

PD route state maintenance approaches
draft-stenberg-pd-route-maintenance-00.txt

Markus Stenberg
<mstenber@cisco.com>

Problem statement

The PD protocol solves the delegation of the prefix but not how to get and maintain routing state for it.

Who does maintenance (options):

- Backend system
- First-hop router (DR)
- Requesting router (DD)

Different approaches (1/2)

- Centralized solution
 - No way for the backend to know DR state
 - No way to push the routes to DR
- DR-based solutions
 - Lease query (on-demand)
 - Requires triggering based on data-traffic / another protocol to maintain state
 - Lease query (anticipatory)
 - Complex to implement, transport debate (TCP)
 - Persistent storage
 - Typically not available

Different approaches (2/2)

- RR-based solutions
 - Routing protocol to the RR
 - Applicable only really in multihoming cases, and even then questions remain
 - Short lifetimes (DHCPv6 T1/T2 values)
 - Current state of the art
 - Keepalive (f.ex. BFD)
 - Considerably less overhead than short lifetimes
 - Layer-2 detection of the link state
 - The cleanest solution, if available

Summary

- There are 8 different ways for maintaining the routing state of the delegated prefixes
 - 4 doable without protocol changes
 - RR-based solutions seem cleanest if available
 - centralized solution not realistic, nor interesting
 - lease queries seem worth pursuing
 - useful in where RR change not possible
 - routing protocol selection options for DHCPv6 may be needed for multihoming cases (otherwise, the routing protocol would need to be statically stored somewhere)