PeeringDB: Still Accurate a Decade Later?

Radu Anghel
R.Anghel@tudelft.nl
Delft University of Technology
The Netherlands

Carlos Gañán
C.HernandezGanan@tudelft.nl
Delft University of Technology
The Netherlands

Yury Zhauniarovich
Y.Zhauniarovich@tudelft.nl
Delft University of Technology
The Netherlands

Abstract

PeeringDB consists of self-reported data and is widely used to facilitate network interconnection, automation, and research, often serving as a ground truth source. However, the accuracy of its data remains a significant concern. In this study, we evaluate the reliability of PeeringDB using a dataset from January 2025. Our analysis finds certain inconsistencies in key data elements. We identified ASNs that remain present in the database despite no longer being assigned. Analysis of self-reported network types reveals that 28.58% of networks either do not report a type or choose not to disclose it, challenging its use to support interconnection decisions. Further comparison of the available network types against their equivalent classifications from ASdb and IPinfo shows considerable disagreement. Manual validation of networks in the "Government" category reveals that only 77.10% of ASNs provided accurate network type. Finally, the consistency check of self-reported prefix counts with routing data from RIPE RIS and CAIDA ASRank indicated significant inconsistencies. Our findings underscore the need for enhanced validation mechanisms to improve the reliability of data provided by PeeringDB for improving its research and operational uses.

CCS Concepts

• **Networks** → **Network measurement**; *Network management*.

Keywords

PeeringDB, Accuracy, Autonomous Systems, Validation

ACM Reference Format:
Radu Anghel, Carlos Gañán, and Yury Zhauniarovich. 2026. PeeringDB: Still Accurate a Decade Later?. In *The 41st ACM/SIGAPP Symposium on Applied Computing (SAC '26), March 23–27, 2026, Thessaloniki, Greece.* ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3748522.3779749

1 Introduction

PeeringDB [37] is a publicly accessible database operated by a non-profit organization with the same name, containing various self-reported information about networks and interconnection that can be used to support peering decisions and strategies. The database includes content generated by its users (affiliated with Autonomous System Numbers (ASNs) / Organizations) related to Internet Exchange Points (IXPs), data centers, and other interconnection facilities. The stated purpose of this database is to facilitate global interconnection of networks. Additional permitted uses for the database are network troubleshooting, abuse reporting, and Internet research and analysis. The information from this database is widely used in research, with a Google Scholar search yielding more than 70 relevant publications only for the year 2024, distributed across various research domains, including network measurements [31, 25, 45, 27, 11], network security [20, 4], and policy [2].

PeeringDB is often assumed to provide highly accurate data and is frequently treated as a ground truth source in academic research [43, 29, 2]. It is also widely integrated into networking software[1]. A particularly relevant example of its operational influence is its integration with IXP Manager [23], a widely used open-source tool for managing IXPs. IXP Manager can automatically configure information about participating networks using data from PeeringDB, including prefix limits [22]. ARouteServer [42] is an open-source tool that automates the generation and testing of advanced BGP route server configurations, primarily for use at IXPs. ARouteServer can ingest data directly from PeeringDB to automatically configure route server policies based on the self-reported attributes of participating networks. This includes information such as maximum prefix limits and routing policy details. This underscores how PeeringDB's data directly influences research outcomes, network operations, and security practices. Any inaccuracies in the data can propagate through subsequent research, potentially affecting the validity of findings in these domains and may introduce further circular dependencies [49].

Since its inception in 2004, PeeringDB has evolved in multiple ways: it has established itself as a non-profit organization, adjusted data formats, and refined the policies governing data management to adapt to the changing Internet landscape. Still, some ambiguities remain in PeeringDB's data structure. For instance, the 'Network Type' field allows users to choose from ten undocumented options, leading to varying interpretations. A discussion in 2023 [41] resulted in a decision to permit multiple self-reported types rather than removing redundant categories. Similarly, PeeringDB's Data Ownership Policy [39] aims to ensure data consistency, but conflicts between data sources persist. In cases where disputes arise, the policy offers guidance for resolution on certain object types but does not extend to the 'net' object, which represents ASNs. As a result, conflicting information may remain in the database unless modified by the parties involved or resolved through pre-authorized automation or PeeringDB's Admin Committee. Finally, PeeringDB has experienced significant growth, expanding from 1,970 registered ASNs in August 2010 to 31,666 ASNs in January 2025. All these factors indicate a high likelihood of a change in the accuracy of the data stored in this database. Thus, *the accuracy of this self-reported information remains an open question and warrants continuous investigation.*

[1] https://docs.peeringdb.com/tools/

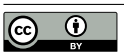

SAC '26, Thessaloniki, Greece

ACM ISBN 979-8-4007-2294-3/2026/03
https://doi.org/10.1145/3748522.3779749

One of the earliest efforts to assess the accuracy of PeeringDB data was conducted in 2014 by Lodhi et al. [28], who found that most networks maintained accurate and up-to-date information. However, the study also identified reduced accuracy of data from ASNs from developing regions. Klöti et al. [24] performed a comparative analysis of several IXP-related datasets, including PeeringDB, Euro-IX, and PCH. They reported that approximately 25% of IXPs listed in PeeringDB contained information that conflicted with the other sources, indicating potential issues with data reliability. More recently, Marder et al. [30] demonstrated that even seemingly verifiable data, such as network presence at public peering exchange points, contains inconsistencies, underscoring the need for external validation mechanisms.

In this paper, we aim to provide a comprehensive and reliable assessment of the data accuracy of specific elements within PeeringDB using its snapshot from January 2025. Ensuring the accuracy of PeeringDB's data is particularly important, as inaccuracies in its records could propagate errors across research studies and operational systems relying on this information. To assess the reliability of the self-reported information in PeeringDB, we analyzed which of its fields can be verified using external sources and performed this comparison. Specifically, as the external data sources we used *NRO Delegated* [35], ASdb [49], *IPinfo* [21], *BGPTools* [7], *RIPE RIS* [44] and *CAIDA ASRank* [9] datasets. Additionally, for some of the properties, we performed manual validation.

Our results challenge the common belief that PeeringDB is a reliable source of data, often considered as the ground truth in research [43, 29, 2]. In reality, the inconsistencies we uncovered paint a more concerning picture. First, we discovered that ASNs (241 as of January 17, 2025) remain listed in PeeringDB despite being no longer assigned. Second, the comparison of types of ASNs, which self-report themselves as belonging to 'Government' and 'Educational/Research', with the types reported by external databases, such as ASdb and IPinfo, revealed that there is a slight agreement between these sources. The manual check of all ASNs self-reported belonging to the 'Government' type, which should not provide space for free interpretation, showed that only 77.10% of networks categorized themselves correctly. Finally, the comparison of prefix counts revealed huge discrepancies with the data obtained from ASRank and derived from RIPE RIS datasets.

Thus, the key contributions of this paper are as follows:

- We analyzed the accuracy of some self-reported information in PeeringDB that may affect the operation of some network tools, e.g., IXP Manager [23], ARouteServer [42], and Peering Manager [19], and may influence research results when this data is considered as ground truth [43, 29, 2].
- We provide insights into the reliability of PeeringDB, uncovering substantial inconsistencies in the self-reported data.
- We analyze the limitations of PeeringDB, offering recommendations for improving its trustworthiness as a source for interconnection data and research applications.

2 Related Work

Researchers frequently use PeeringDB to study the Internet's peering ecosystem, attracted by its public availability and seemingly current information on networks, their presence at IXPs, private peering facilities, and general peering arrangements. However, the accuracy and completeness of this data have been subjects of ongoing investigation and debate. This section briefly traces the evolution of these studies, highlighting their findings and demonstrating the need for continued scrutiny, as undertaken in this paper.

Early work by Lodhi et al. [28] provided an in-depth snapshot analysis of PeeringDB data from August 2013. This study examined business types, traffic volumes, peering strategies, and the presence of networks at various IXPs and private peering facilities. While finding that most networks kept their records updated despite the voluntary nature of participation and the lack of verification mechanisms, the analysis noted lower accuracy for developing regions due to lower membership counts, and questionable IXP-specific properties. However, with PeeringDB reporting 3,392 networks at the time, compared to over 31,666 in January 2025, a 833.55% increase, a significant research gap has emerged due to both the smaller sample size and the dramatic growth of the Internet. Our work seeks to address this gap by revisiting these early findings with current data.

Klöti et al. [24] compared multiple datasets related to IXPs, including PeeringDB, Euro-IX, and PCH. The study highlighted discrepancies between the databases, noting that PeeringDB, in particular, uses separate data sources for IXPs and ASNs participating in those IXPs, which can lead to inconsistencies. Indeed, 25% of IXPs present in PeeringDB had conflicting information between the two sources, implying the existence of data accuracy problems that warrant further examination. These findings motivate the need for improved data validation and reconciliation processes within PeeringDB.

More recent work has directly addressed the data quality concerns within PeeringDB. In [30], PeeringDB was used alongside IXPDB and other sources to infer interconnections between large cloud providers. When encountering discrepancies between the two primary datasets, the authors contacted both PeeringDB and IXPDB. PeeringDB recommended using IXPDB data, as it is updated directly by IXPs, while PeeringDB relies on self-reported and unverified information, a recommendation that highlights a systemic awareness of the challenges within the database. In [10], the authors used PeeringDB and WHOIS information to identify government operated Internet Service Providers, considering PeeringDB to be more accurate than WHOIS.

Further underscoring these challenges, Chen et al. [12] sought to improve the accuracy of CAIDA's AS2ORG sibling relations, using PeeringDB as a reference point. While finding PeeringDB generally more accurate, especially in instances of mergers and acquisitions, the study also noted that correcting discrepancies between the two datasets required manual intervention due to outdated information and lower coverage in PeeringDB. This work provides direct evidence that PeeringDB is not infallible and requires ongoing manual checks, thus reinforcing the need for more research into its accuracy.

Despite these identified limitations, many studies have continued to rely on PeeringDB, often assuming its accuracy. For example, Böttger et al. [8] used PeeringDB to differentiate networks with a large global presence ("hypergiants") from others. The study justified this reliance based on PeeringDB's public nature, popularity, good standing in the network operator community, and the participation of prominent organizations such as Google, Netflix, and

Cloudflare. However, as [8] themselves acknowledge, their assessment of network size and presence relied on self-declared port capacity and traffic profiles, which are inherently subject to potential inaccuracies. This study also highlights the rapid growth of PeeringDB, with 11,596 ASNs participating in January 2018, a 241.9% increase since the paper [28] was published, further emphasizing the need for continuous reevaluation.

Additionally, PeeringDB data, particularly the self-reported network type, were frequently used as a trusted and authoritative source in academic research across several domains, including network measurements [48, 13, 26, 17, 5, 18, 3, 46, 31, 25, 45, 27, 11], network security [47, 33, 15, 16, 20], policy [34, 32, 2], and economics [1]. Researchers relying on this information to classify networks, infer relationships, or contextualize observed behaviors are affected by inaccuracies or inconsistencies of the self-reported data, with results propagating through these studies impacting the validity of the analysis and conclusions.

Finally, PeeringDB remains accepted as a common "ground truth" in academic research, as evidenced by its use in the works of Prehn et al. [43], Loye et al. [29], and Alam et al. [2]. This widespread reliance, combined with the database's rapid evolution and the inherent challenges of self-reported data, necessitates continual reassessment to ensure its continued relevance and accuracy for contemporary Internet research. Our paper contributes to this ongoing effort by providing an up-to-date analysis of PeeringDB data and evaluating the reliability of its key data elements.

3 Methodology

In this section, we describe the methodology we used to check the accuracy of the PeeringDB data. The registration process for PeeringDB networks requires organizations to provide a company name, AS number, and contact information, while additional details, such as Internet Routing Registry (IRR) records, network type, and geographic scope are optional but recommended [38]. Some of this information can be cross-verified, such as the number of IPv4 and IPv6 prefixes, while, other attributes, such as traffic levels, traffic ratios, and geographic scope, are difficult to validate independently. In this work, we focus on the former case using a snapshot of the PeeringDB data from 2025-01-17, which contains the information about 31,666 ASNs, cross-checking the data with the following external information sources: *NRO Delegated* [35], ASdb [49], *IPinfo* [21], BGPTools [7], *RIPE RIS* [44], *CAIDA ASRank* [9].

3.1 External Datasources

NRO Dataset. The Number Resource Organization (NRO) [35] is the coordinating body for the Regional Internet Registries (RIRs) that distribute Internet number resources to their members within their service regions. The NRO publishes multiple datasets to facilitate the research and analysis of IP addresses and ASNs. We use one such dataset, namely the *nro-delegated-stats* from 2025-01-17, which contains data about the region (RIR), status, country of registration of IP addresses and ASNs.

ASdb. ASdb [49] is a dataset that maps Autonomous Systems (identified by ASN) to organizations and classifies them by using an abbreviated version of the North American Industry Classification System (NAICS) for business activities, gathering information from business intelligence databases, website classifiers, and using machine learning algorithms. We include the 2024-01[2] ASdb dataset in our comparison.

IPinfo. IPinfo [21] is a commercial database that provides data about IP addresses, such as geolocation and ASN information. The data is derived from multiple public sources such as Regional Internet Registry information, BGP routing data, and other non-disclosed commercial partnerships. It provides details about Autonomous System Numbers, including the organization's name and type of business. We used the IPinfo IP to ASN dataset from 2025-01-17 to extract the information related to ASNs present in the PeeringDB dataset.

RIPE RIS. The Routing Information Service (RIS) [44] maintained by RIPE NCC provides real-time and historical BGP routing data obtained from many ASes, offering this data to the 26 RIS Route Collectors located at various Internet Exchanges around the globe. Out of 26 Route Collectors, three operate as "multi-hop", meaning that they collect routes not only from networks present at those IXPs but from peers located all over the world. In our research, we downloaded the Multi-Threaded Routing Toolkit (MRT) files from 2025-01-17 (snapshot as of 00:00).

CAIDA ASRank. ASRank [9] is a project of CAIDA that ranks Autonomous Systems (AS) and organizations comprising one or more ASes based on topological data collected by Route Views and RIPE RIS. It ranks ASes by metrics such as the customer cone size, representing the number of direct and indirect customers based on information about the announced IP prefixes. We use the customer cone information available on 2025-01-17 in our comparison with the other data sources.

BGPTools. BGPTools [7] dataset provides information on Autonomous Systems categorization, with specific tags such as "Academic", "Government", and "Home ISP". The "Academic" tag identifies ASes operated by educational institutions, including universities, schools, and national research and education networks (NRENs). The "Government" tag refers to networks managed by national or local governmental bodies, including military or public sector infrastructures, and are typically designed to provide communication systems for public administration and services. The "Home ISP" tag represents networks that deliver residential internet access to a consumer base of at least 50 users, using technologies such as DSL, cable (DOCSIS), or fiber (PON). We used the list of ASNs tagged by BGPTools as "uni", "gov", and "dsl" from 2025-01-17 to compare with PeeringDB.

3.2 Validation Process

Figure 1 shows the PeeringDB attributes, external sources of information, and the steps we follow to validate the data. Before the start of this process, we obtain our primary dataset, PeeringDB, from the CAIDA daily archives for the date 2025-01-17, and extract from it the list of ASNs and the values of their main attributes used during the validation.

Step ①: At this step we validate the status of ASN. According to the PeeringDB Frequently Asked Questions [40] and guidelines for approving net objects [36], unassigned ASNs should not be included

[2]The latest available version at the time of writing.

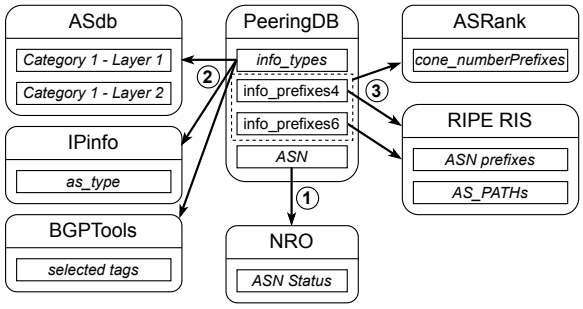

Figure 1: Validation process

in the database. The NRO dataset [35] includes information about *assigned*, *available*, and *reserved* ASNs. If the ASN is unassigned, i.e., have 'available' (not yet assigned) or 'reserved' (not available for registration) status in the NRO dataset, we report this inconsistency. Moreover, for these ASNs we additionally check the values in the PeeringDB fields `rir_status` and `rir_status_updated` that contain the RIR status found by PeeringDB and when PeeringDB checked that status last time respectively.

Step ②: At this step we examine the reliability of self-reported network types, contained in the field `info_types` in PeeringDB. ASNs might skip reporting it or can provide one or multiple (a feature introduced in 2024) type values from 10 possible options. Based on the discussions related to the Network Type definitions [41], the operators of PeeringDB are aware that this data is not accurate and there is uncertainty about who is using this information and for what purposes. Still, no efforts have been made to document what each type means. There were several proposals on how to fix this issue, e.g., provide fewer options to the users or even remove the field. However, as a result of the discussions, this field was kept with all the options and the users were provided with a possibility to choose multiple values.

We compare the values in this field with the data from three external sources: ASdb [49], IPinfo [21] and BGPTools [7]. In the ASdb dataset, we look for the available information related to 'Category 1' - 'Layer 1' and 'Layer 2' for each ASN found in PeeringDB. This information will provide a basis for comparison of the self-declared PeeringDB network types versus the information provided by ASdb. Although the ASdb information could be considered outdated since it was last published in January 2024, we believe that even in the case of mergers and acquisitions, the network type would not change. The cases of deregistration or not yet assigned resources could be determined using the `rir_status` field or the status from the PeeringDB and NRO datasets. From the IPinfo database, we use the value in the `as_type` field and from BGPTools we use the tags representing `Government`, `Academic` and `Home ISP` networks.

Step ③: PeeringDB allows networks to self-declare the "Recommended maximum number of IPv4/6 routes/prefixes to be configured on peering sessions for this ASN". When correctly implemented by both the reporting network and its peers, this measure enhances the security of BGP sessions by mitigating some of the risks associated with large-scale route leaks and prefix hijacks. However, existing literature on the validation of PeeringDB data does not assess the accuracy of these self-reported prefix limits. To address this gap, we compared the self-declared values with the customer cone size (number of prefixes) from CAIDA's ASRank dataset [9] and with the value derived using the information from RIPE RIS [44].

For each ASN in CAIDA's ASRank dataset, we collect the customer cone size data (number of prefixes), contained in the attribute *cone_numberPrefixes*. Note that in the ASRank dataset, the customer cone size shows the total number of IPv4 and IPv6 prefixes, which are reported separately in PeeringDB (in the fields `info_prefixes4` and `info_prefixes6` correspondingly). Therefore, to perform the comparison, we at first get the total number of both IPv4 and IPv6 prefixes in PeeringDB and then compare the obtained value with the one from ASRank.

To obtain the number of prefixes from RIPE RIS data, we at first count the number of unique prefixes originated by the ASN, and then we add to this value the number of prefixes that this ASN transits – the ASN's customer cone – from the perspective of the RIPE RIS peers. To get this last number, we consider all AS_PATHs containing this ASN (only if the ASN is not the RIS peer providing the information – to prevent direct RIS peers from appearing to have large customer cones), and count the number of the corresponding unique prefixes (counting the number of prefixes that this ASN transits).

We analyze inconsistencies in the number of prefixes reported by PeeringDB (self-reported data) and obtained from CAIDA ASRank and RIPE RIS data, looking for differences of more than 20%. This threshold is choosen because the same value is used in some network tools, e.g., IXP Manager [22]. Therefore, the discovered inconsistencies may potentially affect the results produced by these tools.

4 Findings

Since its inception in 2004, PeeringDB has experienced significant growth. Figure 2 shows the number of registered ASNs in PeeringDB from August 2010 to January 2025. As we can see, this number has been continuously growing, expanding from 1,970 registered ASNs in August 2010 to 31,666 ASNs in January 2025. Given the substantial growth and the fact that PeeringDB consists of self-reported user-generated content, the evaluation of its accuracy is essential. In this section, we perform this evaluation following the methodology described in Section 3 using the PeeringDB data snapshot collected as of January 17, 2025, unless otherwise specified.

Examining the geographic distribution of networks registered in PeeringDB across the RIRs, we observe that LACNIC and AFRINIC have the highest levels of participation, with 50.81% and 49.04% of their assigned ASNs appearing in PeeringDB, respectively, with ARIN showing the lowest participation rate at just 16.18%, as shown in Table 1. Although this method does not account for Out-Of-Region ASNs (e.g., RIPE NCC-registered ASNs operating in the United States), such cases represent only about 4% of all ASNs assigned by RIPE NCC and an even smaller share for the other RIRs, allowing us to treat this as a reasonable margin of error.

For a distribution per country, we found that Brazil (BR), the United States (US), and Indonesia (ID) have the highest number

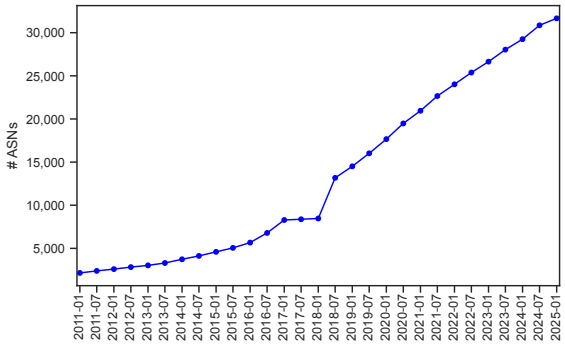

Figure 2: Growth of registered ASNs in PeeringDB

Table 1: ASNs per RIR Present in PeeringDB

RIR	# of Assigned ASNs in NRO	# of ASNs in PeeringDB	%
LACNIC	13,665	6,943	50.81%
AFRINIC	2,506	1,229	49.04%
RIPENCC	39,110	11,159	28.53%
APNIC	29,941	6,757	22.57%
ARIN	32,983	5,336	16.18%
Overall coverage:	118,205	31,666	26.79%

of ASNs present in PeeringDB. Figure 3a presents the per-country distribution of PeeringDB networks based on the number of ASNs. At the same time, if we calculate the percentage of ASNs registered in PeeringDB to the total number of ASNs in a country, the picture changes (see Figure 3b) with some smaller countries with fewer registered ASNs achieving 100% presence.

4.1 Analysis of Unassigned ASNs

As we described in Step ① in Section 3.2, PeeringDB should not contain information about unassigned ASNs. Therefore, using the NRO dataset, we analyzed whether PeeringDB contains unassigned ASNs. Our analysis identifies 241 ASNs that are listed as "reserved" or "available" in the NRO dataset, indicating that they were likely returned to the RIRs and are no longer active. Notably, five of these 241 ASNs have their `rir_status` field in PeeringDB marked as *"ok"* with a `rir_status_updated` timestamp of 2024-06-26. This means that the last check of their status was performed almost half a year ago and at that time it was ok. We also found the presence of a special purpose ASN in PeeringDB – AS112 [14]. This AS, used to soak up leaking DNS queries trying to reverse lookup private IP addresses, is included into PeeringDB due to the collaboration with the Domain Name System Operations, Analysis, and Research Center (DNS-OARC), which aids in coordinating the AS112 project and maintaining its operational stability.

Note that PeeringDB should not contain private-use ASNs. We additionally analyzed PeeringDB for the presence of such ASNs. Expectedly, the most recent dataset analyzed in this study (2025-01-17) does not contain any private-use ASNs. However, historical records (e.g., from 2011-01-14) reveal the presence of this type of ASNs, such as AS65535 indicating that validation of registered ASNs was not performed at the time.

4.2 Analysis of Network Types

As we discussed in Step ② in Section 3.2, PeeringDB provides participating networks a possibility to self-declare their network type: ASNs might skip reporting it, or can provide one or multiple values out of 10 possible options. Out of the total 31,666 ASNs present in the analyzed PeeringDB dataset 9,050 (28.58%) do not declare a network type, 22,026 (69.56%) declare one network type, and only 590 (1.86%) declare more than one type (see Table 2).

Table 2: PeeringDB Network Types counts

# of Types	# of ASNs	% of Total
No type	9,050	28.58%
1	22,026	69.56%
>1	590	1.86%
Total:	31,666	100.00%

Out of the latter, 131 networks self-report themselves belonging to more than three network types, 14 networks – to more than five, and 6 ASNs – to nine or ten of the available options. The records of all of these 6 ASNs have been updated in 2024, so the provided information can be assumed to be fresh. Manual check of those ASNs revealed them to be hosting providers. Manual verification of 14 networks with more than five types showed that those are either also hosting providers or their website is not available (their PeeringDB records were also updated in the past six months). Table 3 shows the number of networks belonging to each type.[3]

Table 3: PeeringDB Network Types

Network Type	# of ASNs	% of Total
Cable/DSL/ISP	11,456	34.95%
no info	9,050	27.62%
NSP	3,830	11.69%
Content	2,413	7.37%
Enterprise	1,929	5.89%
Educational/Research	1,509	4.61%
Network Services	1,064	3.25%
Non-Profit	687	2.10%
Route Server	652	1.99%
Government	131	0.40%
Route Collector	42	0.13%
Total:	32,763	100.00%

Even though multiple options can be selected when self-reporting the network type, some of the available types should not be combined, for example a Route Server (at an IXP) cannot also be a Cable/DSL/ISP or Content network. The cases where such combinations appear are, at best, mistakes caused by misunderstanding the meaning of the undocumented network types or, worse, intentionally misleading. We have identified 28 ASNs with such combinations (all of them are assigned according to the NRO dataset).

[3]Note that the total number (32,763) is higher than the number of ASNs in PeeringDB (31,666) because an ASN may belong to multiple types.

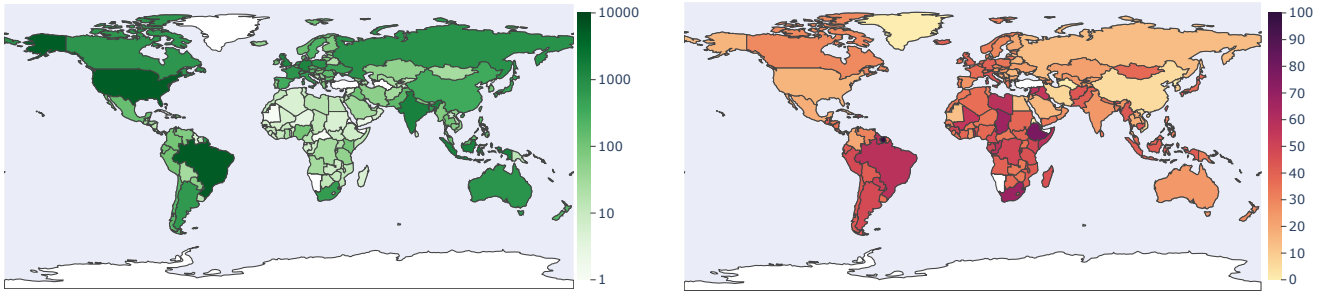

(a) Number per country **(b) % per country**

Figure 3: Geographical distribution of ASNs registered in PeeringDB

Below we analyze ASNs belonging to "Government", "Education/Research" and "Cable/DSL/ISP" network types. For each category, as the first step, we check how many ASNs are assigned by checking their status in the NRO dataset.

Network Types: Government. This category comprises only a small subset of all PeeringDB ASNs (131), with all these networks being assigned according to the NRO dataset (see Table 4). The criteria for a network inclusion into this category should be well-defined, leaving little room for interpretation. However, PeeringDB does not provide guidance about the meaning of the network types that can be selected by operators. For the purpose of this comparison we will define "Government" networks as networks managed by national or local governmental bodies, similarly to the definition used by BGPTools [7].

To assess the accuracy of this self-reported classification, we compared it with three external datasets: ASdb, IPinfo, and BGPTools. Specifically, we matched PeeringDB's "Government" network type with ASdb's Category 1-Layer 1 "Government and Public Administration", IPinfo's `as_type` "government" and BGPTools "gov" tag. Additionally, we conducted manual verification to validate our findings.

Table 4: Network Types: Government

Comparison	# of ASNs	%
PDB Gov	131	100.00%
PDB vs ASdb	52	39.69%
PDB vs IPInfo	35	26.72%
PDB vs BT	9	6.87%
PDB vs ASdb vs IPInfo	18	13.74%
PDB vs ASdb vs IPInfo vs BT	5	3.82%
Manual check	101	77.10%

As shown in Table 4, only 52 (39.69%) ASNs in ASdb are identified as "Government and Public Administration". IPinfo classifies as "government" an even smaller subset of just 35 (26.72%) networks. BGPTools only identifies 9 (6.87%) self-reported government networks. When cross-referencing the PeeringDB with ASdb and IPinfo data sources, we found only 18 (13.74%) ASNs in common, with even lower numbers when including BGPTools. However, by manually checking the websites and registration details of ASNs, we determined that 101 out of 131 ASNs could reasonably be classified as "Government", showing an accuracy of 77.10% that is higher than in ASdb, IPinfo, and in the two combined.

Of the misclassified as Government networks, we found hosting providers and ISPs not owned by the government that could be providing services to government entities, suggesting a misunderstanding of the meaning of the Government network type. We also found two networks appearing as intentionally misleading. These mistakes propagate and grow in importance when they are automatically included in other datasets that trust PeeringDB to be accurate.

Table 5: Countries of Network Types: Government

NRO Country	# of ASNs	%
US	25	19.08%
BR	19	14.50%
ID	17	12.98%
AR	7	5.34%
CA	7	5.34%
Other	56	42.75%

Geographically, the majority of ASNs self-reporting as "Government" originate from the United States (US), Brazil (BR), Indonesia (ID), Canada (CA), and Argentina (AR), as shown in Table 5.

Table 6: Countries with matches vs ASdb and IPInfo: Government

(a) With positive matches

NRO Country	# of ASNs	%
US	7	38.89%
BR	4	22.22%
Other	7	38.89%

(b) With negative matches

NRO Country	# of ASNs	%
US	18	15.93%
ID	17	15.04%
BR	15	13.27%
CA	7	6.19%
AR	7	6.19%
Other	49	43.36%

The ranking of countries based on the geographical distribution of negative matches between PeeringDB networks with the "Government" network type and classifications from ASdb and IPInfo[4] shown in Table 6b compared to the matches in Table 6a do not allow us to draw any conclusions about regional trends of mislabeling due to the low number of networks of this type.

Network Type: Educational/Research. This category includes 1,509 ASNs as shown in Table 3. This number includes 50 (3.31%) ASNs that are not currently assigned, which we exclude from further consideration, leaving 1,459 ASNs for comparison (see Table 7). Similarly to the "Government" network type, we compare the self-reported value with the ASdb, IPinfo and BGPTools datasets. For this comparison, we consider the PeeringDB `info_type` "Educational/Research" to be equivalent to the ASdb Category 1 Layer 1 of "Education and Research", IPinfo as_type of "education" and BGPTools "uni" tag. We do not perform manual verification of the results because the number of ASNs to consider is significantly higher and the criteria of inclusion into this category are more vague. For instance, this network type is frequently chosen by "hobby networks" [6]. This example shows that there is a lack of clear documentation related to the network types in PeeringDB.

Table 7: Network Types: Educational/Research

Comparison	# of ASNs	%
PDB Edu	1,509	100.00%
PDB vs NRO	50	3.31%
PDB vs ASdb	530	35.12%
PDB vs IPInfo	389	25.78%
PDB vs BT	305	20.21%
PDB vs ASdb vs IPInfo	301	19.95%
PDB vs ASdb vs IPInfo vs BT	207	13.72%

In Table 7, we can see that only 301 (19.95%) out of 1,509 ASNs belong to this category according to all three sources. When comparing PeeringDB to ASdb, the agreement is for 530 (35.12%) ASNs, PeeringDB vs IPinfo agrees about 389 (25.78%), and PeeringDB vs BGPTools agree about 305 (20.21%) networks. The relative numbers are similar to those obtained for the "Government" category.

Table 8: Countries of Network Types: Educational/Research

NRO Country	# of ASNs	%
US	317	21.01%
DE	131	8.68%
CN	103	6.83%
GB	89	5.90%
ID	68	4.51%
Other	801	53.08%

Table 8 presents the list of the countries with the highest number of ASNs self-reported as belonging to the Educational/Research network type.

The geographical distribution of networks that match the PeeringDB "Educational/Research" network type in ASdb and IPInfo (Table 9a), compared with the distribution of negative matches (Table 9b), reveals a trend of potential mislabeling in several countries, notably the United States (US), Germany (DE), China (CN), the United Kingdom (GB), and the Netherlands (NL).

[4]We did not consider BGPTools for this comparison because we found the dataset to be incomplete.

Table 9: Countries with matches vs ASdb and IPInfo: Educational/Research

(a) With positive matches

NRO Country	# of ASNs	%
US	94	31.23%
ID	35	11.63%
AR	20	6.64%
BR	12	3.99%
PL	11	3.65%
Other	129	42.86%

(b) With negative matches

NRO Country	# of ASNs	%
US	223	19.26%
DE	122	10.54%
CN	101	8.72%
GB	87	7.51%
NL	64	5.53%
Other	561	48.45%

Network Type: Cable/DSL/ISP. The networks of the type "Cable/DSL/ISP" would benefit the most from peering with other networks. Therefore, not surprisingly, this category is the largest in PeeringDB, with 11,456 (34.95%) networks self-reporting as belonging to it (see Table 3). This number includes 25 (0.39%) ASNs that are currently not assigned according to the NRO dataset, which we exclude from further consideration, giving us 11,411 ASNs for further check (see Table 10). Then, we perform the comparison similarly as for the other two network types described above. We consider the PeeringDB's "Cable/DSL/ISP" network type to be equivalent to the ASdb's Category 1-Layer 2[5] of "Internet Service Provider (ISP)", IPinfo `as_type` of "isp", and BGPTools "dsl" tag.

Table 10: Network types: Cable/DSL/ISP

Comparison	# of ASNs	%
PDB C/D/ISP	11,456	100.00%
PDB vs NRO	45	0.39%
PDB vs ASdb	8,656	75.56%
PDB vs IPInfo	9,397	82.03%
PDB vs BT	1,369	11.95%
PDB vs ASdb vs IPInfo	7,289	63.63%
PDB vs ASdb vs IPInfo vs BT	1,084	9.46%

The results of the comparison are presented in Table 10, where we can see that from the 11,456 ASNs, an agreement between the PeeringDB, ASdb and IPinfo datasets occurs in only 7,289 (63.63%) cases. The percent grows to 82.03% when comparing only PeeringDB to IPinfo and significantly drops to 11.95% when compared to BGPTools, suggesting the BGPTools dataset is incomplete.

Table 11 lists the countries with the highest number of ASNs self-reported as belonging to the Cable/DSL/ISP network type.

A similar comparison of matches (Table 12a) and negative matches (Table 12b) between PeeringDB, ASdb, and IPInfo for the "Cable/DSL/ISP" network type shows that the country rankings remain relatively similar, suggesting a more evenly distributed pattern of potential mislabeling, with no single country standing out.

[5]Layer 1 type is not specific enough for this case.

Table 11: Countries of Network types: Cable/DSL/ISP

NRO Country	# of ASNs	%
BR	3,003	26.21%
US	997	8.70%
IN	888	7.75%
ID	740	6.46%
PL	504	4.40%
Other	5,324	46.47%

Table 12: Countries with matches vs ASdb and IPInfo: Cable/DSL/ISP

(a) With positive matches

NRO Country	# of ASNs	%
BR	2,499	34.28%
US	581	7.97%
IN	443	6.08%
PL	397	5.45%
ID	328	4.50%
Other	3,041	41.72%

(b) With negative matches

NRO Country	# of ASNs	%
BR	504	12.23%
IN	445	10.80%
US	416	10.09%
ID	412	10.00%
JP	157	3.81%
Other	2,188	53.08%

4.3 Analysis of IPv4/IPv6 Prefix Count

In Step ③ in Section 3.2, we describe the approach on how we validate the number of prefixes in PeeringDB. To conduct our comparison, we started with 31,666 ASNs listed in PeeringDB. We exclude from consideration networks that have self-reported their type as "Route Collector" (as they are supposed to passively collect routing information without actively announcing routes) and "Route Server" (as these are Internet Exchange Points (IXPs) that facilitate route exchange between IXP members; while they may handle a large number of routes, they are not visible in the Global Routing Table), and also the ones that belong to both these types. Thus, we excluded 31 "Route Collector", 641 "Route Server" and 11 of networks that belong to both these network types.[6] We then compared the rest ASNs against CAIDA ASRank and RIPE RIS.

Comparison with CAIDA ASRank. As we already mentioned in Section 3.2, CAIDA ASRank does not provide separate customer cone prefix counts for IPv4 and IPv6. Therefore, we sum PeeringDB's `info_prefixes4` and `info_prefixes6` values to compare against ASRank's `cone_numberPrefixes`.

Table 13: PeeringDB self-declared prefix counts compared to CAIDA ASRank

Difference	# of ASNs	%
>0%	22,411	100.00%
>20%	15,365	68.56%
>50%	14,839	66.21%
>200%	11,809	52.69%

The results, summarized in Table 13, indicate significant inconsistencies. With a tolerance of 20%, we find that 15,365 ASNs in PeeringDB have incorrect self-reported prefix counts. Even with a considerably larger tolerance threshold of 200%, the inconsistencies persist, with 11,809 ASNs (52.69% of ASNs with non-zero prefix values) still having inaccurate values. While a small difference (e.g., 20%) could be considered reasonable to account for network growth or small prefix changes without requiring frequent updates, the differences exceeding 200% diminish the security benefits of setting prefix limits for BGP sessions, as they fail to provide an effective safeguard against route leaks or prefix hijacks.

Comparison with RIPE RIS. Using the approach described in Section 3.2, we can analyze separately IPv4 and IPv6 prefix counts. Since the number of ASNs with a non-zero self-reported prefix count is different for each address type, we obtain two sets of ASNs to compare with RIPE RIS, specifically, 21,467 with non-zero prefixes for IPv4 and 18,626 ASNs for IPv6 correspondingly.

Table 14: PeeringDB self-declared prefix counts compared to RIPE RIS

(a) For IPv4

Difference	# of ASNs	%
>0%	21,467	100.00%
>20%	12,873	59.97%
>50%	12,638	58.87%
>200%	9,728	45.32%

(b) For IPv6

Difference	# of ASNs	%
>0%	18,626	100.00%
>20%	14,289	76.72%
>50%	14,534	78.03%
>200%	12,495	67.08%

The comparison between PeeringDB and RIPE RIS (see Table 14a) shows that 59.97% of ASNs report an IPv4 prefix count differing by more than 20%. Furthermore, 45.32% (9,728 ASNs) show differences greater than 200% between self-reported prefix counts compared to RIPE RIS.

For IPv6, the comparison between PeeringDB's `info_prefixes6` and RIPE RIS's IPv6 prefix counts indicates even larger inconsistencies (see Table 14b). A total of 76.72% of ASNs report inaccurate prefix numbers with a difference of more than 20%, while 67.08% self-report values that differ by more than 200%. This comparison further indicates that inaccuracies in self-reported prefix counts are more prevalent for IPv6 than for IPv4.

5 Discussion

Our analysis reveals that there are systemic challenges in maintaining data accuracy within the community-maintained PeeringDB. The presence of unassigned ASNs and inconsistent network type reporting (28.58% missing/undisclosed types) demonstrates fundamental limitations in the current self-reporting mechanisms. These findings have significant implications for both operational network management and Internet governance research that relies on PeeringDB as a primary data source which we discuss next.

Operational and Research Impacts. The propagation of inaccurate data through dependent systems creates cascading reliability issues. For IXPs using PeeringDB for autoconfiguring the prefix-limit of their members [22], even the currently used 20% safety margin [22] may prove insufficient when combined with self-reporting inaccuracies. Our finding that 22.90% of "Government"-classified networks were misidentified illustrates how research analyzing

[6] This gives us 652 "Route Server" and 42 "Route Collector" networks reported in Table 3.

critical infrastructure resilience could draw erroneous conclusions about government-operated networks' peering strategies.

The more than 200% discrepancy between self-reported prefixes and routing-table observations found in 45.32% IPv4 and 67.08% IPv6 networks suggests potential vulnerabilities to route leaks if automation systems rely on PeeringDB data. This aligns with previous findings about IXP presence inaccuracies [30] and reveals a broader pattern of data-quality issues affecting multiple types of information provided by PeeringDB.

Maintainer Incentives and Constraints. PeeringDB's user-driven governance model creates tensions between data availability and accuracy. While the Data Ownership Policy [39] emphasizes user accountability, our findings demonstrate practical limitations in resolving conflicts between data sources, especially for "net" objects. Users maintaining their data face structural challenges:

- **Validation asymmetry**: Automated checks work for verifiable data (such as ASN registration status with a RIR) but lack technical feasibility for subjective fields (Network Type).
- **Participation tradeoffs**: Strict validation could reduce voluntary participation, while leniency permits inaccuracies.
- **Resource constraints**: Manual verification at scale proves impractical given the growing number of ASN entries.

These limitations create data quality issues where database growth (913% increase since 2013) outpaces manual or automatic quality control mechanisms. Rather than enforce strict categorization, the decision to permit users to select multiple network types [41] reflects this tension between usability and accuracy.

Systemic Limitations. Our validation approach using multiple independent academic and commercial sources shows fundamental limitations of the self-reported data:

- **Stale data**: entries showed >6 months since the last update despite RIR status changes (deregistration).
- **Interpretative variance**: Network Type ambiguity leads to inconsistent classification vs. ASdb/IPinfo.

When PeeringDB data is used in research contexts, where it often serves as a source of ground truth, these limitations and inaccuracies can generate incorrect results. The resulting inaccuracies propagate in further research, potentially distorting our understanding of the evolution of Internet topology.

Recommendations. In order to address these challenges, some structural improvements need to be implemented by PeeringDB:

- Develop consensus over what information is useful for the purpose of the database. Network operators should work together to define which data fields are essential and what each field is intended to represent. This process will help ensure that all users interpret and supply information in a consistent and meaningful way, reducing ambiguity and making the database more reliable for both operational and research purposes.
- Implement automatic validation of the self-reported data where possible. Where feasible, automate the cross-checking of self-reported entries with authoritative external sources. For example, regularly compare AS numbers, network types, and prefix counts with reliable registries and routing data. This will help quickly identify and flag inconsistencies, prompting users to review and update their information as needed.
- Periodically perform manual validation of the self-reported data or display more visibly that the data is not in any way validated. For fields that cannot be automatically verified, schedule regular manual reviews—especially for critical or frequently misclassified categories. If validation is not possible, prominently indicate which data is unverified so that users and researchers can take this into account in their analyses and decisions.
- Implement stale-data checks for entries unchanged for longer periods. Regularly identify records that have not been updated in a long time, and prompt the responsible users to confirm or update their entries. Displaying the date of last modification and allowing users to filter by data freshness will further help highlight potentially outdated information.
- Publish guidelines and clear definitions for the Network Type options. Provide comprehensive documentation and standardized definitions for all selectable options, especially for fields like "Network Type." This will reduce the risk of misclassification and ensure that users report their network attributes in a consistent and comparable manner across the database.

While these measures would increase maintenance complexity, they could preserve or increase PeeringDB's utility as "the go-to location for interconnection data". Our findings emphasize that databases based on user-generated content require ongoing validation frameworks rather than static accuracy assumptions, particularly when used for critical infrastructure decisions and research.

6 Conclusions

In this paper, we examined the accuracy of self-reported data in PeeringDB – a publicly accessible database created to facilitate interconnection decisions and strategies. While no dataset is entirely perfectly accurate, our analysis reveals a concerning trend in PeeringDB: the data quality has declined over time. One key factor contributing to this decline is the rapid expansion of the database, with the number of listed networks growing from 1,970 in August 2010 to 31,666 as of January 2025. This increase is an inherent challenge that cannot be mitigated.

However, other factors can and should be addressed. PeeringDB could improve its accuracy by increasing the frequency of automated checks (e.g., updating RIR status) and incorporating objective data from external sources (e.g., prefix counts). Clear, standardized guidelines would also help enhance the quality of information in certain fields, such as network type classification. Additionally, developing independent validation mechanisms for currently unverifiable data – such as traffic volumes, private peering presence, and facility existence – would further strengthen reliability. Without addressing these limitations, PeeringDB data should be used with caution when making interconnection decisions or treating it as ground truth.

References

[1] Md Ibrahim Ibne Alam, Koushik Kar, and Elliot Anshelevich. 2021. Balancing traffic flow efficiency with ixp revenue in internet peering. In *2021 IEEE Global Communications Conference (GLOBECOM)*. IEEE, 1–6.

[2] Md Ibrahim Ibne Alam, Anindo Mahmood, Prasun K Dey, Murat Yuksel, and Koushik Kar. 2024. Meta-peering: automating isp peering decision process. *IEEE Transactions on Network and Service Management*.

[3] Scott Anderson, Loqman Salamatian, Zachary S Bischof, Alberto Dainotti, and Paul Barford. 2022. Igdb: connecting the physical and logical layers of the internet. In *Proceedings of the 22nd ACM Internet Measurement Conference*, 433–448.

[4] Radu Anghel, Yury Zhauniarovich, and Carlos Gañán. 2024. Who's got my back? measuring the adoption of an internet-wide BGP RTBH service. *ACM on Measurement and Analysis of Computing Systems*. ACM POMACS 8, 1, Article 3, (Feb. 2024). doi:10.1145/3639029.

[5] Augusto Arturi, Esteban Carisimo, and Fabián E Bustamante. 2023. As 2 org+: enriching as-to-organization mappings with peeringdb. In *International Conference on Passive and Active Network Measurement*. Springer, 400–428.

[6] benjojo. 2024. Peeringdb feature request: add new network type "personal". PeeringDB, (2024). Retrieved 2025-03-01 from https://github.com/peeringdb/peeringdb/issues/1537.

[7] BGP Tools. 2025. Bgp tools. BGP Tools, (2025). Retrieved 2025-03-01 from https://bgp.tools.

[8] Timm Böttger, Felix Cuadrado, and Steve Uhlig. 2018. Looking for hypergiants in peeringdb. *ACM SIGCOMM Computer Communication Review*, 48, 3, 13–19.

[9] CAIDA. 2025. Caida asrank. CAIDA, (2025). Retrieved 2025-03-01 from https://asrank.caida.org/.

[10] Esteban Carisimo, Alexander Gamero-Garrido, Alex C Snoeren, and Alberto Dainotti. 2021. Identifying ases of state-owned internet operators. In *Proceedings of the 21st ACM Internet Measurement Conference*, 687–702.

[11] Esteban Carisimo, Rashna Kumar, Caleb J Wang, Santiago Klein, and Fabián E Bustamante. 2024. Ten years of the venezuelan crisis-an internet perspective. In *Proceedings of the ACM SIGCOMM 2024 Conference*, 521–539.

[12] Zhiyi Chen, Zachary S Bischof, Cecilia Testart, and Alberto Dainotti. 2023. Improving the inference of sibling autonomous systems. In *International Conference on Passive and Active Network Measurement*. Springer, 345–372.

[13] Ovidiu Dan, Vaibhav Parikh, and Brian D Davison. 2021. Ip geolocation using traceroute location propagation and ip range location interpolation. In *Companion Proceedings of the Web Conference 2021*, 332–338.

[14] DNS-OARC. 2025. As112. DNS-OARC, (2025). Retrieved 2025-03-01 from https://www.as112.net.

[15] Ben Du, Cecilia Testart, Romain Fontugne, Gautam Akiwate, Alex C Snoeren, and Kc Claffy. 2022. Mind your manrs: measuring the manrs ecosystem. In *Proceedings of the 22nd ACM Internet Measurement Conference*, 716–729.

[16] Jan Marius Evang and Ioana Livadariu. 2023. How large is the gap? exploring manrs and iso27001 security management. In *2023 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*. IEEE, 1–6.

[17] Talaya Farasat and Akmal Khan. 2021. Detecting and analyzing border gateway protocol blackholing activity. *International Journal of Network Management*, 31, 4, e2143.

[18] Dimitrios P Giakatos, Sofia Kostoglou, Pavlos Sermpezis, and Athena Vakali. 2022. Benchmarking graph neural networks for internet routing data. In *Proceedings of the 1st International Workshop on Graph Neural Networking*, 1–6.

[19] Guillaume Mazoyer. 2025. Peering manager. (2025). Retrieved 2025-03-01 from https://peering-manager.net/.

[20] Thomas Holterbach, Thomas Alfroy, Amreesh Phokeer, Alberto Dainotti, and Cristel Pelsser. 2024. A system to detect {forged-origin}{bgp} hijacks. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*, 1751–1770.

[21] IPinfo. 2025. Ipinfo - the trusted source for ip data. IPinfo, (2025). Retrieved 2025-03-01 from https://www.ipinfo.io/about.

[22] IXPManager. 2020. Ixp manager tutorial 10 - managing members. INEX, (2020). Retrieved 2025-03-01 from https://www.youtube.com/watch?v=QxltZ1I6lGk.

[23] IXPManager. 2025. Ixp manager: free & open source software platform for ixps. INEX, (2025). Retrieved 2025-03-01 from https://www.ixpmanager.org/.

[24] Rowan Klöti, Bernhard Ager, Vasileios Kotronis, George Nomikos, and Xenofontas Dimitropoulos. 2016. A comparative look into public ixp datasets. *ACM SIGCOMM Computer Communication Review*, 46, 1, 21–29.

[25] Sun Letong, Shi Xingang, Han Fengyan, Yin Xia, Wang Zhiliang, and Zhang Han. 2024. Is it a real cd mismatch in interdomain routing? *arXiv preprint arXiv:2401.11520*.

[26] Jie Li, Shi Zhou, and Vasileios Giotsas. 2021. Performance analysis of multipath bgp. In *IEEE INFOCOM 2021-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 1–6.

[27] Ioana Livadariu, Kevin Vermeulen, Maxime Mouchet, and Vasilis Giotsas. 2024. Geofeeds: revolutionizing ip geolocation or illusionary promises? *Proceedings of the ACM on Networking*, 2, CoNEXT3, 1–21.

[28] Aemen Lodhi, Natalie Larson, Amogh Dhamdhere, Constantine Dovrolis, and Kc Claffy. 2014. Using peeringdb to understand the peering ecosystem. *ACM SIGCOMM Computer Communication Review*, 44, 2, 20–27.

[29] Justin Loye, Sandrine Mouysset, and Katia Jaffrès-Runser. 2023. Emergence of nestedness in the public internet peering ecosystem. *arXiv preprint arXiv:2312.09544*.

[30] Alexander Marder, Kimberly C Claffy, and Alex C Snoeren. 2021. Inferring cloud interconnections: validation, geolocation, and routing behavior. In *Passive and Active Measurement: 22nd International Conference, PAM 2021, Virtual Event, March 29–April 1, 2021, Proceedings 22*. Springer, 230–246.

[31] Fabricio Mazzola, Augusto Setti, Pedro Marcos, and Marinho Barcellos. 2024. Analyzing remote peering deployment and its implications for internet routing. *IEEE/ACM Transactions on Networking*.

[32] Shahzeb Mustafa, Prasun Kanti Dey, and Murat Yuksel. 2022. Peer me maybe?: a data-centric approach to isp peer selection. In *NOMS 2022-2022 IEEE/IFIP Network Operations and Management Symposium*. IEEE, 1–9.

[33] Marcin Nawrocki, Raphael Hiesgen, Thomas C Schmidt, and Matthias Wählisch. 2021. Quicsand: quantifying quic reconnaissance scans and dos flooding events. In *Proceedings of the 21st ACM internet measurement conference*, 283–291.

[34] Ali Nikkhah and Scott Jordan. 2023. Should large isps apply the same settlement-free peering policies to both isps and cdns? In *2023 IEEE 20th Consumer Communications & Networking Conference (CCNC)*. IEEE, 384–391.

[35] NRO. 2025. Number resource organization (nro). NRO, (2025). Retrieved 2025-03-01 from https://nro.net/.

[36] PeeringDB. 2025. Admin committee guidelines and criteria for approving networks, ixps, and facilities. PeeringDB, (2025). Retrieved 2025-03-01 from https://docs.peeringdb.com/committee/admin/approval-guidelines/.

[37] PeeringDB. 2025. Peeringdb. PeeringDB, (2025). Retrieved 2025-03-01 from https://www.peeringdb.com/.

[38] PeeringDB. 2025. Peeringdb howto. PeeringDB, (2025). Retrieved 2025-03-01 from https://docs.peeringdb.com/howtos/.

[39] PeeringDB. 2020. Peeringdb: data ownership policy. PeeringDB, (2020). Retrieved 2025-03-01 from https://docs.peeringdb.com/gov/misc/2020-04-06_PeeringDB_Data_Ownership_Policy_Document_v1.0.pdf.

[40] PeeringDB. 2025. Peeringdb: frequently asked questions. PeeringDB, (2025). Retrieved 2025-03-01 from https://docs.peeringdb.com/faq/.

[41] PeeringDB. 2023. Peeringdb: network type – your input sought. PeeringDB, (2023). Retrieved 2025-03-01 from https://docs.peeringdb.com/blog/network_type_your_input_sought/.

[42] Pier Carlo Chiodi. 2025. Arouteserver. (2025). Retrieved 2025-03-01 from https://github.com/pierky/arouteserver.

[43] Lars Prehn, Franziska Lichtblau, Christoph Dietzel, and Anja Feldmann. 2022. Peering only? analyzing the reachability benefits of joining large ixps today. In *International Conference on Passive and Active Network Measurement*. Springer, 338–366.

[44] RIPE NCC. 2025. Routing information service (ris). RIPE NCC, (2025). Retrieved 2025-03-01 from https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/.

[45] Loqman Salamatian, Calvin Ardi, Vasileios Giotsas, Matt Calder, Ethan Katz-Bassett, and Todd Arnold. 2024. What's in the dataset? unboxing the apnic per as user population dataset. In *Proceedings of the 2024 ACM on Internet Measurement Conference*, 165–182.

[46] Pavlos Sermpezis, Lars Prehn, Sofia Kostoglou, Marcel Flores, Athena Vakali, and Emile Aben. 2023. Bias in internet measurement platforms. In *2023 7th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 1–10.

[47] Lumin Shi, Jun Li, Devkishen Sisodia, Mingwei Zhang, Alberto Dainotti, and Peter Reiher. 2023. Ddos mitigation dilemma exposed: a two-wave attack with collateral damage of millions. In *International Conference on Security and Privacy in Communication Systems*. Springer, 25–44.

[48] Lion Steger, Liming Kuang, Johannes Zirngibl, Georg Carle, and Oliver Gasser. 2023. Target acquired? evaluating target generation algorithms for ipv6. In *2023 7th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE, 1–10.

[49] 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.