

Generic UDP Encapsulation for NVO3

Tom Herbert <therbert@google.com>

Lucy Yong <lucy.yong@huawei.com>

Overview

- UDP encapsulation
- Four byte primary header
 - Version, header length, next protocol, flags
- Message types
 - Control messages (like OAM)
 - Data messages (of an IP protocol)
- Flag-fields like GRE

Salient features

- Foundation
- Next protocol
- Extensibility
- Security

Foundation

- GUE roots lie in GRE
 - GRE is established, well deployed, and **simple**
 - Unfortunately, we hit a wall trying extending GRE
- GUE is generic encapsulation protocol that supports network virtualization
 - Same model of extensibility and simplicity as GRE
 - Header length allows middle box deep parsing
 - Meets isolation and security requirements of NVO3

Next protocol

- IP protocol number
 - Existing IANA number space
 - 8 bits as opposed to 16 bits for Ethertype
- Possibilities
 - IPv4, IPv6, Ethernet, etc
 - ESP, AH, MPLS, etc.
 - Experimental protocols 253 & 254

Extensibility

- How much is needed? (look at comparable protocols)
 - GRE: 13 hdr. bits, 3 32-bit fields
 - IP options: ~20 defined, 10 exp., not well deployed
 - TCP options: ~20 defined, 10 exp., ~5 “private”
- GRE extensibility
 - 48 flag bits, 7 currently defined
 - 16 in primary hdr, 32 in private/extension flags
 - Up to 128 bytes of optional fields
 - Some fields can be repurposed (e.g. security)

Defined/proposed extensions

Defined

- Virtual network identifier
- Security field
- Header checksum
- Remote checksum offload
- Group based policy*

Possibly

- OAM
- Outer/inner TTL mapping
- Congestion control
- Remote segmentation offload

Probably not

- CRC
- Reliability layer
- QoS
- QCN
- Pseudo wire related
- Routing related
- Inband negotiation

GUE security

- Security field
 - Protects VNID, GUE header
 - Anticipate different levels (different field sizes)
 - Simple L2TP-like security cookie defined
- IPsec interaction
 - Header stack: IP|UDP|GUE|ESP|{IP|Ether}
 - All bits created by client are covered
 - GUE header still in outside header for VNID filtering

Request

We would like to ask for WG adoption of Generic UDP Encapsulation as a data plane solution for NVO3.

Thank you!