

# K-Root Nameserver Operations

**Andrei Robachevsky**  
Chief Technical Officer

[andrei@ripe.net](mailto:andrei@ripe.net)

# Outline

- Root Server System
  - What is a root server?
  - Where is the root?
- Anycast Routing
  - The basics
  - Advantages of using anycast
- K.root-servers.net Server
  - History
  - 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 nameservers
  - Most of them selected before 1997
  - Diversity in organisations and location
  - 13 is a practical limit
  - a.root-server.net - m.root-server.net
  - All thirteen are secondary servers
- An average client comes here < 8 times per week

# Root servers and operators

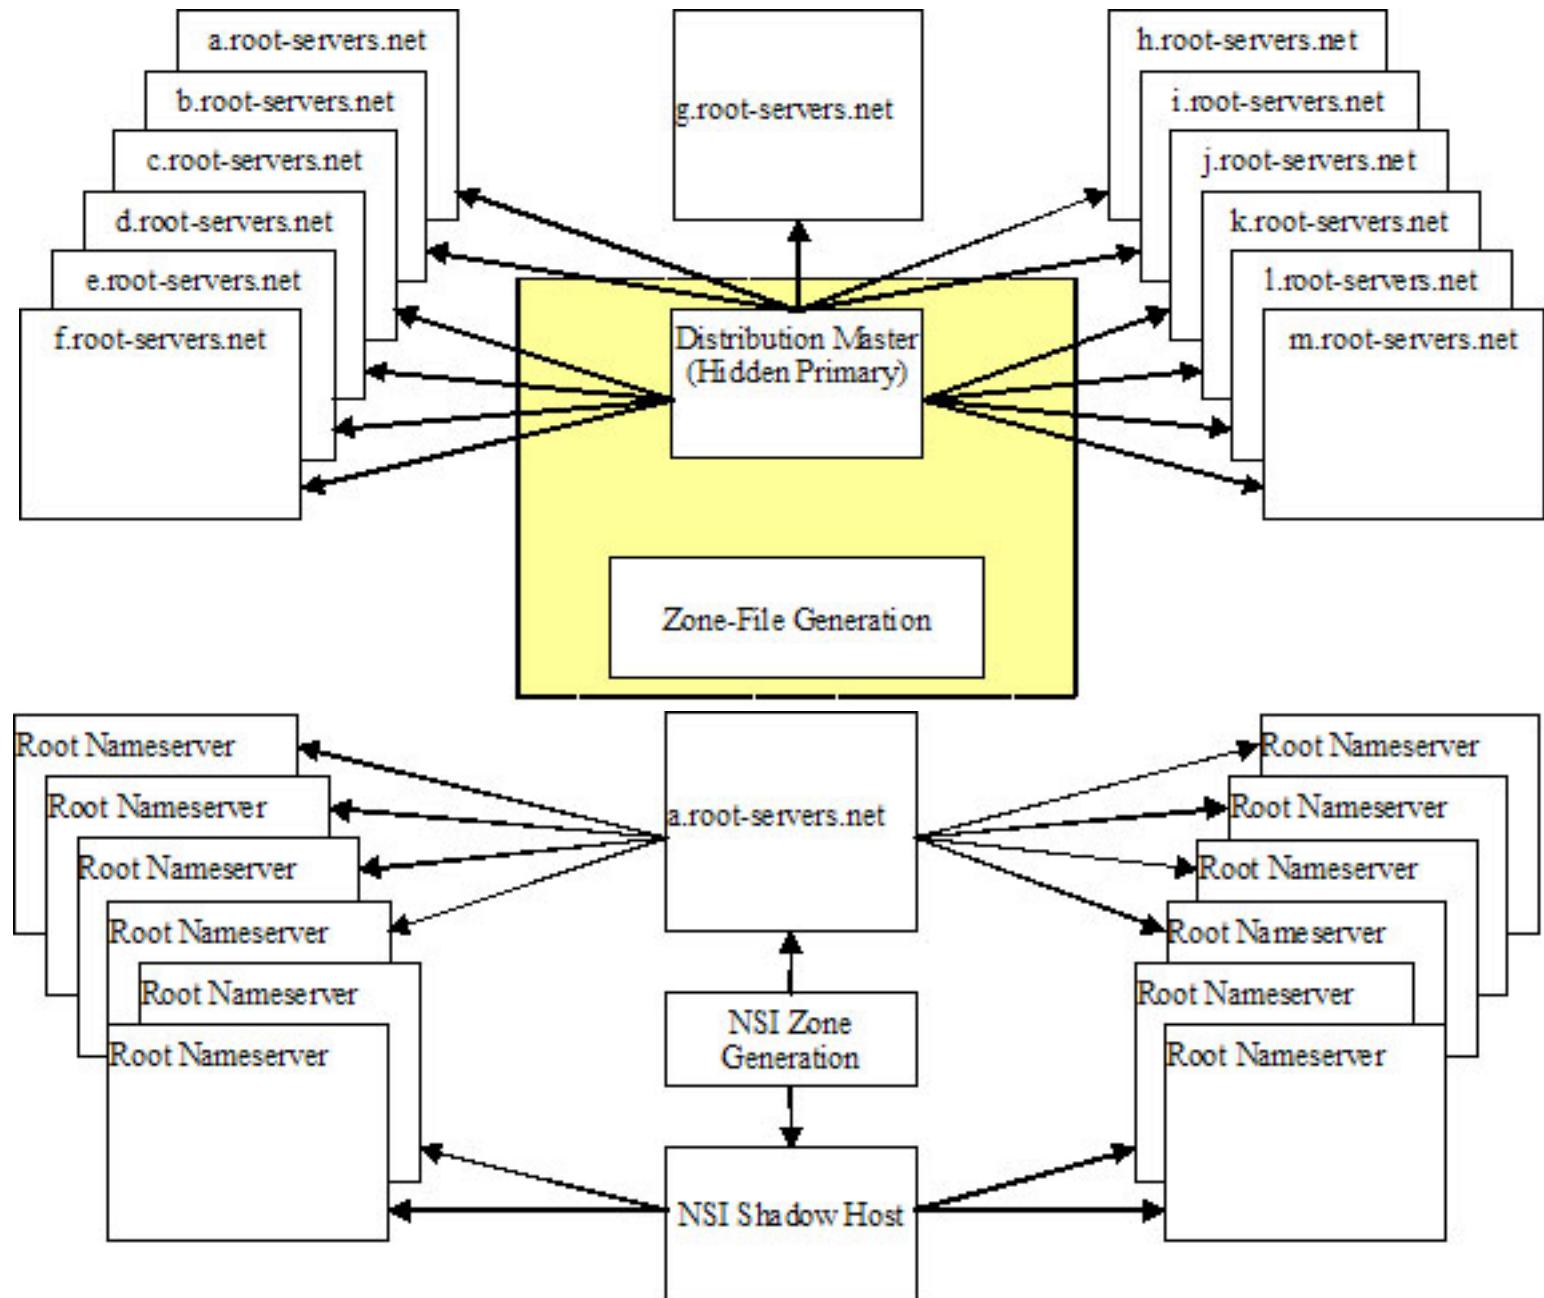
- Thirteen nameservers, selected before 1997
  - a.root-servers.net Verisign
  - b.root-servers.net USC-ISI
  - c.root-servers.net Cogent Communications
  - d.root-servers.net University of Maryland
  - e.root-servers.net NASA
  - f.root-servers.net ISC
  - g.root-servers.net US DoD (DISA)
  - h.root-servers.net 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](http://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 primary nameserver
  - 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 IP address
  - 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 instances: London and Amsterdam

# K-root current status

- 2 instances have been deployed
  - LINX & AMS-IX
  - Average load: 3K/2K qps
- Specification
  - Cluster of 2 servers
  - NSD
  - Sniffer box
  - Global reachability
  - Funded by the RIPE NCC (AMS-IX provides 1G port)

# Future nodes

- 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.)
  - Main beneficiaries are the members of the IX
- Model
  - Hosted by a neutral party (IXP itself?)
  - Open peering policy
  - Fully funded by a hosting party
- Operations
  - Exclusively performed by 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>
  - Distributing Authoritative Name Servers via Shared Unicast Addresses. RFC3258,  
<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/titletoc.html>
  - General Requirements and Guidelines,  
<http://k.root-servers.org/hosting-guidelines-200311.html>
  - Contact at [k-anycast@ripe.net](mailto:k-anycast@ripe.net)



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