United We Stand: Collaborative Detection and Mitigation of Amplification DDoS Attacks at Scale

Daniel Wagner

DE-CIX

Max Planck Institute for Informatics

Daniel Kopp

DE-CIX

Matthias Wichtlhuber

DE-CIX

Christoph Dietzel

DE-CIX

Max Planck Institute for Informatics

Oliver Hohlfeld

Brandenburg University of
Technology

Georgios Smaragdakis

TU Delft

Anja Feldmann

Max Planck Institute for Informatics

ABSTRACT

Amplification Distributed Denial of Service (DDoS) attacks' traffic and harm are at an all-time high. To defend against such attacks, distributed attack mitigation platforms, such as traffic scrubbing centers that operate in peering locations, e.g., Internet Exchange Points (IXP), have been deployed in the Internet over the years. These attack mitigation platforms apply sophisticated techniques to detect attacks and drop attack traffic locally, thus, act as sensors for attacks. However, it has not yet been systematically evaluated and reported to what extent coordination of these views by different platforms can lead to more effective mitigation of amplification DDoS attacks. In this paper, we ask the question: "Is it possible to mitigate more amplification attacks and drop more attack traffic when distributed attack mitigation platforms collaborate?"

To answer this question, we collaborate with eleven IXPs that operate in three different regions. These IXPs have more than 2,120 network members that exchange traffic at the rate of more than 11 Terabits per second. We collect network data over six months and analyze more than 120k amplification DDoS attacks. To our surprise, more than 80% of the amplification DDoS are not detected locally, although the majority of the attacks are visible by at least three IXPs. A closer investigation points to the shortcomings, such as the multi-protocol profile of modern amplification attacks, the duration of the attacks, and the difficulty of setting appropriate local attack traffic thresholds that will trigger mitigation. To overcome these limitations, we design and evaluate a collaborative architecture that allows participant mitigation platforms to exchange information about ongoing amplification attacks. Our evaluation shows that it is possible to collaboratively detect and mitigate the majority of attacks with limited exchange of information and drop as much as 90% more attack traffic locally.

CCS CONCEPTS

• Security and privacy → Network security.

KEYWORDS

Cyberattacks, DDoS, amplification attacks, IXP

ACM Reference Format:

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 *Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security (CCS '21), November 15–19, 2021, Virtual Event, Republic of Korea*. ACM, New York, NY, USA, 18 pages. <https://doi.org/10.1145/3460120.3485385>

1 INTRODUCTION

As our commercial and social activity is increasingly moving online, due to the ongoing pandemic as well [22], cyberattacks are more frequent and devastating [7, 10, 44, 64]. By recent measures, the damage due to cyberattacks in 2020 alone is estimated to one Trillion USD [61], double than the damage in 2018.

Among the most popular cyberattacks are these that target online services. To generate voluminous attack traffic, attackers that are politically or commercially motivated, compromise computers around the globe. These, so-called Distributed Denial of Service (DDoS) attacks, are well orchestrated and typically exploit vulnerabilities of computing systems [49, 50]. In recent years, it is even possible to lease resources or compromised machines for attacks using booter services that are available in the public or dark market [16, 37]. Studies of DDoS attacks [7, 28] have shown that attackers often first test the operation of their attack system by launching low volume attacks, or targeting low profile targets before launching the fully-fledged attack.

In terms of attack volume, recent studies have shown an exponential surge [8, 44]. Until 2012, the largest attack reported was less than 100 Gbps. In recent years, attacks with orders of magnitude higher traffic (up to 2 Tbps), have been reported, e.g., the 2018 memcached attack [3]. At the same time, attacks are becoming more sophisticated. Analysis of recent attacks shows that attackers can generate attacks with hundreds of millions of packets per second (Mpps) [39]. Thus, attackers are not only able to launch voluminous attacks but also attacks that require additional computation

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

CCS '21, November 15–19, 2021, Virtual Event, Republic of Korea.

© 2021 Copyright is held by the owner/author(s).

ACM ISBN 978-1-4503-8454-4/21/11.

<https://doi.org/10.1145/3460120.3485385>

resources from defenders. The most successful of these attacks are *amplification attacks*, i.e., the attacker can create harm to the target that is up to 50,000 times higher than the original attack traffic the attacker generates by utilizing services like DNS or NTP as reflectors [18, 57]. If this trend continues, it is expected that in the next years attacks more than 10 Tbps and multiple Gpps will be reported. At this scale, no single infrastructure provider alone can defend them.

The response from the industry was to introduce DDoS attack detection and mitigation platforms deployed at various locations in the Internet [31, 33]. Among them, traffic scrubbing centers analyze incoming traffic and apply rules to detect DDoS. Detected attack traffic is then dropped locally [5, 24, 33, 63]. The required processing per packet or per flow processing for deep packet inspection increases the detection and mitigation cost of attacks, does not scale well, imposes performance penalties and is vulnerable to evasion tactics [32]. Other techniques, such as Remote Triggered Black-hole filtering [12] are more aggressive and scalable, but require the detection of attack from a separate system. Unfortunately, these coarse-grained mitigation techniques also drop legitimate traffic to the destination under attack, and, thereby cause collateral damage. FlowSpec allows for fine-grained mitigation. However, although it has adopted in intra-domain environments [9, 58], it has not been popular in inter-domain environments as it requires sharing of computational and network resources across independently administrated networks. More recently, finer-grained blackholing has been proposed to address the limitations [20]. Both traffic scrubbing centers and blackholing functionality are present and readily available in peering locations such as Internet Exchange Points (IXPs) [11, 25]. To the best of our knowledge this is not the case for FlowSpec. Today, it is well accepted that the previously mentioned DDoS mitigation techniques are effective. However, their performance has typically been evaluated at a single location [19, 20, 36, 37, 62]. For a small number of attacks, a previous work has utilized publicly available vantage points to infer the efficacy of Remote-triggered Blackholing [25].

In this paper, we investigate whether coordination among attack detection and mitigation platform can be even more effective to: (1) detect and mitigate more amplification DDoS attacks, and (2) drop more amplification DDoS attack traffic locally that otherwise is carried to either be dropped later or to cause harm.

Our main observation is that the distributed nature in which reflectors are exploited for launching reflection DDoS attacks can be leveraged to realize better DDoS detection approaches. This way, infrastructures at different locations in the Internet can act as distributed “sensors” to better capture the global attack activity. Such sensors, in our case Internet Exchange Points, can collectively infer attack activity faster and potentially for even relatively small attacks. Thus, they can signal their peers about ongoing attacks. The contributions of this work can be summarized as follows:

- We establish a collaboration with 11 Internet Exchange Points around the globe to utilize them as vantage points and collect traffic. Over a period of six months, we detect and analyze more than 120k amplifications DDoS attacks. Our analysis shows that more than 50% of the attacks are visible in more than three locations, in many cases in more than 5 locations.

Table 1: Statistics about the 11 IXPs in our study between April 27 to October 5, 2020.

IXP Code	#Members	Peak Traffic (Gbps)	Region	#Sampled Flows (Billions)
CE1	900+	9,000+	Central Europe	1,077.5
CE2	200+	150+	Central Europe	9.9
CE3	200+	150+	Central Europe	3.2
CE4	200+	100+	Central Europe	3.6
NA1	200+	800+	North America	78.8
NA2	75+	150+	North America	16.7
SE1	175+	400+	South Europe	30.5
SE2	75+	100+	South Europe	12.2
SE3	40+	10+	South Europe	2.2
SE4	30+	100+	South Europe	17.9
SE5	20+	50+	South Europe	2.0

- We show that more than 80% of the attacks that send traffic via one of our vantage points are not mitigated because local detection mechanisms aren’t triggered (i) due to local attack traffic thresholds aren’t exceeded or (ii) due to the attack’s multi-protocol profile remaining unseen at a single location.
- We show the critical role that infrastructures, like IXPs, located in the core of the Internet, can play in detecting and mitigating DDoS attacks. We show that around 45% of the reflectors’ traffic is directly transferred from IXP members, and at least 30% of the attack targets are members of these IXPs.
- We develop a lightweight and easy to implement collaborative DDoS Information Exchange Point (DXP) that allows network platforms to report amplification DDoS reflectors or targets thereof to improve DDoS detection and mitigation.
- We estimate the potential benefit of collaborative detection and mitigation. The evaluation of our system shows that it is possible to detect and mitigate in some IXPs up to 90% more attack traffic locally when our collaborative DDoS Information Exchange Point is in operation.

2 DATASETS

For our study we leverage a distributed set of vantage points as well as network data and routing information. Our goal is to analyze all the available data to characterize recent amplification DDoS attacks and assess the potential benefits of collaborative DDoS detection and mitigation.

2.1 Vantage Points

We establish a collaboration with 11 Internet Exchange Points (IXPs) that operate at three regions around the world, namely Central Europe, South Europe, and North America. An IXP is a physical infrastructure with multiple layer-2 Ethernet switches installed in one or multiple peering facilities in one metropolitan area [11]. Network operators can be members of an IXP to exchange traffic with other members using the IXP’s switching fabric. The network members can either place their routers in the same physical location with IXP switches, or exchange traffic using remote peering [48]. Thus, the exchanged traffic can be local or remote. In some very large IXPs, up to 30% of the members are remote. Typically, an IXP has a mix of network members. Many of the large cloud and content providers are members of the IXPs, as well as enterprise networks and regional or national eyeball networks [2, 11]. The size of the IXPs varies in terms of the number of members, i.e., peering

networks, as well as the peak traffic. For an overview of the IXPs in our study, we refer to Table 1.

All the IXPs we collaborate with are already offering DDoS mitigation solutions to their members. They offer BGP blackholing [19] or advanced blackholing [20] to block traffic to specific destinations or transport ports by dropping the traffic at the IXP, as a free service. IXPs are excellent locations for mitigation, as they have a large spare capacity. Thus, they can absorb large attacks at the scale of Tbps. All the IXPs we collaborate with have at least one scrubbing center [31] as member.

2.2 Flow Data

We get access to flow data collected at each of the 11 IXPs. Due to high volume, all these IXPs use the Internet Protocol Flow Information Export (IPFIX) protocol [13] that aggregates information per flow without storing the payload of the packets. For maintaining scalability, they sample packets at the rate of 1:10k.

Passive measurements. Collection of the flow data took place from 27th of April, 2020 to 5th of October, 2020. The total traffic of sampled data is 1.175 Petabytes, that corresponds to approximately 11,750 Petabytes of exchanged traffic in total (all 11 locations).

Active measurements. We collected flow data for self-initiated controlled DDoS attacks. The measurements took place between February 11th and 20th, 2021. During the self-attacks, a total of 4.6 TB was transmitted, resulting in about 340K sampled flows.

2.3 Metadata

IXP member lists. During the time of flow data collection, we have access to the list of members at each of the IXPs. The lists are updated every day, as new members are added (or removed) daily.

IXP route server BGP data. Moreover, we collect routing data from the route server at each IXP during the flow collection period. A route server [55] is a free service offered by all the 11 IXPs to their members. The IXP members have the option to announce prefixes to all the other members of the IXP. This service is very popular with more than 60% of all member networks using the IXP route server with an open peering policy. The IXP network members can announce their prefixes with only one BGP session, instead of establishing one session for each peer network. The Route Server also offers the option to announce prefixes to only some, none or all peers. We have access to both the input and (filtered) output at the 11 IXP's route servers. We analyze the output routes of the route server, i.e., the best path selected for route propagation to the peers. This allows us to derive AS-distance information from the propagated BGP messages. For incoming traffic, however, the AS-distance correctness relies on the assumption of symmetric routing.

Internet routing registers. In addition, we also have access to public collector datasets. We used the Routing Assets Database (RADb) [52] to retrieve mappings from IP blocks to ASNs. This allows us to detect whether an IP address belongs to a peer of one of the IXPs. This further helps to calculate the distance between the IXP and the reflector or target respectively.

3 ETHICAL CONSIDERATIONS

To comply with measurement ethics, we carefully design our study and take a number of measures that we describe next.

Traffic captures. Our study is based on traffic data that the IXPs regularly captured for operational purposes and are in compliance with legal requirements in the respective countries of operation. All traffic traces are aggregated at flow-level and thus do not contain any payload. Additionally, the data is processed and analyzed in-situ at the premise of the IXPs.

Controlled self-attacks. Experience has shown that generating synthetic DDoS traffic in a real-world setup is hardly feasible. To obtain realistic traffic captures, we analyze traffic captures of self-attacks by following ethical guidelines and considerations that have been outlined for similar research studies [34]. Self-attacks run against a specially crafted autonomous system that belongs to the research infrastructure of one of the IXPs. The operator of the IXP executes the attack and several precautions are taken to limit potential negative effects of the attack. First, the IXP ensures that sufficient network bandwidth is available so that the likelihood of members being harmed by the targeted attack is minimized. Second, the IXP uses an experimental AS with no customer traffic and utilizes an unused /24 prefix that is allocated and announced only for the purpose of the experiment. While influence on external infrastructures (reflectors) cannot be completely avoided, the IXP captures the attack traffic to the infrastructure and continuously monitors the traffic sent by each reflector. The average traffic per reflector is typically between 500 kbps and 2 Mbps. The scope is limited by only purchasing the lowest possible low-volume attacks (\$15) and further ensures no attack lasts longer than 5 minutes and the peak traffic is no more than 7 Gbps following recommendations from previous studies [34]. While contracting a booter service is a sensitive matter, the setup originates from a collaboration with law enforcement to study booter services and for operational tests. These include gaining insights into DDoS attack traffic from booters for ensuring operational safety (structure, link dimensioning, etc.), which is relevant at a national level where IXPs are considered critical infrastructure. During the self-attacks experiment, the IXP operator did not receive any complaints.

4 ANATOMY OF DDOS ATTACKS

In this section, we analyze real-world DDoS attacks in detail to assess their visibility across our vantage points. We approach this in two different ways. First, we have an exemplary look into a recent large scale DDoS attack that is well reported by the targeted infrastructure provider. Second, we analyze a set of self-attacks towards our measurement network connected to the IXP infrastructures we collaborate. We run various advertisement scenarios for our attacked IP space to observe how the distribution of DDoS attack traffic is affected.

4.1 A Recent Tbps Reflection Attack

We investigate an exemplary DDoS attack that is one of the largest ever reported in terms of attack traffic volume. The attack took place on June 4th, 2020 and targeted the Content Delivery Network (CDN) Akamai, which is member of all 11 IXPs we have data for. The peak attack traffic was reported to be 1.44 Tbps [30]. The attack

Figure 1: Attack against a large CDN on June 04, 2020. Each row corresponds to a different protocol used for the attack as observed at each IXP.

is not only voluminous in terms of traffic, but also in terms of the number of packets. At the peak, around 385 million packets per second were generated. This makes the attack even more effective as it requires additional defense resources due to the processing of the very high number of packets, especially for traffic scrubbing centers. The attack is very sophisticated as nine different attack vectors are reported. These are multiple TCP and UDP specific attack vectors along with amplification attack vectors, i.e., SSDP, CLDAP and NTP. Our DDoS attack inference approach presented in § 5.1 successfully detected this event as an attack.

By analyzing the network flow data collected at the IXPs we confirm that the attack was visible at 3 of our vantage points. In Figure 1 we plot the reflector to target traffic volume for each IXP where the attack is visible. Note that these IXPs are located at different continents. The attack consists of 3 bursts and ends after a total duration of about 3 hours. The peak traffic observed at our vantage points totals at about 100 Gbps in terms of attack volume and at about 20 Mpps in terms of packets.

Five of the reported attack vectors are clearly visible in our data. The attack had an ON-OFF pattern, as also reported by Akamai [30]. Such patterns are common, as attackers try to avoid detection, and they also switch between attack sources and reflectors. Our data shows that the same target saw multiple smaller DDoS attacks one week earlier, using similar attack vectors. This can be considered a trial DDoS attack, a common pattern used to verify for example function of the reflectors. Furthermore, we observe the same target to be under attack multiple times with more than 1 Gbps of total attack volume across our whole period of observation, after the reported attack. All these attacks use the same aforementioned attack vectors. Thus, we conclude that DDoS attacks are indeed visible at different locations. However, the level of the attack traffic and amount of attack features visible at different locations may differ. As an example, the attack traffic at NA1 and CE2, is quite low compared to CE1. Furthermore, CE1 observes 5 attack vectors, but NA1 and CE2 only 3, respectively. Using DDoS detection practices in one location can lead to slower or no reaction, e.g., because the attack traffic volume is not large enough at a specific location, or the locally visible attack vectors remain under the radar of local defense mechanisms.

Figure 2: Attack traffic from self-attacks per location comparing two announcements scenarios (anycast and unicast).

4.2 Self-Attacks

To gain ground truth for our study, the IXP operator attacks its own measurement infrastructure (for details see § 3) that is connected to 10 of the 11 IXP locations. During the attacks, the observable attack traffic is measured at these different vantage points. By adjusting the advertisement of the measurement network, it is possible to control the direct visibility of the DDoS target at the IXP locations and steer the traffic in different scenarios towards the network. To generate traffic samples from real-world DDoS attacks, a DDoS for hire service, also known as booter service, is utilized. The amplification protocols that are included in the self-attack experiment comprise DNS, NTP, SNMP, CLDAP, SOAP, SSDP, ARD. The attack traffic observed at the measurement infrastructure ranged from 100-600 Mbps for the protocols ARD and SSDP, more than 1 Gbps for SOAP, SNMP and up to 7 Gbps for NTP, DNS, CLDAP. The IXP's measurement infrastructure is attacked with the available amplification vectors with two different advertisement scenarios: (1) advertise measurement network only via the largest IXP, i.e., CE1 (unicast), and (2) advertise our measurement network at all locations (anycast). The self-attack traffic level observed at each location for the specific advertisement scenario is reported in Figure 2. To our surprise, we observe DDoS traffic towards the target at several vantage points even when advertising from a different location. When advertising the measurement network globally via anycast, we can see a stronger attraction of DDoS traffic locally especially for the locations in South Europe. This highlights the geographical distribution of DDoS traffic and the great potential for global mitigation efforts, even for networks that don't advertise their IP space via anycast.

Takeaway: Amplification DDoS attacks are globally distributed across networks, hence visible from different vantage points worldwide. At some, however, only with very low traffic rates that might hinder their local detection. Even for networks which don't employ a distributed infrastructure (e.g., anycast) attack traffic can be visible traversing different locations towards the target. The results show the great potential to more effectively detect and mitigate DDoS attacks.

5 INFERENCE OF DDOS ATTACKS

We focus on DDoS reflection attacks, that are responsible for some of the largest attacks known to date [6]. Their popularity and sophistication has increased the last years [27]. More frequently there are reports highlighting not only the traffic volume of DDoS attack, but also increasingly high numbers of packets per second [27].

5.1 Detecting DDoS Attacks in Flow Traces

Figure 3: DDoS inference approach following [36].

Detection. To identify DDoS reflection attack traffic in the flow-level traces provided by the IXPs, we employ the approach proposed in [36] as shown in Figure 3. We consider an IPv4 address to be under attack, if its inbound traffic exceeds a threshold of $t = 1$ Gbps from more than $k = 10$ (reflector) IPs with the same source port (e.g., a known reflection protocol such as NTP). It is based on the assumption that it is unlikely for an Internet client to receive traffic from many sources with the same source port number (e.g., NTP) at a high traffic rate. We remark and will later show that typical DDoS attacks can generate much larger traffic volumes and involve reflectors in the number of 100s or even 1000s. Yet, prior work [36] has shown this filter to be capable to differentiate between attack and benign traffic, an observation that we confirm in our validation. For our study, we focus on UDP-based amplification attacks.

We remark that the described filter approach [36] was proposed to be applied at a single site only. To be applicable in our multi-site scenario we extend it as follows. First, we define a flow as a septuple: (source (MAC address, IPv4 address, transport port), destination (MAC address, IPv4 address, transport port), IXP code). This ensures, that the same traffic flow traversing multiple IXPs will be captured as individual flows in our data to enable the later analysis of attack traffic visibility at different sites. Second, we define two variants of the detection threshold t . In the first variant (local threshold), the detection is applied to traffic from a single IXP only. In the second variant (global threshold), we detect a DDoS attack if the traffic sum exceeds t over all IXPs. We evaluate these thresholds in § 6.2.

Filtering. To further avoid false-positive classifications, we filter the flow data for traffic having the source transport port set to the well-known port of popular DDoS reflection attack protocols. These (and the associated port number) are the following: Chargen (19), DNS (53), RPC (111), NTP (123), SNMP (161), CLDAP (389), OpenVPN (1194), SSDP (1900), ARMS (3283), WS-Discovery (3702), Device Discovery (10001), Memcached (11211). In addition, we take the packet size in account, as reported in [36]. By this, we populate a new, pre-filtered data set for attack detection purposes, consisting of 4 billion flows belonging to an average of 3TB of traffic exchanged data per day. In order to evaluate the accuracy of this filtering approach, we compare the average packet sizes per source port for benign traffic, attacks as defined by our filtering method and the recorded self-attacks as a ground truth. The results are shown in Figure 4: for any port we have data in the self-attacks, we observe a comparable packet size distribution, which deviates clearly from the benign traffic. In all other cases, we observe a clearly different packet size distribution between attack traffic using our filtering

Figure 4: Validation of filtering approach; we compare the packet size characteristics of benign, attack traffic (as filtered with the described approach), and self-attack traffic for different protocol source ports.

Figure 5: Daily number of observed DDoS amplification attacks and associated traffic volume during the period of our analysis.

method and benign traffic. Our filtering method either selects traffic at the upper end of an Ethernet frame (high amplification factor, e.g., port 11211 for Memcached) or a with very small deviation around a characteristic packet size (e.g., 500 bytes for NTP monlists), which is expected and indicates a correct selection.

Attack Statistics. In Figure 5 we plot the number of attacks and the associated attack traffic volume for the amplifications attacks detected with our methodology. We notice a great variance among the total of 120k detected attacks. For most of the days, there are at least 500 attacks detected, whereas for some days, this number is about 300% higher. However, we did not notice any particular pattern, e.g., day of the week that has more attacks than others. However, we noticed that some of the services (ports) receive more traffic than others. The top ports (attack traffic volume) are: 123 (33%), 389 (30.8%), 53 (27%), 11211 (6%), and 1194 (1.2%). The other ports receive less than 1% of the attack traffic. The aggregated attack traffic volume that is exchanged in the 11 IXPs varies from 1 Terabyte to 4 Terabytes per day.

5.2 Validation: Attack vs. Benign Traffic

Features. To characterize benign and attack network flows, we derive 1,106 features from the flow-level traces. For an exhaustive list, we refer to Appendix A. These features include basic statistics like the duration of an attack or the overall as well as the peak traffic volume (in total, per transport-level protocol and per site). In

Figure 6: PCA analysis: Contributions of different features to the rotation of the first 4 PCAs. This highlights that different IXP members add complementary information.

Figure 7: PCA analysis: Projection of the attack, the self-attack, and benign traffic sample to the first three PCAs. The plot shows 1k samples for the attack and benign traffic, and an ellipsoid to show where the mass of the points fall.

addition, we consider the correlation of the attacks across different transport ports and across sites. Here, we compute the pair-wise correlation matrices for all features across time and then count how often the correlation exceeds the thresholds of 0.7, 0.8, and 0.9. Moreover, we perform local and global attack detection using different thresholds and include the time until attack detection as well as the traffic volume before, during, and after detection as additional features. As the overall traffic volume at the IXPs differs significantly, we also normalize the traffic per IXP site and compute the same feature sets for the normalized traffic volume. We remark that these features are of purely descriptive nature for our analysis and point to Figure 6 for a subset of the available features.

Attack vs. benign traffic. We validate the DDoS inference approach described in § 5.1 by analyzing the features for benign and attack traffic (i.e., traffic that matches our inference approach in § 5.1). We consider traffic within our unfiltered data set to be benign, if the destination IP address did not show up once using our detection mechanism. To reduce the dimensionality of the feature space, we apply a Principle Component Analysis (PCA). A PCA decomposition can be used to project a high-dimensional space to a lower-dimensional space by relying on the initial principle components. In effect it converts a set of values of M possibly correlated variables into a set of K uncorrelated variables, the PCAs. In that regard, PCA is a clustering algorithm for high-dimensional data. We find that a significant number of our features are correlated

since the first 5 PCAs explain more than 25% of the variance and the first 50 more than 75% of the variance.

In Figure 7, we show the projection of our feature set to the first 3 PCA dimensions for both benign (cross) and attack (triangle) traffic. PC1, PC2, and PC3 are the three principal components. We also apply a k-means clustering to 8 clusters on the data and color the corresponding clusters. Note that Figure 7 is zoomed into the 0.01 to 0.99 quantiles for each dimension. We do so as there are a small number of outliers for each dimension (enclosed in the ellipsoid red envelop). We observe that benign and attack traffic can be visually clearly separated, which highlights that their flow-level characteristics in our feature space differ substantially. Moreover, this region only covers data points from three clusters: the red, green, and the blue one. The red and blue ones only contain attacks while the green one contains mainly benign traffic samples. Note, the k-means clustering is a very simple clustering mechanism and with a more sophisticated mechanism it should be easily possible to separate attack from benign traffic samples.

Overall, this analysis suggests that the applied filter successfully detects DDoS attack traffic. Additionally, we manually inspected a random subset of the attack traffic to further support this finding. Thus, we use the DDoS attacks matching our inference approach as data set to inspect the distributed nature of DDoS attacks and illustrate our mitigation approach in the remainder of this paper.

Relevant features per site. We next use the PCA-analysis to understand if the attack traffic features are homogeneous for the different IXPs. Therefore, we consider the contribution of the different features to the different PCAs. More precisely, we look at the rotation values. The rotation per feature and PCA captures which contribution the feature has to this specific PCA. In Figure 6 we show a bar plot of rotation values for the top features for the first four PCAs. This analysis shows, that the feature relevance differs per IXP. It is not only the volume that counts but also where the site is located, what type of member networks there are etc. This highlights that the different IXPs have complementary perspectives on the attack landscape. In the scope of our paper, this suggests that a cooperation of these IXPs in jointly detecting DDoS attacks by exchanging data is beneficial.

6 DETECTING AND MITIGATING THOUSANDS OF DDOS ATTACKS

To understand the challenges and opportunities of combining the views of multiple vantage points to detect and mitigate attacks, we perform a detailed analysis of the more than 120k attacks we inferred with our detection method described in the previous section.

6.1 Challenges in Detecting DDoS Attacks

The detection and, thereby, also the mitigation of DDoS attacks is subject to challenges.

Detection Lag. Recent industry reports show that DDoS attacks are typically of short duration, i.e., less than one hour. For example, Cloudflare [15] reports that 90% of attacks last less than one hour. We observe similar characteristics in our data set (not shown): most of the attacks are relatively short-lived. Indeed, around 70% of the attacks have a duration of 10 minutes or less. 95% of the attacks lasted less than 50 minutes. The short duration of the attack

traffic makes its detection challenging. A detection and mitigation approach that is too slow (e.g., by requiring longer sample periods for stable detection) will, thus, fail to detect a large bulk of the current DDoS attack landscape. By performing a collaborative DDoS detection proposed in this paper, we show later that we can reduce the time required for detection and thereby increase the number of detected DDoS attacks.

Multi-Protocol Attacks. Another challenge is that most of the attacks do not rely on single transport port. Prior work has observed a tendency of DDoS attacks to utilize multiple attack vectors (i.e., amplification protocols) [15]. Thus, simple port-based blocking rules may not suffice in blocking DDoS traffic. In Figure 8, we show the distribution of the number of amplification protocols (ports) used to perform DDoS attacks (sites and combination discussed later). Our main finding is that most attacks involve 3 or more amplification protocols. More than half of the attacks in our data set use more than one amplification protocol. This holds both for short- and long-lived attacks.

Global vs. Local Thresholding. To detect amplification DDoS attacks the typical approach is to use local thresholds. These thresholds are only applied to the local traffic. This can be misleading as only a fraction of the attack traffic (below the local threshold) is routed via one location, but on aggregate the attack traffic yields a large DDoS. To show that this is quite often the case, in Figure 9 we plot the number of attacks detected using different local thresholds. Red bars annotate number of attacks visible at each IXP.

6.2 Opportunities in Detecting and Mitigating DDoS Attacks

Visibility of Attacks at Multiple Sites. A key observation of this paper is that there exists visibility for the same DDoS attack at multiple sites. This is rooted in the fact that these DDoS attacks are executed by abusing a large set of reflectors distributed across many different networks. Thus, given inter-domain routing, the traffic paths from these reflectors to the attacked target can be expected to traverse many different networks. In Figure 8 we show the number of IXP sites at which the benign flows and DDoS attacks in our datasets are visible. In the same figure, we further show the distribution of the (IXP sites, amplification protocol/port) combination. For the benign data, we see most of the flows at a single site, using a single protocol. In contrast, 80% of the attacks are visible at more than one IXP, even if we further restrict this by amplification protocol. Given that the traffic volumes observed at each IXP site vary, attack detection at a single site alone is challenging. Yet, this result shows the opportunity in detecting and mitigating DDoS attacks: if IXPs unite to jointly detect DDoS attacks, the number of detectable attacks increases.

We notice that the majority of the attacks that are visible at the IXPs are often missed for both low and high local thresholds. Only, the very large IXP in our study, CE1, tends to only lose a small fraction, especially when the threshold is small. Our analysis shows that around 80% of the attacks are missed by a large majority of the IXPs, except the very large IXP. Indeed, the very large IXP's view contributes to the global view of the ongoing DDoS attacks both in terms of bytes per second as well as packets per seconds as shown in Figure 10. Thus, although for very large IXPs local

Figure 8: Diversity of ports and sites, and combination of both in attacks and benign data.

Figure 9: Number of attacks detected using local thresholds. Red bars annotate number of attacks visible at each IXP.

Figure 10: Visibility by IXPs on the attack traffic (bytes per second and packets per second).

thresholds can be effective, such thresholds may not be sufficient for typical IXPs. Our study also shows that if local thresholds are set proportionally to the size of an IXP, with reference point a very large IXP, the detection of amplification DDoS attacks can be improved significantly, see Figure 11. However, the false positive rate may increase as well.

Potential Attack Traffic Savings. To estimate the potential attack traffic savings when information about the ongoing attack is shared, in Figure 12 (top) we compare the traffic that could have been detected and blocked at each IXP with a local or a global threshold. The difference is striking especially for the smaller sites and for

Figure 11: Number of attacks detected using local thresholds normalized by the average traffic volume of each IXP. Red bars annotate number of attacks visible at each IXP.

Figure 12: Traffic that could be detected and blocked with global information

various values of local and global threshold. In some cases the missed amplification DDoS attack traffic is close to 100% even for very low thresholds (100 Mbps) as shown in Figure 12 (bottom). It is worth noticing that high (local and even global) thresholds, e.g., 10 Gbps, may have also a negative effect as many of the amplification DDoS attacks do not send traffic at this rate.

Improved Reaction Time. A side benefit of using global information is that the amplification DDoS attack detection time is significantly improved. In Figure 13 we show that more than 80% of the attacks are detectable within 1 minute when the global

Figure 13: Time (in minutes) needed until an attack is detected in at least one site for different thresholds.

threshold is 100 Mbps. Around 70% (resp. 90%) of the attacks are detectable within 1 minute (resp. 5 minutes) with a global threshold of 250 Mbps. This allows to detect even short-lived attacks, i.e., the majority of the attacks that last for less than 10 minutes. Mitigation mechanisms are also more effective as they are activated earlier.

6.3 The Role of IXPs

Spare Capacity. IXPs are located in the core of the Internet. IXPs offer DDoS mitigation services, e.g., blackholing, as a free service to their members. Also, among their members are traffic scrubbing centers. Moreover, IXPs are peering infrastructures with very high capacity. To exemplify, the total capacity of the 11 IXPs we collaborate with is 65 Tbps, while their aggregated peak traffic of is around 11 Tbps. Thus, IXPs have spare capacity for absorbing and dropping even very large amplification attacks, at the scale of Tbps.

Proximity to Reflectors and Targets. To better understand the role that IXPs can play in defending against amplification DDoS attacks, in Figure 14 we plot the fraction of attack traffic that originates from reflectors with the relevant distance to IXPs in our study. To estimate the distance (in AS hops) from the reflector (IP) to an IXP we use routing information at the time of the attack, see § 2.3. Hop 1 refers to reflectors hosted in IXP members. Hop 0 refers to reflectors whose distance we could not estimate with our data. Recall, that the AS of IXPs is not visible in routing tables, thus, estimating the AS distance between reflectors and IXP is a complex task. More than 45% of the attack traffic originates from IXP members. This means that by blocking the traffic at the IXP it is possible to drop attack traffic as close as possible to the reflector of the amplification DDoS. Another 30% of the attack traffic originates from networks that are two hops away from the IXPs, typically customers of the members of the IXP. Again, dropping this traffic will reduce significantly the attack traffic that is routed in the Internet as it stopped close to the source of the attack.

When we turn our attention to the targets of the attacks, see Figure 14, we also observe that a large fraction of the targets is relatively close to the IXPs. Around 30% of the amplification attack traffic targets IPs that are hosted in IXP members. This means that DDoS mitigation solutions can provide significant DDoS protection to IXP members. Moreover, by applying DDoS mitigation at the IXP, it is possible to reduce the AS-distance that attack traffic travels by

Figure 14: Fraction of attack traffic that originates from reflectors and target victim IPs with the relevant distance to the IXPs in our study. (Distance 1: corresponds to IXP members; Distance 0: we could not determine the distance.)

Figure 15: Attack traffic per reflector.

one hop (see Hop 2 in Figure 14) for 35% of the attack traffic and 2-4 hops (see Hop 3-5) for 30% of the attack traffic.

Consolidation of Reflector and Target IPs. Another important observation derived by our analysis is that a relatively small number of reflector IPs are responsible for a large fraction of the attack traffic. In Figure 15 (top) we plot the attack traffic per reflector during for the 120k attacks we studied. Although there were more than 1.93 million reflector IPs identified, the top 1000 of them are responsible for about 40% of the attack traffic. This means that by blocking attack traffic from a relatively small number of reflector IPs yields significant reduction of the attack traffic. For the benign data, we see a similar image, in terms of relative traffic volume. In Figure 15 (bottom) we plot the number of our vantage points (sites) that observe reflector IPs as well as their associated attack traffic. Reflectors that are responsible for 50% of the attack traffic are visible

Figure 16: Attack traffic from reflectors based on the distance from IXPs that receive attack traffic. (Distance 1: IXP members; Distance 0: we could not determine the distance.)

at a minimum of three of our vantage points. This information can be shared across network infrastructures in a joined mitigation effort. In contrast, the benign data shows that less than 1% of the traffic is seen at 5 or more sites, and over 90% of the traffic in only visible at a single site.

In Figure 16 we provide more insights on the consolidation of the reflectors with regards to the distance from the IXP. Again, a handful IPs are responsible for a large fraction of the attack traffic. For example, when we focus on the reflectors hosted in IXP members, some 200 IPs are responsible for more than 50% of the attack traffic that originates from direct members, which is around 22% of the total attack traffic.

7 COOPERATIVE DDoS DETECTION

The previous section demonstrated that a distributed sensing approach can lead to substantial benefits for detecting DDoS attacks. That is, attack detection by combining information from multiple vantage points obtained via a distributed sensing platform. However, the evaluation shown so far is based on the (unrealistic) assumption of perfect information for all participating parties. Such an assumption is only reasonable if, e.g., all vantage points belong to the same network operator. As soon as multiple organizations are involved, data privacy becomes an issue, as the exchanged information may include critical information that may be covered by regulation (e.g., GDPR [1]). Our solution proposal is the concept of a *DDoS Information Exchange Point (DXP)*.

7.1 DDoS Information Exchange Point (DXP)

Next, we describe the DXP concept. The DXP is a central hub for the exchange of DDoS information. The exchanged information can be used by all participating parties for the collaborative detection of DDoS attacks. It follows the idea of IXPs as exchange points for Internet traffic and offers the same economic incentives: the more parties exchange data at the exchange point, the higher the value of participation [11]. In the following, we further introduce the concept and, then, evaluate its benefits based on our data sets:

- *Organizational and technical structure:* the design and interaction of governance and technology forming the DXP concept: § 7.1.1
- *Incentives:* the benefits for all participating members even in the presence of information asymmetry between small and large organizations: § 7.1.2

Figure 17: DDoS Information Exchange Point (DXP) Concept

- *Privacy*: modes for exchanging critical and less critical information and its impact on detection performance: § 7.2

7.1.1 Organizational and technical structure. A naive solution for a DXP would be an open exchange model, where anyone can connect to and provide or consume information about ongoing DDoS attacks. The obvious shortcomings of this approach are (a) the lack of protection against someone that is on purpose undercutting the integrity of the information (e.g., via spam or misinformation) and (b) the lack of privacy protection.

We argue that these problems *cannot* be solved in a *purely technical* manner. While privacy preserving cryptographic approaches such as private set intersection exist (see, e.g., [51]), they involve secure-multiparty computation with significant algorithmic complexity. Furthermore, they assume that certain information can be shared, i.e., the information in the intersection set. Even this information may be subject to GDPR and, thus, its sharing is not a technical problem requiring a technical solution, but rather a legal one. In addition, approaches such as private set intersection do not prevent parties from polluting the system with wrong information. In this regard, decades of distributed/Peer-to-Peer systems research have shown that the problem of spam/misinformation can hardly be prevented in systems without identity management [41, 65]. We, thus, argue that realizing a DXP requires a non-technical solution.

Consequently, the required level of trust for the DXP concept is not feasible without the support of a governance body (see Fig. 17 top). One role model for a governance body are the Regional Internet Registries (RIR), such as RIPE or ARIN, as well as the Mutually Agreed Norms for Routing Security (MANRS) initiative [29] or the Dutch Anti DDoS Coalition [45]. Its most relevant tasks are identity management via memberships and community building, handling abuse cases, standardizing filter and threshold definitions, operation of the necessary technical infrastructure, and definition of a clear, rule-based access model on who may publish and consume what type of information. Misbehavior and spam can be sanctioned in the same manner as at IXPs: by contract termination and, thus, exclusion from the DXP.

For the technological base of the DXP, we envision a publish subscribe system similar to Apache Kafka. The governance body can create topics with certain types of information on the publish

subscribe system and grant read/write access to DXP members for different topics to realize trust schemes.

7.1.2 Incentives for participation. IXPs usually charge a flat price per port capacity, so there is no additional revenue to be generated from DDoS traffic peaks as for transit providers. On the contrary, IXPs have a motivation to block DDoS, as it causes major headaches for the connected customers and leads to increased customer service cost. The DXP can be a very cost effective solution for IXPs if it is implemented as a non-profit community effort (see above), as the operational infrastructure to maintain is minimal. In fact, many IXPs are already organized in non-profit organizations such as EURO-IX, MANRS, or RIRs.

Nevertheless, DDoS mitigation is a mature market with well-established players such as Akamai [4] and Cloudflare [14]. While these companies seem to be competition to the DXP concept, they are in fact not: (1) they target a different set of customers, i.e., mainly enterprises while the DXP targets IXPs and their connected ASes (mainly ISPs, content and cloud providers); (2) they provide a different service, i.e., mitigation up to the application layer including terminating TLS encrypted sessions to protect services, while the DXP is intended to protect network infrastructure from volumetric attacks; (3) their technical approach differs by redirecting all traffic through their infrastructures via BGP or DNS, which is infeasible for large IXPs with more than 10 Tbps peak bandwidth and would mean giving up business for small to medium sized IXPs, as the mitigation provider would take over the task of exchanging traffic. If the DXP concept is in competition with the industry at all, it competes with vendors of security network appliances like Netscout/Arbor DDoS solutions [47], which analyze traffic and can generate filtering recommendations. However, the DXP concept is rather an extension to local detection methods like network security appliances, as their output can be exchanged via the DXP for the benefit of all participating IXPs.

In general, the incentive for DXP participation is characterized by *network effects*, i.e., the value that participants can derive from participating directly depends on the number and types of connected networks and their size. As such the DXP provides the same economic incentives to IXPs for participation as IXPs do for their customers. We have shown that many DDoS attacks are enabled by abusing a large set of distributed reflectors which implies that the attack traffic is also highly distributed. Given the way that inter-domain routing works in the Internet, no exchange point is able to observe the entire attack traffic unless a victim is only connected to the Internet via a single uplink directly at an IXP. We have shown earlier, that the visibility of individual IXPs on a single attack is indeed limited and depends on their size. Since no IXP in our study has full visibility on their own, the collaborative detection of DDoS attacks is beneficial. That is, our evaluation (see § 7.3) shows that networks using data contributed by other DXP members are able to detect more DDoS attacks than they could on their own. As such they have a strong incentive for contributing data to the DXP to be able to access the data of other DXP members given a low cost of DXP participation, which can be realized by the proposed community-driven DXP concept. Nevertheless, if the hardware or other cost for some members is prohibitively high to fight attacks even with the shared information, then they may not

react to signals of attacks and it is up to the consortium to extend the membership of such members or not.

7.2 DXP Realization: Low Trust & High Trust

We next describe two realizations of the DXP model that we later evaluate. The two models differ by the sensitivity of the information that is shared by its participants.

7.2.1 Low Trust Environment. The low-trust exchange model is based on conservative assumptions regarding the exchanged information: only lists of reflector IPs that are locally observed in ongoing DDoS attacks and IXP's peak traffic values are shared with other participants. This DXP model is based on the assumption that reflector IP addresses (e.g., public NTP or DNS servers) are much less privacy sensitive than sharing the targets. Lists of servers that can be used as amplifiers can also be obtained by Internet-wide scans and are publicly published by research projects, e.g., Rapid7 [40] and Censys [21]. Thus, the respective IPs can be considered semi public knowledge. The value added by the DXP is the confirmation that this reflector IP is currently—has recently been—engaged in DDoS attacks and how many other DXP members observed this reflector IP in DDoS attacks. This information can then be used to detect and mitigate DDoS attack traffic by the DXP members, e.g., by blocking or rate limiting reflector systems.

7.2.2 High Trust Environment. The second DXP type realizes a trust-mediated exchange with optimistic assumption on the exchanged (sensitive) information. This realization shares any information exchanged in the low trust scenario and adds the following data: (1) the destination IP of attacked systems and (2) the attack volume received by victim systems per source transport port (which is roughly similar to the attack vector for reflection attacks). Since sharing victim IP addresses can be considered sensitive according to current privacy regulating frameworks (e.g., GDPR), the sharing of this data requires non-technical solutions implemented in legal frameworks. We, thus, assume that the DXP offers (paid) memberships. The membership enables sharing of data subject to NDAs. In addition, we add some computational cost to exfiltrating target IPs at scale by hashing them multiple times with a frequently rotating salt added to the hash. Target IPs are only revealed by the DXP if the total amount of attack traffic exceeds a configurable threshold. The DXP includes a neutral board to handle complaints or spam (i.e., the injection of wrong information), following classical Internet models such as registries or IXPs. Another realization of the high trust environment is feasible, when a single company operates multiple IXPs. In this case the legal requirements for information sharing are significantly smaller.

Notably, both DXP realizations are not mutually exclusive for implementation and can be modeled as two separate publish subscribe topics of the DXP infrastructure. Participation in both environments can be subject to criteria set by the DXP governance body/the community.

For the exchange of traffic to compute the global threshold, each IXP participant uses a hash to consistently anonymize IPs and reports the traffic of each anonymized IP to the DXP with secure communication. The DXP then estimates the global traffic per anonymized IP (source or destination). When the traffic

of an anonymized IP exceeds the global threshold, DXP sends an encrypted alert to each participant. The IXP participant then deanonymizes the IP and takes action.

7.3 DXP Evaluation

Approach. We realize both the low trust and the high trust DXP in a simulation environment based on the observed attack traffic. Additionally, we conduct an experiment with our self-attacks, where we use the low trust DXP to empirically verify how effective this approach can be used to mitigate DDoS attacks. Our attack samples are ideally suited to evaluate not just the DXP concept but also their incentive structure as the involved IXPs differ in size by multiple orders of magnitude just as one would expect in the wild. Thus, our evaluation incrementally adds IXPs to the DXP. Hereby, we order the IXPs according to the overall traffic volume (ranging from CE5—the smallest IXP—to CE1—the largest). Each evaluation starts with a DXP with only a single member, the smallest IXP. We then incrementally add IXPs and evaluate their DDoS detection performance. We end our simulation runs when all 11 IXPs from our data set joined the DXP (i.e., the complete data set is considered).

To simulate the *high* trust DXP, we share the transport port, the target IP addresses and the flow-level traffic volumes between the DXP participants. The detection is then performed by applying the filter described in § 5.1 to the combined attack-level information. That is, a participating IXP can detect an attack even if only a single low-volume attack traffic flow is locally observed.

To realize the *low* trust DXP, each DXP member continuously shares the IP addresses of reflectors that they observe in their locally detected DDoS attacks. This information is cached at all DXP members for 24 hours. The DXP members subscribe to this feed of confirmed reflector IP addresses received via the DXP. The received IP addresses are then used to run the detection approach (§ 5.1), whereby the locally observed traffic by confirmed reflector IPs is multiplied by a boosting factor. This boosting factor can be a constant or depend on the number of members that have observed the reported reflector participating in a DDoS attack. The local attack detection again relies on the approach outlined in § 5.1. This approach effectively boosts the locally observed attack traffic volume and, thus, leads to a faster detection.

7.3.1 Boosting Factor Validation. As the introduced boosting factor modifies the observed traffic volumes, we evaluate to what extent various boosting factors introduce false detections. We investigate boosting factors of 2, 4, 6, 8 and 10 and plot their effects in Appendix B. For our evaluation, we create a mix of 100k benign and attack flows which all are boosted by the aforementioned boosting factors. For the highest boosting factor, i.e., 10, find that 89 additional flows (0.104%) become detected as attacks. As these might be classified as false positives at the first glance, the packet size distribution of the flows harmonizes to what we see in the attack flows. So, it turns out that these were likely false negatives, uncaught by our conservative approach of attack detection.

Improvements in detectable DDoS attacks. We show the percentage of detectable DDoS attacks (relative to all observed attacks) in Figure 18 for both the low (b) and the high (c) trust environment. As baseline, we show the results when a local-only detection is used in (a). In each evaluation, we apply different detection thresholds

Figure 18: Relative: Detectable DDoS attacks for low and high trust DXP setting vs. local detection at a single IXP without DXP.

Figure 19: Absolute: Detectable DDoS attacks for low and high trust DXP vs. local detection at a single IXP without DXP.

Figure 20: Share of detected attack traffic for low and high trust DXP setting vs. local detection at a single IXP without DXP.

$t \in \{50, \dots, 10000\}$ Mbps (see § 5.1). In general, the higher the threshold, the more easily attack traffic can be differentiated from benign traffic. In the local-only detection scenario (Figure 18(a)), we observe that the detectable attacks differ by the IXP size: the larger the IXP in terms of overall traffic, the more attacks can be detected as it is easier to reach the required threshold. Irrespective of its size, all IXPs are able to locally detect some attacks.

In the low-trust DXP (Figure 18(b)), reflector IPs from these locally detected attacks are shared with other DXP members and they locally boost the traffic share of these reflector IPs by a factor of 10. We observe that this sharing of information enables the detection of a substantially larger fraction of attacks as the boosted traffic volume is more likely to reach the required threshold for detection. Yet, the low-trust DXP only boosts attacks that are locally observable. Thus, since the number of attacks visible at each IXP differs, the detection figures are not monotonously increasing with the size of the IXPs. Still, *all* IXPs benefit from participating in the DXP. Yet, the benefit is larger for smaller IXPs.

In Figure 18(c), we show the evaluation for the high trust scenario in which the destination vector of the attack traffic is shared among participants. Given that full information is available, we count an attack as detected if it can be detected based on the shared

information, irrespective of the attack traffic levels at the individual IXPs. Thus, in contrast to the low-trust environment, the number of detectable attacks monotonously increases with the number of participating IXPs. The reason being the attack detection is performed on the combined and complete attack data stream shared via the DXP. This leads to a larger number of detectable attacks and, thus, lower relative ratios than in the low-trust scenario. To compensate for this, we show the same evaluation with the absolute number of attacks for the sake of comparison in Figure 19. This comparison shows that the high-trust DXP indeed leads to a higher number of detectable attacks. Yet, only by a small margin, depending on the respective site. The low-trust DXP may, thus, be a good compromise between privacy aspects and detection performance.

Blocked attack traffic. In Figure 20 we show the share of attack traffic that can be dropped by both the low-trust and the high-trust DXP concept, relative to the overall amount of scaled attack traffic. We observe that the higher attack detection rates by both DXP types directly translate to substantial growth in attack traffic volume that can now be mitigated. Most notably, the low-trust DXP enables >70-80% of the local attack traffic to be mitigated for all DXP sizes. This represents a substantial improvement to a local-only detection without the presence of the DXP. Here, similar

detection figures are reached only for the larger IXPs at the right side of the x-axis and only for the low detection thresholds. For higher detection thresholds ($> 1\text{Gbps}$), the low-trust DXP provides substantial detection performance improvements over the local-only setting. A similar observation is made for the high-trust DXP. As in the previous case, the relative shares are lower, given the higher absolute number of detectable attacks. We further use the information from the low-trust DXP to mitigate our self-attacks. The low-trust DXP provides 130k reflectors from the day before our self-attacks. Within our self-attacks we observe a total of 23k reflectors, from which we knew 932 (5%) from the low-trust DXP. However, this enables us to mitigate 56% of the attack traffic and therefore provides a strong confirmation of our previous findings. Moreover, it shows that this approach can be very beneficial for victims of DDoS attacks, not just large scale network operators.

Takeaways. Both the low- and high-trust DXP models offer substantial improvements for the number of detectable DDoS attacks and mitigatable attack traffic. Even for DXPs with only a few members, participating in a DXP offers substantial improvements and, thus, there are clear incentives for IXPs to join DXPs from the perspective of detection benefits. In this regard, the low-trust DXP model provides good detection performance while keeping the shared data to a bare minimum, i.e., only semi-public reflector IPs and peak traffic values. Given its performance, it provides an attractive model for a practical realization of our concept. Whether the DXP also works from an economic (cost) perspective can generally hardly be evaluated due to differing and secret cost structures of IXPs. Likely, only a real instance of a DXP can show this.

8 RELATED WORK

The industry is currently investigating ways to exchange information among trusted parties to improve routing security towards a more resilient Internet. A global initiative backed by network operators, IXPs, content delivery networks, and cloud providers is MANRS [29]. “MANRS requires collaboration among participants and shared responsibility for the global Internet routing system” by sharing information for validation of network announcements and registries, contact information for emergency situations, and anti-spoofing filters. DOTS [46] introduced requirements for enabling coordinated response to DDoS attacks. Our proposed DDoS Information Exchange Point can be used by participating partners to collectively fight against amplification DDoS attacks.

In the past, systems have been proposed to exchange information among networks to fight against DDoS, e.g., by using blockchain [56] or by introducing accountability to incentivize network operators to isolate sources of attacks in their networks [60]. Other proposed solutions are tailored to a small set of ISPs that are interconnecting with each other and have a relationship of customer-provider or peer [59]. Proposed systems enabled victims of DDoS to request attack monitoring and filtering on demand, and to pay upstream and remote ISPs for the services rendered [54]. Community efforts developed collaborative approaches to detect and neutralize botnets that participate in attacks [17, 38] and build collaborative IP blacklists [23, 35, 42]. These may suffer from shortcomings as they are not well maintained and sufficiently updated [43, 53].

At the national level, an anti-DDoS coalition have been formed. For example, in the Netherlands a national DDoS clearing house [45] is operational for collecting and sharing fingerprints of attacks and suitable mitigation rules among national network providers. Fingerprint extraction is done by dissecting pcaps of attacks, which may be shared through a DDoS database. The approach involves a considerable share of manual work and, to the best of the authors knowledge, there aren't any hard numbers on its efficiency. However, the project takes care of governance requirements and the legal implications of sharing sensitive data. Our approach is not a competitor of this project, but rather quantifies the potential of a distributed DDoS detection mechanism while extending the perspective to international vantage points.

More recent research on DDoS mitigation focused on full-blown scrubbing of traffic with programmable networking hardware, e.g., FPGAs [67] or P4 enabled switches [66]. These approaches aim more at applying fine grained filtering rules at scale to large amounts of traffic while staying flexible in adding, removing, or specifying new filters during operations. These approaches do not tackle distributed sensing of DDoS attacks nor do they tackle sensing at all and, thus, our proposed solution can complement and improve this new generation of DDoS mitigation platforms. The use of programmable networks was also suggested in [26] to enable verifiable in-network filtering for DDoS defense towards making IXPs or other involved infrastructures accountable in case of misbehavior.

9 CONCLUSION

DDoS attacks were first observed twenty years ago, but they are still one of the most serious threats. Amplification DDoS attacks have been repeatedly reported as both frequent and devastating reaching 2Tbps of attack traffic in recent years. In this paper, we show that such amplification attacks are visible at multiple locations in the Internet. Unfortunately, the defense against such attacks is myopic and local today, and, thus, slow to react to attack and not effective, especially for short-lasting ones. We show that coordination in detecting and mitigating such attacks yields significant benefits, especially for smaller network infrastructures. In some cases, more than 80% more attacks and attack traffic can be detected and mitigated. We also show that network infrastructures in the core of the Internet, such as IXPs, are able to drop attack traffic close to the location of the reflector, thus, reducing the distance that attack traffic traverses only to be dropped later and create harm.

Our proposed DDoS Information Exchange Point is easy to realize and suitable for both low- and high-trust settings, where network providers exchange information to collaboratively fight DDoS attacks. Our visibility study, that considers more than 120k amplification attacks, shows that in both settings it is possible to neutralize most of the attacks and drop up to 100% of the attack traffic that is routed via the collaborating networks.

Acknowledgements

This work was partially funded by the German Federal Ministry of Education and Research (BMBF) grants 5G-INSEL 16KIS0691, AIDOS 16KIS0975K and 16KIS0976, and BIFOLD 01IS18025A and 01IS18037A, and by the European Research Council (ERC) Starting Grant ResolutioNet (ERC-StG-679158).

REFERENCES

- [1] 2016. Data protection in the EU, The General Data Protection Regulation (GDPR); Regulation (EU) 2016/679. <https://ec.europa.eu/info/law/law-topic/data-protection/>.
- [2] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. 2012. Anatomy of a Large European IXP. In *ACM SIGCOMM*.
- [3] Akamai. 2018. State of the Internet Security Report (Attack Spotlight: Memcached). <https://www.akamai.com/us/en/multimedia/documents/state-of-theinternet/soti-summer-2018-attack-spotlight.pdf>.
- [4] Akamai. 2021. *Akamai Security Solutions*. <https://www.akamai.com/solutions/security>
- [5] Akamai. 2021. Prolexic Technologies by Akamai. <https://www.akamai.com/us/en/products/security/prolexic-solutions.jsp>.
- [6] Amazon. 2020. *AWS Shield Threat Landscape Report*. https://aws-shield-tlr.s3.amazonaws.com/2020-Q1_AWS_Shield_TLR.pdf
- [7] M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou. 2017. Understanding the Mirai Botnet. In *USENIX Security*.
- [8] C. Labovitz. 2019. Internet Traffic 2009–2019. APRICOT 2019.
- [9] K. Carriello. 2017. Arm Yourself Against DDoS Attacks: Using BGP Flow Specification for Advanced Mitigation Architectures. <http://forum.ix.br/files/apresentacao/>.
- [10] O. Çetin, C. Gañán, L. Altena, T. Kasama, D. Inoue, K. Tamiya, Y. Tie, K. Yoshioka, and M. van Eeten. 2019. Cleaning Up the Internet of Evil Things: Real-World Evidence on ISP and Consumer Efforts to Remove Mirai. In *NDSS*.
- [11] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. 2013. There is More to IXPs than Meets the Eye. *ACM CCR* 45, 5 (2013).
- [12] CISCO. 2005. Remotely Triggered Black Hole Filtering - Destination Based and Source Based. Cisco White Paper, http://www.cisco.com/c/dam/en_us/about/security/intelligence/blackhole.pdf.
- [13] B. Claise, B. Trammell, and P. Aitken. 2013. RFC 7011: Specification of the IPFIX Protocol for the Exchange of Flow Information.
- [14] CloudFlare. 2021. *CloudFlare Comprehensive DDoS Protection*. <https://www.cloudflare.com/ddos/>
- [15] Cloudflare. 2021. *DDoS attack trends for 2021 Q1*. Cloudflare. <https://blog.cloudflare.com/ddos-attack-trends-for-2021-q1/>
- [16] B. Collier, D. R. Thomas, R. Clayton, and A. Hutchings. 201. Booting the booters: Evaluating the effects of police interventions. In *ACM IMC*.
- [17] E. Cooke, F. Jahanian, and D. McPherson. 2005. The Zombie Roundup: Understanding, Detecting, and Disrupting Botnets. In *USENIX Sruti*.
- [18] J. Czyz, M. Kallitsis, M. Gharaibeh, C. Papadopoulos, M. Bailey, and M. Karir. 2014. Taming the 800 Pound Gorilla: The Rise and Decline of NTP DDoS Attacks. In *ACM IMC*.
- [19] C. Dietzel, A. Feldmann, and T. King. 2016. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In *PAM*.
- [20] C. Dietzel, M. Wichtlhuber, G. Smaragdakis, and A. Feldmann. 2018. Stellar: Network Attack Mitigation using Advanced Blackholing. In *ACM CoNEXT*.
- [21] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman. 2015. A Search Engine Backed by Internet-Wide Scanning. In *CCS*.
- [22] A. Feldmann, O. Gasser, F. Lichtblau, E. Pujol, I. Poese, C. Dietzel, D. Wagner, M. Wichtlhuber, J. Tapiador, N. Vallina-Rodriguez, O. Hohlfeld, and G. Smaragdakis. 2021. A Year in Lockdown: How the Waves of COVID-19 Impact Internet Traffic. *Communications of the ACM* 64, 7 (July 2021), 101–108.
- [23] J. Freudiger, E. De Cristofaro, and A. Brito. 2015. Controlled Data Sharing for Collaborative Predictive Blacklisting. In *DIMVA*.
- [24] D. Gillman, Y. Lin, B. Maggs, and R. K. Sitaraman. 2015. Protecting Websites from Attack with Secure Delivery Networks. *IEEE Computer Magazine* 48, 4 (2015).
- [25] V. Giotsas, G. Smaragdakis, C. Dietzel, P. Richter, A. Feldmann, and A. Berger. 2017. Inferring BGP Blackholing Activity in the Internet. In *ACM IMC*.
- [26] D. Gong, M. Tran, S. Shinde, H. Jin, V. Sekar, P. Saxena, and M. S. Kang. 2019. Practical Verifiable In-network Filtering for DDoS Defense. In *IEEE ICDCS*.
- [27] Google. 2020. *Exponential growth in DDoS attack volumes*. www.cloud.google.com/blog/products/identity-security/identifying-and-protecting-against-the-largest-ddos-attacks
- [28] H. Griffioen, K. Oosthoek, P. van der Knaap, and C. Doerr. 2021. Scan, Test, Execute: Adversarial Tactics in Amplification DDoS Attacks. In *ACM CCS*.
- [29] MANRS initiative. 2021. Mutually Agreed Norms for Routing Security. <https://www.manrs.org/>.
- [30] L. Jakober. 2020. Akamai mitigates sophisticated 1.44 Tbps and 385 Mpps DDoS Attack. Akamai Blog, <https://blogs.akamai.com/2020/06/akamai-mitigates-sophisticated-144-tbps-and-385-mpps-ddos-attack.html>.
- [31] M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto, and A. Dainotti. 2017. Millions of targets under attack: a macroscopic characterization of the DoS ecosystem. In *ACM IMC*.
- [32] M. Jonker and A. Sperotto. 2017. Measuring Exposure in DDoS Protection Services. In *IEEE Network and Service Management*.
- [33] M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre, and A. Pras. 2016. Measuring the Adoption of DDoS Protection Services. In *ACM IMC*.
- [34] M. Karami, Y. Park, and D. McCoy. 2016. Stress Testing the Booters: Understanding and Undermining the Business of DDoS Services. In *WWW*.
- [35] S. Katti, B. Krishnamurthy, and D. Katabi. 2005. Collaborating against common enemies. In *ACM IMC*.
- [36] D. Kopp, C. Dietzel, and O. Hohlfeld. 2021. DDoS Never Dies? An IXP Perspective on DDoS Amplification Attacks. In *PAM*.
- [37] D. Kopp, J. Santanna, M. Wichtlhuber, O. Hohlfeld, I. Poese, and C. Dietzel. 2019. DDoS Hide & Seek: On the Effectiveness of a Booter Services Takedown. In *ACM IMC*.
- [38] M. Kühner, T. Hupperich, C. Rossow, and T. Holz. 2014. Exit from hell? Reducing the impact of amplification DDoS attacks. In *USENIX Security*.
- [39] C. Labovitz. 2021. Tracing Volumetric DDoS to its Booter / IPHM Origins. *NANOG* 82.
- [40] Rapid7 Labs. 2021. *Project Sonar*. <https://opendata.rapid7.com/>
- [41] F. G. Mármol and G. Pérez. 2009. Security Threats Scenarios in Trust and Reputation Models for Distributed Systems. *Elsevier Computer Security* 28, 7 (2009).
- [42] L. Melis, G. Danezis, and E. De Cristofaro. 2016. Efficient Private Statistics with Succinct Sketches. In *NDSS*.
- [43] L. Melis, A. Pyrgelis, and E. De. Cristofaro. 2018. On Collaborative Predictive Blacklisting. *ACM CCR* 48, 5 (2018).
- [44] C. Morales. 2018. NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us. <https://www.netscout.com/blog/asert/netscout-arbor-confirms-17-tbps-ddos-attack-terabit-attack-era>.
- [45] No more DDoS Coalition consortium. 2021. National Anti-DDoS-coalition. <https://www.nomoreddos.org/en/>.
- [46] A. Mortensen, T. Reddy, and R. Moskowitz. 2009. DDoS Open Threat Signaling (DOTS) Requirements. IETF RFC 8612.
- [47] NETSCOUT. 2021. *Arbor DDoS Protection*. <https://www.netscout.com/ddos-protection>
- [48] G. Nomikos, V. Kotronis, P. Sermpezis, P. Gigis, L. Manassakis, C. Dietzel, S. Konstantaras, X. Dimitropoulos, and V. Giotsas. 2018. O Peer, Where Art Thou? Uncovering Remote Peering Interconnections at IXPs. In *ACM IMC*.
- [49] E. Osterweil, A. Stavrou, and L. Zhang. 2020. 21 Years of Distributed Denial-of-Service: A Call to Action. *IEEE Computer Magazine* 53, 8 (2020), 94–99.
- [50] E. Osterweil, A. Stavrou, and L. Zhang. 2020. 21 Years of Distributed Denial-of-Service: Current State of Affairs. *IEEE Computer Magazine* 53, 7 (2020), 88–92.
- [51] B. Pinkas, T. Schneider, G. Segev, and M. Zohner. 2015. Phasing: Private Set Intersection Using Permutation-based Hashing. In *USENIX Security Symposium*.
- [52] RADb. 2021. RADb: The Internet Routing Registry. <https://www.radb.net>.
- [53] S. Ramanathan, A. Hossain, J. Mirkovic, M. Yu, and S. Afroz. 2020. Quantifying the Impact of Blacklisting in the Age of Address. In *ACM IMC*.
- [54] S. Ramanathan, J. Mirkovic, M. Yu, and Y. Zhang. 2018. SENSs Against Volumetric DDoS Attacks. In *ACSAC*.
- [55] P. Richter, G. Smaragdakis, A. Feldmann, N. Chatzis, J. Boettger, and W. Willinger. 2014. Peering at Peerings: On the Role of IXP Route Servers. In *ACM IMC*.
- [56] B. Rodrigues, T. Bocek, A. Lareida, D. Hausheer, S. Rafati, and B. Stiller. 2017. A Blockchain-Based Architecture for Collaborative DDoS Mitigation with Smart Contracts. In *AIMS*.
- [57] C. Rossow. 2014. Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In *NDSS*.
- [58] J. Ryburn. 2015. DDoS Mitigation Using BGP Flowspec. *NANOG* 63.
- [59] V. Sekar, N. Duffield, O. Spatscheck, K. van der Merwe, and H. Zhang. 2006. LADS: Large-scale Automated DDoS detection System. In *USENIX Security*.
- [60] D. R. Simon, S. Agarwal, and D. A. Maltz. 2007. AS-Based Accountability as a Cost-effective DDoS Defense. In *HotBots*.
- [61] Z. M. Smith, E. Lofri, and J. A. Lewis. 2020. The Hidden Costs of Cybercrime. McAfee, <https://www.mcafee.com/enterprise/en-us/assets/reports/tp-hidden-costs-of-cybercrime.pdf>.
- [62] K. Subramani, R. Perdisci, and M. Konte. 2020. IXmon: Detecting and Analyzing DRDoS Attacks at Internet Exchange Points. *CoRR abs/2006.12555*.
- [63] T. Vissers, T. Van Goethem, W. Joosen, and N. Nikiforakis. 2015. Maneuvering around clouds: Bypassing cloud-based security providers. In *ACM CCS*.
- [64] A. Welzel, C. Rossow, and H. Bos. 2014. On measuring the impact of DDoS botnets. In *EuroSec*.
- [65] M. Wichtlhuber, S. Bucker, R. Kluge, M. Mousavi, and D. Hausheer. 2016. Of Strategies and Structures: Motif-based Fingerprinting Analysis of Online Reputation Networks. In *IEEE LCN*.
- [66] M. Zhang, G. Li, S. Wang, C. Liu, A. Chen, H. Hu, G. Gu, Q. Li, M. Xu, and J. Wu. 2020. Poseidon: Mitigating volumetric ddos attacks with programmable switches. In *NDSS*.
- [67] Z. Zhao, H. Sadok, N. Atre, J.C. How, V. Sekar, and J. Sherry. 2020. Achieving 100Gbps Intrusion Prevention on a Single Server. In *USENIX OSDI*.

APPENDIX

A FEATURE SET

For analyzing our flow data, we define 1106 features. The following list shows how they can be divided into feature classes, see Tables 2 and 3. Some classes contain multiple features as they are parameterized by IXP and/or by port. Therefore, we also add how many features they contribute for clarity.

The sites are CE1, CE2, CE3, CE4, SE1, SE2, SE3, SE4, SE5, NA1 and NA2. When normalizing by size, we are multiplying the traffic by the relative average traffic volume to the biggest IXP. The source transport protocols are 19 (chargen), 53 (DNS), 111 (RPC), 123 (NTP), 161 (SNMP), 389 (CLDAP), 1194 (OpenVPN), 1900 (SSDP), 3283 (ARMS), 3702 (WS-Discovery), 10001 (Device Discovery), 11211 (Memcached). The thresholds are 100 Mbps, 250 Mbps, 500 Mbps 1 Gbps, 2.5 Gbps, 5 Gbps and 10 Gbps. Every time bin comprises one minute of traffic.

B BOOSTING FACTOR

In this section we show the impact of the boosting factor on (a) the relative number of detectable DDoS attacks in the low trust DXP setting, see Figure 21, (b) the number of detectable DDoS attacks in the low trust DXP setting, see Figure 22, as well as (c) the share of attack traffic that is detectable. We note, that the share of attack traffic as well as the number of detectable DDoS attacks increases as we increase the boosting factor. This underlines the benefit of the DXP.

At the same time one may fear that this increases the potential of false positives, i.e., classifying none attacks as attacks. However, our analysis of applying the boosting factor to benign data reveals that even though a small number of them—less than 100—are identified as attacks using, e.g., a boosting factor of 10, these are likely indeed attacks. We manually checked more than 10 and found that they are indeed likely lower volume DDoS attacks. This underlines that (a) the method for identifying attacks used in this paper is indeed very conservative and (b) that using the DXP with boosting is indeed useful for identifying more of the big DDoS attacks locally as well as identifying others less voluminous ones.

Table 2: List of features used for the PCA analysis.

Feature Class	Feature Count	Description
Sites	1	Number of sites involved in the attack
Ports	1	Number of source transport ports involved in the attack
SitesPorts	1	Sum of source transport ports seen at the sites, where the attack is visible
Dur	1	Total duration of the attack in minutes
DurAttack	1	Duration in minutes where the attack volume is greater than t (In our study: 1 Gbps)
TotalMbps	1	Volume of the attack in Mbps, summed across all sites and all source transport ports
TotalMbpsAttack	1	Volume of the attack in Mbps, summed across all sites and all source transport ports, while the volume is greater than t
TotalPeakMbps	1	Peak of the attack volume in Mbps, summed across all sites and all source transport ports
Peak Mbps	1	Peak of the attack volume in Mbps, single site, single source transport port
TotalMbpsCE1	1	Sum of the attack traffic across all source transport ports in Mbps, seen at site CE1
TotalMbpsAttackCE1	1	Sum of the attack volume across all source transport ports in Mbps, seen at site CE1 while exceeding t
TotalPeakMbpsCE1	1	Peak attack volume across all source transport ports, seen at site CE1, in Mbps
PeakMbpsCE1	1	Peak attack volume of a single source transport port, seen at site CE1, in Mbps
TotalMbpsNoCE1	1	Volume of the attack in Mbps, seen at all sites but CE1, all source transport ports
TotalMbpsAttackNoCE1	1	Volume of the attack in Mbps, seen at all sites but CE1, all source transport ports while exceeding t
TotalPeakMbpsNoCE1	1	Peak volume of the attack in Mbps, seen at all sites but CE1, across all source transport ports
PeakMbpsNoCE1	1	Peak volume of the attack in Mbps, seen at all sites but CE1, across a single transport port
Cor[Site Port]{0.7,0.8,0.9}	6	Counter for correlation of the attack between sites and source transport ports, respectively, being greater than .7, .8, .9, respectively per minute.
TotalMbps[IXP*]	11	Volume of the attack in Mbps, as seen at the 11 sites, all source transport ports, respectively
TotalMbps[PORT*]	12	Volume of the attack in Mbps, summed across all sites, for each of the 12 source transport ports in our study
PeakMbps[IXP*]	11	Peak volume of the attack in Mbps, as seen at the 11 sites, respectively, single source transport port
PeakMbps[PORT*]	12	Peak volume of the attack in Mbps, summed across all sites, for each of the 12 source transport ports in our study
TotalMpps	1	Sum of packets transmitted for the attack across all sites, all source transport protocols, in Mpps
TotalMppsAttack	1	Sum of packets transmitted for the attack across all, all source transport ports, sites while exceeding t , in Mpps
TotalPeakMpps	1	Peak of packets transmitted for the attack, summed across all sites, all source transport ports, in Mpps
PeakMpps	1	Peak of packets transmitted for the attack at any site, single transport port, in Mpps
TotalMpps[IXP*]	11	Sum of packets transmitted across all source transport ports, at the 11 sites, respectively
TotalMpps[PORT*]	12	Sum of packets transmitted at all sites, for each of the 12 source transport protocols in our study
TotalMbpsNorm	1	Volume of the attack, summed across all source transport ports and all sites, normalized by their size

Table 3: List of features used for the PCA analysis (cont.).

Feature Class	Feature Count	Description
TotalMbpsAttackNorm	1	Volume of the attack in Mbps, summed across all source transport ports, all sites, normalized by their size, while exceeding t
TotalPeakMbpsNorm	1	Peak of the attack volume in Mbps, summed across all source transport ports, all sites, normalized by their size
PeakMbpsNorm	1	Peak of the attack volume in Mbps, single source transport port, at a single site, normalized by their size
TotalMbpsNormNoCE1	1	Volume of the attack in Mbps, all source transport ports, seen at all sites but CE1, normalized by their size
TotalMbpsAttackNormNoCE1	1	Volume of the attack in Mbps, all source transport ports, seen at all sites but CE1, normalized by their size, while exceeding t
TotalPeakMbpsNormNoCE1	1	Peak volume of the attack, summed all source transport ports, seen at all sites but CE1, normalized by their size
PeakMbpsNormNoCE1	1	Peak volume of the attack, single source transport ports, seen at all sites but CE1, normalized by their size
TotalMbpsNorm[IXP*]	11	Volume of the attack in Mbps, all source transport ports, as seen at the 11 sites, normalized by their size, respectively
PeakMbpsNorm[IXP*]	11	Peak volume of the attack in Mbps, single source transport port, as seen at the 11 sites, normalized by their size, respectively
Allthresh-Before-[THRESHHOLD*]	7	Volume of traffic across all source ports that belong to an attack, greatest volume of a single site, before the respective threshold was exceeded
Allthresh-Detect-[THRESHHOLD*]	7	Volume of traffic across all source ports that belong to an attack, greatest volume of a single site, while the respective threshold is exceeded
Allthresh-After-[THRESHHOLD*]	7	Volume of traffic across all single source transport ports that belong to an attack, greatest volume of a single site, after the respective threshold is no longer exceeded
Allthresh-Time-[THRESHHOLD*]	7	Amount of time bins for which the attack volume across all source transport ports, greatest of a single site, exceeded the respective threshold
Allthreshnorm-Before-[THRESHHOLD*]	7	Volume of traffic across all source ports that belong to an attack, greatest of a single site, normalized by its size, before the respective threshold was exceeded
Allthreshnorm-Detect-[THRESHHOLD*]	7	Volume of traffic across all source ports that belong to an attack, greatest of a single site, normalized by its size, while the respective threshold is exceeded
Allthreshnorm-After-[THRESHHOLD*]	7	Volume of traffic across all source transport ports that belong to an attack, greatest of a single site, normalized by its size, after the respective threshold is no longer exceeded
Allthreshnorm-Time-[THRESHHOLD*]	7	Amount of time bins for which the attack volume across all source transport ports, greatest of a single site, normalized by its size, exceeded the respective threshold
SiteThresh-[IXP*]-Before-[THRESHHOLD*]	77	Volume of the attack, for every site respectively, single source transport port, before exceeding the respective threshold
SiteThresh-[IXP*]-After-[THRESHHOLD*]	77	Volume of the attack, for every site respectively, single source transport port, after the respective threshold is no longer exceeded
SiteThresh-[IXP*]-Detect-[THRESHHOLD*]	77	Volume of the attack, for every site respectively, single source transport port, while exceeding the respective threshold
SiteThresh-[IXP*]-Time-[THRESHHOLD*]	77	Amount of time bins, for every site respectively, for every threshold, single source transport port, before exceeding the respective threshold
GlobalThresh-[IXP*]-Before-[THRESHHOLD*]	77	Volume of the attack, adding all site's volume to every site respectively, all source transport ports, before exceeding the respective threshold
GlobalThresh-[IXP*]-After-[THRESHHOLD*]	77	Volume of the attack, adding all site's volume to every site respectively, all source transport ports, after the respective threshold is no longer exceeded
GlobalThresh-[IXP*]-Detect-[THRESHHOLD*]	77	Volume of the attack, adding all site's volume to every site respectively, all source transport ports, while exceeding the respective threshold
GlobalThresh-[IXP*]-Time-[THRESHHOLD*]	77	Amount of time bins, when adding all site's volume to the respective site, for every threshold, all source transport ports, while exceeding the respective threshold
SiteThreshNorm-[IXP*]-Before-[THRESHHOLD*]	77	Volume of the attack, for every site, normalized by its size, single source transport port, before exceeding the respective threshold
SiteThreshNorm-[IXP*]-After-[THRESHHOLD*]	77	Volume of the attack, for every site respectively, normalized by its size, single source transport port, after the respective threshold is no longer exceeded
SiteThreshNorm-[IXP*]-Detect-[THRESHHOLD*]	77	Volume of the attack, for every site respectively, normalized by its size, single source transport port, while exceeding the respective threshold
SiteThreshNorm-[IXP*]-Time-[THRESHHOLD*]	77	Amount of time bins, for every site respectively, normalized by its size, for every threshold, single source transport port, before exceeding the respective threshold
Total	1106

(a) Local detection only

(b) Low trust DXP mode: Boosting factor 2

(c) Low trust DXP mode: Boosting factor 4

(d) Low trust DXP mode: Boosting factor 6

(e) Low trust DXP mode: Boosting factor 8

(f) Low trust DXP mode: Boosting factor 10

Figure 21: Relative: Sensitivity of the detectable DDoS attacks in the low trust DXP setting for different boosting factors.

(a) Local detection only

(b) Low trust DXP mode: Boosting factor 2

(c) Low trust DXP mode: Boosting factor 4

(d) Low trust DXP mode: Boosting factor 6

(e) Low trust DXP mode: Boosting factor 8

(f) High trust DXP mode: Boosting factor 10

Figure 22: Absolute: Sensitivity of the detectable DDoS attacks in the low trust DXP setting for different boosting factors.

Figure 23: Sensitivity of the share of the attack traffic detected in the low trust DXP setting for different boosting factors.