Domain Name System Operations
Internet-Draft
Intended status: Informational
Expires: September 28, 2012

J. Livingood
C. Griffiths
Comcast

March 27, 2012

Definition and Use of DNSSEC Negative Trust Anchors
draft-livingood-negative-trust-anchors-01

Abstract

DNS Security Extensions (DNSSEC) is now entering widespread deployment. However, domain signing tools and processes are not yet as mature and reliable as is the case for non-DNSSEC-related domain administration tools and processes. One potential technique to mitigate this is to use a Negative Trust Anchor, which is defined in this document.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction 3
2. Domain Validation Failures 3
3. End User Reaction 4
4. Switching to a Non-Validating Resolver is Not Recommended . . 5
5. Responsibility for Failures 5
6. Definition of a Negative Trust Anchor 6
7. Use of a Negative Trust Anchor 6
8. Managing Negative Trust Anchors 7
9. Comparison to Other DNS Misconfigurations 8
10. Other Considerations 8
  10.1. Security Considerations 8
  10.2. Privacy Considerations 8
  10.3. IANA Considerations 8
11. Contributors 8
12. Acknowledgements 9
13. References 9
  13.1. Normative References 9
  13.2. Informative References 10
Appendix A. Document Change Log 10
Appendix B. Open Issues 10
Authors' Addresses 10

1. Introduction

The Domain Name System (DNS), DNS Security Extensions (DNSSEC), and related operational practices are defined extensively [RFC1034] [RFC1035] [RFC4033] [RFC4034] [RFC4035] [RFC4398] [RFC4509] [RFC4641] [RFC5155].

DNSSEC has now entered widespread deployment. However, domain signing tools and processes are not yet as mature and reliable as is the case for non-DNSSEC-related domain administration tools and processes. As a result, operators of DNS recursive resolvers, such as Internet Service Providers (ISPs), occasionally observe domains incorrectly managing DNSSEC-related resource records. This mismanagement triggers DNSSEC validation failures, and then causes large numbers of end users to be unable to reach a domain. Many end users tend interpret this as a failure of their DNS servers, and may switch to a non-validating resolver or contact their ISP to complain, rather than seeing this as a failure on the part of the domain they wanted to reach.

In the short-term, one potential way to address this is for DNS operators to use a Negative Trust Anchor to temporarily disable DNSSEC validation for a specific misconfigured domain name. This immediately restore access for end users while that domain's administrators fix their misconfiguration. While DNS operators likely prefer not to use this tool, during the global transition to DNSSEC it seems some tool is needed to reduce the negative impact on such operators.

A Negative Trust Anchor should be considered a transitional and temporary tactic which is not particularly scalable and should not be used in the long-term. Over time, however, the use of Negative Trust Anchors will become less necessary as DNSSEC-related domain administration becomes more resilient.

2. Domain Validation Failures

A domain name can fail validation for two general reasons, a legitimate security failure such as due to an attack or compromise of some sort, or as a result of misconfiguration on the part of an domain administrator. As domains transition to DNSSEC the most likely reason for a validation failure will be due to misconfiguration. Thus, domain administrators should be sure to read [RFC4641] in full. They should also pay special attention to Section 4.2, pertaining to key rollovers, which appears to be the cause of many recent validation failures.

In one recent example [DNSSEC Validation Failure Analysis], a specific domain name failed to validate. An investigation revealed that the domain's administrators performed a Key Signing Key (KSK) rollover by (1) generating a new key and (2) signing the domain with the new key. However, they did not use a double-signing procedure for the KSK and a pre-publish procedure for the ZSK. Double-signing refers to signing a zone with two KSKs and then updating the parent zone with the new DS record so that both keys are valid at the same time. This meant that the domain name was signed with the new KSK, but it was not double-signed with the old KSK. So, the new key was used for signing the zone but the old key was not. As a result, the domain could not be trusted and returned an error when trying to reach the domain. Thus, the domain was in a situation where the DNSSEC chain of trust was broken because the Delegation Signer (DS) record pointed to the old KSK, which was no longer used for signing the zone. (A DS record provides a link in the chain of trust for DNSSEC from the parent zone to the child zone - in this case between TLD and domain name.)

3. End User Reaction

End users generally do not know what DNSSEC is, nor should they be expected to at the current time (especially absent widespread integration of DNSSEC indicators in end user software such as web browsers). As a result, end users may incorrectly interpret the failure to reach a domain due to DNSSEC-related misconfiguration as their ISP purposely blocking access to the domain or as a performance failure on the part of their ISP (especially of the ISP's DNS servers). End users may feel less satisfied with their ISP's service, which may make them more likely to switch to a competing ISP. They may also contact their ISP to complain, which of course will incur cost for their ISP. In addition, they may use online tools and sites to complain of this problem, such as via a blog, web forum, or social media site, which may lead to dissatisfaction on the part of other end users or general criticism of an ISP or operator of a DNS recursive resolver.

As end users publicize these failures, others may recommend they switch from security-aware DNS resolvers to resolvers not performing DNSSEC validation. This is a shame since the ISP or other DNS recursive resolver operator is actually doing exactly what they are supposed to do in failing to resolve a domain name, as this is the expected result when a domain can no longer be validated, protecting end users from a potential security threat.

4. Switching to a Non-Validating Resolver is Not Recommended

As noted in Section 3 some people may consider switching to an alternative, non-validating resolver themselves, or may recommend that others do so. But if a domain fails DNSSEC validation and is inaccessible, this could very well be due to a security-related issue. In order to be as safe and secure as possible, end users should not change to DNS servers that do not perform DNSSEC validation as a workaround, and people should not recommend that others do so either. Even if a website in a domain seems to look "normal" and valid, according to the DNSSEC protocol, that domain is not secure. Domains that fail DNSSEC for legitimate reasons may be in control of hackers or there could be other significant security issues with the domain.

Thus, switching to a non-validating resolver to restore access to a domain that fails DNSSEC validation is not a recommended practice, is bad advice to others, is potentially harmful to end user security, and is potentially harmful to DNSSEC adoption.

5. Responsibility for Failures

A domain administrator is solely and completely responsible for managing their domain name(s) and DNS resource records. This includes complete responsibility for the correctness of those resource records, the proper functioning of their DNS authoritative servers, and the correctness of DNS records linking their domain to a top-level domain (TLD) or other higher level domain. Even in cases where some error may be introduced by a third party, whether that is due to an authoritative server software vendor, software tools vendor, domain name registrar, or other organization, these are all parties that the domain administrator has selected and is responsible for managing successfully.

There are some cases where the domain administrator is different than the domain owner. In those cases, a domain owner has delegated operational responsibility to the domain administrator. So no matter whether a domain owner is also the domain administrator or not, the domain administrator is nevertheless operationally responsible for the proper configuration operation of the domain.

So in the case of a domain name failing to successfully validate, when this is due to a misconfiguration of the domain, that is the sole responsibility of the domain administrator.

Any assistance or mitigation responses undertaken by other parties to mitigate the misconfiguration of a domain name by a domain

administrator, especially operators of DNS recursive resolvers, are optional and at the pleasure of those parties.

6. Definition of a Negative Trust Anchor

Trust Anchors are defined in [RFC5914]. A trust anchor should be used by a validating caching resolver as a starting point for building the authentication chain for a signed DNS response. The inverse of this is a Negative Trust Anchor, which creates a stopping point for a caching resolver to end validation of the authentication chain. This Negative Trust Anchor can potentially be placed at any level within the chain of trust and would stop validation at that point in the chain.

7. Use of a Negative Trust Anchor

When a domain has been confirmed to fail DNSSEC validation due to a DNSSEC-related misconfiguration, an ISP or other DNS recursive resolver operator may in some cases use a Negative Trust Anchor for a domain or sub-domain. This instructs a DNS recursive resolver to temporarily NOT perform DNSSEC validation for a specific domain name. This immediately restores access to the domain for end users while the domain's administrator corrects the misconfiguration(s).

In the case of a validation failure due to misconfiguration of a TLD or popular domain name (such as a top 100 website), this could make content or services in the affected TLD or domain to be inaccessible for a large number of users. A Negative Trust Anchor can therefore be useful in the short-term when used on a targeted and time-limited basis. It does not and should not involve turning off validation more broadly, and helps during the transition to DNSSEC as organizations that are new to signing their domains are still maturing their DNSSEC operational practices, alleviating end user issues Section 3 and restoring end user access. However, use of a Negative Trust Anchor should not be automatic in any way, and must involve investigation by technical personnel trained in the operation of DNS servers. Such an investigation must confirm that a failure is due to misconfiguration, as a similar breakage could have occurred if an attacker gained access to a domain's authoritative servers and modified those records or had the domain pointed to their own rogue authoritative servers. Furthermore, a Negative Trust Anchor should be used only for a short duration, perhaps for a day or less.

Finally, a Negative Trust Anchor is used only in a specific domain or sub-domain and would not affect validation at other names up the authentication chain. For example, a Negative Trust Anchor for

zone1.example.com would affect only names within zone1.example.com, and validation would still be performed on example.com, .com, and the root ("."). In another example, a Negative Trust Anchor for example.com would affect only names within example.com, and validation would still be performed on .com, and the root (".")

```
     Root (.)              <======
```

```
                |                       ||
                |                       ||<======>+----+----+    DNSSEC
                |                       ||        |Recursive|   Validation
            TLD (com)             <=====||        |Resolver |   <=========>
                |                        +<------>+----------+
                |                        |                       DNS NTA
                |                        |                    (example.com)
        SUB TLD (example.com)     <------|                     <----------->
                |                        |
                |                        |
                |                        |
                (www.example.com  <-------
```

Figure 1: Negative Trust Anchor Diagram

8. Managing Negative Trust Anchors

This tool is unlikely to be and probably should not be used over the long-term since DNSSEC-related domain administration practices will naturally improve over time. In addition, however, continued and frequent use of Negative Trust Anchors is not scalable since it requires investigation by technical personnel and may involve manual processes, resulting in increased operational overhead (and therefore cost).

While Negative Trust Anchors have proven useful during the early stages of DNSSEC adoption, domain owners are ultimately responsible for managing and ensuring their DNS records are configured correctly Section 5.

Most current implementations of DNS validating resolvers currently follow [RFC4033] on defining the implementation of Trust Anchor as either using Delegation Signer (DS), Key Signing Key (KSK), or Zone Signing Key (ZSK). A Negative Trust Anchor should use domain name formatting that signifies where in a delegation a validation process should be stopped.

9. Comparison to Other DNS Misconfigurations

As noted in Section 5 domain administrators are ultimately responsible for managing and ensuring their DNS records are configured correctly. ISPs or other DNS recursive resolver operators cannot and should not correct misconfigured A, CNAME, MX, or other resource records of domains for which they are not authoritative. Expecting non-authoritative entities to protect domain administrators from any misconfiguration of resource records is therefore unrealistic and unreasonable, and in the long-term is harmful to the delegated design of the DNS and could lead to extensive operational instability and/or variation.

10. Other Considerations

10.1. Security Considerations

End to end DNSSEC validation will be disabled during the time that a Negative Trust Anchor is used. In addition, the Negative Trust Anchor may be in place after the point in time when the DNS misconfiguration that caused validation to break has been fixed. Thus, there may be a gap between when a domain has have been re-secured and when a Negative Trust Anchor is removed. In addition, a Negative Trust Anchor may be put in place by DNS recursive resolver operators without the knowledge of the authoritative domain administrator for a given domain name.

10.2. Privacy Considerations

There are no privacy considerations in this document.

10.3. IANA Considerations

There are no IANA considerations in this document.

11. Contributors

The following people made significant textual contributions to this document and/or played an important role in the development and evolution of this document:

- John Barnitz
- Tom Creighton
- Chris Ganster

12. Acknowledgements

The authors and contributors also wish to acknowledge the assistance of the following individuals or groups. Some of these people provided helpful and important guidance in the development of this document and/or in the development of the concepts covered in this document. Other people assisted by performing a detailed review of this document, and then providing feedback and constructive criticism for revisions to this document, or engaged in a healthy debate over the subject of the document. All of this was helpful and therefore the following individuals merit acknowledgement:

- Your Name Here!

13. References

13.1. Normative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, November 1987.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, March 2005.

[RFC4034] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Resource Records for the DNS Security Extensions", RFC 4034, March 2005.

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

[RFC4398] Josefsson, S., "Storing Certificates in the Domain Name System (DNS)", RFC 4398, March 2006.

[RFC4509] Hardaker, W., "Use of SHA-256 in DNSSEC Delegation Signer (DS) Resource Records (RRs)", RFC 4509, May 2006.

[RFC4641] Kolkman, O. and R. Gieben, "DNSSEC Operational Practices", RFC 4641, September 2006.

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

Existence", RFC 5155, March 2008.

[RFC5914] Housley, R., Ashmore, S., and C. Wallace, "Trust Anchor Format", RFC 5914, June 2010.

13.2. Informative References

[DNSSEC Validation Failure Analysis]
Barnitz, J., Creighton, T., Ganster, C., Griffiths, C., and J. Livingood, "Analysis of DNSSEC Validation Failure - NASA.GOV", Comcast , January 2012, <http://www.dnssec.comcast.net/DNSSEC_Validation_Failure_NASAGOV_20120118_FINAL.pdf>.

Appendix A. Document Change Log

[RFC Editor: This section is to be removed before publication]

-00: First version published as an individual draft.

-01: Fixed minor typos and grammatical nits. Closed all open editorial items.

Appendix B. Open Issues

[RFC Editor: This section is to be removed before publication]

- Decide whether to include intentionally insecure names in this (per Antoin Verschuren and Patrik Wallstrom)

Authors' Addresses

Jason Livingood
Comcast Cable Communications
One Comcast Center
1701 John F. Kennedy Boulevard

Philadelphia, PA 19103
US

Email: jason_livingood@cable.comcast.com
URI: http://www.comcast.com

Chris Griffiths
Comcast Cable Communications
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
US

Email: chris_griffiths@cable.comcast.com
URI: http://www.comcast.com