Project Status Report

Reporting period: 10/01/2021 - 03/31/2022
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), and David Clark (MIT)
Project Manager: Elena Yulaeva
Lead Institution: CAIDA, UCSD
Other Institutions: NSRC (U Oregon)
Cognizant PO:

Table of Contents

1. Summary of project status

A brief summary of project's overall status on technical progress, cost and schedule performance.

Award Duration	Start date: 10/01/2021	Planned close out: 09/30/2024
Project Finish Date	Planned Early Finish:	Estimated Early Finish:
Project Cost	Total project cost: 7,865,527	Estimate-to-Completion: 7,068,707

Cost Contingency	Budgeted contingency: 375,000	Remaining contingency: 375,000
Project %-complete 11%		

2. Near-Term Milestones

Include milestones with the scheduled dates or actual/forecast dates that are in current and the next reporting period, and milestones (with past scheduled dates) that are delayed to future reporting period. (**Completed deliverables have bold font dates**.)

WBS	Subsystem	Milestone	Scheduled Date	Actual date (A) /Forecast Date (F)
1.1	**1.1.1**	Preliminary report on Internet infrastructure security vulnerabilities that the GMI project will gather/manage/share data to address	03/31/2022	**03/31/2022(A)**
	1.1.2.1	Slide deck with detailed information about existing CAIDA datasets used for security research created https://catalog.caida.org/details/media/2022_caida_measurement_data_infrastructure_overview	12/31/2021	**12/20/2021(A)**
	1.2.2.2	Complete data needs report	9/30/2022	9/30/2022(F)
	1.1.2.3	Combine "Internet infrastructure security vulnerabilities" (see 1.1.1) and Data Needs (1.1.2) reports into a single report and post for internal feedback	3/31/2022	**3/31/2022(A)**
	1.1.3	Monitors Requirements Documented	8/31/2022	8/31/2022(F)
	1.1.6	Mailing lists for monitor deployment collaborating partners created	12/31/2021	**12/20/2021(A)**
	1.1.7	Software to support active probing measurements deployed	9/30/2022	9/30/2022(F)
	1.1.8.1	Preferred Network Function Virtualization framework selected	9/30/2022	9/30/2022(F)
1.2	1.2.1	Data storage hardware requirement documented – draft posted for stakeholders' review	9/30/2022	9/30/2022(F)
	1.2.2	DevOps Infrastructure group created	12/31/2021	**12/20/2021(A)**
	1.2.2.1	Data storage systems specifications documented– draft posted for stakeholders' review	9/30/2022	9/30/2022(F)
	1.2.3	Workgroup reports posted 4 weeks after quarterly meetings	04/30/2022	4/30/2022(F)
	1.2.7	Additional data sets identified	9/30/2022	9/30/2022(F)

1.3	**1.3.1.1**	Bring together industry and academic stakeholders to exchange information on data availability, use and accessibility. Quarterly Meetings conducted, reports published. Gmi-traffic group meeting	12/31/2021	**12/17/2021(A)**
	1.3.1.2	Design GMI Science Gateway Interface	9/30/2022	9/30/2022(F)
	1.3.1.3	White paper documenting catalog objectives, goals, architecture, underlying tools	9/30/2022	9/30/2022(F)
	1.3.1.4	Conduct meetings to discuss approaches to integrating non-CAIDA datasets and tools into GMI3S catalog	03/31/2022	**03/31/2022(A)**
	1.3.2.1	Conduct quarterly meeting to understand gap between privacy-preservation techniques and network and security research needs	03/31/2022	**03/31/2022(A)**
	1.3.2.2	Workshops conducted, reports shared with community	03/31/2022	06/2022 (F)
	1.3.2.3	Collect candidate case studies from community	03/31/2022	**03/31/2022(A)**
		Summary of cases identified from literature	9/30/2022	9/30/2022(F)
	1.3.3.2	Biannual meetings with at least two NSF-funded projects	03/31/2022	**03/31/2022(A)**
	1.3.3.4	At least one agreement created and evaluated	03/31/2022	**03/31/2022(A)**
	1.3.4.2	Quarterly meetings conducted	6/30/2022	6/30/2022(F)
1.4	1.4.1	Workshop conducted	02/28/2022	6/30/2022(F)
	1.4.2.1	Virtual collaboration environment launched	12/31/2021	**12/20/2021(A)**
	1.4.2.2	Virtual collaboration environment evaluated and improved	9/30/2022	9/30/2022(F)

3. Technical progress highlight

Highlight of current period technical progress for each subsystem

1.1 Design infrastructure for Data Acquisition.

1.1.1 Our major activity in this reporting period was the work on the report on infrastructure security vulnerabilities (MS 1.1.1). We have completed a first draft[1] which we are now circulating for comments.

Other accomplishments include:

- Claffy and Clark completed the revision of our paper on Measuring the Internet for the Public Interest and submitted it to the *Journal of Information Policy*. This paper makes

[1] Username: `reviewer; Password: MSR1-view-m1-m36`

recommendations aligned with the proposed GMI infrastructure, to solicit feedback from the information policy community. The paper was accepted for publication and will appear on CAIDA's web site in May.

- Clark launched a preliminary study of data and measurement requirements to isolate impairments in Internet access technology. Performance impairments can sometimes be the only indicator of an ongoing attack, but more generally this sort of performance measurement is another data-related challenge the research community has persistently contended with for the last 20 years. The lessons are useful both in terms of data needs and measurement methods, as well as helping to identify the range of issues that influence user quality of experience.
- Clark prepared a short summary of the new security regulation from the EU (NIS-2), which imposed specific operational requirements on various parts of the DNS system. This is relevant in that it may shape the important analytical questions and data needs with respect to DNS security. We will publish a version of this analysis on the web site before August.
- Clark (MIT) and Claffy (UCSD) worked on the next version of our "A Proposal for a Zone of Trust to Improve Security in Internet Routing" paper on improving the security of BGP (MANRS+) which required the development of new analytics to process the CAIDA AS-Rank data. We are exploring how public BGP data can be used to help advance the MANRS initiative, which is an industry organization that has put forward practices intended to improve the security of the Internet. Using the MANRS initiative as a test case helps us understand what data about BGP is useful, what is not available, and what sorts of analysis of available data provide practical insights useful to industry actors.
- Clark, Testart (MIT), and Claffy (UCSD) worked on response to the FCC's Notice of inquiry regarding Internet routing security, due April 13.

1.1.2 We achieved significant progress on the "Data needs report (month 12 MS 1.1.2.2). We identified relevant existing and desirable BGP, DNS, and Certificate Authority Systems primary and derived data and catalogued it in our infrastructure security vulnerabilities report (1.1.1) We will send an initial draft out to the team in April, and integrate feedback, and send the revised report to GMI Strategic Advisory Council (SAC) in May for review and discussion at first SAC meeting in June.

1.1.2.1 As a part of our work on task 1.1.1, Clark started a first catalog of data needs for DNS and BGP security research. Elena and Brad (UCSD) also created a slide deck with detailed information about existing CAIDA datasets used for security research.

1.1.2.3 We invested efforts into further development of CAIDA's Internet science resource catalog which we are using to design and prototype a GMI catalog. We added metadata including fields/variables description, licenses and legal agreements. This catalog currently indexes only CAIDA's data.

1.1.3 We made progress in documenting the monitors requirements (month 11 MS 1.1.3). We addressed monitors requirements for BGP measurements, for 2-way traffic measurements (on a 100G link), and darknet measurements.

1.1.4 Dan Andersen made progress in a development of a prototype 100G capture machine (Related to this project but funded by CCRI project CNS-2120399) to capture 2-way traffic. After capturing and post-processing limited packet traces on the computer/storage server, he was unable to get *wdcap* to work on the storage server. He came up with an alternative plan to move the card over to the storage server directly. This design alternative will limit us to capturing raw packet traces.

1.1.5 We made progress in developing the monitors software:

2-way traffic monitors: Dan Andersen started to work on the software integration to enable 2-way traffic 100GB capture and processing. kc and Dan met with DREN and LLNL teams to discuss the status of the 100GB traffic capture prototype. (Related to this project but funded by CNS-2120399) Dan got the multicast *libtrace* broadcaster working on the 100G capture card. As described in 1.1.4, he encountered hardware design problems which he is now resolving.

Darknet traffic monitors: We worked on resolving the problems with capturing darknet traffic. We started to observe more frequent spikes in the background radiation signal. Our current setup does not allow us to collect data when traffic volume exceeds a certain bitrate. Ricky Mok designed and implemented a new scheme with source IP filtering capabilities. We are currently debugging the code and will add it to our pipeline in the near future. The code will be deployed via Jenkins.

BGP monitors: We explored options for developing new RV monitoring software that would make it easier to maintain and extend. There are two version of BGP data coming out of the RV infrastructure: MRT files, and raw BGP updates. Separately, UO/NSRC has set up a BGP Monitoring Protocol (BMP) platform that uses the IETF BMP standard to generate richer BGP data from each attached peer, which would be transformative for the research community in terms of questions they could answer (e.g., all paths rather than best path for each peer). (Note -- this prototype deployment has fewer peers than standard RV infrastructure.). The software we currently use for this BMP platform is openBMP. While open source, the OpenBMP libraries are written in C, which makes it a higher bar for researchers and IT folks to maintain and extend. We compared and contrasted functionalities of OpenBMP and goBMP, a recent open source software implementation developed by some employees of Cisco as a side project. We met with S. Bezverkhi (Cisco) to confirm the possibilities of implementing crucial RV data management functionalities in goBMP. These functionalities include accommodating the OpenBMP header and message formats used by CAIDA/RV, the ability to introduce new BMP peers without restarting the process, and the ability to specify Kafka topic and group mappings.

We convened weekly technical CAIDA/NSRC zoom meetings during which we discussed issues and next steps of the RV BGP-related tasks. Hervey sends notes out to the team after each meeting. The notes are available by request.

In preparation to onboard new team members for monitors maintenance, David Teach documented the RV maintenance tasks and time/resources needed. Ryan Leonard (NSRC) created a presentation[2] describing the RV data processing pipeline, as part of the effort to reconsider and refactor the design and work flow and to lower the bar for other community members to participate m maintenance.

To ensure redundancy and availability of Kafka services, NSRC started transitioning from a single node Kafka cluster to a 3-node cluster. NSRC compiled a list of key monitoring specs requirements for the Grafana interface, and compiled requirements for codifying the BMP process into Ansible.

1.1.7 Our activity in this reporting period was the deployment of scamper on RouteViews collectors. This functionality facilitates active probing from RouteViews collectors.

- NSRC sub team worked with Matthew Luckie (University of Waikato) to package Scamper for RouteViews CentOS, and Ubuntu based collectors.
- NSRC leveraged their automation tools to deploy Scamper on all 20 modern, internet exchange connected collectors.

[2] Username: reviewer; Password: MSR1-view-m1-m36

1.1.8 NCRS team is now working on prototyping the RV collectors on UO virtual machines. The majority of RV collectors are already running on VMs. NSRC runs mostly a RHEL7/8 and CentOS stack today and are transitioning to standardize on Ubuntu due to recent issues in CentOS upstream. NSRC developed detailed instructions on deploying collectors on VMs.

We began our literature search to familiarize ourselves with the rapidly evolving world of Network Function Virtualization and how it could support our measurement and data management architectures.

1.2 Design Infrastructure for Data Management

1.2.1 As a part of our work on task 1.1.1, David Clark started a first catalog of data needs for DNS and BGP. This work will continue in the next reporting period.

We conducted weekly sysadmin and data admin meetings about evolving data storage needs. Dan and Paul (UCSD CAIDA Sysadmin team) are almost finished with automation of the production of reports on the current data volumes and growth rates. These reports include data on our disks, OpenStack SWIFT, cloud storage, and at NERSC.

Dan Andersen and kc (UCSD) assessed the storage requirements and created purchase orders for required hardware (a database server, a storage server, and a compute server).

1.2.2 (Overlap with 1.2.1 above) Claffy created a DevOps Infrastructure working group to begin discussion of big data management technologies that will inform our hardware requirements and specifications. We hold the first *gmi-dataops* meeting in January 2022. This meeting was mainly informational, we identified interest groups and outlined the goals. We are still waiting (8 months later) for hardware to arrive.
Claffy met with Christine Kirkpatrick (director of the SDSC Research Data Services division, Co-PI of the West Big Data Innovation Hub, and of the Open Storage Network) to discuss best practices in addressing the DevOps challenges in designing hypervisor environments for researcher access.

1.2.3 Yulaeva and Huffaker attended "Leveraging Data Communities to Advance Open Science" NSF workshop where attendees shared their experience with creating and maintaining metadata. We incorporated some of the best practices into CAIDA datasets catalog as described in 1.1.2.3

1.2.4 CAIDA team worked on tools for data curation and documentation. Ricky improved the existing FlowTuple analysis tools. As the volume of the data increases, we found that the processing time of the original library for analyzing FlowTuple data collected by the telescope was longer than the time duration of the data on the VMs provided to researchers (e.g., it takes more than a day to process a day's worth of FlowTuples). The performance bottleneck appeared to be the *avro*-format files parsing. We evaluated the performance of avro parsing libraries in other programming languages, including goavro in Golang, and libavro in C. We further improved the performance by introducing concurrency in the parsing process. The updated analysis tool only required about 1/3 time of the original version.

Paul worked on the Kubernetes (open-source container orchestration engine for automating deployment, scaling, and management of containerized applications) deployment infrastructure that will allow seamless deployment of all future applications including CAIDA catalog and "CAIDA Stats" - a science gateway for network researchers.

1.2.5 Paul started to work on incorporating authentication mechanisms into our public data downloads/queries pipeline. He started by updating the NetlifyCMS/Gitlab Authentication method for *www-caida-cms.* We are still investigating the use of Keycloak for SSO.

1.3 Design Infrastructure for Broad Usability

1.3.1 NSRC is collaborating with Google to store the entire RouteViews dataset (from 1997 to today) in the Google Cloud Platform (GCP) for off-site back-up and retrieval. A new application is being developed with the University of Oregon team and Google to automatically refresh the dataset with the latest BGP routing data.

1.3.1.1 We created a two-page project overview to engage a diverse set of stakeholders.

Claffy met with several industry partners to exchange information on data availability, use, and accessibility. For details see *Table 1*.

Claffy also met with Dr. Debabbi (Concordia University CA) to discuss the telescope data sharing with industry.

1.3.1.2 Brad Huffaker led our work on the GMI3S Science Gateway interface design and prototyping (started as part of the DIBBS project that is almost out of funding). The initial design had three components: user-facing web front-end; back-end flow engine; and back-end coordinator. Our goal was to support user-generated data-processing/analysis modules, which we called flows. We designed a platform where the front end displays a list of existing flows and their attributes, launches execution of those flows, and displays results. The back-end flow engine supports the ability to store, compile, and run a flow graph, and maintains a list of existing flows and data produced by execution of these flows. The back-end coordinator handles user accounts and authentication.

During prototype development we deduced that the existing flow engine could be used as an open proxy into CAIDA and UCSD's address space. This would allow an attacker to access resources using address-based access control lists. We began a redesign that would address this problem.

We also encountered setbacks with the front-end design. Our first idea for the front-end flow engine was to represent the program (that processed and presented the data) as a node graph, with different nodes in the graph representing different data processing components. This representation would allow the user to create a data pipeline without the need to fully understand and implement the back-end process; even non-programmers could use it, or so we hoped. That is, users might assemble the nodes from existing pieces of code and the backend programming would occur when the nodes are "wired" together. Several systems, including Epic Engine’s render graph, use this approach.

We learned in our development of the flow engine and accompanying graphic interface that the wire abstraction (described in the previous paragraph) often prevented the user from understanding or reasoning about the feedback loops needed to generate the desired output. We are adjusting our design after integrating feedback from interviews with potential scientific users.

Another problem we encountered was our use of a linear data processing pipeline, which ran counter to the complexity of relationships across Internet data sets and the complexity of many queries users desired to perform. Supporting the required complexity ran counter to the objective of the graphical node-based flow pipeline creation model -- it became more complex than actual programming, rather than simpler.

The node-based programming platform we created also required a lot of backed development, reproducing systems that likely exist out in the commercial world. We are going to pause on this design effort while we investigate commercial (but hopefully with opensource options) systems, rather than trying to build something from scratch.

In summary, we are still investigating how to balance the tension between exposing the full complexity of the system to the user and keeping it simple enough to avoid overwhelming the user. Our goal was that a data flow pipeline would support complex computation on a large set of data on the user's behalf. To support this complexity, we need to build complexity into the language, which includes generating and storing different versions of the pipeline, different runs against that pipeline, binding each run to the pipeline versions, caching large and long-running datasets. But the resulting complexity confused all our test and evaluation users.

Our next design will hide more of this complexity from users, at least in the default view.

Taking these lessons in mind we wish to build a system that will reduce the complexity of the initial user interface, exposing features incrementally to users. We will present users with data sets and schemas rather than the data processing flow pipelines and facilitate joining of data sets to answer specific questions researchers have presented to us. We will explore use of existing data management technology (e.g., Spark, Hadoop, Hive, Hue) with a smaller logic engine.

1.3.1.3 Huffaker led a team working on prototyping the GMI3S Data Catalog, starting with adding more features to the existing CAIDA catalog (developed as part of DIBBS project ending this year). In October-March REU students added eight new catalog recipes which are instructions (with code) on how to solve various Internet security-related data analysis tasks using existing CAIDA and other datasets and tools. We improved search functionality by adding a relevance score and annotation tags.

1.3.2.1 Clark initiated a discussion with ISI (at USC) about a possible collaboration around the topic of secure multi-party computation. For this purpose, Clark created a list of possible challenge questions for SMC experts, to help launch a more concrete discussion. Clark also discussed with ISI a possible collaboration around the objective of mapping data to knowledge, using our preliminary draft of harms related to DNS as a challenge question.

We identified a software package that implements a version of differential privacy, which is being developed at Harvard. An MIT graduate student in our group downloaded and ran this code. We are making a careful study of the documentation and are scheduling conversations with experts to understand the scope of utility of this code. We have identified a possible dataset to use as a test of this code--a dataset of wireless network performance. We have a called scheduled with the researcher currently working on this code, to discuss possible collaboration.

1.3.3.1 We met with Todd Elmer of Iconnectiv (owner of the CLLI codes data set that we licensed at a great discount this year for research/validation on topology research) on how to improve our current data-sharing agreement.

We reviewed our 2014 data sharing agreement with Farsight, which was recently purchased by DomainTools (December 2021). DomainTools needs to update this agreement, and we need to reconsider it in light of our trying to partner with them to create the Darknet Security Information Exchange using UCSD Telescope data. Other meetings are documented in Table 1.

1.3.3.2 Claffy met with the Pacific Research Platform , a partnership that builds on the optical backbone of Pacific Wave to create seamless research platform that encourages collaboration on a broad range of data-intensive fields and projects) to discuss potential synergy among measurement projects.

Claffy and Yulaeva attended "Keeping Networks Innovative Together (KNIT) Winter '21: A FABRIC Community Workshop".

1.3.3.3 We worked with the DARPA Information Innovation Office to identify their security needs and provide them with relevant telescope darknet data (separate funding to keep current telescope instrumentation operational). We had monthly calls with DREN/Lincoln lab to collaborate on security issues.

1.3.3.4 We hold biweekly meetings with Nancy Maron, Rev-up business analyst and instructor to develop agreements with for-profit organizations interested in using our data.

1.3.3.6 We met monthly with RIPE RIS staff to discuss data sharing issues that CAIDA, NSRC, and RIPE all face.

1.4 Design Infrastructure for outreach

1.4.1 We created *gmi-traffic, gmi-bgp, gmi-dns and gmi-dataops* working groups. We also created the *gmi-interest* mailing list for the industry entities interested in collaborating with us. We convened *gmi-traffic* quarterly meeting (2021-12-17). We will have our 2nd GMI-Traffic meeting on 8 April 2022. We have scheduled our 2nd GMI-BGP meeting for 6 May 2022. We will also launch the GMI-DNS and GMI-ActiveMeasurement workshop series in the next quarter. See Table 2 for more details

1.4.2 We settled on MatterMost as our virtual collaboration environment.

1.4.2.1 We created the following mail lists: *gmi-traffic, gmi-bgp, gmi-dns, gmi-dataops*, and the corresponding MatterMost channels

1.4.3.1 We met with professor Aaron Fraenkel to discuss his class capstones projects and how we can build projects around CAIDA Internet measurements datasets and tools. We will continue this thread in the summer.

1.4.3.2 Claffy made a presentation about CAIDA datasets to the UCSD Systems and Networking group (2022-01-26) We began work to develop a course on Network Infrastructure Data Science (NIDS) but other tasking took priority, this will be postponed to Year 2 of the project.

Claffy met with FBI (2022-03) analysts to help them understand BGP communities in the context of route hijacks they believe may be occurring on Eastern European/Russian networks.

1.5 Project management

1.5.1 We hold weekly management team (CAIDA, MIT, NSRC) meetings. All the meeting notes are available by request. We started to use Jira/Confluence environment for project management

1.5.2 We convened a "Risk Management" meeting with Jim Olds (GMU), risk officer for the SAGE project. We met to discuss PEP management guidance with Helen Taafe (2022-03-01) who provides PEP and risk management for the SAGE project. We updated the PEP and shared it with the team.

1.5.3 We compiled a list of strategic advisory council members and sent email invitations.

4. Issues and major risks

Near-term major risks: It will probably take another two months to get the hardware we ordered, meanwhile we are using our existing capacities. In case we run out of those, we can use the SDSC and U of Oregon facilities till we get all the necessary hardware.

Human Resource Retention. We lost another key technical person on the NSRC/UO side, David Teach, who was Technical Lead for the RouteViews project. He went to AWS, motivated in part by lack of opportunity for career advancement at the university. This is an acute issue and risk for any Internet research infrastructure project today, given tremendous (up to 10X with stock options) opportunity cost of not moving to industry. Teach trained UO network engineer Ryan Leonard on RouteViews operation before Teach left the project and continues to provide technical support as needed. Both Teach and Leonard had a UO network engineering position as their primary job and were only supporting RouteViews at 10-20% of their time. NSRC is planning to hire a new FTE that will be 50% RouteViews supported by this project. In the meantime, NSRC has arranged for technical support from other community members.

5. Cost and performance summary data

Subsystems (L2 or L3 WBS)	Budgeted Cost	Cumulative Actual cost	Invoiced but not paid subcontracts	Equipment committed	Work % completed
1.1	2,731 k	131K	39K	73K	7%
1.2	1,946 k	274K	6K		16%
1.3	1,390 k	239K	12K		17%
1.4	932 k	77K	6K		8%
1.5	491K	76K	0		15%
Project total	7,865K	797K	63K	73K	

6. Approved changes to the project baseline (if any)

Describe your change and briefly state the reason/justification for the change.

N/A

Table 1: Meetings with Industry Partners

Name	Organization	Position	Meeting Date	Data/tools/discipline	Notes
Harold Feld	Public Knowledge	Senior Vice President	2021-12	Data needs for DPA	https://publicknowledge.org/policy/the-case-for-the-digital-platform-act-executive-summary/
Mark Burgess	Consulting working with Arango (until late 2021)	External Consultant	2021-11	Using Arango for ITDK	Outcome: "Exploring the digital barrier reef with CAIDA and ArangoDB", https://mark-burgess-oslo-mb.medium.com/semantic-spacetime-and-data-analytics-aabbb811cb26
Todd Elmer	iconnectiv	Senior Account Director	2021-11	CLLI codes	https://iconectiv.com/commonlanguage-st1t1 owners of CLLI codes data set that we licensed at great discount this year for research/validation on topology research) on how to make our data sharing agreement worth their while
Ben Cartwright-Cox	bgp tools	independent agent	2021-11	bgp.tools	exploring possible inclusion into catalog/gateway
Paul Vixie	Farsight	CEO	2021-12	DNSDB, SIE	His company was purchased by DomainTools, he is "unemployed" but happy to serve on GMI strategic advisory committee. Lots of discussion of difficulty of sustainable business models
Avi Freedman	Kentik	CEO	2021-12	Commercial Network measurement/mgt serviices	discuss joint position for open-src software engineer Kentik and UCSD (as model for GMI staff). He also agreed to serve on advisory committee.
Sentil Kumar	Capital One	Chief Scientist	2022-01-26	Security	Providing commercial access to two-way traffic data
Jason Kahn	Coalition	Director	2022-03	Cyber security	Providing commercial access to telescope data
Dan Ellard	Raytheon BBN	Lead scientist	2022-03-18	Defense	Working with our telescope data, creating new tools
Romain Fontugne	iijlab.net	scientist	2022-03	Internet Health report	How to mirror IHR at CAIDA
Serguei Bezverkhi	CISCO	Technical Leader	2022-02-17	goBMP	transfer from openBMP to goBMP

Phillipa and Lai Yi	Measurements lab	Project Director	2022-01-31	Internet measurements	Data sharing with google
Job Snijders	Fastly	Principal Engineer	Multiple calls	analyzes and architects global networks for future growth.	bgp data needs

Table 2 GMI Groups

Group name	Purpose	Mailing List	Meetings
gmi-traffic		https://mailman.caida.org/mailman/listinfo/gmi-traffic	2021-12-17 2022-04-08
gmi-bgp		https://mailman.caida.org/mailman/listinfo/gmi-bgp	2022-05-06
gmi-routing			
gmi-topology			
gmi-dns		https://mailman.caida.org/mailman/listinfo/gmi-dns	
gmi-ddos			
gmi-policy			
gmi-economics			
gmi-security			
gmi-dataops		https://mailman.caida.org/mailman/listinfo/gmi-dataops	2021-01
gmi-interest	General list for those interested in GMI project	https://mailman.caida.org/mailman/listinfo/gmi-interest	