

dnsop
Internet-Draft
Obsoletes: 3901 (if approved)
Intended status: Best Current Practice
Expires: 4 June 2026

Momoka
WIDE Project
T. Fiebig
MPI-INF
1 December 2025

DNS IPv6 Transport Operational Guidelines
draft-ietf-dnsop-3901bis-08

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 recommends that authoritative DNS servers as well as recursive DNS resolvers support both IPv4 and IPv6. It furthermore provides guidance for how recursive DNS resolvers should select upstream DNS servers, if synthesized and non-synthesized IPv6 addresses are available.

This document obsoletes RFC 3901. (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

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 4 June 2026.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. Terminology	3
3. Name Space Fragmentation	4
3.1. Misconfigurations Causing IP Version Related Name Space Fragmentation	4
3.2. Network Conditions Causing IP Version Related Name Space Fragmentation	6
3.3. Reasons for Intentional IP Version Related Name Space Fragmentation	8
4. Policy Based Avoidance of Name Space Fragmentation	8
4.1. Guidelines for Authoritative DNS Server Configuration . .	9
4.2. Guidelines for Recursive DNS Resolvers	10
4.3. Guidelines for DNS Stub Resolvers	10
5. Security Considerations	11
6. IANA Considerations	11
Acknowledgments	11
References	11
Normative References	11
Informative References	14
Authors' Addresses	15

1. Introduction

Despite IPv6 being first discussed in the mid-1990s [RFC2460], consistent deployment throughout the whole Internet has not yet been accomplished [RFC9386]. Hence, today, the Internet consists of IPv4-only, dual-stack (networks supporting both IP versions), 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. This poses a challenge for such resolvers because they may receive queries for names that have authoritative DNS servers which do not support the same IP version as the resolver.

[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, IPv6 is 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 connectivity 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. This approach 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 means that the name server receives and answers queries 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 means 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 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 set that is authoritative for the name. If it is referred to a name server set that is, based on a referral, only contains name servers that are exclusively reachable via an IP address family the resolver does not support, 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, 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 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 needs to contain both IPv4 and IPv6 GLUE records. A parent needs to serve the corresponding A and AAAA records as ADDITIONAL data when returning the NS record(s) as the referral response [RFC9471].

No A/AAAA record for in-domain NS:

If the parent provides GLUE records for both IP versions but the child zone itself lacks corresponding A or AAAA records for its in-domain name server names, resolution via the missing IP version will fail during delegation revalidation [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 needs to 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 needs to 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 name server 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 ways:

DNS-over-UDP packets requiring fragmentation

When using EDNS(0) to communicate support for DNS messages larger than 512 octets [RFC6891] via traditional DNS-over-UDP transport according to RFC1035 [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 [RFC9715], which provides additional guidance on preventing fragmentation by ensuring that the maximum DNS/UDP payload size does not exceed 1400 octets. This can be accomplished by setting a corresponding EDNS(0) size, with most implementations using a lower EDNS(0) size of 1232 octets following [DNSFlagDay2020], to ensure that generated packets always fit into lower bound of the IPv6 MTU of 1280, as defined in [RFC8200]. Hence, DNS servers MAY opt to set an EDNS(0) size of 1232 octets following [DNSFlagDay2020].

Additionally, DNS servers MAY opt to explicitly not rely on path MTU discovery [RFC4821] or PLPMTUD [RFC8899], by instead using IPV6_USE_MIN_MTU=1 from RFC3542 [RFC3542] to avoid the need to perform path MTU 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 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, transferring the guidance from [RFC9715], setting an MSS of 1388 octets would reduce the impact of this issue. Hence, DNS servers MAY set an MSS of no more than 1388 octets for TCP connections. Similarly, aligned with the recommendations of the [DNSFlagDay2020], DNS servers MAY ensure that a total packet size of 1280 octets is not exceeded by setting an MSS of 1220 octets. Additionally, DNS servers MAY opt to set `IPV6_USE_MIN_MTU=1` from RFC3542 [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 behind a Carrier Grade NAT (CGN) [RFC6888] setup that implements strict timeouts on active sessions, or limits the number of available ports for connections. 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 routed 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 routed on the resolver's network. Similarly, IPv6 connectivity can experience issues when IPv4-IPv6 transition technologies, e.g., NAT64 [RFC6146] on IPv6-mostly networks [RFC9313], or NAT64 connectivity discovered through PREF64 [RFC8781] or DNS64 [RFC7050] on IPv6-only networks are in use. There, the synthesized IPv6 addresses used in 464XLAT [RFC6877] encounter additional PMTU fluctuation due to the difference in header size between IPv4 and IPv6, possibly impacting DNS resolution.

Note: Please note that this document only explicitly discusses DNS-over-TCP and DNS-over-UDP. However, several other transport methods between recursive and authoritative DNS servers exist, including DNS over various encrypted transports. Some of these technologies provide additional mechanisms for preventing the impact of a reduced PMTU or MTU blackholes. Guidance in this document focuses on IP version 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 fragmentation on the IP layer related issues, these measures SHOULD be used for the connection. Otherwise, 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 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, IPv4 and IPv6 have become comparably relevant. 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. To avoid reachability issues, authoritative DNS servers SHOULD use native IPv6 addresses instead of IPv6 addresses synthesized using IPv6 transition technologies for receiving queries. 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], 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 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 performing DNS resolution via IPv4-only and IPv6-only networks respectively.

4.2. Guidelines for Recursive 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 servers, 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. For example, if a recursive DNS resolver is aware of a PREF64 to use for NAT64 [RFC6146], either through static configuration or by discovering it [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 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 on fragmentation avoidance, see also [RFC9715], 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 [RFC3456], DHCPv6 [RFC8415], and/or SLAAC [RFC4862]. 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 synthesized 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 possible DNS servers to stub resolvers, operators SHOULD consider that various implementations can only configure a small set of possible DNS resolvers, e.g., only up to three for libc, and additional resolvers provided may be ignored by clients.

5. Security Considerations

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

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.

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, Philipp S. Tiesel, Mark Andrews, Stefan Ubbink, Joe Abley, Gorrry Fairhurst, Paul Vixie, Lorenzo Colitti, David Farmer, Pieter Lexis, Ralf Weber, Philip Homburg, Marco Davids

Thank you for reading this draft.

References

Normative References

- [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>>.
- [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>>.
- [RFC3456] Patel, B., Aboba, B., Kelly, S., and V. Gupta, "Dynamic Host Configuration Protocol (DHCPv4) Configuration of IPsec Tunnel Mode", RFC 3456, DOI 10.17487/RFC3456, January 2003, <<https://www.rfc-editor.org/info/rfc3456>>.
- [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. Ihen, "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>>.
- [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>>.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, DOI 10.17487/RFC4862, September 2007, <<https://www.rfc-editor.org/info/rfc4862>>.
- [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>>.
- [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>>.

- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8899] Fairhurst, G., Jones, T., T端 xen, M., R端 ngeler, I., and T. V端 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>>.

- [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>>.
- [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>>.

Informative References

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

- [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>>.
- [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>>.
- [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>>.
- [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>.

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