

Light, Camera, Actions: characterizing the usage of IXPs' action BGP communities

Fabricio Mazzola*
UFRGS, Brazil
fmazzola@inf.ufrgs.br

Pedro Marcos*
FURG, Brazil
pbmarcos@furg.br

Marinho Barcellos
University of Waikato, New Zealand
mbarcell@waikato.ac.nz

ABSTRACT

Border Gateway Protocol (BGP) communities, an optional message attribute, allow network operators to tag BGP announcements and act on routing decisions. Although widely used, it is so far unclear how prevalent the different types of communities are and the degree to which the different traffic engineering actions have been used. There are two major reasons for this gap: few community values have been standardised and the limited visibility at route collectors.

In this paper, leveraging the fact that Internet eXchange Points (IXPs) have sets of well-documented BGP communities, we use their BGP Looking Glasses (LGs) to inspect their route servers and shed light on *how* IXP members are using action BGP communities. During twelve weeks, we collected and analysed routing data from eight IXPs worldwide, focusing the analysis on the four largest IXPs. We observe that *i*) over one-third of IXP members (>35.7%) use action communities in at least one route; *ii*) two-thirds (66.6%) of them are intended to avoid propagating routes, mostly to content providers (CPs); *iii*) nearly one-third (31.8%) are targeting Autonomous Systems (ASes) that are *not* present at the IXPs' route servers, resulting in no practical routing effect and only increasing processing and memory storage overheads.

CCS CONCEPTS

• Networks → Network measurement.

ACM Reference Format:

Fabricio Mazzola, Pedro Marcos, and Marinho Barcellos. 2022. Light, Camera, Actions: characterizing the usage of IXPs' action BGP communities. In *The 18th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '22)*, December 6–9, 2022, Roma, Italy. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3555050.3569143>

1 INTRODUCTION

BGP communities [10] is a variable-length message attribute that can be included in BGP updates to convey some routing information or request. ASes can tag routes with communities to either ask other ASes to perform some action (e.g., path-prepend, blackholing, or do not export) regarding the route or to add some information about a given route characteristic (i.e., learned from a customer) which

*Both authors have contributed equally to the paper.

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

CoNEXT '22, December 6–9, 2022, Roma, Italy

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9508-3/22/12...\$15.00

<https://doi.org/10.1145/3555050.3569143>

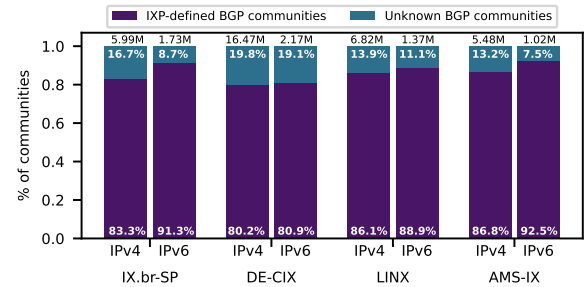


Figure 1: Communities defined by IXP vs. unknown.

can help, for example, improve routing decisions of the AS or its neighbours. While we know that BGP communities are widely used on the Internet, little is known about *how* ASes use them. Earlier studies struggled with the general lack of semantics for all but a tiny fraction of community values, and the poor visibility of *action* communities at route collectors.

In this paper, we use IXPs as vantage points (VPs) to shed light on how ASes use action BGP communities. We leverage IXPs because *i*) they provide a list of available communities and their respective semantic, and *ii*) unlike route collectors, IXP BGP LGs allow us to observe *both* action and informational communities¹. During twelve weeks, we collected and analysed routing data from eight IXPs worldwide. We focused our analysis on the four largest IXPs on the planet (IX.br São Paulo, DE-CIX Frankfurt, LINX London, and AMS-IX Amsterdam) and complemented with insights about other four (DE-CIX Madrid, DE-CIX New York, BCIX Berlin, and Netnod Stockholm). Our study is feasible and relevant because, as shown in Fig. 1, over 80% of the communities we observe, for both IPv4 and IPv6 routes, have a well-defined meaning. Our work is the first to focus on how ASes use action BGP communities.

Our main findings include: action communities are prevalent in IXPs, with over 35.7% (and up to 54.1%) IXP members using action communities in at least one IPv4 route (§5.1); action communities are widely used by ASes, and the number of prefixes they announce and communities they tag are correlated (§5.2); the favourite communities are the ones to avoid another AS (§5.3), and specifically to content providers (§5.4); there exist communities that are ineffective, as they target ASes *not* in the IXP route server (RS) but, surprisingly, are among the most widely used (§5.5); a discussion of implications to operational practices (§5.6).

In addition, we release a twelve-week dataset containing daily snapshots with over 4 billion community instances and a dictionary containing more than 3000 communities, allowing our results to

¹Action BGP communities tend not to be visible at route collectors as ASes usually remove the communities after applying the corresponding action.

Table 1: The eight IXPs in numbers (on the latest snapshot).

IXP	Location	Avg Daily Traffic	# of Members	# of Members at RS		# of Observed Prefixes		# of Observed Routes	
				IPv4	IPv6	IPv4	IPv6	IPv4	IPv6
IX.br-SP	São Paulo, Brazil	9.6 Tbps [33]	2,338 [33]	1,803	1,627	163,981	60,203	282,697	88,652
DE-CIX	Frankfurt, Germany	9.27 Tbps [14]	1,072 [13]	874	711	451,544	65,395	888,478	130,084
LINX	London, United Kingdom	3.8 Tbps [45]	847 [42]	669	508	241,084	62,912	315,215	79,690
AMS-IX	Amsterdam, Netherlands	7.6 Tbps [4]	861 [1]	636	488	252,704	61,528	252,704	61,528
DE-CIX Mad	Madrid, Spain	492 Gbps [18]	214 [17]	151	85	116,237	45,321	125,812	48,711
DE-CIX NYC	New York, USA	941 Gbps [20]	256 [19]	171	145	162,469	48,951	186,983	61,638
BCIX	Berlin, Germany	640 Gbps [6]	145 [5]	88	78	106,249	46,873	111,115	50,569
Netnod	Stockholm, Sweden	1.12 Tbps [48]	187 [47]	127	101	132,179	45,507	150,670	48,874

be fully reproduced and support further research. Next, we briefly revise BGP communities as we position our paper in respect to prior work (§2). Then, we use the dataset we created with routing data collected from IXPs (§3 and §4) to characterise the use of communities in IXPs (§5). We conclude with final remarks (§6).

2 RELATED WORK

The original BGP communities standard [10] defined values for only three communities, essentially providing a way to limit route propagation. In 2008, Donnet and Bonaventure [23] proposed a taxonomy, with three classes: inbound, outbound and blackholing. The inbound referred to tagging announcements with information (e.g. where it was learned), while the outbound referred to communities used for traffic engineering, by influencing route propagation. The third class, blackholing (BH), allowed ASes to drop traffic towards some prefix (as a DDoS defence strategy) [22, 38]. BGP communities can also be grouped in *informational* and *action* communities.

Previous work can be roughly divided into studies on the use/semantics of communities, and their use for measurement studies. In the first group, Dietzel et al. investigated the usage of blackholing in IXPs [21], while Giotsas et al. measured its adoption in the wild [28], and Nawrocki et al. assessed its efficacy against DDoS [46]. While communities are useful for AS' operations, they can also lead to problems. Earlier studies examined its use as a vector of routing attacks [9, 55], and how communities can cause overheads [39]. The (lack of) semantics for community values motivated methodologies for semantics inference [53] and best-effort attempts to build community directories [54]. BGP communities have also been used for inference studies. Examples include finding p2p links at IXP RSes [30], studying RSes in IXPs [52], inferring complex AS relationships [27], mapping peering interconnections to a facility [29], and detecting outages [26].

Comparing our work to prior art, we highlight three studies. In [30], Giotsas et al. collected communities from router server with semantics defined by the IXP in order to infer p2p links. In [52], Philipp et al. examined the role of route servers in IXPs, using communities for some inferences. Krenc et al. [40] observed announcements at BGP collectors (e.g. RIPE and RouteViews) aiming to understand better community usage, but limited to when/how ASes *add* communities to announcements and when they *remove*.

Unlike them, our study aims to clarify how action communities are used by ASes in IXPs. Informational and action communities have different tagging behaviours when comparing IXPs with usual

route propagation. In an IXP, a RS may add informational communities, to be used by ASes. In contrast, action communities will be added by ASes when announcing to the IXP which, in turn, may [37] (and will typically do) perform the action and remove the community from the route. While we cannot observe how the informational communities are going to be used by ASes, we can determine how ASes are using action communities (for traffic engineering). We used this strategy to create a dataset collected from eight IXPs and spanning twelve weeks, with over 4 billion uses of communities in routes, as described next.

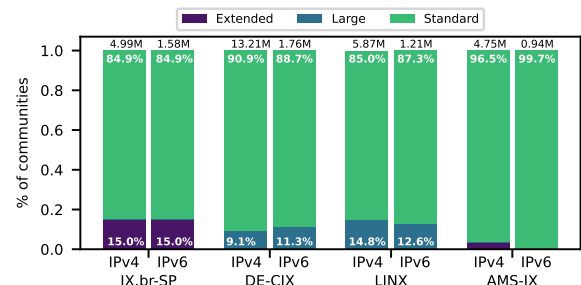


Figure 2: Prevalence of BGP standard communities, while extended and large are relatively small.

3 DATASETS AND SANITATION

We rely on routing data collected from eight IXPs around the world: IX.br Sao Paulo, DE-CIX Frankfurt, LINX London, AMS-IX Amsterdam, BCIX, DE-CIX Madrid, DE-CIX New York, and Netnod Stockholm. To avoid polluting the presentation, we focus the analysis on the first four IXPs in the list (the largest ones) and comment on the other four IXPs whenever relevant. Table 1 presents a summary of key numbers for these IXPs.

Route server data: ASes and routes. Between 19 Jul. 2021 and 4 Oct. 2021 we collected daily snapshots of routing data from the IXP primary IPv4 and IPv6 RSes, using their LG API [2, 8, 16, 35, 43, 49]. Each snapshot consists of a list of member ASes in the RS and a list of routes. We capture the peers with *active* BGP sessions with the RS, regardless whether the AS shares routes or not. The set of ASes in the RS represents, on average, 72.2% and 57.1% of the total IXP members for IPv4 and IPv6, respectively. These ASes were responsible for announcing 111k–888k IPv4 and 48k–130k IPv6 routes (see Table 1) In addition, each snapshot contains a list

of *accepted* routes per AS. The information, captured for every route, includes prefix, next-hop address, AS-Path, and lists of BGP *standard*, *extended*, and *large* communities ².

IXP BGP communities dictionary. To identify the communities defined by IXPs, we build a dictionary combining information from two reliable sources. First, using the LG API, we fetch the RS configuration file containing the semantics of informational and action BGP communities available. To our surprise, when we examined the snapshots, we discovered that this list could be incomplete. We confirm this by comparing the communities in the IXP RSeS with the documentation published at the corresponding IXP website [3, 7, 12, 15, 34, 44, 50]. Therefore, for each IXP, we use the union set between the list from its RS and its website documentation. Our dictionary has 3,183 BGP communities (649 for IX.br-SP, 774 for DE-CIX, 774 for DE-CIX Mad, 774 for DE-CIX NYC, 58 for LINX, 37 for AMS-IX, 50 for BCIX, and 67 for Netnod Stockholm).

Data sanitation. The RS API offers two sets of available routes: *filtered* and *accepted*. Filtered routes are rejected according to rules specified in the router server configuration file. Reasons include bogon prefixes or ASNs, AS paths too long, and prefixes too specific ($>/24$) or too broad ($</8$). These routes are not shared with peers connected to the RS. In contrast, accepted routes are approved by the RS and distributed among ASes via multilateral peering. We focus our analysis on accepted routes, since filtered ones will have no routing impact. For each IXP, we first obtain a summary file with the list of peers, along with the number of routes announced by each peer. Then, for each peer, we collect all the accepted routes. This latter collection process took several hours and was subject to communication failures because of LG instability and/or query rate limits [25]. We inspect all downloaded data and remove from our dataset the snapshots where we found clear “valleys” in the number of members and/or prefixes, i.e. dropped at least 30% from the previous day and returned to previous values in subsequent days. We confirmed with network operators that this behavior was either the result of a failure at the IXP or in our data collection process. The sanitation removed 169 (13.5%) snapshots from our dataset.

Ethics. Our measurements did not involve human subjects. We collect publicly available data from RSeS using the LGs’ API deployed by IXPs. In our automated download process, we kept a single connection to the LG server, to avoiding overloading it.

4 CASTING OUR STARS

We start our analyses by asking the question: how stable are the numbers of BGP communities, in function of ASes, announced prefixes and routes at each IXP RS? We performed two analyses, daily within a week, and weekly for the entire period. Examining the last week, we observed that the variation in the number of communities (as well as ASes, prefixes and routes) in the daily snapshots was under 4%, for all eight IXPs (see Appendix A). Hence, the first snapshot each week (Monday) was used to represent the week. Second, looking at the twelve weekly snapshots, we observed that the median difference between the minimum and the maximum

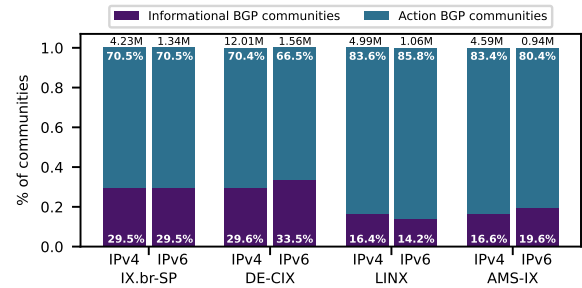


Figure 3: Action vs. informational communities.

values was 5.31%, that is, reasonably stable. For the purpose of the next analyses, we use the most recent snapshot, 4th Oct. 2021. We publish the full dataset and discuss the temporal differences among the snapshots on Appendix A.

Focus on standard BGP communities. We observed that *i*) standard communities (see §2) consistently represented more than 80% of cases in each IXP, and *ii*) even though different types were seen in all IXPs, the usage of *extended* and *large* communities were not consistent among the IXPs. Fig. 2 presents the percentage of IXP BGP communities’ occurrence in routes according to their types for both IPv4 and IPv6. Hence, for the sake of space, we focus now on *standard* communities, leaving the others for future work.

5 BGP COMMUNITIES IN ACTION

We first investigate how often communities are used to perform a traffic engineering action and when they are used to conveying route information (§5.1). Then, we examine whether the BGP communities’ usage is concentrated on a few ASes or more uniformly spread over all the ASes in the IXP RS (§5.2). Next, we break down (§5.3) the action communities into four different categories (i.e., do-not-announce, only-announce-to, prepend-to, blackholing) and analyse which communities are most popular among ASes and which ASes are more targeted (§5.4). We then explore whether all the action BGP communities at IXP routes targeting ASes are indeed necessary, or else could be suppressed to avoid overheads to the RS (§5.5). Finally, we discuss operational implications and possible changes on how ASes are using BGP action communities (§5.6).

5.1 Action is what matters

Our main goal is to understand *how* ASes use BGP communities, i.e. which communities they add when sharing routes with the IXP. An IXP defines sets of action and informational communities, with the informational ones being added by the IXP typically to every route to assist ASes make traffic engineering decisions. While we observe both types in our dataset, inferring how ASes benefit from informational communities is a complex task given the limited visibility of the existing VPs. Thereafter, we focus on action communities.

To confirm the relevance of action communities in the IXP, we analyse the proportion between action and informational communities in the eight IXPs. Fig. 3 shows the results for the four largest ones. We observe that action communities represented at least 66.6% (of the *standard*) communities in each IXP. In the others, we observe that in Netnod Stockholm and BCIX action communities

²We release our dataset and code on <https://github.com/systems-furg/conext-2022-bgp-communities-paper>

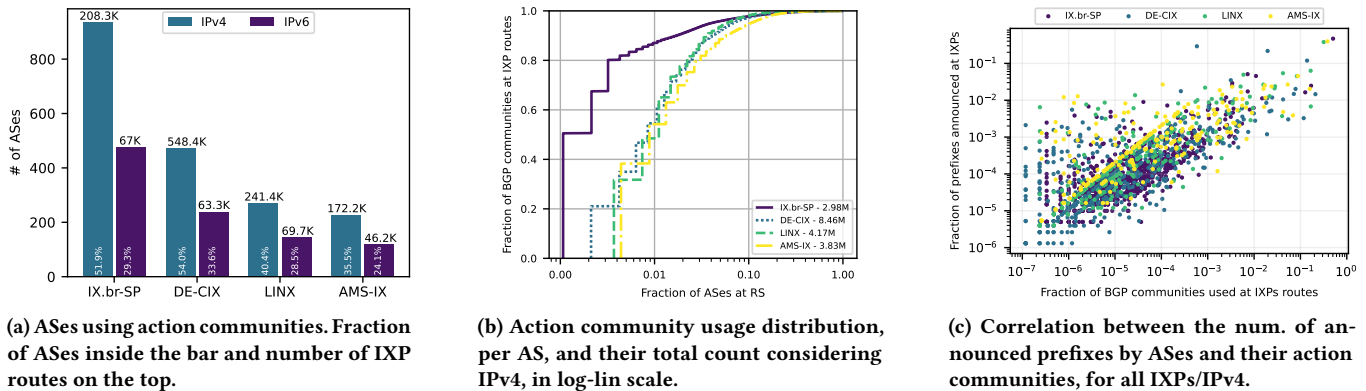


Figure 4: ASes using action communities and corresponding fractions of routes with action communities.

represented more than 95% of the IXP-defined *standard* communities. These observations support our choices for the study.

5.2 Who wants some action?

We now analyse the number of ASes using action communities and the number of routes tagged with at least one action community. We then analyse each AS’s proportion to the set of visible action communities and correlate with their share of routes at the RS.

ASes and routes. In Fig. 4a, each bar represents the number of ASes using action communities in one of the four IXPs. On top of bars, we show the number of routes containing action communities. For this analysis, we consider that an AS uses action communities if at least one of its routes has an action community. Similarly, we consider that a route has an action community if we observe at least one action community. In relative terms, and using data from Table 1, the fraction of members using action communities ranges between 35.5–54% in IPv4 and 24.2%–33.6% in IPv6, with the smallest and largest shares being at AMS-IX and DE-CIX, respectively. When analysing routes, we observe that the percentages are between 61.7% (DE-CIX) and 76.6% (LINX) for IPv4 and 48.7% (DE-CIX) and 85.5% (LINX) for IPv6. We find more significant proportions of routes with action communities than ASes using them, which is expected, as there are some large ASes at IXPs originating thousands of routes. Larger ASes tend to have more complex routing policies, thus likely using more action communities. We detail this next.

Fraction of action communities per AS. We investigate whether the use of action communities is equally distributed among the set of ASes using them. For this, we counted the number of action communities utilised by each AS (i.e., if there are two action communities in a route, we add two to the AS’ count) and divided it by the total number of action communities in all routes of the IXP. We report the results in Fig. 4b. Our findings show that few ASes are accountable for a considerable fraction of the action communities seen at IXP routes. For the three European IXPs, 1% of the ASes on the RS are responsible for around 50–60% of the action communities, while the same AS fraction represents 86% of the action communities at IX.br-SP. On the other extreme, we find that 90% of the ASes are responsible for less than 5% of the communities seen. For IPv6,

the number is even more unbalanced, with less than 1% of the ASes in the RS responsible for 82%–90% of the action communities.

Correlating usage of action communities and share of IXP routes. Lastly, we check if the proportions of announced routes and action communities tagged by an AS are correlated. For each AS, we computed the fraction of routes the AS announces and compared it with the values obtained in the previous analysis. In Fig. 4c, each point represents an AS in a given IXP. We find that, for all IXPs and in most cases, the dots (ASes) are located close to the main diagonal, indicating that they contribute to the number of (IPv4) routes and action communities more or less equally. We also see a group of dots to the upper left but not to the bottom right, indicating that there are large ASes that do not use so many communities, but not small ASes that use many communities. Both observations apply to IPv6 (for all eight IXPs), but with fewer cases.

5.3 ASes’ favorite actions

We now analyse the most common actions ASes request to the IXP RS. First, we assigned each community of our dictionary (see §3) to one of the following four groups: only-announce-to, do-not-announce-to, prepend-to, blackholing. The former two contain BGP communities intended to limit the propagation of a route, either by not exporting it or only exporting it to the specified ASes/regions/facilities, respectively. The latter two are the ones used by ASes to ask the RS to add prepends before propagating the route to the specified ASes/regions/facilities and to blackhole traffic towards the specified prefix, respectively.

ASes using each type of action community. Initially, we look at the popularity of each action BGP community type among ASes. For each IXP, we counted how many ASes used each action community type. We consider that an AS uses a given community type if it tags at least one of its routes with one of the BGP communities that are part of that group. Table 2 presents the results for IPv4 and IPv6, and we comment about IPv4. First, consistently for all IXPs, the most popular type is the do-not-announce-to (27.6%–48.3% depending on the IXP), followed by announce-only-to (6.1%–24.4%). We observe that at DE-CIX the use of blackholing is quite popular (15.7%), in line with the fact that “advanced blackholing” [22] is a said to be a popular IXP service [11]. At the other IXPs, at the time of collection

Table 2: Number and fraction of ASes at the RS using each type of community.

BGP Community	IX.br-SP		DE-CIX		LINX		AMS-IX	
	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6	IPv4	IPv6
Do not announce to	870 (48.3%)	444 (27.3%)	333 (38.1%)	164 (23.1%)	185 (27.6%)	86 (16.9%)	180 (28.3%)	86 (17.6%)
Announce only to	110 (6.1%)	34 (2.1%)	205 (24.4%)	112 (15.7%)	140 (20.9%)	81 (15.9%)	80 (12.6%)	47 (9.6%)
Prepend to	102 (5.7%)	47 (2.9%)	73 (8.3%)	28 (3.9%)	10 (1.5%)	6 (1.2%)	0 (0.0%)	0 (0.0%)
Blackholing	0 (0.0%)	0 (0.0%)	137 (15.7%)	10 (1.4%)	0 (0.0%)	0 (0.0%)	9 (1.4%)	1 (0.2%)

(July–October 2021), we see almost no occurrences of blackholing. The low prevalence of blackholing observed was aligned with the documentation at those IXPs: while IX.br-SP reported *no support* for the blackholing community in 2021 [32], the documentation of AMS-IX and LINX [3, 44] did not mention blackholing. To verify if there were updates, on June 28th 2022, we collected routing data from these two IXPs. There were 1367 and 27 routes with blackholing on AMS-IX and LINX, respectively, which *may* indicate the introduction of support to this community.

Finally, we observe ASes using the prepend-to communities, mostly at IX.br-SP and DE-CIX. On LINX, prepend-to communities were announced on the 29th of June of 2021 [41], only a few weeks before we start capturing data, which can explain the small number of ASes using these communities at LINX. On AMS-IX, prepend-to using *standard communities* is only available when applied to all peers (i.e., prepend to everyone) while *extended communities* need to be used for fine-grained prepending (as we focus on *standard communities*, these are out of the scope).

Number of action communities per type. We also counted the number of times each action community type occurred. We observed that the most popular are do-not-announce-to communities (66.6%–92.0% depending on the IXP), followed by only-announce-to (17.7%–31.4%), then prepend-to (<1.9%), and finally the blackholing community (<0.4%). The corresponding values for IPv6 are, varying among IXPs, shown in the same order: 77.9%–98.5%, 8.4%–22.0%, <0.3% and negligible. We note that while the number of ASes using the blackholing community at DE-CIX is similar to the number of ASes using the prepend-to community, the number of occurrences of each type is significantly different.

5.4 What are ASes tagging about?

We now aim to identify which specific communities appear the most among the IXP route server member’s routes. In each route with a list of communities, we count the occurrences of communities (not types, as in the previous analysis). We highlight the actions most prevalent among routes and the ASes most frequently targeted by them. Fig. 5 shows the top-20 most used action BGP communities.

We observe that, for all IXPs, the most frequent communities are to restrict a route’s access to other peers. On DE-CIX, the most common is the do-not-announce-to community when used to avoid a route’s redistribution to *all* ASes, representing 2.8% of all action communities. For the other three largest IXPs, the top communities have the do-not-announce-to to avoid exporting routes to targeted ASes, Hurricane Electric at IX.br-SP (4.27%), Google at LINX (3.10%), and OVHcloud at AMS-IX (2.83%). For IPv6, the most common targets of do-not-announce-to are Hurricane Electric at IX.br-SP

(5.97%) and LINX (5.87%), OVHcloud at AMS-IX (4.95%), and the European provider Filanco at DE-CIX (4.22%). Apart from the action to avoid, the top 20 contain communities that allow exporting routes to a restricted set of targets. While DE-CIX and LINX have the only-announce-to community to redistribute to *all* as their most common, whereas in IX.br-SP the prevalent use of announce-only-to is to allow redistribution to a few ASes, such as educational networks NIC-Simet, RNP, enterprise Itau, and content provider CDNetworks. In the smaller IXPs, the most popular action communities were to restrict route propagation with a do-not-announce-to community.

Interestingly, there is a considerable intersection among the ASes targeted by action communities in the top 20 of all IXPs. LINX and IX.br, for example, have 14 of the most popular communities aiming to avoid the same ASes. Eight of these networks are large content providers, such as OVHcloud, Google, Akamai, Cloudflare, Netflix, and Edgecast. The intersection for IPv6 is even more prominent, with communities targeting the same 17 ASes at IX.br-SP and AMS-IX. When considering the intersection between the four largest IXPs regarding IPv4, we observe communities to avoid the same six ASes: four content providers (Google, LeaseWeb, Akamai, OVHcloud) and two ISPs (PROLINK, Syntegra Telecom). In IPv6, the intersection between IXPs is again more notable, with nine ASes.

Communities that avoid route redistribution to big content and Internet providers ASes are among the most popular since these networks offer opportunities to exchange large traffic volumes, becoming attractive partners over Private Network Interconnections (PNI) instead of multilateral peering [31]. PNIs are preferred in these scenarios as they can often provide better performance and monitoring over the peering session [24]. When comparing our results with [31], we observe that Google, Akamai, and LeaseWeb still remain among the most avoided ASes at IXPs by means of action communities. Finally, it is revealing to see an increase in the cases of ASes avoiding route redistribution to streaming content providers (e.g. Netflix, Apple) or to cloud providers that host streaming content (e.g. OVHcloud and Edgecast).

5.5 Not all actions are cool

When analysing the frequently targeted ASes, we identify IXP members using many action BGP communities targeting ASes not connected to the IXPs’ RS. We now examine these cases in greater detail in two perspectives. First, we investigate which are the most popular action communities targeting non-IXP RS members. Then, we analyse who are the “culprits”, i.e. ASes tagging their routes with communities targeting non-IXP RS members. Surprisingly, we identify that, for all IXPs, more than 31.8% of the action BGP communities target ASes not connected to the IXPs’ route servers!

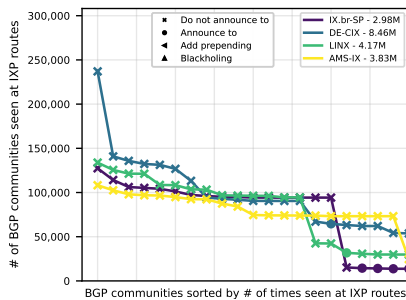


Figure 5: Top-20 most used action communities for each IXP, with total count, for IPv4.

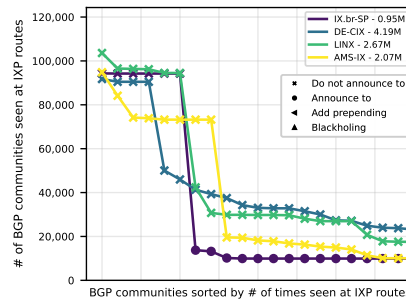


Figure 6: Top-20 non-IXP RS members targeted by action communities, and total count, for IPv4.

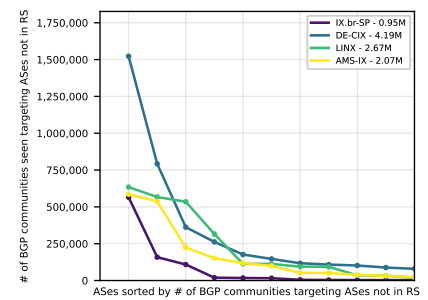


Figure 7: Top-10 ASes targeting action communities to non-IXP RS members for IPv4.

Which are the top action BGP communities targeting non-IXP RS member ASes? This analysis is like §5.4, but considering the smaller subset of action communities targeting non-IXP RS members. Similarly, most action communities targeting non-IXP RS members are aimed at avoiding a few “hot” ASes. In fact, these cases represent some of the most popular ones in the entire set of action communities. As shown in Fig. 6, for IPv4 routes, in IX.br-SP, six action communities targeting non-IXP RS members are part of the top-20 most used action communities in the IXP (see Fig. 5). On DE-CIX, LINX, and AMS-IX we observe four, ten, and eight communities in this condition, respectively. For IPv6, we observe seven, six, nine, and ten action communities targeting non-IXP RS members on the top-20 most popular action communities of IX.br-SP, DE-CIX, LINX, and AMS-IX, respectively. For both IPv4 and IPv6, these communities are mostly targeting content providers.

Surprisingly, we identify large shares of action communities targeting non-IXP RS members. For IPv4, the fractions are 31.8% (0.95M) for IX.br-SP, 49.5% (4.19M) for DE-CIX, 64.3% (2.6M) for LINX, and 54.3% (2.08M) for AMS-IX. For IPv6, we observe 40.3% (0.38M) for IX.br-SP, 40.4% (0.41M) for DE-CIX, 52.6% (0.47M) for LINX, and 45.9.3% (0.35M) for AMS-IX. As the intended target of the actions are *not* members of the RSeS, these BGP communities are achieving no goal other than unnecessary overheads on the RS.

Who is targeting non-IXP RS member ASes? We then analyse which ASes are behind these ineffective action communities. We counted how many times a given AS shares with the IXP RS a route containing an action community that targets a non-IXP RS member. We show in Fig. 7 the top ASes with considering IPv4 routes. We identify that most ASes doing this are large ISPs. Also, we observe ASes acting similarly in multiple IXPs. For example, seven ASes of the Top-10 with more action communities targeting non-IXP RS members are the same on DE-CIX and AMS-IX. Hurricane Electric appears in all IXPs and is responsible for 24.2%-59.4% of these cases in IPv4 routes. We observe a similar behaviour on the other IXPs.

5.6 Operational implications

Intrigued by these findings, we checked our results with network operators and explored two questions. First, *why* operators tag even non-members in the IXP? The practice is to avoid traffic disruptions should a “to-avoid” AS connect to the IXP RS one day. This is in spite of the burden to the RS, which needs to do the filtering. Second, we

discussed *how* this could be avoided and whether an IXP member database would be a solution. While there are databases such as PeeringDB [51] and IXPD [36], they are not updated in real time, which could lead to traffic disruptions. Further, according to an operator, using a database would require constantly processing information “out-of-the-box” (i.e., outside the router), which would increase operational complexity and introduce an additional point of failure. Last, every time an AS on the “to-avoid” list (dis)appeared from the database, the AS would have to send update messages for all of its routes, impacting not only the IXP RS but all its members.

On the IXP side, the RS scrubs the unnecessary BGP communities before propagating the routes, thus limiting any impact to them. A possible alternative to reduce the burden to the RS is for the IXP to filter routes with a large number of communities. DE-CIX, for example, filters routes with “too many communities”. While this approach will not stop ASes from tagging ASes that are not members of the IXP, it creates an incentive for ASes to hygienise their communities.

6 FINAL REMARKS

By leveraging eight IXPs as VPs, and focusing on the the four largest ones, we analysed *how* ASes are using action BGP communities. Overall, we identified that among the standard communities defined by the IXP, the action communities are the majority over the informational ones (at least two-thirds). We find that between one-third and half the members in the eight IXPs use action communities to tag around two-thirds of the IPv4 routes announced at the IXP, and up to 85% of IPv6 routes. We also observed that the most popular action communities are primarily intended to avoid propagating routes to content providers. This practice is likely the result of PNIs between the tagging AS and the content provider. Finally, we identified a large share, between one-third and two-thirds depending on the IXP, of action communities targeting ASes *not present* in the IXP RS. These communities, often intended to avoid ASes, are among the most popular ones. They typically target content providers, and are, in general, added by large ISPs. Discussions with operators revealed that the purpose of these communities is to prevent any disruptions at the operator’s AS should any of these ASes connects to RS. While such practice protects an AS and simplifies operation, it results in unnecessary overheads at the IXP infrastructure.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and our shepherd for their valuable feedback on our paper. We are also very thankful to all network operators for their valuable insights regarding the operational practices about the BGP Communities usage and aspects at IXPs. This study was sponsored (in part) by the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior - Brasil (CAPES) and Grant #2020/05192-9, São Paulo Research Foundation (FAPESP). The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the sponsors.

A STABILITY AMONG SNAPSHOTS

Table 3 is based on the set of seven daily snapshots. It presents, for each of the eight IXPs and for both IPv4 and IPv6, the range of values for members, prefixes, routes and BGP communities. It also shows the percentual difference between the minimum and the maximum values in the set. The highest difference in a week was 3.91%, in the number of communities in AMS-v4.

Table 4, on its turn, shows the corresponding values for the set of twelve snapshots. The highest variation in the three-month period, 18.03%, was observed in the number of communities in DECIXMAD-v4. In most cases, the maximum difference was under 10%, indicating an overall stability for the period.

REFERENCES

- [1] AMS-IX. 2022. AMS-IX Connected networks. <https://www.ams-ix.net/ams/connected-networks>.
- [2] AMS-IX. 2022. AMS-IX Looking Glass. <https://lg.ams-ix.net>
- [3] AMS-IX. 2022. AMS-IX Route Servers. <https://www.ams-ix.net/ams/documentation/ams-ix-route-servers>
- [4] AMS-IX. 2022. AMS-IX Total statistics. <https://www.ams-ix.net/ams/documentation/total-stats>.
- [5] BCIX. 2022. BCIX - BCIX-Members and Peers. <https://www.bcix.de/bcix/members/>.
- [6] BCIX. 2022. BCIX - Traffic Statistics. <https://www.bcix.de/bcix/traffic/>.
- [7] BCIX. 2022. BCIX IXP Manager. <https://www.bcix.de/ixp/content/0/route-servers>
- [8] BCIX. 2022. BCIX Route-Server Looking Glass. <http://lg.bcix.de>
- [9] Henry Birge-Lee, Liang Wang, Jennifer Rexford, and Prateek Mittal. 2019. SICO: Surgical Interception Attacks by Manipulating BGP Communities. In *ACM CCS*. 431–448.
- [10] R Chandra, P Traina, and T Li. 1996. RFC 1997: BGP Communities Attribute. RFC 1997.
- [11] DE-CIX. 2020. DE-CIX - New: Blackholing Advanced and Blackholing Insights. <https://www.de-cix.net/en/about-de-cix/news/new-blackholing-advanced>.
- [12] DE-CIX. 2022. DE-CIX Action BGP Communities. <https://www.de-cix.net/en/resources/service-information/route-server-guides/action-bgp-communities>
- [13] DE-CIX. 2022. DE-CIX Connected networks in Frankfurt. <https://www.de-cix.net/en/locations/frankfurt/connected-networks>.
- [14] DE-CIX. 2022. DE-CIX Frankfurt statistics. <https://www.de-cix.net/en/locations/germany/frankfurt/statistics>.
- [15] DE-CIX. 2022. DE-CIX Informational BGP communities. <https://www.de-cix.net/en/resources/service-information/route-server-guides/informational-bgp-communities>
- [16] DE-CIX. 2022. DE-CIX Looking Glass. <https://lg.de-cix.net>
- [17] DE-CIX Madrid. 2022. DE-CIX Madrid - Connected networks in Madrid. <https://www.de-cix.net/en/locations/madrid/connected-networks>.
- [18] DE-CIX Madrid. 2022. DE-CIX Madrid - Madrid traffic statistics. <https://www.de-cix.net/en/locations/madrid/statistics>.
- [19] DE-CIX New York. 2022. DE-CIX New York - Connected networks in New York. <https://www.de-cix.net/en/locations/new-york/connected-networks>.
- [20] DE-CIX New York. 2022. DE-CIX New York - New York traffic statistics. <https://www.de-cix.net/en/locations/new-york/statistics>.
- [21] Christoph Dietzel, Anja Feldmann, and Thomas King. 2016. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In *PAM Conference*.
- [22] Christoph Dietzel and Matthias Wichtlhuber. 2018. Stellar: Network Attack Mitigation using Advanced Blackholing. In *ACM CoNEXT*.
- [23] Benoit Donnet and Olivier Bonaventure. 2008. On BGP Communities. *ACM CCR* 38, 2 (mar 2008), 55–59.
- [24] Dr. Peering. 2015. Public vs. Private Peering. http://drpeering.net/HTML_IPP/chapters/ch07-0-Public-vs-Private-Peering/ch07-0-Public-vs-Private-Peering.html.
- [25] V. Giotsas, A. Dhamdhere, and k. claffy. 2016. Periscope: Unifying Looking Glass Querying. In *PAM Conference*.
- [26] Vasileios Giotsas, Christoph Dietzel, Georgios Smaragdakis, Anja Feldmann, Arthur Berger, and Emile Aben. 2017. Detecting Peering Infrastructure Outages in the Wild. In *ACM SIGCOMM*. 446–459.
- [27] Vasileios Giotsas, Matthew Luckie, Bradley Huffaker, and kc claffy. 2014. Inferring Complex AS Relationships. In *ACM IMC*. 23–30.
- [28] Vasileios Giotsas, Philipp Richter, Georgios Smaragdakis, Anja Feldmann, Christoph Dietzel, and Arthur Berger. 2017. Inferring BGP blackholing activity in the internet. In *ACM IMC*. 1–14.
- [29] Vasileios Giotsas, Georgios Smaragdakis, Bradley Huffaker, Matthew Luckie, and kc claffy. 2015. Mapping Peering Interconnections to a Facility. In *ACM CoNEXT*. Article 37, 13 pages.
- [30] Vasileios Giotsas, Shi Zhou, Matthew Luckie, and kc claffy. 2013. Inferring Multilateral Peering. In *ACM CoNEXT*. 247–258.
- [31] Vasileios Giotsas, Shi Zhou, Matthew Luckie, and kc claffy. 2013. Inferring Multilateral Peering. In *ACM CoNEXT*. 247–258.
- [32] IX.br. 2021. IX.br - Blackhole no IX.br. <https://forum.ix.br/files/apresentacao/arquivo/1284/Blackhole%20IX%20Forum%2015%20-%20dez-2021.pdf>.
- [33] IX.br. 2022. IX.br - Brazilian Public IXP Project. <https://ix.br/intro>.
- [34] IX.br. 2022. IX.br: BGP Communities Table at IX.br. http://docs.ix.br/doc/communities-table-ix-br-v2_0-22032022.pdf
- [35] IX.br. 2022. IX.br Looking Glass. <https://lg.ix.br>
- [36] IXPDB. 2020. IXPDB. IXP Database. <https://www.ixpdb.net/en/ix-f/ixp-database>.
- [37] Elisa Jasinska, Nick Hilliard, Robert Raszuk, and Niels Bakker. 2016. RFC 7947: Internet Exchange BGP Route Server.
- [38] T. King, C. Dietzel, J. Snijders, G. Doering, and G. Hankins. 2016. RFC 7999: BLACKHOLE Community. RFC 7999.
- [39] Thomas Krenc, Robert Beverly, and Georgios Smaragdakis. 2020. *Keep Your Communities Clean: Exploring the Routing Message Impact of BGP Communities*. 443–450.
- [40] Thomas Krenc, Robert Beverly, and Georgios Smaragdakis. 2021. AS-Level BGP Community Usage Classification. In *ACM IMC*. 577–592.
- [41] LINX. 2021. AS Path Prepending Now Supported on LINX Route Servers. <https://www.linx.net/as-path-prepend-now-supported-on-linx-route-servers/>.
- [42] LINX. 2021. Members by IP/ASN. <https://portal.linx.net/members/members-ip-asn>
- [43] LINX. 2022. LINX Route-Server Looking Glass. <https://alice-rs.linx.net>
- [44] LINX. 2022. Route Servers - LINX Portal. <https://portal-legacy.linx.net/tech-info-help/route-servers>
- [45] LIX. 2022. LINX Statistics. <https://portal.linx.net/services/lans-flow>.
- [46] 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 *ACM IMC*. 435–448.
- [47] Netnod. 2022. Netnod - Connected Networks. <https://www.netnod.se/ix/connected-networks>.
- [48] Netnod. 2022. Netnod - IX Statistics. <https://www.netnod.se/ix/statistics>.
- [49] Netnod. 2022. Netnod Route-Server Looking Glass. <http://lg.netnod.se>
- [50] Netnod. 2022. Route Servers - Netnod. <https://www.netnod.se/ix/route-servers>
- [51] PeeringDB. 2020. PeeringDB. IXPs and colocation database. <https://www.peeringdb.com>.
- [52] Philipp Richter, Georgios Smaragdakis, Anja Feldmann, Nikolaos Chatzis, Jan Boettger, and Walter Willinger. 2014. Peering at Peerings: On the Role of IXP Route Servers. In *ACM IMC*. 31–44.
- [53] Brivaldo A. Silva, Paulo Mol, Osvaldo Fonseca, Italo Cunha, Ronaldo A. Ferreira, and Ethan Katz-Bassett. 2022. Automatic Inference of BGP Location Communities. *ACM SIGMETRICS* 6, 1 (Feb. 2022), 3:1–3:23.
- [54] One Step. 2021. BGP Community Guides. <https://onestep.net/communities/>
- [55] Florian Streibelt, Franziska Lichtblau, Robert Beverly, Anja Feldmann, Cristel Pelsser, Georgios Smaragdakis, and Randy Bush. 2018. BGP Communities: Even More Worms in the Routing Can. In *ACM IMC*. 279–292.

Table 3: Variation in seven daily snapshots (week)

IXP	Members			Prefixes			Routes			Communities		
	Min	Max	Diff%	Min	Max	Diff%	Min	Max	Diff%	Min	Max	Diff%
IX.br-SP-v4	1,723	1,748	1.43%	163,472	164,555	0.66%	281,365	283,049	0.59%	5,095,798	5,141,660	0.89%
IX.br-SP-v6	1,498	1,519	1.38%	58,116	60,198	3.46%	86,113	88,816	3.04%	1,446,707	1,474,654	1.90%
AMS-IX-v4	635	640	0.78%	248,987	252,704	1.47%	248,987	252,704	1.47%	4,905,709	5,105,562	3.91%
AMS-IX-v6	485	491	1.22%	61,528	62,405	1.41%	61,528	62,405	1.41%	1,009,086	1,027,860	1.83%
LINX-v4	632	635	0.47%	244,015	246,531	1.02%	319,556	323,666	1.27%	5,559,238	5,635,387	1.35%
LINX-v6	448	449	0.22%	63,047	63,629	0.91%	79,882	81,007	1.39%	1,123,389	1,138,393	1.32%
DE-CIX-Fra-v4	820	823	0.36%	450,736	458,356	1.66%	883,712	899,377	1.74%	14,760,632	14,885,461	0.84%
DE-CIX-Fra-v6	643	645	0.31%	64,602	65,395	1.21%	128,895	130,084	0.91%	1,893,574	1,906,656	0.69%
BCIX-v4	88	88	0.00%	105,288	106,411	1.06%	110,143	111,432	1.16%	1,651,806	1,673,722	1.31%
BCIX-v6	78	78	0.00%	45,996	46,873	1.87%	49,652	50,569	1.81%	752,573	767,224	1.91%
DE-CIX-NYC-v4	171	174	1.72%	162,004	165,106	1.88%	186,887	191,570	2.44%	2,885,252	2,923,732	1.32%
DE-CIX-NYC-v6	145	146	0.68%	48,305	49,015	1.45%	60,678	61,927	2.02%	1,053,466	1,076,205	2.11%
DE-CIX-Mad-v4	150	152	1.32%	115,788	116,623	0.72%	125,495	126,371	0.69%	2,227,648	2,245,370	0.79%
DE-CIX-Mad-v6	85	86	1.16%	44,549	45,328	1.72%	47,845	48,897	2.15%	758,789	775,848	2.20%
Netnod-v4	126	127	0.79%	132,179	132,197	0.01%	150,670	150,901	0.15%	5,125,979	5,151,156	0.49%
Netnod-v6	100	101	0.99%	45,094	45,507	0.91%	48,463	48,874	0.84%	902,051	908,502	0.71%

Table 4: Variation in twelve weekly snapshots (3 months)

IXP	Members			Prefixes			Routes			Communities		
	Min	Max	Diff%	Min	Max	Diff%	Min	Max	Diff%	Min	Max	Diff%
IX.br-SP-v4	1,652	1,748	5.49%	154,140	164,050	6.04%	241,978	282,697	14.40%	4,327,692	5,141,660	15.83%
IX.br-SP-v6	1,370	1,518	9.75%	57,862	60,203	3.89%	82,486	88,652	6.96%	1,368,582	1,471,665	7.00%
AMS-IX-v4	618	653	5.36%	245,246	265,025	7.46%	245,251	265,030	7.46%	4,929,486	5,206,070	5.31%
AMS-IX-v6	486	495	1.82%	61,187	63,112	3.05%	61,187	63,112	3.05%	955,198	1,032,096	7.45%
LINX-v4	622	640	2.81%	246,014	255,927	3.87%	316,479	329,592	3.98%	5,235,560	5,666,094	7.60%
LINX-v6	427	451	5.32%	59,238	63,734	7.05%	77,319	81,922	5.62%	1,082,610	1,138,393	4.90%
DE-CIX-Fra-v4	815	827	1.45%	444,054	453,847	2.16%	865,946	888,705	2.56%	13,782,937	14,851,619	7.20%
DE-CIX-Fra-v6	635	648	2.01%	62,828	65,395	3.93%	127,234	132,389	3.89%	1,848,666	1,906,656	3.04%
BCIX-v4	85	91	6.59%	98,405	106,351	7.47%	101,719	111,166	8.50%	1,550,217	1,670,622	7.21%
BCIX-v6	76	78	2.56%	45,455	46,873	3.03%	49,236	50,569	2.64%	746,216	767,224	2.74%
DE-CIX-NYC-v4	169	175	3.43%	159,138	164,570	3.30%	175,905	191,097	7.95%	2,604,624	2,915,428	10.66%
DE-CIX-NYC-v6	145	147	1.36%	48,041	51,513	6.74%	59,741	64,033	6.70%	997,500	1,081,904	7.80%
DE-CIX-Mad-v4	148	152	2.63%	103,023	116,237	11.37%	111,125	125,812	11.67%	1,834,093	2,237,424	18.03%
DE-CIX-Mad-v6	81	85	4.71%	43,227	45,321	4.62%	46,214	48,711	5.13%	699,110	773,489	9.62%
Netnod-v4	118	127	7.09%	124,756	132,179	5.62%	142,051	151,081	5.98%	4,853,934	5,151,156	5.77%
Netnod-v6	96	101	4.95%	44,661	45,507	1.86%	47,939	48,874	1.91%	896,846	908,502	1.28%