

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/335823309>

A Taxonomy of Attacks Using BGP Blackholing

Chapter · September 2019

DOI: 10.1007/978-3-030-29959-0_6

CITATIONS

10

READS

588

2 authors:



Loïc Miller

University of Strasbourg

5 PUBLICATIONS 15 CITATIONS

SEE PROFILE



Cristel Pelsser

University of Strasbourg

80 PUBLICATIONS 891 CITATIONS

SEE PROFILE

Some of the authors of this publication are also working on these related projects:



Internet Health Report [View project](#)

A Taxonomy of Attacks using BGP Blackholing

Loïc Miller and Cristel Pelsser

University of Strasbourg, 4 Rue Blaise Pascal, 67081 Strasbourg, France
{loicmiller,pelsser}@unistra.fr

Abstract. BGP blackholing is a common technique used to mitigate DDoS attacks. Generally, the victim sends in a request for traffic to the attacked IP(s) to be dropped. Unfortunately, remote parties may misuse blackholing [57, 29] and send requests for IPs they do not own, turning a defense technique into a new attack vector. As DDoS attacks grow in number, blackholing will only become more popular, creating a greater risk this service will be exploited. In this work, we develop a taxonomy of attacks combining hijacks with blackholing: BGP blackjacks (blackhole hijacks). We show that those attacks effectively grant more reach and stealth to the attacker than regular hijacks, and assess the usability of those attacks in various security deployments. We then find that routing security mechanisms for BGP [30, 31] do not provide an adequate protection against some of those attacks, and propose additional mechanisms to properly defend against or mitigate them.

Keywords: BGP · Security · Blackholing · DDoS · Communities · Hijacks · Leaks

1 Introduction

DDoS attacks are one of the most potent threats to the Internet. With the rise of the Internet of Things (IoT), the number of connected devices is exploding. The potential of a botnet to launch massive Distributed Denial of Service (DDoS) attacks is taking scary proportions [40]. New attack vectors [1, 38] are being discovered and are enabling the largest attacks we have ever seen. In February 2018 for example, Github was under attack, receiving up to 1.3 Tbps of traffic through its CDN, Akamai. Such a high amount of traffic can flood many access links, rendering services behind those links unavailable. These attacks can be motivated by multiple reasons, including but not limited to revenge [27], activism [41], vandalism [42], financial reasons [32] or political reasons [6].

Fortunately, numerous techniques exist to mitigate DDoS attacks [48, 49]. Those techniques can be roughly separated in two categories: proactive mitigation techniques and reactive mitigation techniques. Proactive techniques encompass all the mitigation techniques put in place before an attack happens, like designing protocols with a reduced amplification factor (the amount of traffic one can get in a response compared to the amount of traffic one has to send in a request), or reducing the number of amplifiers available to attackers. Proactive techniques also include response rate limiting, using sessions for UDP, filtering spoofed packets, making use of anycast or even using Access Control Lists.




Fig. 1: BGP Blackholing

DDoS attacks can also be dealt with in a reactive way, by using traffic scrubbing services, where a third party processes the victim’s incoming traffic, detects and mitigates the attack, and then forwards the legitimate traffic to the victim.

While filtering provides a great amount of flexibility, it runs into scalability issues in terms of number of entries and packet rate [29], as well as resources and reaction time [14]. A mitigation technique based on forwarding is thus much more scalable, and this is where BGP blackholing shines.

Blackholing [57, 29] uses the Border Gateway Protocol (BGP) [43] as a means to announce the need for mitigation. BGP is the de-facto inter-domain routing protocol in the Internet, and its primary function is to allow Autonomous Systems (ASes) to communicate with others by exchanging reachability information. More specifically, blackholing is announced via BGP communities [8, 28], optional transitive BGP attributes which are “used to pass additional information to both neighboring and remote BGP peers” [8]. The communities forwarded with an advertisement are interpreted by ASes, which use this information to apply a specific treatment to the route.

Figure 1 depicts blackholing being used to mitigate a DDoS attack. AS 20’s server located at 192.0.2.1/32 is under a DDoS attack going through both its neighbors, AS 10 and AS 30. To mitigate the attack, AS 20 sends an advertisement to AS 30, indicating to blackhole prefix 192.0.2.1/32 by adding the community used to signal blackholing to AS 30, ‘30:666’. The community sent, ‘30:666’ means that AS 30 needs to apply blackholing. In addition to this information, we also usually attach either the NO_EXPORT or the NO_ADVERTISE community to the advertisement, respectively, to keep the scope local to the AS or the router [28].

Blackholing is a very effective mitigation technique [13], but it has a double-edged sword effect: all malicious traffic destined to the blackholed prefix is dropped, but so is legitimate traffic.

The literature highlights shortcomings in BGP communities, namely the lack of standardization and authentication. Firstly, only a handful of communities have a semantic meaning defined in RFCs, the vast majority of them being defined by the AS ‘owning’ them, making them AS-specific [54]. This lack of standardization makes communities harder to classify [15]. In addition, documentation for communities is scattered and incomplete [54]. In the case of blackholing, Giotsas et al. found 307 different community values used to signal blackholing, with an additional 115 labeled as likely [22]. Blackholing is nevertheless frequently used [13], and its use is increasing [22], as it is a very effective way to mitigate DDoS attacks [13, 22]. Even though ASes should keep the scope of blackholing local to the AS or the router, it has been shown that 50% (80%) of

blackhole communities still traverse up to two (four) ASes, with some blackhole communities traversing as many as eleven ASes [54].

Communities are also vulnerable because they can be altered by third parties: “*Because BGP communities are optional transitive BGP attributes, BGP communities may be acted upon or otherwise used by routing policies in other Autonomous Systems (ASes) on the Internet.*” [24]. With other ASes being able to modify communities associated with a BGP advertisement, communities can become a vector of attacks. Solutions to secure Internet routing exist [30, 31], but they focus on securing the AS path, leaving other BGP attributes unprotected. Those solutions also suffer from a lack or absence of deployment, due to the lack of incentives to do so [20].

Attacks trying to falsify BGP attributes to gain an advantage are not new. As BGP is a distributed protocol, lacking authentication of route origins and verification of paths, ASes can advertise illegitimate routes for prefixes they do not own, attracting some or all of the traffic to these prefixes. Those advertisements propagate and pollute the Internet, affecting service availability, integrity, and confidentiality of communications [52]. This phenomenon is called prefix hijacking. In this work, we build on top of prefix hijacking to create new attacks through BGP blackholing: blackjacks. Hijacks and blackjacks are similar, in that they both impact reachability of the affected prefix. However, regular hijacks only poison the ASes near the attacker, whereas blackjacks drop traffic directly at the ASes receiving the advertisement, regardless of AS path length. This means blackjacks have more reach, and are stealthier than simple hijacks.

Considering routing attacks and defenses (Section 2), we construct an attack taxonomy using blackholing as an attack vector (Section 3) and assess the usability of those attacks in different security deployments (Section 4). We then detail good practices and implementations to protect against such attacks (Section 5). Finally, we review related work (Section 6) and conclude in Section 7 by reviewing our contributions and describing the possible perspectives and areas of future work.

2 Background

Prefix hijacking can be caused by misconfiguration [47], or with malicious intent, possibly motivated by retaliation [56], information gathering [34], economical reasons [23] or political reasons [35].

On Figure 2, AS 10 (the victim) advertises a route for the prefix 192.0.2.0/24. The hijacker (AS 40) can fake a direct connection to this network by advertising 192.0.2.0/24 to AS 30. Preferring the shorter AS path, AS 30 will choose a new best route going through AS 40, and forward the advertisement to AS 20. AS 20’s original route is already the best one, so it does not accept the new route and does not forward the advertisement to AS 10.

We base our work on a hijack taxonomy developed in [52], which is based on three dimensions:

- The manipulation of the AS path.
- The affected prefix.
- The way (hijacked) data traffic is treated.




Fig. 2: Prefix hijacking

To illustrate those hijack types, let us reconsider Figure 2, where AS 10 (the victim) owns and legitimately announces 192.0.2.0/24, and AS 40 is the hijacker. For the sake of simplicity, a BGP advertisement is noted as an announced prefix tagged with an AS path. For example, {AS20,AS10 - 192.0.2.0/24} is a BGP advertisement for prefix 192.0.2.0/24 with AS path {AS20,AS10}, originated by the legitimate AS (AS 10). In their paper, they first classify by AS path manipulation, creating three categories of hijacks:

- **Origin AS (or Type-0) hijacking:** The hijacker announces as its own a prefix that it is not authorized to originate (e.g. {AS40 - 192.0.2.0/24}). This type of hijack is sometimes called prefix re-origination, and is the most commonly observed type of hijack [52].
- **Type-N hijacking ($N \geq 1$):** Also called path manipulation in the literature [19, 10, 11]. The hijacker announces an illegitimate path for a prefix it does not own, creating fake adjacencies between ASes. The path contains the ASN of the hijacker as the last hop (e.g. {AS40,AS20,AS10 - 192.0.2.0/24}). Here, AS 40 creates a fake adjacency between itself and AS 20. The position of the rightmost fake link in the forged advertisement determines the type. For example, {AS40,AS10 - 192.0.2.0/24} is a Type-1 hijacking, {AS40,AS20,AS10 - 192.0.2.0/24} is a Type-2 hijacking, etc.
- **Type-U hijacking:** The hijacker leaves the legitimate AS path unaltered (but may alter the advertised prefix). In the case both the AS path and the prefix are left unaltered, the event is not a hijack but rather a traffic manipulation attempt, performed by adding communities to the advertisement for example.

The second discriminant is the affected prefix. The hijacker can either perform an exact prefix hijack, where he announces a path for the same prefix that is announced by the legitimate AS, or he can perform a sub-prefix hijack, where he announces a more specific prefix. In the case of an exact prefix hijack, only the part of the Internet that is close to the hijacker (in terms of AS hops) switches to routes towards the hijacker. In the case of a sub-prefix hijack, the entire Internet traffic is sent towards the hijacker to reach the announced sub-prefix. Note that since most routers do not accept BGP advertisements containing a prefix past a certain length (usually /24) to reduce routing table size, a sub-prefix hijack advertising a /25 or more may not be very effective, as the advertisements will be dropped. There is also the case of squatting, where the hijacker announces a prefix owned but not (currently) announced by the legitimate AS. In this work, we disregard squatting as it is not applicable to blackjack attacks.

The last discriminant is the way the data-plane traffic is handled. Once the hijack is accomplished, the attacker attracts some or all of the traffic originally destined to the hijacked prefix to his own AS. The attacker can then drop the packets (blackhole), impersonate the services tied to the hijacked prefix by responding to the victims (imposture),

eavesdrop on the traffic and forward it back to the victim (interception) [62, 52], and even send spam [59] or carry out other activities.

For example, the hijack depicted in Figure 2 is a Type-0 exact prefix hijack, as AS 40 re-originates 192.0.2.0/24.

In our work, we will classify the attacks only by AS path manipulation and affected prefix, as blackholing attacks have the sole purpose of dropping traffic. Note that this taxonomy can be extended, as it does not cover cases where, for example, the attacker possesses two or more ASes.

Even though techniques to protect oneself against hijacks lack deployment, they still exist and are the go-to solutions to make BGP more secure.

2.1 BGP Routing Security

When receiving an advertisement, a router might want to verify that the included AS path is legitimate. This process is broken down in two validation steps:

- **Origin validation:** Does the origin AS have a right to announce this prefix?
- **Path validation:** Does the sequence of ASes in the AS path reflect the sequence of ASes crossed by this advertisement?

The Resource Public Key Infrastructure Origin validation can be achieved through the Resource Public Key Infrastructure [30]. The RPKI is a distributed, hierarchic public key infrastructure. It allows prefix holders (legitimate holders of IP address space) to emit digitally signed objects, *Route Origin Authorizations (ROAs)*, attesting that a given AS is authorized to originate routes for a set of prefixes.

This way, a given AS can verify that the origin AS present in a given advertisement is authorized to originate the prefix (*Route Origin Validation (ROV)*). While the RPKI provides digitally signed routing objects, it does not sign BGP advertisements, and operates separately from BGP. An advantage of RPKI is that the mapping of prefixes to origin ASes is formally verifiable [37].

BGPsec Path validation can be achieved through BGPsec [31]. BGPsec relies on RPKI as it makes use of certificates.

To secure the path attribute, BGPsec relies on an new optional non-transitive BGP path attribute which replaces the AS_PATH attribute: BGPsec_PATH. The attribute carries digital signatures providing cryptographic assurance that every AS on the path of ASes listed in the advertisement has explicitly authorized the advertisement of the route. BGPsec-compliant BGP speakers (BGPsec speakers) wishing to send BGPsec advertisements to eBGP peers need to possess a private key associated with an RPKI router certificate [46] that corresponds to the BGPsec speakers's ASNs.

Traditional BGP advertisements may still be sent between BGPsec speakers, meaning an attacker can potentially downgrade a BGPsec speaker to regular BGP [33]. BGPsec also does not protect against BGP leaks, which is defined as a violation of the standard model of routing policies, pinpointed by Gao and Rexford [17, 18]. Simply put, the Gao-Rexford model states that ASes have incentives to send traffic along customer routes

(which generate revenue), as opposed to peer routes (which do not generate revenue) or provider routes (which come at a monetary cost). It also models ASes' willingness to transit traffic from one neighbor to another only when paid to do so by a customer. This is important to keep in mind for one of the attacks we define in section 3.

3 Threat Model and Attack Taxonomy

This section is dedicated to the elaboration of an attack taxonomy. We consider a common and general hijacking threat model [50, 52]. An attacker controls a single AS and its border routers. He also has full control of the control plane and the data plane within its own AS. The attacker can arbitrarily manipulate the advertisements that it sends to its neighboring ASes and the traffic that crosses its network. He has no control over advertisements and traffic exchanged between two other ASes.

Even though these attacks can work in numerous configurations, we assume for the sake of explanations that:

Assumption 1 *Every AS uses the Gao-Rexford routing policy model.*

Assumption 2 *Every AS follows the best practices defined in [9] when receiving a blackhole request.*

Those best practices can be summarized as:

1. Set local-preference to 200 (higher preference)
2. Set origin-type to IGP (higher preference)
3. Add the NO_EXPORT community to the advertisement

Following those best practices means that the blackholing advertisement is preferred over other routes and that blackholing is limited to the AS receiving the advertisement.

For the sake of simplicity, a BGP advertisement is noted as an announced prefix tagged with an AS path and communities. For example, {AS20,AS10 - <blackholer AS>:666 - 192.0.2.0/24} is an advertisement for prefix 192.0.2.0/24 with AS path {AS20,AS10}, originated by AS 10, and bearing the blackhole community <blackholer AS>:666, where <blackholer AS> is the AS providing the blackholing service.

Type-0 Blackjack This first attack is also the simplest. Performing a Type-0 blackjack is done by performing a Type-0 hijack and attaching the blackhole community to the advertisement.

Figure 3 shows AS 30 (the victim) advertising a route for 192.0.2.0/24. AS 10 (the attacker) can perform a Type-0 blackjack by re-originating the prefix 192.0.2.0/24, and attaching the blackhole community to the advertisement. Thus, AS 10 sends {10 - 20:666 - 192.0.2.0/24} to its peer. As AS 20 (the blackholer) follows Assumptions 1 and 2, it blackholes traffic destined to 192.0.2.0/24.

This example highlights two main advantages of blackjack attacks:




Fig. 3: Type-0 and Type-N Blackjacks

- **Reach:** The attacker can potentially drop more traffic by sending blackholing advertisements to its neighbors than by hijacking the prefix and blackholing the traffic at his AS. If AS 10 tried to do so, it could not have dropped traffic going through AS 20, as AS 20 would prefer the route going through its customer (AS 30). This is not the case anymore with the blackjack attack, since AS 20 now prefers the advertisement of AS 10 per Assumption 2, thus dropping all traffic destined to the blackholed prefix.
- Blackholing also grants precedence over AS path length, so the longer AS path that generally comes with hijacks is no longer a problem. Considering this, an attacker can effectively target a specific blackholer multiple AS hops away, by using their specific blackhole community value.
- One thing to consider when performing sub-prefix blackjacks with a far away blackholer is that all ASes on the path between the attacker and the blackholer need to forward the advertisement. Since most routers do not accept advertisements containing a prefix past a certain length (usually /24) to reduce routing table size, the blackjack might not reach the blackholer if the targeted prefix is too specific.
- Moreover, when the blackholer applies blackholing, a good practice is to add the NO_EXPORT community, which means that a blackjack targeting a prefix advertised in the Internet will stop the blackholer from advertising this prefix, causing even more disruption. In the case of a sub-prefix blackjack, the prefix will still be advertised, but traffic to the target of the attack will still be dropped at the blackholer even though no routes changed.
- **Stealth:** As the attacker is not the one dropping the traffic, he hides himself better from potential onlookers. Note that it may still be possible to retrieve the source of the attack by looking at the advertisements received by the relevant routers at the time of the attack, even though it might be hard to do so, considering those routers are not likely in the network of the victim.

It is also worth noting that an even stealthier attack is possible, if the blackholer(s) is(are) at multiple hops from the attacker. In this case, not only will the attacker not blackhole the traffic himself, but he will also not be the only one that could have sent a blackhole advertisement, as potentially other ASes could have performed the attack. Since an attacker can target a blackholer that is far away, an attack can propagate far from the source of the attack, increasing the difficulty to detect it and identify the attacker.

In our example, AS 20 is blackholing the traffic, even though it was AS 10 that performed the attack.

A disadvantage of Type-0 blackjack attacks is that some defense mechanisms can detect and counter them. By performing Route Origin Validation, either using IRR records or the RPKI, an AS can effectively know which AS is authorized to announce which prefix. Since in a Type-0 blackjack, the attacker is the origin AS, this type of attack is not effective against ASes performing ROV.

Type-N Blackjack Type-N blackjacks circumvent ROV by creating a false adjacency between the attacker and an AS, in the same way Type-N hijacks work. Indeed, if an AS tries to verify the origin of an AS path, as the origin is legitimate, the AS will deem the origin valid.

Figure 3 depicts a Type-N blackjack. Analogous to our Type-0 blackjack example, AS 30 advertises a route for 192.0.2.0/24. AS 10 can perform a Type-N blackjack by faking an adjacency with AS 30, and attaching the blackhole community to the advertisement. Thus, AS 10 sends {10,30 - 20:666 - 192.0.2.0/24} to its neighbor. As AS 20 follows Assumptions 1 and 2, it blackholes traffic destined to 192.0.2.0/24. A Type-N blackjack has the same reach and stealth properties as a Type-0 blackjack.

Type-N blackjacks can circumvent ROV by creating fake adjacencies, but some defense mechanisms can still detect and counter Type-N blackjacks. By using BGPsec, ASes can verify that the sequence of ASes in the AS path reflects the sequence of ASes crossed by received BGPsec advertisements. In this case, no AS path manipulation is possible, but an attacker can still make use of a subset of Type-U blackjacks.

Type-U Blackjack The Type-U blackjack category regroups all attacks where the AS path is unaltered, meaning the origin AS is authorized to announce the prefix (not like Type-0 blackjacks), and the adjacencies in the AS path reflect real adjacencies (not like Type-N blackjacks).

This category can be broken down into three sub-categories:

- On Path blackjacks.
- On Path blackjacks which violate the Gao-Rexford export rule.
- Not On Path blackjacks.

1. On Path Blackjack (OP) An on path blackjack is characterized by the attacker being on the path of a legitimate advertisement.

Figure 4 depicts an On Path blackjack. Like in the other examples, AS 30 advertises a route for 192.0.2.0/24. AS 10 can perform an On Path blackjack by sending {10,30 -




Fig. 4: On Path Blackjack




Fig. 5: OP-GRV and NOP Blackjacks

20:666 - 192.0.2.0/24} to AS 20. Normally, AS 20 would prefer the route going through its customer (AS 30), however, the blackhole community in the advertisement of AS 10 makes this advertisement preferable to the advertisement of AS 30. Thus, AS 20 will blackhole all traffic destined to 192.0.2.0/24.

2. On path Blackjack with Gao-Rexford Violation (OP-GRV) In this sub-category, the attacker is also on the path of a legitimate advertisement, but violates the Gao-Rexford export rule when propagating the advertisement, imitating the behavior of a BGP leak.

Figure 5 depicts an On Path blackjack which breaks this rule. AS 20 can perform an On Path blackjack with Gao-Rexford violation by sending the advertisement {20,30,40 - 10:666 - 192.0.2.0/24} to its provider (AS 10), making AS 10 blackhole traffic destined to 192.0.2.0/24.

3. Not On Path Blackjack (NOP) The last sub-category contains all other Type-U blackjacks, that is, blackjacks where the attacker is not on the path of an advertisement, but announces a legitimate path. In this sub-category, the origin AS in the AS path is authorized to announce the prefix, the adjacencies in the path reflect real adjacencies, but the attacker is not in the AS path.

Figure 5 gives an example of a NOP blackjack: AS 20 sends the advertisement {40 - 30:666 - 192.0.2.0/24} to AS 30, making AS 30 blackhole traffic destined to 192.0.2.0/24.

On Path blackjacks can be considered stealthier than Type-0 and Type-N blackjacks, as the attacker does not re-originate the prefix and does not create false adjacencies in the AS path. NOP blackjacks are even stealthier, as the attacker is not in the AS path.

Malformed blackjacks This last category contains all blackjacks not covered by the other categories. They correspond to blackjacks where the AS path is malformed, mean-

ing all or some of the links between the ASes of the AS path do not exist and/or the origin AS is neither the attacker nor a legitimate AS. We assessed malformed blackjacks to be of little interest when looking for blackhole-based attacks, so we disregard them in the remainder of the paper.

4 Routing Security Deployments

To protect oneself against such attacks, several routing security mechanisms can be employed. Depending on the adoption rate of such mechanisms in the Internet, those attacks have a variable chance of success. In this section, we will consider five such deployments, each implementing those security mechanisms to different extents:

- **No security:** ASes neither use RPKI nor BGPsec.
- **RPKI (partial):** A subset of ASes uses the RPKI, but no AS is using BGPsec.
- **RPKI (full):** All ASes use the RPKI, but no AS is using BGPsec.
- **BGPsec (partial):** A subset of ASes uses both RPKI and BGPsec. The other ASes either use only RPKI or do not use any security mechanisms.
- **BGPsec (full):** All ASes use both RPKI and BGPsec.

It is important to keep in mind that although we consider multiple security deployments, BGPsec is not deployed at all and RPKI is only partially deployed. The RPKI covered around 5-6% of advertised prefixes in 2015 [26, 61, 19], and covers 13% of advertised prefixes today [39]. Although harder to measure, the deployment status of ROV has also been studied [44, 45] and shows that only a few ASes are currently performing ROV. This means the deployment corresponding the most to a real-life scenario is the ‘RPKI (partial)’ deployment.

Tables 1 and 2 summarize which blackjack attacks can work under those different security deployments, the former against exact prefix blackjacks and the latter against sub-prefix blackjacks. Each row of the table represents a security deployment scenario, and each column represents an attack. Thus, the intersection of a line and a column shows how a particular security deployment fares against a given attack:

-  : The security deployment is resistant to the attack.
-  : The security deployment is not resistant to the attack.
-  : The resistance of the security deployment to the attack is determined by other factors (network topology, where security is deployed, ...)

The next sections go over the different deployments, and describe the attacks possible in each context.

4.1 Fully deployed BGPsec

In this subsection, we consider a situation where *every AS has deployed, and uses, BGPsec and RPKI/ROV according to best practices* [31, 7, 30, 25, 37].

In this deployment, every AS can be assured that the AS path attribute is protected and legitimate in every advertisement they receive, and that the origin AS is authorized

Security Deployment	Type-0	Type-N	NOP	OP	OP-GRV
BGPsec (full)	■	■	■	□	□
BGPsec (partial)	▣	▣	▣	□	□
RPKI (full)	■	□	□	□	□
RPKI (partial)	▣	□	□	□	□
No security	□	□	□	□	□

Table 1: Security deployments against exact prefix blackjacks

Security Deployment	Type-0	Type-N	NOP	OP	OP-GRV
BGPsec (full)	■	■	■	■	■
BGPsec (partial)	▣	▣	▣	■	■
RPKI (full)	■	■	■	■	■
RPKI (partial)	▣	▣	▣	■	■
No security	□	□	□	■	■

Table 2: Security deployments against sub-prefix blackjacks

to announce the prefix. Since the attacker needs to send signed BGPsec advertisements for them to be considered by other ASes, he can only potentially perform either variations of On Path blackjacks. There can also be no sub-prefix blackjacks, since all ASes in this deployment can detect the sub-prefix via the RPKI.

4.2 Partially deployed BGPsec

We now consider a situation where *a subset of AS have deployed, and use, BGPsec and RPKI/ROV according to best practices [31, 7, 30, 25, 37]. The other ASes either use only RPKI/ROV [30, 25, 37] or do not use any security mechanisms.*

Depending on which ASes on the path of the advertisement from the attacker to the blackholer deployed which security mechanisms, multiple cases arise. If ASes on the path implement no security mechanisms, the case can be assimilated to a ‘No Security’ deployment. If at least one of the ASes on the path uses ROV, the case can be assimilated to a ‘RPKI (partial)’ deployment (see Subsection 4.4). Those two cases can also be assimilated in the case of sub-prefix blackjacks.

If at least one of the ASes on the path uses BGPsec and the RPKI/ROV, then the attacker can potentially make use of both On Path attacks, as well as possibly Type-0 blackjacks (see Subsection 4.4). The attacker can also potentially make use of Type-N and NOP blackjacks if he can perform downgrade attacks [33] on the ASes using BGPsec.

For sub-prefix blackjacks in this case, the attacker can potentially use Type-0 blackjacks (see Subsection 4.4). The attacker can also make use of Type-N and NOP sub-prefix blackjacks if he can perform downgrade attacks on the ASes using BGPsec, and the legitimate prefix covering the targeted sub-prefix is either not in the RPKI, or is loose. A prefix is loose “*when not all sub-prefixes of the maximum length allowed by the ROA are advertised in BGP*” [19] (e.g. a ROA allowing a prefix to be advertised up to /24, but the advertised prefix is a /20). The attacker cannot make use of On Path

	Prefix in RPKI		Prefix not in RPKI
	ROA is loose	ROA is not loose	
ROV	Type-N/NOP sub-prefix BJ	■	AS policy
no ROV	□	□	□

Table 3: Security detail of the ‘RPKI (partial)’ deployment against blackjacks

blackjacks in this case, since it would require a prior advertisement of the sub-prefix, which is not possible since we only have one attacker in our attack model.

4.3 Fully deployed RPKI

In this subsection, *every AS has deployed, and uses, RPKI and ROV according to best practices [30, 25, 37].*

Here, every AS can verify the association of the advertised prefix and the AS originating it, which means an attacker can potentially carry out all attacks except Type-0 blackjacks. For sub-prefix blackjacks, no attack is possible since all ASes in this deployment can detect the sub-prefix via the RPKI.

It is important to keep in mind that in a real scenario, it may still be possible to perform Type-0 blackjacks even if ROV is put in place, simply because of the order the router’s filter are applied [54]. Instead of discarding an ‘invalid’ route in case of a Type-0 blackjack, the router might accept the advertisement because blackholing takes precedence.

4.4 Partially deployed RPKI

In this subsection, *a subset of AS have deployed, and use, RPKI and ROV according to best practices [30, 25, 37].*

In this deployment, the attacks potentially usable by an attacker depend on three factors:

- The presence (or absence) of ROV at ASes on the path of the advertisement from the attacker to the blackholer.
- The presence (or absence) of the targeted prefix in the RPKI.
- If the prefix is in the RPKI, the fact that the ROA for the prefix is loose or not.

As you can see in Table 3, if ASes on the path of the advertisement do not enforce ROV, the case can be assimilated to a ‘No Security’ deployment.

Second, if the at least one AS on the path of the advertisement enforces ROV, and the prefix is not in the RPKI, it is up to the AS enforcing ROV to decide what to do (RPKI validation state = ‘unknown’). The AS can either diminish its preference of the route, or drop the route. In the former case, exact prefix blackjacks (of all types) will be possible as the AS classifies all routes to this prefix as ‘unknown’, even the one from the legitimate AS: blackjacks can win the BGP decision process. Sub-prefix blackjacks are also possible (except both On Path variations), and are not even penalized by a diminished preference, as they are more specific than the legitimate advertised prefix. All

in all, for prefixes not in the RPKI, an AS enforcing ROV and lowering preferences for ‘unknown’ route validity states behaves in the same way as an AS not enforcing ROV. In this case, possible attacks are the same as in the ‘No Security’ deployment. If the AS drops ‘unknown’ routes, those attacks are no longer possible, but in the current deployment state of RPKI, dropping ‘unknown’ routes would equate to dropping routes to 87% of the Internet, so for now, a compromise between reachability and security must be made.

Third, if the AS receiving the forged advertisement enforces ROV and the prefix is in the RPKI, two cases arise: either the ROA for the prefix is loose, or it is not. If the ROA is not loose, the deployment can be assimilated to a ‘RPKI (full)’ deployment. If the ROA is loose, in addition to attacks possible in the ‘RPKI (full)’ deployment, an attacker can also perform Type-N and NOP sub-prefix blackjacks within the range of maxLength, as the origin AS will match the asID in the ROA.

4.5 No Security

In this subsection, *ASes do not use any of the aforementioned security mechanisms*. If neither BGPsec nor RPKI and ROV are deployed, an attacker can perform all the attacks of the taxonomy. In the case of sub-prefix blackjacks, an attacker can use all the attacks except On Path blackjacks, since it would require a prior advertisement of the sub-prefix, which is not possible since we only have one attacker in our attack model.

5 Good Practices

We highlight two items having an influence on preventing attacks from the taxonomy:

- **Authorized origin:** The origin is authorized if the association between the origin AS and the prefix is ‘valid’ according to IRRs or the RPKI.
- **Valid path:** The path is considered ‘Valid’ if the AS path reflects the actual path the advertisement went through. This can be verified using BGPsec.

Even if an AS implements both RPKI and BGPsec, it is still vulnerable to both exact prefix On Path blackjacks, as well as possibly Type-0 exact and sub-prefix blackjacks depending on the state of the prefix in relation to the RPKI (see Table 3).

For an AS not to be vulnerable against Type-0 blackjacks, it needs help from third parties, (e.g. another AS registering its prefixes in the RPKI). However, an AS can protect itself against On Path attacks by adding constraints on advertisements it receives.

5.1 Additional Verification Rules

We suggest two verification steps to protect an AS against On Path blackjacks:

- **Legitimate peer:** The peer sending the blackhole advertisement is legitimate if the leftmost AS in the AS path is the ASN specified in the BGP OPEN message that created the session.

- **Direct connection:** The AS sending the blackhole advertisement is directly connected to the local AS. This can be verified by making sure there is only one AS in the AS path.

If an AS can make sure it has a direct connection to the AS sending the blackhole advertisement, it is then only vulnerable to Type-0 and 1-hop NOP blackjacks (e.g. the one in Figure 5) by definition. If this AS can also verify this peer is legitimate and authorized to advertise the prefix, then the AS is protected against all the attacks of the taxonomy without needing BGPsec.

It is worth keeping in mind that at this point, acknowledging the deployment state of RPKI and BGPsec, an AS peering through an IXP virtually has no protection against the attacks, as it must trust the IXP to verify the ‘Legitimate peer’ rule and the route server may not perform ROV¹.

5.2 Additional Good Blackholing Practices

In addition to the rules, other good practices can be put in place. Those good practices help to limit the possible damage caused by an inadvertent blackholing.

A filter for less specific blackholing advertisements The literature specifies that operators should accept blackholing advertisements up to /32 for IPv4, and /128 for IPv6, but does not specify a limit on prefixes which are less specific. We propose that operators reject blackholing advertisements if they are not specific enough, in order to avoid accidental blackholing of large IP blocks.

Acknowledging the distribution of blackholing prefix length [13], we advise to set it to /24, thus only accepting blackholing advertisements from /24 up to /32. This filter can be applied as both an inbound and outbound filter.

Concerning IPv6, observed IXPs put the limit at /19 [12, 16]. The literature does not have any specific information enabling us to determine a good limit for IPv6 blackholing prefix specificity, more research needs to be done.

An outbound filter for more specific blackholing advertisements When using blackholing across AS boundaries, an outbound filter should be set on eBGP peering sessions to deny all prefixes longer than the longest prefix expected to be announced, unless that prefix is tagged with a blackhole community. This does not help with accidental blackholing directly, but prevents an AS from advertising more specific prefixes inadvertently.

Considering some of these good practices might not be applicable depending on the situation, or can constrain the blackhole service too much, we propose using a BGPsec extension as a possible alternative to protect against attacks of the taxonomy.

5.3 A BGPsec Solution

In a full deployment of BGPsec and RPKI, only On Path attacks are still possible. Thus, the goal of integrating communities to BGPsec is to be able to attribute the changes made

¹ This might be changing as several IXPs now seem to implement ROV [3].

to communities to an AS. This attribution is crucial for blackholing, because it allows an AS to accept or reject a blackhole request based on the identity of the AS requesting the blackhole. A blackholing advertisement can then be analyzed, to determine the source of the request, and a decision can be made based on whether or not this AS has a right to blackhole this prefix. Moreover, given an unwanted blackholing event, those responsible for it can be held accountable.

We propose such an extension in [36]. With this extension, as we know which ASes introduced which communities, an AS could simply generate a table associating an AS to a set of prefixes this AS is authorized to blackhole. This table could be populated by RPKI/IRR data, but also manually with trusted peers, or other associations the operator deems relevant. Then, this AS could accept a blackhole request if the AS requesting the blackhole and the prefix in the advertisement matches an association in the table.

6 Related Work

Over the last years, efforts have been made towards characterizing usage and behavior of communities in the Internet. Donnet et al. proposed the first classification of BGP communities [15], and found that community usage increased from 2004 to 2007. The increased popularity of communities has since been established multiple times [13, 21, 22, 54]. Streibelt et al. also found that even though communities are typically relevant only between directly connected ASes, they seem to be propagated beyond, increasing the risks of attacks.

Streibelt et al. also demonstrated that attacks using BGP blackholing are not only possible in theory, but also in an experimental setup and in the wild [54]. In comparison to this paper, they only consider Type-0 blackjacks, so a possible area of future work is to test the other attacks of the taxonomy. Numerous efforts have also been made towards characterizing DDoS attacks, as well as the detection and mitigation techniques that can be used against them [48, 49]. Dietzel et al. study the shortcomings of blackholing, and propose Stellar, an advanced blackholing mechanism [14] which can perform fine-grained blackholing using extended communities as a signaling mechanism.

Finally, BGP hijacking has been studied extensively, to characterize them [5, 58, 59, 51, 52], to detect them [62, 60, 53, 52], or even to conduct further attacks [2, 55, 59, 4]. A possible area of future work is the adaptation of those detection techniques to blackhole-based attacks.

7 Conclusion

In this paper, we construct a taxonomy using blackholing as an attack vector, and assess the usability of those attacks in various security deployments. We also show those attacks have better reach and stealth than regular hijacks. Namely, blackholing takes precedence over AS relationships and AS path length, meaning a blackjack can affect more ASes than hijacks. By using the specific blackhole community value of a blackholer, an attacker can also drop traffic at ASes much further away than hijacks can. As the attacker is not the one dropping traffic and blackjacks may propagate far, blackjacks are stealthier than hijacks.

We also want to draw attention to the fact that since blackjacks make use of the blackholing service of an AS, making this blackholing information publicly available might not be a good idea without proper standardization and security.

Through attacks suited against the considered security mechanisms (RPKI and BGPsec), we highlight the poor state BGP security deployment is in, and suggest additional rules as well as good practices to protect against the attacks of our taxonomy. In a more general way, we want to emphasize the need for BGP community authentication, either through an extension to BGPsec or another mechanism.

As part of our future work, we want to test the attacks not already covered by Streibelt et al. [54] in a real world setting, to demonstrate those attacks can be carried out and present numerous advantages compared to regular hijacks. Another area to investigate is the existence and characteristics of ASes proposing blackholing services to perform blackhole-based attacks, much like open DNS resolvers can be used to carry out DDoS attacks. More work can also be done to adapt hijack detection techniques to blackhole-based attacks.

The feasibility and subtleties of blackjack attacks remain to be studied in a real-world setting. Since BGPsec has yet to be deployed, and there is still little experience with RPKI, those security mechanisms and their limitations can hardly be tested against at this time. Further research is needed to assess the severity of blackjack attacks in the wild, since actual configuration (e.g. blackholing precedence over other policies, community handling, RTBH provider policy, blackholing propagation, ...) might differ from expectations, and from AS to AS.

Acknowledgments

This project has been made possible in part by a grant from the Cisco University Research Program Fund, an advised fund of Silicon Valley Community Foundation.

References

1. Akamai: Memcached-fueled 1.3 Tbps attacks. <https://blogs.akamai.com/2018/03/memcached-fueled-13-tbps-attacks.html> (Mar 2018), [Online; accessed 29-April-2019]
2. Alex Pilosov and Tony Kapela: Stealing The Internet: An Internet-Scale Man In The Middle Attack. <https://www.defcon.org/images/defcon-16/dc16-presentations/defcon-16-pilosov-kapela.pdf> (Aug 2008), [Online; accessed 29-April-2019]
3. Andreas Reuter and Randy Bush and Ethan Katz-Bassett and Italo Cunha and Thomas C. Schmidt and Matthias Wählisch: Measuring Adoption of RPKI Route Origin Validation and Filtering. https://ripe76.ripe.net/presentations/63-rov_filtering_update.pdf (May 2018), [Online; accessed 29-April-2019]
4. Apostolaki, M., Zohar, A., Vanbever, L.: Hijacking bitcoin: Routing attacks on cryptocurrencies. In: 2017 IEEE Symposium on Security and Privacy (SP). pp. 375–392. IEEE (2017)
5. Ballani, H., Francis, P., Zhang, X.: A study of prefix hijacking and interception in the Internet. *ACM SIGCOMM Computer Communication Review* **37**(4), 265–276 (2007)
6. Brewster, T.: Cyber Attacks Strike Zimbabweans Around Controversial Election. <http://www.silicon.co.uk/workspace/zimbabwe-election-cyber-attacks-123938> (Aug 2013), [Online; accessed 29-April-2019]

7. Bush, R.: BGPsec Operational Considerations. BCP 211, RFC Editor (September 2017)
8. Chandra, R., Traina, P., Li, T.: BGP Communities Attribute. RFC 1997, RFC Editor (August 1996)
9. Cisco: Remotely Triggered Black Hole Filtering - Destination Based and Source Based. https://www.cisco.com/c/dam/en/us/products/collateral/security/ios-network-foundation-protection-nfp/prod_white_paper0900aecd80313fac.pdf (2005), [Online; accessed 29-April-2019]
10. Cohen, A., Gilad, Y., Herzberg, A., Schapira, M.: One hop for rpki, one giant leap for bgp security. In: Proceedings of the 14th ACM Workshop on Hot Topics in Networks. p. 10. ACM (2015)
11. Cohen, A., Gilad, Y., Herzberg, A., Schapira, M.: Jumpstarting bgp security with path-end validation. In: Proceedings of the 2016 ACM SIGCOMM Conference. pp. 342–355. ACM (2016)
12. DE-CIX: DE-CIX Blackholing Service. https://www.de-cix.net/_Resources/Persistent/4277e7d4867a78ae923c0f5b3b66d7ff6aeb61f8/DE-CIX-Blackholing-Service.pdf (Jul 2018), [Online; accessed 29-April-2019; Slide 3]
13. Dietzel, C., Feldmann, A., King, T.: Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In: Karagiannis, T., Dimitropoulos, X. (eds.) Passive and Active Measurement. pp. 319–332. Springer International Publishing, Cham (2016)
14. Dietzel, C., Smaragdakis, G., Wichtlhuber, M., Feldmann, A.: Stellar: network attack mitigation using advanced blackholing. In: Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies. pp. 152–164. ACM (2018)
15. Donnet, B., Bonaventure, O.: On BGP communities. ACM SIGCOMM Computer Communication Review **38**(2), 55–59 (2008)
16. France-IX: France-IX Blackholing Service. <https://www.franceix.net/fr/technical/blackholing/> (Jul 2018), [Online; accessed 29-April-2019]
17. Gao, L., Griffin, T.G., Rexford, J.: Inherently safe backup routing with BGP. In: Proceedings IEEE INFOCOM 2001. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society (Cat. No.01CH37213). vol. 1, pp. 547–556 vol.1. IEEE (April 2001). <https://doi.org/10.1109/INFCOM.2001.916777>
18. Gao, L., Rexford, J.: Stable Internet routing without global coordination. IEEE/ACM Transactions on Networking (TON) **9**(6), 681–692 (2001)
19. Gilad, Y., Cohen, A., Herzberg, A., Schapira, M., Shulman, H.: Are We There Yet? On RPKI’s Deployment and Security. IACR Cryptology ePrint Archive **2016**, 1010 (2016)
20. Gill, P., Schapira, M., Goldberg, S.: Let the market drive deployment: A strategy for transitioning to BGP security. ACM SIGCOMM computer communication review **41**(4), 14–25 (2011)
21. Giotsas, V., Dietzel, C., Smaragdakis, G., Feldmann, A., Berger, A., Aben, E.: Detecting peering infrastructure outages in the wild. In: Proceedings of the Conference of the ACM Special Interest Group on Data Communication. pp. 446–459. ACM (2017)
22. Giotsas, V., Smaragdakis, G., Dietzel, C., Richter, P., Feldmann, A., Berger, A.: Inferring BGP blackholing activity in the Internet. In: Proceedings of the 2017 Internet Measurement Conference. pp. 1–14. ACM (2017)
23. Greenberg, A.: Hacker Redirects Traffic From 19 Internet Providers to Steal Bitcoins. <https://www.wired.com/2014/08/isp-bitcoin-theft/> (Aug 2014), [Online; accessed 29-April-2019]
24. Heitz, J., Snijders, J., Patel, K., Bagdonas, I., Hilliard, N.: BGP Large Communities Attribute. RFC 8092, RFC Editor (February 2017)
25. Huston, G., Michaelson, G.: Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs). RFC 6483, RFC Editor (February 2012)

26. Iamartino, D., Pelsser, C., Bush, R.: Measuring bgp route origin registration and validation. In: International Conference on Passive and Active Network Measurement. pp. 28–40. Springer (2015)
27. Kandagatla, N.: Disgruntled ex-employees, DDoS attacks and the revenge of the nerds. <https://www.wittysparks.com/disgruntled-ex-employees-ddos-attacks-and-the-revenge-of-the-nerds/> (Nov 2017), [Online; accessed 29-April-2019]
28. King, T., Dietzel, C., Snijders, J., Doering, G., Hankins, G.: BLACKHOLE Community. RFC 7999, RFC Editor (October 2016)
29. Kumari, W., McPherson, D.: Remote Triggered Black Hole Filtering with Unicast Reverse Path Forwarding (uRPF). RFC 5635, RFC Editor (August 2009)
30. Lepinski, M., Kent, S.: An Infrastructure to Support Secure Internet Routing. RFC 6480, RFC Editor (February 2012), <http://www.rfc-editor.org/rfc/rfc6480.txt>, <http://www.rfc-editor.org/rfc/rfc6480.txt>
31. Lepinski, M., Sriram, K.: BGPsec Protocol Specification. RFC 8205, RFC Editor (September 2017)
32. Leyden, J.: US credit card firm fights DDoS attack. http://www.theregister.co.uk/2004/09/23/authorize_ddos_attack/ (Sep 2004), [Online; accessed 29-April-2019]
33. Lychev, R., Goldberg, S., Schapira, M.: BGP Security in Partial Deployment: Is the Juice Worth the Squeeze? SIGCOMM Comput. Commun. Rev. **43**(4), 171–182 (Aug 2013). <https://doi.org/10.1145/2534169.2486010>, <http://doi.acm.org/10.1145/2534169.2486010>
34. Madory, D.: BackConnects Suspicious BGP Hijacks. <https://dyn.com/blog/backconnects-suspicious-bgp-hijacks/> (Sep 2016), [Online; accessed 29-April-2019]
35. Madory, D.: Iran Leaks Censorship via BGP Hijacks. <https://dyn.com/blog/iran-leaks-censorship-via-bgp-hijacks/> (Jan 2017), [Online; accessed 29-April-2019]
36. Miller, L., Pelsser, C., Cateloin, S.: DDoS, BGP Leaks and Hijack Mitigation Techniques. https://loicmiller.com/documents/hijack_ddos_mitigation.pdf (Aug 2018), [Online; accessed 29-April-2019]
37. Mohapatra, P., Scudder, J., Ward, D., Bush, R., Austein, R.: BGP Prefix Origin Validation. RFC 6811, RFC Editor (January 2013), <http://www.rfc-editor.org/rfc/rfc6811.txt>, <http://www.rfc-editor.org/rfc/rfc6811.txt>
38. Morales, C.: NETSCOUT Arbor Confirms 1.7 Tbps DDoS Attack; The Terabit Attack Era Is Upon Us. <https://www.arbornetworks.com/blog/asert/netscout-arbor-confirms-1-7-tbps-ddos-attack-terabit-attack-era-upon-us/> (Mar 2018), [Online; accessed 29-April-2019]
39. National Institute of Standards and Technology: Global Prefix/Origin Validation using RPKI. <https://rpki-monitor.antd.nist.gov/> (Apr 2019), [Online; accessed 29-April-2019]
40. Newman, L.H.: The Botnet That Broke the Internet Isn't Going Away. <https://www.wired.com/2016/12/botnet-broke-internet-isnt-going-away/> (Sep 2016), [Online; accessed 29-April-2019]
41. Pras, A., Sperotto, A., Moura, G.C.M., Drago, I., Barbosa, R., Sadre, R., Schmidt, R., Hofstede, R.: Attacks by Anonymous WikiLeaks Proponents not Anonymous (2010)
42. Prince, M.: The DDoS That Almost Broke the Internet. <https://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet/> (Mar 2013), [Online; accessed 29-April-2019]
43. Rekhter, Y., Li, T., Hares, S.: A Border Gateway Protocol 4 (BGP-4). RFC 4271, RFC Editor (January 2006), <http://www.rfc-editor.org/rfc/rfc4271.txt>, <http://www.rfc-editor.org/rfc/rfc4271.txt>
44. Reuter, A., Bush, R., Cunha, I., Katz-Bassett, E., Schmidt, T.C., Wählisch, M.: Towards a rigorous methodology for measuring adoption of RPKI route validation and filtering. ACM SIGCOMM Computer Communication Review **48**(1), 19–27 (2018)

45. Reuter, Andreas and Bush, Randy and Cunha, Italo and Katz-Bassett, Ethan and Schmidt, Thomas C and Wählisch, Matthias: Measuring RPKI Route Origin Validation Deployment. <https://rov.rpki.net/> (Apr 2019), [Online; accessed 29-April-2019]
46. Reynolds, M., Turner, S., Kent, S.: A Profile for BGPsec Router Certificates, Certificate Revocation Lists, and Certification Requests. RFC 8209, RFC Editor (September 2017)
47. RIPE NCC: YouTube Hijacking: A RIPE NCC RIS case study. <https://www.ripe.net/publications/news/industry-developments/youtube-hijacking-a-ripe-ncc-ris-case-study> (Mar 2008), [Online; accessed 29-April-2019]
48. Rossow, C.: Amplification Hell: Revisiting Network Protocols for DDoS Abuse. In: NDSS (2014)
49. Ryba, F.J., Orlinski, M., Wählisch, M., Rossow, C., Schmidt, T.C.: Amplification and DRDoS Attack Defense—A Survey and New Perspectives. arXiv preprint arXiv:1505.07892 (2015)
50. Schlamp, J., Holz, R., Jacquemart, Q., Carle, G., Biersack, E.W.: HEAP: reliable assessment of BGP hijacking attacks. *IEEE Journal on Selected Areas in Communications* **34**(6), 1849–1861 (2016)
51. Sermpezis, P., Kotronis, V., Dainotti, A., Dimitropoulos, X.: A survey among network operators on BGP prefix hijacking. *ACM SIGCOMM Computer Communication Review* **48**(1), 64–69 (2018)
52. Sermpezis, P., Kotronis, V., Gigis, P., Dimitropoulos, X., Cicalese, D., King, A., Dainotti, A.: Artemis: Neutralizing bgp hijacking within a minute. *IEEE/ACM Transactions on Networking (TON)* **26**(6), 2471–2486 (2018)
53. Shi, X., Xiang, Y., Wang, Z., Yin, X., Wu, J.: Detecting prefix hijackings in the internet with argus. In: *Proceedings of the 2012 Internet Measurement Conference*. pp. 15–28. ACM (2012)
54. Streibelt, F., Lichtblau, F., Beverly, R., Feldmann, A., Pelsser, C., Smaragdakis, G., Bush, R.: BGP Communities: Even more Worms in the Routing Can. In: *Proceedings of the Internet Measurement Conference 2018*. pp. 279–292. ACM (2018)
55. Sun, Y., Edmundson, A., Vanbever, L., Li, O., Rexford, J., Chiang, M., Mittal, P.: {RAPTOR}: Routing Attacks on Privacy in Tor. In: *24th {USENIX} Security Symposium ({USENIX} Security 15)*. pp. 271–286 (2015)
56. Tomlinson, K.: Cyber battle rages on Internet after arrest of cyber crime suspects. <http://www.archersecuritygroup.com/cyber-battle-rages-internet-arrest-cyber-crime-suspects/> (Sep 2016), [Online; accessed 29-April-2019]
57. Turk, D.: Configuring BGP to Block Denial-of-Service Attacks. RFC 3882, RFC Editor (September 2004)
58. Vervier, P.A., Jacquemart, Q., Schlamp, J., Thonnard, O., Carle, G., Urvoy-Keller, G., Biersack, E., Dacier, M.: Malicious BGP hijacks: appearances can be deceiving. In: *2014 IEEE International Conference on Communications (ICC)*. pp. 884–889. IEEE (2014)
59. Vervier, P.A., Thonnard, O., Dacier, M.: Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks. In: NDSS (2015)
60. Wählisch, M., Maennel, O., Schmidt, T.C.: Towards detecting BGP route hijacking using the RPKI. In: *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*. pp. 103–104. Citeseer (2012)
61. Wählisch, M., Schmidt, R., Schmidt, T.C., Maennel, O., Uhlig, S., Tyson, G.: RiPKI: The tragic story of RPKI deployment in the Web ecosystem. In: *Proceedings of the 14th ACM Workshop on Hot Topics in Networks*. p. 11. ACM (2015)
62. Zheng, C., Ji, L., Pei, D., Wang, J., Francis, P.: A light-weight distributed scheme for detecting IP prefix hijacks in real-time. In: *ACM SIGCOMM Computer Communication Review*. vol. 37, pp. 277–288. ACM (2007)