Project Status Report

Reporting period: 10/01/2024 - 03/31/2025
Project title:
Mid-Scale RI-1 Design Project (M1:DP):
Designing a Global Measurement Infrastructure to Improve Internet Security (GMI3S)
OAC-2131987

Principal Investigator: kc Claffy, Bradley Huffaker (UCSD), David Clark (MIT)
Project Manager: Elena Yulaeva
Lead Institution: CAIDA, UCSD
Other Institutions: NSRC (U Oregon), MIT
Cognizant PO:

1. Summary of project status

Award Duration	Start date: 10/01/2021	Planned close out: **09/30/2025***
Project %-complete	95%	

2. Near-Term Milestones

Below are milestones with the scheduled dates or actual/forecast/revised (A/F/R) dates that are in current reporting period and until the end of the project, and milestones (with past scheduled dates) that are delayed to the next reporting period. (**Completed deliverables have bold font dates**. Red dates have slipped from their original schedule completion dates.) Note that we found the need to revisit some of the milestones as we learned more about other dimensions of the project and as the ecosystem evolved. The date marked with "R" indicates that the WBS element was reopened due to new requirements and was subsequently completed by that day.

WBS	Subsystem	Milestone	Scheduled Date	Actual date (A) /Forecast Date (F)
1.1	**Design Infrastructure for Data Acquisition**			
	1.1.3	**Monitors Requirements documented**	**10/31/2024**	**10/31/2024 (A)**
	1.1.3.7	**DNS monitoring needs compiled**	**10/31/2024**	**10/31/2024 (A)**
	1.1.4	Monitor specification final report	03/31/2025	05/31/2025 (F)
	1.1.4.3	**Post report for public comments**	**01/31/2025**	**01/31/2025 (A)**
	1.1.4.4	Incorporate public comments & publish final document	03/31/2025	05/31/2025 (F)
	1.1.5	**Monitor software Prototyped**	**12/31/2024**	**10/31/2024 (A)**
	1.1.5.5	**DNS data monitoring software prototyped**	**12/31/2024**	**10/31/2024 (A)**
	1.1.7	Evaluation report of Data Acquisition Component published	07/31/2025	07/31/2025 (F)
	1.1.7.3	Evaluation and enhancement of Data Acquisition	05/31/2025	05/31/2025(F)

		Component finished		
	1.1.7.3.1	**BGP data acquisition (GILL) evaluated and enhanced**	**03/31/2025**	**03/31/2025 (A)**
	1.1.7.3.2	Active measurements evaluated and enhanced	05/31/2025	05/31/2025 (F)
	1.1.8	**Virtualization capabilities prototyped**	**1/31/2025**	**01/31/2025(A)**
	1.1.8.3	**Virtualization capabilities documented and evaluated; report published**	**11/30/2024**	**11/30/2024 (A)**
	1.1.8.5	**Scamper containerization completed**	**1/31/2025**	**1/31/2025 (A)**
	1.1.9	Vulnerabilities and Data Needs Report	07/31/2025	07/31/2025 (F)
	1.1.9.9	Newly identified dataset added	04/30/2025	04/30/2025 (F)
	1.1.9.10	Vulnerabilities and Data Needs Report YR4 revision published	07/31/2025	07/31/2025 (F)
1.2	**Design Infrastructure for Data Management**			
	1.2.3	Data and Metadata Standards Final Report	07/31/2025	07/31/2025 (F)
	1.2.3.1	**Develop metadata schema for existing CAIDA datasets**	**11/30/2024**	**11/30/2024(A)**
	1.2.3.4	**Annual revision of data and metadata specifications**	**02/28/2025**	**02/28/2025(F)**
	1.2.3.5	Report on the state-of-the-art metadata generation approaches [1]	07/31/2025	Descoped – moved to EAGER
	1.2.4	Specification of tools for data curation and documentation	07/31/202	07/31/2025 (F)
	1.2.4.3	**Specification of tools for data curation and documentation, report, annually updated**	**03/31/2025**	**03/31/2025 (A)**
	1.2.4.4	Final report on specification of tools for data curation and documentation	07/31/202	07/31/2025 (F)
	1.2.5	Design data and metadata API for access to various datasets	7/31/2025	7/31/2025 (F)
	1.2.5.2	Increase the number of supported data sources, including non-CAIDA datasets	07/31/2025	07/31/2025 (F)
	1.2.5.4	Enhance and document data and metadata APIs, update annually	07/31/2025	07/31/2025 (F)
	1.2.5.4.3	**Year 3(&4) Update of data and metadata APIs**	**03/31/2025**	**03/31/2025 (A)**
	1.2.5.4.4	Final report on data and metadata APIs published	07/31/2025	07/31/2025 (F)
	1.2.6	SDK libraries developed	07/31/2025	07/31/2025 (F)
	1.2.6.3	Libraries evaluated and enhanced; report published	07/31/2025	07/31/2025 (F)
	1.2.7	**Tools for additional data sources integration created**	**03/31/2025**	**03/31/2025 (A)**
	1.2.7.1	**AS path annotations implemented**	**03/31/2025**	**03/31/2025 (A)**
	1.2.7.2	Report on LLM implementation[1]	~~03/31/2025~~	Descoped – moved to EAGER
1.3	**Design Infrastructure for Broad Usability**			
1.3	1.3.1	Data discovery tools prototyped	06/30/2025	06/30/2025 (F)

[1] We descoped the LLM-related tasks (WBS 1.2.3.5, 1.2.7.2) originally outlined in the PEP and transitioned them to a newly funded EAGER project.

	1.3.1.5	Report on integration of automated meta-data/data citation creation into catalog.	06/30/2025	06/30/2025(F)
	1.3.2	**Software for disclosure control developed**	**01/31/2025**	**01/31/2025 (A)**
	1.3.2.7	**Resource Portal prototype deployed (NEW)**	**01/31/2025**	**01/31/2025 (A)**
	1.3.3	**Report on Policy tools**	**03/31/202**	**03/31/2025 (A)**
	1.3.3.8	**New agreements designed and shared**	**01/31/2025**	**01/31/2025 (A)**
	1.3.4	**Case studies on Extensibility**	**10/31/2024**	**10/31/2024 (A)**
	1.3.4.2	**State of Internet report created and shared (1.3.4.2 Conduct meetings and compile data to create community-authored "State of the Internet report")**	**10/31/2024**	**10/31/2024 (A)**
1.4	**Infrastructure for Outreach**			
	1.4.1	**Meetings conducted. Minutes shared**	**02/28/2025**	**02/14/2025(A)**
	1.4.1.1	**Conduct biannual GMI-workshops**	**02/28/2025**	**02/14/2025(A)**
	1.4.3	STEM workforce task completed	02/28/2025	06/30/2025 (F)
	1.4.3.3	Video tutorials on nodes deployment and management created	02/28/2025	06/30/2025(F)
	1.4.4	Quarterly calls conducted; minutes shared	ongoing	**10/31/2024(A)** **01/31/2025(A)** **03/31/2025(A)**
	1.4.5	Project Presentations	ongoing	09/30/2025(F)

3. Executive Summary

This progress report describes the latest advances in CAIDA's NSF-funded MSRI design project, which aims to develop the next generation of Internet measurement infrastructure to support cybersecurity and networking research. Our project focuses on creating innovative platforms and tools for data collection, curation and utilization, particularly targeting data related to the security vulnerabilities within the packet carriage layer of the Internet, which often lead to significant harm. This reporting period has focused on evaluation of the designs via prototyping, proofs-of-concepts, workshops and hackathons. Our spending has dropped as we are in a no-cost extension year with a limited budget. Here we summarize the developments in the six months since the last report.

Active Measurement platform

Our active measurement infrastructure experienced substantial growth, and as of 1 May contains ~300 nodes successfully performing data collection. The expansion included a variety of node types, with the team deploying a mix of Raspberry Pis, virtual machines, and containers across multiple continents. By the end of March, we had deployed 103 containerized monitors, reflecting our shift toward supporting more flexible deployment options (**WBS 1.1.6.5**).

We devoted considerable effort to enhancing the Python module for Scamper with improved documentation and expanded capabilities. We also improved monitoring capabilities, implemented TLS certificates for remote controllers, and improved tracking of node capabilities and restrictions (**WBS 1.1.5.4.1.1**). We released a new version of the active measurement software package integrating these updates (Scamper 20250106) to support the next phase of the project (**WBS 1.1.5.4**).

We designed and began prototyping **Arkmon**—a new portal that allows hosts to manage their Ark monitors (the nodes that they host for the platform). Arkmon is integrated with Dory, CAIDA's current system for monitor management. We migrated Dory's SQLite database to PostgreSQL and developed a shared Python library (arkmon-lib) to manage database access across subsystems. This transition laid the foundation for more consistent and scalable interactions with measurement probe data. (**WBS 1.1.6.5**)

Domain Name System (DNS) Data Platforms

We completed an analysis of DNS monitoring requirements (**Milestone 1.1.3.7**), which resulted in the ACM Internet Measurement Conference presentation and paper: "DarkDNS: Revisiting the Value of Rapid Zone Updates" (https://www.caida.org/catalog/papers/2024_darkdns/darkdns.pdf)

Data Acquisition Component Evaluation

Evaluation of BGP Measurement Design (Milestone 1.1.7.3.1)

Our collaborators at https://www.bgproutes.io/, partially supported by the University of Strasbourg, built on the GILL prototype and other BGP-related tools. Building on the BGP collection system that we designed and prototyped last year (ACM SIGCOMM 2024 Best Paper Award), they are now exploring the possibility of training LLMs to develop a "ChatBGP" tool, specifically designed for monitoring and analyzing Internet routing dynamics based on collected BGP data. We submitted a paper that documents this new system design to ACM SIGCOMM 2025. Their hope is to transition this technology via a startup company they are launching; PI Claffy is serving as an advisor to this effort.

Evaluation of Active Measurement Infrastructure Design (WBS 1.1.7.3.2)

We evaluated our prototype of an Active Measurement Programming Environment and documented the results in a paper accepted to PAM 2025. This environment allows platform operators to (1) define and enforce the types of measurements users are permitted to conduct and communicate those specifications to VP hosts, and (2) provide users with reusable reference implementations of measurement functions to support the creation of more advanced experiments. Through a series of case studies, we demonstrated that this environment significantly reduces the complexity of implementing high-performance measurement experiments, while maintaining precise control over the measurements executed by VPs.

We worked with the Internet2 team to successfully demonstrate how our Active Measurements Infrastructure helps to overcome the limited visibility of routing policies challenges. We submitted a paper documenting this technique to the ACM Internet Measurement Conference.

As the final and most significant evaluation component of this project, we organized and led a successful hackathon at the GMI AIMS workshop (**WBS 1.1.7.3.2**), where attendees worked directly with the measurement platform and software libraries. https://www.caida.org/workshops/aims/2502/hackathon/ This hackathon featured six projects that leveraged this platform and the data it generates. Participants worked on a variety of projects—from ECS scanning techniques and anycast geolocation using traceroute to real-time DNS analysis and integrating measurement data with the Internet Yellow Pages (IYP) knowledge graph. The hackathon was a tremendous amount of work but reinforced our optimism that we are going in the right direction to support the cybersecurity research community.

Virtualization capabilities

We finalized development and publication of Ark Docker images, making them accessible through Docker Hub (**Milestone 1.1.8.3**)

Tools for additional data sources integration (WBS 1.2.7)

Internet Topology Data Kit. As mentioned in the previous report, we have overhauled the software pipeline to support our Internet Topology Data Kit (ITDK) and used it for the first ITDK data set of 2025. We modularized the workflow, replacing hardcoded elements with configurable components to support greater flexibility. We then upgraded the measurement system to use a newer, more efficient way of

handling network connections. We uncovered and addressed issues in IPv6 address classification and optimized the internal coordination framework to improve performance.

Supporting data integration infrastructure NSF/DOD-funded cybersecurity project. We provided infrastructure to support an NSF/DOD-funded Convergence Accelerator project by providing infrastructure for a new platform, designed to deliver cybersecurity insights into global Internet infrastructure. Pathfinder enables users to execute, search, and annotate traceroutes, enriching them with metadata such as inferred organization, country, and router vendor for observed IP addresses (**Milestone 1.2.7**)

Software for disclosure control

We continued enhancing the Resource Access Management (RAM) Portal which handles authentication and data access control (**Milestone 1.3.2.7**). We added new API features such as automatic approval of requests for publicly accessible datasets and automated email notifications to administrators for new restricted data access requests. Additionally, we initiated integration work to add RAM and Keycloak support to all CAIDA publicly available datasets.

Extensibility

We presented "The Age of DDoScovery: An Empirical Comparison of Industry and Academic DDoS Assessments" (https://www.caida.org/catalog/papers/2024_age_of_ddoscovery/age_of_ddoscovery.pdf) paper at the IMC'24 (November 4, 2024) and it was published in the conference proceedings. This paper presents the first comprehensive, longitudinal analysis of direct-path and reflection amplification DDoS attacks—two major threats to Internet infrastructure. Our findings highlight the value of cross-institutional data sharing and offer the strongest empirical case yet for regulatory frameworks that promote such collaboration. (**WBS 1.3.4**)

New data sharing agreements

We signed three new data-sharing agreements and renewed two existing ones. We also drafted a Memorandum of Agreement (MOA) for telescope data sharing with an industry startup and submitted it to UCSD legal counsel for review. (**WBS 1.3.3.8**)

Community engagement

Active measurement infrastructure use by the community

We expanded the use and capabilities of the Ark active measurement infrastructure through new international deployments, tool enhancements, and collaborations with academic and industry partners. Our accomplishments include launching the Local Preference Probe (LPP) study with Internet2, deploying Ark nodes on M-Lab and Vultr platforms, engaging global research networks, and achieving over 50% host opt-in for new measurement primitives.

Collaboration with Open Science Data Federation

We engaged with the OSDF team to begin the process of launching UCSD-NT telescope data on their platform. Our system administrator, Dan Andersen, initiated the upload of encrypted telescope traffic datasets to the OSDF infrastructure

Biannual GMI workshops

We held two GMI-AIMS workshops (GMI-AIMS-4 in Madrid and GMI-AIMS-5 at UCSD) and a 2-day hackathon to advance the design and implementation of global Internet measurement infrastructure. These events gathered researchers and practitioners worldwide to discuss active measurement, BGP visibility, DNS transparency, data accessibility, AI integration, and STEM education (**Milestone 1.4.1.1**)

4. Technical Progress

1.1 Task 1: Design Infrastructure for Data Acquisition

The milestones for the corresponding WBS items have been successfully completed. However, we continued collaborating with other research groups to evaluate prototypes, develop new data sharing methods, explore additional use cases, and sustainability options.

During this reporting period our efforts primarily focused on ensuring the sustainability of this part of the infrastructure. This included actively engaging collaborators to explore and secure external funding (for more details see **Section 1.4.1**)

Traffic: Two-way Passive Traces Monitor

The milestones for the corresponding WBS items were completed, co-funded by "Integrated Library for Advancing Network Data Science - (ILANDS)" (CNS-2120399).

The work on these measurements, data acquisition and sharing during this reporting period was aimed at evaluating this part of the infrastructure to ensure its usefulness and clarity. For details see **WBS 1.1.7.3.4**

BGP Data Monitoring Platform

The milestones for the corresponding WBS items were completed, co-funded by "Integrated Library for Advancing Network Data Science - (ILANDS)"NSF CCRI ILANDS. (CNS-2120399). The milestones were completed, but we continued to iterate on design evaluation and improvements with feedback from the community.

During this reporting period our efforts primarily focused on the outreach tasks as described in **WBS 1.4.1**. Additionally, we utilized the BGP2GO platform prototype (funded by CNS-2120399) to assess how seamlessly our infrastructure components integrate with the authentication framework we developed (see **WBS 1.3.2**). To gain access, users must: (1) Create an account with the CAIDA Services Single Sign-On (SSO) system by providing basic information and completing authentication via Keycloak and (2) Request access to BGP2GO by visiting https://bgp2go.caida.org/

and submitting the request form.

Active Measurement platform

Ark Measurement Infrastructure Growth

During the life of the project, the Ark measurement infrastructure experienced remarkable growth, expanding from 60-70 nodes at the start of 2023 to an impressive 301 active nodes by the end of March. This is a historic high for the project which significantly enhances its global measurement capabilities. The expansion included diverse node types, with the team deploying a mix of Raspberry Pis, virtual machines, and containers across multiple continents. By March, we had deployed 103 containerized monitors, showing a shift toward more flexible deployment options.

Software Development

We developed scripts supporting the installation of Ark nodes.
We released Scamper 20250106 with new features and devoted considerable effort to enhancing the Python module for Scamper with improved documentation and expanded capabilities. He added SVCB support to their measurement toolkit and implemented new HTTP test features including wildcard certificate processing. These improvements supported a successful hackathon at the GMI AIMS workshop (see **WBS 1.1.7.3.2**), where attendees worked directly with the measurement platform and software

libraries.

We also integrated transitive closure techniques for alias resolution into the software infrastructure to support the Internet Topology Data Kit (ITDK) project, enhancing its accuracy. (**WBS 1.1.8.5**)

Ark Infrastructure Improvements

We developed better monitoring and automation systems, implemented TLS certificates for remote controllers, improved tracking of node capabilities and restrictions, and expanded measurement capabilities. An important transition was the shift from NTP to Chrony for time synchronization on most nodes, enhancing the precision of their measurements. We upgraded the certificate management process: updated nginx configurations to support full certificate chains for legacy Ruby clients, migrated certificate templates into the ark-ca package, and generated new-style certificates for spoofer components.
The following diagram shows how the different parts of the Ark server and Ark clients communicate

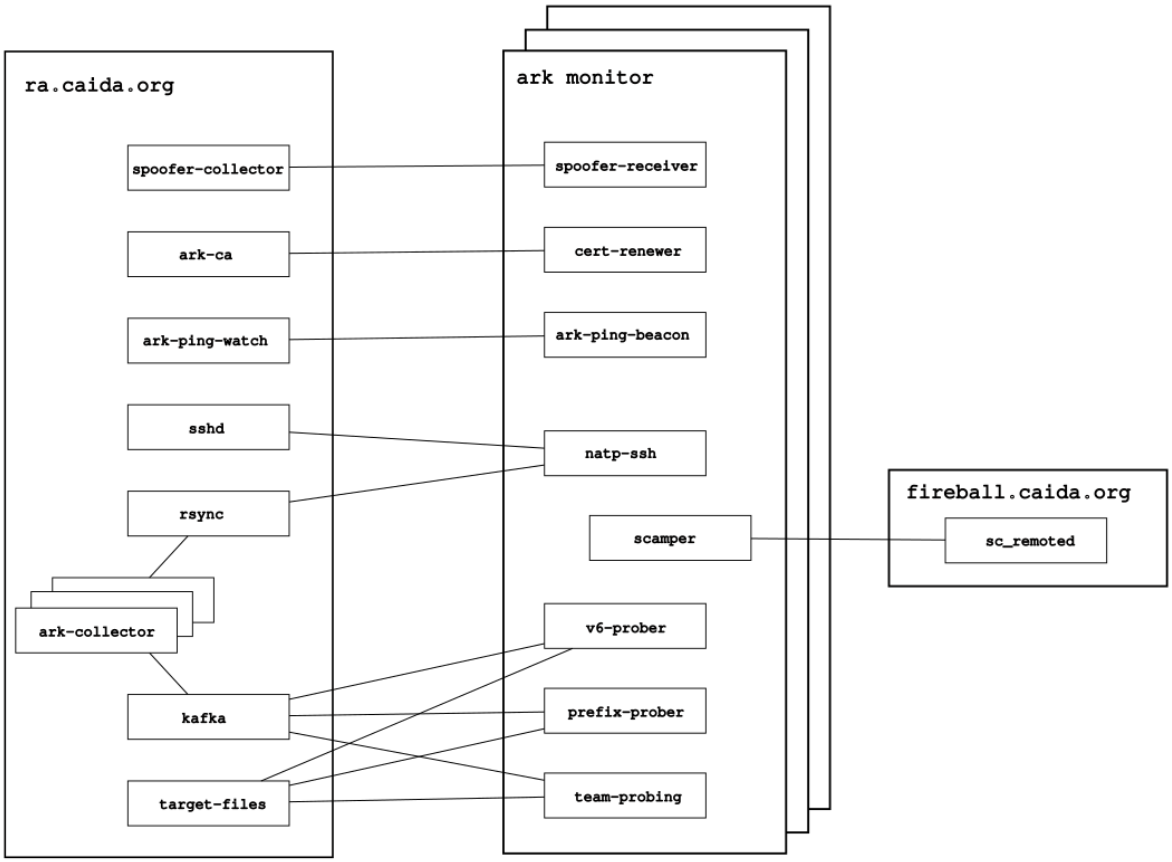

We also worked on spoofer receiver deployment with Twente University, exploring techniques to detect network spoofing capabilities.

Ark Monitors Management System (WBS 1.1.6.5)

We designed and began prototyping Arkmon—a new portal that allows hosts to manage their Ark monitors. Arkmon is integrated with Dory, the current system used by the CAIDA team for monitor management. The schematic below illustrates the new integrated Dory–Arkmon system.

The Dory application has historically maintained the list of Ark monitors, offering a web-based UI and a limited GraphQL API. However, due to insufficient security features, access to Dory was restricted via a

firewall to the CAIDA address space. This limitation prevented external users from viewing or modifying information about their monitors.

In addition, the limited capabilities of the Dory API led approximately half of the scripts interacting with Dory's data to bypass the API entirely and directly access its underlying SQLite database. This direct database access from multiple independent codebases increased the likelihood of bugs and inconsistencies arising from divergent implementations.

To overcome these limitations and modernize the infrastructure, we initiated the design and prototype development of Arkmon with the following primary goals:

- Enable secure external access to monitor data.
- Establish a unified, shared codebase for all database interactions.
- Provide a flexible and user-friendly UI for both hosts and administrators.

We have migrated Dory's SQLite database to PostgreSQL and developed a shared Python library (arkmon-lib) to manage database access across subsystems. This transition has laid the foundation for more consistent and maintainable interactions with monitor data.

We are currently prototyping:

- The Arkmon UI, which will allow hosts to monitor and manage the status of their Ark monitors.
- The Arkmon API, a RESTful service that supports communication between the UI and various Ark data acquisition and processing scripts.

Development steps are documented in CAIDA GitLab.

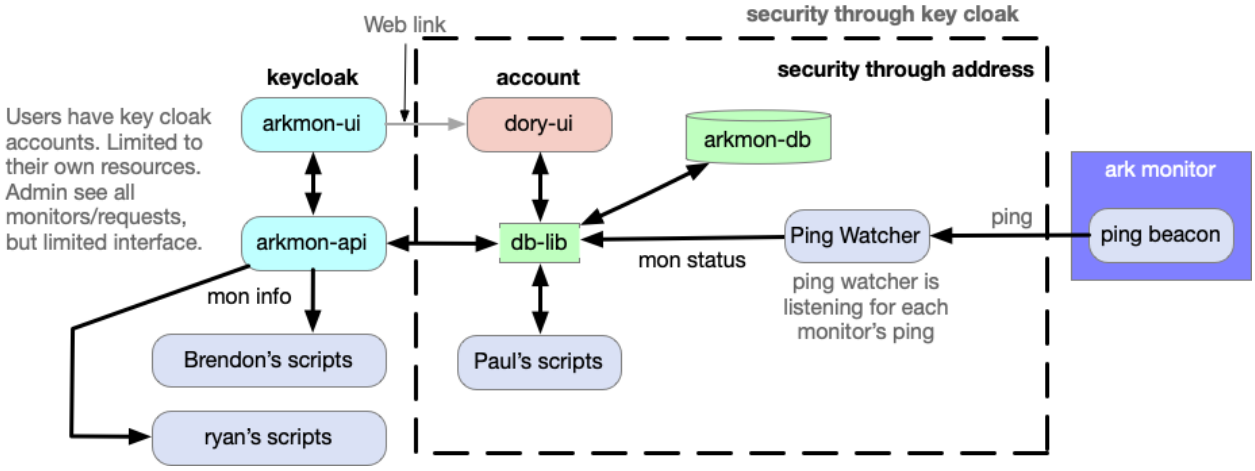

Arkmon-ui dashboard is shown below. Hosts would be able to gain access to their monitors, troubleshoot them if needed, edit monitors details and request new monitor

Monitor details page

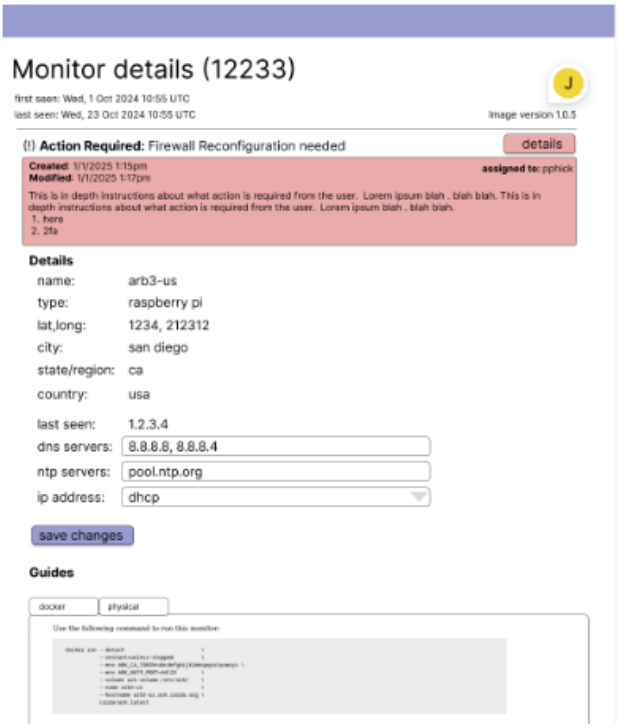

Domain Name System (DNS) Data Platforms

We completed an analysis of DNS monitoring requirements (**Milestone 1.1.3.7**), which resulted in the ACM Internet Measurement Conference presentation and paper DarkDNS: Revisiting the Value of Rapid Zone Update (https://www.caida.org/catalog/papers/2024_darkdns/darkdns.pdf). This work calls for fine-grained, real-time visibility into DNS activities, such as rapid zone updates or leveraging alternative data sources like Certificate Transparency logs. Our approach enhances early detection of transient domains and enables proactive defense against abuse. We emphasize the need to revive rapid zone update mechanisms, supported by frameworks that balance transparency and privacy, allowing trusted parties to mitigate threats effectively.

This paper also discusses several software tools and methodologies for DNS data monitoring that are either suggested or used to address the limitations of traditional approaches that lack ability of estimating the extent of newly registered domains not captured by daily zone snapshots (**Milestone 1.1.5.5**). We also documented a prototype of a pipeline that actively resolves DNS records for newly observed domains. This real-time monitoring enables detection of changes in DNS hosting and name server configurations, capturing domains that may not appear in zone file daily snapshots.

Although our findings highlighted the critical need for a new type of DNS data collection and proposed effective ways to address these requirements, we will not be incorporating this into the Implementation Phase at this time due to budget constraints. However, we remain open to revisiting this in the future and may integrate it into our infrastructure if it is developed by our collaborators and becomes available.

Monitor Specification Report (WBS 1.1.4)

We are still working on this report and are prioritizing it during the next reporting period.

Monitor software Prototyped (WBS 1.1.5)

We completed this WBS (see **Milestone 1.1.5.5** above)

Data Acquisition Component Evaluation (WBS 1.1.7)

BGP Data (WBS 1.1.7.3.1)

We helped our collaborators at https://bgproutes.quest/gill, partially supported by the University of Strasbourg, to build on the GILL prototype to demonstrate the possibility of training LLMs to develop a "ChatBGP" tool, specifically designed for monitoring and analyzing Internet routing dynamics (**Milestone 1.1.7.3.1)**.

Active Measurement (WBS 1.1.7.3.2)

Local Preference Probe (LPP)

We made substantial progress on the LPP demonstration of the use of our software platform to study routing preferences of networks. Working closely with Internet2, we conducted successful experiments demonstrating clear transitions between Commodity and Research & Education network paths. We discussed our results with operators and other researchers, and obtained validation from network operators including AmPath, GWU, SunCorridor, and WiscNet. We presented the results of this work at Internet2's TechEx conference in Boston https://www.caida.org/catalog/media/2024_inferring_relative_route_preference_i2/inferring_relative_route_preference_i2.pdf

Our presentation inspired several other research networks to begin to deploy Ark nodes, which will serve us well going into the Implementation Phase, if funded.

Ark-focused Hackathon

To further evaluate our Active measurement infrastructure, we hosted a two-day Ark-focused hackathon https://www.caida.org/workshops/aims/2502/hackathon/ featuring 6 projects that leveraged Active Measurement software and Ark data. The event was all about testing our infrastructure through creative software development, smart data integration, and innovative analytical approaches. Participants worked on a variety of projects—from ECS scanning techniques and anycast geolocation using traceroute to real-time DNS analysis and integrating measurement data with the Internet Yellow Pages (IYP) knowledge graph. Some teams dove into hands-on tasks like setting up measurement triggers based on background probes, while others focused on fine-tuning data models for the IYP system. There were even inventive explorations into telescope data visualization and analysis, including integrating active measurements to IP addresses observed in telescope traffic. The hackathon was a great opportunity for interdisciplinary collaboration, resulting in new tools, feedback on the infrastructure design, future development ideas, and publications to scientific conferences.

Active Measurement Programming Environment

We worked on a paper (accepted to PAM 2025, https://catalog.caida.org/paper/2024_integrated_active_measurement_programming) that documents and evaluates our prototype of an Active Measurement Programming Environment. This environment enables platform operators to (1) specify the types of measurements that users are permitted to run and communicate with VPs' hosts regarding the specific measurements their vantage points will execute, and (2) provide users with reference implementations of measurement functions that can serve as building blocks for more complex measurement experiments. The paper includes a review of the current technical and usability goals of active measurement infrastructures and presents the prototype's design and deployment within GMI infrastructure.

To illustrate the potential of this environment, we conducted several case studies: (1) identifying the VP with the shortest delay to a given IP address; (2) characterizing the Netflix CDN infrastructure; and (3) reproducing a portion of the Trufflehunter pipeline, which infers the popularity of rare domains by querying large public recursive resolvers operated by Google, OpenDNS, Quad9, and Cloudflare. These use cases demonstrated that our approach significantly lowers the barrier for implementing complex measurement experiments in a high-performance environment while allowing platform operators to precisely define and enforce the types of measurements VPs will perform for site hosts.

Two-way traffic (WBS 1.1.7.3.4)

We surveyed 100 Gb passive trace users, which identified several common applications for these data including:

- Enhancing security monitoring and privacy protection for 5G and beyond cellular networks
- Developing deep generative models (DGMs) for data compression
- Identifying and addressing bottlenecks in high-speed packet analysis systems
- Comparing IPv6 two-way traffic with data from our IPv6 telescope
- Investigating anomaly detection using large language models (LLMs)

These insights will help us tailor our tutorials and further refine analysis infrastructure to meet user needs.

Virtualization capabilities prototype (WBS 1.1.8)

Virtualization capabilities documented and evaluated, report published (Milestone 1.1.8.3)

We finalized development and publication of Ark Docker images, making them accessible through Docker Hub. These container images are based on Debian Bookworm and include Ark-specific packages to emulate physical Ark nodes. They enable the same measurement software and capabilities as physical nodes, supporting seamless reuse of existing workflow configuration process leverages a centralized configuration server and the Dory database for automated setup, eliminating the need for manual customization. This approach allows for quick and secure deployment using IP-based or token-based authentication, with X509 certificates and SSH configuration handled automatically at container startup.

The containers are initialized using a streamlined shell script and managed by the runit service supervisor, ensuring compatibility with Docker and Podman environments. Comprehensive instructions for building, running, and maintaining these images have been documented, including multi-architecture support, network configuration, and certificate handling.

Scamper containerization (Milestone 1.1.8.5)

We prepared a new Scamper Release 20250106, incorporating updates and improvements. Details available here: https://mailman.caida.org/pipermail/scamper-announce/2025-January/000048.html

Updating data needs for securing internet infrastructure report (WBS 1.1.9)

Adding newly identified datasets (WBS 1.1.9.9)

As reported in our M24 report, we indexed IIJ Research Laboratory's Internet Yellow Pages (IYP) into the catalog as an IYP dataset, individual datasets used by IYP, and collection of IYP's datasets. IYP uses a graph database to link a heterogeneous collection of Internet related datasets.

During this reporting period, we developed software capable of resolving ambiguous annotations while preserving all IYP data. Moving forward, we plan to leverage this code to generate informed and accurate annotations for Internet infrastructure researchers.

Vulnerabilities and Data Needs Report (WBS 1.1.9.10)

We are currently analyzing the AIMS workshop presentations and survey responses and will incorporate the findings into the updated version of the report on "Vulnerabilities and Data Needs".

1.2 Design infrastructure for data management

Data and metadata standards (WBS 1.2.3)

As mentioned in our previous 36-month report, our efforts included (a) adding schema annotations to existing datasets and (b) creating a visual representation -- a table within the catalog—to display these annotated datasets. One of the CAIDA REU students is helping us with implementing methods for visualizing the annotated dataset table (**WBS 1.2.4**).

Update data and metadata API (Milestone 1.2.5.4.3)

Our AS Rank UI/API is used by thousands of operators and researchers who often ask for explanation of how AS Rank derived a given ranking (and underlying relationship inferences. In January, our AS Rank

portal experienced an outage due to an unusually high volume of API requests, and it took us a few days to fully restore the service. Although the disruption was unfortunate, it underscored the popularity (and the impact of loss of support) of the infrastructure. We received numerous emails from operators inquiring about the outage, and the issue sparked a discussion thread at NANOG.

Tools for additional data sources integration (WBS 1.2.7)

AS-Rank Refactoring

We began to redesign our AS Rank BGP data analysis infrastructure to be able to show which observed AS paths we used at which steps of the inference process. We documented a new pipeline design and are evaluating how much of this work can fit into the Implementation Proposal **(Milestone 1.2.7.1)**

Internet Topology Data Kit

We have redesigned and reimplemented the analysis software pipeline for the Internet Topology Data Kit (ITDK). This redesign constitutes a transformation from a largely manual and sometimes fragile process to a more automated, reliable, and scalable system that is now easier to maintain.

We started with implementing crucial improvements to the processing pipeline, including bug fixes and updates to vendor data handling. We modularized the workflow, replacing hardcoded elements with configurable components to support greater flexibility. In parallel we automated key IPv6 processes, such as trace selection and DNS resolution. We upgraded the measurement system to use a newer, more efficient way of handling network connections, which required extensive refactoring of several components, including topology inference and latency measurements. Additional issues in IPv6 address classification were uncovered and addressed, improving the integrity of the measurements. We optimized the internal coordination framework to improve runtime performance and monitor utilization in preparation for the next data collection cycle.

Before releasing ITDK-202503, we added checks to catch setup issues early and make the process run more smoothly. During the data collection, we found and fixed bugs with monitor selection and trace filtering. We also improved how vendor probes were handled and made the system more reliable and robust, and the resulting data more accurate.

Pathfinder

We supported an NSF/DOD-funded Convergence Accelerator project with infrastructure to enable a new platform component to provide comprehensive insights into global Internet infrastructure. Pathfinder will enable users to execute, search, and annotate traceroutes, offering enriched metadata such as inferred organization, country, and router vendor for observed IPs.
The platform supports requests via a web interface or API to initiate new traceroutes, annotate existing ones, or query previously collected paths. By integrating data from geolocation services, WHOIS records, BGP paths, and active measurements conducted by Ark monitors, Pathfinder produces unified and detailed annotations that support advanced research and analysis of Internet routing behavior. The diagram below displays the flow of requests from users through Pathfinder. One of our REU students is working on the front end and currently is adding the capability to allow users to manually update the IP geolocation.

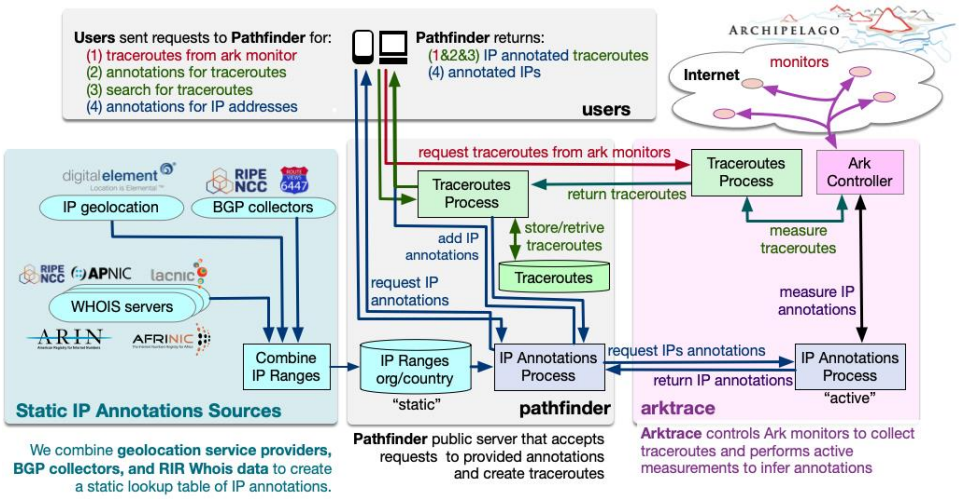

LLM Implementation (WBS 1.2.7.2)
We de-scoped this LLM-related task that we had added to the PEP for a potential supplement, after withdrawing the supplement on NSF guidance. We transitioned this task to a newly funded EAGER project. While the dedicated development of LLM-driven methods will continue under the EAGER effort, some foundational work was carried out as part of this project.

1.3 Design infrastructure for broad usability

Data discovery tools (WBS 1.3.1)
CAIDA data administrator oversaw the REU student-led project focused on improving efficiency and accuracy of cataloging external scientific publications that utilize CAIDA datasets (**WBS 1.3.1.3**). CAIDA REU students, developed ExPub, a Python-based semi-automated tool that extracts metadata from PDFs and publication websites, flags potential publications for curation, and saves metadata for integration with the database and visualization in a Grafana dashboard. The main outcomes of this project include enhanced data quality, new mechanisms to track dataset usage and repeat publications, a human-in-the-loop review pipeline for quality control, and a labeled dataset to train LLMs to infer the utility of datasets and tools from scientific literature. This work was presented at one of the CAIDA GMI meetings.
In a separate effort (**WBS 1.3.1.5**), we supervised two REU students on a project focused on automating the analysis of DNS-related scientific literature to better understand the utility of datasets used in published research. The goals were to classify research papers by DNS-related topics such as malware, denial-of-service attacks, and URL hijacking; extract and analyze how datasets and tools are cited and described; and map each dataset to its role in the research in order to assess its usage and impact. This work resulted in the creation of a growing corpus of manually annotated and LLM-processed DNS papers, along with structured metadata describing the use and citation patterns of various datasets. It also revealed limitations in how current models handle dataset inference across multiple papers, and initiated planning for future LLM-based classifiers that can scale this type of analysis to much larger bodies of scientific literature.

Software for disclosure control (WBS 1.3.2)
We continued working on the Resource Access (authentication) Management (RAM) Portal prototype deployment (**Milestone 1.3.2.7**). In particular, we implemented new portal API functionalities including automatically granting requests for publicly available data access, and automated email notification to admins for new restricted data requests. We also began development work to add RAM and Keycloak support for https://publicdata.caida.org, which hosts the majority of CAIDA's publicly available datasets. We began designing a system to manage access to a large number of resources hosted on a single server. This new functionality is now being integrated into CAIDA's data request processing pipeline.

We engaged two REU students into the RAM portal development. Their work improved the platform's usability, performance, and administrative capabilities by implementing new features, refining the user interface, and ensuring consistency and scalability across the application.

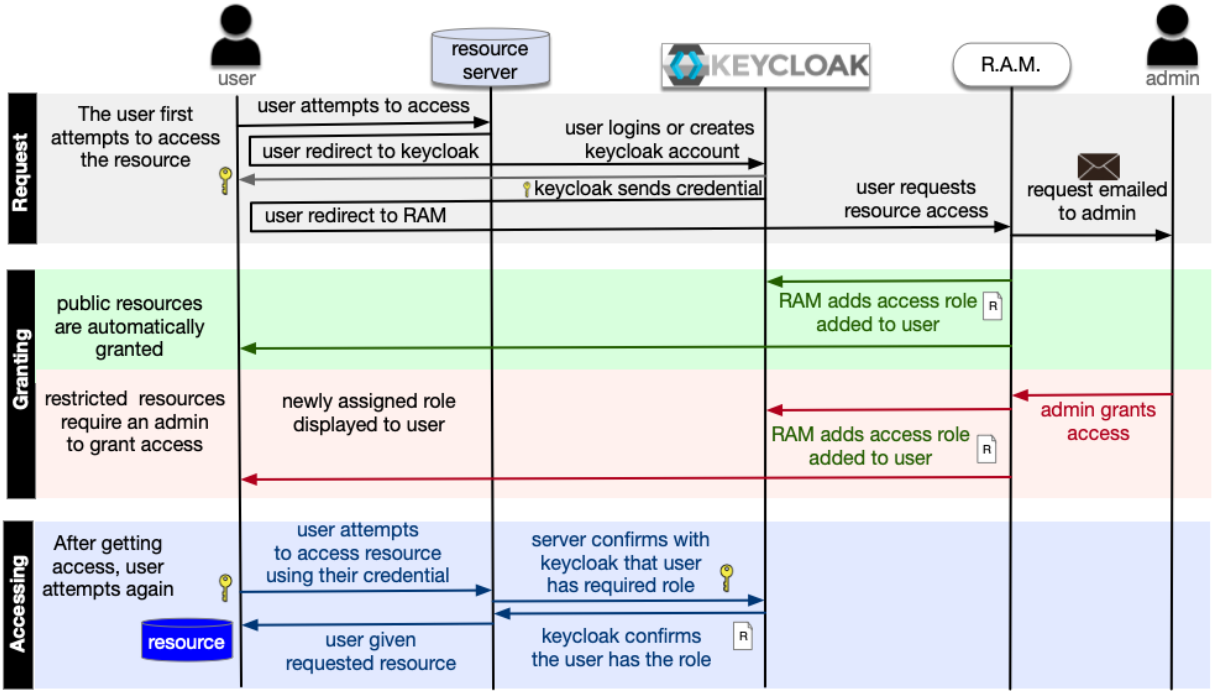

The above figure illustrates the schematic layout of the current prototype of the RAM portal.

Policy Tools (WBS 1.3.3)

New agreements designed and shared (WBS 1.3.3.8)

We drafted and signed a MOA between CAIDA and TU Dresden for telescope data access. This MOA supports the development and maintenance of UCSD-NT Telescope datasets, including tasks such as dataset curation, integrity checks, metadata standardization, long-term accessibility planning, and improving internal data workflows and documentation to ensure reproducibility and usability for external researchers.

We renewed our service agreements with LSU and MITLL which grant the LSU and MITLL teams access to UCSD Network Telescope datasets, with CAIDA providing credentials and maintaining access in line with industry standards and CAIDA Acceptable Use policies. We (UCSD) also initiated MOA negotiations with a new startup for access to telescope data for threat intelligence.

We granted one commercial player a non-exclusive license to use our raw and derived Ark data. The revenue generated supports our ongoing Ark data collection.

Addressing Legal Barriers to Data Sharing

Together with the UCSD OCGA, we navigated a complex process to gain access to the FCC's "Broadband Serviceable Locations" data, licensed through the third-party company CostQuest. This process required significant coordination, including OCGA seeking concurrence and support from UC Legal and Risk Management to file a DA2239 exception request. The primary challenge arose because the FCC's Data Use Agreement (DUA) includes stringent licensing terms, including significant control over how the data is used, and a non-negotiable indemnification clause. At least two other UC campuses (UC Santa Barbara and UC Berkeley) faced similar challenges with this agreement. Ultimately UCSD agreed

to make an exception to their policy and to accept this DUA, provided that: only UC employees use the data under this agreement; no corrections will be sent to the FCC; and UC's activities will not trigger the indemnification clause.

Extensibility case studies (WBS 1.3.4)

We presented "The Age of DDoScovery: An Empirical Comparison of Industry and Academic DDoS Assessments" (https://www.caida.org/catalog/papers/2024_age_of_ddoscovery/age_of_ddoscovery.pdf) paper at the IMC'24 (November 4, 2024) and it was published in the conference proceedings. In this paper we provided the first comprehensive, longitudinal view of direct path and reflection amplification attacks, two dominant classes of DDoS attacks that threaten Internet infrastructure. Our results showcase the collaborative efforts between academia and industry in sharing data across institutional boundaries. We incorporated all available datasets, encompassing far more attacks than any previous study. Our approach to synthesizing these datasets advanced beyond prior efforts in five key areas: the number of macroscopic datasets (ten), the extended observation window (4.5 years compared to months or even three years in earlier works), the inclusion of multiple types of DDoS attacks, the management of large data volumes and usage restrictions through substantial cross-institutional cooperation in processing and normalizing data, and the introduction of a new method to facilitate data sharing by industry.

Beyond synthesizing datasets, our work reinforced the finding that previous studies suffer from significant visibility limitations, providing the strongest empirical foundation to date for regulatory frameworks that encourage data sharing. Moving forward, we aim to motivate others who possess Internet measurement data to contribute to our Research Infrastructure, helping to establish and maintain a view of the Internet transport layer following our established blueprint. (**Milestone 1.3.4.2**)

1.4 Outreach and engagement.

Conduct various meetings with industry and academia (WBS 1.4.1)
We continued our monthly call with DOD and DOD-related stakeholders interested in progress on CAIDA's activities including DREN, Lincoln Labs, ISI.

<u>One-way traffic measurement infrastructure sustainability and partnerships</u>
We conducted meetings with LSU team about the future of the UCSD-NT telescope instrumentation, sustainability models, and help with their IOT dashboard. They presented the current state of the dashboard at the GMI AIMS (Feb 13, 2025) workshop.

We collaborated with TU Dresden team on their paper "Lessons learned from operating a large network telescope" with recommendations for design and sustainability of Internet measurement infrastructure. We submitted this paper to the ACM SIGCOMM 2025 conference (Experience Track) with the following abstract:

> *Network telescopes (aka darknets) collect unsolicited Internet traffic (aka Internet background radiation or IBR), which includes benign and malicious scanning as well as artifacts of spoofed denial-of-service attacks and misconfigured software and networks. Analysis of this traffic has revealed macroscopic insights into security-related events and global network dynamics such as outages. Operating a large-scale network telescope is challenging but often taken for granted, unlike measurement infrastructures in physics. We offer the first study documenting our experiences operating the UCSD Network Telescope, the largest and longest-operating network telescope supporting scientific research. We provide background on the*

history of the telescope and focus on increasing operational challenges as the underlying network evolves. We develop and apply techniques to leverage third party scanning activity to validate the integrity of the data, and to discover misconfigurations in the instrumentation. These insights are crucial for understanding measurement results, which we illustrate using concrete examples. We discuss how our findings generalize to support the expanding ecosystem of other passive techniques, such as honeypots, to track security phenomena.

Active measurement infrastructure use by the community

We designed and executed the Local Preference Probe (LPP)—an active measurement study conducted in collaboration with Internet2—to analyze how BGP route manipulation, specifically AS path prepending, influences routing decisions. Building on his earlier LPP codebase, we developed a new LPP-PRESCAN tool that uses target seeds from the ISI history dataset and the Censys API instead of performing its own ICMP scanning. The updated system processes pre-scan results and probes the same set of targets across nine rounds of varying prepend configurations. In a practice run, the full experiment took approximately 90 minutes, including time for BGP convergence. Results showed clear transitions from commodity to R&E routing paths depending on the prepending strategy and indicated that most R&E networks assign higher local preference to R&E paths than to commodity routes.

We engaged with Case Western Reserve University for geolocation research, and established partnerships with international networks including RedCLARA, RNP, and RIPE NCC. These relationships facilitated the global expansion of the Ark infrastructure, with new nodes appearing in Ghana, Romania, Kosovo, Brazil, Dubai, and many other locations.

We held multiple meetings with the Measurement Lab (M-Lab) team to explore the deployment of Ark nodes within their infrastructure. After thoroughly reviewing M-Lab's policies to ensure full compliance, we successfully initiated the deployment of three Ark containers on M-Lab's sandbox nodes. To validate the setup, we conducted extensive testing using remote-controlled scamper, confirming that all relevant scamper modules were functioning as expected.

We met with Columbia University's reverse traceroute team about integrating Ark VPs and decided to start with a few datacenter nodes due to the 300pps burst traffic requirements.

We collaborated with the University of Twente team to explore the use of Spoofer clients for sending spoofed packets to Vultr-based Ark VMs as part of comparative testing. We configured the Vultr nodes to run both the Spoofer receiver and Ark software, resolving several technical issues in the process. Following this setup, 32 Vultr containers were provisioned for Ark deployment

We improved tools and documentation including refactoring and expanding Scamper's Python module, addressing feature requests and issues raised by measurement researchers.

We reached out to hosts with details about the new capabilities, and to seek their permission to deploy new measurement primitives (DNS, HTTP, UDP probes, spoofing) in addition to the current Internet topology focused measurement primitives (ping, traceroute, alias resolution) that Ark currently supports as part of its macroscopic topology measurement project. To date, more than 50% of hosts have opted in.

BGP measurement infrastructure

On February 11, 2025, we dedicated nearly a full day of the GMI AIMS workshop to presentations and discussions focused on BGP-related data acquisition and analysis. Highlights included:

- Mingwei Zhang, a former CAIDA postdoc now at Cloudflare, presented the latest BGP analysis approaches adopted by Cloudflare.
- Ethan Katz-Bassett from Columbia University shared updates on the PEERINGS testbed. His team is transforming PEERINGS into a comprehensive platform that enables researchers to

conduct Internet routing experiments in environments that closely mimic real-world cloud providers. They have deployed PEERING routers at 30 cloud sites across six continents, established nearly 10,000 BGP peers, and integrated hundreds of Cloudflare CDN sites, significantly expanding the testbed's reach.

- RIPE NCC presented an analysis of the RIPE Atlas anchor mesh following two cable cuts in the Baltic Sea on November 17 and 18, 2024. Their data indicated that only 20%-30% of measured paths experienced latency increases, with no significant evidence of increased packet loss.
- Thomas Krenc (CAIDA/UCSD) conducted a hands-on tutorial session using BGP2GO.
- Hackathon authors presented their project on BGP and active measurements, titled "Trigger real-time (Ark) measurements based on background (Atlas) probes (Traceroute Trigger)."

These discussions and presentations underscored the evolving approaches and challenges in BGP data analysis and measurement, reflecting the active collaboration between industry and academia in this field.

Open Science Data Federation

We had several meetings with the Open Science Data Federation sub-team who will help us launch some UCSD telescope data (a large data set) on the OSDF platform. The sub-team supporting us is led by Christina Koch, Lead Research Computing Facilitator at the University of Wisconsin). This initiative was prompted by the U. Dresden interest in getting the data "cloned" in Europe. Although there is no OSDF cache in Dresden, we are hoping to motivate the German R&E research network infrastructure community to join the OSDF. We are starting with a private namespace of 200TB of (the most recent packet capture) data, provided by OSDF on a trial basis. Access will be multi-user authenticated read. Our first meeting was in January 2025, and in February they agreed to use Missouri for data cache storage to start. We are still waiting for them to send us instructions to generate a token and upload the data; we expect this to advance in April 2025. To manage the privacy concern (sensitive data, users must sign an AUP with us to access), we will encrypt the packet capture data with the same methodology we use to store our archives at DOE's NERSC (OpenSSL with AES encryption).

Conduct biannual GMI-workshops (Milestone 1.4.1.1)

GMI-AIMS-4 workshop

We held the fourth GMI-AIMS workshop (GMI-AIMS-4) on November 7, 2024, right after IMC in Madrid. This meeting brought together over 25 researchers and practitioners from institutions across Europe, North America, and Asia to advance the Global Measurement Infrastructure to Improve Internet Security (GMI3S). We are moving from the design phase to proposing implementation, with an emphasis on active measurement, BGP visibility, DNS transparency, and network telescope infrastructure. As always, the agenda was packed—and fluid—reflecting the dynamic and interdisciplinary nature of this work.

The talks covered everything from DNS blocklists and IPv6 scanning bugs to new tools for anycast and BGP route collection. We discussed how we can make measurement data more usable and shareable, addressing challenges like long-term sustainability, and ensuring that our tools serve not just researchers, but also network operators and policymakers. These workshops continue to be one of the most productive ways we've found to align measurement needs across different parts of the ecosystem and to keep this community moving forward together.

GMI-AIMS-5 workshop

We conducted the GMI-AIMS-5 Workshop, which took place in person at UCSD from February 10 to February 14. Additionally, we led a Hackathon on the weekend preceding the workshop.

The main goals of the workshops are as follows

- Integrate the design phase outcomes with the proposed implementation plan, ensuring alignment with the project's goals and objectives
- Discuss the deployment of a new global measurement platform, data management and accessibility infrastructure, analytics infrastructure, and community engagement/training

workshops
- Explore the integration of the GMI3S infrastructure with national AI resources, and the use of Large Language Models (LLMs) for data annotation and metadata generation
- Finalize the plans for community engagement, workforce training, and STEM education, including the development of a Network Infrastructure Data Science course and video tutorials

The workshop was attended by 82 participants.
Adjacent hackathon focused on active measurement challenges:
https://www.caida.org/workshops/aims/2502/hackathon/.) Several participants created their own project pages to showcase their work (e.g. https://deepakgouda.github.io/IYP_x_Scamper/)
One of the hackathon projects even led to a paper that was accepted to CONEXT.
Web page summarizing hackathon at and summarizing project:
https://www.caida.org/workshops/aims/2502/hackathon/

Meetings with Academic and Industry Stakeholders (WBS 1.4.4)

We continued connecting with academic and industry collaborators to explore ways to support ongoing research and discuss plans for the next phase of the project. We met with Ioana Livadariu from Simula Research (Norway), to discuss how to best support her submarine topology research. Dr Claffy met with Melissa Dark at
https://teachcyber.org/cybersecurity-teaching-resources/curriculum-guidelines/ to discuss the possibility of integrating CAIDA's Internet data science into their classrooms.
We also reached out to several collaborators to discuss letters for the Implementation phase.

Make presentations on various conferences and workshops, submit publications (WBS 1.4.5)

We presented 4 papers that evaluate GMI3S infrastructure components at the IMC'24 (November 4, 2024)
https://www.caida.org/catalog/papers/2024_age_of_ddoscovery/age_of_ddoscovery.pdf
https://www.caida.org/catalog/papers/2024_darkdns/darkdns.pdf
https://www.caida.org/catalog/papers/2024_darksim/darksim.pdf
https://www.caida.org/catalog/papers/2024_sublet_your_subnet/sublet_your_subnet.pdf

Dr Claffy presented GMI3S work at the NITRD October 2024 meeting
https://www.caida.org/catalog/media/2024_designing_global_measurement_infrastructure_nitrd/designing_global_measurement_infrastructure_nitrd.pdf

Dr. Claffy gave a keynote talk to RITE.RO (https://rite.org.ro/en/) meeting in Romania. They have a strong interest in deploying ark mesh there to help them understand interconnection

Matthew Luckie presented work using GMI3S-funded software infrastructure at Internet2 TechEx in Boston, on Dec 11, 2024
https://catalog.caida.org/presentation/2024_inferring_relative_route_preference_i2
We submitted a paper abstract to TPRC53 proposing to give a quick introduction to how Internet measurement tests can be run using the Ark platform, in the hope of persuading a few users from the TPRC community to experiment for themselves.

1.5 Project management

Project Support (WBS 1.5.1)

Periodically review and update PEP (WBS 1.5.1.2)

We descoped the LLM-related tasks (WBS 1.2.3.5, 1.2.7.2) originally outlined in the PEP and transitioned them to a newly funded EAGER project. While the dedicated development of LLM-driven methods will continue under the EAGER effort, the foundational work was carried out as part of this project.

Project Controls (WBS 1.5.2)
We continued close monitoring of project scope, cost, and outcomes quality.

Internal meetings (WBS 1.5.2.1)
We continued our weekly internal project management meetings.

Meetings with subcontractors (WBS 1.5.2.2)
We continued CAIDA/MIT technical meetings

Quality Management (WBS 1.5.3)
As part of the project's quality management efforts, the AIMS workshop provided a valuable opportunity for collaborators to offer direct feedback on infrastructure components, tools, and workflows. The hackathon (https://www.caida.org/workshops/aims/2502/hackathon/) during the workshop served as a real-world evaluation of the infrastructure's usability, reliability, and performance. Participants deployed and tested components in live environments, suggesting improvements, and validating recent enhancements. This collaborative and hands-on approach helped ensure that the infrastructure meets the practical needs of the community and aligns with the project's long-term goals.

List of Acronyms

AIMS	Active Internet Measurements
AMPRNet	AMateur Packet Radio Network (Network 44)
API	Application Programming Interface
Ark	CAIDA Archipelago Measurement Infrastructure
AS	Autonomous System
AUA	Acceptable Use Agreement
AUP	Acceptable Use Policy
AY	At Year
BCP	Baseline Change Proposal
BGP	Border Gateway Protocol
BMP	BGP Monitoring Protocol
CAIDA	The Center for Applied Internet Data Analysis
CDN	Content Delivery Network
CDR	Conceptual Design Report
CENIC	California Educational & Research Network Information Center

CFR	Code of Federal Regulations
CI	Cyber Infrastructure
CISE	Computer and Information Science Engineering
CNS	Computer and Networked Systems
CMP	Configuration Management Plan
CPI	Cost Performance Index
CSAIL	Computer Science & Artificial Intelligence Laboratory
CY	Calendar Year
DDOS	Distributed Denial of Service
DHS	Department of Homeland Security
DNS	Domain Name System
DPM	Deputy PM
DPU	Data Processing Unit
DREN	Defence Research and Engineering Network
DZDB	DNS Zone Database
EAC	Estimate at Completion
eMMC	embedded Multi Media Card
ES&H	Environment, Safety and Health
EVMS	Earned Value Management System
FCC	Federal Communication Commisions
FFRDC	Federally Funded Research and Development Center
FTE	Full Time Equivalent Employee
FY	Fiscal Year
GE	Gigabit Ethernet
GMI3S	Global Measurement Infrastructure to Improve Internet Security
ICANN	Internet Corporation for Assigned Names and Numbers

IHR	Internet Health Report
IIJ	Internet Initiative Japan
IMC	Internet Measurement Conference
IOT	Internet of Things
ISP	Internet Service Provider
IYP	Internet Yellow Pages (IIJ)
KINDNS	Knowledge-Sharing and Instantiating Norms for DNS and Naming Security
L2	Level 2
LFO	Large Facilities Office
LoC	Letter of Collaboration
MANRS	Mutually Agreed Norms for Routing Security
MIT	Massachusetts Institute of Technology
MREFC	Major Research Equipment and Facilities Construction
NIC	Network Interface Card
NIDS	Network Infrastructure Data Science
NIS	Network and Information Security (EU)
NIST	National Institute of Standards and Technology
NITRD	Network and Information Technology Research and Development
NLP	Natural Language Processing
NOG	Internet Network Operators' Group
NSF	National Science Foundation
NSRC	Network Startup Resource Center
NTIA	National Telecommunications and Information Administration
OAC	Office of Advance Cyberinfrastructure
OMB	Office of Management and Budget
PI	Principal Investigator

PII	Personally Identifiable Information
PEP	Project Execution Plan
PM	Project Manager
PMCS	Project Management Control System
PO	Program Officer 5
PRP/NRP	Pacific Research Platform/National Research Platform

RIPE-NCC	Regional Internet Registry for Europe Network Coordination Centre
RIS	Routing Information System
ROA	Routing Origin Authorization
RPKI	Resource Public Key Infrastructure
RV	Route Views
SDSC	San Diego Supercomputer Center
SOW	Statement of Work
SPI	Schedule Performance Index
SSAC	Security and Stability Advisory Committee
TLD	Top-level domain
TPC	Total Project Cost
UO	University of Oregon
USC ISI	USC Information Science Institute
VP	Vantage Point
WBS	Work Breakdown Structure