

International Research Network Connections: Usage and Value Measurement

kc claffy and Josh Polterock
CAIDA/UCSD
kc,josh@caida.org

ABSTRACT

The NSF's International Research Networks Connections (IRNC) Program¹ supports the research and education community by funding globally interoperable high-bandwidth infrastructure, coordinated experimental services, and human network community building. At the January 2013 IRNC principal investigator (PI) meeting, one area of discussion was measurement capabilities and opportunities that could provide NSF and others with greater insight into what science projects are supported internationally through this program. This note summarizes discussions at the meeting, including each project's current measurement activities, and consensus on several concrete directions that would support greater visibility into IRNC infrastructure operations, usage, and value.

1. INTRODUCTION

The National Science Foundation's International Research Networks (IRNC) Program supports global research and education communities by funding globally interoperable high bandwidth networked infrastructures, coordinated experimental services and human network community building. The IRNC program funds two types of projects: production network infrastructure (ProNet) projects, which fund high bandwidth links and associated operational as well as experimental support for data-intensive scientific collaborations; and special projects (SPs), which include measurement activities and experimental functionality such as dynamic circuit technology. The IRNC also funds a special project (NSRC) to cultivate the international research and education network fabric via teaching, training, and technology transfer. The January 2013 Principal Investigators' meeting dedicated a session to discussing existing measurement capabilities of ProNet sites, other measurement activities funded by the IRNC program via special projects, and potential opportunities for collaboration in pursuit of greater empirical visibility into the IRNC infrastructure's operations, usage, and value.

2. PRONET PROJECT MEASUREMENT CAPABILITIES

During the measurement session of the meeting, program manager Kevin Thompson asked each project to summarize existing baseline measurement capabilities across IRNC connections, explain how they are used, and describe possi-

¹<http://www.irnclinks.net/>

ble target measurement capabilities and services beyond the current phase of the IRNC program. This section summarizes each ProNet project's response to this request.

2.1 GLORIAD

Greg Cole presented a new application to support the GLORIAD ProNet project², *GloTop*, which he wrote using the formerly proprietary (recently open-sourced) Live-Code platform from RunRev. This application pulls from a database populated with data from Argus traffic flow measurement software running on commodity Dell servers using a span port (with plans to move to taps). The *GloTop* application displays top talkers, per-protocol and per-application utilization, and a breakdown by science discipline (using local tagging and MaxMind to geolocate IP addresses). Greg's goals in designing this system were to: (1) understand the network requirements of Gloriad's base; (2) identify poor performance of individual applications by continually analyzing per-flow metrics as load, packet loss, jitter and routing asymmetries; (3) mitigate poor performance of applications by identifying fabric weaknesses; and (4) visualize the enormous volume of data.

GLORIAD maintains two primary databases. The first contains flow records of IP addresses, AS numbers, and domains with support tables for mapping IP addresses, countries, scientific disciplines, protocols, services, etc. The second database contains summary tables to enable fast search and retrieval of flow information. Project developers are experimenting with the Argus tools for query and live annotation of flow updates. GLORIAD will provide shell accounts and SQL access to this system, so long as intended use focuses on improving service to customers.

2.2 ACE and TransPac3

John Hicks described operational measurement activities of the ACE and TransPac3 ProNet projects, the focal point of which comes through the Global Research Network Operations Center³. On both TransPac and ACE links, the primary management tools are based on perfSONAR, including perfSONAR Lookup, SNMP MA, Buoy (iperf, bwctl, and OWAMP tester with wrapper for perfSONAR). They also monitor packet loss, conduct netflow (jflow/sflow) data collection and analysis, collect router state via JunOS and publish data via perfSONAR.

They continue to try to keep the two independent perfSONAR implementations compatible: a Perl-based one and

²<http://gloriad.org>.

³<http://globalnoc.iu.edu/atlas.html>

a Java-based one developed in Europe. The GlobalNOC moved from the Java to the Perl implementation for performance and reliability reasons, and to better match system administrator expectations and skill sets. European sites continue to make use of the Java implementation.

On the TransPAC2 link from Tokyo to Los Angeles, they collect SNMP data using *SNAPP* (a high-speed SNMP collector from IU/grNOC) and publish it via perfSONAR⁴. On the ACE link they collect SNMP only, because the circuits plug into existing MANWAN/WIX exchange points. The GRNOC service desk uses Nagios and perfSONAR to determine the status of network devices and the health of the network. The engineering staff at the NOC uses SNMP and OWAMP data to troubleshoot network problems. Through the grNOC portal, researchers can see utilization graphs to determine available bandwidth⁵. The NOC publishes router temperature, CPU, and interface data via perfSONAR, but does not publish IP addresses from netflow/jflow data. They also publish their BGP view available via Route Views⁶. They have assisted CAIDA with contacts in the Asia Pacific region for potential hosting sites for the Ark measurement infrastructure.

2.3 AmLight

AmLight deployed perfSONAR at AMPATH in their Miami node as well as the RedCLARA, Rede Nacional de Ensino e Pesquisa (RNP), and other NRENs in South America. They are observing performance between CTIO and NCSA using perfSONAR in collaboration with the IRIS project and REUNA. AMPATH peers with Route Views and participates in the ROVER testbed.

Julio Ibarra and his team have worked closely with CAIDA this year to discuss, design, purchase, install and configure a passive monitor system running CAIDA's Coralreef software suite and report generator. Once complete, the system will monitor and report on bidirectional traffic flowing over the 10GE AmLight link⁷ between the AMPATH Exchange point at Florida International University in Miami, Florida and the SouthernLight Exchange Point in Sao Paulo (ANSP, RNP, CLARA).⁸ With assistance from CAIDA, AmLight has created custom reporting and has plans to include a new custom report to show top AS flows. Additionally, AmLight makes use of Nagios for alerts and collects and analyzes sflow data. AmLight's measurement data, to date, has been used primarily by collaborators and connectors for operations, troubleshooting, and observation of connection performance.

Working with Professor Luis Lopez, the Director of Academic Network for State of Sao Paulo (ANSP), AmLight assisted CAIDA with deployment of an Ark monitor using the new Raspberry Pi platform.

2.4 Translight/Pacific Wave

Alan Whinery and Celeste Anderson presented updates for the Translight/Pacific Wave Project. As an Internet exchange, TransLight/Pacific Wave⁹ does not control the

⁴<http://dc-snmp.wcc.grnoc.iu.edu/transpac2/>

⁵<http://globalnoc.iu.edu/snapp.html>

⁶<http://www.routeviews.org/>

⁷<http://ampath.net>

⁸The report generator can be found on the web at http://coralreef.ampath.net/cgi-bin/display_report.

⁹<http://www.hawaii.edu/tlpw/home>

endpoints, so would require permission from individual Pacific Wave participants to publish traffic data. However, the endpoints could publish the data toward potential IRNC goals to track top users and usages of IRNC resources. Currently, they publish maps and "speedtest" style reports on inbound/outbound traffic¹⁰.

PerfSONAR deployment around Pacific Wave is the collaborative work of a community. The IRNC grant funded some equipment for Hawaii, Seattle, and Los Angeles, but the large part of operational perfSONAR work is done by staff at various Pacific Wave community institutions, and funded independently. The perfSONAR discussion group "PerfClub" has grown out of the application of the missions of Translight/Pacific Wave to use perfSONAR to measure and test between Pacific Wave peers and resources of interest to each organization. PerfClub holds a monthly conference call, which enjoys participation from individuals from 26 organizations in 5 countries.

David Lassner has assisted CAIDA with contacts for deployment of two Ark nodes in AARnet in Sydney and Perth locations. All four organizations involved in Translight/Pacific Wave (AARNet, University of Hawaii, CENIC, and Pacific Northwest Giga Pop) host Ark nodes. Eleven Pacific Wave participants peer with RouteViews at at least one peering point and several peer at multiple peering points (AARNet, Internet2, APAN, CANARIE, PNWGP, Transpac, CENIC, ESnet, NLR, University of Hawaii, CERNET).

3. OTHER IRNC-FUNDED MEASUREMENT SUPPORTING ACTIVITIES

The IRNC program funds several special projects focused on various types of measurement: performance, topology, routing, and traffic. This section briefly summarizes these projects.

Eric Boyd of Internet2 presented an update on the IRIS project, a software framework to simplify the task of end-to-end network *performance* monitoring and diagnostics. Based on perfSONAR-PS protocols, IRIS facilitates broader deployment of perfSONAR-enabled resources by providing central management of perfSONAR tests, a test point map, network usage portal, and support of the perfSONAR Toolkit. (The recently released perfSONAR-PS Toolkit 3.2.2 includes: improved IPv6 support for all services and tools; Nagios plugins; support for regular traceroute tests; enhanced one-way latency graphs to show jitter; and minor bug fixes. The Toolkit also includes rebuilt updated kernels with web100 support to ensure users receive security updates with NDT and NPAD support.) IRIS also provides a GUI so that a NOC can build meshes of sites/addresses/organizations, and view performance test results, as well as easily locate available perfSONAR test points.

To support topology measurement of the global NREN fabric, the IRNC program has funded five deployments of CAIDA Ark nodes at locations behind IRNC ProNET links, in: 1) Sao Paulo, Brazil, Rede ANSP/Projeto NARA 2) Sydney, Australia, AARNET 3) Perth, Australia, AARNET 4) Serrekunda, GM, QCell and 5) Honolulu, HI, US, University of Hawaii. The IRNC program has also supported Colorado State University's BGPmon/Route Views project with whom several ProNet projects peer to share their BGP-

¹⁰<http://www.hawaii.edu/tlpw/maps>

based view of the Internet.¹¹

The IRNC program has funded software support for ProNet sites interested in using CAIDA's open-source software platform CoralReef to support traffic measurement. CAIDA has focused on improvements to this software suite requested by ProNet PIs, including: added IPv6 support to *crl_flow* and *crl_anf* (backend to the report generator); improved decoding of IPv6 headers; source/destination pairs sorted by volume; several flavors of IPv6 address anonymization; and netflow import to the report generator. In particular, CAIDA has worked closely with ProNet PI Julio Ibarra and his engineer, James Grace to implement Coralreef measurements, the report generator, and customized reporting of top AS Flows.

Finally, the IRNC program funds a unique project at the Network Startup Resource Center (NSRC, based at the University of Oregon) to cultivate collaborations between U.S. scientists and collaborators in developing countries, via technical exchange, engineering assistance, training, equipment, and educational materials. Through their many contacts operating network infrastructure around the world, NSRC has found several new hosting sites for CAIDA Ark monitors to expand visibility into the global Internet (including R&E) infrastructure around the world.

4. APPROACHES TO INCREASING TRANSPARENCY OF IRNC USAGE AND VALUE

Compared to the fine-grained usage accounting of NSF-funded high-performance computing assets, which allow funding agencies to understand resource usage by discipline, a range of privacy and technical issues render visibility into government-funded network resources more problematic. All PIs considered sharing data among ProNET sites to be a worthy activity, however, all expressed concerns over legal, social, and policy issues. The remainder of this session was dedicated to seeking consensus on how additional collaborations and resources could improve not only IRNC infrastructure performance, but also transparency of the IRNC infrastructure's performance, function, and utility. The priorities arrived at by the group were:

1. The perfSONAR stacks currently present the most viable common fabric of resources for global troubleshooting. However, there exist two code-base implementations using the name perfSONAR. The one promoted for deployment in IRNC grants, perfSONAR PS, is primarily written in Perl, and was developed through collaboration between ESNNet, Internet2, and other organizations. In Europe, research networks use perfSONAR MDM, written primarily in Java. Interoperability between perfSONAR-PS and perfSONAR-MDM does not currently support the goals of promoting measurement, testing and troubleshooting across administrative boundaries. Some discussion between Internet2 and GEANT has taken place.
2. The current unsupported state of web100 and NDT is a huge impending problem ("pothole") for the com-

munity. Specifically, NDT is dependent upon web100, which has now been deprecated and Web10G development is only now starting. The community also needs a secure version of NDT+NPAD functionality.

3. A single map of all IRNC-funded network links, ideally with SNMP-based per-node or per-link traffic, would provide high-level visibility into usage of these links, and is technically possible today assuming each IRNC site could provide its SNMP traffic data feeds to a central location.
4. Extending the current IRNC home page¹² to a community (user-editable) wiki that could include the aggregated map described above as well as a single location that consolidates links to the heterogeneous measurement and reporting efforts such as the prototype at <http://wiki.caida.org/irnc/>.
5. Finer-grained analysis of usage, e.g., on boundaries of science projects, requires a common database of mappings of IP addresses to relevant aggregation granularities, e.g., specific projects or laboratories. Standard nomenclature and annotation conventions would allow projects to consistently describe NSF-funded links and demarcation points. Greg Cole's database which he uses for this purpose already contains IP geolocation information, as well as hostnames mapped to laboratories and other attributes. Standardizing community use of this database would be the easiest.
6. Most ProNet project PIs said they were capable of providing several types of traffic-based usage reporting per link: top applications by volume; highest bandwidth flows; traffic matrix (by specific source/destination prefixes or ASes); and loss indications such as packet drops or retransmits. However, PIs expressed continued concerns regarding sharing traffic data not governed by current MOUs with partners. Modifying the MOUs as part of a future IRNC phase agreement would be required to move forward with publishing these types of aggregated reports. One option would be to contract with a central entity to handle advanced measurement capabilities (beyond operational monitoring such as SNMP-based) across all NSF-funded international connections, assuming NSF can effect the required changes to any legal agreements with IRNC awardees and their peering partners.
7. All IRNC ProNet sites have layer three peering with Route Views, but could further promote Route Views peering to downstream partners and encourage Internet exchange facilities supporting NRENs to operate Route Views collectors and Ark topology measurement nodes. Many PIs also considered it a potential best practice to register routing policy, or at least Route Origin Authorizations (ROAs), in an Internet Routing Registry Registering ROAs would enable implementation of BGP security functionality when it becomes available.
8. The increasing support for Software Defined Networks (SDN) and OpenFlow networks by the IRNC ProNets

¹¹The BGPMON project has delved deeper into routing security with BGP Rover, part of secure64.com, which allows sites to implement routing alerts for potential route hijacking events. Dan Massey offered to help any interested IRNC site register their routes in BGP Rover.

¹²<http://www.irnclinks.net/>

has sparked interest in reporting statistics for these networks. For performance measurement and metrics such as flow set up time, number of active flows, and flow arrival rate that do not appear in legacy network metrics, Stanford University has examples online¹³¹⁴. However, for traffic measurement there exists no clear understanding of what that might mean and will require development of new measurement tools and metrics to address the questions posed by PM Kevin Thompson.

9. As a follow on from the Security at the Cyber Border Workshop¹⁵ held last year in Indianapolis, the idea of supporting passive DNS aggregators came up again as a least common denominator activity that could benefit security without risk to end user privacy.

ACKNOWLEDGMENTS. This report was supported by the National Science Foundation grant OCI-0963073.

¹³<https://openflow.stanford.edu/display/SDEP/Management>

¹⁴<http://yuba.stanford.edu/foswiki/bin/view/OpenFlow/Deployment/Measurement>

¹⁵<https://scholarworks.iu.edu/dspace/handle/2022/14070>