The 3rd Workshop on Active Internet Measurements (AIMS-3) Report

kc claffy
CAIDA
kc@caida.org

This article is an editorial note submitted to CCR. It has NOT been peer reviewed. The author takes full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

On February 10-12, 2011, CAIDA hosted the third Workshop on Active Internet Measurements (AIMS-3) as part of our series of Internet Statistics and Metrics Analysis (ISMA) workshops. As with the previous two AIMS workshops, the goals were to further our understanding of the potential and limitations of active measurement research and infrastructure in the wide-area Internet, and to promote cooperative solutions and coordinated strategies to address future data needs of the network and security research communities. For three years, the workshop has fostered interdisciplinary conversation among researchers, operators, and government, focused on analysis of goals, means, and emerging issues in active Internet measurement projects. The first workshop emphasized discussion of existing hardware and software platforms for macroscopic measurement and mapping of Internet properties, in particular those related to cybersecurity. The second workshop included more performance evaluation and data-sharing approaches. This year we expanded the workshop agenda to include active measurement topics of more recent interest: broadband performance; gauging IPv6 deployment; and measurement activities in international research networks.

Categories and Subject Descriptors

C.2.3 [**Network operations**]: Network monitoring; C.2.5 [**Local and Wide-Area Networks**]: Internet; C.2.6 [**Internetworking**]: Standards; C.4.2 [**Performance of Systems**]: Measurement techniques—Active

General Terms

Measurement, Management, Human Factors, Legal Aspects, Standardization, Performance, Verification

Keywords

active measurement, Internet measurement techniques, management techniques, validation

1. MOTIVATION

The AIMS workshop series was established to help stakeholders in Internet active measurement projects to communicate their interests and concerns, and explore cooperative approaches to maximizing the collective benefit of deployed infrastructure and gathered measurements. The final report from the first AIMS workshop [13] outlined open research problems identified by participants, and issued recommendations that could benefit both Internet science and communications policy. These recommendations represent a multi-year roadmap of the landscape with specific suggestions for paths to advance the quality, science, and utility of active Internet measurements. The AIMS workshop series provides a forum to track and evaluate progress on this roadmap, build on previous achievements, refine our understanding of remaining problems and recognize new ones, modifying the course of progress as necessary.

The first two AIMS workshops focused on interaction among researchers involved in recent advances and challenges in active Internet measurement, particularly those related to cybersecurity, as well as exchange of insights among researchers, operators, and policymakers. This year we expanded the agenda to include active measurement topics of more recent interest: broadband performance; gauging IPv6 deployment; and measurement activities in international research networks. This report briefly describes topics covered at the workshop, and reviews progress on previous recommendations [13]. Slides and other materials related to the workshop are available at [8].

2. KEY FINDINGS

The workshop achieved its goals: reviewing objectives, findings, techniques, and plans of various active measurement projects; discussing issues inhibiting progress; identifying priorities, strengths, and weaknesses of measurement efforts; and fostering collaboration, coordination, and data sharing among participants. Geolocation database accuracy and evaluation, infrastructure platforms, topology measurement and mapping techniques including for IPv6, alias resolution, end-to-end performance measurement, and data sharing needs and approaches were among topics discussed at the workshop.

The field of annotated Internet topology mapping continues its slow but steady progress.
Active measurements of the Internet are used to discover, model, and generate realistic macroscopic Internet topologies at various granularities, e.g., IP addresses, routers, autonomous systems (ASes). Advances in measurement, inference, and validation methods, as well as the increasing availability of historical meta-data, have allowed quantitative as well as qualitative improvement in the best available Internet topology maps [9, 16, 20, 25, 24, 27, 32]. Workshop participants presented improved techniques for more efficient adaptive topology probing [6] and for measuring the impact of falsely inferred links [27]. CAIDA is now making available periodic "Internet Topology Data Kits", heavily curated (to AS-level and router-level) data about connectivity and routing gathered from a large cross-section of the global IPv4 Internet over a certain time period (for recent ITDKs, about two weeks). Richer topologies are helping researchers avoid some known pitfalls of infrastructure-based measurements from limited vantage points [10].

Mehmet Engin Tozal and Kamil Sarac from University of Texas at Dallas presented tools for network layer Internet topology construction, including TraceNet, for building subnet-level Internet topologies and XNET, for inferring individual subnets. Mehmet Burak Akgun described a subnet-based Internet topology generator, emphasizing the distinction between the observed degree

distribution and the real degree distribution as it impacts performance dependencies and characteristics. Hakan Kardes and Mehmet Gunes presented Cheleby, a system for probing and resolving subnets, IP aliases, and unresponsive routers to provide sample link-level topologies.

Gabor Vattay presented a talk on Benford-type distributions in Internet data. He argued that the extreme power laws seen in Internet topology degree distribution and traffic correlations might signal a more general scale invariance in the Internet, similar to analogous correlations in other complex networks.

Evaluating the relative performance and accuracy of IP *geolocation* services, i.e., those for determining geographical positions of Internet identifiers (IP addresses, domain names, AS numbers) remains an open challenge, particularly for core infrastructure, e.g. routers. Commercial IP geolocation tools [7] generally use proprietary methodologies, offer coarse granularity, often disagree with each other in location results, and do not yet support IPv6 geolocation. Challenges to comparison across tools include the lack of standard protocols, data formats, ground truth data, and evalution metrics or methodologies. Yuvall Shavitt presented a formal comparision study of geolocation databases undertaken last year, based on a methodology that groups IP addresses to PoPs by considering structure and delay characteristics [37]. Brad Huffaker presented preliminary results of a study introduced at last year's AIMS workshop [19]. He described a new methodology for analyzing inconsistencies across databases observed for IP address blocks attributed to different geographic regions and organization (Autonomous Systems) types. This database comparison is possible without access to significant ground truth geolocation data. Peter Haga presented Spotter [35], a geolocation service that uses triangulation among active RTT measurements between landmarks to provides city-level IP geolocation [35]. Spotter might also be useful in evaluating geolocation databases.

Residential broadband measurement is a much more complex and challenging task than many researchers anticipated. Sound measurement experiments are arduous to implement in real-life conditions, and results are hard to interpret, since numerous parameters vary so widely across ISPs, users, and home networking equipment. Srikanth Sundaresan and Nick Feamster from Georgia Tech discussed various aspects of their new infrastructure called BISMark, which is trying to measure latency, packet loss, jitter, and throughput from residential broadband connections. Hoping to gather representative data on the impact of home network parameters on residential broadband performance, Renata presented and asked participants to spend 3 minutes running LIP6's HomeNet Profiler [39] when at home.

There is still room for much improvement in interaction and technology transfer among the academic, commercial, and government sectors. In the absence of widescale cooperation from ISPs regarding sharing topology data, cross-correlation across independent results obtained by different measurement methods remains the most reasonable validation approach. AIMS-3 participants repeated the strong consensus articulated at previous workshops [11, 13, 26] – that lack of ground truth data for validation of inferences continues to present the most serious bottleneck for rigorous development of Internet science. Informal contacts between researchers and individual providers produce useful sets of ground truth data, but these datasets are limited in scope, not necessarily representative of larger topology maps, and not generally shareable, rendering it difficult to establish or verify the scientific integrity of published results. The validity of topology ground truth can quickly decrease with time, while many analyses need accurate ground truth data concurrent with executed measurements. Both limitations suggest the need for more scalable approaches to validation, but without more formal cooperation from ISPs, improving validation will require closer collaboration across research and infrastructure projects, exchanging tools, data, and measurement methodologies to allow cross-validation and reproducibility of results. Experiences from sharing other types of operational data, e.g., spam blacklists, have potential underexplored applicability to other research data sharing issues.

Edward Rhyne, from DHS S&T, presented a slideset introducing participants to the division's Cyber Security Program Areas. He highlighted Internet mapping and measurement as well as attack modeling as one of the focus areas in the recent BAA [15].

Richard Barnes of BNN Technologies introduced participants to the recently established Measurement, Analysis, and Tools Working Group hosted by the RIPE-NCC. The working group hopes to encourage discussion around measurement issues of interest to the operational community as well as methods.

The international research network community, especially the NSF-funded IRNC networks, provide a valuable potential resource to researchers if NSF can incent these networks to support sharing of network data and measurement-based collaborations as NSF did in the 1980s and 1990s with NSFNET and vBNS. Harika Tandra of the NSF-funded IRNC network infrastructure GLORIAD project presented a new software prototype, DvNoc[18], for helping to manage this network of global links, visualizing the traffic flows that transit the project links, and eventually displaying performance, topology, routing, and account information. The tool also provides an interface for user support and a communications platform for network operations staff.

2.1 Infrastructure updates

One recurring theme of AIMS is for academics, operators, and funding agents to discuss what types of measurement infrastructures are needed, for what purposes, and if there are more effective ways of supporting them. Each AIMS workshop has had a session of talks by operators of active measurement infrastructure projects to provide status updates regarding scope, functionality, and activities of supported platforms, including discussion of ways to fund and cross-fertilize across infrastructure projects.

DIMES lack of support for the new Windows 7 (until recently) stopped its growth at roughly 1000 daily agents producing 4-5 million measurements a day. A new agent which was just released supports all recent Windows OS, as well as various Linux distributions and Mac.

iPlane [1, 29, 30] is a system developed by researchers at UW for Internet-wide performance prediction. iPlane continually measures from hundreds of PlanetLab and public traceroute server vantage points to more than 90% of prefixes at the Internet's edge and then uses the data to inform performance predictions for unmeasured paths.

Archipelago (Ark) [14], CAIDA's active measurement platform, provides infrastructure for the collection of Internet topology and performance data as well as a virtual laboratory for experimentation. The laboratory includes a layer of "middleware" – specifically a coordination and communication facility to support macroscopic distributed asynchronous Internet measurements, shielding researchers from the complexities of managing measurement infrastructure. Ark uses a tuple space implementation to coordinate and buffer asynchronous measurements, which has been useful to projects such as spoofer which may have outages of the backend, in which case Ark queues results for later download. As of April 2011, Ark is composed of 54 nodes capable of flexible probing of IPv4 address space, 16 with IPv6 capability. Ark now supports IPv6 Spoofer measurements described at last year's AIMS workshop [5]. Several of the Ark nodes now implement RADclock [34] which offers microsecond-level time synchronization.

Mikhail Strizhov from Colorado State presented the latest status of their real-time BGP data collection infrastructure, including how to access the data (a publish-subscribe model), formats used (XML, etc.), and potential uses of the data. The software is extensible to easily add new peers and connect feeds together. The data may be particularly interesting to a growing number of active measurement projects, as now one can request BGP data in real-time, simultaneous with active measurements and/or BGP changes could trigger actions by active measurement projects. Ernest McCracken presented related work on NetViews, which geographically visualizes the AS path dynamics reflected in real-time BGP updates.

Robert Kisteleki of RIPE NCC presented their new active measurement network launched in November 2010 – RIPE Atlas. RIPE NCC's goal is to scale up to thousands, potentially tens of thousands, of vantage points and execute simple built-in and user specified measurements. The Atlas nodes are keychain-sized, USB-connected (and powered) and thus trivial to deploy, and perform a small set of measurements (ping and traceroute) to a provided list of targets, and limited DNS queries mostly to root nameservers. Building on many years of experience and lessons learned from RIPE's TTM project, RIPE NCC has developed a system to incent probe hosting by allowing hosting sites to accumulate credits that they can use to request measurements from the system. An explicit non-goal of the Atlas project at this time is any performance evaluation. As of April 2010, RIPE NCC had deployed 350 Atlas nodes.

As envisioned in the CONMI workshop report [12], several infrastructure operators are pursuing approaches to facilitate data exchange and dissemination, including across administrative domains, such as with PerfSonar [2]. Brian Tierney presented on ESnet's expanding PerfSonar deployment and future plans.

TopHat [3] is a measurement platform recently created by UPMC that interconnects to systems of collaborators. The objective is to enable user access to an expanded set of measurement functionality and data via a unified consistent and familiar interface. TopHat currently focuses on serving live data, and allows user-specified querying across some interconnected measurement systems. One application is the support of measurement experiments on PlanetLab Europe thanks to the MySlice utility [33]. TopHat also interconnects with the DIMES platform [36] offering an interface to retrieve measurements from DIMES agents around the world. The current ad-hoc interconnection framework consists of a set of interfaces allowing the description of resources and their measurement capabilities and the exchange of measurement data, complemented by tools for authentication, authorization, and accounting. The next step of the project is to harmonize the interconnection API around a proposed standard, to facilitate interoperability.

Partha Kanuparthy (with Constantine Dovrolis) from Georgia Tech presented the ShaperProbe [21] service hosted on M-Lab which has been running on M-Lab for 2 years [22], and has been run by 100,000+ users. A high-level picture of their data shows traffic shaping signatures implemented in several ISPs. The packet traces are publicly available through M-Lab. They have also presented preliminary efforts on wireless diagnosis inside the home (with Intel Research and CMU), and performance diagnosis in the wide area (Pythia, in collaboration with ESnet and Internet2).

3. PERFORMANCE MEASUREMENTS

We expanded the agenda of AIMS this year to include more talks on performance evaluation in multiple contexts, including residential broadband measurements of recent interest to the U.S. FCC. Many talks were punctuated by discussions over what were appropriate ways to conduct broadband measurements. By the end of the workshop it became clear that one could have at least one workshop entirely dedicated to residential broadband measurements, which are increasing in importance to consumers, government, and industry, but difficult (or impossible) to validate without ground truth data on real networks. Controlled environments are essential for analyzing aspects of traffic measurement in isolation, but without validation against real world background traffic, tools for bandwidth measurement are not likely to be robust in many diverse environments.

Renata's Teixeira and one of her students at UPMC Sorbonne University presented their work on measuring residential broadband performance. First, Oana Goga showed that some tools using default settings underestimate available bandwidth links by more than 50%. She used controlled experiments to demonstrate that the root cause of such underestimation was the limited forwarding rate of some current home gateways. This validation is non-trivial since even employees of gateway hardware vendors cannot get access to configuration parameters and software running the gateways. Renata described and solicited responses to a survey of home network and system performance information annotated with feedback from end-users [38].

Rocky Chang of The Hong Kong Polytechnic University introduced OneProbe [28], a tool and underlying infrastructure that has measured RTT, one-way packet loss, reordering, and capacity on the Hong Kong academic network for the last two years.

Steve Bauer of MIT presented a large-scale re-evaluation of Explicit Congestion Notification (ECN) network and server support, and a new methodology to measure the prevalence of client-side ECN capability [4] – measurements that are particularly relevant given the current widespread implementation of ECN within modern operating systems and proposals to use ECN for not only congestion management, but policy and billing. Initial research finds a significant number of routers incorrectly set or clear the ECN bits in the IP header impeding ECN deployment and use.

David Choffnes of U. Washington discussed a new approach to diagnosing a performance problem using some of the reverse-traceroute technology that his collaborator Ethan Katz-Bassett presented last year. Fabian Bustamante's described his Dasu platform extension to BitTorrent Vuze, which contiuously captures available bandwidth information for BitTorrent connections from over 1,000 ASes and growing. Atef Abdelkefi presented a Principal Components Analysis (PCA) approach to decomposing a delay time series into a smooth periodic (normal) trend and an impulsive sparse burst, and how the performance of the approach is degraded due to the perturbation phenomenon.

4. IPV6 MEASUREMENTS

We had a session of IPv6 measurement talks given by kc claffy, Emile Aben, and Nicholas Weaver. They presented what CAIDA, RIPE, and ICSI, respectively, have learned from previous and ongoing IPv6 measurement activities. Emile described a tool that others can install on web sites to add to the global IPv6 knowledge base, and other IPv6 data analyses RIPE NCC has contributed. kc reviewed four analyses of IPv6 deployment based on four independent data sources: address allocation, active topology probing, DITL DNS traffic to root servers, and two web-based survey of address holders from 2008 and 2009. Nick Weaver described Netalyzer's IPv6 measurement suite, including JavaScript-based IPv4 vs IPv6 vs dual-stacked latency, traceroute, path MTU discovery, network interface enumeration, DNS support for IPv6, and detection of some common IPv6 related failures. Similar to the broadband measurement topic, there was sufficient interest in IPv6 measurements to justify its own workshop, or at least a full day of talks at a future AIMS workshop.

5. RESULTING COLLABORATIONS

Collaborations resulting from the AIMS workshops have facilitated the use of infrastructures and data by other research groups, permitted evaluation of existing and integration of new advanced measurement techniques into operational measurement systems, and enabled exchange of data for cross-validation of various studies. A more specific list includes:

1. Benoit Donnet and collaborators he met at AIMS published a paper (to appear in NGI2011) on cross-validation of mrinfo data using CAIDA's Ark infrastructure [31].
2. CAIDA researchers augmented their 2010 and 2011 IPv4 topology measurement experiments using data from Benoit Donnet's experiments.
3. CAIDA executed a special experiment on Ark following David Choffnes' talk and using Northwestern AquaLab's data. CAIDA also integrated these IP addresses into future versions of Ark's IPv4 topology probing.
4. Rob Beverly presented work that contained a section suggested by Yuvall at AIMS2010 (the set cover).
5. CAIDA and Rob Beverly proposed a (pending) NSF project that will apply Rob's directed probing technique to IPv6 measurements.
6. Young Hyun discussed with Rob how to change Ark on-demand measurement functionality to support his and Steve Bauer's ECN measurements.

7. Rob Beverly and Steve Bauer of MIT pursued a study based on their presented ECN-detection methodology.
8. RIPE NCC continues to import CAIDA data sets (dnsnames) into INRDB.
9. Emile Aben at RIPE and Fabian Bustamante are cooperating on IPv6 adoption measurements.
10. Young Hyun talked with Robert Kisteleki and Emile Aben regarding the use of Atlas probe data to enhance CAIDA's IPv4 topology measurements with additional vantage points and validation support.
11. Mikhail Strizhov of Colorado State University deployed a RIPE ATLAS node at their lab, the only one in their region. He also established a contact at RIPE to test and evaluate BGPMON.
12. Mikhail received several useful new suggested features and directions for BGPMON, and set up a meeting with Matt Zekauskas to discuss Internet2 potentially using BGPMON.
13. CAIDA discussed a follow-on workshop focused on broadband performance measurements later in 2010 with MIT and Georgia Tech, as well as the FCC.
14. Several workshop participants installed and ran Renata Teixeira's hostview tool.
15. Matthew Luckie and Amogh Dhamdhere completed their study to quantify some types of false links in traceroute-based IP topologies [27], much to the relief of Timur Friedman, Renata teixeira, and other users of Paris traceroute.
16. CAIDA and UPMC plan to support Ark on-demand measurements through TopHat's infrastructure interconnection platform.
17. CAIDA is investigating the possibility of incorporating recent DIMES addresses into our next Internet Topology Data Kit analysis process.
18. RIPE and CAIDA are contributing data to World IPv6 Day.

6. WORKSHOP FEEDBACK

We conducted a survey of AIMS-3 participants to rank areas of interest for future workshops. The highest ranked topics were: hybrid BGP/traceroute infrastructures; large scale topology mapping; broadband performance measurement; and IPv6. One consistent piece of constructive feedback was a desire to shorten the length and/or reduce the number of talks in order to expand the amount of time available for group discussions, as has occurred at previous AIMS-1 and -2. The growing interest from the research community and inevitably expanding list of applications of active Internet measurements suggests that for future AIMS we should consider not only the proposed format modifications, but also possibly spin off related workshops. Another option would be to add half a day to the workshop duration (making it a three day workshop) to ensure sufficient time for discussing in depth big emerging topics such as performance measurements and IPv6. CAIDA found the survey feedback very valuable in planning future workshops this and next year.

ACKNOWLEDGMENTS. The workshop was supported by the National Science Foundation via grants CNS-0958547 and OCI-0963073, and by the U.S. Department of Homeland Security (DHS) Science and Technology (S&T) Directorate via contracts N66001-08-C-2029 and NBCHC070133. We thank all participants for their insights.

7. REFERENCES

[1] Iplane. http://iplane.cs.washington.edu.
[2] Performance focused Service Oriented Network monitoring ARchitecture. http://www.perfsonar.net/.
[3] Jordan Auge, Timur Friedman, and Thomas Bourgeau. TopHat. http://www.top-hat.info/.
[4] Steven Bauer, Robert Beverly, and Arthur Berger. Measuring the current state of ecn support in servers, clients, and routers. 2011. In Submission.
[5] Robert Beverly, Arthur Berger, Young Hyun, and k claffy. Understanding the Efficacy of Deployed Internet Source Address Validation Filtering. In *8th ACM SIGCOMM/USENIX Internet Measurement Conference (IMC)*, November 2009.
[6] Robert Beverly, Arthur Berger, and Geoffrey Xie. Primitives for active internet topology mapping: Toward high-frequency characterization. In *Proceedings of the Tenth ACM SIGCOMM/USENIX Internet Measurement Conference (IMC)*, November 2010.
[7] CAIDA. Netgeo ref. http://www.caida.org/utilities/netgeo/.
[8] CAIDA. ISMA 2010 AIMS-3 - Workshop on Active Internet Measurements, 2011. http://www.caida.org/workshops/isma/1102/.
[9] Kai Chen, David Choffnes, Rahul Potharaju, Yan Chen, Fabian Bustamante, and Yao Zhao. Where the Sidewalk Ends: Extending the Internet AS Graph Using Traceroutes From P2P Users. In *Proc. of ACM CoNEXT*, Rome, Italy, December 2009.
[10] David Choffnes and Fabian Bustamante. Pitfalls for testbed evaluations of Internet systems. *SIGCOMM Computing Communications Review*, 40(2), April 2010.
[11] KC Claffy, Emile Aben, Jordan Aug, Beverly, Fabian Bustamante, Benoit Donnet, Timur Friedman, Marina Fomenkov, Peter Haga, Matthew Luckie, and Yuvall Shavitt. The Workshop on Active Internet Measurements (AIMS) Report. *Computer Communication Review*, 40(5), October 2010.
[12] KC Claffy, Mark Crovella, Timur Friedman, Colleen Shannon, and Neil Spring. Community-Oriented Network Measurement Infrastructure (CONMI) Workshop Report. *ACM SIGCOMM Computer Communication Review*, 36(2):41–48, 2006.
[13] KC Claffy, M. Fomenkov, E. Katz-Bassett, R. Beverly, B.Cox, and M. Luckie. The Workshop on Active Internet Measurements (AIMS) Report. *Computer Communication Review*, 39(5), 2009.
[14] KC Claffy, Young Hyun, Ken Keys, Marina Fomenkov, and Dmitri Krioukov. Internet mapping: from art to science. In *DHS Cybersecurity Applications and Technologies Conference for Homeland Security (CATCH)*, 2009.
[15] DHS. 2011 DHS Science and Technology CSRD BAA: Solicitation Number: BAA 11-02. https://www.fbo.gov/index?s=opportunity&mode=form&id=3c71c829bc28fcea61aef3a5e0f58ffe.
[16] Cooperative Association for Internet Data Analysis. Internet Topology Data Kit, April 2010. http://www.caida.org/data/active/internet-topology-data-kit/.
[17] Simson Garfinkel and Lorrie Faith Cranor. Institutional review boards and your research: A proposal for improving the review procedures for research projects that involve human subjects and their associated identifiable private information. 2010. http://simson.net/clips/academic/2010.CACM.IRB.pdf.
[18] GLORIAD. Global Ring Network for Advanced Applications Development: DvNoc. http://viz.gloriad.org/dvnoc/dvnoc.html.
[19] Bradley Huffaker. Geolocation comparison: Caida's geolocation database comparison, May 2011. http://www.caida.org/publications/presentations/2011/nmmc_geocompare/.
[20] Bradley Huffaker, Amogh Dhamdhere, Marina Fomenkov, and KC Claffy. Toward Topology Dualism: Improving the Accuracy of AS Annotations for Routers. In *PAM*, 2010.
[21] Partha Kanuparthy. DiffProbe Shaper Module: Detecting ISP Traffic Rate-Limiting. http://www.cc.gatech.edu/~partha/diffprobe/shaperprobe.html.

The research community needs to introduce and agree upon standards and best practices to promote a diverse and heterogeneous field of Internet active measurements.

	Recommendation	**Current status**
1	*develop standard APIs for measurement systems, standardize tool output, enable tool sharing on different platforms*	Work in progress; BGPMON and TopHat are achieving milestones toward this recommendation.
2	*publicize the best available data, document them as ground truth, provide comprehensive statistical characterization, make these data easily downloadable*	CAIDA released two new ITDK's [16] with as much ground truth validation as feasible. Rob Beverly published anonymized spoofer data. John Heidemann (not at AIMS) has published his IP hitlists. Wide data distribution continues to be complicated by NDAs and lack of resources/priority, but it would help to catalog all the data-sharing occurring as a part of conferences such as IMC, since the situation is improving somewhat.
3	*design flexible, easily extensible measurement infrastructure platforms capable of running various tools and types of measurements at Internet scale*	Ongoing. Ark, TopHat, DIMES, BGPMON and Dasu have made notable progress toward this goal.
4	*provision for continuity of measurements, dissemination of data, with long-term archiving of data to study historical trends*	This task (still) deserves more attention from government, funding, and policy-making agencies. Longitudinal data collection and analysis continue to elude the research community, again based on resource constraints and relative priorities. A Data Management Plan now required by NSF for all new proposals might be a step in the right direction.
5	*maintain no-probe lists based on requests*	Researchers maintain their own *ad hoc* lists, but there is no mechanism in place to share such lists across research groups. This recommendation has received less community support over the years, as different research questions need different types of measurements involving potentially different populations of hosts, thus a single one-size-fits-all no-probe list would unnecessarily constrain the pool of hosts available for a given project.

Table 1: Recommendations on standards.

[22] Partha Kanuparthy. ShaperProbe. `http://measurementlab.net/measurement-lab-tools#tool5`.

[23] Erin Kenneally and KC Claffy. Dialing privacy and utility: a proposed data-sharing framework to advance Internet research. *IEEE Security and Privacy*, 2010. `http://www.caida.org/publications/papers/2010/dialing_privacy_utility/`.

[24] K. Keys. Ip alias resolution techniques. *Computer Communication Review*, January 2010.

[25] Ken Keys, Young Hyun, Matthew Luckie, and KC Claffy. Internet-Scale IPv4 Alias Resolution with MIDAR: System Architecture, May 2011. `http://www.caida.org/publications/papers/2011/midar-tr/`.

[26] D. Krioukov, F. Chung, KC Claffy, M. Fomenkov, A. Vespignani, and W. Willinger. The Workshop on Internet Topology (WIT) Report. *Computer Communication Review*, 37(1), 2007.

[27] Matthew Luckie, Amogh Dhamdhere, KC Claffy, and David Murrell. Measured impact of crooked traceroute. *CCR*, 41(1), 2011.

[28] X. Luo, E. Chan, and R. Chang. Design and implementation of TCP data probes for reliable and metric-rich network path monitoring. In *Proc. USENIX Annual Tech. Conf.*, 2009.

[29] Harsha V. Madhyastha, Tomas Isdal, Michael Piatek, Colin Dixon, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iPlane: An Information Plane for Distributed Services. In *OSDI*, 2006.

[30] Harsha V. Madhyastha, Ethan Katz-Bassett, Thomas Anderson, Arvind Krishnamurthy, and Arun Venkataramani. iPlane Nano: Path Prediction for Peer-to-Peer Applications. In *NSDI*, 2009.

[31] Pascal Mrindol, Benoit Donnet, Jean-Jacques Pansiot, Matthew Luckie, and Young Hyun. Merlin: Measure the router level of the internet. Technical report. `http://www.caida.org/publications/papers/2010/merlin/`.

[32] P. Matray, P. Haga, S. Laki, I. Csabai and G. Vattay. The Network Geography of the Internet. In *INFOCOM*, 2011.

[33] PlanetLab. MySlice: A Slice Management Platform for PlanetLab. `https://www.myslice.info/`.

[34] Julien Ridoux and Darryl Veitch. RADclock. `http://www.cubinlab.ee.unimelb.edu.au/radclock/`.

[35] S. Laki, P. Matray, P. Haga, T. Sebok, I. Csabai and G. Vattay. Spotter: A Model Based Active Geolocation Service. In *INFOCOM*, 2011.

[36] Yuval Shavitt and Eran Shir. Distributed Internet measurements and simulations. `http://www.netdimes.org/`.

[37] Yuval Shavitt and Noa Zilberman. A study of geolocation databases. Technical report. `http://arxiv.org/abs/1005.5674`.

[38] Renata Teixeira, Diana Joumblatt, Jaideep Chandrashekar, and Nina Taft. HostView. `http://cmon.lip6.fr/EMD/Home.html`.

[39] Renata Teixera. HomeNet Profiler. `http://cmon.lip6.fr/hnp/`.

A. Updates to the AIMS-1 Recommendations

The 1st AIMS workshop [13] developed a set of recommendations intended to advance the field of Internet active measurements. While these recommendations represent a multi-year "wish list" of Internet researchers, we offer a review of progress made on them in the last year. Tables 1, 2, and 3 summarize our assessments.

The lack of consistent guidelines for Internet measurement limits the recognized legitimacy and sustainability of Internet measurement systems.

	Recommendation	**Current status**
1	*replace obsolete RFC1262 with a more current document*	Not done, blocked on working group volunteering. No change since AIMS-2.
2	*create a report on ethical guidelines for Internet active measurement research*	A working group associated with the DHS's PREDICT project is finishing a program document attempting to codify ethical guidelines for Internet measurement research. This effort is motivated by the need to advance cybersecurity research while respecting evolving expectations of privacy.
3	*facilitate interaction between Internet researchers and Institutional Review Boards (IRB) that overview and regulate human research activities at individual institutions* [17]	Research groups continue to pursue and several of them have received approval from their IRBs for active measurement. CAIDA's IRB application is posted on http://www.caida.org/home/about/irb/.
4	*identify important research questions/problems in the field of Internet research where macroscopic active measurement can have a positive impact*	Significant progress on this recommendation was made since last year, some described in this report, specifically in broadband performance measurement, IPv6, and efficient macroscopic topology measurements.

Table 2: Recommendations on guidelines.

The research community must increase transparency of Internet measurements and better communicate utility of results to broader communities affected by measurements (legal, political, operators, users).

	Recommendation	**Current status**
1	*create a central easily accessible database of planned or ongoing Internet experiments*	Not done and no progress since AIMS-2. At present, this task can only be done as a volunteer effort, e.g., via a group-editable wiki page. Someone needs incentive (and/or funding) to set it up.
2	*release source code for tools used in publications*	Same as at AIMS-2 - no available survey of how often this happens.
3	*consider other means of communication (i.e., blogs, mailing lists, automated announcements) to keep other communities informed of Internet measurement research experiments*	Same as at AIMS-2 - some research groups do use blogs to announce experiments, as well as continue using operational mailing lists such as NANOG's and RIPE's.
4	*increase visibility and usability of data (including formatting standards), relevance of data to users, and exposure of implications of studies based on data*	Some measurable progress, outlined in several of the workshop talks.
5	*inform debate on clean-slate Internet architecture*	Needs concerted cross-fertilization efforts among communities and funding agencies.
6	*discuss with academics, operators, and funding agencies how many measurement infrastructures are needed, for what purposes, and if there are more effective ways of funding them*	Some progress is made both at and subsequent to AIMS workshop, as documented in report.
7	*enable interaction and technology transfer between three main players in the field of Internet research: academic laboratories, commercial enterprises, and government institutions*	Same as at AIMS-2, still needs policy support, and improved technology transfer methods including data-sharing models. [23]

Table 3: Recommendations on transparency.