# POTENTIAL ROOT SERVER FUTURES

David Conrad
david.conrad@icann.org
ICANN

# AGENDA

- How we got here
- Current status
- What problem are we trying to solve?
- Possibilities

# HOW WE GOT HERE

*(FROM: HTTP://WWW.DONELAN.COM/DNSTIMELINE.HTML)*

- May 1984: First test server (USC-ISIF) run at USC-ISI
- Jul 1984: SRI-NIC (ARPANet: 10.0.0.51, MILNet: 26.0.0.73)
- Jul 1985: ISIB (10.3.0.52) added
- Oct 1985: ISIC (10.0.0.52) and BRL-AOS (192.5.25.82, 128.20.1.2) added
- **Oct 1986: IANA requests more root servers**
- Nov 1986, root servers now:
  - SRI-NIC.ARPA 10.0.0.51 26.0.0.73 ; JEEVES
  - USC-ISIC.ARPA 10.0.0.52 ; JEEVES
  - BRL-AOS.ARPA 192.5.22.82 128.20.1.2 ; BIND
  - USC-ISIA.ARPA 26.3.0.103 ; JEEVES
- Mar 1987: All root servers now use domain names
- Nov 1987: Remove C.ISI.EDU, add GUNTER-ADAM.ARPA, C.NYSER.NET, TERP.UMD.EDU, and NS.NASA.GOV.
- Apr 1990: NS.NIC.DDN.MIL (192.67.67.53) added
- **Jul 1991: NIC.NORDU.NET added**
- Apr 1993: NS.INTERNIC.NET added
- Apr 1994: AOS.BRL.MIL renamed AOS.ARL.ARMY.MIL
- May 1994: KAVA.NISC.SRI.COM removed, NS1.ISI.EDU added
- Sep 1994: NS.ISC.ORG added
- **Aug: 1995: ROOT-SERVERS.NET introduced, existing root servers renamed “A”-”I”**
- Jan 1997: “J” and “K” added, operated by Network Solutions
- Feb 1997: “L” and “M” added, operated by USC-ISI
- May 1997: “K” moved to London, operated by RIPE
- Aug 1997: “M” moved to Tokyo, operated by WIDE

# CURRENT STATUS

- 13 root server letters
  - Operated by 12 organizations (3 non-US) across 466 sites in dozens of countries.
- DNSSEC-signed zone
  - No undetected modifications possible, at least with validating resolvers
- ICANN's RSSAC provides a venue for root server operators and interested stakeholder to coordinate
  - **Not** control

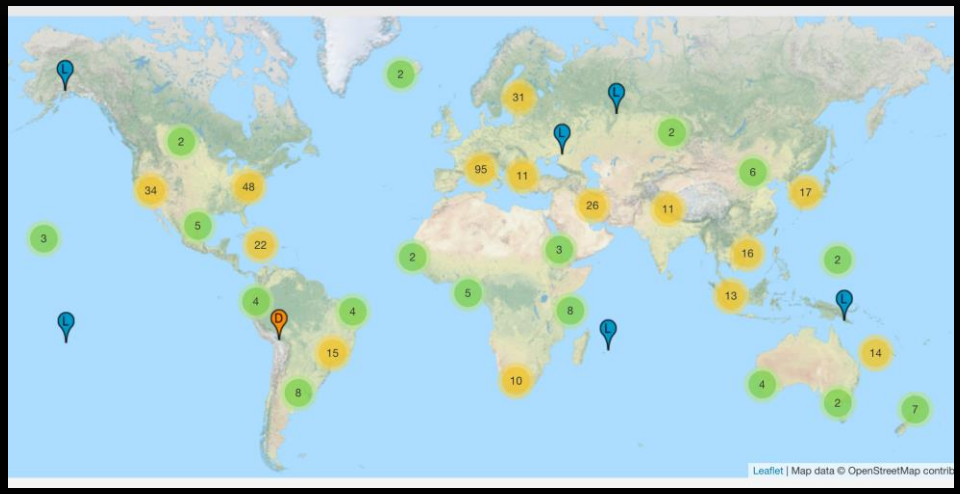


http://root-servers.org

# WHAT PROBLEMS ARE WE TRYING TO SOLVE?

- Distance/time to root server?
  - Particularly important for NXDOMAIN
- Root server overload?
  - E.g., (D)DoS
- Network Partitioning?
  - Inability to reach a root server
- Inappropriate management?
  - Making changes outside of policy

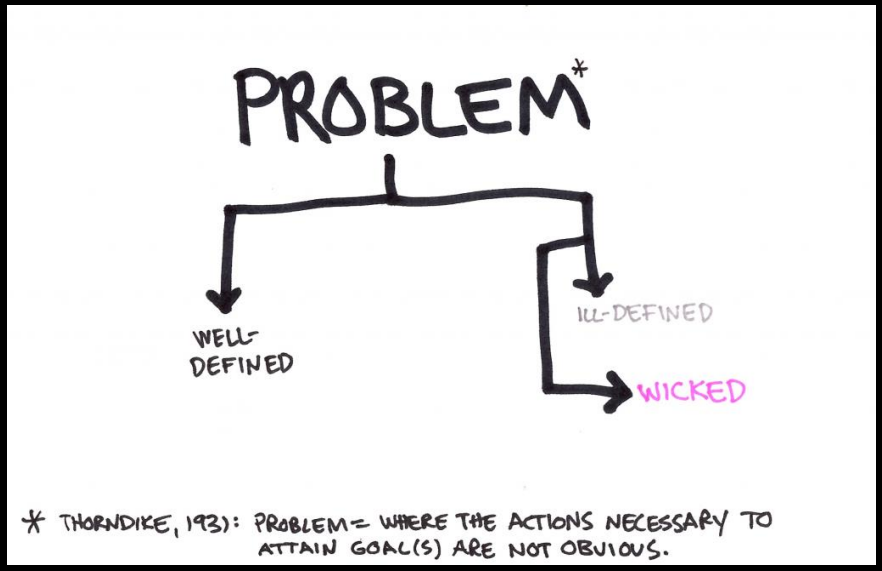

# DISTANCE TO ROOT SERVERS

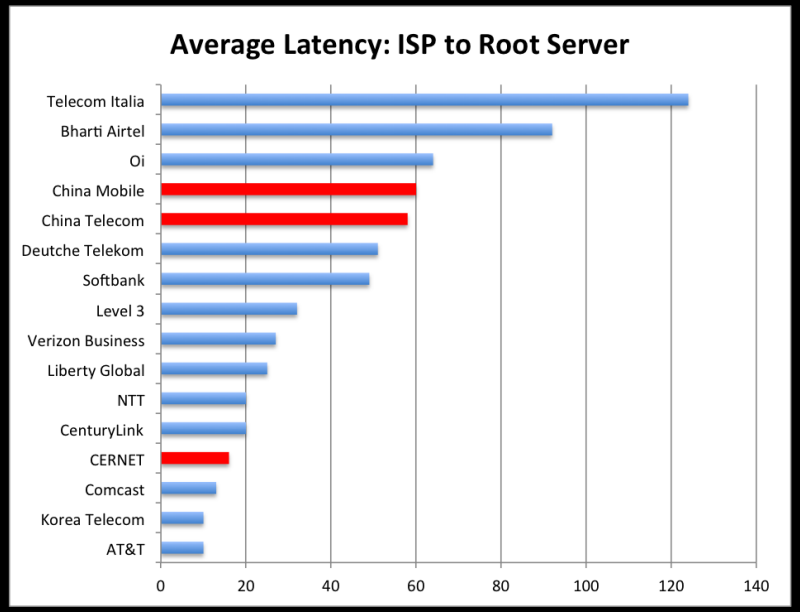


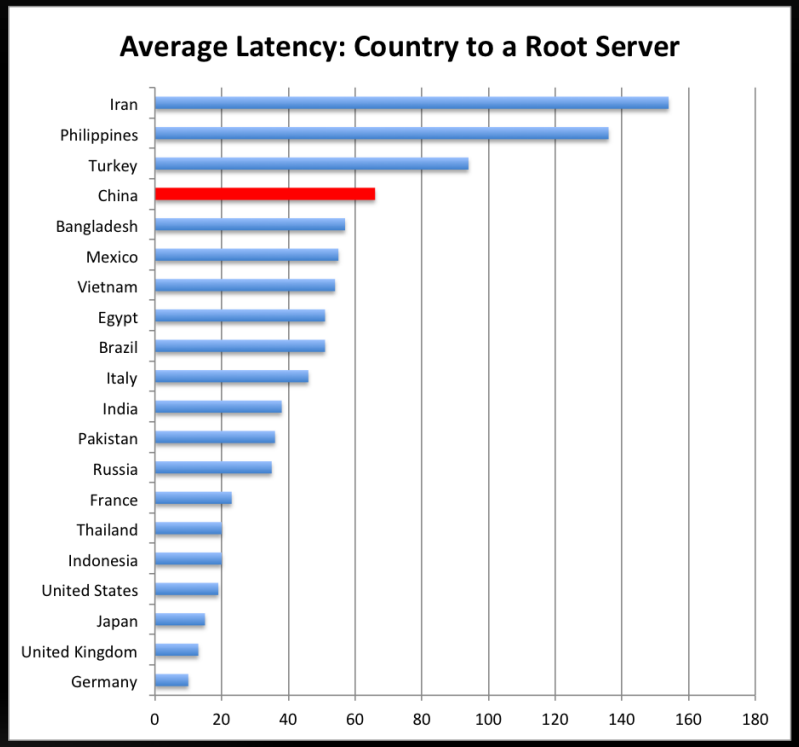


https://blog.thousandeyes.com/comparing-dns-root-server-performance/

# ROOT SERVER OVERLOAD

- Current sustained query load on “L” about 25,000 qps, so…
  - Assume same load on all root servers:
    - 13 x 25,000 = 325,000 qps
  - Current average query about 200 bytes
    - 325,000 x 200 = 65 MBps or 520 Mbps
  - Worst case response: about 1500 bytes
    - 325,000 x 1500 = 487.5 MBps or 3.9 Gbps
- Commodity servers and COTS software can do 200K qps easily
  - A couple of machines on 10GigE at a few IXes

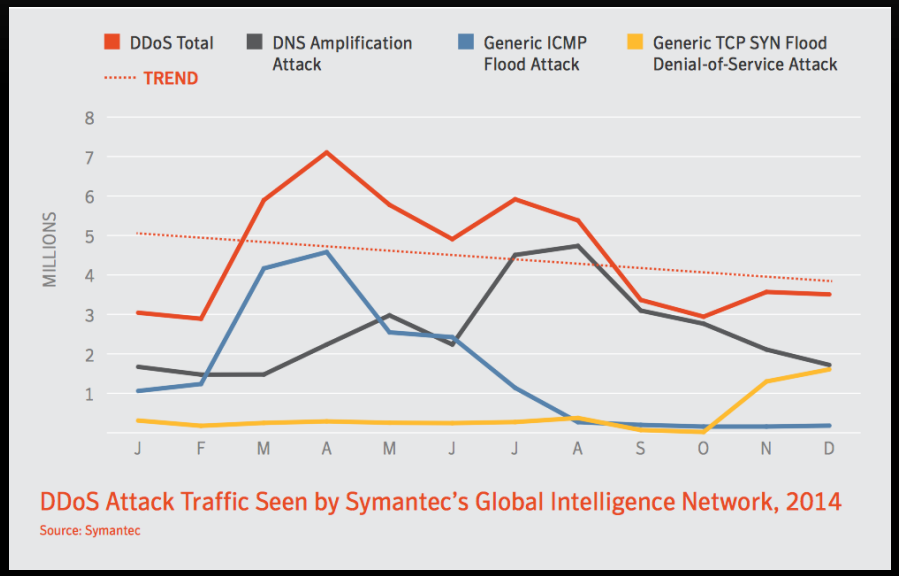

DDoS Attack Traffic Seen by Symantec’s Global Intelligence Network, 2014

Source: Symantec

Maybe Not…

# NETWORK PARTITIONING

- Accidental or malicious breaks in connectivity can remove access to root servers:
  - By root servers: root zone data will go stale
  - By clients: failure to resolve

# INAPPROPRIATE MANAGEMENT

- Examples
  - Serving different answers depending on who asks
  - Out of policy changes to TLDs
- **Not** a problem root servers can solve
  - With DNSSEC, both require resolvers to have different trust anchors
- Root servers are a publication mechanism
  - No editorial control
- With DNSSEC, only the holder of the Zone Signing Key can change zone contents

# POSSIBILITIES

- Add more servers
  - Add more instances
  - Add new letters
- Change the rules
  - “Unowned anycast”
  - Mirroring the root zone

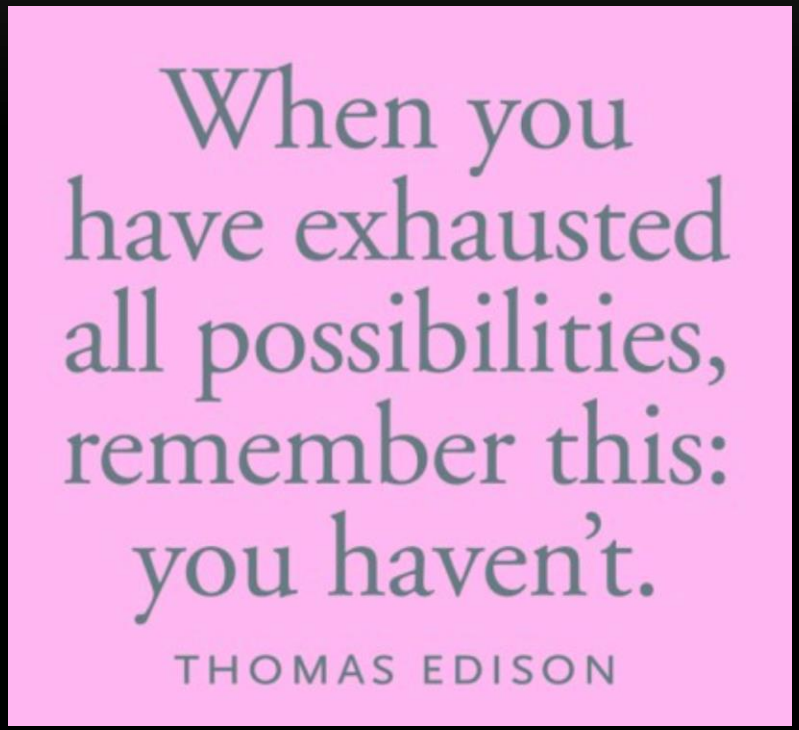

# ADD MORE INSTANCES

- ISC (F), NetNod (I), RIPE (K), and ICANN (L) and possibly others all willing to add instances for pretty much any requester, anywhere
  - Terms and conditions vary
- Requires entering into some sort of agreement with a Root Operator
- No change to protocol required
- Can reduce latency
  - Need to identify locations for new instances
- Can reduce global damage due to DoS
  - Localizes traffic
    - If you're near a lot of sources, too bad
- Can reduce risk of network partition
  - At least for folks outside the partition

# ADD MORE LETTERS

- Stay under 512 byte limit
  - Get rid of root-servers.net, move root servers to “a.”, “b.”, etc.?
  - Get rid of root glue in response Additional section?
- Increase response size
  - Maybe fragmentation isn’t that bad?
  - Move to TCP?
- **Hard problem:**
  - How to decide who operates the new letter?
  - Who decides?

- Does not solve any technical problem by itself
  - It all depends on how the new letter is implemented

# CHANGE THE RULES

- "Unowned Anycast": draft-lee-dnsop-scalingroot
  - Can do this today, but…
    - Potential stale data
    - Potential network management challenges
- Mirror the root zone in resolvers: draft-wkumari-dnsop-root-loopback
  - Can do this today, but…
    - Potential stale data
- Both require improved zone distribution system
  - A Content Delivery Network for DNS
- Statistics/monitoring?

- Both drafts can address latency
  - Moves responder to the end user's ISP or resolver operator
- Both drafts can mitigate DoS
  - The flood would be customer traffic
- Both drafts would reduce the effect of partition
  - At least until the root zone expires

# OTHER POSSIBILITIES?

- Adding more instances addresses latency to root servers, root server overload, and network partition concerns with no protocol changes and no policy development
  - “Mirroring the Root” and “Unowned Anycast” are both a variation of adding more instances
- DNSSEC prevents inappropriate management (assuming global multi-stakeholder management is appropriate)
- DNSSEC means you don’t have to care where you got the root zone.
- Adding more instances does not address non-technical problems.
  - How many root server (letters) do we really need?