

**Use of SRV records in conjunction with HTTP and URIs"
draft-andrews-http-srv-02**

Abstract

The combined use of SRV records for HTTP along with URIs is not as straight forward as it would appear at first glance. This document looks at the issues involved and recommends solutions.

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 November 21, 2014.

Copyright Notice

Copyright (c) 2014 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.

This document may contain material from IETF Documents or IETF

Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

[1.](#) Introduction [3](#)
[2.](#) URIs without an explicit port specification [3](#)
[3.](#) URIs with a explicit port specification [4](#)
[4.](#) Transitioning Considerations [5](#)
[5.](#) Security Considerations [6](#)
[6.](#) IANA Considerations [6](#)
[7.](#) Normative References [6](#)
 Authors' Addresses [6](#)

1. Introduction

Many of today's HTTP sites are virtual, that is they are hosted on a machine that is not known by the name the HTTP site is known by. This leads to the problem of how to rationally give these HTTP sites IP addresses. This has traditionally been done by using CNAMEs [[RFC1034](#)] [[RFC1035](#)] or by using explicit IP address records where CNAMEs are illegal due to restrictions in the DNS.

Both of these solutions have undesired side effects. CNAMEs are not protocol specific. Using IP address records is a logistic nightmare for large servers with many virtual sites. This is becoming a bigger problem as companies move away from identifying their HTTP site with a "www" prefix and just use their delegated domain name, e.g. "http://example.com/".

Using SRV [[RFC2782](#)] records would seem to be a natural solution to this problem in that they are protocol specific and will work where CNAMEs are illegal in the DNS.

There are problems with doing this without thought however in that URIs [[RFC3986](#)] can specify a port and SRV records do specify a port. When this occurs which one do you honour?

In addition to this SRV records provide for load balancing. For most protocols this is straight forward as there will only be a single connection made. For HTTP however there are often many connections made in a session. Should each of these individual connections be load balanced or should the load balancing be on a per session basis?

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

2. URIs without an explicit port specification

If the URI does not explicitly specify a port to connect to, i.e. the URI does not contain a ":<port>" part, there is no port conflict. In this case a client MUST follow the logic specified in [[RFC2782](#)], including the server selection mechanism provided by the priority and weight fields. If SRV records do not exist then the client MUST fall back to looking for IP address records.

Once a server is selected it SHOULD be continued to be used for the rest of the session if possible after an initial connection is made. If a server has multiple addresses the client SHOULD continue to use the same address while possible taking into consideration ttl values

on address records. If connections to this address fail it SHOULD try the other addresses for the server first before attempting other servers.

The use of a SRV record does not affect the contents of the "Host:" field of the HTTP transaction. Its only effect is to potentially change the address and port the client connects to. All other parts of the HTTP transaction are not affected by the presence of a SRV record.

Examples:

Single SRV record:

```
URI:      http://example.com/
SRV RR:   _http._tcp.example.com. SRV  10 0 8080 host1.example.com.
A RRs:   example.com.             A    10.0.0.2
         host1.example.com.       A    10.0.1.1
```

Connect to: 10.0.1.1 port 8080

Multiple SRV records:

```
URI:      http://example.com/
SRV RRs:  _http._tcp.example.com. SRV  10 1 8080 host1.example.com.
         _http._tcp.example.com. SRV  10 3 8080 host2.example.com.
         _http._tcp.example.com. SRV  20 0 8080 host3.example.com.
A RRs:   example.com.             A    10.0.0.4
         host1.example.com.       A    10.0.1.2
         host2.example.com.       A    10.0.2.2
         host3.example.com.       AAAA 1080::8:800:200C:417A
```

Connect to: 10.0.1.2 port 8080 or 10.0.2.2 port 8080 if either is available (the probability of being selected should be 25% for 10.0.1.2 port 8080, and 75% for 10.0.2.2 port 8080); otherwise, try 1080::8:800:200C:417A port 8080.

3. URIs with a explicit port specification

If the URI does explicitly specify a port, other than the default port, to connect to then there is a potential conflict in the port specification between the URI and the SRV records, and the SRV record is ignored. In this case the user agent MUST query for address records for the host name in the URI (instead of SRV records).

If the server has multiple addresses the client SHOULD continue to use the same address while possible taking into consideration ttl

values on address records.

Note [\[RFC3986\], Section 6.2.3](#). Scheme-Based Normalization states that URIs with a port value equal to the default port (80) are identical to those with no port or a empty port.

Examples:

Default port specified:

```
URI:      http://example.com:80/
SRV RR:   _http._tcp.example.com. SRV  10 1 8080 host2.example.com.
A RRs:    example.com.                A   10.0.0.1
          host2.example.com.          A   10.0.2.2
```

Connect to: 10.0.0.2 port 8080

Non-default port specified:

```
URI:      http://example.com:8080/
SRV RR:   _http._tcp.example.com. SRV  10 1 80 host2.example.com.
CNAME RR: example.com.             CNAME host1.example.com.
A RRs:    host1.example.com.        A   10.0.0.1
          host2.example.com.        AAAA 1080::8:800:200C:417A
```

Connect to: 10.0.0.1 port 8080

4. Transitioning Considerations

When transitioning from using a non-SRV solution to using a SRV based solution old, non-SRV aware, clients will continue to look for address records. It may be necessary to use redirection at the HTTP layer to direct these clients to the new servers if the SRV records point to a different <address, port> tuple.

It will also be necessary to continue to provide the existing address / CNAME records until there is a significant percentage of SRV aware clients. Experience has shown that this should be within one to two years of the introduction of the first SRV aware client.

In cases where you are just trying to replace the A or CNAME record referring to a service providers machine with a SRV record the following should suffice.

The service provider is hosting the service on machine.example.net and you are example.com.

```
example.com.          A   <IP address of machine.example.net>
_http._tcp.example.com. SRV 0 0 80 machine.example.net.
```

5. Security Considerations

The authors believe the algorithm described in this document to not cause any new security problems. However care should be taken as SRV and non-SRV aware clients may be directed to different locations.

6. IANA Considerations

A well known label has to be allocated for the first label of the http SRV record. This document has used "_http".

7. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", [RFC 2782](#), February 2000.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.

Authors' Addresses

M. Andrews
Internet Systems Consortium
950 Charter Street
Redwood City, CA 94063
US

Email: marka@isc.org

T. Kottelin
Laivalahden puistotie 10 C 37
Helsinki FIN-00810
Finland

Email: thor@anta.net