Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 26, 2019

A. Ghedini
Cloudflare, Inc.

March 25, 2019

Using Early Data in DNS over TLS
draft-ghedini-dprive-early-data-00

Abstract

This document illustrates the risks of using TLS 1.3 early data with DNS over TLS, and specifies behaviors that can be adopted by clients and servers to reduce those risks.

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 September 26, 2019.

Copyright Notice

Copyright (c) 2019 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

TLS 1.3 [TLS13] defines a mechanism, called 0-RTT session resumption or early data, that allows clients to send data to servers in the first round-trip of a connection without having to wait for the TLS handshake to complete.

This can be used to send DNS queries to DNS over TLS [DOT] servers without incurring in the cost of the additional round-trip required by the TLS handshake, and it can be useful in cases where new DNS over TLS connections need to be established often such as on mobile clients where the network might not be stable, or on resolvers where keeping an open connection to many authoritative servers might not be practical.

However, the use of early data allows an attacker to capture and replay the encrypted DNS queries carried on the TLS connection. This can have unwanted consequences and help in recovering information about those queries. While [TLS13] describes tecniques to reduce the likelihood of a replay attack, they are not perfect and still leave some potential for exploitation.

2. Notational Conventions

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. Early Data in DNS over TLS

TODO: talk more about 0-RTT vs. 1-RTT security properties.

A server can signal to clients whether it is willing to accept early data in future connections by providing the "early_data" TLS extension as part of a TLS session ticket, as well as limit the amount of early data it is willing to accept using the "max_early_data_size" field of the "early_data" extension.

In addition to the mitigation mechanisms mandated in [TLS13] that reduce the ability of an attacker to replay early data, but may not completely eliminate it, a server that decided to offer early data to clients MAY reject early data at the TLS layer, or delay the processing of early data to after the handshake is completed.

If the server rejects early data at the TLS layer, a client MUST forget information it optmisitically assumed about the onnection when sending early data, such as the negotiated protocol [ALPN]. Any DNS queries sent in early data will need to be sent again, unless the client decides to abandon them.

TODO: forbid sending DNS updates in early data (RFC2136)? XFR? Other query types?

4. Security Considerations

4.1. Information Exposure

By replaying DNS queries that were captured when transmitted over early data, an attacker might be able to expose information about those queries, even if encrypted.

For example, it's a common behavior for DNS servers to statefully rotate the order of RRs when replying to DNS queries for an RRSet that contains multiple RRs. If the order of rotation is predictable, replaying a captured early data DNS query and observing the order of RRs in DNS responses before and after the replayed query, might allow an attacker to confirm whether the replayed query targeted a specific name that was suspected of being queried without having to decrypt it.

Servers SHOULD either use fixed ordering for multiple RRs in the same DNS response or shuffle the RRs at random, but MUST NOT use stateful and deterministic ordering across multiple queries.

4.2. Denial of Service

Accepting early data exposes a server to potential denial of service through the replay of queries that might be expensive to handle.

When under load, a server MAY reject TLS early data such that the client is forced to retry them after the handshake is completed.

4.3. Privacy

TODO: linkability (e.g. clients changing network, ...) and more?

4.4. Acknowledgments

This document was heavily inspired by [RFC8470]. Daniel Kahn Gillmor and Colm MacCarthaigh also provided important ideas and contributions.

5. References

5.1. Normative References

[DOT] 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>.

[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>.

[TLS13] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

5.2. Informative References

[ALPN] Friedl, S., Popov, A., Langley, A., and E. Stephan, "Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension", RFC 7301, DOI 10.17487/RFC7301, July 2014, <https://www.rfc-editor.org/info/rfc7301>.

[RFC8470] Thomson, M., Nottingham, M., and W. Tarreau, "Using Early Data in HTTP", RFC 8470, DOI 10.17487/RFC8470, September 2018, <https://www.rfc-editor.org/info/rfc8470>.

Author's Address

Alessandro Ghedini
Cloudflare, Inc.

Email: alessandro@cloudflare.com