Robert C. Hutchinson
Dynamic Ventures

11/24/2008

Fiona Alexander
Associate Administrator
Office of National Telecommunications and Information
Administration
U.S. Department of Commerce
1401 Constitution Avenue, N.W.   Room 4701
Washington D.C. 20230

NTIA Seeks Public Comment Regarding the Deployment of DNSSEC:

I am writing on behalf of myself, an engineer with more than twenty years of experience implementing software and hardware solutions utilizing Internet technology and as participant of ICANN for the past three years.

General Questions Concerning Signing of the Root Zone

- ***Should DNSSEC be implemented at the root zone level? Why or why not? What is a viable time frame for implementation at the root zone level?***

  DNSSEC should be implemented at the root level, as soon as possible. DNSSEC "signing the root" makes a bold statement that the Internet is evolving into a reliable and trustworthy global network capable of hosting all communications. The current DNS system is vulnerable to various spoofing and cache poisoning attacks. While efforts continue to shore-up current DNS with quick-fixes, DNS will become truly reliable when DNSSEC is fully deployed. Signing the root provides a common trust anchor for existing domain names and emerging Internet domain name services.

  A common trust anchor is much easier for TLDs to implement and track changes, rather than the complex environment that will occur if multiple signed roots exist[RFC 2826, "IAB Technical Comment on the Unique DNS Root"]. It's much easier to a TLD to "roll" its key-signing key (KSK) if the root is signed: in that case, the TLD only needs to communicate the new KSK to the root. If the root is not signed, the TLD's KSK is configured as a trust anchor in millions of resolvers and there's the much larger problem of getting all those resolvers to update their local copy of the TLD's key. RFC 5011 automates this process of "trust anchor rollover", by allowing the zone owner to signal a KSK rollover and induce everyone who's configured the KSK as a trust anchor to change their local copy, as well. However, RFC 5011 has not yet been widely deployed (there aren't even any implementations yet, to my knowledge). So it would be complicated at best and risky at worst for a TLD to sign its zone without the root signed.

I recommend signing the root as soon as possible. If a "signed root" test facility is established in early 2009, then it should be possible to sign the root in late 2009 or early 2010.

- ***What are the risks and/or benefits of implementing DNSSEC at the root zone level?***

Risks of signing the root include:

- Getting "signing the root" wrong on the first try – and causing an embarrassing public failure. This risk must be mitigated with testing and proper validation of the operations prior to going live with DNSSEC.
- Cryptographic "blanking" large areas of the Internet to users [the Internet just disappears, because the resolvers version of the zone key becomes out-of-sync with the current zone key].
- Making DNS more vulnerable to amplification DDOS attack.
- Making a public failure of the implementation of DNSSEC. We [the Internet technical standards bodies only get one chance at deployment, so if we screw it up and destroy public confidence, it will be very hard to try again.
- Signing the root may cause many different "alternative" signed roots to emerge because some governments/zones are suspicious of the cryptographic control exerted by a single-signed-root. It is important to communicate that signing the root enables only enables users of the domain name system to verify, via digital signature, that the address of the domain name they have requested was not tampered with by a hacker. The political issue with signing the root revolves around who holds and controls the root zone KSK. No single organization should control the root zone KSK. It's too important for a single organization. The M-of-N technique is a well-developed cryptographic technique to split authority over a key among multiple parties. It's important that those N parties have a common interest in a healthy Internet, and that there be multiple parties. It is conceivable that we start signing the root zone with one set of "N" parties (from M-of-N)

authorizing use of the KSK, and switch to a different set of "N" at a later date. In other words, we should not hold up signing the root while we decide on the perfect set of "N" organizations, because M and N may change down the road.

Benefits of signing the root include:

- Signing the root will "legitimize" DNSSEC. Signing the root will signify to all Internet participants that DNSSEC is technology that must be rolled into their products, and that their customers will demand.
- It is the first step toward full DNSSEC deployment and an end to DNS cache-poisoning attacks.
- Increased trust in DNS. Because of this increased trust, DNS will be used to store and serve additional data that would not otherwise be put there today because of the current lack of security in DNS.
- The issue above regarding TLD KSKs: it's much easier for a TLD to manage its KSK rollovers if the root is signed. Therefore, signing the root may promote adoption of DNSSEC by TLDs.

- ***Is additional testing necessary to assure that deployment of DNSSEC at the root will not adversely impact the security and stability of the DNS? If so, what type of operational testing should be required, and under what conditions and parameters should such testing occur? What entities (e.g., root server operators, registrars, registries, TLD operators, ISPs, end users) should be involved in such testing?***

  InterOp style compatibility testing is needed in the following areas to proof the "operational" capability of DNSSEC:

  - Crypto-invalidations handling in caching servers
  - The DNSSEC testfest /planning/coordination/reporting should last 2009-2011 and be funded by ICANN/IANA.
  - Testing of routers handling DNSSEC packets. ICANN SSAC published a mid-2008 survey of router testing which indicates that there are significant problems with a majority of today's fielded routers. Unfortunately, the bulk of low-end routers that will need either firmware upgrades or software upgrades are owned by small businesses and home offices – which are not well suited to performing the upgrades. ISPs will be reluctant to turn-on DNSSEC if doing so breaks even a small percentage of their customers service. We've already seen a bad example in Sweden when an entire town went off the air because its ISP turned on DNSSEC validation and everyone had the same SOHO router that didn't understand DNSSEC-signed responses and dropped them.
  - While we can never do a complete test of every scenario and have

a 100% guarantee that signing the root won't be harmful, that doesn't excuse us from thoroughly testing DNSSEC from root to end user. I'd like to see a widespread test facility that is "opt in" and aggressively promoted. This would be official sponsored, use the real root zone (but signed), be hosted by the existing root operators and have a definite ending date. The idea would be to encourage as much use of this signed root as possible to find any problems. There have been DNSSEC root test facilities in the past, but none has received enough use to be a benchmark to use to declare that we know enough about what would happen to sign the root.

- A special case for testing is rolling the root Key Signing Key [KSK]. Some advocate that the root zone KSK rollover should be tested, while others believe that this is not necessary for signing the root and starting DNSSEC operations.
  The two schools of thought are:
  - **Yes:** We will eventually roll the root zone KSK, so we might as well test it sooner rather than later, and we should get administrators used to the idea that it does change. We don't want people thinking that the root zone KSK is a "set and forget" because it's not. If the root zone KSK changes and an administrator doesn't change the corresponding root zone trust anchor, his resolver will stop working.
  - **No:** The tools for rolling the root zone KSK are not yet mature (i.e., RFC 5011 has not been widely deployed). We should pick a long-lived (i.e., large) KSK and expect it to live for potentially several years (5-10). After the root is signed, we wait until there is sufficient deployment of RFC 5011 before deciding to roll the key. That way, most resolvers will update their root zone trust anchor automatically. There will still be some who will have to update manually, and some people will miss the rollover, but the chances of reaching more resolvers will be better down the road when KSK rollover is automated[and distributed in subsequent BIND versions], as opposed to now, when it is completely manual. This is another argument for **not** testing KSK rollover now: a KSK rollover now will require entirely manual re-configuration by resolver operators and does not reflect what will really happen operationally in the future when RFC 5011 is deployed and KSK rollover becomes automated for those who have deployed it.

- ***How would the different entities (e.g., root operators, registrars, registries, registrants, ISPs, software vendors, end users) be affected by deployment of DNSSEC at the root level? Are these different entities prepared for DNSSEC at the root zone level and /or are each considering deployment in their respective zones?***

- Root operators: must be running DNSSEC-capable server software (all already are), must be prepared for the increase in bandwidth due to larger responses. Prepared: probably.
- Registrars: must update web sites and databases to hold DNSSEC key material (registrant will submit keys to their registrar, which will then submit keys to the registry to be converted to DS records and signed, which establishes the chain of trust between the registry parent zone and the registrant child zone). For most registrars (certainly the ICANN-accredited ones), this means implementing the EPP extensions for DNSSEC described in RFC 4310. Prepared: no. Supporting DNSSEC may prove to be a competitive advantage.
- Registries: must implement RFC 4310, must update their systems and databases to support storing new DNSSEC information (primarily key material from registrants), must run DNSSEC-capable authoritative name servers, and must have sufficient bandwidth to handle larger responses, must develop DNSSEC policies and procedures (e.g., key rollover), must implement sufficient security to protect their zone's ZSK and KSK. Prepared: a few TLDs already have signed their zone, most are in planning or have announced they will.
- Registrants: must be able to sign their zone and serve their signed zone on DNSSEC-capable name servers. In reality, most registrants do not host their own zones, so they will need to find DNSSEC-capable zone hosting providers, or buy a "DNSSEC in a box" solution. I expect that DNS hosting providers will respond to the market's demand by offering DNSSEC support first as an option/extra charge, and eventually as a check-off item.
- ISPs: must turn on DNSSEC validation in their recursive name servers. (Enterprises, the other large group that operates recursive name servers, must do this as well.) Must track changes to KSKs they configure as trust anchors. This process will be manual at first until wider deployment of RFC 5011 for automated trust anchor rollover. Must be prepared for greater bandwidth utilization and more resources used on recursive name servers (DNSSEC validation is bandwidth-intensive and computationally expensive). Must be prepared to deal with end-user complaints when DNSSEC validation fails on domain names they query. This may be a significant issue for ISP help desks. Prepared: a small minority – Comcast has stated that they will begin deployment in 2009. All ISPs and enterprises will need to upgrade to DNSSEC-capable recursive name server software.
- Software vendors: Need to add DNSSEC support to their products. Widespread support in DNS server software, almost no support at the operating system and application level.
- End users: There's nothing for them to do, at first. Their ISP or enterprise may enable DNSSEC validation, which will protect them from cache poisoning attacks by dropping bad answers, but at the

expense of cryptic error messages in browsers and end user applications. This situation won't be remedied unless operating system vendors and application providers extend their software so that the results of DNSSEC validation can be communicated all the way to the end-user.

- ***What are the estimated costs that various entities may incur to implement DNSSEC? In particular, what are the estimated costs for those entities that would be involved in deployment of DNSSEC at the root zone level?***

Cost of DNSSEC:

- Labor cost to roll-out DNSSEC world-wide in root/authoritative name servers == ~0.5M servers in the system * 4 hrs/update * 2 updates / 2000hrs per year == 2,000 person years of ANS server updates == ~ $100M
- Labor cost to roll-out DNSSEC world-wide in caching name servers == ~5M DNS servers in the system * 4 hrs/update * 2 updates / 2000hrs per year == 20,000 person years of server updates == ~ $1B
- Software upgrade costs?
- Upgrades or replacement of existing access network routers?
- Additional bandwidth to serve the signed resource records?
- Hardware costs == 5.5M * $1500 average equipment cost == ~$8.5B in DNS ANS/Resolver upgrades and replacement for adding signature validation.

Signing the root:

- An important cost to consider is generating, protecting and storing the root zone KSK. No matter what authoritative control [M-of-N or some other] is in use, someone has to hold the KSK. This requires a Hardware Security Module (HSM), multiple physical layers of security, isolated equipment, lots of policy and procedure development and documentation, and expertise. The value of the root zone KSK corresponds to the value of a root key for a certificate authority. Certificate Authorities go to great lengths to protect their key material, because if its security is breached – the validity/value of the issued certificates becomes zero. The importance and expense of good security surrounding KSK storage should not be underestimated. In the case of the root zone KSK, these normally private ceremonies will become complicated by the desire to have the whole key ceremony videotaped and published on the Internet, as part of the transparency, trust and publicity of the event.
  Budget to build, secure and maintain the key signing facilities 24/7/365 could easily exceed $1M/year.

  - The cost to create the new root zone signing procedures – a one-time cost ??
  - The cost to operate the root zone procedures for each update of the root zone file ??

- ***The Department recognizes that the six process flow models discussed in the appendix may not represent all of the possibilities available. The Department invites comment on these process flow models as well as whether other process flow model(s) may exist that would implement deployment of DNSSEC at the root zone more efficiently or effectively.***

- ***Of the six processes flow models or others not presented, which provides the greatest benefits with the fewest risks for signing the root and why? Specifically, how should key management (public and private key sets) be distributed and why? What other factors related to key management (e.g., key roll over, security, key signing) need to be considered and how best should they be approached?***

Proposed Process Flows:

I prefer the compartmentalization of Proposed Process Flow 6; where KSK key generation and handling-control functions are separated into an independent Root Key Operator, overseen by a the M of N and the root zone ZSK is generated and applied by the Root Zone maintainer. This process flow leverages the process flow in use today, and does not rearrange root zone assembly or responsibilities.

DNSSEC root [and all zones] zone security is controlled by Zone Signing Key, and the Key Signing Key. Typical epoch for the Zone Signing Key is monthly and could theoretically be shortened to daily or lengthened to a year. The epoch of a zone Key Signing Key is intended to be one to many years.

The case of the root is very special, because changing the root KSK means creating a whole new chain of trust for the DNS zone hierarchy. Therefore, the generation and distribution of a new root KSK should be a ceremonial event, and should be done only after careful consideration of the cryptographic and logistic impact on the DNS system. Consequently, the best way to oversee the generation and distribution of the root KSK is to separate that function into a unique organization, a Root Key Operator, which is chartered with the limited scope of determining if the root KSK needs to be changed, and of so, generating and communicating the new root KSK.

The Root Zone Maintainer would then be responsible for rolling the root

ZSK on a pre-determined schedule [monthly] and requesting the RKO sign the new root ZSK with the root KSK.

This process flow represents the lowest risk to the integrity of the root zone assembly process, because it adds one additional step in the root assembly process, the signing of the Zone File, which is largely mechanical, and leaves in place the existing root zone assembly processes which have served the Internet community well for the last fifteen years.

- ***We invite comment with respect to what technical capabilities and facilities or other attributes are necessary to be a Root Key Operator.***
  Because the root KSK only needs to be renewed once every two to ten years, ICANN should be encouraged to outsource the root KSK generation and storage to a company with the facilities in place to perform this function. The Root Key Operator should be primarily responsible for initiating the root KSK generation process, verification that the new root KSK is cryptographically correct, and signing new root ZSKs. It would be a huge waste of resources and organizational talent to attempt to duplicate the facilities and procedural expertise of a commercial certificate authority – just to generate one root key set every five or so years.

- ***What specific security considerations for key handling need to be taken into account? What are the best practices, if any, for secure key handling?***
  The RKO must securely sign a new root ZSK with the root KSK. Procedures for secure communications between the RKO and the RZM, for signing the new root ZSK, can be arranged with a combination of Internet secure communications channels like IPSec/SSL and augmented with secondary authentication over email and telephone.

- ***Should a multi-signature technique, as represented in the M of N approach discussed in the appendix, be utilized in implementation of DNSSEC at the root zone level? Why or why not? If so, would additional testing of the technique be required in advance of implementation?***
  The M of N quorum structure proposed is a widely accepted technique for regulating key generation decisions and is enforceable by currently available commercial hardware security modules.

  As stated in the proposals, selecting the individuals or organizations which comprise M and N is a separate political problem. I personally prefer, the N members of the RKO should individually be able to demonstrate knowledge of the DNS system and DNSSEC cryptography, plus collectively be representative of the global Internet community. ICANN, in its capacity as the current holder of the IANA functions contract, should be chartered with the responsibility of defining the N membership rules, and selecting the members.

Thank You for your NOI and interest in improving the Internet,

Robert C. Hutchinson

Internet Product Architect – Dynamic Ventures 2004-2008

HP Server Product Architect – 1997-2003

HP Home Products Architect - IEEE 802.14/DOCSIS 1995-1997

HP OpenView Network Management Product Architect 1993-1995

IETF Desktop Management Task Force – charter member 1994