

How to use URI Resource Records for HTTP
draft-faltstrom-httpbis-dns-01

Abstract

This document describes the reasons why it would be a good thing to use SRV [[RFC2782](#)] or URI [[RFC7553](#)] resource records for HTTP protocol instead of (as of today) just relying on redirects and other mechanisms in the HTTP protocol itself. It also explains how to do it as there are conflicting instructions on how to do it.

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 <http://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 February 12, 2017.

Copyright Notice

Copyright (c) 2016 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 (<http://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	2
2.	URI resource Record	3
3.	HTTP of today	3
4.	The URI Resource Record	4
5.	HTTP using URI Resource Records	5
6.	IANA Considerations	5
7.	Security Considerations	5
8.	Acknowledgements	6
9.	References	6
9.1.	Normative References	6
9.2.	Informative References	6
	Author's Address	7

[1.](#) Introduction

The HTTP protocol has historically been a relatively simple protocol, based on TCP by which a client opens a connection to an IP address, sends a request and get a response. Further evolution of the HTTP protocol has introduced the ability to carry multiple transactions over the same TCP connection and further evolution has made the use of the TCP connection even more efficient. Experimental deployment also exists where UDP is the protocol used.

So far, the IP address used for the peer of the connection has been discovered through a lookup of the hostname in the URI that identifies the resource to be accessed, and similarly the hostname in URIs referred to in the HTTP datastream has been used to locate the host serving these resources (or of course in the HTML or similar formatting in the data returned via HTTP).

This simple model causes a number of issues in deploying HTTP based services; issues that do not exist in other popular protocols like SMTP (for electronic mail) and SIP (for VoIP).

This document tries to explain a few of these weaknesses and why usage of URI and SRV resource record types as a first step before the HTTP connection is established would make deployment, especially at scale, easier and because of that not only life easier for the domain name holder, but also make the connection faster and the experience better.

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#), [[RFC2119](#)].

Faltstrom

Expires February 12, 2017

[Page 2]

2. URI resource Record

After an expert review in February 2011 (see [Appendix A in RFC7553 \[RFC7553\]](#)) IANA allocated RRTYPE 256 for the URI Resource Record Type in the registry named Resource Record (RR) TYPEs and QTYPEs as defined in [BCP 42](#) (at the time [\[RFC6195\]](#)), located at <http://www.iana.org/assignments/dns-parameters>.

Later, in June 2015, [\[RFC7553\]](#) was published explaining how the URI resource record could be used.

3. HTTP of today

HTTP is a protocol that originally was very simple with single transactions over a TCP connection. The client opened the connection, requested a resource and got information back. This has evolved over time and the new HTTP/2 protocol has even more ability; for example multiple resources can be fetched over the same TCP flow using a mechanism that can be regarded as asynchronous.

The basic functionality, though, is the same. A URI is used as an identifier for the resource to be fetched. From the URI a hostname is extracted, an address record (AAAA or A) is fetched and a TCP connection is opened to the address in question. So called Happy Eyeballs mechanisms are sometimes used to open multiple connections in parallel and select the one that works best (from the clients perspective) to maximize the end users experience.

In many cases, though, the resource that is actually to be fetched is not the one that is named with the original URI. In such cases, the HTTP response is a 3xx for a redirect to a different URI, which is then fetched over the same (if possible) or a different TCP connection.

The reasons for such redirects vary, but some common ones are:

- The URI in question does not end with a '/' but the resource (or web server) requires it.

- The URI in question does not start with 'www' but the user did type it (or the other way around). This becomes important in cases where the software serving the content must have one unique hostname component in the URIs it creates and responds to.

- The resource is in reality hosted by some cloud service, and instead of a reverse proxy, a HTTP-level redirect is used to indicate the actual resource location within the cloud service.

The exact URI for the resource is not known, but the identification of the resource is sufficiently known that a so called "well known URI" can be used to explicitly trigger a redirect to the resource itself.

The resource was initially requested using HTTP while the resource is accessible only over HTTPS.

The resource is hosted by a 2nd party and a CNAME is used for the hostname in question. But, as CNAME can not be used at DNS nodes names that already have data, a CNAME can't exist at a zone apex, a limitation not shared by AAAA or A records, necessitating their use. Often address records then go missing at the apex (while the "www.example.com" record is changed into a CNAME for some reason) or the address record is later not updated when the IP address changes. Such situations frequently require multiple address lookups and HTTP redirects before the resource is fetched using the correct intended URI.

In most cases described above, the remembered entry point for human interaction is a string that mostly resembles a domain name, sometimes fully qualified, like "example.com" or "www.example.com". If the string indeed is a domain name it would be beneficial to be able to directly couple it to a complete URI. The URI resource record provides this direct mapping. (If the string is not a domain name, enter the search engines, which are out of scope here.)

4. The URI Resource Record

The URI resource record was approved by the IETF using the then new approval mechanism that did not require an RFC. Later an informational RFC was created that explains the format and usage.

A few things to note:

It might seem that the URI resource record uses the same format for text in RDATA as the TXT resource record, but it does not. The text in the RDATA (the actual URI pointed to) is not a length prefixed string. Such a string would have a too constrained length; 256 octets per . Instead it is defined as the rest of the RDATA field from where the URI starts.

URI inherits prefix label definitions from the ENUM Registry hosted by IANA. This registry, like the registry for ports and services, is not unique, which makes the prefix label to prepend to find "web services" too loosely defined.. This document is clarifying this, see below.

The priority field is used for ordering in what order URIs should be fetched. If the first (lowest number) is not reachable, the 2nd is to be tried etc.

The weight is very seldom used and if not, SHOULD have value 0.

5. HTTP using URI Resource Records

There are multiple ways to set the prefix for the URI resource record by looking at the various tables IANA maintains. There are initiatives to create an ultimate table, like [\[I-D.ietf-dnsop-attrleaf\]](#). This document clarifies what prefix to use when fetching web pages using the HTTP protocol, as the URI specification is ambiguous.

The URI record to look up for the domain example.com is:

`_web._http.example.com.`

When resolving a URI for the web before the HTTP protocol specification is applied to the URI, the URI MUST be rewritten according to the RDATA in the response to the lookup of the URI resource record. If the RRSet returned contains more than one resource record, the records MUST be sorted and tried in accordance with the URI resource record specification.

If no resource record is returned when the URI record is looked up, the HTTP client MUST continue to resolve the URI as it is, without it being rewritten.

Note that whether for example TCP or UDP is to be used is explicitly not specified by setting prefix labels. The decision is instead made according to the HTTP specification in force. The URI resource record only says what the correct URI is to fetch if the goal is to get the web page related to example.com using http.

6. IANA Considerations

Given the registry discussed in [\[I-D.ietf-dnsop-attrleaf\]](#) is created, the value `_web._http` is to be registered for use for normal web services using the HTTP protocol.

7. Security Considerations

Using the URI resource record together with security mechanisms that rely on verification of authentication of hostnames, like TLS, makes it important to choose the correct domain name when doing the comparison, and that the change in what hostname to use is secured by

DNSSEC so that it can be trusted in a similar way as a redirect in HTTP using TLS.

It is specifically the case that although HTTP and HTTPS are different ENUM Service Registrations, only `_web._http` MUST be used for the lookup of the URI resource record for a domain. If HTTPS is the preferred protocol to use, a HTTPS URI is to be returned to the lookup and not a HTTP URI.

8. Acknowledgements

People that have helped me with this document include Daniel Stenberg and Mans Nilsson.

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC7553] Faltstrom, P. and O. Kolkman, "The Uniform Resource Identifier (URI) DNS Resource Record", [RFC 7553](#), DOI 10.17487/RFC7553, June 2015, <<http://www.rfc-editor.org/info/rfc7553>>.

9.2. Informative References

- [I-D.ietf-dnsop-attrleaf]
Crocker, D., "DNS Scoped Data Through '_Underscore' Attribute Leaves", [draft-ietf-dnsop-attrleaf-00](#) (work in progress), March 2016.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", [RFC 2782](#), DOI 10.17487/RFC2782, February 2000, <<http://www.rfc-editor.org/info/rfc2782>>.
- [RFC6195] Eastlake 3rd, D., "Domain Name System (DNS) IANA Considerations", [RFC 6195](#), DOI 10.17487/RFC6195, March 2011, <<http://www.rfc-editor.org/info/rfc6195>>.

Author's Address

Patrik Faltstrom
Netnod

Email: paf@netnod.se