Did F.ROOT-SERVERS.NET really relocate to Beijing?

Peter Losher, Senior Operations Architect
Internet Systems Consortium
NANOG 53 Lightning Talk

Overview of F.ROOT-SERVERS.NET (F)

- F is anycast globally from AS3557
- 55+ sites worldwide
- 4 nodes are considered "global" nodes
  - AS3557 peers with AS1280
- Remainder are considered "local" nodes
  - Serving a local community (IX, NAP, etc.).
  - Restricted by NO_EXPORT community string.
  - AS3557 peers with a per site management AS which then peers with the local peering fabric.
- Traffic wherever possible is driven to a local node.
  - Local nodes have shorter AS paths, etc.
  - Global nodes use AS prepending.
- More info:
  - http://bit.ly/c5Nh5K (ISC blog post explaining how F is routed)

F-Root in Beijing

- ISC operates two instances of F in Beijing.
  - one sponsored by APNIC (PEK1), and another by CNNIC (PEK2).
- They both have connectivity to the national telcos and R&E networks.
- They have their own dedicated management AS.
  - AS23707 (PEK1) & AS55439 (PEK2)
- One is IPv4 only (PEK1), the other is dual stack (PEK2).
- PEK1 is connected to the Beijing NAP where we peer directly.
- PEK2 is connected via CNNIC, who pass on the route with NO_EXPORT to their peers.
- In traffic numbers, Beijing is one of our largest local nodes.

So… whaat happened?

- On October 1st (starting @17:56 UTC) there was a leak of F's IPv6 network block (2001:500:2F::/48) from PEK2 to the IPv6 Internet at large.
- The leak originated from AS37944 (CHINA SCIENCE AND TECHNOLOGY NETWORK).
- Leak propagated via AS3794 to AS6427 (Hurricane Electric) via HKIX.
- Via AS6427 to the world!
- ISC was made aware of the issue ~24 hours later (18:00 UTC on the 2nd). We pulled the IPv6 announcement from PEK2 immediately after verifying the leak.

And NANOG goes wild...

- **FACT:** There was no rewriting of answers.
- **FACT:** Affected less than 0.4% of F's traffic flows.
- **FACT:** No sign of any malicious intent to divert traffic.

- This looks to have been a simple route leak by dropping the NO_EXPORT community string.
- ISC deals with a handful of these leaks annually.

Takeaways

- There needs to be more/better BGP monitoring.
  - Esp for IPv6.
  - ISC is looking at options.
- We/ISC need more peers over IPv6.
  - http://as1280.peeringdb.com/
  - http://www.isc.org/community/peering
- The Internet needs more diverse IPv6 connections.
  - IPv6 routing table needs to be "less flat"
- Deploy DNSSEC
  - Sign your zones.
  - Enable DNSSEC validation on your recursive servers.