# ANYway: Measuring the Amplification DDoS Potential of Domains

17th International Conference on Network and Service Management // Metaverse, the Street

Olivier van der Toorn <o.i.vandertoorn@utwente.nl>
Johannes Krupp <johannes.krupp@cispa.de>
Mattijs Jonker <m.jonker@utwente.nl>
Roland van Rijswijk-Deij <r.m.vanrijswijk@utwente.nl>
Christian Rossow <rossow@cispa.de>
Anna Sperotto <a.sperotto@utwente.nl>

October 14, 2021

University of Twente and CISPA.

UNIVERSITY OF TWENTE. CISPA

# Introduction

# Who am I?

- Ph.D. student from the University of Twente

# Who am I?

- Ph.D. student from the University of Twente
- Works on threat detection using active DNS data

# Who am I?

- Ph.D. student from the University of Twente
- Works on threat detection using active DNS data


Contact details

tide-project.nl

o.i.vandertoorn@utwente.nl

# EvilCorp and Web&Co

- Imagine this town...

# EvilCorp and Web&Co

- Imagine this town...
- Here we have EvilCorp and Web&Co, who are both website hosters.

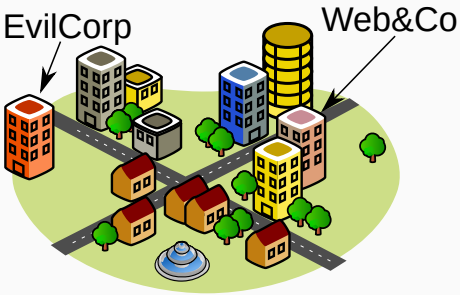

# EvilCorp and Web&Co

- Imagine this town...
- Here we have EvilCorp and Web&Co, who are both website hosters.
- You happen to work for EvilCorp.

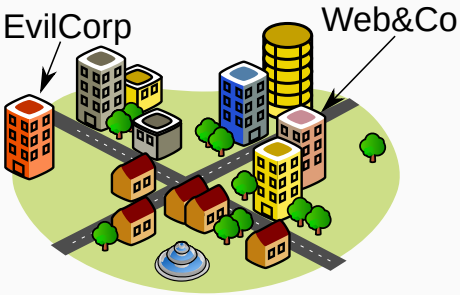

# EvilCorp and Web&Co

- Imagine this town...
- Here we have EvilCorp and Web&Co, who are both website hosters.
- You happen to work for EvilCorp.
- EvilCorp is... well... evil.

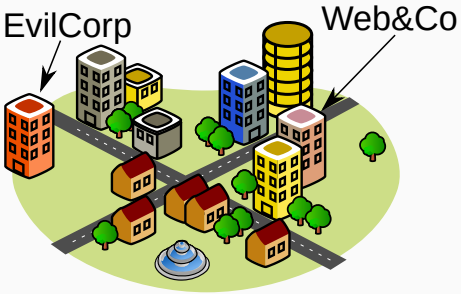

# EvilCorp and Web&Co

- Imagine this town...
- Here we have EvilCorp and Web&Co, who are both website hosters.
- You happen to work for EvilCorp.
- EvilCorp is... well... evil.

- You are tasked with performing a DDoS attack against Web&Co.

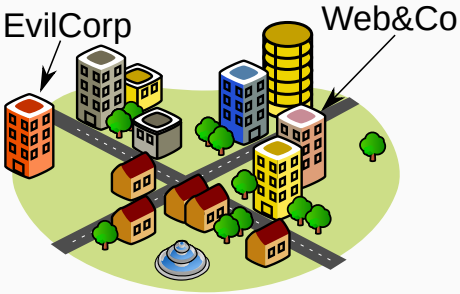

# EvilCorp and Web&Co

- Imagine this town…
- Here we have EvilCorp and Web&Co, who are both website hosters.
- You happen to work for EvilCorp.
- EvilCorp is… well… evil.

- You are tasked with performing a DDoS attack against Web&Co.
- ‘anyway.example.’ is the domain you will use.

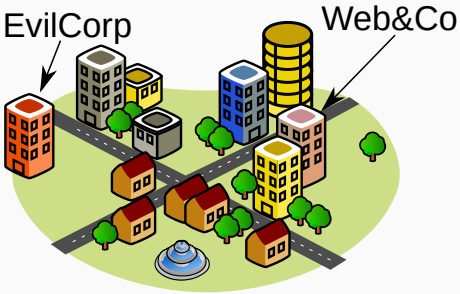


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- Query size (anyway.example.): 43 bytes

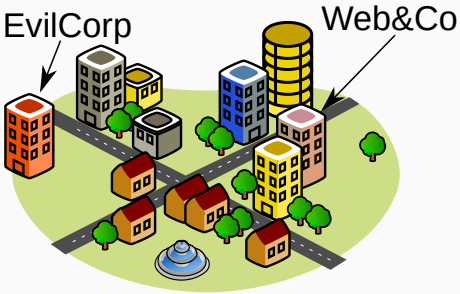


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- Query size (anyway.example.): 43 bytes
- ‘A’ query response size: 59 bytes

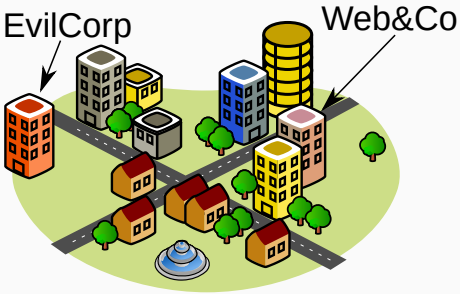


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- Query size (anyway.example.): 43 bytes
- ‘A’ query response size: 59 bytes
- Amplification factor: 1.4

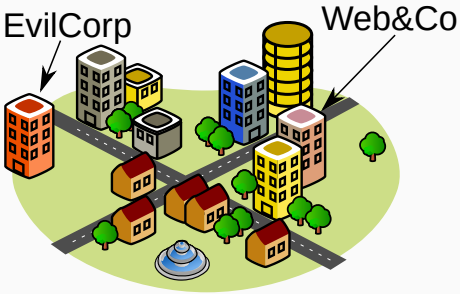


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- Query size (anyway.example.): 43 bytes
- 'A' query response size: 59 bytes
- Amplification factor: 1.4

$$BAF = \frac{len\,(UDP\ payload)\ amplifier\ to\ victim}{len\,(UDP\ payload)\ attacker\ to\ amplifier}$$

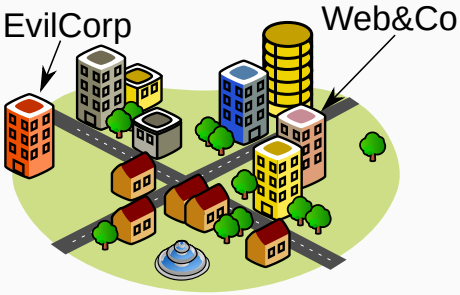


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- Query size (anyway.example.): 43 bytes
- ‘A’ query response size: 59 bytes
- Amplification factor: 1.4

- EvilCorp has a botnet of a 100 bots.
- Each bot is capable of performing 1000 queries per second.

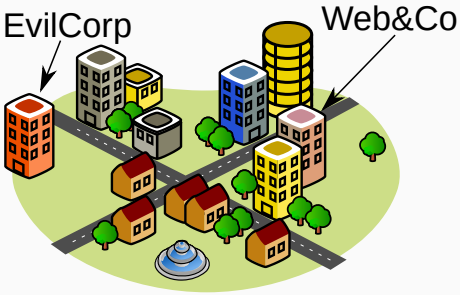


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- Query size (anyway.example.): 43 bytes
- 'A' query response size: 59 bytes
- Amplification factor: 1.4

- EvilCorp has a botnet of a 100 bots.
- Each bot is capable of performing 1000 queries per second.

- Total volume: $(59 * 8) * 100 * 1000 = 47.2\text{Mbit/s}$

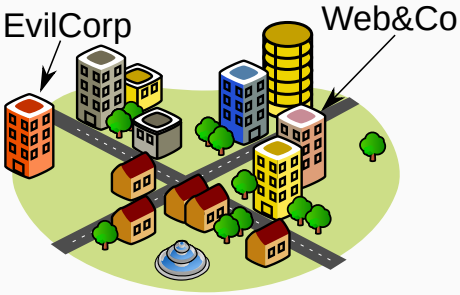


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- Query size (anyway.example.): 43 bytes
- 'A' query response size: 59 bytes
- Amplification factor: 1.4

- EvilCorp has a botnet of a 100 bots.
- Each bot is capable of performing 1000 queries per second.

- Total volume: $(59 * 8) * 100 * 1000 = 47.2\text{Mbit/s}$

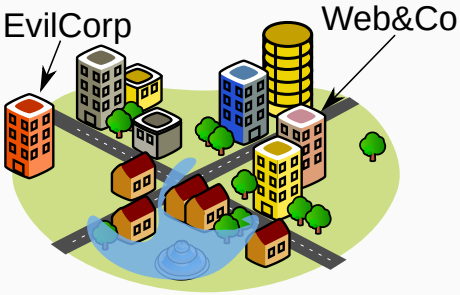


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- You switch from an ‘A’ query to ‘ANY’ queries.

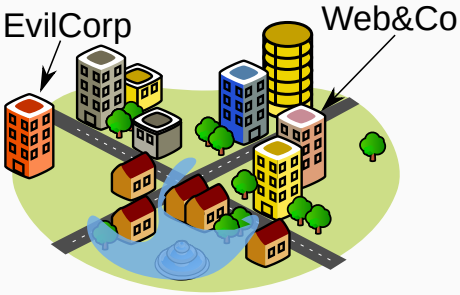


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- You switch from an ‘A’ query to ‘ANY’ queries.
- Query size (anyway.example.): 43 bytes

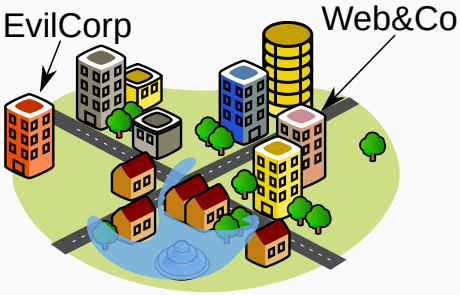


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- You switch from an ‘A’ query to ‘ANY’ queries.

- Query size (anyway.example.): 43 bytes
- ‘ANY’ query response size: 150 bytes

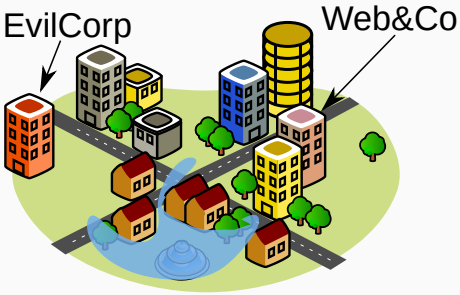


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- You switch from an ‘A’ query to ‘ANY’ queries.

- Query size (anyway.example.): 43 bytes
- ‘ANY’ query response size: 150 bytes
- Amplification factor: 3.5

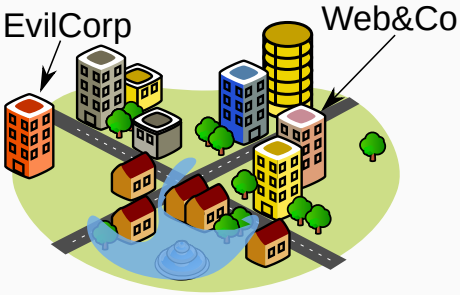


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- You switch from an ‘A’ query to ‘ANY’ queries.

- Query size (anyway.example.): 43 bytes
- ‘ANY’ query response size: 150 bytes
- Amplification factor: 3.5

- Total volume: $(150 * 8) * 100 * 1000 = 120\text{Mbit/s}$

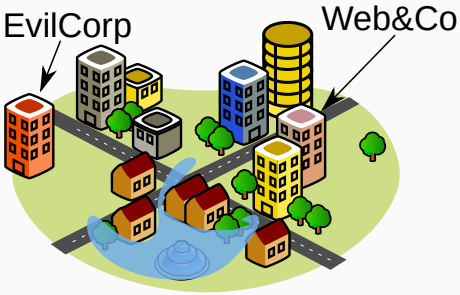


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- You switch from an ‘A’ query to ‘ANY’ queries.

- Query size (anyway.example.): 43 bytes
- ‘ANY’ query response size: 150 bytes
- Amplification factor: 3.5

- Total volume: $(150 * 8) * 100 * 1000 = 120\text{Mbit/s}$

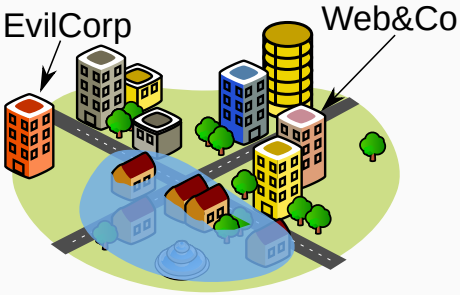


| rdtype | count | size |
|---|---|---|
| A | 1 | 16 |
| AAAA | 1 | 28 |
| MX | 1 | 32 |
| NS | 1 | 31 |

# EvilCorp and Web&Co

- You create a new domain, ‘anyway-ddos.example.’, specially made for DDoS attacks.

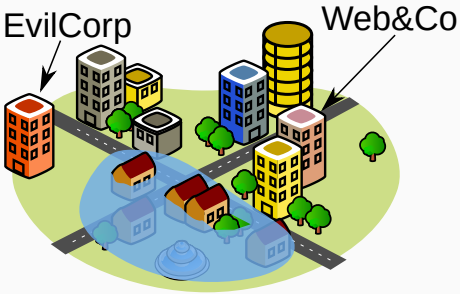

# EvilCorp and Web&Co

- You create a new domain, ‘anyway-ddos.example.’, specially made for DDoS attacks.

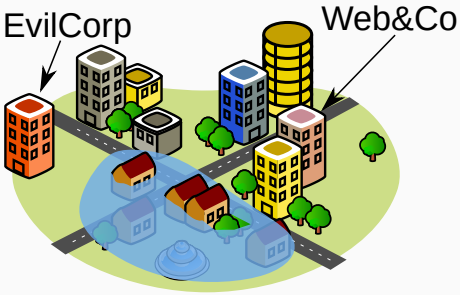


| rdtype | count | size |
|---|---|---|
| A | 30 | 480 |
| AAAA | 30 | 840 |
| MX | 30 | 960 |
| NS | 30 | 930 |

# EvilCorp and Web&Co

- You create a new domain, ‘anyway-ddos.example.’, specially made for DDoS attacks.
- Query size (anyway-ddos.example.): 48 bytes

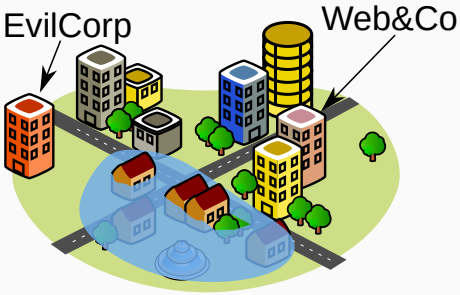


| rdtype | count | size |
|---|---|---|
| A | 30 | 480 |
| AAAA | 30 | 840 |
| MX | 30 | 960 |
| NS | 30 | 930 |

# EvilCorp and Web&Co

- You create a new domain, ‘anyway-ddos.example.’, specially made for DDoS attacks.

- Query size (anyway-ddos.example.): 48 bytes
- ‘ANY’ query response size: 3,258 bytes

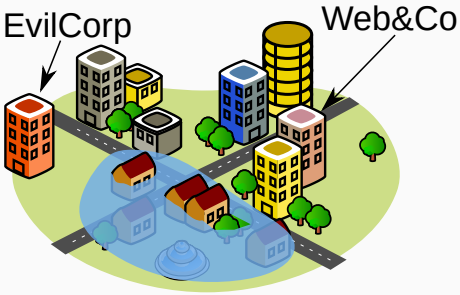


| rdtype | count | size |
|---|---|---|
| A | 30 | 480 |
| AAAA | 30 | 840 |
| MX | 30 | 960 |
| NS | 30 | 930 |

# EvilCorp and Web&Co

- You create a new domain, ‘anyway-ddos.example.’, specially made for DDoS attacks.

- Query size (anyway-ddos.example.): 48 bytes
- ‘ANY’ query response size: 3,258 bytes
- Amplification factor: 67.9

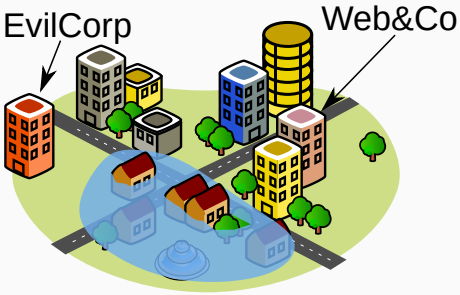


| rdtype | count | size |
|---|---|---|
| A | 30 | 480 |
| AAAA | 30 | 840 |
| MX | 30 | 960 |
| NS | 30 | 930 |

# EvilCorp and Web&Co

- You create a new domain, ‘anyway-ddos.example.’, specially made for DDoS attacks.

- Query size (anyway-ddos.example.): 48 bytes
- ‘ANY’ query response size: 3,258 bytes
- Amplification factor: 67.9

- Total volume:
  $(3258 * 8) * 100 * 1000 = 2.6\text{Gbit/s}$

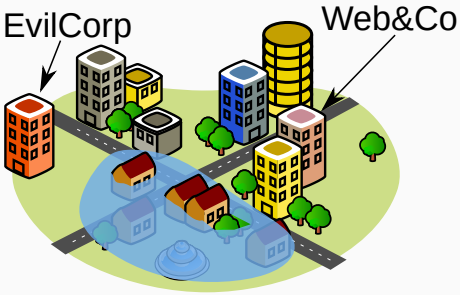


| rdtype | count | size |
|---|---|---|
| A | 30 | 480 |
| AAAA | 30 | 840 |
| MX | 30 | 960 |
| NS | 30 | 930 |

# EvilCorp and Web&Co

- You create a new domain, ‘anyway-ddos.example.’, specially made for DDoS attacks.

- Query size (anyway-ddos.example.): 48 bytes
- ‘ANY’ query response size: 3,258 bytes
- Amplification factor: 67.9

- Total volume:
  $(3258 * 8) * 100 * 1000 = 2.6\text{Gbit/s}$

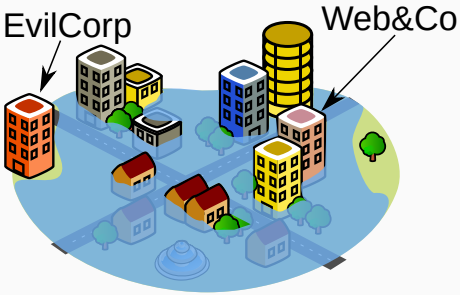


| rdtype | count | size |
|---|---|---|
| A | 30 | 480 |
| AAAA | 30 | 840 |
| MX | 30 | 960 |
| NS | 30 | 930 |

# DDoS for Dummies

- After this presentation you will be able to:

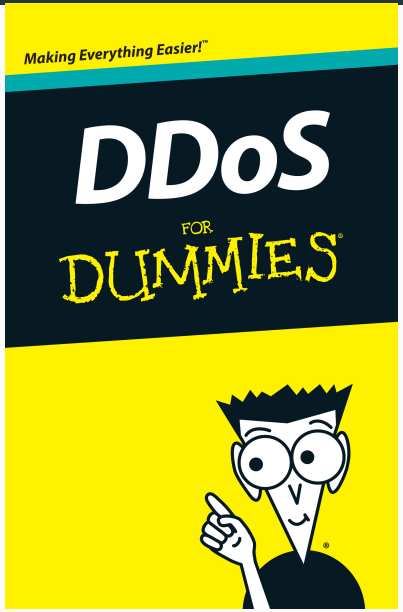

# DDoS for Dummies

- After this presentation you will be able to:
  - Take down your competing web hoster

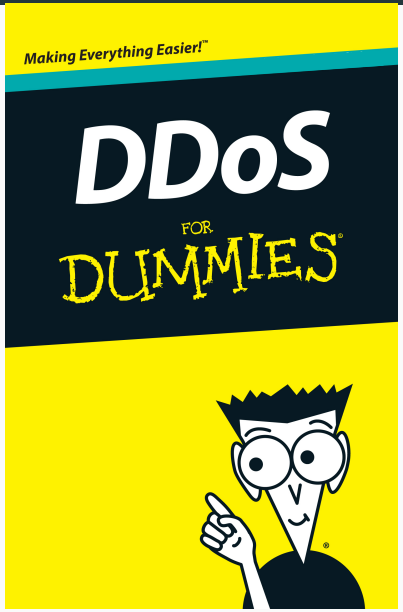

# DDoS for Dummies

- After this presentation you will be able to:
  - Take down your competing web hoster
  - The Internet connection of other gamers

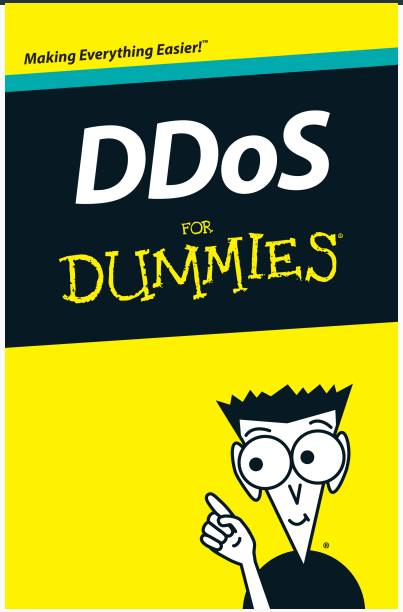

# DDoS for Dummies

- After this presentation you will be able to:
  - Take down your competing web hoster
  - The Internet connection of other gamers
  - And.... more!

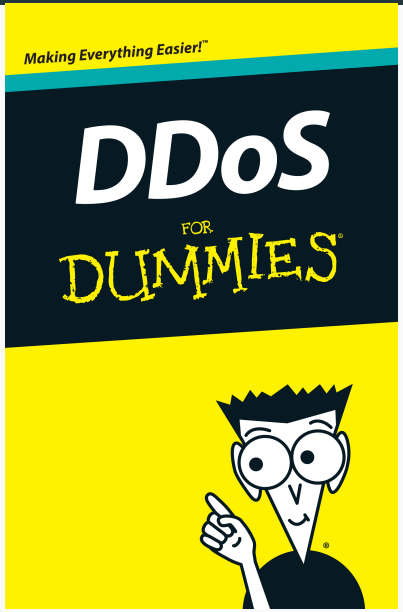

# DDoS for Dummies

- ~~After this presentation you will be able to:~~
  - ~~Take down your competing web hoster~~
  - ~~The Internet connection of other gamers~~
  - ~~And.... more!~~

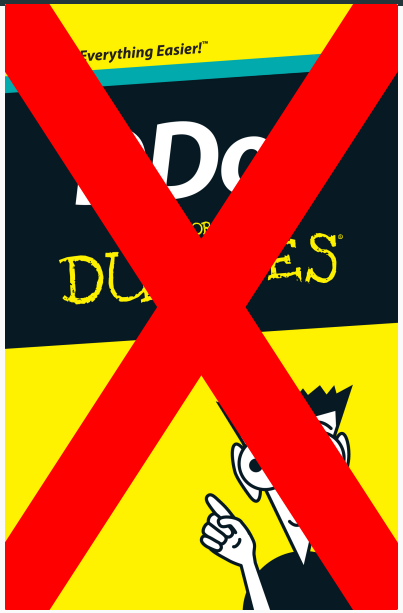

# DDoS for Dummies

- ~~After this presentation you will be able to:~~
  - ~~Take down your competing web hoster~~
  - ~~The Internet connection of other gamers~~
  - ~~And.... more!~~

- ANYway: Measuring the Amplification DDoS Potential of Domains

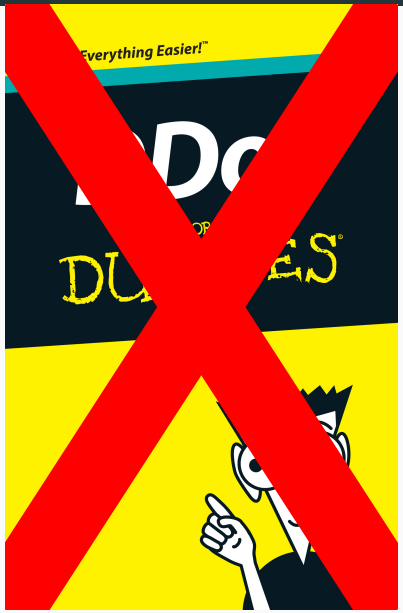

# DDoS for Dummies

- ~~After this presentation you will be able to:~~
  - ~~Take down your competing web hoster~~
  - ~~The Internet connection of other gamers~~
  - ~~And.... more!~~

- ANYway: Measuring the Amplification DDoS Potential of Domains
  - ANY Response Size Estimation

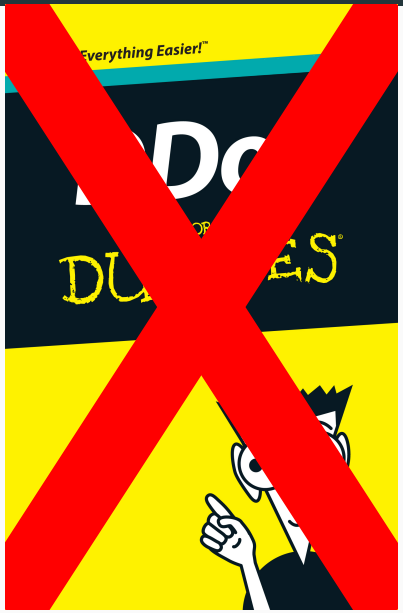

# DDoS for Dummies

- ~~After this presentation you will be able to:~~
  - ~~Take down your competing web hoster~~
  - ~~The Internet connection of other gamers~~
  - ~~And.... more!~~

- ANYway: Measuring the Amplification DDoS Potential of Domains
  - ANY Response Size Estimation
  - Ranking Domains

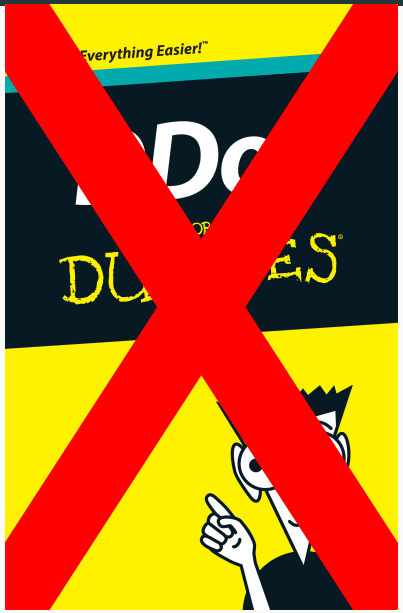

# DDoS for Dummies

- ~~After this presentation you will be able to:~~
  - ~~Take down your competing web hoster~~
  - ~~The Internet connection of other gamers~~
  - ~~And.... more!~~

- ANYway: Measuring the Amplification DDoS Potential of Domains
  - ANY Response Size Estimation
  - Ranking Domains
  - The Impact of Dropping ANY

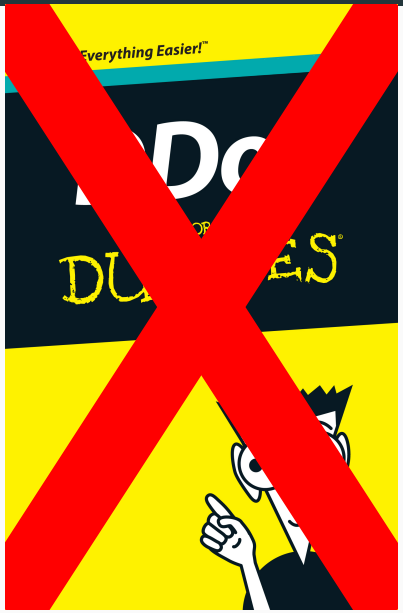

# Measurement based study

# Measurement based study

This work is based on measurements, we mainly use two sources of measurement data.

- AmpPot, for domains used in DDoS attacks.
- OpenINTEL, for the size estimations.

For both datasets we used data from between January 2019 until December 2020.

# AmpPot

The AmpPot project[1] operates a set of geographically and logically distributed amplification DDoS honeypots. These honeypots mimic reflectors for popular, abusable, UDP-based protocols, DNS included.

- Select domains with at least 10 queries during an attack.
- Leaves us with 100 domains used in 448,156 attacks.

[1] Lukas Krämer et al. “AmpPot: Monitoring and Defending Against Amplification DDoS Attacks”. In: *Proceedings of the 18th International Symposium on Research in Attacks, Intrusions, and Defenses - Volume 9404*. RAID 2015. Kyoto, Japan: Springer-Verlag, 2015, pp. 615–636. ISBN: 9783319263618. URL: https://doi.org/10.1007/978-3-319-26362-5_28.

# OpenINTEL

OpenINTEL is an active DNS measurement platform currently measuring over 65% of the DNS name space.

- Measures 236M second-level domains on a daily basis.
- With 12 resource records per domain.
- We used measurement results for the first of every month from January 2019 until December 2020.
- OpenINTEL does not perform, by design, ‘ANY’ queries.

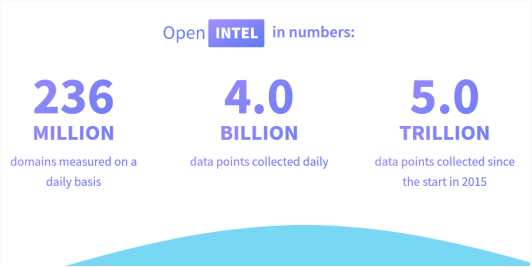


Website

https://openintel.nl/

# ANY Response Size Estimation

# How to estimate ANY response sizes?

- A DNS response consists of the following parts:

# How to estimate ANY response sizes?

- A DNS response consists of the following parts:
  - a header

# How to estimate ANY response sizes?

- A DNS response consists of the following parts:
  - a header
  - the original question

# How to estimate ANY response sizes?

- A DNS response consists of the following parts:
  - a header
  - the original question
  - a response to the question

# How to estimate ANY response sizes?

- A DNS response consists of the following parts:
  - a header
  - the original question
  - a response to the question
- With DNSSEC, the response can be further divided into an answer and signatures.

# How to estimate ANY response sizes?

- A DNS response consists of the following parts:
  - a header
  - the original question
  - a response to the question
- With DNSSEC, the response can be further divided into an answer and signatures.
- The response to an ANY query can be seen as a combination of header, question and a collection of answers and signatures, to answer for all records.

# How to estimate ANY response sizes?

- A DNS response consists of the following parts:
  - a header
  - the original question
  - a response to the question
- With DNSSEC, the response can be further divided into an answer and signatures.
- The response to an ANY query can be seen as a combination of header, question and a collection of answers and signatures, to answer for all records.

**Table 1:** Estimation of DNS response size[2]

| Record type | Equation |
|---|---|
| header size | $= 12 + 4 + \text{len(domain name)} + 1 + 11$ |
| signature size | $= 30 + \text{len(domain name)} + 1 + \text{size(rrsig)}$ |
| A size | $= 12 + 4$ |
| AAAA size | $= 12 + 16$ |
| CAA size | $= 12 + 2 + \text{len(CAA)}$ |
| CDNSKEY size | $= 12 + 4 + \text{sizeof(CDNSKEY)}$ |
| CDS size | $= 12 + 4 + \text{len(CDS)}$ |
| DNSKEY size | $= 12 + 4 + \text{sizeof(DNSKEY)}$ |
| DS size | $= 12 + 4 + \text{len(DS)}$ |
| MX size | $= 12 + 1 + \text{len(mail exchange)} + 1$ |
| NS size | $= 12 + \text{len(nameserver)} + 1$ |
| NSEC3PARAM size | $= 12 + 4 + \text{sizeof(salt)}$ |
| SOA size | $= 12 + 16 + \text{len(mname)} + \text{len(rname)}$ |
| TXT size | $= 12 + \text{len(text)} + 2$ |

[2]The size estimations of 'anyway.example.' were made with Table 1.

# Is estimating ANY size difficult?

Estimating DNS properties is not straight forward. The DNS often adds data to be ‘helpful’.

# Is estimating ANY size difficult?

Estimating DNS properties is not straight forward. The DNS often adds data to be ‘helpful’.

- From our data we cannot see authority or additional sections.

# Is estimating ANY size difficult?

Estimating DNS properties is not straight forward. The DNS often adds data to be ‘helpful’.

- From our data we cannot see authority or additional sections.
- We do not see the use of RFC 8482.

# Validating the estimations

- Zmap scan for open resolvers, resulted in 2,000 resolvers of which 804 were selected.

# Validating the estimations

- Zmap scan for open resolvers, resulted in 2,000 resolvers of which 804 were selected.
- We selected 1,000 domains with an estimated amplification factor larger than eight, but with an estimated response size of fewer than 4,096 bytes.

# Validating the estimations

- Zmap scan for open resolvers, resulted in 2,000 resolvers of which 804 were selected.
- We selected 1,000 domains with an estimated amplification factor larger than eight, but with an estimated response size of fewer than 4,096 bytes.
- We queried each of the resolvers in our set for all the domains in our selection in a randomized order.
  - We set the EDNS0 payload size to 4,096 bytes.
  - We requested a DNSSEC signed answer (DO).
  - And requested for recursive resolution (RD).

# Validation results

- After making sure we could compare our estimations to the measurements, we calculated how much we over- or underestimate.

# Validation results

- After making sure we could compare our estimations to the measurements, we calculated how much we over- or underestimate.

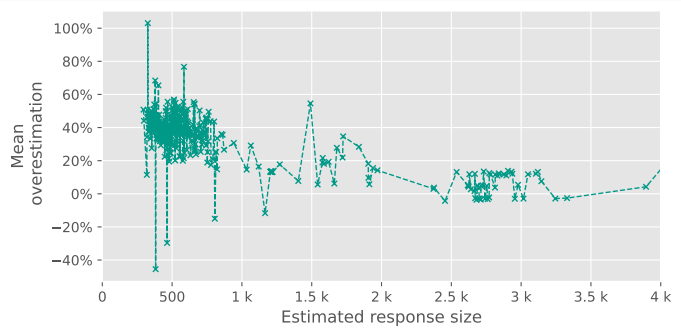


**Figure 1:** Correlation overestimation and estimated size

# Validation results

- After making sure we could compare our estimations to the measurements, we calculated how much we over- or underestimate.
- For ‘smaller’ domains ($< 1,000$ bytes) our estimations are roughly 20%-60% larger.

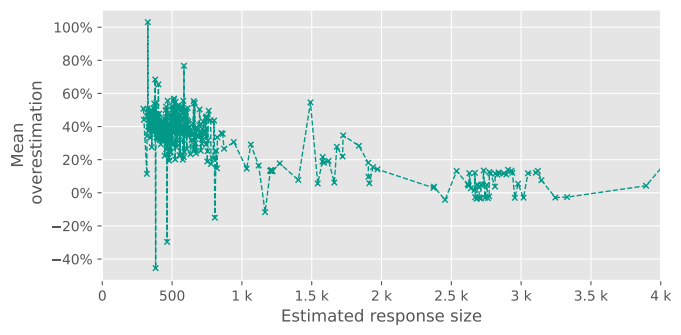


**Figure 1:** Correlation overestimation and estimated size

# Validation results

- After making sure we could compare our estimations to the measurements, we calculated how much we over- or underestimate.
- For ‘smaller’ domains ($< 1,000$ bytes) our estimations are roughly 20%-60% larger.
- For ‘larger’ domains ($> 2,048$ bytes) we see an average overestimation of 5%.

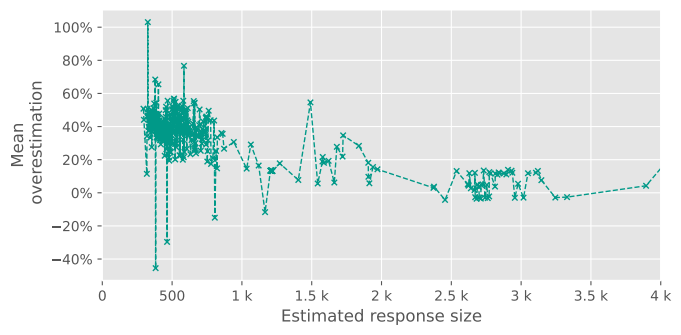


**Figure 1:** Correlation overestimation and estimated size

# Key takeaway

*Key takeaway: Estimating ANY response sizes from active DNS measurements leads to a size overestimation, for large domains, of 5%, making it a viable solution to identify DDoS potent domains.*

# Ranking Domains

# Methodology for ranking domains

- For all domains in OpenINTEL we estimate the ANY response size.

# Methodology for ranking domains

- For all domains in OpenINTEL we estimate the ANY response size.
- Then we rank domains from the largest amplification factor to the smallest.

# Methodology for ranking domains

- For all domains in OpenINTEL we estimate the ANY response size.
- Then we rank domains from the largest amplification factor to the smallest.
- For our stability analysis we take samples of the first of the month between January 2019 and December 2020.

# Methodology for ranking domains

- For all domains in OpenINTEL we estimate the ANY response size.
- Then we rank domains from the largest amplification factor to the smallest.
- For our stability analysis we take samples of the first of the month between January 2019 and December 2020.
- We filter for domains with an amplification factor higher than eight and an estimated response size of below 4,096 bytes.

# Methodology for ranking domains

- For all domains in OpenINTEL we estimate the ANY response size.
- Then we rank domains from the largest amplification factor to the smallest.
- For our stability analysis we take samples of the first of the month between January 2019 and December 2020.
- We filter for domains with an amplification factor higher than eight and an estimated response size of below 4,096 bytes.
- Then we select domains that were present for all 24 samples.

# Methodology for ranking domains

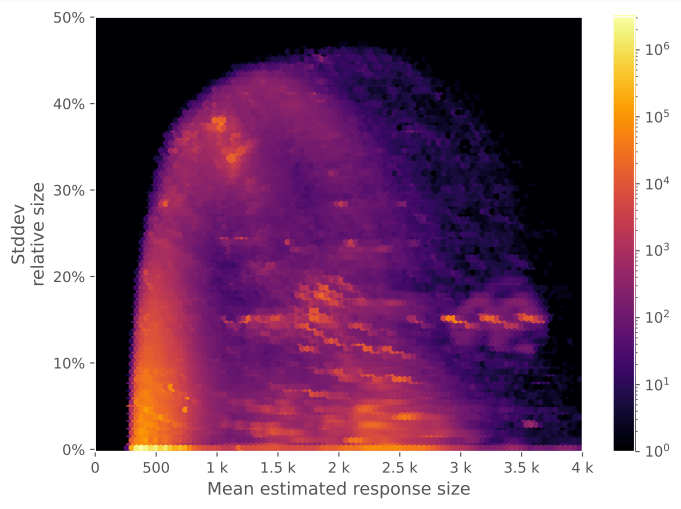


**Figure 2:** Stability of the estimated size

# Methodology for ranking domains

- No clear correlation between estimated response size and standard deviation.

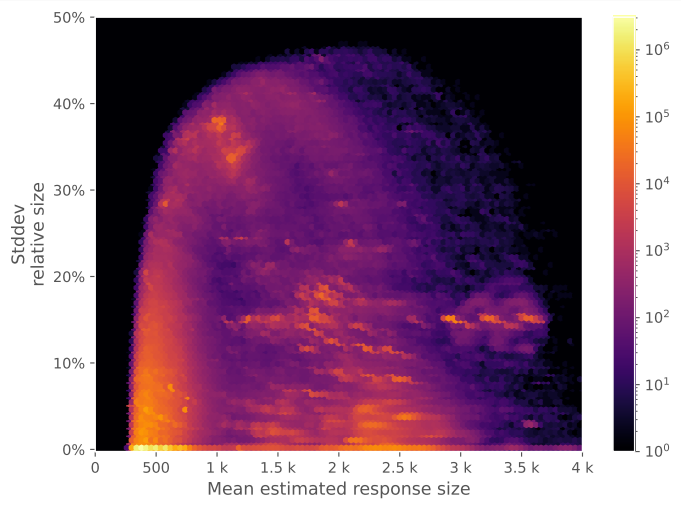


**Figure 2:** Stability of the estimated size

# Methodology for ranking domains

- No clear correlation between estimated response size and standard deviation.
- Two groups:
  - domains with zero standard deviation

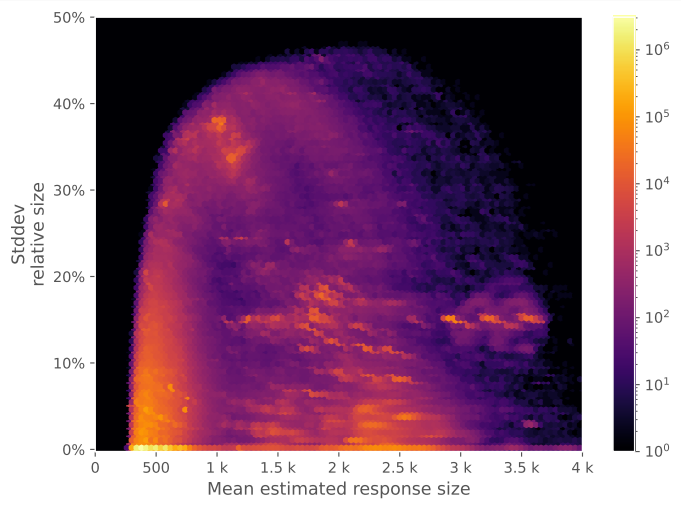


**Figure 2:** Stability of the estimated size

# Methodology for ranking domains

- No clear correlation between estimated response size and standard deviation.
- Two groups:
  - domains with zero standard deviation
  - domains with a non-zero standard deviation

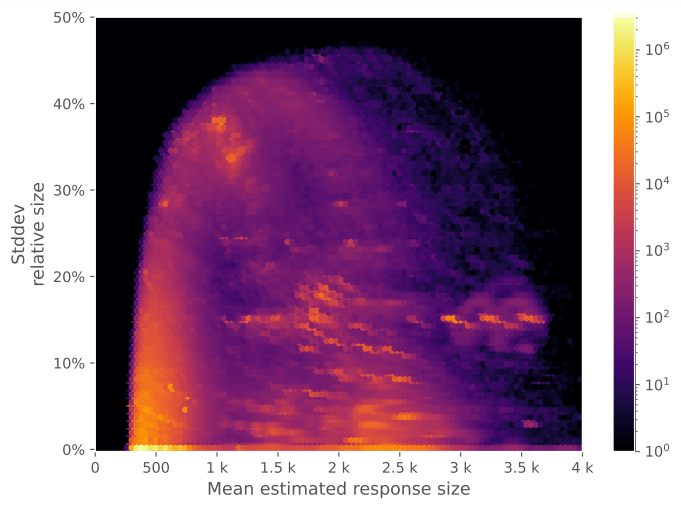


**Figure 2:** Stability of the estimated size

# Methodology for ranking domains

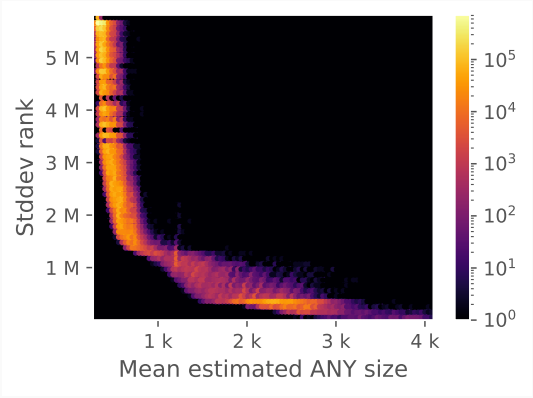


**Figure 3:** Standard deviation of domains without changes in size

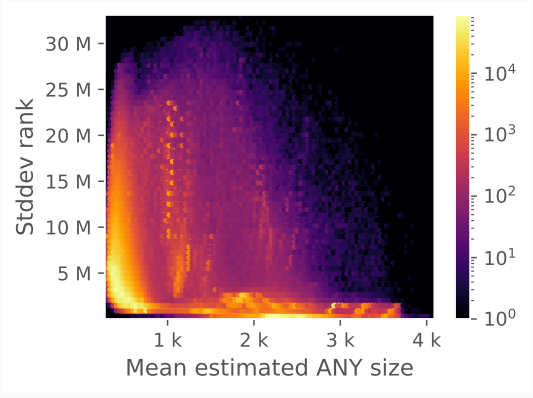


**Figure 4:** Standard deviation of domains with changes in size

# Domain ranking result

- Selecting domains used in attacks.

# Domain ranking result

- Selecting domains used in attacks.

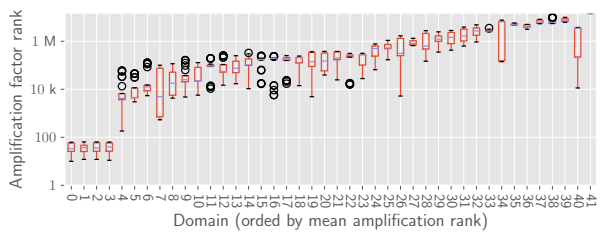

# Domain ranking result

- Selecting domains used in attacks.
- Four domains have reached ranks ten, eleven, and twelve.

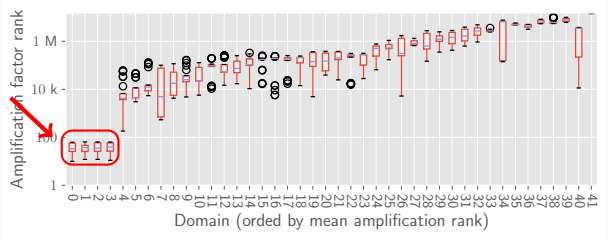

# Domain ranking result

- Selecting domains used in attacks.
- Four domains have reached ranks ten, eleven, and twelve.
- However, there are many domains used in attacks with much lower ranks.

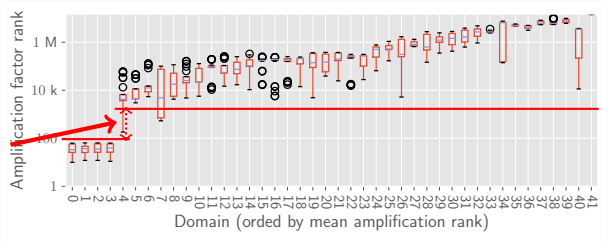

# Key takeaway

*Key takeaway: Domains observed in attacks are among the largest domains available. However, our ranking shows that there are still a sizable number of domains larger than the ones used so far that could easily be exploited.*

# The Impact of Dropping ANY

# How do we estimate the impact of dropping ANY?

- We can adapt our estimation to a single type, rather than combining all types for an ANY query.

# How do we estimate the impact of dropping ANY?

- We can adapt our estimation to a single type, rather than combining all types for an ANY query.
- First, we looked into moving from ANY queries to the ‘next-best’ type per domain.

# How do we estimate the impact of dropping ANY?

- We can adapt our estimation to a single type, rather than combining all types for an ANY query.
- First, we looked into moving from ANY queries to the ‘next-best’ type per domain.
- Second, we looked into moving from ANY queries to a fixed record type.

# Is dropping ANY requests effective?

Moving from ‘ANY’ queries to the next-best type:

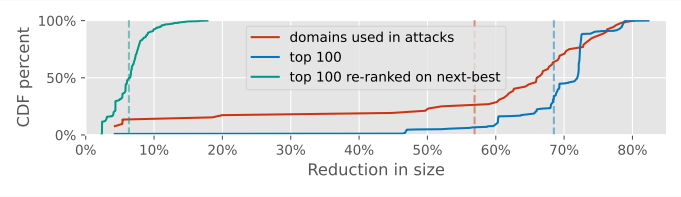


**Figure 5:** Reduction in size by dropping ANY

# Is dropping ANY requests effective?

Moving from ‘ANY’ queries to the next-best type:

- Domains used in attacks have a mean reduction of 57%, with 75% being reduced by 52% or more.

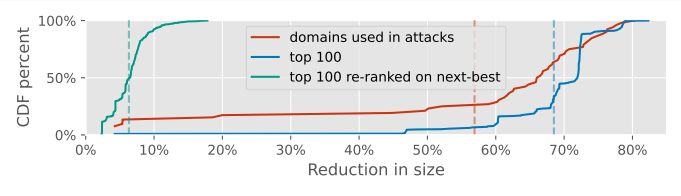


**Figure 5:** Reduction in size by dropping ANY

# Is dropping ANY requests effective?

Moving from ‘ANY’ queries to the next-best type:

- Domains used in attacks have a mean reduction of 57%, with 75% being reduced by 52% or more.
- Domains in the top 100 have a mean reduction of 69%, with 75% being reduced by 68% or more.

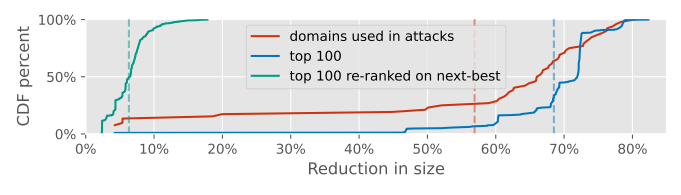


**Figure 5:** Reduction in size by dropping ANY

# Is dropping ANY requests effective?

Moving from ‘ANY’ queries to the next-best type:

- Domains used in attacks have a mean reduction of 57%, with 75% being reduced by 52% or more.
- Domains in the top 100 have a mean reduction of 69%, with 75% being reduced by 68% or more.
- Domains in the ‘new’ top 100 have a mean reduction of 9%, with 75% being reduced by 8% or *less*.

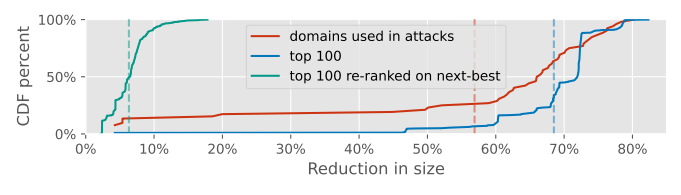


**Figure 5:** Reduction in size by dropping ANY

# Is dropping ANY requests effective?

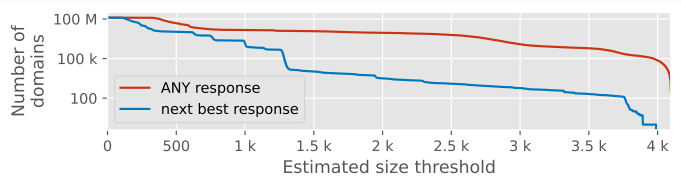


**Figure 6:** Number of domains exceeding the estimated size threshold

# Is dropping ANY requests effective?

- There are still around a thousand domains which are larger than 2,048 bytes without the use of ‘ANY’ queries.

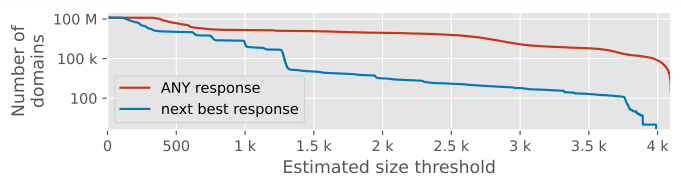


**Figure 6:** Number of domains exceeding the estimated size threshold

# Is dropping ANY requests effective?

Moving from ‘ANY’ query to a fixed record type:

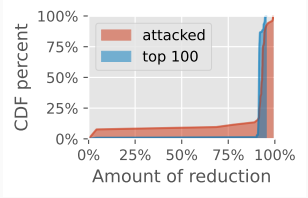


(a) A query type

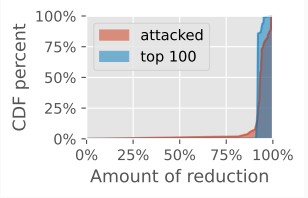


(b) AAAA query type

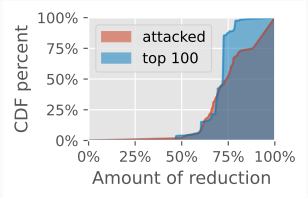


(c) DNSKEY query type

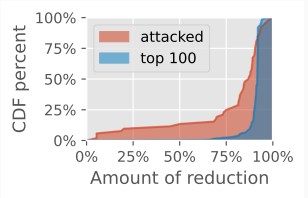


(d) TXT query type

**Figure 7:** Reduction by moving from ANY to a specific query type

# Is dropping ANY requests effective?

Moving from ‘ANY’ query to a fixed record type:

- The query types standing out from this analysis are:

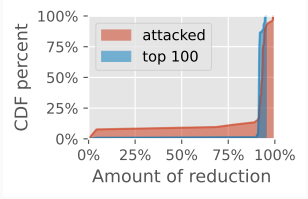


(a) A query type

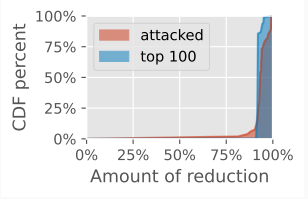


(b) AAAA query type

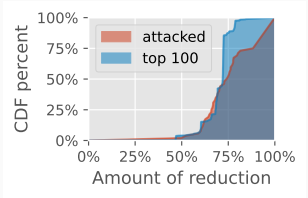


(c) DNSKEY query type

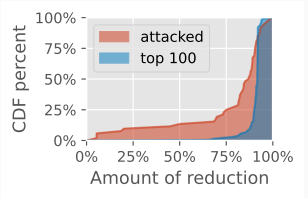


(d) TXT query type

**Figure 7:** Reduction by moving from ANY to a specific query type

# Is dropping ANY requests effective?

Moving from ‘ANY’ query to a fixed record type:

- The query types standing out from this analysis are:
  - DNSKEY: mean reduction of 76%

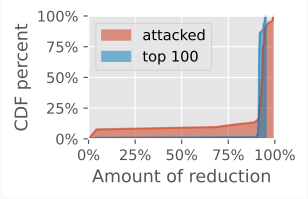


(a) A query type

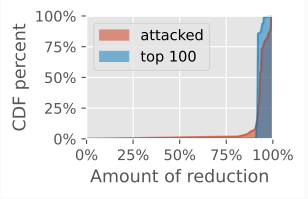


(b) AAAA query type

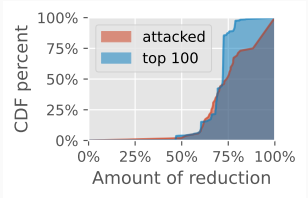


(c) DNSKEY query type

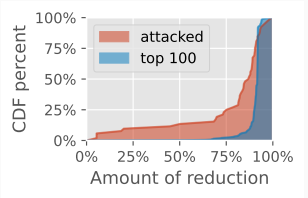


(d) TXT query type

**Figure 7:** Reduction by moving from ANY to a specific query type

# Is dropping ANY requests effective?

Moving from ‘ANY’ query to a fixed record type:

- The query types standing out from this analysis are:
  - DNSKEY: mean reduction of 76%
  - TXT: mean reduction of 79%

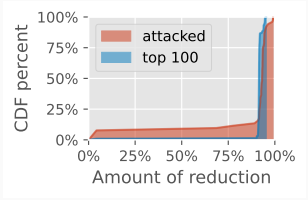


(a) A query type

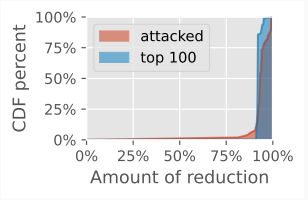


(b) AAAA query type

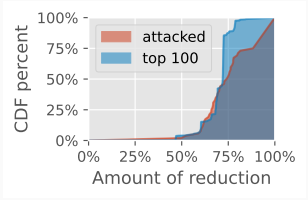


(c) DNSKEY query type

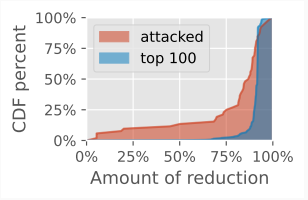


(d) TXT query type

**Figure 7:** Reduction by moving from ANY to a specific query type

# Is dropping ANY requests effective?

Moving from ‘ANY’ query to a fixed record type:

- The query types standing out from this analysis are:
  - DNSKEY: mean reduction of 76%
  - TXT: mean reduction of 79%
- Isn’t ≈80% enough reduction?

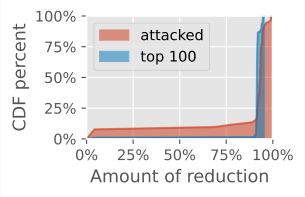


(a) A query type

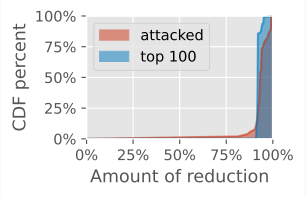


(b) AAAA query type

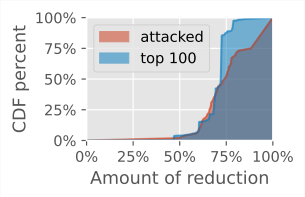


(c) DNSKEY query type

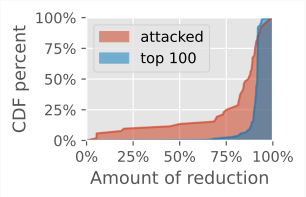


(d) TXT query type

**Figure 7:** Reduction by moving from ANY to a specific query type

# Is dropping ANY requests effective?

Moving from ‘ANY’ query to a fixed record type:

- The query types standing out from this analysis are:
  - DNSKEY: mean reduction of 76%
  - TXT: mean reduction of 79%
- Isn’t ≈80% enough reduction?
- TXT records are the likely candidate to replace ‘ANY’ queries.

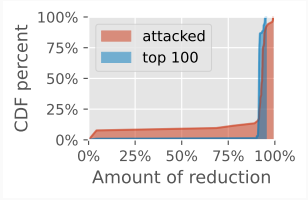


(a) A query type

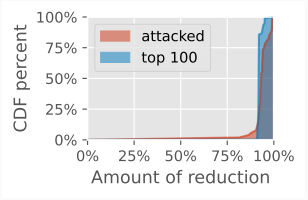


(b) AAAA query type

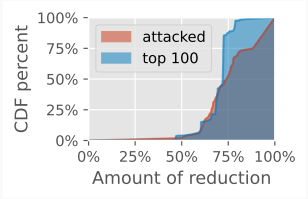


(c) DNSKEY query type

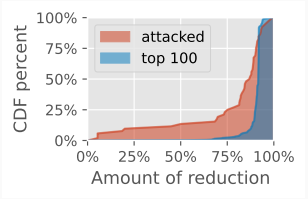


(d) TXT query type

**Figure 7:** Reduction by moving from ANY to a specific query type

# Categorization of TXT records

What is in these TXT records?

**Table 2:** DNS TXT record categories on 2020-12-31.

| Label | # of Records | % of Total | Plot |
|---|---|---|---|
| DNS TXT Records | 3,793 | 100% |  |
| Verification | 1,168 | 31% |  |
| Patterns | 890 | 23% |  |
| Miscellaneous | 698 | 19% |  |
| Encoded | 451 | 12% |  |
| Other | 432 | 11% |  |
| Email | 154 | 4% |  |

# Categorization of TXT records

What is in these TXT records?

- Selected the re-ranked top 100 for this analysis.

**Table 2:** DNS TXT record categories on 2020-12-31.

| Label | # of Records | % of Total | Plot |
|---|---|---|---|
| DNS TXT Records | 3,793 | 100% | 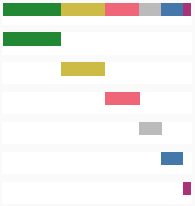 |
| Verification | 1,168 | 31% | |
| Patterns | 890 | 23% | |
| Miscellaneous | 698 | 19% | |
| Encoded | 451 | 12% | |
| Other | 432 | 11% | |
| Email | 154 | 4% | |

# Categorization of TXT records

What is in these TXT records?

- Selected the re-ranked top 100 for this analysis.
- Applied a TXT categorization method from earlier work[3].

**Table 2:** DNS TXT record categories on 2020-12-31.

| Label | # of Records | % of Total | Plot |
|---|---|---|---|
| DNS TXT Records | 3,793 | 100% | 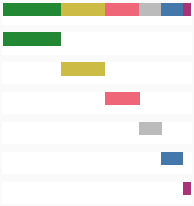 |
| Verification | 1,168 | 31% | |
| Patterns | 890 | 23% | |
| Miscellaneous | 698 | 19% | |
| Encoded | 451 | 12% | |
| Other | 432 | 11% | |
| Email | 154 | 4% | |

[3] Olivier van der Toorn et al. "TXTing 101: Finding Security Issues in the Long Tail of DNS TXT Records". In: *2020 IEEE European Symposium on Security and Privacy Workshops (EuroS PW)*. 2020.

# Categorization of TXT records

Most categories are seen with relatively few records per domain, generally below 20 records. Except:

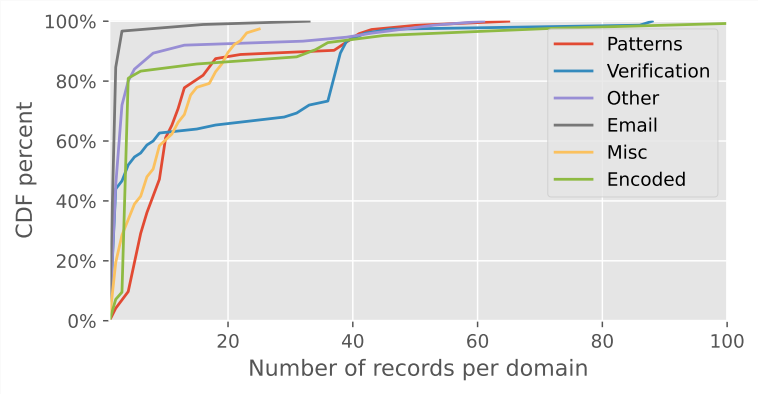


**Figure 8:** Number TXT records per domain

# Categorization of TXT records

Most categories are seen with relatively few records per domain, generally below 20 records. Except:

- Verification records; roughly 31% of domains has 30 records or more

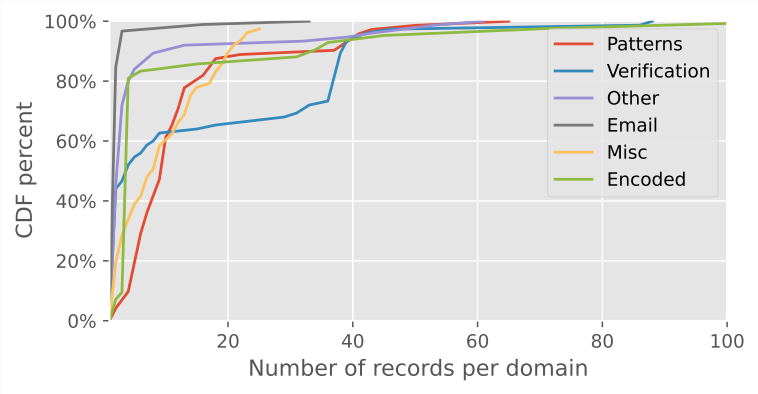


**Figure 8:** Number TXT records per domain

# Categorization of TXT records

- The categories with, on average, the longest records are:

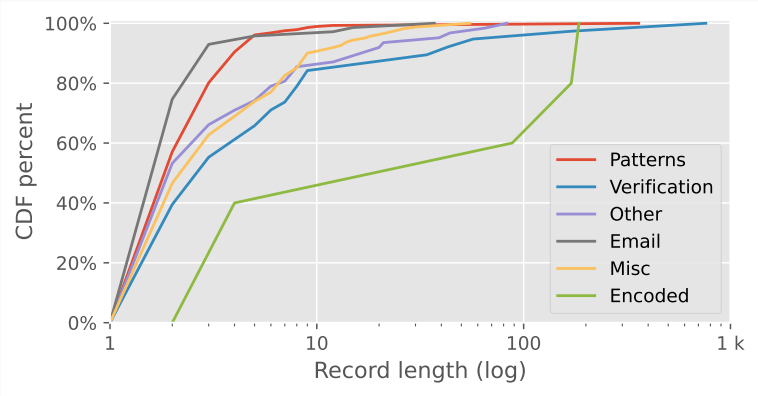


**Figure 9:** TXT record length

# Categorization of TXT records

- The categories with, on average, the longest records are:
  - Encoded; with an average length of 75 characters.

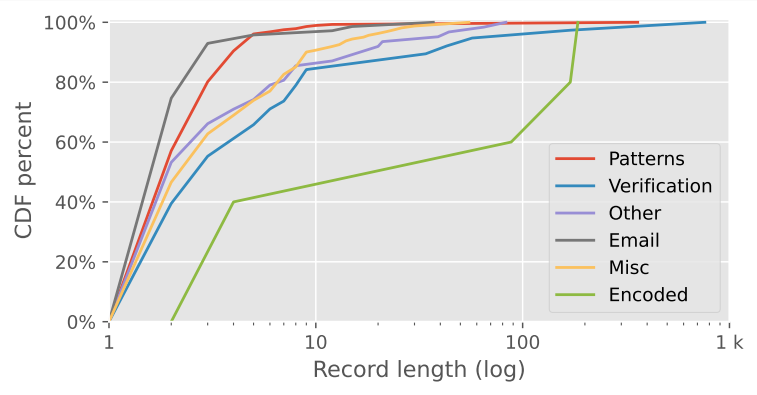


**Figure 9:** TXT record length

# Categorization of TXT records

- The categories with, on average, the longest records are:
  - Encoded; with an average length of 75 characters.
  - Verification: with an average length of 30 characters.

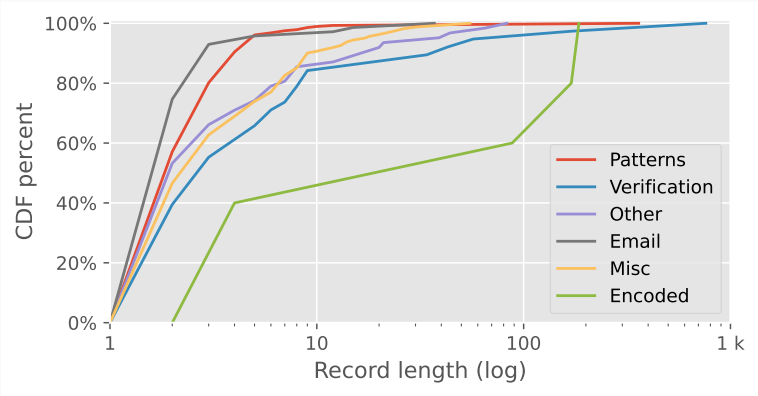


**Figure 9:** TXT record length

# Categorization of TXT records

- The categories with, on average, the longest records are:
  - Encoded; with an average length of 75 characters.
  - Verification: with an average length of 30 characters.
- This view changes, however, when we look at the total contribution instead of individual records.

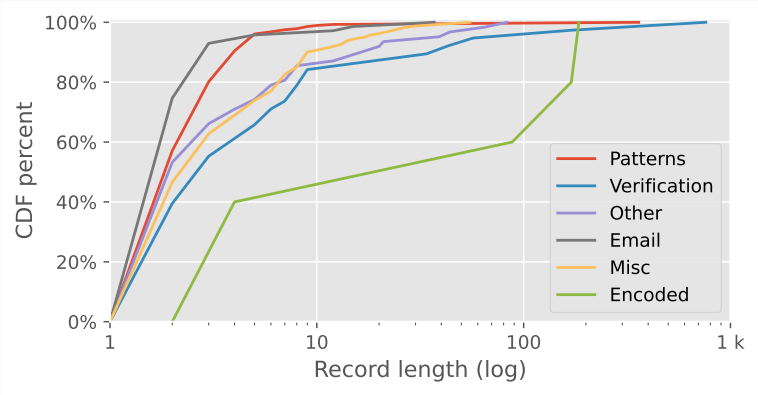


**Figure 9:** TXT record length

# Categorization of TXT records

- The ‘worst’ offenders are:

**Table 3:** DNS TXT record contributions.

| Label | Average Length (bytes) | % of TXT response | # of Domains |
|---|---|---|---|
| Patterns | 2,239 | 65% | 73 |
| Verification | 1,066 | 32% | 76 |
| Email | 1,010 | 35% | 92 |
| Miscellaneous | 888 | 26% | 78 |
| Encoded | 475 | 14% | 43 |
| Other | 389 | 13% | 76 |

# Categorization of TXT records

- The ‘worst’ offenders are:
  - Verification records

**Table 3:** DNS TXT record contributions.

| Label | Average Length (bytes) | % of TXT response | # of Domains |
|---|---|---|---|
| Patterns | 2,239 | 65% | 73 |
| Verification | 1,066 | 32% | 76 |
| Email | 1,010 | 35% | 92 |
| Miscellaneous | 888 | 26% | 78 |
| Encoded | 475 | 14% | 43 |
| Other | 389 | 13% | 76 |

# Categorization of TXT records

- The ‘worst’ offenders are:
  - Verification records
  - Pattern records

**Table 3:** DNS TXT record contributions.

| Label | Average Length (bytes) | % of TXT response | # of Domains |
|---|---|---|---|
| Patterns | 2,239 | 65% | 73 |
| Verification | 1,066 | 32% | 76 |
| Email | 1,010 | 35% | 92 |
| Miscellaneous | 888 | 26% | 78 |
| Encoded | 475 | 14% | 43 |
| Other | 389 | 13% | 76 |

# Categorization of TXT records

- The ‘worst’ offenders are:
  - Verification records
  - Pattern records
  - and Encoded records

**Table 3:** DNS TXT record contributions.

| Label | Average Length (bytes) | % of TXT response | # of Domains |
|---|---|---|---|
| Patterns | 2,239 | 65% | 73 |
| Verification | 1,066 | 32% | 76 |
| Email | 1,010 | 35% | 92 |
| Miscellaneous | 888 | 26% | 78 |
| Encoded | 475 | 14% | 43 |
| Other | 389 | 13% | 76 |

# Categorization of TXT records

- The ‘worst’ offenders are:
  - Verification records
  - Pattern records
  - and Encoded records
- Either because of their relatively long length, or because of the number of records per domain.

**Table 3:** DNS TXT record contributions.

| Label | Average Length (bytes) | % of TXT response | # of Domains |
|---|---|---|---|
| Patterns | 2,239 | 65% | 73 |
| Verification | 1,066 | 32% | 76 |
| Email | 1,010 | 35% | 92 |
| Miscellaneous | 888 | 26% | 78 |
| Encoded | 475 | 14% | 43 |
| Other | 389 | 13% | 76 |

# Key takeaway

*Key takeaway: Dropping responses to ANY queries is an effective way of reducing the response size of domains observed in DDoS attacks and of top ranked domains. However, the RR composition of several domains is such that, even when dropping ANY, a large response (>2,048 bytes) can easily be reached with another record type. Therefore dropping ANY might be only a temporary solution in the fight against DDoS.*

# Conclusions and Operational Considerations

# Operational Considerations

- RFC 8482 may not be the final answer

# Operational Considerations

- RFC 8482 may not be the final answer
- TXT records as a likely candidate
  - We urge DNS operators to carefully check their TXT records.

# Operational Considerations

- RFC 8482 may not be the final answer
- TXT records as a likely candidate
  - We urge DNS operators to carefully check their TXT records.
- Allow zone operators to suspend zones

# Conclusions

- Size estimation may be used as an early-warning system.

# Conclusions

- Size estimation may be used as an early-warning system.
- Linking measurements with observations from DDoS honeypots, we found attackers are already optimizing for the worst domains.

# Conclusions

- Size estimation may be used as an early-warning system.
- Linking measurements with observations from DDoS honeypots, we found attackers are already optimizing for the worst domains.
- However, there is “room for improvement”.

# Conclusions

- Size estimation may be used as an early-warning system.
- Linking measurements with observations from DDoS honeypots, we found attackers are already optimizing for the worst domains.
- However, there is “room for improvement”.
- Dropping support for ‘ANY’ queries decreases the response sizes by $\approx 70\%$ for 75% of the largest 100 domains.

# Conclusions

- Size estimation may be used as an early-warning system.
- Linking measurements with observations from DDoS honeypots, we found attackers are already optimizing for the worst domains.
- However, there is “room for improvement”.
- Dropping support for ‘ANY’ queries decreases the response sizes by $\approx 70\%$ for 75% of the largest 100 domains.
- Still a significant but manageable number of domains with bad amplification even without ‘ANY’ queries.

# Conclusions

- Size estimation may be used as an early-warning system.
- Linking measurements with observations from DDoS honeypots, we found attackers are already optimizing for the worst domains.
- However, there is “room for improvement”.
- Dropping support for ‘ANY’ queries decreases the response sizes by $\approx 70\%$ for 75% of the largest 100 domains.
- Still a significant but manageable number of domains with bad amplification even without ‘ANY’ queries.

# Conclusion

*We fully support RFC 8482, but we fear it is only a matter of time before DDoS attacks achieve the same traffic volume with a query type other than 'ANY'.*

# Conclusion

*We fully support RFC 8482, but we fear it is only a matter of time before DDoS attacks achieve the same traffic volume with a query type other than 'ANY'.*

*As a security community we need to start thinking of additional methods of 'solving' the DDoS problem.*

# Conclusion

*We fully support RFC 8482, but we fear it is only a matter of time before DDoS attacks achieve the same traffic volume with a query type other than 'ANY'.*

*As a security community we need to start thinking of additional methods of 'solving' the DDoS problem.*

Thank you for your time. Any questions?