NCAP Discussion Group | 3 February 2021 | 19:00 UTC

Agenda:

1. Welcome and roll call
2. Update to SOI
3. Update on Study 2
4. .LAN Case
   Study: https://docs.google.com/presentation/d/1LeiVG3t94kqCNPPvMklvQtTJx0iKEMKqvUDPggnnWEY/edit#slide=id.gb590d7797c_0_19
5. .LOCAL Case
   Study: https://docs.google.com/presentation/d/1_luvm2MtGNu8sdj1NT8GVHfRf1tkajIF_RJ1o_r4Ww8/edit#slide=id.gb5b9e08a3d_0_25
6. Next analysis measurements?
7. AOB

Table of Contents

Conflict of Interest discussion: Warren's title has changed from Standards Technical Program Manager to Director of Internet Standards,

Study 2 update: package for Board is complete.

.LOCAL

Slide 1: Daily Query Volume

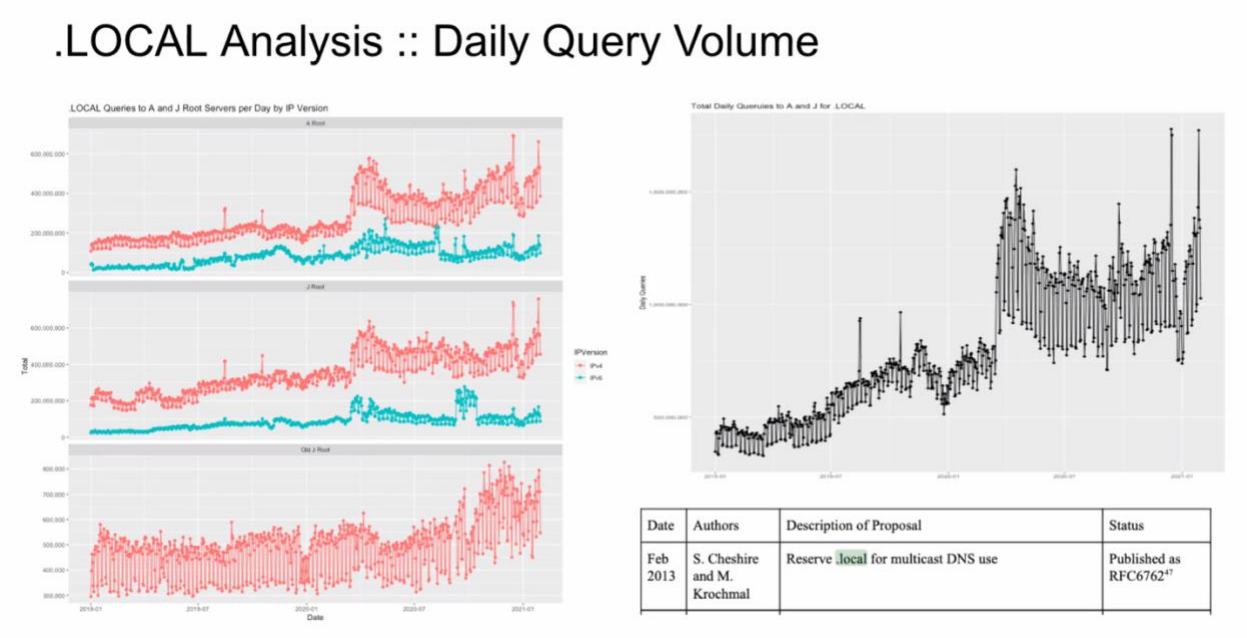

Traffic volumes over 1.5 billion queries per day. Covid issue with transient devices used at home at non-standard DNS environment
SAC113: highlighted that .local was tagged for multicast DNS use

Slide 2: Qtype Distribution

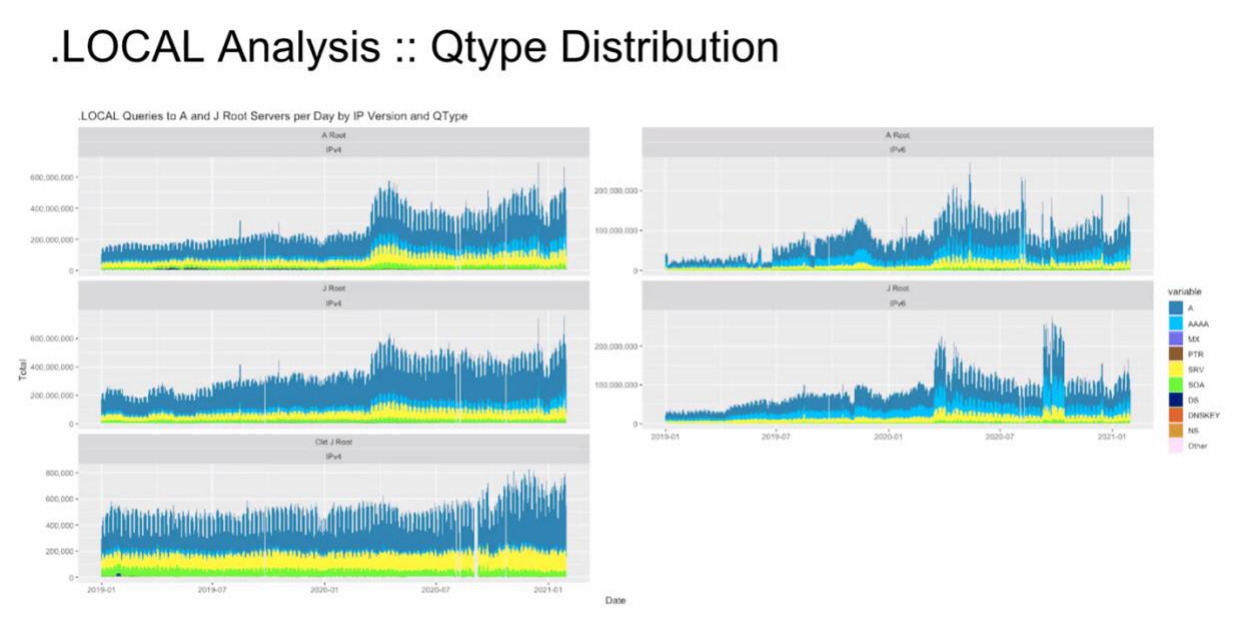

See soa records, different than other strings

Slide 3: Unique Daily Source IPs

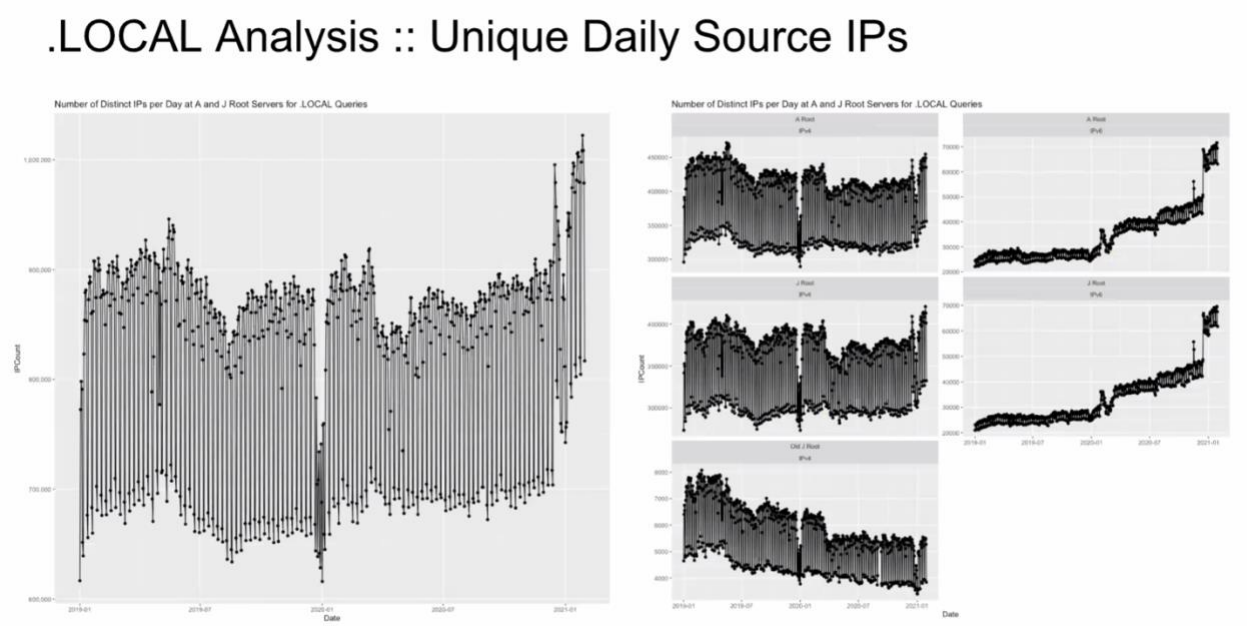

Increased traffic without increase in # of sources (unusual). # of sources just increased recently though.

Rod: doesn't match pandemic pattern. Are source Ips due to broad use of public resolvers, or a large IP infastructure

In future might want to do analysis of known large public ISP recursivers an differentiating their traffic from AS numbers, for example.

STEVE: what might account for this phenomena : An IP address used even once and then there's, an increase that will show up as no change the range of IP addresses, whereas before you had IP addresses that had never been used, presumably that are then included, so If there were a phenomenon which there was kind of a background of all of the IP addresses occurred a little tiny bit. And then you had a massive increase you wouldn't be able to see that phenomena exactly in the way that this is being analyzed so I was just another thought trying to explore why this is behaving as it is.

Slide 4: GEOGRAPHICAL Distribution

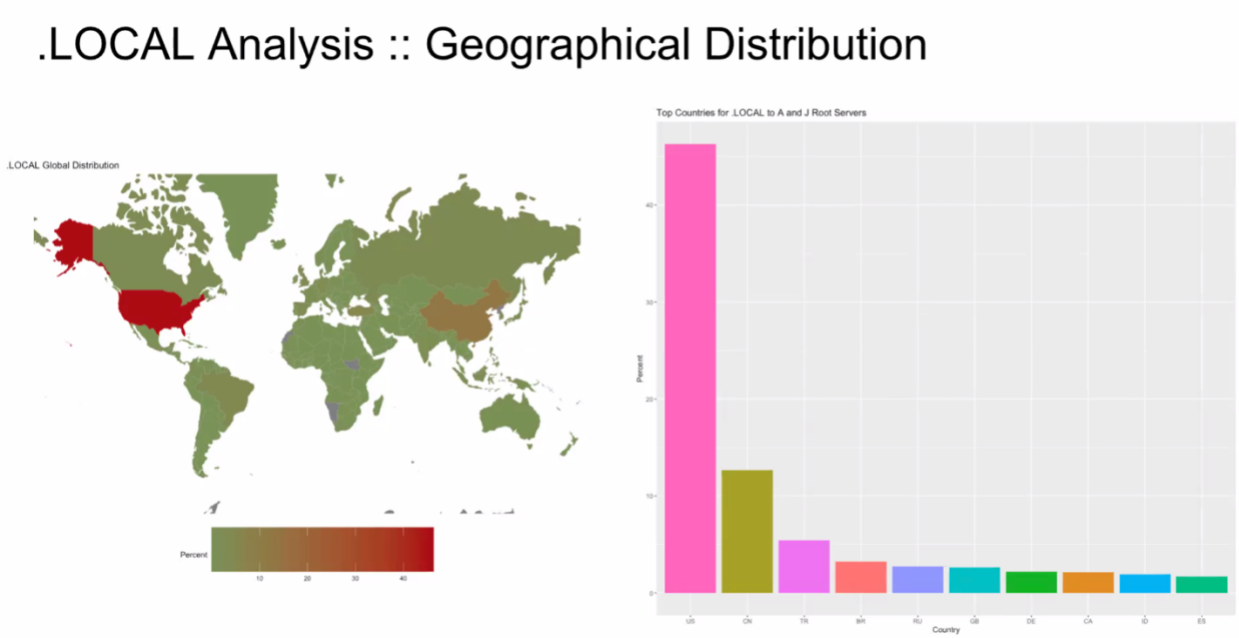

US mostly

Slide 5: ASN Distribution

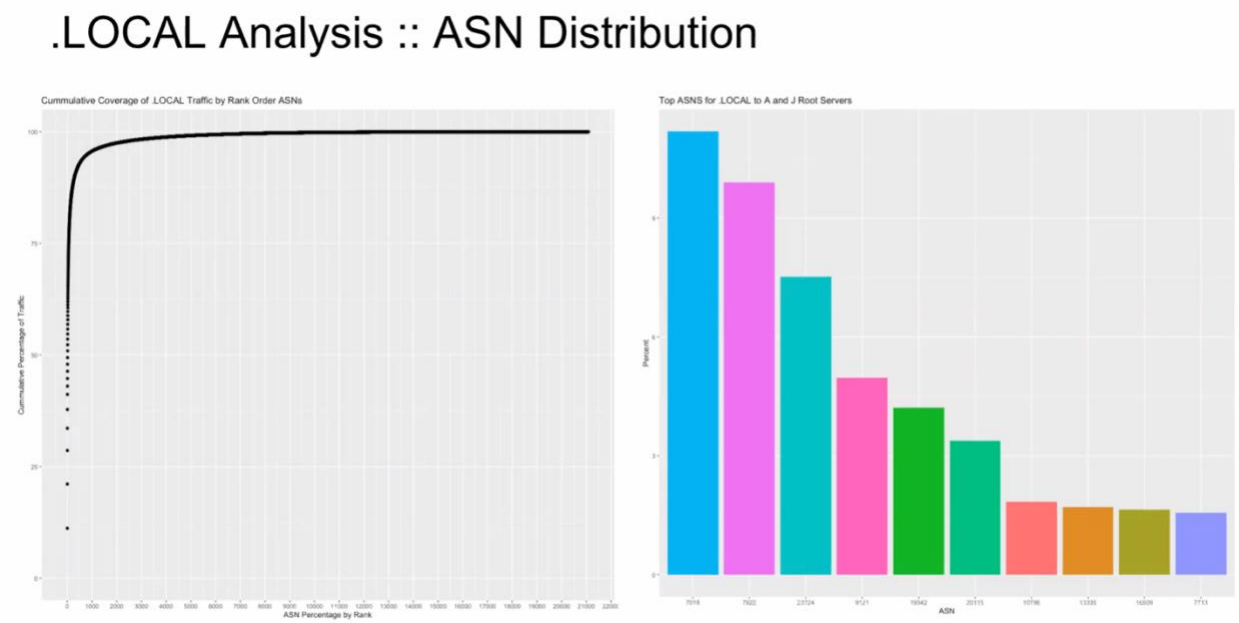

Scale of X axis: other studies dropped at 900 ASNs, but .LOCAL is over 22000 different ASNs- wider spread than other case studies but 50% of traffic comes from top 10 (mostly US/China ISPs)

Slide 6: Root ASN Overlap and IP Growth

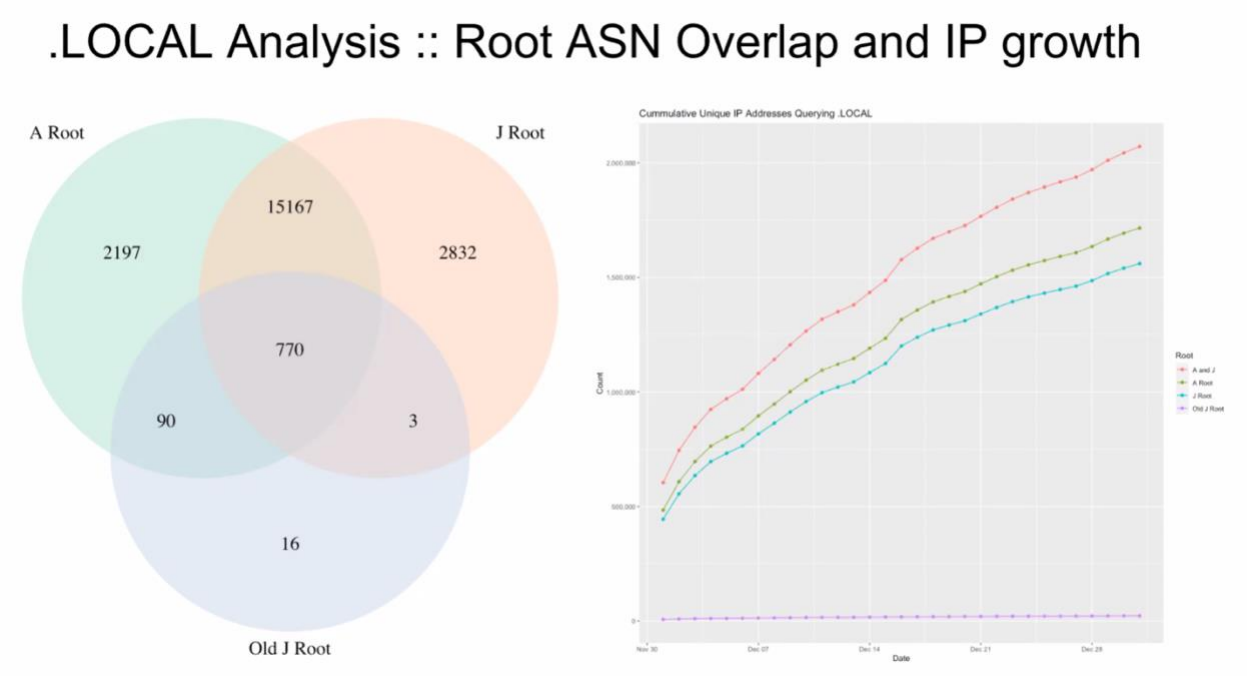

J root always has much larger collection of catchment than A root

Slide 7: Label Analysis

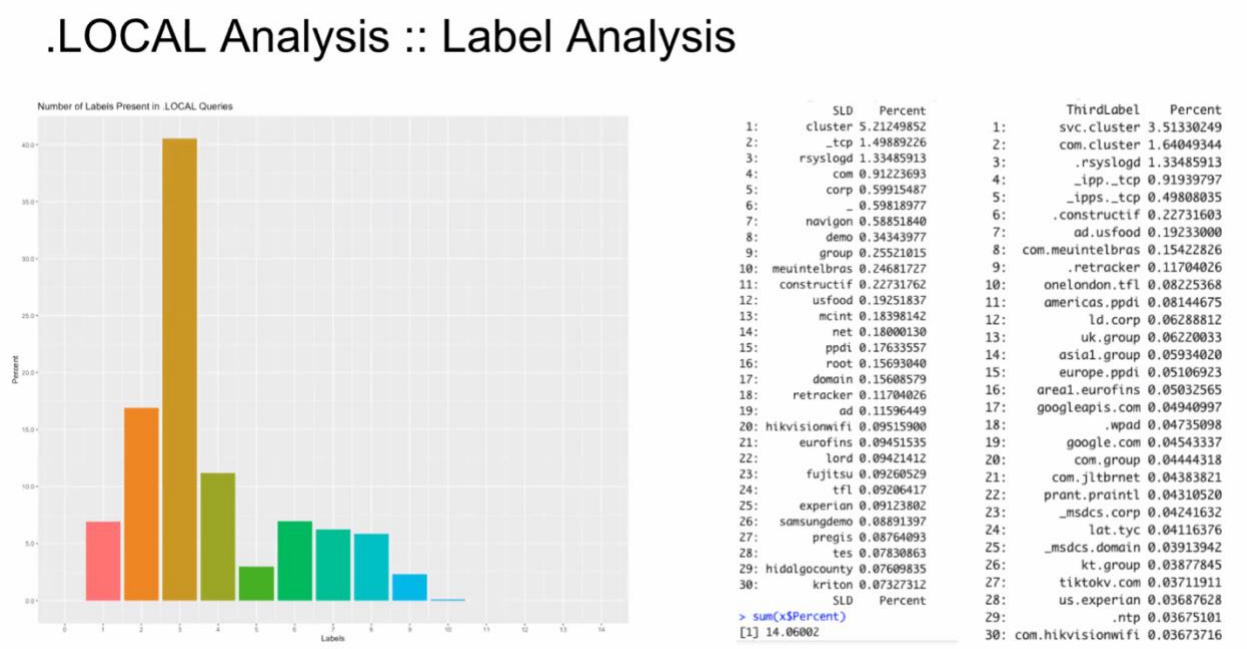

Left: 40% have 3 labels, only 7% had 1 label
(.mail had 50% with 1 label) int
Like .corp – some are anchored under delegated tlds, so you wonder if this is by-product of suffix search list appendage

Right: associated with DNS service discovery protocols

Slide 8: Data Attributes

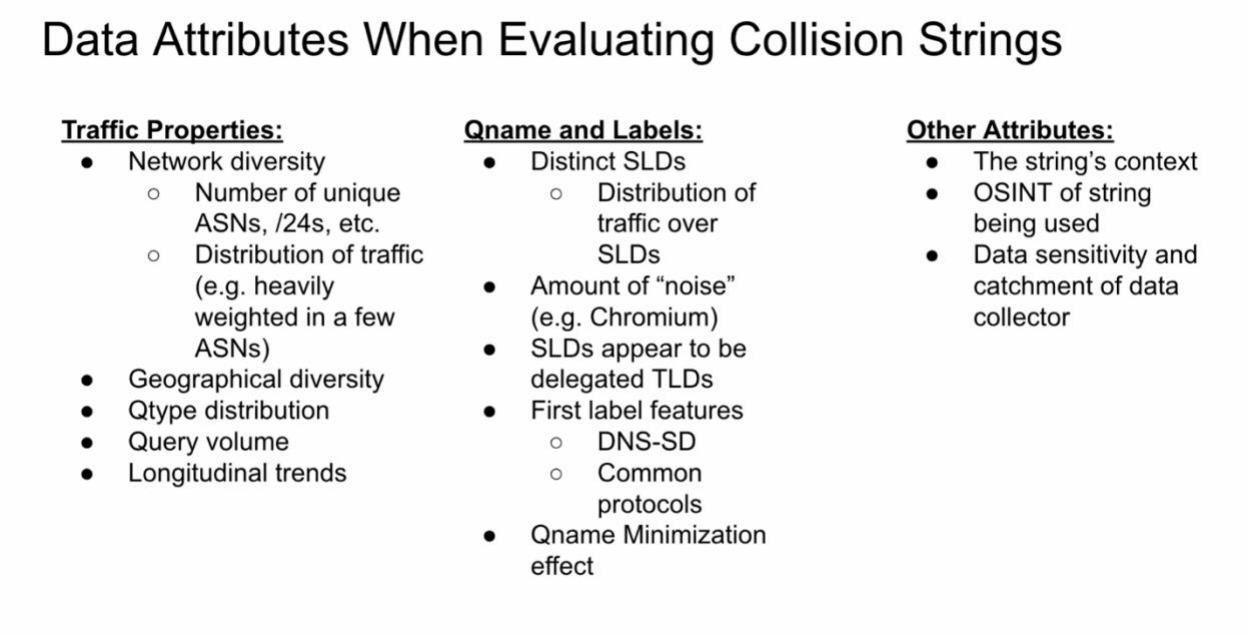

.LAN

Used in open wrt- open source for home routers with wide variety of software features. Uses private use .tld .lan

Slide 1: Daily Query Volume

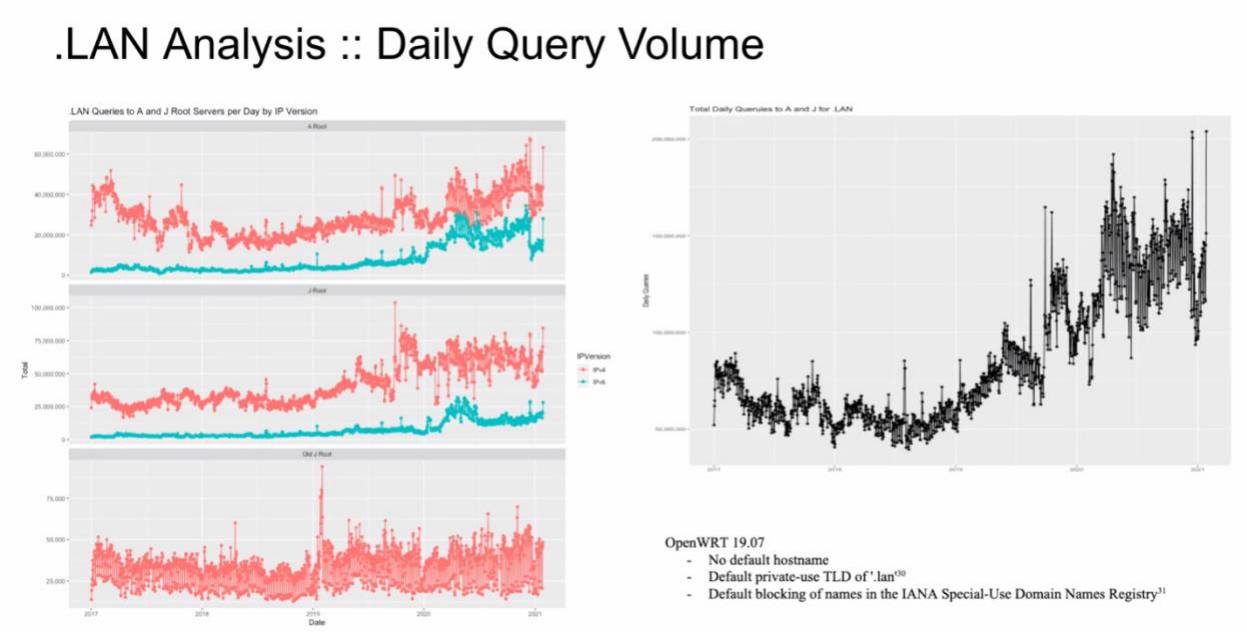

See mar 2020 shift. Decrease at end of 2020 - patch of chronium

Slide 2: Qtype Distribution

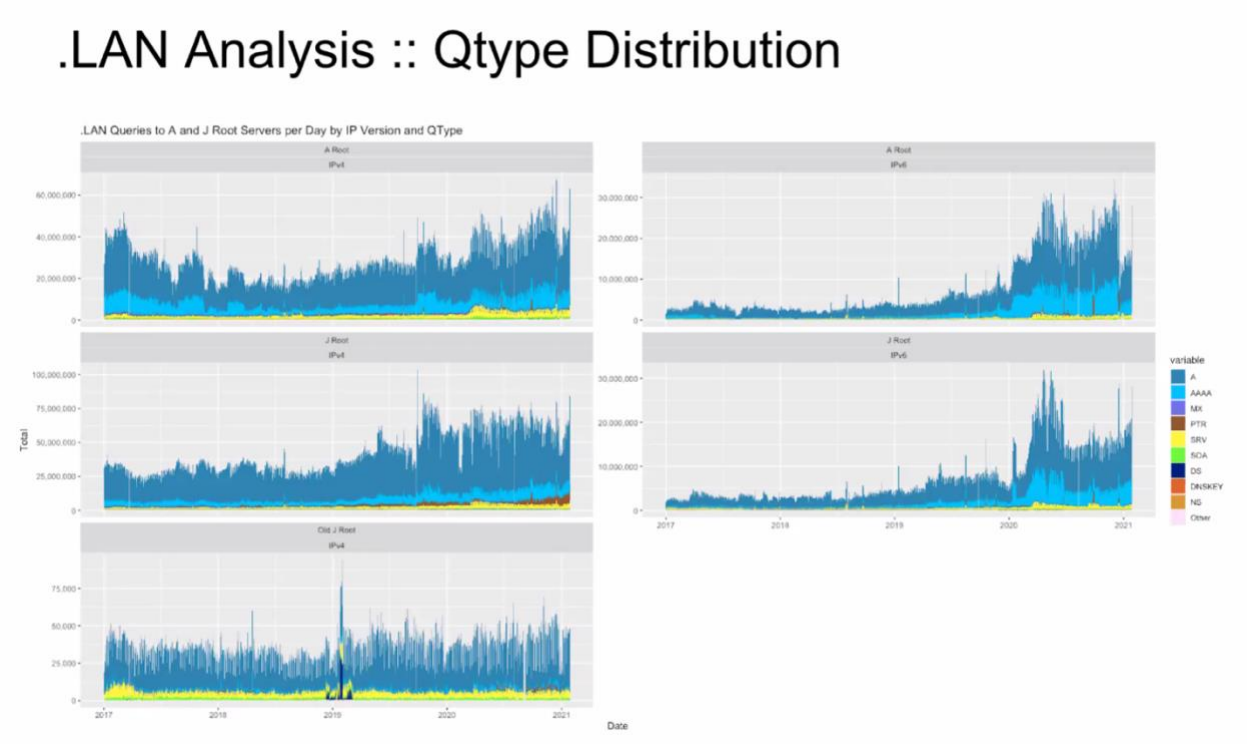

Slide 3: Unique Daily Source IPs

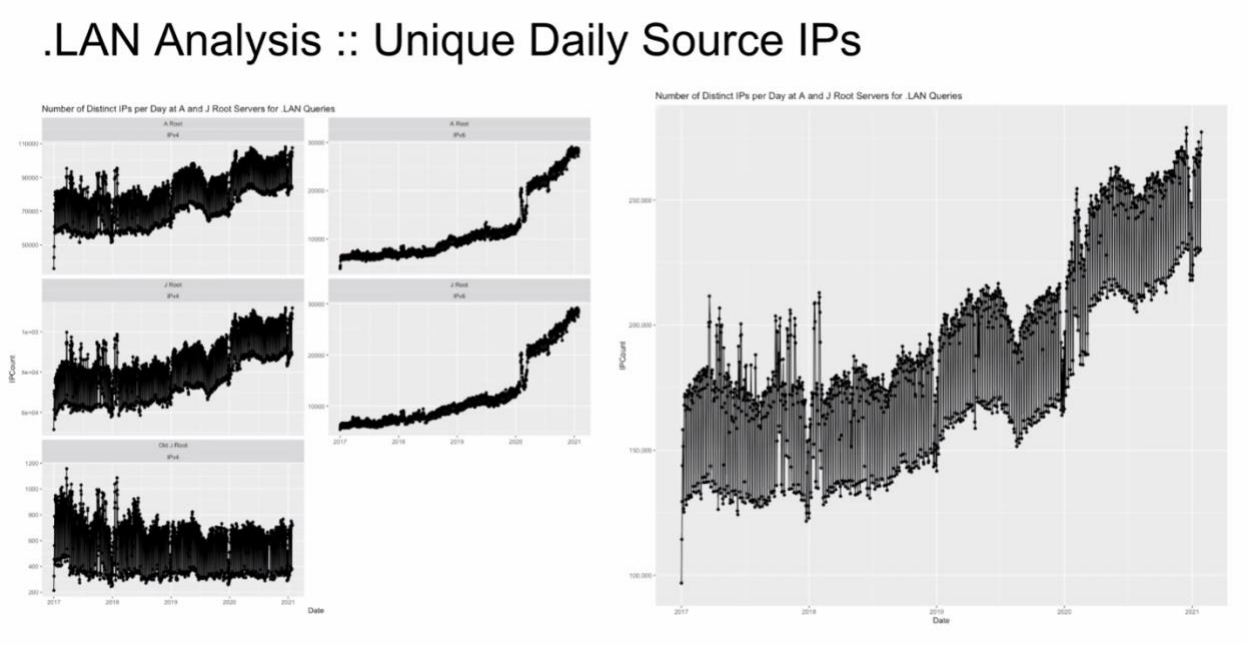

Not as pronounced covid bump

Slide 4: Geographical Distribution

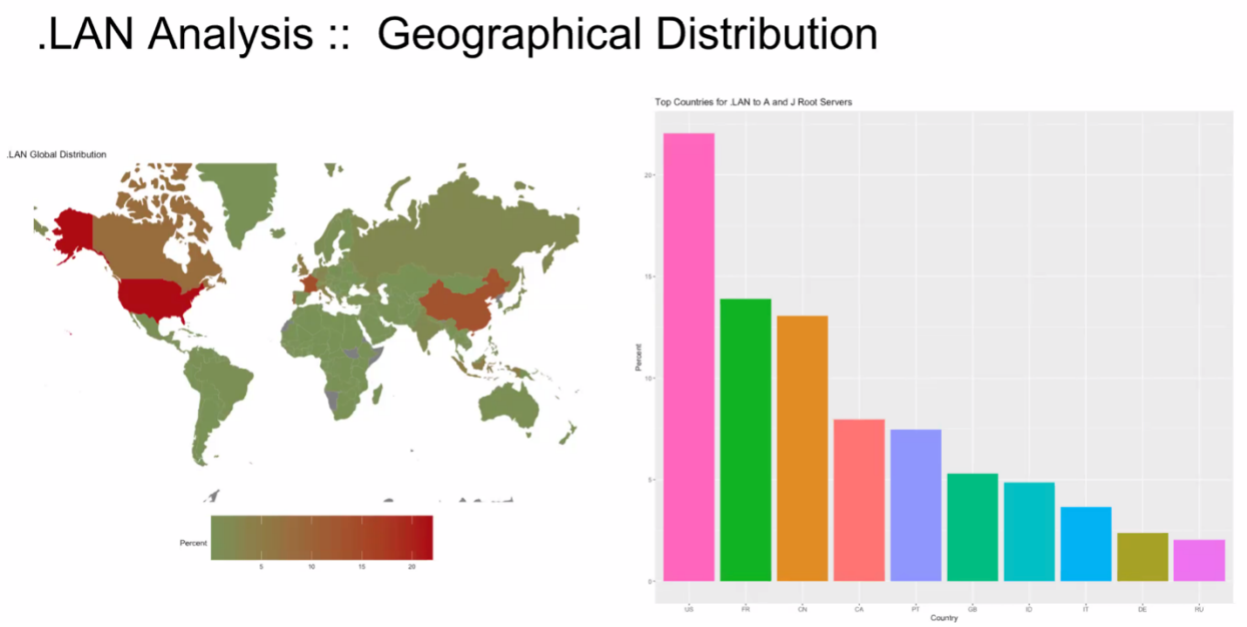

More spread out than others

Slide 5: ASN Disribution

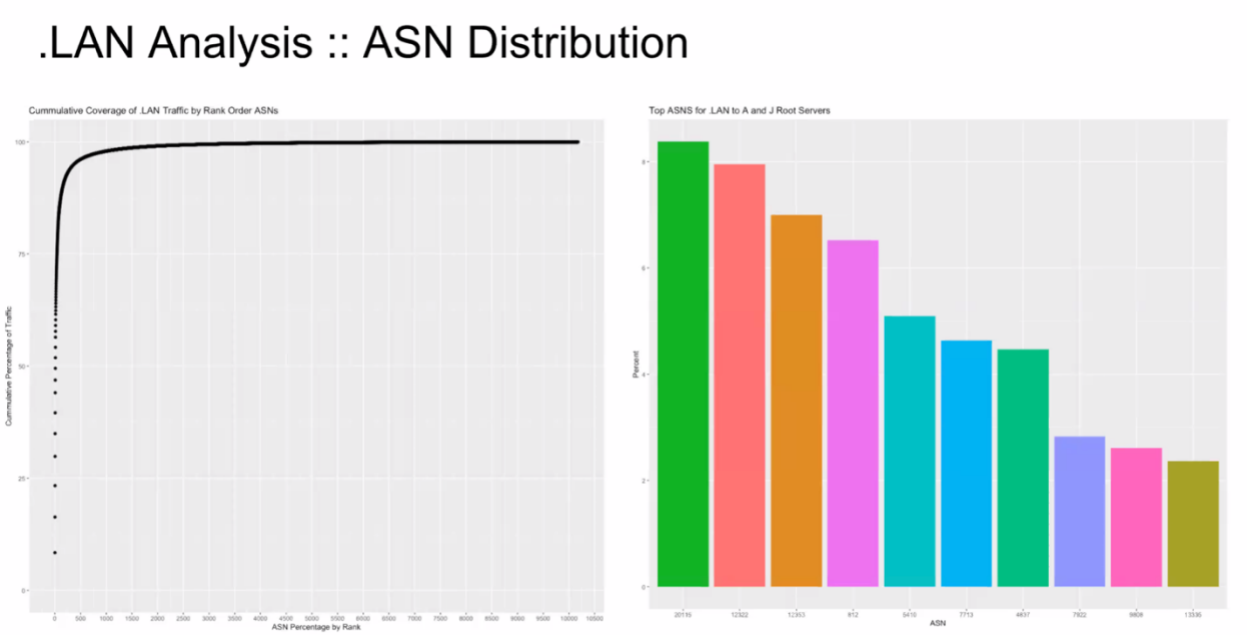

10000 different ASNs sending queires for .lan Q names, top 10 ASNs

Slide 6: Root ASN Overlap and IP Growth

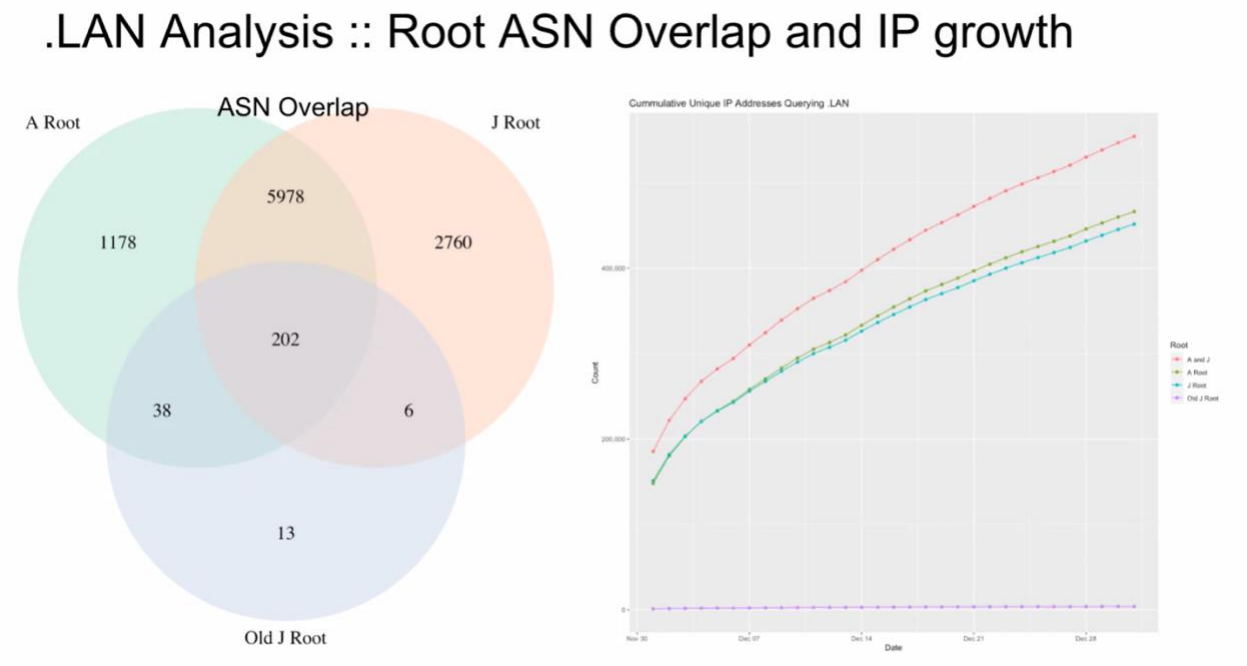

Slide 7: SLD Overlap analysis

.LAN Analysis :: SLD Overlap Analysis

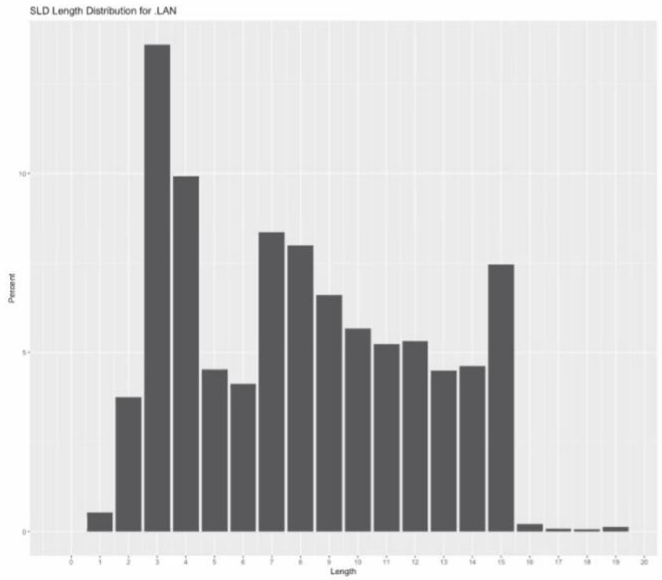

.LAN Names for 12/31/2020

- Unique Qnames: 46,117,456
- Unique SLDs: 37,440,698

Drops off at 15 chars so good portion of traffic probably chromium queries

Slide 8: SLD Analysis

.LAN Analysis :: SLD Analysis

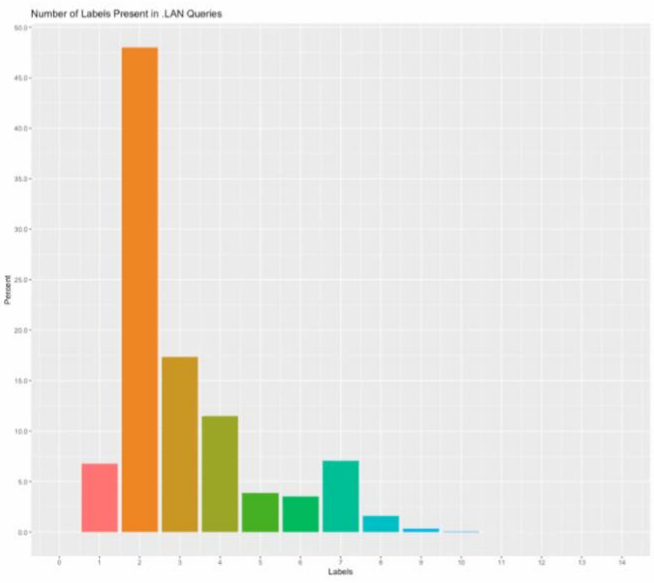

	SLD	Percent
1:	com	7.12973991
2:	arpa	5.18317203
3:	net	0.84312410
4:	wpad	0.75697182
5:	cn	0.63376531
6:	_tcp	0.46920957
7:	_	0.41307506
8:	hl	0.38232998
9:	local	0.37466281
10:	avlcorp	0.37420665
11:	org	0.35474061
12:	_udp	0.33067472
13:	basis	0.30296325
14:	vhgroup	0.29140558
15:	asm	0.22288120
16:	zodiac	0.22274337
17:	sercol	0.19536148
18:	alliance	0.19518921
19:	zebra	0.17299217
20:	isatap	0.15135092
21:	c3connect	0.13618587
22:	workgroup	0.11708910
23:	lan	0.11662245
24:	corp	0.11101439
25:	tsi	0.10818003
26:	lixil	0.10498463
27:	emergent	0.09072142
28:	galaxy	0.08977388
29:	enivest	0.08697547
30:	mobily	0.08560098

```
> sum(x$Percent)
[1] 20.04771
```

	ThirdLabel	Percent
1:	in-addr.arpa	5.15606803
2:	.wpad	0.75620481
3:	qq.com	0.47754414
4:	avl01.avlcorp	0.31909082
5:	_dns-sd._udp	0.31551641
6:	ww.hl	0.30687549
7:	google.com	0.25032900
8:	com.cn	0.24325432
9:	googleapis.com	0.24286857
10:	facebook.com	0.19768126
11:	corp.zodiac	0.15541521
12:	.isatap	0.15135092
13:	snssdk.com	0.14456837
14:	hicloud.com	0.14448897
15:	tiktokcdn.com	0.14359312
16:	ntp.org	0.13207514
17:	baidu.com	0.13203020
18:	vgeu.vhgroup	0.13126917
19:	tiktokv.com	0.12444467
20:	vaillant.vhgroup	0.12275184
21:	_msdcs.workgroup	0.10863245
22:	ksmobile.com	0.10593366
23:	fbcdn.net	0.09152888
24:	yximgs.com	0.09115886
25:	corp.alliance	0.07994349
26:	amemv.com	0.07836976
27:	_aaplcache3._tcp	0.07742597
28:	_aaplcache._tcp	0.07724320
29:	_aaplcache1._tcp	0.07712785
30:	_aaplcache2._tcp	0.07681325

ThirdLabel Percent

2 labels is average for .lan

Slide 9: Label analysis 2

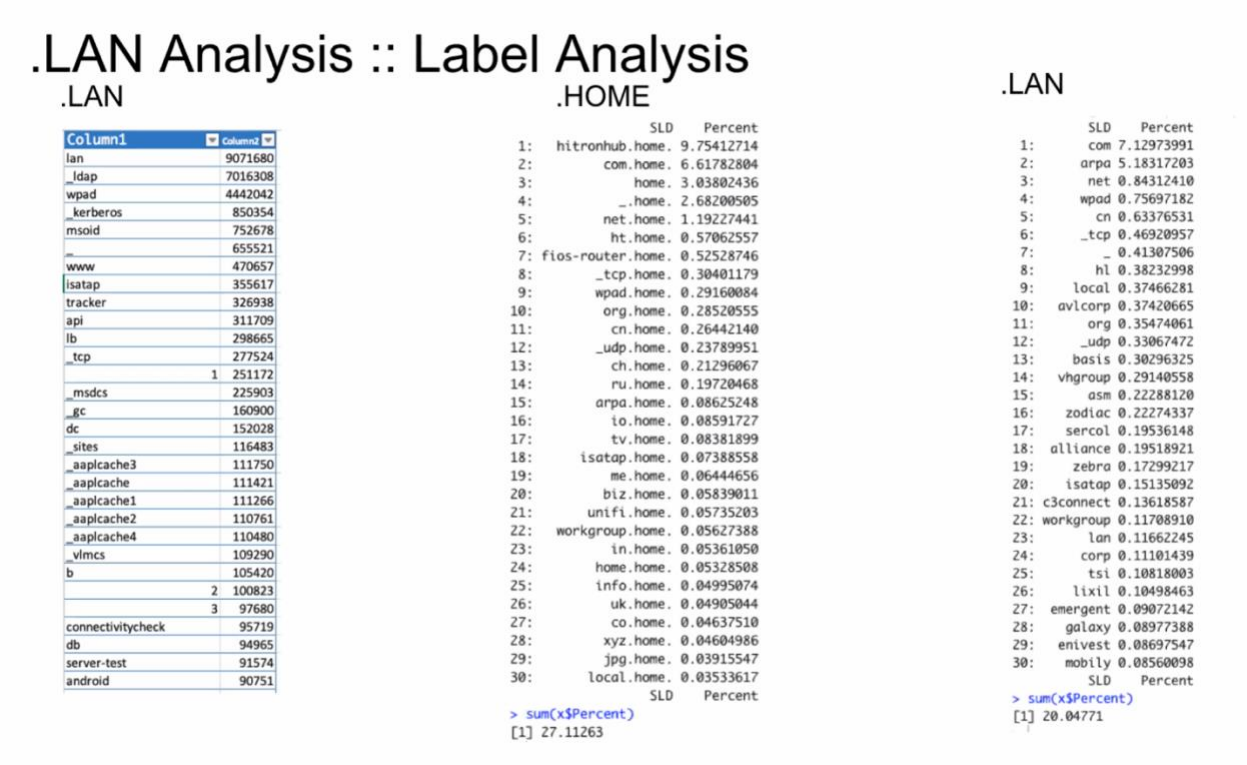

Slide 10:

Data Attributes

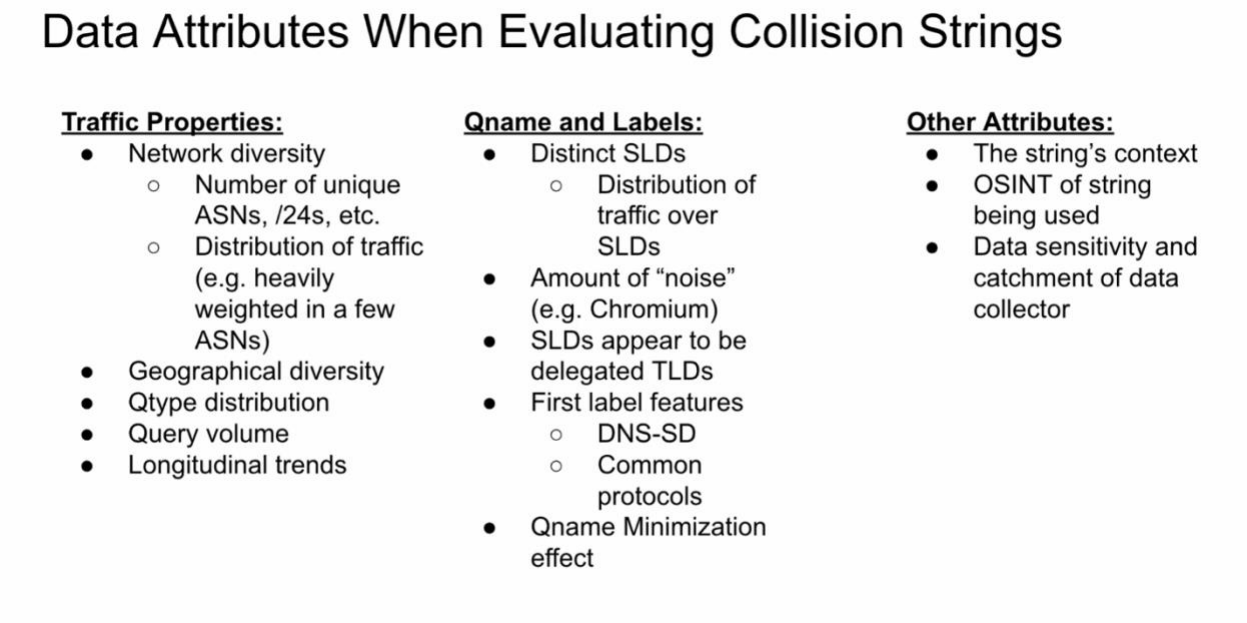

See discussion 48mins into Zoom recording