

Sunset4 Working Group
Internet-Draft
Intended status: Informational
Expires: April 30, 2015

L. Song
Beijing Internet Institute
P. Vixie
Farsight Security, Inc.
D. Ma
ZDNS
October 27, 2014

Considerations on IPv6-only DNS Development
draft-song-sunset4-ipv6only-dns-00

Abstract

Deployment of IPv6-only networks are impacted by assumptions of IPv4-only or dual-stack transition scenarios. For example, these assumptions are in the operations of DNS. This memo is problem statement and hopes to eventually propose a mitigation technique.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 30, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4.e](#) of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	3
3.	Revisit to current situation	3
3.1.	DNS Referral Response Size limitation	3
3.2.	Additional section in IPv4/IPv6 Environments	4
3.3.	DNS proxy	5
4.	Mitigation approach	6
5.	Security Considerations	6
6.	IANA Considerations	6
7.	Acknowledgements	6
8.	References	6
8.1.	Normative References	6
8.2.	URIs	7
	Authors' Addresses	7

[1.](#) Introduction

It's commonly believed that the dual-stack model is the best practice for IPv6 transition in which IPv4 and IPv6 function can work in parallel without mutual interference. Based on this model, IP stacks and applications are expected to be converted into IPv6 smoothly when IPv4 address pool run out. The dual-stack approach gives IPv4/IPv6 capability on end system, network devices, DNS and application servers, but, as a side effect, brings additional problems, such as IPv4 fallback [[RFC6555](#)] or even IPv4/IPv6 competition. This issue makes the dual stack model more complicated to deploy and manage, and overall network less reliable.

To accelerate the transition to a fully connected IPv6 network, IPv6-only experiments [[RFC6586](#)] and IETF standards [[RFC6333](#)], [[RFC7040](#)] are documented. Some techniques verify IPv6 capability and support the IPv6-only deployment. In IPv6-only environments, DNS resolvers or modules are provisioned only with IPv6 address. It is mainly due to three aspects:

- 1) To save more free IPv4 addresses in deploying new DNS resolvers;
- 2) To reduce the cost and risk of management in dual stack environment;
- 3) To follow the inherent requirement in the IPv6 transition scenarios, such as DS-Lite [[RFC6333](#)];

It's worthwhile to mention that the tunnel technology provides an approach that allow IPv6-only network deployment become independent from the rest of the world which makes the IPv6-only strategy much popular. In the IPv6-only network, the ISPs only provision IPv6 address to the end system, network and DNS element via DHCPv6. However, IPv6-only resolver will face an Internet which are partly running in IPv4 only environment and partly in dual-stack, yet with IPv4-preferred paradigm. As a result, the DNS element in IPv6-only environment is suggested to be forwarding requests by relying on the upstream dual-stack DNS recursive server [section 5.5](#) [1] in [\[RFC6333\]](#). However, using the DNS proxy mechanism is a compromise in IPv6 transition context, which still has implicit limitations [\[RFC5625\]](#).

This memo revisits the behavior and implicit inertia of DNS in existing architecture which may hinder the IPv6-only DNS development.

2. Terminology

A: A resource record type used to specify an IPv4 address [\[RFC1034\]](#)

AAAA: A resource record type used to specify an IPv6 address [\[RFC3596\]](#)

EDNS0: Version 0 of Extension mechanisms for DNS [\[RFC6891\]](#)

DNSSEC: DNS Security Extensions [\[RFC4033\]](#)

MTU: Maximum Transmission Unit, the maximum size for a datagram to be forwarded on an interface without needing fragmentation [\[RFC0791\]](#), [\[RFC2460\]](#)

Additional Section: Section in DNS query/response carrying RRs which may be helpful in using the RRs in the other sections [\[RFC1034\]](#).

Note that in this memo the data in additional section is the A/AAAA information of NS server, particular for root zone.

3. Revisit to current situation

3.1. DNS Referral Response Size limitation

Due to the required minimum IP reassembly limit for IPv4, the original DNS standard [\[RFC1034\]](#)[\[RFC1035\]](#) limited the UDP message size to 512 octets. It became an historical and practical hard DNS protocol limit, even after EDNS0 [\[RFC6891\]](#) was introduced to mitigate this issue[\[draft-ietf-dnsop-respsize-15\]](#). This limit presents a problem for zones wishing to (1) add more authority servers or (2) advertise the

IPv6 addresses of newly updated dual-stack NS name servers, or (3) use DNSSEC.

In the context of this memo, the limitation may be relaxed due to the larger base MTU of IPv6 (1280 octets) which is the default for IPv6-only networks.

3.2. Additional section in IPv4/IPv6 Environments

Given there is hard limitation in the DNS referral response size, the implementations preferably decide to keep as much data as possible in the UDP responses no matter it is "critical" or "courtesy" [Appendix B.2 in \[RFC4472\]](#) . It is a typical case in priming exchange between recursive resolver and root server. When a name server resolver bootstrap, it performs the NS lookup for root zone. In the response packet from root server, the additional section is supposed to contain all the A & AAAA records of NS domain name. Ultimately, when all 13 root name servers are assigned IPv6 addresses, the priming response will increase in size to 800 bytes.

There are different strategies for root server operators to choose which RRset (A or AAAA) should be in the additional data if not all of the glue information can be included. Note that in dual-stack environment, IPv4 glue and IPv6 glue of same zone are actually competing for the room of DNS UDP packets. For example, some of DNS root servers prefer to return as many IPv4 glue records as possible. In that case only 2 out of 10 IPv6 glues are included as shown below, irrespective of IPv4 or IPv6 DNS transport.

;; ADDITIONAL SECTION:

```
a.root-servers.net. 518400 IN A 198.41.0.4
b.root-servers.net. 518400 IN A 192.228.79.201
c.root-servers.net. 518400 IN A 192.33.4.12
d.root-servers.net. 518400 IN A 199.7.91.13
e.root-servers.net. 518400 IN A 192.203.230.10
f.root-servers.net. 518400 IN A 192.5.5.241
g.root-servers.net. 518400 IN A 192.112.36.4
h.root-servers.net. 518400 IN A 128.63.2.53
i.root-servers.net. 518400 IN A 192.36.148.17
```



```
j.root-servers.net. 518400 IN A 192.58.128.30
k.root-servers.net. 518400 IN A 193.0.14.129
l.root-servers.net. 518400 IN A 199.7.83.42
m.root-servers.net. 518400 IN A 202.12.27.33
a.root-servers.net. 518400 IN AAAA 2001:503:ba3e::2:30
b.root-servers.net. 518400 IN AAAA 2001:500:84::b
```

In the context of IPv6-only deployments, these glue records are much less optimal. They are based on IPv4 or dual-stack assumptions, where IPv4 is still dominant. It may negatively impact the IPv6 services in IPv6-only deployments.

If the glue set sent in the response is correlated with the IP version of the DNS transport, then the answer, in most cases, will be more optimal. There are two reasons why it is not adopted as an optimization. One is that it breaks the model of independence of DNS transport and resource records [section 1.2 \[2\]](#) in [\[RFC4472\]](#). Another is that it will bring unpredictable risk to the performance and stability of current root server system.

[3.3.](#) DNS proxy

In IPv6-only networking, DNS proxy approach is recommended for IPv6-only DNS element. On one hand, it avoids the difficulty to perform all DNS resolution over IPv6 transport, given that still many networks on Internet are only on IPv4. On another hand, it loses the opportunity to perform a full recursive resolver function via IPv6, at least in Root and TLD level which are mostly IPv6 enabled.

In addition, as described in the beginning of [\[RFC5625\]](#), the DNS proxy function is not an optimal solution to serve the IPv6-only resolver requirement. Large packets caused by priming request or DNSSEC validation packets will be blocked due to the proxy implementation. It is suggested that: "To ensure full DNS protocol interoperability it is preferred that client stub resolvers should communicate directly with full-feature, upstream recursive resolvers wherever possible."

As more and more NS servers updated to IPv6 transport and reachable over the IPv6 Internet, the direct IPv6 resolution will be preferable in IPv6-only resolver. But regarding the long-tail feature of IPv6 adoption in NS servers, certain back-forward compatible mechanism

should be designed, which indeed make an incentive model for IPv6 adoption over IPv4 as well.

4. Mitigation approach

TBD

5. Security Considerations

TBD

6. IANA Considerations

TBD

7. Acknowledgements

TBD

8. References

8.1. Normative References

[I-D.ietf-dnsop-respsize]

Vixie, P., Kato, A., and J. Abley, "DNS Referral Response Size Issues", [draft-ietf-dnsop-respsize-15](#) (work in progress), February 2014.

[I-D.lee-dnsop-scalingroot]

Lee, X., Vixie, P., and Z. Yan, "How to scale the DNS root system?", [draft-lee-dnsop-scalingroot-00](#) (work in progress), July 2014.

[RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), September 1981.

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.

[RFC3596] Thomson, S., Huitema, C., Ksinant, V., and M. Souissi, "DNS Extensions to Support IP Version 6", [RFC 3596](#), October 2003.

- [RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", [RFC 4033](#), March 2005.
- [RFC4472] Durand, A., Ihren, J., and P. Savola, "Operational Considerations and Issues with IPv6 DNS", [RFC 4472](#), April 2006.
- [RFC5625] Bellis, R., "DNS Proxy Implementation Guidelines", [BCP 152](#), [RFC 5625](#), August 2009.
- [RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", [RFC 6333](#), August 2011.
- [RFC6555] Wing, D. and A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", [RFC 6555](#), April 2012.
- [RFC6586] Arkko, J. and A. Keranen, "Experiences from an IPv6-Only Network", [RFC 6586](#), April 2012.
- [RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, [RFC 6891](#), April 2013.
- [RFC7040] Cui, Y., Wu, J., Wu, P., Vautrin, O., and Y. Lee, "Public IPv4-over-IPv6 Access Network", [RFC 7040](#), November 2013.

8.2. URIs

- [1] <http://tools.ietf.org/html/rfc6333#section-5.5>
- [2] <http://tools.ietf.org/html/rfc4472#section-1.2>

Authors' Addresses

Linjian Song
Beijing Internet Institute
2508 Room, 25th Floor, Tower A, Time Fortune
Beijing 100028
P. R. China

Email: songlinjian@gmail.com

Paul Vixie
Farsight Security, Inc.
155 Bovet Road, #476
San Mateo, CA 94402
USA

Phone: +1 650 489 7919
Email: vixie@farsightsecurity.com

Di Ma
ZDNS
Beijing
P. R. China

Email: madi@zdns.cn

