Domain Name System Operations
(dnsop)
Internet-Draft
Expires: January 10, 2006

J. Massar
Unfix / SixXS

July 9, 2005

The _service domain and prefix
draft-massar-dnsop-service-00

Status of this Memo

Copyright Notice

Abstract

This document defines a new domain, _service., which can be used for automatic service configuration and discovery. The associated anycast prefixes can be used to configure a default DNS server, which provides lookups for a local _service. domain but also acts as a (caching) recursive DNS server, thus allowing DNS clients to use this well-known address as their default DNS server as well as to use it to find various well known services, thus avoiding manual configuration.

Table of Contents

1. Requirements notation

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 [RFC2119].

2. Introduction

Currently there are a number of methods of configuring a (caching) recursive DNS server into the resolver code of various clients. Amongst these methods are static configuration, DHCP, IPv6 Router Advertisement options, PPP configuration and a multitude of others.

Next to the configuration of a DNS server, the user of the client will also have to configure various other services, eg Outgoing SMTP server, incoming POP3 or IMAP4 server, HTTP proxy server, FTP proxy server or with the advent of IPv6, one of the various IPv6 tunneling techiques and one day one of the various IPv4 tunneling techniques, to allow IPv6 hosts to contact IPv4 hosts.

All these services now have to be manually configured or using some kind of automation, which is likely to be different for each type of service.

SRV records [RFC2782] defines a method of defining services, it does not however define where these records are located nor handles the case when a service has multiple protocols.

This document specifies a method, which allows vendors to hardcode a well known anycast prefix into their resolving clients. This anycast prefix contains a single well known IP address that runs a (caching) recursive DNS server. [RFC1034] [RFC1035]. This server thus allows looking up of all available domains (example.com, example.org etc). Additionally it is provides lookups for the _service. domain which contains the in this document described domain to be used for autoconfiguration.

3. The _service domain

The _service domain contains PTR records to their respective SRV records. Service names should abide SRV naming rules, they are protocol independent though.

Typical contents of a _service domain.

```
$ORIGIN _service.
@            TXT "Example Networks"
@            RP   helpdesk.example.net. helpdesk.people.example.net.
_website     PTR _https._tcp.example.net.
             PTR _http._tcp.example.net.
_pgpkey      PTR _pgpkey-https._tcp.example.net.
             PTR _pgpkey-http._tcp.example.net.
_imap        PTR _imap._tcp.example.net.
_mailsubmit  PTR _submission._tcp.example.net.
_ntp         PTR _ntp._tcp.example.net.
             PTR _ntp._udp.example.net.
```

The TXT record specifies which organisation is announcing this prefix. This record can be shown in browser functions or to the enduser. The RP record specifies a contact for this service zone.

The above defines a website, available over HTTP and HTTPS, based on the priority and weights given by their SRV records. This allows one to specify that a client must first try the HTTPS variant, if it does not work, or the client does not understand this protocol it can try the HTTP variant. Due to the nature of SRV, these services might be located on different hosts. The NTP service is defined to use either UDP or NTP.

The actual _service. Top Level Domain (TLD) can actually also be a DNAME to the organisations domain, this way the organisation only has to maintain one .service domain.

```
_service. DNAME _service.example.net.
```

This also allows an organisation have multiple domain names or clients domains, simply adding the DNAME to the service domain allows the users to pick any of the domains and the configuration information is already available to the user. This is naturally only required when the user is not inside the anycast range of the organisation.

A _service domain SHOULD be AXFR'able [RFC1995] this to facilitate browsing of the service zone. An organisation MAY opt to decline AXFR's based on their policy.

3.1 DNS Search Path

Lookups in the _service domain should be done according to the DNS search path. Thus if the DNS search path of host is: example.com example.net then the resolver should try the items in:

```
_service.
_service.example.com.
_service.example.net.
```

This allows the _service domain to be located anywhere in the search path of the client. Additionally this also allows one to specify a remote domain and thus having components be configured based on the service values given for that domain.

3.2 Browsing for services

In case the user wants to know what kind of services are available for her, as provided by the local organisation, an application could try to AXFR the _service zones of the domains in the search path and then displaying the available services.

4. The _service anycast address

x.x.x.1 and xxxx::1 (*IANA UPDATE!*) can be hardcoded into any client resolver. These addresses point to the IPv4 and IPv6 variant of a resolver which also provides access to the _service prefix.

On the _service anycast address a full recursive DNS server is responding. It must also provide the lookup facility for the _service. domain. This domain might reside on another DNS server.

This address can be passed to the client using DHCP [RFC2131], other automatic configuration methods or manual configuration.

4.1 The _service prefix

x.x.x.1 is part of x.x.x.1/32 and xxxx::1 is part of xxxx::/64.

The prefix SHOULD only be announced in the local IGP of the organisation. The prefix MAY be announced to other organisations when the two parties agree on this setup. The prefix MUST not be seen globally around the world. Though it is of course possible, having the prefix available to the global internet would not have much of a function as the services are most likely only provided for users of the organisation.

4.2 Monitoring

Any organisation corresponding to this specification must include a monitoring function, to check that the _service is operational. The router must stop injecting the route leading to the server immediately if it detects that the DNS function is not operational.

A remote host should try to query either the TXT or RP record of the to be monitored service zone to see if the DNS server still answers queries. Other methods can of course also be used.

4.3 Discovery and failover

The DNS client resolver send packets to the DNS service by sending them to the anycast address. These packets will reach the closest service provided by their organisation or by another organisation.

When a client does not have connectivity to this prefix, there will be no routing entry for the anycast prefix and thus a destinatation unreachable will be sent to the host. The resolver then learns that the DNS service in question is not available.

When a _service server somehow breaks it should stop announcing the anycast prefix to the local network. At that point, the local IGP will automatically compute a route towards the "next best" _service server. We expect that adequate monitoring tools will be used to guarantee timely discovery of connectivity losses and should allow seamless functionality for the endusers.

5. Security Considerations

This anycast technique introduces an risk, that a rogue router or a rogue AS could introduce a bogus route to their own resolver setup providing rogue _service entries, thereby diverting the traffic to the service they want. Any service using cleartext passwords and having no additional security, eg TLS/SSL, can thus be easily transformed into password collection setups. Care must be taken that nobody can insert a faked _service server into a network.

6. IANA Considerations

IANA will need to mark the _service. domain and the _service sub domain as reserved to be used solely for this purpose.

IANA will need to allocate a IPv6 /64 and a IPv4 /32 for the purposes of having a well known anycast address in which the (caching) recursive DNS server can operate.

The prefix is only required in the organisation itself and should not be carried in intra-domain routing tables. A global prefix is required so that the prefix can be shared between organisations. eg an organisation providing this service to other autonomous systems.

7. References

7.1 Normative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2782] Gulbrandsen, A., Vixie, P., and L. Esibov, "A DNS RR for specifying the location of services (DNS SRV)", RFC 2782, February 2000.

7.2 Informative References

[RFC1995] Ohta, M., "Incremental Zone Transfer in DNS", RFC 1995, August 1996.

[RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

Author's Address

Jeroen Massar
Unfix / SixXS
Hofpoldersingel 45
Gouda 2807 LW
NL

Email: jeroen@unfix.org
URI: http://unfix.org/~jeroen/

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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