DDoS and Collateral Damage risks: are TLDs oversharing DNS infrastructure? (*ongoing work*)

Giovane C. M. Moura, Moritz Müller,
Marco Davids, Cristian Hesselman

IEPG – IETF 100 – Singapore

2017-11-12

SIDN LABS

Introduction

- The DNS comprises one of the core services of the Internet
- Resilience due to good engineering: layers and layers of redundancy
  - Multiple NS records
  - IP anycast
  - Load Balancers

Introduction

- DDoS are becoming cheaper, bigger and more frequent:
  - Dyn DDoS peaked at 1.2TB/s (Mirai Botnet, October 2016) [1]
  - Root DNS DDoS 35Gb/s (Nov 2015) [2]
- In both cases, we have seen **collateral damage:**
  - .nl anycast sites close to Root letters also suffered during DDoS
  - multiple Dyn clients (Spotify, Netflix, etc.) were partially unreachable

SIDN LABS

Introduction

- Collateral damage only happens because *parts of infrastructure* are shared.
  - Name servers/ IP addresses/ pipes/ autonomous systems / datacenters
- Sharing, per se, it is not a problem: big DNS providers are more likely to have more capacity that small operators
- But they have a larger attack surface….

SIDN LABS

Dyn DDoS attack October 2016: 1.2TB/s (Mirai Botnet)

Hackers Used New Weapons to Disrupt Major Websites Across U.S.

By **NICOLE PERLROTH** OCT. 21, 2016

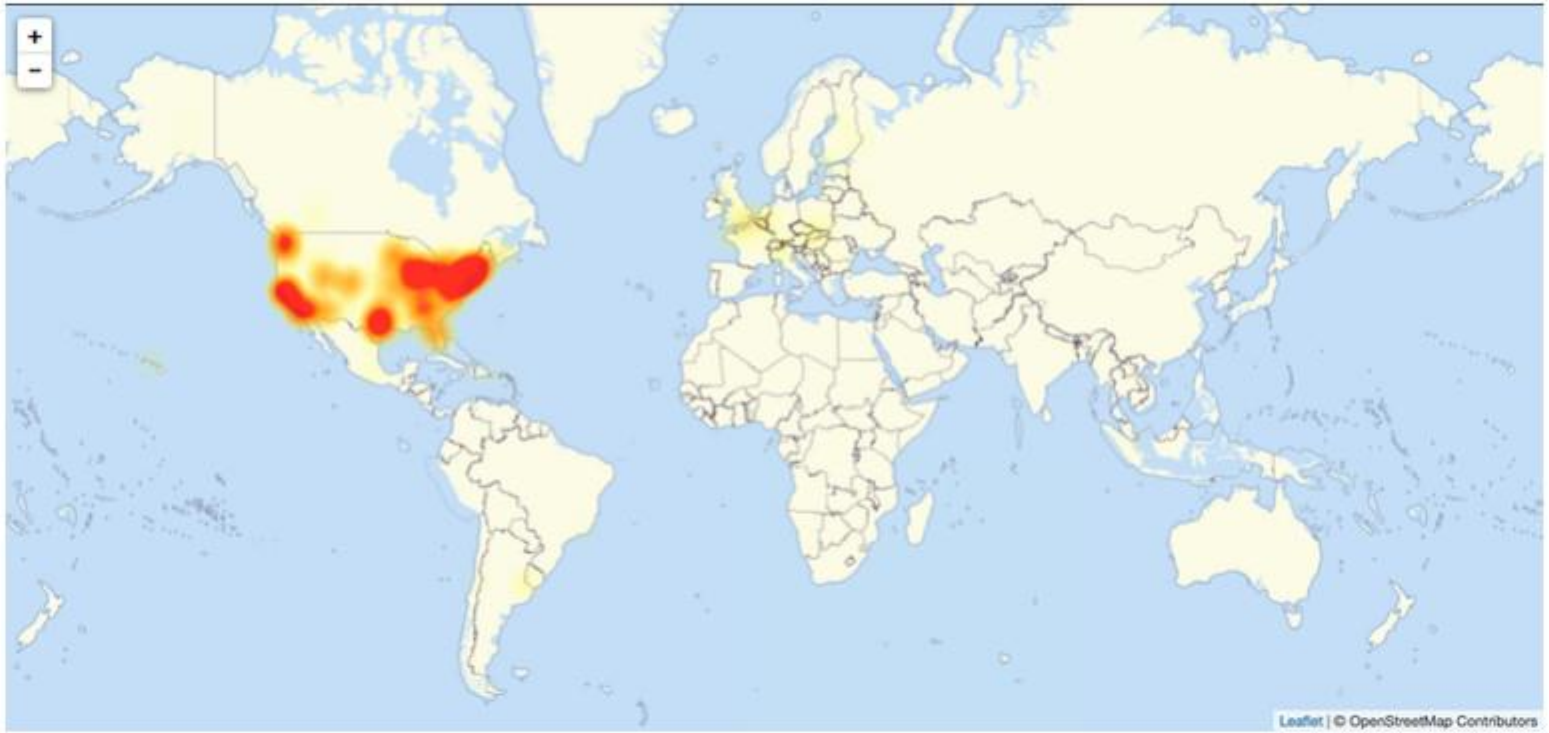

What do to?

- Research question

  How much sharing is in the Root zone?

- Approach: Measurements
- Motivation:
  - Operators: I want to know who my DNS provider shares infra
- Two parts:
  - Root Zone as a whole (Shared IPv4, Ases, IPv6, /24)
  - Individual TLDs (Number of ASes, NSes)

Part 1: Root Zone as a whole

- Look at all NSes, A Records, and how many TLDs share them

Root Zone Analysis: 2014 and 2017

	All		Orig. TLDs		ccTLDs		gTLDs	
Metric	*2014*	*2017*	*2014*	*2017*	*2014*	*2017*	*2014*	*2017*
TLDs	613	1535	7	7	249	247	357	1281
NSes	1569	4326	38	40	1045	1047	608	3396
A Rec	1476	3805	32	34	1009	1003	567	2945
AAAA Recs	900	3364	14	28	510	626	459	2840
ASes(IPv4)	489	511	26	30	445	444	155	214
ASes (IPv6)	NA	241	NA	11	NA	202	NA	127

Table 1. Root Zone: 20140601 and 20170627

- **Original TLDs**: .com,.net, org, mil, gov, edu, int.
- **ccTLDs**: .nl, .ca, and .fr, etc.
- **gTLDs**: all the others, such as .amsterdam and .io.

SIDN LABS

Root Zone Analysis: 2014 and 2017

	All		Orig. TLDs		ccTLDs		gTLDs	
Metric	*2014*	*2017*	*2014*	*2017*	*2014*	*2017*	*2014*	*2017*
TLDs	613	1535	7	7	249	247	357	1281
NSes	1569	4326	38	40	1045	1047	608	3396
A Rec	1476	3805	32	34	1009	1003	567	2945
AAAA Recs	900	3364	14	28	510	626	459	2840
ASes(IPv4)	489	511	26	30	445	444	155	214
ASes (IPv6)	NA	241	NA	11	NA	202	NA	127

Table 1. Root Zone: 20140601 and 20170627

- new gTLDs delegations
- 4.06 IPv4 per ccTLD, **2.29 for gTLDs**
- 1.79 ASv4/ccTLD, **0.16 for gTLDs**

SIDN LABS

Shared NSes

Shared NSes : getting more concentrated due to new gTLDs

(but 612 gTLDs < 10 domains[1])

1: https://ntldstats.com/

2014

NS	Unique TLDs	Examples
demand.gamma.aridns.net.au	165	newGTLDs: zone, works, today
demand.delta.aridns.net.au	165	newGTLDs: zone, works, today
demand.beta.aridns.net.au	165	newGTLDs: zone, works, today
demand.alpha.aridns.net.au	165	newGTLDs: zone, works, today
sns-pb.isc.org	42	ccTLDs: nl, nr, cl, cat
rip.psg.com	19	ccTLDs: sa, lb,uy
d0.cctld.afilias-nst.org	15	ccTLDs: vc, sc,bz
c0.cctld.afilias-nst.info	15	ccTLDs: vc, sc,bz
b0.cctld.afilias-nst.org	15	ccTLDs: vc, sc,bz
a0.cctld.afilias-nst.info	15	ccTLDs: vc, sc,bz

2017

NS	Unique TLDs	Examples
demand.gamma.aridns.net.au	238	newGTLDs: zone, works, today
demand.delta.aridns.net.au	238	newGTLDs: zone, works, today
demand.beta.aridns.net.au	238	newGTLDs: zone, works, today
demand.alpha.aridns.net.au	238	newGTLDs: zone, works, today
ac4.nstld.com	160	newGTLDs:norton, samsclub
ac3.nstld.com	160	newGTLDs:norton, samsclub
ac2.nstld.com	160	newGTLDs:norton, samsclub
ac1.nstld.com	160	newGTLDs:norton, samsclub
l.gmoregistry.net	49	newgTLDs: yokohama, toyota
k.gmoregistry.net	49	newgTLDs: yokohama, toyota

Table 2. Top 10 Shared NS names and number of TLDs

Shared NSes: considering TLDs zone size

2017

NS	Unique TLDs	MeanZone	Median	Total
demand.gamma.aridns.net.au	238	13,688.14	6,542	3,203,026
demand.delta.aridns.net.au	238	13,688.14	6,542	3,203,026
demand.beta.aridns.net.au	238	13,688.14	6,542	3,203,026
demand.alpha.aridns.net.au	238	13,688.14	6,542	3,203,26
ac4.nstld.com	160	435.82	1	62,323
ac3.nstld.com	160	435.82	1	62,323
ac2.nstld.com	160	435.82	1	62,323
ac1.nstld.com	160	435.82	1	62,323
l.gmoregistry.net	49	8,826.75	5	432,511
k.gmoregistry.net	49	8,826.75	5	432,511

Table 3. Top 10 Shared NS names and number of TLDs and Zone Sizes (zone sizes obtained on 20171025 from gtldstats.net, only for gTLDs).

Only new gTLD zone sizes: gtldstats.net

SIDN LABS

Shared IPv4 (A Records)

- Number of zones (TLDs) on each IPv4 A Record
- Growing concentration

2014		2017	
IPv4	# TLDs	IPv4	# TLDs
37.209.198.7	165	37.209.198.7	238
37.209.196.7	165	37.209.196.7	238
37.209.194.7	165	37.209.194.7	238
37.209.192.7	165	37.209.192.7	238
192.5.4.1	45	192.42.176.30	160
147.28.0.39	21	192.42.175.30	160
72.52.71.3	16	192.42.174.30	160
63.243.194.3	16	192.42.173.30	160
38.103.2.3	16	37.209.198.9	49
199.254.62.1	15	37.209.198.4	49

Table 3. Top 10 Shared IPv4 (A Records) and TLDs

SIDN LABS

Shared /24 of A Records

- Obtained IPv4 from each NS
- Aggregate it into /24
- Count unique TLDs
- Imagine a prefix hijack.

2014		2017	
/24	**# TLDs**	**/24**	**# TLDs**
37.209.194.0	176	37.209.194.0	362
37.209.192.0	176	37.209.192.0	362
37.209.198.0	174	37.209.198.0	360
37.209.196.0	174	37.209.196.0	360
193.0.9.0	71	156.154.159.0	179
204.61.21.0	64	156.154.158.0	179
192.5.4.0	51	156.154.157.0	179
194.0.1.0	41	156.154.156.0	179
194.146.106.0	33	156.154.145.0	179
72.52.71.0	25	156.154.144.0	179

Table 3. Top 10 Shared /24 prefixes and of TLDs

Shared ASes of A Records (2017)

- Up to 363 TLDs/AS
- A TLD may use multiple Ases

IPv4		IPv6	
AS	**# TLDs**	**AS**	**# TLDs**
134399	363	134390	363
134395	363	134399	362
134391	363	42	356
134390	363	12008	263
134386	363	12041	226
42	362	19911	223
18210	360	36625	163
134396	360	36616	162
12008	265	36617	161
19911	262	15135	77

Table 4. Top 10 Shared ASes and of TLDs – 2017

SIDN LABS

Part 2: Individual TLDs

- How individual TLDs look like?

Looking at individual TLDs

TLD	Country/Org	NSes	/24	ASes(IPv4)
.kp	North Korea	2	1	1
.pf	French Polynesia	2	1	1
.bb	Barbados	2	2	1
.gf	French Guiana	2	2	1
.sr	Suriname	2	2	1
.dj	Djibouti	2	2	2
.ax	Åland Islands	3	2	1
.bh	Bahrain	4	1	1
.mil	US DoD	6	6	1

Table 6. TLDs with 1 IPv4 AS

Looking at individual TLDs

- .dog, .money: Multiple Origin AS [3,4], and any TLD above line (Anycast)
- Most 4 Nses with < 4 Ases
  6 NSes also popular

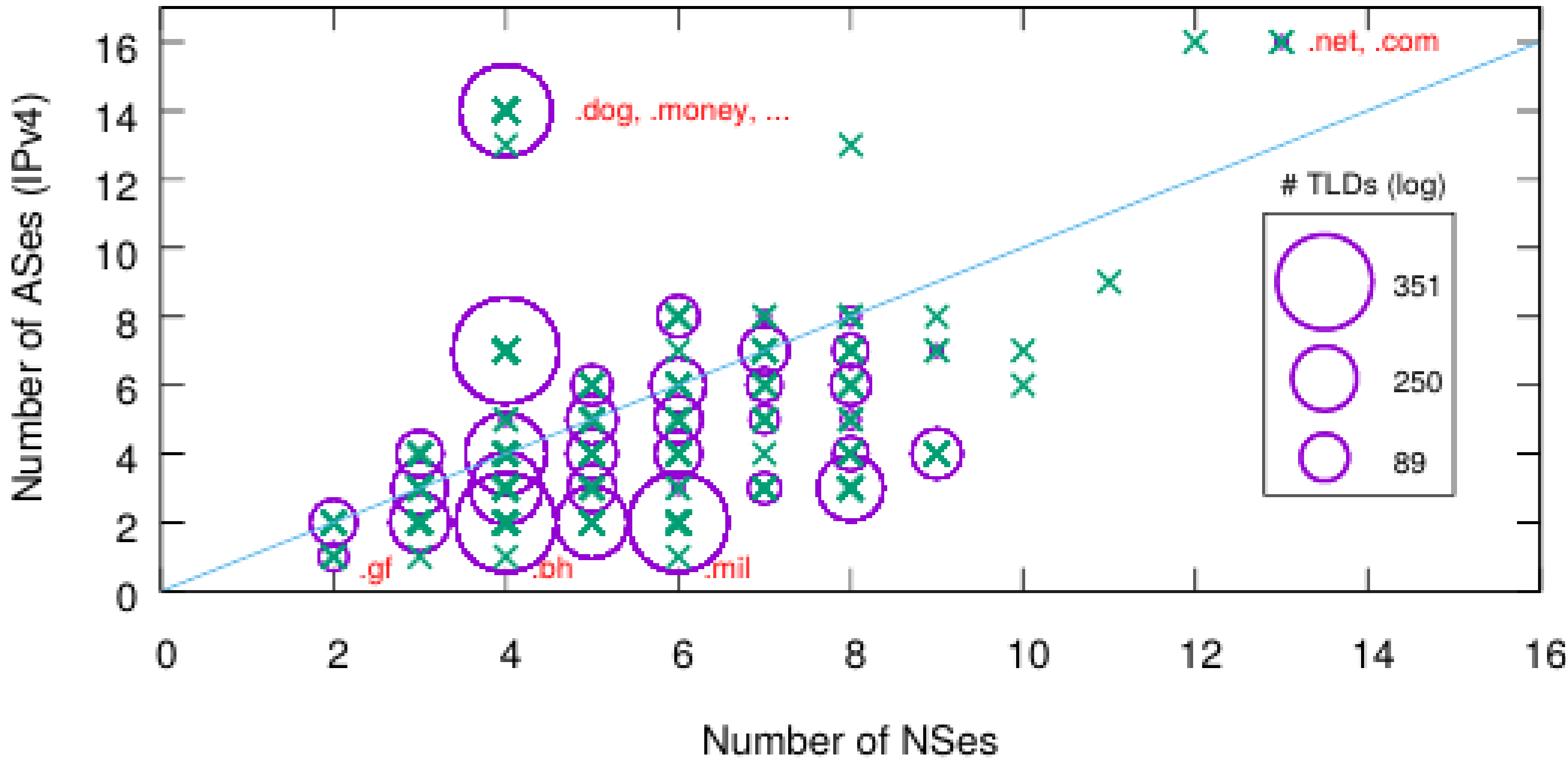

Fig. 4. TLDs: number of NSes vs number of ASes – IPv4

Looking at individual TLDs

- Below line: TLDs with multiple NSes on a single AS (majority)

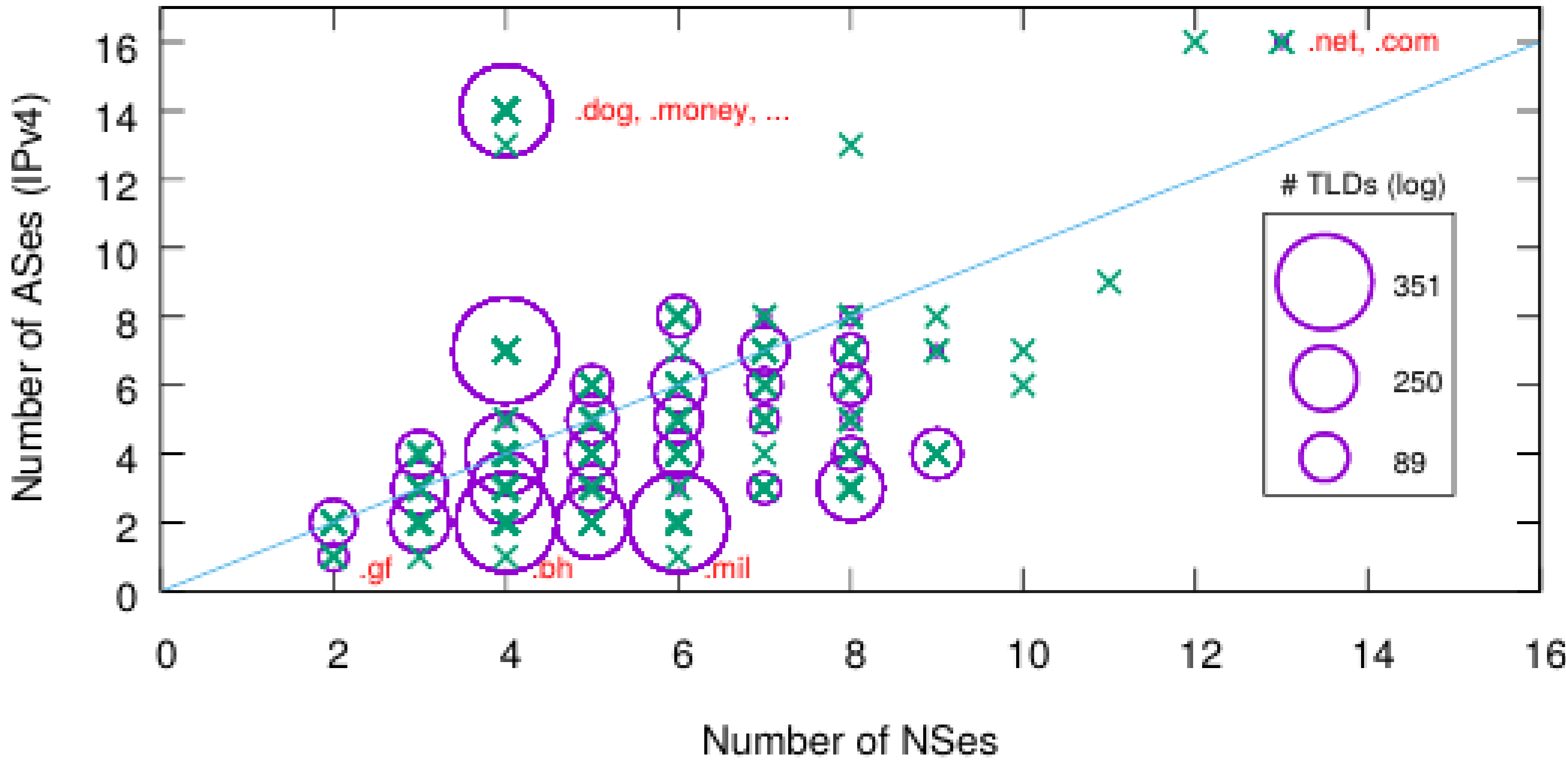

Fig. 4. TLDs: number of NSes vs number of ASes – IPv4

So far:

- ccTLDs have far less sharing than gTLDs
  - Market is very different, 50% new gTLDs (612) have < 10 domains
- 9 TLDs using 1 AS only
- Some /24 having 360+ zones
  - IPv6 subnetting also important, not that straightforward

SIDN LABS

Discussion

- Is this something we should be worried about ?
- What can we possibly do about it?
  - **Ops: mixed 3rd party and in-home anycast services**
- There's many other shared infra we can't possibly measure (pipes, other non DNS services)
- ***Do we need a definition of what is "oversharing"?***

SIDN LABS

Discussion

- We'd like feedback/suggestions on how to move forward

References

1. Perlroth, N.: Hackers Used New Weapons to Disrupt Major Web-sites Across U.S. (2016) . https://www.nytimes.com/2016/10/22/business/internet-problems-attack.html.
2. Moura, G.C.M., de O. Schmidt, R., Heidemann, J., de Vries, W.B., Müller, M.,Wei, L., Hesselman, C.: **Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event**. In: Proceedings of the 2016 ACM Conference on Internet Measurement Conference. (October 2016) 255–270
3. Zhao, X., Pei, D., Wang, L., Massey, D., Mankin, A., Wu, S.F., Zhang, L.: **Analysis of BGP multiple origin AS (MOAS) conflicts**. In: Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement, ACM (2001) 31–35
4. Jacquemart, Quentin, Guillaume Urvoy-Keller, and Ernst Biersack. "**A longitudinal study of BGP MOAS prefixes**." International Workshop on Traffic Monitoring and Analysis. Springer, Berlin, Heidelberg, 2014.

SIDN LABS

Questions?