Global Measurement Infrastructure (GMI) Design Project Monitor Specification Report v1.0.[1]

Abstract

The CAIDA GMI3S (Global Measurement Infrastructure to Improve Internet Security) Design project investigated and developed designs for a new generation of infrastructure to support measurements of the Internet, a new generation of platforms and tools for data curation and utilization, and support for use of Internet measurement data by the research community. While these facilities are relevant to a wide range of measurements, the project focus was Internet infrastructure security, e.g., vulnerabilities and consequential harms that arise in the *packet carriage* layer of the Internet. We focused on four components: the addressing architecture of the Internet, and systems to support address allocation, management, and use; the global Internet routing protocol (BGP); the Domain Name System (DNS), which maps domain names to IP addresses; and the Certificate Authority system, which manages encryption keys for applications. These elements share three key features: their foundational role for all Internet use, the need for collective action to prevent harms, and misaligned incentives to take such action. We also considered denial-of-service (DoS) attacks, which exploit vulnerable nodes and the Internet's packet forwarding function to overload network components in order to prevent proper functioning. Some DoS mitigations also depend on collective operational practices across the ecosystem, not just actions by potential victims.

Based on a 3.5 year NSF-funded (OAC-2131987) project, this document includes a specification of components to support data acquisition and curation. After an initial discussion of harms that derive from vulnerabilities at these packet carriage layers, we structure the remainder of this document according to the data type or measurement modality: BGP data; active measurement; passive traffic measurement; and DNS data. For each measurement type, we review current capabilities and their limitations, requirements articulated by the research community moving forward, and approaches to management of security and privacy concerns. We propose specifications for monitoring components, including approaches to optimize collection and storage of resulting data. We also propose approaches to achieving goals of performance, maintainability, data integrity, standardization, flexibility/extensibility, and FAIR data principles (Findable, Accessible, Interoperable, Reusable) as stakeholders may move forward with a design and prototyping. We also propose several data curation and analytics components to demonstrate and advance the value of the data to national research priorities.

Note: We have proposed an NSF MSRI-R1 Implementation Phase to cover an Implementation Phase for a subset of this Design Effort, namely the active measurement component, and infrastructure to support curation and sharing of several ancillary data sets needed to interpret measurements enabled by the proposed new active measurement platform. We explain these choices in this report.

[1]This material is based on research sponsored by the National Science Foundation (NSF) grant OAC-2131987. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF.

Contents

Chapter 1

Motivation: Data needs to secure the foundations of Internet infrastructure

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 [124,143]. As a **national research priority** [76,180,234,234,235], 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 inevitably shifting 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.

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. To address this gap, the CAIDA GMI3S (Global Measurement Infrastructure to Improve Internet Security) Design Project investigated and developed designs for a new generation of infrastructure to support measurements of the Internet, a new generation of platforms and tools for data curation and utilization, and support for use of Internet measurement data by the research community.

The ultimate goal is 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. The proposed infrastructure would treat measurement, data curation, and accessibility as foundational capabilities. Informed by NSF's Blueprint for National Cyberinfrastructure [107] and aligned with the National AI Research Resource (NAIRR) vision [140], the platform should 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 NSF-funded (OAC-2131987) MSRI 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). These elements, which we call the *packet carriage service* of the Internet, share three key features: their foundational role for all Internet use, the need for collective action to prevent harms, and misaligned incentives to take such action. We developed and evaluated new approaches to data collection and analysis in these areas [3, 6, 7, 33, 58–61, 74, 77–80, 80, 81, 112, 113, 120, 125, 142, 142, 152, 166, 187, 212, 213, 213, 214, 214, 221], the results of which we codify, reference, and contextualize in this deliverable of the Design Phase. We also considered (Distributed) Denial of Service attacks, (D)DoS, which exploit vulnerable nodes as well as the basic packet forwarding function of the Internet to overload a part of the network in order to prevent proper functioning. We included DDoS in our scope because some mitigations depend on collective operational practices across the ecosystem, not just actions by potential victims.

In our Design Phase we designed several new data collection systems, which we codify in this specification document. 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, deliberation with stakeholders across sectors and disciplines reinforced our view that data relevant to these vulnerabilities are inconsistently collected, and inconsistently available to independent researchers. This situation is not new and likely to continue for the foreseeable future. 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 (§2), and the diverse and often unstable sources of data collection and curation, made clear to us that enabling availability of much of the data will require an organization with stable funding at significant scale.

While we have limited our proposed Implementation Phase for the current round of MSRI funding, this document includes specifications of data acquisition and curation components we investigated in our Design Phase. We structure the remainder of this document according to the measurement type: BGP data; active measurement; traffic

data; and DNS data. For each measurement type, we review current capabilities and their limitations, requirements articulated by the research community moving forward, and approaches to management of security and privacy concerns. We propose specifications for monitoring components, including approaches to optimize collection and storage of resulting data. We also propose approaches to achieving goals of performance, maintainability, data integrity, standardization, FAIR data principles (Findable, Accessible, Interoperable, Reusable), and flexibility and extensibility of the proposed infrastructure. Finally we propose several data curation and analytics components to demonstrate and advance the value of the data to national research priorities.

A notable lesson of our Design Phase is that in many cases researchers would prefer not to gather global measurement themselves but instead would rather rely on trustworthy raw and curated 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.

1.1 Established community need

The *scientific justification* of our proposed RI is the need to enable a new generation of critical infrastructure to support Internet 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 [108, 111]. 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. 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** [180] 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** [18] and subsequent update to the national strategic plan on Big Data R&D [16] re-

peatedly 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" [203]. The same year the **National Academy** revisited its report on principles and practices for federal statistical agencies [204]. 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 [234]. 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 [201]. 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 [76, 235].[1]

5. For 12 years, CAIDA's **Active Internet Measurement Systems (AIMS)** workshops [88] have brought researchers, operators, and government stakeholders together to discuss Internet measurement challenges, including the challenge of longitudinal data collection [35, 36, 38–44, 46, 56].

6. For 11 years, CAIDA's **Workshop on Internet Economics (WIE)** workshops [90] have hosted discussions of how to overcome data access barriers for economic and policy researchers seeking to study the Internet [34, 37, 47–55].

7. In 2021, the NSF-sponsored two workshops on **Overcoming Barriers to Internet Research (WOMBIR)** [91]. Participants 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 [108].

8. In 2023, the NSF-sponsored **Workshop on Emerging Research Opportunities at the Intersection of Statistics and Internet Measurement** concluded: "While

[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." [76]*

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. [12]." 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 [102].

9. The **Department of Defense** has funded research programs aimed at enhancing its ability to securely operate over commercial Internet infrastructure [109,184], yet it lacks a dedicated program to support the foundational infrastructure necessary for scientific data collection and topology analytics that underpin this mission.

1.2 Research community benefits

We are proposing infrastructure that will gather data to benefit Internet infrastructure 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.

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 network resilience and broadband service quality across the globe, and how traffic to important services is routed across national boundaries.

Our project targets NSF's articulated research infrastructure goals [107], namely: to improve the process of accessing, integrating and transforming data to knowledge and discovery; 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. We are also specifying infrastructure that 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 [2] needed to navigate other current and future Internet-related harms. The fields of cybersecurity and Internet infrastructure research broadly exemplify these necessary changes in scientific engagement: the need to gather data and test ideas at scale with real-world constraints as part of the research.

1.3 Barriers to collective action to secure the Internet

The packet carriage service of the Internet is a foundation on which every application depends. The designers of every application that operates over the Internet must consider whether and how to attempt to mitigate the harms that may arise due to poor security at the Internet layer. Poor security imposes a cost (or a risk) that every application bears.

But there is a critical difference between end-node vulnerabilities and vulnerabilities in the Internet itself. Organizations that connect to the Internet can take many steps to improve their own security posture, with the help of many published best practices in user authentication, system patching, secure backup, business continuity planning, etc. Individual enterprises can assess their risk profile and invest accordingly.

In contrast, organizations that simply use the Internet are often not in a position to defend themselves from harms that arise from insecurity in these foundational layers of the Internet itself. Such harms may arise in parts of the Internet that are far removed from the firm being harmed, and the harmed firm may have little recourse. Mitigation of the risks to the connected firm depends on the collective action of the providers of the core Internet services. However, those actors in a position to mitigate the vulnerability often have no (or limited) incentive to take the required action. The combination of economic pressures, tensions among competing operational objectives, and problems of coordination raise formidable and persistent challenges to improved security of Internet infrastructure.

In the decentralized space of Internet operations and governance, there is often no coordinating actor with the authority to mandate a specific change in an Internet service, or even the standing to encourage a change. The Internet Engineering Task Force can create a new standard (a process which itself may fail to resolve disagreement), but the creation of a standard does not ensure its uptake. In some cases, a sufficiently powerful centralized actor can set a direction and effectively push a change into the ecosystem (for example, Certificate Transparency) but in many cases progress depends on collective decision-making and commitment. This is problematic for four reasons, aside from the fundamental challenge of misaligned incentives:

- There is often no clear agreement as to what behaviors by different actors actually constitute a malicious act (as opposed to utilizing features of the Internet as they were intended to be used, but to the disadvantage of one or another actor.)
- The Internet (and many of the malicious actions) span jurisdictions.
- There is no actor with the authority to mandate collection of relevant data.
- The Internet protocols were not designed with measurement in mind, and gathering data often depends on opportunistic methods that are at best a compromise.

Organizations connected to the Internet cannot fully protect themselves from poor security, and collective action to improve it is hard. Better data can reveal vulnerabilities, inform assessments of mitigation proposals, and clarify deployment complexities. Thus, improving visibility into these problems is the most urgent need. However, collective action also requires transparency about participation, and when measurement is needed for this, it is crucial to minimize the cost and risk of providing such data.

1.4 Role of data in identifying, assessing, and mitigating harms

Navigation of security threats occurs at five levels: prevention, tactical defense, forensic analysis, strategic mitigation, and longitudinal assessment. All of these require data, the role of which we discuss.

Prevention: Absolute prevention is ideal and sometimes achievable for specific attacks, which then fade from view. However, ongoing monitoring is needed to ensure prevention remains effective, as threats evolve and attackers adopt new methods.

Tactical mitigation: During exploitation of a vulnerability, the immediate question is how is the attacker crafting the attack from the parts provided by the ecosystem? Defenders need timely evidence of the specific attack, details of the attack and the nature of the attacker, etc. The operational reality of most security threat navigation today is attempting to prevent intrusion or compromise using access control lists and/or blacklists. Today, tactical mitigations are typically undertaken by private actors, who often must act with uncertain authority and powers. They also usually operate without access to information that governments might obtain through formal proceeding, but the complexity and delays of such a process are themselves barriers to tactical mitigation.

Forensic: When harms result from an attack, data is essential to assessing the harm. Observable evidence of an attack does not imply the attack was successful, or a material cause for concern. In order to assign a priority to mitigating a vulnerability, we need to establish that the resulting harms are real.

Strategic: Data informs proposed changes to systems and work flows. At first proposed changes are hypothetical. We cannot measure their behavior to see how they will mitigate vulnerabilities. Instead, we use data we have about the system, combined with our best models of how the change will affect the system, to predict the utility of the proposed mitigation. This requires data about the functioning of the system, including the range of benign and malicious actions. Analysts also need data to estimate the magnitude of future harms, to justify deployment of changes to the system.

Longitudinal: Defenders need real-time data, while analysts require historical data to understand trends and predict future harms. Consistent long-term data collection is crucial for strategic mitigation.

Beyond current attack patterns, security planning must anticipate emergent threats exploiting existing vulnerabilities. The rise of ransomware illustrates how previously

known security weaknesses can underlie new, more damaging attack paradigms. This raises ethical and operational questions: should operators address well-known vulnerabilities now to mitigate potential future catastrophes, as part of their duty to maintain overall Internet hygiene?

While initiatives like Mutually Agreed Norms for Routing Security (MANRS) [211] demonstrate willingness among some operators to adopt improved practices, adoption rates and compliance remains an open challenge. When mitigation is costly, voluntary compliance may remain low, suggesting that regulation could be justified to ensure equitable cost distribution among operators. Yet, making a compelling regulatory case requires robust data on harms and persuasive risk models projecting future threats. Without data on actual harms, it is hard to argue which vulnerabilities to address. Measuring harm is difficult; for example, BGP hijacks can be observed over months or years [225], but their true impact remains unclear. This lack of harm data leads to differing views on the urgency of mitigation.

Gathering such data is challenging for two reasons: underreporting and attribution ambiguity. First, victims may choose not to disclose incidents, or may lack knowledge of where or how to report them. For instance, while the U.S. Federal Trade Commission collects fraud reports, the proportion of unreported harms remains unknown. Second, it is often unclear which attack led to an observed harm, complicating analysis and targeted mitigation. This leaves defenders struggling to make a case that the harm is important enough to prioritize among all the other issues that contend for attention.

Moreover, the costs and incentives for proactive security are misaligned. Although Internet operators could improve security through better operational practices, the costs often outweigh direct benefits to them, creating barriers to voluntary action. Given that many attacks have a low success rate but minimal execution cost, continued attack attempts remain rational from an adversary's perspective. Conversely, if the cost of mitigation exceeds the magnitude of harms, a risk management approach, such as cyberinsurance, may be more economically efficient. However, insufficient data prevents a rigorous evaluation of these options.

Another barrier is that the victim may not understand exactly how – or even that – the harm occurred. If a customer is redirected to a malicious web site that steals personal information, this attack may rely on a BGP or DNS hijack. The firm may be able to tell that a customer had personal information stolen, but not how. This lack of data about the methods of attacks drives ongoing disputes about the relative priority of proposed mitigations.

Chapter 2

Mapping vulnerabilities to data

Measuring harms in general is beyond what a technical infrastructure can do, but connecting vulnerabilities to what can be measured and analyzed is feasible and was one of the early tasks in the Design Phase. We collected an inventory of datasets relevant to research related to vulnerabilities in the Internet packet carriage systems. This inventory helped us understand requirements for measurement infrastructure, the features and capacity requirements of platforms for data curation and utilization, and opportunities for collaboration with other groups that collect relevant Internet data.

We identified not only data that are currently collected, and data that might be collected using this new generation of infrastructure, but also data that are collected by other groups, and data that do not now exist (or even where there is no obvious way to collect it) but which would be useful if it were possible to obtain. This planning phase required analysis of barriers to the collection of relevant data, both technical and non-technical. Our inventory covers data that may exist, such as within firms that operate parts of the Internet, but are not currently available to the research community.

To structure our search for relevant datasets, we first identified vulnerabilities in the four foundational systems: addressing, routing, naming, certificates. For each system, we summarize the known or potential vulnerabilities, and possible mitigations to these vulnerabilities. We discuss the role of data at each step, and identify which data plays a role (or *should or* could play a role) in improving the security posture of that system.

Since a mitigation may introduce a new set of vulnerabilities, the design of actions to improve security is iterative, where a vulnerability may suggest possible mitigations, those mitigations may in turn have vulnerabilities, and so on. Improving the security of the Internet requires recognizing the dynamics of the ecosystem, in which actors adapt in response to a given adjustment.

2.1 Addressing architecture vulnerabilities, harms, and mitigations

Endpoint IP addresses are the most fundamental building block of the Internet–they identify the destination to which a packet is to go. Internet routers use the destination address to decide, at each hop across the Internet, how to forward the packet onward. Packets typically include a source IP address in the header, which indicates the originating network address of the sender (though not always uniquely identifying the true endpoint). This information enables the recipient to direct replies back to the sender.

2.1.1 Vulnerabilities

The critical vulnerability of the Internet addressing architecture is the ability to forge (*spoof*) source IP addresses, i.e., putting a false address, rather than the actual sending IP address, in the source address field of the IP packet. Attackers use spoofing to launch a DoS attack against a victim without revealing the attacker's identity or location on the network (§2.5). The harmful consequence is the victim will communicate with the malicious actor as if it were the intended endpoint. Another type of spoofed attack is an *impersonation attack*, where an attacker usurps addresses not allocated to that actor, and attempts to send packets using those addresses. One could use an innocent actor's address space to perform malicious acts that cause the addresses to be blacklisted, denying the original owner reliable use of them. Appropriation of unauthorized addresses is often accomplished before or via a BGP hijack.

In a *reflection/amplification attack*, an attacker uses spoofing to put the victim's IP address in the source IP address field of the packet. Such attacks generally exploit a protocol where a small query produces a large response, giving the attack *amplification*. Protocols used for amplification attacks include DNS, NTP, and memcached. The harmful consequence is that the receiver of the packet (the amplifier) replies to the spoofed address (the victim), sending the victim needless traffic which can overwhelm the victim. This capability is the basis *Denial of Service* attacks.

2.1.2 Mitigation Strategies

The best known but not universally deployed mechanism to prevent such attacks is **Source Address Validation (SAV)**. Under SAV, ISPs verify the source addresses of packets from their customers and discard those with spoofed addresses, following Best Current Practice 38 (BCP 38) [85]). However, SAV suffers from a classic incentive misalignment: ISPs gain little direct benefit from implementing BCP 38, yet risk operational issues if it is misconfigured. While deploying BCP 38 can help prevent some DoS attacks, those attacks may occur far from the implementing ISP's own network. Implementing BCP 38 also carries costs, mainly the ongoing operational burden of configuring and maintain-

ing it. In addition, ISPs often receive no immediate feedback if their configuration fails, leaving spoofed packets undetected.

Role of data: Tracking network compliance with BCP 38 is crucial for mitigation, as networks permitting spoofed source addresses create systemic risks. Publicly identifying non-compliant networks can drive adoption of source address validation (SAV). For cyberinsurance and auditing frameworks, such data enables assessment of a network's security posture, incentivize compliance through premiums or policy requirements, and monitor SAV deployment trends over time and across regions to evaluate collective risk.

2.1.3 Primary data

We use the term *primary data* to describe datasets that directly result from collection. If such data is not gathered at the time of observation, it is typically lost. When collected and archived over time, primary data can support both real-time (tactical) analysis and long-term (longitudinal) studies.

Table 2.1: IP addresses: Primary data

Type	Source	Status	Limitations	Uses
ISPs that do/do not implement BCP 38	CAIDA Spoofer [162]	Available (not currently funded)	Tests must run from inside ISP.	Track/verify compliance with BCP 38.
DSAV Test	BYU IMAAL	Operating (tool more than data)	Primarily checks for spoofed DNS traffic; Single point in time	Auditing of inbound traffic filtering

2.1.4 Derived data

Although there have been one-off studies of how to infer spoofing from traceroute data [151] or IXP traffic data [177], none of these projects have resulted in ongoing sources of derived data regarding spoofing.

Table 2.2: IP addresses: Derived data

Type	Source	Status	Limitations	Uses
The Open Resolver project	Open Resolver Project.	Inactive	Requires DNS forwarder that does not rewrite source address of query. False positives due to sibling ASes.	Detects lack of SAV if authoritative resolver receives query from different ASN than OR-queried forwarder.

2.2 BGP routing architecture vulnerabilities, harms, and mitigations

The Internet is made up of regions called *Autonomous Systems (ASes)* under independent control by their providers. Today, there are ~75K ASes that make up the Internet, most ~70K of which are *stub ASes,* which are networks that do not provide transit to other ASes. The remainder are some form of transit service providers. The Border Gateway Protocol, or BGP, is the global routing protocol that independent networks use to exchange and process routing information that hooks these regions together to make up the global Internet. BGP messages provide autonomous systems (ASes) the information necessary to forward packets to the final AS that hosts the destination.

How BGP Works Each Autonomous System (AS) tells its immediate neighbor ASes the address blocks or *prefixes* (contiguous addresses with common numeric prefix), that it controls and uses. This step is called *originating* a BGP announcement. Each neighbor AS applies its own policies to filter the announcement, often propagating them to its immediate neighbors, which allows the information to spread globally. As the announcement travels, each AS appends its AS number to the announcement, so at any point the message includes the sequence of ASes that define the path back to the originated address block. Every BGP-speaking router then uses the received AS paths to update its own forwarding table that specifies the "best" next hop for sending packets to each destination prefix.

2.2.1 Vulnerabilities

The critical vulnerability with BGP is well-known: a rogue Autonomous System can falsely announce that it originates, or is in the path to, a block of IP addresses it does not control. BGP lacks built-in mechanisms to prevent such false announcements. When routers accept these bogus claims, traffic destined for the hijacked addresses is diverted

to the rogue AS, which can then drop, inspect, manipulate, or send traffic masquerading as those addresses. This attack, known as a *BGP route hijack*, allows a malicious AS to falsify any part of a BGP announcement—including the origin prefix, AS number, or path. Hijacks can target critical Internet services such as name servers, Certificate Authorities, or Regional Internet Registries, potentially causing widespread disruption.

2.2.2 Mitigation Strategies

To prevent origin hijacks, providers can implement *Route Origin Validation (ROV)* by validating the origin prefix and AS announced by their customers. This validation uses *Route Origin Authorizations (ROAs)*, data from the Internet Routing Registry (IRR), or pairwise validation with customers. However, ROV requires prefix owners to have registered ROAs for their prefixes. If these are in place, any AS can drop invalid origin announcements, not just direct providers.

Additionally, providers must verify that the AS number used by a customer is legitimately registered to them, as rogue actors may announce using unassigned or hijacked AS numbers, undermining the integrity of routing and detection efforts.

Finally, BGP *path hijacks*, where attackers announce invalid paths rather than origins, are not detected by simple ROV. Addressing these attacks requires more advanced validation mechanisms to ensure end-to-end routing security [61].

Role of data:

Key questions for assessing the effectiveness of current routing security technologies include: How many providers validate their customers' BGP announcements? What fraction of invalid announcements are dropped? To what extent does this limit the spread of invalid routes across the Internet? How do these answers change over time? While measuring the adoption of Route Origin Authorizations (ROAs) is relatively straightforward, assessing Route Origin Validation (ROV) at scale remains challenging.

There is ongoing debate about the value of blocking simple invalid prefix attacks. One view argues that since attackers can easily shift to more complex invalid path hijacks, mitigating only invalid prefix hijacks has limited benefit. The opposing view suggests that invalid path hijacks are less practical for attackers because such paths become longer than valid announcements as they traverse the Internet, making them less likely to be chosen by routers. This unresolved debate hampers progress in designing effective and comprehensive routing security strategies.

Role of data in understanding path hijack attack surface:

To assess the attack surface, we need to understand how successful these hijacks are or could be. Combining Internet topological maps with hypothetical attacker and victim placements can help identify regions of vulnerability and potential harm. Although AS

topology maps are incomplete, resources such as CAIDA's AS relationship data provide valuable insights. Additionally, the Regional Internet Registries (RIRs)–AFRINIC, APNIC, ARIN, LACNIC, and RIPE NCC–publish current and historical statistics on ROA registrations, enabling analysis of address space coverage.

Most popular applications today use cloud-based services with multiple connection points to the Internet, and often support anycast provisioning of the underlying service. With anycast, a single set of IP addresses are announced from many different locations (instances) on the Internet, to improve performance, resilience, and robustness to failures. While providers know the locations of the anycast instances, this data is generally not public. To gather it, probes can be positioned globally and perform a DNS lookup of the application service to discover the IP address for that service in that region, known as its *catchment*. Smaller catchment regions result in shorter BGP announcements, making them harder to hijack. With catchment maps, analysts can determine the attacker vantage points most effective for hijacking specific services.

Assessing harm also requires understanding user location. For example, a path hijack affecting only users in a distant country poses minimal risk to a regional U.S. bank. For targeted analysis, knowing an application's customer base and catchment locations enables confident assessments. However, evaluating the risk across many services requires data on user connection patterns for a range of applications.

Ultimately, understanding the importance of path hijacks requires building models of attacks and mitigations informed by the best available data.

Role of data in understanding deployment of mitigations

Three proposed but not yet widely deployed mitigations address path hijack vulnerabilities: BGPsec [146], ASPA [10], and the Zone of Trust framework we proposed [61] as part of this MSRI Design Effort. All require ongoing deployment monitoring to evaluate effectiveness. BGPsec provides cryptographic validation of BGP paths. ASPA (AS Path Authorization) verifies AS path segments to prevent route leaks and hijacks. Zone of Trust extends the MANRS [211] framework, creating an incentive-compatible program that protects participant routes. This approach relies on extensive BGP data collection and analysis to ensure conformance with trust zone practices [61].

Another class of routing vulnerabilities stems from potential compromises of Regional Internet Registries (RIRs) and their supporting systems, such as routing registries or ROA software. Understanding these risks requires data to answer critical questions: Do trends indicate increasing targeting of RIR infrastructure? Are RIRs following best practices for securing critical services, e.g., testing resilience of RPKI software? Have they established procedures to detect and remove malicious delegated repositories?

Finally, BGP communities can be exploited for precise interception attacks [17], remote black-holing, traffic steering, and route manipulation—even without prefix hijacking [219]. The flexibility of BGP communities, allowing ASes to assign their own local

meanings, makes interpretation and filtering challenging for recipients. Consequently, operators often propagate these communities without full understanding, increasing the potential blast radius of malicious use. Efforts by the IETF to standardize BGP community practices have had limited success. The security research community would benefit from a dictionary of BGP community values and interpretations, and automated techniques to classify BGP community usage in the wild.

2.2.3 BGP Primary data

Table 2.3: BGP Data Vulnerabilities and Mitigations

Type	**Source**	**Status**	**Limitations**	**Uses**
Collected BGP route announcements	NSRC RouteViews and RIPE RIS	Public. Real time & historical	Limited view of total announcements across net	Detection of hijacks, deriving topology maps
BGP route announcements at small IXPs	PCH BGP data	Public, Real-time & historical	Only updates. Peering rather than full views.	Detection of localized hijacks, understanding peering ecosystem
BGP route announcements	Source: Companies: Akamai, AWS, Google, Kentik	Collected but not shared.	Not available.	Proprietary research/ops
Assertions about valid announcements (ROAs)	RIRs, e.g., Historical: RIPE	Real time available, historical available	Providers may be vulnerable. Data may be erroneous.	Determine validity of a BGP announcement
IRR records	Internet Route Registries.	Available, no complete history	May be vulnerable to attack. Weak authentication	Determine validity of a BGP announcement

Type	Source	Status	Limitations	Uses
Transit topologies for DNS servers	DZDB (TLD zone files)	Available	Daily samples miss short attacks.	Analyze BGP attack surface of DNS infra.
WHOIS DB dumps	RIRs	Available, not public.	CAIDA archives quarterly.	Identify address owners
Peering policies and self-reported presence at facilities	PeeringDB	Available via public API	Data can become stale if ISP does not maintain it.	Validate inference of AS properties, policies, presence at facilities

2.2.4 BGP Derived data

Table 2.4: Derived BGP Data Vulnerabilities and Mitigations

Type	Source	Status	Limitations	Uses
AS relationship	ASrank. Source: CAIDA	Available. Real time, historical	Heuristic inference. Inherits visibility limitations of RV/RIS	deriving topology maps
AS hegemony	IIJ [198] Internet Health Report	Real time and historical available	Heuristic inference. Inherits visibility limitations of RV/RIS	Topology analytics
AS interconnections	Hurricane electric	Available	Visibility limitations of BGP tables.	Topology analysis

Type	Source	Status	Limitations	Uses
AS to owner mapping	CAIDA AS2org	Available	Incomplete underlying (WHOIS) data	Topology analysis
List of ASs that drop invalid BGP announcements	Rovista, RPKI Stats, Cloudflare	Available	Hard to track trends due to multiple factors influencing results.	Predict propagation of invalid routes.
Lists of ASes that announce invalid routes	IIJ Internet Health report	Available, not currently funded.	Limited snapshots, partial BGP/ROA coverage	Tracking ROA adoption
Tactical blocklists	Many: see Section 4.2.	Variably available.	Undisclosed sources. No way to validate.	Allow blocking of traffic from addresses labeled as malicious.
Announcement history	RIPE Stats: Routing	RIPE Stats	Derived from RIPE RIS BGP data	Detect hijacks or leaks.
NIST RPKI dashboard	Source: NIST	Real-time and some historical	Derived from RV data	Track use of ROAs.
BGP community dictionary	CAIDA	BGP Community Dictionary	Old, not updated	Interpretation of BGP data
Valid ROA feed	RIPE	RPKI Validator	Only reflects ROAs under RIPE	Tracking ROA adoption

2.3 Domain Name System vulnerabilities, harms, and mitigations

The Domain Name System, or DNS, performs the essential function of translating higher-level names for endpoints (e.g., `www.example.com`) to the corresponding IP address. An oversimplified model of the DNS involves two stages: registration of a new name, and resolution of that name into an address. In the registration stage, the provider of a web page (or other named resource in the Internet), typically picks an available name in a top-level domain (TLD) of its choice (e.g., .com) and registers that name. A registrant looking to obtain a domain name under .com would contract with a registrar (e.g., Enom) who in turn interfaces with the registry operating .com, Verisign, to query the availability of the domain name and then claim it on behalf of the registrant. On successful purchase of a domain, the registrar is then responsible for the domain until it expires or is transferred by the registrant. In addition to contracts with the registry, registrars also have to be accredited by ICANN

The second stage occurs when a program (such as a browser) encounters a domain name (often as part of a URL) of a resource, and wants to connect to that resource, which requires resolving that name into an address. Computers attached to the Internet usually have software called a *stub* resolver which performs that task. The stub resolver normally contacts a *recursive* resolver to pursue complete resolution of the name. The recursive resolver will take each element of the domain name in turn (hence the term *recursive*) and contact the authoritative name server for that element, to find the address of the server for the next element of the domain name, and finally the address of the resource itself. Thus, given the name `www.example.com`, the recursive resolver will first contact the root name server to find an address for the name server (NS) for the com top-level domain, contact that name server to find the address of the name server for example.com, and then contact that name server to find the IP address of www.example.com.

Many enhancements and details make this work. For example, when a recursive resolver resolves a name (such as com) it will *cache* or remember the result, so it need not repeat the query. A name can map to another name, rather than an address, and the recursive resolver will resolve that name in turn. When the recursive resolver has found the address of the ultimate resource, it will return this value to the stub resolver as the result of the query. Many organizations operate recursive resolvers. Most ISPs operate a recursive resolver for their customers. Large Internet firms also provide a recursive resolver as a service, including Google, Cloudflare, Quad9, and others.

2.3.1 Vulnerabilities in the DNS

The term *vulnerabilities* may not be the best word to describe some of the problems associated with the DNS; a better word might be *abusability*. The design of the DNS makes it easy for anyone to register a domain name, with little to none KYC (Know Your

Customer) due-diligence from the registrars, whether their intended use is malicious or benign. The resulting question is whether it is acceptable to use DNS as a means to thwart malicious behavior, or should it be considered a neutral component in the tension between attack and defense. Both sides are exploiting the features as a tactical element in pursuing their objectives. In this context, we review the many vulnerabilities.

Table 2.5: Table of vulnerabilities, mitigations, and incentive misalignments

Vulnerability	Mitigation	Incentive Misalignment
Service penetration	DNS operators should use best security practices.	Extra cost with no clear benefit
Identity theft	Registrars should use robust methods to authenticate users.	Extra cost with no clear benefit
Operational complexity	DNS providers should provide clear instructions and correctness checkers.	Extra cost with no clear benefit
Plaintext protocols	Use encrypted TCP connection.	Higher latency for DNS queries.
Host misdirection	Users can manually override the default recursive resolver.	Most users have no idea how to do this.
Misrouting of DNS queries	Alternative query protocols like DOH. A browser can choose query protocol and recursive resolver	Selected resolver may not implement desired protections.
BGP hijack of DNS resolver	Use HTTPS and proper key management.	Extra cost with no clear benefit
BGP hijack of name server	DNSSEC (not widely deployed) can provide assurance of authenticity.	Deployment complexity limits uptake.
Malicious name server	Users can register names in well-reputed TLDs.	Lack of user awareness
Cache poisoning	Use DNSSEC.	Complexity limits uptake
Operational complexity of DNSSEC	Provide tools for configuration and checking.	Extra cost with no clear benefit

Malicious name server attacks are rare, and there is limited data on authoritative DNS servers intentionally corrupting DNSSEC chains of trust. Most observed incidents involve misconfigurations or external compromises rather than deliberate malicious behavior by operators.

An important caveat involves countries or networks that deploy rogue or manipulated copies of the root DNS server. While many countries host official ICANN-approved root server mirrors (e.g., local anycast instances) that faithfully replicate the global root zone and maintain DNSSEC integrity, a rogue root server that serves an altered root zone can break DNSSEC protections for users within its reach. In such scenarios, DNSSEC cannot protect resolution integrity because the root of trust itself has been replaced or compromised. When this occurs, domain owners have no recourse at the DNS level to protect the resolution of their domain names for affected users, as DNSSEC assumes trust in the authentic global root key. Alternative protections would need to rely on application-layer security measures such as HTTPS/TLS certificate validation or out-of-band verification mechanisms (§2.4).

More realistic attacks involve *malicious penetration of a Domain Name Server* to modify or add entries to the configured zone. DNS management and configuration complexity contributes to configuration errors, which can allow attackers to take over or manipulate those names. Or, an attacker may be able to steal the credentials of the owner of a domain name, and log in to the registrar using those credentials, effectively controlling the domain and thus any service relying on it.

Another avenue of attack on the basic plaintext query-response protocol of DNS is a *man-in-the-middle* attack. The mitigation is to replace the original query/response protocol with an encrypted TCP connection prevents modification of the communication, but the resulting higher latency for DNS queries may slow the responsiveness of applications.

When a host first connects to the Internet, it receives the address of a recursive DNS resolver to use (usually based on DHCP), which may be a malicious or untrustworthy recursive resolver. Users can manually override the default recursive resolver, but most users have no idea how to do this, or which recursive resolver to pick. In some parts of the world, users may be blocked from picking their own recursive resolver, and blocked from performing their own name resolution, forcing them to use an untrustworthy resolver. Alternative query protocols such as DOH (DNS over HTTP) may make it harder for a restrictive regime to identify and block DNS queries. Also, applications (such as a web browser) can ignore the DNS implementation in the operating system and use its own implementation of a preferred query/response protocol and recursive resolver. However, the browser may choose a recursive resolver that does not implement the desired protections against malicious actions. Users may have no idea what protections they are receiving. The address of the intended recursive resolver can be hijacked, so the user unknowingly connects to a rogue copy of the server. Use of HTTPS and proper key management can reduce this risk, but note that if the CA uses the same hijacked recursive resolver to perform domain validation, the attacker can obtain an apparently-legitimate certificate for the rogue web server. DNSSEC does not protect against this attack, since normally the host trusts the recursive resolver to validate DNSSEC information.

An attack may also *hijack an authoritative name server*, so that the user gets an answer from a malicious variant of the name service. DNSSEC (not widely deployed)

can provide assurance that the answer to a query is authentic. But the cost and complexity of deploying DNSSEC, including user confusion when it fails, has limited uptake of the protocol. Also, an untrustworthy authoritative name server in the chain of trust from the root name server can corrupt the returned result while preserving what appears to be a valid DNSSEC chain of trust. This attack is uncommon enough that there is limited (no?) data on DNS name servers that corrupt DNSSEC chains of trust. The only exception would be countries that deploy their own copy of the root DNS server, which can break the protections of DNSSEC for users in those countries. In such a case, the owner of the domain name has no recourse at the DNS level to protect the integrity of resolution of that name. Today security-conscious users may try to register domain names in a TLD with a history of responsible behavior to minimize risk from this last vulnerability.

Finally, *cache poisoning* occurs when a recursive resolver receives an incorrect response from an authoritative server, and caches that response for use in answering future queries until the TTL of that response expires. Until that time-out occurs, users will be sent to the wrong address. Use of DNSSEC can provide assurance that the answer to a query is authentic, but the increased operational complexity of configuring DNSSEC can lead to operator error and malicious exploits that cause queries to fail, which in turn causes loss of availability. Several organizations have provided tools for configuration and checking of DNSSEC configuration to reduce the risk of these harms.

The vulnerabilities above lead to one of two undesirable outcomes. The first is that the query fails with an error message, and the associated service is unavailable. This outcome leads to frustration, loss of utility, costly complaints, etc. The second is that the user reaches the wrong IP address without knowing. If the user cannot or does not detect that this has happened (see discussion of the CA key management system §2.4), the resulting harm can take many forms. But at the level of the DNS, the harm is that the user ends up talking to the wrong destination. Assessing the final impact of this DNS-level harm depends on many factors beyond the DNS. At a minimum, if the user is sent to the wrong destination, the harm is a loss of availability.

These DNS vulnerabilities arise from its initial design, where the focus was on simplicity, speed of response, and ease of implementation. Lack of attention to security has left a huge attack surface. The FIRST DNS-Abuse Special Interest Group (SIG) has undertaken an effort to organize these vulnerabilities into larger categories that can be mitigated in a systematic way [185].

2.3.2 Mitigation Strategies

Traditional mitigations have taken three forms: hardening resolution, and adding cryptographic authentication (DNSSEC) to the query transaction, and creating (and selling) lists of malicious domains (intended to deceive/defraud the user). Disagreements on the relative importance of these approaches derive from disagreements about what the most important threats are. Resolving these disagreements without concrete data has proven

intractable. However, there is consensus that these mitigations, especially DNSSEC, have created much greater implementation and configuration complexity in the system, increasing the cost to operate services such as recursive resolvers. This complexity brings new vulnerabilities, in that user and operator error create new options for attackers to corrupt the system.

ICANN's KINDNS initiative. ICANN's Knowledge-sharing and Instantiating Norms for DNS and Naming Security (KINDNS) initiative is an effort to strengthen the security and resilience of the Domain Name System (DNS) by promoting voluntary adoption of best operational practices. KINDNS provides operators of authoritative servers, recursive resolvers, and other DNS infrastructure with a curated set of concrete, implementable recommendations designed to mitigate common vulnerabilities, such as inadequate DNSSEC deployment, weak access controls, or insufficient monitoring. Rather than imposing new regulations, the initiative emphasizes knowledge-sharing, capacity-building, and self-assessment, offering training materials, checklists, and a framework to help operators benchmark their practices against community-endorsed standards. By encouraging transparent commitment to these practices, KINDNS aims to foster a culture of collective responsibility in DNS operations—where adherence to proven security norms improves the reliability of the ecosystem as a whole, reduces systemic risk, and builds trust among governments, businesses, and end users.

Sommese *et al.*'s analyses of the KINDNS initiative [213] highlight a critical link between academic research and operational best practices. The KINDNS initiative seeks to codify best practices, but a key barrier is the lack of independent verification of compliance. Without a public, neutral measurement platform, operators have limited incentive to invest in and implement practices that improve collective security. By systematically measuring and visualizing the adoption rates of DNSSEC or anycast over time, the infrastructure should provide the empirical data needed to inform and shape public policy, ultimately transforming a theoretical framework of best practices into a verifiable, measurable reality that strengthens the entire DNS ecosystem.

Risk to longevity of DNS architecture. One long-term reaction to these persistently unsolved vulnerabilities may be a migration away from use of the DNS for name resolution in favor of new alternatives designed with security in mind from the beginning. This outcome becomes more likely as app designers move away from web-based implementations to native application implementations. A web-based app must depend on the name resolution service provided by the browser (which can be the native implementation in the operating system or one in the browser) but a free-standing application is free to use any mechanism it wants to convert a high-level name to an IP address.

2.3.3 Primary data

In contrast to BGP, the DNS has many different sorts of data. There is data about currently registered names, data about how those names are configured, data about who has registered those names, data about usage, and data about abuse. For security researchers, even knowing the registrar for a given domain, or being able to group by registrar on a set of names would be valuable, but access to this type of data requires (generally commercial/contractual) agreement across participating registries to name registrars consistently. GDPR and other privacy regulations have reduced accessibility of this data to researchers.

Table 2.6: DNS: Primary data

Type	Source	Status	Limitations	Uses
Zone files (registered domain names and delegations)	DNS Coffee. Origin: ICANN and registries.	Historical: collected, available	Not all registries make their zone files available. One update per day.	Detection of names suggesting malicious intent.
Rapid Zone Updates (RZU))	Registry	Not widely supported	Not available	Detect attack signals.
Active daily DNS scan	OpenINTEL.	Current, historical. Available	Coverage: 80% of the registered domain namespace. Only CZDS, open/under-NDA ccTLDs + CT Logs domains.	Detect attack signals; map infrastructure changes.
Active DNS scan (ANY, A, AAAA, TXT, MX, CNAME)	Rapid7 fDNS	Current, historical. Available	Coverage: Only CZDS + CT Logs domains. One probe per week.	Same as above.
OpenResolver Census	Shadowserver	Accessible under agreement.	No visibility on private resolvers. Weekly.	Resolver ecosystem and amplification attack studies.
DNS traffic (one day per year)	OARC. Several root servers and TLDs	Available to OARC members	Root-specific view. IPs anonymized. 24 hrs/year.	Name collisions, RSS performance.

Continued on next page

Type	Source	Status	Limitations	Uses
Passive DNS traffic data	DomainTools SIE	Non-commercial use (AUA)	Coverage depends on VPs.	Detect anomalies, domain info
Registration data (WHOIS/RDAP)	Registries, registrars.	Not currently available.	Registrars may have incomplete information on registrants. Privacy issues limit access.	Detect suspicious behavior, e.g., bulk reg.
Enumeration of recursive resolvers.	APNIC Ad measurements, Root, TLD logs	Actively collected data	Incomplete and approximate	Audit conformance to best practices.
Log of queries to recursive resolvers.	Operator of resolver.	May be collected, not available.	Sharing limited by data volume and privacy concerns.	Track interaction with malicious DNS names.
Queries from recursive resolvers.	Source: Domain Tools	Some historical data.	Incomplete picture of query patterns.	What names are queried? Where?
Evidence of malicious DNS names	Many	Variable.	Diverse methods lead to disjoint lists.	Blocking, assessment of abuse.
Pricing data	tld-list.com	Commercial	Limited accuracy	Economic models of ecosystem
Adoption of new protocols (DOH, DNSSEC, etc.)	None	N/A	N/A	Model security of DNS

2.3.4 Derived data

Table 2.7: DNS: Derived data

Type	Source	Status	Limitations	Uses
DNS Databases	Domain Tools DNSDB	Available under agreement	Coverage	Research

Continued on next page

Type	Source	Status	Limitations	Uses
Zone files stats	CAIDA DZDB, dns.coffee	Available as dashboard-/API	Same coverage as available zones (see Table §2.6).	Research
Tactical blocklists	Spamhaus, abuse.ch, Feodo, DShield, FilterLists	Variability available	Derived from undisclosed network monitor sources. No way to validate.	Allow resolvers to block queries to malicious names.
Registrars and registries with many abusive registrations	DAAR, Interisle	available	Derived from blocklists, patterns of registration (limited visibility) DAAR does not provide registrar names.	Coarse visibility of abusive behavior.
Lists of abusive (e.g., phishing) web sites	Large email processors	Not generally public	Inferred from inspection of spam email, etc.	Blocking (e.g., Google Safe Browsing)
Lists of popular web sites	Alexa, Majestic, Tranco, Cisco Umbrella	Variable	Varying methodology to generate lists. Considerable churn.	Modeling collateral harm.
DNSSEC Stats	SWITCH, DNS Resilience	Available as dashboard	Limited to .ch and .li TLDs	DNS resilience studies

2.4 Certificate Authority System vulnerabilities, harms, and mitigations

The Certificate Authority system plays a critical role in the security of Internet services: ostensibly, to provide a final check, after one endpoint has connected to another at a specific IP address, that the entity at that address is the intended one. A *certificate* is an assertion that links a domain name to a public key for that domain. The owner of the domain keeps the corresponding private key, and uses it with a challenge-response protocol that allows anyone to confirm that the domain owner has the private key. The

integrity of this assertion relies on a *certificate authority* to cryptographically sign it, using *its* private key, which is in turn signed by another CA, and so on. The final signature that protects the sequence of signatures is provided by a *root* certificate authority. The public keys of root CAs are included in software such as browsers.

If the CA system works as intended, the vulnerabilities in BGP and the DNS discussed above can at worst lead to a failure of availability. That is, while the CA system cannot ensure a connection reaches the intended destination, it can ideally detect if the connection has reached the wrong destination. Not surprisingly, this causes attackers to target the CA for malicious manipulation. As with the DNS, attacks on the CA system can result in a wide range of harms. Sophisticated attacks often combine abuse of multiple systems, so whatever harm occurs cannot be cleanly associated with a specific vulnerability. ETH Zurich has recently introduced a framework (F-PKI) to allow domain owners to define a policy to specify which CAs have authority to issue certificates for their domain name, and allow clients to choose a policy based on trust levels [31]. This direction is promising, but likelihood of uptake is unclear.

2.4.1 Vulnerabilities

The core vulnerability of Certificate Authorities (CAs) is system compromise. If attackers penetrate a CA's infrastructure, they can gain the ability to issue fraudulent certificates, enabling impersonation attacks and man-in-the-middle interceptions. A well-documented example is the DigiNotar breach in 2011, where attackers compromised the CA and issued hundreds of rogue certificates, including for Google domains, facilitating widespread surveillance [237]. To mitigate such risks, CAs are required to implement industry best practices for operational security and undergo regular independent audits to validate their compliance. The CA/Browser Forum, an industry consortium, oversees CA behavior and enforces accountability by removing untrustworthy CAs from root stores distributed with browsers and operating systems [105].

A CA acting with interests adverse to a targeted service—such as under state coercion—may intentionally issue misleading certificates to facilitate surveillance or interception. While the CA/Browser Forum has the authority to revoke trust in such CAs, detection often depends on external reporting or transparency logs [145, 230].

Another concern is the presence of unexpected CAs in the trusted root store of devices. For example, device manufacturers or distributors (e.g., smartphone vendors or mobile carriers) may install additional root certificates prior to sale. This grants those parties the ability to intercept encrypted traffic for any certificate chains anchored in that added root, creating potential privacy and security risks [106].

Mandated interception. Some firms are legally required to monitor employee behavior (e.g. the brokerage industry must record all conversations with clients), and as part of this may require that employees install an additional root certificate on their work computers so that the employer can intercept and decrypt the communication. Calling

this a vulnerability depends on one's perspective, illustrating a fundamental tension between the goal of privacy and the goal of accountability. The question is whether/how to accommodate this interception within the design of the mechanism (which makes the mechanism explicit and easily a target of abuse) or by forcing the relevant enterprise to break the mechanism.

Imposter names. When users are lured to an imposter website pretending to be a legitimate one, that website normally has a slightly different domain name. The owner of that domain controls it, and can get a valid certificate for that site. The CA system provides no protection in this case. Arguably, this is not a vulnerability of the CA system, but a reflection of an intentional design decision to limit the scope of responsibility of the CA system. The purpose of the CA system is to set up a trustworthy encrypted connection to the server identified by the domain name. It is up to some other actor to decide if the domain name describes where the user meant to go. Evidence suggests that users cannot make this discrimination by looking at the domain name.

Lack of user training; Users may ignore warnings about an invalid certificate and proceed anyway, rendering ineffective the intended protection from the CA. CAs can provide better tools to owners of certificates to automate management and reduce configuration errors, and provide better advice to users about the potential severity of errors.

Attack on certificate issuance. Certain attacks targeting BGP and the DNS can allow an attacker to create an invalid certificate that appears to be legitimate. This vulnerability applies only to the weakest form of certificate, a Domain Validation or DV certificate. Owners of domains could choose to use stronger forms of certificates, such as the Organization Validation or the Extended Validation certificates.

Lack of independent knowledge of certification type. Browsers have no way to know what sort of certificate they should receive. If the owner has obtained an organization-validated (OV) certificate, and the attacker sends a domain-validated (DV) certificate, the browser will accept it. Browsers display information about the type of certificate to the user, but most users have no idea how to interpret it. This vulnerability highlights the importance of preventing the associated attacks on the DNS and BGP.

2.4.2 Mitigation

Mitigation of two of these vulnerabilities run directly into human-computer interaction challenges and incentive misalignment. The first is the problem that users ignore warnings when the browser receives an invalid certificate. Certificate management is complicated, and owners of certificates make errors that cause their certificates to be technically invalid. Users get warnings about these certificates, and are asked to decide whether to proceed. Most users do not know how to assess the risk, but choose to proceed anyway because their objective is to complete the task in question. Almost always the invalid certificate is not malicious, and there is no harm to the user. The users are thus trained to ignore these warnings, and when the user receives a warning about a real malicious

certificate, they ignore the warning, thus completely eliminating the protection hypothetically provided by the CA system.

This reality illustrates a deep issue in the design of security systems. Information security is characterized as having three main goals: confidentiality, integrity and availability. The CA system is designed to detect a malformed certificate (thus in principle protecting confidentiality and integrity), by preventing the intended action from completing, thus presenting the user with a complete failure along the dimension of availability. The design does not give the user any strategy to deal with the loss of availability, except to accept the risk to confidentiality and integrity. Users observably care about availability and choose to proceed. Any mechanism that tries to prevent harm by protecting from loss of confidentiality and integrity but makes no effort to protect from loss of availability is an incomplete solution that will have many negative consequences. However, addressing the problem of availability is complicated, and difficult.

The second vulnerability is perhaps even more fundamental. Conceptually, the role of the CA system is to provide a final check that the end point making the connection has reached the intended service. In principle, it should at least turn failures at the lower layers (the DNS and BGP) into clean failures of availability. However, there is a weakness in the way Domain Validation certificates are issued that threatens this protection. To get a DV certificate, the owner of the domain must demonstrate that they have control over the domain, perhaps by installing a file on the web site. However, by hijacking the address of the web site or the address of the authoritative name server, or by penetrating the registry and changing the information about the location of the web site, an attack can deflect traffic intended for that web site to its rogue copy. By instituting this deflection and then requesting a certificate,the program doing the DV validation will perform the test against the web site controlled by the attacker. The attacker will get a certificate that looks valid in all respects.

There are several lessons. One, which is well understood by attackers, is that the most vulnerable step in a security system is during initial setup, when the end points try to make an initial confirmation that they know who the other parties are. Another lesson is that the DV validation was designed to reduce the complexity of getting a certificate to encourage the use of secure connections on the web. A more complex procedure, such as (perhaps) the one used to get an OV certificate, might not be so vulnerable. However, the complexity and cost of that enhanced validation was a barrier to uptake.

The final consequence of this design is that the CA system cannot protect users from all attacks on the DNS and BGP. The security of each depends on the security of the other, which implies a weak and unpredictable outcome. Pragmatically, the best protection is to position the name server and the service itself close to users (to reduce the chance of effective BGP hijacks), and to deploy strong operational practices to reduce the probability of a social engineering attack on the organization owning the domain to prevent theft of their registry/registrar login credentials. A domain owner who has their login credentials stolen is vulnerable to a wide range of malicious consequences.

2.4.3 Primary data

Table 2.8: Certificate Authority: Primary data

Type	Source	Status	Limitations	Uses
Certificates in active use	Censys,Rapid7	Query interface	Relies on active scanning	Track CA market, security posture, adoption of new standards
Certificates logged in certificate transparency logs	Log providers	Current.	No direct evidence of how certificates are used	Detection of misissuance, transparency
Lists of trustworthy and untrustworthy root CAs	CA/Browser	Available	Subjective	Trust decisions guidance

2.4.4 Derived data

Table 2.9: Certificate Authority: Derived data

Type	Source	Status	Limitations	Uses
Data on CAs and root CAs observed in certificate issuance	CT Logs, CCADB	Avail.	Static analysis	research
Which CAs and root CAs show up in queries	Scans	Avail.	Dynamic analysis	Track harms from excluding untrustworthy CAs?

2.5 Denial of Service attacks

Our discussion of Denial of service (DoS) attacks is different in character from the previous sections, which looked at specific systems that constitute the "transport plumbing" of the Internet. Here we discuss a class of attacks that leverage fundamental aspects of

these systems, most notably that routers will make their best efforts to forward all traffic to the destination IP address in the packet, regardless of the purpose of the traffic.

The term DoS covers a wide range of attacks, with different structure and strategy. Given that the focus of the GMI3S project is on security at the Internet layer, we need some criteria to identify DoS attacks that are within the scope of this study. We limit our focus to DoS attacks that either:

- Exploit a feature of an Internet level service as a part of crafting the attack.
- Attack an Internet service using features or vulnerabilities of that service.
- Have an impact on the Internet layer itself.
- Can be detected and/or mitigated at the Internet layer.

State exhaustion

Many attacks that exploit a feature/vulnerability of a service, e.g., SYN-flood, can be characterized as *state exhaustion* attacks. An example is the SYN-flood attack, where an attacker sends TCP SYN packets, each of which induces the allocation of a block of memory: the Transmission Control Block (TCB) associated with an active TCP connection. A flood of SYNs can exhaust the supply of TCBs, preventing the victim from accepting a legitimate request to open a TCP connection.

Any protocol or mechanism where an incoming message causes an allocation of a resource to create a stateful record can be vulnerable to a state exhaustion attack. Every level of the protocol stack has design features that create a vulnerability to a state exhaustion attack, but many such attacks are outside the scope of this study, based on the four criteria above. In particular, state exhaustion attacks often need a much lower rate of attack packets than a brute-force flood, and may have no observable impact on traffic. As an example, the *slow loris* attack bogs down a web server by sending packets, each of which contains a few more bytes of a GET request, and sending them as slowly as possible, but just fast enough that the receiver does not timeout and reclaim the resources holding state information for the request. This state exhaustion attack succeeds by sending slowly, which minimizes impact and visibility of the attack at the Internet level.

Another form of attack exploiting a feature of a service tries to exhaust the processing resource of the service by sending a query that requires significant processing. An example is the *slow drip* attack against a DNS authoritative name server, in which the attacker sends many requests to resolve a different invalid subdomain of a second level domain name (SLD). Recursive resolvers will not have a cached reply to such requests, and will forward them to the appropriate authoritative name server, which may not have the resources to deal with this flood of requests.

Reflection/amplification

Another DDoS technique that can also achieve state exhaustion is *reflection* and *amplification*. In *reflection*, an attacker sends a packet to an intermediate service with a falsified source address, which causes that service to send a reply to that address, which is actually the final victim [144]. In *amplification,* the attacker crafts a request to that intermediate service that triggers a reply that is larger than the request, so that the rate at which bytes arrive at the final victim is larger than the rate at which the attacker must send the stream of requests. Reflection attacks exploit the fact that ISPs only inconsistently implement Source Address Validation, so attackers can send packets with a forged source address. Amplification attacks exploit specific features of network services, which may (or may not) be essential to their normal operation.

Today, the two most exploited network services used in amplification attacks are the DNS and the Network Time Protocol (NTP). A common exploit using DNS queries is to send queries that trigger larger replies. Attackers scan to find names that trigger large replies, and query for these to amplify an attack, or to exhaust the resources of the server.

The NTP request enabling the most amplification is the "get monlist" request, which returns the identity of the last N time requests, which might be very large. This request was not a part of the normal operation of NTP, but rather more of a debugging tool. Mitigations included deprecating vulnerable commands and encouraging NTP server operators to restrict or disable such features [232]. However, similar to DNS, NTP remains vulnerable where best practices are not fully deployed.

Role of Spoofing

Attackers can falsify (spoof) the source IP address in packets, masking their origin and complicating detection and mitigation. Attacks that exploit this technique are called Randomly Spoofed DoS (RSDoS) attacks. Unlike reflection attacks, the malicious traffic is sent *directly* from the attacking infrastructure towards the victim. Although the IETF standards community has long recommended Source Address Validation (SAV) as a best operational practice, its deployment remains limited due to fundamentally misaligned incentives: networks incur deployment costs but do not directly benefit from the protection, as SAV prevents outbound spoofing rather than protecting the deployer.

Mitigating Amplification Attacks

One mitigation strategy aims to eliminate single-packet interactions that attackers exploit for amplification, replacing them with protocols that require an initial handshake [75]. IETF working groups have proposed TCP-based protocols introducing handshakes and encryption, e.g., DNS over TLS (DoT) and DNS over HTTPS (DoH), as well as DNS over QUIC (DoQ) [127] leveraging UDP but enforcing an anti-amplification factor (typically 3x) to limit abuse. While transitioning DNS from UDP to TCP or QUIC reduces

amplification potential, it does not fully eliminate risks associated with spoofed DoS attacks, and comes with tradeoffs in latency and resource requirements.

Engineers have also explored integrating Reverse Path Forwarding (RPF) into anycast, similar to its use in multicast routing, ensuring symmetric routing paths. This symmetry helps validate source addresses, preventing spoofing-based reflection attacks by requiring senders to use their true addresses to complete the handshake. However, these approaches introduce additional round trips, increasing latency, and require services to maintain per-connection state, exposing them to state exhaustion attacks. Stateless handshakes, such as SYN cookies, can mitigate state exhaustion risks by avoiding pre-allocation of state.

Addressing these vulnerabilities requires balancing security, performance, and deployability. While protocol redesign and SAV can significantly reduce attack surfaces, misaligned incentives and performance penalties remain barriers to universal adoption.

2.5.1 Primary data

Table 2.10: DDoS attack inference: primary data

Type	Source	Status	Limitations	Uses
Passive network telemetry	ISP/IXP	Not available	Not available	analyze attacks
Flow monitoring	DDoS mitigation (scrubbing) providers	Not available	Not available	analyze attacks
Network telescopes	UCSD, Merit	Available	Limited visibility	Infer DDoS with random component
Honeypots/sinkholes	AMPpot	Available	Limited visibility	Track attacks
DDoS 2007 data set	CAIDA	available	old	model DDoS traffic

Network telescope collect Internet Background Radiation traffic which consists of (1): *Scanner* traffic intending to discover hosts and services running on those endpoints. Analyzing scanner traffic to telescopes allows one to estimate the spread of malware and botnets, and revel shifting strategies in scanning operations. (2): *Backscatter* is a side effect of DDoS attacks that utilize random source address spoofing. A victim receiving spoofed packets will respond to the spoofed address if it has capacity to do so. The response is received by the network that owns the address space used as the spoofed source

addresses. Capturing traffic to a large number of IP addresses increases the probability to observe some of those (backscatter) responses. Analysts can thus use backscatter to detect and estimate the magnitude of ongoing DDoS attacks, without requiring the victim to disclose the attack. (3) *Traffic from misconfigured or compromised hosts*. Compromised hosts often contribute scanning traffic and spoofed packets incurring backscatter that is captured by network telescopes. Misconfigured hosts, however, may not be malicious but still contribute Internet Background Radiation by sending packets to an unintended destination. The number and pattern of packets are largely influenced by the specific misconfiguration.

2.5.2 Derived data

Table 2.11: DDoS attack inference: derived data

Type	Source	Status	Limitations	Uses
ISP/IXP traffic data	Not available	Not available		
Threat Intelligence Reports	Proprietary		Track trends	

The most prevalent forms of derived data are reports published on industry web sites describing DDoS attack trends. Motivated by the diffuse scope of DDoS research and reporting, we led a multi-stakeholder (joint industry-academic) analysis to seek convergence across the best available macroscopic views of the relative trends in two dominant classes of DOS attacks – direct-path attacks and reflection-amplification attacks [120]. We led a team that analyzed 24 industry reports to extract trends and (in)consistencies across observations by commercial stakeholders in 2022. We then analyzed ten raw and derived data sets spanning industry and academic sources, across four years (2019-2023), to find and explain discrepancies based on data sources, vantage points, methods, and parameters. Our method included a new approach: we shared an aggregated list of DDoS targets with industry players who returned the results of joining this list with their proprietary data sources to reveal gaps in visibility of the academic data sources. We used these data sources to explore an industry-reported relative drop in spoofed reflection-amplification attacks in 2021-2022. Our study illustrated the value, but also the challenge, in independent validation of security-related properties of Internet infrastructure.

Chapter 3

Interdomain (BGP) routing data

Acknowledgments: Contributions in this section by Thomas Alfroy, Ben Du, Thomas Holterback, John Kemper, Thomas Krenc, Hans Kuhn, Matthew Luckie, Cristel Pelsser, Philip Smith. We include text from the study supported by this project "The Next Generation of BGP Data Collection Platforms", authored by a subset of the above [7].

3.1 Limitations of current BGP measurement capabilities

The study of the global Internet infrastructure relies on BGP data collection platforms (RouteViews [200] and RIPE RIS [181]) that maintain BGP peering sessions with network operators who volunteer to share (sometimes portions of) their routing tables. Originally established decades ago to support operational troubleshooting ("How do others reach my network?"), these systems have become a cornerstone for scientific and operational analysis of the Internet. accessible via either real files on disk.

Our evaluation of the current state of BGP data collection revealed that scaling up data collection to keep pace with the growth of the Internet routing system would require an enormous increase in data volume and number of peers. Collecting global BGP data faces a fundamental cost-benefit trade-off. The information-hiding character of BGP requires collecting routes from as many BGP routers, (vantage points or VPs) as possible. But in practice the BGP protocol extensively propagates connectivity messages, leading to highly redundant (along with significant unique) information coming from each peer. The result is a data set with enormous redundancy and yet dangerous visibility gaps. The platforms' policies to store a snapshot of the aggregated data every few hours, as well as every BGP update received in between these snapshots, exacerbates the storage of redundant data. BGP update received in between these snapshots, exacerbates the storage of redundant data. Continued growth of the Internet ($\approx$ 75k ASes [32] and $\approx$ 1M globally announced prefixes) and increasing connectivity between networks further burden data collection and use [1,139]. Users often resort to sampling the data, e.g., using

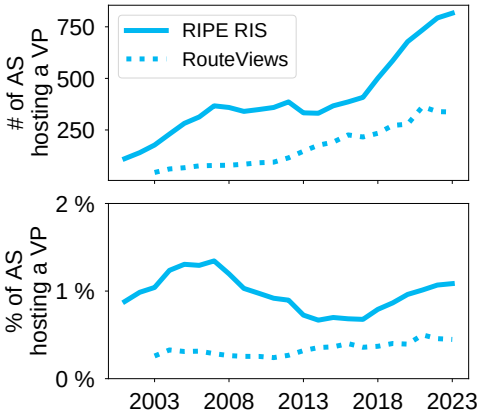

Figure 3.1: Growth in VPs.

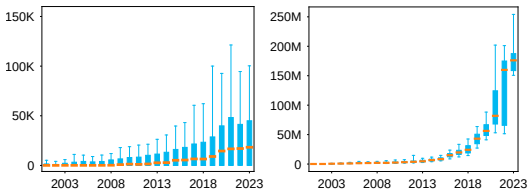

(a) Hourly average # of updates per VP.

(b) Updates per hour among all VPs.

Figure 3.2: RIS and RouteViews collect data volume over time

only a sample of the VPs, neglecting the connectivity uniquely visible to other VPs. The number of BGP routes that the platforms (aggregated) collect every hour has jumped from 82M in 2021 to 160M in 2022 (Figure 3.2). This exponential growth requires a deeper understanding of the data to learn what optimizations make sense. Finally, the manual vetting of new peers also strains platform scalability. The platforms collectively peer with only ≈1% of the observably active ASes on the global Internet. Despite continued addition of peers, RIS and RouteViews's coverage in terms of fraction of ASes they are peering with has remained flat for two decades.

These growing pressures coincide with regulatory concerns about slow progress in deployment of routing security protections [234]. The ensuing public debate has highlighted the importance of these platforms for detecting both accidental and malicious transgressions in the routing system. While significant investment in data collection could accommodate gathering, retention, and sharing orders of magnitude more routing data, resource constraints motivate us to consider a more strategic approach.

3.2 BGP Measurement Research Infrastructure Requirements

Table §3.1 lists a set of high-level system requirements, which we adapted based on feedback from and developments in the research community and industry. Some requirements are in tension with others, and thus designing new infrastructure requires tradeoffs among these design goals:

1. **Expand and optimize coverage: Accommodate 10X more VPs**

   Current BGP data collection strategies do not provide comprehensive visibility of the global routing system; researchers have recently demonstrated how strategically-scoped attacks can evade visibility of current collection systems [170]. Our first

	Goal	Objective
(1)	Expand/optimize coverage	Accommodate 10X more VPs optimally
(2)	Performance	Handle 100x current rates of incoming routing data
(3)	Maintainability	Automated establishment of new peers
(4)	Data Integrity	Handle increasing data rates without loss/error
(5)	Standardization	Support modern protocols/formats, e.g., BMP
(6)	Security	Securing infrastructure against misconfiguration/malice
(7)	Privacy	Respecting privacy of ISPs and vantage points
(8)	Storage	Standard and efficient storage formats and architecture
(9)	Accessibility	Findable, Accessible, Interoperable, Reusable (FAIR)
(10)	Extensibility	Support other measurements from BGP data platforms

Table 3.1: Design goals for BGP monitoring infrastructure

and primary design goal is accommodate a radical increase, e.g., by an order of magnitude, in the number of VPs contributing to public collection systems.

A related and long-standing goal of the BGP measurement platform projects has been to locate new vantage points to maximize visibility of new topology not seen by existing vantage points. Since operation of vantage points has been largely a volunteer practice, the platform operator is not generally in a position to select vantage points at their own discretion, and even if they could one can not generally assume that one can procure a vantage point in the desired location via a commercial service. This need for crowd-sourced volunteer vantage points is a recurring theme in many global Internet measurements.

2. **Performance: Handle 100x current rates of incoming routing data**

   Scaling the number of vantage points will increase data rates. Also, current systems are sometimes beset with orders of magnitude more data from a given vantage points than expected (*noisy peer*), sometimes due to misconfigurations on the VP side. A new BGP measurement platform design must have mechanisms to identify and stop peering with such vantage points, and notify them of the problem.

3. **Maintainability: Automated establishment of new peers** The current RouteViews and RIPE RIS platforms have a significantly manual component to onboarding new peers. A new design should automate this process, with as much intelligent pre-filtering of peers to avoid redundancy as possible.

4. **Data Integrity: Handle increasing data without data loss/errors**

   Sometimes BGP data has errors, and neither of the large platform projects have the resources to monitor the integrity of the data. Data corruption sneaks into MRT

files from sources including storage errors and software bugs. If not detected and understood, data corruption can impair analysis and lead to incorrect results.

5. **Standardization: Support modern protocols/formats, e.g., BMP**

   A modern BGP data collection system would support the recently IETF-standardized BGP Monitoring Protocol (BMP). This requires enhancements to route collector software and expansion of the database architecture, as well as designing and prototyping a new interface to stream telemetry from the global routing system and an API for consuming a steadily growing dataset of routing information.

6. **Security: Securing against misconfiguration and malice**

   BGP update messages reflect configuration changes and link failures for optimal path selection, and the periodic keep-alive. That is, BGP peering sessions are not high-volume in normal operation. However, we have identified recurring patterns of peers that transmit excessive and redundant updates for months, imposing an unnecessary burden on the route collector infrastructure and researchers processing the data. We undertook a study to analyze such noisy peers as part of the design phase [133]. A new measurement infrastructure should include efficient ways to remove this noise without jeopardizing signal in the data.

7. **Privacy: Respecting privacy of ISPs and vantage points**

   BGP data is already aggregated sufficiently to mitigate PII concerns, and those who volunteer peering feeds to the route collector projects are intentionally making their routing data transparent to these systems. The privacy properties of BGP data sharing are well-known, but the newer BGP Monitoring Protocol (BMP) reveals substantially more data, and supporting it will require a privacy impact assessment and ensure the contributing peers are comfortable with it.

8. **Storage: Standard and efficient storage formats and architecture**

   The NSRC RouteViews and RIPE RIS BGP data collection platforms both collect two formats of data. The first is RIB (Routing Information Base) snapshots, which are full routing tables at fixed intervals (every 2 or 6 hours). The second is BGP updates, which are incremental changes in routing (announcements, withdrawals) recorded in near-real time. Data has historically been stored in the MRT (Multi-threaded Routing Toolkit) binary format, a standard for BGP archiving. Files are compressed (e.g., with gzip or bz2) to handle scale — each collector generates gigabytes per day. Given the tremendous growth in BGP data collection volume, efficient storage approaches will be essential.

9. **Accessibility: Findable, Accessible, Interoperable, Reusable (FAIR)**

The FAIR principles of data sharing are critical requirements for NSF-funded research infrastructure, and these principles framed our design contributions. For example, a BGP data collection and sharing infrastructure should support analysis of deployment of the best practices in routing security: adoption of Route Origin Authorization and Route Origin Validation, as well as the emerging ASPA (AS Provider Authorization) [10] or potential alternative proposed schemes [61]. Although adoption of ROAs is straightforward to measure, and many projects do [86, 179], adoption of ROV is more difficult to accurately capture.

10. **Extensibility of platform: Support other measurements from platform**

    We considered the design goal of extensibility and flexibility of the BGP collector platforms to other measurements. In particular, the research community has long expressed the need for vantage points that co-locate data plane (e.g., ping, traceroute, active measurements) and control plane (i.e., BGP route collector) measurements. Such coupled functionality offers both operational and research utility, amplifying active measurement VP coverage to potentially the number of peers interconnecting with the underlying BGP collector. However, other risks can arise if the BGP collectors also execute such active measurements (§3.3.10).

3.3 Proposed Design

The requirements outlined in the previous section led us to initiated a fundamental reconceptualization of public BGP data collection architectures. The result of this effort is a new BGP data collection system design and prototype, *GILL* [7], which will scale to orders of magnitude more vantage points (peers), overcoming a key limitations of the current systems. This redesign of a global BGP monitoring system can overcome some operational scalability limitations with the current RIPE RIS and RouteViews systems. This task was a collaboration with researchers at the University of Strasbourg led by Cristel Pelsser. The lead author of this paper, Cristel's student Thomas Alfroy, completed his 4-month internship at CAIDA in December 2023 We also had one postdoc each in France and UCSD participating in the design. We held weekly meetings to discuss and iterate on the design of this new paradigm for data collection.

We successfully published a early design draft in ACM SIGCOMM HotNets workshop [6], to introduce the data collection method and evaluate its effectiveness in detecting two important phenomena using BGP data: AS-topology mapping and hijacks. Based on feedback on this draft, we We fleshed out the idea and created a prototype culminating in an ACM SIGCOMM submission in January 2024 [7] and associated open source code release. This paper was awarded Best Paper award at ACM SIGCOMM 2024.[1] We

[1] *"Best Paper Awards are presented to authors whose work represents ground-breaking research in their respective areas. By recognizing these select papers for their ingenuity and importance, ACM highlights*

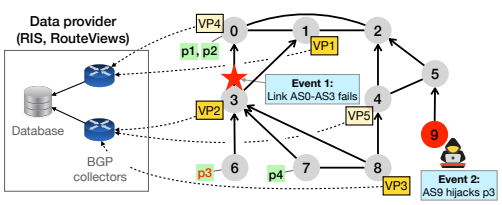

(a) A scenario describing an AS-level topology where a few VPs are deployed and send their routes to a data provider. We focus on two routing events: An outage on the link 0-3 and a hijack launched by *AS9* on p3.

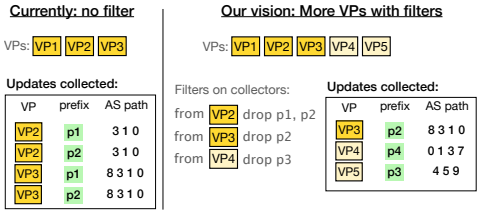

(b) Our overshoot-and-discard approach (right) collects fewer updates but from more useful VPs. These updates better reveal routing events compared to the ones collected with the current approach (left).

Figure 3.3: A scenario that highlights why our overshoot-and-discard approach can be beneficial when collecting BGP updates.

deployed our prototype system where we have invited R&E networks to peer [123]. The peering is entirely automated via a web form.

To support interpretation of the massive amount of collected BGP data, we also designed and prototyped an infrastructure platform that would automatically infer the geolocation semantics of BGP communities. In the meantime, continued pressure on existing infrastructure required streamlining of pipelines and updating of software, which further informed our redesign and in particular the need for automated ingesting of new vantage points. In this section we review how this system and ancillary systems we developed target our specific requirements.

3.3.1 Expand and optimize coverage: scaling number of vantage points

Inspired by tradeoffs made in other disciplines, we proposed a path to accommodating an order of magnitude more data (whether the unit is number of collectors, peers, and/or BGP updates). We proposed a fundamental reconceptualization to how we design public BGP data collection architectures, to use an *overshoot-and-discard* approach that discards redundant data shortly after its collection. Akin to CERN's Large Hadron Collider (LHC) which generates millions of collisions just to see a few interesting particles (e.g., Higgs boson), overshooting BGP data collection will maximize the chance to see interesting routing events, e.g., BGP hijacks. We imagine a world where public BGP data providers could automate deployment of additional VPs, targeting a *moonshot* of peering with one VP in every of the ≈75K ASes participating in the global routing system (even half would be a moonshot!).

Overshooting BGP data collection is only feasible if the system can discard the "less interesting" or "redundant" bits upon acquisition, before it consumes processing or stor-

some of the theoretical and practical innovations that are likely to shape the future of computing." [209]

age resources. In the case of the LHC, fast online algorithms using custom hardware and software discard 99.994% of the likely less interesting collisions [218]. In the case of BGP, predictably redundant properties of BGP data streams [19, 30] suggest that BGP data may be amenable to filtering with minimal loss of information. For example, often several VPs observe BGP routes with similar—sometimes identical—attribute values.

This approach raises the key question: which BGP updates to discard, as missing data inevitably implies loss of information. Defining *redundant* BGP updates depends on context, but we explore a definition that allows us to evaluate our approach in terms of detecting two noteworthy phenomena using BGP data: AS-topology mapping and hijacks. In the case of hijacks, if a hijacked route is either often seen by many VPs [165], or not seen at all e.g., because the hijacker poisoned the AS path to avoid observation [170], then an overshoot-and-discard approach could allow hijack detection systems (e.g., ARTEMIS [206]) to accommodate more VPs (discarding redundant data from them) and thus detect more BGP hijacks.

With European collaborators (in France and Belgium) we designed a system that data providers such as RIS and RouteViews could install to collect BGP routes in an overshoot-and-discard manner. Our design choices were motivated by experimental analyses where we used a probabilistic prediction framework to confirm our hypothesis that BGP routes are highly predictable, and that filtering them carefully leads to minimal loss of information. This approach does not require deploying new hardware nor a software update on the BGP collectors, as filters can simply be configured using *route-maps*, which are already part of the standard BGP implementation used in every router. This approach allows data providers to accommodate more VPs (e.g., using remote peering sessions) without having to upgrade infrastructure.

Executing our overshoot-and-discard raised new research questions such as: how to dynamically update the decisions about which data to discard over time, how to discard routes with minimal infrastructure changes, and how can we ensure that discarding routes does not open new attack vectors [17]. But this approach offers a long-term path toward sustainable scaling of BGP data collection. We showed that a redundancy-aware system consistently improves the accuracy and coverage of studies and tools that rely on BGP data. Our simulations of a scenario where 50% (vs. 2%) of ASes peered with *GILL* tripled the number of peer-to-peer links observed, doubled the number of Internet failures that we could localize, and reduced by 33% the proportion of undetected forged-origin hijacks without processing more data than what RIS and RouteViews do today. We replicated analyses in three studies/tools, in all cases *GILL* improved the accuracy and coverage while processing the same data volume: we inferred more AS relationships (+16%), identified and corrected errors in CAIDA's ASrank dataset, and inferred more forged-origin hijacks (+23%) with $\approx 4\times$ fewer incorrect inferences (i.e., false positives).

Optimizing coverage In the previous work that led to our design, U. Louvain and U. Strasbourg developed a method to best identify the least redundant set of VPs for BGP

and active measurement data collection, building on an recent work [227]. But given the reliance on volunteers to contribute vantage points, our design goal was to avoid the need to be selective about VP selection by scaling to large numbers.

3.3.2 Performance: Handle 100x current rates of incoming routing data

The automation of peering and the ability to algorithmically remove redundant BGP updates fundamentally allow us to scale to orders of magnitude additional peers. Scaling to higher BGP update data rates will be a function of the hardware used in the deployment but the performance bottleneck is likely to be misconfigured routers that are generating a lot of noise. For this reason we undertook a study to support monitoring and repair of such misconfigurations [133].

3.3.3 Maintainability: Automated establishment of new peers

Because our design is targeting a peering session (at least one!) from (theoretically) each AS on the Internet, maintainability of the system was an essential goal. To this end, on-boarding of new VPs is completely automated: operators can connect their BGP routers by submitting a form on the website. GILL automatically configures new peering sessions based on the information provided in the form and new peers are visible on the website within a few minutes. GILL minimizes the risk of fake or misconfigured peering sessions using a two-step authentication scheme: (i) a new participant must send an email to GILL with the AS number provided in the form (ii) once received, GILL cross-checks that the email address of the sender owns that AS according to PeeringDB. In addition to its own peers, GILL also takes as input streams of BGP updates from all RIS VPs using the WebSocket API of RIS Live and all RV VPs using a custom proxy that gathers and gives to GILL the RV data in near real-time.

3.3.4 Data Integrity: Handle increasing data without data loss/errors

The issue of MRT data corruption is orthogonal to the design we have described; we considered this a separate area of study. Unfortunately, collected BGP data is sometimes corrupt. Data corruption sneaks into MRT files (format described in RFC 6396 [20]) from sources including storage errors and software bugs. If not detected and understood, data corruption can impair analysis and lead to incorrect results. To provide tooling to support detection and mitigation of corrupted data, we developed a C program to parse and *explain* corrupted parts of BGP data. This *BGP MRT file explainer tool* fully parses MRT update files, detects and explains errors, and includes a file check mode to quickly determine whether a retrieved MRT file is corrupt, and links to specific sections in RFC's that demonstrate the file's non-compliance with BGP or MRT This open source tool [119] allows both collectors and users of the MRT files to check those files for data corruption before use. The software is fast and efficient enough to do so without impairing the

normal collection pipeline. This tool laid the groundwork for our exploration of more modern compression approaches (§3.3.8).

3.3.5 Standardization: Support modern protocols/formats

Modern BGP Monitoring Protocol (BMP) Neither RouteViews nor RIPE RIS have transitioned to use of BMP for direct peering, due to concerns about sharing more data than a contributing ISP may want to share. The *GILL* algorithms would extend to BMP should providers be willing to share them. In the meantime, to support experimentation with and use of the protocol, RouteViews has created BMP feeds from some collectors, and makes the feeds publicly accessible. The BMP messages follow RFC 7854 [205] and can be consumed in real-time with tools like OpenBMP, GoBMP, BGPstream, or BGPKit. This BMP support allows for capture of not only standard MRT-formatted BGP data but also connection events, and peer-level details. The RouteViews team originally adapted an existing open source BMP implementation (GoBMP), but during the Design Phase they designed and wrote their own BMP parser, Bimper. Bimper is a high-performance message processor that receives BMP routing data and forwards it to Kafka for downstream analysis and storage. It provides real-time monitoring of BGP routing events with Prometheus integration for operational visibility of metrics. The system includes bimperctl, a control utility for managing and monitoring bimper instances, allowing administrators to interact with the service and view connection status information, and manage router connections. This tool would serve as the basis of a future implementation of this infarstructure component.

3.3.6 Security: Securing against misconfiguration and malice

While the BGP data collection platform operators are aware of *noisy peers*, previous work had not systematically quantified its scope and characteristics. In our design phase we examined this phenomenon, analyzing more than 80B BGP updates, finding common attributes and implications for Routeviews' collectors and researchers. In our analysis we found fewer than 1% of peers and sessions were responsible for most noise [133].

3.3.7 Privacy: Respecting privacy of ISPs and vantage points

As described in §3.2, BGP data is aggregated sufficiently to avoid PII concerns. Peering vantage points are intentionally making their routing data transparent to these systems. However, we did find that privacy concerns were/are a key obstacle in moving toward the more modern BMP (BGP Monitoring Protocol) [205]. Specifically, when ISPs establish a BMP session to a collector, they are exporting three types of data that go beyond normal BGP peering:

- Full RIB snapshots (all BGP routes, not just best paths announced to neighbors).

- BGP updates (announcements/withdrawals) from your routers.
- Metadata such as peer IPs, ASNs, next hops, and attributes (communities, MED, local-pref, etc.).

Such data reveals the ISPs internal routing view; indeed this was the objective of the BMP design. But exposure of internal topology, including private peers and backup providers, is typically more than ISPs are comfortable publishing. Similarly, BGP communities often encode operational information (e.g., geolocation, customer IDs, or traffic-engineering hints), disclosure of which may reveal business relationships or infrastructure details. BMP messages may also include router IDs, peer IP addresses, and private ASNs that could be useful to attackers mapping an ISP's infrastructure. Although BMP is used for proprietary monitoring systems, it seems unlikely to see adoption by publicly BGP data sharing projects. After this privacy impact assessment we did not further invest resources in BMP data collection design and prototyping, beyond what RouteViews is already supporting 3.3.5

3.3.8 Storage: Standard and efficient storage formats and architecture

In the second year of the GMI Design project, the RouteViews teams established a collaboration with Google to archive all historical RouteViews data in Google Cloud. Google temporarily made this data available via BigQuery to a small group of GMI researchers for evaluation purposes. e.g,. to provide feedback on how to improve the schema used to ingest the data into BigQuery. This was a volunteer effort by Google and personnel changes left Google unable to participate in accommodating schema changes recommended by the working group. One lesson we can definitely take from the project is the unreliability of volunteer arrangements for significant storage/processing effort as part of the Implementation Phase of the project.

MRT storage format With the *GILL* design we prototyped, we solved the issue of scaling storage as we expanded data collection an order of magnitude by not supporting historical archiving as part of the design. That is, our default design is as a streaming infrastructure, where storage decisions are up to individual researchers who can decide what is important to their research questions. This is a radical paradigm shift but may be appropriate for the radical increase in data generated by an order of magnitude increase in coverage. However, some researchers will want a longitudinal archive, so we investigated another option for scaling storage for existing BGP data platforms. We investigated the possibility of re-architecting the 20+-year MRT data format to encode redundant segments of MRT files. However, early test and evaluation indicated the increased efficiency of the few format was no sufficient to warrant proposing that the community change twenty years of tooling that relies on MRT formats. We conclude that the existing MRT format is the best option for a BGP collection system in the near to medium-term future.

3.3.9 Access: Findable, Accessible, Interoperable, Reusable (FAIR)

Both RouteViews and CAIDA's BGP-data analytics services (including APIs) such as AS Rank, BGPStream, and (MSRI-funded) BGP2Go exemplify the FAIR principles by making global routing data broadly accessible and reusable to the research community. Both projects ensure data is findable by maintaining indexed archives of routing table snapshots and updates, organized by time and collector, and further enriched by metadata and search tools that allow filtering by prefix, ASN, or community. The data is accessible through open repositories and APIs, minimizing technical and policy barriers to entry. By adopting standard formats such as MRT and providing open-source parsing libraries, the platforms promote interoperability, allowing researchers to integrate these datasets seamlessly into diverse analytical toolchains. Finally, they enable reuse by offering long-term archives, comprehensive documentation, and permissive access policies that support reproducibility and comparative research across decades of Internet routing history.

BGPstream transition to BGPkit For years CAIDA has supported and maintained BGPStream, providing the research and operations community with a platform for analyzing BGP data. However, much of the core functionality that BGPStream offered is now available through BGPKit, an actively maintained and modernized toolkit developed and supported as an open source project by Mingwei Zhang (Cloudflare). Given the overlap in capabilities and the fact that BGPkit is robustly maintained, CAIDA is minimizing continued work on sustaining BGPStream in order to focus efforts on other priorities.

BGP2Go to retrieve specific MRT files from repositories In respond to community input and feedback, we designed a specification and prototype for a new platform to help researchers and network operators quickly locate and access the specific MRT files they need—rather than downloading enormous datasets indiscriminately. BGP collectors (like RouteViews) generate vast amounts of routing update data. Without BGP2GO, users must often fetch and sift through gigabytes or terabytes of files to find just the subset relevant to their query. BGP2GO indexes MRT updates against key identifiers like prefixes, Autonomous System Numbers (ASNs), and BGP communities, enabling targeted searches. Users can filter MRT files by resource, time frame, and collector, then either download or stream just the relevant files—perfectly tailored to their needs. After identifying relevant files, BGP2GO offers integrations like BGPStream or bgpreader, allowing live filtering and processing without complete downloads.

Focus on security research. Our three main proposed components – *GILL*, BGPstream/BGPKit, and BGP2Go – help researchers and operators identify invalid BGP announcements such as route hijacks, leaks, or RPKI-invalid routes over time. *GILL*'s removal of redundancy in routing update data makes it easier to spot anomalies like suspicious origin AS changes or sudden, short-lived prefixes. BGPstream and BGPKit provide

an API and SDK to explore historical BGP MRT data, enabling retrospective validation of whether announcements matched RPKI or IRR records at the time. Researchers can track patterns of invalid behavior across networks, identifying recurring offenders or structural weaknesses. Finally, BGP2Go allows targeted inquiry – supporting queries of a specific prefix or ASN to check if an announcement matched RPKI or IRR records at the time. Together, these tools will help the community move from raw, overwhelming BGP feeds toward practical, FAIR-aligned tools that highlight, validate, and contextualize invalid BGP announcements across both real-time and historical perspectives.

3.3.10 Extensibility of platform: Support new measurements

The BGP route collector projects have always focused on the single function of collecting BGP data, which itself was a sufficient challenge. But one extension opportunity has come up repeatedly for years in the research community: vantage points that co-locate data plane (e.g., ping, traceroute, active measurements) and control plane (i.e., BGP route collector) measurements. Such coupled functionality offers both operational and research utility, amplifying active measurement VP coverage to potentially the number of peers interconnecting with the underlying BGP collector. However, other risks can arise if the BGP collectors also execute such active measurements.

As part of a related NSF-funded CCRI effort, we experimented with leveraging existing BGP collector infrastructure to support active network measurements (ping and traceroute), using CAIDA's measurement software *Scamper*. Unfortunately, neither collection platform (RIPE RIS nor RouteViews) considered it prudent to take on the policy development required to allow this new type of measurements on their infrastructure. However, this experience informed our design specification. In particular the focus on independent active measurement infrastructure (§4), so we describe the design and prototyping work we performed to explore this functionality.

Background A RouteViews collector has a transit interface with a globally-routed address configured (192.0.32.10) which the collector uses to transmit archived routing tables to the University of Oregon, as well as an interface in an IXP peering LAN (192.0.2.1) which it uses to establish BGP sessions with members at the exchange (e.g., AS A with 192.0.2.2). To perform an active measurement (ping or traceroute) through AS A in Figure 3.4, the measurement software will form an Ethernet frame with the member peer's MAC address (bb:bb:bb:bb:bb:bb, learned through an ARP request) as the *destination* MAC address, and use the collector's MAC address (aa:aa:aa:aa:aa:aa) as the source. Inside the Ethernet frame is an IP packet, with the collector's transit IP address (192.0.32.10) as the source address, so that the RouteViews collector can receive responses from the Internet. Note that none of the source (IP or MAC) addresses are spoofed; they are assigned to the collector. The destination IP address is set to the measurement target IP address, which will vary according to the measurement objective.

This approach is similar to that used in the PEERING testbed [202]. We prototyped this measurement approach on the RouteViews collectors using *Scamper* [157] (§4.3.10).

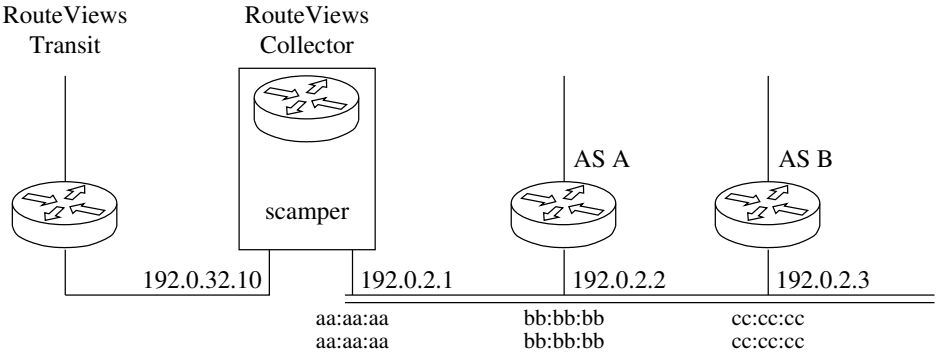

Figure 3.4: Architecture of Scamper use on RouteViews collectors. A collector has a transit interface with a globally-routed address configured, as well as an interface in an IXP peering LAN. We can direct traceroute and ping measurements through a specific IXP member by sending them Ethernet frames directed to the member.

In our evaluation scenario a machine at CAIDA coordinated active measurements from RV peers. This machine used Scamper's remote control mechanism to request measurements and receive results from RouteViews collectors. No user of this service need (or would receive) a login to any RouteViews collector or other infrastructure. We limited Scamper's probing rate on the collectors to a maximum of 100 pps – approximately 6KB/s through the exchange fabric. The responses do not arrive through the exchange; they arrive at the collector through its transit interface, i.e., from the Internet.

The technical component of this capability is straightforward. The more challenging piece is the policy framework. Participating BGP peers may not approve of this enhancement, so we drafted a set of guidelines for how Scamper can responsibly conduct data-plane measurements from RouteViews collectors through RouteViews peers, and implications for both collectors and peers. In addition to these guard rails, a new system will require functionality to allow RouteViews peers to opt-in (for existing peers) or opt-out (for new peers) of use of their peering router for such measurements.

However, there is no current AUP governing any of the RouteViews peering sessions, and incomplete contact information for the peers prevents developing comprehensive AUPs. So the challenge is how to facilitate peers opt-in to enable such measurements on the collectors through willing peers. The RouteViews project decided that at the time there was no path forward for integrating active measurements into the RouteViews infrastructure, but we include it in our design specification as a valuable component of other potential new BGP measurement infrastructures. More realistically, we recognize that most active measurement capability will operate in a separate infrastructure, which informed our focus on active measurement (§4) for the subsequent Implementation Phase.

Chapter 4

Active measurement

Acknowledgments: Contributions in this section by Dan Andersen, Bill Herrin, Paul Hick, Brendon Jones, Matthew Luckie. We have incorporated some text from collaborations with these authors, including: [152].)

Network operators and researchers often require the ability to conduct active measurements of networks from a specific location in order to understand some property of the network. However, 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. 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.

Over the past two years, in consultation with members of the active measurement community, CAIDA designed a next-generation active measurement platform, leveraging experience and residual assets from the existing Archipelago (Ark) platform. In an effort to make a platform that is easier to use for external measurement researchers, while also providing important access control, we designed and implemented a researcher development environment that allows for complex, distributed, and reactive measurements built on a well-defined set of measurement primitives, from a set of distributed VPs.

4.1 Limitations of current active measurement capabilities

Figure 4.1 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

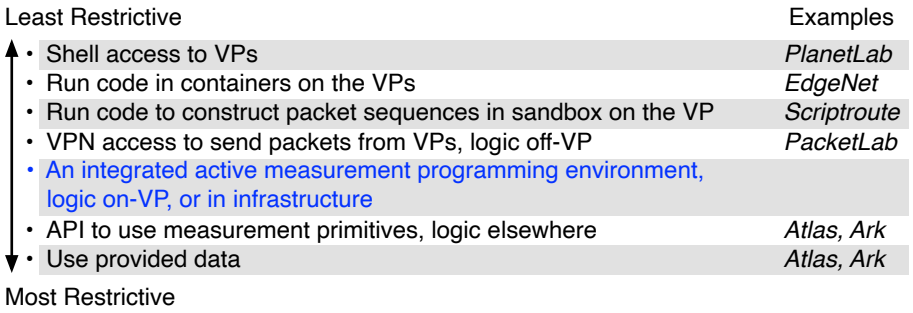

Figure 4.1: Spectrum of active measurement infrastructures.

what the researcher does, and thus assumes significant risk. A step removed from this is VPN-like access: the VP acts as a simple packet forwarder, allowing a researcher to use the node without providing shell access. These solutions allow researchers to craft specific packet sequences that allow for 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.

Previous active measurement research infrastructures

Although the field of active Internet measurement has expanded for three decades, the availability of active measurement infrastructure to support research is scant. As with other types of research infrastructure, without a source of dedicated funding, Internet measurement infrastructures typically do not survive beyond a grant cycle or two [138, 167, 189, 190, 207]. Several of these early (now defunct) active measurement platforms such as Skitter [126], Surveyor [138], AMP [167], NIMI [189, 190], and DIMES [207] provided (primarily simple traceroute topology) data for use by the research community. We do not discuss M-Lab because it provides server-side facilities for client-server active measurements [116], or NLNOG RING because the infrastructure requires a user be an operator at an AS with a participating VP [182].

PlanetLab: In 2002, Peterson *et al.* began deploying PlanetLab, a platform for deploying and managing distributed network services [191]. PlanetLab operators distributed *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. despite its policy restrictions against probing the Internet. At its peak, PlanetLab had systems in ≈700 organizations. PlanetLab shut down in 2020. Subse-

quent attempts to sustain flexible extensible measurement infrastructures have not gained significant traction [67, 116, 147].

Scriptroute: Released in 2003, Scriptroute [216] provided (1) a set of distributed VPs, and (2) a sandboxed scripting environment so that unvetted users could use them. An application programmer wrote Ruby scripts that embedded logic for sending packets and processing received packets. Users found VPs with DNS queries, and uploaded scripts to VPs of interest via an HTTP API, requiring that each VP have a publicly reachable IP address. Each VP's Scriptroute instance protected the VP hosting site from accidental or malicious transgressions by running user scripts in distinct sandboxes that limited the resources and system capabilities available to each script, enforced policy around the types and frequency of packets that each script could send, and matched probes with responses so that each script could only observe responses to packets it sent.

Ark: In 2007, CAIDA began operating the Ark infrastructure to perform comprehensive global topology mapping as well as support third-party experiments on the platform. CAIDA's Ark platform has been one of the most important resources for active Internet measurement for years. The platform consists of a diverse collection of vantage points, including x86 rack-mount systems, Raspberry Pis (versions 2–4), as well as virtual machines and containers. To coordinate measurements across these nodes, CAIDA developed a distributed tuple-space system, Marinda, implemented in Ruby. Marinda allowed CAIDA to orchestrate large-scale global measurements and was central to producing data that has fueled more than 1,200 external publications [101]. However, despite its central role in CAIDA's own operations, no external researchers ever published using Marinda directly to coordinate their own measurements. Researchers could deploy vetted measurement software onto Ark nodes, but this was not straightforward. The infrastructure was highly heterogeneous: a mix of operating systems from different vendors and vintages, combined with multiple CPU architectures. Software installation and maintenance were handled almost entirely manually, which made deployment and upkeep cumbersome and error-prone. Adding or updating a measurement tool sometimes meant dealing with broken dependencies, and deploying a new monitor required a lot of manual setup. As the hardware aged, maintaining the system became more costly and time-consuming. By the time of this GMI3S design project, most Ark nodes were nearing end-of-life. For example, the i386 architecture had been demoted to a tier-2 platform with the release of FreeBSD 13, making continued support increasingly difficult.

RIPE Atlas: operated by RIPE NCC since 2010, Atlas is currently the largest deployed operational active measurement infrastructure, with 13K+ vantage points (VPs) in 4K+ IPv4 (2K IPv6) ASes as of June 2025, representing 5% of routed ASes [197]. The primary mission of RIPE Atlas is diagnosis and troubleshooting to support its operator community, not scientific research. Atlas consists of different types of VPs. The majority (7,697) are small single-board computers with limited CPU, storage, and memory. Atlas also consists of more-powerful *anchors* (794), as well as software VPs (3,620) using the same software as deployed on the single-board computers. Factors in Atlas' success in-

clude (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 then 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. It is challenging to deploy reactive measurements, as it "generally takes a few minutes to get the result of a measurement" [73] and most VPs send 4–12 packets per second [73].

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 [182]. 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.

PacketLab: Proposed in 2017, PacketLab [148] provides a packet-oriented interface for sending and receiving packets via a distributed set of VPs, similar in goal to Scriptroute. PacketLab's architecture includes (1) a controller that provides centralized access to a set of VPs, (2) packet-sending policy enforced through BPF filters, and (3) authentication of measurements through cryptographic certificates. In recent years, the PacketLab authors reported prototype deployment on EdgeNet [238, 239] and implementations of ping, traceroute, DNS lookups, and HTTP requests. A PacketLab implementation of a protocol that uses TLS (such as HTTPS) would be complex, requiring the implementer to marshal packets through a TLS library off the VP.

EdgeNet: In 2017, researchers at Sorbonne began building a software-only platform for deploying distributed network services, motivated by the observation that maintaining and debugging hardware required six full-time people at PlanetLab [27]. Volunteers contribute *software* (VM) nodes to EdgeNet. EdgeNet operators seek to manage the nodes with off-the-shelf software, such as Kubernetes, rather than customize the operating system. Researchers use these software nodes by publishing Docker containers that EdgeNet can deploy on the nodes [67]. EdgeNet's current status is not clear from the website at edgenet.org.

PerfSonar: perfSONAR [229] is the performance Service-Oriented Network monitoring Architecture, a network measurement software platform designed to provide federated measurement of performance across network paths. The research and education networking community has deployed thousands of perfSONAR instances worldwide, but

	Goal	Objective
(1)	Expand/optimize coverage	Accommodate 10X increase in VPs optimally
(2)	Performance	Maximize use while minimizing impact on hosting site
(3)	Maintainability	Automate provisioning and maintenance
(4)	Data Integrity	Monitor workflows for corrupt measurement output
(5)	Standardization	Support modern protocols/formats
(6)	Security	Securing infrastructure against misconfiguration/malice
(7)	Privacy	Respecting privacy of ISPs and vantage points
(8)	Storage	Standard and efficient storage formats and architecture
(9)	Access	Findable, Accessible, Interoperable, Reusable (FAIR)
(10)	Flexibility/Extensibility	Support new measurements

Table 4.1: Design goals for active monitoring infrastructure (similar to BGP monitoring goals)

perfSONAR with a very specific mission: to monitor and diagnose performance problems in research and education (R&E) networks. perfSONAR provides a uniform interface that allows for scheduling of measurements, storage of data in uniform formats, and scalable methods to retrieve data and generate visualizations. The software supports measurement of a few standard network metrics (latency, packet loss, throughput, jitter, traceroute). It is an operational monitoring system and was never intended to support research.

4.2 Active Internet Measurement Infrastructure Requirements

Table §4.1 lists the high-level system requirements for active Internet measurement research infrastructure. These requirements are largely in common with the BGP measurement infrastructure requirements. As with those we derived these requirements based on feedback from and developments in the research community and industry.

1. **Expand Coverage: Scaling up number of vantage points (VPs)**

   Similar to BGP measurement, the platform operator does not generally get to choose vantage points; instead the operator must rely on volunteers donating vantage points to the platform. To achieve this goal our focus was to reduce the cost of deploying and maintaining vantage points, including supporting hardware as well as containerized software nodes. A key lesson from previous work has been the value of using a single operating system that allows for scalable system administration.

2. **Performance: Accommodate complex reactive measurements**

The delay between measurement and result should be small, so that researchers can build complex reactive measurements.

3. **Maintainability: Automate provisioning and maintenance**

   Experience from previous measurement research infrastructures has conveyed the importance of automation for sustainability. Manual registration of node configuration in our legacy node database (hostname, location, point-of-contact, SSH key) was difficult to scale and a barrier to node deployment. Key to future scalable maintenance is a pull-based model, where we publish packages that contain measurement software, and the vantage points self-update their own packages when they are operational without system administrator intervention. Cost-effective maintainability requires using off-the-shelf software when possible, and consistently packaging any custom software that we do write.

4. **Data Integrity: Monitor workflows for corrupt measurement output**

   Another lesson from previous experimental research infrastructures has been the need to check for unexpected data in measurement results, e.g., the measurements should be collected but are not appearing in the archive.

5. **Standardization: Support modern protocols/formats**

   The platform should use as many standardized measurement and system components as possible to support maintainability and interoperability. Specifically, we will implement standards-compliant versions of raw measurement protocol primitives, giving researchers the essential building blocks to create more complex measurements. The measurement libraries themselves should be container-ready, available in packaged form, and extensible through interfaces for adding new primitives. To handle the data pipeline from measurement nodes to the central server, the specification should adopt a modern, widely used message broker, e.g., Kafka. The project should also include a standard Memorandum of Cooperation with site hosts, enabling them to opt in or out of selected measurements.

6. **Security: Protection against misconfiguration and malice**

   An obvious risk is that an experimenter could misuse a primitive in a way that causes harm to VP hosting providers, measurement targets, or other experiments on the platform. For example, a host in a country that censors HTTP, DNS, or TLS could be harmed by measurement traffic that contains keywords that trigger the censor. Similarly, spoofed packets can be used both to test source address validation (SAV) deployment, and for denial of service attacks. The design must have mechanisms to prevent or limit measurements that could be problematic for the site host, including allowing site hosts to opt-out of measurements that they do not want to support.

7. **Privacy: Respecting privacy of ISPs and vantage points**

   Active measurement is generally not a significant privacy concern because it involves sending custom probe traffic rather than collecting sensitive user data. These probes are typically limited in scope, contain no personal identifiers, and designed to measure network performance or topology rather than user behavior. However, inferences about critical infrastructure can yield sensitive information, and a measurement system should control access to data to protect against its misuse.

8. **Storage: Standard and efficient storage formats and architecture**

   One observation from maintaining the Ark infrastructure was premature wear on the nodes' SSD cards from continuous writing of measurement data, so one design goal was to minimize write activity to SD cards.

9. **Flexibility and Extensibility: Support new measurements**

   Using off-the-shelf and easily deployable components will maximize avenues of future deployment. We used components available in Scamper [156] to provide measurement capabilities on VPs, and to support centrally scheduling and receiving of measurements on VPs. Scamper is interoperable, as it builds and runs on a diverse set of operating systems and architectures, 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. We have publicly released our implementation [157] and documentation [104], so that the Internet measurement and operations communities can pursue such extensions. As an example, in 2025 we supported a set of German researchers who wanted to create a new primitive for the platform to support measurement of the new QUIC protocol [114].

4.3 Proposed Design

We designed and prototyped a solution that lies in the middle of the spectrum described in figure 4.1). In particular, we designed and developed a Python-based integrated active measurement programming environment that exposes both a set of distributed VPs, and a set of useful measurement primitives from which to build sophisticated measurement tools. This Python library [104] acts as a bridge to the measurement tools on each Ark node. Figure 4.2 illustrates our architecture; peer-reviewed details of the major components are in [152].

The underlying network consists of globally distributed 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 supports heterogeneous

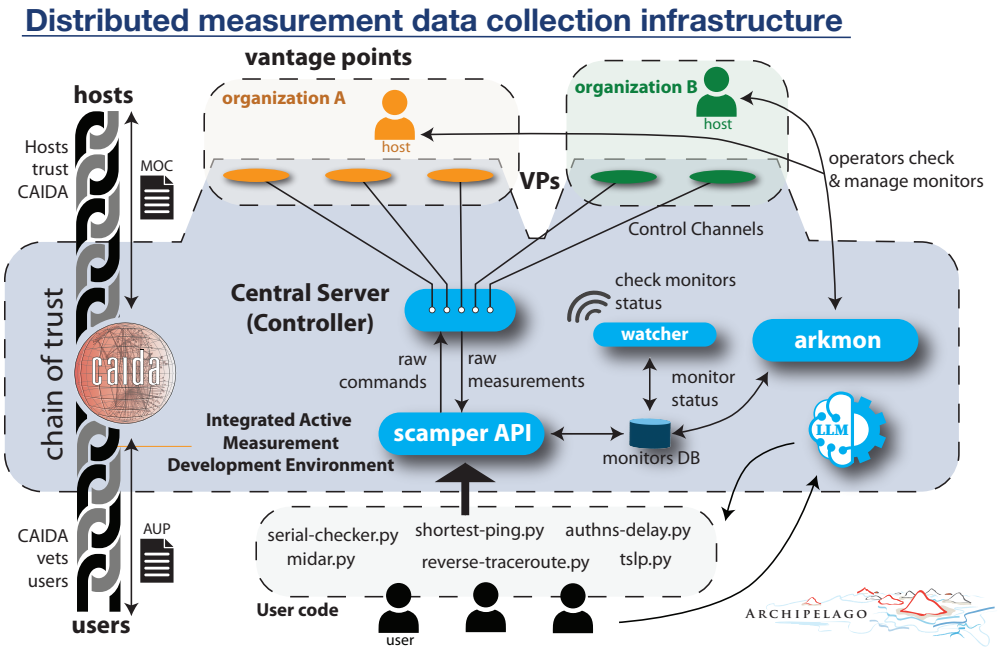

Figure 4.2: *GLIMPSE architecture: VPs connect to a central controller. Scripts access primitives on VPs using an integrated active measurement development environment deployed on, or next to, the controller. The platform is designed to be responsive, interoperable, extensible, and easy to use.*

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 provides 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 have proposed to integrate 300 existing Ark VPs from Ark into this new infrastructure as part of the Implementation Phase of this MSRI effort. Additionally, we proposed to acquire, configure, and deploy 200 new VPs per year in strategically significant locations. Although we proposed to deploy O(1000) VPs in the Implementation 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 intend to expand VP deployment to commercial cloud and mobile vantage points. Each vantage point operates 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 in-

teract with query-response services such as NTP and SNMP, alias resolution methods for identifying which IP addresses belong to the same router, TBIT [168] to infer properties of a remote TCP stack, and packet capture to selectively record specific packets.

A robust measurement software pipeline will 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. Capabilities of the platform can expand to incorporate new measurements in response to evolving researcher needs.

A front-end portal and back-end database will support management of 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 allows hosts to configure VPs they host, and select measurements they are willing to support. The back-end includes a Postgres database and Python library to manage database access across subsystems. This component automates the creation of Debian packages to streamline on-VP software management and ensure consistent deployment and updates across all field vantage points.

The key benefits of this environment to *researchers* are that (1) the environment provides 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 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 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.

In January 2024, we released the launched this measurement environment, including releasing underlying source code, allowing researchers to develop and test their measurements locally before trying them on Ark. We started a series of blog posts [158, 159] explaining how to use these new features, with step-by-step guidance for researchers to make the most out of the improved Ark platform.

4.3.1 Expand Coverage: Scaling up number of vantage points (VPs)

We designed our proposed architecture with a primary goal of lowering the barrier to deploying vantage points, which included support for both hardware and containerized software deployments to accommodate constraints of different hosts. To promote and incentivize deployment, the environment allows platform operators to accurately describe the types of measurements the VPs will do. The environment has 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. We communicate these risks as part of the measurement node on-boarding process.

For hardware VP deployments, we evaluated several different hardware options, but due to the supply chain abundance of Raspberry Pi's, and their prevalence in other research cyberinfrastructure, we chose to stick with Raspberry Pi's for hardware-based Ark nodes in the field. For software deployments, we invested considerable effort to create lean, maintainable software containers to facilitate vantage point expansion (in §4.3.3).

Our experience with prototyping our newly designed architecture and components provided evidence of success in lowering the barriers to expand vantage point coverage. During the life of this GMI3S Design 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 March 2025. The expansion included a mix of Raspberry Pis, virtual machines, and (103) containers across multiple continents.

4.3.2 Performance: Accommodate complex reactive measurements

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. Centralized access to the VP controller interface allows users to run code as close as possible to the VP controller to minimize delay (Figure 4.2).

4.3.3 Maintainability: Automate provisioning and maintenance

Perhaps our most significant goal in the design of a new active measurement system was to modernize how Ark monitors are deployed, configured, and managed. In the past, we had to install software manually across a diverse set of operating systems and hardware types, which made scaling and maintenance increasingly difficult. We redesigned the platform around five components: automated packaging, containerization, automated initialization, a new certificate authority and monitor management.

Ark Docker images We will publish images to the Docker Hub Container Image Library where they are available to users. Both Podman and Docker can download and use images from Docker Hub.[1] The Ark container images mimic the behavior of a physical Ark node, which allows reuse of existing software for performing measurements. Containers run the same measurement software and have the same capabilities as other nodes, and should appear indistinguishable to an outside observer.

It is not feasible to manually customize and publish images with complex site-specific configuration, so the container environment and configuration server provides this information instead. The basic network configuration provided is enough to get Internet access

[1] `https://hub.docker.com/r/caida/ark`.

so that automated systems can reach out to the configuration server, certificate authority, and test controllers.

Containerization Our new platform design will have all software packaged, with files in standard locations, and with packages built automatically from source code repositories. We support a Debian package repository to host the packages (with Ansible scripts to configure them), and GitLab CI scripts to automate building/publishing packages for new releases of software from the CAIDA GitLab instance. We have added these package building and publishing scripts to other repositories we maintain.

Automated Initialization To make deployment truly scalable, we added an automated initialization mechanism. When a container-based node starts, it uses its initial IP address (if known) to contact the primary Ark server. The server checks this against the database, and if valid, issues a JSON Web Token (JWT). Along with configuration details such as hostname and SSH port, the JWT allows the node to obtain an X.509 certificate. This certificate is then used to generate SSH credentials for secure access through the SSH proxy. With this approach, new monitors can configure themselves almost entirely automatically, greatly reducing the need for manual intervention.

Arkmon: Monitor Management Once deployed and initialized, monitors are managed through Arkmon, our new management and monitoring platform. The previous (pre GMI3S project) database for monitor metadata offered a web-based UI and a limited GraphQL API. However, due to insufficient security features, access to this database 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 GraphQL API led approximately half of the scripts interacting with the database to bypass the API entirely and directly access its underlying SQLite database. This direct database access from multiple independent code bases increased the likelihood of bugs and inconsistencies arising from divergent implementations. To modernize the infrastructure, we initiated the design and prototype development of Arkmon with three design requirements: enable secure external access to monitor data; a unified, shared code base for all database interactions; a flexible and user-friendly UI for both hosts and administrators.

We migrated the existing 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. The system includes a web UI that allows hosts to monitor and manage the status of their Ark monitors, and a RESTful service that supports communication between the UI and various Ark data acquisition and processing scripts.

The Arkmon UI will give hosts direct visibility into their monitors. Through a web-

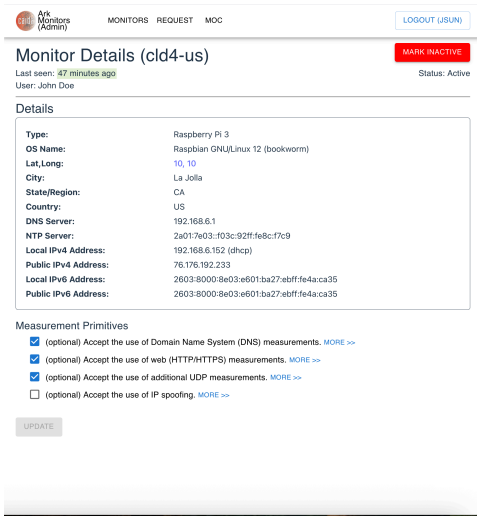

Figure 4.3: Arkmon UI Dashboard: monitor details.

Figure 4.4: Arkmon UI Dashboard: monitors overview.

based dashboard, hosts will be able to check the health of their monitor, troubleshoot issues, update contact or configuration details, and even volunteer to host new monitors. For administrators, the UI will also support tools to manage operator permissions, handle support tickets, and configure monitor parameters such as DNS and NTP servers. A key design principle is to make the experience accessible to monitor hosts, while still providing enough depth for CAIDA staff to maintain and extend the system.

Figure 4.3 and 4.4 illustrate two screenshots of the Arkmon UI dashboard. Hosts will be able to view and update monitor information. The UI will also include a simple workflow for volunteers to request new monitors, and operator/admin tools for managing nodes (Figure 4.4). It will allow configuration of the set of measurement primitives available on each monitor. Implemented as a RESTful service, the Arkmon API provides a secure and standardized channel for communication between the UI and the various Ark data acquisition and processing scripts. The API supports a wide range of operations: submitting new monitor requests (with details such as city, hardware type, latitude/longitude, and mailing address), updating requests, retrieving monitor status, and managing support tickets. This new Arkmon API is secured using Keycloak authentication and will be deployed outside the CAIDA firewall.

The API enables a user to request deployment of a new Ark node by submitting geographic, organizational, and mailing details, which can then be reviewed and approved by administrators. It also allows querying of existing monitors, exposing metadata such

(1)	HTTP queries
(2)	DNS queries
(3)	custom TCP packets
(4)	sourcing of packets with spoofed addresses
(5)	BGP announcements
(6)	topology measurements (ping, traceroute, alias resolution)
(7)	TLS handshakes and STARTTLS
(8)	packet capture of unsolicited traffic received by the VP
(9)	bandwidth measurements
(10)	zgrab-style banner grabs
(11)	arbitrary forwarding of packets delivered to the VP by a service

Table 4.2: Consensus list of initial measurement primitives (May 2023 AIMS consensus)

as node ID, operator, hardware type, organizational affiliation, and configuration state.

4.3.4 Data Integrity: Detect corrupt measurement output

The system will use the Scamper tool for sending probes, which implements probing according to standard protocols, and addresses load balancing and other issues that can distort traditional traceroute results. For example, we will need functionality to detect networks that block traceroute probes, networks that claim to have IPv6 but do not, and patterns of responses that suggest a loop early in the path, which can prevent additional packets from transmitting. Our libraries will adjust parameters to catch phenomena that are causing nodes to collect less data than expected, and compensate for observed idiosyncrasies.

4.3.5 Standardization: Support modern protocols/formats

We have implemented standards-compliant implementations of many raw measurement primitives (Table 4.2) which provide researchers the building blocks to build more complex measurements. For example, researchers could use DNS queries that obtain sets of authoritative nameservers in the resolution path of a domain name that could then be followed with ping and traceroute measurements to identify how close those nameservers are to the resolver, and the path that the resolver takes to the nameserver. These raw measurements overlap; a user could construct their own TCP-based traceroute with custom TCP packets, for example. However, our job is to support the types of measurements that researchers typically use, and provide a degree of flexibility in how the researcher obtains active measurements that answer their questions. Our specification includes a Memorandum of Cooperation with site hosts who can opt-in or opt-out of select measurements.

Our specification relies on components available in scamper [156] to provide mea-

surement capabilities on VPs, and to support centrally scheduling and receiving of measurements on VPs. Scamper is interoperable, as it builds and runs on a diverse set of operating systems and architectures, 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 new measurement primitives.

To support the measurement data pipeline from active measurement nodes to the centralized server, the new specification will use Kafka, which is a modern, actively maintained message broker used already inside CAIDA, as well as by many other organizations. The Spoofer receiver and collector software will use Kafka for communicating messages from probes to the central server.

4.3.6 Security: Protection against misconfiguration and malice

The system must allow measurements only 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.

We also specified a new reverse proxy system to communicate with remote CAIDA-operated nodes behind NATs from our centralized back end (i.e,. the reverse direction from the above paragraph). We designed a completely new approach using secure shell (SSH) port forwarding to establish the reverse proxy. This software is now fully Debian-packaged for deployment on the Raspberry Pi and Ubuntu nodes [135].

Modernized and Automated Certificate Management

To secure authentication across the expanding Ark infrastructure, we designed and deployed a new certificate authority (CA). We investigated which certificate system to adopt, how to transition older nodes into the new framework, and how to address compatibility issues. We ultimately chose the open-source stepCA[2] software to operate a private CA for Ark. Our implementation streamlines 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. We tested certificate issuance and renewal both locally and remotely, and prototyped a pipeline to manage certificates on virtualized Ark nodes. The new CA ensures that every monitor can securely authenticate itself to Ark services, making the system more resilient and easier to scale. It also provides a uniform method for handling authentication across physical, virtual,

[2] https://smallstep.com/certificates/

and containerized nodes, tightening the integration between automated initialization and long-term management. 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. Our tooling will allow certificates with shorter validity periods which limits damage from compromised credentials, and facilitates automation. We will automate all certificate renewal on servers and probes, and automate their initial issuance on container-based VPs with known IP addresses.

4.3.7 Privacy: Respecting privacy of ISPs and vantage points

Unlike passive monitoring, which inspects existing traffic and can expose private communications, active measurements generate synthetic traffic whose content and destinations are explicitly defined by the measurement operator. Our nodes also adhere to published guidelines, such as rate-limiting, target whitelisting, and transparency about measurement goals, which further minimizes potential risks. We also maintain a Memorandum of Understanding with each site host [103]. However, inferences about critical infrastructure can yield sensitive information. We will use our proposed KeyCloak authentication and authorization framework (§10) to govern access to such data at various granularities based on our privacy impact assessments.

4.3.8 Storage: Standard and efficient storage formats and architecture

Our biggest problem with storage on the nodes historically has been premature SDcard wear. Even when we used high endurance microSD flash cards to maximize resilience to heavy write operations, we found that nodes failed in the field. We now provision and deploy VPs with ramdisk partitions to temporarily store data as it is being collected by the VP into a memory-based file system. We use an efficient binary storage format (scamper's warts format) and compress the data with gzip or bzip2 as data is collected. Further, we write smaller compressed contiguous measurement units to the ramdisk, which the central server reassembles before being shared with researchers.

4.3.9 Access: Findable, Accessible, Interoperable, Reusable (FAIR)

Our programming environment provides users with reference implementations of complex measurement functions that act as building blocks for more sophisticated measurements. The Python library interfaces to measurement capabilities present on a collection of remote vantage points, leveraging iterators and generators to efficiently process and yield measurement data streams. The measurements execute on the VPs and the system returns results as objects with normalized field names and field types for consistent data access and type safety. Result classes provide Pythonic interfaces to the data, incorporating methods for common code patterns to simplify data manipulation and analysis.

The library uses Python's datetime and timedelta to accurately represent and manipulate time-related measurement data, enabling precise temporal analysis.

4.3.10 Flexibility and Extensibility: Support new measurements

The system uses the open-source software Scamper [156] 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 new measurement primitives. Our choice of Python for the programming environment was because 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.

To incentivize deployment, the system should perform measurements that provide valuable information back to the hosting AS. Collection at this scale and density will allow us, and the hosting sites, to detect anomalies and security threats. The system will also support a dashboard, providing 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.

Cloud vantage points

Our design incorporates the use of commercial cloud vantage points when it provides sufficient coverage and is cost-effective. We consider an excellent use case for cloud VP deployments to be geolocation measurements: it is low bandwidth (so low-cost), and may provide sufficient diversity of vantage points. If not we can build tools to incorporate other vantage points in tandem with the cloud. The commercial clouds also allow us to experiment with virtualization/containerization of software deployment models.

Extending to Internet Exchange Points (IXPs) Deployment

Our original proposal envisioned scaling the number of vantage points by orders of magnitude, by integrating active measurement capability directly into the BGP collectors operated by the RouteViews project, enabling measurements through the thousands of peers connected to that infrastructure (§3.3.10). We developed and deployed software to test this approach, but ultimately RouteViews leadership informed us that they did not have peer permission for this type of measurement. This limitation meant that we could not rely on RouteViews collectors for active measurement in the foreseeable future.

We then explored an alternative approach of dedicated active measurement nodes at IXPs. The motivation is straightforward: IXP VPs observe and measure the Internet from a perspective quite different from VPs at the network edge. IXPs also present challenges,

since the routing environment is asymmetric and a router will normally install only a single best path, even when many peers offer routes to the same destination. Traditional looking glasses provide only limited visibility in such cases, and critically, they do not allow researchers to run active probes through a specific IXP peer.

We designed a mechanism to measure traffic through individual IXP members.[3] We compared two methods: (1) sourcing measurement traffic using the address space of our transit provider and (2) sourcing traffic using an address supplied directly by the IXP member. The goal was to overcome the limitations of single best-path routing and allow Ark to probe connectivity on a peer-by-peer basis. In exploring this design, we accommodated several operational requirements. We could not rely on IXP-assigned addresses for measurements, since they may not be routable beyond the exchange. Also, the Ark node must present only a single IPv4, IPv6, and MAC address, while appearing to communicate separately with many peers — and without requiring multi-hop BGP configuration on the peers. We proposed deploying a server-class machine that would host one Ark container per connected peer at the IXP. Each container would have its own address from a /24 (IPv4) or /48 (IPv6) block announced by the host. Scamper would originate traffic from the container's address, ensuring that results reflect only the routes advertised by the corresponding peer. Behind the scenes, NAT and source-routing allow us to map each container's traffic onto the host's primary IP. Each container would run its own FRR instance to maintain the peer's BGP session. Meanwhile, the host retains a default route to a transit provider, ensuring that return traffic is captured even if a peer's response fails to make its way back across the exchange.

Four IXPs were interested in participating but the operational cost and complexity of deployment (rack space, co-location costs, transit and peering agreements) proved prohibitive. It may be a component of a future active measurement infrastructure deployment, but we do not recommend this approach as the basis for an NSF-funded Internet measurement infrastructure in the short term. The requirements for cooperation (or colocation fees) are substantially higher than for the approach that we propose.

[3] https://gmi3s.caida.org/outcomes/ixp-active/

Chapter 5

Unsolicited traffic measurement (network telescopes)

Acknowledgments: Contributions in this section by Alex Maennel, Thomas Schmidt, Matthias Waehlisch, Ricky Mok, Max Gao, Raphael Hiesgen, Marcin Nawrocki, Daniel Kopp, Oliver Holhfield, Mattijs Jonker. Some text incorporated from publications with these authors, including [112, 120, 164].

A *network telescopes* is specialized instrumentation that captures unsolicited Internet traffic ("background radiation") directed at unused address space ("darknet") (Figure 5.1). Such instrumentation offers a unique vantage point to observe and analyze a wide range of Internet phenomena at a global scale. Collection of unsolicited traffic faces fewer, although significant, privacy concerns than collection of two-way traffic (§6). In this section we specify traffic monitoring functionality that advances the state of *network telescope infrastructure* to collect unsolicited Internet traffic, i.e., that has no trigger or response.

Over the last two decades, CAIDA (at UCSD) has operated the world's largest Internet traffic observatory (UCSD-NT) to capture Internet background radiation (IBR) from a darknet. The UCSD Network Telescope (UCSD-NT) consists of a globally routed, but lightly utilized /9 and /10 network prefix, which is about 0.4% of the routed IPv4 address space (about 12M IPv4 addresses). This address space contains few legitimate hosts. Inbound traffic to non-existent machines - so called Internet Background Radiation (IBR) - is unsolicited and results from a wide range of events, including misconfiguration (e.g. mistyping an IP address), scanning of address space by attackers or malware looking for vulnerable targets, and backscatter from randomly spoofed denial-of-service attacks. UCSD-NT continuously captures this anomalous traffic discarding the legitimate traffic packets destined to the few reachable IP addresses in this prefix. We archive and aggregate the data and share this valuable resource to network security researchers.

The resulting data has revealed key insights into malicious automated activities, in-

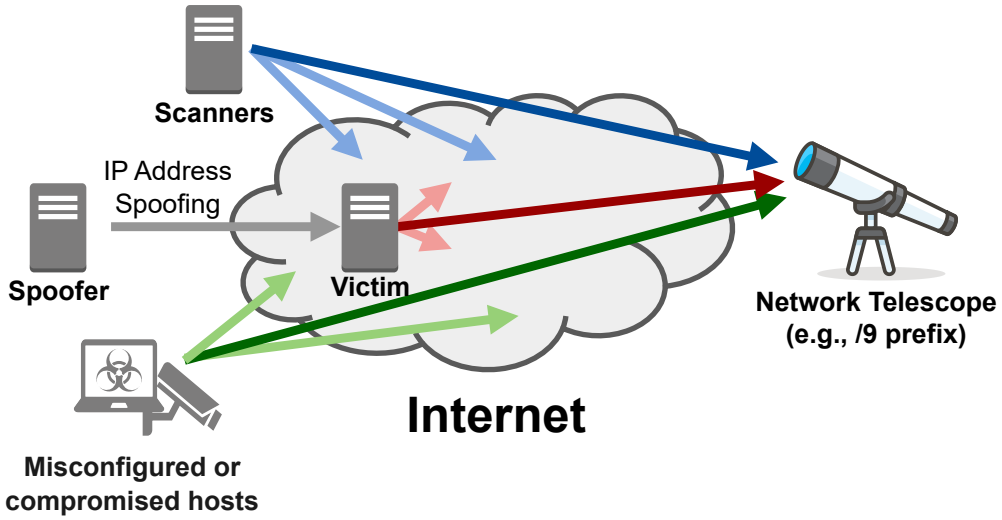

Figure 5.1: A network telescope captures Internet Background Radiation (IBR), i.e., unsolicited traffic from scanners (■), replies to spoofed traffic (■), and traffic from misconfigured or compromised hosts (■). IBR is captured if the destination IP address of the packets belong to the address space of the telescope (bold arrows). *(Source: [164].)*

cluding the spread of Internet worms and viruses [172–174, 217], spoofed-source denial-of-service (DoS) attacks [176], and large-scale botnet behavior [71, 194]. It has also exposed macroscopic events such as Internet blackouts caused by natural disasters [68], infrastructure failures [13], and government censorship [72]. Beyond security incidents, the telescope has informed our understanding of IPv4 address space utilization trends [70] and uncovered bugs and misconfigurations in widely used applications [14]. These insights are not just academic—they carry strategic importance for the stability, security, and economic resilience of the global Internet. UCSD-NT data has contributed to over 300 publications, often in collaboration with external researchers [99, 100]. The telescope has also been a cornerstone for multiple government-sponsored research initiatives [93–95, 98, 162]. While the early work is over 20 years old [175, 188] and has received test-of-time awards [208, 236], telescopes still enable novel scientific research and continually reveal new insights on the stability, security, and economic resilience of the global Internet. During the MSRI Design Phase, we undertook a twenty-year retrospective of our experiences operating the largest network telescope that shares data with the research community [164]. We include (in some cases verbatim, with permission) insights from that study throughout this chapter.

5.1 Limitations of unsolicited traffic measurement capabilities

The three biggest infrastructure challenges of network telescope instrumentation are collection and storage, efficient curation, and sharing large volumes of data. Some network telescopes use partially active (*live*, or *lit*) address space, which brings another challenge: co-existence with operational network traffic.

1. **Storage.** The large of volume of UCSD-NT IBR traffic (>100GB per hour, O(1TB) per day) imposes challenges for researchers to perform data analysis. To manage current storage constraints, UCSD-NT provides the most recent 60-days of pcap files on-site and sends historical pcap files to NERSC HPSS data archive for long-term storage. NERSC provides this storage as part of a U.S.-government funded program for cybersecurity data set.

2. **Compute.** CAIDA has historically provided virtual machines (VMs) to researchers to access live darknet traffic. UCSD-NT uses multicast to broadcast IBR traffic to the VMs. Researchers bring their code to the VMs to analyze real-time IBR traffic. However, the computational power of existing VMs that we provide to researchers for data analysis has long since become insufficient. Each research VM currently has 8 CPU cores, 32GBytes RAM and 100GBytes storage. But the processing time of FlowTuple files in the VMs is longer than the time duration that the files cover, inhibiting real-time analysis. In other words, it takes more than an hour to process an hour of data. This limitation has prompted the requirement for new methods of sharing the data with researchers, e.g., mounting the data on SDSC's HPC systems that researchers can apply to use.

3. **Visibility limitations.** The purely passive approach captures only a few types of security events, and malicious actors evolve their tactics to evade detection. The only type of DDoS attacks that leave backscatter artifacts on telescopes are those that use randomly spoofed source addresses. Some denial-of-service attacks launch from distributed botnets, where there is less need to spoof source IP addresses, in which case backscatter would not appear on a telescope. Attackers can also spoof in a non-random fashion, which will incur an uneven distribution of backscatter across the IPv4 address space, and may cause backscatter traffic to miss any specific telescope lenses. Note that the telescope by default does not send any packets in response, which also limits insight into the traffic it sees. Researchers have also deployed *honeypots* (e.g., [115, 118, 121, 141, 193, 240]) which react to unsolicited traffic to lure further engagement by attackers, yielding attack fingerprints, victim identification, and malware samples.

4. **Address space scarcity.** Decreasing availability of IPv4 address space makes it harder to obtain sufficient address space to establish IPv4 telescope instrumentation. Established telescopes are under pressure as the opportunity cost for owners

	Goal	Objective
(1)	Expand/optimize coverage	Increase number of vantage points
(2)	Performance	Handle growing rates of incoming data
(3)	Maintainability	Automate data pipelines and maintenance
(4)	Data Integrity	Protect against data loss/corruption
(5)	Standardization	Support modern protocols/formats/tools
(6)	Security	Protection against misconfiguration/malice
(7)	Privacy	Respecting privacy of ISPs and vantage points
(8)	Storage	Standard and efficient storage
(9)	Access	Findable, Accessible, Interoperable, Reusable (FAIR)
(10)	Flexibility/Extensibility	Support new measurements

Table 5.1: Design goals for traffic monitoring infrastructure

retaining unused address space increases with IPv4 address prices. Over the last decade, the UCSD-NT and Merit-NT have lost 37% and 97% of their underlying IP addresses, respectively (to high-priced sales).

5. **Lack of IPv6 telescope data.** Both telescope and honeypots face a daunting challenge with the growing use of IPv6. Scanning the vast IPv6 address space, or even small networks, is practically infeasible, so scanners must strategically target likely-active networks, which requires innovative algorithms to generate target hitlists. Furthermore, most existing honeypot implementations support one IPv4 address per instance. Running multiple instances to monitor a large blocks of IPv6 addresses is resource-prohibitive. Researchers have explored deployment of IPv6 network telescopes [220] to understand the evolution of IPv6 scanning activity. While IPv6 address exhaustion is not a concern, IPv6 telescopes bring other challenges such as attracting traffic to the vast and mostly empty space [83, 223].

5.2 Unsolicited Traffic Measurement Requirements

1. **Expand Coverage: Scaling up number of vantage points (VPs)**

   The largest telescopes (UCSD-NT and Merit-NT) have lost significant fractions of their underlying address space, reducing their visibility and scope of their data sets. Address space scarcity has motivated alternative strategies to expand coverage, such as leasing addresses or collecting traffic to unused address space in transit rather than the destination. However, using traffic captured from transit provider links requires accurate identification of IBR traffic within normal two-way traffic. IBR from more predominantly dark networks is also required to validate these new approaches against a baseline.

2. **Performance: Handle growing rates of incoming data**

   Improving in-line packet processing and filter capability will require advanced hardware, e.g., SmartNICs, particularly to intercept packets between the network interface and the operating system of the host machine.

3. **Maintainability: Automate data pipelines and maintenance**

   UCSD-NT maintains a large data processing software pipeline [98] that suffers from technical debt, and substantial opportunity to increase automation and modernization of underlying components.

4. **Data Integrity: Protect against data loss/corruption**

   The UCSD-NT now collects O(1TB) traffic daily, and continued growth in traffic rates pose challenges in maintaining lossless and error-free collection. Moreover, coexistence of this telescope with live Internet traffic requires careful filtering of traffic to subnets in the underlying address space that are legitimately assigned (leased) to users, and therefore not dark. In recent years, UCSD-NT's co-existence with ARDC's address space has grown more complex, as ARDC has expanded its leasing of address space to its members, requiring more frequent updates of the filter list. This complexity demands careful attention to ensure correct filtering of traffic destined toward the telescope, in order to preserve the integrity of the data and research that uses it [120, 164].

5. **Standardization: Support modern protocols/formats/tools**

   For historical reasons, different telescope instrumentation in the community uses different formats for their traffic, flow, or attack inference data, making it challenging, expensive, or impossible to compare them. One goal of new infrastructure would be consensus on which data formats are going to be most useful to the research questions deemed most important to the community.

6. **Security: Securing infrastructure against misconfiguration and malice**

   The system needs to have regular updates and patching, as well as access control to limit access to authorized users. Regular monitoring is necessary to protect against misconfiguration.

7. **Privacy: Respecting privacy of ISPs and vantage points**

   Unsolicited Internet traffic may include sensitive information such as IP addresses that could be linked to individuals or organizations. The system needs to anonymize IP addresses before broad sharing, and limit access to raw data to vetted researchers who have signed an Acceptable User Agreement (AUA) that limits publishing of identifiable details.

8. **Storage: Standard and efficient storage formats and architecture**

   The system should use data compression techniques and aggregation into flows to reduce storage demands. Ideally a new system could leverages the NERSC high performance storage systems to archive older data to cost-effective, slower storage—optimizes resource usage.

9. **Access: Findable, Accessible, Interoperable, Reusable (FAIR)**

   The infrastructure should implement as many options as feasible to provide researchers flexible access to the data, metadata, and derived data sets. Metadata for telescope data sets should be indexed into catalogs used by the cybersecurity community to make them easily discoverable by researchers.

10. **Flexibility/Extensibility: Support new measurements**

    The system should use modular software pipelines that can be reconfigured to filter, process, or analyze new traffic characteristics. The system must also facilitate integration of AI/machine learning models for real-time anomaly detection.

5.3 Proposed Design

Our specification for unsolicited traffic measurement infrastructure includes three components. First, it will extend unsolicited traffic collection to non-dark networks in enterprises and cloud infrastructure, to support both IPv4 and IPv6 traffic collection. Second, the infrastructure will include novel active techniques to attract malicious IPv6 traffic. Third, the system will leverage NAIRR resources to deploy tools to support machine learning (ML)-based time series analytic methods to detect anomalies in IBR traffic. Such tools will efficiently analyze over 200K time series to identify transient or persistent suspicious pattern changes. The system will leverage a new flow data representation, correlate anomalies in different time series, identify potentially affected services, and infer attack origins. These capabilities will facilitate use of machine learning/artificial intelligence (ML/AI) for cyber threat hunting, anomaly detection, and malware analysis. We describe how our approach addresses the specific requirements.

5.3.1 Expand and optimize coverage

Our first capability is to expand the visibility of darknet traffic by capturing traffic toward production networks. The system will leverage network/broadcast IP addresses in each subnet and the addresses assigned to router interface and point-to-point links to form grey-nets, a collection of dark IP addresses that interspersed with active addresses in the same subnets, to capture unsolicited traffic. In some cases, this will require monitoring ultra high capacity (400Gbps) networks of partner networks. It will also require scaling

the darknet and supporting software to handle additional traffic load. We propose to use three types of addresses to form grey-nets:

I. Network address refers to the first IP address in the entire subnet [11], which could represent a subnet (except /31 and /32 networks which use their 1-2 addresses for point-to-point connections [196]), or an IP broadcast address. By default, routers discard traffic to these network addresses [11]. Therefore, the characteristics of network addresses are similar to darknet addresses, *i.e.*, we do not expect any legitimate traffic from or to these addresses.

II. Broadcast address is the last IP address in a subnet [11] intended for IP broadcasts. Currently, only two protocols (DHCP and BOOTP) use IP broadcast addresses within internal networks. Therefore, packets from the Internet toward broadcast addresses are likely unsolicited.

III. Equipment address is assigned to physical or virtual network devices, such as router interfaces, and an endpoint of a point-to-point connection. These addresses might host management services (e.g., SSH, Telnet, SNMP), send/receive routing messages (e.g., OSPF and STP), and respond to ICMP traffic for network diagnosis, but should not carry any application traffic. Furthermore, these management services are often restricted to be accessible only within internal networks. Therefore, ingress traffic from the Internet is also likely unsolicited.

Using greynets as network telescopes not only leverages used address space, but is less likely to be circumvented (avoided) by attackers for two reasons. First, the unused addresses are embedded in the production network that hosts services or end-users [169]. Second, Classless Inter-Domain Routing (CIDR) [110] defines subnets with arbitrary prefix length; greynet addresses could be at any location in the subnet. The last octet of the addresses are not limited to 0 or 255 in the traditional class A/B/C addressing scheme. Analyzing a week of traffic to each IP in one /24 in the UCSD Network Telescope (UCSD-NT), we found no evidence that scanners avoid probing any address in the subnet [164]. Note that subnet assignments could change over time. Scripts will periodically download and parse network configurations from routers to identify greynet IPv4 addresses. We implemented such scripts for our recent NSF-funded CICI STARNOVA project [98].

IPv6 support The next generation telescope should also offer a flexible virtualized environment for researchers to facilitate the rapid development and scalable deployment of distributed **dual-stack (IPv4 and IPv6) telescopes and honeypots**, with novel techniques to attract malicious IPv6 network activities, overcoming the visibility limitations of previous attempts.

5.3.2 Performance: Handle growing rates of incoming data

The telescope traffic volumes (400 GB/hr) imposes challenges on capturing it. We designed new compression algorithms to optimize performance of merging packet capture files [92]. We investigated the use of SmartNICs to improve in-line packet processing and filter capability. We leveraged NVIDIA Bluefield-2 SmartNICs in CloudLab, an NSF-funded testbed, to evaluate the functionalities of the SmartNICs, particularly to intercept packets between the network interface and the operating system of the host machine. We have successfully used the Open vSwitch and DPDK support in the Data Processing Units (DPUs) in the SmartNIC to filter packets/network flows. We propose to adapt it to improve packet processing and annotation of network telescope data.

Using gzip and zstd compression tools improves performance by writing much less data to disk. The current implementation splits packets across 16 streams, which allows for IP address anonymization using multiple *traceanon* processes, improving anonymization performance. The final merged and anonymized files still use gzip compression for compatibility with other tools.

5.3.3 Maintainability: Automate data pipelines and maintenance

We propose several dimensions of maintainability. First we will adopt a modular design for software and hardware components. We will implement comprehensive logging and automated monitoring tools to quickly identify and diagnose issues.

5.3.4 Data Integrity: Protect against data loss/corruption

We undertook an extensive study of data integrity challenges with the telescope, and published the results at ACM SIGCOMM in 2025 [164]. Among the conclusions from that work is the need for continuous monitoring of operational data capture, including leveraging periodic signals from third parties such as benign scanning projects to see if that traffic is received as expected. Other validation tests include whether the number of stored and parsed packets match the number of incoming packets. Our SIGCOMM study provided substantial guidance and recommendations for present and future telescope operators and data users [164].

5.3.5 Standardization: Support modern protocols/formats/tools

Although it is a challenge to find community consensus on a standard flow representation that serves a wide variety of research, we propose a flow data representation that facilitates correlation of anomalies across time series, identification of potentially affected services, and inference of attack origins. The telescope should also leverage as many open source tools as possible. A core component is *Libtrace*, a userspace library for processing of network traffic capture from live interfaces or from offline traces.

5.3.6 Security: Securing infrastructure

The telescope address space should be monitored for hijacks and other events that might impact correct propagation and thus reachability of the prefixes. Open-source tools such as BGPalerter [26, 183] provide this functionality. To harden address space against potential hijackers of IP prefixes, creating ROAs [210] is recommended. ROAs allow other networks to filter invalid BGP announcements of telescope address space based on RPKI route origin validation [171].

Documentation of incidents Since most data consumers will use data retrospectively, incident reports should also be archived and made publicly available via an API, which allows data consumers to analyze the impact of external events on research results.

5.3.7 Privacy: Respecting privacy of ISPs and vantage points

The system will use CryptoPan to perform prefix-preserving anonymization of IPv4 and IPv6 addresses before broad sharing, and limit access to raw data to vetted researchers who have signed the CAIDA Acceptable User Agreement, which limits publishing of identifiable details in reports.

But for many research uses, anonymized data is not sufficient; analyses need real (raw) IP addresses, which are more sensitive. To support such access, the infrastructure must have policy frameworks and institutional agreements that enable trusted collaborators to analyze these data securely. These efforts aim to balance broad research utility with strong safeguards for privacy, security, and compliance. Throughout the Design Phase, we developed and refined policy tools that allow sensitive data to be shared responsibly with collaborators and trusted users. Telescope data is a prime example: it is both extremely high volume and highly sensitive, and historically we only permitted its analysis within CAIDA-managed virtual machines.

To prototype new approaches, we drafted and signed a new *Memorandum of Agreement (MOA)* between CAIDA and TU Dresden covering Telescope data access [23]. This agreement supports the development and maintenance of the UCSD-NT Telescope datasets, with provisions for dataset curation, integrity checks, metadata standardization, long-term accessibility planning, and enhancements to internal workflows and documentation to ensure reproducibility and usability for external researchers.

5.3.8 Storage: Standard and efficient storage formats and architecture

The data collection pipeline includes capturing raw packets and processing them into a more compressed flow record format for archiving. UCSD-NT generates FlowTuples every 5 minutes [96], which are Apache Avro formatted files for compact representation of network flow records (Table 5.3). Each FlowTuple record represents a sequence of

Table 5.2: Traffic metrics, properties, and filters that in combination yield over 200K time series.

Properties	Metrics (per minute)	Filters
Origin ASN	# of packets (PPM)	Unfiltered
Geolocation	# of bytes (BPM)	Non-Spoofed
Protocol number	# of unique source IPs	Spoofed (Derived)
TCP/UDP Destination port	# of unique source ASN	
ICMP type & code	# of unique destination IPs	
Spoofing inference		

packets sharing features, including source IPs, protocol and destination ports. The Corsaro software package computes traffic statistics of the flows (e.g., distribution/frequency of packet sizes, time-to-live value) and annotates each flow with metadata that facilitates analysis (e.g., prefix-to-AS [97], and IP geolocation). This information enables characterization of various types of malicious traffic, including scanning campaigns, without the overhead of analyzing raw packets.

In parallel, the system extracts thousands of time-series statistics directly from the packet headers (Table 5.2), yielding over 200K time series. UCSD-NT applies heuristics [69] to identify traffic with spoofed source IP addresses and implement filters to prevent them from compromising our statistical analysis. We use InfluxDB [131], a time series database, to index the data, and Grafana dashboards [117] to publish interactive visualization [89].

5.3.9 Access: Findable, Accessible, Interoperable, Reusable (FAIR)

We have tested several ways of sharing Telescope (IBR) data with researchers for open source and commercial efforts, and we recommend they all be included in a future telescope data infrastructure:

1. Access to raw historical pcap files from the NERSC archive, again with trusted collaborators who already have an established relationship with NERSC (Lincoln Labs, a DOD FFRDC.) MIT Lincoln Labs has generated at least six papers using data accessed via this mode.

2. A data exporter can send a subset of packets received by the telescope, or curated (reduced from original) event data, to a vetted collaborator/partner over existing infrastructure.

3. We share data through virtual machine (enclave) access through our OpenStack Hypervisor system. This requires that users log into our virtual machines, and that

we limit CPU, memory, and disk usage on a per-user basis.

4. For collaborators who cannot meet their processing requirements using the VM option, the system can provide temporary direct access to a CAIDA compute server, or researchers can leverage NAIRR resources such as ACCESS to process telescope data. To make the data seamlessly available on SDSC HPC systems (Expanse), we integrated Swift object storage containers containing Telescope data with an Amazon S3-compatible setup. This model extends computational capacity far beyond what CAIDA's vir- tual machines can provide and enables researchers with ACCESS allocations to analyze datasets that would otherwise be too large for local resources.

5. A time-series dashboard of statistics of telescope traffic allows researchers to explore this metadata for suspicious events [112].

6. The system should use a community data distribution infrastructure for scientific data to maximize distribution of the data to the research community. A good candidate today is the Open Science Data Federation (OSDF), which is expanding its support for authentication and authorization for access to data. To manage privacy (sensitive data, users must sign an AUP with us to access), the system will encrypt the packet capture data with the same methodology used to store archives at DOE's NERSC (OpenSSL with AES encryption).

The variety of approaches has illustrated to us how much benefit there is to be able to accommodate different needs in accessing the data. But of course, each of these approaches required dedicated IT staff time and attention to make work.

5.3.10 Flexibility/Extensibility: Support new usage modes

The variety of data sharing methods described above should promote considerable flexibility and extensibility with the data, and we also propose to improve the FlowTuple format with additional information and annotations to facilitate use of AI tools on the data.

1. Hostnames contain rich information about IP geolocation [161] and ASNs [163]. We will use reverse DNS to resolve source IPs to hostnames at the time UCSD-NT creates the FlowTuples.

2. Acknowledged benign scanners scan the Internet for research studies or cybersecurity monitoring. They are less likely to be malicious. We will use public collections of known-benign scanning IPs (e.g., [62]) and information provided by the scanner's websites (e.g., [29]) to identify benign scanners by source IPs and hostnames.

3. Scanner implementation provides crucial information about the nature of the traffic. The existing FlowTuple format does not provide TCP/IP header values to enable

use of heuristics [222] to infer scanner implementation (e.g., Hajime, Zmap, Mirai). However, including more header values in the FlowTuple record will significantly increase the file sizes. Instead, we will compute the heuristics at the time of generating the FlowTuple, and add the inference as a new tag in FlowTuple.

4. *Packet payload samples* could help identify the target and intention of the traffic, such as services or vulnerabilities. Providing representative samples can accelerate analysis of anomalies.

Table 5.3: Current information and proposed features (bolded) for FlowTuple [96].

Categories	Information
Time	Timestamp of network flows
Network flow information	Source IPs, destination prefix, destination ports, IP Protocol
Summary of traffic properties	destination IPs, TTLs, (TCP/UDP) source ports, TCP flags, packet sizes, TCP flags, **sample payload**
Source IP Annotation	Source IP geolocation (Maxmind and NetAcuity), Prefix-to-AS, **hostname**
Inference	Spoofed packets, Sent using Masscan/**Hajime/Zmap/Mirai**, **Acknowledged scanners**

Chapter 6

Traffic data: two-way

Internet traffic measurement data is the most sensitive and costly type of data to gather, curate, store, and share. But all harms on the Internet ultimately arrive in some type of traffic – security is thus the primary motivation for extensive investment in traffic monitoring technology.

6.1 Limitations of current two-way traffic measurement capabilities

For decades it has been virtually impossible for researchers to get access to passively collected traffic data from Internet backbone links due to privacy concerns. Commercial ISPs and backbone providers treat traffic data as sensitive intellectual property and a competitive asset. Concerns about customer privacy, security risks, and business confidentiality mean that researchers rarely gain direct access to raw packet traces or flow data. When access is granted, it is often highly aggregated or anonymized to the point that it limits the types of scientific questions that can be asked. This creates a persistent gap between the richness of data held by commercial operators and the limited datasets available to the academic research community.

These constraints against traffic data sharing highlight a critical challenge: academic researchers lack the direct, representative, and longitudinal traffic measurements needed to fully understand the dynamics of commercial Internet backbones. Bridging this gap would require new models of data sharing, privacy-preserving measurement techniques, and stronger partnerships between industry and academia. Doing so is a significantly more substantial investment that is likely to succeed under the NSF MSRI program. As such we do not propose a full specification of a two-way traffic monitor at this time.

However, in this section we describe the traffic collection system we have designed and prototyped during the Design Phase. (This work leveraged a recent NSF-funded CIRC award (NSF CNS-2120399) [84].) Note that the hardware ecosystem moves rather

quickly; the monitor we created was already obsolete by the end of the Design Phase as the links underneath the monitor upgrade to higher bandwidth requiring the need for a higher-bandwidth (400GB) monitor capability.

6.2 Proposed Design (Existing Capability)

Based on trust relationships that have been maintained for over two decades, CAIDA has been able to measure strategic links in the backbone so long as CAIDA could provide funding for the monitor. Since April 2008, CAIDA's passive traces dataset contains traces collected from high-speed monitors on a commercial backbone link, and anonymized for sharing with the research community. Six times in the last 20 years the underlying link was upgraded beyond the scope of the project budget (OC3, OC12, OC48, and OC192, 40GB, and now 100GB). CAIDA's last remaining single last point of public insight into the commercial Internet backbone was lost from January 2019 to February 2024 when we managed to deploy a 100GB monitor on a backbone link. Although that link upgraded to 400GB in 2025 so we had to move the monitor to another (now relatively low bandwidth, non-core) backbone link.

During the Design Phase (as part of the CIRC effort), we deployed a completely new passive traffic monitor on a 100 Gbps backbone link at an IXP in Los Angeles. Using Napatech network cards, we recorded these traces while removing the payload (beyond the layer 4 headers) from all packets. The card interprets various layer 4 headers, including ICMP, ICMPv6, TCP, UDP, SCTP, and GRE, and strips others. For performance optimization, our packet capture architecture utilizes 16 streams, which we then combine into two unidirectional traces.

We anonymize these traces using CryptoPan prefix-preserving anonymization. Previously, the CryptoPan implementation did not support the encryption of bit strings longer than 32 bits, which limited its ability to anonymize 128-bit IPv6 addresses. We now use an updated version of CryptoPan that can anonymize all 128 bits of IPv6 addresses. Our capture and post-processing workflow is thoroughly documented to ensure clarity and reproducibility [22]. We also configured local compute and disk resources to allow proper processing of captured data.

The current data acquisition workflow involves capturing, post-processing (anonymizing), and transferring each monthly snapshot into a Swift storage container for researcher access. This restricted dataset includes the following metadata fields:

- Monitor Name
- Year and month (including a link to a graphical display of breakup by protocol, application, and country)
- Start time of trace (UTC)

- Stop time of trace (UTC)
- Number of IPv4 packets
- Number of IPv6 packets
- Unknown packets (as a fraction of the total number of packets)
- Transmission rate in packets per second
- Transmission rate in bits per second
- Link load (as a fraction of the nominal maximum load for a 100 GB link)
- Average packet size (bytes) (including a link to a graph of the packet size distribution).

We created new trace processing tools by incorporating scripts creating summary statistics similar to Trace Statistics for Passive OC48 and OC192 Traces [21] This publicly available dataset includes cumulative metadata for 100Gb traces, as well as accompanying graphs that display data rates and distributions of IPv4 and IPv6 packet sizes.

Chapter 7

DNS measurement

Acknowledgments: Contributions in this section by Gautam Akiwate, Mattijs Jonker. Raffaele Sommese, Some text incorporated from publications with these authors, including [3, 212–214].

Malicious actors exploit the DNS namespace to launch spam campaigns, phishing attacks, malware, and other harmful activities. Combating these threats requires visibility into domain existence, ownership and name service activity that the DNS protocol does not itself provide. Beyond namespace abuse, adversaries also target the reliability of the DNS ecosystem itself. They launch Denial of Service attacks aimed at disrupting or degrading the resolution process, a fundamental component of Internet communication. In addition, they can misuse security mechanisms such as DNSSEC to amplify attacks against other parts of the Internet infrastructure. There is a vast landscape of vulnerabilities in the DNS, far more than in the other systems we described, due to the higher complexity and much more challenging political economy of the DNS. Similar to user traffic data (§6), there are privacy sensitivities associated with aspects of DNS measurements, most notably personally identifiable information in domain ownership data, and inference of user behavior from domain lookups.

As such, in our Design Phase, we spent more time on assessment of existing DNS harms and data that could facilitate improved study of them, relative to other types of data in the project. Specifically, we considered: (1) active probing of DNS infrastructure; (2) passive DNS measurements (traffic capture of queries and responses); (3) zone files (4) domain blacklist data; (5) logs from DNS servers; (6) registration metadata (e.g., owner, hosting registrar); (7) domain pricing information; (8) evidence of role of DNS in various attack chains, e.g., misdirection; (9) estimates of actual harms due to DNS-related attacks. But as with traffic data, moving forward with a global DNS measurement capability that can move the needle on security research merits merit its own (and would be worth of an) entire MSRI Design project. We discovered and evaluated recent design

advances in this direction, and contributed to extensive evaluation of newly proposed components. We summarize these efforts in this section and refer the reader to the publications we released on these topics [3, 212–214].

Although our findings highlight 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 and policy 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.

7.1 Limitations of current DNS measurement capabilities

We summarize the biggest limitations of current DNS measurement capabilities: fragmentation and access restrictions of efforts; coarse granularity of data collection; privacy, ethical, and legal considerations; and sustainability and coordination challenges.

Fragmentation and access restrictions. Current DNS measurement efforts, while valuable, are generally fragmented and limited in scope. each with different levels of visibility and accessibility. For example, the Centralized Zone Data Service (CZDS) provides access to geld zone files but is limited to daily snapshots [130], subject to ICANN's access policies, and excludes many ccTLDs and private DNS zones. Researchers must apply for access and abide by contractual restrictions, creating barriers to broad and timely use. This fragmentation makes it difficult to build a comprehensive and continuous picture of global DNS activity.

Temporal and resolution limitations. Most existing DNS data collections suffer from limited temporal granularity and retention. For example, one of the best-known sources used today for studying macroscopic attack surfaces in the DNS is ICANN's Centralized Zone Data Service (CZDS) which is limited in coverage (the subset of TLDs under the governance of ICANN contracts) and granularity (24-hour daily snapshots). This coarse granularity obscures short-lived DNS records (e.g., fast-flux hosting, botnet infrastructure) that may exist only for minutes or hours. It also hinders study of other transient phenomena, detection of emerging attacks, or longitudinal studies of operational practices and new protocol deployments. Moreover, ICANN supports no public archive of historical snapshots. CAIDA has helped Ian Foster support an indexed database of CZDS (and other) Top Level Domain (TLD) zone files stretching back over a decade via our DNS Zone Database, but this is an unfunded volunteer effort and cannot scale support for community use of the data.

Privacy and legal constraints. Passive DNS data—collected at resolvers, authoritative servers, or recursive infrastructure—offers richer insight into real traffic, but relies on

opportunistic data collection from client queries but has no control over the temporal spacing of those queries. Several companies gather such data and one (Domain Tools) has made significant effort to share its passive data with with academics for scientific research. DNS queries can reveal user behavior, enterprise infrastructure, or sensitive operational details. Similarly with domain ownership information, extremely valuable for security and DNS abuse research, is in the hands of registrars and registries who have no incentive or direct interest in providing data about domain names. Legislative complications of as the emerging privacy regulations challenge risk assessment.

Sustainability and coordination challenges. As with other Internet measurement projects, many DNS measurement projects are funded by short-term grants or academic initiatives, making long-term sustainability uncertain, e.g., DZDB, dns.coffee. This results in inconsistent coverage, lack of interoperability across datasets, and gaps in metadata documentation. Without centralized coordination or funding, projects often duplicate effort rather than complement each other, limiting efficiency and cumulative impact.

7.2 DNS Measurement Infrastructure Requirements

A future DNS data collection infrastructure to support Internet security research should be built on similar principles as the other data types: an expansive set of vantage points (as DNS resolution can depend on the source IP address of the client); performance; maintainability; data integrity; standardization; security; privacy; storage, access, and flexibility/extensibility. Given the complexity of the DNS ecosystem, we include an additional requirement: a robust querying and analysis platform for historical data.

7.3 Proposed Design

We envision two pillars of a future DNS measurement infrastructure to support security research: an active measurement system; and a zone update sharing component, which will require cooperation of TLD registries.

7.3.1 Active Measurement

The active measurement system component of the infrastructure should be modeled on OpenINTEL [231]. This system will perform a daily, exhaustive scan of the DNS namespace for all known second level domains. The measurement process must collect a wide array of DNS record types to enable diverse research, going beyond simple A and AAAA records. The data dictionary from the OpenINTEL project provides a detailed blueprint for the schema, including critical fields like RTT (round-trip time) and a full suite of DNSSEC-related fields, which are essential for security, performance, and availability

Table 7.1: DNS queries to collect

Record Type	Description	Use
A/AAAA	Mapping a domain name to an IPv4/IPv6 address.	Census/mapping of infrastructure.
MX	Identify mail servers for a domain.	Mapping email service provider ecosystem.
NS	Nameserver records delegate authority for a domain.	Mapping the DNS infrastructure.
DNSKEY, DS, NSEC, NSEC3	DNSSEC records	Measure adoption and effectiveness of DNSSEC.
CAA	Certificate Authority Authorization records.	Measure security practices/policies for TLS certificates.
TXT	Text records.	Understand operational practices, such as DMARC and SPF.

analysis. This active approach provides consistent and reliable data over time, allowing for the observation of long-term Internet evolution and enabling investigation of activities of malicious actors exploiting the namespace and the DNS ecosystem. Table 7.1 shows a detailed list of the record types to be collected and their analytical significance.

7.3.2 DNS Transparency: finer-grained access to zone changes

As we analyzed in the Design Phase [214], a remarkably high concentration of malicious activity is associated with domains that do not live long enough to make it into ICANN's CZDS daily snapshots. We found that the daily snapshots miss at least 1% of newly registered and short-lived domains, which are frequently registered with likely malicious intent. In reducing this critical visibility gap using public sources of data, we demonstrated how more timely access to TLD zone changes can provide valuable data to better prevent abuse [214].

One persistently proposed countermeasure has been finer-grained transparency into (an audit trail of) DNS zone operations – the DNS equivalent of RouteViews and RIPE RIS for BGP data, or Certificate Transparency for certificate data. The idea is not new. In 2007, Verisign proposed and launched a service to share frequent (every five minutes) updates to the .COM and .NET zones [129]. Concerns about the potential for privacy abuses led Verisign to remove access to this DNS update service.

More recently (2017), DNS security experts proposed an effort they call *DNS Transparency* [8, 9], which would enable even finer-grained visibility into zone file changes: a

pub-sub system where DNS operators would share their zone updates in nearly real-time, and interested stakeholders (researchers, security analysts) would subscribe to feeds of these data streams.

We believe that it is time for registrars and registries who want to establish themselves as serious about security to resurrect such a service along with a code of conduct to safeguard against abuses. Given the imminent launch of another round of new gTLDs and the tremendous concerns by security researchers regarding the lack of transparency, it is an ideal time to engage in a design of such transparency and appropriate data disclosure frameworks that protect privacy but allow scientific research into the systemic operational risks of DNS infrastructure. We provide an expanded discussion of this idea in [214]. We believe such and initiative could start with security-conscious and collaborative TLD registries such as .US, .CH, and .NL.

Chapter 8

TLS certificate measurement

Acknowledgments: Contributions in this section by Gautam Akiwate, Mattijs Jonker. Raffaele Sommese, with text incorporated from work with these authors: [214, 215].

The TLS certificate system is the only system we studied that was not part of the original Internet architecture. There are two main types of TLS certificate measurement: active scans of IP addresses to gather TLS certificate information; and certificate data logged in Certificate Transparency (CT) logs. Scanning IP address space for certificate is within reach for a typical researcher and has become quite common due to the release of open source software (zmap [82]) for scalable, performant execution of scan. Censys, a company founded by the author of this software, also performs these scans regularly and shares resulting data with researchers. The second type of data – Certificate Transparency (CT) logs – is a bigger challenge. In this chapter we focus on this latter category of data, and offer a possible path forward.

8.1 Limitations of current TLS certificate measurement capabilities

Certificate Transparency (CT) is a global logging infrastructure designed to improve accountability in the TLS ecosystem. Each CT log is an append-only, cryptographically verifiable record of certificates issued by trusted Certificate Authorities (CAs). The system enables browsers, researchers, and security firms to detect mis-issued or malicious certificates by providing an open, auditable history of certificate issuance. In principle, CT logs offer a powerful mechanism for monitoring the certificate ecosystem and detecting abuse. In practice, however, leveraging this infrastructure presents significant limitations. Every day, more than 2.5 million certificates and 1 million pre-certificates are published across more than 240 active logs. While open-source tools have emerged to support CT log monitoring, such as streaming-based approaches (e.g., Certstream [25])

and batching crawlers (e.g., Axeman [24]), they often suffer from limited maintenance, incomplete coverage, or poor scalability.

Real-time processing of CT logs is particularly difficult. It requires not only continuous ingestion of large data streams but also a clear definition of which indicators of compromise (IoCs) to look for. In practice, most incident investigations are performed in a post-mortem fashion, which means that many relevant signals are only identified after an attack has taken place. Anticipating all possible abuse cases in advance requires both significant infrastructure over-provisioning and strong predictive capabilities.

A well-known example is the mis-issued certificate for Cloudflare's IP address 1.1.1.1 used for their public DNS resolver [134]. Few analysts were monitoring for certificates issued directly to IP addresses. Since certificate issuance overwhelmingly targets domain names rather than numeric IPs, the event initially went unnoticed in CT streams. It was only retrospectively identified and analyzed, showing how even high-profile anomalies can escape real-time detection when they fall outside of expected monitoring heuristics.

Historical CT log analysis presents an equally demanding set of challenges. Collecting, indexing, and querying the data at scale requires substantial computational, network, and storage resources. Public search services such as crt.sh are widely used in the community, but they are not always reliable and often experience outages. Some researchers rely on private efforts, such as those undertaken by the OpenINTEL team, to curate long-term CT datasets and provide consistent access for scientific or operational community.

CT logs also have inherent technical limitations that complicate their use. Each log is independently maintained and may implement slightly different policies or APIs, which makes large-scale aggregation cumbersome. Query support is generally limited, as CT logs are designed as append-only Merkle trees for verification purposes, rather than as searchable databases. This means that efficient lookups, filtering, and correlation across logs require building and maintaining external indexing systems. Furthermore, certificate data itself is noisy: a large fraction of entries are benign or irrelevant, and distinguishing malicious certificates from the background of legitimate issuance is non-trivial.

Together, these factors make CT measurement an area that is both essential for Internet security and difficult to operationalize at scale.

8.2 Proposed Design

A natural path to overcome the limitations discussed above is the development of an indexable and scalable platform for Certificate Transparency (CT) data. We do not propose to build such a system, but only to describe a feasible design direction that could be adopted by research groups, security firms, or community initiatives. The aim of this design would be to provide researchers with an efficient and low-resource means of accessing, filtering, and alerting on both historical and live certificate data, avoiding the need for each research group to maintain its own ingestion pipelines, heavy storage

infrastructure, and ad-hoc scripts.

At the conceptual level, the design would rely on a continuous ingestion layer that consumes certificates from all public CT logs, normalizes their structure, and enriches them with metadata such as parsed Subject Alternative Names (SANs), validity periods, issuer details, and fingerprints. These normalized records could then be exposed through scalable indices optimized for different research tasks. Time-partitioned indices would allow efficient retrospective queries, while inverted indices would enable fast lookups by domain name, IP address, certificate authority, or other attributes.

An important feature of this design scenario is the clear separation between raw data storage and queryable metadata. Raw certificate entries could be archived in compressed columnar formats (e.g., Apache Parquet, ORC) for scalability, while researchers primarily interact with the lightweight indices. Such a system would substantially reduce the computational and storage burden on users: instead of downloading and parsing terabytes of raw CT log data, a researcher could retrieve precisely the subset relevant to a study, such as all wildcard certificates issued by a given CA within a specified time window. This is analogous to the functionality offered by BGP2Go [228] (§3.3.9) which indexes routing data and enables users to download only the MRT files relevant to their query rather than entire archives.

The same design also anticipates real-time monitoring capabilities. A streaming interface, built on top of existing services such as Certstream, could allow users to subscribe to filtered CT data feeds with server-side rules. For example, an analyst could register an alert for any certificate containing a specific IP, or for certificates issued under unusual conditions. Such filtering would significantly lower the barrier for proactive detection, since users would receive only relevant notifications without needing to maintain large-scale ingestion infrastructure.

This proposed design is as an example of how the CT measurement ecosystem could evolve, much as the BGP research community transitioning from raw usage of route collectors to the more modern and actively maintained platforms like BGPStream, BGPKit and BGP2Go. Public search services such as `crt.sh` have demonstrated the value of indexed CT data, but their instability highlights the need for more robust alternatives. By combining scalable ingestion, rich indexing, efficient storage, and community-driven maintenance, this approach would enable more effective incident response and deeper longitudinal studies of the certificate ecosystem. A system of the kind we describe here would provide researchers with a common platform, reducing duplication of effort, ensuring sustainability, and broaden the use of CT data to a much wider set of users.

Chapter 9

AI-enabled analysis components

These proposed infrastructure components contribute to and rely on a interconnected ecosystem of Internet measurement data, tools, and services (Figure 9.1). In this chapter we describe additional components that will illustrate and amplify the value of the research infrastructure by developing and analyzing strategic datasets that reveal vulnerabilities, risks, and security challenges across the global Internet. The raw and derived datasets from the infrastructure components 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 [63–66, 186, 233]. For example, CAIDA and the NSRC RouteViews team already support many external groups who are leveraging BGP data and CAIDA's curation of such data (e.g., prefix2AS, AS Rank) to provide sophisticated security-related analysis and visualization tools. These include IIJ's Internet Health Report [86], IIJ and RIPE's Internet Yellow Pages [87]. Other visualization systems (GRIP) [224], including commercial systems (Cloudflare, Thousand Eyes, Catchpoint). In this section we propose five analytics components that will leverage data from the proposed platforms to support infrastructure and cybersecurity research: AS Rank, Spoofer, Internet Topology Analytics Pipeline (ITAP), Pathfinder, and AI-enabled metadata inference (AIMI).

9.1 AS Rank: Ranking of Autonomous Systems (AS) footprints

CAIDA's AS Rank is a critical resource for understanding the structure, dynamics, and influence of Autonomous Systems (AS) in the global Internet. By providing transparent, rigorously collected, and continuously updated rankings of AS connectivity, it enables researchers, policymakers, and industry stakeholders to monitor resilience, detect vulnerabilities, and assess the concentration of control in the Internet's core infrastructure. We believe continuation of this analysis component is important to a future MSRI, and is not likely to be supported by any other source of funding. Public funding en-

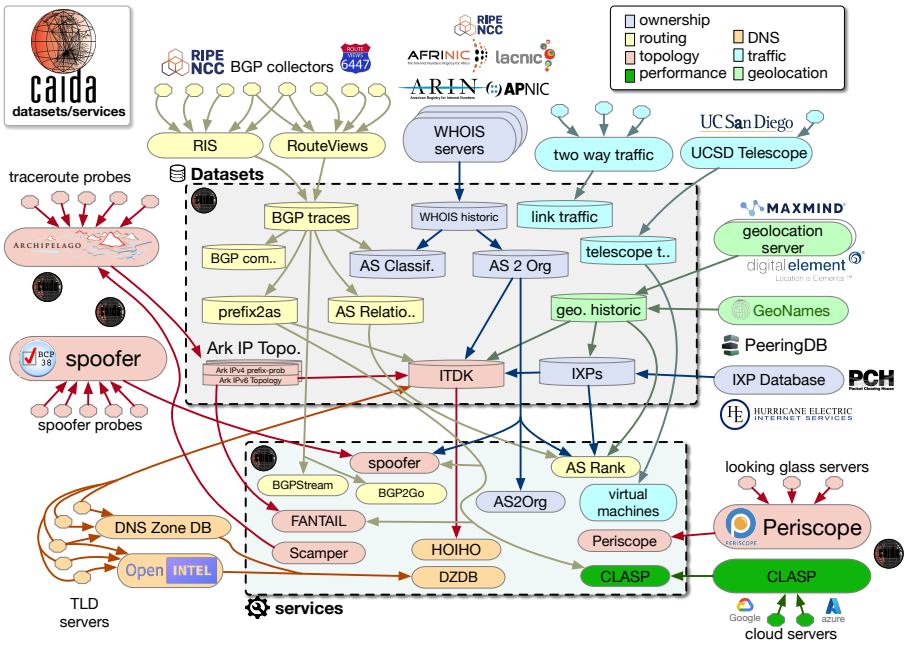

Figure 9.1: GMI Data Acquisition and Curation Architecture

sures that this data remains openly accessible, independent of commercial interests, and aligned with broader national priorities in cybersecurity, digital resilience, and scientific research. Supporting AS Rank through government investment strengthens both the research ecosystem and the nation's ability to safeguard critical communications infrastructure. However, as AS Rank has weathered since its creation and gained utility among Internet operators, flaws and issues with its operation have become apparent. Moreover, the Internet itself has evolved. Our proposed specification for a modern implementation of AS Rank is as follows.

1. **Step 1: Download and Store MRT Files** in an archive. Retry up to 10 times, with 30 minutes between tries. If still failing after 10 tries, abort and report. This process should be restartable.

2. **Extract MRT to Text** and store in byte-sorted order as follows:

```
f [index] [file_path_name]
p [AS|PATH] [CIDR/prefix] [file_index,...] count
r [CIDR/prefix] [AS|PATH] [file_index,...] count
```

   Here, `count` represents the number of times an entry appeared. Since multiple

MRT files from the same source are analyzed, `count` may exceed 1 for persistent announcements.

Additionally, store two byte-sorted index files:

- `[AS|Path] [file offset start]`
- `[CIDR/prefix] [file offset start]`

Store two files of AS paths: one for paths with at least one IPv4 prefix, and one for paths with at least one IPv6 prefix. Input the `[count]` of times a path/prefix combination must have been logged to be considered, to eliminate transients.

3. **Download IRR and RPKI Data** for all ASes appearing in the paths. This data is not currently used, but IRR and RPKI can indicate that AS X is not a customer of AS Y, as providers do not authorize customers to propagate their routes.

4. **Process IRR and RPKI Data** into a simple assertion of the form:

```
f [index] [text description of source]
n [AS X] [AS Y] [index]
```

   This indicates that AS X may announce routes for neighbor AS Y.

5. **Step 5: Compress Files** with a modified `pbzip2` that stores an index of the file position starting each compression block and the associated uncompressed offset, ensuring indexes remain useful. Leave sorted indexes uncompressed for binary search during drill-down investigations. Submit index code upstream to `pbzip2` authors.

6. **Step 6: Process Path Files** to determine customer/provider/peer relationships, extending to include IPv6 alongside existing IPv4. For IPv4 and IPv6, produce a relations list using at least two techniques: pure heuristic with pre-seeded clique; and heuristic with pre-seeded clique and external data hints. For the first heuristic pass, use data manually provided to CAIDA from operator feedback, combined with data from PeeringDB to assign relationships to observed path adjacencies. For additional heuristic passes, add a final gate to to accepting a relationship: IRR/RPKI assertions must not contradict the heuristic.

   Write machine-readable trace logs with three types of entries:

   (a) A path was rejected in whole, with the reason (e.g., AS loop or unregistered AS numbers).

   (b) A path was altered, including the new and old paths and the reason (e.g., known IXP ASes presumed to be route servers are removed from paths).

(c) Each accepted AS relationship, the heuristic assigning it, and the path it was derived from.

(d) Each notable rejected AS relationship, the heuristic considering it, and the path it was derived from. Notable rejections (TBD) include disagreements between heuristic metrics and IRR assertions.

7. **Produce Customer Cones** for IPv4 and IPv6, and for each relation technique.

8. **Produce AS Rank Presentation Data** as compressed text files.

9. **Present on AS Rank Website** including both IPv4 and IPv6 data, and data for each AS Rank technique. Develop website tools to drill back through trace logs to display relevant information.

9.2 Spoofer: identifying networks that allow spoofing

In 2015 a group of network operators defined a set of operational practices that can prevent several types of addressing and routing abuse [211]. This Mutually Agreed Norms for Routing Security (MANRS) initiative[1] depends on (unfunded) infrastructure operated by 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 [57, 192]. In 2018 Luckie *et al.* found that that MANRS participants were no more likely to properly deploy SAV than others [160]. CAIDA's Spoofer measurements have generated many scientific publications [15, 128, 150, 151, 160, 177, 178]. This dataset serves as a tool for detecting, mitigating, and preventing IP spoofing, which is a common vector of cyber-attacks like DDoS, man-in-the-middle attacks, and other malicious activities. 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 platform relies on BGP and active measurement data as described earlier (§3 and §4). This project exemplifies *translational security research* – technical knowledge converted to measurable improvements in infrastructure security.

This infrastructure currently operates on old hardware and several software components need updates. The next generation version of the server software should be deployed on an OpenStack instance, using a packaging software framework so that the platform operator (CAIDA or some other entity) can easily deploy a new instance of the infrastructure in case of hardware failure. Upgrades to this infrastructure will allow

[1] The *Mutually Agreed Norms for Routing Security (MANRS)* [211] 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.

continued publication of the Spoofer data set, which will help researchers and operators continue to analyze and track IP spoofing attacks, techniques, and related vulnerabilities within network traffic.

9.3 Macroscopic Internet topology analytics pipeline (ITAP)

Our specification of this component is currently embedded in documentation in our gitlab repository; we did not have time to complete this specification as part of the Design Effort but intend to pursue it during an Implementation Phase.

An Implementation Phase of the active measurement infrastructure as described in §4 should be accompanied by a new scalable software pipeline to automate the construction of one of our most powerful and scientifically generative datasets—the Macroscopic Internet Topology Data Kit (ITDK) [132]. 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 [4, 5, 28, 122, 136, 137, 149, 153–155, 161, 195, 199, 226]: But current infrastructure constraints limit the coverage, depth, and 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 [45].

9.4 Pathfinder: AI-enabled metadata inference (AIMI)

As we assessed the requirements of future active measurement infrastructure, we realized that the new system needs to support emerging national security research needs by providing comprehensive insights into global Internet infrastructure. Therefore, we designed *Pathfinder*, a path analytics platform that enables users to execute, search, and annotate traceroutes with enriched metadata such as per-hop geolocation, ownership details such as Autonomous System (AS) ownership, and annotations of router manufacturer information where applicable. By integrating data from geolocation services, WHOIS records, BGP paths, and active measurements conducted by widely distributed vantage points, this component will produce detailed annotations that support sophisticated annotations and searching of Internet topology data.

Pathfinder will support multiple types of requests through both a web interface and an API. Users may (1) initiate and annotate a new traceroute between an Ark monitor and a destination; (2) annotate traceroutes that they provide to the system; (3) search annotated traceroutes collected previously, filtering by attributes such as organization, country, or threat category; and (4) annotate individual IP addresses with enriched metadata such as

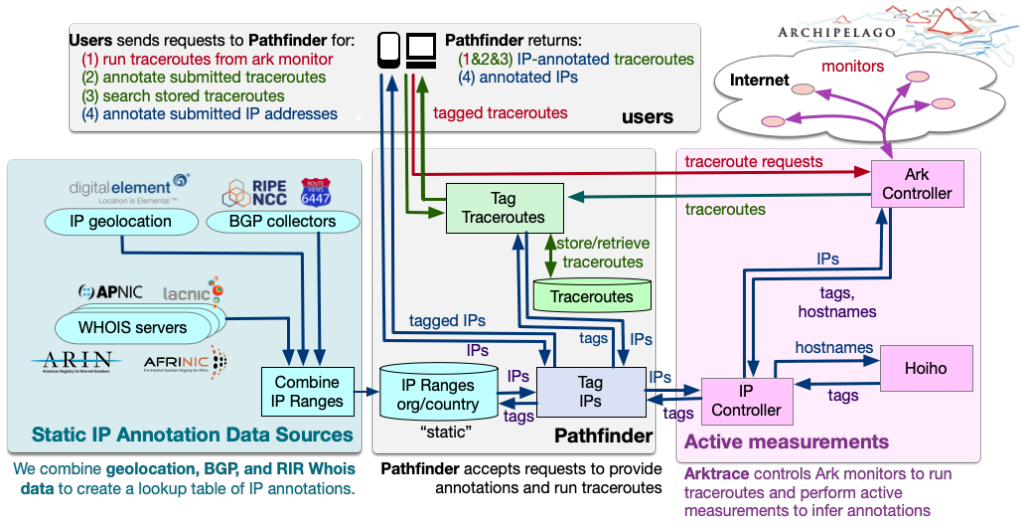

Figure 9.2: *Flow of requests from users through Pathfinder, a proposed system to produce unified and detailed annotations that support advanced research and analysis of Internet routing behavior.*

inferred organization, geographic location, and router vendor information.

The platform will enrich traceroute results by fusing static and dynamic sources of information. Static annotations include WHOIS records, Regional Internet Registry (RIR) data, IP-to-AS mappings, geolocation services, and BGP path information. Dynamic annotations come from active measurements conducted by Ark monitors Pathfinder merges these inputs into a unified annotated dataset, so that each IP address or hop within a traceroute is accompanied by a consistent set of rich metadata that renders a more accurate view of the global Internet. In addition to user-facing features, we designed Pathfinder with extensibility and usability in mind. The architecture integrates a front-end UI for interactive use, an API for programmatic access, and backend services that interface with Ark monitors to run measurements.

Figure 9.2 displays the flow of requests from users through Pathfinder to Arktrace, showing how traceroute execution and annotation are integrated into a unified workflow. Requests enter through the web interface or API, trigger measurement or annotation modules, and return enriched results to users for interactive exploration or further analysis.

Pathfinder will serve as a repository of both system-generated and user-provided traceroutes, which will allow users to revisit and extend past work. For example, a researcher could upload their own traceroutes and annotate them with organizational or geographic metadata, or search for all traceroutes passing through a given AS or country. Combining traceroute execution, annotation, search, and storage into a single system, Pathfinder will support a wide range of scientific workflows.

Table 9.1: External data sources to interpret/analyze DNS measurements

Data Stream	Data Source (Example)	Primary Purpose	Integrated Use Case
Active DNS Measurements	OpenINTEL	Capture daily state of DNS records, including RTT and DNSSEC data.	Identify misconfigurations; measure protocol adoption.
BGP Routing Data	CAIDA, RIPE NCC, RouteViews	Map ASes and AS paths to understand Internet topology.	Link DNS changes to suspicious BGP events.
Certificate Transparency Logs	Various public logs	Track cert. issuance; identify new domains.	Detect malicious domain activity.
WHOIS Records	Registries, RIRs	Link domains/IPs to ownership.	Study full lifecycle of malicious domains.

9.5 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 believe eventually the proposed measurement infrastructure should include capabilities to discover, aggregate, and pre-process structured datasets (e.g., WHOIS, geolocation, BGP, DNS, CT logs, domain registration data (Table 9.1)), and natural language sources (e.g., websites, public reports) to create a data repository of metadata 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, and managing correlations across data sources [18]. We will use an open-source LLM 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 9.3). This component should implement prompt engineering, retrieval-augmented generation (RAG), prompt tuning, and/or fine-tuning for this task.

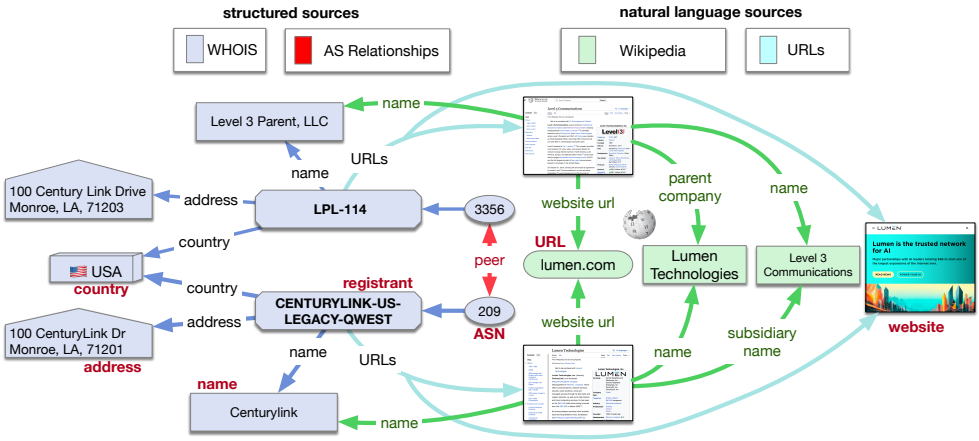

Figure 9.3: *Visual representation of a knowledge graph combining structured data (left) and natural language sources (right). We will train an AI model to extract and use such a knowledge graph to improve inference of Internet infrastructure properties, such as the fact that two ASes (ASN 209 and 3356, center of figure) are owned by the same organization. Researchers have previously relied on heuristic based analysis of structured (but often incomplete and out-of-data) data sources. LLM extraction of the above knowledge graph, including the associated merger, would dramatically improve data accuracy.*

Chapter 10

Data access: authentication

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.

10.1 Limitations of current capabilities

For many years, CAIDA has faced the challenge of identifying external users accessing our publicly available data and tools. Despite thousands of downloads and API calls annually, we only have information about users' IPs. We also manage restricted datasets via a separate access system that requires substantial manual data-administrator effort.

10.2 Authorization and Authentication System Requirements

We identified several requirements:

1. **SSO Capabilities:** Users authenticate once and gain access to all connected applications. The system should support SSO across multiple applications, and integrate with other identity providers (e.g., CILogon), so we can leverage existing user databases and allow users to authenticate via external providers.

2. **Protocol Support:** Keycloak supports OpenID Connect (OIDC), a modern industry-standard protocol for web and mobile SSO. It also supports OAuth 2.0, SAML 2.0, and others.

3. **Built-in User Management:** Keycloak provides a user-friendly admin console that allows us to manage users, groups, roles, permissions, and their credentials. We can configure user registration, password policies, and manage user sessions directly through the admin interface.

4. **Security Features:** Keycloak supports built-in Multi-Factor Authentication (MFA), e.g. TOTP, OTP, WebAuthn. Role-Based Access Control (RBAC) allows us to define roles at a granular level for both users and applications. Keycloak also offers session tracking, automatic session expiration, and user logout management.

5. **API:** The admin and user management functions should be are exposed through REST APIs, allowing us to integrate it with our RAM service and command line tools.

6. **Cost-effective:** We should use an open-source system to avoid licensing costs compared to other commercial IAM solutions like Okta or Auth0.

7. **Cloud-native & Containerized:** We need a system well-suited for modern cloud-native applications, with Docker and Kubernetes support for easy deployment and scaling.

8. **Single Point of Management:** The system should centralize the management of user authentication, which simplifies monitoring, auditing, and enforcing compliance across all integrated applications.

10.3 Proposed Design

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. It extends Keycloak to cover dataset- and request-level access, providing a single platform for managing resources within the infrastructure.

During registration, users authenticate via Keycloak and provide basic information: organization, department, role, and official email. After authentication, they can request access to data and/or tools by submitting a brief data-use justification. Requests for unrestricted data are approved automatically; others require administrator review. Access roles vary by resource type and intended use. Figure 10.1 illustrates the portal components and their interactions.

For administrators, the RAM Portal centralizes all operational work:

- An inbox of access-right requests shows each request with the user's identity and affiliation, resource, requested role, justification, status, timestamps, reviewer, and full history. Most requests for public datasets are approved automatically. For restricted resources, administrators can approve access, deny it with a stated reason, or revoke

previously granted rights. Administrators also retain the ability to revoke access to public resources when necessary.

- A catalog view lists all datasets/services with summary information.
- An accounts directory summarizes each user (AUA status, organization, number of resources with access) and links to a per-user details page.

There are three main components to our SSO system:

- Keycloak, a feature-rich open source Identity and Access Management platform that offers many powerful features for Single Sign-On (SSO) systems.
- `caida_oidc_service`, an optional reverse web proxy that offloads authorization and authentication from the applications
- Resource Access Management (RAM) service, which presents a web interface for users to request access to resources and for admins to vet users and grant access

We considered several approaches for implementing authentication and authorization across services:

1. A library invoked by each service (requires modifying all services).
2. A module in the web server container (requires services to run inside such containers).
3. A reverse proxy between end-user and service (zero changes to existing services, uniform configuration).

At the end we adopted the reverse proxy model, as it requires no modifications to existing services and enables uniform configuration. The resulting component acts as a reverse web proxy that provides the application-side implementation of authentication and authorization. It integrates with CAIDA's identity provider and supports both interactive web browser sessions and programmatic access by scripts and applications. Unauthenticated users are redirected to the identity provider to log in; authenticated users are granted access based on their assigned roles. Scripts and automated tools can use tokens issued by the identity provider to obtain access, including renewable offline tokens where appropriate. It is built on top of OpenResty (Nginx + LuaJIT + useful Lua libraries), including `lua-resty-openidc`, which implements OIDC.

Resource Access Management (RAM)

The RAM service provides a web interface for users to request access and for admins to vet and grant it.

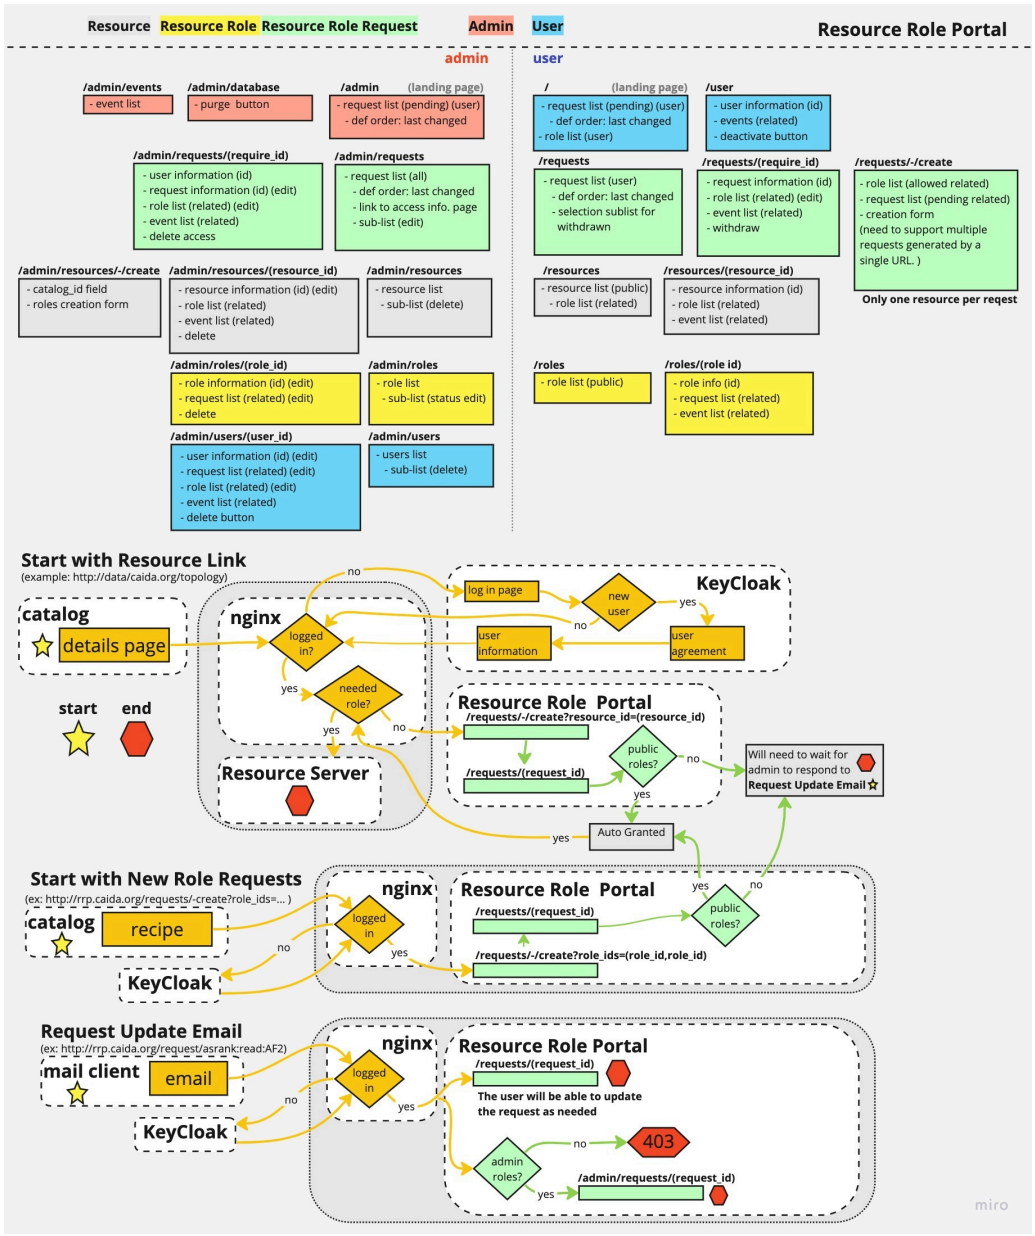

Figure 10.1: Resource Access Management Architecture

- Users are directed to RAM when accessing protected resources.
- Login via Keycloak is required, with additional organizational info.

- Users request roles for specific resources, agree to Acceptable Use Agreement (AUA), and submit their request.
- Public resources: requests are automatically granted (to ensure AUA agreement and usage tracking).
- Restricted resources: requests are reviewed by Data Admins, who may approve, deny, or request more info.
- RAM also supports role expiration: admins can set an expiry, after which roles are revoked unless extended.

From the user's perspective, the process is as follows. A new user registers, signs in once, accepts the current AUA, and then browses a list of datasets, APIs, and applications. Each RAM resource is described in details in CAIDA catalog including its access model (roles), and citation/DOI information. Users submit a brief data-use justification to request a role; most requests are approved automatically according to policy. The portal keeps users informed of request status and provides a clear path back to resource documentation and citation guidance. Alternatively, a user may enter RAM directly from the catalog entry of any resource. The link from the catalog entry leads either to the resource's location on disk or to its API (if the user is already logged into RAM and a corresponding access role is granted), or will prompt the user to log in or register in RAM and request a corresponding role before access is granted.

Chapter 11

The future of Internet measurement

The focus of this project has been security vulnerabilities in the Internet itself, including attacks on the global routing system (BGP route hijacks), abuse of the DNS, vulnerabilities in the Certificate Authority system, lack of best practice in management of Internet addresses, and distributed denial of service (DDoS) attacks. Data on each of these issues is central to understanding the severity of the threats, the magnitude of the resulting harms, trends in mitigation, and options for regulatory intervention or investment into R&D solutions. Underlying data on Internet topology is critical as a baseline to understand how these vulnerabilities manifest and propagate, and how changes in topology may increase or decrease the potential for systemic risks. In this Design Phase we identified a wide variety of needed data types, including:

- ISPs properties: deploy Source Address Validation (SAV), IPv6, DNSSEC, ROV
- BGP routing announcements and associated ROAs (authorizations)
- Ownership of IP addresses and autonomous systems (ASes)
- Peering arrangements at Internet exchanges.
- AS relationships (transit, peering, etc.)
- Tactical blocklists at AS and IP level.
- DNS zone files and changes
- Extent of DNSSEC deployment
- Assessment of actual harms from abuse of BGP
- Census of open DNS resolvers
- DNS query samples (passive DNS)
- DNS name registration information
- Evidence of malicious DNS names
- Registries and registrars with history of facilitating abusive name registration.
- DNS pricing information

- TLS Certificates in use
- Birth and death of certificates.
- Certificates registered in Certificate Transparency logs
- Certificate authorities with history of malicious or incompetent registrations.

We provide this summary list as a quick illustration of the diversity of data that we have identified as relevant and important today. The two key challenges associated with collecting these elements are: the voluntary and international nature of data collection, limiting the value and effectiveness of a national regulatory approach; and the need for historical data to establish baselines and track trends. Together these two challenges imply a third fundamental challenge: finding a sustainable funding model for the effort.

Voluntary and international nature of data collection. Some data collection, such as active measurement and BGP, requires that ISPs volunteer to connect equipment and contribute a vantage point (and BGP data) to the data feed. ISPs are under no obligation to provide this source of data. The success of projects such as RouteViews depends on a level of trust the platform operator has developed with ISPs across the globe. If RouteViews evolved into a more institutionalized organization, perhaps with more substantial funding from the U.S. government, it might erode the trust basis on which RouteViews succeeds today.

This challenge is a key reason that an emphasis in our Design Phase was to codify international partnerships with researchers in Europe and Asia who share our vision of large-scale Internet measurement infrastructure with resulting data being widely shared to enable scientific research. Our international partners have complementary expertise and are supporting synergistic infrastructure projects, funded by their own institutions or national funding agencies.

Need for historical data. Operators are primarily concerned with real-time data—what is happening on the Internet now. Researchers and policy makers are equally likely to want historical data, in order to map trends and assess progress toward objectives. This requirement implies the need for a substantial data repository, continuity of operations, and effective tools for data access. Commercial organizations that collect data intended to support network operations are not likely to undertake this sort of curation. Indeed, they come to CAIDA, Route Views, and RIPE for historical data.

Need for stable funding model. The diverse and often unstable sources of data collection and curation of such data 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.

11.1 Sustainability of data collection in other fields

Table 11.1: Examples of U.S. Domain-Specific Data / Statistical Organizations

Organization	**Funding Sources**	**Headcount/ Budget**	**Contribution**
U.S. National Agricultural Library	USDA	215	Curated agricultural and life-sciences information/archives.
Center for Education Statistics (NCES)	Dept Education	$306M; 84	National education data to drive/evaluate policy.
U.S. National Climatic Data Center	NOAA	N/A	Preserves global climate and environmental records.
Centers for Medicare & Medicaid Services (CMS)	HHS	6,710; $1,516B	Massive health insurance/health data systems.
Software Engineering Institute	DoD (FFRDC)	700 ; $584 M (multi-year)	R&D engine in software, cybersecurity.
National Center for Biotechnology Information (NCBI)	NIH / NLM	—	Critical bioinformatics infrastructure (GenBank, PubMed, genome browsers).
Bureau of Transportation Statistics	DoT	70	Transportation data for policy, safety, economic analysis.
National Centers for Environmental Information (NCEI)	NOAA	500/$71M	Environmental and earth-system data for climate, oceanographic, and geophysical research.
Genomic Data Commons (GDC)	NIH / NCI	—	Cancer genomic and clinical data for reproducible research.
U.S. Census Bureau	D. Commerce	$1.382B ; 8500	Foundational demographic, economic, housing data
Energy Information Administration	DOE	$135M (FY24)	Authoritative data/ forecasts for U.S. energy systems.
National Center for Health Statistics	CDC/HHS	$187.4M (FY24)	National vital statistics, health and survey data underpinning public health analysis.
U.S. Geological Survey (USGS)	DOI	$1.6B; 8000	Geospatial, hydrologic, and hazard data for infrastructure planning and risk mitigation.

The need for organized support for data collection, curation and analysis is recognized in many fields. In fields where the users of this data are primarily academic re-

search institutions, funding for these centers is often funded by a governmental research funding organization. In other cases, the funding may come from a government agency with a more operational charter. Table 11.1 lists several organizations that archive and share data for different societal systems, including the size (measured in head count), the number of years in operation, and the relation to the funding government. They share the characteristic that they have been stably funded for years.

11.2 Institutional history for the Internet

An organization to regularize data collection, curation and analysis of data would not be the first organization that has been created to help shape the Internet. In the early days of the Internet, various institutions were put in place to facilitate, manage or govern activities that were key to its growth. In many cases, these institutions were formalizations of the early stewardship provided by the original creators of the Internet. The Internet Engineering Task Force (IETF) was created to structure and facilitate the development of Internet standards. The Internet Society was later created to provide a legally recognized construct within which the IETF could be housed. ICANN was created to oversee the management and allocation of Internet addresses and Domain names, which had previously been done by Jon Postel, part of the original team of creators.

All such institutions are creatures of their time, with a charter and governance structure established with as much wisdom as the creators could provide. The authority or standing to create them was in large part still located in the informal team of original creators of the Internet.

As the Internet has matured, we have seen the creation of further institutions, and the push of some existing institutions to assert their right to play some role in shaping key activities around the Internet. The International Telecommunications Union (which is older than the UN but is currently housed there) has been pushing for a role in Internet governance. The United Nations set up the Internet Governance Forum as a venue for international stakeholder discussion. More operationally, a consortium of industry actors led by Google set up the Certificate Authority Browser Forum (CA/B) to provide oversight and policing of the CA system. The five Regional Internet Registries have been taking on increasing roles related to the Internet security (for example the RPKI) as well as their original mission allocating of IP addresses. The IETF created a research forum called the Internet Research Task Force.

Some of the decisions taken by these early institutions shaped the landscape of entities that actually provide the core services of the Internet. For example, ICANN created the competitive landscape of registrars and registries for the Domain Name System.

11.3 A possible organizational structure

There are different aspect to the objective we have discussed here, and we conclude that different organizations may be best suited to the different aspects.

Data curation requires infrastructure for storage, staff to maintain that infrastructure, and tools and procedures for the inflow and outflow of data. Setting up perhaps three such institutions across the globe, with the objective that data would be replicated and shared among them, would lessen the sense that one nation was trying to dominate the undertaking.

Data collection requires skills specific to the data being collected. Collecting routing data requires a different set of tools and relationships than collecting DNS data. Different organizations could take on these tasks, exploiting their task-centric skills but relieving them of the complexity of long term data curation and access. Many organizations could undertake similar data collection in different parts of the world and share that data through the data curation organizations.

Planning and conventions will be required to bring all this together. Collective decisions must be made regarding which data is worth collecting and preserving, what the interchange standards are for this data, procedures for making it available, and so on.

Data analysis must be a part of this endeavor. Collecting data with the hope that someone will look at it will not be sufficient to justify this effort. While the objective would be that many organizations, research units, governments and so on will use the data, having one or more groups with the specific mission to demonstrate the utility of the data by actually exploiting it must be a part of the overall scheme. Different existing groups could evolve to take on some of these roles, and focusing on the different roles and their requirements may help to map a path to the future.

Bibliography

[1] Emile Aben. Route Collection at the RIPE NCC - Where are we and where should we go?, Oct 2020. https://labs.ripe.net/author/emileaben/route-collection-at-the-ripe-ncc-where-are-we-and-where-should-we-go/.

[2] David Abramson and Manish Parashar. Translational research in computer science. *Computer*, 52(9):16–23, 2019.

[3] Gautam Akiwate, Raffaele Sommese, Mattijs Jonker, Zakir Durumeric, kc claffy, Geoffrey M. Voelker, and Stefan Savage. Retroactive identification of targeted dns infrastructure hijacking. In *ACM Internet Measurement Conference (IMC)*, Oct 2022.

[4] Taha Albakour, Oliver Gasser, Robert Beverly, and Georgios Smaragdakis. Third Time's Not a Charm: Exploiting SNMPv3 for Router Fingerprinting. In *Proceedings of the 21st ACM Internet Measurement Conference*, IMC '21, page 150–164, New York, NY, USA, 2021. Association for Computing Machinery.

[5] Taha Albakour, Oliver Gasser, Robert Beverly, and Georgios Smaragdakis. Illuminating router vendor diversity within providers and along network paths. In *Proceedings of the 2023 ACM on Internet Measurement Conference*, IMC '23, page 89–103, New York, NY, USA, 2023. Association for Computing Machinery.

[6] Thibault Alfroy, Thomas Holterbach, Thomas Krenc, kc claffy, and Cristel Pelsser. Internet Science Moonshot: Expanding BGP Data Horizons. In *22nd ACM Workshop on Hot Topics in Networks*, Nov 2023.

[7] Thomas Alfroy, Thomas Holterbach, Thomas Krenc, kc claffy, and Cristel Pelsser. The next generation of bgp data collection platforms. In *ACM SIGCOMM 2024 Conference*, Aug 2024.

[8] Tim April and Warren Kumari. Dns transparency, 2017. https://www.internetfire.org/projects/dns-transparency.

[9] Tim April and Warren Kumari. Dns transparency architecture overview, 2022. https://docs.google.com/document/d/1mRou-0lFElXUNnfRWSkyOfbautwOA5I_sAw5oEyZvsE/edit?tab=t.0#heading=h.qyfowjfgh10p.

[10] Alexander Azimov, Eugene Bogomazov, Randy Bush, Keyur Patel, and Job Snijders. Verification of AS PATH Using the Resource Certificate Public Key, Infrastructure and Autonomous System Provider Authorization. https://www.ietf.org/archive/id/draft-ietf-sidrops-aspa-verification-01.txt, July 2019.

[11] Fred Baker. Requirements for ip version 4 routers. RFC 1812, Jun 1995. https://datatracker.ietf.org/doc/html/rfc1812.

[12] 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.

[13] K. Benson, A. Dainotti, KC Claffy, and E. Aben. Gaining insight into as-level outages through analysis of internet background radiation. In *Traffic Monitoring and Analysis Workshop (TMA)*, pages 1–6, Torino, Italy, Apr 2013. TMA 2013.

[14] Karyn Benson, Alberto Dainotti, kc claffy, Alex C. Snoeren, and Michael Kallitsis. Leveraging Internet Background Radiation for Opportunistic Network Analysis. In *Internet Measurement Conference (IMC)*, Oct 2015.

[15] Robert Beverly, Ryan Koga, and kc claffy. Initial longitudinal analysis of ip source spoofing capability on the internet. *Internet Society*, Jul 2013.

[16] Big Data Interagency Working Group. Innovating the Data Ecosystem: An Update of the Federal Big Data Research and Development Strategic Plan. Technical report, National Science and Technology Council, Nov 2024. https://www.nitrd.gov/pubs/Big-Data-Strategic-Plan-2024.pdf.

[17] Henry Birge-Lee, Liang Wang, Jennifer Rexford, and Prateek Mittal. SICO: Surgical Interception Attacks by Manipulating BGP Communities. In *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2019.

[18] Laura J. Biven and Amy Walton. Big data: Pioneering the future of federally supported data repositories, Feb 2022. Workshop Report.

[19] Dionysus Blazakis, Manish Karir, and John S Baras. Bgp-inspect-extracting information from raw bgp data. In *2006 IEEE/IFIP Network Operations and Management Symposium NOMS 2006*, pages 174–185. IEEE, 2006.

[20] Larry Blunk, Manish Karir, and Craig Labovitz. Multi-Threaded Routing Toolkit (MRT) Routing Information Export Format. RFC 6396, IETF, October 2011.

[21] Brendon Jones. Passive 100GB Trace Statistics, 2024. https://www.caida.org/catalog/datasets/trace_stats/.

[22] Brendon Jones and Dan Andersen and Leo Pascual. CAIDA 100GB traffic monitor deployment details, 2023. https://gitlab.caida.org/CAIDA/sysadmin/100g-monitor/.

[23] CAIDA. Memorandum of Agreement Between the Center for Applied Internet Data Analysis (CAIDA) at UCSD and TU Dresden, 2025. https://www.caida.org/funding/msri-gmi3s/reports/TU_Dresden_MoA.pdf.

[24] Cali Dog Security. Axeman is a utility to retrieve certificates from Certificate Transparency Lists (CTLs), 2018. https://github.com/CaliDog/Axeman.

[25] Cali Dog Security. Certificate Transparency Log aggregation, parsing, and streaming service written in Elixir, 2018. https://github.com/CaliDog/certstream-server/.

[26] Massimo Candela. Easy bgp monitoring with bgpalerter. Ripe labs article, RIPE, Apr 2020. https://labs.ripe.net/author/massimo_candela/easy-bgp-monitoring-with-bgpalerter/.

[27] Justin Cappos, Matthew Hemmings, Rick McGeer, Albert Rafetseder, and Glenn Ricart. Edgenet: A global cloud that spreads by local action. In *2018 IEEE/ACM Symposium on Edge Computing (SEC)*, pages 359–360, 2018.

[28] Esteban Carisimo, Caleb J. Wang, Mia Weaver, Fabián Bustamante, and Paul Barford. A hop away from everywhere: A view of the intercontinental long-haul infrastructure. *ACM SIGMETRICS*, 2024.

[29] Censys. Opt Out of Data Collection. https://support.censys.io/hc/en-us/articles/360043177092-Opt-Out-of-Data-Collection, 2022.

[30] Kai Chen, Chengchen Hu, Wenwen Zhang, Yan Chen, and Bin Liu. On the Eyeshots of BGP Vantage Points. In *GLOBECOM '09*, 2009.

[31] Laurent Chuat, Cyrill Krähenbühl, Prateek Mittal, and Adrian Perrig. F-PKI: Enabling Innovation and Trust Flexibility in the HTTPS Public-Key Infrastructure. In *Network and Distributed System Security Symposium (NDSS)*, 2022.

[32] CIDR. CIDR REPORT, 2023. https://www.cidr-report.org/as2.0/.

[33] k claffy and D Clark. Challenges in measuring the Internet for the public Interest. *Journal of Information Policy*, 12, May 2022.

[34] kc claffy. Workshop on internet economics (wie2009) report. *ACM SIGCOMM Computer Communication Review (CCR)*, 40(2), Apr 2010.

[35] kc claffy. The 3rd Workshop on Active Internet Measurements (AIMS-3) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, 41(3), Jul 2011. http://www.caida.org/publications/papers/2011/aims_report/.

[36] kc claffy. The 4th Workshop on Active Internet Measurements (AIMS-4) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, 42(3), Jul 2012. http://www.caida.org/publications/papers/2012/aims4_report/.

[37] kc claffy. Workshop on internet economics (wie2011) report. *ACM SIGCOMM Computer Communication Review (CCR)*, 42(2):110–114, Apr 2012.

[38] kc claffy. The 5th Workshop on Active Internet Measurements (AIMS-5) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, 43(3), Jul 2013. http://www.caida.org/publications/papers/2013/aims5_report/.

[39] kc claffy. The 6th Workshop on Active Internet Measurements (AIMS-6) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, 44(5), Oct 2014. http://www.caida.org/publications/papers/2014/aims6_report/.

[40] kc claffy. The 7th Workshop on Active Internet Measurements (AIMS-7) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, 46(1), Jan 2015. http://www.caida.org/publications/papers/2016/aims2015_report/.

[41] 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/.

[42] 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/.

[43] 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/.

[44] 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/.

[45] kc claffy. Internet data science for cybersecurity, Jan 2023. https://cseweb.ucsd.edu/classes/wi23/cse291-e/.

[46] kc claffy, Emile Aben, Jordan Augé, Robert Beverly, Fabian Bustamante, Benoit Donnet, Timur Friedman, Marina Fomenkov, Peter Haga, Matthew Luckie, and Yuval Shavitt. The 2nd Workshop on Active Internet Measurements (AIMS-2) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, 40(5), Oct 2010. http://www.caida.org/publications/papers/2010/aims_report/.

[47] kc claffy and David Clark. Workshop on internet economics (wie2012) report. *ACM SIGCOMM Computer Communication Review (CCR)*, 43(3):95–100, Jul 2013.

[48] kc claffy and David Clark. Workshop on internet economics (wie2013) report. *ACM SIGCOMM Computer Communication Review (CCR)*, 44(3):116–119, Jul 2014.

[49] kc claffy and David Clark. Workshop on internet economics (wie2014) report. *ACM SIGCOMM Computer Communication Review (CCR)*, 45(3):43–48, Jul 2015.

[50] kc claffy and David Clark. Workshop on internet economics (wie2015) report. *ACM SIGCOMM Computer Communication Review (CCR)*, Jul 2016.

[51] kc claffy and David Clark. Workshop on internet economics (wie2016) report. *ACM SIGCOMM Computer Communication Review (CCR)*, Jul 2017.

[52] kc claffy and David Clark. Workshop on internet economics (wie2017) report. *ACM SIGCOMM Computer Communication Review (CCR)*, Jul 2018.

[53] kc claffy and David Clark. Workshop on internet economics (wie2018) report. *ACM SIGCOMM Computer Communication Review (CCR)*, Apr 2019.

[54] kc claffy and David Clark. Workshop on internet economics (wie2019) report. *ACM SIGCOMM Computer Communication Review (CCR)*, Apr 2020.

[55] kc claffy and David Clark. Workshop on internet economics (wie2020) report. *ACM SIGCOMM Computer Communication Review (CCR)*, Apr 2021.

[56] kc claffy, Marina Fomenkov, Ethan Katz-Bassett, Robert Beverly, Benjamin Cox, and Matthew Luckie. The workshop on active internet measurements (aims) report. *ACM SIGCOMM Computer Communication Review (CCR)*, 39(5), Oct 2009.

[57] kc claffy, Matthew Luckie, and Josh Polterock. Aspire project final report, Oct 2020.

[58] D Clark, S Garfinkel, and k claffy. Differential Privacy, Firm-level Data and the Binomial Pathology. *IEEE Security & Privacy*, 23(1):2–10, September 2024.

[59] D Clark, S Garfinkel, and k claffy. Exploring the Limits of Differential Privacy. In *Telecommunications Policy Research Conference (TPRC)*, September 2024.

[60] D Clark, C Testart, M Luckie, and k claffy. A path forward: Improving Internet routing security by enabling zones of trust. In *Telecommunications Policy Research Conference (TPRC)*, September 2024.

[61] David D. Clark, Cecilia Testart, Matthew Luckie, and kc claffy. A path forward: improving internet routing security by enabling zones of trust. *Journal of Cybersecurity*, 10(1), Dec 2024.

[62] Michael Collins. Acknowledged Scanners. `https://gitlab.com/mcollins_at_isi/acknowledged_scanners`, 2021.

[63] 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`.

[64] Federal Communications Commission. Fcc launches inquiry into internet routing vulnerabilities. `https://docs.fcc.gov/public/attachments/FCC-22-18A1.pdf`, 2022.

[65] 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`.

[66] 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.

[67] Berat Can Şenel, Maxime Mouchet, and Justin Cappos and. EdgeNet: A multi-tenant and multi-provider edge cloud. In *EdgeSys*, pages 49–54, Apr 2021.

[68] Alberto Dainotti, Roman Amman, Emile Aben, and Kimberly Claffy. Extracting benefit from harm: using malware pollution to analyze the impact of political and geophysical events on the internet. *SIGCOMM Comput. Commun. Rev.*, 42, Jan 2013.

[69] Alberto Dainotti, Karyn Benson, Alistair King, kc claffy, Michael Kallitsis, Eduard Glatz, and Xenofontas Dimitropoulos. Estimating Internet Address Space Usage through Passive Measurements. *SIGCOMM Comput. Commun. Rev.*, 44(1), Dec 2014.

[70] Alberto Dainotti, Karyn Benson, Alistair King, Bradley Huffaker, Eduard Glatz, Xenofontas Dimitropoulos, Philipp Richter, Alessandro Finamore, and Alex C. Snoeren. Lost in space: Improving inference of ipv4 address space utilization. *IEEE Journal on Selected Areas in Communications (JSAC)*, 34(6):1862–1876, Jun 2016.

[71] Alberto Dainotti, Alistair King, kc claffy, Ferdinando Papale, and Antonio Pescapè. Analysis of a "/0" Stealth Scan from a Botnet. In *Internet Measurement Conference (IMC)*, Nov 2012.

[72] Alberto Dainotti, Claudio Squarcella, Emile Aben, Kimberly C. Claffy, Marco Chiesa, Michele Russo, and Antonio Pescapé. Analysis of Country-wide Internet Outages Caused by Censorship. In *Internet Measurement Conference (IMC)*, pages 1–18, Berlin, Germany, Nov 2011. ACM.

[73] Omar Darwich, Hugo Rimlinger, Milo Dreyfus, Matthieu Gouel, and Kevin Vermeulen. Replication: Towards a publicly available Internet scale IP. In *IMC*, Oct 2023.

[74] B Degen, B Du, R Mok, R Sommese, M Jonker, R Van Rijswijk-Deij, and k claffy. From Scarcity to Opportunity: Examining Abuse of the IPv4 Leasing Market. In *Network Traffic Measurement and Analysis Conference (TMA)*, page 11, March 2025.

[75] Samuel DeLaughter and Karen Sollins. SYN Proof-of- Work: Improving Volumetric DoS Resilience in TCP. In *2025 IEEE Symposium on Security and Privacy (SP)*, pages 1877–1890, 2025.

[76] White House Office of the National Cyber Director. Roadmap to enhancing internet routing security, Sep 2024.

[77] B Du, G Akiwate, T Krenc, C Testart, A Marder, B Huffaker, A Snoeren, and k claffy. IRR Hygiene in the RPKI Era. In *Passive and Active Measurement Conference (PAM)*, March 2022.

[78] B Du, R Fontugne, C Testart, A Snoeren, and k claffy. Sublet Your Subnet: Inferring IP Leasing in the Wild. In *ACM Internet Measurement Conference (IMC)*, November 2024.

[79] Ben Du, Katherine Izhikevich, Sumanth Rao, Gautam Akiwate, Cecilia Testart, Alex C. Snoeren, and kc claffy. IRRegularities in the Internet Routing Registry. In *ACM Internet Measurement Conference (IMC)*, Oct 2023.

[80] Ben Du, Cecilia Testart, Romain Fontugne, Gautam Akiwate, Alex C. Snoeren, and kc claffy. Mind your manrs: Measuring the manrs ecosystem. In *ACM Internet Measurement Conference (IMC)*, Oct 2022.

[81] Ben Du, Cecilia Testart, Romain Fontugne, Alex C. Snoeren, and kc claffy. Taking the low road: How rpki invalids propagate. In *ACM SIGCOMM Poster*, Sep 2023.

[82] Zakir Durumeric, Eric Wustrow, and J. Alex Halderman. Zmap: Fast internet-wide scanning and its. In *USENIX Security Symposium*, 2013.

[83] Isabell Egloff, Raphael Hiesgen, Maynard Koch, Thomas C. Schmidt, and Matthias Wählisch. A detailed measurement view on ipv6 scanners and their adaption to bgp signals. *Proceedings of the ACM on Networking (PACMNET)*, 3(CoNEXT3), Sep 2025.

[84] Elena Yulaeva. Seeking Beta Users for 100 GB link Anonymized Passive Traces, August 2024. `https://blog.caida.org/best_available_data/2024/08/11/seeking-beta-users-for-100-gb-link-anonymized-passive-traces/`.

[85] P. Ferguson and D. Senie. Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing. RFC 2827, IETF, May 2000.

[86] Romain Fontugne. Internet health report, 2025. `https://ihr.iijlab.net`.

[87] Romain Fontugne and Emile Aben. Internet yellow pages, 2025. `https://iyp.iijlab.net`.

[88] Center for Applied Internet Data Analysis (CAIDA). Active internet measurement systems. `https://www.caida.org/workshops/?workshopserieslisting=AIMS`.

[89] Center for Applied Internet Data Analysis (CAIDA). STARDUST Grafana dashboard. https://explore.stardust.caida.org.

[90] Center for Applied Internet Data Analysis (CAIDA). Workshop on internet economics. https://www.caida.org/workshops/?workshopserieslisting=WIE.

[91] Center for Applied Internet Data Analysis (CAIDA). Workshop on overcoming barriers to internet research. https://www.caida.org/workshops/?workshopserieslisting=WOMBIR.

[92] Center for Applied Internet Data Analysis (CAIDA). Corsaro. https://catalog.caida.org/media/2012_dust_corsaro, 2012.

[93] Center for Applied Internet Data Analysis (CAIDA). Information marketplace for policy and analysis of cyber-risk & trust (impact) project, 2016. https://www.caida.org/projects/impact/.

[94] Center for Applied Internet Data Analysis (CAIDA). Internet outage detection and analysis (ioda), 2018. https://www.caida.org/projects/ioda/.

[95] Center for Applied Internet Data Analysis (CAIDA). Caida's sustainable tools for analysis and research on darknet unsolicited traffic (stardust) project, 2021. https://www.caida.org/projects/stardust/.

[96] Center for Applied Internet Data Analysis (CAIDA). Network telescope documentation: Flowtuple. https://www.caida.org/projects/network_telescope/docs/data/flowtuple/, 2021.

[97] Center for Applied Internet Data Analysis (CAIDA). RouteViews IPv4 Prefix to AS mappings dataset. https://catalog.caida.org/details/dataset/routeviews_ipv4_prefix2as, 2021.

[98] Center for Applied Internet Data Analysis (CAIDA). Caida's scalable technology to accelerate research network operations vulnerability alerts (starnova) project, 2024. https://www.caida.org/projects/starnova/.

[99] Center for Applied Internet Data Analysis (CAIDA). Caida's resource catalog, network telescope collection, papers listing, 2025. https://catalog.caida.org/search?query=links=collection:ucsd_telescope_datasets%20types=paper.

[100] Center for Applied Internet Data Analysis (CAIDA). Caida's resource catalog, network telescope non-caida papers listing, 2025. https://catalog.caida.org/search?query=links%3Dcollection%3Aucsd_telescope_datasets%20types%3Dpaper%20!links%3Dtag%3Acaida.

[101] Center for Applied Internet Data Analysis (CAIDA). Caida's resource catalog, topology data collection, papers listing, 2025. https://catalog.caida.org/search?query=links=collection:caida_topology_datasets%20types=paper.

[102] Center for Applied Internet Data Analysis (CAIDA). Global measurement infrastructure: workshops, 2025.

[103] Center for Applied Internet Data Analysis (CAIDA). Hosting an ark monitor: Archipelago memorandum of cooperation (moc) between hosting sites and caida, 2025. https://www.caida.org/projects/ark/moc/.

[104] Center for Applied Internet Data Analysis (CAIDA). Scamper python documentation. https://www.caida.org/catalog/software/scamper/python/, 2025.

[105] CA/Browser Forum. Baseline requirements for the issuance and management of publicly-trusted certificates. https://cabforum.org/baseline-requirements-documents/, 2018. Accessed: 2025-07-08.

[106] Electronic Frontier Foundation. How root certificates work and why they're dangerous. https://www.eff.org/deeplinks/2019/02/security-compromising-actions-should-have-big-flashing-lights, 2015. Accessed: 2025-07-08.

[107] National Science Foundation. Transforming science through cyberinfrastructure, Feb 2019. Draft of NSF's Blueprint for a National Cyberinfrastructure Ecosystem.

[108] National Science Foundation. Internet measurement research: Methodologies, tools, and infrastructure (imr). https://new.nsf.gov/funding/opportunities/imr-internet-measurement-research-methodologies-tools, 2022.

[109] National Science Foundation and Department of Defense. Nsf convergence accelerator 2022 joint nsf/dod phases 1 and 2 for track g: Securely operating through 5g infrastructure. Technical report, National Science Foundation, Dec 2022.

[110] V. Fuller and T. Li. Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan. https://datatracker.ietf.org/doc/html/rfc4632, Aug 2006.

[111] ESFRI Drafting Group on Research Infrastructures Funding. Esfri report on funding of research infrastructures, Jun 2024. https://zenodo.org/records/14770891.

[112] M Gao, R Mok, E Carisimo, k claffy, E Li, and S Kulkarni. DarkSim: A Similarity-Based Time Series Analytic Framework for Darknet Traffic. In *ACM Internet Measurement Conference (IMC)*, November 2024.

[113] Max Gao, Ricky Mok, and kc claffy. A Scalable Network Event Detection Framework for Darknet Traffic. In *ACM Internet Measurement Conference (IMC) Poster*, Oct 2022.

[114] Nikolas Gauder. Quic in scamper integration update meeting, 2025.

[115] Deutsche Telekom Security GmbH. T-Pot - The all in one multi honeypot platform. `https://github.com/telekom-security/tpotce`.

[116] Google. Measurement lab (m-lab). `http://measurementlab.net`, Jan 2009.

[117] Grafana Labs. Grafana,. `https://grafana.com/grafana/`.

[118] Wonkyu Han, Ziming Zhao, Adam Doupe, and Gail-Joon Ahn. HoneyMix: Toward SDN-based intelligent honeynet. In *Proceedings of ACM SDN-NFV Security*, mar 2016.

[119] Bill Herrin. Caida's mrt tools, 2024. `https://github.com/CAIDA/mrt-tools`.

[120] Raphael Hiesgen, Marcin Nawrocki, Marinho Barcellos, Daniel Kopp, Oliver Hohlfeld, Echo Chan, Roland Dobbins, Christian Doerr, Christian Rossow, Daniel R. Thomas, Mattijs Jonker, Ricky Mok, Xiapu Luo, John Kristoff, Thomas C. Schmidt, Matthias Wählisch, and kc claffy. The Age of DDoScovery: An Empirical Comparison of Industry and Academic DDoS Assessments. In *ACM Internet Measurement Conference (IMC)*, Nov 2024.

[121] Raphael Hiesgen, Marcin Nawrocki, Alistair King, Alberto Dainotti, Thomas C. Schmidt, and Matthias W ahlisch. Spoki: Unveiling a new wave of scanners through a reactive network telescope. In *Proceedings of USENIX Security Symposium*, 2022. `https://catalog.caida.org/paper/2022_spoki`.

[122] Fahad Hilal and Oliver Gasser. Yarrpbox: Detecting middleboxes at internet-scale. *Proc. ACM Netw.*, 1(CoNEXT1), Jul 2023. `https://doi.org/10.1145/3595290`.

[123] Holterback, Thomas and Alfroy, Thomas. GILLnet prototype: BGP Route Collection System, 2024. `https://bgproutes.quest`.

[124] U.S. White House. 2024 report on the cybersecurity posturre of the united states, May 2024. https://bidenwhitehouse.archives.gov/wp-content/uploads/2024/05/2024-Report-on-the-Cybersecurity-Posture-of-the-United-States.pdf.

[125] Bradley Huffaker, Romain Fontugne, Alexander Marder, and kc claffy. On the importance of being an as: An approach to country-level as rankings. In *ACM Internet Measurement Conference (IMC)*, Oct 2023.

[126] Bradley Huffaker, Daniel Plummer, David Moore, and kc Claffy. Topology discovery by active probing. In *SAINT*, Nara City, Japan, Jan 2002. http://www.caida.org/outreach/papers/2002/SkitterOverview/.

[127] Christian Huitema, Sara Dickinson, and Allison Mankin. DNS over Dedicated QUIC Connections. RFC 9250, IETF, May 2022.

[128] Gokay Huz, Steven Bauer, kc claffy, and Robert Beverly. Experience in using mturk for network measurement. In *ACM SIGCOMM Workshop on Crowdsourcing and crowdsharing of Big (Internet) Data (C2B(I)D)*, Aug 2015.

[129] ICANN. VeriSign Application for Registry Service: "Rapid Zone Updates", 2007. https://www.icann.org/en/system/files/files/memo-dns-update-service.pdf.

[130] ICANN. Centralized zone directory service, 2019. https://cdsz.icann.org/.

[131] InfluxData. InfluxDB. https://www.influxdata.com.

[132] CAIDA's Macroscopic Internet Topology Data Kit (ITDK). http://www.caida.org/data/active/internet-topology-data-kit/.

[133] Ebrima Jaw, Thomas Krenc, Moritz Muller, kc claffy, Lambert J.M. Nieuwenhuis, and Cristian Hesselman. Noisy neighbours: Keep the neighbourhood quiet. In *21st International Conference on Network and Service Management*, 2025.

[134] Joe Abley and Thibault Meunier and Vicky Shrestha and Bas Westerbaan. Addressing the unauthorized issuance of multiple TLS certificates for 1.1.1.1, September 2025.

[135] Brendon Jones, Bill Herrin, and Matthew Luckie. natp-ssh, 2023. https://gitlab.caida.org/CAIDA/ark/natp-ssh.

[136] Sangeetha Abdu Jyothi. Solar superstorms: planning for an Internet apocalypse. In *Proceedings of the 2021 ACM SIGCOMM Conference*, 2021. https://doi.org/10.1145/3452296.3472916.

[137] Sangeetha Abdu Jyothi. Characterizing the Role of Power Grids in Internet Resilience, Jun 2023. https://arxiv.org/abs/2306.02502.

[138] Sunil Kalidindi and Matthew J. Zekauskas. Surveyor: An infrastructure for Internet performance. In *INET*, San Jose, CA, Jun 1999.

[139] Robert Kisteleki. RIPE NCC Measurement Data Retention Principles. *RIPE NCC*, 2023. https://labs.ripe.net/author/kistel/ripe-ncc-measurement-data-retention-principles/.

[140] Tess DeBlanc Knowles, Manish Parashar, Lynne Parker, Erwin Gianchandani, Daniela Braga, Mark E. Dean, Fei-Fei Li, Andrew Moore, Michael L. Norman, Frederick H. Streitz, Elham Tabassi, Sethuraman Panchanathan, and Arati Prabhakar. Strengthening and Democratizing the U.S. Artificial Intelligence Innovation Ecosystem: An Implementation Plan for a National Artificial Intelligence Research Resource, Jan 2023. https://www.whitehouse.gov/wp-content/uploads/2023/01/NAIRR-TF-Final-Report-2023.pdf.

[141] Lukas Kramer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. AmpPot: Monitoring and Defending Amplification DDoS Attacks. In *Proceedings of the 18th International Symposium on Research in Attacks, Intrusions and Defenses*, 2015.

[142] T Krenc, M Luckie, A Marder, and k claffy. Coarse-grained inference of bgp community intent. In *ACM Internet Measurement Conference (IMC)*, Oct 2023.

[143] Sarah Krouse, Dustin Volz, Aruna Viswanatha, and Robert 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.

[144] Marc Kuhrer, Thomas Hupperich, Christian Rossow, and Thorsten Holz. Hell of a Handshake: Abusing TCP for Reflective Amplification DDoS Attacks. In *Proceedings of the 2014 USENIX Workshop on Offensive Technologies (WOOT 14)*, 2014.

[145] Ben Laurie, Adam Langley, and Emilia Kasper. Certificate transparency. RFC 6962, IETF, 2013.

[146] Matthew Lepinski and Kotikalapudi Sriram. RFC 8205: BGPsec Protocol Specification, September 2017.

[147] 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.

[148] 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/`.

[149] Shihan Lin, Yi Zhou, Xiao Zhang, Todd Arnold, Ramesh Govindan, and Xiaowei Yang. Latency-Aware Inter-domain Routing, Oct 2024. `https://arxiv.org/abs/2410.13019`.

[150] Q. Lone, M. Luckie, M. Korczyński, H. Asghari, M. Javed, and M. van Eeten. Using crowdsourcing marketplaces for network measurements: The case of spoofer. In *Network Traffic Measurement and Analysis Conference (TMA)*, Jun 2018.

[151] Qasim Lone, Matthew Luckie, Maciej Korczyński, and Michel van Eeten. Using Loops Observed in Traceroute to Infer the Ability to Spoof. In *Passive and Active Measurement*, 2017.

[152] M Luckie, S Hariprasad, R Sommese, B Jones, K Keys, R Mok, and k claffy. An Integrated Active Measurement Programming Environment. In *Passive and Active Measurement Conference (PAM)*, December 2024.

[153] M. Luckie, B. Huffaker, and k. claffy. Learning regexes to extract router names from hostnames. In *ACM Internet Measurement Conference (IMC)*, Oct 2019.

[154] 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)*, Oct 2020.

[155] M Luckie, A Marder, B Huffaker, and k claffy. Learning regexes to extract network names from hostnames. In *Asian Internet Engineering Conference (AINTEC)*, Dec 2021.

[156] Matthew Luckie. Scamper: a scalable and extensible packet prober for active measurement of the Internet. In *ACM SIGCOMM Internet Measurement Conference (IMC)*, 2010.

[157] Matthew Luckie. Scamper, March 2023. `https://www.caida.org/tools/measurement/scamper/`.

[158] Matthew Luckie. Towards a Domain Specific Language for Internet Active Measurement, Jan 2024. https://blog.caida.org/best_available_data/2024/01/16/towards-a-domain-specific-language-for-internet-active-measurement/.

[159] Matthew Luckie. Understanding the Deployment of Public Recursive Resolvers, May 2024. https://blog.caida.org/best_available_data/2024/05/06/understanding-the-deployment-of-public-recursive-resolvers/.

[160] Matthew Luckie, Robert Beverly, Ryan Koga, Kim Keys, Joshua Kroll, and kc claffy. Network Hygiene, Incentives, and Regulation: Deployment of Source Address Validation in the Internet. In *ACM Computer and Communications Security (CCS)*, Nov 2019.

[161] Matthew Luckie, Bradley Huffaker, Alexander Marder, Zachary Bischof, Matthew Fletcher, and kc claffy. Learning to Extract Geographic Information from Internet Router Hostnames. In *ACM SIGCOMM Conference on Emerging Networking Experiments and Technologies (CoNEXT)*, Dec 2021.

[162] Matthew Luckie, Ken Keys, Ryan Koga, Rob Beverly, and kc Claffy. Spoofer source address validation measurement system, 2016. http://spoofer.caida.org.

[163] Matthew Luckie, Alexander Marder, Marianne Fletcher, Bradley Huffaker, and K. Claffy. Learning to extract and use ASNs in hostnames. In *Proceedings of ACM IMC*, 2020.

[164] Alexander Männel, Jonas Mücke, kc claffy, Max Gao, Ricky K. P. Mok, Marcin Nawrocki, Thomas C. Schmidt, and Matthias Wählisch. Lessons Learned from Operating a Large Network Telescope. In *ACM SIGCOMM 2025*, 2025.

[165] MANRS. BGP, RPKI, and MANRS: 2020 in review, 2021. https://www.manrs.org/2021/02/bgp-rpki-and-manrs-2020-in-review/.

[166] Alexander Marder, Zesen Zhang, Ricky Mok, Ramakrishna Padmanabhan, Bradley Huffaker, Matthew Luckie, Alberto Dainotti, kc claffy, Alex C. Snoeren, and Aaron Schulman. Access Denied: Assessing Physical Risks to Internet Access Networks. In *USENIX Security Symposium*, Aug 2023.

[167] Tony McGregor and Hans-Werner Braun. Balancing Cost and Utility in Active Monitoring: The AMP. In *INET*, Jul 2000.

[168] Alberto Medina, Mark Allman, and Sally Floyd. Measuring the Evolution of Transport Protocols in the Internet. *ACM SIGCOMM Computer Communication Review*, 35(2), Apr 2005.

[169] Lihua Miao, Wei Ding, and Haiting Zhu. Extracting Internet Background Radiation from RawTraffic Using Greynet. In *Proceedings of IEEE International Conference on Networks*, 2012.

[170] Alexandros Milolidakis, Tobias Bühler, Kunyu Wang, Marco Chiesa, Laurent Vanbever, and Stefano Vissicchio. On the effectiveness of bgp hijackers that evade public route collectors. *IEEE Access*, 2023.

[171] P. Mohapatra, J. Scudder, D. Ward, R. Bush, and R. Austein. BGP Prefix Origin Validation. RFC 6811, IETF, January 2013.

[172] D Moore, C Shannon, and J Brown. Code-red: a case study on the spread and victims of an internet worm. In *Internet Measurement Workshop (IMW)*, pages 273–284, Nov 2002.

[173] David Moore, Vern Paxson, Stefan Savage, Colleen Shannon, Stuart Staniford, and Nicholas Weaver. Inside the slammer worm. *IEEE Security and Privacy*, 1(4):33–39, Jul 2003.

[174] David Moore and Colleen Shannon. The Spread of the Witty Worm. *IEEE Security and Privacy*, 2(4):46–50, 2005.

[175] David Moore, Colleen Shannon, Douglas J. Brown, Geoffrey M. Voelker, and Stefan Savage. Inferring Internet Denial-of-Service Activity. *ACM Transactions on Computer Systems*, 24(2), May 2006. `http://doi.acm.org/10.1145/1132026.1132027`.

[176] David Moore, Geoffrey M. Voelker, and Stefan Savage. Inferring Internet Denial-of-Service Activity. *Usenix Security Symposium*, 2001.

[177] 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.

[178] L. Müller, M. Luckie, B. Huffaker, kc claffy, and M. Barcellos. Spoofed Traffic Inference at IXPs: Challenges, Methods and Analysis. *Computer Networks*, 182, Aug 2020.

[179] National Institute for Stantards and Technology. RPKI Deployment Monitor. `https://rpki-monitor.antd.nist.gov/`.

[180] National Science and Technology Council. Federal Cybersecurity Research and Development Strategic Plan: Implementation Roadmap, Dec 2019. https://www.nitrd.gov/pubs/Federal-Cybersecurity-RD-Strategic-Plan-2019.pdf.

[181] RIPE NCC. Ripe's routing information service. http://www.ris.ripe.net/.

[182] NLNOG. Ring. https://ring.nlnog.net/.

[183] NTT Global IP Network. BGPalerter, 2025.

[184] Department of Defense. Open Programmable Secure 5G (OPS-5G) Initiative. Technical report, Office of the Under Secretary of Defense for Research and Engineering, 2023. https://www.darpa.mil/research/programs/open-programmable-secure-5g.

[185] Forum of Incident Response and Security Teams (FIRST). DNS Abuse Techniques Matrix. https://www.first.org/global/sigs/dns/DNS-Abuse-Techniques-Matrix_v1.1.pdf.

[186] U.S. Government Accountability Office. GAO Information Technology Cybersecurity Team, Aug 2020. https://blog.gao.gov/2019/03/12/gaos-information-technology-cybersecurity-team/.

[187] Leo Oliver, Gautam Akiwate, Matthew Luckie, Ben Du, and kc claffy. Stop, DROP, and ROA: Effectiveness of Defenses through the Lens of DROP. In *ACM Internet Measurement Conference (IMC)*, Oct 2022.

[188] Ruoming Pang, Vinod Yegneswaran, Paul Barford, Vern Paxson, and Larry Peterson. Characteristics of Internet Background Radiation. In *Proceedings of the 4th ACM SIGCOMM Conference on Internet Measurement*, IMC '04, 2004.

[189] V. Paxson, J. Mahdavi, A. Adams, and M. Mathis. An architecture for large-scale Internet measurement. *IEEE Communications Magazine*, 36(8):48–54, 1998.

[190] Vern Paxson, Andrew Adams, and Matt Mathis. Experiences with NIMI. In *Passive and Active Measurement (PAM)*, Hamilton, New Zealand, Apr 2000.

[191] Larry Peterson, Andy Bavier, Marc E. Fiuczynski, and Steve Muir. Experiences building PlanetLab. In *OSDI*, Nov 2006.

[192] Josh Polterock. Spoofer Surpasses One Million Sessions and Publishes Final Report, Oct 2020. https://blog.caida.org/best_available_data/2020/10/13/spoofer-surpasses-one-million-sessions-and-publishes-final-report/.

[193] Morteza Safaei Pour, Joseph Khoury, and Elias Bou-Harb. HoneyComb: A darknet-centric proactive deception technique for curating IoT malware forensic artifacts. In *Proceedings of IEEE/IFIP Network Operations and Management Symposium*, Apr 2022.

[194] Emmanouil Raftopoulos, Eduard Glatz, Xenofontas Dimitropoulos, and Alberto Dainotti. How Dangerous Is Internet Scanning? A Measurement Study of the Aftermath of an Internet-Wide Scan. In *Traffic Monitoring and Analysis Workshop (TMA)*, Apr 2015.

[195] Alagappan Ramanathan, Rishika Sankaran, and Sangeetha Abdu Jyothi. Xaminer: An internet cross-layer resilience analysis tool. *Proc. ACM Meas. Anal. Comput. Syst.*, 8(1), Feb 2024.

[196] A. Retana, R. Whiteand V. Fuller, and D. McPherson. Using 31-bit prefixes on IPv4 point-to-point links. `https://datatracker.ietf.org/doc/html/rfc3021`, Dec 2000.

[197] RIPE NCC. RIPE Atlas - Coverage. `https://atlas.ripe.net/coverage/`.

[198] Romain Fontugne. AS Hegemony: AS Hegemony: A Robust Metric for AS Centrality, 2024. `https://www.iijlab.net/en/members/romain/pdf/romain_sigcomm2017.pdf`.

[199] Sina Rostami, Tiago Heinrich, and Taha Albakour. Poster: An Investigation into Internet-Facing Router Services. In *Proceedings of the 2024 ACM on Internet Measurement Conference*, IMC '24, 2024.

[200] RouteViews. University of Oregon Route Views Project. http://www.routeviews.org/.

[201] Public Safety and Homeland Security Bureau. Docket 22-90: In the matter of secure internet routing, 2022. `https://www.federalregister.gov/documents/2022/03/11/2022-05121/secure-internet-routing`.

[202] Brandon Schlinker, Todd Arnold, Italo Cunha, and Ethan Katz-Bassett. PEERING: Virtualizing BGP at the Edge for Research. In *Proceedings of the 15th International Conference on Emerging Networking Experiments and Technologies*, CoNEXT '19, 2019.

[203] Henning Schulzrinne and Marie-José Montpetit. Nsf broadband research 2020 report, Jun 2021. `https://www.cs.columbia.edu/~hgs/papers/2021/NSF-BBRR.pdf`.

[204] National Academy of Sciences. Principles and practices for a federal statistical agency, edition 7. https://www.nap.edu/resource/25885/P&P%207%20Higlights.pdf.

[205] John Scudder, Rex Fernando, and Stephen Stuart. BGP Monitoring Protocol (BMP). RFC 7854, IETF, Jun 2016.

[206] Pavlos Sermpezis, Vasileios Kotronis, Petros Gigis, Xenofontas Dimitropoulos, Jae Hyun Park, Danilo Cicalese, Alistair King, and Alberto Dainotti. ARTEMIS: Neutralizing BGP hijacking within a minute. *Transactions on Networkign*, 2018.

[207] Yavul Shavitt and Eran Shir. DIMES: let the Internet measure itself. *Computer Communication Review*, 35(5):71–74, 2005.

[208] ACM SIGCOMM. Imc test of time award, 2022. Accessed: 2025-01-30.

[209] ACM SIGCOMM. Best paper award, 2025. https://www.acm.org/conferences/best-paper-awards.

[210] Job Snijders, Ben Maddison, Matt Lepinski, Di Kong, and Stephen Kent. A Profile for Route Origin Authorizations (ROAs). RFC 9582, IETF, May 2024.

[211] Internet Society. Mutually agreed norms for routing security. https://www.manrs.org/.

[212] R Sommese, k claffy, R Van Rijswijk-Deij, A Chattopadhyay, A Dainotti, A Sperotto, and M Jonker. Investigating the impact of DDoS attacks on DNS infrastructure. In *ACM Internet Measurement Conference (IMC)*, October 2022.

[213] R Sommese, M Jonker, and k claffy. Observable KINDNS: Validating DNS Hygiene. In *ACM Internet Measurement Conference (IMC) Poster*, October 2022.

[214] Raffaele Sommese, Gautam Akiwate, Antonia Affinito, Moritz Müller, Mattijs Jonker, and kc claffy. DarkDNS: Revisiting the Value of Rapid Zone Update. In *ACM Internet Measurement Conference (IMC)*, Nov 2024.

[215] Raffaele Sommese, Roland van Rijswijk-Deij, and Mattijs Jonker. This Is a Local Domain: On Amassing Country-Code Top-Level Domains from Public Data. *SIGCOMM Comput. Commun. Rev.*, 54(2), Aug 2024.

[216] Neil Spring, David Wetherall, and Tom Anderson. Scriptroute: A Public Internet Measurement Facility. In *4th USENIX Symposium on Internet Technologies and Systems*, Mar 2003.

[217] Stuart Staniford, David Moore, Vern Paxson, and Nicholas Weaver. The top speed of flash worms. In *ACM Workshop on Rapid Malcode (WORM)*, pages 33–42. ACM, 2004.

[218] Michelle Starr. Less Than 1% of Large Hadron Collider Data Ever Gets Looked at. Published on ScienceAlert, 2018. https://www.sciencealert.com/over-99-percent-of-large-hadron-collider-particle-collision-data-is-los

[219] Florian Streibelt, Franziska Lichtblau, Robert Beverly, Anja Feldmann, Cristel Pelsser, Georgios Smaragdakis, and Randy Bush. BGP Communities: Even More Worms in the Routing Can. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, 2018.

[220] Stephen D. Strowes. Debogonising 2a10::/12: Analysis of One Week's Visibility of a New /12. In *4th Network Traffic Measurement and Analysis Conference (TMA)*, 2020.

[221] N Sultana, H Bang, E Yulaeva, R Mok, k claffy, and R Mortier. Survey on Packet Filtering. In *ACM SIGCOMM*, July 2024.

[222] Akira Tanaka, Chansu Han, Takeshi Takahashi, and Katsuki Fujisawa. Internet-wide scanner fingerprint identifier based on TCP/IP header. In *Proceedings of IEEE International Conference on Fog and Mobile Edge Computing*, 2021.

[223] Hammas Bin Tanveer, Echo Chan, Ricky K. P. Mok, Sebastian Kappes, Philipp Richter, Oliver Gasser, John Ronan, Arthur Berger, and kc Claffy. Unveiling ipv6 scanning dynamics: A longitudinal study using large scale proactive and passive ipv6 telescopes. *Proceedings of the ACM on Networking*, 3(CoNEXT3), Sep 2025.

[224] Georgia Tech. Global Routing Intelligence Platform (GRIP), 2023. https://grip.inetintel.cc.gatech.edu.

[225] Charles Testart, Philipp Richter, Alistair King, Alberto Dainotti, and David Clark. Profiling BGP Serial Hijackers: Capturing Persistent Misbehavior in the Global Routing Table. In *ACM Internet Measurement Conference (IMC)*, Oct 2019.

[226] Kedar Thiagarajan, Esteban Carisimo, and Fabián E. Bustamante. POSTER: Revealing Hidden Secrets: Decoding DNS PTR Records with Large Language Models. In *ACM SIGCOMM 2024: Posters*, 2024.

[227] Thomas Alfroy, Thomas Holterbach and Cristel Pelsser. MVP: Measuring Internet Routing from the Most Valuable Points. *Internet Measurement Conference*, 2022. https://cristel.pelsser.eu/publication/alfroy-2022/alfroy-2022.pdf.

[228] Thomas Krenc and Jennifer Sun. BGP2Go: BGP metadata database and User Interface, 2024. https://bgp2go.caida.org/.

[229] Brian Tierney, Jeff Boote, Eric Boyd, Aaron Brown, Maxim Grigoriev, Joe Metzger, Martin Swany, Matt Zekauskas, Yee-Ting Li, and Jason Zurawski. Instantiating a global network measurement framework. Technical Report LBNL-1452E, LBNL, January 2009. http://acs.lbl.gov/~tierney/papers/perfsonar-LBNL-report.pdf.

[230] Liam Tung. Google boots China's main digital certificate authority CNNIC . https://www.zdnet.com/article/google-banishes-chinas-main-digital-certificate-authority-cnnic/.

[231] U. Twente, SIDN, NLNET Labs, and SURFnet. Openintel. https://openintel.nl.

[232] US-CERT. NTP Amplification Attacks Using CVE-2013-5211, 2014. https://www.cisa.gov/news-events/alerts/2014/01/13/ntp-amplification-attacks-using-cve-2013-5211.

[233] U.S. Congressional Subcommittee on Intelligence and Emerging Threats and Capabilities. H.R. 6395—FY21 National Defense Authorization Bill, 2020.

[234] 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.

[235] U.S. White House. National Cybersecurity Strategy Implementation Plan, Jul 2023. https://bidenwhitehouse.archives.gov/wp-content/uploads/2023/07/National-Cybersecurity-Strategy-Implementation-Plan-WH.gov_.pdf.

[236] USENIX Association. USENIX Test of Time Awards. https://www.usenix.org/conferences/test-of-time-awards.

[237] Wikipedia. The Penetration of the DigiNotar CA, 2021. https://en.wikipedia.org/wiki/DigiNotar.

[238] Tzu-Bin Yan, Yuxuan Chen, Anthea Chen, Zesen Zhang, Bradley Huffaker, Ricky Mok, Kirill Levchenko, and kc claffy. Packetlab: tools alpha release and demo. In *Proceedings of the 22nd ACM Internet Measurement Conference*, IMC '22, page 766–767, New York, NY, USA, 2022. Association for Computing Machinery.

[239] Tzu-Bin Yan, Zesen Zhang, Bradley Huffaker, Ricky Mok, kc Claffy, and Kirill Levchenko. Marionette measurement: Measurement support under the packetlab model. In Cecilia Testart, Roland van Rijswijk-Deij, and Burkhard Stiller, editors, *Passive and Active Measurement*, 2025.

[240] Feng Zhang, Shijie Zhou, Zhiguang Qin, and Jinde Liu. Honeypot: A Supplemented Active Defense System for Network Security. In *Parallel and Distributed Computing, Applications and Technologies (PDCAT)*, Sep 2003.