Project Summary

Internet research critically depends on measurement, but effective Internet measurement raises several daunting issues for the research community and funding agencies. There is increasing awareness that obtaining a better understanding of the structure and dynamics of Internet topology, routing, workload, performance, and vulnerabilities calls for large-scale distributed network measurement infrastructure. But to be viable, such measurement must overcome many challenges: logistical, financial, methodological, technical, legal, and ethical. CAIDA has been navigating these challenges with modest success for ten years, providing data sets to the Internet research and operational community in support of Internet science [1]. However, CAIDA's measurement infrastructure now faces a financial crisis. There is no dedicated source of funds to support macroscopic Internet measurement, particularly in pursuit of rigorous scientific validation of a vast amount of currently unvalidated network research. CAIDA proposes to upgrade both of our current measurement infrastructures (passive and active) to provide the research community data from the wide area Internet that will target the need for validation of current and proposed efforts in large-scale network modeling, simulation, empirical analysis, and architecture development.

Our infrastructure upgrade has two components. First, CAIDA's (and, to our knowledge, the Internet's) last remaining single last point of public insight into the commercial Internet backbone was lost this year. The OC-48 link that a commercial Internet backbone allows us to monitor (at the packet header level) upgraded in February to OC192, well out of our financial ability to monitor. This is devastating news for the Internet research community, who no longer has any source of Internet backbone data upon which to base empirically grounded research. CAIDA proposes to build a set of 10 passive monitoring platforms to capture data in support of both technical and public policy questions of vital interest to the health of the Internet. We already have agreements in place to allow access to strategic infrastructural locations in backbone providers, tools for anonymizing the data, and are now only blocked on funding for the monitors themselves.

Second, CAIDA's (and, to our knowledge, the Internet's) most comprehensive macroscopic Internet topology measurement infrastructure will be out of funding in January 2006, and has been limping along for years without repairs, upgrades or expansions despite the Internet's relentless topological growth. CAIDA proposes to upgrade this infrastructure based on research commuity feedback from the NSF-funded CONMI workshop [2] so that the data gathered will meet the needs of as much of the Internet research community as possible at lowest overall cost.

Maintaining funding for such measurement infrastructure past the span of a given funded research project has eluded the Internet research community. This failure has a huge negative impact on one of CISE's critical interests: supporting not only scientifically sound Internet research but all large-scale networking research that requires empirical validation. **We propose the acquisition of infrastructure that not only directly addresses this failure, but promises measurable progress toward restoring the intellectual strength of a wide range of Internet modeling, simulation, analysis, and theoretical research activities currently occurring without any validation against the real world.** The measurement data gathered from **the proposed infrastructure will enable a wide breadth of CISE-funded and CISE-relevant projects, in at least four categories: support for validation of scientific research; development of new measurement technology; evaluation of proposed future Internet architectures; and empirical answers to questions of critical national security and public policy importance.** CAIDA's long-standing relationships with Internet operational and governance organizations will continue to improve our understanding of the landscape and enable us to better meet the measurement and analysis needs of the network research community as well as the larger Internet.

Project Description

1 Overview: Measurement history, status, problems, and threats

1.1 Historical context of Internet measurement

As the era of the NSFnet Backbone Service came to a close in April 1995, the community lost the ability to rely on what was the only set of publically available statistics for a large national U.S. backbone. The transition to the commercial sector essentially eliminated the public availability of statistics and analyses at a macroscopic national level.

Obstacles to the collection and analysis of traffic data on the commercial Internet pose not only formidable technical and engineering challenges but also include often more daunting legal (privacy), logistical, and proprietary considerations. Data acquisition is further complicated by the networks upgrading to new technologies that are generally prohibitively more expensive or impossible for research groups to monitor. Indeed, statistics collection functionality takes resources directly away from forwarding of packets/frames, which tends to drive commercial providers toward switches from vendors who sacrifice such functionality in exchange for forwarding performance. As a result, core backbone routers simply do not have the functionality to gather fine-grained data needed to support scientifically sound modeling, simulation, and analysis efforts. In combination these issues leave the Internet research community continually starved for data.

For ten years, the NSF has partially funded CAIDA to try to address this problem, by supporting CAIDA in the development and deployment of Internet monitoring hardware and software, as well as supporting curation, analysis, and research based on the resulting data. CAIDA has the only source of OC-48 commercial backbone data provided to the research community, under appropriate non-disclosure and acceptable use agreements.[1]

However, NSF and other U.S. federal agency support has allowed CAIDA to make measurable progress at a goal more directly in line with NSF's critical interest: supporting the scientific community with data to improve the scientific integrity of Internet research. It is this goal, achievable within our current context, that motivates this proposed infrastructure upgrade. Specifically, we want to expand our ability to serve the network research community with data essential to improving the currently abysmal state of Internet science.

Furthermore, more recently the Department of Homeland Security has recognized the need to support the calibration of cybersecurity tools in real world environments, and is in the process of launching the PREDICT Project [3] to allow researchers to request datasets to assist their research into cyber defense technologies, products, models and strategies. The DHS has facilitated progress in the legal and privacy aspects of infrastructure data access, specifically addressing the concerns of ISPs who want to support if not collaborate with the research community but are constrained by privacy laws or policies. But the DHS does not currently fund any measurement infrastructure of the kind proposed here.

In the last decade CAIDA has learned a great deal about both technical and non-technical

[1] We acknowledge that CAIDA had a larger goal of stimulating the formation of an industry-wide neutral body empowered (by its constituent members) to serve as a clearinghouse for analysis and aggregation of potentially sensitive data into formats safe for sharing, as well as to take macroscopic Internet measurements to quantify performance and workload dynamics and trends. Not only did CAIDA not succeed at this loftier goal, but several other attempts at organized data-sharing consortiums, including by non-regulatory federal government agencies, have failed as well, for similar reasons. Namely, while technical measurement challenges exist, the non-technical aspects (legal, economic, privacy, ethical) quickly became, and have remained for a decade, the persistent obstacles to progress in this area.

obstacles to measurement, and we have established relationships and participated in projects with other agencies which continue to improve our understanding of the measurement landscape and enable us better to meet the needs of the network research community as well as the larger Internet.

1.2 Limitations of previous and existing infrastructure

CAIDA is strongly supportive of community-oriented measurement infrastructure, and considers the proposed infrastructure acquisition to be a step in the direction toward providing infrastructure that can satisfy the needs of as much of the scientific community as possible. We have experience in trying to accommodate the needs of a large and diverse research community, and in particular where the most difficult obstacles arise. For active measurement infrastructure, the paramount concerns are how to coordinate measurement requests from a large community of researchers, and how to ensure integrity of the data if gathered by an unknown party. For passive measurement infrastructure, the paramount concerns are cost of hardware for high speed trace collection, and protected access to trace data. In both cases, a community-oriented approach will be necessary to accommodate as many needs of the community as possible as cost-effectively as possible. And yet, given the limited funding currently available to invest in measurement infrastructure, a radical departure from the current approach is premature. We describe relevant current infrastructures for macroscopic Internet measurement and how we propose to maintain their strengths and mitigate their weaknesses as we support the community in developing a more integrated longer term strategy.

1.2.1 Active measurement infrastructures

Several infrastructures have supported or are currently supporting active measurement for research: fully dedicated, such as Skitter [4]; shared with AUPs but more easily accessed, such as PlanetLab [5] or RON [6]; and fully decentralized, such as NetDIMES [7]. Other infrastructures that supported limited active measurement went away after the funding period ended (Surveyor [8], AMP [9], NIMI [10]). The skitter project runs out of funding in January 2006. Each project has different costs, advantages, and weaknesses. There is no consensus on a single correct model for supporting network measurement, and while there is general support for integration of component infrastructures into a higher level platform to be used by the larger Internet community, there is no available agent or funding source to support such an endeavor. Considerable obstacles to such an open platform were discussed at the NSF-funded CONMI workshop [2]: preventing unscrupulous use of the data to facilitate infrastructure attack; forestalling attempts to manipulate measurements to mischaracterize ISP network structure and scope. Further, measurements that generate bad press for ISPs or difficulty for ISP operators may lead quickly to countermeasures that try to deceive or block network measurement efforts.

In this project we offer a compromise for our active measurement infrastructure: support limited flexibility of the infrastructure, in particular implementing more intelligent probing techniques described in the literature [11]. We will continue to make periodic traceroute-derived AS adjacency matrices available for research, education and non-commercial purposes.

1.2.2 Passive measurement infrastructures

For years it has been virtually impossible for researchers to get access to passive (sniffed) data from Internet backbone links due to privacy reasons.[2] Based on trust relationships that have been maintained for 10 years, CAIDA has been able to measure strategic links in the backbone so long is it could provide funding for the monitor. Four times in the last ten years this backbone infrastructure has upgraded beyond the scope of the budget CAIDA has for monitoring (OC3, OC12, OC48, and OC192). As of March 2005 there is no available data on Internet backbone links, and so researchers can no longer analyze Internet backbone workloads. The infrastructure we propose to acquire for this project would solve this ominous problem.

1.2.3 End-user (peer production) measurement infrastructures

Several projects have drawn on the inspiration of SETI@home to develop client-based measurement software for use in a peer production [12] model, taking advantage of end users volunteering their hosts to the measurement infrastructure by downloading and executing measurement software that sends gathered data or statistics back to the project's central processing site. The NSF-sponsored CONMI workshop [2] discussed the challenges faced by three projects in this area: traceroute@home [11]; DIMES [7]; and NETI@home, a passive measurement platform with especially daunting challenges. For active measurements, the biggest challenges are: deployment to insure low impact on the infrastructure; prevention of use of tools for dDOS attacks; accountability of measurement source in case of operational problems; analysis of bias due to self-selection of sources (by volunteers); and validation of the integrity of resulting data. Client-side passive measurement infrastructures have all the same problems as active ones but in addition they bear formidable privacy challenges. For certain measurement questions, such as provider or application prevalence on the macroscopic Internet, there is substantial material incentive to manipulate a macroscopic Internet measurement system, so trust in the integrity of the measurement is essential.

In general, we do not yet understand the methodological problems of scaling Internet measurement to the peer production mode. We propose a strategic approach toward a model for peer production of Internet measurements, that begins with cross-validation of client-produced data with more trusted measurements from controlled infrastructure. Indeed, we hope the next few years can serve as a transition period for the community to determine if we can achieve the same integrity from client-side infrastructures that we now can from controlled infrastructures.

2 Infrastructure to be acquired

CAIDA proposes to build a set of 10 passive OC192/10GE (core Internet backbone) monitoring platforms to directly capture data in support of both technical and public policy questions of vital interest to the health of the Internet. We already have agreements in place to allow access to strategic infrastructural locations in backbone providers, donated colocation space for measurement infrastructure, tools for anonymizing the data, and are now only blocked on funding for the monitors themselves. These monitors will allow research access to backbone links across the continental United States, as well as links connecting the US to Europe and Asia. The monitors will be manufactured by Endace (www.endace.com). This is an OC192/STM-64/10GB monitor.

[2] There are exceptions for researchers who worked directly for a backbone ISP, e.g., Sprint, ATT, but the number of such researchers has diminished over the years due to the meager financial situation of the ISPs.

One FTE system administrator will be needed to build, install, and maintain the passive monitoring system. This systems administrator will deploy ten passive monitoring systems, configure, calibrate, and maintain sensitive measurement cards remotely, and monitor the tapped network links to ensure maximum availability for research efforts.

CAIDA also proposes to replace active measurement monitors that are more than five years old with current infrastructure so that the data gathered will meet the needs of as much of the Internet research community as possible at lowest overall cost to NSF. The active boxes we will be using are Dell PowerEdge 750s.

One half-time programmer will develop and maintain drivers and software for collecting measurements on passive and active systems. As traffic measurement card firmware and link encapsulation change, software to perform reliable measurement collection must continue to evolve. This software is a critical component of the measurement system, as it preserving information about any physical loss or corruption of data and ensures that timestamps applied to packets are preserved accurately. As link speeds exceed bandwidth to disk, this software allows longer-term data collections by using sampling techniques tuned to disk bandwidth with provable error bounds to support subsequent research.

We will use existing storage facilities to store the traces, and index them in CAIDA's Internet Measurement Catalog [13] to be available for public use in late 2005.

3 Research and education activities enabled

A variety of CISE-funded and CISE-relevant projects will be enabled by the measurement data gathered from the proposed infrastructure. We separate these into four classes of projects: support for validation of Internet scientific research, development of new measurement technology, research into future Internet architectures, and empirical research of critical national security and public policy importance.

3.1 Support for validation of Internet modeling, simulation, analysis, and measurement

1. **congestion control models and simulations.** Sally Floyd has carefully described the kind of data required to provide legitimate validation for congestion control modeling [14, 15]. Specifically, proposed congestion control schemes, including those for emerging VOIP and real-time streaming protocols, need to be tested in realistic traffic scenarios, which means researchers need accurate data on: distribution of per-packet round-trip times, relevant to fairness and delay/throughput tradeoffs; distribution of per-packet sequence numbers and of connection sizes, both relevant to burstiness of aggregate traffic. Effective modeling will also require statistics on drop rates as a function of packet and burst size.

   This infrastructure will directly support the needs of the congestion control research community by providing traces of aggregate traffic flows from core Internet backbone links, by far the most difficult but critical scenarios against which to test.

2. **evaluation of proposed transport protocols.** Transport protocol research also needs realistic data on the degree of synchronization of loss events between two TCP flows on the same path. TCP-friendly rate control [16] is a particularly timely example since the IETF is trying to standardize on transport protocols for voice that will be most effective and efficient on the wide-area Internet. Evaluating such proposals requires data comparing drop rates for

large-packet TCP, small-packet TCP, and small-packet UDP on the same path, and paths that transit the core will be of critical importance to such evaluation.

3. **Internet topology models.** Many key publications of Internet topology analysis and modeling [17, 18, 19, 20, 21, 22] have relied on NLANR [23] passive data, funding for which was discontinued in 2005, and RouteViews [24] BGP data. However, several researchers have independently shown that BGP-based topology measurements yield results that are different [25, 26] from those based on traceroute-based measurements [4]. Resolving incongruities within and across these data sources is still an active area of research. Data from the proposed infrastructure will support not only the networking but also on physical science communities, who are actively using results of Internet topology measurements [17, 27, 28, 18, 19, 29, 20, 21, 30, 31, 32, 33, 34, 35, 36, 37, 38, 22].

   In particular, there is still a lively debate in the Internet research community regarding whether an infrastructure with few ($\sim$50) sources but many ($\sim$1M) destinations will sufficiently capture the relevant parameters of topological structure and dynamics. Resolving this debate will be impossible without more strategic measurements of the wide-area Internet, targeted specifically toward validation of proposed hypotheses..

   In addition to extensive measurement and cross-validation, discovering the fundamental laws that govern Internet topology growth will require understanding the economic forces that promote interconnection. Important work has begun in this area [39], and the proposed infrastructure will allow further validation of such developing economic models.

4. **Internet reverse engineering.** CISE is funding a sizable project called "A Shared Facility for Internet Reverse Engineering" [40] draws on Spring *et al.*'s 2003 work [41], but there is no discussion of how the reverse engineered knowledge will be validated. Their objective is ambitious: "the public availability of an unprecedented data set: an annotated map of the entire Internet, complete with a rich set of link, router and operational attributes", but there is no existing infrastructure to support validation of the data gathered, nor has it been made clear how the data will be gathered. The proposed infrastructure could support this project in a number of ways, and CAIDA has offered to support this project with validation as much possible with the available infrastructure.

5. **Internet path characterization and diagnosis.** Diagnosis of Internet path characteristics and problems is a persistently daunting challenge, since the phenomena being measured are occuring in infrastructure not under the control or influence of the person performing the diagnosis. A variety of tools for bandwidth estimation have been developed, but validating their accuracy across real-world paths remains challenging [42] since paths with known link characteristics are required for validation. We have been able to use our current topology measurement infrastructure, in cooperation with Abilene, to evaluate the performance and accuracy of some of these tools, but continued access to active and passive monitors will allow an unprecedented level of validation of bandwidth estimation tools.

   Mahajan *et al.* [43] introduces an architecture for user-level Internet path diagnosis and a practical tool (*tulip*) to diagnose paths in the current Internet. Tulip diagnoses reordering, loss and significant queuing events by leveraging well deployed but little exploited router features. The availability of dedicated active measurement probes can provide an opportunity for validation, refinement, and extension of such tools, and co-location of active probes with

passive monitors along a path can allow integrated techniques that will improve the integrity of the models and tools.

6. **Inference of Internet routing policies.** Recent techniques for inferring business relationships between ASs have yielded maps that have few 'invalid BGP paths' in the terminology of Gao [44]. However, although these algorithms are improvements over their predecessors, some relationships inferred by these newer algorithms are incorrect, leading to the deduction of unrealistic AS hierarchies. CAIDA has investigated this problem, discovered what caused it, and generalized the problem of AS relationship inference as a multi-objective optimization problem with node-degree-based corrections to the original objective function of minimizing the number of invalid paths. We solved the generalized version of the problem using the semidefinite programming relaxation of the MAX2SAT problem. Keeping the number of invalid paths small, we obtained a more veracious solution than yielded by previous heuristics. Validation of routing policy inference techniques is still an essential component of measuring progress in the field of topology analysis, and a dedicated infrastructure for active probing measurement will support this progress.

3.2 Support for development of new measurement technology

1. **Active measurement tools.** CAIDA has three main goals with its own topology measurement activities: provide the Internet topology data to the community [45, 46]; analyze the statistical properties of Internet topologies [47]; and construct equilibrium models for Internet graphs. And yet it important to acknowledge that the accuracy of all current tools used for wide area macroscopic topology measurement [4, 7, 48] is under dispute in the research community [47, 49]. As stated in section 1.2.3, CAIDA seeks to support the research community in using this infrastructure to explore the use of more efficient and validated algorithms for large-scale topology discovery, starting with implementing the intelligent probing techniques described in the traceroute@home project [11]. The proposed infrastructure will also provide a source of data against with to validate existing client-side topology measurement projects [7] which have thus far had no validation.

   CAIDA is also undertaking the development of a probing tool that will characterize topology at a POP (point of presence) granularity. Pop-level traceroute (PLT) is analogous to traceroute except working at the POP level. PLT will show the POPs in the forwarding path from oneself to a given destination. PLT will work by segmenting the router-level forwarding path obtained with standard traceroute into POPs based on several heuristics and measurement. PLT will provide a glimpse into AS topology and AS peering at the granularity of POPs. Indeed, without progress on PLT techniques the network research community will be unable to make breakthrough progress in modeling of intradomain and interdomain routing, including inference and analysis of fine-grained peering relationships. More broadly, research on PLT is a requisite step toward ultimately developing an Internet-wide pop-level map, which will allow empirically grounded answers to a multitude of relevant architectural and policy questions about the Internet. This infrastructure will provide a place to test and refine CAIDA's and others' pop-level topology probing tools.

2. **Passive measurement tools.** Good performance under extreme workloads and isolation of resource consumption of concurrent jobs are perennial design goals of computer systems ranging from multitasking servers to network routers. CAIDA has developed a specialized

system that computes multiple summaries of IP traffic in real time and achieves robustness and isolation between tasks in a novel way: by automatically adapting the parameters of the summarization algorithms [50]. In traditional systems, anomalous network behavior such as denial of service attacks or worms can overwhelm the memory or CPU, making the system produce meaningless results exactly when measurement is needed most. In contrast, our measurement system reacts by gracefully degrading the accuracy of the affected summaries.

We have evaluated many existing algorithmic solutions for computing traffic summaries, as well as two of our own solutions that offer better memory versus accuracy tradeoffs and have more predictable resource consumption. We have also evaluated the actual implementation of a protoype system that combines the best of these algorithms. The proposed infrastructure provides an ideal place to test and further improve these implementations.

3. **Trajectory Sampling**

   Trajectory sampling [51] is a measurement technique that allows the inference of a traffic flow via observation of the trajectories of a subset of all packets in a flow as it traverses a path. This approach samples packets based on a hash function computed over the packet content, allowing reconstruction of packet trajectories. Experimenting with these novel techniques requires instrumentation of multiple nodes along a path, a circumstance the research community does not have access to without dedicated measurement infrastructure for this purpose, e.g., what we propose here.

4. **IETF's packet sampling protocol development**

   The IETF's Packet Sampling (PSAMP) working group [52] is in the process of defining a standard set of capabilities for network elements to sample packets at high speed but with rich filtering semantics. The goal of this effort is to have network devices support multiple parallel packet samplers, each with independently configurable filters, reports, and export. The proposed infrastructure would support experimentation and refinement of packet sampling algorithms under realistic traffic loads, in pursuit of developing functionality that can eventually be deployed standard on device line cards.

3.3 Support for empirical research of critical public policy questions

There are several areas where empirical macroscopic Internet measurement could directly inform critically important public policy debates.

1. **Workload characterization and calibration.** Various measurements over the years have offered confirmation of the heavy-tailed distribution of Internet flow sizes, but this knowledge has had little impact on operational reality. There is growing interest in capturing and analyzing Internet traffic characteristics in pursuit of insights that will positively affect operations, but opportunities to gather such data have diminished almost completely since the privatization of the infrastructure. Since its establishment CAIDA has been able to measure strategic links in the backbone so long is it could provide funding for the monitor. Four times in the last ten years this backbone infrastructure has upgraded beyond the scope of the budget CAIDA has for monitoring (OC3, OC12, OC48, and OC192). As of March 2005 there is no available data on Internet backbone links, and so researchers can no longer analyze Internet backbone workloads. If researchers are lucky enough to gather workload data from local

campuses, they have no means to calibrate it with commercial Internet backbone traffic. The proposed infrastructure addresses this problem.

2. **Analysis of the spread of Internet worms.** It has been clear since 1988 that self-propagating code can quickly spread across a network by exploiting homogeneous security vulnerabilities. However, the last few years have seen a dramatic increase in the frequency and virulence of such worm outbreaks. For example, the Code-Red worm epidemics of 2001 infected hundreds of thousands of Internet hosts in a short period - incurring enormous operational expense to track down, contain, and repair each infected machine. In response to this threat, considerable effort has focused on developing technical means for detecting and containing worm infections before they can cause such damage.

   CAIDA used its empirical topology data to investigate how well worm containment approaches would work on actual Internet infrastructure. Using a combination of analytic modeling and simulation, we described how various design factors impact the dynamics of a worm epidemic and, conversely, the minimum engineering requirements necessary to contain the spread of a given worm. While our analysis cannot provide definitive guidance for engineering defenses against all future threats, we demonstrated the lower bounds that any such system must exceed to be useful today. Unfortunately, our results suggest that there are significant technological and administrative gaps to be bridged before an effective defense can be provided in today's Internet. Data from the proposed measurement infrastructure will support continued research into the viable protections and responses of Internet infrastructure to virulent worms.

3. **Analysis of the Internet Service Provider hierarchy.** Analysis of the the Internet Service Provider (ISP) hierarchy is critical to a deeper understanding of technical, economic and regulatory aspects of the Internet inter-domain routing system. As part of CAIDA's research agenda to measure and analyze macroscopic Internet structure, we have developed and refined our procedure to rank Autonomous Systems (AS Rank) by their location in the Internet hierarchy [53]. Our ranking relies upon AS relationship information that we discover using our new inference algorithms [54], rooted in economic AS relationships. Specifically, we rank each AS as a function of the number of IP prefixes advertised by this AS, its customer ASes, their customers ASs, and so on. **This analysis is critically dependent on the accurate collection of a vast amount of Internet topology data. Without the infrastructure proposed in this project, we know of no other objective analysis of commercial ISP coverage of the Internet's topology.**

4. **Prevalence and growth of P2P file sharing.** Since the emergence of peer-to-peer (P2P) networking in the late '90s, P2P file shsaring applications have multiplied, evolved and established themselves as the leading growth application of Internet traffic workload. The presence of this traffic on the Internet has catapulted into the focus of public policy, legislative, and ethical discussions in the mainstream media, including U.S. Supreme Court case (MGM vs. Grokster). In particular, the RIAA and others have published claims that P2P file sharing traffic on the Internet is dropping as a result of the RIAA's criminal lawsuits [55, 56]. These conclusions are based on no actual Internet traffic data, but rather on telephone surveys or software downloads of a single P2P application.

   CAIDA has invested considerable effort in the objective analysis of P2P traffic identification and growth, in an attempt to ground policy discussions with as accurate empirical data and

the most scientifically sound analysis possible. The identification of P2P file sharing traffic is challenging from a scientific perspective. In contrast to first-generation P2P networks which used well-defined port numbers, current P2P applications have the ability to disguise their existence through the use of arbitrary ports (application identifiers). As a result, reliable estimates of P2P traffic require examination of packet payload, a methodological land mine from legal and privacy perspectives. Indeed, access to user payload is often rendered impossible by these factors, inhibiting trustworthy estimation of P2P traffic growth and dynamics. To deal with these issues, CAIDA has developed a systematic methodology to identify P2P flows at the transport layer [57], i.e., based on connection patterns of P2P networks, without relying on packet payload. We believe our approach is the first method for characterizing P2P traffic using only knowledge of network dynamics rather than any user payload. To evaluate our methodology, we also developed a payload technique for P2P traffic identification, by reverse engineering and analyzing the nine most popular P2P protocols, and demonstrating its efficacy with the discovery of P2P protocols in our traces that were previously unknown to us. Our results, including a follow-up study [58], demonstrate that P2P traffic continues to grow unabatedly, contrary to claims made in the popular media [55, 56]. The availability of neutral data sources and objective analyses from the most aggregated, statistically multiplexed points in the network, i.e., the core backbones, will inform one of the most important public policy questions of the 21st century: will either litigious or technical (DRM) approaches succeed in thwarting the explosion of P2P file sharing on the global Internet?

3.4 Support for research into future Internet architectures

As the National Science Foundation considers its role in supporting the research and development of new and innovative Internet architectures [59, 60], it will be important to draw on lessons that we have learned about the current Internet, in particular its persistent weaknesses, of which measurement is a huge one. It would be unwise to undertake such an ambitious venture without clearly identifying the roots of the problems with the current Internet. Research areas include: validating or refuting [implicit] assumptions about the current network (traffic, naming, routing, security) that are driving its evolution; and applying what we have learned from Internet measurement to the design of new network architectures, including how to facilitate the kind of measurement needed to support architectural goals [61].

Without applied experience with real-world Internet measurement, researchers will be addressing such questions in a vacuum. The proposed infrastructure will allow a larger fraction of the research community access to data that would inform discussion of measurement questions that are increasingly relevant to not just the Internet, but future large-scale networks.

4 Limitations of proposed infrastructure

We recognize that even if this project is funded, it will only facilitate insights into small pieces of the larger picture. The proposed infrastructure does not cover wireless data, IGP data, traffic matrix data. The active probing will run into the same methodological problems that all probing does: filtering, layer-2 technologies that hide layer 3 information, e.g., MPLS, and inconsistent non-standard behavior in responding to probes that can mislead topology inferences. But ten years of Internet research has made it clear that gathering enough data to improve the state of Internet science at a broad scale would take substantially more funding than CISE has available, so for now

we seek strategic areas that can offer us the greatest potential for a wide range of scientific progress at the lowest cost to the taxpayers.

5 Relationship to other CISE and CAIDA projects

Consistent with goals of the solicitation, this project will: (1) increase the recognition of the research group (CAIDA) in the host institution (SDSC) and national community (academic, government, and industry); (2) increase industry participation and increase the opportunities for cooperation in the CISE research community; and (3) enable the proposers as well as the broader community to undertake important work that would not be possible without the infrastructure. Table 1 lists a sample of projects that will be enabled by the proposed infrastructure.

Table 1: Relationships between community projects and proposed infrastructure

research area	community project
validation	congestion control models and simulation
	evaluation of proposed transport protocols
	internet toplogy models
	internet reverse engineering
	path characterization and diagnosis
	inferring internet routing policies
measurement technology	topology measurement tools
	passive fbw measurement tools
	IETF's PSAMP sampling protocol development
empirical analysis	workload characterization
	application cross-section
	analysis of spread of Internet worms
	relative ISP coverage of Internet topology
	prevalence and growth of P2P fi le sharing
architectural	measurements to support future Internet architectures

6 Why CAIDA is the most appropriate team for this project

CAIDA is recognized as a world leader in Internet measurement and data analysis, and has provided several landmark studies of Internet performance, workload, and topology issues [62]. CAIDA has years of experience in development, implementation, and evaluation of measurement infrastructure, as well as with anonymization and analysis tools for the gathered data. CAIDA's longstanding trust relationships with many Internet service providers and equipment vendors facilitate monitor deployment and informed analyses. To technical, operational, and policy communities, CAIDA is among the most trusted sources of objective measurement tools and analyses.

With regard to software infrastructure support, CAIDA has recently developed a new tool, *whale* [63], for tracking data sets, including backup timestamps, why and with what configuration a data set was collected, whether the data has been made publically available, and automatic generation and posting of statistics about collected data sets each night to CAIDA's data page [1]. *Whale* also maintains monthly and cumulative totals of per-project and overall data collected, allows comparison of filesystem trees to verify backups, and stores md5 checksums to ensure data integrity.

For each packet trace, *whale* stores statistics about the number of packets, bytes, and flows in passive traces, such that we can easily locate traces by such attributes. *Whale* also supports flexible grouping of files by filesets, projects, and experiments, where files can be part of multiple projects to support correlation across heterogenous data sets.

Housed at SDSC, CAIDA represents a unique combination of talents and facilities necessary to achieve the proposed goals. Moreover, the involvement of the DHS into this area (via the PREDICT project [3] in which CAIDA participates) has triggered significant progress with the legal and privacy issues, which will further increase the potential accessibility and thus utility of this data to the community.

6.1 Integration of Research and Education: Workshops

Each year of the project, CAIDA will host a workshop specifically targeted at measurement needs for validation of modeling and simulation, bringing together researchers, preferentially graduate students, with measurement experts and operators who can provide insights into collected data sets. These workshops will provide an opportunity to present research using the data collected from the infrastructure, get feedback on operational idiosyncrasies of the data, and discuss new or changing measurement needs to support validation tasks for the following year.

6.2 Integration of Research and Education: Internships

CAIDA offers a unique opportunity for students to gain experience with massive (and messy) datasets that require both sound methodology and efficient management of computing resources to analyze. During summers, CAIDA hosts several graduate students from other institutions in order to acquaint Internet researchers with our data sets and to advance mutually beneficial collaboration. We expect to have 2-4 summer students each year of the grant, at least half of them from institutions other than UCSD. CAIDA also encourages sabbatical visits from industry engineers [64].

6.3 Integrating Diversity into CAIDA Activities

Based at UC, San Diego, CAIDA has a strong record of integrating diversity into our activities. Since July 1999, the composition of our 54 paid interns has included 15 females, 30 Asian, and 2 Hispanic students. Our 17 volunteer interns in that same period have included one female and 7 Asian students.

7 Results from prior support

1. **CAIDA: Cooperative Association for Internet Data Analysis.** ANI-9711092. $3,199,580. Sep 1997 - Aug 2002. (Claffy) This collaborative undertaking brings together organizations in the commercial, government, and research sectors. CAIDA provides a neutral framework to support cooperative technical endeavors, and encourages the creation and dissemination of Internet traffic metrics and measurement methodologies. Results of this collaborative research and analytic environment can be seen on published web pages on the CAIDA web site www.caida.org. CAIDA also develops advanced Internet measurement and visualization tools.

2. **Internet Atlas.** ANI-99-96248 $468,834. Jan 1999 - Dec 2002. (Claffy) This effort involves developing techniques and tools for mapping the Internet, focusing on Internet topology, performance, workload, and routing data. A gallery that assesses state-of-the-art in this nascent sector is published on the web.

3. **Correlating Heterogeneous Measurement Data to Achieve System-Level Analysis of Internet Traffic Trends.** ANI-0137121, $1,013,794 Sep 2002 - Aug 2005 (Claffy and Moore) As it grows, the Internet is becoming more fragile in many ways. The complexity in managing or repairing damage to the system can only be navigated with sustained understanding of the evolving commercial Internet infrastructure. The research and tools proposed under this effort lead to such insights. In particular, richer access to data will facilitate development of tools for navigation, analysis, and correlated visualization of massive network data sets and path specific performance and routing data that are critical to advancing both research and operational efforts. concentration of administration of Internet infrastructure.

4. **Routing and Peering Analysis for Enhancing Internet Performance and Security.** ANI-0221172, $882,999 Oct 2002 - Sep 2005 (Claffy) CAIDA performs topology analysis and characterizes sources of growth and instability of the routing system, applying graph theory and comparing combinatorial approaches for identifying strategic locations in the macroscopic Internet.

5. **Quantitative Network Security Analysis.** CCR-0311690, $384,183 Aug 2003 - Jul 2005. (Moore) Much information about the state of large-scale malicious activity on the Internet is anecdotal. Under this grant, we are developing a combination of network analysis techniques and network measurement infrastructure to analyze large-scale Internet security threats, such as denial of service attacks or Internet worms. In addition to our own research and analysis of these events, datasets of interesting events collected by the UCSD Network Telescope have been made available to other researchers.

6. **New Directions in Accounting and Traffic Measurement.** ANI-0137102, $649,754 Sep 2002 - Aug 2006 (Moore) As network link bandwidths increase, the ability to measure every single packet meaningfully in an operational setting decreases. To assist, we have developed several novel techniques for generating accurate measurement reports which degrade gracefully under adverse network traffic conditions. One of these approaches, Adaptive NetFlow, was designed to be implementable in routers and produce reports which are essentially the same as those typically collected operationally today.

7. **SCI: ITR-(NHS+EVS)-(dmc+SIM): Improving the Integrity of Domain Name System (DNS) Monitoring Trends.** SCI-0427144, $3,397,981 Sep 04 - Aug 06 (Claffy) This project helps to address National and Homeland Security recommendations by the President's Critical Infrastructure Protection Board to develop a 'cyberspace network operations center (NOC)'. The long-term mission of this proposal - to provide data needed to support DNS research - also has relevance to the real Internet and how it supports economic prosperity and a vibrant civil society. Indeed, the data, models, communications analysis, and simulation functionalities to be provided have the potential to dramatically improve the quality of the lens with which we view the Internet as a whole.

8. **NeTS-NR Toward Mathematical Rigorous Next-Generation Routing Protocols for Realistic Network Topologies.** CNS-0434996, $900,000 Oct 04 - Sep 07 (Claffy and Kri-

oukov) CAIDA proposes to open a new area of research focused on applying key theoretical routing results in distributed computation to extremely practical purposes, i.e. fixing the Internet. Our agenda is ambitious, but firmly justified by a set of several previous results, all spectacularly unexpected, which have revealed a huge gap in our fundamental understanding of data networks. Our agenda has three related and clearly defined tasks: 1) execute the next step on the path toward construction of practically acceptable next-generation routing protocols based on mathematically rigorous routing algorithms; 2) validate the applicability of the above algorithms against several sources of real Internet topology data; 3) build and evaluate a model for Internet topology evolution, which reflects fundamental laws of evolution of large-scale networks.

References

[1] Cooperative Association of Internet Data Analysis, "CAIDA Internet Data Repository."

[2] M. Crovella and J. Michel, "NSF sponsored workshop: Community-Oriented Network Measurement Infrastructure." http://www.caida.org/outreach/workshops/conmi/, March 2005.

[3] D. of Homeland Security, "PREDICT project: Protected Repository for Defense of Infrastructure against Cyber Threats." http://www.predict.org/.

[4] CAIDA, "skitter active measurement tool." http://www.caida.org/tools/measurement/skitter.

[5] L. Peterson, T. Anderson, D. Culler, and T. Roscoe, "A blueprint for introducing disruptive technology into the Internet," in *Hotnets*, pp. 59–64, 2002.

[6] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris, "Resilient overlay networks," in *SOSP*, pp. 131–145, 2001.

[7] Y. Shavitt and E. Shir, "DIMES: Let the Internet Measure Itself," tech. rep., Tel-Aviv University EE Department, 2005.

[8] S. Kalidindi and M. J. Zekauskas, "Surveyor: An infrastructure for Internet performance measurements," in *INET'99*, June 1999.

[9] "Active Measurement Project." http://amp.nlanr.net/.

[10] V. Paxson, A. Adams, and M. Mathis, "Experiences with NIMI," in *Passive and Active Measurement*, Apr. 2000.

[11] J. I. Alvarez-hamelin, A. Barrat, M. Crovella, B. Donnet, T. Friedman, M. Latapy, P. Raoult, and A. Vespignani, "Traceroute@home project," August 2005.

[12] Y. Benkler, "Coase's penguin, or, linux and the nature of the firm," *The Yale Law Journal*, vol. 112, 2002.

[13] CAIDA, "Internet Measurement Data Catalog Project." http://www.datcat.org/.

[14] S. Floyd, "Requirements for Simulation and Modeling Tools," August 2005.

[15] S. Floyd, "Questions about the Internet," May 2002.

[16] S. Floyd and E. Kohler, "TCP Friendly Rate Control (TFRC) for Voice: VoIP Variant," July 2005.

[17] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On power-law relationships of the Internet topology," in *ACM SIGCOMM*, pp. 251–262, 1999.

[18] H. Tangmunarunkit, R. Govindan, S. Jamin, S. Shenker, and W. Willinger, "Network topology generators: Degree-based vs. structural," in *ACM SIGCOMM*, pp. 147–159, 2002.

[19] Q. Chen, H. Chang, R. Govindan, S. Jamin, S. J. Shenker, and W. Willinger, "The origin of power laws in Internet topologies revisited," in *IEEE INFOCOM*, 2002.

[20] H. Chang, S. Jamin, and W. Willinger, "Internet connectivity at the AS level: An optimization driven modeling approach," in *Proceedings of MoMeTools*, 2003.

[21] R. Albert and A.-L. Barabási, "Statistical mechanics of complex networks," *Reviews of Modern Physics*, vol. 74, pp. 47–97, 2002.

[22] S. N. Dorogovtsev and J. F. F. Mendes, *Evolution of Networks: From Biological Nets to the Internet and WWW*. Oxford: Oxford University Press, 2003.

[23] National Laboratory for Applied Network Research. `http://www.nlanr.net/`.

[24] David Meyer, "University of oregon route views project." http://www.routeviews.org/.

[25] Y. Hyun, A. Broido, and k claffy, "Traceroute and BGP AS path incongruities," in *Cooperative Association for Internet Data Analysis (CAIDA)*, 2003. `http://www.caida.org/outreach/papers/2003/ASP/`.

[26] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz, "Towards an accurate ASlevel traceroute tool," in *ACM SIGCOMM*, 2003.

[27] G. Siganos, M. Faloutsos, P. Faloutsos, and C. Faloutsos, "Power-laws and the AS-level Internet topology," *ACM/IEEE Transactions on Networking*, vol. 11, no. 4, pp. 514–524, 2003.

[28] T.Bu and D. Towsley, "On distinguishing between Internet power law topology generators," in *IEEE INFOCOM*, 2002.

[29] W. Willinger, R. Govindan, S. Jamin, V. Paxson, and S. Shenker, "Scaling phenomena in the Internet: Critically examining criticality," *PNAS*, vol. 99, no. Suppl. 1, pp. 2573–2580, 2002.

[30] S. Zhou and R. J. Mondragon, "Towards modelling the Internet topology - the Interactive Growth Model," in *Proceedings of the 18th International Teletraffic Congress (ITC18)*, (Berlin), 2003. `http://arxiv.org/abs/cs.NI/0303029`.

[31] A. Vázquez, R. Pastor-Satorras, and A. Vespignani, "Internet topology at the router and Autonomous System level," `http://arxiv.org/abs/cond-mat/0112400`.

[32] A. Vázquez, R. Pastor-Satorras, and A. Vespignani, "Large-scale topological and dynamical properties of the Internet," *Physical Review E*, vol. 65, no. 06, p. 066130, 2002. `http://arxiv.org/abs/cond-mat/0112400`.

[33] A. Capocci, G. Caldarelli, R. Marchetti, and L. Pietronero, "Growing dynamics of Internet providers," *Physical Review E*, vol. 64, p. 35105, 2001.

[34] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, "Pseudofractal scale-free web," *Physical Review E*, vol. 65, no. 06, p. 066122, 2002.

[35] S.-H. Yook, H. Jeong, and A.-L. Barabási, "Modeling the Internet's large-scale topology," *PNAS*, vol. 99, pp. 13382–13386, 2002.

[36] K.-I. Goh, B. Kahng, and D. Kim, "Fluctuation-driven dynamics of the Internet topology," *Physical Review Letters*, vol. 88, p. 108701, 2002.

[37] G. Caldarelli, A. Capocci, P. D. L. Rios, , and M. A. M. noz, "Scale-free networks from varying vertex intrinsic fitness," *Physical Review Letters*, vol. 89, p. 258702, 2002.

[38] G. Caldarelli, P. D. L. Rios, and L. Pietronero, "Generalized Network Growth: from microscopic strategies to the real Internet properties." `http://arxiv.org/abs/cond-mat/0307610`.

[39] Emanuele Giovannetti and Alessio D'Ignazio and Joerg Lepler, " Report on Vertical and Geographical Market Boundaries for the European Internet ," tech. rep., University of Cambridge CoCombine Project Report, 2005.

[40] T. Anderson and D. Wetherall, "A shared facility for internet reverse engineering," 2004. NSF awarded NeTS proposal.

[41] N. Spring, D. Wetherall, and T. Anderson, "Reverse engineering the internet," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 1, pp. 3–8, 2004.

[42] A. Shriram, M. Murray, Y. Hyun, N. B. e, A. Broido, M. Fomenkov, and k claffy, "Comparison of Public End-to-end Bandwidth Estimation Tools on Hig h-Speed Links," in *PAM 2005, to appear*, Apr. 2005.

[43] R. Mahajan, N. Spring, D. Wetherall, and T. Anderson, "User-level internet path diagnosis," in *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles*, (New York, NY, USA), pp. 106–119, ACM Press, 2003.

[44] L. Gao and F. Wang, "Inferring and characterizing internet routing policies," in *ACM SIGCOMM Internet measurement workshop*, april 2003.

[45] CAIDA, "As-level topology data extracted from continuous traceroute (skitter) measurements."

[46] CAIDA, "Router-level topology data extracted from continuous traceroute (skitter) measurements."

[47] P. Mahadevan, D. Krioukov, M. Fomenkov, B. Huffaker, X. Dimitropoulos, and kc claffy, "Lessons from three views of the internet topology: Technical report," tech. rep.

[48] N. Spring and R. Mahajan, "rocketfuel: an ISP topology mapping engine," 2002. `http://www.cs.washington.edu/research/networking/rocketfuel/`.

[49] R. Teixeira, K. Marzullo, S. Savage, and G. M. Voelker, "In search of path diversity in isp networks," in *IMC '03: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, (New York, NY, USA), pp. 313–318, ACM Press, 2003.

[50] C. Estan, K. Keys, D. Moore, and G. Varghese, "Building a better netflow," in *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, (New York, NY, USA), pp. 245–256, ACM Press, 2004.

[51] N. G. Duffield and M. Grossglauser, "Trajectory sampling for direct traffic observation," in *SIGCOMM '00: Proceedings of the conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, (New York, NY, USA), pp. 271–282, ACM Press, 2000.

[52] CAIDA, “Ietf PSAMP working group.”

[53] CAIDA, “Ranking of internet service providers by observed topology,” 2005.

[54] X. Dimitropoulos, D. Krioukov, B. Huffaker, kc claffy, and G. Riley, “Inferring as relationships: Dead end or lively beginning?,” in *4th Workshop on Efficient and Experimental Algorithms*, 2005.

[55] PC World, “Are the RIAA’s Lawsuits Working?,” January 2004.

[56] Pew Internet Life, “The RIAA lawsuits against online music file sharers appear to have had a devastating impact,” January 2004.

[57] Thomas Karagiannias and Andre Broido and Michalis Faloutsos and kc claffy, “Transport Layer Identification of P2P Traffic,” in *Internet Measurement Conference*, October 2004.

[58] Thomas Karagiannias and Andre Broido and Nevil Brownlee and Michalis Faloutsos and kc claffy, “Is P2P dying or just hiding,” in *Globecom*, November 2004.

[59] N. workshop report, “Overcoming barriers to disruptive innovation in networking,” tech. rep., January 2005.

[60] M. Baard, “Net pioneer wants new internet,”

[61] “fima@caida.org, mailing list for discussion of Future Internet Measurement Architectures.”

[62] CAIDA. `http://www.caida.org/outreach/papers/`.

[63] Emile Aben, “Whale, CAIDA’s data set catalog management software,” August 2005.

[64] CAIDA, “Caidabbatical program.” `http://www.caida.org/home/jobs/caidabbatical.xml`.