



K-Root Name Server Operations

Andrei Robachevsky

andrei@ripe.net



Outline

- Root Server System – brief update
 - Architecture
 - Current locations
 - Anycast deployment
- K.root-servers.net Server
 - Major milestones
 - Current status
 - K-Anycast deployment



Root Server System

- Provides nameservice for the **root zone**
 - Root DNS node with pointers to the authoritative servers for all top-level domains (gTLDs, ccTLDs).
- Thirteen name server operators
 - Selected by IANA
 - Diversity in organisations and location
 - 13 is a practical limit
 - a.root-server.net ÷ m.root-server.net - equal publishers
 - All thirteen are authoritative servers for the root zone
- An average client comes here < 8 times per week



Root servers and operators

- Thirteen nameservers, selected before 1997
 - a.root-servers.net Verisign
 - b.root-servers.net USC-ISI
 - c.root-servers.net Cogent Communications
 - d.root-servers.net University of Maryland
 - e.root-servers.net NASA
 - f.root-servers.net ISC
 - g.root-servers.net US DoD (DISA)
 - h.root-servers.net US DoD (ARL)
 - i.root-servers.net Autonomica
 - j.root-servers.net Verisign
 - k.root-servers.net RIPE NCC
 - l.root-servers.net ICANN
 - m.root-servers.net WIDE Project
- Look at www.root-servers.org

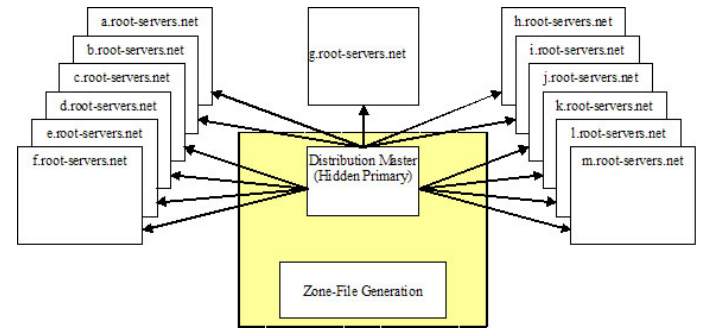


Location of 13 DNS Root Servers (pre-anycast era)

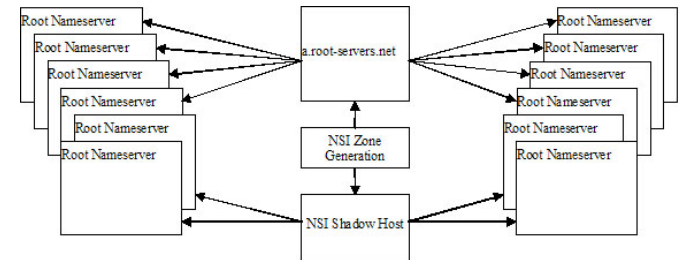


Evolution of Root System Architecture

- Public primary nameserver
 - a.root-servers.net primary
 - Other 12 are secondary
 - NSI generates the zone (Verisign since 2000)



- Enhanced architecture (2002)
 - Hidden distribution master
 - All 'letter' servers are equal
 - Authenticated transactions between the servers (TSIG)
- Wide deployment of anycast (2003)





Anycasting

- Point-to-point communication between a single client and the “nearest” destination server
 - Basics described in RFC 1546 in 1993
- “Cloning” a server
 - Multiple locations
 - Same operator
 - Same IP address belonging to the operator
 - Identical data
- Benefits
 - Distribution
 - Performance
 - Resilience
 - Redundancy



Location of 13 DNS Root Servers (spot the differences)





Location of 13 DNS Root Servers (spot the differences)

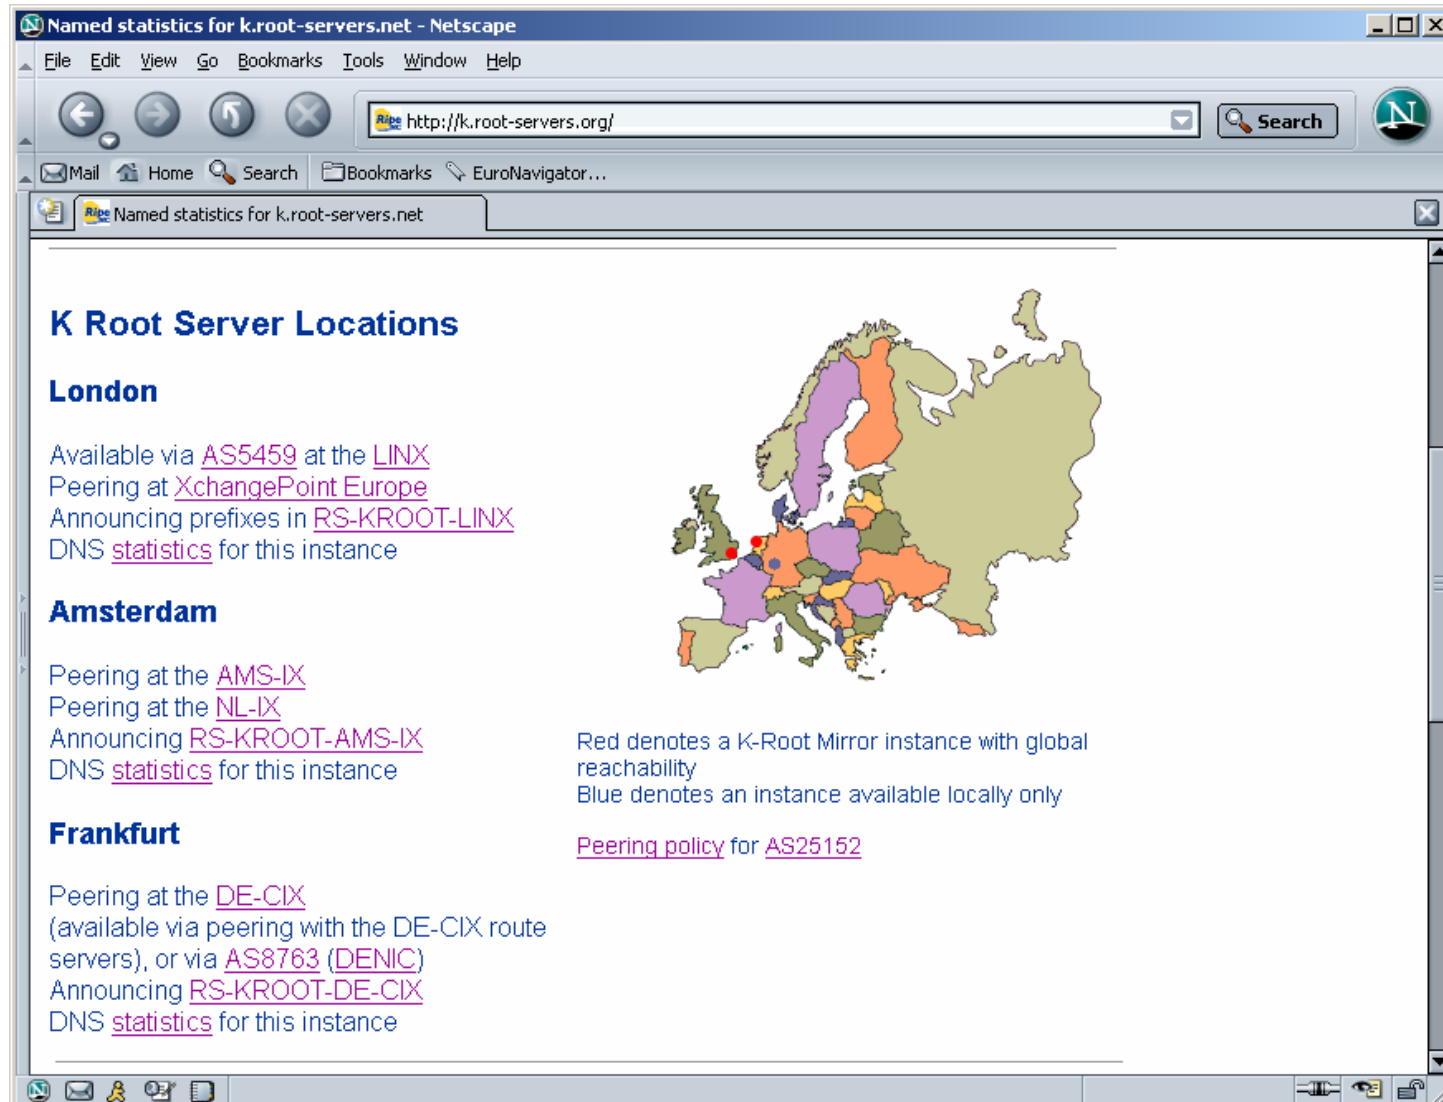




K-root Milestones

- Operated by RIPE NCC since May 1997
 - Hosted by LINX in London
- Running NSD since February 2003
 - Increased software diversity and performance
- Anycast since July 2003
 - Two global instances: London and Amsterdam
- Wider anycast deployment (2004)
 - 3-5 global nodes
 - 10-15 local nodes
 - Frankfurt, 19 January 2004

K-root Locations



K Root Server Locations

London

Available via [AS5459](#) at the [LINX](#)
Peering at [XchangePoint Europe](#)
Announcing prefixes in [RS-KROOT-LINX](#)
DNS [statistics](#) for this instance

Amsterdam

Peering at the [AMS-IX](#)
Peering at the [NL-IX](#)
Announcing [RS-KROOT-AMS-IX](#)
DNS [statistics](#) for this instance

Frankfurt

Peering at the [DE-CIX](#)
(available via peering with the DE-CIX route servers), or via [AS8763 \(DENIC\)](#)
Announcing [RS-KROOT-DE-CIX](#)
DNS [statistics](#) for this instance

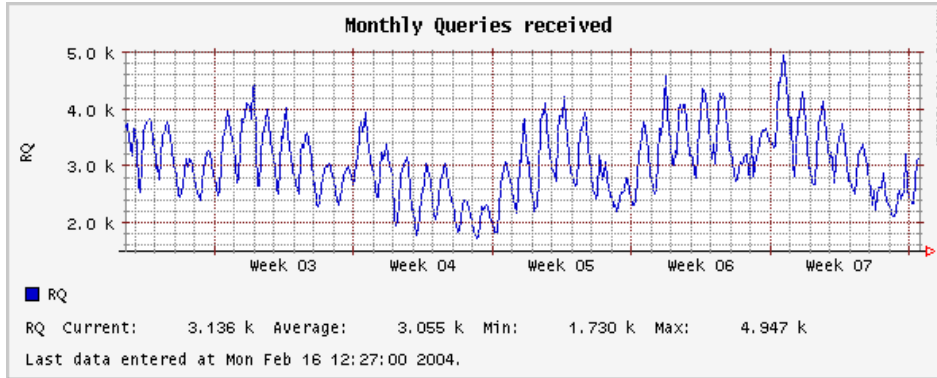
Red denotes a K-Root Mirror instance with global reachability
Blue denotes an instance available locally only

[Peering policy](#) for [AS25152](#)




K-root Statistics


London



Amsterdam



Frankfurt





“Local” Mirror Instances

- Objectives
 - Improving access to K for a significant ISP community
 - Isolating impact of an “external” DDoS
 - Localising impact of a “local” DDoS
- Location
 - Well connected points with significant ISP community (IXP, etc.)
 - Improved responsiveness for the members of the IX
 - Improved resilience of the whole system for others
- Model
 - Hosted by a neutral party
 - Open peering policy
 - Fully funded by a hosting party
- Operations
 - Exclusively performed by the RIPE NCC



“Global” Mirror Instances

- Ideally located at topologically equidistant places
 - In practice there are not so many choices
- Globally reachable
 - But less preferable than “local” mirror instances
- Powerful in terms of connectivity and CPU
 - Have to sustain DDoS and local nodes failures
- The same management model as for local nodes
 - RIPE NCC is the operator
- Different funding model
 - No distinguished group of local beneficiaries
 - Costs are shared between the host and the RIPE NCC

More Information

- Root operators
 - <http://www.root-servers.org>
 - [http://\[a-m\].root-servers.org](http://[a-m].root-servers.org)
 - <http://dnsmon.ripe.net>
- Root server analysis
 - <http://www.caida.org/projects/dns-analysis/>
- Anycasting
 - [Host Anycasting Service, RFC1546, http://www.ietf.org/rfc/rfc1546.txt](http://www.ietf.org/rfc/rfc1546.txt)
 - [Distributing Authoritative Name Servers via Shared Unicast Addresses. RFC3258, http://www.ietf.org/rfc/rfc3258.txt](http://www.ietf.org/rfc/rfc3258.txt)



More Information (cont.)

- K-root
 - <http://k.root-servers.org>
- K-root anycasting
 - Distributing K-Root Service by Anycast Routing of 193.0.14.129, RIPE- 268,
<http://www.ripe.net/ripe/docs/ripe-268.html>
 - General Requirements and Guidelines,
<http://k.root-servers.org/hosting-guidelines-200311.html>
 - Contact at k-anycast@ripe.net



<http://www.ripe.net/presentations>