University of Glasgow | School of Computing Science

Multimedia Congestion Control: Circuit Breakers for RTP Sessions

draft-ietf-avtcore-rtp-circuit-breakers-09

Colin Perkins – University of Glasgow
Varun Singh – Aalto University

Changes since the -08 draft

- Interval for the RTCP timeout circuit breaker
- Enforce minimum sending rate below which circuit breaker will not trigger
- Adapt time-to-trigger based on the RTCP reporting interval

Interval for RTCP Timeout Circuit Breaker

- When sending, if no RTCP is received for 3 RTCP reporting intervals, then terminate the session
- Issue: this depends on the RTCP reporting interval chosen by the remote system, which is not known
- Solve using fixed RTCP reporting interval:
  - “When calculating the timeout, the deterministic RTCP reporting interval, Td, without the randomization factor, and using the fixed minimum interval of Tmin=5 seconds, MUST be used” (was “SHOULD be used”)
  - “To reduce the risk of premature timeout, implementations SHOULD NOT configure the RTCP bandwidth such that Td is larger than 5 seconds”
    - Plus equivalent for RTP/AVPF profile using T_rr_interval
  - Matches timeout behaviour in RFC 3550 as updated by Section 6.1.4 of draft-ietf-avtcore-rtp-multi-stream

Minimum Sending Rate

- For media timeout and congestion circuit breakers, disable circuit breaker when sending less than one packet per RTT
- Some applications send few packets per reporting interval, and granularity of reported loss fraction is poor – change prevents accidental triggers

- Open issue: do we need to specify a minimum RTT or other bounds?

Adaptive Time to Trigger (1)

- Previous versions could trigger media timeout or congestion circuit breaker after 3 RTCP reports, irrespective of RTCP reporting interval
- This version is adaptive, based on RTCP reporting interval:
  - Timeout if sending media and RTCP reports show no data received for CB_INTERVAL reporting intervals
  - For the congestion circuit breaker, when >CB_INTERVAL reports have been received, on each report calculate weighted average loss fraction over last CB_INTERVAL reports, and derive a TCP throughput estimate based on that average; terminate if actual sending rate >10x estimate
  - CB_INTERVAL = min(floor(3+(2.5/Td)), 30) RTCP reporting intervals

Adaptive Time to Trigger (2)

If CB_INTERVAL = 3:

Average loss rate, $L_{avg} = (L_1 * D_1 / D_T) + (L_2 * D_2 / D_T) + (L_3 * D_3 / D_T)$

Unlike previous versions, trigger when the average loss rate exceeds a threshold, rather than if loss rate in each interval exceeds threshold

Adaptive Time to Trigger (3)

- Scale time-to-trigger so higher rate sessions trigger circuit breaker more rapidly
  - CB_INTERVAL = min(floor(3+(2.5/Td)), 30) RTCP reporting intervals
  - Non-linear, so sessions with lower reporting intervals take proportionally more reporting intervals to trigger:

```
Td =  0.016    CB_INTERVAL =  30   time to trigger =  0.48
Td =  0.033    CB_INTERVAL =  30   time to trigger =  0.99
Td =  0.100    CB_INTERVAL =  28   time to trigger =  2.80
Td =  0.500    CB_INTERVAL =   8   time to trigger =  4.00
Td =  1.000    CB_INTERVAL =   5   time to trigger =  5.00
Td =  2.000    CB_INTERVAL =   4   time to trigger =  8.00
Td =  5.000    CB_INTERVAL =   3   time to trigger = 15.00
Td = 10.000    CB_INTERVAL =   3   time to trigger = 30.00
```

Next Steps

- Feedback from the working group
- Simulate the new mechanisms to ensure they work as expected
- Hopefully can proceed to working group last call before IETF 93