Internet Draft
draft-ietf-dnsop-v6-name-space-fragmentation-01.txt
March 2002
Expires in six months

Johan Ihren
Autonomica AB

IPv4-to-IPv6 migration and DNS namespace fragmentation

Status of this Memo

Abstract

This memo documents some problems forseen in transitioning from a IPv4-only DNS hierarchy via a long period of mixture to an IPv6-mostly situation sometime in the future. The mixture period is expected to be very long, and hence design choices should very much take this into account, rather than just regard the transition as a relatively short period of pain.

The main problem with transition that this paper focus on is what to do about the namespace fragmentation that may result from certain DNS data only being available over one type of transport (i.e. v4 or v6) which is thereby likely unavailable to hosts that can cannot utilize that transport.

Two orthogonal issues are identified and discussed: deployment and use. The former while technically simple holds certain dangers that should be avoided. The "use" (as in performing DNS lookups) is much more complicated, and a suggested roadmap for this is presented.

1. Terminology

The key words "MUST", "SHALL", "REQUIRED", "SHOULD", "RECOMMENDED", and "MAY", when used un uppercase, in this document are to be interpreted as described in RFC 2119 [RFC2119].

The phrase "v4 name server" indicates a name server available over IPv4 transport. It does not imply anything about what DNS data is served. Likewise, "v6 name server" indicates a name server available over IPv6 transport. In general this document only discuss transport issues and does not care exactly what is transported.

2. Introduction to the problem of namespace fragmentation

With all DNS data only available over IPv4 transport everything is simple. IPv4 resolvers can use the intended mechanism of following referrals from the root and down while IPv6 resolvers have to work through a "translator", i.e. they have to use a second name server on a so-called "dual stack" host as a "forwarder" since they cannot access the DNS data directly. This is not a scalable solution.

With all DNS data only available over IPv6 transport everything would be equally simple, with the exception of old legacy IPv4 name servers having to switch to a forwarding configuration.

However, the second situation will not arise in a foreseeable time. Instead, it is expected that the transition will be from IPv4 only to a mixture of IPv4 and IPv6, with DNS data of theoretically three types of availability, depending on whether it is available only over IPv4 transport, only over IPv6 or both.

The latter is the best situation, and a major question is how to ensure that it as quickly as possible becomes the norm. However, while it is obvious that some DNS data will only be available over v4 transport for a long time it is also obvious that it is important to avoid fragmenting the namespace available to IPv4 only hosts. I.e. during transition it is not acceptable to break the namespace that we presently have available for IPv4-only hosts.

2.1. Namespace fragmentation vs. unreachability.

Something that is presently not clear is whether it is actually necessary to provide access to the "Internet namespace" as defined by what is visble on the public v4 Internet also on v6 transport.

The reason for the unclarity is that if one regards "the Internet" as the largest set of nodes that have a mutual 1-1 reachability for any pair of nodes over IP and adjust the "Internet namespace" to fit this set, then there is by definition no need to bridge or do any special tricks (since they can all reach each other anyhow).

On the other hand, if we regard "the Internet" as the set of nodes that share a namespace that we can refer to as "the Internet namespace" regardless of whether they can all reach each other or not, then we have to ensure that this namespace is accessible to every node, regardless of its available transport.

It is out of scope for this document to make a choice between the two alternatives, and therefore the rest of this document has to work from the assumption that the same namespace should, if possible, be made available to all nodes that claim to be part of the Internet.

3. Consequences of deploying a "IPv6 root name server"

If and when a root name server that is accessible over IPv6 transport is deployed it will immediately for the first time become possible to change IPv6-only name servers to a "native configuration", i.e. to a configuration where they follow referrals directly from the root (which is now accessible to them because of the v6 transport).

However, initially they will typically quite soon get a referral to a name server only available over IPv4 transport, and this will be impossible to follow, since there is no common transport available. Therefore the name it is trying to lookup will not get resolved and the result is that the v6-only name server cannot lookup the same set of domain names that its v4-only counterpart can.

This is fragmentation of the namespace.

Regardless of how this problem is handled it is important to realize that at first it will only concern the namespace as viewed from an IPv6-host. I.e. the IPv4 namespace will not (initially) be fragmented, and an important question is possibly how to keep it unfragmented.

4. A taxonomy of alternatives to avoid fragmentation.

4.1. Ignore the problem.

It is possible to ignore the fragmentation issue. Whether that is an acceptable choice or not has to be very carefully considered. Is it reasonable to allow v4 only hosts to over time lose access to parts of the Internet namespace just because they are not "IPv6-aware"?

4.2. DNS transport bridging.

By providing some sort of "DNS transport bridging", i.e. create a fallback mechanism that enables a name server with only one type of transport to reach a name server only available over the other transport via some sort of proxy service it would be possible to unify the DNS zones available on each transport into a common namespace.

The general consensus is that it is not possible to design such a bridging solution that works in both directions. However, it may be possible to design one that allows v6 clients to query v4 servers. See for instance [DNS-opreq] and [DNS-proxy] for more detailed discussions.

4.3. Policy based avoidance of fragmentation.

Today there are only a limited number of DNS zones on the public Internet that are only available over v6 transport, and they can mostly be regarded as "experimental". However, as soon as there is a root name server available over v6 transport it is reasonable to expect that it will become more common with v6-only zones over time.

Such a development would erode the Internet namespace as viewed from an v4-only client. There are obviously strong reasons to find a mechanism to avoid this happening.

4.3.1. Requirement of zone reachability over IPv4 transport.

To ensure that all zones remain available over IPv4 transport one method would be to require that nameservers authoritative for a zone as part of the zone validation process ensure that there are IPv4 address records available for the name servers of any child delegations within the zone).

I.e. the future policy could be:

> "Every delegation point delegated to nameservers available over v6 transport should have the same availability requirements for servers over both v4 and v6 transport as v4 only zones have over v4 transport.

I.e. if the parent requires "multiple nameservers" for a child, then the requirement becomes "multiple nameservers available over v4 transport plus multiple nameservers available over v6 transport"

I.e. for given the domain EXAMPLE.COM with the following data

```
$ORIGIN example.com.
child.example.com.          IN      NS      ns.example.com.
child.example.com.          IN      NS      dns.autonomica.se.
ns.example.com.             IN      A       1.2.3.4
```

the delegation of CHILD.EXAMPLE.COM is to the two name servers "ns.example.com" and "dns.autonomica.se". The first name server, "ns.example.com", obviously has an IPv4 address (as shown by the "glue" record on the last line).

However, "ns.example.com" may have additional addresses assiciated with it. Also there is no way for the server loading the zone to know the address(es) of "dns.autonomica.se". Therefore, to find out all the publicly available addresses they have to be queried for.

To ensure this the authoritative server will have to lookup the address records of the name servers that are part of any "delegation" points in the zone. However, this operation is very costly for large, delegation-dense zones and therefore it is likely that compromises a la

* only validate on the master (this is likely always good practice)

* validate as an offline process (i.e. not part of the zone loading)

* only validate at time of delegation

* never validate

Clearly, as validation is relaxed the amount of errors will increase, so the sum of pain as usual remains mostly constant.

4.3.2. Zone validation for non-recursive servers.

Non-recursive authoritative servers are name servers that run without ever asking questions. A change in the zone validation requirements that force them to query for the addresses of name servers that are part of delegations in the zone change this, since they now have to query for these addresses.

However, the main reason that it is important to be able to run without asking questions is to avoid "caching" possibly bogus answers. This need can be managed by requiring that a non recursive name server throw away the looked up address information after having used it for validation of the delegations in the zone.

4.3.3. Future requirement of zone reachability over IPv6 transport.

The immediate need for clarified policies for delegation is to ensure that IPv4 namespace does not start to fragment. Over time, however, it is reasonable to expect that it may become important to add a similar requirement to IPv6 namespace.

I.e. an even more refined policy possible at some point in the future would be:

> "Every delegation point should have at least one name server for the child zone reachable over IPv4 transport (i.e. should have an A record) and at least one name server reachable over IPv6 transport (i.e. should have e.g. an AAAA record)".

4.3.4. Implementation issues for new zone validation requirements.

Exactly what action should be taken when a zone does not validate is not immediately clear. Immediate alternatives include:

a) fail the entire parent zone (the extreme case, not suggested)

b) load the zone but remove the delegation that failed validation (also drastic, and not suggested)

c) load the entire zone but issue a warning message about the

delegation that failed validation (more reasonable)

Implementations should make it configurable what action to take. In the case of registries that have a business realtion to the child zone it is also in principle possible to work on the deployment of child zones over v6 transport by cost diffentiation for the customer.

5. Overview of suggested transition method.

By following the steps outlined below it will be possible to transition without outages or lack of service. The assumption is that the site has only v4 name servers or possibly v4 name servers plus v6 name server in a forwarding configuration. All DNS data is on the v4 name servers.

1) Do not change the method of resolution on any (recursive) name server. I.e. v4 servers go to the root and follow referrals while v6 servers go to their translator/forwarder which lookup the name and return the end result.

2) Start serving authoritative DNS data on v6 transport by providing name servers with v6 transport serving the zones. Add v6 address information to to the zones and as glue at the parent zone. Note that it is of crucial importance that the zone should have the same contents regardless of whether it is the v4 version or the v6 version. Anything else will lead to confusion.

4) Wait for the announcement of the DNS root zone being available from a v6 name server.

5) Ensure that the entire path from the root down to the domain in question is reachable over both IPv4 and IPv6 transport.

When this is accomplished it it possible to begin a migration of the lookup of selected services to be available over IPv6 (i.e. typically by adding a IPv6 address record, eg. AAAA record, for a server of some sort).

6. Security Considerations

Much of the security of the Internet relies, often wrongly, but still, on the DNS. Thus, changes to the characteristics of the DNS may impact the security of Internet based services.

Although it will be avoided, there may be unintended consequences as a result of operational deployment of RR types and protocols already approved by the IETF. When or if such consequences are identified, appropriate feedback will be provided to the IETF and the operational community on the efficacy of said interactions.

7. Summary.

The namespace fragmentation problem is identified and examined at some length.

A solution based upon a change in the validation method of delegation points is suggested. This will both help keep the v4 namespace unfragmented and may also help speed up deployment of DNS hierarchy in v6 space.

9. References

[RFC1034] Domain names - concepts and facilities.
P.V. Mockapetris.

[RFC1035] Domain names - implementation and specification.
P.V. Mockapetris.

[RFC2826] IAB Technical Comment on th Unique DNS Root

[DNS-proxy] draft-durand-dns-proxy-00.txt
Alain Durand

[DNS-opreq] draft-ietf-ngtrans-dns-ops-req-02.txt
Alain Durand

A. Authors' Address

Johan Ihren
Autonomica
Bellmansgatan 30
SE-118 47 Stockholm, Sweden
johani@autonomica.se