

Network Working Group
Request for Comments: 2219
BCP: 17
Category: Best Current Practice

M. Hamilton
Loughborough University
R. Wright
Lawrence Berkeley Laboratory
October 1997

Use of DNS Aliases for Network Services

Status of this Memo

This document specifies an Internet Best Current Practices for the Internet Community, and requests discussion and suggestions for improvements. Distribution of this memo is unlimited.

Abstract

It has become a common practice to use symbolic names (usually CNAMEs) in the Domain Name Service (DNS - [RFC-1034, RFC-1035]) to refer to network services such as anonymous FTP [RFC-959] servers, Gopher [RFC-1436] servers, and most notably World-Wide Web HTTP [RFC-1945] servers. This is desirable for a number of reasons. It provides a way of moving services from one machine to another transparently, and a mechanism by which people or agents may programmatically discover that an organization runs, say, a World-Wide Web server.

Although this approach has been almost universally adopted, there is no standards document or similar specification for these commonly used names. This document seeks to rectify this situation by gathering together the extant 'folklore' on naming conventions, and proposes a mechanism for accommodating new protocols.

It is important to note that these naming conventions do not provide a complete long term solution to the problem of finding a particular network service for a site. There are efforts in other IETF working groups to address the long term solution to this problem, such as the Server Location Resource Records (DNS SRV) [RFC-2052] work.

1. Rationale

In order to locate the network services offered at a particular Internet domain one is faced with the choice of selecting from a growing number of centralized databases - typically Web or Usenet News "wanderers", or attempting to infer the existence of network services from whatever DNS information may be available. The former approach is not practical in some cases, notably when the entity seeking service information is a program.

Perhaps the most visible example of the latter approach at work is in the case of World-Wide Web HTTP servers. It is common practice to try prefixing the domain name of an organization with "http://www." in order to reach its World-Wide Web site, e.g. taking "hivnet.fr" and arriving at "http://www.hivnet.fr." Some popular World-Wide Web browsers have gone so far as to provide automatic support for this domain name expansion.

Ideally, the DNS or some complementary directory service would provide a means for programs to determine automatically the network services which are offered at a particular Internet domain, the protocols which are used to deliver them, and other technical information. Unfortunately, although much work has been done to develop said directory service technologies and to define new types of DNS resource record to provide this type of information, there is no widely agreed upon or widely deployed solution to the problem - except in a small number of cases.

The first case is where the DNS already provides a lookup capability for the type of information being sought after. For example: Mail Exchanger (MX) records specify how mail to a particular domain should be routed [RFC-974], the Start of Authority (SOA) records make it possible to determine who is responsible for a given domain, and Name Server (NS) records indicate which hosts provide DNS name service for a given domain.

The second case is where the DNS does not provide an appropriate lookup capability, but there is some widely accepted convention for finding this information. Some use has been made of Text (TXT) [RFC-1035] records in this scenario, but in the vast majority of cases a Canonical Name (CNAME) or Address (A) record pointer is used to indicate the host or hosts which provide the service. This document proposes a slight formalization of this well-known alias approach.

It should be noted that the DNS provides a Well Known Services (WKS) [RFC-1035] lookup capability, which makes it possible to determine the network services offered at a given domain name. In practice this is not widely used, perhaps because of the absence of a suitable programming interface. Use of WKS for mail routing was deprecated in the Host Requirements specification [RFC-1123] in favour of the MX record, and in the long term it is conceivable that SRV records will supersede both WKS and MX.

2. A generic framework

Our approach to dealing with aliases for protocols is straightforward. We define a standard set of DNS aliases for the most popular network services that currently exist (see the "Special Cases" section below). For protocols that are not explicitly listed in this document, the protocol specification must propose a name.

3. Special cases

Special Cases:

Alias	Service
archie	archie [ARCHIE]
finger	Finger [RFC-1288]
ftp	File Transfer Protocol [RFC-959]
gopher	Internet Gopher Protocol [RFC-1436]
ldap	Lightweight Directory Access Protocol [RFC-1777]
mail	SMTP mail [RFC-821]
news	Usenet News via NNTP [RFC-977]
ntp	Network Time Protocol [RFC-1305]
ph	CCSO nameserver [PH]
pop	Post Office Protocol [RFC-1939]
rwhois	Referral WHOIS [RFC-1714]
wais	Wide Area Information Server [RFC-1625]
whois	NICNAME/WHOIS [RFC-954]
www	World-Wide Web HTTP [RFC-1945]

4. (Ab)Use of the DNS as a directory service

The widespread use of these common aliases effectively means that it is sometimes possible to "guess" the domain names associated with an organization's network services, though this is becoming more difficult as the number of organizations registered in the DNS increases.

It should be understood by implementors that the existence of a DNS entry such as

```
www.hivnet.fr
```

does not constitute a registration of a World-Wide Web service. There is no requirement that the domain name resolve to an IP address or addresses. There is no requirement that a host be listening for

HTTP connections, or if it is, that the HTTP server be running on port 80. Finally, even if all of these things are true, there can be no guarantee that the World-Wide Web server will be prepared to honor requests from arbitrary clients.

Having said this, the aliases do provide useful "hints" about the services offered. We propose that they be taken in this spirit.

The conventions described in this document are, essentially, only useful when the organization's domain name can be determined - e.g. from some external database. A number of groups, including the IETF, have been working on ways of finding domain names given a set of information such as organization name, location, and business type. It is hoped that one or more of these will eventually make it possible to augment the basic lookup service which the DNS provides with a more generalized search and retrieval capability.

5. DNS server configuration

In the short term, whilst directory service technology and further types of DNS resource record are being developed, domain name administrators are encouraged to use these common names for the network services they run. They will make it easier for outsiders to find information about your organization, and also make it easier for you to move services from one machine to another.

There are two conventional approaches to creating these DNS entries. One is to add a single CNAME record to your DNS server's configuration, e.g.

```
ph.hivnet.fr. IN CNAME baby.hivnet.fr.
```

Note that in this scenario no information about ph.hivnet.fr should exist in the DNS other than the CNAME record. For example, ph.hivnet.fr could not contain a MX record.

An alternative approach would be to create an A record for each of the IP addresses associated with ph.hivnet.fr, e.g.

```
ph.hivnet.fr. IN A 194.167.157.2
```

It isn't a simple matter of recommending CNAMEs over A records. Each site has its own set of requirements that may make one approach better than the other. RFC 1912 [RFC-1912] discusses some of the configuration issues involved in using CNAMEs.

Recent DNS server implementations provide a "round-robin" feature which causes the host's IP addresses to be returned in a different order each time the address is looked up.

Network clients are starting to appear which, when they encounter a host with multiple addresses, use heuristics to determine the address to contact - e.g. picking the one which has the shortest round-trip-time. Thus, if a server is mirrored (replicated) at a number of locations, it may be desirable to list the IP addresses of the mirror servers as A records of the primary server. This is only likely to be appropriate if the mirror servers are exact copies of the original server.

6. Limitations of this approach

Some services require that a client have more information than the server's domain name. For example, an LDAP client needs to know a starting search base within the Directory Information Tree in order to have a meaningful dialogue with the server. This document does not attempt to address this problem.

7. CCSO service name

There are currently at least three different aliases in common use for the CCSO nameserver - e.g. "ph", "cso" and "ns". It would appear to be in everyone's interest to narrow the choice of alias down to a single name. "ns" would seem to be the best choice since it is the most commonly used name. However, "ns" is also being used by DNS to point to the DNS server. In fact, the most prevalent use of "ns" is to name DNS servers. For this reason, we suggest the use of "ph" as the best name to use for CCSO nameservers.

Sites with existing CCSO servers using some of these aliases may find it desirable to use all three. This increases the likelihood of the service being found.

As noted earlier, implementations should be resilient in the event that the name does not point to the expected service.

8. Security Considerations

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 non-existent address, or a passive attack such as an intruder's server which masquerades as the legitimate one.

Work is ongoing to remedy this situation insofar as the DNS is concerned [RFC-2065]. In the meantime it should be noted 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 situations. Examples of these would be on-line financial transactions, and any situation 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" and packet sniffing.

9. Conclusions

The service names listed in this document provide a sensible set of defaults which may be used as an aid in determining the hosts which offer particular services for a given domain name.

This document has noted some exceptions which are either inherently unsuitable for this treatment, or already have a substantial installed base using alternative aliases.

10. Acknowledgements

Thanks to Jeff Allen, Tom Gillman, Renato Iannella, Thomas Lenggenhager, Bill Manning, Andy Powell, Sri Sataluri, Patrik Faltstrom, Paul Vixie and Greg Woods for their comments on draft versions of this document.

This work was supported by UK Electronic Libraries Programme (eLib) grant 12/39/01, the European Commission's Telematics for Research Programme grant RE 1004, and U. S. Department of Energy Contract Number DE-AC03-76SF00098.

11. References

Request For Comments (RFC) documents are available from <URL:ftp://ftp.internic.net/rfc> and numerous mirror sites.

- [ARCHIE] A. Emtage, P. Deutsch. "archie - An Electronic Directory Service for the Internet", Winter Usenix Conference Proceedings 1992. Pages 93-110.
- [PH] R. Hedberg, S. Dorner, P. Pomes. "The CCSO Nameserver (Ph) Architecture", Work in Progress.
- [RFC-768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, August 1980.

- [RFC-793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.
- [RFC-821] Postel, J., "Simple Mail Transfer Protocol", STD 10, RFC 821, August 1982.
- [RFC-954] Harrenstien, K., Stahl, M., and E. Feinler, "NICNAME/WHOIS", RFC 954, October 1985.
- [RFC-959] Postel, J., and J.K. Reynolds, "File Transfer Protocol", STD 9, RFC 959, October 1985.
- [RFC-974] Partridge, C., "Mail routing and the domain System", STD 14, RFC 974, January 1986.
- [RFC-977] Kantor, B., and P. Lapsley, "Network News Transfer Protocol", RFC 977, February 1986.
- [RFC-1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.
- [RFC-1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC-1123] Braden, R., "Requirements for Internet hosts - application and support", STD 3, RFC 1123, October 1989.
- [RFC-1288] Zimmerman, D., "The Finger User Information Protocol", RFC 1288, December 1992.
- [RFC-1305] Mills, D., "Network Time Protocol (Version 3) Specification, Implementation", RFC 1305, March 1992.
- [RFC-1436] Anklesaria, F., McCahill, M., Lindner, P., Johnson, D., Torrey, D., and B. Albert, "The Internet Gopher Protocol (a distributed document search and retrieval protocol)", RFC 1436, March 1993.
- [RFC-1590] Postel, J., "Media Type Registration Procedure", RFC 1590, March 1994.
- [RFC-1625] St. Pierre, M., Fullton, J., Gamiel, K., Goldman, J., Kahle, B., Kunze, J., Morris, H., and F. Schiettecatte, "WAIS over Z39.50-1988", RFC 1625, June 1994.
- [RFC-1700] Reynolds, J.K., and J. Postel, "ASSIGNED NUMBERS", STD 2, RFC 1700, October 1994.

- [RFC-1714] Williamson, S., and M. Kesters, "Referral Whois Protocol (RWhois)", RFC 1714, November 1994.
- [RFC-1777] Yeong, W., Howes, T., and S. Kille, "Lightweight Directory Access Protocol", RFC 1777, March 1995.
- [RFC-1912] Barr, D., "Common DNS Operational and Configuration Errors", RFC 1912, February 1996.
- [RFC-1939] Myers, J., and M. Rose, "Post Office Protocol - Version 3", STD 53, RFC 1939, May 1996.
- [RFC-1945] Berners-Lee, T., Fielding, R., and H. Nielsen, "Hypertext Transfer Protocol -- HTTP/1.0", RFC 1945, May 1996.
- [RFC-2052] Gulbrandsen, A., and P. Vixie, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2052, October 1996.
- [RFC-2065] Eastlake, D., and C. Kaufman, "Domain Name System Security Extensions", RFC 2065, January 1997.

12. Authors' Addresses

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

E-Mail: m.t.hamilton@lut.ac.uk

Russ Wright
Information & Computing Sciences Division
Lawrence Berkeley National Laboratory
1 Cyclotron Road, Berkeley
Mail-Stop: 50A-3111
CA 94720, USA

E-Mail: wright@lbl.gov