

Hunting BGP Zombies in the Wild

Porapat Ongkanchana
The University of Tokyo
pora@hongo.wide.ad.jp

Romain Fontugne
IJJ Research Lab
romain@ijj.ad.jp

Hiroshi Esaki
The University of Tokyo
hiroshi@wide.ad.jp

Job Snijders
Fastly
job@fastly.com

Emile Aben
RIPE NCC
emile.aben@ripe.net

ABSTRACT

As the key component of Internet's inter-domain routing, BGP is expected to work flawlessly. However, a recent study has revealed the presence of BGP zombies: Withdrawn prefixes that are still active in routing tables and that can cause routing issues. That study used experimental prefixes with scheduled withdrawals (BGP beacons). In this study we aim at detecting BGP zombies for any prefixes announced on the Internet. To that end we study characteristics of withdrawn messages, and devise a method to differentiate withdraw messages corresponding to local topological changes to those standing for prefixes withdrawn by their origin AS. Based on this classification we study the occurrence of zombies in the wild in six years of BGP data. We find over 6.5 millions zombies, among those we confirm that 94% report incoherent states and caused 468 potential routing loops. Our study also reveals that noisy prefixes, long AS paths, and ASes announcing a large number of prefixes are more prone to zombies.

1 INTRODUCTION

The Border Gateway Protocol (BGP) is the inter-domain routing protocol of the Internet. Routers all across the world exchange reachability information using this protocol, and it is of the utmost importance that these operations are timely and correctly executed. However, BGP has no mechanisms to ensure the integrity of exchanged information which makes BGP vulnerable to different types of attack and mishaps [2, 10, 16, 18], and also discrepancies, such as BGP zombies [5].

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 ACM 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.

ANRW '21, July 24–30, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8618-0/21/07...\$15.00

<https://doi.org/10.1145/3472305.3472315>

Also known as stuck route or ghost route, BGP zombies can emerge when an AS stops announcing an IP prefix. Ideally, the AS sends withdrawal messages to all its peers and the messages propagate throughout the Internet triggering the removal of corresponding entries in all routers' Routing Information Base (RIB). However, a recent study [5] has shown that this basic BGP operation sometimes fails and causes BGP zombies, that is active RIB entries for withdrawn prefixes. This past study focused solely on a few experimental prefixes that are withdrawn at scheduled time slots (BGP beacons [11, 15]) and omits all other prefixes that are in use on the Internet. As network operators have subsequently reported issues related to BGP zombies [5], including outages [13] and the difficulty to pinpoint zombies root causes [9, 17], an analysis of BGP zombies in the wild is needed to better understand the extent of BGP zombies on the Internet.

In this paper we aim to quantify the impact of BGP zombies on the numerous prefixes in use on the Internet. To detect the presence of BGP zombies in the wild, we study characteristics of BGP withdrawal messages and devise a technique to differentiate messages caused by local topological changes to those representing prefixes withdrawn from the origin AS. Then, using this technique we quantify the extent of BGP zombies for any prefix in use from 2014. We believe this study can assist network operators to troubleshoot routing discrepancies and improve the integrity of the Internet. Our main contributions are:

- We analyze 6 years of historical BGP data to characterize and classify different types of withdrawal messages.
- We devise a simple technique to detect zombies for any prefix (as opposed to beacon prefixes [5]) which can be easily implemented by network operators.
- Using this detection technique, we identify BGP zombies across 6 years of BGP data and document their characteristics as well as their impact on popular Internet services.
- We report a total of 6.5M BGP zombies (i.e. pair <BGP peer, prefix>), 88% of which are IPv4 prefixes.
- We validate the detection method by revealing incoherent states between route collector peers and comparing the result with past research.

- We show that zombies are pervasive and also observed for popular content providers.
- We uncover how AS routing characteristics, such as number of announced prefixes, average path length, and number of update messages can contribute to the emergence of zombies.
- We report 468 potential routing loops and 77k detours caused by BGP zombies.

2 BGP ZOMBIES

2.1 Terminology

In this paper we use the same terminology as the one established in [5]. A BGP zombie refers to an active RIB entry for a prefix withdrawn from its origin AS. Zombie peers and zombie ASes are used to describe BGP peers and ASes whose RIBs contains BGP zombies. We refer to a group of BGP zombies for the same prefix and occurring at approximately the same time as a zombie outbreak. Consequently, a zombie outbreak may contain multiple zombie peers and zombie ASes but corresponds only to one prefix. We measure the zombie outbreak size by counting the number of zombie peers involved.

2.2 Related Work

Fontugne et al. [5] have conducted the first thorough investigation on BGP zombies. They have confirmed the existence of BGP zombies, while providing in depth analysis on various characteristics. They shown that zombies are not uncommon and even large transit ASes can be affected by BGP zombies. But these results are only for RIPE’s Routing Information Service (RIS) BGP beacon prefixes. The main purpose of our study is to extend previous study and investigate the characteristics of BGP zombies in the wild. There is fundamental differences between the two studies: (1) The study with BGP beacons looked at a small set of prefixes that are periodically withdrawn, our study is more general by looking at regular prefixes announced on the Internet, but for which the exact withdrawal times are unknown. This implies that we need to devise a technique to detect when a prefix is withdrawn by its origin AS. Also our analysis covers a much larger number of Internet prefixes. (2) The number of regular prefixes is several orders of magnitude higher than beacon prefixes hence only scalable methods can be considered. (3) Beacon prefixes accommodate no host and no traffic, thus the impact of zombies for beacons is limited. Network operators have however reported that zombies for regular prefixes may trigger customer complaints [5] and outages [13].

3 DATA SET

For this study, we analyze historical BGP data collected by RIPE RIS between January 2014 and December 2019. We select 10 days each month, from the 10th to the 20th, and all route collectors that operated for at least half of the measurement period (i.e. rrc00, rrc01, rrc03-07, rrc10-16, rrc18-21). In order to ease data manipulation, we use only RIS data in this study, this dataset includes over 300 full-feed peers for IPv4 and IPv6. We have examined a total of 720 days of BGP data. Note that, we do not aggregate this data in any way, all 72 groups of 10 days data are analyzed separately. From this data set we calculate the number of active peers per prefix, the key metric for this study.

Active peers $A_p(t)$ refers to the set of routers which are announcing the prefix p at time t . We add a router to set $A_p(t)$ if and only if that router’s most recent BGP update message for prefix p in prior to time t was announcing a route (i.e. not a withdraw message). The total number of active peers varies from one prefix to another because BGP peers are exposed to a different sets of prefixes. To ease our analysis we compute the fraction of active peers $n_p(t)$, such as $n_p(t) = \frac{|A_p(t)|}{\max_i(|A_p(i)|)}$ with $\max_i(|A_p(i)|)$ the maximum number of active peers for prefix p observed across a 10-day batch. This metric ranges between 0 and 1 which respectively represents a prefix that is withdrawn by all peers and a prefix seen by a maximum number of peers. We compute $n_p(t)$ for each prefix seen in the data set described above with a temporal granularity of 15 minutes.

For this study we filter out prefixes with a very low visibility, that is prefixes consistently announced by a small number of peers. The definition of BGP zombies in this case is ill-defined because the propagation of these prefixes is intentionally limited. Thus we filter out prefixes where $\max_i(|A_p(i)|) < 100$ (20% of all monitored prefixes). This value is empirically set in function of the number of peers in our data set (over 300 full-feed peers for IPv4 and IPv6).

4 HUNTING BGP ZOMBIES

BGP zombies emerge when a prefix is withdrawn and a router fails to reflect this change in its routing table. These are the two fundamental information we need in order to detect BGP zombies. Routing table changes for RIS peers are directly available in our dataset but inferring prefix withdrawn in the wild is challenging for the four following reasons: (1) **Withdrawal can happen at any time.** Unlike beacon prefixes, regular prefixes are independently managed by their origin AS and can be withdrawn at any point in time. (2) **Withdrawal propagation time is varying and unpredictable.** Past research [6, 11] has shown that the propagation time of withdrawal messages is significantly fluctuating. These variations are mainly due to path hunting

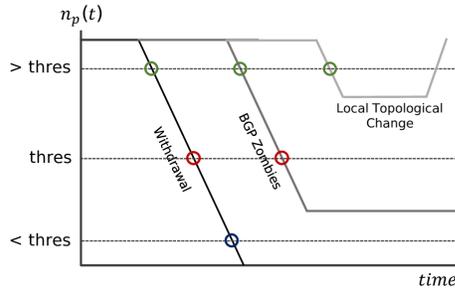


Figure 1: The overview of how different threshold values would affect the withdrawal detection.

and noise reduction techniques (e.g. MRAI and Route Flap Damping [4, 6, 8, 12, 19]) and are hardly predictable. (3) **Local topological changes.** Withdraw messages are observed in the case of local changes although the origin AS has not withdrawn the prefix. This is, for example, due to reconfigurations of networks between RIS peers and the monitored prefix. (4) $n_p(t)$ is not null when zombies emerge. By definition zombie peers have an active entry for a withdrawn prefix, thus the number of active peers $n_p(t)$ is not always dropping to 0 when the prefix is withdrawn by its origin AS.

4.1 Withdrawal scope

Comprehending when a prefix is globally withdrawn and how long it took for this information to propagate is key for our detection algorithm. The principles of the algorithm are simple, we make the assumption that a prefix p is withdrawn if its number of active peers, $n_p(t)$, has dropped below a certain value and stay low for an extended period of time. As shown in Figure 1, $n_p(t)$ drops to 0 when the withdrawal has propagated to all peers and there is no zombie. In the case of zombies, we expect a similar $n_p(t)$ drop but it stabilizes at a low, non-null value, for an extended period of time. For local topological changes, we expect $n_p(t)$ to drop then rise again and stabilize to a value close to 1. To detect zombies we thus face the trade-off of setting a threshold value $thres$ that is low enough to avoid most local topological changes and high enough to detect all zombies.

In order to select a suitable threshold value, we investigate the typical $n_p(t)$ drops that happen between two stable states (excluding complete withdraws where $n_p(t)$ reaches 0). A stable state is defined as a constant $n_p(t)$ value for more than one hour and compute the maximum $n_p(t)$ drop as follows:

- (1) At $t = t_a$, the number of active peers $n_p(t_a)$ was constant for more than one hour (stable) and over 0.9.
- (2) At t , $t_a < t < t_b$, the number of active peers $n_p(t)$ has changed and does not stay at the same value for more than one hour (unstable).
- (3) At $t = t_b$, $n_p(t_b)$ becomes again stable.

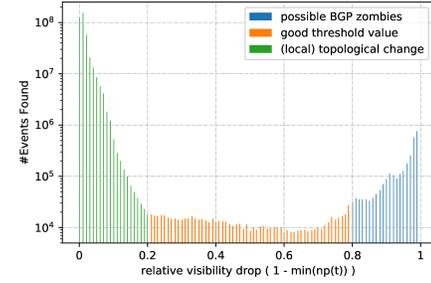


Figure 2: Distribution of maximum $n_p(t)$ drop (i.e. $1 - \min(n_p(t))$) when withdrawal messages are observed.

- (4) The maximum $n_p(t)$ drop is equal to $1 - \min(n_p(t))$ where $t_a < t < t_b$ and $\min(n_p(t)) \neq 0$

Figure 2 depicts the distribution of the maximum $n_p(t)$ drops observed in our data set. We observe two typical range of values, (0,0.2) and (0.8,1). The smaller drops between 0 and 0.2 represent mostly local changes, while the larger ones between 0.8 and 1 represent significant topological changes and potential BGP zombies. Based on these results, one should select a threshold value between 0.2 and 0.8. For this study, we arbitrarily set this threshold to 0.5, meaning that we ignore events that affect less than 50% of all observed peers for each prefix. As discussed in the following, we further filter selected events based on their temporal characteristics.

4.2 Withdrawal propagation time

To discriminate BGP zombies from large topological changes we investigate the temporal dynamics of events. For BGP zombies we expect a $n_p(t)$ drop that stabilize at a low value after BGP convergence, then depending on network operators actions $n_p(t)$ could either drop to 0 or raise again to 1. On the other hand, significant topological changes are characterized by a $n_p(t)$ drop quickly followed by a $n_p(t)$ rise, both happening during BGP convergence. Understanding prefix withdrawal propagation time is key to differentiate both events. Past research has shown that withdrawals usually lasts a few minutes [11] and in the case of Route Flap Damping up to one hour [8]. From our data set we estimate typical withdrawing time, T_w , by looking at the time duration of $n_p(t)$ drops, that is $t_b - t_a$, necessary for prefix p to be completely withdrawn ($n_p(t_b) = 0$ and $n_p(t_a) > thres$). Figure 3 shows that for both IPv4 and IPv6 more than half of the prefixes are withdrawn within 15 minutes (our smallest time resolution). We also observe that a small fraction of withdrawals take hours to complete which is probably due to the presence of BGP zombies. For this study we conservatively select withdrawals that last more than $T_w = 90min$ and search for zombies only in these events.

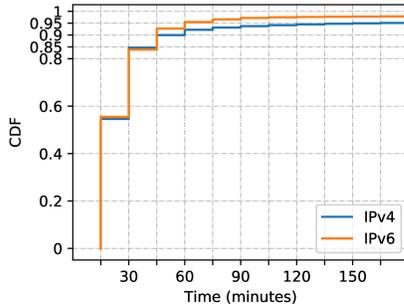


Figure 3: Distribution of the time necessary for prefixes to be globally withdrawn.

4.3 Zombie detection

The above observations constitute the core of our BGP zombie detection algorithm. BGP zombies are reported at time t for prefixes where the fraction of active peers $n_p(t - 90)$ drops below 0.5, but is not reaching $n_p(t) = 0$ within the next 90 minutes. BGP zombies are not reported if $n_p(t)$ is quickly going down to 0. In this case we infer that the prefix was successfully withdrawn by all RIS peers. And if $n_p(t)$ is going back above 0.5 then we classify this as a topological change. This simple method is easy to implement by network operators and, as shown in the following, provides efficient detection.

5 ZOMBIES IN THE WILD

Using the above zombie detection algorithm and our 6-year data set we found a total of 6.5M BGP zombies (88% are for IPv4 prefixes) from 486k outbreaks. In the following, we evaluate the proposed detection method by confirming the incoherent prefix state between RIS peers during zombie outbreak (§5.1) and past reports for BGP beacons (§5.2). We then compare zombie outbreaks for prefixes managed by popular content providers (§5.3) and reveal relationships between routing characteristics and zombie outbreaks (§5.4). Finally, we investigate the detrimental effects of BGP zombies on the routing infrastructure (§5.5).

5.1 State coherence between RIS peers

To validate that detected zombies are indeed erroneous RIB entries we now investigate RIS RIB entries, or lack thereof, during zombie outbreaks. The AS path found in a RIB entry represents the set of ASes that should be traversed to reach a certain prefix, hence this prefix is expected to be known by all ASes along the path. For BGP zombies, however, we expect that a zombie peer advertises an AS path that includes ASes that have no route to the corresponding prefix. These incoherent prefix states along the AS path corroborate the presence of BGP zombies.

In order to check for state coherence along zombie paths we need routing information from all ASes along advertised paths. In practice this comprehensive analysis is not possible with our data set, we can only look at state coherence across all RIS peers. We found that on average zombie outbreaks have 31.3% of AS path with no RIS peers along the path. For the remaining 68.7% of detected zombies we observe that 94.7% report incoherent states, we could thus verify that these are indeed zombies. The 5.3% of paths with coherent states are not conclusive because it could be due to several zombie RIS peers along the paths. For 99% of these paths the RIS peers are only one (80%) or two (19%) hops away, suggesting that they are likely to be part of the same zombie outbreak. Routing information from ASes closer to the origin AS is required to ensure that these are indeed paths with coherent states.

Looking at zombies for prefixes originated by RIS peers (1.1% of all outbreaks) allow us to estimate the fraction of misclassified zombies. We found that 97.6% of these zombies are indeed withdrawn by their origin AS. The few cases (2.4%) where the origin AS has not withdrawn the prefix but we detected BGP zombies illustrate that our detection method rarely classifies large topological changes as zombies.

In summary, given that state incoherence is not observed when prefixes are successfully withdrawn, that 94.7% of detected zombie paths have provably incoherent states, and that only a few detected paths are misclassified, we believe the proposed algorithm is effective for zombie detection.

5.2 Beacons and noisy prefixes

We further validate our results by comparing them with previous reports for RIS beacons and noticed interesting singularities for these prefixes. The 27 RIS beacon prefixes monitored in past research [5] accounts for 3.22% of all outbreaks detected in our dataset. This significant number of outbreaks for such small number of prefixes suggests that noisier prefixes, like beacons, are more prone to BGP zombies. We thus investigate the relationship between the number of zombie outbreaks and the number of BGP updates per prefix. To ease computation we focus only on prefixes that have at least 10 outbreaks per 10-day measurement period in 2018 and 2019. Figure 4 shows that the number of outbreaks increases with the number of BGP update messages for these prefixes. For IPv4, the Spearman correlation between these two quantities is $\rho = 0.6$ which confirms a non-negligible relationship between the number of update messages and the number of zombie outbreaks. In addition we found that IPv4 beacons are quite outstanding in our results as they produce a lot more zombies than regular IPv4 prefixes (Fig. 4). For IPv6, we also observe a lot of outbreaks for beacons but some regular prefixes have even more updates and zombie outbreaks.

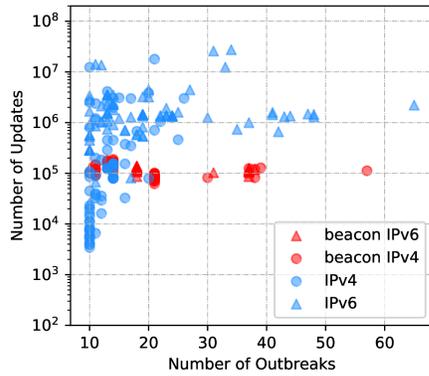


Figure 4: Number of zombie outbreaks and number of update messages for prefixes with more than 10 outbreaks per 10-day measurement period (2018-2019).

We confirmed that these regular prefixes are occurring irregularly in our data set, whereas beacons are seen in all measurement periods. Given the small number of beacon prefixes and their frequent appearance in the most impacted prefixes, we argue that the frequent zombie outbreaks found for beacon prefixes is not representative of what we can expect for regular prefixes. This is an important point to keep in mind when interpreting results from past study [5].

5.3 Zombies for popular content networks

To illustrate the prevalence of BGP zombies in regular prefixes we now focus on popular content networks. We investigate the frequency of BGP zombies for 42 ASes that commonly appear in the top 25 of Alexa, Umbrella, and Majestic lists [14]. Figure 5 shows the results for the top 15 ASes whose prefixes are consistently reported in our results for 2018-2019. These ASes are sorted by their median number of monthly outbreaks per prefix, this is hereafter referred as the zombie ranking. For both IPv4 and IPv6, we found that Akamai (AS16625 and AS20940) prefixes are generating the highest number of zombies. For AS20940, we observe IPv4 zombies for 22 (out of 24) measurement periods. This is an order of magnitude higher than what we record for some other large content providers, such as Google (AS15169, not even in the IPv4 top 15) which has only zombies in three measurement periods for IPv4 and one measurement period for IPv6. To understand these discrepancies we select relevant routing characteristics for these ASes and cross-reference them with the emergence of zombies.

5.4 Routing characteristics and zombies

Due to the erratic emergence of zombies in routers, we expect the number of zombies to be proportional to the number of prefixes announced by an AS. That is an AS announcing

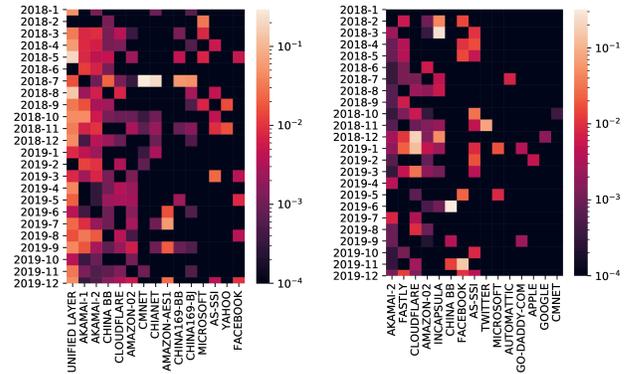


Figure 5: BGP zombies for popular ASes from Jan. 2018 to Dec. 2019 and for IPv4 (left) and IPv6 (right). Colors show the average number of outbreaks per announced prefix.

Table 1: Ranking of popular content networks according to prevalence of zombie outbreaks

AS	zombie rank	prefix rank	path rank
46606 Unified Layer	1	13	3
16625 Akamai	2	3	1
20940 Akamai	3	2	7
4134 China BB	4	7	15
13335 Cloudflare	5	6	12

more prefixes is more likely to have one of these prefixes turn into zombies. Similarly, we suppose that the probability of zombie emergence increases with the length of announced AS paths as these paths are likely involving more routers, i.e. components that can contain bugs [5, 13]. Based on these intuitions, we investigate the relation between the number of announced prefixes and AS path length to the occurrence of zombies.

From the 15 ASes of Figure 5 left plot, we compute two other rankings based on the number of announced prefixes per AS and the average path length from RIS peers to these ASes. Table 1 shows these ranking values for the top 5 zombie rank ASes and reveals that these ASes either announce a large number of prefixes or have the longest AS paths to RIS peers. For example, Akamai (AS16625), has the longest average AS path length and announces a considerably high number of prefixes (ranked third in terms of the number of prefixes). On the other hand, Amazon (AS16509) ranked sixth, despite announcing the most prefixes. The paths to these prefixes are usually short (ranked 12 for path length).

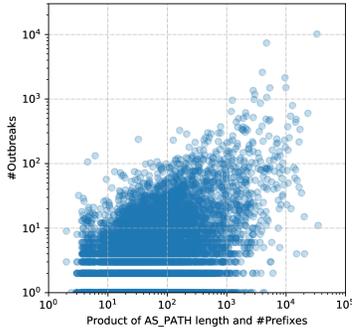


Figure 6: Number of outbreaks vs. number of prefixes and average path length from Jan. 2018 to Dec. 2019.

To better understand the contribution of both attributes to the emergence of zombies, we compute the Spearman correlation between these three quantities. Figure 6 shows the relation between the number of outbreaks and the product of average path length and number of prefixes per AS for all zombie outbreaks in 2018 and 2019. The correlation coefficient between these two metrics is $\rho = 0.40$, which is higher than the correlation of the number of outbreaks with only average path length $\rho = 0.03$ or only the number of prefixes $\rho = 0.39$. This indicates that the emergence of zombies for an AS is mainly related to the number of announced prefixes while path length is secondary.

5.5 Impact of BGP Zombies

BGP zombies misdirect affected routers to peers that are sometimes undesirable. This may create detours that make routes longer than expected and that may resemble to hijacking [13]. In our data set we found 77k zombies where the second hop in the AS path is different from the one found in the legitimate AS path of the covering prefix. Meaning that RIS peers in these cases are likely directing traffic to backup paths hence not to the preferred paths. For 51k zombies we found that the origin AS is different than the one found for the covering prefix. Hence, zombie AS may even misinterpret the origin of certain IP blocks. Our manual inspection of these results reveal that many of these zombies are prefixes delegated to customer ASes that have been withdrawn.

In certain cases detours caused by zombies create routing loops. This happens, for example, when AS_A has a zombie for prefix (e.g. 10.0.0.0/24) and its peer AS_B has no zombie but a valid route for the covering prefix (e.g. 10.0.0.0/16). Then, if the zombie path contains the pair $\langle AS_A AS_B \rangle$ and the valid path from AS_B contains the pair $\langle AS_B AS_A \rangle$, a routing loop will occur and traffic to the zombie prefix will not reach the destinations. To quantify the emergence of

routing loops caused by detected BGP zombies, we retrieve for all RIS peers the AS paths corresponding to prefixes covering detected zombies and search for routing loops. As in Section 5.1, this analysis is limited by the number of RIS peers and their location. We found 468 potential routing loops where zombie paths contain a pair $\langle AS_A AS_B \rangle$ and other RIS peers report a pair $\langle AS_B AS_A \rangle$ in paths of covering prefixes. But we have not enough BGP vantage points to confirm that AS_A is indeed infected and AS_B is not. Inferring this information would require the use of machine learning or different type of measurement in near-real time (e.g. traceroute). We leave this task for future work, for examples of zombie-caused routing loops observed by network operators we recommend [13].

6 DISCUSSION

The results presented in this paper have several implications for the networking community. In regard to the increasing size of routing tables and the corresponding concerns about routers resources limitations, our results show that BGP zombies contribute to routing tables inflation, but less than estimated earlier, based on BGP beacons [5].

In addition, as the number of zombies is increasing with the number of announced prefixes, an extensive use of prefix deaggregation [3] may be detrimental in terms of BGP zombies. Similarly, as we have shown that BGP churn is another important factor, the use of BGP optimizer that generates a lot of update messages may also be detrimental.

Finally peering policies and IP space management have a certain impact on BGP zombies. We discovered that shorter paths are less susceptible to BGP zombies, this is evident for large-scale networks that provide complete connectivity at each peering (Google [7]) and anycasted networks (Fastly).

7 CONCLUSION

In this paper, we extended prior study by investigating BGP zombies in the wild. We have analyzed 6 years of BGP data to reveal common prefix withdrawal patterns and then implemented a BGP zombie detection algorithm based on our observations. Using the detection algorithm we found that BGP zombies are not uncommon, over 6.5 million zombies has been observed in our data set. We confirmed that detected BGP zombies report incoherent states and discovered that BGP beacons are particularly prone to zombies. We found that BGP zombies are also present for popular web services and especially ASes that announces a lot of prefixes and that are reached through long AS paths. Finally, we discussed the impact of zombies on the routing infrastructure. For reproducibility purposes our source code is publicly available [1].

REFERENCES

- [1] zombie-hunter: Tool for analyzing BGP data and find BGP zombies. <https://github.com/pora49494/zombie-hunter>.
- [2] S. Cho, R. Fontugne, K. Cho, A. Dainotti, and P. Gill. Bgp hijacking classification. 2019.
- [3] L. Cittadini, W. Mühlbauer, S. Uhlig, R. Bush, P. Francois, and O. Maennel. Evolution of internet address space deaggregation: myths and reality. *IEEE Journal on Selected Areas in Communications*, 28(8):1238–1249, 2010.
- [4] A. Fabrikant, U. Syed, and J. Rexford. There’s something about mrai: Timing diversity can exponentially worsen bgp convergence. In *2011 Proceedings IEEE INFOCOM*, pages 2975–2983. IEEE, 2011.
- [5] R. Fontugne, E. Bautista, C. Petrie, Y. Nomura, P. Abry, P. Goncalves, K. Fukuda, and E. Aben. BGP Zombies: an analysis of beacons stuck routes. In *Passive and Active Measurement (PAM’20)*, pages 197–209, 2019.
- [6] A. García-Martínez and M. Bagnulo. Measuring bgp route propagation times. *IEEE Communications Letters*, 23(12):2432–2436, 2019.
- [7] Google. Peering. <https://peering.google.com/#/options/peering>, (Accessed on June 2020).
- [8] C. Gray, C. Mosig, R. Bush, C. Pelsser, M. Roughan, T. C. Schmidt, and M. Wählisch. BGP Beacons, Network Tomography, and Bayesian Computation to Locate Route Flap Damping. In *Proc. of ACM Internet Measurement Conference (IMC)*, New York, 2020. ACM. Accepted for publication.
- [9] I.-D. R. (idr) mailing list. TCP & BGP: Some don’t send terminate BGP when holdtimer expired, because TCP recv window is 0. https://mailarchive.ietf.org/arch/msg/idr/L9nWFBpW0Tci0c9DGfMoqC1j_sA/, 2020.
- [10] T. Kitabatake, R. Fontugne, and H. Esaki. BLT: A Taxonomy and Classification Tool for Mining BGP Update Messages. In *2018 IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*, April 2018.
- [11] Z. M. Mao, R. Bush, T. G. Griffin, and M. Roughan. Bgp beacons. In *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement, IMC ’03*, page 1–14, New York, NY, USA, 2003. Association for Computing Machinery.
- [12] Z. M. Mao, R. Govindan, G. Varghese, and R. H. Katz. Route flap damping exacerbates internet routing convergence. In *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 221–233, 2002.
- [13] P. Malachowski. Zombie routes, PLNOG Q3. <https://www.slideshare.net/atendesoftware/bgp-zombie-routes>, 2020.
- [14] J. Naab, P. Sattler, J. Jelten, O. Gasser, and G. Carle. Prefix top lists: Gaining insights with prefixes from domain-based top lists on dns deployment. In *Proceedings of the Internet Measurement Conference, IMC ’19*, page 351–357, New York, NY, USA, 2019. Association for Computing Machinery.
- [15] RIPE NCC. Current RIS Routing Beacons. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/current-ris-routing-beacons>, (Accessed on June 2020).
- [16] P. Sermpezis, V. Kotronis, P. Gigis, X. Dimitropoulos, D. Cicalese, A. King, and A. Dainotti. Artemis: Neutralizing bgp hijacking within a minute. *IEEE/ACM Transactions on Networking*, 26(6):2471–2486, 2018.
- [17] J. Snijders and B. Cartwright-Cox. Border Gateway Protocol 4 (BGP-4) Send Hold Timer. Internet-Draft draft-spaghetti-idr-bgp-sendholdtimer-00, Internet Engineering Task Force, Apr. 2021. Work in Progress.
- [18] P.-A. Vervier, O. Thonnard, and M. Dacier. Mind your blocks: On the stealthiness of malicious bgp hijacks. In *NDSS*, 2015.
- [19] C. Villamizar, R. Chandra, and R. Govindan. Bgp route flap damping. RFC 2439, RFC Editor, November 1998.