IPv6 & recursive resolvers: How do we make the transition less painful?

March 24th, 2010
IETF 77

Igor Gashinsky (igor@yahoo-inc.com)
Infrastructure Architecture

Overview of the problem

- IPv6 rollout may not impact production IPv4
- Rolling out dedicated IPv6 hostnames is not a good long-term solution
  - Good for early adopters, not good for general public
- Today, enabling AAAA on the production hostnames would adversely impact IPv4 reachability
  - 0.078% of users drop off the grid
    - Assuming a user base of 600M, that's **470K** users that you broke!
  - Additionally, client time-outs for IPv4 fallback when AAAA fails is between 21 and 186 **seconds**
  - That's a lot of breakage!
  - This is a barrier for a lot of content players

What can we do about it?

- Don't roll out IPv6
  - – Not very practical
- Roll out IPv6, accept the breakage
  - – Not very realistic
- Prefer A over AAAA
  - – This ship has already sailed, unfortunately
- Work with OS/app vendors to fix IPv6 issues
  - – Awful long lead times/upgrade cycles
- Don't let users with broken IPv6 connectivity know about AAAA records
  - – Sounds good, **how do we do that?**

How do we accomplish this?

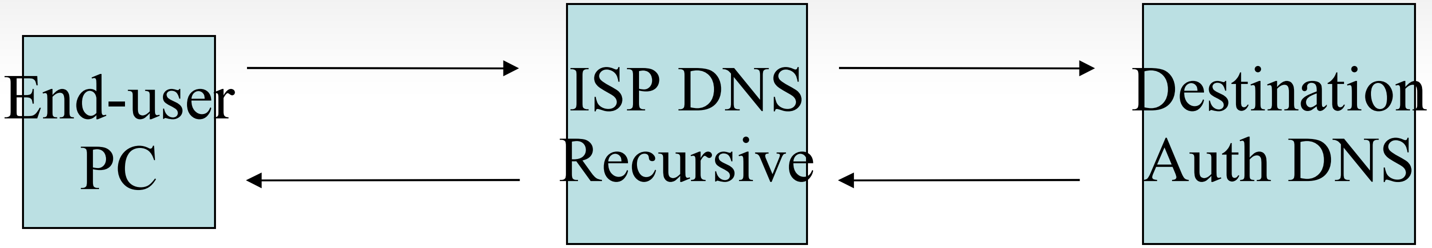

- Options:
  - Auth DNS server does not return AAAA for queries via Recursive server if it detects “broken” IPv6 users behind it
    - Requires a lot of instrumentation to set-up
    - Collateral damage for working IPv6 users
  - ISP DNS Recursive servers does not return AAAA for users who have broken IPv6 connectivity
    - How to accurately measure working users when you are not an endpoint?
  - ISP DNS recursive server only returns AAAA for users who have known working IPv6 connectivity
    - OK, sounds too good to be true, how does that work?

How do we accomplish this (2)

- Only way of **knowing** the user has working IPv6 connectivity, is if the AAAA query came over IPv6!
  - Proposed solution:
    - ISP must roll out native IPv6 on their network, and have IPv6-addressable recursive servers deployed
    - Hand out IPv6 && IPv4 recursive server addresses to the end-users
    - Return 0 answers for AAAA if, and only if:
      - Query comes over Ipv4
      - “A” record exists for same name
      - DNSSEC is not used
    - Auth DNS server now only has to worry about IPv6 reachability to the Recursive server
      - A lot easier to resolve problems at the ISP level then with individual end-users
      - A few broken IPv6 users don’t adversely impact everyone else

What does this do?

- Benefits:
  - Allows for IPv6 reachability issues to be resolved between NOCs
  - Less support calls for “what is this IPv6 thing that broke my internets?!?!?!”
  - Fewer “brokenness” with deploying IPv6 = more people may deploy it sooner
- Side-effects:
  - Trust -- now we have recursive servers modifying authoritative records
  - This effectively turns off IPv6 for OS’s that can only do DNS queries over IPv4 (ie Windows XP)

- QUESTION: Is this worth pursuing further?

Feedback from previous forums

- This idea has been presented at NANOG, ISOC IPv6 Roundtable and OARC over the past year, and the feedback that we received so far:
  - This is a **really** ugly hack.
  - People however think this **may be necessary** to get widespread IPv6 adoption
  - Needs ability to restrict behavior based on ACL
    - allow AAAA to get through for selected v4 addresses
    - stop it from getting through for selected v6 addresses

    Some of this is to make various 6RD deployments work

Status

- BIND:
  - In mainline after 9.7.0b2
  - disable-aaaa-on-v4-transport (yes | no | break-dnssec);
  - Upon receipt of a query for an AAAA record:
    - If the request has DNSSEC turned on (DO or AD bit set), return the record as requested.
    - If the request comes in over IPv6 transit, return the record as requested.
    - If the request is over IPv4 and an A record exists at the same label, respond with NOERROR but with 0 answers, forcing the client to fall back to an A record query.
- PowerDNS & Secure64
  - Also looking at implementing this

Questions?

- Email:
  - igor@yahoo-inc.com
    - or
  - jfesler@yahoo-inc.com