Progress of DNS Security Deployment in the Federal Government

Scott Rose
NIST

Abstract

In 2008, the US Federal government mandated that all Federal Executive Branch owned DNS zones must deploy DNSSEC. Initial deployments lagged and often error prone, and in response, the DNSSEC Tiger Team was formed to aid deployment and develop a system to monitoring system. The results showed a significant increase in deployment as well as a reduction in errors. When errors were detected, the time it took to resolve the problem was also reduced.

This paper discusses the history of DNSSEC in the gov domain, the types of errors seen, and how they were reported. This paper concludes with a set of lessons learned that would apply to other large domains or groups wishing to make DNSSEC a requirement for operation in members' zones.

1. Introduction

The DNS Security Extensions (DNSSEC) is a collection of extensions to the DNS to provide source authentication and integrity protection. DNSSEC does this by adding digital signatures to DNS data, which clients can validate using public keys also stored in the DNS. The DNSSEC specification was initially published in 2005 by the IETF in RFC 4033[1], RFC 4034[2] and RFC 4035[3].

DNSSEC is deployed on a per-zone basis and uses the existing DNS hierarchy to establish chains of trust from parent zone to child zones. Parent zones vouch for the DNSSEC status of delegated child zones by using a special Delegation Signer (DS) Resource Record (RR). The presence of this RR gives the client information about the key used by the client. Lack of this RR means that the child zone is either not signed, or in the process of deployment and the child zone administrator does not consider their deployment to be production ready. Often the last step in deployment of DNSSEC in a given zone is to upload the key information for the zone to its parent zone. This is done out of band of the DNS, usually through a registrar web portal or similar. For example, the registrar website for the gov Top-Level Domain (TLD) is https://www.dotgov.gov/ and allows registered administrators to upload and request key data to be published for their delegation (e.g. dnsops.gov).

This linking of security from parent to child makes the deployment of DNSSEC at the DNS root and TLD's important, as clients with the DNS root public key and/or TLD public keys would be able to validate the widest set of possible DNS responses. Child zones under TLD's can sign their zones and simply upload their key material to their parent zone without having to go through the effort of publishing their public keys for all clients to obtain.

2. Background and Deployment Drivers

Though the current specification was published in 2005, initial deployment of DNSSEC was scant, with few zones being signed. DNSSEC got a boost in interest with the disclosure and publication of the so-called Kaminksy attack, presented at Def-Con in August 2008[4]. Also that month, the US Office of Management and Budget (OMB) issued OMB-08-23 "Securing the Federal Government's Domain Name System Infrastructure" [5] which set deadlines for the deployment of DNSSEC by Federal agencies. This was often referred to as the "OMB DNSSEC mandate". The deadlines given in the memo were for the gov TLD used by the Federal government to deploy DNSSEC by January 2009 and every Federally owned second level zone be signed by December 2009. Federal agencies only make up roughly 20-25% of the gov TLD, the remaining delegations belonging to state and local governments, and American Indian tribes (Native Sovereign Nations) neither of which falls under the mandate.

Parallel to the OMB mandate was the addition of DNSSEC to the Federal Information Security Management Act (FISMA) controls. FISMA requires each Federal agency, and private entities that posses and process Federal information to have an established IT security policy for each system, and includes a set of checklist items (called controls) that are recommend or required for Federal systems, depending on the risk factor. Deployment of

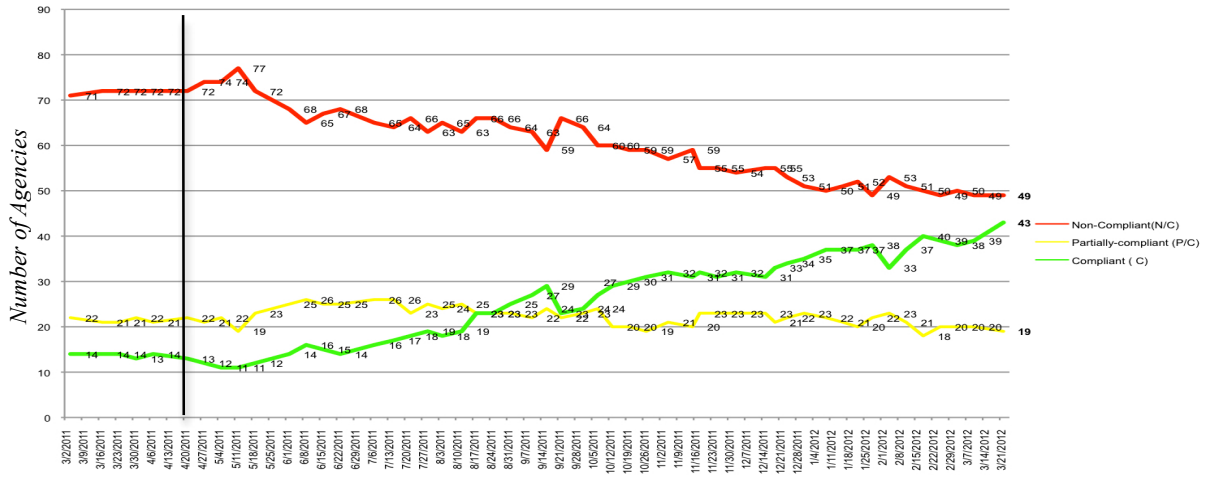

Figure 1 DNSSEC Deployment by agency in the gov TLD per week (taken from a sample DHS compliance report)

DNSSEC was added to the controls for Federal systems to deploy, which applied to all Federal zones, not just second level zones as mandated in OMB-08-23

3. USG DNSSEC Tiger Team

Even with the OMB DNSSEC mandate and FISMA drivers, deployment of DNSSEC by Federal agencies was slow. The gov TLD was DNSSEC signed by February 2009, missing the deadline by 1 month, but only a fraction of agencies signing their zones within the agency deadline [6]. This is partially due to a lack of coordination within the Federal space, as agencies have traditionally been very independent in establishing and managing their DNS. The operation of DNS by Federal agencies varies from internally operated by Federal employees (or contractors) to outsourcing to commercial hosting services, so no "one-size-fits-all" approach could be applied across all agencies.

In response to the slow rate of adoption, the Federal CIO Council chartered the DNSSEC Tiger Team to address the issue. The Tiger Team initially met in April of 2011 and met monthly to coordinate efforts to discuss deployment barriers and address concerns. The Tiger Team consisted of volunteer participants from various agencies and members of DHS Federal Network Security (FNS) tasked with monitoring deployment of DNSSEC in the gov TLD. NIST participated in the Tiger Team first as subject matter experts, and later as co-chair of the Tiger Team.

The Tiger Team helped increase the number of signed zones through a program of training, communication and monitoring. The team collected a set of training material for distribution to all Federal DNS administrators, as well as funding a computer based training course made free for all government employees. An internal, government only email discussion list was created for administrators to discuss issues, roadblocks and tips. Finally, a DHS program to regularly monitor DNSSEC deployment was established that would send weekly reports to administrators (via the email discussion list) and monthly reports to agency CIO's.

As seen in Figure 1, the Tiger Team has had a measureable impact on the number of signed zones. In the figure, the black vertical line in April indicates when the Tiger Team first convened. A few weeks after the Tiger Team formed, the DNSSEC compliance monitoring and reporting program started at DHS. From then, the number of DNSSEC compliant agencies increased. On March 26th, 2012 (the latest date in Figure 1), 910 unique Federal zones were signed and chained from the gov TLD, or 54% of all Federal zones.

One other interesting thing to note is that the total number of Federal zones in the gov TLD have also decreased during this period. There are two factors to this trend: First, OMB memo 11-24 [7] issued a call for agencies to reduce the number of Federal domain names and websites in order to reduce possible citizen confusion. Secondly, the DNSSEC requirement gave many agencies a chance to re-visit its DNS infrastructure and offered an excuse to take inventory and remove zones that were no longer needed or desired.

4. Types of Errors Seen, and Efforts to Remediate Them

DNSSEC involves a new set of operations to traditional DNS operations. Rushing to deploy can result in errors, which can be worse than not deploying. NIST performed daily scans of the known Federal gov name space and recorded any seen errors that would cause a client to reject DNS responses from a zone. These errors can be broken down into five basic categories:

- **NoSigs**: The parent zone claims the zone is signed (i.e. a DS RR is present), but the zone does not have DNSSEC signatures over the zone data. Clients would expect signatures, and failure to obtain them in responses results in an error.

- **ExpiredSigs**: The parent zone claims the zone is signed, and the zone has signatures over its data, but the signatures have expired. The client would reject these responses as invalid.

- **SigsPriortoInception**: The parent zone claims the zone is signed, but the signatures in the zone have a creation date set in the future. This means that a validating client would reject these signatures, as they are not valid. This error is likely due to a clock error on the system used to generate signatures.

- **BadKeyRollover**: The zone has recently changed its keyset, but the parent zone was not informed and continues to publish the old DS RR. Clients see a key mismatch between a trusted source (the parent), and the zone, and a validation error results.

- **DSPointstoPre-publishedKey**: The zone is in the process of changing its keyset, but the parent zone has a DS RR for the (possibly) new key instead of the current key. This error is rarely seen, and was largely due to an appliance implementation error that has since been patched, but still seen in production services.

Other errors were seen during this period, but they were not directly DNSSEC related, so they were not included. These errors were typically network related issues, or system problems that rendered the entire zone unreachable for all clients, not just for those performing validation.

Not having a DS RR in the parent zone is not considered an error, as this will not result in the response being rejected by a validating client. These zones are often called "islands" since they are signed, but often can't be validated unless clients have some means of obtaining the zones' keys in a trusted manner.

Figure 2 below shows the DNSSEC errors seen per day during the last five months as Figure 1. The errors are color coded to show each category.

From the above figure, it is clear that while the number of errors changes over time, the two categories that make up the majority of the errors are ExpiredSigs and BadKeyRollover. That is likely due to the fact that these are often manual operations done by administrators or (in the case of BadKeyRollover) require a human to perform at least part of the operation (i.e. interact with a registrar).

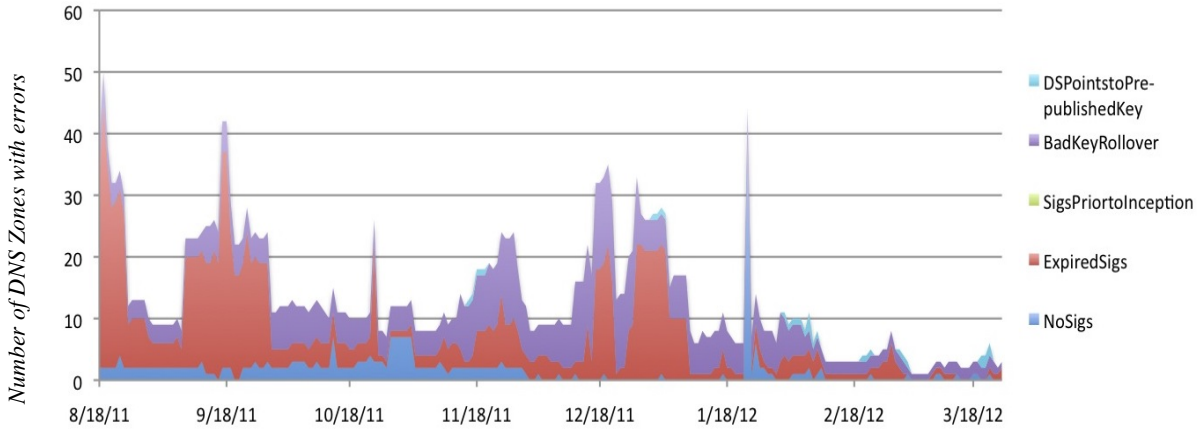

Figure 2: DNSSEC errors seen per day in gov domains

Further evidence that these errors are due to the dependence of manual operations is the spike seen during the holiday months (i.e. Aug/Sept and Nov/Dec). During these months, DNS administrators are often on leave, or dedicating their time to end-of-year tasks and DNSSEC maintenance operations (like DNSSEC re-signing) are overlooked. The number of errors seen drops post December, indicating that administrators are back at work and resolving the issues seen over the holidays.

It is also interesting to note that while the overall number of signed zones have increased (Figure 1), the number of errors seen have stayed the same or decreased. The monitoring and reporting system in place is often credited with this trend. As both IT management and administrators are informed as to the state of their DNS, errors are being caught sooner and problems resolved quicker.

5. How Monitoring Helps

It has been shown that the DHS monitoring program has increased the number of signed zones in the Federal DNS name space, but also has helped shorten the time between the initial failure and remediation of the problem. Analyzing the number and types of errors in the April of 2011 (after the Tiger Team was formed but before the monitoring program started), shows the following breakdown of errors as seen in Table 1.

Error seen	Num. Errs	Min. Days	Max Days	Avg. Days
NoSigs	41	1	20	2
ExpiredSigs	21	1	27	6
SigsPriorto-Inception	1	9	9	9
BadKeyRollover	3	1	27	14
DSPointstoPre-publishedKey	6	1	27	9

Table 1: An analysis of errors seen in March 2011

The large number of errors and the wide range of days until the errors are addressed reflect the wide range of quality in DNSSEC operations in the Federal name space. Interestingly, the majority of the errors are ExpiredSigs and NoSigs. It is believed, based on some anecdotal reports that the NoSig errors are due to failures in configuring name servers to properly serve DNSSEC responses. BIND (the most widely used software for authoritative DNS servers) requires manual changes to its configuration file before loading and sending DNSSEC replies. DNSSEC processing is turned off by default.

The length of time required to resolve some of these errors seems very long. In the most extreme cases, problems were not addressed for weeks. This is likely attributed to the lack of outside clients asking or checking for DNSSEC signed responses. It is known that very few (if any) clients were performing DNSSEC validation in March of 2011. Even today, only one major US ISP (Comcast) provides DNSSEC validation for all customer DNS queries. There is also the problem of how to report issues when DNSSEC validation fails. Unlike problems with web pages or other services, there is no common standard "dns@example.com" type email addresses to report problems. The RNAME field in the SOA RR is cited as the place for this address (RFC 1035 [8]), but it is often not used correctly or redirects to a mailbox that is rarely (if ever) checked. Often, problems are resolved by asking on email distribution lists for a point of contact, which often relies on finding the right audience or searching for a help desk at the zone's registrar or hosting service.

Performing the same analysis in April of 2012 shows a marked improvement; not just in the number of errors, but also in the time it took for the errors to be resolved. The improvement can be seen in Table 2.

Error seen	Num. Errs	Min. Days	Max Days	Avg. Days
NoSigs	6	1	1	1
ExpiredSigs	4	1	4	2
SigsPriorto-Inception	0	0	0	0
BadKeyRollover	2	3	7	12
DSPointstoPre-publishedKey	3	3	3	3

Table 2: An analysis of error seen in March 2012

This improvement can be attributed to the monitoring program that not only regularly checks the DNSSEC validity of zones, but also reports directly to agency administrators and managers the status of their zones.

This feedback is seen as critical to the success of deployments. The second factor is the increase in knowledge sharing between DNS administrators that have resulted in improved operations in agencies. More agencies are automating regular DNSSEC operations such as resigning and portions of the key rollover process. This is done using a range of options from simple scripts and open source tools to purchasing automated appliances or outsourcing of operations to contracted parties.

The current situation is not immune to problems, however. The most famous example is the incident in January of 2012 when Comcast customers could not reach nasa.gov servers because the nasa.gov zone incorrectly performed a key rollover. The after action report from the Comcast perspective was published in their blog [9] which contains a detailed description of the problem and gives details about how they react to DNSSEC validation failures in order to maintain service for their customers.

6. Lessons Learned

The experience with the Tiger Team in deploying DNSSEC across the Federal Executive branch highlight several key lessons to consider for large organizations and communities when they seek to deploy[10]. Perhaps the most important lesson doesn't directly involve DNS at all; that is, knowing whom in another organization is responsible for DNS (and by extension, DNSSEC), and who to contact when something goes wrong. The first issue the Tiger Team had was identifying responsible points of contact (either administrators or network managers) in each agency. This is especially difficult in agencies that outsourced their network operations.

This becomes more important if an error is detected. By policy, the gov zone does not have a thick WHOIS, and does not list email addresses of points of contact. There is a registrar help desk that is publically accessible (at http://www.dot-gov.gov/), but few Internet users and network managers outside of the government know of its existence. Outside of the government, a delegation's WHOIS information may be out of date, or not list the actual day-to-day operator of the problem zone. Domain name owners should take every step to insure that their WHOIS information is current as well as having a valid email address in the SOA RR that is monitored by operations staff.

However, often an end-user does not seek out points of contact to report errors, so zone administrators must be pro-active. The Comcast-NASA.gov incident illustrated the point that the majority of end users are not aware of how DNS or DNSSEC works, and will instead vent their frustrations on social media sites first, and their ISP's help desk second.

The second lesson is that it is easier to detect one's own problems than to react to them when learning from outside sources. Regular monitoring from an outside point of view can alert administrators to a potential problem before they become a larger issue. Some external sites provide a snapshot view of the DNSSEC status of a given zone, but it is trivial to set one up to monitor a zone's own DNSSEC status. This could be extended to warn of potential failures (i.e. signatures that may expire soon) as well.

The third lesson is it helps to have a forum for community members to hold candid discussions on roadblocks, challenges and ask questions about DNSSEC deployment. NIST created the "gov-dns" mailing list on behalf of the DNSSEC Tiger Team to be used by USG administrators or contractors directly supporting a zone in the gov TLD. The purpose for the restriction was to give administrators a forum to ask questions about Federal policy (e.g. key size and rollover frequency) or questions they may feel uneasy asking in a public DNS operations forum (i.e. questions about a specific DNS server implementation). Admins who noticed a particular issue with a .gov zone would also use the list to call attention to the issue if they could not reach a zone POC directly. While not perfect, it was one of the ways USG zone administrators were able to coordinate responses to DNSSEC errors. The forum was also served as the outlet for summary compliance reports from DHS to administrators to show progress of DNSEC deployment in the gov TLD.

Lastly, DNSSEC requires a different operational approach than traditional DNS, as there are more time dependent operations such as re-signing of zone data and routine key changes (depending on an organization's security policy). These operations can be easily automated and tools are available (both open source and COTS) to aid administrators. These tools, coupled with tasking backup administrators for coverage during holidays or leave periods will reduce the risk of the most frequent types of DNSSEC errors encountered (e.g. ExpiredSigs).

7. Current Status and Future Efforts

The USG DNSSEC Tiger Team is no longer meeting,

but the monitoring program set up by DHS FNS continues to scan Federal domains and issue reports. The rate of DNSSEC deployment continues to increase and the frequency of errors continues to remain low and stable at less than ten zones identified as having errors seen on a daily basis (1%-2% of all signed Federal zones).

DNSSEC is seen as an enabling technology, not just a means to protect traditional name to address translation. Secondary to DNSSEC deployment, the USG Tiger Team sought to increase the use of various email authentication techniques that rely on the DNS in some way, such as Sender Policy Framework (SPF) [11] by Federal agencies. These technologies publish email policy information in the DNS, which can be digitally signed by DNSSEC just like any other type of DNS information. It is still too early to conclusively say how a signed DNS infrastructure can be used to build trust in other applications and services.

8. Acknowledgements

The author would like to acknowledge key individuals who were vital to the increased deployment of DNSSEC in the gov domain. These individuals are Earl Crane, DHS and chair of the DNSSEC Tiger Team, Valeri Stoyanov of DHS FNS and the other members of the USG DNSSEC Tiger Team.

9. References

[1] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose, DNS Security Introduction and Requirements. RFC 4033, March 2005.

[2] R. Arends, et al. Resource Records for DNS Security Extensions. RFC 4034, March 2005

[3] R. Arends, et al. Protocol Modifications for the DNS Security Extensions. RFC 4035, March 2005.

[4] D. Kaminksy. BackOps 2008: It's the end of the Cache as we Know it. http://www.slideshare.net/dakami/dmk-bo2-k8

[5] Securing the Federal Government's Domain Name System Infrastructure. Office of Management and Budget Memoranda 08-23. http://www.whitehouse.gov/sites/default/files/omb/memoranda/fy2008/m08-23.pdf

[6] William Jackson. DNSSEC spreads slowly through government domains. *Government Computer News Magazine*. Sept 23 2010. http://gcn.com/articles/2010/09/23/governmemt-slow-to-deploy-dnssec.aspx

[7] Implementing Executive Order 13571 on Stream lining Service Delivery and Improving Customer Service. Office of Management and Budget Memoranda 11-24. http://www.whitehouse.gov/sites/default/files/omb/memoranda/2011/m11-24.pdf

[8] P. Mockapetris. DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION. RFC 1035, November 1987.

[9] Analysis of DNSSEC Validation Failure Comcast – DNS Engineering http://www.dnssec.comcast.net/DNSSEC_Validation_Failure_NASAGOV_20120118_FINAL.pdf

[10] CONSIDERATIONS AND LESSONS LEARNED FOR FEDERAL AGENCY IMPLEMENTATION OF DNS SECURITY EXTENSIONS AND E-MAIL AUTHENTICATION. Federal CIO Council Whitepaper. Nov. 2011 http://www.cio.gov/documents/DNSSEC_and_E-Mail_Authentication_Considerations_and_Lessons_Learned.pdf

[11]M. Wong. Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1. RFC 4408. April 2006.