# DNS Zombie Queries

Geoff Huston & Joao Damas
APNIC Labs


Street Art: Banksy


https://xkcd.com/1361/

# Why is the DNS so interesting?

- Because everything you do on the net starts with a call to the DNS
  - If someone could see your stream of DNS queries then they could probably assemble a detailed profile of you and interests and activities
- Do we have any evidence of DNS data mining?
  - Data miners don’t disclose their sources as a rule
- How about something related:
  - Do we have any evidence of DNS stalking?

# What if…

- I gave you an absolutely unique DNS name to resolve:
  - The name never existed before now
  - The name will never be used again
  - The name includes the time when the name was created
- If I am the authoritative server for the name's zone then I should see your efforts to resolve the name
- Then I should never see the name as a resolution query ever again

  Unless you have attracted a digital stalker who performs re-queries of your DNS names!

# What if DNS query labels contained usable information?

06s-0640-uedb11641-c13-s1565985257-i0a000001-0.ap.dotnxdomain.net

# What if DNS query labels contained usable information?

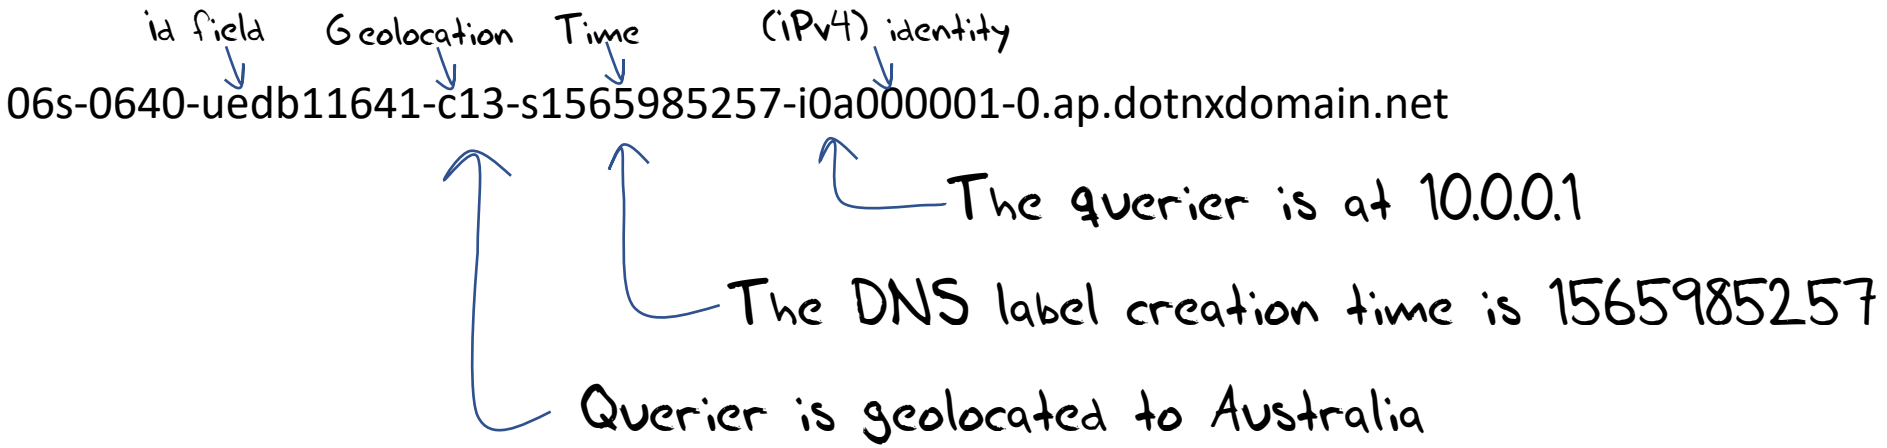


Then we could distinguish between the initial query and various re-query events!

# DNS Re-query Rate

- Over the past 12 months, a minimum of 6% - 8% of daily query totals are zombie queries, asking the same query more than 30 seconds after the initial query – and some days its as high as 70%

Daily Query Counts

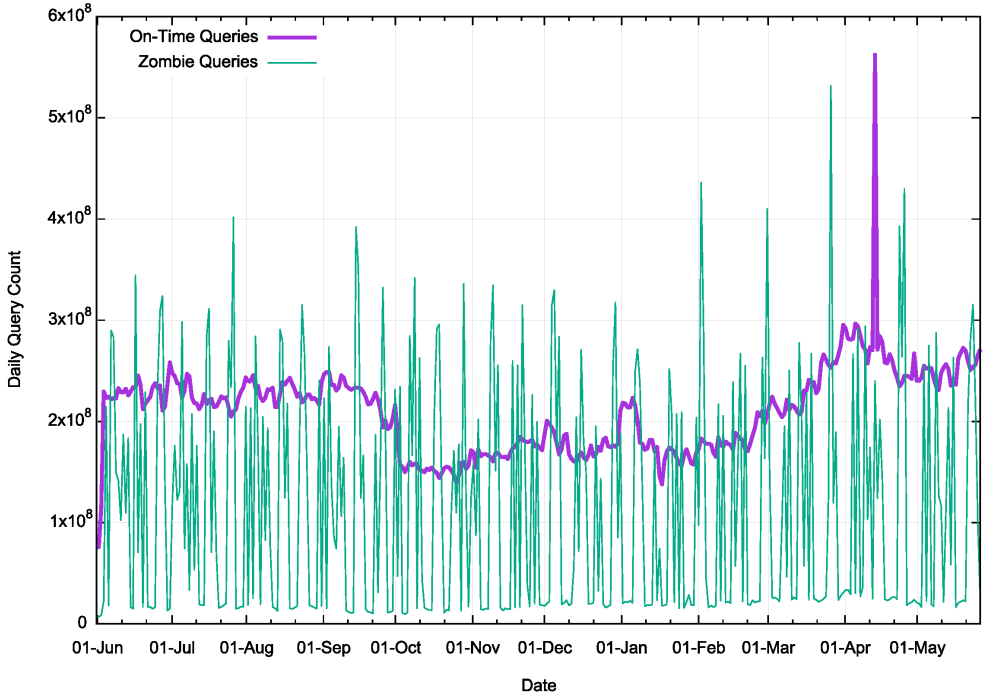


Daily Query to Zombie Ratios

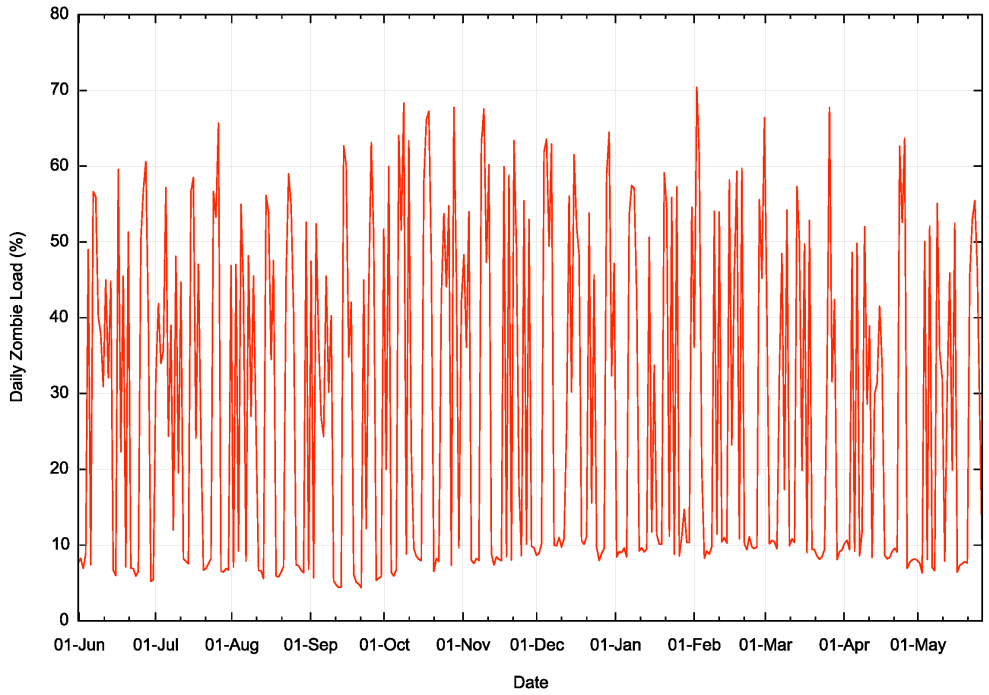

# DNS Zombies

- There are two kinds of DNS zombie behaviours
  - Rapid tracking zombies that appear to track users in real time
  - Bulk replay zombies that replay DNS queries at a very high rate in bursts

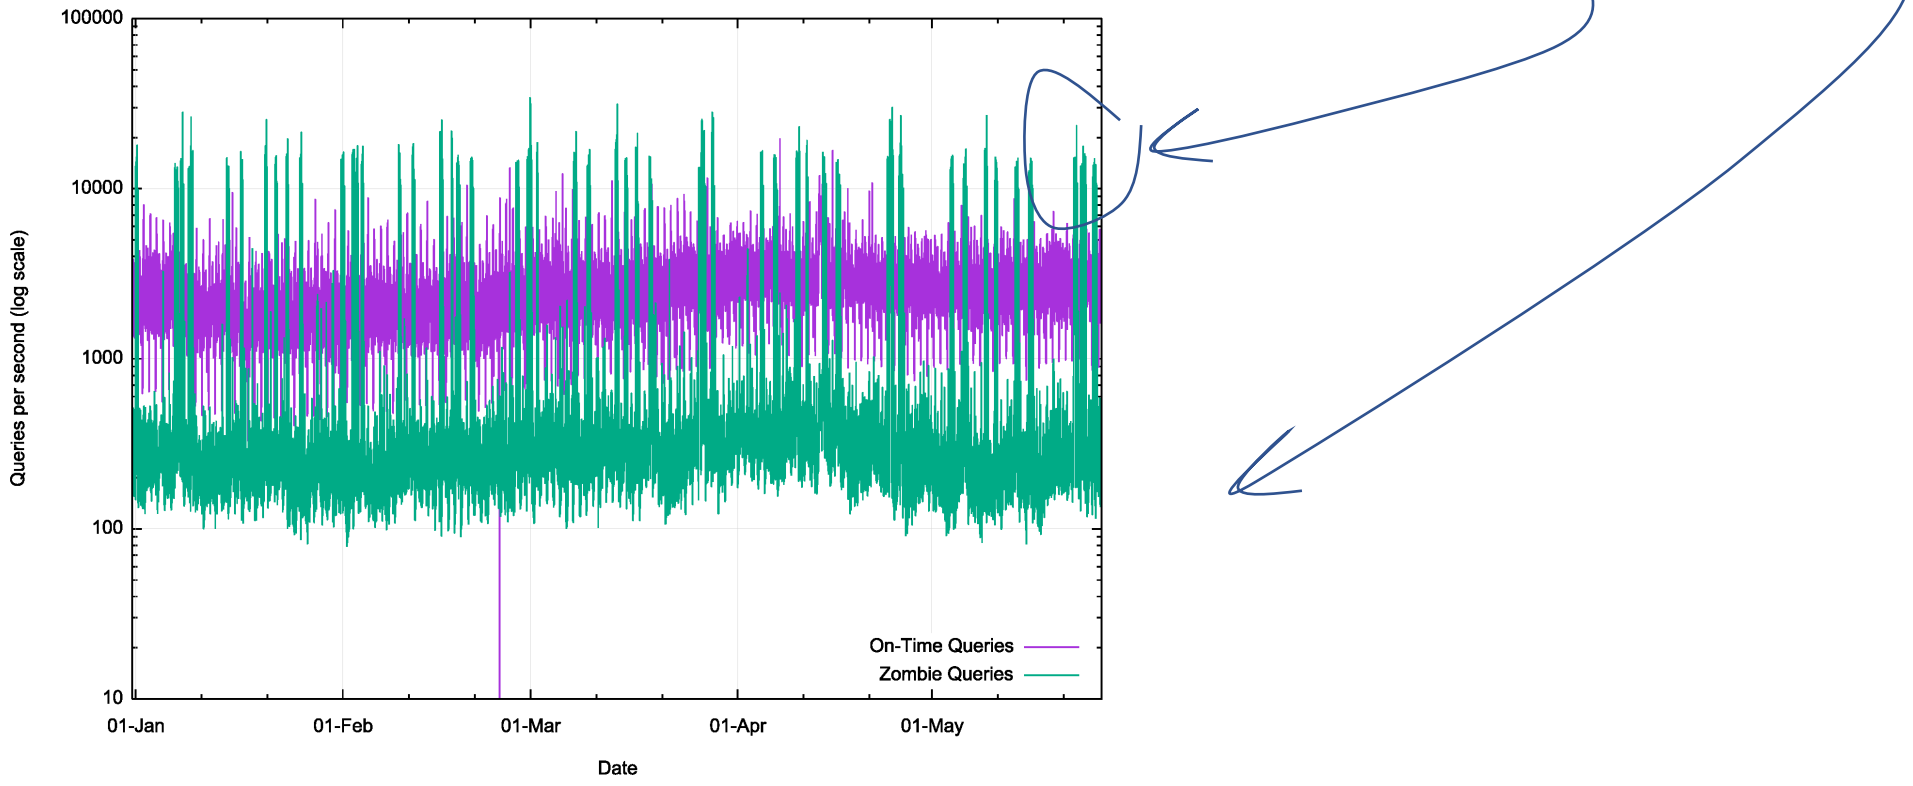

# DNS Stalking by Query Age

DNS stalking uses both really recent data and more than year-old data!

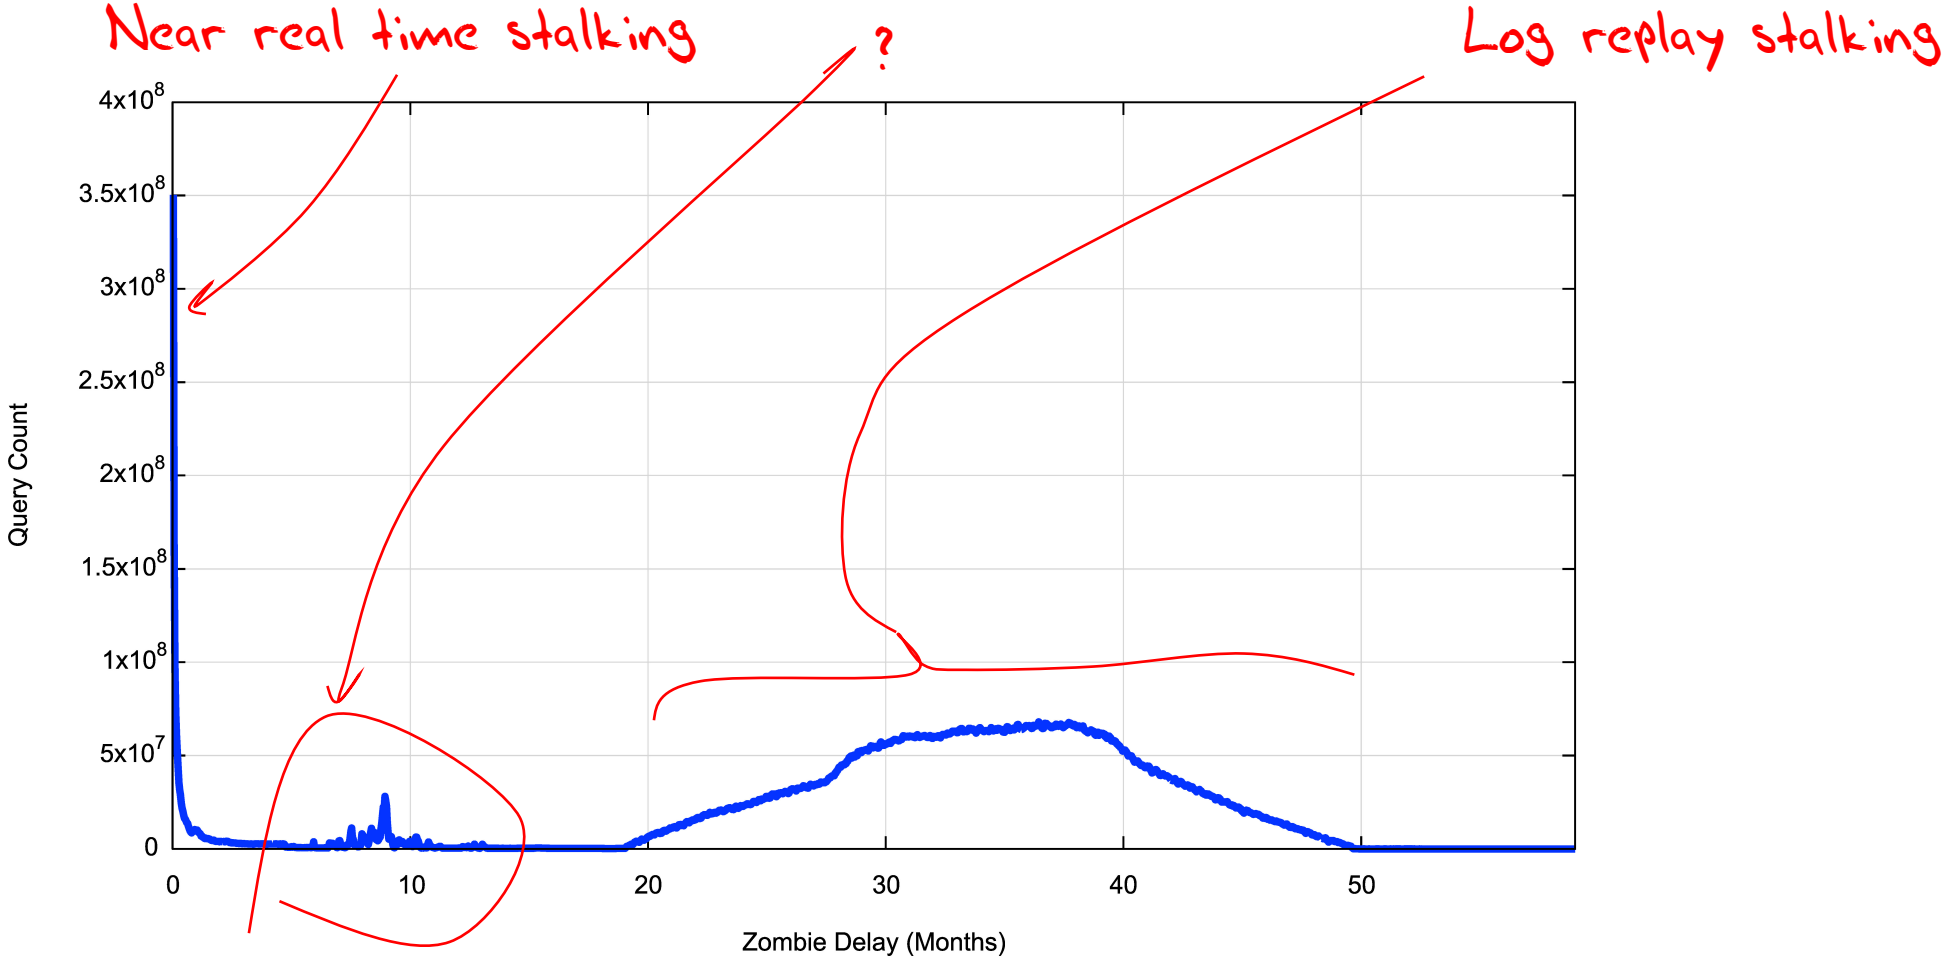

# DNS Stalking by Query Age

What about the first 24 hours?

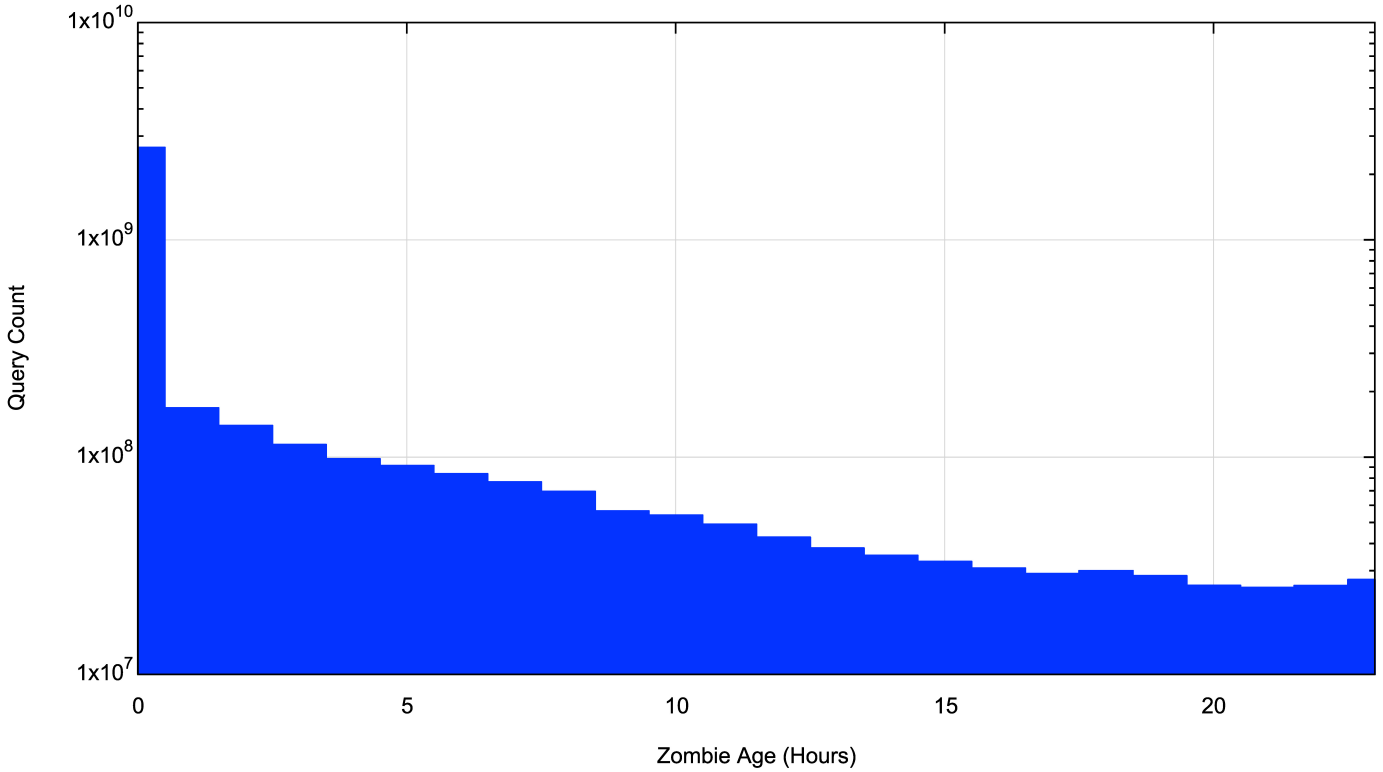

# Do Zombie experiments repeat?

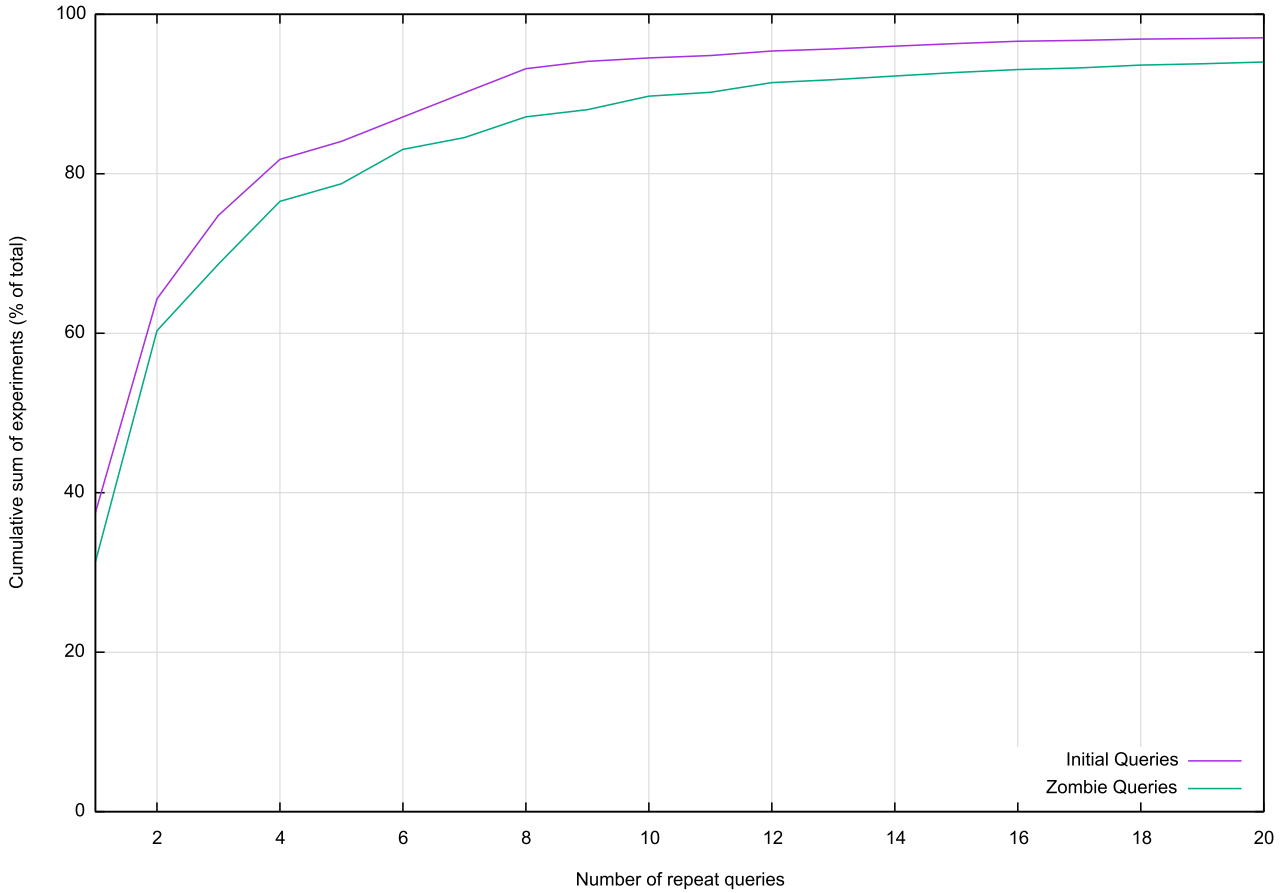


Data gathered across January – May 2020

Most zombies have a similar repeat distribution to the original query set for low volume repeats.

80% of experiments generated 4 or fewer queries in the first 30 seconds, while 78% of zombie experiment queries were repeated 4 or fewer times.

The outliers are one name that was queried 117K times in 30 seconds and one zombie name that was queried 2,171,541 times over 150 days

# Do Zombie queries repeat?

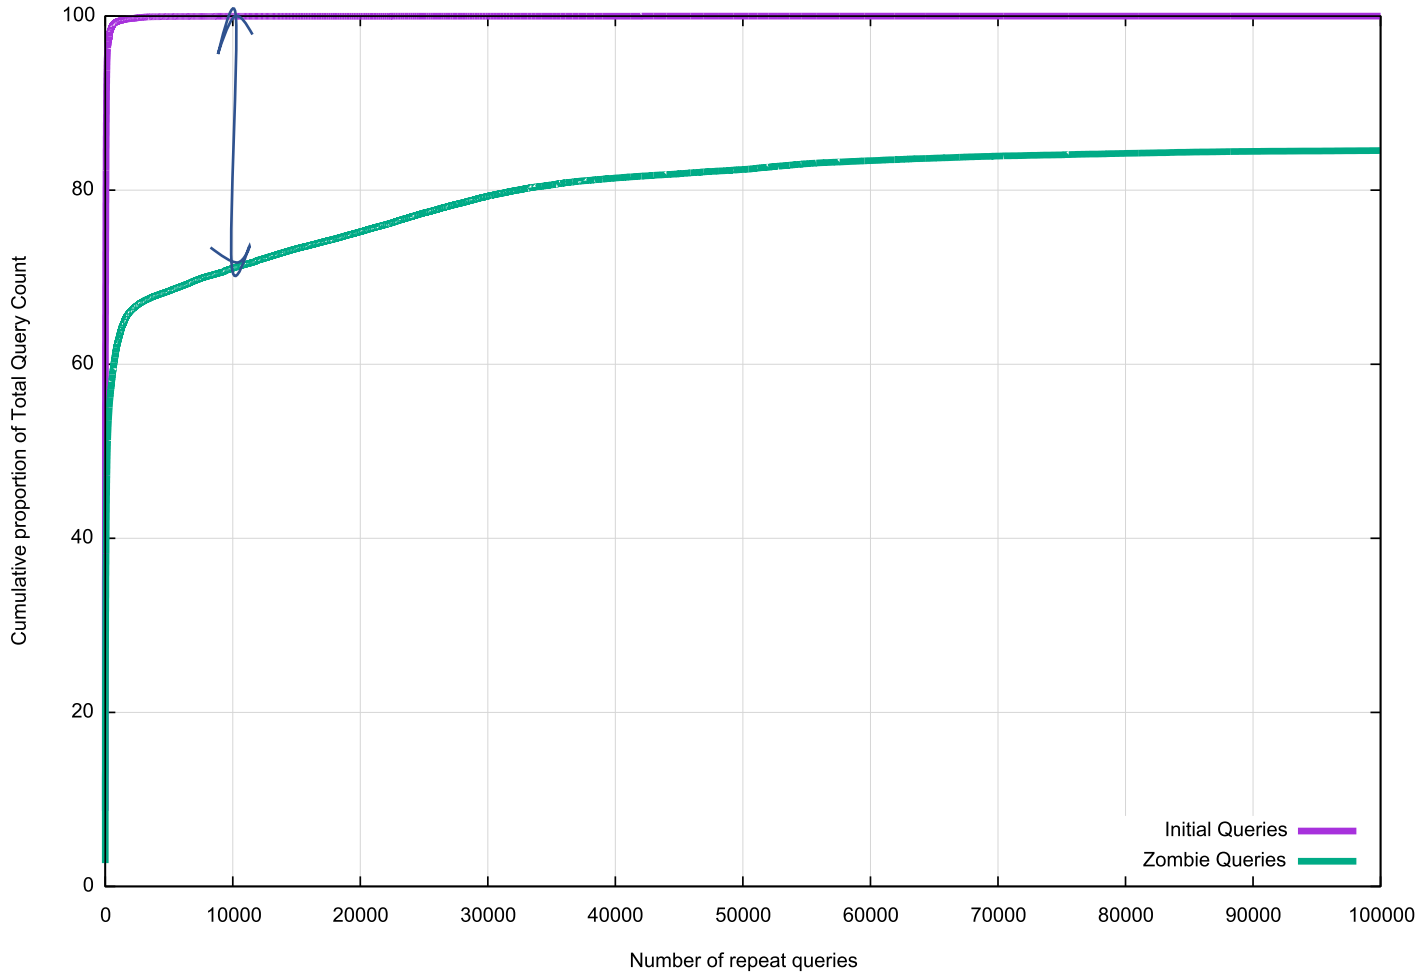


Data gathered across January – May 2020

This shows the distribution of repeated queries.

There is a long tail of high repeat count zombie queries.

1/3 of all zombie queries are part of a repeat sequence of between 10K to 2M repeats  over this 150 day interval

# Top Stalkers by Resolver Origin

| | ASN | Query Count | AS Name |
|---|---|---|---|
| 1 | 15169 | 13,471,892,786 | GOOGLE, US, |
| 2 | 4837 | 469,202,187 | China UNICOM, CN |
| 3 | 797 | 300,313,314 | AMERITECH-AS, US, |
| 4 | 55836 | 144,125,017 | Reliance, Jio, Infocomm, Limited, IN |
| 5 | 45271 | 93,948,908 | Idea Cellular, IN, |
| 6 | 28573 | 84,488,457 | CLARO, BR |
| 7 | 16509 | 76,664,272 | AMAZON-02, US |
| 8 | 13335 | 71,876,261 | CLOUDFLARENET, US |
| 9 | 4134 | 53,559,644 | CHINANET, CN, |
| 10 | 7922 | 51,454,865 | COMCAST-7922, US |
| 11 | 14618 | 47,017,864 | AMAZON-AES, US |
| 12 | 38266 | 43,033,362 | Vodafone India, IN |
| 13 | 12322 | 32,788,750 | PROXAD, FR |
| 14 | 9808 | 27,734,921 | Guangdong Mobile, CN |
| 15 | 36692 | 25,723,872 | OPENDNS, US |
| 16 | 3462 | 22,779,266 | HINET, TW |
| 17 | 327931 | 20,640,108 | Optimum-Telecom-Algeria, DZ, |
| 18 | 2860 | 15,619,160 | NOS_COMUNICACOES, PT |
| 19 | 6799 | 15,271,784 | OTENET-GR, Athens, GR, |
| 20 | 8190 | 13,699,795 | MDNX, GB |
| 21 | 7018 | 13,199,438 | ATT-INTERNET4, US |
| 22 | 58542 | 13,170,596 | CHINA TELECOM-TIANJIN, CN |
| 23 | 25019 | 12,699,618 | SAUDINETSTC-AS, SA |
| 24 | 24445 | 11,501,883 | Henan Mobile, CN |
| 25 | 53813 | 11,111,416 | ZSCALER, US |

Data gathered across January – May 2020

# Top Stalkers by Resolver Origin

| | ASN | Query Count | Avg Delay (hours) | AS Name |
|---|---|---|---|---|
| 1 | 15169 | 13,471,892,786 | 26,355 | GOOGLE, US, |
| 2 | 4837 | 469,202,187 | 29 | China UNICOM, CN |
| 3 | 797 | 300,313,314 | 6,107 | AMERITECH-AS, US, |
| 4 | 55836 | 144,125,017 | 26 | Reliance, Jio, Infocomm, Limited, IN |
| 5 | 45271 | 93,948,908 | 1 | Idea Cellular, IN, |
| 6 | 28573 | 84,488,457 | 4 | CLARO, BR |
| 7 | 16509 | 76,664,272 | 3,556 | AMAZON-02, US |
| 8 | 13335 | 71,876,261 | 607 | CLOUDFLARENET, US |
| 9 | 4134 | 53,559,644 | 163 | CHINANET, CN, |
| 10 | 7922 | 51,454,865 | 27 | COMCAST-7922, US |
| 11 | 14618 | 47,017,864 | 7,315 | AMAZON-AES, US |
| 12 | 38266 | 43,033,362 | 4 | Vodafone India, IN |
| 13 | 12322 | 32,788,750 | 7 | PROXAD, FR |
| 14 | 9808 | 27,734,921 | 104 | Guangdong Mobile, CN |
| 15 | 36692 | 25,723,872 | 1,607 | OPENDNS, US |
| 16 | 3462 | 22,779,266 | 233 | HINET, TW |
| 17 | 327931 | 20,640,108 | 42 | Optimum-Telecom-Algeria, DZ, |
| 18 | 2860 | 15,619,160 | 51 | NOS_COMUNICACOES, PT |
| 19 | 6799 | 15,271,784 | 13,216 | OTENET-GR, Athens, GR, |
| 20 | 8190 | 13,699,795 | 876 | MDNX, GB |
| 21 | 7018 | 13,199,438 | 3,196 | ATT-INTERNET4, US |
| 22 | 58542 | 13,170,596 | 167 | CHINA TELECOM-TIANJIN, CN |
| 23 | 25019 | 12,699,618 | 148 | SAUDINETSTC-AS, SA |
| 24 | 24445 | 11,501,883 | 26 | Henan Mobile, CN |
| 25 | 53813 | 11,111,416 | 2 | ZSCALER, US |

Data gathered across January – May 2020

I've added an average delay value (hours) - the time between the original name and the replay query

Some resolvers are re-querying very old names, while others are performing re-query soon after the original query

# Same Day Stalking

- Now lets look at just one day of DNS query data
- And also limit the zombie filter to only look at re-queries that are between 30 and 86,400 seconds (i.e. 1 day zombies)
- While this data may include a component of log replay activity, its more likely to reflect near real time DNS query capture and replay

# Same Day Stalking: Whose customers are being watched?

Data gathered 3 June 2020

| ASN | Exps | Zombie | | Country |
|---|---|---|---|---|
| **135377** | 20 | 187 | 90.3% | UHGL-AS-AP UCloud (HK) Holdings Group Limited |
| **25820** | 59 | 139 | 70.2% | IT7NET |
| **17929** | 285 | 499 | 63.6% | JPMC-AP JPMorgan Chase APTI |
| **268880** | 47 | 72 | 60.5% | Maximus Net |
| **6674** | 59 | 81 | 57.9% | NATIONAL-BANK-OF-GREECE, GR |
| **197296** | 44 | 60 | 57.7% | MAXTEL, CZ |
| **7743** | 348 | 443 | 56.0% | AS-7743, US |
| **132203** | 693 | 872 | 55.7% | TENCENT-NET-AP-CN Tencent Building, CN |
| **58543** | 987 | 1,217 | 55.2% | CHINATELECOM-GUANGDONG-IDC Guangdong, CN |
| **34569** | 1,144 | 1,276 | 52.7% | NETWORX-BG Online Direct, BG |
| **7726** | 70 | 78 | 52.7% | FITC-AS, US |
| **198668** | 100 | 108 | 51.9% | TLAPNET, CZ |
| **24442** | 70 | 70 | 50.0% | BMW-AS-AP BMW Asia Technology Sdn Bhd, DE |
| **46313** | 71 | 70 | 49.6% | WAL-MART4, CA |
| **28665** | 386 | 379 | 49.5% | Predlink Rede de Telecomuniccoes Ltda, BR |
| **263929** | 96 | 93 | 49.2% | E. S. DAMASCENO EIRELI - ME, BR |
| **203629** | 57 | 55 | 49.1% | TERRAHOST, NO |
| **46313** | 112 | 107 | 48.9% | WAL-MART4, US |
| **134027** | 262 | 247 | 48.5% | SAMSUNG DATA SYSTEMS, IN |
| **4816** | 392 | 360 | 47.9% | CHINANET-IDC-GD China Telecom (Group), CN |
| **20426** | 151 | 137 | 47.6% | PWC-AS, US |
| **34400** | 2,713 | 2,447 | 47.4% | ASN-ETTIHADETISALAT, SA |
| **31250** | 90 | 81 | 47.4% | ONLINEDIRECT-AS, BG |
| **35432** | 864 | 775 | 47.3% | CABLENET-AS, CY |
| **12722** | 70 | 62 | 47.0% | RECONN, RU |
| **8200** | 67 | 59 | 46.8% | UPLINK-AS, KZ |
| **268483** | 135 | 118 | 46.6% | speed net redes servicos comunicacao ltda-me, BR |
| **20255** | 115 | 98 | 46.0% | Tecnowind S.A., UY |
| **797** | 139 | 118 | 45.9% | AMERITECH-AS, US |

This data was gathered for a single day, looking at the counts of ‘normal’ and ‘zombie’ queries, where the zombie query reflect a DNS name created in the previous 24 hours that was re-queried more than 30 seconds after the original query

# Same Day Stalking: Where are these stalked users?

Data gathered 3 June 2020

| CC | Exps | Zombie | % | | Country |
|---|---|---|---|---|---|
| AF | 259 | 169 | 39.5% | | Afghanistan |
| YE | 980 | 409 | 29.4% | | Yemen |
| DM | 274 | 110 | 28.6% | | Dominica |
| IE | 404 | 141 | 25.9% | | Ireland |
| SA | 67,133 | 19,530 | 22.5% | | Saudi Arabia |
| JP | 658 | 190 | 22.4% | | Japan |
| JP | 1,746 | 493 | 22.0% | | Japan |
| TC | 199 | 49 | 19.8% | | Turks and Caicos Islands |
| YE | 3,460 | 777 | 18.3% | | Yemen |
| BR | 219 | 47 | 17.7% | | Brazil |
| CY | 3,908 | 830 | 17.5% | | Cyprus |
| CN | 382,669 | 77,424 | 16.8% | | China |
| MG | 1,248 | 245 | 16.4% | | Madagascar |
| NL | 614 | 115 | 15.8% | | Netherlands |
| SG | 1,119 | 202 | 15.3% | | Singapore |
| UZ | 10,649 | 1,903 | 15.2% | | Uzbekistan |
| AF | 1,834 | 321 | 14.9% | | Afghanistan |
| HK | 1,994 | 326 | 14.1% | | Hong Kong Special Administrative Region of China |
| IN | 7,165 | 1,168 | 14.0% | | India |
| IR | 29,903 | 4,724 | 13.6% | | Iran (Islamic Republic of) |
| AI | 326 | 45 | 12.1% | | Anguilla |
| DJ | 423 | 58 | 12.1% | | Djibouti |
| MW | 651 | 83 | 11.3% | | Malawi |
| SG | 2,518 | 299 | 10.6% | | Singapore |
| CZ | 15,183 | 1,707 | 10.1% | | Czech Republic |
| SO | 5,705 | 622 | 9.8% | | Somalia |
| ID | 229 | 24 | 9.5% | | Indonesia |

This data was gathered for a single day, looking at the counts of ‘normal’ and ‘zombie’ queries, where the zombie query reflect a DNS name created in the previous 24 hours that was re-queried more than 30 seconds after the original query

# Same Day Stalking: Who's Asking?

| Resolver Address | Zombie Query Count | Origin AS | | Origin AS |
|---|---|---|---|---|
| 112.110.90.192 | 716,016 | 45271 | | ICLNET-AS-AP Idea Cellular Limited, IN |
| 60.215.138.162 | 179,565 | 4837 | | CHINA169-BACKBONE CHINA UNICOM China169 Backbone, CN |
| 60.215.138.171 | 167,115 | 4837 | | CHINA169-BACKBONE CHINA UNICOM China169 Backbone, CN |
| 1.38.3.167 | 151,522 | 38266 | | VODAFONE-IN Vodafone India Ltd., IN |
| 111.6.229.46 | 150,159 | 24445 | | CMNET-V4HENAN-AS-AP Henan Mobile Communications Co., CN |
| 111.6.229.47 | 149,520 | 24445 | | CMNET-V4HENAN-AS-AP Henan Mobile Communications Co., CN |
| 1.38.3.168 | 125,343 | 38266 | | VODAFONE-IN Vodafone India Ltd., IN |
| 39.102.45.112 | 97,908 | 37963 | | CNNIC-ALIBABA-CN-NET-AP Hangzhou Alibaba Advertising Co., CN |
| 83.223.39.126 | 95,164 | 30247 | | VOONAMI, US |
| 83.223.50.125 | 91,360 | 48921 | | VEROTEL-AS, NL |
| 129.45.117.142 | 84,692 | 327931 | | Optimum-Telecom-Algeria, DZ |
| 123.57.1.38 | 83,086 | 37963 | | CNNIC-ALIBABA-CN-NET-AP Hangzhou Alibaba Advertising Co., CN |
| 129.45.67.143 | 83,075 | 327931 | | Optimum-Telecom-Algeria, DZ |
| 129.45.106.22 | 78,470 | 327931 | | Optimum-Telecom-Algeria, DZ |
| 129.45.68.119 | 78,368 | 327931 | | Optimum-Telecom-Algeria, DZ |
| 39.102.52.129 | 72,091 | 37963 | | CNNIC-ALIBABA-CN-NET-AP Hangzhou Alibaba Advertising Co., CN |
| 123.56.255.23 | 71,778 | 37963 | | CNNIC-ALIBABA-CN-NET-AP Hangzhou Alibaba Advertising Co., CN |
| 125.5.210.212 | 53,277 | 7629 | | EPLDT-AS-AP 5F L.V. Locsin Bldg, PH |
| 60.215.138.163 | 52,401 | 4837 | | CHINA169-BACKBONE CHINA UNICOM China169 Backbone, CN |
| 60.215.138.166 | 51,909 | 4837 | | CHINA169-BACKBONE CHINA UNICOM China169 Backbone, CN |
| 3.6.17.111 | 51,182 | 16509 | | AMAZON-02, US |
| 60.215.138.244 | 50,899 | 4837 | | CHINA169-BACKBONE CHINA UNICOM China169 Backbone, CN |
| 60.215.138.229 | 46,160 | 4837 | | CHINA169-BACKBONE CHINA UNICOM China169 Backbone, CN |
| 89.107.115.1 | 46,047 | 39812 | | Kamensk-Uralsky, RU |
| 60.215.138.164 | 45,417 | 4837 | | CHINA169-BACKBONE CHINA UNICOM China169 Backbone, CN |

Data gathered 3 June 2020

This data was gathered for a single day, looking at the counts of 'normal' and 'zombie' queries, where the zombie query reflect a DNS name created in the previous 24 hours that was re-queried more than 30 seconds after the original query

# Same Day Stalking: Who's Asking Again?

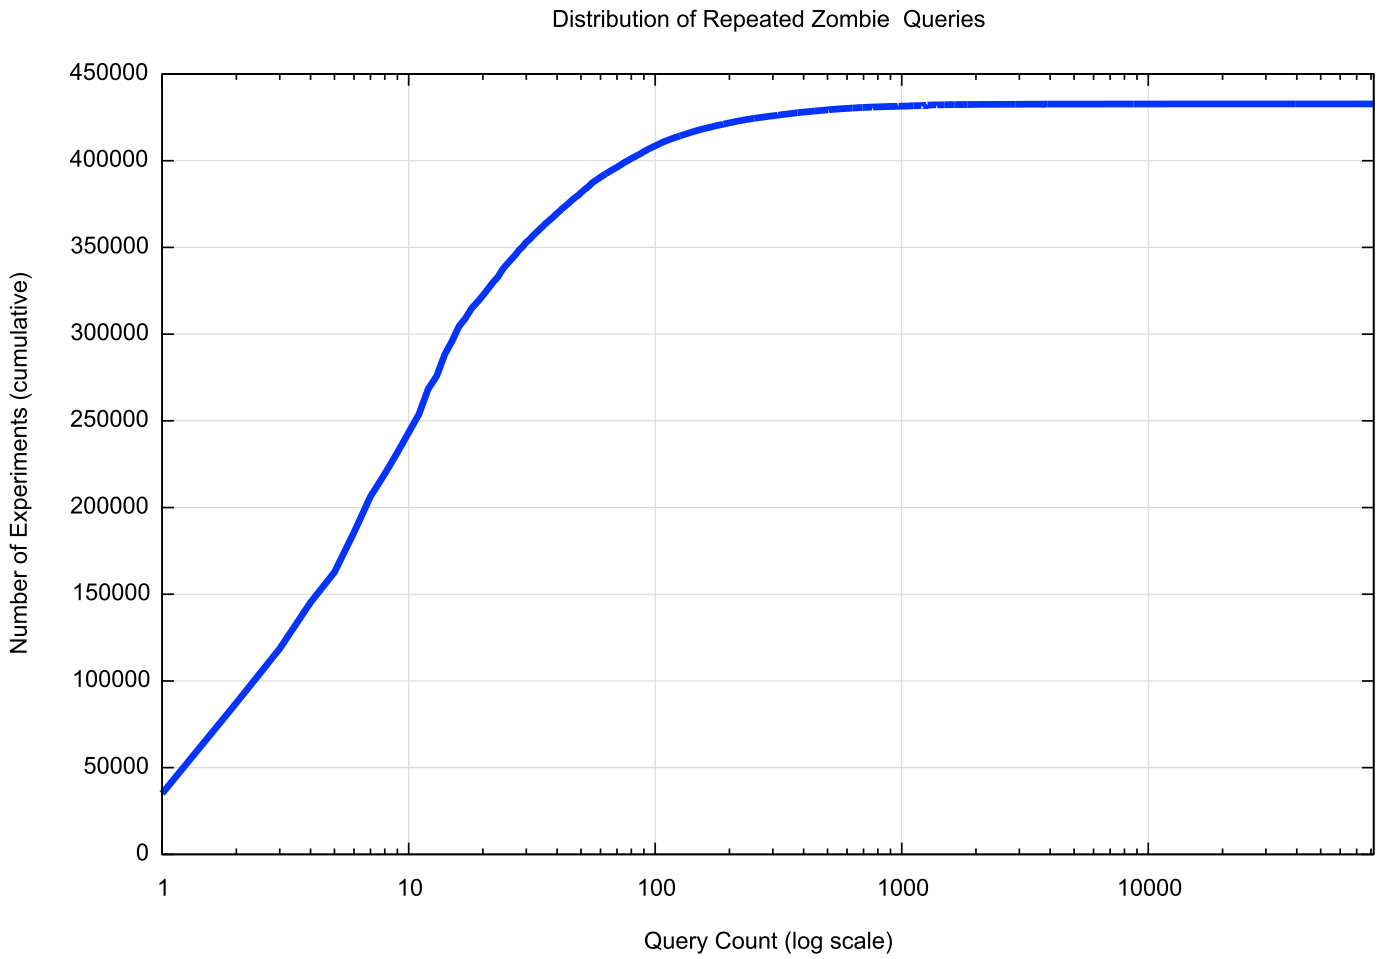


Data gathered 3 June 2020

This data was gathered for a single day, looking at the counts of 'normal' and 'zombie' queries, where the zombie query reflect a DNS name created in the previous 24 hours that was re-queried more than 30 seconds after the original query.

This shows the distribution of repeated zombie queries in the 24 hour period

90% of experiments attracted repeated zombie queries

50% of experiments attracted 10 or more repeated zombie queries

Max of 82,061 queries

This data set is just a small glimpse into the larger picture of DNS log capture and replay activity

Thanks!