DNSSEC at Scale

Dani Grant | DNS @ CloudFlare

CloudFlare

- Authoritative DNS provider (includes DNSSEC for free)
- 4M+ domains
- 40+ billion queries per day
- 76 edge locations in 40 countries (growing)

DNSSEC at Scale

1. Elliptic Curves
2. Negative Answers
3. Registrar and Registry Support

Elliptic Curves

Speed and Size

Background

- CloudFlare mitigates large DDoS attacks (often 400M+ pps)
- DDoS is sometimes done through DNS amplification (small DNS query returns large DNS answer)
- Signed zones with large signature sizes are good for attackers doing amplification attacks

Elliptic Curves: Small Packet Size

- CloudFlare uses ECDSA to keep key and signature sizes small
- Almost all DNS answers CloudFlare returns are < 512 bytes, even with DNSSEC

Why does ECDSA have smaller key sizes?

Energy to break 228 bit RSA key vs. 228 bit ECDSA key

RSA:

Energy to boil a teaspoon of water

Energy to break 228 bit RSA key vs. 228 bit ECDSA key

RSA:
same as boiling a teaspoon of water

ECDSA:
boiling all the water on earth

CLOUDFLARE

Comparing DNSKey Answers

ietf.org. 985 IN DNSKEY 256 3 5 AwEAAdDECajHaTjfSoNTY58WcBah1BxPKVIHBz4IfLjfqMvium4lgKtK ZLe97DgJ5/NQrNEGGQmr6fKvUj67cfrZUojZ2cGRizVhgkOqZ9scaTVX
NuXLM5Tw7VWOVIceeXAuuH2mPIiEV6MhJYUsW6dvmNsJ4XwCgNgroAmX hoMEiWEjBB+wjYZQ5GtZHBFKVXACSWTiCtddHcueOeSVPi5WH94Vlubh HfiytNPZLrObhUCHT6k0tNE6phLoHnXWU+6vpsYpz6GhMw/R9BFxW5Pd
PFIWBgoWk2/XFVRSKG9Lr61b2z1R126xeUwvw46RVy3hanV3vNO7LM5H niqaYclBbhk=
ietf.org. 985 IN DNSKEY 257 3 5 AwEAAavjQ1H6pE8FV8LGP0wQBFVL0EM9BRfqxz9p/sZ+8AByqyFHLdZc HoOGF7CgB5OKYMvGOgysuYQloPlwbq7Ws5WywbutbXyG24IMWy4jijlJ
UsaFrS5EvUu4ydmuRc/TGnEXnN1XQkO+walT4cLtrmcWjoY8Oqud6lDa Jdj1cKr2nX1NrmMRowIu3DIVtGbQJmzpukpDVZaYMMAm8M5vz4U2vRCV ETLgDoQ7rhsiD127J8gVExjO8B0113jCajbFRcMtUtFTjH4z7jXP2ZzD
cXsgpe4LYFuenFQAcRBRlE6oaykHR7rlPqqmw58nIELJUFoMcb/BdRLg byTeurFlnxs=
ietf.org. 985 IN RRSIG DNSKEY 5 2 1800 20170213210526 20160214200831 45586 ietf.org. lv7deO/DZ+5Q6mZa9NsT4QQ7ibFU5s73yv7+gHoRyhis/3JmsMy8NIA9
7xoQcYhw1kYNqIgJYZ39XbKcmLyxVG9lzIMFcJOWcWA7QZQ8dW7IbQ4Z /jm8tuoXWWCmO9m1MgSwYfpuPz6IELh8czNylHuG+RZJn1t31wlOnet/ xUDrM5btKotJFeYKAEyVPiuC5N3+R3icd8U96lS1ybKCkXVzbcaDMBNe
r21/avPL7ympHeDiR4ubSTJ4xHr0pg5wCusZS0VRrKMPZrYrW/XW1gWl qRlyY/i4rxl9xyaBiP39eD7B7JvyyRTJObsnjpdd1blchM+DLLzl/7q1 y/vFXw==
ietf.org. 985 IN RRSIG DNSKEY 5 2 1800 20170213210642 20160214200831 40452 ietf.org. J3FK20+dp6Dy8QnDE4xlv9LJroKfrYQIa4i+ymYWulZqL0GQhElkkfLb
vyjMrNoVPhKjzNiBobFZDgjhFBDur9GONuWMkM4isBc4gBAKgNrirmh7 963HJ+ngsgHsfRTUHp27lSTgPw/SaxrUOz5JJJytNvr6eTilsKHgtpaP Xn44E210XQd5ak71//xY2/yCNJHjN3zH41Z0ipDG8UIITWzScFRZcEA+
9frDMBwiv7M9CBbOBeMNDAZXXa6JjkuASROmNIu8mU2XRa+Q8yDnYfF1 1r7JrdASF+zLlrxBX0HHjWtCjn+GvEoPDDTDN6J9oDHImt8WH6Tmt57h oluC+g==

RSA: 1181 bytes

cloudflare.com. 3574 IN DNSKEY 257 3 13 mdsswUyr3DPW132mOi8V9xESWE8jTo0dxCjjnopKl+GqJxpVXckHAeF+ KkxLbxILfDLUT0rAK9iUzy1L53eKGQ==
cloudflare.com. 3574 IN DNSKEY 256 3 13 koPbw9wmYZ7ggcjnQ6ayHyhHaDNMYELKTqT+qRGrZpWSccr/lBcrm10Z 1PuQHB3Azhii+sb0PYFkH1ruxLhe5g==
cloudflare.com. 3574 IN RRSIG DNSKEY 13 2 3600 20160310040015 20160110040015 2371 cloudflare.com. kgH/lAYN5endrnFAfJsNZPJHQvcYXqOLHDgrkhMXwvVJzyac/892fFwa
r5jo6u/57JnMJTCGF3P+YHmLiBKE1w==

ECDSA: 313 bytes

ECDSA is fast

...important when you are computing 56.9 billion signatures a day.

Speeding up ECDSA in Go

- Native implementation in assembler (by Vlad Krasnov)
- 21x speed improvements
- Now part of standard Go crypto library as of Go 1.6
- Takes CloudFlare 0.0001 seconds to sign a DNS record

	Before	After	Speedup
ECDSA Sign	1,015,006 ns/op	48,741 ns/op	**20.8x**
ECDSA Verify	3,086,282 ns/op	146,991 ns/op	**21.0x**

Negative Answers

Saving Compute

Two problems with negative answers

1. Requires authoritative server to return previous and next name
2. 2 NSEC + 2 NSEC RRSIG to say one thing

The trouble with previous and next name.

Background on CloudFlare DNS technology

- In house DNS server in Go called RRDNS
- No concept of zone file, instead SQL database of DNS records
- Business logic in DNS, we dynamically generate answers on the fly

The problem with previous + next name

1. No zone file, so requires sorted search of the database
2. Dynamic answers make previous and next name hard
3. NSEC exposes zone info (and NSEC3 can be dictionary attacked)

RFC4470 White Lies

- Randomly generate previous and next name for NSEC
- Helps prevent zone walking and extra database lookups

The trouble with 2 NSEC to say 1 thing.

RFC4470 White Lies

- Still, two separately signed NSEC records to say one thing

CloudFlare “Black Lies” for NXDOMAIN

- The next name is always \000.[themissingname]
- One NSEC per answer

```
cloudflare.com.         1799  IN    SOA   ns3.cloudflare.com. dns.cloudflare.com. 2020905521 10000 2400
604800 3600
bogus.cloudflare.com.   3599  IN    NSEC  \000.bogus.cloudflare.com. RRSIG NSEC
cloudflare.com.         1799  IN    RRSIG SOA 13 2 86400 20160309213638 20160307193638 35273 cloudflare.
com. mgx1FncjVdOpWhMOqm6+kcPBi/6zC8LF00ccG3DA1RNiI6hXmrqnFiUg dsngBT3VYo0+8AsZ1l0vJiopCdNoTw==
bogus.cloudflare.com.   3599  IN    RRSIG NSEC 13 3 3600 20160309213638 20160307193638 35273 cloudflare.
com. 8nbevvyI/RsSjunQzjlPkIHphiAOu5gti+aj2ucBx3Nhc7cnaHtJbJ5C dFrOF7eoZuPeiegf0KTtMyhAYp3tWQ==
```

Comparing Negative Answers

ietf.org. 1799 IN SOA ns0.amsl.com. glen.amsl.com. 1200000317 1800 1800 604800 1800
ietf.org. 1799 IN RRSIG SOA 5 2 1800 20170213210533 20160214200831 40452 ietf.org. P8XoJx+SK5nUZAV/IqiJrsoKtP1c+GXmp3FvEOUZPFn1VwW33242LVrJ
GMI5HHjMEX07EzOXZyLnQeEvlf2QLxRIQm1wAnE6W4SUp7TgKUZ7NJHP dgLr2gqKYim4CI7ikYj3vK7NgcaSE5jqIZUm7oFxxYO9/YPz4Mx7COw6 XBOMYS2v8VY3DICeJdZsHJnVKlgl8L7/yqrL8qhkSW1yDo3YtB9cZEjB
OVk8uRDxK7aHkEnMRz0LODOJ10AngJpg9LrkZ1CO444RhZGgTbwzN9Vq rDyH47Cn3h8ofEOJtYCJvuX5CCzaZDlnBsjq9wNAiNBgIQatPkNriR77 hCEHhQ==
ietf.org. 1799 IN NSEC ietf1._domainkey.ietf.org. A NS SOA MX TXT AAAA RRSIG NSEC DNSKEY SPF
ietf.org. 1799 IN RRSIG NSEC 5 2 1800 20170213210816 20160214200831 40452 ietf.org. B9z/JJs30tkn0DyxVz0zaRlm4HkeNY1TqYmr9rx8rH7kC32PWZ1Fooy6
16qmB33/cvD2wtOCKMnNQPdTG2qUs/RuVxqRPZaQojIVZsy/GYONmlap BptzgOJLP7/HOxgYFgMt5q/91JHfp6Mn0sd218/H86Aa98RCXwUOzZnW bdttjsmbAqONuPQURaGz8ZgGztFmQt5dNeNRaq5Uqdzw738vQjYwppfU
9GSLkT7RCh3kgbNcSaXeuWfFnxG1R2SdlRoDICos+RqdDM+23BHGYkYc /NEBLtjYGxPqYCMe/7lOtWQjtQOkqylAr1r7pSI2NOA9mexa7yTuXH+x o/rzRA==
www.apps.ietf.org. 1799 IN NSEC cloudflare-verify.ietf.org. A RRSIG NSEC
www.apps.ietf.org. 1799 IN RRSIG NSEC 5 4 1800 20170213210614 20160214200831 40452 ietf.org. U+hEHcTps2IC8VKS61rU3MDZq+U0KG4/oJjIHVYbrWufQ7NdMdnY6hCL
OmQtsvuZVRQjWHmowRhMj83JMUagxoZuWTg6GuLPin3c7PkRimfBx7jI wjqORwcuvpBh92A/s/2HXBma3PtDZl2UDLy4z7wdO62rbxGU/LX1jTqY FoJJLJfJ/C+ngVMIE/QVneXSJkAjHV96FSEnreF81V62x9azv3AHo4tl
qnoYvRDtK+cR072A5smtWMKDfclr2fl11TAGIyhR55yAiollPDEz5koj BfMstC/JXVURJMM+1vCPjxvwYzTZN8iICf1AupyyR8BNWxgic5yh1ljH 1AuAVQ==

NSEC: 1094 bytes

cloudflare.com. 1799 IN SOA ns3.cloudflare.com. dns.cloudflare.com. 2020742566 10000 2400 604800 3600
blog.cloudflare.com. 3599 IN NSEC \000.blog.cloudflare.com. RRSIG NSEC
cloudflare.com. 1799 IN RRSIG SOA 13 2 86400 20160220230013 20160218210013 35273 cloudflare.com. kgjtJDuuNC/yX8yWQpol4ZUUr8s8yAXZi26KWBI6S3HDtry2t6LnP1ou
QK10Ut7DXO/XhyZddRBVj3pIpWYdBQ==
blog.cloudflare.com. 3599 IN RRSIG NSEC 13 3 3600 20160220230013 20160218210013 35273 cloudflare.com. 8BKAAS8EXNJbm8DxEI1OOBba8KaiimIuB47mPlteiZf3sVLGN1edsrXE
+q+pHaSHEfYG5mHfCBJrbi6b3EoXOw==

Black Lies: 357 bytes

Problems with NODATA:

- Would have to search the database for existing types (CPU expensive)
- Not always possible because of dynamic answers

CloudFlare “Black Lies” for NODATA

- Set all the types, except for the type you asked for
- When you ask for TXT:

```
blog.cloudflare.com.  3599 IN  NSEC \000.blog.cloudflare.com. A WKS HINFO MX
~~TXT~~ AAAA LOC SRV CERT SSHFP IPSECKEY RRSIG NSEC TLSA HIP OPENPGPKEY SPF
```

- When you ask for MX:

```
blog.cloudflare.com.  3599 IN  NSEC \000.blog.cloudflare.com. A WKS HINFO  ~~MX~~
TXT AAAA LOC SRV CERT SSHFP IPSECKEY RRSIG NSEC TLSA HIP OPENPGPKEY SPF
```

Registrar and Registry Support

The Last Mile

Registrar and Registry Support

- Big difference between being an Internet Standard and being adopted in practice
- DNSSEC is required by ICANN registrar agreement
- DNSSEC Algorithm 13 (ECDSA) has been a standard for years
- Still, many registrars + registries do not have support

*"In order **to enable the DNSSEC, the domain name must be under [the registrar's] DNS management** which means the domain will need to be moved to our servers. The changes [i.e. adding the DS] have not been completed and this request has been closed."*

“I talked to support at the registrar and they said that I would need to ***enter the DS record with you since my DNS is hosted here****.”*

Registrar Support: “The DNSSEC option is not yet operational, we still don’t provide support for it. ”

Registrant: “So if I add my DS record and ***it says ‘DNSSEC Active’, DNSSEC won’t really be active****?”*

Registrar Support: “Exactly.”

Have added support since our launch:

- Norid (.no)
- SIDN (.nl)
- Eurid (.eu)
- eNic (.eu)
- NZRS (.nz)
- NIC.br (.br)
- DNSimple
- Hover
- Internet.bs
- OVH
- Metaname

UNINETT Norid

DNSSEC: Discount calculation on 20 November

On Friday 20 November we will count the number of signed domain names per registrar and calculate a discount which will be credited each registrar. More than half of the .no zone is now signed.

Since we launched DNSSEC in December 2014, a number of registrars have started using the technology. As of today, 55 per cent of the .no zone was signed.

We encourage more registrars to sign there domain names, preferably before 20 November so they can be a part of the next discount calculation. Registrars with a lot of domains should sign them in batches over a couple of days. We suggest that you do the signing process during normal working hours (Monday-Friday 08:00-16:00 CEST) so that we can be of assistance if problems should emerge.

Please note that we do not credit amounts below NOK 200. Therefore, each registrar must have at least 67 signed domain names in order to get the discount.

https://www.ietf.org/id/draft-latour-dnsoperator-to-rrr-protocol-00.txt

Network Working Group
Internet-Draft
Intended status: Informational
Expires: April 21, 2016

J. Latour
CIRA
O. Gudmundsson
Cloudflare, Inc.
P. Wouters
Red Hat
M. Pounsett
Rightside
October 19, 2015

Third Party DNS operator to Registrars/Registries Protocol
draft-latour-dnsoperator-to-rrr-protocol-00.txt

Abstract

There are several problems that arise in the standard Registrant/Registrar/Registry model when the operator of a zone is neither the Registrant nor the Registrar for the delegation. Historically the issues have been minor, and limited to difficulty guiding the Registrant through the initial changes to the NS records for the delegation. As this is usually a one time activity when the operator first takes charge of the zone it has not been treated as a serious issue.

CLOUDFLARE

Interested in getting involved?

dnssec-integration@cloudflare.com

Questions?

@thedanigrant

dani@cloudflare.com