Reachability and Failure Detection

draft-ietf-shim6-failure-detection-02.txt (to appear)
draft-ietf-shim6-reach-detect-00.txt

Jari Arkko

Iljitsch van Beijnum

Presentation Outline

- Design Decisions
- Strawman Protocols
- Issues

Design Decisions

Definitions and Background

- Available addresses
- Locally operational address pairs
- Operational address pairs
- Unidirectionally operational address pair
- Primary path

Design Decisions

- Multi6 does not go to the area of the configuration modules or protocols -- we shall not reinvent DHCP, and we shall believe what ND tells us
- Own addresses learned locally, peer addresses are communicated
- Multi6 only works as a fail-over, Erik's model:
  - Separate hosts don't share locators to same peer
  - A pair of communicating hosts can have multiple contexts with independent locator choices
- FBD is chosen for simplicity
- Sender chooses outgoing address pair (independently from the other direction)

Strawman Protocols

Reachability protocol (1/4)

If you send payload packets and receive something, assume path is OK

Reachability protocol (2/4)

If you are not sending or receiving payload packets, assume path is OK

A

B

Reachability protocol (3/4)

If you are receiving payload packets but not sending, assume path is OK but send a keepalive

Reachability protocol (4/4)

If you are sending payload packets but not receiving anything, assume the path is broken!

A B

Summary

Outgoing payload?	Incoming Payload	Incoming Keepalive	ACTION
No	No	No	-
No	No	Yes	-
No	Yes	No	Send keepalive
No	Yes	Yes	Send keepalive
Yes	No	No	Start exploration
Yes	No	Yes	-
Yes	Yes	No	-
Yes	Yes	Yes	-

Other Considerations for Reachability

- All activity considered within a time slot *t*
  - Selection of *t* is a tradeoff between failover delay and overhead
  - Listening interval needs to be larger than sending
  - Peers need to be synced about this!
- Higher-level reachability test under extended periods of inactivity?
  - I'm assuming this is combined with context liveness ping
- There's no easy way to measure RTT
  - Except adding TLVs to payload packets

Packet Formats - Keepalive

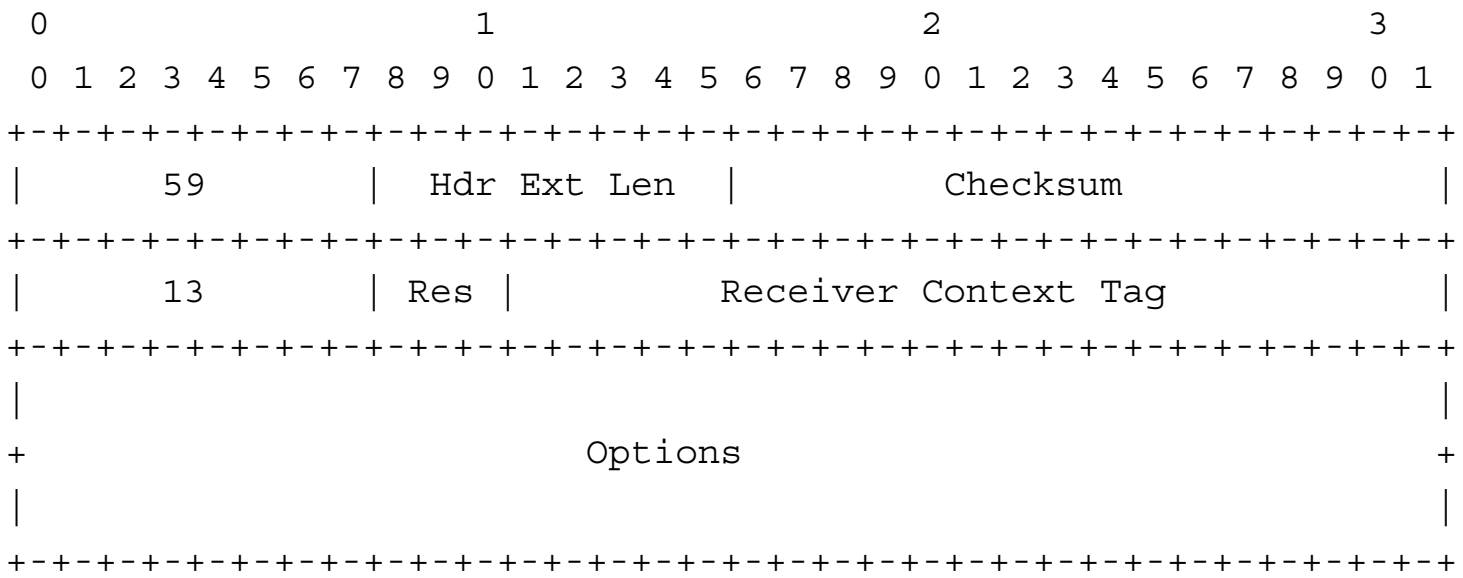

(Also need a time interval agreement in state setup messages.)

Exploration Design

- The chosen reachability scheme implies that it may be just one end who knows there's a problem
  - Given unidirectional connectivity, both ends need to explore => exploration triggered by other end's exploration
- End of exploration explicitly signalled
- Exploration messages need to carry results from other ends's messages

Exploration Algorithm

Initiate exploration if reachability indicates failure, or if peer is performing it

res = empty; sqn = 0; t = 0; succ[sqn] = false

while true **do case**
  when receive explore message:
    res = res | *(create explore result)*
    if opts has an explore results such that result.S && succ[result.sqn] **then**
      exit(OK)
    if opts has explore result **then**
      succ[sqn] = true
  when timeout t**:**
    pair[sqn] = *(select most promising pair)*
    send explore message (src/dst=pair[sqn],S=succ[sqn],opts=res)
    sqn++; succ[sqn] = succ[sqn-1]
    t = *(select suitable timeout)*
done

Exploration protocol

A's locators: A1, A2

B's locators: B1, B2

Current primary pairs are A1 -> B1 and B1 -> A1; now a problem has developed and A1->B2 & B1->A2 are only working pairs

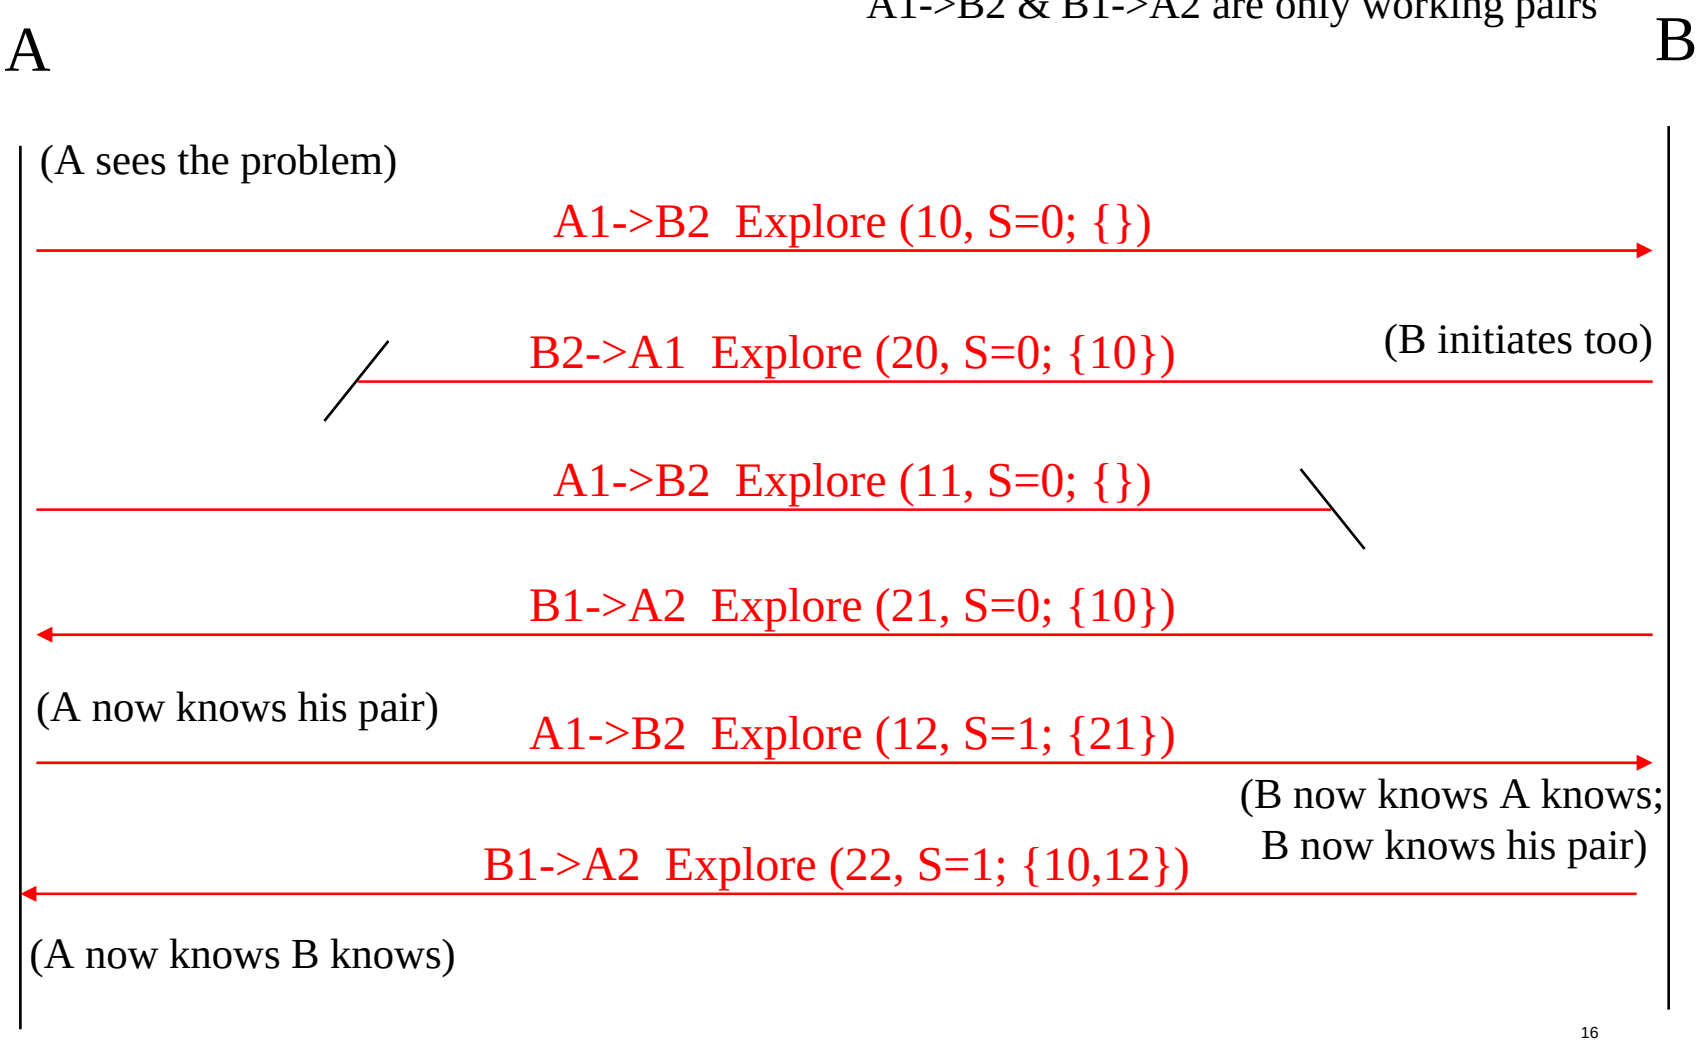

Packet Formats - Exploration

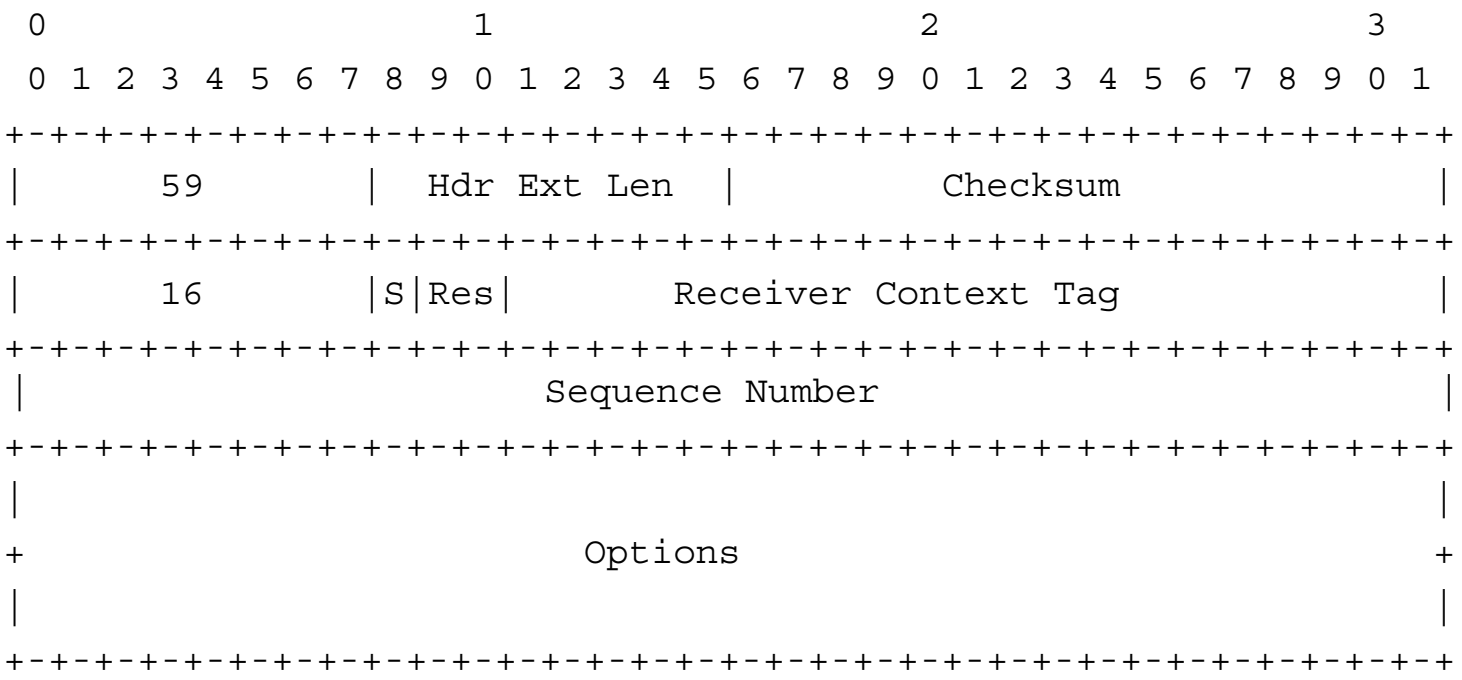

Options in Explore Messages

- Explorer Result
  - You have seen explore message no. N
  - You received it on address X, with the source address of the packet being Y
- S bit
  - I have found my pair but will continue for your benefit (could also be inferred)
- Nonce, ULID
  - Included optionally to perform RR

Issues

Remaining Issues

- Does the termination algorithm work? Need for the S bit?
- Load balancing, best path selection would need something additional
- When does reachability end and exploration begin? Why different messages?
- Interaction between LLU and exploration?
- Any use for ULP feedback?
- Wireless and wired requirements may differ
- Exploration and unequal link costs