

In the loop: A measurement study of persistent routing loops on the IPv4/IPv6 Internet

Markus Maier^{a,*}, Johanna Ullrich^{a,b}

^a Research Group Security and Privacy, Faculty of Computer Science, Universität Wien, Kolingasse 14-16, 1090 Wien, Austria

^b SBA Research, Floragasse 7, 1040 Wien, Austria

ARTICLE INFO

Keywords:

Distributed denial of service
IPv4/IPv6 Internet transition
Internet Protocol
Internet routing
Misconfiguration
Autonomous systems

ABSTRACT

Routing loops forward packets over the same set of routers again and again. The packets, if caught in such a loop, do not only miss the intended destinations, but might also congest links between or even overload involved routers and render additional destinations whose (non-looping) paths include these routers unreachable. Given their potential impact on performance and availability, the situation of long-lasting (persistent) routing loops in today's Internet is unknown. Comprehensive measurement studies of this phenomenon (in the IPv4 Internet) date back to 2005; studies considering the successor protocol IPv6 – already accounting for more than one third of the Internet's total traffic – are lacking.

In this paper, we conduct a comprehensive measurement study to determine the status quo of persistent routing loops in today's Internet — the first-ever considering IPv6, and the first for IPv4 since 2005. We carefully extended the methodology from 2005, including multiple successive measurements, and adapted it for IPv6. Our results reveal that routing loops are still a matter of concern: in total, we found 23,208 persistent loops in IPv4 and 30,090 in IPv6, rendering 0.91% (IPv4), resp. 2.20% (IPv6), of the current Internet – as announced in BGP – unreachable. Another 7.18% (IPv4), resp. 23.00% (IPv6), are at risk to become unreachable in presence of an attack, yielding an overall higher threat potential for the IPv6 protocol. In comparison to the 2005 study, the situation has become more complex: As a consequence of IPv4 address scarcity, the number of ex ante unreachable addresses has decreased by 19.81% (despite the fact that the number of BGP announced addresses has more than doubled); at the same time, the number of addresses endangered by persistent routing loops has sharply increased (+1,907.58%) due to individual routers serving more addresses.

1. Introduction

Routing loops forward packets over the same set of routers again and again, caused by inconsistencies in the Internet's routing configuration. Two types of routing loops can be distinguished. (I) *Transient routing loops* are caused by the convergence of routing protocols. As a consequence of topology changes (e.g., outage of a network link), a router's forwarding table is automatically updated and the respective information are propagated to other routers. This takes time; routers might have inconsistent views of the network and cause routing loops. As soon as the routers' configuration is consistent again, the routing loops are remedied. Transient routing loops arise regularly on the Internet; however, they last for relatively short periods of time (usually below one minute). (II) *Persistent routing loops*, caused by (manual) configuration errors, remain for longer periods of time, negatively impacting network traffic for months or even years.

Routing loops are highly undesired artifacts, particularly when long-lasting, due to their negative effects on performance and availability, as depicted in Fig. 1: First, Network S is shadowed, i.e., packets towards its addresses loop over routers A and B before eventually being discarded. Consuming the routers' or the connecting link's capabilities, the loop imperils Network I which is reachable over router B. If much traffic gets caught in the loop, the router A/B or their connecting link is overwhelmed – rendering the destinations in Network I unreachable – or, in a milder form, packet forwarding takes more time.

Although their potential impact is significant, the current situation of routing loops on the Internet remains unknown: The last comprehensive measurement study by Xia et al. [1] dates back to 2005. Since then, the Internet and its population have significantly changed, and the study thus has to be considered outdated. In 2019, Rütth et al. [2] investigated error messages that had been collected as a byproduct of

* Corresponding author.

E-mail addresses: markus.maier@univie.ac.at (M. Maier), johanna.ullrich@sba-research.org (J. Ullrich).

URLs: <https://informatik.univie.ac.at/> (M. Maier), <https://www.sba-research.org> (J. Ullrich).




Fig. 1. Shadowed and imperiled networks: Traffic towards network S is caught in a routing loop – rendering S unreachable – and might overwhelm router B, additionally rendering network I unreachable.

regularly conducted Internet-wide scans. Among others, this included ICMP Time Exceeded messages, which indicate the existence of routing loops. Li et al. [3] developed a new way to discover IPv6 periphery routers, and by occasion found more than 128,000 routing loops at the Internet edge. Recently, Nosyk et al. [4] demonstrated how special routing loops (involving a middlebox as one hop) combined with DNS-based DDoS attacks, could result in amplification factors of up to $6 \cdot 10^8$.

Such results strongly suggest that loops still play a role on today's Internet, but further details on their persistence and characteristics were beyond these studies' scope. Also, previous measurements – with the exceptions of Li et al. [3] and Nosyk et al. [4] – only considered the IPv4 Internet, although the successor protocol IPv6 already accounts for more than one third of the Internet's total traffic [5]. These numbers are expected to increase further in the future.

In this paper, we conduct a comprehensive measurement study, including both protocols, IPv4 and IPv6, to determine the status quo of persistent routing loops. We carefully extend the approach of Xia et al. [1], encompassing multiple successive measurements and adapting it for the characteristics of IPv6. We collect data through more than eight weeks of active measurements, thus, for the very first time, shedding light on routing loops in the IPv6 Internet, and highlighting differences between IPv4 and IPv6. Beyond, we were able to trace back the developments of IPv4 routing loops over time by comparing our results with those from Xia et al. [1]. Our results show that routing loops remain an open challenge for both protocols. In total, we found 23,208 (IPv4) resp. 30,090 (IPv6) persistent routing loops, rendering 0.91% resp. 2.20% of the address space, as currently announced in BGP, unreachable. Another 7.18% (IPv4) resp. 23.00% (IPv6) are at risk to become unreachable when under attack.

Through this paper, we make the following contributions:

- We expanded the methodology by Xia et al. [1] to meet current requirements of Internet measurements (e.g., with regard to tools, measurement rates, ICMP rate limiting) and to facilitate IPv6 measurements for the very first time. Thereby, it was our goal to retain comparability between IPv4 and IPv6, and the 2005- and 2022 measurement results.
- We performed the first-ever measurement study on persistent routing loops on the IPv6 Internet, and the first comprehensive campaign on IPv4 since 2005. For this purpose, we collected data over a period of more than eight weeks, and from three vantage points in Vienna, Sydney, and Virginia.
- We developed analysis capabilities to handle the massive amount of IPv6 data (1.8 TB) and conducted a thorough analysis on the nature of today's routing loops. Therefore, we correlated our data with other data sets like CAIDA's *Routeviews IP Prefix to AS* data set [6,7], *PeeringDB* [8] and *ASdb* [9], facilitating previously unknown ways of evaluation and presentation of results.

- We revealed high dynamics of routing loops in the IPv6 Internet that are caused by the more generous address assignment practices in IPv6, i.e., enabling the assignment of public address(es) for each device. In this context, we discussed whether the current definition of persistent routing loops – i.e., mandating the same number of involved routers, and congruence of their IP addresses over time – should be modified to account for changed behavior in IPv6.
- We concluded that a simple mitigation of persistent routing loops is possible in most cases. More than 90% of the persistent loops are within a single autonomous system (AS), i.e., the involved routers are under joint control. This assigns clear responsibilities for the mitigation of persistent routing loops, which should be straightforward as no external coordination among different organizations is needed.

The remainder of this paper is structured as follows: Section 2 discusses related work. Section 3 defines persistent routing loops and further terminology used in our work. Section 4 describes our methodology, including its limitations. It is followed by Section 5 identifying persistent loops in the collected data, providing an overview of our results and a discussion of the high dynamics in IPv6 routing loops. Section 6 analyzes the characteristics of persistent loops in more detail. Section 7 maps our data with additional data sets, namely the *Routeviews IP Prefix to AS* data set, *PeeringDB* and *ASdb*, to gain further understanding of persistent routing loops and their potential root causes. Section 8 discusses our findings, and Section 9 concludes this paper.

2. Related work

Paxson [10] performed traceroute measurements using 37 Internet sites to investigate end-to-end routing behavior on the Internet. In the years 1994 and 1995, 0.13% resp. 0.16% of the measurements revealed loops, out of which 50% were observed for at least ten hours. In 2002, Hengartner et al. [11] conducted an offline analysis of four tier-1 Internet backbones and identified routing loops based on packet replicas crossing the same link multiple times with decreasing TTL. This way, 4318 transient, i.e., short-lived, routing loops were detected. Xia et al. [1] complemented this approach by focusing on persistent, i.e., long-lived, routing loops and conducted tracerouting for each /24 network prefix announced in BGP in 2005. The study shed light on routing loop characteristics such as length, location, hazardous potential, and root causes. R  th et al. [2] analyzed ICMP Time Exceeded responses that were received as a by-product in regularly conducted *ZMap* scans and performed traceroutes to the respective destination addresses. Li et al. [3] developed a methodology for the detection of routers at the Internet periphery, and found more than 128,000 routing loops at the IPv6 Internet's edge; however, the lifetime of these loops is unclear. Recently, Nosyk et al. [4] demonstrated how special routing loops, involving a middlebox as a hop, in combination with DNS-based DDoS attacks could result in amplification factors of up to $6 \cdot 10^8$; through Internet measurements, the authors identified 115 such loops in IPv4. These results suggest that loops still play a significant role in today's Internet – an assumption that is further supported by recent approaches on real-time loop detection [12], however, further details on their persistence and characteristics were beyond these previous studies' scope. Summarizing, preliminary work is either more than 15 years old [1,10,11] – and most likely outdated due to the Internet's evolution – or more recent [2,4] but not providing a detailed analysis.

From a methodological perspective, our approach appears similar to certain CAIDA measurement initiatives; however, we could not use their publicly provided data sets for multiple reasons: The *IPv4 Routed /24 Topology Dataset* [13] traceroutes a random address in each /24 prefix and is, thus, comparable to our full scan. The distinct address, which is probed in tracerouting changes, however, from measurement

to measurement and interferes with the identification of persistent routing loops, requiring multiple traceroutes to the very same address. This drawback has already been considered in the context of exploiting persistent loops in order to detect the absence of source address validation [14].

The regularly collected *Ark IPv6 Topology Dataset* [15] is too coarse-grained as only two random addresses per BGP-announced prefix are probed. As the number of probes is independent of the prefix length, it leads to varying degrees of granularity. The *IPv6 Routed /48 Topology Dataset* [16], probing each /48 network, is comparable to our full scan, but has been collected only once in 2014/15. Considering the rapid deployment of IPv6 in recent years, it has to be considered outdated. Furthermore, it does not allow to check for routing loops' persistence due to lacking additional measurements.

3. Terminology

This section provides an overview on routing, routing loops and related terminology. For the sake of comparability with previous work, our terminology is consistent with Xia et al. [1].

Routing loops: From source to destination, an IP packet traverses a set of routers $R = (r_1, r_2, r_3, \dots, r_N)$. If a router r_i appears at least twice in this set, i.e., $r_i = r_j$ for $i \neq j$, the path contains a routing loop $L = (r_i, r_{i+1}, \dots, r_{j-1})$ of length $j - i$. Routing loops might be transient, i.e., short-lived, or persistent. Either way, packets do not reach the intended destination and travel over the same set of routers again and again.

Shadowed addresses: If traffic towards a destination address d_S is caught in a routing loop L , we refer to d_S as a “shadowed address”. Routers which are included in a routing loop towards address d_S might also be involved in successfully forwarding traffic towards other addresses.

Imperiled addresses: If a destination address d_I is reachable via the routers $R = (r_1, r_2, r_3, \dots, r_N)$ and one of these routers $r_i \in R$ is also part of a routing loop towards another address, i.e., $r_i \in L$, address d_I is considered an imperiled address. In principle, imperiled addresses are reachable but bear the risk of DoS. Traffic towards the shadowed address d_S is caught in the loop L and might overwhelm the involved routers, including router r_i . The latter is involved in path R , eventually rendering d_I unreachable.

Dark addresses: A shadowed address d_S which can be used to attack an imperiled address d_I is also referred to as “dark address”.

TTL and Hop Limit: IPv4 and IPv6 define a TTL (Time To Live) resp. Hop Limit field in their protocol header to prevent packets from endlessly cycling in routing loops. Their value is chosen by the packet's sender and decremented by 1 at each forwarding router. When its value reaches 0, the packet is discarded. The highest possible value is 255; however, operating systems tend to use lower values by default. In case a router discards a packet, it returns an ICMPv4 or ICMPv6 Time Exceeded message to the sender [17,18]. According to the RFCs, this behavior is optional for IPv4/ICMPv4 and mandatory for IPv6/ICMPv6. TTL (IPv4) and Hop Limit (IPv6) serve the same purpose; thus, we use these terms interchangeably.

Tracerouting: For network diagnosis, the path from source to destination is traced by leveraging the behavior described above. The source sends a message towards the destination with a TTL value of 1. The first router along the path decrements the value to 0 and responds with a Time Exceeded message revealing its interface address. Then, the source sends a message with $TTL = 2$ triggering the second router to respond, and so forth.

Traceroute tools are provided as part of the operating system, but these implementations are typically stateful, i.e., they wait for a




Fig. 2. Pull-Up route: Following its default route, router B returns traffic to addresses outside of 2001:db8:100:1::/64 to router A, causing a routing loop. A pull-up route would discard this traffic.

response (or timeout) for $TTL = n$ before sending the packet for $TTL = n + 1$. Hence, measurements take time. For large-scale measurements, stateless alternatives like *yarrp* [19,20] are preferred. They infer all necessary information from the ICMP replies, allowing packet rates of up to 100,000 pps. Additionally, destination addresses and TTL values are randomly permuted to prevent overloading of individual routers, and Paris tracerouting [21] is implemented. This ensures that all requests of a trace take the same path, even when encountering flow-based load balancers.

ICMP rate limiting: Routers are limited in the number of self-originated ICMPv4/ICMPv6 messages [18]. Typical values are 10, 100 and 1000 pps [22]. This behavior impairs high-speed tracerouting as it triggers high amounts of ICMP Time Exceeded messages.

Pull-up route: A pull-up route (or Null route) discards packets towards a defined prefix and might prevent routing loops. In Fig. 2, router A forwards traffic towards prefix 2001:db8:100::/48 to B. However, router B is only configured for the longer prefix 2001:db8:100:1::/64. In absence of a pull-up route, a packet towards an address outside of the configured /64 prefix arrives at router B, and will – due to the configured default route – be returned to router A, causing a routing loop. In case of a configured pull-up route at router B, such a packet will be discarded.

IPv4 and IPv6: The main difference between IPv4 and IPv6 is the address length (IPv4: 32 bit, IPv6: 128 bit). While the functionality of routing is in principle consistent among the different versions, address length has an impact on global address planning and assignment. IPv4 suffers from address scarcity, thus multiple hosts usually share a single global address to connect with the Internet. In contrary, IPv6 assigns end sites (e.g. private households) at least a /64 network prefix, at least; more common are /48 for business and /56 for residential customers [23]. Thanks to these 2^{64} addresses, each and every device is assigned its own globally reachable address. This generous practice aims at compact IPv6 routing tables and is intended to lighten the operational burdens of routers.

4. Measurement design

Our methodology, as depicted in Fig. 3, has *exhaustive tracerouting* of the Internet as the first step. In a next step, the identified routing loop candidates are checked for their persistence over time. Confirmed persistent routing loops are then fed back into the analysis of the *exhaustive tracerouting* data set in order to identify shadowed and imperiled networks. Further measurements investigate if the persistent routing loops are (i) observable from different vantage points, and (ii) also found towards other addresses in the investigated prefixes. A measurement timeline is shown in Fig. 4, and an overview of the measurement dates in Table 1.

This section provides details on *exhaustive tracerouting* (Section 4.1), and persistence measurements (Section 4.2). Next, we discuss our method of loop detection (Section 4.3), our input data set as inferred




Fig. 3. Measurement design: Persistent loops are fed back into the analysis of the T_E (exhaustive tracerouting) data set to identify shadowed and imperiled networks.




Fig. 4. Measurement timeline: The measurement campaigns were conducted from three vantage points in Vienna, Sydney, and Virginia.

Table 1

Measurement dates: The measurement campaigns were conducted in spring 2022 over a period of more than eight weeks.

	Protocol	Start date	End date
Exhaustive scan	v4	2022-02-17	2022-02-17
	v6	2022-02-07	2022-02-26
Persistence over time	v4/v6	2022-03-17	2022-03-21
Persistence over location	v4/v6	2022-03-17	2022-03-21
Persistence over prefix	v4/v6	2022-04-05	2022-04-07

from BGP (Section 4.4), the maximum measurement rate (Section 4.5), the measurement campaign's limitations (Section 4.6), and eventually provide a comparison of our methodology with Xia et al. [1] (Section 4.7).

4.1. Exhaustive tracerouting

In our first step, the routed address space is exhaustively tracerouted at a certain prefix granularity. We decided to measure one random address per /24 (IPv4) resp. /48 (IPv6) prefix for the following reasons: (I) It is best practice that /24 [24] resp. /48 [25] are the most specific prefixes which are announced by BGP, and, indeed, (II) an analysis of BGP announcements – as of 2022-02-05 and provided by CAIDA, with their public Prefix to ASN list – shows that only a minority of the announced prefixes are more specific (IPv4: 0.66%, IPv6: 1.14%); and (III) our IPv4 measurements comply with Xia et al. [1] and allow to compare the current state of the Internet with its state in 2005.

The data set resulting from *exhaustive tracerouting* is referred to as T_E . We rely on *yarrp* [19,20], a high-speed measurement tool for tracerouting which operates in a stateless and asynchronous way. *Yarrp* is available for both IP versions and permutes destination IP addresses and TTL values to prevent router overloading; beyond, it implements Paris tracerouting to ensure that all packets of a trace follow the same path, even in presence of flow-based load balancing. The measurements were conducted from our local measurement server, having a dedicated 10Gbit connection and residing in a Vienna data center. In total, we measured 20 days for IPv6, and a single day for IPv4.

4.2. Persistence measurements

Persistence measurements investigate whether the identified persistent routing loop candidates from T_E are persistent over time and location. In addition, it is measured if persistent loops are also prevalent towards other addresses in the same /24 resp. /48 prefix. Therefore, the following additional measurements were conducted.

Persistence over time: For each routing loop candidate which was identified in T_E , we choose up to five destination addresses that have experienced this loop and included them into this persistence measurement. If a loop shadowed more than five addresses in *exhaustive tracerouting*, we randomly choose five addresses from this set of addresses. The identified addresses are then tracerouted – again using *yarrp* – in 12-hour interval, i.e., at 12am and 12pm on five consecutive days, leading to a total of ten measurements from our vantage point in Vienna.

The idea behind this measurement is that temporary loops as caused during routing protocol convergence are typically gone after a few minutes [11]. If the loops are, however, active over multiple days, they can be considered persistent. Precisely, we consider loops to be persistent if they appear in *exhaustive tracerouting* and in all ten rounds of the *persistence over time* measurements. All loops found in these measurements form data set P_T ; this is independent of their persistence.

Persistence over location: Complementing the measurements from Vienna, the *persistence over time* measurement were also conducted at two additional vantage points in Sydney and Northern Virginia, both instances in the Rackspace.¹ cloud The measurements were shifted in time to prevent mutual interference of our measurements, preventing routers to run into ICMP rate limits. The measurements in Virginia started at 1am and 1pm, and Sydney at 2am and 2pm. These measurements allowed to investigate whether the same routing loops are tangible from different points of the Internet. In this measurement, it is our goal to determine whether a loop is also observable from another location. Thus, it is sufficient if a loop appears once to confirm its visibility from a distinct vantage point.

Persistence over prefix: A final measurement investigated whether persistent routing loops are not only present to the single random address that has been probed in the previous measurements, but also to other addresses in the respective /24 (IPv4) and /48 (IPv6) prefixes. Therefore, an additional 50 addresses were randomly generated for each shadowed prefix and probed in a single measurement. Apparently, an IPv6 /48 prefix allows more addresses than a /24 in IPv4. However, probing 50 out of 256 IPv4 address is roughly comparable to probing 50 out of 256 /56 prefixes in IPv6: /56 prefixes are typically assigned to residential customers [23, 4.2.2]; in IPv4, such a residential customer receives a single address.

Due to time constraints, we decided to use the network scanner *ZMap* [26] instead of *yarrp*. *ZMap* was modified to include *ICMP Time Exceeded* messages in the evaluation. The rationale behind this decision is as follows: As we are solely interested in whether packets towards the additional addresses are caught in a loop, a single request is sufficient. In case of a loop, a Time Exceeded message will be returned. With this method, it is only feasible to detect presence of a loop; however, it is not guaranteed that it is the same as prevalent towards the other shadowed destinations – measured in *exhaustive tracerouting* or *persistence over time* measurements – in this prefix.

¹ <https://www.rackspace.com/>

4.3. Loop detection

Packet handling and loop detection is the same for *exhaustive tracerouting*, *persistence over time measurements* and *persistence over location measurements*: Yarrp stores all responses, providing a full picture of responses; thus, it is also feasible to determine hops that did not reply. If traces encompass Echo Replies, indicating reachability of the probed address, or ICMP Destination Unreachable messages, indicating the absence of a prefix or address, we exclude them from further processing. From the remainder traces, we check for addresses that appear multiple times at different TTL values, and define their first occurrence as the beginning, and the address before their second occurrence as the end of the loop. If we cannot identify all addresses (e.g. due to a missing Time Exceeded message) of a loop, we exclude the respective trace from our analysis. The same holds for traces showing new IP addresses after a loop. Loops are considered equivalent if they have the same length, i.e., the number of involved router addresses remains constant, and include the same router addresses.

For *persistence over prefix*, we have adapted ZMap to store the probed destination address in combination with the information on the receipt of a Time Exceeded message. If such a Time Exceeded message has been received, we infer that there is a loop towards the probed destination. In principle, a Time Exceeded message could also be returned in case the (non-looping) path to the destination is longer than the TTL that is set by the sender. This is however unlikely as we checked typical paths lengths towards reachable nodes; they are lower than 64 — the TTL used in our ZMap measurements.

4.4. BGP data set

The IPv4- and IPv6 prefixes, serving as input for T_E , were inferred from CAIDA's Prefix to AS mapping data set [6,7] – as provided on 2022-02-05 – and processed as follows: Overlapping prefixes were identified, and the shortest prefixes were further included into processing in order to guarantee comprehensive coverage of the routed Internet. Then, all /24 (IPv4) resp. /48 (IPv6) networks within these prefixes were inferred, a random address within each prefix was generated.² Eventually, this line of action resulted in 11,996,245 (IPv4) resp. 5,500,185,205 (IPv6) measured addresses.

For IPv6, one /16 prefix dedicated to the transition technology 6to4 [27] and another eight prefixes with a length between /19 and /21 were found in the CAIDA data set, posing a significant measurement effort. Consequently, we excluded the 6to4 prefix from our measurements. For the remaining prefixes we manually reduced them by choosing four /24 prefixes to be included in the our measurements. The decision is based on responses that are received upon ICMPv6 Echo Requests; the latter are sent to the first and a random address in each /32 using ZMap [26].

The CAIDA dataset is build upon multiple collectors around the globe. If one prefix is announced by multiple ASes (e.g. for anycast functionality), this information is kept throughout the analysis, and respective addresses will be attributed to all ASes announcing the respective prefix.

4.5. Measurement rate

On the one hand, measurement rates should be high in order to finish in reasonable time; this is particularly relevant for IPv6. On the other hand, routers might run into rate limits due to high-speed measurements and refrain from sending ICMP Time Exceeded messages, negatively impacting our results. In consequence, we had to find a measurement rate balancing efficiency and rate limiting.

² We checked the generated addresses to guarantee that they are neither network addresses, first addresses in a network nor broadcast addresses.




Fig. 5. Maximum measurement rate wrt. to served prefixes: For IPv6, small routers ($l = 10$ pps) serving up to 2^{18} /48 prefixes (equivalent to /30 or longer prefix) and large routers ($l = 100$ pps) serving up to 2^{21} (equiv. to /27) are unlikely to be affected by rate limits.

First, we calculated t_{min} ; it defines the minimum measurement time that is necessary to prevent a router – responsible for forwarding traffic towards p /24 (IPv4) resp. /48 (IPv6) prefixes – from ICMP rate limiting. In case of faster measurements, the router would not send ICMP Time Exceeded messages, that are a prerequisite to detect routing loops, anymore, affecting our results in a negative way. Our measurement campaign traceroutes a single address in each /24 resp. /48 prefix (*exhaustive tracerouting*), and a router thus has to origin one Time Exceeded messages per prefix, i.e. $m = 1$ ppp.³

Typical values for rate limits l of routers are 10, 100 and 1000 pps [22]. For good measure, the router should still be able to respond to others, and we restricted ourselves to a reduced rate limit $s \cdot l$, including a safety margin s (e.g., 1%). Concluding, the minimum measurement time t_{min} is calculated as follows:

$$t_{min} = \frac{p \cdot m}{s \cdot l} \quad (1)$$

With t_{min} , we are now able to infer the maximum measurement rate r_{max} of our campaign. We traceroute n /24 resp. /48 prefixes, and tracerouting one of these prefixes requires h packets, each with a different TTL value. Consequently, the maximum measurement rate r_{max} is calculated as follows:

$$r_{max} = \frac{n \cdot h}{t_{min}} = \frac{n \cdot h \cdot s \cdot l}{p \cdot m} \quad (2)$$

Assuming $s = 1\%$ and $h = 25$, Fig. 5 shows the maximum measurement rate r_{max} in dependence of m . For IPv6, we chose r_{max} to be 50,000 pps: Small routers ($l = 10$ pps) serving up to 2^{18} /48 prefixes (equivalent to IPv6 /30 or longer prefix) are unlikely to be affected by rate limiting; the same holds for large routers (100 resp. 1000 pps) serving up to 2^{21} (equiv. to /27) resp. 2^{24} (equiv. to /24) prefixes. For IPv4, high measurement rates are not as critical due to the limited address space. We decided for 5000 pps, preventing small routers serving up to 2^{12} (equiv. to /12) prefixes and large routers serving up to 2^{15} (equiv. to /9) resp. 2^{19} /24 prefixes (equiv. to /5) from rate limiting.

4.6. Limitations

Our results on routing loops and shadowed/imperiled networks have to be considered as the lower bounds for their actual prevalence on the Internet as routers might refrain from sending Time Exceeded messages for multiple reasons: First, such messages are optional for IPv4/ICMPv4 [28]; and while they are mandatory for IPv6/ICMPv6 [18], manual tracerouting towards well-known addresses shows that certain routers do not follow this specification. Second,

³ ppp=packets per prefix.

individual routers might still run into ICMP rate limits, despite all efforts to mitigate them. This particularly holds for routers more towards the core of the Internet forwarding traffic to many destinations. At the same time, a routing loop involving such a router might be prevalent towards multiple /24 resp. /48 prefix, and one of the many traces towards these prefixes will reveal all hops of this loop.

4.7. Comparison of methodology

On the one hand, one of our goals is the comparability of our results for the IPv4 Internet with those from 2005 and, consequently, to remain compliant with the methodology of Xia et al. [1]. On the other hand, we aim to comprehensively measure the IPv6 Internet for the very first time and therefore have to adapt the existing methodology to handle the latter's protocol characteristics. The remainder paragraphs highlights the differences between the methodology of Xia et al. and ours:

(I) Xia et al. measured the first address and a random address in each prefix. As overall measurement time would exceed a month for IPv6, we decided to measure a single random address per prefix. While the methodology of Xia et al. is prefix-centric, our approach is loop-centric as the following differences emphasize: (II) For the *persistence over time measurements*, Xia et al. re-measured all prefixes experiencing a routing loop to check for their persistence. For determination of persistence it is however sufficient to rediscover the loop towards a single prefix, motivating the reduction of probed prefixes. Thus, we restricted re-measurement to only five prefixes per unique loop instead of all its shadowed prefixes, again significantly reducing our measurement effort. (III) We continued likewise for *persistence over location measurements* as it is sufficient that a loop reappears at least once in the measurements from another vantage point. Thus, we measured five prefixes per unique loop from additional the two vantage points; whereas, Xia et al. measured four addresses in 4894 prefixes, that were (randomly) chosen among the total of 135,973 prefixes, from an additional four vantage points.

For *persistence over prefix measurements*, we measured – like Xia et al. – 50 additional addresses per prefix. However, we refrained from tracerouting them and instead relied on ZMap scanning. If a Time Exceeded message is received, a loop is considered to be existent towards the probed destination – though, we cannot guarantee that it is the same loop as discovered in the previous steps of our measurement. Yet, this line of action allows us to measure all shadowed prefixes, Xia et al. – relying on tracerouting these 50 additional addresses – remeasured only 3705 prefixes from a total of 135,973.

Xia et al. did not explicitly state the distinct tool used for tracerouting. From the statements on non-responding routers, we assume that the authors used tracerouting tools as integrated in major operating systems like Linux or Windows. In comparison, we used yarrp [19]. It is stateless facilitating faster scanning, shuffles destination addresses to prevent routers from overloading, and also implements Paris tracerouting to consider flow-based load balancing appropriately.

For exhaustive tracerouting, Xia et al. measured 11 million traces in 16 days resulting in a measurement rate of approx. eight traces per second. We measured with 5000 pps resp. 50,000 pps, resulting in approx. 200 resp. 2000 traces per second, in order to finish the measurements within reasonable time, particularly relevant for the high amount of IPv6 addresses. In return, we have to include ICMP rate limiting behavior of routers, effectively limiting the number of Time Exceeded messages originated by such node, in our methodology, see Section 4.5.

Table 2

Synopsis on results: Since 2005, the number of shadowed IPv4 prefixes has decreased by 19.81%, while imperiled prefixes increased by 1,907.58%, enhancing risk of individual prefixes.

	Xia et al. [1]	Our measurements	
Measurement campaign			
Year	2005	2022	2022
Protocols	IPv4	IPv4	IPv6
Prefix granularity	/24	/24	/48
Destination prefixes	5,499,518	11,996,245	5,500,185,205
Results			
Shadowed prefixes	135,973	109,178	121,234,603
Imperiled prefixes	42,887	860,991	1,265,045,930
Routing Loops	–	23,208	30,090
Involved routers	–	42,035	40,565




Fig. 6. In comparison to IPv4, IPv6 networks are at a higher risk to be shadowed (IPv4: 0.91%, IPv6: 2.20%) and imperiled (IPv4: 7.18%, IPv6: 23.00%).

5. Persistent routing loops

As defined by Xia et al. [1], a persistent loop appears in *exhaustive tracerouting* (data set T_E), as well as in all of the ten *persistence over time measurements* (data set P_T). Loops are considered equivalent if they have the same length, i.e., the number of involved router addresses remains constant, and include the same router addresses. Traces containing unknown routers – a consequence of not returning ICMP Time Exceeded messages – are excluded from our analysis.

Synopsis of the results: The analysis of T_E revealed 34,971 (IPv4) resp. 161,284 (IPv6) persistent routing loop candidates that were further checked for persistence by tracerouting them an additional ten times over a period of five days. In total, 23,208 (IPv4) resp. 30,090 (IPv6) routing loops were found to be persistent. These persistent loops were fed back into the analysis of T_E to determine the number of shadowed and imperiled prefixes. 109,178 (IPv4) resp. 121,234,603 (IPv6) prefixes were shadowed, and 860,991 (IPv4) resp. 1,265,045,930 (IPv6) imperiled. This means that 0.91% (IPv4) resp. 2.20% (IPv6) of the destination address space remained unreachable (“shadowed”) from our vantage point, and 7.18% (IPv4) resp. 23.00% (IPv6) were threatened by routing loops prevalent on the Internet (“imperiled”). In total, 42,035 (IPv4) resp. 40,565 (IPv6) unique routers were involved in the persistent loops. An overview encompassing our results and the results from 2005 by Xia et al. [1] is provided in Table 2. Fig. 6 depicts the relative share of shadowed and imperiled networks in the IPv4- resp. IPv6 Internet.

IPv4 development since 2005: According to the BGP announcements, the number of IPv4 destination networks, which are serving as input to the measurements, has increased by 118.13% since 2005. Despite this growth, our results show that today's amount of shadowed prefixes is 19.81% lower than in 2005. The imperiled networks have increased by 1,907.58%. The latter implies a greatly enhanced risk for individual IPv4 prefixes to be subjected to adversarial misuse of routing loops. While 0.78% of the announced prefixes were at risk to be imperiled in 2005, this share is now at 7.18%. Unfortunately, Xia et al. [1] neither




Fig. 7. 81.50% (28,500/34,971) of the IPv4 loop candidates were found again in the persistence measurements; for IPv6, this holds for only 25.55% (41,211/161,284).

provide a number on unique persistent loops nor on involved routers; thus, no statements on their development over time are possible.

The reason for this development might be the Internet's growth rate, accompanied by IPv4 address exhaustion. First, this might motivate a sparse use of IPv4 addresses for hosts and services, guaranteeing their reachability and effectively reducing the share of shadowed addresses. Second, routing tables might have become larger and more complex due to address ranges that were not announced in 2005, but also due to their more fine-grained assignment. In consequence, individual loops might nowadays affect more addresses than in 2005.

Comparison IPv4/IPv6: For both protocols, the number of persistent loops (23,208 vs. 30,090) and involved routers (42,035 vs. 40,565) remains in the same order of magnitude. This is insofar outstanding as the amount of investigated destination networks differs by a factor of 458.49 (IPv4: 11,996,245, IPv6: 5,500,185,205), and the number of persistent routing loop candidates found in T_E by a factor of 4.61 (IPv4: 34,971, IPv6: 161,284).

With regard to the destination networks, the numbers are explained by IPv6's generous address assignment practices: Individual IPv6 routers serve a vastly increased address space in comparison to IPv4; however, the number of IPv6-capable hosts does not excessively exceed their IPv4 counterparts (if at all), and large parts of the address space remains inactive.

Considering routing loop candidates, the discrepancies among protocol versions indicate higher dynamics of the IPv6 routing infrastructure, and become also apparent in the Venn diagrams (see Fig. 7), depicting the number of all loops found in T_E , those found in the P_T , and their overlap.⁴

For IPv4, 81.50% (28,500/34,971) of the loops from T_E were also found in P_T ; only 2,506 additional loops were found in P_T . For IPv6, the picture is different: First and foremost, only 25.55% (41,211/161,284) of the loops in T_E were found again in P_T ; beyond, P_T revealed an additional 62,633 loops which had not been present in the measurement before. The latter number is 51.98% larger than the actual overlap. As P_T probed only target destinations which had already been shadowed in T_E , our numbers indicate that multiple loops towards the same destinations were found.

We identified three potential reasons for these dynamics: (I) In comparison to its predecessor, IPv6 still experiences steady deployment and reconfiguration of networks and hosts, leading to continuous changes in the routing infrastructure. (II) Devices which shared a global IPv4 address are assigned globally reachable IPv6 addresses of their own. In the case of multiple devices serving the same purpose (e.g., load balancing), the replies are returned by one or the other instance and consequently contain different source addresses, causing – according to our definition – distinct routing loops. (III) The generous IPv6 address space also motivates regular address changes over time (e.g., Prefix Rotation [29,30], Privacy Extension [31,32]). This leads to a situation in which responses of a single host contain different sources addresses, seemingly causing multiple loops in succession.

⁴ The overlap includes all loops that were seen at least once in the persistence measurements; thus, it exceeds the number of persistence loops that have to be prevalent ten times.




Fig. 8. IPv6 loops and persistent loops wrt. measurement time (*exhaustive tracerouting*). The black bar indicates the time span of the IPv4 measurement.

Routing loop dynamics in IPv6: Fig. 8 shows the cumulative amount of total and persistent loops over the time span of *exhaustive tracerouting*. Apart from jumps caused by a single AS, the numbers are steadily increasing for total loops and persistent loops. In this specific AS, eight routers appear to form loops with many other routers in this AS. Overall, it is apparent that IPv6 loops show high dynamics as only a fraction thereof are considered to be persistent, i.e., are alive for longer time periods.

To gain further understanding of the underlying reasons for the observed IPv6 dynamics, we investigated destination prefixes which were persistently shadowed – i.e., experienced a routing loop in T_E and P_T , but yielded multiple loops – and grouped these loops in “loop sequences”. While IPv4 reveals only 579 such sequences, we found 11,857 unique sequences for IPv6. They encompass between two (IPv4: 450, IPv6: 4511) and eleven loops (IPv4: 3, IPv6: 14), i.e., each of our probes revealed a different loop towards the destination.

For IPv6, 11,081 sequences contained only loops of the same length, and they typically only differ in a single address (9993). In 9086 (76.63%) of the cases, these differing addresses resided in the same /64, suggesting close vicinity of the replying hosts. Only 458 (3.86%) of these sequences replied with addresses differing in all 8 bytes of the interface identifier, suggesting limited deployment of the Privacy Extension or similar schemes for routers (reason III). Conversely, the addresses of 7857 (66.26%) sequences differed only in the last byte, a strong indicator for load balancing or similar behavior (reason II). Finally, only 776 (6.54%) sequences revealed loops of different lengths, suggesting manual or automatic reconfiguration of the routing infrastructure (reason I). In comparison, 507 IPv4 sequences contained only loops of the same length, typically also differing in a single address (399, 68.91%). In 242 cases (41.80%), these differing addresses resided in the same /24, i.e., vicinity among IPv4 addresses is not as prevalent as in IPv6.

Summarizing, IPv6 does not only reveal more loop sequences in comparison to the total amount of persistent loops, but, proportionally, also more of these IPv6 sequences contain close addresses (IPv4: 41.80%, IPv6: 76.63%). In most cases, the addresses only differ by a single byte. This suggests wide prevalence of load balancing in IPv6, causing high dynamics, thus, it might be reasonable to include such loops in the set of persistent loops, potentially increasing the total amount of IPv6 persistent loops by up to 30.20%. Therefore, we suggest




Fig. 9. Most loops (IPv4: 64.70%, IPv6: 58.43%) resp. routers (IPv4: 63.57%, IPv6: 55.19%) shadow only a single prefix. Only 29.23% (IPv4) resp. 36.46% (IPv6) of the routers, that are involved in loops, imperil other prefixes.

further investigations to understand these scenarios in detail and to differentiate them properly. For this work, we continue with the stricter definition of Xia et al. [1] – mandating equivalence of the involved addresses for a persistent loop – also to allow for comparability with previous work. In this context, we emphasize once more that our measurements reveal a lower bound on the prevalence of routing loops on today's Internet, see Section 4.6.

6. Routing loop characteristics

In this section, we elaborate on the characteristics of persistent routing loops in the IPv4- resp. IPv6 Internet in more detail. Where available, we contrast our results on IPv4 with those from Xia et al. [1], thereby shedding light on the evolution of the Internet since 2005.

Shadowed prefixes: Traffic towards shadowed prefixes is caught in persistent routing loops, rendering these networks unreachable. In total, we found 109,178 (IPv4) resp. 121,234,603 (IPv6) such prefixes, representing 0.91% (IPv4) resp. 2.20% (IPv6) of the investigated address space. Fig. 9 depicts the number of shadowed networks per loop resp. per router. Most loops (IPv4: 64.70%, IPv6: 58.43%) and routers (IPv4: 63.57%, IPv6: 55.19%) shadow only a single prefix. In general, IPv6 loops resp. routers affect more prefixes; the maximum amount of shadowed prefixes per loop is 27,069 (IPv4), and 12,085,972 (IPv6). For IPv6, we discovered accumulations at powers of 2, e.g., 2^6 , 2^8 , 2^{14} , and 2^{16} . As the granularity of our measurements is /48 prefixes, this suggests that IPv6 loops render entire /42, /40, /34, and /32 prefixes unreachable.

Imperiled prefixes: In our measurements, we identified 860,991 (IPv4) resp. 1,265,045,930 (IPv6) imperiled prefixes, representing 7.18% (IPv4) resp. 23.00% (IPv6) of the probed address space. Fig. 9 shows the number of imperiled prefixes per routing loop resp. router. 52.87% (IPv4) resp. 30.20% (IPv6) of the persistent routing loops do not imperil a single address. From a router perspective, these numbers are even more nuanced. 70.77% (IPv4) resp. 63.54% (IPv6) of the routers do not pose any threat; in other words, a total of only 12,288 (IPv4) resp. 14,790 (IPv6) routers threaten other destinations. At this point, we want to emphasize that these routers are not necessarily the culprit; loops arise from misconfigurations and, from our perspective, it remains unclear which router(s) involved in a specific loop is/are the root cause. For IPv6, a sharp increase appears between 2^8 and 2^9 in Fig. 9, suggesting that multiple /40 are imperiled by individual loops.

Comparison: Back in 2005, Xia et al. [1] report that only 24.1% of the persistent loops shared routers with non-looping destinations, leading to imperiled prefixes. In other words, 75.9% of the loops did not pose any danger for other destination addresses. Today, this number is significantly decreased, potentially caused by the more densely populated IPv4 Internet: only 52.87% of the IPv4 loops do not imperil




Fig. 10. 61.74% of the IPv4 loops shadow more addresses than they imperil; this only holds for 42.42% of the IPv6 loops.




Fig. 11. Densities of shadowed prefixes are higher, i.e., shadowed prefixes are closer to each other than their imperiled counterparts affected by the same loop.

other addresses. For IPv6, this number is even lower, namely 30.20%, and might be caused by the large address ranges served by individual routers.

Dark prefixes: Dark prefixes are shadowed prefixes that might be exploited in an attack to imperil addresses. We determined them by summing up the shadowed addresses of loops threatening other prefixes. In total, our measurements revealed 85,362 (IPv4) resp. 121,096,767 (IPv6) dark addresses, i.e., 78.19% resp. 99.89% of the shadowed networks actually threaten other networks.

Comparison: Xia et al. [1] report 25,019 dark addresses, representing 18.4% of the shadowed prefixes, a figure that is clearly exceeded by our measurements in 2022. The underlying reason might be a more densely populated Internet with routers forwarding traffic towards more destinations than in 2005. If one of these shadowed addresses causes a loop, many other destinations depending on the involved routers become unreachable, making the shadowed address a dark address.

Ratio of imperiled to shadowed prefixes: With n_I being the number of imperiled addresses, and n_S the number of shadowed addresses, we define r as the ratio of imperiled to shadowed addresses. For a distinctive loop, it is calculated as follows:

$$r = \frac{n_I}{n_S} \quad (3)$$

There is at least one shadowed address per loop, i.e., $n_S \geq 1$; thus, r is defined for all loops.

Fig. 10 shows the results for the persistent loops: While the majority of IPv4 loops (61.74%) shadows more prefixes than it imperils, i.e., $r < 1$, this does not hold for its IPv6 counterparts (42.42%). $r = 1$ holds for 9.50% (IPv4) resp. 3.45% (IPv6) of the loops, i.e., these loops shadow as many addresses as they imperil. Summarizing, IPv4 loops tend to shadow, IPv6 loops to imperil.

Address density: In order to see whether the /24 resp. /48 networks which are shadowed resp. imperiled by the same loop are close to each other, an address density d was defined. Through determining the number of equal bits b among these addresses, the longest network prefix encompassing all shadowed resp. imperiled addresses is identified. The ratio of the amount of shadowed/imperiled addresses n and the total




Fig. 12. Loops involving one or two routers are common for both protocols. Only 2.85% resp. 3.65% of Loops are of length 3 or longer.




Fig. 13. Amplification factor as perceived from vantage point Vienna: Traffic is typically amplified by a factor of 121.5 resp. 122 (median).

number of potential addresses in the shared prefix provides the address density. The granularity g of our measurements is /24 (IPv4) resp. /48 (IPv6), and considered accordingly in the number of feasible addresses.

$$d = \frac{n}{2^{g-b}} \quad (4)$$

If all prefixes within a $/b$ network are shadowed or imperiled, density becomes $d = 1.0$; otherwise, it is lower.

Our results are depicted in Fig. 11 and show that the densities are lower for imperiled- than for shadowed networks. This implies that networks which are shadowed by individual loops are closer than those imperiled by a loop; this holds for both protocol versions. However, the disparity is more nuanced for IPv6. According to our results, exhaustively shadowed/imperiled prefixes with $d = 1.0$ (or close to 1.0) seem rare, with the exception of shadowed IPv6 prefixes. Our results on density have to be considered as lower bounds – similarly to the amount of routing loops or shadowed/imperiled addresses (see Section 4.6 on limitations) – and actual density rates might be higher.

Routing loop length: Fig. 12 depicts the length of the identified persistent routing loops, i.e., the number of involved routers. For both protocols, loops including one (IPv4: 14.76%, IPv6: 6.15%) or two (IPv4: 82.39%, IPv6: 91.03%) router addresses are most common; the longest loops encompass 13 routers for both protocols. The distribution of lengths is roughly comparable among IPv4 and IPv6.

Routing loops render shadowed destinations unavailable and, foremost, threaten the reachability of imperiled networks. Since packets trapped in a loop traverse a router multiple times, they consume more resources than necessary, potentially overwhelming the involved routers. The amplification factor a defines how often a packet towards a shadowed address traverses such a router. It is dependent on the Hop Limit set by the sender – an adversary would aim to maximize amplification and thus choose the maximum value 255 for both protocols – the number of hops between the sender and the first router in the loop – herein referred to as loop distance l_d – and the loop length l :

$$a = \frac{255 - l_d}{l} \quad (5)$$




Fig. 14. For Time Exceeded Replies, IPv6 prefixes are typically totally shadowed (26.62%), while IPv4 fall into the extreme. They are either totally shadowed (26.06%) or show no additional shadowed addresses (27.02%).

Fig. 13 shows the persistent loops' amplification factors as observed from our vantage point in Vienna; its median value is 121.5 (IPv4) resp. 122 (IPv6), implying that a moderate traffic of 15 Mbps – typical for a 4k movie stream – results in a total of 1.83 Gbps.

Comparison: In their 2005 measurements, Xia et al. [1] identified loops encompassing between 2 to 16 routers. Loops with two hops clearly dominated (89%); another 10% had a length between 3 and 9. Most notably, however, no loops of length 1 were reported. These loops were filtered from the process as, according to the authors, only a few traces contained such cases. Our measurements clearly show a significant presence of such loops (IPv4: 3,425, IPv6: 1,852). The underlying root cause might be today's more widespread deployment of middle-boxes [33] for purposes such as firewalling, deep packet inspection, or censorship. As of 2002 such middleboxes were still considered “a recent phenomenon” [34]. Similar to our results, Xia et al. reported a typical amplification factor of 120. An alternative reason for these behavior is Multiprotocol Label Switching (MPLS).

MPLS is a label-based routing protocol. If a TTL reaches 0 at a MPLS router, the message is forwarded to the next IP-capable router, then returning an ICMP Time Exceeded messages to the sender of the original packet. Thus, multiple MPLS routers on a path might result in multiple Time Exceeded messages from the same IP address, and consequently be considered as a one hop loop in our data. An analysis of our data shows that 673 loops, representing 10.6% of all one hop loops in IPv4, that might be caused by MPLS. For IPv6, yarrp does not collect MPLS labels; consequently, we cannot provide any insights on such artifacts.

Persistence over prefix: In each prefix (/24 for IPv4 and /48 for IPv6) shadowed by a persistent loop, we probed 50 additional, randomly chosen addresses to see whether the whole prefix is affected. The receipt of ICMP Time Exceeded messages indicates the presence of a loop towards the probed address. Using the network scanner ZMap, however, we cannot check if the same loop, as observed in previous measurements, is present.

Fig. 14 shows how many ICMP Time Exceeded message are received per probed prefix, revealing notable differences between the protocols: For IPv6, prefixes returning 50 Time Exceeded messages – implying an entirely shadowed /48 – are most common; for IPv4, we see that prefixes either refrain from sending Time Exceeded messages at all – indicating that the remaining addresses in the /24 are not shadowed – or reply with 50 such messages.

Comparison: Xia et al. [1] report that in 1% of the IPv4 prefixes no additional address was shadowed, while in 67% all probed addresses were shadowed. Today's measurements are in stark contrast: the number of totally shadowed prefixes decreased (26.06%), whereas the share of prefixes with no additionally shadowed addresses increased (27.02%). Both numbers appear to reflect the Internet's growth over the last 17 years, including more fine-grained address assignment practices and a higher node density.

Table 3

Vantage points: 74.17% (IPv4) resp. 77.35% (IPv6) of the persistent loops identified in Vienna were also observed from our Sydney- and Virginia vantage points.

	Vienna	Observed in		
		Sydney	Virginia	All locations
IPv4	23,208	18,542	20,188	17,215
	100%	79.89%	86.99%	74.18%
IPv6	30,090	24,691	26,572	23,274
	100%	82.06%	88.31%	77.35%

Table 4

Vantage points: 79.66% (IPv4) resp. 84.59% (IPv6) of the prefixes that are shadowed by persistent loops identified in Vienna were also shadowed from our Sydney- and Virginia vantage points.

	Vienna	Observed in		
		Sydney	Virginia	All locations
IPv4	38,377	32,582	34,683	30,572
	100%	84.90%	90.37%	79.66%
IPv6	121,465	105,589	114,992	102,749
	100%	86.93%	94.67%	84.59%

Vantage points: From an adversary's perspective, it is beneficial to identify persistent routing loops from multiple locations on the Internet. This facilitates distributed attacks, e.g., as conducted by a botnet, and allows to easily reach high traffic volumes. We investigated whether the persistent loops observed from our vantage point in Vienna are also discernible from other locations, namely Sydney and Virginia (see Section 4.2 for more details). For our analysis, it is sufficient that a persistent loop – as found in Vienna – is observed at least once in Sydney and/or Virginia. The results are presented in Table 3: 79.89% (IPv4) resp. 82.06% (IPv6) of the Viennese loops were also observable in Sydney, and 86.99% (IPv4) resp. 88.31% (IPv6) in Virginia. A potential reason for the reduced number of loops found at other vantage points might be anycast addresses. Their prefixes are, depending on geographic location, announced by different ASes, and requests from different vantage points are routed towards different destinations. In *persistence over time measurements*, we certainly detected loops that are exclusively seen at a certain vantage point (Vienna: 3440/7815, Sydney: 2395/5337, Virginia: 2148/9794); it appears as this behavior is more common in IPv6 than in IPv4.

Comparison: Xia et al. [1] had randomly chosen 4894 shadowed prefixes from their data set and probed them from four additional vantage points. Depending on the location, between 4262 (87.02%, Europe) and 4543 (92.83%, US East Cost) prefixes were found to be shadowed from other vantage points as well.

Based on shadowed prefixes, Xia et al.'s numbers are not directly comparable with ours, as we count loops. Thus, we added Table 4 to show how many of the prefixes which are shadowed by persistent loops – as observed from Vienna – were also shadowed from the other vantage points: 84.90% of the targets from Sydney and 90.37% from Virginia, yielding comparable numbers to Xia et al. [1].

7. Organizational and regional characteristics

In this section, we combine our results with additional data sets to gain further insights. In particular, we mapped IP addresses to their autonomous systems (ASes), relying on CAIDA's Routeviews IP Prefix to AS data set [6,7] (Section 7.1), and then classify these ASes by their purpose of operation using PeeringDB [8] and ASdb [9] (Section 7.2).

Table 5

Persistent loops and involved ASes: Loops in a single AS dominate (IPv4: 89.99%, IPv6: 87.97%), evincing clear responsibilities for their existence.

	Xia et al.	IPv4	IPv6
1 AS	94.27%	91.86%	91.06%
(preceding router in same AS)	(67.06%)	(50.54%)	(48.58%)
(preceding router in other AS)	(27.21%)	(41.32%)	(42.48%)
2 ASes	5.35%	7.88%	8.59%
≥ 3 ASes	0.38%	0.26%	0.35%
Total	100%	100%	100%

Table 6

Persistent loops and shadowed/imperiled ASes (IPv4): In IPv4, loops shadowing only involved ASes and imperiling no ASes are dominant (45.65%).

		Imperil			Total
		No ASes	Invol. ASes	Other ASes	
shadow	Invol. ASes	45.65%	30.52%	6.53%	82.70%
	Other ASes	7.22%	0.59%	9.48%	17.29%
	Total	52.87%	31.11%	16.01%	100.00%

Table 7

Persistent loops and shadowed/imperiled ASes (IPv6): In comparison, loops shadowing and imperiling only involved ASes are dominant in IPv6 (55.29%).

		Imperil			Total
		No ASes	Invol. ASes	Other ASes	
shadow	Invol. ASes	28.72%	55.29%	6.61%	90.12%
	Other ASes	1.49%	0.08%	8.32%	9.88%
	Total	30.20%	55.37%	14.43%	100.00%

7.1. Autonomous systems

An autonomous system (AS) consists of networks under joint control and is uniquely identified by its autonomous system number (ASN). CAIDA's Routeviews IP Prefix to AS data set allows to map IP addresses to their AS, i.e., clustering addresses controlled by the same entity.

Loops and involved ascs: In a first step, we investigated the number of ASes that are typically involved in a loop. This allows us to determine whether routing loops are caused internally or due to interface problems among ASes. Therefore, we assigned the router addresses to ASes and counted the number of unique ASNs per loop. Loops with routers which could not be attributed to an AS were excluded from our analysis. The results are shown in Table 5: 91.86% (IPv4) resp. 91.06% (IPv6) of the loops involve only one AS. The dominance of single-AS loops clearly evinces responsibilities: While the distinct (technical) root causes for these loops remain unclear, the loops can be attributed to individual organizations which are responsible for their mitigation.

Comparison: Table 5 also includes the numbers reported by Xia et al. [1]. In 2005, 94.27% of the loops involved only one single AS. This means that the dominance of single-AS loops was already prevalent, even though their number has slightly decreased since then.

Shadowed and imperiled ASes: If a loop shadows or imperils only ASes that are also involved in the loop, both the root cause of a loop and its (potential) aftermath are related to the same authorities. The latter are thus able to remove the loop in order to prevent damage to their networks. The situation is different if affected ASes are not involved in the loop. Cause and effect are connected to different authorities, and the first might not be (fully) motivated to resolve loops in order to benefit the latter.

In our second step, we mapped the shadowed and imperiled addresses of each loop to their ASN and checked whether these numbers

Table 8

Classification error: Routers at AS borders might be incorrectly assigned, and classification errors are bound to 7.17%/8.34% resp. 0.56%/0.12%.

	Xia et al.	IPv4	IPv6
Destination AS is involved	87.44%	82.98%	90.88%
(1 address in destination AS)	3.78%	7.17%	8.34%
(≥2 addresses in dest. AS)	83.66%	75.81%	82.53%
No address in destination AS	12.56%	17.02%	9.12%
(Preceding router in dest. AS)	1.47%	0.56%	0.12%
(Preced. router not in dest. AS)	11.09%	16.46%	9.00%
Total	100.0%	100.0%	100.0%




Fig. 15. Misclassification at AS Boundaries: Assuming addresses of AS 1 on the shared link and routers responding with their inbound interface, the loop appears to involve both ASes despite being an internal loop of AS 2.

correspond with those assigned to the router addresses in the previous step, see [Tables 6](#) (IPv4) and [7](#) (IPv6) for our results. We distinguish between loops threatening (I) no ASes, (II) only involved ASes, and (III) other ASes. All loops shadow at least one single AS, including the probed destination address revealing the respective loop — this means that loops shadowing no ASes are inexistent.

On the positive side, most loops do not imperil any (IPv4: 52.87%, IPv6: 30.20%) resp. only involved ASes (IPv4: 31.11%, IPv6: 55.37%). For IPv4, loops shadowing only involved ASes and imperiling no ASes dominate (45.65%); for IPv6, those shadowing and imperiling only involved ASes represent the majority (55.29%). We assume that these results are a consequence of lacking pull-up routes in combination with the vast IPv6 address space. In the absence of pull-up routes, traffic towards inactive addresses is not filtered and might be routed back via the default route, causing a routing loop and imperiling the active part of the same network.

On the negative side, 16.01% (IPv4) resp. 14.43% (IPv6) of the loops imperil other ASes. We consider these loops to be high-risk since a sudden attack might lead to DoS of ASes which are not involved in the loop. To fix such an incident would require coordination with other stakeholders and potentially cause noticeable downtime of services.

Classification errors: [Table 8](#) shows whether loops with two hops or more involve addresses in the destination AS, i.e., in the same AS as the probed destination address. It appears that routing loops are close to the destinations: Most loops (IPv4: 82.98%, IPv6: 90.88%) include at least one address in the destination AS, 74.04% (IPv4) resp. 79.64% (IPv6) of the loops are solely involving addresses in the destination ASes.

Mapping IP addresses, as revealed by tracerouting, to ASes bears the risk of misclassification at AS boundaries [\[35,36\]](#). On the connecting link, border router interfaces are assigned addresses from a shared address range that is provided by one of the ASes. Beyond, routers differ how they choose the source address in self-originated ICMP Time Exceeded messages. While the majority of routers decide for the address of the inbound interface, a minority of 1.7% to 5.8% [\[37\]](#) prefer the outbound interface as a source address when responding. For example, the loop in [Fig. 15](#) – assuming the use of addresses of AS 1 on the shared link and routers answering with their inbound interfaces – would be classified to involve both ASes despite being an internal loop of AS 2. If only one address of an AS is involved it might be a misassigned




Fig. 16. Most AS are involved in a single loops (62.92%/53.42%), and operated two routers involved in loops (45.07%/38.72%).




Fig. 17. For both protocols, the number of imperiled addresses is higher than the number of shadowed, even though more nuanced for IPv6.

border router. If two or more router addresses in an AS are found, the chance of incorrect attribution becomes negligible. [Table 8](#) facilitates to determine the error of classification: 7.17% (IPv4) resp. 8.34% of the loops encompass only a single address in the destination AS, and might be incorrectly assigned.

Another way to assess potential misclassification is to investigate the address directly preceding a loop. If all addresses of a loop are in the destination AS, but the preceding address is in another AS, the first router in the loop might be incorrectly assigned, and might actually belong to another than the destination AS. These cases are however rare (IPv4: 0.56%, IPv6: 0.12%). The described classification errors are aligned with those in Xia et al. [\[1\]](#).

Comparison: [Table 8](#) also includes the results of Xia et al. [\[1\]](#). For IPv4, the amount of loops involving the destination AS has decreased, while the number of those with a single address in the destination AS has doubled (Xia et al.: 3.78%, our results: 7.04%), effectively also increasing the estimation of incorrect assignment. The reason might be found in the Internet's specialization (e.g., specialized transit- and content networks) – resulting in less routers operated by peripheral networks.

Proneness of ASes: We investigated whether certain ASes are more prone to routing loops than others; therefore, we depicted the number of loops resp. routers found per ASN, see [Fig. 16](#).

Typically, an AS is involved in a single loop (IPv4: 62.92%, IPv6: 53.42%) and encompasses two routers (IPv4: 45.07%, IPv6: 38.72%). There is a limited number of ASes serving high numbers of loops resp. routers. The maximum is 730 (IPv4) resp. 4,531 (IPv6) loops per ASN, i.e., these ASes are involved in 3.15% (IPv4) resp. 15.06% (IPv6) of all loops.

[Fig. 17](#) depicts the number of shadowed resp. imperiled addresses per AS; overall, the numbers of imperiled addresses is higher than of shadowed addresses. The maximum is 32,735 (IPv4) resp. 12,096,324 (IPv6) shadowed prefixes per ASN; and 76,961 (IPv4) resp. 46,304,433 (IPv6) imperiled prefixes per ASN. For imperiled IPv6 addresses, we can see a sharp increase at 2^{13} in [Fig. 17](#).

IPv4/IPv6 overlap: [Table 9](#) depicts the total number of ASes that encompass at least one router, one shadowed resp. imperiled address per

Table 9

ASes with router, shadowed and imperiled Addresses: ASes simultaneously involved in IPv4 and IPv6 are non-dominant, suggesting independence of the routing loop phenomena in IPv4 and IPv6.

	IPv4	IPv6	Both
ASes with routers involved in loops	7,470	4,264	1,634
ASes with shadowed addresses	9,585	8,228	1,727
ASes with imperiled addresses	14,630	23,955	4,792




Fig. 18. IPv4 loops are predominantly found in NSP networks (33.12%), IPv6 loops in Cable/DSL/ISP networks (44.37%) suggesting that the latter are based more towards the Internet's periphery.

protocol and their overlap. For example, 7470 ASes include at least one single IPv4 router address that is involved in a loop, and 4264 do so for an IPv6 router address. Thereof, 1634 ASes include IPv4 as well as IPv6 router addresses. Summarizing, the overlap among the protocols is apparent but not dominant. Thus, we conclude that the phenomena of IPv4 and IPv6 loops are largely independent of each other and might even arise from different root causes.

7.2. Network classification

In a next step, we assigned each AS a category describing its purpose/economic sector, relying on the publicly accessible databases PeeringDB and ASdb. With this information, we are able to determine in which type of networks loops arise resp. which type of networks are shadowed/imperiled.

Classification with PeeringDB: Fig. 18 illustrates the relative shares of network types,⁵ as provided by PeeringDB, among destination-, shadowed-, imperiled- and router addresses as well as loops for both protocols. 7.97% (IPv4) resp. 8.64% (IPv6) of the loops span multiple ASes (see Table 5) and might lead to multiple network types; the latter loops are considered in a separate category “mixed”. For ASes unavailable in the database, we added the option Unclassified.

Unclassified ASes: PeeringDB relies on self-reporting of AS operators, and reflects the network's purpose best possible. Nevertheless, it suffers a main drawback: between 15.39% and 37.01% of the total population remain unclassified, see Fig. 18. Thus, we contrasted the PeeringDB-based classification scheme with another database, namely ASdb [9], to gain further understanding of these unclassified ASes.




Fig. 19. Contrasting PeeringDB (left) and ASdb (right) classification: Unclassified ASes in PeeringDB appear to be predominantly peripheral networks (e.g. enterprise networks) not interested in peering.

Fig. 19 contrasts the classification according to PeeringDB with ASdb for all ASes appearing in our measurements in a Sankey diagram⁶. From the figure, we conclude that ASes unclassified by PeeringDB but classified in ASdb are predominantly assigned to (non-)ICT industries (52.99%) and to a lesser extent to ISPs (16.79%), i.e., the networks that remain unclassified in PeeringDB are rather found towards the Internet's edge (peripheral networks) than in its core (transit networks). Considering the purpose of PeeringDB, easing of peering decision, this makes sense as they are typically not interested in this kind of information [38]. With this additional information, we are now able to interpret Fig. 18 with regard to the position of loops or shadowed/imperiled addresses in the Internet.

Destination addresses: The relative share of network types among the destination addresses can be seen in Fig. 18. A comparison of IPv4- and IPv6 destination addresses reveals a decreasing share of addresses in NSP networks (IPv4: 29.59%, IPv6: 21.39%); in turn, the share of Cable/DSL/ISP-, Content-, Enterprise- and Unclassified ASes increases (IPv4: 21.15%, IPv6: 30.49%), reflecting the change in address assignment practices among the protocols. With IPv4, peripheral networks rarely used publicly reachable addresses, shared public addresses wherever possible (e.g., NAT [39], CGNAT [40], Name-Based Virtual Hosts serving Multiple Domains on a Web server [41]), and used private addresses internally [42]. The plethora of available IPv6 addresses allows for generous assignment practices in all parts of the Internet, decreasing the relative share of NSP-related addresses in comparison to peripheral networks.

Shadowed and imperiled addresses: An even more pronounced decrease of NSP addresses is observed for shadowed addresses (IPv4: 47.37%, IPv6: 9.34%). We assume a mix of the following underlying root causes, but cannot exactly pinpoint the extent of their contribution: First, unreachability of IPv4 addresses in peripheral networks might be detected soon as it would affect multiple services due to address sharing. With IPv6, this pressure might have been lifted, increasing the relative share of shadowed addresses in non-NSP networks. Second, IPv4/IPv6 transition appears to cause more substantial changes in peripheral networks, e.g., the removal of NAT, public reachability of all hosts, or the large number of available addresses. Lacking experience of their administrators might result in non-optimal network configurations; for example, non-existent pull-up routes for inactive addresses might result in routing loops. Finally, IPv4 address scarcity leads to fine-grained address assignments; however, BGP typically filters prefixes longer than

⁵ The options for network type in PeeringDB are Cable/DSL/ISP, Content, Educational/Research, Enterprise, Government, Network Services, Non-Profit, Not Disclosed, NSP, Route Collector and Route Server.

⁶ For readability, we combined multiple categories: (I) We combined PeeringDB's Network Services-, Non-Profit-, Route Collector- and Route Server classification – due to their limited numbers – as *others*. (II) We summarized the high number of different industries in ASdb into classes *non-ICT industry* and *ICT industry*, the latter containing all related subclasses with the exception of Internet Service Providers (ISPs) which represent a class of their own.

/24 as reported by RIPE NCC [43]. We might thus have incorrectly attributed addresses of smaller networks to their NSPs.

For imperiled addresses, NSPs again play a larger role. This is caused by their routers being involved in routing loops, also imperiling (large) parts of the NSPs' address ranges. In IPv6, the majority of such addresses is found in unclassified ASes (IPv4: 17.76%, IPv6: 37.01%) and Cable/DSL/ISPs (IPv4: 21.20%, IPv6: 26.43%).

Loops and involved router addresses: Fig. 18 shows that IPv4 loops and routers are predominantly found in NSPs (loops: 33.12%, routers: 36.94%). In IPv6, this is shifted and Cable/DSL/ISPs account for 44.37% of the loops and 47.14% of the routers. These insights suggest that loops in IPv6 are – in comparison to those found in IPv4 – located towards the edge of the Internet. It appears that this is a consequence of the increased address space, allowing publicly reachable addressing of peripheral networks. It remains, however, unclear whether these loops are new in IPv6 or analogous loops exist in internal IPv4 networks, thus remaining invisible to our measurements.

Proneness of network types: Serving as a baseline, we compared the relative shares of destination addresses with those of the shadowed-, imperiled- and router addresses to infer whether certain network types are more or less prone to routing loops (Fig. 18). Content ASes, representing 5.00% (IPv4) resp. 6.11% (IPv6) of the destination addresses, are only marginally involved in loops (IPv4: 2.11%, IPv6: 1.73%) and router addresses (IPv4: 2.34%, IPv6: 2.56%). This might indicate that the respective operators optimize their networks, effectively preventing routing loops. The opposite is observed for ISPs and NSPs. For both protocols, they account for more loops (IPv4: 29.35%, IPv6: 32.98%) than their share of destination addresses (IPv4: 26.28%, IPv6: 25.34%). At the same time, ISPs are underrepresented among the affected networks and appear to suffer less from the consequences of their loops. This also holds true for NSPs in IPv6.

8. Discussion

Our results show that persistent routing loops are still of relevance, both in the IPv4- and IPv6 Internet. While the orders of magnitude for persistent routing loops are comparable (IPv4: 23,208, IPv6: 30,090), we were able to pinpoint significant differences between the protocols. Most notably, 0.91% resp. 7.18% of the IPv4 Internet are shadowed resp. imperiled; for its successor protocol IPv6 those numbers are higher, namely 2.20% resp. 23.00%, indicating a higher threat potential. The following paragraphs discuss the most relevant findings of our measurements.

Effort of IPv6 measurements: When comparing Xia et al. [1] with our IPv4 measurements, it becomes evident that the effort necessary for conducting such measurements has significantly decreased since 2005. Back then, *exhaustive tracerouting* took 16 days to complete, whereas our measurements – including 118% more prefixes as input – finished in less than a day. IPv6 measurements, however, still pose a considerable challenge with regard to time, memory, and storage as well as require a well thought-through design.

Preventing routers from ICMP rate limiting is key to gain sound results, posing multiple challenges: First, measurement speed has to be throttled to prevent rate limits; however, excessive throttling would lead to long measurement times. Balancing the demands, *exhaustive tracerouting* effectively took 20 days. Second, the responses arrive out of order, resulting in a 1.8 TB large data set, which brings further difficulties for the analysis. The individual traceroutes have to be reconstructed from responses from all over the data set; furthermore an analysis of the imperiled prefixes per loop requires to keep sets of prefixes per router, easily exhausting the 256 GB of memory on our server. Finally, the runtime of our analysis scripts has also been a key issue. Moving the analysis from Python scripts to Rust, among other optimizations, reduced the runtime up to a factor of 20.

IPv4 address scarcity: For IPv4 addresses, the chance of being shadowed has clearly decreased since 2005, and it seems that this positive effect must be attributed to IPv4 address scarcity. As of today, each and every address is needed to fulfill the increasing network demands and frequently shared among multiple services; thus, there might be almost no potential for addresses to become shadowed (i.e., unreachable) for longer time periods.

Even though the Internet has grown significantly in the past 15+ years – reflected by the more than doubled /24 destination prefixes inferred from BGP announcements – the absolute number of shadowed prefixes has decreased by 19.71%, and we assume a similar positive effect with regards to persistent routing loops. Due to a lack of numbers on unique persistent loops and involved routers from 2005, we are not able to confirm this assumption.

Persistently shadowed by dynamic loops: In the IPv6 Internet, we discovered more than 11,000 destination prefixes that were persistently shadowed, but by different routing loops. Such loops towards the same destination were typically of the same length, differing only in a single address, and these addresses were commonly found to be from the same /64. This phenomenon is probably caused by load balancing, and impacts Internet measurements. In our case, a single path with load balancing is considered as multiple paths.

The question which therefore arises is whether the heuristic for persistent loop detection – currently mandating the same loop length and congruence of the involved IP addresses – has to be adapted for IPv6 to include such scenarios. If so, the total amount of persistent IPv6 loops would be up to 30% higher. In comparison, this behavior is rarely seen for IPv4 (<600 cases), presumably because load-balancing hosts share a public IP address, and pretend to be a single host. In terms of future research, we recommend to investigate this – to the best of our knowledge previously unknown – behavior in more detail.

Danger of persistent loops: Persistent routing loops are clearly undesired, however, the crucial question is if they are really harmful for Internet operation. Most loops shadow only ASes that are involved in the loop and imperil none (IPv4) or only involved (IPv6) ASes. Thus, we consider these loops, at 76.17% (IPv4) resp. 84.01% (IPv6), to be low risk. If these shadowed resp. imperiled addresses should become reachable, the responsible authority is able to reconfigure accordingly.

On the other hand, 16.01% (IPv4) resp. 14.43% (IPv6) of the loops imperil other ASes and thus have to be considered high risk. An attack, exploiting these routing loops and overwhelming routers, would cause Denial of Service in uninvolved ASes; mitigation would require coordination between different entities, probably causing noticeable downtime of services. Another 7.81% (IPv4) resp. 1.57% (IPv6) have to be considered medium-risk loops as they shadow other ASes but do not imperil. Rendering those shadowed addresses reachable requires coordination with other parties; however, there is typically less time pressure in such scenarios than in the mitigation of an ongoing attack.

Still, further aspects have to be considered: First, universally high amplification factors (median 122) generate high traffic volumes to easily overwhelm routers and connecting links. Second, routing loops are observed from multiple locations around the world, i.e., even higher traffic volumes could be generated by distributed attacks. Despite the fact that only every 6th resp. 7th loop has to be considered high-risk, they are able to imperil 14,630 (IPv4) resp. 23,955 (IPv6) uninvolved ASes, representing 20.02% (IPv4) resp. 85.04% (IPv6) of all ASes currently found in BGP announcements.

On the positive side, it must be pointed out that more than 90% of loops involve only a single AS, assigning clear responsibilities to their operators for their existence and remedy.

Internet core vs. Periphery: Persistent loops in IPv4 are predominantly found in NSPs, whereas their IPv6 counterparts reside in Cable/DSL/ISP networks, i.e., the former are found in the Internet's core, and the latter towards its periphery.

On the one hand, this might be caused by the more generous address practices in IPv6: In peripheral networks, devices which shared an address in IPv4 are now assigned their own public IPv6 address(es), decreasing the relative share of addresses attributed to NSPs.

On the other hand, the IPv6 transition resembles more modifications in peripheral networks than core networks, i.e., removal of address sharing and internal addressing, address abundance instead of scarcity, multiple and regularly changing addresses per host, and loops might arise from following IPv4 (i.e., legacy) customs for IPv6. One of these customs seems to be the lack of pull-up routes. Such routes filter traffic towards inactive addresses; otherwise, this traffic might be routed back to the Internet by a default route, eventually causing a routing loop. In IPv4, such pull-up routes are rarely necessary (see above on IPv4 address scarcity); however, for IPv6 with its vast inactive address spaces, this situation is quite different. We have feedback from administrators of three ASes with routing loops. All of them confirmed the loops' existence and the lack of a pull-up route; one of them even configured a pull-up route after our discussion. The responses confirm our hypothesis on pull-up routes; though, this very limited set of responses cannot be considered representative for the Internet as a whole.

9. Conclusion

We conducted a comprehensive and fine-grained measurement study on persistent routing loops — the first-ever considering the next-generation Internet Protocol IPv6, and the first for IPv4 since 2005, painting a clear picture of this phenomenon's prevalence and characteristics in today's Internet.

We discovered 23,208 persistent routing loops in IPv4 and 30,090 in IPv6, rendering 0.91% (IPv4) resp. 2.20% (IPv6) of the address space — as currently announced in BGP — unreachable. Beyond, adversaries might exploit the identified loops to amplify traffic, thus congesting links between or even overwhelming the involved routers and render additional addresses that are connected via these resources unreachable. In total, we found 7.18% (IPv4) resp. 23.00% (IPv6) of the currently used address space to be vulnerable. Even worse, loops typically cause traffic amplification by a factor of 122 (median), rendering attacks quite easy. Loops are observable from other vantage points, enabling adversaries to collaborate in distributed attacks.

Finally, we emphasize positive aspects, particularly on the remedy of persistent routing loops: For both protocols, more than 90% of the loops are within a single autonomous system, i.e., the involved routers are under joint control (typically of an ISP or NSP), and the respective organizations can manage such situations without the involvement of third parties. Additionally, only 1 in 4 (IPv4) resp. 1 in 6 (IPv6) loops threatens — i.e., shadows or imperils — autonomous systems which are not involved in the loop. In this light, we advise to mitigate persistent routing loops best possible, as they are a powerful tool for DoS attacks whose exploitation affects all Internet services of the targeted addresses.

Funding

This material is based upon work partially supported by (1) the Christian-Doppler-Laboratory for Security and Quality Improvement in the Production System Lifecycle; the financial support by the Austrian Federal Ministry for Digital and Economic Affairs, the National Foundation for Research, Technology and Development and the Christian Doppler Research Association are gratefully acknowledged; (2) SBA Research (SBA-K1), a COMET Center within the framework of COMET — Competence Centers for Excellent Technologies Programme and funded

by BMK, BMDW, and the province of Vienna. The COMET Programme is managed by FFG; (3) Project 877110 2big2fail funded by the Program “BRIDGE 1” (FFG); (4) Project DynAISEC FO999887504 funded by the Program “ICT of the Future” — an initiative of the Austrian Ministry of Climate Action, Environment, Energy, Mobility, Innovation and Technology; (5) Project “Handling Data from IPv6 Scanning” funded by the Next Generation Internet (NGI) Initiative.

CRediT authorship contribution statement

Markus Maier: Methodology, Software, Validation, Investigation, Visualization, Writing — review & editing. **Johanna Ullrich:** Conceptualization, Validation, Resources, Writing — original draft, Writing — review & editing, Supervision, Project administration, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.


Data availability

Data will be made available on request.

References

- [1] J. Xia, L. Gao, T. Fei, A measurement study of persistent forwarding loops on the internet, *Comput. Netw.* 51 (2007) 4780–4796.
- [2] J. Rüth, T. Zimmermann, O. Hohlfeld, Hidden treasures — recycling large-scale internet measurements to study the internet's control plane, in: *Passive and Active Measurement*, Springer International Publishing, Cham, 2019, pp. 51–67.
- [3] X. Li, B. Liu, X. Zheng, H. Duan, Q. Li, Y. Huang, Fast IPv6 network periphery discovery and security implications, in: 2021 51st Annual IEEE/IFIP International Conference on Dependable Systems and Networks, DSN, IEEE, 2021, pp. 88–100, <http://dx.doi.org/10.1109/DSN48987.2021.00025>, <https://ieeexplore.ieee.org/document/9505062/>.
- [4] Y. Nosyk, M. Korczyński, A. Duda, Routing loops as mega amplifiers for dns-based ddos attacks, in: *Passive and Active Measurement*, Springer International Publishing, Cham, 2022, pp. 629–644.
- [5] Google IPv6 statistics, 2022, Google, <https://www.google.com/intl/en/ipv6/statistics.html>. (Accessed 30 June 2022).
- [6] CAIDA, RouteViews IPv4 Prefix to AS mappings, 2022a, https://catalog.caida.org/details/dataset/routeviews_ipv4_prefix2as. (Accessed 30 September 2021).
- [7] CAIDA, RouteViews IPv6 Prefix to AS mappings, 2022b, https://catalog.caida.org/details/dataset/routeviews_ipv6_prefix2as. (Accessed 30 September 2021).
- [8] PeeringDB, The interconnection database, 2022, <https://www.peeringdb.com/>. (Accessed 30 September 2021).
- [9] M. Ziv, L. Izhikevich, K. Ruth, K. Izhikevich, Z. Durumeric, Asdb: A system for classifying owners of autonomous systems, in: *Proceedings of the 21st ACM Internet Measurement Conference, IMC '21*, Association for Computing Machinery, New York, NY, USA, 2021, pp. 703–719, <http://dx.doi.org/10.1145/3487552.3487853>.
- [10] V. Paxson, End-to-end routing behavior in the internet, *IEEE/ACM Trans. Netw.* 5 (1997) 601–615.
- [11] U. Hengartner, S. Moon, R. Mortier, C. Diot, Detection and analysis of routing loops in packet traces, in: *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement, IMW '02*, 2002, pp. 107–112, <http://dx.doi.org/10.1145/637201.637217>.
- [12] J. Kučera, R.B. Basat, M. Kuka, G. Antichi, M. Yu, M. Mitzenmacher, Detecting routing loops in the data plane, in: *Proceedings of the 16th International Conference on Emerging Networking Experiments and Technologies, CoNEXT '20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 466–473, <http://dx.doi.org/10.1145/3386367.3431303>.
- [13] CAIDA, The IPv4 routed /24 topology dataset, 2022, https://www.caida.org/catalog/datasets/ipv4_routed_24_topology_dataset/. (Accessed 15 November 2021).
- [14] Q. Lone, M. Luckie, M. Korczyński, M. van Eeten, Using loops observed in traceroute to infer the ability to spoof, in: *Passive and Active Measurement*, Springer International Publishing, Cham, 2017, pp. 229–241.
- [15] CAIDA, Ark IPv6 topology dataset, 2022a, https://catalog.caida.org/details/dataset/ipv6_allpref_topology. (Accessed 30 September 2021).

- [16] CAIDA, The IPv6 routed /48 topology dataset, 2022b, https://www.caida.org/catalog/datasets/ipv6_routed_48_topology_dataset/. (Accessed 15 November 2021).
- [17] J. Postel, Internet control message protocol, RFC 792, IETF, 1981, <http://tools.ietf.org/rfc/rfc0792.txt>.
- [18] A. Conta, S. Deering, M. Gupta, Internet control message protocol (ICMPv6) for the internet protocol version 6 (IPv6) specification, RFC 4443, IETF, 2006, <http://tools.ietf.org/rfc/rfc4443.txt>.
- [19] R. Beverly, Yarrp'ing the internet: Randomized high-speed active topology discovery, in: Proceedings of the 2016 Internet Measurement Conference, IMC '16, Association for Computing Machinery, New York, NY, USA, 2016, pp. 413–420, <http://dx.doi.org/10.1145/2987443.2987479>.
- [20] R. Beverly, R. Durairajan, D. Plonka, J.P. Rohrer, In the IP of the beholder: Strategies for active IPv6 topology discovery, in: Proceedings of the Internet Measurement Conference 2018, IMC '18, Association for Computing Machinery, New York, NY, USA, 2018, pp. 308–321, <http://dx.doi.org/10.1145/3278532.3278559>.
- [21] B. Augustin, X. Cuvelier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, R. Teixeira, Avoiding traceroute anomalies with paris traceroute, in: Proceedings of the 6th ACM SIGCOMM on Internet Measurement, IMC '06, ACM Press, 2006, p. 153, <http://dx.doi.org/10.1145/1177080.1177100>.
- [22] P. Alvarez, F. Oprea, J. Rula, Rate-limiting of IPv6 traceroutes is widespread: measurements and mitigations, 2017, <https://www.ietf.org/proceedings/99/slides/slides-99-maprg-rate-limiting-of-ipv6-traceroutes-is-widespread-measurements-and-mitigations-01.pdf>.
- [23] RIPE NCC, Best current operational practice for operators: IPv6 prefix assignment for end-users - persistent vs non-persistent, and what size to choose, 2017, <https://www.ripe.net/publications/docs/ripe-690>. (Accessed 07 October 2021).
- [24] RIPE NCC, RIPE routing working group recommendations on route aggregation, 2006, <https://www.ripe.net/publications/docs/ripe-399>. (Accessed 15 November 2021).
- [25] RIPE NCC, RIPE routing working group recommendations on IPv6 route aggregation, 2011, <https://www.ripe.net/publications/docs/ripe-532>. (Accessed 15 November 2021).
- [26] Z. Durumeric, E. Wustrow, J.A. Halderman, Zmap: Fast internet-wide scanning and its security applications, in: 22nd USENIX Security Symposium, USENIX Security 13, USENIX Association, Washington, D.C., 2013, pp. 605–620, <https://www.usenix.org/conference/usenixsecurity13/technical-sessions/paper/durumeric>.
- [27] P. Traina, D. McPherson, J. Scudder, Autonomous system confederations for BGP, RFC 3065, IETF, 2001, <http://tools.ietf.org/rfc/rfc3065.txt>.
- [28] D. Crocker, Telnet byte macro option, RFC 729, IETF, 1977, <http://tools.ietf.org/rfc/rfc0729.txt>.
- [29] E. Rye, R. Beverly, K.C. Claffy, Follow the scent: Defeating IPv6 prefix rotation privacy, in: Proceedings of the 21st ACM Internet Measurement Conference, IMC '21, Association for Computing Machinery, New York, NY, USA, 2021, pp. 739–752, <http://dx.doi.org/10.1145/3487552.3487829>.
- [30] S.J. Saidi, O. Gasser, G. Smaragdakis, One bad apple can spoil your IPv6 privacy, ACM SIGCOMM Comput. Commun. Rev. 52 (2022).
- [31] T. Narten, R. Draves, S. Krishnan, Privacy extensions for stateless address autoconfiguration in IPv6, RFC 4941, IETF, 2007, <http://tools.ietf.org/rfc/rfc4941.txt>.
- [32] F. Gont, S. Krishnan, T. Narten, R. Draves, Temporary address extensions for stateless address autoconfiguration in IPv6, RFC 8981, IETF, 2021, <http://tools.ietf.org/rfc/rfc8981.txt>.
- [33] K. Bock, A. Alaraj, Y. Fax, K. Hurley, E. Wustrow, D. Levin, Weaponizing middleboxes for TCP reflected amplification, in: 30th USENIX Security Symposium, USENIX Security 21, USENIX Association, 2021, pp. 3345–3361, <https://www.usenix.org/conference/usenixsecurity21/presentation/bock>.
- [34] B. Carpenter, S. Brim, Middleboxes: Taxonomy and issues, RFC 3234, IETF, 2002, <http://tools.ietf.org/rfc/rfc3234.txt>.
- [35] Z.M. Mao, J. Rexford, J. Wang, R.H. Katz, Towards an accurate as-level traceroute tool, in: Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM '03, Association for Computing Machinery, New York, NY, USA, 2003, pp. 365–378, <http://dx.doi.org/10.1145/863955.863996>.
- [36] R. Motamedi, B. Yeganeh, B. Chandrasekaran, R. Rejaie, B.M. Maggs, W. Willinger, On mapping the interconnections in today's internet, IEEE/ACM Trans. Netw. 27 (2019) 2056–2070.
- [37] A. Marder, M. Luckie, B. Huffaker, K. Claffy, Vrfinder: Finding outbound addresses in traceroute, Proc. ACM Meas. Anal. Comput. Syst. 4 (2020).
- [38] A. Lodhi, N. Larson, A. Dhamdhere, C. Dovrolis, K. Claffy, Using peeringdb to understand the peering ecosystem, SIGCOMM Comput. Commun. Rev. 44 (2014) 20–27.
- [39] P. Srisuresh, M. Holdrege, IP network address translator (NAT) terminology and considerations, RFC 2663, IETF, 1999, <http://tools.ietf.org/rfc/rfc2663.txt>.
- [40] J. Weil, V. Kuarsingh, C. Donley, C. Liljenstolpe, M. Azinger, IANA-reserved IPv4 prefix for shared address space, RFC 6598, IETF, 2012, <http://tools.ietf.org/rfc/rfc6598.txt>.
- [41] A.H.S. Project, Name-based virtual host support, 2022, <https://httpd.apache.org/docs/current/vhosts/name-based.html>. (Accessed 28 June 2022).
- [42] Y. Rekhter, B. Moskowitz, D. Karrenberg, G.J. d. Groot, E. Lear, Address allocation for private internets, RFC 1918, IETF, 1996, <http://tools.ietf.org/rfc/rfc1918.txt>.
- [43] S. Strowes, RIPE labs, visibility of IPv4 and IPv6 prefix lengths in 2019, 2019, http://labs.ripe.net/author/stephen_strowes/visibility-of-ipv4-and-ipv6-prefix-lengths-in-2019/. (Accessed 09 May 2021).



Markus Maier is a Ph.D. student at University of Vienna, Austria, Research Group Security and Privacy. He received his B.Sc. and M.Sc. at Technical University of Vienna in Software Engineering & Internet Computing. His research interests include routing, network measurement and network security.



Johanna Ullrich is a key researcher at SBA Research, Austria, leading the Networks and Critical Infrastructures Security Research Group, and a researcher of the Christian Doppler Laboratory for Security and Quality Improvement in the Production System Lifecycle (University of Vienna). She received a Ph.D. sub auspiciis praesidentis from TU Wien. She was awarded the Research Prize of the Dr. Maria Schaumayer Foundation and nominated for the Hedy Lamarr Prize twice. Her research focuses on network security, particularly measuring experiments and IPv6. She has proven that the IPv6 Privacy Extension as specified in RFC 4941 and implemented in major operating systems was vulnerable.