Domain Name System Operations
Internet-Draft
Intended status: Best Current Practice
Expires: May 4, 2016

X. Lee
H. Li
H. Zhang
P. Zuo
CNNIC

November 2015

An approach to improve recursion performance
draft-lee-dnsop-recursion-performance-improvement-00

Abstract

A recursive DNS server generally uses random port numbers to send outbound requests to protect against DNS spoofing attacks. Due to the limitation of operation system, a process typically can only open numerable file descriptors simultaneously. This limit reduces recursion performance of resolvers. This draft offers an approach to improve both recursion performance and security for recursive servers.

Status of This Memo

Copyright Notice

1. Introduction

A recursive DNS server generally uses random port numbers to send outbound requests to avoid cache poisoning. This is also essential to protect against Dan Kaminsky’s DNS attack. Due to the limitation of operation system, a process typically can only open numerable file descriptors simultaneously. For example, the limitation on Linux is 1024 by default. Although this configuration could be modified by operation system operators, there is still a limit for maximum port number (65535).This limitation not only reduces recursion performance, but also makes resolvers vulnerable to attackers. Suppose that a hacker sends thousands of queries for domains which change irregularly and are actually not exist, the resolver must start corresponding recursive requests to authoritative servers as these domains are not cached, and soon, this resolver will not be able to generate more outbound requests because no more file descriptor can be open.

This draft proposes an approach to solve this problem. A resolver should reuse a group of fixed port numbers for outbound requests. In this case, the resolver could improve recursion performance greatly as it avoids limitation of maximum file descriptors. As for security, the resolver should add an extra recursive identifier(RQID) in EDNS0 record in outbound requests. Authoritative servers should copy this RQID in EDNS0 record to the response packet. The resolver then match up the reply using RQID option. In this case, this approach improves security because a hacker can hardly correctly guess the randomized 32bit RQID besides the Transaction ID in DNS message header.

2. Protocol changes

This draft uses an EDNS0 ([RFC6891]) option to include recursive query ID(RQID) in DNS messages. The option is structured as follows:

```
             +0 (MSB)                +1 (LSB)
             +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
         0:  |                   OPTION-RCODE                |
             +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
          2: |                   OPTION-LENGTH               |
             +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
          4: |                                               |
             +            (RQID)Recursive Query ID           |
          6: |                                               |
             +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

- (Defined in [RFC6891]) OPTION-CODE, 2 octets, for RQID is 10.

- (Defined in [RFC6891]) OPTION-LENGTH, 2 octets, use 4 as recommended.

- RQID, a 32 bit identifier(recommended) assigned by the resolver that generates a recursive query. This RQID should be copied by authoritative server to the corresponding reply message and then be used by the resolver to match up the reply.

3. Stub Resolver Considerations

This approach is aimed to work between recursive servers and authoritative servers, a stub resolver by itself should determine if it has to support this RQID option.

4. Recursive Server Considerations

A resolver supporting RQID option should reuse fixed port numbers to send recursive queries to improve performance. The port number could be configured by users and the RQID function should also be configurable.

If the RQID function is enabled, a resolver should behave as follows: When sending a query, it generates a random 32 bit RQID (recommended) in the EDNS0 record as described above. This RQID option indicates the resolver reuses its port to send recursive queries and expect the authoritative server to copy the RQID option in the responses. When receiving a response, it should check RQID option in EDNS0 record in the response to match up the reply. If the response contains no EDNS0 record or RQID option, the resolver itself should determine if to accept this reply. For compatibility, it is recommended to accept these replies. This could increase the risk of cache poisoning, but in most cases, the resolver should be secured by other equipments(firewalls etc).

If a recursive server receives a query containing RQID option from stub resolver, it should copy this option in the reply.

5. Authoritative Server Considerations

If an authoritative server does not support RQID function, it just ignores RQID option in EDNS0 record.

If a authoritative server supports RQID function, it should copy the RQID option in the reply.

6. Performance Considerations

The recursive query performance should be greatly improved as the resolver reuses port numbers to avoid the operation system limit of maximum file descriptors.

7. Security Considerations

This draft proposes an approach to use a RQID(32 bit as recommended) option to match up DNS replies. If both the recursive and authoritative server support this option, the risk of cache poisoning is much lower than previous protocol.

For compatibility, it is recommended to accept DNS replies which contain no RQID option.

8. IANA Considerations

IANA is requested to assign the option code 10 for the RQID Option Code in the EDNS0 meta-RR.

9. References

[RFC1035] Mockapetris, P., ""DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION"", November 1987, <http://www.rfc-editor.org/info/rfc1035>.

[RFC6891] Damas, J., Graff, M., and P. Vixie, ""Extension Mechanisms for DNS (EDNS(0))"", April 2013, <http://www.rfc-editor.org/info/rfc6891>.

Authors’ Addresses

Xiaodong Lee
CNNIC
4 South 4th Street,Zhongguancun,Haidian District
Beijing, Beijing 100190
China

Phone: +86 10 5881 3020
Email: lee@cnnic.cn

Hongtao Li
CNNIC
4 South 4th Street,Zhongguancun,Haidian District
Beijing, Beijing 100190
China

Phone: +86 10 5881 3164
Email: lihongtao@cnnic.cn

Haikuo Zhang
CNNIC
4 South 4th Street,Zhongguancun,Haidian District
Beijing, Beijing 100190
China

Phone: +86 10 5881 3163
Email: zhanghaikuo@cnnic.cn

Peng Zuo
CNNIC
4 South 4th Street,Zhongguancun,Haidian District
Beijing, Beijing 100190
China

Phone: +86 10 5881 2629
Email: zuopeng@cnnic.cn