Dear Ms. Alexander,

I'd like to offer the following comments to the NTIA:

The original DNS had no security measures. Over the years we have reused protocol features (like the ID field) to make life more difficult for the attackers, as well as adding technology to protect updates, zone transfers, and the like. However, there's a growing threat due to simple cache poisoning attacks, a la Kaminsky. There's also the fact that even the Kaminsky patches are insecure on high speed networks, and the elephant in the room: the fact that if an attacker can divert or wiretap the transactions between DNS entities (server to server, server to resolver) THERE'S ZERO DEFENSE. If a malicious sysadmin chooses to alter data THERE'S ZERO DEFENSE. It's important that we recognize that it's time to address this problem.

Additional DNS security is needed to face known attacks and attacks we can reasonably expect in the future, not to mention the possibility of novel attacks or attacks that blend DNS attacks with other attack methods. Some believe, myself included, that a higher level of DNS security will also enable new uses of the DNS, expanding its utility.

For decades, one school of DNS thought has maintained that we "should not overload the DNS". Sometimes meaning that we should not rely on the DNS to carry information that needs to be accurate, and sometimes meaning that political decisions about what data should be in the DNS (e.g. keys) should be enforced at the technical standards level. Both these arguments are now shown to be as bogus the mortgage backed securities risk management. The website addresses carried by the DNS deserve the protection of strong security in order to protect the general public; whether we meant to or not, the DNS carries data that is critical to the user's security. Secondly, the DNS should be viewed as a delivery mechanism that is insensitive to what it carries, and no limitation should be placed on what it carries, and if DNSSEC goes forward, it's really a PKI, whether we choose to call it that or not.

We should not muddle political and technical choices. We should also question several of the assumptions on which DNSSEC is based. We need security that is understandable both in principle and in implementation.

Are there alternatives? Yes. We can digitally sign the DNS data, or protect the paths between DNS entities, or use additional defenses to the Kaminsky attack which we have found to be effective. Opportunistic encryption between DNS entities, in one form or another can reduce the problem, particularly if extended to the user machine. There's a basic duality between protecting the connections between DNS elements and signing the data itself. The best choice for the future is to see the two as complimentary, rather than in competition. However, digital signature technology is the long term solution.

To be clear, digital signatures for DNS data are an essential part of a better Internet, but DNSSEC is a highly imperfect implementation.

So how do we move forward? We need to recognize that there are both political and technical problems to be solved, and pay careful attention to their real nature.

ICANN and others would undoubtedly like to have sole control of the root signing. This should not be allowed. This is not required by the technology in general, and systems designed around this principle are simply broken by design and don't answer the needs of the network. Everywhere I went during recent trips abroad, I heard skepticism about the US's, and ICANN's hold over the DNS. It's time to break that political problem, by having distributed control, in one form or another. ICANN needs oversight, and the root key signing represents a good time for the US to share that burden with others.

Even leaving aside the sensitive issue of control of the root, the assumption that there should be a single trust anchor is incorrect. Rather than regarding TARs or DLVs as temporary scaffolding, we should see them as the first generation of tools and work on the next generation of technical facilities that make administration in the presence of multiple trust anchors convenient and scalable. Any large organization, whether a commercial company or government, should design its DNS infrastructure to work securely for internal use regardless of whether any exterior domains (including the root) are signed or unsigned, or even if exterior certification make the interior data appear not authentic, nonexistent, or anything else. This means that such organizations will want at least two trust anchors that may disagree or be in conflict. It's logical to assume that extranets may continue to expand the number of trust anchors in place.

What are the significant problems that lie ahead?

- The DNSSEC protocols are still somewhat in flux. It appears that several DNSSEC details are broken or not well understood.
- The difficulty of administration, coupled with the fact 50% of existing domains have small configuration errors, coupled with the fact that small errors can cause DNSSEC tragedies.
- Applications and specialized hardware (e.g. load balancers that use DNS) will need TLC to do DNSSEC right.
- Some think DNSSEC is necessary to answer threats that may emerge any day. But it's unlikely that, absent a $1 billion DNS attack, DNSSEC can be deployed to the majority of the Internet in under 5 years.

What should NTIA do?

In an ideal world, we should undertake a 2-3 year effort to deconstruct and rebuild DNSSEC in a focused way, with the goals of understandability, generality, and ease of deployment. But NTIA doesn't have this choice.

So, in a cynical but enlightened way, we want DNSSEC and the root signing to go forward. We know that we may simply get an evolutionary dead end, but the failed experience will at least get us general recognition of the flaws in the base assumptions of the current DNSSEC. We need a bit of evolution here which may lead to DNSSEC2. What we don't need is another 15 years of intelligent design.

The key aspects should be:

- Distribution of control over the root key, perhaps implemented by the “key splitting” idea contained in one of the existing proposals.
- Policies to further the secure distribution of DNS data via protection of the connections between DNS entities, and the availability of mirrored copies of TLDs to large communication providers.
- Work on contingency plans to deal with a DNS cache poisoning or other crisis.
- Get a study of the operational experience and likely future of DNSSEC performed by NAE or a similar organization. Make sure that includes expertise from the DNSSEC community, but is dominated by representatives of the broader Internet community.

Sincerely,

Paul V. Mockapetris

Chairman & Chief Scientist
Nominum