Operating a Root Name Server

A Long Term Perspective

2022-06-23 APNIC. By Lars-Johan Liman

Many moons ago ...

- Some TCP/IP experiments were run in the hostile X.25 land referred to as "Europe".
- One by one, cells of connectivity were painstakingly merged, to the dismay of the national monopoly telcos.
- The DNS was invented and deployed 1987–1988.
  - Root name servers was suddenly a thing.
- NORDUnet was an early adopter, and after a trial period, it was appointed as root-server operator in 1991.

A long journey begins ...

- On July 28, 1991, the records

```
.                 IN NS    nic.nordu.net.
nic.nordu.net.    IN A     192.36.148.17
```

  were introduced in the root zone – along with some others.

- That's now more than 30 years ago ...

… for us both!

- The 31 years isn't the scary number …
- … it's the 30.5 years that I've been part of it …

Runestone from the 11th century, Uppsala, Sweden

What was it like in July 1991?

THE (singular!) Root Name Server in Europe was a workstation on a desk in an empty office.

- Sun 4/65 (SPARCstation 1+)
- LSI L64801 (25 MHz)
- 40 MB of RAM
- 10Base2 Ethernet (10 Mbps half duplex, coax cable)
- SunOS 4.0.3

The Netnod museum

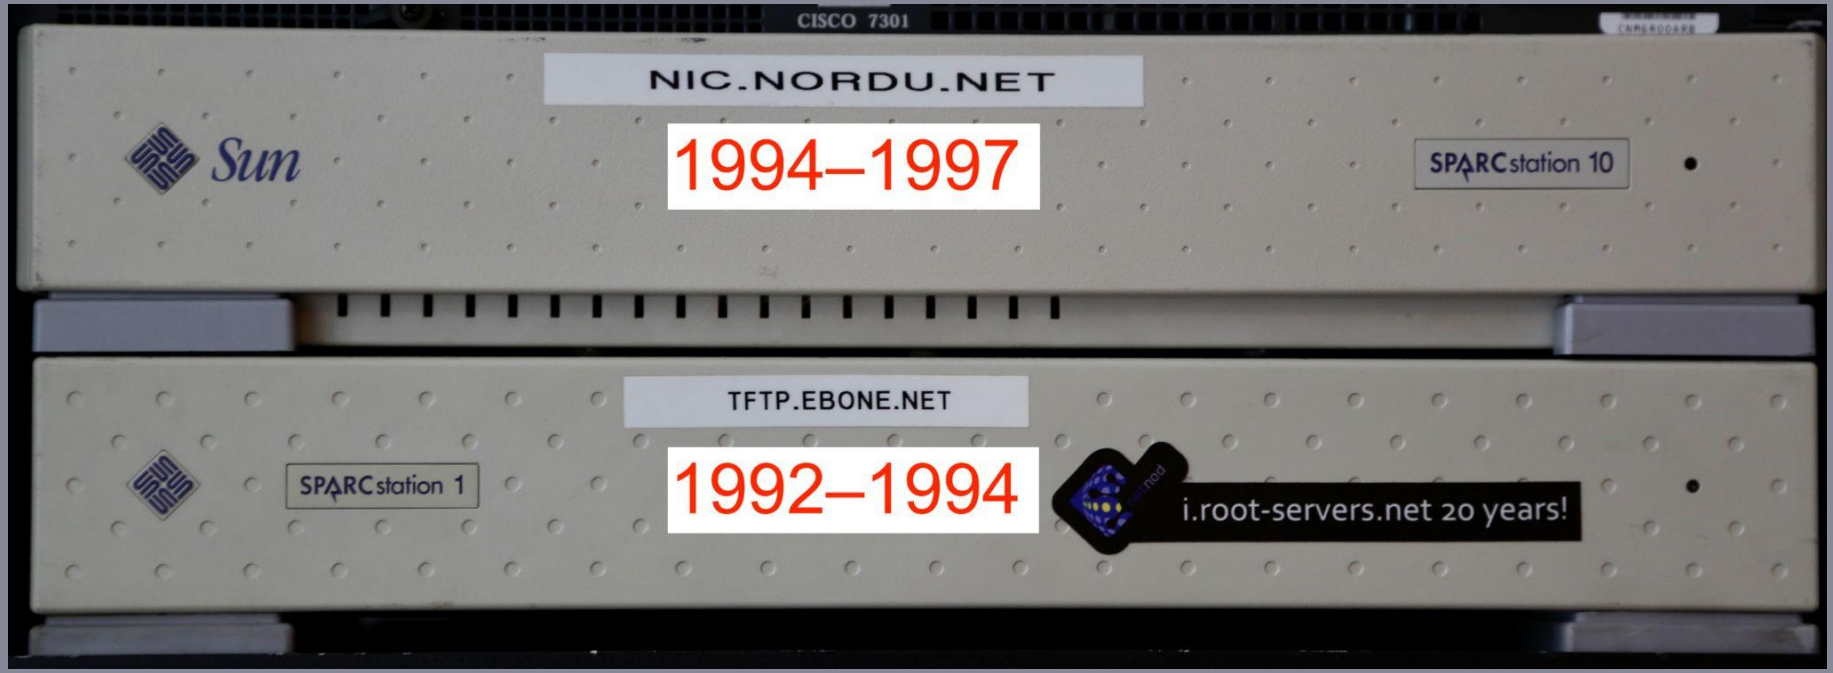

Functions 1992

- BIND version 4.8
- Both authoritative and full service resolver(!)
  - Queries were logged to the console
    You could read them all as they progressed in front of you …
- Served (IIRC) root, COM., NET., ORG., EDU., MIL., ARPA.
  - Also SE. and some other domains.

- FTP server
  - Official mirror of the RFC series archive
  - NORDUnet document archive

Challenges in 1992

- Getting the zone files!
  - Manual endeavour a 1–2 times per week.
  - Major packet loss to the USA. (30 % was not unusual, I remember seeing > 50 %.)
  - Had to use FTP. DNS zone transfer wasn't stable enough.
- Zone content mix (modern term: cache pollution)
  - Cache and authoritative were "in the same pool"
  - Served several "levels" (e.g., root + SE.) from the same server.
  - Zone borders were not enforced. You could put any data in any zone.
  - Lots of human errors and "lack of clarity in thought processes …"

The 1990s – politics

- The Internet picks up speed … and business!
- Politics enters the field.
  - Acceptable Use Policies (AUP) haunt the net.
  - Distribution of root servers becomes an issue.
- 1995: root servers were renamed
  - Saving space in the 512b-limited priming responses.
  - Making room for more servers.
  - nic.nordu.net ⇒ i.root-servers.net
  - {j, k, l, m}.root-servers.net were established.

1990s – IXs, AUPs, and renumbering

- Risks of allegations of traffic filtering to root servers.
- Stockholm D-GIX established in 1993/1994 at the KTH.
  - Idea: move the server to the IX and its own autonomous system.
- Problem: IP address was part of NORDUnet's backbone.
- Renumber I-root ⇔ renumber NORDUnet? 🤔
  - Hard-coded in 100.000s (?) root hints files?
  - 30–40 IP links in NORDUnet?
  - Easy! Renumber NORDUnet!

1990s: separation of services

- Network:
  - 1998: i-root is moved to Netnod IX (still operated by KTHNOC at the time).

- Memory:
  - 1992: 40 MB of RAM (Sun 4/65)
  - 1999: 4 GB of RAM (DEC AlphaServer DS20)
  - TLDs enabled their own services. One by one they left, including SE.

The turn of the century

- 1998: Jon Postel dies.
  - Result: all root-server operators meet for the first time ever.
  - Covert meeting.
  - All agree: IANA is the source of data.
- 1998: ICANN is created
  - RSSAC is one of the initial advisory committees.

- The Y2K problem …
  - Test runs
  - Staff on site

Gradually Improving Service

Over years the Root Server Operators (RSOs) have collaborated and coordinated their work to improve the service in several areas:

- Stability
- Reachability
- Security
- Data integrity

RSO motto: "Diversity is good!"

Various type of stability

- Server stability
- Service reachability
- Software diversity
- Policy stability
- Financial stability

This is how we do it at Netnod!

Addressing server stability

- Anycast
- Provisioning
  - Dual systems for critical parts.
  - VPN-based transport.
  - Machine-generated configurations.
- Statistics
  - RSSAC002
- RZERC

Service reachability

- Anycast
  - Thousands of peering relationships.
  - Extensive use of route servers.
  - Dialogue with site hosts.
  - Extensive use of peering-db.
  - Machine generated configurations.
- RPKI
- Internal know-how from IX experience

Software diversity

- DNS experience
  - "It's not the soldier with the best weapon who wins the battle. It's the soldier who knows his weapon best." (Swedish Army Capt. Hedberg, 1984)
- Operating several different types of DNS software.
- Maintain close relationships with software vendors.
  - Explain any special requirements for root service.
- Participate in industry-related meetings.
  - OARC, IETF, RIPE DNS WG
  - Operational fora.

Policy stability

- Global DNS policy affects the root more than other services.
- ICANN
  - IANA/PTI/CSC
  - RSSAC
  - RZERC
  - SSAC
  - RSS GWG

- ITU (New IP)
- EU (NIS2)
- IETF (DNSOP WG)
- RIPE
- CENTR
- Euro-IX, EPF, GPF, …

- We take part actively!

Financial stability

- Operating a root server requires neutrality and independence.
- No a-priori source of income.
- The Netnod way:
  - Provide commercial DNS service (primarily for TLDs) from other parts of the same platform.
  - The multitude of customers alleviates the possible problem with undue influence that comes with a single source of money.

Deployment challenges

Deploying servers

- Finding dependable and knowledgeable hosts.
- Finding sites that strike the cost–benefit sweet spot.
- Communicating with remote hands whom you have never met.
- Logistics and shipping.
  - "Customs. Bl--dy customs …" 🙂

Provisioning challenges

- Zone propagation.
- Status reporting and monitoring.
- Statistics: collecting and reporting.
- Slow admin links.

Service challenges

- Does the software do the right thing?
- … in a timely fashion?
- Performance is an issue.
- The root is a singular case. Sometimes software breaks in spectacular ways …
  - Support engineer: "You ***do*** know that you have to add the compile time switch '--enable-root-server' to make it serve 'the dot', right …?"

Operational challenges

- Distributed denial of service attacks
- Registry attacks
- The general onslaught of blackhat badness
  - Brute force login attacks
  - Phishing e-mail
  - … etc, etc.

Political challenges

To provide everyone in the Internet community across the public and private sector with a clear understanding of how the root server system works, and to make it clear that its purpose is technical infrastructure and nothing else.

Our goal is to ensure that any decisions that may impact the running of the root service are based on facts and not misconceptions.

Anecdotes?

- Reading queries off the screen ...
- Falling off the edge of the net ...
- Customs and their forms ...
- Routing issues …
  - Anycast has its shortcomings.

netnod.se

Lars-Johan Liman, M.Sc.
Senior Systems Specialist
Netnod Internet Exchange

https://www.netnod.se/