

10+ Years of Responsible Alternate Rootism

Paul Vixie, CEO
Farsight Security, Inc.
Yokohama, 2015-10-30

But What Is A DNS Root?

- Every zone has designated *apex* nameservers
 - Since mid-1990's there have been 13 of these
- Each such name server has some addresses
 - For example, an IPV4 and/or an IPv6 address
- The names and addresses are in *hints* file
 - But this is only used for the initial *priming* query
- Change is possible:
 - We've renamed servers (e.g., A..L)
 - We've renumbered servers

1999 or so: Wide Area Anycast

- Kato@M was already anycasting in the local area
 - Two *instances*, one data center, two IXP fabrics
- Vixie@F was already anycasting in the local area
 - Two *instances*, one OSPF mesh, used “ECMP”
- So, Vixie@F added an *instance* in Madrid
 - Had to renumber `clock.isc.org` first
 - Huge controversy! (So, added Beijing)
- Same zone, servers, addresses, operators
 - Just more reachability (thank you, BGP)

2002: AS112 Unowned Anycast

- All wide area anycast to date had used *owned* address space, with specific, known operators
- Manning@B had set up name servers for 10.in-addr.arpa (et al), for RFC 1918 PTR
 - Purpose was to draw “junk” traffic away from roots
- Vixie@F then registered AS112 and 192.175.48/24 and placed them under DNS-OARC control
 - IANA changed 10.in-addr.arpa (et al) nameservers
 - Anyone, anywhere, can operate AS112 DNS

2005: Alternate Rootism?

- Assertion: a *namespace* can be served by semantically equivalent but distinct *zones*
- So, to add IPv6, rootop Vixie suggested that IANA create an *advanced services* root zone
 - Huge controversy! (So, we didn't do it.)
- Would also have helped get IDN and DNSSEC out there faster
 - But instead, the zone remains == the namespace

~2013: ICANN ITI Panel

- Panelist Vixie suggested *hierarchical anycast* in the AS112 model, to make the root zone H.A.
 - Huge controversy!
- With DNSSEC, namespace piracy would fail
 - But, DNSSEC isn't universally deployed
 - And, there would be many unknown operators
 - Most root service would go unmeasured
 - Problems would be hard to diagnose/correct
- So, for now, the zone == the namespace, still

2015: Yeti DNS Project

- Created our own DNSSEC keys + hints file
- Operating three distribution masters (peers)
- 14 public server operators have volunteered
- A couple dozen RDNS operators have joined
- Now we can science the sh*t out of the root:
 - E.g., GOST, rapid ZSK roll, server add/delete, server renumber, hierarchical anycast, RFC 5011 KSK roll, load/stress tests, 30+ servers, *etc.*

The End

Questions, comments?