

Router Info Advertisement

draft-zzhang-rtgwg-router-info-01

Kevin Wang, Jeffrey Zhang – Juniper

Changwang Lin – New H3C

Niranjan Vaidya - Broadcom

IETF 121, Dublin

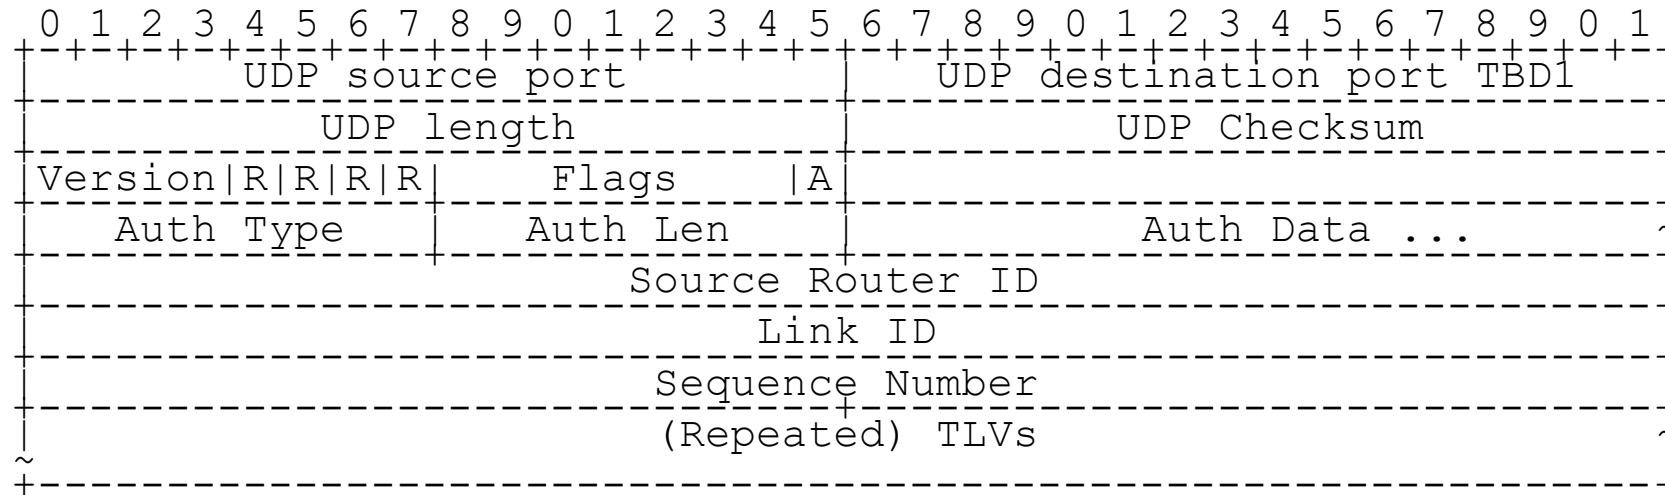
Background

- Two IETF119 presentations/drafts
 - draft-wang-idr-next-next-hop-nodes-00
 - Consider next-next-hop's load information when load-balancing traffic
 - draft-liu-rtgwg-path-aware-remote-protection
 - Quick notification of remote down links for fast reroute purposes
- IETF 120 draft/presentation
 - draft-cheng-rtgwg-adaptive-routing-framework
 - Consider next-next-hop's load information when load-balancing traffic
- All assume fast flooding of link/neighbor/path information
 - Outside routing protocols

Goals and Non-Goals

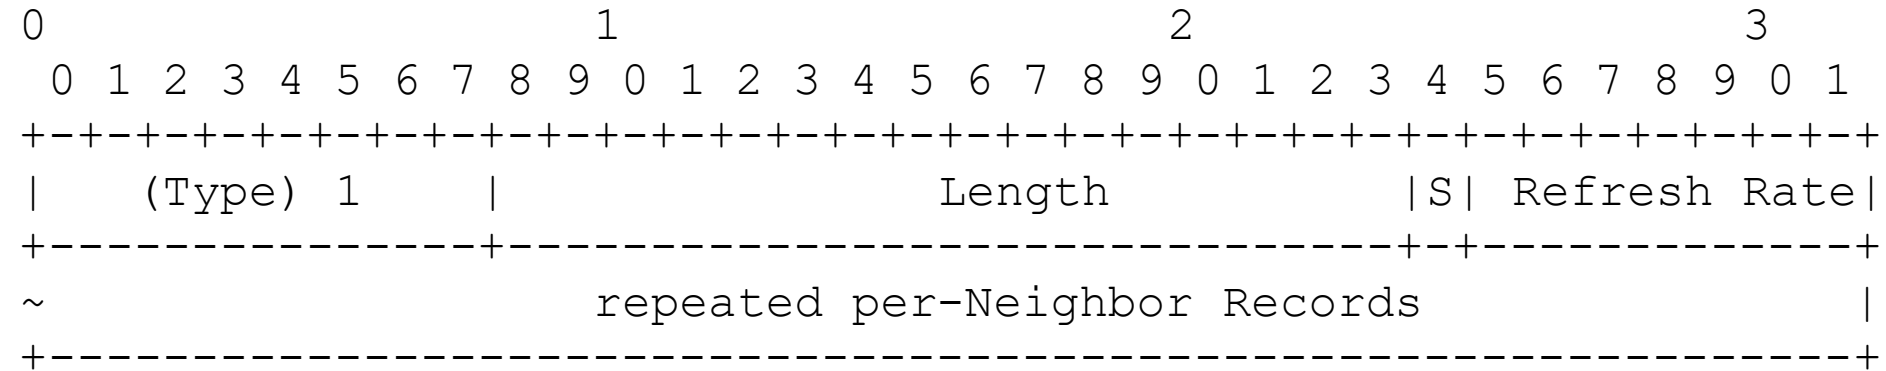
- Specify a generic way of flooding router information
 - Outside routing protocols
 - Easily extensible
 - Hardware friendly
 - Certain information may be sourced and handled in ASIC
- Non-goals
 - Re-flooding of received advertisements
 - E.g., B re-floods C's advertisement to A
 - B may advertise information that it learns/derives from C's advertisement but that is orthogonal to the flooding mechanism defined in this document
 - Specifying how often the information is flooded and how the information is retained/used

UDP-based




- Fragmentation is allowed; though unlikely used when timely flooding is needed
 - Multiple messages are used instead
- TLV-based
 - Typically a message only has one TLV for simpler hardware-based handling
- Authentication optional

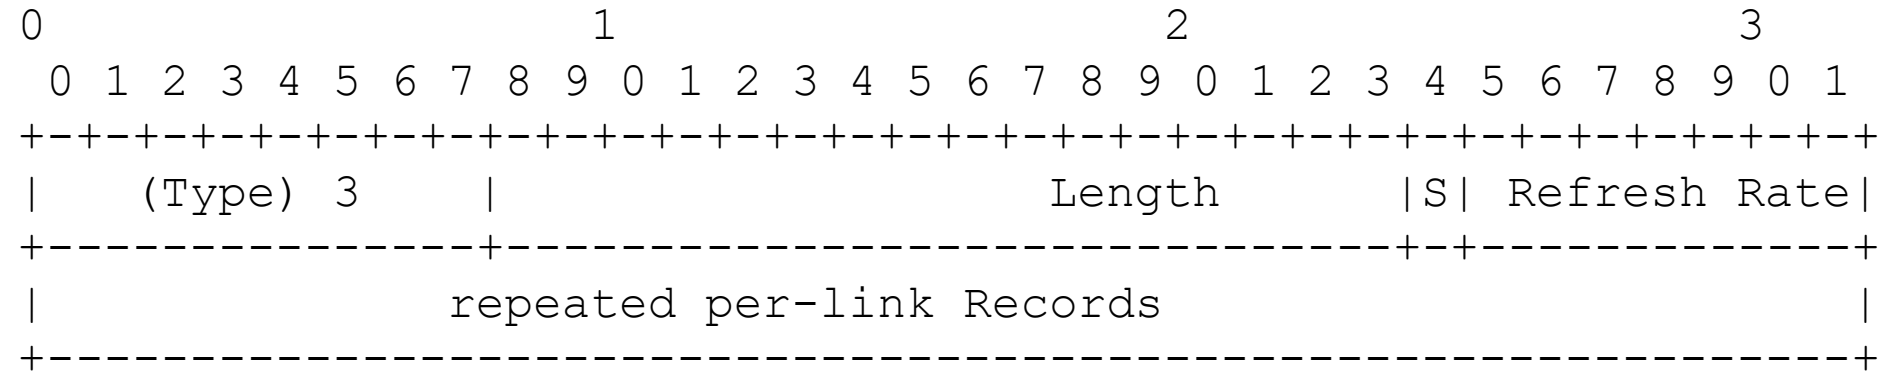
Neighbor Path Info TLV




The per-Neighbor record:



Link Info TLV



The per-link record:



Encoded Info

- -00 only defines value 0
 - Indicating the link/path is down
 - Used for draft-liu-rtgwg-path-aware-remote-protection
- -01 defines new semantics
 - R-bits: reserved
 - U-bit: up/down
 - Quality: 16 levels of path/link quality
 - If the value is N, then up to $(N+1)/16$ of total BW remains

```
  0 1 2 3 4 5 6 7
+--+--+--+--+--+--+
|R|R|R|U|Quality|
+-----+
```

Acknowledgment Option Removed

- Refreshed advertisements remove the need for acknowledgment
 - Bursty refreshes for critical information and then slows down
- Acknowledgement is for hard-state, which incur the burden of adjacency/session management

Mailing List Discussions

- Relationship with LMP/L3DL
 - Both require adjacency/session management and are not designed for fast stateless advertisements
- Clarification on one-hop advertisement
 - Not to all routers in the network through re-advertisement
 - Though the “hop” could be an overlay hop
 - E.g., advertising to all PEs
 - Or even to all routers via multicast (but not via re-advertisement)

Next Steps

- New revision with additional TLVs
 - And co-authors: Jeff Tantsura (Nvidia) and Yisong Liu (China Mobile)
- Security Considerations
- Discussions, comments, and suggestions requested