An internet-wide scan of root-hints

Roy Arends
Pricipal Research Scientist

ICANN tech-day

13 June 2022

What are root hints

- The names and IP addresses of the authoritative name servers for the root zone, so the software can bootstrap the DNS resolution process.
- For many pieces of software, this list comes built into the software. This file is often used in priming.
- https://www.internic.net/domain/named.root

What is this scan?

- A scan of the IPv4 address space for SOA records for the root zone
- What is returned in the authority and additional section may look like root-hints such as NS and A/AAAA records, but these are the result of the priming process, and not necessarily root-hints.
- We’re going to ignore these NS/A/AAAA records and start by observing the SOA record.
- As such, we’re trying to understand what resolvers use as root servers.

Expectations

- Since the 3rd of October 2001, the SOA RNAME for the IANA root zone is “nstld.verisign-grs.com”
  - The RNAME is the domain name representing the administrator’s mailbox (email address)
- We expect that the bulk of resolvers use IANA root-servers (a..m.root-servers.net) and will return an SOA with an RNAME containing “nstld.verisign-grs.com”
- The version of the root zone is encoded in the SOA serial number. The convention is that a date and a daily version number is used.
- We expect that the bulk of resolvers that use IANA root-servers have a SOA serial number that is at most a day behind the current version.

Setup

- We send a DNS message send with the following characteristics:
- It is a request (QR=0),
- For a standard query (OPCODE=QUERY)
- with a single question
  - (qname: empty label (aka root label), qtype:SOA, qclass=IN)
- No Extended DNS
- All header bits 0
- the 16 bit identifier in the DNS message is 0

Selecting targets

- Naïve approach is to send queries to the entire IPv4 space
  - Minus the multicast, experimental, RFC1918 space, etc.
- Better: select routable address ranges from a route-view.
  - From the Oregon route-views archive.
  - Minus DNS-OARC's don't probe list.
- This is about 80% of the entire IPv4 address space.
- We use zmap to send queries. It uses an allow list (our targets) and a block list (don't probe list) and a hexadecimal string to represent the DNS query
  - Zmap's DNS module contains a bug, report has been sent.

Results

- Statistics
  - 3,445,927,936 (3.4B) queries sent.
  - 10,140,034 (10M) responses received.
  - Response rate of approximately 0.3%.
- About 3,198,067 responses had the wrong identifier (not 0).
- An additional 566,322 were duplicates.
- There are a variety of reasons we have received these.
  - Mainly hosts forwarding a message or bouncing a packet to a resolver.
- 6,375,645 responses remaining

Results

RCODE	VOLUME
REFUSED	3716978
NOERROR	2276319
SERVFAIL	313343
NXDOMAIN	40235
syntax_error	22902
FORMERR	4246
NOTIMP	1388
NOTAUTH	223
NOTZONE	7
NXRRSET	3
YXRRSET	1

Results

- The 2,276,319 NOERROR responses are useful for this survey
- Found an interesting “bug”:
  - 875 responses did not have the QR bit set.
  - QR=0 implies request, not a response.
  - Time to dust off an old IETF draft
    - QR clarify was a simple IETF draft that indicated that a response to a request MUST have the QR bit set
    - And a server or resolver MUST not respond to a response.
    - Ignoring these rules may lead to a DDoS attack using infinite loops.

Results

- About 45,153 responses had the TC bit set
  - indicates that the response was truncated
  - more likely a simple denial of service mitigation technique.
- About 623,441 responses had the RD bit set
  - We never set the RD bit as we don’t want recursion.
- We’re going to ignore these message for now, as they contain no additional information or have caused additional recursion.
- 1,237,020 responses contained an SOA record

Results

RNAME	Volume	Percentage
nstld.verisign-grs.com	1,147,566	92.8
*.hostgator.com	38,145	3.1
hostmaster	20,612	1.7
*.bluehost.com	13,500	1.1
root	3,209	0.3
*.hostgator.in	2,057	0.2
*.hostgator.com.br	1,895	0.1
support.localhost	886	0.1
*.webhostbox.net	813	0.1
*.ehosts.com	581	0.05

Results

- RNAME: nstld.verisign-grs.com MNAME distribution

MNAME	Volume
a.root-servers.net.	1,147,559
n.root-servers.net	4
cache1-main.mtl1.rogerstelecom.net.	2
v.root-servers.net.	1

- The 7 without a.root-servers.net are statistically insignificant
  - Cache1-main contains an old root-zone (2012033001) without apex DNSSEC records and new NS records
  - V.root-servers.net: old as well, resigned with different DNSKEYS
  - N.root-servers.net is a private root-zone

Results

- SOA Serial distribution

SOA SERIAL	Volume
Recent	1,147,139
Other	406
2019103000	1
2018041600	1
2017092880	2
2016091200	1
2016061401	2
2016053100	2
2012061900	1
2012041813	2
2012041809	1
2012021400	1

Conclusion

- A fair amount of brokenness, which was expected.
  - The QR-clarify draft will be brought back to life
  - Stale configurations leads to old versions of root-zones
  - Large amount of consumer routers forward or bounce DNS requests
- Large amount of DNS hosting providers use private root hints.
- Most implementations use the IANA root hints.
- No diaspora of intentional alternative root hints.
  - There are some alternative root hints, but they are limited to some authoritative servers.
  - No indication that they are used by a significant set of resolvers.

Questions?

Thank You and Questions

Visit us at **icann.org**

 @icann

 facebook.com/icannorg

 youtube.com/icannnews

 flickr.com/icann

 linkedin/company/icann

 slideshare/icannpresentations

 soundcloud/icann