

The Age of DDoSDiscovery: An Empirical Comparison of Industry and Academic DDoS Assessments

Raphael Hiesgen
HAW Hamburg
Hamburg, Germany

Marcin Nawrocki
NETSCOUT
Westford, MA, USA

Marinho Barcellos
U of Waikato
Hamilton, New Zealand

Daniel Kopp
DE-CIX
Frankfurt am Main, Germany

Oliver Hohlfeld
University of Kassel
Kassel, Germany

Echo Chan
Akamai/Hong Kong PolyU
Hong Kong, China

Roland Dobbins
NETSCOUT
Westford, MA, USA

Christian Doerr
Hasso Plattner Institute
Potsdam, Germany

Christian Rossow
CISPA
Saarbrücken, Germany

Daniel R. Thomas
University of Strathclyde
Glasgow, Scotland

Mattijs Jonker
University of Twente
Enschede, The Netherlands

Ricky Mok
CAIDA/UC San Diego
La Jolla, CA, USA

Xiapu Luo
Hong Kong PolyU
Hong Kong, China

John Kristoff
NETSCOUT/UIC
Westford, MA, USA

Thomas C. Schmidt
HAW Hamburg
Hamburg, Germany

Matthias Wählisch
TU Dresden
Dresden, Germany

kc claffy
CAIDA/UC San Diego
La Jolla, CA, USA

Abstract

Motivated by the impressive but diffuse scope of DDoS research and reporting, we undertake a multistakeholder (joint industry-academic) analysis to seek convergence across the best available macroscopic views of the relative trends in two dominant classes of attacks – direct-path attacks and reflection-amplification attacks. We first analyze 24 industry reports to extract trends and (in)consistencies across observations by commercial stakeholders in 2022. We then analyze ten data sets spanning industry and academic sources, across four years (2019-2023), to find and explain discrepancies based on data sources, vantage points, methods, and parameters. Our method includes a new approach: we share an aggregated list of DDoS targets with industry players who return the results of joining this list with their proprietary data sources to reveal gaps in visibility of the academic data sources. We use academic data sources to explore an industry-reported relative drop in spoofed reflection-amplification attacks in 2021-2022. Our study illustrates the value, but also the challenge, in independent validation of security-related properties of Internet infrastructure. Finally, we reflect on opportunities to facilitate greater common understanding

of the DDoS landscape. We hope our results inform not only future academic and industry pursuits but also emerging policy efforts to reduce systemic Internet security vulnerabilities.

CCS Concepts

• **Networks** → Denial-of-service attacks; **Network measurement**; • **Social and professional topics** → Governmental regulations.

Keywords

DDoS; Reflection-Amplification Attacks; Direct-Path Attacks

ACM Reference Format:

Raphael Hiesgen, Marcin Nawrocki, Marinho Barcellos, Daniel Kopp, Oliver Hohlfeld, Echo Chan, Roland Dobbins, Christian Doerr, Christian Rossow, Daniel R. Thomas, Mattijs Jonker, Ricky Mok, Xiapu Luo, John Kristoff, Thomas C. Schmidt, Matthias Wählisch, and kc claffy. 2024. The Age of DDoSDiscovery: An Empirical Comparison of Industry and Academic DDoS Assessments. In *Proceedings of the 2024 ACM Internet Measurement Conference (IMC '24)*, November 4–6, 2024, Madrid, Spain. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3646547.3688451>



This work is licensed under a Creative Commons Attribution International 4.0 License.

IMC '24, November 4–6, 2024, Madrid, Spain
© 2024 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-0592-2/24/11
<https://doi.org/10.1145/3646547.3688451>

1 Introduction

Distributed Denial-of-Service (DDoS) attacks were first reported around 2000 [22, 143] and continue to cause substantial damage, with cycles of new attack strategies and novel mitigation approaches. While hundreds of scientific studies and proposals have provided

academic perspectives (e.g., [66, 115, 117, 155, 156]), the more impactful developments have been commercial, where the need to mitigate the harms of DDoS led to the creation of a vibrant DDoS mitigation market worth over US\$ 1.5 Billion [60] and multiple practical attempts to deter a broad range of attacks [52, 58, 96, 160]. The concentration of content services among a few heavily provisioned network infrastructures has also provided some protection against the threat of DDoS, at the cost of strengthening their oligopolies. This combination of industry forces that benefit from DoS prevalence has arguably reduced the motivation for collective action to remediate the underlying DDoS threat and its root causes.

While academic projects have attempted longitudinal analysis of DDoS trends, gaining a consensus view of the state of the DDoS landscape has proven elusive. We undertake an extended multi-stakeholder analysis to pursue such consensus. We focus on direct-path attacks and reflection-amplification attacks, two dominant classes of attacks. We define each attack class and approaches to detecting and mitigating them (§2). We analyze the state of industry reporting on this topic, reviewing 24 reports to extract trends and inconsistencies across them (§3). We then describe the range of (raw and derivative) data sources available and challenges in comparing them, confirming that different detection approaches, and even the same approach using different parameters and vantage points, will yield different inferences of attack scope, duration, and impact (§4).

We analyze ten data sets covering a four years' period (2019–2023) to explore discrepancies based on data sources, vantage points, methods, and parameters (§6, §7). In the process, we use academic data sources to explore an industry-reported relative drop in spoofed reflection-amplification attacks following a concerted industry effort to encourage deployment of source address validation (SAV). However, we find more differences than similarities across data sets. Table 1 summarizes partial inconsistencies visible across various DDoS observatories used in this paper, and also among industry reports (from \approx 2022). Our work reinforces findings of [117] that singular data sources may have serious visibility limitations, which provide the strongest empirical grounding to date for regulatory framing to share data. Our four contributions are:

- (1) We taxonomize information extracted from industry reports characterizing DDoS phenomena in 2022-2023, which we publish as supplementary knowledge base, including an archive of the reports (§3).
- (2) We quantitatively compare ten data sources over four years, spanning honeypots, IXPs, and edge networks, including industry and academic vantage points. To our knowledge this is the largest correlation of longitudinal DDoS data ever published (§5,6).
- (3) We propose and execute a new approach to facilitating a degree of industry transparency, by aggregating academic sources and sharing them with industry players who then return the results of joining these shared data sets with their proprietary data sources to indicate gaps in visibility of the academic data (§7).
- (4) We propose several recommendations to facilitate scientific study of the DDoS landscape and of whether proposed mitigations are effective. We introduce possible self-regulatory

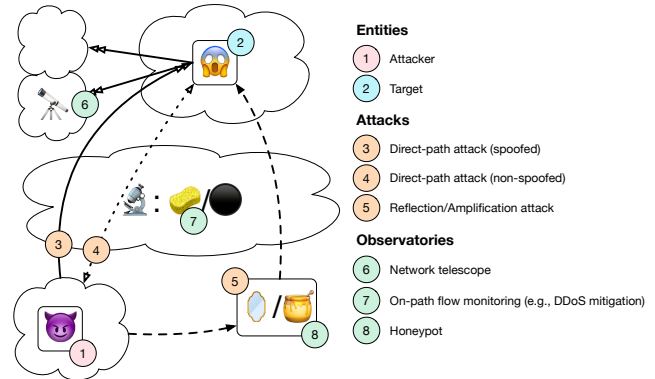


Figure 1: Three DDoS attack types: Direct-path spoofed (solid line), direct-path non-spoofed (dotted line), and typically spoofed reflection-amplification (dashed).

approaches, other potential regulatory developments, and roles for researchers (§9).

Our empirical study clearly shows that the assessment of DDoS trends is challenging and that collaboration between research and industry is needed to gain sound insights. Because regulators are now showing vigorous interest in policy intervention to reduce systemic Internet security vulnerabilities [50, 51], we hope that the results of this paper help inform and guide not only academic and industry efforts but also future policy decisions.

2 Definitions, Detection, Defense

We describe the two most prevalent classes of DDoS attacks—(i) direct-path (spoofed and non-spoofed) attacks and (ii) reflection-amplification attacks—and methodologies to observe them. For a broader classification of DDoS attacks, we refer to prior work [46, 102, 105, 163].

2.1 DDoS Attack Models

Figure 1 illustrates direct-path and reflection-amplification attacks. In either cases, an attacker ① aims to overwhelm a target host or target service ②, or saturate its uplink.

Direct-path attacks. Attackers send packets directly to the target. If the source address is *spoofed* [14] ③, the target sends responses to the hosts with the spoofed addresses. An example is the well-known *SYN-flood attack* [48], where an attacker sends TCP SYN packets, each of which induces a memory allocation related to the TCP connection. The spoofed source addresses are often chosen randomly, leading to the term *randomly-spoofed DDoS (RSDoS) attacks*.

Another type of direct-path attack uses *non-spoofed* source addresses ④ to establish many sustained connections with a server [40]. This state exhaustion attack also minimizes impact on the sending network and visibility of the attack.

Reflection-amplification attacks. The attacker indirectly sends traffic to a target via a *reflector* ⑤. The reflector typically produces large responses to small requests, i.e., it *amplifies* [155]. Amplifiers allow attackers to subject their target to massive amounts of traffic

Table 1: Data comparison results: Partially inconsistent views among DDoS data observatories used in this paper measuring decreasing ▼ (< -5% in 4 years), increasing ▲ (> 5% in 4 years), and steady ◆ trends of attack types in 2019–2023. The surveyed industry reports from ≈ 2022, which usually compare relative share of attacks, similarly provide inconsistent views. Here, numbers in braces indicate the number of reports out of 24 surveyed reports.

Attack Type	Observatories Used in This Paper (2019-2023)								Industry Reports (#) (≈ 2022)
	Network Telescopes		Flow Data			Honeypots			
	UCSD	Orion	Netscout	Akamai	IXP	Hopscotch	AmpPot	NewKid	
Direct-path	▲	▲	▲	◆	▲	n/a	n/a	n/a	▲(5), ▼(0)
Reflection-Ampl.	n/a	n/a	▲	◆	▼	▼	◆	▲	▲(2), ▼(3)

by sourcing much less traffic from their own network, reducing the impact as well as likelihood of detection of the attack.

Enabling platforms. Attackers today frequently rent facilities on *botnets* or dedicated infrastructure, both of which are robust to take-downs and hide the attacker’s identity. *Bullet-proof hosters* (BPH) are service providers that avoid responding to law enforcement requests and are lenient with acceptable use. *Booter* services perform DDoS attacks for a fee. They are surprisingly resilient to takedowns or after takedown often return shortly on a new website [31, 83].

2.2 DDoS Observatories

We describe three types of measurement that allow inferences about DDoS attacks: network telescopes, flow monitoring, and honeypots.

***Network telescopes (NT)** ⑥ are passive measurement platforms that collect packets sent to large blocks of unused IP space. Network telescopes achieve visibility of attack preparation in the form of scans for open reflectors or vulnerable hosts. Telescopes also collect artifacts of certain types of attack execution, *i.e.*, replies to randomly spoofed packets in RSDoS attacks. By identifying typical response packets and thresholds, network telescopes enable inference of scope, prevalence, and duration of RSDoS attacks [76, 107]. Network telescopes do not generally observe evidence of reflection-amplification attacks, since the spoofed address in the attack traffic is not random, but rather the address of the intended target. The attack traffic is also destined to the reflector, rather than broadly toward the Internet where a telescope might observe it.

***On-path flow monitoring by DDoS mitigation providers.** Observing DDoS traffic toward victims requires a vantage point ⑦ on the path to the target. DDoS identification can rely on manual inference, deep-packet inspection, aggregate packet- or flow-level statistics [82], or more complex machine learning approaches [177]. IXPs, CDNs, or specialized DDoS protection service (DPS) providers may protect their customers by *scrubbing* ◆ or *blackholing* ● attack traffic (§2.3). These vantage points directly observe ongoing attacks, but access to such data is limited to the network owner or commercial DDoS mitigators serving that network.

***Honeypots (HP)** ⑧ emulate a vulnerable host to learn more about the behavior of an attacker [117, 144]. The level at which a honeypot interacts with a presumed attacker ranges from one packet to full compromise and access to a service. To support the study of DDoS, honeypots may try to appear as reflectors for common protocol vectors, *e.g.*, DNS or NTP. To avoid participating in attacks, honeypots stop engaging with a probing source after some

sending threshold is reached. Several honeypot platforms have been operational for years (§6), *e.g.*, AmpPot [84], AmpPotMod [135], Hopscotch [167], NewKid [68], and HPI [66].

2.3 Prevention and Mitigation of DDoS

We review approaches to prevention and mitigation of the two DDoS types we study: disabling amplifiers, anti-spoofing campaigns, booter takedowns, and filtering of attacks.

Prevention: Disabling reflection-amplification vectors. Academic and industry efforts to identify and decommission open servers that support reflection and amplification [90, 116, 123, 126, 161] have had limited success. Some services are easier to decommission than others, *e.g.*, operators of NTP servers can disable a specific command that enables extraordinary amplification (`get monlist`) but is not of operational use. DNS servers are less amenable to such curation of function. DNS operators must consider more complex configuration changes, such as rate limiting, filtering (*e.g.*, ANY requests), or truncating large responses. The result is a long-standing persistence of amplification vectors [82, 116].¹

Prevention: Promotion of source address validation (SAV). Operators have pursued efforts to reduce the number of networks that allow source IP address spoofing, the basis of all spoofed DDoS attacks. One such effort is the Spoofer measurement project [96], which identifies networks that allow spoofing and assists with remediation of this vulnerability. This project relies on users to download software from CAIDA’s website and launch it in the background on their laptop; the software tests each new network visited for the ability to spoof. This volunteer crowdsourced approach yields limited measurement coverage.

For many years groups have undertaken various efforts to eliminate sources of spoofing in networks [47]. DDoS mitigation providers reported a successful concerted effort since 2021 by the global Internet operational community to reduce spoofing [9, 127, 133].² Netscout reported a 17% decrease in reflection-amplification attacks (which leverage spoofable networks) in 2022 compared with 2021, in their view a direct result of this concerted effort [128].³ Our

¹Authors of [116] continue measuring the prevalence of transparent DNS forwarders, showing a drop in mid-2023. <https://odns.secnw.net/data>

²“...the lower global backbone impact was largely due to an industry wide antispoofing initiative – the DDoS Traceback Working Group.” [133], *s.a.* [101].

³“In 2022 [...] a momentous 17 percent global decrease in reflection/amplification attacks was observed when compared with 2021.” [128].

observatories also saw drops between 2021 and 2022, of varying magnitudes (see Figure 3 and discussion in §6.2).

Prevention: Takedown of botter services by law enforcement. Other efforts by academia, industry, and law enforcement have focused on shutting down DDoS-for-hire services [18, 31, 49, 83]. Booters often reappear within a few months under different domains, seemingly leaving little long-term reduction from the effort [31]. Our analysis of observed reflection-amplification attacks (§6.2) shows no lasting downward trend subsequent to recent publicized large-scale takedowns [18, 138, 171]. Law enforcement have also used targeted messaging campaigns (Google search result ads for keywords like ‘botter’) to raise awareness that these websites offer illegal services. This approach may be a cheap and effective way of reducing DDoS [31], but its effectiveness remains unclear [106], as do the ethics of this approach [30].

Mitigation: Filtering attack traffic. Filtering ongoing attacks at the victim’s network allows for a close loop between detection and mitigation but limits the scope of mitigation since the victim network must still receive attack traffic. A more effective approach is a *scrubbing service*, where a third party (e.g., IXP, CDN, or DPS) uses deep-packet inspection or application proxies to identify and block DDoS attacks at network/application layers and forward the sanitized traffic to their customer’s networks [45, 78, 159, 169, 177]. A coarser-grained approach is remote triggered black hole (RTBH) filtering [63, 77, 82, 92, 113], where a target (victim) remotely triggers the dropping of traffic to a whole IP prefix when one or more addresses in that prefix is under a DDoS attack. Blackholing risks collateral damage [77, 113].

Mitigation: Standardization efforts to support cooperative filtering. Standards for DDoS defense and prevention are partially documented as IETF best current practice (BCP) [12, 165], informational RFCs [23, 92], and standards-track RFCs [3, 20]. Other proposed standards with weaker operational roots did not gain traction [67, 74, 108, 150]. Operator groups have published their own BCPs [100], and two ISPs presented a bilateral *DDoS peering* framework to allow mutual filtering of DDoS traffic between peers [131, 152]. In 2020, DE-CIX proposed a technical and governance framework to facilitate DDoS-related data sharing among Internet Exchange Points (IXPs) [176]. The Netherlands Anti-DDoS Coalition (ADC) has pursued a similar concept at national scale [8]. Team Cymru has created a service to facilitate relaying of destination-based remote triggered black holing (RTBH) signals between ASes [39, 75]. Deployment, scaling, and sustaining such collaborative efforts have proven challenging [7].

3 Analyzing DDoS Industry Reports

The DDoS mitigation industry publishes reports about the state of DDoS, identifying trends and alerting decision-makers about the need to deploy appropriate DDoS protection. The primary purpose of these reports is to promote use of a DDoS mitigation approach. They also offer a glance at data usually not accessible to researchers. Since industry and academia have diverging views [117], we were motivated to take a close look at published threat reports. We dissected 24 reports of 22 vendors to contrast numbers and trends, laying a foundation to compare industry with academic perspectives. We undertook and now publish a thorough *structural analysis*

of these reports as a supplementary artifact [13]; due to space constraints, we only summarize highlights here.

Our method to survey industry reports on DDoS. DDoS reporting from industry is fragmented, scattered among periodic reports, blogs, related educational resources, and talks. We limit our analysis to written content, which we call “reports”. We collected reports from companies listed in a related market survey [151] and excluded global threat reports without DDoS content (e.g., [38, 59, 141]), DDoS reports without DDoS data (e.g., [56, 153]), and DDoS assessments before 2022 (e.g., [11, 132]). We considered available reports from all major DDoS mitigation providers: A10 [1], Akamai [4], Arelion [9], Cloudflare [29], Comcast [33], Corero [34], DDoS-Guard [41, 42], F5 [53], Huawei [72], Imperva [73], Kaspersky [80], Link11 [95], Lumen [99], Microsoft Azure [166], NBIP [121], Netscout [124], NexusGuard [130], Nokia [134], NSFocus [136], Qrator [148], Radware [149], and Zayo [180]. Most reports were released early 2023 and focus on 2022. Appendix E provides details.

Presentation style. Industry reports are unlike scientific papers, typically using vague language and lacking clarity about data analysis methodologies. They vary substantially in format and organization, from full documents to web blogs to infographics. Some reports cover DDoS exclusively (e.g., [26, 34, 53, 72, 124]), while others report on a range of malicious activities (e.g., [33, 134, 148]). Technical depth spans from superficial trends to in-depth analyses that explain the vectors and methods to launch attacks. Not even the most detailed reports clearly explain the methodologies to identify attacks or discuss limitations of their analyses. Industry reports also do not contextualize the findings in terms of overall traffic patterns so that readers cannot judge whether attacks are growing in proportion with other properties, e.g., user base. Reports mix absolute and relative values, depending on the message to be emphasized. For example, a high increase (e.g., 500%) in some form of attack may actually represent a small absolute change [125].

What concerns us most is that some reports cherry-pick numbers to convince readers about the increase of DDoS attacks and the damage they cause. Most reports highlight the growing impact of DDoS attacks, but when the data suggests a decrease in severity, the message is less clear. Marketing concerns may lead to revision of reports prepared by technical staff, to present observations in a way that is more aligned with business interests.

Metrics used by reports. The reports we analyzed used a range of numbers to illustrate the DDoS attack landscape (in 2022) compared to previous periods. The attack attributes frequently reported by industry papers are: *count*: per period and attack type; *size*: peak packet rate, peak bandwidth, or attack volume; *duration*: in minutes/hours (e.g., “most attacks under 10 min”); *vectors*: protocol/packets used (e.g., TCP SYN and DNS amplification); *methods*: carpet-bombing [130], pulse-wave; *vector instances*: number of hosts that can send attack packets; and *context*, e.g., cyberwarfare or hacktivism. Some reports include information about the use of multi-vector attacks, attack repetition, use of botnets, targeted industries (e.g., finance sector, IT, education), and geolocation of attack sources and targets.

Analysis period. Most reports focus on one year, comparing with the previous year or sometimes a few years back if highlighting a trend. Generally, metrics show oscillations when the period used in

a report is a quarter or a month. Comparing short periods may be misleading, but may illustrate cyclic behaviors, *e.g.*, use of a vector every few months.

Comparing findings. Companies generally reported an overall increase in DDoS attacks. Among the exceptions, F5 indicated a decrease of 9.7% in total attacks [53], while Arelion reported a “dramatic” reduction of DDoS activity [9]. Arelion associated the drop with a decrease in UDP spoofed attacks, following actions of an “industry-wide anti-spoofing initiative” [9] (§2.3), and despite some increase in direct path attacks. Netscout also reported a drop in reflection-amplification attacks, and Akamai reported a decrease in CharGEN, SSDP and CLDAP-based attacks, typical amplification vectors. Several providers, namely Cloudflare, F5, Imperva, NBIP, Netscout, NexusGuard, and Radware, reported substantial increases in application-layer (L7) attacks, *e.g.*, via HTTP/S. Consistency across most reports is the dominance of UDP-based vectors, predominantly UDP flooding.

Summary. The reports we reviewed reflect a variety of data sources and analysis methods. Industry reports can provide directions for researchers to perform more scientific explorations, but they do not present or claim a scientific contribution. Given pecuniary interests and perspectives of vendors that publish reports, our community should consider them with care. Reported results can complement scientific studies but still do not provide a complete picture. Our analysis of these reports (see Appendix E and [13]) further inspired our development of a more rigorous framework for comparison.

4 Challenges in Comparing DDoS Data

Finding consensus on the DDoS landscape faces three obstacles: vantage point limitations; definitions and detection methods; and inhibitions on data sharing.

Vantage point limitations. Characteristics of observed DDoS attacks vary by observation point and method (§2.2). DDoS inference in the Internet core, *e.g.*, from flow data at IXPs, more likely underestimates attack length and volume of observed attacks, since some (or all) attack traffic may transit paths other than the IXP. In contrast, a vantage point at a DDoS protection service (DPS) or the Internet edge, *e.g.*, a victim’s network, observes only packets targeting or originating from the observed network(s). Measuring close to the target can lead to more accurate detection, but obstacles to sharing such data. Concerns regarding privacy or reputation may prevent its use for establishing a consensus view. Researchers have developed aggregated vantage points (honeypots and telescopes) that have lower obstacles to sharing data. Honeypots observe reflection-amplification attacks when an attacker selects the honeypot as a reflector for the attack. Telescopes observe RSDoS backscatter, and if monitoring a large enough segment of address space, they achieve high visibility of such attacks independent of target location.

Definition of attack. Distinguishing between natural traffic peaks and attack traffic is challenging. Honeypots need to discern scanning [84] and testing by attackers from actual attacks. No single definition can accurately capture characteristics of all attacks. Appropriate thresholds depend on attack type, protocol [117], and

observatory [82]. Some attacks require additional inference and sanitization steps, including aggregation across multiple sensors (§5).

Data sharing obstacles. Prohibitions on multi-lateral (vs bi-lateral) data sharing create obstacles to mitigation. Personal trust between individual members does not automatically translate to trust in procedures and governance of a larger collaborating group [61, 64]. The key challenge is establishing a data governance model and legal agreements to support it. Some established such agreements successfully [24, 174].

Longitudinal trend bias. We used normalized attack counts per week (§6), without considering growth in traffic, customers, or measurement coverage. Normalization reduces data sharing concerns and helps control for coverage variation among data sources.

5 Data Corpus Used in This Study

We analyze ten data sets from seven observatories: 2 network telescopes (🔭), 3 DDoS mitigation providers (🛡️), and 2 honeypots (🍷). Table 2 summarizes information about these observatories such as the types of attack they measure: direct path (DP), reflection-amplification (RA), or randomly-spoofed DoS (RSDoS), a subset of DP attacks (§2.2). We compare attack data across 4.5 years (2019 to mid-2023). To our knowledge this is the largest correlation of DDoS data sets not only including academic but also industry sources. Ethical considerations are discussed in Appendix A.












Each observatory captures different traffic, and thus may not see the same attack at the same intensity, or at all. This distinction is often a function of the vantage point (§4). Honeypots and telescopes are essentially end points in some portion of attack traffic, whereas industry (traffic flow-based) solutions sit somewhere on the network path, perhaps toward the target endpoint. Flow data will include not just attack-induced traffic but legitimate traffic which can offer a baseline for comparison and inference.

Another concern is that observatories might interfere with each other’s visibility. For example, an observed but quickly mitigated randomly-spoofed direct-path attack might not reflect packets into a network telescope. Partially mitigated attacks may affect the attack proportions inferred by other observatories (*e.g.*, length, volume).

Finally, the attack detection strategy, including threshold parameters, defines how an observatory identifies DDoS attack in its traffic. Lack of ground truth data on attacks prevents confident estimates of DDoS detection accuracy. Academic sensors may overestimate attacks, *e.g.*, if honeypots mistakenly interpret scans as attacks. In contrast, industry efforts to mitigate large attacks may miss short or low-volume attacks that nonetheless harm a target, *e.g.*, an uplink of an end user.

🔭 **Telescopes observe (backscatter from randomly-spoofed) direct path attacks.** We used data from two of the largest, longest-running IPv4 network telescopes: Merit’s *ORION project* [103] with ≈500k IPv4 addresses and UCSD’s *UCSD-NT* operated by CAIDA [21], spanning a lightly utilized /9 and /10 network, *i.e.*, ≈12M IPv4 addresses. Telescopes (🔭 in Figure 1) observe backscatter from direct path attacks that use randomly-spoofed source addresses. Assuming an approximately random distribution of spoofed IP addresses, larger telescopes will receive more attack traffic and can thus detect smaller attacks. Using the parameters suggested in [107], *UCSD-NT* and *ORION* can detect DDoS events with attack rates


Table 2: The observatories used in this research vary in collection methods and attack detection strategies. Honey pots use different flow identifiers, see [117]. (Location: Geographically & Topologically distribution.)


Platform	Type	Attack	Loc.	Coverage	Attack Definition		
					Flow Identifier	Timeout	Threshold
UCSD NT		RSDoS	US	12M IPs	protocol, src IP	300s	≥ 25 pkts, $\geq 60s^2$
ORION NT		RSDoS	US	500k IPs	protocol, src IP	300s	≥ 25 pkts, $\geq 60s^2$
Netscout Atlas (RA)		DP	G/T	proprietary	Hand-craft flow identifiers & thresholds		
Netscout Atlas (DP)		RA	G/T	proprietary	Hand-craft flow identifiers & thresholds		
Akamai Prolexic (RA)		RA	G/T	proprietary	Hand-craft flow identifiers & thresholds		
Akamai Prolexic (DP)		DP	G/T	proprietary	Hand-craft flow identifiers & thresholds		
IXP BH (RA)[82]		DP	G/T	proprietary	UDP, ampl. src port		≥ 10 IPs, > 1 Gbps
IXP BH (DP)[82]		RA	G/T	proprietary	TCP		≥ 10 IPs, > 100 Mbps
AmpPot [84]		RA	G/T	≈ 30 IPs	Src IP, src port, dst IP, dst port	60 min	≥ 100 pkts
Hopscotch [167]		RA	G/T	65 IPs	Src IP, dst IP, dst port	15 min	≥ 5 pkts
NewKid [68]		RA	BR	1 IP	Src prefix, dst IP, [dst port] ¹	1 min	≥ 5 pkts, $[\geq 2$ ports] ¹

¹ NewKid uses two thresholds, one for mono-(dst port) and for multi-protocol (≥ 2 ports) attacks. ² See Appendix J for RSDoS inference details.

of 0.026 Mbps and 0.60 Mbps in 5 minutes, respectively. (With the same assumptions, a /20 telescope could detect attacks of ~ 70 Mbps in 5 minutes.)

For both telescopes we had the raw traffic data available. We used the RSDoS-detection algorithm developed at CAIDA (based on [107]) on both data sets to identify attacks, which were the basis for our analysis. Appendix J has details of CAIDA’s current algorithm.

 **Monitored flow data can include both attack types.** We had access to DDoS attack counts from two DDoS mitigation providers (⑦ in Figure 1). Both observe traffic in an on-path network. The first data set, *IXP Blackholing*, contains daily counts of attacks identified for traffic that was blackholed by a European IXP (method in [82]). The second data set contains daily attack counts observed by Netscout, which receives anonymized DDoS attack statistics from more than 500 ISPs and 1500 enterprises worldwide. These industry sources do not share data that would reveal anything about a customer suffering an attack. We received the daily attack counts separated by attack type (reflection-amplification and direct-path). Netscout also provided counts for spoofed and non-spoofed attacks in the DP attack data. The third data set was collected by Akamai Prolexic, a DDoS protection service (DPS) that detects and mitigates attacks in traffic transiting its AS. It includes weekly attack counts.

 **Honey pots observe reflection-amplification attacks.** Honey pots observe DDoS attacks when their sensors are selected as amplifiers (⑧ in Figure 1). We used data from two academic honey pots: *Hopscotch* [167] and *AmpPot* [84]. Although *AmpPot* has ≈ 70 IPs allocated, it responds from only ≈ 30 , so it can associate attacks with previous scans based on which sensors it revealed [86].

Both *Hopscotch* and *AmpPot* provided observed attack counts and metadata (target, length, and (estimated) packet counts). We used algorithms developed by CCC [167] to process both the *Hopscotch* and *AmpPot* data in the same way. We aggregated attacks seen at multiple sensors into one event, including carpet-bombing [68] against many IPs, which a single sensor may not see (Appendix I introduces our improved logic for detecting these attacks, which we shared back with CCC).

We also have access to data from *NewKid* [68], but due to its single sensor (Table 2) the weekly attack counts were erratic. For completeness, we include it in Appendix D but exclude it from our long-term trend analysis §6.

Data aggregation. The plots in §6 are based on attack counts for each observatory aggregated per week, *i.e.*, new attacks observed each day, summed up to weekly totals. We normalized values to the median attack count of the first 15 weeks. We used a similar normalization as prior work [57] with an extended normalization period to fit the irregular nature of DDoS attacks. This approach defines a common metric and allows data providers to keep absolute counts private. Plots in §7 are based on distinct targets seen by each observatory per day, *i.e.*, (date, IP address) tuples. The time series count daily tuples and sum them up to weekly totals.


6 Comparing Long-term DDoS Trends

We now inspect the various DDoS detection data sets in detail and analyze how closely their findings correlate. Our data aggregation is described at the end of §5. For visualizing overall trends, we evaluated the exponentially weighted moving average (EWMA) of attack counts with a span of 12 weeks, and linear regression lines starting in 2019 to 2022 (respective slopes reported in legends). Note that these data sets do not allow us to distinguish between missing data and the absence of attacks unless otherwise noted.


6.1 Direct-path Attacks

Figure 2 shows normalized attack counts for observatories of direct path attacks. (Missing data: ORION in 2019Q3-Q4, IXP in Jan 2019.)


The telescopes observed a similar rise in attacks (Figures 2(a), 2(b)). They repeatedly saw short peaks that at least tripled attack counts; these peaks did not coincide in time. ORION saw its largest peaks in the first half of 2022 with smaller peaks in 2019Q2 and mid-2021. In contrast, UCSD saw its largest peak in 2023Q2 and small peaks in each year. In 2021 both telescopes saw an increase in attacks until summer, followed by a mild decrease until the end of the year. Although ORION attacks peaked in 2022Q1 and Q2 the




(a) Network Telescope: ORION.




(b) Network Telescope: UCSD. The peak in 2023 reaches 27.



(c) Flow Data: Netscout Atlas.




(d) Flow Data: Akamai Prolexic.




(e) Flow Data: IXP Blackholing. The peak in 2022 rises to 101.


Figure 2: Normalized weekly direct-path attack counts (to median of first 15 weeks as a baseline, highlighted in grey) show a growth in attacks over 4.5 years. Four observatories (ORION, UCSD, Akamai, Netscout) saw an upward trend in 2023 while one (IXP) saw a downward trend. Note y-axis scales differ.




(a) Honeypot: Hopscotch.




(b) Honeypot: AmpPot.



(c) Flow Data: Netscout Atlas.



(d) Flow Data: Akamai Prolexic.



(e) Flow Data: IXP Blackholing.

Figure 3: Normalized weekly reflection-amplification attack counts (to median of first 15 weeks as a baseline, highlighted in grey) show varying behavior over 4.5 years. The most striking similarity is the rise in attacks in 2020, and subsequent drop across 2021. Attacks rise again in 2023, except for Hopscotch. Red dashed lines mark DDoS takedowns by law-enforcement.

subsequent downward trend continued throughout 2022. UCSD trends remained positive for all sub-periods.

We identified three possible reasons for these divergent views of randomly-spoofed DoS attacks. (i) The UCSD telescope is roughly 20x larger than ORION, so ORION will see fewer packets of the same attack—provided the spoofed source addresses are uniformly distributed (§5), which makes it harder to distinguish attacks. (ii) Attacks may be too short-lived to rotate source addresses across the entire address space. (iii) Some attackers might exclude known telescopes from their address selection. Past studies showed that even telescopes of the same size will observe different IBR [178] and that telescopes in topological proximity, which is not the case here, capture more similar observations [142].

Netscout Atlas observed a relatively stable growth of attacks with an exception in 2021 (Figure 2(c)). Netscout presumably has a stable customer base with persistent sensors that monitor traffic for attacks. The first of two peaks, in 2020Q2, overlapped with peaks at UCSD and the IXP but with different dynamics and amplitudes. The peak in 2021Q1 partially overlapped with small peaks at the IXP and Akamai. One consistent aspect of the IXP and Netscout data (both based on observing two-way user traffic) is that relative attack counts reached a peak during the first half of the year (2019–2022) followed by a valley. The overall trend remained upward. However, visibility of attacks was limited to customers willing to share data.

Akamai Prolexic observations differed from other DP observatories as they exhibited a slight downward trend over the total period (Figure 2(d)). Attacks detected at Akamai remained relatively stable, hence the small y-axis scale. A valley in the second half of 2019 was followed by a high at the beginning of 2021 and a subsequent downward trend in 2021. A rise in attacks throughout 2020 was also observed by ORION and Netscout, but with more pronounced amplitudes. While peaks in 2021 rose by a factor of ≈ 1.5 the baseline, attacks decline overall, which was jointly observed with the IXP. While there were several peaks in 2022, the minima dropped below $\approx 0.5\times$ of the baseline. Attacks rose again in 2023, but even the peaks remained below $1.3\times$ of the baseline.

The IXP observed an increase in direct-path attacks (Figure 2(e)). Attack counts were more erratic, often dropping to zero. Attack counts jumped $\approx 10\times$ from their baseline in the first half of 2020 and 2021, and $\approx 30\times$ in 2023, but dropped for the second half of each year. The closest overlap with the telescopes was the higher activity in mid-2021, which was only a small elevation ($\approx 2\times$) in the UCSD data. An increase in activity was common to UCSD in 2020Q1/Q2 and ORION in 2022Q1/Q2. Note the IXP data reflects traffic for which customers requested blackholing. It is thus a lower-bound of direct-path attacks passing this IXP and may depend on IXP customer actions.

Trends in direct-path attacks. The linear regressions show that four of the five observatories experienced an upward trend over the full measurement period. A few peaks correlate across multiple data sets, albeit at different amplitudes. In 2023, three observatories saw a slight upward trend (UCSD, Akamai, Netscout) while ORION saw no trend, and the IXP saw a downward trend. These divergent observations across vantage points reflect the limited and disparate coverage of each sensor instrumentation (§7).

6.2 Reflection-amplification Attacks

Figure 3 shows the evolution of reflection-amplification attacks in our data sets. As in Figure 2, the y-axis range shows the difference in observed attacks as a factor of the normalized week count. We marked dates of known DDoS takedown operations with red dotted lines in these plots. Per seizure warrants, these happened on Dec 13, 2022, and May 4, 2023 [138].

The honeypots in our study observed a pronounced growth of attacks in 2020 (Figure 3(a),3(b)). Hopscotch recorded most attacks early in 2020, when Netscout and IXP counts also increased but with different relative amplitudes. In contrast, AmpPot saw its highest attack peaks later in 2020, mysteriously when Hopscotch peaks declined. Notably, all honeypots (HP) observed a decline in attacks from late 2021 until mid-2022, when both observed a spike not visible at the industry observatories (Netscout, Akamai, IXP). This downward trend is consistent with industry data (Figure 3(c), 3(d),3(e)) and industry efforts to deploy SAV (see discussion in §2.3).

Netscout Atlas exhibited a stable, mild upward trend with a pronounced rise until 2021, quick declines in early 2021 and early 2022, and a slow rise starting in Q1 of 2022 (Figure 3(c)).

Akamai Prolexic saw only small variations in attacks until 2020Q3, when attacks increased with a peak above $2\times$ its baseline in 2021Q1 (Figure 3(d)). The subsequent peak in 2021Q1 coincides with a period of high attack counts for Netscout, the IXP, and AmpPot. Like others, Akamai saw a decrease in attacks in the first half of 2021. However, the peaks in 2021Q4 are unique to Akamai. After dropping to $\approx 0.5\times$ in late 2022, attacks increased again.

The IXP observed a slow decline of attacks through most of 2019, followed by a steep increase until 2020Q2 (Figure 3(e)). Hopscotch and Netscout similarly observed a rise in attacks during 2020Q2, although at lower amplitudes. 2020Q4 was a second period of high attacks with a subsequent decline that continued (punctuated by bursts of attacks) until 2023. While the decrease in 2021Q2 was also observed by Netscout and Akamai, attack counts at the IXP decreased until the turn to 2023. After a low at the turn to 2023, attacks increase again but stayed below the baseline. This time series was more stable than IXP direct-path attacks (Figure 2(e)).

Trends in reflection-amplification attacks. No pair of time series exhibits similar behavior for the whole period, but all five vantage points showed the increase in attacks in 2020 followed by a decrease 2021. While Akamai, the IXP, and AmpPot saw this downward trend continue through 2022, Netscout and Hopscotch saw a flatter trend that year. Finally, attacks rose again through 2023 except for Hopscotch, which only saw a short peak in Q1. There were also short periods (3–6 months), in which two or more time series proceeded similarly, *i.e.*, (i) Hopscotch, AmpPot, IXP 2019Q4, (ii) Hopscotch, Netscout, IXP in 2020Q2, (iii) AmpPot, Netscout, Akamai, IXP (dip in mid 2021), (iv) Hopscotch and AmpPot (peak mid-2022), and (v) all observatories had a low in January 2023.

DDoS-service takedowns by law enforcement. Arrests and infrastructure seizures should have an immediate effect on attacks [31]. Two DDoS-takedown efforts during our observation time left an indeterminate footprint. The first in late 2022Q4 was followed by immediate, (small) valleys at the turn to 2023 in all four graphs. In contrast, the takedown in 2023Q2 was only followed




Figure 4: Normalized weekly attack counts observed at our 10 vantage points. Direct-path (DP) attacks (top 5 rows) increased in 2022 while reflection-amplification (RA) attacks (bottom 5 rows) had highest intensities during 2020 and declined thereafter.




Figure 5: The relative share of reflection-amplification (RA) and direct-path (DP) attacks observed by Netscout per week shows a shift toward DP attacks. The horizontal line indicates 50%, i.e., equal share. The dotted vertical line marks the latest crossing of the 50% mark.

by valleys in Figure 3(d), Figure 3(e), and Figure 3(a). Even if these changes were caused by takedowns, their impact on DDoS trends remained insignificant in our time series.

6.3 Trends Across Attack Types

We now take a comparative view of trends across all attack types, even though the different methods and coverage in measurements make this challenging.

Shifts between attack types. Figure 4 combines the time series from Figure 2 and Figure 3 into a single heatmap. Colors represent normalized attack counts. We arrange observatories by attack type: direct-path (top) and reflection-amplification (bottom). Overall, we find that time series data of the same attack type—with the exception of Akamai—are more positively correlated, but periods of anti-correlation also exist. Most direct-path attacks trended toward increased attack counts in early 2022, although at different intensities, before observation diverged.

In contrast, Akamai saw higher attack counts during 2019 and 2020 followed by a downward trend until 2023. Reflection-amplification attacks showed higher intensities between 2020Q2 and the end of 2021Q2 but lacked a clear trend toward the end. Again, Akamai

differed as attack counts increased later with the highest peak in 2021Q4, and continuously show smaller peaks throughout 2022. Even for the same attack type and measurement method, we can see opposite trends, such as for AmpPot and Hopscotch in 2023.

When comparing across observatory types, coherence lowers as the observed attack types differ. Netscout, which measures both direct path and reflection-amplification attacks at the same platform, observes a relative shift toward direct path attacks based on absolute attack counts (Figure 5). The dotted line marks the shift in 2021Q2, which matched the downward trend of RA attacks (Figure 3(c)). This roughly echoes our observation (Figure 4) that RA attacks were relatively high in the first two years while direct-path attacks increased relatively in the latter years. In contrast, Akamai consistently reported a larger share of direct-path attacks throughout the entire period.

Correlations between attack trends. Figure 6 shows Spearman correlations between pairs of observatories—a linear value between 1 (correlation) and -1 (anti-correlation). We applied the Spearman correlation because it calculates a monotonic correlation that is less susceptible to outliers than Pearson, a linear correlation. We calculated correlations for the normalized data (left) and the weighted moving average (EWMA) (right). Correlations with p-values (bottom part of the figure) above 0.05 are considered statistically insignificant and have their font greyed out.

We found a low to modest correlation between Netscout (DP) and direct-path (DP) attack observatories (ORION 0.2, UCSD 0.33, IXP (DP) 0.39) – except Akamai (p-value >0.05). Akamai (DP) showed a low correlation with ORION (0.20), but exhibited high p-values in correlations with other direct-path time series (0.06-0.72). Other DP observatories correlated weakly (0.16-0.20). Correlations between the EWMA were more pronounced. ORION and IXP (DP) showed the same medium correlation as Netscout and other observatories (0.43). The correlation between UCSD and ORION was statistically insignificant (p-value: 0.2). Akamai (DP) was anti-correlated with UCSD (-0.18) and the IXP (DP) (-0.45).

Our RA observatories showed higher correlations than the DP observatories, with low to moderate pairwise correlations (0.17 to 0.59). Hopscotch had statistically insignificant correlations with




Figure 6: Spearman correlation: Platforms that observe the same kind of DDoS event (direct-path or reflection-amplification) show higher correlation. The Akamai direct-path time series is an exception as it correlated with reflection-amplification observatories. The top two graphs show the Spearman correlation for the normalized data on the left and the EWMA on the right. A grey font marks correlation coefficients with a p-values above 0.05. The bottom graphs show the respective p-values.

Netscout (RA) (p-value: 0.12) and Akamai (RA) (p-value: 0.75). The EWMA lines correlated more strongly (0.26-0.69) when significant. For EWMA, Hopscotch and Akamai (RA) had a low p-value and exhibit an anti-correlation (-0.20). The p-value of Akamai (RA) and Netscout (RA) increased leaving the correlation insignificant.

Two observatories stood out in correlations across attack types: (i) Akamai (DP) showed positive correlations with observatories of the opposite (RA) type (0.27-0.56). (ii) ORION showed positive correlations with four of the five RA observatories: Netscout (RA) (0.39), Akamai (RA) (0.38), IXP (0.23), and AmpPot (0.32). Additionally, the two Netscout time series had a medium correlation (0.48). Analysis between observatories showed low to medium anti correlations (-0.14 to -0.36). In addition to two positive correlations with ORION and Akamai (DP), the Akamai (RA) time series only had statistically insignificant correlations with the other three direct-path observatories. Two more correlations were statistically insignificant: Netscout (RA) & UCSD and AmpPot & Netscout (DP).

Apart from Akamai, time series of the same attack type tended to correlate more strongly, *i.e.*, the attack counts evolved in similar ways. This correlation did not hold for all pairs of observatories within each group, implying that different observatories of the same group do not uniformly observe the same attack events—consistent with our earlier findings. While our trend lines were more consistent within the direct-path group the correlation analysis shows higher values for the RA observatories.

We cross-checked our results by calculating the Pearson correlation. It confirmed our results, even if correlation values varied

slightly: correlation was stronger among RA observatories and weaker among DP observatories. Appendix F provides quarterly pairwise correlations and similarly shows stronger correlations among observatories that observe the same attack type.

Why does Akamai see different trends? The data from Akamai Prolexic represents attack events from traffic that transited its AS. Customers must own a prefix that can be rerouted through the Prolexic AS for attack mitigation. This requirement will affect attack methodologies and trends in their data. Netscout and the IXP also observe with biasing characteristics. This reality is an inherent challenge to understanding a competitive private sector landscape with no standardized reporting of the analyzed phenomena.

Pandemic. As lockdowns drove people to spend more time online [57], DDoS attacks increased and diversified [43, 54]. The observatories in our study saw an increase in both DP (Figure 2(b),2(e), 2(d), 2(c)) and RA (Figure 3(c),3(e), 3(a), 3(b)) attacks during 2020.

7 Analyzing DDoS Targets

Our trend analysis (§6) showed rough similarities within attack types. We now examine how targets of DDoS appeared across observatories. During a 4.5-year observation period, IP addresses can change owners while others rotate due to dynamic assignment. We used the tuple (*attack start date, target IP address*) (see §5) to identify a target unless otherwise noted, and deduplicated the resulting set, obtaining 28,474,161 distinct targets or 14,565,588 distinct IP addresses.

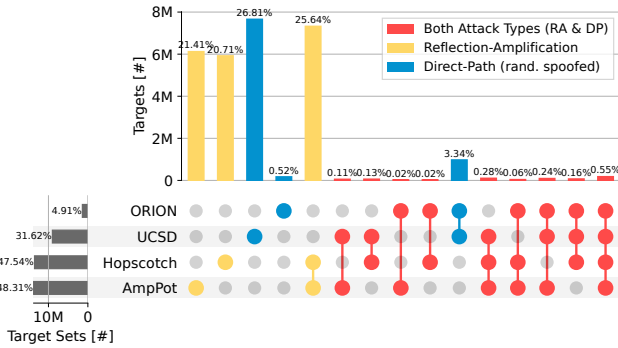


Figure 7: The UpSet plot displays targets that were exclusively seen in intersection sets of observatories on the same day. Colors signify the attack types in each set. Observatories of the same type (RA, DP) see large shares of overlapping targets. Only 0.55% of targets were observed by all four observatories.

7.1 Attackers Often Chose One Attack Type

We first compared targets across observatories to answer the two questions: (i) How many distinct targets did each observatory see?, and (ii) how large was the target overlap between all four observatories? The UpSet plot [94] in Figure 7 answers both questions. It shows the number of distinct targets (date, IP address) seen by each observatory in the left bar plot. These shares are not exclusive, *i.e.*, they sum up to more than 100% of distinct targets. The top bar plot quantifies the target intersections, *i.e.*, the targets exclusively observed by the combination of observatories marked in the matrix below. Each bar shows absolute counts and additionally displays its share among all distinct targets at the top.

Both honeypots (HP) saw nearly the same number of targets (left bars: roughly 48%). ORION saw an order of magnitude fewer targets than the HPs and six times fewer than UCSD, which monitors 22 times more IP addresses.

Hopscotch and AmpPot each uniquely observed $\approx 21\%$ of all targets and another 25% together (three yellow bars in the top bar plot), more than the respective individual shares. UCSD had the highest share of uniquely observed distinct targets ($>26\%$). The limiting factor for a larger overlap is the low target count observed by ORION. Except for UCSD (14%) the share of overlapping targets across observatories of the same kind (NT: UCSD & ORION, HP: AmpPot & Hopscotch) was over 50% (AmpPot 57%, Hopscotch 56%, ORION 87%). AmpPot, for example, shared 57% of the targets it observed with Hopscotch. For reflection-amplification attacks, this means that attackers likely selected reflectors from multiple HPs. Still, both honeypots observed a significant share of exclusive targets. Randomly-spoofed DoS attacks that produce enough packets to be visible in ORION are likely to produce enough packets to appear in larger telescopes. But even with its size, UCSD NT did not observe all targets.

Nawrocki et al. [117] examined the effect of different attack definitions used by honeypots and found a 15%–45% difference in attack targets. Our data sets are based on already processed traffic




Figure 8: Target tuples (date, IP address) observed by all four observatories (ORION, UCSD, Hopscotch, AmpPot) during our measurement period, summed up per week. These targets reflect the 0.55% target-most set in Figure 7. The dashed line counts new targets as CDF on the right axis.

data, which does not allow us to make statements about the effect of attack detection thresholds.

Highly-visible targets. A small fraction of targets (1.57%) was affected by multiple attack types, a third of these (0.55% or 155,663 targets or 97,470 distinct IP addresses) at all four observatories. Figure 8 shows the time series of these 0.55% targets as a stacked area plot (left axis). The blue area signifies new targets, *i.e.*, IPs seen for the first time, while the orange area counts recurring targets. The dotted line plots target growth over time as a CDF (right axis). The four observatories continuously saw new shared DDoS targets, most of which appearing between 2020Q4 and 2021Q2. This time series does not resemble any individual trend line (Figure 2 and 3), but honeypots saw periods of higher attack counts during these times. Leaving out the relatively small ORION, the three remaining observatories saw an additional 0.28% overlapping targets (50% more). This still adds up to fewer than 1% of all targets (Figure 7).

The largest share of highly-visible targets belonged to OVH (AS16276, 18.8%), followed by Hetzner (AS24940, 5.1%) and Amazon (AS16509, 2.69%). Including these three, 7 of our top 10 most targeted ASes belong to hosters. Hosters usually offer DDoS-protection-as-a-service, which may lead attackers to use multiple attack vectors, *e.g.*, RA and DP attacks, to overcome defenses.

In a two-year study (2015-17) Jonker et al. [76] examined the overlap between AmpPot and UCSD NT. They found 4.5% (282k) shared IP addresses, half of which were hit by attacks at the same time. In our data, this overlap is lower, *i.e.*, 1.18%–2.9% of the IP addresses. Jonker et al. also saw OVH as the main target (12.3%), followed by China Telecom (14th largest share of targeted IPs in our data) and China Unicom/AS4837 (7th). See Appendix H for the top 10 ASes by number of highly-visible targets.

7.2 Target Overlap with Industry

Netscout. We extended the target overlap analysis to industry with a focus on federated attack inference. Netscout compared targets from academia (Figure 7) to their baseline data set—constituting approximately 28% of all Netscout alerts. The shares of confirmed targets are presented in Figure 9. There are important caveats when comparing Netscout and academic attack observations. First, Netscout’s anonymized data limits complete confirmation. Second, Netscout excludes attack alerts below the product-defined “medium” threshold, which may prevent confirmation of less severe attacks

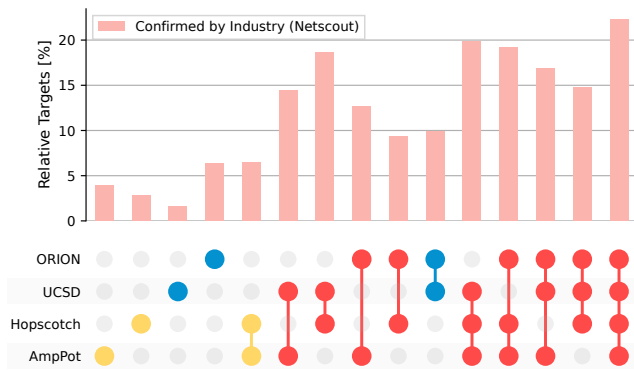


Figure 9: Relative share of targets observed by academia and confirmed by industrial baseline data from Netscout. Super-sets are shown in Figure 7. Netscout baseline data shows the largest relative overlap with the targets seen by all four observatories. These are likely large, multi-vector attacks.

seen by academic observatories. Lastly, Netscout observations are made at intermediate systems on network paths of customers that contribute alert feedback as opposed to the endpoint nodes that are party to all attack events observed. Netscout targets had a 2%-6% overlap with each *individual* observatory, but a 20% overlap with the set of targets seen by all academic observatories, *i.e.*, 20% of the 0.55% in Figure 8. Larger, multi-vector attacks were more likely seen from all vantage points.

We assessed how many targets inferred by Netscout were also observed by academia. With 23% of the baseline data set this is a substantial but partial view. No academic observatory independently saw all shared targets, with overlaps of 15.2%, 13.6%, 5.7%, and 3.1%, respectively. This intersection analysis compares attack type-specific observatories with *all* DDoS attacks as observed by Netscout, which generally will result in lower intersections.

Akamai. The overlap analysis with targets in the network prefix of Akamai showed minor overlaps (<0.25%), about 100× lower than Netscout. See Appendix G for a detailed plot. A small overlap is consistent with the focus of the Akamai dataset on their ASN, which advertises a subset of prefixes on the Internet. Together, academic research observatories saw 33% of the Akamai target set. Both honeypots saw a larger share of the targets ($\approx 20\%$) than the respective telescopes ($\approx 7\%$). These findings underscore the value of federated DDoS inference.

Previous work. Nawrocki et al. [117] found a 3% overlap in targets between Hopscotch (RA) and a commercial DoS mitigation provider (RA) (Nov 21 and Jan 22). Figure 9 shows a similar $\approx 3\%$ overlap between Hopscotch and Netscout, which is a lower bound as mixed attack type shares are not considered. A 7-month study in 2019–20 [82] found 33% overlap of targets between IXP blackholing and honeypot observations.

7.3 Target Overlap over Time

Figure 10 shows the overlap in observed targets per day (summed up per week) between telescopes and honeypots. Blue and orange

are the respective observatories while the green line marks the overlapping subset.

Telescopes. The telescopes observed 32.23% of all targets, 95% (30.66%) of these *only* the telescopes observed Figure 10(a). Matching the observation from Figure 7, UCSD observed most targets seen by ORION. Unlike their trend graphs (Figure 2(a), 2(b)), their target lines show similar behavior: a downward trend towards 2020, an upward trend until 2021Q3, and a low point at the turn of 2023 followed by a small upward trend. The UCSD line has many more spikes. The high in mid-2021 also appears in the attack trend graphs (Figure 2(a), 2(b)).

Honeypots. Both honeypots (HP) saw the same order of magnitude of targets (Figure 10(b)). Their observations add up to 69.33% of all unique targets and 67.76% of targets they observed exclusively. While the shared targets were usually a subset of both HP target sets, there was a period from mid-2021 to mid-2022 where AmpPot observed most of the targets seen by Hopscotch.

The target graphs show a similar behavior to the respective attack trends (compare Figure 3(a), 3(b)). We marked the DDoS takedowns once more with red dotted lines. All observatories saw a small valley around the first takedown, followed by a rise in targets.


An attacker must select a honeypot as a reflector for that honeypot to observe the attack. Differences in protocol support across honeypots will affect the composition of attacks they see. As an example, AmpPot observed more targets attacked via CHARGEN while Hopscotch saw more targets attacked via CLDAP until mid-2020. For protocols such as QOTD, RPC, and NTP both had largely overlapping target sets.

8 Related Work


The last decade has yielded a vast range of DDoS studies, of which we can only present a snapshot here.

Attack characterization. Studies have characterized DDoS attacks [76, 82, 157] as well as abusable protocols [89, 91, 114, 115, 118, 140, 155, 158, 175]. They quantified the state of amplifier deployment [116], identified scan infrastructure [86], and booter services [87], measured the adoption of DDoS services [78] enabling flow telemetry-based traceback [88], examined detection methods [15, 104, 137], quantified attacks [16, 62, 65, 76, 167], discovered new attack vectors [17, 19, 82, 83, 110, 112, 115], and criminal techniques, tactics, and procedures (TTPs) [66, 69, 70, 83]. Few research efforts moved beyond a single class of attacks [76, 77] or arrived at insights by studying the overlap between independent datasets [117]. To our knowledge, we provide the first comprehensive, longitudinal trend analysis of the two most concerning DDoS attack types, enabled by vantage points from research and industry.

Mitigation. Prior studies have focused on proactive measures and reactive responses. These include the efficacy of operational controls, such as blackholing [71, 77, 113, 177], traffic engineering [154], or anycast [109]. Other work measured the adoption of protection techniques and resulting resilience [78, 154], and explored mitigation options in improved protocol design [155], collaborative detection and response [88, 176], and law enforcement intervention [31, 83, 106]. Our work analyzed attack trends, which informs where to focus mitigation techniques.



(a) The UCSD telescope observed most targets seen by ORION.



(b) Both honeypots uniquely observed a large share of targets.

Figure 10: Weekly observed targets (unique per day) for telescopes (a) and honeypots (b) and respective overlapping observations.

Open challenge. A graphical taxonomy of recent related work is presented in Appendix C. While each study contributes to the overall DDoS landscape, our community has not yet arrived at a common and comprehensive understanding of DDoS. We aim to narrow this gap with our data-driven analysis involving multiple communities, and advocate steps to support a science of DDoS assessment in the future.

9 Implications for Public Policy

Persistence and prevalence of DDoS poses a systemic risk to the ecosystem, and researchers and operators do not understand the current extent of harms or effectiveness of mitigations. The breadth and scope of this study suggests the inherent limits to what the academic community can provide on its own.

Growing awareness of the obstacle of opacity has motivated governments to advocate for more transparency. The EU NIS2 directive (2021) requires essential and important entities to notify their competent authority of any incident that “significantly impacts” provision of their service [51]. Similar U.S. regulations are in discussion [44]. These developments offer an opportunity to inform consideration of how mandated data sharing of DDoS incidents could occur to ensure its utility to the regulatory objective.

In particular, mandated reporting does not usually involve publishing the information or even sharing it with independent researchers for study. Peer-reviewed research that concretely substantiates the need for specific frameworks for data sharing will be essential to ensuring academics can contribute to public policy efforts to advance Internet security. Unexplored details include many that we directly address in this paper: definition of incidents and their impact; data formats to accommodate comparisons; disclosure controls technologies and access policies to allow rigorous independent analyses. We hope this work can guide regulators to consider (and facilitate) the role of academic research in informing and leveraging reporting regulations that can advance scientific understanding of the DDoS landscape.⁴

Measurement of spoofing. Spoofing persists as a key vector of DDoS attacks, and remediation of this vulnerability requires knowing which networks allow spoofing. Attempts to measure source address validation (SAV) have struggled with sustainability for decades, *e.g.*, [97]. Efforts to internalize this negative externality

⁴A shorter term goal for industry transparency would be for companies or a third party to publish and share historic versions of industry reports, rather than only the most recent, sometimes limited in distribution.

– “naming and shaming”, procurement guidelines, voluntary code of conducts – have had limited impact [96].

We see two paths forward. First, a requirement for transparency regarding SAV deployment, similar to other recent ISP transparency requirements [172, 173]. This direction would require a sustained operational measurement infrastructure to support auditing and impact assessments. Although long controversial, the RIPE Atlas measurement system could support such measurements [2]. Second, industry and governments could collaborate to promote SAV measurement capability as a default on end user equipment, a compromise on making SAV deployment itself a default.

Availability measurement. Breaches of confidentiality are now accepted as sufficiently important to mandate reporting, but regulatory attention to requiring data on availability is in early stages.⁵ A non-governmental approach could rely on financial sector auditors, *e.g.*, PCI-DSS, to add availability to its auditing framework.

10 Summary and Future Work

We provided the first comprehensive, longitudinal view on direct-path and reflection amplification attacks, two dominant classes of DDoS attacks that threaten Internet infrastructure. Our results document joint forces of academia and industry sharing data across institutional boundaries. We included all datasets that were made available to us, covering far more attacks than any other study. Our approach to synthesizing datasets expanded beyond prior efforts in five dimensions: number of macroscopic datasets (ten), observation window (4.5 years vs months or even 3 years in previous works), multiple types of DDoS attacks, volume and usage restrictions of data (which required substantial cooperation across institutions in processing and normalizing data), and a new method to facilitate sharing by industry. Beyond the synthesis of datasets, our work reinforced findings that previous studies have serious visibility limitations, which provide the strongest empirical grounding to date for regulatory framing to share data.

Our artifacts include a living compilation of facts from industry reports characterizing DDoS phenomena in 2022–2023, and a mindmap taxonomy of academic DDoS studies over the last few years (Appendix B). The next step is to motivate others who have DDoS data to contribute to our effort to establish and maintain a holistic view of the DDoS ecosystem, following this blueprint.

⁵In the UK, consumer ISPs must reimburse customers £9.33 for each calendar day where the service is unavailable [170], telecoms providers must notify Ofcom of availability incidents [139], and banks must notify the Financial Conduct Authority [10].

Acknowledgments

We thank our shepherd and reviewers for helpful feedback. We gratefully acknowledge all partner teams from industry and academia for sharing their data. This work was partly supported by National Science Foundation grant CNS-2212241 and OAC-2319959, the German Research Foundation (DFG) within the project ReNO (#511099228), and the German Federal Ministry of Education and Research (BMBF) within the project PRIMeNet.

References

- [1] A10. 2022. 2022 A10 Networks DDoS Threat Report. <https://www.a10networks.com/resources/reports/2022-ddos-threat-report/>
- [2] Emile Aben. 2016. [atlas] What is 'iwantbcp38compliancecetesting' user tag? ripe-atlas – RIPE Network Coordination Centre. <https://www.ripe.net/ripe/mail/archives/ripe-atlas/2016-January/002581.html>
- [3] Paul Aitken, Benoît Claise, and Brian Trammell. 2013. Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information. RFC 7011. <https://doi.org/10.17487/RFC7011>
- [4] Akamai. 2022. The Relentless Evolution of DDoS Attacks. <https://www.akamai.com/blog/security/relentless-evolution-of-ddos-attacks>
- [5] Akamai. 2023. DDoS Attacks in 2022: Targeting Everything Online, All at Once. <https://www.akamai.com/blog/security/ddos-attacks-in-2022-targeting-everything-online>
- [6] Alibaba Cloud. 2021. DDoS Attack Statistics and Trend Report by Alibaba Cloud. https://www.alibabacloud.com/blog/ddos-attack-statistics-and-trend-report-by-alibaba-cloud_597607
- [7] Radu Anghel, Swaathi Vetrivel, Elsa Turcios Rodriguez, Kaichi Sameshima, Daisuke Makita, Katsunari Yoshioka, Carlos H. Gañán, and Yury Zhauniarovich. 2023. Peering into the Darkness: The Use of UTRS in Combating DDoS Attacks. In *European Symposium on Research in Computer Security (ESORICS)*. Springer-Verlag, Berlin, Heidelberg, 23–41.
- [8] Anti-DDoS-Coalition. 2023. Dutch National Anti-DDoS-coalition. <https://www.nomoreddos.org/en/>
- [9] Arelion. 2023. Arelion DDoS Threat Landscape report 2023. <https://www2.arelion.com/wp-securityreport2023>
- [10] Financial Conduct Authority. 2019. *Interpreting the data*. Financial Conduct Authority. <https://www.fca.org.uk/data/mandated-voluntary-information-current-account-services/interpreting-data>
- [11] AWS. 2021. AWS Shield Threat Landscape Review: 2020 Year-in-Review. <https://aws.amazon.com/blogs/security/aws-shield-threat-landscape-review-2020-year-in-review/>
- [12] F. Baker and P. Savola. 2004. BCP 84, RFC 3704: Ingress Filtering for Multihomed Networks. <https://www.rfc-editor.org/info/bcp84>
- [13] Marinho Barcellos, Raphael Hiesgen, Marcin Nawrocki, Daniel Kopp, Oliver Hohlfeld, Echo Chan, Roland Dobbins, Christian Doer, Christian Rossow, Daniel R. Thomas, Mattijs Jonker, Ricky Mok, Xiapu Luo, John Kristoff, Thomas C. Schmidt, Matthias Wählisch, and KC Claffy. 2024. DDoS Industry Reports Repository. <https://ddoscovery.github.io/>
- [14] S. M. Bellovin. 1989. Security Problems in the TCP/IP Protocol Suite. *SIGCOMM Comput. Commun. Rev.* 19, 2 (April 1989), 32–48.
- [15] Agathe Blaise, Mathieu Bouet, Vania Conan, and Stefano Secci. 2020. Detection of zero-day attacks: An unsupervised port-based approach. *Computer Networks* 180 (2020), 107391. <https://doi.org/10.1016/j.comnet.2020.107391>
- [16] Norbert Blenn, Vincent Ghiëtte, and Christian Doerr. 2017. Quantifying the Spectrum of Denial-of-Service Attacks through Internet Backscatter. In *Proc. of the ARES (Reggio Calabria, Italy)*. ACM, New York, NY, USA, Article 21, 10 pages. <https://doi.org/10.1145/3098954.3098955>
- [17] Kevin Bock, Abdulrahman Alaraj, Yair Fax, Kyle Hurley, Eric Wustrow, and Dave Levin. 2021. Weaponizing Middleboxes for TCP Reflected Amplification. In *Proc. of USENIX Security*. USENIX Association, Berkeley, CA, USA, 3345–3361. <https://www.usenix.org/conference/usenixsecurity21/presentation/bock>
- [18] Brian Krebs. 2023. Feds Take Down 13 More DDoS-for-Hire Services. <https://krebsonsecurity.com/2023/05/feds-take-down-13-more-ddos-for-hire-services>
- [19] Renée Burton. 2019. Characterizing Certain DNS DDoS Attacks. *CoRR abs/1905.09958* (2019), 25 pages. arXiv:1905.09958 <http://arxiv.org/abs/1905.09958>
- [20] C. Loibl and S. Hares and R. Raszuk and D. McPherson and M. Bacher. 2020. Dissemination of Flow Specification Rules. <https://www.rfc-editor.org/rfc/rfc8955>
- [21] CAIDA. 2012. The UCSD Network Telescope. Website. https://www.caida.org/projects/network_telescope/ Last Access: Nov 2023.
- [22] R K C Chang. 2002. Defending against flooding-based distributed denial-of-service attacks: A tutorial. *IEEE Communications Magazine* 40, 10 (Jan. 2002), 42–51. <https://doi.org/10.1109/MCOM.2002.1039856>
- [23] Benoît Claise. 2004. Cisco Systems NetFlow Services Export Version 9. RFC 3954. <https://doi.org/10.17487/RFC3954>
- [24] Richard Clayton, Julia Powles, and Cambridge University Legal. 2016. *Cambridge Cybercrime Centre: Legal framework*. Cambridge Cybercrime Centre. <https://www.cambridgecybercrime.uk/data.html>
- [25] Cloudflare. 2022. Cloudflare DDoS threat report 2022 Q3. <https://blog.cloudflare.com/cloudflare-ddos-threat-report-2022-q3>
- [26] Cloudflare. 2022. Cloudflare DDoS threat report for 2022 Q4. <https://blog.cloudflare.com/ddos-threat-report-2022-q4/>
- [27] Cloudflare. 2022. DDoS Attack Trends for 2022 Q1. <https://blog.cloudflare.com/ddos-attack-trends-for-2022-q1/>
- [28] Cloudflare. 2022. DDoS Attack Trends for Q2 2022. <https://blog.cloudflare.com/ddos-attack-trends-for-2022-q2/>
- [29] Cloudflare. 2023. Cloudflare DDoS Trends Report Q1 2023. https://cf-assets.www.cloudflare.com/slt3lc6tev37/4CvITDALVKaap3iwrNOWxl/f9a653dacc12d3635c1a1955b59a7b91/BDES-4486_Q1-2023-DDoS-Trends-Report-Letter.pdf
- [30] Ben Collier, Gemma Flynn, James Stewart, and Daniel Thomas. 2022. Influence advertising: Exploring practices, ethics, and power in the use of targeted advertising by the UK state. *Big Data & Society* 9, 1 (2022), 1–13. <https://doi.org/10.1177/20539517221078756>
- [31] Ben Collier, Daniel R. Thomas, Richard Clayton, and Alice Hutchings. 2019. Booting the Booters: Evaluating the Effects of Police Interventions in the Market for Denial-of-Service Attacks. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 50–64. <https://doi.org/10.1145/3355369.3355592>
- [32] COMCAST. 2021. Comcast Business DDoS Threat Report 2021. https://i.crn.com/sites/default/files/ckfinderimages/userfiles/images/crn/custom/2022/Comcast_LC_Q2_2022_DDOS_Threat_Report.pdf
- [33] COMCAST. 2023. 2023 Comcast Business Cybersecurity Threat Report. https://business.comcast.com/community/docs/default-source/default-document-library/ccb_threatreport_071723_v2.pdf?sfvrsn=c220ac01_2
- [34] Corero. 2023. 2023 DDoS Threat Intelligence Report. <https://www.juniper.net/content/dam/www/assets/analyst-reports/us/en/2023/corero-ddos-threat-intelligence-report.pdf>
- [35] Corero. 2023. How Have DDoS Attacks Evolved Over the Last 10 Years? <https://www.corero.com/ddos-attack-evolution/>
- [36] Corero. 2023. The Shifting Landscape of DDoS Attacks. <https://www.corero.com/shifting-landscape/>
- [37] Craig Labovitz. 2021. Tracing DDoS End-to-End in 2021. https://www.youtube.com/watch?v=TP3H_GeLl-0
- [38] CrowdStrike. 2023. Global Threat Report. <https://www.crowdstrike.com/global-threat-report/>
- [39] Team Cymru. 2023. Unwanted Traffic Removal Service. <https://www.teamcymru.com/ddos-mitigation-services>
- [40] Evan Damon, Julian Dale, Evaristo Laron, Jens Mache, Nathan Land, and Richard Weiss. 2012. Hands-on Denial of Service Lab Exercises Using SlowLoris and RUDY. In *Proc. of the InfoSecCD*. ACM, New York, NY, USA, 21–29. <https://doi.org/10.1145/2390317.2390321>
- [41] DDoS-Guard. 2023. DDoS Attack Trends in 2022. <https://ddos-guard.net/en/blog/ddos-attack-trends-2022>
- [42] DDoS-Guard. 2023. DDoS-Guard Analytical Report on DDoS Attacks for 2022. <https://ddos-guard.net/info/protect?id=40954>
- [43] Anderson Bergamini de Neira, Burak Kantarci, and Michele Nogueira. 2023. Distributed denial of service attack prediction: Challenges, open issues and opportunities. *Computer Networks* 222 (Jan 2023), 1–27. <https://www.sciencedirect.com/science/article/pii/S1389128622005874>
- [44] Homeland Security Department. 2024. Cyber Incident Reporting for Critical Infrastructure Act (CIRCIA) Reporting Requirements. <https://www.federalregister.gov/documents/2024/04/04/2024-06526/cyber-incident-reporting-for-critical-infrastructure-act-circia-reporting-requirements>
- [45] Christoph Dietzel and Matthias Wichtlhuber. 2018. Stellar: Network Attack Mitigation using Advanced Blackholing. In *Proc. of ACM CoNEXT*. ACM, New York, NY, USA, 152–164. <https://doi.org/10.1145/3281411.3281413>
- [46] Christos Douligeris and Aikaterini Mitrokotsa. 2004. DDoS attacks and defense mechanisms: classification and state-of-the-art. *Computer Networks* 44, 5 (2004), 643–666. <https://doi.org/10.1016/j.comnet.2003.10.003>
- [47] Ben Du, Cecilia Testart, Romain Fontugne, Gautam Akiwate, Alex C. Snoeren, and kc claffy. 2022. Mind Your MANRS: Measuring the MANRS Ecosystem. In *Proc. of ACM IMC (IMC '22)*. ACM, New York, NY, USA, 716–729. <https://doi.org/10.1145/3517745.3561419>
- [48] W. Eddy. 2007. *TCP SYN Flooding Attacks and Common Mitigations*. RFC 4987. IETF. <https://doi.org/10.17487/RFC4987>
- [49] Elliott Peterson and Cameron Schroeder. 2023. Dismantling DDoS: Lessons in Scaling. <https://www.blackhat.com/us-23/briefings/schedule/#dismantling-ddos---lessons-in-scaling-31408>

- [50] European Commission. 2022. Regulation of the European Parliament and of the Council on horizontal cybersecurity requirements for products with digital elements and amending Regulation (EU) 2019/1020. <https://digital-strategy.ec.europa.eu/en/library/cyber-resilience-act>
- [51] European Union. 2021. The NIS2 Directive: A high common level of cybersecurity in the EU. [https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI\(2021\)689333](https://www.europarl.europa.eu/thinktank/en/document/EPRS_BRI(2021)689333).
- [52] EUROPOL. 2022. Global crackdown against DDoS services shuts down most popular platforms. <https://www.europol.europa.eu/media-room/newsroom/global-crackdown-against-ddos-services-shuts-down-most-popular-platforms>
- [53] F5. 2023. F5 DDoS Attack Trends 2023. <https://www.f5.com/labs/articles/threat-intelligence/2023-ddos-attack-trends>
- [54] Olufunso I. Falowo, Murat Ozer, Chengcheng Li, and Jacques Bou Abdo. 2024. Evolving Malware and DDoS Attacks: Decadal Longitudinal Study. *IEEE Access* 12 (Mar 2024), 39221–39237. <https://doi.org/10.1109/ACCESS.2024.3376682>
- [55] Fastly. 2023. Cyber 5 Threat Insights. <https://www.fastly.com/blog/cyber-5-threat-insights>
- [56] Fastly. 2023. What Is a DDoS Attack? <https://www.fastly.com/learning/what-is-a-ddos-attack>
- [57] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Posee, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapiador, Narseo Vallina-Rodriguez, Oliver Hohlfeld, and Georgios Smaragdakis. 2020. The Lock-down Effect: Implications of the COVID-19 Pandemic on Internet Traffic. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 1–18. <https://doi.org/10.1145/3419394.3423658>
- [58] P. Ferguson and D. Senie. 2000. *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*. RFC 2827. IETF. <https://doi.org/10.17487/RFC2827>
- [59] FORTINET. 2023. Global Threat Landscape Report. <https://global.fortinet.com/lp-en-ap-2023globalthreatlandscape-H1>
- [60] Future Market Insights (FMI). 2023. DDoS Protection Market. <https://www.futuremarketinsights.com/reports/ddos-protection-market>
- [61] Thomas Geras and Thomas Schreck. 2023. Sharing Communities: The Good, the Bad, and the Ugly. In *Proc. of the 2023 ACM SIGSAC Conference on Computer and Communications Security* (, Copenhagen, Denmark.) (CCS '23). Association for Computing Machinery, New York, NY, USA, 2755–2769. <https://doi.org/10.1145/3576915.3623144>
- [62] Vincent Ghiette and Christian Doerr. 2018. How Media Reports Trigger Copycats: An Analysis of the Brewing of the Largest Packet Storm to Date. In *ACM SIGCOMM Workshop on Traffic Measurements for Cybersecurity (WTMC)*. ACM, New York, NY, USA, 8–13. <https://doi.org/10.1145/3229598.3229606>
- [63] Vasileios Giotsas, Georgios Smaragdakis, Christoph Dietzel, Philipp Richter, Anja Feldmann, and Arthur Berger. 2017. Inferring BGP Blackholing Activity in the Internet. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 1–14. <https://doi.org/10.1145/3131365.3131379>
- [64] L. Gommans, J. Vollbrecht, B. Gommans-de Bruijn, and C. de Laat. 2015. The Service Provider Group framework: A framework for arranging trust and power to facilitate authorization of network services. *Future Generation Computer Systems* 45 (2015), 176–192.
- [65] Harm Griffioen and Christian Doerr. 2020. Quantifying TCP SYN DDoS Resilience: A Longitudinal Study of Internet Services. In *IFIP Networking*. IEEE, Piscataway, NJ, USA, 217–225.
- [66] Harm Griffioen, Kris Oosthoek, Paul van der Knaap, and Christian Doerr. 2021. Scan, Test, Execute: Adversarial Tactics in Amplification DDoS Attacks. In *Proc. of ACM CCS*. ACM, New York, NY, USA, 940–954. <https://doi.org/10.1145/3460120.3484747>
- [67] Yuhei Hayashi, Meiling Chen, and Li Su. 2023. Use Cases for DDoS Open Threat Signaling (DOTS) Telemetry. RFC 9387. <https://doi.org/10.17487/RFC9387>
- [68] Tiago Heinrich, Rafael R. Obelheiro, and Carlos A. Maziero. 2021. New Kids on the DRDoS Block: Characterizing Multiprotocol and Carpet Bombing Attacks. In *Proc. of PAM*. Springer International Publishing, Cham, 269–283. https://doi.org/10.1007/978-3-030-72582-2_16
- [69] Raphael Hiesgen, Marcin Nawrocki, Alistair King, Alberto Dainotti, Thomas C. Schmidt, and Matthias Wählisch. 2022. Spoki: Unveiling a New Wave of Scanners through a Reactive Network Telescope. In *Proc. of 31st USENIX Security Symposium*. USENIX Association, Berkeley, CA, USA, 431–448. <https://www.usenix.org/system/files/sec22-hiesgen.pdf>
- [70] Raphael Hiesgen, Marcin Nawrocki, Thomas C. Schmidt, and Matthias Wählisch. 2022. The Race to the Vulnerable: Measuring the Log4j Shell Incident. In *Proc. of Network Traffic Measurement and Analysis Conference (TMA)* (Neschede, Netherlands). IFIP, Laxenburg, MD, Austria, 1–9.
- [71] Nico Hinze, Marcin Nawrocki, Mattijs Jonker, Alberto Dainotti, Thomas C. Schmidt, and Matthias Wählisch. 2018. On the Potential of BGP Flowspec for DDoS Mitigation at Two Sources: ISP and IXP. In *Proc. of ACM SIGCOMM. Poster Session*. ACM, New York, NY, USA, 57–59. <https://doi.org/10.1145/3234200.3234209> 2nd price at ACM student research competition.
- [72] Huawei. 2023. Global DDoS Attack Status and Trend Analysis in 2022. <https://e.huawei.com/en/material/networking/security/0c561b8fd2d342999cd402bcecf6d452>
- [73] Imperva. 2023. The Imperva Global DDoS Threat Landscape Report 2023. <https://www.imperva.com/resources/resource-library/reports/ddos-threat-landscape-report-2023/>
- [74] International Telecommunications Union (ITU). 2003. X.805 : Security architecture for systems providing end-to-end communications. <https://www.itu.int/rec/T-REC-X.805-200310-I/en>
- [75] John Kristoff. 2015. An Internet-wide BGP RTBH service.
- [76] Mattijs Jonker, Alistair King, Johannes Krupp, Christian Rossow, Anna Sperotto, and Alberto Dainotti. 2017. Millions of Targets under Attack: A Macroscopic Characterization of the DoS Ecosystem. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 100–113. <https://doi.org/10.1145/3131365.3131383>
- [77] Mattijs Jonker, Aiko Pras, Alberto Dainotti, and Anna Sperotto. 2018. A First Joint Look at DoS Attacks and BGP Blackholing in the Wild. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 457–463. <https://doi.org/10.1145/3278532.3278571>
- [78] Mattijs Jonker, Anna Sperotto, Roland van Rijswijk-Deij, Ramin Sadre, and Aiko Pras. 2016. Measuring the Adoption of DDoS Protection Services. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 279–285. <https://doi.org/10.1145/2987443.2987487>
- [79] Kaspersky. 2022. Kaspersky DDoS Attacks in Q2 2022. <https://securelist.com/ddos-attacks-in-q2-2022/107025/>
- [80] Kaspersky. 2022. Kaspersky DDoS Attacks in Q3 2022. <https://securelist.com/ddos-report-q3-2022/107860/>
- [81] Kaspersky. 2022. Kaspersky DDoS Report in Q1 2022. <https://securelist.com/ddos-attacks-in-q1-2022/106358/>
- [82] Daniel Kopp, Christoph Dietzel, and Oliver Hohlfeld. 2021. DDoS Never Dies? An IXP Perspective on DDoS Amplification Attacks. In *Proc. of PAM*. Springer International Publishing, Cham, 284–301. https://doi.org/10.1007/978-3-030-72582-2_17
- [83] Daniel Kopp, Matthias Wichtlhuber, Ingmar Poese, Jair Santanna, Oliver Hohlfeld, and Christoph Dietzel. 2019. DDoS Hide & Seek: On the Effectiveness of a Booter Services Takedown. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 65–72. <https://doi.org/10.1145/33555369.3355590>
- [84] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. 2015. AmpPot: Monitoring and Defending Amplification DDoS Attacks. In *Proc. of RAID*. Springer Verlag, Berlin, Heidelberg, N.Y., 615–636. https://doi.org/10.1007/978-3-319-26362-5_28
- [85] John Kristoff. 2022. The DDoS Threat Landscape Report (NANOG 86). https://storage.googleapis.com/site-media-prod/meetings/NANOG86/4488/20221017_Kristoff_The_2022H1_Ddos_v1.pdf
- [86] Johannes Krupp, Michael Backes, and Christian Rossow. 2016. Identifying the Scan and Attack Infrastructures Behind Amplification DDoS Attacks. In *Proc. of the ACM SIGSAC CCS* (Vienna, Austria) (CCS '16). ACM, New York, NY, USA, 1426–1437. <https://doi.org/10.1145/2976749.2978293>
- [87] Johannes Krupp, Mohammad Karami, Christian Rossow, Damon McCoy, and Michael Backes. 2017. Linking Amplification DDoS Attacks to Booter Services. In *Proc. of the RAID*, Marc Dacier, Michael Bailey, Michalis Polychronakis, and Manos Antonakakis (Eds.). Springer International Publishing, Cham, 427–449.
- [88] Johannes Krupp and Christian Rossow. 2021. BGPeek-a-Boo: Active BGP-based Traceback for Amplification DDoS Attacks. In *Proc. of IEEE Euro Security & Privacy*. IEEE, Piscataway, NJ, USA, 423–439. <https://doi.org/10.1109/EuroSP51992.2021.00036>
- [89] Mirjam Kühne and John Kristoff. 2014. NTP Reflections. <https://labs.ripe.net/author/mirjam/ntp-reflections/>
- [90] Marc Kühner, Thomas Hupperich, Christian Rossow, and Thorsten Holz. 2014. Exit from Hell? Reducing the Impact of Amplification DDoS Attacks. In *Proc. of USENIX Security*. USENIX Association, Berkeley, CA, USA, 111–125. <https://www.usenix.org/conference/usenixsecurity14/technical-sessions/presentation/kuhner>
- [91] Marc Kühner, Thomas Hupperich, Christian Rossow, and Thorsten Holz. 2014. Hell of a Handshake: Abusing TCP for Reflective Amplification DDoS Attacks. In *Proc. of the USENIX WOOT* (San Diego, CA) (WOOT'14). USENIX Association, Berkeley, CA, USA, 4.
- [92] Warren "Ace" Kumari and Danny R. McPherson. 2009. Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF). RFC 5635. <https://doi.org/10.17487/RFC5635>
- [93] Black Lotus Labs. 2021. Tracking UDP Reflectors for a Safer Internet. <https://blog.lumen.com/tracking-udp-reflectors-for-a-safer-internet/>
- [94] Alexander Lex, Nils Gehlenborg, Hendrik Strobelt, Romain Vuillemot, and Hanspeter Pfister. 2014. UpSet: Visualization of Intersecting Sets. *IEEE Transactions on Visualization and Computer Graphics* 20, 12 (2014), 1983–1992. <https://doi.org/10.1109/TVCG.2014.2346248>
- [95] LINK11. 2023. LINK11 DDoS-REPORT 2022. <https://www.link11.com/en/download/ddos-report-2022/#download-detail-form>
- [96] Matthew Luckie, Robert Beverly, Ryan Koga, Ken Keys, Joshua A Kroll, and K Claffy. 2019. Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In *ACM SIGSAC Conference on*

- Computer and Communications Security (ACM CCS)*. ACM, New York, NY, USA, 465–480. <https://doi.org/10.1145/3319535.3354232>
- [97] Matthew Luckie, Ken Keys, Ryan Koga, Rob Beverly, and kc claffy. 2016. Spoofer Source Address Validation Measurement System. <http://spoofer.caida.org>
- [98] Lumen. 2022. Lumen Quarterly DDoS Report Q3 2022. <https://assets.lumen.com/is/content/Lumen/lumen-quarterly-ddos-report-q3-2022?Creativeid=6f6d4450-a936-4f14-9121-6a7b8f292392>
- [99] Lumen. 2022. Lumen Quarterly DDoS Report Q4 2022. <https://blog.lumen.com/q4-2022-lumen-ddos-quarterly-report/>
- [100] M3AAWG. 2017. M3AAWG Initial Recommendations: Arming Businesses Against DDoS Attacks. <http://www.m3aawg.org/DDoS-Recommendations-Business>
- [101] M3AAWG. 2023. Scholl Receives 2023 M3AAWG J.D. Falk Award for IP Spoofing Mitigation. <https://www.m3aawg.org/blog/2023JDFAward-TomScholl>
- [102] Mousa Taghizadeh Manavi. 2018. Defense mechanisms against Distributed Denial of Service attacks: A survey. *Computers & Electrical Engineering* 72 (2018), 26–38. <https://doi.org/10.1016/j.compeleceng.2018.09.001>
- [103] Inc. Merit Network. 2024. ORION Network Telescope: Observatory for cyber-Risk Insights and Outages of Networks. Website. <https://www.merit.edu/initiatives/orion-network-telescope/> Last Access: Nov 2023.
- [104] Jorge Merlino, Mohammed Asiri, and Neetesh Saxena. 2022. DDoS Cyber-Incident Detection in Smart Grids. *Sustainability* 14 (02 2022), 2730. <http://dx.doi.org/10.3390/su14052730>
- [105] Jelena Mirkovic and Peter Reiher. 2004. A taxonomy of DDoS attack and DDoS defense mechanisms. *ACM Sigcomm Computer Communication Review* 34, 2 (April 2004), 39–53. <https://doi.org/10.1145/997150.997156>
- [106] Asier Moneva and Rutger Leukfeldt. 2023. The effect of online ad campaigns on DDoS-attacks: A cross-national difference-in-differences quasi-experiment. *Criminology & Public Policy* 22, 4 (2023), 869–894. <https://doi.org/10.1111/1745-9133.12649>
- [107] D Moore, C Shannon, D Brown, G Voelker, and S Savage. 2006. Inferring Internet Denial-of-Service Activity. *ACM Transactions on Computer Systems* 24, 2 (May 2006), 115–139. <https://doi.org/10.1145/1132026.1132027>
- [108] Mortensen et al. 2007. *DDoS Open Threat Signaling (DOTS) requirements*. RFC 8612. IETF. <https://doi.org/10.17487/RFC4987>
- [109] Giovane C.M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Muller, Lan Wei, and Cristian Hesselman. 2016. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 255–270. <https://doi.org/10.1145/2987443.2987446>
- [110] Giovane C. M. Moura, Sebastian Castro, John Heidemann, and Wes Hardaker. 2021. TsuNAME: Exploiting Misconfiguration and Vulnerability to DDoS DNS. In *Proc. of ACM IMC (Virtual Event) (IMC '21)*. ACM, New York, NY, USA, 398–418. <https://doi.org/10.1145/3487552.3487824>
- [111] Giovane C. M. Moura, John Heidemann, Moritz Müller, Ricardo de O. Schmidt, and Marco Davids. 2018. When the Dike Breaks: Dissecting DNS Defenses During DDoS. In *Proc. of ACM IMC (Boston, MA, USA)*. ACM, New York, NY, USA, 8–21. <https://doi.org/10.1145/3278532.3278534>
- [112] Giovane C. M. Moura, Cristian Hesselman, Gerald Schapman, Nick Boerman, and Octavia de Weerd. 2020. Into the DDoS maelstrom: A longitudinal study of a scrubbing service. In *5th International Workshop on Traffic Measurements for Cybersecurity (WTMC 2020)*. IEEE, Piscataway, NJ, USA, 550–558. <https://doi.org/10.1109/EuroSPW51379.2020.00081>
- [113] 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 *Proc. of ACM IMC*. ACM, New York, NY, USA, 435–448. <https://doi.org/10.1145/3355369.3355593>
- [114] Marcin Nawrocki, Raphael Hiesgen, Thomas C. Schmidt, and Matthias Wählisch. 2021. QUICsand: Quantifying QUIC Reconnaissance Scans and DoS Flooding Events. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 283–291. <https://doi.org/10.1145/3487552.3487840>
- [115] Marcin Nawrocki, Mattijs Jonker, Thomas C. Schmidt, and Matthias Wählisch. 2021. The Far Side of DNS Amplification: Tracing the DDoS Attack Ecosystem from the Internet Core. In *Proc. of ACM IMC*. ACM, New York, NY, USA, 419–434. <https://doi.org/10.1145/3487552.3487835>
- [116] Marcin Nawrocki, Maynard Koch, Thomas C. Schmidt, and Matthias Wählisch. 2021. Transparent Forwarders: An Unnoticed Component of the Open DNS Infrastructure. In *Proc. of ACM CoNEXT*. ACM, New York, NY, USA, 454–462. <https://doi.org/10.1145/3485983.3494872> Continued data collection: <https://odns.secnov.net/data>.
- [117] Marcin Nawrocki, John Kristoff, Chris Kanich, Raphael Hiesgen, Thomas C. Schmidt, and Matthias Wählisch. 2023. SoK: A Data-driven View on Methods to Detect Reflective Amplification DDoS Attacks Using Honeypots. In *Proc. of IEEE Euro Security & Privacy (Delft, Netherlands)*. IEEE, Piscataway, NJ, USA, 576–591. <https://doi.org/10.1109/EuroSP57164.2023.00041>
- [118] Marcin Nawrocki, Pouyan Fotouhi Tehrani, Raphael Hiesgen, Jonas Mücke, Thomas C. Schmidt, and Matthias Wählisch. 2022. On the Interplay between TLS Certificates and QUIC Performance. In *Proc. of ACM CoNEXT*. ACM, New York, NY, USA, 204–213. <https://dl.acm.org/doi/10.1145/3555050.3569123>
- [119] NBIP. 2023. DDoS Attack Figures from the First Quarter 2023. <https://www.nbip.nl/wp-content/uploads/2023/04/NBIP%20-%20Infographic%20-%20DDoS%20data%20-%202023%20Q1.pdf>
- [120] NBIP. 2023. DDoS Attack Figures from the Fourth Quarter 2022. <https://www.nbip.nl/wp-content/uploads/2023/01/NBIP%20-%20Infographic%20-%20DDoS%20data%20-%20Q4%202022%20%5BEN%5D.pdf>
- [121] NBIP. 2023. DDoS Attack Figures from the Second Quarter 2023. <https://www.nbip.nl/wp-content/uploads/2023/07/NBIP-Infographic-DDoS-data-Q2-2023-EN.pdf>
- [122] Netscout. 2021. NETSCOUT Threat Intelligence Report 2H 2021. https://www.netscout.com/sites/default/files/2022-03/ThreatReport_2H2021_WEB.pdf
- [123] Netscout. 2022. TP240PhoneHome Reflection/Amplification DDoS Attack Vector. <https://www.netscout.com/blog/aser/tp240phonehome-reflectionamplification-ddos-attack-vector>
- [124] Netscout. 2023. 5th Anniversary DDoS Threat Intelligence Report: Unveiling the New Threat Landscape. <https://www.netscout.com/threatreport/wp-content/uploads/2023/04/Threat-Report-2H2022.pdf>
- [125] Netscout. 2023. NETSCOUT DDoS Attack Vectors and Methodology. <https://www.netscout.com/resources/threat-report/threat-intelligence-report-ddos-attack-vectors-and-methodology>
- [126] Netscout. 2023. Service Location Protocol (SLP) Reflection/Amplification Attack Mitigation Recommendations. <https://www.netscout.com/blog/aser/slp-reflectionamplification-ddos-attack-vector>
- [127] Netscout. 2023. Unveiling the New Threat Landscape. <https://www.netscout.com/threatreport/ddos-threat-intelligence-report/>
- [128] NETSCOUT. 2023. Unveiling the New Threat Landscape. <https://web.archive.org/web/20230413213001/https://www.netscout.com/threatreport/ddos-threat-intelligence-report/#global-defense>
- [129] NEXUSGUARD. 2023. DDoS Statistical Report for 1HY 2023. <https://www.nexusguard.com/threat-report/ddos-statistical-report-for-1hy-2023>
- [130] NEXUSGUARD. 2023. DDoS Statistical Report for 2022. <https://blog.nexusguard.com/threat-report/ddos-statistical-report-for-2022>
- [131] Nimrod Levy and John Schiel and John A Schiel. 2017. Bi-lateral Security Management Framework (aka DDoS peering). https://pc.nanog.org/static/published/meetings/NANOG71/1447/20171003_Levy_Operationalizing_Isp_v2.pdf
- [132] Nokia. 2022. Nokia Deepfield Network Intelligence Report DDoS in 2021. https://onestore.nokia.com/asset/211059?_ga=2.234339031.813264975.1691960553-1225881009.1691960553
- [133] Nokia. 2022. The Changing DDoS Threat Landscape. <https://www.nokia.com/networks/security/ddos-security/the-changing-ddos-threat-landscape/>
- [134] Nokia. 2023. Nokia Threat Intelligence Report 2023. <https://www.nokia.com/networks/security-portfolio/threat-intelligence-report/>
- [135] Arman Noroozian, Maciej Korczyński, Carlos Hernandez Gañan, Daisuke Makita, Katsunari Yoshioka, and Michel Van Eeten. 2016. Who Gets the Boot? Analyzing Victimization by DDoS-as-a-Service. In *Proc. of RAID*. Springer Verlag, Berlin, Heidelberg, N.Y., 368–389. https://doi.org/10.1007/978-3-319-45719-2_17
- [136] NSFOCUS. 2023. 2022 Global DDoS Attack Landscape Report. <https://nsfocusglobal.com/company-overview/resources/2022-global-ddos-attack-landscape-report/>
- [137] Riyadh Rahef Nuiaa, Selvakumar Manickam, Ali Hakem Alsaedi, and Es-rra Saleh Alomari. 2022. Enhancing the Performance of Detect DRDoS DNS Attacks Based on the Machine Learning and Proactive Feature Selection (PFS) Model. *IAENG International Journal of Computer Science* 49, 2 (2022), 14 pages.
- [138] Central District of California. 2023. SEIZURE WARRANT BY TELEPHONE OR OTHER RELIABLE ELECTRONIC MEANS. United States District Court. <https://krebsonsecurity.com/wp-content/uploads/2023/05/Booter-seizure-warrant-Tucows.pdf>
- [139] Ofcom. 2003. General statement of policy under section 105Y of the Communications Act 2003. , 43 pages. https://www.ofcom.org.uk/_data/assets/pdf_file/0030/253677/General-statement-of-policy-under-section-105Y-of-the-Communications-Act-2003.pdf
- [140] Eric Osterweil, Pouyan Fotouhi Tehrani, Thomas C. Schmidt, and Matthias Wählisch. 2022. From the Beginning: Key Transitions in the First 15 Years of DNSSEC. *Transactions on Network and Service Management (TNSM)* 19, 4 (December 2022), 5265–5283. <https://doi.org/10.1109/TNSM.2022.3195406>
- [141] Palo Alto. 2023. Unit 42 INCIDENT RESPONSE REPORT 2022. https://www.paloaltonetworks.com/content/dam/pan/en_US/assets/pdf/reports/2022-unit42-incident-response-report-final.pdf
- [142] Ruoming Pang, Vinod Yegneswaran, Paul Barford, Vern Paxson, and Larry Peterson. 2004. Characteristics of Internet Background Radiation. In *Proc. of the 4th ACM SIGCOMM conference on Internet measurement (Taormina, Sicily, Italy)*. ACM, New York, NY, USA, 27–40. <http://doi.acm.org/10.1145/1028788.1028794>
- [143] Vern Paxson. 2001. An Analysis of Using Reflectors for Distributed Denial-of-Service Attacks. *ACM Sigcomm Computer Communication Review* 31, 3 (2001), 38–47. <https://doi.org/10.1145/505659.505664>

- [144] Niels Provos and Thorsten Holz. 2008. *Virtual Honeypots. From Botnet Tracking to Intrusion Detection* (2nd ed.). Addison-Wesley, Upper Saddle River, NJ.
- [145] Qrator. 2022. Q1 2022 DDoS Attacks and BGP Incidents. https://blog.qrator.net/en/q1-2022-ddos-attacks-and-bgp-incidents_155/
- [146] Qrator. 2022. Q2 2022 DDoS attacks and BGP incidents. <https://qratorlabs.medium.com/q2-2022-ddos-attacks-and-bgp-incidents-efe7e5c1395a>
- [147] Qrator. 2022. Q3 2022 DDoS attacks and BGP incidents. https://blog.qrator.net/en/q3-2022-ddos-attacks-and-bgp-incidents_158/
- [148] Qrator. 2023. Q4 2022 DDoS Attacks and BGP Incidents. https://blog.qrator.net/en/q4-2022-ddos-attacks-and-bgp-incidents-report_163/
- [149] Radware. 2023. Radware Global Threat Analysis Report 2022. <https://www.radware.com/2022-2023-global-threat-analysis-report/>
- [150] Raju Rajan, Jim Boyle, Arun Sastry, Ron Cohen, David Durham, and Shai Herzog. 2000. The COPS (Common Open Policy Service) Protocol. RFC 2748. <https://doi.org/10.17487/RFC2748>
- [151] Research and Markets. 2023. Global DDoS Protection & Mitigation Security Market Report to 2027: Players Include CloudFlare, Corero, DDoS-Guard, Fastly and Fortinet. Website. <https://www.prnewswire.com/news-releases/global-ddos-protection--mitigation-security-market-report-to-2027-players-include-cloudflare-corero-ddos-guard-fastly-and-fortinet-301752182.html>
- [152] Rich Compton and Thomas Bowlby and Taylor Harris and Pratik Lotia. 2019. eBGP Flowspec Peering for DDoS Mitigation. https://pc.nanog.org/static/published/meetings/NANOG75/1887/20190219_Compton_Ebgp_Flowspec_Peering_v1.pdf
- [153] RioRey. 2015. RioRey Taxonomy DDoS V2.9. https://static1.squarespace.com/static/5548bab5e4b08ecb6652391c1c/5d8d0538f8cc9e3295187a76/1569523017682/RioRey_Taxonomy_DDoS_V2.9.pdf
- [154] A S M Rizvi, Leandro Bertholdo, João Ceron, and John Heidemann. 2022. Anycast Agility: Network Playbooks to Fight DDoS. In *Proc. of USENIX Security*. USENIX Association, Boston, MA, 4201–4218. <https://www.usenix.org/conference/usenixsecurity22/presentation/rizvi>
- [155] Christian Rossow. 2014. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *Proc. of NDSS*. Internet Society, Reston, VA, USA, 15 pages. <https://doi.org/10.14722/ndss.2014.23233>
- [156] Fabrice J. Ryba, Matthew Orlinski, Matthias Wählisch, Christian Rossow, and Thomas C. Schmidt. 2015. *Amplification and DRDoS Attack Defense – A Survey and New Perspectives*. Technical Report arXiv:1505.07892. Open Archive: arXiv.org. <http://arxiv.org/abs/1505.07892>
- [157] Ravjot Singh Samra and Marinho Barcellos. 2023. DDoS2Vec: Flow-level characterisation of volumetric DDoS attacks at scale. *Proc. ACM Netw.* 2, CoNEXT (Dec. 2023), 25 pages.
- [158] Matthew Sargent, John Kristoff, Vern Paxson, and Mark Allman. 2017. On the Potential Abuse of IGMP. *ACM Sigcomm Computer Communication Review* 47, 1 (Jan 2017), 27–35. <https://doi.org/10.1145/3041027.3041031>
- [159] Kyle Schomp, Onkar Bhardwaj, Eymen Kurdoglu, Mashooq Muhaimen, and Ramesh K. Sitaraman. 2020. Akamai DNS: Providing Authoritative Answers to the World’s Queries. In *Proc. of ACM SIGCOMM*. ACM, New York, NY, USA, 465–478. <https://doi.org/10.1145/3387514.3405881>
- [160] ShadowServer. 2023. DDoS | The Shadowserver Foundation. <https://www.shadowserver.org/topics/ddos/>
- [161] ShadowServer. 2023. The Shadowserver Foundation: Network Reporting. <https://www.shadowserver.org/what-we-do/network-reporting/>
- [162] Shane Alcock and Alistair King. 2010. Corsaro Version 3, flow analysis tools. <https://github.com/CAIDA/corsaro3/>
- [163] Stephen M. Specht and Ruby B. Lee. 2004. Distributed Denial of Service: Taxonomies of Attacks, Tools, and Countermeasures. In *Proc. of the PADS*. ISCA, Winona, MN, USA, 543–550.
- [164] Splunk. 2023. Denial-of-Service Attacks: History, Techniques & Prevention. https://www.splunk.com/en_us/blog/learn/dos-denial-of-service-attacks.html
- [165] K. Sriram, D. Montgomery, and J. Haas. 2020. BCP 84, RFC 8704: Enhanced Feasible-Path Unicast Reverse Path Forwarding. <https://www.rfc-editor.org/info/bcp84>
- [166] Microsoft Azure Network Security Team. 2023. 2022 in Review: DDoS Attack Trends and Insights. <https://www.microsoft.com/en-us/security/blog/2023/02/21/2022-in-review-ddos-attack-trends-and-insights/>
- [167] Daniel R Thomas, Richard Clayton, and Alastair R Beresford. 2017. 1000 days of UDP amplification DDoS attacks. In *2017 APWG Symposium on Electronic Crime Research (eCrime)*. IEEE, Piscataway, NJ, USA, 79–84. <https://doi.org/10.1109/ECRIME.2017.7945057>
- [168] Daniel R. Thomas, Sergio Pastrana, Alice Hutchings, Richard Clayton, and Alastair R. Beresford. 2017. Ethical issues in research using datasets of illicit origin. In *Proc. of ACM IMC* (London, UK). ACM, New York, NY, USA, 445–462. <https://doi.org/10.1145/3131365.3131389>
- [169] Tony Miu Tung, Chenxu Wang, and Jinhe Wang. 2018. Understanding the Behaviors of BGP-based DDoS Protection Services. In *Proc. of NSS*. Springer International Publishing, Cham, 463–473.
- [170] U.K. Office of Communications (Ofcom). 2023. *Automatic compensation: What you need to know*. Ofcom. <https://www.ofcom.org.uk/phones-telecoms-and-internet/advice-for-consumers/costs-and-billing/automatic-compensation-need-know>
- [171] US Attorney’s Office. 2023. Federal Authorities Seize 13 Internet Domains Associated with ‘Booter’ Websites that Offered DDoS Computer Attack Services. <https://www.justice.gov/usao-cdca/pr/federal-authorities-seize-13-internet-domains-associated-booter-websites-offered-ddos>
- [172] U.S. Federal Communication Commission. 2018. Restoring Internet Freedom. 33 FCC Rcd 311 (1).
- [173] U.S. Federal Communication Commission. 2022. Broadband Consumer Labels. <https://www.fcc.gov/broadbandlabels>
- [174] T. van den Hout, C. Hesselman, R. Poortinga, R. Yazdani, M. Jonker, C. Papachristos, P. De Lutiis, M. Baltatu, and B. Rodrigues. 2022. *DDoS Clearing House Cookbook, CONCORDIA Deliverable D3.6*. CONCORDIA. <https://ddosclearinghouse.eu/cookbook> Accessed on 15 July 2023.
- [175] Olivier van der Toorn, Johannes Krupp, Mattijs Jonker, Roland van Rijswijk-Deij, Christian Rossow, and Anna Sperotto. 2021. ANYway: Measuring the Amplification DDoS Potential of Domains. In *Proc. of the CNSM*. IEEE, Piscataway, NJ, USA, 500–508. <https://doi.org/10.23919/CNSM52442.2021.9615596>
- [176] Daniel Wagner, Daniel Kopp, Matthias Wichtlhuber, Christoph Dietzel, Oliver Hohlfeld, Georgios Smaragdakis, and Anja Feldmann. 2021. United We Stand: Collaborative Detection and Mitigation of Amplification DDoS Attacks at Scale. In *Proc. of ACM CCS*. ACM, New York, NY, USA, 970–987. <https://doi.org/10.1145/3460120.3485385>
- [177] Matthias Wichtlhuber, Eric Strehle, Daniel Kopp, Lars Prepens, Stefan Stegmueller, Alina Rubina, Christoph Dietzel, and Oliver Hohlfeld. 2022. IXP Scrubber: Learning from Blackholing Traffic for ML-Driven DDoS Detection at Scale. In *SIGCOMM*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3544216.3544268>
- [178] E. Wustrow, M. Karir, M. Bailey, F. Jahanian, and G. Huston. 2010. Internet Background Radiation Revisited. In *Proc. of ACM IMC* (Melbourne, Australia). ACM, NY, USA, 62–74. <https://doi.org/10.1145/1879141.1879149>
- [179] Zayo. 2022. A Look at Recent DDoS Attacks and the Cyberattack Landscape in 2022 So Far. <https://www.zayo.com/resources/ddos-attack-trends/>
- [180] Zayo. 2023. Protecting Your Business From Cyber Attacks: The State of DDoS Attacks DDoS (Insights From Q1 & Q2, 2023). <https://go.zayo.com/zayo-ddos-protection-ebook/>

A Ethics

We considered the following ethical concerns.

Harmful data collection. Our analysis is based on data sets extensively used in prior work, not of itself a justification [168], however, we do no additional harm. DDoS *honeypots* deploy safeguards to avoid participating in attacks while still collecting as much information about attackers as possible (often reducing attack traffic). *Network telescopes* are purely passive instruments and do not participate in any activity. *On path flow monitoring* data from IXPs and DDoS mitigation providers is already collected for operational purposes, which are supported by the insights from research reusing the data. We worked with a wide range of stakeholders to minimize unanticipated risks [168].

Leaking personal information. Our analysis focuses on aggregate trends and correlations. We do not reveal individual IP addresses, or any personal information. Correlations with customer information was done by a party already trusted by the customer without leaking information into the collected data set. Some datasets are shared in a controlled way with researchers by their original sources [24] enabling reproducibility while maintaining safeguards.

Data anonymization. The data from industry does not include any personal identifying information (PII) as we received attack counts of different granularity that we aggregate to weeks. While we have access to attack event data for the research observatories, we do not include any PII information in the paper. As such, no anonymization was required to present our results.

Table 3: List of documents assessed.

Company	Included	Omitted
A10	[1]	
Akamai	[4, 5]	
Alibaba Cloud		[6]
AWS		[11]
Arelion	[9]	
Cloudflare	[26]	[25, 27–29]
Comcast	[33]	[32]
Corero	[34]	[35, 36]
CrowdStrike		[38]
DDoS-Guard	[41, 42]	
F5	[53]	
Fastly		[55, 56]
Fortinet		[59]
Huawei	[72]	
Imperva	[73]	
Kaspersky	[80]	[79, 81]
LINK11	[95]	
Lumen	[99]	[93, 98]
Microsoft	[166]	
NBIP	[120]	[119, 121]
Netscout	[124]	[122, 125]
NexusGuard	[130]	[129]
Nokia	[134]	[37, 132]
NSFocus	[136]	[133]
Palo Alto		[141]
Qrator	[148]	[145–147]
Radware	[149]	
RioRey		[153]
Splunk		[164]
Zayo	[180]	[179]

B Artifacts

We contribute two artifacts that categorize related work on DDoS attacks. We believe this work can serve the community as a reference.

Industry: DDoS report survey. Throughout the DDoS industry, companies report on their observation on DDoS attacks. We present a deep dive into reports released around 2022 in §3 alongside supplemental materials in Appendix E and online [13].

Academia: related work taxonomy. Work on DDoS in research has been extensive. §8 presents a condensed overview that is systematically categorized based on research field and data sets in Appendix C.

C Taxonomy “Mindmap” of DDoS Literature

We provide a graphical (“mindmap”) taxonomy of the extensive relevant literature in the last few years on the two dominant classes of attacks that we study (Direct-Path and Reflection/Amplification attacks). Figure 11 illustrates the many dimensions of the problem that have received focused attention. The figure is not exhaustive – there are hundreds of other papers on DDoS, which we expect will

mostly fall within the themes included within this taxonomy. Our takeaway from this analysis is that while there is an abundance of research activity related to DDoS, there is little effort to find a convergent position on how effective current defenses are against the threat.

D Honeypot: NewKid

Figure 12 shows attack counts observed by the NewKid honeypot, which consists of a single sensor whose observations are erratic, although still reasonably consistent with the macroscopic data sets we analyzed. For example, periods of high attacks around mid-2020 and mid-2022 match peaks in attacks observed by Hopscotch Figure 3(a) and AmpPot Figure 3(b).

E Summary of Industry Reports

Table 3 lists vendors that published reports, and highlights the reports we discuss in this paper. We created an extensive table [13] containing information we extracted from industry reports. The table is a living artifact of this work; we invite interested community members (including report authors!) to expand the table with additional reports, or annotations of existing reports. We imagine this table as a potential tool to support pursuit of some measure of *community consensus* regarding the DDoS ecosystem. Further, we hope it provides some incentive for industry players to engage in conversation to elucidate their results. An accompanying git repository [13] contains the table in PDF and other formats and all the files used as sources for our analysis.

The columns in this comprehensive table [13] include: vendor, Document Description, title, format, period of analysis, attack counts, reflection/amplification vectors, direct path vectors, other vectors, changes during 2022, attack duration/size, attack intensity, carpet bombing, multi-vector, target, vantage points/sources.

F Quarterly Correlations

§6.3 discusses the correlations of long-term trends and summarizes them in Figure 6. Figure 14 shows a coarse-grained view on pair-wise correlations. Each box summarizes quarterly Spearman correlations (18 values over 4.5 years) for the pair of observatories marked in the matrix below. Vertical bars mark the median and the circles the mean. The coloring scheme matches the UpSet plots in §7, *blue*: direct-path observatories, *yellow*: reflection-amplification observatories, and *red*: correlations across attack types.

Most correlations are not stable across quarters, *i.e.*, the whiskers cover a large part of the possible correlations (-1 to +1). Two exceptions are the pair AmpPot & IXP Blackholing (RA) and the pair Akamai Prolexic (DP) & Akamai Prolexic (RA), which only have outliers in the anti-correlation range, but generally show a high correlation with each other.

The whiskers for box that correlate observatories of the same attack type tend to have less variance. While they reach less into anti-correlations overall, they often mix periods of correlation and anti-correlation. The boxes for mixed correlations tend to be focused more around 0, although the DDoS mitigation services (Netscout, Akamai) both show higher correlations among their own two time series.

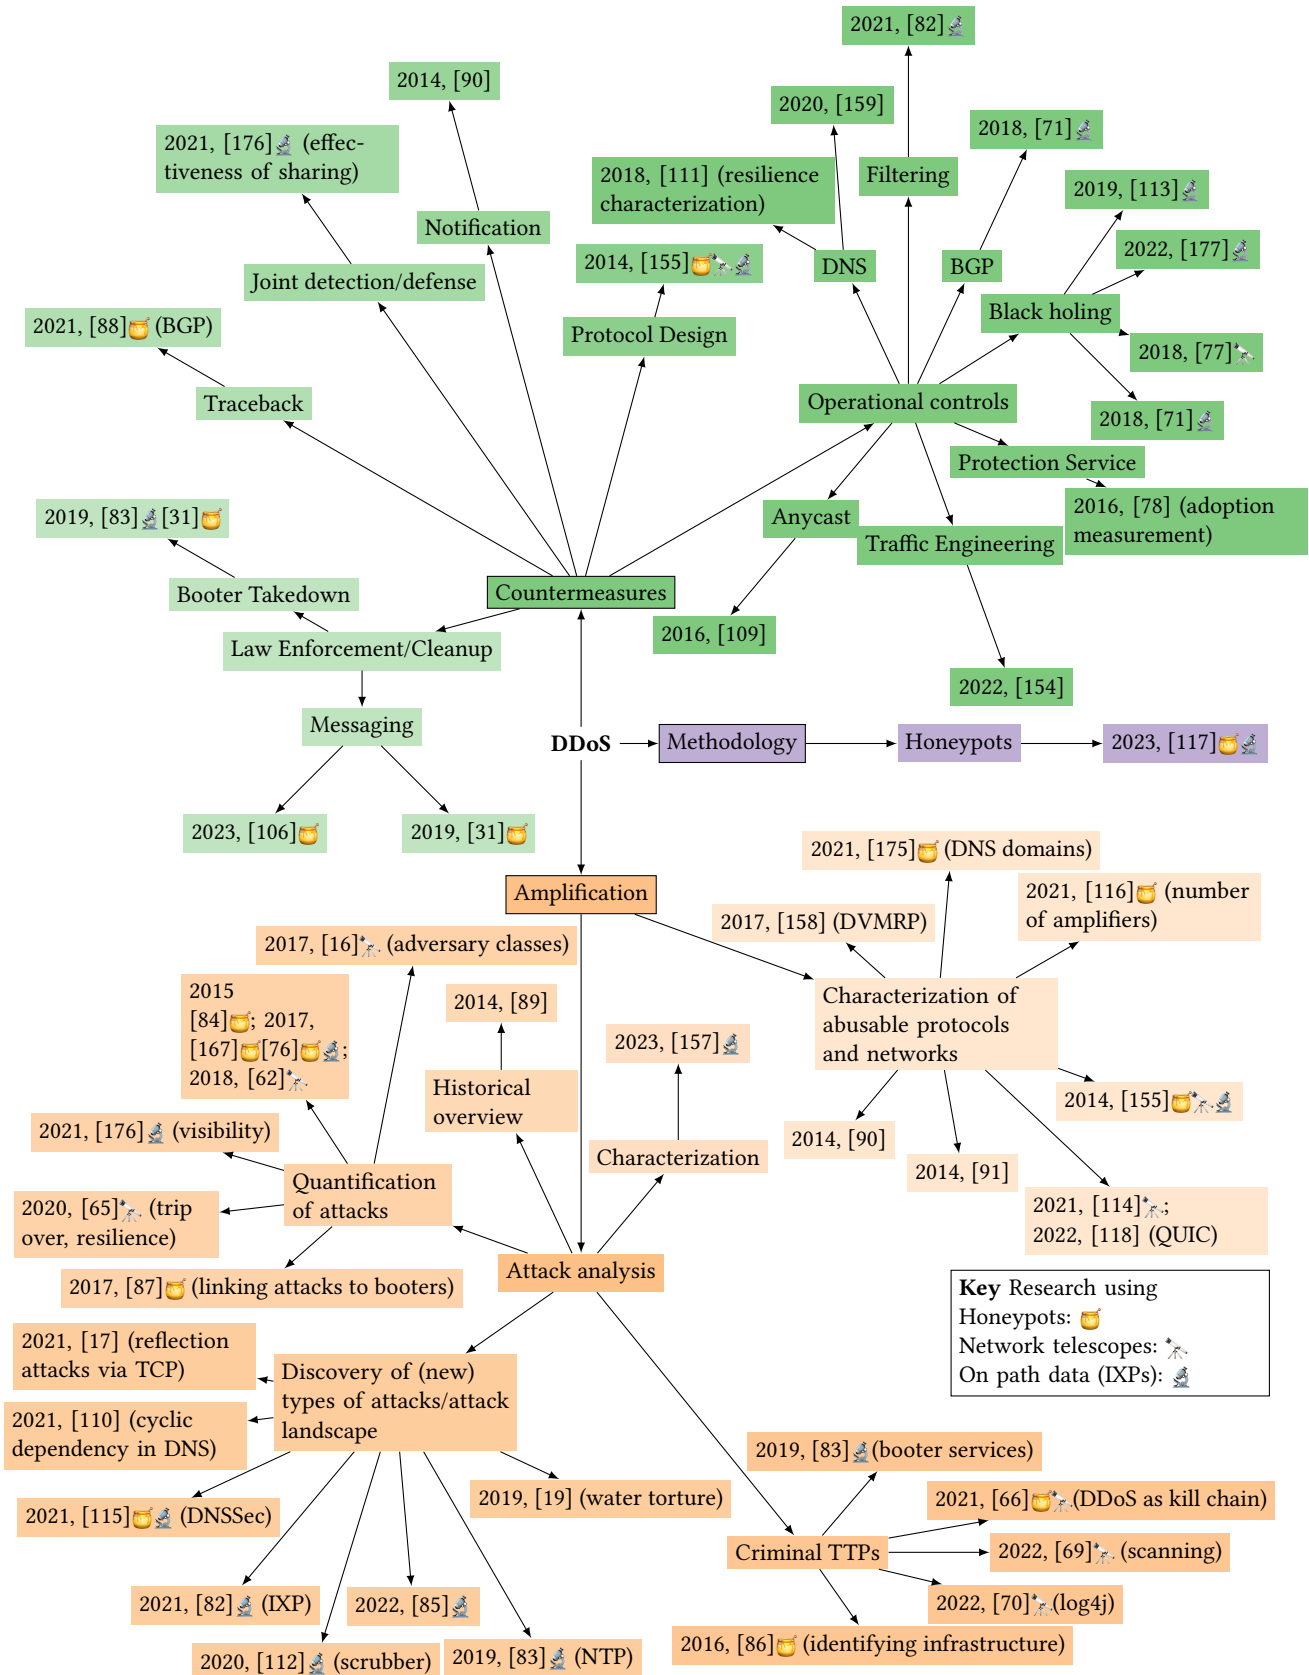


Figure 11: Taxonomy of related work in DDoS.

Rank	Provider	ASN	Tuples	Share
1	OVH	16276	28,769	18.80%
2	Hetzner	24940	7872	5.14%
3	Amazon	16509	4123	2.69%
4	Microsoft	8075	3125	2.04%
5	Google	396982	2898	1.89%
6	Cloudflare	13335	2427	1.59%
7	China Unicom	4837	2421	1.58%
8	Digitalocean	14061	2081	1.36%
9	Nuclearfallout	14586	1885	1.23%
10	Alibaba	37963	1847	1.21%

Table 4: Top 10 ASes that were observed in all our four research observatories (ORION, UCSD, Hopscotch, AmpPot).




Figure 12: Normalized attack trends observed by of the NewKid honeypot. The peak in mid 2022 rises up to 33.




Figure 13: Relative share of targets observed by academic observatories and confirmed by industrial baseline data from Akamai.

G Akamai Target Overlap

Figure 13 shows how many of the DDoS target tuples (date, IP address) observed by academia (Figure 7) were also observed by Akamai. Similar to Netscout (Figure 9) the shares were higher for attacks visible at multiple observatories. These are likely highly visible attacks.

H Overlap of Targeted ASes

Table 4 shows the top 10 ASes observed as DDoS targets (date, IP) at ORION, UCSD, Hopscotch, and AmpPot (§7.1). All are labeled as

hosting ASes except for Microsoft (business), China Unicom (ISP), and Alibaba (business).

I Detecting Carpet Bombing/Prefix Attacks in Honeypot Data

Spreading one attack over many IP addresses (variously called “carpet bombing”, “prefix attacks” or “horizontal attacks”) makes it more difficult to block, or even identify as a single attack. Honeypots in particular struggle with this challenge because they are generally comprised of many IP addresses, and attackers deploy a range of strategies to select target IP addresses (randomized, sequential, bursty) and so honeypots may not observe even one packet for every targeted address. Our approach to inferring (effectively reconstructing) such attacks builds on prior work [167, Appendix A-C] to aggregate attacks within the same IP prefix. Our method essentially finds the longest BGP-routed prefix (from /11 to /28) that covers the attack.

Our approach does not aggregate attacks that span multiple IP address block allocations (from Regional Internet Registry (RIR) data). This means that when an attacker targets many blocks of IP addresses that are allocated to the same AS, because they are targeting an ISP rather than a customer, this is recorded as many attacks rather than a single attack. Such attacks targeted Brazil using SSDP in mid-2022 causing the spikes in Figure 3(a) and 3(b). Algorithms that detect attacks targeting entire ASes might remove this noise. We provided the details of our algorithm in a script we shared with the Cambridge Cybercrime Center for use in their honeypot infrastructure; they will make it available to researchers on request [24].

J RSDoS Inference from Network Telescopes

We provide implementation details and references for the RSDoS analysis in §6. We used CAIDA’s Corsaro tools [162], which is based on [107]. Corsaro uses a *flow identifier* to group packets into flows. A packet *threshold* and *timeout* discern which of these flows are part of an attack and when this attack stops.

- (1) **Flow identifier:** The tuple (protocol, source IP) identifies a flow. In code, packets are matched in two steps: (i) the protocol selects a hashmap, (ii) the source IP identifies the flow within it. Source and destination ports are aggregated as part of the data rather than part of the key⁶.
- (2) **Threshold:** To be considered an attack a flow must have a minimum of 25 packets from a single source IP and last for 60 seconds. Additionally, an attack flow must (at one point) meet a packet rate of at least 30 packets across a 60-second window, which slides every 10 seconds⁷.
- (3) **Timeout:** Corsaro counts packets in intervals of 300 seconds. After an interval with no new packets an attack flow is finished.⁸

⁶https://github.com/CAIDA/corsaro3/blob/master/libcorsaro/plugins/corsaro_dos.c#L1014

⁷<https://github.com/CAIDA/corsaro3/wiki/DoS-Plugin#configuration>

⁸https://github.com/CAIDA/corsaro3/blob/master/libcorsaro/plugins/corsaro_dos.c#L1251

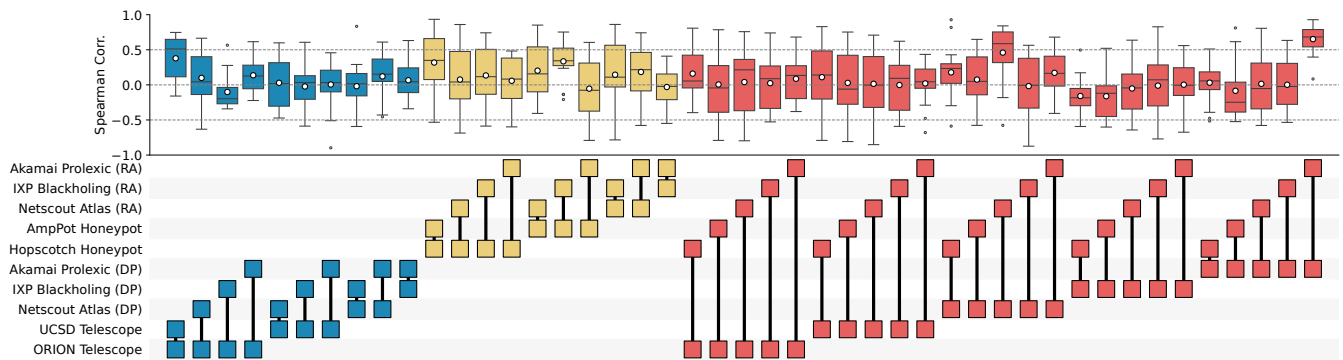


Figure 14: Quarterly pair-wise correlations across observatories from academia and industry. Vertical bars mark the median and the circles the mean.

Once both thresholds (packet count, packet rate) have been met (at any point in the flow) that flow counts as an attack for the rest of its lifetime. Any number of packets is enough to maintain it until the flow times out because no new packets arrive. (We are

not sure why the authors evolved the code this way.) The default values mentioned are explicitly set in the Corsaro config file used in the analysis (https://github.com/CAIDA/corsaro3/blob/master/libcorsaro/plugins/corsaro_dos.c#L1265).