

Internet Engineering Task Force
Internet-Draft
Intended status: Experimental
Expires: 25 October 2026

S. Bortzmeyer
Afnic
W. Toorop
NLnet Labs
B. Farrokhi
Quad9
M. Rahman
The FreeBSD Foundation
O. Sur^筆
Internet Systems Consortium
O. Moerbeek
PowerDNS
23 April 2026

Synchronizing caches of DNS resolvers
draft-bortzmeyer-dnsop-poisonlicious-03

Abstract

Network of cooperating and mutually trusting DNS resolvers could benefit from cache sharing, where one resolver would distribute the result of a resolution to other resolvers. This document standardizes a protocol to do so.

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 25 October 2026.

Copyright Notice

Copyright (c) 2026 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
1.2. Terminology	3
2. The protocol	3
3. IANA Considerations	4
4. Security Considerations	4
5. Privacy Considerations	4
6. Operational Considerations	5
7. Related and future work	5
7.1. Related work	5
7.2. Future work	5
7.2.1. Actual measurement	5
7.2.2. Negative answers	5
7.2.3. ECS	5
7.2.4. Authentication of messages	6
7.2.5. Dispatching of messages	6
7.2.6. Transport of messages	6
7.2.7. Packing of messages	6
7.2.8. Different responses	6
8. References	6
8.1. Normative References	6
8.2. Informative References	7
Acknowledgements	8
Authors' Addresses	8

1. Introduction

When an organisation operates a big network of DNS resolvers [RFC1034] [RFC1035], for instance for an important public resolver (Section 6 of [RFC9499]), it may be a performance improvment to distribute the result of the resolution process between the resolvers. This document standardizes how to to do so, using unicast messages to a set of pre-configured peers.

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.

1.2. Terminology

Network of resolvers

A set of resolvers working together under the same administration

Peer (or peer resolver)

One of the other resolvers in the network

Originating resolver

A resolver sending data to its peers in the network

Receiving resolver (or receiving peer)

A resolver receiving data from one of its peers in the network

Resolver

As used in Section 6 of [RFC9499]

2. The protocol

When completing a successful DNS resolution, the resolver transmits a DNS message (with the Q/R bit set, since it is a response) to the pre-configured peers, authenticating with TSIG [RFC8945]. No acknowledgment is sent or expected. To save work, the resolver MAY send the data only if the TTL is higher than some predefined value.

The resolver must send only data that it is sure of (for instance by DNSSEC validation or because it came with the AA bit from the queried server). Since all of the network of resolvers are in the same organizational domain, they MUST agree on the same policy for this assessment.

Negative answers ([RFC9499], section 3) MUST NOT be transmitted to peers.

Messages of this protocol are distinguished from other DNS messages by the TSIG key they use (which must therefore be specific to this protocol).

This message MAY be the message received by the resolver from the authoritative name servers or it MAY be a new message with data composed from data already obtained by the resolver, for instance without the question section (this would help a bit with privacy Section 5).

The EDNS section MUST be a new one, created to fit the needs of successful transmission to the peer.

Each peer then MAY store the data in its cache. The peer is not supposed to do DNSSEC validation (there is not always all the necessary data in the message). After all, the goal is to save work for the peers, so Section 5.4.1 of [RFC2181] does not apply here. (Remember all peers trust each other, and have a consistent policy. The data is as trustworthy as if you validated it yourself.) The receiver MAY cache only what is in the Answer section.

3. IANA Considerations

None. [RFC-Editor: you may delete this section]

4. Security Considerations

The integrity and authenticity of the cached data is of course critical. DNSSEC would help but it is not yet universally deployed and, moreover, the peer resolvers should not have to redo the validation. So, trust between the peer resolvers is expected because it is the only way for the receiver to be sure of the data. One way to do so is to have all of the peers under the same organisational authority, as mandated here.

For the same reason, the channel between peers must be protected, preferably with cryptography (currently, TSIG is mandatory). ACL and other network techniques are of course useful.

Encryption is less important than authentication since we transmit only public data. Nevertheless, it is better to be sure that the channel between the peers is not open to snooping.

5. Privacy Considerations

Confidentiality is currently out of scope for this document. The communication between the originating resolver and its receiving peers could be encrypted, for instance with DoT [RFC7858] but it is not otherwise specified.

If the originating resolver sends the original question section in its messages to receiving peers, it can have privacy consequences [RFC9076]. These consequences are limited since all the peers are under the same administration, anyway. The originator MAY remove this section or replace it with dummy data.

6. Operational Considerations

It is reminded that all resolvers in the network need to trust each other, being in the same administrative domain. This specification is not meant to be deployed between unrelated resolvers.

The network of peer resolvers have to be configured out-of-band before. The way to do it is out-of-scope for this specification.

7. Related and future work

7.1. Related work

[I-D.hl-dnsop-cache-filling] describes a mechanism to fill DNS caches with data. The format is, like in this document, standard DNS as seen on the wire.

7.2. Future work

7.2.1. Actual measurement

Measuring the efficiency of caching optimizations is hard because the performance depend a lot on the actual queries sent and their timeline. It would be interesting to test on real-world data to measure the improvement brought by this technology.

7.2.2. Negative answers

If in the future, we may allow negative answers to be sent, which would permit the receivers to use [RFC8020] and/or [RFC8198] to synthesize negative answers.

Be careful of the risk of overloading receiving peers for instance when there is a dictionary attack.

7.2.3. ECS

Today, we don't transmit the EDNS record since EDNS is not end-to-end. But ECS [RFC7871] may be an important information for the participating resolvers.

7.2.4. Authentication of messages

Today, we authenticate only with TSIG [RFC8945]. SIG(0) [RFC2931] or DoT [RFC7858] may be interesting improvements in the future.

7.2.5. Dispatching of messages

Today, messages following this specification are sent to the proper handler based on the TSIG key used. It could be an interesting alternative to use a separate transport port instead.

7.2.6. Transport of messages

Messages could be transmitted in long-lived TCP sessions, too.

If there are 1,000 servers, sending 1,000 messages, or having a full mesh of 1,000 TCP connections may be too much. It may be interesting to replace the unicast messages by multicast [RFC5110] (the issues of multicast on the public Internet do not apply here since we envision work under only one organisation).

Other protocols may be considered such as MQTT [MQTT] which is well suited for publish-by-one/consume-by-many, raw protocol buffers [GPB] or dnstap [dnstap] over them.

7.2.7. Packing of messages

It could be interesting to optimize by packing the data in a C-DNS [RFC8618] flow, sent with TCP (with TLS) or QUIC. (Of course, other formats/protocols are possible.)

7.2.8. Different responses

When the authoritative servers send different replies depending on the client, the various peers may send different (and under-optimized) responses to a receiving peer.

8. References

8.1. 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>>.
- [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>>.
- [RFC8945] Dupont, F., Morris, S., Vixie, P., Eastlake 3rd, D., Gudmundsson, O., and B. Wellington, "Secret Key Transaction Authentication for DNS (TSIG)", STD 93, RFC 8945, DOI 10.17487/RFC8945, November 2020, <<https://www.rfc-editor.org/info/rfc8945>>.
- [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>>.

8.2. Informative References

- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", RFC 2181, DOI 10.17487/RFC2181, July 1997, <<https://www.rfc-editor.org/info/rfc2181>>.
- [RFC2931] Eastlake 3rd, D., "DNS Request and Transaction Signatures (SIG(0)s)", RFC 2931, DOI 10.17487/RFC2931, September 2000, <<https://www.rfc-editor.org/info/rfc2931>>.
- [RFC5110] Savola, P., "Overview of the Internet Multicast Routing Architecture", RFC 5110, DOI 10.17487/RFC5110, January 2008, <<https://www.rfc-editor.org/info/rfc5110>>.
- [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>>.
- [RFC7871] Contavalli, C., van der Gaast, W., Lawrence, D., and W. Kumari, "Client Subnet in DNS Queries", RFC 7871, DOI 10.17487/RFC7871, May 2016, <<https://www.rfc-editor.org/info/rfc7871>>.
- [RFC8020] Bortzmeyer, S. and S. Huque, "NXDOMAIN: There Really Is Nothing Underneath", RFC 8020, DOI 10.17487/RFC8020, November 2016, <<https://www.rfc-editor.org/info/rfc8020>>.

- [RFC8198] Fujiwara, K., Kato, A., and W. Kumari, "Aggressive Use of DNSSEC-Validated Cache", RFC 8198, DOI 10.17487/RFC8198, July 2017, <<https://www.rfc-editor.org/info/rfc8198>>.
- [RFC8618] Dickinson, J., Hague, J., Dickinson, S., Manderson, T., and J. Bond, "Compacted-DNS (C-DNS): A Format for DNS Packet Capture", RFC 8618, DOI 10.17487/RFC8618, September 2019, <<https://www.rfc-editor.org/info/rfc8618>>.
- [RFC9076] Wicinski, T., Ed., "DNS Privacy Considerations", RFC 9076, DOI 10.17487/RFC9076, July 2021, <<https://www.rfc-editor.org/info/rfc9076>>.
- [I-D.hl-dnsop-cache-filling] Hoffman, P. E. and M. Larson, "Additional Method for Filling DNS Caches", Work in Progress, Internet-Draft, draft-hl-dnsop-cache-filling-00, 2 March 2018, <<https://datatracker.ietf.org/doc/html/draft-hl-dnsop-cache-filling-00>>.
- [MQTT] OASIS, "MQTT Version 5.0", 2019, <<https://docs.oasis-open.org/mqtt/mqtt/v5.0/os/mqtt-v5.0-os.docx>>.
- [GPB] Google Developers, "Protocol Buffers", <<https://protobuf.dev/>>.
- [dnstap] Edmonds, R., "dnstap", 2014, <<https://dnstap.info/>>.

Acknowledgements

Original idea at the DNS hackathon (RIPE-NCC / Netnod / DNS-OARC) in march 2025 at the Netnod office in Stockholm.

Authors' Addresses

Stéphane Bortzmeyer
Afnic
7 avenue du 8 mai 1945
78280 Guyancourt
France
Email: bortzmeyer+ietf@nic.fr
URI: <https://www.afnic.fr/>

Willem Toorop
NLnet Labs
Science Park 400
1098 XH Amsterdam
Netherlands
Email: willem@nlnetlabs.nl
URI: <https://nlnetlabs.nl/>

Babak Farrokhi
Quad9
Werdstrasse 2
CH-8004 Zürich
Switzerland
Email: babak@farrokhi.net
URI: <https://quad9.net/>

Moin Rahman
The FreeBSD Foundation
3980 Broadway St
Boulder, CO 80304
United States of America
Email: bofh@freebsd.org
URI: <https://freebsd.foundation.org/>

Ondřej Surina
Internet Systems Consortium
Czech Republic
Email: ondrej@isc.org

Otto Moerbeek
PowerDNS
Netherlands
Email: otto.moerbeek@powerdns.com