

# **UDP Usage Guidelines for Application Designers**

**draft-ietf-tsvwg-udp-guidelines-02**

**Lars Eggert & Gorry Fairhurst**

**TSVWG Meeting**

**IETF-69, Chicago, IL, USA**

**2007-7-26**

# Content & History

- guidelines to the designers of applications and application-layer protocols that use unicast UDP
  - presented at IETF-68, adopted as WG item shortly thereafter
  - list discussion has resulted in three revisions since IETF-68
  - contents
    - (1) congestion control
    - (2) message sizes
    - (3) reliability
    - (4) checksum use
    - (5) middlebox traversal
- } new since IETF-68

# Baseline Guideline

unchanged from IETF-68

- apps SHOULD use TCP, SCTP or DCCP whenever they can
- congestion control, message size determination and reliability are difficult to get right
- if used correctly, more featureful transports aren't as heavyweight as often claimed
- if you can't use those transports, use UDP according to the rest of these guidelines

# Congestion Control Guidelines

unchanged from IETF-68

- apps doing UDP bulk transfers → SHOULD use TFRC or TCP-like windowing
- apps that send a small number of messages → SHOULD maintain an RTT estimate and limit themselves to 1 outstanding message per RTT
  - loss looks like long RTT sample

# Congestion Control Guidelines (2)

unchanged from IETF-68

- apps that can't maintain an RTT estimate → SHOULD use a conservative fixed timer and exponentially back it off under loss
  - e.g., 500ms, such as SIP & GIST
- apps that can't detect loss → SHOULD use a more conservative fixed timer
  - e.g., 3 seconds, such as TCP SYN retransmit

# Message Size Guidelines

unchanged from IETF-68

- apps SHOULD NOT send messages larger than the path MTU
- either implement path MTU discovery
- or use IP-layer path MTU information
- or don't send anything larger than the minimum path MTU
  - IPv6 → 1280 bytes
  - IPv4 →  $\min(\text{1st-hop-MTU}, 576 \text{ bytes})$

# Reliability Guidelines

unchanged from IETF-68

- apps should be aware that UDP does not provide
  - reliability
  - duplication protection
  - reordering protection
- apps **SHOULD** be robust in the presence of such events

# Checksum Guidelines

NEW SINCE IETF-68

- IPv4 apps SHOULD use checksums (they're optional in RFC 793)
  - IPv6 apps MUST use checksums anyway
- if data integrity is of importance, SHOULD use stronger checksums on the transmitted data object
- apps that can tolerate data corruption MAY use UDP-Lite (RFC 3828)



# Middlebox Traversal Guidelines

NEW SINCE IETF-68

- apps should implement robust session handling that lets them recover from disappearing middlebox state
- apps MAY in addition send periodic keepalives every 2 minutes
  - keepalives don't invalidate the need for robust session handling
  - keepalive transmission is governed by congestion control

# Status

- authors think -02 is reasonably complete, modulo two issues
  - (1) guideline for keepalive recommendation – what value?
  - (2) congestion control over the entire traffic to a destination
- would like to forward this for early review to other areas, once the WG has come to consensus on these two issues