draft-ietf-dnsop-3901bis

Status and Next Steps

Tobias Fiebig[1]

[1] Max-Planck Institut für Informatik

Momoka Yamamoto[2]

[2] Independent

RFC3901

- DNS IPv6 Transport Operational Guidelines
- Published in 2004
- Anecdotally written in response to some closed-network IPv6 root server experiments
- Bottom line:
  - All recursive DNS servers SHOULD be IPv4-only or dual-stack (IPv4 only allowed, IPv6 not required)
  - All zones SHOULD have at least one IPv4 name server (IPv6 implicitly optional)
  - Zone validation SHOULD ensure that there is one IPv4 authoritative

RFC3901 Impact

- IANA TLD requirements only mandate IPv4 (`https://www.iana.org/help/nameserver-requirements`)
- Most TLDs only mandate IPv4
- IPv6 adoption in the DNS is lacking behind (ca. 60% of zones)[a] [b]

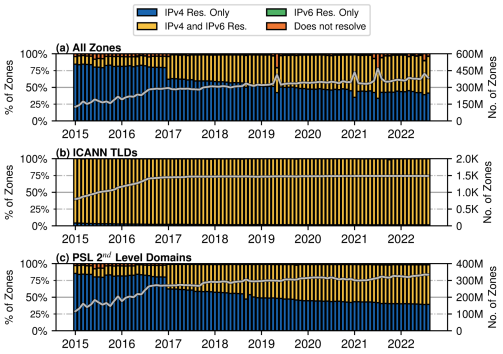

[a] Steurer, Florian, Anja Feldmann, and Tobias Fiebig. "A Tree in a Tree: Measuring Biases of Partial DNS Tree Exploration." PAM, 2025.

[b] Streibelt, Florian, et al. "How ready is DNS for an IPv6-only world?" PAM, 2023.

Why not put v4/v6 on equal footing then?

There is a 10 year+ discussion on this topic, that can be summarized as:

- IPv6 does not allow on-path fragmentation
- PMTUD is notoriously ‘difficult’
- There are high fragment drop rates for IPv6 fragments on the Internet, making DNS-over-UDP break for (especially) DNSSEC (for `DNSKEY` records)
- TCP fallback does not always work
- DNSSEC needs even more TCP fallback (for `DNSKEY` records)
- See also (from 2015): https://insinuator.net/2015/11/some-notes-on-the-drop-ipv6-fragments-vs-this-will-break-dnssec-debate/

Comments during DNSOP at 121 and on the ML echo the same concern w.r.t. fragmentation being an issue for IPv6.

Measuring DNS Resolution

- Resolve the Google CrUX Top 10M IPv4 only/IPv6 only/dual-stack
- Use a ‘minimally covering NS set’, i.e., the set of NS sets in which each NS set is represented with one zone (1x DNSSEC, 1x without DNSSEC)
- Implement various MTU scenarios:
  - MTU1500
  - MTU1280 on-link (with working PMTUD)
  - MTU1280 on-link (with broken PMTUD)
  - MTU1280 on-path (with broken PMTUD)
- Once via specific Tier-1, once via other upstreams

! All data is available on https://dns-mtu.measurement.network/ (past 7 days) and https://data.measurement.network/

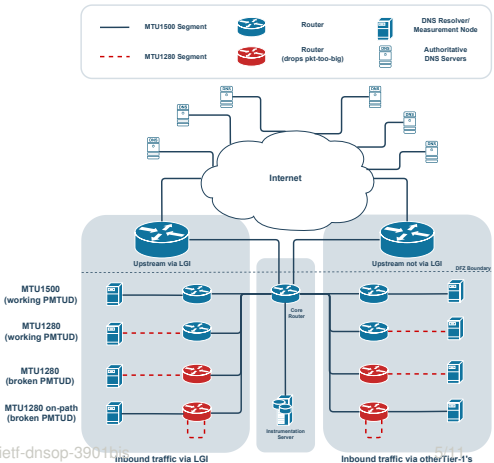

IPv4 vs. IPv6

- Generally comparable (usually within ca. 3%)
- `SERVFAIL` for No DNSSEC around 33%, but $\leq$ 10% for DNSSEC
- Some IPv6 packets get fragment-header-sized bigger (UDP packet too big @EDNS0 1232)
- A small ($\leq$0.15%) baseline of broken IPv6 NS sets
- *More* IPv6 fragments for IPv6 @EDNS0 4096 for DNSSEC
- Given broken PMTUD, v4 has more `TIMEOUT`s with DNSSEC @EDNS0 1232

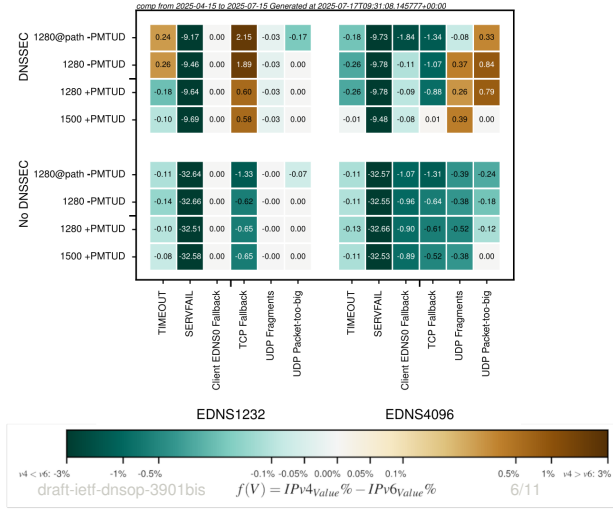

Differences in Perspective: DNSSEC

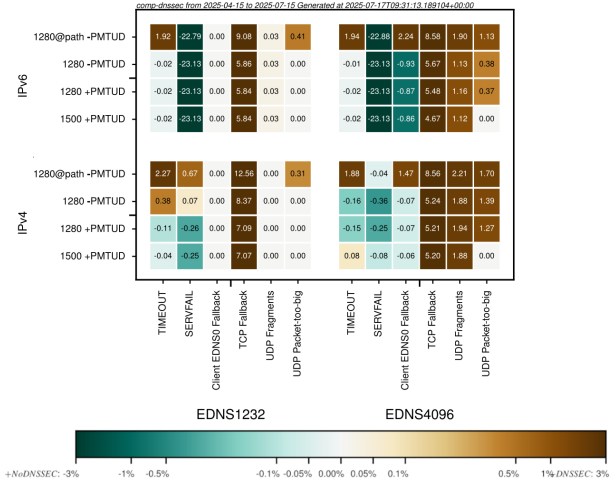

- Significantly higher TCP fallbacks (as expected) for v4 and v6
- Less `SERVFAIL` for DNSSEC
- More fragments than for No DNSSEC

High Fragment Drop Rates

- Discussed first measurement data at IEPG during 122
- Identified a Tier-1 back then dropping IPv6 fragments on transit
- Tier-1 put in *a lot* of effort to have this fixed before the next IETF ;-)
- It seems like there are slightly more IPv6 fragments for broken MTU cases

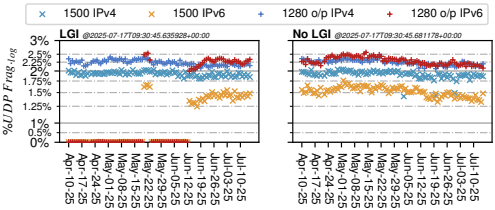

Differences in Perspective: No Fragments

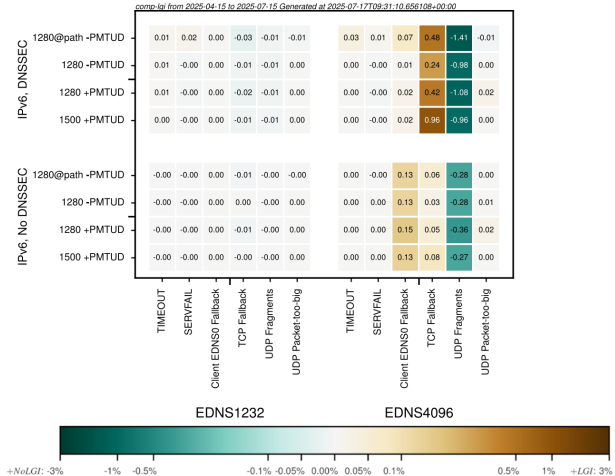

- Following best practices (EDNS0 at 1232) no notable difference
- Mildly ($\leq$1%) more TCP fallbacks
- A bit more EDNS0 fallback ($\leq$0.10%); Even a bit more for No DNSSEC ($\leq$0.15%)
- Difference in `TIMEOUT`s is marginal; $\leq$0.5% for an artificially created MTU blackhole without PMTUD with DNSSEC @EDNS0 4096

Summary

- Fragmentation does not play a major role in IPv6 for the DNS
- Misconfigured IPv6 DNS has *an order of magnitude* more impact on `TIMEOUT`s than fragment dropping
  - Better guidance on use of IPv6 (SHOULD) and authoritative testing (SHOULD) should help here
- IPv4 does not necessarily ‘perform better’ for DNS resolution
- People following RFC9715, doing IPv6 for DNS, and running NS sets on which at least one domain has DNSSEC seem to coincide

State of the Draft / Discussion Points

- States SHOULD for v4 *and* v6 on authoritatives and for zone testing
- Explicitly referenced RFC9715, also highlighting the implied max MSS that SHOULD be used for v4 and v6
- Also: Added guidance for (stub) resolvers:
  - Allow IPv6 only cases if a transition technology is used ensuring IPv4 resolvability of zones
  - Guidance for stub resolvers to prefer native IPv6 resolvers over PREF64 synthesized ones if both are available

Question: What to add, what to remove, and how should we proceed?