

TCP ECN

Experience with enabling ECN on
the Internet

Padma Bhooma
Apple

Using ECN from client side

- Apple enabled negotiation of TCP ECN (RFC 3168) from the client-side for the first time on iOS and macOS!
- TCP ECN negotiation is enabled on
 - 5% of randomly selected connections over Wi-Fi / Ethernet in iOS 9 and macOS El Capitan
 - 50% of randomly selected connections over Wi-Fi / Ethernet and a few cellular carriers in iOS 10 and macOS Sierra

ECN Benefits

- Reduce packet loss in the Internet
- Promote Smart Queue Management
- Reduce Buffer bloat
- Improve user experience
- Reference: [draft-ietf-aqm-ecn-benefits-08](#)

Good News!

- It works!
- No problems reported from customers after using ECN on 50% of randomly selected TCP connections on all Apple devices

Heuristics for detecting broken middle boxes

- Middle boxes could treat ECN enabled SYN or data packets differently
- Apple devices have a few heuristics to detect these anomalies
- As a fallback, devices avoid using ECN on those network attachment (path) for a limited period of time after detecting an anomaly

CE marking on every packet

- Impact: Performance degradation
- Heuristic is triggered when:
 - CE marking is seen on Non-ECN connections
 - 7 out of 10 initial packets are marked with CE on ECN enabled connections
- Frequency
 - Fixed on the ISP where it was reported initially in Germany
 - Not seen on any other ISP in the world

Packet Reordering

- Impact: Performance degradation
- Heuristic is triggered when:
 - PAWS drop (RFC 7323) is seen due to reordering
 - Reordering is detected using TCP Selective acknowledgements

TCP SYN loss

- Impact: Adds another RTO during connection establishment
- Heuristic is triggered when:
 - More than 2 successive ECN negotiating SYNs are lost
- ECN is disabled conservatively even if the SYN loss is due to other reasons

RST on first data packet

- Impact: Connection dropped
- Heuristic is triggered when:
 - More than 2 consecutive connections receive a RST for the first data packet after successfully establishing an ECN connection
- Frequency: Seen rarely, still need a metric to quantify the impact

Connection drop after multiple retransmissions

- Impact: Connection failed
- Heuristic is triggered when:
 - More than 4 successively established ECN enabled connections fail to send data after multiple retransmissions
- Frequency: Seen rarely, still need a metric to quantify the impact

More about the heuristics

- Heuristics disable ECN conservatively for limited time even if the underlying problem is not due to ECN marking
- Heuristics are triggered rarely from the data collected so far
- Later, we plan to remove all these temporary heuristics and stop trying to accommodate the few remaining defective middle boxes that mishandle the ECN bits

Networks with CE marking

- Percentage of reports that have seen any CE marking on any of the ECN enabled connections in a 12 hour period

Country	Percentage
United States	0.2
China	1
Mexico	3.2
France	6
Argentine Republic	30

- Marking was mainly seen on the uplink

Performance comparison between ECN and Non ECN connections

- Path characteristics measured over 12 hour period
 - RTT average and variance
 - Percentage of out-of-order bytes
 - Percentage of packet retransmissions
 - Connection drops
 - Percentage of reordered packets

- Every device reports these path characteristics seen on ECN and Non-ECN connections as two separate metrics during a period of 12 hours
- Our goal is to see that ECN connections perform no worse than Non-ECN connections
 - Makes it safe for Apple to continue using ECN
 - Offers clear competitive advantage for ISPs that do Smart queueing with support for ECN

Using Heat maps for data analysis


- Generating Heat maps with ECNOn metric on Y-axis and ECNOff metric on X-axis
- Each cell has a count of the number of reports that fall into that (ECNOff, ECNOn) values as (X,Y) co-ordinates
- Each cell is color-coded using the count of records on that cell

Maximum Heat

ECN Off Out of order byte percentage

Scale
1 cell = 1 percent

ECN On Out of order byte percentage



Color to report count mapping

White	<= 0
Dark Red	1 - 105
Light Red	106 - 1,000
Yellow	1,001 - 10,000
Light Green	10,001 - 100,000
Med Green	100,001 - 500,000
Green	500,001 - 1,000,000
Dark Green	1,000,001 - 10,000,000
Light Blue	10,000,001 - 20,000,000
Med Blue	20,000,001 - 50,000,000
Dark Blue	50,000,001 - 152,660,669

Comparison of Out-of-order byte percentage


- 3 Billion reports from the field were used for this data representation
- Maximum heat is around 0 - 2 percentage
- Report count (heat) is distributed evenly or symmetrically along the diagonal for 0-100 percentage points
- Indicates that the percentage of out-of-order bytes is not worse for ECN connections than it is for Non-ECN connections

Maximum Heat

ECN Off Round Trip Time

Scale
1 cell = 10 ms

ECN On Round Trip Time



Color to report count mapping

White	<= 0
Dark Red	1 - 772
Light Red	773 - 1,000
Yellow	1,001 - 10,000
Light Green	10,001 - 100,000
Med Green	100,001 - 500,000
Green	500,001 - 1,000,000
Dark Green	1,000,001 - 10,000,000
Light Blue	10,000,001 - 20,000,000
Med Blue	20,000,001 - 50,000,000
Dark Blue	50,000,001 - 57,561,044

Comparison of Round Trip Time

- 3 billion reports from the field were used for this representation
- Maximum heat is around 40 - 50 ms
- Again, there is symmetric distribution of report counts (heat) along the diagonal for different values of RTT
- Indicates that average RTT is not any worse for ECN connections because of ECN marking

Summary

- ECN negotiation is enabled on 50% of TCP connections on all Apple devices on Wi-Fi and ethernet interfaces
- No problems reported from the customers
- Observed increasing adoption
- Now is the time for deploying SQM with ECN marking in the network and enabling ECN negotiation on all servers