



CLOUDFLARE®

DNSSEC

~~DNS~~

# A Review of CloudFlare

# Development and Deployment

Martin J. Levy @ CloudFlare

(based on work by Filippo Valsorda, Jono Bergquist & Ólafur Guðmundsson)

# Introduction

# CloudFlare DNS (the background)

- How big?
  - 2+ million domains
  - Authoritative for 40% of Alexa top 1 million
  - 43+ billion DNS queries/day
    - Second to only Verisign
- 63+ Anycast datacenters globally



# CloudFlare DNS offerings

- DNS for customers
  - UI based access; heavily linked to CDN/DDoS services
- DNS for partners
  - API based access; heavily linked to resold CDN/DDoS services
- DNS as a secondary service (vDNS offering)
  - Operates as an authoritative NS for TLDs (or significant domains)
    - Looks like a classic secondary service

# CloudFlare Goals & Solution

- DNSSEC at web scale
  - Scalable // DNSSEC for entire CloudFlare customer base
  - Simple // make it easy to consume
  - DNSSEC shouldn't be for power users only! It should be for everyone!
- DNS & DNSSEC software structure for this large scale deployment
  - CloudFlare wrote our own DNSSEC systems (scale & speed dictated this)
  - CloudFlare uses modern crypto and sign-on-the-fly at the edge

# CloudFlare Goals & Solution

- Changing the rules in order to deploy DNSSEC at large scale
  - Modifying and extending existing protocols to automate registrar interactions
    - Necessary to enable ease of use and deployment
    - Documented in RFCs or drafts (and code provided on github)
  - CloudFlare operates as a third-party DNS operator
    - i.e. Do not exit is many registration models
    - We are not the registrar or registry for most of these zones

Scale

# Why CloudFlare needs live signing

- Lots (lots!) of small, light traffic zones
- Heavily distributed network (45+ datacenters)
- **Dynamically generated records**
- Zone walking protection



# Issues with live signing

- Speed!
- Negative answers
- Key management

## Constraints

Keep size small, and don't require full zonefiles

Our solutions!

Speed

# CloudFlare's DNS(SEC) overview

- RRDNS is our in-house DNS server written in Go
- Resilient against attacks and abuse
- No zonefiles, records are pulled from a global distributed database
- Full featured (dynamic answers, CNAME flattening, ...)
- DNSSEC is just a “filter” applied to the answer

# Solving speed (and size): ECDSA P-256

- ECDSA (Elliptic Curve Digital Signature Algorithm) P-256 signatures
  - > 3x faster than RSA1024
- Measured on OpenSSL 1.0.2 on our servers
- We (Vlad Krasnov) ported OpenSSL ASM to Go
- 21x speedup for the sign: <https://go-review.googlesource.com/#/c/8968/>
- Bonus: small signatures, small keys, modern crypto!
- Supported by most validators, working on registrars

<https://tools.ietf.org/html/rfc6605>



# Solving speed (and size): ECDSA P256

RSA:  
1181 BYTES

```
ietf.org. 1800 IN DNSKEY 256 3 5 AwEAAAdDE
Uj67cfrZUojZ2cGRizVhgk0qZ9scaTVX NuXLM5Tw7VW0VIceeXAuuH2mPIiEV6M
WH94Vlubh HfiytNPZLr0bhUCHT6k0tNE6phLoHnXWU+6vpsYpz6GhMw/R9BFxW
ietf.org. 1800 IN DNSKEY 257 3 5 AwEAAavj
oPlwbq7Ws5WywbutbXyG24lMwy4jijlJ UsaFrS5EvUu4ydmuRc/TGnEXnN1XQk0
vz4U2vRCV ETLgDoQ7rhsiD127J8gVExj08B0113jCajbFRcMtUtFTjH4z7jXP2ZzD
ietf.org. 1800 IN RRSIG DNSKEY 5 2 1800 201
i3nTYvsuTFKqEou4Smku5Up01giVp s0pdDRwvei5g2HC8VK/nKHDhcotNR2unawRvA
z7mS8M NLgysKQMEZqJHfZhARZeSNIuK/QpRJhBX9UQYrv6IJ/215WqdL6C6aeB fYe
yX4Pnm09TtrpduZQqz120v+8nMITf4HJnSj7EvPN AxmCXg==
```

```
LjfqMvium4lgKtK ZLe97DgJ5/NQrNEGGQmr6fKv
:jjBB+wjYZQ5GtZHBFKVXACSWTiCtddHcue0eSVPi5
!6xeUwww46RVy3hanV3vN07LM5H niqaYc1Bbhk=
sZ+8AByyqFHLdZc Ho0GF7CgB50KYMvG0gysuYQ1
'nX1NrmMRowIu3DIVtGbQJmzpukpDVZaYMMAm8M5
'qqmw58nIELJUFOmcb/BdRLg byTeurFlnxs=
ietf.org. dp001u/mE0ZmcergtT4RA5DdV8E
YHee0MTVeHqk6YeyyiFvCL1XMLt3jj4/G3pjo
uEnn8uLXnXT1RdthZbnY g5yZReSwb4jVYQKC
```

ECDSA:  
305 BYTES

```
filippo.io. 3600 IN DNSKEY 257 3 13 DGpDkudNu/XQT1Km
QkXFtKcFZPxHGV07qSTIcDXS33/WtT8UUG7LyxAg KznsRSFEhiQVR53E69/E57IFm8b6Zw==
filippo.io. 3600 IN DNSKEY 256 3 13 koPbw9wmYZ7ggcjn
Q6ayHyhHaDNMYELKTqT+qRGrZpWScrr/lBcrm10Z 1PuQHB3Azhii+sb0PYFkH1ruxLhe5g==
filippo.io. 3600 IN RRSIG DNSKEY 13 2 3600 20150523
162528 20150422162528 42 filippo.io. KGjopS+z5rsK++grfGMuA2a1/vQ9S5tBX0Jq
ZbeTOYB0hfHG7S16hqR1 xfoibSJA1BiX5r9Ujo5YVU/NE1H0TQ==
```

# Solving speed (and size): ECDSA P256

## Standard Go crypto:

BenchmarkSingleSignECDSA	832,295 ns/op
BenchmarkSingleSignRSA	6,003,261 ns/op

## Go with Vlad's changes:

BenchmarkSingleSignECDSA	<b>60,806</b> ns/op
BenchmarkSingleSignRSA	3,124,274 ns/op

<https://blog.cloudflare.com/go-crypto-bridging-the-performance-gap/>

# Negative Answers



# Solving negatives: “Black Lies”

- To answer a NXDOMAIN normally we need:
  - Database lookups for previous and next name
  - 2 or 3 signatures (NSEC/NSEC3) - slow and big!
  - Previous and next name disclosure

# Solving negatives: “Black Lies”

```
mipappstg.comcast.com. 3600 IN NSEC mmgr.comcast.com. CNAME RRSIG NSEC
mipappstg.comcast.com. 3600 IN RRSIG NSEC 5 3 3600 20150508165102 2015050
1134602 39162 comcast.com. 0jKZ/h3bkK/AXs0kkg2Cbd13+aabCnCnp0sW9QHSrX8xcD04+SdxYx+E
F6PtFUUYh0KA8u9dcir7nkqI2Et326oAPuV8gbY6cLB8sFTceK6Fz0V0 /cIXrZyggy/VPf82FuBcoZsQnAb
erV0sI6RRbwjatPW65Wlo1bqKBrr9 Z7Q=
comcast.com. 3600 IN NSEC 208.20.10.201.comcast.com. A NS SOA
MX TXT RRSIG NSEC DNSKEY
comcast.com. 3600 IN RRSIG NSEC 5 2 3600 20150508165102 2015050
1134602 39162 comcast.com. TdPdnLkg5Zf12/rgskPWG194L+WigPn4AUD59p0qaX/T1fDmXU0g7WXH
38RORuUGmBmu7HSqzCekxJf1S//4ohw07NP3gSTz5dtW6co0Hw1E5n0 XaW+5nQC7pSBBjxa99DrUtPtpk6
2WACXuug/6A61FcIov0ppknsU1/12 fsQ=

;; Query time: 344 msec
;; SERVER: 127.0.0.1#53(127.0.0.1)
;; WHEN: Thu May 07 14:05:56 BST 2015
;; MSG SIZE rcvd: 736
```

# Solving negatives: “Black Lies”

- RFC 4470 introduces “white lies” for online signing:
  - Generate a NSEC on the name’s immediate predecessor, covering up to the successor (RFC4471)
  - Same with the wildcard
  - Solves: zone walking, database lookups
  - Still, 2 signatures to say one thing :(



# Solving negatives: “Black Lies”


- Our solution: true lies. Just sign a NOERROR.
- Place a NSEC on the name, cover until the successor, set only the NSEC and RRSIG bits

```
missing.filippo.io.      3587      IN        NSEC      \003.missing.filippo.io. RRSIG NSEC
missing.filippo.io.      3587      IN        RRSIG     NSEC 13 3 3600 20150507190048 201505
05170048 35273 filippo.io. Fb/xInfArVCMJWBDBqsbBPxiKsC1ueUyBFGi5lAHbjRBGAGm8sKDJx/l
YA01bKYzJep3dRgQw5hS89JukD+m8w==
```

# Solving negatives: “Black Lies”

```
missing.filippo.io.      3587      IN          NSEC       \003.missing.filippo.io. RRSIG NSEC
missing.filippo.io.      3587      IN          RRSIG      NSEC 13 3 3600 20150507190048 201505
05170048 35273 filippo.io. Fb/xInfArVCMJWBDBqsbBPxiKsC1ueUyBFGi51AHbjRBGAGm8sKDJx/1
YA01bKYzJep3dRgQw5hS89JukD+m8w==
```

```
:: Query time: 0 msec
:: SERVER: 127.0.0.1#53(127.0.0.1)
:: WHEN: Wed May 06 19:01:01 BST 2015
:: MSG SIZE rcvd: 363
```





# Solving negatives: “Black Lies”

- 1 signature op, no db lookup or zone walking
- The entire answer fits 512 bytes (actually, < 400!)
- End-user behavior is unchanged

```
missing.filippo.io.      3587      IN        NSEC      \003.missing.filippo.io. RRSIG NSEC
missing.filippo.io.      3587      IN        RRSIG     NSEC 13 3 3600 20150507190048 201505
05170048 35273 filippo.io. Fb/xInfArVCMJWBDBqsbBPxiKsC1ueUyBFGi5lAHbjRBGAGm8sKDJx/l
YA01bKYzJep3dRgQw5hS89JukD+m8w==
```

# Solving negatives: the “NSEC shotgun”

- But. To answer a missing type on an existing name, we still need to query the database for the NSEC bitmap
- That’s not even always possible! (Dynamic answers)

```
filippo.io. 3600 IN NSEC \003.filippo.io. A NS SOA MX TXT AAAA RRSIG  
NSEC DNSKEY
```

# Solving negatives: the “NSEC shotgun”

- Step back: what is a NSEC? A denial of existence.
- “The types not in the bitmap don’t exist”
- So, let’s make a “minimally covering” one.  
By setting all possible bits in the bitmap!

```
filippo.io. 3600 IN NSEC \003.filippo.io. A NS SOA WKS HINFO MX TXT  
AAAA LOC SRV CERT SSHFP IPSECKEY RRSIG NSEC DNSKEY TLSA HIP OPENPGPKEY  
SPF
```



# Solving negatives: the “NSEC shotgun”

- Asked for TXT and there’s no TXT? Set all the other bits that might exist.
- The NSEC is a valid denial for TXT, and is useless for an attacker that wants to replay it for other queries.

filippo.io. 3600 IN NSEC \003.filippo.io. A NS SOA WKS HINFO MX ~~TXT~~  
AAAA LOC SRV CERT SSHFP IPSECKEY RRSIG NSEC DNSKEY TLSA HIP OPENPGPKEY  
SPF

# Key Management

# Solving keys: centralized DNSKEY sets

- It's live-signing, you need the ZSK at the edge (for now)
- Protect the KSK: keep it in a safe central auditable machine, distribute the signed DNSKEY sets to edges
- Short regular RRSIG validity, longer for DNSKEY
- Prepared to roll the ZSK fast at any time

# Solving keys: global ZSK and KSK

- No reason to have millions of ZSKs and KSKs:
  - all would be used/stored/rolled together
- Use a single KSK and a single ZSK with multiple names

```
filippo.io. 3600 IN DNSKEY 256 3 13 koPbw9wmYZ7ggcjnQ6ayHyhHaDNMYELKTqT  
+qRGrZpWSccr/lBcrm10Z 1PuQHB3Azhii+sb0PYFkH1ruxLhe5g==
```

```
cloudflare-dnssec-auth.com. 3600 IN DNSKEY 256 3 13 koPbw9wmYZ7ggcjnQ6a  
yHyhHaDNMYELKTqT+qRGrZpWSccr/lBcrm10Z 1PuQHB3Azhii+sb0PYFkH1ruxLhe5g==
```

“DS” – Simplify

# How long does it take to ?

- Post a new selfie on Facebook and all your friends to be notified
  - few seconds (this is **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  $\text{MAX}(\text{Parent NS TTL}, \text{Child NS TTL})$
- Transfer a domain from one registrar to another one?
  - 1 sec ... **5 days**
- DNSSEC key rollover
  - **many DAYS (your-mileage-may-vary)**

# 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

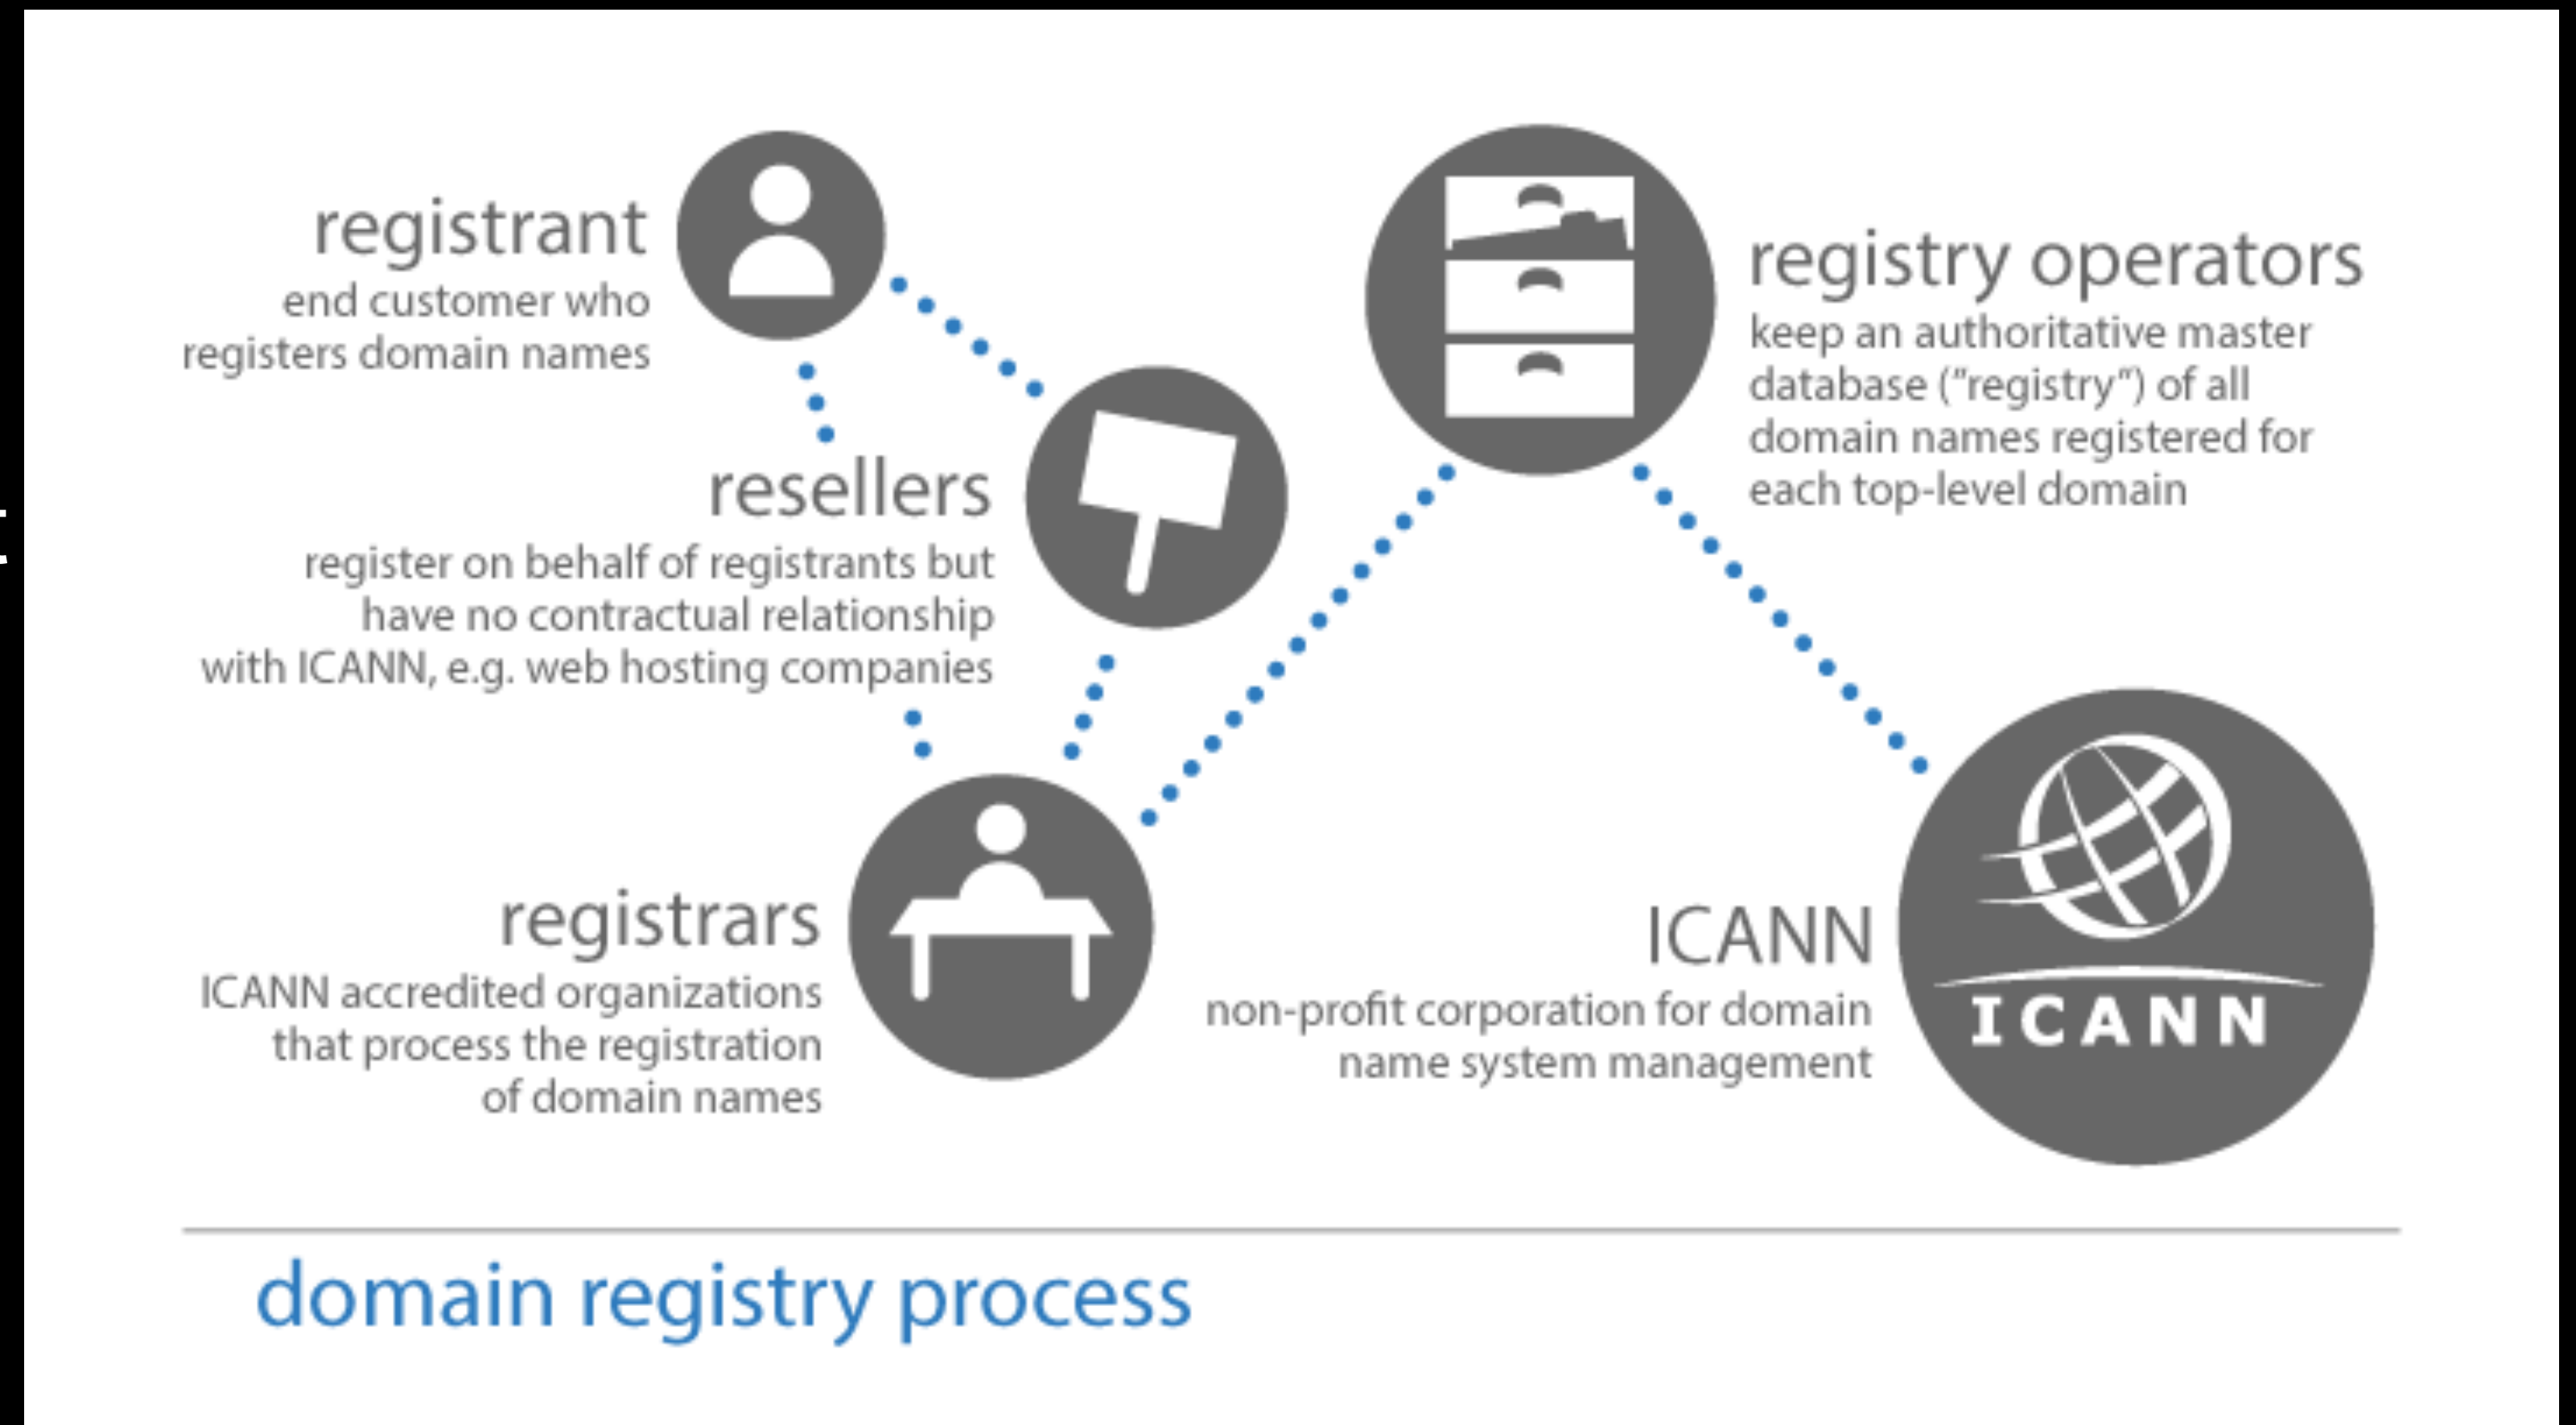
# Third party DNS operator (3-DNS)

- Definition: An entity contracted by “owner” of the domain to operate DNS on their behalf.
- Who: 3-DNS Operators include CDNs, 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



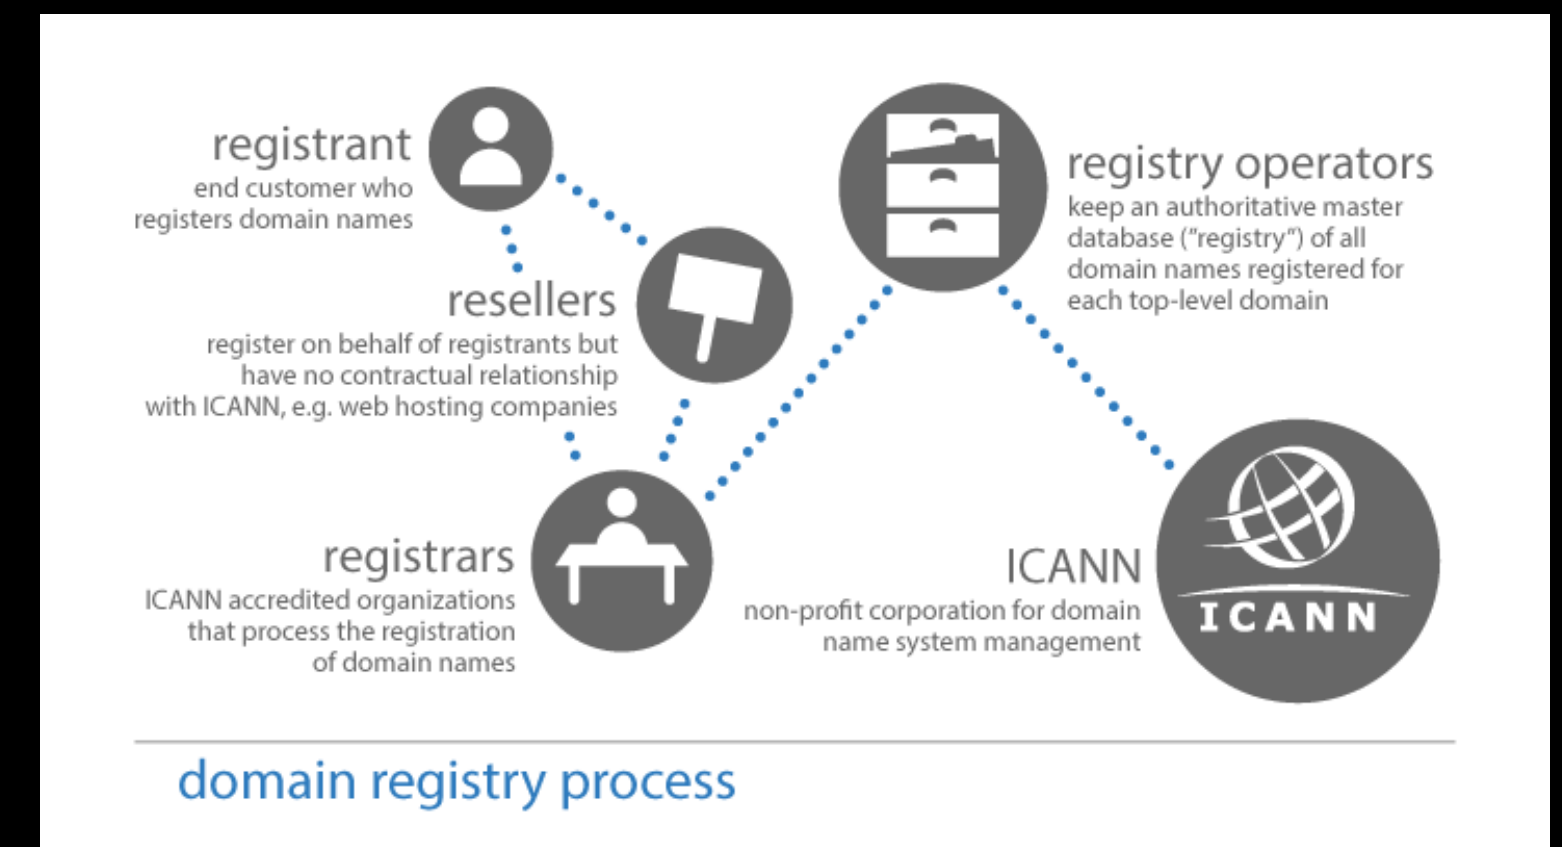
# What info does 3-DNS want to maintain?

- 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 !!



# 3-DNS as registrars?

- 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?

- #1 In-band signaling
  - 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 – #2 Registry System interface

- 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
  - Registrars/Resellers
  - Registries
- Hence: Updates can take place at Internet speed

# 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  $\leq$  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  $\leq 2H$
  - Give 3-DNS access to maintain Delegation information
- Bonus: get registries and registrars to support new DNSSEC algorithms by default in particular algorithm 13 ECDSA

ANY queries

# Deciding to Neuter “ANY” queries

- An ANY query is a bad idea
  - Amplification, Information leaks, Non reliable responses, Expensive
- Applications (and people) assume ANY returns ALL records of all types
  - Firefox had a version that used ANY to retrieve A & AAAA in one query



<https://tools.ietf.org/html/draft-ogud-dnsop-any-notimp-00>  
<https://blog.cloudflare.com/deprecating-dns-any-meta-query-type/>

# Responses to neutering “ANY” queries

- Positive!
  - “We have this problem”
  - “We spend too much on bandwidth because of ANY queries”
  - “Yes stop this information leak”
- Negative!
  - “You are hurting Firefox and Qmail”
  - “you are idiots !!!!”
  - “I use ANY to debug my systems all the time!!!”

# The qmail issue

- On DNSOP mailing list D. J. Bernstein wrote an explanation as to what Qmail is doing
  - Translation: Qmail uses ANY as a probabilistic optimization
  - Will fall back to normal resolution if ANY does not yield “useful” answer
- Hence: CloudFlare will not break qmail

[https://mailarchive.ietf.org/arch/msg/dnsop/kXSApuM4i0WLolo3\\_OhrCcAZ-cc](https://mailarchive.ietf.org/arch/msg/dnsop/kXSApuM4i0WLolo3_OhrCcAZ-cc)

# Why does CloudFlare care about “ANY”

- Expensive and complex to enumerate all RR Type for a name
  - We hate big answers
  - Sometimes not even available => incomplete answers
- Deploying DNSSEC with on-line signing on the edge at massive scale
  - Waste of effort to sign all the RR types the query origin does not care about

# CloudFlare implemented “ANY”

```
$ dig +nocmd +nostats ANY cloudflarestatus.com @fred.ns.cloudflare.com
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 56815
;; flags: qr rd; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0
;; WARNING: recursion requested but not available

;; QUESTION SECTION:
;cloudflarestatus.com.      IN  ANY

;; ANSWER SECTION:
cloudflarestatus.com. 3789 IN HINFO "Please stop asking for ANY" "See draft-jabley-dnsop-refuse-any"

$
```

<https://tools.ietf.org/html/draft-jabley-dnsop-refuse-any-01>



# CloudFlare implemented “ANY”

- No customers use HINFO in their zones → No need for new type
- We can generate this on the fly early in the processing
  - No need for multiple database lookups, discovery of all types, or multiple signatures
  - Simplified our code as we can remove ANY processing from various parts
- Cached as-is by resolvers → stops retries
- Accepted by resolvers → doesn't break ... applications

"We have time for just one long-winded, self-indulgent question that relates to nothing we've been talking about."



# Summary – Questions & Answers

IXP peering information  
at PeeringDB

AS13335

Martin J. Levy, Network Strategy  
@mahtin / @cloudflare  
<http://www.cloudflare.com/>