Comparing DNS zones

Stéphane Bortzmeyer
<stephane+blog@bortzmeyer.org>

First publication of this article on 8 September 2010

https://www.bortzmeyer.org/canonicalize-zones.html

For a long time, network administrators have been comparing DNS zone files, to see the changes. A specially well-watched is the root zone, which is available on line <ftp://rs.internic.net/domain/root.zone.gz>. Now that many zones, including the root, are signed with DNSSEC, how to compare them meaningfully ? We certainly do not want a simple resigning to trigger a difference.

The issue was for instance raised by Paul Hoffman on the dns-operations mailing list <https://lists.dns-oarc.net/pipermail/dns-operations/2010-August/005987.html>. Many possible solutions have been suggested in the resulting thread.

There are two ways to strip the zone file of the things that one can see as ”meaningless”. One is converting it to a canonical format and one is to strip the things you are not interested in. We'll see that you need both, in order to address all the cases.

OK, first, the canonicalization. There are two tools to do so. To test them, let's first create a simple DNS zone file (en ligne sur https://www.bortzmeyer.org/files/example-01.db) for the TLD .example. Now, let's try the tool shipped with BIND :

```
% named-compilezone -i none -o example-01.db-canonical example example-01.db
example-01.db:3: no TTL specified; using SOA MINTTL instead
zone example/IN: loaded serial 2010090804
dump zone to example-01.db-canonical...done
OK
```

We can check the new file (en ligne sur https://www.bortzmeyer.org/files/example-01.db-canonical) and/or see with diff the changes : comments removed, SOA record on one line, TTL made explicit, @ expanded, etc. More important, the canonical version is now independant of small changes (such as added or removed white spaces) in the ”source” file. But now let's use a zone file signed with DNSSEC (en ligne sur https://www.bortzmeyer.org/files/example-02.db.signed). The canonicalisation normalizes the DNSSEC records like the RRSIG but do not suppress them, which means that a simple resigning of the zone will produce a huge (and probably meaningless) diff. We should add a filtering step or try another tool.

The excellent ldns <http://www.nlnetlabs.nl/projects/ldns/> library has a tool named ldns-read-zone which can perform canonicalization :

```
% ldns-read-zone -c -z example-01.db
...
```

Do note that this tool (unlike `named-compilezone`) cannot get the zone name from the outside, it has to be in the zone itself. The above example does more or less the same thing than `named-compilezone` but there is another option, `-s`, which strips DNSSEC records. Then, we have a zone file which is almost usable for diff : the only exception is the SOA record which, in some cases (like the current root) changes every day even if there is no actual change. Finally, we have to use filtering.

Filtering can be done alone but it may miss some transformations in the zone file. Here is an example of `grep` and a big regexp to ”clean” a zone file :

```
% grep -E -v ’;(File (start|end)|(End of file)|(serial))|ˆ[ˆ[:space:]]+[[:space:]]+[0-9]+[[:space:]]+IN[[:sp
```

You can find other regexps in the thread on dns-operations mentioned above.

My final strategy was to use canonicalization + filtering, with a much simpler regexp, just to remove SOA records. My shell script looks like :

```
CLEAN_REGEXP=’ˆ[ˆ[:space:]]+[[:space:]]+[0-9]+[[:space:]]+IN[[:space:]]+(SOA)[[:space:]]’
...
# Canonicalize and strip DNSSEC transient data (such as signatures)
gunzip --to-stdout root.zone.gz | \
    ldns-read-zone -s -c -z > root.zone-canonical.new
grep -E -v $CLEAN_REGEXP root.zone-canonical > $TMPROOTOLD
grep -E -v $CLEAN_REGEXP root.zone-canonical.new > $TMPROOTNEW
if [ ! -z "‘cat $DIFF‘" ]; then
    # New version, do something
```

This is the script which is behind the root zone history published as the canonical version `<https://viewvc.generic-nic.net/viewvc.cgi/NIC-generique/iana/root.zone-canonical?root=R%26D&view=log>` and the original unmodified version `<https://viewvc.generic-nic.net/viewvc.cgi/NIC-generique/iana/root.zone?root=R%26D&view=log>`.

Other possible tools include, again in ldns, `ldns-compare-zones`, a zone file comparison tool, but which do not seem to allow ignoring DNSSEC records, and yazvs `<http://yazvs.verisignlabs.com/>` (which I did not test since it depends on a Perl module which is not yet in Debian).

Thanks to Ray Bellis, Marco Davids, Ond[Caractère Unicode non montré [1]]ej Sur[Caractère Unicode non montré], Joe Abley and Duane Wessels for ther ideas, suggestions and regexps.

[1] Car trop difficile à faire afficher par LaTeX