

Multimedia Congestion Control: Circuit Breakers for Unicast RTP Sessions

draft-ietf-avtcore-rtp-circuit-breakers-16

Colin Perkins – University of Glasgow
Varun Singh – callstats.io

Status

| Date | Status |
|---------------|--|
| February 2016 | IETF last call started |
| March 2016 | Submitted -14 to address comments from Ben Campbell and the OPS directorate |
| April 2016 | Submitted -15 to address comments from Spencer Dawkins |
| June 2016 | Submitted -16 to address comments from Stephen Farrell, Alissa Cooper, and Mirja Kühlewind |

Changes in -14:

- Address comments from Ben Campbell & OPS dir:
 - Section 4: clarify that the RTP circuit breaker SHOULD NOT be disabled on networks that might be subject to congestion, even if the peer doesn't support RTCP, unless there is some other way of detecting congestion
 - Section 4.3: clarify that the fraction lost field from a received SR/RR block is recorded
 - Section 4.4: clarify that the time period for congestion to be considered as persistent is on the order of seconds or 10s of seconds
 - Section 4.5: suggest implementations monitor RTCP reception reports to give a UI cue that problems are occurring, before the RTP circuit breaker triggers
 - Section 8: expand discussion of grouping when different DSCP values are used in a bundled group
 - Minor editorial corrections, fix references, etc.

Changes in -15:

- Address comments from Spencer Dawkins:
 - Section 4.5: “RTP flows halted by the circuit breaker SHOULD NOT be restarted automatically unless the sender has received information that the congestion has dissipated, or can reasonably be expected to have dissipated” – expanded discussion, giving example of a mobile device that changes location as “reasonably be expected”
 - Section 7: clarify why ECN-CE marked packets SHOULD be treated as lost, rather than MUST be treated as lost due to the unreliability of ECN marking on some paths

Changes in -16:

- Address comments from Stephen Farrell, Alissa Cooper, and Mirja Kühlewind:
 - Abstract and Introduction: clarify what it means to “operate within the envelope defined by this memo”
 - Section 4.5: clarify that it’s RTP flows sharing a 5-tuple that are stopped when the circuit breaker triggers, rather than the entire call
 - Section 7: expand discussion of reaction to ECN-CE marks, by changing “ECN-CE marked packets SHOULD be treated as if they were lost when calculating if the congestion-based RTP circuit breaker” into significantly more nuanced recommendation, based on likely future evolution of ECN response

Section 7:

There is no consensus on what would be the correct response of the congestion circuit breaker (Section 4.3) to ECN-CE marked packets. The guidelines in [RFC3168] and [RFC6679] are that the response to receipt of an ECN-CE marked packet needs to be essentially the same as the response to a lost packet for congestion control purposes. Since the RTP congestion circuit breaker responds to the same congestion signals, this suggests that it ought to consider ECN-CE marked packets as lost packets when calculating the TCP throughput estimate to determine if the congestion circuit breaker triggers.

More recent work, however, has suggested that the response to an ECN-CE mark ought to be less severe than the response to packet loss.

For example, the TCP ABE proposal

[I-D.khademi-tcpm-alternativebackoff-ecn] makes the argument that TCP congestion control ought to back-off less in response to an ECN-CE mark than to packet loss, because networks that generate ECN-CE marks tend to use AQM schemes with much smaller buffers. For RTP

congestion control, both NADA [I-D.ietf-rmcat-nada] and SCREAM

[I-D.ietf-rmcat-scream-cc] suggest responding differently to ECN-CE marked packets than to lost packets, for quality of experience reasons, but make different proposals for how the response ought to change. Such proposals would imply that a different circuit breaker threshold be used for congestion signalled by ECN-CE marks than for congestion signalled by packet loss, but unfortunately they offer no clear guidance on how the threshold ought to be changed.

Finally, there are suggestions that forthcoming AQM proposals

[I-D.briscoe-aqm-dualq-coupled] might mark packets with ECN-CE in a significantly more aggressive manner than at present. Any such deployment would likely be incompatible with deployed TCP implementations, so is not a short-term issue, but would require significant changes to the congestion circuit breaker response.

Given the above issues, implementations MAY ignore ECN-CE marks when determining if the congestion circuit breaker triggers, since excessive persistent congestion will eventually lead to packet loss that will trigger the circuit breaker. Doing this will protect the network from congestion collapse, but might result in sub-optimal user experience for competing flows that share the bottleneck queue, since that queue will be driven to overflow, inducing high latency. If this is a concern, the only current guidance is for implementations to treat ECN-CE marked packets as equivalent to lost packets, whilst being aware that this might trigger the circuit breaker prematurely in future, depending on how AQM and ECN deployment evolves. Developers that implement a circuit breaker based on ECN-CE marks will need to track future developments in AQM standards and deployed ECN marking behaviour, and ensure their implementations are updated to match.

Response to ECN-CE marks

Two open questions:

- 1) Should persistent excessive congestion signalled by ECN-CE marks trigger a circuit breaker?
 - Should circuit breaker only be triggered by packet loss?
 - The -16 draft states that congestion circuit breaker “MAY ignore ECN-CE marks” – is this appropriate?
 - Mailing list and offline discussion – some people are unhappy with this statement

Response to ECN-CE marks

Two open questions:

2) If the circuit breaker is to respond to ECN-CE, how should it respond?

- RFC 3168 and RFC 6679: response to ECN-CE “MUST be essentially the same” as response to packet loss – the -15 draft reflected this
- L4S BoF proposing alternative ECN semantics, with higher marking rate
- Short queues maintained by CoDel and PIE need different TCP response (Naeem Khademi’s PhD and draft-khademi-alternativebackoff-ecn-00)
- RMCAT congestion control candidates have different ECN-CE response to TCP (NADA and SCReAM respond differently to each other)
- Unclear what is the right response

Response to ECN-CE marks

- What should this draft say about the response to ECN-CE? Proposal:

7. Impact of Explicit Congestion Notification (ECN)

The use of ECN for RTP flows does not affect the RTCP timeout circuit breaker (Section 4.1) or the media timeout circuit breaker (Section 4.2), since these are both connectivity checks that simply determinate if any packets are being received.

At the time of this writing, there's no consensus on how the receipt of ECN feedback will impact the congestion circuit breaker (Section 4.3) or indeed whether the congestion circuit breaker ought to take ECN feedback into account. A future version of this memo is expected to provide guidance for implementors.

For the media usability circuit breaker (Section 4.4), ECN-CE marked packets arrive at the receiver, and if they arrive in time, they will be decoded and rendered as normal. Accordingly, receipt of such packets ought not affect the usability of the media, and the arrival of RTCP feedback indicating their receipt is not expected to impact the operation of the media usability circuit breaker.

Next Steps

- If agreement, submit updated draft for processing by IESG
- Discussion will continue in TSVWG, leading to a future revision of the draft once ECN response is better understood