Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 11, 2015

K. Fujiwara
JPRS
A. Kato
Keio/WIDE
March 10, 2015

Aggressive use of NSEC/NSEC3
draft-fujiwara-dnsop-nsec-aggressiveuse-00

Abstract

DNS highly depends on cache, however, cache usage of non-existence information was limited to exact matching. This draft proposes the aggressive use of NSEC/NSEC3 resource record, which is able to express non-existence of range of names authoritatively. With this proposal, shorter latency to many of negative response is expected as well as some level of mitigation of random sub-domain attacks (referred to as "Water Torture" attacks). And more, non-existent TLD queries to Root DNS servers will decrease.

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 September 11, 2015.

Copyright Notice

Table of Contents

1. Introduction

While negative (non-existence) information of DNS caching mechanism has been known as DNS negative cache [RFC2308], it requires exact matching in most cases. Assume that "example.com" zone doesn't have names such as "a.example.com" and "b.example.com". When a full resolver receives a query "a.example.com" , it performs a DNS resolution process, and eventually gets NXDOMAIN and cache it into its negative cache. When the full resolver receives another query "b.example.com", it doesn't match with "a.example.com". So it will send a query to one of the authoritative servers of "example.com". This was because the NXDOMAIN response just says there is no such name "a.example.com" and it has no ability to tell that there is no such name "b.example.com".

By the way, DNSSEC [RFC4035] [RFC5155] has been practically deployed recently. Two resource record types (NSEC and NSEC3) are used for authentic non-existence. For a zone signed with NSEC, it may be possible to use the information carried in NSEC resource records to indicate that the range of names where no valid name exists. Such use is discouraged by Section 4.5 of RFC 4035.

This document proposes to make a minor change to RFC 4035 and the full resolver can use NSEC/NSEC3 resource records aggressively.

2. Problem Statement

Random sub-domain attacks (referred to as "Water Torture" attacks) send many non-existent queries to full resolvers. Their query names consist of random prefixes and a target domain name. As a result, the negative cache does not work well and target full resolvers result in sending queries to authoritative DNS servers of the target domain name.

When number of queries is very large, the full resolver's outstanding queue will be full, and then, the full resolver will drop queries from both users and attackers.

The countermeasures performed at present are rate limiting and disabling name resolution of target domain names.

3. Possible Solution

If the target domain names are DNSSEC signed, aggressive use of NSEC/NSEC3 resource records solves the problem.

DNSSEC defines NSEC resource record. Section 4.5 of [RFC4035] shows that "In theory, a resolver could use wildcards or NSEC RRs to generate positive and negative responses (respectively) until the TTL or signatures on the records in question expire. However, it seems prudent for resolvers to avoid blocking new authoritative data or synthesizing new data on their own. Resolvers that follow this recommendation will have a more consistent view of the namespace".

To reduce non-existent queries to authoritative DNS servers, the countermeasure is to relax this restriction.

Then, DNSSEC enabled full resolvers MAY use NSEC/NSEC3 resource records to generate negative responses until their effective TTLs or signatures on the records in question expire.

This technique is called as "NSEC/NSEC3 aggressive negative caching" in Unbound [Unbound] TODO file. Unbound has aggressive negative caching code in its DLV validator.

The full resolver need to check the existence of wildcards. If the cache does not have an NSEC/NSEC3 resource record whose range includes a wildcard ('*') in a zone which the query name belongs to, wildcard may exist in the zone, then, the aggressive use of NSEC/ NSEC3 cannot be applied and the full resolver need to send the query to authoritative DNS servers.

If the zone has a wildcard and it is in full resolver's cache, the full resolver may generate positive responses from the wildcard in the cache.

This approach is effective for DNSSEC signed zones with NSEC or NSEC3, except zones with NSEC3 Opt-Out.

NSEC/NSEC3 aggressive negative caching works as follows. When the query name has a matching NSEC or NSEC3 resource records in the cache and there is no wildcard in the zone which the query name belongs to, a full resolver is allowed to respond with NXDOMAIN error immediately.

The matching procedure may be applied to all ancestor domain names of the query name.

This function needs care on the TTL value of negative information because newly added domain names cannot be used while the negative information is effective. RFC 2308 states the maximum number of negative cache TTL value is 10800, and this value is reasonable small but still effective for the purpose of this document.

It can eliminate significant amount of DNS queries when an attacker tries to send large number of DNS queries by using randomly generated names.

The same discussion is applicable for wildcards. If a query name is covered by NSEC or NSEC3 resource records in the cache and there is a covering wildcard, full resolvers can use wildcards to generate

positive responses until wildcard and NSEC/NSEC3 resource records in the cache are effective.

Aggressive use of wildcards requires aggressive use of negative information because there may be other domain names.

4. Possible side effect

Aggressive use of NSEC/NSEC3 resource records may decrease queries to Root DNS servers.

People may generate many typos and they tend to generate DNS queries. Some implementations leak non-existent TLD queries whose second level domain are different each other. Well observed TLDs are ".local" and ".belkin". With this proposal, it is possible to return NXDOMAIN to such queries without further DNS recursive resolution process. It may reduces round trip time, as well as reduces the DNS queries to corresponding authoritative servers, including Root DNS servers.

5. Another option

The proposed technique is applicable to zones where there is a NSEC record to each owner name in the zone even without DNSSEC signed. And it is also applicable to full resolvers without DNSSEC validation. Full resolvers can set DNSSEC OK bit in query packets and they will cache NSEC/NSEC3 resource records. They can apply aggressive use of NSEC/NSEC3 resource records without DNSSEC validation.

It is highly recommended to sign the zone, of course, and it is recommended to apply DNSSEC validation of NSEC record prior to cache it in the negative cache.

6. IANA Considerations

This document has no effect on IANA registries.

7. Security Considerations

Newly registered resource records may not be used immediately. However, choosing suitable TTL value will mitigate the problem and it

is not a security problem.

It is also suggested to limit the maximum TTL value of NSEC resource records in the negative cache to, for example, 10800 seconds (3hrs), to mitigate the issue. Implementations which comply with this proposal is suggested to have a configurable maximum value of NSEC RRs in the negative cache.

Aggressive use of NSEC/NSEC3 resource records without DNSSEC validation may cause security problems.

8. References

8.1. Normative References

[RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", RFC 2308, March 1998.

[RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

[RFC5155] Laurie, B., Sisson, G., Arends, R., and D. Blacka, "DNS Security (DNSSEC) Hashed Authenticated Denial of Existence", RFC 5155, March 2008.

8.2. Informative References

[Unbound] NLnet Labs, "Unbound DNS validating resolver", <http://www.unbound.net/>.

Authors' Addresses

Kazunori Fujiwara
Japan Registry Services Co., Ltd.
Chiyoda First Bldg. East 13F, 3-8-1 Nishi-Kanda
Chiyoda-ku, Tokyo 101-0065
Japan

Phone: +81 3 5215 8451
Email: fujiwara@jprs.co.jp

Akira Kato
Keio University/WIDE Project
Graduate School of Media Design, 4-1-1 Hiyoshi
Kohoku, Yokohama 223-8526
Japan

Phone: +81 45 564 2490
Email: kato@wide.ad.jp