Detecting and Characterizing Exposed BGP Routers

Shyam Krishna Khadka[*], Suzan Bayhan[*], Ralph Holz[†*], Cristian Hesselman[‡*]
[*]University of Twente, The Netherlands
[†]University of Münster, Germany
[‡]SURF, The Netherlands
{s.k.khadka, s.bayhan, r.holz, c.e.w.hesselman}@utwente.nl

***Abstract*—RFC 7454 on BGP operations and security recommends that network operators restrict access to TCP port 179, which routers use to receive reachability messages from their peers. Reports from commercial companies indicate that around 283k IP addresses expose TCP 179 to the Internet, potentially increasing their susceptibility to denial-of-service attacks. This paper presents a detailed insight into this underexplored problem of exposed BGP routers using an Internet-wide TCP scan of port 179. We analyze device responses at both the application layer via BGP OPEN messages and transport layer (successful TCP connections, TCP RSTs, or silent drops). We identify 141,313 IP addresses that accept TCP connections and exchange OPEN messages, corresponding to 20,432 routers across 4,194 Autonomous Systems (ASes). From the OPEN messages, we infer each router's IP addresses, the router's type (border or internal), and assess their criticality based on connected ASes and interface IP addresses. Using an external SNMPv3 dataset, we further examine vendor distribution among exposed routers. Our study offers new insights into router exposure: we identified 1,127 border routers, 16,124 internal routers, and 4,194 ASes associated with them, including critical routers in ASes. These insights can help network operators, policymakers, and security researchers assess the security of BGP routers and their associated ASes on the Internet.**

***Index Terms*—BGP, router, TCP protection**

I. Introduction

The Border Gateway Protocol (BGP) is the core routing protocol of the Internet, operating at the application layer over TCP port 179. Peer routers use BGP to exchange reachability messages, which is essential for the correct operation of the Internet and all services that use it, such as the Domain Name System (DNS), web services, databases, and IP telephony.

RFC 7454 [1] on BGP operations and security therefore recommends that network operators restrict access to port 179 to prevent it from being exposed to attacks. For example, an attacker could flood an open port 179 with requests to open a BGP session, causing the router's performance to degrade because the router allocates CPU, memory, and session state for each request. Such denial of service attacks bear particular risk if the router is part of a core network with a high peer count or if the attack traffic interferes with legitimate BGP sessions.

Unfortunately, network operators still expose TCP port 179 to the Internet. For example, Shadowserver reported 115 routers with open BGP ports in July 2023 [2]. Similarly, Shodan's BGP report on 15 December 2025 shows 283,995 IP addresses (IPs) with open BGP ports [3], although it does not indicate how many routers these represent.

The gap we aim to fill with this paper is the lack of understanding of the scale and characteristics of such "exposed routers". Existing sources, such as Shodan and similar companies, primarily identify devices that respond on TCP port 179, but do not provide a detailed characterization of how these IPs respond to BGP and TCP messages, as done in the present work. For example, they do not systematically map IPs to routers or distinguish border routers from internal routers. Furthermore, while prior research on exposed BGP routers is limited, with only one earlier study dating back to 2010 [4], most subsequent work has instead focused on securing BGP itself through mechanisms such as RPKI-based Route Origin Authorization (ROA) [5], Route Origin Validation (ROV) [6], and Autonomous System Provider Authorization (ASPA) [7].

Understanding the scale and characteristics of exposed routers offers important insights for multiple stakeholders. One example is the Mutually Agreed Norms for Routing Security (MANRS)+ working group [8], which seeks to enhance routing security through compliance and audits of Autonomous Systems (ASes). MANRS+ can use a better understanding of exposed routers to assess to what extent their member ASes implement the MANRS+ metric "BGP session protection". Additionally, understanding exposed routers helps researchers infer Internet topology and identify critical routers whose instability could have a large impact on Internet routing.

In this paper, we design a methodology that finds exposed routers through active measurements of TCP port 179 across the entire IPv4 address space and analyzes their operational behavior in detail using full packet captures. We emphasize that we do not aim to detect or exploit vulnerabilities of BGP daemons. Rather, our goal is to assess the extent to which Internet-reachable routers expose their control planes to BGP traffic from unknown sources, which might put them at risk.

Our contributions and main findings are as follows:

- We develop "Scan-179", a tool that scans port 179 across all IPv4 addresses. Scan-179 provides full packet captures of how IPs respond to BGP OPEN and TCP handshake messages, and maps these IPs to routers using BGP responses. Scan-179 is open source[1], allowing researchers

978-3-903176-83-6

[1]Source code: https://github.com/shyamkkhadka/scan-179

to reproduce our work and use it for future work.
- We conduct the first study to characterize exposed routers in terms of their types (border or internal), criticality (e.g., based on the number of their connected ASes and number of interfaces), their associated ASes, and exposure to unsolicited SNMPv3 requests. We demonstrate that we can reliably infer such properties using a minimal and non-intrusive BGP message exchange.
- We identify 141,313 interface IPs belonging to 20,432 exposed routers associated with 4,194 ASes that allow a remote peer to initiate a BGP session. These routers vary in criticality: we classify 1,127 as border routers, 16,124 as internal routers, and 3,181 could function as either.

The remainder of this paper is structured as follows. Section II provides background on exposed BGP routers. Section III reviews related work, and Section IV describes our methodology. Section V presents our findings on exposed routers and their characteristics, followed by a discussion in Section VI. Section VII concludes and outlines future work.

II. Background on exposed BGP routers

We discuss our definition of an exposed BGP router (Section II-A) and the risks that they may introduce for network operators and the Internet (Section II-B). We also discuss the protocol interactions that establish a BGP session (Section II-C) and the protections that the operators can use to restrict access to port 179 (Section II-D).

A. Definition of exposed BGP router

We define an exposed BGP router as a router whose BGP service on TCP port 179 is reachable from any IP address on the Internet and responds to BGP OPEN messages. In this work, we consider a router as exposed if it returns any valid BGP response, including OPEN, KEEPALIVE, or NOTIFICATION messages. The exposed routers allow an adversary to establish a TCP connection from anywhere on the Internet, without requiring cryptographic credentials or spoofing the IP address of a legitimate peer.

BGP OPEN is the first message sent by a BGP router after a TCP connection has been successfully established on port 179. It contains version, ASN, BGP identifier, and optional parameters. An example of a BGP OPEN message using a reserved ASN and a sample BGP Identifier IP for documentation purposes is as follows:

```
Message #1:
  Type: OPEN
  Length: 29
  Version: 4
  Hold Time: 90
  ASN: 64511
  BGP Identifier: 192.0.2.0
  Optional Parameters Length: 16
  Optional Parameters:
    Parameter 1:
      Type: 128 (Capability)
      Length: 8
      Value:
        - Capability: Route Refresh (RR)
        - Capability: 4-byte ASN (AS4)
```

BGP Identifier (BID): a unique 32-bit identifier assigned to a router within an AS [9], [10].

Autonomous System Number (ASN): a 2-octet or 4-octet identifier carried in the BGP OPEN message to identify the sender's AS [9], [11]. The ASN may be either public or non-public. The non-public ASN could be private or reserved and does not appear on the public Internet as defined in RFCs 5398 [12] and 7300 [13].

Optional parameters: The optional parameters are supported capabilities and extensions of a router, such as Multi-Protocol BGP (MP-BGP) and Route Refresh (RFC 9072 [14]).

B. Risk associated with exposed BGP routers

We identify three risks that an exposed router introduces, which exist because exposed routers accept TCP connections and process BGP OPEN messages from arbitrary Internet devices. They arise prior to full BGP session establishment.

Information leakage: BGP OPEN exchanges expose sensitive control-plane metadata, including AS numbers, BGP Identifiers, and supported capabilities. This information is typically not observable externally and can be leveraged for router fingerprinting, vendor identification, topology mapping, and targeted reconnaissance. For example, the authors of [15] demonstrate that network and transport-layer fingerprints can be used to identify router vendors across the Internet and to map router deployments.

Exposure to implementation vulnerabilities: Certain BGP implementations may process OPEN messages before fully validating critical fields. For example, vulnerabilities disclosed at Black Hat 2023 (CVE-2022-40302 [16] and CVE-2022-43681 [17]) showed that Free Range Routing (FRR), an open-source routing software suite, performed insufficient validation of certain OPEN message attributes. As a result, malformed messages could trigger denial-of-service conditions [18].

Control-plane resource exhaustion: Repeated partial attempts to establish BGP sessions force routers to allocate CPU, memory, and session state, even without completing a full BGP session [19]. This can potentially degrade router performance, interfere with legitimate peering sessions, and potentially cause operational instability.

C. How remote adversaries attempt to set up a BGP session

A remote adversary can attempt to establish a BGP session with a router by following the normal Finite State Machine (FSM) defined in RFC 4271. The adversary acts as the active peer (client) and initiates the connection, while the exposed router acts as the passive peer (server) listening on port 179.

The adversary first sends a TCP SYN. The BGP router then responds with a SYN-ACK, and the client completes the handshake with an ACK. If the TCP connection succeeds, the router sends a BGP OPEN message to the client, and we consider the router to be exposed. The router may later send a NOTIFICATION message to cease the connection.

If the client were a legitimate peer, the routers would then validate each other's OPEN messages by checking the ASN, BGP version, and capabilities. If they are valid, they

exchange KEEPALIVE messages to confirm mutual agreement on session parameters. The BGP session is now fully active, and routers exchange UPDATE messages to advertise or withdraw routes. Periodic KEEPALIVE messages maintain session liveness.

D. Existing protections for TCP 179

To protect TCP port 179, a combination of operator filtering techniques, as well as implementation-level and protocol-inherent mechanisms, can be applied as follows:

1. Access Control List (ACL). A network administrator implements an ACL to accept TCP connections from only the configured IPs (BGP peers), while dropping traffic from other IPs (see RFC 7454 and RFC 6192). Optionally, the administrator can also impose rate limits on traffic from legitimate peers.

2. Generalized TTL Security Mechanism (GTSM). With GTSM (RFC 5082), a network administrator configures a TTL threshold based on the expected hop distance to each BGP peer. Packets with a TTL below this threshold are dropped before TCP processing, preventing the establishment of TCP state for unauthorized IPs.

3. TCP authentication. The router verifies incoming TCP SYN segments using a cryptographic mechanism configured by network administrators, either through TCP MD5 signatures (RFC 2385) or the TCP Authentication Option (TCP-AO, RFC 5925). If authentication fails, the packet is discarded, and the TCP connection is not established.

4. TCP connection handling. This mechanism, implemented in the router's operating system kernel, denies state allocation for incoming TCP SYN messages if they come from an unpermitted source or exceed pre-defined resource limits (see RFC 9293 and RFC 6192). The router actively terminates the attempt with a TCP RST, which prevents unauthorized state exhaustion at the transport layer.

5. BGP protocol validation. This represents a protocol-inherent mechanism defined in RFC 4271. After completing the TCP handshake, the router responds with a BGP OPEN followed by a NOTIFICATION or a NOTIFICATION alone if it does not recognize the ASN or BID, thus terminating the BGP session.

III. Related work

The work that comes closest to ours is that by Cavedon et al. [4]. They conducted a study in 2010 in which they first scanned TCP port 179 using Nmap and then developed a custom BGP speaker to probe approximately 2.2 million IP addresses, determining which devices were willing to establish a BGP session. The authors reported that 45% of the devices that completed the TCP three-way handshake immediately closed or reset the connection, while 0.4% (7,904 devices) successfully parsed the BGP OPEN message but declined the peering session by returning a NOTIFICATION message. In contrast to their approach of listing only IPs, we go further by analyzing responses to BGP OPEN messages, characterizing them in terms of routers representing those IPs, associated ASes, and the router vendors. Also, we terminate the connection immediately after receiving the OPEN message, without completing session establishment. This reduces the intrusiveness of the measurement while still providing insight into the presence and behavior of the BGP routers.

Another closely related work is by Albakour et al. [20], which identifies BGP-speakers using a ZMap [21] and ZGrab2 [22] pipeline by first finding devices with TCP port 179 open and then completing the TCP handshake. In contrast, our Scan-179 tool explicitly sends the initial BGP OPEN message from the probing side, as we observe that not all routers initiate BGP message exchange on their own after TCP establishment. Similar to their approach, we use the ASN in the BGP OPEN message together with optional parameters to group multiple IP addresses belonging to the same router. However, while their work focuses on alias resolution, our study goes beyond router grouping by characterizing exposed BGP routers, mapping them to hosting ASes, classifying their roles as border or internal routers, and estimating their criticality.

Czyz et al. [23] show that IPv6 routers and servers often expose high-value services (e.g., SSH, Telnet, SNMP, BGP) more than their IPv4 counterparts due to inconsistent security policies. For the case of BGP, they look at the IPs that respond to TCP 179. However, we extend this approach of scanning TCP 179 by sending BGP OPEN messages to responsive addresses, mapping them to routers and ASes, and focusing exclusively on IPv4 BGP routers.

A blog post [24] published in October 2023 reports on experiments to determine how routers block unsolicited BGP messages by default. The author of the blog evaluated 4 BGP platforms: Cisco Nexus 9300v, Cisco IOSv, Arista, and FRR routing, and revealed that Arista EOS provides comparatively stronger default protections than other platforms. The Arista EOS did not respond to unsolicited TCP SYN packets for BGP sessions, and its built-in iptables rules dropped unexpected BGP packets early in the packet-processing path. In contrast, the other platforms were more permissive in their default handling of such traffic, highlighting significant variation in vendors' security-by-default postures.

IV. Methodology

Our methodology begins by identifying BGP routers using ZMap and Scan-179 (Sections IV-A and IV-B), adhering to the ethical considerations outlined in Appendix A. We conducted a one-shot Internet-wide IPv4 measurement campaign between 2 and 5 February 2026 from a vantage point located in Enschede, the Netherlands. Next, we analyze the collected data by extracting metadata for each IP, including response categories, ASNs, router identification, router types, and vendor classification. We filter out likely honeypots before forming the final dataset, which is used for analysis in Section V.

A. Identifying possible BGP routers using Zmap

Since there is no authoritative list of BGP routers and BGP runs on port 179, we conduct an Internet-wide IPv4 scan for

port 179 using ZMap [21]. ZMap is a popular high-speed scanner that can probe the entire IPv4 address space efficiently. We use ZMap to identify devices that respond to TCP SYN with SYN-ACK on port 179. Some devices may drop SYN probes due to filtering, but our scans still allow us to identify routers that respond to unsolicited BGP connection attempts.

We conduct our own scans because we need the full packet captures of responding devices to study their behavior in detail. For example, with our own scans, we can observe whether devices respond with TCP RST packets, silently drop the probes, or proceed with BGP-level message exchanges. In contrast, platforms such as Censys only provide lists of IPs that respond on port 179 and do not include the packet-level information needed to analyze these behaviors.

B. Scan-179

Our custom-built Python-based tool (Scan-179) sends BGP OPEN messages through TCP connections that it established with the devices obtained from a ZMap scan. To maximize device discovery, we set a conservative 30-second socket timeout, ensuring that delayed TCP or BGP responses are captured.

Next, we send a syntactically valid BGP OPEN message to IPs identified by ZMap as potential BGP speakers using Scan-179. This RFC 4271–compliant OPEN message includes BGP version 4, our institution's ASN, a 90-second hold time, a BGP Identifier set to a publicly routable IP address from our institution, and five optional capabilities.

We conservatively select five BGP optional capabilities to resemble a real BGP router and reduce the likelihood that targets classify our probes as scanning. These include IPv4 and IPv6 Unicast (RFC 4760), reflecting common dual-stack deployments; Route Refresh (RFC 2918), which allows routing information to be updated without resetting sessions; Extended Messages (RFC 8654), enabling support for larger UPDATE messages; Graceful Restart with a 120-second timer (RFC 4724), used to preserve forwarding state during control-plane restarts; and 4-byte ASN support (RFC 4893), ensuring compatibility with the modern ASN space. Capability codes are assigned by the Internet Assigned Numbers Authority [25].

C. Categorizing BGP responses

We distinguish three categories of devices that respond to our BGP OPEN messages in a way that is compliant with the BGP specification in RFC 4271. These devices allow our messages to reach the router's control plane, which is primarily responsible for maintaining routing tables, processing BGP updates, and enforcing routing policies rather than controlling access. As a result, this layer is inherently challenging to secure against unauthorized or malicious connections.

BGP OPEN with KEEPALIVE. We found that some IPs responded to our BGP OPEN with a KEEPALIVE. In this case, our probe sent a TCP FIN to terminate the connection. This behavior is the riskiest among the observed behaviors because the device progresses beyond the OPEN stage to establish a BGP session with our probe, exposing ASN, BID, and capabilities.

BGP OPEN with NOTIFICATION. The device responds with a BGP OPEN message followed by a NOTIFICATION (Error Code 6, Subcode 5), per RFC 7421. This behavior exposes control-plane information such as the BID and ASN, but the risk is lower than in the KEEPALIVE case because the device terminates the connection immediately. Specifically, it sends a FIN to close the TCP session and prevent further progression of the BGP session establishment process, and our probe subsequently sends a RST to fully terminate the connection.

BGP NOTIFICATION only. The device promptly rejects the unsolicited BGP OPEN message by responding with a BGP NOTIFICATION, without sending its own OPEN, and immediately closes the connection. Although the device exposes TCP port 179 and allocates some control-plane resources, the associated risk is lower than in cases where the OPEN message is accepted. This is because the BGP session does not progress.

D. Categorizing non-BGP responses

We also received the responses that allowed TCP connections but returned messages that did not comply with RFC 4271. For instance, some devices actively terminate the session by sending TCP FIN or RST packets, while others silently discard our BGP OPEN message without returning a TCP RST or a BGP NOTIFICATION message. In these cases, we cannot determine whether the responding device is a BGP router or an intermediary device that intercepts our BGP OPEN message and generates a reply on behalf of the destination device. One example is a Juniper firewall [26] patent that Juniper implemented in their secure service gateway proxy [27]. It protects against unwanted connections through tarpitting, where the firewall responds to incoming SYN packets with a zero-window SYN-ACK [28]. After the client completes the TCP handshake, the firewall initiates a separate SYN toward the backend server to establish the actual connection.

E. Categorizing Pre-TCP-handshake responses

We observed that some devices did not complete the TCP handshake. These pre-TCP-handshake responses are socket timeout and connection refused. A socket timeout indicates that no response was received to our TCP SYN probe within the timeout of 30 seconds. We can not determine the underlying cause based on our measurements, but it may be the result of packet filtering by firewalls, the silent dropping of packets, or a lack of connectivity. In the case of a connection refused case, the device actively responds with a TCP RST. For devices that are real BGP routers, this behavior may indicate the presence of defenses such as ACLs, GTSM, connection handling, or TCP authentication, as described in Section II-D.

F. Associating ASes with BGP responders

Our methodology maps each IP address that responds with a BGP OPEN or a BGP NOTIFICATION message (see Section IV-C) to its ASN to identify the network operator that manages the corresponding router.

1) ASNs of OPEN responders: If the BGP responder uses a public ASN in its OPEN response (optionally followed by a KEEPALIVE or NOTIFICATION), we consider that ASN as the one managing the router, as per RFC 4271. If the ASN is non-public (e.g., private, or reserved), we map the responder's IP address to an ASN using IPinfo (2 February 2026) [29]. IPinfo performs this mapping by analyzing BGP routing data and Regional Internet Registry (RIR) information. This approach allows us to associate the IP with the AS that provides its network connectivity.

2) ASNs of NOTIFICATION-only responders: If the BGP responder only returns a NOTIFICATION, we infer its ASN by mapping the responder's IP address to an ASN using IPinfo. Since the IP address may belong to an upstream provider or a peer, as per RFCs 3021 [30] and 1518 [31], this mapping provides a reasonable approximation of the AS in which the responding router is observed. As an illustrative example, we observed that an IP address in SURF's network (Dutch national research and education network provider) responded with a NOTIFICATION message. Upon inquiry, SURF confirmed that the IP was assigned to a peer and that such address assignments are common in interconnections between networks.

G. Mapping OPEN responders to routers

Since routers may have multiple interfaces, multiple OPEN responder IPs may belong to the same router. If those IPs have the same public ASN and the same BID in their OPEN messages, they belong to the same router. Because, as per RFC 6286, the BID must be unique per router within an AS, combining the BID with the ASN allows us to identify routers.

For BGP responder IPs using a private ASN, the BID and private ASN alone are insufficient. This is because the same combination can appear in multiple ASes' internal networks as private ASNs are used internally. In these cases, we additionally include the associated public ASN obtained in Section IV-F. According to RFC 6286, a BID must be unique within an AS to prevent session collisions and routing loops. By combining the public ASN, private ASN, and BID, we form a composite key that distinguishes internal routers across different ASes. This key provides a reasonable grouping of routers that likely belong to the same AS.

In addition to following our methodology based on RFCs 4271 and 6286, we provide further evidence for our router grouping in two ways: (a) examining the optional parameters in their BGP OPEN messages, and (b) analyzing latencies from Scan-179. Firstly, for each identified router, all IPs grouped under the same router reported the same set of optional parameters in their OPEN responses. This consistency indicates that these IPs belong to the same router. However, identical BGP capabilities do not necessarily imply that two routers are the same router. In our measurements, we observed 36 patterns of optional parameters across 20,432 routers, showing that optional parameters alone may not always uniquely distinguish routers. To improve identification, we therefore also examine the response time of each responder to our scan and compare latencies across IPs believed to belong to the same router.

For 95% of routers, latency differences of the IPs within a router were 4ms or less, indicating those IPs likely refer to the same router. Only about 5% of routers showed larger differences, up to 16ms, which may be due to variations in Internet paths. Since our scan took place over 78 hours, measurements for different IPs of the same router were sometimes collected hours apart. During that time, routing paths or network load could have changed, contributing to the observed latency differences. While latency alone cannot definitively identify routers, particularly distinct routers in proximity, these observations provide supporting evidence for grouping IPs within the same router.

H. Classifying router types

We classify the exposed routers as border or internal to assess their operational impact. This classification is primarily guided by RFCs 5398, 6996, and 7300, which define non-public ASNs intended for internal use. We map our observations to these RFC guidelines for the classification. In some cases where we cannot classify the category of the router, we consider them as unclassified.

Border router. We classify a router as a border router if we observe a public ASN in its OPEN message, and it participates directly in eBGP peering with at least a public ASN different from its own. Instability in such routers can directly affect inter-domain routing. To determine whether a router participates in eBGP peering with another public ASN, we check whether it has at least one interface IP address that belongs to a different public ASN. This serves as an indication that the router is connected to another AS. This heuristic is based on common operational practices in directly connected eBGP deployments, where the two peering routers typically share a subnet assigned by one of the participating ASes (often a /30 or /31, consistent with RFC 3021). As a result, at least one interface IP on the link belongs to the address space of one of the ASes involved.

We evaluated this heuristic with the CAIDA AS Rank dataset [32] and through discussions with a network operator in the Netherlands. When two border routers belonging to different ASNs are directly connected, they reveal the relationship between the two ASNs. CAIDA's AS Rank infers such relationships from traceroute and BGP data. Among the 1,127 routers we identified as border routers, we extracted the ASNs of their connected neighbors and compared the resulting AS-to-AS links with CAIDA's inferred relationships. Of the 2,572 AS-to-AS links observed from these routers, 1,719 (67%) matched CAIDA's dataset. The remaining differences likely reflect AS relationships that are not visible from CAIDA's vantage points, which rely on a limited number of public BGP collectors and traceroutes. These results provide supporting evidence for the heuristic, but do not constitute full validation.

Internal router. We classify a router with a non-public ASN (private, or reserved) in its OPEN message as an internal router. As per RFCs 5398, 6996, and 7300, the non-public

ASNs are meant for internal purposes, and therefore, routers configured with those ASNs can not participate in eBGP peering with another public ASN. Since those routers are not connected to any other AS, they do not affect inter-domain routing. However, they affect routing within that AS. Because non-public ASNs do not require registration, ASes commonly use them for internal management and may assign them to customers to establish eBGP sessions without the customers having a public ASN. When such routes are propagated to the public Internet, the AS replaces the private ASN with its own public ASN [33]–[35].

Unclassified router. This type of router has a public ASN in its OPEN message, but all its interface IPs map to the same ASN. As a result, we cannot determine whether they function as border or internal routers based on our measurements, and therefore classify them as unclassified. This ambiguity arises from two possible scenarios. First, the router may be a border router participating in external BGP (eBGP). However, the router uses IP addresses from its own ASN on the interconnection link, providing its address space to the peer AS. This prevents the observation of interface IPs belonging to a different ASN. Second, the router may be an internal router participating only in internal BGP (iBGP) within the same AS. Due to these indistinguishable cases, we cannot reliably classify its type.

I. Finding router vendors

We identify the vendors of OPEN responders to evaluate whether they also accept unsolicited SNMPv3 connections. To accomplish that, we query Shodan's API for SNMPv3 responses using the IP addresses of OPEN responders. We use SNMPv3 responses because prior work has shown that SNMPv3 replies provide persistent identifiers and detailed device information that can be used to fingerprint routers and determine vendor characteristics at Internet scale [36]. We acknowledge that not all routers respond to unsolicited SNMPv3 requests. We can only obtain vendor information for a subset of OPEN responders. Also, operators should restrict SNMP access to authorized management stations only, as per operational security best current practices (RFC 4778).

J. Excluding possible honeypots

We exclude potential (BGP) honeypots based on observed behavior. We check whether the BGP responder's OPEN message contains the same ASN and BID as the one our scanner uses in its BGP OPEN request message. Such responses are unlikely to originate from legitimate remote routers, since a BGP router advertises its own ASN and BID in the OPEN message rather than copying those of its peer (RFC 4271). We identified 1,627 IPs exhibiting this behavior. We queried Shodan and found that these IPs expose numerous open ports with a minimum of 151 ports, and are explicitly labeled with the tag "honeypot". This confirms they are likely honeypots.

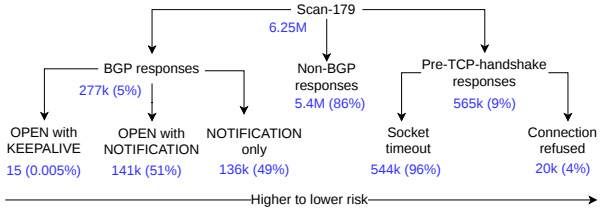

Fig. 1: Distribution of observed response types with corresponding counts and percentages.

V. Results: Exposed BGP routers and their characteristics

We analyze the exposed routers we found with our methodology and discuss their characteristics: their type (border or internal), criticality, the distribution of ASes by router type, and the vendors of exposed routers. Our results are based on our Scan-179 run of 2 February 2026.

A. Exposed BGP routers and their ASes

Figure 1 summarizes the results of our Scan-179 run (Section IV-B). It shows the number and percentages of BGP, non-BGP, and pre-TCP-handshake responses for the 6.25M IP addresses we obtained from ZMap (Section IV-A). Risk decreases from left to right, with OPEN with KEEPALIVE responses having the highest risk.

About 5% (277,408) of devices responded with a valid BGP message, thus revealing exposed BGP routers. Among these, 141k responded with an OPEN message, representing 20,432 exposed BGP routers. 136k devices only sent a NOTIFICATION message without exchanging an OPEN message (more in Section V-F). Using the methodology from Section IV-F, we identified 4,194 ASes with exposed BGP routers.

For 86% of the devices in our dataset (5.4M), we received non-BGP responses. The remaining 9% of the devices in our dataset (565k), TCP connections could not be established. These two types of devices do not allow us to reliably determine whether they are BGP routers. Consequently, the remainder of this study focuses on exposed routers that exchange OPEN messages (the two left-most branches of Figure 1).

B. Border and internal routers

We use the methodology of Section IV-H to classify exposed BGP routers that exchange OPEN messages into border or internal routers. We found 1,127 border routers and 16,124 internal routers. The border routers collectively interconnect 1,496 ASes. We could not classify the router type of 3,181 exposed routers (type "unclassified").

Figure 2 indicates that internal routers constitute the majority of exposed devices in our dataset. Since the total number of each router type on the Internet is unknown, the smaller number of exposed border routers in our dataset does not necessarily imply that they are generally more secure than internal routers. Furthermore, the exposure of internal routers indicates non-standard operational behavior because they are

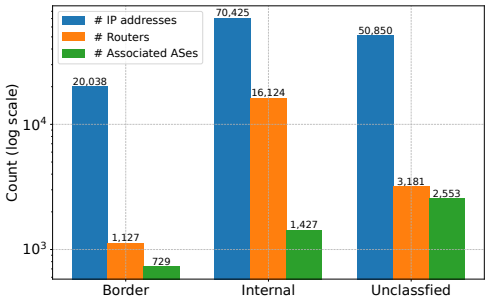

Fig. 2: Distribution of IPs, routers, and associated ASes across border, internal, and unclassified routers.

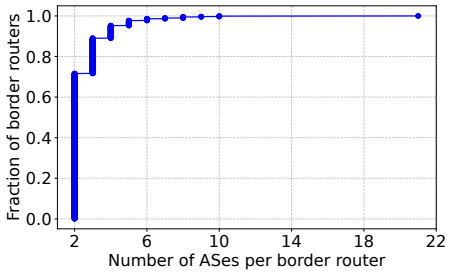

Fig. 3: Distribution of the number of ASes connected to each border router. Routers connected to more ASes are likely more critical in inter-domain routing.

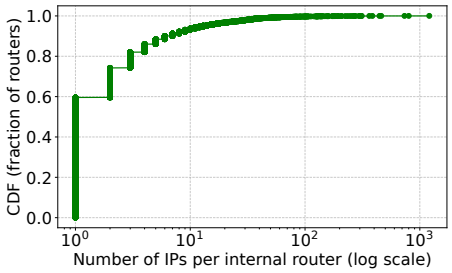

Fig. 4: Distribution of IPs per internal router. Routers having more IPs are likely more critical in their AS.

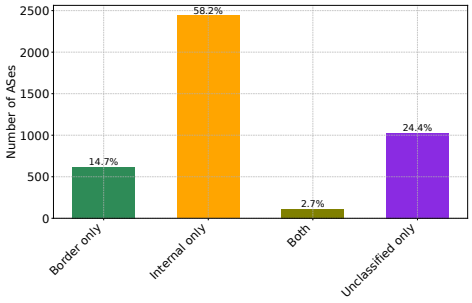

Fig. 5: Distribution of ASes by router type: Border only, Internal only, Both, and Unclassified only.

intended to only handle intra-AS traffic and are generally not expected to accept OPEN messages from other public ASNs.

C. Router criticality

We assess the criticality of exposed border routers by considering a border router more critical if it connects to more ASes, which we use as a simple heuristic for their relevance in inter-domain routing. For internal routers, we consider a router more critical within its AS if it has a larger number of interface IPs, as this likely indicates it aggregates multiple internal networks.

For the purpose of estimating criticality, we conservatively treat unclassified routers as internal routers and measure their importance using the number of interface IPs as a proxy metric. This allows us to assess the potential impact of these routers, while acknowledging that some may, in fact, be border routers with relatively few IPs. As a result, our approach provides an approximate, conservative view of their potential importance without assuming that these routers actively participate in inter-domain routing.

Figure 3 shows that around 75% of exposed border routers connect to only one other AS, while roughly 20% connect to up to three ASes. A border router associated with a research and education network in the US connects to 21 ASes, suggesting this exposed router is likely a critical border router.

Figure 4 shows the distribution of IPs per exposed internal router. 60% have only one IP address, and around 90% have fewer than 10. The remaining 10% have substantially more, up to 1,208 IPs. These routers may be critical within their AS, as they aggregate many subnets. The top three ASes are an ISP in the Philippines with 1,208 IPs (non-public ASN), a telecommunications provider in Grenada with 801 IPs (public ASN), and an ISP in Vietnam with 735 IPs (non-public ASN).

D. Distribution of ASes by router type

We analyze the distribution of ASes by router type to determine how many ASes expose only border routers, only internal routers, or both. The objective of this analysis is to assess whether exposure of port 179 differs between border and internal routers, potentially reflecting differences in security configurations for this port by router type. “Unclassified only” indicates ASes with routers whose type could not be determined. Unlike our criticality analysis (Section V-C), we consider unclassified routers separately to avoid bias in classifying AS-level exposure.

Figure 5 shows that the fraction of ASes in our dataset that only expose internal BGP routers is larger than those that only expose border routers (around 58% and 15% respectively). This suggests that stricter security controls are in place for port 179 on border routers. Furthermore, 112 ASes have both border and internal routers exposed. These counts represent a lower bound, as some routers in the unclassified category may belong to either type (border or internal). The ASes having both border and internal routers appear to apply similar security configurations that do not restrict access to port 179 on all router types.

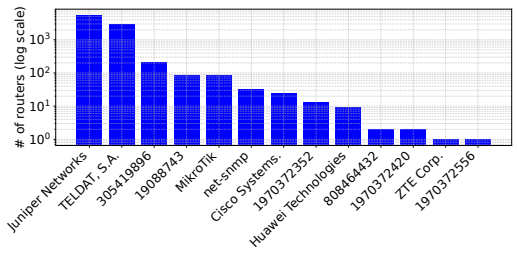

Fig. 6: Vendors of open-responder routers. Numbers indicate vendors not using a private enterprise number from IANA [37].

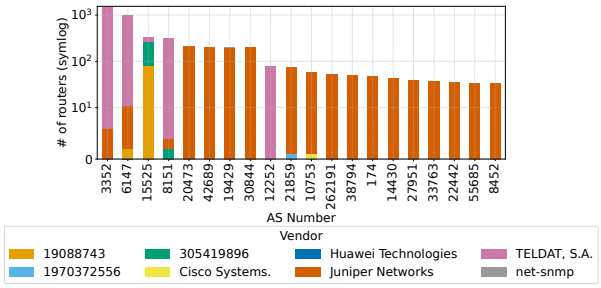

Fig. 7: Vendors across the top 20 ASes ranked by the highest number of known router vendors.

E. Vendor distribution

Using the methodology of Section IV-I, we identified the vendors of 8,871 (43.4%) of 20,432 exposed routers and found a total of 13 distinct vendors. This SNMPv3-based method classifies the majority of them as Juniper routers (5,448) or Teldat routers (2,956).

From the ASes with identifiable vendors, we selected the top 20 ASes with the highest number of exposed routers. Our goal was to examine whether these routers are concentrated among a particular vendor and to determine if ASes apply similar security configurations for port 179 across all observed router vendors. Figure 7 shows that most exposed routers in these ASes come from a single vendor, with Juniper being the most common. Only 6 ASes deploy routers from up to 3 different vendors. This suggests that, based on our study, exposure of routers is not vendor-specific, but rather that these ASes appear to apply similar security configurations for port 179 restriction on BGP routers across vendors.

We identified 2,421 ASes with routers responding to our probes on both the SNMPv3 and BGP ports, increasing the risk in these ASes. The risk is because SNMPv3 allows vendor fingerprinting, while exposed BGP ports reveal information about routers, including their roles (border or internal), which can facilitate more effective targeted attacks.

F. IPs sending NOTIFICATION without OPEN

We observed that 136,095 IPs respond with a BGP NOTIFICATION message without completing the exchange of OPEN messages. Those IPs are distributed across 3,420 ASNs, which may either directly host the routers associated with the responding IPs, or they are the ASNs' peers as explained in Section IV-F. We cannot reliably determine the true ASN of the responding router because NOTIFICATION messages do not carry the identifying information of an OPEN message. We consider such routers to pose a lower risk than routers that exchange OPEN messages, as they terminate the session early with a NOTIFICATION message. Nevertheless, they still process unsolicited BGP messages and allocate control-plane resources to generate the response, which may expose them to control-plane resource exhaustion attacks.

VI. Discussion and Limitations

Using our *Scan-179* tool and our analysis of BGP responses, we observe that only about 5% of the IPs responding on port 179 represent actual exposed BGP routers. These correspond to 20,432 routers managed by 4,194 ASes, suggesting that a non-trivial number of routers do not follow the RFC 7454 recommendation to restrict access to port 179 from unknown sources.

We further identify 1,127 exposed border routers connected to 1,496 neighboring ASes, highlighting their criticality in interdomain routing and the potential effects that these routers' open BGP ports might have on other ASes. Additionally, 112 ASes expose both internal and border routers, suggesting that these ASes allow external access to port 179 on both router types in their networks.

We observe a dominance of exposed Juniper routers, which account for approximately 61% of the routers (8,871) whose vendors we were able to identify using SNMPv3 data. However, this observation should not be interpreted as indicating that these vendors are inherently less secure, as their prevalence may simply reflect their broader deployment across operational networks. This finding contrasts with the results reported in [36], where SNMPv3-based router fingerprinting indicated Cisco as the dominant vendor in 2021. Notably, that study considered all SNMPv3 ports, regardless of whether they operated BGP.

The insights into exposed BGP routers and their characteristics are useful for multiple stakeholders. For example, MANRS+ could use these findings to assess the extent to which its member ASes implement the *BGP session protection* metric. Information about exposed and potentially critical routers may also help network operators better understand their network exposure and guide improvements in router configuration and access control. In addition, data on router interfaces (e.g., the number of IPs per router) may assist researchers in studying router interconnections and inferring relationships between ASes. More broadly, such information can contribute to a better understanding of both intra-AS and inter-AS Internet topology.

We contacted 12 network operators to better understand the phenomenon of exposed routers and validate our findings. One operator confirmed our IP to ASN mapping results and stated that BGP's default protocol behavior provides sufficient port 179 protection, and that it therefore does not require additional protection. We speculate that other possible reasons

that operators do not further enhance the security of port 179 are complex vendor hardening guidelines that are difficult to interpret [24], and software limitations that prevent the deployment of port 179 security mechanisms.

Limitations. We acknowledge the following five limitations of our study.

1) We lack ground-truth data for validation. Consequently, our heuristic-based router grouping and router classification using ASN visibility (e.g., private vs. public), interface IP ownership, and IP-to-router mapping may not always accurately reflect operational deployments, potentially leading to misclassified routers or incorrect ASN attribution. This also limits our ability to infer router protection mechanisms from observed response behavior, such as socket timeouts, and timing between TCP handshake and BGP OPEN. Obtaining such ground truth at Internet scale remains challenging, and we are not aware of an alternative for studying exposed BGP routers at the scale of our study.

2) We use IPinfo to map IPs to ASNs in two cases: when non-public ASNs appear in BGP OPEN messages, and when IPs respond with NOTIFICATION-only messages. Errors or outdated information in IPinfo can lead to incorrect ASN attribution and affect router grouping and classification, despite IPinfo's widespread use in network measurement research. In some cases, routers with private ASNs may be connected to multiple upstream ASes and assigned IP addresses from those ASes. We associate such routers with each observed ASN, which may lead to duplicate counting across ASes. We expect such cases to be very rare, as private ASNs are typically used within a single provider environment.

3) Our methodology cannot distinguish routers that actively initiate BGP OPEN messages after a TCP handshake from those that respond only after receiving an OPEN from our probe.

4) Our study is limited to IPv4 and only captures routers that respond to our probes. We cannot distinguish whether non-responsive IPs correspond to unreachable devices, non-routers, non-BGP routers, or protected BGP routers. Extending the study to IPv6 would require different hitlist-based measurement strategies due to its large address space.

5) Our measurements provide a snapshot of exposed BGP routers from a single point in time and a single European vantage point. Router visibility and exposure may vary over time and across probing locations due to operational changes, routing policies, or filtering.

VII. Conclusions and future work

We present the first study to identify exposed BGP routers on the Internet, which violate recommended security practices (RFC 7454) by accepting unsolicited BGP OPEN messages. Our tool *Scan-179* detects these routers through Internet-wide scans on port 179 and analyzes the BGP OPEN responses based on full packet captures. We also characterize the routers we identified by type (border or internal), criticality, the ASes that manage them, and their vendors.

We identified 20,432 exposed BGP routers across 4,194 ASes: 1,127 border routers, 16,124 internal routers, and 3,181 either. While such insights about exposed routers can advance Internet measurement research and improve understanding of network deployments, they also highlight potential security risks. For example, our method can infer critical routers of ASes that adversaries could exploit for reconnaissance. We therefore recommend that network operators strengthen router configurations to reduce exposure and prevent leakage of sensitive operational information.

Based on our dataset of exposed routers, their IP addresses, configured ASNs, and associations with 4,194 ASes, we aim to construct router-level and AS-level subsets of the Internet topology. This effort can complement existing datasets, such as CAIDA AS relationship [38], and provide additional insights into Internet structure and connectivity. In addition, we plan to engage with more network operators to understand the causes of exposed BGP routers, share our results at forums such as RIPE and NANOG, and validate our findings. We also plan to extend our measurements using geographically distributed vantage points and longitudinal measurements to better capture temporal and location-dependent variations in exposed BGP router exposure..

Acknowledgment

We want to thank our shepherd, Luigi Iannone, and anonymous reviewers for their valuable feedback on our paper. This research was funded by the Dutch Research Council (NWO) as part of the projects CATRIN (NWA.1215.18.003) and UPIN (CS.004). CATRIN is part of NWO's National Research Agenda (NWA).

References

[1] J. Durand, I. Pepelnjak, and G. Döring, "BGP Operations and Security," RFC 7454, Feb. 2015. [Online]. Available: https://www.rfc-editor.org/info/rfc7454

[2] "BGP Usage Report," Last accessed 15-December-2025. [Online]. Available: https://www.shadowserver.org/what-we-do/network-reporting/open-bgp-service-report/

[3] "BGP Usage Report," Last accessed 15-December-2025. Data retrieved via query for product:bgp port:"179". [Online]. Available: https://www.shodan.io/

[4] L. Cavedon, C. Kruegel, and G. Vigna, "Are BGP routers open to attack? an experiment," in *International Workshop on Open Problems in Network Security*. Springer, 2010, pp. 88–103.

[5] J. Snijders, B. Maddison, M. Lepinski, D. Kong, and S. Kent, "A Profile for Route Origin Authorizations (ROAs)," RFC 9582, May 2024. [Online]. Available: https://www.rfc-editor.org/info/rfc9582

[6] G. Huston and G. Michaelson, "Validation of Route Origination Using the Resource Certificate Public Key Infrastructure (PKI) and Route Origin Authorizations (ROAs)," RFC 6483, Feb. 2012. [Online]. Available: https://www.rfc-editor.org/info/rfc6483

[7] A. Azimov, E. Bogomazov, R. Bush, K. Patel, J. Snijders, and K. Sriram, "BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects," Internet Engineering Task Force, Internet-Draft draft-ietf-sidrops-aspa-verification-24, Oct. 2025, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/draft-ietf-sidrops-aspa-verification/24/

[8] "MANRS," Last accessed 03-March-2026. [Online]. Available: https://manrs.org/

[9] Y. Rekhter, S. Hares, and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271, Jan. 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4271

[10] E. Chen and J. Yuan, "Autonomous-System-Wide Unique BGP Identifier for BGP-4," RFC 6286, Jun. 2011. [Online]. Available: https://www.rfc-editor.org/info/rfc6286
[11] Q. Vohra and E. Chen, "BGP Support for Four-Octet Autonomous System (AS) Number Space," RFC 6793, Dec. 2012. [Online]. Available: https://www.rfc-editor.org/info/rfc6793
[12] G. Huston, "Autonomous System (AS) Number Reservation for Documentation Use," RFC 5398, Dec. 2008. [Online]. Available: https://www.rfc-editor.org/info/rfc5398
[13] J. Haas and J. Mitchell, "Reservation of Last Autonomous System (AS) Numbers," RFC 7300, Jul. 2014. [Online]. Available: https://www.rfc-editor.org/info/rfc7300
[14] E. Chen and J. Scudder, "Extended Optional Parameters Length for BGP OPEN Message," RFC 9072, Jul. 2021. [Online]. Available: https://www.rfc-editor.org/info/rfc9072
[15] T. Albakour, O. Gasser, R. Beverly, and G. Smaragdakis, "Illuminating router vendor diversity within providers and along network paths," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023, pp. 89–103.
[16] National Vulnerability Database, "CVE-2022-40302 Detail," 2023, Last accessed 11-May-2026. [Online]. Available: https://nvd.nist.gov/vuln/detail/CVE-2022-40302
[17] NVD, "CVE-2022-43681 Detail," 2023, Last accessed 11-May-2026. [Online]. Available: https://nvd.nist.gov/vuln/detail/CVE-2022-43681
[18] "Route to Bugs: Analyzing the Security of BGP Message Parsing," Last accessed 10-November-2025. [Online]. Available: https://i.blackhat.com/BH-US-23/Presentations/US-23-dosSantos-Route-to-Bugs-Analyzing-the-Security-of-BGP.pdf
[19] D. Dugal, C. Pignataro, and R. Dunn, "Protecting the Router Control Plane," RFC 6192, Mar. 2011. [Online]. Available: https://www.rfc-editor.org/info/rfc6192
[20] T. Albakour, O. Gasser, and G. Smaragdakis, "Pushing alias resolution to the limit," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023, pp. 584–590.
[21] Z. Durumeric, E. Wustrow, and J. A. Halderman, "ZMap: Fast Internet-wide scanning and its security applications," in *22nd USENIX Security Symposium*, 2013.
[22] "ZGrab2," Last accessed 11-May-2026. [Online]. Available: https://github.com/zmap/zgrab2
[23] J. Czyz, M. Luckie, M. Allman, and M. Bailey, "Don't forget to lock the back door! A characterization of IPv6 network security policy," in *Network and Distributed Systems Security (NDSS)*, 2016.
[24] "Will Network Devices Reject BGP Sessions from Unknown Sources?" Last accessed 16-December-2025. [Online]. Available: https://blog.ipspace.net/2023/10/reject-unknown-bgp-session/
[25] IANA, "Capability Codes," Last accessed 11-February-2026. [Online]. Available: https://www.iana.org/assignments/capability-codes/capability-codes.xhtml
[26] N. G. Shetty, C. K. Ojha, R. Katsuri, V. S. Rajaram, G. Krishna, and V. B. Ramachandra, "TCP proxying of network sessions mid-flow," U.S. Patent 9,438,699 B1, Sep., 2016, issued September 2016.
[27] Juniper Networks, "Flow-based and packet-based processing user guide for security devices," Last accessed 03-March-2026. [Online]. Available: https://www.juniper.net/documentation/us/en/software/junos/flow-packet-processing/topics/topic-map/security-srx-devices-processing-overview.html
[28] L. Izhikevich, R. Teixeira, and Z. Durumeric, "LZR: Identifying unexpected internet services," in *30th USENIX Security Symposium (USENIX Security 21)*, 2021, pp. 3111–3128.
[29] IPinfo, "IPinfo - IP Address Data for Developers & Enterprises," 2026, accessed: 2026-03-24. [Online]. Available: https://ipinfo.io/data/ip-asn
[30] A. Retana, D. R. McPherson, R. White, and V. Fuller, "Using 31-Bit Prefixes on IPv4 Point-to-Point Links," RFC 3021, Dec. 2000. [Online]. Available: https://www.rfc-editor.org/info/rfc3021
[31] Y. Rekhter and T. Li, "An Architecture for IP Address Allocation with CIDR," RFC 1518, Sep. 1993. [Online]. Available: https://www.rfc-editor.org/info/rfc1518
[32] M. Luckie, B. Huffaker, k. Claffy, V. Giotsas, and R. Oliveira, "AS relationships, customer cones, and validation," in *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 2013, pp. 243–256.
[33] Cisco, "IP Routing: BGP Configuration Guide - Removing Private AS Numbers from the AS Path in BGP," Last accessed 11-March-2026. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bgp/configuration/xe-16/irg-xe-16-book/removing-private-as-numbers-from-the-as-path-in-bgp.html?utm_source=chatgpt.com
[34] "BGP 4-Byte AS Numbers — Junos OS — Juniper Networks," Last accessed 20-January-2026. [Online]. Available: https://www.juniper.net/documentation/us/en/software/junos/bgp/topics/topic-map/4-byte-as-numbers.html
[35] "Solved: BGP Private ASN use cases - Cisco Community," Last accessed 19-January-2026. [Online]. Available: https://community.cisco.com/t5/routing-and-sd-wan/bgp-private-asn-use-cases/m-p/4854852
[36] T. Albakour, O. Gasser, R. Beverly, and G. Smaragdakis, "Third time's not a charm: exploiting SNMPv3 for router fingerprinting," in *Proceedings of the 21st ACM internet measurement conference*, 2021, pp. 150–164.
[37] "Private Enterprise Numbers (PENs)," Last accessed 22-December-2025. [Online]. Available: https://www.iana.org/assignments/enterprise-numbers/
[38] X. Dimitropoulos, D. Krioukov, M. Fomenkov, B. Huffaker, Y. Hyun, k. claffy, and G. Riley, "AS relationships: inference and validation," vol. 37, no. 1, p. 29–40, Jan. 2007. [Online]. Available: https://doi.org/10.1145/1198255.1198259

APPENDIX A
ETHICAL CONSIDERATIONS

Our active measurements (see Sections IV-A and IV-B) follow the best practices outlined in [21]. We limited the probing rate and used a blocklist. Our probe packets comply with TCP and BGP standards to maintain network stability. We conducted our measurements from a dedicated server with an informative reverse DNS record. We provided contact information to address questions or opt-out requests. We received only one request to be excluded from our scans. Our active scans were approved by our Institutional Review Board (IRB), ensuring that our methods meet ethical standards and minimize potential harm to network operators.