draft-ietf-tsvwg-rfc5405bis-03

L Eggert
G Fairhurst
G Shepherd

UDP Usage Guidelines

TSVWG Meeting IETF-93, Prague, July 2015

Changes in -03

o Mentioned crypto hash for integrity protection (From Magnus.)

o Mentioned PCP (From Magnus.)

o More accurate text on secure RTP (From Magnus.)

o Section on ECN (Gorry, with help from Mirja)

o Section on RTT Implications on CC (Gorry with help from Magnus)

o Note about other IP protocols (E. Nordmark, rtg encaps guidance)

o Note on reordering between sessions (ECMP, rtg encaps guidance)

o Reworked text on need for off-path data protection (port usage)

Section 3.1.3: RTT

o TCP, SCTP react directly after detecting congestion

- Often not the case for UDP-based apps.
- Can be much more than 1 RTT to change sending rate.

Applications using UDP SHOULD implement a congestion control scheme that provides a prompt reaction to congestion signals (e.g., by adjusting the sending rate within the next RTT).

Applications that do not reduce their rate within one RTT after detecting congestion MUST calculate a safe sending rate, e.g. based on the total time it takes the application to react to congestion, rather than only the measured RTT.

Any implemented congestion control scheme SHOULD result in bandwidth (capacity) use that competes fairly with TCP within an order of magnitude.

Section .1.5 ECN over UDP

o Specifies requirements for using ECN

MUST provide a *method to determine* (e.g., negotiate) … ECN …

A receiver … MUST check the ECN field...

MUST provide *feedback* of congestion information …

MUST provide an *appropriate congestion reaction* ...

SHOULD detect *network paths* that do not support the ECN field ...

A sender is encouraged to provide a mechanism to detect and react

appropriately to *misbehaving receivers*...

Key new recommendations

MAY implement ECN;

a specific set of application (3.1.5)

mechanisms are REQUIRED if ECN is used.

for QoS-enabled paths:

MAY choose not to use CC: (3.1.7)

SHOULD implement a transport circuit breaker

SHOULD NOT rely solely on QoS for their capacity (3.1.8)

non-CC controlled flows SHOULD implement a transport circuit breaker

for non-IP tunnels or rate not determined by traffic:

SHOULD perform CC or use circuit breaker

SHOULD perform CC or use circuit breaker (3.1.9)

SHOULD restrict types of traffic transported by the tunnel

SHOULD use a randomized source port or equivalent (5.2)

technique, and, for client/server applications, SHOULD

ensure responses from src address match request

SHOULD validate payload in ICMP messages

(i) Changes planned for -04

o Section 3.1.7: Pre-provisioned or Reserved Capacity

- expect a strong applicability statement

- example from Section 5 of RFC 7510 (MPLS/UDP)

- last para in Section 3.1.1 may point to 3.1.7

o Update Zero UDP Checksum for IPv6 (David).

- Need to set the scope for usage.

- Need for mechanisms (UDP-Lite or a encapsulation checksum).

- Checksum implications of NAPT/NAT for IPv6 ;-)

(ii) Changes planned for -04

o Incorporate new text on Tunnel Encapsulation.

- Tunnel MTU (Joe & David) & refer to draft-ietf-intarea-tunnels.

- Likely need to reorder tunnel section.

o Check TSVWG consensus on RTT recommendations.

o Add new recommendations in summary table.

Next Steps

New revision published after IETF

Would like to see WGLC.

Thanks to all who provided feedback, including: David Black, Mirja Kuehlewind, Joe Touch, Magnus Westerlund

Details: ECMP

Applications that use multiple transport ports need to be robust to reordering between sessions. Load-balancing techniques within the network, such as Equal Cost Multipath (ECMP) forwarding can also result in a lack of ordering between different transport sessions, even between the same two network endpoints.