# A First Look into Long-lived BGP Zombies

Iliana Xygkou
Cisco ThousandEyes
Heraklion, Greece
Georgia Institute of Technology
Atlanta, USA
ixygkou3@gatech.edu

Antonis Chariton
Cisco ThousandEyes
Zürich, Switzerland
acharito@cisco.com

Xenofontas Dimitropoulos
Cisco ThousandEyes
Heraklion, Greece
chxenofo@thousandeyes.com

Alberto Dainotti
Georgia Institute of Technology
Atlanta, USA
dainotti@gatech.edu

## Abstract

BGP is the de facto protocol used to manage a network's reachability on the Internet. Network operators announce and withdraw their prefixes on BGP to enable or to prevent communication towards their origin network, respectively. However, the withdrawal of a prefix could fail to propagate totally in the Internet and routes towards withdrawn prefixes could remain in the routing tables of routers. These routes are called stuck or zombie BGP routes, and their persistence can lead to performance degradation, or even partial or complete outage. In this paper, we first revisit existing work on BGP zombies using RIPE RIS beacons, identify the *double-counting* discrepancy, and revise the methodology to address this problem and detect zombies more accurately. Second, we point out limitations of the RIPE RIS beacons with respect to their periodicity, lack of diversity, and noise, and introduce and deploy our own beacons, which address these limitations. Using our beacons and the revised methodology, we analyze the lifespan of BGP zombies. We show that zombie routes can persist in RIBs for days, weeks, or even months. Furthermore, we document that BGP zombies can be announced months after their original withdrawal, affecting new ASes. Finally, we discuss interesting cases of long-lived zombie outbreaks that affected large ISPs with hundreds of ASes in their customer cones.

## CCS Concepts

• **Networks** → **Network measurement**; **Routing protocols**.

## Keywords

BGP, Routing Security, BGP Zombies.

**ACM Reference Format:**
Iliana Xygkou, Antonis Chariton, Xenofontas Dimitropoulos, and Alberto Dainotti. 2025. A First Look into Long-lived BGP Zombies. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25), October 28–31, 2025, Madison, WI, USA.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3730567.3764469


ACM ISBN 979-8-4007-1860-1/2025/10
https://doi.org/10.1145/3730567.3764469

## 1 Introduction

One commonly discussed problem among network operators is stuck or *zombie* BGP routes, which falsely indicate that a prefix is still reachable even though the origin AS has withdrawn the associated route. Zombie routes can occur due to misconfigurations, software bugs, or even BGP protocol flaws [2, 27] that prevent routers from withdrawing or updating routes properly in their BGP routing tables. Zombie BGP routes can potentially lead to suboptimal routing decisions, network instability, and disruptions in traffic flow within a network, causing operational issues such as performance degradation and outages. Furthermore, zombie routes unnecessarily increase the size of the global BGP table, consuming more memory and processing power.

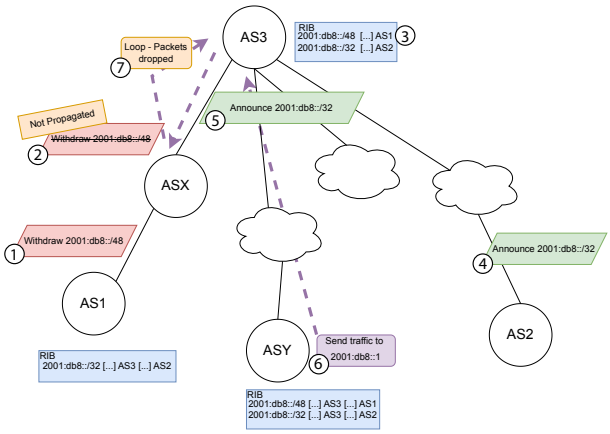


**Figure 1: Example of partial outage due to a zombie route.**

In Fig. 1, we provide an example of a zombie route that could lead to a partial outage. AS1 advertises only 2001:db8::/48, even though it owns its covering prefix 2001:db8::/32. At some point, AS1 sells the /32 prefix to AS2, and thus ① AS1 stops advertising the /48. ASX despite removing the prefix from its routing table ② fails to propagate the withdrawal further to AS3, and thus ③ AS3 will retain the route in its table. Once ④ AS2 starts announcing the /32, ⑤ its announcement will propagate to the rest of the ASes. However, since the /48 route of AS1 did not get completely withdrawn and instead remained in the dominant AS3 (e.g., Tier

1, IXP), the latter will forward the Internet traffic destined to the /48 route based on the zombie route towards AS1, causing a partial outage to AS2. For example, ⑥ if a user within ASY starts sending traffic towards the IP address 2001:db8::1, because of Longest Prefix Matching, the traffic will be forwarded along the path *ASY [...] AS3 ASX AS1*. However, ASX has only the route towards the prefix 2001:db8::/32, and therefore when the traffic reaches it, ⑦ it will forward it back to AS3 creating a loop. Eventually, the routers will drop the packets when they decrement the Hop Count or Time to Live field to zero.

Previous research by Fontugne *et al.* [4, 17] has explored the presence and characteristics of BGP zombies both in a dedicated infrastructure and in the wild. They first find that over the course of 5.5 months, there had been 5,115 zombie outbreaks for the RIPE RIS beacons [4, 15]. They then study 6 years of historical BGP data and report 486k zombie outbreaks for arbitrary prefixes [17].

In this paper, we first reproduce the BGP zombie analysis of Fontugne *et al.* [4] and identify a discrepancy due to double-counting BGP zombies over multiple beacon intervals, which we show can significantly overestimate the total number of zombies. We therefore introduce a revised zombie detection methodology that relies solely on RIPE RIS raw data. Second, we highlight key limitations of methods based on RIPE RIS beacons due to their periodicity, lack of prefix diversity, and noise. To address these limitations, we introduce a new beaconing methodology and deploy it to advertise 96 different IPv6 prefixes per day from an origin AS to more than 1,700 directly connected networks.

Next, we ask the question: What is the lifespan of BGP zombies? How long do they stay in the routing tables? Using our methodology and beacons, we conduct the first study of long-lived BGP zombies and show that zombie routes can persist in RIBs for long periods. Due to unavailability of IPv4 space, our beacons consist only of IPv6 prefixes. We find that 31.4% of the detected zombies remain alive for more than three hours. In addition, zombies can persist for a very long time, up to 8.5 months! Furthermore, we document for the first time that ASes affected by these stuck routes can (re)announce them at a later stage thus affecting new ASes even months after the initial withdrawal. We call this phenomenon BGP zombie *resurrection*. Finally, we discuss interesting cases of persistent zombie outbreaks and pinpoint their root causes.

In summary, we make four contributions: *(i)* we reproduce previous work and identify methodological issues that we address; *(ii)* we highlight key limitations of the RIPE RIS beacons and introduce and deploy a new beaconing methodology with IPv6 prefixes to address these limitations; *(iii)* we conduct the first IPv6-only analysis of the lifespan of BGP zombies; and *(iv)* document for the first time the BGP zombie *resurrection* phenomenon.

## 2 Related Work

In literature, BGP beacons are used as controlled, scheduled prefix announcements and withdrawals to actively measure and analyze Internet routing dynamics. Key BGP beacon projects include the extensive RIPE Routing Information Service (RIS) providing geographically diverse, scheduled beacons [15] for general research, and specialized initiatives like NLnet Labs' RPKI beacons [16] and NLNOG RING's Large BGP Communities beacon [9], designed to test specific security mechanisms or protocol features. Since BGP handles prefixes independently from whether they attract traffic or not, beacons help identify routing issues such as stuck routes that will likely affect actual traffic. BGP beacons have been used to study convergence time, route flap damping, path visibility, policy impacts, and BGP anomalies like zombies [4, 7, 8, 10–12].

Researchers [4, 17] have studied the BGP zombies phenomenon to a certain extent. In a first work, Fontugne *et al.* [4] detect BGP zombies using RIPE RIS beacons. The RIPE RIS beacons are IPv4/IPv6 prefixes that specific RIPE RIS collectors announce every four hours and withdraw two hours later. Fontugne *et al.* define a stuck or *zombie BGP route* or simply *BGP zombie* as a prefix route that remains in the RIBs of some routers despite the prefix's withdrawal by the origin AS. Consistently, they define a *zombie outbreak* as the set of all the zombie routes of the same prefix within the same defined time interval. They kept track of the state of RIPE RIS beacons for all RIPE RIS peers and detected a zombie if a withdrawn beacon was stuck in at least one RIPE RIS peer after 1 hour and 30 minutes. They find that over the course of approximately 5.5 months (split in 3 continuous intervals) there had been 5,115 zombie outbreaks for the RIPE RIS beacons. The RIPE RIS beacons were 14 IPv6 and 13 IPv4 prefixes at the time of their experiments. They further executed timely traceroutes towards the stuck prefixes from RIPE Atlas probes [24] located in ASes from the AS paths of zombie routes. Based on the AS paths of both the withdrawn and the stuck routes, for each zombie outbreak the authors build the AS graph and classify unknown ASes as zombie ASes or normal ASes. They validate the outcome using the traceroute results and they make their software and traceroute results publicly available [6].

In [17], Ongkanchana *et al.* advance their findings by analyzing prefix withdrawals in the wild and detecting how many of them lead to BGP zombies. They study 6 years of historical BGP data and apply a simple threshold heuristic based on the withdrawal propagation time and RIPE RIS peers visibility in order to classify BGP withdrawals belonging to local topological changes or complete prefix withdrawals. They report 486k zombie outbreaks and observe that 3.22% of them are due to the RIPE RIS beacons. They thus argue that noisy prefixes such as beacons are more prone to get stuck than regular prefixes. Our approach takes this issue into account by reusing the same beacon prefixes less frequently and using fresh new prefixes.

In recent work [1], Anahory *et al.* propose Route Status Transparency (RoST). This design enables ASes to verify a route's status through a public transparency repository and detect suppressed withdrawals, allowing them to eliminate BGP zombie routes.

Our work is the first study to analyze the lifespan of BGP zombies. For this purpose, we develeop a new beaconing methodology. Moreover, we identify a limitation in the methodology of [4] that leads to double-counting zombies, and reproduce their results with a revised methodology that addresses this limitation. Contrary to [17], we focus on BGP zombies produced from beacon prefixes, rather than in the wild. Using BGP beacons we obtain accurate results, since we know exactly when they are announced and withdrawn.

# 3 Replication of previous study

## 3.1 Methodology

Similarly to [4], we identify a zombie outbreak if a withdrawn beacon is still in at least one RIPE RIS peer 1 hour and 30 minutes after its withdrawal. Our methodology differs from [4] in three ways. First, we track the removed or present state of beacon prefixes in RIPE RIS peers solely from RIPE RIS raw data [23], which facilitates historical state monitoring at message-level granularity. Instead, previous research used the real-time RIPEstat looking glass [14] to be able to execute timely traceroutes. Second, we detect and remove duplicate BGP zombies, *i.e.,* zombies that persist along multiple beacon intervals and are counted multiple times. Third, we detect and remove noisy RIPE RIS peers. Previous work has not tackled the problems of double-counting zombies and noisy peers.

*1. Reconstructing the state of a prefix.* Previous work [4] used the RIPEstat looking glass [14] to identify stale prefixes in real time and later filtered out false positives based on RIPE RIS raw data. The RIPEstat looking glass is a RIPE service that provides the routing state of RIPE RIS peers. However, the RIPEstat looking glass is a "black box" service. We do not know how the state is computed, if it is updated in real time, and what is the precise update delay. For example, if the service state is updated with a delay of a few minutes, then checking the state of a fully withdrawn prefix before the service is updated would lead to false positives. Finally, based on our direct communication with RIPEstat engineers, over time the service has gone through updates that could have affected the previous work's results [19–22]. Instead, in our zombie detection analysis we use solely archived RIPE RIS raw data to process all the related historical BGP data at message-level granularity for most accurate results. From all the available RIPE RIS peers, we collect the BGP UPDATE messages from RIPE RIS raw data [23] associated with the RIPE RIS beacons. Additionally, we collect the STATE messages which report changes of the state of the session between a RIPE RIS peer and a RIPE RIS collector. With these, we are able to reconstruct the state of a prefix (present or removed) at any RIPE RIS peer at a specific time point. We divide the BGP UPDATE messages into 4-hour intervals that start at the RIPE RIS beacons announcement times. We process each interval independently, *i.e.,* without any prior knowledge about the routing state of the beacon prefixes. As a result, we ignore any stale RIB entries from previous announcements, *i.e.,* a zombie route of a prefix announced at 00:00 should not affect our results for the same prefix announced at 04:00.

*2. Eliminating double-counting.* To uniquely identify zombie outbreaks, we also take into account the *Aggregator IP Address* BGP attribute. This attribute is populated in RIPE RIS beacons as "10.$x$.$y$.$z$", where $x$, $y$ and $z$ represent the 24-bit count of the number of seconds between midnight UTC on the 1st day of the month and the time of the BGP announcement. E.g., we found a BGP announcement for a RIPE RIS beacon with the following attributes:

- *Timestamp on which the collector received the announcement message*: 2018-07-19 02:00:02 UTC
- *Aggregator IP Address*: 10.19.29.192, which translates to 1,252,800 seconds after 2018-07-01 in the best case scenario [1]. This indicates that the announcement was first originated on 2018-07-15 12:00 UTC, *i.e.,* 3.5 days before the one of the current 4-hour time interval (2018-07-19 00:00 UTC).

If we observe a stuck route belonging to a previous announcement, we do not consider this a new zombie, as it should have already been counted in a previous interval. Without taking into account the Aggregator IP Address, one would overestimate the number of zombies. In the above example, a single stuck route would be counted as 21 separate zombie outbreaks. In Section 3.2, we show that double-counting has a significant impact on the estimated number of outbreaks.

*3. Removing noisy peers.* We detected an outlier RIPE RIS peer that caused an abnormally high number of zombies. For reasons and with methodology explained in Sec. 3.2 we ignore this peer in the rest of our reproduction analysis to avoid overestimating the number of zombies.

In summary, our methodology has the following differences from that of [4]: *(i)* we rely solely on RIPE RIS raw data to track the state of beacons, *(ii)* we use the BGP Aggregator IP Address attribute to eliminate double-counting BGP zombie outbreaks, and *(iii)* we filter a noisy RIPE RIS peer.

## 3.2 The impact of double-counting and of a noisy peer

We replicate the analysis of BGP zombies characteristics in [4]: the zombie emergence rate, AS path lengths of normal and zombie routes, and the number of concurrent zombie outbreaks (see Appendix B.2 for details). In this section, we analyze the impact of double-counting and of a noisy peer using archived BGP UPDATE messages collected by RIPE RIS for the same three time periods as the previous study by Fontugne *et al.* [4]: 2018-07-19 – 2018-08-31, 2017-10-01 – 2017-12-28, 2017-03-01 – 2017-04-28.

First, we discuss the impact of filtering with the BGP Aggregator IP Address attribute. In the columns "With double-counting" and "Without double-counting" of Table 1, we present our results *(i)* if we do not filter with the Aggregator IP Address BGP attribute, and *(ii)* if we filter with it. Filtering has a major effect on detected zombie outbreaks, a reduction of 21.36%, which incidentally shows that stuck routes can remain for several days. The impact of this process is evident: both IPv4 and IPv6 zombie outbreaks experience a significant reduction of 57.8% and 31% respectively, for the time period of July-August 2018, whereas in the case of the 2017 periods the corresponding reduction decreases to 32.76% for IPv4 and is minimal for IPv6.

| Period (#visible prefixes) | With double-counting | | Without double-counting | |
|---|---|---|---|---|
| | IPv4 | IPv6 | IPv4 | IPv6 |
| Jul 19 - Aug 31, 2018 (7126) | 536 | 745 | 226 | 514 |
| Oct 01 - Dec 28, 2017 (14336) | 705 | 1378 | 478 | 1370 |
| Mar 01 – Apr 28, 2017 (9556) | 1781 | 610 | 1319 | 610 |

**Table 1: Comparison of the estimated numbers of zombie outbreaks with and without double-counting, for each time period of the experiment in [4]**

[1]The BGP attribute is relative to the beginning of each month. It is possible that the announcement was originated on the 15th day of any previous month before July 2018.

BGP collection platforms such as RIPE RIS [25] and RouteViews [28] consist of numerous route collectors that establish BGP sessions with volunteer ASes, called peers, and collect and store the exchanged BGP UPDATE messages. Unfortunately, there have been reported cases with misbehaving peers that accidentally propagate bogus routes that pollute the BGP data [3] [2] or with peers using BGP features not yet supported by the collectors resulting in corrupted records such as FRR not supporting ADD-PATH encodings in MRT [26]. In our analysis, we found that the RIPE RIS peer AS16347, Inherenet Adista SAS, connected to the collector RRC21, has a probability of approximately 42.8% of having a zombie IPv6 prefix. This probability remains high for IPv6 (42.6%) even after we eliminate double-counting. Since the remaining peers have an average probability of 1.58% of having a zombie IPv6 prefix, we consider this RIPE RIS peer AS an outlier and exclude it from the rest of our analysis to prevent overestimating zombies. We provide detailed numbers in Appendix B.2 and in Table 4 in Appendix C.

## 4 Our BGP beacons methodology

In this section, we describe and implement a new BGP beacon methodology to periodically announce and withdraw prefixes. With respect to the format of the beacon prefixes, we applied two different approaches which affected the prefixes' announcement frequency and varied the extensiveness of the study of a zombie's lifetime. We will first outline the reasons that led us to implement our own methodology on beacon prefixes instead of using RIPE RIS beacons.

*Periodicity.* All RIPE RIS beacons announce the same prefixes every 4 hours, whereas our beacon recycles its prefixes every 24 hours and every 15 days in our first and second approach respectively. This enables us to study temporal properties of stuck routes: how long a stuck route would survive (*e.g.*, multiple days), and whether any changes to the prefix visibility can be observed over a long period of time without triggers from the origin AS.

*Multitude and diversity.* RIPE beacons are only 22 IPv6 prefixes and 3 IPv4 prefixes, while we advertise and track 96 different IPv6 prefixes per day. This greater number of announced prefixes increases diversity and clarity about which prefix becomes stuck, since our beacons do not interact with each other within the course of at least one day. Moreover, with more unique prefixes we raise the likelihood that at least one becomes a zombie, thereby improving our ability to uncover potential issues. Due to the global unavailability of IPv4 prefixes at the time of the experiments, we use only IPv6 ones. With current market prices, we would require over $500,000 worth of assets to conduct the equivalent experiment in IPv4.

*Noise and proneness to zombies.* As discussed in Sec. 2, researchers claimed that noisy prefixes such as the RIPE RIS beacons are more prone to zombies (and thus not representative) because they are announced and withdrawn multiple times per day for many years [17]. On the contrary, our prefixes appear on BGP for the first time in our study, and we recycle them at most once per day. This way, they better approximate common BGP prefix withdrawals.

*Experimental setup.* Our BGP beacons are IPv6 /48 prefixes and belong to an (already advertised) /29 IPv6 prefix from a personal AS, *AS210312*. These /48 prefixes have not been advertised in the past before our experiments. *AS210312* announced the prefixes from *all* its available Points of Presence (i.e., not selectively), to more than 1,700 directly connected networks. We also registered an appropriate RPKI ROA to make sure the prefixes are RPKI-valid and thus not filtered due to Route Origin Validation. Every 15 minutes (at :00, :15, :30, :45) we announce a different prefix, which we withdraw 15 minutes later and re-announce only after respectively 24 hours or 15 days in our first and second approach (recycle time). The timestamp of the announcement is encoded in the bits of the prefix (resembling a *BGP Clock*) using a different format depending on the recycling approach: "`2a0d:3dc1:(HHMM)::/48`", (where HHMM is the timestamp of the announcement) for 24-hours recycled prefixes, and "`2a0d:3dc1:(HH)(minute+day%15)::/48`" for 15-days recycled ones[3]. We ran experiments with the first approach from 2024-06-04 11:45 to 2024-06-10 09:30 UTC, and with the second one from 2024-06-10 11:30 to 2024-06-22 17:30 UTC. A 15-days recycle period means that an announcement (and withdrawal) of a beacon prefix can wipe out a stuck route only after 15 days, thus allowing us to to detect and analyze zombie routes that persist for a week or more. We note that our experiments concluded before the expiration of the 15-day interval, allowing us to study long-lived zombies. Nevertheless, infrequent beacon re-announcement is a key feature of a long-term beacon service (Sec. 6) that enables studying zombies persisting for up to 15 days.

## 5 Long-lived BGP zombies

In this section, we first detect zombies caused by our beacons over an 18-day interval using RIPE RIS raw data [23]. Due to limited resources, we do not include BGP data from RouteViews [28] peers acknowledging the potential omission of zombie routes. We collect BGP updates related to our advertised prefixes, and examine if a prefix gets stuck on a RIPE RIS peer. We consider a prefix as stuck at a RIPE RIS peer if after a defined threshold since the prefix's withdrawal from the originating AS, the last BGP update is not a withdrawal. Consistently with prior work [4, 17], we conservatively identify as zombies only routes that were stuck for at least 90 minutes. However, differently from these prior studies, we are able to capture long-lived zombies that last up to months, i.e., way beyond the 2-hours limit imposed by the re-announcement after withdrawal performed by the RIPE RIS beacon infrastructure. We acknowledge the limitation—common to all literature—of excluding from our scope potential short-lived zombies that might last less than 90 minutes. Second, we process the RIBs of all RIPE RIS peers to study how long detected zombies stay alive over approximately a year, from 2024-06-04 to 2025-05-09. RIPE RIS publishes RIB dumps of all RIPE RIS peers every 8 hours and therefore provides a coarser granularity that allows us to scale our analysis.

In our results, we find that BGP zombies still appear regularly. In Fig. 2, we show with respect to the value of the variable threshold the total number of zombie outbreaks (right axis) and the percentage of the beacon announcements that lead to zombie outbreaks (left axis) for *(i)* all RIPE RIS peers, and *(ii)* all RIPE RIS peers except

[2] It is possible for the—often longstanding— router configurations of BGP peering sessions with RouteViews/RIS collectors to not be monitored as carefully as those used for production traffic, e.g., by not being updated consistently with other policy changes within the network, thus causing misconfigurations.

[3] The second approach had an implementation bug: On some days, 2 out of the 96 different prefixes to be announced would be the same. For example: on 2024-06-15 the prefix of 00:30 and 03:00 would be the same: 2a0d:3dc1:30::/48. In these cases, we study only the latter prefix.

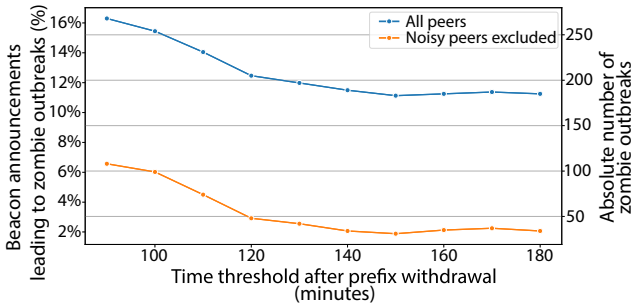


**Figure 2: Number of zombie outbreaks (right axis) and the percentage of the beacon announcements that lead to zombie outbreaks (left axis) for *(i)* all RIPE RIS peers, and *(ii)* all RIPE RIS peers except for the noisy peers, with respect to the time threshold (minutes) after the beacons' withdrawal.**

for three noisy peers. In the case of the noisy peers' exclusion (orange line), we observe that the proportion and total number of BGP zombies stabilize at approximately 2% and 34 over time, respectively. These numbers keep decreasing after the threshold value of 90 minutes, at which point the values are 6.6% and 108. Therefore, with the extended interval of three hours, only 31.4% of the zombies seen at the 90 minute point remain alive.

We observe that three outlier RIPE RIS peer routers[4] from two peer ASes of collector RRC25, AS211380 (SIMULHOST-AS Simulhost Limited, GB), and AS211509 (Rudakov Ihor, UA) contain zombie routes for at least 6.88% of the beacon announcements even 3 hours after the beacons' withdrawal[5]. We provide detailed numbers in Table 5 in Appendix C. Furthermore, in Fig. 2 we highlight that after 3 hours since the beacons' withdrawal, the remaining ~ 670 RIPE RIS peers collectively contribute only 34 zombies. This number grows significantly to more than 170 when we include the three outlier peers. Thus, we present our results both with and without the noisy peers and focus on the latter to avoid zombies' overestimation.

In Fig. 3, we show the CDF of the duration of zombie outbreaks (that last at least one day) for *(i)* all RIPE RIS peers and *(ii)* all RIPE RIS peers except for the three noisy peers. It is evident that stuck routes can persist for a very long time, up to 8.5 months! Regarding the orange line *(ii)* , we observe that all outbreaks that last exactly ~ 35 – 37 days are visible from a single peer *2a0c:b641:780:7::feca* from AS207301 and the next AS in the path of the stuck route is a noisy peer AS211509. We note that additional zombies can persist even longer, but may not be visible from the RIPE RIS peers.

In addition, on 2024-06-22 19:49 UTC we removed the ROA that protected our beacon prefixes. Thus, beacon routes after that point are RPKI invalid. Fig. 3 shows that there are ASes with zombie routes that do not evict the RPKI invalid routes from their RIB even after an expected ROA deletion delay has passed [5]. These ASes do not perform ROV yet, or their ROV implementation is flawed or does not comply with RPKI standards [13].

[4]The IP addresses of the peer routers are 2a0c:9a40:1031::504, 2001:678:3f4:5::1, and 176.119.234.201. The last IP address is IPv4 because this particular peer is exchanging IPv6 Address Family data over an IPv4 BGP session.

[5]Note that these 3 noisy peers refer to our beacon experiment and are different from the noisy peer in the replication analysis in Sec. 3.2.

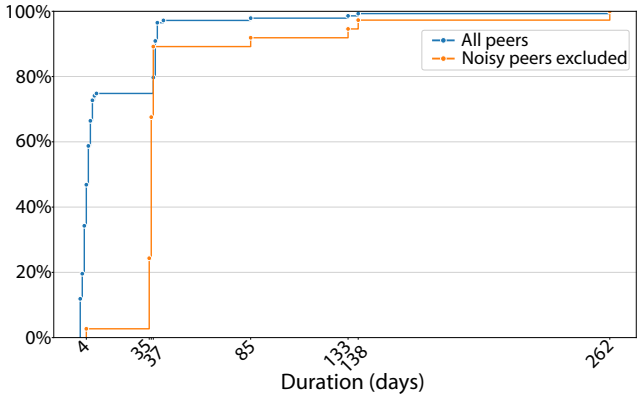


**Figure 3: CDF of the duration of zombie outbreaks (considering only those lasting at least 1 day) for *(i)* all RIPE RIS peers, and *(ii)* all RIPE RIS peers except for the noisy peers. X-axis values are displayed only for line *(ii)* .**

## 5.1 Resurrected Zombies!

In Fig. 2, even though we would expect the proportion to be a decreasing function due to the propagation of withdrawal messages, we notice an increasing trend after 160 minutes. We find that a small amount of prefixes was withdrawn by specific peers before 150 minutes, but appeared again 20 minutes later (170 minutes after the prefixes' withdrawal) due to the reception of a new announcement. Interestingly, all newly stuck routes share the common subpath "4637 1299 25091 8298 *210312*". This indicates that AS4637 (Telstra Global, HK), a Tier 2 provider in Hong Kong with ~ 6000 ASes in its customer cone could be the root cause of these late emerged zombie routes in the RIPE RIS peers. Furthermore, in Fig. 3 all outbreaks that lasted ~ 35 – 37 days became visible by the RIPE RIS peers ~1 month after our last beacon withdrawal, indicating that zombie routes can be announced to new ASes even long after their withdrawal.

We call this phenomenon *resurrection* of zombie routes and are the first to document it in the literature. Potential reasons for such incidents are router bugs, BGP session resets, changes in filters, some of which occur during normal operation. If a downstream session of an infected router is reset, new announcements are generated for these stuck prefixes [18].

A notable example of a zombie route that was resurrected is *2a0d:3dc1:1851::/48*. In Fig. 4 we show the periods during which this prefix was withdrawn and then resurrected. The prefix was fully withdrawn by all RIPE RIS peers on 2024-06-21, but then appeared again in a RIPE RIS peer's RIB a week later on 2024-06-29 without a new beacon announcement! Furthermore, the zombie route was visible for ~3 months (until 2024-10-04) before its withdrawal by the RIPE RIS peer and then re-appeared ~2 months later for another ~3.3 months in the same peer (from 2024-11-29 to 2025-03-11). In total, the prefix was stuck for ~ 8.5 months. The path was "61573 28598 10429 12956 3356 34549 8298 *210312*".

## 5.2 Cases of ASes that were impacted

In this section, we point out two cases of persistent zombie outbreaks that we observe with our beacons to show that zombie

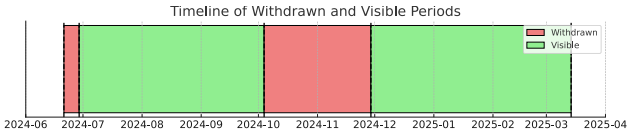


**Figure 4: Timeline of a BGP zombie prefix becoming invisible and resurrecting twice over a period of 8.5 months in the RIPE RIS peers.**

outbreaks *(i)* can have a significant impact in terms of the number of infected RIPE RIS peers and the customer cone size of the likely responsible AS, and *(ii)* can persist for a very long time.

We attempt to pinpoint the ASes that are possibly responsible for these zombie outbreaks using the concept described in [4]. The AS graph constructed based on the AS Paths of the zombie routes has the typical form of a "palm tree". Starting from the origin AS (root) there is a single chain of ASes which eventually branches out to multiple subtrees. The last AS of this chain could be the one propagating zombie routes. We note that this methodology attempts to pinpoint the cause of a zombie outbreak solely based on AS Paths. Moreover, the identified AS is not necessarily responsible for the zombie outbreak, since the previous AS in the path could have failed to propagate the withdrawal to it, or there can be "invisible" ASes like Internet Exchange Point Route Servers. We leave the improvement of the root cause AS inference algorithm and the characterization of root cause ASes as future work.

*Impactful zombie.* The prefix *2a0d:3dc1:2233::/48* stayed in the RIBs of 24 peer routers and 21 peer ASes even after 3 hours since its withdrawal. All routes shared the common subpath "33891 25091 8298 *210312*", and AS33891 (Core-Backbone GmbH, DE), a Tier 2 provider in Germany with ~ 2100 ASes in its customer cone, might have caused these stale routes. The beacon disappeared from all the peers' RIBs 4 days later.

*Extremely long-lived zombie.* The prefix *2a0d:3dc1:163::/48* remained in the RIBs of two peers AS9304 and AS17639 for ~4.5 months (from 2024-06-18 until 2024-11-03), and in the RIB of a third peer AS142271 for ~4 months (from 2024-06-23 until 2024-10-25). All routes share the common subpath "9304 6939 43100 25091 8298 *210312*", indicating that AS9304 (HGC Global Communications Limited, HK) with more than ~750 ASes in its customer cone, could be responsible for this exceptionally persistent zombie outbreak.

## 6 Discussion & Future Work

*Comparing and combining RIPE RIS beacons & our beacons.* Due to the unavailability of IPv4 space, our beacons consist only of IPv6 prefixes. IPv4 prefix offers only a limited number of bits for timestamp encoding and has only a few more specific prefixes (up to /24) that can be used as beacons. Thus, a compact encoding schema of the announcement time is necessary to maximize space utilization. Using both RIPE RIS beacons and our beacons (enhanced with IPv4 prefixes), we plan to explore how specific characteristics—such as the number and geographic distribution of origin ASes, the frequency of announcements, and the address family of the beacons—affect the BGP zombies phenomenon and whether their combination (including BGP data from RouteViews peers [28]) could facilitate more accurate detection of the causes of the outbreaks.

*Real-time detection of BGP zombies.* In this paper, we use historical BGP data to detect and analyze zombie outbreaks that occurred during a specific period of time. However, a stuck route can immediately affect the way traffic is routed. Real-time detection of a zombie outbreak and identification of the AS causing it, will notify the network operators of the infected ASes to examine and resolve the issue more quickly in order to minimize the adverse effects of the stuck routes.

*Our beacons as a long-term service.* We have presented this work at network operator meetings and received feedback: operators acknowledged the current existence of the phenomenon, expressed interest in discovering where and why zombies occur, and requested continued operation of our beacons. Running our beacons as a long-term service will provide research data for both real-time and historical analysis of the BGP zombies and other routing phenomena. While there are various possible causes for the appearance of BGP zombies—*e.g.,* previous work identified a software bug in the handling of a BGP peer with a 0 sized TCP window [2, 27];others have speculated about bugs and issues associated with BGP optimizers and route reflectors [4]—timely empirical evidence to help pinpoint them is still unavailable. In the future we could deploy an online detection platform for this type of phenomena, which will facilitate the prompt investigation of zombie routes and their causes.

## 7 Conclusion

In this paper, we quantified the emergence of BGP zombies both with RIPE RIS beacons and our new beaconing methodology. We revisited previous research work [4] and revised their methodology to obtain more accurate results. By relying on RIPE RIS raw data [23] and excluding a noisy peer, we found 12.51% more zombie outbreaks. However, when we filtered out older zombie routes with the Aggregator IP Address attribute, we found in total 13% fewer zombie outbreaks. This highlights the impact of the detection methodology on the final assessment of zombie outbreaks. Furthermore, we examined the prevalence of BGP zombies in today's Internet with our beacons and showed that even three hours after their withdrawal 2% of the prefixes remained in the RIBs of RIPE RIS peers. For the first time, we presented the concept of zombie resurrection where ASes can be infected again with zombie routes despite their initial withdrawal and without a new beacon announcement. Finally, we presented two cases where the root cause ASes are impactful and the zombie routes persisted for days or even months, despite the timely removal of the covering ROA. Besides the straightforward consequences of network disruptions and instability, such incidents also raise security concerns around the implementation of ROV.

## 8 Acknowledgments

This scientific paper is partially supported by the Onassis Foundation - Scholarship ID: F ZT078-1/2023-2024. This entire work has been completed while Iliana had a full time internship at Cisco ThousandEyes. All BGP prefix announcements and withdrawals have been provided independently by Antonis using personal resources, such as `AS210312`, `AS4601`, and `2a0d:3dc1::/32`. Finally, we would like to thank Romain Fortugne and Emile Aben for their insightful feedback on our replication analysis.

## References


[1] Yosef Edery Anahory, Jie Kong, Nicholas Scaglione, Justin Furuness, Hemi Leibowitz, Amir Herzberg, Bing Wang, and Yossi Gilad. 2025. Suppressing BGP Zombies with Route Status Transparency. In *22nd USENIX Symposium on Networked Systems Design and Implementation (NSDI 25)*. USENIX Association, Philadelphia, PA, 1349–1366. https://www.usenix.org/conference/nsdi25/presentation/anahory

[2] Ben Cartwright-Cox. 2021. Hunting down the stuck BGP routes. https://blog.benjojo.co.uk/post/bgp-stuck-routes-tcp-zero-window. Accessed: 2024-08-07.

[3] Eric. 2022. What's going on with AS147028? https://mailman.nanog.org/pipermail/nanog/2022-July/219943.html. NANOG Mailing List, July 12, 2022. Accessed: 2025-05-05.

[4] Romain Fontugne, Esteban Bautista, Colin Petrie, Yutaro Nomura, Patrice Abry, Paulo Goncalves, Kensuke Fukuda, and Emile Aben. 2019. BGP Zombies: An Analysis of Beacons Stuck Routes. In *Passive and Active Measurement*, David Choffnes and Marinho Barcellos (Eds.). Springer International Publishing, Cham, 197–209.

[5] Romain Fontugne, Amreesh Phokeer, Cristel Pelsser, Kevin Vermeulen, and Randy Bush. 2023. RPKI Time-of-Flight: Tracking Delays in the Management, Control, and Data Planes. In *Passive and Active Measurement: 24th International Conference, PAM 2023, Virtual Event, March 21–23, 2023, Proceedings*. Springer-Verlag, Berlin, Heidelberg, 429–457. https://doi.org/10.1007/978-3-031-28486-1_18

[6] Romain Fortugne. 2018. BGP Zombie: tools and data. https://github.com/romain-fontugne/BGPzombiesSSL/. Accessed: 2024-08-07.

[7] Caitlin Gray, Clemens Mosig, Randy Bush, Cristel Pelsser, Matthew Roughan, Thomas C. Schmidt, and Matthias Wahlisch. 2020. BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping. In *Proceedings of the ACM Internet Measurement Conference* (Virtual Event, USA) *(IMC '20)*. Association for Computing Machinery, New York, NY, USA, 492–505. https://doi.org/10.1145/3419394.3423624

[8] Craig Labovitz, Abha Ahuja, Abhijit Bose, and Farnam Jahanian. 2000. Delayed Internet routing convergence. In *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication* (Stockholm, Sweden) *(SIGCOMM '00)*. Association for Computing Machinery, New York, NY, USA, 175–187. https://doi.org/10.1145/347059.347428

[9] Large BGP Communities Project. 2016. Milestone: Large BGP Communities Beacon Prefixes. Large BGP Communities Project Blog. http://largebgpcommunities.net/2016/beacon/ Accessed: 2025-04-21.

[10] Jun Li, Randy Bush, ZM Mao, Timothy Griffin, Matthew Roughan, Daniel Stutzbach, and Eric Purpus. 2006. Watching data streams toward a multi-homed sink under routing changes introduced by a BGP beacon. In *Passive & Active Measurement (PAM)*.

[11] Z. Morley Mao, Randy Bush, Timothy G. Griffin, and Matthew Roughan. 2003. BGP beacons. In *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement* (Miami Beach, FL, USA) *(IMC '03)*. Association for Computing Machinery, New York, NY, USA, 1–14. https://doi.org/10.1145/948205.948207

[12] Zhuoqing Morley Mao, Ramesh Govindan, George Varghese, and Randy H. Katz. 2002. Route flap damping exacerbates internet routing convergence. In *Proceedings of the 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications* (Pittsburgh, Pennsylvania, USA) *(SIGCOMM '02)*. Association for Computing Machinery, New York, NY, USA, 221–233. https://doi.org/10.1145/633025.633047

[13] Prodosh Mohapatra, John Scudder, David Ward, Randy Bush, and Rob Austein. 2013. BGP Prefix Origin Validation. RFC 6811. https://doi.org/10.17487/RFC6811

[14] RIPE NCC. 2024. RIPEstat: BGP Looking Glass. https://stat.ripe.net/widget/looking-glass. Accessed: 2024-08-07.

[15] RIPE NCC. 2024. Routing Beacons. https://ris.ripe.net/docs/routing-beacons/. Accessed: 2024-08-07.

[16] NLnet Labs. 2020. Route Origin Validation measurements. NLnet Labs Research Projects. https://nlnetlabs.nl/research/projects/ Accessed: 2025-04-21.

[17] Porapat Ongkanchana, Romain Fontugne, Hiroshi Esaki, Job Snijders, and Emile Aben. 2021. Hunting BGP zombies in the wild. In *Proceedings of the 2021 Applied Networking Research Workshop* (Virtual Event, USA) *(ANRW '21)*. Association for Computing Machinery, New York, NY, USA, 1–7. https://doi.org/10.1145/3472305.3472315

[18] Yakov Rekhter, Susan Hares, and Tony Li. 2006. A Border Gateway Protocol 4 (BGP-4). RFC 4271. https://doi.org/10.17487/RFC4271

[19] RIPE NCC. 2018. RIPE NCC Member Update: April 2018. https://mailman.ripe.net/archives/list/ncc-announce@ripe.net/message/OBRGS7ZRD236GO7EIQ67ORTIHBA6W6JS/. Accessed: 2025-05-05.

[20] RIPE NCC. 2018. RIPE NCC Member Update: August 2018. https://mailman.ripe.net/archives/list/ncc-announce@ripe.net/message/KOPSEQUMVZMSZGZPHXVYFPNF3AO5HMUA/. Accessed: 2025-05-05.

[21] RIPE NCC. 2018. RIPE NCC Member Update: February 2018. https://mailman.ripe.net/archives/list/ncc-announce@ripe.net/message/BA726EYPGXB73K4SQAD4MFR5HZDUZLM4/. Accessed: 2025-05-05.

[22] RIPE NCC. 2018. RIPE NCC Member Update: March 2018. https://mailman.ripe.net/archives/list/ncc-announce@ripe.net/message/UCLRWRTJX32UDTXRIRRUMSKYAKNLMOJF/. Accessed: 2025-05-05.

[23] RIPE NCC. 2024. RIS Raw Data: MRT Files. https://ris.ripe.net/docs/mrt/. Accessed: 2024-08-07.

[24] RIPE NCC. 2025. RIPE Atlas. https://atlas.ripe.net/. Accessed: 2025-05-09.

[25] RIPE NCC. 2025. RIPE Routing Information Service (RIS). https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/. Accessed: 2025-05-05.

[26] Aftab Siddiqui. 2022. *RFC 7911 – What happens when routers do not speak the same language*. https://manrs.org/2022/04/rfc-7911-what-happens-when-routers-do-not-speak-the-same-language/ Accessed: 2025-05-05.

[27] Job Snijders, Ben Cartwright-Cox, and Yingzhen Qu. 2024. Border Gateway Protocol 4 (BGP-4) Send Hold Timer. RFC 9687. https://doi.org/10.17487/RFC9687

[28] University of Oregon RouteViews Project. 1995. RouteViews Project. Project Website. https://www.routeviews.org/ Accessed: 2025-04-21.


# Appendix

## A Ethics

This work does not raise ethical issues. For the experiments, the authors advertised unused address space from *AS210312*: both *AS210312* and the address space are under their administrative control. In addition, the stuck routes did not interfere with non-owned routes in the routing table, and did not affect the functionality of the routers.

## B Previous Study [4] Reproduction

We reproduce the results of [4] using archived BGP UPDATE messages collected by RIPE RIS for the same three time periods as [4]: 2018-07-19 – 2018-08-31, 2017-10-01 – 2017-12-28, 2017-03-01 – 2017-04-28.

### B.1 Comparison of results

In the column "Study" of Table 2, we present the estimated numbers of zombie outbreaks of the previous study [4]. It is evident that there are discrepancies between the results of [4] and ours even if we ignore the *Aggregator IP Address* BGP attribute. In total, we find 640 more zombie outbreaks (an increase of 12.51%), and the greatest difference of 321 (an increase of 83.6%) is observed for IPv4 beacons during the period of October-December 2017.

We investigated more into these disparities and counted how many zombie routes we miss (including the ones from the noisy peer), and conversely how many zombie routes the previous research misses based on the data they made publicly available on GitHub [6]. In Table 3, we report the differences in terms of missed zombie routes and missed zombie outbreaks. We note that surprisingly, each side misses zombie routes and outbreaks that the other reports.

### B.2 Reproduction

In Figures 5 and 6, we replicate the analysis of the Fig. 3 in [4]: Fig. 5 shows the CDF of the likelihood of a $< beacon, peerAS >$ pair to have a zombie route (zombie emergence rate), and Fig. 6 shows the CDF of the AS path length of the routes before the beacons' withdrawal (normal path) *(i)* in peers that withdraw the prefix (normal peers), *(ii)* in peers that do not withdraw the prefix (zombie peers), and of the stuck routes after the beacons' withdrawal (zombie path).

| Period | Study [4] | | With double-counting | | Without double-counting | | Total visible prefixes |
|---|---|---|---|---|---|---|---|
| | IPv4 | IPv6 | IPv4 | IPv6 | IPv4 | IPv6 | |
| 2018-07-19 – 2018-08-31 | 520 | 686 | 536 | 745 | 226 | 514 | 7126 |
| 2017-10-01 – 2017-12-28 | 384 | 1202 | 705 | 1378 | 478 | 1370 | 14336 |
| 2017-03-01 – 2017-04-28 | 1732 | 591 | 1781 | 610 | 1319 | 610 | 9556 |

**Table 2: Comparison of the previous study [4] with the estimated numbers of zombie outbreaks with double-counting and without double-counting for each time period of the experiment.**

| Study [4] | | | | Our results | | | |
|---|---|---|---|---|---|---|---|
| Missing zombie routes | | Missing zombie outbreaks | | Missing zombie routes | | Missing zombie outbreaks | |
| IPv4 | IPv6 | IPv4 | IPv6 | IPv4 | IPv6 | IPv4 | IPv6 |
| 4956 | 4374 | 616 | 308 | 22110 | 15169 | 230 | 54 |

**Table 3: Number of missing zombie routes and missing zombie outbreaks in the results of [4] and our results.**

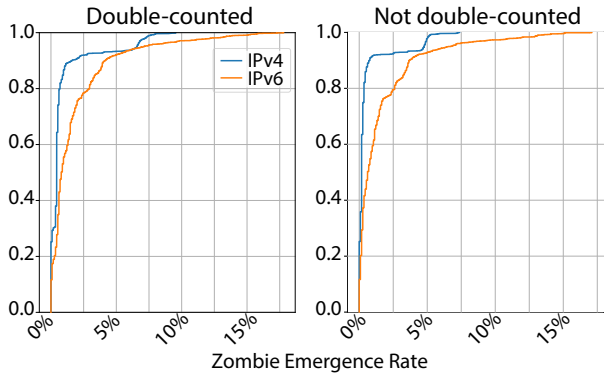


**Figure 5: CDF of likelihood of a <RIPE RIS beacon, peerAS> to have a zombie route (zombie emergence rate) with double-counting and without double-counting.**

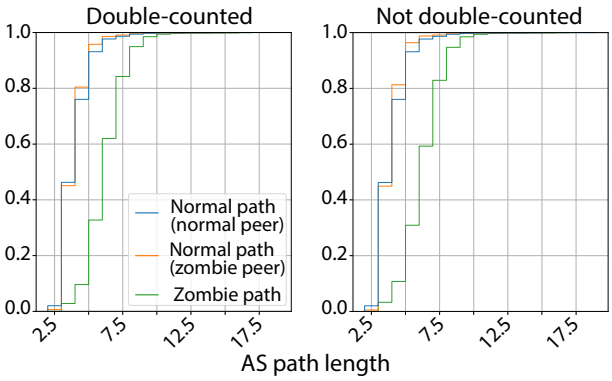


**Figure 6: CDF of the AS path length of the routes before the beacons' withdrawal (normal path) *(i)* in peers that withdraw the prefix (normal peers), *(ii)* in peers that do not withdraw the prefix (zombie peers), and of the stuck routes after the beacons' withdrawal (zombie path). Graphs are shown with double-counting and without double-counting.**

We find that in accordance with the authors of [4] BGP zombies are not highly frequent at RIPE RIS peers. Nevertheless, we outline statistical findings that indicate that BGP zombies are even rarer than previously reported. First, without filtering with the BGP attribute we observe that a significant portion, 18.76% of the $< beacon, peerAS >$ pairs shows no zombie occurrences at all, while 50% of pairs are less than 0.52% likely to cause a zombie route with an average value of 0.88% for IPv4 beacons and 1.82% for IPv6 beacons. The corresponding numbers after filtering the results the BGP attribute are formed as follows: 50% of pairs are less than 0.26% likely to cause a zombie route with an average value of 0.54% for IPv4 beacons and 1.58% for IPv6 beacons.

Furthermore, in Fig. 6 we confirm the authors' conclusions that the AS paths that remain in the RIBs of zombie peer ASes are of longer length. This means that the BGP route selection initially did not elect these routes as the optimal ones, but they emerged in the routing tables during the path hunting process after the withdrawal of the beacons. However, our numbers are significantly higher than the ones in previous work. We find that 96.1% of the paths of the IPv4 zombie routes (95.54% if we filter the results with the BGP attribute) are different from the ones before the associated beacon's withdrawal. We observe similarly high numbers in the case of IPv6 zombies as well, 90.03% and 79.61% respectively.

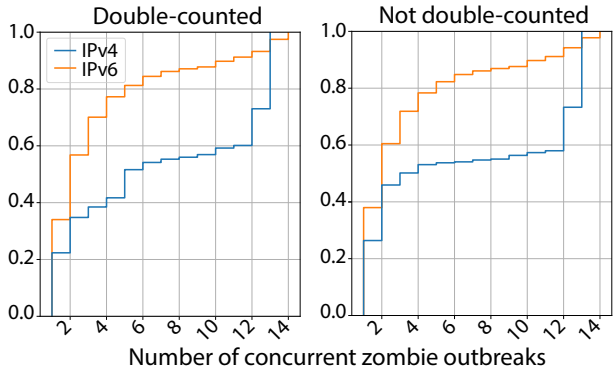


**Figure 7: CDF of the number of concurrent zombie outbreaks with double-counting and without double-counting.**

In Fig. 7 we replicate the analysis of Figure 4.b of [4] that shows the CDF of the number of concurrent zombie outbreaks at multiple RIPE RIS peers. As expected, we find that 22.35% of IPv4 and 34.04% of IPv6 zombie outbreaks occurred singly (after filtering the results with the BGP attribute, 26.38% and 37.97% respectively), whereas 26.96% of IPv4 zombie outbreaks emerged simultaneously for all beacon prefixes (after filtering the results with the BGP attribute, 26.71%).

## C Noisy peers

| With double-counting | | Without double-counting | |
|---|---|---|---|
| IPv4 | IPv6 | IPv4 | IPv6 |
| **mean** | **mean** | **mean** | **mean** |
| 0.044 | 0.4284 | 0.0018 | 0.426 |
| **median** | **median** | **median** | **median** |
| 0.0043 | 0.4705 | 0.0017 | 0.467 |

**Table 4: Mean and median value of the likelihood of the pair <RIPE RIS beacon, AS16347> to have a zombie route for *(i)* IPv4 and *(ii)* IPv6 beacon prefixes, and *(i)* without applying the double-counting filter and *(ii)* with applying it.**

| Peer Address (ASN) | zombie routes | Perc. | zombie routes | Perc. |
|---|---|---|---|---|
| | 1:30h | | 3h | |
| 176.119.234.201 (211509) | 163 | 9.91% | 149 | 9.06% |
| 2001:678:3f4:5::1 (211509) | 163 | 9.91% | 149 | 9.06% |
| 2a0c:9a40:1031::504 (211380) | 115 | 7% | 113 | 6.88% |

**Table 5: Absolute number of zombie routes and percentage of the beacon announcements that led to zombie routes in the noisy RIPE RIS peer ASes 211509 and 21380 after *(i)* 1.5 hours and *(ii)* 3 hours after the beacons' withdrawal.**