RIPE NCC
RIPE NETWORK COORDINATION CENTRE

Introduction to the Root Server System

Martin Pels | 23 June 2022 | APNIC Webinar

The DNS Hierarchy

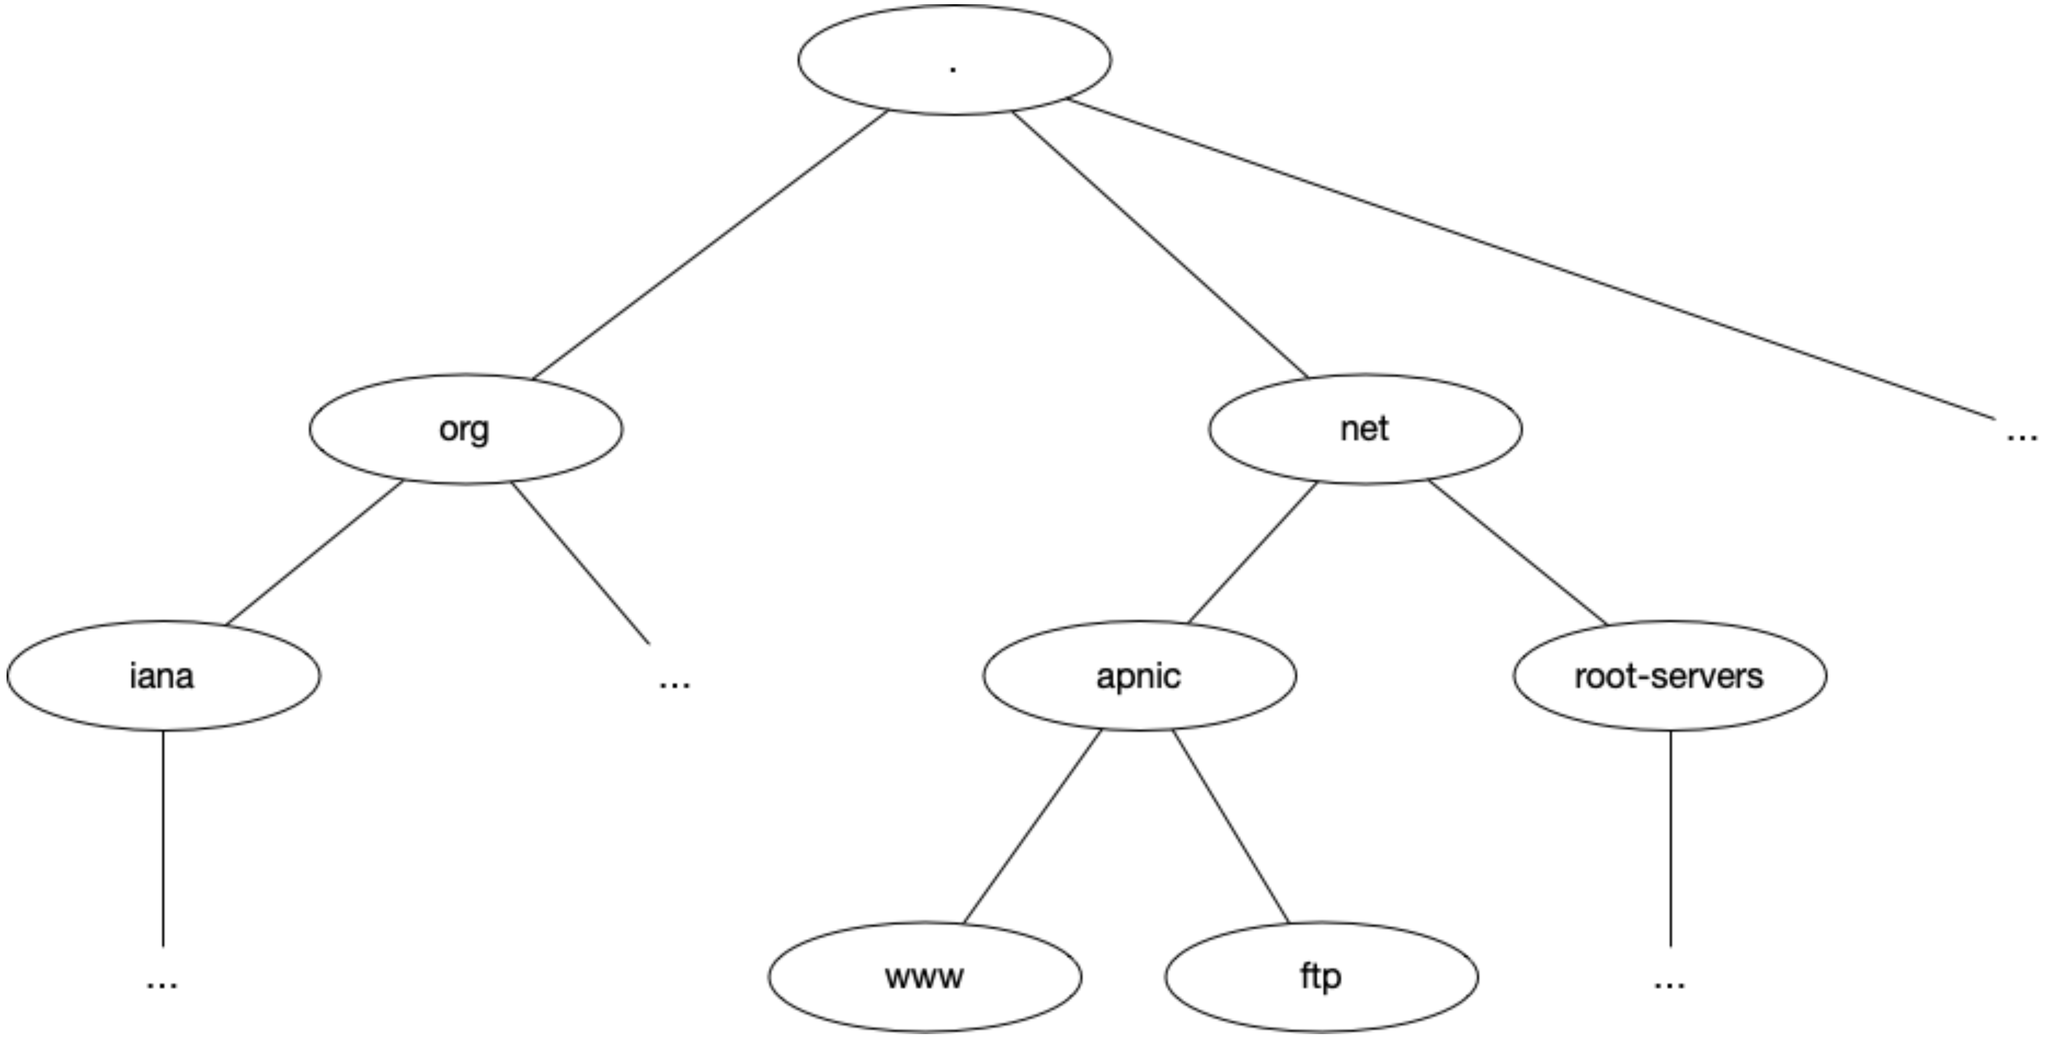

The Root Servers

- 13 letters (*[a-m].root-servers.net*)
- Serve '*.*' and '*root-servers.net.*'
- All but one also serve '*.arpa.*'
  - Via '**.ns.arpa.*', as defined in RFC9120

```
; <<>> DiG 9.11.3-1ubuntu1.17-Ubuntu <<>> . NS
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 429
;; flags: qr rd ra ad; QUERY: 1, ANSWER: 13, AUTHORITY: 0, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;.                              IN      NS

;; ANSWER SECTION:
.                       1759    IN      NS      b.root-servers.net.
.                       1759    IN      NS      c.root-servers.net.
.                       1759    IN      NS      d.root-servers.net.
.                       1759    IN      NS      e.root-servers.net.
.                       1759    IN      NS      f.root-servers.net.
.                       1759    IN      NS      g.root-servers.net.
.                       1759    IN      NS      h.root-servers.net.
.                       1759    IN      NS      i.root-servers.net.
.                       1759    IN      NS      j.root-servers.net.
.                       1759    IN      NS      k.root-servers.net.
.                       1759    IN      NS      l.root-servers.net.
.                       1759    IN      NS      m.root-servers.net.
.                       1759    IN      NS      a.root-servers.net.

;; Query time: 0 msec
;; SERVER: 185.12.64.2#53(185.12.64.2)
;; WHEN: Mon Jun 13 11:51:30 CEST 2022
;; MSG SIZE  rcvd: 239
```

Anycast

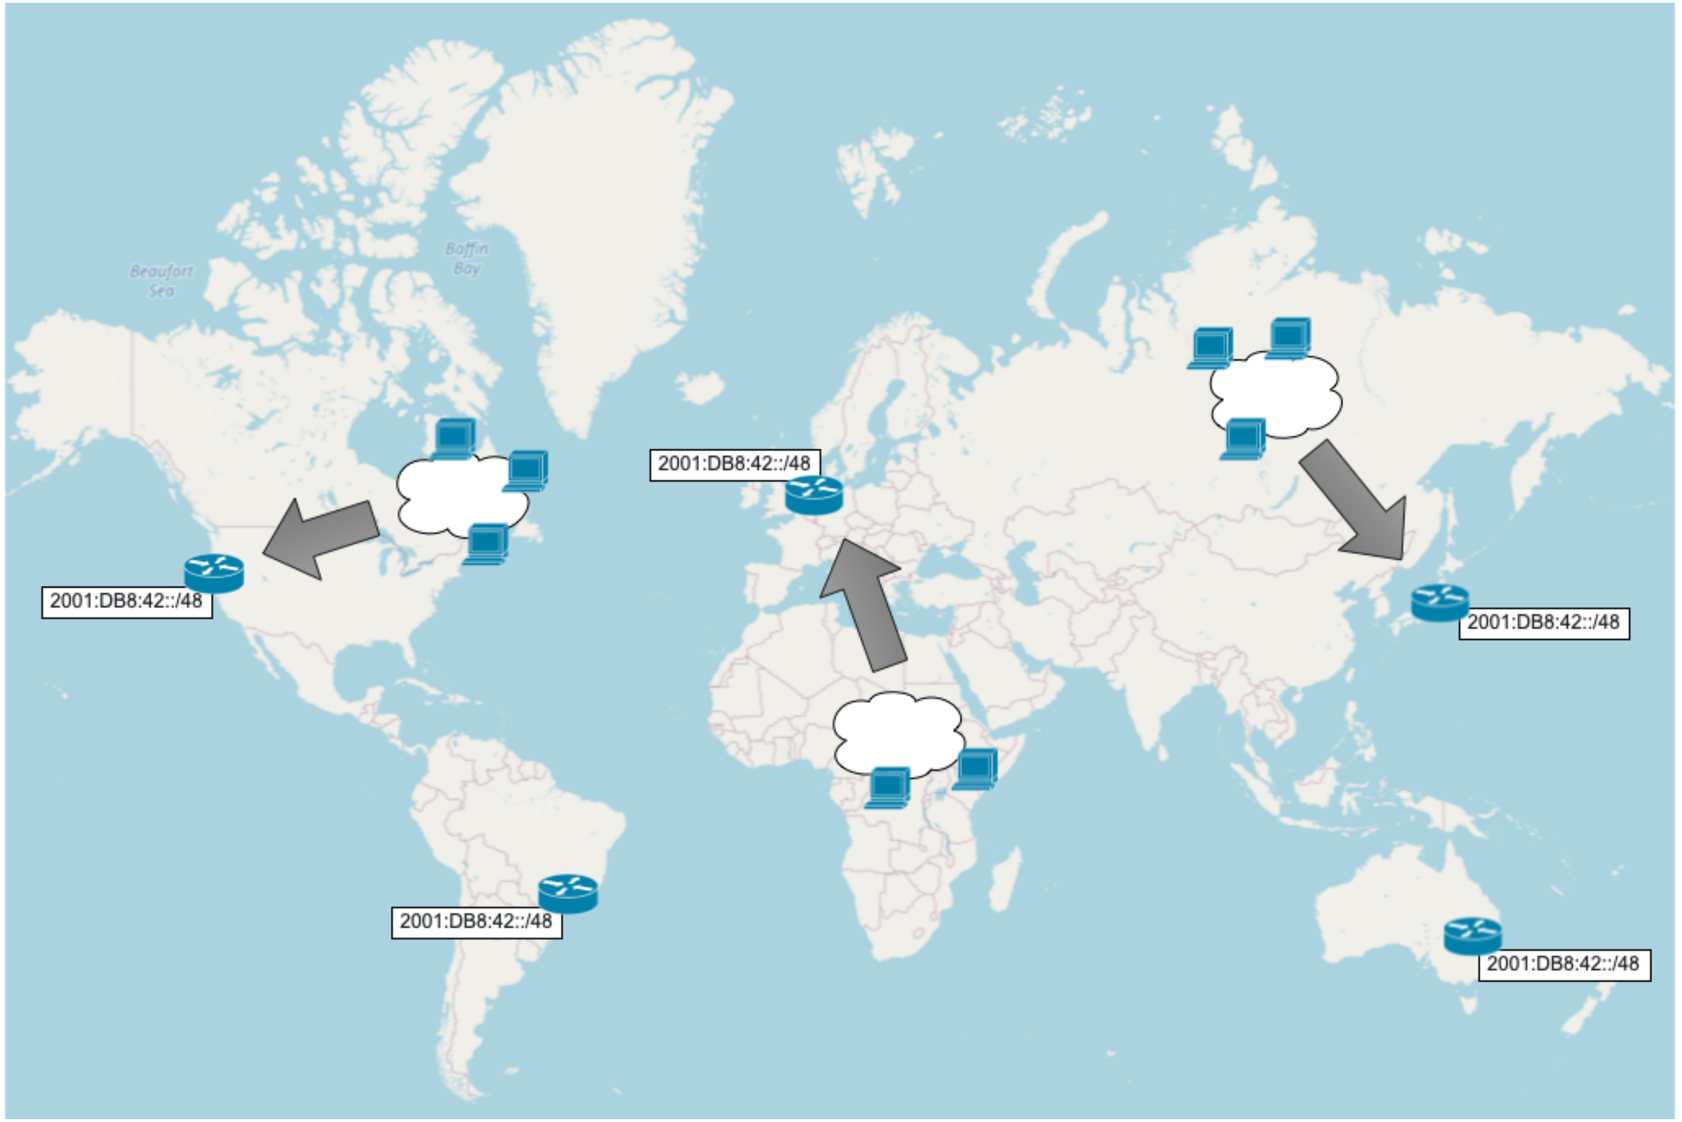
Beaufort Sea
Baffin Bay
2001:DB8:42::/48
2001:DB8:42::/48
2001:DB8:42::/48
2001:DB8:42::/48
2001:DB8:42::/48

Root Server Instance Types

- Global
  - Internet-wide visibility
  - Instance chosen based on BGP best-path
  - Goal: Increasing path diversity, resilience of RSS
- Local
  - Limited visibility
  - Constrained by NO_EXPORT or filtering policy
  - Goal: Improving performance in local area

RFC8806 - Hyperlocal Instance

- Authoritative DNS server instance
- Runs on same machine as recursive resolver
- Serves an up-to-date copy of the root zone to the recursor

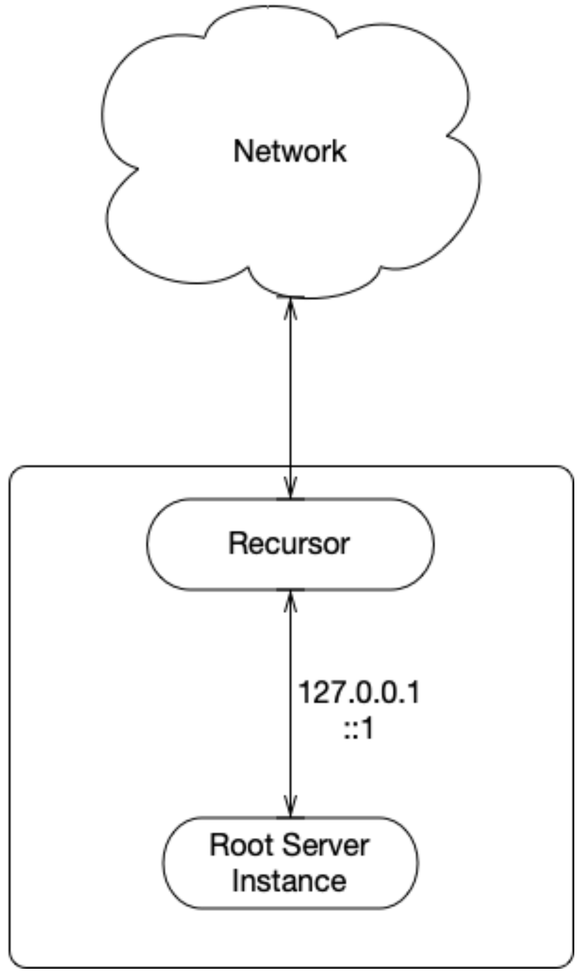

Root Server Operations

The Root Server System

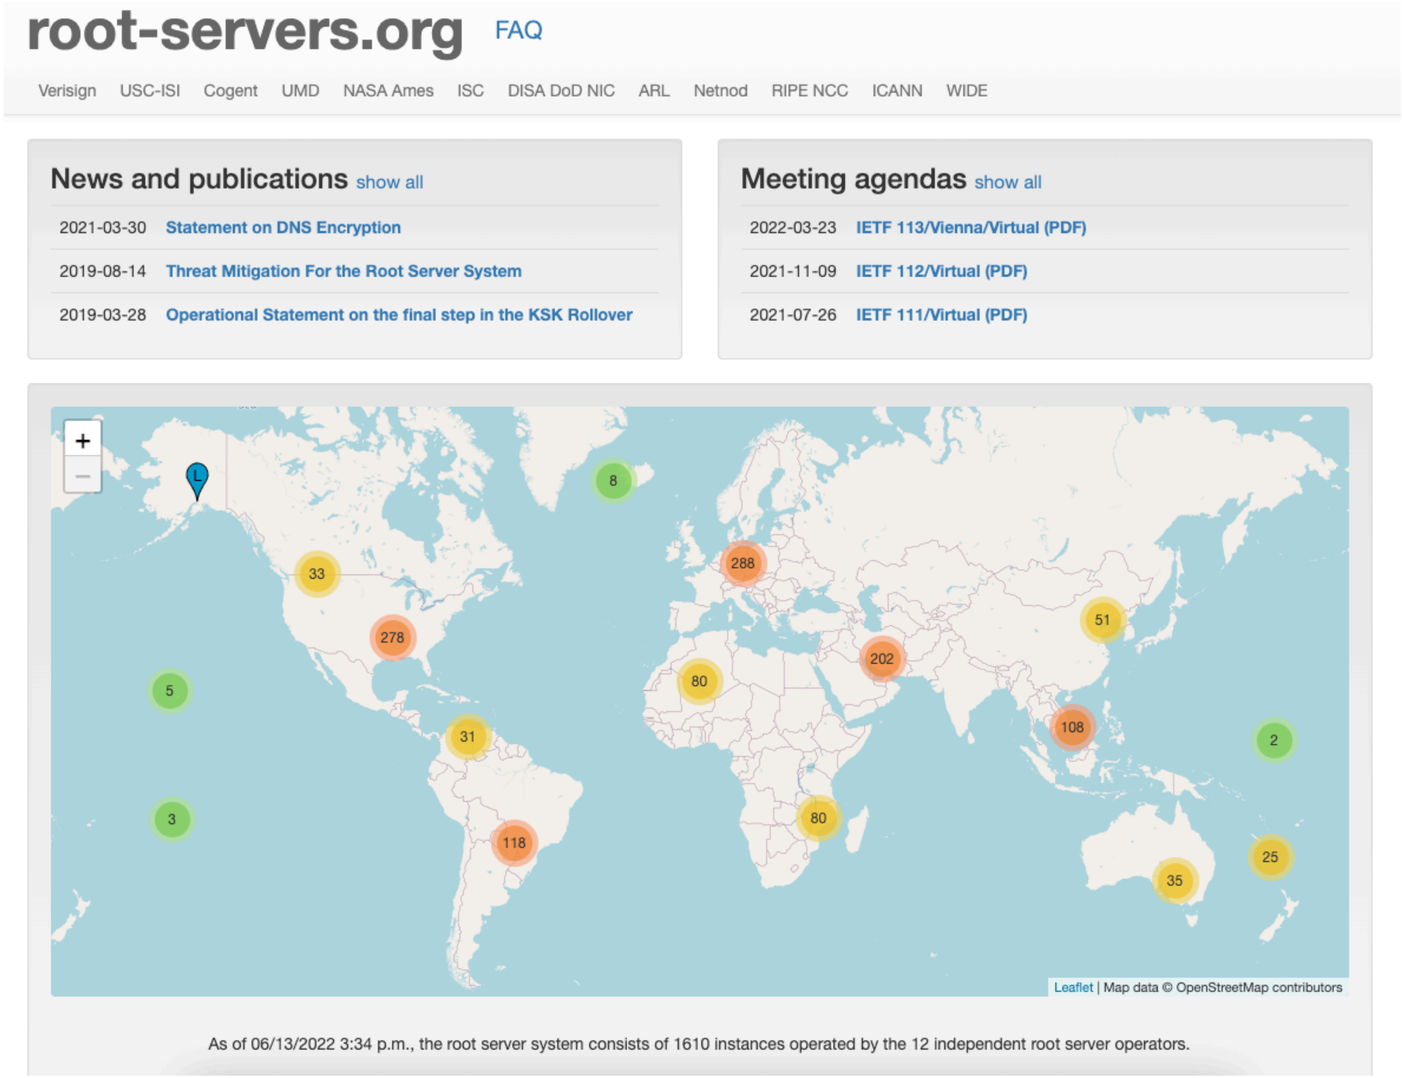

Operating Principles (RSSAC055)

- The root server system must:
  - Have a single, globally unique root, sourced by IANA
  - Be stable, reliable and resilient
  - Consist of diverse components and software
  - Evolve based on technical need, as defined by the IETF
- Root server operators must:
  - Be committed to the common good of the Internet
  - Collaborate and engage with stakeholders
  - Be transparent
  - Be autonomous and independent
  - Be neutral and impartial

Technical Requirements (RFC7720)

- Protocol requirements
  - Implement the core DNS protocol
  - Support IPv4 and IPv6
  - Support UDP and TCP
  - Support UDP checksums
  - Support DNSSEC and EDNS(0)
- Deployment requirements
  - Answer queries from any host with a valid IP address
  - Must serve the unique root zone

K-root

- Operated by the RIPE NCC
- First instance at London Internet Exchange (LINX) in 1997
- Anycast since 2003
- https://k.root-servers.org/

ASN	25152
IPv4	193.0.14.0/24
IPv6	2001:7fd::/48

K-root - Core Nodes

- Amsterdam - AMS-IX, NL-ix
- Frankfurt - DE-CIX
- London - LINX LON1/LON2, LONAP
- Miami - Equinix Miami (NOTA)
- Tokyo - JPNAP, DIX-IE

K-root - Hosted Nodes

- IXP instance
  - Directly connected to one or more IXPs
  - Peering with route servers to maximise reach
  - Global reachability
- Single instance
  - Single upstream: Only connected to sponsoring AS
  - Global or Local reachability (case-by-case)

K-root

Measuring the Root

RSSAC002

- Common set of measurement parameters
  - Number of queries and responses
  - Response type and size distribution
  - Number of sources seen
  - Latency in publishing root zone data
- Collected by all RSOs
- Published on https://rssac002.root-servers.org/

RSSAC002 Measurements

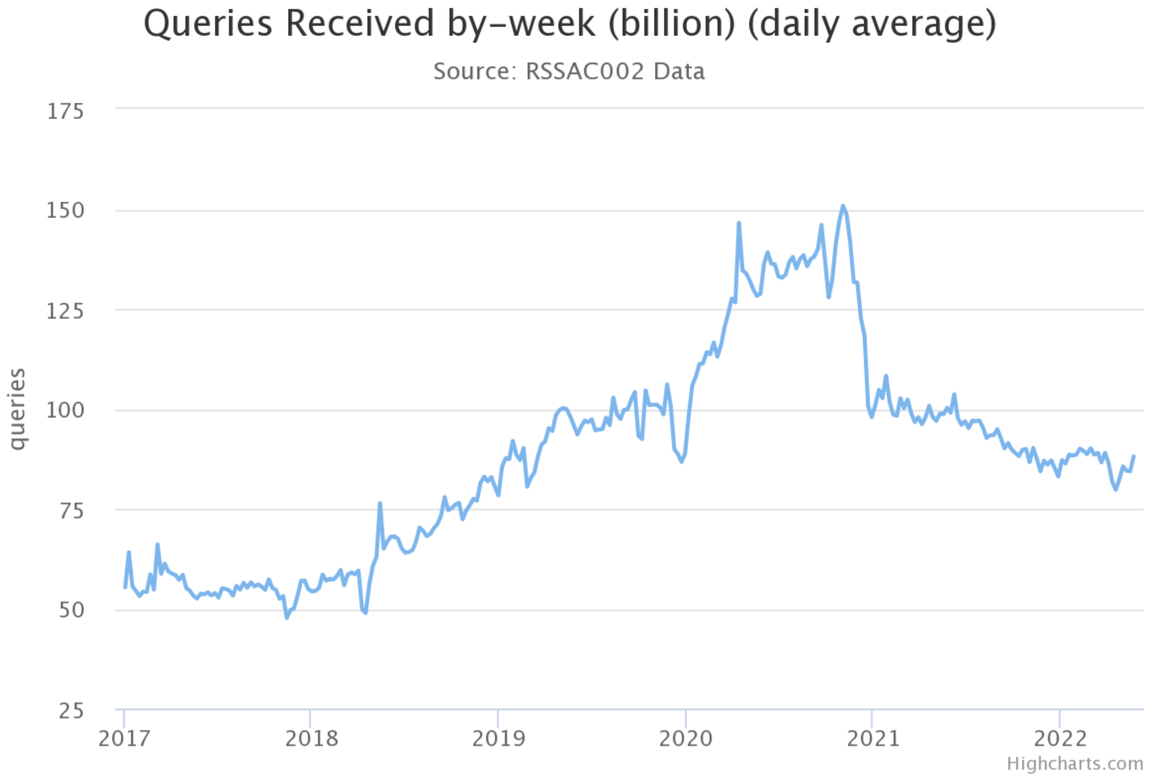

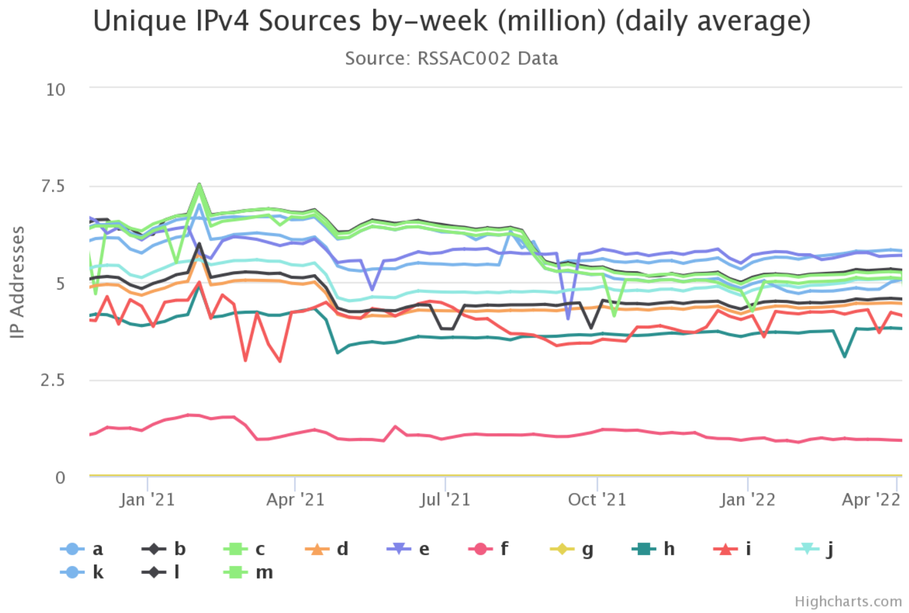

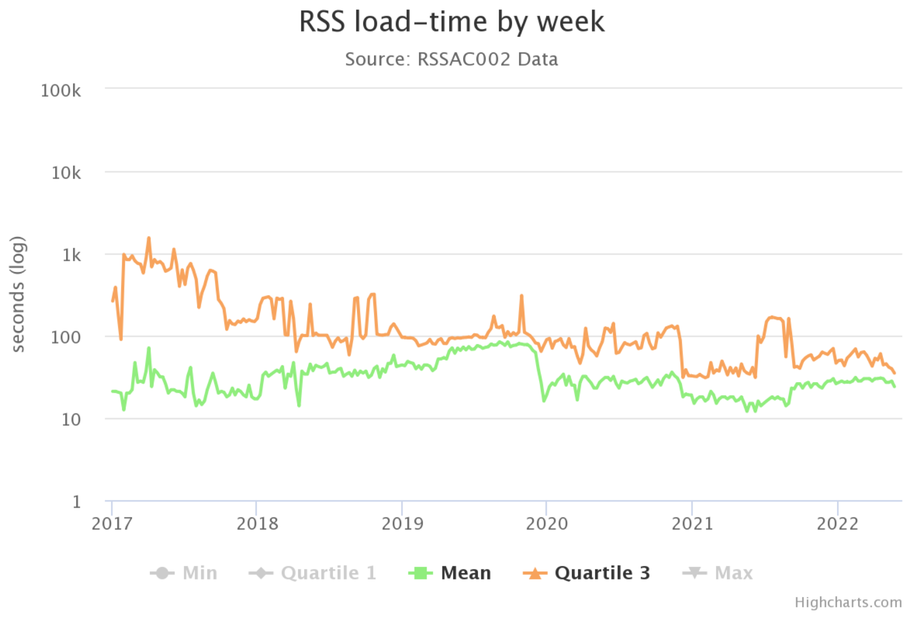

DNSMON

- DNS measurements of all root name-servers (+ several TLDs)
- UDP/TCP and IPv4/IPv6
- Uses RIPE Atlas Anchors

https://atlas.ripe.net/dnsmon/

ISC Atlas-vis

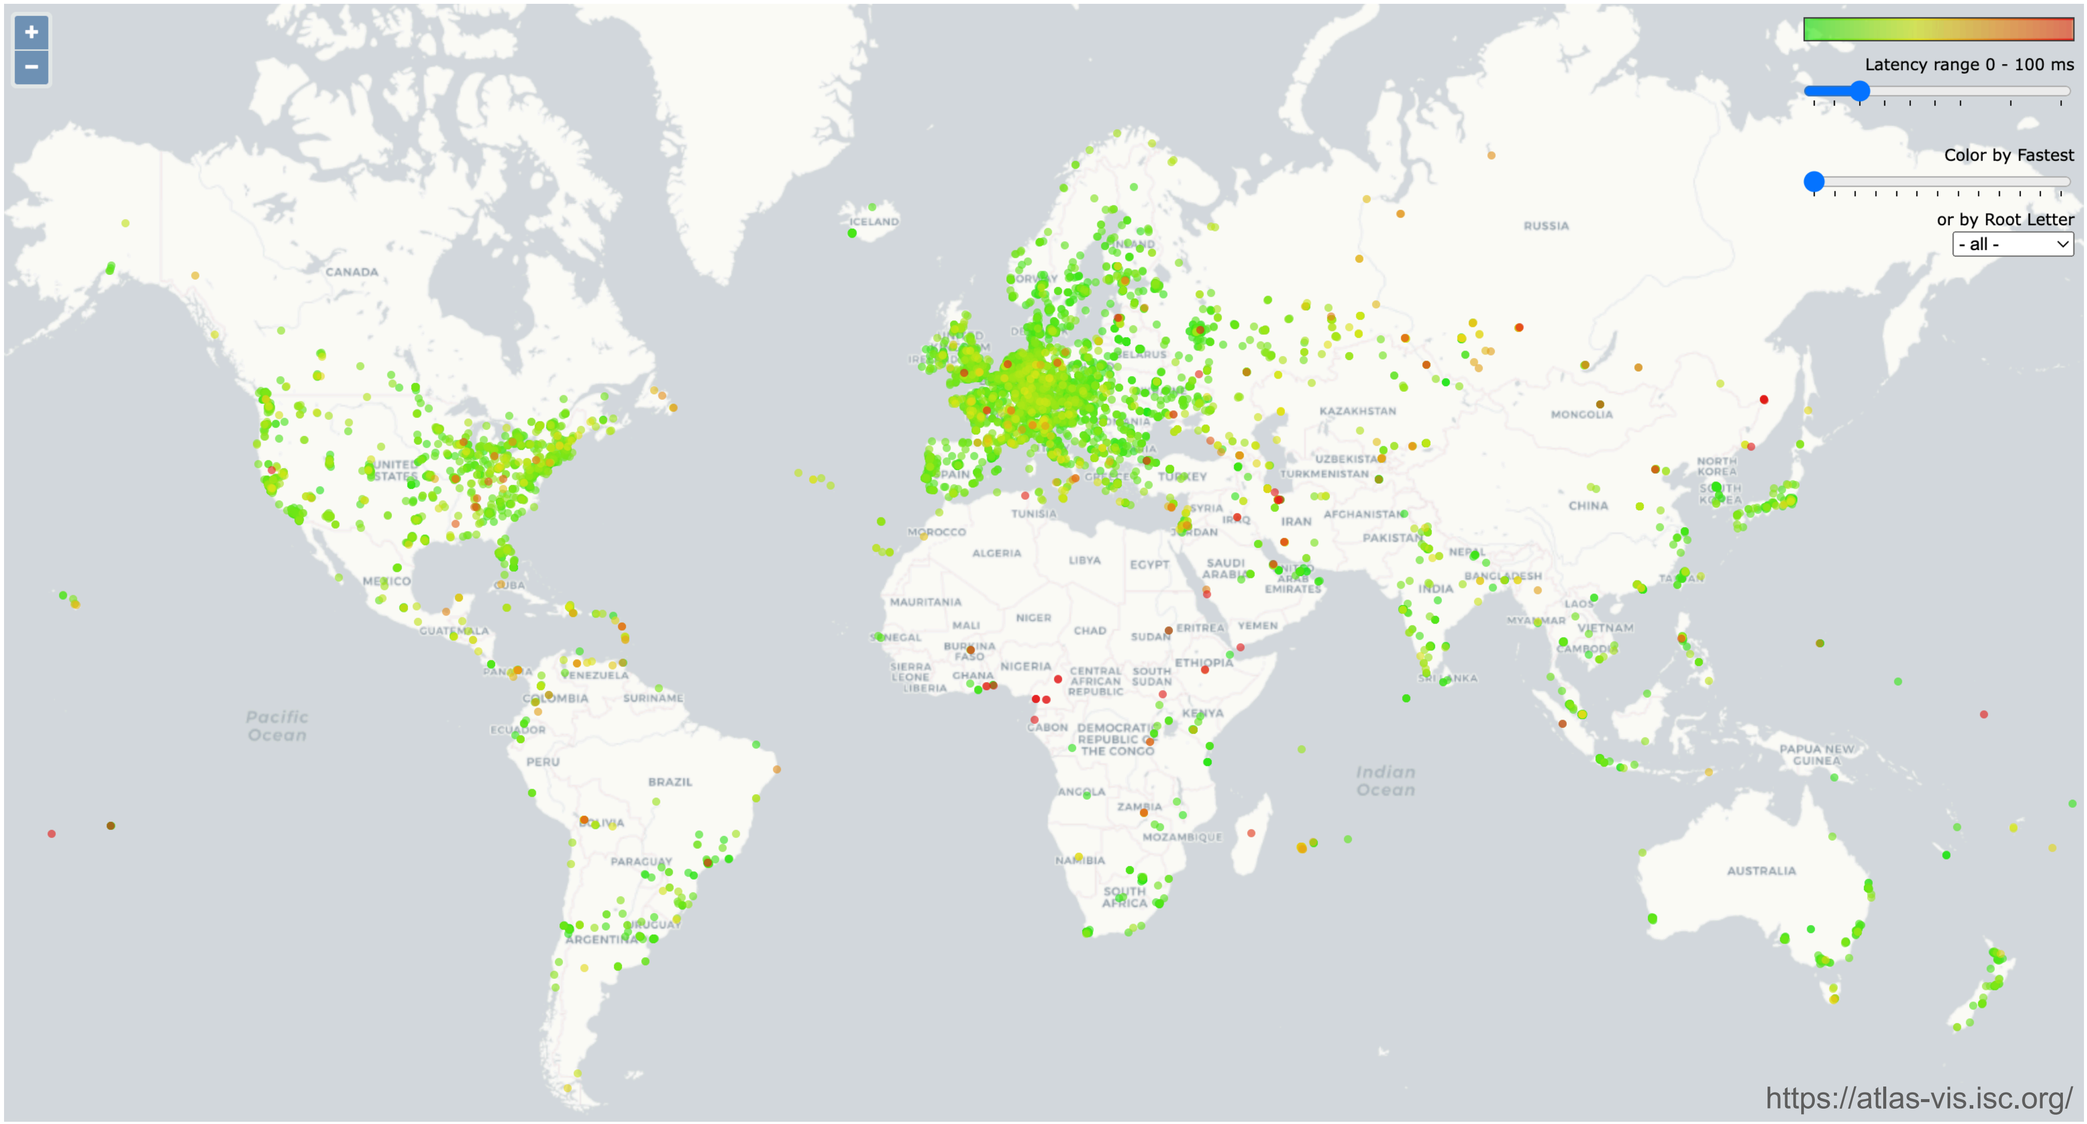

ISPs and the Root

How can you help?

Do Not Use the Root

- Implement RFC8806 (hyperlocal instance)
- Host an AS112 instance (https://www.as112.net/)
- Resolvers: Configure DNSSEC + RFC8198 (Aggressive NSEC)

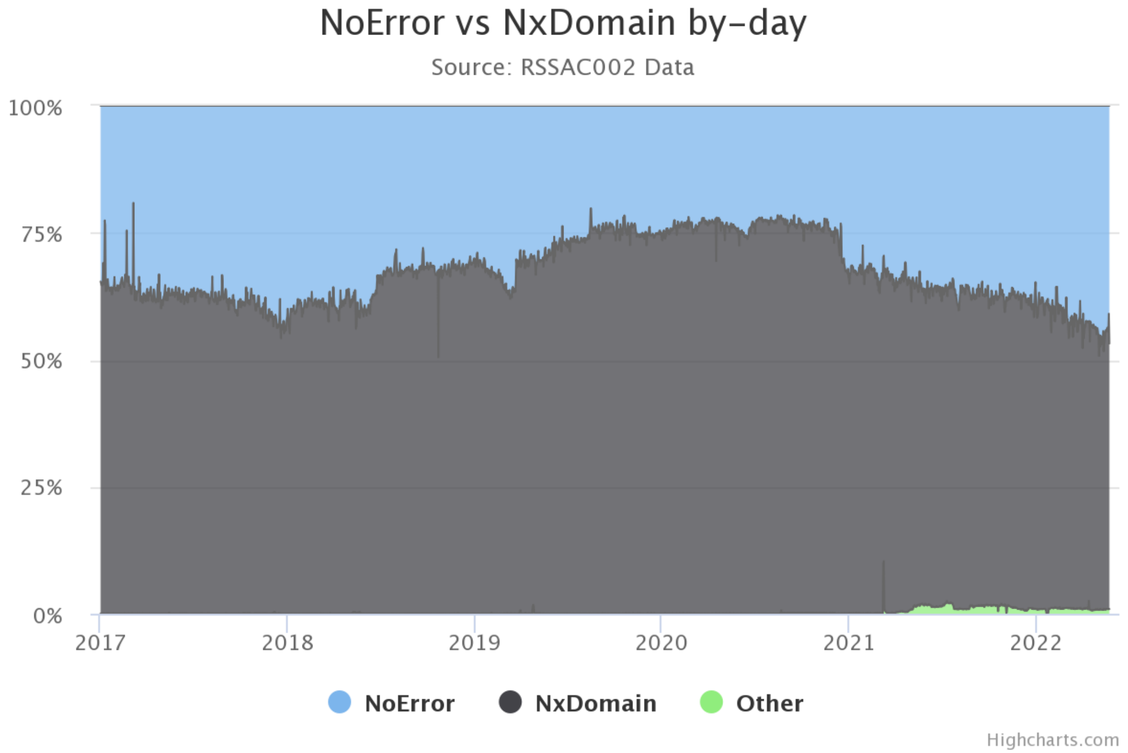

Peering with Root Servers

- Reduces latency
- Increases path diversity
- Available at many IXPs
  https://peeringdb.com/

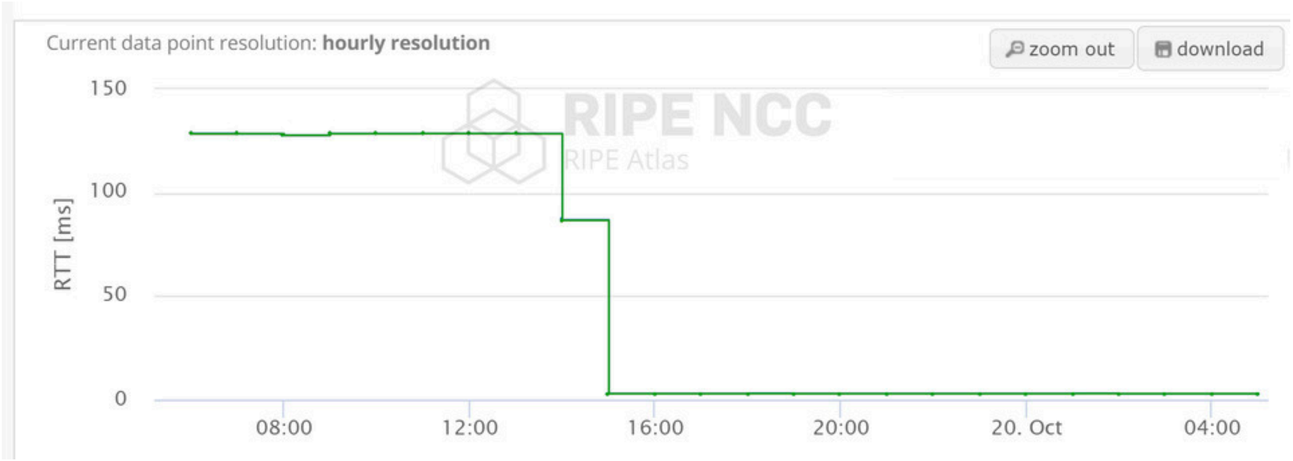

Hosting a Root Server Instance

- Benefits
  - Reduces latency (depending on location)
  - Increases resilience of the RSS
- Requirements
  - Professionally run colocation environment
  - Stable (BGP) network connectivity
  - Detailed requirements vary per RSO

Questions

mpels@ripe.net
www.root-servers.org