International Journal of Network Management

WILEY

SPECIAL ISSUE PAPER OPEN ACCESS

Mitigating BGP Route Leaks With Attributes and Communities: A Stopgap Solution for Path Plausibility

Nils Höger | Nils Rodday | Gabi Dreo Rodosek

Research Institute CODE, Universität der Bundeswehr München, Bavaria, Germany

Correspondence Nils Höger (NilsHoeger@pm.me)

Received: 9 September 2024 | **Revised:** 24 December 2024 | **Accepted:** 9 January 2025

Funding: The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the program "Souverän. Digital. Vernetzt.". Joint project 6G-life, project identification number: 16KISK002.

Keywords: autonomous system provider authorization | border gateway protocol | down only | only to customer | path plausibility | routing security

ABSTRACT

The Border Gateway Protocol (BGP) is known to have serious security vulnerabilities. One of these vulnerabilities is BGP route leaks. A BGP route leak describes the propagation of route announcements beyond their intended scope, violating the Gao-Rexford model. Route leaks may lead to traffic misdirection, causing performance issues and potential security risks, often due to mistakes and misconfiguration. Several potential solutions have been published and are currently greatly discussed within the Internet Engineering Task Force (IETF) but have yet to be widely implemented. One approach is the Autonomous System Provider Authorization (ASPA). In addition to these new approaches, there are also efforts to use existing BGP functionalities to detect and prevent route leaks. In this paper, we implement the Down Only (DO) Community and Only to Customer (OTC) Attribute approaches, using them isolated and in conjunction with ASPA. Our research indicates that implementing a DO/OTC deployment strategy focusing on well-interconnected ASes could significantly reduce route leaks. Specifically, we observed mitigation of over 98% of all route leaks when DO and OTC were deployed by the top 5% of the most connected ASes. We show that combining DO/OTC and ASPA can greatly enhance ASPA's route leak prevention capabilities.

1 | Introduction

The Border Gateway Protocol (BGP) [1] allows Autonomous Systems (AS) to exchange reachability and routing information with each other. Despite its wide usage, the protocol still has significant security flaws and is vulnerable to both intentional and unintentional misconfiguration [2]. The most prominent vulnerabilities are BGP route leaks [3, 4] and BGP prefix hijacking [2].

Route leaks represent policy violations in which ASes deviate from the expected behavior outlined in the Gao-Rexford model [5] by incorrectly forwarding a route received from a provider or peer to another provider or peer. As providers prefer updates sent by their customers for financial reasons, this misstep can cause the leaked route to gain higher priority at the provider, potentially overwhelming the leaking AS with traffic and rendering the original BGP announcement's source unreachable. Furthermore, this breach of policy can create a financially unfavorable scenario for the leaking AS [6]. BGP prefix hijacking occurs when a malicious AS falsely advertises ownership of IP prefixes, redirecting traffic intended for the legitimate owner to itself [2].

Several solutions have been proposed to address BGP's vulnerabilities, such as the path validation algorithms Pretty Secure

Abbreviations: AS, Autonomous System; ASPA, Autonomous System Provider Authorization; BGP, Border Gateway Protocol; BGPsec, Border Gateway Protocol Security; DO, Down Only; IETF, Internet Engineering Task Force; OTC, Only to Customer; psBGP, Pretty Secure BGP.

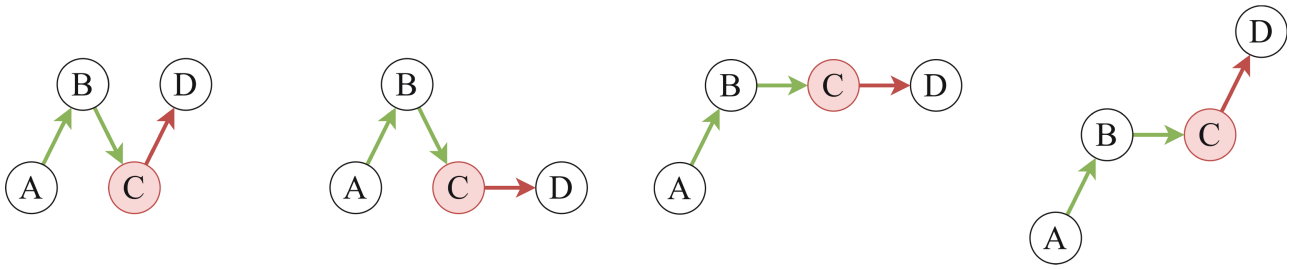

(a) Provider to provider leak (b) Provider to peer leak (c) Peer to peer leak (d) Peer to provider leak

FIGURE 1 | Possible route leak scenarios.

BGP (psBGP) [7] in 2007 or Border Gateway Protocol Security (BGPsec) [8] in 2017. While the solutions protect against numerous vulnerabilities, they also have many disadvantages that severely limit performance and usability. For this reason, there has been no widespread use of such approaches to date. An alternative to the established but little-used approaches is path plausibility algorithms. Their pros and cons are the subject of ongoing discussions within the Secure Inter-Domain Routing Operations (SIDROPS) [9] working group of the Internet Engineering Task Force (IETF) [10]. One of these algorithms is Autonomous System Provider Authorization (ASPA) [11]. While BGPsec proves the integrity of an AS path cryptographically through concatenated signatures, ASPA enables ASes to check the plausibility of a path by generating and checking cryptographically signed objects in the Resource Public Key Infrastructure (RPKI). The RPKI is an existing globally distributed database standardized by the IETF [12, 13]. Unlike BGPsec, ASPA does not need cryptographic signature validation during the BGP decision-making process. The validation can be outsourced to dedicated systems, which reduces the computational effort for routers. ASPA is a fast and effective method for detecting route leaks and forged origin hijacks, but it offers weaker security guarantees than BGPsec [14].

In addition to the two fields of path plausibility and path validation algorithms, there are also other efforts to detect and prevent route leaks, even across multiple hops. Here, we look at two approaches: the BGP Only to Customer (OTC) [15] Attribute and the BGP Down Only (DO) Community [16]. Both approaches mark BGP updates based on the relationship between the sender and the next hop. Such behavior makes it possible to recognize if a message leaves its intended scope and a route leak occurs, even multiple hops away [15].

1.1 | Contributions

This work builds upon our earlier findings [17] regarding the effectiveness of preventing route leaks using different deployment approaches for ASPA and AS-Cones, another path plausibility algorithm. In this paper, we assess the implementation and effectiveness of tagging BGP messages with OTC and DO as independent deployment methods and in conjunction with ASPA. Specifically, our contributions are as follows:

1. We extend the Python-based BGP Security Simulations (BGPsec Sim) [18] to support marking BGP messages according to the OTC and DO ingress and egress rules.
2. We use random and selective deployment methods and conclude which deployment strategy works best against route leaks.
3. We combine the marking of BGP messages with the partial deployment of ASPA to showcase how a combined deployment could reduce the risk of route leaks.

This paper is structured as follows: Section 2 provides the background of our work, including ASPA, BGP roles, DO, and OTC. Section 3 presents related work and Section 4 details our methodology. Section 5 presents our results for route leak scenarios with sole DO/OTC deployment. Section 6 presents our previous results on deploying ASPA in different deployment strategies, and Section 7 presents our newly obtained results for combined deployment with ASPA. Finally, we offer a brief discussion in Section 8, while Section 9 summarizes our work.

2 | Background

This section explains BGP route leaks and the Gao–Rexford model. It also explains the functionality of ASPA and BGP Roles and how DO and OTC work.

2.1 | BGP Route Leaks

BGP traffic follows business relationships in interdomain routing as providers charge their customers for transit. Since a business is required to have a positive cash flow to sustain its operation, the routing should align with these relationships [19]. Violating policies or these relationships can result in misrouted traffic and financial loss and affect the global stability and reliability of the internet [20, 21]. The Gao–Rexford model [5] defines three rules for valley-free routing to prevent these violations:

1. If an AS receives a route from a customer, it forwards it to all peers regardless of the relationship.
2. If an AS receives a route from a peer, it may only forward it to its customers.
3. If an AS receives a route from a provider, it may only forward it to its customers.

Based on these rules, there are four different route leak scenarios, which RFC7908 [21] also describes in detail:

1. An AS learns a route from one provider and forwards it to another provider.
2. An AS forwards routes that it has learned from its provider to a peer.
3. An AS forwards routes that it has learned from a peer to a peer
4. An AS learns the route from a peer and forwards it to its provider.

Route leaks can occur due to intentional and unintentional misconfigurations. The majority of route leak cases can be attributed to human errors [20]. Figures 1a–d shows the different scenarios.

2.2 | ASPA

ASPA [11] is a path plausibility algorithm that can mitigate route leaks and forged-origin prefix hijacks. ASes generate and store ASPA objects in the RPKI. Each object authorizes the provider of an AS to propagate a prefix, and its validity can be checked cryptographically as the creating AS signs it. To check the plausibility of a received AS path, ASes can check whether their respective customers have authorized the ASes on the path via ASPA objects. This plausibility check does not have to occur directly on the routers but can be outsourced to dedicated machines, leading to improved performance. This outsourcing of validation is referred to as out-of-band. Figure 2a,b shows how ASPA can prevent hijacks and route leaks even with partial deployment. In Figure 2a, AS_F can check which provider AS_B has authorized. Since AS_B only authorizes AS_E in its ASPA object, AS_D should not propagate the prefix of AS_B. On the other hand, Figure 2b shows a route leak scenario in which AS_E incorrectly forwards a route learned from the provider to another provider. If AS_G now checks the ASPA objects of the ASes on the route, it can detect that the route consists of a downstream path followed by an upstream path. This indicates that the described route leak scenario is occurring. Compared to path validation algorithms such as BGPsec, ASPA offers a lower level of security, as the algorithm can only check a path's plausibility, not its integrity. However, due to the out-of-band validation mechanism, this path plausibility approach does not produce any additional overhead in BGP traffic, unlike path validation algorithms [11].

As of August 2024, 73 out of the over 117.000 ASes worldwide have created an ASPA object in the RPKI [22, 23].

2.3 | BGP Roles

Azimov et al. [15] published an approach to assigning roles to External BGP (eBGP) sessions. Only four relationships between the roles are allowed, as shown in Table 1. The roles do not necessarily correspond to the business relationships but are based on the restrictions described in the Gao-Rexford model. The relationships indicate which data may be sent from the respective ASes to their neighbors, which reduces complexity and can reduce the risk of accidental route leaks. ASes can establish their relationship using the BGP OPEN message. If the roles of two ASes do not match Table 1, no session is established. If one of the ASes has not configured a role, a session can be established (loose mode) or not (strict mode), depending on the implementation.

2.4 | BGP OTC Attribute

In 2022, Azimov et al. [15] published an approach to prevent route leaks using role assignment. After establishing the roles, the ASes mark the subsequently propagated routes according to the relationship. Azimov introduced the OTC attribute to mark the message. This attribute "is an optional transitive Path Attribute of the UPDATE message" [15], which should ensure that routes sent once to a peer, customer, or Route Server (RS) customer are only sent on to customers. The OTC attribute contains an AS number, which is determined according to the following ingress rules[1]:

TABLE 1 | BGP roles and allowed relationships [15].

Role	Allowed remote AS role
Provider	Customer
Customer	Provider
Route server	Route server–client
Route server–client	Route server
Peer	Peer

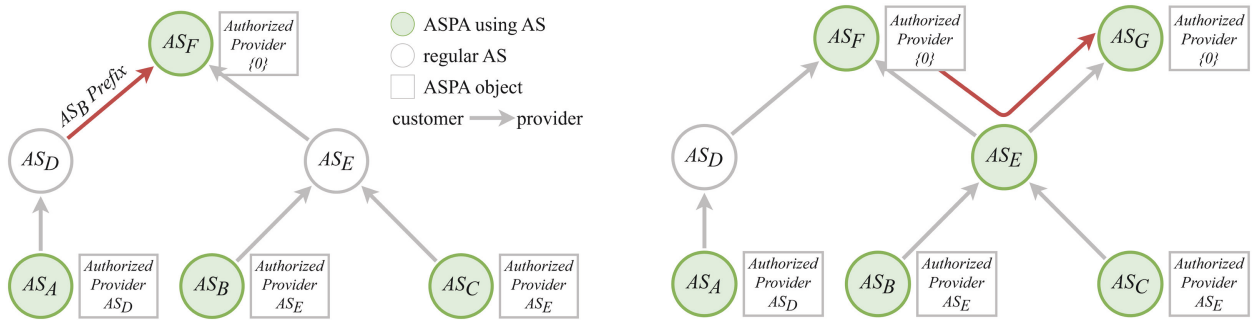

(a) Hijack detection with ASPA

(b) Route leak prevention with ASPA

FIGURE 2 | ASPA in partial deployment.

- If an AS receives a route from a customer or RS customer and an OTC attribute is present, this indicates a route leak. In such cases, the AS rejects the route to prevent the leak from spreading further.
- If an AS receives a route from a peer and there is an OTC attribute that contains AS numbers (ASN) that are not equal to the peer's ASN, it is also a route leak and will not be accepted.
- If an AS receives a route from a provider, peer or RS, it must add the remote AS's ASN to the OTC attribute if it does not already exist.

In addition to these rules, there are also the following egress procedures:

- If an AS advertises a route without an OTC attribute to a customer, peer, or RS customer, it must add its own ASN to the attribute.
- An AS may not propagate routes to providers, peers, or RSes if an OTC attribute is present.

ASes can detect route leaks immediately after following these rules. Moreover, they can detect leaks multiple hops away from the leaking AS. If, for example, an AS adds its ASN before forwarding to a peer, the message continues over several hops, and another AS sends the route to a provider, the provider can still detect the route leak and discard the route.

Even though the OTC attribute reliably detects route leaks, it is essential to remember that it only helps with accidental route leaks. An attacker can remove the attribute before forwarding the message to the following AS, and it will not be noticed.

2.5 | BGP Down Only (DO) Community

In January 2024, Sriram et al. [16] published another approach to prevent, detect, and mitigate route leaks. Again, this approach uses BGP roles for sessions between BGP peers and marks BGP messages. They use the same ingress and egress rules to mark the messages as the OTC approach. However, unlike OTC, Sriram et al. use BGP Large Communities to parse the marking, not BGP Attributes. They created a new Community called the DO Community to store the marking. Communities have a higher risk of being dropped. However, routers that want to access the data stored in a Community will not require a software upgrade, while accessing Attributes will require one [16].

3 | Related Work

There are few other publications on preventing route leaks besides the DO, OTC, and ASPA approaches presented here. Another approach with a similar function to ASPA is AS-Cones [24]. As with ASPA, ASes can create cryptographic AS-Cones objects to store in the RPKI. These signed objects include a list of trusted customers and their respective prefixes. ASes can check the validity of the BGP messages path by iterating over the individual hops of the path and checking whether these ASes have been certified as trustworthy by their respective provider. In contrast to ASPA, this is a top-down approach. Snijders et al. [24] have discontinued their work on this approach, and the draft has now expired, so we focus on a combined approach with ASPA.

In addition to new approaches for preventing route leaks, Morris et al. [25] have introduced a new method called BGP-iSec. This method addresses various mechanisms to make OTC resilient against different attacks. Their approach enhances the BGPsec protocol. Among other things, BGP-iSec introduces integrity-protected OTC fields, a new UP attribute to complement OTC, and a Providers-Cone Identification (ProConID), an RPKI object similar to ASPA objects. BGP-iSec strengthens route prevention mechanisms against different attacker models. In contrast to this work, the BGP-iSec paper focuses on improving BGPsec and not purely on using OTC for route leak prevention and the combination with ASPA.

This work is an extension of our previous work [17], where we compared the deployment strategies of ASPA and AS-Cones using the same simulation environment, BGPsec Sim [18]. We concluded that even though AS-Cones has comparable route leak prevention capabilities as ASPA with fewer participating ASes, only ASPA can prevent forged-origin prefix hijacks. We have looked at different deployment strategies and concluded that a top-down approach is most effective. Furthermore, we showed that the effectiveness of ASPA hinges on the active involvement of the victim AS in generating ASPA objects.[2]

4 | Methodology

This paper has two objectives. Firstly, we want to implement DO and OTC and measure their effectiveness in route leak detection and prevention in different deployment strategies. Secondly, we want to examine how DO/OTC can support ASPAs' route leak detection. We use a simulation environment from Brand and Posen [26, 27], which was extended by Rodday et al. [17, 18]. We extend the existing implementation with DO and OTC and use different scenarios to test their effectiveness.

To evaluate DO and OTC in a realistic network, we use the Center for Applied Internet Data Analysis (CAIDA) as_rel2 dataset, dated May 2024 [28]. This dataset includes connected ASes and their relationships, allowing us to create a graph within our test environment with 76.801 ASes. We divide the ASes into three categories and follow the definition of our previous work [29]:

- Tier-1: These ASes have no provider and only customer or peer relationships. They are located at the top of the network topology. Out of the CAIDA dataset, 0.18% of all ASes can be classified as Tier-1.
- Tier-2: These ASes have provider, customer, and peer relationships with other ASes. 15.37% of the ASes in the CAIDA dataset are tier-2 ASes.

- Tier-3: These ASes only have provider and peer relationships. They represent the leaves of the topology and make up the remaining 84.45% of the dataset's ASes.

All our tests are based on the same route leak scenario. We start by creating a graph based on the CAIDA data. Each AS in this graph is assigned a default policy. This policy ensures that the AS accepts and forwards routes based on the rules of the Gao-Rexford model. Furthermore, each AS prefers routes based on the following sequence:

1. Local preference (customer over peer and peer over provider)
2. Path length
3. Next hop ASN

After creating the graph with the default policy ASes, we randomly choose a pair of ASes, an attacker and a victim. We assign a new policy to the attacker that accepts and forwards all BGP update messages. This new policy deliberately creates route leaks. The victim is the starting point for our tests. After initialization and adaptation of the attacker policy, the simulation begins to iterate through the victim's neighboring ASes. An UPDATE is generated for each of these ASes, and it is checked whether the UPDATE matches the sender's egress rules and the receiver's ingress rules. If this is the case, the victim AS announces the route. Each AS that receives an update in this way repeats the procedure. As a result, BGP updates are propagated piece by piece, with the victim as the origin. Once all possible paths have been found, the next step of the evaluation begins. We iterate over all ASes of the graph and check whether one of the stored routes in the routing information database (RIB) violates the Gao-Rexford model and thus represents a route leak. The sum of all route leaks is then output. The number of route leaks measured can vary significantly depending on where the attacker is located within the graph, how high its connectivity is, and how the topology is built.

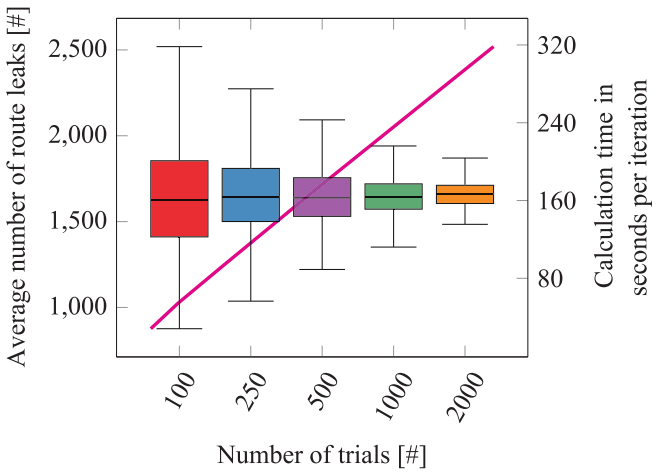

FIGURE 3 | Number of route leaks for different number of trials.

To achieve valid and statistically significant results, we have to repeat our process, starting with selecting the attacker and victim and ending with the output of the number of route leaks and creating an average amount of route leaks detected. Figure 3 shows how the average attacker success rate behaves with increasing repetition. The deviation becomes smaller and smaller as the number of trials increases. However, the calculation time required increases simultaneously linearly. With 1000 repetitions, we achieve a good trade-off between calculation time and precision, which allows us to achieve reliable data in a manageable amount of time. Algorithm 1 shows the pseudocode for this procedure. We have saved the 1000 randomly generated trails we use for our measurement and made them available on our GitHub repository [30].[3]

4.1 | Implementation

Our first step was to extend the BGP message in our test environment to support the data fields needed for OTC and DO. Since our test environment uses simplified BGP messages and no actual router instances, there is no difference in whether we implement attributes or communities. The test environment stores both data fields in a simplified route object similarly, and all ASes in our simulation will pass all of these objects' fields. Therefore, we implemented a general data field to store the marking based on the OTC/DO ingress and egress rules. In addition to extending the route object, we created an OTC and DO policy that respects both approaches' rules and marks routes accordingly. Furthermore, we duplicated the ASPA policy for our combined ASPA and OTC/DO test cases. We extended the new policy to use the OTC and DO ingress and egress rules combined with the ASPA policy.

Algorithm 1 Pseudeocode for our route leak test

```
1: procedure RouteLeakTest
2:     create_graph()
3:     assign_default_policy_to_all_ASes()
4:     for i ← 0 to 999 do
5:         create_attacker_victim_pair()
6:         attacker.policy ← RouteLeakPolicy()
7:         find_routes_from_victim()
8:         count_route_leaks()
9:     end for
10: end procedure
```

5 | Only-to-Customer and Down-Only Deployment

Our tests examined two deployment strategies: random selection of ASes and selection based on degree of connectivity. This section describes our procedure and results for both approaches.

5.1 | Random Selection

Our evaluation starts with a deployment rate of 0% on all tiers, meaning all ASes have a default policy and only the attacker AS receives the route leak policy. The proportion of ASes using OTC/DO was then changed step by step, as Figure 4 illustrates. For better visibility, Figure 4 uses 10% increments instead of 5%.

We divided the ASes of the graph into different tiers and considered these tiers as a whole. In order to consider the influence of the different tiers, we first increased the deployment of OTC/DO on tier 1 from 0 to 100% in 5% increments. For each iteration, we selected a random quantity from the set of ASes of tier 1 according to the percentage and changed their policy to OTC/DO. Then, we changed the deployment of tier 2, adding that for each iteration on tier 2, we also mapped the deployment of tier 1 from 0% to 100%. Finally, we increased the deployment for tier 3 in 5% steps and again increased tiers 2 and 1 for each step as described above to gain every possible OTC/DO deployment combination.

Figure 5a shows the results for isolated deployment on each tier. We observe that tiers 1 and 2 significantly influence the number of route leaks, whereas OTC/DO deployment on tier 3 ASes has no effect. The use on tier 1 reduces the average number of route leaks from nearly 1800 to fewer than 1700 by nearly 6%. OTC/DO on tier 2 ASes, on the other hand, has an even more significant effect. Here, we have a reduction from nearly 1800 to under 300, that is, over 84% fewer route leaks. However, it should also be noted that 139 ASes adjusted their policy at tier 1, compared with 11,804 ASes at tier 2.

Figure 6a shows the number of route leaks with increasing tier 2 and constant tier 1 deployment. Figure 6b shows the results for tier 3. We find here that increasing tier 1 deployment amplifies the effect of deployment on tiers 2 and 3. With a constant 50% for tier 1, OTC/DO on tier 2 can reduce the average number of route leaks by over 88% and tier 3 by over 20%. However, we keep in mind that preventing over 26% of leaks with tier 3 was only achieved by a full OTC/DO deployment on 64,858 ASes. If we compare the deployment of tiers 2 and 3 with a constant 100% deployment on tier 1, we find that it is now possible to detect almost all route leaks with tier 2 deployment. Tier 3 deployment also has a more significant impact and can almost halve the number of route leaks. The advantage of the random deployment strategy is that it reflects a more realistic deployment scenario. However, it can also distort the data. Figure 6a,b demonstrates that there has been another increase after a steady decrease in route leaks. This is because a new random set of ASes was chosen for each iteration. In other words, the set for 70% deployment may not contain all ASes that were already changed at 60% deployment. This can lead to one set having more ASes with higher connectivity than the other, causing the number of unrecognized route leaks to continue increasing, even with increasing deployment.

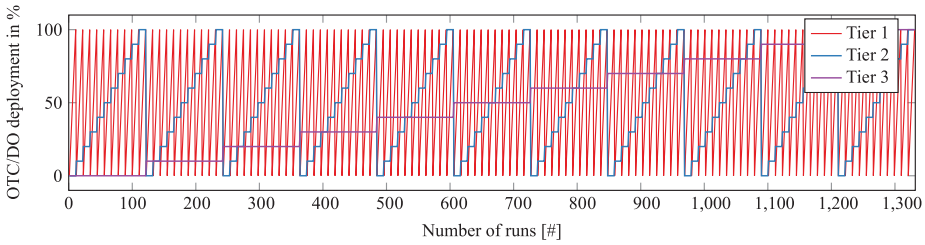

FIGURE 4 | OTC/DO deployment increments for each tier.

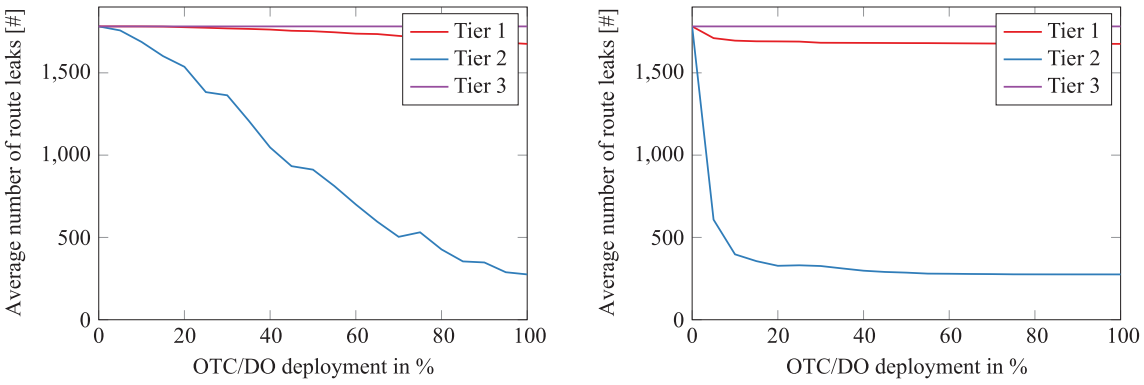

(a) Random deployment for tier 1 - tier 3

(b) Top-ISP deployment for tier 1 - tier 3

FIGURE 5 | Deployment separated by tiers.

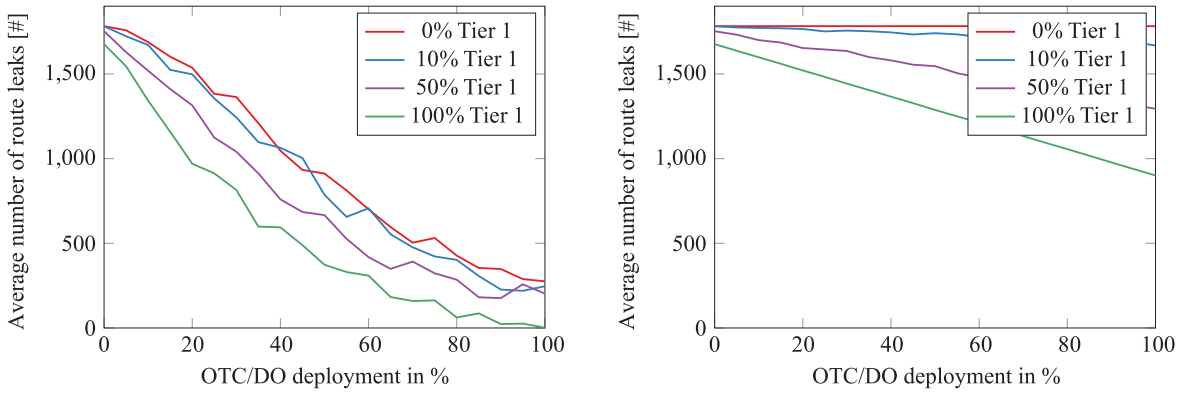

(a) Tier 2 deployment

(b) Tier 3 deployment

FIGURE 6 | Tiers 2 and 3 deployment with constant tier 1 deployment.

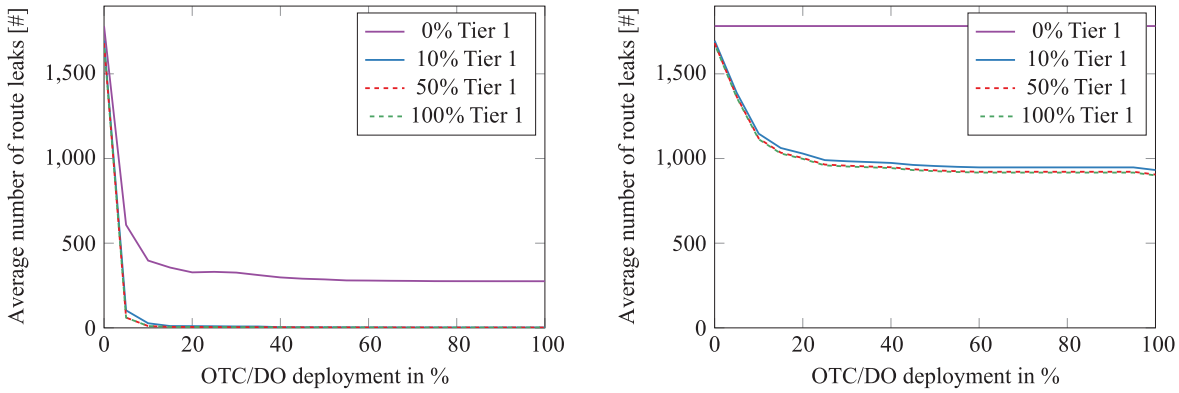

(a) Tier 2 deployment

(b) Tier 3 deployment

FIGURE 7 | Tiers 2 and 3 deployment with constant tier 1 deployment.

5.2 | Top-ISPs

The top-ISP deployment method is similar to the random deployment, with one key difference. Instead of selecting a random quantity from the tiers, we arrange the ASes within each tier based on their connectivity and prioritize the best-connected ASes. This approach ensures that ASes with the highest connectivity are chosen first, allowing us to quickly observe the effects of OTC and DO when used at critical nodes. In previous work, Rodday et al. [17] showed that this selection strategy proved highly effective. In contrast to random deployment, the top-ISP approach enables us to better identify and analyze the full impact of OTC/DO.

Figure 5b shows the impact of isolated deployment on different tiers. Tiers 1 and 2 significantly affect the number of route leaks, while tier 3 has no impact. Tier 1 deployment reduces the average number of route leaks by 5,3%, while tier 2 deployment detects 80% of all route leaks.

Figure 7a illustrates the deployment on tier 2 with a constant deployment on tier 1. The data indicate that almost all route leaks can be identified with 100% OTC/DO usage at tier 2 and 10% deployment on tier 1. Comparing these data with Figure 5b, it is apparent that even a 10% deployment at tier 1 significantly benefits the effects of OTC and DO on other tiers. Notably, the top 10% of ISPs on tier 1 only represents 14 ASes, so even a small number of altered ASes can have a substantial impact. If we compare the results for 10% tier 1 deployment with constant 50% or 100% deployment, we see that the increased OTC/DO usage on tier 1 barely increases the average number of detected route leaks. Figure 7b displays the deployment for tier 3 ASes with different constant deployment scenarios for tier 1. Like the tier 2 deployment, tier 1 greatly affects the impact of OTC/DO for tier 3. Full deployment on tiers 3 and 1 at 10% or higher can cut the number of route leaks by half. However, similar to the effects in Figure 2, the improvement from 10% to 100% tier 1 deployment is only marginal.

Figure 8a,b shows a closer examination of tier 3. Both figures show that tier 3 deployment, with increased OTC/DO usage on tiers 1 and 2, only prevents a fraction of route leaks. It can be seen here that the top 5% of tier 3 ASes with the highest connectivity, that is, around 3000 ASes, have the most significant influence on the average number of route leaks. Overall, detection of all route leaks is only possible with 100% OTC/DO usage on all tiers, as Figure 8a shows.

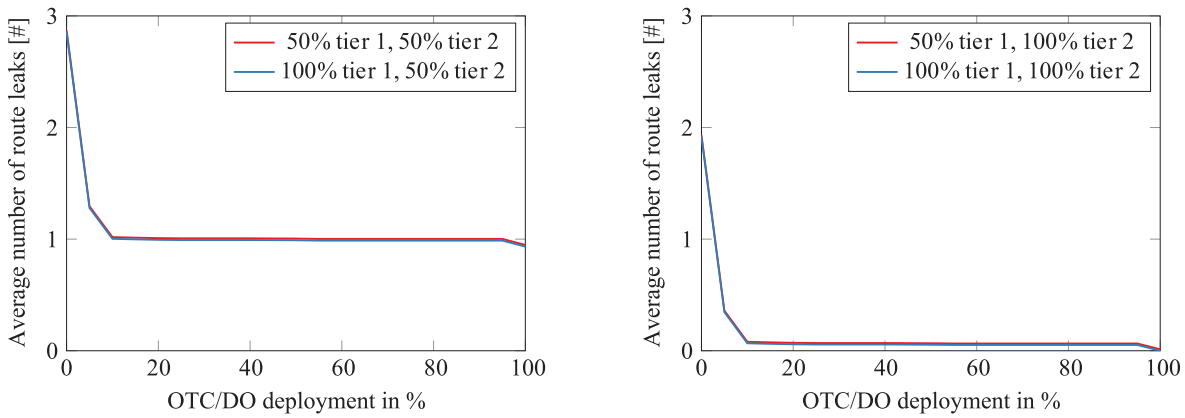

(a) Tier 3 deployment with tier 2 at 50%

(b) Tier 3 deployment with tier 2 at 100%

FIGURE 8 | Tier 3 deployment with constant tiers 1 and 2 deployment.

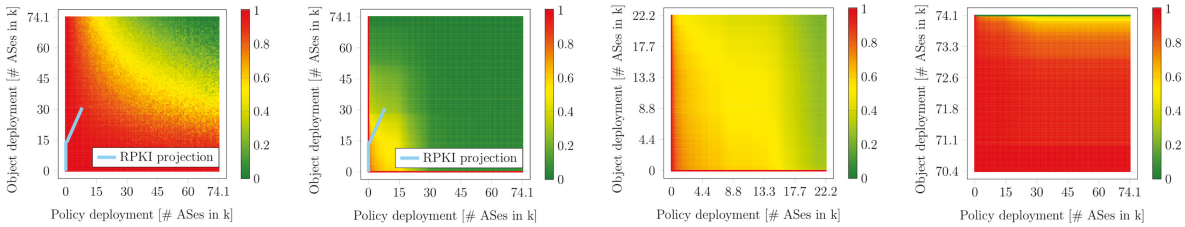

(a) Random object and policy deployment.

(b) Top-down object and policy deployment.

(c) Lower left enlarged from Figure 9b.

(d) Bottom-up object and top-down policy deployment.

FIGURE 9 | ASPA deployment scenarios. Objects and policies are deployed in all ASes.

6 | ASPA Deployment

This section presents the results of our previous work [17]. We evaluate ASPA's route leak mitigation capabilities under various deployment strategies. We analyze the impact of policy deployment and object creation on ASPA's ability to detect and prevent route leaks, providing insights into the optimal deployment approaches for maximizing its benefits.

Figure 9a shows our simulation results for random ASPA object and policy deployment. We represent the *route leak success rate* by showing the mean value over 1000 trials on a color-coded scale. Red indicates that a route leak still reached most ASes, while the further the color moves via yellow to green, the fewer ASes have been affected by the leak. The x-axis and y-axis represent the absolute amount of ASes deploying the security solution. Note that the ASes creating the object and the ASes chosen for policy deployment do not necessarily need to be the same.

We observe a gradual increase in benefits as more objects and policies are deployed within the topology. The benefits achieved are negligible below a threshold of 15k ASes for object deployment and below a threshold of 25k ASes for policy deployment. ASPA starts to show more benefit the more ASes participate in object and policy deployment. Overall, about 45k ASes would have to adopt ASPA to have a meaningful impact on routing security with a random deployment strategy, which is unlikely to happen.

Instead of anticipating random deployment, we are interested in a deployment scenario that provides significantly better protection with fewer objects and fewer ASes participating in filtering. To this end, we simulate a top-down selection strategy in Figure 9b. ASes deploying ASPA objects and policies are chosen by their out-degree. Therefore, larger ASes start with deploying the security solution, while smaller ASes deploy at a later stage. We observe that ASPA is much more beneficial in mitigating route leaks with such a selection strategy. This behavior can be explained by the fact that larger ASes are positioned in-between many paths and, therefore, have a more significant impact when filtering. Also, ASPA requires tier 1 providers to issue ASPA objects with *AS0* in them to show their participation in ASPA, stating that they do not have any upstreams. This is important for the algorithm as otherwise *no attestation* will be assumed for that particular edge in the graph, and the route including that edge will become *unknown* instead of *invalid* in most cases. An unknown route will be accepted, while an invalid route will be rejected by an AS deploying the policy.

Since we observe in Figure 9b a significant improvement that only requires relatively little deployment of ASPA in large ASes to be fruitful, we take a closer look at the lower left corner considering the largest 22k ASes in Figure 9c. Significant benefits in route leak mitigation can already be achieved with only very few large ASes deploying objects and some hundred ASes deploying the policy. At a deployment stage of 8k ASes (10.8%) deploying the objects and 5k ASes (6.7%) deploying the policy, we already

see that 50% of route leaks are successfully mitigated. The *y*-axis shows that the largest providers are the most important ones to deploy ASPA objects. Results are improving as more ASes continue to deploy ASPA objects. We also observe that filtering routes by deploying the ASPA policy has more of an impact compared to deploying ASPA objects and relying on other ASes to perform the filtering.

Since ASPA objects contain relationship information from a customer towards a provider, we show the opposite object creation strategy in Figure 9d. Here, objects are created in a bottom-up fashion, starting with the smallest ASes and moving towards larger ones, but policy deployment remains in a top-down fashion, starting with the largest ASes and moving towards smaller ones. The benefit of such a deployment strategy is only evident in the top 1% of ASes deploying objects. This is again due to the fact that the largest ASes are in-between many paths and those ASes not having ASPA objects renders many routes unknown instead of invalid, leading to the propagation of route leaks.

Overall, a selective ASPA deployment strategy from top-down proves to be much more beneficial than a random deployment strategy. Our results for the bottom-up selection strategy emphasize the fact that deployment in large ASes is essential for the success of ASPA.

To provide an estimate of how beneficial ASPA deployment would be considering the current state of RPKI deployment, we added an RPKI projection line into the Figure 9a,b. The blue line resembles the past 11 years of RPKI deployment from July 1, 2012 until September 30, 2023. We explain our methodology in obtaining the underlying data for growth patterns in Appendix A. We observe that with the current state of RPKI deployment, a random selection strategy would yield no benefit while with a top-down selection strategy route leaks would be reduced to a success rate of below 50%.

7 | Combined Deployment

This section describes the combined deployment of OTC/DO and ASPA. Like the top-ISP deployment method, this approach selects ASes ranked by their connectivity. Instead of the OTC/DO policy, our simulation assigns the combined policy to each selected AS and creates an ASPA object simultaneously. Therefore, every AS using this new policy marks messages according to the DO and OTC policies and creates an ASPA object. Each AS forwards and accepts routes only if both ASPA and OTC/DO policy ingress or egress rules are fulfilled.

Figures 10a to 11a show the sole deployment for each tier, with sole OTC/DO and ASPA deployment and combined deployment of both approaches. One can observe that sole ASPA deployment on tier 1 has no effect, and when combined with OTC/DO, the effects remain the same as with sole OTC and DO deployment. This result is expected due to the bottom-up functionality of ASPA. As the tier 1 ASes in our simulation have no providers, their ASPA objects are correspondingly empty and do not authorize anyone.

If, on the other hand, we look at Figure 10b and examine the pure deployment for tier 2, we can see that ASPA already has an effect here. The combination with OTC/DO further intensifies the effect. Figure 11a shows the tier 3 deployment. Unlike DO and OTC, ASPA slightly impacts the number of route leaks and can mitigate roughly 0.8% of all route leaks.

Figure 11b shows tier 3 deployment with static 10% usage on tier 1. One can observe that OTC and DO's route leak prevention capabilities surpass those of ASPAs. A combined approach can enhance these capabilities by 1.75%.

In summary, OTC/DO detects significantly more route leaks than ASPA. Using both approaches together typically has the same effect on preventing route leaks as using OTC or DO alone. Joint deployment with ASPA rarely enhances this effect.

8 | Discussion

The primary distinction between DO and OTC is that DO utilizes BGP Communities to label BGP messages, while OTC uses BGP Attributes. Attributes are more likely to parse non-upgraded routers [16]. However, routers requiring access to this

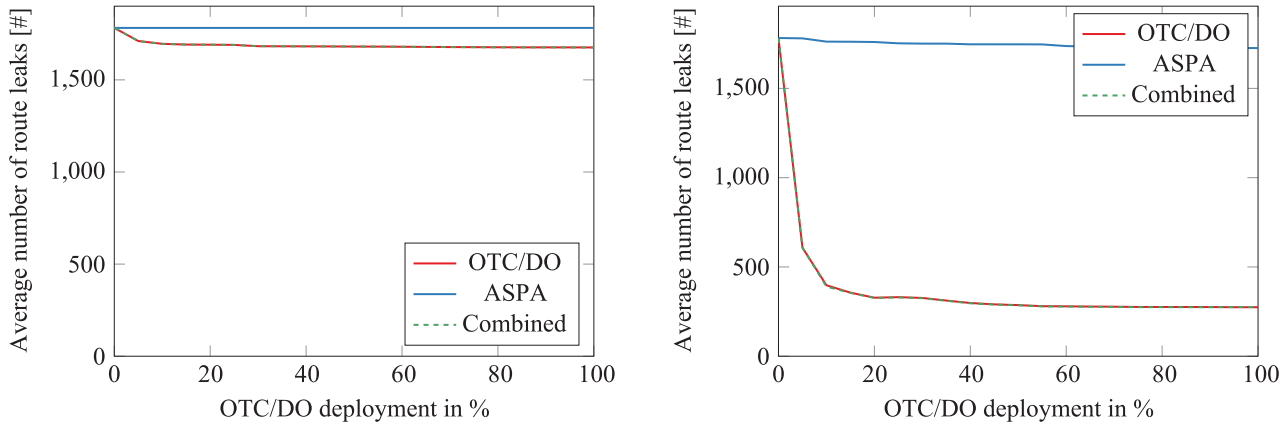

(a) Tier 1 deployment

(b) Tier 2 deployment

FIGURE 10 | Combined deployment for tiers 1 and 2.

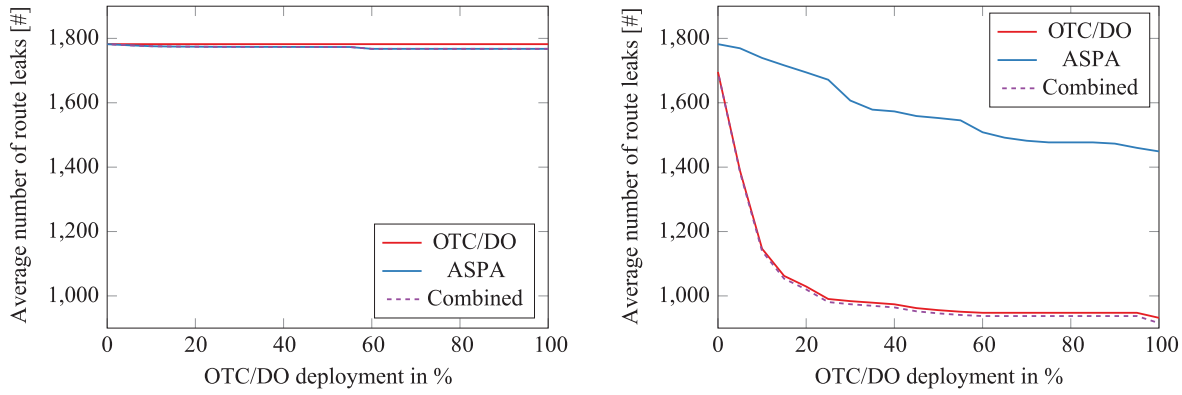

(a) Tier 3 deployment

(b) Tier 3 deployment with tier 1 at 10%

FIGURE 11 | Combined deployment for tier 3.

attribute will need a software upgrade, potentially hindering the widespread adoption of OTC. Nevertheless, one must remember that methods such as ASPA and BGPsec also require an upgrade.

On the other hand, the situation is different in BGP communities. While there is an increased risk of these being dropped by routers that have not yet been adopted, they offer the advantage that routers wanting to use DO do not require a software upgrade. Therefore, DO has a lower entry barrier compared to OTC. Communities are already used for many purposes and are somewhat overloaded. Therefore, Azimov et al. [15] have reserved the Large Community <TBD1> class for DO.

Based on the marking method, we recommend using OTC. The reduced risk of dropping is a massive advantage, especially for partial deployment. We have demonstrated that we can prevent most route leaks by adapting the ASes on tier 1 and well-connected ASes on tier 2. As a result, the number of systems requiring upgrades is relatively low. Moreover, deployment on tier 1 ensures that OTC significantly impacts tiers 2 and 3, which is why we recommend a top-down deployment strategy, starting with the top ISPs. If one looks beyond our measurement results to real-life deployments, one can already see the first uses of OTC. We searched the saved data of the RouteViews [31] project for July and August 2024 and found a total of 0.33% OTC usage among all routes, covering over 87,000 prefixes. However, we were unable to detect DO usage.

Our measurements show that route leak prevention is more substantial with OTC than with ASPA. Furthermore, ASPA demonstrates greater complexity than OTC due to its reliance on the RPKI infrastructure and the necessity of generating ASPA objects. However, this complexity translates to enhanced security against potential attacks. While OTC markings can be easily removed, ASPA objects are inherently more resistant to malicious manipulation or deletion. By combining both approaches, better route leak prevention by OTC can be used in parallel with the path plausibility of ASPA. In addition, even if the OTC field is deleted or modified, ASes can still detect route leaks using ASPA. We therefore recommend that ASes that use ASPA also use OTC. Furthermore, Rodday et al. [17] also demonstrated that ASPA, like OTC, works best with a top-down deployment approach.

8.1 | Security Considerations

While our findings demonstrate the effectiveness of OTC/DO, it is essential to note that the OTC Attribute and the DO Community are not secure. This means that an attacker or an inattentive user could remove this marking and potentially prevent the next ASes from detecting a possible route leak, something that would not be possible with ASPA. OTC/DO, when used alone, cannot prevent such behavior. However, solutions like BGP-iSec [25] have been proposed to address these vulnerabilities and have demonstrated their effectiveness.

Besides the potential for removing the OTC/DO fields, our evaluations have considered deployments of up to 100%, which are not practically feasible. Consider, for instance, the introduction of RPKI and Route Origin Authorizations (ROAs). It took over a decade for ROAs to cover more than half of all IPv4 routes, a milestone only reached in May 2024 [32]. We can anticipate a similar pace of adoption for new approaches like OTC and DO. Therefore, it is unlikely that OTC/DO fields will reach 100% deployment in the near future. Instead, we can expect a gradual increase in adoption over several years.[4]

8.2 | Limitations

Our findings offer important insights into the effectiveness of OTC and DO, individually and in combination with ASPA. However, it is crucial to acknowledge the limitations of our tests and results. These limitations can be categorized into three areas: the simplification of our simulation, the incompleteness of the CAIDA dataset we used, and the accuracy of our results.

We utilized a simplified simulation with reduced routers and BGP message instances to analyze large-scale topologies and diverse deployment scenarios. While necessary, this simplification excludes factors like prefix filters, complex routing policies, and potential removal of the OTC/DO markings. Our findings provide a valuable baseline, but real-world scenarios can be more complicated.

While extensive, the CAIDA dataset we used is incomplete and does not fully represent the global routing infrastructure. It does

not contain all ASes and their relationships, private peerings, or route servers. Despite these limitations, it offers a sufficiently large and realistic simulation topology.

The number of trials was balanced against computational cost, impacting the precision of our results. The effectiveness of route leak prevention mechanisms depends on deployment scenarios and the specific route-leaking AS. While our repeated evaluations provide a benchmark, real-world deployments will exhibit variability due to factors like topology, routing policies, and other dynamic conditions.

9 | Conclusion

Route leaks are still a significant vulnerability in BGP. While new methods for preventing route leaks have been proposed and are currently being widely discussed, there has been limited real-world implementation. This work enhanced a Python simulation environment with two new route leak prevention mechanisms: DO and OTC. We assessed the effectiveness of both approaches in a route leak scenario and conducted testing for random, top-ISP, and ASPA-combined deployment. Our research shows that DO and OTC effectively minimize route leaks, offering simple and efficient protection against them. The only difference between DO and OTC is the marking, not the ingress and egress rules. OTC is currently in use and has a lower risk of being dropped, so we recommend OTC over DO. Compared to other methods like BGPsec or ASPA, OTC requires minimal effort. Routes only need to be marked and checked for marking without any additional validation, resulting in no extra overhead. OTC can effectively protect against accidental route leaks. However, it does not guard against attackers who modify or remove the marking. Therefore, we recommend implementing ASPA alongside OTC. This way, even if the OTC marking is manipulated, ASPA's route leak prevention capabilities can be a backup. We have shown that when used by highly connected ASes, OTC has a rapid onset of action. As soon as ASes with high connectivity use this approach in real life, route leaks can be reduced quickly.

Acknowledgements

We thank the anonymous reviewers for their constructive feedback. We also thank Tim Meyer for his work on the initial ASPA implementation and Klement Hagenhoff for his ideas on the graphical representation of measurement data. Moreover, we thank Sriram Kotikalapudi and the members of the NIST Internet Technologies Research Group for their time for discussion. Also, we thank Ralph Holz for his insights into certificate transparency. The authors acknowledge the financial support by the Federal Ministry of Education and Research of Germany in the program "Souverän. Digital. Vernetzt.". Joint project 6G-life, project identification number: 16KISK002.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The datasets generated and analyzed during the current study are available in our GitHub [30] repository.

Endnotes

[1] **Abbreviations:** ASN, AS number; eBGP, External BGP; RS, Route Server.

[2] **Abbreviations:** CAIDA, Center for Applied Internet Data Analysis; ProConID, Providers-Cone Identification.

[3] **Abbreviation:** RIB, Routing Information Database.

[4] **Abbreviation:** ROA, Route Origin Authorization.

References

1. G. I. Taylor and A. E. Green, "Mechanism of the Production of Small Eddies From Large Ones," *Proceedings of the Royal Society of London. Series A-Mathematical and Physical Sciences* 158, no. 895 (1937): 499–521, https://doi.org/10.1098/rspa.1937.0036, http://rspa.royalsocietypublishing.org/content/158/895/499.

2. K. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A Survey of BGP Security Issues and Solutions," *Proceedings of the IEEE* 98, no. 1 (2010): 100–122, https://doi.org/10.1109/JPROC.2009.2034031.

3. M. S. Siddiqui, D. Montero, M. Yannuzzi, R. Serral-Gracia, and X. Masip-Bruin, "Route Leak Identification: A Step Toward Making Inter-Domain Routing More Reliable," in *10th DRCN Conference*, (IEEE, 2014): 1–8.

4. S. L. Murphy, "BGP Security Vulnerabilities Analysis," (2006), RFC 4272.

5. L. Gao and J. Rexford, "Stable Internet Routing Without Global Coordination," *IEEE/ACM Transactions on Networking* 9, no. 6 (2001): 681–692, https://doi.org/10.1109/90.974523.

6. K. Sriram, D. Montgomery, D. McPherson, E. Osterweil, and B. Dickson, "Problem Definition and Classification of BGP Route Leaks," (2016).

7. T. Wan, E. Kranakis, and P. C. van Oorschot, "Pretty Secure BGP, PSBGP.," (2005), Citeseer.

8. M. Lepinski and K. Sriram, "Bgpsec Protocol Specification," (2017).

9. "Sidr Operations," https://datatracker.ietf.org/wg/sidrops/about/, Accessed: 2024-08-08.

10. N. Rodday, Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, T. C. Schmidt, et al., "The Resource Public Key Infrastructure (rpki): A Survey on Measurements and Future Prospects," *IEEE Transactions on Network and Service Management* 21, no. 2 (2024): 2353–2373.

11. A. Azimov, E. Uskov, R. Bush, J. Snijders, R. Housley, and B. Maddison, "A Profile for Autonomous System Provider Authorization," draft-ietf-sidrops-aspa-profile-18. Internet Engineering Task Force, (2024).

12. P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein, "BGP Prefix Origin Validation," 6811. IETF, (2013).

13. M. Lepinski and S. Kent, "An Infrastructure to Support Secure Internet Routing," Request for Comments, (RFC Editor, 2012).

14. N. Umeda, T. Kimura, and N. Yanai, "The Juice Is Worth the Squeeze: Analysis of Autonomous System Provider Authorization in Partial Deployment," *IEEE Open Journal of the Communications Society* 4 (2023): 269–306.

15. A. Azimov, E. Bogomazov, R. Bush, K. Patel, and K. Sriram, "Route Leak Prevention and Detection Using Roles in UPDATE and OPEN Messages," (2022), RFC 9234.

16. K. Sriram and A. Azimov, "Methods for Detection and Mitigation of BGP Route Leaks," Internet-Draft draft-ietf-grow-route-leak-detection-mitigation- 10, Internet Engineering Task Force, Work in Progress.

17. N. Rodday, G. D. Rodosek, A. Pras, and R. van Rijswijk-Deij, "Exploring the Benefit of Path Plausibility Algorithms in BGP," in *Noms*

2024-2024 IEEE Network Operations and Management Symposium, (IEEE, 2024): 1–10.

18. N. Rodday, "BGP Security Simulations," (2023), https://github.com/nrodday/NOMS-24.

19. B. Wijchers and B. Overeinder, "Quantitative Analysis of BGP Route Leaks," (2014), http://ripe69.ripe.net/presentations/157-RIPE-69-Routing-WG.pdf, Accessed: 2024-08-08.

20. S. Abd El Monem, A. Khalafallah, and S. I. Shaheen, "BGP Route Leaks Detection Using Supervised Machine Learning Technique," in *2020 2nd Novel Intelligent and Leading Emerging Sciences Conference (NILES)*, (2020): 15–20.

21. K. Sriram, D. Montgomery, D. R. McPherson, E. Osterweil, and B. Dickson, "Problem Definition and Classification of BGP Route Leaks," (2016), RFC 7908.

22. "Rpki Console," https://console.rpki-client.org/aspa.html, Accessed: 2024-08-08.

23. "Regional Internet Registries Statistics," https://www-public.imtbs-tsp.eu/~maigron/rir-stats/rir-delegations/world/world-asn-by-number.html, Accessed: 2024-08-08.

24. J. Snijders, M. Stucchi, and M. Aelmans, "RPKI Autonomous Systems Cones: A Profile To Define Sets of Autonomous Systems Numbers to Facilitate BGP Filtering," Internet-Draft draft-ietf-grow-rpki-as-cones-02, Internet Engineering Task Force; 2020. Work in Progress.

25. C. Morris, A. Herzberg, B. Wang, and S. Secondo, "BGP-ISEC: Improved Security of Internet Routing Against POST-ROV Attacks," in *Network and Distributed System Security (ndss) Symposium 2024*, (Internet Society, 2024).

26. J. Posen and W. Brand, "BGP Security Simulations," (2020), https://github.com/jimpo/bgpsec-sim.

27. W. Brand and J. Posen, "A Reproduction of "Jumpstarting BGP Security With Path-End Validation"," (2020).

28. "As Relationships," https://www.caida.org/catalog/datasets/as-relationships/, Accessed: 2024-08-08.

29. N. Rodday, L. Kaltenbach, I. Cunha, R. Bush, E. Katz-Bassett, G. D. Rodosek, et al., "On the Deployment of Default Routes in Inter-Domain Routing," in *Proceedings of the ACM Sigcomm 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet*, TAURIN'21, (New York, NY, USA: Association for Computing Machinery, 2021): 14–20.

30. N. Höger, "Integration of BGP Roles Into the BGPSIM Environment," (2024), https://github.com/nhoeger/bgpsim.

31. "University of Oregon Routeviews Project," https://www.routeviews.org/routeviews/, Accessed: 2024-08-08.

32. "RPKI Monitor," https://rpki-monitor.antd.nist.gov/ROV, Accessed: 2024-08-08.

33. A. Reuter, R. Bush, I. Cunha, E. Katz-Bassett, T. C. Schmidt, and M. Wählisch, "Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering," *ACM SIGCOMM Computer Communication Review* 48, no. 1 (2018): 19–27.

34. T. Hlavacek, A. Herzberg, H. Shulman, and M. Waidner, "Practical Experience: Methodologies for Measuring Route Origin Validation," in *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (dsn)*, (IEEE, 2018): 634–641.

35. G. Huston and J. Damas, "Measuring Route Origin Validation," (2020), https://www.potaroo.net/ispcol/2020-06/rov.html, [Online; accessed 16-October-2020].

36. B. Cartwright-Cox, "The Year of RPKI on the Control Plane," (2019), https://blog.benjojo.co.uk/post/the-year-of-rpki-on-the-control-plane, [Online; accessed 10-January-2020].

37. N. Rodday, I. 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 *Proc. of network traffic measurement and analysis conference (tma). ifip*, (2021).

38. 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 ACM Internet Measurement Conference (IMC)*, (Montreal: ACM, 2023).

Appendix A

ASPA Deployment Forecast

To forecast the deployment of ASPA and discuss within which time frame it might become useful, we use historical data on the RPKI deployment. We assume for RPKI object creation two linear functions. These estimations do not resemble reality with 100% accuracy. From July 1, 2012 to December 31, 2019, we assume 0.007% growth per day; from January 1, 2020 to December 31, 2022, we assume 0.01768% growth per day. During the first year of the RPKI, not many ROA were created, which is why our model starts at July 1, 2012.

RPKI policy deployment, also called ROV, is hard to measure as measurements are based on the inference of private router configurations. Many publications deal with different measurement methods to reliably measure ROV [33–38]. Either ROV research is capable of pinpointing filtering AS [33, 37] but is only executed once and therefore does not allow the inference of a growth rate, or methodologies are not capable of pinpointing [35, 36, 38] and measure the overall benefit of ROV including collateral benefit and therefore the RPKI protection rate, which does not translate to AS that actually deployed ROV.

To obtain a growth rate for ROV deployment, we use data from the latest publication [38]. It includes collateral benefit, and the actual deployment rate might be lower. By joining RPKI object creation and policy deployment rates, we can extrapolate on the benefit of ASPA and AS-Cones deployment. We observe the RPKI deployment rate within Figure 9 as the cyan bar. The RPKI projection is based on 11 years of deployment efforts.