Date: 14 December 2009

To: Human Research Protections Program Office, MC 0052

Re: Project #091710S - Cooperative Association for Internet Data Analysis (CAIDA)

Dear Committee:

This is a reply to your letter dated 2 December 2009 requesting clarification/revision of the enclosed Research Project Application. The relevant revisions to the Research Application have been underlined as per your request.

Your committee asked for clarification on what controls are instituted in the Research Plan to eliminate the risk of intervention/interaction with individuals, and relatedly, how we mitigate the potential risk of disclosure and also inform the IRB of potential problems that might arise as a result of the data collected by this research. We begin by making explicit what was apparently not clear in our original application: in the course of the conducted research, we contend that our research does not obtain data through intervention or interaction with individuals, nor does our research involve obtaining private information that is individually identifiable. For these reasons, we advocate that this research does not involve human subjects. The data we obtain does not involve contact with or manipulations of any human subjects' environments (in this case Internet traffic), and there are no direct or indirect communications or interpersonal contacts between the researchers and human individuals. The information we collect does not pertain to behavior of persons in contexts where they can reasonably expect that they will be observed or recorded, nor is the obtained data provided for specific purposes by an individual or that which s/he can reasonably expect will not be made public. And, the identity of a subject is not readily ascertained by the researchers.

We nevertheless aim to appropriately address your concerns based on the text of the Research Plan you reviewed. Specifically, we minimize any plausible risk of intervening and interacting with individuals by virtue of the data we collect and disclose. We also augment any minimal risk with a privacy sensitive data management infrastructure that prohibits interaction with persons. This control framework also addresses the Committee's second concern related to how we mitigate the potential risk of disclosure that might arise as a result of the data collected by this research.

The purpose of our research data collection and sharing is to understand the 'anatomy and physiology' of macroscopic traffic, and not for engaging persons or pursuing knowledge in the social, behavioral or psychological sciences. Our data collectors are instrumented to collect network layer traffic data that does not contain individually identifiable information such as biometric, financial, health or other unique identifiers. The data collected may contain Internet Protocol Addresses (IPA) or payload from anomalous traffic, neither of which is private information that is individually identifiable.

The IPA data that we collect is not identifiable private information such that the identity of a user is or may readily be ascertained by the researchers. In this way, we preempt the risk of interaction/intervention with persons by not collecting data that could be readily used to engage individuals. IPA directly identifies an interface to a specific device connected to the Internet. As described in the *Methodology* section, IPA often does not identify a specific host on the Internet due to the use of Network Address Translation (NAT) devices and other firewalls which functionally obscure source and/or destination computers by collating them under one IP address. Further, even though every device on the Internet is assigned a unique IPA, not every device points to a host operated by a an identifiable human subject, but rather, may also attach to a 'dumb' machine performing administrative operations to induce traffic to flow across the Internet. Most of those responding IPAs that we collect are not end hosts operated by individuals, but rather, are intermediate computers (routers) that automatically intermediate traffic along its paths. This means that these IPAs are incapable of being mapped to individual human users.

In addition, the selection of the destination hosts (identified by IP address) is random, with the only deliberate target being the broad, routed networks on the Internet. This approach ensures that we sample from all routed network prefixes so that our topology measurements are complete and accurately reflect the full scope of the commodity Internet. The destination addresses within each routed network prefix are arbitrarily selected and probed approximately every 48 hours (one probing cycle), with each probing cycle involving the selection of another randomly selected host. We do this random probing within a large scope of possible hosts (i.e., there are roughly 250,000 routed network prefixes that we separate into roughly 7.4 million targeted networks each containing 256 IP addresses). At the time of writing, approximately 6.5% of probed IP addresses respond. Furthermore, a large portion of IPAs associated with human users are assigned in a dynamic manner, meaning that a user's IPA while on the Internet today or even several hours in the future may not be the same. So, IPA are not static and permanent -- features which are paramount to individual identification.

Although an IPA that identifies a device could possibly be mapped back to a specific computer that is operated by a human user, such links are not readily ascertainable by researchers. Records that link an IPA to a user account, as opposed to an automated device, are maintained by the Internet Service Provider (ISP) through

which the device is registered, and those records are protected from disclosure by state and federal laws. To link an IPA with an identifiable individual, a researcher would need to compel an ISP to disclose the mapping to a registered account name via a court order or subpoena. This would necessitate the filing of a civil lawsuit or criminal charges for a violation of the law. Moreover, these records linking the IPA to a registered account are still not dispositive of the actual users' identity since multiple persons can access the Internet using a known person's account while maintaining anonymity. Certainly there are legitimate and publicly-available search tools that any person could query to obtain the name of the registrant of an IPA. However, these records map back only to the level of the ISP, and not to a specific device, user account or individual users associated with the specific network activity that may be collected in the traffic. So, inherent to the issuing of IPAs, if the computer is identified by an IPA, private information that is individually identifiable is not readily ascertainable by researchers at CAIDA or external researchers to whom we share the data.

As for the payload data from passive collection (i.e, Network Telescope), there is a slight possibility that inadvertent and random information related to an individual could be collected in those traces, however the controls we describe immediately below provide the accountability and enforcement to minimize any probable risk that there may be interaction/intervention with an individual. More importantly, we contend that payload data does not invoke human subjects concerns under the reasoning that the data is from non-legitimate traffic in ungoverned network space that does not serve the public. By its nature, such data is outside of the bounds of normal and expected network communications and as such does not invoke the need for authorization to observe and collect. Any payload information that may flow across our monitors does not have a reasonable expectation that it will be observed, recorded or made public. Persons' normal, daily communications and transactions traffic are not designed or intended to flow through the network on which we place our Network Telescope monitors.

In addition to controlling against intervention/interaction with persons by the choice of data that we collect, we augment any minimal risk with a privacy sensitive data management infrastructure that prohibits interaction with persons. This means that for both IPA and payload data, we further control against interaction/intervention with an individual by not targeting or tracking any IPAs or identifying data that are collected in our traffic measurement. Specifically, probing or other communication with the IPAs are prohibited by policy and binding agreements, as are any attempt to connect any of the data to an identifiable individual. This multifaceted framework – the Privacy Sensitive Sharing (PS2) framework also serves to mitigate the potential risk of disclosure of any data that relates to an identifiable individual. It involves proactive and reactive means by which we diminish the linking of an IPA to an individual and ensures that risk of future misuse by disclosing datasets to other researchers is minimized.

CAIDA has since its inception over ten years ago, and prior to IRB oversight, implemented and revised various techniques, policies and procedures for preventing and mitigating risks associated with the collection, analysis, storage and use of measurement data for network research. To continue responsible and risk sensitive research related to evolving efforts described in our Research Application, in 2009 we implemented the PS2 framework. The PS2 is a hybrid control framework -- it integrates privacy-enhancing technologies with an enforceable collection of policies by applying standard privacy principles (i.e., the Fair Information Practices) in coordination with methods that provably enforce those requirements. It is thus an umbrella framework that creates a repeatable, transparent, and scalable set of risk controls, in contrast to previous case-by-case and compartmentalized approach to risk management. The result is a dynamic set of privacy and confidentiality risk controls that address the people, policies and technologies in the lifecycle of our research activities. The PS2 controls against risk in each dataset by compartmentalizing against the user type and function (who), conditions for accessing the dataset (why), and use and disclosure obligations (where, when, how). The PS2 controls include: data authorization, internal oversight, compliance with laws, adherence to acceptable research purposes, limitation of access to data, disclosure controls, audit tools, and data security.

The primary technology component of the PS2 is anonymization. For active data measurements, we mitigate risk of correlating IPAs with individual users by modifying the IPAs using commonly-accepted anonymization techniques. This allows for distinguishing among groups of users while preventing individual identification of a user on the network. In addition to these technical controls, we further mitigate risk of intervention/interaction with persons by executing Acceptable Use Agreements with internal and external researchers who access the data. Prior to releasing any data to a requesting researcher, CAIDA initially controls against potential unauthorized use and disclosure risks by requiring that researchers provide descriptions of intended usage. Researchers who are granted access assert that as a condition to accessing data they agree to abide by the use and disclosure restrictions. These include prohibiting researchers from interacting/intervening with any IPA in the requested dataset.

Specifically, this means that attempting to connect to, probe or in any other way contact a machine or machine administrator associated with an IPA in the dataset are prohibited, as are any attempts to resolve any of the data to an individual whether by reverse engineering, decryption or otherwise. Researchers using the data are also restricted from distributing the data beyond authorized users and are required to anonymize, aggregate or summarize any data that may contain personally identifiable information, IPA, network names, and domain names prior to any publication to protect the privacy of end users. Also, the Acceptable Use Agreement

prohibits users of the datasets from de-identifying any anonymized data. And, researchers also agree to use appropriate and reasonable care in safeguarding access to and preventing unauthorized use of the data. The data is secured from unauthorized users by access controls associated with a detailed research user profile which we maintain in a records log and monitor regularly. We use best efforts to enforce compliance with these restrictions via an audit policy that requires the researchers to report a summary of the research and any findings, publications, or URLs using the data to CAIDA at the conclusion of the research, or semi-annually.

Finally, we have revised our PS2 control framework to include an 'incident response' mechanism, whereby researchers who fail to abide by the data use and disclosure restrictions are required to notify CAIDA of the nature and circumstances surrounding the violation. Having been put on notice, CAIDA will work with researcher to remediate any violations, and will promptly notify the IRB of potential new risks or harms that may result.

Application to the University of California, San Diego
Human Research Protections Program Office
For Review by the UCSD Institutional Review Boards

Project #091710S
Cooperative Association for Internet Data Analysis (CAIDA)

Human Research Protections Program (0052)
University of California, San Diego
La Jolla, CA 92093-0052

31 October 2009

Dear Committee on Investigations Involving Human Subjects:

Attached please find the full Social and Behavioral Human Research Protections Application by The Cooperative Association for Internet Data Analysis (CAIDA) that we submit for expedited review. This application was previously approved under expedited review effective 11/1/2008 - 10/31/2009. This updated application reflects substantively identical research with the exception of one methodological change for which we describe and provide a risk assessment therein.

This application describes the minimal threats posed to human subjects in the course of our research activities, as well as the technical, policy and administrative risk controls CAIDA has implemented to adhere to and advance privacy and confidentiality standards for research involving human subjects. These risk controls are encompassed within our Privacy Sensitive Sharing (PS2) framework, a structure that fosters transparent, accountable and privacy-protective research activities. Risk-benefit thresholds are assessed and maintained with the due care of both domain-specific researchers and the ordinary prudent person under objectively reasonable standards.

In light of the minimal risks attached to the information involved in this research and the interconnected risk control structure, we suggest that the potential risks and threats to the privacy and confidentiality of human subjects are minimal, unremarkable, and tolerable relative to the potential contributions to generalizable scientific knowledge we are pursuing. Specifically, the only data collected, maintained or disclosed that may present a potential risk are: (1) Internet Protocol Addresses (IPA) of the source and destination computers involved in network research measurement; and, (2) payload from the Network Telescope.

Regarding the former, an IPA identifies an interface to a specific device connected to the Internet that may or may not link back to a specific human subject operator. There is no comprehensive and incontrovertible executive, judicial or legislative authoritative pronouncement in the United States that has defined IPA alone to be individually identifiable information. Further, there is no intervention or interaction with the IPAs, only collections of measurements of Internet traffic that may contain these electronic device markers. To be precise, there is a possibility that with enough of the right type(s) of additional data, IPA could be mapped back to a specific computer, user account, organization, or person. For IPA data to render individually identifiable

information, a requester could undertake and satisfy legal requirements to compel an Internet Service Provider to disclose the mapping to a registered account name (which itself does not necessarily map to an individual since multiple, unidentifiable persons can access the Internet using a single account). Otherwise, a person's identity could possibly be inferred by match-linking IPA data with another data source of direct or indirect personal identifiers.

Secondly, the Network Telescope collects payload for which there is a slight chance that inadvertent and random information related to an identifiable individual exists, however there is no interaction/intervention with any such data.

Notably, even if the research data (e.g., IPA and Telescope payload) can be linked to individually identifiable information, the context within which it was collected may not have a reasonable expectation of privacy. In other words, there are cases where the data may not be private and individual Internet users cannot reasonably expect that it is not is being observed, recorded or otherwise subject to public disclosure. Given the collection of an extremely limited quality and quantity of such information, the difficulty in correlating the data to an individual (i.e., the effort, resources and coordination costs), and the expectations attached to public behavior, our research collection and disclosure presents little potential risk of legal liability, reputation damage, psychological or physical harm, or economic damage to an identifiable human subject.

To further diminish any identification and disclosure risks posed by identifiable data from network traffic that may/may not carry an expectation of privacy, we implement multiple layers of additional risk controls via the PS2 hybrid technology-policy platform. We reduce the risk of correlating IP addresses with individual users by modifying the IPAs using commonly-accepted methods to remove identifying particulars (anonymization), as well as by simply filtering out portions of traffic data known to contain personally identifying information. In addition, CAIDA requires both internal and external researchers to provide descriptions of intended usage and assert to abide by acceptable data use agreements (AUP) that restrict the distribution, use and disclosure of shared research data that carries any privacy or confidentiality risk to an identifiable subject.

For the above reasons, we contend that the controls implemented are proportionate to the reasonably foreseeable risk while not obstructing our research goal of improving generalizable knowledge produced through network science. The risks described above are minimal, and present a privacy impact that is no greater than that which individuals face in their normative use of the Internet. Finally, this application presents issues of first impression in the application of human subjects review to network traffic context. Our risk assessment and due diligence in conferral with similarly situated researchers and evaluators supports our suggestion that this research is eligible for IRB exemption status as research involving the observation of public behavior (45 CFR

46.101(b)). Whether this research activity qualifies as exempt; is deemed not to involve human subjects; or upon substantive review is determined to not present a prohibitive risk in light of research benefits, we embrace full transparency with the IRB. We see great value in engaging dialogue and formal review so that we may collectively establish a framework for all stakeholders to achieve their respective goals while protecting their respective interests. We maintain interest in the opportunity to extend or enhance the training course on protection of human research subjects to include a network research context.

The Cooperative Association for Internet Data Analysis (CAIDA)
Application to the University of California, San Diego
Human Research Protections Program
For Review by the Institutional Review Board

1. FACILITIES.

CAIDA's primary facilities for managing and coordinating our research and measurements occur at the San Diego Supercomputer Center on the UCSD campus. Our research involves both actively (via probes) and passively (via taps) measuring the Internet, using measurement servers located in facilities around the globe.

2. DURATION.

CAIDA intends this application to address various methods by which we measure the Internet. The group conducts limited duration measurement experiments and data collection events as well as ongoing efforts that span more than ten years in order to recognize longer term trends. Due to the nature of these longitudinal studies, we hope to maintain ongoing approval for an indefinite period of time, filing supplemental and/or renewal information to this application to reflect any changes in the research design and as deemed necessary by the IRB.

3. SPECIFIC AIMS.

CAIDA provides data, tools and analyses promoting the engineering and maintenance of a robust, global Internet infrastructure. CAIDA investigates both practical and theoretical aspects of the Internet, with particular focus on topics that:

* are macroscopic in nature and provide enhanced insight into the function of Internet infrastructure worldwide,
* improve the integrity of the field of Internet science,
* improve the integrity of operational Internet measurement and management, and
* inform science, technology, and communications public policies.

4. BACKGROUND AND SIGNIFICANCE.

For over ten years CAIDA has undertaken various approaches to narrowing a gap that now impedes the field of network research as well as telecommunications policy: a dearth of available empirical data on the public Internet since the infrastructure has undergone privatization.

As an Internet data analysis and research group largely supported with public funding to apply measurement

and analysis toward understanding and solving globally relevant Internet engineering problems, we accept a responsibility to seek, analyze, and communicate the salient features of the best available data about the Internet.

5. PROGRESS REPORT/PRELIMINARY STUDIES.

As mentioned above, CAIDA has conducted measurement experiments and coordinated collection events for over ten years. The list of publications resulting from the collection, maintenance, and analysis of these datasets is too numerous to catalog in this application. The complete list of publications, presentations, and visualizations can be found on the CAIDA web site at http://www.caida.org/publications/.

To date, we do not know of any untoward effects on any individual(s) as a result of our measurements or data collection, distribution, or publication efforts.

6. RESEARCH DESIGN AND METHODS.

CAIDA conducts measurements of the Internet using two distinctly different methods: active and passive. The following addresses these two methods for data collection, data analysis and data interpretation, respectively. In neither method do we publish data that reveals any personal information about users.

(a) Active methods: CAIDA's Macroscopic Topology Project [1] actively measures the connectivity and latency data for a wide cross-section of the commodity Internet. These measurements contribute to the study of the topology or graph of the Internet.

CAIDA maintains a set of monitors (servers), hosted in sites around the globe, that send probes to trace the route that packets travel on their way through the numerous routers to random destination hosts in the Internet address space. The monitors collect the responses and send them to a central server located at SDSC on the UCSD campus. These monitors act in teams, coordinated by the central server, to distribute the work of probing the millions of destinations.

The collected research data does not obtain information about or involve interaction or intervention with an individually identifiable living human subject. Further, the selection of the destination hosts (identified by IP address) is random, with the only deliberate target being the broad, routed networks on the Internet. This approach ensures that we sample from all routed network prefixes so that our topology measurements are complete and accurately reflect the full scope of the commodity Internet. The destination addresses within

each routed network prefix are arbitrarily selected and probed approximately every 48 hours (one probing cycle), with each probing cycle involving the selection of another randomly selected host. We do this random probing within a large scope of possible hosts (i.e., there are roughly 250,000 routed network prefixes that we separate into roughly 7.4 million targeted networks each containing 256 IP addresses). At the time of writing, approximately 6.5% of probed IP addresses respond.

Any devices that do receive probes are likely to be intermediary devices that effectively mask identification of a specific individual's computer (i.e. IP addresses often resolve to router or NAT (network address translation) machines; and Internet service providers frequently use DHCP (dynamic host configuration protocol) which renders IP addresses non-static and random). In other words, it is likely that there will not be a human on a computer at the end of the random probe. In the unlikely event that the randomly selected address is connected to an end user's computer, that address only identifies the computer, not any of the user accounts associated with the computer or the specific human who operates the device. Furthermore, probing is discontinuous, which means that there is no continuous tracking, repeated monitoring or surveillance of the randomly chosen address. It is worth noting that we use a standard, widely accepted diagnostic protocol, Internet Control Message Protocol (ICMP), for conducting the probes.

Finally, the substance of the collected probe responses does not contain any contents of the packets collected from the destination, only the traffic ("transactional routing control") data. To clarify further, the Internet is a packet switched network in which the basic unit of transmission is a "packet." Each packet contains a "header" and a "payload." Much like a letter sent through the U.S. Postal Service, the payload of each packet is like the contents of the letter and the header acts like the envelope. In a computer network, applications generate the payload of packets (the "contents") and the operating system generates the header (the "addressing information"). For these preceding reasons, there is no "selecting" or "enrolling" a human subject.

Once collected, we copy this traffic data back to the central server where we store and bundle based on the probing cycle. We then annotate the data with information from the Domain Name Service (DNS) which maps the IP address to a human readable domain (network) name. We also generate derived datasets that describe the links between the systems (Autonomous Systems, AS) that route the collected traffic, their inferred relationships with each other, and their taxonomy. We publish these datasets as well as our interpretation of this data in the form of a visualization topology map of the AS-level Internet graph [11].

(b) Passive methods: CAIDA collaborates with organizations that provide local and wide-area network infrastructure. Through these collaborations, we have explicit authorization to passively "tap" heavily aggregated links that provide data packet transport to and from local, statewide, national and international

research and education networks as well as the commodity Internet. [3] The tap involves instrumenting these links with specialized measurement equipment to collect packets, anonymize IP addresses, and analyze packet header traces acquired from these networks. Additionally, we publish web-based reports of aggregated traffic statistics on the monitored link. [12] None of the published information contains any personal information about users.

When we measure a network link, we capture all of the packet headers, or a statistically representative sample of them, using a passive network tap that splits the traffic and provides a copy of all the packets to a host that records the data. The packet headers contain IP addresses that can be used to identify the originating computer (the "source address") and the destination computer (the "destination address") of each packet.

The packet header traces are used to build empirical models of network traffic. These models apply to a variety of analyses such as understanding how applications use networks and how such use changes over time. Also, these models assist researchers who need to generate synthetic network traffic to test and experiment with new hardware and software under conditions representative of real world network conditions.

In addition to authorized passive header collection with collaborators, we employ the UCSD Network Telescope. A network telescope (aka a black hole, an Internet sink, or a darknet) typically has few or often no real computers attached to it and carries almost no legitimate traffic. It serves research value as a monitoring point for anomalous traffic that comprises a significant portion of Internet activity. The network telescope may capture phenomena from a wide range of events, including misconfiguration, malicious scanning of address space, backscatter from random source denial-of-service attacks, and automated spread of malicious software. [4, 5, 6] The network telescope is thus a tool to help researchers identify root causes of this anomalous traffic and has already proven successful in uncovering denial-of-service attack victims and tracking the automated spread of worms. We intend to deploy a single host to store approximately thirty (30) days of data from the Telescope and make those “live” trace files – data made available within an hour of captured – to vetted researchers in accordance with our PS2 framework.

7. HUMAN SUBJECTS.

In the course of the proposed research, we contend that our research does not obtain data through intervention or interaction with individuals, nor does our research involve obtaining private information that is individually identifiable. For these reasons, we advocate that this research does not involve human subjects. The data we obtain does not involve contact with or manipulations of any human subjects’ environments (in this case

Internet traffic), and there are no direct or indirect communications or interpersonal contacts between the researchers and human individuals. The information we collect does not pertain to behavior of persons in contexts where they can reasonably expect that they will be observed or recorded, nor is the obtained data provided for specific purposes by an individual or that which s/he can reasonably expect will not be made public. And, the identity of a subject is not readily ascertained by the researchers.

If any information is incidentally or unintentionally obtained which may be individually identifiable, such information is not private in that it carries no reasonable expectation that it will not be observed or recorded, but nevertheless the subject population would include all potential users of the Internet regardless of geography, age, sex, race, ethnicity or health status. In such a circumstance, the data would exist within aggregated Internet links which involves no intentional or consequential targeting of any specific individual or demographic population. Further, controls are instituted to eliminate the risk of intervention/interaction with any individuals.

8. INFORMED CONSENT.

As described in Section 6, there is no explicit or implicit recruitment, enrolling, selecting, intervention or interaction with, or identification of human subjects. Therefore, plans or procedures for obtaining informed consent are not applicable. Even assuming that the described research activities were deemed to involve the identification of individual human subjects by virtue of identifying a person from an IP address, CAIDA is practically precluded from soliciting consent from those individuals for two reasons, one technical and one methodological. Technically, we measure backbone network traffic links for aggregated populations of millions of users, and it is not viable to request consent, nor is it correct to presume that every IP address correlates to a human subject. Methodologically, informing users of monitoring experiments, e.g., posting placards on computer kiosks on campus stating "This terminal is being monitored as part of a research project" will create a biased sample of usage. Since such bias would invalidate the research, we do not expect to request or receive individual waivers or consents.

9. POTENTIAL RISKS.

The datasets we will collect and make available to researchers represent activities of a large sample of Internet communication traces that may contain data indirectly related to random and arbitrary individuals. Different types of data will be collected, each with its own degree of risk (or non-risk). For the Institutional Review Board to adequately assess this risk, we provide below a general description of how the normal functioning of the Internet generates datasets that can be aggregated and published. In general, we expect that all information

that could potentially identify or link an individual with a certain portion of the data will be modified to remove any risk to individuals.

The data that is generated in the normal functioning of the Internet can include three data types that may have human subject sensitivities:

1. Internet protocol addresses (IPA) of the source and destination computers.

IPA is not individually identifiable information, such that the identity of a user is or may readily be ascertained by the researchers, there is no interaction or intervention with a person who may be indirectly associated with an IPA, and in many cases the IPA data has no expectation of privacy that would cause an individual to reasonably expect that it is not being observed or recorded. The collection and disclosure of IPA is essential to realizing the benefits of this proposed network research activity.

The IPA data that we collect is not identifiable private information such that the identity of a user is or may readily be ascertained by the researchers. In this way, we preempt the risk of interaction/intervention with persons by not collecting data that could be readily used to engage individuals. As described in *Methodology*, IPA often does not identify a specific host on the Internet due to the use of Network Address Translation (NAT) devices and other firewalls which functionally obscure source and/or destination computers by collating them under one IP address. Further, even though every device on the Internet is assigned a unique IP address, not every device points to a host operated by a an identifiable human subject, but rather, may also attach to a 'dumb' machine performing administrative operations to induce traffic to flow across the Internet. Finally, a large portion of IP addresses associated with human users are assigned in a dynamic manner, meaning that a user's IP address while on the Internet today or even several hours in the future may not be the same. So, IP addresses are not static and permanent, features that are paramount to individual identification. While IP addresses could possibly be mapped to a specific computer, or even a specific user account, it is not individually identifiable information as per the criteria that the identity of the subject is readily ascertainable by the researcher. This information would need to be cross-matched with records from various, privately stored and independently secured Internet Service Providers' records which would require deliberate, time and resource-intensive legal process to uncover.

Although IPA that identifies a device could possibly be mapped back to a specific computer that is operated by a human user, such links are not readily ascertainable by researchers. Records that link an IPA to a user account, as opposed to an automated device, are maintained by the Internet Service Provider (ISP) through which the device is registered, and those records are protected from disclosure by state and federal laws. In order to link an IPA with an identifiable individual, a researcher would need to compel an ISP to disclose the mapping to a

registered account name via a court order or subpoena. This would necessitate the filing of a civil lawsuit or criminal charges for a violation of the law. Moreover, these records linking the IPA to a registered account are still not dispositive of the actual users' identity since multiple persons can access the Internet using a known person's account while maintaining anonymity. Certainly there are legitimate and publicly-available search tools that any person could query to obtain the name of the registrant of an IPA. However, these records map back only to the level of the ISP, and not to a specific device, user account or individual users associated with the specific network activity that may be collected in the traffic. So, inherent to the issuing of IPAs, if the computer is identified by an IPA, private information that is individually identifiable is not readily ascertainable by researchers at CAIDA or external researchers to whom we share the data. Given the extremely limited quality and quantity of information, and the difficulty in correlating the IP address to an individual, the collection and disclosure risk presents little potential risk of social stigmatization, psychological injury, or physical harm.

In addition to the above justification for why the proposed activities do not present a risk to human subjects, the set of network data which we will make available (telescope, or darknet traffic described in Section 6) contains IP addresses that are not considered private information that users would reasonably expect to be observed or recorded. Specifically, these IP addresses are collected in the context of a network that is ungoverned network space that does not serve the public. This network does not contain normal, daily communications and transactions (legitimate traffic,) but rather, traffic which by its nature is presumed to predominantly involve either illegitimate activity (i.e., misconfiguration, malicious scanning of address space, backscatter from random source denial-of-service attacks, and/or automated malicious software) or is part of the public record (i.e., Ham radio traffic). This 'illegitimate traffic' is analogous to network pollution in that it is outside the bounds of normal and expected network communications because there are no computers on that network to communicate with the traffic. Thus, there is no "communication" to which a reasonable expectation of privacy can attach.

2. Application payload information.

In general, this information could include IP addresses, Uniform Resource Locators (URLs) from websites, the content (information concerning the substance, purpose or meaning of a communication) of email messages or website communications. CAIDA does not collect any payload content information from legitimate Internet traffic, but instead we strip payload from the packet header traces we do collect. CAIDA does collect payload information from our "darknet" (UCSD Network Telescope), which is Internet space that is ungoverned and does not serve the public. Monitoring unexpected traffic arriving at a network Telescope enables research on remote network events such as various forms of computer attacks, infection of hosts by Internet worms, and network scanning, that would otherwise be impossible to conduct without collecting this payload. As stated immediately

above, we justify the collection of payload from this network under the reasoning that: (a) while there is a chance that it may contain data that links to an identifiable individual, there is no interaction or involvement with an associated identifiable person; and, (b) the data is from non-legitimate traffic in an ungoverned network that does not serve the public. By its nature, such data is beyond the bounds of normal and expected network communications and as such does not invoke the need for authorization to observe and collect. Because all traffic collected from this Telescope network are impermissible communications, any observable payload information that may flow across our monitor does not have a reasonable expectation of privacy/of not being made public. For these reasons, we assess that there is little potential risk of legal liability, reputation damage, psychological or physical harm, or economic damage to an identifiable human subject, and compelling research benefits.

10. RISK MANAGEMENT.

CAIDA collects and publishes datasets in both categories mentioned in Section 9 with careful attention to preventing any unreasonable risk to identifiable human subjects. In addition to the minimal data threat discussed in Section 9 CAIDA utilizes an additional layer of risk prevention and mitigation. In addition to controlling against intervention/interaction with persons arising from collected data, we augment any minimal risk with a privacy sensitive data management infrastructure that prohibits interaction with persons. This means that for both IPA and payload data, we further control against interaction/intervention with an individual by not targeting or tracking any IPAs or identifying data that are collected in our traffic measurement. Specifically, probing or other communication with the IPAs are prohibited by policy and binding agreements, as are any attempt to resolve any of the data to an identifiable individual. This multifaceted framework – the Privacy Sensitive Sharing (PS2) framework also serves to mitigate the potential risk of disclosure of any data that relates to an identifiable individual. It involves proactive and reactive means by which we diminish the linking of an IPA to an individual and ensures that risk of future misuse by disclosing datasets to other researchers is minimized.

CAIDA has since its inception over ten years ago, and prior to IRB oversight, implemented and revised various techniques, policies and procedures for preventing and minimizing risks associated with the collection, analysis, storage and use of measurement data for network research. To continue responsible and risk sensitive research related to evolving efforts described in our Research Application, in 2009 we implemented the PS2 framework. The PS2 is a hybrid control framework -- it integrates privacy-enhancing technologies with an enforceable collection of policies by applying standard privacy principles and obligations (i.e., the Fair Information Practices) in coordination with techniques that provably enforce those requirements. It is thus an umbrella framework

that creates a repeatable, transparent, and scalable set of risk controls, in contrast to previous case-by-case and compartmentalized approach to risk management. The result is a dynamic set of privacy and confidentiality risk controls that address the people, policies and technologies in the lifecycle of our research activities. The PS2 controls against any risk in each dataset by compartmentalizing against the user type and function (who), conditions for accessing the dataset (why), and use and disclosure obligations (where, when, how). The PS2 controls include: data authorization, internal oversight, compliance with laws, adherence to acceptable research purposes, limitation of access to data, disclosure controls, audit tools, and data security.

The primary technology component of the PS2 is anonymization. For active data measurements we mitigate risk of correlating IPAs with individual users by modifying the IPAs using commonly-accepted anonymization techniques. This allows for distinguishing among groups of users while preventing individual identification of a user on the network. In addition to these technical controls, we further eliminate risk of intervention/interaction with persons by executing Acceptable Use Agreements with internal and external researchers who access the data. Prior to releasing any data to a requesting researcher, CAIDA initially controls against potential unauthorized use and disclosure risks by requiring that researchers provide descriptions of intended usage. Researchers who are granted access assert that as a condition to accessing data they agree to abide by the use and disclosure restrictions. These include prohibiting researchers from interacting/intervening with any IPA in the requested dataset.

Specifically, this means that attempting to connect to, probe or in any other way contact a machine or machine administrator associated with an IPA in the dataset are prohibited, as are attempts to resolve the data to an individual whether by reverse engineering, decryption or otherwise. Researchers using the data are also restricted from distributing the data beyond authorized users and are required to anonymize, aggregate or summarize any data that may contain personally identifiable information, IPA, network names, and domain names prior to any publication to protect the privacy of end users. Also, the Acceptable Use Agreement prohibits users of the datasets from de-identifying any anonymized data. And, researchers also assent to use appropriate and reasonable care in safeguarding access to and preventing unauthorized use of the data. The data is secured from unauthorized users by access controls associated with a detailed research user profile which we maintain in a records log and monitor regularly. We use best efforts to enforce compliance with these restrictions via an audit policy that requires the researchers to report a summary of the research and any

findings, publications, or URLs using the data to CAIDA at the conclusion of the research, or semi-annually. Also, we restrict access to users from export-restricted countries. We follow the guidelines provided by the Export Administration Regulations (EAR), International Traffic in Arms Regulations (ITAR), and the Office of Foreign Assets Control (OFAC).

Finally, we have revised our PS2 control framework to include an 'incident response' mechanism, whereby researchers who fail to abide by the data use and disclosure restrictions are required to notify CAIDA of the nature and circumstances surrounding the violation. Having been put on notice, CAIDA will work with researcher to remediate any violations, and will promptly notify the IRB of potential new risks or harms that may result.

11. POTENTIAL BENEFITS.

As national utility infrastructures become intertwined with emerging global data networks, the stability and integrity of the two have become synonymous. This connection, while necessary, leaves network assets vulnerable to the rapidly moving threats of today's Internet. These new threats have impact beyond the scope of the individual enterprise, not only infecting vulnerable hosts (i.e., the individual computers on the network) with malicious code but also denying service to legitimate network users. Fast spreading worms have disrupted financial institutions and emergency services. Inadvertent routing configuration changes have crashed national ISP networks. The enterprise, its upstream ISP(s), and the global Internet community address these threats differently because each has a separate view of the network. Unfortunately, research into Internet-wide or infrastructure level attacks is hampered by a lack of macroscopic datasets. While researchers can often study individual packet traces and compromised machine forensics, these datasets rarely reflect system-level behaviors. Available datasets are often fragmented or difficult to correlate because of missing meta-data or wildly disparate time frames. As part of its participation in the PREDICT community, CAIDA will help address this gap by providing a repository of rich, correlated datasets representing Internet scale behaviors, which will enable qualified cybersecurity researchers to test and prototype novel attack mitigation techniques. Data available from this virtual repository will include both infrastructure data and data from distributed forensic tools.

Our proposal to prototype a model for real-time sharing of Internet traffic data via the Network Telescope is an important attempt to diminish the methodological problems that beset network science due to the dearth of empirical data. While researchers at CAIDA and in the network security community are still working to identify the variety sources of that send traffic into the network telescopes and their utility has been well-established

(i.e., for years researchers have made use of telescopes to identify denial-of-service attack victims, track the spread of Internet worms, and to analyze aberrant packets on otherwise empty address space) no existing or legacy project are providing data to the research community. Our fundamental innovation to data-sharing, informed by a policy-backed and risk-sensitive sharing framework, is seminal to enabling effective and results-oriented cybersecurity and network research.

12. RISK/BENEFIT RATIO.

Based on Section 9: Potential Risks and Section 10: Risk Management, we contend that the controls implemented are proportionate to the risk while not obstructing our research goals. The risks described above are minimal, and represent an impact on privacy that is no greater than the risk that individuals face in their regular use of the Internet. The datasets included in the repository represent large aggregations and pose no additional risk to individual privacy. The datasets that do have the potential to affect individual privacy will be anonymized prior to distribution to researchers using standard anonymization techniques described, and therefore represent an acceptable risk. As mentioned above, we hope the availability of these datasets will strengthen research in network security, workload, traffic classification, performance, topology discovery, and routing.

We have engaged an information technology law advisor, view the risk assessment as an ongoing and iterative process, and invite inquiries and dialogue with the IRB to help establish a framework for all stakeholders to achieve their respective goals while protecting their respective interests.

13. BIBLIOGRAPHY.

1. CAIDA's Macroscopic Topology Project, http://www.caida.org/projects/macroscopic/

2. J. Xu, J. Fan, and M. H. Ammar. Prefix-Preserving IP Address Anonymization: Measurement-based Security Evaluation and a New Cryptography-based Scheme. In EEE ICNP, 2002.

3. CAIDA Internet Data -- Passive Data Sources, http://www.caida.org/data/passive/

4. DeBaecke, D., "Denial of Service Tools and Techniques", University of Memphis 2006 http://umdrive.memphis.edu/ddebaeck/public/Denial%20of%20Service%20Tools%20and%20Techniques.ppt

5. Rajab, M., Monrose, F., Terzis, A., "Worm Evolution Tracking via Timing Analysis", ACM WORM 2006

6. Zesheng Chen; Chuanyi Ji, "Measuring Network-Aware Worm Spreading Ability", INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE Volume , Issue , 6-12 May 2007 Page(s):116 - 124

7. kc claffy, "Ten Things Lawyers Should Know About Internet Research", CAIDA, 2008 http://www.caida.org/publications/papers/2008/lawyers_top_ten/

8. kc claffy, "According to the Best Available Data" http://blog.caida.org/

9. Simson L. Garfinkel, "IRBs and Security Research: Myths, Facts and Mission Creep", USENIX, UPSEC, 2008 http://www.usenix.org/events/upsec08/tech/full_papers/garfinkel/garfinkel_html/

10. CAIDA:Publications:Bibliography:Networking:Anonymization http://www.caida.org/publications/bib/networking/bytopic/index.xml#anonymization

11. Visualizing IPv4 Internet Topology at a Macroscopic Scale http://www.caida.org/research/topology/as_core_network/

12. Passive Network Monitors, http://www.caida.org/data/realtime/passive/

14. OTHER FUNDING. Indicate whether this project is supported by federal, state, or another source. Provide the UCSD grant number and inclusive dates of support. If you have indicated on the face sheet that there is NO funding support for this project, you will need to explain just how the project is to be supported.

2006-0375 - NSF CNS 0551542 CRI: Toward Community Oriented Network Measurement Infrastructure 9/1/06 - 8/31/11 2007-2459 - DHS NBCHC070133 PREDICT Contract 8/1/07 - 7/31/12 2008-0644 - DHS SPAWAR Contract N66001-08-C-2029 Leveraging the Science and Technology of Internet Mapping for Homeland Security 3/21/08 - 9/20/10 2008-0444

15. CONFLICT OF INTEREST. PRINCIPAL INVESTIGATOR'S STATEMENT OF ECONOMIC INTEREST (Form 730-U). [See Attached]

16. Copies of questionnaires,* survey instruments, testing instruments that are not part of standard clinical or educational practice, must be submitted with the application.

We have no questionnaires or surveys to submit at this time.