

Domain-Based Application Service Location Using URIs and
the Dynamic Delegation Discovery Service (DDDS)

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

The purpose of this document is to define a new, straightforward Dynamic Delegation Discovery Service (DDDS) application to allow mapping of domain names to URIs for particular application services and protocols. Although defined as a new DDDS application, dubbed U-NAPTR, this is effectively an extension of the Straightforward NAPTR (S-NAPTR) DDDS Application.

Table of Contents

1. Introduction	3
2. Straightforward URI-Enabled NAPTR (U-NAPTR)	3
2.1. Permitted Flags	3
2.2. Permitted Regular Expressions	4
3. Sample U-NAPTR DNS Records	4
4. Formal Definition of U-NAPTR Application of DDDS	5
4.1. Application Unique String	5
4.2. First Well Known Rule	5
4.3. Expected Output	5
4.4. Flags	5
4.5. Service Parameters	5
4.5.1. Application Services	6
4.5.2. Application Protocols	6
4.6. Valid Rules	6
4.7. Valid Databases	7
5. IANA Considerations	7
6. Security Considerations	8
7. Acknowledgements	8
8. References	8
8.1. Normative References	8
8.2. Informative References	9

1. Introduction

The purpose of this document is to define a new, straightforward Dynamic Delegation Discovery Service (DDDS) [7] application to allow mapping of domain names to URIs for particular application services and protocols. This allows the "lookup" of particular services available for given domains, for example.

Although this is defining a new and separate DDDS Application, dubbed U-NAPTR, it is built from the same principles as the Straightforward NAPTR (S-NAPTR) application, specified in [2]. This specification is not an update of S-NAPTR, but the reader is encouraged to review that document for extensive coverage of motivation and implementation considerations.

S-NAPTR provides for application service location that does not rely on rigid domain naming conventions. It is deemed "straightforward" in part because it rules out the use of regular expressions in NAPTR records (for the S-NAPTR DDDS Application). However, that also rules out the possibility of providing a URI as the target of DDDS resolution. A number of applications, specified (e.g., [9]) and proposed, find the restriction too limiting, making S-NAPTR a near miss to suit their needs.

This U-NAPTR is effectively a modest extension to S-NAPTR, to accommodate the use of URIs as targets, without allowing the full range of possible regular expressions in NAPTR records.

2. Straightforward URI-Enabled NAPTR (U-NAPTR)

This document assumes the reader is familiar with the S-NAPTR specification [2]. The intention of U-NAPTR is to provide everything that S-NAPTR does, except that it allows the use of the "U" flag in the NAPTR record, and a specific form of REGEXP.

2.1. Permitted Flags

U-NAPTR permits the same flags as S-NAPTR ("S", "A", or empty), plus the "U" Flag. For the U-NAPTR DDDS Application, the presence of the "U" Flag in the NAPTR record indicates the REGEXP field must be populated (and, consequently, the REPLACEMENT field is empty). The regular expression in the REGEXP field must be of the limited form described below, and the result of the regular expression evaluation will be a URI that is the result of the DDDS resolution.

2.2. Permitted Regular Expressions

U-NAPTR permits regular expressions of a form that does a complete replacement of the matched string with a URI, expressed as a constant string. This is essentially a dodge around the fact that the REPLACEMENT field in NAPTR is required to produce only a fully qualified domain name (and, therefore, cannot be used for a URI).

The specific allowed syntax for U-NAPTR regular expressions is:

```
u-naptr-regexp = "!.*!"<URI>!"
```

where <URI> is as defined in STD 66 [8], the URI syntax specification.

With this limited form of regular expression, applications using U-NAPTR need not implement full regular expression parsers.

3. Sample U-NAPTR DNS Records

In the sample NAPTR RRs for example.com shown below, "WP" is the imagined application service tag for "white pages", and "EM" is the application service tag for an imagined "Extensible Messaging" application service.

```
example.com.
; ; order pref flags
IN NAPTR 100 10 "" "WP:whois++" ( ; service
    "" ; regexp
    bunyip.example.com. ; replacement
)
IN NAPTR 100 20 "s" "WP:ldap" ( ; service
    "" ; regexp
    _ldap._tcp.myldap.example.com. ; replacement
)
IN NAPTR 200 10 "u" "EM:protA" ( ; service
    "!.*!protA://someisp.example.com!" ; regexp
    "" ; replacement
)
IN NAPTR 200 30 "a" "EM:protB" ; service
    "" ; regexp
    myprotB.example.com. ; replacement
)
```

4. Formal Definition of U-NAPTR Application of DDDS

This section formally defines the DDDS Application, as described in [7].

4.1. Application Unique String

The Application Unique String is a fully qualified domain name (FQDN) for which an authoritative server for a particular service is sought.

4.2. First Well Known Rule

The "First Well Known Rule" is identity -- that is, the output of the rule is the Application Unique String, the FQDN for which the authoritative server for a particular service is sought.

4.3. Expected Output

The expected output of this Application is the information necessary to connect to authoritative server(s) (host, port, protocol, or URI) for an application service within a given domain.

4.4. Flags

This DDDS Application uses only 3 of the Flags defined for the URI/URN Resolution Application [5]: "S", "A", and "U". No other Flags are valid. If a client obtains a NAPTR RR for a U-NAPTR-using application that contains any other flag, that NAPTR RR should be ignored and processing continues with the next record (if any).

These flags are for terminal lookups. This means that the Rule is the last one and that the flag determines what the next stage should be. The "S" flag means that the output of this Rule is a FQDN for which one or more SRV [3] records exist. "A" means that the output of the Rule is a domain name and should be used to lookup address records for that domain. "U" means that the output of the Rule is a URI that should be resolved in order to obtain access to the described service.

Consistent with the DDDS algorithm, if the Flag string is empty the next lookup is for another NAPTR record (for the replacement target).

4.5. Service Parameters

Service Parameters for this Application take the form of a string of characters that follow this ABNF [1]:

```

service-params = [ [app-service] *(":" app-protocol)]
app-service    = experimental-service / iana-registered-service
app-protocol   = experimental-protocol / iana-registered-protocol
experimental-service    = "x-" 1*30ALPHANUMSYM
experimental-protocol   = "x-" 1*30ALPHANUMSYM
iana-registered-service = ALPHA *31ALPHANUMSYM
iana-registered-protocol = ALPHA *31ALPHANUMSYM
ALPHA                = %x41-5A / %x61-7A ; A-Z / a-z
DIGIT                = %x30-39 ; 0-9
SYM                  = %x2B / %x2D / %x2E ; "+" / "-" / "."
ALPHANUMSYM         = ALPHA / DIGIT / SYM
; The app-service and app-protocol tags are limited to 32
; characters and must start with an alphabetic character.
; The service-params are considered case-insensitive.

```

Thus, the Service Parameters may consist of an empty string, just an app-service, or an app-service with one or more app-protocol specifications separated by the ":" symbol.

Note that this is similar to, but not the same as the syntax used in the URI DDDS application [5]. The DDDS DNS database requires each DDDS application to define the syntax of allowable service strings. The syntax here is expanded to allow the characters that are valid in any URI scheme name (see [8]). Since "+" (the separator used in the RFC3404 service parameter string) is an allowed character for URI scheme names, ":" is chosen as the separator here.

4.5.1. Application Services

The "app-service" must be an IANA-registered service; see Section 5 for instructions on registering new application service tags.

4.5.2. Application Protocols

The protocol identifiers that are valid for the "app-protocol" production are standard, registered protocols; see Section 5 for instructions on registering new application protocol tags.

4.6. Valid Rules

Permitted rules are substitution rules and regular expressions of the following syntax (i.e., a regular expression to replace the domain name with a URI):

```
u-naptr-regexp = "!.*!"<URI>!"
```

where <URI> is as defined in STD 66 [8], the URI syntax specification.

4.7. Valid Databases

At present, only one DDDS Database is specified for this Application. [4] specifies a DDDS Database that uses the NAPTR DNS resource record to contain the rewrite rules. The Keys for this database are encoded as domain names.

The First Well Known Rule produces a domain name, and this is the Key that is used for the first lookup -- the NAPTR records for that domain are requested.

DNS servers MAY interpret Flag values and use that information to include appropriate NAPTR, SRV, or A records in the Additional Information portion of the DNS packet. Clients are encouraged to check for additional information but are not required to do so. See the Additional Information Processing section of [4] for more information on NAPTR records and the Additional Information section of a DNS response packet.

5. IANA Considerations

This document does not itself place any requirements on IANA, but provides the basis upon which U-NAPTR-using services can make use of the existing IANA registries for application service tags and application protocol tags (defined in RFC 3958 [2]).

As is the case for S-NAPTR, all application service and protocol tags that start with "x-" are considered experimental, and no provision is made to prevent duplicate use of the same string. Use them at your own risk.

All other application service and protocol tags are registered based on the "specification required" option defined in [6], with the further stipulation that the "specification" is an RFC (of any category).

There are no further restrictions placed on the tags other than that they must conform with the syntax defined above (Section 4.5).

The defining RFC must clearly identify and describe, for each tag being registered:

- o Application protocol or service tag
- o Intended usage
- o Interoperability considerations

- o Security considerations (see Section 6 of this document for further discussion of the types of considerations that are applicable)
- o Any relevant related publications

The defining RFC may also include further application-specific restrictions, such as limitations on the types of URIs that may be returned for the application service.

6. Security Considerations

U-NAPTR has the same considerations for security as S-NAPTR; see Section 8 of [2]. U-NAPTR has the additional consideration that resolving URIs (from the result of the DDDS resolution) has its own set of security implications, covered in the URI specification (in particular, Section 7 of [8]). In essence, using DNSSEC, client software can be confident that the URI obtained using U-NAPTR is indeed the one specified by the administrator of the domain from which it was retrieved; but the validity of the service reached by resolving that URI is a matter of URI resolution security practices.

7. Acknowledgements

Thanks to Martin Thomson, John Klensin, Bernard Aboba, Alfred Hoenes, Dan Romascanu, Suresh Krishnan, and Lars Eggert for reviewing earlier versions and catching errors!

8. References

8.1. Normative References

- [1] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 4234, October 2005.
- [2] Daigle, L. and A. Newton, "Domain-Based Application Service Location Using SRV RRs and the Dynamic Delegation Discovery Service (DDDS)", RFC 3958, January 2005.
- [3] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.
- [4] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database", RFC 3403, October 2002.

- [5] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part Four: The Uniform Resource Identifiers (URI)", RFC 3404, October 2002.
- [6] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

8.2. Informative References

- [7] Mealling, M., "Dynamic Delegation Discovery System (DDDS) Part One: The Comprehensive DDDS", RFC 3401, October 2002.
- [8] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", RFC 3986, STD 66, January 2005.
- [9] Malamud, C., "Attaching Meaning to Solicitation Class Keywords", RFC 4095, May 2005.

Author's Address

Leslie L. Daigle
Cisco Systems
13615 Dulles Technology Drive
Herndon, VA 20171
US

E-Mail: ledaigle@cisco.com; leslie@thinkingcat.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

