Please publish this new draft as
draft-ihren-dnsop-v6-name-space-fragment-00.txt

Johan Ihren

Internet Draft Johan Ihren
draft-ihren-dnsop-v6-name-space-fragment-00.txt Autonomica
November 2001
Expires in six months

IPv4-to-IPv6 migration and DNS name space 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 name space 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 roadmap for this is presented.

1. Terminology

The key words "MUST", "SHALL", "REQUIRED", "SHOULD", "RECOMMENDED", and "MAY" 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.

2. Introduction to the problem of name space 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 categories 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 name space available to IPv4 only hosts. I.e. during transition it is not acceptable to break the name space that we presently have available for IPv4-only hosts.

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 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 so-called "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 looked up and the result is that a v6-only name server cannot lookup the same names that its v4-only counterpart can.

There are two available methods of addressing this problem:

a) ignore it, i.e. don't solve the problem, but put the effort into helping deployment along so that the problem will shrink over time.

b) provide some sort of "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. See for instance [DNS-opreq] and [DNS-proxy] for discussions.

Regardless of how this problem is handled it is important to realize that it only concerns the fragmented name space in IPv6. I.e. the IPv4 name space is not (yet) fragmented, and a more important question is possibly how to keep it unfragmented.

4. Policy based avoidance of name space fragmentation.

Today there are only a few 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.

This would not be a good development, since this will fragment the previously unfragmented IPv4 name space and there are strong reasons to find a mechanism to avoid it.

4.1. Requirement of IPv4 address for at least one name server.

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 would be:

> "Every delegation point should have at least one name server for the child zone reachable over IPv4 transport".

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.

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.

4.2. 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. Future requirement of IPv6 address for at least one name server.

The immediate need for clarified policies for delegation is to ensure that IPv4 name space 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 name space.

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 an AAAA record)".

4.4. 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 zone

b) load the zone but remove the delegation that failed validation

c) load the entire zone but issue a warning message about the delegation that failed validation.

A likely implementation will make it configurable what action to take.

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 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 mirroring DNS data into v6 by providing v6 name servers serving the zones. Add v6 address information to to the zones and as glue at the parent zone. Note that it is important 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 AAAA record for a server of some sort).

6. How to deploy DNS hierarchy in v6 space.

The main problem with changing the DNS data so that it will become available over both IPv6 and IPv4 transport is one of scale. There are too many name servers and too many DNS zones for any kind of forced migration to be aven remotely possible.

The way of achieving deployment is by providing domain owner with

a) a reason to deploy

b) a method to deploy

c) a way of verififying the correctness of the resulting configuration

6.1. A reason to deploy.

It is important to the migration process that zones migrate to become available over v6 transport (as well as v4 transport). But it is difficult to actually require such deployment too early in the migration process.

Over time, however, it will become more reasonable to add such a requirement. One likely method to do this will be by updating the requirements for proper zone validation as was outlined above.

6.2. How to deploy DNS data.

Assuming the owner of the DNS domain has access to both IPv4 and IPv6 address space that is globally routed. The steps to take are then

a) identify all name servers that will serve the DNS domain, with their IPv4 and/or IPv6 addresses

b) arrange for a suitable method of zone synchronization

c) announce the new set of servers to the parent zone, including possible new IPv6 glue

It is recommended that the name servers run on single stack machines, i.e. machines that are only able to utilize either IPv4 transport or IPv6 transport, but not both.

A common recommendation (mostly orthogonal to IPv6 transition issues) is that authoritative name servers only serve data, i.e. they do not act as caching resolvers. That way, since they operate in non-recursive mode, they will not have any cache, and hence will not be able to give out wrongful answers based upon errors in the cache.

Since the announced name servers are single stack, the primary master from which they fetch zone data will typically have to be dual stack or otherwise some other method of data transfer has to be arranged.

7. 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.

8. Summary.

The name space 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 name space 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