DNS flag days (plural!)

2019 and beyond

2019-05-21 RIPE 78
Petr Špaček <petr.spacek@nic.cz>, CZ.NIC
Ondřej Surý <ondrej@isc.org>, ISC

Outline

- Motivation
- What is DNS flag day?
- 2019 - wrap up
- 2020? - early heads up

Motivation: Does DNS just work?

- Problem #1: DNS is complex (200 RFCs!)
- Hard to implement
- People make implementation mistakes
- Vendors add workarounds to improve interoperability
- With workarounds, it "just works"

Motivation: Workarounds ... so what

- Problem #2: DNS workarounds ossify
- Workarounds interact with
  - Standard protocol
  - Other workarounds!
- Workarounds **from 1999** causing breakage in **2018!**
- Breakage/cost incurred on compliant players
- **No incentive for non-compliant players to fix things**

DNS flag day: theory

- Trash pick up day!
- Software vendors + big DNS operators cooperate
- Workarounds get removed on certain date
- Shifts costs to non-compliant players
- Compliant players do nothing

DNS flag day

2019

DNS flag day: 2019 in practice

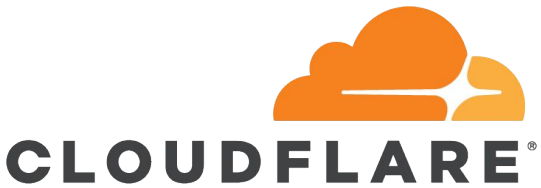

2019: Recap

- First time in history
- A lot of fear
- Misunderstandings
- Reach out campaign
- News articles
- Measurements

2019: T_0 - 3 months, sample 23 M domains

Mode	Permissive (<= 2018)	Strict (2019+)
OK	48.61 %	
Compatible	23.37 %	
High latency	13.15 %	7.48 %
Dead	14.87 %	20.55 %
Breakage		+5.68 %

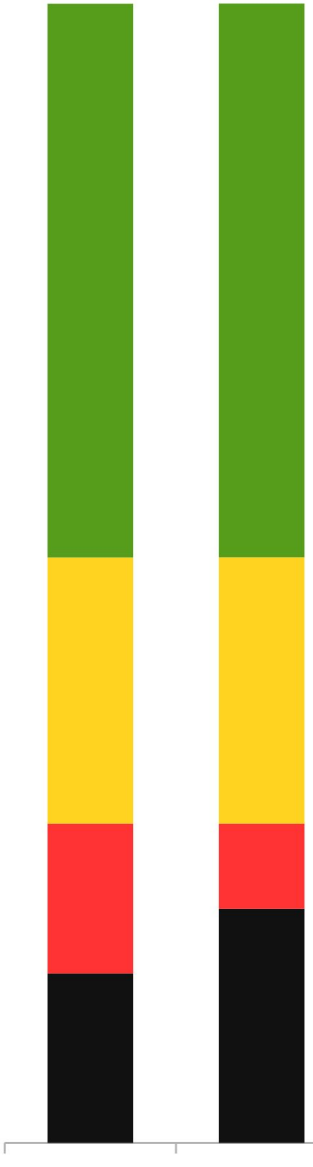

2019: T_0 - 3 months: clusters of breakage

provider domain	breakage	# broken
hichina.com.	35.78 %	469 611
dnspod.com.	25.66 %	336 797
myhostadmin.net.	5.04 %	66 208
xincache.com.	4.82 %	63 246
dnspod.net.	3.27 %	42 881
dnsdun.net.	2.85 %	37 435
gmoserver.jp.	2.71 %	35 595
registrar-servers.com.	1.64 %	21 533
alidns.com.	1.63 %	21 369
metaregistrar.nl.	1.20 %	15 762

∑

66 %

∑

85 %

Prepare for impact

2019: Did it work?

- **It did work**
- Cooperative community
- Vast majority domains fixed
- Remaining domains largely unused (parking ...)
- Support lines remained silent
- No measurable problems
- **Big thank you to all involved players!**

2019: Lessons learned

- We **can improve** Internet at global scale
  - As long as we cooperate
- Communication was a problem
  - Missing communication channel to operators
  - Thus this presentation!

DNS flag day **2020**

2020: Motivation

- **IP fragmentation does not work**
  - https://tools.ietf.org/html/draft-bonica-intarea-frag-fragile-03
  - http://www.potaroo.net/ispcol/2017-08/xtn-hdrs.html
- If IP fragmentation works, it is not secure enough
  - Research by Kazunori Fujiwara
    https://indico.dns-oarc.net/event/31/contributions/692/
- -> UDP is unsuitable for **large DNS messages**
- Operational issues around the globe

2020: Goal

- Eliminate operational issues caused by fragments
- Improve security of DNS
  - Also, think of domain validation ...

2020: Eliminating fragments

- **For large DNS answers** switch to TCP
  - No change for small answers - UDP
- Existing standards
  - DNS over TCP in RFC 7766 and predecessors
  - Default EDNS buffer size ~= 1220 (= never fragment)
- Non-compliance on several levels
  - Authoritative - do not listen on TCP
  - Authoritative - do not honor EDNS buffer size
  - Recursive (ignores TC=1)

2020: Advantages of TCP

- Hides IP fragmentation issues
- Harder to spoof
  - Low-throughput high-value services
    - CA domain validation
    - DNSSEC bootstrapping (CDS/CDNSKEY)
- Preparation for DNS-over-TLS

2020: Authoritative side (operations)

- **Honor RFC 7766 - DNS Transport over TCP**
- Answer on TCP port 53
  - **Check your firewall**, too!
- EDNS buffer size ~= 1220 to avoid fragmentation
  - Defaults in software will reflect this
- Authoritative MUST NOT send oversized answers
  - Standard compliant software does not require changes

2020: Resolver side (operations)

- **Honor RFC 7766**
- Answer on TCP port 53
  - **Check your firewall**, too!
- EDNS buffer size ~= 1220 to avoid fragmentation
  - Defaults in software will reflect this
- Resolvers MUST support fallback from UDP to TCP
  - Standard compliant software does not require changes

2020: Preliminary measurement

- ~ 7 % domains on servers not accepting TCP
  - Not all domains are equal
  - Includes parked domains etc.
- Breakage is very concentrated
- 1 operator > 70 %
- 9 operators > 90 %

TCP on auths in May 2019, 34 M domains, 59 TLDs

Mode	TCP as last instance	TCP required
OK	67.52 %	67.52 %
High latency	12.83 %	5.76 %
Dead	19.65 %	26.72 %
Breakage		+7.07 %

email, solutions, tel, date, review, one, link, services, company, agency, group, guru, news, network, photography, studio, jobs, business

net, co, xyz, se, cz, loan, online, club, site, icu, nz, shop, ltd, cl, mobi, app, live, pro, website, space, nu, fun, store, win, tech, men, life, blog, stream, world, dev, wang, bid, rocks, cat, tokyo, xxx, today, design, trade, xin

Top ten: TCP-broken providers in May 2019

	provider domain	breakage	# broken
∑	hichina.com	67.84 %	1 610 817
	name-services.com	6.74 %	160 070
	foundationapi.com	3.66 %	86 970
70 %	xincache.com	2.63 %	62 479
	alidns.com	2.16 %	51 309
	123-reg.co.uk	2.04 %	48 411
∑	domainparkingserver.net	1.69 %	40 036
	ztomy.com	1.27 %	30 238
	mytrafficmanagement.com	1.23 %	29 285
90 %	myhostadmin.net	1.05 %	24 856

2020: Testing manually

- **Tools with nice UI are coming**
- Manual test - all queries must succeed
  - $ dig +tcp @auth_IP yourdomain.example.
  - $ dig +tcp @resolver_IP yourdomain.example.
  - $ dig @resolver_IP test.knot-resolver.cz. TXT

2020: Test resolver configuration

- BIND
  - options { edns-udp-size 1220; };
- Knot Resolver
  - net.bufsize(1220)
- PowerDNS Recursor
  - edns-outgoing-bufsize=1220
- Unbound
  - server:
    edns-buffer-size: 1220

2020: What's missing

- Exact date
  - Measurements in progress
  - Targeting February 2020 - 9 months from now
- Exact EDNS buffer size value
  - 1220, 1232, 1280, ...
  - Will go into software defaults (there's **no time based trigger**)
- None of these change the principle
  - **DNS over TCP must work**

2020: Get in touch

- Web https://dnsflagday.net/
- Twitter https://twitter.com/dnsflagday
- Announcements: **https://lists.dns-oarc.net/mailman/listinfo/dns-announce**
- Questions: dns-operations@lists.dns-oarc.net
- Talk to us this week
  - NOGs around?

Questions?