CloudFlare®

DNSSEC and DNS Proxying

DNS is hard

- at scale
- when you are a huge target

CloudFlare DNS

- is big

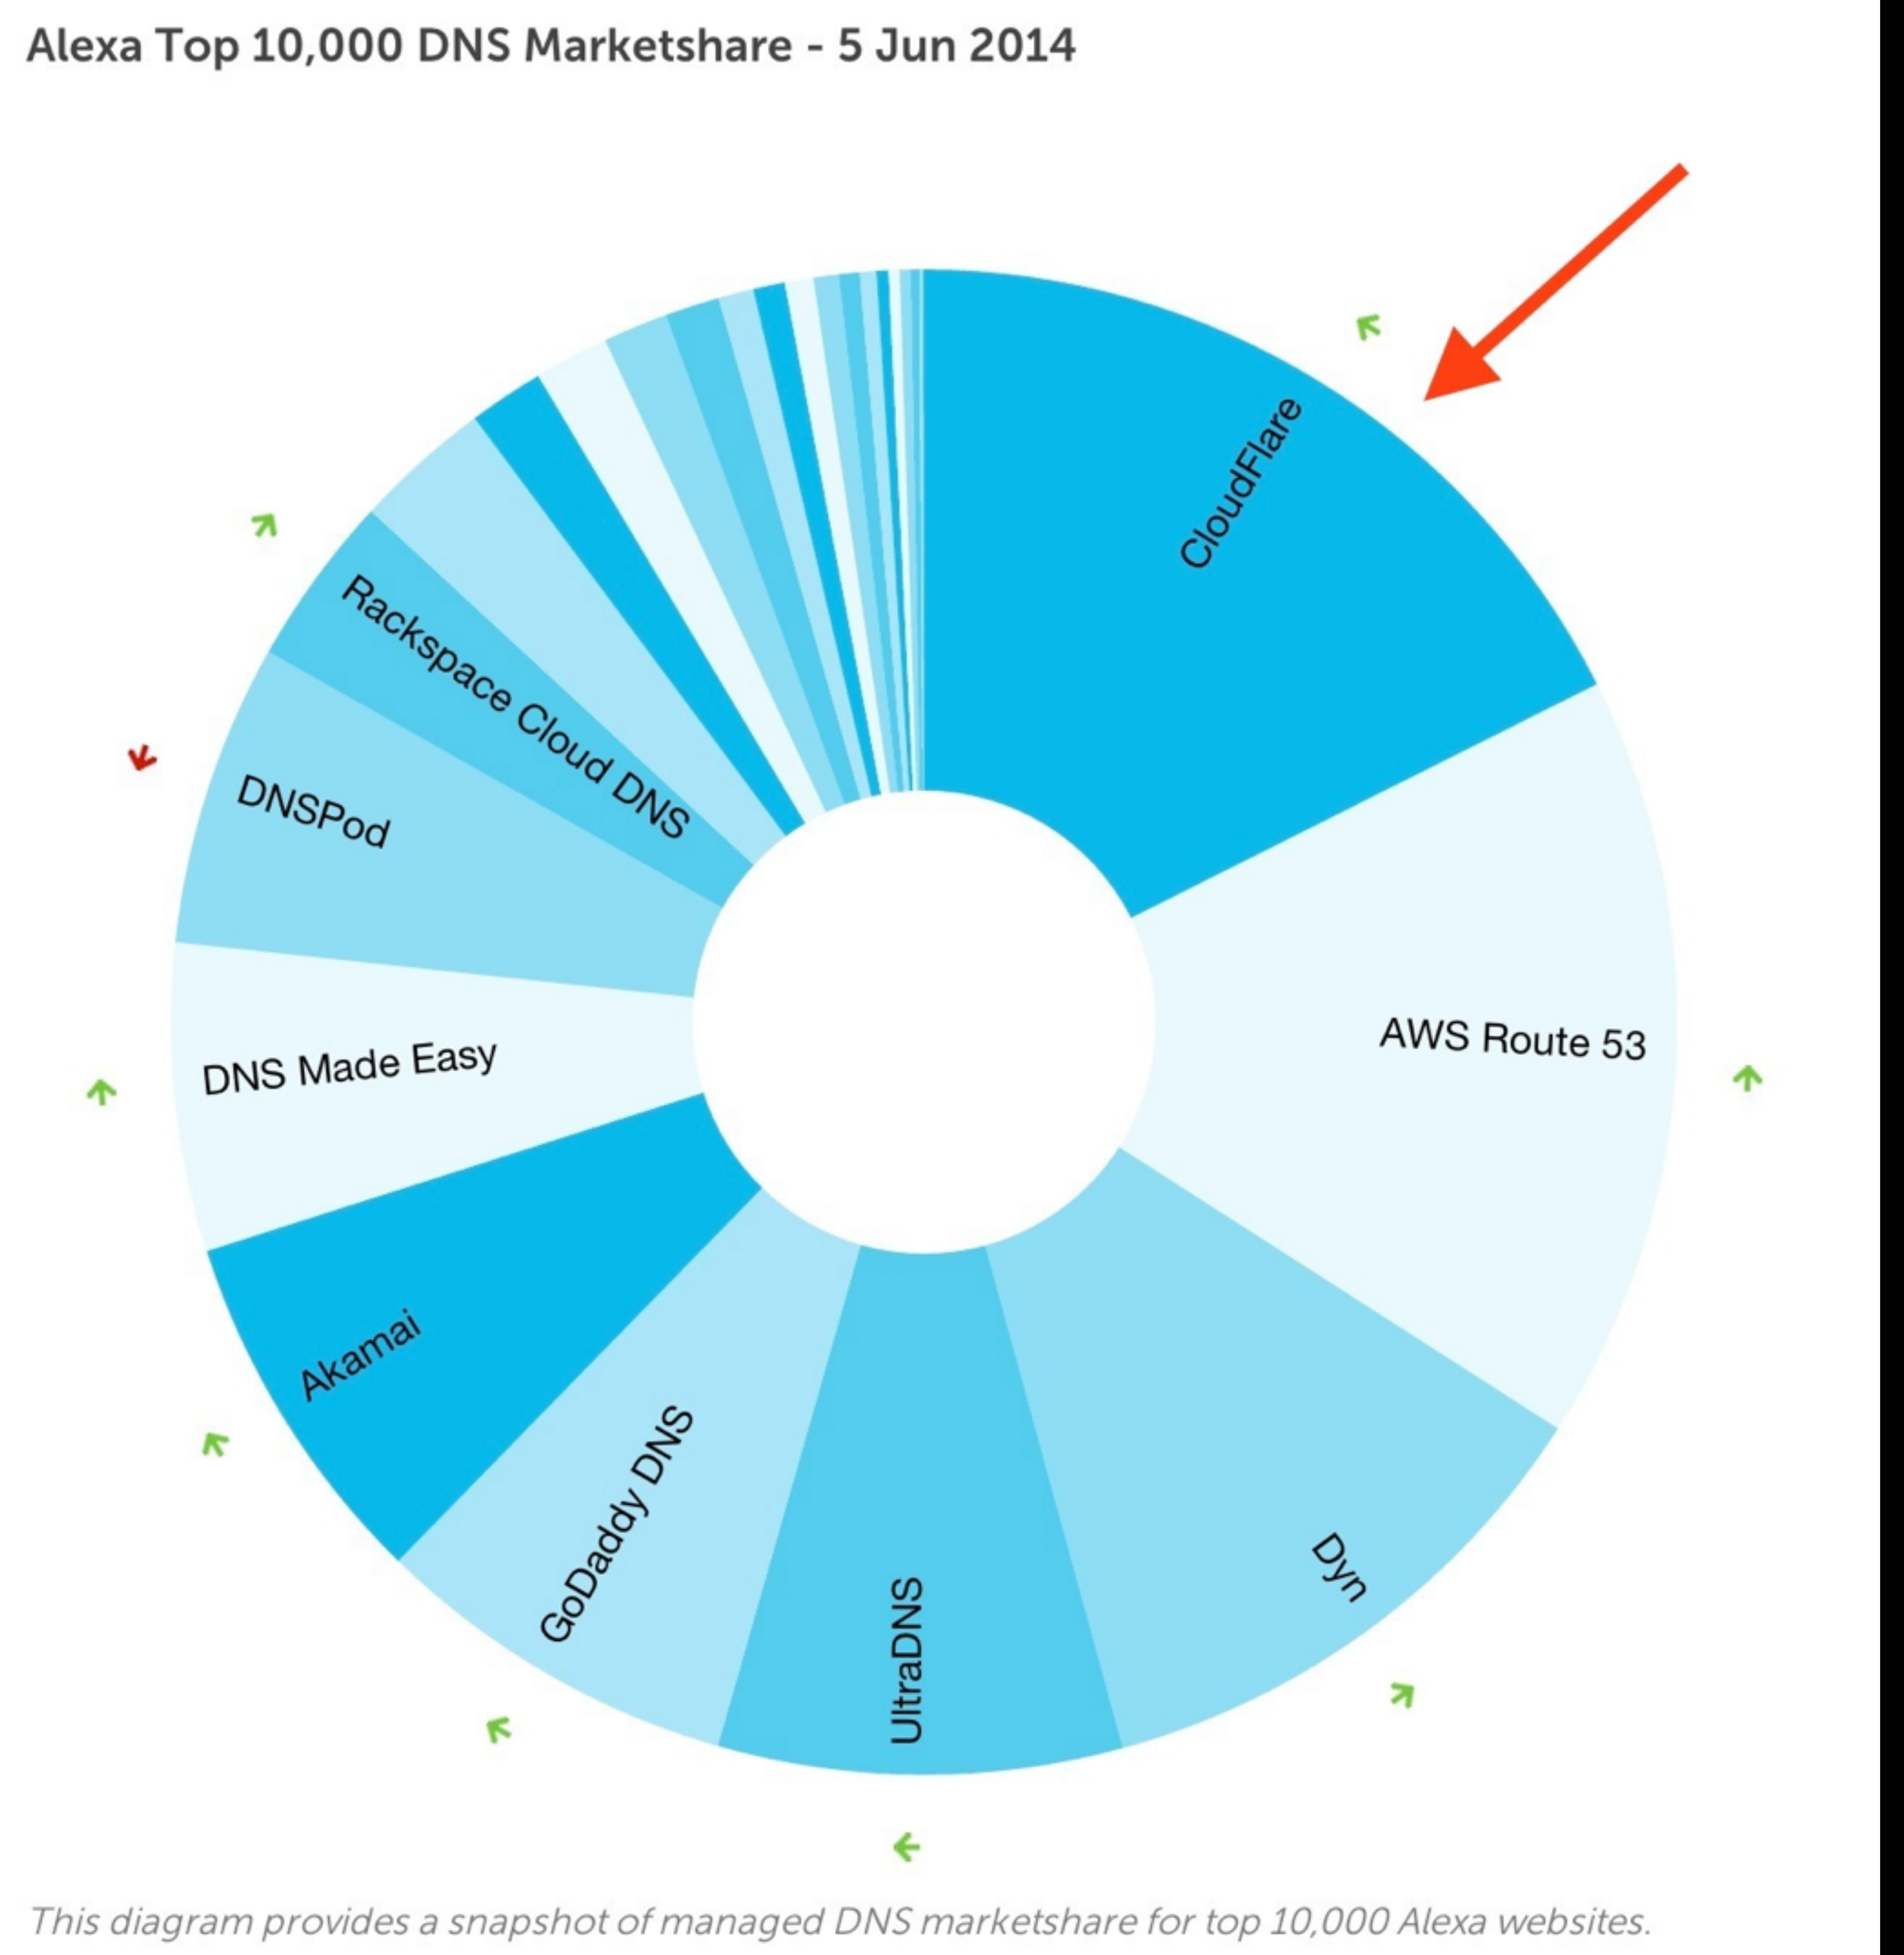

Alexa Top 10,000 DNS Marketshare - 5 Jun 2014

This diagram provides a snapshot of managed DNS marketshare for top 10,000 Alexa websites.

CloudFlare DNS

- is fast

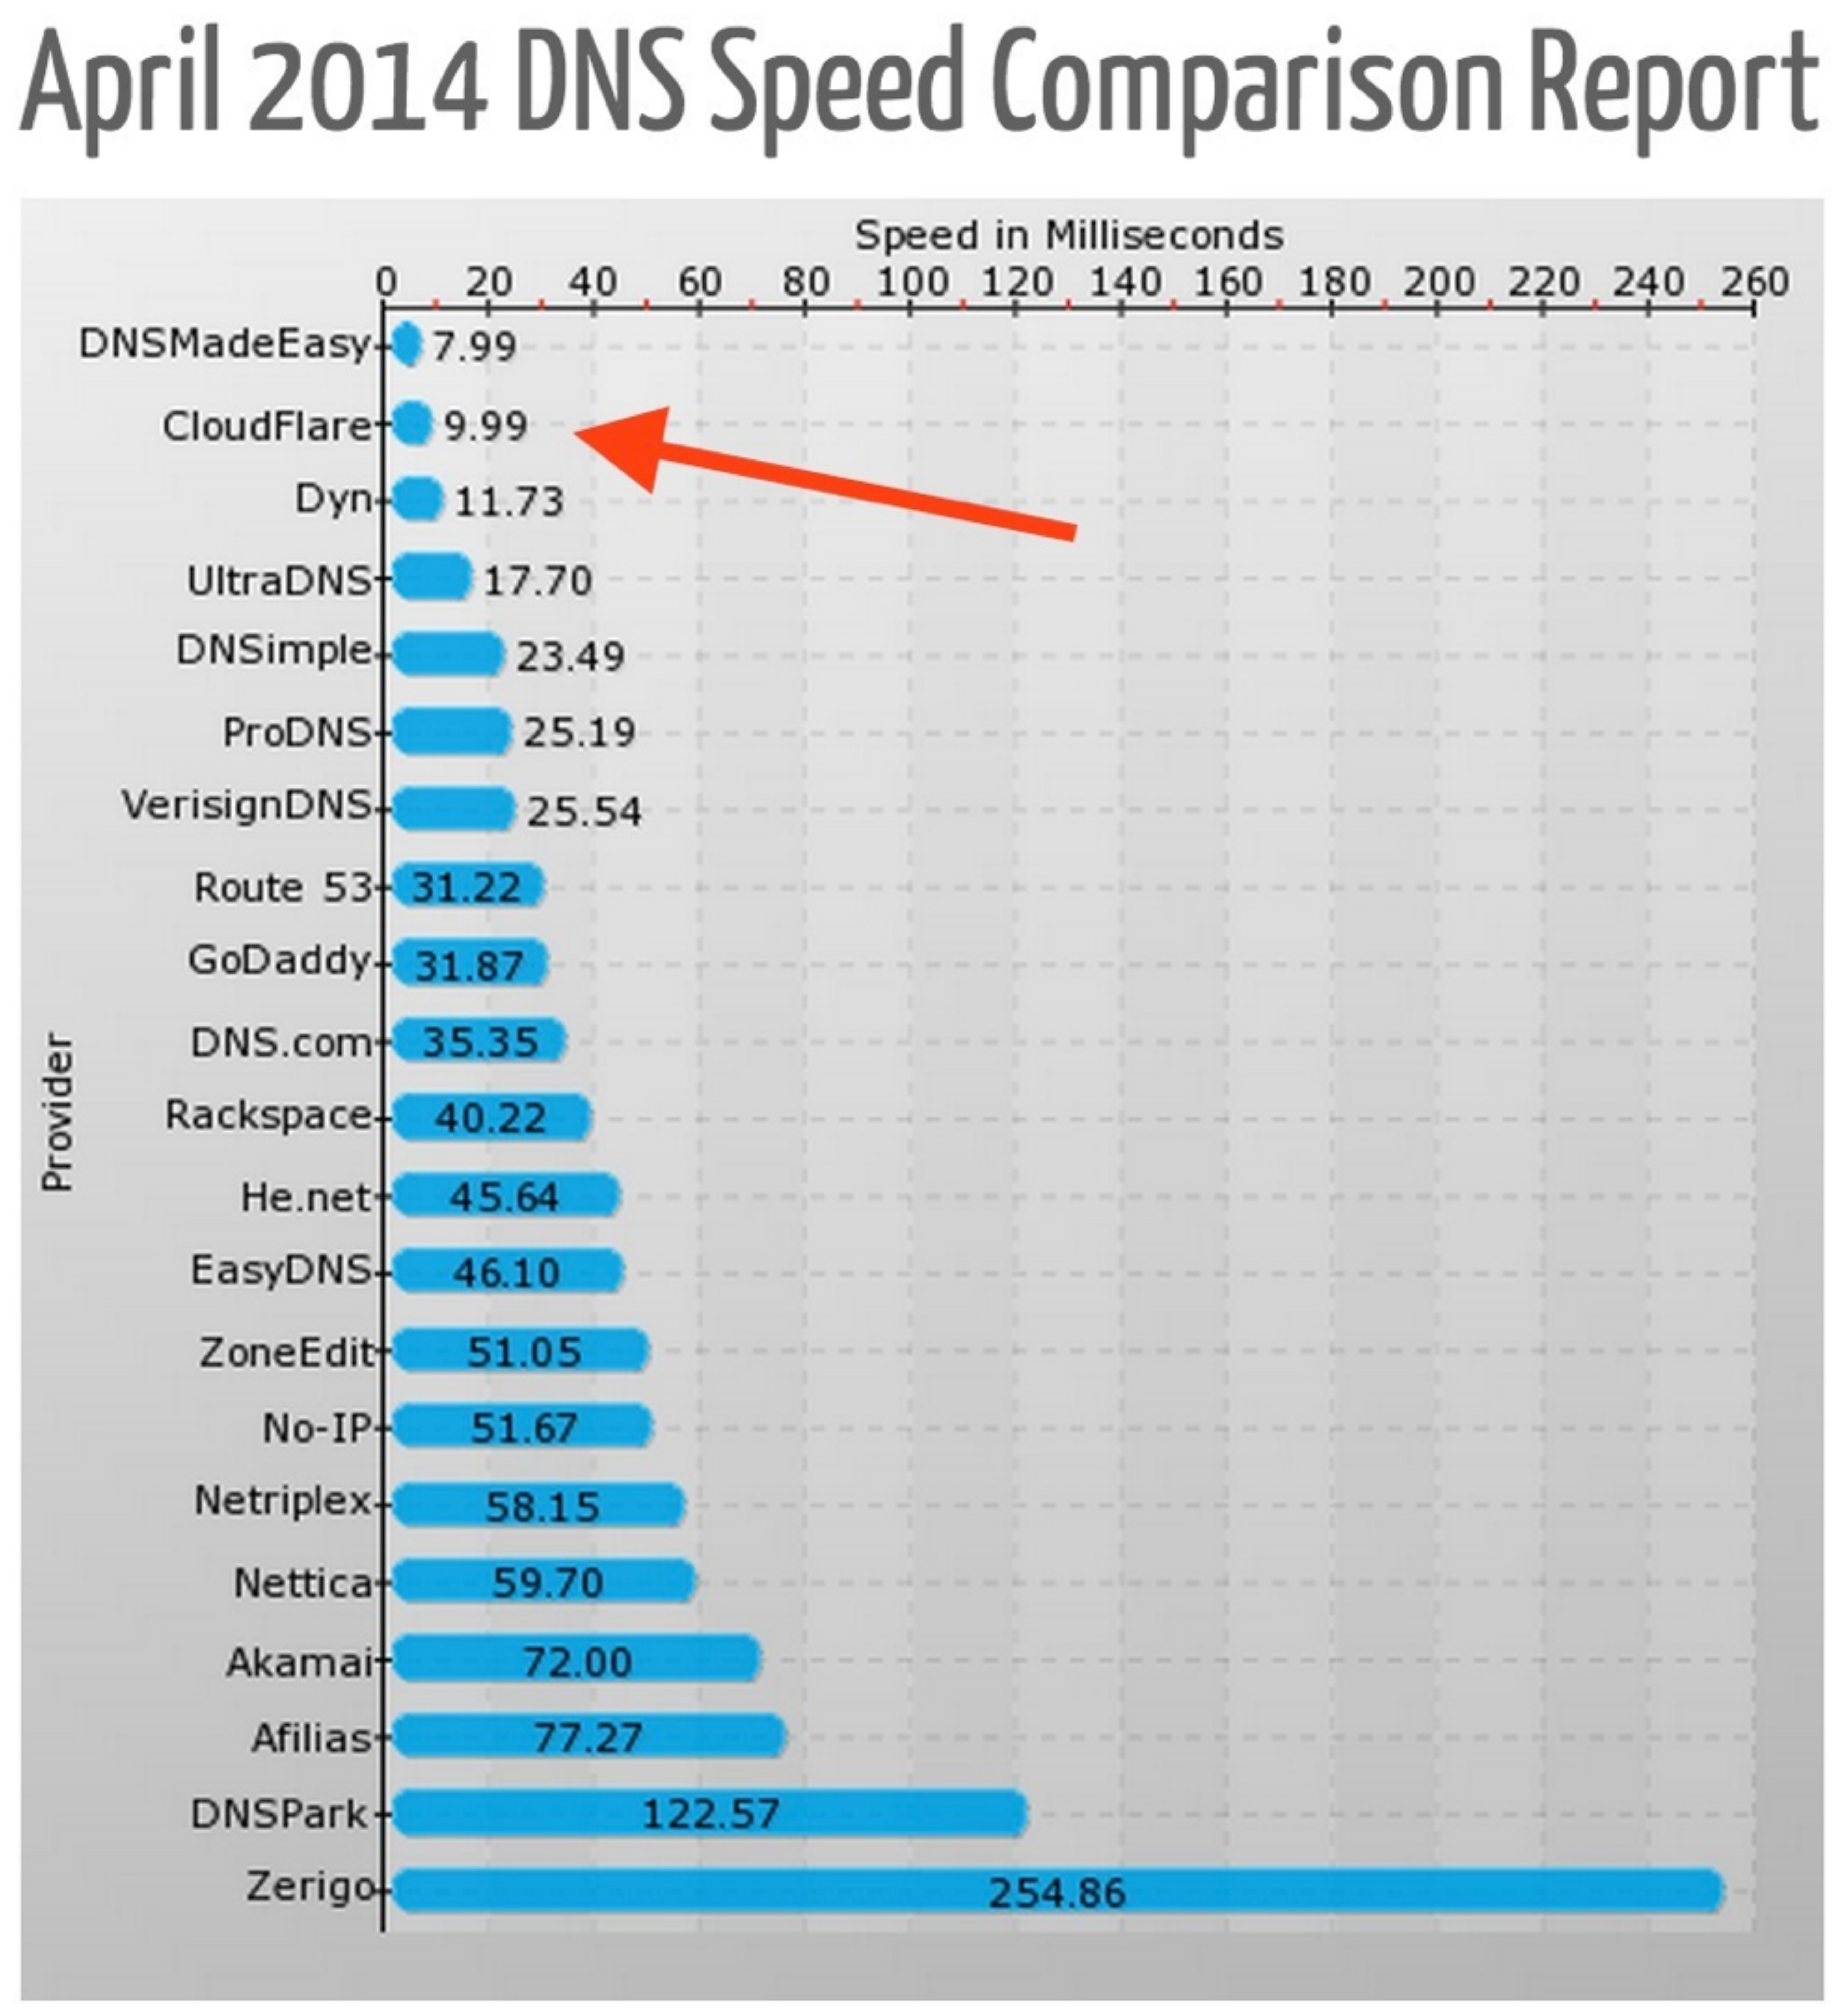

CloudFlare DNS

- is always under attack

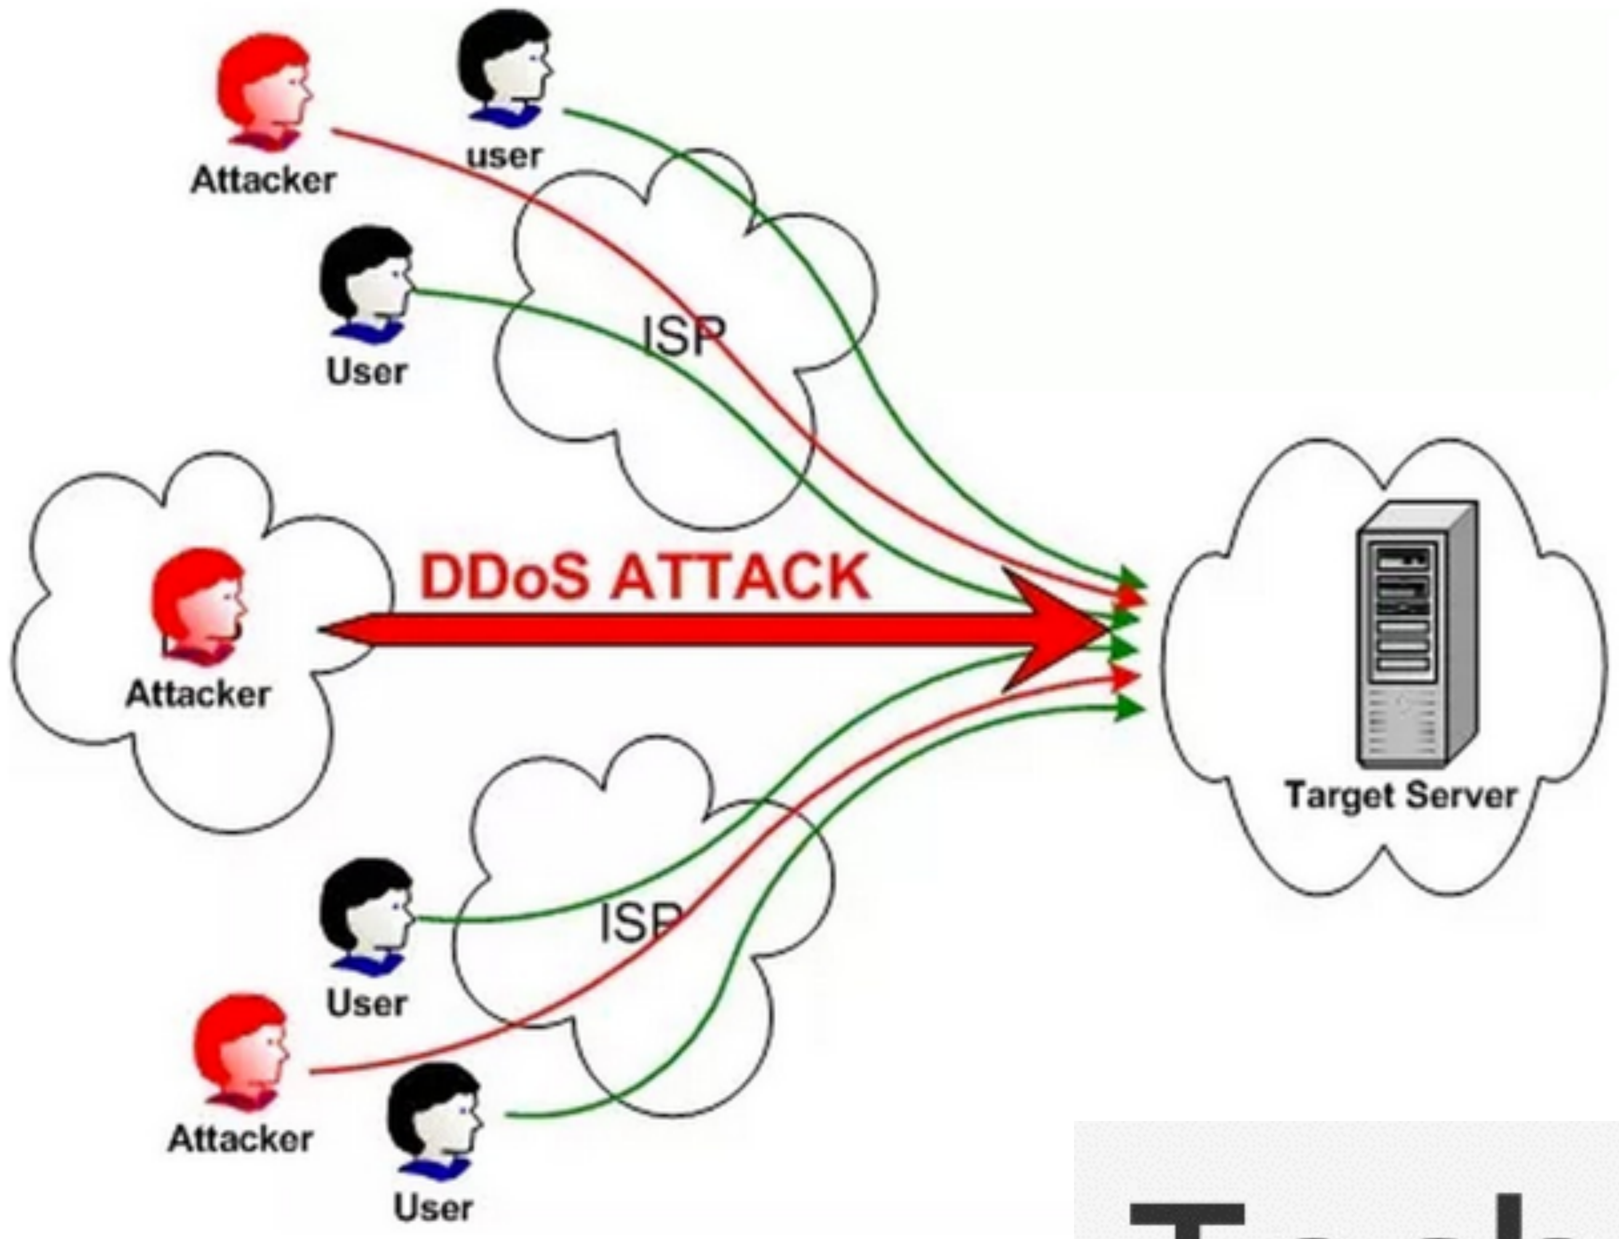

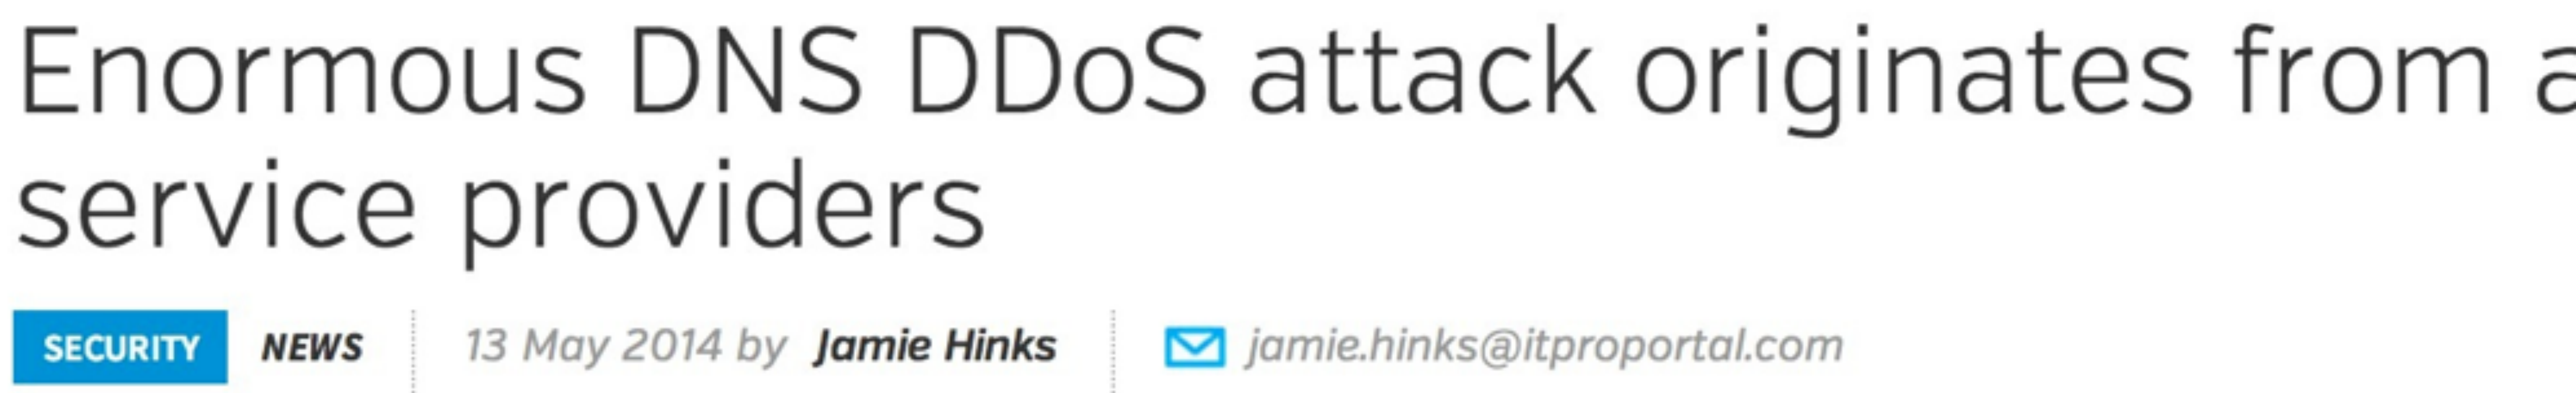
Enormous DNS DDoS attack originates from a
service providers

SECURITY NEWS 13 May 2014 by Jamie Hinks jamie.hinks@itproportal.com

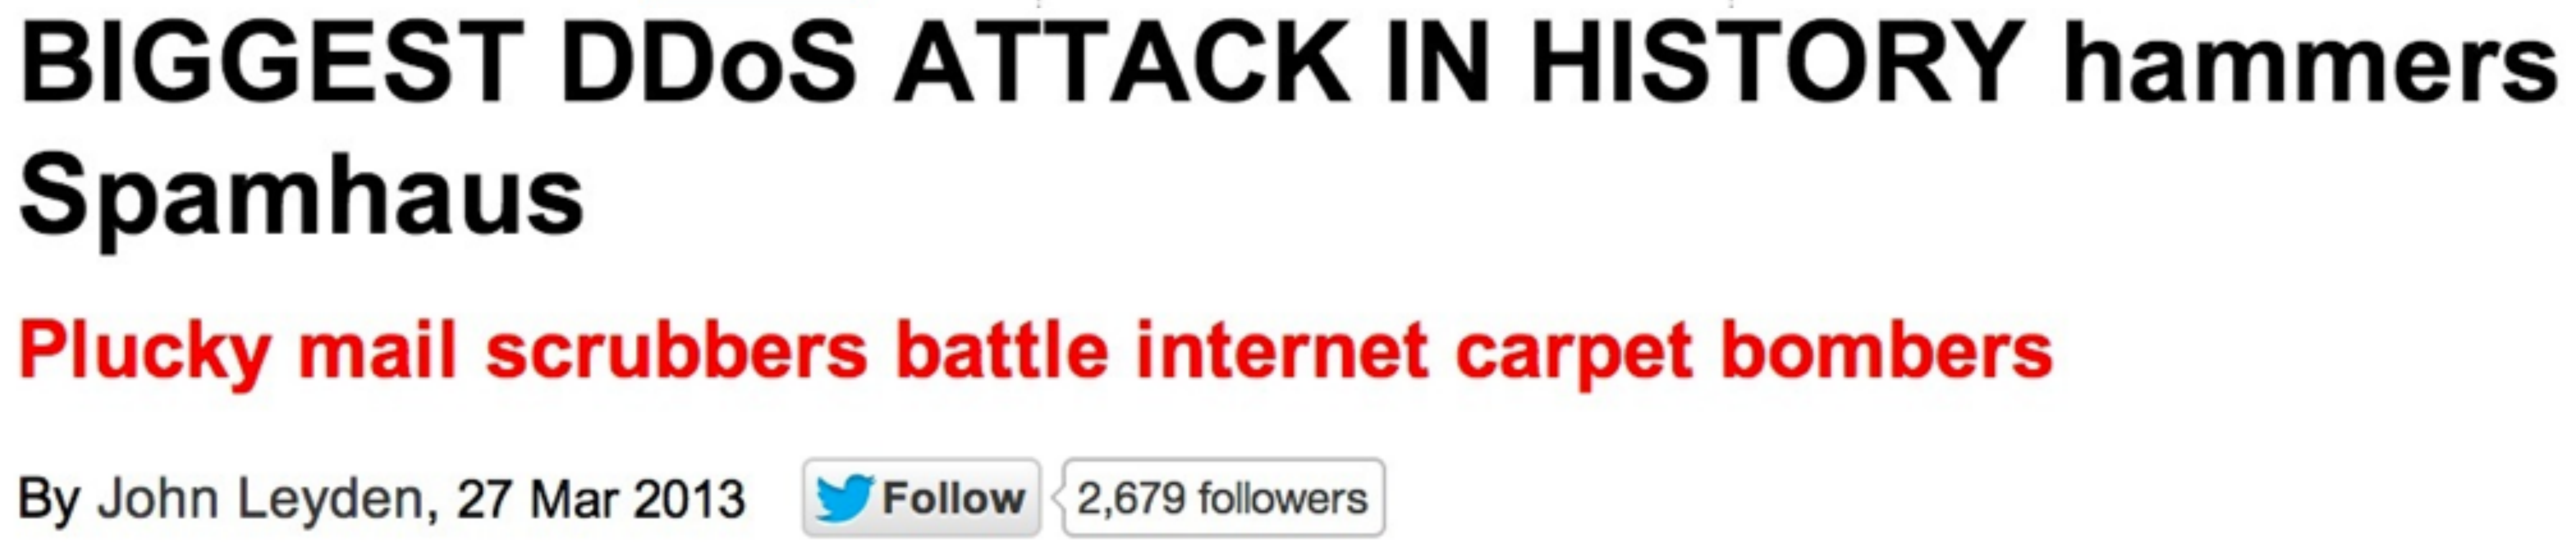
BIGGEST DDoS ATTACK IN HISTORY hammers Spamhaus

Plucky mail scrubbers battle internet carpet bombers

By John Leyden, 27 Mar 2013 Follow 2,679 followers

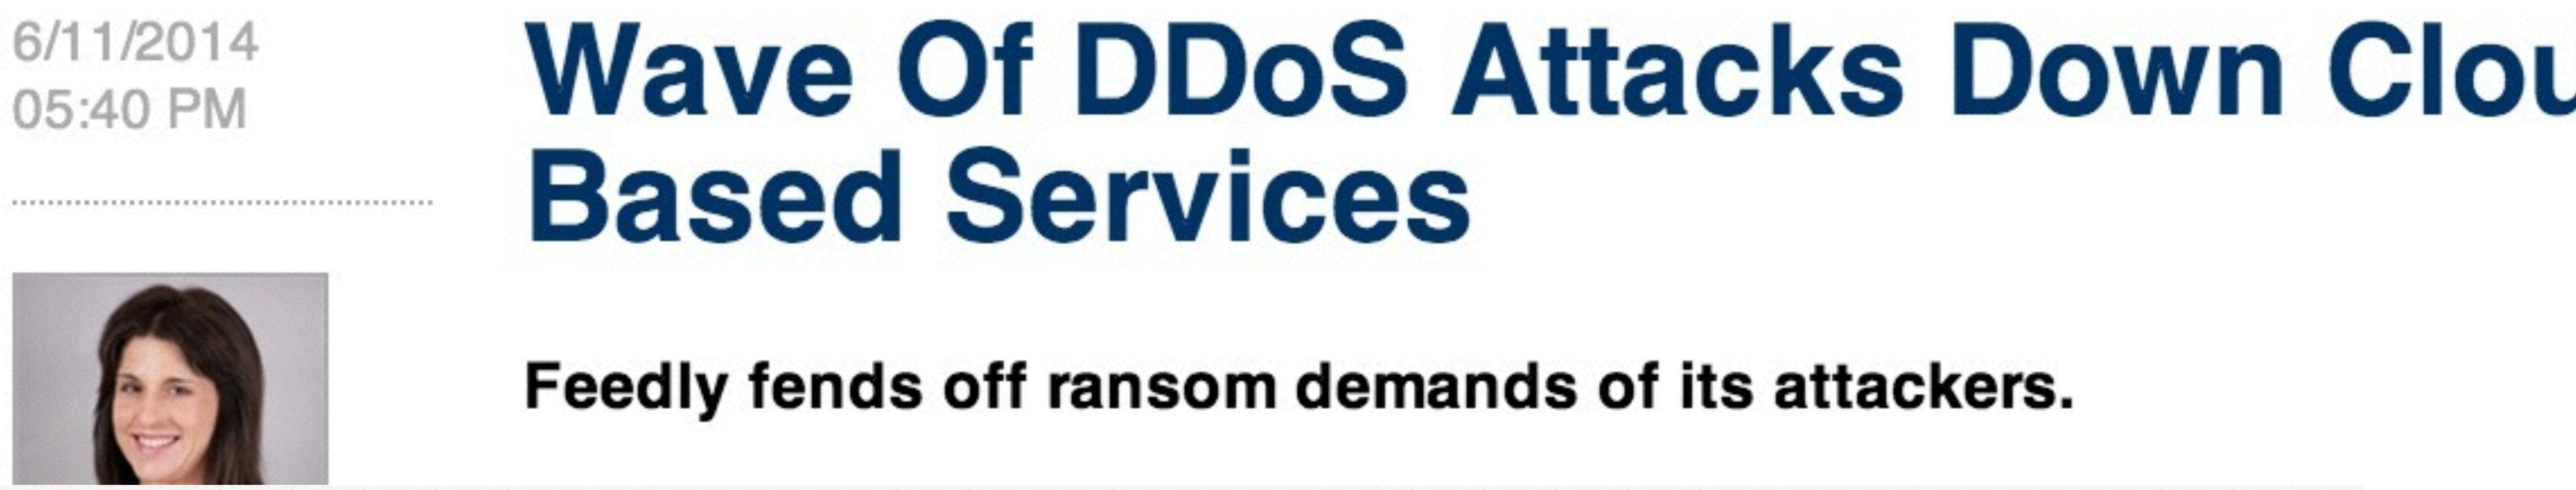
6/11/2014
05:40 PM

Wave Of DDoS Attacks Down Clou
Based Services

Feedly fends off ransom demands of its attackers.

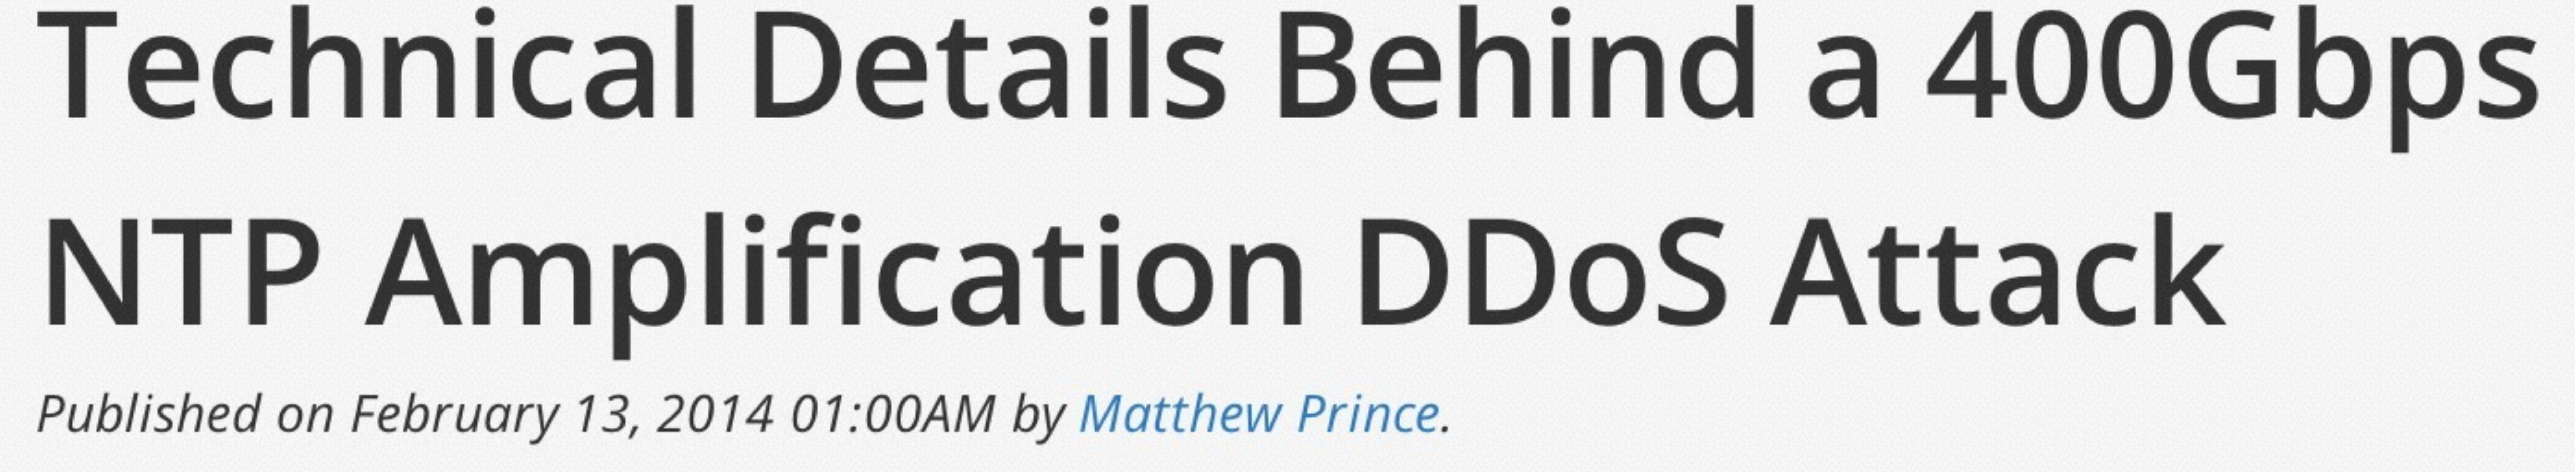
Technical Details Behind a 400Gbps NTP Amplification DDoS Attack

Published on February 13, 2014 01:00AM by Matthew Prince.

CloudFlare

- A secure reverse proxy for http(s)
  - Change your SOA to us
  - We will point your A records to us
- We need internal and external DNS to keep track

CloudFlare

- DNS Resolver

- Q: Who is something.com? → CloudFlare External DNS
- A: CloudFlare Proxy IP

CloudFlare

- Web browser

- Hi something.com, get me index.html → CF Proxy IP
- CF proxy: do I have index.html cached? No.
- CF proxy: who is something.com, really? → CF Internal DNS
- CF Internal DNS: origin IP → CF proxy
- CF proxy: Hi something.com, get me index.html → Origin IP
- Origin IP: index.html → CF proxy
- CF proxy: index.html → Web browser

CloudFlare External DNS

- Deals with attempted DDoS constantly
- Huge DNS floods of legitimate requests
  - 50+ million packets per second to one location
- Large volumetric reflection attacks
  - 300+ Gbps DNS reflection (2013, Spamhaus)
  - 400+ Gbps NTP reflection (2014)

CloudFlare External DNS

- Standard RRL not enough, need special filters
  - String matching
  - Length matching
  - Statistical approach: heavy hitters
  - Regular expressions

CloudFlare External DNS

- Other special feature: **CNAME flattening**
- Following CNAME records is slow
- Can’t CNAME the zone apex

- Solution: Follow CNAME chain, transform into A or AAAA record

What to do?

- How did we solve HTTP DDoS?
  - Anycast and a reverse proxy (nginx)

- How do we solve DNS DDoS?
  - Write your own DNS server? Maybe
  - Create a DNS reverse proxy? YES

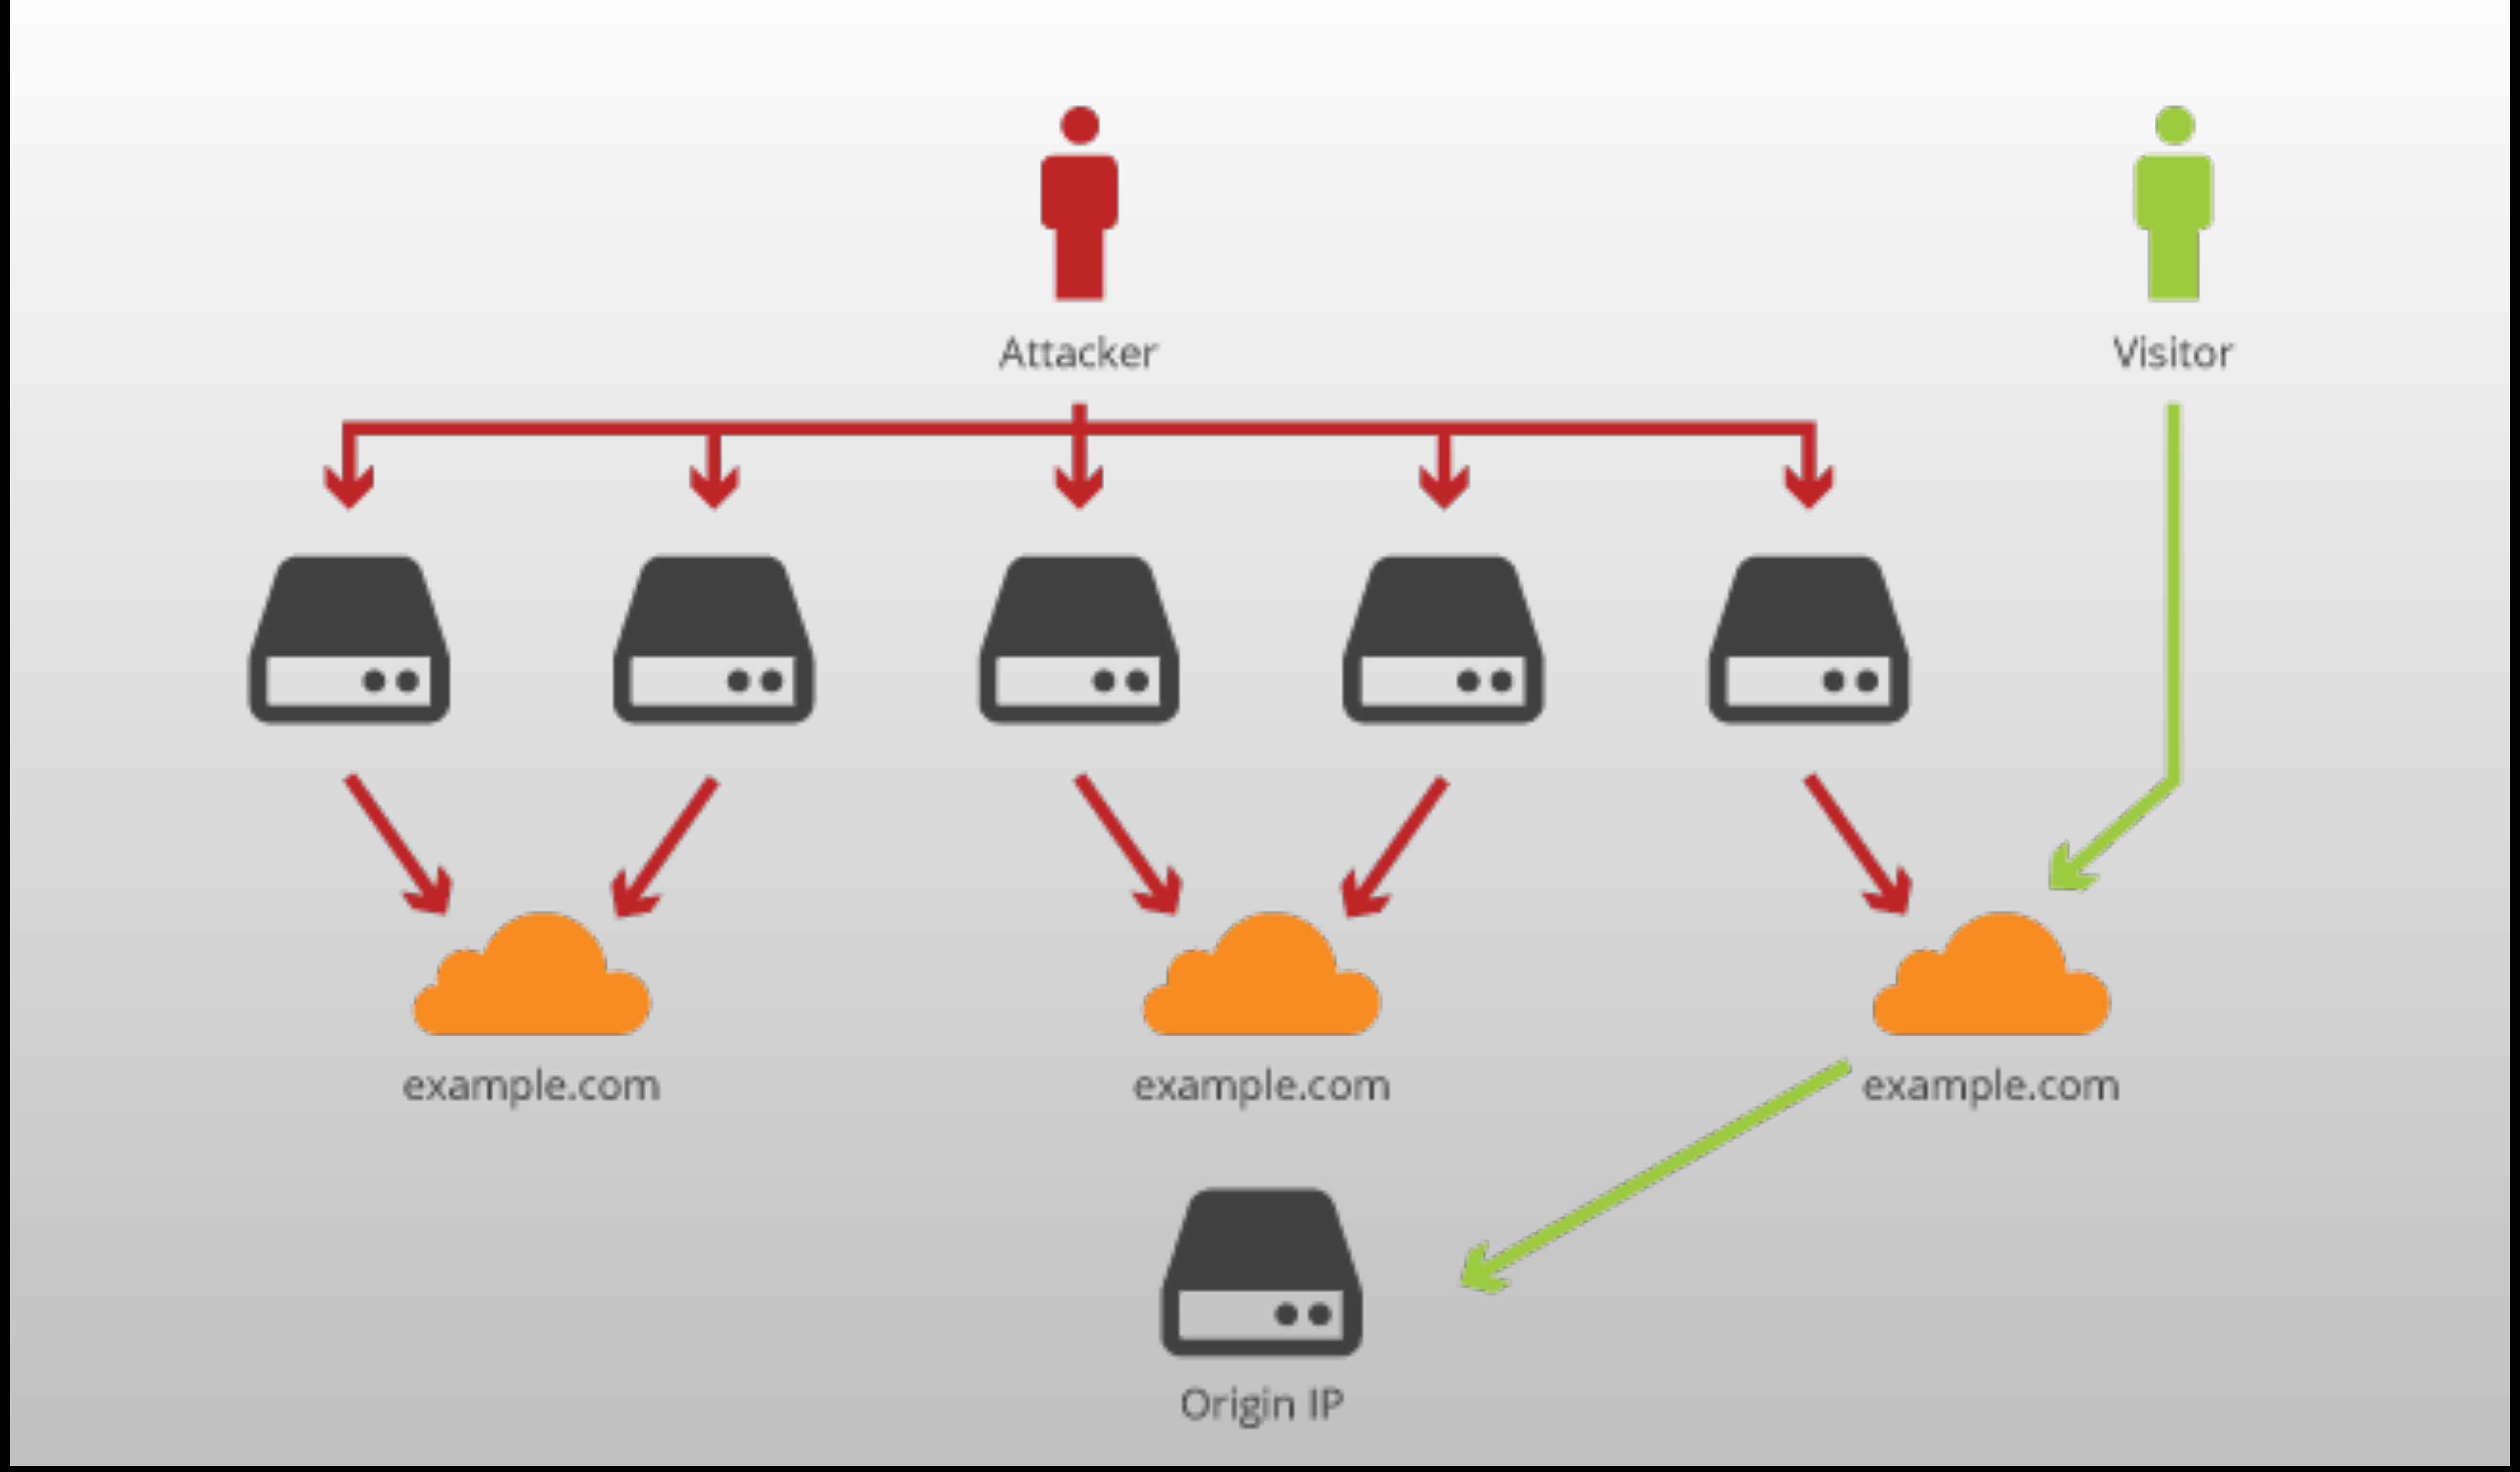

What to do?

- RRDNS: a DNS reverse proxy in Go

- Why Go?
  - compiled language gives great performance
  - built-in concurrency
  - easy to write, maintain, and make modular

What does it do?

- Acts as a transparent reverse proxy in front of an authoritative server

- Not a recursive nameserver
- Filters bad/spoofed requests, caches, load balances
- Returns the authoritative bit
- Responses look like ones from authoritative server

More advantages

- Highly dynamic
- Does not use zone files
- Automation reduces cost for operator

How we use it

- RRDNS handles both internal and external DNS
- Filter model inspired by nginx
  - SSL
  - WAF
  - Business logic
  - Cache
  - Upstream

How we use it

- RRDNS filter
  - front-line rate limit filtering
    - length & string matching, heavy hitter, IP reputation, geolocation, truncation test, etc.
  - request type filtering (limit to A, AAAA, CNAME, MX, etc.)
  - caching layer
  - optional authoritative module (for internal DNS)
  - upstream DNS resolution (for cache misses and CNAME resolution)

Where does DNSSEC fit in?

- Do it yourself behind the reverse proxy
- Let RRDNS take care of it

Pure Proxy DNSSEC

- Upstream manages all DNSSEC related data
- NSEC or NSEC3 records computed and served by upstream
- CloudFlare Internal DNS upstream:
  - Centralized offline signing with zone distribution over encrypted KV store

- Problems: CNAME flattening signatures unavailable
- Questions: Should proxy validate signatures from upstream?

Zone Enumeration

- NSEC or NSEC3 records computed offline
- Zone enumeration possible with NSEC
- Offline dictionary attack with NSEC3

- We want zone privacy, and CNAME flattening
- Solution: Live signing

Hybrid DNSSEC

- Upstream creates full DNSSEC zone (including NSEC3 records)
- Centralized offline signing with zone distribution over encrypted KV store
- KSK, ZSK1 used for offline signing (long lived)
- ZSK2 used for online signing of CNAME and NSEC3 white lies (short-lived)
- Under DDoS
  - serve real NSEC3 record
  - disable CNAME flattening

DNS Reverse Proxy as a service

- Large authoritative nameservers need Cloud DDoS protection, acceleration, caching
- Put CloudFlare/RRDNS in front

- What if they don't want to set up DNSSEC?
- Use RRDNS live signing!

Live DNSSEC

- Upstream creates regular non-DNSSEC zone
- KSK created centrally, DNSKEY RRSIG distributed to edge
- ZSK created centrally, distributed to edge servers via TPM binding
- ZSK used for live signing of all records
  - Flattened CNAME and NSEC3 white lies
- Live signatures stored in shared cache within a colocation
- CloudFlare integration with registrar

Result

- Authoritative servers get DDoS protection and acceleration
- Works with already integrated DNSSEC solution
- Or flip a switch and get DNSSEC automatically

Conclusion

- DNS is hard
- DNSSEC is hard
- Special problems require custom solutions

- Let us do DNSSEC for you
- But first: we have lots of work to do