Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 12, 2017

P. Hoffman
ICANN

April 10, 2017

Running DNS in Existing HTTP/2 Connections
draft-hoffman-dns-in-existing-http2-00

Abstract

Intermediaries such as governments and ISPs spoof DNS responses, and block DNS requests to particular recursive resolvers, for a variety of reasons. They spoof by capturing traffic on port 53, or by redirecting port 853 traffic in the hopes that the client is using opportunistic encryption. They block if they know the address of a resolver that they don't like, such as public resolvers that give honest answers.

This document describes how to run DNS service over existing HTTP/2 connections over TLS, such as those being used for HTTP for basic web service. This design prevents intermediaries from spoofing DNS responses, and makes it impossible for intermediaries to block the use of those recursive resolvers without blocking the desired HTTP connections. It also prevents intermediaries or passive observers from seeing the DNS traffic. This design is meant for communication between a DNS stub resolver and a DNS recursive resolver.

Table of Contents

1. Introduction

HTTP/2 [RFC7540] over TLS is now used widely by many web sites. Large web sites who care about good DNS resolution service (that is, DNS resolution that is not subject to getting wrong answers from intermediaries) might want to offer DNS resolution on the same servers as those running HTTP/2 over TLS. Running DNS over existing HTTP/2 over TLS connection prevents intermediaries from spoofing DNS responses, and makes it impossible for intermediaries to block the use of those recursive resolvers without blocking the desired HTTP connections.

This document covers only the use case of getting DNS service once a HTTP/2 over TLS connection is already set up. That means that the user already has some DNS service before getting to the DNS resolver that is running in the existing HTTP/2 connection. That original DNS service might be standard DNS running on port 53 ([RFC1035]), or DNS-in-TLS running on port 853 ([RFC7858]), or even DNS in its own HTTP/2 over TLS connection that could be defined in the future. Regardless, this document is describing a second DNS service for the user, one

that was bootstrapped by running DNS in a way that might have been spoofed by an intermediary.

A beneficial effect of using DNS over existing HTTP/2 over TLS connections after using DNS over port 53 is that the DNS messages are then encrypted.

A parallel document, [draft-hoffman-dns-in-existing-quic], covers approximately the same use cases as this one, but describes how to carry DNS in QUIC streams.

2. DNS in Existing HTTP/2 over TLS Connections

**** This section, which is the meat of the protocol, is completely tentative. People might have strong opinions on how to best run DNS over HTTP/2. The choice of using a new frame is an early guess for a protocol that meets the desing objectives given above; the HTTPBIS WG might have (much) better alternatives. For example, reserved streams might be a better idea than a new type of frame. ****

This document defines a new type of HTTP/2 frame, "DNS".

DNS in HTTP/2 is run as a stream of DNS frames. The DNS stub resolver opens an HTTP/2 stream if it is not already open. The stub resolver then sends DNS wire-format requests ([RFC1035]), and the recursive resolver sends wire-format requests in the same stream. The wire format used is that for DNS over UDP (not with the extra two-octet header defined in [RFC1035] for TCP). Either side can close the HTTP/2 stream for DNS whenever they wish.

2.1. HTTP/2 DNS Frame Definition

DNS frames (type=0xTBD) convey variable-length sequences of octets associated with a DNS message. One or more DNS frames are used, for instance, to carry a DNS request or response payload.

DNS frames MAY also contain padding. Padding can be added to DNS frames to obscure the size of messages. Padding is a security feature; see Section 4.

The format of the DNS frame is:

Figure 1: DNS frame format

The DNS frame contains the following fields:

Pad Length: An 8-bit field containing the length of the frame padding in units of octets. This field is conditional (as signified by a "?" in the diagram) and is only present if the PADDED flag is set for the frame.

DNS message: The wire-format of the message. The wire format used is that for DNS over UDP (not with the extra two-octet header defined in [RFC1035] for TCP).

Padding: Padding octets that contain no application semantic value. This is handled identically to padding in the DATA frame in [RFC7540].

The DNS frame uses the END_STREAM and PADDED frame flags, identically to the DATA frame in [RFC7540].

DNS frames MUST be associated with a stream. If a DNS frame is received whose stream identifier field is 0x0, the recipient MUST respond with a connection error of type PROTOCOL_ERROR.

DNS frames are subject to flow control identical to the DATA frame in [RFC7540].

2.2. Service Discovery

The DNS stub resolver discovers whether the HTTP/2 server with the exisiting connection supports DNS resolution by attempting to open a DNS stream in the HTTP/2 connection. Because opening a HTTP/2 stream requires sending protocol data, the stub resolver needs to pick a DNS request to use as a probe for DNS resolution service. The stub resolver might send a request for data it actually wants, or it could send a request that it does not care about, such as the A record for example.com.

3. IANA Considerations

This section will eventually have a request to assign a new value, TBD, to the "HTTP/2 Frame Type" registry.

4. Security Considerations

Running DNS over existing HTTP/2 over TLS connections relies on the security of the TLS connections themselves.

A beneficial effect of using DNS over existing HTTP/2 over TLS connections after using DNS over port 53 is that the DNS messages are then encrypted.

*** Copy some text about the uses (and abuses) of padding from Section 10.7 of RFC 7540 here. ***

5. References

5.1. Normative References

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <http://www.rfc-editor.org/info/rfc1035>.

[RFC7540] Belshe, M., Peon, R., and M. Thomson, Ed., "Hypertext Transfer Protocol Version 2 (HTTP/2)", RFC 7540, DOI 10.17487/RFC7540, May 2015, <http://www.rfc-editor.org/info/rfc7540>.

[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, <http://www.rfc-editor.org/info/rfc7858>.

5.2. Informative References

[draft-hoffman-dns-in-existing-quic] Hoffman, P., "Running DNS in Existing QUIC Connections", 2017, <https://tools.ietf.org/id/draft-hoffman-dns-in-existing-quic>.

Author's Address

Paul Hoffman
ICANN

Email: paul.hoffman@icann.org