Evolving the Root Zone technical checks

Kim Davies
Director, Technical Services

ICANN 53, Buenos Aires, Argentina

The basics

- ICANN conducts a set of technical checks for each zone change (e.g. root zone)
- These are repeated at several intervals throughout the life of a change request
- All current tests are fully automated
- Any issues identified are reported to customer, and they are asked to remedy them
- Failed tests are automatically repeated every few hours, or customers can force a re-test
- Customers can ask to proceed despite a specific failed check by providing rationale to IANA staff

    Subject matter expert internally reviews such requests to see if they make sense

How we got here

Current set of technical checks are the result of public consultation in 2006

https://www.icann.org/news/announcement-2006-08-18-en

Community contributed feedback, including both ccTLD and gTLD registries

Current set of requirements: http://iana.org/help/nameserver-requirements

Codified into Root Zone Management System (RZMS) and support tools

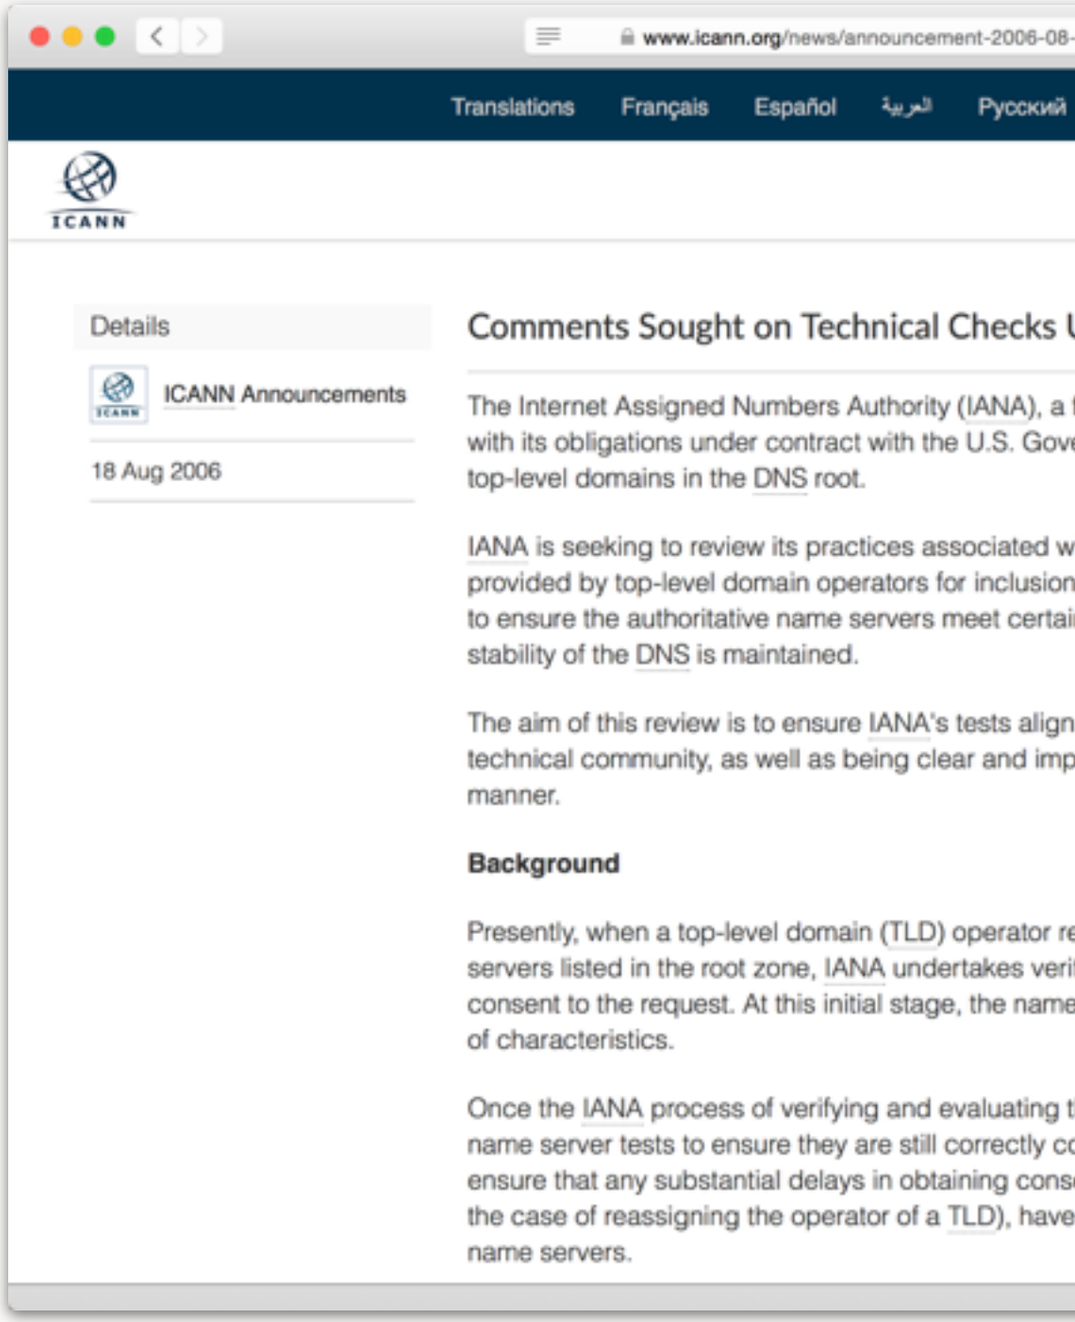

The current test suite

Current tests (The Basics)

- **Minimum 2 nameservers**
  … that don’t share IP addresses

- **Valid hostnames**
  … that comply with RFC 1123

- **Answer authoritatively**
  … must respond with the AA-bit set to the apex of the child zone

Current tests (Network connectivity)

- **Nameservers must be reachable**
  … must respond over port 53 using both UDP and TCP

- **Network Diversity**
  … must be in two topologically separate networks, defined as not sharing the same origin AS. Assessed through inspection of routing tables (RIPE RIS, Cymru, etc.)

- **No prohibited networks**
  … must not be tunnels, private networks, etc.

Current tests (Consistency)

- **Consistency between zone glue and authoritative zone**
  IP addresses for glue in parent must match A/AAAA in authoritative zone for hosts

- **Consistency between delegation and zone**
  NS set for listing in parent must list NS set for listing in apex of child

- **Consistency between authoritative name servers**
  Each authoritative name server for zone must return same NS and SOA at apex

Current tests (Prevent other breakage)

- **Referrals do not truncate**

  Parent referrals must fit on a 512-byte packet (i.e. non-EDNS0 UDP packet limit). Payload must fit the maximum QNAME, plus the complete NS set, plus at least 1 glue record for each supported transport

- **Don’t provide open recursive name service**

  Don’t answer to queries you aren’t authoritative for.

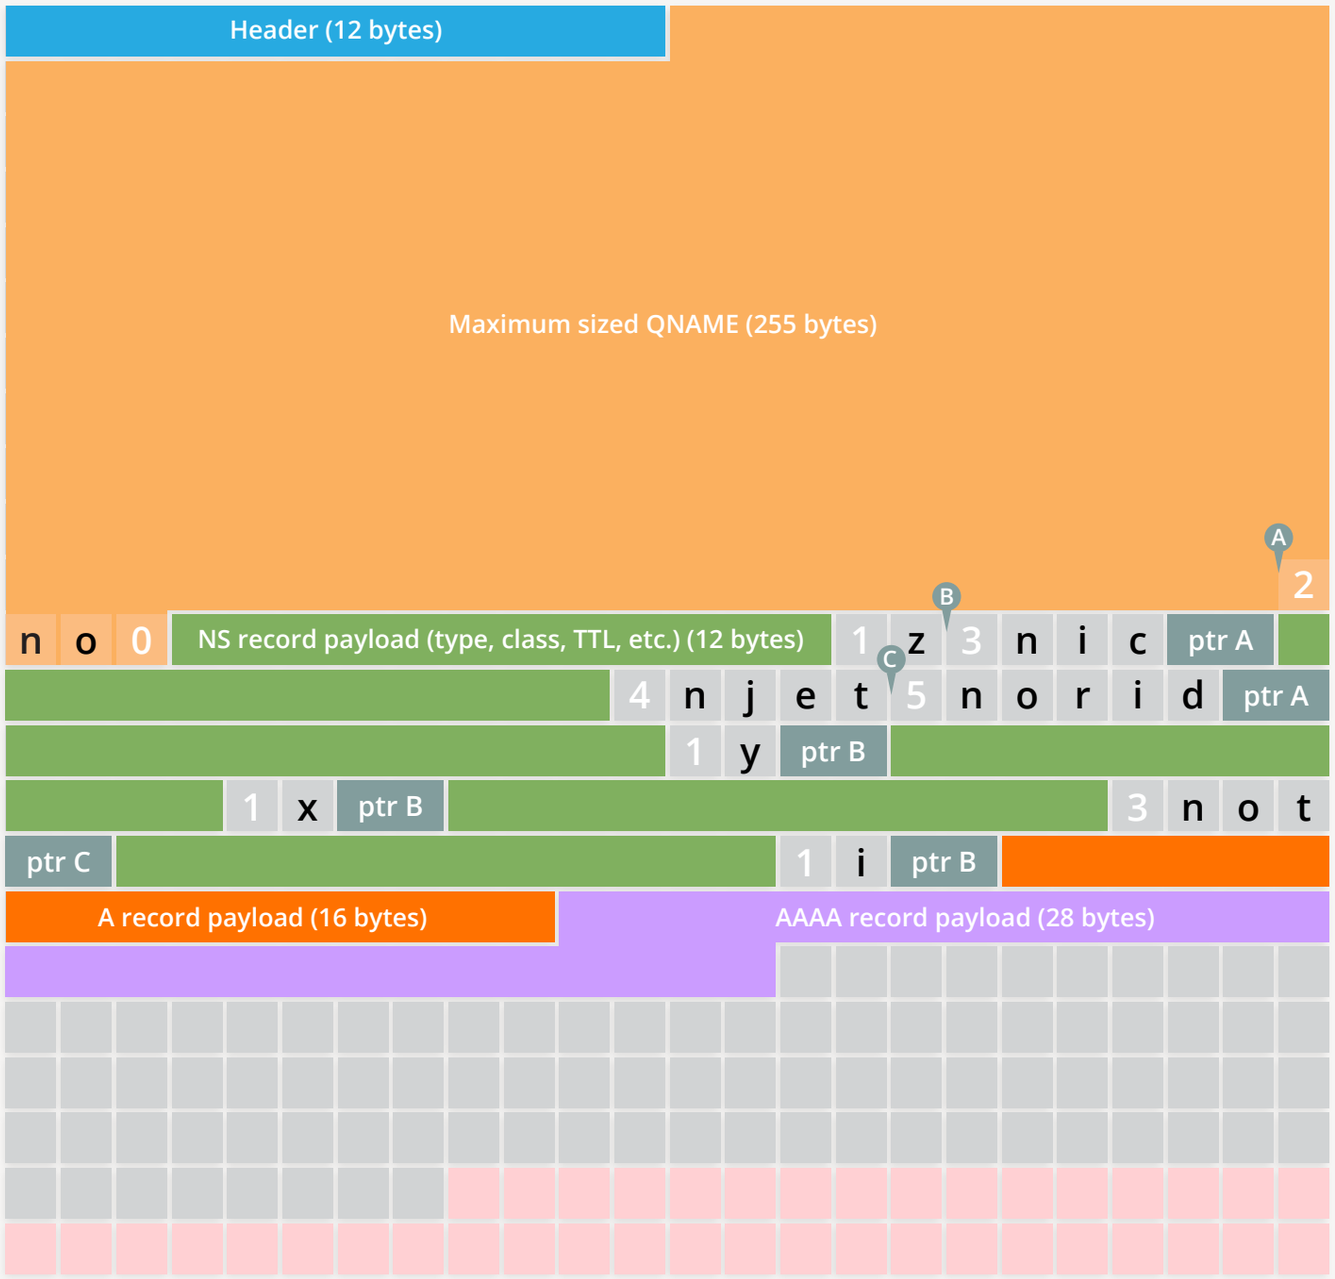

Current tests (DNSSEC)

- **DS record format**
  Hash of correct, length, type etc. Must be a supported type.

- **Matching DNSKEY**
  Must have a DNSKEY in zone apex that matches each DS record provided

- **Validation of RRSIG**
  Validate the RRSIG for the apex of the zone using the DS record set

Things we've seen

Network Diversity

Increasingly seeing a TLDs name server infrastructure operated by a single party

> Working assumption 10 years ago is it is good practice to have at least two distinct vendors for resiliency.

Appeal is often “it uses Anycast, so it’s OK”

> Not just seeking to protect against failure in the physical topology, but things like broken announcements and business failure

Some vendors obtain a second AS operated by same party as the first, nominally meeting diversity test

Consider the need to identify unskilful operators that put everything in one basket

DS record issues

TLDs wishing to list inactive “standby” DS records

- Purports to be an off-line key that would be switched in an emergency
- Can not be verified against a matching DNSKEY
  - Base assumption has been all root zone data can be correlated/confirmed with other data in the DNS
  - IANA has had invalid standby keys submitted, explicitly confirmed by TLDs as being valid, to be identified as invalid afterward

DS records pointing to keys without the SEP-bit set

- Validates fine, meets our rules, but is it what they really wanted to do?
- Upon querying the customer, answer was “yes”
- In the cases where this has been submitted, customer has been notified and decided to proceed.

SOA consistency

Zones that change too quickly, and propagate too slowly, to ever see it in a fully coherent state

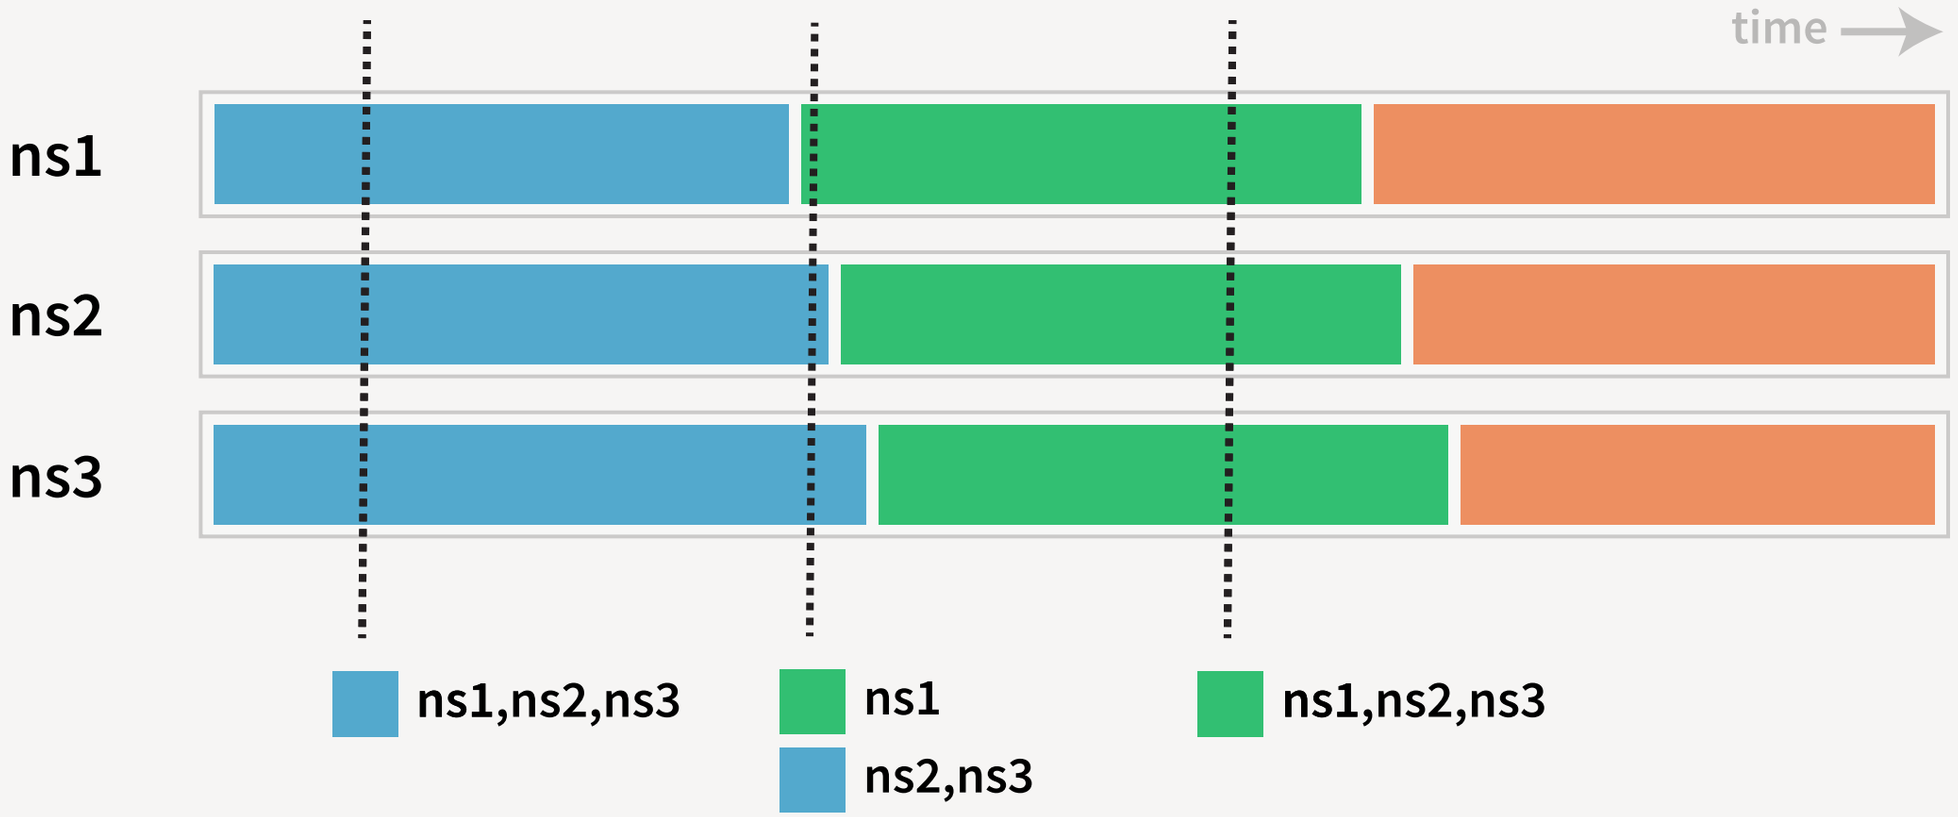

SOA consistency

Zones that change too quickly, and propagate too slowly, to ever see it in a fully coherent state

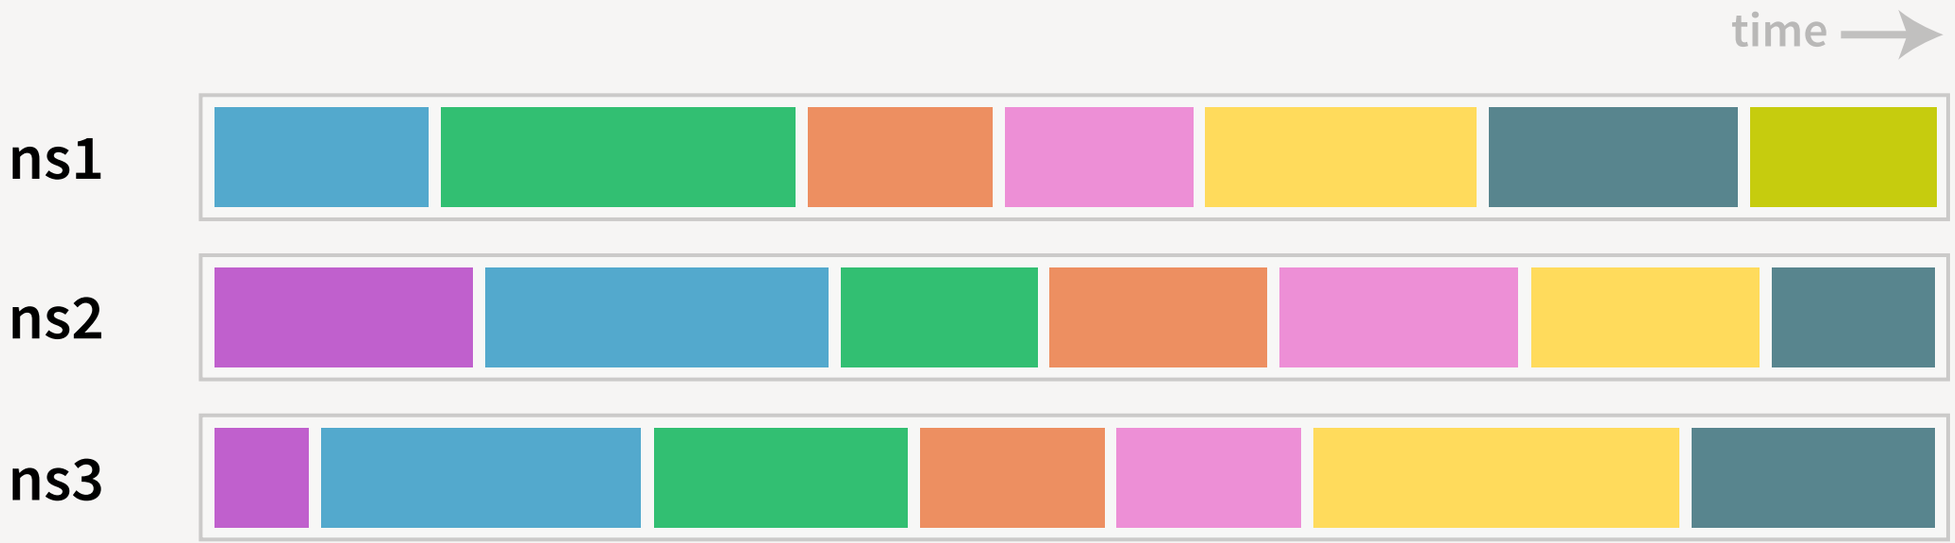

Other feedback

Expand tests to check for protocol compliance

> “ICANN should be testing and blocking [TLDs] until these blocks are removed.”
>
> “We have ICANN checking query rates and uptimes but not protocol basics (like answering all non meta query types) prior to letting new TLDs go live. … ICANN and the TLDs should be showing leadership in this area.”

Treat IPv4 and IPv6 the same

> IPv6 currently optional in IANA tests, but mandatory for gTLDs per contract

What could we do?

Remember

Checks need to accommodate all top-level domains, regardless of skill level

- These checks represent the only place to have a minimum level of technical compliance applied across all TLDs
- Many TLDs have no SLAs or other agreements with ICANN
- Some TLDs still have their entire infrastructure sitting in a single room

Revise technical checks?

Anticipate a public comment period soliciting structured feedback, similar to 2006

Some specific ideas to consider

1. How to test for “loose coherence” in a fully automated way?
2. Is there an improved network diversity test that allows single origin AS?
3. What is proper expectation for DS records and standby keys?
4. Add support for more DNSSEC algorithms?

   … or skip testing requirement for unimplemented DNSSEC algorithm/hash types?

Introduce technical check waivers?

Identify checks that may be waived

Only a subset of checks are potential candidates for allowing a TLD to skip the particular test

Provide a mechanism for TLDs to put a waiver on file

Noting the risks and opt-out reason

Update RZMS

Skip over tests?
Make them non-blocking or skippable?

Apply for permanent waiver

Certain technical configurations will often fail our technical checks. If you have a c
that regularly fails the technical checks, you may opt to have us automatically skip
tests. Choosing these permanent waivers should be considered carefully as enab
can mask legitimate problems that we are trying to identify to ensure the stable o
your domain.

Permanent waivers

X **Waive serial coherency check**

Waive this requirement if your technical configuration updates the zones so regularly that the entire set is not never fully synchronised. Only registries that update their zones multiple times per minute need to consider this option. Using this option on a zone that updates less regularly will mask problems with your zone propagation.

☐ **Waive DNSKEY must match DS record**

Waive this requirement if you list standby keys in the root zone which are not represented in the apex of your zone. Using this option gives us no way of verifying your DS record is valid. Use with extreme care.

Improved implementation with clearer communication

System output can be obtuse/ insufficient

Rewriting the whole architecture of the technical check process to support better reporting of issues identified

Clearer output via email and web

Verbose debug logging of test runs available for TLDs to access via self-service portal

Remove reliance on third-party tools (weird recursor caching bug, etc.)

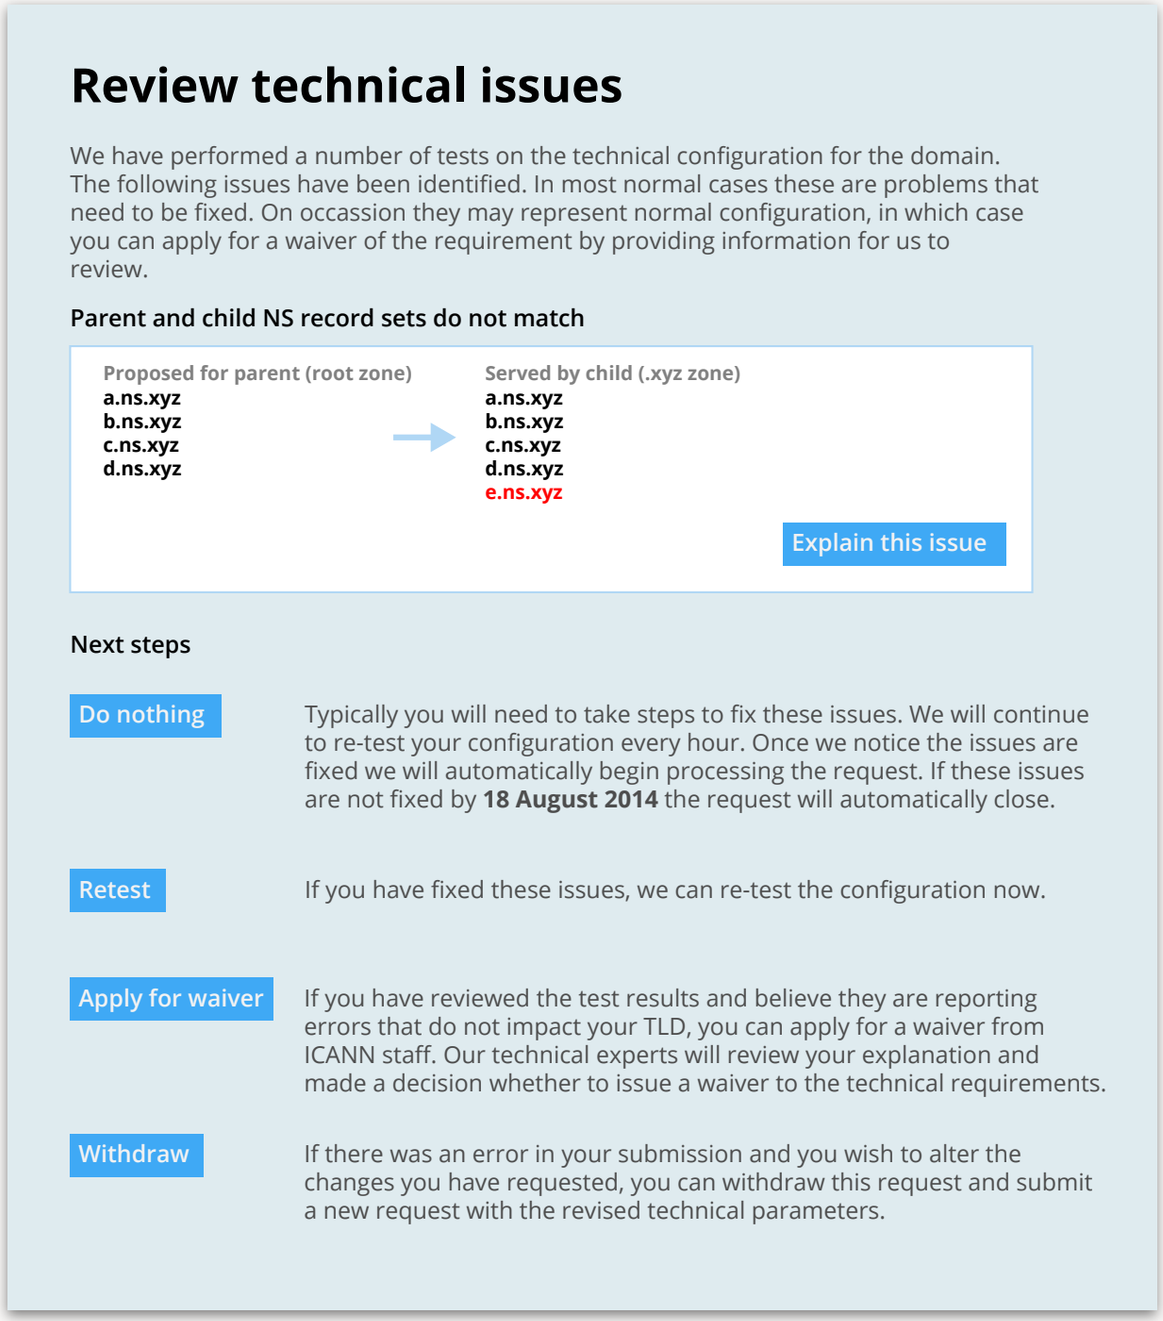

Notification of issues

IANA can regularly perform checks for TLDs

- Notify TLDs of new issues as a courtesy
- Provide link to easily trigger a pre-populated change
- Manage notifications via self-service portal

Provide self-checking tools

- Provide reference implementations of checks
- Our code and/or profiles for Zonecheck etc.

```
$ lincoln ar
Domain: ar (parent: .)
Current authorities from parent (from zone file with serial 2015062001):
    - a.dns.ar (200.108.145.50, 2801:140::10)
    - ar.cctld.authdns.ripe.net (193.0.9.59, 2001:67c:e0::59)
    - c.dns.ar (200.108.148.50, 2801:140:10::10)
    - ctina.ar (200.16.97.17)
    - ns2.switch.ch (130.59.138.49, 2001:620::1b:5054:ff:fe74:8780)
Proposed authorities:
    (None provided, assuming no change.)
Changes between current and proposed:
    (No difference between current and proposed.)
Siblings affected by glue changes:
    (None.)
Issue summary:
    1. ns2.switch.ch[2001:620::1b:5054:ff:fe74:8780/udp] did not respond to queries, nameservers must be reachable (2.3)
$5 2015-06-20T20:57:09.256341
$
```

Other ideas

RFC 7344 (CDS/CDNSKEY) support

Poll for keys, triggers invitation to create a matching change request

Skipping supplemental technical check

Has the second test become superfluous?

Can retest only if longer than *x* days since first test

Self-service testing

Open implementation

Community requests for more testing

Could be informational (non-blocking)

What's next?

Good ideas welcome at any time

kim.davies@icann.org

Public comment period

Structured feedback mechanism to provide evidence of evolution required

Technical work already underway in RZMS

Major revision underway to implement various other changes

Plan to implement new technical checking platform in that release

Thanks!