Internet Engineering Task Force (IETF)
Request for Comments: 7553
Category: Informational
ISSN: 2070-1721

P. Faltstrom
Netnod
O. Kolkman
ISOC

June 2015

The Uniform Resource Identifier (URI) DNS Resource Record

Abstract

This document describes the already registered DNS resource record (RR) type, called the Uniform Resource Identifier (URI) RR, that is used for publishing mappings from hostnames to URIs.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7553.

Table of Contents

1. Introduction

This document explains the use of the Domain Name System (DNS) for the storage of URIs [RFC3986] and how to resolve hostnames to such URIs that can be used by various applications using the URI resource record type. For resolution, the application needs to know both the hostname and the protocol that the URI is to be used for. The protocol is registered by IANA.

Historically, uses of the DNS to map a domain name to a URL have relied on the Naming Authority Pointer (NAPTR) RRTYPEs [RFC2915] and then on the Dynamic Delegation Discovery System (DDDS) [RFC3401] application framework with the DNS as a database as specified in RFC 3404 [RFC3404]. This has a number of implications such as the fact the RRSet returned will contain all URIs "connected" with the owner, and not only the ones related to a specific service.

The URI resource record specified in this document enables the querying party to do the equivalent of selecting which of the NAPTR records one is interested in and have only those returned. This is possible because data in the service field of the NAPTR record is included in the owner part of the URI resource record type. It is also the case that as the URI resource record type includes the target URI directly as part of the RDATA, it is very easy to extract the correct target URI, instead of applying rewrite rules as in NAPTR.

Querying for URI resource records is not replacing querying for NAPTR resource records (or use of S-NAPTR [RFC3958]). Instead, the URI resource record type provides a complementary mechanism to be used when one already knows what service field is interesting. With it, one can directly query for the specific subset of the otherwise possibly large RRSet returned when querying for NAPTR resource records.

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, RFC 2119 [RFC2119].

2. Applicability Statement

In general, it is expected that URI records will be used by clients for applications where the relevant protocol to be used is known, but, for example, an extra abstraction is needed in order to separate a domain name from a point of service (as addressed by the URI). One example of such a situation is when an organization has many domain names but only one official web page.

Applications need to know the specific service to prepend the hostname with. Using repetitive queries for URI records is not a replacement for querying for NAPTR records according to the NAPTR (DDDS) or S-NAPTR algorithms. NAPTR records serve the purpose of discovering the various services or the URIs (for looking up access points for a given service). These are two very different kinds of needs.

3. DNS Considerations

Using prefix labels, such as underscored service tags, for a specific owner name may cause a counter-intuitive effect when the owner name is a wildcard name. For example, _s2._s1.*.example.net is not a wildcard name and cannot be used to return a synthesized answer for a query name of _s2._s1.a.example.net. See Section 4.5 of RFC 4592 [RFC4592] for more details. Besides, underscored service tags used for the URI RR (based on the "Service Name and Transport Protocol Port Number Registry") may have slightly different semantics than service tags used for underscored prefix labels that are used in combination with other (yet unspecified) RR types. This may cause subtle management problems when delegation structure that has developed within the context of URI RRs is also to be used for other RR types. Because the service labels might be overloaded, applications should carefully check that the application-level protocol is indeed the protocol they expect.

Subtle management issues may also arise when the delegations from service to sub-service labels involve several parties and different stakeholders.

4. The Format of the URI RR

This is the presentation format of the URI RR:

```
Owner name TTL Class URI Priority Weight Target
```

The URI RR does not cause any kind of Additional Section processing.

4.1. Owner Name, Class, and Type

The URI owner name is subject to special conventions.

Just like the SRV RR [RFC2782], the URI RR has service information encoded in its owner name. In order to encode the service for a specific owner name, one uses service parameters. Valid service parameters are those registered by IANA in the "Service Name and Transport Protocol Port Number Registry" [RFC6335] or as "Enumservice

Registrations [RFC6117]. The Enumservice Registration parameters are reversed (i.e., subtype(s) before type), prepended with an underscore (_), and prepended to the owner name in separate labels. The underscore is prepended to the service parameters to avoid collisions with DNS labels that occur in nature, and the order is reversed to make it possible to do delegations, if needed, to different zones (and therefore providers of DNS).

For example, suppose we are looking for the URI for a service with ENUM Service Parameter "A:B:C" for host example.com. Then we would query for (QNAME,QTYPE)=("_C._B._A.example.com","URI").

As another example, suppose we are looking for the URI for a service with Service Name "A" and Transport Protocol "B" for host example.com. Then we would query for (QNAME,QTYPE)=("_A._B.example.com","URI").

The type number for the URI record is 256.

The URI resource record is class independent.

The URI RR has no special Time-to-Live (TTL) requirements.

4.2. Priority

This field holds the priority of the target URI in this RR. Its range is 0-65535. A client MUST attempt to contact the URI with the lowest-numbered priority it can reach; URIs with the same priority SHOULD be selected according to probabilities defined by the weight field.

4.3. Weight

This field holds the server selection mechanism. The weight field specifies a relative weight for entries with the same priority. Larger weights SHOULD be given a proportionately higher probability of being selected. The range of this number is 0-65535.

4.4. Target

This field holds the URI of the target, enclosed in double-quote characters (’"’), where the URI is as specified in RFC 3986 [RFC3986]. Resolution of the URI is according to the definitions for the Scheme of the URI.

Since the URI will not be encoded as a <character-string> (see Section 3.3 of RFC 1035 [RFC1035]), there is no 255-character size limitation.

The Target MUST NOT be an empty URI ("").

4.5. URI RDATA Wire Format

The RDATA for a URI RR consists of a 2-octet Priority field, a 2-octet Weight field, and a variable-length Target field.

Priority and Weight are unsigned integers in network byte order.

The remaining data in the RDATA contains the Target field. The Target field contains the URI as a sequence of octets (without the enclosing double-quote characters used in the presentation format).

The length of the Target field MUST be greater than zero.

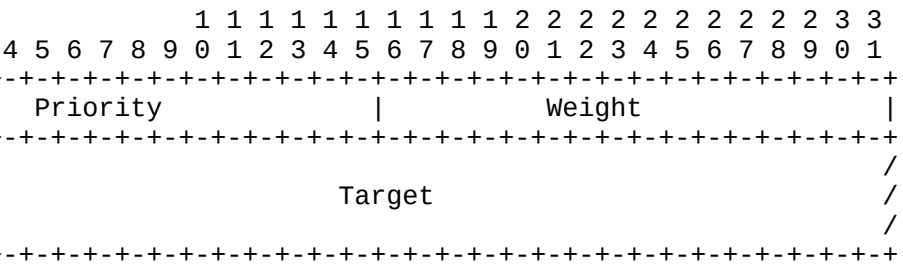

5. Usages

5.1. Example: FTP Server in the example.com Domain

An organization has the domain names example.com and example.net, and their FTP archive is at ftp://ftp1.example.com/public. Given the service name "ftp" and transport protocol "tcp" (from the IANA "Service Name and Transport Protocol Port Number Registry"), the following URI resource records could be made available in the respective zones (example.com and example.net):

```
$ORIGIN example.com.
_ftp._tcp    IN URI 10 1 "ftp://ftp1.example.com/public"

$ORIGIN example.net.
_ftp._tcp    IN URI 10 1 "ftp://ftp1.example.com/public"
```

5.2. Relation to S-NAPTR

The URI resource record type is not a replacement for the S-NAPTR. It is instead an extension and the second step of the S-NAPTR resolution can resolve a URI resource record instead of using SRV records and yet another algorithm for how to use SRV records for the specific protocol.

```
$ORIGIN example.com.
;;                  order pref flags
           IN NAPTR 100   10   "D"   "EM:ProtA"  ( ; service
                    ""                            ; regexp
                    _http._tcp.example.com. )     ; replacement

_http._tcp IN URI   10 1 "http://www.example.com/path"
```

5.3. Relation to U-NAPTR

The URI resource record type, together with S-NAPTR, can be viewed as a replacement for U-NAPTR [RFC4848]. The URI resource record type is only interesting when one know a base domain name, a protocol, and a service so that one can compose the record to look up. NAPTR records of any kind are used to look up what services exist for a certain domain, which is one step before the URI resource record is used.

5.4. Relation to SRV

The URI resource record type can be viewed as a replacement for the SRV record. This is because it, like the SRV record, can only be looked up if one knows the base domain, the protocol, and the service. It has a similar functionality and uses the same registry for service names, but instead of returning a hostname and port number, the URI record returns a full URI. As such, it can be viewed as a more powerful resource record than SRV.

6. IANA Considerations

6.1. Registration of the URI Resource Record Type

After an expert review in February 2011 (see Appendix A), IANA allocated RRTYPE 256 for the URI resource record type in the registry named "Resource Record (RR) TYPEs" (as defined in [BCP42], which was RFC 6195 at the time but has since been replaced by RFC 6895) located at <http://www.iana.org/assignments/dns-parameters>.

IANA has updated the reference for this registration to refer to this RFC.

6.2. Registration of Services

No new registry is needed for the registration of services as the Service Name, Transport Protocol Port Numbers, Enumservices and the DNS SRV Service Type registries are also used for the URI resource record type.

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 ensure that the change in the hostname to be used is secured by DNSSEC so that it can be trusted in a similar way as a redirect in HTTP using TLS.

If, for example, the URI resource record is not signed with the help of DNSSEC and then validated successfully, trusting the non-signed URI will effectively lead to a downgrade attack.

The basic mechanism for successful use of URI works as follows:

1. Announce that example.com is hosted at example.org (with some URL) in DNS.

2. Secure the URI resource record with DNSSEC. This is best done by doing validation in the application doing the lookup, but it could also be done in the local recursive resolver or in the trusted recursive resolver also doing validation. All are according to the local trust policy.

3. Verify the TLS (for example) certificate for the connection to example.org matches, i.e., use the hostname in the URI and not the hostname used originally when looking up the URI resource record.

4. If needed, do application-layer authentication, etc., over the then encrypted connection.

It is also possible that the URI in the resource record type has errors in it. Applications using the URI resource record type for resolution should behave similarly as if the user typed (or copied and pasted) the URI. At least it must be clear to the user that the error is not due to any error from his side.

One SHOULD NOT include userinfo (see "User Information", Section 3.2.1 of RFC 3986 [RFC3986]) in a URI that is used in a URI resource record as DNS data must be viewed as publicly available information.

8. References

8.1. Normative References

[BCP42] Eastlake 3rd, D., "Domain Name System (DNS) IANA Considerations", BCP 42, RFC 6895, April 2013, <http://www.rfc-editor.org/info/bcp42>.

[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>.

[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>.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <http://www.rfc-editor.org/info/rfc3986>.

[RFC6117] Hoeneisen, B., Mayrhofer, A., and J. Livingood, "IANA Registration of Enumservices: Guide, Template, and IANA Considerations", RFC 6117, DOI 10.17487/RFC6117, March 2011, <http://www.rfc-editor.org/info/rfc6117>.

[RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", BCP 165, RFC 6335, DOI 10.17487/RFC6335, August 2011, <http://www.rfc-editor.org/info/rfc6335>.

8.2. Informative References

[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>.

[RFC2915] Mealling, M. and R. Daniel, "The Naming Authority Pointer (NAPTR) DNS Resource Record", RFC 2915, DOI 10.17487/RFC2915, September 2000, <http://www.rfc-editor.org/info/rfc2915>.

[RFC3401] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part One: The Comprehensive DDDS", RFC 3401, DOI 10.17487/RFC3401, October 2002, <http://www.rfc-editor.org/info/rfc3401>.

[RFC3403] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database", RFC 3403, DOI 10.17487/RFC3403, October 2002, <http://www.rfc-editor.org/info/rfc3403>.

[RFC3404] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Four: The Uniform Resource Identifiers (URI)", RFC 3404, DOI 10.17487/RFC3404, October 2002, <http://www.rfc-editor.org/info/rfc3404>.

[RFC3597] Gustafsson, A., "Handling of Unknown DNS Resource Record (RR) Types", RFC 3597, DOI 10.17487/RFC3597, September 2003, <http://www.rfc-editor.org/info/rfc3597>.

[RFC3958] Daigle, L. and A. Newton, "Domain-Based Application Service Location Using SRV RRs and the Dynamic Delegation Discovery Service (DDDS)", RFC 3958, DOI 10.17487/RFC3958, January 2005, <http://www.rfc-editor.org/info/rfc3958>.

[RFC4592] Lewis, E., "The Role of Wildcards in the Domain Name System", RFC 4592, DOI 10.17487/RFC4592, July 2006, <http://www.rfc-editor.org/info/rfc4592>.

[RFC4848] Daigle, L., "Domain-Based Application Service Location Using URIs and the Dynamic Delegation Discovery Service (DDDS)", RFC 4848, DOI 10.17487/RFC4848, April 2007, <http://www.rfc-editor.org/info/rfc4848>.

[RFC5507] IAB, Faltstrom, P., Ed., Austein, R., Ed., and P. Koch, Ed., "Design Choices When Expanding the DNS", RFC 5507, DOI 10.17487/RFC5507, April 2009, <http://www.rfc-editor.org/info/rfc5507>.

Appendix A. The Original RRTYPE Allocation Request

On February 22, 2011 IANA assigned RRTYPE 256 for the URI resource record based on a request that followed the procedure documented in [BCP42] (which was RFC 6195 at the time but has since been replaced by RFC 6895). The DNS RRTYPE PARAMETER ALLOCATION form as submitted to IANA at that time is replicated below for reference.

Note: Although "ownername" should be "owner name", "ownername" has been preserved below because it was part of the original request form submitted to IANA.

A. Submission Date:

May 23, 2009

B. Submission Type:

[X] New RRTYPE
[ ] Modification to existing RRTYPE

C. Contact Information for submitter:

Name: Patrik Faltstrom
Email Address: paf@cisco.com
International telephone number: +46-8-6859131
Other contact handles:
(Note: This information will be publicly posted.)

D. Motivation for the new RRTYPE application?

There is no easy way to get from a domain name to a URI. Some mechanisms exists via use of the NAPTR [RFC3403] resource record. That implies quite complicated rules that are simplified via the S-NAPTR [RFC3958] specification. But, the ability to directly look up a URI still exists. This specification uses a prefix based naming mechanism originated in the definition of the SRV [RFC2782] resource record, and the RDATA is a URI, encoded as one text field.

See also above (Section 1).

E. Description of the proposed RR type.

The format of the URI resource record is as follows:

Ownername TTL Class URI Priority Weight Target

The URI RR has service information encoded in its ownername. In order to encode the service for a specific ownername one uses service parameters. Valid service parameters used are either Enumservice Registrations registered by IANA, or prefixes used for the SRV resource record.

The wire format of the RDATA is as follows:

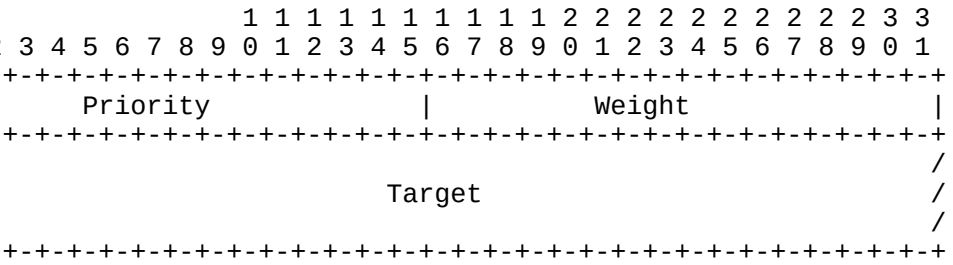

F. What existing RRTYPE or RRTYPEs come closest to filling that need and why are they unsatisfactory?

The RRTYPE that come closest is the NAPTR resource record. It is for example used in the DDDS and S-NAPTR algorithms. The main problem with the NAPTR is that selection of what record (or records) one is interested in is based on data stored in the RDATA portion of the NAPTR resource record. This, as explained in RFC 5507 [RFC5507], is not optimal for DNS lookups. Further, most applications using NAPTR resource records uses regular expression based rewrite rules for creation of the URI, and that has shown be complicated to implement.

The second closest RRTYPE is the SRV record that given a prefixed based naming just like is suggested for the URI resource record, one get back a port number and domain name. This can also be used for creation of a URI, but, only URIs without path components.

G. What mnemonic is requested for the new RRTYPE (optional)?

URI

H. Does the requested RRTYPE make use of any existing IANA Registry or require the creation of a new IANA sub-registry in DNS Parameters?

Yes, partially.

One of the mechanisms to select a service is to use the Enumservice Registry managed by IANA. Another is to use services and protocols used for SRV records.

I. Does the proposal require/expect any changes in DNS servers/resolvers that prevent the new type from being processed as an unknown RRTYPE (see [RFC3597])?

No

J. Comments:

None

Acknowledgements

Ideas on how to split the two different kinds of queries, "What services exists for this domain name" and "What is the URI for this service", came from Scott Bradner and Lawrence Conroy. Other people that have contributed to this document include Richard Barnes, Leslie Daigle, Victor Dukhovni, Olafur Gudmundsson, Philip Hallam-Baker, Ted Hardie, Sam Hartman, Evan Hunt, John Klensin, Peter Koch, Eliot Lear, Andy Newton, Mark Nottingham, Penn Pfautz, Jinmei Tatuya, Willem Toorop, and Nico Williams.

Cisco is acknowledged as Mr. Faltstrom's employer at the time this document was developed.

The NLnet Labs is acknowledged as Mr. Kolkman's employer at the time this document was developed.

Authors' Addresses

Patrik Faltstrom
Netnod

EMail: paf@netnod.se

Olaf Kolkman
Internet Society

EMail: kolkman@isoc.org