

DNS Security (DNSSEC) Experiments

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 IETF Trust (2007).

Abstract

This document describes a methodology for deploying alternate, non-backwards-compatible, DNS Security (DNSSEC) methodologies in an experimental fashion without disrupting the deployment of standard DNSSEC.

Table of Contents

| | |
|--|---|
| 1. Overview | 2 |
| 2. Definitions and Terminology | 2 |
| 3. Experiments | 2 |
| 4. Method | 3 |
| 5. Defining an Experiment | 4 |
| 6. Considerations | 5 |
| 7. Use in Non-Experiments | 5 |
| 8. Security Considerations | 5 |
| 9. References | 6 |
| 9.1. Normative References | 6 |
| 9.2. Informative References | 6 |

1. Overview

Historically, experimentation with DNSSEC alternatives has been a problematic endeavor. There has typically been a desire to both introduce non-backwards-compatible changes to DNSSEC and to try these changes on real zones in the public DNS. This creates a problem when the change to DNSSEC would make all or part of the zone using those changes appear bogus (bad) or otherwise broken to existing security-aware resolvers.

This document describes a standard methodology for setting up DNSSEC experiments. This methodology addresses the issue of coexistence with standard DNSSEC and DNS by using unknown algorithm identifiers to hide the experimental DNSSEC protocol modifications from standard security-aware resolvers.

2. Definitions and Terminology

Throughout this document, familiarity with the DNS system (RFC 1035 [5]) and the DNS security extensions (RFC 4033 [2], RFC 4034 [3], and RFC 4035 [4]) is assumed.

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 RFC 2119 [1].

3. Experiments

When discussing DNSSEC experiments, it is necessary to classify these experiments into two broad categories:

Backwards-Compatible: describes experimental changes that, while not strictly adhering to the DNSSEC standard, are nonetheless interoperable with clients and servers that do implement the DNSSEC standard.

Non-Backwards-Compatible: describes experiments that would cause a standard security-aware resolver to (incorrectly) determine that all or part of a zone is bogus, or to otherwise not interoperate with standard DNSSEC clients and servers.

Not included in these terms are experiments with the core DNS protocol itself.

The methodology described in this document is not necessary for backwards-compatible experiments, although it certainly may be used if desired.

4. Method

The core of the methodology is the use of strictly unknown algorithm identifiers when signing the experimental zone, and more importantly, having only unknown algorithm identifiers in the DS records for the delegation to the zone at the parent.

This technique works because of the way DNSSEC-compliant validators are expected to work in the presence of a DS set with only unknown algorithm identifiers. From RFC 4035 [4], Section 5.2:

If the validator does not support any of the algorithms listed in an authenticated DS RRset, then the resolver has no supported authentication path leading from the parent to the child. The resolver should treat this case as it would the case of an authenticated NSEC RRset proving that no DS RRset exists, as described above.

And further:

If the resolver does not support any of the algorithms listed in an authenticated DS RRset, then the resolver will not be able to verify the authentication path to the child zone. In this case, the resolver SHOULD treat the child zone as if it were unsigned.

Although this behavior isn't strictly mandatory (as marked by MUST), it is unlikely for a validator to implement a substantially different behavior. Essentially, if the validator does not have a usable chain of trust to a child zone, then it can only do one of two things: treat responses from the zone as insecure (the recommended behavior), or treat the responses as bogus. If the validator chooses the latter, this will both violate the expectation of the zone owner and defeat the purpose of the above rule. However, with local policy, it is within the right of a validator to refuse to trust certain zones based on any criteria, including the use of unknown signing algorithms.

Because we are talking about experiments, it is RECOMMENDED that private algorithm numbers be used (see RFC 4034 [3], Appendix A.1.1. Note that secure handling of private algorithms requires special handling by the validator logic. See "Clarifications and Implementation Notes for DNSSECbis" [6] for further details.) Normally, instead of actually inventing new signing algorithms, the recommended path is to create alternate algorithm identifiers that are aliases for the existing, known algorithms. While, strictly speaking, it is only necessary to create an alternate identifier for the mandatory algorithms, it is suggested that all optional defined algorithms be aliased as well.

It is RECOMMENDED that for a particular DNSSEC experiment, a particular domain name base is chosen for all new algorithms, then the algorithm number (or name) is prepended to it. For example, for experiment A, the base name of "dnssec-experiment-a.example.com" is chosen. Then, aliases for algorithms 3 (DSA) and 5 (RSASHA1) are defined to be "3.dnssec-experiment-a.example.com" and "5.dnssec-experiment-a.example.com". However, any unique identifier will suffice.

Using this method, resolvers (or, more specifically, DNSSEC validators) essentially indicate their ability to understand the DNSSEC experiment's semantics by understanding what the new algorithm identifiers signify.

This method creates two classes of security-aware servers and resolvers: servers and resolvers that are aware of the experiment (and thus recognize the experiment's algorithm identifiers and experimental semantics), and servers and resolvers that are unaware of the experiment.

This method also precludes any zone from being both in an experiment and in a classic DNSSEC island of security. That is, a zone is either in an experiment and only possible to validate experimentally, or it is not.

5. Defining an Experiment

The DNSSEC experiment MUST define the particular set of (previously unknown) algorithm identifiers that identify the experiment and define what each unknown algorithm identifier means. Typically, unless the experiment is actually experimenting with a new DNSSEC algorithm, this will be a mapping of private algorithm identifiers to existing, known algorithms.

Normally the experiment will choose a DNS name as the algorithm identifier base. This DNS name SHOULD be under the control of the authors of the experiment. Then the experiment will define a mapping between known mandatory and optional algorithms into this private algorithm identifier space. Alternately, the experiment MAY use the Object Identifier (OID) private algorithm space instead (using algorithm number 254), or MAY choose non-private algorithm numbers, although this would require an IANA allocation.

For example, an experiment might specify in its description the DNS name "dnssec-experiment-a.example.com" as the base name, and declare that "3.dnssec-experiment-a.example.com" is an alias of DNSSEC algorithm 3 (DSA), and that "5.dnssec-experiment-a.example.com" is an alias of DNSSEC algorithm 5 (RSASHA1).

Resolvers MUST only recognize the experiment's semantics when present in a zone signed by one or more of these algorithm identifiers. This is necessary to isolate the semantics of one experiment from any others that the resolver might understand.

In general, resolvers involved in the experiment are expected to understand both standard DNSSEC and the defined experimental DNSSEC protocol, although this isn't required.

6. Considerations

There are a number of considerations with using this methodology.

1. If an unaware validator does not correctly follow the rules laid out in RFC 4035 (e.g., the validator interprets a DNSSEC record prior to validating it), or if the experiment is broader in scope than just modifying the DNSSEC semantics, the experiment may not be sufficiently masked by this technique. This may cause unintended resolution failures.
2. It will not be possible for security-aware resolvers unaware of the experiment to build a chain of trust through an experimental zone.

7. Use in Non-Experiments

This general methodology MAY be used for non-backwards compatible DNSSEC protocol changes that start out as or become standards. In this case:

- o The protocol change SHOULD use public IANA allocated algorithm identifiers instead of private algorithm identifiers. This will help identify the protocol change as a standard, rather than an experiment.
- o Resolvers MAY recognize the protocol change in zones not signed (or not solely signed) using the new algorithm identifiers.

8. Security Considerations

Zones using this methodology will be considered insecure by all resolvers except those aware of the experiment. It is not generally possible to create a secure delegation from an experimental zone that will be followed by resolvers unaware of the experiment.

Implementers should take into account any security issues that may result from environments being configured to trust both experimental and non-experimental zones. If the experimental zone is more

vulnerable to attacks, it could, for example, be used to promote trust in zones not part of the experiment, possibly under the control of an attacker.

9. References

9.1. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.
- [3] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.
- [4] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", RFC 4035, March 2005.

9.2. Informative References

- [5] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [6] Weiler, S. and R. Austein, "Clarifications and Implementation Notes for DNSSECBis", Work in Progress, March 2007.

Author's Address

David Blacka
VeriSign, Inc.
21355 Ridgetop Circle
Dulles, VA 20166
US

Phone: +1 703 948 3200
EMail: davidb@verisign.com
URI: <http://www.verisignlabs.com>

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

