

Indicating Resolver Support of DNSSEC

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

In order to deploy DNSSEC (Domain Name System Security Extensions) operationally, DNSSEC aware servers should only perform automatic inclusion of DNSSEC RRs when there is an explicit indication that the resolver can understand those RRs. This document proposes the use of a bit in the EDNS0 header to provide that explicit indication and describes the necessary protocol changes to implement that notification.

1. Introduction

DNSSEC [RFC2535] has been specified to provide data integrity and authentication to security aware resolvers and applications through the use of cryptographic digital signatures. However, as DNSSEC is deployed, non-DNSSEC-aware clients will likely query DNSSEC-aware servers. In such situations, the DNSSEC-aware server (responding to a request for data in a signed zone) will respond with SIG, KEY, and/or NXT records. For reasons described in the subsequent section, such responses can have significant negative operational impacts for the DNS infrastructure.

This document discusses a method to avoid these negative impacts, namely DNSSEC-aware servers should only respond with SIG, KEY, and/or NXT RRs when there is an explicit indication from the resolver that it can understand those RRs.

For the purposes of this document, "DNSSEC security RRs" are considered RRs of type SIG, KEY, or NXT.

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].

2. Rationale

Initially, as DNSSEC is deployed, the vast majority of queries will be from resolvers that are not DNSSEC aware and thus do not understand or support the DNSSEC security RRs. When a query from such a resolver is received for a DNSSEC signed zone, the DNSSEC specification indicates the nameserver must respond with the appropriate DNSSEC security RRs. As DNS UDP datagrams are limited to 512 bytes [RFC1035], responses including DNSSEC security RRs have a high probability of resulting in a truncated response being returned and the resolver retrying the query using TCP.

TCP DNS queries result in significant overhead due to connection setup and teardown. Operationally, the impact of these TCP queries will likely be quite detrimental in terms of increased network traffic (typically five packets for a single query/response instead of two), increased latency resulting from the additional round trip times, increased incidences of queries failing due to timeouts, and significantly increased load on nameservers.

In addition, in preliminary and experimental deployment of DNSSEC, there have been reports of non-DNSSEC aware resolvers being unable to handle responses which contain DNSSEC security RRs, resulting in the resolver failing (in the worst case) or entire responses being ignored (in the better case).


Given these operational implications, explicitly notifying the nameserver that the client is prepared to receive (if not understand) DNSSEC security RRs would be prudent.

Client-side support of DNSSEC is assumed to be binary -- either the client is willing to receive all DNSSEC security RRs or it is not willing to accept any. As such, a single bit is sufficient to indicate client-side DNSSEC support. As effective use of DNSSEC implies the need of EDNS0 [RFC2671], bits in the "classic" (non-EDNS enhanced DNS header) are scarce, and there may be situations in which non-compliant caching or forwarding servers inappropriately copy data from classic headers as queries are passed on to authoritative servers, the use of a bit from the EDNS0 header is proposed.

An alternative approach would be to use the existence of an EDNS0 header as an implicit indication of client-side support of DNSSEC. This approach was not chosen as there may be applications in which EDNS0 is supported but in which the use of DNSSEC is inappropriate.

3. Protocol Changes

The mechanism chosen for the explicit notification of the ability of the client to accept (if not understand) DNSSEC security RRs is using the most significant bit of the Z field on the EDNS0 OPT header in the query. This bit is referred to as the "DNSSEC OK" (DO) bit. In the context of the EDNS0 OPT meta-RR, the DO bit is the first bit of the third and fourth bytes of the "extended RCODE and flags" portion of the EDNS0 OPT meta-RR, structured as follows:



Setting the DO bit to one in a query indicates to the server that the resolver is able to accept DNSSEC security RRs. The DO bit cleared (set to zero) indicates the resolver is unprepared to handle DNSSEC security RRs and those RRs MUST NOT be returned in the response (unless DNSSEC security RRs are explicitly queried for). The DO bit of the query MUST be copied in the response.

More explicitly, DNSSEC-aware nameservers MUST NOT insert SIG, KEY, or NXT RRs to authenticate a response as specified in [RFC2535] unless the DO bit was set on the request. Security records that match an explicit SIG, KEY, NXT, or ANY query, or are part of the zone data for an AXFR or IXFR query, are included whether or not the DO bit was set.

A recursive DNSSEC-aware server MUST set the DO bit on recursive requests, regardless of the status of the DO bit on the initiating resolver request. If the initiating resolver request does not have the DO bit set, the recursive DNSSEC-aware server MUST remove DNSSEC security RRs before returning the data to the client, however cached data MUST NOT be modified.

In the event a server returns a NOTIMP, FORMERR or SERVFAIL response to a query that has the DO bit set, the resolver SHOULD NOT expect DNSSEC security RRs and SHOULD retry the query without EDNS0 in accordance with section 5.3 of [RFC2671].


Security Considerations

The absence of DNSSEC data in response to a query with the DO bit set MUST NOT be taken to mean no security information is available for that zone as the response may be forged or a non-forged response of an altered (DO bit cleared) query.

IANA Considerations

EDNS0 [RFC2671] defines 16 bits as extended flags in the OPT record, these bits are encoded into the TTL field of the OPT record (RFC2671 section 4.6).

This document reserves one of these bits as the OK bit. It is requested that the left most bit be allocated. Thus the USE of the OPT record TTL field would look like



Acknowledgements

This document is based on a rough draft by Bob Halley with input from Olafur Gudmundsson, Andreas Gustafsson, Brian Wellington, Randy Bush, Rob Austein, Steve Bellovin, and Erik Nordmark.

References

- [RFC1034] Mockapetris, P., "Domain Names - Concepts and Facilities", STD 13, RFC 1034, November 1987.
- [RFC1035] Mockapetris, P., "Domain Names - Implementation and Specifications", STD 13, RFC 1035, November 1987.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2535] Eastlake, D., "Domain Name System Security Extensions", RFC 2535, March 1999.
- [RFC2671] Vixie, P., "Extension Mechanisms for DNS (EDNS0)", RFC 2671, August 1999.

Author's Address

David Conrad
Nominum Inc.
950 Charter Street
Redwood City, CA 94063
USA

Phone: +1 650 381 6003
EMail: david.conrad@nominum.com

Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

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

