DNSOP T. Saraj, Ed.
Internet-Draft M. Yousaf
Intended status: Standards Track Riphah Institute of Systems Engineering
Expires: November 12, 2018 A. Qayyum
Capital University of Science and Technology
May 11, 2018

IVIPTR: Resource Record for DNS
draft-tariq-dnsop-iviptr-01

Abstract

This document proposes a new DNS Resource Record IVIPTR which provides the capability to resolve the IPv4 address to IPv6 address and IPv6 address to IPv4 address. This document assumes that the reader is familiar with all the concepts and details discussed in Domain Names Concepts and Facilities [RFC1034] , Domain Names - Implementation and Specification [RFC1035]

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 https://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 November 12, 2018.

Table of Contents

1. Introduction

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 [RFC2119].

The current DNS standard does not support to resolve IPv4 address to IPv6 address and IPv6 address to IPv4 address. For example, if a user program initiate a query for AAAA resource record against an IPv4 address of a domain, the current DNS will return a negative answer normally with RCODE(3)-Non-Existent Domain. Using the current DNS standard, a user program can resolve IPv6 address for a desired IPv4 address by the process as in figure-01:

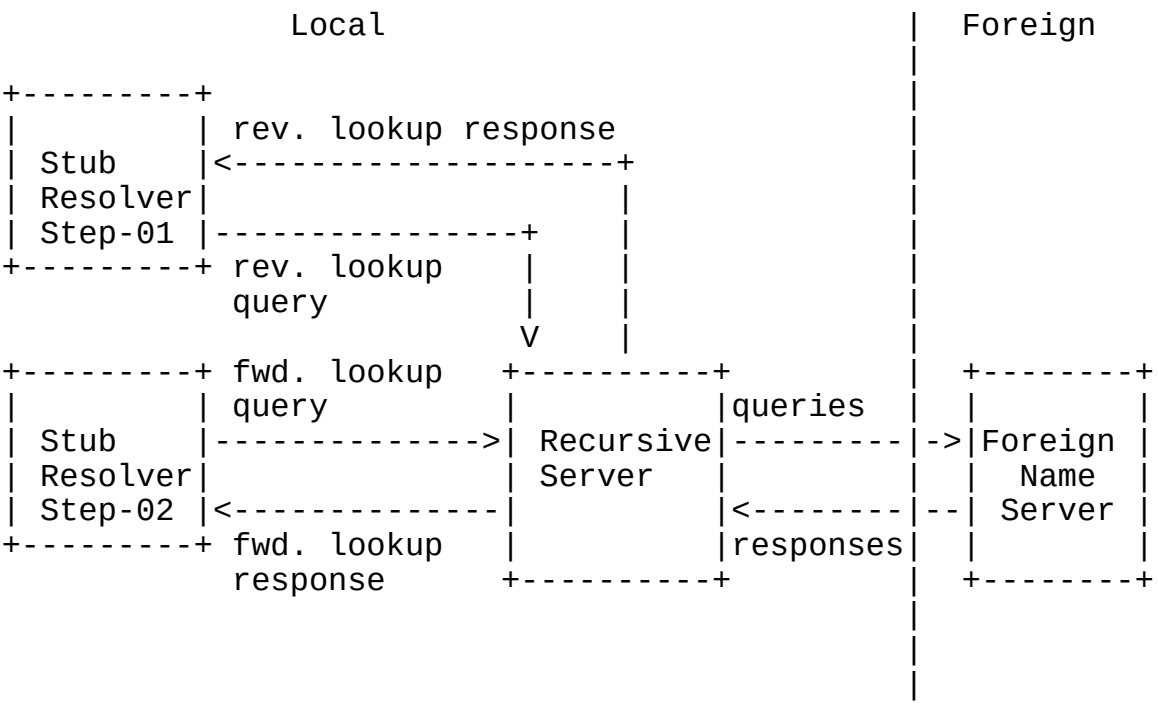

Figure 1

1. The stub-resolver in Step-01, sends a reverse lookup query for an A record to the recursive server to resolve the corresponding fully qualified domain name from the Foreign Name Server

2. The Foreign Name Server returns the PTR resource record against the query to the recursive server, which is responded back to the Stub-resolver as a response.

3. For the received domain name in Step-01, the stub-resolver in Step-02, sends a forward lookup query to recursive server to resolve the corresponding AAAA resource record from the Foreign Name Server.

4. The Foreign Name Server returns the AAAA resource record against the query to the recursive server, which is responded back to the Stub-resolver as a response.

Here, the bottleneck in this process is that now a days, mostly domains has different PTR records for a corresponding A or AAAA resource record. In this case the aforementioned process in figure-01 is not suitable. Also, this process requires to make changes to the Stub-resolver functionality to pursue the aforementioned process. Even, if the Stub-resolver functionality is modified it will work only if a single domain name is used for both A and AAAA record. The proposed solution (IVIPTR) is that, when the

Stub-resolver send a query to the recursive server for resolving AAAA record against an IPv4 address and vice versa, it will respond with the desired resource record (RR) without depending upon a Fully Qualified Domain Name(FQDN) knowledge on Stub-resolver. The term IVI in the proposed IVIPTR resource record is borrowed from one of the IPv4/IPv6 transition mechanisms address translation algorithm [RFC6219].

2. Motivation and Usecases

IPv4 is the principal protocol being used for communication in most of the organizations. Primarily, the need of IVIPTR RR in DNS evolved in a lab environment during the translation of IPv4 security rules to IPv6 security rules in a network security component (Firewall). This section discuss four usecases for the proposed DNS resource record.

2.1. Usecase-01: Firewall Automation

In network security components, mostly traffic monitoring is done through rule based filtering. An organization may enable IPv6 for certain reasons such as:

1. Functionality testing of a newly developed application with IPv6.

2. Performance and compatibility testing of application with IPv6.

3. Or, the organization has decided to keep their network on dual stack from onwards for transition purpose etc.

As a result the security guys has to maintain dual security rules for both Inbound and Outbound network traffic. This can be done by manually configuring the security rules in all network security components for the newly enabled Internet protocol IPv6. Mistakenly, configuring any security rule can result in an undesired consequences.

To automate the security configuration process in a network, there is a need to resolve IPv6 address for a corresponding IPv4 address against every security rule in a network security component (Firewall). The only resource in any network available for this automation process is the DNS. Currently in DNS, there is no such mechanism that can return IPv6 address of a domain if IPv4 address is known or vice versa. The IVIPTR Resource Record conceived as a solution to the problem for resolving IPv6 address if IPv4 address is known or IPv4 address if IPv6 address is known.

There may exist IPv4 or IPv6 address in network security components rules, which does not belong to any fully qualified domain name (FQDN) and thus, are out of the scope of this work. The presence of this IVIPTR Resource Record in the reverse zone file of an authoritative name server can result in automating a number of service for enabling them to reconfigure their security rules for the newly enabled address family protocol i.e. IPv4 or IPv6.

2.2. Usecase-02: Promoting IPv6 Usage

When accessing service such as FTP for a domain say example.com, a user can connect to the server by either:

1. ftp example.com

2. Or, ftp 192.168.0.1

For the second FTP access mechanism, the IVIPTR RR will help to retrieve the IPv6 address against the IPv4 address of the FTP server. Further, the user application will use the newly retrieved IPv6 for connectivity instead of the given one to promote the usage of IPv6 as the priority Internet address for connectivity.

2.3. Usecase-03: Customized Debugging Utilities

Debugging utilities such as traceroute can be customized in such a way that it will give detailed response. For example if a user gives a traceroute command as:

traceroute++ 192.168.0.1 or traceroute++ example.com

Thus, the output will be both PTR record and IVIPTR record.

2.4. Usecase-04: Spam Filtering

When applying spam filtering policy for a mail server such as mail.example.com, the IVIPTR can be helpful in providing additional details such as:

If filtering is performed on IPv4 address, the same can be done for IPv6 address for the corresponding mail server

3. The IVIPTR Resource Record

The IVIPTR RR has mnemonic IVIPTR and type code TBD (decimal). The IVIPTR RR has the following format:

<OWNER> <TTL> <CLASS> IVIPTR <IVI target >

The OWNER is either unqualified or fully qualified domain name depending upon the configuration of reverse zone file optional directive $ORIGIN. The TTL and CLASS fields are the same as for all other PTR records in the reverse zone file. As for the usecases discussed in the previous section the fact of IVIPTR RR usage, it is to be believed that this resource record will not be required to access frequently or in some cases just once, so one can set a smaller TTL value for this resource record to facilitate the recursive server by reducing the cache from unnecessary increase.

IVIPTR is the new RR type that points to a fully qualified domain name (FQDN) i.e. IVI target in a reverse zone file. The <IVI target> from onwards for simplicity written as <target> SHOULD be a fully qualified domain name (FQDN).

The presence of <IVIPTR RR> in a reverse zone can be elaborate by considering the domain example.com. Realistically, most of the times labels in a domain name for an IPv4 and IPv6 glue record are different. There are two possible scenarios for configuration of forward lookup zone file.

3.1. Ideal Scenario

An ideal scenario for a forward lookup zone file would be the one in which, labels in a domain name are same for both IPv4 and IPv6 glue records as:

```
; zone file for example.com

x.example.com.  IN A 192.168.0.1

x.example.com.  IN AAAA 2001:DB8:0::1
```

3.2. Non-Ideal Scenario

A non-ideal scenario for a forward lookup zone file would be the one in which, labels in a domain name are slightly different for both IPv4 and IPv6 glue records as:

```
; zone file for example.com

x.example.com.  IN A 192.168.0.1

x6.example.com.  IN AAAA 2001:DB8:0::1
```

The use of IVIPTR RR is effective only against forward lookup zone file Non-Ideal configuration scenario. Although, it will cause no issue with the Ideal scenario except additional processing overhead.

The representation of IVIPTR RR against both the IPv4 and IPv6 addresses would be as discussed in the next two sub-sections respectively.

3.3. Reverse zone file for IPv4 network prefix

When configuring a reverse zone file for example.com of IPv4 network prefix, the representation of IVIPTR RR type would be as:

```
; reverse zone file example.com for IPv4

1.0.168.192.IN-ADDR.APRPA.  IN PTR x.example.com.

1.0.168.192.IN-ADDR.ARPA.  IN IVIPTR x6.example.com.
```

3.4. Reverse zone file for IPv6 network prefix

When configuring a reverse zone file for example.com of IPv6 network prefix, the representation of IVIPTR RR type would be as:

```
; reverse zone file example.com for IPv6

1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.IP
6.ARPA.  IN PTR x6.example.com.

1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.8.b.d.0.1.0.0.2.IP
6.ARPA.  IN IVIPTR x.example.com.
```

For the purpose of simplicity in the forward lookup zone file, there is no referral RR Type such as CNAME is listed. In case of presence of any referral record in the forward lookup zone file the <IVI target > in both of the reverse zone files SHOULD be the same as the CNAME < target > in the forward lookup zone file. Thus, IVIPTR MUST follow the rule of robustness principle discussed in section 3.6.2 of [RFC1034] to avoid extra indirections in accessing information.

4. Query Processing

The IVIPTR follow the top level RR format and semantics as defined in the section 3.2.1 of RFC 1035 [RFC1035].

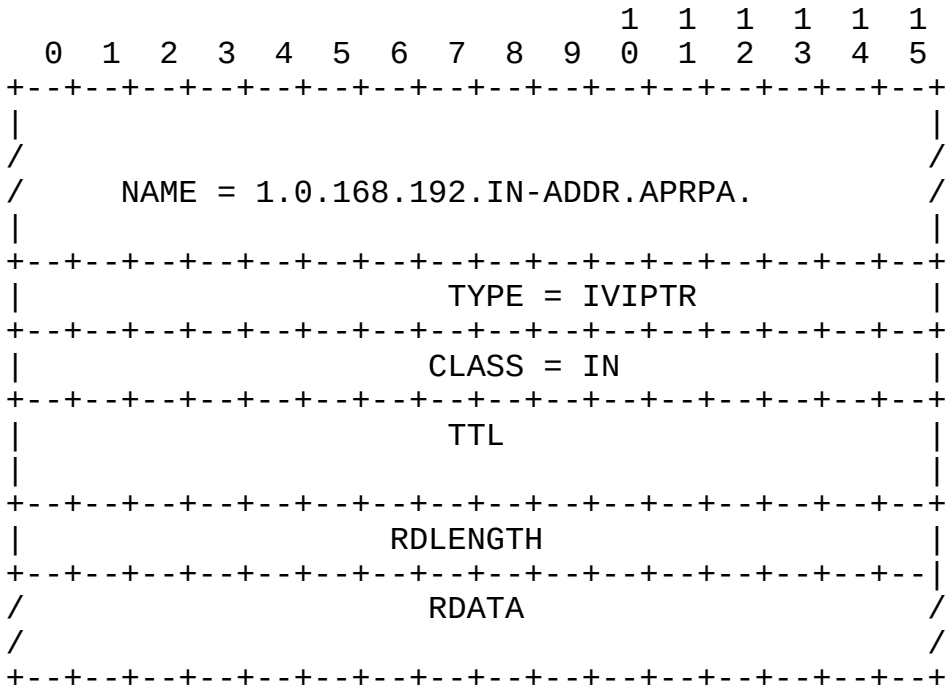

Figure 2

Where:

NAME: the owner name, same as in any reverse lookup query.

TYPE: the two octets field containing the IVIPTR RR TYPE code.

CLASS: two octets containing the RR IN CLASS code value 1.

TTL: the time interval in seconds that the resource record may be cached before the source of the information again to be contacted.

RDLENGTH: specifies the length of RDATA field.

RDATA: A variable length string of octets that represents the <IVI target> resource. The resource depends on the owner in the NAME field of the query.

The query processing is same as any other DNS query except that when the recursive server receives the response for the IVIPTR RR, first it will cache the response like any other resource record and then it will form a new query based on the rules in the sub-sections of this section.

4.1. Client Query: Case-01

If the original query NAME field contains IPv4 representation and TYPE field is IVIPTR then:

1. Upon receiving the response at the recursive server, it SHOULD form a new query.

2. The NAME field of the new query SHOULD be mapped appropriately in the desired format to the IVIPTR target in RDATA resource.

3. The TYPE field for the new query SHOULD be AAAA.

4. This query will be resolved as any other forward lookup query. Upon receiving the response which will contain AAAA RR type target, the recursive server will place this in the answer section of the original query request from client. The IVIPTR RR SHOULD cause no additional section processing.

5. In case of failure or any error the standard error response will be send back to the stub-resolver against the original query request.

4.2. Client Query: Case-02

If the original query NAME field contains IPv6 representation and TYPE field is IVIPTR then:

1. Upon receiving the response at the recursive server, it SHOULD form a new query.

2. The NAME field of the new query SHOULD be mapped appropriately in the desired format to the target in RDATA resource.

3. The TYPE field for the new query SHOULD be A.

4. This query will be resolved as any other forward lookup query. Upon receiving the response which will contain A RR type target, the recursive server will place this in the answer section of the original query request from client. The IVIPTR RR SHOULD cause no additional section processing.

5. In case of failure or any error the standard error response will be send back to the stub-resolver against the original query request.

5. Security Considerations

On a security-aware name server, while resolving the IVIPTR the query processing involves a forward lookup on recursive server in both Section 4.1 and section 4.2 when the new query is formed. The forward lookup in both the cases SHOULD comply completely with the DNSSEC on a security-aware name server and stub-resolver.

6. Acknowledgements

7. IANA Considerations

8. Normative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <https://www.rfc-editor.org/info/rfc1034>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC6219] Li, X., Bao, C., Chen, M., Zhang, H., and J. Wu, "The China Education and Research Network (CERNET) IVI Translation Design and Deployment for the IPv4/IPv6 Coexistence and Transition", RFC 6219, DOI 10.17487/RFC6219, May 2011, <https://www.rfc-editor.org/info/rfc6219>.

Authors' Addresses

Tariq Saraj (editor)
Riphah Institute of Systems Engineering
Aga Khan Road, Sector F-5/1
Islamabad, Federal Capital 44000
Pakistan

Phone: 00923345755556
Email: tariqsaraj@gmail.com

Muhammad Yousaf
Riphah Institute of Systems Engineering
Aga Khan Road, Sector F-5/1
Islamabad, Federal Capital 44000
Pakistan

Email: Muhammad.Yousaf@riu.edu.pk

Amir Qayyum
Capital University of Science and Technology
Kahuta Road, Islamabad Expressway
Islamabad, Federal Capital 44000
Pakistan

Email: aq@cust.edu.pk