
Generic UDP Encapsulation for IP Tunneling

draft-ietf-tsvwg-gre-in-udp-02

Lucy Yong lucy.yong@huawei.com

July 2014 Toronto CA

Update in v02

- Address the comments from London's meeting
 - UDP checksum
 - Congestion in network
 - Applicability of the solution
 - Others

UDP Checksum Rules

- Tunnel Ingress MAY set zero value in UDP checksum in IPv4 network
 - tunnel egress SHOULD perform the UDP checksum verification upon receiving a packet with a non-zero checksum
 - UDP checksum SHOULD be used in the environment where the UDP packets may be mis-delivered
- Tunnel SHOULD perform UDP checksum in IPv6 network
 - alternative: use IPv6 label as flow entropy if UDP checksum raises a performance concern
 - will change “SHOULD” to “MUST” in next version

Congestion Handling Rules

- If ONLY IP traffic is carried by a tunnel, no need to apply any congestion control mechanism at tunnel endpoints.
- If the end points don't provide congestion control, but the tunnel is used where congestion on the underlying IP network is mitigated by some form of end to end traffic engineering or scheduling, additional congestion control at tunnel endpoints may be unnecessary.
- In the absence of either, a congestion control mechanism SHOULD be implemented at the tunnel ingress and egress.

* Tunnel congestion control mechanism is outside the scope of the draft

Applicability

- GRE-in-UDP encapsulation may be used within a Service Provider (SP) network and/or DC network domain, or inter-domain and Internet
 - Different environments require different tunnel end point process that are specified in the doc.

Others

- GRE-in-UDP is a single tunnel mechanism where both GRE and UNP headers are required
- Tunnel ingress SHOULD perform the fragmentation on a packet before the encapsulation and factor in both GRE and UDP header bytes in the effective MTU size
 - Not perform the fragmentation will cause the packets exceeding network MTU size to be dropped in the network
- Minor editing

Next Steps

- Address Gorry's comments on this version
 - give an example of random constant value selection for UDP source port in the case where tunnel ingress can't get flow entropy from the payload
 - use “MUST” instead of “SHOULD” for requesting use of UDP checksum in IPv6 network
 - more concise text for congestion description; use some text in [RFC5405]
 - others
- Solicit other comments and address them prior to Hawaii meeting

Acknowledgement

Special Thanks to Joe Touch and Gorry Fairhurst for the review, discussions, and suggestions on this update