Network Architectures and Services
Faculty of Informatics
Technische Universität München, Germany

DoS Amplification Attacks - Protocol-Agnostic Detection of Service Abuse in Amplifier Networks

Timm Böttger†, Oliver Gasser, Lothar Braun,
Felix von Eye, Helmut Reiser, Georg Carle

Technische Universität München and
Leibniz Supercomputing Centre

† now Queen Mary University of London

Queen Mary
University of London

Technische Universität München

Agenda

1. On Amplification Attacks
2. Detection Approach
3. Evaluating the Approach
4. Conclusion

ON AMPLIFICATION ATTACKS

Motivation

- Why is DoS detection important?

- Spamhaus DoS amplification attack
  - March 2013
  - 75 Gbit DoS traffic, mostly DNS
  - Made Spamhaus email blocklist unavailable

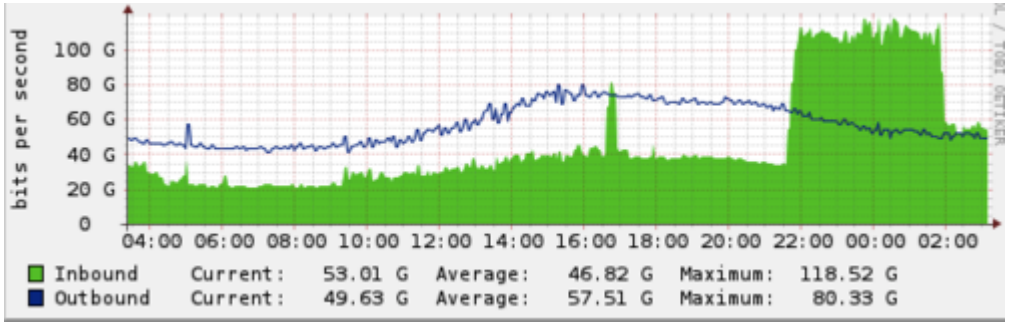

https://blog.cloudflare.com/the-ddos-that-knocked-spamhaus-offline-and-ho/

Example of an Amplification Attack

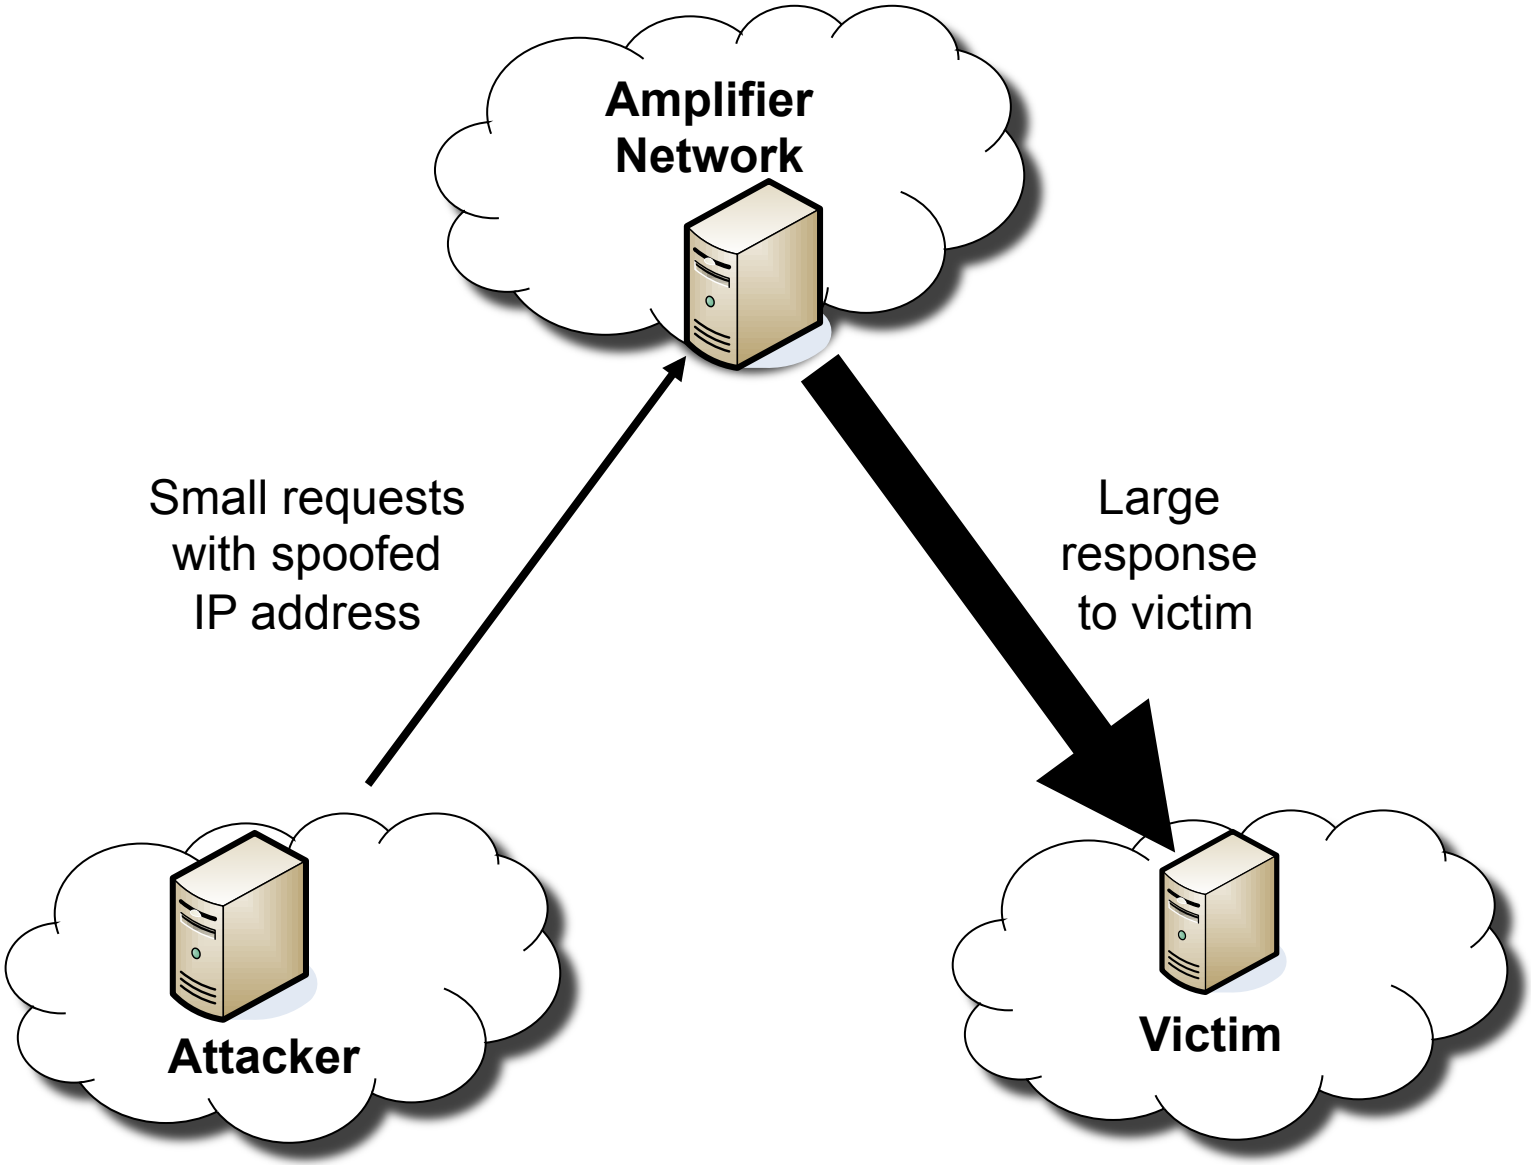

Example of an Amplification Attack

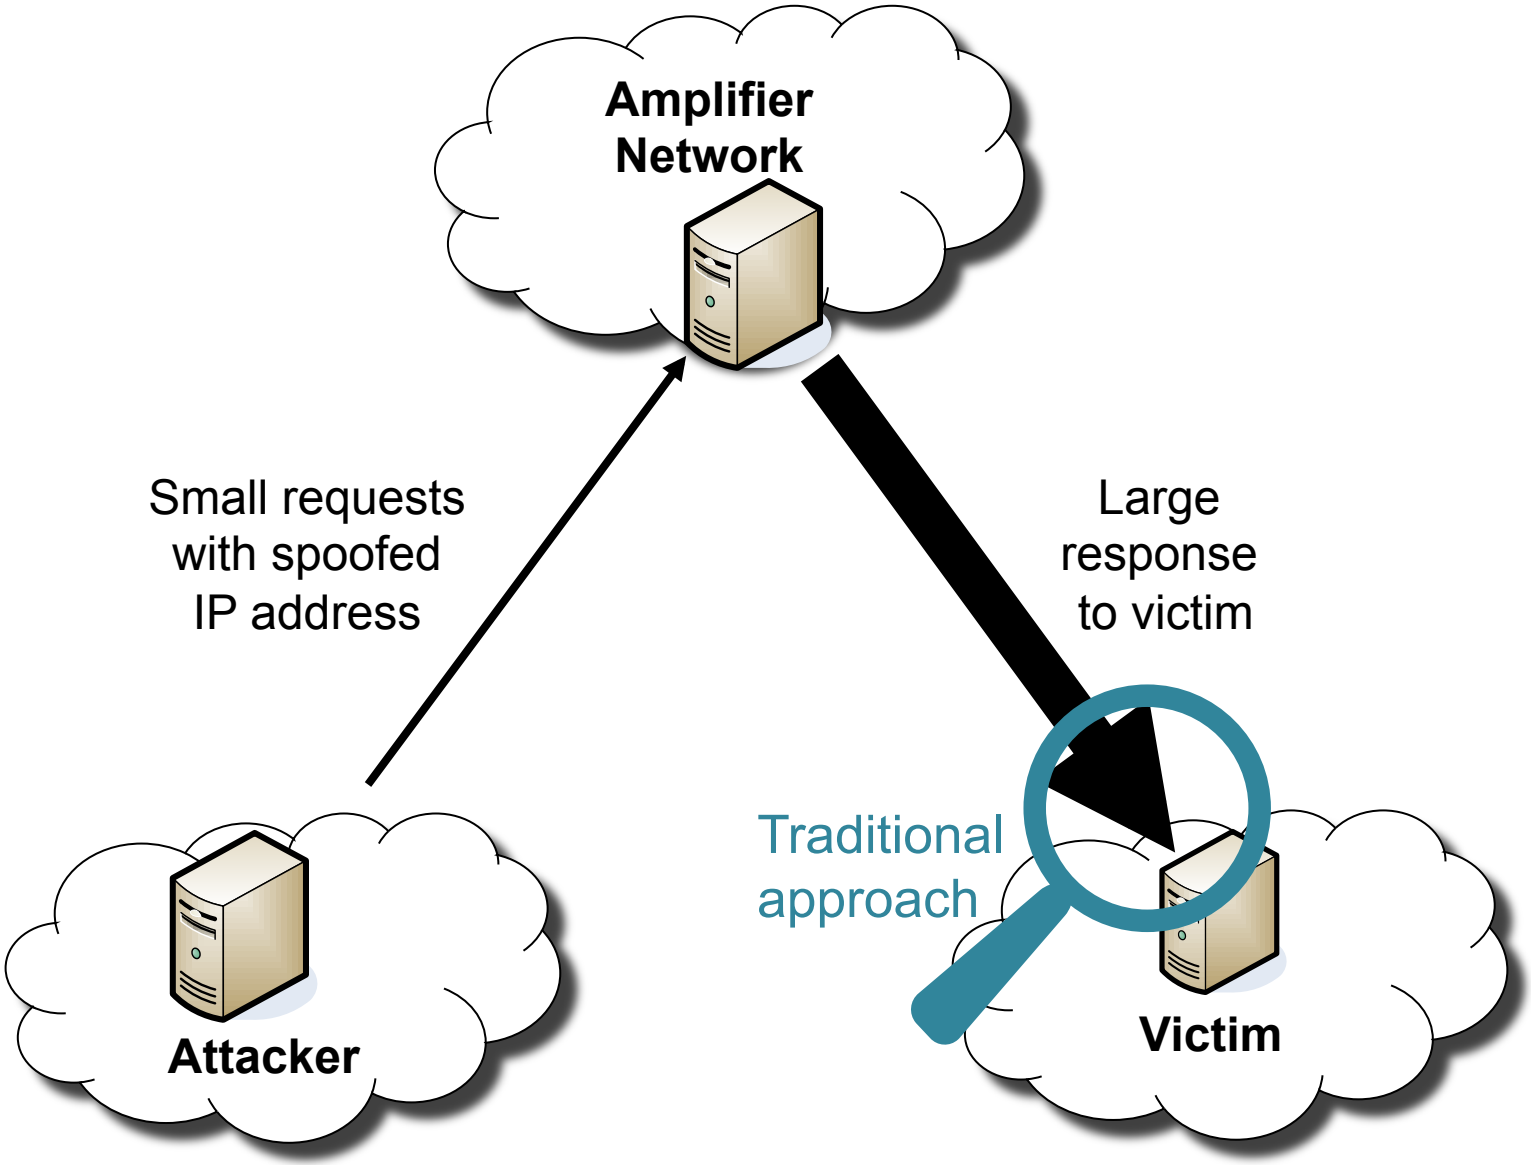

Example of an Amplification Attack

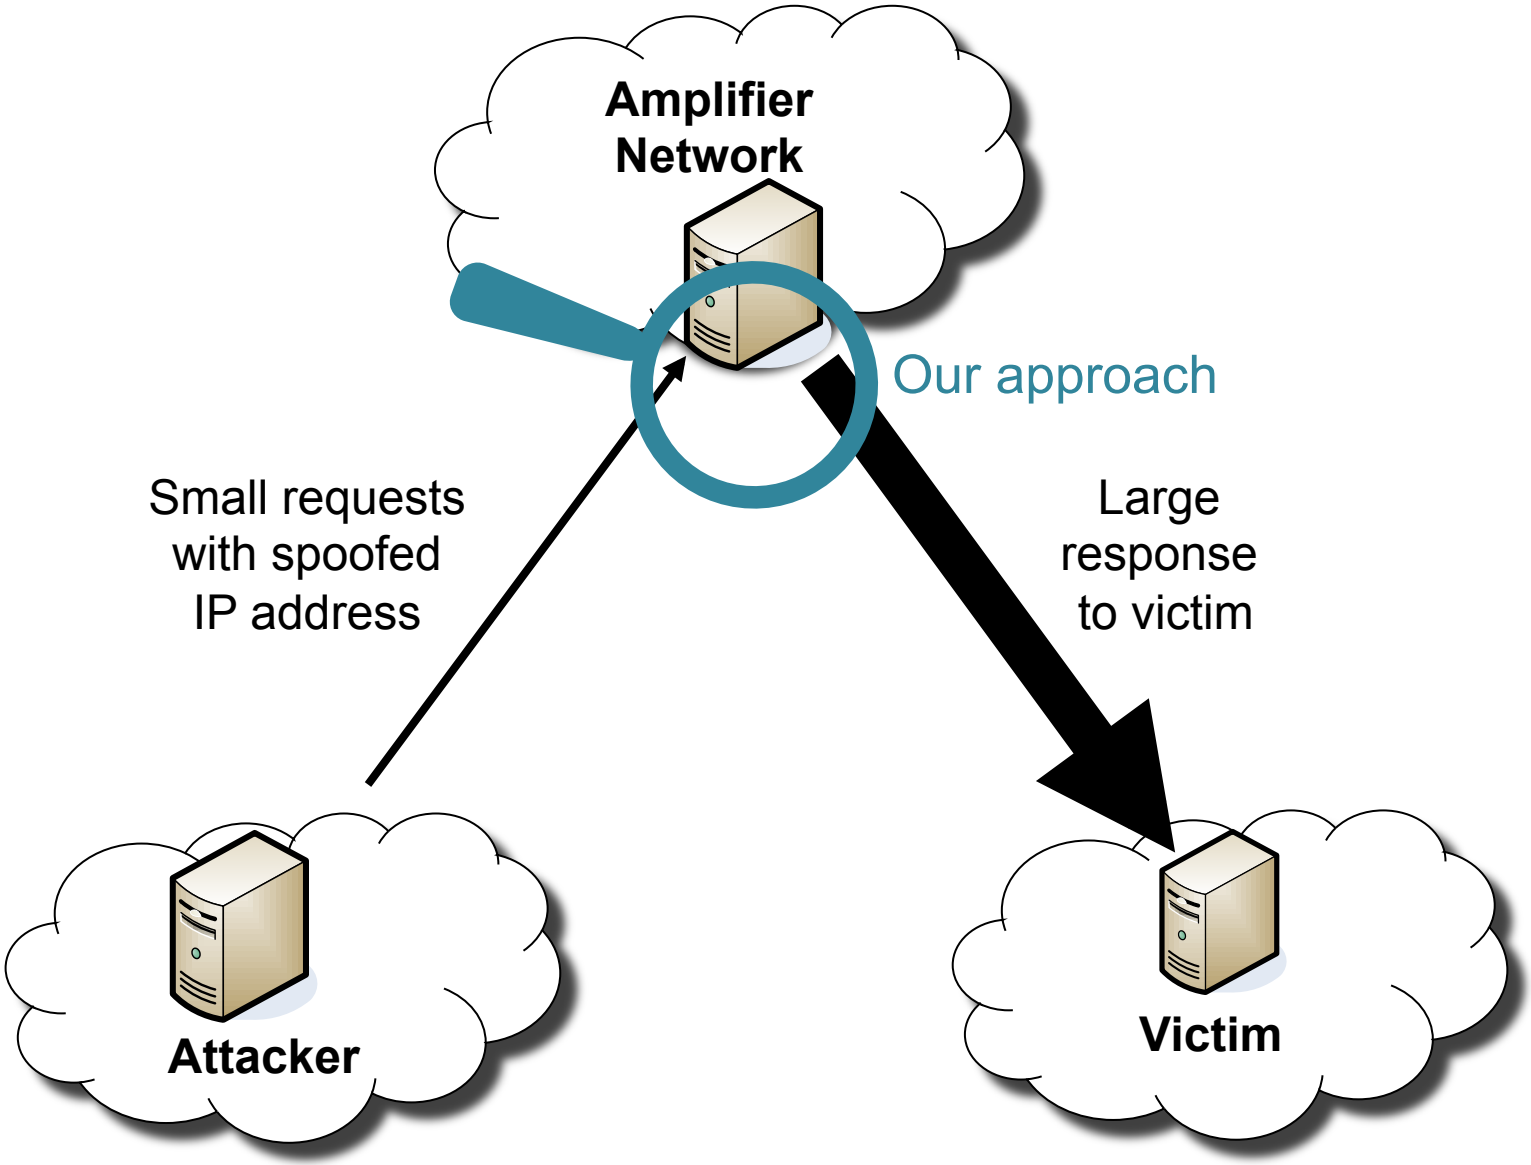

Example of an Amplification Attack

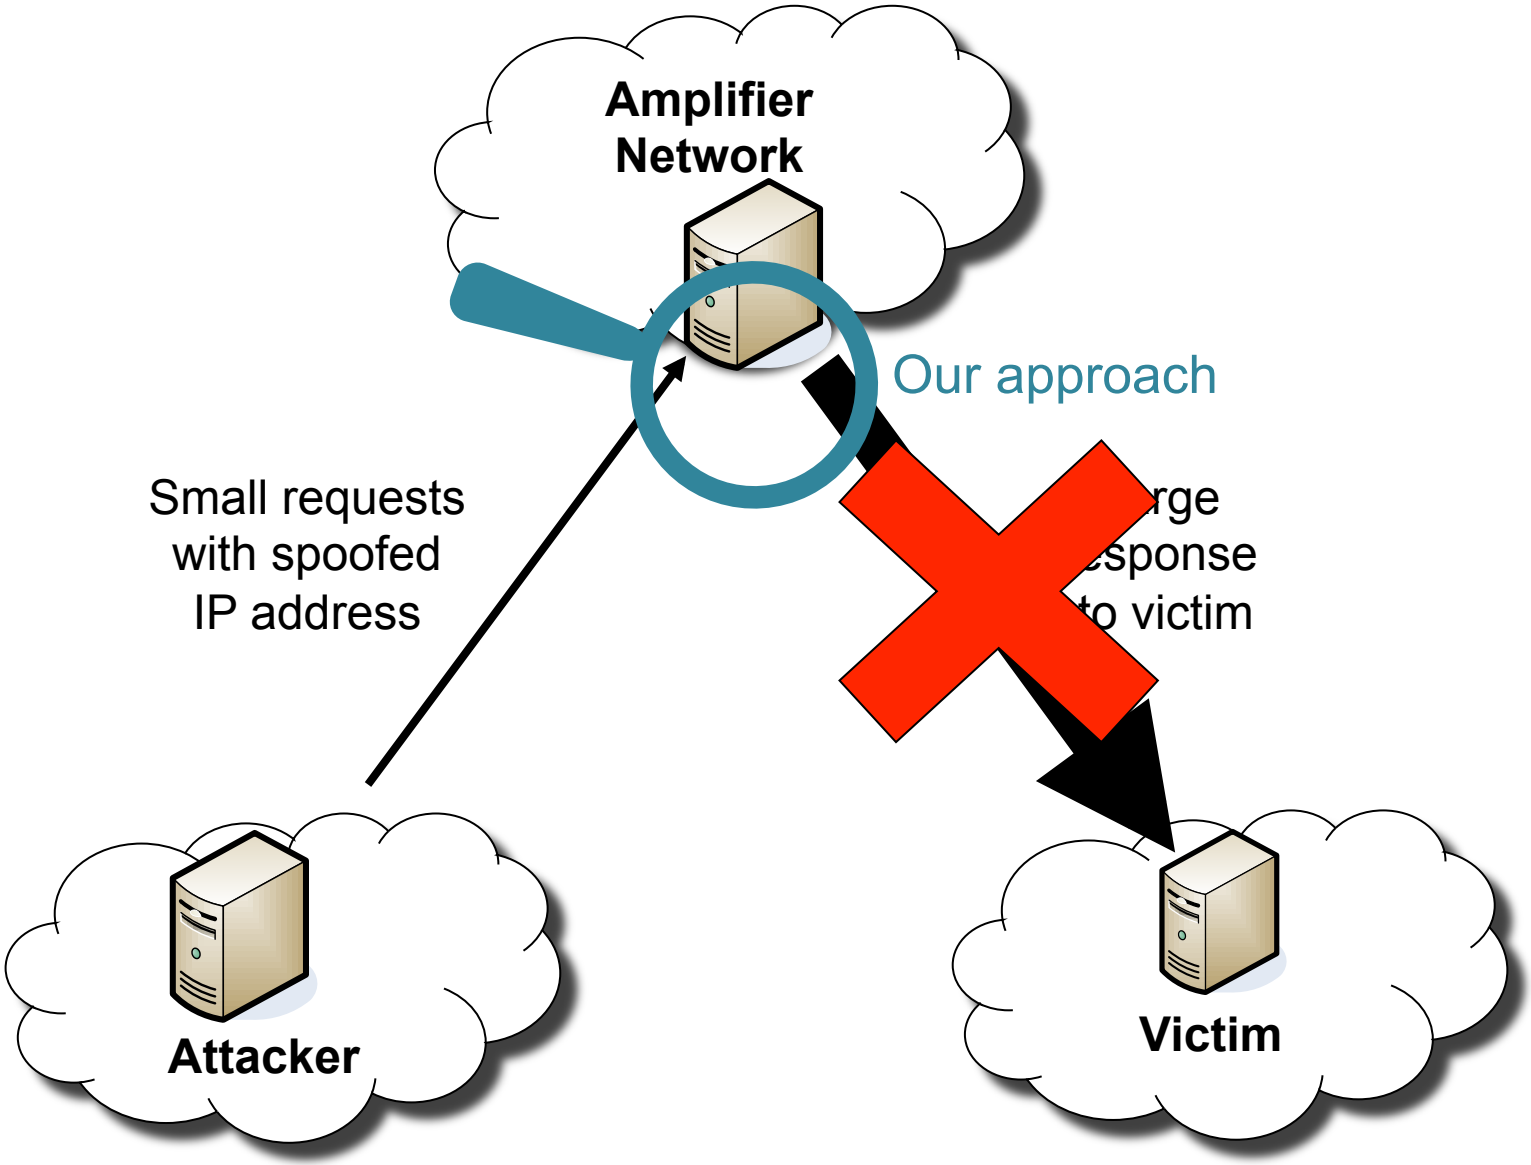

DETECTION APPROACH

Previous Work

- Christian Rossow from RUB (now Saarland U) worked on amplification attack detection as well
- Christian Rossow. *"Amplification Hell: Revisiting Network Protocols for DDoS Abuse"*, NDSS 2014.
- His work provided insightful ideas for amplification attack detection in amplifier networks
- We extended and generalized his approach to arbitrary ports and protocols

Aggregation to Pairflows

- Amplifier very likely is a server
- Server – Client communication:
  - Server uses fixed port
  - Client can change port with each request

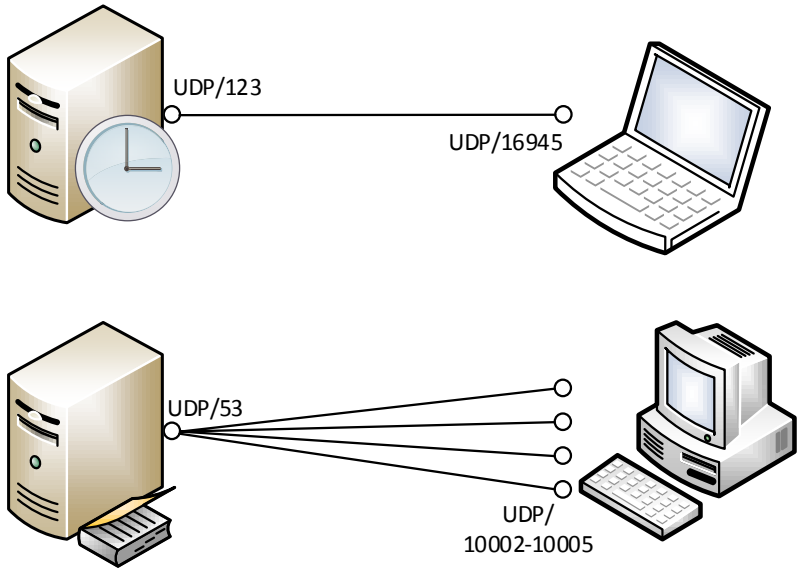

- Aggregate to pairflows

Making it Protocol-Agnostic

- Thresholds proposed by Rossow:
  - ≥ 10 MB traffic in 10 minutes
  - Server sends 5 times more than it receives (BAF ≥ 5)
- Works well if applied to fixed ports
- Generalizing to arbitrary ports introduces false positives
- We need additional detection criteria

The Attacker's Point of View

- Send requests to amplifier and expect large answer
- Challenges for the attacker
  - Attacker cannot see server's replies
  - Can not use requests which require shared state → UDP
  - No direct proof of successful attack
- Attacker preferably uses requests validated before
- Typically a small number of different requests

Detectable Attack Traffic Properties

- Asymmetric traffic: small request, large response
- Similar payload between packets in each direction
- Similar packet size between packets in each direction
- Victim does not expect amplifier's responses traffic: ICMP port unreachable

Further Attack Traffic Properties

- Unsolicited messages
  - Problem solved if matching between request and response done inside network
- IP spoofing
  - Integral part of attack
  - Filtering spoofed packets would mitigate attack
  - BCP 38
- Both properties difficult to detect in general

EVALUATING THE APPROACH

Measurement Setup

- Implemented detection mechanism in Suricata IDS
- Automated post processing of detection results

- Multiple measurement runs at Internet gateway of Munich Scientific Network (MWN)
  - MWN connects Munich's universities, student resident halls, research institutes
  - Avg. 2.6 Gbit/s incoming, 1.5 Gbit/s outgoing
  - Total 1200 TB inbound, 730 TB outbound in one month

Three Measurement Runs

	Run #1	Run #2	Run #3
Duration (in h)	144	96	24
Total Bytes Sent	7,340.66 GB	3,425.62 GB	734.67 GB
Total Packets Sent	6,589,456,476	3,208,724,852	674,865,692
Total Pairflows Reported	77,693	45,747	10,974
Unique Server-Port-Client Triples	22,428	14,567	4,058
Unique Server-Port Pairs	3,324	1,682	504
Unique Servers	530	309	204

Deriving Detection Thresholds

- No ground truth on attacks
- Conduct own attacks to derive detection thresholds
  - Measurement run 1
  - NTP, DNS, SNMP, Chargen, SIP, QOTD, BitTorrent
- Derive detection threshold such that all these attacks are detected

Payload Similarity

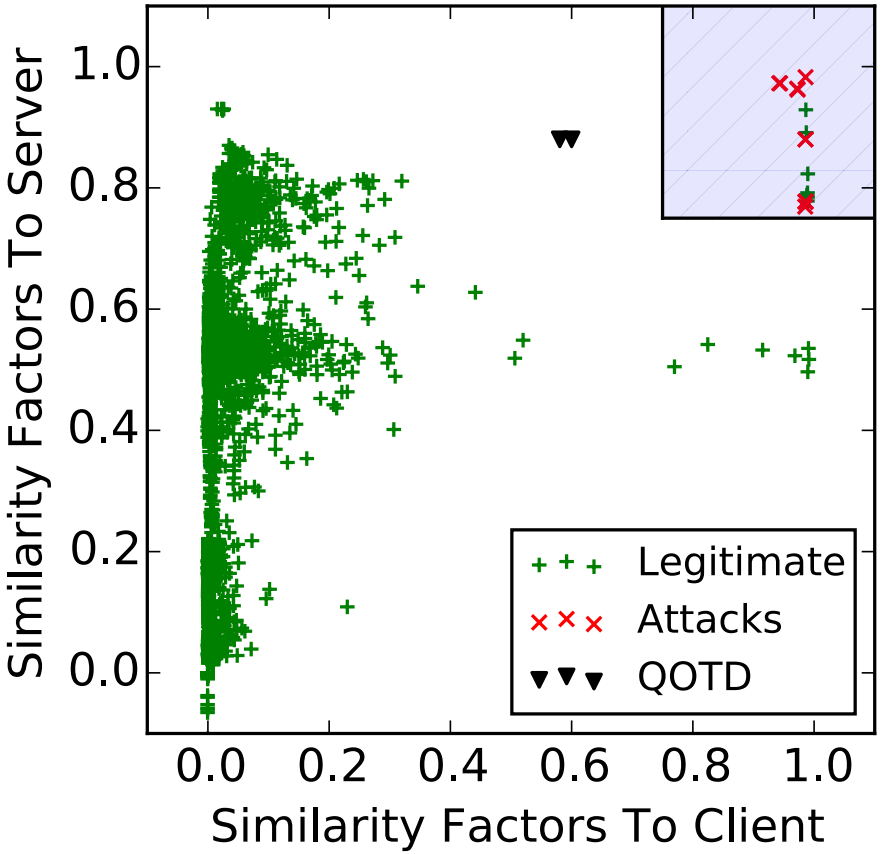

- Comparing payload similarity by compressing packets
- High compression ratio → high similarity
- Attack traffic shows high similarity in packet payload

Differences in Packet Sizes to Server

- Looking at packets directed to server
- Attack traffic showed smaller difference in packet size
- Set detection threshold to 25 bytes

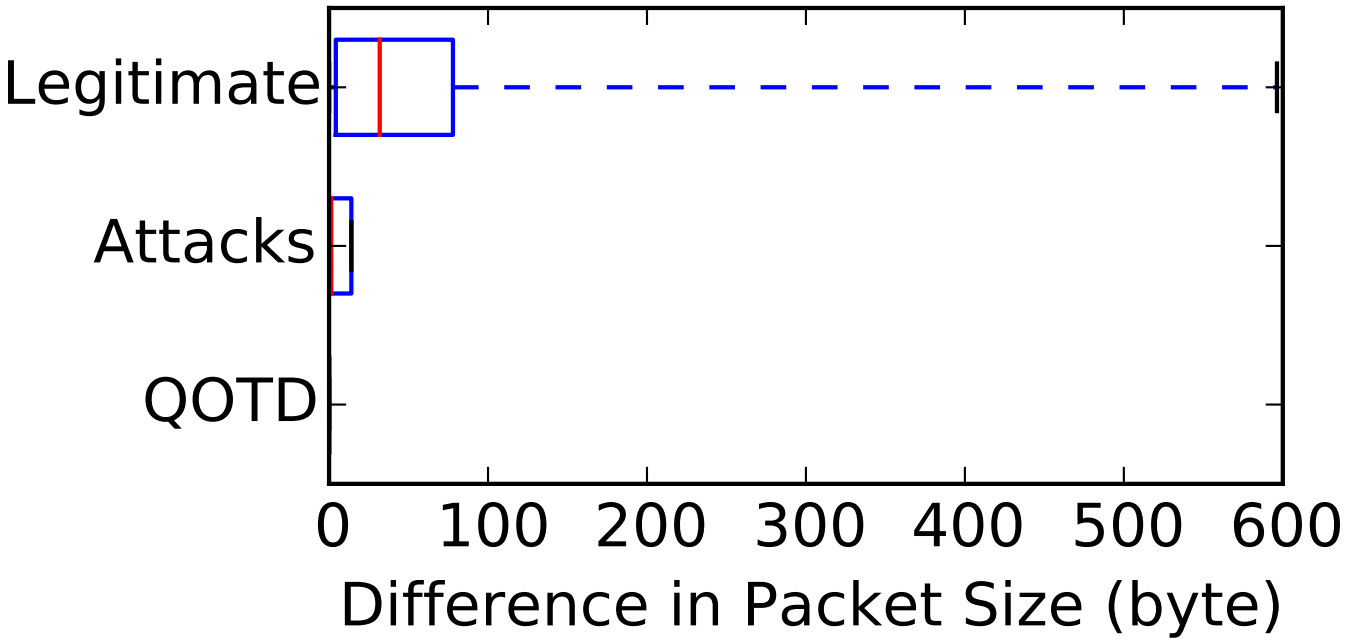

Differences in Packet Sizes to Client

- Looking at packets directed to client
- Attack traffic showed smaller difference in packet size
- Set detection threshold to 25 bytes

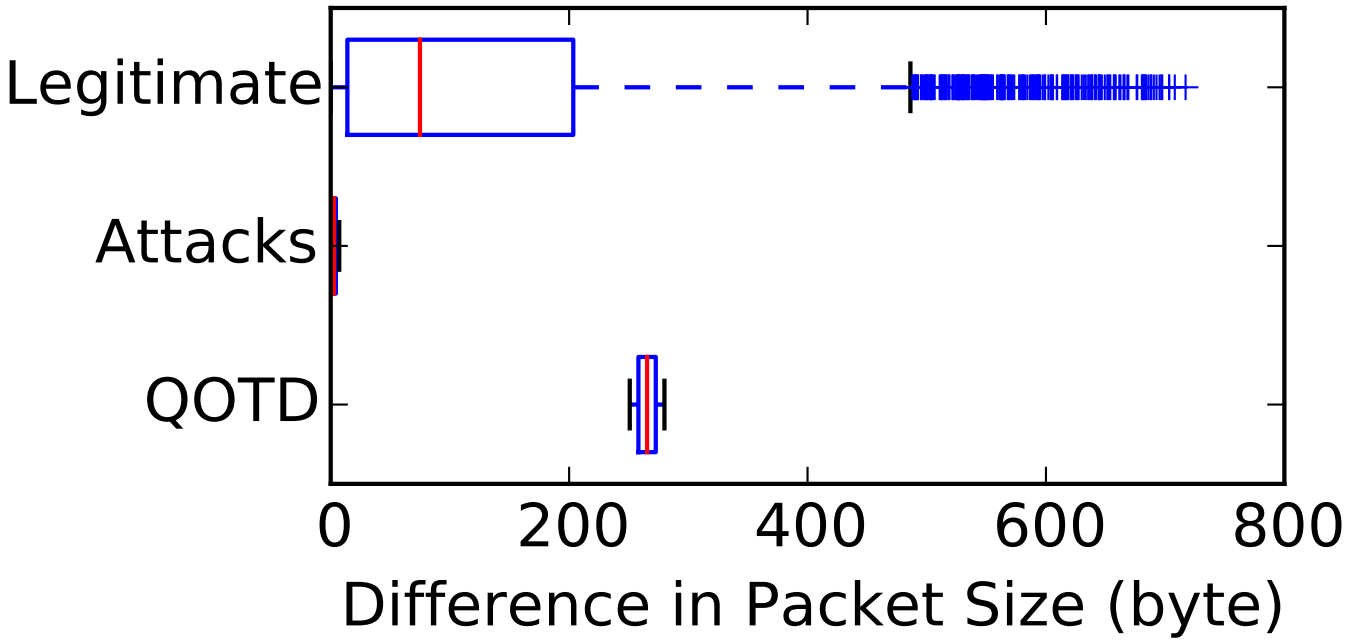

ICMP Port Unreachable Replies

- Criterion seems generally usable
- But only small number of attacked hosts sent ICMP messages
- Excluded criterion from attack detection engine

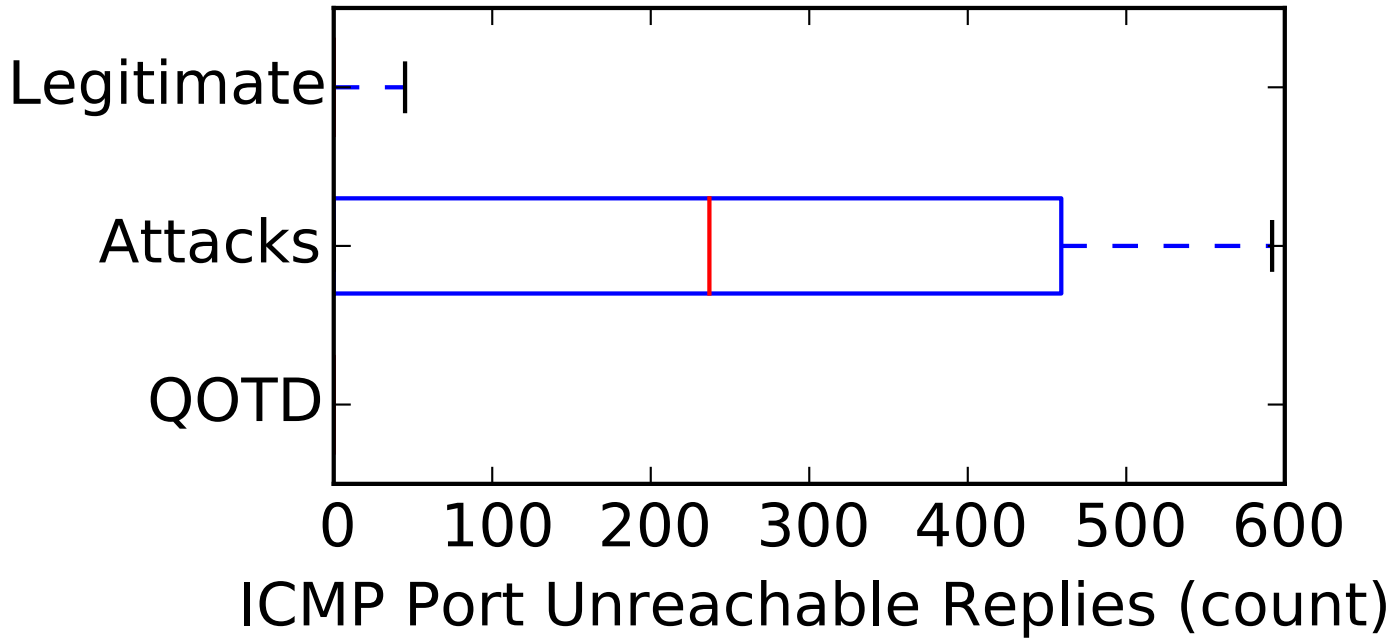

Evaluation Results

	Run #1	Run #2	Run #3
Duration (in h)	144	96	24
BAF identified services	3,324	1,682	504
BAF identified alarms	22,428	14,567	4,058
True positive alarms	277	30	18
False positive alarms	3	9	0
True negative alarms	22,149	14,534	4,041
False negative alarms	0	0	0

CONCLUSION

Conclusion

- Amplification attacks remain a threat to the availability of Internet services
- We identified attack properties and derived a detection approach
- Multiple measurement runs proved its viability
- More details in the paper ☺

Acknowledgements

- This work has been supported by the German Federal Ministry of Education and Research (BMBF) under support code 01BY1203C, project Peeroskop, and 16BP12304, EUREKA project SASER, and by the European Commission under the FP7 project EINS, grant number 288021.

Technische Universität München

SASER

SEVENTH FRAMEWORK PROGRAMME

Federal Ministry of Education and Research

ENDEAVOUR

Queen Mary University of London

- This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 644960.

Questions

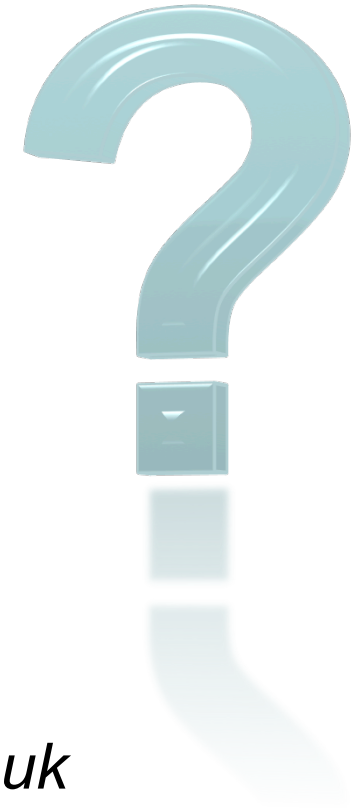

Timm Böttger
boettget@eecs.qmul.ac.uk

Oliver Gasser
gasser@net.in.tum.de

Queen Mary University of London

Technische Universität München

RELATED WORK

Related Work

- Fathi Özavci: VoIP Wars : Return of the SIP, 2013
- Jerome Harrington: A BitTorrent-Driven Distributed Denial-of-Service Attack, 2007
  - Mainly focusing on describing attack vectors, no focus on detection

- Georgios Kambourakis: Detecting DNS Amplification Attacks, 2007
- Changhua Sun, Efficient and Low-Cost Hardware Defense Against DNS Amplification Attacks, 2008
  - Focus on the victim's network, not the amplifier's one

EVADING DETECTION

Evading Detection

- Attacker can reduce amount of traffic
  - Attack is weakened, desirable outcome
- Attacker can use multiple amplifiers to stay below individual detection thresholds
  - 10MB/10 min allow only up to 136 kbit/s
- Attacker can undermine similarity detection by sending garbage traffic
  - Sending random messages will lower amplification factor