CLOUDFLARE®

Keeping DNS parents and children in sync at Internet Speed!

Ólafur Guðmundsson
olafur@cloudflare.com

How long does it take to ?

- Post a new selfie on Facebook and all your friends to be notified
    - **few seconds <= INTERNET SPEED**
- **For a new domain to appear in the DNS?**
    - less than 5 minutes in ICANN TLD's, random in others
- Move domain from one DNS operator to another?
    - long time limited by MAX(Parent NS TTL, Child NS TTL)
- Transfer a domain from one registrar to another one?
    - 1 sec … 5 days

WHO ?

- Who can change the delegation information in parent?
  - The registrant, and registrar when registrar is also DNS operator
  - Outside TLD registrations==> organizational policies apply.
- Who gets blamed when things do not work as expected?
  - The entity closest to “customer”
- Who is at fault ?
  - Publisher or publisher agent

Recent example: HBOnow.com

- Affected: Customers behind DNSSEC validating DNS resolvers
- Blamed: Comcast and ISP’s for resolution failure i.e. blocking
- Root cause: HBO for not checking the domain was DNSSEC bogus
- Time to full recovery: 1 day to purge DS from all caches after HBO made a change in .com registration system
- Mitigation: temporary enable negative trust anchor by resolvers operators
- Side effect: Lots of non-polite Facebook and Twitter posts

Third party DNS operators or 3-DNS

Third party DNS operator (3-DNS)

- Definition: An entity contracted by “owner” of the domain to operate DNS on their behalf.
- Who: 3-DNS Operators include CDN’s, DNS specialists, Appliance vendors, friends, etc.
- Millions of domains are operated by 3-DNS
  - Many “important” domains are operated by 3-DNS
  - Some domains use vanity DNS server names, but routing/traceroute do not lie :-)

Domain Registry model:

- Includes Registries, Registrars, Resellers and Registrants.
- when developed did not envision 3-DNS

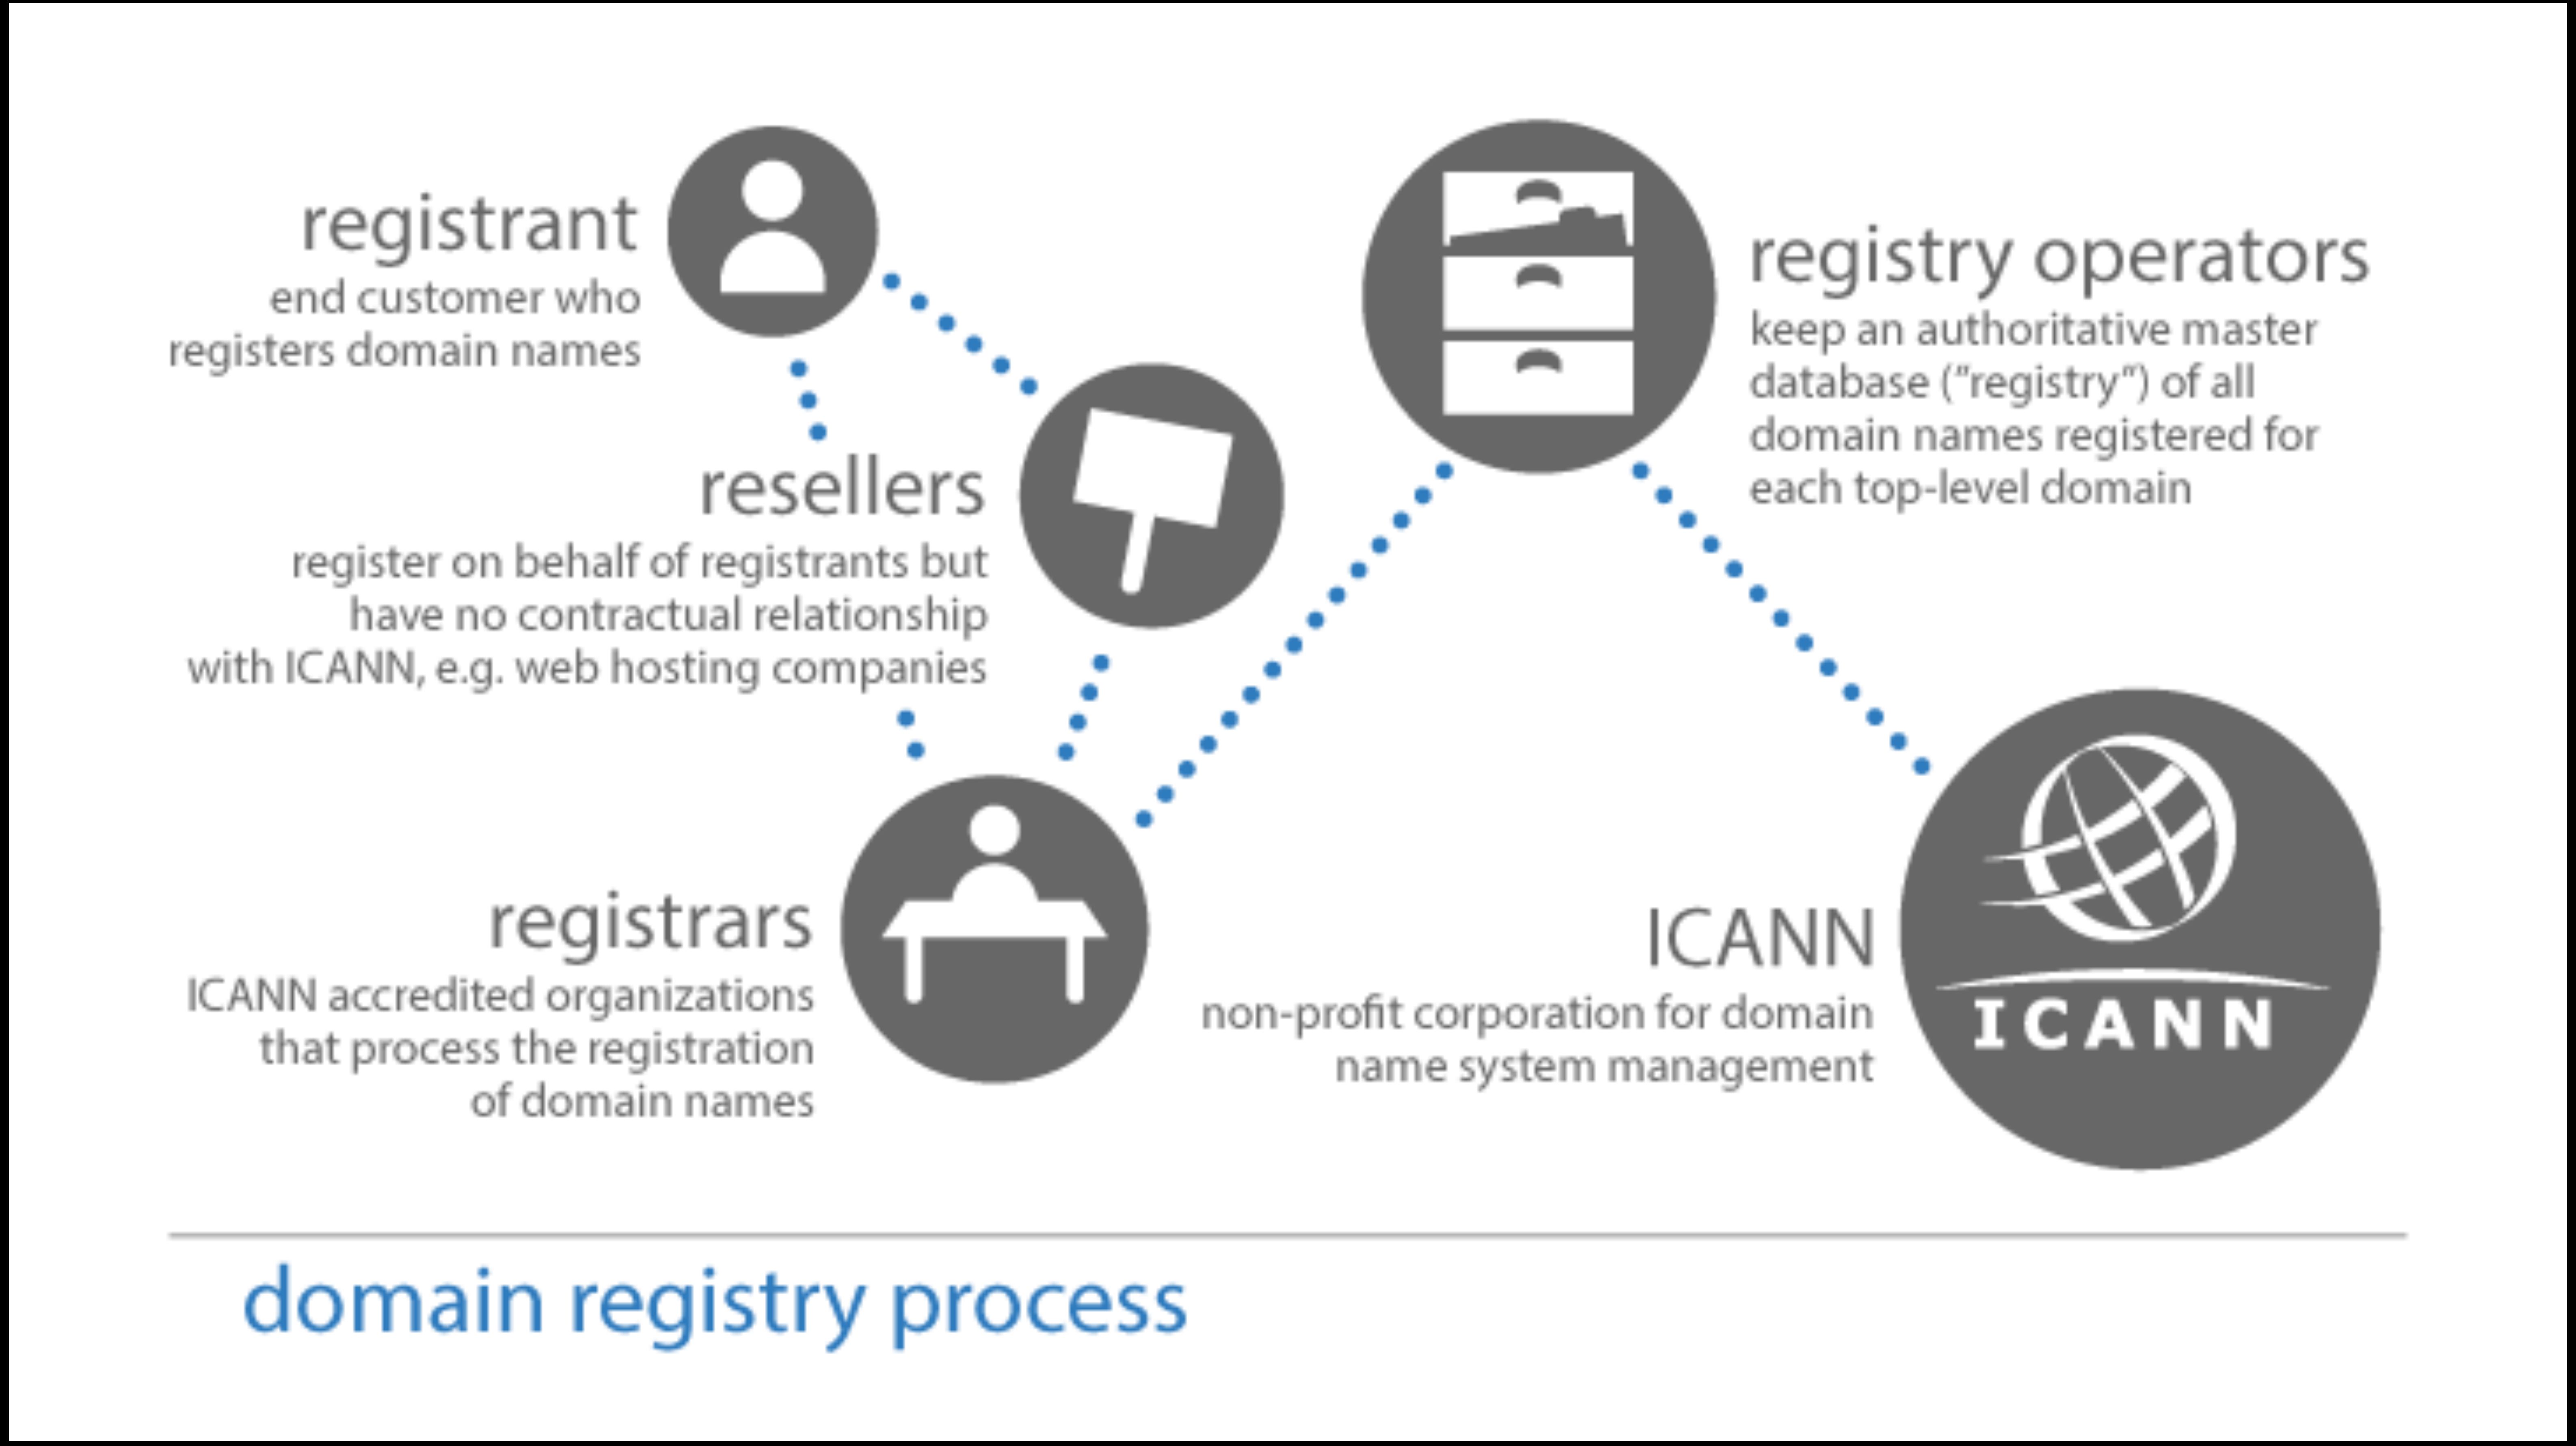

registrant

end customer who registers domain names

resellers

register on behalf of registrants but have no contractual relationship with ICANN, e.g. web hosting companies

registrars

ICANN accredited organizations that process the registration of domain names

registry operators

keep an authoritative master database ("registry") of all domain names registered for each top-level domain

ICANN

non-profit corporation for domain name system management

domain registry process

What info do 3-DNS want to maintain?

- NS records
- DS records
- A/AAAA records
    - need to be able to look up if glue is registered, add and delete.

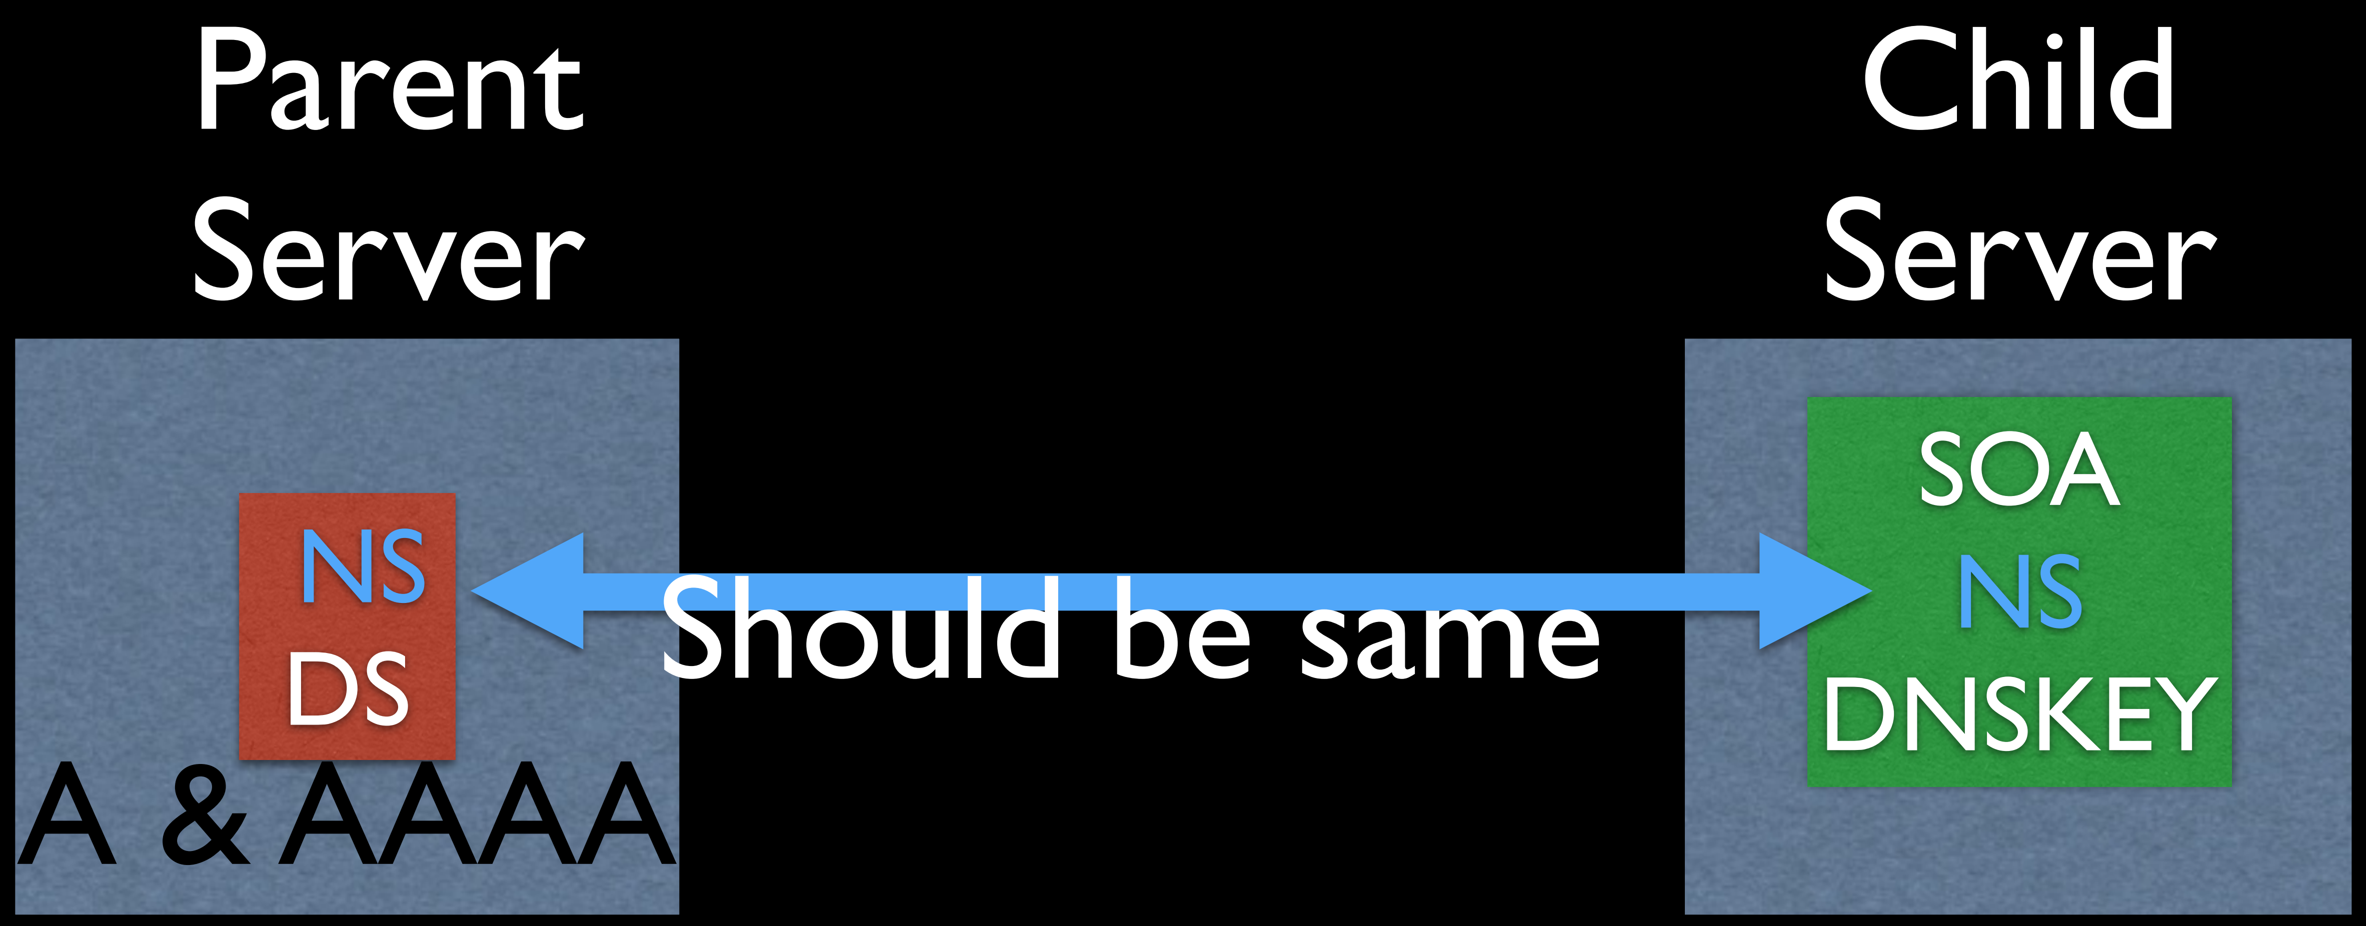

What happens today?

- To change information in parent Registrant has to be in the loop
  - Not reliable, registrant may or may not take action
  - Not timely
  - Cut & Paste errors happen.
- Registrant can give access to registration account to 3-DNS ==> BAD idea

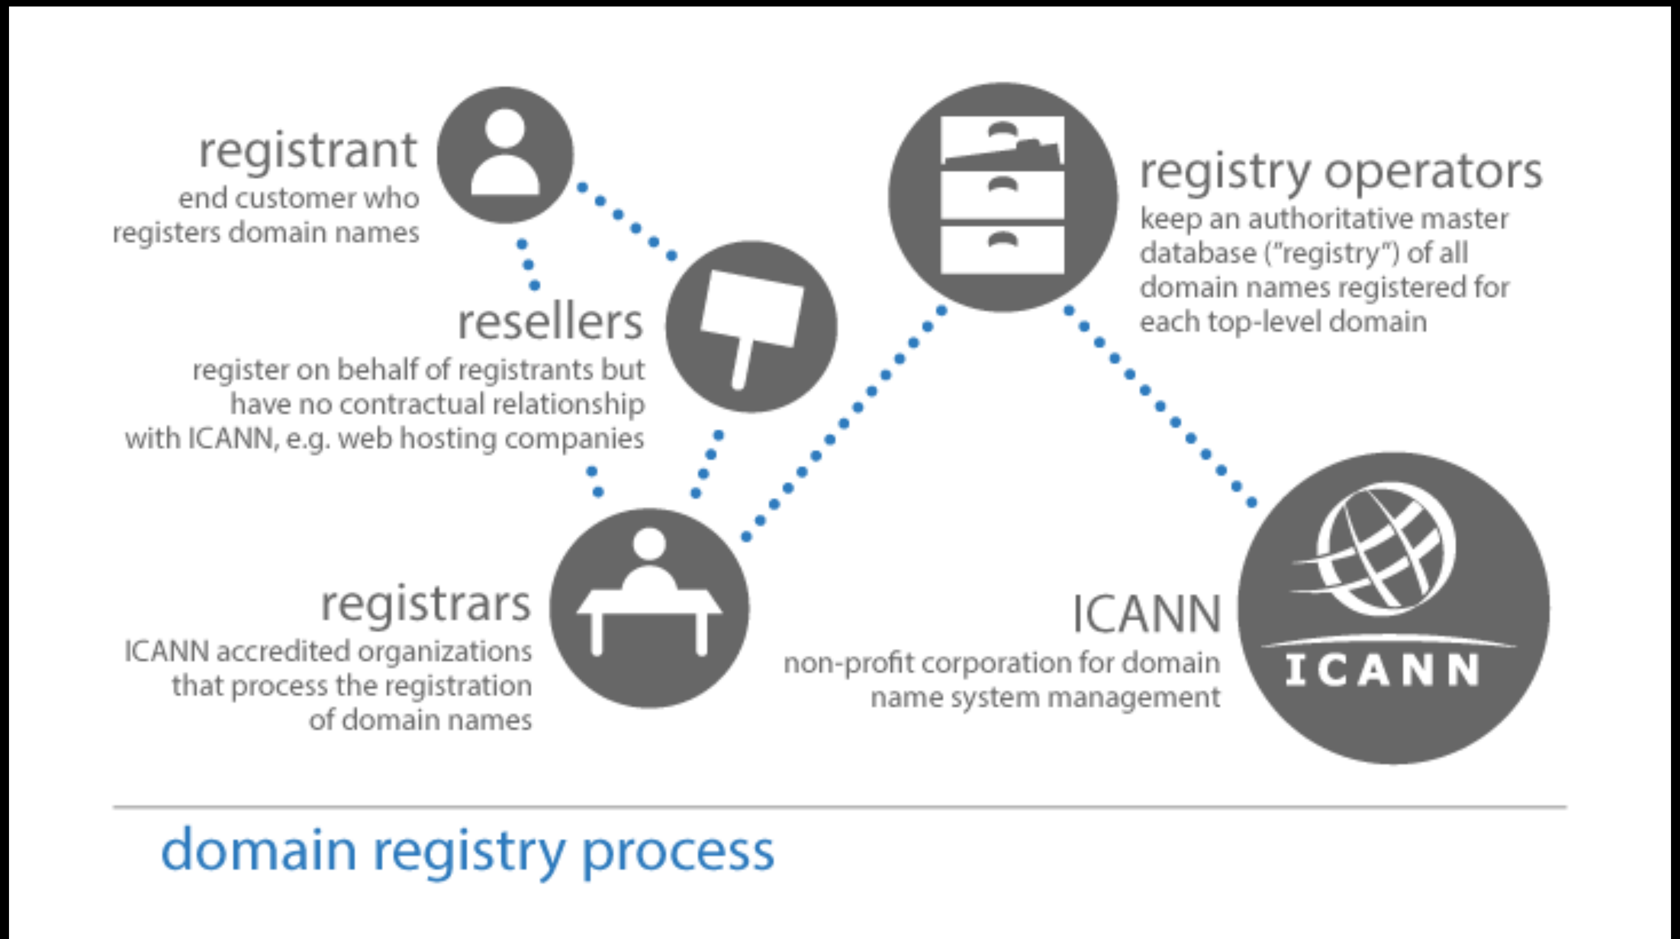
registrant
end customer who registers domain names
resellers
register on behalf of registrants but have no contractual relationship with ICANN, e.g. web hosting companies
registrars
ICANN accredited organizations that process the registration of domain names
registry operators
keep an authoritative master database ("registry") of all domain names registered for each top-level domain
ICANN
non-profit corporation for domain name system management
ICANN
domain registry process

3-DNS as registars?

- Addresses part of the problem
  - Hard to become registrar in all ccTLD’s
  - Registrars/resellers are frequently partners with 3-DNS

What is desired by 3-DNS?

- Ability to gain authenticated permission to maintain delegation information for customers
- Ability to learn where to change information and connect there
    - WHOIS has last century contact information when it has any, frequently unusable

How can this be done?

- When DNSSEC is enabled
  - Child zone can advertise what the contents of NS and DS should be
    - via NS and CDS/CDNSKEY records when DNSSEC is present [RFC7344]
    - Not specified how to tickle right parental agent.
    - Not possible to say do it NOW!!

Vision

- If 3-DNS gets authenticated and authorized to make changes to
  - NS/DS/glue for specific domain, these changes can be injected into registration systems via
    - Registars/Resellers
    - Registries
- ==> updates can take place at Internet speed

What is preventing ?

- ICANN and most ccTLD polices do not include 3-DNS as participants in the domain name industry, thus no relationships
- No way to grant limited access [in most systems].
- Current players worry about implications
- Protocols need to be specified
- Systems need to be updated

Current steps

- Make Domain Name industry and customers aware of the issues
- Motivate people to start thinking about changes.
- Motivate development and deployment
- Start experimenting ccTLD’s are the place to experiment

DNS TTL’s

The meaning of DNS TTL

- TTL == Time To Live ==> This tells resolver “you can cache these records no longer than this”.
- Caching resolvers will honor this within reason i.e. apply upper and lower limits as well as tossing of records when cache is full
- Non-Caching Resolver (mostly Forwarders) will only keep during one operation

The effects of TTL’s

- Long TTL ==> perceived stability
  ==> changes take long time
- Short TTL ==> frequently more query traffic
  ==> requires auth servers to be accessible at all times
- During change answers are inconsistent across the Internet

TTL Cost and Benefits

- Do not always line up nicely
- Goal: Fast Global visibility
  - when all resolvers in the have the “current” version of a RRset
  - Two sequential propagation delays
    - Authoritative servers: From Primary to last secondary
    - Resolvers: For old data to expire in all resolvers

Myth: TTL Only affects zone Operator

- For short TTL
  - Resolvers do more lookups
  - Users may see increased latency but more current content
    - Low TTL must be backed up by good DNS service
- For Long TTL
  - Resolver may return out-of-date answers
  - Diagnosing and Fixing problems is harder and takes longer

Delegation Records

- Affected RRsets: NS (both sides), DS, DNSKEY
  - NS is supposed to match but not always true
  - DS to DNSKEY mismatch==>
    - bogus domain in validating resolvers
    - works fine in non-validating
- For changes in delegation Parent must be updated
- Summary: Child is hostage of parent TTL and delays when updating delegation info

State of TLD TTL’s

- Table is summative for sample of TLD’s
- For big TLD’s the situation is worse
    - <=7200 cz/ch/nl/us
    - 86400 org/info/de/cn/dk/es/eu/be/..
    - 172800 com/net/io/uk
    - 345600 ru

TTL Range	NS Fraction	DS Fraction
0..2H	30%	28%
2H..24H	57%	62%
>24H	13%	10%

Goal: DNS operators change < 4 hours

- Assume Changes in parent take less than 1 hour
- Operations:
  - provision new operator
  - change NS in parent and old operator (if possible)
  - wait for resolvers
- Precondition: Child and Parent NS
  - TTL <= 2 hours

Goal: DNSSEC KSK rollover in 6 hours

- Assume changes in TLD’s take less than 1 hour
- Operations:
  - update DNSKEY and/or DS;
  - switch KSK signing key;
  - purge old DS and DNSKEY records (Not in critical path)
- Child DNSKEY set < 1 hour TTL
- Child and Parent NS + DS sets TTL <= 2 hours

Call for Action

- Start discussion on what the right goals and policies are
- Proposed goals:
  - Get TLD’s to adopt lower TTL <= 2H
  - Give 3-DNS access to maintain Delegation information
- Bonus: get registries and registrars to support new DNSSEC algorithms by default in particular ALG-13 ECDSA

Comments