Marcos Sanz
DENIC eG
Gerhard Winkler
NIC.AT

Internet Draft
Document: draft-sanz-whois-srv-00.txt

Expires: October 2003

April 2003

Using DNS SRV records to locate whois servers

Status of this Memo

Abstract

Whois servers are used to locate administrative, technical and security contacts for given IP addresses, domain names or other network objects associated with an organisation, e.g. AS numbers. While usually Top Level Domain (TLD) registries run a whois server, there is no generic name for it and it may not even be obvious that the TLD registry's whois server is the right one to ask, since there are TLDs where registration takes place under specialised second level domains (e.g. UK, AT). The Regional Internet Registries (RIR) also provide whois service as part of their coordination task.

All this can be solved by central "master" or "meta" whois servers, which keep track of all new and changing servers and refer to the DNS registries' or RIRs' whois servers.

This document proposes a DNS-based approach which eliminates the need for a central master repository and works down to lower levels in the hierarchy. It is the intent to locate a whois server as close to the

target (in terms of hierarchy) as possible, while preserving the opportunity to locate higher level servers for escalation purposes.

Conventions used in this document

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 RFC-2119 [2].

Other terms used in this document are defined in the DNS specification RFC-1034 [3].

Table of Contents

1. Format

The general format of DNS SRV records is documented in RFC 2782:

```
_Service._Proto.Name TTL Class SRV Priority Weight Port Target
```

Therefore the simplest format of an SRV record to locate a whois server is:

```
_nicname._tcp     IN   SRV   0 0 43 whois.nic.example.
```

The symbolic name of the service is defined as "nicname" (case insensitive) and the protocol is TCP based, as per RFC 954 [4].

Priority and Weight have a value of 0 in the example above just for readability purposes.

Target and Port (in the example "whois.nic.example." and "43") have to be substituted with the values the administrator has chosen for the whois server.

2. Usage

The service record functionality is meant as an extension to the existing whois service and not as a new service. If there is a whois server running for a specific domain, such an SRV record can be defined. When used for looking up information about a domain, whois clients can do DNS lookups for SRV records, and can use the retrieved target information to point their whois queries accordingly. This kind of client is called "SRV-cognizant" or "SRV-aware" whois client.

It is imaginable that this functionality could be extended for other purposes (like IP address space allocation), but this remains open for a future discussion.

3. Domain search strategy

There are two different approaches in general. They both have advantages and disadvantages and will be discussed below.

3.1 Top-Down model

The whois client parses the domain name to be looked up. Then the client issues a DNS query for "_nicname._tcp" (QTYPE="SRV", QCLASS="IN") in the TLD of that domain.

If the answer is positive, the whois client processes the returned SRV record(s) according to the algorithm defined in RFC 2782 [5] in order to discover the whois server to be queried. The whois client targets now the original whois query to the identified whois server.

Regardless of the existence/absence of SRV records at the TLD of the

domain (or at any other level), the whois client SHOULD continue querying for SRV records in the subdomains of the previous original domain name, up to the point where that domain name itself is reached. Any returned SRV record does not provide any information about the existence/absence of a service with the same name on subdomains or zones above or below.

For instance:

If the whois client has to look up the domain "very.weird.example.", in order to locate the corresponding whois server, it CAN do following DNS queries looking for SRV records:

QNAME="_nicname._tcp.example.", QTYPE="SRV", QCLASS="IN"

QNAME="_nicname._tcp.weird.example.", QTYPE="SRV", QCLASS="IN"
QNAME="_nicname._tcp.very.weird.example.", QTYPE="SRV", QCLASS="IN"

Regardless of the existence/absence of DNS search lists, if the Top-Down model approach is used, this search strategy should be applied.

3.2 Bottom-Up model

The whois client takes the complete name, including the leaf element, and issues a DNS query for "_nicname._tcp" (QTYPE="SRV", QCLASS="IN") in the corresponding domain.

If the answer is positive, the whois client processes the returned SRV record(s) according to the algorithm defined in [5] in order to discover the whois server to be queried. The whois client targets now the original whois query to the identified whois server.

If the answer is not positive the client strips the leftmost element from the name and the query process is repeated; so it walks the DNS tree upwards.

This process is repeated until a SRV record is found or the TLD is reached.

Clients SHOULD continue the search after they have got a positive answer to look for more additional answers.

To avoid unnecessary load on the DNS root servers, a client MUST NOT

ask for a whois server for the root domain, i.e. it MUST NOT issue queries for an SRV at "_nicname._tcp.".

For instance:

If the whois client has to look up the name "www.very.weird.example.", in order to locate the corresponding whois server, it CAN do following DNS queries looking for SRV records:

QNAME="_nicname._tcp.www.very.weird.example.", QTYPE="SRV", QCLASS="IN"
QNAME="_nicname._tcp.very.weird.example.", QTYPE="SRV", QCLASS="IN"
QNAME="_nicname._tcp.weird.example.", QTYPE="SRV", QCLASS="IN"
QNAME="_nicname._tcp.example.", QTYPE="SRV", QCLASS="IN"

Regardless of the existence/absence of DNS search lists, if the Bottom-Up model approach is used, this search strategy should be applied.

3.3 Conclusion

The Top-Down model follows the idea that information of domains is stored at a central place (relative within a TLD) as it is handled like a global resource. This resource is centrally managed and delegated and the delegation information is a critical element of the resource data.

The Bottom-Up model follows the common idea that information should be looked up as close as possible to the requested object of the query. This goes much more for a decentralized structure of storing information (e.g. organisations could setup their internal whois server for storing local data).

The strategy recommended for domain search clients is that it does not stop at the first positive answer independent of Top-Down or Bottom-Up strategy.

Clients MAY allow to switch between both strategies.

A general purpose client SHOULD default to the Bottom-Up model but specific clients heavily used for domain name lookups SHOULD use the Top-Down model to reduce DNS load and unnecessary lookups.

4. Clarifications

The SRV-cognizant whois client MUST NOT modify the domain name to be looked up in the whois server, independently of the domain source of the SRV record.

In the absence of a whois protocol whose specification calls for the use of other weighting information, the field Weight in the SRV record keeps the standard meaning specified in [5].

As defined in [5] the client SHOULD abort if it finds a record like:

```
_nicname._tcp     IN   SRV   0 0 0 .
```

This means the SRV processing SHOULD be aborted at that level, since that record is an explicit statement that the service is not supported there. But nothing avoids the client to search for other SRV records above or below that level.

There is no definition of which target should be used by an SRV-cognizant whois client if no whois server could be discovered by means of SRV records. The client MAY try addressing the whois query to "whois".<domain> (cf. RFC 2219 [6]). The use of a default whois server is local dependent.

5. Authority

There is no authority which defines who should run a whois server. At present, ICANN requires the operation of whois servers by registries of gTLDs, and best practice guidelines for ccTLDs recommend the operation of such a service as well. This means, most of the SRV-cognizant whois clients would already get an SRV record after the first DNS query when following the Top-Down strategy described in this document. However, if the client decides searching for SRV records below that level, more than one whois server could be discovered. There is no authority, and obviously no algorithm, that defines which whois server or whois answer is the right one.

6. Related Work at IETF

[7] describes the requirements for the directory services of Internet

registries (specifically, domain name registries), which are not specific to any protocol. [7] requires these services to use DNS in order to determine the authoritative source of information about domain names.

[8] describes an architectural framework for locating and retrieving information about network resources using LDAP. Although based on a different application level protocol, this document aligns with the query processing model for domains described in [8].

Security Considerations

The same security considerations as defined in [5] should apply.

There is no discussion on security, data protection and privacy relating to the contents of the whois server in this paper. This is a responsibility of the whois server operator and has nothing to do with a mechanism that describes how whois servers can be discovered.

The strategies described in this document could allow an organisation, by means of DNS query logging, to find out who is issuing whois queries about them even without operating a whois server themselves.

The strategy described in section 3.2 could allow an organisation to misdirect whois requests to their own whois server containing false information or no information at all.

An SRV-cognizant whois client should always display, together with the whois data, the whois server it is getting its data from.

References

1 Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

2 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

3 Mockapetris, P., "Domain names - concepts and facilities", RFC 1034, November 1987

4 Harrenstien, K., "NICNAME/WHOIS", RFC 954, October 1985

5 Gulbrandsen, A., "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000

6 Hamilton, M., "Use of DNS Aliases for Network Services", BCP 17, RFC 2219, October 1997

7 [ldap-whois] Hall, E., "The Internet Resource Query Service and the WHOIS Resource Schema", draft-hall-ldap-whois-02, work in progress

8 [crisp-req] Newton, A., "Cross Registry Internet Service Protocol (CRISP) Requirements", draft-ietf-crisp-requirements-02, work in progress

Acknowledgments

We would like to thank Linus Corin, Kim Davies and Peter Koch among others for their useful input.

Author's Addresses

Marcos Sanz
DENIC eG
Wiesenhuettenplatz 26
D-60329 Frankfurt/Main, Germany
Email: sanz@denic.de

Gerhard Winkler
Vienna University Computer Center / NIC.AT
Universitaetsstrasse 7
A-1100 Vienna, Austria
Email: gerhard.winkler@univie.ac.at