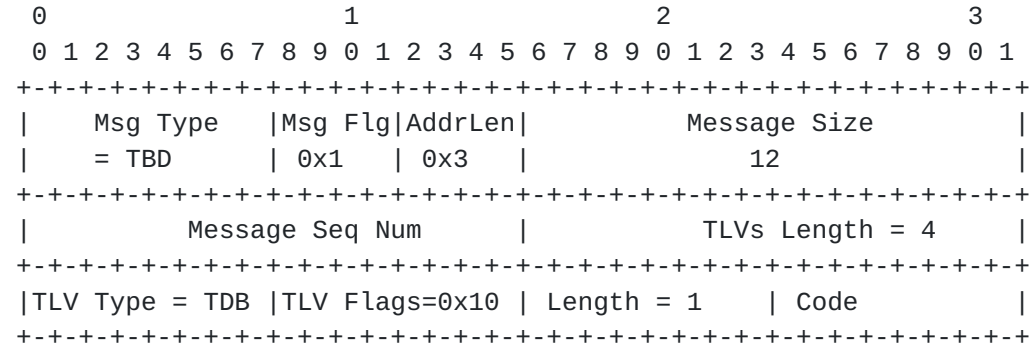

Peer Termination ACK
Message Type - TBD

Message Flags	- Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.
Message Address Length	- 0x3
Message Size	- 12
Message Sequence Number	- A 16-bit unsigned integer field containing the sequence number in the corresponding Peer Termination Message being acknowledged.
TLV Block	- TLV Length = 4. Status TLV

15. Neighbor Up Message

The client sends the Neighbor Up message to report that a new potential routing neighbor has been detected. A Neighbor Up ACK Message is required to confirm a received Neighbor Up. The sender of the Neighbor Up Message is free to define its retry heuristics in event of a timeout.

The Neighbor Up Message contains the following fields:

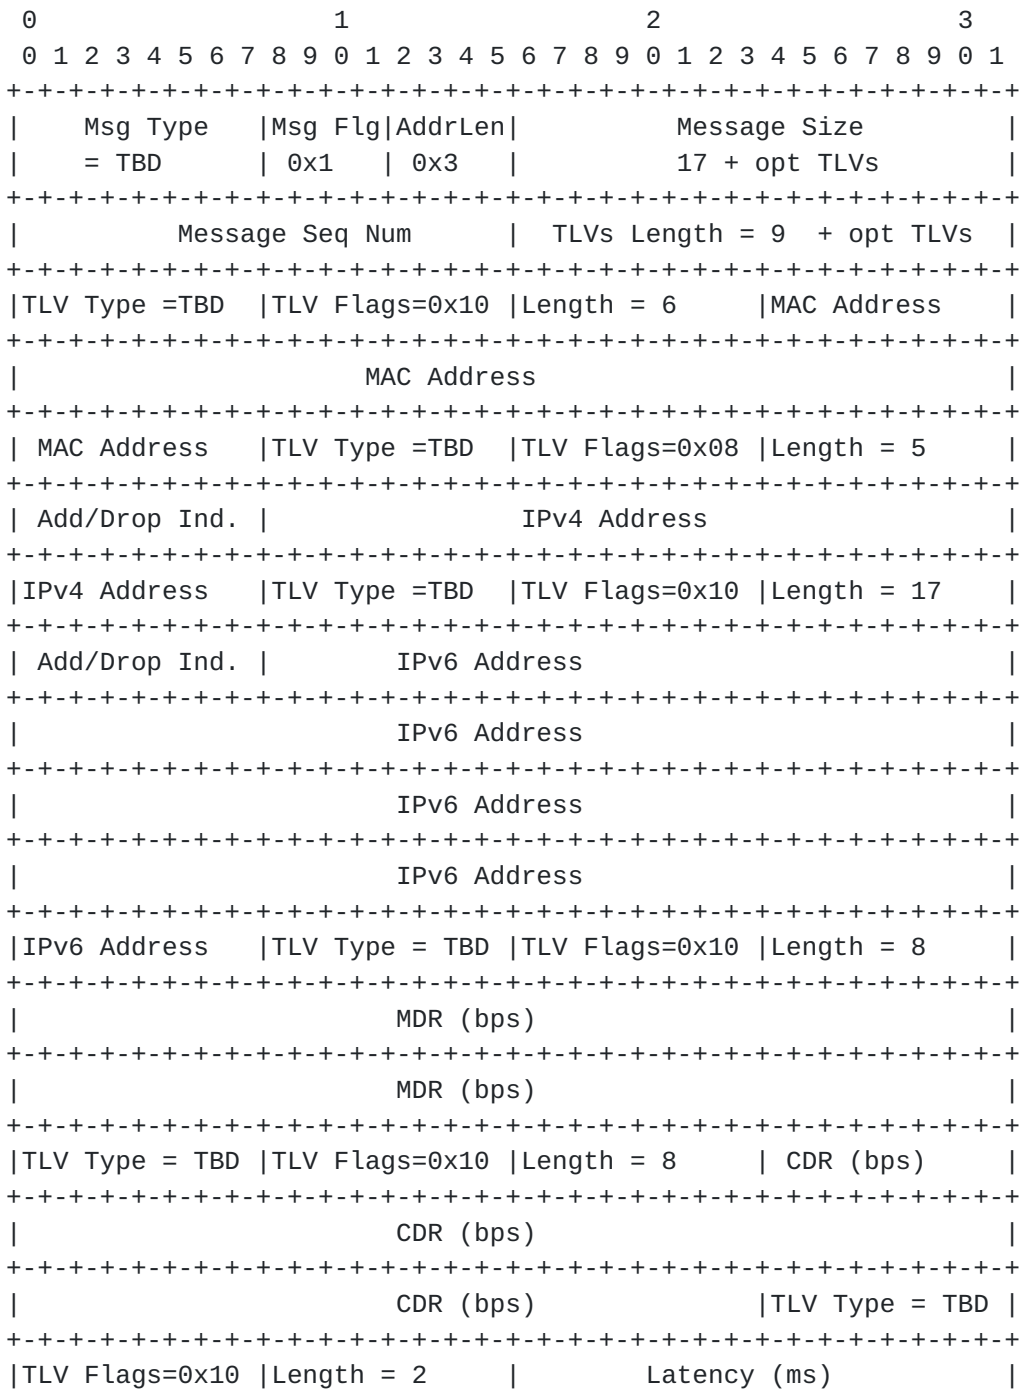

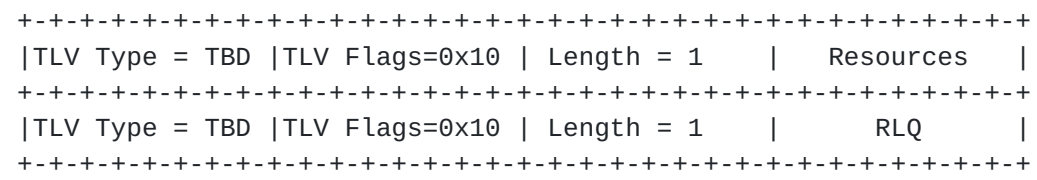
|TLV Type = TBD |TLV Flags=0x10 | Length = 1 | Resources |
|TLV Type = TBD |TLV Flags=0x10 | Length = 1 | RLQ |

Neighbor Up Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 17 + optional TLVs

Message Sequence Number - A 16-bit unsigned integer field containing a sequence number generated by the message originator.

TLV Block - TLV Length: 9 + optional TLVs.
MAC Address TLV (MANDATORY)
IPv4 Address TLV (OPTIONAL)
IPv6 Address TLV (OPTIONAL)
Maximum Data Rate TLV (OPTIONAL)
Current Data Rate TLV (OPTIONAL)
Latency TLV (OPTIONAL)
Resources TLV (OPTIONAL)
Relative Link Factor TLV (OPTIONAL)

16. Neighbor Up ACK Message

The router sends the Neighbor Up ACK Message to indicate whether a Neighbor Up Message was successfully processed.

The Neighbor Up ACK message contains the following fields:

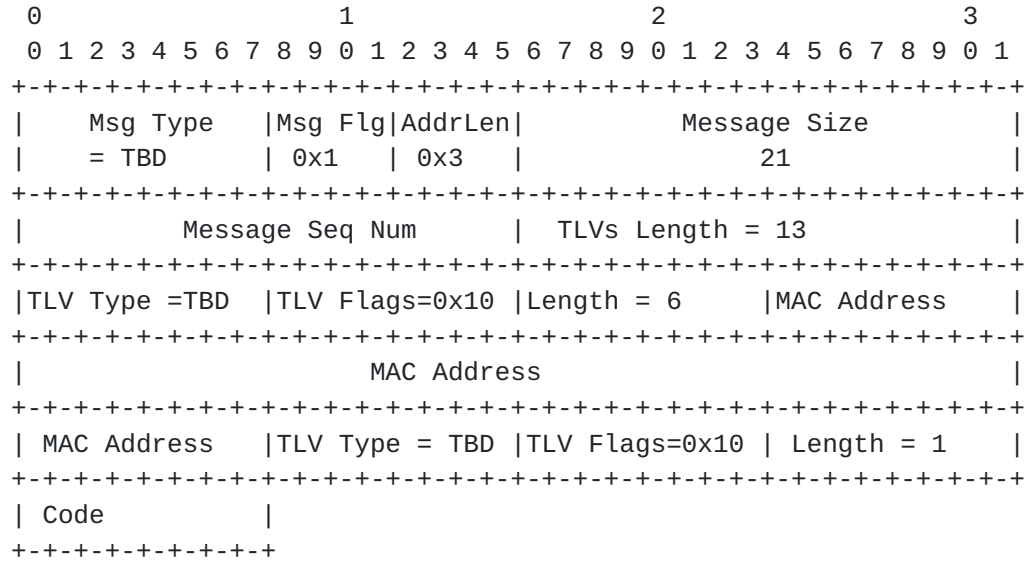

Neighbor Up ACK

Message Type - TBD

Message Flags	- Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.
Message Address Length	- 0x3
Message Size	- 21
Message Sequence Number	- A 16-bit unsigned integer field containing the sequence number from the Neighbor Down Message that is being acknowledged.
TLV Block	- TLV Length: 13 MAC Address TLV (MANDATORY) Status TLV (MANDATORY)

17. Neighbor Down Message

The client sends the Neighbor Down message to report when a neighbor is no longer reachable from the client. The Neighbor Down message MUST contain a MAC Address TLV. Any other TLVs present MAY be ignored. A Neighbor Down ACK Message is required to confirm the process. The sender of the Neighbor Down message is free to define its retry heuristics in event of a timeout.

The Neighbor Down Message contains the following fields:

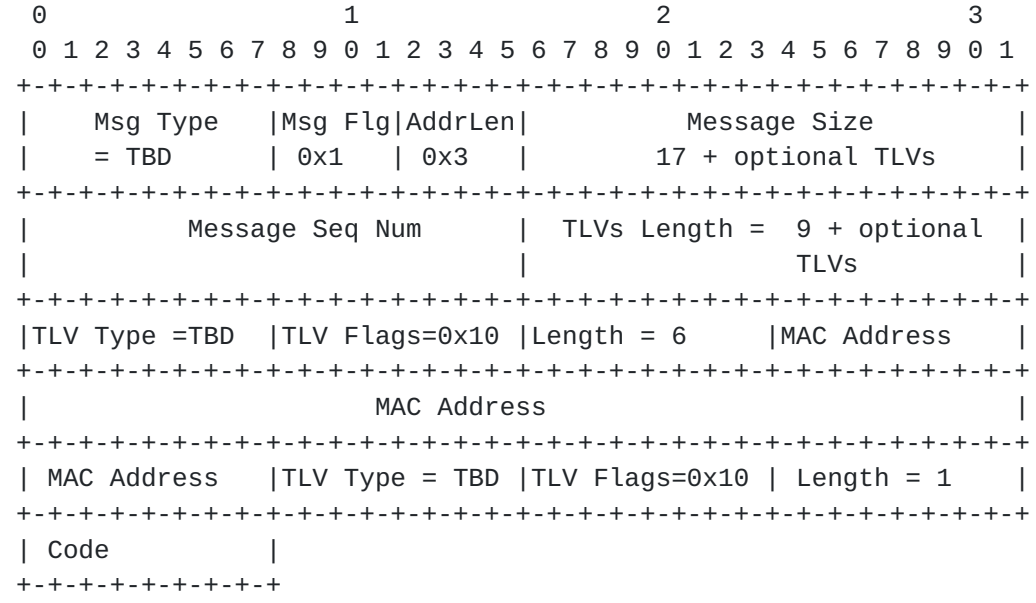

Neighbor Down Message Type - TBD

Message Flags	- Set to 0x1 (bit 3, mhasseqnum bit

is set). All other bits are unused and
MUST be set to '0'.

Message Address Length - 0x3

Message Size - 17 + optional TLVs

Message Sequence Number - A 16-bit unsigned integer field containing a sequence number generated by the message originator.

TLV Block - TLV Length: 9 + optional TLVs
MAC Address TLV (MANDATORY)
Status TLV (OPTIONAL)

18. Neighbor Down ACK Message

The router sends the Neighbor Down ACK Message to indicate whether a Neighbor Down Message was successfully processed.

The Neighbor Down ACK message contains the following fields:

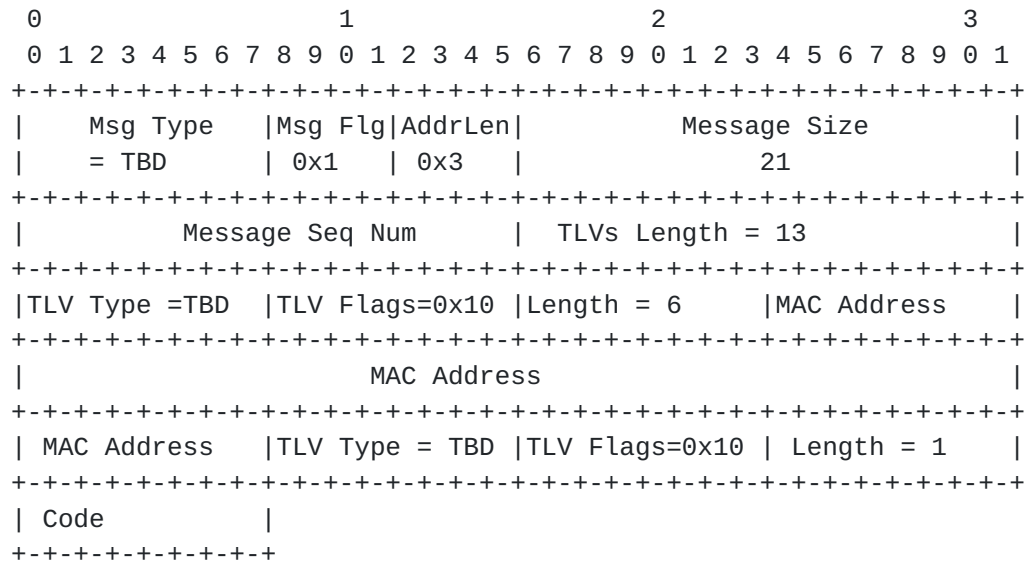

Neighbor Down ACK

Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 21

Message Sequence Number - A 16-bit unsigned integer field containing the sequence number from the Neighbor Down Message that is being acknowledged.

```
   TLV Block                - TLV Length:  13
                              MAC Address TLV (MANDATORY)
                              Status TLV (MANDATORY)
```

19. Neighbor Update Message

The client sends the Neighbor Update message when a change in link metric parameters is detected for a routing neighbor.

The Neighbor Update Message contains the following fields:

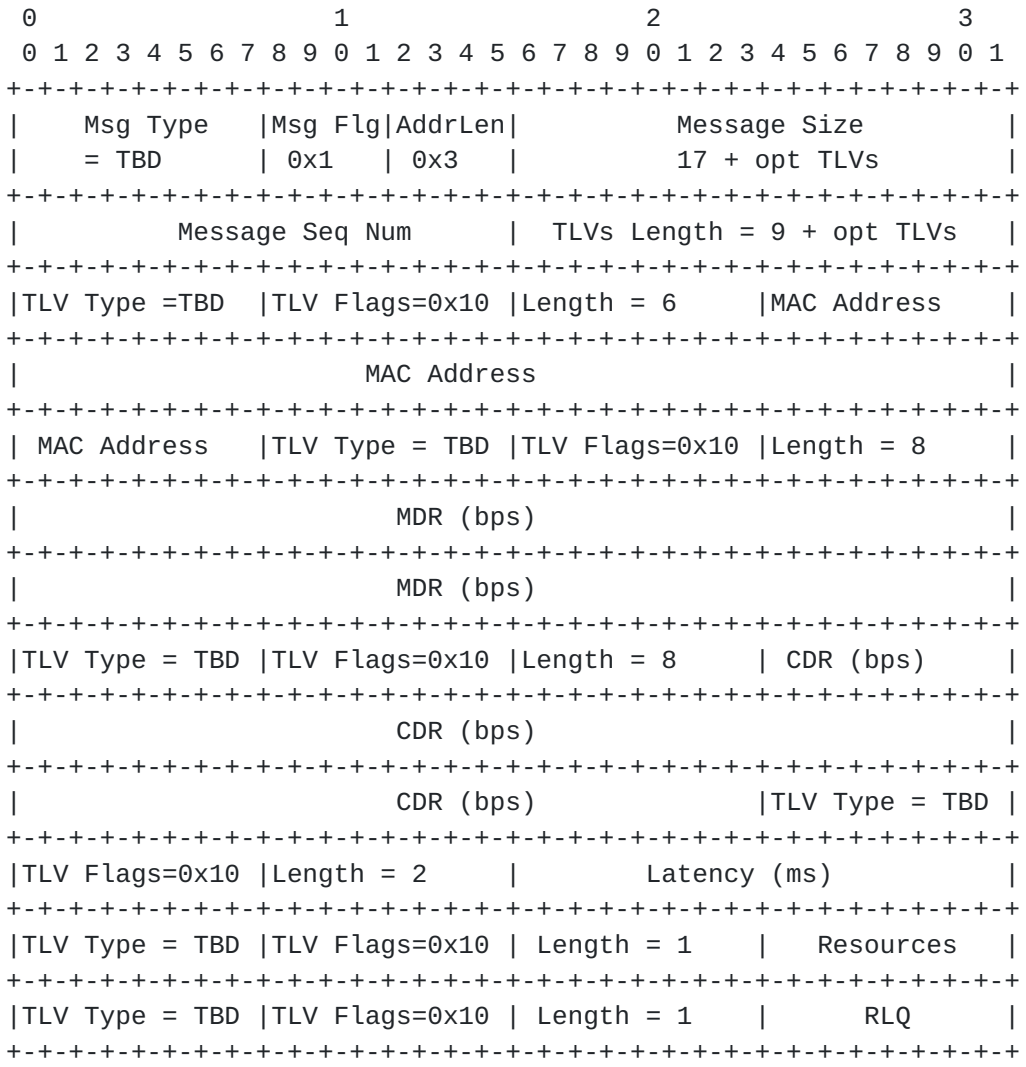

Neighbor Update Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 17 + optional TLVs

Message Sequence Number - A 16-bit unsigned integer field

containing a sequence number,
generated by the message originator.

TLV Block - TLVs Length - 9 + optional TLVs.
MAC Address TLV (MANDATORY)
Maximum Data Rate TLV (OPTIONAL)
Current Data Rate TLV (OPTIONAL)
Latency TLV (OPTIONAL)
Resources TLV (OPTIONAL)
Relative Link Quality TLV (OPTIONAL)

20. Neighbor Address Update Message

The client sends the Neighbor Address Update message when a change in Layer 3 addressing is detected for a routing neighbor.

The Neighbor Address Update Message contains the following fields:

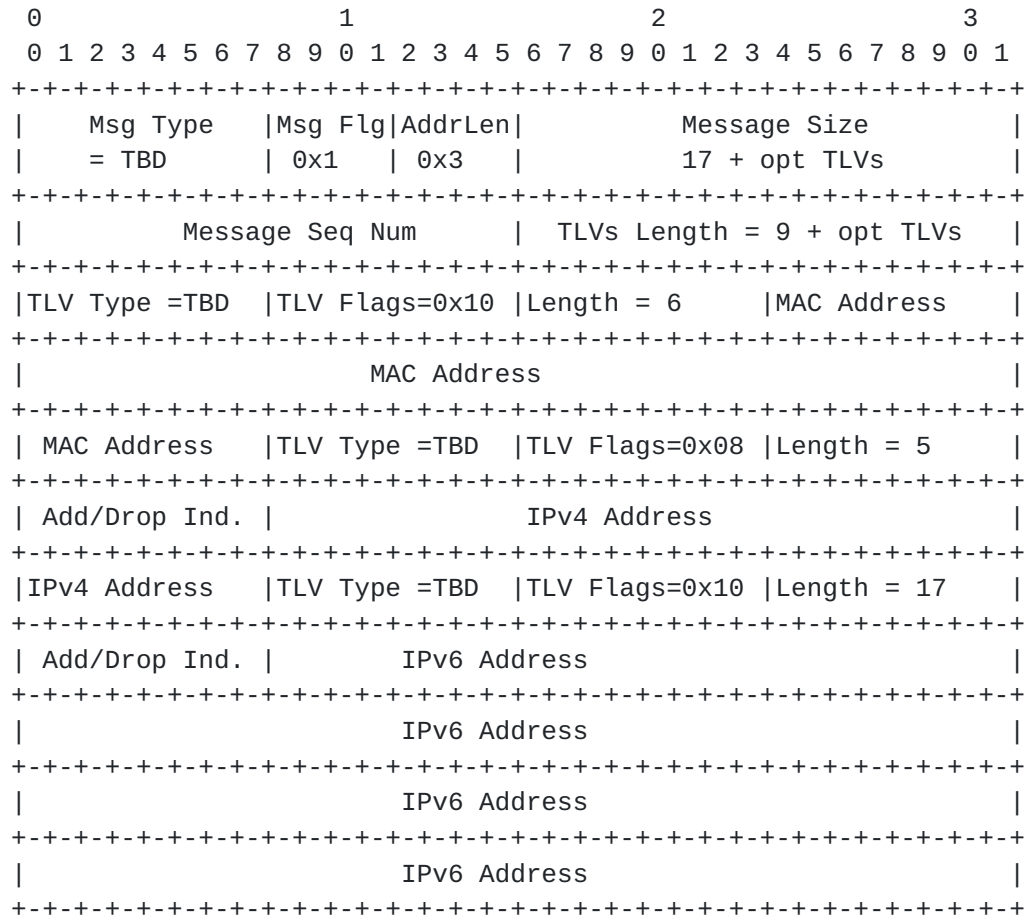

Neighbor Address Update
Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are

unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 17 + optional TLVs

Message Sequence Number - A 16-bit unsigned integer field containing a sequence number, generated by the message originator.

TLV Block - TLVs Length - 9 + optional TLVs.
MAC Address TLV (MANDATORY)
IPv4 Address TLV (OPTIONAL)
IPv6 Address TLV (OPTIONAL)

21. Neighbor Address Update ACK Message

The router sends the Neighbor Address Update ACK Message to indicate whether a Neighbor Address Update Message was successfully processed.

The Neighbor Address Update ACK message contains the following fields:

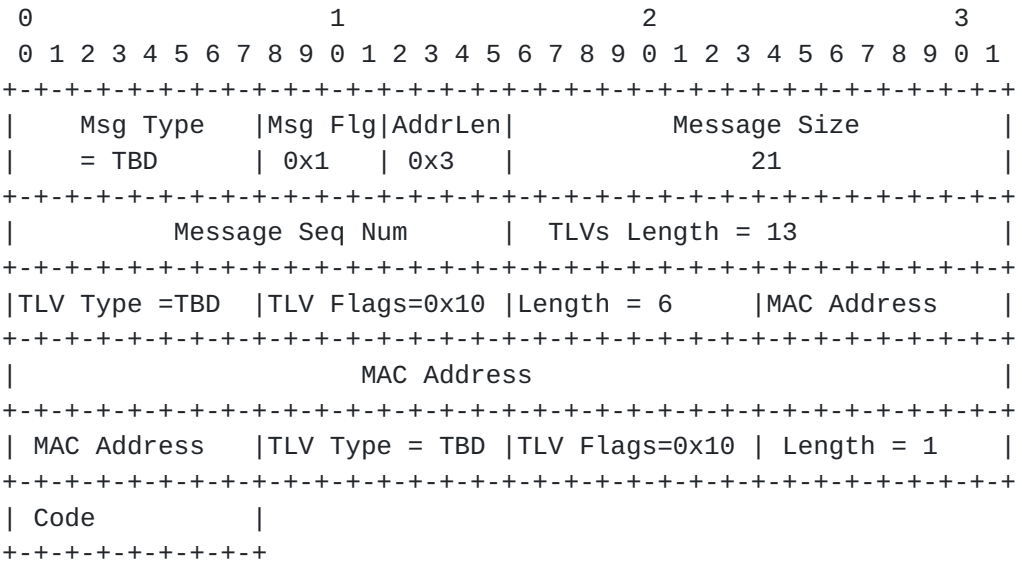

Neighbor Address Update ACK Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 21

Message Sequence Number - A 16-bit unsigned integer field containing the sequence number from the Neighbor Down Message that is being acknowledged.

TLV Block - TLV Length: 13
MAC Address TLV (MANDATORY)
Status TLV (MANDATORY)

22. Heartbeat Message

A Heartbeat Message is sent by a peer every N seconds where no messages have been received on the DLEP session. The message is used by peers to detect when a DLEP session partner is no longer communicating. When the (N * 2) timeout expires, the peer should initiate DLEP session termination procedures.

The Heartbeat Message contains the following fields:

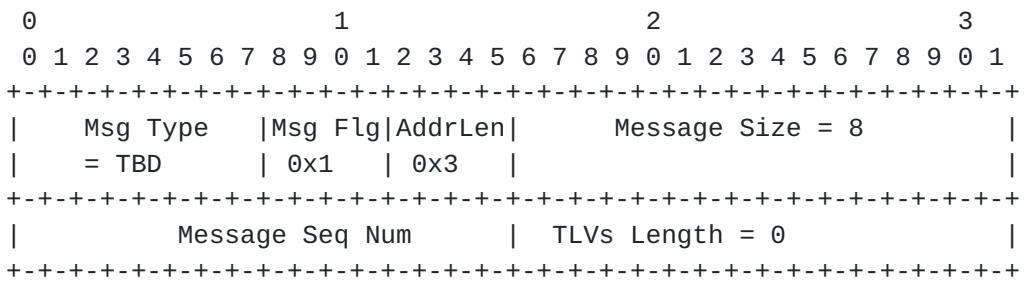

Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and SHOULD be set to '0'.

Message Address Length - 0x3

Message Size - 8

Message Sequence Number - A 16-bit unsigned integer field containing a sequence number generated by the message originator.

TLV Block - TLV Length = 0

23. Link Characteristics Request Message

The Link Characteristics Request Message is sent by the router to the modem device when the router detects that a different set of transmission characteristics is necessary (or desired) for the type of traffic that is flowing on the link. The request contains either a Current Data Rate (CDR) TLV to request a different amount of bandwidth than what is currently allocated, a Latency TLV to request that traffic delay on the link not exceed the specified value, or both. A Link Characteristics ACK Message is

required to complete the request. Implementations are free to define their retry heuristics in event of a timeout. Issuing a Link Characteristics Request with ONLY the MAC Address TLV is a mechanism a peer MAY use to request metrics (via the Link

Characteristics ACK) from its partner.

The Link Characteristics Request Message contains the following fields:

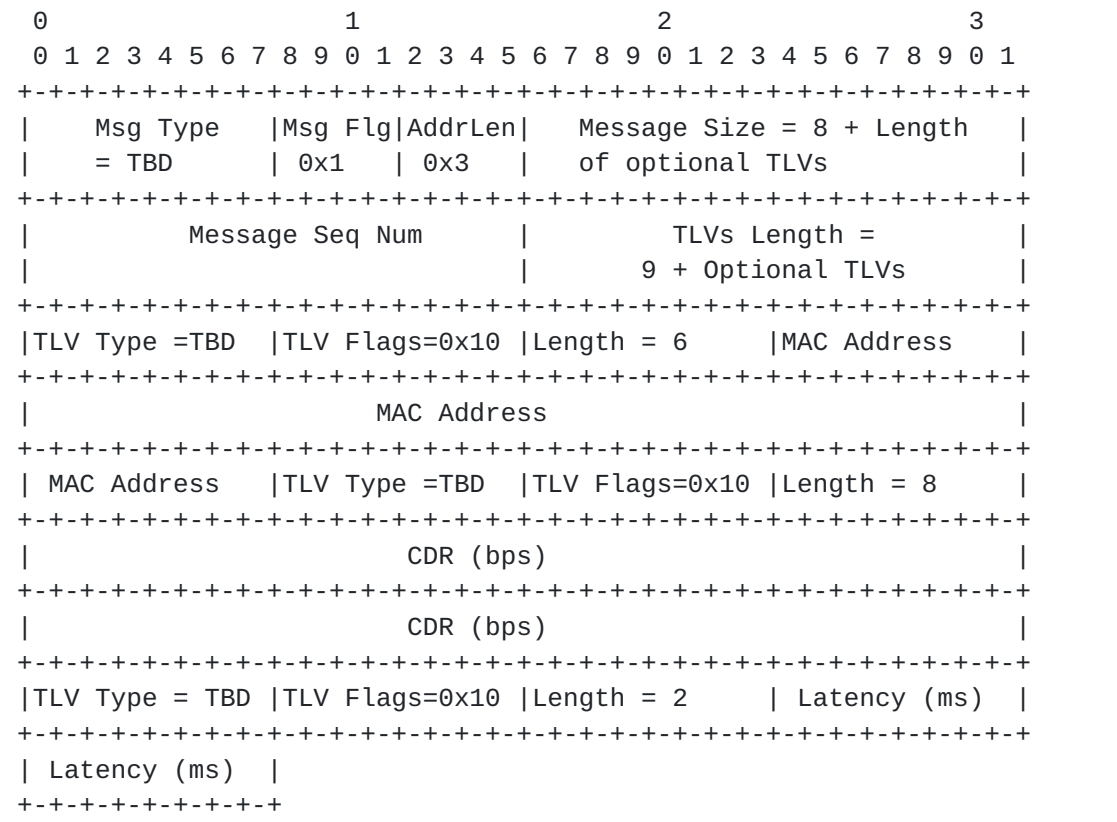

Message Type - TBD

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 8 + length of optional (Current Data Rate and/or Latency) TLVs

Message Sequence Number - A 16-bit unsigned integer field containing a sequence number generated by the message originator.

TLV Block - TLVs Length

MAC Address TLV (MANDATORY)

Current Data Rate TLV - if present, this value represents the requested data rate

in bits per second (bps). (OPTIONAL)

Latency TLV - if present, this value represents the maximum latency, in milliseconds, desired on the link. (OPTIONAL)

24. Link Characteristics ACK Message

The Link Characteristics ACK Message is sent by the client to the router letting the router know the success (or failure) of the requested change in link characteristics. The Link Characteristics ACK message SHOULD contain a complete set of metric TLVs. It MUST contain the same TLV types as the request. The values in the metric TLVs in the Link Characteristics ACK message MUST reflect the link characteristics after the request has been processed.

The Link Characteristics ACK Message contains the following fields:

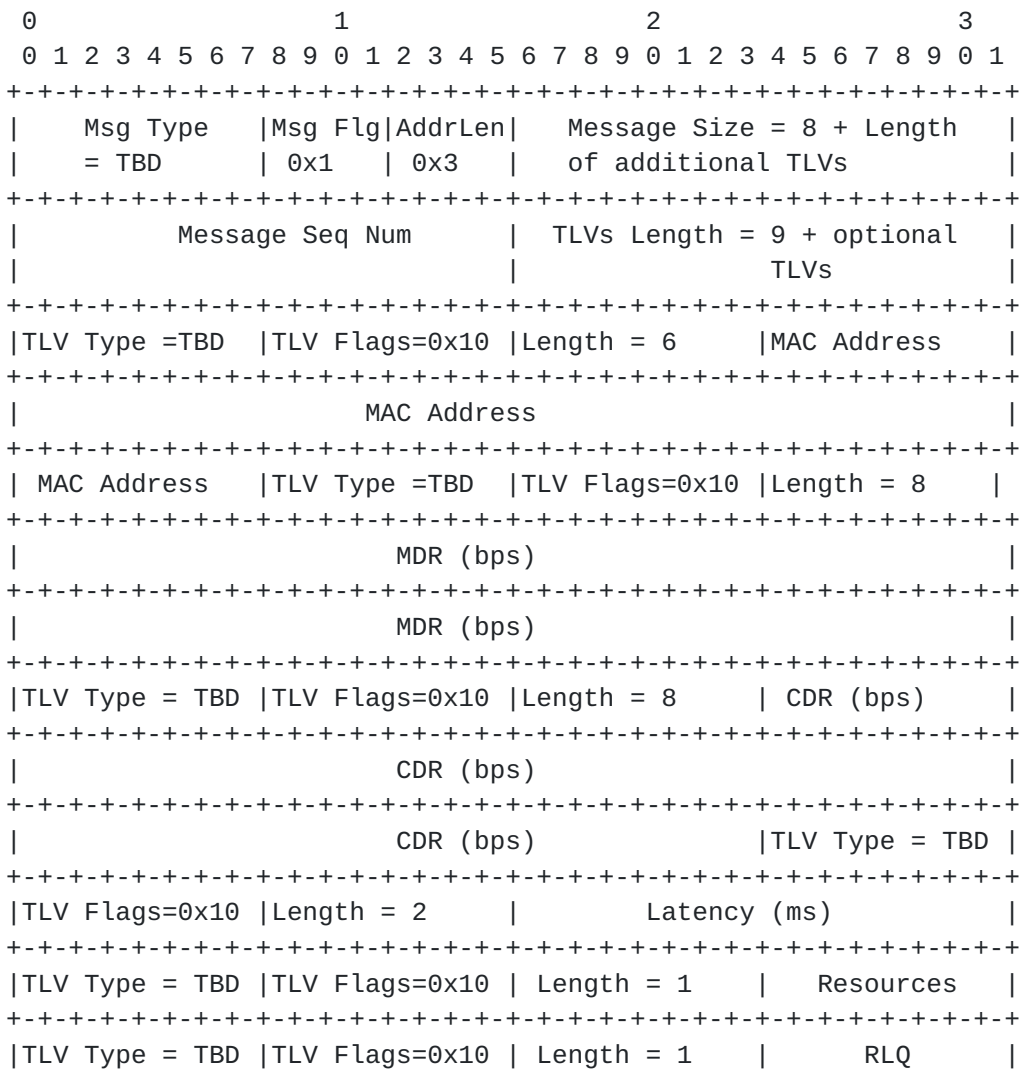

```
  +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

   Message Type            - TBD
```

Message Flags - Set to 0x1 (bit 3, mhasseqnum bit is set). All other bits are unused and MUST be set to '0'.

Message Address Length - 0x3

Message Size - 8 + length of optional (Current Data Rate and/or Latency) TLVs

Message Sequence Number - A 16-bit unsigned integer field containing the sequence number that appeared on the corresponding Link Characteristics Request message.

TLV Block - TLVs Length = 9 + Optional TLVs

MAC Address TLV (MANDATORY)

Maximum Data Rate TLV (OPTIONAL)

Current Data Rate TLV - if present, this value represents the NEW (or unchanged, if the request is denied) Current Data Rate in bits per second (bps). (OPTIONAL)

Latency TLV - if present, this value represents the NEW maximum latency (or unchanged, if the request is denied), expressed in milliseconds, on the link. (OPTIONAL)

Resources TLV (OPTIONAL)

Relative Link Quality TLV (OPTIONAL)

25. Security Considerations

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

26. IANA Considerations

This section specifies requests to IANA.

26.1 TLV Registrations

This specification defines:

- Twelve TLV types which must be allocated from the 0-223 range of the "Assigned Packet TLV Types" repository of [RFC5444].

- Seventeen Message types which must be allocated from the 0-127 range of the "Assigning Message TLV Types" repository of [RFC5444].

26.2 Expert Review: Evaluation Guidelines

For the registries for TLV type extensions where an Expert Review is required, the designated expert SHOULD take the same general recommendations into consideration as are specified by [RFC5444].

26.3 Packet TLV Type Registrations

The Packet TLVs specified below must be allocated from the "Packet TLV Types" namespace of [RFC5444].

- Identification TLV
- DLEP Version TLV
- Peer Type TLV
- MAC Address TLV
- IPv4 Address TLV
- IPv6 Address TLV
- Maximum Data Rate TLV
- Current Data Rate TLV
- Latency TLV
- Resources TLV
- Relative Link Quality TLV
- Status TLV

26.4 Message TLV Type Registrations

The Message TLVs specified below must be allocated from the "Message TLV Types" namespace of [RFC5444].

- Attached Peer Discovery Message
- Detached Peer Discovery Message
- Peer Offer Message
- Peer Update Message
- Peer Update ACK Message
- Peer Termination Message
- Peer Termination ACK Message
- Neighbor Up Message
- Neighbor Up ACK Message
- Neighbor Down Message
- Neighbor Down ACK Message
- Neighbor Update Message
- Neighbor Address Update Message
- Neighbor Address Update ACK Message
- Heartbeat Message

- Link Characteristics Request Message
- Link Characteristics ACK Message

27. Appendix A.

Peer Level Message Flows

*Modem Device (Client) Restarts Discovery

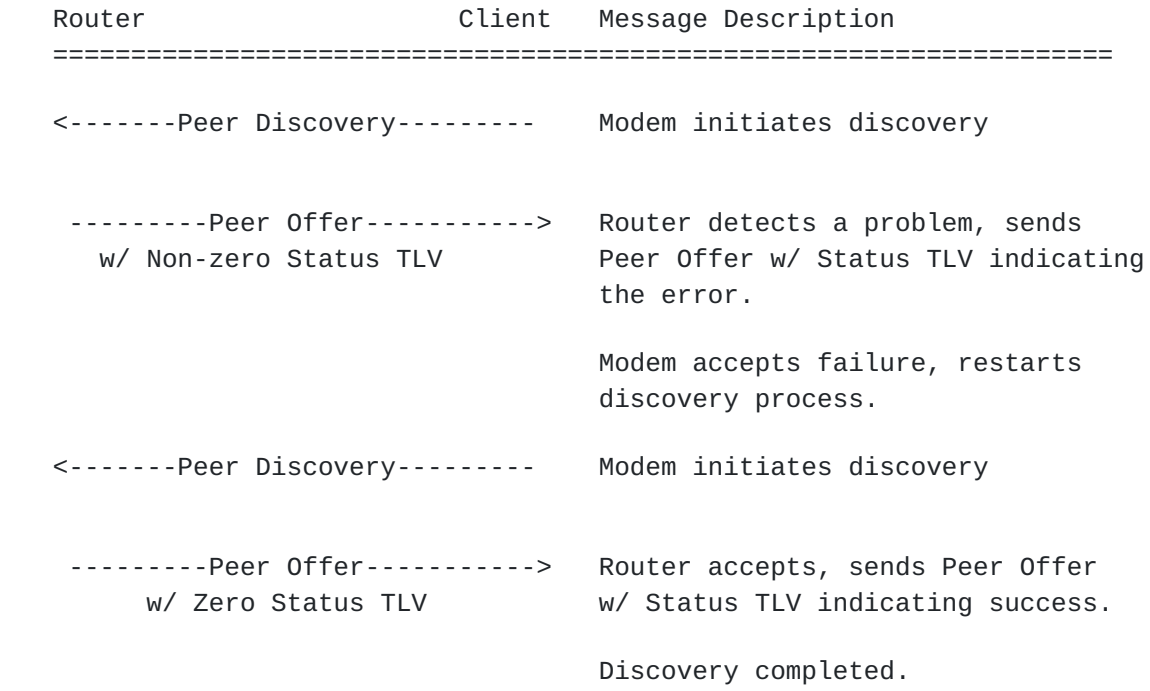

*Modem Device Detects Peer Offer Timeout

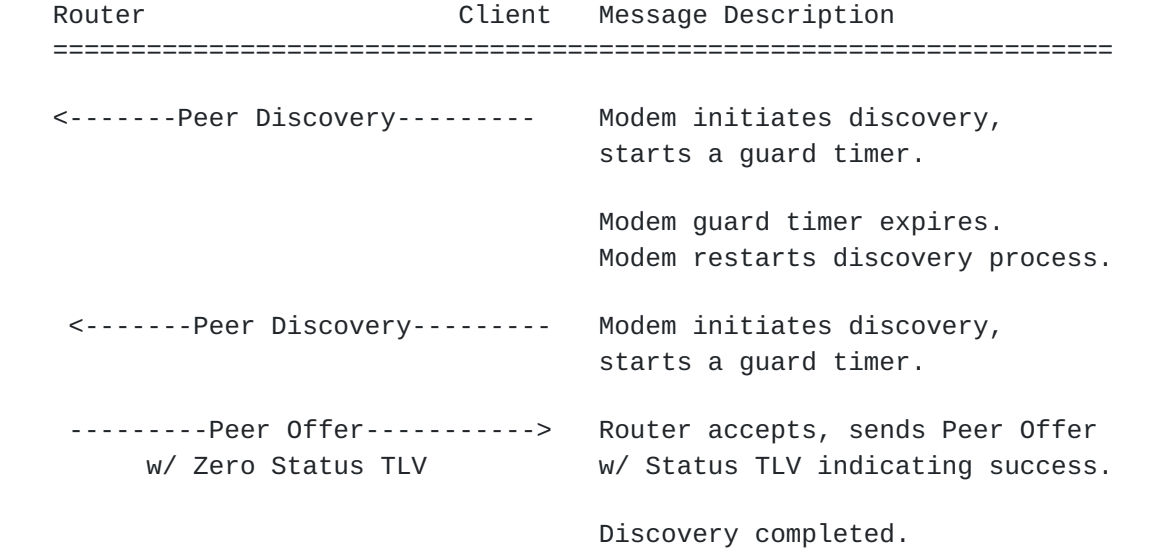

*Router Peer Offer Lost

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

*Discovery Success

Router	Client	Message Description
<-------Peer Discovery---------		Modem initiates discovery
---------Peer Offer---------->		Router offers availability
-------Peer Heartbeat--------->		
<-------Peer Heartbeat---------		
-------Peer Heartbeat--------->		
<==============================>		Neighbor Sessions
<-------Peer Heartbeat---------		
-------Peer Heartbeat--------->		
--------Peer Term Req--------->		Terminate Request
<--------Peer Term Res---------		Terminate Response

*Router Detects a Heartbeat timeout

```
   Router                    Client   Message Description
   ========================================================================

   <-------Peer Heartbeat---------

    -------Peer Heartbeat--------->

      ||---Peer Heartbeat---------

           ~ ~ ~ ~ ~ ~ ~

    -------Peer Heartbeat--------->

      ||---Peer Heartbeat---------
                                    Router Heartbeat Timer expires,
                                    detects missing heartbeats. Router
                                    takes down all neighbor sessions
                                    and terminates the Peer association.

    ------Peer Terminate --------->  Peer Terminate Request

                                    Modem takes down all neighbor
                                    sessions, then acknowledges the
                                    Peer Terminate

   <----Peer Terminate ACK---------  Peer Terminate ACK
```

*Modem Detects a Heartbeat timeout

```
   Router                    Client   Message Description
   ========================================================================

   <-------Peer Heartbeat---------

    -------Peer Heartbeat------||

   <-------Peer Heartbeat---------

           ~ ~ ~ ~ ~ ~ ~

    -------Peer Heartbeat------||

   <-------Peer Heartbeat---------
                                    Modem Heartbeat Timer expires,
                                    detects missing heartbeats. Modem
```

takes down all neighbor sessions
and terminates the Peer association.

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

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

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

*Peer Terminate (from Modem) Lost

```
   Router                    Client   Message Description
   ================================================================

     ||------Peer Terminate--------   Modem Peer Terminate Request

                                      Router Heartbeat times out,
                                      terminates association.

    --------Peer Terminate------->    Router Peer Terminate

    <-----Peer Terminate ACK------    Modem sends Peer Terminate ACK
```

*Peer Terminate (from router) Lost

```
   Router                    Client   Message Description
   ================================================================

    -------Peer Terminate-------->    Router Peer Terminate Request

                                      Modem HB times out,
                                      terminates association.

    <------Peer Terminate--------     Modem Peer Terminate

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

Neighbor Level Message Flows

*Modem Neighbor Up Lost

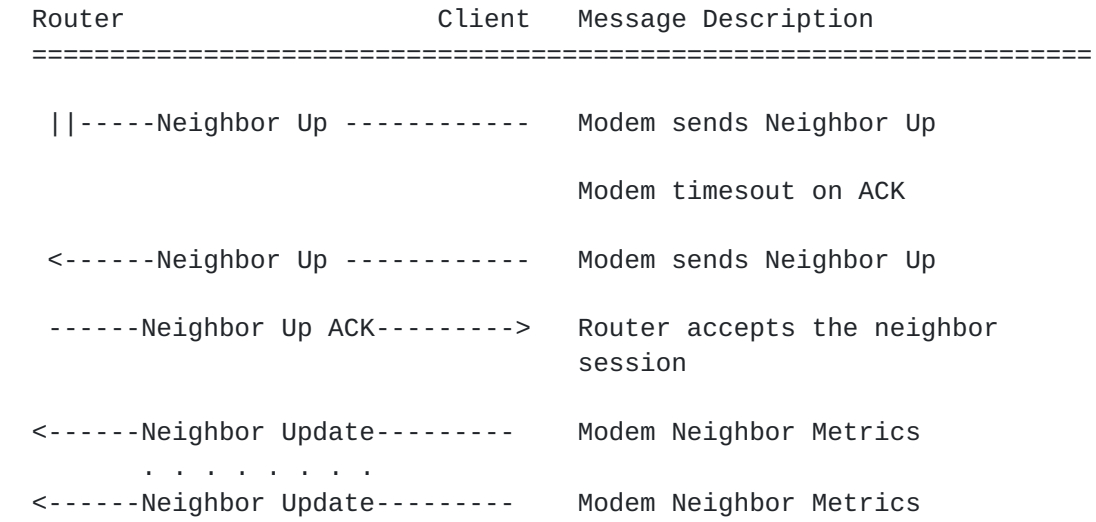

*Router Detects Duplicate Neighbor Ups

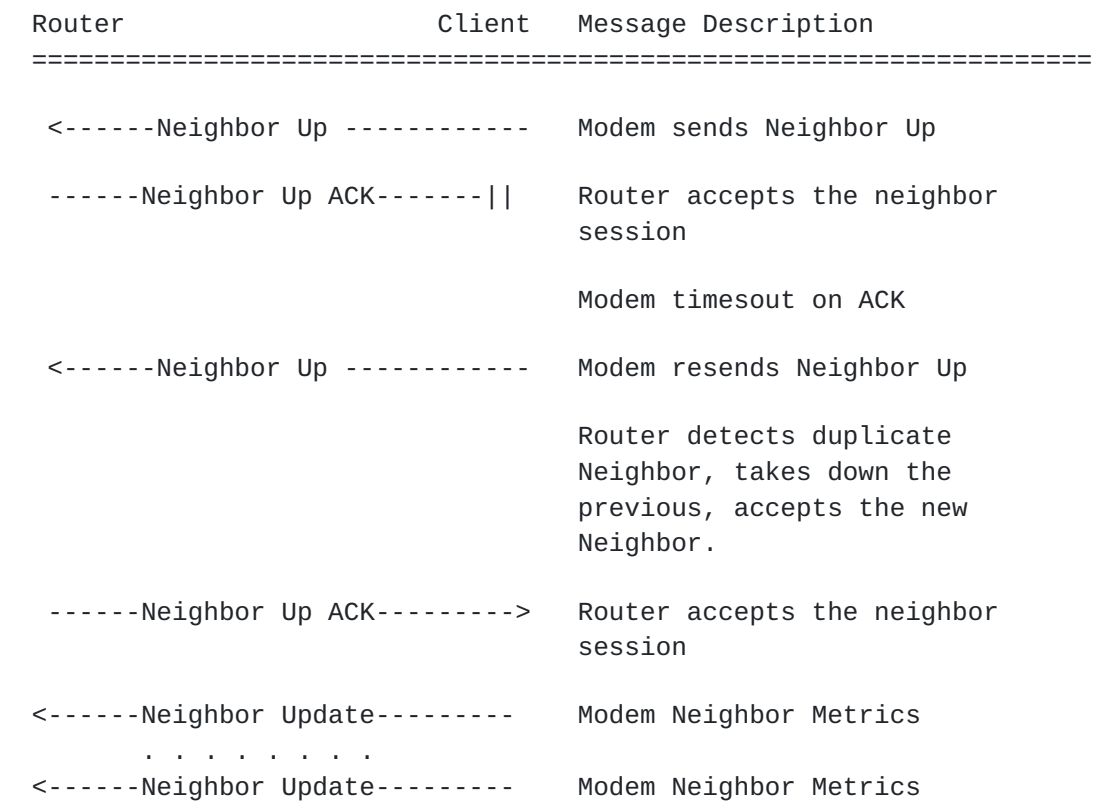

*Neighbor Up, No Layer 3 Addresses

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

*Neighbor Up with IPv4, No IPv6

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

*Neighbor Up with IPv4 and IPv6

Router	Client	Message Description
	<------Neighbor Up ------------	Modem sends Neighbor Up with the IPv4 and IPv6 TLVs
	------Neighbor Up ACK--------->	Router accepts the neighbor session

```
   <------Neighbor Update---------    Modem Neighbor Metrics
          . . . . . . . . .
   <------Neighbor Update---------    Modem Neighbor Metrics
```

*Neighbor Session Success

```
   Router                    Client   Message Description
   ========================================================================

    ---------Peer Offer----------->   Router offers availability

    -------Peer Heartbeat--------->

   <------Neighbor Up -----------      Modem

    ------Neighbor Up ACK-------->     Router

   <------Neighbor Update---------     Modem
          . . . . . . . . .
   <------Neighbor Update---------     Modem

                                       Modem initiates the terminate

   <------Neighbor Down ----------     Modem

    ------Neighbor Down ACK------->    Router

                                       or

                                       Router initiates the terminate

    ------Neighbor Down ---------->    Router

   <------Neighbor Down ACK-------     Modem
```

Acknowledgements

The authors would like to acknowledge the influence and contributions of Chris Olsen and Teco Boot.

Normative References

[RFC5444] Clausen, T., Ed,. "Generalized Mobile Ad Hoc Network (MANET) Packet/Message Format", RFC 5444, Februar, 2009.

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

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate
Requirement Levels", RFC 2119, March 1997.

Informative References

[DTLS] Rescorla, E., Ed,. "Datagram Transport Layer Security", RFC 4347, April 2006.

Author's Addresses

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

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

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

Darryl Satterwhite
Cisco
170 West Tasman Drive
San Jose, CA 95134
USA
Email: dsatterw@cisco.com

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