Domain Name System Operations
Internet-Draft
Intended status: Standards Track
Expires: 24 August 2025

B. M. Schwartz
Meta Platforms, Inc.
P. Sood
Google
J. Todd
Quad9

20 February 2025

Standardized Query Name for DNS Resolver Reachability Probes
draft-sst-dnsop-probe-name-00

Abstract

This specification standardizes DNS names that should be used for checking connectivity to a DNS server.

About This Document

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

Status information for this document may be found at https://datatracker.ietf.org/doc/draft-sst-dnsop-probe-name/.

Source for this draft and an issue tracker can be found at https://github.com/bemasc/probe-name.

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 24 August 2025.

Copyright Notice

Table of Contents

1. Introduction

In the Domain Name System (DNS, [RFC1034]), clients normally send queries to a recursive resolver in order to receive the resolution results. However, some clients also send queries merely to check that a response is received at all. We call these queries "DNS probes". DNS probes are used for many reasons:

* To determine if the network is working at all.

* To detect if a particular address family is connected to the global internet.

* To check that the client is able to reach the resolver's IP address.

* To assess the reliability and performance of the network path between the client and the resolver.

* To establish which transport protocols (e.g., UDP, TCP, TLS [RFC7858], QUIC [RFC9250]) are available.

* To confirm that the DNS resolver itself is operational.

When sending a DNS query, the client must choose a QNAME. Popular QNAME values for probes include:

* Names owned by the entity performing the probe.

* Names used by prominent, high-reliability internet services.

* Names operated at the direction of prominent internet organizations such as the IETF (e.g., "example.com", [RFC2606]).

* Names that form an essential part of the internet infrastructure.

These choices are pragmatic, but they also present a number of downsides for the client:

* The response could be delayed if the selected name is not in cache.

* The probe will return unneeded RDATA, wasting bandwidth.

* Depending on the success criteria, the probe could report a spurious failure

  - if the selected name is removed, or experiences an outage.

  - if the resolver experiences an interruption on its outbound link.

* A distinctive QNAME could enable unwanted fingerprinting of the client by the resolver or a network adversary.

These popular types of QNAME also present some downsides for the resolver operator:

* The probe may cause the resolver to do more work than necessary, especially when the selected name is not in cache.

* The operator cannot distinguish probe queries from ordinary queries, limiting their understanding of how their service is being used.

This specification registers a Special-Use Domain Name for DNS probing to avoid these downsides.

2. Conventions and Definitions

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.

3. Client Requirements

Clients SHOULD set the QNAME on probe queries to "probe.resolver.arpa.".

Clients SHOULD set the QTYPE to A or AAAA. Clients that use other QTYPEs are at a higher risk of implementation fingerprinting due to the distinctive QTYPE.

Clients SHOULD NOT set the "DNSSEC OK" flag. Setting this bit causes more work for the resolver, but does not provide any benefit to the client.

Clients SHOULD avoid any stub caching, as this would cause probe results to be out of date.

Clients MAY set the "Recursion Desired" flag to either value. Setting this flag to 0 reduces load on resolvers that do not implement this specification.

3.1. Using getaddrinfo

Clients MAY perform probe queries using a high-level DNS query interface such as getaddrinfo ([RFC3493], Section 6.1). Note that implementations of getaddrinfo and similar interfaces often employ a stub cache, resulting in probe results that may not be fresh. These implementations typically retry queries across several servers until one responds successfully, so the result may not be attributable to a specific resolver and cannot be used to assess the network’s latency or packet loss rate.

Clients that use getaddrinfo for probes SHOULD call it with the following input parameters:

* nodename = "probe.resolver.arpa."
* servname = null
* hints = null or all zeros except for .ai_family

Clients SHOULD interpret the getaddrinfo return value as follows:

* EAI_NONAME: The probe has succeeded.

* EAI_AGAIN: The probe has failed.

* 0: The resolver is misconfigured (returning addresses where there should be none).

* Any other error: The client system is misconfigured.

4. Server Requirements

Upon receiving a query with a QNAME of "probe.resolver.arpa.", DNS servers MUST return a valid NXDOMAIN response from the "resolver.arpa." locally-served zone [RFC9462].

5. Security Considerations

If a resolver operator applies rate limits to queries, it SHOULD NOT exclude "probe.resolver.arpa" from such limits. Queries for this name could still be used as part of a high query rate attack.

6. IANA Considerations

6.1. Special Use Domain Name "probe.resolver.arpa"

This document calls for the addition of "probe.resolver.arpa" to the Special-Use Domain Names (SUDN) registry established by [RFC6761].

6.2. Domain Name Reservation Considerations

In accordance with Section 5 of [RFC6761], the answers to the following questions are provided for this document:

1) Are human users expected to recognize these names as special and use them differently? In what way?

No. This name is principally intended to be useful to resolver operators, and should never be seen by ordinary users.

2) Are writers of application software expected to make their software recognize these names as special and treat them differently? In what way?

Yes. Writers of DNS resolver monitoring software are expected to categorize queries for this name as distinct from ordinary user-generated queries.

3) Are writers of name resolution APIs and libraries expected to make their software recognize these names as special and treat them differently? If so, how?

No. Stub resolvers process this name in the ordinary fashion.

4) Are developers of caching domain name servers expected to make their implementations recognize these names as special and treat them differently? If so, how?

No. This name is subject to ordinary caching logic.

5) Are developers of authoritative domain name servers expected to make their implementations recognize these names as special and treat them differently? If so, how?

No. Queries for this name are not intended to reach authoritative domain name servers.

6) Does this reserved Special-Use Domain Name have any potential impact on DNS server operators? If they try to configure their authoritative DNS server as authoritative for this reserved name, will compliant name server software reject it as invalid? Do DNS server operators need to know about that and understand why? Even if the name server software doesn’t prevent them from using this reserved name, are there other ways that it may not work as expected, of which the DNS server operator should be aware?

This name has no special impact on DNS server operators beyond those already implied by the status of "resolver.arpa." as a Locally Served Zone.

7) How should DNS Registries/Registrars treat requests to register this reserved domain name? Should such requests be denied? Should such requests be allowed, but only to a specially-designated entity?

This name is inside an existing Locally Served Zone ("resolver.arpa."), so the question of registration requests is moot.

7. References

7.1. 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/rfc/rfc1034>.

[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/rfc/rfc2119>.

[RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <https://www.rfc-editor.org/rfc/rfc6761>.

[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/rfc/rfc8174>.

7.2. Informative References

[RFC2606] Eastlake 3rd, D. and A. Panitz, "Reserved Top Level DNS Names", BCP 32, RFC 2606, DOI 10.17487/RFC2606, June 1999, <https://www.rfc-editor.org/rfc/rfc2606>.

[RFC3493] Gilligan, R., Thomson, S., Bound, J., McCann, J., and W. Stevens, "Basic Socket Interface Extensions for IPv6", RFC 3493, DOI 10.17487/RFC3493, February 2003, <https://www.rfc-editor.org/rfc/rfc3493>.

[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/rfc/rfc7858>.

[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/rfc/rfc9250>.

[RFC9462] Pauly, T., Kinnear, E., Wood, C. A., McManus, P., and T. Jensen, "Discovery of Designated Resolvers", RFC 9462, DOI 10.17487/RFC9462, November 2023, <https://www.rfc-editor.org/rfc/rfc9462>.

Acknowledgments

TODO acknowledge.

Authors' Addresses

Benjamin M. Schwartz
Meta Platforms, Inc.
Email: ietf@bemasc.net

Puneet Sood
Google
Email: puneets@google.com

John Todd
Quad9
Email: jtodd@quad9.net