

On the Deployment of Default Routes in Inter-domain Routing

Nils Rodday
RI CODE, Universität der
Bundeswehr München
Munich, Germany
nils.rodday@unibw.de

Lukas Kaltenbach
RI CODE, Universität der
Bundeswehr München
Munich, Germany
lukas.kaltenbach@unibw.de

Italo Cunha
Universidade Federal de Minas Gerais
Belo Horizonte, Brazil
cunha@dcc.ufmg.br

Randy Bush
Arrcus / IJ
USA
randy@psg.com

Ethan Katz-Bassett
USC / Columbia University
New York, USA
ethan@ee.columbia.edu

Gabi Dreö Rodosek
RI CODE, Universität der
Bundeswehr München
Munich, Germany
gabi.dreo@unibw.de

Thomas C. Schmidt
HAW Hamburg
Hamburg, Germany
t.schmidt@haw-hamburg.de

Matthias Wählisch
Freie Universität Berlin
Berlin, Germany
m.waehlich@fu-berlin.de

ABSTRACT

In this paper, we argue that the design of a responsible Internet requires a clear understanding of the current state of deployment. This work sheds light on default routing in the Internet, a routing strategy that reduces control but may help to increase availability when forwarding packets. We revisit and extend two common methodologies based on active measurements to increase coverage and accuracy. Our results show larger differences in the results between the methodologies. We confirm that default route deployment strongly correlates with the customer cone size of autonomous systems and that smaller networks are more likely to deploy a default route. Our data will help to better assess the deployment of other protocols such as RPKI route origin filtering.

CCS CONCEPTS

• **Networks** → **Network measurement**;

KEYWORDS

Default Route, DFZ, path-poisoning, RIPE Atlas, middleboxes

ACM Reference Format:

Nils Rodday, Lukas Kaltenbach, Italo Cunha, Randy Bush, Ethan Katz-Bassett, Gabi Dreö Rodosek, Thomas C. Schmidt, and Matthias Wählisch. 2021. On the Deployment of Default Routes in Inter-domain Routing. In *ACM SIGCOMM 2021 Workshop on Technologies, Applications, and Uses of a Responsible Internet (TAURIN '21)*, August 23, 2021, Virtual Event, USA. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3472951.3473505>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

TAURIN '21, August 23, 2021, Virtual Event, USA

© 2021 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-8639-5/21/08...\$15.00

<https://doi.org/10.1145/3472951.3473505>

1 INTRODUCTION

Selecting the next hop is a crucial step when forwarding packets, in particular in inter-domain routing when packets transit networks. Autonomous Systems (ASes) decide on a potential next hop towards a given destination prefix by accepting announcements from neighboring ASes. ASes may explicitly choose specific next hops by applying fine-grained Border Gateway Protocol (BGP) policies.

A default route is the last resort when forwarding packets because it captures all destinations that are not covered by a more specific prefix in the forwarding table. On the positive side, this reduces resource requirements as entries can be aggregated, and provides a fallback in case no route exists. On the negative side, it limits control and may forward packets to networks that are unable to or will not carry the packets towards the destination. For example, introducing default routes in the default free zone (DFZ) would lead to imbalanced and thus unfair routing. An attacker may even misuse default routes by sending packets towards destinations that are not available (*i.e.*, no covering prefix is announced in BGP), which unnecessarily increases traffic volume on upstream or peering links. In case of misconfigured or malicious traffic redirections, default routes undermine security mechanisms such as RPKI-based route origin validation [19].

The deployment of default routes in the Internet core conflicts with the design goals of the Responsible Internet paradigm [9], *controllability*, *accountability*, and *transparency*. When a default route is deployed, routing behavior diverges from common expectations. Packets that would be dropped because of a missing route are instead forwarded using the default route and connectivity is provided. Other members of the inter-domain routing infrastructure lack transparency since a default route is usually not exported to neighbors or public routing dumps. Default routes could lead to more traffic than expected to be sent towards a peering partner, which is the reason why some ASes specifically state in their peering policies that default routes should not be used towards their own AS. In the context of resilient inter-domain traffic exchange, accepting default routes is not recommended [30].

To measure default routes in inter-domain routing, two independent methodologies have been proposed in prior work, which we extend. (i) *Path-poisoning*, introduced by Bush *et al.* [1] and (ii) *not-announced prefix*, introduced by Hlavacek *et al.* [11].

Contributions. In detail, we make the following contributions:

- (1) We extend prior measurements based on the *path-poisoning* methodology by considering IPv6 and by using dual poisoning of the prepended AS-path to increase efficiency.
- (2) We add RIPE Atlas vantage points to the *path-poisoning* methodology where available. This allows us to not only identify a default route but also to infer its direction.
- (3) We increase coverage of the *not-announced prefix* methodology by adding NLNOG vantage points.
- (4) We introduce a threshold to decide on the ratio of routes before an AS will be flagged deploying a default route.
- (5) We present our findings and up-to-date results as part of an ongoing measurement study on our website <https://www.defaultroutes.net>. All datasets are publicly available.

The remainder of this paper is structured as follows: We introduce the measurement methodologies and setup in Section 2. We present our results in Section 3, discuss default routing in broader context in Section 4, and conclude in Section 5.

2 METHODOLOGIES AND MEASUREMENT SETUP

In this work, we extend two methodologies to identify ASes that deploy default routing. The first methodology is based on poisoning AS paths, the second methodology sends probe packets towards prefix space that is withdrawn. Both methodologies require that no covering prefix of the target address is announced. A covering prefix would provide a routing table entry, allowing the tested AS to be able to route traffic towards the peer it learned the route from. Hence, default route measurements would falsely identify connectivity. We tested for a covering prefix manually before starting the experiments.

2.1 Path-Poisoning

Overview. The concept of path-poisoning to identify default routes was first introduced by Bush *et al.* in 2009 [1]. The idea is twofold. First, ensure that the AS under test does not install a route to an experiment prefix but all other ASes are able to do. This ensures reachability of the experiment prefix except from the AS under test. Second, send probe packets from this experiment address space to an address that belongs to a prefix range that is announced by the AS under test. When a reply occurs, we can conclude that the AS under test reaches the experiment prefix via an alternative path. As no covering prefix of the experiment prefix is available, the AS must deploy default routing.

To make the experiment prefix available but ensure that the AS under test does not include a route to the prefix, path-poisoning is used. In detail, the origin AS of the experiment prefix artificially adds the ASN under test to the AS-path. This poisoning triggers BGP

loop prevention at the target AS, which prevents further consideration in the best path selection process. Hence, the announcement is discarded.


Extension. Originally, the methodology was only testing one AS at a time. Since the authors of [1] had a /16 prefix range at their disposal, which was used in 256 /24 chunks, time constraints were of less importance. For our ongoing study, we are able to utilize four /24 as well as four /48 prefix ranges for IPv4 and IPv6, respectively. A study covering all 72,004 ASes from the CAIDA AS relationship dataset [3] with four /24 prefix ranges and an average time of three hours per test would last about six years. To speed up the process, we decided to improve the efficiency by testing two ASes at a time, prepending both of them to the AS-path. This step halves the time for a full campaign. Since we want to avoid creating artificial links between two third-parties within the routing infrastructure that might not exist already, we always inject our own AS inbetween. Our testbed allows to path-poison up to two ASes at a time. If a testbed allows for more, efficiency could be further increased.

Two problems might arise from poisoning multiple ASes during one measurement run. First, if the tested AS as well as its upstream to which the default route points are poisoned in parallel, it remains ambiguous which of the two ASes (or both) deploy default routing. To eliminate this problem, we always check that the two ASes which are poisoned in parallel are not related, *i.e.*, do not exhibit a relation based on the CAIDA AS relationship dataset [3]. Second, filtering of the announcement is twice more likely to happen, according to [29]. However, in our measurements we did not observe any negative side effects.


Setup. We used the PEERING testbed [28] to announce our IPv4 and IPv6 prefix ranges. Our control server runs the PEERING testbed client, which connects via OpenVPN tunnels towards the PEERING Amsterdam Point of Presence (PoP). Figure 1a illustrates the PEERING testbed from which we announce our prefix ranges. AS3 receives the announcement and forwards it to its peers. While AS3 and AS4 incorporate the new route into their Route Information Base (RIB), AS1 and AS2 would be dropping the route announcement due to the loop prevention mechanism. As a result, AS1 and AS2 would not have connectivity to the announced prefix range, while AS3 and AS4 would be able to send packets our way.

In total, the updated measurement methodology consists out of eight phases:

- 1) **Announce prefix, wait 20 min.** The prefix is announced via the PEERING testbed and 20 minutes wait time are allocated for prefix propagation. Within that time the new route should have been received by most ASes.
- 2) **Look-ahead test with ZMap.** The look-ahead test is performed to find active hosts within the target AS and use those as reference points for later querying during the poisoned measurement phase. We made use of ZMap, which is optimized for Internet-wide network surveys [4, 7]. Since ZMap uses raw-sockets, many packets can be sent out in a very short period of time. We exploited this fact to scan subnets that were announced by the target AS. Once a reply was received we save this IP address as a reference point. To even further speed-up the process and reduce traffic for the PEERING testbed, we used IPv4 [10] and IPv6 hitlists [6]. We began with



(a) Path-poisoning methodology. AS1 and AS2 are poisoned in the announced AS-path, triggering BGP loop prevention at these ASes. AS3 and AS4 have connectivity towards the PEERING testbed.



(b) Not-announced prefix methodology. AS1 and AS2 are sending traceroutes towards a withdrawn prefix range. AS4 does not host a probe and cannot be tested.

Figure 1: Overview of two methodologies and our setup to measure default routes.

searching for active IPs within those hitlists, if nothing could be found we performed a full subnet scan with ZMap.

3) Look-ahead test with RIPE Atlas. We added a look-ahead test with RIPE Atlas, where vantage points within the tested AS were available. The methodology was originally designed to only function in an outgoing manner, e.g. hosts within the tested AS would be probed from the control server before and after path-poisoning. Hence it only allowed a binary answer whether there was a default route present, or not. We extend upon this work by also utilizing RIPE Atlas probes [23], where available. From the RIPE Atlas probes we send traceroutes towards our control server (opposite direction of measurement as before) to obtain information about the next AS hop. If we find a next AS hop, we are able to tell where the default route points to. To schedule the measurements we used the RIPE Cousteau library [20], for parsing we utilized RIPE Sagan [21].

4) Withdraw prefix, wait 90 min. The prefix is withdrawn and we wait 90 minutes to avoid triggering Route Flap Damping (RFD) upon reannouncing the prefix [8].

5) Announce poisoned prefix, wait 20 min. The prefix is reannounced, but the tested ASes are now poisoned within the AS-path of the BGP announcement, see Figure 1a. Again, we wait for 20 minutes to give BGP routers ample time to propagate the announcement.

6) Validation test with ZMap. The previously stored IP addresses from the look-ahead test are queried again with ZMap. If they are still reachable, a default route must be present. Otherwise, no default route is installed.

7) Validation test with RIPE Atlas. We also repeat the process with the RIPE Atlas probes. If the traceroute is able to reach the upstream, a default route is present. With this extension we are also able to determine the direction of the default route. If it fails, no default route is installed.

8) Withdraw prefix, wait 90 min. The prefix is withdrawn and a wait time of 90 minutes has to expire to avoid RFD before reusing the prefix for additional measurements.

2.2 Not-announced Prefix

Overview. This method was introduced by Hlavacek *et al.* in 2020 [11]. It is based on the assumption that an AS will not forward packets to a destination if no covering prefix exists, except if a default route is available. The idea is to issue traceroute measurements from the AS under test towards an IP address that belongs to an IP prefix that is not announced, as shown in Figure 1b. If the traceroute reaches routers in an upstream of the AS under test, we conclude that the AS under test deploys default routing.

The overall advantage of this methodology compared to *path-poisoning* is that the method is much faster. No prefix announcements are required, hence less resources are required from the experimenter. Complexity is reduced but at the expense of visibility. A typical measurement run takes around 30 minutes and covers all ASes hosting a vantage point. The drawback is that this method requires access to a vantage point inside the AS under test to initiate traceroute measurements. Previous work was based on RIPE Atlas, which currently covers 3,699 ASes.

Extension. An AS is considered deploying default routing if one traceroute of any single probe inside the AS reaches an upstream or peer. The authors do not discuss potential inconsistencies, which we observed during our measurements with multiple probes within the same AS. We consider inconsistent behavior across probes by adding a threshold value that allows the experimenter to decide at which ratio an AS is considered having a default route.

Setup. We extend upon prior work by adding the set of NLNOG vantage points, which cover 467 ASes in total, increasing the coverage of RIPE Atlas by 193 ASes for IPv4 and 234 ASes for IPv6. NLNOG provides a ring-all command that is designed to execute commands on all nodes simultaneously, e.g., running a traceroute from all nodes. This was, however, not usable for our purpose. We developed a tool that connects to all nodes and executes the traceroute command. Additionally, results were saved in a RIPE Atlas compatible way. The benefit is that other researchers using RIPE Atlas who also have access to NLNOG can now reuse existing

Table 1: Comparison of results of different methodologies measuring default routes.

AS Tier	Not-announced Prefix						Path-Poisoning			
	RIPE Atlas			NLNOG			PEERING Testbed			
	# Tested ASes	Default Routes		# Tested ASes	Default Routes		# Tested ASes	Default Routes		
IPv4										
1	11	(0.30%)	0 %	3	(0.64%)	0 %	15	(0.91 %)		33.33 %
2	1909	(53.30%)	23.31 %	306	(65.66%)	17.97 %	1001	(61.26%)		51.45 %
3	1662	(46.40%)	41.14 %	157	(33.70%)	39.49 %	618	(37.83%)		66.02 %
<i>Sum</i>	3585	(100%)	31.52 %	466	(100%)	25.11 %	1634	(100%)		56.79 %
IPv6										
1	10	(0.61%)	10 %	3	(0.64%)	33.33 %	15	(0.91%)		53.33 %
2	974	(59.46%)	25.67 %	306	(65.66%)	16.99 %	1002	(61.24%)		32.83 %
3	654	(39.93%)	39.76 %	157	(33.70%)	33.76 %	619	(37.85%)		21.49 %
<i>Sum</i>	1638	(100%)	31.2 %	466	(100%)	22.75%	1636	(100%)		28.73 %

code and easily extend coverage via additional probes. Our tool is available on Github [24]. In our study, we consider IPv4 and IPv6.

2.3 Classification of Network Tiers

Prior work [1] used the UCLA dataset [17] to categorize ASes. This dataset is no longer available. We create our own dataset to classify ASes into one of three tiers. We follow a lean, common definition: Tier 1 ASes have global connectivity without the need to purchase any connectivity from other participants. Tier 2 ASes peer for free with some ASes but need to buy transit from others to reach all portions of the Internet. Tier 3 ASes are stub networks that solely purchase from other ASes to obtain connectivity.

In order to map ASes to tiers, we briefly analyze the CAIDA AS relationship dataset [3] as well as the most recently published ProbLink dataset [31] to verify whether they can serve as input datasets. Both methodologies [12, 13] rely on publicly available BGP collector data from RIPE RIS [22] and RouteViews [26] to feed their algorithms to infer AS relationships. ProbLink claims to be 27% more accurate for inferring complex relationships.

We observed 72,004 and 44,695 unique ASes in CAIDA and ProbLink datasets. A discussion with the authors of [12] revealed that CAIDA's dataset is based on BGP data from multiple days, while ProbLink's inferences are calculated using BGP data from a single day. This leads to significantly reduced coverage. Therefore, we decided to use the CAIDA AS relationship dataset for our classification algorithm.

We process the data as follow: First, we label all ASes that form the input clique in the CAIDA AS relationship dataset as tier 1. ASes belonging to the clique are verified by CAIDA based on ground-truth. This subset comprises 19 ASes. Second, we label all ASes that either serve as provider for other ASes or maintain peering connections as tier 2. Third, we label all remaining ASes as tier 3. The resulting dataset gives us 19 (0.03 %) tier 1, 11,325 (15.73 %) tier 2, and 60,660 (84.25 %) tier 3 ASes.

For comparison, the UCLA approach used in [1], which distinguishes large, small, and stub networks, obtained 255 (0.08 %) large ASes, 1,361 (4.11 %) small ASes, and 31,517 (95.12 %) stub ASes. Our recent dataset classifies more tier 2 ASes, possibly since peering relationships between ASes are more common nowadays, leading to a flattening Internet [2, 32].

3 RESULTS

We summarize our results in Table 1. Since the *path-poisoning* methodology is very time consuming, every AS was only tested once, starting from December 2020. Our measurements are ongoing but already allow for relative comparison. The *not-announced prefix* methodology, on the other hand, can be rapidly reproduced, as testing can be done in parallel for all probes with RIPE Atlas and NLNOG. We ran those measurements on five consecutive days in February 2021 and found results to be the same for all days.

3.1 Preliminaries

Consistency check. For the *not-announced prefix* methodology we observed inconsistent behavior across multiple probes within a single AS, which led to manual investigation. For example, within AS3320, we used 190 RIPE Atlas probes but only four were able to reach the upstream (or peer), each via a different peer. Prior work [11] does not discuss this problem and labels an AS as deploying default routing if at least a single probe is able to forward towards the not-announced prefix. In our evaluation, we introduce a threshold that allows to decide when an AS is flagged. For each AS, this threshold is the minimum number of probes that need to indicate default routing in an AS compared to the overall number of probes in this AS. We flag an AS as deploying default routing if this threshold is exceeded.

Figure 2 illustrates the impact of the threshold on the overall results for the *not-announced prefix* methodology. Different probes in an AS may yield conflicting results in cases of partial deployment




Figure 2: Relative ratio of ASes, which we have identified as deploying default routes (y-axis), depending on a probe consistency threshold (x-axis). The number of ASes decreases slightly when an increased number of probes per AS exhibit the same outcome. We observe a decline at 0.5 because some ASes host only two probes with contrary results.

of default routes or incorrect identification of the border between the AS under test and the neighboring AS [14]. We observe a sharp decline at 0.05 (by definition). The decline at 0.5 can be explained by some ASes only hosting two probes. If those probes lead to conflicting results, the impact on the overall results is relatively large. In our experiments, we use a threshold of 0.55, requiring the majority of probes to reach our experiment prefix, in order to be flagged as having a default route.

For the *path-poisoning* methodology an AS is flagged deploying default routing if at least one probe has been answered during the poisoning phase.

Coverage. RIPE Atlas covers 11 of 19 tier 1 ASes, 1909 tier 2 ASes and 1662 tier 3 ASes. In contrast to this, NLNOG covers relatively more tier 2 ASes compared to tier 3. This is not very surprising, since NLNOG is a collaboration platform between network operators and mostly operators involved in the community are expected to participate. While many RIPE Atlas probes are IPv4 only, NLNOG requires nodes to support both IPv4 and IPv6.

Middlebox identifications. In some cases, we observed replies that claimed to have successfully reached the PEERING testbed even for a prefix range that has not been announced. We checked how many probes are located behind a middlebox that impacts our measurement results.

We announced a /24 prefix range via PEERING and selected 10,105 RIPE Atlas probes that are IPv4 capable, currently connected, and also have a public IPv4 address entry in the probe’s meta data [16]. Simultaneously, we capture the incoming traffic at our PEERING client. The public IP addresses of the probes are required since many probes are behind a Network Address Translation (NAT). 9,530 probes actually participated, with 9,474 probes claiming to have successfully reached PEERING. We confirmed 9,045 (95.47 %) probes, based on captured traffic at the client.




Figure 3: Presence of default routes in relation to Customer Cone Size based on RIPE Atlas data and a threshold of 0.55. We observe a higher deployment ratio of default routes when customer cone size is smaller. IPv4 and IPv6 results are almost identical.

322 (3.4 %) probes were not visible in our packet captures but received replies from the PEERING router immediately upstream of our client and must, therefore, have reached our prefix. 107 (1.13 %) probes could not be confirmed. All measurements for those 107 probes exhibit the same topology properties: Short paths, few private hops, and then the final destination (without going through the PEERING router). We conclude that 107 RIPE Atlas IPv4 probes are placed behind a middlebox that replies to ICMP or TTL-limited packets instead of forwarding the packet to its desired destination.

3.2 Comparison of Vantage Points, Methodologies, and Prior Work

Comparison of RIPE Atlas and NLNOG vantage points. We observe a reduced identification rate for IPv4 and IPv6 for all NLNOG probes compared to RIPE Atlas probes, see Table 1. 273 ASes provide probes for both RIPE Atlas and NLNOG: 250 (91.6 %) have identical results, while 23 (8.4 %) ASes are flagged differently. This behavior is consistent for different measurement runs.

Comparison of not-announced prefix and path-poisoning methodologies. 601 ASes were covered in case of IPv4 for both methodologies. 271 (45.11 %) ASes exhibit the same results. Of the remaining fraction, 162 (26.95 %) ASes show different results although tested from the same RIPE Atlas probe. The cause might be the time difference, since the not-announced prefix measurements finish at the same time, while *path-poisoning* measurements were executed throughout the course of five months and remain ongoing. 103 (17.13 %) ASes hosted RIPE probes whose traceroutes did not leave the AS but successfully replied to ICMP echo requests during the poisoning. The remaining 65 (10.81 %) ASes exhibit different results since the threshold for the *not-announced prefix* methodology prevented the inference of a default route, while for *path-poisoning*

a single probe leaving the AS was counted as deploying default routing.

Comparison with prior work. Prior work [1] found that 74.8 % of ASes were able to consistently reply to probe packets from the experiment prefix during the AS path-poisoning, 20.4 % of the ASes were not able to reply successfully, and 4.3 % answered for some IPs but not all. Table 1 shows that we found 57.25 % of ASes to be responsive during the poisoning while the others were not. However, the fraction of tier 3 ASes that were tested in the previous study was much higher, weighting in on the overall percentage. Broken down into large ISPs, small ISPs, and stub they found 17.1 %, 44.5 %, and 77.1 % to deploy default routes, respectively.

Hlavacek *et al.* [11] found that 768 (46.37 %) out of 1656 RIPE Atlas probes were hosted in networks using default routes. We covered all available ASes within RIPE Atlas and found the fraction of ASes having default to be lower (31.52%). On the one hand, the selection of RIPE probes might induce differences between the two studies, on the other hand, the threshold that we introduced sanitizes the results and only flags an AS as deploying a default route if the threshold was reached.

Overall, we confirm previous results that found default routes to be prevalent in smaller ASes. Figure 3 illustrates this based on data retrieved from the *not-announced prefix* methodology.

Ongoing measurements. All results are made public at our website <https://www.defaultroutes.net> to allow fellow researchers the use of our data. Moreover, the website allows to adjust the described threshold. Our scripts and results can be found on GitHub [24].

4 DISCUSSION

We consider this work as a first step to better understand rather unexpected routing behavior. Default routes may serve well in intra-domain routing scenarios but they conflict with controllability, transparency, and accountability in inter-domain settings. Any deployed default route ensures that data is forwarded towards a destination that is—based on BGP principles—not available because no covering IP prefix was announced or accepted. Default routes have economical and security consequences. Often, routes that reduce monetary costs are preferred. Default routes may undermine this model. Even worse, in case of not-announced prefixes, they lead to traffic forwarding that is not required. This can be misused by attackers. Moreover, it conflicts with the transparency goal, since it is not clear why connectivity persists while a prefix was withdrawn.

Controllability. By announcing (or withdrawing) IP prefixes a BGP peer controls potential incoming traffic. Announcing an IP prefix to a BGP neighbor signals that a feasible route is available for use by this neighbor [18]. Whether this route is accepted or dropped is the decision of the neighbor but the announcing AS indicates that it is willing to forward traffic to this prefix. Any default route undermines this implicit agreement and limits control of the announcing AS. Depending on the peering relations, this might have economical implications. In a customer-provider relation, a customer deploying a default route will still pay for traffic. In case of a peer to peer agreement, it will lead to imbalanced traffic shares.

Accountability. In many inter-domain measurement scenarios, default routes conflict with the common assumption that connectivity for a prefix is only available if this prefix (or a covering prefix)

has been announced. For example, studies that identify BGP zombies [5] need reliable data to determine whether a BGP zombie is present or a default route is responsible for providing connectivity.

RPKI Route Origin Validation (ROV) data plane measurements are another example [25]. Here, two prefixes are announced on the control plane during the experiments and the RPKI Route Origin Authorization (ROA) state of one of the prefixes is changed to trigger ROV within the AS under test. Traceroute deviations are recorded and conclusions regarding deployment of ROV drawn. If default routes are present they will provide connectivity while the actual experiment prefix was dropped by the AS. A false negative will be introduced since the AS is performing ROV and dropping invalid prefixes but jeopardizing its own security implementations with the presence of default routes. In their RPKI implementation guidelines [15], some hardware manufacturers explicitly highlight that ROV is only applicable to ASes that are default-free. The deployment of default routes limits the intention of deploying ROV. Therefore, default routes should be removed if there is broader agreement on the usefulness of origin validation among ASes.

5 CONCLUSION AND FUTURE WORK

In this paper, we analyzed the current state of default routing in inter-domain routing. We extended prior work and compared two different methodologies, *path-poisoning* and *not-announced prefix* probing. We discussed inconsistencies in the results and introduced a threshold to give more flexibility while conducting experiments. We found smaller networks to be more likely to deploy default routing. When probing for default routes from the experiment prefix to the AS under test with path-poisoning, we found relatively more default routes for IPv4 compared to IPv6.

In future work, we want to analyze the implications of default routes on BGP measurement results. We will start with RPKI route origin validation.

Artifacts. All artifacts of this paper are available on Github [24] to support reproducible research [27].

ACKNOWLEDGMENTS

The authors would like to thank Yuchen Jin for his feedback on the difference in AS classification dataset sizes and our classification approach. We would also like to thank Florian Steuber, Steven Bellovin, and Clemens Mosig for discussions. We also thank the anonymous reviewers and our shepherd Joel Sommers for their constructive feedback. We thank the RIPE NCC team for lifting our measurement limits to give us the freedom necessary to perform large scale traceroute campaigns and the RIPE Atlas community for supporting our research with credits necessary to schedule the measurements. Moreover, we thank the NLNOG Ring for providing access to their infrastructure.

This work was partly supported by the European Union within the Horizon 2020 project *CONCORDIA* (grant agreement No 830927), and the German Federal Ministry of Education and Research (BMBF) within the projects *PIVOT* and *PRIMEnt*.

REFERENCES

- [1] Randy Bush, Olaf Maennel, Matthew Roughan, and Steve Uhlig. 2009. Internet Optometry: Assessing the Broken Glasses in Internet Reachability. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement*. ACM, 242–253.

- [2] Ignacio Castro, Juan Camilo Cardona, Sergey Gorinsky, and Pierre Francois. 2014. Remote Peering: More Peering without Internet Flattening. In *Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies*. 185–198.
- [3] Center for Applied Internet Data Analysis (CAIDA). 2021. *The CAIDA AS Relationships Dataset, 20210401.as-rel2.txt*. Retrieved Jan 16, 2021 from <https://publicdata.caida.org/datasets/as-relationships/serial-2/>
- [4] Zakir Durumeric, Eric Wustrow, and J Alex Halderman. 2013. ZMap: Fast Internet-wide scanning and its security applications. In *22nd {USENIX} Security Symposium ({USENIX} Security 13)*. 605–620.
- [5] Romain Fontugne, Esteban Bautista, Colin Petrie, Yutaro Nomura, Patrice Abry, Paulo Gonçalves, Kensuke Fukuda, and Emile Aben. 2019. BGP zombies: An analysis of beacons stuck routes. In *International Conference on Passive and Active Network Measurement*. Springer, 197–209.
- [6] Oliver Gasser, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczynski, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. 2018. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *Proceedings of the 2018 Internet Measurement Conference* (Boston, MA, USA). ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3278532.3278564>
- [7] Github. 2021. *ZMap v6 Fork*. Retrieved Feb 15, 2021 from <https://github.com/tumi8/zmap>
- [8] Caitlin Gray, Clemens Mosig, Randy Bush, Cristel Pelsser, Matthew Roughan, Thomas C Schmidt, and Matthias Wählisch. 2020. BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping. In *Proceedings of the ACM Internet Measurement Conference*. 492–505.
- [9] Cristian Hesselman, Paola Grosso, Ralph Holz, Fernando Kuipers, Janet Hui Xue, Mattijs Jonker, Joeri de Ruiter, Anna Sperotto, Roland van Rijswijk-Deij, Giovane CM Moura, et al. 2020. A Responsible Internet to Increase Trust in the Digital World. *Journal of Network and Systems Management* 28, 4 (2020), 882–922.
- [10] Internet Address Hitlist. 2021. *The ANT Lab: Analysis of Network Traffic, Internet Hitlist it93w*. Retrieved Jan 13, 2021 from https://ant.isi.edu/datasets/readmes/internet_address_hitlist_it93w-20210202.README.txt
- [11] Tomas Hlavacek, Amir Herzberg, Haya Shulman, and Michael Waidner. 2018. Practical Experience: Methodologies for Measuring Route Origin Validation. In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 634–641.
- [12] Yuchen Jin, Colin Scott, Amogh Dhamdhere, Vasileios Giotsas, Arvind Krishnamurthy, and Scott Shenker. 2019. Stable and Practical {AS} Relationship Inference with ProbLink. In *16th {USENIX} Symposium on Networked Systems Design and Implementation ({NSDI} 19)*. 581–598.
- [13] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and KC Claffy. 2013. AS Relationships, Customer Cones, and Validation. In *Proceedings of the 2013 conference on Internet measurement conference*. 243–256.
- [14] A. Marder, M. Luckie, A. Dhamdhere, B. Huffaker, kc claffy, and J. M. Smith. 2018. Pushing the Boundaries with bdrmapIT: Mapping Router Ownership at Internet Scale. In *Proceedings of the 2018 Internet Measurement Conference*.
- [15] Niels Rajjer Melchior Aelmans. 2019. *Day One: Deploying BGP Routing Security*. Juniper Networks Books. Retrieved June 26, 2021 from https://www.juniper.net/documentation/en_US/day-one-books/DO_BGP_SecureRouting2.0.pdf
- [16] RIPE NCC. 2021. *RIPE Atlas probes metadata*. Retrieved May 1, 2021 from <https://ftp.ripe.net/ripe/atlas/probes/archive/2021/05/>
- [17] Ricardo V Oliveira, Beichuan Zhang, and Lixia Zhang. 2007. Observing the Evolution of Internet AS Topology. In *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*. 313–324.
- [18] Y. Rekhter, T. Li, and S. Hares. 2006. *A Border Gateway Protocol 4 (BGP-4)*. RFC 4271. IETF.
- [19] Andreas Reuter, Randy Bush, Italo Cunha, Ethan Katz-Bassett, Thomas C Schmidt, and Matthias Wählisch. 2018. Towards a Rigorous Methodology for Measuring Adoption of RPKI Route Validation and Filtering. *ACM SIGCOMM Computer Communication Review* 48, 1 (2018), 19–27.
- [20] RIPE NCC. 2019. *RIPE Atlas Cousteau*. <https://github.com/RIPE-NCC/ripe-atlas-cousteau>. [Online; accessed 10-January-2020].
- [21] RIPE NCC. 2019. *RIPE Atlas Sagan*. <https://github.com/RIPE-NCC/ripe.atlas.sagan>. [Online; accessed 10-January-2020].
- [22] RIPE NCC. 2020. *RIPE Routing Information Service (RIS)*. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>. [Online; accessed 16-October-2020].
- [23] RIPE NCC Staff. 2015. *RIPE Atlas: A global Internet Measurement Network*. *Internet Protocol Journal* 18, 3 (2015).
- [24] Nils Rodday. 2021. *TAURIN-21 workshop paper artifacts*. Retrieved June 26, 2021 from <https://github.com/nrodday/TAURIN-21>
- [25] Nils Rodday, Italo Cunha, Randy Bush, Ethan Katz-Bassett, Gabi Dreo Rodosek, Thomas C. Schmidt, and Matthias Wählisch. 2021. Revisiting RPKI Route Origin Validation on the Data Plane. In *Proc. of Network Traffic Measurement and Analysis Conference (TMA)*. IFIP. accepted for publication.
- [26] Oregon RouteViews. 2013. University of Oregon RouteViews project. *Eugene, OR*. [Online]. Available: <http://www.routeviews.org> (2013).
- [27] Quirin Scheitle, Matthias Wählisch, Oliver Gasser, Thomas C. Schmidt, and Georg Carle. 2017. Towards an Ecosystem for Reproducible Research in Computer Networking. In *Proc. of ACM SIGCOMM Reproducibility Workshop*. ACM, New York, NY, USA, 5–8.
- [28] Brandon Schlinker, Kyriakos Zarifis, Italo Cunha, Nick Feamster, and Ethan Katz-Bassett. 2014. PEERING: An AS for Us. In *Proceedings of the 13th ACM Workshop on Hot Topics in Networks*. ACM, 18.
- [29] Jared M Smith, K Birkeland, T McDaniel, and M Schuchard. 2020. Withdrawing the BGP Re-Routing Curtain. In *Network and Distributed System Security Symposium (NDSS)*.
- [30] Kotikalapudi Sriram and Doug Montgomery. 2019. *Resilient Interdomain Traffic Exchange: BGP Security and DDoS Mitigation*. NIST Special Publication 800-189. National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-189>
- [31] Jin Yuchen. 2021. *ProbLink AS Relationships Dataset, 20210401.txt*. Retrieved Jan 12, 2021 from <https://yuchenjin.github.io/problink-as-relationships/>
- [32] Zhengwei Zhao and Jingping Bi. 2013. Characterizing and Analysis of the flattening Internet Topology. In *2013 IEEE Symposium on Computers and Communications (ISCC)*. IEEE, 000219–000225.