

DLEP operational experience and proposed updates, revisited

Ronald in 't Velt, IETF 124, MANET WG

Slides 3 through 17 from IETF 110
presentation by Rick Taylor and Lou
Berger, included with their approval

Revisiting a presentation from IETF 110

- **Why again?**

Almost all issues identified are still relevant, but nobody is working on them.

- **Why now?**

Recently, radio manufacturers / waveform implementers have expressed interest in DLEP. However, some have also complained that RFC 8175 is difficult to read and/or ambiguous. There has been talk of drafting a tighter spec outside the IETF.

DLEP operational experience and proposed updates

Rick Taylor and Lou Berger

IETF 110, MANET WG

Outline

Since RFC 8175 (DLEP) was published a lot of operational and implementation experience has been gained, and that has highlighted some issues and inconsistencies.

These really need to be addressed by the WG.

Sorry Stan and Bo.

Issues

- Lack of clarity on metrics
 - Experts spend a lot of time explaining the text to implementers and operators.
- Lack of clarity on multicast behaviour
 - Highlighted by Lou at IETF-105
<https://datatracker.ietf.org/doc/slides-105-manet-dlep-multicast-support-discussion/>
- Incorrect behaviour when responding to Peer Discovery signal
 - Errata raised: <https://www.rfc-editor.org/errata/eid6472>
- Transaction handling rules cause unnecessary session resets
 - This *may* cause interruption of the data plane – bad.
- Excessive session resets due to benign errors

Clarity of metrics

- Maximum Data-rate
 - Not clear that this is the maximum theoretical data-rate of a link, given the current modem configuration/mode, e.g. 11Mbps for 802.11b.
- Relative Link Quality
 - Needs a much clearer definition.
 - Rick describes this as a measure of stability – “How hard the modem is working to maintain the CDR”, i.e. an indication of the probability that the CDR will suddenly change; e.g. Current BER vs FEC mode.
- Resources
 - The only ‘node’ rather than ‘link’ metric.
 - Unclear whether this is the Resources at the remote modem, the remote router, or both?
 - Does the router need to send Destination Announce messages with its Resources metric to the modem so the aggregated value can be propagated across the radio net?

Multicast behaviour

- There is not always a 1-to-1 mapping between multicast MAC and multicast IP address
 - It is not explicit that the multicast IP address **MUST** be included when this is not the case.
- Should the multicast Destination messages include the unicast IP addresses of the receivers?
 - This makes a lot of sense, and at/after IETF 105 Stan was in agreement.
 - This may require the router to announce interest (per [RFC8175] Appendix C.3)
 - Alternately modem could perform IGMP snooping and/or support PIM, BIER, NORM, etc.
 - As usual, how the peer router and modem synchronise their state (about who is watching IGMP) is out of scope of DLEP.
- What about IPv4 broadcast addresses, and by extension, Link-Layer broadcast?

Transaction sequencing

The RFC is explicit about the Request-Response transaction model:

- If a destination-related Request is in-flight, then any received message that is not the corresponding Response causes a session terminating error.
- A transaction applies to a single Destination, i.e. multiple requests may be in progress for different destinations.

However, there are several cases where a non-Response message may be legitimately received while expecting a Response:


- Asynchronous Destination Update.
- Simultaneous Request from both peers, e.g. Destination Announce and a Destination Up.
- Router Destination Down
- Session reset processing

In these cases, a session termination (and corresponding teardown of any routes associated with the session) could cause unnecessary interruption of the data-plane, i.e. “DLEP makes the network less stable”.

Cases identified by looking for terminate in <https://github.com/mit-ll/LL-DLEP/blob/master/Peer.cpp>


Transaction sequencing

Case 1:




Transaction sequencing

Case 2:




Transaction sequencing

Case 3:




Transaction sequencing

Case 4:



Transaction sequencing

Case 5:



Transaction sequencing fixes

Can we define safe behaviour for when an unexpected message can be ignored?

- Do we need to define “idempotent” messages?
- Defining an explicit list of ignorable messages may interact badly with extensions.
- Some messages are valid to be processed rather than ignored, e.g. Destination Updates about a destination in the process of being changed.

Excessive Session Resets

- RFC8175 Section 12.1 has simple, but very strict rules on handling of errors:

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.

- These rules even apply when the message does not impact state!
 - This can lead to session resets
 - See earlier comments about possible impact to data traffic

Examples In Code

- Some example cases of benign cases that trigger resets (May not be a full list):
 - Destination down of unknown destination
 - Why reset, destination state is already removed?
 - Update of unknown destination
 - Is previously discussed case 3 the only source of this?
 - Termination if status is empty

Proposed Next Steps

- We address the Errata.
- We suggest that the WG publish an “update” to RFC 8175, rather than 8175-bis.
 - Any fixes/clarifications to issues mentioned here can be backwards compatible and interoperable with existing DLEP implementations.

Fast forward to November 2025

Next steps:

- Discussion (here and on the ML):
 - Should we be aiming for an 8175-bis or an “updates 8175” document?
 - Who wants to work on this?
 - How to lead external interested parties to the IETF?
 - What about DLEP Extensions? (E.g., RFC 8703)
- Plan an interim (online) meeting on this topic, with Rick participating?