dnsop
Internet-Draft
Obsoletes: 3901 (if approved)
Intended status: Best Current Practice
Expires: 15 August 2026

Momoka
WIDE Project
T. Fiebig
MPI-INF

11 February 2026

# Operational Guidelines for DNS Transport in Mixed IPv4/IPv6 Environments
draft-ietf-dnsop-3901bis-16

## Abstract

This document provides guidelines and documents Best Current Practice for operating authoritative DNS servers, recursive resolvers and stub resolvers in a mixed IPv4/IPv6 environment. This document recommends that both authoritative DNS servers and recursive resolvers support IPv4 and IPv6. It also provides guidance on how recursive DNS resolvers should select upstream DNS servers, including when IPv4-embedded IPv6 addresses are available.

This document obsoletes RFC 3901.

## 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 since the mid-1990s [RFC2460], consistent deployment throughout the whole Internet has not yet been accomplished [RFC9386]. Hence, the Internet still consists of IPv4-only, dual-stack (networks supporting both IP address families), and IPv6-only 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 connectivity. This poses a challenge for such resolvers because they may receive queries for names whose authoritative DNS servers do not support the same IP address family as the resolver itself.

[RFC3901] was initially written at a time when IPv6 deployment was not widespread, focusing primarily on maintaining name space continuity within the IPv4 landscape. Two decades later, not only is IPv6 widely deployed, it is also becoming the de facto standard in many areas, such as mobile and access networks and data center underlays. Furthermore, since 2012, IPv6 support being required for all IP-capable nodes has been established as a best current practice [RFC6540]. This document broadens the scope of [RFC3901] recommending IPv6 connectivity for authoritative DNS servers, recursive resolvers, and stub resolvers.

This document provides:

* Guidance on IP address family-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 transition and coexistence setups may mitigate some of the DNS resolution issues in a mixed IP address family Internet, making DNS data accessible over both IPv4 and IPv6 is the most robust and flexible approach. This approach allows resolvers to retrieve the information they need without requiring intermediary translation or encapsulation services, which may introduce additional failure cases.

Refer to Appendix A for an overview of the main changes since [RFC3901].

## 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 that provides either authoritative or recursive DNS services and is reachable via IPv4. This does not imply anything about the DNS data served, but rather that the name server receives and answers queries over IPv4.

IPv6 name server:
: A name server that provides either authoritative or recursive DNS services and is reachable via IPv6. This does not imply anything about the DNS data served, but rather that the name server receives and answers queries over IPv6.

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

Effective PMTU
: The effective Path Maximum Transmission Unit (PMTU) is the largest IP packet size (in octets) that can successfully traverse a network path from source to destination without requiring fragmentation.

## 3. Name Space Fragmentation

When a resolver looks up a name, it starts at the root and follows referrals until it reaches a name server set that is authoritative for the name. However, if the referrals lead to a name server set that only contains name servers reachable via an IP address family not supported by the resolver, the resolver is unable to continue DNS resolution.

If this occurs, the DNS has effectively fragmented due to mismatching IP address family support between the recursive DNS resolver and the authoritative DNS server.

With the deployment of both IPv4 and IPv6, name space 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 misconfigurations that make a zone unresolvable by either IPv4-only or IPv6-only resolvers. The latter is often hard to identify because the impact of misconfigurations affecting one IP address family (IPv4 or IPv6) may be hidden in a dual-stack setting. In the worst case, where both IP address families must be fully supported by a resolver, a specific name may only be resolvable via dual-stack enabled resolvers.

### 3.1. Misconfigurations Causing IP Address Family Related Name Space Fragmentation

Even when an administrator assumes that they have enabled support for a specific IP address family on their authoritative DNS server, various misconfigurations may break the DNS delegation chain of a zone for that IP address family, preventing any of its records from being resolved by clients that only support that IP address family. Such misconfigurations may remain undetected if most clients can successfully fall back to the other IP address family.

The following name-related misconfigurations can cause broken delegation for one IP address family:

No A/AAAA records for NS names:
: If all of the NS resource records (RR) for a zone in their parent zone have either only A RRs or only AAAA RRs, then resolution via the other IP address family 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 needs to contain both IPv4 and IPv6 glue records. A parent needs to serve the corresponding A and AAAA RRs in the additional section when returning the NS RRs as the referral response [RFC9471].

No A/AAAA RR for in-domain NS:
: If the parent provides glue records for both IP address families but the child zone itself lacks corresponding A or AAAA RRs for its in-domain NS' names, resolution via the missing IP address family will fail during delegation revalidation (see, e.g., [I-D.ietf-dnsop-ns-revalidation]).

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

Parent zone not resolvable via one IP address family:
: For a zone to be resolvable via an IP address family the parent zones up to the root zone needs to be resolvable via that IP address family as well. Any zone not resolvable via the concerned IP address family 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 RR but also via an RR providing the IP address of a name server that is not configured to answer queries via that IP address family [V6DNSRDY-23].

Finally, at the time of this writing, addresses (A or AAAA RRs) for a delegation's authoritative name servers are the only type of glue defined for the DNS. In the future, alternative, yet related, delegation systems may be available, where other considerations apply.

### 3.2. Network Conditions Causing IP Address Family 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 are packets requiring fragmentation given a reduced path MTU (PMTU) and MTU discards, 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 ways:

DNS-over-UDP packets requiring fragmentation
: When using EDNS(0) to communicate support for DNS messages larger than 512 octets [RFC6891] via conventional DNS-over-UDP transport according to [RFC1035], an IP packet carrying a DNS response may exceed the PMTU for the path to a resolver. If an authoritative DNS server does not follow [RFC9715], i.e., honors EDNS(0) sizes larger than 1232 octets, it will try to fragment the packet according to the discovered PMTU. Such packets mostly occur for DNSKEY responses with DNSSEC [RFC4034].

In general, DNS servers SHOULD follow [RFC9715], which provides additional guidance on preventing fragmentation. [RFC9715] suggests setting an upper bound for received EDNS(0) sizes of 1400 octets to avoid the need for fragmentation. However, the [DNSFlagDay2020] initiative suggests using an upper bound EDNS(0) size of only 1232 octets, which is also adopted by most implementations. Setting the upper bound at 1232 octets ensures that generated packets do not exceed 1280 octets, i.e., the minimum MTU for IPv6 [RFC8200] which avoids IPv6 host fragmentation by the server. Hence, for clarity, the present document specifically notes that clients MAY use an EDNS(0) size of 1232 octets as well.

Additionally, e.g., as an additional precaution or because the DNS implementation in use does not support limiting the effective EDNS(0) size, DNS servers MAY opt to explicitly not rely on path MTU discovery [RFC4821] or PLPMTUD [RFC8899]. It can do so, for example, by setting IPV6_USE_MIN_MTU=1 from [RFC3542] to avoid the need to perform PMTU discovery.

DNS-over-TCP packets requiring fragmentation

A resolver can for various reasons also initiate connections via TCP for resolution to an authoritative server. However, similar to the case of DNS-over-UDP, DNS-over-TCP may encounter MTU discards if PMTUD is not possible on a given path. This can occur, for example, if PMTUD related ICMP/ICMPv6 messages are dropped (i.e., cannot be returned to the sender) or if the size communicated in these messages is incorrect (i.e., an on-path device alters packets' size). Under these conditions, the MSS honored by the authoritative DNS server leads to IP packets exceeding the effective PMTU of the path taken by responses. 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.

[RFC8200] recommends that IPv6 nodes implement Path MTU Discovery in order to discover and take advantage of path MTUs greater than 1280 octets. Usually, when a transport protocol can use PMTU (or PLPMTUD [RFC8201] or Datagram PLPMTUD [RFC4821] [RFC8899]) this SHOULD be used to determine an effective PMTU.

However, as DNS benefits from low latency, and performing PMTU (or PLPMTUD [RFC8201] or Datagram PLPMTUD [RFC4821] [RFC8899]) could lead to DNS requests timing out before the effective PMTU can be established by the server. Furthermore, at the time of writing

most DNS messages fit into less than 1280 octets [DNSv6MTU], which means that the benefits of being able to leverage a larger effective PMTU only effect corner cases, e.g., requests for DNSKEY RRs. Additionally, not having to rely on PMTUD benefits DNS' time budget, as the time needed for PMTUD could already exceed the timeout budget for DNS resolution, i.e., could prevent resolution for cases where PMTUD is needed all together.

Hence, DNS servers SHOULD configure the maximum response size to avoid fragmentation or on-path discarding of packets larger than the effective PMTU. For TCP, this can be accomplished by restricting the used maximum segment size (MSS), either by the host limiting the MSS on its own, or by rewriting the MSS field in packets during a TCP handshake.

Therefore, it is RECOMMENDED that DNS servers set a Sender MSS (MSS_S) of no more than 1388 octets for TCP connections. Setting this MSS ensures that packets do not exceed a size of 1448 octets, i.e., the same packet size recommended to avoid fragmentation for DNS-over-UDP packets in [RFC9715]. Furthermore, to provide additional clarity similar to the above guidance on UDP, DNS servers MAY ensure that a total packet size of 1280 octets is not exceeded by setting the Sender MSS (MMS_S) to 1220 octets, as suggested by the [DNSFlagDay2020] initiative, see Section 3.7.1 of [RFC9293].

Additionally, e.g., as an additional precaution or because the DNS implementation in use does not support limiting the effective MSS size, DNS servers MAY opt to explicitly not rely on path MTU discovery [RFC4821] or PLPMTUD [RFC8899]. It can do so, for example, by setting IPV6_USE_MIN_MTU=1 from [RFC3542].

Broken IP Connectivity at the Resolver

Similar to authoritative servers, (stub) recursive resolvers may face broken IP connectivity for either IPv4 or IPv6:

IPv4 connectivity for a DNS resolver may experience issues, e.g., if the resolver is deployed behind a Carrier Grade NAT (CGN) [RFC6888] that implements strict timeouts on active sessions, or limits the number of available TCP and UDP ports numbers for connections below the number required by the multiple connections necessary during recursive DNS resolution. Similarly, [RFC1918] addressing may be in use on the resolver, while address translation is not performed, or, similar to the case for IPv6, when the DNS resolver has a global IPv4 address, but that address is not forwarded on the resolver's network.

IPv6 connectivity for a DNS resolver may experience issues, if, e.g., a client has been assigned a global unicast IPv6 address, but IPv6 traffic is not forwarded on the resolver's network. Also, a resolver may only have received an [RFC4193] unique local IPv6 unicast address (ULA), which does not allow it to reach global addresses without translation. Similarly, IPv6 connectivity can experience issues when IPv4-IPv6 transition technologies like NAT64 [RFC6146] on IPv6-mostly networks [RFC9313] are in use, where the use of NAT64 can be, e.g., discovered through PREF64 in Router Advertisements (RAs) [RFC8781] or DNS64 [RFC7050]. There, the synthesized IPv6 addresses used in, e.g., 464XLAT [RFC6877] encounter additional PMTU fluctuation due to the difference in header size between IPv4 and IPv6, possibly impacting DNS resolution.

Note: This document only explicitly discusses DNS-over-TCP and DNS-over-UDP. However, several other transport methods between recursive and authoritative DNS severs exist, including DNS over various encrypted transports. Some of these technologies provide additional mechanisms for preventing the impact of a reduced PMTU or MTU discards. Guidance in this document focuses on IP address family support, and questions of the underlying transport protocol (TCP or UDP). If DNS servers use an additional protocol layer, e.g., DNS-over-TLS [RFC7858] or DNS-over-QUIC [RFC9250], for their communication, and that protocol supports additional measures to prevent issues related to fragmentation on the IP layer, these measures SHOULD be used for the connection. If the protocol is not resilient to IP layer fragmentation related issues by default, the above guidance for TCP and UDP based connections SHOULD be applied analogously.

### 3.3. Reasons for Intentional IP Address Family 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 RRs for NS names. Based on a 2023 study, 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 address exhaustion progresses, IPv6 adoption is expected to increase.

## 4. Policy Based Avoidance of Name Space Fragmentation

With the final exhaustion of IPv4 address pools in RIRs, e.g., [RIPEV4], and the progressing deployment of IPv6, IPv4 and IPv6 have become comparably relevant. Yet, while it is observed that the first zones becoming exclusively IPv6 resolvable, there is still a major portion of zones solely relying on IPv4 [V6DNSRDY-23]. Hence, dual-stack connectivity is still instrumental to be able to resolve zones and avoid name space fragmentation.

Having zones served only by name servers reachable via one IP address family would fragment the DNS. Hence, the need for 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 (Section 4.1 of [RFC1034]). Typically, these servers are geographically diverse and operate under different routing policies [RFC2182], as also mirrored by, e.g., the IANA requirements for TLD authoritative name servers [IANANS]. To prevent DNS name space fragmentation, at least two IPv4-reachable and two IPv6-reachable name servers MUST be configured for a zone. A single name server that is reachable over both IPv4 and IPv6 counts once per address family. Specifically, key requirements for a zone are:

IPv4 adoption:
: To maintain name space continuity, every DNS zone MUST be served by at least two authoritative DNS servers providing services via IPv4. Furthermore, the delegation configuration of an NS (Resolution of the parent, resolution of sibling domain names, glue) MUST NOT rely on IPv6 connectivity being available.

IPv6 adoption:
: To maintain name space continuity, every DNS zone MUST be served by at least two authoritative DNS servers providing services via IPv6. To avoid reachability issues, authoritative DNS servers MUST NOT use IPv4-embedded addresses [RFC6052] (including IPv4-Mapped IPv6 addresses and deprecated IPv4-Compatible addresses [RFC4291]) for receiving queries. Furthermore, the delegation configuration of an NS (Resolution of the parent, resolution of sibling domain names, glue) MUST NOT rely on IPv4 connectivity being available.

Consistency:
: Both IPv4 and IPv6 transports MUST 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 reductions in the effective PMTU (e.g., when NAT64 is used (Section 7 of [RFC7269])). Therefore, IP fragmentation SHOULD be avoided by following guidance on maximum DNS payload sizes [RFC9715]. Furthermore, as per Section 5 of [RFC7766], DNS-over-TCP MUST be available as a fall-back option, instead of relying on fragmented UDP packets. Similar to the guidance in [RFC9715], authoritative DNS servers MAY set an MSS of either 1388 (analogous to [RFC9715]) or 1220 (analogous to the [DNSFlagDay2020] suggestions) in TCP sessions carrying DNS responses.

To prevent name space fragmentation, zone validation processes SHOULD ensure that:

* There are at least two IPv4 address records and two IPv6 address records 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 performing DNS resolution via IPv4-only and IPv6-only networks respectively.

## 4.2. Guidelines for Recursive DNS Resolvers

To ensure robust DNS resolution even when facing namespace fragmentation, every recursive DNS resolver SHOULD be dual-stack. Exceptions apply if one of the below methods to prevent namespace fragmentation are in place.

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 servers or uses a configuration where it forwards queries failing IPv6-only DNS resolution to a dual-stack recursive DNS resolver (i.e., a resolver that is also able to perform DNS resolution over IPv4). If a

recursive DNS resolver is aware of a PREF64 to use for NAT64 [RFC6146], either through static configuration or by discovering it (e.g., [RFC8781]), it MAY synthesize IPv6 addresses for remote authoritative DNS servers.

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 dual-stack recursive DNS resolver (i.e., a resolver that is also able to perform DNS resolution over IPv6).

Finally, when responding to recursive queries (i.e., a query with the RD bit set [RFC1035]), a DNS resolver SHOULD follow the above guidance on fragmentation avoidance (Section 4.1) for communication between authoritative DNS servers and recursive DNS resolvers analogously.

### 4.3. Guidelines for DNS Stub Resolvers

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. Furthermore, a stub DNS resolver has to rely on recursive DNS servers discovered for the local network, e.g., using DHCPv4 [RFC2131], DHCPv6 [RFC8415], and/or router advertisements [RFC8106]. In that case, the stub resolver may obtain multiple different IPv4 and IPv6 DNS resolver addresses to use.

To prioritize different IPv4 and IPv6 DNS resolver addresses, a stub resolver SHOULD follow [RFC6724]. However, a stub DNS resolver SHOULD NOT utilize IPv4-embedded IPv6 addresses if it is able to identify them as such, e.g., by having discovered the PREF64 in use for the network [RFC8781].

When providing multiple DNS servers to stub resolvers, network operators have to consider that, at the time of writing, various implementations can only configure a small set of possible DNS resolvers, e.g., only up to three for libc [MAN], and additional resolvers provided may be ignored by clients. Hence, when providing more than three DNS servers to stub resolvers, operators SHOULD ensure that no more than two recursive DNS servers supplied to clients are unable to perform dual-stack DNS resolution, and at least one of the supplied recursive DNS servers is able to perform dual-stack DNS resolution. If this is not done, a client might select a subset of recursive DNS servers that leads to address family based namespace fragmentation.

## 5. Security Considerations

The guidelines described in this memo introduce no new security considerations into the DNS protocol itself.

Nevertheless, corner cases exist where forwarding queries requiring an IP address family for resolution that is not supported by the initial resolver lead to an infinite forwarding loop, under the following conditions:

* Two resolvers handle queries for a set of clients, each of these resolvers support one and only one address family that is distinct from the address family supported by the other resolver;

* Both resolvers are configured to forward queries requiring DNS resolution via the IP address family they do not support to the other; and

* A query for a zone that is not resolvable via IPv4 and not resolvable via IPv6 is received.

In such a case, a query for the non-resolvable zone would be endlessly forwarded between these resolvers.

To prevent such cases, single-stack recursive DNS resolvers SHOULD be configured to forward queries they cannot resolve due to lacking support for one address family to dual-stack recursive DNS resolvers. Furthermore, recursive DNS resolvers MUST NOT be configured to forward queries to DNS resolvers that are configured to forward queries to them in the first place.

Recommendations for recursive and stub resolvers rely on a correctly discovered PREF64. Security issues may materialize if an incorrect PREF64 is used. Hence, guidance from [RFC9872] on securely discovering PREF64 SHOULD be followed.

Preventing fragmentation according to the guidance in this document may increase load on DNS servers, as more TCP fallbacks might be required. While measurements have shown this to be (at the time of writing) in the range of 3-5% of connections [DNSv6MTU], operators SHOULD monitor the actual impact on their servers when implementing guidance from this document to detect unexpected load increases early on.

## 6. IANA Considerations

This document requests IANA consider updating its technical requirements for authoritative DNS servers to require both IPv4 and IPv6 addresses for each authoritative server [IANANS], in accordance with its processes for reviewing and revising these procedures.

## 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, Philipp S. Tiesel, Mark Andrews, Stefan Ubbink, Joe Abley, Gorry Fairhurst, Paul Vixie, Lorenzo Colitti, David Farmer, Pieter Lexis, Ralf Weber, Philip Homburg, Marco Davids, Mohamed Boucadair, Thomas Fossati, Aihua Guo, Bernie Volz, David Dong, Roman Danyliw, テ詠ic Vyncke

Thank you for reading this draft.

The authors furthermore express their thanks towards the authors of [RFC3901], Alain Durand and Johan Ihren, and provide their original acknowledgements verbatim below:

This document is the result of many conversations that happened in the DNS community at IETF and elsewhere since 2001. During that period of time, a number of Internet drafts have been published to clarify various aspects of the issues at stake. This document focuses on the conclusion of those discussions.

The authors would like to acknowledge the role of Pekka Savola in his thorough review of the document.

## References

### Normative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <https://www.rfc-editor.org/info/rfc1034>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[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>.

[RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <https://www.rfc-editor.org/info/rfc4291>.

[RFC4821] Mathis, M. and J. Heffner, "Packetization Layer Path MTU Discovery", RFC 4821, DOI 10.17487/RFC4821, March 2007, <https://www.rfc-editor.org/info/rfc4821>.

[RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", RFC 6052, DOI 10.17487/RFC6052, October 2010, <https://www.rfc-editor.org/info/rfc6052>.

[RFC6724] Thaler, D., Ed., Draves, R., Matsumoto, A., and T. Chown, "Default Address Selection for Internet Protocol Version 6 (IPv6)", RFC 6724, DOI 10.17487/RFC6724, September 2012, <https://www.rfc-editor.org/info/rfc6724>.

[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>.

[RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, RFC 8200, DOI 10.17487/RFC8200, July 2017, <https://www.rfc-editor.org/info/rfc8200>.

[RFC8899] Fairhurst, G., Jones, T., Tテシxen, M., Rテシngeler, I., and T. Vテカlker, "Packetization Layer Path MTU Discovery for Datagram Transports", RFC 8899, DOI 10.17487/RFC8899, September 2020, <https://www.rfc-editor.org/info/rfc8899>.

[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>.

[RFC9293] Eddy, W., Ed., "Transmission Control Protocol (TCP)", STD 7, RFC 9293, DOI 10.17487/RFC9293, August 2022, <https://www.rfc-editor.org/info/rfc9293>.

[RFC9471] Andrews, M., Huque, S., Wouters, P., and D. Wessels, "DNS Glue Requirements in Referral Responses", RFC 9471, DOI 10.17487/RFC9471, September 2023, <https://www.rfc-editor.org/info/rfc9471>.

[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

[DNSFlagDay2020] "DNS flag day 2020", <https://dnsflagday.net/2020/>.

[DNSv6MTU] Fiebig, T. and A. Feldmann, "'How I learned to stop worrying and love IPv6': Measuring the Internet's Readiness for DNS over IPv6", October 2025, <https://doi.org/10.1145/3730567.3764439>.

[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-11, 19 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-dnsop-ns-revalidation-11>.

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

[MAN] Linux, "resolv.conf(5) 寰Linux manual page", 2025, <https://man7.org/linux/man-pages/man5/resolv.conf.5.html>.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G. J., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <https://www.rfc-editor.org/info/rfc1918>.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, DOI 10.17487/RFC2131, March 1997, <https://www.rfc-editor.org/info/rfc2131>.

[RFC2182] Elz, R., Bush, R., Bradner, S., and M. Patton, "Selection and Operation of Secondary DNS Servers", BCP 16, RFC 2182, DOI 10.17487/RFC2182, July 1997, <https://www.rfc-editor.org/info/rfc2182>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <https://www.rfc-editor.org/info/rfc2460>.

[RFC3542] Stevens, W., Thomas, M., Nordmark, E., and T. Jinmei, "Advanced Sockets Application Program Interface (API) for IPv6", RFC 3542, DOI 10.17487/RFC3542, May 2003, <https://www.rfc-editor.org/info/rfc3542>.

[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>.

[RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <https://www.rfc-editor.org/info/rfc4193>.

[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>.

[RFC6540] George, W., Donley, C., Liljenstolpe, C., and L. Howard, "IPv6 Support Required for All IP-Capable Nodes", BCP 177, RFC 6540, DOI 10.17487/RFC6540, April 2012, <https://www.rfc-editor.org/info/rfc6540>.

[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>.

[RFC6888] Perreault, S., Ed., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", BCP 127, RFC 6888, DOI 10.17487/RFC6888, April 2013, <https://www.rfc-editor.org/info/rfc6888>.

[RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <https://www.rfc-editor.org/info/rfc7050>.

[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>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8106] Jeong, J., Park, S., Beloeil, L., and S. Madanapalli, "IPv6 Router Advertisement Options for DNS Configuration", RFC 8106, DOI 10.17487/RFC8106, March 2017, <https://www.rfc-editor.org/info/rfc8106>.

[RFC8201] McCann, J., Deering, S., Mogul, J., and R. Hinden, Ed., "Path MTU Discovery for IP version 6", STD 87, RFC 8201, DOI 10.17487/RFC8201, July 2017, <https://www.rfc-editor.org/info/rfc8201>.

[RFC8415] Mrugalski, T., Siodelski, M., Volz, B., Yourtchenko, A., Richardson, M., Jiang, S., Lemon, T., and T. Winters, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 8415, DOI 10.17487/RFC8415, November 2018, <https://www.rfc-editor.org/info/rfc8415>.

[RFC8781] Colitti, L. and J. Linkova, "Discovering PREF64 in Router Advertisements", RFC 8781, DOI 10.17487/RFC8781, April 2020, <https://www.rfc-editor.org/info/rfc8781>.

[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>.

[RFC9250] Huitema, C., Dickinson, S., and A. Mankin, "DNS over Dedicated QUIC Connections", RFC 9250, DOI 10.17487/RFC9250, May 2022, <https://www.rfc-editor.org/info/rfc9250>.

[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>.

[RFC9872] Buraglio, N., Jensen, T., and J. Linkova, "Recommendations for Discovering IPv6 Prefix Used for IPv6 Address Synthesis", RFC 9872, DOI 10.17487/RFC9872, September 2025, <https://www.rfc-editor.org/info/rfc9872>.

[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-Ga単叩n, 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>.

## Appendix A. Changes Since [RFC3901]

The following changes have been made to the guidance published in [RFC3901]:

* Expanded the terminology section, also taking considerations from [RFC9499] into account.

* Expanded namespace fragmentation, independently discussing IP address family related namespace fragmentation, network condition based namespace fragmentation, and intentional namespace fragmentation.

* Now recommends the use of IPv4 and IPv6 for authoritative DNS servers, instead of leaving IPv6 optional.

* Now recommends testing IPv4 and IPv6 resolvability when delegating zones, instead of only testing IPv4 resolvability.

* Added guidance on handling IP layer fragmentation.

* Added guidance for IP address family handling for recursive and stub resolvers.

## Authors' Addresses

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

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