No Easy Way Out — A Study of Single-homed ISPs

Paper #374, 6 pages body, 9 pages total

Anonymous Author(s)

ABSTRACT

In 2023, Italy's largest internet service provider (ISP), TIM (formerly Telecom Italia), was struck by a network outage which affected thousands of users and lasted nearly five hours. The root cause was a connectivity problem at its only upstream autonomous system (AS), Telecom Italia Sparkle. Apparently, even large ISPs depend on a single upstream. This paper quantifies how widespread this behavior is on today's Internet and identifies other at-risk countries. First, we analyze BGP data and show that seemingly single-homed networks are not uncommon. Focusing on large eyeball networks, we perform active measurements using bidirectional traceroutes and BGP poisoning to confirm if they are indeed single homed. Our analysis concentrates on 247 ASes from 163 countries for IPv4 and 141 ASes from 113 countries for IPv6. We found that BGP data is mostly accurate and we confirmed that 70% of analyzed ASes in IPv4 and 80% in IPv6 are indeed single homed. This also implies that 78% of analyzed countries in IPv4 and 81% in IPv6 have at least one large single-homed provider. We find that a provider and its upstream often belong to the same company and discover other incumbents, like TIM, that are single homed.

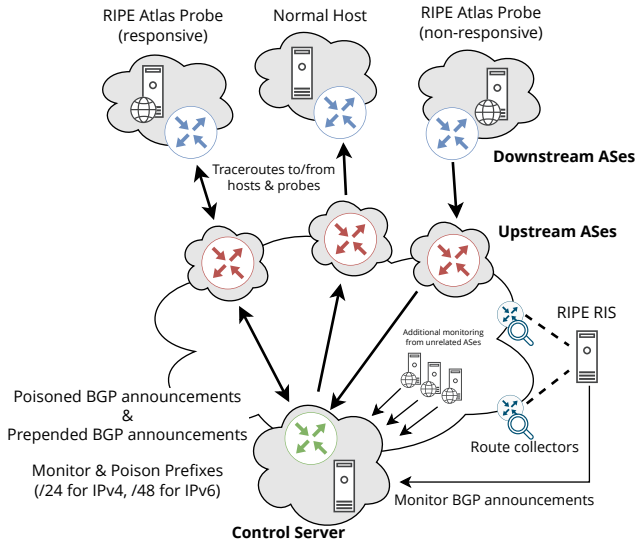

Figure 1: Overview of the measurement infrastructure.

1 INTRODUCTION

Billions of humans all over the world rely on internet service providers (ISPs) to connect to the Internet and expect uninterrupted access as part of their daily life. However, even established ISPs are not safe from outages, as was the case with TIM, Italy's largest ISP, in 2023 [27]. This outage lasted nearly five hours and affected one third of Italy's Internet users. The root cause lay not within TIM itself though, but at its only upstream network, Telecom Italia Sparkle. This relationship is also visible in BGP: to reach IP prefixes announced by TIM, one must traverse Sparkle.

Motivated by this example we set out to explore how widespread this pattern — a large ISP with only a single upstream (i.e., single homed; Section 2) — is on today's Internet. We use BGP data [2, 22] as the starting point to find potential single-homed autonomous systems (ASes) (Section 3). This data is used frequently in research to measure the Internet's resilience [5, 12, 25]. However, BGP provides a partial view that only reveals the "best" path towards IP prefixes.

We bridge the gap between BGP and the true network topology with active measurements (Section 4). We confirm if an AS has truly only one upstream, as claimed by BGP, with a combination of traceroute measurements and BGP poisoning (Section 4.1). In a nutshell, we run traceroutes from an experimental IP prefix to eyeball ASes and then poison their upstreams (see Figure 1). The poisoning causes upstreams to lose connectivity to the prefix and if the eyeball ASes lose connectivity as well, we confirm that they are single homed. Our measurements focus on large eyeball ASes where outages would directly impact many users (Section 5).

Our results (Section 6) show that BGP data is mostly accurate. We confirm that many eyeball ASes are indeed single homed and that even the largest providers are no exception. In addition, these cases are not localized, but affect countries all over the world.

In summary, this paper makes the following contributions:

- Overview of single-homed ASes in BGP (IPv4 & IPv6).
- Validation of BGP data with active measurements for 247 ASes from 163 countries for IPv4 and 141 ASes from 113 countries for IPv6.
- Characterization of 173 confirmed single-homed eyeball ASes for IPv4 (113 IPv6) by their size, relationship to upstream, and location.

2 TERMINOLOGY

First, we introduce the terminology used throughout the paper. Our analysis focuses on single-homed ASes: Networks that only use a single transit provider. We call the single-homed AS the *downstream* and the transit provider the *upstream* (see Figure 1), which together form a *pair*. In BGP the upstream is seen as the preferred way of the downstream to reach the wider Internet. The downstream may still have connections to peer and customer ASes, but if the upstream fails, the downstream loses connectivity to the rest of the Internet.

To detect potentially single-homed ASes we use AS Hegemony results [10], which are based on BGP and quantify the dependency between ASes on the Internet as a percentage. If AS **A** has a 100% dependency on another AS **B**, virtually all paths towards **A** traverse through **B**.

AS Hegemony also reveals dependencies between ASes that are not direct neighbors. In this case our definition of single homed differs from the one network operators traditionally use: We include all ASes that have a 100% dependency on another AS, even if they are not direct neighbors. In some cases, a downstream depends on multiple upstreams that are arranged in a chain making it especially vulnerable to outages. In Section 6.4 we show that an outage of the indirect upstream can indeed cause the downstream to lose connectivity.

3 SINGLE-HOMED NETWORKS IN BGP

We quantify the prevalence of single-homed networks as seen in BGP and identify the most prominent upstreams. We analyze a snapshot of AS Hegemony [15] and extract all AS pairs with a dependency of 100%.

Our analysis shows that single-homed networks and indirect dependencies are not uncommon. There are around 60k pairs consisting of 43k unique downstreams and 7k upstreams present in IPv4, and 37k pairs of 22k downstreams and 4k upstreams in IPv6. Indirect dependencies account for around 28% of IPv4 pairs and 41.5% of IPv6 pairs, highlighting the need to include them in the analysis.

The pairs of IPv4 and IPv6 are mainly disjoint. 81% of IPv4 pairs and 70% of IPv6 pairs only occur in their respective address family. Around 10k downstreams are present in both IPv4 and IPv6, however, 2366 of these have a slightly different set of upstreams and for 884 the sets are disjoint.

We also look at the upstreams with the most dependent downstreams and find that there are staggering differences between IPv4 and IPv6. Figure 2 shows the five upstreams with the most downstreams. For IPv4 we see an expected picture of Tier-1 providers (e.g., Lumen AS3356, Hurricane Electric AS6939), which have up to 2450 downstreams.

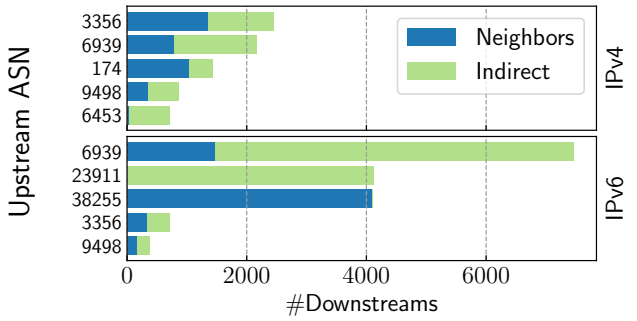

Figure 2: The five upstreams with the most downstreams for IPv4 and IPv6.

For IPv6, Hurricane Electric's impact is remarkable with 7470 downstreams (1483 direct; 5987 indirect). Standing out, and a prime example of a dependency chain, are AS38255 (China Education and Research Network) and AS23911 (China Next Generation Internet Beijing IX). All of AS38255's 4097 neighbor downstreams also indirectly depend on AS23911.

4 METHODOLOGY

One goal of this paper is to confirm if a single-homed AS as seen by BGP is really single homed. The basic idea is to run traceroutes from a prefix controlled by us to a target in the downstream AS. We then poison the upstream AS using BGP poisoning, which forces the upstream to drop routes to the prefix. If the target stops responding to traceroute probes, we know that the downstream is single homed, as the return packets cannot be delivered.

4.1 BGP Poisoning

BGP poisoning is a technique that leverages the routing loop detection of BGP[1] to prevent a remote AS from accepting a route announcement. Figure 3 shows an example for an AS **X** announcing a prefix, which propagates via the best path **U Y X** to the downstream AS **D**. The alternate path **B Y X** would not be shared with route collectors by **D**. However, **X** can insert **U** into the AS_PATH attribute of the announcement. When **U** receives this announcement, it will detect its AS number in the path and refuse the route to prevent a routing loop. As a consequence, **U** will not announce a route to the prefix to **D**, which promotes the alternate path to the best path. However, if no alternate path exists, **D** cannot reach the prefix and is thus single homed.

4.2 Measurement Infrastructure

We now explain our measurement infrastructure using Figure 1. We use one AS and two /24 (/48) prefixes for IPv4

[1] See Section 9.1.2. of RFC 4271 [21]

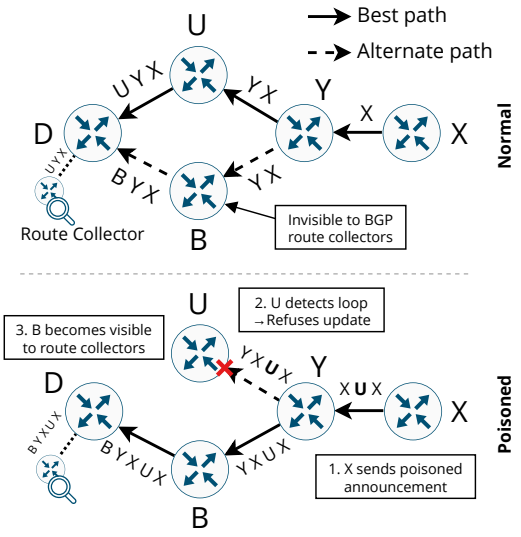

Figure 3: Overview of BGP Poisoning. *Normal*: Only best paths are visible to BGP route collectors. *Poisoned*: A poisoned BGP announcement can force a change in the topology and reveal previously hidden paths.

(IPv6), which are exclusively used for the purpose of this measurement. All prefixes have correct IRR entries and valid RPKI ROAs configured. The *poison* prefix is announced in turn with and without a poisoned AS path. We use the *monitor* prefix to confirm that our AS is reachable at all times in order to prevent false positive results. For this purpose, the monitor prefix should exhibit the same announcement pattern as the poison prefix (e.g., to account for Route Flap Dampening [20, 29]). Therefore, we announce both prefixes in lockstep but for the monitor prefix we mimic the `AS_PATH` attribute change using AS-path prepending. In addition, we use all route collectors of RIPE RIS [22] to monitor the poisoned announcement propagation. Also, we check connectivity to both prefixes from Atlas probes in 90 unrelated ASes.

One IP from each prefix is assigned to the control server, which sends traceroutes from each interface to multiple targets in the downstream AS. In total, we target up to ten destinations in the downstream and favor destinations from different IP prefixes. If available, we also request traceroutes from up to ten RIPE Atlas probes [23] located in the downstream to our prefixes. The selected target set is discussed in Section 5.1.

4.3 Unpoisonable Upstreams

We found that the local transit providers of our experimental AS do not accept poisoned announcements for a certain set of ASes, mostly consisting of Tier-1 providers (for the full list see Appendix B). We suspect that they are protected by peer locking [18]. However, we managed to perform measurements for most of these upstreams nonetheless by using BGP communities [16].

We use a BGP community to signal our transit providers that they should not announce the poisoned prefix to the targeted upstream, hence having the same effect as BGP poisoning. This works for two reasons: Since all targeted upstreams are large networks (1) our transit providers peer with them and (2) they do not accept the announcement via a different neighbor, because they are transit free. We manually confirmed that the targeted upstreams lose connectivity with a combination of Atlas probes and BGP looking glasses.

4.4 Measurement Timing

The measurement runs in hour-long cycles that are divided into three sequential phases of 15, 30, and 15 minutes each. In the first phase both prefixes are in their *normal* state (unpoisoned & not prepended) to ensure reachability from the downstream to the experimental AS. At the start of the second phase, the control server sends the poisoned and prepended announcements, moving to the *poisoned* state. In the third phase, the control server returns the prefixes to their normal state and we confirm that connectivity recovers.

Effectively, we send a BGP announcement every 30 minutes per prefix, which prevents blocking of announcements due to Route Flap Dampening [20, 29] and leaves enough time for the announcement to propagate. The control server and Atlas probes perform traceroute measurements every five minutes during the entire cycle. To minimize the impact on the targets and prevent rate limiting, we opted for a low measurement frequency and apply conservative traceroute parameters, which are detailed in Appendix C.

4.5 Quality Assurance

To ensure reliable results, we mark pairs as *Inconclusive* if they behave inconsistent, have too few targets, or lack visibility of the upstream in traceroute.

If, for a single target host, three or more traceroutes to/from the monitor prefix fail, the target is excluded. If as a consequence the remaining number of targets falls below three, the pair is marked as inconclusive.

In addition, if the upstream is not visible as a traceroute hop during the normal state, the pair is excluded as well. We observed that this mostly occurs because the upstream does not reply to traceroute probes or does not announce any IP prefixes in BGP, making an IP-to-AS mapping impossible.

4.6 Categorization

After each measurement cycle we categorize the pair based on the majority behavior of its traceroute targets. A pair

is called *unidirectional* if only traceroutes from the control server to the downstream are available, or *bidirectional* if Atlas probes provide return path information. If connectivity is lost, the pair is *single homed*. If connectivity remains and the pair is bidirectional it can be either *multi homed* or *unknown*. The pair is categorized as *multi homed* in case we see a path change or as *unknown* if we do not. In the unknown case, it is likely that the downstream has a default route to the upstream. Finally, if connectivity remains and the pair is unidirectional, it is also classified as unknown since we cannot confirm the existence of an alternate path.

5 MEASUREMENT CANDIDATES

We first discuss the set of eyeball ASes used for the measurement followed by an explanation of the traceroute target selection. Next, we outline the AS pairs that are filtered due to unpoisonable upstreams (Section 4.3) and quality assurance (Section 4.5). We conclude with a short discussion of the measurement consistency.

5.1 Target Selection

AS Pairs. Due to the plethora of potential single-homed ASes and inspired by our motivating example, we focus only on large eyeball ASes in our active measurement. Thus, we only include downstreams that cover at least 5% of the population of any country according to the APNIC population estimate [4]. In addition, we lower the AS hegemony threshold to 95% to include pairs that are potentially single homed, but have no perfect dependency due to measurement noise. For example, TIM has a permanent <1% dependency on an Akamai DDoS mitigation network (AS32787) [14] causing the dependence to Sparkle to drop below 100%. These selection criteria result in 410 pairs (192 upstreams, 326 downstreams) for IPv4 and 402 pairs (174 up, 289 down) for IPv6.

Traceroute. For our traceroute measurements we select up to ten targets from each downstream. We extract responsive IPs from CAIDA's Ark topology datasets for IPv4 [7] and IPv6 [8], yielding 16.6M IPs covering 53.7k ASes for IPv4. For IPv6 we additionally use a snapshot of the IPv6 Hitlist [1, 11, 26, 30], resulting in a total of 18.6M IPs covering 24.3k ASes.

We also extract responsive IPs from OpenINTEL's active DNS measurements [19, 28]. We prefer these targets since they are hosts inside the network, as opposed to router interfaces, which might reside at the border of an AS. This process adds 1.6M IPs covering 41.1k ASes for IPv4 and 589k IPs covering 8.8k ASes for IPv6.

For IPv4, we find five pairs where no responsive target is available, but we can fulfill the goal of ten targets for 360 (93%) of the remaining pairs (cf. Section 5.2). For 100 pairs (25%) all ten targets are in different /24 prefixes. Overall we achieve a median number of 8 distinct /24 prefixes per pair.

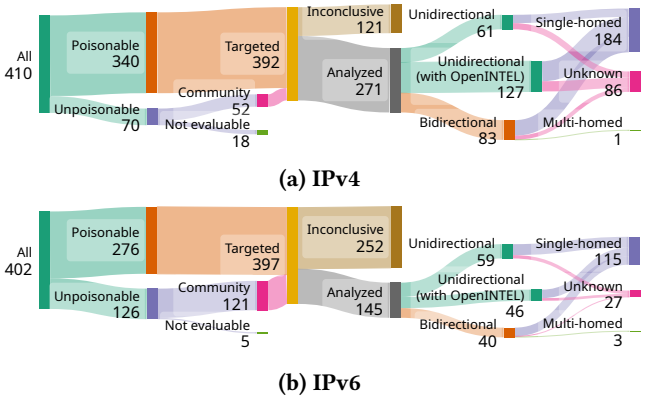

Figure 4: Breakdown of steps required to analyze pairs.

For IPv6 the target selection is more challenging, possibly caused by a lower adoption and a larger search space. A total of 69 pairs had no responsive targets, but we still achieve ten targets for 251 (76%) of the remaining pairs. We get ten distinct /48 prefixes for 62 pairs (19%) and find a median number of 6 prefixes per pair in total.

5.2 Quality Assurance

AS Pair Filtering. Figure 4 breaks down the steps from the initial set to the final categories. The initial set contains 410 and 402 pairs for IPv4 and IPv6 respectively.

The *Unpoisonable* set consists of 16 upstreams responsible for 70 (IPv4) and 126 (IPv6) downstreams. However, we managed to perform measurements for 13 (IPv4) and 14 (IPv6) upstreams by using BGP communities (Section 4.3).

Moving to the *Targeted* set, as mentioned before we find 5 pairs with no responsive target for IPv4 and 69 pairs for IPv6. The remaining pairs in the *Inconclusive* set are caused by factors explained in Section 4.5 and some corner cases.

Measurement Consistency. Since we base the categorization on the majority of traceroute targets (Section 4.6), we inspect how representative the majority is. For 89% (IPv4) and 92% (IPv6) of pairs we confirmed that all targets showed consistent behavior, and for the remaining pairs on average 82% (IPv4) and 81% (IPv6) of targets agreed. Different per-prefix routing policies could be one reason for inconsistent behavior. Since inconsistent targets represent a small fraction, we assume it is safe to ignore them.

Although the data presented in this paper is only based on one measurement round, we repeated the measurement for IPv4 three times over the duration of one month with a randomized set of targets and found the results to be stable. While there were some pairs that were inconclusive in one of the iterations, caused by targets going offline or a change in dependency, 249 pairs were analyzed in all three iterations

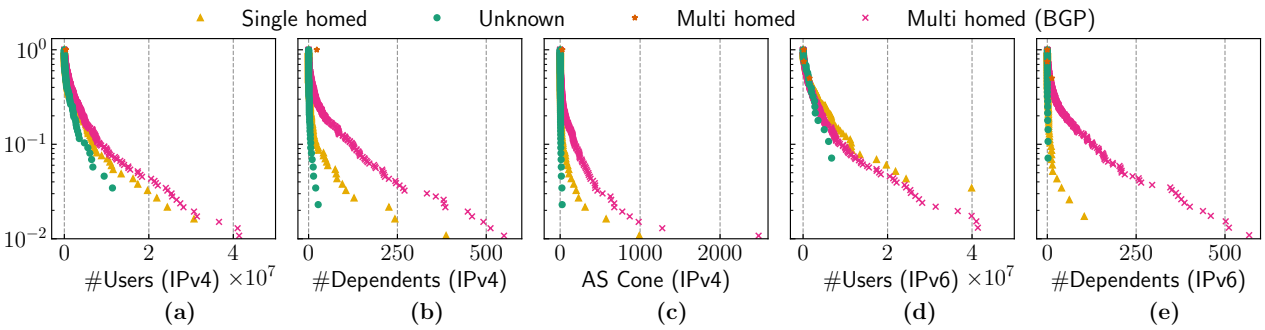

Figure 5: Downstreams by different size metrics. AS Cone is unavailable for IPv6.

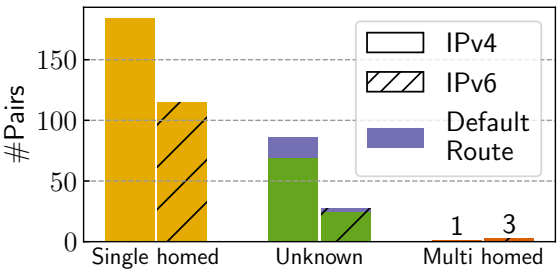

Figure 6: Category distribution. Marked pairs in the *Unknown* category likely use a default route.

and 94% of these remained in the same category. We conclude that the analysis is stable given sufficient targets.

6 RESULTS

We now discuss the final results of our study. First, we give an overview of the categories and show that even large eyeball networks can be single homed. Next, we compare IPv4 and IPv6 and highlight prominent upstreams. We finish with a discussion of single-homed downstreams that serve a particularly large customer base and show that many countries have at least one large single-homed ISP. We publish the detailed results at [3].

6.1 Confirmed Single-homed Networks

The analyzed set of ASes consists of 271 pairs (134 upstreams, 247 downstreams) for IPv4 and 145 pairs (89 up, 141 down) for IPv6. The category distribution shown in Figure 6 reveals that the majority of analyzed ASes is indeed single homed.

While 184 pairs (68%) behave single homed for IPv4, almost a third of pairs does not. Unfortunately, most of these are based on unidirectional measurements and we can not distinguish between a default route and an alternate path. For 17 pairs (highlighted purple in Figure 6) Atlas probes were available in the downstream and we could confirm that a default route was installed, as no path change was observed. Finally, there is only one confirmed case of an alternate path in IPv4 (described in Appendix D).

For IPv6 the relative share of single-homed networks is larger with a total of 115 pairs (79%). We see 27 pairs that keep responding out of which at least two have installed a default route and three instances of potential default routes.

6.2 Downstreams by Size

We now inspect the downstream size using different metrics and look for a relationship between the size of a network and its upstream diversity. We show CCDFs of the number of users [4], dependent networks [10], and AS Cone [9, 17] in Figure 5. For comparison we also show data points for large eyeball ASes (with the same population threshold) that are multi homed in BGP, i.e., without 100% dependency on another AS.

We make one key observation: the largest eyeball ASes seen as single homed in BGP, are indeed single homed (Figure 5a and d). While this might seem surprising, for networks of this size keeping an alternate connection with sufficient capacity on standby is not economically viable.

However, we also note that the largest ASes are multi homed, painting a good picture for the Internet topology. The difference is particularly pronounced for networks with many dependents (Figure 5b and e). This also matches our expectations, since these networks are usually acting as transit providers and therefore should offer diverse connectivity to their customers.

6.3 Comparison of IPv4 and IPv6

Comparing the analyzed pairs of IPv4 and IPv6 we see that there is a large set of 173 disjoint IPv4 pairs, matching the observation of Section 3. However, there are only 47 pairs unique to IPv6 and the remaining 98 pairs overlap. Within the overlapping pairs, 9 belong to different categories.

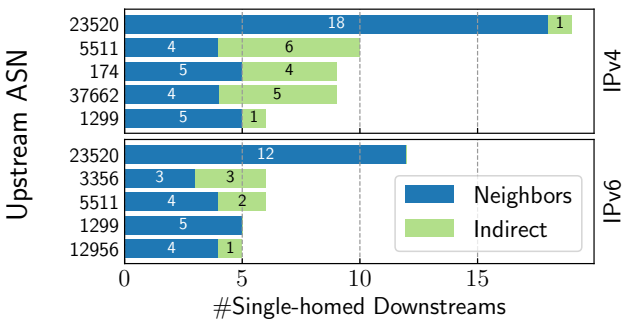

Figure 7: Upstreams with the most downstreams.

Comparing only the downstreams, we find an overlap of 99 ASes of which only 11 have completely or partly different upstreams in IPv4 and IPv6. Overall, if a network is present in both IPv4 and IPv6, it usually behaves the same, indicating that providers do not have different policies in place.

6.4 Prominent Upstreams

We repeat the analysis of Section 3 to find the upstreams that have the most single-homed downstreams (Figure 7) and see a similar pattern emerge.

First, we note that upstreams with multiple single-homed downstreams are common: 39% of upstreams for IPv4 and 33% for IPv6 have more than one downstream. Figure 7 shows the five upstreams with the most single-homed downstreams and we observed that this number can grow quite large. Columbus Networks (AS23520) has the most single-homed downstreams with 19 in IPv4 and 12 in IPv6, with 10 being present in both. These are mostly ISPs for countries in South America and we find four (IPv4) and three (IPv6) downstreams that belong to Columbus Networks itself.

Finally, a manual inspection revealed three types of upstreams: upstreams with mostly subsidiary downstreams (i.e., from the same company), with unrelated downstreams, and a mix of both. Cogent (AS174), West Indian Ocean Cable Company (AS37662) and Lumen (AS3356) have no subsidiaries in their downstreams. The opposite are Arelion (AS1299) and Telxius Cable (AS12956) which almost exclusively have subsidiaries (see also appendix Table 3).

6.5 Countries with Single-homed Networks

In this final section we show that single-homed eyeball ASes exist all over the world. In total we analyzed eyeball ASes in 163 countries for IPv4 and 113 countries for IPv6 and found single-homed ASes for 127 countries (78%; IPv4) and 91 countries (81%; IPv6). While some of these countries are relatively small, our findings show that users in many countries are served by ASes that rely on a single upstream for connectivity.

Table 1: Example countries with single-homed ISPs. Incumbents are marked with an asterisk.

Country	Company	AS Pair		Pop. Cov.
Australia	Telstra*	4637	1221	46.54%
France	Orange*	5511	3215	35.70%
Japan	NTT*	2914	4713	12.07%
UAE	Etisalat*	8966	5384	69.05%
USA	Verizon	701	6167	9.30%

We also discovered that the relationship between the up- and downstream in these cases is often subsidiary, i.e., both ASes belong to the same company. A sample of countries and incumbent providers is shown in Table 1 together with their population coverage. Some notable examples include NTT for Japan, Orange for France, and Telstra for Australia. A more extensive list is in appendix Table 3. Although this list is not exhaustive, it shows that Italy with TIM is not a special case and similar problems could affect many countries all over the world.

7 RELATED WORK

Bush et al. [6] analyzed IPv4 Internet reachability via ping and traceroute measurements. As part of their analysis they used BGP poisoning to discover hidden upstreams by poisoning the visible upstreams of an AS. While they only analyze IPv4 they come to a similar conclusion, namely that a majority of inspected ASes has no hidden upstream. We show that this behavior is not limited to small ASes, but does affect large eyeball ASes as well.

Building on this work, Rodday et al. [24] investigated the prevalence of default routes on the Internet. They extended the approach with Atlas probes and also consider IPv6. 57% of their inspected ASes for IPv4 and 29% for IPv6 use default routes, which is higher than our share of *Unknown* results. This is possibly explained by a different choice of ASes, but their data is not available anymore, which prevents us from comparing our datasets in more detail.

8 CONCLUSION

This paper presented a study of single-homed ISPs. We found that single-homed ASes are common in BGP and verified with active measurements, focusing on large eyeball ASes, that BGP data is mostly accurate. We characterized the measured ASes by size and found that even the largest ASes have no secondary upstream. There are upstreams with multiple single-homed downstreams, some of which are not direct BGP neighbors. Finally, we discovered that single-homed ISPs are not limited to a certain region, but exist globally. To aid future research, we publish our detailed measurement results at [3].

A ETHICS

We perform active measurements concerning three parties: Our experimental AS, RIPE Atlas probes, and other hosts in the downstream AS. Our AS and prefixes are exclusively used for the purpose of this measurement and thus no other traffic is affected by the poisoned announcements. We produce a small amount of BGP churn (96 updates per day), which is negligible compared to the 180k updates the Internet sees each day [13]. Atlas probe operators voluntarily participate in the measurement platform and have the option to disable traceroute responses, in which case we do not target them. For traceroute targets in general, we use conservative parameters and timings (Appendix C) so that the target hosts only receive 72 packets (5.3kB) in one hour. In summary, to the best of our knowledge our research does not raise any ethical concerns.

B UNPOISONABLE UPSTREAMS

There are 16 upstreams for which the transit providers of our experimental AS do not accept poisoned announcements, but we were able to perform measurements for 13 (IPv4) and 14 (IPv6) upstreams by using BGP communities.

One interesting upstream was Hurricane Electric (AS6939), since it was the only network with different behavior for IPv4 and IPv6. For IPv4 they accept an alternate route via Arelion (AS1299), but not in IPv6. This might hint at the different roles of Hurricane Electric in IPv4 and IPv6.

Table 2: Upstreams for which poisoned announcements are not propagated. A checkmark indicates that the alternative BGP community approach worked.

Name	ASN	IPv4	IPv6
Cogent	174	✓	✓
Verizon	701	✓	✓
Vodafone Global Network	1273	✗	✗
Arelion	1299	✓	✓
NTT Global IP Network	2914	✓	✓
GTT Communications	3257	-	✓
Deutsche Telekom	3320	✓	-
Lumen	3356	✓	✓
PCCW Global	3491	✓	✓
Orange / OpenTransit	5511	✓	✓
TATA Communications	6453	✓	✓
Zayo	6461	✓	✓
Telecom Italia Sparkle	6762	✓	✓
Liberty Global	6830	✓	✓
Hurricane Electric	6939	✗	✓
Telxius Cable	12956	✓	✓
Google	15169	✗	✗

C MEASUREMENT PARAMETERS

We used conservative timing and measurement parameters to ensure a sound measurement.

The BGP announcement interval of 30 minutes is based on the recommendations specified in RFC 7196 [20], which is more aggressive than Cisco and Juniper defaults, but caused no problems as confirmed by our BGP monitoring.

The traceroute parameters were chosen for a slower but more reliable execution:

- `-n`: No DNS lookup
- `-N 3`: Only 3 parallel probes
- `-z 0.5`: 0.5 seconds wait between probes
- `-w 5`: Fixed 5 second timeout (no adaptive timeout)

D CONFIRMED MULTI-HOMED NETWORKS

For IPv4, Elisa Eesti (AS2586), an AS of the Finnish company Elisa and second largest eyeball network in Estonia, has a 100% dependency on Elisa's main network (AS6667). When poisoned, we observed a path change to Arelion (AS1299), which was already visible in the path before indicating that the downstream simply bypassed the upstream.

For IPv6, Cgates (AS21412), a Lithuanian network, changed from Hurricane Electric (AS6939) to Arelion, which seems to be a completely new path. In contrast, Reunicable (AS37002), also a downstream of Hurricane Electric, switched to Vodafone Global Network (AS1273) and simply bypassed the upstream similar to the case in IPv4. Finally, CGI Norge (AS25225) changed from Next Layer Telecommunications (AS1764) to A1 Telekom Austria (AS8447)

REFERENCES

[1] [n. d.]. IPv6 Hitlist Service. https://ipv6hitlist.github.io/ Accessed: 2024-03-16.

[2] [n. d.]. University of Oregon Route Views Project. http://www.routeviews.org/

[3] 2024. Redacted to adhere to anonymity guidelines.

[4] APNIC. 2024. Visible ASNs: Customer Populations (Est.). https://stats.labs.apnic.net/aspop

[5] Alfred Arouna, Ioana Livadariu, Azan Latif Khanyari, and Ahmed Elmokashfi. 2023. On Large-Scale IP Service Disruptions Dependencies. In *International Conference on Network and Service Management (CNSM'23)*. IEEE, 1–5. https://doi.org/10.23919/cnsm59352.2023.10327794

[6] Randy Bush, Olaf Maennel, Matthew Roughan, and Steve Uhlig. 2009. Internet Optometry: Assessing the Broken Glasses in Internet Reachability. In *Internet Measurement Conference (IMC'09)*. ACM, 242–253. https://doi.org/10.1145/1644893.1644923

[7] CAIDA. 2007. Ark IPv4 Routed /24 Topology. https://doi.org/10.21986/CAIDA.DATA.ARK-IPV4-TRACEROUTE Dates used: 2023-07-02, 2023-08-30, 2023-11-30, 2023-12-01, 2024-03-08.

[8] CAIDA. 2008. Ark IPv6 Topology Dataset. https://doi.org/dataset/ipv6_allpref_topology Dates used: 2024-04-01, 2024-04-05, 2024-04-06.

Table 3: Selection of countries with single-homed ISPs in a subsidiary structure. Countries where the ISP is an incumbent are marked with an asterisk.

Country		Upstream		Downstream	Pop. Coverage
Australia*	4637	Telstra Global	1221	Telstra Corporation	46.54%
Brazil	4230	Claro	28573	Claro (NET)	10.12%
DRC	36994	Vodacom	37453	Vodacom Congo	27.55%
Estonia	1299	Arelion	3249	Telia Eesti	40.97%
Finland*	6667	Elisa Corporation	719	Elisa Finland	32.64%
Finland	1299	Arelion	1759	Telia Finland	24.98%
France*	5511	Orange / OpenTransit	3215	Orange France	35.70%
Greece*	12713	OTEGlobe	6799	OTEnet	41.69%
India	64039	Jio	55836	Reliance Jio Infocomm	48.17%
Italy*	6762	Telecom Italia Sparkle	3269	Telecom Italia / TIM	17.93%
Japan*	2914	NTT Global IP Network	4713	NTT Communications (OCN)	12.07%
Namibia*	20459	Telecom Namibia	36996	Telecom Namibia	38.96%
New Zealand*	4648	Spark New Zealand	4771	Spark New Zealand	36.00%
Norway	1299	Arelion	41164	Telia Norge	7.70%
Oman*	8529	Zain Omantel International	28885	Omantel	54.30%
Poland	5511	Orange / OpenTransit	5617	Orange Polska	22.33%
Portugal*	8657	MEO	3243	MEO Residential	30.12%
Saudi Arabia*	39386	Saudi Telecom Company	25019	Saudi Telecom Company	19.55%
Spain*	12956	Telxius Cable	3352	Telefónica Spain	23.44%
Sweden*	1299	Arelion	3301	Telia Sweden	30.40%
Türkiye*	9121	Türk Telekom	20978	Türk Telekom	7.67%
UAE*	8966	Etisalat	5384	Etisalat	69.05%
United Kingdom*	6830	Liberty Global	5089	Virgin Media	16.58%
USA	701	Verizon	6167	Verizon	9.30%

[9] CAIDA. 2013. AS Rank: A ranking of the largest Autonomous Systems (AS) in the Internet. https://asrank.caida.org/ Date used: 2024-03-01.

[10] Romain Fontugne, Anant Shah, and Emile Aben. 2018. The (Thin) Bridges of AS Connectivity: Measuring Dependency Using AS Hegemony. In *Passive and Active Measurement Conference (PAM'18)*. Springer, 216–227. https://doi.org/10.1007/978-3-319-76481-8_16

[11] Oliver Gasser, Quirin Scheitle, Pawel Foremski, Qasim Lone, Maciej Korczynski, Stephen D. Strowes, Luuk Hendriks, and Georg Carle. 2018. Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists. In *Internet Measurement Conference (IMC'18)*. ACM, 364–378. https://doi.org/10.1145/3278532.3278564

[12] Bradley Huffaker, Romain Fontugne, Alexander Marder, and kc claffy. 2023. On the Importance of Being an AS: An Approach to Country-Level AS Rankings. In *Internet Measurement Conference (IMC'23)*. ACM, 52–65. https://doi.org/10.1145/3618257.3624798

[13] Geoff Huston. 2024. BGP in 2023 — BGP updates. https://blog.apnic.net/2024/01/10/bgp-in-2023-bgp-updates/

[14] IHR. 2024. AS Hegemony for AS3269. https://ihr.iijlab.net/ihr/api/hegemony/?timebin=2024-03-01T00:00:00Z&af=4&asn=6762,32787&originasn=3269

[15] IHR Archive. [n. d.]. AS Hegemony Scores. https://ihr-archive.iijlab.net/ihr/hegemony/ Dates used: 2024-03-01 (IPv4), 2024-04-01 (IPv6).

[16] Tony Li, Ravi Chandra, and Paul S. Traina. 1996. BGP Communities Attribute. RFC 1997. https://doi.org/10.17487/RFC1997

[17] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and kc claffy. 2013. AS Relationships, Customer Cones, and Validation. In *Internet Measurement Conference (IMC'13)*. ACM, 243–256. https://doi.org/10.1145/2504730.2504735

[18] NTT. 2016. Deployment of NTT "Peer Locking" route leak prevention mechanism. http://instituut.net/~job/peerlock_manual.pdf

[19] OpenINTEL. 2016. Open Access Datasets. https://data.openintel.nl/data/ Dates used: 2024-03-01 (IPv4), 2024-04-01 (IPv6).

[20] Cristel Pelsser, Randy Bush, Keyur Patel, Prodosh Mohapatra, and Olaf Maennel. 2014. Making Route Flap Damping Usable. RFC 7196. https://doi.org/10.17487/RFC7196

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

[22] RIPE NCC. [n. d.]. Routing Information Service (RIS). https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/

[23] RIPE NCC Staff. 2015. RIPE Atlas: A Global Internet Measurement Network. *The Internet Protocol Journal* 18, 3 (Sept. 2015), 2–26.

[24] Nils Rodday, Lukas Kaltenbach, Italo Cunha, Randy Bush, Ethan Katz-Bassett, Gabi Dreo Rodosek, Thomas C. Schmidt, and Matthias Wählisch. 2021. On the Deployment of Default Routes in Inter-domain Routing. In *Workshop on Technologies, Applications, and Uses of a Responsible Internet (TAURIN '21)*. ACM, 14–20. https://doi.org/10.1145/3472951.3473505

[25] Internet Society. 2023. Internet Resilience Index. https://pulse.internetsociety.org/resilience

[26] Lion Steger, Liming Kuang, Johannes Zirngibl, Georg Carle, and Oliver Gasser. 2023. Target Acquired? Evaluating Target Generation Algorithms for IPv6. In *Network Traffic Measurement and Analysis Conference (TMA'23)*. IEEE, 1–10. https://doi.org/10.23919/TMA58422.2023.10199073

[27] Massimiliano Stucchi. 2023. Italy's Internet Outage a Perfect Storm. https://pulse.internetsociety.org/blog/italys-internet-outage-a-perfect-storm

[28] Roland van Rijswijk-Deij, Mattijs Jonker, Anna Sperotto, and Aiko Pras. 2016. A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements. *IEEE Journal on Selected Areas in Communications* 34, 6 (June 2016), 1877–1888. https://doi.org/10.1109/JSAC.2016.2558918

[29] Curtis Villamizar, Ravi Chandra, and Ramesh Govindan. 1998. BGP Route Flap Damping. RFC 2439. https://doi.org/10.17487/RFC2439

[30] Johannes Zirngibl, Lion Steger, Patrick Sattler, Oliver Gasser, and Georg Carle. 2022. Rusty Clusters? Dusting an IPv6 Research Foundation. In *Internet Measurement Conference (IMC'22)*. ACM, 395–409. https://doi.org/10.1145/3517745.3561440