

GUE & ILA update

draft-ietf-nvo3-gue-01
draft-herbert-nvo3-ila-01

Tom Herbert <tom@herbertland.com>

Update overview

- GUE
 - Extensions
 - Enhancements/clarifications for middlebox interactions
 - Support for TLVs in common space?
- ILA
 - Support added to Linux
 - Commencing canary testing and deployment
 - Presenting in v6 ops also

Protocol extensions

Defined

- VNID
- Security field
- Header checksum
- Remote checksum offload
- Fragmentation
- Payload transform
- Session identifier

Possibly

- Passive OAM
- Outer/inner TTL mapping
- Congestion control
- Group based policy
- Segmentation offload

Probably not

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

Middlebox interactions

- Middlebox identification of GUE packets
 - Destination port not enough
 - UDP magic numbers (draft-herbert-udp-magic-numbers-01)
- ICMP error handling
 - Cannot use always source port to find sender
- Stateful middleboxes
 - Session identifier option to allow middleboxes to track “flow state” (draft-herbert-gue-session-id-00)

Potential support for TLVs?

- Could add new option specifying format of private data to allow “infinite” flexibility
- Examples
 - TLVs (like in Geneve)
 - CBOR for SPUD
- No effect on rest of the protocol
- Complexity tradeoffs

Identifier locator addressing

- Network virtualization with encapsulation
- Split IPv6 address into locator and identifier
- Transparent to network and devices
- Use cases
 - Address per task. Simplifies job scheduling
 - Task mobility (virtualizing the data center)
 - L3oL3 alternative for VMs

Identifier Locator Addressing (ILA)

- -01 version
 - Describe application NVo3 Architecture
 - Only use ILA in destination addressed
 - Address per task rationale
- Implementation in Linux 4.1
 - Transparent to network and HW
 - ~2% perf. hit. Acceptable for scaled deployment

Deployment status

- Proceeding with canary testing
- Design for address assignment
- Design/implementation of ILA “router”

Thankyou!