draft-fujiwara-dnsop-resolver-update-00

Kazunori Fujiwara
fujiwara@jprs.co.jp
IETF 97 dnsop WG

Summary of resolver-update-00

- RFC 1034 specifies parent side NS RRSet (=referrals) creates zone ‘cut’ and ‘new zone’
  - “That is, parent zones have all the information needed to access servers for their children zones.” (Quoted from RFC 1034, section 4.2.1)
- However, parent side NS RRSet may be overwritten by child zone apex NS RRSet
  - Glue records are overwritten by authoritative data
  - RFC 2181 ranking data specifies the overwrite
- Proposal: (simplified) new resolver algorithm
  - Only use referral + glue records (+ additional name resolution for out-of-bailiwick name server name) to iterate
  - Resolvers answer authoritative data only
  - (Update RFC 1034 and RFC 2181)

Problems happened by overwrite

- Unstable name resolution
  - First name resolution uses parent side NS RRSet
  - Next name resolution uses child side NS RRSet
  - If they are different, resolution results may change
- "Ghost Domain Names: Revoked Yet Still Resolvable" reported in 2012
  - Assume a resolver caches and uses zone apex NS RRSet, and the parent side NS RRSet is removed.
  - The resolver send queries to name servers specified by zone apex NS RRSet and update NS RRSet by the NS RRSet attached in the authority section of the answer.
  - Resolvers may not check the existence of the parent side NS RRSet and the domain name will remain in the resolvers
    - if the parent NS RRSet has already been removed.

Effects to existing systems/protocols

- No effect to authoritative servers
- No effect to existing resolvers
  - Gradual deployment is possible
- No effect to qname minimisation
  - because answers from authoritative servers don’t change
- No effect to DNSSEC
  - DNSSEC validates authoritative data
  - DNSSEC does not validate referrals

Comments from list, IPR

- Weak agree
  - support motivation (jinmei)
  - NS mismatch is an issue (yao)
- Child NS RRs should be used (jinmei, rharolde, ondrej)
  - if child NS RRSet is stable
- Too detail algorithm (jinmei)
- RFC 1034 has another text (marka)
- Why two caches ? (stephen, edlewis)
- How to retrieve both data (stephen)
- Some implementation does not fill authority section (ondrej)
- IPR disclosure from patent author
  - https://datatracker.ietf.org/ipr/2907/
    - (https://patents.google.com/patent/US7769826B2/)

Next steps

- Remove current detailed algorithm
- Focus on problem collection and proposal of requirements
  - Parent NS vs Child NS
  - mismatch
  - (New?) Requirement: Resolvers MUST answer authoritative data