Observations on Checksum Errors in DNS UDP Messages

Duane Wessels

October, 2011

VERISIGN

Motivation

- Artem Dinaburg wrote about “bitsquatting” mid-2011
- Bit-level errors in domain names result in misdirected traffic.
  - MICROSOFT.COM -> MICROSMFT.COM
- Dinaburg’s data show 96% of bit errors happen prior to DNS resolution.
- **Up to 4% happened during resolution**

- What evidence does Verisign see of bit errors?
- Whereas Dinaburg’s work focused on errors in query name, here we focus on messages with bad UDP checksusm.

Typical DNS/UDP/IPv4 Message

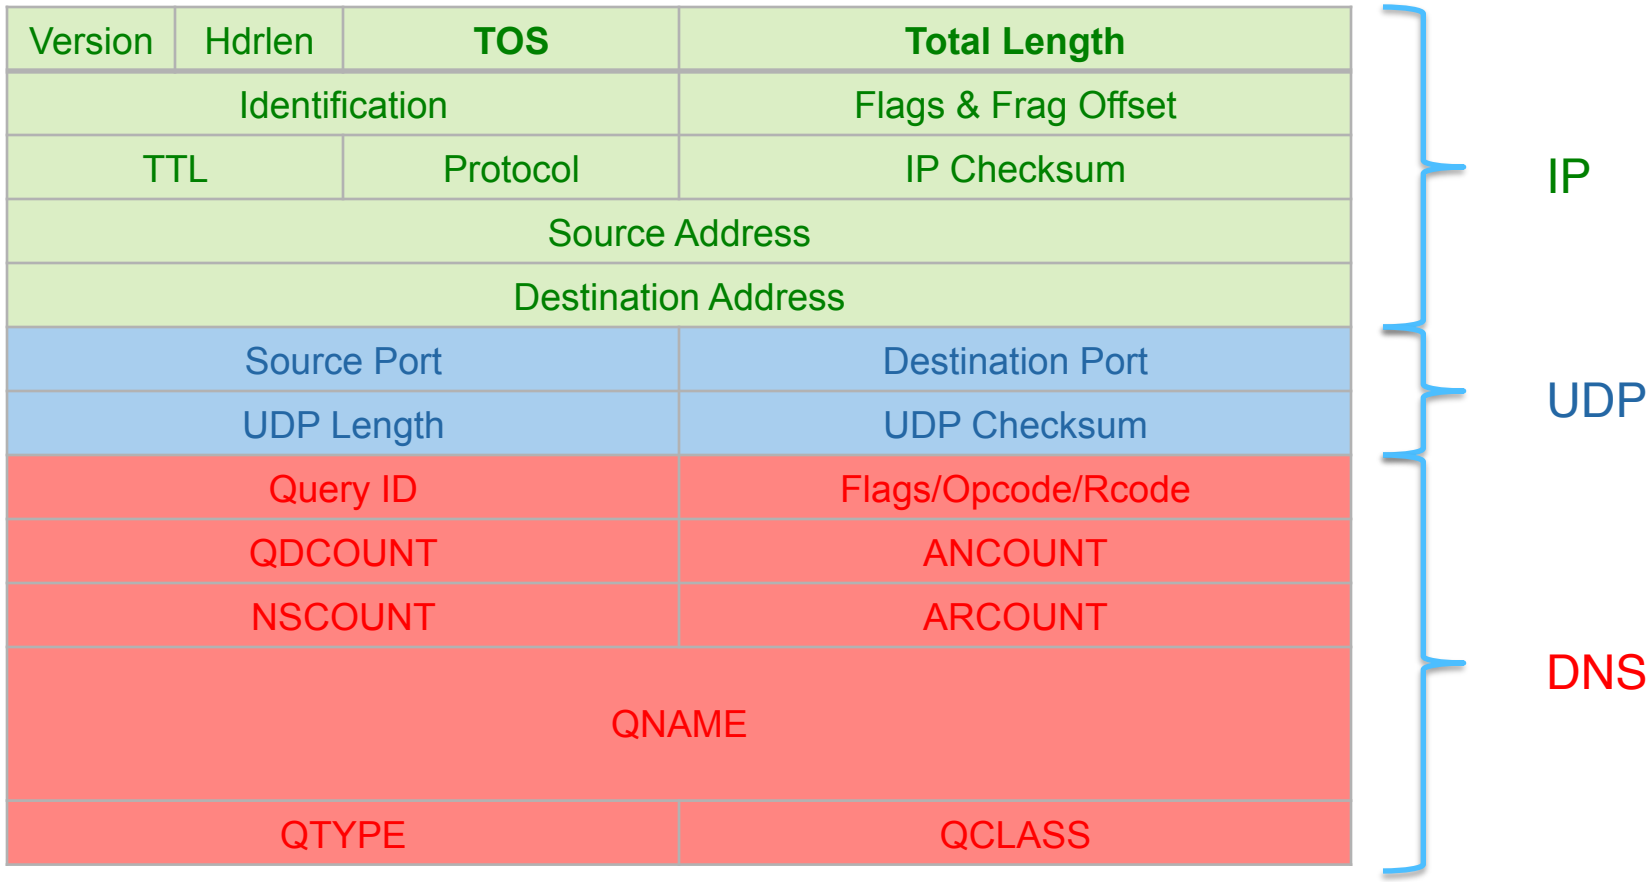

UDP Checksum

- RFC 768:

  Checksum is the 16-bit one's complement of the one's complement sum of a pseudo header of information from the IP header, the UDP header, and the data, padded with zero octets at the end (if necessary) to make a multiple of two octets.

- In other words:
  - Sum everything as 16-bit values
  - Overflow is wrapped around and added back in
  - Flip all bits in the answer, which is equivalent to saying

    checksum = -checksum

UDP Checksum

- Checksumming optional for senders
- Set to 0x0000 when disabled
  - If not disabled, and calculated checksum is 0x0000, then it is transmitted as 0xFFFF instead.

- Receiver includes sender's checksum in its calculation, which has the nice property that a correctly received message checksums to 0x0000.

UDP Checksum Calculation (Sender)

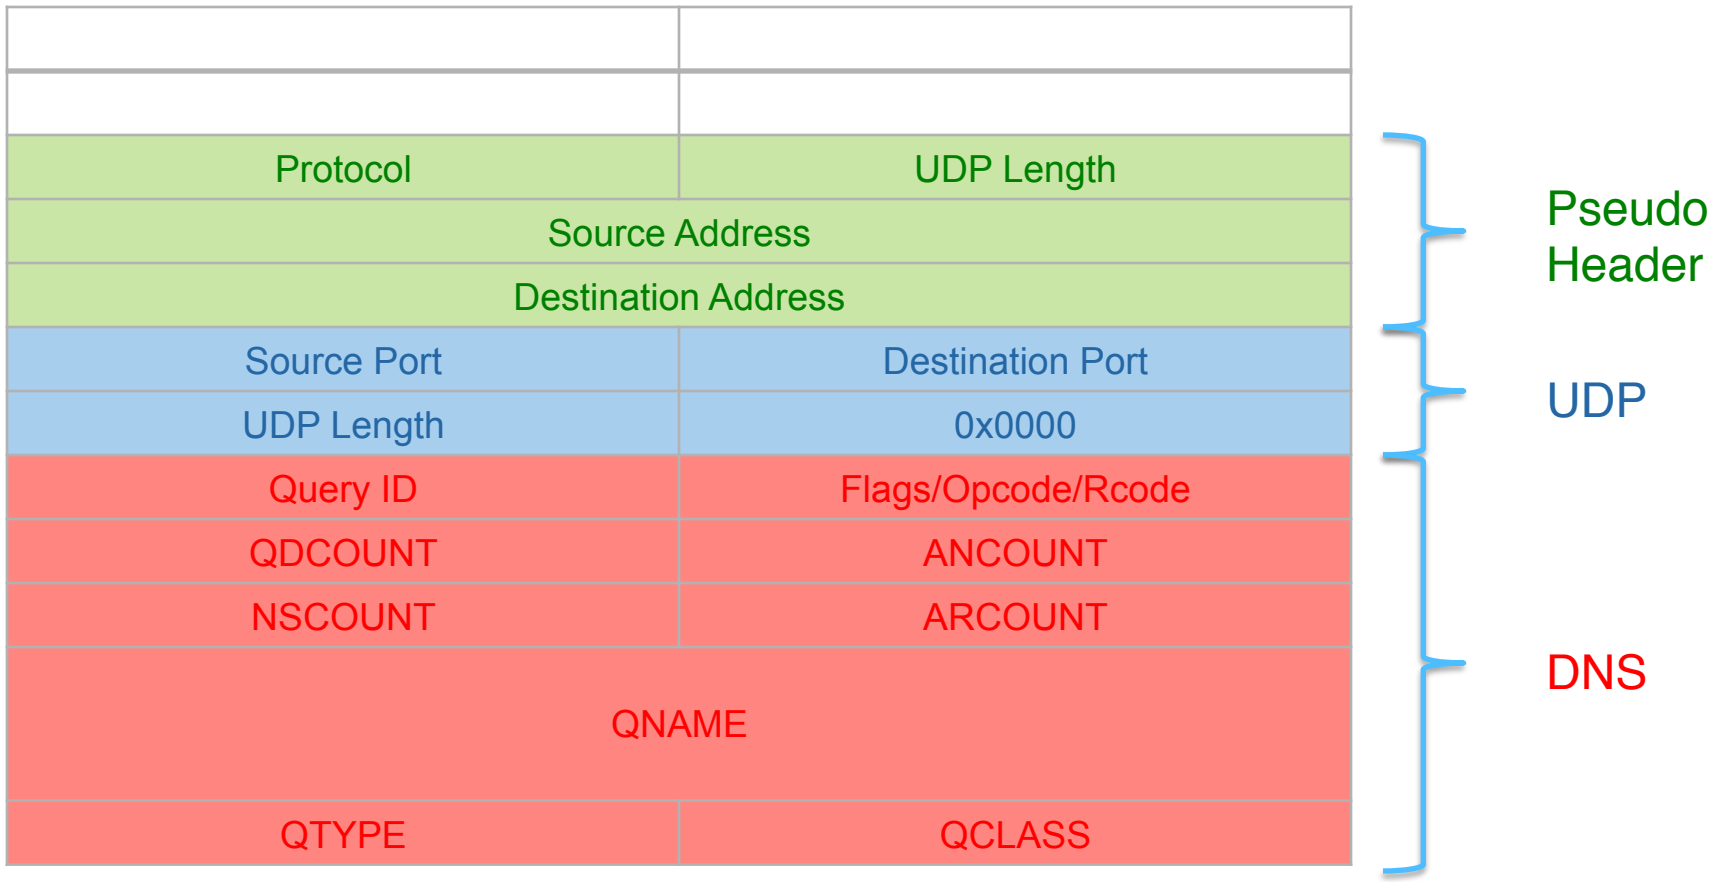

Simple Checksum Examples

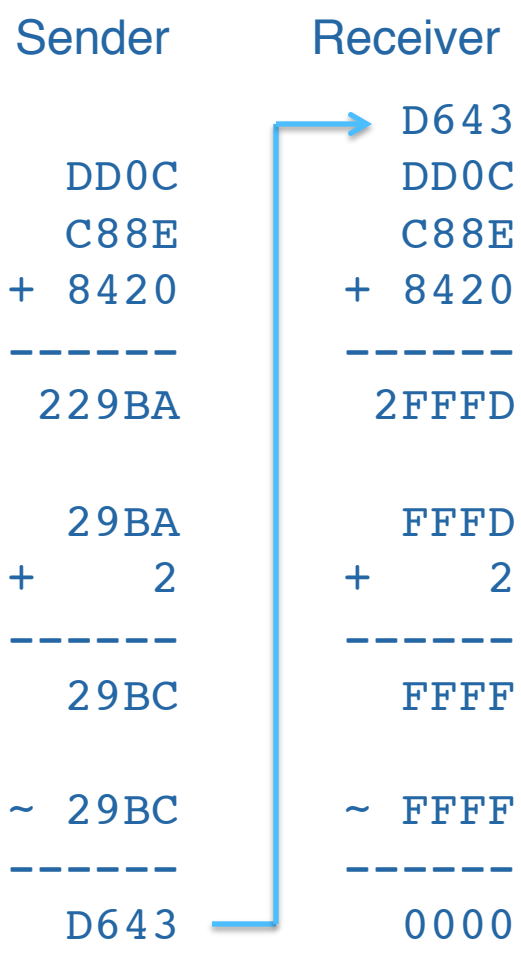

Checksums and Bitflips

- If a single bit is flipped during transmission, the difference in checksums tells which position (mod 16) the flipped bit was in.
- Thus, if a receiver-calculated checksum is equal to 2^n or 65535-2^n, it is *likely* that a bit in position *n* was flipped.

Sender	Receiver	Receiver
	7A9D	7A9D
0377	0377	0377
7777	777F	7773
+ 0A74	+ 0A74	+ 0A74
------	------	------
8562	10007	FFFB
~ 8562	0007	~ FFFB
------	+ 1	------
7A9D	------	0004
	0008	
	~ 0008	

	FFF7	

Multi-bit Flips Can Also Look Like Single Flips

Sender	Receiver
	974F
1234	1234
5678	5678
+ 0004	+ 0002
------	------
~ 68B0	~ FFFD
------	------
974F	0002

0100 ➔ 0010
Two bits flipped

Sender	Receiver
	9753
1234	1234
5678	5678
+ 0000	+ FFFE
------	------
~ 68AC	~ FFFE
------	------
9753	0001

15 bits flipped

Can we detect multiple bitflips?

- Not if the same position is flipped more than once.
- Yes, if multiple, independent bits are flipped in the same "direction."

DATA

VERISIGN

Data Source

- 24 hours of packets captured on backbone SPAN feed
  - 2011-09-08 (UTC)
- 5 sites (NYC3, SFO1, DFW2, IAD3, LON3)

- 13,031,158,230 Queries

How often are checksums disabled?

$$\frac{16{,}021{,}232}{13{,}031{,}158{,}230} = 0.13\%$$

How often are checksums wrong?

$$\frac{131,412}{13,031,158,230} = .001008\%$$

$$PER \cong 10^{-5}$$

PER = Packet Error Rate

Number of Wrong Bits in Bad Checksums

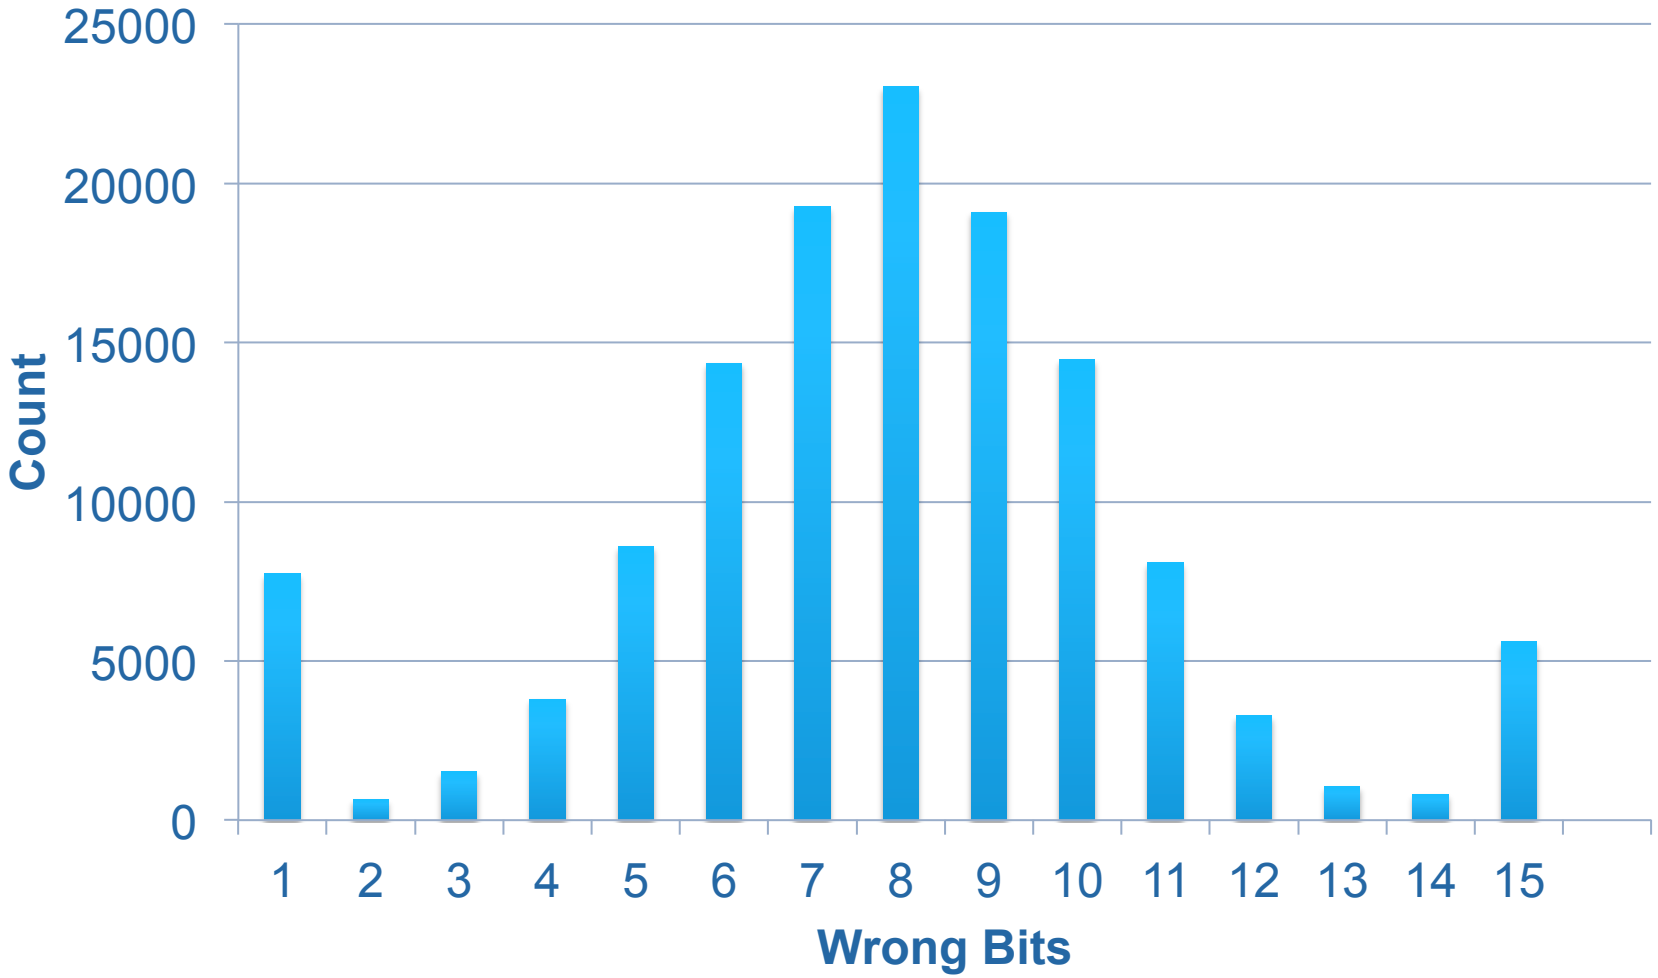

Sum of Two Distributions

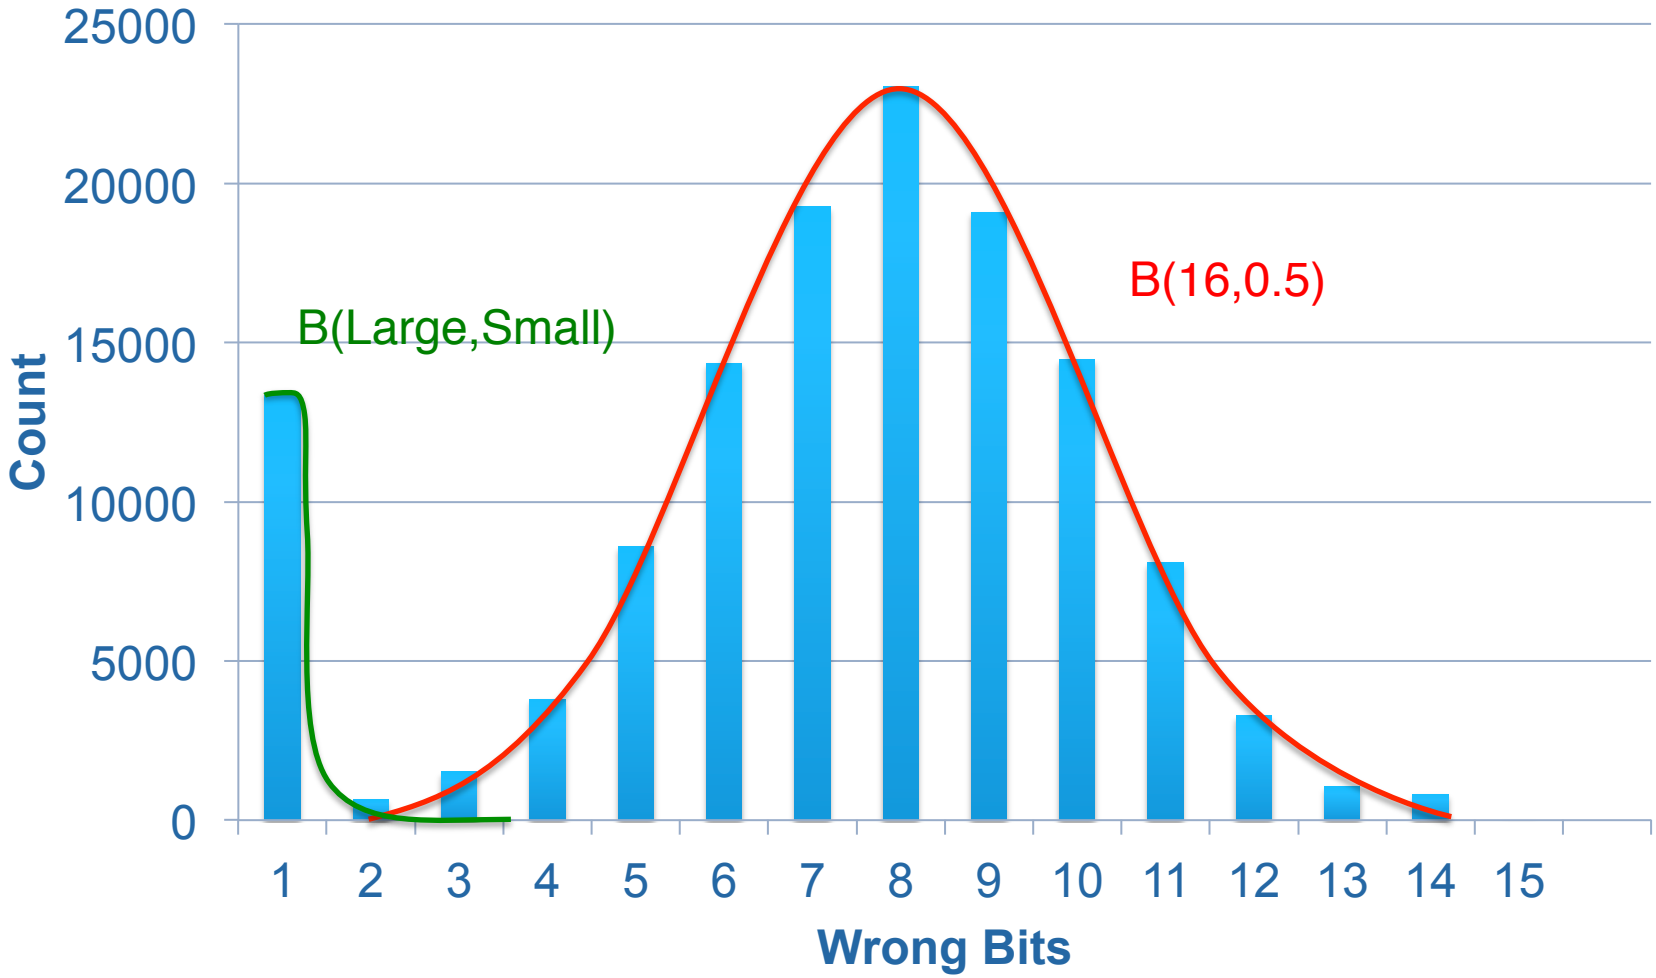

B(n,p) = Binomial Distribution

Binomial Distribution

- “the binomial distribution is the discrete probability distribution of the number of successes in a sequence of n independent yes/no experiments, each of which yields success with probability p.” --wikipedia

- B(n,p)
- n = number of trials (bits)
- p = probability of success (a bitflip)

Error Rate due to Single Bit Flips

- Assume 15-flipped-bits are really 1-flipped-bit in disguise.
- Number of single bit flips: 13,345
- Ignore other multi-bit flips
- What's the probability that a DNS message is corrupted by a single-bit transmission error?

$$\frac{13,345}{13,031,158,230} = 0.00000102$$

$$PER \cong 10^{-6}$$

Packet Error Rate to Bit Error Rate

Mean DNS (query) message size: 58 bytes, or 464 bits

$$10^{-6} = PER$$

$$10^{-6} = 1 - (1 - BER)^{N}$$

$$BER = 1 - \sqrt[N]{1 - 10^{-6}}$$

$$N = 464$$

$$BER = 2 \times 10^{-9}$$

Error Rate due to Multi Bit Flips

$$\frac{118,067}{13,031,158,230} = 0.00000906$$

$$PER = 9.06 \times 10^{-6}$$

In other words

- 10% of checksum errors due to single bit flips
  - Transmission errors?
- 90% of checksum errors due to … ?
  - Incorrect checksum algorithms?
  - Bad checksum offloaders?
  - Data-altering middleboxes?
  - Checksum-altering middleboxes?
  - Transmission noise bursts?

Is Our Channel Clean?

- We receive packets via backbone SPAN ports
- Perhaps errors are introduced between our servers and the packet collector?
- Let's look at DNS/UDP responses generated by our servers…

Responses	Errors	PER
12,957,669,961	1	7.7×10^{-11}

- Sample size too small.
- Receiver checksum = DFFF (single bit flip)

Which Bit Gets Flipped (single bit errors)

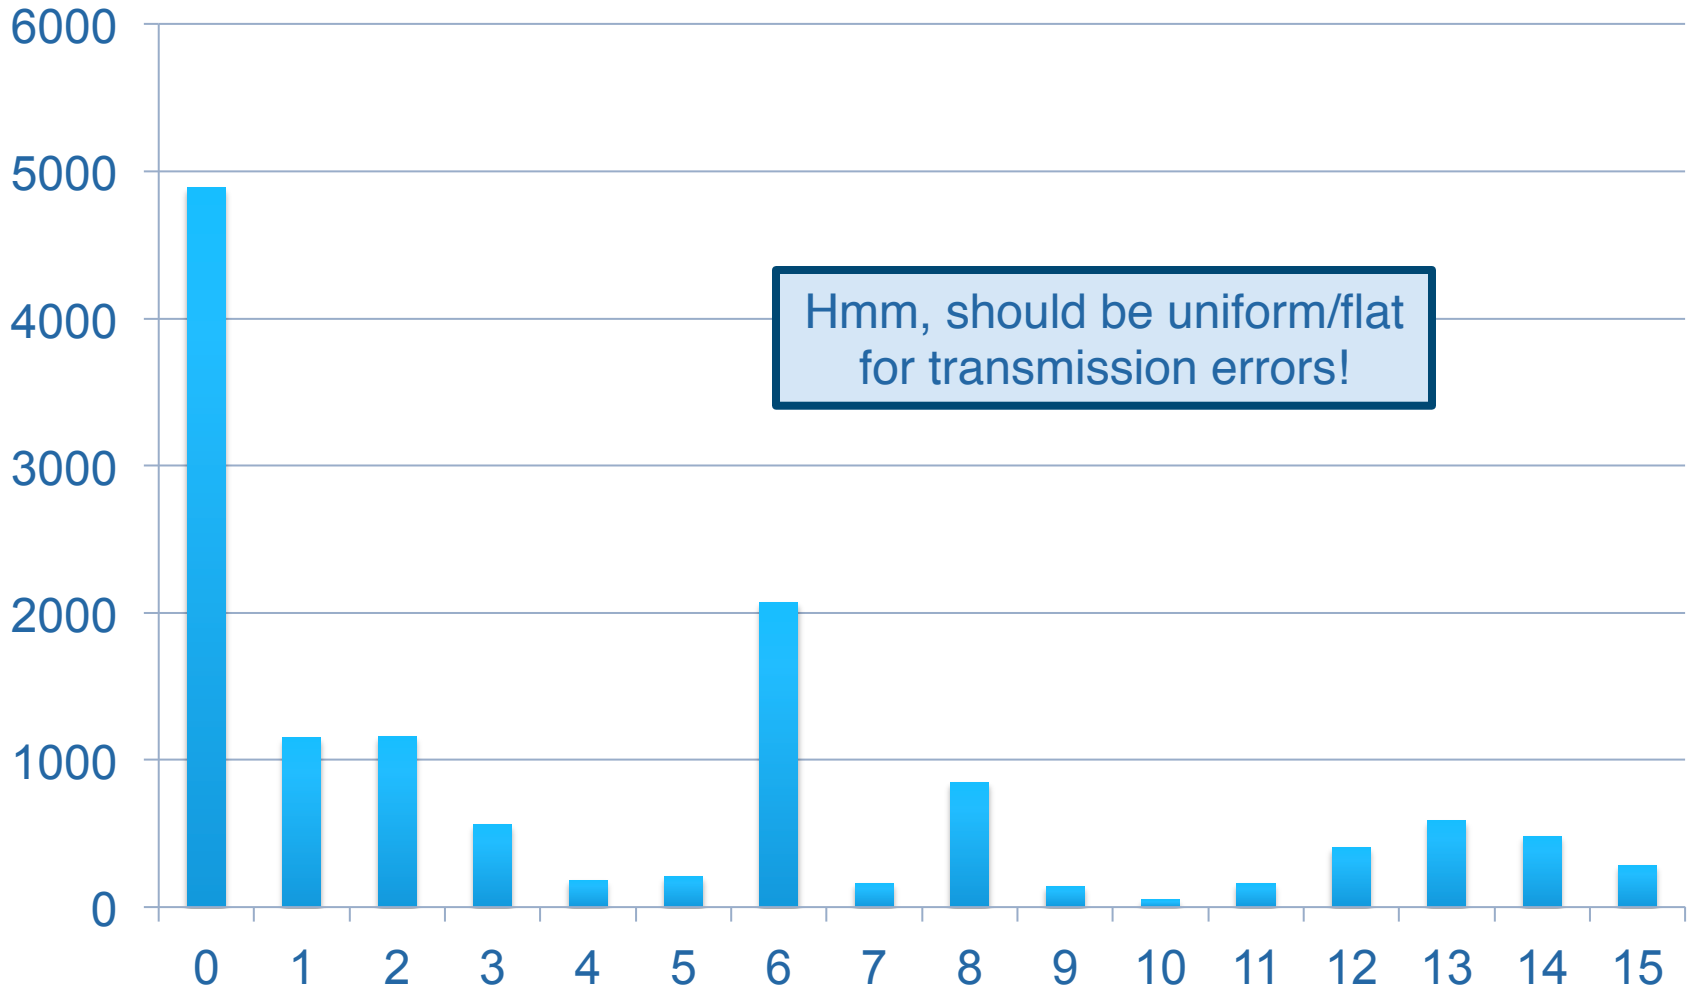

Number of Checksum Errors per Source IP

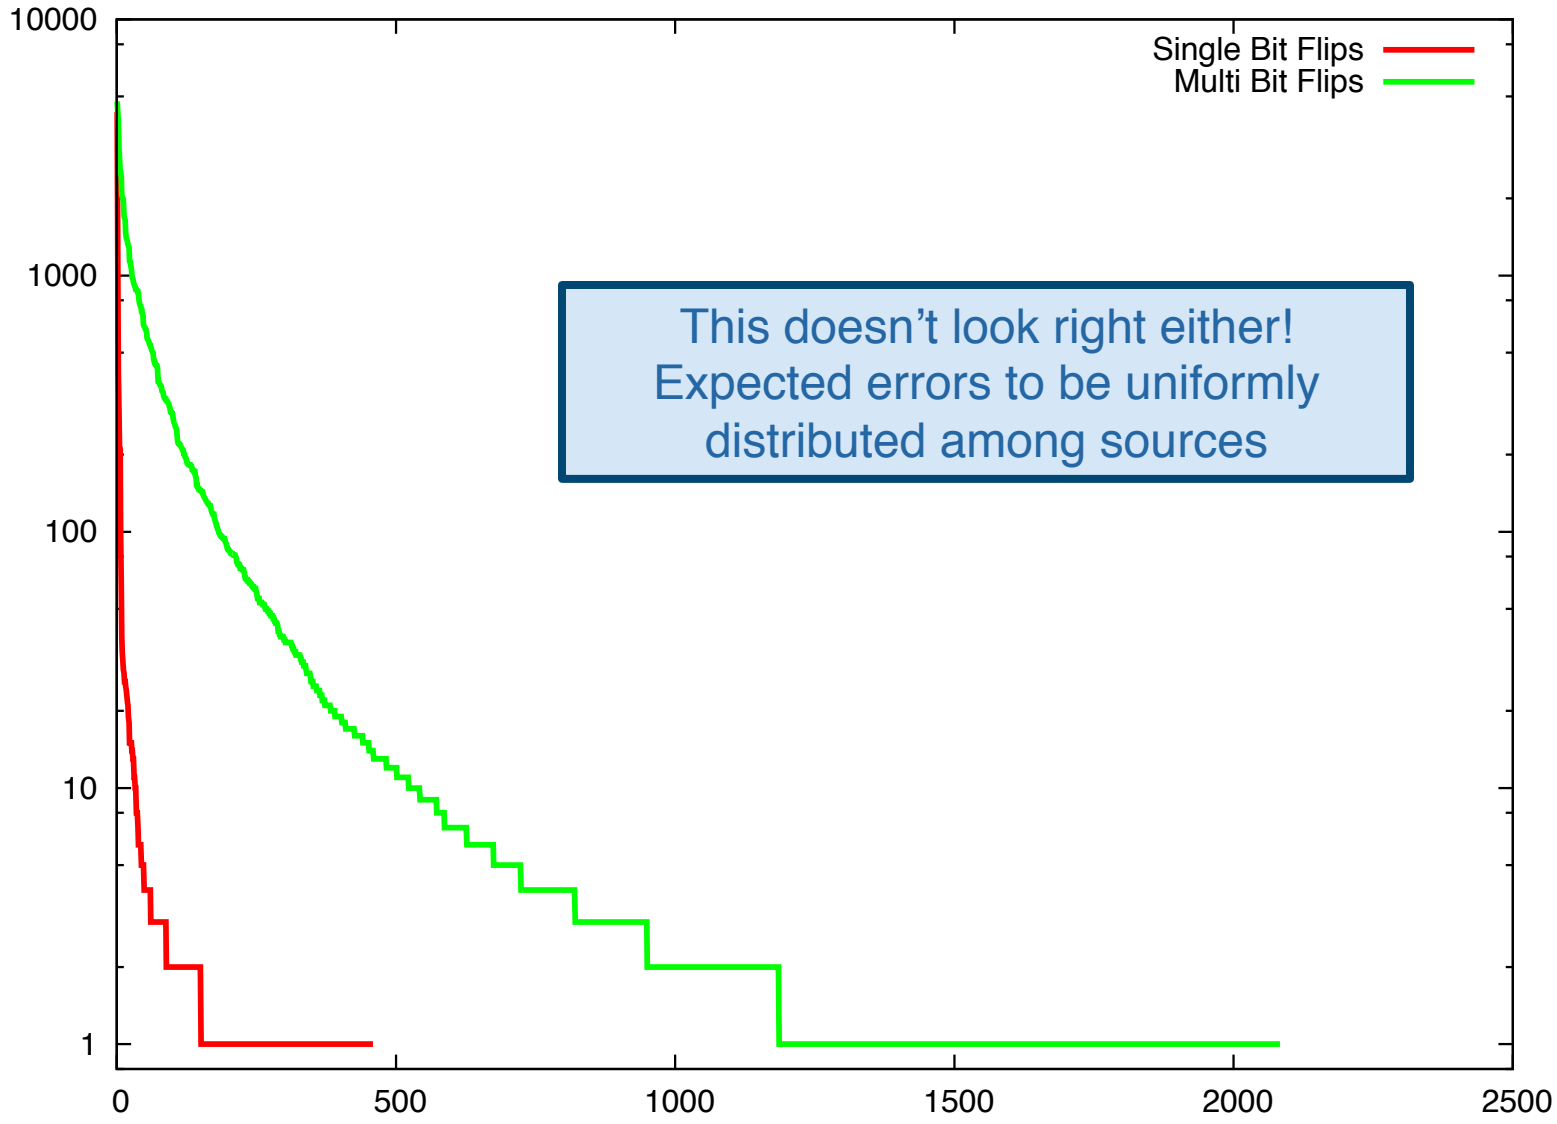

Fraction of Checksum Errors per Source IP

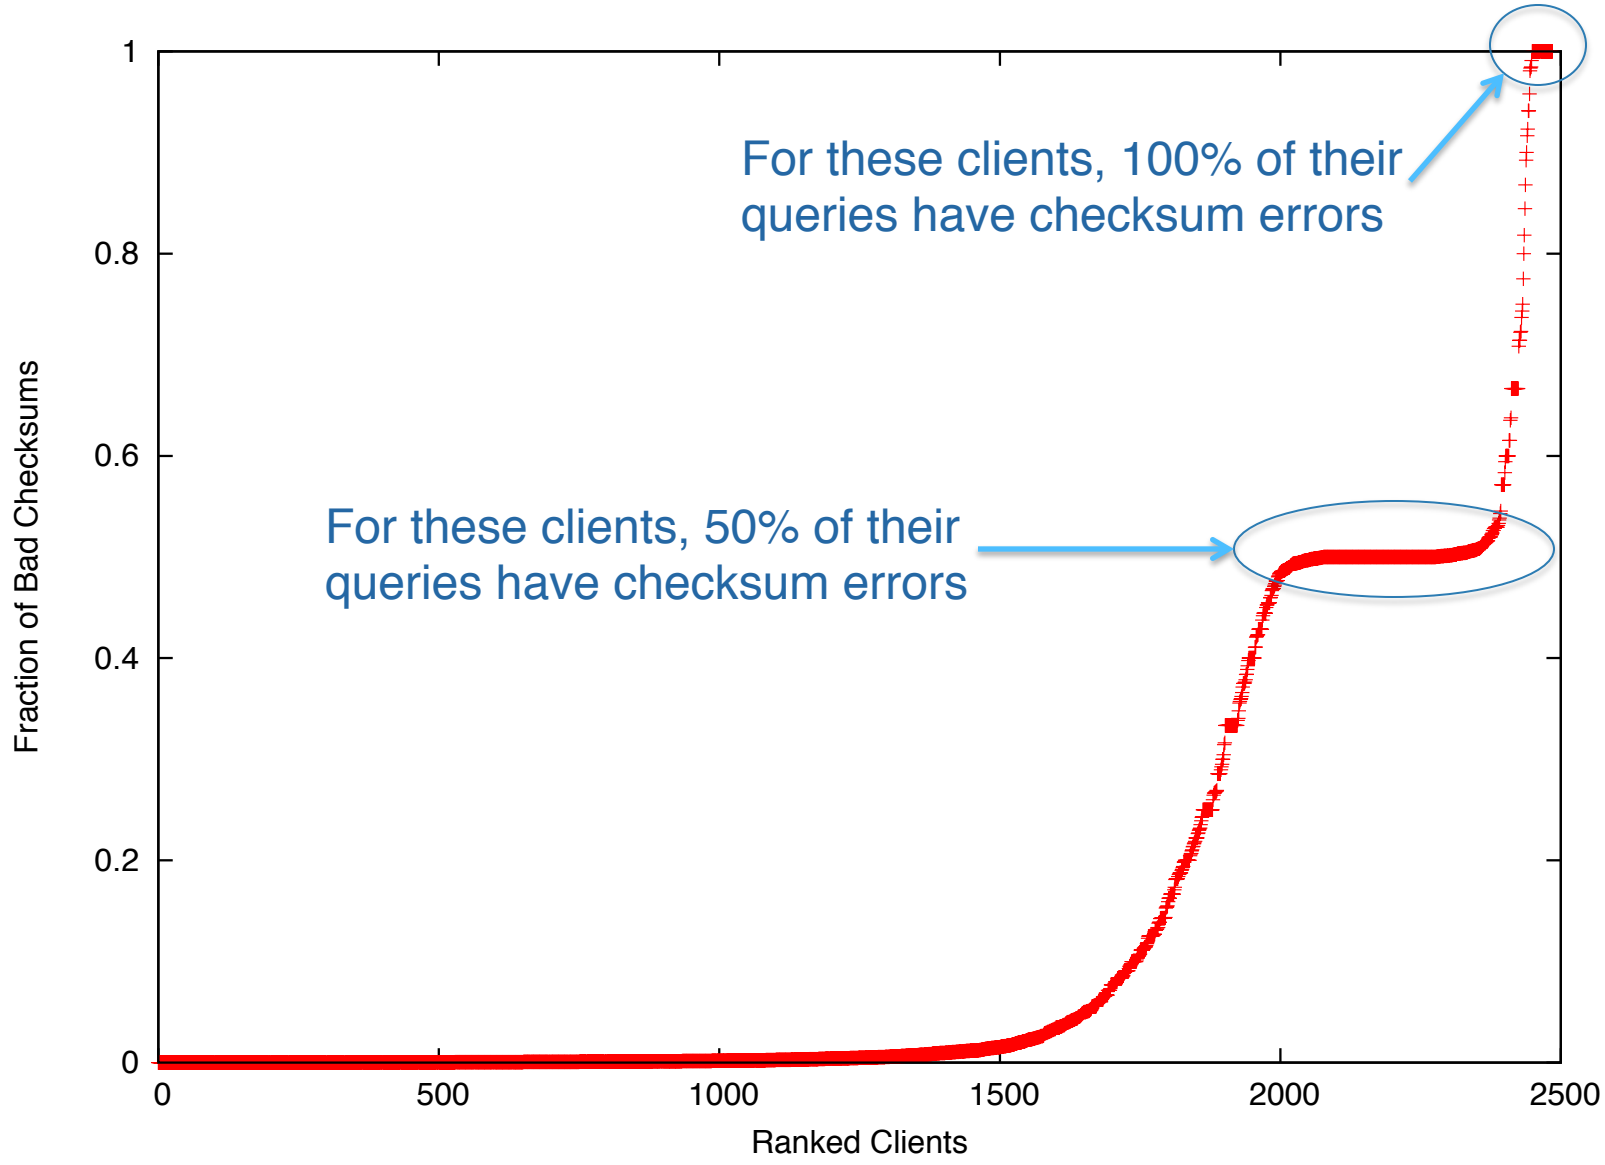

Examples

66.114.50.63

- 357 messages with bad checksums
- 0 messages with good checksums
- Most for ?.root-servers.net and ?.gtld-servers.net
- Usually, but not always, off by 1-bit
- Usually, but not always, for AAAA

```
00:04:45.977276 IP 66.114.50.63.32768 > 192.58.128.30.53: 3111% AAAA? m.gtld-servers.net. (36)
00:05:09.994614 IP 66.114.50.63.32768 > 192.58.128.30.53: 37477% AAAA? l.gtld-servers.net. (36)
00:05:09.994860 IP 66.114.50.63.32768 > 192.58.128.30.53: 56099% AAAA? c.gtld-servers.net. (36)
00:05:09.994861 IP 66.114.50.63.32768 > 192.58.128.30.53: 43707% AAAA? d.gtld-servers.net. (36)
00:05:09.994862 IP 66.114.50.63.32768 > 192.58.128.30.53: 15202% AAAA? i.gtld-servers.net. (36)
00:05:09.994986 IP 66.114.50.63.32768 > 192.58.128.30.53: 20215% AAAA? e.gtld-servers.net. (36)
00:05:09.994987 IP 66.114.50.63.32768 > 192.58.128.30.53: 32095% AAAA? j.gtld-servers.net. (36)
00:05:11.995590 IP 66.114.50.63.32768 > 192.58.128.30.53: 52463% AAAA? A.ROOT-SERVERS.NET. (36)
00:05:11.995591 IP 66.114.50.63.32768 > 192.58.128.30.53: 56664% AAAA? C.ROOT-SERVERS.NET. (36)
00:05:11.995592 IP 66.114.50.63.32768 > 192.58.128.30.53: 61100% AAAA? D.ROOT-SERVERS.NET. (36)
00:05:11.998389 IP 66.114.50.63.32768 > 192.58.128.30.53: 57738% AAAA? L.ROOT-SERVERS.NET. (36)
00:05:11.998429 IP 66.114.50.63.32768 > 192.58.128.30.53: 63608% AAAA? M.ROOT-SERVERS.NET. (36)
00:05:11.998430 IP 66.114.50.63.32768 > 192.58.128.30.53: 48723% AAAA? E.ROOT-SERVERS.NET. (36)
00:05:11.998542 IP 66.114.50.63.32768 > 192.58.128.30.53: 24474% AAAA? F.ROOT-SERVERS.NET. (36)
```

- Messages look normal; where is the error?

1.9.4.66

- Every query is duplicated
- One of the duplicates has wrong checksum
- 400 other sources like this

```
00:01:16.340277 IP 1.9.4.66.65314 > 192.42.93.30.domain: 21064 A? symantec.com. (30)
00:01:16.340278 IP 1.9.4.66.65314 > 192.42.93.30.domain: 21064 A? symantec.com. (30)
00:01:16.340386 IP 192.42.93.30.domain > 1.9.4.66.65314: 21064- 0/6/2 (233)
00:04:26.451905 IP 1.9.4.66.53372 > 192.26.92.30.domain: 4532 A? www.isupolitiksemasa.net. (42)
00:04:26.451906 IP 1.9.4.66.53372 > 192.26.92.30.domain: 4532 A? www.isupolitiksemasa.net. (42)
00:04:26.452052 IP 192.26.92.30.domain > 1.9.4.66.53372: 4532- 0/2/2 (129)
00:04:26.479518 IP 1.9.4.66.53372 > 192.42.93.30.domain: 4532 A? www.isupolitiksemasa.net. (42)
00:04:26.479521 IP 1.9.4.66.53372 > 192.42.93.30.domain: 4532 A? www.isupolitiksemasa.net. (42)
00:04:26.481201 IP 192.42.93.30.domain > 1.9.4.66.53372: 4532- 0/2/2 (129)
00:04:33.471917 IP 1.9.4.66.51252 > 192.42.93.30.domain: 48411 A? ns.second-ns.com. (34)
00:04:33.471933 IP 1.9.4.66.51252 > 192.42.93.30.domain: 48411 A? ns.second-ns.com. (34)
00:04:33.472040 IP 192.42.93.30.domain > 1.9.4.66.51252: 48411- 0/3/1 (124)
00:05:38.536049 IP 1.9.4.66.49769 > 192.26.92.30.domain: 43629 A? g1.panthercdn.com. (35)
00:05:38.536455 IP 1.9.4.66.49769 > 192.26.92.30.domain: 43629 A? g1.panthercdn.com. (35)
00:05:38.536551 IP 192.26.92.30.domain > 1.9.4.66.49769: 43629- 0/2/2 (103)
00:06:23.624261 IP 1.9.4.66.58598 > 192.26.92.30.domain: 48897 A? mmi.explabs.net. (33)
00:06:23.624363 IP 1.9.4.66.58598 > 192.26.92.30.domain: 48897 A? mmi.explabs.net. (33)
00:06:23.624456 IP 192.26.92.30.domain > 1.9.4.66.58598: 48897- 0/3/3 (157)
00:09:17.086740 IP 1.9.4.66.52698 > 192.42.93.30.domain: 31885 A? www.adobe.com. (31)
00:09:17.086741 IP 1.9.4.66.52698 > 192.42.93.30.domain: 31885 A? www.adobe.com. (31)
00:09:17.086974 IP 192.42.93.30.domain > 1.9.4.66.52698: 31885- 0/4/4 (203)
```

63.146.170.241

- Source sent 205,439 queries
- 5 had bad checksums
  - Single bit
  - Position 7 or 15

```
01:50:39.118685 IP 63.146.170.241.44889 > 192.55.83.30.53: 46274 [1au] MX? latingraf.com. (42)
18:32:06.849303 IP 63.146.170.241.51292 > 192.26.92.30.53: 54241% [1au] A? ns3.covad.com. (42)
01:19:39.607501 IP 63.146.170.241.44309 > 192.42.93.30.53: 51316 [1au] MX? brightroll.com. (43)
01:32:39.348499 IP 63.146.170.241.10057 > 192.42.93.30.53: 44656 [1au] MX? tampatrib.com. (42)
05:48:08.852933 IP 63.146.170.241.13068 > 192.42.93.30.53: 26515 [1au] MX? corwinford.com. (43)
```

- Messages look normal?

12.168.71.66

- 135 Queries, 114 bad checksums
- All ending with bluecoat.com
- Single/double bitflips, positions 12,13,14

```
00:29:58.244777 IP 12.168.71.66.39965 > 198.41.0.4.domain: 43797+ A? av-begister.bluecoat.com. (42)
00:30:59.237624 IP 12.168.71.66.40036 > 198.41.0.4.domain: 39957+ A? av-Begister.bluecoat.com. (42)
00:31:51.241279 IP 12.168.71.66.40154 > 198.41.0.4.domain: 44309+ A? hb.bluecoat.com. (33)
00:32:51.232745 IP 12.168.71.66.40236 > 198.41.0.4.domain: 36373+ A? smtp.bluecoat.com. (35)
00:57:16.193311 IP 12.168.71.66.42116 > 198.41.0.4.domain: 44821+ A? download.bluecoat.com. (39)
00:58:17.196258 IP 12.168.71.66.42169 > 198.41.0.4.domain: 53269+ A? dowNload.bluecoat.com. (39)
01:22:20.151177 IP 12.168.71.66.43878 > 198.41.0.4.domain: 50453+ A? serFices.bluecoat.com. (39)
04:46:03.804729 IP 12.168.71.66.57706 > 198.41.0.4.domain: 59925+ A? av-Register.bluecoat.com. (42)
08:28:59.426377 IP 12.168.71.66.21790 > 198.41.0.4.domain: 4118+ A? dow^load.bluecoat.com. (39)
08:30:00.423890 IP 12.168.71.66.21842 > 198.41.0.4.domain: 4374+ A? dowNload.bluecoat.com. (39)
08:54:52.380529 IP 12.168.71.66.23358 > 198.41.0.4.domain: 1046+ A? hb.bluecoat.com. (33)
08:55:52.385696 IP 12.168.71.66.23436 > 198.41.0.4.domain: 1302+ A? smtp.bluecoat.com. (35)
09:06:13.362615 IP 12.168.71.66.24262 > 198.41.0.4.domain: 2070+ A? av-begister.bluecoat.com. (42)
09:07:14.359314 IP 12.168.71.66.24344 > 198.41.0.4.domain: 6422+ A? av-Register.bluecoat.com. (42)
09:26:52.327399 IP 12.168.71.66.25746 > 198.41.0.4.domain: 2838+ A? hb.bluecoat.com. (33)
09:27:52.331050 IP 12.168.71.66.25811 > 198.41.0.4.domain: 3094+ A? smtp.bluecoat.com. (35)
```

84.235.6.32/27

- Too aggressive with 0x20 hack
- Spills over into qdcount, qtype, qclass fields
- ...but why is checksum wrong?
- Affects small percent of queries from these sources

src	cksum_is	cksum_ex	qd	an	ns	ar	type	class	name	nflip	pos
84.235.6.43	0xfddf	0xbddf	16385	0	0	0	1	1	iocTGNoWrmzWas.COm	1	14
84.235.6.43	0xac9a	0x6c9a	16385	0	0	0	1	1	www.tUbeseXMOvieS.coM	1	14
84.235.6.43	0x3c82	0xfc81	16385	0	0	0	1	1	wWW.bLOGoutILs.COM	1	14
84.235.6.43	0xf4bc	0xb4bc	1	0	0	0	28	1	pro.hit.GEmIus.pL.aFaqEre.COm	1	14
84.235.6.43	0xe80a	0xa80a	1	0	0	0	1	65	TidSPFtjPerHUuw.nEt	1	14
84.235.6.43	0x06e4	0xc6e3	1	0	0	0	16385	1	uyVaAHllvd.afaqE2e.cOm	1	14
84.235.6.43	0x0533	0xc532	1	0	0	0	25423	27904	<Unknown extended label>	1	14
84.235.6.43	0xcc97	0x8c97	16385	0	0	0	1	1	ns2.DigIzAAl.net	1	14
84.235.6.43	0xb76b	0x776b	1	0	0	0	25455	27911	<Unknown extended label>	1	14
84.235.6.43	0x6c21	0x2c21	1	0	0	0	16385	1	NqMMKBfgxv.afaqe2e.COm	1	14
84.235.6.43	0xbc7e	0x7c7e	16385	0	0	0	1	1	tEiKIQgzob.AfAqE2E.cOm	1	14
84.235.6.43	0x9d71	0x5d71	1	0	0	0	16385	1	ZAlSJpojse.aFaQE2e.com	1	14
84.235.6.43	0x8974	0x4974	1	0	0	0	16385	1	www.Horror-ExTREme.Com	1	14
84.235.6.43	0x50c3	0x10c3	16385	0	0	0	1	1	bAcKuPsMtp.p1Tp.com	1	14
84.235.6.43	0xb886	0x7886	1	0	0	0	1	16385	fly.CHeapsHARIng.CoM	1	14
84.235.6.43	0x1fb2	0xdfb1	1	0	0	0	16385	1	LkiEFbxueq.aFAQE2E.com	1	14
84.235.6.43	0x1618	0xd617	16385	0	0	0	1	1	zEbra.lIVecHaTNow.Com	1	14
84.235.6.43	0x6009	0x2009	16385	0	0	0	1	1	www.samMySgAMebOx.com	1	14
84.235.6.43	0xf36d	0xb36d	16385	0	0	0	1	1	VzVrpekgio.AFaQE2E.CoM	1	14

Numerous

- qclass=8193 and qname length = 18 is common
- Always bit 13 flipped
- 300 queries like this

src	cksum_is	cksum_ex	qd	an	ns	ar	type	class	name	nflip	pos
186.200.45.42	0x4024	0x2024	1	0	0	1	1	8193	www.googlelabs.com	1	13
186.201.201.54	0xe964	0xc964	1	0	0	0	1	8193	ns2.p31.dynect.net	1	13
186.201.201.54	0x21c6	0x01c6	1	0	0	0	1	8193	f1g1ns1.dnspod.net	1	13
186.201.201.66	0xc4e9	0xa4e9	1	0	0	0	1	8193	www.culturamix.com	1	13
186.202.21.40	0x9a99	0x7a99	1	0	0	0	1	8193	ns2.controlhop.com	1	13
186.224.0.18	0x3193	0x1193	1	0	0	0	1	8193	web17.dnsgeral.com	1	13
186.224.0.20	0x8d69	0x6d69	1	0	0	0	1	8193	ns5.brazilmall.net	1	13
186.238.86.7	0x0f6e	0xef6d	1	0	0	0	1	8193	i89jn2nte.y88a3r5l	1	13
186.238.86.7	0x78a9	0x58a9	1	0	0	0	1	8193	ytimg.l.google.com	1	13
186.238.86.7	0x4d1f	0x2d1f	1	0	0	0	1	8193	ad.foxnetworks.com	1	13
187.9.11.50	0xff6a	0xdf6a	1	0	0	1	1	8193	d.ns.swiftcdn1.com	1	13
187.10.52.205	0xf76f	0xd76f	1	0	0	0	15	8193	suresportsbets.com	1	13
187.11.35.189	0x4013	0x2013	1	0	0	0	15	8193	wrightbrothers.net	1	13
187.11.132.47	0xebda	0xcbda	1	0	0	0	15	8193	nightclubitems.com	1	13
187.11.137.192	0x4410	0x2410	1	0	0	1	1	8193	b.gtld-servers.net	1	13
187.11.185.103	0xd457	0xb457	1	0	0	1	28	8193	f.gtld-servers.net	1	13
187.11.193.88	0x5272	0x3272	1	0	0	1	1	8193	b.root-servers.net	1	13
187.11.204.49	0xc4cb	0xa4cb	1	0	0	1	1	8193	img.cancaonova.com	1	13
187.33.176.6	0x086e	0xe86d	1	0	0	1	1	8193	www.moovielive.com	1	13
187.34.197.173	0xad96	0x8d96	1	0	0	0	1	8193	dns7.inatserve.net	1	13

Observations

- Multi-bit checksum errors surprisingly common
- Single-bit checksum errors probably not due to transmission errors
- Hundreds of clients send duplicated queries, with bad checksums on the duplicates.

~~Questions?~~

Answers?