

Project Description: Mid-scale RI-1 (M1:IP): Global Laboratory for Internet Measurement, Probing, and Security Evaluations (GLIMPSE)

1 Objective: Closing a critical gap in scientific Internet research

We propose a Mid-Scale RI-1 Implementation Project (M1:IP) to build an integrated research infrastructure that fills a critical gap in U.S. science and engineering: the absence of a platform for rigorous, scientific measurement of the global Internet.

The Internet, a vital backbone of modern society, faces relentless threats to its security and integrity. Safeguarding its availability, stability, and trustworthiness is a critical challenge for the U.S. government, intensified by adversarial actors leveraging the Internet in geopolitical conflicts [1, 2]. As a **national research priority** [3, 4, 5, 6, 7], Internet security demands robust data to drive solutions. Researchers, technologists, governments, and societal advocates urgently need a deeper, measurement-driven understanding of the Internet ecosystem to counter these threats effectively. Yet, accessible, relevant data is elusive.

Effective Internet measurement is inherently complex, facing interdisciplinary challenges across engineering, economics, law, and policy. Fragmented control and misaligned incentives—especially among private network operators wary of independent measurement—limit data sharing and impede collective understanding of Internet structure and security. While some private companies share data, such access depends on business priorities. The Internet’s decentralized design and fragmented private control prevent comprehensive, system-wide understanding of risks to national security and hinder U.S. scientific and engineering research. To address this gap, we propose a Mid-Scale RI-1 project to build a shared, independent measurement infrastructure that supports reproducible research, fosters scientific collaboration, and advances NSF’s mission to promote secure, open, and innovative cyberinfrastructure.

Understanding and securing the Internet begins with measurement. However, today, the research community still lacks the sustained, large-scale, and observational infrastructure to support measurement effectively. We propose to address this challenge by developing a next-generation research infrastructure that treats measurement, data curation, and accessibility as foundational capabilities. Informed by NSF’s Blueprint for National Cyberinfrastructure [8] and aligned with the National AI Research Resource (NAIRR) vision [9], our platform will integrate sustainable data acquisition, advanced AI capabilities, and scalable tools for processing, storage, metadata, and discovery. By democratizing access to a trustworthy Internet measurement platform and expertly curated data, our approach will enable transformative research across disciplines and advance the scientific basis for—and thus drive progress in—cybersecurity, policy, and infrastructure resilience.

Our Implementation Project will enable the following Research Community Benefits: technical and operational cybersecurity work force training; data-intensive methods for assessing security, stability, and resilience (SSR) properties of networks; and scientific and engineering advances to navigate other current and future harms to critical Internet infrastructure.

We structure our project in five tasks. The first task is to deploy a new global measurement infrastructure that can transform our ability to measure and analyze properties of the global Internet related to security and resilience. Our second task includes infrastructure for data management: data curation and documentation; and techniques for efficient data sharing, discovery, use, and dissemination. Our third task is to deploy proven analytics infrastructure and demonstrate its capabilities with case studies that yield strategic security-relevant global data sets. Our fourth task is to integrate the research infrastructure with national AI resources, to maximize the impact

of the infrastructure and the resulting data. The fifth task is to scale community and researcher engagement with the infrastructure, including STEM work force development activities that train students how to use the infrastructure and its data products.

We are nearing completion of an in-depth Design Phase, focused on measurement and data needs to navigate security vulnerabilities in the Internet’s packet carriage layer, specifically IP addressing, routing, DNS and certificate authority systems. These systems are foundational, require collective action to secure, and suffer from misaligned incentives to do so. In this Phase we designed several new data collection systems, which we codify in a specification document that will be a deliverable of the Design Phase. However, implementing all of the measurement infrastructure that we designed would require more resources than is available in the MSRI R1 (or even the R2) program. Furthermore, after much deliberation with stakeholders across sectors and disciplines, we recognize that data relevant to these vulnerabilities are inconsistently collected, and inconsistently available to independent researchers. While much data is collected, in many cases we lack the information necessary to support action toward improved Internet security. The wider variety of needed data types that we identified in the Design Phase [10], and the diverse and often unstable sources of data collection and curation, made clear to us that enabling availability of much of the data that would support translational security research will require an organization with stable funding at significant scale, such as an FFRDC. We have substantial thoughts on this direction that are beyond the scope of this proposal.

Thus, in consultation with research, industry, and government stakeholders, we distilled our Proposed Implementation to the types of data that can be collected by direct measurement of the Internet, while still providing critically needed data to stakeholders. Our most advanced achievement in the Design Phase was our prototype implementation of a Python-based integrated active measurement programming environment that exposes both a set of distributed *vantage point* (VPs), and a set of useful measurement primitives from which to build sophisticated measurement tools.

The key benefits of this design to *researchers* are that (1) the environment can provide reference implementations of measurement primitives that are difficult to implement correctly, making the environment useful especially for novice programmers, (2) the environment allows researchers to focus on the logic that ties a series of measurements together in an experiment, and (3) the logic is close to the VP, reducing experiment latency.

The key benefit to a *site host* is that the environment makes it difficult for a researcher to cause harm, intentionally or not, as researchers are restricted to the available measurements. The environment allows the *platform operator* to describe to the hosting site how researchers can use their VPs. However, researchers must rely on the environment maintainers and platform operators to expose useful measurement primitives and to keep the environment current with modern systems and evolving Internet protocols. These requirements make the Mid-Scale Research Infrastructure the ideal program to support our Implementation.

Through the Design Phase, we codified deep international partnerships with researchers in Europe and Japan who share our vision of large-scale Internet measurement infrastructure with resulting data being widely shared to enable scientific research. Our international partners (U. Twente in the Netherlands and Internet Initiative Japan (IIJ)) have complementary expertise and are supporting synergistic infrastructure projects, funded by their own institutions or national funding agencies. Our Project Execution Plan Outline (Supplementary Document) provides details of the roles and responsibilities of our international partners, along with details of our project management strategy. We also have international research infrastructure partners participating in measurement vantage point deployment.

2 Intellectual merit

The *scientific justification* of our RI is the need to enable a new generation of critical infrastructure to support security, stability, and resilience research. As a country we have spent over 30 years with no dedicated research infrastructure to support long-term scientific research on the Internet. Past efforts to study Internet infrastructure vulnerabilities have relied on fragmentary measurement with limited funding, and no ability plan for stable collection, curation, and analysis [11, 12]. Individual researchers devise clever ways to gather data, and with luck find and publish important results describing a moment in time. But individual researchers cannot sustain measurement for decades, or usually even beyond the life of a (typically 2-3 year) grant. This project will provide a vetted and validated infrastructure that supports collection, curation, archiving, and expanded sharing of data needed to advance critical scientific research on the Internet infrastructure.

2.1 Established community need

Although the field of Internet measurement has expanded for three decades, the availability of measurement infrastructure to support research is scant. Without a source of dedicated funding, Internet measurement infrastructures typically do not survive beyond a grant cycle or two [13, 14, 15, 16, 17]. The PlanetLab platform for deploying and testing distributed network services [18] was used by the measurement community for over a decade, despite its policy restrictions against probing the Internet, but shut down in 2020. Subsequent attempts to sustain flexible extensible measurement infrastructures have not gained significant traction [19, 20, 21].

Although not focused on research use of its data, RIPE Atlas is the largest deployed operational active measurement infrastructure available to researchers, with 13K+ vantage points (VPs) as of June 2025 (§3). RIPE (Europe’s Regional IP address Registry (RIR)) achieved this scale in part because they heavily restricted the types of measurements conducted on the VPs to mitigate risk to VP hosting sites. The deployment model makes it challenging to deploy reactive or Internet-scale measurements, as it “generally takes a few minutes to get the result of a measurement” [22]. The primary mission of RIPE Atlas is diagnosis and troubleshooting to support its operator community, not scientific research.

After three decades with no mid-scale research infrastructure dedicated to scientific research on the Internet itself, many U.S. government and research community reports have advocated for U.S. government investment into longitudinal measurement infrastructure:

1. NITRD’s 2019 **Federal Cybersecurity R&D Strategic Plan** [3] emphasized the importance of evidence-based evaluations and measurements in cybersecurity research, and recommended that the Federal Government prioritize basic and long-term cybersecurity research, including the development of sound scientific foundations and formal, reproducible, and quantifiable methods for assessing the efficacy of cybersecurity solutions.
2. NITRD’s 2022 workshop report on **federally supported data repositories** [23] and subsequent update to the national strategic plan on Big Data R&D [24] repeatedly emphasized building trustworthy data ecosystems that ensure integrity, security, and ethical use of scientific data, especially their critical role in AI/ML and other data-intensive research.
3. In 2021, amid continued concerns on the validity of data on Internet broadband availability in the wake of the pandemic, an **NSF-funded workshop** recommended that “**NSF or NTIA should fund data hosting infrastructure** that makes broadband-related data sets available for queries at scale” [25]. The same year the **National Academy** revisited its report on principles and practices for federal statistical agencies [26]. The Internet remains the only critical infrastructure with no federal agency dedicated to its resilience and accountability.

4. Concerns over slow progress on implementation of Internet routing security practices led the **U.S. Federal Communications Commission (FCC)** to issue a February 2022 Notice of Inquiry into potential regulatory interventions that could reduce the severity of the routing security threat to U.S. networks [4]. Virtually all public comments, including those from industry, agreed that the U.S. government should invest more in capabilities to scientifically measure and analyze the global routing system [27]. Subsequently, in 2024, the **U.S. White House (National Cyber Director)** articulated a clear need for Internet routing-focused measurement and monitoring infrastructure to facilitate global Internet security [6, 7].¹
5. For 12 years, CAIDA’s **Active Internet Measurement Systems (AIMS)** workshops [28] have brought researchers, operators, and government stakeholders together to discuss existing Internet measurement system challenges. A perennial topic at these workshops has been the challenge of supporting longitudinal data [29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39].
6. For 11 years, CAIDA’s **Workshop on Internet Economics (WIE)** workshops [40] have hosted discussions of how to overcome data access barriers for economic and policy researchers seeking to study the Internet [41, 42, 43, 44, 45, 46, 47, 48, 49, 50, 51].
7. In 2021, the NSF-sponsored two **Workshop on Overcoming Barriers to Internet Research (WOMBIR)** workshops [52]. Participants heavily emphasized the need for dedicated infrastructure to support longitudinal studies of the Internet. This workshop led NSF to launch a new program to fund Internet Measurement Research [53].
8. In 2023, the NSF-sponsored **Workshop on Emerging Research Opportunities at the Intersection of Statistics and Internet Measurement** concluded: “While there are groups that collect, archive and make data available to the community (e.g., CAIDA), there is a broader need to identify and support longitudinal data collection, archival and distribution. We recommend working toward community consensus around what longitudinal datasets should be valuable to the community and funding/institutional methods to sustain support for collection, archival and distribution of those datasets. [54].” Our Design Project held quarterly meetings and annual workshops to develop such community consensus on which measurement capabilities and data sets to prioritize in the proposed project [55].
9. The **Department of Defense** has consistently funded research programs aimed at enhancing its ability to securely operate over commercial Internet infrastructure [56, 57], yet it lacks a dedicated program to support the foundational infrastructure necessary for scientific data collection and comprehensive topology analytics that underpin this mission.

2.2 Research community benefits

Our project targets NSF’s articulated research infrastructure goals [8], namely: to improve the process of accessing, integrating and transforming data to knowledge and discovery; to enable new usage modes to address multi-disciplinary and cross-domain scientific objectives; to address emerging scientific data challenges such as real-time, streaming data, data discovery and delivery, data fusion, integration and interoperability; and enhancing data accessibility and utility. By integrating advanced computational resources, our infrastructure will support AI-enabled Internet measurement experiments and data analysis. Implementing our proposed infrastructure will enable data-intensive methods for assessing security, stability, and resilience properties of network infrastructure, accelerating the *translational* scientific and engineering advances needed to

¹“Research-funding agencies (e.g., NSF) should continue to fund development of Internet routing-focused measurement, monitoring, and alerting technology to facilitate U.S. and global Internet routing security deployment efforts. Funding should support government entities, academic institutions, and independent subject matter experts equipped to measure progress, develop solutions, and...address the next generation of threats and solutions.” [7]

navigate other current and future Internet-related harms [58].

The resulting data will benefit Internet SSR research in the areas of: security vulnerabilities, including detection and mitigation of BGP and DNS hijacking, as well as fingerprinting malicious IP addresses; discovering and modeling Internet interconnection structure, including mapping and finding choke points in router-level topology, submarine links, cellular, and satellite links; performance and stability dynamics including fault diagnosis, and estimating the effect of damage to specific links or segments; classification of network and path properties; traffic sovereignty and infrastructure geolocation; and validation of new measurement methodologies (§6).

Beyond the security threats that are our primary science drivers, the data will contribute to a broad range of disciplines that now depend on data about the Internet, including network science, socioeconomic studies, international relations, and political science. These topics include questions of how network resilience and broadband service quality varies across the globe, and how traffic to important services is routed across national boundaries.

3 Operational context for proposed research infrastructure

Active Internet measurement is not a zero-risk activity, and access to Internet measurement vantage points typically requires navigating trust relationships among actors involved in deploying, operating, and using the infrastructure. Operators of vantage points (VPs) must balance VP capability against who gets access: the more capable a vantage point, the riskier it is to allow access. The hosting site incurs risk in hosting a VP, and has to trust that the platform operator will use the VP in ways that do not harm the hosting network. The platform operator incurs significant risk when they allow researcher access to the platform. These trust requirements inhibit deployment of active measurement infrastructure, impeding progress in the field of Internet measurement.

Figure 1 illustrates a spectrum of access models for active measurement infrastructure, ordered from least to most restrictive. The least restrictive solutions grant researcher access directly to the VP, either bare-metal, or within a container. The platform operator can restrict access with process and capability limits, but has little other control over what the researcher does, and thus assumes significant risk. A step removed from this approach is VPN-like access, where the VP acts as a simple packet forwarder, allowing a researcher to use the vantage point without providing shell access. These solutions allow researchers to craft specific packet sequences that allow inference based on how the receiver reacts. More restrictive solutions do not allow access to the VPs, or do not allow researchers to construct their own packet sequences. The most restrictive solutions provide raw data, which relies on the platform operator knowing the needs of the measurement community *a priori*, or provide access to a restricted set of tests via an API. The utility of the platform hinges on the usefulness of the data, the provided tests, and responsiveness of the API.

Least Restrictive	Examples
• Shell access to VPs	<i>PlanetLab</i>
• Run code in containers on the VPs	<i>EdgeNet</i>
• Run code to construct packet sequences in sandbox on the VP	<i>Scriptroute</i>
• VPN access to send packets from VPs, logic off-VP	<i>PacketLab</i>
• An integrated active measurement programming environment, logic on-VP, or in infrastructure	
• API to use measurement primitives, logic elsewhere	<i>Atlas, Ark</i>
• Use provided data	<i>Atlas, Ark</i>
Most Restrictive	

Figure 1: *Spectrum of active measurement infrastructures. We propose, implement, and deploy an integrated active measurement programming environment for executing defined measurement primitives on-VP or in infrastructure.*

Several early (now defunct) active measurement platforms such as Skitter [59], Surveyor [13], AMP [14], NIMI [15, 16], and DIMES [17] provided (primarily simple traceroute topology) data for use by the research community. We do not discuss M-Lab because it provides server-side (not client-side) facilities for client-server active measurements [20].

PlanetLab: In 2002, Peterson *et al.* began deploying PlanetLab, a platform for deploying and managing distributed network services [18]. PlanetLab nodes were *customized* Linux-based hardware systems to research and education organizations. The customizations included (1) virtual *slices* isolated from other slices running on the same system, (2) the ability to use socket APIs that typically required root privileges, and (3) management software. The measurement community made extensive use of PlanetLab, although its policy did not permit Internet-wide measurements. At its peak, PlanetLab had systems in ≈ 700 organizations. PlanetLab shut down in 2020.

Scriptroute: Released in 2003, Scriptroute [60] provided (1) a set of distributed VPs, and (2) a sandboxed scripting environment so that unvetted users could use them. Each Scriptroute instance protected the VP hosting site by running user scripts in distinct sandboxes that limited the resources and system capabilities available to each script.

Ark: In 2007, CAIDA began operating the Ark infrastructure to perform comprehensive global topology mapping and to support third-party experiments. As of June 2025, and due to our Design Phase (§4), this infrastructure consists of ≈ 300 VPs in 39 countries and 198 ASes. The infrastructure includes x86 rack-mount systems and Raspberry Pi devices, as well as virtual machines (VMs) and containers. Ark has generated data used by at least 500 publications by non-CAIDA researchers [61]. With effort, researchers could deploy vetted measurement software on the vantage point, but deployment was cumbersome: a mix of operating systems, both vendors and vintages, and a mix of CPU architectures. Software installation and maintenance has been primarily a manual job, which makes it extremely challenging, motivating the current proposal.

RIPE Atlas: operated by RIPE NCC since 2010, Atlas is currently the largest deployed operational active measurement infrastructure available to researchers, with $\approx 13\text{K}$ VPs in $\approx 4\text{K}+$ IPv4 ($\approx 2\text{K}$ IPv6) ASes as of June 2025, representing 5% of routed ASes [62]. Atlas has different types of VPs. Most (8K) VPs are small single-board computers with limited CPU, storage, and memory. Atlas also has more-powerful *anchors* (≈ 800), as well as software VPs ($\approx 4\text{K}$) using the same software as on the single-board computers. Factors in Atlas’ success include: (1) the VPs were cheap to produce, (2) RIPE restricts the types of measurements conducted on the VPs to mitigate risk to volunteers, (3) these primitives provide useful building blocks, (4) volunteers are incentivized to deploy VPs because they gain credits that enable them to conduct measurements from other Atlas VPs, and (5) RIPE subsidizes Atlas through RIR fees. Atlas exposes simple measurement primitives through their web-based API that allows users to conduct ping, traceroute, and selected DNS and NTP queries. Users schedule measurements through the API, and fetch the results when they become available. To accomplish a complex measurement, the user must parse the raw data, and then issue new requests through the API. Deploying reactive measurements is challenging, as it “generally takes a few minutes to get the result of a measurement” and most VPs send 4–12 packets per second [22].

NLNOG RING Launched in 2010 by the Dutch Network Operators Group (NLNOG), the RING is a measurement platform of ≈ 700 VPs (as of 2025) hosted by participating network operators across ≈ 400 ASes, primarily in Europe but with global reach [63]. RING VPs are Linux servers offering programmability (subject to resource limits and community guidelines), which support a range of measurements including ping, traceroute, and DNS probing. The trust model limits access to a smaller community; a user must be an operator at a participating AS, and its capabilities are bounded by the need to submit and propagate scripts across VPs, which can induce timing variability and limit reproducibility.

4 Results of design phase

As described in §1, our Design Phase considered a broad set of security vulnerabilities and consequential harms that arise in the underlying systems of the Internet: addressing, routing, DNS, and transport layer security (TLS). We developed and evaluated new approaches to data collection and analysis in these areas [64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 75, 76], the results of which we codify in a deliverable of the Design Phase. However, the wide range of data types that we identified [10], and often unstable sources of data collection and curation, made clear that transitioning all of these designs to sustained implementation phases would require resources beyond the scope of the MSRI program.

Thus, as we transition to this Implementation Phase, we scaled down our aspirations, concentrating on data that can be collected by direct measurement of the Internet. In our Design Phase, based on feedback from community stakeholders over the course of three years of workshops and online meetings, we developed a solution that lies in the middle of the spectrum (Figure 1) to optimize the tradeoff between capability and security. We designed and prototyped a Python-based integrated programming environment that exposes both a set of distributed VPs, and a set of useful measurement primitives from which to build sophisticated measurement tools. The environment makes it difficult for a researcher to cause harm, intentionally or not, as researchers are restricted to the available measurements. We benefited from significant uptake from interested researchers and partners who used and evaluated our prototype deployment.

During the Design Phase, the Ark measurement infrastructure expanded from ≈ 60 vantage points in early 2023 to an 301 active VPs by June 2025. The expansion included a mix of Raspberry Pis, virtual machines, and containers across multiple continents. Our prototype VP management system (§5.1.1) facilitated deployment of 103 containerized nodes, including multiple fleets: NRP (56), Vultr (32), DREN (8), M-Lab (3), and a few one-off home deployments.

Another lesson of our Design Phase is that many researchers would prefer not to gather global measurement themselves but rely on trustworthy raw data archives and derivative data sets. Therefore, systems for indexing and searching data archives, and systems for processing and curating data in ways that researchers have most requested or performed themselves for single studies, are both integral to our RI. These infrastructure components will accelerate discovery and scientific advances in our understanding of Internet infrastructure.

We developed a tutorial demonstrating how researchers can leverage our newly created Python library [77], which provides programmatic access to the measurement capabilities of Ark VPs [78, 79]. We published sample code that runs swiftly on the Ark platform, showcasing how experimenters can easily integrate these capabilities into their workflow [80].

As the final and most significant evaluation component of this project, we organized a hackathon at our February 2025 workshop where attendees worked directly with the measurement platform and software libraries [81]. This hackathon featured six projects that leveraged the prototype RI and the data it generates. Projects objectives included: measuring reachability of RPKI invalid prefixes (that should be filtered); mapping cross-border content delivery of nationally popular content; geolocation of anycast instances; new techniques for measurement of DNS infrastructure properties; and triggering active measurements based on a background signal, e.g., BGP dynamics. The hackathon outcomes, including several submitted papers and open source code repositories, reinforced our confidence in our direction for the RI to support the scientific research community. We describe these and other use cases in §6.

5 Implementation project agenda

We structure our project in five tasks (Figure 2): deploying a new global measurement platform; deploying data management and accessibility infrastructure; deploy analytics infrastructure that demonstrate the infrastructure’s capabilities; integration of the resulting data with national AI resources; and scaling community engagement with the infrastructure.

The global measurement platform will consist of a network of hardware and software-based measurement vantage points, connected to the Internet from heterogeneous networks and geographic regions, designed to systematically collect active measurement data to transform our ability to infer security, stability, and resilience properties of network infrastructure. Essential to our infrastructure is an integrated active measurement programming environment (§3) that performs two crucial functions: (1) allows a platform operator to specify the measurements that a user can run, and therefore to communicate to the host what their vantage point will do, and (2) provides users with reference implementations of complex measurement functions that act as building blocks to more sophisticated measurements.

There are four design goals for this development environment: (1) easy to use, (2) performant, (3) site-host transparent, and (4) interoperable and extensible.

Easy to Use: We will provide Python interfaces to measurement capabilities present on a collection of remote vantage points. The environment executes the measurements on the VPs, and provides the results as objects. Python is extensively used in the measurement community, both in academia and industry, with a large set of modules available for users to re-use.

Performant: To minimize delay between measurement and result, thus enabling complex measurements, our environment will have an event-driven API, where results return to the researcher’s code as they arrive. We will provide centralized access to the VP controller interface, where code runs as close as possible to the VP controller to minimize delay.

Site-host transparent: The environment will allow platform operators to accurately describe the types of measurements the VPs will do. Our environment will have measurement primitives that let us precisely describe the type of traffic that the site host should expect to see, and communicate risks around each of the available measurement primitives.

Interoperable and Extensible. We will use the open-source software scamper [82] to provide measurement capabilities on VPs, and to schedule and receive measurements on VPs. Scamper has been maintained, improved, and extended for more than 20 years, runs on many different operating systems and architectures, including mobile phones, has few (all optional) external dependencies, can run inside containers, and is available in packaged form. Crucially, scamper is extensible, and provides interfaces to add measurement primitives to those already present.

To incentivize deployment, we will design measurements to provide valuable information back to the hosting AS. Collection at this scale and density will allow us, and the hosting sites themselves, to detect operational anomalies and security threats.




Figure 2: Five tasks in GLIMPSE project.

Distributed measurement data collection infrastructure




Figure 3: GLIMPSE architecture: VPs connect to a central controller. Scripts access primitives on VPs using an integrated active measurement development environment deployed on, or next to, the controller. The platform is designed to be responsive, interoperable, extensible, and easy to use.

5.1 Task 1: Construction of global measurement infrastructure

5.1.1 Distributed measurement data collection infrastructure

This task will consist of constructing a globally distributed network of active measurement vantage points, designed to collect Internet security-related data: DNS and application layer vulnerabilities, topological structure and bottlenecks (single points of failure) including mapping to router and Layer 2 infrastructure, etc. The platform will support heterogeneous deployment configurations, including hardware, virtualized software, and mobile vantage points, enabling broad participation from research and education (R&E) networks, IXPs, cloud services, and individual researchers. It will provide comprehensive, continuous data collection at scale, supporting near-real-time insights into vulnerabilities in global Internet infrastructure and facilitating experimental deployments by vetted researchers.

We will integrate 300 existing Ark VPs from Ark into GLIMPSE. Additionally, we will acquire, configure, and deploy 200 new VPs per year in strategically significant locations. Although we propose to deploy $O(1000)$ VPs in this phase, our platform is designed to support significantly higher VP capacity, allowing (and we intend) for scalable expansion well beyond this initial deployment. In later years of the Implementation Phase we will expand VP deployment to commercial cloud and mobile vantage points. Each vantage point will operate an instance of Scamper, our measurement software library. Scamper contains implementations of measurement primitives: traceroute and ping for simple IP topology and delay measurements, DNS lookups for resolving

names, HTTP(S) to interact with web servers, UDP probes to interact with query-response services such as NTP and SNMP, alias resolution methods for identifying which IP addresses belong to the same router, TBIT [83] to infer properties of a remote TCP stack, and packet capture to selectively record specific packets. We elaborate on this component in §5.2.1.

In parallel, we will develop a robust measurement software pipeline to support comprehensive data collection across vantage points. This includes building new measurement primitives, implementing centralized control interfaces for both researcher-driven and ongoing measurements, as well as coordinating the transfer and storage of results. We will expand the capabilities of the platform to incorporate new measurements in response to evolving researcher needs.

We will create a front-end portal and back-end database to manage the full lifecycle of each VP, including recording and managing VP metadata such as location, technical and administrative contacts, hosting network, and custom parameters. The portal will allow hosts to configure VPs they host, and select measurements they are willing to support. The back-end will include a Postgres database and Python library to manage database access across subsystems. This component will automate the creation of Debian packages to streamline on-VP software management and ensure consistent deployment and updates across all field vantage points.

Certificate Management. We will deploy a new certificate authority (CA) to support authentication of Ark VPs. Our implementation will streamline configuration of the SSH proxy service we use to access remote VPs. Trusting this certificate authority removes the need to install public keys from each VP on the SSH proxy server. This approach improves security by restricting the proxy user to establishing a port-forward, with no ability to log in or run remote processes, regardless of how the server is configured. We will automate all certificate renewal on servers and probes, and automate their initial issuance on container-based VPs with known IP addresses.

Permissioning System. We will develop a system to allow measurements in accordance with hosting site preferences to give the owner of the vantage point sufficient control over what we do with it to prevent its use for experiments outside their comfort zone. The system will connect to the VP management system described earlier to record each hosting site’s measurement preferences. The VP database will store the list of allowed measurement primitives, and generate metadata for the controller and VPs to limit the measurements to those allowed. This robust enforcement mechanism will ensure that these preferences are adhered to across all field deployments.

Statistics dashboard for VP hosts. To incentivize deployment of VPs and facilitate their use, we will develop a statistics dashboard, providing insights and aggregated metrics to VP hosts. The dashboard will display statistics on measurement time ranges, reachable IP addresses and networks, RTT and path length distributions, and other performance metrics.

Expansion to commercial partners, clouds, IXPs, and cellular (Contingency). Our infrastructure will incorporate use of commercial cloud vantage points when it provides sufficient coverage and is cost-effective, e.g., for geolocation measurements, which are low bandwidth (thus low-cost), and provide additional diversity of vantage points where necessary. To manage cost, we can provision these VPs dynamically as needed for specific experiments to manage expenses.

Figure 3 illustrates our architecture; peer-reviewed details of the major components are in [84].

5.2 Task 2: Infrastructure to support data management, access, analysis, discovery

Broadening the community that can effectively use the research infrastructure requires reducing the learning curve for researchers, especially students and early career academics. Task 2 thus focuses on making the RI and datasets it gathers accessible and easy to use. This component will handle ingestion, processing, and curation of collected data, ensuring secure and ethical data




Figure 4: VP coordination classes (left), primitives (middle), and result classes (right) in our integrated active measurement development environment. (Details in [84]).

handling while maintaining data integrity. It includes APIs, metadata tagging, indexing, and data discovery tools, as well as leveraging complementary data sets that facilitate data interpretation.

5.2.1 Active measurement programming environment

In the Design Phase, we designed and prototyped a Python module to allow us to provide native Pythonic interfaces for scamper’s active measurement capabilities. Our module exposes two broad collections of classes. The first collection consists of interfaces for interacting with VPs. The second collection consists of interfaces for interacting with measurement results, which normalize the methods and attributes across different measurement types. (Details in [84]).

With valuable feedback from our Design Phase, we will extend and thoroughly document this user-friendly, Python-based programmatic environment to enable researchers to access and effectively use the measurement capabilities of the VPs. We will (continue to) design this environment for ease of use and will extend it based on researcher input and evolving needs. It will include reference implementations of active measurement primitives that we will deploy on VPs around the world (Task 1). It will include capabilities to schedule, interpret, and store measurements. Additionally, we will integrate this programmatic environment with Jupyter Notebooks (Task 5), providing researchers with a seamless and interactive interface for analysis and experimentation.

5.2.2 Internet measurement query system

In the first two years, we will operationalize a query system designed to search and analyze extensive measurement logs and results via Elasticsearch. We consider this system essential to the RI, a strategic response to regular queries from Ark hosts that contact us when their VP performs a measurement to an address that is incidentally on a threat intelligence feed, and their internal monitoring systems flag the VP as possibly compromised. Investigating such activity and communicating with the VP host that the measurement was benign is sufficiently time-consuming that it is a scale-limiting factor. This query system will provide transparency to hosts and enable them to answer their own questions about what their vantage point is doing. Beyond this operational use, we plan for this system to enable more efficient and insightful exploration of network measurement data to support and inform a wide range of network research questions.




Figure 5: Schematic layout of the current prototype of the RAM portal.

5.2.3 Resource access management

Security-related data about critical infrastructure can be sensitive, and CAIDA has spent decades applying and evolving disclosure control technologies and policies for the data that we collect and steward. For the proposed RI, we will develop a secure and efficient Data Resource Access Management (RAM) Portal to enable authorized users to access data resources, including databases, flat files, and APIs, through a centralized, user-friendly interface. This portal will incorporate role-based access control (RBAC) for precise permission management, robust audit logging for accountability, and seamless integration with existing authentication and authorization systems to ensure security and compliance.

We have designed a system with three main components (Figure 5): Keycloak, a feature-rich open source Identity and access management platform; an OpenID-capable reverse web proxy that offloads authorization and authentication from each supported application; and the Resource Access Management (RAM) service, which provides a user-friendly web interface for users to request access to resources and for admins to vet users and grant access. Keycloak supports integration with other identity providers (e.g., CILogon), so we can leverage existing user databases and allow users to authenticate via external providers. Keycloak also supports OpenID Connect (OIDC), OAuth 2.0, SAML 2.0, modern industry-standard protocols that could support future interoperability with other cyberinfrastructure or data systems.

Upon visiting the Resource Portal, the user will be required to login to Keycloak. In addition to the standard information required by Keycloak (name, email), our Keycloak is configured to require information about the user’s organization. Users can also browse available resources in RAM and request any of them. The RAM request form will ask the user for information about their intended use of the infrastructure or data, and require the user agree to an Acceptable Use Agreement (AUA). Their responses allow us to ensure that they are requesting the most appropriate data for their intended use. (In our experience researchers sometimes do not request the appropriate CAIDA data set for their stated research question, so we redirect them.) This frame-

work also allows us to track usage of data, which provides valuable feedback to guide future measurement campaigns.

5.3 Task 3: Internet security analytics platform and case studies

This component focuses on assessing the value of the research infrastructure by developing and analyzing strategic datasets that reveal vulnerabilities, risks, and security challenges across the global Internet. These datasets will support scientific use cases aimed at evaluating the national security posture of critical infrastructure systems, providing crucial insights for strengthening resilience and addressing emerging threats [85, 86, 87, 88, 89, 90].

5.3.1 Vulnerability that allows spoofed DoS attacks

In 2015 a group of network operators defined a set of operational practices that can prevent several types of addressing and routing abuse [91]. This Mutually Agreed Norms for Routing Security (MANRS) initiative² depends on (unfunded) infrastructure from CAIDA to verify compliance with the requirement that operators do *source address validation* (SAV). Persistent lack of source address validation represents one of many failures of market forces to incentivize best security practices in the Internet ecosystem [92, 93].³ In 2018 Luckie *et al.* found that that MANRS participants were no more likely to properly deploy SAV than others [94]. CAIDA’s Spoofer measurements have generated many scientific publications [95, 96, 97, 98, 99, 100, 101]. But perhaps most notable is that this project exemplifies *translational security research* – technical knowledge converted to measurable improvements in infrastructure security. A common use of this tool has been to help operators diagnose their SAV configurations, a function the private sector has had no incentive to provide.

This infrastructure is operating on aged hardware and has several software components in need of update. We will transition the server software to an OpenStack instance in CAIDA infrastructure, and will create a packaging software framework so that we can easily deploy a new instance of the infrastructure in case of hardware failure. Upgrades to this infrastructure will allow us to continue to publish our spoofing data set, which will help analyze and track IP spoofing attacks, techniques, and related vulnerabilities within network traffic. This dataset serves as a tool for detecting, mitigating, and preventing IP spoofing, which is a common vector of cyberattacks like DDoS, man-in-the-middle attacks, and other malicious activities.

5.3.2 Macroscopic Internet topology analytics pipeline

We will develop a scalable infrastructure to automate the construction of one of our most powerful and scientifically generative datasets [102, 103, 104, 105, 106, 107, 108, 109, 110, 111, 112, 113, 114, 115]: the Macroscopic Internet Topology Data Kit (ITDK). This data kit captures critical information about the global structure of the Internet, including network connections, router ownership, geolocation, router vendors, and other macro-level metrics that facilitate study of interconnection and infrastructure resilience. But current infrastructure constraints limit the coverage, depth, and

²The *Mutually Agreed Norms for Routing Security* (MANRS) [91] initiative includes four practices: (1) Prevent propagation of illegitimate routes from customers or one’s own network; (2) Maintain correct contact information for addresses in public databases. (3) Document intended routing policy in public routing registry. (4) Prevent traffic with spoofed source IP address from leaving one’s network.

³Incentive misalignment of this practice represents a classic negative externality in the Internet infrastructure market: networks that allow spoofing save on their own operational costs, while imposing costs on others (in the form of attacks and attack risk).




Figure 6: Flow of requests from users through Pathfinder, a proposed system to produce unified and detailed annotations that support advanced research and analysis of Internet routing behavior.

accuracy of the ITDK annotations. We will expand coverage of the data kit, and enrich it with security, stability, and resilience (SSR)-related annotations, such as performance indicators, and physical facilities through which paths transit. We will create a user-friendly interface to interact with the data, based on the prototype interface we created during the Design Phase to serve our Internet data science course [116]. Fully implementing our prototype system for automatically building the ITDK, historically an intensive manual task, will enable us to create quarterly updates of this dataset by Year 3.

5.3.3 Path security threat evaluation platform

We will develop a comprehensive real-time capability to perform detailed Internet path analytics, providing inferences that annotate network routes from source to destination. This dataset will include per-hop geolocation, ownership details such as Autonomous System (AS) ownership, and annotations of router manufacturer information where applicable. We will also create a platform to enable users to execute, search, and annotate traceroutes, offering enriched metadata such as inferred organization, country, and router vendor for observed IPs. By integrating data from geolocation services, WHOIS records, BGP paths, and active measurements conducted by widely distributed Ark vantage points, this component will produce unified and detailed annotations that support advanced research and analysis of Internet routing behavior.

5.4 Task 4: Artificial Intelligence (AI) solutions

This task will focus on development and establishment of resources to facilitate AI-based research using the GLIMPSE RI and the gathered data. Specifically, we will use SDSC's Expanse, Voyager, and NRP HPC resources to build and tune AI tools and implement AI-driven solutions to ensure the data's impact is maximized for AI advancements.

5.4.1 Open-source LLM deployment

We will deploy a state-of-the-art open-source large language model (LLM) to support development of the AI tools. We will deploy the model on GPUs dedicated to this project, and integrate it into SDSC’s Nautilus platform, which is part of the National Research Platform (NRP). We will implement access to the LLM via an API and a chat UI. To optimize system performance, we will develop a monitoring environment to track usage and load.

We will also have access to LLMs deployed on NRP. Some of these models were set up as part of a separate effort initiated by members of our team to provide LLM capabilities for SDSC. We will use these open-weights models, which currently include LLama-3.3-70B, LLama3.2-90B, and Gemma-3-27B, for evaluating different models to determine those best suited for the RI uses [117]. For tasks requiring model parameter updates, we will use AI-optimized accelerators on Voyager and GPUs on Expanse. Voyager and Expanse are both HPC resources hosted at SDSC with some nodes dedicated as NAIRR Pilot resources (See Facilities supplementary document).

5.4.2 AI-enabled metadata mapping and validation

A key challenge in Internet infrastructure security research is creating meta-data that allows researchers and operational analysts to map millions of measurements to security-relevant properties such as network ownership, geographic location, interconnection (economic) relationships, business type, and hardware vendors (some of which the U.S. government does not trust.). Traditional methods of metadata extraction cannot address the challenges posed by myriad external natural language sources (such as company websites) relevant to Internet infrastructure security.

We will discover, aggregate, and preprocess structured datasets (e.g., WHOIS, geolocation) and natural language sources (e.g., websites, public reports) to create a data repository of meta-data of unprecedented accuracy, coverage, and AI-readiness of our most popular metadata data sets. Challenges include how to best represent data for various AI/ML use cases, how to correct for noise, managing correlations across data sources [23]. We will use a state-of-the-art open-source LLM (Task 1.4.1) to extract and infer metadata from natural language sources, aligning outputs with existing structured datasets to improve inference of security-relevant infrastructure properties, such as which ASes are owned by the same organization (Figure 7). We will implement prompt engineering, retrieval-augmented generation (RAG), prompt tuning, and/or fine-tuning for this task.

5.4.3 AI-enabled use of research infrastructure

To facilitate utilization of the research infrastructure, we will develop an agentic AI framework with conversational AI capabilities. With this framework, users will be able to easily access data and make optimal use of the RI through simple, natural language queries. The AI agent will autonomously locate, retrieve, and integrate relevant data based on task-specific goals to deliver responses to user queries.

For example, rather than having to know the technical details of Scamper or the RI, a researcher could pose queries such as: “How do I find which authoritative nameserver that a client in a network or country would use for a given zone?” or “How do I find determine which clients are served web content from a Cloudflare cache?” or “How do I find which anycast instances share the same last hops in traceroute paths?” The AI agent will determine which measurements are relevant, write and execute code to generate those measurements (including optimizing efficiency and performance) based on requirements and evolving documentation of the underlying Scamper

sets and capabilities of other infrastructure. These modules will support integration with curriculum across STEM and cybersecurity programs, including for high schools. Our goal is to scale expertise in the U.S. cybersecurity work force, equipping students and researchers with skills to analyze Internet measurement data responsibly, contributing to a sustainable knowledge base within the community. We will integrate data products into the next version of PI Claffy’s Internet Data Science for Cybersecurity course, developed during the Design Phase of this project.

5.5.2 Community-driven RI evaluation and feedback

Our workshops and hackathon during the Design Phase provided opportunities for researchers to experiment with our prototypes as we deployed them, and provided valuable feedback that we integrated into our Design and Prototyping activities. As with most software development infrastructure projects, continual feedback is critical to optimizing investment of resources. We structure this task to maximize this feedback, and we will leverage funding from other sources to increase opportunities for scientific and translational research on the RI as we build it. In particular, we will host annual workshops to facilitate scientific engagement, gain feedback on the RI implementation, sharing of lessons learned, and discuss potential improvements and re-prioritization of tasking. We will also use these engagements as opportunities to showcase use cases and promote uptake of use cases in the national interest, including as capstone projects for data science and engineering courses.

6 Demonstration of use cases for infrastructure

We review previous and ongoing use cases for the infrastructure in the areas of security, stability, resilience (SSR), and threat assessment of Internet infrastructure, and validation of new measurement methods and architectures. We also highlight how the RI can explicitly support reproducibility, a persistent challenge in the field of Internet measurement.

6.1 Security, resilience, and threat assessments

In the last year of the Design Phase, researchers used the prototype RI to analyze: performance and resilience of global anycast critical infrastructure [118], and the uptake of routing security best practices [119] (see Georgia Tech LOC). One of our hackathon projects focused on triggering active measurements in real-time based on a background signal, e.g., BGP dynamics or the Internet Yellow Pages (see IJ LOC) [120], illustrating the value of leveraging other data sources to seed measurements on the RI [81].

Identifying adversarial components of communication paths. Concern over the security of defense-related Internet communications has led to U.S. government (NSF/DOD)-funded research programs [121], and a proposal to pursue discovery of such SSR properties [122, 123, 124]. Our densely distributed global measurement RI will enable this emerging field of analytics to provide situational awareness of threats along communications paths [122, 125],

Mapping cross-border content delivery of nationally popular content. Understanding where content is served from has become important, due to rising concerns about traffic sovereignty, cross-border dependencies, and geopolitical implications of content delivery infrastructure. To investigate content locality, a UCSD PhD student measured paths from 70 Ark vantage points in 28 countries to the top-1000 popular websites for those countries. Analysis of the resulting data confirmed the well-known dominance of U.S.-headquartered Content Delivery Networks (CDNs), but also revealed regional hosting patterns and international infrastructure dependencies [126].

6.2 New measurement methods and architectures

Equally important in the field of Internet science is the ability to test and evaluate new measurement methods and architectures. We have already supported several such endeavors. We supported the PacketLab project (NSF CNS-1903612) to evaluate the feasibility and limitations of a fundamentally new type of active measurement vantage point when operating infrastructure at scale [127]. Columbia University (see LOC) has requested that Ark support its Reverse Traceroute capability [128, 129] as its current use of Measurement Lab is in jeopardy due to Google scaling back its support. More recently, we are supporting researchers at Princeton and Johns Hopkins (see JHU LOC) to deploy Scamper on Android devices, and Raspberry Pis with cellular modems installed. Our goal is to contribute such nodes to the Ark infrastructure.

We collaborated with Internet2 to leverage our software infrastructure to develop a new measurement method that combines BGP and active probing to infer relative route preference policies of ASes. Internet2 used this method to examine routing policies of research and education (R&E) network members. This work led to a new proposed project to develop a security-focused routing observatory and operational support system designed to ensure that routing policies align with the security and integrity goals of the U.S. science ecosystem [130].

In 2025, we have also developed and integrated new measurement capabilities in response to researcher needs for large-scale DNS measurement experiments: specifically to study public DNS resolvers and root DNS server performance and behavior (see BYU LOC). These efforts have already led to new open source code repositories [131], as well as a top-tier conference publication [132] (see U Twente LOC). To support rigorous systematic measurement of layer 2 (e.g., MPLS) topology, we worked with Johns Hopkins to provide Donnet’s MPLS measurement capability [133] in the new platform. This work allowed a recent (May 2025) characterization of MPLS use across the global Internet [134]. Researchers interested in pursuing a global characterization of QUIC servers implemented a QUIC measurement primitive in Scamper [135], which we are now integrating into the main branch.

6.3 Reproducibility of measurements

Angst is increasing over the state of reproducibility in our field; one advantage of a stable measurement infrastructure is its accommodation of reproducibility mindsets. Several groups have used the prototype infrastructure to explicitly target reproducibility challenges. During the Design Phase, we demonstrated the ability of the prototype RI to reproduce methods to fingerprint router vendors and aliases [136, 137], and to perform MPLS tunnel detection [133] (previously mentioned). Another team at our hackathon used the infrastructure to reproduce previous work on the caching behavior of public DNS resolvers [138] (see LSU LOC). This method was originally applied to infer the popularity of rare domains through queries to large public recursive resolvers operated by Google, OpenDNS, Quad9, and Cloudflare. These experiments are difficult to support on existing measurement platforms, as they require distributed, coordinated probing facilities that provide fine-grained control of measurements. More recently, a U.K. researcher requested to use our prototype platform to reproduce her 2018 survey of common Maximum Transmission Unit (MTU) values [139], needed by QUIC protocol implementers who seek to send the largest possible packet sizes to increase throughput [140].

Table 1: Deliverables Gantt chart

		year 1	year 2	year 3	year 4
1 GLIMPSE (Claffy)					
1.1 Distributed measurement data collection infrastructure		Jones			
1.1.1 Node Hardware procurement and deployment	J.				
1.1.2 Node Measurement Software Pipeline	J.				
1.1.3 Node Management Software (Backend)	J.				
1.1.4 Measurement Permissioning System	J.				
1.1.5 Ark Node Management System (Frontend)	J.				
1.1.6 Node statistics dashboard for Ark hosts	J.				
1.1.7 Global Expansion of VPs	J.				
1.1.8 Expand to commercial cloud nodes	J.				
1.1.9 Expansion: Mobile Nodes	J.				
1.2 Infrastructure for Data Management, Discovery, Usability		Huffaker			
1.2.1 Active Measurement Programming Environment	H.				
1.2.2 Internet Measurement Query System	H.				
1.2.3 ASRank	H.				
1.2.4 AS Reachability Analytics	H.				
1.2.5 Resource Access Portal (RAM)	H.				
1.3 Security Analytics Platforms		Luckie			
1.3.1 Spoofing Vulnerability Detection Platform	L.				
1.3.2 Macroscopic Internet Topology Analytics Pipe	L.				
1.3.3 Path Threat Evaluation Platform	L.				
1.4 AI Resources for Security Data Analytics		Nguyen			
1.4.1 Open-source LLM Deployment	N.				
1.4.2 AI-Enabled Metadata Mapping and Validation	N.				
1.4.3 AI-Enabled Use of Research Infrastructure	N.				
1.5 Outreach and Community Engagement		Ma			
1.5.1 Ark Host Onboarding Dashboard	M.				
1.5.2 Legal Framework	M.				
1.5.3 Knowledge Base	M.				
1.5.4 Extended capabilities for infrastructure and	M.				
1.5.5 STEM modules to support cybersecurity and CI	M.				
1.5.6 Community-Driven RI Evaluation and Feedback	M.				

References

- [1] U.S. White House, “2024 Report on the Cybersecurity Posture of the United States,” May 2024. <https://www.whitehouse.gov/wp-content/uploads/2024/05/2024-Report-on-the-Cybersecurity-Posture-of-the-United-States.pdf>.
- [2] S. Krouse, D. Volz, A. Viswanatha, and R. McMillan, “U.S. Wiretap Systems Targeted in China-Linked Hack,” *Wall Street Journal*, Oct 2024. <https://www.wsj.com/tech/cybersecurity/u-s-wiretap-systems-targeted-in-china-linked-hack-327fc63b>.
- [3] National Science and Technology Council, “Federal Cybersecurity Research and Development Strategic Plan: Implementation Roadmap,” December 2019. <https://www.nitrd.gov/pubs/Federal-Cybersecurity-RD-Strategic-Plan-2019.pdf>.
- [4] U.S. Federal Communications Commission, “NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing,” Feb. 2022. <https://www.fcc.gov/ecfs/document/1022806680214/1>.
- [5] U.S. Department of Homeland Security Cybersecurity and Infrastructure Security Agency (CISA), “NOTICE OF INQUIRY. PS Docket No. 22-90. In the Matter of Secure Internet Routing,” Feb. 2022. <https://www.fcc.gov/ecfs/document/1022806680214/1>.
- [6] U.S. White House, “National Cybersecurity Strategy Implementation Plan,” July 2023.

- [7] White House Office of the National Cyber Director, “Roadmap to Enhancing Internet Routing Security,” September 2024.
- [8] “Transforming Science Through Cyberinfrastructure,” Feb. 2019. Draft of NSF’s Blueprint for a National Cyberinfrastructure Ecosystem.
- [9] T. D. Knowles, M. Parashar, L. Parker, E. Gianchandani, D. Braga, M. E. Dean, F.-F. Li, A. Moore, M. L. Norman, F. H. Streitz, E. Tabassi, S. Panchanathan, and A. Prabhakar, “Strengthening and democratizing the u.s. artificial intelligence innovation ecosystem: An implementation plan for a national artificial intelligence research resource,” tech. rep., National Artificial Intelligence Research Resource Task Force, Washington, D.C., Jan 2023. Accessed: 2025-05-24.
- [10] D. Clark and k. claffy, “Toward a Theory of Harms in the Internet Ecosystem,” in *Telecommunications Policy Research Conference (TPRC)*, Sep 2019.
- [11] N. S. Foundation, “Internet Measurement Research: Methodologies, Tools, and Infrastructure (IMR).” <https://new.nsf.gov/funding/opportunities/imr-internet-measurement-research-methodologies-tools>, 2022.
- [12] ESFRI Drafting Group on Research Infrastructures Funding, “ESFRI Report on Funding of Research Infrastructures,” june 2024. <https://zenodo.org/records/14770891>.
- [13] S. Kalidindi and M. J. Zekauskas, “Surveyor: An infrastructure for Internet performance measurements,” in *INET*, (San Jose, CA), June 1999.
- [14] T. McGregor and H.-W. Braun, “Balancing cost and utility in active monitoring: The AMP example,” in *INET*, (Yokohama, Japan), July 2000.
- [15] V. Paxson, J. Mahdavi, A. Adams, and M. Mathis, “An architecture for large-scale Internet measurement,” *IEEE Communications Magazine*, vol. 36, no. 8, pp. 48–54, 1998.
- [16] V. Paxson, A. Adams, and M. Mathis, “Experiences with NIMI,” in *PAM*, (Hamilton, New Zealand), Apr. 2000.
- [17] Y. Shavitt and E. Shir, “DIMES: let the Internet measure itself,” *Computer Communication Review*, vol. 35, no. 5, pp. 71–74, 2005.
- [18] L. Peterson, A. Bavier, M. E. Fiuczynski, and S. Muir, “Experiences building PlanetLab,” in *OSDI*, (Seattle, WA), pp. 351–366, Nov. 2006.
- [19] B. C. Şenel, M. Mouchet, J. Cappos, O. Fourmaux, T. Friedman, and R. McGeer, “EdgeNet: A multi-tenant and multi-provider edge cloud,” in *EdgeSys*, pp. 49–54, Apr. 2021.
- [20] “Measurement Lab.” <http://measurementlab.net>, Jan 2009.
- [21] K. Levchenko, A. Dhamdhere, B. Huffaker, k. claffy, M. Allman, and V. Paxson, “PacketLab: A Universal Measurement Endpoint Interface,” in *ACM Internet Measurement Conference (IMC)*, Nov 2017. <https://packetlab.github.io/>.
- [22] O. Darwich, H. Rimlinger, M. Dreyfus, M. Gouel, and K. Vermeulen, “Replication: Towards a publicly available Internet scale IP geolocation dataset,” in *IMC*, pp. 1–15, Oct. 2023.
- [23] Big Data Interagency Working Group (BD IWG), “Big data: Pioneering the future of federally supported data repositories,” tech. rep., Networking and Information Technology Research and Development (NITRD) Program, February 2022. Workshop Report.
- [24] Networking and Information Technology Research and Development (NITRD) Program, “Innovating the data ecosystem: An update of the federal big data research and development strategic plan,” tech. rep., National Science and Technology Council, Washington, D.C., Nov 2024. Prepared by the NITRD Big Data Interagency Working Group.
- [25] Henning Schulzrinne and Marie-José Montpetit, “NSF Broadband Research 2020 Report,” June 2021. <https://www.cs.columbia.edu/~hgs/papers/2021/NSF-BBRR.pdf>.
- [26] National Academy of Sciences, “Principles and Practices for a Federal Statistical Agency, Edition 7.” <https://www.nap.edu/resource/25885/P&P%207%20Highlights.pdf>.

- [27] Public Safety and Homeland Security Bureau, "Docket 22-90: In the Matter of Secure Internet Routing," 2022.
- [28] Center for Applied Internet Data Analysis, "Active Internet Measurement Systems." <https://www.caida.org/workshops/?workshopserieslisting=AIMS>.
- [29] k. claffy, M. Fomenkov, E. Katz-Bassett, R. Beverly, B. Cox, and M. Luckie, "The Workshop on Active Internet Measurements (AIMS) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 39, Oct 2009.
- [30] kc claffy, E. Aben, J. Augé, R. Beverly, F. Bustamante, B. Donnet, T. Friedman, M. Fomenkov, P. Haga, M. Luckie, and Y. Shavitt, "The 2nd Workshop on Active Internet Measurements (AIMS-2) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 40, Oct. 2010.
- [31] kc claffy, "The 3rd Workshop on Active Internet Measurements (AIMS-3) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 41, July 2011.
- [32] kc claffy, "The 4th Workshop on Active Internet Measurements (AIMS-4) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 42, Jul 2012.
- [33] kc claffy, "The 5th Workshop on Active Internet Measurements (AIMS-5) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 43, Jul 2013.
- [34] kc claffy, "The 6th Workshop on Active Internet Measurements (AIMS-6) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 44, Oct 2014.
- [35] kc claffy, "The 7th Workshop on Active Internet Measurements (AIMS-7) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 46, Jan 2015. http://www.caida.org/publications/papers/2016/aims2015_report/.
- [36] kc claffy, "The 8th Workshop on Active Internet Measurements (AIMS-8) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, Oct 2016. http://www.caida.org/publications/papers/2016/aims2016_report/.
- [37] kc claffy, "The 9th Workshop on Active Internet Measurements (AIMS-9) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, Oct 2017. http://www.caida.org/publications/papers/2017/aims2017_report/.
- [38] kc claffy, "The 10th Workshop on Active Internet Measurements (AIMS-10) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, Oct 2018. http://www.caida.org/publications/papers/2018/aims2018_report/.
- [39] kc claffy, "The 11th Workshop on Active Internet Measurements (AIMS-11) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, Jul 2019. http://www.caida.org/publications/papers/2019/aims2019_report/.
- [40] Center for Applied Internet Data Analysis, "Workshop on Internet Economics." <https://www.caida.org/workshops/?workshopserieslisting=WIE>.
- [41] k. claffy, "Workshop on Internet Economics (WIE2009) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 40, Apr 2010.
- [42] k. claffy, "Workshop on Internet Economics (WIE2011) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 42, pp. 110–114, Apr 2012.
- [43] k. claffy and D. Clark, "Workshop on Internet Economics (WIE2012) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 43, pp. 95–100, Jul 2013.
- [44] k. claffy and D. Clark, "Workshop on Internet Economics (WIE2013) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 44, pp. 116–119, Jul 2014.
- [45] k. claffy and D. Clark, "Workshop on Internet Economics (WIE2014) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, vol. 45, pp. 43–48, Jul 2015.
- [46] k. claffy and D. Clark, "Workshop on Internet Economics (WIE2015) Report," *ACM SIGCOMM Computer Communication Review (CCR)*, Jul 2016.

- [47] k. claffy and D. Clark, “Workshop on Internet Economics (WIE2016) Report,” *ACM SIGCOMM Computer Communication Review (CCR)*, Jul 2017.
- [48] k. claffy and D. Clark, “Workshop on Internet Economics (WIE2017) Report,” *ACM SIGCOMM Computer Communication Review (CCR)*, Jul 2018.
- [49] k. claffy and D. Clark, “Workshop on Internet Economics (WIE2018) Report,” *ACM SIGCOMM Computer Communication Review (CCR)*, Apr 2019.
- [50] k. claffy and D. Clark, “Workshop on Internet Economics (WIE2019) Report,” *ACM SIGCOMM Computer Communication Review (CCR)*, Apr 2020.
- [51] k. claffy and D. Clark, “Workshop on Internet Economics (WIE2020) Report,” *ACM SIGCOMM Computer Communication Review (CCR)*, Apr 2021.
- [52] Center for Applied Internet Data Analysis, “Workshop on Overcoming Barriers to Internet Research.” <https://www.caida.org/workshops/?workshopserieslisting=WOMBIR>.
- [53] N. S. Foundation, “Internet measurement research: Methodologies, tools, and infrastructure (imr),” 2022.
- [54] Paul Barford and Madison Tony Ng, “NSF Workshop on Emerging Research Opportunities at the Intersection of Statistics and Internet Measurement Final Report,” 2023. https://pages.cs.wisc.edu/~pb/imr_workshop_final.pdf.
- [55] Center for Applied Internet Data Analysis, “Global Measurement Infrastructure: workshops,” 2025.
- [56] D. of Defense, “Open programmable secure 5g (ops-5g) initiative,” tech. rep., Office of the Under Secretary of Defense for Research and Engineering, 2023.
- [57] N. S. Foundation and D. of Defense, “NSF Convergence Accelerator 2022 Joint NSF/DOD Phases 1 and 2 for Track G: Securely Operating Through 5G Infrastructure,” tech. rep., National Science Foundation, December 2022.
- [58] D. Abramson and M. Parashar, “Translational research in computer science,” *Computer*, vol. 52, no. 9, pp. 16–23, 2019.
- [59] B. Huffaker, D. Plummer, D. Moore, and k. Claffy, “Topology discovery by active probing,” in *2002 Symposium on applications and the Internet (SAINT 2002)*, (Nara City, Japan), pp. 90–96, Jan. 2002.
- [60] N. Spring, D. Wetherall, and T. Anderson, “Scriptroute: A public Internet measurement facility,” in *4th USITS*, Mar 2003.
- [61] CAIDA, “CAIDA’s Resource Catalog, Topology Data Collection, Papers Listing,” 2024. https://catalog.caida.org/search?query=links=collection:caida_topology_datasets.
- [62] RIPE NCC, “Ripe atlas - coverage.” <https://atlas.ripe.net/coverage/>.
- [63] NLNOG, “Ring.” <https://ring.nlnog.net/>.
- [64] T. Alfroy, T. Holterbach, T. Krenc, and C. Pelsser, “The Next Generation of BGP Data Collection Platforms,” in *ACM SIGCOMM 2024 Conference*, Aug 2024.
- [65] T. Krenc, M. Luckie, A. Marder, and k. claffy, “Coarse-grained Inference of BGP Community Intent,” in *ACM Internet Measurement Conference (IMC)*, Oct 2023.
- [66] B. Du, K. Izhikevich, S. Rao, G. Akiwate, C. Testart, A. Snoeren, and k. claffy, “IRRegularities in the Internet Routing Registry,” in *ACM Internet Measurement Conference (IMC)*, Oct 2023.
- [67] R. Sommesse, G. Akiwate, A. Affinito, M. Müller, M. Jonker, and k. claffy, “DarkDNS: Revisiting the Value of Rapid Zone Update,” in *ACM Internet Measurement Conference (IMC)*, Nov 2024.

- [68] B. Huffaker, R. Fontugne, A. Marder, and k. claffy, "On the Importance of Being an AS: An Approach to Country-Level AS Rankings," in *ACM Internet Measurement Conference (IMC)*, October 2023.
- [69] D. Clark, C. Testart, M. Luckie, and k. claffy, "A path forward: improving Internet routing security by enabling zones of trust," *Journal of Cybersecurity*, vol. 10, December 2024.
- [70] B. Du, C. Testart, R. Fontugne, A. Snoeren, and k. claffy, "Taking the Low Road: How RPKI Invalids Propagate," in *ACM SIGCOMM Poster*, September 2023.
- [71] A. Marder, Z. Zhang, R. Mok, R. Padmanabhan, B. Huffaker, M. Luckie, A. Dainotti, k. claffy, A. Snoeren, and A. Schulman, "Access Denied: Assessing Physical Risks to Internet Access Networks," in *USENIX Security Symposium*, August 2023.
- [72] B. Du, C. Testart, R. Fontugne, G. Akiwate, A. Snoeren, and k. claffy, "Mind Your MANRS: Measuring the MANRS Ecosystem," in *ACM Internet Measurement Conference (IMC)*, October 2022.
- [73] R. Sommesse, M. Jonker, and k. claffy, "Observable KINDNS: Validating DNS Hygiene," in *ACM Internet Measurement Conference (IMC) Poster*, October 2022.
- [74] G. Akiwate, R. Sommesse, M. Jonker, Z. Durumeric, k. claffy, G. Voelker, and S. Savage, "Retroactive Identification of Targeted DNS Infrastructure Hijacking," in *ACM Internet Measurement Conference (IMC)*, October 2022.
- [75] M. Gao, R. Mok, and k. claffy, "A Scalable Network Event Detection Framework for Darknet Traffic," in *ACM Internet Measurement Conference (IMC) Poster*, October 2022.
- [76] L. Oliver, G. Akiwate, M. Luckie, B. Du, and k. claffy, "Stop, DROP, and ROA: Effectiveness of Defenses through the lens of DROP," in *ACM Internet Measurement Conference (IMC)*, October 2022.
- [77] Center for Applied Internet Data Analysis (CAIDA), "Scamper python documentation." <https://www.caida.org/catalog/software/scamper/python/>, 2025. Accessed: 2025-06-01.
- [78] M. Luckie, "Towards a domain specific language for internet active measurement." https://blog.caida.org/best_available_data/2024/01/16/towards-a-domain-specific+-language-for-internet-active-measurement/, Jan 2024. Accessed: 2025-05-31.
- [79] "A Python-based Integrated Active Measurement Programming Environment on Ark." https://catalog.caida.org/presentation/2025_python_based_integrated_active_gmi, February 2025.
- [80] M. Luckie, "Understanding the deployment of public recursive resolvers." https://blog.caida.org/best_available_data/2024/05/06/understanding-the-deployment-of-public-recursive-resolvers/, May 2024. Accessed: 2025-05-26.
- [81] CAIDA, "CAIDA GMI AIMS Hackathon 2025." <https://www.caida.org/workshops/aims/2502/hackathon/>.
- [82] M. Luckie, "Scamper: a scalable and extensible packet prober for active measurement of the Internet," in *ACM SIGCOMM Internet Measurement Conference (IMC)*, 2010.
- [83] A. Medina, M. Allman, and S. Floyd, "Measuring the Evolution of Transport Protocols in the Internet"," *ACM SIGCOMM Computer Communication Review*, vol. 35, Apr. 2005.
- [84] M. Luckie, S. Hariprasad, R. Sommesse, B. Jones, R. Mok, and K. Claffy, "An Integrated Active Measurement Programming Environment," in *Passive and Active Measurement*, 2025.
- [85] GAO Information Technology Cybersecurity Team, "WatchBlog: Following the Federal Dollar," 2019. Authorized by the National Defense Authorization Act signed august 2020.

- [86] U.S. Congressional Subcommittee on Intelligence and Emerging Threats and Capabilities, “H.R. 6395—FY21 National Defense Authorization Bill,” 2020.
- [87] U.S. Federal Communication Commission, “Final Report: U.S. Anti-Bot Code of Conduct (ABCs) for Internet Service Providers (ISPs),” Mar. 2012. <https://transition.fcc.gov/bureaus/pshs/advisory/csric3/CSRIC-III-WG7-Final-Report.pdf>.
- [88] Federal Communications Commission, “Communications Security, Reliability and Interoperability Council,” 2019. <https://www.fcc.gov/about-fcc/advisory-committees/communications-security-reliability-and-interoperability-council-0>.
- [89] F. C. Commission, “FCC Launches Inquiry into Internet Routing Vulnerabilities.” <https://docs.fcc.gov/public/attachments/FCC-22-18A1.pdf>, 2022.
- [90] U.S. Federal Communications Commission, “Reporting on Border Gateway Protocol Risk Mitigation Progress; Secure Internet Routing.” <https://www.govinfo.gov/content/pkg/FR-2024-06-17/pdf/2024-13048.pdf>.
- [91] Internet Society, “Mutually Agreed Norms for Routing Security.” <https://www.manrs.org/>.
- [92] J. Polterock, “Spoofers Surpasses One Million Sessions and Publishes Final Report,” Oct 2020.
- [93] k. claffy, M. Luckie, and J. Polterock, “ASPIRE Project Final Report,” October 2020.
- [94] M. Luckie, R. Beverly, R. Koga, K. Keys, J. Kroll, and k. claffy, “Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet,” in *ACM Computer and Communications Security (CCS)*, Nov 2019.
- [95] L. Müller, M. Luckie, B. Huffaker, k. claffy, and M. Barcellos, “Spoofed traffic inference at IXPs: Challenges, methods and analysis,” *Computer Networks*, vol. 182, Aug 2020.
- [96] R. Beverly, R. Koga, and k. claffy, “Initial Longitudinal Analysis of IP Source Spoofing Capability on the Internet,” *Internet Society*, Jul 2013.
- [97] G. Huz, S. Bauer, k. claffy, and R. Beverly, “Experience in using MTurk for Network Measurement,” in *ACM SIGCOMM Workshop on Crowdsourcing and crowdsharing of Big (Internet) Data (C2B(ID))*, Aug 2015.
- [98] Q. Lone, M. Luckie, M. Korczyński, and M. van Eeten, “Using Loops Observed in Traceroute to Infer the Ability to Spoof,” in *Passive and Active Measurement Conference (PAM)*, Mar 2017.
- [99] Q. Lone, M. Luckie, M. Korczyński, H. Asghari, M. Javed, and M. van Eeten, “Using Crowdsourcing Marketplaces for Network Measurements: The Case of Spoofers,” in *Network Traffic Measurement and Analysis Conference (TMA)*, Jun 2018.
- [100] L. Müller, M. Luckie, B. Huffaker, k. claffy, and M. Barcellos, “Challenges in Inferring Spoofed Traffic at IXPs,” in *ACM SIGCOMM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, Dec 2019.
- [101] M. Luckie, R. Beverly, R. Koga, K. Keys, J. Kroll, and k. claffy, “Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet,” in *ACM Computer and Communications Security (CCS)*, Nov 2019.
- [102] S. A. Jyothi, “Solar superstorms: planning for an Internet apocalypse,” in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference, SIGCOMM ’21*, (New York, NY, USA), p. 692–704, Association for Computing Machinery, 2021.
- [103] M. Luckie, B. Huffaker, and k. claffy, “Learning Regexes to Extract Router Names from Hostnames,” in *ACM Internet Measurement Conference (IMC)*, Oct 2019.
- [104] M. Luckie, A. Marder, M. Fletcher, B. Huffaker, and k. claffy, “Learning to Extract and Use ASNs in Hostnames,” in *ACM Internet Measurement Conference (IMC)*, October 2020.

- [105] M. Luckie, A. Marder, B. Huffaker, and k. claffy, "Learning Regexes to Extract Network Names from Hostnames," in *Asian Internet Engineering Conference (AINTEC)*, December 2021.
- [106] M. Luckie, B. Huffaker, A. Marder, Z. Bischof, M. Fletcher, and k. claffy, "Learning to Extract Geographic Information from Internet Router Hostnames," in *ACM SIGCOMM Conference on emerging Networking EXperiments and Technologies (CoNEXT)*, December 2021.
- [107] F. Hilal and O. Gasser, "Yarrpbox: Detecting middleboxes at internet-scale," *Proc. ACM Netw.*, vol. 1, Jul 2023.
- [108] T. Albakour, O. Gasser, R. Beverly, and G. Smaragdakis, "Third time's not a charm: exploiting snmpv3 for router fingerprinting," in *Proceedings of the 21st ACM Internet Measurement Conference, IMC '21*, (New York, NY, USA), p. 150–164, Association for Computing Machinery, 2021.
- [109] T. Albakour, O. Gasser, R. Beverly, and G. Smaragdakis, "Illuminating router vendor diversity within providers and along network paths," in *Proceedings of the 2023 ACM on Internet Measurement Conference, IMC '23*, (New York, NY, USA), p. 89–103, Association for Computing Machinery, 2023.
- [110] S. Rostami, T. Heinrich, and T. Albakour, "Poster: An investigation into internet-facing router services," in *Proceedings of the 2024 ACM on Internet Measurement Conference, IMC '24*, (New York, NY, USA), p. 775–776, Association for Computing Machinery, 2024.
- [111] K. Thiagarajan, E. Carisimo, and F. E. Bustamante, "Poster: Revealing hidden secrets: Decoding dns ptr records with large language models," in *Proceedings of the ACM SIGCOMM 2024 Conference: Posters and Demos, ACM SIGCOMM Posters and Demos '24*, (New York, NY, USA), p. 10–12, Association for Computing Machinery, 2024.
- [112] E. Carisimo, C. J. Wang, M. Weaver, F. Bustamante, and P. Barford, "A hop away from everywhere: A view of the intercontinental long-haul infrastructure," *ACM SIGMETRICS*, 2024.
- [113] S. A. Jyothi, "Characterizing the role of power grids in internet resilience," 2023.
- [114] S. Lin, Y. Zhou, X. Zhang, T. Arnold, R. Govindan, and X. Yang, "Latency-aware inter-domain routing," 2024.
- [115] A. Ramanathan, R. Sankaran, and S. Abdu Jyothi, "Xaminer: An internet cross-layer resilience analysis tool," *Proc. ACM Meas. Anal. Comput. Syst.*, vol. 8, Feb. 2024.
- [116] "Internet Data Science for Cybersecurity," Jan 2023. <https://cseweb.ucsd.edu/classes/wi23/cse291-e/>.
- [117] National Research Platform, "NRP-Managed LLMs." <https://nrp.ai/documentation/userdocs/ai/llm-managed/>.
- [118] R. Sommesse, L. Bertholdo, G. Akiwate, M. Jonker, R. van Rijswijk-Deij, A. Dainotti, K. Claffy, and A. Sperotto, "Manycast2: Using anycast to measure anycast," in *Proceedings of the ACM Internet Measurement Conference, IMC '20*, (New York, NY, USA), p. 456–463, Association for Computing Machinery, 2020.
- [119] Kumar, Deepak, "Tracing RPKI Invalid Propagation with Internet Yellow Pages and Scamper," tech. rep., Georgia Tech, 2025. https://deepakgouda.github.io/IYP_x_Scamper/.
- [120] Romain Fontugne, 2025. <https://iyp.iijlab.net>.
- [121] National Science Foundation, "NSF Convergence Accelerator 2022 Joint NSF/DOD Phases 1 and 2 for Track G: Securely Operating Through 5G Infrastructure," 2022. <https://new.nsf.gov/funding/opportunities/nsf-convergence-accelerator-2022-joint-nsfdod-phases-1-2>.
- [122] CAIDA, "Automated Verification of Internet Data Paths," 2024. <https://avoid.caida.org>.

- [123] A. Marder, J. Larrea, K. Claffy, E. Kline, and K. Jamieson, "AVOID: Automatic Verification Of Internet Data Paths," *MILCOM*, 2024.
- [124] A. Marder, J. Larrea, K. Claffy, E. Kline, K. Jamieson, B. Huffaker, L. Thurlow, and M. Luckie, "REVEAL: Real-time Evaluation and Verification of External Adversarial Links," *MILCOM*, 2024.
- [125] Stephanie Stamm, "Ukrainians Work Through Blackouts, Internet Outages as Russia Targets Power Grid," *Wall Street Journal*, March 2023. https://www.wsj.com/articles/ukrainians-work-through-blackouts-internet-outages-as-russia-targets-power-grid-reflink=integratedwebview_share.
- [126] H. Hariprasad, M. Luckie, R. Fontugne, M. Tashiro, and k. claffy, "Popular Here, Hosted Where? Mapping the Infrastructure Behind Nationally Popular Content," in *ACM SIGCOMM Internet Measurement Conference (IMC)*, Oct 2025. submitted.
- [127] T.-B. Yan, Z. Zhang, B. Huffaker, R. Mok, k. Claffy, and K. Levchenko, "Marionette measurement: Measurement support under the packetlab model," in *Passive and Active Measurement (C. Testart, R. van Rijswijk-Deij, and B. Stiller, eds.)*, (Cham), pp. 77–105, Springer Nature Switzerland, 2025.
- [128] E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. Anderson, "Reverse traceroute," in *Proceedings of the 7th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, (San Jose, CA, USA), pp. 219–234, USENIX Association, 2010.
- [129] K. Vermeulen, E. Gurmericililer, I. Cunha, D. Choffnes, and E. Katz-Bassett, "Internet scale reverse traceroute," in *Proceedings of the 22nd ACM Internet Measurement Conference, IMC '22*, (New York, NY, USA), p. 694–715, Association for Computing Machinery, 2022.
- [130] k. Claffy, S. Wallace, and M. Luckie, "CICI:TCR: Routing Operations Observational Technology: Building to Ensure Efficacy of Research (ROOTBEER)," 2025. submitted to NSF's Cybersecurity Innovations for CyberInfrastructure CICI program, available upon request.
- [131] Sattler, Patrick and Zirngibl, Johannes and Hilal, Fahad and Gasser, Oliver and Vermeulen, Kevin and Carle, Georg and Jonker, Mattijs, "Ecseptional github repository," 2025.
- [132] P. Sattler, J. Zirngibl, F. Hilal, O. Gasser, K. Vermeulen, G. Carle, and M. Jonker, "Ecseptional dns data: Evaluating nameserver ecs deployments with response-aware scanning," *Proc. ACM CoNEXT*, vol. abs/2412.08478, 2025.
- [133] Y. Vanaubel, J.-R. Luttringer, P. Mérindol, J.-J. Pansiot, and B. Donnet, "TNT, Watch me Explode: A Light in the Dark for Revealing MPLS Tunnels," in *2019 Network Traffic Measurement and Analysis Conference (TMA)*, pp. 65–72, 2019.
- [134] J. Huddleston, M. Luckie, and A. Marder, "Replication: Detecting Invisible MPLS Tunnels Over Internet Paths," 2025.
- [135] Gauder, Nikolas, "QUIC in Scamper Integration Update Meeting," 2025.
- [136] T. Albakour, O. Gasser, R. Beverly, and G. Smaragdakis, "Third time's not a charm: Exploiting SNMPv3 for router fingerprinting," in *IMC*, pp. 150–164, Nov. 2021.
- [137] T. Albakour, O. Gasser, R. Beverly, and G. Smaragdakis, "Illuminating router vendor diversity within providers and along network paths," in *IMC*, pp. 89–103, Oct 2023.
- [138] A. Randall, E. Liu, G. Akiwate, R. Padmanabhan, G. Voelker, S. Savage, and A. Schulman, "Trufflehunter: Cache snooping rare domains at large public dns resolvers," in *ACM Internet Measurement Conference (IMC)*, 10 2020.
- [139] A. Custura, G. Fairhurst, and I. Learmonth, "Exploring usable path mtu in the internet," in *2018 Network Traffic Measurement and Analysis Conference (TMA)*, pp. 1–8, 2018.
- [140] S. Cheshire, J. Touch, and M. Bagnulo, "Packetization layer path mtu discovery for datagram transports." RFC 8899, Sept. 2020.