draft-haas-bfd-large-packets

Jeff Haas (jhaas@juniper.net)
Albert Fu (afu14@bloomberg.net)
IETF Montreal 18th July 2018

Motivation

- MTU issue may occur without any indication of link/protocol issue, as protocol hello/keepalive packets are small
- In redundant topology, troubleshooting MTU issue is time consuming due to multiple ECMP paths
  - e.g. Traffic between two end points may have more than 16 ECMP paths in typical Core/Distribution/Access design
- Current network typically involves multiple routing protocols (e.g. OSPF, ISIS, BGP, MPLS etc.)
  - Only ISIS has automatic MTU detection mechanism
  - Routing protocol hellos are handled by control plane, hence unable to use sub-second timer for fast detection

An automatic and fast mechanism for detection of MTU issue is highly desirable in high performance network.

Why use BFD for MTU detection?

- Routing protocols (e.g. OSPF/ISIS/BGP) on modern high performance routers all support using “light-weight” BFD for faster failure detection
- BFD function may be supported on distributed hardware and independent of control plane
  - Enable fast sub-second failure detection
  - Best practice routing protocol design leverages BFD
- BFD failure detection (due to connectivity or MTU) will bring client routing protocols down, providing fast automatic traffic diversion
- Alarm for BFD failure will enable immediate problem detection and follow up

bfd.PaddedPduSize

Current

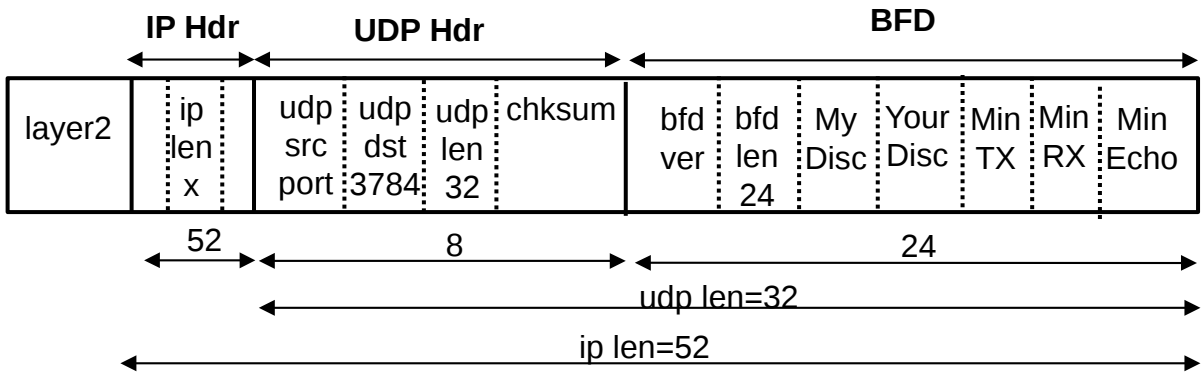

Proposal

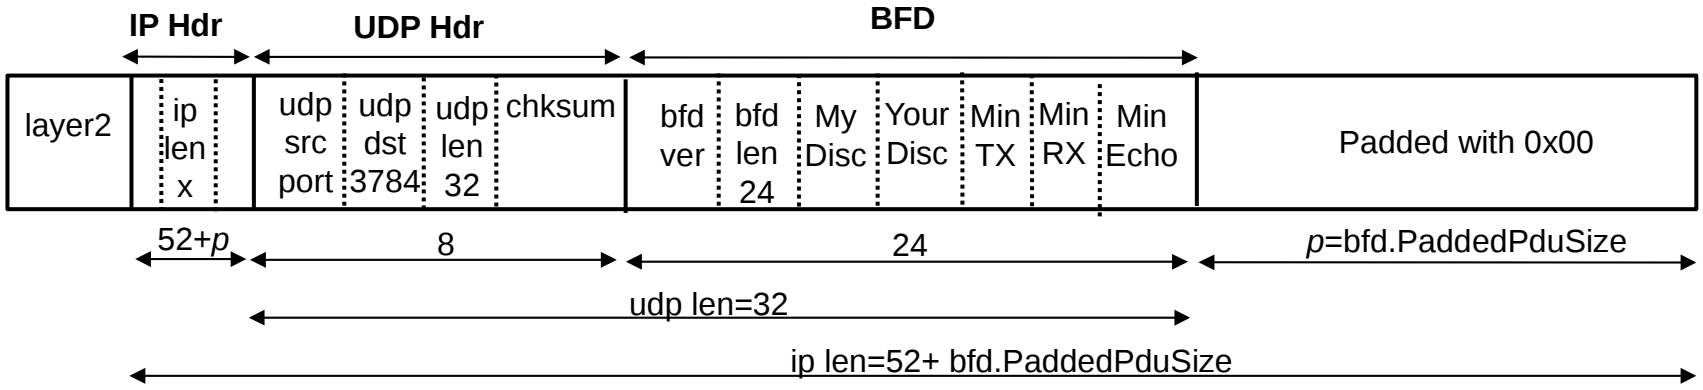

Sample User Cases	Max IP Payload	Bfd.PaddedPduSize (*p*)
Internet Core	1,500	1,448 (1,500-52)
MPLS VPN Core (up to 3 labels)	1,500	1,460 (1,512-52)

Questions?