Encapsulating IP in UDP

draft-xu-intarea-ip-in-udp-03

Xiaohu Xu (Huawei)

Rajiv Asati (Cisco)

Tom Herbert (Facebook)

Lucy Yong (Huawei)

Yiu Lee (Comcast)

Yongbing Fan (China Telecom)

Iljitsch van Beijnum (Institute IMDEA Networks)

IETF95, Buenos Aires

Introduction

- [RFC5640] describes a method for improving the load balancing efficiency in a network carrying Softwire service [RFC5565] over L2TPv3 [RFC3931] and GRE [RFC2784] encapsulations.
  - It requires core routers to perform hash calculation on the "load- balancing" field contained in tunnel encapsulation headers (i.e., the Session ID field in L2TPv3 headers or the Key field in GRE headers), which is not widely supported by existing core routers.
- Since most existing core routers already support balancing IP traffic flows based on the hash of five-tuple of UDP packets, why not carry Softwire service over UDP tunnels?

IP-in-UDP Encapsulation

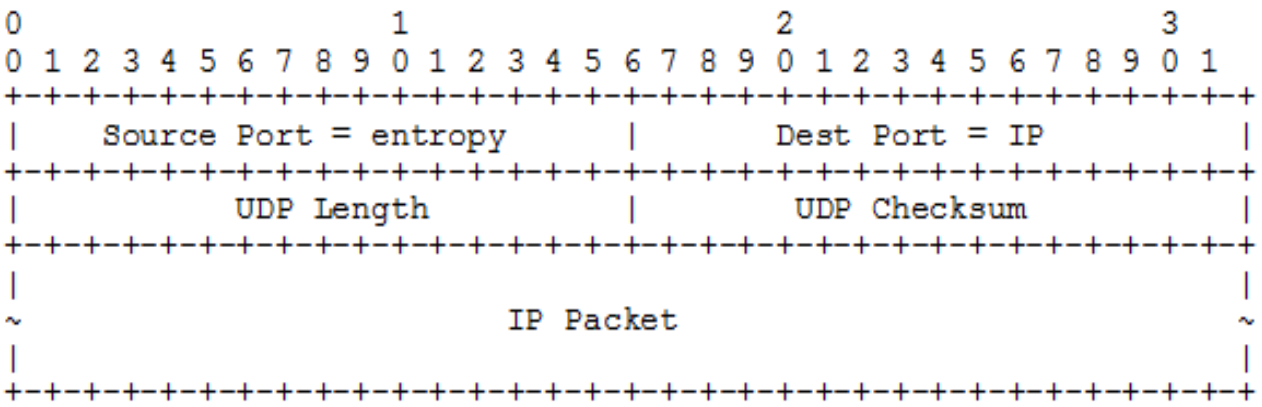

- Source Port: This field contains a 16-bit entropy value that is generated by the encapsulator to uniquely identify a flow.
- Destination Port: This field is set to a value (TBD1) allocated by IANA to indicate that the UDP tunnel payload is an IP packet.
- Checksum: Follow [RFC0768] and [RFC2460] unless one of the exceptions that allows use of UDP zero-checksum mode (as specified in [RFC6935]) applies.

Congestion Control Consideration

- As specified in [RFC5405]: "IP-based traffic is generally assumed to be congestion-controlled, i.e., it is assumed that the transport protocols generating IP-based traffic at the sender already employ mechanisms that are sufficient to address congestion on the path. Consequently, a tunnel carrying IP-based traffic should already interact appropriately with other traffic sharing the path, and specific congestion control mechanisms for the tunnel are not necessary".
- Since IP-in-UDP is only used to carry IP traffic which is generally assumed to be congestion controlled by the transport layer, it generally does not need additional congestion control mechanisms.

Next-steps

- The use of UDP tunnels with source port being used as an entropy field has been widely accepted in the IETF community.
  - LISP [RFC6830]
  - VXLAN [RFC7348]
  - MPLS-in-UDP [RFC7510]
  - TRILL-in-UDP [draft-ietf-trill-over-ip-05]
  - NSH-in-UDP [draft-kumar-sfc-nsh-udp-transport]
  - VXLAN-GPE [draft-ietf-nvo3-vxlan-gpe]
  - GENEVE [draft-ietf-nvo3-geneve]
  - GUE [draft-ietf-nvo3-gue]
  - …
- Solicit WG adoption of this doc (i.e., IP-in-UDP).