dnsop
Internet-Draft
Obsoletes: 3901 (if approved)
Intended status: Best Current Practice
Expires: 29 January 2026

Momoka. Y
WIDE Project
T. Fiebig
MPI-INF

28 July 2025

DNS IPv6 Transport Operational Guidelines
draft-ietf-dnsop-3901bis-03

Abstract

This memo provides guidelines and documents Best Current Practice for operating authoritative DNS servers as well as recursive and stub DNS resolvers, given that queries and responses are carried in a mixed environment of IPv4 and IPv6 networks. This document expands on RFC 3901 by recommending that authoritative DNS servers as well as recursive DNS resolvers support both IPv4 and IPv6. It furthermore provides guidance for how recursive DNS resolver should select upstream DNS servers, if synthesized and non-synthesized IPv6 addresses are available.

This document obsoletes RFC3901. (if approved)

Discussion Venues

This note is to be removed before publishing as an RFC.

Source for this draft and an issue tracker can be found at https://github.com/ietf-wg-dnsop/draft-ietf-dnsop-3901bis.

Status of This Memo

Table of Contents

1. Introduction

Despite IPv6 being first discussed in the mid-1990s [RFC1883], consistent deployment throughout the whole Internet has not yet been accomplished [RFC9386]. Hence, today, the Internet is a mixture of IPv4, dual-stack (networks connected via both IP versions), and IPv6 networks.

This creates a complex landscape where authoritative DNS servers might be accessible only via specific network protocols [V6DNSRDY-23]. At the same time, DNS resolvers may only be able to access the Internet via either IPv4 or IPv6. This poses a challenge for such resolvers because they may encounter names for which queries must be directed to authoritative DNS servers with which they do not share an IP version during the name resolution process.

[RFC3901] was initially written at a time when IPv6 deployment was not widespread, focusing primarily on maintaining name space continuity within the IPv4 landscape. Now, nearly two decades later, with IPv6 not only widely deployed but also becoming the de facto standard in many areas, this document seeks to expand the scope of RFC3901 by recommending IPv6 compatibility for authoritative DNS servers, as well as recursive and stub DNS resolvers.

This document provides guidance on:

* IP version related name space fragmentation and best-practices for avoiding it.

* Guidelines for configuring authoritative DNS servers for zones.

* Guidelines for operating recursive DNS resolvers.

* Guidelines for stub DNS resolvers.

While transitional technologies and dual-stack setups may mitigate some of the issues of DNS resolution in a mixed protocol-version Internet, making DNS data accessible over both IPv4 and IPv6 is the most robust and flexible approach, as it allows resolvers to reach the information they need without requiring intermediary translation or forwarding services which may introduce additional failure cases.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Terminology

This document uses DNS terminology as described in [RFC9499]. Furthermore, the following terms are used with a defined meaning:

IPv4 name server:
A name server providing DNS services reachable via IPv4. It does not imply anything about what DNS data is served, but requires DNS queries to be received and answered over IPv4.

IPv6 name server:
A name server providing DNS services reachable via IPv6. It does not imply anything about what DNS data is served, but requires DNS queries to be received and answered over IPv6.

Dual-stack name server:
A name server that is both an "IPv4 name server" and also an "IPv6 name server".

3. Name Space Fragmentation

A resolver that tries to look up a name starts out at the root, and follows referrals until it is referred to a name server that is authoritative for the name. If somewhere down the chain of referrals it is referred to a name server that is, based on the referral, only accessible over a transport which the resolver cannot use, the resolver is unable to continue DNS resolution.

If this occurs, the DNS has, effectively, fragmented based on the recursive DNS resolver’s and authoritative DNS server’s mismatching IP version support.

In a mixed IP Internet, namespace fragmentation can occur for different reasons. One reason is that DNS zones are consistently configured to support only either IPv4 or IPv6. Another reason is due to misconfigurations that make a zone unresolvable by either IPv4 or IPv6-only resolvers. The latter cases are often hard to identify, as the impact of misconfigurations for only one IP version (IPv4 or IPv6) may be hidden in a dual-stack setting. In the worst case, a specific name may only be resolvable via dual-stack enabled resolvers.

3.1. Misconfigurations Causing IP Version Related Name Space Fragmentation

Even when an administrator assumes that they have enabled support for a specific IP version on their authoritative DNS server, various misconfigurations may break the DNS delegation chain of a zone for that protocol and prevent any of its records from resolving for clients only supporting that IP version. These misconfigurations can be kept hidden if most clients can successfully fall back to the other IP version.

The following name related misconfigurations can cause broken delegation for one IP version:

No A/AAAA records for NS names:
If all of the NS records for a zone in their parent zone have either only A records or only AAAA records, then resolution via the other IP version is not possible.

Missing GLUE:
If the name from an NS record for a zone is in-domain, i.e., the name is within the zone or below, a parent zone must contain both IPv4 and IPv6 GLUE records, i.e., a parent must serve the corresponding A and AAAA records as ADDITIONAL data when returning the NS record(s) as the referral response.

No A/AAAA record for in-domain NS:
If an NS record of a child zone, either provided by the parent or from the child zone’s apex, points to a name in the NS RDATA that is in-domain but the name does not contain corresponding A or AAAA record(s) in the child zone, name resolution via the concerned IP version will fail even if the parent provides GLUE, when the recursive DNS resolver revalidates the delegation path [I-D.ietf-dnsop-ns-revalidation].

Zone of sibling domain NSes not resolving:
If the name from an NS record for a zone is sibling domain, the corresponding zone must be resolvable via the IP version in question as well. It is insufficient if the name pointed to by the NS record has an associated A or AAAA record correspondingly.

Parent zone not resolvable via one IP version:
For a zone to be resolvable via an IP version the parent zones up to the root zone must be resolvable via that IP version as well. Any zone not resolvable via the concerned IP version breaks the delegation chain for all its children.

The above misconfigurations are not mutually exclusive.

Furthermore, any of the misconfigurations above may not only materialize via a missing Resource Record (RR) but also via an RR providing the IP address of a nameserver that is not configured to answer queries via that IP version [V6DNSRDY-23].

3.2. Network Conditions Causing IP Version Related Name Space Fragmentation

In addition to explicit misconfigurations in the served DNS zones, network conditions may also influence a resolver's ability to resolve names in a zone. The most common issue here are packets requiring fragmentation given a reduced path MTU (PMTU) and MTU blackholes, i.e., packets being dropped on-path due to exceeding the MTU of the link to the next-hop without the sender being notified. This can manifest in the following way:

DNS-over-UDP packets requiring fragmentation
: When using EDNS(0) to communicate support for DNS messages larger than 512 bytes [RFC6891], an IP packet carrying a DNS response may exceed the PMTU for the path to a resolver. If an authoritative DNS server and does not follow [RFC9715], i.e., honors EDNS(0) sizes larger than 1232 bytes, it will try to fragment the packet according to the discovered PMTU. Such packets mostly occur for DNSKEY responses with DNSSEC [RFC4034].

  If the requesting resolver is unable to process fragments, or if fragments are filtered on-path, resolution will fail over UDP. These issues are more prevalent for IPv6, as it no longer allows on-path hosts to fragment packets. Therefore, working Path MTU Discovery (PMTUD) is essential for IPv6 DNS-over-UDP packets to be fragmented to a size that allows them to traverse all segments on a path.

  [RFC9715] provides guidance on preventing this issue by always using a maximum EDNS(0) size of 1232 bytes.

DNS-over-TCP packets requiring fragmentation
: If DNS resolution over UDP fails, or if a packet exceeds the communicated EDNS(0) size, a resolver should fall back to DNS resolution over TCP. However, similar to the case of DNS-over-UDP, DNS-over-TCP may encounter MTU blackholes, especially on IPv6, if PMTUD does not work, if the MSS honored by the authoritative DNS server leads to IP packets exceeding the PMTU. In that case, similar to the case of DNS-over-UDP, DNS resolution will time out when the recursive DNS resolver did not receive a response in time.

  [RFC9715] does not provide explicit guidance on mitigating this issue. However, similar to the guidance in [RFC9715], setting an MSS of 1240 bytes for IPv4 and 1220 bytes for IPv6 would similarly mitigate this issue.

Broken IPv6 Connectivity at the Resolver

Similar to authoritative servers, (stub) recursive resolvers may face broken IPv6 connectivity, e.g., if a client has been assigned a global unicast IPv6 address, but IPv6 traffic is not routed on the resolver's network. Furthermore, broken IPv6 connectivity may be encountered when IPv4-IPv6 transition technologies, e.g., NAT64 on IPv6-mostly networks [RFC9313] are in use. There, the synthesized IPv6 addresses used in XLAT encounter additional PMTU fluctuation due to the difference in header size between IPv4 and IPv6.

3.3. Reasons for Intentional IP Version Related Name Space Fragmentation

Intentional IP related name space fragmentation occurs if an operator consciously decides not to deploy IPv4 or IPv6 for a part of the resolution chain. Most commonly, this is realized by intentionally not listing A/AAAA records for NS names. At the time of writing, the share of zones not resolvable via IPv4 is negligible, while a little less than 40% of zones are not resolvable via IPv6 [V6DNSRDY-23]. However, as IPv4 exhaustion progresses, IPv6 adoption will have to increase.

4. Policy Based Avoidance of Name Space Fragmentation

With the final exhaustion of IPv4 pools in RIRs, e.g., [RIPEV4], and the progressing deployment of IPv6, there no longer is a "preferred" IP version. Yet, while we now observe the first zones becoming exclusively IPv6 resolvable, we also still see a major portion of zones solely relying on IPv4 [V6DNSRDY-23]. Hence, at the moment, dual stack connectivity is instrumental to be able to resolve zones and avoid name space fragmentation.

Having zones served only by name servers reachable via one IP version would fragment the DNS. Hence, we need to find a way to avoid this fragmentation.

The recommended approach to maintain name space continuity is to use administrative policies, as described in this section.

4.1. Guidelines for Authoritative DNS Server Configuration

It is usually recommended that DNS zones contain at least two name servers, which are geographically diverse and operate under different routing policies [IANANS]. To reduce the chance of DNS name space fragmentation, it is RECOMMENDED that at least two name servers for a zone are dual stack name servers. Specifically, this means that the following minimal requirements SHOULD be implemented for a zone:

IPv4 adoption:
: Every DNS zone SHOULD be served by at least one IPv4-reachable authoritative DNS server to maintain name space continuity. The delegation configuration (Resolution of the parent, resolution of sibling domain names, GLUE) MUST NOT rely on IPv6 connectivity being available. As we acknowledge IPv4 scarcity, operators MAY opt not to provide DNS services via IPv4, if they can ensure that all clients expected to resolve this zone do support DNS resolution via IPv6.

IPv6 adoption:
: Every DNS zone SHOULD be served by at least one IPv6-reachable authoritative DNS server to maintain name space continuity. The delegation configuration (Resolution of the parent, resolution of sibling domain names, GLUE) MUST NOT rely on IPv4 connectivity being available.

Consistency:
: Both IPv4 and IPv6 transports should serve identical DNS data to ensure a consistent resolution experience across different network types.

Avoiding IP Fragmentation:
: IP fragmentation has been reported to be fragile [RFC8900]. Furthermore, IPv6 transition technologies can introduce unexpected MTU breaks, e.g., when NAT64 is used [RFC7269]. Therefore, IP fragmentation SHOULD be avoided by following guidance on maximum DNS payload sizes [RFC9715] and providing TCP fall-back options [RFC7766]. Furthermore, similar to the guidance in [RFC9715], it is RECOMMENDED that authoritative DNS servers sets an MSS of 1240b for IPv4 and 1220b for IPv6 in TCP sessions carrying DNS responses.

Note: To prevent namespace fragmentation zone validation processes SHOULD ensure that:

* There is at least one IPv4 address record and one IPv6 address record available for the name servers of any child delegation within the zone.

* The zone’s authoritative servers follow [RFC9715] for avoiding fragmentation on DNS-over-UDP.

* The zone’s authoritative servers support DNS-over-TCP [RFC9210].

* The zone’s authoritative servers can be reached via IPv4 and IPv6 when when performing DNS resolution via IPv4-only and IPv6-only networks.

4.2. Guidelines for DNS Resolvers

Every recursive DNS resolver SHOULD be dual stack.

While the zones that IPv6-only recursive DNS resolvers can resolve are growing, they do not yet cover all zones. Hence, a recursive DNS resolver MAY be IPv6-only, if it uses a transition mechanism that allows it to also query IPv4-only authoritative DNS server or uses a configuration where it forwards queries failing IPv6-only DNS resolution to a recursive DNS resolver that is able to perform DNS resolution over IPv4.

If a recursive DNS resolver runs in a network that uses XLAT [RFC6877], and the recursive DNS resolver is aware of the used PREF64 [RFC6146], it SHOULD synthesize mapped IPv6 addresses for remote authoritative DNS servers directly for DNS resolution, instead of relying on the socket translation layer of the operating system. A recursive DNS resolver SHOULD prefer non-synthesized IPv6 addresses over synthesized IPv6 addresses based on a PREF64. Additionally, the PREF64 in use MAY also be statically configured for the DNS resolver.

Similarly, a recursive DNS resolver MAY be IPv4-only, if it uses a configuration where such resolvers forward queries failing IPv4-only DNS resolution to a recursive DNS resolver that is able to perform DNS resolution over IPv6.

Finally, when responding to recursive queries sent by stub DNS resolvers, a DNS resolver SHOULD follow the above guidance for communication between authoritative DNS servers and recursive DNS resolvers analogously.

4.3. Guidelines for DNS Stub Resolvers

In general, DNS Stub Resolvers SHOULD follow the same guidance as outlined for recursive DNS resolvers when they are deployed to a dual-stack network not using IPv4-IPv6 transition techniques. Contrary to authoritative DNS servers and recursive DNS resolvers, stub DNS resolvers are more likely to find themselves in either an IPv6 mostly or IPv4 only environment, as they are usually run on end-hosts / clients.

For IPv4 only environments, a stub DNS resolver has to rely on the provided recursive DNS server following guidance in this document. However, in an IPv6 mostly scenario, the environment might appear similar to a dual-stack scenario, and the host running the stub DNS resolver may receive multiple IPv4 and IPv6 addresses for possible DNS resolvers to use via different protocols (DHCPv4, DHCPv6, SLAAC).

Hence, when a host running a stub DNS resolver receives addresses for IPv4 and IPv6 recursive DNS resolver to use, it SHOULD prioritize reachable IPv6 recursive DNS resolvers. If the host is aware of a PREF64 being used, it SHOULD NOT use IPv4 recursive DNS resolvers for which it also received the corresponding mapped address in the PREF64, but instead SHOULD only use the supplied IPv6 address. Additionally, if the host receives multiple IPv6 reachable recursive DNS resolvers and is aware of a PREF64 being in use, it SHOULD prioritize recursive DNS resolvers outside the PREF64.

5. Security Considerations

The guidelines described in this memo introduce no new security considerations into the DNS protocol or associated operational scenarios.

6. IANA Considerations

This document requests IANA to update its technical requirements for authoritative DNS servers to require both IPv4 and IPv6 addresses for each authoritative server [IANANS].

Acknowledgments

Valuable input for this draft was provided by: Bob Harold, Andreas Schulze, Tommy Jensen, Nick Buraglio, Jen Linkova, Tim Chown, Brian E Carpenter, Tom Petch

Thank you for reading this draft.

References

Normative References

[I-D.ietf-dnsop-ns-revalidation]
Huque, S., Vixie, P. A., and W. Toorop, "Delegation Revalidation by DNS Resolvers", Work in Progress, Internet-Draft, draft-ietf-dnsop-ns-revalidation-10, 25 June 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-ns-revalidation-10>.

[RFC1883] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 1883, DOI 10.17487/RFC1883, December 1995, <https://www.rfc-editor.org/info/rfc1883>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3901] Durand, A. and J. Ihren, "DNS IPv6 Transport Operational Guidelines", BCP 91, RFC 3901, DOI 10.17487/RFC3901, September 2004, <https://www.rfc-editor.org/info/rfc3901>.

[RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, DOI 10.17487/RFC4034, March 2005, <https://www.rfc-editor.org/info/rfc4034>.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <https://www.rfc-editor.org/info/rfc6146>.

[RFC6877] Mawatari, M., Kawashima, M., and C. Byrne, "464XLAT: Combination of Stateful and Stateless Translation", RFC 6877, DOI 10.17487/RFC6877, April 2013, <https://www.rfc-editor.org/info/rfc6877>.

[RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <https://www.rfc-editor.org/info/rfc6891>.

[RFC7766] Dickinson, J., Dickinson, S., Bellis, R., Mankin, A., and D. Wessels, "DNS Transport over TCP - Implementation Requirements", RFC 7766, DOI 10.17487/RFC7766, March 2016, <https://www.rfc-editor.org/info/rfc7766>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC9210] Kristoff, J. and D. Wessels, "DNS Transport over TCP - Operational Requirements", BCP 235, RFC 9210, DOI 10.17487/RFC9210, March 2022, <https://www.rfc-editor.org/info/rfc9210>.

[RFC9715] Fujiwara, K. and P. Vixie, "IP Fragmentation Avoidance in DNS over UDP", RFC 9715, DOI 10.17487/RFC9715, January 2025, <https://www.rfc-editor.org/info/rfc9715>.

Informative References

[I-D.ietf-v6ops-ipv6-only-resolver]
Yamamoto, M. and Y. Toyota, "IPv6-only Capable Resolvers Utilising NAT64", Work in Progress, Internet-Draft, draft-ietf-v6ops-ipv6-only-resolver-00, 23 October 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-v6ops-ipv6-only-resolver-00>.

[IANANS] IANA, "Technical requirements for authoritative name servers", <https://www.iana.org/help/nameserver-requirements>.

[RFC7269] Chen, G., Cao, Z., Xie, C., and D. Binet, "NAT64 Deployment Options and Experience", RFC 7269, DOI 10.17487/RFC7269, June 2014, <https://www.rfc-editor.org/info/rfc7269>.

[RFC8900] Bonica, R., Baker, F., Huston, G., Hinden, R., Troan, O., and F. Gont, "IP Fragmentation Considered Fragile", BCP 230, RFC 8900, DOI 10.17487/RFC8900, September 2020, <https://www.rfc-editor.org/info/rfc8900>.

[RFC9313] Lencse, G., Palet Martinez, J., Howard, L., Patterson, R., and I. Farrer, "Pros and Cons of IPv6 Transition Technologies for IPv4-as-a-Service (IPv4aaS)", RFC 9313, DOI 10.17487/RFC9313, October 2022, <https://www.rfc-editor.org/info/rfc9313>.

[RFC9386] Fioccola, G., Volpato, P., Palet Martinez, J., Mishra, G., and C. Xie, "IPv6 Deployment Status", RFC 9386, DOI 10.17487/RFC9386, April 2023, <https://www.rfc-editor.org/info/rfc9386>.

[RFC9499] Hoffman, P. and K. Fujiwara, "DNS Terminology", BCP 219, RFC 9499, DOI 10.17487/RFC9499, March 2024, <https://www.rfc-editor.org/info/rfc9499>.

[RIPEV4] RIPE NCC, "The RIPE NCC has run out of IPv4 Addresses", November 2019, <https://www.ripe.net/publications/news/about-ripe-ncc-and-ripe/the-ripe-ncc-has-run-out-of-ipv4-addresses>.

[V6DNSRDY-23]
Streibelt, F., Sattler, P., Lichtblau, F., Hernandez-Gan, C., Gasser, O., and T. Fiebig, "How Ready is DNS for an IPv6-Only World?", March 2023, <https://link.springer.com/chapter/10.1007/978-3-031-28486-1_22>.

Authors' Addresses

Momoka Yamamoto
WIDE Project
Email: momoka.my6@gmail.com

Additional contact information:

山本 桃歌
WIDE Project

Tobias Fiebig
Max-Planck-Institut fuer Informatik
Campus E14
66123 Saarbruecken
Germany
Phone: +49 681 9325 3527
Email: tfiebig@mpi-inf.mpg.de