GOV multi-signer transition with NSEC/NSEC3

Christian Elmerot
Engineering Manager, Authoritative DNS
Cloudflare

Domain migration between providers

Unsigned zones are easy, update NS at registrar/parent and you’re done

Signed zones

- Unsign, update NS, resign
- Owner holds the keys
  - Update NS - Usually secondary transfer/pre-signed
  - Multi-signer DNSSEC (model 1) - provider holds ZSK
- Provider holds the keys
  - (Export/import keys, update NS)
  - Multi-signer DNSSEC (model 2)

RFC 8901, Multi-Signer DNSSEC Models

Not your average domain migration

Transitioning a TLD, like GOV, between providers requires additional considerations

- Migrating registry and related functionality
- Nameserver transition
- DNSSEC transition

Going unsigned is not really a realistic option!

Coordination and cooperation is key

SLD domain owners can “force” through a transition - on a technical level with little to no issues

TLD transitions holds the potential to disrupt all the delegated domains unless done with consideration - in coordination between the providers

Good coordination requires good cooperation

Multi-signer DNSSEC transition

Migrating a TLD is between providers is essentially a transitioning into multi-signer DNSSEC configuration and subsequent transition out to single signer again

RFC 8901, Multi-Signer DNSSEC Models, model 2 for secure zone transfers

- Incoming provider's DNSKEYs are added
- Incoming provider's DS is added to the parent (root)
- Joint operation can begin by adding the new nameservers to the set

... reverse for the outgoing provider's NS, DS, and DNSKEYs

... except this requires both providers to use the same DNSSEC configuration

DNSSEC configuration mismatch

Going into this process Verisign used RSA/SHA256 (algorithm 8) with NSEC3 for the GOV zone

Cloudflare use ECDSA P-256 (algorithm 13), live signing, with NSEC and Compact Denial of Existence

Not a single common parameter there to begin with

DNSSEC work needed

We (Cloudflare) needed to come up with a realistic proposal for how to go forward

Options:

1. Roll algorithm and use NSEC/NSEC3 during transition
2. Switch GOV to NSEC prior to migration, roll algo during the process
3. Switch GOV to algo 13 prior to migration, use NSEC/NSEC3 during transition
4. Switch GOV to algo 13 and NSEC prior to migration, prior to migration
5. We add support for signing with algorithm 8 (Cloudflare)
6. 5 + Add support for NSEC3 (Cloudflare)

Staring at the road signs

Least amount of work rolling algorithm and dual use NSEC/NSEC3 during the transition

Does resolving the GOV zone still work in such a scenario?

Testing needed

Test dual provider using different algorithms with NSEC/NSEC3

Test dual provider using NSEC and NSEC3 at the same time (same algorithm)

Test the waters

Test resolvers in the wild using RIPE Atlas probe measurements

Important to note that there is always a small amount of background noise so it’s important to establish a baseline.

Baseline was established using DO=0

Major thank you to RIPE staff allowing the use of many more probes to perform measurements for this than the default allowed

Sink or swim - the results

Algorithm roll during transition

Tests performed: positive answer, NODATA answer, NXDOMAIN lookup

Each provider uses either algo 8 or 13 to sign

SERVFAIL 0,49% (Baseline ~0,07%)

# of Probes with complete failures (all configured resolvers SERVFAIL) higher than other setups

So what was failing?

Look closer into the SERVFAIL results...

… oops my own Atlas probe SERVFAIL

So looking closer at the SERVFAIL results I find probe #11603
… my own probe’s resolvers SERVFAIL!

Slightly embarrassing but good for debugging!

Unbound, with configuration option: `harden-algo-downgrade` enabled
will SERVFAIL when RRSIGs for all configured DNSKEY algorithms are not present in the response. All in RFC compliance.

A bit too many resolvers use this configuration so algo roll during the transition was out as an option

NSEC+NSEC3 results

SERVFAIL results 0,13% only slightly higher than baseline
Very few probes fail on all resolvers

Only negative answers can see impact

Results indicate a valid configuration to use during the transition

Agreed DNSSEC configuration during transition

- Cloudflare implements live signing using algorithm 8.
  Algo rollover pushed until after transition is completed
- Both NSEC and NSEC3 in use by respective provider at the same time!
- RFC 8901 supports this:
  `Since authenticated-denial responses are self-contained, NSEC and NSEC3 can be used by different providers to serve the same zone.`

Transition successfully completed

- No reported issues during the multi-signer stage of the transition
  To our knowledge, this is the first time a domain of this significance has transitioned using NSEC+NSEC3
- Successful completion of the transition thanks to the testing, cooperation, and coordination of Verisign's and Cloudflare's teams together with CISA
- GOV zone currently fully served by Cloudflare's platform using live signing with algorithm 8 and NSEC Compact Denial of Existence
- Algorithm roll from 8 to 13 to happen soon (Q1 2024)

TL;DR

- NSEC+NSEC3 during domain transitions work well and was successfully used during the .GOV TLD transition
- Rolling algorithm during a domain transition needs community work.
  RFC update is required by not requiring answers to contain signatures from all configured algorithms as long as a supported and valid signature is found

  This is also vital to improve multi-signer operations

Verisign NSEC/NSEC3 Tests

- Single second level domain whose signing we control
- Authoritative name servers
  - A – BIND
  - B – BIND
  - C – Cloudflare
- Experiments

1. A & B sign with NSEC3
2. A & B sign with NSEC3, but B has invalid signatures
3. A with NSEC3, B with NSEC
4. A with NSEC3, C with NSEC compact denial of existence

- Resolvers
  - Google Public DNS
  - Cloudflare Public DNS
  - Quad9 Public DNS
  - Neustar Public DNS
  - BIND 9.11.4
  - Unbound 1.6.6
  - PowerDNS Recursor 4.1.16
  - Knot Resolver 5.3.2
  - RIPE Atlas probes
- Queries
  - 50% NOERROR
  - 50% NXDOMAIN

Results

- Experiment 1
  - Shows “before” state and verifies test validity
  - Even distribution of queries to authoritative servers
  - Small amount of SERVFAIL responses from RIPE Atlas
- Experiment 2
  - Invalid signatures on one server, verifies technique
  - Google, Cloudflare, Quad9, PowerDNS, RIPE Atlas resolvers have increase in SERVFAIL.
  - Others retry and return NOERROR.
- Experiment 3
  - Mixed NSEC/NSEC3
  - Good distribution between A and B name servers
  - Good distribution of NSEC and NSEC3 in responses
- Experiment 4
  - Introduce compact denial of existence
  - NXDOMAIN responses become NODATA