

Domain Name System Operations
Internet-Draft
Intended status: Informational
Expires: March 29, 2015

P. Ebersman
Comcast
C. Griffiths
Dyn
W. Kumari
Google
J. Livingood
Comcast
R. Weber
Nominum

September 25, 2014

Definition and Use of DNSSEC Negative Trust Anchors
draft-livingood-dnsop-negative-trust-anchors-00

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.

This document discusses Trust Anchors for DNSSEC and defines a Negative Trust Anchor, which is potentially useful during the transition to ubiquitous DNSSEC deployment. These are configured locally on a particular instance of a validating DNS recursive resolver and can shield end users of such a resolver from the DNSSEC-related authoritative name server operational errors that appear to be somewhat typical during the transition to ubiquitous DNSSEC deployment. Negative Trust Anchors are intended to be temporary, and should not be distributed by IANA or any other organization outside of the administrative boundary of the organization locally implementing a Negative Trust Anchor. Finally, Negative Trust Anchors pertain only to DNSSEC and not to Public Key Infrastructures (PKI) such as X.509.

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 March 29, 2015.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Definition of a Negative Trust Anchor	4
3.	Limited Time and Scope of Use	4
4.	Domain Validation Failures	5
5.	End User Reaction	5
6.	Switching to a Non-Validating Resolver is Not Recommended . .	6
7.	Responsibility for Failures	6
8.	Use of a Negative Trust Anchor	7
9.	Managing Negative Trust Anchors	8
10.	Removal of a Negative Trust Anchor	9
11.	Comparison to Other DNS Misconfigurations	9
12.	Intentionally Broken Domains	10
13.	Other Considerations	10
13.1.	Security Considerations	10
13.2.	Privacy Considerations	11
13.3.	IANA Considerations	11
14.	Acknowledgements	11
15.	References	12
15.1.	Normative References	12
15.2.	Informative References	13
Appendix A.	Document Change Log	13
Appendix B.	Open Issues	14
	Authors' Addresses	16

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](#)] [[RFC6781](#)] [[RFC5155](#)].

This document discusses Trust Anchors for DNSSEC and defines a Negative Trust Anchor, which is potentially useful during the transition to ubiquitous DNSSEC deployment. These are configured locally on a particular instance of a validating DNS recursive resolver and can shield end users of such a resolver from the DNSSEC-related authoritative name server operational errors that appear to be somewhat typical during the transition to ubiquitous DNSSEC deployment. Negative Trust Anchors are intended to be temporary, and should not be distributed by IANA or any other organization outside of the administrative boundary of the organization locally implementing a Negative Trust Anchor. Finally, Negative Trust Anchors pertain only to DNSSEC and not to Public Key Infrastructures (PKI) such as X.509. [REFERENCE NECESSARY?]

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 to 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 restores 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. 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.

3. Limited Time and Scope of Use

As noted in [Section 1](#), the use of Negative Trust Anchors should be temporary. These are key recommendations pertaining to this practice:

1. The general practice of using Negative Trust Anchors should be limited to the transition to widespread deployment of DNSSEC (including signing of domain names and validation in DNS recursive resolvers). Thus, the practice of using Negative Trust Anchors should not be permanent.
2. During this transition phase when Negative Trust Anchors may be useful, the use of a particular Negative Trust Anchor should be temporary and in most cases limited to no more than 1 day. Thus, the use of an individual Negative Trust Anchor should be strictly time limited and very short in duration.
3. So that the use of Negative Trust Anchors remains temporary and useful only during a transition to widespread DNSSEC deployment, the use and distribution of individual Negative Trust Anchors should not be centralized, beyond the borders of one organization's operational unit. Thus, no organization should endeavor to create and centrally distribute Negative Trust Anchors to other organizations as was the case with positive Trust Anchors prior to the signing of the root.
4. As noted in [Section 12](#), organizations that utilize Negative Trust Anchors should not add a Negative Trust Anchor for any intentionally broken domain.
5. As noted in [Section 8](#), 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.

4. 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 [\[RFC6781\]](#) 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.)

In addition, it is possible that some DNSSEC validation failures could arise due to differences in how different software developers interpret DNSSEC standards and/or how those developers choose to implement support for DNSSEC. For example, it is conceivable that some domain may be DNSSEC signed properly, and Unbound-based DNS recursive resolvers will validate the domain but those using BIND or Nominum's Vantio software may fail to validate a domain.

5. 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.

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

As noted in [Section 5](#) 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.

7. 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.

8. 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 as noted in [Section 5](#) 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.

Technical personnel should also confirm that the domain is not intentionally broken, such as for testing purposes as noted in [Section 12](#). 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. In addition, personnel should make a reasonable attempt to contact a domain for which a Negative Trust Anchor may be used, and preferably prior to implementing it.

Furthermore, a Negative Trust Anchor **MUST** only be used for a short duration, such as for a day or less. Implementors **SHOULD** set an end time and date associated with any Negative Trust Anchor. Implementors **SHOULD** in most cases limit the maximum duration to one day, meaning the Negative Trust Anchor will be removed or invalidated from the point of implementation, plus 86,400 seconds. However, there may be corner cases where a Negative Trust Anchor is needed for a longer period of time. Optimally this time and date is set in a DNS recursive resolver's configuration, though in the short-term this may also be achieved via other systems or supporting processes.

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 (".")



Figure 1: Negative Trust Anchor Diagram

9. 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 7](#).

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.

Different DNS recursive resolvers may have different configuration names for a Negative Trust Anchor. For example, Unbound calls their configuration "domain-insecure" [\[Unbound-Configuration\]](#)

[10.](#) Removal of a Negative Trust Anchor

As explored in [Section 13.1](#), if a Negative Trust Anchor is still in place after the point in time when the DNS misconfiguration that caused validation to break has been fixed, this could be problematic. It is therefore recommended that implementors should periodically or even continuously attempt to validate the domain in question, for the period of time that the Negative Trust Anchor is in place, until such validation is again successful. (Obviously a Negative Trust Anchor could be removed prior to validation succeeding again, alleviating an implementor of the need to continuing to test validation separate from their normal operations.)

Once validation is again successful, a Negative Trust Anchor should be removed as soon as is reasonably possible. Optimally this is automatic, though it may also be achieved via other systems or supporting processes.

[11.](#) Comparison to Other DNS Misconfigurations

As noted in [Section 7](#) 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.

12. Intentionally Broken Domains

Some domains, such as `dnssec-failed.org`, have been intentionally broken for testing purposes [[Measuring-DNSSEC-Validation-of-Website-Visitors](#)] [[Netalyzr](#)]. For example, `dnssec-failed.org` is a DNSSEC-signed domain that is broken. If an end user is querying a validating DNS recursive resolver, then this or other similarly intentionally broken domains should fail to resolve and should result in a SERVFAIL error. If such a domain resolved successfully, then it is a sign that the DNS recursive resolver is not fully validating.

Organizations that utilize Negative Trust Anchors should not add a Negative Trust Anchor for any intentionally broken domain.

Organizations operating an intentionally broken domain may wish to consider adding a TXT record for the domain to the effect of "This domain is purposely DNSSEC broken for testing purposes".

13. Other Considerations

13.1. 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.

End users of a DNS recursive resolver or other people may wonder why a domain that fails DNSSEC validation resolves with a supposedly validating resolver. As a result, implementors should consider transparently disclosing those Negative Trust Anchors which are currently in place or were in place in the past, such as on a website [[Disclosure-Example](#)]. This is particularly important since there is currently no special DNS query response code that could indicate to end users or applications that a Negative Trust Anchor is in place. Such disclosures should optimally include both the data and time that the Negative Trust Anchor was put in place and when it was removed.

13.2. Privacy Considerations

There are no privacy considerations in this document.

13.3. IANA Considerations

There are no IANA considerations in this document.

14. Acknowledgements

Several people made contributions of text to this document and/or played an important role in the development and evolution of this document. This in some cases included performing a detailed review of this document and then providing feedback and constructive criticism for future revisions, or engaging in a healthy debate over the subject of the document. All of this was helpful and therefore the following individuals merit acknowledgement:

- Joe Abley
- John Barnitz
- Tom Creighton
- Marco Davids
- Brian Dickson
- Patrik Falstrom
- Tony Finch
- Chris Ganster
- Olafur Gudmundsson
- Peter Hagopian
- Wes Hardaker
- Paul Hoffman
- Shane Kerr
- Murray Kucherawy
- Warren Kumari

- Rick Lamb
- Marc Lampo
- Ted Lemon
- Ed Lewis
- Antoin Verschuren
- Paul Vixie
- Patrik Wallstrom
- Nick Weaver
- Ralf Weber

15. References

15.1. 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.

- [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.
- [RFC6781] Kolkman, O., Mekking, W., and R. Gieben, "DNSSEC Operational Practices, Version 2", [RFC 6781](#), December 2012.

15.2. 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](http://www.dnssec.comcast.net/DNSSEC%20Validation%20Failure%20NASAGOV%2020120118%20FINAL.pdf)>.
- [Disclosure-Example]
Comcast, "faa.gov Failing DNSSEC Validation (Fixed)", Comcast , February 2013, <<http://dns.comcast.net/index.php/entry/faa-gov-failing-dnssec-validation-fixed>>.
- [Measuring-DNSSEC-Validation-of-Website-Visitors]
Mens, J., "Is my Web site being used via a DNSSEC-validator?", July 2012, <<http://jpmens.net/2012/07/30/is-my-web-site-being-used-via-dnssec-validator/>>.
- [Netalyzr]
Weaver, N., Kreibich, C., Nechaev, B., and V. Paxson, "Implications of Netalyzr's DNS Measurements", Securing and Trusting Internet Names, SATIN 2011 SATIN 2011, April 2011, <<http://conferences.npl.co.uk/satin/presentations/satin2011slides-Weaver.pdf>>.
- [Unbound-Configuration]
Wijngaards, W., "Unbound: How to Turn Off DNSSEC", June 2010, <http://unbound.net/documentation/howto_turnoff_dnssec.html>.

Appendix A. Document Change Log

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

Individual-00: First version published as an individual draft.

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

Individual-02: Simple date change to keep doc from expiring. Substantive updates planned.

Individual-03: Changes to address feedback from Paul Vixie, by adding a new section "Limited Time and Scope of Use". Changes to address issues raised by Antoin Verschuren and Patrik Wallstrom, by adding a new section "Intentionally Broken Domains" and added two related references. Added text to address the need for manual investigation, as suggested by Patrik Falstrom. Added a suggestion on notification as suggested by Marc Lampo. Made several additions and changes suggested by Ralf Weber, Wes Hardaker, Nick Weaver, Tony Finch, Shane Kerr, Joe Abley, Murray Kucherawy, Olafur Gudmundsson.

Individual-04: Moved the section defining a NTA forward, and added new text to the Abstract and Introduction per feedback from Paul Hoffman.

Individual-05: Incorporated feedback from the DNSOP WG list received on 2/17/13 and 2/18/13. This is likely the final version before the IETF 86 draft cutoff date. Updated references to [RFC6781](#) to [RFC6781](#), per March Davids.

Individual-06: Added more OPEN issues to continue tracking WG discussion. No changes in the main document - just expanded issue tracking.

Individual-07: Refresh document - needs revision and rework before IETF-91. Planning to add more contributors.

WG-00: Renamed at request of DNSOP co-chairs, added co-authors

[Appendix B](#). Open Issues

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

Determine whether [RFC 2119](#) language should be used or not when describing things like the duration of a NTA.

The DNSOP WG should discuss whether a 1 day limit is reasonable, whether a different time (more or less than 1 day, such as 1 hour or 1 week) should be specified, or whether no time should be specified (just a recommendation that it SHOULD generally be limited to X).

Olafur Gudmundsson has suggested that we may want to consider whether a non validatable RRSIG should be returned or not when a NTA is in

place. This was raised in the context of NLnet Labs' DNSSEC-Trigger, which apparently acts like forwarding stub-validator. He said, "The reason for this is if NTA strips signatures the stub-validator thinks it is under attack and may a) go into recursive mode to try to resolve the domain, getting to the right answer the long way. b) Give the wrong error "Missing signatures" instead of the real error. If all the validator does is not to set the AD bit for RRsets at and below the NTA, stub-resolvers (and cascading resolvers) should be happy."

Determine whether an informative reference to X.509 in the Introduction is necessary.

Is it desirable to say that NTAs should not be distributed across organizational boundaries?

Per Warren Kumari, add examples to appendix. "it would be very helpful to actually show how this is used, with e.g. an example in an Appendix, for -insert favorite resolver here-. The document contains a lot of really useful content about why you might use one, how to minimize damage, etc but (IMO) doesn't do a great job of explaining how to actually do so". Rick Lamb and Joe Abley also agreed on the need for this.

Per Rick Lamb, "it might be useful to have [section 2](#) "Definition .." make that clear for slow people like me - that the NTA is not an RR and is more of a configuration. Maybe simply replacing "placed" with "implemented" in [section 2](#)? "This NTA can potentially be -placed/implemented- at any level within the chain of trust"

Per Olafur Gudmundsson, address fact that ALL authoritative name servers must be working. "[section 10](#) you talk about possible early removal the NTA when validation succeeds but there may be instances where validation succeeds when using a sub-set of the authoritative servers thus NTA should only be removed if all servers are providing "good" signatures."

Per Olafur Gudmundsson, "Furthermore what to do if some names work but others do not, for example I remember a case where the records at the apex worked but all names below the apex were signed by a key not in the DNSKEY RRset, thus it is possible that either human or automated checks may assume there is no problem when there actually is one. What this is bringing to my mind is maybe you want a new section with guidelines on how to test for failures and in what cases failure justifies NTA and what tests MUST pass before preemptive removal of an NTA."

Per Olafur Gudmundsson, "Also should there be guidance that removal of NTA should include cleaning the caches of all RRsets below the name?"

Reference and text per Ed Lewis: One thing that seems to need repeating from time to time is this passage in [RFC 4033](#). ... In the final analysis, however, authenticating both DNS keys and data is a matter of local policy, which may extend or even override the protocol extensions defined in this document set. See [Section 5](#) for further discussion. A responsibility (one of many) of a caching server operator is to "protect the integrity of the cache." DNSSEC is just a tool to help accomplish that. It carries ancillary data that a local cache administrator may use to filter out undesired responses. DNSSEC is not an enforcement mechanism, it's a resource. When I see folks voice opinions that DNSSEC's recommended operation has to strictly followed, my gut reaction is that these folks have forgotten the purpose of all of our efforts. We don't secure protocols to make things work better. We don't operate the DNS because we like to run a well run machine. We don't run the Internet for the fun of it. (Some might enjoy running it, that's job satisfaction to some extent.) At the end of the day all that matters is that what is being done benefits society. We run the Internet to enrich society. We prefer a well run DNS because it saps less resources than a poorly run DNS. We prefer secure protocols so that people don't become victims (in some sense of the word). Make it work. Do what it takes to make it work. "Local policy" rules.

Per David Conrad: I'd suggest that in the BCP/RFC/whatever, in addition to recommending that NTAs be time capped and not written to permanent storage, it should also recommend NTAs be written as specifically as possible. (Should be obvious, but doesn't hurt to reiterate I suppose).

Per Ralf Weber: Informing the domain owner on the validation failure. There should be a section in the document that the operator deploying an NTA has to inform the domain owner of the problem. (JL note: would prefer to say operator SHOULD take reasonable steps to notify the domain owner, etc.)

Authors' Addresses

Paul Ebersman
Comcast
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
US

Email: paul_ebersman@cable.comcast.com
URI: <http://www.comcast.com>

Chris Griffiths
Dyn
150 Dow Street
Tower Two
Manchester, NH 03101
US

Email: cgriffiths@gmail.com
URI: <http://www.dyn.com>

Warren Kumari
Google
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: warren@kumari.net
URI: <http://www.google.com>

Jason Livingood
Comcast
One Comcast Center
1701 John F. Kennedy Boulevard
Philadelphia, PA 19103
US

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

Ralf Weber
Nominum

Email: Ralf.Weber@nominum.com
URI: <http://www.nominum.com>

