SIDN Labs
https://sidnlabs.nl

Author Version

Title: Unpacking Forged-Origin BGP Hijack Inferences

Authors: E. Jaw, C. Pelsser, M. Müller, L. Nieuwenhuis, C. Hesselman,

Published in: Proceedings of the 2026 22nd International Conference on Network and Service Management (CNSM)

©IFIP, (2026). This is the author's version of the work. It is posted here by permission of IFIP for your personal use. Not for redistribution. The definitive version was published in the proceedings of 2026 22nd International Conference on Network and Service Management (CNSM).

Unpacking Forged-Origin BGP Hijack Inferences

Ebrima Jaw[†], Cristel Pelsser[§], Moritz Müller[†*], Lambert Nieuwenhuis[†], Cristian Hesselman[†‡]
[†]University of Twente, Enschede, The Netherlands [*]SIDN Labs, Arnhem, The Netherlands
[‡]SURF, Utrecht, The Netherlands [§]UCLouvain, Belgium

***Abstract*—The Border Gateway Protocol remains susceptible to forged-origin BGP hijacks, in which attackers manipulate AS paths using legitimate origin ASes to evade RPKI route-origin validation. The Detection of Forged-Origin Hijacks (DFOH) system addresses this vulnerability by analyzing AS paths and inferring potential forged-origin hijacks in near real time. DFOH has operated for three years and flagged more than 29 million suspicious events, enabling the first extensive analysis of its inferences. In our study, we investigate how legitimate operational routing practices and misconfigurations contribute to these inferences, thereby providing DFOH users with guidance for judging hijack severity. We identify BGP announcement patterns that explain why certain AS links are inferred as suspicious and quantify their prevalence among DFOH alerts. In particular, we identify plausible operational explanations for 30.99% of the suspicious events, including RPKI-authorized origins and operator configuration errors. Finally, we provide recommendations to improve the reliability of forged-origin detection methods and the interpretation and management of suspicious inferences.**

***Index Terms*—BGP, forged-origin hijacking, BGP hijack detection, routing security, operational routing practices**

I. Introduction

The Border Gateway Protocol (BGP) is the Internet's interdomain routing protocol, enabling Autonomous Systems (ASes) to exchange reachability information (e.g., that 192.0.2.0/24 is reachable via the AS path {1 2 3}). However, BGP lacks built-in authentication mechanisms, allowing networks to falsely announce prefixes belonging to other ASes (prefix hijacks) [1], [2]. The Resource Public Key Infrastructure (RPKI) [3] enables operators to publish Route Origin Authorizations (ROAs) that specify which ASes are authorized to originate prefixes. Network operators use these ROAs for Route Origin Validation (ROV) to mitigate prefix hijacks and configuration errors. ROV validates the origin AS but does not protect against forged-origin hijacks [1].

Forged-origin BGP hijacks involve attackers manipulating an announcement's AS path while retaining the legitimate origin AS, thereby evading RPKI-based origin validation [1], [4]. For example, assume AS4 legitimately owns 192.0.2.0/24, and AS3 is an attacker. AS3 can falsely originate the route 192.0.2.0/24 $\rightarrow$ {1 2 **3** 4}, creating a fake AS link with AS4. Such announcements can redirect traffic intended for AS4 without triggering route-origin validation [1], [5], potentially enabling traffic interception [5]. Detecting forged-origin hijacks is particularly challenging because the origin AS remains legitimate [1], [6]. Recent incidents in operational networks demonstrate that this threat persists despite the growing adoption of RPKI [7], underscoring the need for effective monitoring systems. Holterbach et al. address this challenge with DFOH, which detects forged-origin hijacks from BGP data [1]. The system combines selected topological features with machine learning to infer suspicious forged-origin hijacks and provide interpretable alerts for operators and researchers.

DFOH has operated since 2022, providing daily inferences of suspicious new AS links and their AS paths. The accumulated observations enable us to study these inferences at Internet scale. These alerts provide valuable signals of potentially anomalous behavior but require operational assessment and investigation. To date, prior studies have not systematically quantified the extent to which operational routing practices and misconfigurations influence such inferences. We therefore focus on DFOH suspicious inferences to understand the operational factors that contribute to them and help researchers and operators interpret, prioritize, and triage suspicious events. Specifically, we investigate the question: *How do operational routing practices and operator misconfigurations contribute to forged-origin hijack alerts?* We address this question by analyzing three years of DFOH inferences and identifying those with plausible operational explanations. These findings inform the interpretation and prioritization of DFOH alerts and provide practical recommendations to strengthen forged-origin detection systems. More concretely, our contributions are:

- We conduct the first longitudinal analysis of DFOH suspicious forged-origin hijack alerts to investigate the operational factors influencing its inferences.
- We identify and quantify operational explanations for suspicious forged-origin hijack alerts and show that they explain approximately 31% of DFOH inferences.
- We provide insights and recommendations to strengthen forged-origin hijacking detection, helping researchers and operators build more robust systems.

II. BGP: Security Risks and Countermeasures

BGP enables ASes to exchange reachability information and construct interdomain routes. However, it also allows other networks to announce, accept, and propagate arbitrary routes [1], [8], thereby enabling forged-origin hijacks [1], [4], which can divert traffic while evading RPKI-ROV [1], [6].

A. Countermeasures and their Limitations

Resource Public Key Infrastructure: RPKI uses ROAs to cryptographically authorize origin ASes for IP prefixes [3]. RPKI validation provides limited protection against origin and subprefix hijacks but cannot mitigate forged-origin hijacks [1].

Autonomous System Provider Authorization (ASPA): ASPA is a new RPKI-based mechanism that validates

customer-to-provider relationships using signed ASPA objects that list each AS's authorized upstreams [9]. This validation helps prevent route leaks and some forged-origin hijacks, but it cannot prevent malicious AS-path manipulation by authorized transit providers, and its deployment remains in an early stage.

BGP Security (BGPsec): BGPsec extends RPKI by requiring each AS to cryptographically sign UPDATE messages to prevent AS-path manipulation [10], including forged-origin hijacks. However, BGPsec deployment remains negligible.

Despite advances in RPKI, ASPA, and BGPsec, limited deployment leaves the Internet vulnerable to forged-origin hijacks and highlights the need for reliable detection methods.

B. DFOH: Architecture and Overview

DFOH is a machine-learning-based system that continuously monitors BGP updates from route collectors and analyzes observed AS paths to detect forged-origin BGP hijacks [1]. We provide a high-level overview and refer the reader to the original paper for details [1]. First, DFOH compares newly observed AS links against its historical topology graph [1], [11], [12]. A new AS link is an AS adjacency (u, v) observed for the first time. For example, the AS path $\{1\ 2\ 3\ 4\}$ contains the links (1,2), (2,3), and (3,4). If DFOH has previously observed (1,2) and (2,3), but not (3,4), it treats (3,4) as a new AS link and classifies it as *legitimate* or *suspicious*.

Next, DFOH computes topological, peering, AS-path, and directionality features for each new AS link, capturing properties such as conformance with routing policies (e.g., Gao-Rexford valley-free [13]) and bidirectional propagation. DFOH then combines these features with a Random Forest classifier to distinguish legitimate and suspicious links. The model is trained on existing AS links and simulated nonexistent links as suspicious cases, and is periodically retrained using topology-balanced sampling to reduce bias toward stub ASes.

DFOH identifies an attacker AS and one or more victim ASes for each newly observed link by aggregating the announcements carrying that link. It infers attacker-victim roles at the announcement level, where the victim corresponds to the origin AS [1]. For each new link, DFOH labels the leftmost AS as the attacker. For example, given the AS path $\{1\ 2\ 3\ 4\}$ where (3,4) is the new link, DFOH labels AS3 as the attacker and AS4 as the victim, since AS4 appears as the origin AS [14]. DFOH then classifies the event as a *Type-X* hijack, where *X* denotes the attacker's position relative to the victim.

DFOH has operated since 2022, accumulating several years of suspicious inferences that enable our longitudinal analysis of these inferences and their operational explanations.

C. AS path poisoning

BGP poisoning is a *non-standardized* traffic-engineering technique in which an AS deliberately inserts selected ASNs into an announced AS path to influence route propagation [6], [15]. The inserted ASNs trigger BGP loop detection, causing the poisoned ASes to reject the route [6].

A common practice is for the poisoner to append its ASN after the poisoned ASes [6], [15]. In the example below, AS58367 reappears after AS15964, indicating that AS58367 is poisoning AS15964, whose routers should reject the route after observing their own ASN in the AS path [6], [16].

$$\underbrace{20932\ 1299\ 2914\ 3927}_{\text{upstream ASes}}\quad\underbrace{58367\ 15964\ 58367}_{\text{poisoned segment}}$$

Consequently, poisoning can create AS links that do not reflect underlying business relationships or Internet topology [17], potentially triggering suspicious DFOH inferences.

D. Terminology

We refer to each DFOH-inferred attacker-victim assignment as a *case* and each corresponding BGP announcement labeled *legitimate* or *suspicious* as an *event*. We use *AS link* instead of DFOH's original term *new edge*. A single AS link can correspond to multiple cases and events. We use *legitimate* and *suspicious* strictly as DFOH classification labels and do not interpret suspicious events as confirmed forged-origin hijacks.

III. Data Sets and Methodology

The following subsections summarize our approach to identifying operational explanations for suspicious events.

A. Identifying Suspicious and Legitimate Events

We download 1,140 daily DFOH snapshots covering 2022-01-01 to 2025-02-18, with five missing days. Each snapshot contains newly observed AS links, their inferred status (*legitimate* or *suspicious*), and the corresponding attacker and victim ASes. We also download the corresponding daily BGP announcements containing the newly observed links and AS paths associated with each case, enabling us to investigate the operational factors underlying DFOH's suspicious inferences.

Since DFOH does not annotate announcements introducing new AS links, we match its inferred cases to the daily announcement data. For each announcement, we match the reported AS link to an inferred case from the same day [1] and determine whether DFOH labels it as *legitimate* or *suspicious*.

Because DFOH can infer multiple victims for an AS link, we verify that the two ASes forming the link appear consecutively in the AS path and that the origin AS corresponds to one of the listed victims. We then use DFOH's attacker–victim metadata to identify the attacker AS and compute the attack type as the hop distance between the attacker and victim [1].

B. Attacker-Victim AS Relationship

Although DFOH incorporates topology-derived features such as AS degree, it does not explicitly use inferred AS business-relationship data. We therefore investigate whether these relationships can explain some suspicious events. We hypothesize that attacker-victim pairs in suspicious events should not have an existing business relationship when the inference is made [18], [19]. Using CAIDA's monthly AS-relationship datasets [19], we determine whether attacker-victim pairs involved in suspicious events from 2022 to 2025 have customer-provider or peering relationships. For each event, we use the corresponding monthly dataset to determine

whether an attacker-victim relationship existed at the time of inference. Section §IV-B presents our findings.

C. ROA Covering Attacker ASes

Legitimate routing practices, including customer-provider relationships, sibling ASes, and multi-homing, may produce announcements that resemble forged-origin hijack behavior. DFOH does not use RPKI information when classifying new links. We therefore investigate whether DFOH-inferred attacker ASes are authorized to originate the affected prefixes. We parse ROAs from daily RPKI snapshots across all five RIRs (January 2022-February 2025) [20] and build IPv4 and IPv6 prefix tries to determine whether each attacker AS is authorized to originate the affected prefix using the corresponding daily ROAs. Because daily snapshots may not exactly match the ROAs in effect when an event occurred, our validation reflects the closest available RPKI state. We classify each attacker-victim pair by its RPKI validation state: **Valid** when the attacker AS is authorized to originate the affected prefix, **Invalid** when a covering ROA exists but does not authorize the attacker AS, and **NotFound** when no covering ROA exists. Table II summarizes the results.

D. BGP Poisoning Incidents

BGP poisoning is widely used for traffic engineering, DDoS mitigation, censorship evasion, and Internet measurement [6], [21]. We analyze DFOH-flagged suspicious events by examining repetitions of attacker and victim ASes in poisoned AS paths to distinguish patterns consistent with documented operational practices [21] from ambiguous poisoning events.

Attacker-Victim Patterns in Poisoned AS Paths: We identify a recurring pattern in which DFOH infers the same AS as both the attacker and the victim (Pattern I: *attacker equals victim*).

$$\underbrace{29049\,12389\,4766\,4775}_{\text{upstream ASes}}\ \underbrace{14499}_{\text{attacker}}\ \overbrace{64254}^{\text{poisoned AS}}\ \underbrace{14499}_{\text{victim}} \quad \text{(I)}$$

This pattern is consistent with documented traffic-engineering practices to influence route selection or suppress upstream paths during maintenance or failures [6], [15], [17]. The resulting adjacency (14499-64254) deviates from DFOH's previously observed topology and triggers a suspicious inference. Although such paths could support forged-origin attacks [1], [16], the repeated ASN creates a longer and less attractive route. The documented poisoning pattern and DFOH's inference of the same AS as both attacker and victim provide evidence of a plausible operational explanation for Pattern (I).

We also observe repetition patterns involving the attacker AS, victim AS, or both. Although these repetition patterns are consistent with documented BGP poisoning practices [6], the same patterns can result from forged AS-path manipulation [5], [16]. We therefore treat all patterns other than Pattern (I) as ambiguous and further evaluate whether the inferred attacker and victim have an existing business relationship using the methodology in Section §III-B. Section §IV-D presents the results.

E. Fat-Finger Errors

DFOH removes repeated ASNs caused by prepending [22] before analyzing AS paths, but this normalization can introduce ambiguity when operators make typographical errors (fat-fingers) [8], [22]. For example, assume AS64496 originates 192.0.2.0/24 and prepends itself three times with a typographical error: (192.0.2.0/24 → [64515 **64469** 64496 64496 64496]). DFOH removes the repeated ASNs and treats AS64469 as distinct, creating a previously unseen AS link. Since AS64469 differs from AS64496 by a single edit, it may represent a typographical error. We use the Optimal String Alignment (OSA) distance, which considers insertions, deletions, substitutions, and adjacent transpositions [22], to identify single-edit deviations in suspicious DFOH inferences.

We restrict this analysis to Type-2 hijacks, which account for 22.33M (76.93%) of suspicious events and contain a single intermediate AS (X) between the inferred attacker (A) and victim (V) [1], [8]. This structure (A, X, V) lets us attribute a single-edit deviation to a specific ASN, whereas other hijack types involve multiple intermediate ASes. We exclude deviations requiring multiple edits, such as `266954` ↔ `269654`, because they cannot be attributed to a single-edit typo [22].

F. DDoS Protections

Benign on-demand DDoS mitigation can temporarily make a scrubbing provider appear as an upstream while cleaning and rerouting traffic for the victim [23], [24]. This can introduce previously unseen links between the scrubber and victim AS that forged-origin detectors such as DFOH may classify as suspicious [1]. We obtain all 100 ASes identified by **bgp.tools** [25] as DDoS-mitigation providers and examine whether they contribute to DFOH's suspicious inferences. For each suspicious event, we identify cases where DFOH infers a known provider as the attacker and the provider appears in the new AS link. We then count suspicious events for each provider and across all providers to quantify the contribution of potential DDoS mitigation to DFOH's suspicious events. Such events may provide a plausible operational explanation involving DDoS scrubbing rather than evidence of confirmed forged-origin hijacking. Section §IV-F presents our findings.

IV. Results

The following subsections present the results of our analysis and identify operational explanations for suspicious events.

A. Overview: DFOH's New Links and Cases

Table I summarizes the results of Section §III-A. From January 2022 to February 2025, DFOH inferred 73K suspicious AS links, compared with 472K legitimate links. These links appeared in 29.0M suspicious and 24.7M legitimate events, resulting in 377K suspicious and 485K legitimate cases. After AS-path deduplication, suspicious events decreased from 29.0M to 11.6M unique AS-path events, whereas legitimate events decreased from 24.7M to 23.5M (Table I). This indicates substantially greater repetition among suspicious events. Finally, suspicious cases involved ≈11.2K unique attackers

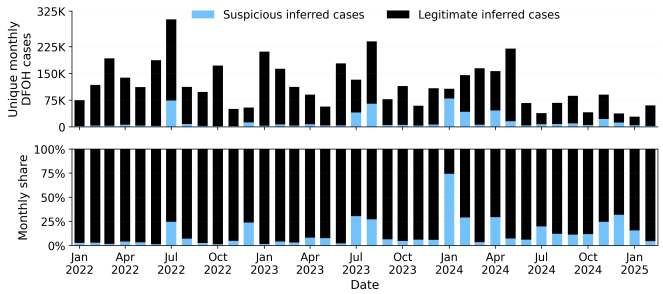


Fig. 1: DFOH monthly suspicious and legitimate inferences. A few months exhibit pronounced peaks, with ≈75% of the January 2024 spike attributed to the AS link (59105, 209306).

and 88.8K victims, indicating that suspicious inferences were concentrated among fewer attacker than victim ASes.

TABLE I: Statistics of DFOH's daily new AS links and cases

Start Date: 2022-01-01	End Date: 2025-02-18 (1.14K daily snapshots)			
	Suspicious Cases		Legitimate Cases	
Total (unique) events	29.0M (11.6M)		24.7M (23.5M)	
Total cases	377.0K		484.8K	
Total new AS links	73.0K		472.0K	
Unique DFOH victims	88.8K		95.1K	
Unique DFOH attackers	11.2K		14.8K	
	IPv4	IPv6	IPv4	IPv6
Unique prefixes	320.4K	116.8K	547.4K	132.3K
Total prefixes	17.8M	11.2M	20.2M	4.5M
Unique peer IPs	716	613	719	258
Total peer IPs	18.2M	10.8M	20.4M	4.3M

DFOH Inferred Cases: We deduplicate events by attacker-victim pair to avoid overcounting cases observed from multiple route collector peers. Figure 1 shows the resulting monthly counts. Overall, DFOH inferred an average of 14.28K suspicious and 103.22K legitimate cases per month, corresponding to ≈12% and 88% of monthly inferences, respectively.

The following subsections focus on identifying plausible operational explanations for DFOH suspicious cases.

B. Suspicious Events with Existing AS Relationships

Most DFOH-inferred suspicious events involved attacker-victim pairs with no inferred business relationship at the time of the event. However, ≈17.32% of the 29.0M suspicious events involved pairs with an existing peering or provider-customer relationship. Peering relationships accounted for most of these events, with provider-customer relationships representing a much smaller fraction. Similarly, pairs with an inferred business relationship accounted for 17.11% of the 11.6M unique suspicious events and ≈17.50% of the 73K unique suspicious AS links.

These findings show that a nontrivial fraction of suspicious inferences occurred between ASes with existing business relationships, providing important operational context. Consequently, incorporating AS-relationship information into forged-origin hijack monitoring systems (e.g., DFOH) could improve their ability to reliably identify suspicious events.

C. RPKI-Authorized Suspicious Events

Table II shows that most inferred attackers had invalid RPKI status for the affected prefixes, consistent with DFOH identifying them as potential attackers. However, 2.1M (≈7.20%) suspicious events had valid ROAs authorizing the inferred attackers to originate the affected prefixes. These 2.1M events involved 2.21K attackers and 5.75K affected prefixes. This finding shows that RPKI information provides important operational context for interpreting suspicious inferences [26].

TABLE II: RPKI status summary for DFOH cases

Start Date: 2022-01-01 End Date: 2025-02-18				
	Suspicious Cases		**Legitimate Cases**	
	Victims	Attackers	Victims	Attackers
Valid	17.8M (61.7%)	2.1M (7.2%)	12.3M (50.1%)	425K (1.7%)
NotFound	10.5M (36.2%)	10.5M (36.2%)	11.8M (48.1%)	11.8M (48.1%)
Invalid	605K (2.1%)	16.3M (56.6%)	443K (1.8%)	12.3M (50.2%)
Unique *prefix RPKI status* pairs for DFOH attacker and victim ASes				
Valid	221K (47.5%)	6K (1.2%)	337K (44.5%)	6K (0.8%)
NotFound	232K (49.9%)	232K (50.1%)	406K (53.6%)	406K (53.8%)
Invalid	13K (2.7%)	226K (48.7%)	14K (1.9%)	342K (45.4%)

Among these 2.1M events, only ≈10.23% involved attacker-victim pairs with inferred business relationships, primarily provider-customer relationships, whereas ≈89.90% had no inferred business or peering relationship. The large proportion without inferred relationships suggests that conventional AS-relationship inference may miss operational arrangements underlying valid-ROA events, motivating further examination of the responsible ASes.

AS44570 (Xodo Technology) [27] accounted for 891.97K (≈42.78%) of the 2.1M valid-ROA events and 873.92K (≈46.70%) of the 1.87M without inferred relationships. The latter represented ≈98% of events where AS44570 was the inferred attacker. AS44570 provides IP transit and leasing with RPKI and LOA/ROA support [27], potentially explaining its authorization to originate the affected prefixes. All affected prefixes were IPv6 and involved 195 victim ASes, primarily small hobby or educational networks likely routed through AS44570 without formal business or peering relationships. These observations are consistent with delegated BGP support and related operational practices associated with its transit [28] and IP-leasing services [27], providing a plausible operational explanation for these suspicious inferences.

D. Interpreting BGP Poisoning-induced Events

Table III summarizes the poisoning-induced suspicious events. We identified 413.56K such events, accounting for ≈1.43% of all suspicious DFOH-inferred events. These events involved 902 attacker ASes, ≈2K victim ASes, and 5.26K affected prefixes. Of these, 277.77K (≈67.16%) matched Pattern (I), where DFOH inferred the same AS as both attacker and victim. We analyzed the remaining 135.79K events using the AS-relationship methodology in Section §III-B.

Of the remaining 135.79K events, 65.92K (48.55%) had an inferred peering or business relationship between the attacker and victim. Combined with the 277.77K Pattern (I) events, these account for 343.69K (83.11%) operationally explainable poisoning-induced events. These findings show that most poisoning-induced suspicious events have plausible operational explanations consistent with documented BGP poisoning practices [5], [6], [15], [17].

TABLE III: Summary of suspicious poisoning-induced events.

	Poisoning Events	AS Relationship
(I) Attacker is victim	277.77K (67.16%)	–
(II) Remaining patterns	135.79K (32.84%)	65.92K (48.55%)
Operationally explainable	343.69K (83.11%)	
Remaining events	69.88K (16.89%)	
Total poisoning events	413.56K (100%)	

DFOH inferred three researcher-operated ASes (AS3927, AS58367, and AS10233) as attackers in 141.90K (≈34.30%) poisoning-induced events. The operators confirmed that these events were research-driven BGP poisoning experiments aimed at identifying backup links [21]. These experiments created previously unseen AS links as part of confirmed research activity, demonstrating that BGP experiments can trigger DFOH's suspicious inferences [1].

E. Fat-Finger–Induced Suspicious Events

We identified 676 unique one-edit incidents resulting in 128.19K suspicious events (≈0.44% of all DFOH-inferred suspicious events). Insertions and substitutions accounted for most events, with 45.58K (≈35.56%) and 42.14K (≈32.87%), respectively. Deletions accounted for 23.25K (≈18.14%) and transpositions for 17.23K (≈13.44%). Figure 2 shows that substitutions concentrate around adjacent digit pairs (e.g., 1→2, 5→6, and 9→8), a pattern consistent with typographical errors [22]. These findings indicate that operator mistakes can trigger DFOH suspicious inferences.

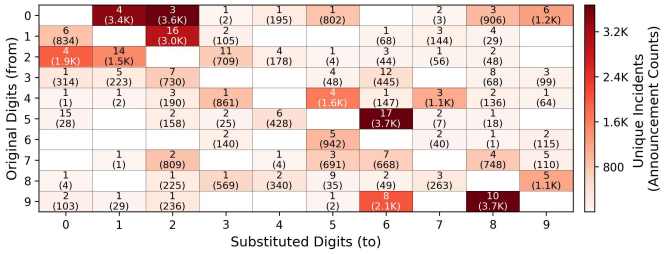


Fig. 2: Unique incidents and announcement counts for single-edit substitutions. Only 25.96% (478) of vantage points observed these events, while 74.04% (1.4K) did not, suggesting they are largely localized, possibly due to misconfiguration.

F. Suspicious Events Involving DDoS-Mitigation Providers

We identified 2.03M suspicious events in which DFOH inferred a DDoS mitigation provider as the attacker, corresponding to ≈7.00% of all suspicious events. Of the 100 providers considered, 92 appeared as inferred attackers, with the top six accounting for 1.06M (52.01%) events. AS1299 and AS3223 contributed 250.3K (≈12.31%) and 222.0K (10.92%), respectively, while AS59796, AS137409, AS2914, and AS60068 collectively contributed another 28.78%. The remaining 86 providers accounted for 47.99%, showing that these events extend beyond the dominant providers.

Among the 2.03M suspicious events involving DDoS mitigation providers, 1.72M (84.76%) involved attacker-victim pairs with no inferred business relationship, while the remaining 15.24% involved pairs with a provider-customer or peering relationship. The predominance of pairs without inferred relationships is consistent with on-demand DDoS mitigation, where a provider may appear in the AS path only when customer traffic is redirected for scrubbing. Prior work similarly shows that DDoS scrubbers can resemble benign hijackers [2], [29], suggesting that DDoS mitigation may contribute to suspicious DFOH inferences.

Finally, we expect DDoS mitigation events to exhibit broad visibility because effective mitigation requires attracting traffic from many networks. We therefore quantify visibility as the fraction of vantage points observing at least one corresponding announcement and find that ≈51.77% observed these events. This broad visibility is consistent with the large-scale traffic attraction expected during DDoS mitigation.

G. Attacker-Victim Pairs Between Sibling ASes

We use CAIDA's AS-to-Organization dataset to identify DFOH-inferred attacker-victim pairs operated by the same organization at the time of each event. We identify 616.5K suspicious events and 1.1K suspicious new links involving sibling ASes (Table IV). Of these, 326.1K events were identified exclusively by the sibling-AS heuristic and not by any of the other five heuristics. For example, DFOH inferred AS38255 as an attacker against more than 170 victim ASes operated by China Education and Research Network (CERNET), including AS142087, AS142088, AS142089, and AS142090. These findings show that missing organizational ownership can cause sibling AS relationships to be misinterpreted as suspicious.

H. Operationally Explained Inferences

Table IV summarizes the suspicious events and new links identified by our heuristics. Overall, the heuristics identified 9.0M (30.99%) suspicious events and 21.40K (29.32%) suspicious new links as operationally explainable. Of the operationally explainable events, 87.8% were identified by exactly one heuristic and 12.2% by multiple heuristics. Consequently, the per-heuristic counts in Table IV are not mutually exclusive.

We examine how individual heuristics change over time. RPKI-based explanations peak at 30.51% of suspicious events in May 2022 and decline to 2.44% in August 2023, with BGP poisoning showing a similar trend. This suggests that some operationally explainable suspicious events become less prevalent over time, potentially reflecting increasing RPKI deployment [30] and periodic DFOH retraining [1].

TABLE IV: Operationally explainable suspicious events

	Suspicious events	Suspicious new links
AS Relationship	5.0M (17.3%)	12.53K (17.2%)
RPKI-Based	2.1M (7.2%)	5.4K (7.4%)
Same Organization	616.5K (2.1%)	1.1K (1.6%)
DDoS Scrubbers	2.0M (7.0%)	4.3K (5.9%)
BGP Poisoning	343.69K (1.2%)	2.5K (3.4%)
Fat-finger Errors	128.19K (0.44%)	684 (0.94%)
Breakdown of operationally explainable inferences (unique events)		
Single method	7.9M (27.1%)	18.65K (25.6%)
Multiple methods	1.1M (3.9%)	4.3K (5.8%)
Operationally explainable	**9.0M (30.99%)**	**21.40K (29.32%)**
Remaining inferences	20.0M (69.01%)	51.59K (70.68%)

Overall, our heuristics identify plausible operational explanations for nearly one-third of DFOH-inferred suspicious events and new links, providing context for their interpretation and prioritization. Our heuristics did not identify the remaining 69.01% of suspicious events, and they should not be interpreted as confirmed forged-origin hijacks.

V. Related Work

Unlike forged-origin BGP hijacks [1], mis-origin BGP hijacks have been extensively studied [2], [22], [29], [31], [32]. In 2018, Testart et al. [29] coined the term "serial hijackers" to refer to ASes that repeatedly hijack address blocks from other networks and classified them using machine learning. They showed that DDoS scrubbers could mislead their classifier [29]. The authors of [2] extended this work and found that a non-negligible fraction of the alleged serial hijackers were still active on the Internet. However, neither study quantified the operational factors that may explain suspicious inferences.

Schulmann et al. [26] fill this gap by introducing Learning Origin Validation (LOV) to quantify benign root cases for RPKI-invalids. They reported that about 79% of the daily RPKI invalid announcements are benign. Cho et al. [22] combined heuristic methods with AS hegemony to classify BGP hijacks into typos, prepending mistakes, AS path forging, and origin changes. They demonstrate how standard routing practices, such as AS path prepending, and human errors can lead to events that resemble hijacks. However, none of these studies examined how operational practices affect detection of forged-origin BGP hijacks. Chen et al. proposed a prototype anomaly detection method (BEAM) that models the unique properties, so-called *routing roles*, of ASes to detect anomalies such as origin changes and forged-origin hijacks [33].

Giaconia et al. [34] demonstrated that public BGP data can be manipulated in ways that affect topology-based detection metrics used by systems such as BEAM [33] and DFOH [1]. In contrast, we examine how operational routing practices and misconfigurations, such as BGP poisoning experiments and typographical errors, influence DFOH suspicious inferences, and how information such as RPKI-authorized origins can provide additional context for interpreting those inferences.

VI. Discussion, Implications, and Limitations

The increasing complexity of the inter-domain routing ecosystem [35], driven by diverse operational, economic, and business practices [30], [36], [37], continues to challenge reliable BGP hijack detection [1], [29], [31]. Our results show that operational factors explain $\approx 31\%$ of DFOH-inferred suspicious events and 29.32% of suspicious new links. These findings indicate that while previously unseen AS links help identify anomalous behavior, they are insufficient to distinguish malicious activity from known operational practices.

Our findings also highlight the limitations of relying solely on publicly observable routing information. For example, a July 2022 remote IXP connection led DFOH to infer numerous suspicious links and attacker-victim pairs [38]. Similarly, CDNs, Anycast deployments, and DDoS mitigation providers can announce the same prefixes from multiple locations using customized routing policies [39]. Specifically, DFOH inferred AS26415 as an attacker against 51 VeriSign-operated Anycast ASes, including AS396622 and AS396623, and also inferred suspicious attacker-victim pairs among sibling ASes operated by the same organization (Section IV-G). These examples show how operational practices and relationships not captured in public routing data can create previously unseen AS links that trigger topology-based detection systems such as DFOH.

Accordingly, incorporating business relationships, RPKI-authorized origins, measurement-driven practices such as BGP poisoning, and misconfigurations such as fat-finger errors could improve forged-origin hijack detection and the interpretation of suspicious inferences. Research-driven BGP experiments are particularly relevant because they can introduce unusual AS paths and previously unseen links that resemble forged-origin hijacks. Accounting for such behavior could help distinguish unconventional routing practices from malicious activity and improve detection reliability.

Finally, our analysis relies primarily on publicly available BGP data, which provides only a partial view of operational relationships and Internet topology. Remote peering, private interconnections, and short-lived business relationships may therefore remain unobserved. Moreover, without ground truth, we cannot conclusively classify suspicious events as benign or malicious. Thus, our heuristics identify plausible operational explanations rather than confirmed benign activity.

VII. Conclusions and Recommendations

We analyzed three years of DFOH-inferred suspicious forged-origin hijack events to investigate how operational routing practices and misconfigurations influence its inferences. Our analysis shows that operational factors explain 30.99% of suspicious events and 29.32% of suspicious new links. These results demonstrate that previously unseen AS links, while valuable indicators of anomalous behavior, are insufficient to distinguish malicious activity from operational routing practices. Incorporating business relationships, organizational ownership, RPKI-authorized origins, and operator misconfigurations can provide the operational context needed

to assess suspicious inferences, thereby improving the effectiveness and reliability of forged-origin hijack detection. More broadly, robust forged-origin hijack detection requires both effective algorithms and greater transparency from researchers and operators about routing policies, operational relationships, and experimental measurement activities.

REFERENCES

[1] T. Holterbach, T. Alfroy, A. Phokeer, A. Dainotti, and C. Pelsser, "A system to detect forged-origin bgp hijacks," in *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. Santa Clara, California, USA: USENIX Association, 2024, pp. 1751–1770.

[2] E. Jaw, M. Müller, C. Hesselman, and L. Nieuwenhuis, "Serial bgp hijackers: A reproducibility study and assessment of current dynamics," in *2024 8th Network Traffic Measurement and Analysis Conference (TMA)*. Dresden, Germany: IEEE, 2024, pp. 1–10.

[3] R. Bush, "Origin validation operation based on the resource public key infrastructure (RPKI)," Internet Engineering Task Force, Request for Comments RFC 7115, 2014-01, num Pages: 11. [Online]. Available: https://datatracker.ietf.org/doc/rfc7115

[4] J. Furuness, A. Herzberg, and R. Morillo. (2025, January) Securing BGP ASAP: ASPA and other Post-ROV Defenses.

[5] H. Birge-Lee, Y. Sun, A. Edmundson, J. Rexford, and P. Mittal, "Bamboozling certificate authorities with BGP," in *27th USENIX Security Symposium (USENIX Security 18)*, 2018, pp. 833–849. [Online]. Available: https://tinyurl.com/3d6uztdz

[6] J. M. Smith, K. Birkeland, T. McDaniel, and M. Schuchard, "Withdrawing the BGP re-routing curtain: Understanding the security impact of BGP poisoning through real-world measurements," in *Proceedings 2020 Network and Distributed System Security Symposium*. San Diego, CA: Internet Society, 2020, pp. 1–15.

[7] B. Herdes, B. Walters, and M. Zhang. Enforcing the first AS in BGP AS PATHs. Cloudflare. [Online]. Available: https://blog.cloudflare.com/enforce-first-as-bgp/

[8] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti, "ARTEMIS: Neutralizing BGP Hijacking Within a Minute," *IEEE/ACM Transactions on Networking*, vol. 26, no. 6, pp. 2471–2486, Dec. 2018. [Online]. Available: https://ieeexplore.ieee.org/document/8481500

[9] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram, "BGP AS_path verification based on autonomous system provider authorization (ASPA) objects," Internet Engineering Task Force, Internet Draft draft-ietf-sidrops-aspa-verification-24, 2025, num Pages: 21.

[10] M. Lepinski and K. Sriram, "BGPsec protocol specification," Internet Engineering Task Force, Request for Comments RFC 8205, 2017-09, num Pages: 45. [Online]. Available: https://tinyurl.com/2wu72yk6

[11] RIPE NCC, "Routing Information Service (RIS)," https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris, 2023.

[12] University of Oregon. (2023) Routeviews project. [Online]. Available: https://www.routeviews.org/routeviews/

[13] L. Gao and J. Rexford, "Stable internet routing without global coordination," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 681–692, 2001-12. [Online]. Available: https://tinyurl.com/yerxkwnj

[14] T. Holterbach, T. Alfroy, A. Phokeer, A. Dainotti, and C. Pelsser. (2024) Forged_origin_hijacks_detection/runner. Santa Clara, California, USA. [Online]. Available: https://tinyurl.com/nhrday2m

[15] E. Katz-Bassett, C. Scott, D. R. Choffnes, Í. Cunha, V. Valancius, N. Feamster, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy, "LIFEGUARD: practical repair of persistent route failures," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 4, pp. 395–406, 2012.

[16] H. Birge-Lee, L. Wang, J. Rexford, and P. Mittal, "SICO: Surgical interception attacks by manipulating BGP communities," in *Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security*. London United Kingdom: ACM, 2019-11-06, pp. 431–448. [Online]. Available: https://tinyurl.com/yn3csr8k

[17] M. Luckie, "Spurious routes in public BGP data," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 3, pp. 14–21, 2014. [Online]. Available: https://dl.acm.org/doi/10.1145/2656877.2656880

[18] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy, and G. Riley, "AS relationships: inference and validation," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 1, pp. 29–40, 2007-01-22. [Online]. Available: https://doi.org/10.1145/1198255.1198259

[19] CAIDA. (2013) AS relationships. [Online]. Available: https://www.caida.org/catalog/datasets/as-relationships/

[20] RIPE NCC. (2024) Rpki snapshots/. [Online]. Available: https://ftp.ripe.net/rpki/

[21] M. Tashiro. (2025) Backup-links experiments. [Online]. Available: https://github.com/m-appel/backup-links/blob/main/README.md

[22] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill, "BGP hijacking classification," in *2019 Network Traffic Measurement and Analysis Conference (TMA)*. Paris, France: IEEE, Jun. 2019, pp. 25–32. [Online]. Available: https://ieeexplore.ieee.org/document/8784511/

[23] Akamai. (2025) DDoS attack protection and mitigation. [Online]. Available: https://www.akamai.com/products/prolexic-solutions

[24] Radware, "DDoS protection: Techniques, types & 7 solutions to know in 2024," https://tinyurl.com/bddbe4ek.

[25] B. Cartwright-Cox. (2025) bgp.tools tagged networks with ddos mitigations appliances. [Online]. Available: https://bgp.tools/tags/ddosm

[26] H. Schulmann and S. Zhao, "Learning to identify conflicts in RPKI," in *Proceedings of the 20th ACM Asia Conference on Computer and Communications Security*, ser. ASIA CCS '25. New York, NY, USA: Association for Computing Machinery, 2025, pp. 1490–1505.

[27] Xodo. (2025) IP space-xodo technology, LLC: Enterprise network solutions. [Online]. Available: https://www.xodo.net/ip-space.php

[28] Xodo. (2025) IP transit-xodo technology, LLC: Enterprise network solutions. [Online]. Available: https://www.xodo.net/ip-transit.php

[29] C. Testart, P. Richter, A. King, A. Dainotti, and D. Clark, "Profiling BGP serial hijackers: Capturing persistent misbehavior in the global routing table," in *Proceedings of the Internet Measurement Conference*, ser. IMC '19. New York, NY, USA: Association for Computing Machinery, 2019-10-21, pp. 420–434.

[30] W. Li, Z. Lin, M. I. Ashiq, E. Aben, R. Fontugne, A. Phokeer, and T. Chung, "RoVista: Measuring and analyzing the route origin validation (ROV) in RPKI," in *Proceedings of the 2023 ACM on Internet Measurement Conference*. Montreal QC Canada: ACM, 2023, pp. 73–88. [Online]. Available: https://tinyurl.com/3hkvyvm4

[31] L. Qin, D. Li, R. Li, and K. Wang, "Themis: Accelerating the detection of route origin hijacking by distinguishing legitimate and illegitimate moas," in *Proceedings of the 31st USENIX Security Symposium (USENIX Security 22)*. Boston, MA: USENIX Association, 2022, pp. 4509–4524. [Online]. Available: https://tinyurl.com/2k3x34c9

[32] M. Wählisch, O. Maennel, and T. C. Schmidt, "Towards detecting BGP route hijacking using the RPKI," *ACM SIGCOMM Computer Communication Review*, vol. 42, no. 4, pp. 103–104, Aug. 2012. [Online]. Available: https://doi.org/10.1145/2377677.2377702

[33] Y. Chen, Q. Yin, Q. Li, Z. Liu, K. Xu, Y. Xu, M. Xu, Z. Liu, and J. Wu, "Learning with semantics: Towards a semantics-aware routing anomaly detection system," in *Proceedings of the 33rd USENIX Security Symposium (USENIX Security '24)*. Boston, MA: USENIX Association, 2024, pp. 5143–5160.

[34] A. Giaconia, M. Tran, L. Vanbever, and S. Vissicchio, "Is crunching public data the right approach to detect BGP hijacks?" 2025-07-27. [Online]. Available: http://arxiv.org/abs/2507.20434

[35] M. Caesar and J. Rexford, "Bgp routing policies in isp networks," *IEEE Network*, vol. 19, no. 6, pp. 5–11, Nov. 2005. [Online]. Available: https://ieeexplore.ieee.org/document/1541715

[36] N. Rodday, Í. Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, and M. Wählisch, "Revisiting rpki route origin validation on the data plane," in *IFIP Network Traffic Measurement and Analysis Conference (TMA)*. Virtual Event: IFIP, 2021, pp. 1–8.

[37] T. Hlavacek, H. Shulman, N. Vogel, and M. Waidner, "Keep your friends close, but your routeservers closer: Insights into RPKI validation in the internet," 2023. [Online]. Available: http://arxiv.org/abs/2303.11772

[38] NANOG. (2022) nanog: Re: What's going on with AS147028? [Online]. Available: https://seclists.org/nanog/2022/Jul/118

[39] Cloudflare. (2025) Cloudflare magic transit. [Online]. Available: https://developers.cloudflare.com/magic-transit/

VIII. APPENDIX

A. Ethics

This study uses only publicly available BGP and DFOH inference datasets and raises no ethical concerns. We contacted the identified AS operators, who confirmed the BGP poisoning experiments and provided the experiment website.