
Public Review for
**Hyper-Specific Prefixes: Gotta Enjoy the
Little Things in Interdomain Routing**

Khwaja Zubair Sediqi, Lars Prehn, Oliver Gasser

This paper investigates the presence of high-specific prefixes (HSP) on the BGP Internet routing during the last decade. These prefixes are more-specific than /24 (/48) for IPv4 (IPv6) and are commonly filtered by Autonomous Systems operators. In this paper, the authors provide a thorough analysis of the data, shedding light on the usage of these HSPs based on an analysis of snapshots from route collectors. First, the number of HSPs have increased over time, reaching 10% of all prefixes today. Second, HSPs can be both short- and long-lasting. Third, most HSPs are due to blackholing, leaked internal peering routes, or address block relocation. Interestingly, HSPs are more likely to be hosting SMTP/HTTP(s) services. Overall this paper offers a nice contribution to the understanding of the BGP universe, with a clear message and a nice quantification of the phenomenon. The authors clearly present and motivate the work, offering also to not experts a nice view of the routing complexity of the nowadays internet.

Public review written by
Marco Mellia
Politecnico di Torino

Hyper-Specific Prefixes: Gotta Enjoy the Little Things in Interdomain Routing

Khwaja Zubair Sediqi
MPI-INF
zsediqi@mpi-inf.mpg.de

Lars Prehn
MPI-INF
lprehn@mpi-inf.mpg.de

Oliver Gasser
MPI-INF
oliver.gasser@mpi-inf.mpg.de

ABSTRACT

Autonomous Systems (ASes) exchange reachability information between each other using BGP—the de-facto standard inter-AS routing protocol. While IPv4 (IPv6) routes more specific than /24 (/48) are commonly filtered (and hence not propagated), route collectors still observe many of them.

In this work, we take a closer look at those “hyper-specific” prefixes (HSPs). In particular, we analyze their prevalence, use cases, and whether operators use them intentionally or accidentally. While their total number increases over time, most HSPs can only be seen by route collector peers. Nonetheless, some HSPs can be seen constantly throughout an entire year and propagate widely. We find that most HSPs represent (internal) routes to peering infrastructure or are related to address block relocations or blackholing. While hundreds of operators intentionally add HSPs to well-known routing databases, we observe that many HSPs are possibly accidentally leaked routes.

CCS CONCEPTS

• **Networks** → **Network protocols**;

KEYWORDS

BGP, routing, hyper-specific prefixes.

1 INTRODUCTION

Autonomous Systems (ASes) use the Border Gateway Protocol (BGP) to announce prefixes to their peers [39]. Each BGP-speaking router of an AS can decide to accept or reject incoming announcements based on the prefix itself, the AS path, or other attributes that are attached to a route (e.g., BGP community values). Due to this concept, every single AS (and, in fact, also all its routers) may have a unique viewpoint into the Internet’s routing ecosystem [45].

Many popular BGP guidelines recommend the rigorous filtering of prefixes that encompass only a few addresses [11, 12, 29, 33, 34, 49, 50] and, hence, those prefixes have been shown to propagate neither far nor reliably [51]. While the possible reasons for announcing these types of prefixes are broad and range from traffic engineering over multi-homing configurations to prefix-hijack prevention [7, 17], the boundary for announcements which are deemed “widely acceptable” are

usually considered to be a /24 prefix in IPv4 and a /48 prefix in IPv6.

In this paper, we perform an in-depth analysis of prefixes that are more specific than those boundaries, i.e., /25 to /32 IPv4 prefixes and /49 to /128 IPv6 prefixes. We refer to those prefixes as **hyper-specific prefixes** (HSPs, see Appendix A for more details) and analyze their prominence in the global routing ecosystem, the functions that they serve, and whether they represent intentional or accidental announcements. More specifically, we make the following main contributions:

Observability. We perform a decade long analysis of HSPs as seen by 67 route collectors (see §2). We find that the number of HSPs has increased substantially since 2010 and peaked in 2018 at around 115K IPv4 and 18K IPv6 prefixes. While we observe that especially HSPs which are announced consistently for an entire year are visible by hundreds of collector peers, the average HSP can only be seen by a handful of them.

Use Cases & Functions. We analyze potential use cases of HSPs by combining insights from different analyses of CIDR sizes, BGP communities, and service hit rates across multiple years (see §3). We find that IPv4 HSPs mostly represent (internal) routes towards peering subnets and blackholing, whereas IPv6 HSPs are mainly used for address block relocations and, in substantially fewer cases, blackholing. We further find that HSPs are unlikely to contain many end hosts and that they are rarely used for traffic engineering.

Intended or Accidental Use. We compare the HSPs visible in BGP with those that were explicitly entered into routing databases—in particular, the Internet Routing Registries (IRR) and Resource PKI (RPKI)—to investigate intended or accidental use of HSPs (see §4). We find that while thousands of ASes explicitly specify their intent to use HSPs, many HSPs likely represent accidentally leaked routes.

The Future of HSPs. We discuss how the research and operator communities could make use of HSPs in the future. Finally, we publish a dashboard providing up-to-date HSP statistics to help AS operators in detecting leaked internal routes at <https://hyperspecifics.io>.




Figure 1: Growth of HSPs and HSP origin ASes as visible in all feeder ASes (left) and a consistent set of feeder ASes (right).

2 OBSERVABILITY

We begin our exploration of hyper-specific prefixes by analyzing their current and past presence in the Internet’s routing ecosystem.

In particular, we examine the routing information from hundreds of globally distributed ASes—called “feeder ASes” or “route collector peers”—collected by the Isolario [20], RIPE RIS [44], and Routeviews [46] projects. Starting from January 2010, we generate snapshots consisting of a week of RIB and update files every three months until October 2021. We provide further details about the choice of this window size in Appendix B. We employ various filtering steps to sanitize the data from, e.g., announcements of unallocated Internet resources, certain noisy origin ASes¹, or temporarily misconfigured feeder ASes. We also reached out to operators of noisy origin ASes. Two of these operators were not aware of this problem, but addressed it quickly upon our notification. A comprehensive list with justifications for the individual steps can be found in Appendix E.

First, we investigate the evolution of HSPs from January 2010 to October 2021. Figure 1 shows the number of hyper-specific prefixes (lines) and ASes that originate them (bars) over time. Looking at the left sub-plot, we observe that the number of seen HSPs (despite being noisy) consistently increases throughout the eleven years. We see more than 10k IPv6 and 100k IPv4 HSPs by the end of 2021, i.e., approximately one-tenth of all visible prefixes are hyper-specific (see Appendix C for further details). Relative to the increase in HSPs, we also observe an increase of ASes that originate them, with 584 and 2.5K ASes announcing hyper-specific prefixes via IPv6 and IPv4 by the end of 2021, respectively.

¹These ASes announced either (1) an extraordinary high number of HSPs (i.e., 100 or more times higher than in other snapshots) or (2) HSPs in an extraordinary high number of anchor prefixes for a limited time.

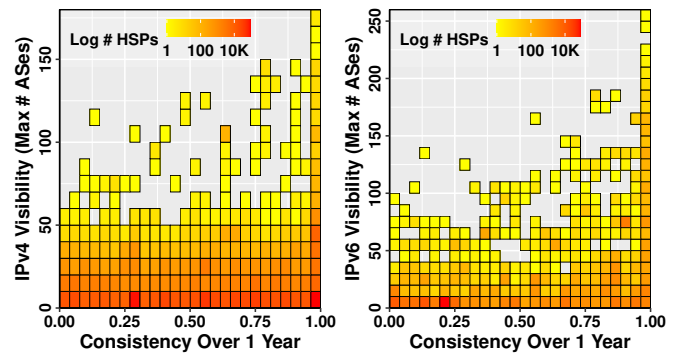


Figure 2: Heatmap showing HSP visibility and consistency for IPv4 (left) and IPv6 (right).

Given that the route collector projects acquired feeder ASes within our observation period, the increasing trend could simply be a sampling error. To test this hypothesis, we replicate the analysis using only data from the 105 IPv4 and 45 IPv6 feeder ASes that are consistently peering with route collectors throughout all snapshots. While our observations remain similar for IPv6, there are two changes for IPv4: (1) the number of hyper-specific prefixes that can be seen by a consistent set of ASes appears more stable (if any trend exists, it remains hidden behind the massive fluctuations); and (2) despite an initial increase, the number of ASes originating HSPs stagnates after 2016. Therefore, the number of IPv4 HSPs does not show a constant increase over time, but rather we observe more IPv4 HSPs due to an increase in feeder ASes at route collector projects.

This hypothesis check leads to another observation: When shrinking the set of feeder ASes, the number of HSPs and their respective origin ASes drops substantially (note the different y-axes for the left and right subplot of Figure 1). To improve our understanding of this insight, we analyze the visibility of HSPs, i.e., by how many peers each HSPs is seen. At the same time, we want to understand what causes the substantial fluctuations in the number of HSPs; hence, we also analyze their consistency, i.e., the fraction of time for which the prefix was seen by at least one feeder AS. Given that a one-week observation period would not provide much insight into consistency patterns, we conduct this analysis using data from the entirety of 2020. We first read the RIB snapshots from January 1, 2020 and then apply all updates for the whole year sequentially. By tracking the state of each routing table on a per-update basis, we can extract consistency in seconds granularity.

Figure 2 reports the visibility of an HSP on the y-axis against its consistency on the x-axis. For both heatmaps—IPv4 (left) and IPv6 (right)—each cell represents groups of

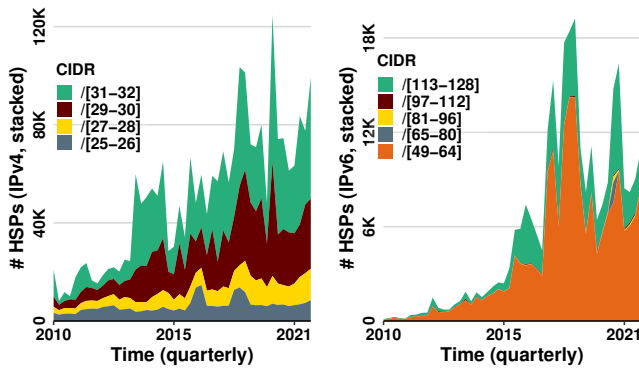


Figure 3: HSPs per CIDR size over time.

ten feeder ASes on the y-axis and two weeks of time on the x-axis. We first observe that there is no particular consistency trend: While some HSPs can only be observed for less than two weeks, others can be observed throughout the entire year. Our second observation is that the vast majority of hyper-specific prefixes can only be observed by a small number of collector peers, although we do also observe HSPs being visible during the entire year by hundreds of peers. This observation aligns with the restricted propagation characteristics of HSPs reported by previous blog posts [1, 2, 51] and observed by our own active experiments (an in-depth description of the experiments, their analysis, and subsequent results can be found in Appendix D). We hypothesize that the substantial fluctuations in the number of totally observed HSPs is a result of these two observations; the restricted propagation of HSPs might inflate the importance of the individual placement of feeder ASes and HSP origin ASes, and the tens of thousand of short-lived HSPs might cluster around certain real-world events, such as DDoS attacks or data center outages.

In summary, we observe that the presence of hyper-specific prefixes in the Internet’s routing ecosystem has increased through the last decade and HSPs make up about one-tenth of all the prefixes that are observed by route collectors. In IPv4 the increase in HSPs is driven by an increment in feeder ASes, whereas in IPv6 we see an increase also for a constant set of feeder ASes. While most HSPs only propagate locally, some of them are globally visible and can be consistently observed throughout an entire year.

3 USE CASES & FUNCTIONS

Given their past and current presence in the global routing system, we want to get a deeper understanding of the functions that hyper-specific prefixes potentially serve. As a first step in this direction, we use the fact that specific CIDR




Figure 4: Hit rate comparison of HSPs vs. IPv4-wide.

sizes often hint towards certain use cases. Consider the following example: If an AS wants to defend one of its servers against an ongoing DDoS attack, it may use blackholing announcements. Up to 98 % of these announcements are /32 (/128) IPv4 (IPv6) prefixes, i.e., they only cover the specific addresses of the attacked servers [9, 10, 15]. Larger CIDR sizes are rarely used for blackholing, as they would impair the services running on non-attacked servers as well, i.e., they would introduce unnecessary collateral damage [32]. Using similar lines of reasoning, we rely on the following associations between CIDR sizes and intended use cases: We associate (1) /25 and /26 IPv4 prefixes with traffic engineering (e.g., selective announcements [4, 37]), (2) /29 and /30 IPv4 prefixes with (Point-to-Point) peering subnets (i.e., the subnets needed to form inter-AS connections) [40], (3) /31 and /32 IPv4 prefixes with blackholing [9, 10, 15], (4) /49 to /64 IPv6 prefixes with address block reassignments [35], and (5) /113 to /128 IPv6 prefixes again with blackholing².

Figure 3 shows the number of IPv4 (left) and IPv6 (right) HSPs over time colored by their respective CIDR size groups. We first observe that the overall trends are stable over time. In IPv4, we observe that the most common CIDR size is /31–/32, i.e., the most prominent use case seems to be blackholing. Yet, we also observe that /29–/30 HSPs are comparably common; hence, many HSPs may actually represent peering subnets. Given that only about 10 % of HSPs have a CIDR size of /25 or /26, we believe that traffic engineering is a rare use case. For IPv6, we mainly observe the /49–/64 CIDR size range that we associate with address block relocations. In some ASes we also observe instances of /64s being used by hypergiants for off-nets [14]. We further observe a small fraction of /113–/128 CIDR sizes that we associate with blackholing. The share of blackholing HSPs is smaller in IPv6 compared to IPv4, which is in line with reports that blackholing in IPv6 makes up

²In private conversations a large European IXP confirmed that around 90 % of all blackholed IPv6 prefixes fall into the /113 to /128 prefix range.

less than 2 % compared to IPv4 [15, 32]. Those observations also explain some of the fluctuations that we observed in the previous section—blackholing events, and their subsequently announced prefixes, are often short-lived [32] and subsequently can cause substantial changes in the number of unique HSPs seen throughout a week.

As our CIDR-based analysis only provides us with hints on the actual usage, we now also analyze the services hosted in hyper-specific prefixes. For this analysis, we leverage archived scanning data from Rapid7’s Open Data platform [38] for 2019, 2020, and 2021. Rapid7 frequently scans the entire routed IPv4 address space³ for more than 100 well-known TCP and UDP ports. To compare regular with hyper-specific prefixes, we rely on the difference in protocol hit rate, i.e., we compare the fraction of responding hosts and total tested hosts⁴ on a per-protocol basis. We observe that four out of the top five protocols with the highest hit rate for regular and HSP prefixes overlap; BGP is only present in the HSP top five while CWMP is only present in the IPv4-wide top five. For those six protocols, Figure 4 shows a the relative difference of hit rates between regular and hyper-specific prefixes, where a positive value indicates an increase of hit rate in hyper-specific prefixes. While HTTP and HTTPS overall only see an increase of +100 %, we observe strong differences when drilling down on a per-CIDR level: When considering only /32 prefixes, HTTP’s hit rate increases by more than +500 % compared to its hit rate for IPv4-wide scans—which substantiates the association of the /32 CIDR size for blackholing. Even more pronounced than HTTP(S), SMTP and BGP see increases of up to +500%. When digging deeper we further observe that BGP is mainly prevalent in /30 and /29 prefixes, which underlines that these sizes might be dedicated to routing infrastructure. In contrast, we observe the only hit rate decrease (of more than 90%) for CWMP—a protocol used to remotely manage customer-premises equipment (CPE) devices such as home routers [52].

Finally, we investigate BGP communities attached to HSP announcements. BGP communities are used for many different reasons, such as information tagging, blackholing, or route redistribution. The most common BGP communities attached to hyper-specific prefixes are route steering or prepending instructions. In our analysis we look for BGP communities which are specifically used for blackholing




Figure 5: BGP communities distribution for HSPs.

(BH) [23] or restrict route propagation (RES)⁵. Figure 5 shows the use of BGP communities among HSPs from snapshots between 2019 and 2021. The bars indicate the median share of HSPs with the respective community, the whiskers denote the standard deviation over time. The “Any” keyword is used to specify groups of community targets, e.g., “Any RES” describes all prefixes that have any restriction community attached (i.e., it refers to the union of prefixes with “NO_ADV” community and prefixes with “NO_EXP” communities); similarly, the “Any Comm.” bar refers to the highest aggregation, i.e., prefixes for which we saw any community attached. As we can see, 60% of all IPv4 HSPs and almost three quarters of IPv6 HSPs come with some form of BGP communities. The vast majority of these communities is, however, not related to blackholing or restricting propagation. Only about 13% and 7% of prefixes can be associated with blackholing for IPv4 and IPv6, respectively. The by far most popular blackholing community is X:666. Moreover, we see no propagation restriction communities (“no advertise” or “no export”) in IPv6 and only about 0.5% in IPv4. Furthermore, we see that RES communities are a subset of BH communities, hinting that operators do not want their blackholing prefixes to propagate. Blackholing is therefore one contributor of HSPs, but blackholing communities are not present on the majority of HSP announcements. We note that the blackholing communities that we see at route collector peers is a lower bound: Blackholing communities—similar to other communities—could be cleaned along the path but the prefix itself could continue to propagate [24].

³Except for prefixes on their blacklist which were explicitly requested by network operators.

⁴Given that Rapid7 does not publish the state of their blacklist, we assume that all (at the time of the scan) routed IP addresses were tested. Additionally, we focus on analyzing what services are prominent in HSPs. We can not ensure that Rapid7 (or its upstream) does in fact receive the HSP announcements, as information about their probing vantage points and routing is not available.

⁵We also test for communities such as NOPEER or NO_EXPORT_SUBCONFED, but these are not prevalent among HSPs.




Figure 6: Visibility of origin ASes across data sets.

In summary, we observe that for IPv4 many and for IPv6 some HSPs are likely related to blackholing activities due to the used HSP prefix size. We find concrete evidence for 7–13% of HSPs explicitly tagged with blackholing communities. While we also observe many HSPs dedicated to routing infrastructure (e.g., peering subnets or address relocations), we observe that hyper-specific prefixes rarely contain any CPE devices.

4 INTENDED OR ACCIDENTAL USE?

Now that we have a basic understanding of the use cases of HSPs, we want to analyze whether HSPs are used intentionally or accidentally by ASes and their operators. If operators take the time and effort to explicitly enter hyper-specific prefixes into voluntarily-maintained databases, then it is likely that they plan to use them. Hence, we look at the Resource Public Key Infrastructure (RPKI) and Internet Routing Registry (IRR) operator databases.

We use private, three-monthly IRR snapshots [19] between January 1, 2017, and October 7, 2021, which contain information about routing policies. The RPKI database contains legally binding mappings between Internet resources and ASes. We use daily snapshots of the RPKI database [43] from April 1, 2015, until October 7, 2021, generated by Chung et al. [5] to verify the validity of HSP announcements by ASes.

While we extract HSPs directly from the route(6) objects contained in the IRR databases, the Route Origin Authorization (ROA) objects in the RPKI snapshots describe CIDR size ranges [18]. Hence, a ROA can explicitly describe an HSP when both the minimum and maximum prefix length are hyper-specific, or implicitly when only the maximum prefix length is hyper-specific. When extracting HSPs and their origins from the RPKI database, we rely solely on explicit definitions as these clearly represent the desire to use HSPs (as all covered prefixes are hyper-specific). As implicit definitions might describe the future—but not necessarily




Figure 7: ROV status for HSPs

current—use of HSPs (e.g., an AS might currently announce a /24 but has already entered a currently unused max-length of /25), we decide to ignore them. We compare the HSPs on those two databases against the HSPs visible via BGP route collectors.

Figure 6 shows the number of unique origin ASes for both IPv4 and IPv6 within each dataset over time. We classify those origin ASes available in more than one dataset into the “Multiple” category. Our first observation is that for both IPv4 and IPv6, the IRR dataset contains the largest fraction of HSP origin ASes. While this might imply that network operators tend to actually use HSPs, it is well-known that route objects can become stale given that the database is only maintained on a voluntary basis [48]. Yet, some entities, e.g. certain IXP Route Servers [8], require route objects in the IRR database to redistribute prefixes (i.e., HSPs). Even for the RPKI database we observe hundreds of explicitly defined HSPs⁶. Notably, for the last snapshot in October 2021, implicit HSPs would have increased the number of RPKI origin ASes from 294 to 990 for IPv4 and from 172 to 794 for IPv6, respectively. Beyond these intentional HSPs, we also observe that many of the HSPs from Route Collectors have no entries in operator databases, hence, they could potentially represent accidental announcements or misconfigured route collector sessions that leak internal routes.

While it is hard to link malicious intent to a more-specific announcement (since it could be, e.g., an address leasing agreement [36] or traffic engineering of sibling ASes [13]), we want to understand if the visible HSPs in the BGP are legitimate prefix advertisements by valid origin ASes or associated with possible prefix hijacks. Therefore, we perform route origin validation (ROV) of HSPs and its origin AS by checking them against the ROA records from the RPKI dataset. If a ROA covers the address space described by the

⁶Most of these HSPs are also in the BGP data set and hence end up in the multiple class.

prefix, then this prefix can violate the ROA in two ways: it can be too specific—which we mark as “Invalid (Length)”—and it can be announced by a different origin—which we mark as “Invalid (Origin).” If both of these conditions are met at the same time, we mark a prefix as “Invalid (Both).” If none of these conditions are met, we consider the prefix as “Valid.” Notably, we observe that 22 % of IPv4 and 19 % of IPv6 HSPs have a covering ROA entry (median percentages across snapshots in 2020 and 2021).

Figure 7 shows that legitimate ASes, i.e., the valid and invalid length categories together, advertise around 75 % of all HSPs. With an average of 25 % and peaking to around 50 % in 2016, 2017, and 2019, IPv6 has a higher percentage of valid HSPs than IPv4. The HSPs with invalid length form the largest group in IPv4, and mostly the second largest group in IPv6. The third largest group of HSPs has the “Invalid (Both)” ROV state, while the invalid origin category forms a minor fraction of HSPs’ ROV state. Legitimate ASes advertise around 75 % of HSPs, which indicates that HSPs are not majorly associated with BGP prefix hijacks. Beyond malicious ASes, the “Invalid (Origin)” and “Invalid (Both)” status could also be caused by not properly entered sibling ASes [13] or from a DDoS Protection Service (DPS) [21]. We analyze how many hyper-specific prefixes are originated from a DPS as identified by Jin et al. [21] and find that only around 1 % of HSPs in IPv4 and IPv6 are related to DPS companies.

In summary, we observe that for both IPv4 and IPv6, hundreds of ASes intentionally entered hyper-specific prefixes into operator databases. Yet we also see that many of the HSPs that are visible from route collectors have no respective entries and are likely related to the accidental announcement or disclosure of internal routes. This is further substantiated by the observation that most HSPs are actually ROV invalid since they are more specific than intended by their covering ROA entry.

5 DISCUSSION

Research Community. While many HSPs seem to be intentional, we also observe a large number that potentially represent leaked internal routes. While the task of reconfiguring a leaking router ultimately belongs to the feeder AS’ operators, we believe that the maintainers of route collector projects play a vital role when it comes to raising awareness for the existing problems. To support and guide this process, we publish and maintain a dashboard that provides up-to-date HSP statistics as well as a rankings of the top HSP contributors at <https://hyperspecifics.io>. Beyond fixing potential leakage errors, we believe that studying the potential correlations between hyper-specific prefixes and their less-specific counter parts may lead to new insights into the routing optimizations used by ASes.

Operator Community. Even though various guides [11, 12, 29, 33, 34] recommend strict filtering of HSPs, we find that many hyper-specific prefixes propagate to 100 or more collector peers. After discussing our results with thirteen operators from different types of networks, we believe that the limited filtering is often a result of popular customer requests. The operator of a major transit network told us that their network recently (throughout Summer 2020) changed from the filtering of all IPv4 HSPs to only filtering prefixes more specific than /28; this shift enabled (especially new and small) customer networks to perform basic traffic engineering despite a limited address allocation⁷.

This opens up the question whether operators should filter HSPs in the first place. We believe that for IPv6 the answer is a resounding “yes”. Given that there is no shortage of IPv6 addresses and obtaining new blocks is virtually free (compared to the high costs of obtaining IPv4 addresses), we do not see any reason to loosen the current filtering guidelines. For IPv4, we think that the answer should be more nuanced. While loosening the filtering guidelines allows even small ASes to perform traffic engineering, it would also further increase the routing table size. Hence, we believe that shifting the acceptable boundaries by a few CIDR sizes (e.g., /26 or /28) might be an agreeable compromise.

6 RELATED WORK

In this section, we report on related work in the areas of hyper-specific prefix analysis and prefix deaggregation.

HSP Analysis: Previous research in this area consists mostly of blog posts. In 2014, Aben and Petrie report on an experiment where they announced /24, /25, and /28 IPv4 prefixes and ran RIPE Atlas measurements to them [1]. Their findings show that HSPs are visible for at most 20 % of RIPE RIS peers [44] with route objects slightly improving the visibility. The RIPE Atlas experiments lead to similar results with fewer than 15 % of probes reaching their targets. One year later, Aben and Petrie revisit the propagation of hyper-specific prefixes and find a marginal increase of a few percent [2]. In 2017, Strowes and Petrie conclude that not much has changed regarding hyper-specific prefix propagation and at most one fourth of all BGP peers receive those announcements [51].

Prefix Deaggregation: In 2002, Bu et al. first characterize prefix deaggregations and the reasons for them, e.g., traffic engineering, multi-homing, and address fragmentation [3]. Meng et al. report in 2005 that even newly assigned address space is deaggregated and that the deaggregation rate of prefixes increases over time [31]. In 2010, Cittadini et al. [7] report that more than 10 % of ASes deaggregate their prefixes while around 1 % of ASes announce more than 10 prefixes

⁷This is a direct result of the current IPv4 Address exhaustion and the subsequently inflated prices [36].

for each address block they got assigned. Lutu et al. present a simulation model that estimates that origin ASes can reduce their transit cost by 5 % by using more-specific announcements [26–28]. Notably, the authors neither focused on IPv6 nor on hyper-specific prefixes. In 2016, Krenc and Feldmann analyze the address delegations realized via prefix deaggregations and report on delegations from customers to providers or between unrelated ASes (often involving CDNs) [25]. In 2017, Huston analyzes the prevalence and different types of more-specific prefix announcements in the Internet as an effect of prefix deaggregation [17]. His taxonomy attributes MSPs to three different root causes, hole punching (different origin AS), traffic engineering (same origin AS, but different AS path), and overlay (same AS path). He concludes that the former two play a useful role for network operators, while the usefulness of overlay more-specific prefixes could be argued about. Huston did not specifically investigate the effect of hyper-specific prefixes.

To the best of our knowledge, this paper presents the first scientific analysis of hyper-specific prefixes by providing an in-depth look into the prevalence and possible root causes for HSPs in the wild.

7 CONCLUSION

In this paper, we analyzed the presence of hyper-specific prefixes in the Internet’s ecosystem throughout the last decade. While we found an overall increase in the number of HSPs, most of them can only be observed by a few route collector peers. Yet, there are still plenty of HSPs that propagate to hundreds of route collector peers and can be consistently observed throughout an entire year. Inspired by those findings, we took a closer look at the function that these prefixes serve. For IPv4, we observed that HSPs are mainly associated with blackholing and infrastructure announcements (e.g., routes to peering subnets). While we only found limited evidence for any connection to traffic engineering, we observed that hyper-specific prefixes are less likely to contain end-user devices. For IPv6, we observe that almost all hyper-specific prefixes are related to address block reassignments, with only a small fraction representing blackholing. Even though we have seen that hundreds of networks use HSPs intentionally, we attributed even more cases to the accidental “leakage” of internal routes. Finally, we discussed the current state of HSPs from an academic as well as an operator point of view.

REFERENCES

- [1] Emile Aben and Colin Petrie. 2014. Propagation of Longer-than-24 IPv4 Prefixes. <https://labs.ripe.net/Members/emileaben/propagation-of-longer-than-24-ipv4-prefixes>
- [2] Emile Aben and Colin Petrie. 2015. Has the Routability of Longer-than-24 Prefixes Changed? <https://labs.ripe.net/Members/emileaben/has-the-routability-of-longer-than-24-prefixes-changed>
- [3] Tian Bu, Lixin Gao, and Don Towsley. 2002. On characterizing BGP routing table growth. In *Global Telecommunications Conference, 2002. GLOBECOM’02. IEEE*, Vol. 3. IEEE, 2185–2189.
- [4] Rocky KC Chang and Michael Lo. 2005. Inbound traffic engineering for multihomed ASs using AS path prepending. *IEEE network* 19, 2 (2005), 18–25.
- [5] Taejoong Chung, Emile Aben, Tim Bruijnzeels, Balakrishnan Chandrasekaran, David Choffnes, Dave Levin, Bruce M Maggs, Alan Mislove, Roland van Rijswijk-Deij, John Rula, et al. 2019. RPKI is coming of age: a longitudinal study of RPKI deployment and invalid route origins. In *Proceedings of the Internet Measurement Conference*. 406–419.
- [6] Cisco. 2021. Removing Private AS Numbers from the AS Path in BGP. https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xr-16/irg-xe-16-book/removing-private-as-numbers-from-the-as-path-in-bgp.html
- [7] Luca Cittadini, Wolfgang Mühlbauer, Steve Uhlig, Randy Bush, Pierre Francois, and Olaf Maennel. 2010. Evolution of Internet address space deaggregation: Myths and reality. *IEEE Journal on Selected Areas in Communications* 28, 8 (2010), 1238–1249.
- [8] DE-CIX. 2021. BGP announcement filtering. <https://www.de-cix.net/en/locations/frankfurt/route-server-guide>
- [9] Christoph Dietzel, Anja Feldmann, and Thomas King. 2016. Blackholing at ixps: On the effectiveness of ddos mitigation in the wild. In *International Conference on Passive and Active Network Measurement*. Springer, 319–332.
- [10] Christoph Dietzel, Matthias Wichtlhuber, Georgios Smaragdakis, and Anja Feldmann. 2018. Stellar: network attack mitigation using advanced blackholing. In *Proceedings of the 14th International Conference on emerging Networking EXperiments and Technologies*. 152–164.
- [11] Gert Döring. 2013. IPv6 BGP filter recommendations. <https://www.space.net/~gert/RIPE/ipv6-filters.html>
- [12] J. Durand, I. Pepelnjak, and G. Doering. 2015. BGP Operations and Security. RFC 7454 (Best Current Practice). <https://doi.org/10.17487/RFC7454>
- [13] Lixin Gao. 2001. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on networking* 9, 6 (2001), 733–745.
- [14] Petros Gigis, Matt Calder, Lefteris Manassakis, George Nomikos, Vasileios Kotronis, Xenofontas Dimitropoulos, Ethan Katz-Bassett, and Georgios Smaragdakis. 2021. Seven Years in the Life of Hypergiants’ Off-Nets. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*. 516–533.
- [15] Vasileios Giotsas, Georgios Smaragdakis, Christoph Dietzel, Philipp Richter, Anja Feldmann, and Arthur Berger. 2017. Inferring BGP blackholing activity in the internet. In *Proceedings of the 2017 Internet Measurement Conference*. 1–14.
- [16] 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*. 492–505.
- [17] Geoff Huston. 2017. BGP more specifics: routing vandalism or useful? <https://blog.apnic.net/2017/06/26/bgp-specifics-routing-vandalism-useful/>
- [18] G. Huston and G. Michaelson. 2012. Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs). RFC 6483 (Informational). <https://doi.org/10.17487/RFC6483>
- [19] irr.net. 2021. List of Routing Registries. <http://www.irr.net/docs/list.html>
- [20] Isolario 2021. Isolario Project. Available at <https://isolario.it/>.
- [21] Lin Jin, Shuai Hao, Haining Wang, and Chase Cotton. 2018. Your Remnant Tells Secret: Residual Resolution in DDoS Protection Services.

- In *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 362–373.
- [22] Juniper. 2021. Understanding BGP Confederations. <https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/bgp-confederations-for-scaling.html>
- [23] T. King, C. Dietzel, J. Snijders, G. Doering, and G. Hankins. 2016. *BLACKHOLE Community*. RFC 7999. RFC Editor.
- [24] Thomas Krenc, Robert Beverly, and Georgios Smaragdakis. 2021. AS-level BGP community usage classification. In *Proceedings of the 21st ACM Internet Measurement Conference*. 577–592.
- [25] Thomas Krenc and Anja Feldmann. 2016. BGP prefix delegations: a deep dive. In *Proceedings of the 2016 Internet Measurement Conference*. 469–475.
- [26] Andra Lutu, Marcelo Bagnulo, Cristel Pelsser, Kenjiro Cho, and Rade Stanojevic. 2015. An analysis of the economic impact of strategic deaggregation. *Computer Networks* 81 (2015), 147–163.
- [27] Andra Lutu, Marcelo Bagnulo, and Rade Stanojevic. 2012. An economic side-effect for prefix deaggregation. In *2012 Proceedings IEEE INFOCOM Workshops*. IEEE, 190–195.
- [28] Andra Lutu, Cristel Pelsser, Marcelo Bagnulo, and Kenjiro Cho. 2013. The aftermath of prefix deaggregation. In *Proceedings of the 2013 25th International Teletraffic Congress (ITC)*. IEEE, 1–8.
- [29] MANRS. 2021. Prefix filter configuration tools. <https://www.manrs.org/isps/guide/filtering/>
- [30] Alexander Marder, Matthew Luckie, Amogh Dhamdhere, Bradley Hufaker, KC Claffy, and Jonathan M Smith. 2018. Pushing the boundaries with bdrmapit: Mapping router ownership at Internet scale. In *Proceedings of the Internet Measurement Conference 2018*. 56–69.
- [31] Xiaoqiao Meng, Zhiguo Xu, Beichuan Zhang, Geoff Huston, Songwu Lu, and Lixia Zhang. 2005. IPv4 address allocation and the BGP routing table evolution. *ACM SIGCOMM Computer Communication Review* 35, 1 (2005), 71–80.
- [32] Marcin Nawrocki, Jeremias Blendin, Christoph Dietzel, Thomas C Schmidt, and Matthias Wählisch. 2019. Down the black hole: dismantling operational practices of BGP blackholing at IXPS. In *Proceedings of the Internet Measurement Conference*. 435–448.
- [33] NLNOG. 2021. Filtering Small Prefixes. https://bgpfilterguide.nlno.org/guides/small_prefixes/
- [34] NOCTION. 2021. BGP Prefix Filtering. <https://www.noction.com/knowledge-base/bgp-prefix-filtering>
- [35] Ramakrishna Padmanabhan, John P Rula, Philipp Richter, Stephen D Strowes, and Alberto Dainotti. 2020. DynamIPs: Analyzing address assignment practices in IPv4 and IPv6. In *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies*. 55–70.
- [36] Lars Prehn, Franziska Lichtblau, and Anja Feldmann. 2020. When wells run dry: the 2020 IPv4 address market. In *Proceedings of the 16th International Conference on emerging Networking EXperiments and Technologies*. 46–54.
- [37] Bruno Quoitin, Cristel Pelsser, Louis Swinnen, Olivier Bonaventure, and Steve Uhlig. 2003. Interdomain traffic engineering with BGP. *IEEE Communications magazine* 41, 5 (2003), 122–128.
- [38] Rapid7. 2021. Rapid7 Open Data. <https://opendata.rapid7.com/>
- [39] Y. Rekhter (Ed.), T. Li (Ed.), and S. Hares (Ed.). 2006. A Border Gateway Protocol 4 (BGP-4). RFC 4271 (Draft Standard). <https://doi.org/10.17487/RFC4271> Updated by RFCs 6286, 6608, 6793, 7606, 7607, 7705, 8212, 8654.
- [40] A. Retana, R. White, V. Fuller, and D. McPherson. 2000. Using 31-Bit Prefixes on IPv4 Point-to-Point Links. RFC 3021 (Proposed Standard). <https://doi.org/10.17487/RFC3021>
- [41] RIPE. 2021. RIPE Atlas Measurement Dumps - [2021/05/17-23]. <https://data-store.ripe.net/datasets/atlas-daily-dumps/>
- [42] RIPE NCC. 2021. RIPE Atlas measurement platform. <https://atlas.ripe.net/>
- [43] RIPE NCC. 2021. RIPE RPKI Snapshots. <https://ftp.ripe.net/rpki/>
- [44] RIPE-RIS 2021. RIPE Routing Information Service. <http://www.ripe.net/ris/>.
- [45] Matthew Roughan, Walter Willinger, Olaf Maennel, Debbie Perouli, and Randy Bush. 2011. 10 lessons from 10 years of measuring and modeling the internet’s autonomous systems. *IEEE Journal on Selected Areas in Communications* 29, 9 (2011), 1810–1821.
- [46] Routeviews 2021. Routeviews Project – University of Oregon. Available at <http://www.routeviews.org/>.
- [47] Brandon Schlinker, Todd Arnold, Italo Cunha, and Ethan Katz-Bassett. 2019. PEERING: Virtualizing BGP at the Edge for Research. In *Proceedings of the 15th International Conference on Emerging Networking Experiments And Technologies*. Orlando, FL.
- [48] Georgos Siganos and Michalis Faloutsos. 2004. Analyzing BGP policies: Methodology and tool. In *IEEE INFOCOM 2004*, Vol. 3. IEEE, 1640–1651.
- [49] P. Smith and R. Evans. 2011. RIPE-532 - RIPE Routing Working Group Recommendations on IPv6 Route Aggregation. <https://www.ripe.net/publications/docs/ripe-532>
- [50] P. Smith, R. Evans, and M. Hughes. 2006. RIPE-399 - RIPE Routing Working Group Recommendations on Route Aggregation. <https://www.ripe.net/publications/docs/ripe-399>
- [51] Stephen Strowes and Colin Petrie. 2017. BGP Even-More Specifics in 2017. https://labs.ripe.net/Members/stephen_strowes/bgp-even-more-specifics-in-2017
- [52] The Broadband Forum. 2018. TR-069: CPE WAN Management Protocol. https://www.broadband-forum.org/technical/download/TR-069_Amendment-6.pdf
- [53] CAIDA UCSD. 2021. The IPv4 Routed /24 Topology Dataset - [2021/05/17-23]. <https://data.caida.org/datasets/topology/ark/ipv4/probe-data/>
- [54] CAIDA UCSD. 2021. The IPv6 Routed /48 Topology Dataset - [2021/05/17-23]. <https://publicdata.caida.org/datasets/topology/ark/ipv6/probe-data/>
- [55] Maya Ziv, Liz Izhikevich, Kimberly Ruth, Katherine Izhikevich, and Zakir Durumeric. 2021. ASdb: a system for classifying owners of autonomous systems. In *Proceedings of the 21st ACM Internet Measurement Conference*. 703–719.

A MSPS VS. HSPS

In this section, we want to briefly contrast the definitions of More-Specific Prefixes (MSPs) and Hyper-Specific Prefixes (HSPs). P is an MSP of P' when the address space that P describes is entirely contained in P' , e.g., $1.0.0.0/24$ is an MSP of $1.0.0.0/22$. In contrast, we call a prefix hyper-specific if its CIDR size is larger than $/24$ or $/48$ for IPv4 and IPv6, respectively. While labelling a prefix as an MSP requires another (covering) prefix, the HSP label relies entirely on the CIDR size of a given prefix and does not require a second, related prefix. Notably, many—but not all—hyper-specific prefixes are also MSPs of less-specific prefixes. As the definitions of MSPs and HSPs are very different, further classifications of HSPs (as in, e.g., Geoff Huston’s blogpost [17]) are not directly applicable to HSPs.

B ROUTE COLLECTOR CONSISTENCY

In order to analyze representative route collector snapshots of the three RC projects Isolario [20], RIPE RIS [44], and Routeviews [46], we first analyze their consistency over time. To estimate the consistency, we initially retrieve data for all days in 2010, 2013, 2016, and 2020. For each day, we download the first routing information base (RIB) snapshot as well as all available update messages produced by each RC. If an update file is missing, we, additionally, download the first available RIB snapshot after the missing update file. After extracting the HSPs for each day, we analyze consistency as the fraction of HSPs seen at day $n + w + 1$ that are also visible within the observation period $[n, n + w]$. Notably, we try all possible window size positions, i.e., $n \in \{0, \dots, d - w - 1\}$ where d is the number of days in the given year.




Figure 8: Impact of window size on visible HSPs.

Figure 8 shows the mean as well as the interquartile range (IQR) across all possible n for window size w between 1 and 60 days for IPv4 and IPv6 HSPs in 2020. We observe that a seven-day window allows us to achieve a consistency of 97 %

and 98 % for IPv4 and IPv6, respectively. Notably, further expanding the window size to 60 days would only increase the consistency by $\sim 0.5\%$. Given that we now have a snapshot aggregation window, we still need to pick a snapshot interval. When comparing the number of visible HSPs for different snapshot intervals, we observe that a three-month interval provides an optimal balance: While the number of data points is still capable of capturing all visible trends in more-frequent snapshot intervals, the reduced amount of data (i.e., only seven days every three months) still allows us to perform computationally expensive observations for the entire decade promptly.

C FURTHER ANALYSIS

How prominent are HSPs? To understand the prevalence of hyper-specific prefixes, we aggregate the routing tables of all collector peers and compare the distribution of prefixes depending on CIDR sizes. Figure 9 shows those distributions as stacked bar plots for each snapshot. We observe that up to 13 % (in 2015) and 25 % (in 2018) of totally visible prefixes are hyper-specific for IPv4 and IPv6, respectively. Yet, the usual contribution of HSPs is approximately 10 % for most months. Note that this does not mean that any single routing table contains that many HSPs on its own.

How visible are HSP? To further elaborate on this point, Figure 10 shows the number of hyper-specific prefixes per IPv4 (left) and IPv6 (right) snapshot separated based on the number of route collector peers that can see them. For IPv6, we observe that most hyper-specific prefixes can be seen by two or more peers, with around a fifth of all HSPs being visible by 11+ peers for most snapshots. Similar to the previous plot, we again observe a peak of ($\sim 20K$) hyper-specific prefixes at around 2018. While we are not able to account this peak to a single factor, we observe that the increase is rather uniform across collector peers, origin ASes, intermediate ASes, and address space and, hence, is unlikely to stem from a measurement artifact or some local misconfiguration. When comparing the situation before and after the peak, we still can see an increase from $\sim 7K$ HSPs in 2016 to $\sim 11K$ HSPs in 2021. In contrast to IPv6, many HSPs in IPv4 can only be seen by one peer. While we observe few HSPs that can be seen by 100+ peers, the vast majority of HSPs can only be seen by 10 or less peers. Even though the number of low-visibility HSPs strongly fluctuates between snapshots, it increases rather continuously across many snapshots. Both such characteristics are significantly less pronounced for IPv4 HSPs that can be seen by 6+ peers. This difference may be accountable to various reasons including the association of a prefix to a certain function or a prefix’s lifetime.

HSP aggregation. ASes often have economical incentives to keep their BGP routing table size low. To realize this




Figure 9: HSP prefix contribution over time

goal, some ASes aggregate (multiple) more-specific routes into a single less-or-equally-specific route [7]. If an anchor-prefix results from aggregating prefixes with different CIDR sizes (prefix-based aggregation), we know that one of such pre-aggregation prefixes must have been hyper-specific. Yet, confidently identifying such aggregations is challenging. According to RFC 4271 [39], a router *MAY* set the AGGREGATOR field when it performs prefix-aggregation—which can serve as indication that *some* form of aggregation must have happened. Thus, we first extract all routes for anchor-prefixes which have the AGGREGATOR field set. At this stage, our selected routes might be a result prefix-based aggregation or the aggregation of different routes—e.g., with different AS_PATH attributes (path-based aggregation)—for the same prefix (or both). To reduce the likelihood of falsely identifying HSP usage due to path-based aggregation, we rely on the ATOMIC_AGGREGATE field as well as the presence of AS_SET elements in the AS_PATH attribute. A router *SHOULD* set the ATOMIC_AGGREGATE field if the newly generated AS_PATH attribute of the post-aggregation route does not contain all AS numbers present in the pre-aggregation routes, e.g., the paths *AB* and *AC* can be aggregated to *AB* (which hides the existence of *C*). If the ATOMIC_AGGREGATE field is not set, ASes often use AS_SETs to signal path-aggregation, e.g., the paths *AB* and *AC* can be aggregated to *A{B, C}* (where {...} denotes the AS_SET containing all ASes after *A*). As the ATOMIC_AGGREGATE field and AS_SETs indicate path-based aggregation, we remove all anchor-routes that contain at least one of them.

Where does HSP aggregation happen? Now that we have a set of anchor-prefixes that are likely the result of prefix-based aggregations, we can analyze how close to the origin HSPs are aggregated. We compare the AS number in the AGGREGATOR field with the AS_PATH and differentiate between the following cases: (1) **Origin**—the origin itself performed the aggregation, (2) **On-path**—an AS within the AS path that is not the origin performed the aggregation,




Figure 10: HSPs by # Peer ASes over time

and (3) **Off-path**—some AS that does not occur in the AS path performed the aggregation⁸. Figure 11 shows the number of anchor prefixes in each class over time. Notably, the figure also contains the class **Multiple** that contains anchor prefixes for which there are multiple paths with inconsistent classes. We observe that the vast majority of anchors are actually aggregated at the origin with only few hundreds of anchors being aggregated on-path. Origin and off-path (especially AGGREGATOR fields with private ASNs) aggregation often occurs due to the use of BGP confederations [6, 22] where the AS is internally split into multiple private sub-ASes. Depending on how an AS border router handles the aggregation of internal confederation routes, it might either correctly set the external AS number or leak the internal confederation AS Number in the AS_PATH or AGGREGATOR attribute. Notably, those HSP routes are likely not available to other ASes (including neighbors of the origin).

Projected actual usage. While our IRR snapshots produced actual HSPs, our final prefix-aggregation and ROAs only produced a list of anchor-prefixes that is likely to contain HSPs. Therefore, we decided to analyze the potential extent of HSP usage on the basis of anchor-prefixes. Figure 12 shows the number of IPv4 (left) and IPv6 (right) anchor-prefixes per data set (stacked) over time. Notably, the aggregated class only contains on-path aggregated anchor prefixes and the RPKI class only contains anchor prefixes for explicit HSP ROAs. The “multiple” class covers those entries that are visible via multiple data sources. We observe that the current route collector infrastructure misses roughly one-third of the of the anchor prefixes that potentially contain HSPs. We further observe a less noisy, linear increase in the number of anchor prefix for which HSPs are visible compared to the raw count of visible HSPs. Notably, some part of this increase can potentially be accounted to the increasing numbers of route collectors and route collector peers over time.

⁸This class also includes reserved AS numbers.




Figure 11: Position of HSP Aggregation




Figure 12: HSP anchors across data sets

Who uses hyper-specific prefixes? We leverage the “AS Classification Inferences” dataset described in ASDB [55] to classify ASes as Content, Education, Hypergiant, ISP (Stub), ISP (Transit), Tier 1, and Others. Figure 13 compares the classes of all BGP-visible ASes (left) to HSP origin ASes (right) over time. We find that in contrast to all origin ASes, HSP origins are more likely to be ISP (Transit) ASes. Interestingly, the majority of Tier 1 ASes is also originating HSPs. During the period of January 2019 until October 2021, we identify between 12 and 15 of the total 19 Tier 1’s as HSP origins. In contrast to the high share of Tier 1 HSP origins, we find that most hypergiants do not originate HSPs.

D REAL-WORLD EXPERIMENT

Does BGP reflect control plane reachability? Finally, we want to understand how much the lack of additional BGP vantage points impacts our observations on reachability. Hence, we configure a real-world experiment using the PEERING testbed [47] in which we announce an anchor prefix as well as multiple hyper-specific prefixes. Once those prefixes have converged, we run traceroutes from RIPE Atlas [42] probes and compare their resulting paths to those visible at route collectors.

Vantage points & resources. The PEERING testbed allocates Internet resources (specifically, IPv4/IPv6 address space and AS numbers) to its users based on approved experiment proposals. Once allocated, users can announce those resources via the testbed’s infrastructure. Given that the PEERING testbed strongly relies on third party resources (e.g., for hosting infrastructure), announcements must be designed carefully to not cause trouble or irritation for other network operators. For our experiment we use the address ranges 184.164.240.0/23 and 2804:269c:4::/46. More specifically, we utilize 184.164.240.0/24 and 2804:269c:4::/48 as anchor prefixes (i.e., they represent our control group) and announce HSPs only from the remaining address space⁹.

RIPE Atlas [42] is a measurement platform with probing devices (henceforth called probes) all over the world. To maximize probing coverage and minimize probing load, we choose at most one probe per AS. To reduce the likelihood of probe outage, we select only probes that are not tagged with system-problematic tags¹⁰. We further validate that selected IPv4 and IPv6 probes are tagged with system-ipv4-works, and system-ipv6-works, respectively. If an AS hosts multiple probes, we prefer dual-stack probes (such that we can use a consistent probe for our IPv4 and IPv6 measurements) over anchor probes (i.e., better equipped probes) over any other probes. If we still have multiple choices, we pick the probe that is tagged with the highest stability tag (e.g., system-ipv4-stable-90d). Our final probe set consists of 3097 probes distributed across 2990 IPv4 and 1239 IPv6 ASes.

Experimentation environment. The PEERING testbed currently has a total of 180 IPv4 and 152 IPv6 neighboring ASes. Yet, most neighbors do not support/redistribute HSPs. We identify supportive neighbors by iteratively announcing a /25 or /49 prefix from our allocated address space through each neighbor and analyzing the resulting update stream from RIPE RIS and Routeviews. Since, at this point, we only care about a “life sign” (i.e., whether or not *any* update was received) rather than full convergence, we adopt a short announcement cycle: We announce a prefix at the beginning of every full hour and withdraw it 30 minutes later¹¹. We identify a set of 8 IPv4 and 9 IPv6 neighboring ASes that

⁹In particular, we announce 184.164.241.0/25, 184.164.241.128/28, 184.164.241.255/32, 2804:269c:5::/49, 2804:269c:6::/64, 2804:269c:6:8000::/65, and 2804:269c:7::/128.

¹⁰tags: system-flash-drive-filesystem-corrupted, system-v1, system-no-flash-drive, system-flash-drive-bad-or-too-small, system-firewall-problem-suspected, system-trying-to-connect system-readonly-flash-drive, system-no-controller-connection, system-bad-firmware-signature, system-flakey-connection, system-flakey-power, system-flash-drive-problem-detected, and system-v2

¹¹These experiments ran between the May 1, 2021 and the May 3, 2021.




Figure 13: HSP origin AS classification over time.




Figure 14: PEERING testbed propagation results

redistribute HSPs. Notably, those ASes are distributed across 4 and 3 geographically separate Points of Presence.

Technical realization. Throughout May 21 and 22, 2021, we announce /24, /25, /28, and /32 IPv4 prefixes and /48, /49, /64, /65, /128 prefixes through a single neighbor at the beginning of every even hour. After the announcement, we wait 40 minutes to allow the prefix to converge¹². After those 40 minutes, we run active measurements for 10 minutes, and then withdraw the prefixes again. Notably, we choose 70 minutes between a withdrawal and the next announcement on purpose such that we out-wait the expiration of potential Route Flap Damping hold-down timers, which have been shown to usually expire after 60 or less minutes [16].

During our 10 minute active measurement period, we run paris-traceroutes from all probes towards either the network address or the first non-network address of all prefixes (which are configured to be pingable). To reduce the dependence of our results on the underlying protocol, we simultaneously issue ICMP, TCP, and UDP probing. To keep the induced load for the RIPE Atlas platform as well as for the peering testbed manageable, we reduce the number of probing packets used per per hop by paris-traceroute from RIPE ATLAS' default of 3 packets to one packet. Notably, as the resulting load still exceeds the default limitations (e.g., for measurement results per day) for a single RIPE Atlas account, we coordinate our probing efforts with the RIPE Atlas team who generously raised the limits for our experiments.

We map traceroutes to AS Paths using the state-of-the-art mapping tool bdrmapit [30]. As bdrmapit requires a large

¹²During previous experiments we observed that usually the 95th percentile of updates reach the collector peers already in the first 15 minutes.

corpus of traceroutes as input to perform well, we use traceroute data from CAIDA's IPv4 Prefix-Probing data set [53], CAIDA's IPv4 Routed /24 Topology Dataset [53], CAIDA's IPv4 Routed /48 Topology Dataset [54], and RIPE's hourly archives of Atlas traceroutes [41] between May 17, 00:00 and May 24, 00:00. For all the other inputs (e.g., prefix-to-origin mappings or business relationship inferences) we use recent snapshots from the recommended data sources. Finally, we use bdrmapit's output to map our successful (i.e., only those that actually reached the respective target host) traceroutes to AS paths.

Comparison. Figure 14 compares the the number of ASes (aggregated over all iterations) that (1) hosted Atlas probes that reached the target (ATLAS_SOURCE, yellow), (2) appeared along the path between ATLAS_SOURCE ASes and the Peering Testbed (ATLAS_PATH, dark red), (3) are visible from route collector peers (BGP, gray). The most drastic observation is that hyper-specific prefixes see a very sharp drop in reachability. Even the best performing CIDR size, /25, only reached ~15 % of of the ASes that are reached by its respective anchor prefix. Especially for IPv6 we observe that most PEERING neighbors redistribute our prefixes (including the anchor prefix) only towards their customers, hence, some of our Atlas probes are unable to reach the peering testbed even for the anchor prefix. We further find that the more-specific the prefix gets, the less likely it propagates. This finding is interesting as most recommended filtering guides [11, 12, 29, 33, 34] treat all hyper-specific CIDR sizes equally. Our third observation is that the reachability reflected by route collector peers substantially underestimates

data plane reachability. While we are able to observe approximately a third of the total ASes for our /48 prefix via BGP, this fraction lies at around 14 % for our /24 prefix.

E FILTERING PIPELINE

When an AS peers with a Route Collector, the router that feeds the collector may provide all routes that are not removed during (or before) egress filtering. Hence, misconfigured egress filters can lead to misinterpretations. For our

analysis, we filter out HSPs which are originated by feeder AS directly connected to a route collector. However, we use the HSP if it has been propagated to at least 2 AS hops, including feeder AS. In addition, we filter all private, reserved, multicast, and experimental IP prefixes. Furthermore, we also filter prefixes originated by a private AS. Finally, we remove the HSPs we identify as outliers during the data cleaning process. Appendix F provides detail information on HSPs we have filtered out.

F APPLIED DATA ISOLATION RULES

Timeframe	Filter name	Filter Details	Reason
entire period	Private Origin ASes 2 Bytes	Origin AS number from 64512 to 65534	private IPv4 ranges.
entire period	Private Origin ASes 4 Bytes	Origin AS number from 4200000000 to 4294967294	private IPv4 ranges.
entire period	Private IPs	IPv4 Private IP ranges	private IPv4 ranges.
entire period	Class D and E	IPv4 Prefixes > 223.x.x.x	IPv4 multicast and class E IP ranges.
entire period	Abnormal Prefixes	for IPv4 prefix > /32 for IPv6 prefix > /128	abnormal IPv4 prefixes. abnormal IPv6 prefixes
entire period	No Origin Internal	Routes having no origin AS Feeder AS is the Origin AS	AS-internal routes.
2015/10/01-07	IPv4 Noisy Origins	Origin AS == 9498	routes from particular origin AS.★
2016/10/01-07	IPv4 Noisy Origins	Origin AS == 36937	routes from particular origin AS.★
2017/04/01-07	IPv4 Noisy Origins	Origin AS == 9498	routes from particular origin AS.★
2019/07/01-07	IPv4 Noisy Origins	Origin AS 7122	routes from particular origin AS.★
entire period	IPv4 Noisy Origins	Origin AS 12400	routes from particular origin AS.★
2016/07/01-07	IPv4 Noisy Peer AS	Peer AS 35908	routes from particular peer AS.★
2017/01/01-07	IPv4 Noisy Peer AS	Peer AS 60924 and 27630	routes from particular peer AS.★
2017/10/01-07	IPv4 Noisy Peer AS	Peer AS 37497	routes from particular peer AS.★
2018/10/01-07	IPv4 Noisy Peer AS	Peer AS 14361	routes from particular peer AS.★
2019/01/01-07	IPv4 Noisy Peer AS	Peer AS 262757	routes from particular peer AS.★
2020/04/01-07	IPv4 Noisy Peer AS	Peer AS 268430	routes from particular peer AS.★
2021/04-07/01-07	IPv4 Noisy Peer AS	Peer AS 398465	routes from particular peer AS.★
2021/01-10/01-07	IPv4 Noisy Peer AS	Peer AS 203125	routes from particular peer AS.★
2020/04-07/01-07	IPv4 Noisy Peer AS	Peer AS 268430	routes from particular peer AS.★
entire period	IPv6 Noisy Origins	Origin AS 4761	routes from particular origin AS.★
2017/07/01-07	IPv6 Noisy Origins	Origin AS 17451 and 45899	routes from particular origin AS.★
2019/04/01-07	IPv6 Noisy Origins	Origin AS 7713	routes from particular origin AS.★
2021/07/01-07	IPv6 Noisy Origins	Origin AS 8100	routes from particular origin AS.★
2018/07/01-07	IPv6 Noisy Peer AS	Peer AS 199036	routes from particular peer AS.★

Table 1: Applied filtering and isolation rules. ★: these ASes contributed/announced either (1) an extraordinary high number of HSPs (i.e., 100 or more times higher than in other snapshots) or (2) HSPs in an extraordinary high number of anchor prefixes for a limited time.