Internet Engineering Task Force
INTERNET-DRAFT
September 6, 2001
Expires March 7, 2002

Alain Durand
SUN Microsystem
Johan Ihren
Autonomica AB

NGtrans IPv6 DNS operational requirements and roadmap

draft-ietf-ngtrans-dns-ops-req-02.txt

Status of this memo

Abstract

This document describes IPv6 DNS operational requirements and deployment roadmap. It is the result of discussion from members of the IPv6, NGtrans, DNSops and DNSext working groups. The DNS is looked as a critical part of the Internet infrastructure and is used for much more purposes than name to address resolution. Thus a smooth operation of the DNS is critical in the IPv6 transition.

Discussion of this memo should happen in the NGtrans mailling list.

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 RFC 2119 [RFC2119].

1. IPv6 transition: Dual-stack vs IPv6 only nodes vs IPv6-mostly nodes

1.1. Incremental deployment on existing network.

This needs to be done without disturbing IPv4 service. This strategy relies heavily on dual-stack nodes and tunnels. It is foreseen that this scenario is likely to happen in corporate networks.

1.2. Large scale deployment of new infrastructure

This scenario envision large to very large networks where public IPv4 address space is not available and private address is not practical. Nodes in this scenario will very likely be IPv6-only or IPv6-mostly (getting an IPv4 address only on demand). This scenario is likely to happen in the wireless/3G world. Note that those networks will still need to communicate with the rest of the Internet.

Given the two above scenarios, the requirements discussed in this memo are not targeted at transitioning the DNS from IPv4 only to IPv6 only, but more at the transition of IPv4 only to a mixed environment, where some systems will be IPv4 only, some will be IPv6 only and others will be dual-stacked.

2. Issues with the standard resolution process

2.1 Following the referral chain

The caching resolver that tries to lookup an name starts out at the root, and follows referrals until it is referred to a nameserver that is authoritative for the name. If somewhere down the chain of referrals it is referred to a nameserver that is only accessible over a type of transport that is unavailable, a traditional nameserver is unable to finish the task.

When the Internet moves from IPv4 to a mixture of IPv4 and IPv6 it is only a matter of time until this starts to happen and the complete DNS hierarchy starts to fragment into a graph where authoritative nameservers for certain nodes are only accessible over a certain transport. What is feared is that a node using only a particular version of IP, querying information about another node using the same version of IP can not do it because, somewhere in the chain of servers accessed during the resolution process, one or more of them will only be accessible with the other version of IP.

2.2 Examples of problems for an IPv6 only resolver

This problem shows for IPv6 only resolver trying to fetch data from a zone that is served by IPv6 servers when somewhere in the referral chain, the list of name servers pointed at does not contain any IPv6 reachable server.

Hints for the root:

```
X.ROOT-SERVERS.NET IN A 100.100.100.101
X.ROOT-SERVERS.NET IN AAAA 3ffe:ffff:100:100::1
```

In the root zone:

```
org. IN NS dot-org.X.ROOT-SERVERS.NET
dot-org.X.ROOT-SERVERS.NET IN A 100.100.100.102
```

In the .org zone:

```
foobar.org. IN NS ns.foobar.org
ns.foobar.org IN A 200.200.200.201
ns.foobar.org IN AAAA 3ffe:ffff:200:200::201
```

In the foobar.org zone:

```
www.foorbar.org IN AAAA 3ffe:ffff:200:200::202
```

Although the zone foobar.org and the root are served by an IPv6 server, an IPv6 only resolver can not resolve www.foobar.org because there is no IPv6 server for the parent zone .org.

2.3 Examples of problems for an IPv4 only resolver

Another instance of the problem shows for an IPv4 only MTA trying to send mail to someone in an IPv6 only domain which has made provision to have an IPv4 reachable MX.

In the .org zone:

```
foobar.org. IN NS ns.foobar.org
ns.foobar.org IN AAAA 3ffe:ffff:200:200::201

3rd_party_dualstack_mail.org. IN NS ns.3rd_party_dualstack.org.
ns.3rd_party_dualstack.org. IN A 100.100.100.103
```

in the foobar.org zone:

```
foobar.org IN MX 10 mail6.foobar.org.
foobar.org IN MX 20 mail4.3rd_party_dualstack.org.
mail6.foobar.org. IN AAAA 3ffe:ffff:200:200::202
```

in the 3rd_party_dualstack_mail.org zone:

```
mail4.3rd_party_dualstack.org. IN A 100.100.100.104
```

An IPv4 only host cannot get the information about the IPv4 MX relay mail4.3rd_party_dualstack_mail.org because the foobar.org zone is not served by an IPv4 DNS server.

3. IPv4 constraints

Due to the very large IPv4 deployment phase, any solution that will require any change either on the resolver binary or resolver configuration is out of scope.

4. IPv6 constraints

When IPv6-only and IPv6-mostly devices send DNS queries, per configuration/implementation they have little choice but to use IPv6 transport. For example, it may be too expensive, too slow or just not possible to get an IPv4 address for just that purpose. With growing IPv6-only areas it will become increasingly cumbersome for such installations to maintain their DNS data only on remote systems that are not available over local transport.

5. Requirements

The DNS is looked as a critical part of the Internet infrastructure and is used for much more purposes than name to address resolution. Thus, failures such as described earlier are to be prevented, and some kind of bridging between the two world is necessary.

However, it should be noted that, even though bridging has to work both ways, it is not strictly necessary to use the same technique in each direction. That is, it is perfectly acceptable to build two different mechanisms, one to enable IPv4 only hosts to query IPv6 only DNS servers and one for IPv6 only hosts to query IPv4 only DNS servers.

However, it is not clear if more than one bridging systems in a particular direction is a good thing or not.

5.1 IPv4 requirements

Also, it may be the case that no bridging solution is possible to bridge an unmodified IPv4 resolver to an IPv6 only name server. In that case, there will be a requirement that any zone be served by at least one IPv4 DNS server.

5.2 IPv6 requirements

The bridging systems that enable an IPv6 resolver to query data on an IPv4 server will have to be in place for a long time, will be a key part of the IPv6 transition and will heavily be used.

Thus, the bridging systems MUST have good scaling properties.

The IPv6 resolver MUST have a way to discover the bridging systems. This discovery mechanism MUST also have good scaling properties.

The bridging systems MUST be able to bridge any zones and, in particular in the absence of IPv6 name servers for the root, the bridging systems MUST be able to bridge the root.

6. Roadmap for DNS service in a mixed environment IPv4/IPv6

6.1 Bridging system

A robust, scalable bridging system MUST be in place prior to large scale deployment IPv6 DNS deployment.

6.2 Root name service accessible via IPv6

The first DNS query a caching resolver will send is directed to a root name server. This, if the configuration of the bridging system is derived automatically from the DNS itself, there is a strong requirement to make root name service available over IPv6 transport. If the configuration is derived any other way or is done manually, there is a possibility to operate the system without an IPv6 accessible root in certain cases. However, as this document does not want to preclude any particular implementation of the bridging systems at this point, it is highly recommended that some IPv6 enable root name server be in place as early as possible. It is an important step to show that IPv6 DNS deployment is possible.

6.3 TLDs servers accessible via IPv6

Having the capability to query a root name server using IPv6 is just the first step. The next one is to query a TLD for a NS record pointing to a domain name. Again, although not strictly necessary from a technical perspective, it is important to make sure that some TLD servers are accessible from the beginning via IPv6 so at least some label strings are resolvable with IPv6 transport without resorting to any fall-back mechanism.

Also note that great care should be taken when adding IPv6 glue in the TLD delegation by the root.

6.4 IPv6 glue at TLD registries.

Whenever glue is needed, it is necessary for domains delegated from a TLD to be able to specify an IPv6 name server address to the TLD registry. This is not so much a technical issue as it is a question of management and procedures.

6.5 Reverse path DNS servers

Reverse DNS queries should also be supported in IPv6, for the same reasons as direct queries. Today's resolvers do reverse nibbles queries under the ip6.int tree. [RFC3152] has deprecated ip6.int, thus reverse DNS queries MUST be moved to ip6.arpa. So, as in 6.2 and 6.3, although not strictly speaking a technical requirement, it is important to have at least one server for ip6.arpa accessible via IPv6.

7. Security considerations

Any bridging system, acting as open relay, could be misused to create denial of service attacks on external DNS servers. Some provision should be made in the design of those relay to deal with this issue.

8 Editor address

Alain Durand
SUN Microsystems, Inc
901 San Antonio Road
MPK17-202
Palo Alto, CA 94303-4900
USA
Mail: Alain.Durand@sun.com

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

9. References

[RFC2119] Bradner, S.,
"Key words for use in RFCs to Indicate Requirement Levels",

BCP 14, RFC 2119, March 199

[RFC3152] Bush, R.,
"Delegation of IP6.ARPA",
RFC 3152, August 2001