



Delft University of Technology

Document Version

Final published version

Licence

CC BY

Citation (APA)

Albakour, T., Osali, F., Franke, M., & Smaragdakis, G. (2026). Detecting IPv4 Subnets in the Wild. In S. Krishnan, & T. C. Schmidt (Eds.), *ANRW 2026 - Proceedings of the 2026 ACM/IRTf Applied Networking Research Workshop* (pp. 82-88). (ANRW 2026 - Proceedings of the 2026 ACM/IRTf Applied Networking Research Workshop). Association for Computing Machinery (ACM). <https://doi.org/10.1145/3822163.3827935>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Detecting IPv4 Subnets in the Wild

Taha Albakour
Max Planck Institute for
Informatics
talbakou@mpi-inf.mpg.de

Fariba Osali*
Max Planck Institute for
Informatics
fosali@mpi-inf.mpg.de

Max Franke
Technische Universität
Berlin
m.franke@tu-berlin.de

Georgios Smaragdakis
Delft University of
Technology
g.smaragdakis@tudelft.nl

Abstract

In this paper, we present methodologies based on two Internet Control Message Protocol (ICMP) message types to measure subnet deployments on the Internet. First, we exploit a peculiar behavior in certain router implementations, together with variations in the interpretation of protocol specifications, to infer prefix boundaries and thus prefix lengths through remotely probing the Internet with specially crafted ICMP Echo requests. Second, we evaluate the extent to which hosts continue to respond to the deprecated ICMP Address Mask request. We assess the applicability of these methods across devices from different network vendors and discuss protocol quirks that arise from lenient interpretations of the specifications. By combining both methods, we evaluate the consistency of our approach and present a dataset comprising 3.8M subnets. This dataset covers 15% of BGP announced prefixes across more than 20k autonomous systems. Using this dataset, we examine the uniformity of host responsiveness within autonomous systems and across subnet sizes, finding that smaller subnets tend to exhibit greater uniformity. Ultimately, our work demonstrates the feasibility of subnet inference, at a reasonable Internet scale.

CCS Concepts

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

Keywords

Internet Measurement; Measurement Techniques.

ACM Reference Format:

Taha Albakour, Fariba Osali, Max Franke, and Georgios Smaragdakis. 2026. Detecting IPv4 Subnets in the Wild. In *ACM/IRTF Applied Networking Research Workshop (ANRW '26)*, July 20, 2026, Vienna, Austria. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3822163.3827935>

1 Introduction

Understanding how Internet resources are used and deployed is crucial for studying the Internet structure and improving measurement methodologies. While prior efforts examined various aspects of the Internet, e.g., reachability [4, 8, 16], service deployments [5, 24, 47], and network paths [6, 11, 32, 37] and their characteristics [12, 22, 23], subnet deployments remain opaque.

Subnets are subdivisions of a larger network prefix. They are defined by boundaries: the network and broadcast addresses [33].

Subnetting practices allow for the deployment of subnets smaller than the recommended BGP announcement size of /24, and up to /32. Measuring subnets on the Internet is challenging, however, without visibility into network configurations, motivating the need for indirect methods based on observable protocol behavior.

Knowledge of subnet deployments has broad potential applications in both research and operations. From a research perspective, it improves our understanding of IPv4 resource utilization and offers a new perspective on Internet topology. For example, identifying /30 and /31 subnets—which are typically used for point-to-point links—enriches topology inference [20, 43]. Similarly, traceroutes toward more specific network prefixes can increase path discoverability [45]. From an Internet-scanning perspective, topological scanning reduce traffic load [29]. This reduction can be significant, especially as prefixes tend to host similar services [35] or indicate the existence of specific ones [25]. From an operational perspective, detecting subnets can improve the accuracy of abuse handling and blocklisting by eliminating unnecessary over-blocking when a potentially malicious address is detected within a large prefix. It can also provide insight into Internet Service Providers' (ISPs) network deployments, enabling third parties (e.g., CDNs) to better adapt their policies. In fact, RFC 9977 proposes publishing prefix-length deployment information, similar to geofeeds, whereby operators would publish prefix-length files via the WHOIS database [17]. A notable challenge for such a solution, however, is trust, as there is currently no mechanism to verify its correctness.

In this paper, we introduce two methodologies based on Internet Control Message Protocol (ICMP) message types to measure subnet deployments on the Internet. Our main method leverages a peculiar behavior in routers implementation. We determine subnet boundaries by remotely probing the network and broadcast addresses with a specially crafted ICMP Echo requests. While ICMP Echo has been used in measurement studies for decades, our work highlights nuances of source address selection in ICMP Echo and exploits behavioral variation that arise from differing interpretations of the standard. In addition, we employ the deprecated ICMP Address Mask request, in which responding hosts disclose the subnet mask, thereby identifying the division between the network and host portions of an IP address. Utilizing this request complement our approach and offers insights into the status of this obsolete ICMP messages. Our contributions can be summarized as follows:

- We present a unique ICMP Echo-based methodology that utilizes the network stack implementation in a popular network vendor to detect subnet deployments. By probing the Internet with specially crafted ICMP Echo requests, we identify 2.05M subnets.
- We assess the extent to which Internet hosts respond to the deprecated ICMP Address Mask request and use these responses to complement our first approach, identifying an additional 1.8M subnets.

* Also affiliated with Saarland University.



This work is licensed under a Creative Commons Attribution 4.0 International License. ANRW '26, Vienna, Austria

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2876-1/2026/07

<https://doi.org/10.1145/3822163.3827935>

- We validate the consistency of our observations and examine vendor-specific behavior using a fingerprinting technique. We highlight variation in the applicability of our method against different router vendors.
- We compile the first-of-its-kind dataset of subnet deployments using our methodology and examine its coverage across BGP-announced prefixes and autonomous systems, as well as the consistency of host behavior.
- To support future research, we will continuously run these measurements and make our tools and dataset publicly available at <https://github.com/talbakour/echos-subnet>.

2 Background

The Internet Control Message Protocol (ICMP), a protocol for sending diagnostic and error messages, specifies several types of control messages for Internet Protocol version 4 (IPv4) [33]. In this section, we provide background and historical context for the relevant ICMP message types.

2.1 ICMP Message Types

We begin by briefly introducing the two relevant ICMP message types to our work: *ICMP Echo* and *Address Mask*.

Echo. ICMP Echo message (ICMP type 8) and Echo Reply (ICMP type 0), are commonly used to verify hosts reachability [38]. A host receiving data in an Echo request returns the exact data in the reply.

Address Mask. ICMP Address Mask request (ICMP type 17) and reply (ICMP type 18) are originally specified in RFC 950 in 1985 as part of the initial introduction of subnetting [33]. These messages are designed for use on Local Area Networks (LANs), enabling a host to determine the network address mask (commonly referred to as the netmask) of its local network. The request is intended to be broadcast by a host that has just joined the network. A gateway that receives the request would respond with a reply. The reply includes the network's address mask. Both the request and the reply use the same format. However, the Address Mask value in the request is all zeros, whereas the reply contains the actual network mask.

2.2 Interpretation of Standards

ICMP Echo messages have been part of the Internet Protocol Suite since their initial specification by Jon Postel in RFC 792 in 1981 [38]. Since then, the specification has been steadily updated and modified, most recently by RFC 6918 in 2013 [19]. Because the original specification is over 40 years old—predating the introduction of normative keywords in BCP 14 (1997) [7]—it contains some ambiguities. One of them relates to how Echo replies should be constructed. The specification states: *“The address of the source in an echo message will be the destination of the echo reply message. To form an echo reply message, the source and destination addresses are simply reversed, the type code changed to 0, and the checksum recomputed”*. This statement appears to be interpreted as a recommendation rather than a *strict* requirement, as hosts can reply from a source other than the intended destination. Further, the behavior of special addresses, such as network and broadcast addresses, is unspecified and therefore remains, to this day, implementation-dependent. Notably, deviating from the lenient standard can have practical operational use cases. For example, it may ensure reachability or

provide a stable host identity. Indeed, popular vendors provide configuration for such use cases. They allow the configuration of the custom source address, e.g., the loopback address [10, 26, 27]. As we show in Section 4, many hosts on the Internet send Echo replies with a source address different from the original Echo request's destination. It is worth noting, however, that such mismatches can also arise from various policies and network settings, including load balancers, multi-homed egress selection, anycast endpoints, and address aliasing. In this work, we focus particularly mismatches caused by popular router vendor's handling of requests targeting network and broadcast addresses.

2.3 Address Mask Deprecation Status

The functionality of address mask messages was later superseded by mechanisms such as DHCP and CIDR, making the use of ICMP to learn a netmask obsolete. Accordingly, this ICMP message type, along with several others, was deprecated by RFC 6918 in 2013 [19]. It should be noted, however, that this RFC only specifies changes to the IANA registry entries for the relevant types (i.e., 17 and 18), marking them as deprecated. The RFC does not specify how devices should handle these ICMP messages, such as dropping them. Thus, responding to Address Mask requests remains standard-compliant and depends on the implementation. Despite the deprecation and vendors disabling Address Mask reply by default, many hosts still respond to Address Mask requests, enabling *remote* detection of subnet deployments.

3 Detecting Subnets on the Internet

In this section, we discuss our approaches to detecting subnets on the Internet. We begin by presenting an ICMP Echo-based technique for identifying subnet boundaries. Additionally, we complement this approach by examining deployment of the deprecated ICMP Address Mask, further expanding our dataset coverage.

3.1 Finding Boundaries with ICMP Echos

Network and broadcast addresses are two special IPv4 addresses that are part of every IPv4 subnet up to /30. The network address is the first address in a prefix (or subnet), identifying the network itself, whereas the broadcast address is the last address in a prefix and is used to send data to all devices within a local subnet simultaneously. These addresses define the boundaries of a subnet and cannot be assigned to hosts. The behavior of probing such addresses, i.e., by sending an Echo request from an external network, is not well-defined and can vary across vendors and implementations. We develop our subnet boundary detection technique by examining how devices from different vendors behave when receiving packets targeting network and broadcast addresses in a lab environment.

Lab Testing. To better understand the *default* handling of network and broadcast addresses across major vendors, we test Cisco 2901 and Juniper MX150 series devices, as well as a virtualized MikroTik RouterOS instance. Note that the Juniper operating system (JunOS) is based on FreeBSD, while MikroTiks RouterOS is based on Linux. Both, therefore, potentially inherit certain behaviors from their respective underlying open-source operating systems. In contrast, the Cisco device runs an IOS image, which is a proprietary OS. For each device, we configure subnets of varying sizes and probe both




Figure 1: An illustration of vantage point impact on subnet discovery.

the network and broadcast addresses with ICMP Echo requests, recording the observed behavior. On Juniper, the device responds to ICMP Echo requests sent to the network or broadcast address by using that address as the source of its Echo reply. In contrast, MikroTik simply ignores the request and sends no reply. For Cisco devices, the reply source is the IP address of the interface on which the request is received, effectively revealing the subnet boundaries. We refer to the source address of the reply as the *Responder Address*. This behavior appears to stem from ambiguity in the RFC specification, as the destination address of the request and the source address of the reply do not match. Nevertheless, it provides a means of measuring subnet boundaries when a subnet is reachable via a Cisco router.

Due to our limited access to hardware, we consider only three implementations. As such, our observations may not capture all possible behaviors, even within a single vendor, as implementations can vary across platforms (e.g., Cisco IOS, IOS-XR, etc.) and software versions. Nevertheless, prior work shows that Cisco is the most popular router vendor on the Internet [1, 2], potentially indicating a relatively wide applicability of our approach.

From Testing to Measurement. Our method relies on the fact that routers differ in how they handle Echo requests toward network and broadcast addresses. We exploit implementation differences, particularly in Cisco devices, and attempt to identify hosts with responsive network and broadcast addresses based on our lab observations. To detect a mismatch between the request destination and the reply source, we craft an ICMP Echo request in which we encode the destination address in the payload. We enumerate the entire routable IPv4 address space using ZMap [15]. We use the `icmp_echo_time` module, as it already encodes the destination address in each request payload. For each reply, we compare the source address with the echoed payload.

Filtering Responses. If the encoded destination address in the payload does not match the source address of the reply, we mark the encoded address as a potential subnet boundary. We refer to the source address of such a reply as a *Responder Address*, which may belong to a different subnet from the original destination. We calculate the subnet size by examining the original destination addresses encoded in each *Responder* Echo Reply, where any mismatch suggests a potential subnet boundary. We only consider a candidate subnet valid if both its network and broadcast addresses are present in our dataset. We discard a subnet if any additional address within the candidate range triggers a mismatch response from a different address, as this implies the presence of a usable host rather than a boundary. We also discard addresses that lack a matching boundary pair or pairs that do not correspond to a valid

Table 1: ICMP Echo probing overview: the number of Echo Replies, Responder Address (source of mismatched request destination/reply source), and subnets, per Vantage Point (VP).

VP	Echo Replies	Responder Addr.	Detected Subnets
US	382M	1.05M	1.79M
DE-1	355M	1.01M	1.54M
DE-2	384M	990k	1.6M
AU	343M	850k	1.52M
Union	396M	1.16M	2.05M

subnet. Lastly, we keep the most specific non-overlapping subnet from the remaining set. These filters, while conservative, ensure that we retain only likely boundary pairs, reducing false inferences caused by responses from non-boundary addresses.

Vantage Point Consideration. To maximize our coverage and investigate the impact of vantage points (VPs) on subnet discovery, we run ICMP Echo request measurements concurrently from multiple vantage points. We use four vantage points: one in the United States (US), two in Germany (DE-1 & DE-2), and one in Australia (AU), deployed across research networks and two cloud providers. Depending on routing and egress routers, multiple vantage points can increase the coverage of detected subnets. Figure 1 shows how different vantage points can affect subnet discovery. Probing subnet `192.168.0.0/30` from VP 2 reaches it via a Juniper router. In this case, the Echo Reply has a source address of `192.168.0.0`. In contrast, probing the same network address from VP 1 elicits a response from `192.168.1.1`. Using multiple vantage points increases paths and router diversity and, thus the number of detected subnets. All vantage points use public IPv4 addresses and are directly accessible from the Internet, without firewalls or NATs. Such configurations could otherwise block responses if the source address differs from the original destination.

3.2 Probing for Address Mask

Conceptually, detecting subnets using Address Mask requests is straightforward: a host either replies, thereby revealing its subnet, or simply drops the request. We scan the Internet for responsive hosts and infer the corresponding subnet sizes from the returned masks (see Section 2). Similar to our ICMP Echo-based approach, we probe the entire routable IPv4 address space using ZMap. We implement a custom ZMap module that supports sending Address Mask requests. Unlike the Echo request measurement, we run this experiment from a single vantage point, as multiple vantage points do not directly affect our methodology. Nevertheless, we repeat the measurement twice to ensure consistency.

4 Active Measurement Results

In this section, we provide an overview of our measurement results. We discuss overall responsiveness to both ICMP Echo and ICMP Address Mask requests. Further, we discuss the number of detected subnets and the impact of vantage points on the ICMP Echo-based technique. We run all measurements in November 2025, following best practices as described in [14, 36].

ICMP Echo Replies. In Table 1, we provide an overview of the ICMP Echo measurement results. We exclude ICMP error messages because they are not relevant to our analysis. When comparing the responses across vantage points, we observe non-negligible variation in the total number of ICMP Echo replies received by each. DE-2 reports the highest number of responses, with more than 384M replies. Similarly, we observe a comparable number of responses from the US VP. In contrast, the remaining VPs report slightly lower counts, with 355M and 343M replies for DE-1 and AU, respectively. For the majority of replies across all VPs, the source address matches the encoded address in the payload. However, depending on the vantage point, between 3.9M and 5.2M replies have a source address that does not match the address in the payload. We refer to the corresponding source IPs as “Responder Addresses” and observe a total of 1.16M unique such addresses. Note that a single *Responder Address* can be observed in multiple mismatches. Despite the large difference in total replies between DE-1 and the other VPs, the number of Responder Addresses is similar, with the US VP recording the highest count. This set of replies serves as candidates for our subnet detection pipeline. To ensure data quality, we apply the filters described in Section 3.1, discarding 44% (511k) of the *Responder Addresses*. Since each responder may map to multiple probe destinations, this also eliminates their corresponding original destination addresses, resulting in the removal of 1.5M addresses across all vantage points. In total, we detect 2.05M subnets, with the US vantage point contributing the largest share at 1.79M subnets.

Vantage Point Impact. Next, we investigate how the vantage point affects subnet discovery using the ICMP Echo-based method. Our methodology relies on reaching a subnet via a specific ingress point that shows the desired behavior. In Figure 2, we visualize the overlap between the sets of subnets discovered from each vantage point, and the dot-and-line matrix encodes which combination of vantage points contributes to each intersection. The vertical bars show the intersection sizes. Surprisingly, only 55% of subnets are discovered by all four vantage points simultaneously. Furthermore, 5.9%, 4.4%, 2.3%, and 1.5% are visible only from US, DE-1, DE-2, and AU, respectively. We detect the remaining subnets using either two or three vantage points. In other words, if we used only AU as our vantage point, we would miss 26.09% of the subnets detected when combining all vantage points. Similarly, using only the US vantage point results in missing 12% of the subnets detected across all vantage points. Overall, these findings highlight the importance of using multiple vantage points for similar measurements.

ICMP Address Mask. Despite its deprecated status, we find a non-negligible number of responsive hosts to ICMP Address Mask requests. We conduct two measurements within a week of each other in November 2025. We identify 2.47M responsive hosts in the first measurement, and 2.46M in the second. Of these, 2.31M IPs remain stable across both scans. In total, we receive 72,534 invalid responses, with 72,260 containing no netmask and 274 containing an invalid netmask. Overall, these translate to 1.82 M additional detected subnets.




Figure 2: Upset plot of subnets discovered per vantage point using the ICMP Echo-based method.

5 Vendor Patterns and Validation

Our limited lab test (see Section 3.1) reveals that the ICMP Echo approach is only applicable to Cisco devices. To confirm this in the wild, we attempt to fingerprint the vendor by scanning the *Responder Addresses* for SNMPv3, as described in [1]. Additionally, we attempt to fingerprint hosts that are responsive to Address Mask requests to explore whether certain vendors are more likely to respond to such probes.

ICMP Echo Responder Addresses. Overall, we identify the vendor for 289k *Responder Addresses*, representing 44% of the total. Indeed, we find that Cisco is the most prevalent vendor, accounting for 98.2% of those addresses. However, we also observe the presence of Huawei and Adtran, with a combined share of less than 1.03%. The remaining addresses, totaling 2.5k, correspond to a long tail of 47 vendors. These results not only support the observation from our lab experiment but also provide indirect validation of our findings. Note, however, that the SNMPv3 approach provides only a vendor fingerprint, with no additional information about the operating system or target platform. Consequently, we cannot verify whether the ICMP Echo approach applies universally to all Cisco devices.

Address Mask. We identify the vendor for 686k responsive addresses, representing 27% of the total. Here, we find the top five vendors, namely H3C (30%), Huawei (24%), Ruijie (20%), NEC (17%), and Nokia (5%), account for around 95% of the identified addresses. The remaining 4.7% comes from 76 different vendors. In contrast to the ICMP Echo method, where Cisco dominates, we identify only 223 Cisco addresses. Given this concentration, we speculate that some vendors still allow Address Mask replies by default. To validate this assumption, we examine the available vendor documentation. Both Ruijie and Cisco appear to disable address mask replies by default [10, 41], whereas Nokia allows them [34]. For the remaining three vendors, we could not find public documentation stating the default behavior. However, they do provide a mechanism to enable or disable address mask replies. We also examine whether Address Mask responsive address are known router addresses by comparing them against CAIDA’s ITDK dataset [9]. As of 08. 2025, the ITDK dataset contains 5.3M router addresses. However, only




Figure 3: Number of detected subnets per size and method.

286k of those are responsive to the Address Mask Request, suggesting that responsiveness may not be exclusive to network elements such as routers.

Validation from a Network Operator. We examine our research network and identify five /30 subnets. To verify the results, we contact the local network operator, who confirm their accuracy and, drawing on their experience, supports our lab observation that this behavior is specific to Cisco devices.

6 Preliminary Analysis

In this section, we discuss the detected subnets and inter-method consistency. We explore subnet coverage within BGP prefixes, investigate subnet liveness uniformity, and discuss potentially mis-configured broadcast addresses.

Subnets Overview. With the ICMP Echo approach across all vantage points, we detect 2.05M subnets. Similarly, we detect 1.82M subnets with the Address Mask approach. We combine both approaches, resulting in 3.8M unique subnets. In Figure 3, we plot the number of detected subnets and highlight the contribution of each method (note the log scale). The majority of subnets are /30s, with more than 1M and 820k subnets detected using ICMP Echo and Address mask, respectively. Furthermore, we observe more than 363k /32s and 101k /31s with the Address Mask method but none with Echo, as those subnets omit network and broadcast addresses, which are required for this method (recall Section 3.1). Overall, the Address mask approach contributes one third of the detected subnets for each subnet size between /20 and /30. Both /31 and /30 are commonly used for point-to-point links between routers [39]. However, /32s can server multiple purposes. For instance, RFC 6752 discusses their use for loopback interfaces [28], while prior work highlights their usage for BGP blackholing [13, 18]. Notably, routers are more likely to elicit an ICMP response, which could explain the higher fraction of small subnet sizes in our dataset. Further, such subnets can potentially reveal hidden links between routers that may not be visible with, e.g., traceroute. We also find 1,096 subnets with sizes ranging from /19 (537 subnets) to /17 (56 subnets), but exclude them from the figure.

Methods Consistency. To examine the consistency of our results, we compare the subnets detected by each method. We observe minimal overlap between the two approaches, with only 25k subnets found by both. We assess whether the detected subnets overlap by measuring how well the ICMP Echo-based subnets are covered by the Address Mask approach. From this analysis, we find that 6.9k Echo-based subnets (i.e., 0.3%) are contained within an Address Mask subnet. This shows that combining both approaches




Figure 4: ECDF of subnets coverage in BGP prefixes.

not only increases coverage but also enables more granular subnet size detection.

6.1 Subnets Coverage

BGP Prefixes. To assess how well our subnet dataset covers the IPv4 address space, we compare the subnets against BGP prefixes using a RouteViews snapshot from November 5, 2025 [40]. Our dataset maps to 159.8k BGP prefixes, representing 15% of the 1.07M prefixes in the snapshot. Figure 4 shows the fraction of coverage within each BGP prefix. We find that most BGP prefixes have low overlap with our dataset. Overall, we observe that half of the BGP prefixes having less than 10% of their address space covered by a subnet. Furthermore, less than 20% of the BGP prefixes have 50% or more of their addresses in our dataset. Notably, 8.3k BGP prefixes show coverage of 80% or more. We then examine Autonomous System (AS) coverage. We detect at least one subnet in 21.8k ASes. We detect a single subnet in 30% of ASes, whereas in 2% (432 ASes), we detect more than 1k subnets, with over 363k subnets in a single AS. Overall, our analysis indicates that the detected subnets tend to originate from unevenly distributed BGP prefixes and autonomous systems rather than uniformly across the entire address space. Although the coverage appears modest, this is expected, as our methods have inherent limitations. We rely on vendor-specific ICMP Echo behavior, which is implementation-dependent, or on the largely deprecated Address Mask Request. Consequently, our results are strongly influenced by the vendor ecosystem on the Internet and cannot be applied to every network.

DNS PTR Records. Regional Internet Registries (RIRs) and DNS Pointer Records (PTRs) can offer insights into subnet deployments. For example, if PTR records exist for .1 through .30 but not from .31 onwards, one could infer a /27 subnet. We therefore compare our inferred subnets against the <https://rir-data.org/dataset> [3]. The dataset contains more than 1.2M prefixes, with 97.9% of size /24, 2.07% corresponding to /16s, and only 22 prefixes more specific than a /24. We find minimal exact overlap, with 18.4k prefixes, all of which are /24, matching our dataset. However, our dataset contain 461k more-specific subnets for 110k prefixes. Note, however, that such comparison may not be reliable due to incomplete PTR records, historical artifacts, or operators deliberately avoiding PTR record creation. We therefore leave a more thorough cross-dataset analysis for future work.

6.2 Subnet Hosts Liveness Uniformity

Subnet-level structure often reflects uniform responsiveness behavior, as addresses within the same prefix tend to share similar




Figure 5: % of ICMP Echo replies per subnet size.

characteristics [25, 35]. We examine this from the perspective of host uniformity in responsiveness to ICMP Echo requests by calculating the percentage of ICMP Echo responsive hosts within each subnet. Overall, we find 40% of subnets to be fully responsive across all addresses, while only 2% are entirely unresponsive. To investigate this further, we plot an ECDF for each subnet size in Figure 5. Across all sizes, excluding /30 and /31, 25% of subnets are not responsive to ICMP Echo requests. Conversely, 55% and 75% of /30s and /31s, respectively, are fully responsive. The remaining subnets vary, with only a small fraction being fully responsive. Overall, we find smaller subnets to be more uniformly responsive.

Next, we examine whether such uniformity is consistent on the Autonomous System (AS) level. We select the twelve largest ASes in our dataset and plot an ECDF of the percentage of responsive addresses per subnet identified in each AS across all subnet sizes (see Figure 6). In seven out of the ten ASes, 50% or more of subnets are fully responsive. Notably, in AS 4713, 97% of its 48k subnets in our dataset are fully responsive. We find subnet responsiveness to be more uniform at the AS level than in the per-size subnet view shown in Figure 5.

6.3 Direct Broadcast

Direct broadcast, when enabled on misconfigured networks, can be abused for amplification attacks [46]. Our dataset provides a unique perspective for investigating such misconfigurations. In particular, if the same destination address triggers responses from multiple addresses within the same subnet, we can assume that the address is either a direct broadcast address or misconfigured. Across vantage points and subnet sizes, we detect between 73 and 84 subnets exhibiting this behavior. While most subnets responded from only two addresses, some responded from up to 15 different addresses. Note, however, that we consider only ICMP Echo Replies and leave examining other protocols and services for future work.

7 Related Work

Measurement studies on various aspects of Internet topology are an active area of research. However, subnet detection remains a relatively under-examined aspect, as subnet boundaries are not typically visible in the data plane, while BGP prefixes aggregate actual deployments. Early work by Tozal et al. [43, 44] introduces a set of heuristics for constructing a router-level Internet topology collector capable of identifying addresses that belong to the




Figure 6: % of ICMP Echo replies for top ASNs.

same subnet. Similarly, Gunes et al. [21] propose techniques to classify IP addresses as members of the same subnet. Lee et al. [31] further propose Hobbit, a method that aggregates IPv4 addresses into larger blocks based on their topological proximity. Unlike our work, these efforts rely on analytical processing of traceroute data and require accompanying path measurements. Sediqi et al. [42] analyze hyper-specific, e.g., /30s and /29s, IPv4 prefixes as observed in route collectors. They measure the prevalence of these prefixes and their use cases. However, these prefixes are not the norm, as BGP often contain prefixes up to /24. In contrast, our work does not rely on BGP data to detect subnets. Instead, we employ an empirical approach. More recent work by Koch et al. [30] investigates IPv6 subnet router anycast probing with the goal of identifying active IPv6 prefixes. Contrary to our work, they do not attempt to detect subnet sizes. We note, however, that our methodology is applicable only to IPv4. To the best of our knowledge, our work is the first to detect subnets, empirically, at scale.

8 Conclusion

In this paper, we introduce two methodologies for detecting subnets at scale. We present a new ICMP Echo-based technique and investigate the deprecated ICMP Address Mask message type, identifying 3.8M subnets of varying sizes. We evaluate the applicability of our methods across different network vendors and attribute the observed variation to implementation differences. We analyze the resulting dataset for its coverage of BGP-announced prefixes and find that it spans 15% of BGP prefixes across more than 20k ASNs. Our findings and dataset provide a new perspective for future Internet topology research. In future work, we envision using our dataset and techniques to enhance Internet mapping efforts, improving path discoverability and alias resolution, and providing insights into service similarity and subnet deployment practices.

Acknowledgements

We thank our reviewers for their constructive feedback, and our shepherd, Tijay Chung, for guiding the revision process. This work was supported by the European Union under the Horizon Europe Programme as part of the project SafeHorizon (Grant Agreement #101168562) and by the Federal Ministry of Research, Technology and Space of Germany in the programme of “StartUpConnect” Project QUICast, project identification number 16KIS2650.

References

- [1] Taha Albakour, Oliver Gasser, Robert Beverly, and Georgios Smaragdakis. 2021. Third Time's not a Charm: Exploiting SNMPv3 for Router Fingerprinting. In *Proceedings of the 21st ACM Internet Measurement Conference*. 150–164.
- [2] Taha Albakour, Oliver Gasser, Robert Beverly, and Georgios Smaragdakis. 2023. Illuminating router vendor diversity within providers and along network paths. In *Proceedings of the 2023 ACM on Internet Measurement Conference*. 89–103.
- [3] Alfred Arouna, Ioana Livadariu, and Mattijs Jonker. 2023. Lowering the Barriers to Working with Public RIR-Level Data. In *Proceedings of the 2023 Applied Networking Research Workshop (ANRW '23)*.
- [4] Shehar Bano, Philipp Richter, Mobin Javed, Srikanth Sundaresan, Zakir Durumeric, Steven J Murdoch, Richard Mortier, and Vern Paxson. 2018. Scanning the internet for liveness. *ACM SIGCOMM Computer Communication Review* 48, 2 (2018), 2–9.
- [5] Benjamin Beurdouche, Karthikeyan Bhargavan, Antoine Delignat-Lavaud, Cédric Fournet, Markulf Kohlweiss, Alfredo Pionti, Pierre-Yves Strub, and Jean Karim Zinzindohoue. 2017. A messy state of the union: Taming the composite state machines of TLS. *Commun. ACM* 60, 2 (2017), 99–107.
- [6] Robert Beverly. 2016. Yarrp'ing the Internet: Randomized high-speed active topology discovery. In *Proceedings of the 2016 Internet Measurement Conference*.
- [7] S. Bradner. 1997. Key words for use in RFCs to Indicate Requirement Levels. RFC 2119. IETF. <https://www.rfc-editor.org/rfc/rfc2119.txt>
- [8] Randy Bush, James Hiebert, Olaf Maennel, Matthew Roughan, and Steve Uhlig. 2007. Testing the reachability of (new) address space. In *Proceedings of the 2007 SIGCOMM workshop on Internet network management*. 236–241.
- [9] CAIDA. 2026. ITDK: Internet Topology Data Kit. <https://doi.org/10.21986/CAIDADATA.ARK-ITDK>. Dates used: 2025-08. Accessed: 2026-06. doi:10.21986/CAIDADATA.ARK-ITDK
- [10] Cisco Systems, Inc. 2014. *Cisco IOS XR IP Addresses and Services Command Reference, Release 4.3.x*. Cisco Systems, Inc. Accessed: 2025-12-11. https://www.cisco.com/c/en/us/td/docs/routers/xr12000/software/xr12k_r4-3/addr_serv/command/reference/b_ipaddr_cr43xr12k/b_ipaddr_cr42xr12k_chapter_01000.html
- [11] Kimberly Claffy, Young Hyun, Ken Keys, Marina Fomenkov, and Dmitri Krioukov. 2009. Internet mapping: from art to science. In *2009 Cybersecurity applications & technology conference for homeland security*. IEEE, 205–211.
- [12] Gregory Detal, Benjamin Hesmans, Olivier Bonaventure, Yves Vanaubel, and Benoit Donnet. 2013. Revealing middlebox interference with tracebox. In *Proceedings of the 2013 conference on Internet measurement conference*.
- [13] Christoph Dietzel, Anja Feldmann, and Thomas King. 2016. Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild. In *Passive and Active Measurement*.
- [14] D Dittrich and E Kenneally. 2012. *The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research*. Technical Report. U.S. Department of Homeland Security. doi:paper/2012_menlo_report_actual_formatted
- [15] Zakir Durumeric, Eric Wustrow, and J Alex Halderman. 2013. ZMap: Fast Internet-wide Scanning and its Security Applications. In *22nd USENIX Security Symposium*.
- [16] Xun Fan and John Heidemann. 2010. Selecting representative IP addresses for Internet topology studies. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. 411–423.
- [17] Oliver Gasser, Randy Bush, Massimo Candela, and Russ Housley. 2026. *Publishing End-Site Prefix Lengths*. Technical Report. doi:10.17487/RFC9977
- [18] Vasileios Giotas, Georgios Smaragdakis, Christoph Dietzel, Philipp Richter, Anja Feldmann, and Arthur Berger. 2017. Inferring BGP blackholing activity in the internet. In *Proceedings of the 2017 Internet Measurement Conference*.
- [19] F. Gont and C. Pignataro. 2013. *Formally Deprecating Some ICMPv4 Message Types*. RFC 6918. IETF. <https://www.rfc-editor.org/rfc/rfc6918.txt>
- [20] Jean-Francois Graillet, Fabien Tarissan, and Benoit Donnet. 2016. TreeNET: Discovering and connecting subnets. In *8th International Workshop on Traffic Monitoring and Analysis (TMA)*.
- [21] Mehmet H Gunes and Kamil Sarac. 2007. Inferring subnets in router-level topology collection studies. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*. 203–208.
- [22] Fahad Hilal, Taha Albakour, Oliver Gasser, and Kevin Vermeulen. 2026. Unpacking Internet Ossification: A Large-Scale Study of Path-Impairing Middleboxes Across IPv4 and IPv6. In *International Conference on Passive and Active Network Measurement*.
- [23] Fahad Hilal and Oliver Gasser. 2023. Yarrpbox: Detecting middleboxes at internet-scale. *Proceedings of the ACM on Networking* (2023).
- [24] Liz Izhikevich, Renata Teixeira, and Zakir Durumeric. 2021. {LZR}: Identifying unexpected internet services. In *30th USENIX security symposium (USENIX Security 21)*. 3111–3128.
- [25] Liz Izhikevich, Renata Teixeira, and Zakir Durumeric. 2022. Predicting IPv4 services across all ports. In *Proceedings of the ACM SIGCOMM 2022 Conference*.
- [26] Juniper Networks. 2024. SRX Self-Generated Traffic Not Taking Egress Interface IP as Source. Juniper Support Portal Knowledge Base. Accessed: 2025-12-11. <https://supportportal.juniper.net/s/article/SRX-Self-Generated-Traffic-Not-Taking-Egress-Interface-IP-as-Source>
- [27] Juniper Networks. 2025. *Junos OS CLI Reference*. Juniper Networks. Accessed: 2025-12-11. <https://www.juniper.net/documentation/us/en/software/junos/cli-reference/cli-reference.pdf>
- [28] A. Kirkham. 2012. *Issues with Private IP Addressing in the Internet*. RFC 6752. IETF. <https://www.rfc-editor.org/rfc/rfc6752.txt>
- [29] Johannes Klick, Stephan Lau, Matthias Wählisch, and Volker Roth. 2016. Towards Better Internet Citizenship: Reducing the Footprint of Internet-wide Scans by Topology Aware Prefix Selection. In *Proceedings of the 2016 Internet Measurement Conference*.
- [30] Maynard Koch, Raphael Hiesgen, Marcin Nawrocki, Thomas C Schmidt, and Matthias Wählisch. 2025. Scanning the IPv6 Internet Using Subnet-Router Any-cast Probing. *Proceedings of the ACM on Networking* (2025).
- [31] Youndo Lee and Neil Spring. 2016. Identifying and aggregating homogeneous ipv4/24 blocks with hobbit. In *Proceedings of the 2016 Internet Measurement Conference*. 151–165.
- [32] Matthew Luckie, Young Hyun, and Bradley Huffaker. 2008. Traceroute probe method and forward IP path inference. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*. 311–324.
- [33] Jeffrey Clifford Mogul and Jon Postel. 1985. Internet Standard Subnetting Procedure. RFC 950. doi:10.17487/RFC0950
- [34] Nokia. [n. d.]. *IES interface commands*. Nokia 7210 SAS 22.9.R1 Issue 01, accessed 2026-04-17. https://infocenter.nokia.com/public/7210SAS229R1A/topic/com.nokia.DEK_Services_Guide/ies_interface_c-d1772e2961.html
- [35] Fariba Osali, Khwaja Zubair Sediqi, and Oliver Gasser. 2025. Sibling Prefixes: Identifying Similarities in IPv4 and IPv6 Prefixes. In *Proceedings of the 2025 ACM Internet Measurement Conference*. 100–119.
- [36] Craig Partridge and Mark Allman. 2016. Ethical considerations in network measurement papers. *Commun. ACM* 59, 10 (Sept. 2016), 58–64. doi:10.1145/2896816
- [37] Vern Paxson. 1996. End-to-end routing behavior in the Internet. *ACM SIGCOMM Computer Communication Review* (1996).
- [38] Jon Postel. 1981. Internet Control Message Protocol. RFC 792. <https://datatracker.ietf.org/doc/html/rfc792>
- [39] A. Retana, R. White, V. Fuller, and D. McPherson. 2000. *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*. RFC 3021. IETF. <https://www.rfc-editor.org/rfc/rfc3021.txt>
- [40] Route Views. 2025. RouteViews project. <https://www.routeviews.org/routeviews/index.php/collectors/>
- [41] Ruijie Networks. 2022. *Ruijie RG-WLAN Series Access Points RGOS Command Reference Release 11.90B7P4*. <http://www.ruijiedestek.com/wp-content/uploads/2022/05/Ruijie-RG-WLAN-Series-Access-Points-RGOS-Command-Reference-Release-11.90B7P4.pdf>
- [42] Khwaja Zubair Sediqi, Lars Prehn, and Oliver Gasser. 2022. Hyper-specific prefixes. *ACM SIGCOMM Computer Communication Review* 52 (2022), 20 – 34.
- [43] M Engin Tozal and Kamil Sarac. 2010. Tracenet: an internet topology data collector. In *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. 356–368.
- [44] M Engin Tozal and Kamil Sarac. 2011. Subnet level network topology mapping. In *30th IEEE International Performance Computing and Communications Conference*. IEEE, 1–8.
- [45] Bulut Ulukapi, Anna Sperotto, and Ralph Holz. 2025. Towards understanding middlebox deployments in dutch ases: Impact of ip sampling size. In *Proceedings of the 2025 Applied Networking Research Workshop*.
- [46] Ramin Yazdani, Yevheniya Nosyk, Ralph Holz, Maciej Korczyński, Mattijs Jonker, and Anna Sperotto. 2023. Hazardous echoes: the dns resolvers that should be put on mute. In *2023 7th Network Traffic Measurement and Analysis Conference (TMA)*. IEEE.
- [47] Johannes Zirngibl, Philippe Buschmann, Patrick Sattler, Benedikt Jaeger, Juliane Aulbach, and Georg Carle. 2021. It's over 9000: analyzing early QUIC deployments with the standardization on the horizon. In *Proceedings of the 21st ACM Internet Measurement Conference*. 261–275.