Internet Engineering Task Force   Christian Huitema
INTERNET DRAFT   Susan Thomson
                                                              Bellcore
<draft-ietf-ipngwg-aaaa-03.txt>   February 6, 1998

DNS Extensions to support IP version 6

Status of this Memo

Abstract

This document defines the changes that need to be made to the Domain Name System to support hosts running IP version 6 (IPv6). The changes include a new resource record type to store an IPv6 address, a new domain to support lookups based on an IPv6 address, and updated definitions of existing query types that return Internet addresses as part of additional section processing.

1.Introduction

Current support for the storage of Internet addresses in the Domain Name System (DNS)[1,2] cannot easily be extended to support IPv6 addresses[3] since applications assume that address queries return 32-bit IPv4 addresses only.

To support the storage of IPv6 addresses we define the following extensions:

- o A new resource record type is defined to map a domain name to an IPv6 address.
- o A new domain is defined to support lookups based on address.
- o Existing queries that perform additional section processing to locate IPv4 addresses are redefined to perform additional section processing on both IPv4 and IPv6 addresses.

The changes are designed to be compatible with existing software. The existing support for IPv4 addresses is retained. Transition issues related to the co-existence of both IPv4 and IPv6 addresses in DNS are discussed in [4].

This memo proposes an incompatible extension to the specification in RFC 1886, and a departure from current implementation practices. The changes are designed to facilitate network renumbering.

2. NEW RESOURCE RECORD DEFINITION AND DOMAIN

A new record type is defined to store a system's IPv6 address, or addresses. The new record contains the least significant bits of the host's IPv6 address. When the number of significant bits is lower than 128, the record also contains the domain name of another IPv6 system, which typically describes a complete link, or a complete site. The most significant bits will be copied from the IPv6 address of that system. If that system has several IPv6 addresses, the low bits of the host address will be combined with each prefix of the several addresses, resulting in as many IPv6 addresses for the host.

A system may need several records if it is connected to several domains, as would be the case, for example, of a site connected to several providers, or of a host connected to different links.

2.1 AAAA record type

The AAAA resource record type is a new record specific to the Internet class that stores the lower bits of a single IPv6 address and the name of a domain where to fetch the higher bits.

The value of the type is 28 (decimal).

Note that we decide here to reuse the name and code specified in RFC 1886. The record format has been changed, but it has been Changed in a compatible way. Essentially, we have added to The old format an optional "tail". Updated systems will be capable of reading the old records. Old systems, however, will only be capable of using the new records if they decide to use the first 128 bits and ignore the remainder. This will be discussed in the "transition" section.

2.2 AAAA data format

```
+------------------+--------------+----------------------+
|   Ipv6 address   | Pre. length  | Domain name of prefix |
|     (128 bits)   |  (1 octet)   | (variable, 0..255)    |
+------------------+--------------+----------------------+
```

The data portion of the AAAA record contains three fields:

- o A 128 bit IPv6 address, encoded in network byte order (high-order byte first).

- o a prefix length, encoded as one single octet, with value in the range [0..128].

- o the domain name of the prefix, encoded as a domain name, possibly compressed as specified in [3]. (The compression of the domain name saves space, but may cause problems if servers that don't understand the AAAA type cache this record.)

The domain name component shall not be encoded if the length of the prefix is zero.

2.3 AAAA query

An AAAA query for a specified domain name in the Internet class returns all associated AAAA resource records in the answer section of a response.

A type AAAA query does perform additional section processing, by returning the AAAA records associated to the domain names mentioned in the domain's AAAA records.

2.4 Textual format of AAAA records

The textual representation of the data portion of the AAAA resource record used in a master database file is composed of three fields separated by white spaces:

- o the textual representation of the host's IPv6 address as defined in [3],

- o a prefix length, represented as a decimal number,

- o a domain name.

The domain name may be absent if the prefix length is zero.

2.5 IP6.INT Domain

A special domain is defined to look up a record given an address. The intent of this domain is to provide a way of mapping an IPv6 address to a host name, although it may be used for other purposes as well. The domain is rooted at IP6.INT.

An IPv6 address is represented as a name in the IP6.INT domain by a sequence of nibbles separated by dots with the suffix ".IP6.INT". The sequence of nibbles is encoded in reverse order, i.e. the low-order nibble is encoded first, followed by the next low-order nibble and so on. Each nibble is represented by a hexadecimal digit. For example, the inverse lookup domain name corresponding to the address

    4321:0:1:2:3:4:567:89ab

would be

b.a.9.8.7.6.5.0.4.0.0.0.3.0.0.0.2.0.0.0.1.0.0.0.0.0.0.0.1.2.3.4.IP6.INT.

4. MODIFICATIONS TO EXISTING QUERY TYPES

All existing query types that perform type A additional section processing, i.e. name server (NS), mail exchange (MX) and mailbox (MB) query types, must be redefined to perform both type A and type AAAA additional section processing. These new definitions mean that a name server may add any relevant IPv4 addresses and any relevant IPv6 addresses available locally to the additional section of a response when processing any one of the above queries.

5. TRANSITION FROM RFC-1886 TO NEW FORMAT

The new specification of the AAAA record allows domain managers to only specify the lower bits of the Ipv6 address in the AAAA record. The upper, or most significant, bits, will be derived from the AAAA record of the prefix. This new format is designed to better support network renumbering and network multi-homing, while preserving some degree of compatibility with the existing records.

5.1 Typical usage of the new format.

Let's take the example of a site, X, that is multi-homed to two "intermediate" providers A and B. The provider A is itself multi-homed to two transit providers, C and D. The provider B gets its transit service from a single provider, E. Using the Ipv6 "aggregatable addresses" format, C, D and E are represented by the top level aggregates (TLA) '2CCC', '2DDD', and '2EEE', respectively. C assigns to A the "next level aggregate" '00CA', D assigns '00DA' to A, and E assigns '00EB' to B. A assigns to X the subscriber identification '001A', and B assigns the subscriber identification '034B'. As a result, the site X inherits three address prefixes:

- o 2CCC:00CA:001A/48 from A, for routes through C,
- o 2DDD:00DA:001A/48 from A, for routes through D,
- o 2EEE:00EB:034B/48 from B, for routes through E,

Lets suppose that S is a station in the site X, that it is connected to the link number 1 in this site, and that it uses the interface identifier (UID) '1234:5678:90AB:CDEF'. In our configuration, this station will be configured with three addresses:

- o 2CCC:00CA:001A:0001:1234:5678:90AB:CDEF
- o 2DDD:00DA:001A:0001:1234:5678:90AB:CDEF
- o 2EEE:00EB:034B:0001:1234:5678:90AB:CDEF

In the DNS, we will assume that the site X is represented by the domain name XXX.COM, while A, B, C, D and E are represented by A.NET, B.NET, C.NET, D.NET and E.NET. In each of these domains, we create a subdomain "IP6" that will hold the corresponding prefixes. The station S is identified by the domain name S.XXX.COM. Using the new format, we have to enter the following records in the DNS:

```
S.XXX.COM   AAAA ::1:1234:5678:90AB:CDEF 48 IP6.XXX.COM

IP6.XXX.COM AAAA 0:0:001A:: 28 IP6.A.NET
IP6.XXX.COM AAAA 0:0:034B:: 28 IP6.B.NET

IP6.A.NET   AAAA 0:00CA:: 16 IP6.C.NET
IP6.A.NET   AAAA 0:00DA:: 16 IP6.D.NET
```

```
IP6.B.NET   AAAA 0:00EB:: 16 IP6.E.NET

IP6.C.NET   AAAA 2CCC::

IP6.D.NET   AAAA 2DDD::

IP6.E.NET   AAAA 2EEE::
```

An immediate observation is that we will only manage one record per station, or more precisely per interface, instead of one per prefix per station. The new format allows us to easily support renumbering.

5.2 Support for network renumbering

Network renumbering occurs when a site changes providers, when a provider switches to a new provider, or when either the site or one of its providers decides to "multihome". The new format enables domain managers to manage these events by updating a single DNS record.

Suppose, in our example, that the site X decides to replace its Providers, A and B, by a direct connection to the transit network C. A single DNS update would do the trick, replacing the two AAAA records in "IP6.XXX.COM" by:

```
IP6.XXX.COM AAAA 0:0123:4567: 16 IP6.C.NET.
```

There will be no need to modify the individual records of the site's stations. As a consequence, the TTL of the station's record can be set to a large value, and the switching can be prepare by simply reducing the TTL value of the site's prefix record.

Note that in most cases the switching will be organized in two phases. The connection to the new provider will be installed, a new site AAAA record will be added for that new connection. After a transition period, the site will be disconnected from its previous providers, and the old records will be removed from the DNS.

Similarly, if the network provider B decide to switch transit provider, say from E to D, it will only have to update its own AAAA network records.

Obviously, the DNS updates will have to be synchronized with the address configuration and router renumbering procedures. This should be specified in a separate memo.

5.3 Transition strategies

The new AAAA format is an extension of the format specified in RFC 1886. Systems have already be deployed that implement RFC 1886. These old systems will not be able to understand the new format, while updated systems will still be capable of understanding the old records. This suggest a two-phase transition strategy:

1) develop resolvers that understand the new record format, but ban actual usage of the new format in the DNS, except for test purposes.

2) when the new resolvers have been deployed, start usage of the indirection capabilities provided by the new format.

5.4 Transition of the inverse tree

A characteristic of the new format is that a given prefix information is stored in exactly one place. The current "IP6.INT" domain does not share this desirable characteristic: inverse trees have to be replicated for each prefix. Moreover, it can be argue that the inverse name format is quite ugly. The IPNG working group examined several proposals that tried to solve these problems, but could not settle on any of them. In fact, it is clear that the problem is difficult, and that "better" solution require a change in the DNS specifications.

The DNS working group is examining such changes. Some proposals will allow "virtual links" within the naming tree at a zone level -- such links are currently only allowed for individual entries, using CNAME records. Other proposal would allow "binary" labels. None of these proposals is yet final.

When these new capabilities are defined, we expect to revisit the structure of the inverse tree.

6. SECURITY CONSIDERATIONS

The AAAA and PTR records can be secured by using the DNS security procedures. The signature of the AAAA record only proves that the record is genuine, i.e. has been inserted in the DNS by the manager of the specified domain. The signature of the NS and SOA records in the inverse tree can be used to check the validity of the address delegation.
Because the address will be composed by combining several records, the security level will be determined by the weakest of these records. Managers should take this in consideration when deciding to use references to prefixes.

7. ACKNOWLEDGEMENTS

Many of the ideas here were developed during a discussion between the authors, Robert Elz, Olafur Gudmundsson, Jim Bound, Bill Manning, Bob Fink, Mike O'Dell, Matt Crawford, Ken Powell, Bob Hinden and Steve Deering.

8. REFERENCES

[1] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, RFC 1034, USC/Information Sciences Institute, November 1987.

[2] Mockapetris, P., "Domain Names - Implementation and Specification", STD 13, RFC 1035, USC/Information Sciences Institute, November 1987.

[3] Hinden, R., and S. Deering, Editors, "IP Version 6 Addressing Architecture", RFC 1884, Ipsilon Networks, Xerox PARC, December 1995.

[4] Gilligan, R., and E. Nordmark, "Transition Mechanisms for IPv6 Hosts and Routers", Work in Progress.

[5] Huitema C., and S. Thomson, "DNS Extensions to support IP version 6." RFC 1886.

Authors' Addresses

Susan Thomson
Bellcore
MCC 1C259B
445 South Street
Morristown, NJ 07960
U.S.A.

Phone: +1 201-829-4514
EMail: set@bellcore.com

Christian Huitema
Bellcore
MCC 1J236B
445 South Street
Morristown, NJ 07960
U.S.A.

Phone: +1 201-829-4266
EMail: huitema@bellcore.com