

From: Jim Reid <jim@rfc1035.com>
To: <DNSSEC@ntia.doc.gov>
Date: Mon, Nov 24, 2008 12:42 PM
Subject: Response to NTIA NoI on DNSSEC-signing the DNS root

I write these remarks in a purely personal capacity. They do not represent the views of any other organisation I may be associated with or any affiliations I have. I speak as someone who has almost 20 years of experience with the Domain Name System, first as an administrator and later in a number of more specialised roles including participation in the development of the DNSSEC protocol.

I welcome the NTIA consultation and encourage all parties who have a role in the co-ordination of the DNS root to work towards a prompt and generally accepted introduction of a signed root zone. This is essential to a wider uptake of Secure DNS (DNSSEC) and therefore the securing of a critical component of Internet infrastructure.

My response is in two parts. First there are some general comments about the framework and attributes surrounding a signed root that I would like to see. I believe these are generally supported by the Internet community. This is followed by answers to the questionnaire which NTIA/DoC published in the Federal Register.

Secure DNS, DNSSEC is the only viable way to protect the DNS by guaranteeing the authenticity and integrity of DNS responses. It therefore follows that deployment of DNSSEC is crucial for securing this fundamental component of critical Internet infrastructure. The lack of progress towards a signed root is one of the major reasons why DNSSEC deployment has stalled to date. While it is essential for the root to be signed, this should not be done in haste. On the other hand it cannot be unduly delayed either. The technical and operational aspects of DNSSEC are ready. Signing the root is now largely a procedural and political matter, though there are minor operational details that need to be worked out. These will depend on which of the proposed solutions, if any, is adopted.

The most important considerations for signing the root are: (a) not to destabilise the co-ordination of the root system; (b) maintain and enhance the stability and security of the DNS; (c) produce a solution that has the widest possible acceptance and trust. This requires a process that is seen to be open, fair and transparent so that there is confidence in all aspects of a signed root: security, reliability, stability, trust, integrity and so on.

Although DNSSEC is not about control, there may be concerns about that from some quarters. It would be advisable for NTIA to take account of these perceptions, even if they are not expressed through this consultation. From a geopolitical perspective it would be advantageous to use some form of shared key signing key so that no one entity has ultimate authority over that key. It is also essential that whoever holds the root keys (or parts of these keys) can be shown to be trustworthy and to act in the best interests of the Internet as a

whole. This is likely to be the way to ensure the signed root is globally accepted.

It is my opinion that the zone signing key should be held by the organisation which generates the root zone file. However it is IANA which has the established relationship with TLD operators and Sponsoring Organisations. It would therefore be best if those TLD operators engage with IANA whenever they deploy DNSSEC. Those existing channels can be extended to accommodate that more easily than introducing another entity or data flow into the co-ordination of the root zone. Adding a new process or new entities will add complexity and create extra data paths that need to be authenticated and validated. This would provide more opportunities for error.

If IANA is to handle keying material from TLDs, this would imply that they should hold the root zone signing key and generate the root zone. That would mean a minor but acceptable deviation to current practice where Verisign generate and distribute the zone file. As a non-profit and neutral organisation operating the fundamental Internet registry, IANA is by far the obvious choice for generating a signed root zone file anyway.

While I encourage the root to be signed promptly, I appreciate that some procedural and technical details may need fine-tuning in the light of actual operational experience. Therefore, a degree of flexibility in the root signing system will be necessary. Whatever processes and entities are chosen today may need to be reviewed at some point in the future. It may be helpful to propose an interim solution/approach to signing the DNS root, deploy it and then have a further consultation before any lasting decisions are made. This could also help to achieve consensus from those who may have reservations about DNSSEC or the current arrangements for co-ordinating changes to the DNS root.

Questions on DNSSEC Deployment Generally

In terms of addressing cache poisoning and similar attacks on the DNS, are there alternatives to DNSSEC that should be considered prior to or in conjunction with consideration of signing the root?

Not really. Minor tweaks to the DNS protocol may help to minimise the impact of DNS spoofing attacks. However these cannot eliminate them because of design weaknesses in the original DNS protocol. The IETF has conducted an exhaustive analysis of various approaches to securing the DNS over many years. Its bottom-up consensus-driven processes have arrived at DNSSEC as the only viable option for securing the DNS by providing the means to guarantee the integrity and authenticity of DNS responses.

What are the advantages and/or disadvantages of DNSSEC relative to other possible security measures that may be available?

There are no other viable security measures that can be generally deployed to protect the DNS. Although there are some disadvantages in deploying DNSSEC such as cost and complexity, these are greatly outweighed by the benefit of having DNS traffic authenticated and validated for those who wish to do so. This is an inevitable security trade-off: better security usually costs more. In addition to securing critical Internet infrastructure -- a laudable goal in itself -- deployment of DNSSEC will be an enabler for new services and security applications because these would be able to benefit from validated and authenticated data from the DNS.

What factors impede widespread deployment of DNSSEC?

One of the most obvious factors is the absence of a unique and generally accepted trust anchor: the signed root. Others include the somewhat primitive tools for signing and debugging available today. The lack of widespread DNSSEC expertise is also a concern. To some extent these considerations underline a chicken-and-egg problem. Lack of uptake of DNSSEC has not encouraged the development of better tools and increased training or awareness. This in turn has dissuaded operators and the wider Internet industry such as hardware and software vendors from investing in DNSSEC deployment. This vicious circle could be broken if the root was signed. That would demonstrate to the Internet community and the broader public that DNSSEC was now ready for general use instead of the current perception that it is largely the preserve of a small group of engineers and technically sophisticated operators.

Wider deployment of Secure DNS will also require new application programmer interfaces (APIs) so that software developers can exploit the validation and authentication features that DNSSEC offers. These should emerge as soon as it becomes clear that the root zone will get signed, indicating that DNSSEC is clearly moving out of the laboratory and into the mainstream.

What additional steps are required to facilitate broader DNSSEC deployment and use? What end user education may be required to ensure that end users possess the ability to utilize and benefit from DNSSEC?

Improved tools for DNS administrators would help DNSSEC deployment. More training and outreach would also help. Although there have been many workshops and initiatives to raise awareness of DNSSEC, these have been hampered by the lack of deployment. Signing the root would provide the essential foundation to build from.

It would also be desirable for procurement guidelines to be produced and possibly some sort of independent DNSSEC conformance certification developed so that vendors could demonstrate their products and services were DNSSEC ready. Ideally, this could be comparable to the government and business initiatives that nurtured POSIX and TCP/IP compliance in systems procurements since the 1980s.

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?

Emphatically, yes. Signing the root is long overdue. The DNSSEC protocol works and has been well tested. A few TLDs and operators have been using it in production for some time, years in the case of .se. It is now time for DNSSEC to be seen to move away from testbeds into general production use. Signing the root is the first step in that process.

From a purely technical perspective, the root could be signed today. However there are procedural and contractual issues that need to be resolved, particularly with respect to integration of DNSSEC key management into the co-ordination and administration of the root zone. These could and should be completed in less than a calendar year. It should be possible to have a globally accepted signed DNS root by the end of 2009. I urge NTIA and the others involved to meet that challenge.

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

The risks of not implementing DNSSEC at the root are quite serious. Those who wish to use DNSSEC will be forced to continue to adopt ad-hoc and somewhat brittle solutions to acquire and maintain up to date trust anchors and install these in their DNS setups. This was confirmed by the statement from the RIPE community to ICANN following the Tallinn RIPE meeting in May 2007. These ad-hoc measures are likely to delay the wider use of DNSSEC and discourage others, particularly TLD operators, from deploying it. In the worst case, this could fragment DNSSEC into islands of trust that cannot communicate reliably or easily transition to a globally recognised signed root.

A generally accepted signed root will be the cornerstone of DNSSEC deployment. This underpins everything. Without it, DNSSEC will not flourish. Signing the root would be the foundation for this technology to advance and get widely used.

Is additional testing necessary to assure that deployment of DNSSEC at the root will not adversely impact the security and stability of the DNS?

No. There has been plenty of testing to date and there is no need for any more. Some tests and testbeds have concentrated on protocol aspects and interoperability. Others have focused on operational issues. Most of these tests have been open to anyone to participate and their results have been freely published. There may well be a need for an independent assessment of the processes and procedures for whatever scheme is chosen for signing the root. This should not and must not be confused with a perception that further testing of DNSSEC at the root is necessary.

One area of concern is the potential for an emergency rollover of the

root zone's key-signing key. In my opinion, this is significant but not likely to be operationally damaging. Those who do not use DNSSEC (or the signed root as their DNSSEC trust anchor) will not be affected. Those who do rely on the signed root as their trust anchor will have an obvious interest in using the correct KSK, just as competent DNS administrators update their name server configurations whenever one of the root name servers gets renumbered. A change of root KSK should not impose much of an administrative burden on those who choose to use that key as their trust anchor. They should in any event appreciate that the root KSK will need to be changed from time to time and factor that into their procedures for rolling out DNSSEC and managing their DNS infrastructure.

How would implementation of DNSSEC at the root zone impact DNSSEC deployment throughout the DNS hierarchy?

This would be the catalyst to get DNSSEC deployed. Although a small number of top-level domains (TLDs) have either signed their zones or announced plans to do so, many have not. The wait-and-see approach has been understandable given the lack of progress to date towards a signed root. Signing the root will mean the waiting is over and these TLDs should become more pro-active. That in turn will encourage others in the domain name business such as registrars, service providers, DNS vendors, application developers, domain name holders and end users to take their initial steps towards the deployment of DNSSEC. Signing the root will provide a strong justification for those who have not yet invested in DNSSEC to make that decision.

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?

Adoption of DNSSEC is inherently an opt-in process. Those who do not wish to participate will not be affected by DNSSEC. Those who do will deploy the technology when it makes sense for them to do so: i.e. when there are business justifications, tools and systems have been installed, appropriate training and customer service arrangements are in place, etc, etc. Some of the above actors are ready to deploy. Others have already done so. Signing the root will encourage others to follow. They may not be ready today. However if the root is not signed, they may well never be ready to use DNSSEC or make any effort to become ready.

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?

It is hard to answer that question in the general case because DNS costs are usually hidden in an organisation. Even so, the most likely cost will be a small increase in administrative overheads. These would be minimised by having a globally recognised unique trust anchor

instead of managing a variety of ad-hoc solutions. In addition, there will be administrative costs for signing zones and managing their keys. These should not be a significant burden, though better tools would help. DNSSEC validation will increase the load on resolving servers. This is likely to mean adding more servers or upgrading them: something that should be readily accommodated in the routine equipment procurement life-cycle.

The costs for those involved in the management and co-ordination of the root zone are unlikely to be significant. The signed root will be larger than an unsigned one and there will be an increase in the volume of data that the root server operators have to handle. These should not have a noticeable impact on root server load or traffic levels. There will be extra costs and overheads for signing the root zone, managing keying material, using secure facilities and tamper-proof hardware, providing transparent processes and so on. These are containable within the existing budgets of the organisations that would presumably be involved in the production and distribution of a signed root zone.

One potential hazard here is ICANN's intention to create large numbers of new top-level domains. This could impose a large burden on the organisation which signs the root and the root servers. If these proposals go ahead, it could compromise the security and stability of the DNS system as a whole because of the extra load needed to sign and propagate a greatly increased root zone. A significant expansion of the root zone needs to be carefully weighed against the impact of deploying DNSSEC.

Operational Questions Concerning Signing of the Root Zone

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 process 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?

In my opinion, proposals 1, 2, 3 and 6 are better than the others because they provide a clear separation between the holder of the zone signing and key signing keys. These have different roles in DNSSEC and for the root zone it makes sense for these roles to be separated. Access to the KSK will only be needed infrequently while the ZSK will be used on a perhaps daily basis to generate a new version of the signed root zone because of changes to TLD delegations. Any of these proposals would be acceptable from a technical or operational

perspective. Proposal 6 is best on political grounds because it provides a capability for the KSK to be distributed in a way that could prevent it from being under the control of one organisation or jurisdiction. This is likely to be a very important consideration if the signed root is to be accepted globally.

The process for deploying DNSSEC at the root zone should as far as practically possible follow the way in which changes to the root are managed: namely IANA handles the change request, NTIA/DoC approves it and Verisign arrange for the updated zone to be distributed to the root servers. IANA already has relationships with TLD operators and Sponsoring Organisations, so it makes sense for IANA to handle keying material from TLDs when they implement Secure DNS. In fact IANA has a secure system ready for handling conventional updates to the root zone, so extending this to handle a TLD's key signing key should be straightforward. This would imply that IANA should hold the root zone signing key and be responsible for generating the root zone file. Other arrangements would require more data paths and processes, which would create more scope for confusion and complexity because of the need for extra authentication and validation steps.

It will be critical for the success of a signed DNS root that the process has wide support, both from the Internet community and from other stakeholders. Some level of community involvement and consultation should be provided so the views and advice of stakeholders can be represented to those co-ordinating the signed root zone. Ideally this should follow the established mechanisms of bottom-up, consensus driven decision making and consultation that has prevailed in Internet institutions. Community involvement will be an essential component of engendering trust in the signed root. Other confidence-building measures should also be provided. These could include but are not limited to: open, fair and transparent processes; regular independent reviews and audits; liaison and outreach activities; and public reports on DNS signing operations.

We invite comment with respect to what technical capabilities and facilities or other attributes are necessary to be a Root Key Operator.

Root key operators must have a good understanding of DNS and information security. This should include experience in operational and procedural issues surrounding core Internet and/or security infrastructure. They should have demonstrated experience at providing stable, secure service. They should also have sufficient resources (staff, expertise and finance) to look after the keys appropriately. This implies the root key operators should have stable funding and be based in jurisdictions that are underpinned by a stable government and legal system.

Root key operators should not have conflicts of interest or perceived conflicts of interest: for example by having ancillary businesses which overlap with the responsibilities of holding (parts of) the root key(s). Neutral, non-profit organisations would be preferable since they would not be subject to a company's obligations to its

shareholders if these were not aligned with operations of the root key(s).

The root key operators should be trusted by the Internet community and wider stakeholders and be proven to act responsibly in the best interests of the security and stability of the Internet. The root server operators are obvious candidates to be root key operators because they demonstrably have these attributes. Ideally the root key signing key could be shared between them under the M of N proposal suggested by NTIA. Other possible candidates as root key operators would be the Regional Internet Registries. They are widely trusted and respected.

What specific security considerations for key handling need to be taken into account? What are the best practices, if any, for secure key handling?

Procedures for handling the keys should follow best practices from comparable security sensitive environments: eg banking and the military. These should be published openly and subject to regular external scrutiny, both for audit purposes and for threat analysis. The objectives here are primarily transparency and trust: the Internet community, governments, regulators and the general public must have confidence in the way keying material is managed and protected. This requires full disclosure of the system, except of course for the keys themselves. Security through obscurity should be avoided: the only information that should be confidential in any security system is the key. Some use of tamper-proof hardware would be desirable. The suggestions made in this area by both the ICANN/IANA and Verisign proposals are satisfactory.

Adequate backup measures should be in place to recover from a damaged or lost hardware containing the root keys (or part of the key). Consideration should be given to providing a second site to act as a backup or standby key repository, ideally in a different geographic region and legal jurisdiction. The system for managing the root keys should have flexibility to allow the location and organisation(s) involved to be transferred with minimal disruption to DNSSEC operations. This would ensure robustness and continuity in the event that export controls or other legislative measures make it impractical to distribute the keys (or parts of the keys) from a particular jurisdiction.

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.

A multi-signature approach should be used for the root zone's KSK. DNSSEC does not care how the KSK is generated and managed, so there is no distinction in purely protocol terms between a key held by one organisation and one that is shared between organisations. However for a number of practical and geopolitical reasons, it would be advisable for the KSK to not be in the hands of any single organisation or

jurisdiction, no matter how benign. An N of M approach would need some testing on matters of process and procedure such as the key-signing ceremony. This would be analogous to how keying material is generated and distributed between key holders in other arenas such as the banking industry. These techniques are quite mature and well tested, so applying an N of M solution to the root KSK should be straightforward.