

*afnic*

*ENT was here !!!*

*DNS-OARC 25, Dallas, 10/2016*

*afnic*

# *About Afnic in brief*

- ✓ Non-profit association founded in 1998.
- ✓ Operates 6 ccTLD (.fr/.re/.pm/.tf/.yt/.wf).
- ✓ Back-End Registry for 14 gTLD (.paris/.bzh/...).
- ✓ More than 3 million domain names.
- ✓ DNSSEC was introduced in September 2010.
- ✓ About 10% of domain names have a DS published.

# *DNS/DNSSEC at Afnic in brief*

- ✓ Keys are stored in AEP Keyper HSMs.
- ✓ **OpenDNSSEC** is used to **manage Keys**.
- ✓ **Bind** is used to **sign zones**.
- ✓ Home-made scripts are used to control/synchronize/update/distribute/... zones/keys.
- ✓ Zones are **dynamically updated** (every 10 minutes).
- ✓ NSEC3 with salt changed several times/month.
- ✓ Keys and salt are **not** shared amongst TLDs.
- ✓ **KSK** rollover **every 2 years** (we use standby keys, so 2 DS are present in root but only one KSK is published). **ZSK** rollover every **2 months**.

# *Afnic, we have a problem !!!*

- ✓ Friday, May 13th 2016, we receive a detailed alert from Florian Maury, **ANSSI**.
- ✓ « We can not reach some « `gouv.fr` » domains ».
- ✓ « DNSSEC validation is broken ».
- ✓ « At least one instance of your anycasted server `d.nic.fr` does not reply properly in case of denial of existence answer ».
- ✓ « SERVFAIL is replied when querying « `gouv.fr DS` » ».

# ANSSI Details

- ✓ **A**gence **N**ationale de la **S**écurité des **S**ystèmes d'**I**nformation.
- ✓ The National authority in the area of cyberdefence network and information security.
- ✓ Main missions: Prevention, Defence, Information.
  - ✓ Transversal actions covering French governmental organizations, critical operators and the general public.
  - ✓ Guidance provider on network security topics (DNS, BGP, DDoS prevention, ...).
  - ✓ ...
- ✓ ... Well... I guess this is an alert that we have to consider seriously...

*afnic*

# What we can see with *DNSViz*

### Notices

RRset status

Insecure (2)

Secure (2)

DNSKEY/DS/NSEC status

Secure (9)

Non\_existent (1)

Delegation status

Secure (1)

### Notices

Errors (2)

- gouv.fr/A (NODATA): No NSEC RR(s) were returned to validate the NODATA response. (194.0.9.1, 2001:678:c::1, UDP\_0\_EDNS0\_32768\_4096)
- gouv.fr/AAAA (NODATA): No NSEC RR(s) were returned to validate the NODATA response. (194.0.9.1, 2001:678:c::1, UDP\_0\_EDNS0\_32768\_4096)

### DNSKEY legend

[Full legend](#)

- SEP bit set
- Revoke bit set
- Trust anchor

**See also**

### DNSSEC Authentication Chain

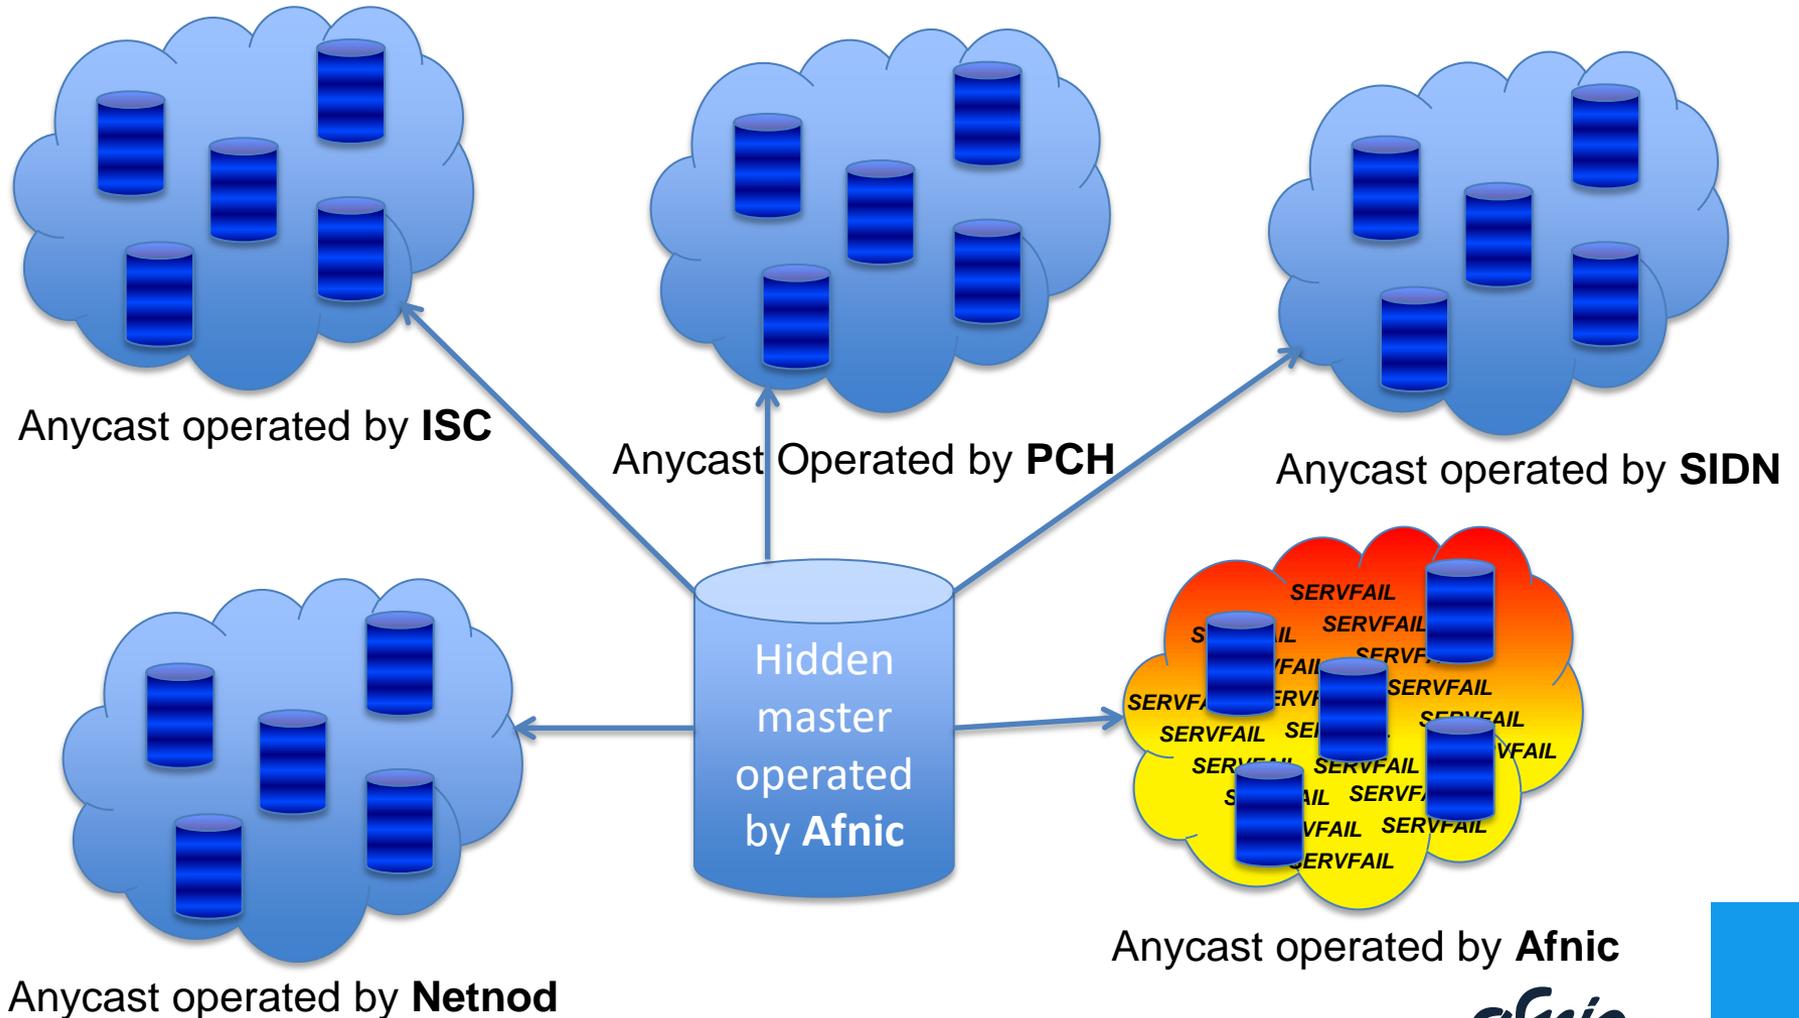
Download: [png](#) | [svg](#)

The diagram shows a DNSSEC authentication chain for the domain gouv.fr. It is divided into two parts: a top part for the parent domain 'fr' and a bottom part for the child domain 'gouv.fr'.  
1. Top part (fr): A trust anchor DNSKEY (alg=8, id=19036) points to a child DNSKEY (alg=8, id=16665). This child DNSKEY points to two DS records (digest alg=2).  
2. Bottom part (gouv.fr): One of the DS records from the 'fr' domain points to a DNSKEY (alg=8, id=20122). This DNSKEY points to two child DNSKEYs (alg=8, id=38025 and alg=8, id=36264). The id=36264 key points to two NSEC3 records, which in turn point to NODATA responses for 'gouv.fr/A' and 'gouv.fr/AAAA'. A central 'fr/SOA' record is also shown. A dashed oval represents a DNSKEY (alg=8, id=35095) that is not part of the active chain.  
Timestamps: (2015-03-09 22:11:48 UTC) for the top part, and (2015-03-09 22:13:35 UTC) for the bottom part.

# *Fix validation issue first...*

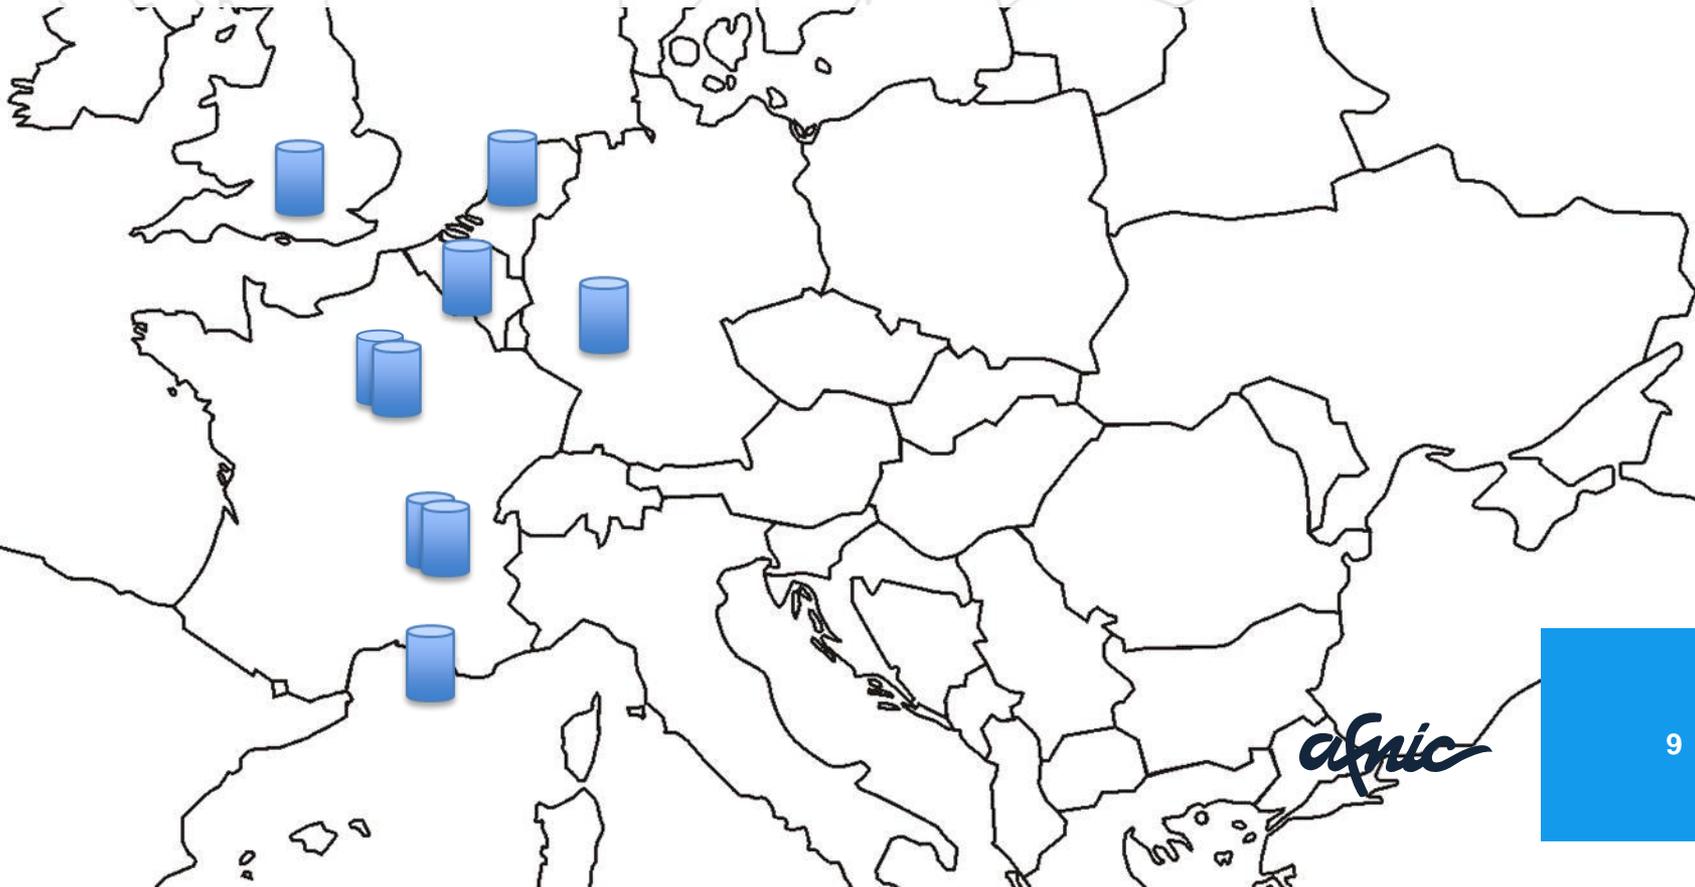
- ✓ First we find out that **all instances** of our anycasted server d.nic.fr have the **same issue**. All of them, but **only** them...
- ✓ Then we realize that all of them run **NSD**. Authoritative servers run by **Afnic** with Bind do not have the issue.
- ✓ We decide to stop all NSD servers and switch to Bind as long as we don't know what is wrong with NSD.
- ✓ This switch solves the validation issue...
- ✓ ... but we dont know yet what was wrong and we are about to discover more issues than expected...

# DNS Configuration of Afnic zones



# *DNS Configuration of Afnic zones*

- ✓ Except for one server located in African region, all instances of our anycasted server are located in Europe.



# Understanding

- ✓ .fr zone is error free... The zone file has been checked by several tools.
- ✓ Why `gouv.fr` ?
  - ✓ « `gouv.fr` » is an Empty Non Terminal domain. In `.fr` and `.re` zones, we have 8 ENT. Only 4 of them have the validation issue.
  - ✓ `.fr` and `.re` zones are signed. NSEC3+Opt-out, keys are in HSMs, zones are signed with Bind, and are dynamically updated. Because it is an ENT, « `gouv.fr` » is also signed.
  - ✓ Broken ENT have **no** signed subzones.
  - ✓ NSEC3 records are missing when these 4 broken ENT are queried on all `d.nic.fr` instances (when they run NSD).

# *NSEC<sub>3</sub>+ENT*

- ✓ Explained in RFC 5155 and 7129.
- ✓ « exact-match » vs « covering ».
- ✓ NSEC<sub>3</sub>+opt-out+ENT is not common. RFC 5155 has an errata just for this case which is not clear in RFC.
- ✓ ... But which one is supposed to be correct ?

# *What RFCs say...*

- ✓ RFC 5155: 7.1. Zone signing

Each empty non-terminal **MUST** have a corresponding NSEC3 RR, **unless the empty non-terminal is only derived from an insecure delegation covered by an Opt-Out NSEC3 RR.**

- ✓ RFC 7129: 5.1. Opt-Out

When using Opt-Out, names that are an insecure delegation (and empty non-terminals that are only derived from insecure delegations) don't require an NSEC3 record.

- ✓ According to this, my understanding is that “covering” scenario seems to be a reasonable one...

# *What RFCs say...*

- ✓ RFC 7129: 5.1 Opt-Out

A recently discovered corner case (see RFC Errata ID 3441 [[Err3441](#)]) shows that not only those delegations remain insecure but also the empty non-terminal space that is derived from those delegations.

[...]

A way to resolve this contradiction in the specification is to always provide empty non-terminals with an NSEC3 record, even if it is only derived from an insecure delegation.

- ✓ ... Well... it seems that “exact-match” scenario is also to be considered !?!?
- ✓ So, for NSEC3+Opt-Out cases, there are two ways to deal with.

# NSEC3 (1/4)

- ✓ NSEC3 « covering » case when it works
  - ✓ Answer contains **APEX SOA record signed** (2 records).
  - ✓ A **signed NSEC3 record** corresponding to APEX (2 records).
  - ✓ A **covering NSEC3 record signed** (2 records).

```
; <<>> DiG 9.10.0-P2 <<>> +dnssec presse.fr
fr.      5400 IN  SOA nsmaster.nic.fr. [...]
fr.      5400 IN  RRSIG SOA 8 1 172800 [...]
1S8KM6Q1E2Q3BLANMJ2MAJ298ICR430G.fr. 5400 IN NSEC3 1 1 1 F8F8832D [...]
1S8KM6Q1E2Q3BLANMJ2MAJ298ICR430G.fr. 5400 IN RRSIG NSEC3 8 2 5400 [...]
DNCS28UODU1BDVAFBFUUG9EKMURU9G55.fr. 5400 IN NSEC3 1 1 1 F8F8832D [...]
DNCS28UODU1BDVAFBFUUG9EKMURU9G55.fr. 5400 IN RRSIG NSEC3 8 2 5400 [...]

> /usr/sbin/nsec3hash F8F8832D 1 1 presse.fr
DNCTF5G2SETQRN9EHDP3CMN68D8UAMFB (salt=F8F8832D, hash=1, iterations=1)
> /usr/sbin/nsec3hash F8F8832D 1 1 fr
1S8KM6Q1E2Q3BLANMJ2MAJ298ICR430G (salt=F8F8832D, hash=1, iterations=1)
```

# NSEC3 (2/4)

- ✓ NSEC3 « exact-match » case when it works (this test was done after because we need to change the signer to have an « exact-match » case).
- ✓ Answer contains **APEX SOA record signed** (2 records).
- ✓ An « **exact-match** » **signed NSEC3 record** (2 records).

```
; <<>> DiG 9.10.0-P2 <<>> +dnssec asso.re
re. 5400 IN SOA nsmaster.nic.fr. hostmaster.nic.fr. [...]
re. 5400 IN RRSIG SOA 8 1 172800 [...]
1LK7IQUQSEKN1IC76FP5032PMVJ4D5N2.re. 5400 IN NSEC3 1 1 1 F8F8832D [...]
1LK7IQUQSEKN1IC76FP5032PMVJ4D5N2.re. 5400 IN RRSIG NSEC3 8 2 5400 [...]

> /usr/sbin/nsec3hash F8F8832D 1 1 asso.re
1LK7IQUQSEKN1IC76FP5032PMVJ4D5N2 (salt=F8F8832D, hash=1, iterations=1)
```

# NSEC3 (3/4)

- ✓ NSEC3 « covering » case when it does not work (we query for A record in this example).
- ✓ Answer contains APEX SOA record signed (2 records).
- ✓ NSEC3 closest encloser is missing.

```
; <<>> DiG 9.10.0-P2 <<>> +dnssec com.re  
re.      5400 IN  SOA nsmaster.nic.fr. hostmaster.nic.fr. [...]  
re.      172800 IN RRSIG  SOA 8 1 172800 [...]
```

# NSEC3 (4/4)

- ✓ Case when it does not work when we query for a DS and Bind is the signer... But in a different way ?!?!?
- ✓ Answer contains **APEX SOA record signed** (2 records).
- ✓ A **signed NSEC3 record corresponding to APEX** (2 records).
- ✓ But... NSEC3 closest encloser is missing.

```
; <<>> DiG 9.10.0-P2 <<>> +dnssec com.re DS
td0qaj9tl88l4h0od5prv6n1pccktlg1.re. 5400 IN NSEC3 1 1 1 F8F8832D [...]
td0qaj9tl88l4h0od5prv6n1pccktlg1.re. 5400 IN RRSIG NSEC3 8 2 5400 [...]
re.      5400 IN  SOA nsmaster.nic.fr. hostmaster.nic.fr. [...]
re.      5400 IN  RRSIG SOA 8 1 172800 [...]

> /usr/sbin/nsec3hash F8F8832D 1 1 re
TD0QAJ9TL88L4H0OD5PRV6N1PCCKTLG1 (salt=F8F8832D, hash=1, iterations=1)
```

# *Bind vs OpenDNSSEC*

- ✓ dnssec-signzone (Bind signer): uses « covering » method.
- ✓ Idns library (OpenDNSSEC signer): uses « exact-match » method.
- ✓ If we use ODS instead of Bind tool, the issue does not exist. But Bind is not wrong and we can not change our DNS/DNSSEC publication system easily (ODS does not support Dynamic Updates for instances). There is no reason to switch to ODS now.

# *Evil Combo*

- ✓ Empty Non Terminal Sub-TLD.
- ✓ NSEC3+opt-out.
- ✓ Bind as a signer.
- ✓ ENT with no signed subzones.
- ✓ **NSD** as authoritative server.

# *Fix the tools*

## ✓ NSD

- ✓ Few days after the issue report, a fix was available in the SVN repository and new version of NSD was released which fixed the bug (NSD 4.1.10 and NSD 3.2.22).

## ✓ DNSviz

- ✓ Did not follow errata 3441 on RFC 5155 (fixed since).

# *Google public DNS issue*

- ✓ Even with **Bind** instead of **NSD**, it does not work if we try to resolve using Google Public DNS 8.8.8.8 server.
  - ✓ We still have a SERVFAIL if « `gouv.fr` » is queried. This is still a validation issue.
  - ✓ But it works when we query « DNS over HTTP » service from Google.
- ✓ We find out that GPD does not work properly when a ENT is queried.

# *Fixing Google public DNS issue*

- ✓ We decide to transform ENT in non-ENT domains.
- ✓ We add a TXT record in all of them (the 8 we have). We are obliged to modify our scripts to deal with that exception.
- ✓ The TXT says « **ENT was here !!!** ».... And IT WORKS, no more SERVFAIL.
  - ✓ NSEC3 « covering » record is replaced by an NSEC3 « exact-match » record.
- ✓ Other TLDs have that kind of record, but we do not know if it addresses the same issue. Like co.jp (TXT « co.jp. »), ...
- ✓ We found several similar ENT in hosted zones, which have been fixed by registries since.
- ✓ Do we have to « fix » all ENTs until GPD has not fixed 8.8.8.8 ?

# Statistics

- ✓ 1365 TLDs in root zone.
- ✓ 11,5% are not signed.
- ✓ 19% use NSEC.
- ✓ 69,5% use NSEC3.
  - ✓ 88,5% use opt-out option (841 TLDs).
- ✓ 841 TLDs may host « broken » GPD ENT...

# *Lessons learned*

- ✓ We need to have different DNS software distributions ready to run on our authoritative servers. It was already initiated, but we decided to increase the priority of this project (« genetic diversity »).
- ✓ We have decided that from now there will be Bind/NSD/Knot-DNS servers on all our authoritative servers.
  - ✓ We are working on a configuration script that will allow us to create all configurations files from a unique source.
  - ✓ We work on processes to switch from one DNS software to another easily.
- ✓ We need to keep the capability to manually handle data in zone files even if they are generated.
- ✓ We need to add specific non regression tests on corner cases when upgrading DNS servers software.

# *Lessons learned*

- ✓ Why query an ENT like `gouv.fr` ?
  - ✓ Zone cut search during resolver chain of trust creation ?
  - ✓ Qname minimisation ? (we know it was not the case here)
  - ✓ Name misspelling in an email, web, ... ?
- ✓ Whatever, it's our clients right to query `gouv.fr`, especially as queries on domains in `gouv.fr` zone had correct answers...
- ✓ We had to fix our servers because we must be able to correctly answer all queries but we don't really know why this old but unknown issue had suddenly such an impact...

# What to do next ?

- ✓ Wait for a fix from Google (issue has been reported #1506).  
"We still are tracking this issue, but unfortunately it has fallen below a number of critical operational issues that need to be addressed to make sure that we can continue to provide the 100% uptime service level objective (for non-DNSSEC-edge-case scenarios, at any rate) that we have maintained so far. I'll certainly update this when we have a chance to work on it."
- ✓ In the meantime if you want to prevent any kind of DNSSEC validation issue to your customers:
  - ✓ Check if you are using NSEC3+opt-out.
  - ✓ Check your zones to find if there are ENTs.
  - ✓ Add fake records to transform ENT into non-ENT or switch to a signer that generates NSEC3 « corresponding » records.
- ✓ `dig @8.8.8.8 unsigned.ent.zft-root.rd.nic.fr`

*Thank you !*

*afnic*

[www.afnic.fr](http://www.afnic.fr)

[Vincent.Levigneron@afnic.fr](mailto:Vincent.Levigneron@afnic.fr)

*afnic*