Network Working Group
Internet-Draft
Expires: August 31, 2001

M. Kosters
Network Solutions, Inc.
March 2, 2001

DNSSEC Opt-in for Large Zones
draft-kosters-dnsext-dnssec-opt-in-01.txt

Status of this Memo

Copyright Notice

Abstract

In order for DNSSEC to be deployed operationally with large zones and little operational impact, there needs to be included a mechanism that allows for the separation of secure versus unsecure views of zones. This needs to be done in a transparent fashion that allows DNSSEC to be deployed in an incremental manner. This document proposes the use of an extended RCODE to signify that a DNSSEC-aware requestor may have to re-query for the information, if and only if, the delegation is not yet secure. Thus, one can maintain two views of the zone and expand the DNSSEC zone as demand warrants.

Table of Contents

1. Introduction

DNS is an unsecure system. The key features that gives DNS its power can also be its chief weaknesses. One feature is the facility to delegate branches of information from one set of servers to another. Currently, this is done in a non-cryptographically verified way that allows spoofing attacks. For example, an alternative domain registry called AlterNIC exploited this vulnerability to redirect www.netsol.com and www.internic.net websites to their own website in July 1997 that gained widespread exposure. If this delegated information had been cryptographically verified, this attack would not have been able to occur.

In recent years, there has been much work within the IETF regarding DNS security. There are a number of RFCs that integrate public key technology within DNS to enable cryptographically-verified answers. To this end, three new resource record types (RR's) have been defined:

- KEY -- a public key of the zone
- SIG - a signature of an accompanying RR
- NXT - a negative response record

Within the zone, each authoritative RR will have accompanying SIG RR's that can be verified with the KEY RR of the zone. Each KEY RR can be verified hierarchically with a SIG RR from the direct parent zone. For unsecure delegations, a null-KEY RR is inserted in the parent zone. Finally, NXT RR's and their accompanying SIG RR's are issued in the case of a negative reply.

As a zone maintainer, transitioning to a secure zone has a high overhead in the following areas:

KEY RR
: At a delegation point, the zone maintainer needs to place a NULL key and accompanying SIG RR's when the child zone is not known to be secure.

NXT RR
: Each delegation needs to be lexigraphically ordered so that a NXT RR can be generated and signed with SIG RR's. For large zone operators, generating the zone file is a very time consuming process. In the resolution process, NXT lookups require that the server replace efficient hash structures with a lexigraphically ordered search structure that degrades lookup performance. This lookup performance is a critical element for a high-query rate DNS server.

Thus, the net effect is when one initially secures a zone as defined

in RFC2535[4], the net overhead is massive because of the following

factors:

1. Zone ordering and maintenance for large zones is difficult and expensive.
2. Adding null-KEY RR's, NXT RR's and their accompanying SIG RR's for unsecure delegations will consume large amounts of memory (6x the current memory requirements).
3. Having a less efficient look-up algorithm to provide answers to queries will degrade overall performance.
4. Very little initial payoff (anticipate only a small fraction of delegations to be signed. This equates to less than 1% over the first six months).
5. Unsecured delegations are more expensive at the parent than secure delegations (NULL KEY).

2. Rationale

As DNSSEC is initially deployed, it is anticipated that DNSSEC adoption will be slow to materialize. It is also anticipated that DNSSEC security resolution will be top down. Thus for DNSSEC to be widely adopted, the root zone and GTLD zones will need to be signed. Based on the implications previously listed, a large zone maintainer such as the administrator of COM, needs to create an infrastructure that is an order of magnitude larger than its current state with very little initial benefit.

This document proposes an alternative opt-in approach that minimizes the expense and complexity to ease adoption of DNSSEC for large zones by allowing for an alternate view of secured only delegations.

3. Protocol Additions

The opt-in proposal allows for a zone operator to maintain two views of its delegations - one being non-DNSSEC and the other being DNSSSEC aware. The non-DNSSEC view will have all delegations - both secured and non-secured. The DNSSEC aware view will only have secured delegations. It is assumed that neither view will have any innate knowledge of the other's delegations. Thus, the cost of securing a zone is proportional to the demand of its delegations with the added benefit of no longer having to maintain NULL KEY RRs for unsecure delegations.

On the server side, identification of the zone being opt-in will be identified by using one of the reserved bits of the flags section within the KEY RR for that particular zone [note - the actual bit needs yet to be selected out of reserved bits 4-5 or 8-11].

On the client side, the client MUST be identified by sending a option-code of RETRY-NO-SEC-AWARE within the OPT RR RDATA to ensure

that it can accept and understand the RETRY-NO-SEC RCODE. The RETRY-NO-SEC-AWARE option-code MUST have an option-length value of zero with no option-data. The RETRY-NO-SEC-AWARE option-code will be determined by IANA.

To determine which view each DNS query packet is to be queried against, there is a simple algorithm to be followed:

1. The DNSSEC view is to be queried when the DO bit is set within the EDNS0 OPT meta RR as indicated in [6] Additionally,
2. The DNSSEC view is to be queried when the query type is SIG, KEY, or NXT and the RRs added match the query name and query type.

If the query does not follow either case (1) or (2), the non-DNSSEC view MUST be consulted by default.

Since the DNSSEC view will have a subset of the actual delegations of that zone, it will not be able to respond to an unsecured delegation. To that end, one of two things will happen:

1) If the client has been identified as RETRY-NO-SEC-AWARE, a new extended RCODE MUST be set within the EDNS OPT RR for the resolver to retry again with the DO bit not set. This RCODE is referred to as "RETRY-NO-SEC" (RS). In the context of the EDNS0 OPT meta-RR, the RS value will be determined by IANA.

Setting the RS RCODE in a response indicates to the resolver that the resolver is retrying the query again without the DO bit set. The behavior of the authority and additional records section being populated should be the same using the RS RCODE as the RCODE being set to NXDOMAIN. Therefore, the resolver will be able to verify that the answer does not exist within the secure zone since the NXT RR will be sent in the Authority section. To avoid caching, the server SHOULD set the TTL on the NXT RR to 0.

2) If the client has been identified as not being RETRY-NO-SEC-AWARE, the server itself MUST consult the non-secure view to compile the answer and respond back to the client. If the RR exists, the answer will show up normally with in the Answer and Additional sections and the NXT RR's within the Authority section along with the KEY RR and its SIG in the Additional section. If the RR does not exist, RCODE will be set to NXDOMAIN with the NXT RR will be sent in the Authority section along with the KEY RR and its SIG in the Additional section . Again, to avoid caching, the server SHOULD set the TTL on the NXT RR to 0.

Note that latter case should be used during the transition of moving to clients that understand the RS RCODE only. It should not be

viewed as a permanent solution and may deprecated in a short period of time.

Example:

Consider a zone with the secure names 3, 6, and 9, and unsecure names 2, 4, 5, 7, and 8.

Unsecured zone Contents:

```
@ SOA
2 NS
3 NS
4 NS
5 NS
6 NS
7 NS
8 NS
9 NS
```

Secured zone Contents:

```
@ SOA, SIG SOA, NXT(3), SIG NXT
3 NS, SIG NS, NXT(6), SIG NXT
6 NS, SIG NS, NXT(9), SIG NXT
9 NS, SIG NS, NXT(@), SIG NXT
```

1. Query for 5 RR type A with EDNS0 DO bit set along with the RETRY-NO-SEC-AWARE option code, the response would return with the extended RCODE RS bit set:

```
RCODE=RS
Authority Section:
  SOA, SIG SOA, 3 NXT(6), SIG NXT
Additional Section:
  KEY, SIG KEY
```

The source would then retry without the EDNS0 DO bit set which would return an answer as defined in RFC1035[2].

2. Query for 5 RR type A with EDNS0 DO bit only, the response would return with the following:

```
RCODE=NOERROR
Answer Section:
  5 NS
Authority Section:
```

```
    3 NXT(6), SIG NXT
  Additional Section:
    KEY, SIG KEY
```

3. Query for 55 RR type A with EDNS0 DO bit set along with the RETRY-NO-SEC-AWARE option code, the response would return with the extended RCODE RS bit set:

```
  RCODE=RS
  Authority Section:
    SOA, SIG SOA, 3 NXT(6), SIG NXT
  Additional Section:
    KEY, SIG KEY
```

The source would then retry without the EDNS0 DO bit set which would return an answer as defined in RFC1035[2]. The subsequent 1035 answer would contain a RCODE of NXDOMAIN since the domain 55 does not exist.

4. Query for 3 RR type KEY without EDNS DO bit set. The response would return with an answer as defined in RFC2535[4].

5. Query for 3 RR type A, with EDNS0 DO bit set, the response would be the same as defined in RFC2535[4].

4. Security Considerations

This draft is different and separate from RFC2535[4] in that it allows for secured delegation paths to exist but does not allow for secure answers to unsecured delegations at the parent level. Increased exposure will be marginal given that the children are unsecure.

5. IANA Considerations

1) Allocation of a bit within the reserved portion of the KEY RR to indicate that the zone is an opt-in zone.

2) Allocation of the most significant bit of the RCODE field in the EDNS0 OPT meta-RR is required.

3) Allocation of an option-code within the OPT RR to indicate that the client can understand the new RCODE.

6. Acknowledgements

This document is based on a rough draft by Brian Wellington, and input from Olafur Gudmundsson.

References

[1] Mockapetris, P.V., "Domain names - concepts and facilities", RFC 1034, STD 13, Nov 1987.

[2] Mockapetris, P.V., "Domain names - implementation and specification", RFC 1035, STD 13, Nov 1987.

[3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, BCP 14, March 1997.

[4] Eastlake, D., "Domain Name System Security Extensions", RFC 2535, March 1999.

[5] Vixie, P., "Extension Mechanisms for DNS (EDNS0)", RFC 2671, August 1999.

[6] Conrad, D. R., "Indicating Resolver Support of DNSSEC (work in progress)", August 2000.

Author's Address

Mark Kosters
Network Solutions, Inc.
505 Huntmar Park Drive
Herndon, VA 22070
US

Phone: +1 703 948-3362
EMail: markk@netsol.com
URI: http://www.netsol.com

Full Copyright Statement

Acknowledgement

Funding for the RFC editor function is currently provided by the Internet Society.