

DNSSEC Key Deletion Issue

March 14th, 2011

vincent.levigneron@afnic.fr

www.afnic.fr – afnic@afnic.fr

Plan

- **Key numbers of the publication process**
- **AFNIC DNSSEC specifications**
- **Key deletion process in AFNIC zones**
- **Focus on Bind Private Record**
- **First DNSSEC outage in november**
- **Second DNSSEC outage in february**
- **Thanks to the community**
- **What happens next**
- **Lessons learned**
- **Get 3 for the price of 2...**

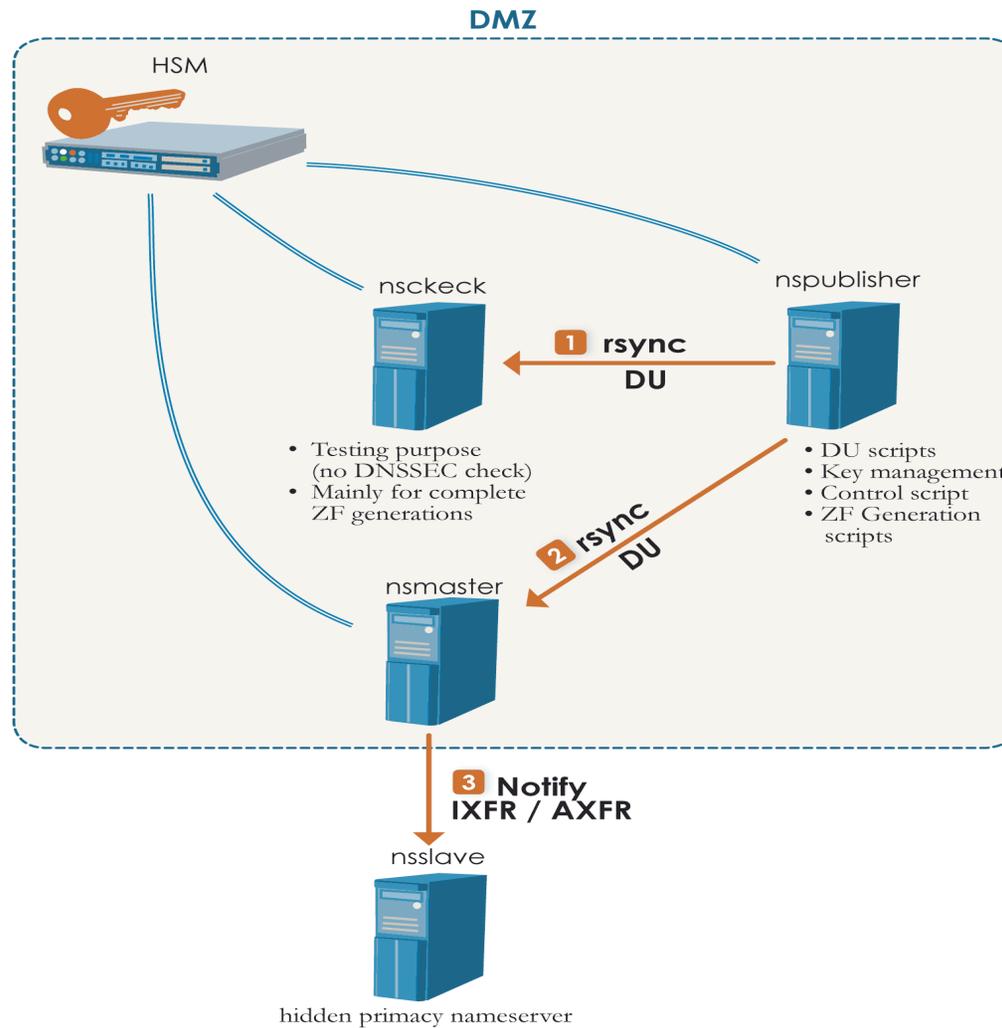
Key numbers of the publication process

- **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 nearly 2 millions domains.**
- **fr zone contains 4 250 000 Ressource Records.**
- **No DS records yet (registration of DS should be launched in one month. EPP/RFC5910 implementation finished, tests in progress).**

AFNIC DNSSEC specifications [1/4]

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

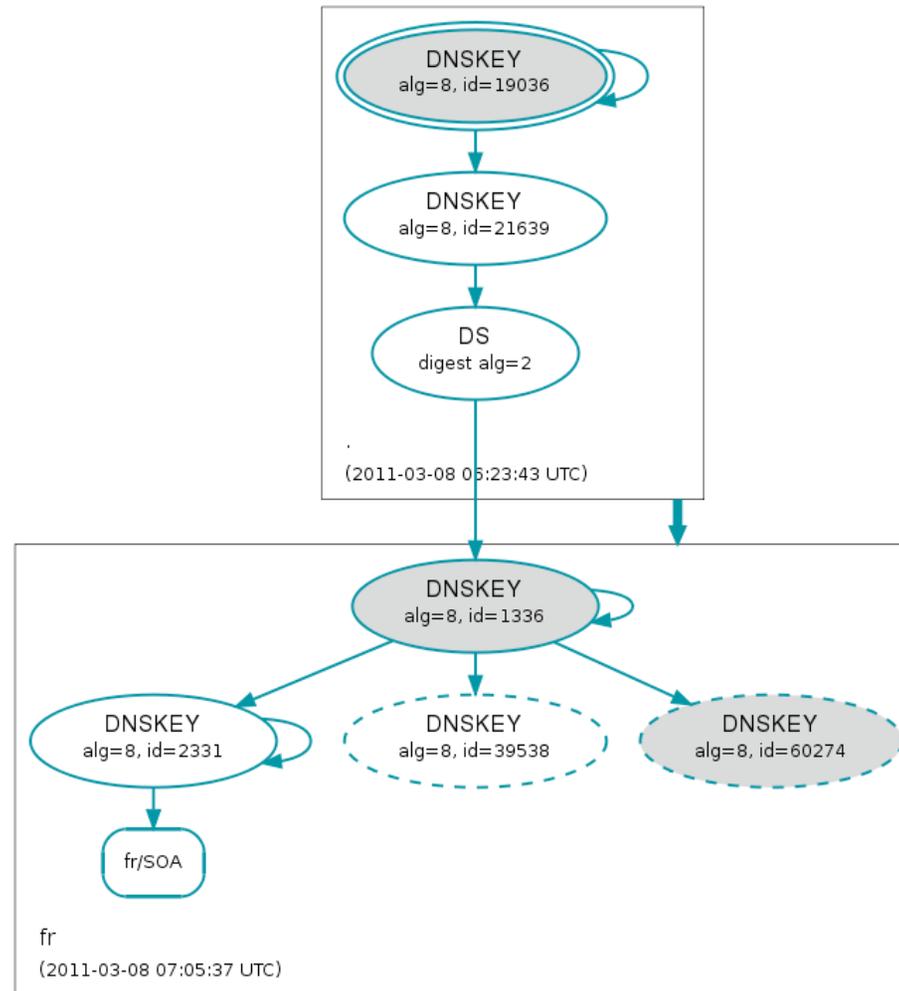
AFNIC DNSSEC specifications [2/4]



AFNIC DNSSEC specifications (closer look on fr zone) [3/4]

- **While there are more than 4 millions records...**
- **... there are only 17 NSEC3 records and 35 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 every hour and once a week there is a complete zonefile generation for administrative purposes.**
- **Dynamic Updates is only used for delegations (NS/A/AAAA).**
- **All key/signature stuff is based on “automatic signing” Bind capabilities (no Dynamic Updates in this case).**

AFNIC DNSSEC specifications (DNSViz on fr zone) [4/4]



Key deletion process in AFNIC zones

- **We use very large timings.**
 - When a key becomes inactive, it is deleted one month later.
 - 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.**

Focus on Bind Private Record

- **Described in ARM (4.9.4 Private-type Records).**
- **5 octets TYPE 65534 record is used to know the state of a signing process.**
 - “If the first octet is non zero then the record indicates that the zone needs to be signed with the key matching the record, OR that all signatures that match the record should be removed“
 - In this case, the final octet indicates when signing is complete (non zero value if it's finished).
- **This record causes us some troubles in what we called the “TYPE65534 Bug”... (more official reference is [ISC-Bugs #23232] 😊)**

First DNSSEC outage in november (quick look on a big mess)

- **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.**
- **Each element, separately, seems obvious, but when all happens at the same time... It's a big mess and it's a lots of confusion to solve all the problems in once...**
- **And Yes, we also had the TYPE65534 Bind Bug 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 (detailed view) [1/4]

- **Situation just before the key is deleted from the zone fr (serial was 2222240887).**
 - The zone signing key with keytag 43893 is still in the DNSKEY RRset but there are no more RRSIG records generated with that key.
 - There are no TYPE65534 record in the zone. There has been a complete zonefile generation few days before and no key operation since this time, so it's a normal situation.
- **Then it's now time for key deletion.**
 - Our script updates corresponding key files and execute an rndc sign zone.
- **This is what we do for each key state transition.**

Second DNSSEC outage in february (detailed view) [2/4]

- **What was expected from our point of view when this key switched from "inactive" state to "delete" state in zone with serial 2222240888...**
 - 1/ DNSKEY RR corresponding to keytag 43893 needed to be removed.
 - 2/ DNSKEY RRset signature needed to be updated.
 - 3/ Serial should be incremented.
 - 4/ SOA signature had to be updated.
- **...In less than the blink of an eye...**

Second DNSSEC outage in february (detailed view) [3/4]

- **We just forgot the TYPE65534 RR that appeared in the zone.**
 - fr. 0 IN TYPE65534 \# 5 08AB750100
 - 08AB750100 → means that signing process with key 43893 is not finished (?!?!?)
 - No RRSIG for this record.
- **The Typemap of the NSEC3 RR corresponding to the Apex should have been modified.**
 - This is not correct but this doesn't prevent the validating resolvers to validate.
- **The BIG issue appears NOW in zone with serial 2222240889 when Bind sign the TYPE65534 record and add TYPE65534 to the Typemap of the previous NSEC3 RR without updating the signature which makes it invalid.**
 - **THE FR ZONE BECAME INACCESSIBLE TO ANY VALIDATING RESOLVERS...**

Second DNSSEC outage in february (detailed view) [4/4]

- **If we had wait, several hours later we would have had a new revision of the zone...**
 - Without TYPE65534 RR.
 - With a NSEC3 with a correct Typemap.
 - With a new signature (a correct one) for this NSEC3.
 - Of course we couldn't wait but this behaviour was confirmed on our lab testbed.
- **We noticed that during this long period, while there are no visible changes in the zone, Bind does lot's of RBT calculation. When this task is over, the zone is valid again.**
- **That's why we didn't notice anything for the other zones. They are very small and this RBT calculation takes seconds, not hours.**

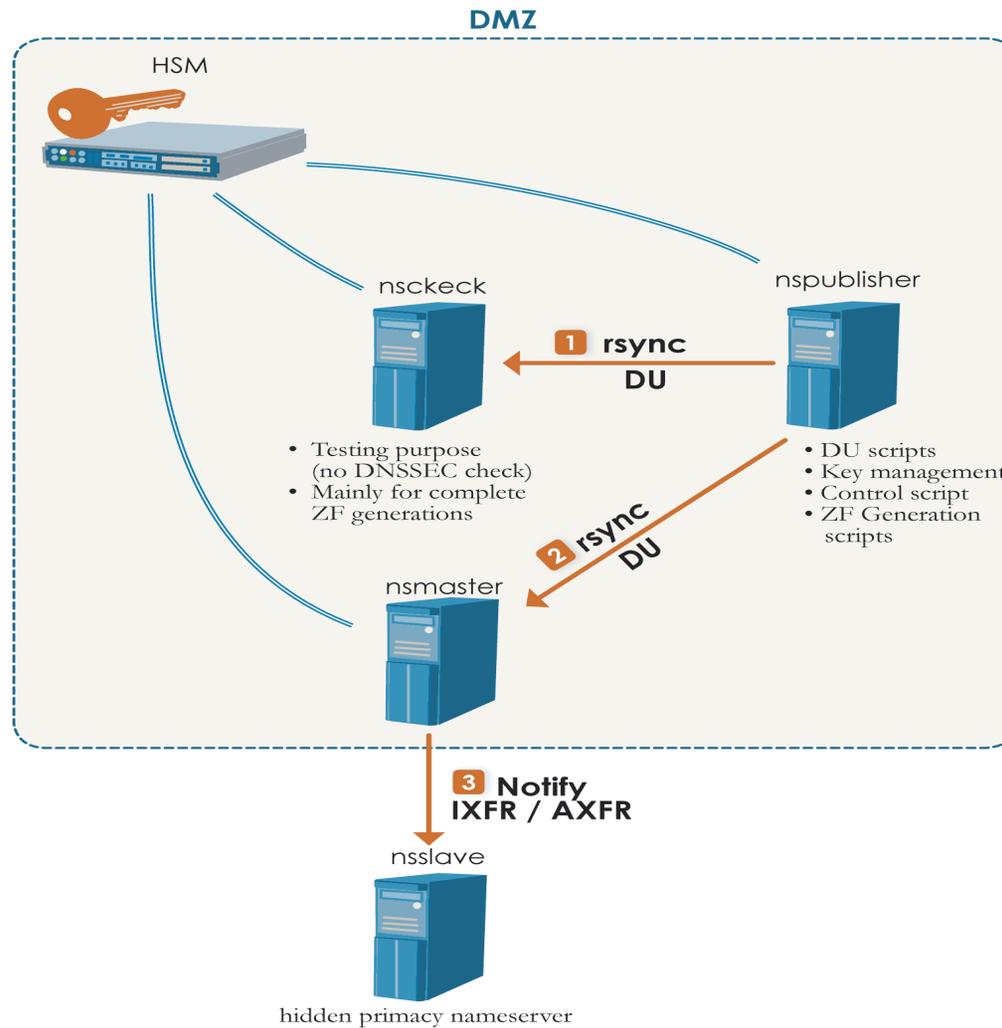
Thanks to the community

- **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). Our tests showed us this bug doesn't exists any longer.
 - We also found a bug in Unbound, the patch should be published soon.
 - We also had good feedbacks on our search for a zone verification tool. Idns (developped at the NLnet Labs) for instance is really promising while not fast enough (for the moment). Not easy to find a tool able to deal with a “big” zone.

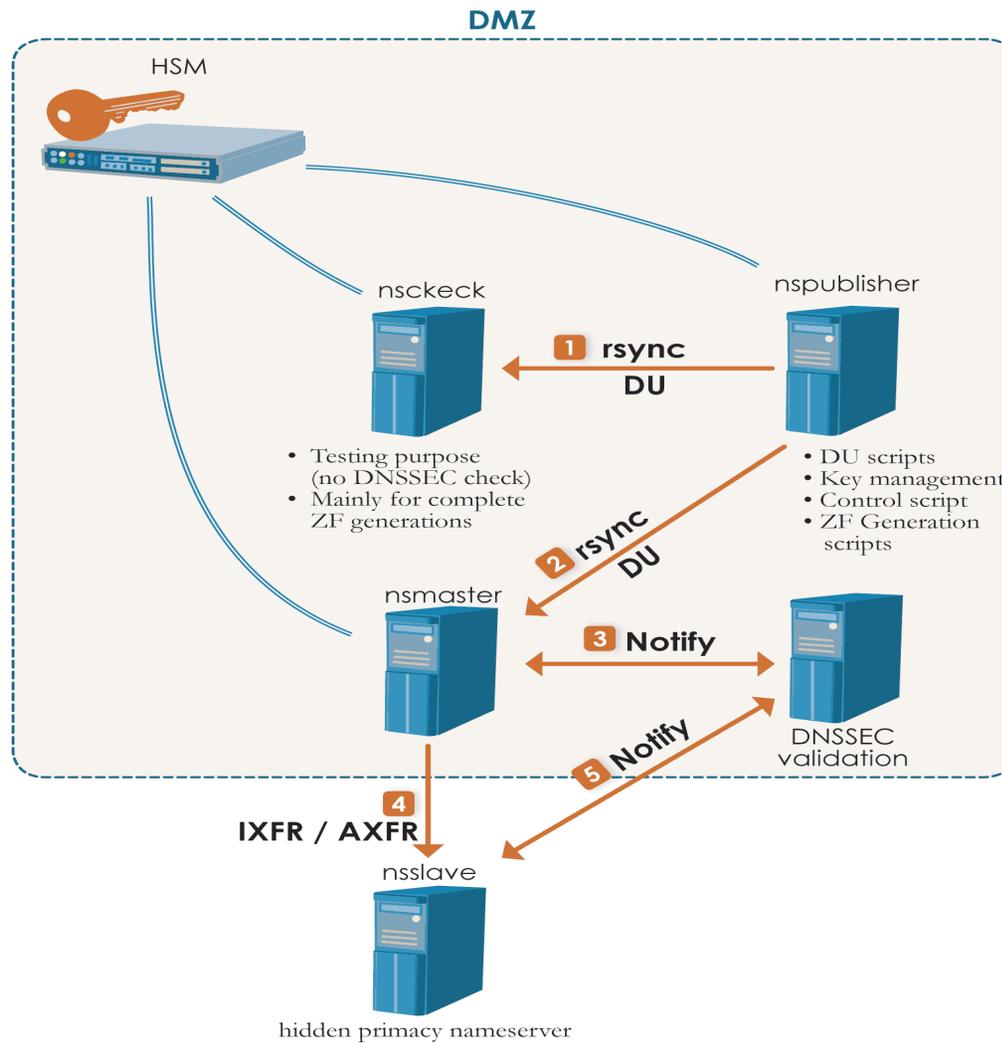
What happens next

- **There will be other issues... Unfortunately...**
- **Patched version of Bind 9.7.3 has been deployed.**
- **We are modifying our system to have a better control over zone changes.**
- **The zone is now validated before it is sent to our hidden master (not yet finished but should be operational before the next fr key rollover deletion phase).**
- **A new Notify Proxy Server controls this.**

Publication System Evolution (Before)



Publication System Evolution (After)



Lessons learned

- **DNSSEC, is still young.**
 - Teams training is essential.
 - “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 issue. We just missed a version of fr zone.
- **Monitor, monitor and monitor again...**
- **Provide as fast as possible transparent information was much appreciated.**
 - <http://operations.afnic.fr/en/2011/02/18/study-and-action-plan-following-the-incident-with-validating-resolvers-on-12-february-2011.html>

Get 3 for the price of 2 (Bonus track) [1/2]

- **First of all, sorry, if we were not focused on your very interesting presentations yesterday, but...**
- **...AFNIC decided to do something very special for this OARC/ccNSO TechDay.**
- **Just for you, we had a “Third DNSSEC outage”...**
- **We are still analyzing it but it seems we have found a new combo.**
- **It's too early to give conclusions but these are some details of the problem.**

Get 3 for the price of 2 (Bonus track) [2/2]

- **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, it seems, we had 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 “seems” it also “decided” to remove all other signatures of the apex ?!?!?

Questions...

