Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: January 9, 2020

A. Grover
P. Saint-Andre
Mozilla
July 8, 2019

DNS Resolver-Based Policy Detection Domain
draft-grover-add-policy-detection-00

Abstract

This document specifies the behavior that is expected from the Domain Name System with regard to DNS queries for the special-use domain name 'TBD.arpa' and designates this domain as a special-use domain 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 January 9, 2020.

Table of Contents

1. Introduction

Content-control software can be used to filter (i.e., block) web requests that the user, the user's guardian, or the network operator deems objectionable or outside the usage policy of the network. Blocked resource categories can include advertisements, explicit content, known malware, and government-unapproved material, along with many others.

One way to implement content control that does not rely on software or settings on the end-user's computing device is DNS-based content filtering, which examines a client's initial DNS request for the domain providing a resource and then either returns no result or returns an alternate result so that the user is presented with an explanation that filtering has taken place.

DNS-based policy such as content filtering is often built into a network's configured DNS recursive resolver. In addition to blocking a request, the resolver may also log the request for use by the network administrators.

A network operator might wish to provide, or might be obligated to provide, a filtering policy to users of its network. Because such a policy is often enforced by the network operator's default resolver, the use of a technology such as DNS over HTTPS (DoH) [RFC8484] or DNS over TLS (DoT) [RFC7858] can result in bypassing local policies. If the user agent can check for the presence of a policy, this could be used as a signal that the network operator wishes its resolver to be used as a condition of using the network, and that DoH or DoT should be disabled.

At present, there is no standardized mechanism for the user or user agent to identify the presence of a policy on a network's default

resolver without making a request that could trigger the policy and logging thereof, which could have undesirable side-effects.

Therefore, this document defines such a mechanism by defining a so-called "canary domain" that is an instance of Special-Use Domain Names [RFC6761]. DNS requests for this domain would return different results when a DNS-based policy is in place, allowing for the detection of the policy in a consistent way by user agents.

2. Terminology

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. Design Goals and Constraints

This canary domain has been defined with the following design goals and constraints in mind:

- Minimize the risk of exposing personal information
- Ensure that the canary request cannot be mistaken for a user-initiated DNS request
- Ensure that the technique is not specific to any given user agent, policy, or resolver service
- Ensure that the technique is easy to implement for user agents and resolvers

4. Behavior

Resolvers implementing a policy modify the result for the reserved domain 'TBD.arpa', which can be observed by clients to determine if a policy is present.

If a policy exists, the resolver MUST return NXDOMAIN [RFC1035]. If policy is not present, DNS lookup will be successful (i.e., not NXDOMAIN). (This could perhaps resolve to an actual host with a web page managed by IANA, similar to example.com [RFC6761].)

5. Domain Name Reservation Considerations

This section specifies considerations for systems involved in domain name resolution when resolving queries for the reserved domain 'TBD.arpa', in accordance with [RFC6761].

1. Users: Users may invoke command-line DNS lookup tools to resolve the domain, for the purposes of determining if a DNS-based policy is present.

2. Application Software: Application software doing automated lookups are the primary targets of this domain name reservation. Applications can attempt to resolve this name in order to determine if a DNS-based policy is present.

3. Name Resolution APIs and Libraries: Caching servers MUST NOT treat this name as special, unless they implement a policy, in which case they MUST return NXDOMAIN.

4. Caching DNS Servers: Caching servers MUST NOT treat this name as special, unless they implement a policy, in which case they MUST return NXDOMAIN.

5. Authoritative DNS Servers: Authoritative servers other than those supporting the '.arpa' TLD MUST respond to queries for this name with NXDOMAIN.

6. DNS Server Operators: Operators SHOULD ensure that any caching DNS server with a policy on their network properly responds to this name with NXDOMAIN.

7. DNS Registries/Registrars: The defined name is a subdomain of the '.arpa' top-level domain, which is operated by IANA under the authority of the Internet Architecture Board according to the rules established in [RFC3172]. There are no other registrars for '.arpa'.

6. IANA Considerations

IANA is requested to record the domain name 'TBD.arpa' in the "Special-Use Domain Names" registry. See Section 5 for the completed registration template.

[[NOTE TO RFC EDITOR: please change `TBD` to the name assigned by IANA. The name 'dns-content-policy-detection' is suggested.]]

7. Security Considerations

Although a DNS resolution request for the 'TBD.arpa' domain can reveal whether the user or application wishes to detect the presence of DNS-based policy, such a request is relatively neutral compared to a request for a domain that might be subject to a policy.

8. References

8.1. 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>.

[RFC3172] Huston, G., Ed., "Management Guidelines & Operational Requirements for the Address and Routing Parameter Area Domain ("arpa")", BCP 52, RFC 3172, DOI 10.17487/RFC3172, September 2001, <https://www.rfc-editor.org/info/rfc3172>.

[RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <https://www.rfc-editor.org/info/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/info/rfc8174>.

8.2. Informative References

[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>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

Appendix A. Acknowledgements

Thanks to Martin Thomson for his feedback.

Authors' Addresses

Andy Grover
Mozilla

Email: agrover@mozilla.com
URI: https://mozilla.com/

Peter Saint-Andre
Mozilla

Phone: +1 720 256 6756
Email: stpeter@mozilla.com
URI: https://mozilla.com/