Internet-Draft <u>draft-ietf-svrloc-advertising-05.txt</u> Expires in six months Ryan Moats AT&T Martin Hamilton Loughborough University February 1999

Advertising Services (Providing information to support service discovery) Filename: <u>draft-ietf-svrloc-advertising-05.txt</u>

Status of This Memo

This document is an Internet-Draft and is in full conformance with all provisions of <u>Section 10 of RFC2026</u>. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/lid-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Abstract

This document proposes a solution to the problem of finding information about that services are being offered at a particular Internet domain, based on deployment experience with the Netfind White Pages directory software.

This approach makes it possible to supply clients with more information than the DNS aliases that have been widely deployed in this role - notably the port numbers being used by servers. However, it is not without problems, and we have tried to take account of these.

1. Rationale

There is no one single way of discovering the network services and application protocols supported at a particular Internet domain. The Domain Name System (DNS - [1,2]) provides some basic facilities for finding the hosts that offer particular services, such as DNS servers themselves (NS records), and mail exchangers (MX records). It does not provide a mechanism for locating arbitrary servers of arbitrary protocols, or a search capability.

In addition to the name and IP address(es) of the host offering the service, clients also sometimes require further information to make effective use of the service - e.g. TCP or UDP port numbers, protocol version information, and information about the protocol options supported by the server. Another example would be the organization's base within the X.500 [3] Directory Information Tree (DIT), that is needed for X.500 browsing and searching.

At the time of our initial draft, it was common practice to "hint" at the services and protocols offered at a given domain via DNS aliases. For example, the alias "www.internic.net" implies that the HTTP server for the domain "internic.net" is running on TCP port 80 of the machine (or machines) that answer to the name "www.internic.net." A slight formalization of this approach [4] has been accepted. Other schemes have been suggested for explicitly registering services either by DNS extensions, as in [5], or by dedicated directory services such as X.500.

Several mechanisms have been suggested to address the problem of finding this service information, ranging from new DNS record types to dedicated directory services. No single one of these would-be solutions has as yet gained the competitive edge. If the lengthy gestation period to date is anything to go by, it seems that we can expect even more delay before there is any widespread deployment unless there is a "killer application" that forces the issue.

2. Where Netfind has gone before

The Netfind software $[\underline{6}, \underline{7}]$ follows what has been proposed in <u>RFC 1588</u> [$\underline{8}$]: using URLs [$\underline{9}$] for passing directory service information to clients. It uses stylized Text (TXT) record encoding within the DNS and currently understands the following "White Pages" URLs:

White Pages URL	Information
wp-noop://	This site should not be visited
wp-dap:// <sb></sb>	X.500 search base for the site

[Page 2]

```
e.g. wp-dap://o=Loughborough%20University,c=GB
wp-ph://host/port
                         Suggests CCSO nameserver [10]
wp-whois://host/port
                         Suggests WHOIS [11] server
wp-smtp-expn-finger://host Use the SMTP [12] EXPN command,
                           and the finger [13] protocol
wp-smtp-expn://host/port
                         Suggests the SMTP EXPN command
wp-finger://host/port
                         Suggests the finger protocol
wp-telnet://host/port
                         Suggests a text based info
                           service that should be used
                           via telnet [<u>14</u>]
_____
```

Note that the notation "protocol://host/port" is used, rather than the "protocol://host:port" format that is being standardized for generic URLs.

Note also that these URL schemes have not all been standardized, although wp-ph and wp-whois may be accommodated by translation to the widely supported Internet Gopher Protocol [15].

<u>3</u>. A simple interim solution

In this document, we propose that the "service:" URL scheme as described in [<u>17</u>] be encoded in DNS NAPTR records [<u>18</u>] as a solution. With this scheme, software agents would do a DNS lookup on <service>.<domain name>. The NAPTR record associated with <service>.<domain name> would have the following syntax:

<service> IN NAPTR [preference] [weight] "u" "" "!^.*\$!<service
url>!" .

Most of these fields are discussed in [18]. Of note is the regular expression field, which is the central feature of this proposal. As URLs routinely use the "/" character to donate hierarchy, the regular expression should be delimited by the some other character. For purposes of this document, we propose using "!", although another character can be used so long as the service URL does not contain that character. Finally, the regular expression must conform to the requirements of [18].

Further, the Service URL scheme supports the concept of abstract service types that are useful for white pages service advertisement. For general white pages discovery, we propose that software agents do a DNS lookup on wp.<domain name>. Here, the NAPTR records would contain URLs of the "wp:" abstract service type as documented in [19]. For example, the NAPTR records for wp.lut.ac.uk could be written as

[Page 3]

3.1 Why not TXT records?

Previous versions of this document proposed using TXT records for storing the URL information. Rather than continue to propose using TXT records, we have decided to propose using NAPTR records for storing this information as this is the type of information that NAPTR was intended to hold when it was developed.

<u>4</u>. Further details and usage scenarios

4.1. Finding "White Pages" information

This case is already catered for by the Netfind "wp-" prefix. To advertise their White Pages services explicitly, a site would create one or more TXT records under both wp and the service being advertised, e.g.

whois IN NAPTR 10 0 "u" "" "!^.*\$!service:wp:whois://whois.lut.ac.uk/!"

ph IN NAPTR 10 0 "u" "" "!^.*\$!service:wp:ccso://cso.lut.ac.uk/2!" .

Another example showing the possibility of multiple protocols for accessing a service would be (the domain for this example is aecom.yu.edu):

```
ns IN NAPTR 10 0 "u" ""
"!^.*$!service:wp:gopher://gopher.aecom.yu.edu/2!" .
ns IN NAPTR 20 0 "u" "" "!^.*$!service:wp:http://www.middlebury.edu/
cgi-bin/WebPh?other_ph_servers!" .
ns IN NAPTR 30 0 "u" ""
"!^.*$!service:wp:http://faker.ncsa.uiuc.edu:8080/
cgi-bin/phfd!" .
```

```
It is envisaged that this information could be used in some
```

scenarios. Assuming their Internet domain is already known, mail user agents with integrated support could offer to do directory service lookups to determine a correspondent's address from their

Expires 8/31/99

[Page 4]

name, to verify the contents of address books, and to determine alternative email addresses should delivery fail. This last technique might also be applied by lower level mail delivery software.

4.3. Public key lookup

Attempts to build a scalable infrastructure for the distribution of public key information, in particular for the public keys of individuals, have been hampered by the lack of a convention that could be used to suggest the public key servers for a site or organization.

For these examples, we postulate an abstract service type "keys:", e.g.

keys IN NAPTR 10 0 "u" "" "!^.*\$!service:keys:finger://mrrl.lut.ac.uk!"

```
lut.ac.uk IN NAPTR 10 0 "u" ""
".^.*$!service:keys:finger://mrrl.lut.ac.uk!" .
```

It does not, however, address the issue of public key (certificate) format. It is expected that this would be taken care of by format negotiation in the protocol or protocols being used to do the lookup.

Public key lookup would be of immediate use in software that has integrated support for public key authentication, signing and encryption - e.g. mail and news user agents.

<u>4.4</u>. Finding "Yellow Pages" information

By "Yellow Pages" we mean a catch-all category: information about services offered that do not fall into any of the above categories. For this, we propose using the "yp:" abstract service type described in [<u>19</u>].

For example, consider the case of a machine that is running a HTTP server - but not on the IANA registered default port (80)

www IN A 204.179.186.65 IN A 198.49.45.10 IN A 192.20.239.132 IN NAPTR 10 0 "u" "" "!^.*\$!service:yp:http:// www.ds.internic.net:8888/!" .

yp IN A 204.179.186.65 IN A 198.49.45.10 IN A 192.20.239.132 IN NAPTR 10 0 "u" "" "!^.*\$!service:yp:http://

Expires 8/31/99

[Page 5]

www.ds.internic.net:8888/!" .

This "Yellow Pages" mechanism provides a means for DNS maintainers to effectively register the existence of their major network services. This can have a variety of uses - e.g. the service information is available to any "web crawler" type applications that might choose to index it, and to interactive applications such as World-Wide Web browsers, that might use it to override their default behavior.

<u>4.5</u> Finding "Directory Agent" information

The Service Location Protocol [20] provides the ability to search for services according to their characteristics, as opposed to solely by type. This is useful to clients that need to discover a particular service in a case where a domain offers more than one service of the same type and the services are not identical. Clients can discover what types of services are supported, the attributes of those services and can obtain service URLs by issuing structured queries. Thus, a service may be discovered by description as opposed to by name.

To do this, the Service Location Protocol defines a scheme for Directory Agent discovery. The term "directory" in this context refers to a directory of services as opposed to a directory of "white pages" information that is used elsewhere in this document. A site may wish to present services to hosts outside its domain may elect to set up a Directory Agent (with the remote registration features turned off, see [20]) outside its firewall. A client supporting the service location protocol would then make queries for individual services inside the domain. The Directory Agent would be found via the following DNS entries:

5. Limitations of this approach

Note that older DNS servers may not support the NAPTR record type. This is because support for SRV and NAPTR have only recently been added to the BIND code base.

Some resolvers are not capable of requesting a NAPTR record, or not capable of generating any DNS lookup requests other than a simple

[Page 6]

address lookup. NAPTR records can actually be requested by setting the question type in the request to 35 (decimal), regardless of the symbolic names defined by the stack's resolver code. Implementing more advanced resolver functionality when the stack only provides address lookup requires a little work, but sample code is freely available.

The size limitations on DNS packets will have some effect on the number of URLs that can be associated with a domain name using NAPTR records. Response packets are subject to truncation if they grow to over 576 bytes.

Characters that are illegal in URLs must be escaped, for example:

"service:wp:ldap://ldap.lut.ac.uk/o=Loughborough%20University%20of %20Technology,c=GB"

Domain name compression is normally used to reduce the size of the response packet needed for a given domain name. Clearly, this will not be possible on arbitrary strings embedded within the response packet.

Widespread use of NAPTR records in the role proposed by this document would increase the amount of information held in nameserver caches, and in particular might cause problems where negative cacheing is concerned. Consequently we suggest that clients use them as a fall back mechanism if more conventional methods, such as DNS aliases, prove unproductive.

7. Security Considerations

Since this draft proposes to use DNS for storage of URL information, all the normal security considerations for applications that depend on the DNS apply. The DNS is open to many kinds of "spoofing" attacks, and it cannot be guaranteed that the result returned by a DNS lookup is indeed the genuine information. Spoofing may take the form of denial of service, such as directing of the client to a nonexistent address, or a passive attack such as an intruder's server that masquerades as the legitimate one.

Work is ongoing to remedy this issue insofar as the DNS is concerned [16]. In the meantime, note that stronger authentication mechanisms such as public key cryptography with large key sizes are a pre-requisite if the DNS is being used in any sensitive environment. Examples of these would be on-line financial transactions, and any scenario where privacy is a concern - such as the querying of medical records over the network. Strong encryption of the network traffic may also be advisable, to protect against TCP connection "hijacking"

[Page 7]

and packet sniffing.

There are some additional considerations that are specific to URLs. Specifically, client applications should be wary of URLs that direct them to alternative Internet domains and/or unusual port numbers. They should also be proactive when passing URLs to external programs, to ensure that the user's environment is not exposed to malevolent meta-characters. Finally, implementors should take care to avoid buffer overruns when processing these DNS response packets.

8. Conclusions

Whilst far from ideal, we believe the approach outlined in this document does provide a workable interim solution to the problem of locating the network services offered at a particular Internet domain - particularly when used in combination with DNS aliases, as outlined in [4]. Suitable DNS server software is already widely deployed, and client support may be implemented without any great difficulty.

It is debatable whether any of this is strictly necessary. Certainly there is less work involved in adding a few lines to an existing DNS server configuration than in setting up a whole new directory service, such as X.500. From this point of view, a new DNS resource record type or types would perhaps address the problem more effectively, but it may be some time before any new types are widely deployed.

9. Acknowledgments

Special thanks to Erik Guttman for his help with the service location example, information on the "service:" scheme, as well as much e-mail in working out the service schemes proposed here. Thanks to Tim Howes, Sri Sataluri and members of the IETF SVRLOC and IDS working groups for their comments on earlier drafts of this document. This document is partially supported by the National Science Foundation, Cooperative Agreement NCR-9218179, the UK Electronic Libraries Programme (eLib) grant 12/39/01, and the European Commission's Telematics for Research Programme grant RE 1004.

The format used for representing Netfind White Pages URLs within the DNS was originally defined by Mike Schwartz, with help from Carl Malamud and Marshall Rose. The Netfind work was supported in part by grants from the National Science Foundation, the Advanced Research Projects Agency, Sun Microsystems' Collaborative Research Program, and AT&T Bell Laboratories.

Some of the points in the security considerations section were drawn from $[\underline{4}]$.

[Page 8]

10. References

Request For Comments (RFC) and Internet Draft documents are available from numerous sites.

- [1] P. V. Mockapetris. "Domain names concepts and facilities," <u>RFC 1034</u>. November 1987.
- [2] P. V. Mockapetris. "Domain names implementation and specification," <u>RFC 1035</u>. November 1987.
- [3] C. Weider, J. Reynolds, S. Heker. "Technical Overview of Directory Services Using the X.500 Protocol," <u>RFC 1309</u>. March 1992.
- [4] M. Hamilton, R. Wright. "Use of DNS Aliases for Network Services," <u>RFC 2219</u>, October 1997.
- [5] A. Gulbrandsen, P. Vixie. "A DNS RR for specifying the location of services (DNS SRV)," <u>RFC 2052</u>. October 1996.
- [6] M. F. Schwartz. "Netfind Support for URL-Based Search Customization," June 28, 1994. <URL:ftp://ftp.cs.colorado.edu/pub/cs/distribs/ netfind/Netfind.WP.URLs>
- [7] M. F. Schwartz, C. Pu. "Applying an Information Gathering Architecture to Netfind: A White Pages Tool for a Changing and Growing Internet," University of Colorado Technical Report CU-CS-656-93. December 1993, revised July 1994. <URL:ftp://ftp.cs.colorado.edu/pub/cs/techreports/ schwartz/Netfind.Gathering.txt.Z>
- [8] J. Postel, C. Anderson. "White Pages Meeting Report," <u>RFC 1588</u>. February 1994.
- [9] T. Berners-Lee, L. Masinter & M. McCahill. "Uniform Resource Locators (URL)," <u>RFC 1738</u>. December

[Page 9]

1994.

- [10] R. Hedberg, S. Dorner, and P. Pomes. "The CCSO Nameserver (Ph) Architecture," <u>RFC 2378</u>, August 1998.
- [11] K. Harrenstien, M. K. Stahl, E.J. Feinler. "NICNAME/WHOIS," <u>RFC 954</u>. October 1985.
- [12] D. Crocker. "Standard for the format of ARPA Internet text messages," <u>RFC 822</u>. August 1982.
- [13] D. Zimmerman. "The Finger User Information Protocol," <u>RFC 1288</u>. December 1992.
- [14] J. Postel, J.K. Reynolds. "Telnet Protocol specification," <u>RFC 855</u>. May 1983.
- [15] F. Anklesaria, M. McCahill, P. Lindner, D. Johnson, D. Torrey & B. Albert. "The Internet Gopher Protocol (a distributed document search and retrieval protocol)," <u>RFC 1436</u>. March 1993.
- [16] D. E. Eastlake 3rd, C. W. Kaufman. "Domain Name System Security Extensions," <u>RFC 2065</u>, January 1997.
- [17] E. Guttman, C. Perkins, J. Kempf, "The service: URL Scheme," Internet Draft (work in progress), October 1997.
- [18] R. Daniel, M. Mealling. "Resolution of Uniform Resource Identifiers using the Domain Name System," <u>RFC 2168</u>, June 1997.
- [19] R. Moats, "The "wp" and "yp" Abstract Services," Internet Draft (work in progress), February, 1997.

[Page 10]

[20] J. Veizades, E. Guttman, C. Perkins, S. Kaplan, "Service Location Protocol," <u>RFC 2165</u>, June 1997.

<u>11</u>. Authors' addresses

Ryan Moats AT&T 15621 Drexel Circle Omaha, NE 68135-2358 USA

Phone: +1 402 894-9456 EMail: jayhawk@att.com

Martin Hamilton Department of Computer Studies Loughborough University of Technology Leics. LE11 3TU, UK

Email: m.t.hamilton@lut.ac.uk

[Page 11]