Internet Engineering Task Force
INTERNET-DRAFT
March 1, 2002
Expires Sep. 2, 2002

Alain Durand
SUN Microsystems
Johan Ihren
Autonomica AB

# NGtrans IPv6 DNS operational requirements and roadmap

draft-ietf-ngtrans-dns-ops-req-04.txt

## Status of this memo



## Abstract

This document describes IPv6 DNS operational requirements and deployment roadmap steps. It is the result of discussion from members of the IPv6, NGtrans, DNSop 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 mailing 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. DNS issues in a mixed IPv4/IPv6 environment

IPv4 and IPv6 are two versions of the same original concept, but they are not "binary compatible". That is, a datagram send by one version of IP cannot be received by the other. Several things can go wrong when operating DNS in a mixed environment IPv4 and IPv6.

## 1.1 Following the referral chain

The caching resolver that tries to lookup a 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.

## 1.2 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.

### 1.3 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.

## 2. Fundamental requirements

### 2.1 Uniqueness of the DNS root

[RFC2826] requires the existence of a globally unique public name space derived from a unique root. This root is valid for both IPv4 and IPv6.

---

Requirement 1:

The public DNS has a unique root valid for IPv4 & IPv6.

---

### 2.2 DNS should be an IP version agnostic application

Although DNS is regarded as a key component of the Internet infrastructure, it is an application at layer 7 of OSI model and should be independent from particular protocol choice at the network

layer. Some record type, like CNAME or MX are clearly IP version agnostic. Even data like A, AAAA or PTR records contained in the DNS may be relevant to particular applications requesting then regardless of the IP version used during the queries. Also, [RFC2826] states, "A DNS name can be passed from one party to another without altering the semantic intent of the name." So, this is not because a particular host can only communicate with a certain version of IP that it should be prevented to query information regarding the over version of IP. Another way of saying this is to say that the DNS data are independent of the particular version of IP used to carry them.

---

Requirement 2:

Any node SHOULD be able to query any data from the DNS regardless of the IP versions used for the transport of the queries and responses issued by the various parties in place.

---

There is ongoing discussion to know if queries should get the same answer regardless of the IP version used during the process. This, of course, would not apply to records in the additional sections. See [NAT-PT-ISSUES] and [INTERACTION] section 5.2.1.

## 2.3 Transition is a long journey

It is usually believed that transition can happen simultaneously following two main scenarios.

- 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.

- 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). 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.

It is generally admitted that, the burden of transition should be placed on the new IPv6 systems and their local IPv6 infrastructure. Ad-hoc administrative practices such as a local dual stack resolver or locally Local dual stack resolver or locally administered NAT-PT translator [RFC2766] could enable networks where some dual stacks node are available to query IPv4 only DNS servers. (Note that NAT-PT would have to be modified for that purpose as it translate AAAA queries into A queries, see [NAT-PT-ISSUES].) Administrative practices requiring any zone served by IPv6 only servers to be also served by IPv4 servers would enable IPv4 only resolvers to perform DNS queries for those zones.

However, the requirements described here are looking at solving the long term problem. Although dual stack networks will be common in the early days of transition, IPv6 only networks would eventually be a reality and solutions describe above would not be practical.

---

Requirement 3:

A global approach IS REQUIRED to enable networks operating with only one version of IP to query zones of the public DNS that are only served by systems operating only with the over version of IP.

---

The choice and the details of this approach are beyond the scope of this document and should be discussed in the DNSop and DNSext working groups. It can be the case that communication can be achieved via a set of agreed administrative procedures. It can be also the case that a general purpose, ubiquitous translator will be the right thing or that a DNS specific solution must be developed. If new pieces of protocols are needed in the resolvers, due to the extraordinary amount of time it takes to define then, implement them, test them, ship them into existing products and get them deployed, works should start as soon as possible.

## 3. Global approach requirements

Even though communication 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 havew two different approaches, one to enable IPv4 only hosts to query IPv6 only DNS servers and one for IPv6 only hosts to query IPv4 only DNS servers. It is also possible that part of the approach consists of a set of administrative procedures required to operate DNS zones.

### 3.1 IPv4 contraints

Due to the very large IPv4 deployment phase, any solution that will require any change either on binaries or configurations on every IPv4 resolvers is out of scope.

### 3.2 Scaling

The aproach that enable a resolver to query data from a server which use a different IP version will have to be in place for a long time. It will be a key part of the general IPv6 transition and will heavily be used.

---
Requirement 4:

Whatever approach will be chosen SHOULD have good scaling properties.
---

### 3.3 Scaling even more

Auto configuration is the tendency for end systems. If the global approach involves resolvers connecting to intermediary systems, Resolver SHOULD have a way to discover those components. This discovery mechanism SHOULD also have good scaling properties.

---
Requirement 5:

If the agreed approach include discovery of intermediary components, the discovery mechanism SHOULD have good scaling properties.
---

### 3.3 Scope

The agreed solution SHOULD be able to bridge any zones. In particular, until there is an IPv6 root name server, the communication systems SHOULD be able to bridge the IPv4 root.

---
Requirement 6:
All zones (even the root) SHOULD be reachable.
---

### 3.4 Security matters

Being a critical piece of the Internet infrastructure, the DNS is a potential value target and thus should be protected. Great care should be used not to introduce new security issues.

---
Requirement 7:

The solution SHOULD NOT introduce new security hazards.
---

### 3.5 Bridging from IPv4

Although the details are beyond the scope of this document, it may be the case that there is no general solution to allow an unmodified IPv4 resolver to query an IPv6 only name server. In that would be the case, the IPv4 to IPv6 communication approach could consist of an operational procedure:

---
Possible operational procedure to bridge from IPv4 to IPv6:

Any zone SHOULD be served by at least one IPv4 DNS server.
---

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

### 4.1 Communication system

A communication system or a set of administrative proceduresr satisfying all the above requirements SHOULD be in place as early as possible to allow large scale IPv6 only DNS deployment.

---
Roadmap step 1:

A robust, scalable communication system and/or set of administrative procedures should be defined, agreed and put in place as soon as possible.
---

### 4.2 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 system

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.

---

Roadmap step 2:

The root SHOULD have at least one IPv6 name server.

---

### 4.3 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 the mechanims described above.

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

---

Roadmap step 3:

Each TLD zone SHOULD have at least one IPv6 name server.

---

### 4.4 IPv6 glue at TLD registries.

Whenever glue is needed, it is necessary for domains delegated under a TLD to be able to specify an IPv6 name server address to the TLD registry. This is not so much a protocol issue but a management and procedural issue.

---

Roadmap step 4:

Domains registering under TLDs SHOULD be able to specify IPv6 glue wherever they are specifying IPv4 glue today.

---

### 4.5 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, although

again not strictly speaking a technical requirement, it is important to have at least one server for ip6.arpa accessible via IPv6.

---

Roadmap step 5:

The ip6.arpa zone SHOULD have at least one IPv6 server.

---

## 5. 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.

## 6 Authors addresses


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


## 7. References


[INTERACTION] Baudot, A. and al,
"Interaction of transition mechanisms",
draft-ietf-ngtrans-interaction-00.txt, Work in progress.

[NAT-PT-ISSUES] Durand, A.,
"Issues with NAT-PT DNS ALG in RFC2766",
draft-durand-natpt-dns-alg-issues-00.txt, Work in progress.

[RFC2026] Bradner, S.,
"The Internet Standards Process -- Revision 3",
BCP 9, RFC 2026, October 1996

[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

[RFC2826] Internet Architecture Board,
"IAB Technical Comment on the Unique DNS Root",
RFC 2826, May 2000

[RFC2766] Tsirtsis, G., Srisuresh, P.,
"Network Address Translation - Protocol Translation (NAT-PT)",
RFC 2766, February 2000

## 8. Changes since -03

- remove the name "briging system" wherever possible
- add a open issue on wherever or not queries should get
  the same answer regardless of the Ip version used
  during the process.
- add refereces to [NAT_PT_ISSUES] and [INTERACTION]