DNS IPv6 Transport Operational Guidelines

The (IPv6) future is now

Momoka Yamamoto[1]

[1]WIDE Project

Tobias Fiebig[2]

[2]Max-Planck Institut für Informatik

2024-11-04

Recap: RFC3901

- Goal: Prevent Namespace Fragmentation due to IPv6 Islands
  - *“every recursive name server SHOULD be either IPv4-only or dual stack”*
  - *“every DNS zone SHOULD be served by at least one IPv4-reachable authoritative name server”*
- Problems:
  - If you like it or not, IPv6 is ‘there’
  - With more and more IPv6 mostly/IPv6 only networks, we may as much have IPv4 islands as we were afraid of IPv6 only islands in 2004
  - The recommendations trickle down through the stack (see e.g. auth. server requirements)

draft-momoka-dnsop-3901bis-06

- Goal: Prevent Namespace Fragmentation due to AFI Islands
  - every recursive name server SHOULD be dual stack
  - every DNS zone SHOULD be served by at least one IPv4-reachable and one IPv6-reachable authoritative name server
    - ... and the two should be consistent
- Several comments and suggestions from the ML integrated
- Problems:
  - Yes
- Prefer IPv6 over IPv4:
  - No.This is about making both equal netizens in the DNS.

draft-momoka-dnsop-3901bis-06

- Goal: Prevent Namespace Fragmentation due to AFI Islands
  - every recursive name server SHOULD be dual stack
  - every DNS zone SHOULD be served by at least one IPv4-reachable and one IPv6-reachable authoritative name server
    - ... and the two should be consistent
- Several comments and suggestions from the ML integrated
- Problems:
  - Yes
- Prefer IPv6 over IPv4:
  - No.This is about making both equal netizens in the DNS.

draft-momoka-dnsop-3901bis-06

- Goal: Prevent Namespace Fragmentation due to AFI Islands
  - every recursive name server SHOULD be dual stack
  - every DNS zone SHOULD be served by at least one IPv4-reachable and one IPv6-reachable authoritative name server
    - ... and the two should be consistent
- Several comments and suggestions from the ML integrated
- Problems:
  - Yes
- Prefer IPv6 over IPv4:
  - No.This is about making both equal netizens in the DNS.

draft-momoka-dnsop-3901bis-06

- Goal: Prevent Namespace Fragmentation due to AFI Islands
  - every recursive name server SHOULD be dual stack
  - every DNS zone SHOULD be served by at least one IPv4-reachable and one IPv6-reachable authoritative name server
    - ... and the two should be consistent
- Several comments and suggestions from the ML integrated
- Problems:
  - Yes
- Prefer IPv6 over IPv4:
  - No.This is about making both equal netizens in the DNS.

draft-momoka-dnsop-3901bis-06

- Goal: Prevent Namespace Fragmentation due to AFI Islands
  - every recursive name server SHOULD be dual stack
  - every DNS zone SHOULD be served by at least one IPv4-reachable and one IPv6-reachable authoritative name server
    - ... and the two should be consistent
- Several comments and suggestions from the ML integrated
- Problems:
  - Yes
- Prefer IPv6 over IPv4:
  - No.This is about making both equal netizens in the DNS.

draft-momoka-dnsop-3901bis-06

- Goal: Prevent Namespace Fragmentation due to AFI Islands
  - every recursive name server SHOULD be dual stack
  - every DNS zone SHOULD be served by at least one IPv4-reachable and one IPv6-reachable authoritative name server
    - ... and the two should be consistent
- Several comments and suggestions from the ML integrated
- Problems:
  - Yes
- Prefer IPv6 over IPv4:
  - No. This is about making both equal netizens in the DNS.

draft-momoka-dnsop-3901bis-06

- Goal: Prevent Namespace Fragmentation due to AFI Islands
  - every recursive name server SHOULD be dual stack
  - every DNS zone SHOULD be served by at least one IPv4-reachable and one IPv6-reachable authoritative name server
    - ... and the two should be consistent
- Several comments and suggestions from the ML integrated
- Problems:
  - Yes
- Prefer IPv6 over IPv4:
  - No. This is about making both equal netizens in the DNS.

draft-momoka-dnsop-3901bis-06

- Goal: Prevent Namespace Fragmentation due to AFI Islands
  - every recursive name server SHOULD be dual stack
  - every DNS zone SHOULD be served by at least one IPv4-reachable and one IPv6-reachable authoritative name server
    - ... and the two should be consistent
- Several comments and suggestions from the ML integrated
- Problems:
  - Yes
- Prefer IPv6 over IPv4:
  - No**.**This is about making both equal netizens in the DNS.

IPv6 DNS is *slow*

- Yes. No. Maybe;
- Steurer, Florian, et al. *The Roots Go Deep: Measuring '.' Under Change.*, Proc. of the ACM Internet Measurement Conference, 2024
  - Routing can have an impact on RTT to NS, esp. when anycasted; Sometimes IPv4 is faster... sometimes IPv6.
- Streibelt, Florian, et al. *How Ready is DNS for an IPv6-Only World?* Conference on Passive and Active Network Measurement, 2023
  - Misocnfigured IPv6 AUTH exist, having to fail-back increases resolution time
- Solutions:
  - Start running it to find and fix the gaps

IPv6 DNS is *slow*

- Yes. No. Maybe;
- Steurer, Florian, et al. *The Roots Go Deep: Measuring '.' Under Change.*, Proc. of the ACM Internet Measurement Conference, 2024
  - Routing can have an impact on RTT to NS, esp. when anycasted; Sometimes IPv4 is faster... sometimes IPv6.
- Streibelt, Florian, et al. *How Ready is DNS for an IPv6-Only World?* Conference on Passive and Active Network Measurement, 2023
  - Misocnfigured IPv6 AUTH exist, having to fail-back increases resolution time
- Solutions:
  - Start running it to find and fix the gaps

IPv6 DNS is *slow*

- Yes. No. Maybe;
- Steurer, Florian, et al. *The Roots Go Deep: Measuring '.' Under Change.*, Proc. of the ACM Internet Measurement Conference, 2024
  - Routing can have an impact on RTT to NS, esp. when anycasted; Sometimes IPv4 is faster... sometimes IPv6.
- Streibelt, Florian, et al. *How Ready is DNS for an IPv6-Only World?* Conference on Passive and Active Network Measurement, 2023
  - Misocnfigured IPv6 AUTH exist, having to fail-back increases resolution time
- Solutions:
  - Start running it to find and fix the gaps

UDP and IPv6

- The history of IPv6 and PMTUD is... *“a story full of misunderstandings...”*
- There are a lot of ways to hold IPv6 *wrong*
- Solutions:
  - I-D.ietf-dnsop-avoid-fragmentation
  - Start doing things so things get fixed
  - TCP

UDP and IPv6

- The history of IPv6 and PMTUD is... *“a story full of misunderstandings...”*
- There are a lot of ways to hold IPv6 *wrong*
- Solutions:
  - I-D.ietf-dnsop-avoid-fragmentation
  - Start doing things so things get fixed
  - TCP

UDP and IPv6

- The history of IPv6 and PMTUD is... *“a story full of misunderstandings...”*
- There are a lot of ways to hold IPv6 *wrong*
- Solutions:
  - I-D.ietf-dnsop-avoid-fragmentation
  - Start doing things so things get fixed
  - TCP

UDP and Transport Protocols

- DoU, DoTCP, DoT, DoQ, DoH, DoHQ, DoCarrierPidgeon
- Happy Eyeballs has some ideas about AFI selection
- HEv3 draft-pauly-v6ops-happy-eyeballs-v3-02 has some about transport selection
- DNS tends to have the idea that HE should not have so many ideas
- HE-approach does not work for DNS
- Solutions:
  - Also add transport/AFI selection/preference to the document (We're gonna do our OWN HE, with(out) fragmentation and browsers.)
  - Expand the BCP?

UDP and Transport Protocols

- DoU, DoTCP, DoT, DoQ, DoH, DoHQ, DoCarrierPidgeon
- Happy Eyeballs has some ideas about AFI selection
- HEv3 draft-pauly-v6ops-happy-eyeballs-v3-02 has some about transport selection
- DNS tends to have the idea that HE should not have so many ideas
- HE-approach does not work for DNS
- Solutions:
  - Also add transport/AFI selection/preference to the document (We're gonna do our OWN HE, with(out) fragmentation and browsers.)
  - Expand the BCP?

UDP and Transport Protocols

- DoU, DoTCP, DoT, DoQ, DoH, DoHQ, DoCarrierPidgeon
- Happy Eyeballs has some ideas about AFI selection
- HEv3 draft-pauly-v6ops-happy-eyeballs-v3-02 has some about transport selection
- DNS tends to have the idea that HE should not have so many ideas
- HE-approach does not work for DNS
- Solutions:
  - Also add transport/AFI selection/preference to the document (We're gonna do our OWN HE, with(out) fragmentation and browsers.)
  - Expand the BCP?

UDP and Transport Protocols

- DoU, DoTCP, DoT, DoQ, DoH, DoHQ, DoCarrierPidgeon
- Happy Eyeballs has some ideas about AFI selection
- HEv3 draft-pauly-v6ops-happy-eyeballs-v3-02 has some about transport selection
- DNS tends to have the idea that HE should not have so many ideas
- HE-approach does not work for DNS
- Solutions:
  - Also add transport/AFI selection/preference to the document (We're gonna do our OWN HE, with(out) fragmentation and browsers.)
  - Expand the BCP?

Expand the BCP?!

- Resolvers are different then AuthNS
- Transport is different from Network
- Selection might differ
- Small badges are better (and easier to maintain) than an all-eggs-basket
- There are additional things that belong into the same ... basket (I-D.ietf-dnsop-avoid-fragmentation)
- Options:
  - Keep things as they are, and update RFC3901
  - Take a comprehensive approach...

Expand the BCP?!

- Resolvers are different then AuthNS
- Transport is different from Network
- Selection might differ
- Small badges are better (and easier to maintain) than an all-eggs-basket
- There are additional things that belong into the same ... basket (I-D.ietf-dnsop-avoid-fragmentation)
- Options:
  - Keep things as they are, and update RFC3901
  - Take a comprehensive approach...

Expand the BCP?!

- Resolvers are different then AuthNS
- Transport is different from Network
- Selection might differ
- Small badges are better (and easier to maintain) than an all-eggs-basket
- There are additional things that belong into the same ... basket (I-D.ietf-dnsop-avoid-fragmentation)
- Options:
  - Keep things as they are, and update RFC3901
  - Take a comprehensive approach...

We have ‘comprehensive approach’ at home...

Eggs... er... Documents to put into the bask... er... BCP

- ID.1 Guidance on IP AFI Support and Selection for Authoritative Nameservers
- ID.2 Guidance on IP AFI Support and Selection for Recursive and Stub Resolvers
- ID.3 Guidance on Transport Protocol Support and Selection for Authoritative Nameservers
- ID.4 Guidance on Transport Protocol Support and Selection for Recursive and Stub Resolvers
- ID.5 I-D.ietf-dnsop-avoid-fragmentation
- ID.? ...

If we’d want this (and I kind of like it), this needs more people to get writing...

Should we adopt or should we draft... *sing*

Summary:

- BCP91 needs some dusting
- There are problems with IPv6 in DNS, but we will not get rid of them without some changes in our practices
- draft-momoka-dnsop-3901bis-06 moves in that general direction in a ‘just update RFC3901’-way
- We could also get out the big bat and hit on problems beyond AFI (Transport etc.)

Options:

- Adopt draft-momoka-dnsop-3901bis-06 and update RFC3901[1]
- Split into smaller documents and fill BCP91 with a couple more atomic documents[2]

[1] This is an explicit request for a call for adoption :-)

[2] If you like this option, ask yourself: Which part would *I* contribute?

Should we adopt or should we draft... *sing*

Summary:

- BCP91 needs some dusting
- There are problems with IPv6 in DNS, but we will not get rid of them without some changes in our practices
- draft-momoka-dnsop-3901bis-06 moves in that general direction in a ‘just update RFC3901’-way
- We could also get out the big bat and hit on problems beyond AFI (Transport etc.)

Options:

- Adopt draft-momoka-dnsop-3901bis-06 and update RFC3901[1]
- Split into smaller documents and fill BCP91 with a couple more atomic documents[2]

[1] This is an explicit request for a call for adoption :-)

[2] If you like this option, ask yourself: Which part would *I* contribute?