Response to the NTIA Notice of Inquiry

November 24, 2008

REQUEST FOR COMMENT

The Department seeks comments on DNSSEC deployment and a signed root generally, as well as specific details, comments, and evaluations of the various process flow models proposed or other process flow models that may otherwise be technically feasible to implement DNSSEC at the root zone level. Please include an analysis of the risks, benefits, and impacts of each process flow on the DNS security and stability generally. This analysis should include whether there are security weaknesses or strengths with each process flow model, whether there are methods or suggestions that will increase security and efficiency, and/or whether any alternative process flow models exist that may be preferable to those described in the appendix.

> *Autonomica wishes to thank the NTIA for its open consultation on DNSSEC deployment and a signed root.*
>
> *As the operator of* `i.root-servers.net` *(one of Internet's 13 DNS root name servers) and as one of the early DNSSEC adopters, we believe that the successful deployment of DNSSEC is crucial for the continued stability and security of the Internet and this is contingent upon a signed DNS root zone.*
>
> *We have from the very early beginnings of DNSSEC believed that DNSSEC was essential in making the DNS system more reliable and thereby less susceptible to fraudulent behaviour. Autonomica staff has been heavily involved with the standards development of DNSSEC, testing of early implementations, as well as operating the public name servers for the Swedish Top Level Domain ".SE" – the first TLD to be signed.*

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?

  > *No. The technical community has been working on these issues for a long time, and has converged on DNSSEC as the solution. We strongly encourage the use of DNSSEC to mitigate these issues.*
  >
  > *Furthermore, while alternative actions and methods may have a mitigating effect on the problem, we are not aware of any solution that covers as many aspects of the problem as DNSSEC does. They each contribute with a small improvement for a specific problem, but DNSSEC has more of an over-all effect on many aspects of the problem. We therefore believe that efforts on making DNSSEC work are well spent, and that resources diverted to other solutions will only delay the deployment of the better solution.*

Address	Electronic mail	Telephone	Facsimile	Corporate ID
Bellmansgatan 30 SE–118 47 Stockholm, Sweden	`info@autonomica.se`	+46 8 – 562 860 00	+46 8 – 442 09 67	556586–9467

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

  The Internet is an open network comprised of innumerable parts, where each participant contributes resourses – be that an Internet Service Provider, a datacenter with server farms, or the DSL router and desktop computer at home; they are all part of the Internet. It is virtually impossible to predict which way through the system one's packets will take, and that goes for DNS packets as well. In addtion, the DNS system is astonishingly distributed, and the process of resolving a DNS name may lead to packets being sent to various corners of the world, passing through systems owned by a vast number of Internet citizens. One's immediate conclusion must be that packets can be read and modified by a lot of people during their travel across the network.

  In order to secure the DNS packets, so that the information therein can be trusted upon, it must be made sure that they can traverse any part of the internet without the risk of packet modification going unnoticed. The only way to do this, is to use keys and encryption to generate digital signatures. For the DNS, use of symmetric crypto systems ("one to one") is impossible, due to the distributed nature of the DNS database. It simply doesn't scale. One cannot have each Internet user maintain a peer-to-peer relationship with every DNS server in the world. The number of such relationships that would need to be maintained is just staggering (in the order of 10 raised to the power of 12). The only way forward is to use crypto systemes with asymmetric *algorithms ("one to many"). For well over 10 years, the Internet engineering community has spent time on trying to figure out a good way to use them together with the DNS. The result is DNSSEC.*

  Any mechanism that doesn't *use asymmetric crypto systems cannot be anything else than temporary band-aid, as technology will catch up with it, and new methods of circumventing such mechanisms will be invented by miscreants. Digital signatures are indeed also just a barrier, but with well designed crypto systems, the height of the barrier is significant, and well defined, i.e., it is known what reasourses it takes to climb that barrier, and the height can be adjusted (by changing crypto algorithms or key lengths) and judged for appropriateness.*

  Almost all methods that claim to improve DNS security have the problem that they try to utilise odd corners of the protocol, where the specifications are or have been fuzzy, or which implementors have simply not bothered to implement the details of. It is possible to increase the level of robustness by refining details in these corners, but that will not change the fundamental problem that traditional DNS is not designed with security in mind. Yes, they may be quicker and cheaper to implement, but they will not make DNS substantially more secure.

  DNSSEC may seem cumbersome and complex, but it has a well designed

Address
Bellmansgatan 30
SE–118 47 Stockholm, Sweden

Electronic mail
info@autonomica.se

Telephone
+46 8 – 562 860 00

Facsimile
+46 8 – 442 09 67

Corporate ID
556586–9467

strength in the realms of security – vetted by not only DNS engineers, but also by security experts. We see no viable alternatives when it comes to substantially increasing the security and robustness of the DNS.

- What factors impede widespread deployment of DNSSEC?

  The absence of a signed root zone directly contributes to the development of inferior alternatives, thereby confusing the community and jeopardising the long term success of DNSSEC deployment. In our view, a further politicisation of the issue of signing the root zone may jeopardise DNSSEC deployment in other parts of the hierarchy.

- 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?

  Wide deployment of DNSSEC requires deployment of validation (in the service providers resolvers) as well as deployment of signed zones. Without a single trust anchor (i.e., a signed root zone) in the secure hierarchy of DNSSEC, the incentive for providing signed zones lower in the hierarchy is low.

  It's also worth noting that due to privacy concerns, many Top Level Domains have been waiting for an update to the DNSSEC protocol specification called `NSEC3` to become available, and have signaled that once software that supports `NSEC3` becomes available on the market, they will deploy DNSSEC. With more signed zones becoming available, the incentive for validating zones will increase, but deployment of validation without a single trust anchor is laborious and error prone. So in short, the availability of a signed root zone is a key gating factor for wider DNSSEC deployment.

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?

  Yes. DNSSEC should be deployed in the root zone to allow a reversal of the present trend of decreasing trust in the Domain Name System. While DNSSEC may cause various other benefits over time, the crucial issue is trust in identity. This because the Internet in general is moving in a direction where trust in identity is no longer implicit, but has to be earned (by providing means of verification).

  The time frame should be "without delay" from an engineering point of view.

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

  There are two primary benefits. The first one is that a signed root zone allows validating resolvers to only track one single "trust anchor" rather

Address
Bellmansgatan 30
SE–118 47 Stockholm, Sweden

Electronic mail
info@autonomica.se

Telephone
+46 8 – 562 860 00

Facsimile
+46 8 – 442 09 67

Corporate ID
556586–9467

than separate trust anchors for every signed sub-tree. This significantly lowers the amount of effort needed in the validator end and will therefore improve the DNSSEC value proposition.

The second benefit is that a signed root zone signals to the community that DNSSEC is finally ready, that the Internet community agrees that DNSSEC is the answer to the DNS trust concerns and that there are no other alternatives around the corner.

- 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?

  We believe that enough testing of DNSSEC has been done. It has been deployed in production environments and has proved to work. It will always be possible to argue the benefits of further testing, but at some point the advantages of further testing will no longer outweigh the disadvantages of further delays. It is our belief that we have reached this point.

  Furthermore, while it may be argued that more experience with some parts of the DNSSEC specification (e.g., key rollovers) would be good, it is important to note that it is quite possible to continue such testing and software development in parallel with already having deployed DNSSEC in the root zone. One alternative to achieve this is to not do any key rollovers at all and another is to make key rollovers less sensitive by not accepting so-called signed delegations.

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

  A signed DNS root zone would remove a crucial obstacle to wide-scale deployment of DNSSEC. It would encourage the deployment of DNSSEC throughout the DNS hierarchy and thereby contribute to greater trust in the DNS in general.

- 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?

  Autonomica is one of the root operators and in the case of `i.root-servers.net` *we are fully prepared for deployment of DNSSEC in the root zone and have been so for several years. The same is true for the other root operators.*

  As for other entities, they may be affected in two different ways.

  First and foremost they may be positively affected as part of DNSSEC deployment gaining in importance which may increase trust in the Domain

Name System in general, and also generate new applications and business opportunites.

Secondly, they may be negatively affected in terms of potential operational problems. With our experience in DNSSEC deployment, which includes significant testing and many deliberations, we cannot foresee any such operational problems affecting the DNS users.

- 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?

  For Autonomica the additional cost to providing service for a DNSSEC signed root zone as opposed to the traditional unsigned root zone is zero.

  Registries and registrars would incur noticable costs in deploying DNSSEC in their respective zones. But this cost is not tied to DNSSEC deployment in the root zone as such, as it will be a function of entirely separate decisions. However, deploying DNSSEC in the root zone will make such decisions easier to make, as a DNSSEC signed root zone will improve the value proposition of DNSSEC for other zones.

  Software vendors will over time in all likelihood choose to incorporate DNSSEC verification as a method of adding value to their products.

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?

  We do not recognise any of the models as a perfect model, simply because such a solution cannot possibly be found. We make the following observations that may influence the choice of overall model.

  I. *DNSSEC is about data authenticity and integrity and not about control. Today's system with four involved players (i.e., the IANA function, the Dept. of Commerce, Verisign, and the root name server operators) works, and should any one of them stop to fulfill their undertakings,*

Address
Bellmansgatan 30
SE–118 47 Stockholm, Sweden

Electronic mail
info@autonomica.se

Telephone
+46 8 – 562 860 00

Facsimile
+46 8 – 442 09 67

Corporate ID
556586–9467

they would suffer severe penalties – from governments, from Internet business conductors, from Internet users, from voters, etc. They can each be said to exercise varying amount of control over the root zone.

There is no reason to believe that the introduction of DNSSEC would change the control model in any significant way. The role of the DNSSEC technology is to preserve the integrity of the DNS information as it traverses the Internet; it is not to "vouch" for the correctness of the data. It should rather be viewed as a "seal", the role of which is only to indicate that "this information has not been tampered with while in transit".

II. *The introduction of DNSSEC to the root zone must be made in such a way that it is accepted internationally in the broadest sense of the word. It is vital that the organisations chosen to perform the various tasks receive trust from all parts of the Internet community. Trust is something that is earned by showing long term dedication for duties and respect for others, and deep trust cannot be established in a short period of time.*

One consequence of this is, that it is unsuitable to try to create a new organisation to fulfill any of the DNSSEC roles on short terms. A new organisation would have to operate for a long period of time (probably in the order of several years) before the Internet community would be able decide whether it wants to trust in the organisation, its processes, and its staff. Such a build-up of trust in a new organisation would introduce a much unwanted delay in the overall process of deploying DNSSEC in the root zone, and we have already stated the urgency in getting this done.

We therefore support an approach where existing organisations are evaluated by the level of trust they receive from the Internet community today, and where the outcome of the evaluation is used as input when selecting a future organisational model. In this process, care should be taken to spread responsibilities over more than one organisation, to maintain a balance, and limit the susceptibility to mistakes and outside pressure to make decisions.

While we recognise that the challenge of finding the most appropriate model for the signing of the root zone may be difficult, we wish to stress the importance of this process not taking an undue amount of time, thereby further delaying the signing of the DNS root.

- We invite comment with respect to what technical capabilities and facilities or other attributes are necessary to be a Root Key Operator.
- What specific security considerations for key handling need to be taken into account? What are the best practices, if any, for secure key handling?

Address
Bellmansgatan 30
SE–118 47 Stockholm, Sweden

Electronic mail
info@autonomica.se

Telephone
+46 8 – 562 860 00

Facsimile
+46 8 – 442 09 67

Corporate ID
556586–9467

As with any large system, it is virtually impossible to get all details right from the start, and since the world is a place of constant change, the chosen model for signing the root must be able to evolve with time.

To facilitate that, one quickly arrives at the conslusion that it must either *be possible to transfer the root Key Signing Key (KSK) from one party to another,* or *it must be possible to "rotate" the KSK (i.e., issue a new key and change the active key to the new one).*

In this context it's important to note that the DNS, as many other systems, has a producer side, and a consumer side. What we are discussing here is to provide *DNSSEC at the producer side. This does not force any consumer to "buy the product", but it gives those who want to the take advantage of the security extensions to do so.*

There must be a limit to how far the producer can be expected to take responsibility for how the consumer uses the product. We must expect users who decide to take advantage of DNSSEC to be able to follow the instructions, e.g., to update the trust anchor with a new one, once the KSK is rotated. Just to make a parallel: you need to fill up gasoline in your car from time to time; else it won't take you further. It's not the car manufacturers' problem to make sure that you remember to fill gasoline, or to tell you where the gasoline stations are located.

The above said, we come to the conclusion that it must be possible to rotate the KSK. The current methods may be crude, but it is possible. While the Internet community, wich the engineering corps at the forefront, continue to develop better and more automated ways to rotate the KSK (and consequentially the trust anchors), the root can be signed, and the KSK rotated manually.

To futher limit the impact of introducing DNSSEC in the root zone, we propose to sign the root zone, but to start by not signing any delegations. This way, there will be no secure delegations yet, but there will be a KSK, there will be signatures, and there can be KSK rotation. Processes for this can now be established. None of this will impact the existing signed TLDs (nor the unsigned TLDs), since there will be no displayed security connections between them and the root. Once the processes are in place and trusted, signed delegations can be introduced, and we will finally have reched the goal of making the most significant part of the DNS system trusted.

The ability to rotate the KSK is also important to facilitate development of the signing process. As we have already stated, speed of deployment is of essence. If the KSK is allowed to change, a first implementation of a signed root need not be perfect, and one can start with a relatively simple method which involves few parties and simple key handling – e.g., some process which involves (a subset of) the current four root zone manatement players. That will be relatively quick to implement and easy to test. In parallel to this, newer and better processes – possibly with new organisations

Address
Bellmansgatan 30
SE–118 47 Stockholm, Sweden

Electronic mail
info@autonomica.se

Telephone
+46 8 – 562 860 00

Facsimile
+46 8 – 442 09 67

Corporate ID
556586–9467

– can be developed, using more elaborate key schemas, and with more parties involved. Hence agreements must be carefully designed to allow for this.

The public part of the root Key Signing Key, used by client resolvers to validate the signed DNS data, must be distributed as widely as possible, and in a reliable manner. The key must be so well known that it becomes impossible to even try to forge it. One idea in addition to more obvious methods like wide-spread announcements, is to make use of already existing local associations where trust hierarchies other than DNS could be built and used (such as ISOC with its chapters, the Regional Internet Registries, Internet Service Provider organisations, etc.).

Obviously, when information is transferred between parties, transfer must not occur without proper authentication and integrity validation, and in the case of private keys, also that appropriate measures are taken to avoid dissemination of the information to inappropriate parties.

Should the root KSK not *be allowed to rotate, we end up with a more or less permanent root KSK. This will lead to a number of disadvantages:*

- *It will be next to impossible to transfer the private part of the KSK from one party to another. The suspicion that the first party retains a copy of the key will always linger. Such a retained key could be used to sign data that is not supposed to be signed.*
- *The above will lead to a situation where the party chosen to hold the private KSK perpetuates itself in that role.*
- *If the private part of the KSK is stored in a Hardware Storage Module (HSM), it may prove inappropriate to store it together with other keys used for other purposes, since that may make it impossible to transfer just the DNS root key to some other environment.*
- *If the key is stored in a unique HSM that is owned by one particular player, that player may claim that ownership, thereby keeping the key as a "hostage". Ownership of the equipment that holds the KSK will have to be very carefully negotiated.*
- Should *the root KSK be compromised, it will be next to impossible to recover the system to a secure mode, since the processes for key rotation will not have been tested and vetted.*

- 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 use of a multi-signature technique would indeed have some advantages. It could relieve some of the political anxiety that seems to surround the issue. While we are in favour of such solutions, we want to stress that the

Address
Bellmansgatan 30
SE–118 47 Stockholm, Sweden

Electronic mail
info@autonomica.se

Telephone
+46 8 – 562 860 00

Facsimile
+46 8 – 442 09 67

Corporate ID
556586–9467

timeliness of deployment of a signed root must take priority over the "political correctness" of the methods used. The most important characteristic of the chosen process for signing the root must be that it can be altered in the future. As a consequence it is less relevant which signing method is used to begin with. Choosing something simple to start with, to "get things going" and to enable domain name holders and the corresponding DNS users to start exchanging trustworthy DNS information is the important step. The signing method and key handling need not be perfect from day one, if *it is done in a way that allows for change. Improvement over time is the engineering principle that most often generates stable and reliable end results.*

We also want to point out that there is a significant difference between the process outlined in "flow 6" in your appendix, and the proposal from Verisign on how to handle root signing. In "flow 6" the text suggests the root Key Signing Key that is split between a group of organisations, e.g., the public root server operators, in such a way that each participant holds a fragment of the actual *key. In the Verisign proposal the* actual key *is stored* in its entirety *in Verisign's Hardware Security Module (HSM), and the fragments (or tokens) mentioned in their proposal are needed to* access *(or "unlock") the actual key, which, again, sits in its entirety in the HSM module. Transferring the actual key to some other party may prove very difficult.*

Then again, transferring a key that is fragmented and split between N parties also has its problems, but at least, if the fragmenst are transferred one by one in bilateral transactions, no party has access to enough fragments to create its own authority over the signatures in the root.

Address
Bellmansgatan 30
SE–118 47 Stockholm, Sweden

Electronic mail
info@autonomica.se

Telephone
+46 8 – 562 860 00

Facsimile
+46 8 – 442 09 67

Corporate ID
556586–9467