Internet Engineering Task Force
INTERNET-DRAFT
Feb 20, 2004
Expires Aug 19, 2005

S. Yamamoto
H. Yokota
KDDI R&D Labs
C. Williams
KDDI Labs USA
A. Durand
no affiliation

Service Discovery using NAPTR records in DNS
<draft-yamamoto-naptr-service-discovery-01.txt>

Status of this Memo

Abstract

This document describes a method to store and retrieve local configuration information from the DNS using NAPTR records in the reverse path DNS tree. It works for both IPv4 in-addr.arpa and IPv6 ip6.arpa.

1. Introduction

DHCP is the protocol of choice to pass configuration information and perform service discovery in well defined environment. However, defining and getting new options for DHCP is a slow process, as it requires not only standardization steps but also need to be

implemented in all potential clients and server.

This memo proposes a new approach based on NAPTR records in the reverse path tree of the DNS. Defining new options for experimental purposes can be done with very little to no code change in the clients and none on the server.

This protocol can be deployed independently of DHCP and only requires the operation of a regular DNS server.

This protocol is also suitable when the administrative authority who manages the service is different from the administrative authority who manages DHCP. This is true in particular in deployment where local DHCP servers do not communicate with the DHCP server run by the entity that manages the service to be discovered.

Using the reverse path DNS tree instead of the forward path DNS tree has three major advantages:

- it does not require to discover the domain name used by the entity managing the service,
- it does not require to reserve any label,
- it matches nicely the underlying topology.

2. NAPTR Record

2.1 NAPTR

NAPTR records are defined in RFC3403 [1]. The format of the NAPTR RR, whose DNS type code is 35, is:

```
   NAPTR order        16 bits
         preference   16 bits
         flags        character-string
         service      character-string
         regexp       character-string
         replacement  domain-name
```

2.2 Defining Services

When defining a new type of configuration or a new service to be discovered, one has only to standardize the different relevant NAPTR parameters, the most important being the name of the "service" tag. For the sake of illustration, the following services are defined.

2.2.1 Isatap

The isatap router service discovery within a site can be done using the following record:

```
   flags = "",
   service = "isatap",
   regexp =  "",
```

```
   replacement = Fully Qualified Domain Name of the isatap router
```

For example:

```
   flags = "",
   service = "isatap",
   regexp =  "",
   replacement = "r21.example.com"
```

2.2.2 Tunnel Broker

The Tunnel Broker discovery within an ISP can be done using the following record:

```
   flags = "",
   service = "TB",
   regexp =  "",
   replacement = Fully Qualified Domain Name of the Tunnel Broker
```

For example:

```
   flags = "",
   service = "TB",
   regexp =  "",
   replacement = "tb.example.com"
```

2.3 Populating the DNS

The administrative authority in charge of the service to be discovered using this method will populate the reverse path DNS tree associated to the address space it controls with the relevant records.

For example, a site deploying isatap will put isatap NAPTR records for every single node of the site in the reverse path DNS tree in the form:

```
9.1.6.10.in-addr.arpa. 0 IN NAPTR 10 10 "" isatap "" r21.example.com
```

In another example, an ISP deploying a tunnel broker service will put TB NAPRT records for every single node in the reverse path DNS tree for all its customers in the form:

```
9.1.6.1.in-addr.arpa. 0 IN NAPTR 10 10 "" TB "" tb.example.com
```

The administrative entity in charge of the reverse path DNS tree can use several methods to populate the tree with NATPR records. It can use scripts to generate the zone, use wildcards or use some extension to the auto-generation methods present in most DNS servers.

When wildcards are used, they are only working on the last level of delegation. That is, if there is a zone delegated under the zone

where the wildcard is placed, that zone won't be covered by the wildcard. In practice, this means putting a wildcard in every terminating zone. This is not a problem in the reverse tree, as those zones usually already exist at the subnet boundaries for the PTR records and are most of the time populated via scripts. Note that using wildcards does not prevent to populate more specific address with different NAPTR records, as long as they are on the same zone.

Note also that wildcard are record type agnostic, that is if there is already another record present in the zone, like a PTR, wildcards cannot be used. In practice, this is not a show stopper, as, if this is the case, there is certainly an automated script in place to manage those PTR and the solution is to update that script to also manage the NAPTR records.

When a very large number of NAPTR records have to be generated, an alternative to wildcards is to have the DNS server dynamically generate the corresponding records on demand according to predefined rules.

The entire tree does not have to be populated. An ISP could, for example, only populate the records for tunnel brokers for the IP addresses of its customers who actually subscribed for the service.

Note:

When several services are to be discovered using this method, several NAPTR records would be created per node in the reverse path DNS tree representing an IP address, as many as the number of services to be discovered. It is actually possible to have several NAPRT records for the same service, the querying host would then decide which one(s) to use.

3. Discovering Services

When a client node wants to discover a given service, it creates a corresponding NAPRT DNS query for its IP address and send it as a regular DNS query. For example, the node 10.6.1.9 trying to discover its isatap router will send the following query:

```
query(type=NAPTR, node=9.1.6.10.in-addr.arpa.)
```

and then filter all the responses to retain those which service field is equal to "isatap". The fully qualified domain name (FQDN) of the isatap router to use will be contained in the replacement fields. Note: several NAPTR records could match, and then the node will end up with as many potential isatap routers to try. Mapping this FQDN to an IP address will required a supplementary DNS request for an A record for that FQDN.

A similar algorithm can be used by clients willing to discover their ISP tunnel broker. The NAPTR query would be the same, but this time, the client will filter the responses to retain those which service

field is equal to "TB"

4. Operational Considerations

This methods works well when finding the name of a server is enough to complete the service discovery bootstrap phase. As most of the time the DNS data is publicly readable, no sensitive data should be place within those NAPTR records.

DNS administrator concerned about not revealing to the outside world details about their internal service configuration can use two face DNS servers to only server those NAPTR records internally.

In the case where NAT and private address space are expected to be used, the proposed mechanism does not work very well.

The administrative authority in charge of the service to be discovered could pre-populate the RFC1918 [2] private address space with the relevant NAPTR records. This assumes that the client behind the NAT will use the DNS server that is provided by the entity managing the service to be discovered. This may or may not be a valid assumption depending on the situation. Also, doing this loses one of the main benefits of this proposal. Unless the client are redirected to a DNS server that is topologically close to them, it is difficult to return information that are tailored for specific customers.

Another alternative suggested to deal with NAT is to first discover the outside, global address of the NAT box, using STUN [5] for example. However, this would move the problem of tunnel end point discovery to the one of STUN server discovery, which is not really much of an improvement.

5. Defining a new DNS record type as an alternative to NATPR

Following IAB recommendations in [4], it might make sense not to overload the use of NATPR but to define a new record type, specially for the purpose of tunnel end point discovery. The format of that new record could be simplified and look like:

```
5.6.7.8.in-addr.arpa. IN TEP tep.example.com
```

or

```
5.6.7.8.in-addr.arpa. IN TEP 1.2.3.4
```

The later would have the advantage to remove one step in the resolution process by having the IPv4 address directly available. It might be stored as a 32 bit value but displayed as an IPv4 dotted decimal address.

The way those records would be stored in the DNS would be the same. However, having one record type define per type of service (TEP,

ISATAP,...) would greatly reduce the workload on the client side parsing all the information returned by the server.

6. IANA Considerations

The definition of NAPTR service fields should be standardized at IETF and recorded with IANA. A special category should be created for that. Service fields used in this memo are there only to server as examples an in no way should be used like this.

7. Security Considerations

Administrator concerned about the security of the discovery mechanism discussed here should deploy DNSsec [3]. Limiting the propagation of DNS data linked to this mechanism to "internal" customers as described in section 4 is also a good way to limit security risks. Also, as DNS data always end up leaking, one should refrain from placing sensitive information in the DNS.

8. Authors Addresses

Shu Yamamoto
KDDI R&D Labs
Kamifukuoka-Shi
Saitama 356-8502
Japan
Phone: 81 (49) 278-7894
EMail: shu@kddilabs.jp

Hidetoshi Yokota
KDDI R&D Labs
Kamifukuoka-Shi
Saitama 356-8502
Japan
Phone: 81 (49) 278-7894
EMail: yokota@kddlabs.co.jp

Carl Williams
KDDI Labs USA
Palo Alto, CA 94301
USA
Phone: +1 650 279 5903
EMail: carlw@kddilabs.com

Alain Durand
no affiliation
Mail: alain@tycool.net

8. Normative References

[1] RFC3403. Dynamic Delegation Discovery System (DDDS) Part Three: The Domain Name System (DNS) Database. M. Mealling. October 2002.

8. Informative References

[2] RFC1918. Address Allocation for Private Internets. Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot, E. Lear. February 1996.

[3] RFC2535. Domain Name System Security Extensions. D. Eastlake 3rd. March 1999.

[4] draft-iab-dns-choices-00.txt, P. Faltstrom, R. Austein, Design Choices When Expanding DNS, October 2004.

[5] RFC3489. STUN - Simple Traversal of User Datagram Protocol (UDP) Through Network Address Translators (NATs). J. Rosenberg, J. Weinberger, C. Huitema, R. Mahy. March 2003.

10. Copyright Statement