DNSSEC INDUSTRY COALITION

Symposium Report

Signed Root Deployment: Framing the Issues

Issues and Proposed Solutions

11–12 June 2009

Reston, Virginia, USA

Special note

It has been three and half years since a two-day Symposium of members of the DNSSEC Industry Coalition met to consider issues surrounding the signing of the DNS root. The Symposium's report, titled "Signed Root Deployment: Framing the Issues," was written in 2009 but remained an unpublished working draft until now.

We are now publishing that unfinished draft, which includes Coalition members' comments on specific sections of the draft. These comments show the state of the debate at a time when the root had not yet been signed and there were multiple concerns about whether doing so would have undesirable side effects. Some of the potential problems outlined in the draft report are in the past, but one issue in particular, key rollover, remains unresolved. Hence, we believe this paper still makes a contribution.

The three areas the Symposium focused on were:

Unintended consequences

Few of the problems listed in this section have been realized; there has been no real impact on name servers, and while there have been instances of firewalls and routers not handling DNSSEC queries and responses properly, these are problems with DNSSEC in general (or the devices in particular) rather than being connected with the signing of the root. These issues are unlikely to become concerns with any rollover of the root key.

Key distribution and use

Similarly, there appear to have been few problems with key distribution, as hardware and software vendors continue to successfully download the root key. What is less clear is whether hardware and software vendors, as well as others, will acquire a new root key in the event that it is changed.

Key rollover

There has been little movement on this subject in the past three years, and there are still no procedures for, nor has there been testing of, either regular or unscheduled root-key rollover.

This creates a potential vulnerability as DNSSEC is increasingly adopted; therefore, a major purpose of this report is to urge vigorous testing of key rollover in order to help the community evolve and "shake down" its products and systems.

Steve Crocker
December 2012

Table of contents

Symposium overview

In the interest of advancing the awareness and adoption of DNSSEC, a group of interested parties formed the DNSSEC Industry Coalition, a global network of registries, registrars, DNS software implementers and other industry experts whose mission is to work collaboratively to facilitate widespread adoption of Domain Name Security Extensions (DNSSEC) and streamline the implementations across domain name registries. Members work together to establish a consistent set of tools and applications, shared best practices, specifications, and shared nomenclature. The Coalition was founded by the Public Interest Registry, which operates .org, in August 2008, and includes a variety of Internet community stakeholders.

This Coalition coordinates industry experts and relevant stakeholders to spur interest in the deployment and adoption of DNSSEC. One of the most important events in the deployment process will be the signing of the top level of the DNS hierarchy, the root zone. There has been considerable attention to the signing process, particularly which organization will hold which of the two types of keys and what technology and procedures will be used to generate and use these keys, but there has been very little attention to the impact on the rest of the community. Towards that end, the Coalition organized a symposium to review the technical and operational issues that will affect DNS operators and users when the root in signed.

Symposium structure

The symposium was hosted at the Google office in Reston, Virginia on 11-12 June 2009. The structure of the event reflected its tactical focus: to discuss and identify potential and perceived issues with the Domain Name System (DNS) and DNSSEC deployment due to signing the DNS root zone. The Symposium was segmented into sections organized around the processes, technology and testing to achieve a seamless, successful deployment.

The discussion was organized from the premise the root will be signed and distributed by a combination of ICANN, VeriSign, the U.S. Department of Commerce and the other root operators. The process for this to happen—including policy, contractual concerns, software, hardware, etc.—were excluded from this discussion, as they are currently under consideration by those parties. Thus, the focus of the symposium was on the post-signing adoption and organized according to the following topics:

- *Avoiding Unintended Consequences:* Also known as the “Day 1” problem, this subject covers the potential, immediate impact on the DNS (both non-DNSSEC and DNSSEC) after the root is signed and signed delegations are in the root zone. It includes, software, hardware, and traffic concerns throughout the Internet.
- *Key Distribution and Use:* The potential processes for publishing the key to vendors, ISPs, and users, manual and automated efforts for finding the key, and notification of changes.
- *Key Rollover:* As a matter of stability and security, the key should be periodically updated. This session analyzed the circumstances and proposed frequency of key rollover and considered what preparations and testing are needed during the initial deployment process to support future key changes, including the possibility of unplanned or emergency changes.

With a tactical focus and aggressive agenda, the symposium attendance was kept to a small group of experts, reflecting numerous thought-leaders, advocates and stakeholders interested in the deployment of DNSSEC. In alphabetical order, Symposium attendees included:

Joe Abley

Ray Bellis

Melinda Clem

David Conrad

Steve Crocker

Whitfield Diffie

Michael Graff

Becky Granger

Chris Griffiths

Olafur Gudmundsson

Ashley Heineman

Peter Koch

Olaf Kolkman

Rick Lamb

Matt Larson

Ed Lewis

Bill Manning

Douglas Maughan

Ram Mohan

Thierry Moreau

Russ Mundy

Vivian Neou

Lauren Price

Tim Polk

Scott Rose

Jakob Schlyter

John Schnizlein

Shyam Seshadri

Suzanne Woolf

Kelvin Yiu

The attendees participated on personal title and under the Chatham House rule at varying levels—providing comments, making presentations, leading or moderating topics, and sharing relevant data. All topics and discussions reflect the personal, informed opinions of the attendees but not necessarily that of their respective organizations.

Symposium outcome

The participants engaged in an intensive two-day discussion, identified issues that warrant further examination and agreed that subsequent meetings will be held to advance these issues. As noted above, the purpose of the Symposium was not to restrict participation or involvement, rather to expedite discussion and achieve tangible results to present to the larger community. Accordingly, this report reflects the culmination of the Symposium's discussions. It follows the structural, topical outline, summarizing the presentations, comments and recommendations for each subject. The report concludes with the recommended next steps for advancing deployment based on the results of each subject area presented.

All questions and comments about this report should be directed to the DNSSEC Industry Coalition (*info@dnsseccoalition.org*).

Avoiding unintended consequences

Maintaining stability and maximizing the security benefits of signing the root requires an understanding of what could go wrong and what advance actions can be taken to mitigate problems. The first session of the Symposium focused on how to sign the root and launch DNSSEC with minimal surprise or harm to the Internet community.

The session was divided into three issues: (1) understanding the impact of the behavior of fielded software around the DNSSEC-okay (DO) bit and compliance with the current standard; (2) the immediate impact of the appearance of signed answers for queries from the root on firewalls and routers; and (3) anticipated traffic impact on nameservers. These issues are illustrative of the primary and immediate known effects on the Internet, and are not to be perceived as comprehensive. To add to this discussion, traffic data on the recently signed .org TLD was presented.

Steve Crocker 8/13/12 11:25 AM

Comment [1]: Not sure I understand these two sentences. Need to review. Suresh also commented: How do we identify the comprehensive list of issues?

Issue 1: DO-bit compliance

There was a problem for DNSSEC in that much of the installed base of DNS resolvers would either not understand DNS answers larger than 512 bytes (which many DNSSEC answers are) or the new record types used for DNSSEC, with possible negative side effects to older code. A decision was made that only resolvers that could affirm their ability to handle large answers containing DNSSEC data would be able to query for it. EDNS0 had been adopted some time before as a way for DNS resolvers to negotiate larger packet sizes and signal various other states. The mechanism adopted for signaling that a resolver could accept DNSSEC information was to define an option bit, referred to as DO, and set it in queries. This is a configuration option that can be turned on or off in compliant resolvers, although much of the installed base of DNS software sets DO=1 by default.

Today, on the “L” root, roughly 62% of queries have the DO bit set; roughly 2% with DO=1 also offer the server a response buffer size of 512 bytes, in violation of the requirement in RFC 3226 that the offered buffer size be at least 1220 bytes[1]. A positive answer could easily exceed 512 bytes, but the real impact here concerns negative answers (NXDOMAIN): for the root, an unsigned negative answer is 88 bytes, but a *signed* negative answer is 618 bytes. In cases where the offered buffer size is overrun, fallback to TCP is the most likely result, which suggests an increase to 2% in the number of queries to the root servers that will result in TCP connections being opened.

DNSSEC response sizes will likely exceed the limit set in RFC 3225.

Thus, with the signing of the root and the resulting requirement for larger data-transfer capacities creates two immediate problems:

(1) An immediate spike in various loads, which could have a negative impact on some root servers, especially those that do not employ anycast, and;

(2) Devices (e.g., NATs, firewalls and filtering routers) that are sensitive to EDNS0 responses, unusual RR types, DNS over TCP, or large packet sizes, may drop responses to root queries.

The impact on TLD servers is not fully known and hasn’t been highlighted as a problem on the signed top-level domains .se, .cz and .gov. There was some discussion of this, including the observation that the root generates a dramatically high proportion of negative answers under ordinary workloads and thus may well respond differently.

Steve Crocker 8/3/09 3:28 PM

Comment [2]: Suresh asked: What are the concerns?

Mitigation possibilities

Four possible mitigation solutions for the DO-bit problem were presented in this session and several discussed at length throughout other sessions in the symposium. The solutions include:

[1] OARC “Day in the life” study of the root server impact.

- *Amend validator implementation.* Under this scenario, there would be no large-scale change in the implementation scheme, only a modification of RFC 3226 to remove the buffer-size limit.
- *Amend nameserver implementations.* For a resolver, if buffer <1220, turn off DO; for an authoritative server, if buffer <1220 and DO=1, return an error.
- *Strengthen the servers.* Increase both hardware and bandwidth at the root and TLD servers to handle the increase in TCP traffic.
- *Sign a parallel root*. Rather than signing the existing root zone and serving data from the existing root servers, create a parallel set of root servers (potentially at the next sequential IP address for each participating root server) and create a parallel root infrastructure (not an alternative namespace) that allows resolver operators to opt into DNSSEC by simply changing the root hints file to point to the new IP addresses for the parallel set of root servers.

The general consensus of the attendees was to amend RFC 3226 to modify the buffer limitations as a reasonable start at solving the DO-bit problem. This session saw significant discussion of the concept of establishing a parallel infrastructure, and again during the closing session of defining next steps. Arguments for and against such a construct were presented with an action item generated regarding a fuller analysis and presentation to the broader Internet community in the future.

Issue 2: Firewall and router impacts

A summary of findings of potential DNSSEC impacts on SOHO and consumer-grade firewalls and routers was presented. A recent study[2] looked at 24 SOHO gateways and firewalls that work both as DNS proxies and as pure routers. The DNSSEC response structure looks very different from traditional DNS responses because it uses the larger EDNS0 packet format and includes RR types defined specifically for DNSSEC. An intermediate or middle-box device such as those discussed here needs to be able to recognize this data as legitimate DNS answers, since many such devices will drop data they don't understand. Potential breakage from DNSSEC for DNS clients that send queries and get responses through these devices depends on whether they default to compatible settings or can be configured to handle DNSSEC data properly.

A summary of the devices and their respective issues is as follows:

- *Stub resolver, not security-aware.* While these devices do not automatically experience problems from DNSSEC queries, they are prone to problems with large packet sizes and subsequently break when receiving those packets. Note that these types of stub resolvers will not automatically receive such packets.
- *Stub resolver, security-aware but not validating.* When we start sending queries with the DO bit set, the device requires EDNS0 handling, and must be able to handle large RRsets in the response. Some answers will already fail in these devices, but it is not because of DNSSEC; it's simply because they're too large. Breakage is possible due to mishandling of EDNS0, and minor breakage is possible due to mishandling of AD responses.
- *Fully validating stub resolver.* The only difference anticipated between these devices and a non-security-aware device is that where the validation has failed at the recursive resolver, the records come back instead of a 'SERVFAIL' error message. In these scenarios, problems are very likely due to large RRsets but less likely for packets <~1470 bytes; there are possible problems due to mishandling of EDNS0, and minor breakage is possible due to mishandling of CD/AD bits.
- *Fully recursive, security-aware resolver.* For devices running BIND or Unbound with no intervening firewall, the likelihood of problems is very low, as the queries are sent out directly to the authoritative nameservers without involving an intermediate cache.

Steve Crocker 8/3/09 2:12 PM
Comment [3]: I don't understand this "solution"

Steve Crocker 8/13/12 11:25 AM
Comment [4]: Suresh also commented that this needs more explanation. He says it's not clear what the limit refers to.

Suzanne Woolf 8/3/09 11:07 AM
Comment [5]: FWIW I'm not sure this is true. The limits in 3226 make reasonably good sense I thought, the problem is with a negotiation process that ends up not complying with it (min. bufsize is 1220 for DNSSEC, our problem is at 512).

Steve Crocker 8/13/12 11:25 AM
Comment [6]: Suresh says: I don't see how these issues are "Day 1" problems for a signed root, since in each case the problem will be visible *today* if we have DNSSEC-aware clients/resolvers sending packets as mentioned. A more interesting case, I think, is where you have a middle-box in front of an authoritative name server that contains signed zones. Today, if the middle-box does something funny, the impact is only felt for those validators have configured a trust anchor for the zone in question. When the root is signed (and when there is a changing of trust leading from the root to the signed zone n question), potentially *all* validators will haves accessing t this zone even if they haven't explicitly configured a TA for the zone.

Steve Crocker 8/13/12 11:25 AM
Comment [7]: Suresh asks: If RD=1 and CD=0, should the full response still be sent back to the stub? (I think the full response is sent today, but should this be the case?)

Suzanne Woolf 8/3/09 11:07 AM
Comment [8]: Sanity check: does a non-validating stub ever set AD?

Steve Crocker 8/13/12 11:25 AM
Comment [9]: Suresh comments: As in the typical MTU for the last mile? Crocker says, I think we need to add something about why 1470 is mentioned.

Suzanne Woolf 8/3/09 11:07 AM
Comment [10]: Sanity check this too. The whole set of cases here needs reviewed against Ray's paper and/or the sensibilities of an implementor about whether the spec really requires this behavior if we're going to assert it does.

[2] The study report can be found at http://public.icann.org/files/ssac/sac035.pdf

Were the root signed tomorrow, it is anticipated that the only devices that would experience immediate problems are those running a security-aware resolver that (because of DHCP) forwards through the proxy to the upstream recursive resolver, where the device can not handle EDNS0 or the data size. In short, a device breakdown will only happen if you send DO=1 queries to a router or to a SOHO gateway's proxy that is not properly set up to forward the incoming replies.

For non-aware devices, problems will be addressed by policy enforced at the recursive resolver. At the end-user level, because many SOHO users are not running BIND or similar software or sending root queries, problems will be limited and a device will only have a problem if it sends a DO query. There are also few anticipated problems at the user level until user devices themselves become security-aware.

Suzanne Woolf 8/3/09 11:07 AM
Comment [11]: Did we really determine what the policy should be? Or only that we need one to prevent TCP fallback and loss of responses?

The discussion concluded by listing other use cases that should be examined for failure problems, including Wi-Fi hotspots or hotel Wi-Fi networks that intercept DNS in the presence of a smart stub; instances where there is a validator in a browser and no clear path on Port 53; and on VPNs when DHCP provides the resolver's address.

Issue 3: Traffic impact on nameservers

The results of an informal test were presented to frame the topic of potential changes in traffic to root servers. The study analyzed impact on recursive servers at the moment the root gets signed. Existing recursive nameservers that already set the DO bit will experience signatures coming back from the root. The impact is a function dependent on the traffic patterns of the specific recursive name server and on the traffic patterns of its ISP.

When looking[3] at the least-specific label of a query stream, and assuming the TLD label is not cached so the query must go to the root, the study found that 99.95% of all queries lead to a referral and only 0.02% of these answers are not cached. When we look at the number of queries that a recursive nameserver sends to the root, these numbers translate into the following: for existing TLDs, 0.51% are generating traffic, i.e. this is the traffic stream that goes from the recursive name server to the root. For root, this is roughly 0.01% of relevant traffic; the other 99.49% of the traffic coming from the recursive nameserver is for names for which there is no delegation in the root zone.

Steve Crocker 8/3/09 3:45 PM
Comment [12]: I've edited this paragraph and hopefully it makes sense now.

It is the queries for nonexistent names that generate the problems, because these queries will be answered with a signed answer, and the answers that are returned are most likely not going to be cached (now or ever).

Suzanne Woolf 8/3/09 11:07 AM
Comment [13]: Review this point also: exactly what are we claiming is not cached?

Steve Crocker 8/13/12 11:25 AM
Comment [14]: Suresh also asked: why? Aren't they cached for the SAO-min time?

When you consider the packet sizes for all these queries and the increase in the traffic towards the root for these types of questions, then the increase in traffic because of DNSSEC (with respect to 1024-bit keys and for referrals) is a factor of 1.65. The increase in the amount of traffic to answer queries going to the root for nonexistent domains is estimated to be 5.12 times the current rate. If the response is signed by a 2048-bit key, combined traffic is anticipated to increase by a factor of 8.

The impact on the traffic between the stub resolvers and the large recursive nameservers they use is estimated by assuming that 10% of the queries by the stub resolver to the recursive nameserver are done with the DO set. The increase in traffic load between those is only 1.016 times the current rate, a negligible increase.

Steve Crocker 8/3/09 3:46 PM
Comment [15]: Suesh comments: I think henceforth we need to do all assessments using 100% also.

For a recursive nameserver that resolves 30,000 queries/second (typical for a large ISP), the amount of extra traffic is about 69 kilobytes/second. This is to be compared to an overall traffic load of around 10 megabytes/second.

[3] 1Reference: NLnet Labs 2009-002: Resolver Analysis for a Signed Root; Wouter Wijngaards and Olaf Kolkman, June 2009 http://www.nlnetlabs.nl/downloads/publications/rs_analysis_06.pdf

Thus, the effects on the recursive nameserver are noticeable, but the overall increase of traffic on the total query stream is anticipated to be relatively small.

Looking at these numbers from the root's perspective, the predicted combined impact on all resulting activity on the recursive nameserver is an eightfold increase. This number seems to be in conflict with earlier measurements as performed in RIPE 352[4].

Suzanne Woolf 8/3/09 11:07 AM

Comment [16]: I believe the 8x number applies to the total bw required for queries and responses once root responses to queries with the DO bit set include DNSSEC data signed with a 2048-bit key. I need that assertion sanity checked and then this needs to be rephrased.

Steve Crocker 8/13/12 11:25 AM

Comment [17]: I was also puzzled by this paragraph. Rereading it again, I think everything that needs to be said is covered in the preceding few paragraphs. I recommend deleting this paragraph completely. The reference to the RIPE report is important and needs to be somewhere. It can be moved into one of the earlier paragraphs. I read that report carefully, played with the numbers to make sure I understood them, and interacted with Olaf and Wouter.

There was an informal consensus that further work should be done, specifically on analyzing traffic impacts on the transaction between recursive resolvers and the root servers when signed answers from the root are commonplace; this was listed as an item requiring further discussion and planning.

Sample query data from .org

On 2 June 2009, .org became the first open generic TLD to successfully sign its zone with DNSSEC; to date, .org is the largest domain registry to implement DNSSEC. This deployment was done after careful planning by .org, The Public Interest Registry, and Afilias, the domain's registry operator, and with unanimous approval of the ICANN board. In its first two weeks, there were no substantive problems or issues reported by .org domain holders or the greater community.

The following chart was presented to the Symposium as empirical evidence of the increase in TCP query rates:

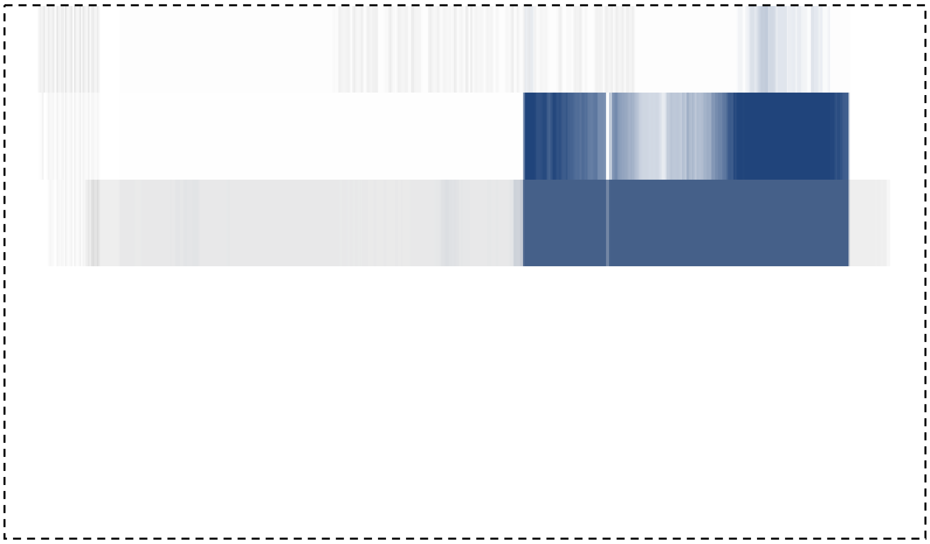

Figure 1: 10-minute sample of TCP rates before and after the signing of .org

The chart depicts 10-minute samples of query rates across the three days surrounding the event of signing .org. A large portion of the queries result in name error or NXDOMAINs. TCP traffic increased by 600 times and the causes of this are being investigated. Some decline is anticipated as negative caching was not enabled during this period, but by how much is unknown.

Steve Crocker 8/3/09 3:47 PM

Comment [18]: Suresh asks: how was this disabled?

Afilias demonstrated it is well-provisioned to handle the query traffic with no adverse effects. The .org example highlights the need for TLD operators to be over-provisioned as a matter of stability, especially those who are deploying NSEC3 and are employing strong encryption key sets.

[4] ftp://ftp.ripe.net/docs/ripe-352.pdf

Mitigation and problem resolution

With a number of potential issues defined, a group discussion on how to avoid the problems, prioritization, and next steps moved the group to the tactical realm[5]. A few general considerations were posited to frame the discussion: If it is assumed that perfection is not possible—something, somewhere will break—what is the acceptable failure rate? The issues were analyzed in two passes: do we understand failures well enough to proffer a solution for resolving a problem, and if so, what is the solution; if not, what needs be done to define the way forward?

Steve Crocker 8/3/09 3:48 PM
Comment [19]: Suresh asks: Acceptable failure rate of what?

The group believed minimizing failures was an important goal, which was best described as protecting the integrity and stability of the current environment. Many opined on how to define failure. The primary lines of discussion looked at where avoiding problems was possible and where we could focus efforts to "absorb" issues and potentially mitigate them throughout the Internet. It was assumed that there will be some problems because, even with a phased approach, there will be active participants looking for signed keys as well as those with security-aware devices who passively receive query streams their infrastructure or network cannot handle.

Steve Crocker 8/3/09 3:49 PM
Comment [20]: Suresh says: not sure what this means.

DNSSEC was designed to enable the gradual use of signed records by voluntary validators. The bug of deployed recursive resolvers already requesting signature data without adequate buffers may increase traffic much more sharply than intended. Careful monitoring of the actual impacts will be necessary. While there was some discussion of "turning back", this line of thought is highly impractical—once the root is signed, there is virtually no way to truly un-sign it. As such, the focus must remain on identifying potential issues and developing risk-mitigation strategies so we can continue to innovate and achieve the ultimate goal of adding great security to the DNS.

Steve Crocker 8/3/09 3:50 PM
Comment [21]: Suresh comments: possible, yes, But impractical unless things are truly broken.

Instant traffic impact

With respect to minimizing risk and shifting problem mitigation, the idea was raised to focus resources at the TLD level rather than requiring end users and ISPs to address device and software issues. We know there is an immediate problem with resolvers set to 512 bytes. Most resolvers have an EDNS heuristic that starts with an EDNS buffer of 4096, and if that sequence fails or times out, the resolver retries with 512 bytes—and if that sequence fails, it tries with EDNS0 off. If the resolver gets a TC response at any point during this query stream, it will fall back to querying via TCP. Because we know the NXDOMAIN on a signed root zone will likely exceed 512 bytes, this will trigger a TCP query.

Steve Crocker 8/3/09 2:53 PM
Comment [22]: I don't know what "at the TLD level" means in this context.

This brings the issue back to the devices with a greater likelihood of failure. As discussed (without disagreement), it's almost certain that SOHO-based CPE type devices have a very low probability of failure as they have DHCP-assigned addresses, with a resolver that's sitting on the outside of the CPE device. The potential problems are with the higher-end firewalls and NAT boxes where the resolver or the authoritative nameserver is on the inside; these are usually within enterprise networks that are managed using much more stringent security policies that will not allow fragmented responses or the use of DNS over TCP. We have already established that a signed TLD absorbs significantly more TCP queries, so this addition as a result of the 512-limitation is an immediate problem.

Steve Crocker 8/13/12 11:25 AM
Comment [23]: I'm not sure I understood the intent of this last part. Do enterprises prohibit DNS over TCP? Require? I may have garbled this even worse than it was. In any case, it wasn't clear and needs to be fixed.

A few questions remained from this discussion:

- Where is the tipping point: dealing with EDNS0 or the large packet size?
- What is the largest response today's root would generate with DNSSEC enabled?
- Could the expected problems be isolated to TLDs that have large key sets?
- Will the pain be felt on the authoritative side or pushed to the edge?

Steve Crocker 8/3/09 3:51 PM
Comment [24]: Suresh says: The scope of this comment is not clear.

Steve Crocker 8/13/12 11:25 AM
Comment [25]: I did not understand this. Suresh also commented: How does this help solve the problem?>

Steve Crocker 8/3/09 3:52 PM
Comment [26]: Suresh asks: Again, how does this knowledge help?

[5] This open discussion session was held after the presentations for both the unintended consequences and key distribution and use sessions took place.

The general consensus of the group was that this is a known, quantifiable issue that must be addressed immediately to avoid adverse consequences surrounding the planned signing of the root. While the root and .com/.net (as well as .org as demonstrated in the successful signing) are well-provisioned and can withstand the potential query load, there is no reason not to solve this problem and avoid the increased load on the root.

Parallel signed infrastructure

A detailed discussion of establishing a parallel signed infrastructure commenced as an alternative resolution scheme for this and other potential problems. Arguments were presented both for and against this methodology. The arguments for focused on creating an opt-in universe that would isolate activity to those participating—the consenting actors who have the foresight and ability to make the necessary changes, actively seek keys, and make appropriate policy changes in their networks. It was noted that we do not have a test environment, and no substantive data to understand or plan for consequences, only a notion that there could be negative root and end-user impacts.

The arguments against a parallel signed infrastructure included that this scheme doesn't avoid problems, only isolates them, and does not achieve the full intention of further securing the DNS; rather it only secures a portion of the DNS. It was noted that because there have been deployments at the TLD level, there is no empirical evidence to suggest why we can not assume similar activity (traffic, participants, problems) at the root level. The cost and burden on root operators was highlighted as an additional concern of establishing a parallel environment. Also it was noted that a parallel signed infrastructure raises several 'political' issues; it may not have the same quality or geographical spread as the current root-server infrastructure, causing DNSSEC to be a First World privilege and it may be confused with an alternative namespace. And finally, questions were raised about how the parallel infrastructure would persist, as it conceivably would need to, into perpetuity.

There was no resolution on this discussion; rather, the discussion was slated to continue in a larger forum (see the action item detailed in "Next Steps" below).

Topic summary: Avoiding unintended consequences

This session highlighted several important considerations in the deployment of DNSSEC, from RFC compliance issues to potential device failures to provisioning. There was no "show stopper" presented or discussed, that is, no potential consequence was deemed so dire that signing the root should be delayed.[6] Rather, the session demonstrated meaningful responses and plans to address the known negative possibilities. In summary, the following items were noted as requiring follow-up discussion, testing, and analysis:

- A full examination of the pros and cons of establishing a parallel, signed root-server infrastructure
- A process for removing the DO bit buffer limitation
- Full analysis of root-server traffic impacts under a statistically valid study to obtain a consistent look across all (root, TLD, name) servers

[6] Delayed with respect to the current plans (announced on 3 June 2009) under development between ICANN, NTIA and VeriSign for DNSSEC root deployment by the end of this year.

Key distribution and use

The second Symposium topic covered the practical activities of key distribution: how the root key will be distributed and used in resolvers. This session looked at the operational and policy aspects of notification and publication (both automated and manual), notification schemes, and distribution. Key distribution from the ISP perspective was discussed at length, from the views of an ISP, a browser operator, and software vendors.

Proposed publication schemes

The process of key distribution can be summarized into two publication actions: (1) initiate changes and distribute information in a timely manner, and (2) provide a means for those who operate resolvers that use the keys to obtain them and know with certainty that they are valid. There are a variety of manual and automated means for the notification and collection process, as well as corresponding problems to mitigate. Inherent in any discussion of publication is the notion of testing—creating a system to check whether the notification and distribution system works, and that it is accurate and easily understandable by man or machine. Further, we should treat the distribution of the key as a serious matter, one that can always be improved, possibly via better security that becomes available or through clearer notification mechanisms.

For any publication process to be successful, buy-in is essential. The publication scheme must be developed via public procedures by experts from the community. All procedures should be documented and key use should follow audited procedures. Stakeholders should have the ability choose to publish (trust) and distribute, or not. At all times, we should be working with the community to maximize trust.

Publication participants

If we assume that the issue of who holds the KSK does not have to be contentious as long as the Internet community at large can trust the KSK holder, we can also utilize existing trust models, e.g., a public key infrastructure, software distribution (e.g., Windows Update), and name space delegation (the DNS chain of delegations). Accordingly, a proposal was made to define a schema to identify the actors in the DNSSEC trust chain[7]:

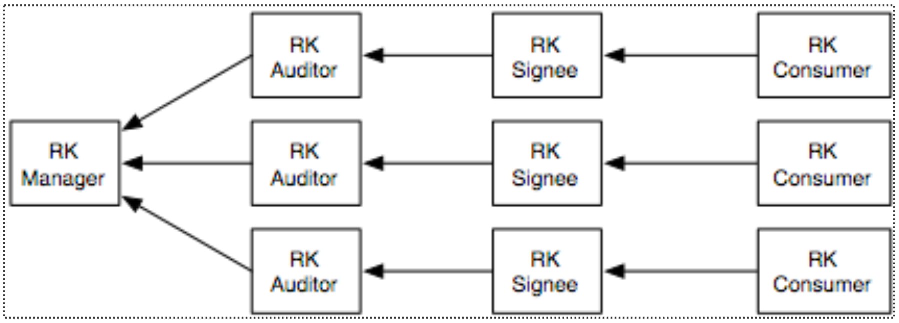

Figure 2: Proposed schema for defining participants and stakeholders

The actors and their roles are defined as follows:

[7] Images and definitions in this section were provided by Jakob Schlyter and Patrik Fältström of kirei.

- *RK Manager, the entity that manages the root key:* Operates according to its policy and practice statement; allows for auditors to be present upon key generation and to make audits of key usage.
- *RK Auditor, an entity executing an audit of the RK Manager on behalf of the RK Signee:* Checks that the RK Manager operates according to its policy and practice statement; several RK signees may use the same—or multiple—auditors.
- *RK Signee, an entity signing the RK, preferably someone trusted by many:* Initiates audits of the RK Manager; if the audit is satisfactory, the Signee signs and publishes the audited set of root keys as an RK Bundle; the signature is the proof that the Signee believes the RK Manager operates as described.
- *RK Consumer, someone using the RK as a DNSSEC trust anchor:* Gets the signed RK Bundle from one or many trusted RK Signees; if enough signatures are valid, installs the root keys as trust anchors.

We see existing companies and organizations fitting into this model per Figure 3:

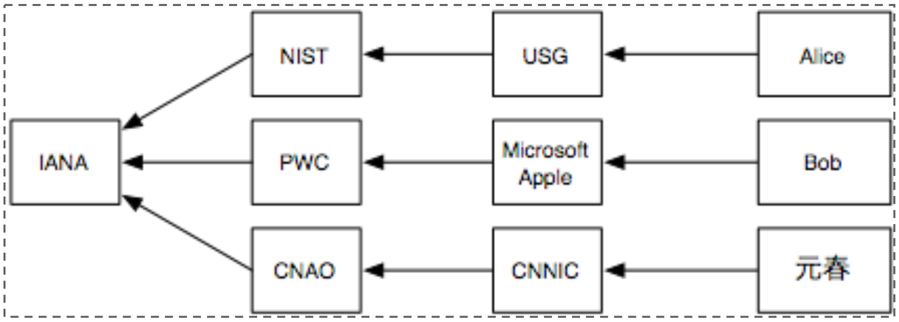

Figure 3: Participant schema showing existing organizations

This depicts a model for distributing bundles (which can contain fingerprints—DS records—of the current root keys or possibly public-key fingerprints of future keys not yet assigned) whereby the RK Manager generates keys, Auditors (operating on behalf of the RK Signees) audit the key-generation and -usage policies and procedures, the RK Signee signs the RK and distributes the result as a signed RK Bundle, and the RK Consumers can trust one or more RK Signees.

Publication processes

The notification of key availability can be obtained in numerous ways, as can the authentication and receipt by users. This process can be defined and tested today to be readily available upon the signing of the root. Available options for notification and validation today include:

- *DNS:* to distribute DNSKEY data
- *Web/FTP sites:* the ICANN and IANA websites, e.g., ftp.internic.net or ftp.rs.internic.net
- *Software updates:* root hints, pushed vendor updates
- *One or more online sources of key hashdata:* Regional Internet Registry websites, dnssec-deployment.org and dnssec.net, one or more of the network-operator specific forums such as NANOG, electronic editions of various news websites, vendor websites, and results from popular search engines

- *Use of print media where other means to verify key are not available:* announcements and ads in widely read, trusted publications
- *Use of some weighted K-of-N scheme in order to reconcile differences*

Given the existing format and proven methodology of sharing information via the root hints file, the question was raised as to whether this vehicle should be used for key notification. While this process is ubiquitous with respect to availability, its use is limited to a subset of the Internet community that understands it and knows how to use it today. Further, because of differences in the format of the key publication data, the root hints file is likely incompatible for pushing the appropriate key information.

To establish and test a publication process in an open and transparent manner, the following sequence was proposed:

1. *Define a human-friendly hash format, similar to s/key, for DNSSEC keys and get vendors to implement it.*
2. *Define a consistent way for the root zone to make its trust-anchor information available via online sources in some trustworthy manner. Identify the specific websites to be used and the entity that will vouch for the data.*
3. *Identify a reasonable number of online publication points and monthly print publications willing to carry hash information for current root keys.*
4. *Assess the key-change propagation delay for various sources of root key hash material and assign some basic weighting scheme based on their timeliness and reliability.*
5. *Define a mechanism that allows resolver programs to fetch and install an initial trust anchor collection from a set of online sources based on some (configurable) K-of-N scheme.*
6. *Get vendors of DNSSEC validator software to expose details of the trust anchors currently in use, to facilitate checking against out-of-band sources.*
7. *Identify the maximum effective key-rollover frequency for the root zone, taking into consideration constraints imposed by graceful key-rollover protocols and typical EDNS0 message sizes.*
8. *Test the root-zone publication process against one or more RFC 5011 compliant trust anchor management utilities. Evaluate the need for additional tools that may be required to detect emergency root key-rollover events.*
9. *Have the root commit to shipping more than one key the first time and commit to a rollover schedule for the first few rollovers.*

With respect to testing the publication process, the following parameters must be addressed: Does the scheduled rollover function correctly? What are resolvers' expectations for an unscheduled key rollover? How often, technically, can keys be changed? How long will it typically take for a change to propagate to different resolver deployments (i.e., where resolver deployments = ISP, enterprise, localhost, DNSSEC-capable applications)?

The publication schema session concluded with a call to consider the possibility of what happens if key information is published incorrectly, to further discuss the notion of standby keys, and to develop a formal plan based on the presentations described above.

ISP perspective: key receipt for resolvers

The task of distributing keys through ISPs and nameserver software providers was the next topic of conversation. This interactive session shared the experience of those with end-user contacts—either ISP or software customers. Three common concerns were identified: the ability to reach customers, a timeline for implementation and updates, and the balancing of customer needs.

This discussion focused on validation and trust. Existing relationships between end users and their ISPs and software vendors is based on an existing level of trust. Since this framework of trust is already in place, these vendors could continue to deploy or ready their services for DNSSEC. In this framework, the relevant questions are: How do validators (the vendors, in this construct) handle key acceptance? How does the validator originally obtain keys? How does it roll keys over and when or how does it determine that it needs to? What is the disposal process for old keys? The first of these questions is covered herein, the distribution issue is touched upon in a summary of RFC 5011, and the rollover discussion is presented in the next section of this report.

For each vendor, customer contact was a concern since they do not necessarily have an authoritative list of users of their services. Further, except for instances where there is a direct, current relationship, vendors do not know which customers have an infrastructure design or devices that could be directly impacted by DNSSEC. Thus, for customers who cannot be reached and are not getting regular updates, there is a potential breakdown. For some software customers, the root-hints file is an old yet still viable means of sending information, but there are concerns about potentially causing any breakdown of the existing systems by employing this method.

The timeline for deployment was another concern for vendors addressing customer needs; it could take as much as a year to have conversations with clients and others to develop an implementation plan. To address this concern, vendors will possibly have a staged approach. This phasing could include shipping subsequent software versions with DNSSEC disabled, making enabling it a conscious act on the part of their customers.

Discussion next settled around the roles and responsibilities relating to balancing customer needs. Given the more direct end-user relationship with customers, it was posited that the responsibility for enabling or facilitating DNSSEC deployment lies with DNS service providers, not software vendors. Because the software vendors do not control devices or keys, software vendors can only provide a toolset for management, not be authoritative. Thus, it is the ISP that is in the trust path to the end user and that must work in tandem with software vendors to ensure seamless transitions for their customers.

ISPs recognize the potential traffic impacts and the issues surrounding authoritative and recursive resolver sizing. Their testing will look into the impacts on recursive resolvers as well as on authoritative nameservers. Also, they are highly aware of data growth, the impacts that the size of signed data will have on networks, and the resulting bandwidth needs. The next logical step in preparing for root deployment involves testing validation on signed TLDs.

The relevant standard that addresses a model for validating new keys as you discover them in the DNSKEY RRset is RFC 5011. It defines a standardized process for tracking keys. One known concern is that it requires update probing every 15 days, and requires a 30-day period before keys become effective or ineffective—which creates rollover or updating problems. For example, if a device has shut down during a key push and reboots with stale information, that device may not be able to resolve and there will be no way to be online and fix the problem. The draft requires testing and potentially refinement, and may subsequently need to be republished within the IETF.

This subsection concluded by identifying points for follow-up, including the needs for further testing and review of RFC 5011, increased conversation between ISPs and software vendors, and a more detailed plan on distribution processes.

Lessons learned from Microsoft's certificate-authority (CA) process

As a point of comparison, Microsoft presented its current processes for updating browsers and operating software for the root CA program. Today there are 95 CAs in Windows, with a total of about 265 root certificates, and 30 of these are government CAs. There were several aspects of this process to be considered for the key distribution and rollover process, ranging from validity timeframes, authentication, notification and updates, and operational experiences.

There are several aspects of establishing a root CA that *may* be relevant to the DNSSEC model:

- The root CA must issue certificates to the general public.
- The CA must undergo audits (e.g., WebTrust, ETSI TS 102 042, or ISO 21188 audits, or for government CAs, self-asserted equivalent audits); these audits must be successfully conducted every 12 months in order for the CA to remain in the program.
- Root certificates must be valid for at least eight years at the time of their submission.
- CAs must submit root certificates over email.

From an operational perspective, the CA rollover process highlights several issues and lessons learned. Root CA rollover is a slow process and several factors can hinder updates, most frequently the fact that not all browsers or devices support dynamic root updates. Also, CAs usually try to get their root into as many platforms and devices as possible, but the old root is used until new root is ubiquitous, thereby causing a lag and providing no definitive means for CAs to ensure that the new root is ubiquitous.

Users also face update challenges; for example, they cannot always connect to the Internet, for all manner of reasons; enterprise users do not always have Internet access; or a Wi-Fi client might need a root CA before it is able to connect to the Internet. An additional enterprise challenge is that those customers who want to trust only some CAs have a difficult time ensuring the newest roots are included in their environments. Finally, government CAs operate according to local law; this presents conflicts for vendors that do not accept government CAs according to a consistent criteria.

Another consistent challenge is that software is not always current; in Microsoft's experience, individuals and enterprises are not actively or even semi-regularly installing updates. Only massive attacks and viruses inspire large numbers of users to perform updates, but these are often specific patches, not full system updates or others of the type that include CA rollover or updates. Most end users do not have to install a patch or take any manual action; there is an auto-update when Explorer or other services are launched, which updates the CA list on devices that have shipped. Enterprise administrators can override this process, potentially giving those users inaccurate information.

With respect to DNSSEC deployment at Microsoft, given the current shipping schedule and updates that are in process, there is no immediate opportunity for introducing DNSSEC technology into their software or operating system. At this time, it would take several months to initiate the process for integrating something into their systems. This is driven by a strict goal of stability: Whatever updates Microsoft ships for a live production system may not knowingly break anything.

Additional key distribution issues

With the issues surrounding key distribution outlined, the discussion turned to identifying problems to be solved, tabled or discussed further. The first order of business was to reach consensus that the root-hints distribution mechanisms are not sufficient for the trust anchor as they are not a ubiquitous solution for all users. With that consensus, and the understanding that no progress could really be made on the software end until the key distribution process is defined, the discussion considered alternatives.

Combining two proven methods was offered as a possible solution, where IANA publishes the KSK on an SSL-secured website with a CA that everybody acknowledges has been created using a CA that is in every browser. By using PGP to sign it with a key with signatures, there would be two very strong cryptographic assertions as to the key's authenticity. With such a scheme, the publication focus turned to securing the secret portion of the KSK and ensuring that it would not be compromised. For further levels of security, it was suggested that the key should be generated differently from the hints file.

Steve Crocker 8/13/12 11:25 AM

Comment [27]: Why are we wandering into this topic? I would delete this sentence entirely.

With this potential structure, the final set of questions revolved around the integrity of the key-generation process, or auditing the generation and rollover process. One issue was that third-party audits of the company performing key generation would constitute a regular (potentially annual) expense to someone, likely ICANN.

The proposal for developing a different but similarly situated root hints type distribution process was generally accepted as a reasonable, staged launch point. It would be more than adequately secure from a learning standpoint and sufficient for vendors like Microsoft to begin building processes.

Key rollover

The final early deployment aspect of signing the root is the concept of key rollover: defining the circumstances under which the root key would change on a regular schedule as well as under abnormal circumstances, such as in the case of a compromise. This session looked at the processes that are defined or under development, evaluated comparable processes for thoughts on best practices and intervals, and discussed the testing process. This discussion also included the ability to have standbys or backup keys to use in either rollover scenario. Key rollover was generally accepted to be an essential piece of DNSSEC deployment.

How the rollover process could work

The key-rollover process features very few actions, but involves very many considerations to ensure its integrity and transparency. The first considerations involve how long a key must have been published before it's universally trusted, and how it may be removed. Considering the proper intervals is important as the number of keys available will affect how long a key can remain published. Because RFC 5011[8] allows for spare keys, we must also consider how many keys are active at once. Each of these parameters affects the size of the set. The key-rollover process is essential, though, as we know cryptologic algorithms tend to become weaker over time and stronger resources for penetrating them will become available.

Figure 4 was offered[9] as a visual representation of the key lifecycle (where blue indicates "not visible," green indicates "in DNSKEY set" and red indicates "revoked in DNSKEY set").

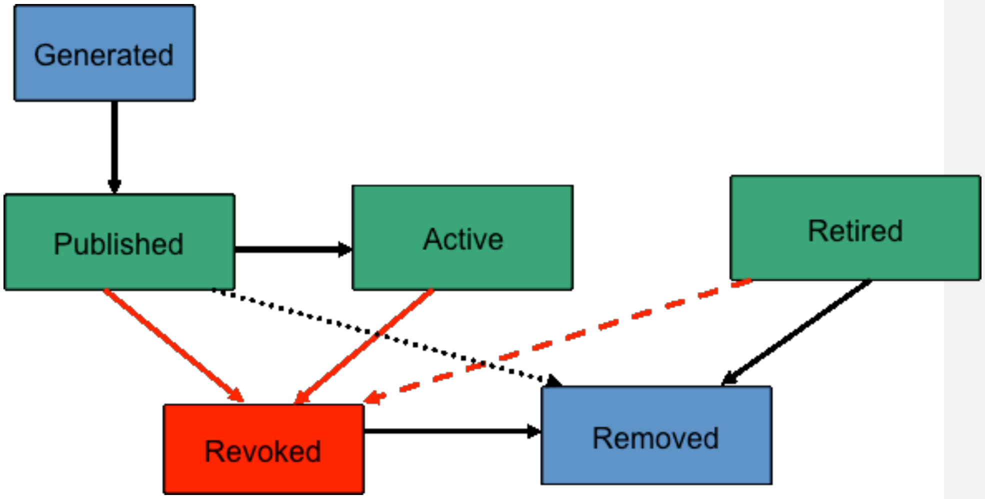

Figure 4: State of a key

Keys in Active and Revoke states generate signatures, and the keys in Retired state may have signatures in the system. This is predicated on the concept that it is inadvisable to remove a key before its last signature expires unless there is a problem with that key; keys in a Revoked state need to sign the DNSKEY set to prove the key has been revoked. Further, only keys in Active and Revoked states may generate signatures. With respect to removal, a key is revoked by having the key sign the DNSKEY set with a Revoke bit set on the key. The revoke operation is a "super" action, a definitive

[8] RFC 5011; http://www.rfc-editor.org/rfc/rfc5011.txt

[9] Per presentation "Root Key Rollover Process" by Olafur Gudmundsson, Shinkuro.

and immediate move, as the key can never be brought back from Revoked state. While keys in a Retired state may have old signatures present (and strictly speaking, keys in a Removed state can also have signatures), that should be viewed as an error as those keys are no longer valid.

When we discuss the root keys, we must also consider space versus time tradeoffs. The speed of key rollover in the DNS protocol depends on:

- *Trust-anchor propagation delays to the furthest corner*
- *How many keys and signatures are present per algorithm and role*
  - *One KEY + two RRsigs (1+2)*
  - *Two KEYs + one RRsig (2+1)*
  - *Two KEYs + two RRsig (2+2)*
  - *One KEY + one RRsig does not work for rollover (1+1) but can be used in between*
- *Number of RRsigs affects all answers*
- *Number of KEYs only affects DNSKEY answers*
- *KSK only: How quickly trust anchors are updated*

It was proposed that only the KSK keys sign the DNSKEY RRset as there is no value in having the zone signing keys do it. This is to reduce the size of the DNSKEY answer. Further, no one should configure ZSKs as trust anchors. As the KSK can only be trusted based on external factors, there are several quality considerations to address: If you have your key on the hardware, what happens when the hardware reaches the end of lifetime; can you still access the key, can you transfer it to new hardware? What about the distribution timing; can new keys be configured before the old ones are retired? What is the process for boxes that are offline (i.e., being shipped) during a key rollover? Some vendors need a longer lead time than others to make updates to reflect the key rollover; is a simple policy publication notice/notification sufficient? Each of these questions frames the key rollover process.

Several factors were discussed relative to key-rollover policy, several of which require testing, review and/or policy statements. These include:

- *What are the key's contents?*
- *How frequently should the root-zone key be rolled—each month, each year, each decade?*
- *How long can a key remain in a published state? An active state? A revoked state?*
- *What happens when multiple algorithms are in use and a key is revoked?*

Three rollover process were compared: the RFC 5011 process, the manual process (e.g., print and online publications, checking sites), and software updates. A recommendation was made to support RFC 5011, in which a new KSK can be phased in two months after publication and the old KSK retired soon thereafter and removed one signature lifetime later. RFC 5011 works with all other distribution mechanisms, and is the most ubiquitous of the choices presented. However, there are issues yet to be addressed: (1) RFC 5011 requires all validators to run a periodic checking process, and not everyone will build this into their architecture; (2) it does not address a device that is not turned on during the rollover; and (3) it may not provide trust history.

This discussion closed with a review of the issues concerning rollover failure and others yet to be addressed. Consensus must be reached on how fast the KSK can be rolled, how long that should take, the relationship between failures and key-rollover speed, and how frequently the KSK should be rolled. Each of these issues carries operational and adoption consequences. Finally, gauging a rollover's success requires that some criteria be established. This could be a mix of end-user and traffic-related metrics including the number of complaints to operators, blog/email list complaints, adoption on new sites, and changes in DO-bit traffic. These measures of success must also factor into the testing and reviewing process.

This topic closed with a more detailed discussion of RFC 5011 and a call for further testing and review.

Testing the rollover process

With the assumptions that the root will be signed and that regular key-rolling is a prudent security measure, a test plan was offered for scheduled transitions of the KSK. While the details have not been defined, the assumption is that the KSK will be rolled every one to five years. While manual updates or software-pushed updates are a possibility, a proposal for an automated plan was presented in RFC 5011.

The challenges for a rollover process were defined as:

- *If key rollover is infrequent, how can we be sure it will work?*
- *How many resolvers are involved?* (We don't know exactly, only that it is a huge number.)
- *Will it work? Who will get the new key smoothly? What will happen to the others?*

Steve Crocker 8/13/12 11:25 AM

Comment [28]: Suresh comments: We need to do this multiple times in the first couple of years to make sure that the process works properly. It's probably better to find (and fix) real operational issues early on instead of running into them later on when we really want to do a rollover.

A process for testing and planning is needed to proceed; it was suggested that such activities be completed before a full deployment—planned well ahead of time, publicized, involving multiple resolver operators—and that data be measured and then gathered into a report. A practical proposal was offered: to avoid any delays to the root signing, if we focus testing on the signed example TLDs used in the IDN process (as opposed to existing signed TLDs, e.g., Sweden, the Czech Republic), we can develop a smooth testing process and move to the real TLDs knowing we won't jeopardize stability. This phased process expedites the rollover testing without risking the current environment and would entail signing some of these example TLDs and using them for testing.

If we are to proceed with this test by doing multiple key generation and publication, and serving a signed root with example TLDs, this would not preclude continuing testing multiple rollovers at a fairly fast pace as well as after real data is in those TLDs. One dilemma is with the language in RFC 5011 today where a single rollover takes two months—a time too long for stress testing. Given this and previous discussions, it seems reasonable to change the hold-down time limitation for the testing process. In addition to this test environment, it is also reasonable to establish a testbed installation, so that someone may test the rollover on a new device or software before a product or service is released. This could be established in the same way as the OT&E environment that registrars use for accreditation.

This testing process is for a single algorithm. While there is no formal discussion for testing algorithm rollover, it seems reasonable to conclude that an algorithm rollover would be very similar to the KS rollover.

Emergency key rollover

We must plan for an emergency key rollover because there is always a chance—even though it may be years away—of compromise or because some underlying system has failed. The emergency rollover is a plan for an unplanned event and must be executed as quickly as possibly given the dire circumstances of its need. This reinforces a requirement noted in discussions of both the planned key-rollover process and the publication schemes: for quick, ubiquitous, trusted notification. It also raises a somewhat philosophical question of how a user can trust the new key publication during a period in which trust has been broken.

As the current automated rollover method, RFC 5011 was again discussed at length as the vehicle on which to model an emergency rollover. By having largely identical processes for planned and emergency rollovers, we can minimize the possibility for mistakes. The first part of the emergency process that would need to be different from RFC 5011 would involve notification—how the Internet community is made aware that a new key is not only available, but essential to obtain.

One aspect of RFC 5011's design makes it viable as the basis of an emergency plan, but also highlights what must be modified in creating an effective emergency process. RFC 5011 allows for a standby key that you can hold and push at your discretion—basically, a "hot" spare key that is held offline that could be used in the event of a root-key compromise. In theory, we could enable multiple back-up keys; however, a concern was raised about the security of this design, especially the possibility of compromise of the spare(s). More importantly, policy modifications would be required as RFC 5011 mandates a 30-day wait for a new key to be obtained—an unworkable delay in the event of emergency rollover.

There was consensus that this was an area to be worked on in tandem with other action items related to the planned key-rollover process.

Topic summary: culmination of the key-rollover discussion

Steve Crocker 8/13/12 11:25 AM

Comment [29]: Added "Topic Summary" to match previous sections. Check for consistency throughout.

Several items were noted as needing further review prior to root deployment. These involved both testing and policies and processes to be drafted, vetted and finalized to cover all aspects of a key life-cycle. Open items include:

- *Refining the key-rollover process:*
  - Rehearse both the regular/policy roll and whatever backup ("emergency") facility is established
  - Test RFC 5011
  - Draft a transparency policy that is then simplified into an end-user document
- *Establishing a key lifespan.* Proposals ranged from one to 10 years.
- *Coordinating publication efforts to ensure that widespread notification is available for emergency rollover.*

Special attention must be paid at the ISP and end-user levels. For TLD operators, these processes can be executed with relatively short scripts in a timely manner; but those parties managing networks and directly servicing end users could have greater difficulty pushing out updates and initiating key rollover procedures, especially emergency procedures.

There was general consensus that we must approach key rollover as a continual process to be refined and enhanced, rather than as a static policy that does not consider the creation of stronger cryptographic algorithms or new software and new devices.

Conclusion of the Symposium

The Symposium closed with two important activities: summarizing what we learned and what would be done, and secondly, evaluating the event's level of success.

Findings and action items

Throughout each session we captured lists of action items, further discussion points and general considerations that were essential for advancing DNSSEC deployment; many of these have been presented in their respective sections and called out accordingly. The closing section focused on developing actionable tasks and deliverables so that the community can make substantive progress.

Toward that end, Figure 5 summarizes the activities the Symposium agreed were essential and immediate next steps. It captures the activities required to move items from concept to final policy and/or technology, and defines ownership of or participation in that process.

	PLANNING	EVALUATION	DRAFTING
1	**Parallel System Report**—Full description and pros/cons—Key distribution and rollover plan. Owner: David Conrad/ICANN, IETF	**Root Production Plan**—Resolve EDNS-512 limitations and TCP traffic implications. Owner: IETF, NOGs	**Recommendation Report** summarizing comments and testing
2	**Key Rollover Testbed**—Plan for development of a testbed. Owner: Olafur Gudmundsson, Shinkuro	**RFC 5011 Testing**—Testing plan and parameters, publish results and recommendations. Owner: Shinkuro	If needed, **revised RFC 5011**
3	**Key Distribution Notification plan**—Definition of roles—Publication schemes for planned and emergency rollover. Owner: NTIA, ICANN, VeriSign	Appropriate community review and test	Formal publication of **Notification Plan**
4	**Key Rollover plan**—Definition of roles and procedures—Part of existing efforts by NTIA, ICANN, VeriSign and other community members	Comment and review by relevant community including NIST, ENISA, cryptological community, etc.	**Final Rollover Plan** and **appropriate end-user documentation**
5	**Secure Delegation Provisioning plan**. Owner: IANA	**Test interface to the ITAR** Owner: IANA	**Final Delegation Policy**

Figure 5: Current activities toward advancing the deployment of DNSSEC

A general consideration is to be held constant throughout each phase of this work and future work on deployment of DNSSEC: Always aim to *protect the current environment*. Equally important is the idea that all items noted will culminate in public reports that will be made available for comment; this effort should remain an open and transparent process that allows for community review.

It is important to note that some of these activities are currently under way, specifically the efforts toward implementing DNSSEC at the root. As noted in the Department of Commerce press announcement of 3 June 2009, NTIA, ICANN and VeriSign are actively proceeding with plans to implement DNSSEC at the root this year in a manner that maintains the stability and security of the DNS, through a process that includes regular consultation with the technical community.

These planning processes will consider efforts by this Symposium and other relevant parties; for example, NTIA and NIST are currently drafting high-level technical requirements that will be made available for review by the technical community in the coming weeks. Additionally, ICANN and VeriSign are developing draft testing and implementation plans, which will also be made available for appropriate technical community input.

With respect to completion of the items in Figure 5, each task is on a unique schedule, but generally targeted for completion within the next six to eight months. A status meeting, likely a conference call or brief meeting, will be held in the next six months to gauge the status of each deliverable, and this would occur prior to a subsequent, second symposium.

Evaluation of the Symposium

Completing the action items and articulating clearly defined next steps brought the substantive portion of the meeting to a successful close. Polling of attendees yielded agreement that the Symposium was a success and that subsequent meetings should take place. Group discussion noted areas for improvement, which items should stay constant in subsequent meetings, and logistical matters. Consistent with its coordinating role, ICANN volunteered to host the second official meeting. This next session will be open to a larger forum (at a time within the next nine months and at a location to be determined) and will advance the topics addressed herein.

Final remarks

The DNSSEC Industry coalition formally thanks all participants for taking the time to attend and actively participate in this inaugural Symposium. The knowledge, energy and enthusiasm they brought to the event generated thoughtful, active and productive discussion and a clear and definitive plan to advance DNSSEC deployment.

We also wish to thank Google, specifically its staff at the Reston office, for providing a well-equipped meeting facility, technical support and refreshments throughout the event.

Finally, a very special thank you to the planning committee for their numerous efforts that made the Symposium possible: Steve Crocker, Olaf Kolkman, Lauren Price and Suzanne Woolf.

Glossary of abbreviations

Auth Key	A public key that a security-aware resolver has verified and can therefore use to authenticate data
CA	Certificate authority
DNSSEC	DNS Security Extensions
DS	Delegation signer, a cryptographic hash of the key-signing key
IANA	Internet Assigned Numbers Authority
ICANN	Internet Corporation for Assigned Names and Numbers
IETF	Internet Engineering Task Force
ISP	Internet Service Provider
KSK	Key-signing key: An authentication key that corresponds to a private key used to sign one or more other authentication keys for a given zone
NIST	National Institute of Standards and Technology (U.S.)
NOG	Network operator group
NTIA	National Telecommunications & Information Administration (part of the U.S. Dept. of Commerce)
NSEC3	DNSSEC Hashed Authenticated Denial of Existence
RRset	Resource Record set
RRSig	Resource Record signature
“Signed”	A zone whose RRsets are signed and that contains properly constructed DNSKEY, Resource Record Signature (RRSIG), Next Secure (NSEC), and (optionally) DS records
SOHO	Small-office/Home-office
ZSK	Zone-signing key: an authentication key that corresponds to a private key used to sign a zone