

Day in the Life of RIPE Atlas: Operational Insights and Applications in Network Measurements

Yevheniya Nosyk
KOR Labs
France

Malte Tashiro
IJJ Research Laboratory / SOKENDAI
Japan

Qasim Lone
RIPE NCC
The Netherlands

Robert Kisteleki
RIPE NCC
The Netherlands

Andrzej Duda
Université Grenoble Alpes
France

Maciej Korczyński
Université Grenoble Alpes
France

ABSTRACT

Network measurement platforms are increasingly popular among researchers and operators alike due to their distributed nature, simplifying measuring the remote parts of the Internet. RIPE Atlas boasts over 12.9 K vantage points in 178 countries worldwide and serves as a vital tool for analyzing anycast deployment, network latency, and topology, to name a few. Despite generating over a terabyte of measurement results per day, there is limited understanding of the underlying processes. This paper delves into one day in the life of RIPE Atlas, encompassing 50.9 K unique measurements and over 1.3 billion results. While most daily measurements are user-defined, it is built-ins and anchor meshes that account for 89% of produced results. We extensively examine how different probes and measurements contribute to the daily operations of RIPE Atlas and consider any bias they may introduce. Furthermore, we demonstrate how existing measurements can be leveraged to investigate censorship, traceroute symmetry, and the usage of reserved address blocks, among others. Finally, we curate a set of recommendations for researchers using the RIPE Atlas platform to foster transparency, reproducibility, and ethics.

1 INTRODUCTION

The Internet is the biggest network of networks in the world, interconnecting more than 82.3 K IPv4/IPv6 autonomous systems that advertise over 1 million BGP routing prefixes [119]. Given this unprecedented scale, it is no surprise that the Internet is constantly affected by network outages that stem from cyberattacks, natural disasters, hardware or software failures, as well as human errors. Such events may have a global impact and far-reaching consequences for millions of end users.

One of the prominent network outages happened in 2016 when the Mirai botnet attacked the Dyn Domain Name System (DNS) operator. It resulted in rendering unavailable some of the most popular websites, such as those of Amazon, GitHub, and Netflix, to name a few [7]. More recently, Facebook and its associated services, Instagram and WhatsApp, went offline when one of the “routine maintenance jobs” disconnected the Meta network hosting DNS servers from the Internet [54]. It is estimated that the outage cost the company about \$65 million [18]. The understanding and mitigation of such disruptions is more efficient when one can accurately measure how far the problem propagated and which parts of the Internet were affected. It is also crucial to proactively identify any vulnerability or misconfiguration that may become an issue in the future.

Network measurements can be conducted from external or internal vantage points, each offering unique insights into network performance. External vantage points evaluate network reachability and end-to-end performance, identifying issues such as connectivity problems and high latency. However, they often lack the depth needed to uncover the root causes of these issues. In contrast, internal vantage points are essential for gaining a comprehensive understanding of network operations. They enable detailed monitoring of traffic patterns and facilitate the identification of internal issues, such as routing misconfigurations and bandwidth bottlenecks.

To fully leverage the strengths of internal vantage points, developing a distributed infrastructure that supports a variety of custom measurements across multiple networks is crucial. Such a system would not only streamline automation and simplify result retrieval but also significantly enhance overall network visibility. Examples of such systems are proxy networks, advertising tens of millions of vantage points available to their clients. Researchers previously used them to test the capabilities of DNS resolvers [26], DNS time-to-live (TTL) violations [15], and traffic manipulations [25]. Yet, valid concerns have been raised as to whether participating IP addresses are sourced ethically. In particular, Mi et al. [78] found some of the residential proxies to be compromised Internet of Things devices. NLNOG RING [88] is another initiative introduced by the Netherlands Network Operator Group. Any operator joining the measurement network with its own virtual machine gains shell access to all the participating servers. However, the platform is only open to participants.

RIPE Atlas stands out as one of the most widely known Internet measurement platforms, appreciated by researchers and operators alike. Powered by a network of volunteers running small measurement devices (probes), it constantly collects data about the topology, connectivity, and reachability of the global Internet. The platform is open for anyone to launch custom measurements, provided a user possesses a sufficient amount of RIPE Atlas credits. If in need, researchers can request them by contacting the RIPE Atlas team directly. For that reason, the platform has been extensively used in the research community to study network latency [9, 21, 22, 30, 40, 64, 75, 76, 90], topology [6, 10, 56, 93, 120, 124], and anycast [63, 77, 83, 103, 114, 122, 125, 126] among many other topics.

By 2024, RIPE Atlas has grown into a large-scale network of 12.9 K probes, spread over 178 countries and 4 K ASes. Since its initial launch in 2010, the platform has generated almost one petabyte of data [112], with a rate of 1.3 billion measurement results per day

in early 2024. Such a scale made the Atlas team reflect on how to best store all the invaluable data [62]. However, there has not been a study that systematically analyzes the measurements being run and the phenomena they can reveal. Our paper fills this research gap and provides an extensive overview of the RIPE Atlas platform, with its building blocks, measurements, and use cases. Overall, our contributions are as follows:

- We analyze one day in the life of RIPE Atlas, comprising 50.9 K measurements and 1.3 billion results generated by 12.9 K probes. We uncover that anchors produce almost 70% of daily results despite regular users launching the highest number of individual measurements.
- We present several exploratory case studies that leverage RIPE Atlas data to provide insights into network behavior and infrastructure. By analyzing a single day’s worth of RIPE Atlas data, we highlight the unique value of this daily snapshot dataset in uncovering patterns and trends that inform network performance and reliability. Our findings demonstrate the practical applications of these data in real-world scenarios, offering valuable implications for network operators and researchers alike.
- Based on the key findings of this paper and the discussions with the RIPE Atlas Team, we curate the list of best practices for all researchers wishing to use the platform. We call for transparency, reproducibility, ethical conduct, and encourage thorough exploration of RIPE Atlas data to uncover valuable insights for the network community.

The rest of the paper is organized as follows. Section 2 describes the inner workings of the RIPE Atlas, while Section 3 illustrates the use of the platform by researchers. Section 4 presents the day-in-the-life dataset that we analyze in the remainder of the paper. Section 5 discusses anchoring, built-in, and user-defined measurements. Section 6 proposes a set of recommendations for a better use of RIPE Atlas, Section 7 deals with ethics, and Section 8 concludes the paper.

2 BACKGROUND ON RIPE ATLAS

This section provides an overview of the RIPE Atlas measurement platform. We first describe the overall system design and its main building blocks. We then discuss the types of supported measurements and how one can launch them.

2.1 Building Blocks

The core element of the RIPE Atlas platform is a *probe*, capable of running in its host network and executing various types of Internet measurements. Probes come in two flavors—hardware and software. The former are small low-power devices (Lantronix XPort Pro, TL-MR 3020, or NanoPi NEO Plus2, depending on the version [58]) following the install-and-forget principle, thus requiring little to no maintenance from the probe owner. On the contrary, software probes are set up and updated independently on the existing infrastructure (e.g., virtual machines or physical servers). RIPE NCC receives applications from potential probe hosts and accepts those that would contribute to the better coverage of the platform. In return, probe owners gain *credits* that can be used to run custom

measurements. As of February 2024, there are 12.9 K connected probes.

Anchor is an enhanced probe, providing two main services: i) executing a large number of measurements and ii) being measurement targets for other probes and anchors. Accomplishing these two goals raises higher requirements for the host network and the underlying hardware. Anchors are expected to be highly available, so they are not suitable for small home or office networks behind firewalls. Instead, Internet Service Providers (ISPs), Internet eXchange Points (IXPs), or big cloud providers are natural candidates for hosting such systems. Note that, unlike regular probes, anchors are explicitly required to respond to pings and traceroutes. To reward the owners for operating anchors, they earn ten times more credits than regular probe hosts.

Originally, all the anchors were meant to be hardware devices, and hosts were committed to purchasing the specific hardware approved by the Atlas team. Non-profit organizations could additionally benefit from several rounds of sponsorship, helping to set up anchors in underrepresented networks [28]. However, RIPE NCC relaxed the hardware requirement in 2018 and introduced Virtual Machine Anchors that contributed to a large expansion of the infrastructure [29]. As of February 2024, 810 anchors are connected to the RIPE Atlas platform. They constitute a full mesh and perform a series of measurements between themselves to serve as a baseline for the state of Internet connectivity.

2.2 Measurements

The RIPE Atlas *measurement* is the key unit of work performed by all the probes and anchors. Every device comes with a series of *built-in* measurements performed towards well-known targets, pre-defined by the Atlas team. At the same time, one can launch custom *user-defined* measurements via the web interface¹ or the API.² Each measurement is defined by its type, participating probes/anchors, targets, and timing.

Built-in measurements [98] generate an important body of data about global Internet connectivity. They are executed on all the connected probes every 4 minutes to once per day, some of them up and running since 2010. Every 240 seconds, pings are sent to all 13 DNS root servers and 7 measurement targets from the RIPE Atlas infrastructure. Traceroutes are executed towards the same destinations every 30 minutes in addition to topology scans running every 15 minutes. DNS measurements are duplicated over both TCP and UDP transport protocols, with each target receiving SOA DNS and 4 CHAOS-class TXT requests for `version.bind`, `hostname.bind`, `id.server`, `version.server`. In addition, probes resolve some of the popular and random domain names. The frequency of DNS measurements greatly varies depending on the particular query sent.

The Transport Layer Security (TLS) certificates of `www.ripe.net` and `atlas.ripe.net` are retrieved once per day, and the HTTP measurements towards `http://www.ripe.net/favicon.ico` and `http://ip-echo.ripe.net/` are executed every 24 and 1 hour, respectively. If supported by the host network, all the measurements are run in IPv4 and IPv6 address spaces. Moreover, all the probe connection

¹<https://atlas.ripe.net>

²<https://atlas.ripe.net/docs/apis/rest-api-manual/>

and disconnection events are also logged to give an overview of the probe uptime. Built-in measurements are public by default, and the results can be freely accessed by any interested party.

RIPE Atlas anchors form a distinct group of targets for so-called anchoring measurements. Every anchor receives a ping, a traceroute, and an HTTP GET request from all the other anchors (thus forming a full mesh) and a subset of regular RIPE Atlas probes. The measurements are run every 4 to 30 minutes to create a reliable picture of the global Internet connectivity. Whenever available, all the measurements are run over IPv4 and IPv6. Note that the list of anchoring measurements is not static and gets updated every time a new anchor is added to the mesh.

One may need to perform a more fine-grained measurement, e.g., resolve a custom domain name or ping a particular destination. In this case, built-ins would be of little help. Instead, any registered user can launch a custom measurement, provided they have a sufficient amount of credits. There are currently 6 supported measurement types, including ping, traceroute, DNS, NTP, TLS, and HTTP [99]. Note that the last one is restricted to target only RIPE Atlas anchors. By default, the web interface proposes to run a measurement on 50 random probes located worldwide. It is also possible to restrict the participants to a specific AS, an IP prefix, a country, or a region.

If known, a particular probe ID can also be specified as well as an existing measurement ID to reuse the same set of participants. Each measurement can either be one-off or recurring (e.g., every 5 minutes). Based on its definition, each measurement has an associated cost in credits, proportional to the load placed on the probes themselves and the RIPE Atlas infrastructure in general. Finally, certain limits in place prevent a single user from abusing the platform. One cannot run more than 100 measurements simultaneously, use more than 1000 probes per measurement, and spend more than 1 M credit points per day. All the special requests to bypass the quotas and limitations are considered by the Atlas team on a case-by-case basis.

3 RELATED WORK

Since the launch of RIPE Atlas in 2010 and its introduction to the research community in 2015 [102], the platform has been widely used in more than a thousand scientific publications [46]. To tackle this amount of related work in the field efficiently, we adopt the approach of Scheitle et al. [108]. Specifically, we have surveyed all the papers published in the past 5 years (2019-2023) at relevant conferences in the areas of Internet measurements (PAM, ACM IMC, TMA), networking (ACM CoNEXT, ACM/IRTf ANRW, IEEE INFOCOM, ACM SIGCOMM), and security (USENIX Security, IEEE S&P, IEEE Euro S&P, ACM CCS, NDSS). We looked for publications that relied on RIPE Atlas in their methodology and selected 79 papers for further analysis.

Given that the RIPE Atlas platform is mostly considered to be a network monitoring tool, it is of no surprise that 89% of selected publications appeared in Internet measurement and networking venues. Consequently, it was largely underrepresented in the security community. Researchers mostly performed traceroute, DNS, and ping experiments, only occasionally referring to other scan types. RIPE Atlas was used to study anycast deployment [63, 77, 83,

103, 114, 122, 125, 126], latency [9, 21, 22, 30, 40, 64, 75, 76, 90], topology [6, 10, 56, 93, 120, 124], and IP geolocation [17, 27, 42, 69, 105]. Apart from those, researchers tackled a broad number of networking topics, including anomalies and outages [36, 55, 80], Resource Public Key Infrastructure (RPKI) [23, 39, 49, 67, 104, 111], router fingerprinting [4, 5], Border Gateway Protocol (BGP) [38, 66, 113], traceroute analysis [45, 47, 81], IPv6 [53, 116], routing loops [3], submarine connectivity [37, 68], and satellite [79, 92]. They used both built-in and user-defined measurements with a slight preference towards the latter.

DNS measurements constituted another important part of related work covering a wide range of topics such as DNS Security Extensions (DNSSEC) [86, 87], QNAME minimization [73, 121], caching [1, 84, 96, 110], DNS-over-TCP [74], encryption [24, 33], EDNS(0) Client Subnet [2, 19, 107], fragmentation [85], cyclic dependencies [82], lame delegations [115], and infrastructure [44, 118]. However, the need to query custom domain names implied setting up user-defined measurements.

The remaining measurement types are a small fraction of related work. Notably, different RIPE Atlas measurement types (TLS, DNS, traceroute) allowed researchers to analyze censorship from multiple angles. For example, they examined DNS interception [70, 89, 97], rogue root certificates [94], censorship circumvention [65], and more broadly, the exposure to traffic observation and tampering [43]. The latter two relied on TLS measurements that were rarely seen in other existing work. Finally, probe (dis)connection events were used [16, 95] to identify IP address reuse.

Given the number of measurements running on the platform and the associated load on RIPE Atlas probes, valid concerns were raised about the accuracy of the results. Holterbach et al. [51] found that i) measurements running concurrently on the same probe may increase the delay by milliseconds and ii) measurements on different probes can become desynchronized despite being launched at the same time. Bajpai [14] further summarized the potential issues that may be encountered by RIPE Atlas users, such as rate limits, per-hop latency aggregation, and a biased probe distribution across ASes/countries. Researchers specifically highlighted that vantage point locations may significantly influence the obtained measurement results [8, 109].

Our analysis of the related work shows several patterns of RIPE Atlas usage among researchers. First, measurements are skewed towards the main three types (namely traceroutes, DNS, and ping) while others remain underrepresented. Second, user-defined measurements are used in more than half of the cases because built-ins cannot address all the research questions. Finally, security researchers rely considerably less on RIPE Atlas than the networking community.

4 DAY IN THE LIFE OF RIPE ATLAS

This section provides an overview of one day in the life of RIPE Atlas. We carefully select a 24-hour window to analyze the active probes, anchors, and measurements in which they participated.

4.1 Dataset

When describing one day in the life of RIPE Atlas, it is crucial to evaluate whether the chosen dataset is a representative snapshot of

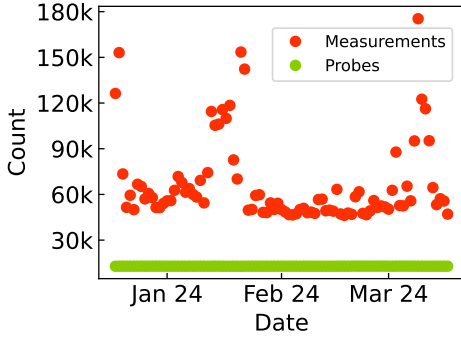


Figure 1: Ongoing measurements and active RIPE Atlas probes per day between Jan 1, 2024 and Mar 31, 2024.




Figure 2: Cumulative distribution of dates of the first connection of probes and anchors.




Figure 3: Cumulative distribution of total uptimes in days for probes and anchors.

the system operations. We assess it using two distinct metrics: the number of ongoing measurements and the number of connected probes per day.

Considering the major growth of the platform since 2010, we refrain from the naive approach of computing the mean values over the past 14 years. Instead, we focus on a more recent period of time spanning the first three months of 2024. Figure 1 presents the number of connected probes and ongoing measurements between January and March 2024. Overall, we see a stable distribution of probes, fluctuating between 12.8 K and 12.9 K per day with a mean of 12,888. However, trends differ for ongoing measurements, where the numbers range from 46 K to 175 K. The mean of such a distribution will inevitably be skewed towards a few outliers. Therefore, we compute its median of 56 K measurements per day. Finally, we choose the date for which the numbers of ongoing measurements and connected probes fall within the range of the median absolute deviation, set to 7 K for measurements and 14 for probes. One of the days that satisfies these criteria is February 21, 2024. We note that while the analysis of data for any short period is not necessarily illustrative of the overall system use, we argue that the chosen day is a representative snapshot of the platform operation in 2024. We download all the active measurements during this period, all the generated results, and the probe metadata using the RIPE Atlas API.




Figure 4: Number of RIPE Atlas probes and anchors connected to the platform on Feb 21, 2024, per IP prefix.




Figure 5: Number of RIPE Atlas probes and anchors connected to the platform on February 21, 2024, per AS.

4.2 Probes and Anchors

On the analyzed day, 36.9 K measurement devices (probes and anchors) are known to the RIPE Atlas platform. More than 20 K of them are abandoned—they have not been connected to the system for at least several months. Almost 4 K probes/anchors are temporarily disconnected, for example, possibly due to transient network failures, probe hardware not functioning properly [52], or owners unplugging the device for any reason. The remaining 12,082 probes and 810 anchors are active, forming the basis for our subsequent analysis.

4.2.1 First Connection Dates and Uptimes. Figure 2 shows the cumulative distribution of the dates of the first connection of the currently active probes and anchors. In general, probes have been steadily added to the network since the first days of RIPE Atlas, 96 of them being connected since 2010—they have been generating invaluable longitudinal data for the last 14 years. The great majority of active anchors have joined the system since 2018, which coincides with the introduction of virtual machine anchors. The oldest one has been running since 2013. We note, however, that the first connection dates do not necessarily signify that devices have been constantly connected ever since. Figure 3 additionally plots the total uptimes of connected probes and anchors measured in days. We can see that these findings are generally consistent with the expected uptimes based on the first connection dates. For example, the few probes first connected in 2010 have a total uptime of over 4.8 K days (over 13 years).

The consistent availability of RIPE Atlas probes and anchors highlights the network’s reliability and stability. Probes that have remained connected for extended periods provide essential longitudinal data crucial for historical analysis. The consistency between their initial connection dates and ongoing uptime indicates near-continuous operation, ensuring reliable and uninterrupted data collection, essential for comprehensive insights into global Internet dynamics.

4.2.2 IP Connectivity. More than 12% of connected probes are private, thus not revealing any information about their underlying IP connectivity. Focusing on public probes only, we see 5.6 K (46.5%) of them being dual-stack, while the remaining 4.8 K (40%) and 146 (1.2%) are IPv4-only and IPv6-only, respectively. Anchors show a substantially higher support of IPv6, with 744 (92%) being fully

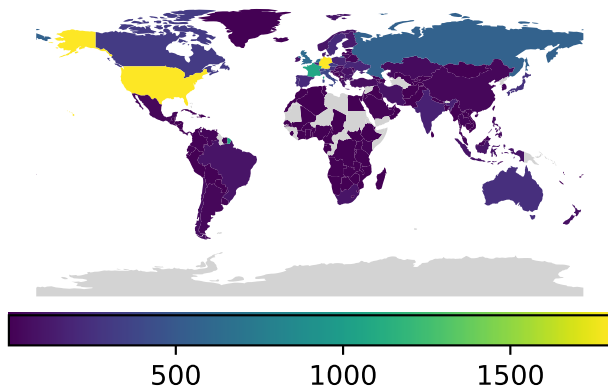


Figure 6: Number of connected RIPE Atlas probes and anchors per country in February 21, 2024.

dual-stack. To further understand the distribution of these measurement vantage points, we compute the number of anchors and probes per IP prefix and AS. Figures 4 and 5 show the cumulative distributions for IPv4 and IPv6 address spaces. Overall, most IP prefixes and ASes with RIPE Atlas measurement vantage points contain exactly one probe or anchor, regardless of the address family. Anchors tend to be better distributed than probes, with the best coverage attained among IPv4 prefixes, for which 97.84% contain only one anchor. In all other cases, each prefix contains at most 8 anchors. When aggregating the numbers at the AS level, we see some decrease in the diversity: the ratio of ASes with one anchor is 87-88%, and at most, one probe is hosted in 67% of ASes. The long tail also shows some ASes with a high concentration of probes, mostly Deutsche Telekom, Comcast Cable Communications, Free, and Orange. As they are major telecommunications operators in their corresponding countries, we hypothesize that probes may be hosted in the home networks of their owners.

The prevalence of private probes among RIPE Atlas measurements limits visibility into their IP connectivity, potentially skewing the analysis at the network level. The analysis of public probes reveals a notable dual-stack support among anchors compared to probes, highlighting varying network capabilities that researchers must consider when performing measurements across IPv4 and IPv6 address spaces.

4.2.3 Geographical Distribution. We finally assess the geographical distribution of RIPE Atlas anchors and probes connected to the system as of February 21, 2024. Figure 6 plots the number of probes and anchors per country. Overall, the platform has measurement devices in 178 countries worldwide. Germany and the United States host substantially more vantage points than any other country, both accounting for 28% of probes and anchors. At the other extreme, 32 countries only have one probe or anchor. Generally, there is a strong bias towards Europe and North America, while RIPE Atlas remains underrepresented in other parts of the world.

Key takeaways. Despite being under-represented in certain regions, RIPE Atlas boasts a wide coverage. Probes and anchors have

been steadily joining the measurement network since 2010, serving as highly available and reliable baselines for analyzing the Internet. Their increasing support of both IP address families makes them suitable for any experiments requiring dual-stack support.

4.3 Measurements

The 12.9 K active probes and anchors participated in 50,885 RIPE Atlas measurements, generating more than 1.3 billion individual results and 1.1 TB of raw data. Tables 1,2,3 detail these statistics per anchoring, built-in, and user-defined measurements.

4.3.1 Measurement Types. The highest number of ongoing measurements (39 K or 76.7%) were user-defined, most of them being pings and traceroutes. Pings are one of the most convenient ways to test the reachability of a particular host at a particular instant, thus being extremely popular among network operators. Overall, platform users launched all types of measurements available to them. Anchors contribute to the second largest chunk of ongoing measurements, but we recall that i) only anchors are targets of them, and ii) three types of measurements are performed (ping, traceroute, and HTTP). The remaining 251 measurements are built-in and are, therefore, preconfigured to run on each connected probe. More than half of them are DNS CHAOS-class TXT queries targeting all DNS root servers. We can also note that two types of measurements—specifically, probe (dis)connection events and traffic—are only applicable to built-ins. The latter is an opt-in feature that probe hosts can choose to collect statistics on the traffic received by the probe, for example, the number of received bytes or packets.

4.3.2 Measurement Results. One RIPE Atlas measurement may generate zero or more results, depending on the number of participating probes and periodicity. Out of 50.9 K ongoing measurements, 2.3 K did not generate a single result. In particular, 340 of them did not have a single participating probe, possibly due to the inability of the platform to accommodate the probe request, which may happen if a long-running measurement was defined with a set of probes not connected to the RIPE Atlas network anymore. Other cases included measurements launched and executed before the 24-hour analysis window but marked as active for a short period of time on February 21, 2024. We might have encountered active recurring measurements scheduled to run less frequently than every 24 hours, thus giving no results during the chosen day. Other 5.8 K measurements generated exactly one result object, meaning that they were either one-offs or recurring with the frequency of no more than once every 86.4 K seconds. In either case, one probe per measurement participated.

A great majority of ongoing measurements generated two or more results, which becomes apparent when considering the average ratio of measurement results to measurements—over 26 K to 1. Due to the full mesh between all the anchors, they lead in the number and size of generated results despite being the most limited chunk of measurements in terms of sources, destinations, and types. Built-ins come second with over 283 M results and 182 GB of raw data. Even though only 251 measurements generated them, we recall that they are frequent (running up to every 4 minutes) and scheduled on all the connected probes. User-defined measurements generated the least amount of data.

Table 1: All the anchoring measurements executed, results (in millions), and the data generated (in gigabytes) on February 21, 2024.

Type	Anchoring		
	Measurements	Results (M)	Size (GB)
Ping	3,868	648.15	287.46
Traceroute	3,841	172.08	419.26
DNS	-	-	-
HTTP	3,881	86.37	30.91
Probe (dis)connection	-	-	-
TLS	-	-	-
NTP	-	-	-
Traffic	-	-	-
Total:	11,590 (22.8%)	906.6 (67.5%)	737.63 (68.7%)

Table 2: All the built-in measurements executed, results (in millions), and the data generated (in gigabytes) on February 21, 2024.



Type	Built-in		
	Measurements	Results (M)	Size (GB)
Ping	38	127.77	54.63
Traceroute	44	21.51	50.79
DNS	158	126.91	74.95
HTTP	4	0.47	0.23
Probe (dis)connection	2	6.14	1.1
TLS	4	0.07	0.25
NTP	-	-	-
Traffic	1	0.68	1.0
Total:	251 (0.5%)	283.55 (21.1%)	182.95 (17.0%)

4.3.3 Creation Times. We next analyze whether the 50.8 K measurements observed on February 21, 2024, were first launched on the same day or they had been running for some time. Figure 7 presents the distribution of measurement creation dates, showing that 240 built-ins have been running since the first days of RIPE Atlas in 2010, collecting results from all the connected probes ever since. A significantly larger chunk of ongoing measurements (17 K or 33.5%) were first executed during the 24-hour analysis window. Moreover, most of them were also stopped the same day, suggesting that an important part of daily active measurements is either one-off or short-term. Figure 8 additionally plots the number of measurements created per hour (UTC) on February 21, 2024. There was a burst of 3,276 traceroute measurements between approximately 13:00 UTC and 16:00 UTC, seemingly belonging to the same campaign, as all the descriptions started with “Traceroute to”. They concern 295 domain names with up to 15 measurements per domain, however without any clear pattern or any domain-specific selection strategy.

4.3.4 Measurement Origins. We wanted to know *who* launched all 39 K user-defined measurements, keeping in mind that this information is not publicly available and there is no mapping between measurement IDs and end users. Instead, we refer to the description field available for each measurement. If not using the default string provided by the system, it may give a good indication of who is behind a particular measurement. It turns out that some of the user-defined measurements were in fact launched to back up certain RIPE NCC services. For example, DNSMON [20] is behind 4,435 executed measurements to assess the quality of

Table 3: All the user-defined measurements executed, results (in millions), and the data generated (in gigabytes) on February 21, 2024.

Type	User-defined		
	Measurements	Results (M)	Size (GB)
Ping	19,225	33.81	15.21
Traceroute	10,831	31.96	59.95
DNS	7,898	83.15	65.9
HTTP	3,881	86.37	30.91
Probe (dis)connection	-	-	-
TLS	670	2.98	11.86
NTP	166	0.2	0.13
Traffic	-	-	-
Total:	39,044 (76.7%)	153.55 (11.4%)	153.6 (14.3%)

**Figure 7: Dates when measurements active on February 21, 2024 were first created.****Figure 8: Creation times of measurements that were launched on February 21, 2024.**

service provided by root servers and selected top-level domains (TLDs). DomainMON [117] (1,684 measurements) is a similar tool that one can set up to monitor nameservers of its own domains. We also see 1,527 measurements triggered by RIPE IPmap [34], an active geolocation API relying on RIPE Atlas ping measurements under the hood. Focusing on external measurements, there are 598 pings, traceroutes, and DNS queries with the description of “HE Network Tools Site”. A closer look reveals that they are launched from <https://bgp.he.net/traceroute/> – a measurement toolkit run by Hurricane Electric [60]. We then see a number of ping measurement campaigns targeting the infrastructure of DNS operators, ISPs, content delivery networks (CDNs), and one regional Internet registry, among others.

4.3.5 Measurement Definitions. Before launching a measurement, a RIPE Atlas user needs to choose one or more probes by supplying either i) a probe ID, ii) a geographic area, iii) a country, iv) an AS, v) an IP prefix, or vi) an existing measurement ID to reuse the same list of participants. Focusing on 39 K user-defined measurements, we analyze the participant selection criteria requested by end users. Figure 9 plots the distribution of participation request types. Probe IDs are by far the most common way to choose measurement sources, although the two geographical selectors are also largely used. The chosen area was rarely a particular part of the world but rather “worldwide” and the US, Germany, and France were the most requested source countries.




Figure 9: Criteria used by RIPE Atlas users to choose probes for measurements.

Despite the possibility of freely choosing participating probes, the probe may be disconnected or too busy to accommodate more measurements at the time. Figure 10 plots the distribution of the number of selected probes and the number of actually participating probes. Few user-defined measurements (337) did not have a single participating probe, most of them being requested with probe IDs. We further verified that a great majority of requested probes were not connected on the day of the measurement, therefore being unable to process the requests. Almost 20 K measurements were executed with exactly the same sources as defined and the remaining 19 K has a subset of desired participants.

4.3.6 Measurement Targets. Finally, we assess the distribution of measurement targets, i.e., the systems being tested. Pings were sent towards 8 K unique hosts, the top 5 being one DNS provider from China and 4 TikTok CDN servers. They represented 20% of all daily ping measurements, as they participated in 4.9 K measurements in total. They are also the most common measurement destinations across all types of measurements. Traceroutes were distributed towards 6.5 K destinations, without having any significant outliers.

DNS measurements differ from other types as their targets are domain names rather than Internet hosts. However, the two pseudo domain names (namely `hostname.bind` and `version.bind`) leading the ranking are not globally resolvable, and the responses are, therefore, to be provided by the destination DNS servers. Despite HTTP measurements only targeting the RIPE Atlas infrastructure, we found 7 hosts outside the list of normally allowed destinations, including two regional Internet registries, two connectivity monitoring domains from a leading operating systems vendor, one lifestyle blog, one data center, and one domain belonging to a search engine operator. TLS and NTP measurements targeted 560 and 137 unique hosts, respectively, almost always being unique per measurement. Interestingly, 80 user-defined measurements have overlapping types and targets with some of the built-ins. While they generated a negligible number of results (4.9 M) with respect to all produced in a single day, we argue that built-ins could have been used instead.

Key takeaways. A great majority of ongoing measurements in a day are user-defined, but it is the anchoring meshes and built-ins that generate almost 89% of daily results. Overall, more than 39 K measurements were launched by end users in a day, mostly running on particular requested probes. However, some of them are




Figure 10: Distribution of the number of requested and actually participating probes.




Figure 11: Difference of traceroute lengths computed based on anchoring measurements.

not standalone and back up certain RIPE NCC services as well as external measurement platforms, a grey area in the usage of RIPE Atlas.

5 MEASUREMENT USE CASES

The previously described dataset contains more than 1.3 billion individual measurement results. In this section, we present a series of exploratory use cases observed on February 21, 2024, illustrating how RIPE Atlas data can uncover emerging patterns and open new avenues for research.

5.1 Anchoring Measurements

The majority of results generated by the RIPE Atlas platform in a single day come from anchoring measurements. Below, we describe how to use them to analyze traceroute symmetry.

5.1.1 Traceroute Symmetry. The full mesh traceroute measurements performed between all the anchors offer a rare opportunity to analyze the symmetry of paths on the Internet, as the anchor measurements provide a large set of bidirectional traceroutes in close temporal proximity. The asymmetry of Internet paths has gained attention in the past [31, 48, 120] and was studied both via the analysis of BGP data and active measurements using traceroute. Understanding path symmetry is useful for troubleshooting cases like routing problems that occur on the reverse path or path asymmetry that leads to differences in latency affecting latency-sensitive applications like video conferencing.

While there are studies that employed RIPE Atlas for their traceroute measurements [31], they only scheduled their own measurements. Thus, they faced issues with the daily credit limit and had to make compromises for their measurement configuration: Instead of a full mesh measurement, which would only include 112 probes due to the credit limit, they opted to perform pairwise measurements between 4 K random probes, resulting in 2 K pairs. While anchors were not widely deployed back then, we can now leverage the data generated by 810 anchors without having to schedule a single measurement on our own.

We find that from the 810 active anchors, only 770 generated valid traceroutes during our chosen day. In total, the measurements result in 292.7 K anchor pairs and almost 28 M traceroutes suitable for analysis. A detailed analysis of symmetry is beyond the scope of this paper, but we provide a preliminary estimation by comparing

the length of traceroutes in both directions. If the length is different, the traceroutes cannot be symmetric. We use two length metrics: i) the number of IP hops and ii) the number of unique ASes traversed by each traceroute. The analysis reveals that only 21% of traceroutes could be symmetric, having the same length in terms of IP hops in both directions (see Figure 11). The majority of traceroutes is longer by one hop or more in one direction, making symmetry impossible. However, considering the length at the AS granularity, around 50% of traceroutes traverse the same number of ASes in both directions and are potentially symmetric.

5.2 Built-in Measurements

Built-in measurements, which have been running on every connected RIPE Atlas host since 2010, offer a wealth of long-term data for exploration. In this section, we highlight practical use cases that illustrate the potential of these measurements, revealing insights into network behavior.

5.2.1 Manipulation of the Popular Domain Resolution. One of the built-in measurements queries local resolvers to obtain A and TXT resource records of 50 popular domain names. Previous work already relied on user-defined RIPE Atlas measurements to study DNS resolution manipulation [89]. Below, we show how one can detect response injection using built-ins only.

We first consulted the Open Observatory of Network Interference (OONI) [11] project and identified that 17 out of 50 queried domains show strong indications of blocking in at least one country. For each of them, we parsed the responses to A requests and extracted the returned IPv4 addresses. We cannot, however, create one-to-one mappings between the domains and A records as the DNS load balancing systems return different IP addresses depending on the location and the time of the query. Therefore, we estimate a range of allowed responses by first resolving each domain directly at <https://dns.google> (Google Public Resolver) and then mapping responses to ASes [12]. We consider a response as valid if it falls within the range of ASes seen in responses from Google Public Resolver. While this non-strict approach may potentially underestimate the amount of blocking (i.e., IP belonging to an observed AS is not, in fact, a correct address for a particular domain), it does highlight the cases for which the answer is incorrect. We also note that other strategies exist to restrict access to a particular domain at the DNS level, for example, returning the NXDOMAIN or SERVFAIL response codes. We argue, however, that the codes are indistinguishable from intermittent failures, while response injection is a strong indication of deliberate manipulation.

The highest number of probes received injected responses for two Meta services: `facebook.com` and `instagram.com`. They were blocked on at least 69% of probes from China and Iran, mostly receiving globally routable IPs that belong to well-known organizations such as Twitter, Dropbox, or Amazon. Such a behavior is consistent with the effects of the Great Firewall of China [50]. The two countries restricted access to other entertainment domains such as `youtube.com`, `reddit.com`, `twitter.com`, and `netflix.com`, as it can be seen on more than half of the countries' probes. The third most blocked domain was `vk.com`, a Russian social media network with 200 affected RIPE Atlas probes, 60 of which were from Iran alone (68% of the country probes). The two countries

formally blocking this domain name, namely Ukraine and Latvia, have 18% and 37% of probes receiving injected responses, respectively. Some of the other domain names triggering injected responses are `google.com`, `google.co.in`, and `linkedin.com`—they were blocked on the majority of probes located in China. Similarly, `t.co` and `wordpress.com` were not reachable for the majority of RIPE Atlas probes located in Iran.

Overall, we demonstrate how RIPE Atlas built-in measurements can be used to study DNS interception, manipulation, and censorship of popular domains, complementing some of the existing tools like OONI. They can give insight into blocking strategies used on a per-probe basis, e.g., injecting globally routable or private IP addresses. When interpreting results, it is crucial to take into account the number of probes per country to avoid false inference of nationwide blocking.

5.2.2 Unallocated 240/4 in the Wild. The IPv4 space has many unallocated or reserved addresses with the largest block of 240/4. Two proposals suggested repurposing 240/4 into unicast [41] or private [123] address spaces. Neither proposal was adopted though, as critics argued that these additional addresses would be quickly exhausted, emphasizing the need to transition to IPv6. Previous work revealed that some organizations internally use the IP address space allocated to others but never announced on the Internet [106]. Below, we analyze all the built-in traceroutes collected on 21 February 2024 to see whether the same phenomenon is observed for unallocated address ranges.

Overall, we identified 1.7 M traceroutes containing 240/4 addresses. Table 4 further lists all the autonomous systems hosting the probes from which the traceroutes originate. Similarly to previous reports [72], the majority of them (92.7%) originate from two Amazon ASes – AS16509 and AS14618. We saw unallocated IPs in the first and second hops, indicating the constant internal use of this prefix by Amazon. Similar behavior was observed at traceroutes originating from Telefonica Spain (AS3352), where hops subsequent to 240/4 remained inside the same AS. Despite the previous hops timing out, we can still conclude that Telefonica uses the 240/4 address space internally. Traceroutes originating from the remaining autonomous systems (AS8728, AS577, AS398721) seem to merely traverse the Amazon AS that makes 240/4 appear in some of the transit hops.

The use of the unallocated/squatted address space is strongly discouraged, as it might render parts of the Internet unavailable should the concerned IP addresses become globally routable. RIPE Atlas measurements reveal that at least two organizations use the unallocated 240/4 address block internally. Another key finding is that all the networks we observed in the traceroutes do not have proper filtering in place, as packets with addresses in the 240/4 prefix should be dropped.

5.2.3 Unspecified ::/128 as a Source Address. Some of the IP addresses are called special use and are not intended to be assigned to hosts. One such example is `::/128`, called *unspecified* and prohibited from being used neither as a source nor as a destination [32]. Nevertheless, it was reported that some of the traceroutes collected by RIPE Atlas contain hops with `::/128` source addresses [61]. We verify whether the same behavior persisted on the day of the analysis.

Table 4: The number of traceroutes containing 240/4 hops per originating autonomous systems.

AS Number	Organization	Country	Number of Traceroutes
16509	Amazon.com, Inc.	US	1,601,758
14618	Amazon.com, Inc.	US	106,888
3352	TELEFONICA DE ESPANA S.A.U.	ES	10,128
8728	AS INFONET	EE	8,502
577	Bell Canada	CA	458
398721	Cogeco Connexion inc	CA	30

On February 21, 2024, RIPE Atlas gathered 95 M IPv6 traceroute results. We found that 334 K out of them (0.35%) include at least one hop with the `::/128` source address. All traceroutes originated from 161 probes, including 1 probe from NeuStyle (AS4508) responsible for 256 K (76%) cases observed. This probe has `::/128` consistently appearing as the first hop. In the remaining cases, we mapped the IP addresses of previous and subsequent hops to autonomous system numbers. We reveal several traces with AS24186 (RailTel Corporation of India Ltd.) surrounding the `::/128` hop. Previously, it was AS9198 (JSC Kazakhtelecom) that was seen before and after the `::/128` hop. Although we do not see the same AS in our dataset, we observed some instances where the hop preceding the address belonged to either AS12389 or AS20485, and the subsequent hop was in AS9198 indicating that the phenomenon is still persistent in this network.

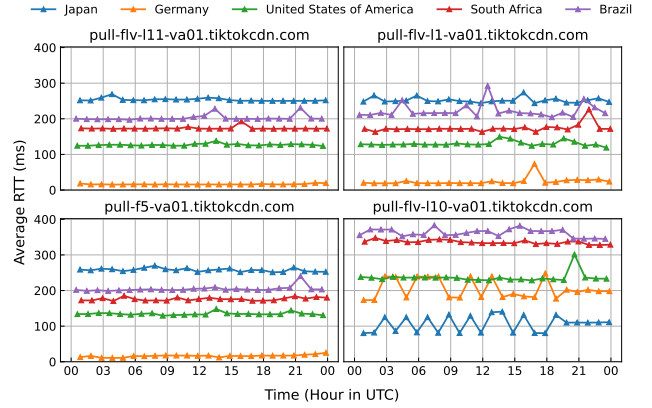
Overall, our analysis reveals that some ASes violate the requirements of the RFC 4291 and route packets with unspecified source addresses. This behavior also reveals the lack of Source Address Validation (SAV) [71] in these networks.

5.3 User-Defined Measurements

The two previously discussed groups of measurements are well-known and have been defined by the RIPE Atlas team. Below, we look at some of the user-defined measurements to understand what they may reveal.

5.3.1 TikTok CDN. We observed one noticeable user-defined measurement that involved periodic pings to four TikTok CDN servers. They account for 25% (4,863) of user-defined ping measurements but generated only 26.5 K results (less than 0.1%) because of the way the measurements are set up: the four servers are targeted from 50 countries in hourly cycles. However, for reasons unknown to us, the creators of the campaign chose to create individual one-off measurements instead of one long-running periodic measurement per country. Our best guess is that the creators use a custom scheduling setup, supported by the fact that even though measurements are run per country (indicated by a “Ping from *country*” description), they do not use the country selection provided by Atlas, but request individual probes by ID instead. Thus, the measurements also significantly contribute to the probe source distribution shown in Figure 9. We speculate that—for the most part—using periodic measurements instead of one-off would have been more economical for this purpose.

A rough search reveals that this campaign has been running since June 2023 and created 2.5 M measurements since then. The frequency (hourly) and the type of measurement (ping) make this dataset useful to track long-term changes in RTT from different

**Figure 12: Average RTT from Japan, Germany, United States, South Africa, and Brazil to four TikTok CDN servers.**

countries to the CDN servers. Figure 12 presents the average RTT observed from five countries during our analyzed day. It is apparent that the connectivity was stable during this day and the server `flv-110` seems to be located in the Asian region, since the average RTT from Japan is the lowest for this server, whereas the country order is stable for the other three targets. A deeper analysis of this dataset might reveal intermittent connectivity problems from certain countries or long-term changes in RTT.

5.3.2 IXP Country and Regional Jedi Measurements. RIPE Atlas team set up a number of traceroute meshes that aim at analyzing the connectivity between probes located inside the same country and region. These measurements offer critical information on the paths taken, including out-of-country detours and the IXPs the traffic traverses.

Country-level meshes are executed on the first day of each month. These are regular and up-to-date measurements that can be used to analyze how traffic patterns differ between countries. For example, they reveal the presence of 17 IXPs operating in Germany and 14 IXPs in the Netherlands. In the United States, they identify 66 unique peering LANs in probe paths, associated with 45 unique IXPs. Conversely, no IXPs are identified in Uzbekistan, China, and Iran. However, the traceroutes between ASes in China and Uzbekistan predominantly pass through the incumbent providers `Uzbektelecom` and the China Backbone network for the majority of the traceroutes.

Regional meshes have been collecting traceroutes in Southeast Europe, Central Asia, Latin America, and the Middle East on the 15th of every month since January 2024. These measurements originate from several countries within each region, highlighting the involvement of intermediary countries and ASes. We therefore focus on the data collected in the Central Asia region, including Uzbekistan, Kazakhstan, Iran, Kyrgyzstan, and Tajikistan. Overall, we see that Central Asian countries heavily rely on intermediaries outside the region to communicate with other hosts inside. Russia is by far the country that received the highest ratio of transit traffic, ranging from 42.4% of traceroutes in Kazakhstan to as many as 72.6% in Tajikistan and 80.7% in Iran. The autonomous systems involved are

PJSC Rostelecom (AS12389), PJSC “Vimpelcom” (AS3216), and Inetcom LLC (AS35598). Azerbaijan received the second highest amount of traffic, always transiting through Delta Telecom Ltd (AS29049) from Uzbekistan (19.2% of traceroutes), Kazakhstan (24.2%), and Iran (46.9%). We do see a much smaller number of cases in which traffic goes to ASes registered in as far as Sweden, Italy, and the United States.

Based on the regional traceroute meshes collected by RIPE Atlas, Central Asia is highly dependent on several foreign countries and ASes to transit its traffic. Such a lack of connectivity options is generally expensive, not flexible, and may result in a single point of failure. The five aforementioned countries will benefit from consolidating efforts to build a more diverse internal interconnection.

6 GUIDELINES FOR USING RIPE ATLAS

RIPE Atlas is a powerful tool that executes diverse measurements and generates over a billion results daily. As a measurement community reliant on its capabilities, we have a unique opportunity to optimize our use of its resources. Therefore, we offer guidelines for researchers to enhance their usage of RIPE Atlas, which we proposed to and extensively discussed with the RIPE Atlas team.

6.1 Optimize the Infrastructure Usage

RIPE Atlas operates with a set of safeguards to ensure that no single user or measurement jeopardizes the whole platform. Apart from adhering to internally enforced measurement constraints, researchers must consider the volume of generated data and avoid performing unnecessary measurements, especially for extended periods beyond those necessary to address their research questions.

One might feel compelled to launch a new measurement campaign as soon as the research question is clearly defined. However, we recommend stepping back and carefully considering whether any existing measurements can serve the purpose instead. The platform offers a variety of built-ins and user-defined measurements (documented in Measurement Bundles [100] repository), which already run on RIPE Atlas. Reusing existing data helps avoid redundant efforts. Each research paper using RIPE Atlas must include a statement addressing whether built-in measurements were considered to answer the research question and justify the need for any new measurements.

Other examples of suboptimal behaviors have already been provided [59], such as failing to use a single recurring measurement instead of multiple one-offs or polling for results more frequently than they are generated. Carefully reading the relevant documentation will help avoid such non-optimal behaviors.

6.2 Consider Ethics

In addition to relevant guidelines and best practices for the measurement community [35, 91], authors should carefully consider specific ethical implications related to the RIPE Atlas platform to minimize unforeseen consequences for all users.

If a new measurement is deemed necessary, researchers must assess its ethical aspect from multiple angles. First, the Terms and Conditions [101] of the platform must serve as a starting point to understand the expectations from all participants. In particular, it states that probe owners agree that the installed devices are used for

performing measurements and obtaining the results. Yet, one must ensure that no risk is put on the hosts. For example, avoid sending DNS requests to sensitive domain names that may be banned under certain jurisdictions such as related to gambling or adult content. We also recommend that researchers consult the RIPE community discussion on measurement ethics [57].

6.3 Encourage Reproducibility

RIPE Atlas offers various mechanisms, such as tags and descriptions, to facilitate finding relevant measurements and foster reproducibility. We strongly advise researchers to assign unique tags and provide comprehensive descriptions for all new measurements. They, along with measurement IDs, should be referenced in scientific publications for easy data access. Furthermore, documenting new user-defined measurement campaigns in the Measurement Bundles repository is encouraged to inform others and promote the reuse of existing measurements whenever feasible.

6.4 Explore the RIPE Atlas Data

RIPE Atlas is often used to investigate specific research questions or validate hypotheses. However, given its extensive measurement dataset, there may be numerous events and insights beyond one’s immediate focus. We encourage researchers to explore existing RIPE Atlas measurements thoroughly and uncover findings that could benefit the network community. We firmly believe that this rich dataset holds significant potential for discovering new insights and phenomena.

6.5 Host a Probe

Participants around the globe play a crucial role in expanding the global presence of RIPE Atlas by hosting probes—essential components of the measurement network. We strongly encourage researchers and network operators to consider applying to host either a hardware or a software probe, which can be installed and managed on their own infrastructure. Active involvement from users in underrepresented regions is particularly encouraged. For detailed instructions, interested researchers can visit the official RIPE Atlas web page.

7 ETHICS AND REPRODUCIBILITY

Large-scale Internet measurements typically do not fall under the purview of Institutional Review Boards. Therefore, researchers must carefully weigh the advantages and disadvantages before initiating any new measurement campaign. To aid in this process, the community has developed a set of best practices to adhere to [13, 35, 91].

This paper strictly adheres to the RIPE Atlas Terms and Conditions regarding data retrieval, analysis, and presentation. Specifically, all network configurations, measurements, and probes discussed are publicly accessible, as agreed upon by probe hosts in §4.3 and §4.2 of the terms. We solely use existing datasets to uncover various phenomena observed on the analysis day, ensuring an unbiased exploration. Recognizing the significant data volume, we carefully refer to relevant documentation to prevent overloading the RIPE Atlas infrastructure with our requests.

8 CONCLUSIONS

RIPE Atlas stands as one of the largest network measurement platforms worldwide, executing 50.9 K unique measurements and generating over 1.3 billion results daily. This unprecedented scale presents challenges in understanding its diverse user base, methodologies, and objectives.

We examined a typical day in the life of RIPE Atlas to uncover usage patterns. The majority of daily measurements are user-defined, with nearly half being one-off tests. Anchors contribute significantly, generating almost 70% of daily results, surpassing both built-in (21.1%) and user-defined measurements (11.4%). Pings, traceroutes, and DNS tests dominate, with HTTP, TLS, and NTP lagging behind. End users primarily perform one-off measurements and do not heavily burden the system.

Probes and anchors have steadily expanded the RIPE Atlas network since 2010, providing robust baselines for Internet analysis. Around half of the public probes support IP dual-stack, while 92% of anchors are fully dual-stack. Despite uneven representation across regions, RIPE Atlas spans 178 countries and over 4,000 ASes.

Given the data volume, RIPE Atlas proves invaluable for diverse research and operational questions. With exclusively existing datasets, we have demonstrated how anchoring, built-in, and user-defined measurements can illuminate traceroute symmetry, address space usage, DNS censorship, and more.

Lastly, after extensive discussions with the RIPE Atlas team, we propose guidelines for researchers leveraging RIPE Atlas: careful network usage, full reproducibility, maximal reuse of existing measurement data, and rigorous ethical considerations.

REFERENCES

- [1] Akhavan Niaki, A., Marczak, W., Farhoodi, S., McGregor, A., Gill, P., Weaver, N.: Cache Me Outside: A New Look at DNS Cache Probing. In: PAM. Springer International Publishing (2021)
- [2] Al-Dalky, R., Rabinovich, M., Schomp, K.: A Look at the ECS Behavior of DNS Resolvers. In: IMC. ACM (2019)
- [3] Alaraj, A., Bock, K., Levin, D., Wustrow, E.: A Global Measurement of Routing Loops on the Internet. In: PAM. Springer Nature Switzerland (2023)
- [4] Albakour, T., Gasser, O., Beverly, R., Smaragdakis, G.: Third Time's Not a Charm: Exploiting SNMPv3 for Router Fingerprinting. In: IMC. ACM (2021)
- [5] Albakour, T., Gasser, O., Beverly, R., Smaragdakis, G.: Illuminating Router Vendor Diversity Within Providers and Along Network Paths. In: IMC. ACM (2023)
- [6] Anderson, S., Salamatian, L., Bischof, Z.S., Dainotti, A., Barford, P.: iGDB: Connecting the Physical and Logical Layers of the Internet. In: IMC. ACM (2022)
- [7] Antonakakis, M., April, T., Bailey, M., Bernhard, M., Bursztein, E., Cochran, J., Durumeric, Z., Halderman, J.A., Invernizzi, L., Kallitsis, M., Kumar, D., Lever, C., Ma, Z., Mason, J., Menscher, D., Seaman, C., Sullivan, N., Thomas, K., Zhou, Y.: Understanding the Mirai Botnet. In: USENIX Security. USENIX Association (Aug 2017)
- [8] Appel, M., Aben, E., Fontugne, R.: Metis: Better Atlas Vantage Point Selection for Everyone. In: TMA. IFIP (2022)
- [9] Arnold, T., Gürmerçiler, E., Essig, G., Gupta, A., Calder, M., Giotsas, V., Katz-Bassett, E.: (How Much) Does a Private WAN Improve Cloud Performance? In: IEEE INFOCOM. IEEE (2020)
- [10] Arnold, T., He, J., Jiang, W., Calder, M., Cunha, I., Giotsas, V., Katz-Bassett, E.: Cloud Provider Connectivity in the Flat Internet. In: IMC. ACM (2020)
- [11] Arturo Filastò and Jacob Appelbaum: OONI: Open Observatory of Network Interference. In: 2nd USENIX Workshop on Free and Open Communications on the Internet. pp. 1–8. USENIX Association (2012)
- [12] Asghari, H.: pyasn (2020), <https://github.com/hadiasghari/pyasn>
- [13] Bailey, M., Kenneally, E., Maughan, D., Dittrich, D.: The Menlo Report. IEEE Security & Privacy 10(02) (Mar 2012)
- [14] Bajpai, V., Eravuchira, S.J., Schönwälder, J.: Lessons Learned From Using the RIPE Atlas Platform for Measurement Research. SIGCOMM Comput. Commun. Rev. 45(3) (jul 2015)
- [15] Bhowmick, P., Ashiq, M.I., Deccio, C., Chung, T.: TTL Violation of DNS Resolvers in the Wild. In: PAM. Springer-Verlag (2023)
- [16] Böck, L., Levin, D., Padmanabhan, R., Doerr, C., Mühlhäuser, M.: How to Count Bots in Longitudinal Datasets of IP Addresses. In: NDSS. ISOC (2023)
- [17] Bremner-Barr, A., Hay, D., Meyuhas, B., Danino, S.: It's Not Where You Are, It's Where You Are Registered: IoT Location Impact on MUD. In: ANRW. ACM (2023)
- [18] Brown, A.: Facebook Lost About \$65 Million During Hours-Long Outage. <https://www.forbes.com/sites/abrambrown/2021/10/05/facebook-outage-lost-revenue/?sh=565a1a88231a> (2021)
- [19] Calder, M., Fan, X., Zhu, L.: A Cloud Provider's View of EDNS Client-Subnet Adoption. In: TMA. IEEE (2019)
- [20] Candela, M.: Visualising DNS Issues with DNSMON. https://labs.ripe.net/author/massimo_candela/visualising-dns-issues-with-dnsmon/ (2024)
- [21] Candela, M., Gregori, E., Luconi, V., Vecchio, A.: Dissecting the Speed-of-Internet of Middle East. In: IEEE INFOCOM Workshops. IEEE (2019)
- [22] Charyyev, B., Gunes, M.H.: Latency Characteristics of Edge and Cloud. In: TMA Lightning Talks Session. IFIP (2020)
- [23] Chen, W., Wang, Z., Han, D., Duan, C., Yin, X., Yang, J., Shi, X.: ROV-MI: Large-Scale, Accurate and Efficient Measurement of ROV Deployment. In: NDSS. ISOC (2022)
- [24] Chhabra, R., Murley, P., Kumar, D., Bailey, M., Wang, G.: Measuring DNS-over-HTTPS Performance Around the World. In: IMC. ACM (2021)
- [25] Chung, T., Choffnes, D., Mislove, A.: Tunneling for Transparency: A Large-Scale Analysis of End-to-End Violations in the Internet. In: IMC. ACM (2016)
- [26] Chung, T., van Rijswijk-Deij, R., Chandrasekaran, B., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Wilson, C.: A Longitudinal, End-to-End View of the DNSSEC Ecosystem. In: USENIX Security. USENIX Association (2017)
- [27] Darwich, O., Rimlinger, H., Dreyfus, M., Gouel, M., Vermeulen, K.: Replication: Towards a Publicly Available Internet Scale IP Geolocation Dataset. In: IMC. ACM (2023)
- [28] Davies, A.: 2018 Campaign to Sponsor 10 RIPE Atlas Anchors. https://labs.ripe.net/author/alun_davies/2018-campaign-to-sponsor-10-ripe-atlas-anchors/ (2018)
- [29] Davies, A.: RIPE Atlas Anchors 400+. https://labs.ripe.net/author/alun_davies/ripe-atlas-anchors-400/ (Jan 2019)
- [30] Davisson, L., Jakovleski, J., Ngo, N., Pham, C., Sommers, J.: Reassessing the Constancy of End-to-End Internet Latency. In: TMA. IFIP (2021)
- [31] De Vries, W., Santanna, J.J., Sperotto, A., Pras, A.: How Asymmetric Is the Internet?: A Study to Support the Use of Traceroute. In: Intelligent Mechanisms for Network Configuration and Security. Springer International Publishing (2015)
- [32] Deering, D.S.E., Hinden, B.: IP Version 6 Addressing Architecture. RFC 4291 (Feb 2006), <https://www.rfc-editor.org/info/rfc4291>
- [33] Doan, T.V., Tsareva, I., Bajpai, V.: Measuring DNS over TLS from the Edge: Adoption, Reliability, and Response Times. In: PAM. Springer International Publishing (2021)
- [34] Du, B., Candela, M., Huffaker, B., Snoeren, A.C., Claffy, K.: RIPE IPmap Active Geolocation: Mechanism and Performance Evaluation. SIGCOMM Comput. Commun. Rev. 50(2) (may 2020)
- [35] Durumeric, Z., Wustrow, E., Halderman, J.A.: ZMap: Fast Internet-Wide Scanning and Its Security Applications. In: USENIX Security. USENIX Association (2013)
- [36] Enayet, A., Heidemann, J.: Internet Outage Detection Using Passive Analysis. In: IMC. ACM (2022)
- [37] Fanou, R., Huffaker, B., Mok, R., Claffy, K.C.: Unintended Consequences: Effects of Submarine Cable Deployment on Internet Routing. In: PAM. Springer International Publishing (2020)
- [38] Fontugne, R., Bautista, E., Petrie, C., Nomura, Y., Abry, P., Goncalves, P., Fukuda, K., Aben, E.: BGP Zombies: An Analysis of Beacons Stuck Routes. In: PAM. Springer International Publishing (2019)
- [39] Fontugne, R., Phokeer, A., Pelsser, C., Vermeulen, K., Bush, R.: RPKI Time-of-Flight: Tracking Delays in the Management, Control, and Data Planes. In: PAM. Springer Nature Switzerland (2023)
- [40] Fontugne, R., Shah, A., Cho, K.: Persistent Last-mile Congestion: Not so Uncommon. In: IMC. ACM (2020)
- [41] Fuller, V.: Reclassifying 240/4 as usable unicast address space. Internet-Draft draft-fuller-240space-02, Internet Engineering Task Force (Mar 2008), <https://datatracker.ietf.org/doc/draft-fuller-240space-02/>, work in Progress
- [42] Gamero-Garrido, A., Belding, E., Choffnes, D.: Using Reverse IP Geolocation to Identify Institutional Networks. In: IMC. ACM (2022)
- [43] Gamero-Garrido, A., Carisimo, E., Hao, S., Huffaker, B., Snoeren, A.C., Dainotti, A.: Quantifying Nations' Exposure to Traffic Observation and Selective Tampering. In: PAM. Springer International Publishing (2022)
- [44] Gao, Z., Venkataramani, A.: Measuring Update Performance and Consistency Anomalies in Managed DNS Services. In: IEEE INFOCOM. IEEE (2019)
- [45] Giotsas, V., Koch, T., Fazzion, E., Cunha, I., Calder, M., Madhyastha, H.V., Katz-Bassett, E.: Reduce, Reuse, Recycle: Repurposing Existing Measurements to Identify Stale Traceroutes. In: IMC. ACM (2020)

- [46] Google Scholar: "RIPE Atlas". https://scholar.google.com/scholar?hl=en&as_sdt=0%2C5&q=RIPE+Atlas&btnG= (May 2024)
- [47] Gouel, M., Darwich, O., Mouchet, M., Vermeulen, K.: Poster: Towards a Publicly Available Framework to Process Traceroutes with MetaTrace. In: IMC. ACM (2023)
- [48] He, Y., Faloutsos, M., Krishnamurthy, S., Huffaker, B.: On Routing Asymmetry in the Internet. In: IEEE GLOBECOM. IEEE (2005)
- [49] Hlavacek, T., Cunha, Í., Gilad, Y., Herzberg, A., Katz-Bassett, E., Schapira, M., Schulmann, H.: DISCO: Sidestepping RPKI's Deployment Barriers. In: NDSS. ISOC (2020)
- [50] Hoang, N.P., Niaki, A.A., Dalek, J., Knockel, J., Lin, P., Marczak, B., Crete-Nishihata, M., Gill, P., Polychronakis, M.: How Great is the Great Firewall? Measuring Chinas DNS Censorship. In: USENIX Security. USENIX Association (Aug 2021)
- [51] Holterbach, T., Pelsner, C., Bush, R., Vanbever, L.: Quantifying Interference between Measurements on the RIPE Atlas Platform. In: IMC. ACM (2015)
- [52] Homburg, P.: Troubleshooting RIPE Atlas Probes: USB Sticks. https://labs.ripe.net/author/philip_homburg/troubleshooting-ripe-atlas-probes-usb-sticks/ (Jul 2016)
- [53] Hou, B., Cai, Z., Wu, K., Su, J., Xiong, Y.: 6Hit: A Reinforcement Learning-based Approach to Target Generation for Internet-wide IPv6 Scanning. In: IEEE INFOCOM. IEEE (2021)
- [54] Janardhan, S.: More details about the October 4 outage. <https://engineering.fb.com/2021/10/05/networking-traffic/outage-details/> (2021)
- [55] Jia, Y., Kuzmanovic, A.: Perceiving Internet Anomalies via CDN Replica Shifts. In: IEEE INFOCOM. IEEE (2019)
- [56] Kiedanski, D., Nogueira, M., Grampin, E.: Youtube Traffic From the Perspective of a Developing Country: the Case of Uruguay. In: IEEE INFOCOM Workshops. IEEE (2019)
- [57] Kistelevi, R.: Ethics of RIPE Atlas Measurements. <https://labs.ripe.net/author/kistel/ethics-of-ripe-atlas-measurements/> (2016)
- [58] Kistelevi, R.: RIPE Atlas and IoT. <https://www.ripe.net/media/documents/20170921-iot-atlas.pdf> (Sep 2017)
- [59] Kistelevi, R.: Observations on RIPE Atlas API "Anti-Patterns". <https://labs.ripe.net/author/kistel/observations-on-ripe-atlas-api-anti-patterns/> (2023)
- [60] Kistelevi, R.: The Role of Aggregators in RIPE Atlas. <https://labs.ripe.net/author/kistel/the-role-of-aggregators-in-ripe-atlas/> (2023)
- [61] Kistelevi, R., Lone, Q., Stam, M.: The Curious Case of Packets From ::. <https://labs.ripe.net/author/kistel/the-curious-case-of-packets-from/> (May 2022)
- [62] Kistelevi, R., de Weerd, P.: RIPE NCC Measurement Data Retention Principles. <https://labs.ripe.net/author/kistel/ripe-ncc-measurement-data-retention-principles/> (Nov 2023)
- [63] Koch, T., Katz-Bassett, E., Heidemann, J., Calder, M., Ardi, C., Li, K.: Anycast in Context: a Tale of Two Systems. In: SIGCOMM. ACM (2021)
- [64] Koch, T., Yu, S., Agarwal, S., Katz-Bassett, E., Beckett, R.: PAINTER: Ingress Traffic Engineering and Routing for Enterprise Cloud Networks. In: SIGCOMM. ACM (2023)
- [65] Kon, P., Gattani, A., Saharia, D., Cao, T., Barradas, D., Chen, A., Sherr, M., Ujcich, B.E.: NetShuffle: Circumventing Censorship with Shuffle Proxies at the Edge. In: IEEE S&P. IEEE Computer Society (May 2024)
- [66] Li, J., Zhou, S., Gotsas, V.: Performance Analysis of Multipath BGP. In: IEEE INFOCOM Workshops. IEEE (2021)
- [67] Li, W., Lin, Z., Ashiq, M.I., Aben, E., Fontugne, R., Phokeer, A., Chung, T.: RoVista: Measuring and Analyzing the Route Origin Validation (ROV) in RPKI. In: IMC. ACM (2023)
- [68] Liu, S., Bischof, Z.S., Madan, I., Chan, P.K., Bustamante, F.E.: Out of Sight, Not Out of Mind: A User-View on the Criticality of the Submarine Cable Network. In: IMC. ACM (2020)
- [69] Livadariu, I., Dreiholz, T., Al-Selwi, A.S., Bryhni, H., Lysne, O., Bjørnstad, S., Elmokashfi, A.: On the Accuracy of Country-Level IP Geolocation. In: ANRW. ACM (2020)
- [70] Loe, A.F., Quaglia, E.A.: You Shall Not Join: A Measurement Study of Cryptocurrency Peer-to-Peer Bootstrapping Techniques. In: CCS. ACM (2019)
- [71] Lone, Q., Frik, A., Luckie, M., Korczyński, M., van Eeten, M., Gañán, C.: Deployment of Source Address Validation by Network Operators: A Randomized Control Trial. In: IEEE S&P (2022)
- [72] Lone, Q., Gilmore, J., Schoen, S.D., Täht, D., Aben, E.: 240/4 As Seen by RIPE Atlas. <https://labs.ripe.net/author/qasim-lone/240-4-as-seen-by-ripe-atlas/> (Aug 2022)
- [73] Magnusson, J., Müller, M., Brunstrom, A., Pulls, T.: A Second Look at DNS QNAME Minimization. In: PAM. Springer-Verlag (2023)
- [74] Mao, J., Rabinovich, M., Schomp, K.: Assessing Support for DNS-over-TCP in the Wild. In: PAM. Springer-Verlag (2022)
- [75] Martin, N., Dogar, F.: Divided at the Edge - Measuring Performance and the Digital Divide of Cloud Edge Data Centers. *Proc. ACM Netw.* **1**(CoNEXT3) (Nov 2023)
- [76] Mazzola, F., Marcos, P., Castro, I., Luckie, M., Barcellos, M.: On the Latency Impact of Remote Peering. In: PAM. Springer International Publishing (2022)
- [77] McQuistin, S., Uppu, S.P., Flores, M.: Taming Anycast in the Wild Internet. In: IMC. ACM (2019)
- [78] Mi, X., Feng, X., Liao, X., Liu, B., Wang, X., Qian, F., Li, Z., Alrwais, S., Sun, L., Liu, Y.: Resident Evil: Understanding Residential IP Proxy as a Dark Service. In: IEEE S&P (2019)
- [79] Michel, F., Trevisan, M., Giordano, D., Bonaventure, O.: A First Look at Starlink Performance. In: IMC. ACM (2022)
- [80] Milolidakis, A., Fontugne, R., Dimitropoulos, X.: Detecting Network Disruptions at Colocation Facilities. In: IEEE INFOCOM. IEEE (2019)
- [81] Mouchet, M., Vaton, S., Chonavel, T.: Poster Abstract: A Flexible Infinite HMM Model for Accurate Characterization and Segmentation of RTT Timeseries. In: IEEE INFOCOM Workshops. IEEE (2019)
- [82] Moura, G.C.M., Castro, S., Heidemann, J., Hardaker, W.: TsuNAME: Exploiting Misconfiguration and Vulnerability to DDoS DNS. In: IMC. ACM (2021)
- [83] Moura, G.C.M., Heidemann, J., Hardaker, W., Charnethikul, P., Bulten, J., Ceron, J.M., Hesselman, C.: Old but Gold: Prospecting TCP to Engineer and Live Monitor DNS Anycast. In: PAM. Springer International Publishing (2022)
- [84] Moura, G.C.M., Heidemann, J., Schmidt, R.D.O., Hardaker, W.: Cache Me If You Can: Effects of DNS Time-to-Live. In: IMC. ACM (2019)
- [85] Moura, G.C.M., Müller, M., Davids, M., Wullink, M., Hesselman, C.: Fragmentation, Truncation, and Timeouts: Are Large DNS Messages Falling to Bits? In: PAM. Springer International Publishing (2021)
- [86] Müller, M., Thomas, M., Wessels, D., Hardaker, W., Chung, T., Toorop, W., Rijswijk-Deij, R.v.: Roll, Roll, Roll your Root: A Comprehensive Analysis of the First Ever DNSSEC Root KSK Rollover. In: IMC. ACM (2019)
- [87] Müller, M., Toorop, W., Chung, T., Jansen, J., van Rijswijk-Deij, R.: The Reality of Algorithm Agility: Studying the DNSSEC Algorithm Life-Cycle. In: IMC. ACM (2020)
- [88] NLNOG: Introduction - NLNOG RING. <https://ring.nlnog.net> (2024)
- [89] Nosyk, Y., Lone, Q., Zhauniarovich, Y., Gañán, C.H., Aben, E., Moura, G.C.M., Talizadehkhoo, S., Duda, A., Korczyński, M.: Intercept and Inject: DNS Response Manipulation in the Wild. In: PAM. Springer Nature Switzerland (2023)
- [90] Ozcan, S., Livadariu, I., Smaragdakis, G., Griwodz, C.: Longitudinal Analysis of Inter-City Network Delays. In: TMA. IEEE (2023)
- [91] Partridge, C., Allman, M.: Ethical Considerations in Network Measurement Papers. *Commun. ACM* **59**(10) (Sep 2016)
- [92] Raman, A., Varvello, M., Chang, H., Sastry, N., Zaki, Y.: Dissecting the Performance of Satellite Network Operators. *Proc. ACM Netw.* **1**(CoNEXT3) (Nov 2023)
- [93] Raman, P., Flores, M.: Building Out the Basics with Hoplets. In: PAM. Springer-Verlag (2021)
- [94] Raman, R.S., Evdokimov, L., Wurstraw, E., Halderman, J.A., Ensafi, R.: Investigating Large Scale HTTPS Interception in Kazakhstan. In: IMC. ACM (2020)
- [95] Ramanathan, S., Hossain, A., Mirkovic, J., Yu, M., Afroz, S.: Quantifying the Impact of Blocklisting in the Age of Address Reuse. In: IMC. ACM (2020)
- [96] Randall, A., Liu, E., Akiwate, G., Padmanabhan, R., Voelker, G.M., Savage, S., Schulman, A.: Trufflehunter: Cache Snooping Rare Domains at Large Public DNS Resolvers. In: IMC. ACM (2020)
- [97] Randall, A., Liu, E., Padmanabhan, R., Akiwate, G., Voelker, G.M., Savage, S., Schulman, A.: Home is Where the Hijacking Is: Understanding DNS Interception by Residential Routers. In: IMC. ACM (2021)
- [98] RIPE Atlas: Built-in Measurements. <https://atlas.ripe.net/docs/built-in-measurements/> (2024)
- [99] RIPE Atlas: Starting your own Measurements (User-defined Measurements). <https://atlas.ripe.net/docs/getting-started/user-defined-measurements.html> (2025)
- [100] RIPE Atlas Community: Measurement Bundles. <https://github.com/RIPE-Atlas-Community/ripe-atlas-tips-and-tricks/wiki/Measurement-Bundles> (2024)
- [101] RIPE NCC: RIPE Atlas Service Terms and Conditions. <https://www.ripe.net/about-us/legal/ripe-atlas-service-terms-and-conditions/> (2020)
- [102] RIPE NCC Staff: RIPE Atlas: A Global Internet Measurement Network. *The Internet Protocol Journal* **18** (01 2015)
- [103] Rizvi, A.S.M., Bertholdo, L., Ceron, J., Heidemann, J.: Anycast Agility: Network Playbooks to Fight DDoS. In: USENIX Security. USENIX Association (Aug 2022)
- [104] Rodday, N., Cunha, Í.S., Bush, R., Katz-Bassett, E., Rodosek, G.D., Schmidt, T.C., Wählisch, M.: Revisiting RPKI Route Origin Validation on the Data Plane. In: TMA. IFIP (2021)
- [105] Rye, E.C., Beverly, R.: IPv6SeeYou: Exploiting Leaked Identifiers in IPv6 for Street-Level Geolocation. In: IEEE S&P. IEEE Computer Society (2023)
- [106] Salamati, L., Arnold, T., Cunha, I., Zhu, J., Zhang, Y., Katz-Bassett, E., Calder, M.: Who Squats IPv4 Addresses? *SIGCOMM Comput. Commun. Rev.* **53**(1) (Apr 2023)
- [107] Sattler, P., Aulbach, J., Zirngibl, J., Carle, G.: Towards a Tectonic Traffic Shift? Investigating Apple's New Relay Network. In: IMC. ACM (2022)

- [108] Scheitle, Q., Hohlfeld, O., Gamba, J., Jelten, J., Zimmermann, T., Strowes, S.D., Vallina-Rodriguez, N.: A Long Way to the Top: Significance, Structure, and Stability of Internet Top Lists. In: IMC. ACM (2018)
- [109] Sermpezis, P., Prehn, L., Kostoglou, S., Flores, M., Vakali, A., Aben, E.: Bias in Internet Measurement Platforms. In: TMA. IEEE (2023)
- [110] Shafir, L., Afek, Y., Bremler-Barr, A., Peleg, N., Sabag, M.: DNS Negative Caching in the Wild. In: SIGCOMM Posters and Demos. ACM (2019)
- [111] Shulman, H., Vogel, N., Waidner, M.: Poster: Insights into Global Deployment of RPKI Validation. In: CCS. ACM (2022)
- [112] Silveira, F.V.: Reducing the RIPE NCC's Data Centre Footprint. https://labs.ripe.net/author/felipe_victolla_silveira/reducing-the-ripe-nccs-data-centre-footprint/ (2024)
- [113] Smith, J.M., Birkeland, K., McDaniel, T., Schuchard, M.: Withdrawing the BGP Re-Routing Curtain: Understanding the Security Impact of BGP Poisoning through Real-World Measurements. In: NDSS. ISOC (2020)
- [114] Sommesse, R., Bertholdo, L., Akiwate, G., Jonker, M., van Rijswijk-Deij, R., Dainotti, A., Claffy, K., Sperotto, A.: MANycast2: Using Anycast to Measure Anycast. In: IMC. ACM (2020)
- [115] Sommesse, R., Moura, G.C.M., Jonker, M., van Rijswijk-Deij, R., Dainotti, A., Claffy, K.C., Sperotto, A.: When Parents and Children Disagree: Diving into DNS Delegation Inconsistency. In: PAM. Springer (2020)
- [116] Strowes, S.D., Aben, E., Wilhelm, R., Obser, F., Stagni, R., Formoso, A.: Debonising 2a10: : /12: Analysis of One Week's Visibility of a New /12. In: TMA. IFIP (2020)
- [117] Taylor, S.: RIPE Atlas: DomainMON is Here. https://labs.ripe.net/author/suzanne_taylor_muizin/ripe-atlas-domainmon-is-here/ (2015)
- [118] Thiagarajan, K., Kumar, R., Bustamante, F.E.: Poster: A Peek Backstage: Organizations in DNS Resolver Hierarchies. In: SIGCOMM. ACM (2023)
- [119] University of Oregon: Route Views Project (2020), <http://www.routeviews.org/routeviews/>
- [120] Vermeulen, K., Gurmericliler, E., Cunha, I., Choffnes, D., Katz-Bassett, E.: Internet Scale Reverse Traceroute. In: IMC. ACM (2022)
- [121] de Vries, W.B., Scheitle, Q., Müller, M., Toorop, W., Dolmans, R., van Rijswijk-Deij, R.: A First Look at QNAME Minimization in the Domain Name System. In: PAM. Springer International Publishing (2019)
- [122] Wei, L., Flores, M., Bedi, H., Heidemann, J.S.: Bidirectional Anycast/Unicast Probing (BAUP): Optimizing CDN Anycast. In: TMA. IFIP (2020)
- [123] Wilson, P., Michaelson, G.G., Huston, G.: Redesignation of 240/4 from "Future Use" to "Private Use". Internet-Draft draft-wilson-class-e-02, Internet Engineering Task Force (Sep 2008), <https://datatracker.ietf.org/doc/draft-wilson-class-e/02/>, work in Progress
- [124] Zhang, Z., Marder, A., Mok, R., Huffaker, B., Luckie, M., Claffy, K.C., Schulman, A.: Inferring Regional Access Network Topologies: Methods and Applications. In: IMC. ACM (2021)
- [125] Zhou, M., Zhang, X., Hao, S., Yang, X., Zheng, J., Chen, G., Dou, W.: Regional IP Anycast: Deployments, Performance, and Potentials. In: SIGCOMM. ACM (2023)
- [126] Zhu, J., Vermeulen, K., Cunha, I., Katz-Bassett, E., Calder, M.: The Best of Both Worlds: High Availability CDN Routing Without Compromising Control. In: IMC. ACM (2022)