

The Root Server System

Johan Ihrén
`i.root-servers.net`

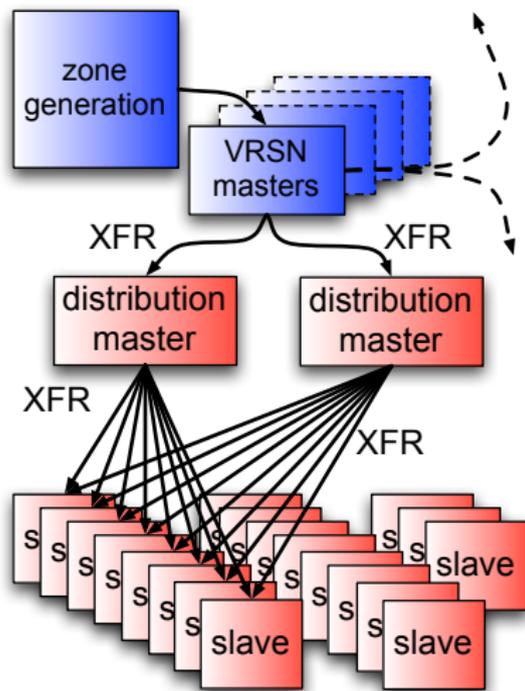
March 15, 2011

What Is A Root Name Server?

- ▶ A very simple server with a rather important task
 - ▶ To respond to queries from the entire Internet with the referral information needed for the client to locate the “final answer”
 - ▶ Root name servers **very rarely** provide “final answers”, the vast bulk of the responses are referrals
- ▶ The interface to a root name server is a “signed zone transfer” where we get the complete database from which to provide answers (i.e. referrals)
 - ▶ Because it is signed, we know that the database is complete and unchanged, i.e. exactly the data that we **should** publish
- ▶ Root name servers do not participate in defining the content of the root zone (that is IANAs and Verisigns task), they **only** deal with publication of the zone they receive

Zone Distribution

- ▶ The zone is transferred from Verisign to the root nameservers via a “zone transfer”
- ▶ Because of anycast and many root nameservers consisting of a number of sites the zone transfer hierarchy is often several layers deep
- ▶ There are different alternatives to increase redundancy and the operators do not coordinate but instead choose different strategies
 - ▶ to maximize diversity



Where are Root Name Servers Located?

Basically everywhere.



Where are Root Name Servers Located?

Basically everywhere.

- ▶ There are multiple instances on all continents except Antarctica
- ▶ Before anycast it made sense to try to identify “optimal” locations. But with anycast that no longer matter in practice
- ▶ We try to deploy wherever it makes sense and we find a solution
 - ▶ When or if there is a request from Antarctica we will certainly try to figure out how to solve that problem
- ▶ There is rarely any coordination between the different root operators on deployment issues
 - ▶ The reason is our usual mantra of “diversity”, i.e. we want to avoid locking in to a shared view of exactly where it makes sense to deploy

Queries to One Root Name Server (a 30ms snapshot)

```
02:19:00.547 query: 3.36.3.10.in-addr.arpa IN PTR -
02:19:00.548 query: lanark.WIROOS.COM IN A -EDC
02:19:00.548 query: lanark.WIROOS.COM IN AAAA -EDC
02:19:00.549 query: bu1690.influ8115kz.1x5m2750.jf IN A -EDC
02:19:00.549 query: ns1.telkom.net.id IN AAAA -EDC
02:19:00.550 query: kiu.wiroos.com.ar IN AAAA -EDC
02:19:00.552 query: ForestDnsZones.sp.local IN SOA -ED
02:19:00.553 query: 131.76.34.174.in-addr.arpa IN PTR -EDC
02:19:00.553 query: 35.69.17.209.in-addr.arpa IN PTR -EDC
02:19:00.553 query: 23.54.17.193.in-addr.arpa IN PTR -EDC
02:19:00.557 query: dizell.home IN SOA -ED
02:19:00.557 query: b._dns-sd._udp.RT-G32 IN PTR -ED
02:19:00.559 query: florik.ru IN A -ED
02:19:00.561 query: _641_34_7 IN A -ED
02:19:00.561 query: ms.seznam.cz IN A -EC
02:19:00.561 query: ns.seznam.cz IN A -EC
02:19:00.565 query: _ldap._tcp.Default-First-Site-Name._sites.gc._msdcs.mobilt...
02:19:00.566 query: _731_18_2 IN A -ED
02:19:00.566 query: l.root-servers.net IN A -
02:19:00.567 query: ns2.iasi.rdsnet.ro IN A -
02:19:00.568 query: NSFEXCHANGE.nsfad.local IN A -ED
02:19:00.573 query: ns.sdjnptt.net.cn IN A -EDC
02:19:00.573 query: caseydesign.com IN MX -EDC
02:19:00.573 query: ns1.sdqdppt.net.cn IN A -EDC
02:19:00.574 query: yrs-cassa19 IN SOA -EDC
02:19:00.576 query: e8h6rj:67.71yt IN A -ED
```

Queries to One Root Name Server (Junk More junk's snapshot)

```

02:19:00.547 query: 3.36.3.10.in-addr.arpa IN PTR -EDC
02:19:00.548 query: lanark.WIROOS.COM IN A -EDC
02:19:00.548 query: lanark.WIROOS.COM IN AAAA -EDC
02:19:00.549 query: bu1690.inf1u8115kz.1x5m2750.jf IN A -EDC
02:19:00.549 query: ns1.telkom.net.id IN AAAA -EDC
02:19:00.550 query: kiu.wiroos.com.ar IN AAAA -EDC
02:19:00.552 query: ForestDnsZones.sp.local IN SOA -ED
02:19:00.553 query: 131.76.34.174.in-addr.arpa IN PTR -EDC
02:19:00.553 query: 35.69.17.209.in-addr.arpa IN PTR -EDC
02:19:00.553 query: 23.54.17.193.in-addr.arpa IN PTR -EDC
02:19:00.557 query: dizell.home IN SOA -ED
02:19:00.557 query: b._dns-sd._udp.RT-G32 IN PTR -ED
02:19:00.559 query: florik.ru IN A -ED
02:19:00.561 query: _641_34_7 IN A -ED
02:19:00.561 query: ms.seznam.cz IN A -EC
02:19:00.561 query: ns.seznam.cz IN A -EC
02:19:00.565 query: _ldap._tcp.Default-First-Site-Name._sites.gc._msdcs.mobilt...
02:19:00.566 query: _731_18_2 IN A -ED
02:19:00.566 query: l.root-servers.net IN A -
02:19:00.567 query: ns2.iasi.rdsnet.ro IN A -
02:19:00.568 query: NSFEXCHANGE.nsfad.local IN A -ED
02:19:00.573 query: ns.sdnptt.net.cn IN A -EDC
02:19:00.573 query: caseydesign.com IN MX -EDC
02:19:00.573 query: ns1.sdqptt.net.cn IN A -EDC
02:19:00.574 query: yrs-cassa19 IN SOA -EDC
02:19:00.576 query: e8h6rj:67.71yt IN A -ED

```

Queries to One Root Name Server (as snapshot)

```

02:19:00.547 query: 3.36.3.10.in-addr.arpa IN PTR -EDC
02:19:00.548 query: lanark.WIROOS.COM IN A -EDC
02:19:00.548 query: lanark.WIROOS.COM IN AAAA -EDC
02:19:00.549 query: bu1690.inf1u8115kz.1x5m2750.jf IN A -EDC
02:19:00.549 query: ns1.telkom.net.id IN AAAA -EDC
02:19:00.550 query: kiu.wiroos.com.ar IN AAAA -EDC
02:19:00.552 query: ForestDnsZones.sp.local IN SOA -ED
02:19:00.553 query: 131.76.34.174.in-addr.arpa IN PTR -EDC
02:19:00.553 query: 35.69.17.209.in-addr.arpa IN PTR -EDC
02:19:00.553 query: 23.54.17.193.in-addr.arpa IN PTR -EDC
02:19:00.557 query: dizell.home IN SOA -ED
02:19:00.557 query: b._dns-sd._udp.RT-G32 IN PTR -ED
02:19:00.559 query: florik.ru IN A -ED
02:19:00.561 query: _641_34_7 IN A -ED
02:19:00.561 query: ms.seznam.cz IN A -EC
02:19:00.561 query: ns.seznam.cz IN A -EC
02:19:00.565 query: _ldap._tcp.Default-First-Site-Name._sites.gc._msdcs.mobilt...
02:19:00.566 query: _731_18_2 IN A -ED
02:19:00.566 query: l.root-servers.net IN A -
02:19:00.567 query: ns2.iasi.rdsnet.ro IN A -
02:19:00.568 query: NSFEXCHANGE.nsfad.local IN A -ED
02:19:00.573 query: ns.sdnptt.net.cn IN A -EDC
02:19:00.573 query: caseydesign.com IN MX -EDC
02:19:00.573 query: ns1.sdqptt.net.cn IN A -EDC
02:19:00.574 query: yrs-cassa19 IN SOA -EDC
02:19:00.576 query: e8h6rj:67.71yt IN A -ED

```

You get the picture...

Even more junk

Queries to One Root Name Server, cont'd

But in spite of the huge fraction of the queries that are junk...

- ▶ ...we respond to junk queries just as willingly as to “good queries”
- ▶ To the best of our ability
- ▶ It is not our role to decide what is a query worth responding to or not (although it may seem obvious which are junk queries)
- ▶ That said, it would of course be **much appreciated** if some of the broken queries that should not exist ceased to exist
 - ▶ But we have to be realistic, and while not holding our breath waiting for that improvement, we deal with all the cruft we get

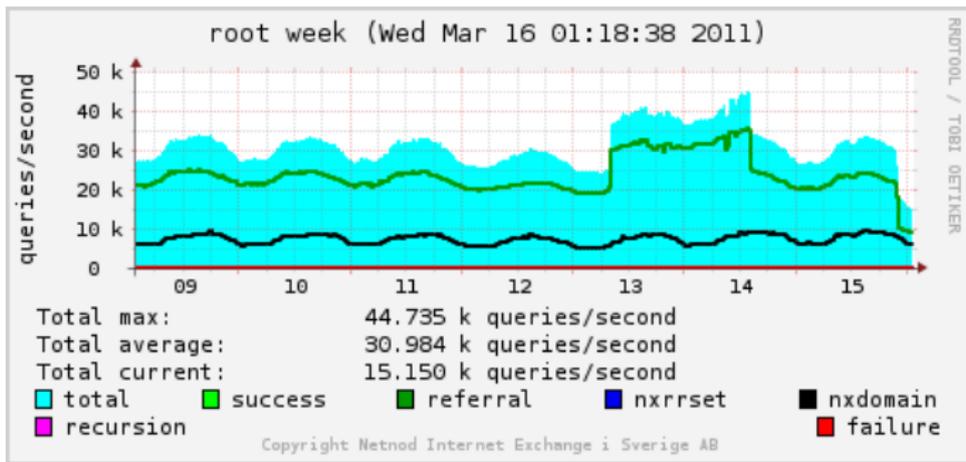
Statistics and Data Collection

- ▶ All the root operators have had their own data collection systems for a long time, but in preparation for signing the root zone last year these systems were more coordinated
 - ▶ to facilitate collection of comparable data for all the root nameservers
 - ▶ ... which has been done already in collaboration with DNS-OARC
- ▶ As an example, in the case of `i.root-servers.net`, there is dedicated hardware for packet capture and data storage on every site
 - ▶ ... and we're working on automation of fan-out of queries to this data and aggregation of results to facilitate central analysis without having to bring all the data home

Internal Monitoring

All operators monitor their servers (obviously), but we all use different systems for this

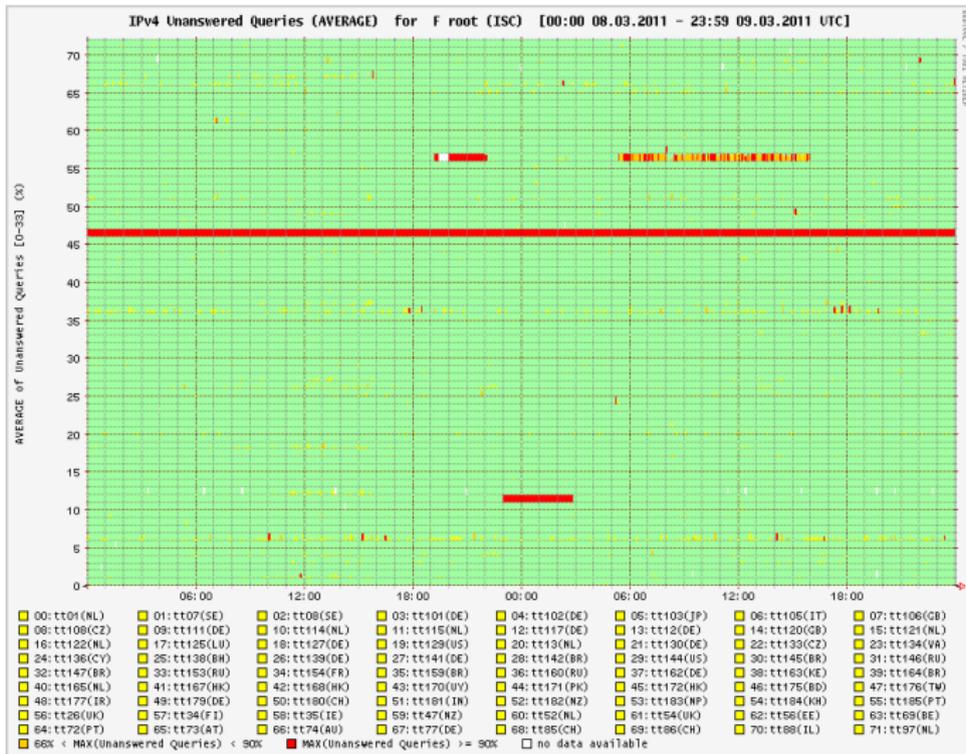
- ▶ **Internal** monitoring deals with the servers themselves and the queries they receive, etc
- ▶ This is queries per second for `i.root-servers.net` for the last week:



External Monitoring

- ▶ **External** monitoring deals with the performance of the service we provide, as observed from the outside
 - ▶ For anycast installations this is quite difficult to achieve, as you need many “probes” around the Internet to ensure that you hit all individual sites
- ▶ For this reason, in addition to some operators having their own probe networks, there is also external monitoring of our performance done by other parties, e.g. the RIPE NCC `dnsmon` service
 - ▶ We know that the world is watching and that every blip will be noticed. . .
 - ▶ Makes it very easy to stay honest

External Monitoring, cont'd



Thanks!

Questions?

`johani@autonomica.se`

(or any other root operator that you may find lurking in the corridors)