Recycling Large-Scale Internet Measurements to Study the Internet's Control Plane

Jan Rüth, Torsten Zimmermann, Oliver Hohlfeld

appeared at PAM 2019

https://icmp.netray.io
https://www.comsys.rwth-aachen.de/

Montreal / IETF105 MAPRG, July 2019

What and Why?

- **For the past years, I have been scanning the Internet**
  - IETF 101 (London): I presented about the gQUIC deployment
  - We scan a lot: DNS, HTTP/2, TLS, TCP, Cryptominers

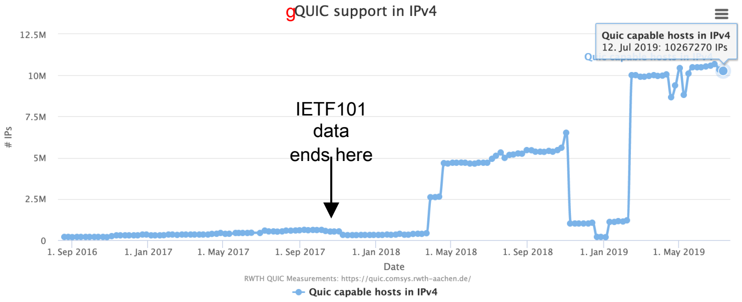

IETF101 data ends here

RWTH QUIC Measurements: https://quic.comsys.rwth-aachen.de/

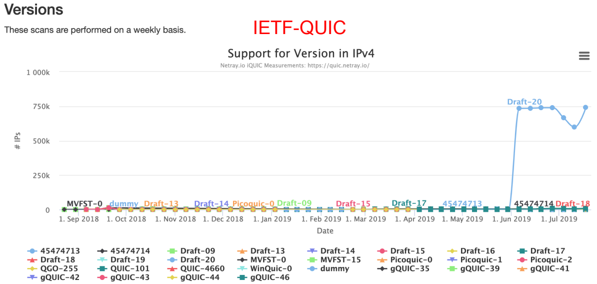

Versions

These scans are performed on a weekly basis.

IETF-QUIC

Netray.io iQUIC Measurements: https://quic.netray.io/

What do we scan but actually aren’t looking for?

Lets Study ICMP!

- **Idea: Let’s use our scans to study Internet Control Messages**

- **In one week we got**
  - 637,500,000 ICMP messages
  - from 171,000,000 different IPs out of
  - 53,000 autonomous systems

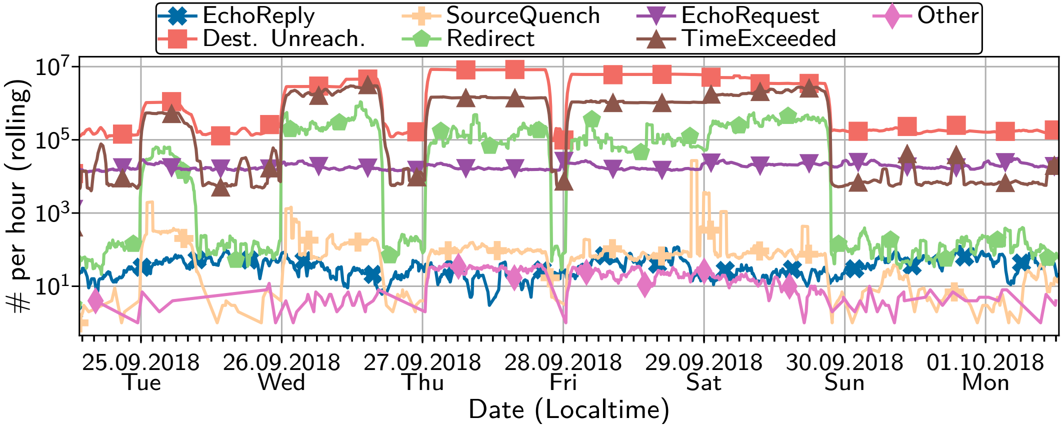

Lets Study ICMP!

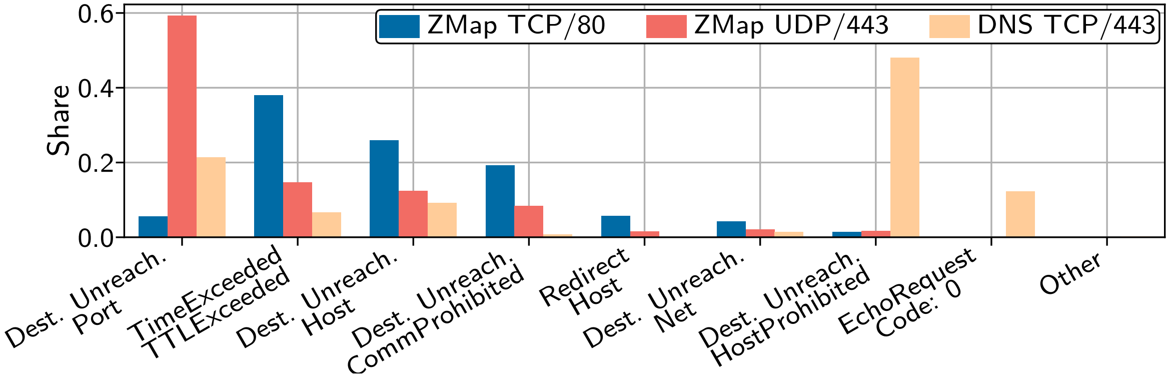

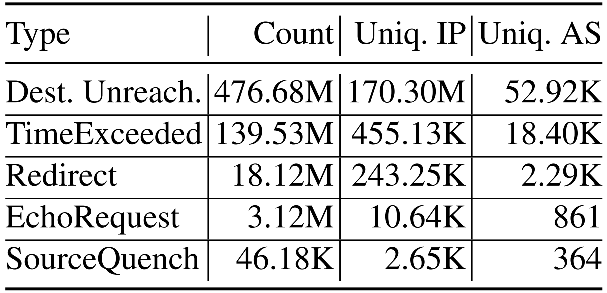

Type	Count	Uniq. IP	Uniq. AS
Dest. Unreach.	476.68M	170.30M	52.92K
TimeExceeded	139.53M	455.13K	18.40K
Redirect	18.12M	243.25K	2.29K
EchoRequest	3.12M	10.64K	861
SourceQuench	46.18K	2.65K	364

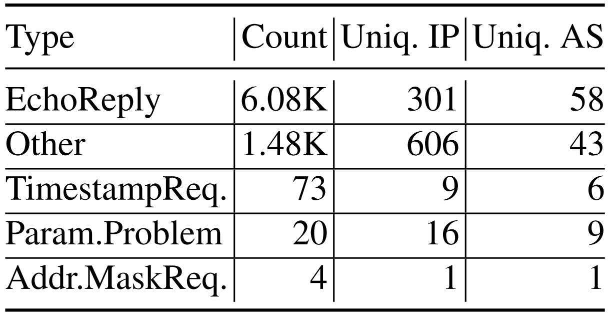

Type	Count	Uniq. IP	Uniq. AS
EchoReply	6.08K	301	58
Other	1.48K	606	43
TimestampReq.	73	9	6
Param.Problem	20	16	9
Addr.MaskReq.	4	1	1

- **ICMP replies not uniform wrt. Protocol/Port**
- **ICMP port unreachable for TCP**

ICMP Redirects

- **Wait, we should not get these: Redirects**
  - Used to signal a better path if (RFC1812 (from 1995 ☺))
    - The packet is being forwarded out the same physical interface that it was received from,
    - **The IP source address in the packet is on the same logical IP (sub)network as the next-hop IP address**, and
    - The packet does not contain an IP source route option
- **18.12M redirects**
  - 105.78K network redirects (RFC1812: *MUST NOT* send)
    - 238 different ASes affecting nearly 19k different destinations (20 have A-record in our DNS data)
  - 18.01M host redirects
    - 2.20K ASes affecting ~400k destinations (900 have A-record in our DNS data)
  - 2.7K unique redirects to private address space

ICMP Source Quench

- **Source Quench (SQ): ECN's grandparent**
  - Sent by router when congested → sender should reduce rate
  - Research: Is unfair and blind throughput-reduction attacks possible
  - IETF: don't do it (1995) and ignore it (2012)!
  - Most OSes ignore it since 2005

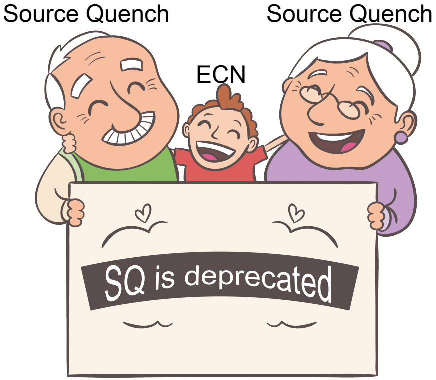

- **2.65K unique IPs located in 364 ASes issue SQ messages**
  - Very few SQs not from the destination AS
  - 53 IPs found in A-records of our DNS data subject to SQ-generation

- **Most network hardware vendors have removed SQ**
  - Between 2000 – 2010
  - **It takes decades to remove features from the Internet!**

Unreachable: Time Exceeded

- **Fragment reassembly time exceeded on IP fragmentation (7.31K)**
  - How large are our probes?
    - QUIC probes ~1300 byte: could trigger fragmentation
      - Do we set the DF-bit? ZMap by default does not
  - 26.66K *fragmentation needed and DF set* messages

- **TTL exceeded when path too long (139.52M)**
  - Quoted when dropped: 97% TTL=1, 2.4% TTL=0, and everything else, MPLS?
  - What TTL do we set?
    - ZMap: 255 hops
    - Linux Stack: 64 hops

Routing Loops

- **We performed**
  - ~27M traceroutes to
  - ~612K different /24 subnets from
  - ~28K ASes

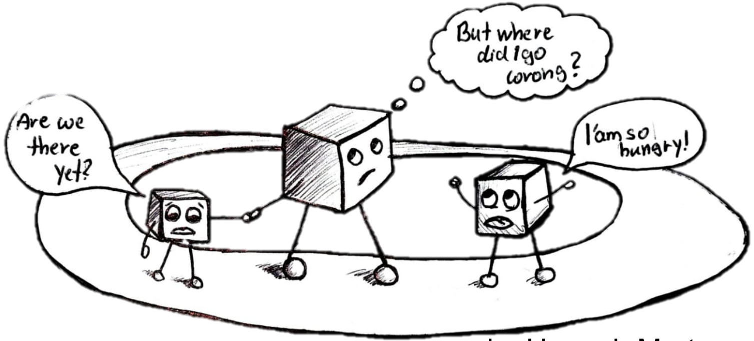

by Hannah Mertens

- **439K subnets from 19.8K ASes are unreachable due to a loop**
  - 167K different loops in 13.9K ASes
  - 136K have IPs for all routers involved in the loop
    - 13% (17.7K) already cover all different ASes paths involved
    - 4.8K cross AS boundaries

Routing Loops

- **Are the loops persistent?**
  - Compare traceroutes two weeks apart
  - Loops from roughly 150 ASes disappear
  - Still: 404K subnets unreachable

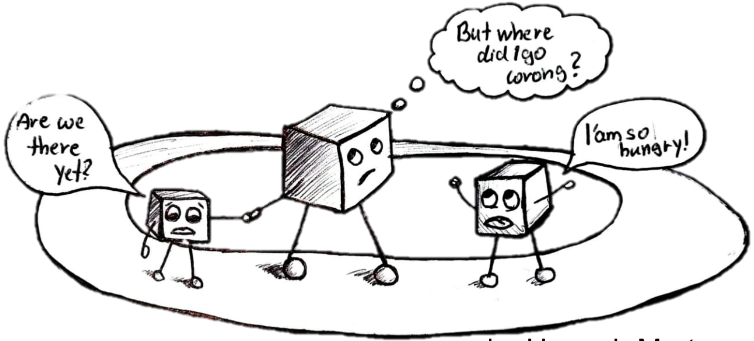

by Hannah Mertens

- **We found loops at our upstream ISP (German Research Network)**
  - We contacted them
  - They confirmed the loops
  - They fixed the loops
  - Root cause
    - Manually configured static routes at one router (R1) towards R2
    - R2 no idea how to forward, forwards to default (R1), …

Conclusion

- **The Internet is full of deprecation and badly configured systems!**
  - More odd things in the paper: https://arxiv.org/abs/1901.07265
- **There seem to be lots of routing loops**
  - Better mapping to interdomain loops desirable
- **We provide an evolving dataset**
  - If you need, we can provide live stream access to the data, contact me ☺
- **https://icmp.netray.io**

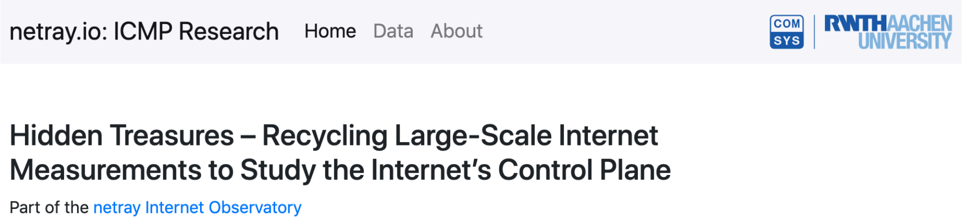

Hidden Treasures – Recycling Large-Scale Internet Measurements to Study the Internet's Control Plane

Part of the netray Internet Observatory

THANK YOU

IP Quotation

- **Quoted IP packets: D. Malone and M. Luckie. Analysis of ICMP Quotations. In *PAM*, 2007.**
  - Most quoters (87.60%) quote 28 bytes, the minimum in RFC 792
  - Some quoters (8.60%) quote 40 bytes

- **Our data (2018)**
  - 180.25M unique source IP/payload length combinations (generating the quote)
  - 76% are longer than 40 bytes
  - 24% are exactly 28 byte long
  - 1.06M destination addresses (in the quote) are in reserved address space
    - E.g., generated behind NATs

Unreachable Hosts

- **Unreachability largest fraction of ICMP messages**
  - How persistent?
    - Host and Network
  - Compare Thu to Fri
    - Both (UDP/443)
  - And Thu to Thu + 1 week

Type	Code	Count
Dest. Unreach.	Port	256.72M
TimeExceeded	TTLExceeded	139.52M
Dest. Unreach.	Host	107.15M
	CommProhibited	71.70M
	HostProhibited	23.07M
	Net	17.94M
	Protocol	51.04K

Type	Code	Count
Dest. Unreach.	Frag.Needed	26.66K
	NetProhibited	26.28K
TimeExceeded	Frag.Reassembly	7.31K
Dest. Unreach.	HostUnknown	336
	NetTOS	25
	NetUnknown	6
	SourceIsolated	2

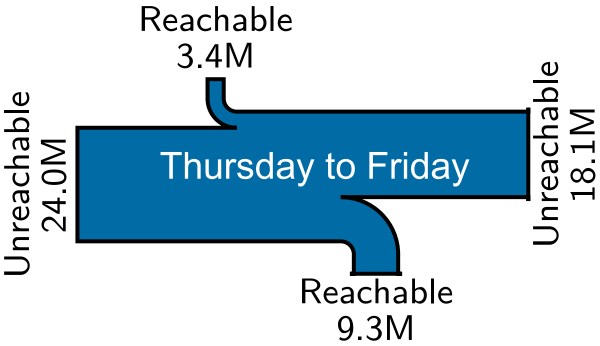

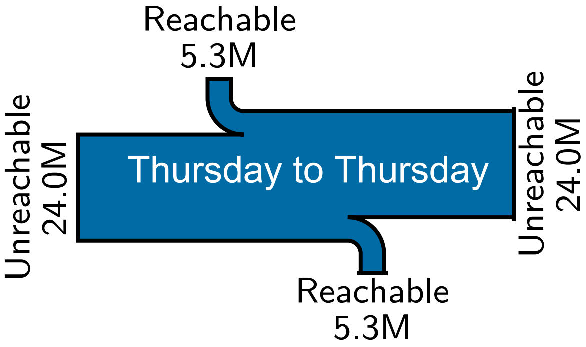

ICMP Echos?!

- **What we expected: Echo Requests**
  - Our infrastructure is regularly hit by pings
  - 10.57K unique IPs out of 840 ASs
  - IDSs?

- **What we did not expect: Echo Replies**
  - We do not generate ICMP! These replies flow towards us!
  - All directed towards our DNS resolvers
  - Contain quoted IP+UDP+DNS query response packets destined to us
  - Source IP: active DNS servers
    - When manually doing a lookup, no ICMP but two different DNS responses
    - IP stacks differ significantly → DNS Spoofer?