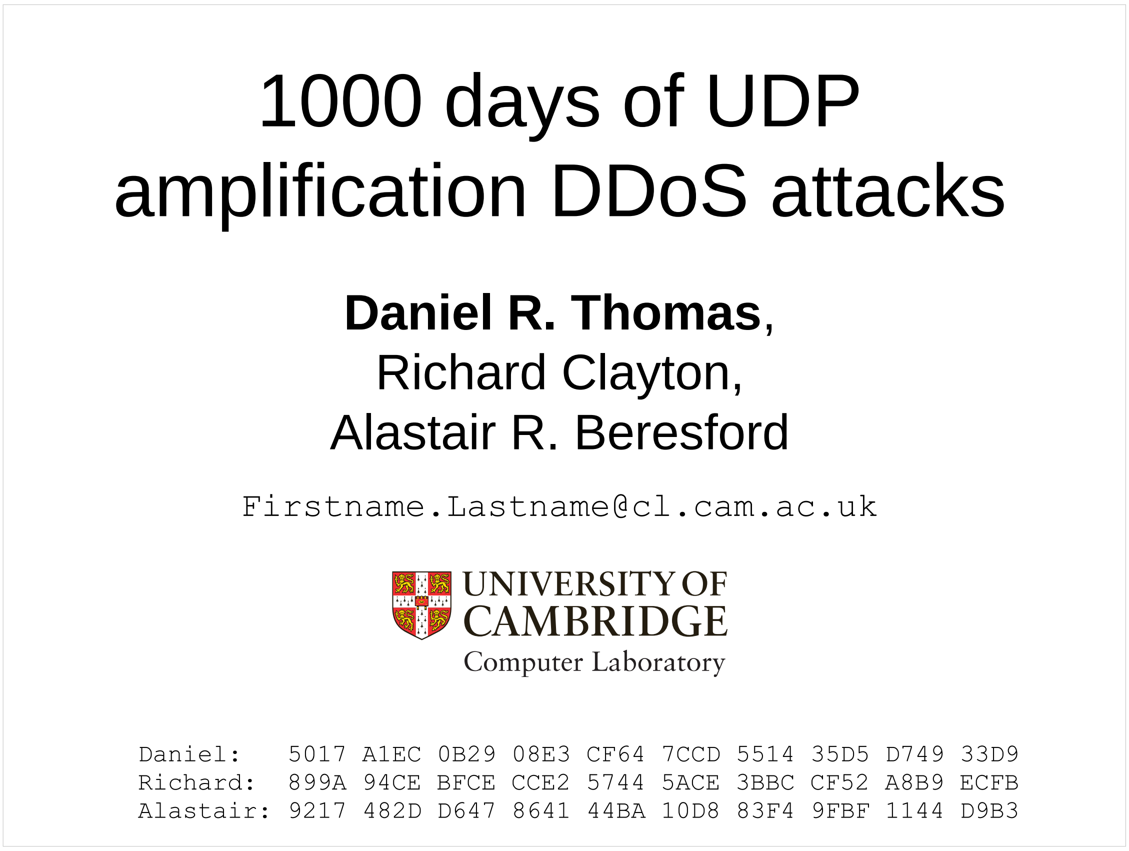

1000 days of UDP amplification DDoS attacks

Daniel R. Thomas,
Richard Clayton,
Alastair R. Beresford

Firstname.Lastname@cl.cam.ac.uk

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

```
Daniel:    5017 A1EC 0B29 08E3 CF64 7CCD 5514 35D5 D749 33D9
Richard:   899A 94CE BFCE CCE2 5744 5ACE 3BBC CF52 A8B9 ECFB
Alastair: 9217 482D D647 8641 44BA 10D8 83F4 9FBF 1144 D9B3
```

Hello, I am Daniel Thomas from the Cambridge Cybercrime Centre which is hosted by the Computer Laboratory.

We have been using honeypots to collect data on UDP amplification Distributed Denial of Service attacks for over 1000 days.

I will describe some of what we have learnt from this data

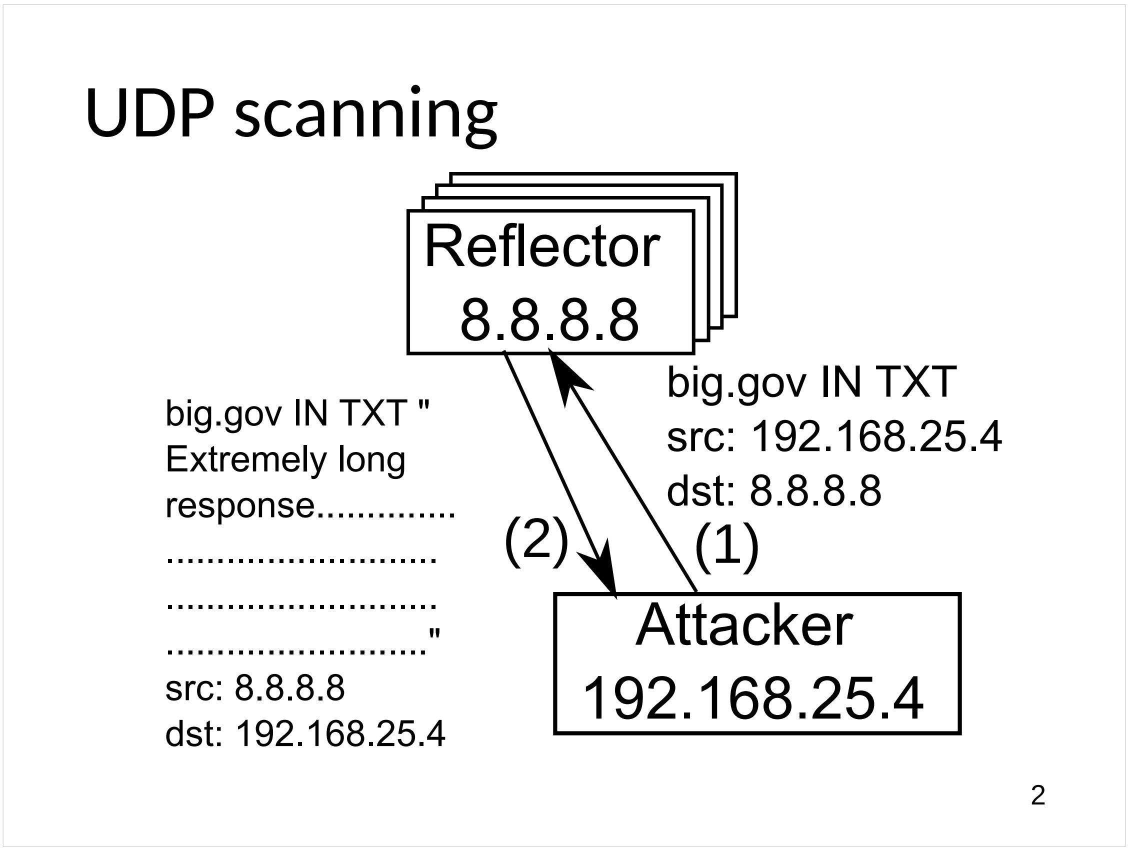

UDP scanning

Reflector
8.8.8.8

big.gov IN TXT
src: 192.168.25.4
dst: 8.8.8.8

(1)

(2)

big.gov IN TXT "
Extremely long
response..............
....
....
.........................."
src: 8.8.8.8
dst: 192.168.25.4

Attacker
192.168.25.4

- To conduct UDP amplification DDoS attacks the attacker first needs to find reflectors it can use to reflect off.
- To do this it uses UDP in a standard way, sending out UDP packets and collecting the responses.
- In this example it sends out a DNS packet, and when it finds a real reflector it gets a response back.

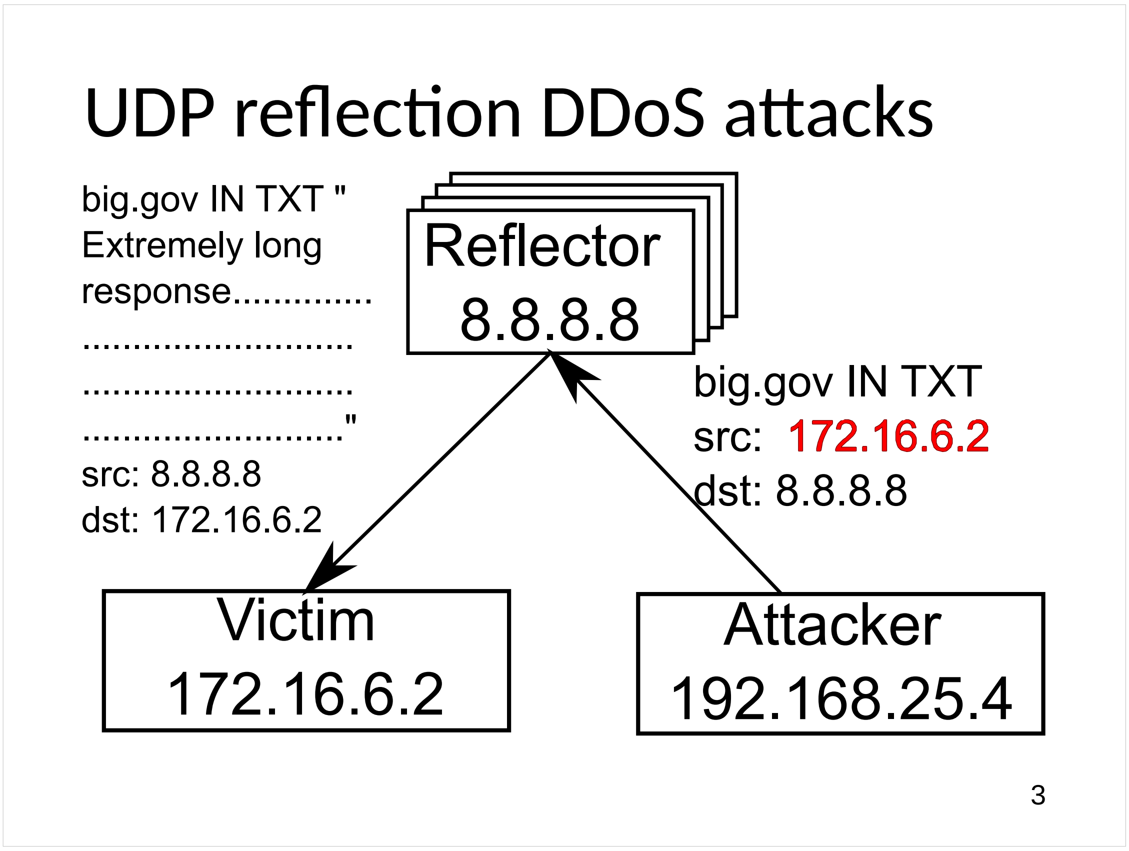

UDP reflection DDoS attacks

big.gov IN TXT "
Extremely long
response.............
....
....
.........................."
src: 8.8.8.8
dst: 172.16.6.2

Reflector
8.8.8.8

big.gov IN TXT
src: 172.16.6.2
dst: 8.8.8.8

Victim
172.16.6.2

Attacker
192.168.25.4

- UDP reflection DDoS attacks exploit the fact that UDP (unlike TCP) does not verify the source IP address with a 3 way handshake. Hence, if an attacker can spoof the source IP address on the packets they send then the response will go to their victim.
- In this example the attacker sends a DNS query to a resolver but spoofs the source IP address as the victim IP address. The much larger response goes to the victim.
- The attacker can repeat this many times and over thousands of resolvers. This results in a large volume of traffic to the victim. The victim does not know the address of the attacker.
- Most of the attacks using this method are from booters: DDoS as a service.

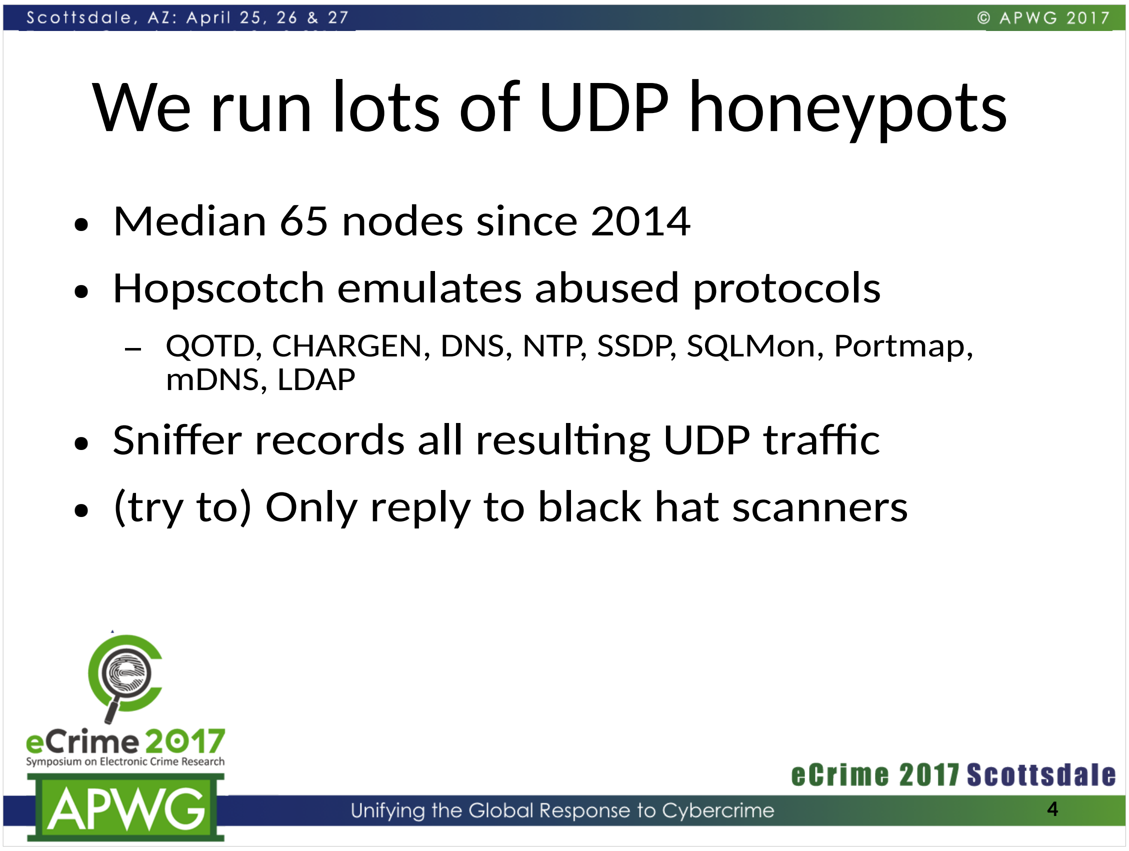

Scottsdale, AZ: April 25, 26 & 27

© APWG 2017

We run lots of UDP honeypots

- Median 65 nodes since 2014
- Hopscotch emulates abused protocols
  - QOTD, CHARGEN, DNS, NTP, SSDP, SQLMon, Portmap, mDNS, LDAP
- Sniffer records all resulting UDP traffic
- (try to) Only reply to black hat scanners

eCrime 2017 Symposium on Electronic Crime Research

APWG

eCrime 2017 Scottsdale

Unifying the Global Response to Cybercrime

Since March 2014 we have been running UDP honeypots.

A small program called hopscotch emulates UDP protocols that are abused in UDP reflection attacks.

Another small program called sniffer records UDP traffic.

Hopscotch aims to only reply to black hat scanners and so when it has seen more than a handful of packets from the same destination it stops responding. The nodes also collaborate to report victims.

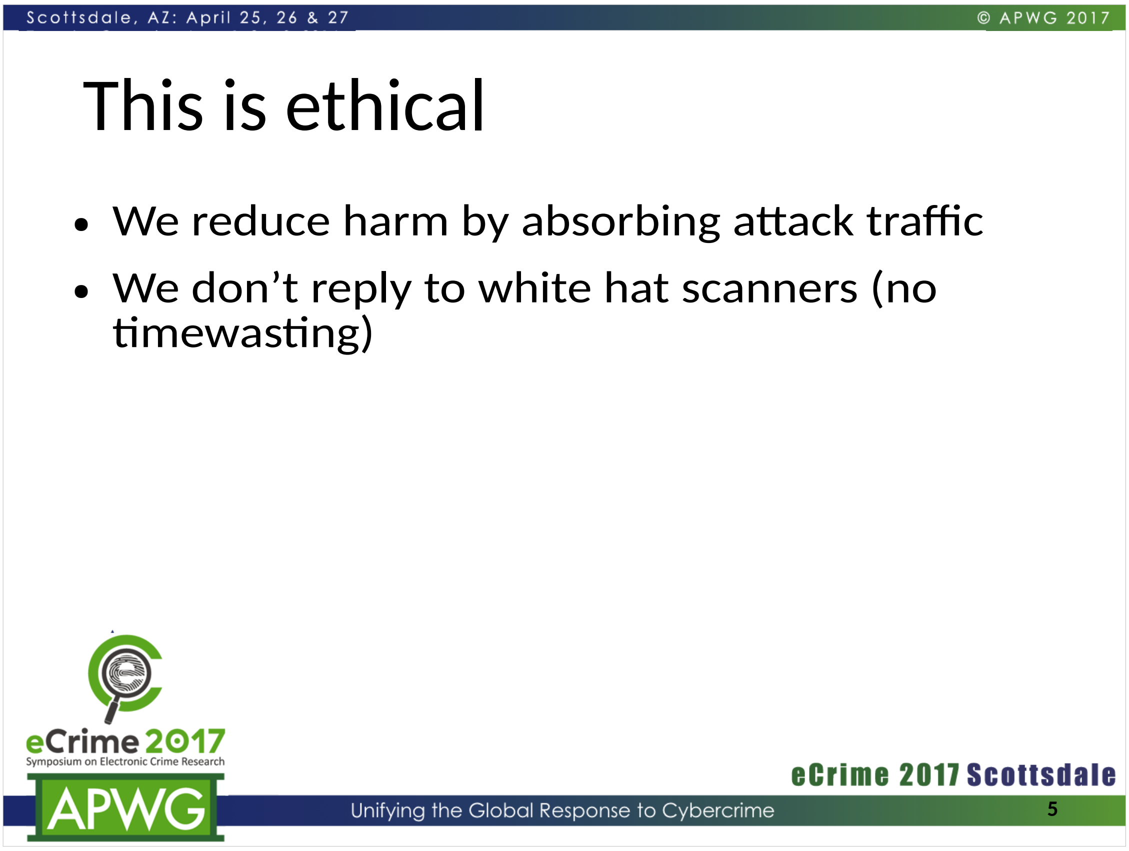

This is ethical

- We reduce harm by absorbing attack traffic
- We don’t reply to white hat scanners (no timewasting)

eCrime 2017 Symposium on Electronic Crime Research

APWG

We followed our institutions ethical procedure.

Running these honeypots reduces harm as when an attacker uses our honeypots to attack their victim their victim will receive rather less traffic than they would have if the attacker had used one of the many real reflectors.

To avoid wasting white hat’s time we never reply to their scanners so they don’t report us as being reflectors.

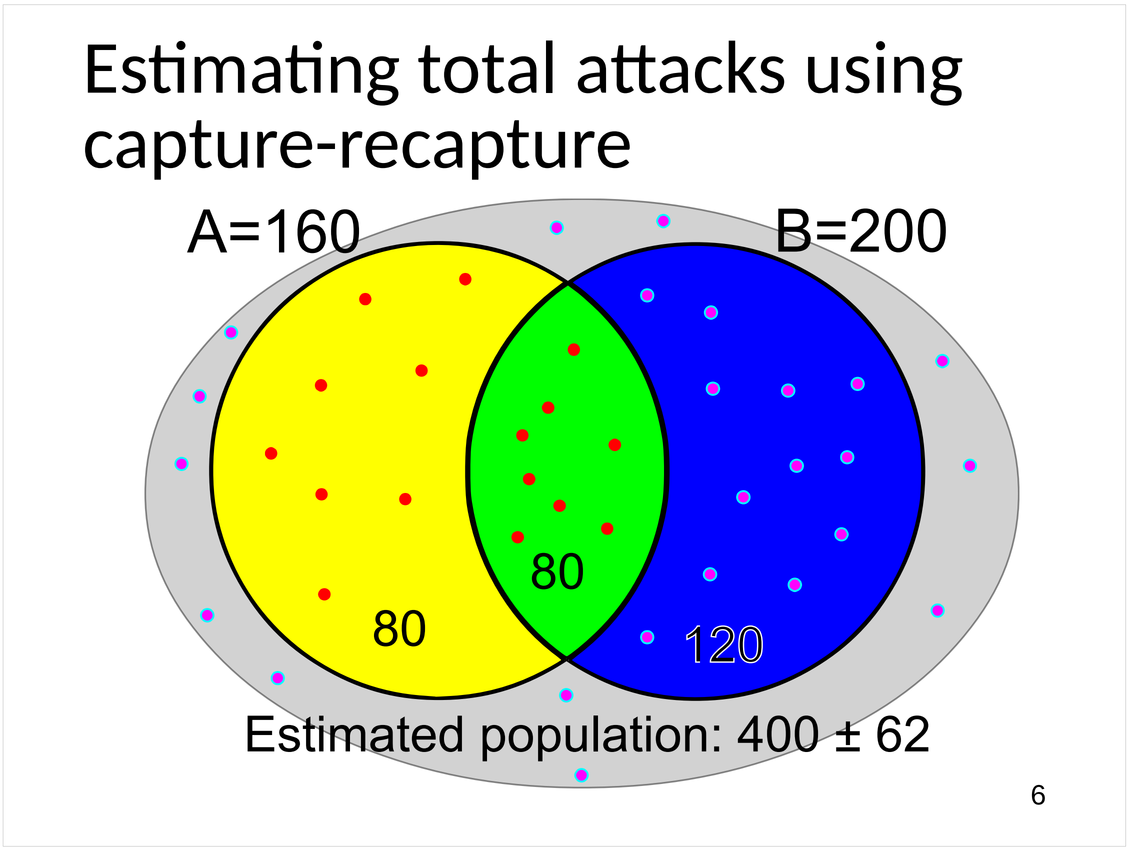

Estimating total attacks using capture-recapture

Estimated population: 400 ± 62

With these sensors we can see some attacks, but we want to know how many attacks there were, including the attacks we did not observe.

We can do this using the capture-recapture technique originally developed for ecology.

On day A we go fishing in a lake and catch 16 fish, mark them and return them to the lake, on day B we go fishing and catch 20 fish, of which 8 were marked as being previously caught. From this we can estimate that there are 40 fish in the lake.

We can then use this to estimate the total number of UDP attacks. We can split our sensors into two groups, A and B and look at the number of attacks that each detected and the size of the overlap.

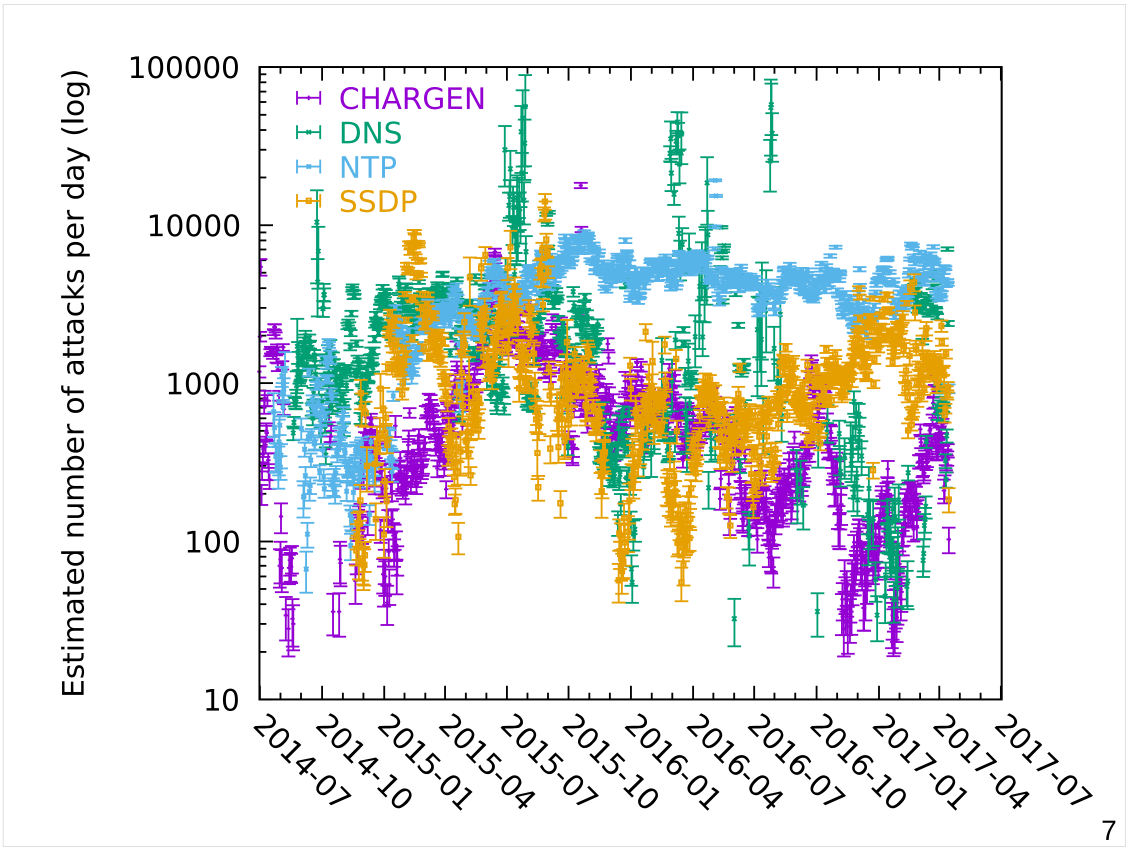

This graph shows the estimated total number of attacks per day for the four most used protocols.

It shows substantial changes in the number of attacks being made with each protocol over time.

Protocols go in and out of fashion.

SSDP is becoming more fashionable again after a period when it was much less widely used.

NTP has remained consistently popular and DNS has varied a lot.

There was a paper that examined data from before the start of our measurement period and concluded that NTP was declining in popularity. Our longitudinal study shows that protocols go in and out of fashion. Just because it stops being used so much doesn’t mean it won’t come back.

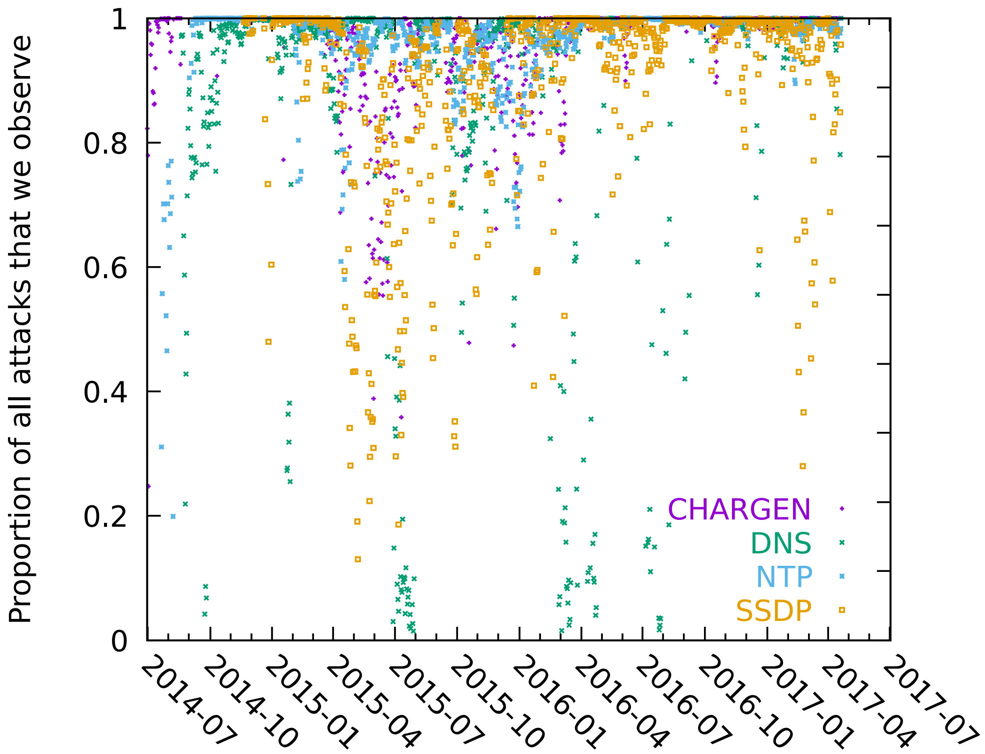

This graph shows the proportion of the estimated total number of attacks that we observe each day. In general we have very good coverage, seeing almost all attacks. However, on some days we do rather worse, particularly for DNS and SSDP.

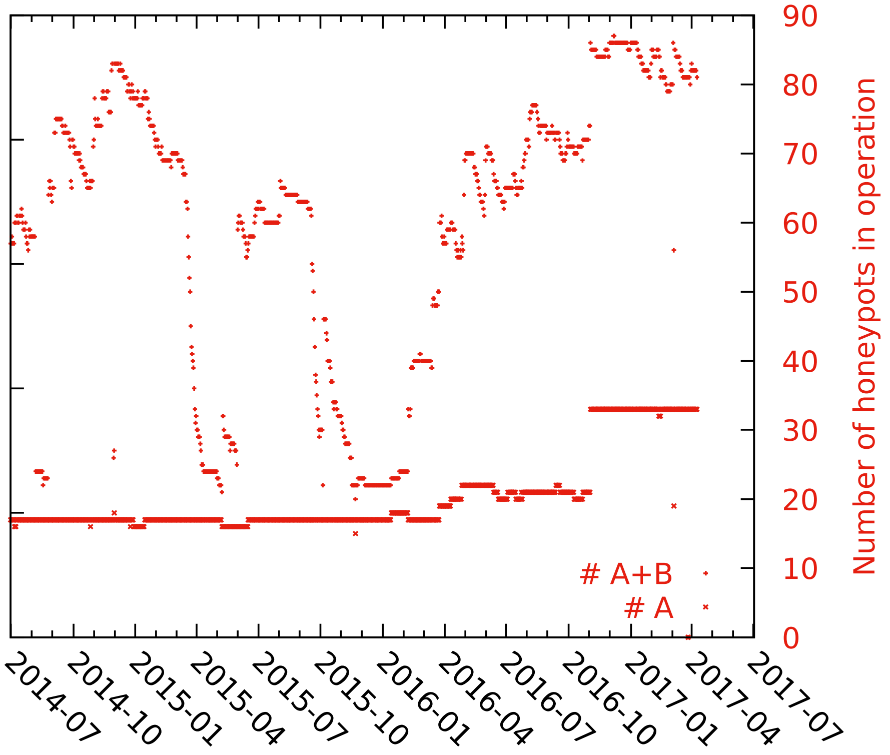

This also shows both the total number of honeypots we had in operation and the number in the A set used for capture-recapture. It varies over time as a result of our main contributor ceasing to share data with us and our rebuilding our own network of sensors.

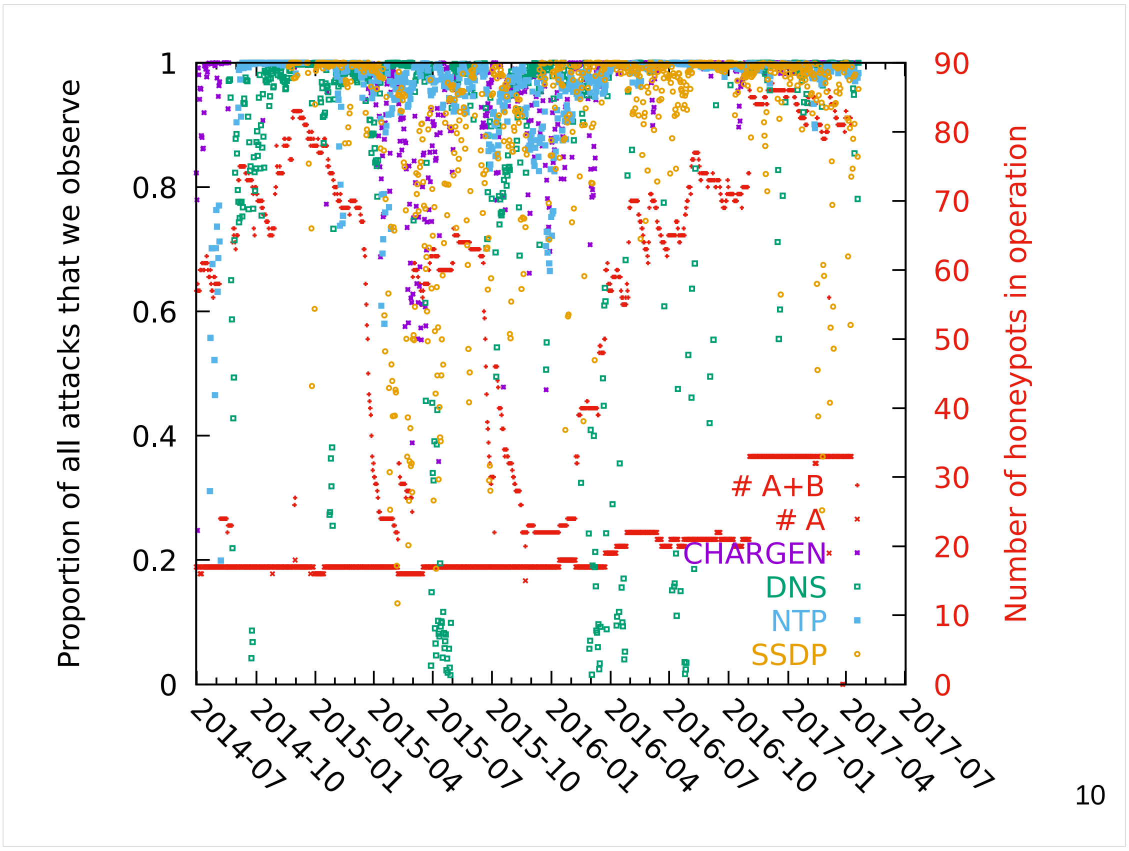

As you might expect there is correlation between the number of honeypots in operation and the proportion of attacks that we observe.

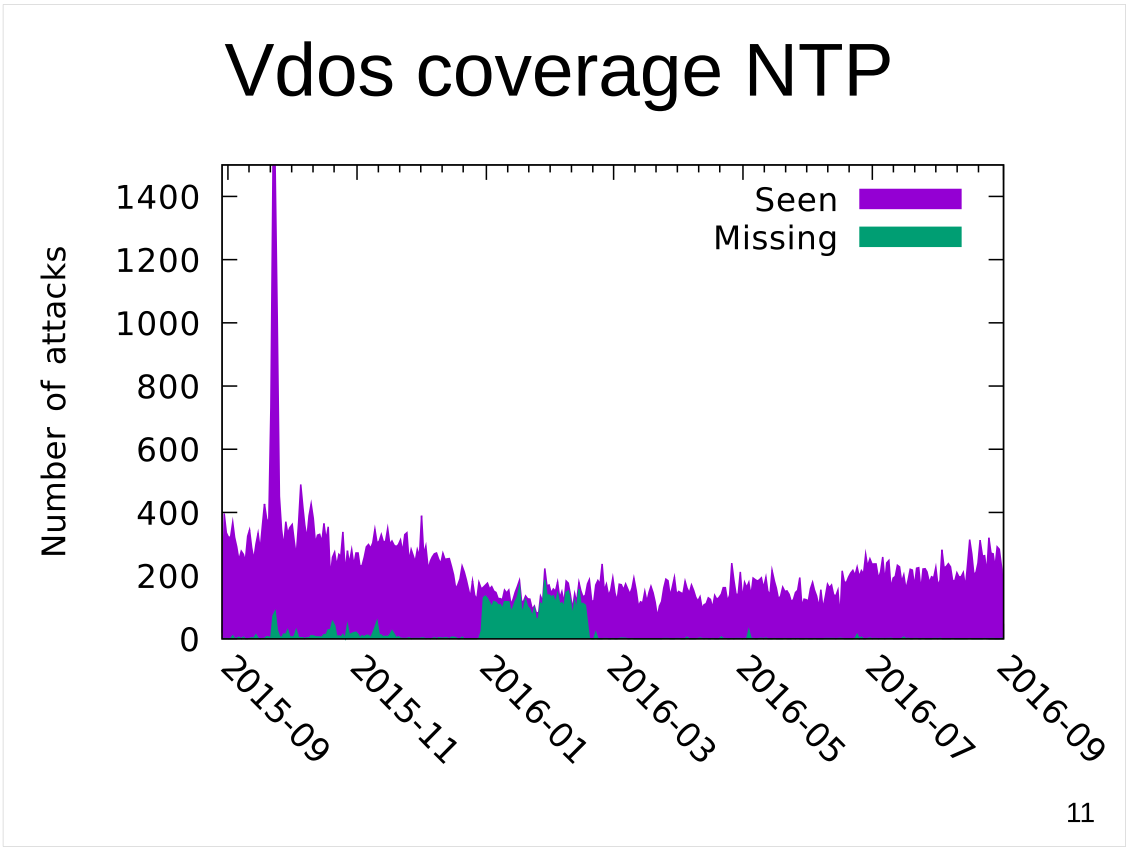

Vdos coverage NTP

In the paper we check to see what proportion of attacks that different booters are recorded as making appear in our database.

Mostly we see the same proportion of attacks for a particular protocol on a particular booter, however for the Vdos booter that was recently shut down we observe changes in behaviour over time.

This is a stacked plot.

For NTP we mostly saw almost all attacks but as this graph shows there was a period in early 2016 when we saw almost none of the attacks.

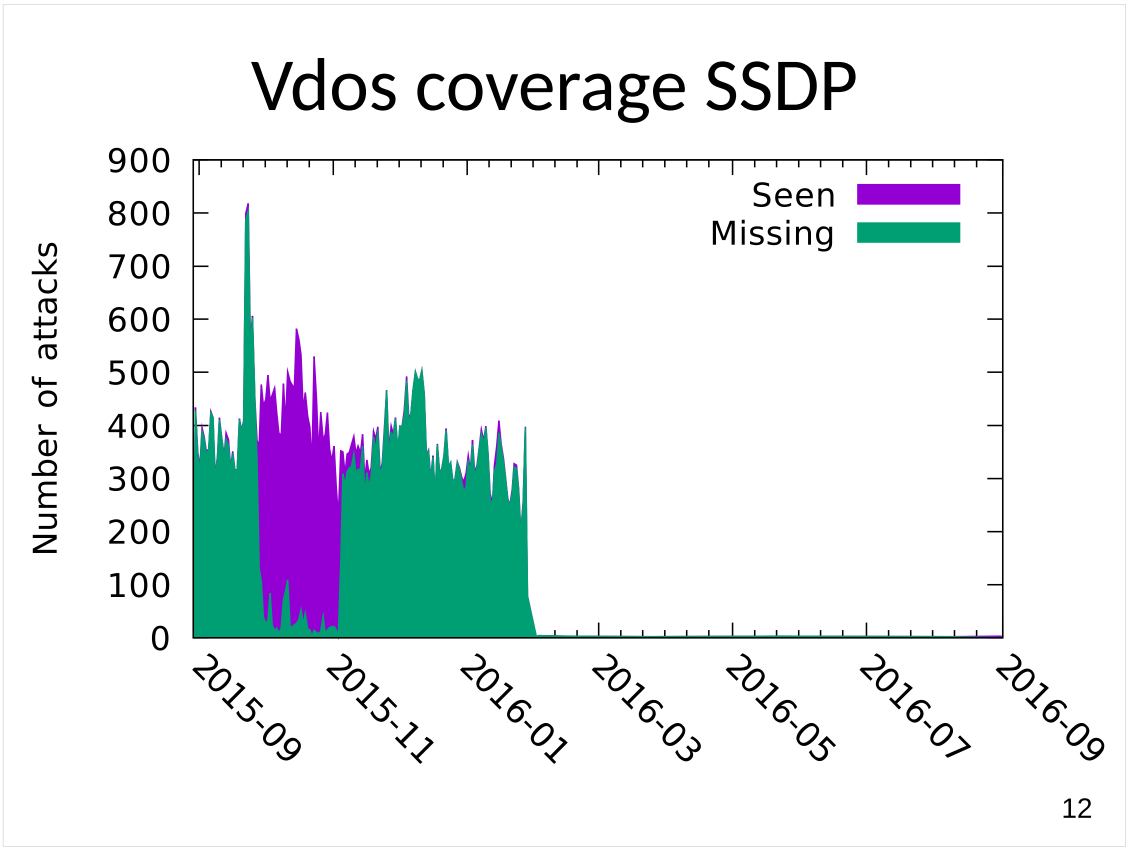

Vdos coverage SSDP

For SSDP we see the reverse, mostly we see none of the attacks but for a period in 2015 we see almost all attacks. In 2016 Vdos stopped using SSDP.

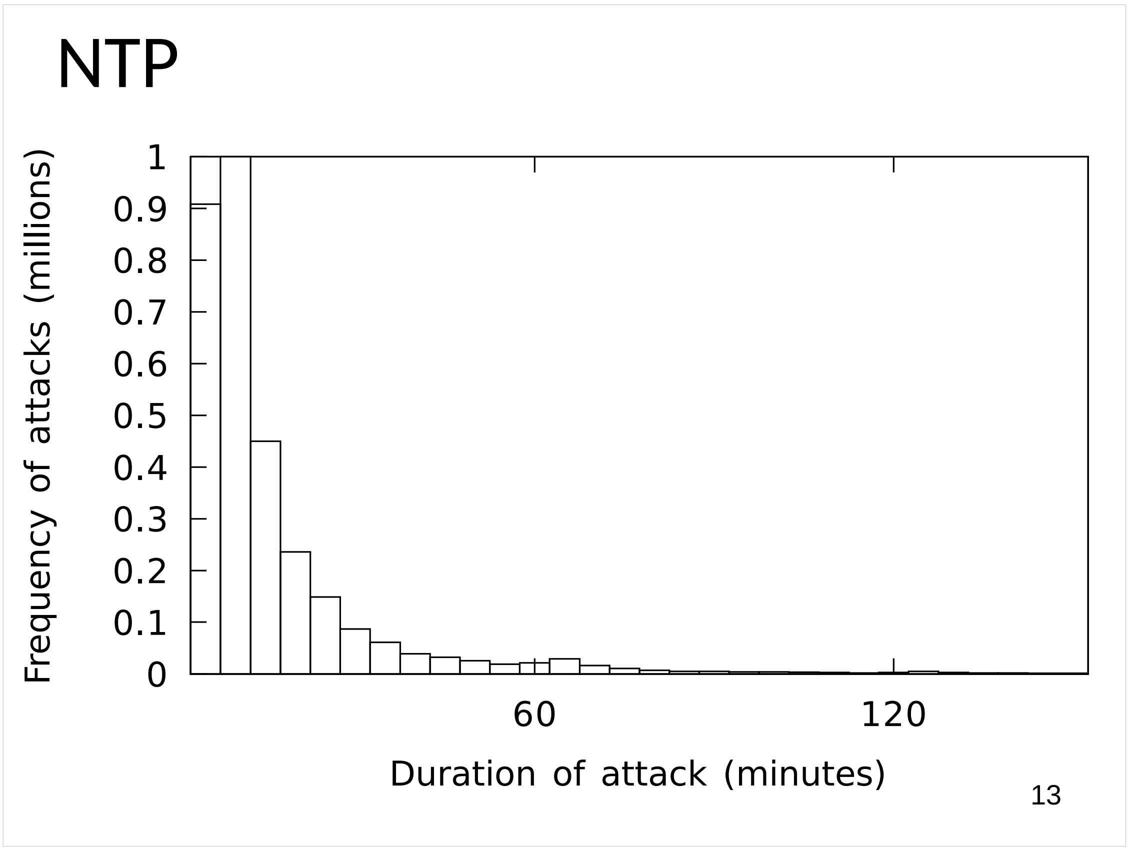

NTP

Different attacks take different lengths of time.

This graph shows the frequency of different attack lengths, clearly most attacks are short.

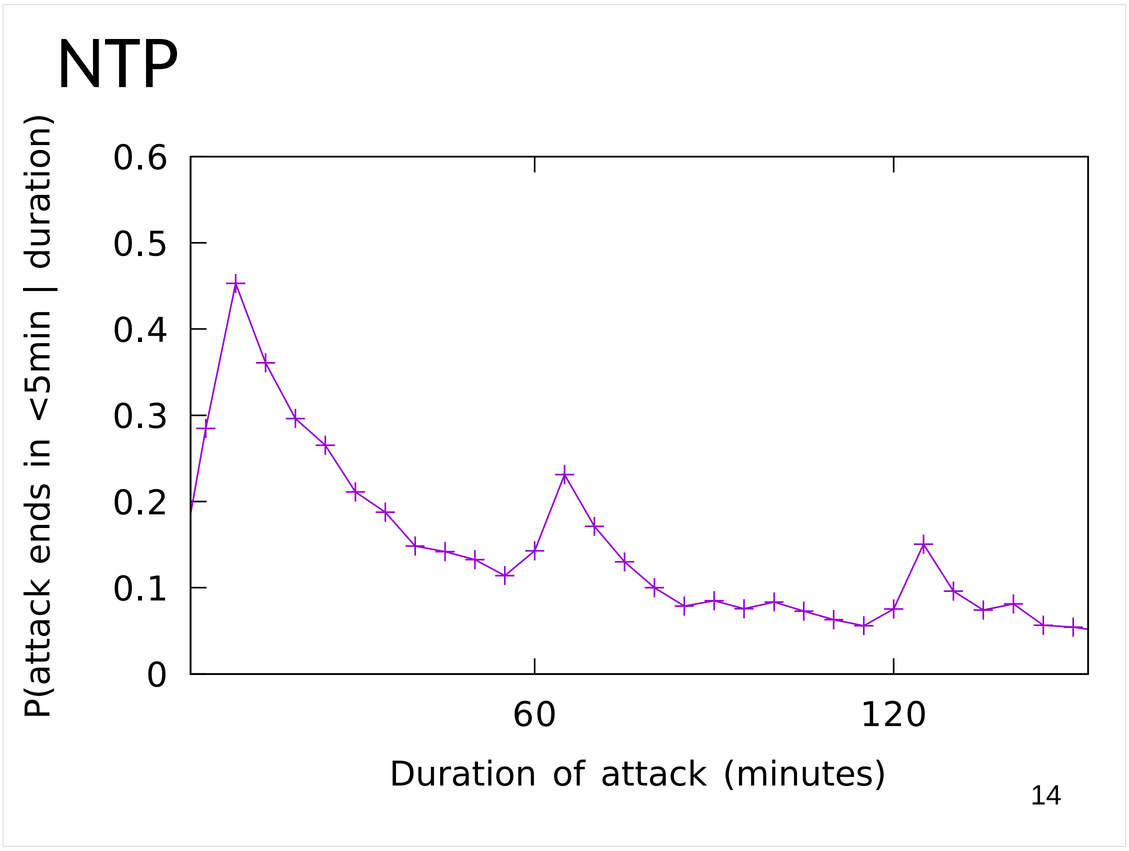

NTP

If site operators know that they are being DDoSed they might like to know how long the attack is likely to last. Hence we computed the death curve.

The death curve for attacks shows, given an attack has lasted for a certain time, what is the probability it will stop in the next 5 minutes. This shows the decreasing probability with length, the longer an attack has gone on the less likely it is stop soon. However, there are also peaks at 60 and 120 minutes, caused by these being default attack lengths on booters.

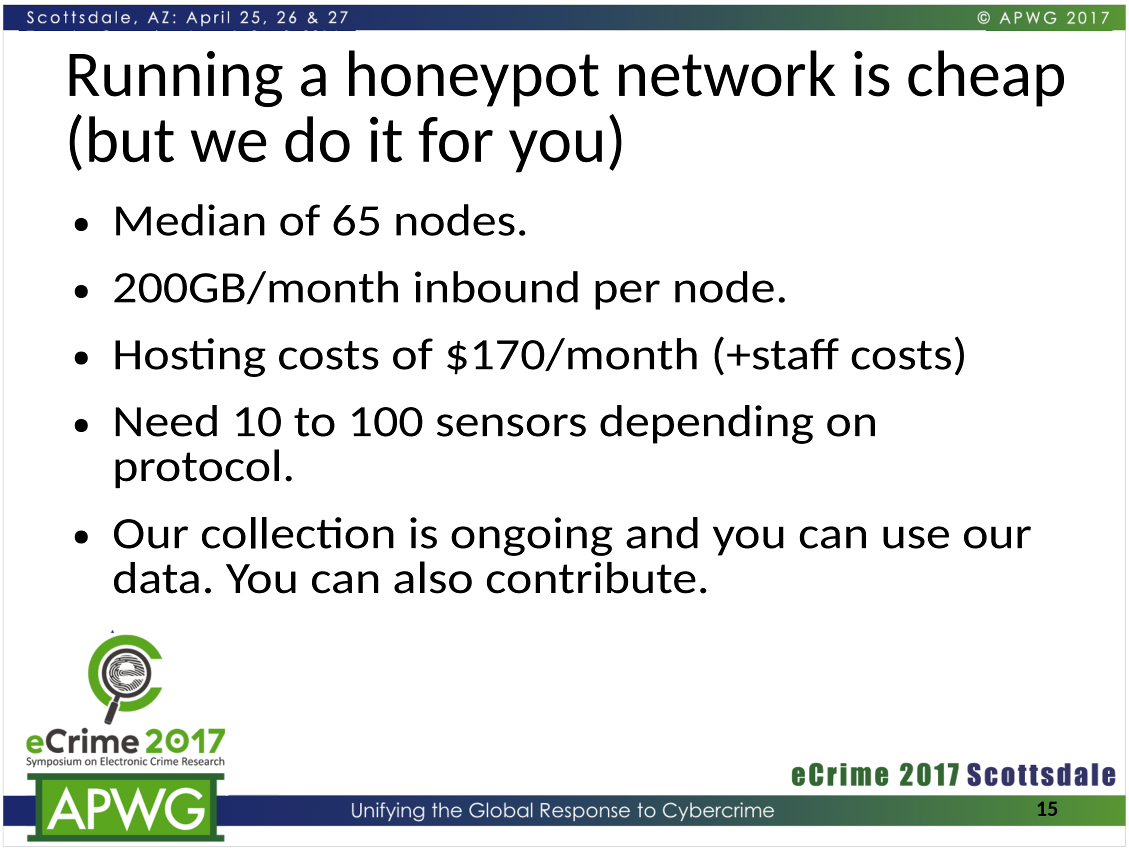

Scottsdale, AZ: April 25, 26 & 27

© APWG 2017

Running a honeypot network is cheap (but we do it for you)

- Median of 65 nodes.
- 200GB/month inbound per node.
- Hosting costs of $170/month (+staff costs)
- Need 10 to 100 sensors depending on protocol.
- Our collection is ongoing and you can use our data. You can also contribute.

eCrime 2017

Symposium on Electronic Crime Research

APWG

eCrime 2017 Scottsdale

Unifying the Global Response to Cybercrime

15

- 10 QOTD, 10 LDAP, 32 MDNS, 100 DNS, 45 Portmap, 32 SQLMon, 55 CHARGEN, 80 SSDP, and 70 NTP sensors.
- If you want to examine the data yourself then please get in touch.
- If you have IP addresses and want to host some nodes for us then we would love to hear from you.

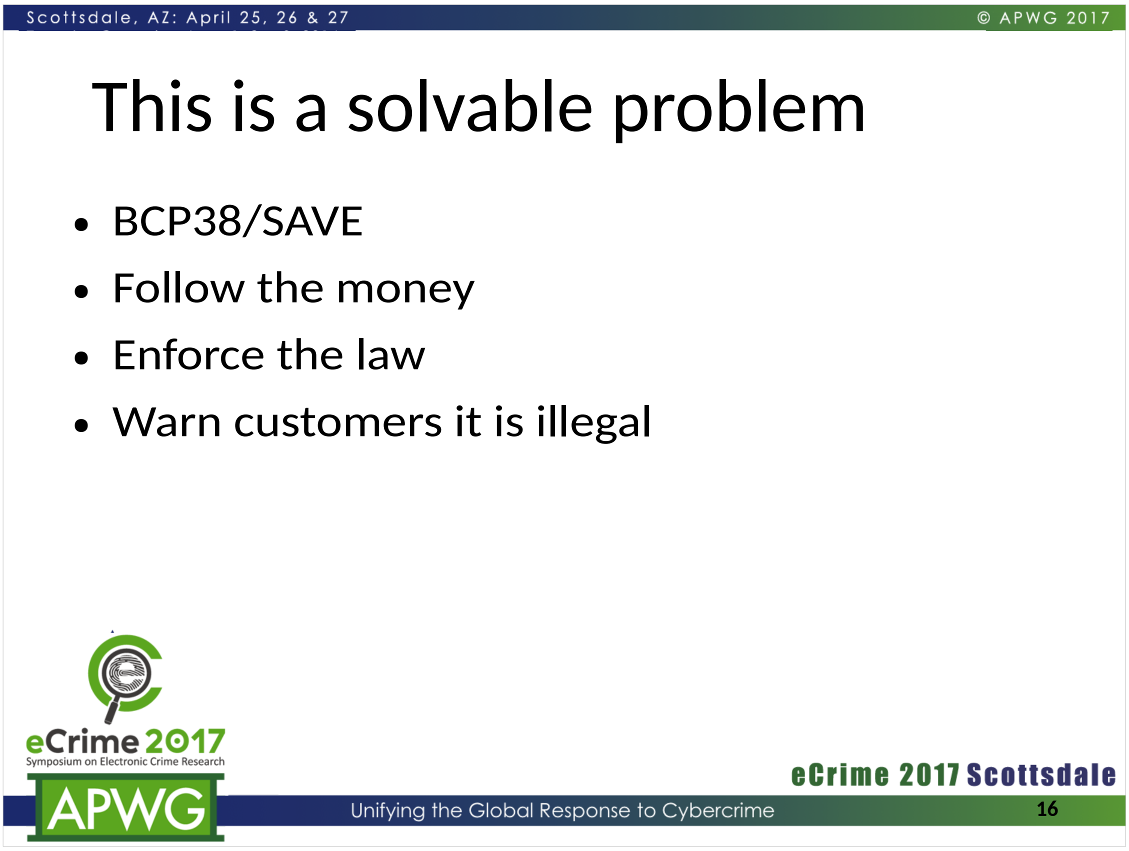

Scottsdale, AZ: April 25, 26 & 27 © APWG 2017

This is a solvable problem

- BCP38/SAVE
- Follow the money
- Enforce the law
- Warn customers it is illegal

eCrime 2017 Symposium on Electronic Crime Research

APWG

eCrime 2017 Scottsdale

Unifying the Global Response to Cybercrime

CAIDA’s spoofer prober project measures compliance with BCP38.

Paypal has made a big impact on booter revenue.

Lots of arrests have been made.

Booter users don’t all realise fully that what they are doing is illegal.

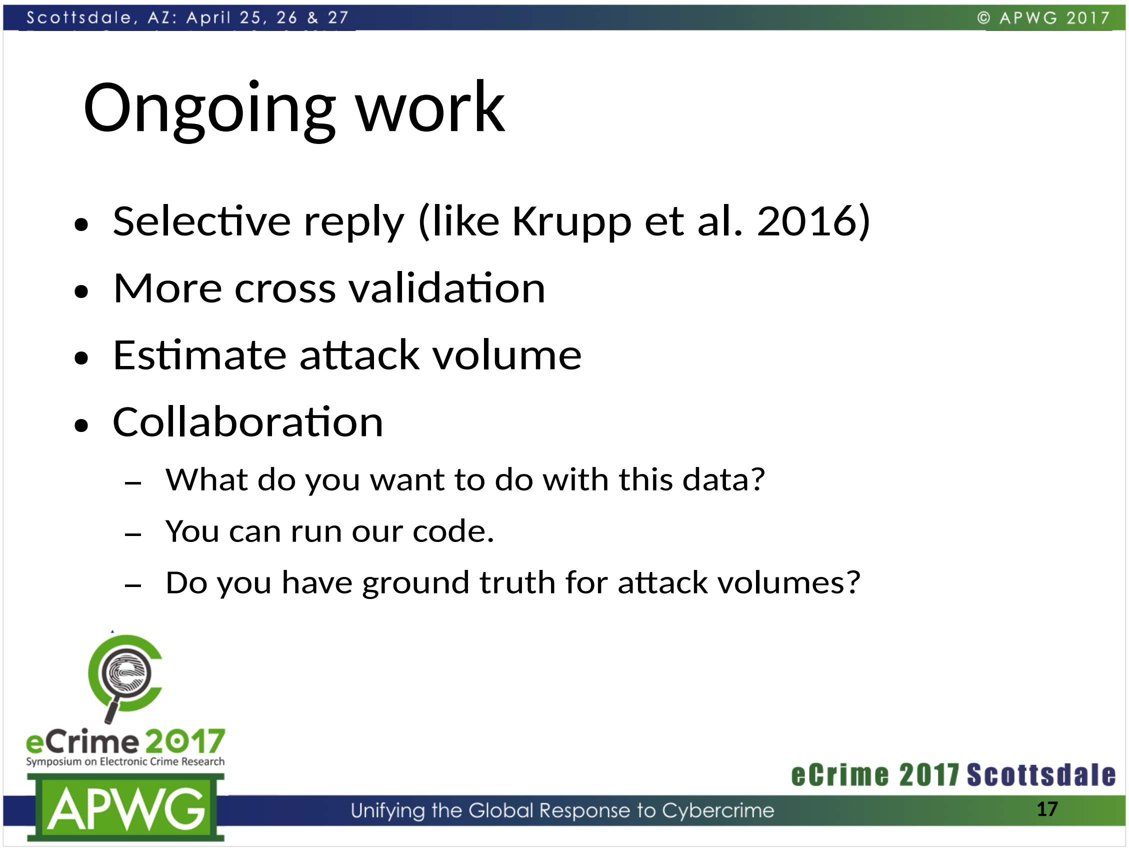

Ongoing work

- Selective reply (like Krupp et al. 2016)
- More cross validation
- Estimate attack volume
- Collaboration
  - What do you want to do with this data?
  - You can run our code.
  - Do you have ground truth for attack volumes?

We can alter whether we reply dynamically based on the source IP address this can then be used to tie subsequent attempted attacks back to the scanner that found the honeypot. This then allows attacks using the same scanner data to be tied together.

We want to do more cross validation with other data to check how representative our results are.

You can use our data. What do you want to do with it?

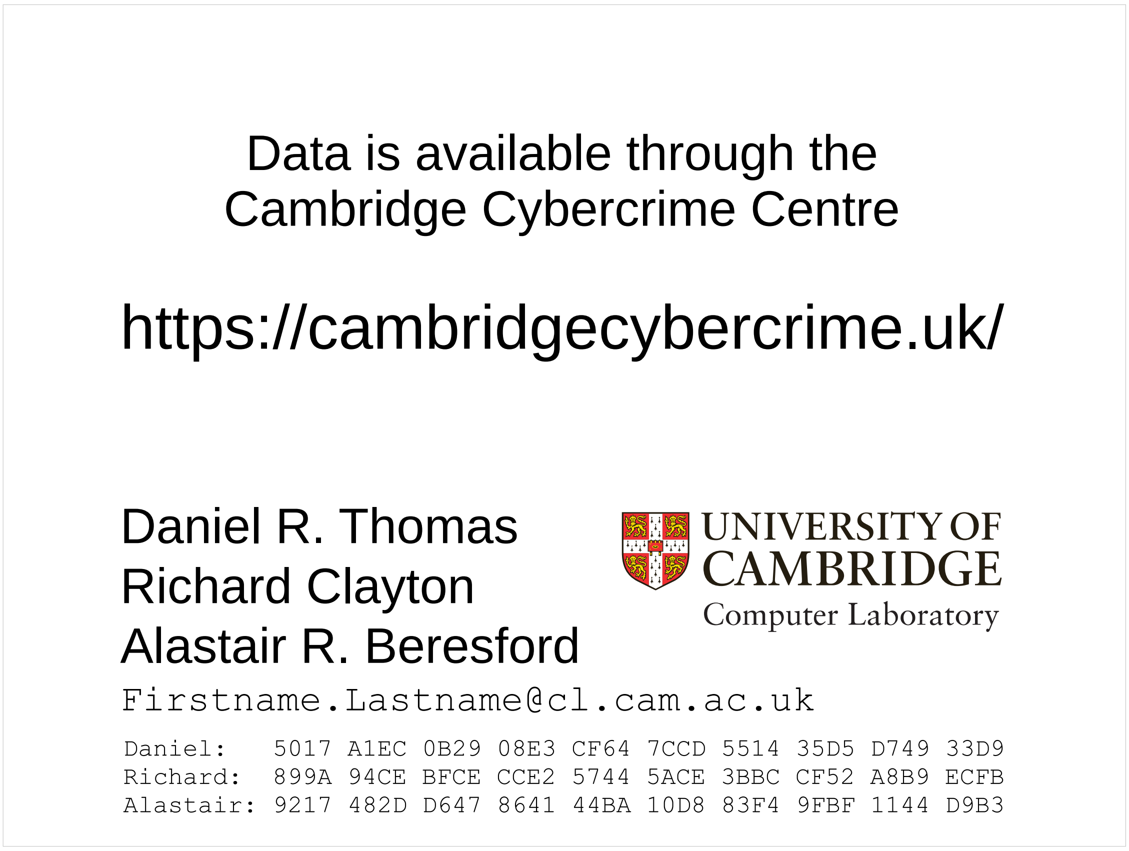

Data is available through the
Cambridge Cybercrime Centre

https://cambridgecybercrime.uk/

Daniel R. Thomas
Richard Clayton
Alastair R. Beresford

UNIVERSITY OF CAMBRIDGE
Computer Laboratory

Firstname.Lastname@cl.cam.ac.uk

```
Daniel:    5017 A1EC 0B29 08E3 CF64 7CCD 5514 35D5 D749 33D9
Richard:   899A 94CE BFCE CCE2 5744 5ACE 3BBC CF52 A8B9 ECFB
Alastair:  9217 482D D647 8641 44BA 10D8 83F4 9FBF 1144 D9B3
```

Thank you.
Questions?