ECN Deployment Observations

draft-heist-tsvwg-ecn-deployment-observations

Pete Heist, Jonathan Morton

Overview

- For three weeks, we gathered data on ECN at an ISP’s upstream AS border router. Why?
- More data needed for ECN engineering
  - ECN endpoint deployments
  - ECN marking middleboxes
  - Unexpected uses of ECN field
- Informative, not authoritative

ISP Info

- Cooperative ISP
- Location: Czech Republic
- 660 members
- 861 active IP addresses
- ~5.6 TB/day, 58kpps mean

Collection Method (iptables-ecn)

`iptables-ecn repo: https://github.com/heistp/iptables-ecn/`

- Linux iptables and about 40 ipsets
- Pros: ipsets usable in production, fewer privacy concerns
- Cons: packet counts by IP or IP/port, not flow
- Stats analysis in Go (`ecn-stats.go`)

ipset example
ECT(0) out by IP:

10.45.0.4	32453
10.45.0.9	1717
10.45.0.14	9792
...	...

Observed ECN Endpoint Activity

Clients

- 319.5M TCP SYNs
- 1.44% of SYNs ECN
- 382/681 IPs (56%) saw ECN SYN-ACK
- Suggests low but widespread usage

Servers

- 4.6M ECN SYNs out
- 3.3M ECN SYN-ACKs in
- Suggests high acceptance rate

Detecting AQM Activity

IP	Flags	ECT(0) from WAN	CE from WAN	ECE from LAN	ECT(0) from LAN	CE from LAN	ECE from WAN
10.45.9.88		17970	0	0	0	0	431
10.45.64.3	AK	2909975	36	13348	245614	0	45
10.45.140.73		6036	510	551	1918	0	520
10.45.230.25	A	4560825	3132	18481	290819	0	0
10.45.242.146	A	894737	21	25	85268	0	44

"Possible" AQM Activity Criteria:

- ECT(0) nonzero in both directions
- AND ECE nonzero in either direction
- AND ECE:CE ratio >= 2:1 OR ECE meets same criteria after "anomaly levelling"

There are likely:

- False positives and negatives
- Missed AQMs (need ECN flow and congestion to find them)

AQM Activity for Negotiated TCP ECN Flows

- 90 ECN negotiating IPs saw CE or ECE, 71 from “possible AQMs”
- Of the 71...

	Known AQMs	Unknown, Possible AQMs	Total
# of IPs	38	33	71
Percentage	60.3% (of 63 with known AQM)	10.3% (of 319 without known AQM)	18.7% (of 382 that negotiated ECN flows)

ECN Codepoints on Non-TCP Protocols

- About 0.053% of 43 billion Non-TCP packets had nonzero ECN codepoints
- Many marking ratios not consistent with ECN
- Marking proportion higher from WAN, even with 10:1 ratio of traffic from WAN:LAN
- 6.4 of 6.6 million ECT(1) marks from a single user IP

Direction	CE	ECT(0)	ECT(1)
From LAN	59	26692	28
From WAN	2838929	9562002	6632561

Possible Reasons for ECN on Non-TCP

- Tunneled ECN traffic: can’t be established definitively
- QUIC-ECN: one IP/dstport pair to udp:443 with bi-directional ECT(0) marks (4603 from WAN, 1883 from LAN), it’s possible
- Misuse of the ECN field likely:
  - For historical reasons (obsolete RFC1349)
  - Inadvertently (not shifting DSCP left two bits)
  - Maliciously

Thank You

Anyone care to repeat this experiment? ☺

iptables-ecn repo: https://github.com/heistp/iptables-ecn/

pete@heistp.net