This work has been accepted for publication in PAM 2026. The final published version will be available via Springer.

Detecting and Characterizing DDoS Scrubbing from Global BGP Routing: Insights from Five Leading Scrubbers

Shyam Krishna Khadka[1 [0000−0003−3660−6226]], Suzan Bayhan[1 [0000−0001−6662−704X]], Ralph Holz[2,1 [0000−0001−9614−2377]], and Cristian Hesselman[3,1 [0000−0002−7052−9300]]

[1] University of Twente, Enschede, The Netherlands
[2] University of Münster, Münster, Germany
[3] SIDN Labs, Arnhem, The Netherlands
{s.k.khadka, s.bayhan, r.holz, c.e.w.hesselman}@utwente.nl

Abstract. Many scrubbers use the Border Gateway Protocol (BGP) to route Distributed Denial of Service (DDoS) traffic to their infrastructure, allowing them to drop the DDoS traffic and forward legitimate traffic to the Autonomous Systems (ASes) the scrubber protects. Despite their importance, the prevalence and operational behaviors of BGP-based DDoS scrubbing services remain poorly understood, such as the extent to which protected ASes always have a scrubber on their path or activate a scrubber on-demand when an attack occurs. We bridge this gap by detecting scrubbing activations and deactivations in public BGP data, where they manifest themselves as a scrubber dynamically appearing as the first upstream of an origin AS or as an origin AS for a particular prefix. We use 30 days of BGP data from the RIS route collectors, focusing on the global top five scrubbing providers, such as Cloudflare and Akamai. We also characterize their behavior, including protection modes, on-demand mitigation strategies, and RPKI/IRR practices. We find that prefixes that always use a scrubber are dominant compared to those that activate a scrubber on-demand. We also observe that 48% of the prefixes that scrubbers temporarily originate during an attack are not covered by valid RPKI ROAs (12.5% Invalid and 35.5% Notfound), which highlights a potential operational gap in current scrubbing practices regarding routing security. These insights are conservative because we only consider public BGP data and AS path changes that are most likely to be scrubbing events (e.g., those observed by two or more route collector peers). We believe our work is useful for security researchers and policymakers, for instance, to better understand DDoS protection levels of ASes in a particular country or region.

Keywords: DDoS · Scrubbing · BGP.

1 Introduction

With the rise of Distributed Denial of Service (DDoS) attacks, many commercial DDoS "scrubbing" providers have come into existence. They often use the Border

Gateway Protocol (BGP) to redirect DDoS traffic to their infrastructure and then drop the DDoS traffic and forward legitimate traffic to the Autonomous Systems (ASes) the scrubber protects. We refer to such ASes as "protected ASes" and to their prefixes as "protected prefixes".

Scrubbers offer two modes of protection: always-on and on-demand [1, 9, 53, 25, 47]. The first means that a protected AS always routes all of its inbound traffic through a scrubber, and the protected AS does not have to make any changes in BGP to activate/deactivate the scrubber. In the case of on-demand protection, the protected AS dynamically activates and deactivates the scrubbing service by adding it to the path during a DDoS attack and removing it afterwards. As a result, the scrubber only occasionally appears on the path to the protected AS.

Different ASes might use different types of protection depending on their requirements, such as costs, operational complexity, and criticality of their services. We spoke with several scrubbers, and one of the top ones indicated that always-on is the default model. In contrast, other scrubbers have on-demand as their default setup (e.g., Nawas [41]), where the scrubber dynamically appears and disappears as an upstream in AS paths and the protected AS originates its own prefixes. Others report on the use of on-demand scrubbers that originate a protected prefix rather than the protected ASes being the origin [54, 7, 29].

Despite the presence of a large number of BGP-based DDoS scrubbers globally, a significant research gap remains because we have no insight into the prevalence of always-on and on-demand protection on the Internet. Khadka et al. [31] study global adoption of BGP-based scrubbing for the top five providers using public BGP data, but explicitly limit themselves to cases where the protected AS remains the origin and do not distinguish on-demand from always-on. Also, we know little about the operational characteristics of on-demand scrubbers, for instance, in terms of frequency and duration of their activation and how they handle routing security objects, such as in the Resource Public Key Infrastructure (RPKI) and the Internet Routing Registry (IRR). For example, a scrubber's RPKI practices are important when it originates the BGP announcements for protected prefixes. This is because protected AS must authorize the scrubber's AS Number (ASN) in its Route Origin Authorizations (ROA) [5] or otherwise ROA-validating ASes [24] might reject the route toward the scrubber, undermining DDoS mitigation.

Obtaining these insights is useful for multiple stakeholders. For example, the Mutually Agreed Norms for Routing Security (MANRS)+ working group could use them for their "DDoS Attack Prevention" metric. This metric reflects the DDoS protection practices of an AS [38], and is needed to implement MANRS+' stricter routing security compliance and auditing controls. Also, researchers can use insights into BGP-based DDoS scrubbers to gain a better understanding of the operation of the scrubber, for instance, to better distinguish BGP anomalies from DDoS scrubbing events. Additionally, insights into on-demand DDoS scrubbing can improve BGP monitoring platforms such as GRIP [22] and Radar [44], which network operators use to detect and prevent routing incidents like leaks or hijacks. For example, if the operators can more reliably distinguish legiti-

mate scrubbing events from actual hijacks, they can reduce the number of false positives in their results.

We answer the following Research Questions (RQs) to obtain these insights:

- **RQ1:** Which type of scrubbing is more dominant on the Internet: always-on or on-demand?
- **RQ2:** What are the characteristics of on-demand scrubbers? For example, to what extent do they originate a protected prefix on behalf of the ASes they protect, and to what extent do they propagate BGP announcements for protected prefixes as an upstream of a protected AS?
- **RQ3:** What are the RPKI and IRR management practices of scrubbers that originate the prefixes of their protected ASes?
- **RQ4:** What hints can we get about DDoS attacks based on the BGP behavior of on-demand scrubbers? For example, how long do DDoS attacks last, and how often do they take place?

To address these research questions, we design a method to detect scrubbing activations and deactivations in public BGP data, for instance, to capture a scrubber dynamically appearing as the first upstream (hereafter referred to simply as *upstream*) of an origin AS or as a scrubber becoming an origin AS for a particular prefix. We use the BGP data from the RIS project, which has 21 active route collectors gathering BGP data from about 1,300 BGP peers. We analyze 30 days of BGP routing data, combining daily Routing Information Bases (RIBs) snapshots with detailed analysis of BGP updates of 5-minute granularity. We focus on the global top five DDoS scrubbing providers: Cloudflare, Akamai Prolexic, Vercara (formerly Neustar), Imperva, and Radware.

Our detection method is intentionally conservative: it identifies only those events in public BGP data that are most likely attributable to scrubbing activations and deactivations. For example, we require that two or more route collector peers observe an AS path change that signals the activation of a scrubber, such as when the scrubber starts to originate the protected prefix rather than the protected AS. This approach aims to prioritize accuracy and minimize false positives, increasing confidence that the observed cases reflect actual DDoS scrubbing (de)activations.

Our contributions and main findings are as follows:

- We develop and implement a methodology that detects scrubbing by using BGP dumps and BGP updates, thereby providing insights into DDoS scrubbing and its operational practices. Our study focuses on the global top five scrubbers on the Internet, but it can be used for any scrubber as long as the ASN it uses for scrubbing is known.
- We observe that the number of prefixes using always-on protection is higher than those using on-demand protection, with 11,408 vs. 5,649 prefixes, corresponding to 0.8% and 0.4% of all routed prefixes as of 30 May 2025. We note that the actual number of on-demand protected prefixes may be higher, since some prefixes might not have been activated during our study period.

- We map out the different protection modes of on-demand scrubbers. We find that there are about 10 times more cases where a scrubber appears as an upstream provider of a protected AS than cases where a scrubber originates the prefix of a protected AS, with 1,070 and 104 prefixes observed, respectively. We also identify 4,475 prefixes that are likely to change their origin to scrubbers during DDoS protection, based on our ROA analysis of scrubber ASNs.
- We find 48% of prefixes that the scrubbers temporarily originate for DDoS protection are either RPKI Invalid (12.5%) or do not have records (35.5%) in the Route Origin Authorization (ROA) objects, which highlights a potential operational gap in current scrubbing practices regarding routing security.

The remainder of this paper is structured as follows. Section 2 provides information about always-on and on-demand scrubbing. We discuss related work in Section 3 and introduce our methodology in Section 4. We present our results scrubbing (de)activation dynamics and their characteristics in Section 5, and discuss our results in Section 6. We end with conclusions and future work in Section 7.

2 Background on DDoS Scrubbers

A DDoS scrubber connects to its protected AS through mechanisms such as GRE tunnels [48], dedicated links [1], or peering arrangements [8], which the protected AS uses to advertise its routes to the scrubber. DDoS scrubbers provide two modes of protection: always-on and on-demand [1, 9, 47, 53, 25]. We divide the latter into two subtypes: on-demand scrubbing for a prefix activated by changing the origin of the prefix to the scrubber or by changing the upstream of the origin AS of the prefix in an AS path [9, 46, 26]. In both scrubbing modes, only the inbound traffic comes through the scrubber, while the outbound traffic from the protected AS continues to exit normally through its ISP toward the Internet [15]. The activation and deactivation of on-demand scrubbing can be manual, where a protected AS modifies BGP configurations directly [57]; fully automated, as in Cloudflare's Magic Transit [11]; or partially automated, triggered by flow-monitoring alerts and executed via tools such as API calls to scrubbing providers [2, 56].

2.1 Always-on scrubbing

In always-on scrubbing, the protected AS always routes its traffic through a scrubbing service, such that the scrubber's AS always appears on the path toward the protected AS. In this case, a scrubber AS appears as the upstream of the protected AS in AS paths.

Consider a real-world example from our analysis, involving an AS path with ASNs *[513 25091 25091 13335 24864]* for prefix *2.58.145.0/24*, with *AS24864* the originating AS, and *AS13335* a known scrubber (Cloudflare). If *AS13335* always appears as an upstream for the prefix *2.58.145.0/24*, we say the prefix uses always-on mode of scrubbing and is a protected prefix.

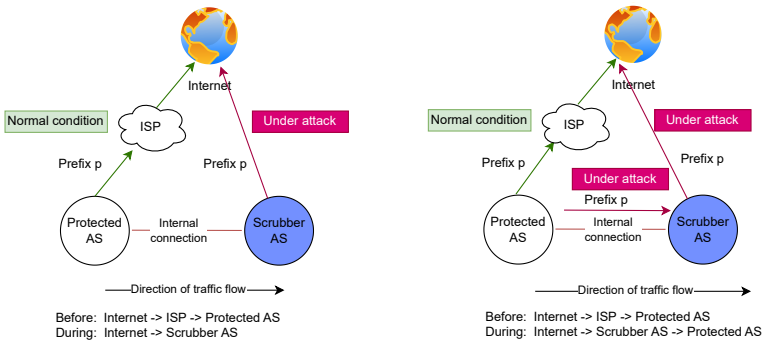

(a) Origin-change (b) Upstream-change

Fig. 1: Two ways of on-demand scrubbing we distinguish: a) Origin-change, where the scrubber originates the protected prefix, and b) Upstream-change, where a scrubber comes as an upstream of the protected AS.

2.2 On-demand scrubbing through origin changes

In on-demand scrubbing through origin changes, a scrubber AS originates the protected AS' prefixes on behalf of the protected AS during a DDoS attack (as shown by the red arrow in Figure 1a). Because this action involves a temporary change of the origin of the protected prefix, we call this mode "origin-change".

In scrubbing through origin changes, the scrubber ASN appears as the origin of a protected prefix in the BGP AS path during a DDoS attack. As a result, all Internet traffic first goes to the scrubber, which delivers clean traffic to the protected AS using an internal connection. A few research papers discuss this method [54, 7], but they lack detailed empirical evidence of inferring such cases.

As an example, consider AS path with ASNs *[3492 52025 32097 12200 33070]* for prefix *148.62.86.0/24*, where *AS33070* is the originating AS. In our data, we observed scrubber *AS19905* (Vercara) starting to originate that prefix, presumably to absorb a DDoS attack. As a result, the AS path changes to *835 3257 19905* for prefix *148.62.86.0/24*. After around 35 minutes, Vercara stops originating the prefix, and the origin reverts to *AS33070* (or to another ASN different from the scrubber that is authorized to originate), presumably when the DDoS attack is over.

For routing security, the prefix owner must authorize the scrubber's ASN in its ROAs [5] to prevent routing anomalies and allow scrubbing to operate safely [5]. ROAs are cryptographically signed objects in RPKI specifying which AS may originate which prefixes and up to what length. Without such ROAs, ROA-validating ASes may reject the scrubber's announcements, undermining mitigation. A prefix is RPKI-valid if a ROA exists that authorizes the scrubber's ASN and the announced length does not exceed the ROA's maximum; it is

Invalid if the ROA specifies a different ASN or a shorter length. A Notfound state indicates no covering ROA.

Similarly, ASes using IRR-based filtering must ensure their route objects are updated to include the scrubber's ASN. A route object specifies the IP prefix and originating ASN that an AS intends to announce in BGP. Since many network operators have not yet deployed RPKI-based filtering, IRR-based filtering remains the most widely used method, even among networks participating in the MANRS routing security initiative [37, 19].

2.3 On-demand scrubbing through upstream changes

In on-demand scrubbing through upstream changes, the protected AS originates its BGP prefix and advertises the prefix to a scrubber for mitigation during a DDoS attack, as shown in Figure 1b. In this way, the scrubber appears as an upstream of its protected AS instead of its Internet Service Provider (ISP). We therefore use the term "upstream-change" in this paper to refer to this mode of scrubbing.

As an example of upstream-change-based scrubbing, consider the AS path with ASNs *[2914 9744 45753]* for the originated prefix *182.16.18.0/24*, which we found in our data. First, the originating *AS45753* changes its upstream to scrubber *AS19905* (Vercara), presumably because of a DDoS attack. As a result, we see a new AS path for the prefix with ASNs *[2652 1299 2914 19905 45753]*. After around 47 minutes (presumably when the attack is over), the origin *AS45753* reverts its upstream to *AS9744*. In this example, we found that the protected AS used the same upstream before and after the attack, but they might differ if the AS has multiple upstream providers.

A protected AS can implement upstream-changes in three ways [46, 26, 21], which are essentially traffic engineering techniques [17]: i) the protected AS withdraws its announcement to its ISP and advertises it toward the scrubber, ii) the protected AS prepends its ASN on the AS path toward its ISP, causing the traffic to flow through the scrubber during the attack because of the shorter path toward the scrubber, and iii) the protected AS announces a more specific prefix (e.g. /24) toward the scrubber and a less specific prefix (e.g. /22) toward its regular ISP, thus deaggregating the a larger prefix into smaller ones to control routing and traffic flows. As a result, the routing will choose the route towards the scrubber.

3 Related Work on Scrubbing

We are unaware of prior work that provides insight into the prevalence of different scrubbing services, on-demand scrubbing services, and the operational practices of on-demand scrubbers based on public BGP and RPKI data. In general, the aspects of DDoS scrubbing that use BGP are relatively underexplored.

Khadka et al. [31] investigate the model in which the scrubber appears as an upstream provider of a protected AS. However, their study does not consider

on-demand protection or differentiate between on-demand and always-on modes, as it samples data only one day per month. While the DDoS scrubbing providers they examine overlap with ours, the research questions and methodology differ in important ways. In contrast, we sample BGP data every five minutes, allowing us to distinguish on-demand from always-on protection, detect cases where a scrubber appears as the origin of a protected prefix, and assess the RPKI and IRR status of these prefixes.

Tung et al. [55] propose a method for distinguishing DDoS attacks from other BGP anomalies and for analyzing scrubber behavior during such attacks. However, their study is limited to a single attack affecting three prefixes of one AS, and only identifies ASes that had previously experienced attacks. In contrast, our method detects DDoS scrubbing activity, providing insights in different modes of scrubbing regardless of prior attacks.

Jonker et al. [29] state that a DPS (DDoS Protection Service) provider announces an IP subnet of its protected AS, but they primarily focus on diverting traffic to the scrubber using the DNS rather than providing empirical evidence for BGP-based scrubber behavior. Testart et. al [54] state that DDoS scrubbers originate the prefixes of protected ASes during an attack, but their focus is on identifying and characterizing ASes that repeatedly and intentionally hijack IP prefixes.

Chung et. al [7] analyze the RPKI-based ROA deployment history. The paper looks at one snapshot of Routeviews data and finds 15 prefixes that are announced by DDoS protection ASes but belong to other ASes. Authors classify such announcements as "wrong" BGP announcements. As their primary focus is not on exploring RPKI practices of scrubbers, the study does not analyze the BGP updates that we perform to identify scrubbing cases.

Livadariu et al. [32] analyze the use of RPKI in conjunction with Remotely Triggered Black Hole (RTBH) filtering at Internet Exchange Points (IXPs), which is related but different from DDoS-handling using scrubbers. In addition, their study does not examine the RPKI practices of scrubbing providers. Chung et al. [7] study the deployment and coverage of RPKI longitudinally, where they examine a single routing snapshot, looking at 3 DDoS mitigation providers: Verisign, Neustar (now Vercara), and Level 3.

4 Methodology to Detect and Characterize Scrubbing

The goal of our methodology is to detect always-on and on-demand scrubbing (de)activations and to map out the properties of on demand-scrubbers, using public BGP data, such as BGP updates from RIS and RPKI data. We start the design of our methodology by identifying the five leading scrubbers and their ASNs, which we selected for our study. Next, we describe our approach for detecting on-demand scrubbing.

We classify on-demand scrubbing into "same-day" and "cross-day" events. Same-day scrubbing occurs when activation and deactivation signals for a prefix happen within a single day (24 hours), while cross-day scrubbing spans two con-

Table 1: ASNs of scrubbers and the references to verify them.

Scrubbers	Cloudflare	Akamai	Vercara	Imperva	Radware
ASNs	13335	32787	19905	19551	198949
References	[9, 42, 14]	[1, 10, 21, 39]	[53, 30]	[58]	[31]

secutive days (48 hours). Both categories include origin-change and upstream-change (de)activations (Sections 2.2 and 2.3). We make this distinction for methodological simplicity: same-day events are detected directly from BGP updates within a 24-hour window, whereas cross-day events require first comparing RIBs across consecutive days to identify prefixes that disappear the next day. This approach reduces the sample size for subsequent BGP update analysis by focusing only on these disappearing prefixes. We elaborate on the motivation for this approach further below. To facilitate reproducibility, we make our source code and analysis scripts publicly available[4].

4.1 Identifying scrubbers and their AS numbers

We focus on five leading scrubbers: Cloudflare, Akamai Prolexic, Vercara, Imperva, and Radware. We use them because they are among the global top five in terms of bandwidth [18] capabilities to mitigate DDoS attacks, and are also recognized as leading scrubbers in the 2021 Forrester wave market analysis [23]. Prior research has similarly considered these scrubbers [31, 29, 28].

Following the approach of Khadka et al. [31], we determine the ASNs that the five scrubbers use for scrubbing by inspecting multiple sources: documentation of the scrubbers; checking AS path changes involving Akamai, Cloudflare, and Vercare as scrubbers against known past DDoS incidents; blogs; and contacting operators from Akamai, Cloudflare, and Radware. We also identify "sibling" ASNs, which are other ASNs of the five scrubbers that they do not use for scrubbing. We employ the CAIDA AS rank API [6], augmenting it with bgp.tools [3] to capture siblings missed by CAIDA, for this purpose.

We summarize the ASNs of scrubbers in Table 1, which includes the sources (e.g., blogs, past incidents) we used to confirm the identified ASNs. For example, we determined that Akamai's ASN is 32787 for DDoS protection, which we also verified from their website [1], and by validating past DDoS scrubbing events [21, 39, 30].

4.2 Detecting always-on scrubbing

We consider a prefix as protected by always-on scrubbing if its AS path has a scrubber as an upstream (second last AS in the path) for 30 days in a row, as seen in BGP Routing Information Bases (RIBs) data. This definition infers

[4] https://github.com/shyamkkhadka/scrubbing_detection

always-on protection based on our limited observation window and the typical minimum subscription length of 30 days reported in the product sheets of several scrubbers [36, 52, 51], but it does not provide direct evidence of uninterrupted protection.

We collect RIBs data containing AS paths with the scrubber ASN anywhere on the path using a daily snapshot from all RIS route collectors at 00:00:00 GMT from 1st May to 30th May, 2025. The RIBs data provides a snapshot of the global routing data as seen by all RIS public route collectors at 8-hour intervals, one of which is at 00:00:00 GMT.

Next, we find the records where the scrubber ASN appears as the second last AS in the AS path for all 30 days. This indicates that traffic always passes through the scrubber before reaching the protected AS. We drop any routes where the origin AS matches the scrubber ASN or any of its sibling ASNs, thus ensuring that only prefixes originating from non-scrubber ASNs are included.

4.3 Detecting same-day on-demand scrubbing

A same-day scrubbing is an on-demand scrubbing which is activated and deactivated within the same day. Figure 2 shows an overview of our methodology to detect same-day on-demand scrubbing for the upstream-change case. Our objective is to determine whether a prefix was scrubbed on a given day rather than how frequently it was scrubbed on a day. We therefore do not track individual cycles of activation, deactivation, or reactivation of a prefix in a day. Instead, we focus on detecting and characterizing scrubbing events at the daily level.

We consider a window of 24 hours (from 00:00:01 GMT until 23:59:59 GMT) because DDoS attacks are often short-lived. For example, Cloudflare's Q2 2025 notes that 92% of attacks are mitigated within 10 minutes [12]. Similarly, Netscout reports that only 1.88% of attacks last more than 12 hours [43]. A short-lived DDoS attack can, however, still span the boundary between two days, for instance, if it occurs a few minutes before and after midnight. We term such cases as cross-day scrubbing, which in our analysis is limited to the attacks spanning at most two consecutive days, and explain it in Section 4.4.

Daily BGP RIBs: We first collect daily snapshots of BGP RIBs data at 00:00:00 GMT (red lines in Figure 2) and filter out private prefixes, default routes, reserved, and unallocated prefixes using Team Cymru's bogon lists [4], ensuring that only valid BGP prefixes are included in our analysis.

Next, we extract the AS paths of protected prefixes that have one of the five scrubbers we study as an origin or as an upstream. These prefixes are potentially being scrubbed at the time of the snapshot, which we use as a reference point to spot new scrubber activations in the next 24 hours, both activations based on origin changes and on upstream changes. In the example of Figure 2), our daily snapshot tells us that scrubber Vercara acts as an upstream for prefixes p1, p2, and p3 at 00:00:00 GMT on day D_1. In such cases, we also check that the origin is not a sibling of the scrubber, ensuring that we identify prefixes belonging to other ASes rather than the scrubber itself.

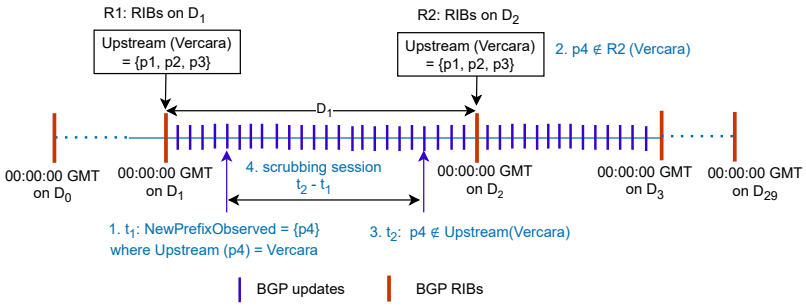

Fig. 2: Steps to detect on-demand scrubbing activated for prefix p4 at t_1 and deactivated at t_2 on day D_1 (same-day scrubbing). The protected AS activates the scrubber for prefix p4 through an upstream change.

Activation signals: Next, we identify scrubber activation signals up to 23:59:59 on day D_1 by monitoring BGP updates that the route collectors dump every five minutes (blue lines in Figure 2). We use BGP updates instead of RIBs dumps because RIS collectors generate RIBs data only every eight hours, whereas BGP updates provide all routing changes at five-minute intervals, capturing any route changes that occur within each interval.

We say that scrubbing was activated for a prefix p during day D_1 if (1) p was not in the RIBs snapshot of D_1 and (2) either one of the five scrubbers began originating announcements for p before 23:59:59 (scrubber activation through an origination change) or if the scrubber appeared as an upstream in p's AS path before 23:59:59 (activation through an upstream change). Step 1 in Figure 2 illustrates this step for prefix p4.

For example, in our data we observed a prefix *182.16.18.0/24* where scrubber *AS19905* (Vercara) appeared as an upstream in AS path *[29504, 15935, 174, 3257, 19905, 45753]* at 11:17:45 on 2025-05-12 in the BGP update files. Since this prefix was not in the RIBs snapshot of 00:00:00 GMT on 2025-05-12, we say scrubbing was activated for it on that day.

De-activation signals: We consider a scrubber to deactivate for a particular prefix p on day D_1 when (1) the scrubber stops originating p and the protected AS or an AS different from the scrubber starts announcing p (deactivation through origin change) before 23:59:59 on D_1 or (2) when the scrubber is not an upstream any longer for p before 23:59:59 on D_1 (deactivation through upstream change). We determine whether such a deactivation event has occurred by examining the RIBs snapshot of day D_2 and comparing it to BGP updates of D_1. For example, prefix p4 is no longer in the RIBs snaphot of D_2 (step 2 in Figure 2), which means that scrubbing was deactivated for it somewhere in day D_1. In such cases, we analyze BGP updates on D_1 to identify the precise time at which the scrubber disappeared from the AS paths (step 3). In our data, we observed that the scrubber (*AS19905*) is no longer an upstream for the prefix

182.16.18.0/24 on the path toward *AS45753* on 2025-05-13, which our method flags as a deactivation.

Scrubbing session: We define a scrubbing session as the time elapsed between the first observation of an activation signal for a particular prefix by a route collector peer and the last observation of that change (step 4 in Figure 2). For example, we observed a BGP update where scrubber *AS19905* first appeared on the AS path *[48362, 1299, 6453, 19905, 45753]* for prefix *182.16.18.0/24* on 2025-05-12 at 11:17:45 GMT (activation based on upstream change). The last time we saw it was at 11:32:49 GMT on that same day (deactivation), which means the presumable scrubbing session (and DDoS attack) for that prefix is around 15 minutes.

Accuracy of (de)activation signals: We require that at least two route collector peers observe the BGP updates corresponding to an activation or deactivation event before we flag it as such, to avoid relying on a single collector that may have stale data. For example, we saw the scrubbing activation signal for *182.16.18.0/24* at 215 collector peers. Our rationale is that seeing the same (de)activation event at multiple collectors increases the probability that a scrubber indeed started or stopped scrubbing because the BGP changes spread across the Internet. We limit the threshold to two collector peers to minimize operational effort.

We ignore scrubbing sessions of less than 1 minute because they could be route flapping cases involving a scrubber, where an AS originates its prefix to a scrubber and withdraws rapidly over a short period of time. We choose 1 minute because the recommended BGP Minimum Route Advertisement Interval (MRAI) is 30 seconds [27], and network operators can change that interval. The MRAI timer controls how frequently a BGP speaker can send updates to a peer, enforcing a minimum interval between consecutive advertisements of the same prefix.

4.4 Detecting cross-day on-demand scrubbing

We detect scrubbing that activates on a certain day and deactivates the following day (cross-day) by combining coarse and fine-grained analyses of BGP data, as shown in Figure 3. We use a 48-hour window for our analysis, comparing RIBs from two consecutive days. Although the number of days to consider is not fixed, this window is sufficient to capture relevant scrubbing events [12, 43].

We use the same concepts as in the case of same-day scrubbing (Section 4.3) in terms of methods for BGP data collection, scrubbing session, (de)activation signals, and accuracy of (de)activation signals. The differences are about examining changes in BGP RIBs and the activation signal in cross-day scrubbing.

We consider scrubbing to be activated for a prefix p before day D_1 if (1) p was in the RIBs snapshot of D_1 but not in RIBs snapshots of D_0 and D_2 and (2) either one of the five scrubbers began originating announcements for p before 00:00:00 on D_1 (scrubber activation through an origination change) or if the scrubber appeared as an upstream in p's AS path before 00:00:00 D_1 (activation

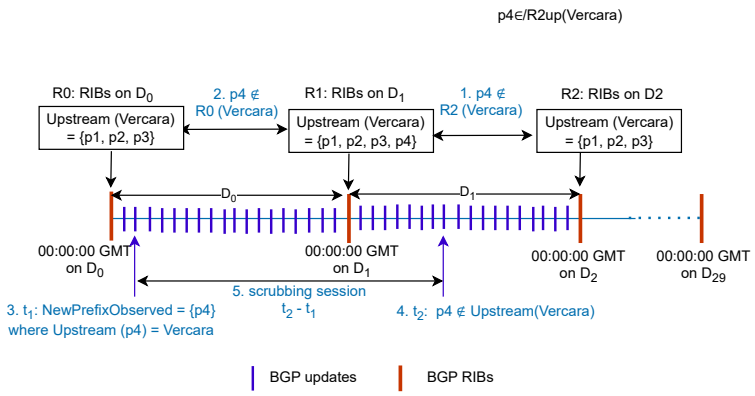

Fig. 3: Steps to detect on-demand scrubbing activated for prefix p4 on day D_0 and deactivated on day D_1 (cross-day scrubbing). The protected AS activates the scrubbing for p4 through an upstream-change.

through an upstream change). The detection method involves the following two steps.

Changes in BGP RIBs: As a coarse-grained analysis, we examine changes in the snapshots on reference day D_1 compared to those of days D_0 (previous day) and D_2 (following day) to identify the prefixes whose origin or upstream has changed. Steps 1 and 2 in Figure 3 illustrate this step for prefix p4. This prefix might be being scrubbed because it is in the snapshot of D_1, but not in the snapshot of D_0 and D_2. We investigate that further in the next step. As an example, prefix *46.184.88.0/24* appeared in our analysis with upstream *AS19905* in the RIBs snapshot collected on 2025-05-03 at 00:00:00 GMT. We checked the prefix's upstream on 2025-05-02 and 2025-05-04 at 00:000:00 GMT, where it had a different upstream (*AS47794*) on both days, which means it might have been scrubbed using an upstream-change-based activation.

Activation signals: Next, our fine-grained analysis involves monitoring BGP updates for the prefixes that are potentially being scrubbed before and after the D_1 snapshot, as identified in the previous step. We analyze BGP updates from 00:00:00 GMT on day D_0 to 23:59:59 GMT on day D_1, to find the activation and deactivation times for prefix p4 (t_1 and t_2 in Figure 3, respectively), which are the times the scrubber (dis)appeared as an origin or as an upstream in BGP updates. Steps 3 and 4 in Figure 3 illustrate this process.

In our data, we observed that the first time a route collector saw the prefix *46.184.88.0/24* with scrubber *AS19905* as its upstream was 2025-05-02 03:19:13 GMT (step 3), and the last time was 2025-05-03 23:59:35 GMT (step 4). We flag this prefix as a scrubbed prefix across days starting on 2025-05-02.

4.5 Identifying potential origin-change using ROA configurations

In addition to detecting dynamic scrubbing (de)activation through BGP origin changes (see Sections 4.3 and 4.4), we also identify prefixes that are likely to use origin change-activated scrubbing by analyzing their ROA objects. For this analysis, we examine ROA objects published by the five RIRs (RIPENCC, APNIC, ARIN, LACNIC, and AFRINIC) in which the scrubbers are registered as the origin ASN for prefixes. We then check how many of these prefixes have ROAs listing a different ASN as the origin. We use the ROA data on 30th May 2025, the end date of our study period. We exclude any ASNs that are siblings of the scrubber.

Our method is grounded in the operational requirement that only ASNs authorized in a prefix's ROAs can validly originate it under RPKI. Hence, if a prefix has a ROA authorizing the scrubber's ASN as well as ROAs for the protected AS' ASN, then this strongly indicates that the scrubber originates the prefix during DDoS mitigation. Such ROAs allow the scrubber to announce the prefix during an attack while remaining compliant with RPKI and ROA-based routing security standards. Consequently, origin change-activated scrubbing requires protected ASes to create these ROAs, as documented in scrubbers operation guides, RFCs, and mailing lists [16, 35, 13, 20, 40].

4.6 Analyzing RPKI and IRR management practices of scrubbers

Finally, our methodology aims to characterize the RPKI and IRR management practices of scrubbers that protected ASes activate through origin changes. This is important because ROA-validating ASes may drop announcements with RPKI statuses Invalid or Notfound (see Section 2.2) [24]. In addition, ASes that perform IRR-based filtering may also reject announcements if they cannot retrieve a matching IRR record, which might reduce the effectiveness of scrubbing.

We examine the prefixes that we obtained from our same-day and cross-day scrubbing detection methodologies (see Section 4.3 and Section 4.4) to characterize such practices.

RPKI practices. We assess if an origin-change activated scrubber properly implements ROA management practices by validating the ROAs at the dates when scrubbing was activated for a prefix, with the outcome being RPKI-valid (practices properly implemented), RPKI-Invalid (practices not properly implemented), or RPKI-Notfound (practices not implemented). We use the RPKI archive for this purpose, a publicly available repository of cryptographically signed ROA objects collected from all RIRs [49]. We analyze and verify the ROA records from the five RIRs .

IRR practices. For a scrubber-originated prefix that has RPKI statuses Invalid or Notfound, we also determine whether an ASN was registered as an origin in IRRs by examining the prefix's route objects in the IRRs operated by the five RIRs and by inspecting RADb (Routing Assets Database) using the RADb API [45]. While the five RIRs collectively cover all global IP address allocations and thus provide broad visibility, their IRRs alone are insufficient to

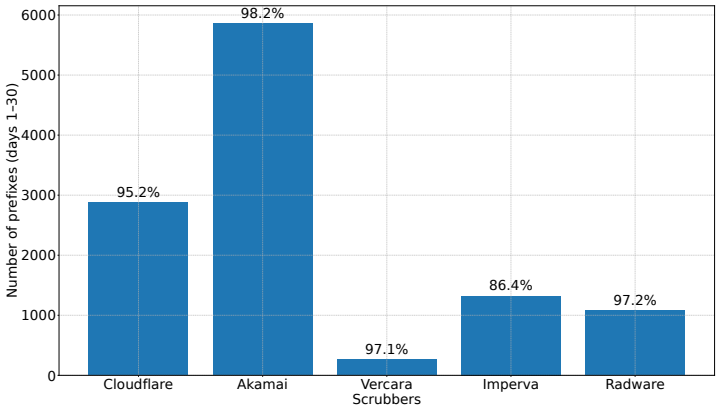

Fig. 4: Always-on protected prefixes. The y-axis shows the number of prefixes that retain the scrubber on their AS paths each day over our 30-day study period. The percentages above the bars indicate the share of always-on protected prefixes relative to prefixes for which the scrubber appeared as an upstream on the first day.

capture all existing route object registrations because many operators maintain entries in third-party IRRs for operational or historical reasons. To improve coverage and completeness, we include RADb, one of the largest and oldest public IRRs operated by Merit Network.

5 Results on Detecting and Characterizing Scrubbing

This section presents our findings on detecting and characterizing scrubbing, which we obtained by applying our methodology to the global top five scrubbers. Our study period is 2025-05-01 (1st May) until 2025-05-30 (30th May).

5.1 Always-on scrubbing

Figure 4 shows the number of prefixes using always-on protection, as identified through our methodology (Section 4.2). We observe that Cloudflare and Akamai account for the largest number of always-on protected prefixes, which are 2,876 and 5,861, respectively. Imperva and Radware have a more moderate number of always-on protected prefixes (1,325 and 1,082). In contrast, Vercara covers only 264 prefixes in always-on mode, indicating a more limited deployment.

The bars in Figure 4 show how many prefixes consistently retained a scrubber on their AS paths throughout the 30-day observation window. The proportions of always-on protected prefixes relative to the total number of protected prefixes for a particular scrubber on the first day are 95.2%, 98.2%, 97.1%, 86.4%, and 97.2% for the five scrubbers, respectively. For example, for Cloudflare, we find

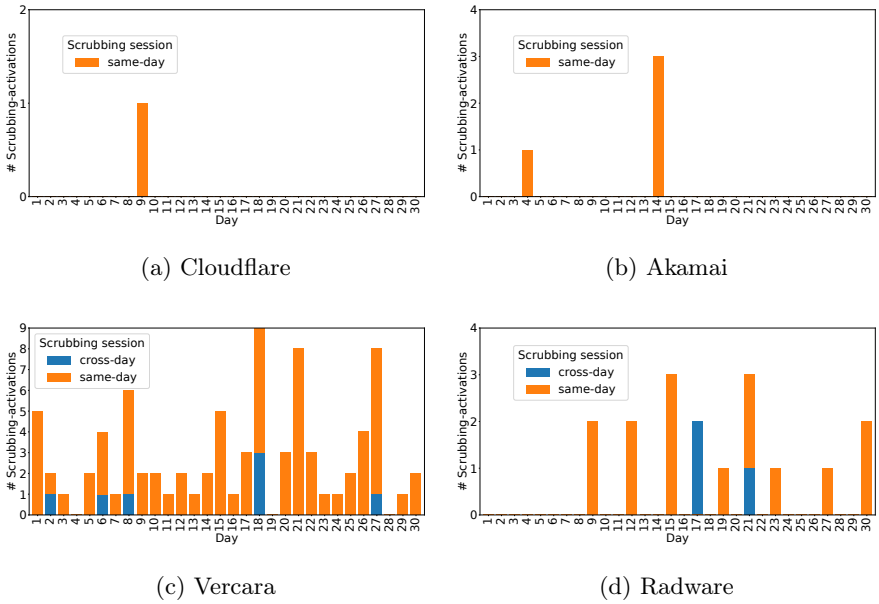

(a) Cloudflare (b) Akamai

(c) Vercara (d) Radware

Fig. 5: Origin-change-based scrubbing activations for four scrubbers. We did not find any origin-change activations with Imperva. The y-axes use different scales for readability.

2,876 prefixes that were always scrubbed in our study period, which corresponds to 95.2% of the prefixes that Cloudflare protected on 1st May.

These results confirm that the majority of prefixes remain protected for at least one month, consistent with our definition of always-on scrubbing described in Section 4.2. The only notable deviation is observed for Imperva, of which 86.4% had that scrubber on their AS path for 30 consecutive days.

Key takeaway: We identify 11,408 prefixes with always-on protection from the five scrubbers, with an average of 94.8% of prefixes present on the first day remaining continuously protected over the 30-day period.

5.2 On-demand scrubbing: (de)activation through origin changes

We address two research questions in this section. First, we investigate how widely origin change-activated scrubbing is being used (RQ2 in Section 1), for which we use the methodology we discussed in Sections 4.3 and 4.4. Next, we determine what the ROA and IRR management practices of origin change-based scrubbers are (RQ3) using the methodology from Section 4.6.

Figure 5 shows the number of prefixes that were protected using origin change-activated scrubbing for our four scrubbers, with same-day scrubbing and cross-day scrubbing for the 30 days of our study period. We found no such cases for Imperva, so we did not plot them. Cloudflare and Akamai each have a very

small number of prefixes that they protect using origin-change activations: 1 and 4 prefixes, which were activated and deactivated on the same day. We did not observe any cases of cross-day scrubbing for them.

We observe a large number of origin-change activations involving Vercara and Radware, with 82 and 17 total activations, respectively. Of these, same-day and cross-day scrubbing account for 75 and 7 activations for Vercara, and 14 and 3 activations for Radware.

Key takeaway: Out of the five scrubbers, the protected prefixes of Vercara and Radware seem to use the origin-change model mostly. We suspect that the origin-change-based activations are used in a limited way, as we did not observe this behavior for the other three major scrubbers.

RPKI management practices of scrubbers. We examine the ROA status of all prefixes on the Internet to determine whether scrubbers were authorized to originate them, using the methodology we described in Section 4.6.

Figure 6a shows the distribution of RPKI statuses for four of our five scrubbers that originate prefixes of protected ASes (Imperva has no prefixes protected by origin-change activations). Cloudflare and Akamai have a very low number of protected prefixes they originate: 1 and 4, respectively (top of the bars).

We observe that 51% of the prefixes originated by Vercara for scrubbing are RPKI Valid, while 16% (13 prefixes) are RPKI Invalid, and 33% (27 prefixes) are NotFound. All of the Invalid cases are due to an origin AS mismatch, meaning that the prefix originated by the scrubber does not have a ROA authorizing the scrubber's ASN as the origin. For Radware, we find a larger proportion of NotFound prefixes compared to valid ones, accounting for 53% and 47%, respectively, with no Invalid cases observed. *Key takeaway: Out of 104 scrubbed prefixes, 52% of scrubbed prefixes were RPKI Valid, 12.5% are invalid, while 35.5% were not found, indicating that the scrubbers do not strictly follow RPKI/ROA practices. This may reduce mitigation effectiveness, as ROA filtering [5] ASes could filter routes from these scrubbers.*

IRR management practices of scrubbers. Among the five scrubbers we studied, we observed that Akamai, Vercara, and Radware originated prefixes with RPKI-Notfound or Invalid status. For Vercara, out of 40 prefixes with RPKI status Invalid or Notfound, we found that 16 (40%) were registered in at least one of the six IRRs we examined, while 24 (60%) were absent from all six IRRs. Radware originated 8 prefixes with Invalid or Notfound RPKI status, 6 of which were registered in one of the IRRs. Overall, we observed that 24 out of 50 prefixes (48%) that underwent scrubbing and had Invalid or Notfound RPKI status were registered in at least one of the six IRRs, whereas 52% were not present in any of them.

We speculate on two reasons for such a high percentage of prefixes not found in IRRs. First, they might be registered to other IRRs beyond our study. Second, the scrubbers might have some private agreements (e.g., Letter of Autho-

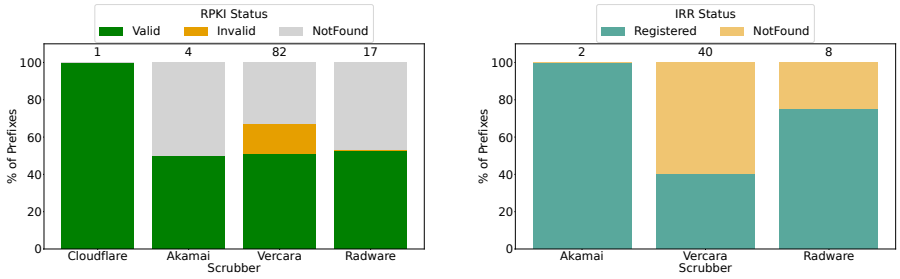

(a) RPKI status (b) IRR status

Fig. 6: (a) RPKI status of prefixes protected by origin-change-based activations. The numbers at the top of the bars indicate the total number of prefixes originated by each scrubber. (b) IRR registration details of prefixes having RPKI-Invalid or Notfound statuses in five RIRs and RADb. Numbers at the top bars indicate the total number of prefixes originated by each scrubber with RPKI-Invalid or Notfound statuses.

rization [33, 34]) with their protected ASes and their upstreams to originate the protected ASes' prefixes under an attack.

Key takeaway: Only 48% of the total number of prefixes that underwent scrubbing are registered into one of the 6 IRRs we examined. Since the scrubbers are not authorized to originate these prefixes based on either ROA or IRR data, such routes may have been dropped during a DDoS attack by ASes performing ROA or IRR-based filtering, potentially limiting scrubbing effectiveness.

On-demand scrubbing from ROA configurations. We use our methodology of Section 4.5 to identify prefixes that are likely to use on-demand scrubbing based on the origin-change model. Figure 7 shows the distribution of ROAs in which prefixes are registered with both scrubbers and other ASNs as origins across five RIRs. Among the scrubbers, we observe that Vercara has the highest number of such prefixes across all five RIRs (amber bar in the figure), which potentially shows it protects many prefixes that use origin-change-based on-demand scrubbing. We counted the number of such prefixes in all five RIRs and aggregated them. In our case of Vercara, we see that the ROAs registered in five RIRs is the maximum (3,237 prefixes), having ROAs for both the scrubber and other ASNs. Radware comes next with 935 such prefixes registered in the five RIRs.

Figure 7 also complements our findings in Section 5.2, where we observed a high number of origin change-based activations for Vercara and Radware from BGP data. We observe a very low number of such cases for Akamai and Cloudflare, indicating low usage of origin change-based on-demand scrubbing cases by them. We aggregate such prefixes across five RIRs and find a total of 4,475 prefixes that are likely to use origin-change-based on-demand scrubbing.

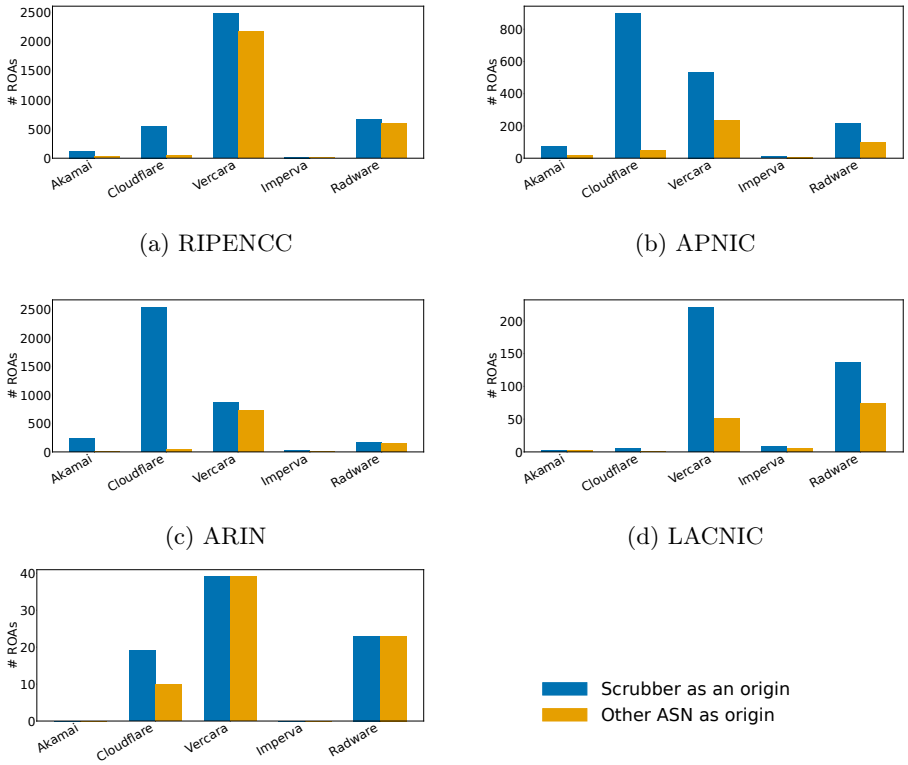

(a) RIPENCC (b) APNIC

(c) ARIN (d) LACNIC

(e) AFRINIC

Fig. 7: Distribution of ROAs listing the prefixes with scrubber as the origin in five RIRs, highlighting how many of these prefixes also have ROAs registered under other ASNs, indicating potential for origin-change during on-demand scrubbing.

Key Takeaway: We identify up to 4,475 prefixes with ROAs authorizing scrubbers, suggesting that these prefixes likely employ origin-change–based on-demand scrubbing. Notably, Vercara and Radware account for the majority of such cases, with 3,237 and 935 prefixes, respectively.

5.3 On-demand scrubbing: (de)activation through upstream changes

We assess the extent to which a scrubber appears as an upstream of a protected AS for scrubbing purposes, which gives the prefixes that were upstream change-activated. Figure 8 shows that Cloudflare regularly engages in this type of scrubbing, with the daily number of scrubbed prefixes ranging from a minimum of 2 to a maximum of 16. We observe that the frequency of activation

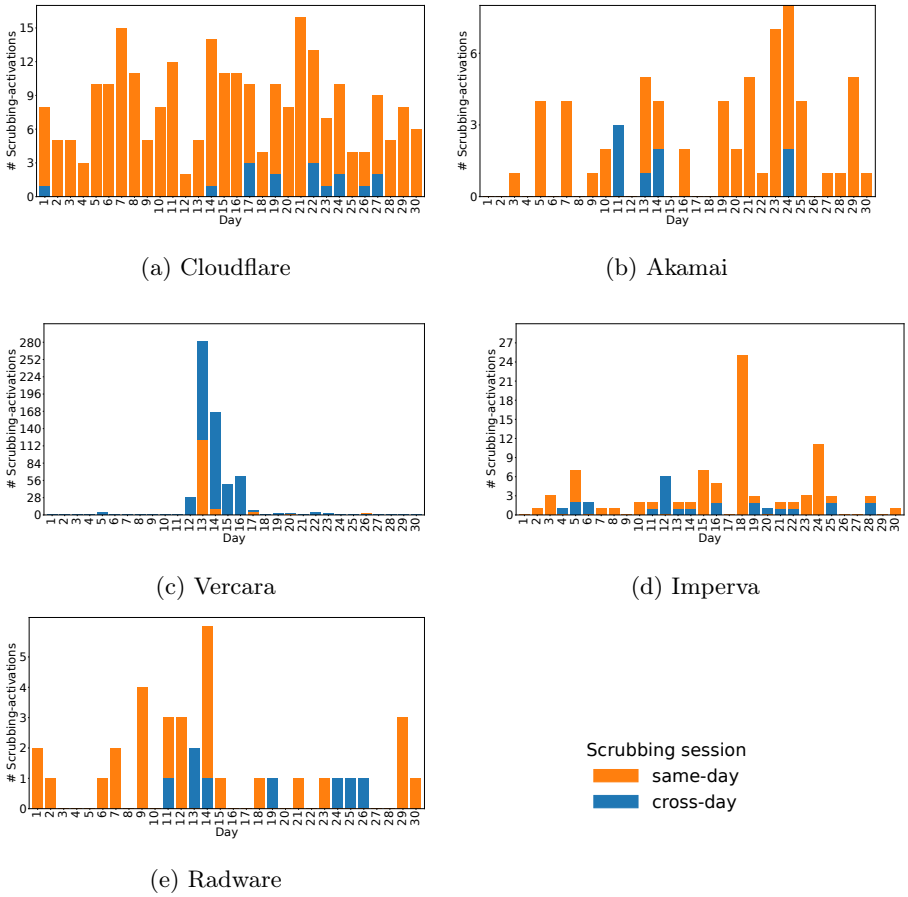

(a) Cloudflare

(b) Akamai

(c) Vercara

(d) Imperva

(e) Radware

Fig. 8: Upstream-change-based scrubbing activations for five scrubbers. The y-axes use different scales for readability.

and deactivation cycles for protected ASes of Cloudflare is higher than the other scrubbers. Akamai and Imperva follow Cloudflare in terms of regular scrubbing, as we did not see any activations for them for some days. Vercara showed a high number of prefixes scrubbed on the 13th May: 282 prefixes. Radware also scrubbed a relatively low number of prefixes: 38 prefixes in 30 days.

Key Takeaway: We observe daily DDoS scrubbing activity in our study period that follows the scrubbing activation model based on upstream-changes, with Cloudflare and Akamai accounting for the largest share of daily scrubbing activity.

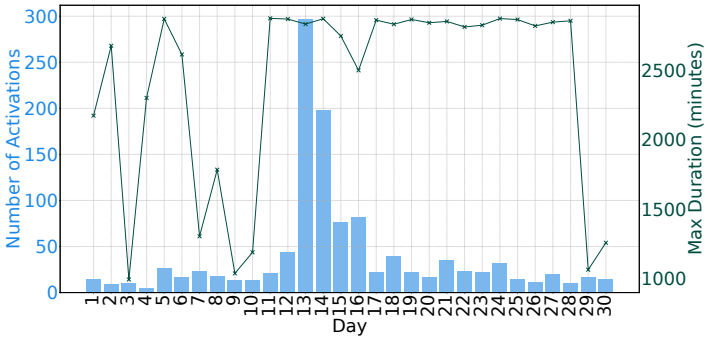

Fig. 9: Aggregated scrubbing events across all five scrubbers, showing daily activation counts and the maximum scrubbing session duration.

5.4 Hints about DDoS attacks from behavior of on-demand scrubbing

This section provides an initial characterization of DDoS attacks mitigated by on-demand scrubbers by examining the frequency and maximum duration of scrubbing activations. We then analyze the distribution of scrubbing sessions with their durations, total scrubbing activations, and the top five scrubbed prefixes during our study period. To do this, we aggregate scrubbing events across all five scrubbers, including both origin-change and upstream-change activations that occur within the same day or across consecutive days (cross-day) over the 30-day study period. While these scrubbing durations offer useful hints about the periods during which prefixes may be under mitigation, they should be interpreted with caution, as scrubbing can remain active for reasons unrelated to actual attack activity (e.g., operator caution, delayed deactivation, human error, contractual arrangements, or pricing and notification policies).

Frequency: We define scrubbing frequency as the number of prefixes that are activated per day in our 30-day study period. Although we found cases where a protected AS activates scrubbing for the same prefix multiple times a day, we do not consider these repeated activations for determining frequency. This is because our focus is on detecting and characterizing scrubbing events at the daily level rather than tracking the (de)activation cycles of individual prefixes within a day, and also simplifies our methodology.

The frequency of scrubbing activations indicates the number of DDoS attacks in a day.

Maximum duration: The maximum duration of a scrubbing session indicates the maximum duration of the DDoS attacks that the scrubber presumably handled during activation. We consider determining the exact activation and deactivation times, and thus the precise duration of the mitigation, as future work.

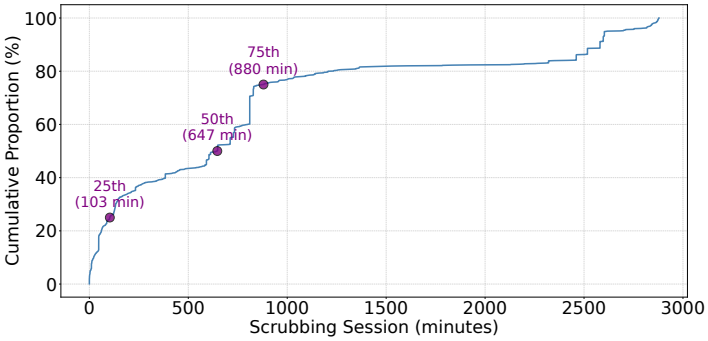

Fig. 10: Cumulative distribution of scrubbing sessions.

Figure 9 highlights the daily frequency of scrubbing activations and the maximum scrubbing session durations. We observe that the maximum duration of scrubbing in our 30-day study period is 2878.68 minutes (48 hours) for three prefixes that were originated by *AS28846* on 2025-05-11. Since we limited our study window to 48 hours, we missed scrubbing sessions that take longer, but these are rare as we discussed in Section 4.3.

The maximum duration of a scrubbing session within a day in our study period is 1,346 minutes (22 hours and 26 minutes) for prefix *66.248.170.0/24* on 2025-05-20. The minimum duration of the scrubbing session was 1 minute for the prefix *2401:e380:10::/48* originated on 2025-05-21 at 05:49:55. We observed 291 scrubbing activations on 2025-05-13, which is the maximum in our study period.

Distribution of scrubbing sessions with their durations: Figure 10 shows the cumulative distribution of the durations of scrubbing sessions. We see that 25% of them lasted less than 1 hour and 45 minutes (around 103 minutes). This indicates that short-lived attacks are quite common, and many scrubbing operations resolve quickly. Half of the scrubbing sessions ended within about 10 hours and 45 minutes (around 647 hours), while the other half persisted longer. Additionally, we see 25% of scrubbing sessions lasted longer than around 15 hours (around 880 minutes). These findings also validate our assumption of considering 24 hours as the maximum duration of a DDoS attack in Section 4.3. While some scrubbers report mitigation within 10 minutes [12, 43], we find that AS operators often keep scrubbing active longer. We learnt from NaWas [41] that this is done as a precaution, beyond the immediate duration of the attack.

Total scrubbing activations. Table 2 shows the number of scrubbing activations that the five scrubbers faced in a month, along with the number of origin-change-activated and upstream-changed activations. We see that Vercara handled the largest number of activations (703), suggesting they handled the most DDoS attacks. The next highest is Cloudflare (250 activations). Radware went through the lowest number activations (55) for the prefixes it protects.

Table 2: Origin-change and upstream-change activations for the five scrubbers.

Scrubbers	#Origin-changed	#Upstream-changed	#Total
Cloudflare	1	249	250
Akamai	4	66	70
Vercara	82	621	703
Imperva	0	96	96
Radware	17	38	55

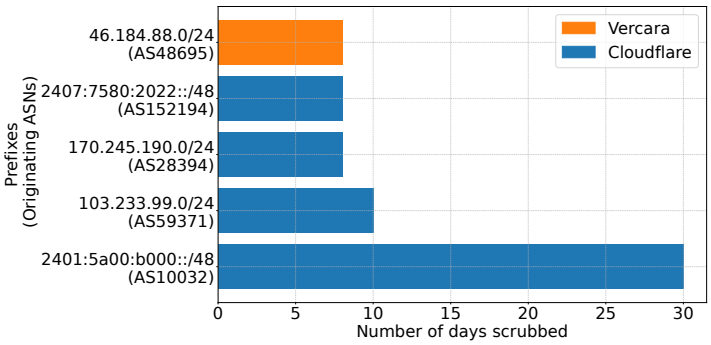

Fig. 11: The top five prefixes in terms of the number of days they underwent scrubbing, their originating ASes, and their scrubbers.

Top five scrubbed prefixes: Figure 11 shows the top five protected prefixes in terms of the number of days that they underwent on-demand scrubbing in our study period. We observe that Cloudflare handled the majority of the top prefixes with the longest scrubbing time, while Vercara was responsible for a smaller subset. The prefix *2401:5a00:b000::/48* underwent upstream-change-based scrubbing at least once in the 30 days of our study period, with different durations. The originating ASN of that prefix under normal conditions (not scrubbing) is *AS10032* (BDX DC Services (HK) Limited), which is a data center company.

Key Takeaway: We observe that the daily duration of scrubbing sessions ranges from a minimum of 1 minute to a maximum of approximately 22 hours for same-day scrubbing, while the maximum is around 48 hours for cross-day scrubbing, with 25% of sessions lasting less than 1 hour and 45 minutes.

6 Discussion and Limitations

First, we answer the four research questions (RQs) that we raised in Section 1. Then, we present the limitations of our study.

RQ1: Dominance of DDoS scrubbing protection types. Our study of a 30-day period and based on the global top five scrubbers shows that there

are more protected prefixes that use always-on than on-demand scrubbing. We observe that the number of always-on protected prefixes is 11,408. We find 1,174 prefixes use on-demand protection, which we inferred by analyzing BGP data. In addition, we assess that the number of prefixes that are likely to use on-demand scrubbing is 4,475 based on our analysis of ROA records. Among these, 54 prefixes with RPKI-valid status underwent scrubbing in the origin-change cases we analyzed. In total, we thus identify 5,649 prefixes that either use on-demand scrubbing (from BGP data) or are likely to use so (from ROA records).

RQ2: Characteristics of on-demand scrubbing. We observe on-demand scrubbing occurring daily on the Internet. Among the two models of on-demand protection (origin-change and upstream-change), we find that upstream-change is more prevalent, with 104 and 1,070 prefixes, respectively. This differs somewhat from prior studies [54, 29, 7], which primarily discuss the origin-change model of on-demand scrubbing. We see that only 2 scrubbers (Vercara and Radware) out of five in our study mostly change the origin of prefixes of their protected ASes under a DDoS attack. One reason for the dominance of the upstream-change model could be that it gives the owner AS of a prefix flexibility and control to originate their prefix themselves, rather than delegating the prefix origination task to the scrubber AS. In this way, the protected AS also does not need to create ROAs for scrubbers for its prefixes.

RQ3: RPKI and IRR management practices. We observe that 52% of the prefixes that underwent the case of origin-change have RPKI Valid status, while the remaining 48% of the prefixes have RPKI Invalid(12.5%) or Notfound (35.5%)status, which shows that the scrubbers that mostly follow the origin-change model do not strictly follow RPKI and ROA practices. This might affect their DDoS mitigation services and put their protected ASes at risk because other ASes on the Internet that filter Invalid RPKI announcements can drop such BGP announcements. Additionally, we observe 48% of total prefixes that underwent scrubbing are registered in IRRs. The remaining 52% are not found in IRR records. Some ASes that rely on IRR-based filtering may reject or ignore the scrubbed routes, reducing the effectiveness of the mitigation. We suspect that some private agreements exist between the scrubber and the protected AS, allowing the scrubber to originate the protected AS's prefixes. Without such agreements, the routes originated by scrubbers for scrubbing would likely be dropped by other BGP peers of the scrubber that enforce RPKI or IRR-based filtering.

RQ4: Hints about DDoS attacks from scrubbing. We observe daily scrubbing sessions with scrubbing durations ranging from as short as 1 minute to as long as 22 hours in a day, which hints at a wide range of DDoS attack durations that the scrubbers handle. Notably, 25% of the scrubbing sessions have a duration of less than 1 hour and 45 minutes, which illustrates that short-lived DDoS attacks occur frequently. These insights into DDoS attacks could inform regulators and policymakers about the DDoS-protection levels of ASes within a region or a country.

Limitations. There are inherent limitations to using publicly available BGP data. For example, limited coverage and bias in the data collection of RIS route collectors [50], which means we miss some scrubbing-protected prefixes in our analysis. Also, some inaccurate records in IRR data [19] affect our findings about IRR management practices of the scrubbers.

Beyond these inherent dataset limitations, we acknowledge the other limitations related to our methodology. First, we lack ground truth data about DDoS scrubbing to validate our findings about the prefixes that underwent scrubbing. Getting data from scrubbers or from their protected ASes would help us to assess our methodology and validate our findings. Our second limitation is that we might misclassify a prefix that a scrubber originates or a protected AS advertises its announcement to a scrubber temporarily for its testing purpose, as a protected prefix.

7 Conclusion and Future Work

We have presented a methodology that detects scrubbing by combining BGP dumps and BGP updates, thereby providing insights into always-on and on-demand DDoS scrubbing for the global top five scrubbers. Our methodology tracks the changes in BGP paths to infer on-demand scrubbing. We characterize on-demand scrubbing by investigating its different modes of operation and the RPKI and IRR management practices of scrubbers.

We show that a scrubber does not always originate its protected ASes' prefixes to handle an attack. The scrubbers also perform mitigation by allowing the protected ASes to originate their prefixes themselves, with the scrubber coming as the first upstream provider of the protected AS. In fact, we observe that this model is more prevalent during our study period, as reflected in the BGP data.

As future work, we plan to evaluate the effectiveness of scrubbing by measuring the time between the start of a DDoS attack and scrubbing activation, as well as between attack termination and scrubbing deactivation. We plan to get the DDoS attack datasets from scrubbers or from network operators who are using the scrubbing service to know the precise DDoS attack start and termination time.

Acknowledgment

We want to thank the anonymous reviewers for their valuable feedback on our paper. This research was funded by the Dutch Research Council (NWO) as part of the projects CATRIN (NWA.1215.18.003) and UPIN (CS.004). CATRIN is part of NWO's National Research Agenda (NWA).

Ethical Considerations

Our work raises no ethical concerns. Our analysis relies on publicly available datasets.

References

1. Akamai: Prolexic - Comprehensive DDoS Attack Protection, https://www.akamai.com/resources/product-brief/prolexic, Last accessed 13-November-2025
2. Akamai Services Descriptions | Akamai, https://www.akamai.com/site/en/documents/corporate/akamai-services-descriptions.pdf, Last accessed 19-November-2025
3. BGP.tools, https://bgp.tools/
4. Bogon Reference HTTP, https://www.team-cymru.com/bogon-reference-http, Last accessed 10-October-2025
5. Bush, R.: Origin Validation Operation Based on the Resource Public Key Infrastructure (RPKI). RFC 7115 (Jan 2014). https://doi.org/10.17487/RFC7115, https://www.rfc-editor.org/info/rfc7115
6. CAIDA: AS Rank (March 2024), https://doi.org/10.21986/CAIDA.DATA.AS-RANK
7. Chung, T., Aben, E., Bruijnzeels, T., Chandrasekaran, B., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Rijswijk-Deij, R.v., Rula, J., et al.: RPKI is coming of age: A longitudinal study of RPKI deployment and invalid route origins. In: Proceedings of the Internet Measurement Conference. pp. 406–419 (2019)
8. Cloudflare Docs: About | cloudflare network interconnect docs (2024), https://developers.cloudflare.com/network-interconnect/about/
9. Cloudflare Magic Transit Docs: Advertise prefixes (2024), https://developers.cloudflare.com/magic-transit/how-to/advertise-prefixes/, Last accessed 14-November-2025
10. Cloudflare Radar: AS32787 overview (2024), https://radar.cloudflare.com/as32787
11. Flow-based monitoring for Magic Transit, https://blog.cloudflare.com/flow-based-monitoring-for-magic-transit/, Last accessed 19-November-2025
12. Hyper-volumetric DDoS attacks skyrocket: Cloudflare's 2025 Q2 DDoS threat report , https://blog.cloudflare.com/ddos-threat-report-for-2025-q2/, Last accessed 28-September-2025
13. Get started · Cloudflare BYOIP docs, https://developers.cloudflare.com/byoip/get-started/, Last accessed 10-October-2025
14. BGP zombies and excessive path hunting, https://blog.cloudflare.com/going-bgp-zombie-hunting/, Last accessed 20-November-2025
15. Magic Transit Reference Architecture, https://developers.cloudflare.com/reference-architecture/architectures/magic-transit/, Last accessed 20-November-2025
16. RPKI and Cloud DDoS Protection - Corero Network Security, https://www.corero.com/rpki-and-cloud-ddos-protection/, Last accessed 14-October-2025
17. Darwich, O., Pelsser, C., Vermeulen, K.: Detecting traffic engineering from public bgp data. In: International Conference on Passive and Active Network Measurement. pp. 307–334. Springer (2025)
18. DDoS providers, https://x.com/eastdakota/status/1937828176056521011/photo/1, last accessed 10-Sep-2025
19. Du, B., Izhikevich, K., Rao, S., Akiwate, G., Testart, C., Snoeren, A.C., Claffy, K.: Irregularities in the internet routing registry. In: Proceedings of the 2023 ACM on Internet Measurement Conference. pp. 104–110 (2023)

20. Gilad, Y., Goldberg, S., Sriram, K., Snijders, J., Maddison, B.: The Use of maxLength in the Resource Public Key Infrastructure (RPKI). RFC 9319 (Oct 2022). https://doi.org/10.17487/RFC9319, https://www.rfc-editor.org/info/rfc9319
21. GitHub Survived the Biggest DDoS Attack Ever Recorded, https://www.wired.com/story/github-ddos-memcached/, Last accessed 14-October-2025
22. Global Routing Intelligence Platform , https://grip.inetintel.cc.gatech.edu/, last accessed 17-Sep-2025
23. Holmes, D.: The forrester wave™: DDoS mitigation solutions, q1 2021 (2021), https://allofsecurity.pl/wp-content/uploads/2021/03/The-Forrester-Wave-DDoS-Mitigation-Solutions-Q1-2021.pdf, Last accessed 14-January-2025
24. Huston, G., Michaelson, G.: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs). RFC 6483 (Feb 2012). https://doi.org/10.17487/RFC6483, https://www.rfc-editor.org/info/rfc6483
25. Imperva: Advanced DDoS Protection & Mitigation Services, https://www.imperva.com/products/ddos-protection-services/, Last accessed 11-March-2025
26. Imperva Documentation Portal, https://docs.imperva.com/bundle/cloud-application-security/page/introducing/network-ddos-protection.htm, Last accessed 14-June-2025
27. Jakma, P.: Revisions to the BGP 'Minimum Route Advertisement Interval'. Internet-Draft draft-ietf-idr-mrai-dep-04, Internet Engineering Task Force (Sep 2011), https://datatracker.ietf.org/doc/draft-ietf-idr-mrai-dep/04/, work in Progress
28. Jonker, M., Sperotto, A., Pras, A.: DDoS Mitigation: A measurement-based approach. In: NOMS 2020-2020 IEEE/IFIP Network Operations and Management Symposium. pp. 1–6. IEEE (2020)
29. Jonker, M., Sperotto, A., van Rijswijk-Deij, R., Sadre, R., Pras, A.: Measuring the adoption of DDoS protection services. In: Proceedings of the 2016 Internet Measurement Conference. pp. 279–285 (2016)
30. BGP Flowspec Doesn't Suck. We're Just Using it Wrong., https://www.kentik.com/blog/bgp-flowspec-doesnt-suck-were-just-using-it-wrong/, Last accessed 17-July-2025
31. Khadka, S., Bayhan, S., Holz, R., Shokoohi, S., Barcellos, M.: A First Look at the Adoption of BGP-based DDoS Scrubbing Services: A 5-year Longitudinal Analysis. In: Proceedings of the 2025 International Conference on Network and Service Management (2025)
32. Livadariu, I., Fontugne, R., Phokeer, A., Candela, M., Stucchi, M.: A tale of two synergies: Uncovering rpki practices for rtbh at ixps. In: International Conference on Passive and Active Network Measurement. pp. 88–103. Springer (2024)
33. Letter of Agency (LOA) · Cloudflare BYOIP docs, https://developers.cloudflare.com/byoip/concepts/loa/, Last accessed 10-October-2025
34. L3/L4 DDoS Mitigation | F5 Distributed Cloud Technical Knowledge, https://docs.cloud.f5.com/docs-v2/ddos-and-transit-services/how-tos/network-firewall/l3l4-ddos-mitigation, Last accessed 10-October-2025
35. Lumen RPKI Guide, https://assets.lumen.com/is/content/Lumen/rpki-customer-notification?Creativeid=1c967bb7-1321-4be8-92c0-74097e840849, Last accessed 14-October-2025

36. MANRS: DataSheet Emergency OnboardingDDoS, https://real-sec.com/wp-content/uploads/2022/02/DataSheet_Emergency_OnboardingDDoS.pdf, Last accessed 06-October-2025
37. MANRS: MANRS, https://manrs.org/, Last accessed 01-September-2025
38. MANRS: MANRS+ Controls, https://manrs.org/wp-content/uploads/2023/12/MANRSPlus_Controls.pdf, Last accessed 01-September-2025
39. Mike Hicks: Akamai Prolexic Routed Outage Analysis, https://www.thousandeyes.com/blog/akamai-prolexic-routed-outage-analysis, Last accessed 05-February-2025
40. nanog: Re: It can be challenging to advise DDoS mitigation subscribers on their RPKI-ROA needs, https://seclists.org/nanog/2024/Oct/73, Last accessed 10-October-2025
41. NBIP, https://www.nbip.nl/en/nawas/faq/, Last accessed 06-October-2025
42. Phil Gervasi: How Kentik Visualizes the BGP Propagation of a DDoS Mitigation (2022), https://www.kentik.com/blog/how-bgp-propagation-affects-ddos-mitigation/, Last accessed 14-November-2024
43. How Long Does a DDoS Attack Last? , https://www.netscout.com/blog/how-long-does-ddos-attack-last, Last accessed 28-September-2025
44. Cloudflare Radar, https://radar.cloudflare.com/, last accessed 17-Sep-2025
45. RADb , https://www.radb.net/, Last accessed 28-September-2025
46. Choosing the Best Diversion For Your Needs (2019), https://support.radware.com/app/answers/answer_view/a_id/1018554/related/1, Last accessed 05-December-2025
47. Radware Doc: DDoS Protector Cloud Service. https://www.radware.com/getattachment/bfc20642-47c5-4e1a-adb4-40350695541e/ds-checkpoint-ddos-protector-cloud-service.pdf.aspx (2024), Last accessed 12-November-2025
48. Radware Support: How to setup GRE tunnels (2019), {https://support.radware.com/app/answers/answer_view/a_id/1018552/~/how-to-setup-gre-tunnels}, Last accessed 05-December-2024
49. RIPE RPKI, https://ftp.ripe.net/rpki/, last accessed 17-Aug-2025
50. Sermpezis, P., Prehn, L., Kostoglou, S., Flores, M., Vakali, A., Aben, E.: Bias in internet measurement platforms. In: 2023 7th Network Traffic Measurement and Analysis Conference (TMA). pp. 1–10. IEEE (2023)
51. DDoS Hyper: DDoS Protection Solution | Lumen, https://www.lumen.com/en-us/security/ddos-hyper.htm.html, Last accessed 10-October-2025
52. eSecurity Planet, https://www.esecurityplanet.com/products/distributed-denial-of-service-ddos-protection-vendors/, Last accessed 10-October-2025
53. Team, Vercara: UltraDDoS protect - FAQs (2024), https://vercara.com/resources/ultraddos-protect, last accessed 19-November-2025
54. Testart, C., Richter, P., King, A., Dainotti, A., Clark, D.: Profiling BGP serial hijackers: capturing persistent misbehavior in the global routing table. In: Proceedings of the Internet Measurement Conference. pp. 420–434 (2019)
55. Tung, T.M., Wang, C., Wang, J.: Understanding the behaviors of BGP-based DDoS protection services. In: Man Ho Au et al. (ed.) Network and System Security. pp. 463–473. Springer International Publishing (2018)
56. DDoS Protection Services | Cloud-based DDoS Mitigation | UltraDDoS Protect, https://vercara.digicert.com/ddos-protection, Last accessed 19-November-2025

57. Release Note: DDoS Protection for Networks - Manually divert your ranges | Imperva Cyber Community, https://community.imperva.com/discussion/release-note-ddos-protection-for-networks-manually-divert-your-ranges, Last accessed 19-November-2025
58. Wallace Lee: DDoS Protection for Networks: Combatting Local Preference from ISPs | Imperva (2020), https://www.imperva.com/blog/ddos-protection-for-networks-combatting-local-preference-from-isps/, Last accessed 16-July-2025