

Network Security and Measurement

- Distributed Denial of Service (DDoS) -

Prof. Dr. Thomas Schmidt

<http://inet.haw-hamburg.de> | t.schmidt@haw-hamburg.de

Agenda

DDoS attacks

Examples of DDoS attacks

Spoofing and spoofing detection

Blackholing & Filtering

How it works

DDOS ATTACKS

Distributed Denial of Service - DDoS

Method to interrupt or take down a service using multiple, coordinated machines

- First DDoS in July 1999 took down a network at U of Minnesota using worm Trin00
- Since then, a rich ecosystem developed that delivers medium size attacks on demand (Booters) *and* massive attacks above 1 Tbps (2.3 Tbps, Feb'20 on AWS)

Attacks are generated following economical, political, or personal reasons

Distributed Denial of Service - DDoS


DDoS attacks can be operated from a botnet – a group of infected machines, or from one source that misuses Internet infrastructure

Method to interrupt or take down a service using multiple, coordinated machines


- First DDoS in July 1999 took down a network at U of Minnesota using worm Trin00
- Since then, a rich ecosystem developed that delivers medium size attacks on demand (Booters) *and* massive attacks above 1 Tbps (2.3 Tbps, Feb'20 on AWS)

Attacks are generated following economical, political, or personal reasons


Reflection attack with spoofing



Amplification attack with reflection



Amplification attack with reflection



Popular amplifying protocols

Applications

- CLDAP (389)
- Apple Remote (5900)
- Memcached (11211)
- Bittorrent (many)
- Quake3 (27960)
- Steam (27015)

Network services

- DNS (53)
- NTP (123)
- Netbios (137)
- SSDP (1900)
- ICMP

From the long history of DDoS

EXAMPLES OF ATTACKS

Smurf attack

Ping flooding attack assisted by the network
Attacker sends ICMP echo request to (remote) network broadcast address, using the spoofed source address of the victim

Simple attack mechanism exploited in the early days of DDoS attempts

DNSPod attack in China (2009)

81 servers instrumented to bring down the DNSPod DNS service in China

Initial outages caused a cascading effect: The popular video platform Baofeng.com relied on DNSPod and its video player Storm created a request storm, unintentionally overloading DNSPod

Original motivation: Two competing game service providers trying to impair each other

Effect: Network outages in several Chinese provinces affecting millions of users

Operation Ababil (2011 – 2013)

Multilateral attack focused on American financial institutions

Exploited a PHP vulnerability on various powerful Web servers, via which the attack was launched

Operated in isolated phases, reached up to 70 Gbps attack traffic

Attackers claimed to be Arab cyber fighters, but specialists believe that plain criminals tried to keep bank IT-people busy while working on orthogonal intrusion attacks

Spamhouse attack (2013)



Record breaking attack that achieved the rate of 300 Gbps against Spamhouse:

“It Almost Broke the Internet”

Amplification and reflection attack using NTP servers

Mitigated by Cloudflare using anycast, but did cause major pressure for peers and at LINX

Initiated by a teenage hacker from Britain who was paid for the purpose

Mirai botnet (2016)

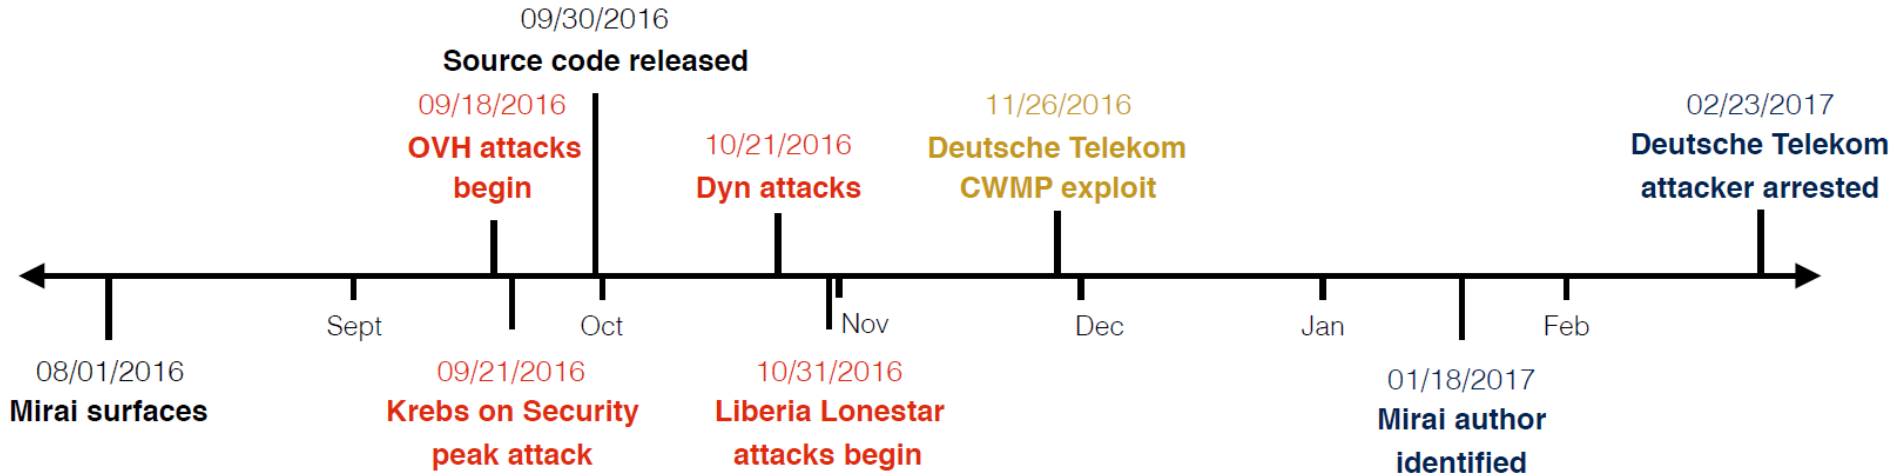
Statelessly scanned the Internet for ARC-based IoT devices to infect and create a botnet

Mirai intrudes systems trying 62 default usernames/passwords

Operated a series of very large attacks, one brought down the registrar DYN, others Liberian Internet providers, and Telekom home routers

Originators were young entrepreneurs of a startup for DDoS mitigation services

Timeline of the Mirai attack



Post Mirai

Mirai source code was published early and inspired copycats

An entire ecosystem evolved: Numerous variants were generated in different contexts

Mutating Mirai-type software is continuously populating the net: New botnets, new intents, new exploits – a toolbox to continuously threatening the Internet

Enabling reflection attacks

SPOOFING


IP spoofing

Spoofed packets include an incorrect source IP address


Bogon filters help partially

Ingress filters are not well deployed

IP spoofing enables amplification attacks, a major threat



GITHUB SURVIVED THE BIGGEST DDoS ATTACK EVER RECORDED



The DDoS That Knocked Spamhaus Offline (And How We Mitigated It)

20 Mar 2013 by Matthew Prince.

At CloudFlare, we deal with large DDoS attacks every day. Usually, these attacks are directed at large companies or organizations that are reluctant to talk about their details. It's fun, therefore, whenever we have a customer that is willing to let us tell the story of an attack they saw and how we mitigated it. This is one of those stories.

Spamhaus

Yesterday, Tuesday, March 19, 2013, CloudFlare was contacted by the non-profit anti-spam organization Spamhaus. They were suffering a large DDoS attack against their website and asked if we could help mitigate the attack.


Security

BIGGEST DDoS ATTACK IN HISTORY hammers Spamhaus

Plucky mail scrubbers battle internet carpet bombers

By John Leyden 27 Mar 2013 at 17:03

124 SHARE



Anti-spam organisation Spamhaus has recovered from possibly the largest DDoS attack in history.

IP spoofing enables amplification attacks

The IETF early advised to deploy filters that prevent use of spoofed IP source addresses

- Best Current Practice: BCP 38 prescribes network ingress filtering

Unfortunately, deployment is heterogeneous

How to identify networks that allow for spoofing?

Challenges

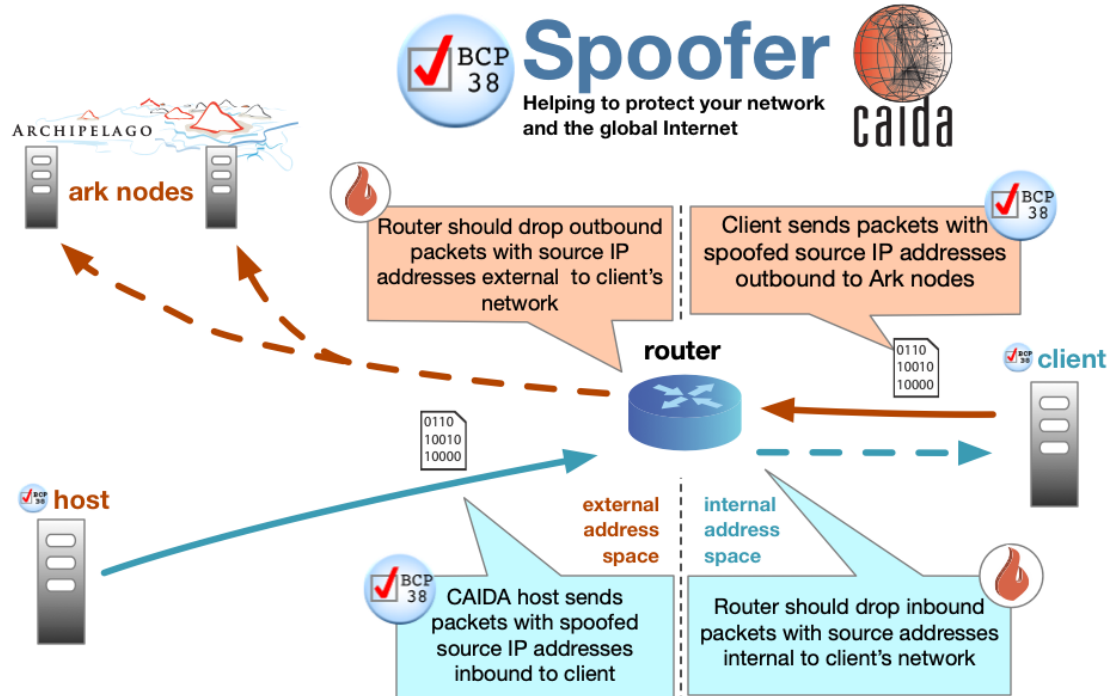
Sending spoofed packets require admin rights

No way to induce spoofed packets


Idea

Volunteers that run a “spoofer” test program

<https://www.caida.org/projects/spoofers>



How to determine where filter is deployed?



Possible outcomes

1. Blocked because of source address filters
2. Sometimes operating systems block spoofed packets, even when raw Ethernet is used
3. NATs rewrite the source address
4. Dropped because of non-spoofed reasons (e.g., congestions)
5. Packet arrives

Literature

Robert Beverly, Arthur Berger, Young Hyun, and k claffy.
[Understanding the efficacy of deployed internet source address validation filtering](#). In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement (IMC '09)*. ACM, New York, NY, USA, 356-369.
<https://doi.org/10.1145/1644893.1644936>

Understanding the Efficacy of Deployed Internet Source Address Validation Filtering

Robert Beverly
MIT CSAIL
rbeverly@csail.mit.edu

Arthur Berger
MIT CSAIL
awberger@csail.mit.edu

Young Hyun
CAIDA
younggh@caida.org

k claffy
CAIDA
kc@caida.org

ABSTRACT

IP source address forgery, or “spoofing,” is a long-recognized consequence of the Internet’s lack of packet-level authenticity. Despite historical precedent and filtering and tracing efforts, attackers continue to utilize spoofing for anonymity, indirection, and amplification. Using a distributed infrastructure and approximately 12,000 active measurement clients, we collect data on the prevalence and efficacy of current best-practice source address validation techniques. Of clients able to test their provider’s source-address filtering rules, we find 31% able to successfully spoof an arbitrary, routable source address, while 77% of clients otherwise unable to spoof can forge an address within their own /24 subnetwork. We uncover significant differences in filtering depending upon network geographic region, type, and size. Our new tracerfilter tool for filter location inference finds 80% of filters implemented a single IP hop from sources, with over 95% of blocked packets observably filtered within the source’s autonomous system. Finally, we provide initial longitudinal results on the evolution of spoofing revealing no mitigation improvement over four years of measurement. Our analysis provides an empirical basis for evaluating incentive and coordination issues surrounding existing and future Internet packet authentication strategies.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design; C.2.3 [Computer Communication Networks]: Network Operations

General Terms

Measurement, Experimentation, Security

Keywords

Source address validation, IP spoofing, filtering

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC '09, November 4–8, 2009, Chicago, Illinois, USA.
Copyright 2009 ACM 978-1-60558-776-7/09/11...\$10.00

1. INTRODUCTION

The Internet architecture includes no explicit notion of packet-level authenticity. A long-recognized [30] consequence of this weakness is the ability to forge or “spoof” IP packet headers. While willing and able networks implement various ad-hoc authentication techniques, history demonstrates that malicious users probe far, and capitalize on, any ability to spoof. A common attack vector is to spoof source IP addresses, to enable anonymity, indirection, and amplification exploits (e.g. [4, 33, 44]).

As good Internet citizens, many operational networks implement source address validation best common practices. Ingress address filtering [18, 45] and unicast reverse path forwarding (uRPF) checks [5] are effective against source spoofing. In practice however, implementation of such techniques is often limited by multi-homing, route asymmetry, lengthy ad-hoc filter list maintenance, and router design. More importantly, current anti-spoofing filtering techniques are hindered by incentive and coordination problems. A provider can follow all best practices and still receive anonymous, malicious traffic from third-parties who do not properly filter. Protection from spoofed traffic using existing practice requires global coordination, a difficult, expensive, and unenforceable goal. As a result, previous research [29, 34] and recent attacks [33] demonstrate that source address spoofing has remained a viable attack vector. Moreover, despite two-decade old exploits [7], new source spoofing based attacks continue to emerge; we review three in §2.

This paper seeks to understand the real-world efficacy of Internet source address filtering best practices. We leverage a widely distributed measurement infrastructure [21] in conjunction with active client measurements to facilitate this understanding. We tailor our probing of the network to infer the extent of different types of filtering. In addition, we develop and use tracerfilter, a novel tool for determining the in-network location of source address filtering. We significantly extend an initial study [10] with the following new contributions and findings:

1. Use of the Ark [21] global distributed measurement infrastructure as active probe reception points. Ark facilitates path-based analysis and tomography over disparate (e.g. commercial, academic, etc) routes.
2. Analysis of ~12,000 unique tests which reveal significant differences between the filtering encountered by clients based on geographic region, network type, and network size.

Mitigating DDoS attacks

BLACKHOLING & FILTERING

How to mitigate DDoS

Remove traffic as early as possible

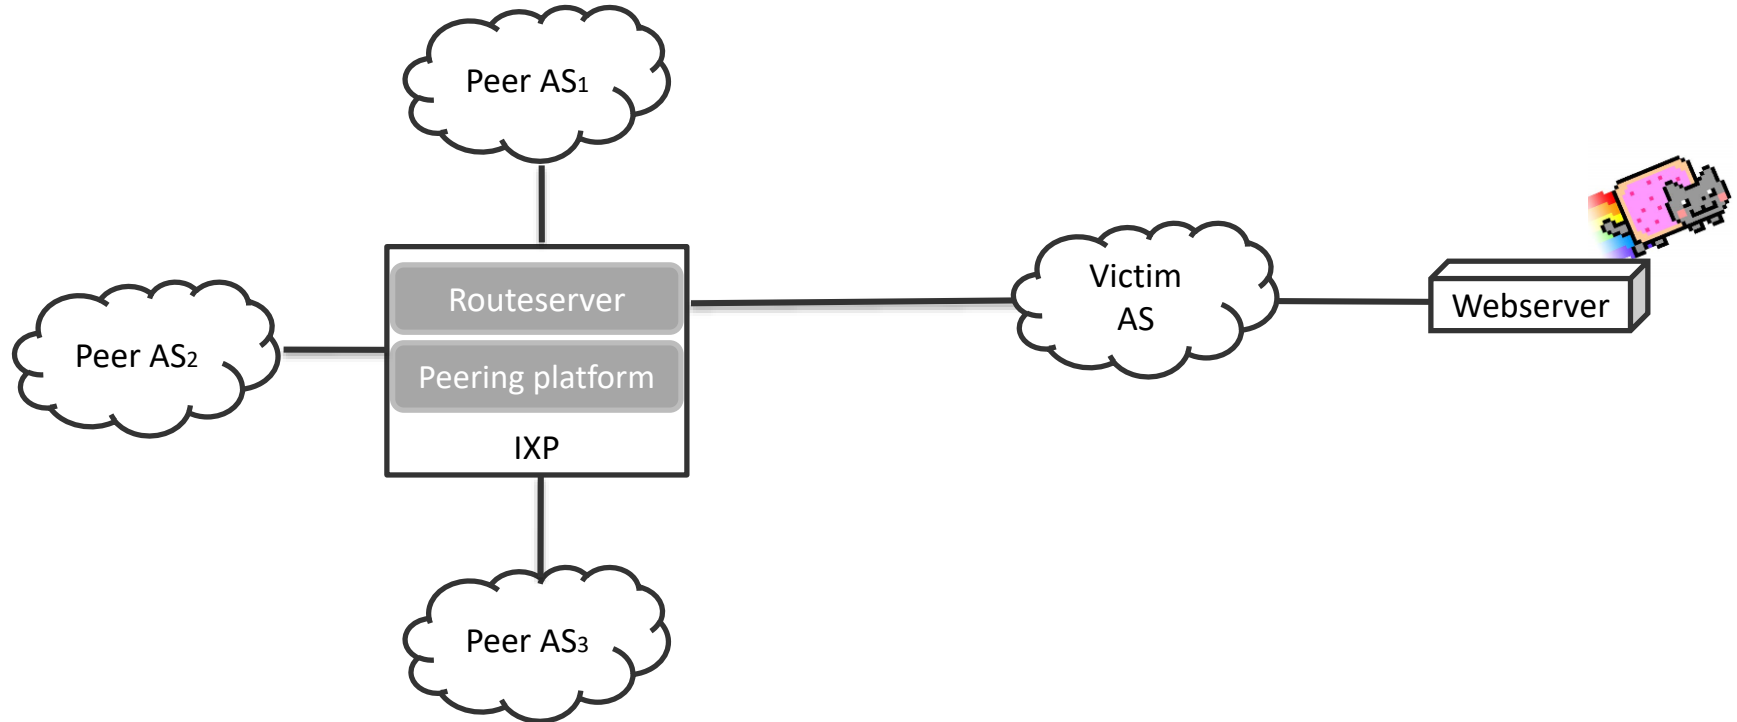
- Easiest at the source – but often impossible
- At transit? – violates business model
- Idea: At IXPs – two options
 - Blackholing
 - Fine-grained filtering
- Blackholing can be remotely triggered by announcing a BGP community
- Filtering can be implemented via BGP FlowSpec



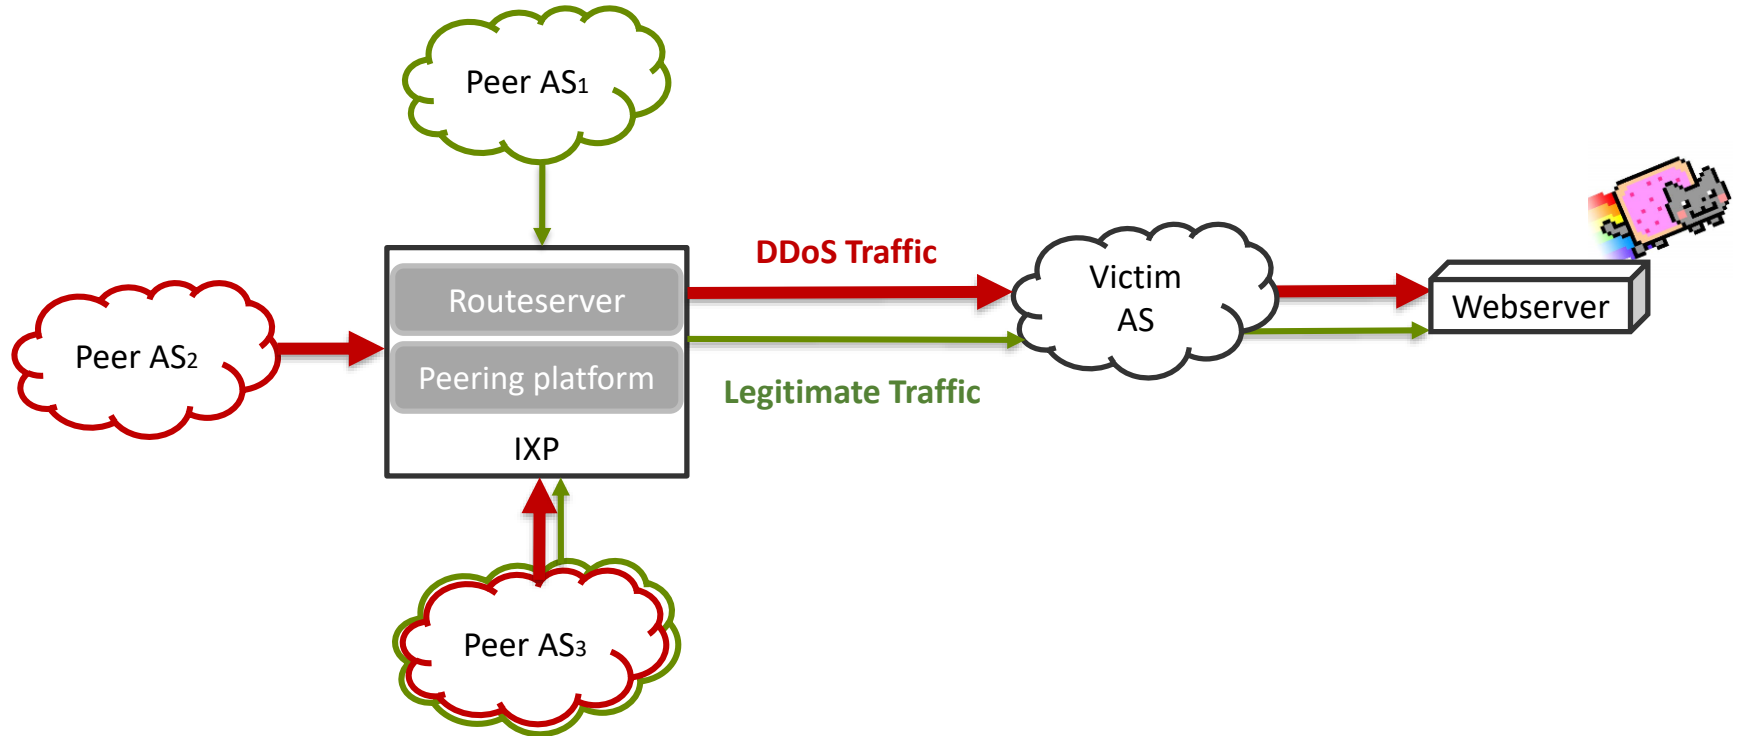
https://en.wikipedia.org/wiki/Black_hole#/media/File:Black_hole_-_Messier_87_crop_max_res.jpg

How does BGP Blackholing work at IXPs?

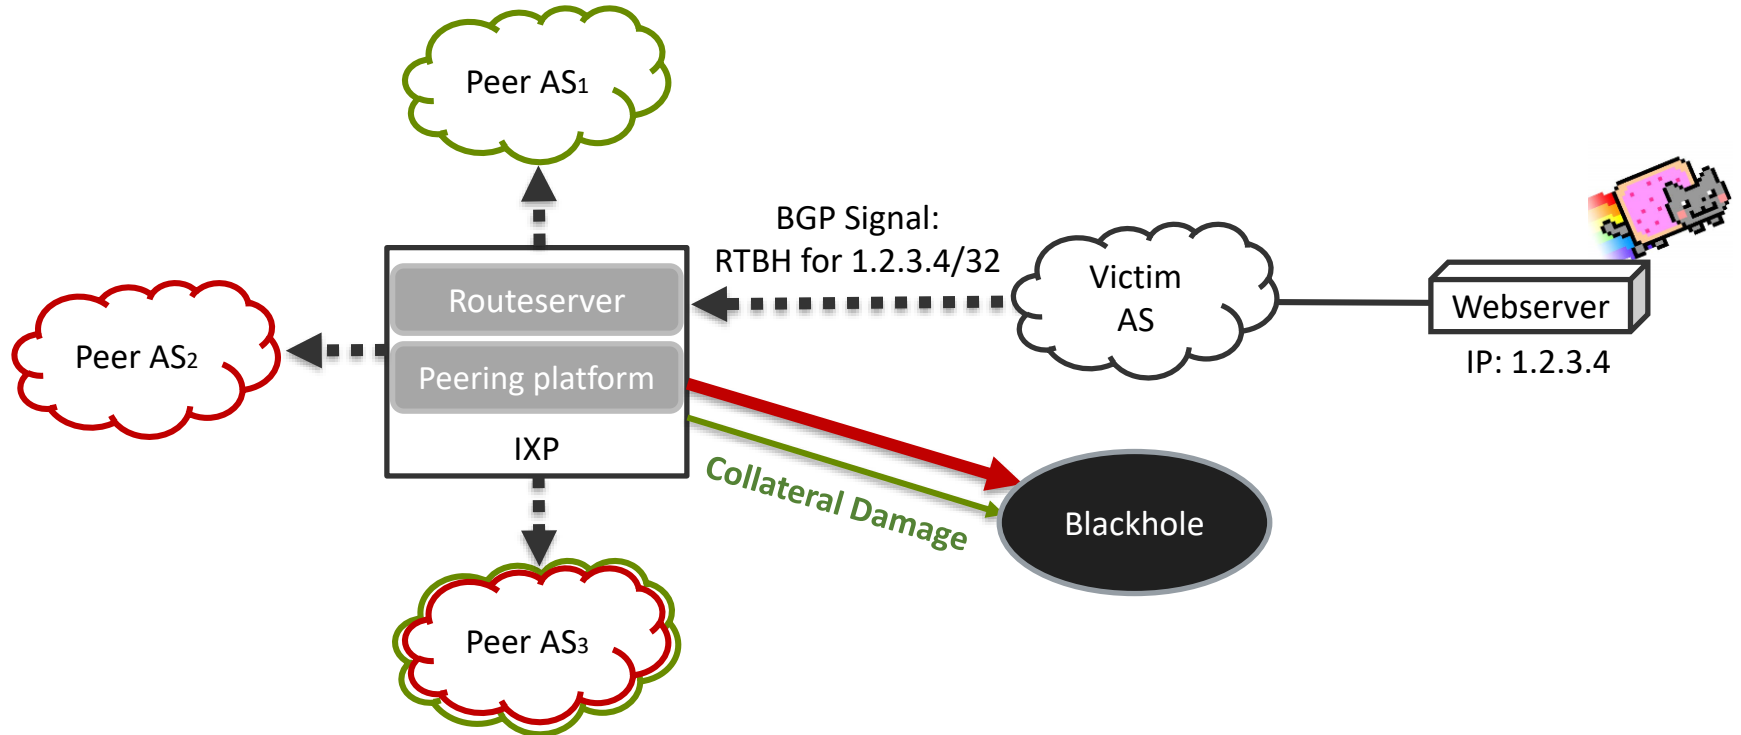
Remotely-Triggered Blackholing at IXPs



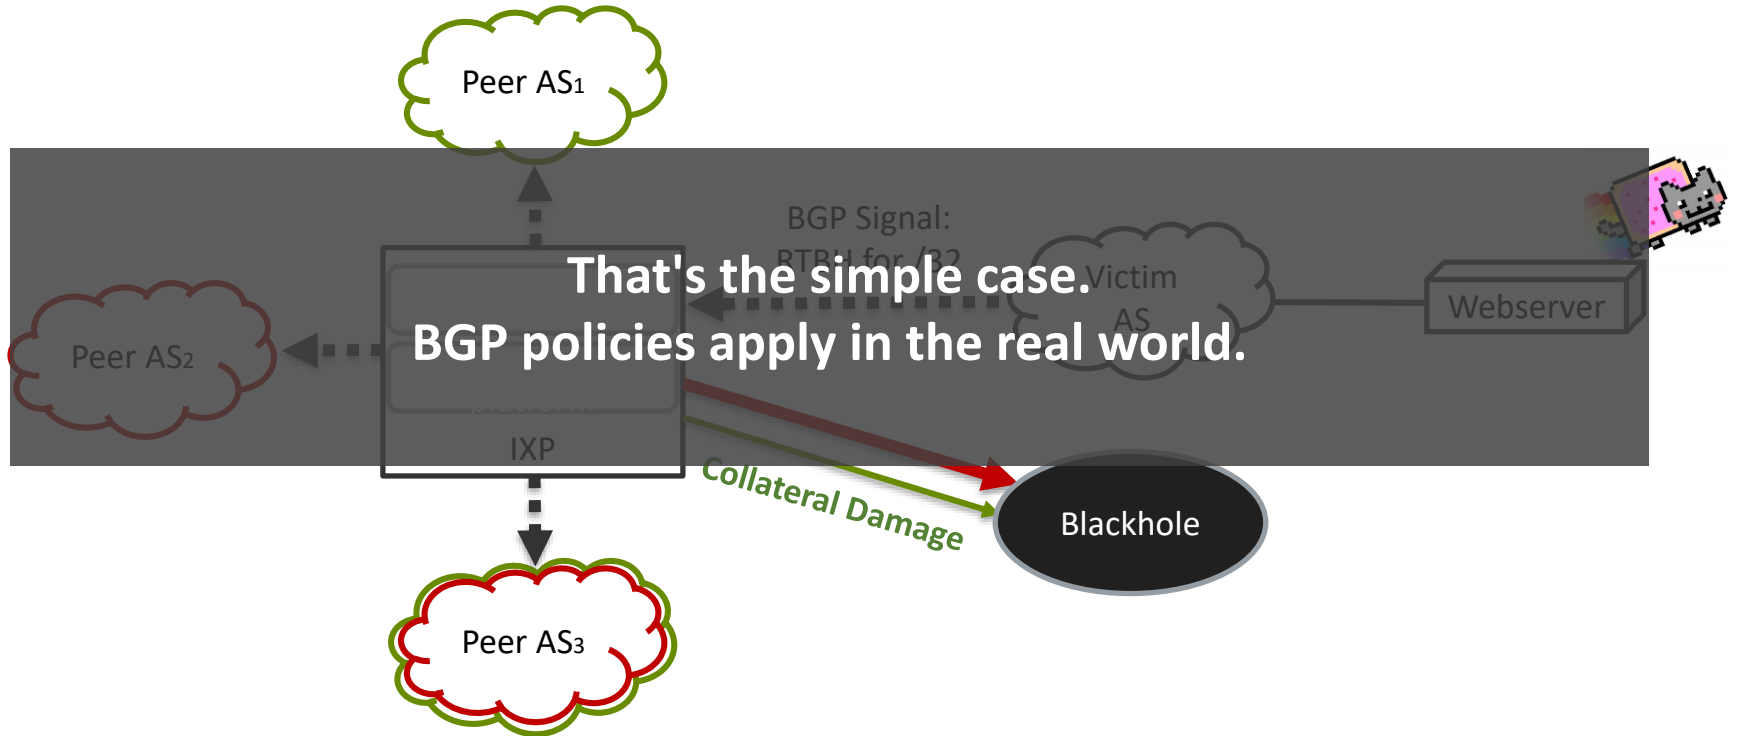
Remotely-Triggered Blackholing at IXPs



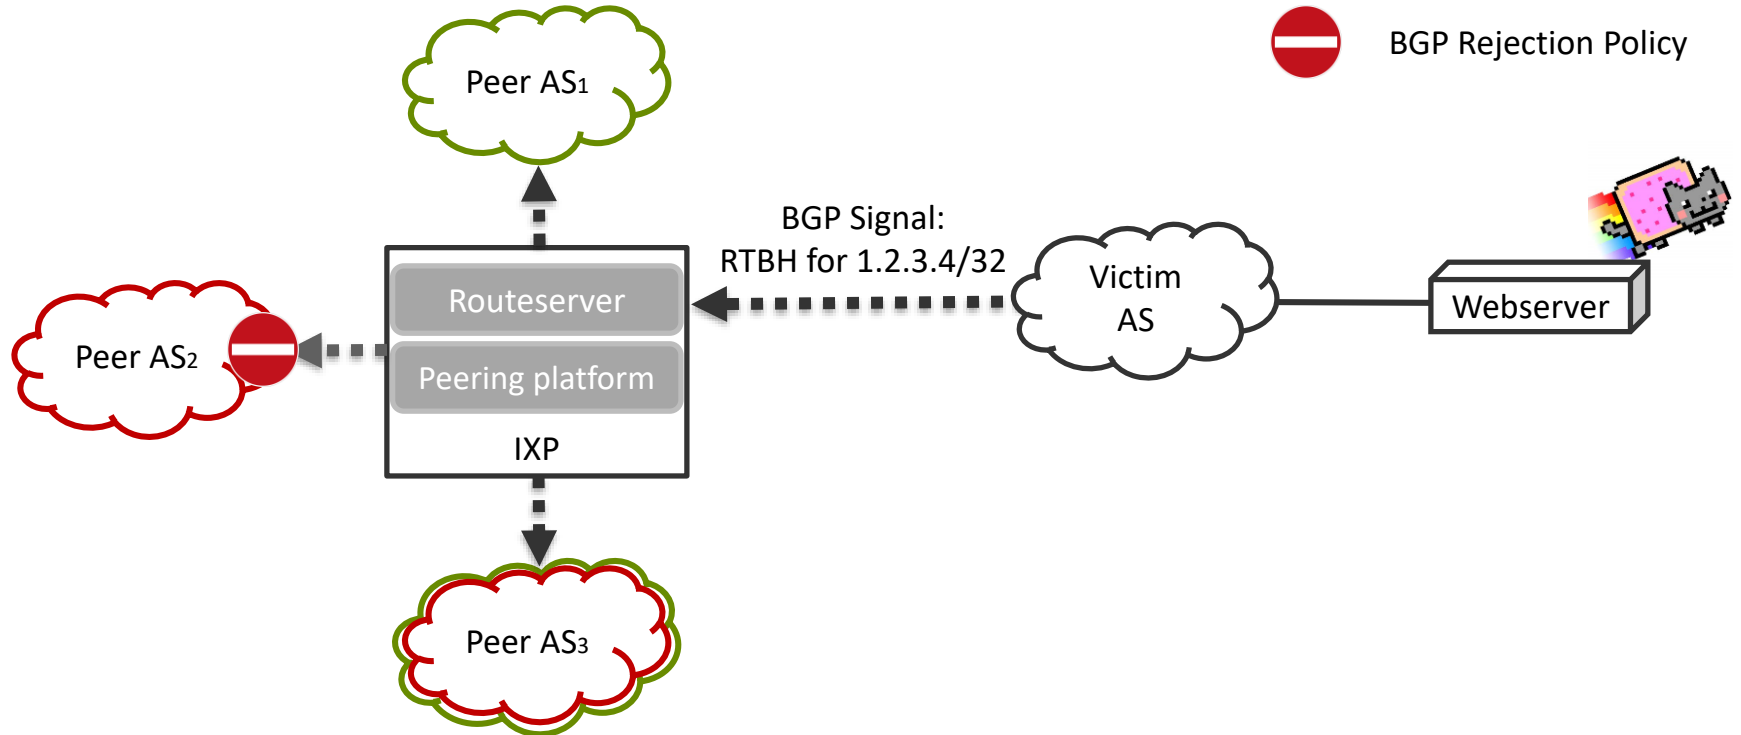
Remotely-Triggered Blackholing at IXPs



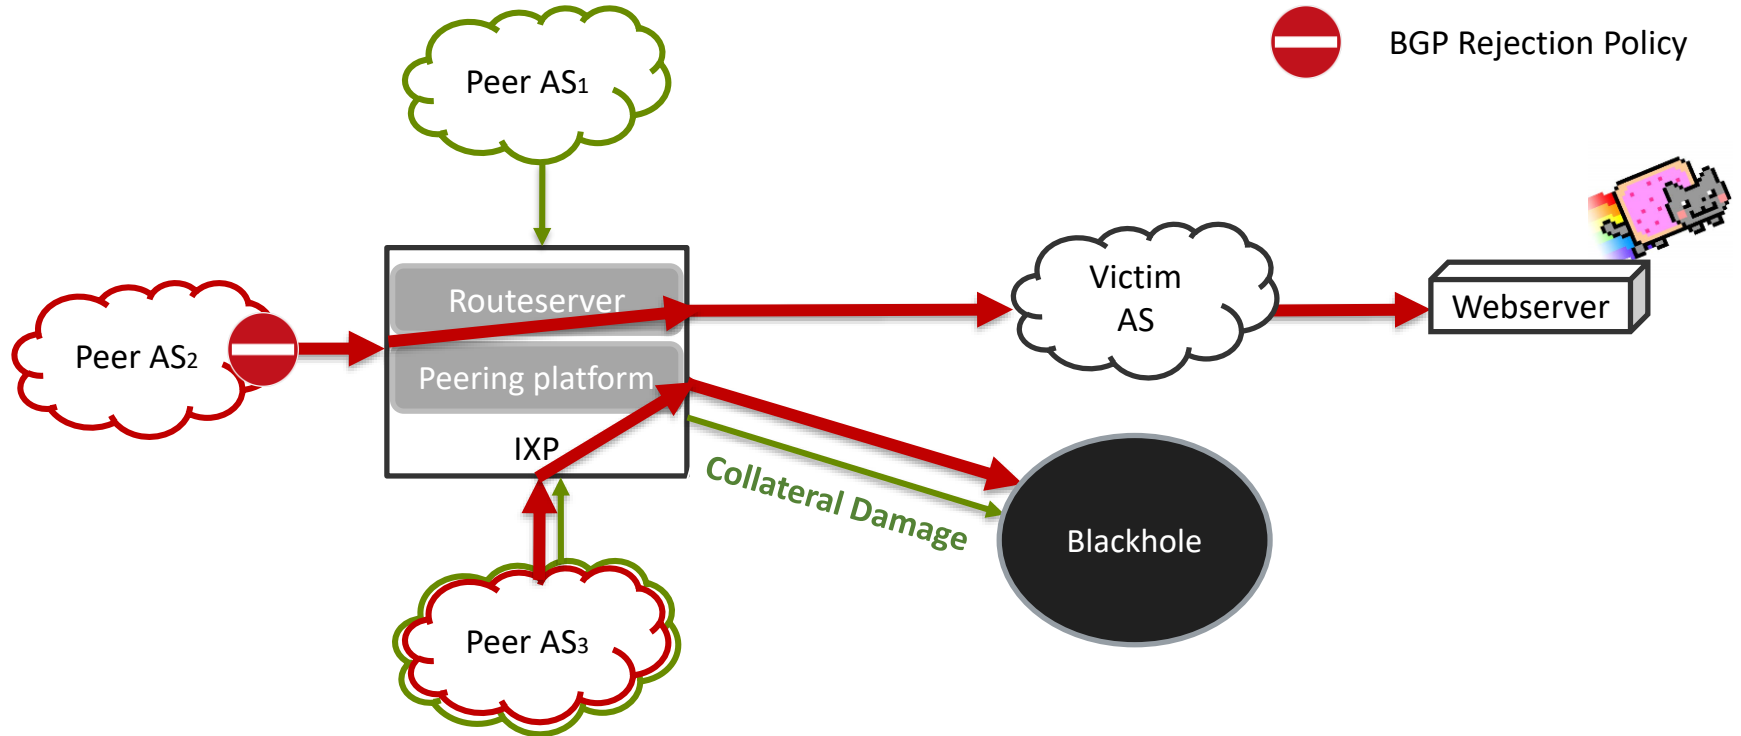
Remotely-Triggered Blackholing at IXPs



Remotely-Triggered Blackholing and BGP Policies




Remotely-Triggered Blackholing and BGP Policies



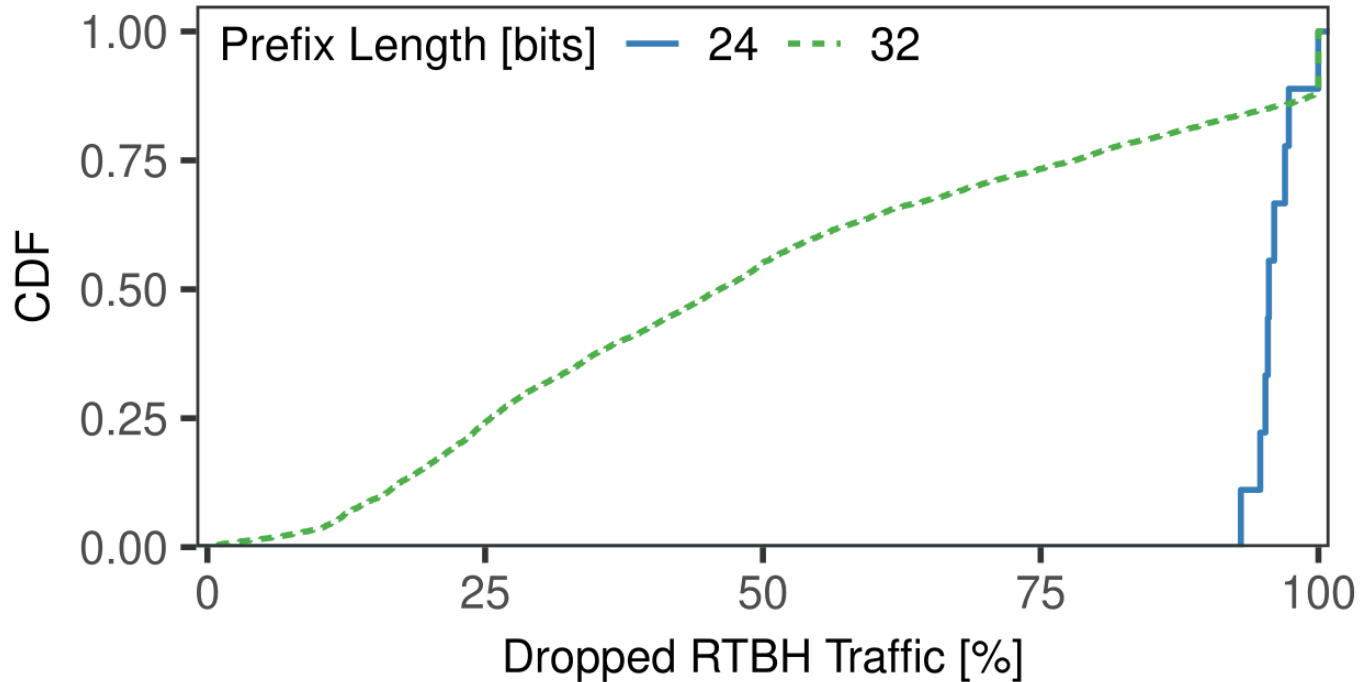
Mitigation practice

How effective is DDoS blackholing
at a large IXP?


Successful mitigation depends on the announced RTBH prefix length



Successful mitigation depends on the announced RTBH prefix length



Successful mitigation depends on the announced RTBH prefix length



Mitigation practice

How fast do IXP members react to DDoS events?

Analysis of **72 hours before** an RTBH Event

Use a sliding window algorithm (EWMA) to infer whether one of the **monitored features** exhibits an anomalous peak:

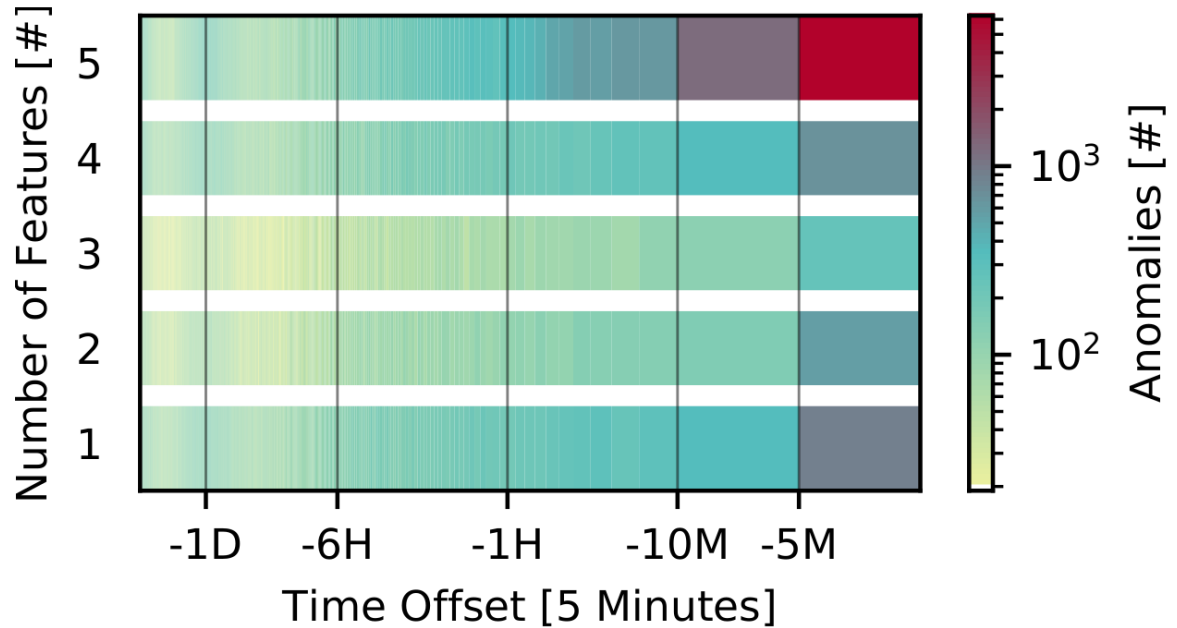
Amplification Attacks

TCP SYN Attacks

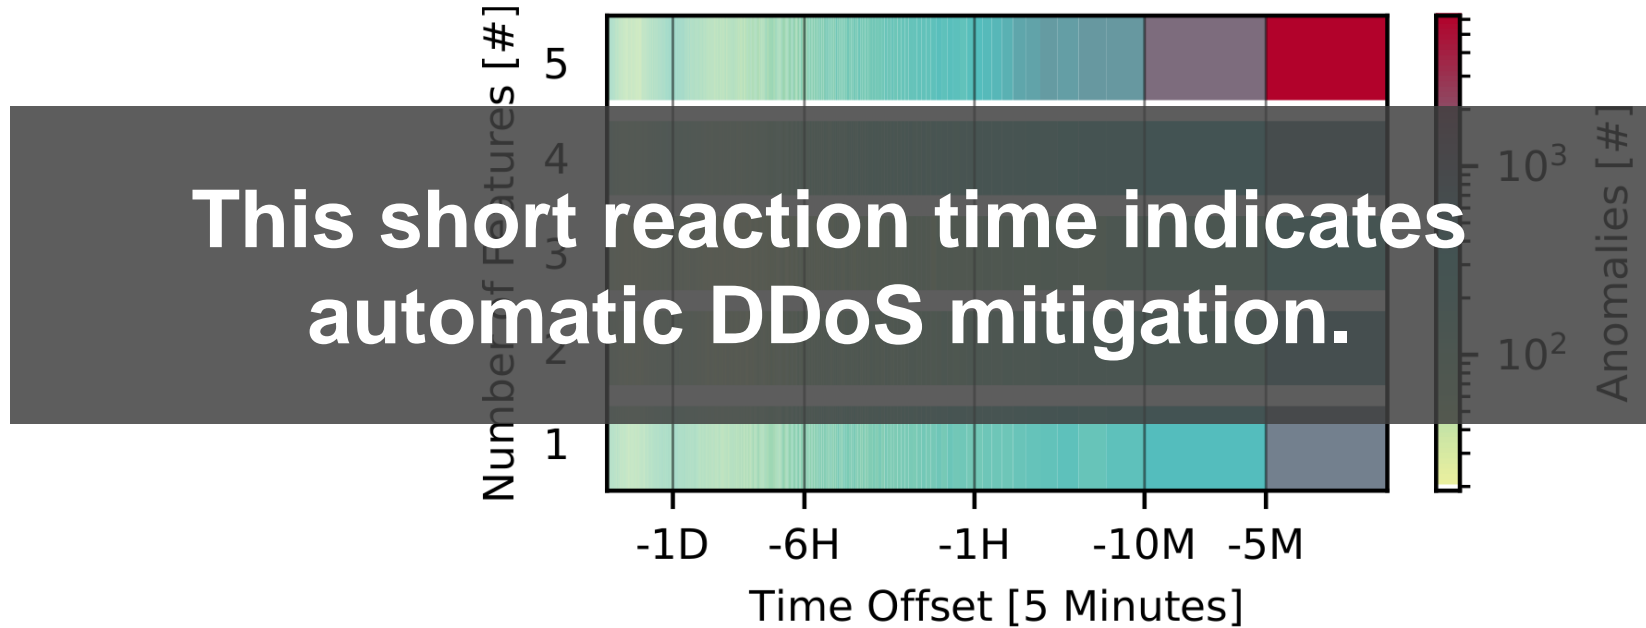
GRE Floods

- i. number of packets
- ii. number of unique destination ports
- iii. number of flows
- iv. number of unique source IP addresses
- v. number of non-TCP flows

Most anomalies occur up to 10 minutes before an RTBH Event



Most anomalies occur up to 10 minutes before an RTBH Event





Can we configure fine-grained filtering?

RTBH - Pro and Con

THE GOOD

RTBHs drop DDoS traffic early in the network.

THE UGLY

RTBHs complete the attack, the victim is unreachable.

RTBH - Pro and Con

THE GOOD

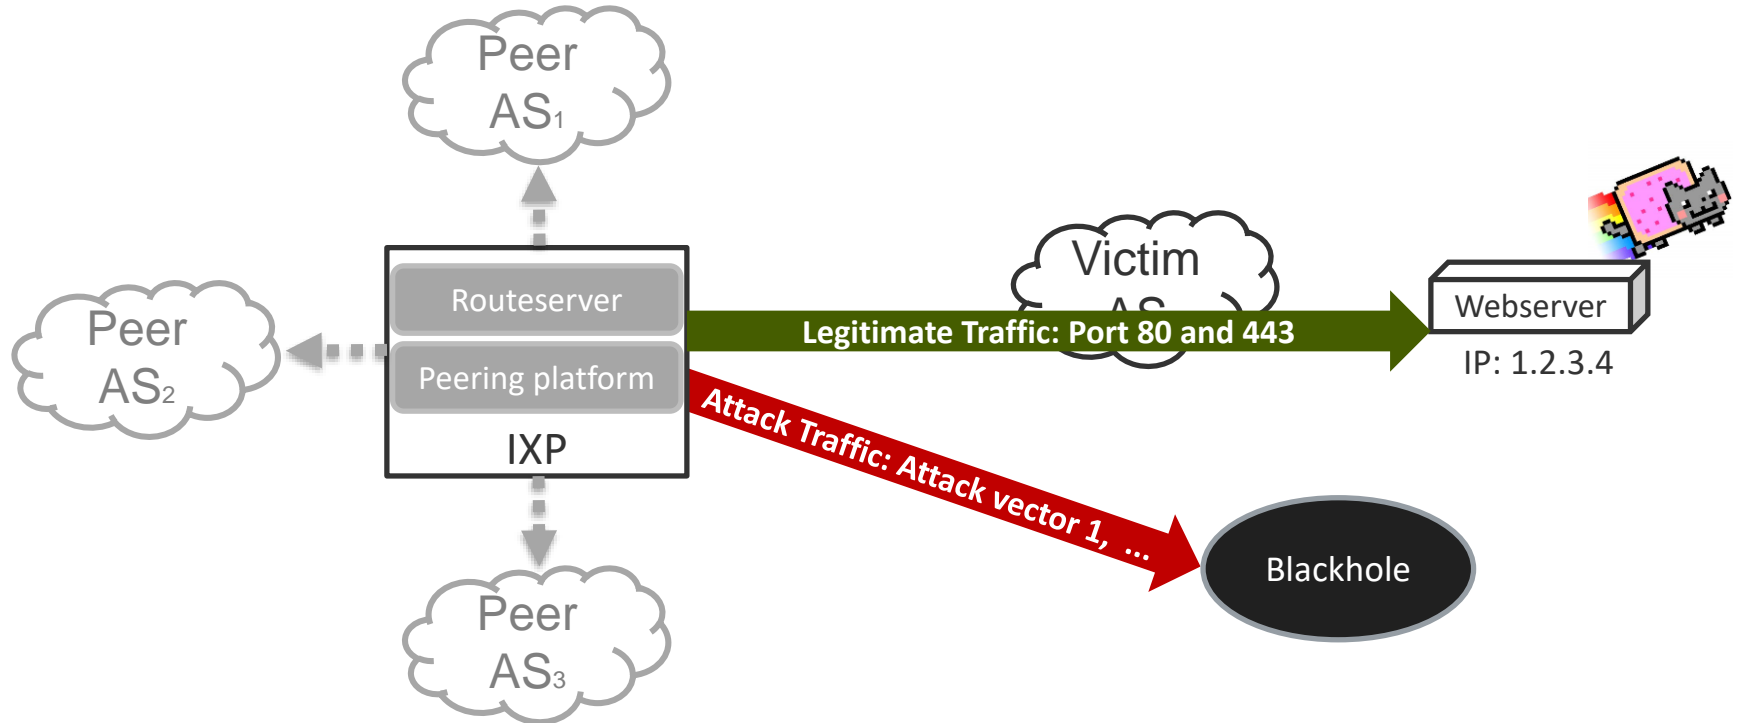
RTBHs drop DDoS traffic early in the network.

THE UGLY

RTBHs complete the attack, the victim is unreachable.

Fine-grained filtering would keep a service reachable.

Whitelisting vs. blacklisting of ports



Challenge

We cannot whitelist client traffic, because client traffic is highly variable.


RadViz Projection

Visualizing
multidimensional
port information allows a
classification into clients
and servers



RadViz Projection

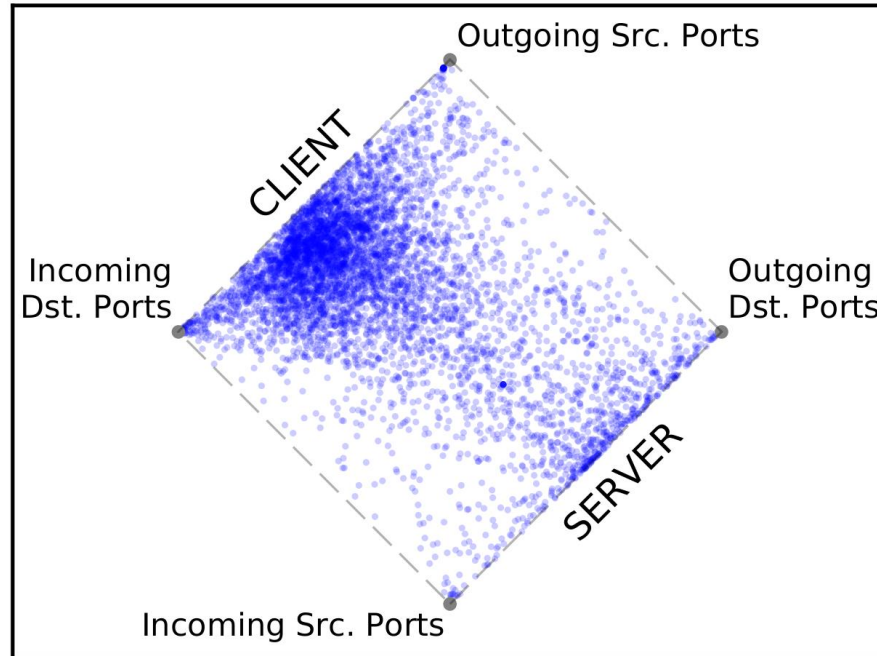
Visualizing multidimensional port information allows a classification into clients and servers



**number of
different
destination
ports**


**number of different
source ports**

Many blackholed IP addresses exhibit high port fluctuations



Many blackholed IP addresses exhibit high port fluctuations

Most of the protected IP addresses are clients.



Cross-validation using PeeringDB

Type	Clients	Server
# Hosts	4057	1036
Content	2%	34%
Cable/DSL/ISP	60%	14%
NSP	14%	13%
Enterprise	1%	1%
Unknown	23%	38%

Cross-validation using PeeringDB

**Most clients located in DSL networks.
PeeringDB supports our classification.**

Type	Clients	Server
# Types	4057	1036
Content	2%	34%
Cable/DSL/ISP	66%	14%
NSP	14%	13%
Enterprise	1%	1%
Unknown	23%	38%

Potentials of fine-grained whitelisting?

Clients are often affected by BGP Blackholing.

Whitelisting of regular, expected traffic patterns **is not an option.**

Fine-Grained Blacklisting

Fine-grained filtering based on source-ports is very effective and potentially saves legitimate traffic!

Filter example: CharGEN/19, DNS/53, NTP/123

Literature

Marcin Nawrocki, Jeremias Blendin, Christoph Dietzel,
Thomas C. Schmidt, Matthias Wählisch,
[Down the Black Hole: Dismantling Operational
Practices of BGP Blackholing at IXPs,](#)
In: *Proc. of ACM SIGCOMM Internet Measurement
Conference (IMC)*, p. 435–448, Amsterdam, 2019.
<https://doi.org/10.1145/3355369.3355593>

Down the Black Hole: Dismantling Operational Practices of BGP Blackholing at IXPs

Marcin Nawrocki
marcin.nawrocki@fu-berlin.de
Freie Universität Berlin
Germany

Jeremias Blendin
jeremias.blendin@de-cix.net
DE-CIX
Germany

Christoph Dietzel
christoph@mpi-inf.mpg.de
DE-CIX / MPI for Informatics
Germany

Thomas C. Schmidt
t.schmidt@haw-hamburg.de
HAW Hamburg
Germany

Matthias Wählisch
m.waehlich@fu-berlin.de
Freie Universität Berlin
Germany

ABSTRACT

Large Distributed Denial-of-Service (DDoS) attacks pose a major threat not only to end systems but also to the Internet infrastructure as a whole. Remote Triggered Black Hole filtering (RTBH) has been established as a tool to mitigate inter-domain DDoS attacks by discarding unwanted traffic early in the network, e.g., at Internet eXchange Points (IXPs). As of today, little is known about the kind and effectiveness of its use, and about the need for more fine-grained filtering.

In this paper, we present the first in-depth statistical analysis of all RTBH events at a large European IXP by correlating measurements of the data and the control plane for a period of 104 days. We identify a surprising practice that significantly deviates from the expected mitigation use patterns. First, we show that only one third of all 34k visible RTBH events correlate with indicators of DDoS attacks. Second, we witness over 2000 blackhole events announced for prefixes not of servers but of clients situated in DSL networks. Third, we find that blackholing on average causes dropping of only 50% of the unwanted traffic and is hence a much less reliable tool for mitigating DDoS attacks than expected. Our analysis gives also rise to first estimates of the collateral damage caused by RTBH-based DDoS mitigation.

CCS CONCEPTS

• Security and privacy → Denial-of-service attacks; • Networks → Public Internet.

KEYWORDS

DDoS, BGP, IXPs, RTBH

ACM Reference Format:

Marcin Nawrocki, Jeremias Blendin, Christoph Dietzel, Thomas C. Schmidt, and Matthias Wählisch. 2019. Down the Black Hole: Dismantling Operational Practices of BGP Blackholing at IXPs. In *Internet Measurement Conference (IMC '19)*, October 21–23, 2019, Amsterdam, Netherlands. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3355369.3355593>

1 INTRODUCTION

The Border Gateway Protocol (BGP) is used to exchange IP prefix reachability information between Autonomous Systems (ASes) to form the global Internet. Yet, one BGP application has the opposite effect in practice: Signaling Remotely Triggered Black Hole filtering (RTBH) through BGP requests a neighboring AS to discard traffic destined towards an owned IP prefix. The most prominent and well-established use case for RTBH filtering is the mitigation of volumetric Distributed Denial-of-Service (DDoS) attacks. Recent attacks peak beyond multiple Tbps (Terabit per second) [23]. DDoS attacks build upon simple to exploit IP address spoofing [??] in combination with amplification characteristics of network protocols such as NTP, DNS, or cLDAP [4, 12]. These attacks deplete network bandwidth to suppress legitimate traffic towards a destination IP. In consequence, a network or web service is not reachable anymore. Still, DDoS attacks do not only cause damage at the attacked system itself, but can also overwhelm the infrastructure of intermediate or upstream networks [31]. Such collateral damage often impairs common customers badly.

Intermediate ASes mitigate the collateral damage of DDoS traffic passing through their infrastructure by signaling RTBHs to their neighbors that specifically cover the target address of the DDoS attack. Thereby, volumetric attack traffic is dropped before it