

# Generic UDP Encapsulation

draft-ietf-nvo3-gue-04

Tom Herbert <tom@herbertland.com>

Lucy Yong <lucy.yong@huawei.com>

Osama Zia <osamaz@microsoft.com>

# Goal

An efficient, extensible, and generic encapsulation mechanism to facilitate packet transport in data center networks for non-virtualization as well as virtualization use cases.

# GUE's roots are in GRE

- GRE is established, well deployed, & **simple**
- Unfortunately, we've hit the wall in trying to extend GRE
- GUE as a “successor” to GRE
  - Retain same model of simplicity and extensibility
  - Allow more opportunity to extend the protocol
  - A few other “improvements”

# Features

- Flag-fields like GRE for extensibility
- Header length allows middle box deep parsing
- IP protocol number indicates next header
- UDP encapsulation to facilitate ECMP
- Data messages as well as control messages (e.g. OAM)
- Security to provide integrity or authentication of header
- Checksum like UDP-lite, tunnel fragmentation
- Hardware friendliness considerations
- Support for network virtualization
- L2, L3, L4 encapsulation

# GUE version 0

Source port	Destination port
Length	Checksum

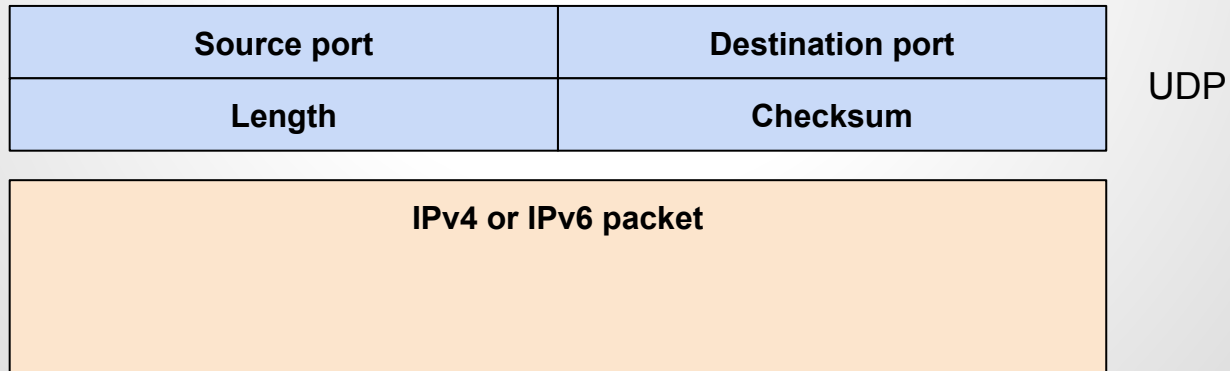
UDP

Ver	C	Hlen	Proto/ctype	Flags
-----	---	------	-------------	-------

Fields (optional)
Private data (optional)

# GUE version 1

- Direct IP encapsulation
- Version 01 coincides with IPv4/v6 version numbers 0100 and 0110
- Header compression, no need to define IP/UDP



# Changes in -04 (input from Adrian Farrel RTG Dir review)

- Remove E bit flag extensions field
- Remove magic number description
- Renamed inner flow identifier to flow entropy
- Described “legal” protocol numbers
- User defined control types
- Expanded IANA considerations
- Defer GUE checksum to extensions draft
- Took out checksumming for L4 (described in TOU)
- Added text to security, reference to security extensions

# Extensions for Generic UDP Encapsulation

draft-herbert-gue-extensions-00

Tom Herbert <tom@herbertland.com>

Lucy Yong <lucy.yong@huawei.com>

Fred Templin <fltemplin@acm.org>



# Summary


Consolidates the “fundamental” set of extensions for GUE.

# Included extensions

- Checksum option
- Fragmentation option
- Security and payload transform options
- Remote checksum offload


# Checksum options

- UDP-lite like checksum
- Covers whole GUE header
- Optional n bytes of GUE payload
- Includes a pseudo header



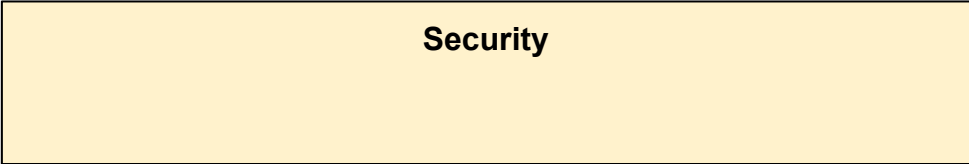
# Fragmentation option

- Fragmentation as part of encapsulation
- Motivated by RFC4459
- Fragment packet, each fragment gets GUE encapsulation
- 40 bit identifier



# Security option

- Authenticating GUE headers
- 64, 128, 256 bit field
- Meaning agreed by end points
- Allows different algorithms (cookies, secure hash etc.)



Security


# Payload transform option

- For encryption, compression, etc.
- DTLS encryption defined
- Type field for different transforms
- Payload type holds protocol number of clear text payload

Type	Payload type	Reserved
------	--------------	----------

# Remote checksum offload options

- Method to leverage checksum offload capabilities of NICs
- Offload outer UDP checksum, inner checksum can be deduced
- Options gives start of checksum coverage and where to write



# Int-area request

Please take up these as WG items:

**draft-ietf-nvo3-gue-04**

**draft-herbert-gue-extensions-00**

Thankyou!