Open Resolver Project revisited

Jared Mauch
jared@puck.nether.net

Quick recap

- Open resolvers bad
  - Hopefully we all remember this
- Been scanning for them since 2013-Mar-24
  - Scans happen Sunday mornings at 0 UTC
  - NTT continues to provide me machine/VM space despite employer change
- Scan Data is available
  - Website has been broken until Nov 1 2019
- Sometimes scans didn’t work
  - temporarily broken VM (Usually my fault)
  - hopefully obvious in graphs

Quickly, what you care about

- Trends are mostly the right direction

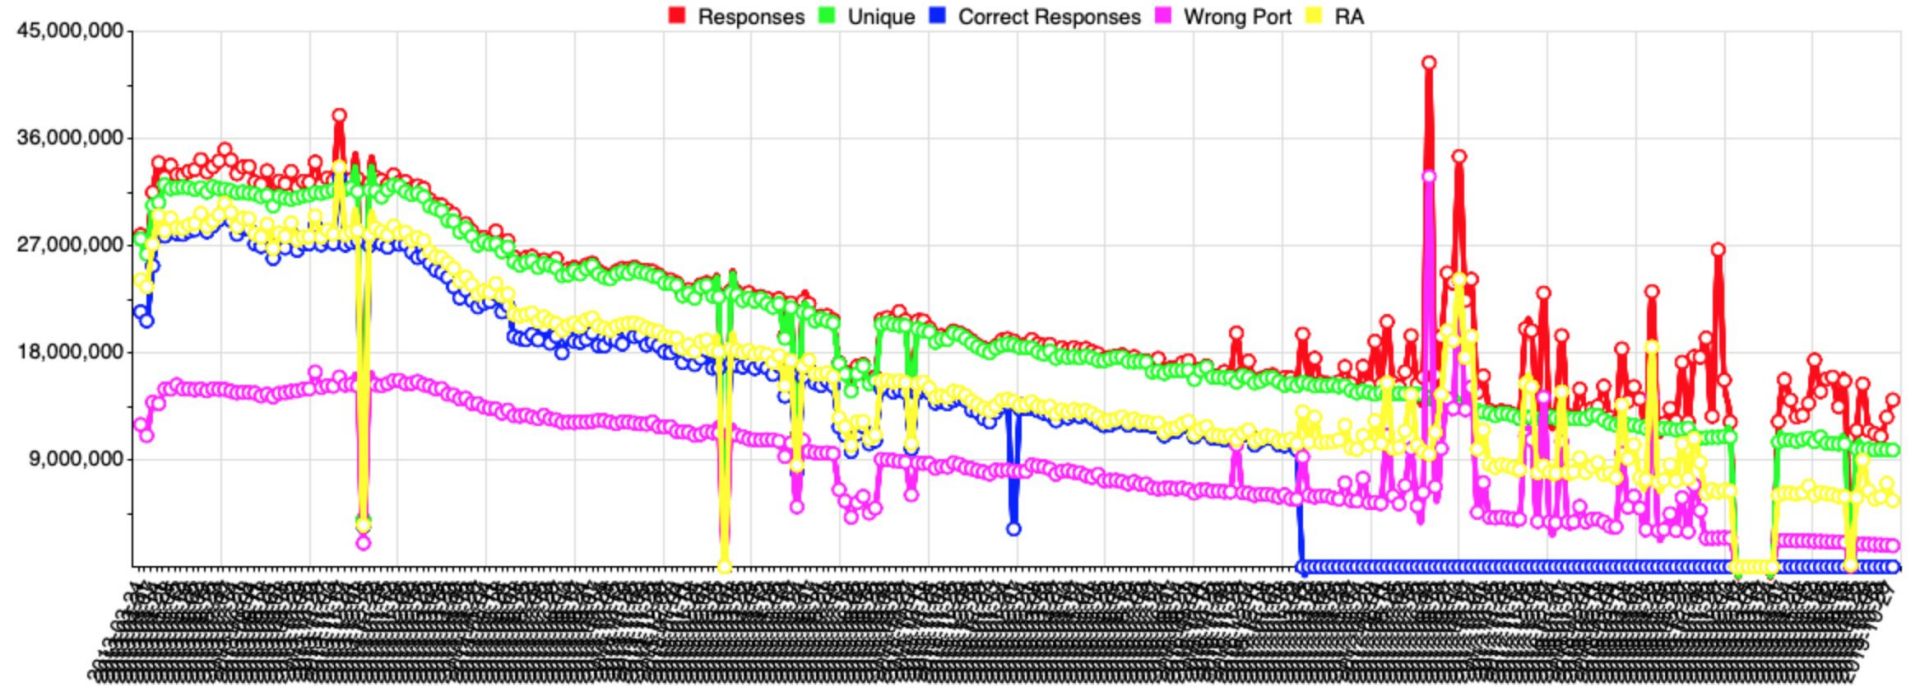

Detailed breakdown

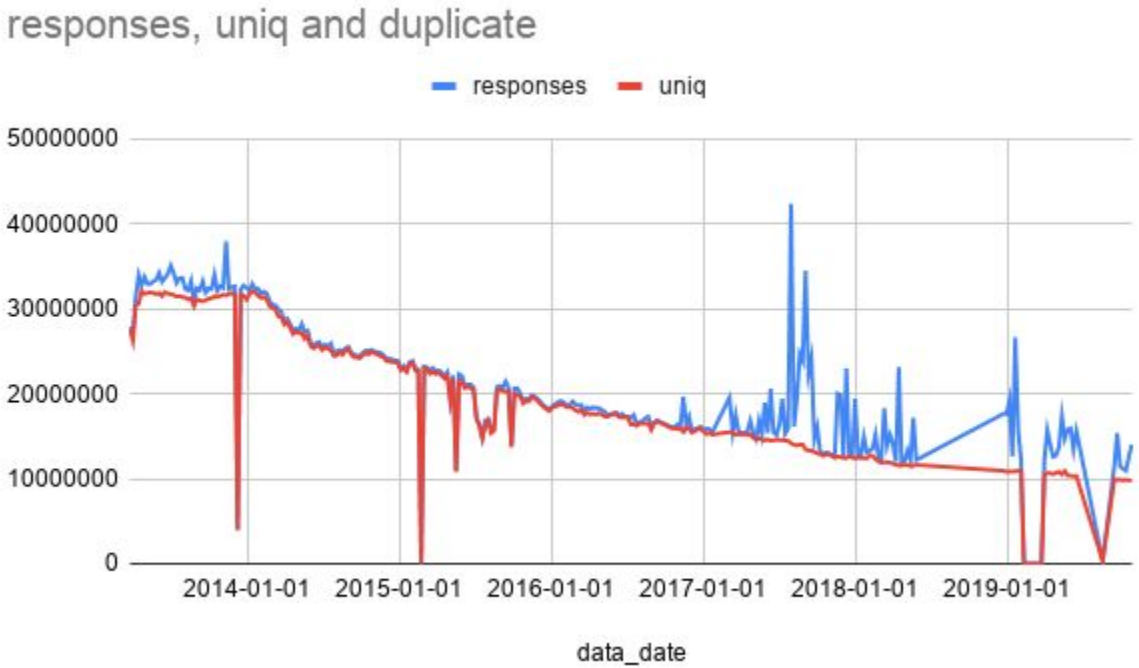


unique
40000000
30000000
20000000
10000000
0
2014-01-01
2015-01-01
2016-01-01
2017-01-01
2018-01-01
2019-01-01

Alternate IP responded

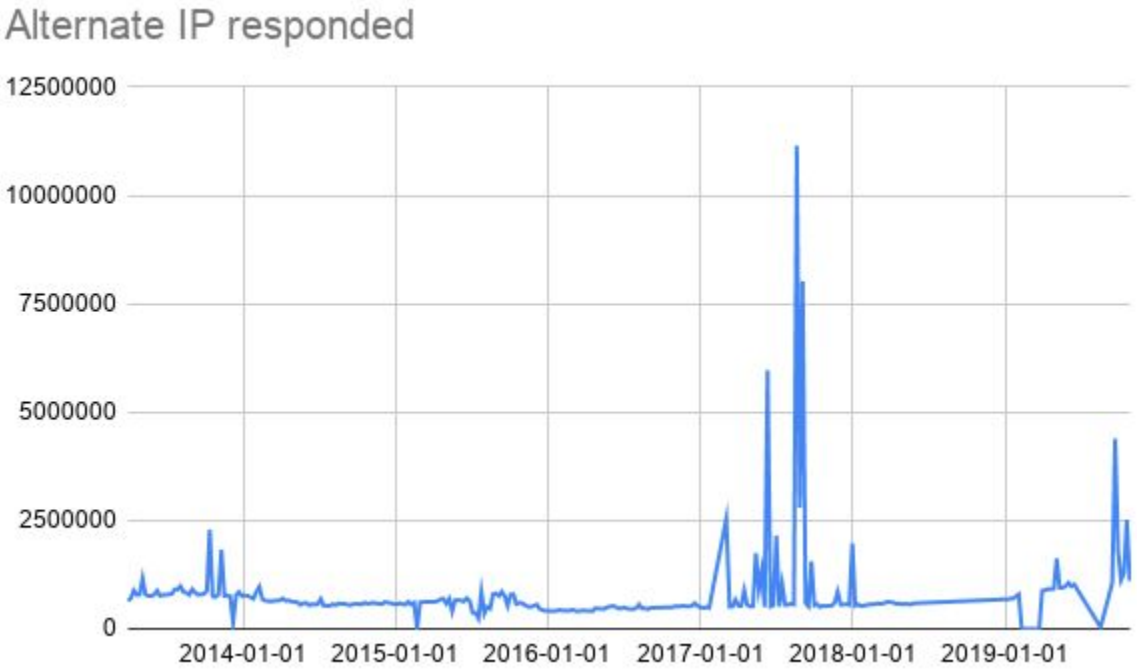

Non-port 53 reminder

- Many CPE are Linux based
- Use iptables
- Install udp/53 forwarder rule to their configured DHCP service
- Not constrained to their inside interface
- Packet flow:
  - Scanning Host
  - CPE
  - PAT (copies IP source to destination, skips NAT rule as it’s not “inside” interface)
  - Forwards to DNS server
  - Reply to scanning host

DNS Reply but via non-port 53

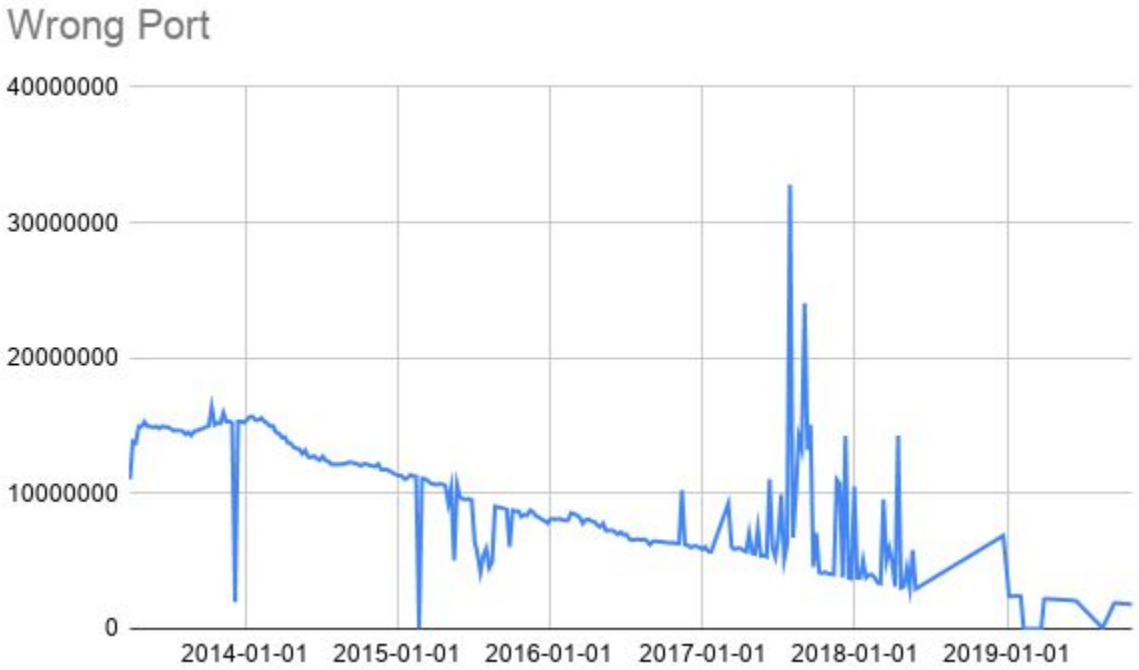

RA=1

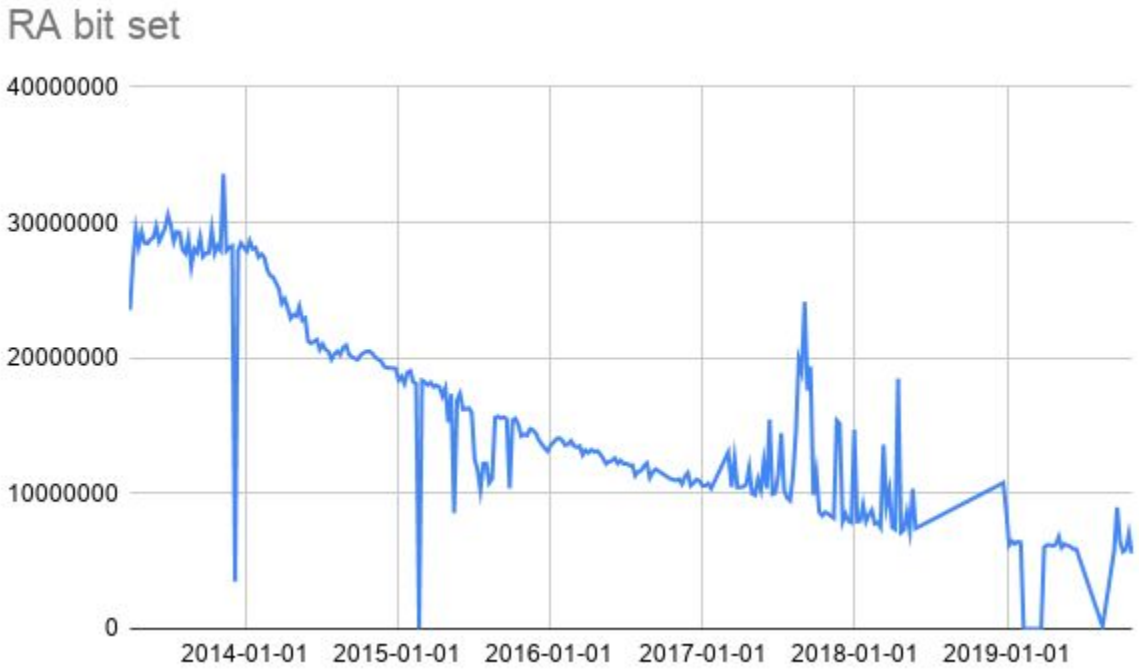
RA bit set
40000000
30000000
20000000
10000000
0
2014-01-01
2015-01-01
2016-01-01
2017-01-01
2018-01-01
2019-01-01

REFUSED


REFUSED
5000000
4000000
3000000
2000000
1000000
0
2014-01-01
2015-01-01
2016-01-01
2017-01-01
2018-01-01
2019-01-01

Observations - revisited

- Overall trend seems to be flattening out
- Still very much a reduction to 9.8M unique IPs from 27.5M
- Some data even shows increases
- Alternate IP responding increased to 1.1M from 790k
- non-Port 53 responses (dysfunctional NAT) increased to 1.7M from 1.5M

More data available upon request

- I log the timestamp, replying port and full packet response
- Much easier to share some parsed data
- I need to go back and reprocess some files as the website change seems to have broken the “correct response” check

Quick lessons learned

- Make sure you use source control even for your “quick scan hack”
- Your original parsers may not scale well over time
  - I rewrote the statistics processor in python vs perl
  - Lower memory footprint, runs slightly faster
- Beware database upgrades
  - Postgresql version upgrades have caused issues a few times
- Ensure data rotation isn’t a manual task
  - Website outages mostly due to limited disk space on web server
- Still get regular requests for data from PhD students
  - I need to be better at responding to them and collecting derivative papers

Questions?

jared@puck.nether.net