

Dialing privacy and utility: a proposed data-sharing framework to advance Internet research

Erin E. Kenneally and Kimberly Claffy
Cooperative Association for Internet Data Analysis (CAIDA)
University of California, San Diego *
erin,kc@caida.org

1. INTRODUCTION

We re-visit the common assumption that privacy risks of sharing Internet infrastructure data outweigh the benefits, and suggest that we have a window of opportunity in which to apply methods for undertaking empirical Internet research that can lower privacy risks while achieving research utility. This window of opportunity lies in public re-examination of the assumption that the privacy risks of sharing network measurement data outweigh the benefits, and for stakeholders to self-regulate in the interests of building social capital and informing legal and judicial regimes. By sharing we mean any deliberate exchange, disclosure, or release of lawfully possessed data by a Data Provider (DP) to one or more Data Seekers (DS).

The current default, defensive posture to not share network data derives from the purgatory formed by the gaps in regulation and law, commercial pressures, and evolving considerations of both threat models and ethical behavior. The threat model from not data sharing is necessarily vague, as damages resulting from knowledge management deficiencies are beset with causation and correlation challenges. More fundamentally, we lack a risk profile for our communications fabric, partly as a result of the data dearth. Notably, society has not felt the pain points that normally motivate legislative, judicial or policy change – explicit and immediate body counts or billion dollar losses. Admittedly, the policies that have given rise to the Internet’s tremendous growth and support for network innovations have also rendered the entire sector opaque, unamenable to objective empirical macroscopic analysis, in ways and for reasons disconcertingly resonant with the U.S. financial sector before its 2008 meltdown. The opaqueness, juxtaposed with this decade’s proliferation of Internet security, scalability, sustainability, and stewardship issues, is a cause for concern for the integrity of the infrastructure as well the information economy it supports.

Internet research stakeholders have an opportunity to

tip the risk scales in favor of more protected data sharing by proactively implementing appropriate management of privacy risks. We seek to advance this objective by outlining a model – the Privacy-Sensitive Sharing (PS2) framework – that can effectively manage privacy risks that have heretofore impeded more than ad hoc or nod-&-a-wink data exchanges. Our model integrates privacy-enhancing technologies with a policy framework that applies proven and standard privacy principles and obligations of data seekers and data providers, in coordination with techniques that implement and enforce those obligations. We evaluate this framework along two primary criteria: (1) how they well the policies and techniques address privacy risks; and, (2) how well policies and techniques achieve utility objectives. We also include a case study showing how we apply the principles and techniques of the framework to share network operational data for use in cybersecurity R&D.

2. CHALLENGES AND MOTIVATIONS

Historically, Internet data of interest to network researchers has included IP topology data, traffic traces including traffic to unused address space, full packet captures of DDOS, worm, or botnet communications, exported flow records, and exterior and interior routing table data [8]. Our collective use of and dependence on the Internet continually grow, and accordingly so does the range of disciplines which must study aspects as scientifically as possible. An expanding range of Internet data of potential interest and utility to an expanding domain of researchers must bring with it deeper consideration of privacy in the collaboration and data sharing models, especially between industry and academia.

The strategic challenge is similar to other domains: how to balance utility goals with privacy risks for data seekers (DS) and data providers (DP). Internet researchers and systems security personnel are generally DS – entities seeking to share, responsibly disclose, acquire or otherwise exchange real world data. Researchers have argued that greater access to real network traffic datasets would “cause a paradigmatic shift in computer security research.” [1] While data providers (DP) acknowledge

*This work is sponsored by the U.S. Department of Homeland Security (DHS) Science and Technology (S&T) Directorate NBCHCC040159.

the potential benefits of sharing, they are sufficiently uncertain about the privacy-utility risk that they yield to a normative presumption that the risks outweigh potential rewards.

Implicit incentives to share measurement data exist, but their implementations have mostly floundered. Data sharing relationships that occur are market-driven or organically developed. Unsurprisingly then, there are no widespread and standard procedures for network measurement data exchange. Inconsistent, ad hoc and/or opaque exchange protocols exist, but measuring their effectiveness and benefit is challenging. Consequently it is difficult to justify resources for the research and collaboration costs that incentivize a sharing regime. On the other hand, the high cost of independently acquiring datasets is a motivation for re-use where possible. Transactional and opportunity costs include administrative and legal permissions, hardware and software support for instrumentation, and human capital to monitor instrumentation, and manage and curate data [1].

Privacy is difficult to quantify, as is the utility of measurement-based research. Both variables are dynamic and lack normative understanding among both domain professionals and the general citizenry. As fields of study, privacy and network science are both hindered by the absence of: common vocabulary, open and configurable reference models, uniform means of analysis, common sets of use cases, and unsurprisingly, any standard cost (liability) accounting or ROI formulas. A circular conundrum is that the risk-averse data provider needs utility demonstrated before data is released, and the researcher needs data to prove utility.

The rational predilection against sharing is strengthened by an uncertain legal regime and the social costs of sensationalism-over-accuracy-driven media accounts in cases of anonymized data being reverse engineered. While there are no procedures or regulatory framework to foster widespread exchange, there is also no framework that prohibits it either. Although there is interest in efficient and widespread sharing of measurement data, it hangs against a backdrop of legal ambiguity and flawed solution models. This backdrop and our experiences with data-sharing inform the privacy-sensitive sharing framework (PS2) we propose.

2.1 An Uncertain Legal Regime

Under the U.S. and European Union (EU) regulatory regimes, the concept of personally identifiable information (PII) is central to privacy law and data stewardship. It is commonly defined as information which can be used alone or in combination with other information to distinguish or trace an individual's identity[10]. Unlike the EU model which allots overarching protection for PII, the U.S. protects PII across a patchwork

of caselaw, state and federal industry-specific laws covering health data, financial data, education data, employment data, insurance records, government-issued records, credit information, and cable and telephone records.

At its core, there is ambiguity over fundamental concepts upon which privacy risk assessment turns. First, privacy presumes identity so unless identity is defined in relation to network data artifacts, the notions of privacy and PII are already disjointed in the Internet data realm. Both the legal and Internet research communities acknowledge that the concept of PII in Internet data is not clear – its definition is context-dependent, both in terms of technology and topology. Further, the ability to link network data to individuals – as well as the cost of doing so – changes over time as technologies and protocols evolve. Yet, PII is fundamental to interpreting and applying many laws.

For example, binary and blanket characterizations of Internet protocol addresses (IPA) or uniform resource locators (URLs) as PII (or not) are necessarily inaccurate because neither alone can capture the range of privacy risks. In practice there is little functional differentiation between these traffic components and other, traditionally protected PII, yet the related legal treatment of IPAs and URLs is far less consistent. A more accurate risk assessment depends on context, i.e, meta-data about who collected it and how they use, disclose, and dispose of the traffic data.

The risk management challenge lies in the linguistic incongruity between the legal and technical discourse about traffic data – its definitions, semantic classifications and interpretations. Officers of the court often associate IPAs with a greater privacy risk than URLs based on our ability to link IPAs to an individual via service provider account records. This distinction is artificial (albeit not totally unfounded) since both types of data reference a device or virtual location rather than an individual, and many URLs directly reveal much more user information than an IP address.

Furthermore, this legal-technical gap exposes privacy risks with network operational data insofar as many laws do not explicitly allow for research use of network data [4], and there is no bright line caselaw applying their respective exceptions to the context of sharing Internet data for research.

2.2 Flawed Technology Models

Most data-sharing efforts by the networking research community focus on improving privacy enhancing technologies (PET) to solve the privacy problem, with anonymization commanding the bulk of the attention. A typical research approach is to enumerate the possibly privacy-sensitive information present in network traffic traces, and then implement a technical, typically cryptographic,

solution to replace this information completely or partially with synthetic identifiers, normally implemented by encrypting or otherwise removing all or part of identifiers.

Since privacy risk is influenced by evolving contexts associated with relationships between people, data, technology and institutions, solely technical solutions are inherently insufficient to balance the privacy/utility trade-off. Technical researchers may rightly ask why we would predicate sharing architectures on ambiguous, unquantifiable and fallible human trust enforced by law and policy, if we can build trust through technology. The response is simple: while a purely technical approach may significantly ameliorate privacy risk, it largely fails to render empirically grounded answers to most questions being asked about the Internet today.

For example, while anonymization schemes can enhance the privacy of IPA in shared network traces, if it removes the ability to do any sort of geographic or topological analysis, the research utility of that data for studying DDoS modus operandi is dramatically reduced. A policy control framework enables the technical dials to allow for more privacy risk if a specific use justifies it. For example, traces protected by prefix-preserving anonymization may be subject to re-identification risk or content observation risk, but policy controls can help data providers minimize the chances that sensitive information is misused or wrongfully disclosed [3].

2.3 Reactive Top-Down Policy

Strategies to incentivize sharing by amending or enacting legislation merit consideration. However, regulation, especially in the technology arena, is largely reactive rather than making proactive, fundamental adjustments to predictable difficulties. Further, the length of the legislative policy cycle, confluence of variables involved in changing law, and unpredictable change agents are not amenable to immediate solutions that interested stakeholder DS and DPs can execute. A legislative solution means awaiting the infamous change agent: body counts or billion dollar losses that result from the lack of ground truth about the structure and function of networks that comprise our critical communications infrastructure.

3. SHARING RISKS AND BENEFITS

3.1 Who, What, When

Who is at risk when network data is shared?

Entities potentially at risk when network traffic is shared include: persons who are identified or identifiable in network traffic, researchers, and network providers (NP) such as ISPs, backbone providers, and private network owners. In addition to legal liabilities and ethical re-

sponsibilities, researchers and their institutions also risk withdrawal of data and/or funding as a result of privacy leakage. Society also bears costs associated with misinformation, mistrust, and internalizing behavioral norms that may result from privacy harms.

Which traffic data components are privacy-relevant?

We call a *first-order identifier* one which functionally distinguishes an individual: first and last name, social security number, government-issued and other account identifiers, physical and email addresses, certain biometric markers, and possibly the same information about immediate family. A *second-order identifier* could be an IPA, machine access code (MAC) address, host name, birthdate, phone number, zip code, gender, and financial, health, or geographic information. These indirect identifiers can also include aggregated or behavioral profile information such as IP header information, which can reveal the applications used, how often, and from which machines. Indirect identifiers also include URL click streams, which can reveal information about the content of communications, including search terms.

Under what conditions do these data types pose risk?

Network traffic data can present a privacy risk when information in packets and flow records can directly expose non-public information about persons – such as health, sexual orientation, political affiliation, religious affiliation, criminal activity, associations, behavioral activities, physical or virtual location; or, organizations – such as intellectual property, trade secrets or other proprietary information. Indirect risk exposure can occur when data is correlated (linked) with other public or private data, such as the case with IPAs of worm-infected and thus vulnerable hosts. Network data can also yield mistaken attributions and inferences about behavior.

The privacy risk across time may also vary – the threat may be immediately manifest upon disclosure of data, or it may be a latent risk which is held in abeyance until some future condition arises. Lack of transparency between the DP and DS regarding the shared data's nature, scope, and lineage is invariably a condition that enhances risk.

3.2 Privacy Risks of Internet Research – Laws and Courts of Public Opinion

It is impractical to enumerate all laws that may affect privacy risk, but such inventorying is not prerequisite to capture the foreseeable risks of network data sharing. It is sufficient to note that legal liability or ethical obligations underlie each privacy risk. In the U.S., privacy-related legal liabilities can derive from the federal Constitution (most notably the Fourth Amendment), federal law and regulation, contract law, tort law (e.g., invasion of privacy), state law equivalents, and organizations' privacy policies. Beyond legal risks, violations of ethical obligations can create normative harms that

implicate reputation and financial damages. Dismissing ethical obligations as discretionary and unenforceable overlooks how ethical violations are treated by public opinion, and also ignores the fact that many laws are informed by ethical norms [7].

Public disclosure is the act of making data readily available to the general public via publication or posting on the web. The privacy risks of sharing data containing PII which is subsequently displayed on the web are obvious and incontrovertible. More common and challenging are publicly available network traces and activity logs which reveal identifying information about infected hosts. Such disclosure raises the risk that unpatched or vulnerable hosts will be further exploited, thus creating security and reputation risks for individuals and organizations.

Accidental or malicious disclosure is the act of making information or data available to a third party(s) as a result of inadequate data protection. AOL provided a quintessential example in 2006 when they released an anonymized data set of search queries that were linked back to users conducting the searches using public metadata. These persons were then exposed in the NYT. [2]

Compelled disclosure to third parties risk arises with the obligations attendant to controlling data, such as having to respond to subpoenas requesting data disclosure in lawsuits. The RIAA campaign to massively subpoena ISPs and universities in an attempt to identify copyright infringers is a notorious example. To avoid such risk, many entities (including research organizations) have chosen not to retain data, thereby also losing operational and research value.

Government disclosure involves the release of data to government entities. An infamous example is the disclosure of call data records by major telecommunications carriers to the National Security Agency around 2007 [6]. Release to the government introduces another level of risk involving civil rights and liberties, such as imprisonment and restrictions on speech and associations.

Misuse of user or network profiles arises with network traffic that contains information about proprietary or security-sensitive network architectures or business operations. Advancing traffic and topology analysis, data mining and classification techniques can derive sensitive information from seemingly benign traffic data, and thereby reveal user behaviors, associations, preferences or interests, which attackers, advertisers, or content owners can then exploit. Network operators themselves may use such information for network management, illustrated by Comcast's recent throttling of BitTorrent traffic.

Inference misuse risk involves synthesizing first-order or second-order identifiers to draw inaccurate inferences

about a person's behavior or identity that leads to damage or harm.

Re-identification /De-anonymizing misuse risk.

Re-identification/de-anonymization, involves reversing data anonymization or masks to link an obfuscated identifier with its associated person. Shared anonymized data poses a misuse risk because it is variably vulnerable to re-identification attacks using public or private information whose availability is beyond the knowledge or control of the original or intermediate data provider [12]. Anonymized data may not immediately expose PII, but any time a piece of de-identified data has been linked to first order identifying information, other anonymous aspects of the obfuscated data are easier to de-anonymize. Aggregation or statistical techniques for anonymization are not immune to re-identification risk.

Examples of reidentification risk are the 2007 Netflix prize incident [9], and a similarly embarrassing episode of re-identification within the Internet research community [1].

De-anonymization risk bears special consideration in the growing incongruity around PII. DPs face increasing legal and societal pressures to protect the expanding amounts of PII they amass for legitimate business purposes. Yet, DPs are under equal pressure from the marketplace to uncover and exploit PII in order to better connect supply and demand, and increase profit margins on their goods and services. DPs will turn to anonymization to avoid triggering privacy laws that exempt aggregate or anonymized data.

Like the arms race between exploits and defenses in the systems security arena, de-anonymization techniques will likely become commoditized to support investigative reporting, law enforcement, business intelligence, research, legal dispute resolution, and the presumed criminal threatscape.

3.3 Utility of Internet Measurement

The benefits of network research derive from the value of empirical network science [5], which includes a better understanding of the structure and functions of networks that comprise critical Internet infrastructure. Network researchers and funding agencies struggle to establish a science agenda, partly due to their lack of visibility into the infrastructure, but also because the field is younger and less well-defined than traditional scientific disciplines.

The following are criteria against which to measure, evaluate and communicate the benefits of sharing network data for research:

- The objective for sharing the data promotes social welfare or generalizable knowledge.
- The data is not already being shared, or if it is, there remains a qualitative need for sharing be-

tween other DS and DPs

- The research could not be conducted without the data.
- The scientific methodology using the data is transparent, objective, and repeatable relative to any privacy controls that are implemented.
- Research results can be acted upon meaningfully.
- Research results can be integrated with business processes, such as situational awareness of critical infrastructure or operational security.

Research in network measurement that could satisfy the above criteria include:

- information and network security questions regarding system threats, including characterizing baseline and anomalous workloads, modeling malware, developing effective strategies to deal with threats.
- macroscopic analysis of Internet topology; understanding the how the evolution of the network is affecting the efficiency and capabilities of the underlying routing, transport, and naming protocols.
- understanding the effect of the prevalence and growth of new applications on Internet workload, topology, and infrastructure economics.
- validation of traffic, congestion control, and performance assumptions, models, and analyses, both for current and proposed new technologies.
- development and evaluation of new tools and algorithms, including measurement and sampling techniques.

4. PS2 FRAMEWORK: ELEMENTS, EXECUTION, AND EVALUATION

We describe the Privacy-Sensitive Sharing Framework and then evaluate the model’s ability to address the privacy risks outlined in 3.2 and the utility criteria in 3.3. Recognizing that privacy risk management is a collective action problem, our PS2 framework contains this risk by conveying the collection, use, disclosure and disposition controls over to the DS coincident with the shared data. This framework contemplates that the privacy risks associated with shared data are contagious – if the data is transferred, some degree of responsibility for containing the risk lies with both provider and seeker of data.

4.1 Elements of PS2

The PS2 is a structured framework for describing privacy risks and controls to support and implement functional privacy requirements. It serves three purposes: an analytical tool for assessing the risk posture of the proposed data sharing; a basis for establishing privacy management (technical and policy) controls; and a template for developing operational solutions to balancing privacy and utility in data sharing.

While not anchored on specific regulation, the com-

ponents of our framework are rooted in principles and practices that underlie privacy laws and policies on both the national and global levels. The Fair Information Practices (FIPS) are considered de facto, international standards for information privacy and address collection, maintenance, use, disclosure, and processing of personal information [11]. The PS2 framework – a hybrid of policy and technical controls – applies these principles to the context of Internet research, allowing navigation of data disclosure and misuse risks, and serving as a touchstone for legally and ethically defensible data sharing.

- Authorization – Internal authorization to share requires explicit agreement between the DP and DS. This may require direct consent from individuals identifiable in network traffic or via proxy consent with the DP.¹
- Oversight – The DP and DS should engage external oversight of the proposed sharing, such as from an Institutional Review Board (IRB).
- Transparency – The DP and DS should be open and in agreement over the collection, use, disclosure, objectives and obligations and associated with shared data. For example, data-sharing terms might require that the algorithms be public but that the data and/or conclusions remain protected, or vice versa [13].
- Compliance with applicable law(s) – Collection, use and disclosure of data should comport to a reasonable if not case-law precedented interpretation of laws that speak directly and clearly to sharing risks about proscribed behaviors or mandated obligations.
- Purpose adherence – The data should be used consistent with the documented goal for why it is being shared.
- Access limitations – The shared data should be restricted from those who do not have a need and right to access the shared data.
- Use specification and limitation – Unless otherwise agreed, the DP should prohibit merging or linking data that would create or enhance privacy risk.
- Collection and Disclosure Minimization – The DP should collect and disclose only the data that is necessary to achieve the research goals, and eliminate extraneous data that carries a privacy risk. Prominent privacy-sensitive techniques include:
 - A. Deleting/filtering sensitive data.
 - B. Deleting/filtering part(s) of the sensitive data.
 - C. Anonymizing/hashing/de-identifying all or parts of the sensitive data.
 - D. Aggregating or sampling.

¹Consent requirements for Internet traffic monitoring are unresolved, but will no doubt be a part of forthcoming legal, policy and community decisions.

- E. Mediation analysis /human proxy – this is a sandbox approach that involves “sending the code to the data” rather than releasing sensitive data for analyses.
 - F. Aging the data – traffic data that is de-sensitized by virtue of being non-current, i.e., no longer contains a direct or indirect identifier that poses a risk of harm.
 - G. Size/quantity limitation – this entails minimizing the quantity of traces shared.
 - H. Multiple layers of anonymization.
- Audit tools – Techniques for provable compliance with policies for data use and disclosure, e.g., secure audit logging via a tamper-resistant, cryptographically protected device connected to but separate from the protected data, accounting policies to enforce access rules on protected data.
 - Redress mechanisms – Procedures to address harms from inappropriate data use or disclosure, including a feedback mechanism to support correction of datasets and/or erroneous conclusions.
 - Data and analysis quality assurances – Awareness by the DS and DP of inference confidence levels associated with the data.
 - Security – Controls should reasonably ensure that sensitive PII is protected from unauthorized collection, use, disclosure, and destruction.
 - Training – Those who are authorized to engage the data should be educated and made aware of the privacy principles and controls associated with the data.
 - Impact assessment – Sharing dynamics should consider potential collateral effects on stakeholders affected by the data, and seeks methods that do no further harm.
 - Transfer to third parties - This should be prohibited unless equivalent data control obligations are transferred, relative to the disclosure risks associated with that data.

4.2 Execution of PS2

The technical controls of the PS2 are self-contained, although they need to be identified and enforced. The policy controls require an execution vehicle, such as a bi/multilateral Memoranda of Understanding (MOU) or Memoranda of Agreement (MOA), a model contract, or a binding organizational policy. For lower risk sharing situations, a unidirectional Acceptable Use Policy AUP may be cost-preferential to negotiated bilateral agreements. Mutual and explicit consent to engage policy and technical controls provides an enforceable standard and certainty that can serve as a safe harbor for liability under many data privacy laws.

4.3 Evaluation of PS2

PS2/Privacy Risk	Public Disclosure	Compelled Disclosure	Malicious Disclosure	Government Disclosure	Misuse	Inference Risk	Re-ID Risk
Authorization	✓			✓			
Transparency						✓	✓
Law Compliance	✓	✓	✓	✓	✓		
Access Limitation	✓		✓	✓			
Use Specification	✓			✓	✓		
Minimization	✓	✓	✓	✓	✓	✓	
Audit Tools							
Redress						✓	
Oversight	✓				✓		
Data Quality						✓	
Security	✓		✓		✓		
Training/Education	✓			✓	✓		
Impact Assessment						✓	✓

Table 1: PS2 policy and technical components evaluated against privacy risks. The ✓’s indicate that the particular PS2 component in the row addresses the privacy risk enumerated in the corresponding column. (*Minimization* refers to the techniques in Table 2.)

The PS2 framework offers a template for assessing and developing operational solutions for balancing privacy risks and utility rewards when sharing data for research. It can help an oversight committee determine whether possible risks are justified, by specifically asking the user to assess sharing risks against technical and policy controls, as well as to assess the achievement of utility goals against those controls. For the prospective DP, the assessment will assist the determination whether or not to participate.

Table 1 illustrates whether the privacy risks are mitigated by the primary components of PS2. The ✓’s indicate that the particular PS2 policy component in the row mitigates against the privacy risk enumerated in the corresponding column. The table illustrates that the policy control component of PS2 leaves gaps in addressing the full range of privacy risks. Further, it suggests that the technical control component (minimization techniques) (Section 4.1) can, however, address all privacy risks. The implication is that a purely technical sharing framework is sufficient to address privacy risks, and therefore a policy control backdrop is superfluous. However, evaluating the technical minimization controls against the utility goals in Table 2 illustrates

Minimiz.Tech./Utility Need	Is Purpose Worthwhile?	Is there a need?	Is it already being done?	Are there alternatives?	Is there a scientific basis?	Can results be acted upon?
Not Sharing	X	X	X	X	X	X
Filter All	X	X	X	X	X	X
Filter Part	X	X		X	X	
Anonymize	X	X	X	X	X	
Aggregate	X	X	X	X	X	
Mediate (SC2D)	X					
Age Data	X	X	X	X	X	
Limit Quantity	X	X	X	X	X	
Layer Anonymization	X	X	X		X	

Table 2: PS2 technical component (minimization controls) evaluated against utility needs. The X’s indicate where the minimization technique impedes the research utility goal.

the weakness of a one-dimensional technical approach. This weakness is unsurprising, since data minimization techniques intentionally obfuscate information often essential to most Internet research. These utility gaps can be modulated (“dialed down”) by engaging the policy components of PS2. In short, a purely technical approach breaks down along the utility dimension, and the pure policy approach may leave too high privacy risk exposure, justifying a hybrid framework that covers both privacy risks and utility goals. We note that evaluation of the framework should also consider practical issues such as education costs, whether new privacy risk(s) are introduced, whether control(s) are forward-looking or also address legacy privacy risks, and possible free rider problems created by DPs who choose not to share.

5. PS2 CASE STUDY: NETWORK TELESCOPE

To promote cooperative analysis of Internet traffic and performance and advance the state of cybersecurity research, we have implemented the policy-supported and risk-sensitive PS2 data sharing framework. We recently applied this framework to a then-new mode of data sharing: real-time sharing of Internet traffic data observed at the network telescope. A network telescope is a segment of routed IP address space on which little or no legitimate traffic exists. Each such chunk of address space provides a unique and continuous view of anomalous, unsolicited Internet traffic to no legitimate

destination.

Observing traffic from a network telescope allows visibility into a wide range of security-related events, including misconfiguration (e.g. a human being mis-typing an IP address), malicious scanning of address space by hackers looking for vulnerable targets, backscatter from random source denial-of-service attacks, and the automated spread of malicious software called Internet worms. The primary obstacles to sharing telescope data are privacy and security concerns. Because viruses and worms may involve the installation of backdoors that provide unfettered access to infected computers, telescope data may advertise these vulnerable machines.

CAIDA previously addressed the privacy risk of releasing victim host IPAs and unexpected but occasional payload content with strict filtering and anonymization disclosure controls. This implementation of privacy risk controls came at research utility costs, which two events in 2009 motivated us to re-examine: the Conficker worm outbreak, and a new storage cost allocation structure in our organization. In 2009 we transitioned from a model of static trace sharing and indefinite storage of data on CAIDA servers, to a model of real-time data sharing with vetted researchers, storing only a 30-day window of history. This new model aims to allow researchers access to a telescope observatory during a worm outbreak, where raw traces containing target addresses and payload that could enable autopsy of the structure and function of cybersecurity threats.

Consistent with the PS2 framework, we use Acceptable Uses Policies (AUP) and disclosure control techniques to guide this shift in our data-sharing approach. We implement transparency by clearly describing the dataset and its use obligations on our public website. We obtain explicit consent to abide by the stated responsibilities by requiring each researcher (DS) to complete and execute a data request form which includes acknowledging data use terms prior to receiving access. External oversight is addressed by our university’s Institutional Review Board, which certifies that the datasets are collected and made available in accordance with the principles of respect for persons, beneficence and justice as relevant to human subjects. We institute purpose specification by obtaining explicit webform acknowledgment from the DS that s/he will use the dataset solely for the stated research purposes. We enforce access to the dataset(s) and authorization to use it by application review, approval and communication of acquisition instructions by CAIDA administrators. This review includes restricting access to DS from export-restricted countries. DS also consent to use appropriate and reasonable care in safeguarding access to and preventing unauthorized use of the data. We obtained legal advice to ensure sharing methods comply with laws and policies related to data privacy, confidentiality and pro-

tection.

Another element of the PS2 is the impact assessment. CAIDA researchers and administrators considered possible harm to individuals or organizations, as well as the likelihood of achieving the growing needs of security research by only releasing and storing completely sanitized versions of static, periodic datasets. To the extent possible within the real time strategy, privacy sensitivities were addressed with loosened disclosure controls (anonymization of any identifying payload). Our AUP backstopped re-identification risk by requiring that researchers agree to make no attempts to reverse engineer, decrypt, or otherwise identify the original IP addresses collected in the trace.

As discussed in 4.3, our original disclosure control strategy largely ameliorated the privacy risks of disclosing victim IPAs and payload, but to the detriment of security research utility needs. The speed, scope, and strength of automated malicious software demand effective real-time sources of data that matches the dynamics of the threat. Studying a worm in situ requires real time traffic access, including raw victim host IPAs, and payload data. None of these needs were supported by our original disclosure control strategy.

The PS2 hybrid framework allowed CAIDA to realize utility goals in a risk-sensitive manner, by dialing down the technical disclosure controls, and relying on policy components to close resulting privacy gaps. Specifically, we collected and shared telescope data as raw (un anonymized) traces, with payload (content). Rather than mitigate the risk of releasing victim IPA by anonymization and wholesale deletion of security-relevant data, CAIDA revisited the privacy impact assessment, loosened the technical disclosure controls, and tightened its use and disclosure obligations in the AUP. CAIDA used the considerations enumerated in Table 2 and Section 3.3 to enable transparent and reproducible scientific research of a critical infrastructure security event *while still in progress*, a contribution to the security field not possible with previously available data sets.

The privacy risks associated with this dataset were such that CAIDA could effectively manage them by relaxing the technical anonymization barriers to research utility, and constricting DS use and disclosure via the policy component. For example, the DS are prohibited from attempting to connect to, probe, or in any other way interacting or intervening with a machine or machine administrator identified in the dataset, without permission from CAIDA. For any publication or other disclosure of non-anonymized data, the DS is obligated to anonymize or aggregate IP addresses, network names, and domain names unless obtaining written authorization from CAIDA to do otherwise. We (do our best to) enforce compliance with these restrictions with an audit policy that requires the DS to report a sum-

mary of the research and any findings, publications, or URLs using the data to CAIDA at the conclusion of the research, or semi-annually.

6. CONCLUSIONS

The Privacy Sensitive Sharing (PS2) framework considers practical challenges confronting security professionals, network analysts, systems administrators, researchers, and legal advisors. It embodies the proposition that privacy problems are exacerbated by a shortage of transparency surrounding the who, what, when, where, how and why of sharing privacy-sensitive information. The PS2 enables transparency as a touchstone of data-sharing.

The PS2 offers a consistent, transparent and replicable evaluation methodology for risk-benefit determinations rather than relying on subjective, opaque and inconsistent evaluations that turn on *trust-me* decision metrics. The PS2 is a hybrid approach: a policy framework that applies proven and standard privacy principles between the data seekers and data providers, coordinated with technologies that implement and provably enforce those obligations. We evaluated this framework along two main criteria: (1) how well the policies and techniques address privacy risks; and, (2) how well policies and techniques achieve utility objectives.

We hope this framework helps network measurement advocates use this window of opportunity to experiment with models that effectively manage privacy risks which have heretofore impeded more than ad hoc or nod-&-a-wink data exchanges. Because the principles underlying data stewardship are domain-agnostic, the PS2 principles can also help prevent a proliferation of infamous poster cases [2, 9, 1] across disciplines.

By taking proactive and ethically defensible steps to transparently engage sharing models like PS2, we can more practically influence policy and law at these crossroads. The alternative is to wait for a legislative reaction to catastrophe, at which time a window of opportunity will have closed. Information security controls were initially considered a liability (from a cost perspective) until regulations rendered lack of security a compliance liability. We anticipate circumstances to reveal that rather than data-sharing being a risk, *not* sharing data is a liability. We offer the PS2 as a tool to help move the community mindset in that direction as productively and safely as possible.

7. REFERENCES

- [1] ALLMAN, M., AND PAXSON, V. Issues and etiquette concerning use of shared measurement data. In *IMC* (2007).
- [2] BARBARO, M., AND T. ZELLER, J. A Face is Exposed for AOL Searcher No. 4417749. *New York Times* (Aug 2006).

- [3] BURKHART, M., SCHATZMANN, D., TRAMMEL, B., BOSCHI, E., AND PLATTNER, B. The role of network trace anonymization under attack. *ACM SIGCOMM Comp. Comm. Rev.* (2009).
- [4] BURSTEIN, A. Amending the ECPA to Enable a Culture of Cybersecurity Research. *Harvard Journal of Law & Technology* 22, 1 (2008), 167–222.
- [5] C. B. DUKE, *et al.*, Ed. *Network Science*. The National Academies Press, Washington, 2006.
- [6] CAULEY, L. NSA has massive database of Americans’ phone calls. *USA Today* (May 2006).
- [7] CENTER FOR DEMOCRACY AND TECHNOLOGY. CDT’s Guide to Online Privacy, 2009.
- [8] CROVELLA, M., AND KRISHNAMURTHY, B. *Internet Measurement: Infrastructure, Traffic and Applications*. John Wiley and Sons, Inc., 2006.
- [9] NARAYANAN, A., AND SHMATIKOV, V. Robust De-anonymization of Large Sparse Datasets. *IEEE Symposium on Security and Privacy* (2008).
- [10] NATIONAL INSTITUTE OF STANDARDS AND TECHNOLOGY. Guide to Protecting the Confidentiality of Personally Identifiable Information, January 2009.
- [11] OECD. Guidelines on the protection of privacy and transborder flows of personal data, 1980.
- [12] PORTER, C. De-Identified Data and Third Party Data Mining: The Risk of Re-Identification of Personal Information. *Silder Journal of Law, Communication, and Technology*, 3 (2008).
- [13] SWIRE, P. A Theory of Disclosure for Security and Competitive Reasons: Open Source, Proprietary Software, and Government Agencies. *Houston Law Review* 42, 5 (January 2006).