Exploiting Vulnerabilities at IXP Route Servers to Perform Stealth BGP Hijacks

Gabby Rimlinger
LIP6, Sorbonne Université, LINCS
Paris, France

Joaquim Pereira
University of Waikato
Hamilton, New Zealand

Matthieu Gouel
nxthdr
Annecy, France

Olivier Fourmaux
LIP6, Sorbonne Université
Paris, France

Timur Friedman
LIP6, Sorbonne Université, LINCS
Paris, France

Pedro Macros
FURG
Rio Grande, Brazil

Ronaldo A. Ferreira
UFMS
Campo Grande, Brazil

Cristel Pelsser
UCLouvain
Louvain-la-Neuve, Belgium

Kevin Vermeulen
LIX, CNRS, École Polytechnique
Palaiseau, France

ABSTRACT

Internet Exchange Points (IXPs) are critical Internet infrastructure that interconnect tens of thousands of Autonomous Systems (ASes). To support scalable multilateral route exchange and fine-grained routing control, IXPs provide services such as route servers for scalable route dissemination and BGP communities for selective advertisement. Route servers enable both scalability and expressive routing policies, but some of their design choices can be exploited by a malicious actor. In this paper, we identify two of them: route-server path-hiding mitigation and the deployment of multiple independent route servers. Combined with well-known hijack techniques, these design choices allow an attacker to make multiple routes to the same prefix co-exist at the IXP, and strategically disseminate malicious routes to different subsets of peers, to increase their attack surface and the number of potentially vulnerable prefixes. We validate the feasibility of our attacks across three large IXPs, and show that these attacks increase the number of polluted ASes by 28% to 366%, and the number of vulnerable prefixes by 41% to 61%, depending on the IXP, compared to prior work. Moreover, we show that the IXP environment makes it easier to perform interception attacks than in other settings and allows such attacks to be invisible to public BGP collector peers. We also propose a novel data-plane detection technique based on the Layer-2 IXP architecture. Finally, drawing on discussions with IXP operators, we provide practical recommendations to improve route security and visibility at IXPs.

CCS CONCEPTS

• **Networks → Network security**.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

ACM ISBN 978-1-4503-XXXX-X/2018/06…$15.00
https://doi.org/XXXXXXX.XXXXXXX

KEYWORDS

IXP security, BGP

ACM Reference Format:
Gabby Rimlinger, Joaquim Pereira, Matthieu Gouel, Olivier Fourmaux, Timur Friedman, Pedro Macros, Ronaldo A. Ferreira, Cristel Pelsser, and Kevin Vermeulen. 2018. Exploiting Vulnerabilities at IXP Route Servers to Perform Stealth BGP Hijacks. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX).* ACM, New York, NY, USA, 16 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

Internet exchange points (IXPs) play a major role in today's Internet routing by enabling Autonomous Systems (ASes) to exchange large volumes of traffic efficiently and at low cost [7]. ASes exchange routes using the Border Gateway Protocol (BGP). To support scalable multilateral peering, IXPs commonly deploy route servers, which are specialized BGP speakers that collect routes from participating networks and redistribute them to other members, thereby eliminating the need for a full mesh of bilateral sessions. Beyond this basic functionality, modern IXPs provide increasingly sophisticated features to deal with the growing number and diversity of participants. For instance, IXPs often deploy multiple route servers for scalability and redundancy, support fine-grained routing control through BGP communities, and maintain per-client Routing Information Bases (RIBs) to allow more complex route filtering and distribution.

While these advanced features improve flexibility and scalability, they may create opportunities for malicious actors to exploit BGP vulnerabilities, such as prefix hijacking, in which attackers announce routes to IP prefixes they do not own [74, 81, 82]. The goal of such hijacks is generally to attract traffic destined to the victim network and allow the attacker to disrupt services (denial-of-service), eavesdrop on communications (traffic interception), or manipulate traffic flows for economic or strategic advantage (e.g., traffic redirection or blackholing) [18, 27, 51].

In this work, we identify two IXP route server design choices (or, more precisely, one design choice and one deployment choice, called both design choices in the rest of the paper), that an attacker can exploit. Our work fundamentally differs from prior efforts in

that it does not introduce new individual BGP hijack primitives such as subprefix announcements [12], AS-path manipulation [53], or BGP communities [13]. Instead, it shows how the combination of these two IXP route server design choices, with well-known BGP hijacks primitives, can lead to attacks that increase the attack surface and the number of vulnerable prefixes announced at the IXPs compared to prior work.

The first identified route-server exploitable design choice arises from per-client path distribution, an IXP route server feature designed to mitigate a dysfunction called "path hiding" [43] (§2.4.1). Path hiding occurs when a route server receives multiple routes to the same prefix but forwards only the best route to a subset of participants while excluding the others (e.g., when the best route carries excludes BGP communities that prevent its announcement to some participants). In this case, the excluded participants receive no route to the prefix, even though alternative routes exist at the IXP.

To mitigate path hiding, route servers maintain per-client RIBs and perform best-path selection only after per-client filtering, so paths excluded for a client are not part of its best-path selection. An attacker can exploit this behavior by announcing a route to excluded participants while withholding it from those receiving the legitimate route or from peers connected to route collectors to make the attack invisible to public collector peers. Path-hiding mitigation [43] is a must-have feature for IXPs: a major European IXP reported to us, in a private communication, that several IXP participants requested this capability after observing that a transit provider announced routes that were selected as best paths by the route server but selectively exported them only to specific participants, which prevented the other participants from learning alternative paths.

The second exploitable design choice comes from the deployment of multiple route servers, which is a common practice used by IXPs to ensure resilience. When these route servers are independent, i.e., do not share routes between each other, this allows an attacker to exploit this separation to announce distinct malicious routes (one per route server) to different subsets of IXP members that are vulnerable to different types of BGP hijacks, in order to increase the attack surface.

All of the attacks we describe in this paper evade control-plane detection mechanisms based on public BGP collector peers [36, 75, 76], by using BGP communities offered by route servers to exclude ASes that would forward the malicious routes to a public collector. Also, to avoid data-plane monitoring systems based on connectivity [84], the attacker can simply forward the traffic back to the victim using the Layer-2 architecture of IXPs. Finally, in a setup where the legitimate prefix is announced from a single location at the IXP, these attacks are also robust to latency-based techniques [73], as traffic redirection within an IXP might incur negligible delays.

We validate the feasibility and evaluate the impact of these attacks through experiments at three major IXPs, where we perform controlled, ethical attacks, and measure their attack surface. These attacks can potentially increase the number of polluted ASes, i.e., that adopt a malicious route, by 29%, 28%, and 366%, depending on the IXP, compared to prior attacks, affecting thousands of ASes at all the tested IXPs. Moreover, we find that 82%, 73%, and 90% of the prefixes announced at the three IXPs are potentially vulnerable to at least one of our attacks, representing an absolute increase of 60%, 49% and 61% compared to prior attacks. We also provide new insights into how route visibility, including looking glass access, influences attack detectability.

Beyond demonstrating the feasibility of these attacks, we provide operational insights into the practicality of launching such attacks. We show that an attacker can obtain the required capabilities without privileged access by joining an IXP as a regular participant, configuring routing policies, and deploying infrastructure (e.g., a VM) within the IXP.

Finally, we propose countermeasures to mitigate these attacks, including Layer-2 monitoring mechanisms and operational recommendations for IXP operators. We discussed the feasibility and effectiveness of these defenses with IXP operators to ensure that they are practical steps to reduce the attack surface introduced by modern IXP features.

2 BACKGROUND

We provide essential background on the BGP protocol, explain how routes are exchanged at IXPs via route servers, some security measures and design choices taken by IXPs at their route servers. Among the latter are the path-hiding phenomenon and its mitigations, as well as why IXPs deploy multiple route servers.

2.1 BGP and Path Preference

BGP is the protocol that allows ASes to exchange routes across the Internet. An AS advertises to its neighbors that it has a route to a prefix by making route announcements. The neighbor of an AS can be classified as *peer*, *customer*, or *provider* based on their commercial relationships [30]. Although more complex arrangements exist [30], these three suffice for our purposes. Under widely adopted stability policies, an AS propagates routes learned from its customers to all neighbors, whereas routes learned from peers or providers are propagated only to its customers [31]. The *customer cone* of an AS is the set of ASes reachable by following only provider-to-customer links.

In the context of this paper, to understand the attacks, readers need only to understand the basic route preference mechanism IP and BGP use when an AS receives multiple routes to the same destination. First, a more specific prefix takes precedence over a less specific one (e.g., a /24 prefix is preferred over a /23), as IP uses the longest-matching prefix for a destination IP address. Second, when multiple routes to the same prefix are received, BGP selects a single best route by sequentially applying its tie-breaking rules in the following order: highest 'local preference', shortest AS-path length, lowest 'origin type', etc. [65]. Some attacks in this paper exploit the AS-path length tie-breaker to increase route preference: when an AS receives both a malicious and a legitimate route to a prefix, it selects the malicious one if it has a shorter AS path and all preceding tie-breakers are equal. A common scenario that enables attackers to announce a shorter AS path is when a legitimate AS performs AS-path prepending, a technique typically used for traffic engineering in which an AS deliberately repeats its AS number multiple times to artificially increase the path length and make its route less attractive.

2.2 IXP Route Servers

In traditional BGP, an AS establishes individual BGP sessions with each neighbor to exchange routes, a practice called bilateral peering. At IXPs, however, maintaining individual connections with many peers[1] becomes prohibitively complex. IXPs therefore offer members the option to establish a single BGP session with a central entity, called a route server (RS), which collects and distributes routes among peers. This approach is called multilateral peering [33]. Beyond simplifying peering, route servers typically provide mechanisms that allow IXP members to control how their routes are propagated. In particular, an AS can instruct the route server to advertise a route only to selected peers or to exclude specific peers from receiving it. This control is implemented through BGP communities, which are tags that an AS can add to a route to carry information or influence how the route is handled and propagated.

Although these communities are not standardized, many route servers adopt consistent semantics: `0:ASN` excludes a specific AS (identified by its ASN – Autonomous System Number), `RS-ASN:ASN` includes a particular AS, `0:RS-ASN` excludes all ASes, and `RS-ASN:RS-ASN` includes all ASes, where `RS-ASN` denotes the IXP route server's ASN. A peer can for instance perform selective advertisement by combining `0:RS-ASN` with explicit include communities. With peer-to-peer relationships at IXPs, these communities enable fine-grained control over route dissemination, since stability-preserving policies require ASes to propagate peer-learned routes only to their customers [31].

2.3 Security Measures at IXP Route Servers

IXPs implement several security measures to prevent the propagation of incorrect or malicious routes. To understand current practices, we manually reviewed the public documentation of the route servers of ten major IXPs across five continents (including the two largest by number of peers, according to IXPtracker [42]) [2, 22, 24, 25, 35, 40, 41, 44, 48, 52, 55]. Route servers accept routes with only valid prefixes (i.e., non-bogon, non-martian, and non-private) [19, 66] from public ASNs which comply with prefix-length limits (≤ 24), have a valid next-hop corresponding to the peer's interface and AS on the IXP LAN, avoid large transit ASes in the path, respect AS-path length limit, and remain within the allowed prefix quotas. In addition, IXPs perform verification from two external databases: IRR and RPKI.

IRR-based filtering: The Internet Routing Registry (IRR) [38] is a database where ASes register routing objects (e.g., AS-SETs) that specify which prefixes they and their customers are authorized to announce. IXPs use this information to validate whether a peer is allowed to originate a given prefix. Some IXPs automatically drop an announcement that is IRR invalid, while others just tag the route with a BGP community indicating its invalid status, partly because IRR is not sanitized [23]. However, this mechanism can be bypassed, as an attacker can falsify the AS-SET to include the ASes that legitimately own the targeted prefixes. The three tested IXPs where we tested our attacks trust IRR data, allowing us to run our attacks (§§ 3.1 and 5.1).

[1]We use the terms peer, member, and participant interchangeably to refer to an AS connected to an IXP.

RPKI/ROV: The Resource Public Key Infrastructure (RPKI) publishes cryptographically signed Route Origin Authorizations (ROAs) that specify which ASes are authorized to originate a given prefix, with an eventual maximum prefix length (maxLength parameter). Using Route Origin Validation (ROV), ASes validate received routes by checking whether the origin AS and the advertised prefix are consistent with the corresponding ROA. Routes that satisfy these checks are classified as *valid*. All our attacks respect the origin AS criteria, and only the ones based on subprefix hijacks can be prevented by the ROA maxLength prefix parameter (or the absence of it, which is equivalent to setting a maxLength equal to the prefix length).

2.4 Route Server Exploitable Design Choices

2.4.1 Path Hiding. Path hiding is a route server (RS) dysfunction that causes an IXP to not transmit a viable path to a peer that ought to receive it. It arises from the conflict between two aspects of an RS's role in filtering the paths it receives. On the one hand, it applies the BGP decision process and selects, for each destination prefix, a single best path to communicate to its peers. On the other hand, it enforces its peers' selective communication policies, allowing each of them to specify which others should or should not receive their path announcements. When the best path comes from a peer that has excluded another from receiving it, the RS honors the exclusion and does not send that path to the excluded peer. However, the RS has already ruled out the other candidate paths during its best-path selection. Even if one of the unselected paths originates from a peer willing to share it with the excluded peer, that otherwise viable path is no longer eligible for forwarding. As a result, the RS does not announce any route to the destination prefix to the excluded peer. The viable path has been therefore "hidden" by the RS.

The path-hiding problem was identified over a decade ago [67], and RFC 7947 [43] proposed three mitigation approaches. In the Multiple Route Server RIBs approach, the RS maintains a separate routing information base (RIB) for each peer and executes the BGP decision process based only on the paths that the peer is eligible to receive. The Diverse BGP Paths approach applies the observation [64] that multiple BGP sessions can be used to distribute multiple viable paths, beyond the best path for a prefix. Finally, the BGP ADD-PATH approach uses BGP's additional path extension [83] to send multiple paths for a prefix via a single session. All three approaches solve the path-hiding problem through a per-client path distribution mechanism, which allows multiple routes to the same prefix to co-exist in the IXP, creating opportunities for an attacker to announce a malicious route that does not have to compete with legitimate ones (§4.1).

2.4.2 Deployment of Multiple Route Servers. To ensure resilience, IXPs typically deploy at least two route server instances, so that if one route server becomes unavailable due to maintenance or failure, members can continue exchanging routes through the second instance, preserving the connectivity. IXP members typically maintain BGP sessions with all available route servers simultaneously, both to ensure redundancy and maintain a complete view of the routes exchanged at route servers. The deployment of multiple route servers is a common IXP design choice, but when the route

servers are independent, it creates an opportunity for an attacker to announce several malicious routes, one to each route server (§4.2).

3 ATTACK OVERVIEW

3.1 Threat Model

We consider an adversary that operates as a regular member of an IXP. The attacker controls one AS, can establish BGP sessions with the IXP route server, and can announce arbitrary routes to allowed prefixes, tagged or not with BGP communities, subject to the same operational constraints as any other participant. We detail how we became a regular member of the three tested IXPs and the operational steps to perform successful attacks in Section 5.1, without requiring any privileged access.

3.2 Attack Goals and Success Conditions

The attacker's goal is to perform an interception attack, i.e., intercept traffic destined to a victim's prefix while avoiding being detected by detection systems that rely on either control-plane visibility through public BGP collectors or data-plane signals such as connectivity disruption or latency inflation. Consequently, if malicious announcements are not visible to collectors and the victim's connectivity and performance remain unaffected, detection becomes significantly harder. Our attacks are designed to satisfy these conditions (§4).

The attacker also wants to maximize its attack surface [53], or run a targeted interception [13]. We show that both are actually possible at IXPs.

The specific vulnerabilities opened by the route server design choices and possible announcements used to achieve these goals are described in Section 4.

4 ATTACKS AT IXP ROUTE SERVERS

We present how an attacker can exploit the two identified route server design choices. The first one is path-hiding mitigation, which creates a "per-client" route distribution at the route server. This allows an attacker to hijack routes by selectively advertising a malicious route to (i) peers that do not receive the victim's announcements due to BGP exclude communities and (ii) peers that receive the victim's routes with long AS paths, while (iii) simultaneously excluding peers that forward routes to BGP collectors. We say a peer forwards routes to a BGP collector if the peer or one of its customers propagates the route to a BGP collector upon receiving it. Because the legitimate route still propagates normally toward the peers that forward routes to BGP collectors, visibility of the legitimate route is unaffected, while the malicious route remains invisible to public collector peers and, consequently, does not trigger alarms.

The second one is the deployment of multiple independent route servers, which allows an attacker to send as many malicious routes per prefix as there are route servers, to amplify its attack surface. This amplification is particularly effective because IXP peers are not uniformly vulnerable to all BGP hijacking techniques, especially with respect to how their routes are observable to BGP collectors. Both attacks comply with all the security measures adopted by IXPs (§2.2).

Table 1 gives an overview of the conditions for an attacker to perform the different attacks, according to the victim's route properties, the attacker announcement, the required route-server behavior, the ROA properties on the victim's prefix, the evasion of public collector peers, and how to preserve connectivity. Note that remote peering does not fundamentally change the attacks, as the remote peers are also connected remotely to the route server, with the same application of the route-server policies.

4.1 Exploiting the Path-hiding Mitigation

We present two possible attacks exploiting path-hiding mitigation using BGP exclude communities and the shorter AS-path length techniques, called *path-hiding excluded peers* and *path-hiding shorter AS path* attacks.

4.1.1 Path-hiding Excluded Peers. In this attack, illustrated in Figure 1a, the fact that the victim V has availed itself of the legitimate option to exclude some ASes at a route server creates a vulnerability that attacker A can exploit.

Here, the route announced by V is accompanied by the exclude community 0:4, and so the route server only propagates it to AS 3, and the attacker A. AS 9 learns it indirectly as a customer of AS 3, as does route collector X, which has peered with AS 9. The same effect would be obtained by V specifying, through an include community, that only ASes 3 and A should receive the route.

To obtain visibility on V's route with exclude communities, A can use the route-server looking glass or its own BGP session with the route server. However, the latter cannot be used to observe V's routes with include communities (unless A is included), or routes where A is itself excluded.

In this attack, A announces the same prefix to the route server, but with AS path [A-V] so that it will intercept traffic destined towards V. To conceal this fact from the ASes that have received the legitimate route, it excludes them, in this example, by employing exclude communities 0:3 and 0:V. The attacker's announcement is RPKI-valid because the origin AS is V, the legitimate origin.

Assuming the route server implements path-hiding mitigation, both routes coexist at the IXP but are propagated to disjoint sets of peers. V's route reaches ASes A, 3, and 9, and route collector X, while A's route reaches AS 4 and its customers, ASes 7 and 8. Note that AS 4 does not pass A's route to AS 5, and so the route does not propagate to AS 6 or route collector Y. This behavior is consistent with valley-free routing [31], under which AS 4 will not export to a provider (AS 5) a route that it has learned from a peer (AS A).

The attack is complete when A takes steps to evade data-plane detection. Attacker A forwards to V any packets that it has intercepted that are destined for the prefix that V has advertised. It does so by rewriting the Layer-2 headers to use V's MAC address as the destination and its own MAC address as the source.

Latency-based detection techniques' efficiency depends on the difference in latency that AS 4 (or any AS susceptible to adopting the malicious route) experiences to reach V's prefix before and after the attack. Before the attack, this latency depends on which path AS 4 uses to reach V - it may route through the IXP or via a different path entirely. After the attack, AS 4 adopts the malicious route and forwards traffic to the attacker's VM at the IXP, so the relevant latency becomes that between AS 4 and the attacker's

Table 1: Conditions to run each attack.

	Path-hiding excluded peers (§4.1.1)	Path-hiding shorter AS path (§4.1.2)	Multiple independent RS (§4.2)
Required victim route properties	Excludes ≥ 1 AS (exclude/include communities)	AS-path length > 2	Vulnerable to ≥ 2 different types of attacks (e.g., PH excluded peers, PH shorter AS path, subprefix ± NO_EXPORT [12])
Attacker announcement	Same prefix to excluded peers that are not collector-forwarding peers.	Same prefix with a shorter AS path to peers that are not collector-forwarding peers	One malicious announcement per RS affecting different sets of ASes
Required RS behavior	Implements path-hiding mitigation	Implements path-hiding mitigation	Deployment of ≥ 2 independent route servers
ROA condition on the victim's prefix	None	None	No ROA or ROA max-length > announced prefix length, to conduct a subprefix attack
Public collector visibility condition	Exclude collector-forwarding peers		
Interception condition	The attacker forwards the traffic back to the victim by rewriting the Layer-2 MAC addresses: Source to its own MAC address; Destination to the one of the victim if the victim is directly connected to the IXP, or to the one of its provider if not.		

machine. Note that AS 4 is not necessarily directly connected to the IXP, making this latency potentially non-negligible. To assess the risk of detection, the attacker should therefore estimate: (1) the latency between AS 4 and V's prefix before the attack, for instance by pinging V's prefix from a vantage point in AS 4; and (2) the latency between AS 4 and the attacker's VM at the IXP, for instance by pinging AS 4 from the attacker's machine. If (1) is higher than (2), AS 4 was reaching V via a higher-latency path than through the IXP. The attack actually reduces the latency AS 4 experiences, and the attacker can add artificial delay to minimize any observable difference. If instead (2) is higher than (1), the attack increases AS 4's latency to V, which may trigger detection; the attacker can then decide whether to proceed based on the magnitude of the difference. Finally, if V's prefix is only announced at the IXP, and if AS 4 was accepting the route, then AS 4 was already routing to AS V through the IXP before the attack, and our attack incurs no latency increase.

Note that although Figure 1a shows ASes 4, 7, and 8 adopting the attacker's route, they will not necessarily do so if they receive alternative routes to the same prefix, e.g., from other providers or peers not shown in the figure. Attacker A can seek to compete with any alternative routes by advertising a shorter path (e.g., [A-V]), though shorter paths may increase visibility in the route server's looking glass (§5.2.5).

If path-hiding mitigation were not implemented, the route server would select either the legitimate or the malicious route: in the first case, the attack fails; in the second, legitimate peers receive no route, potentially triggering alarms from those peers due to loss of connectivity to the victim. This latter scenario is equivalent to the known attack using BGP exclude communities [13]. Our attack instead expands the attack surface by reaching peers excluded from the legitimate route without disrupting connectivity for those that still receive it or its visibility at BGP collectors.

Requirements: The route server must implement path-hiding mitigation, and the victim must announce a route that excludes at least one AS.

4.1.2 Path-hiding Shorter AS Path. In this attack, illustrated in Figure 1b, V suffers from a vulnerability when the prefix that it announces has to go through at least two ASes before reaching the route server, allowing attacker A to announce a shorter path. This attack exploits the fact that AS-path length is a route selection criterion.

In this example, the victim V announces a prefix to its provider AS 1, which propagates it to AS 2, and subsequently to the IXP route server. At the route server, all members receive the route with an AS path of length 3: [2, 1, V]. To launch its attack, the attacker A announces the same prefix to the route server with a shorter AS path [A-V]. So as to avoid alerting the victim's provider, AS 2, or making the attack visible through public route collector X, A employs exclude communities 0:3 and 0:2. As a result, AS 4 and its customers receive both the legitimate route and the attacker's route. Because the attacker's AS path is shorter, those ASes will select it whenever they use AS-path length as their BGP tie breaker. This allows A to attract and intercept their traffic destined towards the prefix that V had advertised.

Path-hiding mitigation is required for the attack to be invisible to public collector peers. Without it, the route server would select the attacker's route as the single best path and suppress the victim's route. Because A's announcement contains the exclude community 0:3, AS 3 would receive no route to the prefix advertised by V and thus lose connectivity to that block of addresses. This anomaly would be detectable, as the route collector X would receive no route to that prefix, even though AS 2 had exported the legitimate route to the route server.

The attacker can apply the same methodology as described for the first attack to evade data-plane detection techniques.

Note that this attack is also possible if V is directly present at the IXP provided that V employs AS-path prepending, which similarly allows A to offer a shorter path.

Requirements: The route server must implement path-hiding mitigation, and the legitimate route to the victim prefix must have an AS path longer than two hops.

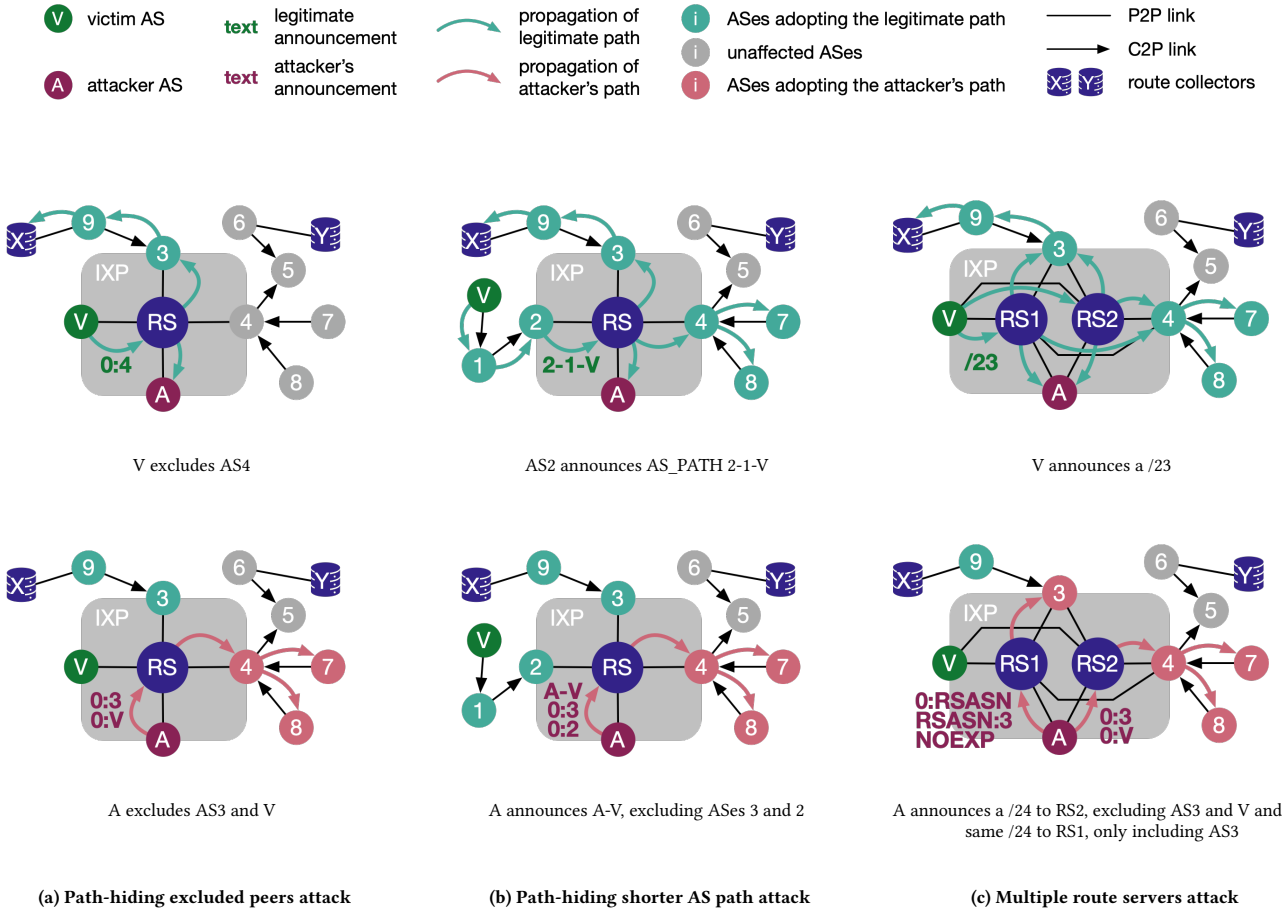

Figure 1: Three interception attacks at an IXP: the legitimate announcement (top) followed by attacker A's announcement (bottom). The attacker uses path-hiding mitigation at the route server to make routes to the same prefix coexist, combined with routes that exclude peers (left) or have long AS paths (middle), and can exploit the deployment of multiple route servers to send several malicious routes and increase the attack surface (right). To perform the interception, the attacker can re-route the intercepted traffic using the L2 architecture of the IXP.

4.2 Exploiting the Deployment of Multiple Independent Route Servers

We describe how an attacker can exploit the presence of multiple independent route servers to announce one route to a subset of peers that do not forward the route to BGP collectors, and another one to those that would forward it. We call this attack a *multiple route server attack*. The attack builds on the subprefix + NO_EXPORT community attack of Birge-Lee et al. [12], which we adapt to the IXP setting (§4.2.1). On its own, this adapted attack already allows an attacker to attract traffic from ASes without being detected by BGP collectors, even when those ASes have customers connected to a BGP collector. We then show (§4.2.2) how combining it with a second, complementary announcement sent through a different route server further expands the set of ASes whose hijacked traffic is undetected.

4.2.1 Adapting the Birge-Lee et al. Attack for IXPs. If an attacker announces a subprefix to a transit AS and tags it with the NO_EXPORT community, the transit AS's customers will not see the route and therefore will not forward it to a BGP collector [12]. Because these customers still receive the legitimate route to the covering prefix, they send their traffic to their provider. The provider, which adopts the malicious route to the subprefix, then forwards the traffic to the attacker. For this attack to succeed, the victim's prefix must **not** be covered by a route origin authorization (ROA) whose maxLength equals the prefix length. In other words, the prefix must either not be protected by ROA or covered by an ROA that allows the announcement of subprefixes, i.e., an ROA with maxLength greater than the prefix length.

Adapting this attack to an IXP requires identifying which peers support the NO_EXPORT community. Birge-Lee et al. identified supporting ASes by manually inspecting the documentation of four tier-1 networks. At an IXP, however, there can be hundreds of peers, so we instead design a measurement-based approach. Our approach alternately announces a route with and without the NO_EXPORT community while sending pings to IP addresses behind each IXP

peer. For a given peer, if a sufficient fraction of its clients lose connectivity or change their provider when the community is set, we conclude that the peer supports NO_EXPORT. We describe the experimental setup in more detail later (§5.2.2). The technique does not need to be perfect: an attacker can always add or remove a peer from the announcement on a test prefix if unsure whether that peer forwards the malicious route to a BGP collector.

We close with two remarks. First, this attack does not require that path-hiding mitigation (§4.1) be employed at the IXP since, from the route server's point of view, the subprefix and the victim's prefix are two distinct prefixes. Second, an attacker can also carry out a subprefix attack without using the NO_EXPORT community by announcing the route only to peers that do not forward it to BGP collectors.

4.2.2 Combining Announcements across Route Servers. An attacker can exploit the deployment of multiple independent route servers by sending a different announcement to each one, as illustrated in Figure 1c: (1) through one route server, the attacker sends the announcement to the peers that *do not* forward routes to a BGP collector. This announcement can for instance use a standard subprefix attack without the NO_EXPORT community, as in Figure 1c, but could also be one of the attacks based on path hiding described above; (2) through the other route server, the attacker sends a subprefix announcement tagged with the NO_EXPORT community to the subset of peers that *do* forward routes to a BGP collector and support NO_EXPORT.

In the example of Figure 1c, the victim V announces a /23 prefix, with a maxLength parameter of 24. The attacker A announces the subprefix without the NO_EXPORT community to AS 4 through one route server, and the same subprefix with the NO_EXPORT community to AS 3 through the other. AS 3, AS 4 and their customers are all routed to A, yet the malicious route does not appear at any BGP collector.

The data-plane evasion technique is the same as in the attacks described previously.

Requirements: The IXP must operate multiple independent route servers that do not share routes between each other. Moreover, if one of the malicious announcements is a subprefix announcement, the victim's prefix must **not** be covered by a route origin authorization (ROA) with no maxLength or whose maxLength equals the prefix length.

5 EVALUATION

Overview: We first show that obtaining the setup required to launch our attacks does not require any privileged or special access, confirming their practical feasibility (§5.1). We then demonstrate our attacks with our experimental setup, where we play both the victim and the attacker (§5.2). Compared to prior work, our attacks increase the attack surface, i.e., the number of ASes adopting the malicious route, by 29%, 28%, and 366% (§5.3) across the tested IXPs. Our attacks represent an absolute increase of the number of potentially vulnerable prefixes by 60%, 49% and 61% compared to the Subprefix + NO_EXPORT attack [12] and by 52%, 41% and 61% compared to the AS-path poisoning attack [53] (§5.4). All our attacks are interception attacks that do not break connectivity and evade BGP collectors.

5.1 Experimental Setup

To demonstrate the feasibility of our attacks in the real world, we conducted experiments at three major IXPs, two in Europe and one in the Americas, while designing all tests to avoid disrupting normal IXP operations. At each site, we deployed the same architecture: two VMs connected directly to the IXP's LAN, one acting as the victim (announcing the hijacked routes) and the other as the attacker (announcing the malicious routes). Since we own two ASNs, each VM operates under a distinct ASN and maintains BGP sessions with the IXP's route servers. Each VM has two interfaces: one connected to the LAN of the IXP, and the other one connected to a transit IXP provider to allow us to connect to the VMs and send probes to prefixes that are not reachable via the IXP.

We also own a /23 prefix registered to the victim AS for experimentation. We configured an ROA for the /23 prefix with maxLength 24 to permit valid subprefix announcements, in particular for the multiple route servers attack. The IXPs are anonymized as IXP-EU-1 and IXP-EU-2 for the IXPs in Europe and IXP-AM for the IXP in the Americas. The three IXPs operate two route servers, and implement path-hiding mitigation techniques, in particular per-client RIB or ADD-PATH.

Note that, in practice, an attacker would need only a single VM for the attacker, which makes the real-world barrier to entry even lower. Our experiments were run during the period of January-April 2026.

IXP-EU-1: We rented the VMs from an intermediate provider. For provisioning the first VM, we provided standard administrative information: a legal entity, ASN, AS-SET, NOC contact details, organization type (in our case, non-profit), organization name, contact person, and the maximum number of IPv4/IPv6 prefixes we were planning to announce. For the second VM, we simply requested an additional instance from our intermediate provider and supplied the same information as for the first VM. We also added the AS of this second VM, configured as the victim, to the AS-SET of the first one, configured as the attacker. The IXP automatically retrieved IRR data to validate the prefixes we were allowed to announce and required only that the AS-SET be "up to date", without further verification.

IXP-EU-2: We requested the same setup as at IXP-EU-1 through a direct contact at IXP-EU-2. Because this contact was a privileged partner, we informed them that we would conduct security experiments at the IXP, although we did not disclose the details of our attacks beforehand. As with IXP-EU-1, the IXP automatically retrieved IRR data (our AS-SET) to validate the prefixes we were authorized to announce. Consequently, our experiments did not rely on privileged access and would have worked identically through the standard onboarding process.

IXP-AM: The process was similar to that of IXP-EU-2, with standard provisioning and no privileged access required.

5.2 Feasibility of the Attacks

We first emulate a realistic attacker who seeks to avoid detection by public BGP collectors. To do so, we identify IXP peers that either operate a RouteViews [69] or RIPE RIS [68] collector, or have such collectors within their customer cones, and exclude them from our route announcements. To determine the customer cone of an

IXP peer, we use the CAIDA AS-Relationships dataset [16, 49]. This dataset may be incomplete or inaccurate: an AS erroneously reported as a client of a peer would lead to underestimate the surface because we would exclude the peer from the announcement, whereas a missing client of a peer just makes the route visible at the collector. We adopt the following iterative approach: After each announcement, we inspect public collectors and exclude any additional peer that forwarded our route (i.e., the AS immediately following the attacker in the AS path) until the attack is no longer visible.

For all the attacks, the order in which we set the different BGP communities did not change the results.

5.2.1 *Exploiting the Path-hiding Mitigation.*

Path-hiding excluded peers: The victim announces our prefix only to the route server with AS path [V] while excluding a subset of ten peers randomly selected among those with an open peering policy in PeeringDB [1] and without BGP collectors in their customer cones. An open peering policy at an IXP indicates that an AS is willing to freely exchange traffic with any other participant at that IXP, without restrictive prerequisites. Consequently, we expect these peers to accept the attacker's announcements. The attacker AS A then announces the same prefix to the route server with AS path [A-V], tagging the route with communities that exclude all peers except those previously excluded by the victim (i.e., the randomly selected open-peering peers), so that the route is propagated only to those peers. To validate the adoption of the malicious route, we send ICMP probes from the victim's VM using a source address within the hijacked prefix and check whether the replies return to the attacker's VM. Across all three IXPs, the ten targeted peers responded to the attacker, confirming that path-hiding mitigation was in effect.

Path-hiding shorter AS path: The victim announces its route with two prepends, yielding AS path [V-V-V]. The attacker then advertises the same prefix with AS path [A-V], excluding peers with BGP collectors in their customer cones. Because the attacker's path is shorter, it is preferred once the BGP decision process reaches the AS-path tie breaker. This attack succeeded at all three IXPs: 36%, 43%, and 52% of the peers at IXP-EU-1, IXP-EU-2, and IXP-AM adopt the malicious route, respectively. Peers that do not adopt it typically enforce restrictive routing policies (e.g., non-open routing policy) or still prefer the victim's path. We analyze this attack surface in more detail in Section 5.3.

Path-hiding mitigation implementations: The three IXP operators confirmed that they were using one of the three techniques in RFC7947, in particular per-client RIB or ADD-PATH. The successful implementation of the attacks demonstrates these path-hiding mitigation techniques can be used by an attacker. There is no reason to think Diverse BGP Paths would be different, as in all three, export policy is applied per client (§2.4.1).

Detecting path-hiding mitigation at scale: Measuring path-hiding mitigation for more IXPs without being connected at the IXPs: assuming one only has access to public data, one needs to identify RS-originated routes from them, which is not trivial. For instance, operators could have a single configuration for both bilateral and multilateral announcements, keeping route server communities that would appear in public BGP data, although some of the routes come from bilateral peerings. Monitoring the route-server looking glass is also not sufficient to demonstrate path-hiding mitigation, as IXPs can export all route-server-received routes to the looking glass before (per-client) best-path selection. However, we expect most IXPs to deploy it, as it was standardized in RFC 7947 in 2016 and is commonly requested by clients to avoid missing routes (§7).

5.2.2 *Exploiting the Deployment of Multiple Independent Route Servers.*

We implement the methodology described in Section 4.2.1 to determine the subset of peers that support the NO_EXPORT community, by performing the following experiments, five times in a row. We announce a prefix without the NO_EXPORT community, send pings to the 5 million responsive IP addresses of the ISI hitlist [29] and map each IP address to its AS using routeviews [69]. For each response, we extract its MAC address, which tells us the IXP peer that forwarded the response to us (each peer is identified by its distinct MAC address on the IXP's Layer-2 fabric). This allows us to map each target with a peer of the IXP when the route is announced without the NO_EXPORT community. Then, we withdraw the announcement, wait 15 minutes, and announce the route again, but this time with the NO_EXPORT community. We send the same pings, and map each IP address to their IXP peer, or to no peer if the IP address becomes unresponsive.

First, for each peer, we check whether our route was propagated to a BGP collector. Peers that forwarded the route to a collector are confirmed **not** to support NO_EXPORT. For all remaining peers, we compute the fraction of targets that either became unresponsive or switched to a different IXP peer when we set the NO_EXPORT community. Some targets change behavior even for confirmed non-supporting peers, due to unrelated path changes elsewhere in the Internet, corresponding to noise in our measurement. We therefore set a noise threshold equal to the highest such fraction observed among confirmed non-supporting peers: if a peer known not to support NO_EXPORT can cause this much change by coincidence, fractions at or below that level are uninformative. Any peer whose fraction exceeds this threshold is classified as supporting NO_EXPORT. These thresholds are 66%, 1.4%, and 46% at the three IXPs.

Again, the technique does not need to be perfect, and the thresholds are here to reduce the number of configurations to test for an attacker. In reality, if an AS still forwards the announcement to a BGP collector peer whereas it was categorized as supporting the NO_EXPORT community, the attacker can always remove that peer from the announcement, and vice versa, using their own prefixes. Using these thresholds, our malicious announcement was never forwarded to the BGP collectors at the three IXPs.

With the set of ASes supporting the NO_EXPORT community, we can perform the attack. The victim announces our /23 prefix with AS path [V]. The attacker advertises two routes to the different route servers: (1) a more specific /24 subprefix with AS path [A-V], to peers with no BGP collectors in their customer cones. (2) a more specific /24 subprefix with AS path [A-V] with the NO_EXPORT community, to peers with BGP collectors in their customer cones.

The combination of the two routes significantly improves the number of peers that adopt the route: 41%, 51%, and 49% of receiving

peers at IXP-EU-1, IXP-EU-2, and IXP-AM adopt the malicious route respectively, versus 11%, 15% and 6% compared to the Subprefix + NO_EXPORT attack alone.

5.2.3 Invisibility to BGP Public Collector Peers and Data-plane Detection Techniques. For all attacks, we forward the intercepted traffic back to the victim to preserve connectivity and monitor traffic at the victim's VM to confirm successful interception. We also check that no public BGP collector observes the malicious path using BGPstream [58]. In all cases and across all IXPs, the attacks remain invisible to collectors while successfully intercepting traffic, so techniques based on visibility at collector peers [36, 75], or on connectivity loss [84] do not work.

For latency-based detection techniques, in our setup, the attacks do not increase latency (§4.1.1), as both the malicious and the legit routes are announced from the same single location. Therefore, in our setup, such techniques cannot detect our attacks. However, in the case of a route being announced from multiple locations (e.g., anycast), some ASes adopting the malicious route could experience increased latency to reach the victim's prefix if they switch to a further anycast site. Whether or not state-of-the-art techniques, such as HiDe [73], would detect the attack, depends on the severity of the latency increase and the location of the different anycast sites. HiDe reports that under ideal conditions, it should be able to detect hijacks when the latency increase is as low as 5 ms. However, under real conditions, the authors found that RTTs were actually noisy (Section 3.5 of [73]), so in practice it is not easy for an operator to detect the attacks solely based on latency variation.

5.2.4 Route-server Behaviors. Executing the path-hiding excluded peers attack requires identifying which ASes were excluded from specific announcements, so that the attacker can subsequently target them with a malicious route. However, IXPs differ in how much information their route servers expose, which affects how easily an attacker can infer these ASes. For instance, the route server at IXP-EU-1 propagates the route-server communities to peers, which allows them to observe the excluded ASes. In contrast, at IXP-EU-2 and IXP-AM the route servers do not forward these communities to peers. However, at all three IXPs the communities remain visible in the route server looking glass.

Depending on the route server's behavior, an attacker can either monitor BGP updates directly from it (e.g., using the BGP monitoring protocol), or query its looking glass through public API [4], to identify vulnerable prefixes.

5.2.5 Visibility at Looking Glasses to Detect Attacks. Although route server looking glasses are not commonly used by hijack monitoring systems, we examine whether they could aid in detecting our attacks. Whether they are useful, however, depends on their implementation, as some display only the best route for each prefix [5, 26, 56], whereas others expose all available routes, directly [20] or via APIs [39]. The route server looking glasses at two of the IXPs we studied only show the best route. Under this behavior, they can reveal the path-hiding shorter AS path and multiple route servers attacks using subprefix. In the first one, the malicious route is preferred over the legitimate one; in the second, the malicious announcement targets a more specific prefix and therefore appears independently of the covering route. In contrast, the path-hiding excluded peers attack may remain undetected. The attacker can deliberately advertise a less-preferred route that is propagated only to selected peers and never becomes the route server's best path, and thus does not appear in the looking glass. If the route server looking glass shows all the routes, which is what was implemented by the third IXP and what we recommend (§6), then the malicious route is always visible.

5.2.6 Private Data to Help Detection. For control plane evasion, our threat model considers the attack successful if no malicious route appears in the public BGP collector peers. We nonetheless describe how other non-public vantage points can help detect our attacks. First, some operators have access to private BGP collector peers, e.g., from private companies [28], or from a privileged access to all the routes received by the route server (if the looking glass does not provide it publicly. Second, an IXP operator can monitor a traffic increase for a particular prefix at the IXP to help detect a hijack. Indeed, the set of ASes that was originally not routed to the IXP to reach the prefix (e.g., in an anycast deployment) and that start sending traffic to the IXP because they adopted the malicious route will cause an increase of traffic for that prefix.

5.2.7 Multipath BGP. We evaluate how multipath interacts with our attacks. The expected behavior is that when an AS installs both a legit route from an AS and a malicious route from another AS, it should perform load balancing between the two. With only two ASes under our control, we cannot set up an experiment with one victim AS, one attacker AS, and a third AS that receives and installs both routes with multipath BGP. Instead, we verify that an AS enabling multipath BGP that receives two equal routes via two different ASes actually performs load balancing between the two. There is no reason why the behavior should change between two legit routes, and one legit and one malicious route.

We enable multipath BGP on one AS (AS M) (and set no local pref values), and announce multiple routes from our other AS (AS O): One route to AS O's Tier-1 provider (via a bilateral peering at the IXP), and one direct route to AS M via the RS prepended once. Both routes have the same length and are received, considered equal, and installed, by AS M. Then, we perform pings from AS M with different source IP addresses to vary the flow identifier and check the load balancing behavior. Indeed, we observed that AS O received packets either from the Tier-1 provider, or directly from AS M, using the MAC addresses of the packets to know their origin.

5.3 Potential Attack Surface

We estimate the potential attack surface of our attacks, i.e., the portion of the Internet that could adopt the malicious routes. We compare the path-hiding shorter AS path and multiple route server attacks with the subprefix + NO_EXPORT attack [12]. We deliberately do not compare with AS-PATH poisoning techniques [53] because it would have required poisoning tens (or hundreds) of ASes, which violates some IXPs' and operators' best practices [57]. We analyze the path-hiding excluded peers attack separately (§5.3.2) because its attack surface is prefix-dependent. For this attack, we report percentages rather than raw numbers to preserve the anonymity of the IXPs.

Table 2: Number of IP targets, ASes, and population that adopted the malicious route for different attacks exploiting path-hiding mitigation (PH), and combined with a subprefix + NO_EXPORT community attack exploiting the deployment of multiple independent route servers.

	IXP-EU-1			IXP-EU-2			IXP-AM		
	IP	AS	Pop.	IP	AS	Pop.	IP	AS	Pop.
Path-hiding shorter AS path	110k	1558	165M	34k	1030	99M	15k	3294	126M
Multiple route servers (PH Shorter AS path and subprefix + NO_EXPORT)	113k	2312	526M	31k	1446	383M	14k	3062	120M
Multiple route servers (Subprefix and subprefix + NO_EXPORT)	**161k**	**5725**	**957M**	**61k**	**3170**	**476M**	**18k**	**3556**	**171M**
Subprefix + NO_EXPORT [12]	61k	4409	871M	40k	2481	423M	5185	762	27M

5.3.1 Path-hiding Shorter AS Path and Multiple Route Servers Attacks. We quantify the attack surface using three metrics: (1) the number of IP addresses adopting the malicious route; (2) the number of ASes adopting the malicious route; (3) the estimated population (i.e., number of users) affected by the malicious route.

To compute these metrics, we send ICMP echo requests to responsive destinations in the ISI hitlist [29, 80], which contains one responsive IP address in each of the 5M /24 prefixes, covering 73k ASes. Because the VM's interface connected to the IXP LAN can only reach prefixes learned via the route server, we send probes from the VM's interface connected to a transit provider, which has access to the full Internet routing table. We say that an IP address is affected if the reply returns through the attacker's interface on the IXP LAN. We then map each IP address to its AS using Routeviews. An AS is considered affected if at least one of its addresses is affected. Finally, we map ASes to population estimates using the validated [70] APNIC dataset [8]. This methodology provides a reasonable approximation. Indeed, we find that when one IP address of an AS adopts the malicious route, there are between 75% and 90% of the ASes for which this is also true for all of the IP addresses in these ASes, over the different types of attacks at the three IXPs.

Table 2 summarizes the results. The results should be interpreted as reasonable relative estimates rather than exact values. Two factors may bias them in either direction. First, we probe only responsive hosts, so non-responsive destinations may also be affected, causing us to underestimate the attack surface. Second, the prefix of our victim is announced from a single location, whereas a real victim may advertise the same prefix from multiple locations, which could reduce the number of ASes that adopt the malicious route and therefore lead to overestimation.

There are two takeaways: First, announcing multiple malicious routes to different route servers does increase the attack surface from just announcing a subprefix + NO_EXPORT community. At IXP-EU-1, it goes from 61k IP addresses corresponding to 4409 ASes and 871M users to 161k IP addresses, 5725 ASes, and 957M users, or an increase of 29% in terms of polluted ASes. The increases of the number of polluted ASes vary for other IXPs, with 28% at IXP-EU-2 and 366% at IXP-AM.

Second, if a subprefix attack is not possible because a prefix is a /24 or protected by a maxLength parameter, one can nonetheless perform a path-hiding shorter AS path attack, which attracts, at IXP-EU-1, traffic from 110k IPs. It corresponds to 1,558 ASes and approximately 165M users. Results are slightly lower at IXP-EU-2, but higher at IXP-AM (3,294 ASes). This difference reflects the fraction of peers excluded to evade BGP collectors and maintain the attack invisible to public collector peers: 28% at IXP-EU-1, 31% at IXP-EU-2, and only 7% at IXP-AM, resulting in a larger attack surface at IXP-AM.

5.3.2 Path-hiding Excluded Peers. This attack must be evaluated on a per-prefix basis. At IXP-EU-1, at most eight such ASes are excluded per route. At IXP-AM, 97% of routes exclude fewer than 10, although a small number exclude up to 45. These numbers increase when we consider ASes with BGP collectors in their customer cones. Include communities exhibit a similar pattern. At IXP-AM, 60% of announcements that use include communities result in 99% of peers being excluded. However, such announcements are rare, representing only 3% of all routes tagged with include or exclude communities. Overall, 5% of peers at IXP-EU-1 and 13% at IXP-AM are excluded in at least one announcement while having no BGP collector in their customer cones.

5.4 Potentially Vulnerable Prefixes

We next assess which prefixes are susceptible to hijacking. A prefix may be vulnerable to any combination of the three attack types, including all three. Specifically, we classify a prefix as potentially vulnerable if the best route to it meets any of the following conditions: (1) for the path-hiding excluded peers attack, its route is tagged with include or exclude communities; (2) for the path-hiding shorter AS path attack, the AS-path length of the route announced at the IXP is greater than two; and (3) for the multiple route servers attack using the subprefix + NO_EXPORT, the prefix is covered by an ROA whose maxLength exceeds the prefix length.

We use the word "potentially" deliberately. Meeting these conditions does not guarantee that a hijack will succeed, as a more attractive route may exist elsewhere and prevent the malicious route from being selected, depending on competing paths, local preference, peer policy, or route filtering.

We also compute the number of potentially vulnerable prefixes for attacks using Subprefix + NO_EXPORT[12], which needs condition (3) to work, and using AS path poisoning [53]. For the latter, the attacker must announce a route with a minimal length of three [attacker AS, poisoned ASes, origin AS], which corresponds to the most favorable scenario, when only one AS needs to be poisoned to evade BGP collectors. The minimal AS-path length announced by the attacker is of length three, so the prefix is considered vulnerable if the AS-path length of the best route to it is of length four or more.

Table 3 shows the fraction of potentially vulnerable prefixes depending on the IXP and the attack. There are two takeaways: first, overall, a significant number of prefixes is vulnerable to at least one attack, with 82%, 73%, and 90% of the prefixes announced at the

Table 3: Fraction of vulnerable prefixes for different types of attacks. PH stands for Path Hiding.

	Potentially vulnerable prefixes		
	IXP-EU-1	IXP-EU-2	IXP-AM
PH excluded peers	53%	x	76%
PH shorter AS path	59%	62%	54%
Multiple RS (Subprefix)	22%	24%	29%
Union	**82%**	**73%**	**90%**
Subprefix + NO_EXPORT [12]	22%	24%	29%
AS path poisoning [53]	30%	32%	29%

three IXPs, with each of the attacks based on path hiding affecting more than half of the prefixes at each IXP. Second, \revised[R-2]{these attacks are possible on a significantly higher number of prefixes compared to attacks from prior work, with 22%, 24%, and 29% for the subprefix + NO_EXPORT attack, and 30%, 32%, and 29% for the shorter AS path using AS path poisoning}.

Targeting specific prefixes: While the previous analysis quantifies the aggregate number of vulnerable prefixes, an attacker may instead target specific services [13], e.g., websites that handle sensitive information. To evaluate this risk, we examine whether potentially vulnerable prefixes host government services in the countries where our three IXPs operate. From an attacker's perspective, the objective is to determine whether users access these services through the IXP and whether the corresponding prefixes are vulnerable. We obtain government websites from prior works datasets [45] and resolve their DNS names using RIPE Atlas probes located in the country of the IXP of interest, using their local resolver. For each resolved address, we check whether it falls within a previously identified vulnerable prefix. When it does, we perform traceroutes to check whether the path traverses the IXP, i.e., whether an IP address from the IXP LAN appears in the path. If both conditions hold, we classify the website as potentially vulnerable.

At IXP-EU-1, we identify 796 government websites across 99 prefixes as potentially vulnerable, compared to 303 websites across 47 prefixes at IXP-EU-2 and 224 websites across 113 prefixes at IXP-AM. These users originate from 45 ASes for IXP-EU-1, 19 for IXP-EU-2, and 55 for IXP-AM. Interestingly, we also observe users in the country of IXP-EU-1 that reach three government websites through IXP-EU-2.

5.5 Qualitative Comparison with Prior Attacks

After the quantitative comparison with prior attacks on attack surface and the number of vulnerable prefixes, we compare our attacks with three recent BGP hijacking attacks [12, 13, 53], along three qualitative dimensions: (1) stealthiness with respect to control-plane monitoring with public collector peers; (2) stealthiness with respect to data-plane monitoring; and (3) compliance with RPKI/ROV. We also discuss operational cost, as prior works explicitly seek to minimize it [13].

Recall that our attacks are invisible to BGP public collector peers and evade techniques based on connectivity, as they are interception attacks. For latency-based techniques, whether our attacks increase latency depends on the prefix deployment. In our setup where the prefix is announced from a single place at one IXP, our attacks do not increase latency.

SICO [13]: SICO is an interception attack in which the attacker connects to two providers: one to hijack traffic from its customers and another to maintain connectivity to the victim. The attack is resilient to data-plane connectivity monitoring but was not designed to evade public BGP collectors.

One might expect that BGP communities could be used to hide the attack from public BGP collectors. However, this intuition does not directly apply to SICO's setting, where restricting route visibility using BGP communities might interfere with the conditions required to preserve traffic interception.

Figure 3 of the SICO paper illustrates three attack scenarios and the mechanisms used to ensure that at least one provider continues to forward traffic to the victim. In the first scenario (Figure 3a of [13]), that attacker is connected to two providers, A and B, which are interconnected and connected also to a Tier-1 provider. The victim is also connected to the same Tier-1 provider but has no direct connection to providers A or B. The victim announces a route to the prefix to the Tier-1 provider. To carry out the attack, the attacker announces a malicious route to provider A and tags it with a NO_EXPORT_ALL or a NO_EXPORT_SELECT BGP community to prevent propagation to provider B and the Tier-1 provider. If the attacker uses a NO_EXPORT_ALL community, the route is accepted only by provider A and is not propagated to its customers. As a result, the visibility of the attack depends on whether provider A peers with a BGP collector that belongs to the same AS confederation [14, 47]. Moreover, because provider A's customers do not receive the route, the attack surface is significantly reduced, so it restricts the set of ASes that can be targeted in the context of a targeted interception. If the attacker instead uses NO_EXPORT_SELECT, the situation does not improve significantly. Although providers may support suppressing announcements to peers (Table 10 of [13]), they typically do not offer mechanisms to suppress announcements to specific customer ASes. Consequently, SICO has limited control over the visibility of the malicious route. In the remaining two scenarios shown in Figure 3 of [13], the malicious route is likewise propagated to the customers of provider A, which leads to the same limitation.

SICO is also not robust to latency-based detection techniques [73], unless the attacker is geographically and topologically close to the victim. Finally, SICO incurs higher operational costs, as the attacker must pay transit fees for both receiving and forwarding hijacked traffic. In contrast, our attacks intercept traffic at peer-to-peer links at IXPs and therefore avoid transit costs.

Birge-Lee et al. [12]: In this attack, the attacker connects to a provider that is shared by the victim, the target ASes, and the attacker itself (Figure 1 of [12]). The attacker announces a more specific prefix to the shared provider and tags the route with a NO_EXPORT_ALL community to prevent its propagation to the provider's customers and peers. As a result, the provider continues to receive traffic destined to the victim but forwards the traffic destined to the more specific prefix to the hijacker. To maximize the attack surface, the attacker seeks to connect to a large provider (e.g., a Tier-1).

This attack evades public BGP collectors but does not evade data-plane detection techniques. Moreover, the paper does not describe a variant of the attack that forges the origin AS, so the attack does not succeed in the presence of ROV. Finally, as acknowledged in [12], the ROA maxLength attribute can prevent such subprefix attacks.

Milolidakis et al. [53]: This attack uses AS-PATH poisoning to evade BGP collectors by inserting collector ASNs into the path. When the AS of a collector receives such a route, BGP's loop-prevention mechanism discards it, making the route invisible to the collector. Although effective in principle, the evaluation considers AS paths of length at most five (Section VII of [53]), which is insufficient to evade the large number of collectors observed in practice. Moreover, the paper does not experimentally demonstrate resilience against data-plane detection techniques. Finally, although the attack complies with RPKI/ROV, the use of long AS paths may trigger filtering by networks that enforce maximum path-length limits, which limits the attack's deployability.

6 DETECTION AND RECOMMENDATIONS

We propose techniques and recommendations to better detect attacks at IXPs, both at the data and control planes.

6.1 Data Plane

A potential victim present at an IXP can flag attacks by continuously pinging its peers' addresses and monitoring the source MAC addresses of returning ICMP echo reply packets. Pings to a given address will consistently generate replies with the same MAC source address: that of the peer's interface on the IXP's LAN. A change in this source MAC address would indicate that the return path now passes through a different peer. This could be benign, but it could also be a sign that the peer now appearing as the MAC-layer source has hijacked one of the victim's prefixes and is forwarding intercepted traffic, including the ping replies, to the victim in an effort to maintain connectivity. Our technique assumes that a peer cannot spoof the MAC address of another peer, a constraint that is recommended as best practice [50], easy to enforce by IXPs at L2, and applied by the three IXPs we study. A victim that is not directly present at the IXP could delegate this monitoring to a trusted provider that is present at the IXP, authorizing it to issue pings on its behalf using source IP addresses from the victim's advertised prefixes.

Figure 2 illustrates the approach. V sends regular ICMP echo requests, sourced from each of its advertised prefixes, to every peer at the IXP. Here, V is pinging AS 4. If AS 4 is not compromised (Figure 2a), the echo replies travel directly across the IXP's data plane from AS 4's router to V's, and V can confirm this by checking the source MAC address (shown as MAC@4). If, however, AS 4 has been compromised to redirect traffic destined to V's prefix toward an attacker A (Figure 2b), this traffic includes the echo replies, which first reach A. In an attempt to go undetected, A forwards all intercepted traffic onward to V, including the echo replies. But V will notice that the source MAC address on those replies has changed (shown as MAC@A). When V detects such a change, it should investigate whether A is legitimately announcing V's prefix, for instance by checking A's AS-SET or contacting A directly. This technique is semi-automated, requiring some manual investigation, but it provides a useful early warning signal for suspicious forwarding behavior.

We verified that our technique is practical by checking five conditions: (1) all peers that the victim wishes to monitor are reachable via ICMP echo request and respond with distinct MAC addresses; (2) when we perform our attack, the source MAC address of packets from peers that have adopted our malicious route changes from the victim's perspective, while it remains unchanged for peers that have not; (3) the peer's MAC addresses do not often change, and cannot be spoofed; (4) false positives (benign MAC address changes), i.e., a peer that changes its next hop to the monitored prefix to another peer of the IXP, should be rare or easy to identify as benign, and (5) the technique is scalable and has a small traffic overhead.

Conditions (1) and (2) were satisfied at the three IXPs, and confirmed by their operators. For condition (3), IXP operators confirmed that the MAC address table was very stable, strict filtering on MAC addresses was enforced, and a port of the IXP fabric was only allowed to emit traffic with its assigned MAC address.

For condition (4), we deployed the technique during three days in one of our ASes at the three IXPs, pinging each peer every minute. Across thousands of peers, we observe 0, 1, and 138 MAC changes at IXP-EU-1, IXP-EU-2, and IXP-AM. Although 138 changes can represent a too important overhead for operational teams, these changes at IXP-AM were triaged automatically as benign, as the IXP-AM route-server looking glass, which shows all received routes for a prefix, showed no other route to our prefix, and no malicious one by extension. For the one case at IXP-EU-2, manual investigation showed it was a routing change from a peer switching to its provider, also connected at IXP-EU-2, to reach our prefix, so a benign change. This analysis advocates for multiple-route visibility at LGs to allow automatic triage of false positives and save investigation time.

For condition (5), in the worst case, each peer send pings to all other peers, for all of their vulnerable prefixes (§5.4), plus the subprefixes (up to /24) contained in any vulnerable prefix that is not protected (no ROA of ROA with maxLength > prefix length). The median (90th percentile) probing rate per operator at IXP-EU-2 (worst of the three) is 26 (3.9k) packets/s at a 10-minute granularity, which is negligible for most ASes. Still, if this is too much overhead for ASes announcing many prefixes (e.g., large providers), they can adapt the frequency by prefix and peer importance, or traffic volume. Moreover, if the LG reports all RS-received routes, the data-plane method can be used only for further investigation, to reduce overhead even more. Altogether, the technique incurs a total overhead of 660 Mb/s, which is small for a multi-Tb/s IXP.

As a final note, one IXP noted that deploying this mechanism directly on peers' border routers could be cumbersome. The technique can be adapted: instead of pinging from the border router itself, a server connected to the router's LAN can send the pings, targeting an IP address of the peer's router at the IXP other than the one on the LAN (which is only reachable locally), and use port mirroring to see the MAC addresses of the echo replies received by the router's LAN. This adaptation would simply require IXP members to share this alternative IP address and make it pingable.

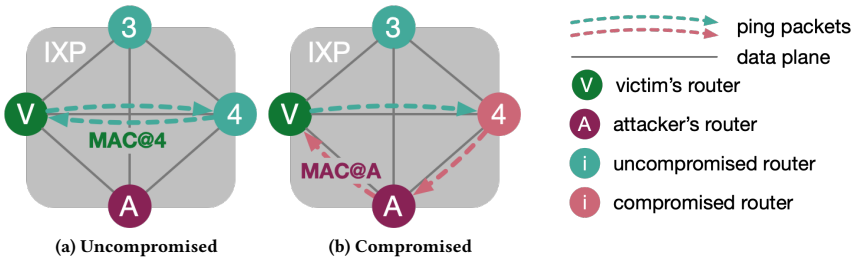

Figure 2: Detection of suspicious forwarding behavior at the IXP based on MAC address change.

6.2 Control Plane

Our first recommendation is that route server looking glasses should expose not only the best route for a prefix, but all routes learned by the route server, as already implemented at some IXPs [20, 39]. Full route visibility would make malicious routes observable in all cases (§5.2.5). Alice looking glass, a popular implementation, already proposes the option to expose all the routes, and so we recommend enabling it. We also recommend that operators monitor the prefixes they announce at the router server using the looking glass, which provides local and IXP-specific visibility. Because looking glasses often enforce query rate limits [78], our data plane method for detecting suspicious forwarding behavior can trigger targeted looking glass queries only when anomalies are suspected. Moreover, we recommend that IXP route servers peer with public BGP collectors, as already done at two of the three IXPs we studied. Such peering allows IXPs to directly observe announcements that exclude public collectors. Finally, route servers should avoid forwarding include and exclude communities to peers. This policy, already adopted at two of the three IXPs we evaluated, makes it harder to identify vulnerable prefixes (§5.3.2).

Our second recommendation is based on a common usage of exclude communities [6]: An AS announces two types of routes at the IXP: one to the peers with which it has bilateral peerings, and one via the route server that excludes those peers. In this scenario, we recommend assigning a higher local preference to the route learned via the bilateral peering, in case a malicious route announced at the route server tries to compete with it, which is the case for attacks exploiting the path-hiding mitigation.

Our third recommendation is that IXPs should limit the number of communities that an AS can put on a route, or at least flag them (e.g., with a specific community). Some IXPs' route servers already reject these routes[21]. In our attacks, an attacker needs to use on average 146, 196, and 198 exclude communities at IXP-EU-1, IXP-EU-2, and IXP-AM to evade public collector peers. These numbers are significantly higher than average numbers at IXPs: On 12 IXPs providing a looking glass that programmatically exposes the routes received by their route servers, a route has on average 10.8 communities per route across the IXPs (although a few routes with hundreds of communities exist). If an attacker uses include communities instead of exclude communities to evade public collector peers, rejecting routes with too many communities could still help to reduce the number of ASes adopting the malicious route.

Finally, we recommend the adoption of ASPA (Autonomous System Provider Authorization) [9]. ASPA enables routers to validate whether an AS path is valid, by checking customer-provider relationships that ASes cryptographically publish. In our attacks, if the two ASes surrounding the malicious AS (the peers of the IXP) have an ASPA object, then the AS path would be invalid and rejected, as none of the surrounding ASes would have the malicious AS listed in its providers. However, if only one of the peers deploys ASPA, the AS path cannot be invalidated.

In detail, let us consider ASPA validation of AS paths observed between two peers (A) and (B) at an IXP. Let (M) be a malicious AS that attempts to insert itself between them, yielding a path fragment (A ; . . . ; M ; . . . ; B), where (M) is not a provider of either peer.

ASPA validation identifies the max-up-ramp (the longest initial segment of the AS path consistent with provider-to-customer relationships) and the min-down-ramp (the shortest terminal segment of the AS path consistent with customer-to-provider relationships). A path is classified as Invalid if these segments overlap (i.e., max-up-ramp >= min-down-ramp), indicating a violation of ASPA-consistent valley-free structure.

If both (A) and (B) publish ASPA objects, the insertion of (M) necessarily creates an invalid provider-customer relationship at one or both adjacent edges. Since (M) is not listed as a provider by either peer, neither ((A, M)) nor ((M, B)) can be validated consistently. Consequently, ASPA validation yields overlapping max-up-ramp and min-down-ramp segments, and the path is classified as **Invalid**.

If only one peer (e.g., (A)) publishes an ASPA object, validation is incomplete. While the edge adjacent to (A) is constrained by its ASPA object, the segment near (B) remains unconstrained. As a result, the max-up-ramp is bounded, but the min-down-ramp may extend through the unverifiable portion of the path without contradiction. In this case, the path may be classified as **Unknown** rather than **Invalid**, and thus is not reliably rejected.

7 DISCUSSION WITH IXPS

We disclosed our findings to the three IXPs where we tested our attacks. We summarize the main points of the discussions.

First, all IXPs were already aware of the path-hiding problem and the existing mitigation mechanisms [43]. In particular, IXP-EU-2 reported that this functionality had been explicitly requested by some of their clients. They described a case in which a large transit provider announced a prefix using BGP communities to propagate

the route to a specific peer while excluding others, thereby leaving those clients without any route to the prefix at the IXP.

Second, all IXPs agreed that the attacks we describe are real operational threats and acknowledged that no straightforward mitigation is currently available. IXP-EU-1 also mentioned that they cannot thoroughly validate all information contained in their clients' AS-SETs due to the scale of the data involved and the difficulty of distinguishing legitimate from illegitimate entries in the IRR.

Third, the IXPs agreed that our recommendations are reasonable and worth further consideration. Two of them initiated internal discussions to assess deployment costs and to define appropriate operational responses in the event of an alert. On the control plane, IXP-EU-1 suggested developing a best current practice (BCP) RFC to encourage full route visibility for a prefix via the looking glass. On the data plane, the IXPs expressed interest in the deployment or our detection technique.

8 RELATED WORK

BGP attack techniques: A substantial body of prior work has explored increasingly sophisticated techniques to launch attacks on BGP routing [11, 13, 34, 53, 59]. These attacks rely on mechanisms such as selective route advertisement [11], the use of BGP communities [13], announcing a route with a shorter AS path [34, 53], or announcing a more specific prefix [12]. They are also designed to evade control plane [12, 53] or data plane [13] detection systems. We provide a detailed qualitative comparison with three state-of-the-art attacks in Section 5.5.

Again, our attacks differ from this prior work in that they do not introduce new BGP hijack primitives, but rather combine these different hijack primitives with the two identified IXP common route server design choices, i.e., path-hiding mitigation and the deployment of multiple independent route servers, to enable attacks with a significantly bigger attack surface and a higher number of potentially vulnerable prefixes.

Detection systems: Existing detection systems rely on control-plane techniques [17, 36, 46, 62, 75, 79], data-plane techniques [15, 73, 84, 85], or both [37, 63, 72, 77].

Control plane techniques detect anomalies such as Multiple-Origin AS (MOAS) events [46], spurious links inferred from topology violations [17, 75], or deviations in routing behavior learned via machine learning models [36], or edit distance between routing trees to the victim and the attacker [79]. All these techniques require visibility of the malicious route in the control plane, which our attacks deliberately avoid. Data plane techniques monitor connectivity disruptions [84, 85] or latency changes [73] to infer attacks. Our interception attacks preserve connectivity and can eventually be engineered to avoid measurable latency inflation (§4.1.1). Hybrid approaches [61, 77] combine control-plane and data plane signals, and therefore inherit the same limitations.

IXP architecture: Prior work studied route-server features [67], peering relationships [32, 33, 71], traffic volumes at large IXPS [3], IXP route propagation [60], or IXP blackholing techniques (RTBH) [54]. For the RTBH announcements, these very specific (and accepted by route servers) prefixes, could potentially be added to the set of possible malicious routes that an attacker at the IXP could use. We leave this as future work.

9 CONCLUSION

In this paper, we identified two route-server design choices, path-hiding mitigation and deployment of multiple independent route servers, which allows an attacker to perform BGP hijacks at IXPs that are invisible to public BGP collector peers while keeping the connectivity. Then, we demonstrated the feasibility of these attacks in three major IXPs, and showed that these attacks have a significantly higher attack surface and a higher number of vulnerable prefixes compared to prior work. Finally, to improve the detection of these attacks, we presented a new practical MAC based data-plane technique, and proposed some immediately actionable recommendations at the control plane.

Acknowledgements: We thank the IXPs that let us perform our experiments. We used Grammarly for proofreading. This work was supported by FAPERGS grant 25/2551-0000971-5, CNPq grants 312371/2026-8, 420934/2023-5, and 407302/2025-5, CGI.br IETF Program, a grant from the French Ministry of Armed Forces, and by the "Wallonie 2021-2027" program FEDER CyberGalaxia.

.

REFERENCES

[1] 2026. PeeringDB. https://www.peeringdb.com.
[2] NAP Africa. 2026. NAP Africa RPKI. https://www.napafrica.net/technical/rpki-handy-hints/.
[3] Bernhard Ager, Nikolaos Chatzis, Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. 2012. Anatomy of a large European IXP. In *Proc. ACM SIGCOMM*. 163–174.
[4] Yasin Alhamwy, Bashar Khoulani, and Oliver Hohlfeld. 2025. Crawling Alice Looking Glasses at IXPs to Quantify BGP Route Diversity. In *Proc. ACM SIGCOMM (Posters and Demos)*. 121–123.
[5] AMS-IX. 2026. AMS-IX Looking Glass. https://lg.ams-ix.net.
[6] AMS-IX. 2026. AMS-IX Route Servers. https://www.ams-ix.net/ams/documentation/ams-ix-route-servers.
[7] AMS-IX. 2026. Total statistics at AMS-IX. https://www.ams-ix.net/ams/documentation/total-stats.
[8] APNIC. 2026. AS population dataset. https://stats.labs.apnic.net/aspop/.
[9] Alexander Azimov, Eugene Bogomazov, Randy Bush, Keyur Patel, Job Snijders, and Kotikalapudi Sriram. 2025. BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects. IETF draft ASPA.
[10] Michael Bailey, David Dittrich, Erin Kenneally, and Doug Maughan. 2012. The menlo report. *IEEE Security & Privacy* 10, 2 (2012), 71–75.
[11] Hitesh Ballani, Paul Francis, and Xinyang Zhang. 2007. A study of prefix hijacking and interception in the Internet. In *Proc. ACM SIGCOMM*.
[12] Henry Birge-Lee, Maria Apostolaki, and Jennifer Rexford. 2025. Global bgp attacks that evade route monitoring. In *Proc. PAM*. 335–357.
[13] Henry Birge-Lee, Liang Wang, Jennifer Rexford, and Prateek Mittal. 2019. Sico: Surgical interception attacks by manipulating bgp communities. In *Proc. ACM CCS*. 431–448.
[14] Jay Borkenhagen, Randy Bush, Ron Bonica, and Serpil Bayraktar. 2019. Policy Behavior for Well-Known BGP Communities. RFC 8642.
[15] Tobias Bühler, Alexandros Milolidakis, Romain Jacob, Marco Chiesa, Stefano Vissicchio, and Laurent Vanbever. 2023. Oscilloscope: Detecting BGP hijacks in the data plane. *arXiv preprint arXiv:2301.12843* (2023).
[16] CAIDA. 2026. AS Rank - Autonomous System Rankings. https://asrank.caida.org/.
[17] Ying-Ju Chi, Ricardo Oliveira, and Lixia Zhang. 2008. Cyclops: the AS-level connectivity observatory. *ACM SIGCOMM Computer Communication Review* 38, 5 (2008), 5–16.
[18] Cloudflare. 2025. Cloudflare 1.1.1.1 incident on July 14, 2025. https://blog.cloudflare.com/cloudflare-1-1-1-1-incident-on-july-14-2025/.
[19] Team Cymru. 2026. Bogon Reference BGP. https://www.team-cymru.com/bogon-reference-bgp.
[20] DEC-IX. 2026. DEC-IX Looking Glass. https://www.de-cix.net/en/services/looking-glass.
[21] DEC-IX. 2026. DEC-IX Looking Glass API. https://lg.de-cix.net/api/v1/config.
[22] DEC-IX. 2026. DEC-IX Route Servers. https://docs.de-cix.net/article/traw7kch04-de-cix-frankfurt-globepeer-route-server-guide.

[23] Ben Du, Gautam Akiwate, Thomas Krenc, Cecilia Testart, Alexander Marder, Bradley Huffaker, Alex C Snoeren, and Kimberly C Claffy. 2022. IRR Hygiene in the RPKI Era. In *Proc. PAM*. 321–337.

[24] EQUINIX. 2026. EQUINIX route server. https://docs.equinix.com/internet-exchange/peering/ix-prod-mlpe/.

[25] EQUINIX. 2026. EQUINIX RPKI implementation. https://docs.equinix.com/internet-exchange/peering/ix-rpki/.

[26] Equinix. 2026. IEquinix Looking Glass. https://lg.equinixmetal.com.

[27] Cisco Thousand Eyes. 2020. Why Rostelecom's Route Hijack Highlights the Need for BGP Security. https://www.thousandeyes.com/blog/rostelecom-route-hijack-highlights-bgp-security.

[28] Thousand Eyes. 2026. Thousand Eyes: bgp-and-route-monitoring. https://www.thousandeyes.com/fr-fr/solutions/bgp-and-route-monitoring.

[29] Xun Fan and John Heidemann. 2010. Selecting representative IP addresses for Internet topology studies. In *Proc. ACM IMC*. 411–423.

[30] Lixin Gao. 2001. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on networking* 9, 6 (2001), 733–745.

[31] Lixin Gao and Jennifer Rexford. 2001. Stable Internet routing without global coordination. *IEEE/ACM Transactions on networking* 9, 6 (2001), 681–692.

[32] Vasileios Giotsas, George Nomikos, Vasileios Kotronis, Pavlos Sermpezis, Petros Gigis, Lefteris Manassakis, Christoph Dietzel, Stavros Konstantaras, and Xenofontas Dimitropoulos. 2020. O peer, where art thou? Uncovering remote peering interconnections at IXPs. *IEEE/ACM Transactions on Networking* 29, 1 (2020), 1–16.

[33] Vasileios Giotsas, Shi Zhou, Matthew Luckie, and Kc Claffy. 2013. Inferring multilateral peering. In *Proc. ACM CoNEXT*. 247–258.

[34] Sharon Goldberg, Michael Schapira, Peter Hummon, and Jennifer Rexford. 2010. How secure are secure interdomain routing protocols. In *Proc. ACM SIGCOMM*.

[35] HK-IX. 2026. HK-IX route policy. https://www.hkix.net/hkix/route-policy.htm.

[36] Thomas Holterbach, Thomas Alfroy, Amreesh Phokeer, Alberto Dainotti, and Cristel Pelsser. 2024. A system to detect {Forged-Origin}{BGP} hijacks. In *Proc. USENIX NSDI*. 1751–1770.

[37] Xin Hu and Z Morley Mao. 2007. Accurate real-time identification of IP prefix hijacking. In *Proc. IEEE S&P*. 3–17.

[38] IRR. 2026. IRR website. https://irr.net.

[39] IX.br. 2026. IX.br Looking Glass. https://lg.ix.br/.

[40] IX.br. 2026. IX.br route server. https://ix.br/route-server-and-commmunities.

[41] IX.nz. 2026. IX.nz peering. https://ix.nz/service/peering/.

[42] IXP tracker. 2026. IXP tracker. https://pulse.internetsociety.org/en/ixp-tracker/.

[43] Elisa Jasinska, Nick Hilliard, Robert Raszuk, and Niels Bakker. 2016. Internet Exchange BGP Route Server. RFC 7947. https://doi.org/10.17487/RFC7947

[44] JPNAP. 2026. JPNAP IXP option. https://www.jpnap.net/en/ix/option.

[45] Rashna Kumar, Esteban Carisimo, Lukas De Angelis Riva, Mauricio Buzzone, Fabián E Bustamante, Ihsan Ayyub Qazi, and Mariano G Beiró. 2024. Of Choices and Control-A Comparative Analysis of Government Hosting. In *Proc. ACM IMC*. 462–479.

[46] Mohit Lad, Daniel Massey, Dan Pei, Yiguo Wu, Beichuan Zhang, and Lixia Zhang. 2006. PHAS: A Prefix Hijack Alert System.. In *Proc. USENIX Security*, Vol. 1. 3.

[47] Tony Li, Ravi Chandra, and Paul S. Traina. 1996. BGP Communities Attribute. RFC 1997.

[48] LINX. 2026. LINX Route Servers. https://community.linx.net/exchange-docs-oo8vcsp0/post/linx-route-servers-information-xCXmq6SqZUpC80k.

[49] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and KC Claffy. 2013. AS relationships, customer cones, and validation. In *Proc. ACM IMC*. 243–256.

[50] MANRS. 2020. IXP peering platform: an environment to take care of. https://manrs.org/2020/12/ixp-peering-platform-an-environment-to-take-care-of/.

[51] MANRS. 2022. KlaySwap – Another BGP Hijack Targeting Crypto Wallets. https://manrs.org/2022/02/klayswap-another-bgp-hijack-targeting-crypto-wallets/.

[52] Megaport. 2026. Megaport IX features. https://docs.megaport.com/ix/overview-features/.

[53] Alexandros Milolidakis, Tobias Bühler, Kunyu Wang, Marco Chiesa, Laurent Vanbever, and Stefano Vissicchio. 2023. On the effectiveness of bgp hijackers that evade public route collectors. *IEEE Access* 11 (2023), 31092–31124.

[54] Marcin Nawrocki, Jeremias Blendin, Christoph Dietzel, Thomas C Schmidt, and Matthias Wählisch. 2019. Down the black hole: dismantling operational practices of BGP blackholing at IXPs. In *Proc. ACM IMC*. 435–448.

[55] IXP NG. 2026. IXP NG documentation. https://ixp.net.ng.

[56] NL-IX. 2026. NL-IX Looking Glass. https://lg.nl-ix.net.

[57] NLNOG. 2026. Guide on long paths. bgpfilterguide.nlnog.net/guides/long_paths/.

[58] Chiara Orsini, Alistair King, Danilo Giordano, Vasileios Giotsas, and Alberto Dainotti. 2016. BGPStream: a software framework for live and historical BGP data analysis. In *Proc. ACM IMC*. 429–444.

[59] Alex Pilosov and Tony Kapela. 2008. Stealing the Internet: An Internet-scale man in the middle attack. *NANOG-44, Los Angeles, October* (2008), 12–15.

[60] Lars Prehn, Franziska Lichtblau, Christoph Dietzel, and Anja Feldmann. 2022. Peering only? analyzing the reachability benefits of joining large ixps today. In *Proc. PAM*. 338–366.

[61] Lancheng Qin, Dan Li, Ruifeng Li, and Kang Wang. 2022. Themis: Accelerating the detection of route origin hijacking by distinguishing legitimate and illegitimate {MOAS}. In *Proc. USENIX Security*. 4509–4524.

[62] Jian Qiu, Lixin Gao, Supranamaya Ranjan, and Antonio Nucci. 2007. Detecting bogus BGP route information: Going beyond prefix hijacking. In *Proc. SecureComm*. 381–390.

[63] Tongqing Qiu, Lusheng Ji, Dan Pei, Jia Wang, Jun (Jim) Xu, and Hitesh Ballani. 2009. Locating Prefix Hijackers using LOCK.. In *Proc. USENIX Security*. 135–150.

[64] Robert Raszuk, Rex Fernando, Keyur Patel, Danny McPherson, and Kenji Kumaki. 2012. *Distribution of diverse BGP paths.* Technical Report.

[65] Yakov Rekhter, Tony Li, and Susan Hares. 2006. RFC 4271: A border gateway protocol 4 (BGP-4).

[66] Yakov Rekhter, B Moskowitz, Daniel Karrenberg, GJ de Groot, and Eliot Lear. 1996. Rfc1918: Address allocation for private internets.

[67] Philipp Richter, Georgios Smaragdakis, Anja Feldmann, Nikolaos Chatzis, Jan Boettger, and Walter Willinger. 2014. Peering at peerings: On the role of IXP route servers. In *Proc. ACM IMC*. 31–44.

[68] RIPE NCC. 2026. RIPE RIS. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/.

[69] Route Views Project. 2026. Route Views Project - BGP Data. https://www.routeviews.org/routeviews/.

[70] Loqman Salamatian, Calvin Ardi, Vasileios Giotsas, Matt Calder, Ethan Katz-Bassett, and Todd Arnold. 2024. What's in the Dataset? Unboxing the APNIC per AS User Population Dataset. In *Proc. ACM IMC*. 165–182.

[71] Loqman Salamatian, Kevin Vermeulen, Italo Cunha, Vasilis Giotsas, and Ethan Katz-Bassett. 2024. metAScritic: Reframing AS-Level Topology Discovery as a Recommendation System. In *Proc. ACM IMC*. 337–364.

[72] Johann Schlamp, Ralph Holz, Quentin Jacquemart, Georg Carle, and Ernst W Biersack. 2016. HEAP: reliable assessment of BGP hijacking attacks. *IEEE Journal on Selected Areas in Communications* 34, 6 (2016), 1849–1861.

[73] Satadal Sengupta, Hyojoon Kim, Daniel Jubas, Maria Apostolaki, and Jennifer Rexford. 2025. Data-Plane Telemetry to Mitigate Long-Distance BGP Hijacks. *arXiv preprint arXiv:2507.14842* (2025).

[74] Pavlos Sermpezis, Vasileios Kotronis, Alberto Dainotti, and Xenofontas Dimitropoulos. 2018. A survey among network operators on BGP prefix hijacking. *ACM SIGCOMM Computer Communication Review* 48, 1 (2018), 64–69.

[75] Pavlos Sermpezis, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Danilo Cicalese, Alistair King, and Alberto Dainotti. 2018. ARTEMIS: Neutralizing BGP hijacking within a minute. *IEEE/ACM transactions on networking* 26, 6 (2018), 2471–2486.

[76] Tal Shapira and Yuval Shavitt. 2022. AP2Vec: an unsupervised approach for BGP hijacking detection. *IEEE Transactions on Network and Service Management* 19, 3 (2022), 2255–2268.

[77] Xingang Shi, Yang Xiang, Zhiliang Wang, Xia Yin, and Jianping Wu. 2012. Detecting prefix hijackings in the internet with argus. In *Proc. ACM IMC*. 15–28.

[78] Markus Stubbig. 2019. Looking Glass Command Set. RFC 8522.

[79] Yinxiang Tao, Chengwan Zhang, Changqing An, Shuying Zhuang, Jilong Wang, and Congcong Miao. 2025. Ares: Comprehensive Path Hijacking Detection via Routing Tree. In *Proc. USENIX Security*. 803–821.

[80] ISI team. 2026. ISI hitlist. https://ant.isi.edu/datasets/ip_hitlists/.

[81] Cecilia Testart, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. 2019. Profiling BGP serial hijackers: capturing persistent misbehavior in the global routing table. In *Proc. ACM IMC*. 420–434.

[82] Pierre-Antoine Vervier, Olivier Thonnard, and Marc Dacier. 2015. Mind Your Blocks: On the Stealthiness of Malicious BGP Hijacks.. In *Proc. NDSS*.

[83] Daniel Walton, Alvaro Retana, Enke Chen, and John Scudder. 2016. *Advertisement of multiple paths in BGP.* Technical Report.

[84] Zheng Zhang, Ying Zhang, Y Charlie Hu, Z Morley Mao, and Randy Bush. 2008. iSPY: Detecting IP prefix hijacking on my own. In *Proc. ACM SIGCOMM*. 327–338.

[85] Changxi Zheng, Lusheng Ji, Dan Pei, Jia Wang, and Paul Francis. 2007. A lightweight distributed scheme for detecting IP prefix hijacks in real-time. In *Proc. ACM SIGCOMM*.

A ETHICAL CONSIDERATIONS

With the demonstration that the IXPs have practical and actionable defenses against the attacks presented in our paper, following how to perform ethical measurement research[10], and from the discussions that we had with the IXP operators, we decided to publish our work. A longer discussion is available in the extended version of the paper[**?**].

B OPEN SCIENCE

We make our code and datasets publicly available while preserving IXPs anonymity at https://i62nextcloud.tm.kit.edu/index.php/s/2omctPH2yJfkpZM. A longer description of the artifacts is available on the extended version of the paper[**?**].