ECN in QUIC - Questions Surfaced

Magnus Westerlund

Relevant drafts:
draft-ietf-quic-transport-13
draft-ietf-quic-recovery-13

ECN in QUIC Overview

- Packets with ECT or ECN-CE marks acknowledged in ACK_ECN Frame
  - Counters for the markings types
  - Immediate ACK on ECN-CE mark
- Per direction verification of ECT
  - At Start of Connection
  - At Connection Migration
  - Not-ECT will result in ACK frame
- Continuous Verification
- ECN Blackhole Mitigation
  - Optional: Retransmission timeout (RTO) -> retransmit without ECT
  - Implementation freedom

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                     Largest Acknowledged (i)                ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          ACK Delay (i)                      ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        ECT(0) Count (i)                     ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        ECT(1) Count (i)                     ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        ECN-CE Count (i)                     ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                       ACK Block Count (i)                   ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          ACK Blocks (*)                     ...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The ACK_ECN Format

ECN related discussions in QUIC WG

- Optimizing the ACK format
  - https://github.com/quicwg/base-drafts/issues/1439
- Continuous Verification and ACK Loss
  - https://github.com/quicwg/base-drafts/issues/1481
  - Resulted in text changes for (-14)
- Detecting lying Receivers
  - https://github.com/quicwg/base-drafts/issues/1426

Q1: Suppression of ECN values in Packet Duplicates

- QUIC never retransmits the same Packet Number
- On-the-Side Attack
  - Attacker (A) gets a tap on B->C flow from R
  - A modifies ECN field to CE
  - Sends it to C with B as source address
- To mitigate A from reducing B's congestion window
  - C reports ECN only for first packet that arrived
- Missing a CE mark in legit duplicates
  - Delays congestion response to next marked packet

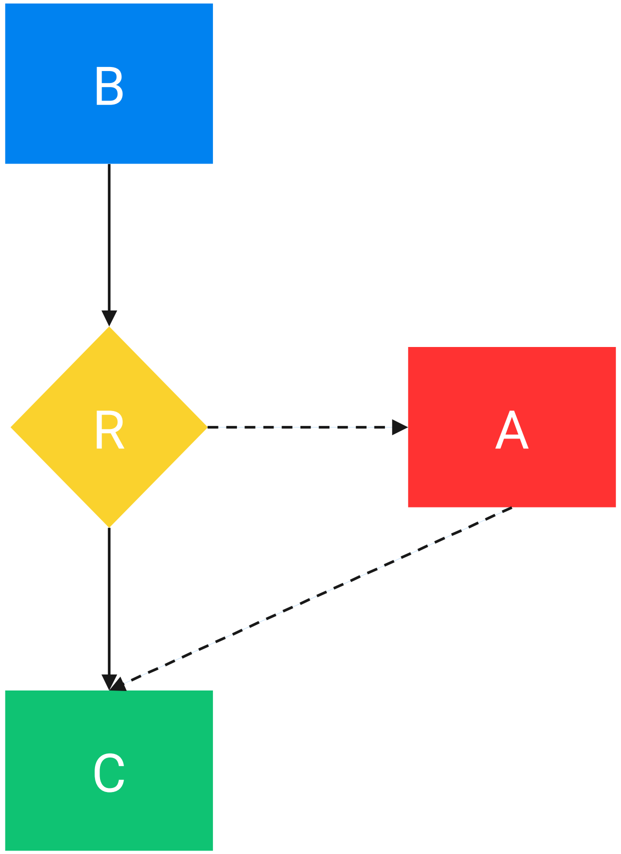

Q2: Will ECT(0) and ECT(1) be mixed in one packet flow?

- Question arose in ACK format discussions
- If a flow will only use one of the ECT code points 0 or 1
  - Build solution utilizing that assumption
  - Signal what will be used
- Is there a need to detect network nodes changing the markings?
  - ECT(0) to ECT(1)
  - ECT(1) to ECT(0)
- If they change, should ECN be turned off?

- RFC 8311 Experimental Types:
  - Congestion Response Differences
  - Congestion Marking Differences
  - TCP Control Packets and Retransmissions
- L4S will use only ECT(1)
- Using alternating ECT markings appear to require
  - Running two parallel controllers
  - Have feedback information for the two sub-flows
- Is this correct?

Q3: Detecting Cheating Receivers

- Sender-side detection of cheating receivers:
  - Receiver that fails to report ECN-CE marks
    - To gain increased throughput
  - Sender marks occasionally a sent packet with ECN-CE from start
  - Sender ignores the CE mark if reported
  - If not reported turn off ECN
- What frequency of test markings are acceptable or allowed?
  - Sender side CE marks can hide real ones
  - A general recommendation would be good
- Related resources:
  - RFC 3168 – Security Discussion
  - RFC 8311 – Declaring Nonce Historic
  - RFC 3540 – ECN Nonce

Q4: Delayed Acknowledgement and ECN

- QUIC allows delayed acknowledgment
- ECN-CE Immediate Acknowledgement
  - Rapid response to Congestion Event
- But what is required for additional ECN-CE marks during the recovery period?
  - Could be delayed while in recovery
    - Will not affect congestion state
  - ECN-CE marks after recovery ends
    - New Recovery period
    - Counters don’t give explicit indication of packet numbers marked

- Currently all ECN-CE marked are sent as immediate ACK
  - Unnecessary many Acknowledgements
- Alternatives
  - Frequent enough acknowledgement
    - Discussion of scaling delay of ACK to a maximum of RTT/4
      - Was implemented in several stacks
  - Use explicit CE reporting so sender knows which Packet Number was marked
  - Provide Receiver with information about when sender exists recovery

Q5: Utility of Detailed CE information

- Using bit vector to provide per packet CE vs ECT information
  - Suggested in discussion of Optimizing the ACK format
- Useful to handle Q4 issues
- What other benefits exists applicable in QUIC?