Network Working Group
Internet-Draft
Intended status: Informational
Expires: May 20, 2020

J. Levine
Taughannock Networks

November 17, 2019

Signaling That an Authoritative DNS server offers DoT
draft-levine-dprive-signal-02

Abstract

DNS resolvers that wish to use DNS over TLS to authoritative servers (ADoT) need some way to tell whether server offers DoT. This document describes some ways that a server might signal that it uses DoT.

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 May 20, 2020.

Table of Contents

1. Introduction

The Domain Name System[RFC1034] [RFC1035] uses a directed presumably acyclic graph of servers to provide authoritative answers to queries. The link from one server to the next is provided by an NS record in the zone on the upper server that points to the lower server. For zones signed with DNSSEC, the upper server zone contains DS records that contain hashes of signing keys in DNSKEY records in the zone on the lower server.

2. General observations

Depending on your threat model it may be a problem if an intermediate party can intercept the signal and force a DoT client to use unencrypted DNS.

The probe query would generally use query minimization to limit leakage of the requested name. Even so, if a server handles many zones, this leaks the name of the zone being probed.

Some zones have servers run by multiple operators. (The DNS root is a well known example.) It is possible that some of the servers will offer ADoT and some will not. Some of the schemes below handle per-server signals, some don't.

In several of the following schemes, the client probes the server to see whether it offers ADoT. In those cases, the client presumably remembers what servers it's probed so there's only one probe per server.

3. Signaling methods

This is a working list of possible signaling methods. Just because they're in the list doesn't mean that anyone thinks they're a good idea.

3.1. EDNS0 option

We define a new EDNS0 option edns-adot. The client sends an edns-adot option in its request, and the server responds with a value of 0 or 1 to say whether it supports ADoT. The option could be served by the upper level server along with the NS records, which avoids the extra probe, or by the lower level server.

This is easy to implement, but since the OPT isn't signed, it's subject to downgrade attacks. If served by the upper level server, there's no per-server indication, but also no extra round trip.

3.2. DNSKEY flags

A DNSKEY [RFC4034] at the apex of the zone signals that ADoT is available. The simplest approach would be to use one of the unassigned DNSKEY flags to indicate that the zone is expected to be served over ADoT. This is resistant to downgrade, since the DNSKEY is signed, but there's no per-server indication. DNSSEC clients have to fetch the DNSKEY records anyway so there's no extra round trip. Since nobody has ever used DNSKEY records with flag values other than 0, 256, and 257, some software may fail if it sees other flag values.

3.3. Other DNS records

A variation of the previous approach is to put some other kind of DNSSEC signed record at the zone apex that lists nameservers expected to support ADoT, either yet another overloaded TXT record or a new RRTYPE. The list of names would presumably have to be names already listed in NS records (but see the next section.) This provides per-server indication, and is backward compatible, but it makes the DNS Camel sad.

The signal record could also be a signed record in the parent next to the NS records, such as the DSPKI record in [DOTINSECURE].

Another variation puts the signal record at the rDNS name for a nameserver's IP address.

3.4. SRV records

Servers could publish SRV records for ADoT service discovery. ADoT cleints would use the servers identified by SRV instead of the NS servers.

This is downgrade resistant, backward compatible, and allows per-server signalling, even allowing non-standard port numbers. There is potentially an extra round trip for the SRV lookup and more if the name of the servers aren't the same as the NS servers. The number of round trips could be limited if servers provide the SRV and related A/AAAA records as additional data in responses to DNSKEY lookups. It might lead to unpleasant resolution loops if if SRV records use out of bailiwick nameservers.

```
_domain-s._tcp.blah1.example. IN SRV 10 0 853 ns.blah2.example.
. . .
_domain-s._tcp.blah2.example. IN SRV 10 0 4242 ns.blah1.example.
```

3.5. DANE TLSA

If ADoT servers all have DANE secured TLS certificates, the TLSA record can be the ADoT signal.

Publishing a TLSA record is straightforward if a zone is already DNSSEC signed. It's downgrade-resistant, allows per-server signals, and there's no extra round trip beyond what's needed to do the DANE validation.

```
_853._tcp.ns1.blah.example  IN TLSA . . ."
```

3.6. Special server names

Any server that supports ADoT has a name starting with the four characters "XS--". All names starting with two letters other than "XN" and two dashes were reserved when IDNs were invented, so these names are unlikely to collide with any existing names. These are not IDNs, they're just funny looking ASCII names, and you can't do "XN--XS--blah" or anything like that.

This is backward compatible, downgrade resistant, needs no extra round trip, and allows per-server signals. It doesn't allow server names to be IDNs which should not be a big problem since DNS server names are not generally shown to users, although it may confuse people who believe that anything with two dashes must be an IDN.

The Camel is also not crazy about it.

4. References

4.1. References - Normative

[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>.

[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>.

4.2. References - Informative

[DOTINSECURE] Bretelle, M., "DNS-over-TLS for insecure delegations", March 2019, <https://tools.ietf.org/html/draft-bretelle-dprive-dot-for-insecure-delegations-00>.

Author's Address

John Levine
Taughannock Networks
PO Box 727
Trumansburg, NY 14886

Email: standards@taugh.com
URI: http://jl.ly