The Root Zone from A to Z

Naela Sarras, naela.sarras@icann.org

Kim Davies, kim.davies@iana.org

Terry Manderson, terry.manderson@icann.org

David Conrad, david.conrad@icann.org

Roy Arends, roy.arends@icann.org

21 April 2021

ICANN Managed Root Server

The Techie Bits

Terry Manderson
Snr Director, Security and Network Engineering

The Root Zone from A to Z

21 April 2021

IMRS Drivers – Security, Stability, Resiliency

1. More of the same
   - Figure out the best places to put new instances
   - Deploy more IMRS clusters
   - Deploy more IMRS singles
   - Get people to turn on DNSSEC in validating resolvers
     - Utilize “NSEC Aggressive Use”

David Conrad, CTO

- **More “Clusters”, More “Singles”**
  - What is the best fit for a location?
  - How many of the “13” do you already have? No need to “collect them all”!
  - Are the right building blocks there?
    - Connectivity
    - Datacenter
    - IPv6 & BCP38 support
- **Real world challenges**
  - No such thing as an open cheque book – responsible use of resources!
  - Ideally, reduce last mile from recursive resolvers to root server instances
  - 24x365 in both operations and security

Single v Cluster

Single

- Hosted by Host Orgs*
- Operated by ICANN
- One machine
- 1RU – 4RU
- Dependent on needs
- Entry point $US2K
- Robust, “locked down”
- Want to host? Contact your local ICANN GSE person!

Vs

Cluster

- Only in ICANN Datacenters
- Lots of machines and routers and switches
- At least a full rack
- Multiple Internet providers
- Extreme levels of physical and logical security
- Stop gap … if all else fails

The Details

- IMRS <-> ICANN Managed Root Server
  - The “letter” isn’t important
  - What we do, and how we approach it is important
  - Actual label is “L.ROOT-SERVERS.NET”
- Techie details
  - ASN: 20144
  - IPv4: 199.7.83.42
  - IPv6: 2001:500:9F::42
- Transparency (this is a key ethos of IMRS)
  - Stats stats and more stats
  - https://stats.dns.icann.org

The Details: Transparency

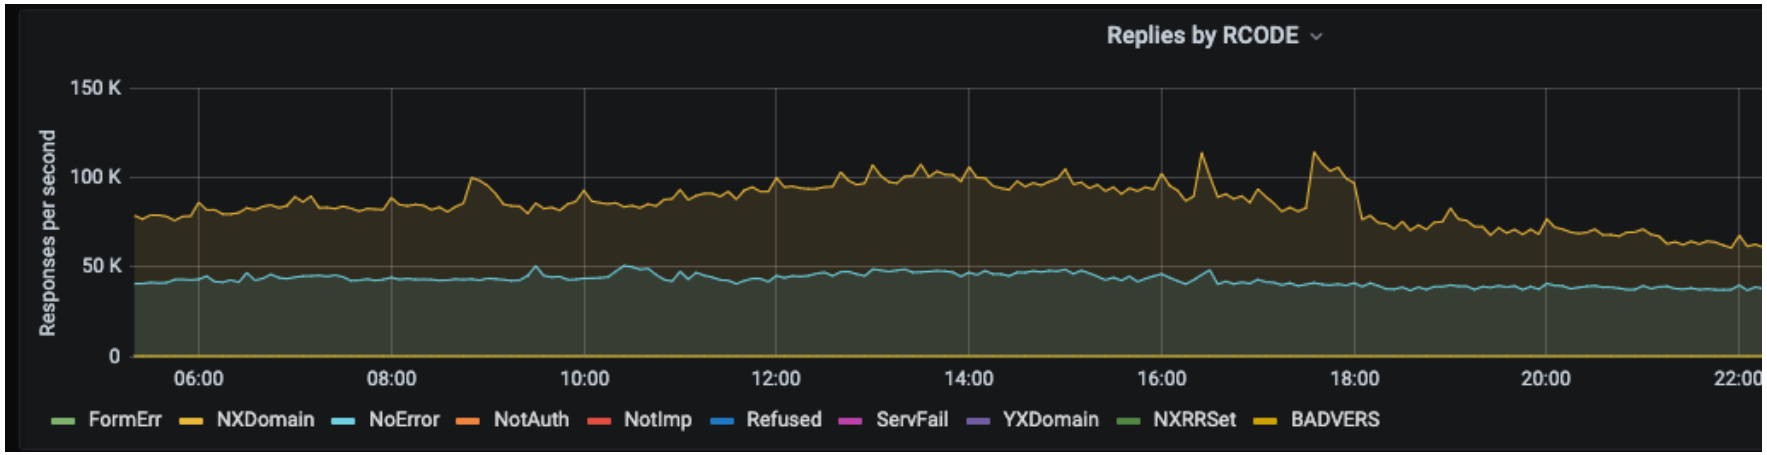

- https://stats.dns.icann.org
- Open Source: http://dns-stats.org/
- Uses the C-DNS format RFC8618 (https://tools.ietf.org/html/rfc8618)

The Details: 187 Instances in 84 countries

How many instances?

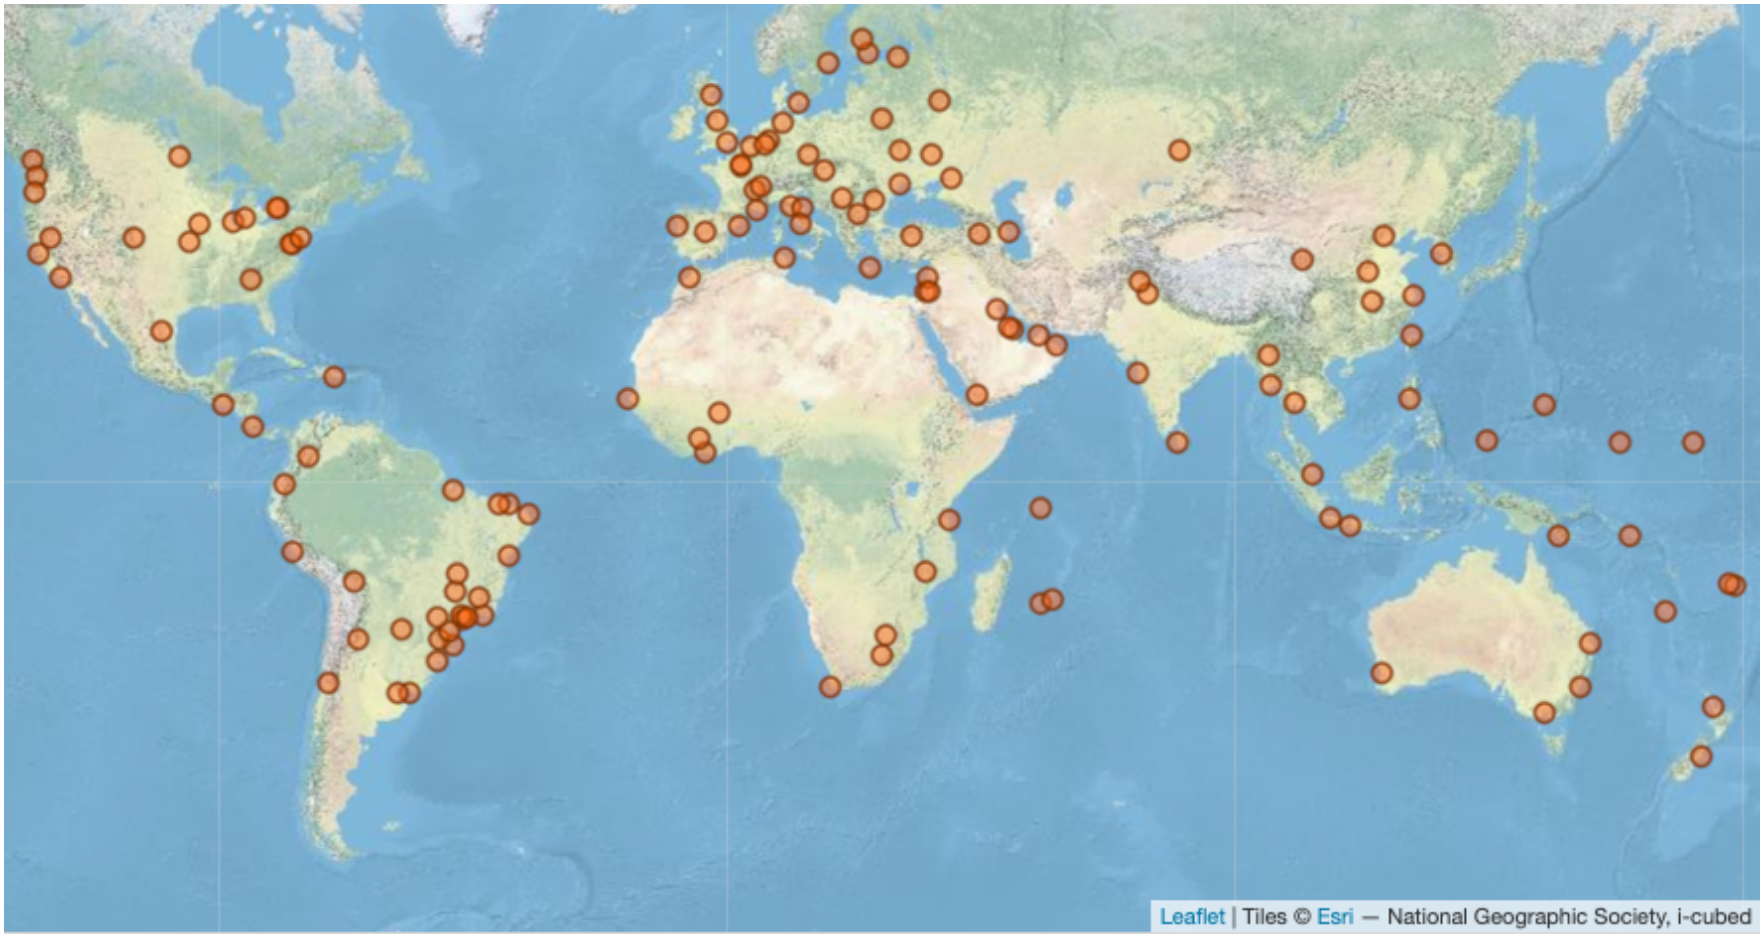

https://www.dns.icann.org/imrs/locations/

Engage with ICANN

Thank You and Questions

Visit us at **icann.org**
Email: noc@dns.icann.org

 @icann

 facebook.com/icannorg

 youtube.com/icannnews

 flickr.com/icann

 linkedin/company/icann

 soundcloud/icann

 instagram.com/icannorg

ICANN Root Service Strategy

Preparing for the Worst

David Conrad
ICANN Chief Technology Officer

North America Stakeholder Webinar: The Root Zone from A to Z

21 Apr 2021

A Few Definitions

For the purposes of this presentation, I'll be using these terms

Root Server	One of 13 "identities" (i.e., an IPv4 and IPv6 pair) associated with an authoritative name server specified in the root zone and the root hints file, which serves root zone data. Acts as "an entry point to the root server system cloud."
Root Server Operator	The organization responsible for managing the authoritative name servers (and related infrastructure) of one of the root servers.
Root Server System	The set of all 13 root servers.
Root Service	The service by which queries for root zone data are answered, typically (but not exclusively) offered by the root server system. Critical for normal Internet operation.
IMRS	The ICANN Managed Root Server, one of the 13 root server identities. Historically known as "L.ROOT-SERVERS.NET".

Background

- The Internet Today
  - DoS attack capacity increasing **exponentially**
    - ”IoT = Internet of Threats”
    - Cost to attackers negligible
  - Defending against DoS has non-zero cost
    - At the root, traditional solution: throw bandwidth/CPU at it, usually in the form of “anycast instances”
  - Increasing risk of data compromise
  - Increasing concerns about data privacy

Exponential growth in DDoS attack volumes
Oct 2020 – Google Cloud Blog

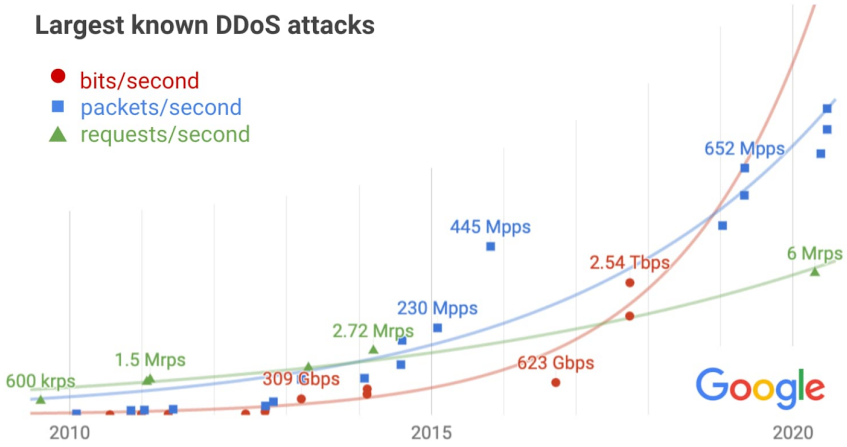

https://cloud.google.com/blog/products/identity-security/identifying-and-protecting-against-the-largest-ddos-attacks

- **The Problem**
  - Assuming current trends continue or accelerate, what can ICANN org do to help minimize the future risk of an attack disabling/compromising root service?
- **The Constraints**
  - ICANN does not have infinite funding
  - ICANN has control over exactly 1/13 of the root service infrastructure

Strategy – A Rough Summary

See OCTO-016: "*ICANN's Root Name Service Strategy and Implementation*"
- Being revised based on public comment

1. **More of the same**
   - Figure out the best places to put new instances
   - Deploy more IMRS clusters
   - Deploy more IMRS singles
   - Encourage people to turn on QNAME Minimization and DNSSEC in resolvers
     - Utilize "NSEC Aggressive Use"
   - Constantly monitor and improve performance

2. **Explore options**
   - Can we use "the cloud"?
     - Probably not, but needs more study

3. **Look to the future**
   - Track technology developments and implement where appropriate
   - "Decentralize all the things!" aka "hyperlocal root service"

RFC 8198: NSEC Aggressive Use

- If a zone is signed with NSEC (not NSEC3), like the root
  - Query for a non-existent name returns cacheable information about the range in which names do not exist
- Very good defense against "non-existent name" DoS attack
  - Rare but can be effective
- Enabling DNSSEC validation is easy in most resolvers
  - Might protect customers
  - Removes one type of DoS attack
  - Increases the usefulness of DNSSEC
    - Enables an alternative PKI

If you run a resolver, please turn on DNSSEC validation

Warren Kumari@Google (at IEPG) wrote:

- May 12, 2016 (a Friday afternoon), Colin Petrie / Kaveh Ranjbar from RIPE poked me: "Google is suddenly sending K-root way more junk queries, e.g 'nq0nnjzba-fn.357.225.340.251'. It burns us, please make it stop…"

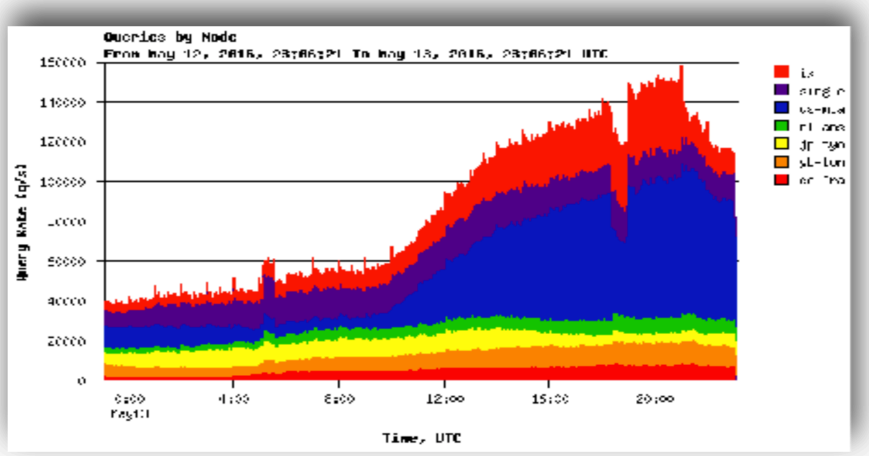

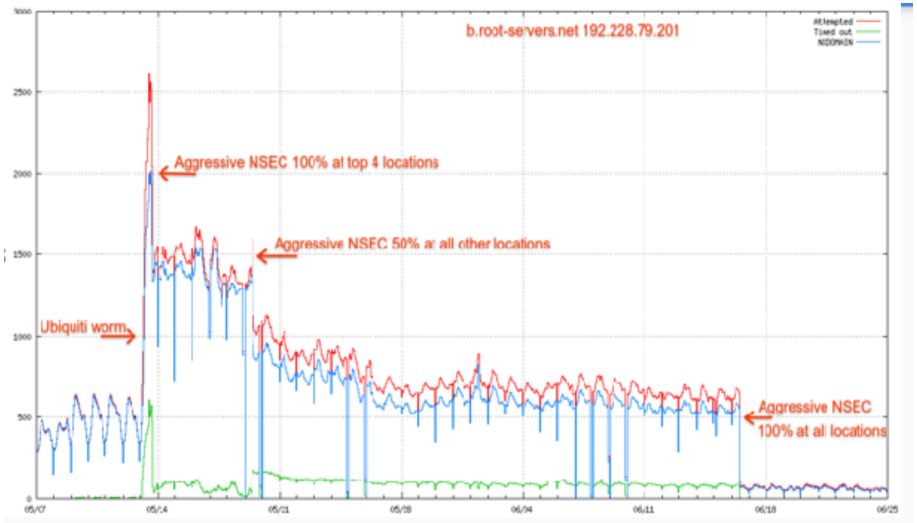

Looking to the Future: Hyperlocal Root Service

- RFC 8806: “Running a Root Server Local to a Resolver”
- Pros
  - Reduces load on the root server system
  - Lower latency for queries needing root data
    - Particularly helpful for non-existent names
  - Improves resilience
  - Increases privacy
  - Better aligns service provision with service funding
    - People benefitting from the service are paying (perhaps indirectly) for the service
- Cons
  - Possible misconfiguration has more interesting impact
  - Reduces telemetry to folks who monitor root traffic
  - Harder to change if we need to
    - As we saw with the KSK rollover
  - Need a better way to make the root zone available and ensure its integrity
- ICANN OCTO will be publishing a technical analysis Real Soon Now

Summarizing…

- ICANN's Board approved the strategy discussed here to help mitigate future risks related to threats to root service.
  - Most of the strategy is focused on increased decentralization, either through anycast or hyperlocal root service.
- The risks are driven by increased **use** and **abuse** of the DNS
  - **Use** can probably be handled the traditional way: throw money at the problem
  - **Abuse** is the real risk given proliferation of insecure devices
- ICANN org operates 1/13th of the root server infrastructure and we have limited funds, so any risk mitigation strategy must take this into account.
  - In the near- to medium-term, continuing what we're doing is probably the right choice
  - In the longer term, decentralization like hyperlocal root service may be necessary
    - If current attack trends continue or accelerate

Engage with ICANN

Thank You and Questions

Visit us at **icann.org**
Email: octo@icann.org

@icann

facebook.com/icannorg

youtube.com/icannnews

flickr.com/icann

linkedin/company/icann

soundcloud/icann

instagram.com/icannorg

Hyperlocal Root Zone

A Collective Term for Using the Root Zone Locally

Roy Arends
Principal Research Scientist, ICANN's Office of the CTO

21 April 2021

Hyperlocal

“relating to or focusing on matters concerning a small community or geographical area.” *(Oxford English Dictionary*)

- Used in the context of local news and weather forecast provisioning
- Now more generally used in the context of provisioning data pertaining to locally used applications.
  - weather apps, local maps, local services, etc.
- Hyperlocal root zone: resolver uses a locally available root zone instead of root-servers

Hyperlocal

- Concept is not new
- Not invented by ICANN
- Suggested by Paul Mockapetris in 2003
  - Suggested by many since
- Researched in 2004 by David Malone: “Hints or Slaves”
- Many “user-group” questions throughout the last 10 years on how to do this
- Operators already do this
- Time for a technical analysis

Hyperlocal Impacts the Resolver in Various Ways

- Query privacy
- Root zone integrity
- Query latency
- Telemetry
- Operational complexity

Query Privacy

- DNS servers are observers (RFC6973)
  - an entity that can observe and collect information from communications, potentially posing privacy threats
  - DNS data is collected passively at observation points (passive DNS)
  - DNS data is kept for a long time and distributed to third parties
  - No transparency how DNS query data is collected, stored, processed, analyzed, used, shared, and sold
  - Query minimization and aggressive negative caching helps to preserve privacy
- A hyperlocal root zone avoids the need to send queries to root-servers
- A query not sent is a query that can’t be collected

Root Zone Integrity

- The bulk of records in the root zone are not DNSSEC signed
  - None of the delegation point NS records and glue records have signatures
- There is no transport security between root-servers and resolvers
- A hyperlocal root zone provides better integrity than individual responses coming from root servers.
  - Provided that the root zone is securely retrieved or securely checked
  - Currently with HTTPS, PGP signatures or TSIG (via LocalRoot)
  - Future: DNSSEC validated ZONEMD records

Query Latency

- A query to the root zone is often a resolver's first query in a series, blocking the rest of the series
  - This only happens sporadically though, when the information is not available in cache
- About 68% of queries to the root return NXDOMAIN
  - Chrome browsers send a large amount of nonce-labels, which causes a lot of processing
  - Responses will be cached, causing memory consumption in caching resolvers
  - Root-servers spend a lot of time answering these queries.
  - Google is working to fix this
- Hyperlocal root zone lowers latency, causing better throughput for all queries.

Reduced Telemetry

- DITL data provides a lot of fertile ground for DNS research
- Some interesting telemetry data, such as deployment of new features, v4/v6, UDP/TCP ratios will be lost
  - However, they could be observed elsewhere

Elements of Deployment

- Availability, or “Where am I going to get it?”
  - Root Server Operators? IANA? Root Zone Maintainer?
- Transport , or “How am I going to get it?”
  - FTP, HTTPS, AXFR?
- Integrity, or “How do I know it is correct?”
  - ZONEMD+DNSSEC, PGP, TLS…
- Timely Updates, or “How do I make sure that I use the latest”
  - Notify is handy, but I should check anyway
- Fallback Mechanism, or “What do I do when it fails?”
  - Make sure to use them root hints again.

Operational Complexity

- Current security provisioning is cumbersome
  - LocalRoot offers TSIG, but a shared secret doesn't scale well
- TLS certificates are guaranteed by Certicom, not IANA
  - Internic.net uses HTTPS
- PGP is cumbersome in an automated environment
  - How to roll the PGP key…
- Local disk management, simple file write rights, cronjob management
  - For hand-rolled deployments
- Some of this is addressed by modern implementations
  - Each implementation has its own method
- Cryptographic zone file integrity check remains an issue
  - . . . until ZONEMD is deployed

Hyperlocal Deployment Methods

- Resolver serves authoritative data
  - Clients may not see AD bit on root zone content from the resolver
  - LocalRoot ships this configuration
- Resolver uses a local authoritative server for the root zone
  - On the network, on loopback, or as an internal “mirror zone”
  - RFC8806 has this configuration. Bind uses “mirror zone”
- Resolver primes the cache with the root zone
  - Times out nicely, re-prime once a day
  - Knot resolver does this

Conclusion

- Hyperlocal root zone is not new, and has been deployed for years
- Recent software makes a hyperlocal root zone deployment easier
- There are benefits, such as better integrity, privacy, and latency
- There are drawbacks
  - such as less telemetry at observation points
  - additional operational complexity
- There is work to be done to make a hyperlocal root zone
  - Deployment more secure (ZONEMD)
  - More available (maybe via a pool of root-zone publishers)

Engage with ICANN – Thank You and Questions

One World, One Internet

Visit us at **icann.org**


@icann


facebook.com/icannorg

youtube.com/icannnews


flickr.com/icann


linkedin/company/icann


slideshare/icannpresentations


soundcloud/icann