

IP Forwarding in Deep Space: Implementation

Jean-Philippe Dionne

Context

- When a link is down (because orbital dynamics or else), IP packets on intermediate nodes should be temporarily stored instead of dropped
- Preferable if no need to change kernel/base IP stack
- Intermediate node does not retransmit packets. Transport is taking care of the issue.

Method 1

- Linux:
 - use the tc queuing discipline « plug »
 - Two commands to control packet flow: hold and release.
 - when interface/route changes state, set the qdisc plug accordingly

Method 2

- Have a virtual interface (eg tun) in user space, get packets and do storage as needed
- when interface/route changes state or intercepting an ICMP host unreachable, forward or store

Other Methods

- eBPF
- DPDK (Data Plane Development kit)

Status

- Both methods implemented.
 - Tun method in a mix of Go and C.
 - Route change/Link down is detected by receiving an internal ICMP host unreachable from the kernel
- Tested on a single node using ip netns and docker containers
- Next steps:
 - larger tests with quic stack and app
 - Use netlink events to react to route changes