IPng Working Group
Internet Draft

Matt Crawford
Fermilab
November 17, 2000

Discovery of Resource Records Designating IPv6 Address prefixes
<draft-ietf-ipngwg-prefix-rr-disc-00.txt>

Status of this Memo

Abstract

The A6 resource record type [A6] was introduced to store IPv6 addresses in a manner which facilitates prefix changes and assignment of addresses from multiple prefixes. In order to allow use of dynamic DNS updates while still respecting whatever prefix hierarchy may be in use in a site's "reverse" DNS zone, a method is needed for discovering the name(s) of the A6 record(s) which specify an address prefix.

This memo specifies such a method of prefix name discovery.

1. Introduction

The A6 resource record type [A6] was introduced to store IPv6 addresses in a manner which facilitates prefix changes and assignment of addresses from multiple prefixes. In order to allow use of dynamic DNS updates while still respecting whatever prefix hierarchy may be in use in a site's "reverse" DNS zone, a method is needed for discovering the name(s) of the A6 record(s) which specify an address prefix.

This memo specifies such a method. No new protocols or DNS record types are involved -- only a convention for storing the required information and a procedure for finding it.

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 [KWORD].

2. Prefix Name Storage

Recall from [A6] that address-to-name mapping information may be stored in a subzone of IP6.ARPA, or in another zone reached by a chain of one or more DNAME records. Nodenames are stored in PTR records in such a zone. Extending that custom, we specify that prefixes are to be named in PTR records in the same way. For a prefix "P" of length "L" bits there should be a PTR whose RDATA contains the owner name of an A6 record suitable for designating the prefix P/L, and this PTR record is to be stored so that it will be returned by a DNS query for the QNAME \[P/L].IP6.ARPA (possibly after resolving intervening DNAMEs [DNAME]).

Since the purpose of prefix name discovery is to facilitate dynamic registration by hosts of their IPv6 addresses in DNS, only the names of "longest" prefixes need to be discoverable. Accordingly, this example will show just a prefix which is not subnetted further.

Building on the example from [A6], section 5.2.3, the addition of the required PTR record is shown below.

```
$ORIGIN X.EXAMPLE.
N            A6 64 ::1234:5678:9ABC:DEF0 SUBNET-1.IP6
SUBNET-1.IP6 A6 48 0:0:0:1::  IP6
             PTR   SUBNET-1.IP6        ; added record
IP6          A6 48 0::0       SUBSCRIBER-X.IP6.A.NET.
IP6          A6 48 0::0       SUBSCRIBER-X.IP6.B.NET.
$ORIGIN IP6
\[x0001/16]                   DNAME   SUBNET-1
\[x123456789ABCDEF0].SUBNET-1 PTR     N.X.EXAMPLE.
```

Notice that the owner and RDATA are the same. This is a consequence of a somewhat arbitrary choice. The new record could equally well have been

```
\[x0001/16].IP6.X.EXAMPLE.  PTR   SUBNET-1.IP6.X.EXAMPLE.
```

It cannot be determined by inspecting an A6 DNS record whether that

record is meant to specify all the trailing bits of a 128-bit IPv6 address or merely a prefix. Inclusion of the trailing bits does not preclude its being pointed to as a prefix by some other A6 record. Nevertheless, a human or automated zone maintainer will generally know the intended purpose of each A6 record and which one should be named in a PTR for prefix name discovery.

3. Prefix Name Discovery

If a process wishing to do prefix name discovery has the prefix itself available (as opposed to a full address of which an unknown initial portion is the prefix), the prefix can be looked up directly. Otherwise, two heuristics are available.

First, it is possible that looking up a PTR record based on the full IPv6 address, as would be done for ordinary address-to-name mapping, will yield a PTR record containing a hostname. That hostname will then be the owner of an A6 record. If its prefix length field is non-zero, its prefix name field will contain the desired name.

Otherwise, looking up a PTR record will fail, returning an authoritative name error no no data of the requested type. There will be a set of DNAME records in the answer section of the reply. The last of these DNAMEs will indicate where to start looking for the required PTR record. First its target should be tried, then its owner. An especially persistent implementation can then prepend one bit at a time from the portion of the IPv6 address not mapped by the DNAME records to the target name, looking for a PTR record which was not at a DNAME cut point of its own. An authoritative name error is a stopping signal for this search.

4. Security Considerations

No security concerns are raised by this specification beyond those which apply to all uses of the DNS.

5. References

[A6] Crawford, M. and C. Huitema, "DNS Extensions to Support IPv6 Address Aggregation and Renumbering", RFC 2874, July 2000.

[KWORD] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," RFC 2119.

[DNAME] Crawford, M., "Non-Terminal DNS Name Redirection", RFC 2672, August 1999.

6. Authors' Addresses

Matt Crawford
Fermilab
MS 368
PO Box 500
Batavia, IL 60510
USA

+1 630 840-3461
crawdad@fnal.gov