Results from wide testing of ECN

HOPSRG
IETF 94, November 2015, Yokohama
Tommy Pauly, Apple Inc

- Apple’s deployment of ECN
- How we measure ECN support
- Results
  - ECN-incompatible networks
  - Support for ECN negotiation
  - Support for ECN marking

Deployment

- Enabled in betas of iOS 9 and OS X El Capitan (11)
- Disabled in released versions
- Enabled again in betas of iOS 9.2 and OS X El Capitan (11.2)
    - Enabled by default for Wi-Fi and Ethernet
    - Enabled on cellular for select carriers

Measuring ECN Support

- Aggregated, anonymous data collection
  - # of attempted ECN negotiations
  - # of successful ECN negotiations
  - # of ECN markings on connections
- Targeted testing
  - Open connections to well-known servers
  - Measure negotiation success on various networks

Measuring ECN Support

Many new metrics being collected in the most recent betas, including:

- Negotiations on IPv4 vs. IPv6
- Negotiations on Cellular vs. Wi-Fi
- Fallback due to SYN or SYN-ACK loss
- Excessive reordering on ECN connections
- CE being marked on majority of packets
- RTT comparison ECN vs. non-ECN

ECN Incompatibility

How many networks block or mistreat ECN connections?

- Very few. New metrics should help determine a more precise percentage.
- Two categories
  - Misuse of ECN bits (TOS bits)
  - Performance degradation

ECN Incompatibility: IPSec

In iOS 9 and OS X El Capitan, we added support for RFC 6040, “Tunneling of Explicit Congestion Notification”.

- Replaced RFC 3168 and RFC 4301
- Describes behavior for moving ECN markings between inner and outer IP packets within IPSec tunnels

ECN Incompatibility: IPSec

RFC 6040

3. Summary of Pre-Existing RFCs

> On decapsulation, if the inner ECN field is Not-ECT the outer is ignored. RFC 3168 (but not RFC 4301) also specified that the decapsulator must drop a packet with a Not-ECT inner and CE in the outer.

4. New ECN Tunneling Rules

> If the inner ECN field is Not-ECT and the outer ECN field is CE, the decapsulator **MUST** drop the packet.

ECN Incompatibility: IPSec

RFC 6040

4.2. Default Tunnel Egress Behavior

Arriving Inner Header	Arriving Outer Header			
	Not-ECT	ECT(0)	ECT(1)	CE
Not-ECT	Not-ECT	Not-ECT(!!!)	Not-ECT(!!!)	<drop>(!!!)
ECT(0)	ECT(0)	ECT(0)	ECT(1)	CE
ECT(1)	ECT(1)	ECT(1) (!)	ECT(1)	CE
CE	CE	CE	CE(!!!)	CE

ECN Incompatibility: IPSec

RFC 6040

4.2. Default Tunnel Egress Behavior

Arriving Inner Header	Arriving Outer Header			
	Not-ECT	ECT(0)	ECT(1)	CE
Not-ECT	Not-ECT	Not-ECT(!!!)	Not-ECT(!!!)	<drop>(!!!)
ECT(0)	ECT(0)	ECT(0)	ECT(1)	CE
ECT(1)	ECT(1)	ECT(1) (!)	ECT(1)	CE
CE	CE	CE	CE(!!!)	CE

ECN Incompatibility: IPSec

During the first month of the release of iOS 9 and OS El Capitan, we discovered that one ISP marked the CE bits on every packet in its network

- IPSec could be negotiated, but all ESP packets were dropped by the device, as per RFC 6040
- All customer reports were from a single ISP, so this behavior seems isolated

Marking CE on every packet would also cause ECN-negotiated TCP connections to be throttled

ECN Incompatibility: Performance

Negotiating ECN on some networks causes a throughput degradation of 10-30%

- Not all causes have been identified
- Some causes are due to packets taking different routes based on ECN bits

ECN Incompatibility: Performance

RFC 3168
6.1.5 Retransmitted TCP packets

> This document specifies ECN-capable TCP implementations **MUST NOT** set either ECT codepoint (ECT(0) or ECT(1)) in the IP header for retransmitted data packets...

RFC 1323
4.2.1 Basic PAWS Algorithm

> If there is a Timestamps option in the arriving segment and SEG.TSval < TS.Recent and if TS.Recent is valid (see later discussion), then treat the arriving segment as not acceptable... and drop the segment.

ECN Incompatibility: Performance

Client

Server

Seq. N
ECT
Timestamp 1

2s buffer

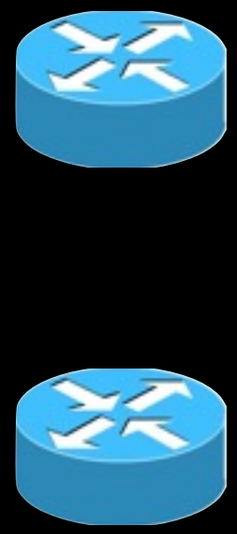

100ms buffer

ECN Incompatibility: Performance

Client

Server

Seq. N
ECT
Timestamp 1

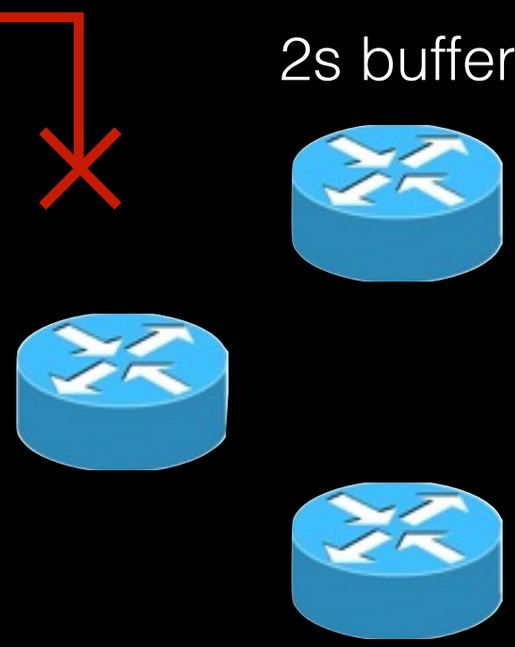

ECN Incompatibility: Performance

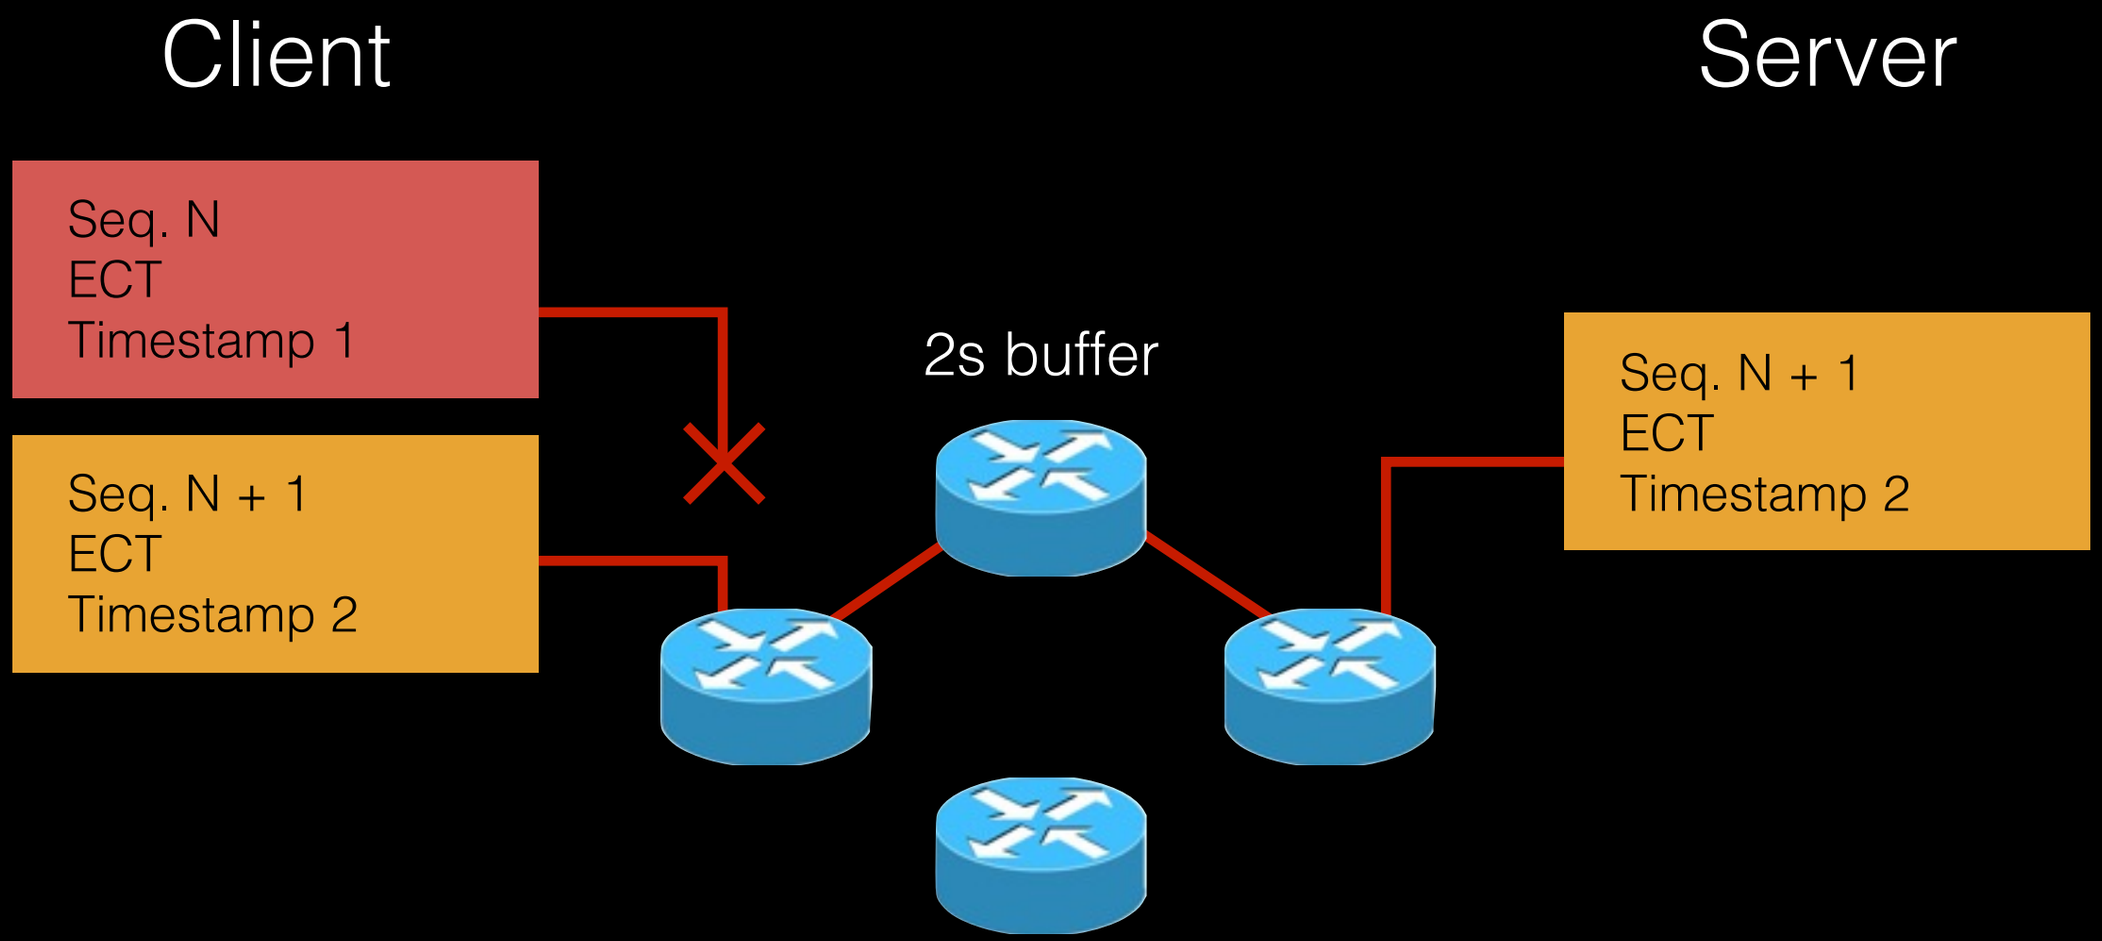

ECN Incompatibility: Performance

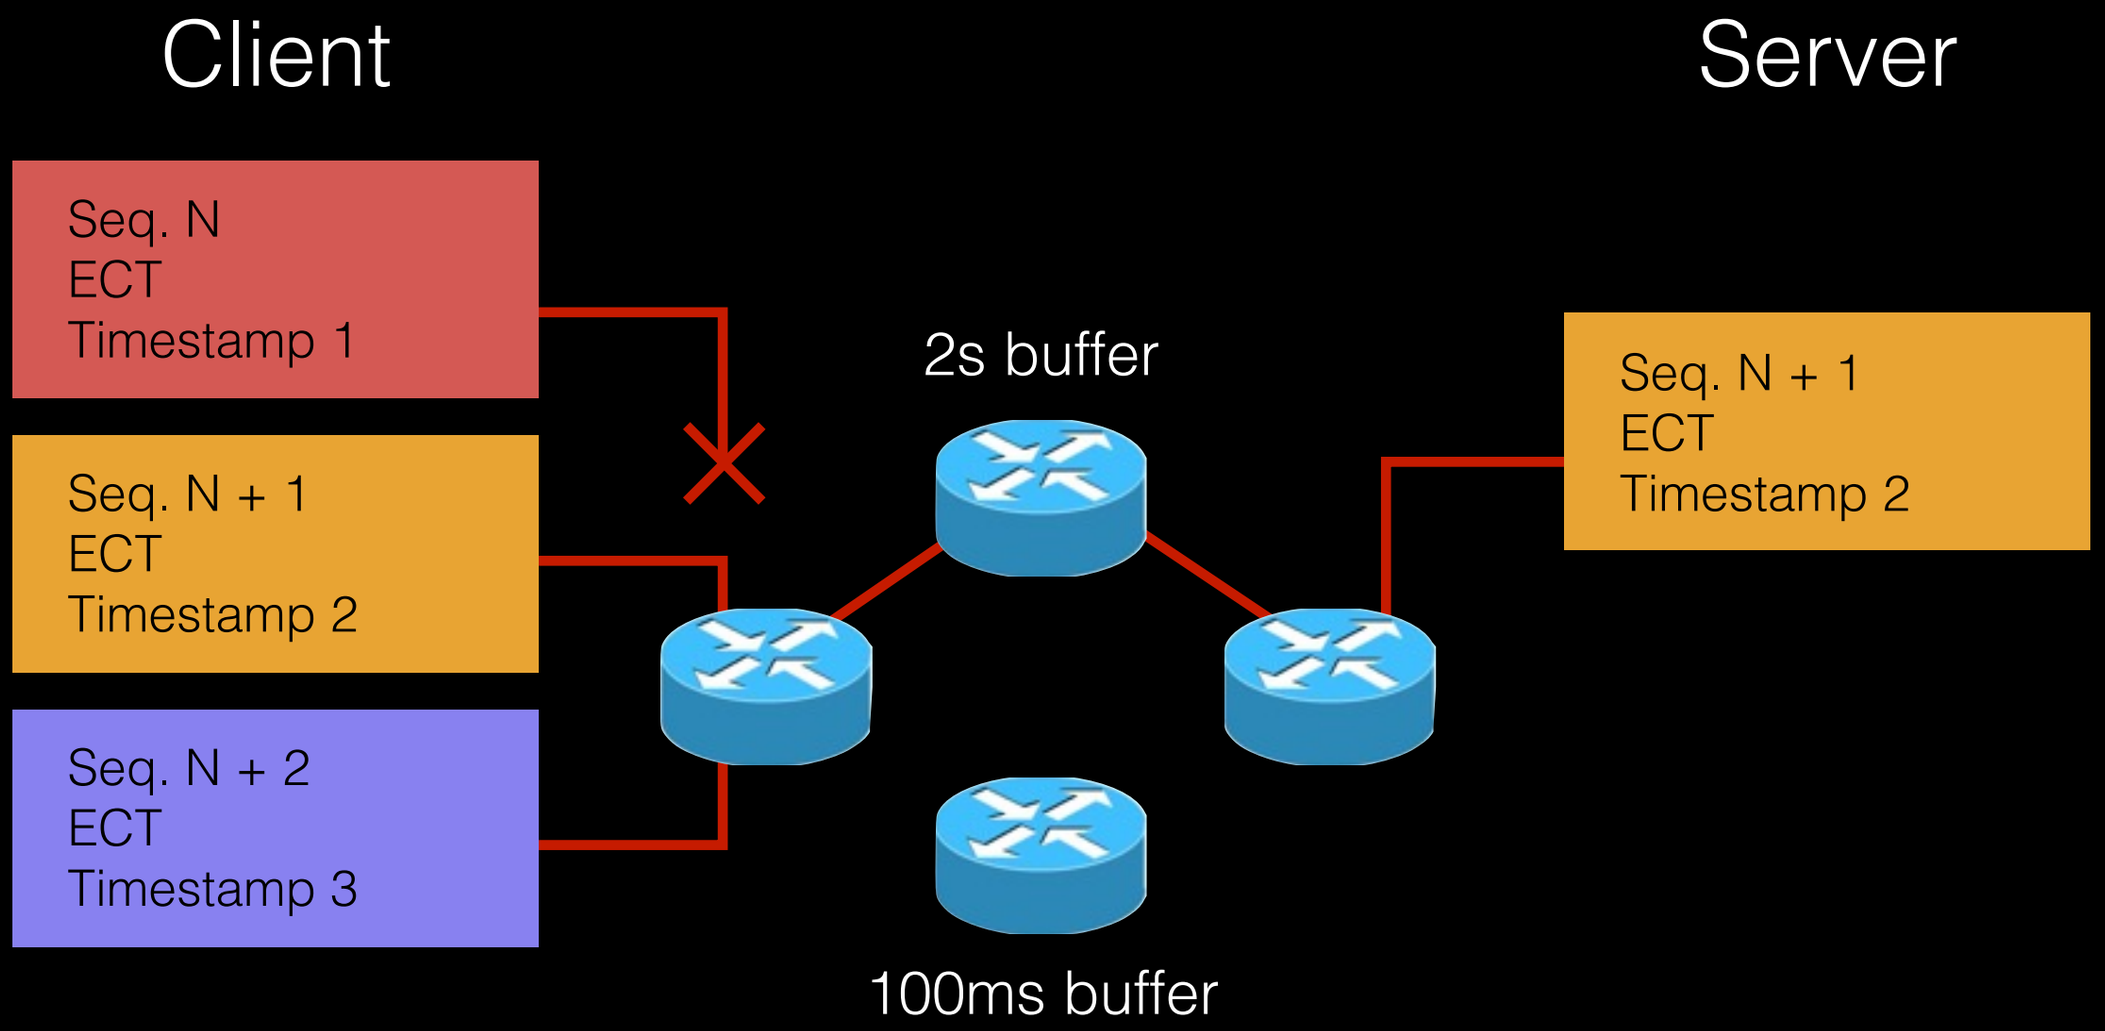

ECN Incompatibility: Performance

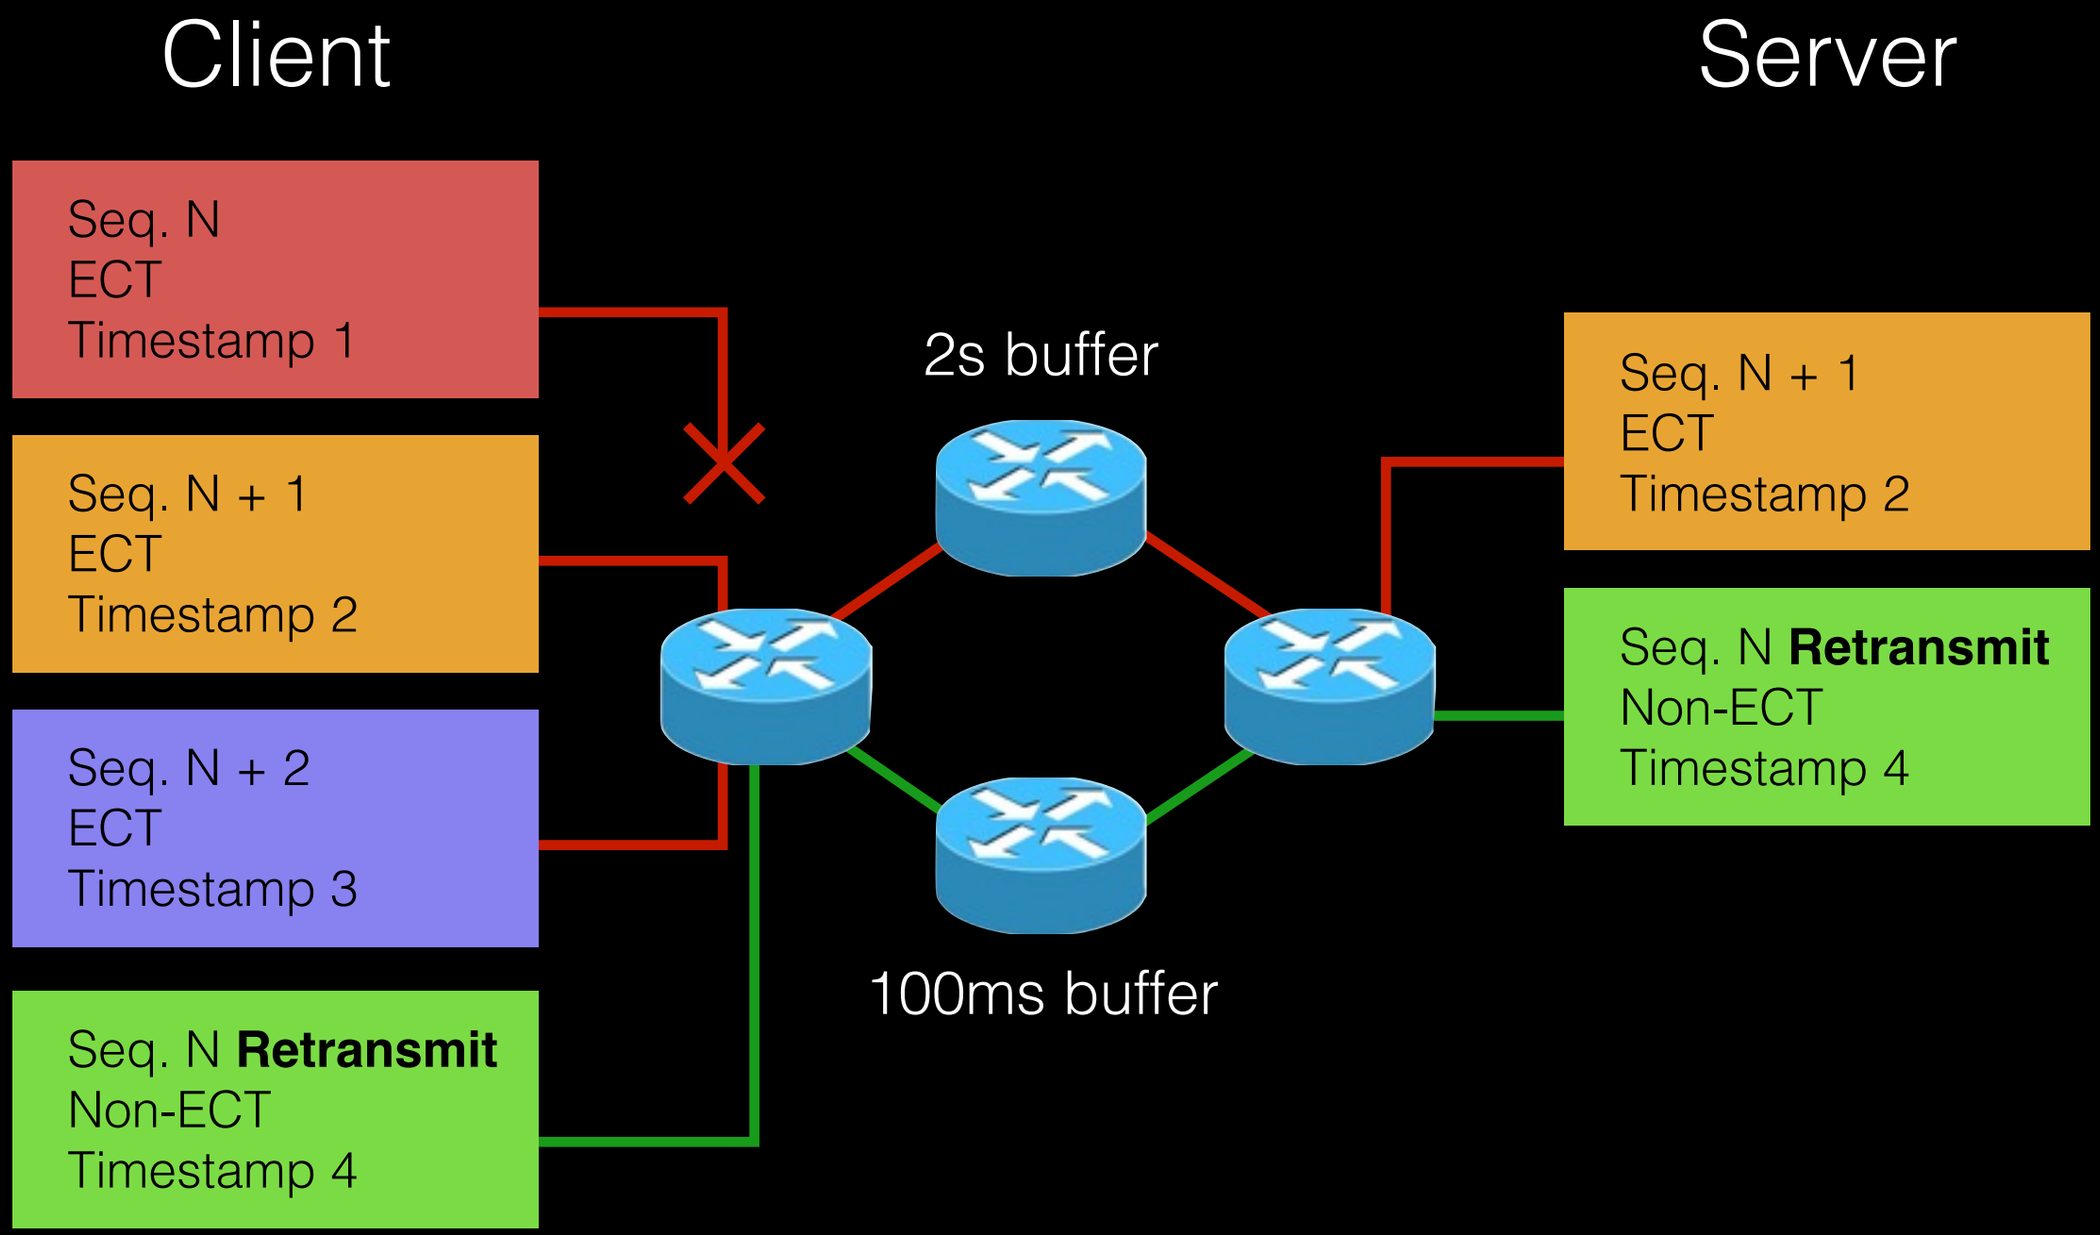

ECN Incompatibility: Performance

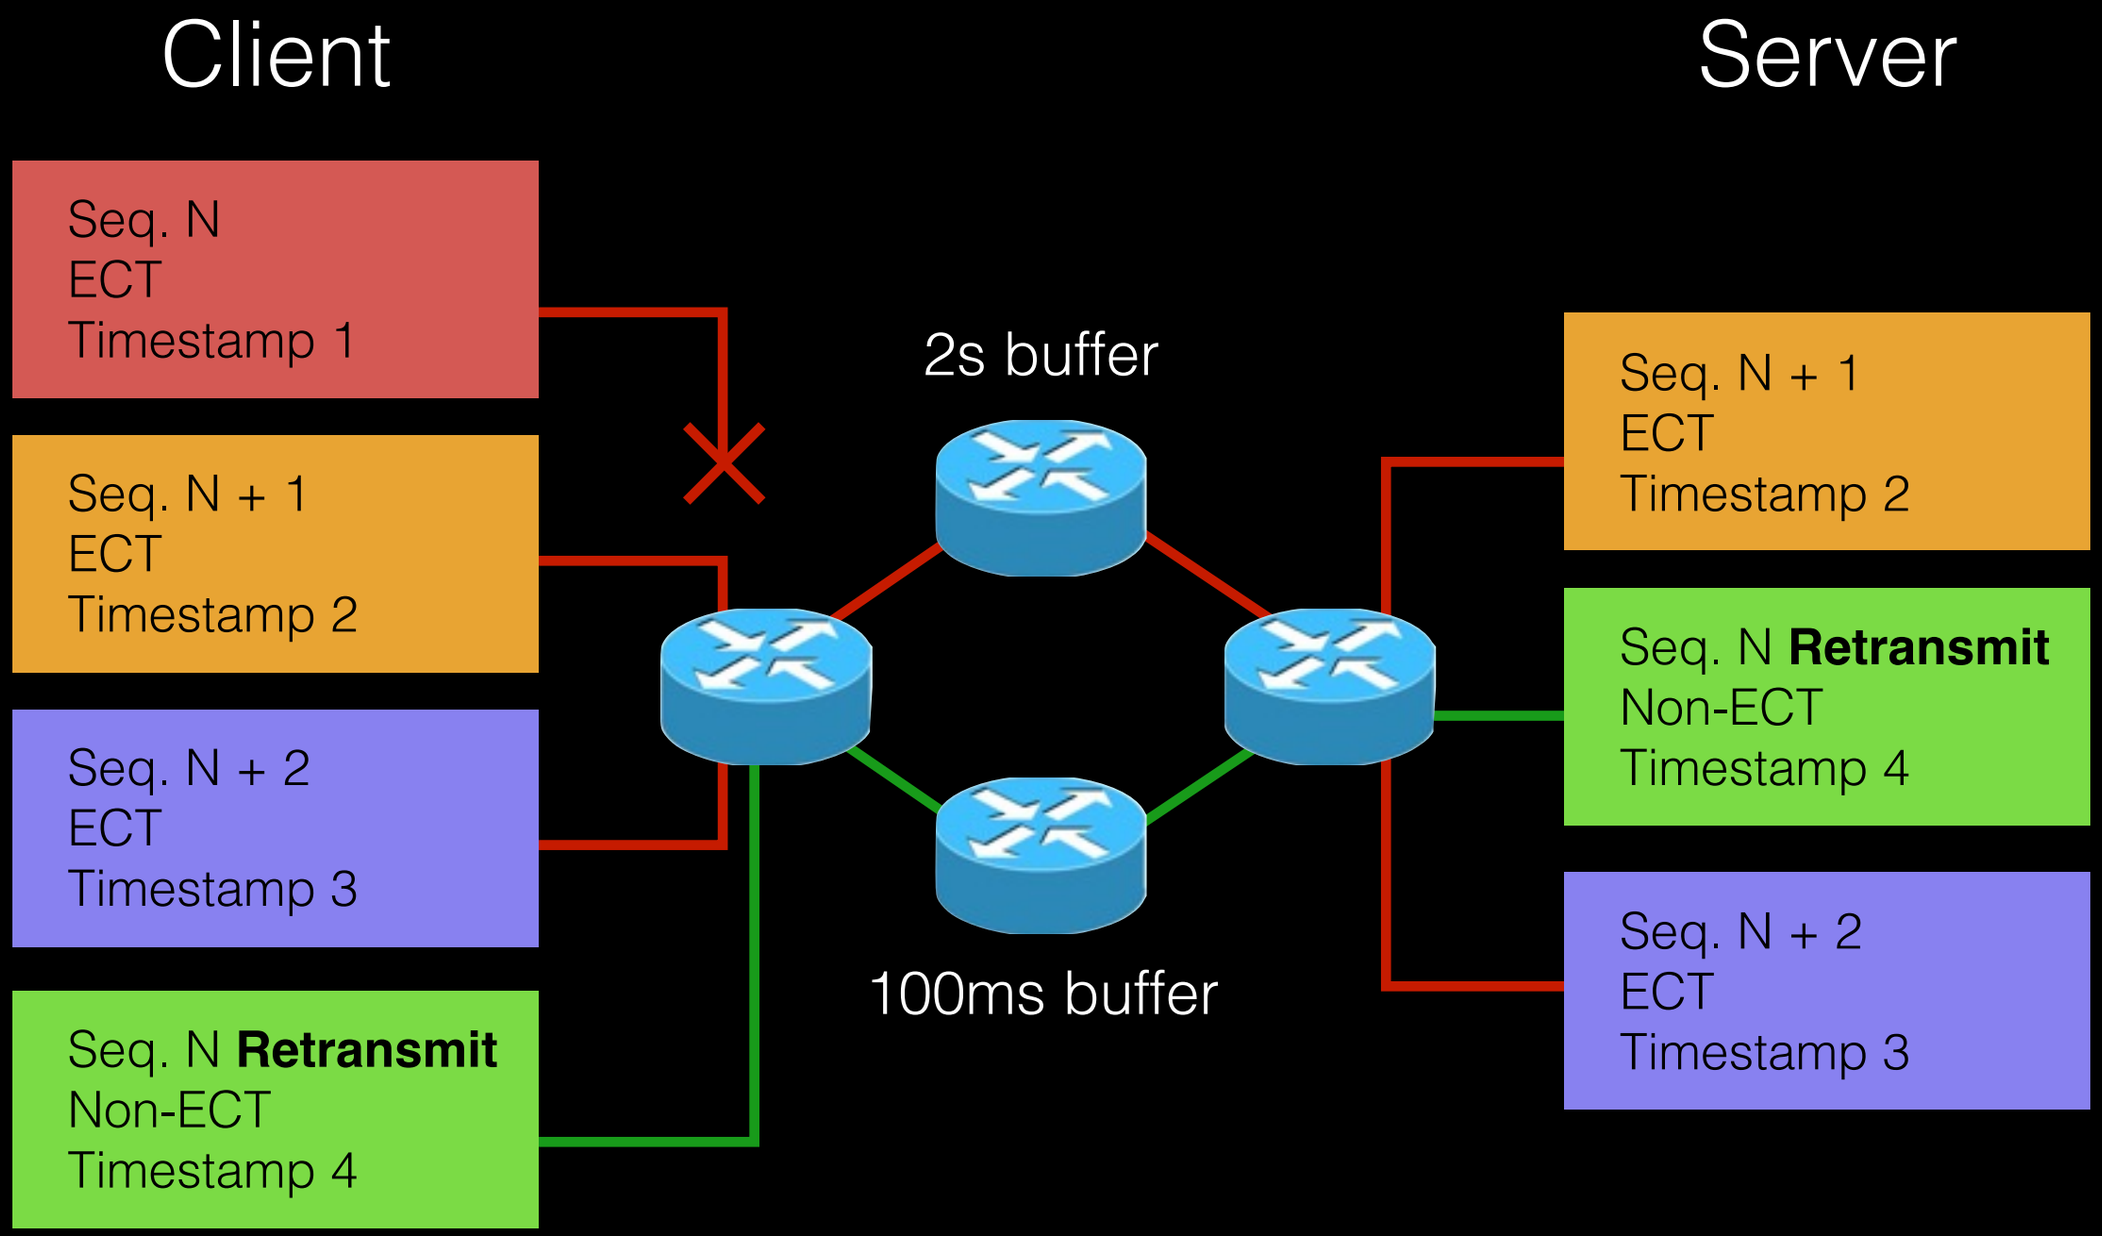

ECN Incompatibility: Performance

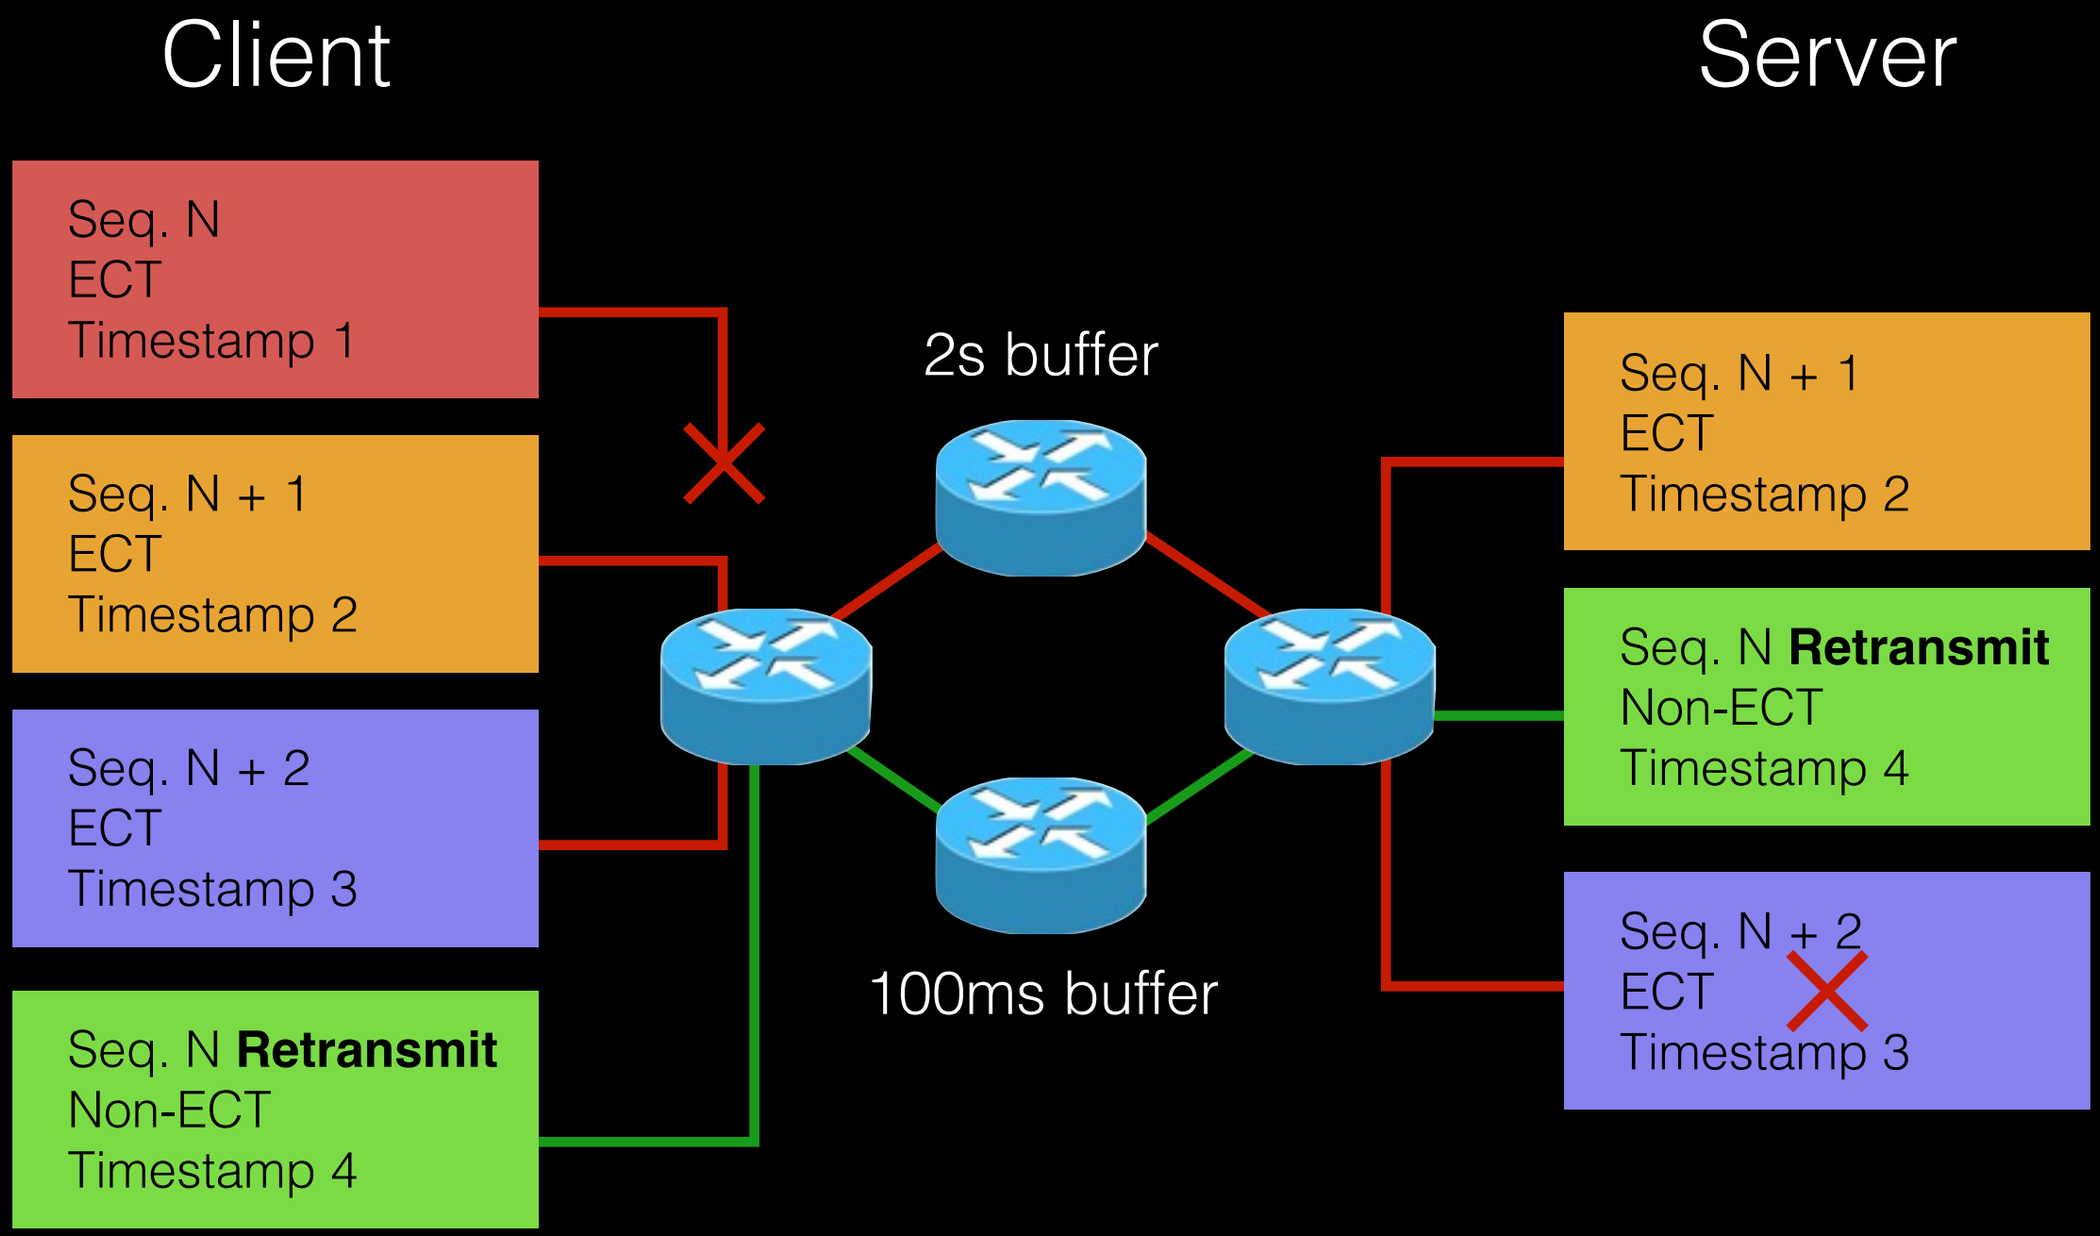

ECN Negotiation

Table 2. ECN negotiation statistics, of 581,711 IPv4 hosts and 17,028 IPv6 hosts, all vantage points, 27 Aug - 9 Sep 2014, compared to previous measurements.

IPv4		IPv6		2011	2012	
hosts	pct	hosts	pct	pct[5]	pct[2]	Description
326743	**56.17%**	**11138**	**65.41%**	11.2%	29.48%	**Capable of negotiating ECN**
324607	55.80%	11121	65.31%	–	–	...and always negotiate
2136	0.37%	17	0.11%	–	–	...sometimes negotiate, of which...
107	0.02%	1	0.01%	–	–	negotiation depends on path
27	0.02%	0	0.00%	–	–	sometimes reflect SYN ACK flags
248791	43.23%	3961	26.23%	82.8%	70.52%	Not capable of negotiating ECN
2013	0.35%	83	0.48%	–	–	...and reflect SYN ACK flags
6177	1.06%	1929	11.33%	–	–	Never connect with ECN (see §3.1)

Enabling Internet-Wide Deployment of Explicit Congestion Notification (Trammell, et al.)

At IETF 93, we reported that 20-30% of TCP connections from iOS and OS X negotiated ECN

ECN Negotiation

Many of the servers that support ECN negotiation in the Alexa top 1000 are Linux servers that support ECN by default.

These generally do not include the top CDN servers that distribute media, which may have the most to gain from ECN.

- Most video streaming does not support ECN
- HBOGO and some Amazon Video do negotiate ECN!

ECN Marking

To see the benefits of ECN, we need bottleneck routers to start marking CE on congestion.

Our aggregated measurements saw some CE marking, but very little (may be noise). New metrics will help determine which markings are legitimate.

ECN Marking

Carrier Networks

- We are working with several carriers to enable marking on their networks

Home ISPs

- No known support for marking at this time

We hope to see progress by the next IETF!

ECN Results Summary

ECN-Incompatible Networks

- 1 ISP marks CE on every packet
- Several ISPs and carrier networks experience reduced performance

ECN Negotiation

- 20-30% of iOS and OS X connections negotiate ECN
- 2 media streaming CDNs negotiate ECN

ECN Marking

- 0 networks reliably mark
- Working with carrier networks to enable marking