A First Joint Look at DoS Attacks and BGP Blackholing in the Wild

Mattijs Jonker
Aiko Pras (UTwente)
Alberto Dainotti (CAIDA / UC San Diego)
Anna Sperotto (UTwente)

UNIVERSITY OF TWENTE.

caida

Denial-of-Service attacks

- A conceptually simple, yet effective class of attacks

  … that have gained a lot in popularity over the last years

  … are also offered "as-a-Service" (Booters)

- Some well-known incidents stipulate threat/risks
  - e.g., attacks on Dyn & GitHub (memcached)


New world record DDoS attack hits 1.7Tbps days after landmark GitHub outage

Memcached denial-of-service attacks are getting bigger by the day, according to new analysis.

By Liam Tung | March 6, 2018 -- 12:34 GMT (04:34 PST) | Topic: Security

- DoS has become one of the **biggest** threats to Internet **stability & reliability**

BGP blackholing

- Is a technique that can be used to mitigate DoS attacks
- Leverages the BGP control plane to drop network traffic
- BGP communities are used to signal blackholing requests
  - by “tagging” prefix announcements with <asn:*value*>
  - 666 is is a common *value* for blackholing
- Is very “coarse-grained”, meaning all network traffic destined to a prefix is indiscriminately dropped

A missing piece of the puzzle

Given its coarse-grained nature, we wonder if blackholing is used only in extreme cases

A clear understanding of how blackholing is used in practice when DoS attacks occur is missing

We use large-scale, longitudinal (3y) data sets on DoS attacks and blackholing to get more insights into operational practices

Part 1: Blackholed Attacks

UCSD Network Telescope [data set 1/3]

- A large, **/8** network telescope operated by UC San Diego
- Captures backscatter from DoS activity in which source IP addreses are *randomly and uniformly spoofed*
- We use the classification methodology by Moore et al. to infer DoS attacks [1]

[1] Moore et al.,“Inferring Internet Denial-of-service Activity”, in ACM TOCS 2006

Amplification Honeypots [data set 2/3]

- Honeypots

  ... mimick reflectors abused in *reflection* attacks (e.g., NTP)

  … try to be appealing to attackers by offering large amplification

  … capture attempts at reflection

- We use logs from 24 honeypot instances that are geographically & logically distributed
  - From the AmpPot project (Christian Rossow, CISPA) [1]

[1] Krämer al.,"AmpPot: Monitoring and Defending Against Amplification DDoS Attacks", in RAID 2015

Inferred blackholing events [data set 3/3]

- Scan BGP collector data for blackholing activity, using public BGP data: RIPE RIS and UO Route Views
- Use BGPStream framework for BGP data analysis [1]
- Match BGP updates against dictionary of known BH communities [2]

[1] Orsini et al., "BGPStream: A Software Framework for Live and Historical BGP Data Analysis", in IMC 2016

[2] Giotsas et al., "Inferring BGP blackholing activity in the internet", in IMC 2017

BGP STREAM

Measurement systems placement

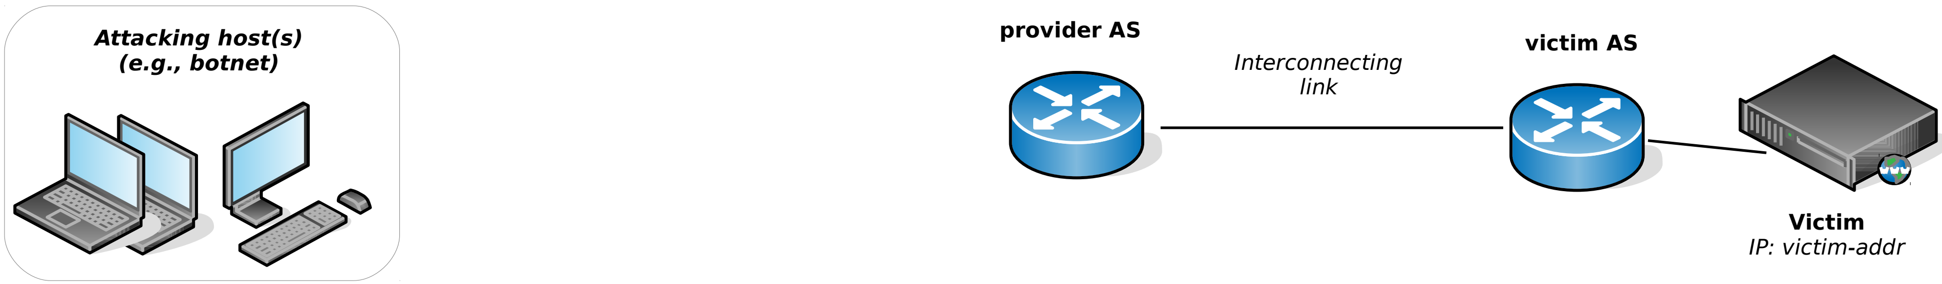

Measurement systems placement

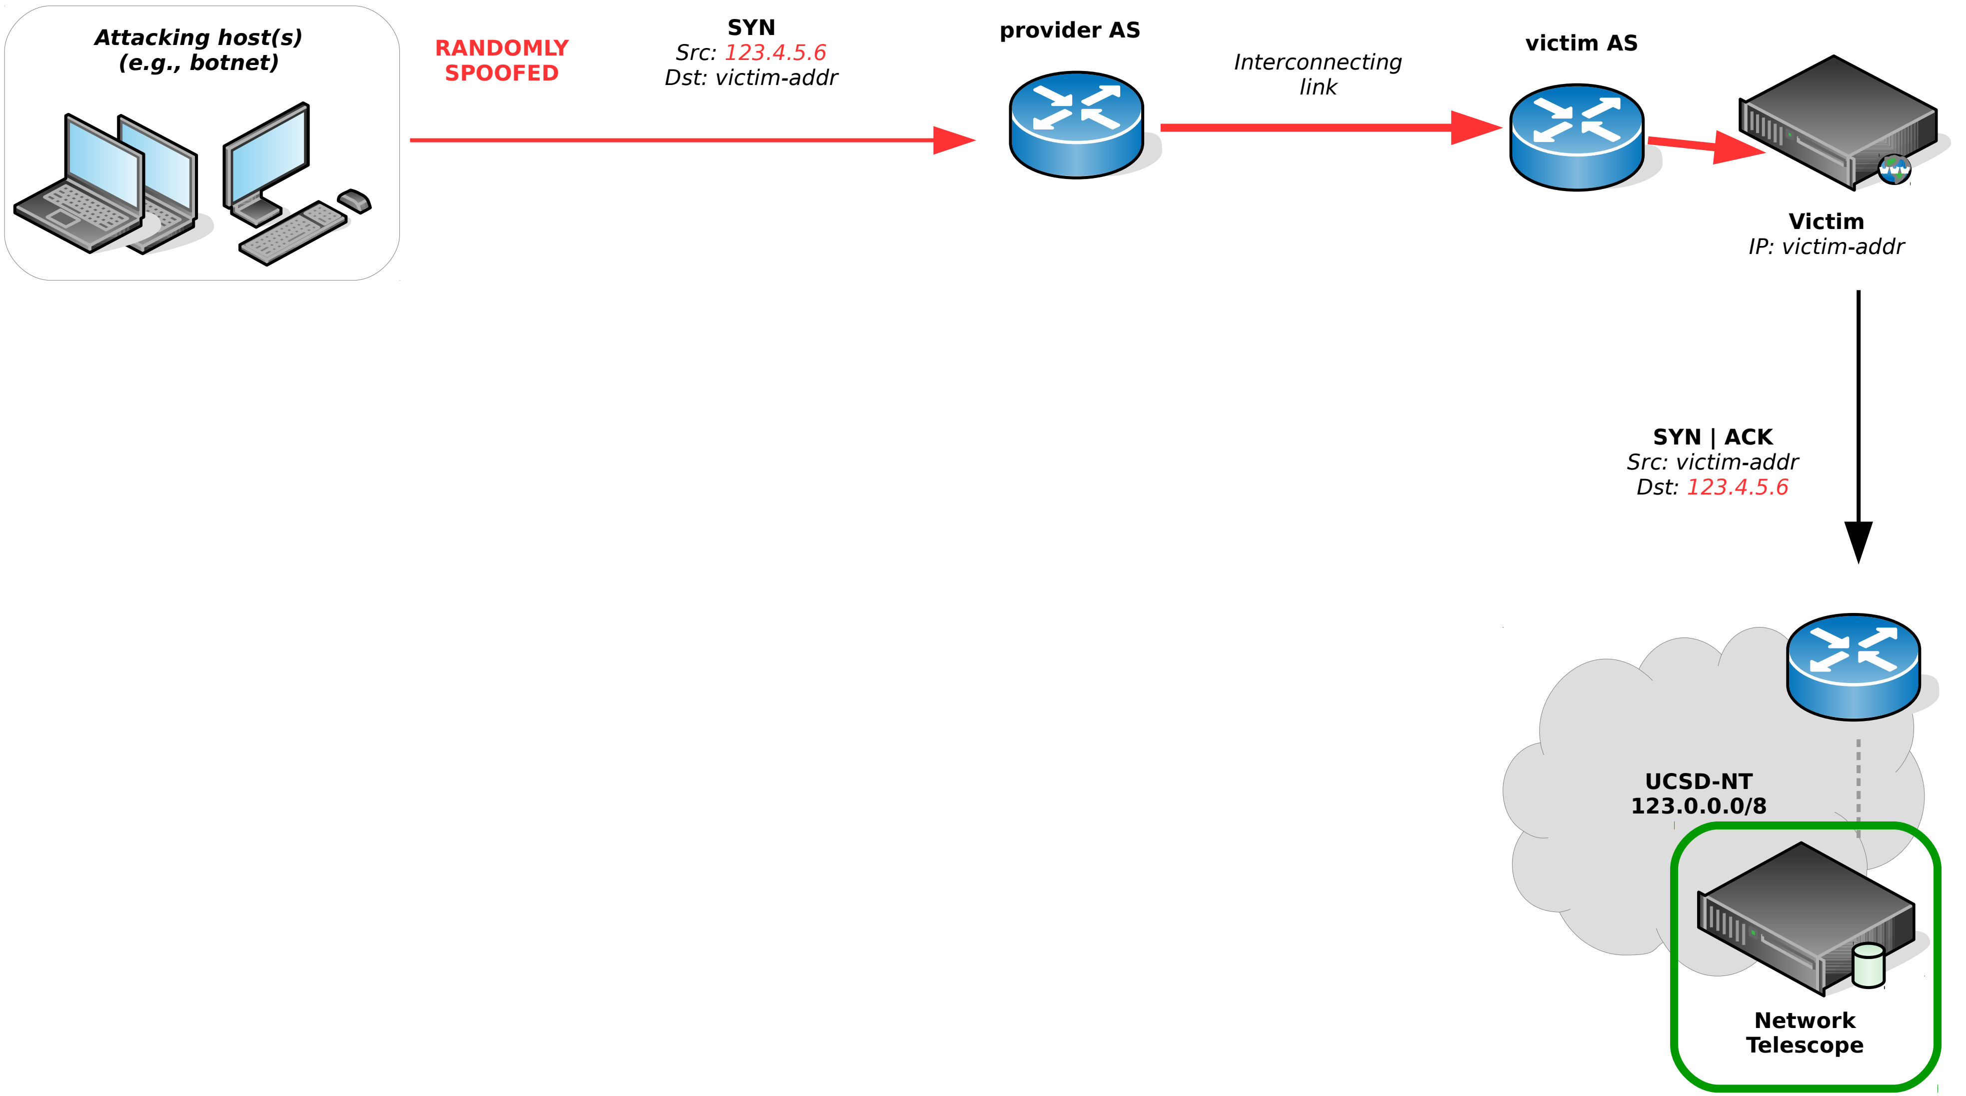

Measurement systems placement

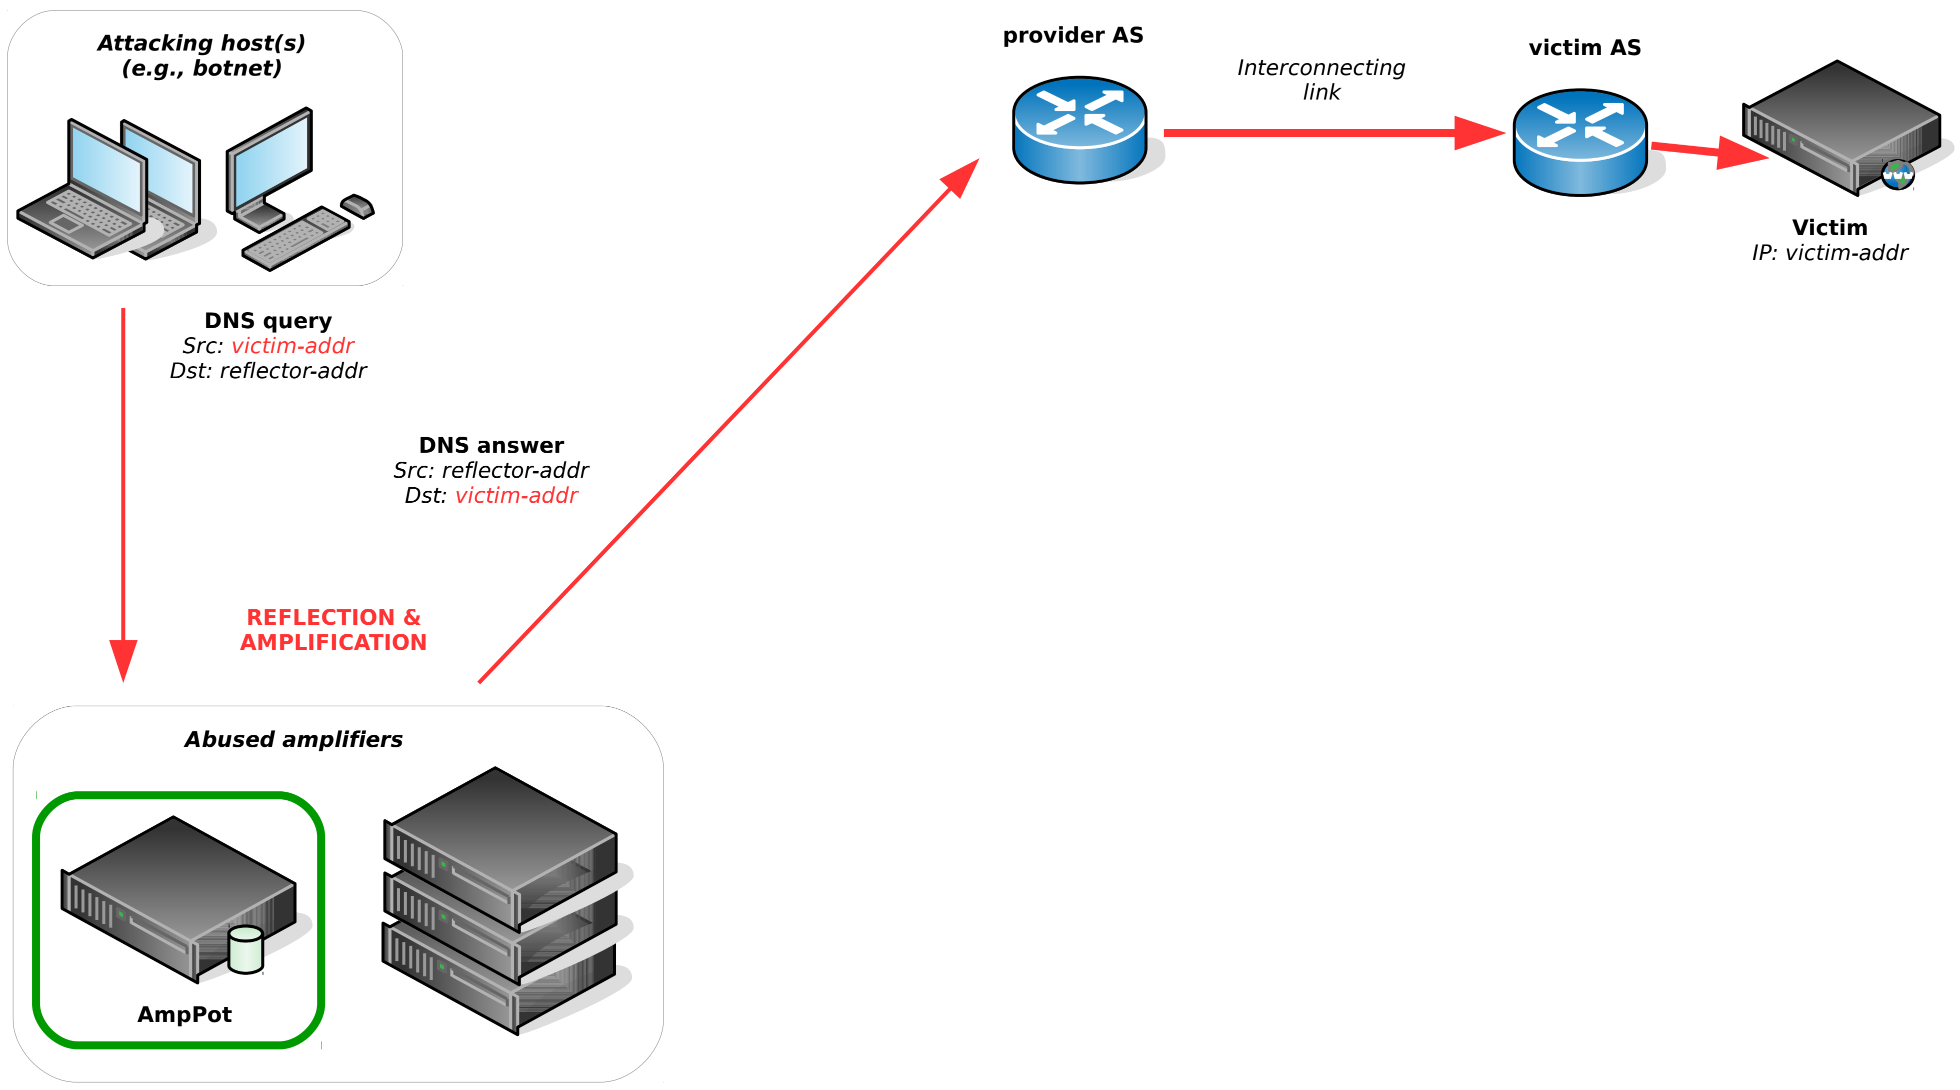

Measurement systems placement

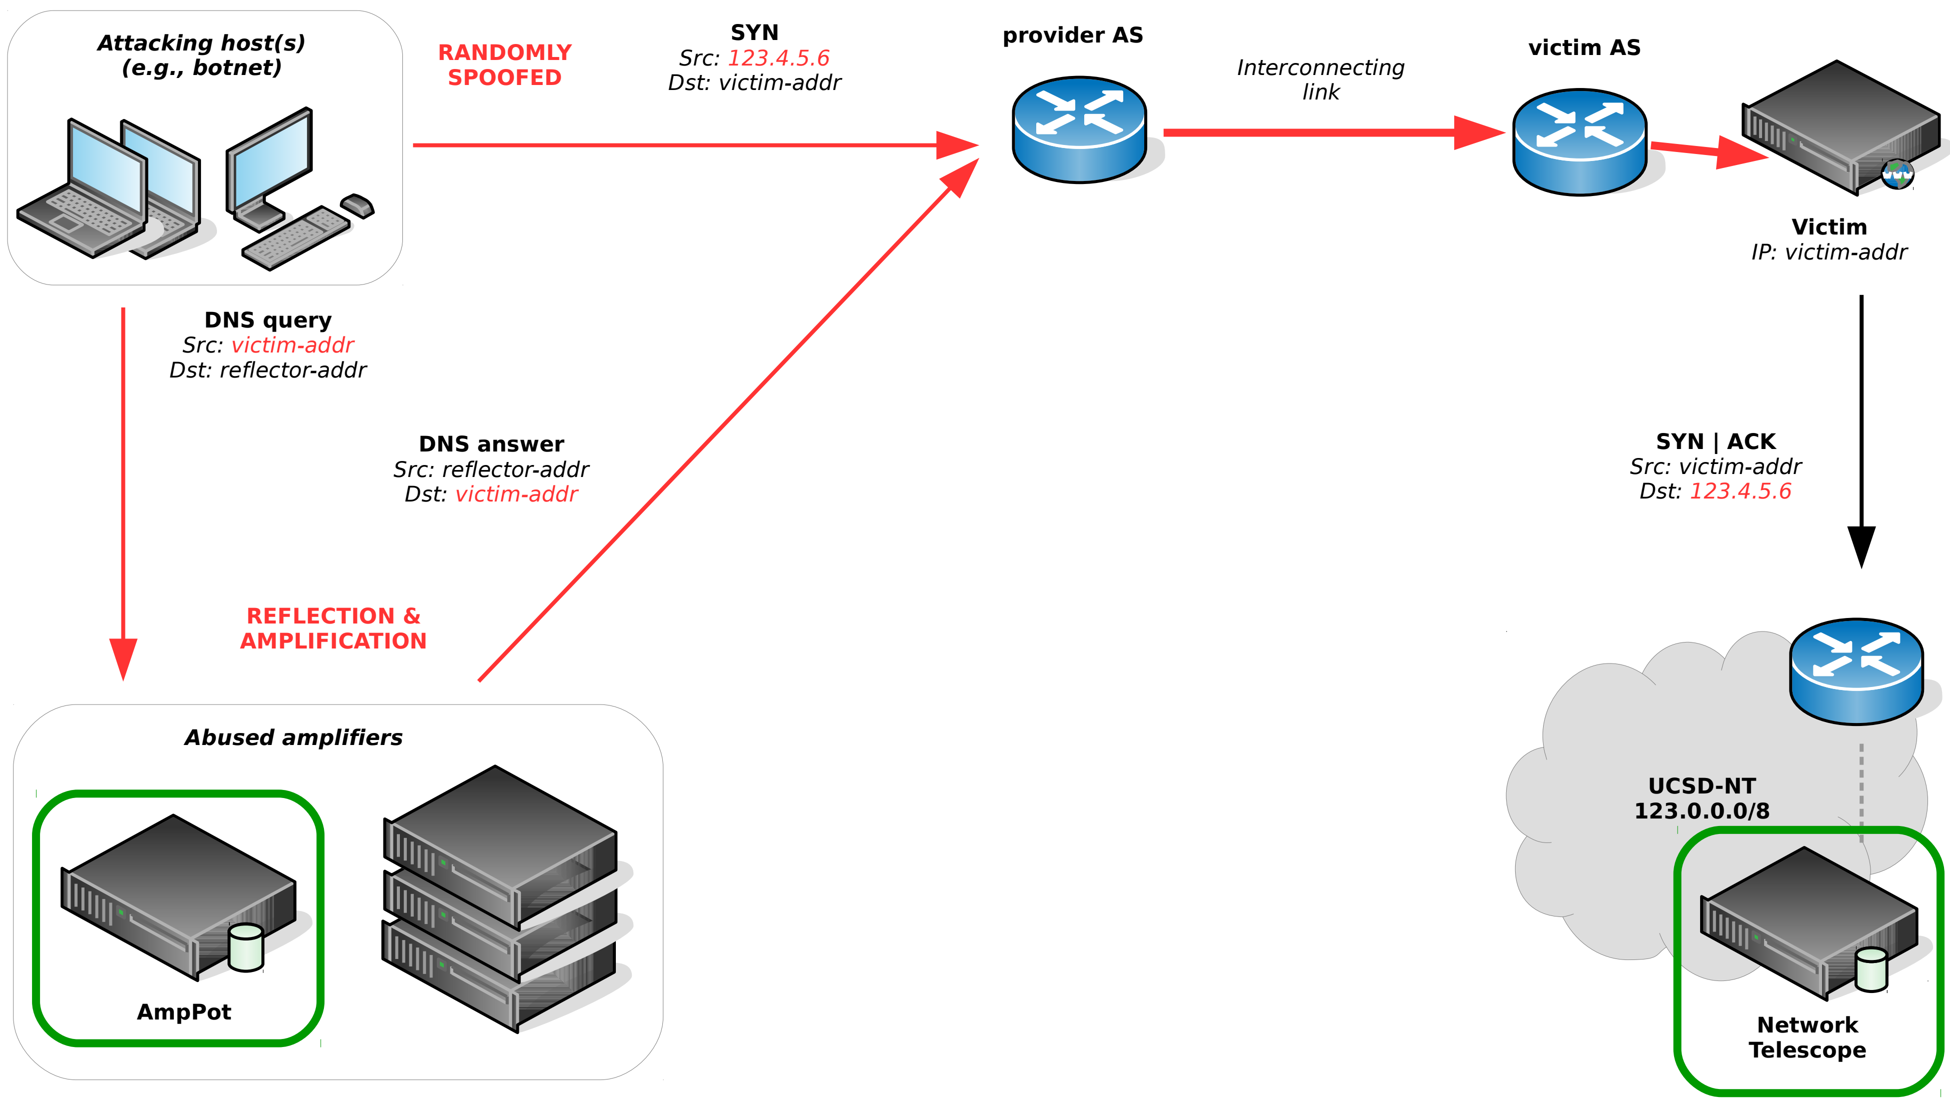

Measurement systems placement

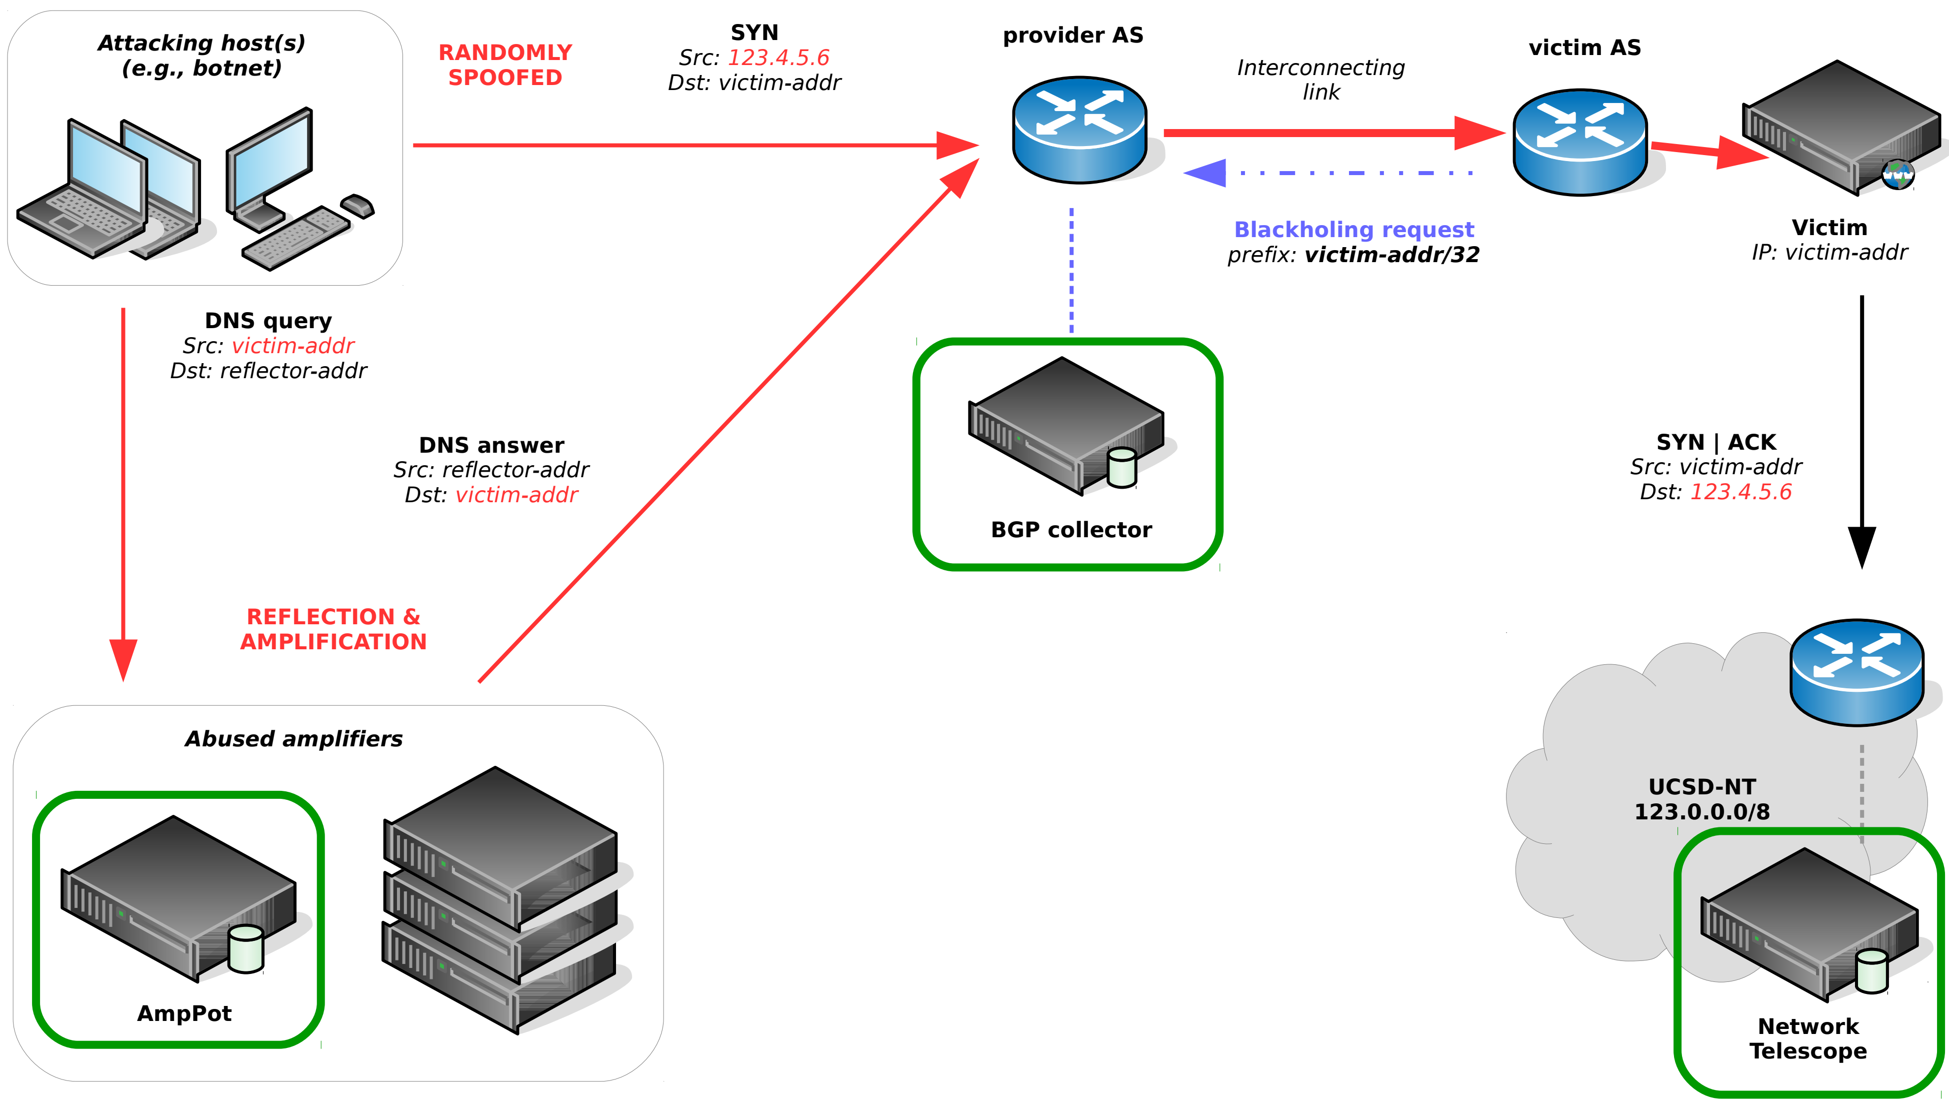

Attacks are mitigated within minutes

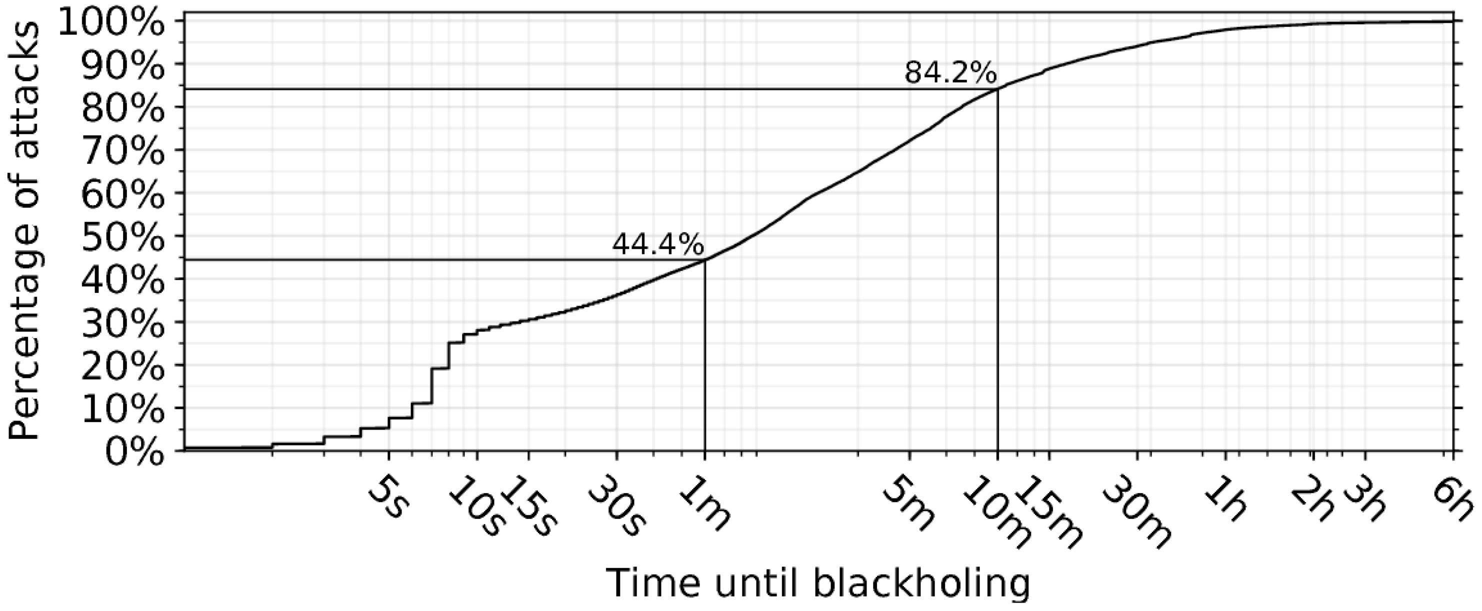

- More than half of attacks mitigated within minutes
  - 84.2% within **ten** minutes
  - takes longer than **six** hours for only 0.02%
- Suggest use of automated, rapid detection and mitigation

Blackholing endures after attacks end

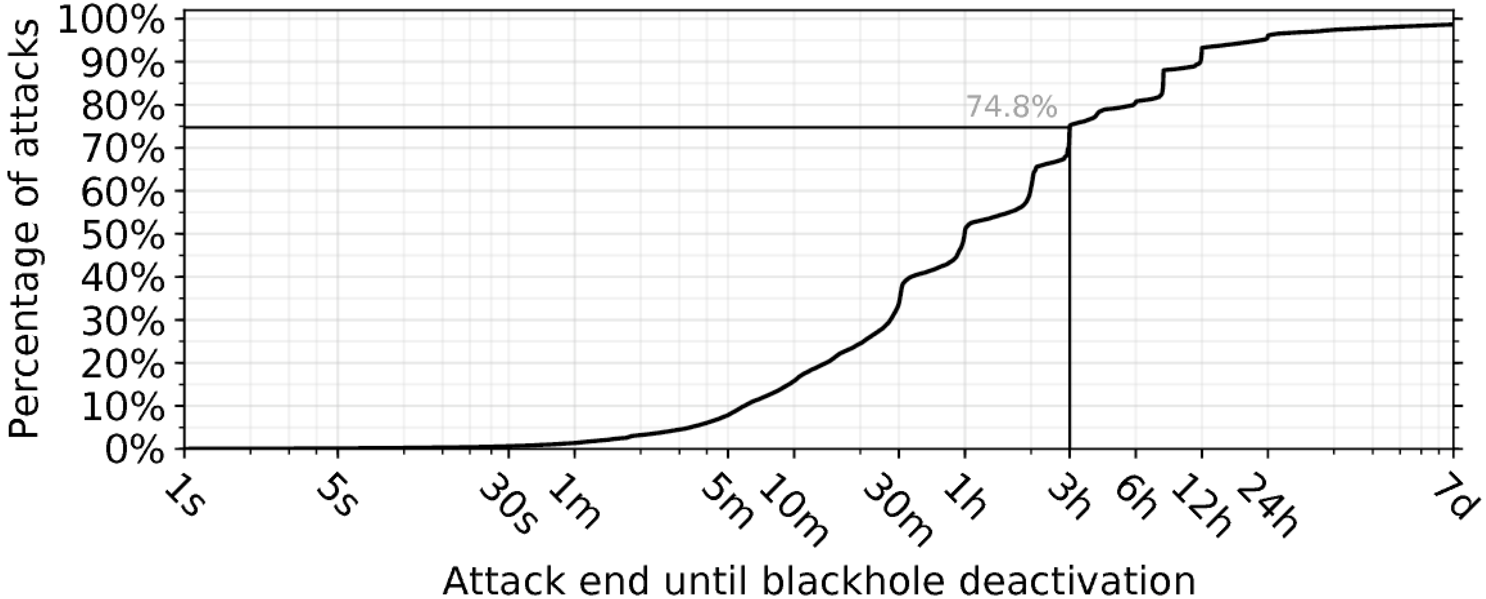

- Deactivated within **three** hours following 74.8% of BH’d attacks
- For 3.9% it takes more than **24** hours
  - Suggests lack of automation in recovery
- Side effects of coarse-grained technique extend well beyond duration of attack

Less intense attacks are also BH’d

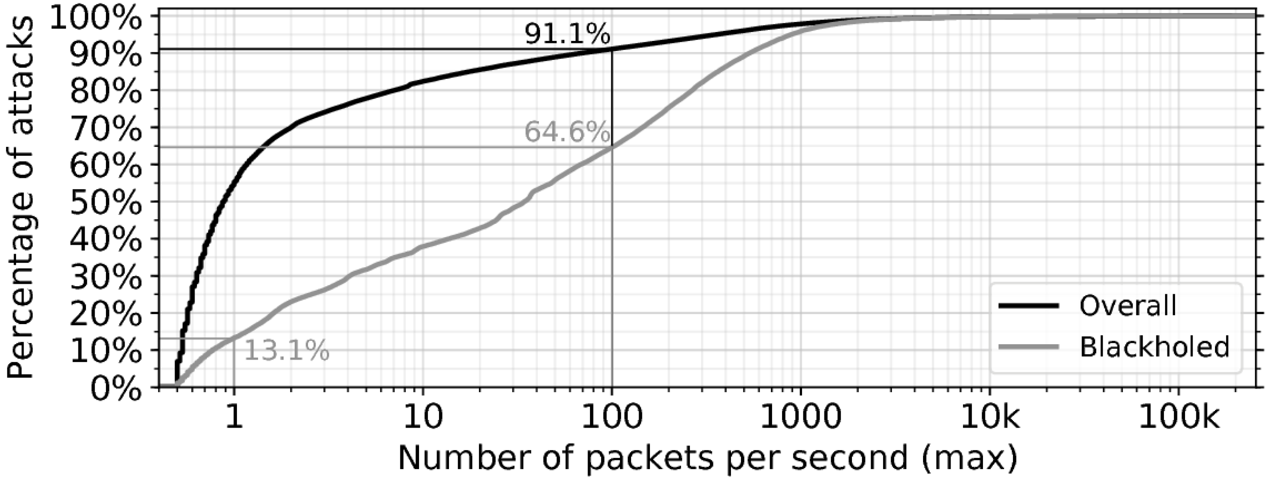

- ~2/3rd of ***BH’d*** attacks (against ~9/10th of ***all*** attacks) have an intensity of up to ~300Mbps (100pps),
- 13.1% see at most 3Mbps (1pps), showing that operators take drastic measures for less intense attacks
- Similar findings for *reflection attacks* (see paper)
- Results confirm Moore et al. methodology at scale (USENIX ‘01)
- Corroborates our previous finding of ~30k attacks/day (IMC ‘17) [1]

[1] Jonker et al., “Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem”, in IMC 2017

Attacks we do not see

- Match blackholing events with preceding attacks

source	#BH events	#BH'd prefixes
UCSD-NT ∪ AmpPot	363.0k / 1.3M (27.8%)	45.2k / 146.2k (30.9%)

- We match 27.8% of BH events with DoS attacks
- Results do not allow us to infer the fraction of other types of attacks (e.g., direct and unspoofed)
- However, highlights that *reflection* and *randomly spoofed* DoS represent a significant share of DoS that operators had to deal with

Part 2: Service Collateral

DNS Measurements [data set 1/2]

- Large dataset of active DNS measurements
- Provides mappings from IPv4 to:
  - *Websites (www. → A RR)*
  - *Mail exchangers* (MX → A)
  - *Authoritative nameservers* (NS → A)
- We use .com, .net & .org (~50% of global namespace)

type	**#prefixes**	**#names associated**		
		overall	**no-alt**	**ratio**
Web	13.7k (9.3%)	782k	670k	0.86
Mail	2247 (1.5%)	180k	177k	0.98
NS	1176 (0.8%)	10k	10k	0.99

Open INTEL

Reactive measurements [data set 2/2]

- Reactively measure blackholed /32s
  - Upon BH *activation* (i.e., announcement) and *deactivation* (i.e., withdrawal/re-announcement)
  - Subject to various heuristics (max 4 in /24, spacing, ...)
- Use RIPE Atlas to send traceroutes
  - From probes in *peer*, *customer* & *provider* networks
- Scan a handful of IANA-assigned ports
  - For Web, mail and DNS
  - From a single VP

RIPE NCC
RIPE Atlas

BGP STREAM

Inferring blackhole (in)efficacy

Port probes

- Exclusively open state on *deactivation* → infer efficacy
- Open on *activation* → infer inefficacy
- Other cases → inconclusive

Traceroutes

- Exclusively last_hop_is_destination on *deactivation* → infer efficacy
- last_hop_is_destination on *activation* → infer inefficacy

Port probe inferences

response	#service		
	Web	Mail	DNS
a ⋃ d	2886	464	528
a ⋂ d	6.98%	8.41%	11.36%
a \ d	0.38%	0.43%	0.76%
d \ a	92.64%	91.16%	87.88%

- Jointly, we infer efficacy in 95.25% of “coverable” cases
- The **a \ d** category is near-zero, which supports the chosen methodology

Trace route inferences

Probe network	**#groups**	**inferrence**		
		Efficacy	Inefficacy	∩
peer	5.0k	29%	8%	1.0%
provider	5.4k	29%	6%	0.8%
customer	2.0k	17%	8%	2.1%

- Jointly, we infer efficacy significantly more often than inefficacy
- But our “coverage” is limited (i.e., last hops never respond)

Corroborated Service Collateral

type	#prefixes	#corroborated names	#affected
Web	734	30916	
Mail	107	3533	522
NS	46	323	708

- Unreachable for the duration of the blackhole
  - At least for part of the Internet
- However
  - MTA retries may simply incur a delay
  - Cache mechanism may mitigate NS issues

Conclusions

- We started addressing the lack of understanding in how blackholing is used in practice when DoS attacks occur
  - *e.g., we wondered if blackholing is used only in extreme cases*
- Although we only provide first insights, our findings show:
  - Rapid reaction times suggest frequent use of automation
  - Excessive retention times suggest lack of automated recovery
  - Less intense attacks are also mitigated
- Preliminary augmentation with complementary measurements
  - Enabled us to corroborate BH (in)efficacy
  - “coverage” is limited (e.g., due to observation delays, firewalls)
- Future work
  - We linked only 28% of blackholing to attacks!
  - Improve reactive measurements (e.g., path or last hop analyses)

Questions
?

Mattijs Jonker
m.jonker@utwente.nl
linkedin.com/in/mattijsj/
mattijsjonker.com

BACKUP SLIDES

Previous study [1/2]

DoS characterization at scale

- Integrates data from a large darknet, honeypots and a platform for DNS measurements
- Finds macroscopic and detailed insights about DoS attacks
  - ~30k attacks daily, Internet-wide
  - Affecting many networks and /24 blocks
  - Various attack types are sometimes launched simultaneously against the same target
  - Migration to cloud-based protection occurs faster following more intense attacks

Jonker et al., "Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem", in IMC 2017

Previous study [2/2]
Blackholing activity at scale

- Systematically studies BGP blackholing at scale

  … using large public and private BGP routing data sets
- Finds detailed insights that relate to, among others:

  … the adoption of blackholing over time

  … effects on the data plane

  … operational practices

Giotsas et al., "Inferring BGP blackholing activity in the internet", in IMC 2017

Data sets

Attacks: 28 million in total

source	#events	#targets	#ASNs
UCSD-NT $\cup$ AmpPot	28.1M	8.6M	36.9k
UCSD-NT $\cap$ AmpPot	447.6k	0.2M	9.2k

- Blackholing events: 1.3 million in total

#BH events	#prefixes	#origins
1.3M	146.2k	2.7k

Blackholed attacks [1/2]

- Match attacks with succeeding mitigation through BH

  … by requiring BH prefix to “cover” attacked /32

  … and cap at 24h

source	#attacks	#targets	#ASNs
UCSD-NT ⋃ AmpPot	456.0k / 28.1M (1.6%)	70k / 8.6M (0.8%)	2.5k
UCSD-NT ⋂ AmpPot	18.4k / 447.6k (4.1%)	5.7k / 6.0M (3.3%)	0.8k

- Small percentages suggest noise, but:
  - Small attack intensities trigger BH (later)
  - We can observe BH only for a subset of ASes/targets
  - 2.5k ASes involved significant, but BH use might not be largely widespread
- Joint attacks (⋂) appear more likely to be BH’d

Blackholed attacks [2/2]

- Match blackholing events with preceding attacks

source	#BH events	#BH'd prefixes
UCSD-NT ∪ AmpPot	363.0k / 1.3M (27.8%)	45.2k / 146.2k (30.9%)

- We match 27.8% with attacks
- Results do not allow us to infer the fraction of other types of attacks (e.g., direct and unspoofed)
- However, highlights that *reflection* and *randomly spoofed* DoS represents a significant share of DoS that operators had to deal with

Observation delay