Key Deletion Issues and other DNSSEC stories

June 22th, 2011

vincent.levigneron@afnic.fr

Plan

- Key numbers of the publication process
- AFNIC DNSSEC specifications
- Key deletion process in AFNIC zones
- First DNSSEC outage in November 2010
- Second DNSSEC outage in February 2011
- Third DNSSEC outage in March 2011
- Thanks to the community
- What happened next
- Our Proxy Server
- Lessons learned
- DNSsexy

Key numbers of the publication process

- **AFNIC registry operates 6 ccTLDs (fr/re/pm/tf/wf/yt).**
- **Each zone is signed (DNSSEC was introduced in september 2010). Zone Signing Keys are rolled over every 2 months.**
  - NSEC3+opt-out.
- **fr zone is the largest one with more than 2 millions domain names (since less than 2 months).**
- **fr zone contains 4 400 000 Ressource Records.**
- **Very few DS records yet (registration of DS has been launched less than 2 months for .fr and .re).**

AFNIC DNSSEC specifications [1/2]

- **OpenDNSSEC is only used for Key Management.**
- **AEP Keyper HSM are used for Key storage.**
- **Bind (auto-dnssec allow; option set) do all the signature stuff (with HSM).**
  - Version 9.7.1-P2 was first used.
  - Version 9.7.3 deployed after second outage.
- **Homemade synchronisation script to create V1.3 Bind key files from ODS data.**

AFNIC DNSSEC specifications (closer look on fr zone) [2/2]

- **While there are more than 4 millions records…**
- **… there are less than one hundred NSEC3 and RRSIG records.**
- **2 KSKs (one published, the other active).**
- **2 or 3 ZSKs at a time (one published ready to be used, one active, and if we are just after a key rollover, the previous active key is still published while inactive).**
- **Zone is dynamically updated (RFC 2136) every hour and once a week there is a complete zonefile generation mainly for administrative purposes.**
- **Dynamic Updates is only used for delegations (NS/A/AAAA. And DS since 2 months).**
- **All key/signature stuff is based on "automatic signing" Bind capabilities (no Dynamic Updates in this case).**

Key deletion process in AFNIC zones

- **We use very large timings.**
  - When a key becomes inactive, it is deleted one month later (that's why you'll often notice more than 2 ZSK at a time during a key rollover while querying for .fr DNSKEY RRset).
  - When a key is deleted, we purge/archive key files 3 days later (it was just one hour during the first outage we had in November, it has been increased after that event).
- **When a key is about to be deleted, we are sure there are no RRSIG left corresponding to this key.**

First DNSSEC outage in November 2010

- **During key deletion we had a network issue making our HSM unreachables.**
- **The error was not well detected, so the publication process didn't stop as expected.**
  - Zone was not updated. Key with “delete” state was still present (while inactive).
- **OpenDNSSEC to Bind synchronization process (homemade script) decided to purge the key files one hour after it was supposedly deleted.**
- **Then, Bind couldn't process Dynamic Updates.**
- **We also had a Bind “Private Record” Bug ([ISC-Bugs #23232]) we are about to describe... But we were so focused on the other parts of the system we discovered that... 2 months later...**

Second DNSSEC outage in February

- **This should have been a boring key deletion operation.**
- **But we had an unexpected behaviour from a Bind private record usage. Record TYPE 65534 is used to give the state of a signing process.**
- **What was expected… In less than the blink of an eye…**
  - 1/ DNSKEY RR corresponding to deleted key needed to be removed.
  - 2/ DNSKEY RRset signature needed to be updated.
  - 3/ Serial should be incremented.
  - 4/ SOA signature had to be updated.
- **But, there was a bug in Bind (patched since) that lead to a bad signature on Apex NSEC3 record.**
- The .fr zone became inaccessible to any validating resolvers…
- **In fact, the problem also occurred with the other zones we are dealing with. But, with small zones, the problem is "visible" for less than a second… For fr zone, it lasts hours….**

Third (and last ?) DNSSEC outage in March 2011

- **In the morning, there were key states transitions on zone fr. It worked.**
- **Following Dynamic Updates were well processed.**
- **Then Bind "decided" to modify it's private records. But at the same time, we had (again) a HSM reachability issue. The published zone was not correct.**
  - A new TYPE65534 record has been added to the corresponding RRset.
  - SOA RR has been modified.
  - But 2 RRSIG are missing, the one for the TYPE65534 RRset and the one for the SOA RR.
  - There's already a patch, but not yet applied or check if it fix this problem
- **Is RRSIG missing for SOA a big problem ?**
  - Answer is "No ! It's bad but it could be worst".
- **We have 2 NSD nameservers amongst our slaves. And in this very special case, NSD did something "unexpected" (the behaviour is the same for the 2 nameservers).**
  - It also "decided" to remove all other signatures of the apex ?!?!?

Thanks to the community

- **Yes, thank you...**
  - Our monitoring system failed (we didn't check NSEC3 RR), first alerts came from you.
    - Those of you who already use validating resolvers were not able to send us emails. In this case, social networking is a good mean to communicate between registry and the community (Twitter/direct phone call/…).
  - ISC provided a patch very fast (patch 3020).
  - We also found a bug in Unbound, the patch has been published.
  - We also had good feedbacks on our search for a zone verification tool. ldns for instance was promising while not fast enough. Not easy to find a tool able to deal with a “medium size” zone. Validns, while very young and still under development was able to deal with zones with millions of entries very fast.

What happened next

- There will be other issues… Unfortunately… Even if we patch tools…
- Just after OARC/ICANN meeting, we had crisis meeting to decide if we should remove AFNIC ccTLDs DS from the root…
- It was decided that we could keep DS if we could finish the implementation of our proxy and deploy it within the next 2 weeks.
- It was also decided to postponed the the implementation of the service for inserting DS records into the AFNIC zones which was the next step of our DNSSEC project.
- We have modified our system to have a better control over zone changes.
- The zone is now validated before it is sent to our hidden master.
- A new Notify Proxy Server controls this.
- Eventually, we launched our DNSSEC-aware version of Zonecheck as well as a new version of EPP server (RFC 5910/DS Data interface).

Our Proxy Server [1/2]

- **First objective of our Proxy Server is to coordinate all different DNS publication processes (Dynamic Updates, Key Management, Complete Zonefile Generation, the Proxy itself, Zonefile transfer between internal nameserver and hidden primary nameserver, …). This prevents, for instance, Bind from modifying private records during other DNS processes.**
- **Second objective is to add a validation step in this proxy. There are, mainly, two ways for that:**
  - Do a complete zone DNSSEC validation (it's a long term plan).
  - Just check Apex records and some specific ones (it would have been enough to detect the outages we had).
- **If there is a problem, the Proxy stop the publication system and prevent the transfer to our hidden primary nameserver.**
- **… next steps**
  - Automatize recovery system
  - Integrate zonefile revision system in the proxy

Our Proxy Server [2/2]

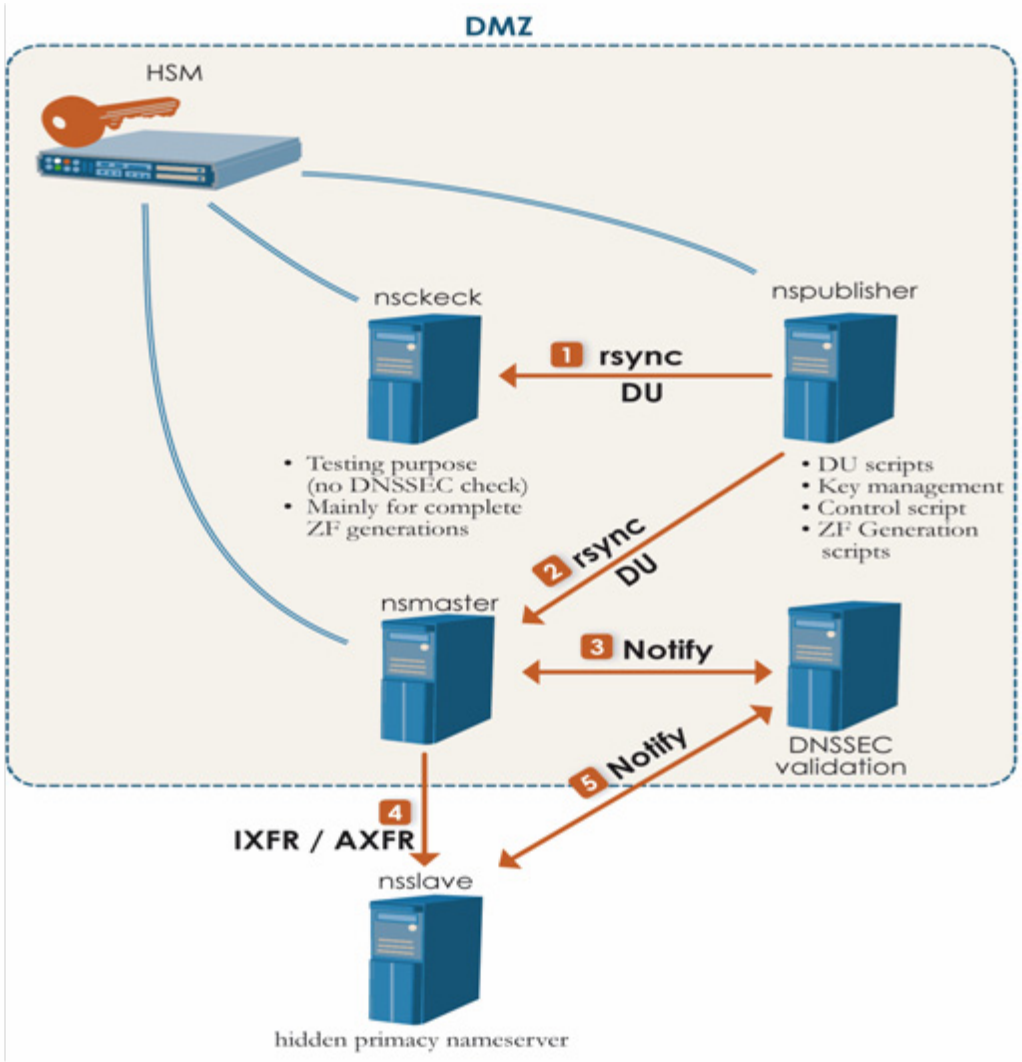

Lessons learned

- **Regarding different aspects, DNSSEC, is still young.**
  - It was still possible to find bugs in the most common DNNSEC tools (Bind, Unbound, NSD, …). The good point is that they are patched very fast.
  - Teams training is essential. Mastering DNS doesn't mean you'll deal with DNSSEC easily.
  - "DNSSEC specialists" are still mandatory when problems occur.
  - Few fieldproven tools available (zone size is often a problem).
- **Keep all zonefile revisions (it would have been impossible to find the bug without that).**
  - With Dynamic Update + Automatic Signing, it's not that obvious.
  - Hopefully we had deployed a zone versioning system few time before the issues. We just missed a version of fr zone.
- **Monitor, monitor and monitor again…**
- **Provide, as fast as possible, transparent information was much appreciated. All details for each outage were published on our public website.**
- **The good news, is… "we are not alone"…**

DNSsexy

- **Ripe also had issues with DNSSEC.**
- **First discussions about the need of a DNSSEC verification tool on OARC mailing-list.**
- **With other interested parties in this topic, a first meeting was held to gather first requirements during OARC meeting.**
- **Informal meeting during IETF-80.**
- **Discussions are still in progress. If interested, join the mailing-list hosted by the NLnet Labs**
  - http://nlnetlabs.nl/mailman/listinfo/dnssexy

Questions…