f.root-servers.net

NZNOG 2

Joe Abley <jabley@isc.org>

DNS

- The Domain Name System is a huge database of resource records
  - globally distributed, loosely coherent, scaleable, reliable, dynamic
  - maps names to various other objects

Resolving www.isc.org

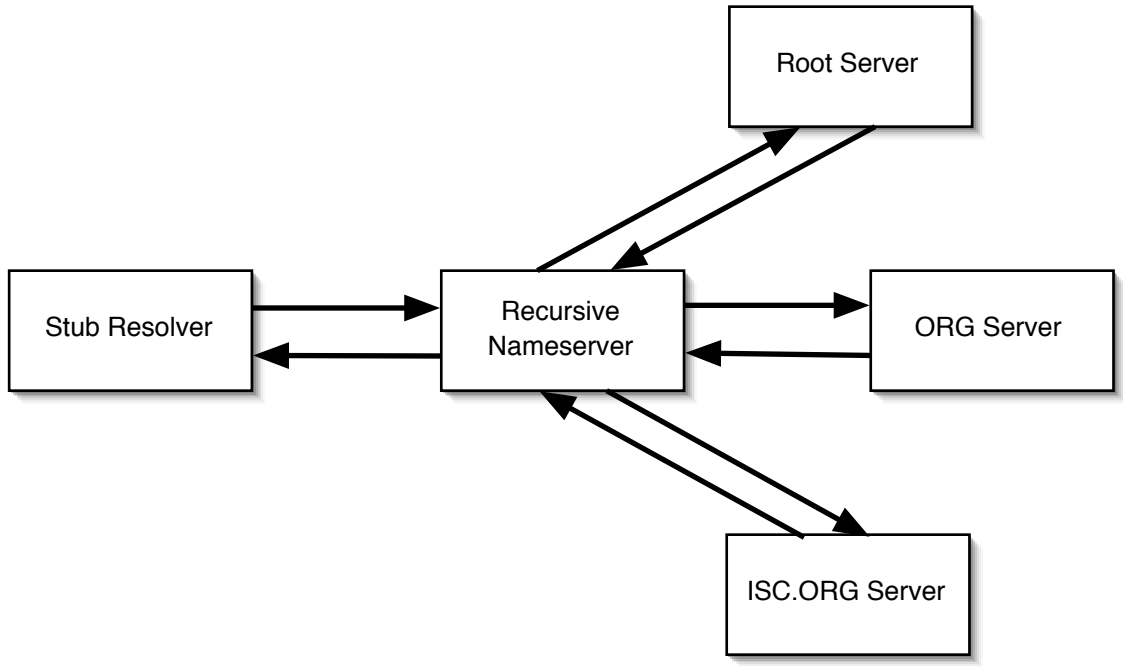

Root Servers

- Every recursive nameserver needs to know how to reach a root server
- Root servers are the well-known entry points to the entire distributed DNS database
- There are 13 root server addresses, located in different places, operated by different people
- http://www.root-servers.org/

The Root Servers

A.ROOT-SERVERS.NET	Verisign Global Registry Services	Herndon, VA, US
B.ROOT-SERVERS.NET	Information Sciences Institute	Marina del Rey, CA, US
C.ROOT-SERVERS.NET	Cogent Communications	Herndon, VA, US
D.ROOT-SERVERS.NET	University of Maryland	College Park, MD, US
E.ROOT-SERVERS.NET	NASA Ames Research Centre	Mountain View, CA, US
F.ROOT-SERVERS.NET	Internet Software Consortium	Various Places
G.ROOT-SERVERS.NET	US Department of Defence	Vienna, VA, US
H.ROOT-SERVERS.NET	US Army Research Lab	Aberdeen, MD, US
I.ROOT-SERVERS.NET	Autonomica	Stockholm, SE
J.ROOT-SERVERS.NET	Verisign Global Registry Services	Herndon, VA, US
K.ROOT-SERVERS.NET	RIPE	London, UK
L.ROOT-SERVERS.NET	IANA	Los Angeles, CA, US
M.ROOT-SERVERS.NET	WIDE Project	Tokyo, JP

Challenges on the Root

- There have been a number of attacks on the root servers
- Distributed denial of service attacks can generate a lot of traffic, and make the root servers unreachable for many people
- Prolonged downtime would lead to widespread failure of the DNS

Widespread Failure

- Probability of the entire DNS system failing is low
  - the most important data in the DNS (records which are frequently queried) are cached
- Regional failure is more likely
  - e.g. loss of international connectivity, bulk probe traffic from worms

f.root-servers.net

- Has a single IP address (192.5.5.241)
  - no change there
- Requests sent to 192.5.5.241 are routed to different nameservers, depending on where the request is made from
  - this behaviour is transparent to devices which send requests to F

Routing

- Most traffic on the Internet is unicast
  - packets have a single destination
- Some traffic is multicast
  - packets are directed to multiple destinations
- Traffic to f.root-servers.net is anycast
- packets are directed to a single instance of F, but different queries (from different places) may land on different instances

Anycast Routing

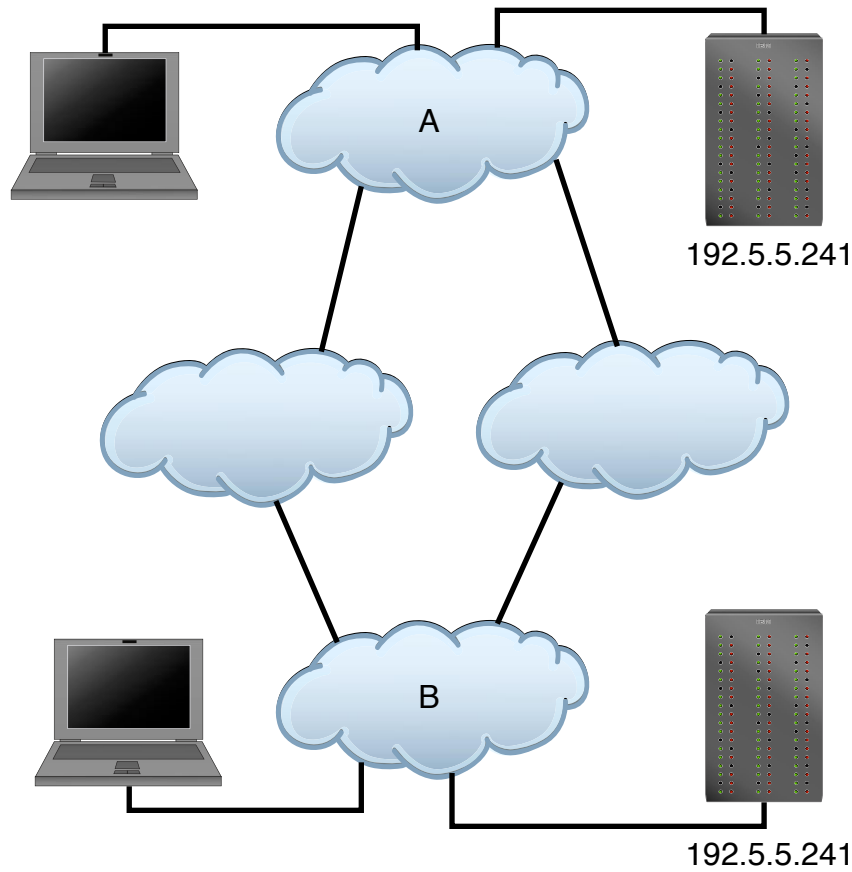

Hierarchical Anycast

- Some of the F root nameserver nodes provide service for 192.5.5.241 to the entire Internet (global nodes)
  - very large, well-connected, secure and over-engineered nodes
- Others provide service for 192.5.5.241 to a particular region (local nodes)
  - smaller

Hierarchical Anycast

- Architecture described in an ISC Technical Note
  - http://www.isc.org/tn/

Failure Modes

- If a local node fails, queries to 192.5.5.241 are automatically routed to a global node
- If a global node fails, queries are automatically routed to another global node
- Catastrophic failure of all global nodes results in continued service by remote nodes within their catchment areas

Sponsorship

- ISC is a non-profit company
- Equipment, colo, networks for remote nodes are paid for by a sponsor
- All equipment is operated by ISC engineers
- The sponsor covers the ISC's operational costs of running the remote node

Deployment Status

- Two global nodes
  - Palo Alto, CA, US
  - San Francisco, CA, US

Deployment Status

- Five local nodes
  - Hong Kong
  - Madrid, Spain
  - New York, NY, USA
  - San Jose, CA, USA
  - Los Angeles, CA, USA

Deployment Status

- Six! Six local nodes
  - Hong Kong
  - Madrid, Spain
  - New York, NY, USA
  - San Jose, CA, USA
  - Los Angeles, CA, USA
  - **Auckland, New Zealand**

Deployment Targets

- 10 local nodes live by the end of 2003
- 20 more in 2004

For More Information

- Contact ISC
  - Paul Vixie <vixie@isc.org>
  - Joe Abley <jabley@isc.org>
- Contact APNIC
  - Paul Wilson <dg@apnic.net>