

NETSAGE: OPEN PRIVACY-AWARE NETWORK MEASUREMENT, ANALYSIS, AND VISUALIZATION SERVICE

Alberto Gonzalez¹, Jason Leigh¹, Sean Peisert², Brian Tierney², Andrew Lee³, and Jennifer M. Schopf³

1: University of Hawaii Mānoa, emails agon@hawaii.edu, leighj@hawaii.edu

2: Lawrence Berkley National Lab, emails speisert@ucdavis.edu; bltierney@es.net

3: Indiana University: emails leea@iu.edu, jmschopf@iu.edu

Paper type

Technical Paper

Abstract

NetSage is a project to develop a unified open, privacy-aware network measurement, and visualization service to address the needs of today's international networks. Modern science is increasingly data-driven and collaborative in nature, producing petabytes of data that can be shared by tens to thousands of scientists all over the world. The National Science Foundation-supported International Research Network Connection (IRNC) links, have been essential to performing these science experiments. Recent deployment of Science DMZs [Dart, E. et al., 2013], both in the US and other countries, is starting to raise expectations for data throughput performance for wide-area data transfers. New capabilities to measure and analyze the capacity of international wide-area networks are essential to ensure end-users are able to take full advantage of such infrastructure.

NetSage will provide the network engineering community, both US domestic and international, with a suite of tools and services to more deeply understand: 1) the current traffic patterns across IRNC links, and anticipate growth trends for capacity-planning purposes; 2) the main sources and sinks of large, elephant flows to know where to focus outreach and training opportunities; and 3) the cause of packet losses in the links and how they impact end-to-end performance.

Keywords

Network Measurement, Infrastructure, International, Tools, Visualization, Analytics, National Science Foundation, International Research Networks Connection.

1. Introduction

NetSage is being designed and developed to support a network measurement service for use with the NSF International Research Network Connection (IRNC)-funded backbone and exchange point services, as depicted in Figure 1. NetSage will provide a unified view of the traffic on the links, and assist stakeholders in identifying congestion and bottlenecks. The data collection and visualization service are being developed to directly to respond to explicit questions that have been provided by the user community.

NetSage is focused on building and deploying advanced measurement services that will benefit science and engineering communities in a number of ways, focusing on:

- Better understanding of current traffic patterns across IRNC links, and the ability to better understand growth trends for capacity-planning purposes;
- Better understanding of the main sources and sinks of large, elephant flows, to know where to focus attention on outreach and training; and
- Better understanding of where packet loss is occurring, whether or not the loss is caused by congestion or other issues, and the impact of this on end-to-end performance.

NetSage services will provide an unprecedented combination of passive measurements, including SNMP data [Case, J. et al., 1990], flow data [Claise, B., 2004], and tstat-based traffic header analysis [Marco M., et al., 2005], as well as

active measurements, mainly perfSONAR [Tierney, Brian et al., 2009] , to create longitudinal network performance data visualizations.

In section 2 of the paper we describe our intended NetSage users and use cases, in section 3 we give a brief overview of the NetSage monitoring architecture, in the following sections: 4, 5 and 6 we describe more in depth the different parts of the architecture. In section 4 we describe NetSage data sources, in section 5 we describe the NetSage archive and in section 6 we describe the NetSage visualization approach.

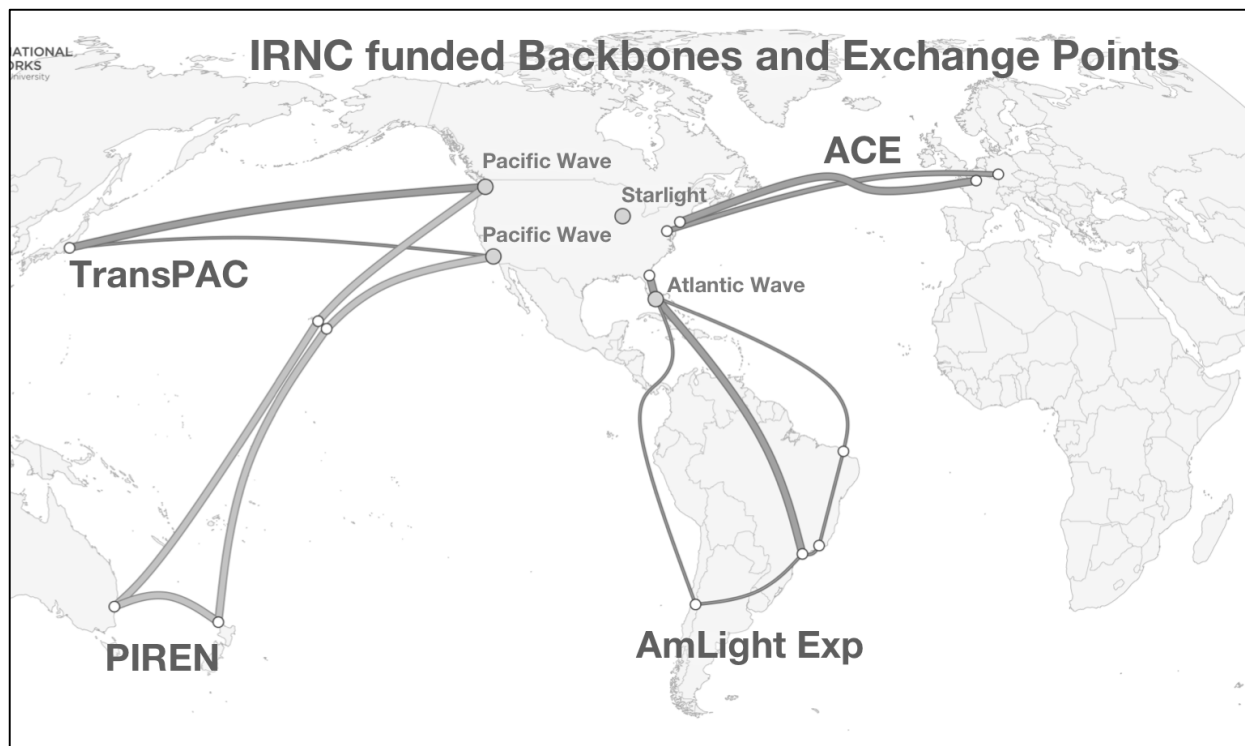


Figure 1: The current NSF IRNC-funded backbones and exchange points.

2. NetSage Users and Use Cases

For monitoring to be effective, it must begin with a clear understanding of the intended audience and the types of insight needed by that audience. The NetSage project defines four sets of end users for their data and visualization services:

1. **Project oversight managers.** These include National Science Foundation program managers and other higher-level users, and they may be concerned with issues such as showing that the traffic on the links are exhibiting broad societal relevance
2. **The IRNC Network Operation Center (NOC) and other network operators.** Network operators may use the NetSage data to understand performance issues or identify developing problems.
3. **Project planners for the IRNC-funded backbone networks and exchange points.** Owners of the backbones and exchange points may want to use NetSage data to do capacity planning or to understand points of contention. The IRNC-funded backbones include: TransPAC4 [TransPAC4, 2015], AmLight Exp [AmLight, 2015], PIREN (Pacific Islands Research and Education Networks)[PIREN, 2015], ACE (America Connects to Europe) [ACE, 2010]; and the network exchange points include: AtlanticWave [AtlanticWave, 2006], StarLight [StarLight, 2015], and Pacific Wave [Pacific Wave, 2004], as shown in Figure 1.
4. **Application engagement staff.** Several projects have funded application engagement staff in order to assist end users in getting better performance over the funded circuits or exchange points. NetSage data may be able to be used to identify some of these classes of users.

Collectively the insight requested by these stakeholders include: the desire to better understand the current traffic patterns across IRNC links, and to understand growth trends for capacity planning purposes; the desire to understand the location and cause of packet loss and its impact on end-to-end performance; and the desire to better understand the main sources and sinks of “elephant” flows to focus end-user outreach and training efforts to enable them to best leverage the network resources. These three categories of inquiries and example questions are presented in Table 1.

With a clear understanding of the audience and their inquiries, it is now possible to focus data gathering efforts. Table 1 also summarizes the tools that NetSage will use to gather the types of data needed in order to answer the specific questions of our target audience.

Goals	Insight Desired	Data Required
Understand traffic patterns for capacity planning	<ul style="list-style-type: none"> • What is the max, min, average bandwidth used between links? • What is the duration and are there any periodic patterns or peak periods? • Which exchange points or networks are congested? 	<ul style="list-style-type: none"> • Throughput data (SNMP) • Packet Headers (Flow Data)
Understand the impact, location, and cause of packet loss	<ul style="list-style-type: none"> • Is loss due to the client, the link, or the exchange points? • How much does packet loss diminish the overall throughput of a flow? • Do losses tend to be due to flows exceeding capacity or network infrastructure problems? • How long do packet loss problems persist before detection, source identification, and resolution? 	<ul style="list-style-type: none"> • Packet loss (tstat) • Packet Headers (tstat) • Bandwidth (perfSONAR) • Latency (perfSONAR)
Understand the nature of elephant flows that use the links	<ul style="list-style-type: none"> • What is the max, min, average duration of elephant flows? • How many elephant flows tend to occur at the same time? • Which exchange points or networks service the most elephant flows? • Are the flows bursty or experiencing sustained traffic? 	<ul style="list-style-type: none"> • Packet Headers (Flow Data) • Packet Headers (tstat) • Packet loss (tstat)

Table 1: Directed questions being used to guide NetSage development.

3. NetSage Monitoring Architecture

NetSage has defined a basic 3-layer architecture, shown in Figure 2, similar to many other monitoring approaches. The first level consists of various data sources. The testpoints that we deploy are capable of running at 10Gbps or higher, support both IPv4 and IPv6, support layer 2 circuit technologies, and do not impact the production traffic. These data sources will combine multiple, different active and passive measurements, as detailed in Section 4.

The second level shows the NetSage archive service. We are using the Time Series Data System (TSTS) to implement a shared archive with the IRNC NOC, as detailed in Section 5. This service provides a standard interface to upload data from the various data sources, as well as a standard interface to the higher level services.

The third level is a suite of data analysis and visualization tools that address the questions outlined by NetSage end users in Table 1. This is detailed in Section 6.

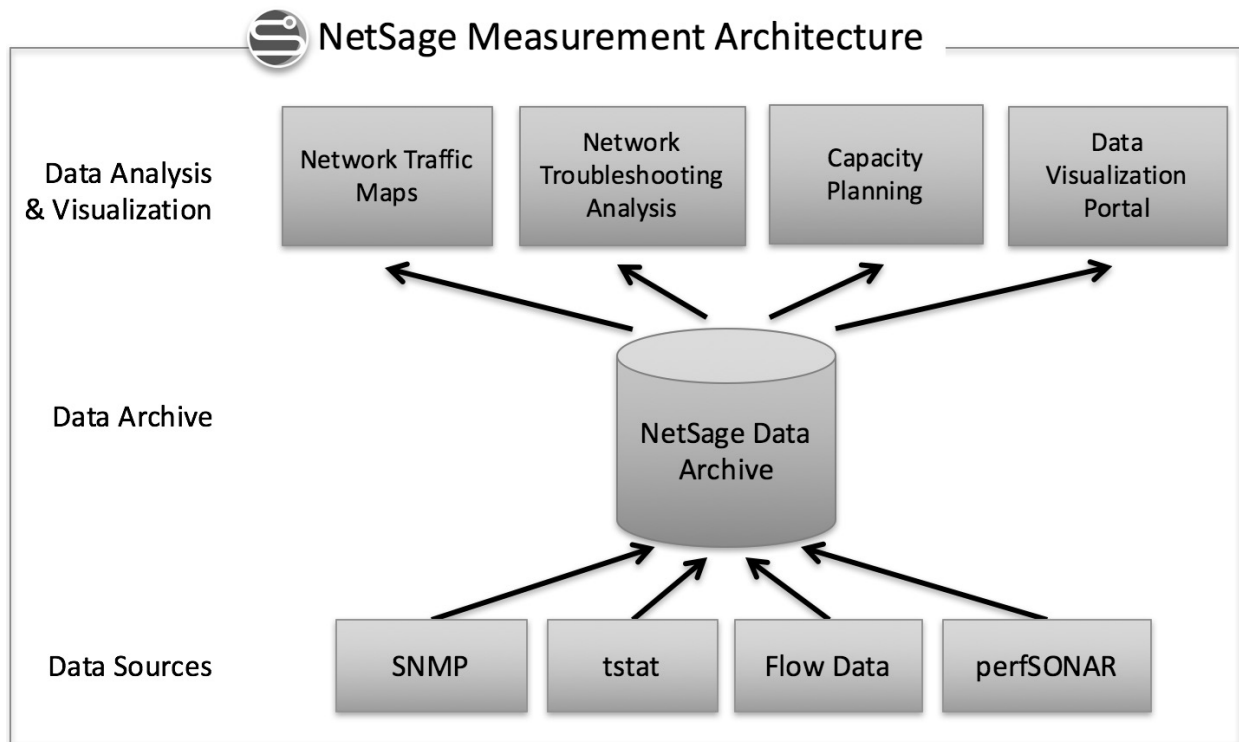


Figure 2: The current NetSage monitoring architecture features a set of data sources, a data archive, and a suite of analysis and visualization tools.

4. NetSage Data Sources

The first layer of the NetSage architecture consists of the different data sources that are being used by the project. These are a combination of passive measurements, such as SNMP, flow data, and data from packet header inspection, and active measurements, such as perfSONAR. Note that for “sensitive” data, such as flow data, NetSage anonymizes the data on site before forwarding it to the central database.

The Simple Network Management Protocol (SNMP) is an application-layer protocol defined in RFC1157 [Case, J., et al., 1990] for collecting and organizing information about managed devices on IP networks. SNMP is commonly used by routers and switches to monitor networks for conditions that warrant administrative attention. This data is commonly collected and openly archived by most R&E networks. We began by collecting SNMP data since each of the backbones and exchange points were already collecting this data and archiving the data sets in public archives. Figure 3 is an example of SNMP data as shown by the open source tool SNAPP [SNAPP, 2016], maintained by the GlobalNOC.

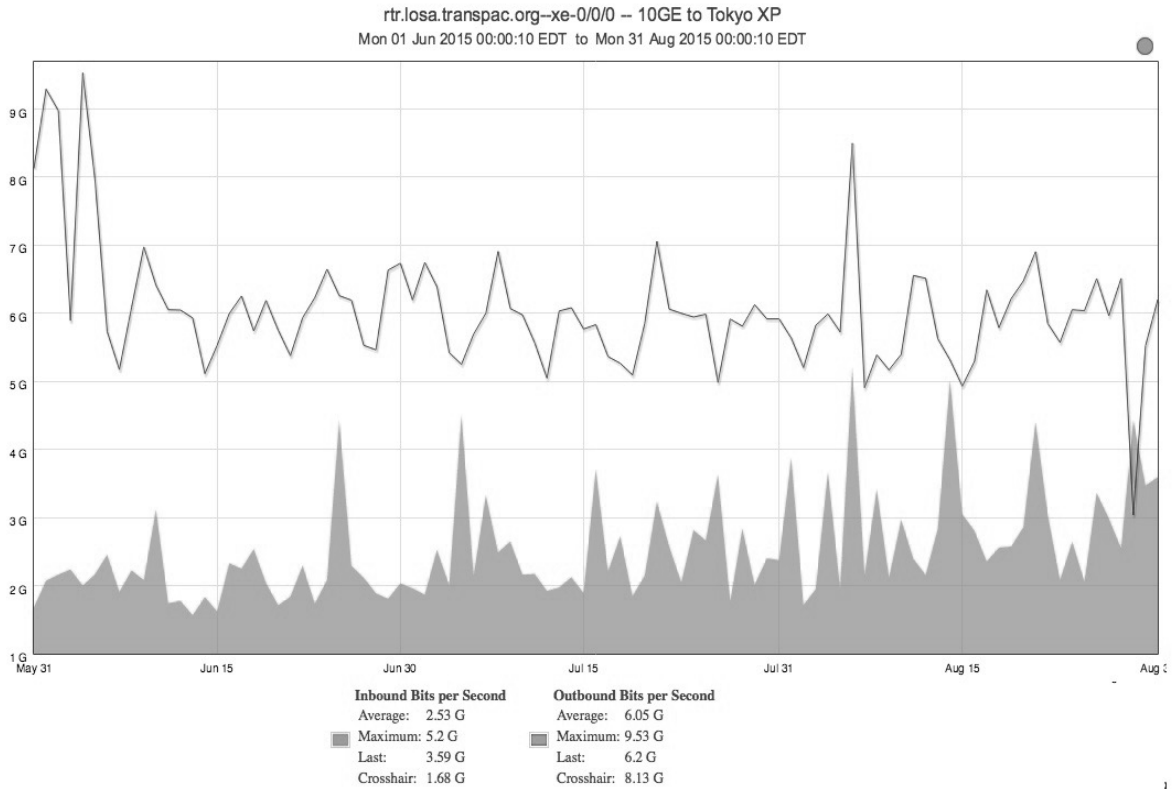


Figure 3: SNAPP graph showing SNMP aggregate traffic data (maximum daily values) for the 10 Gbps TransPAC circuit between Los Angeles and Tokyo.

A second set of passive data can be acquired by means of packet header inspection tools that examine the headers of a network flow and pulls out valuable data from them, without touching the payload of the message. Initially, the NetSage project studied the possibility of using Bro [Paxson, V., 1991] for this purpose. However, after analyzing its performance in a test lab, we found that the Bro TCP analyzer was very CPU intensive, and that large numbers of packets were dropped when even moderate numbers of flows were analyzed.

Instead, we are currently deploying tstat [Marco M., et al., 2005][tstat, 2008] for our packet header inspection tool. Tstat is part of the EUMeasurement Plane (mplane) FP7 project developed by Munafó and Mellia at Politecnico di Torino, and can be used to analyze either real-time or captured packet traces. It rebuilds each TCP connection by looking at the TCP header in the forward and reverse direction. Tstat reports a number of useful TCP statistics, including congestion window size and number of packets retransmitted, which can be used to analyze the health and performance of the link. After running lab tests, we have deployed tstat on one backbone link to further test its scalability and setup, and in Year 2 of the project will be extending this data source.

Another source of passive data for networks is related to flow data collection. Depending on the hardware, this might be NetFlow [Claise, B., 2004], sFlow [Phaal, P. et al., 2001], or IPFIX [Claise, B. et al., 2013]. Flow data will allow us to be able to answer several of the questions desired by our end users: which science domains are using the networks, what do elephant flows look like, etc. However, due to the sensitive nature of this data, we are working on data privacy polices (see Section 7) to be put in place before collecting this information.

The NetSage measurement services will include active measurements as well, starting with perfSONAR. The perfSONAR toolkit [perfSONAR, 2009][Tierney, Brian et al., 2009], is a network measurement toolkit designed to provide federated coverage of paths and help to establish end-to-end usage expectations. We have developed a perfSONAR exporter tool that pulls data from an open perfSONAR MA and inserts it into TSDS. As with SNMP

and Flow data, most of the IRNC backbones were already performing tests using perfSONAR to collect bandwidth and throughput data. We have set up our own set of tests on these hosts to unify the data across the full project. The current IRNC perfSONAR dashboard of tests, available at <http://data.ctc.transpac.org/maddash-webui/index.cgi?dashboard=Netsage%20Mesh>, shows the results of testing across IRNC hosts located at backbones and exchange points for each IRNC link.

One of the cornerstone principles of NetSage is a well understood and documented privacy policy when working with non-public network monitoring data. In this community, for example, flow data is seen as very sensitive. To that end, our initial data collection was for data, which is already shared publicly – SNMP and perfSONAR. Our default is to collect as little data as necessary in order to perform the analyses. For example, we will never collect payload data, only packet headers. Data is anonymized on site, and on-site data logs are limited in time. We do not publish raw data, only aggregated statistics of that data, which provides an additional layer of privacy as the behavior of individuals is rolled up and masked with the behavior of the other members of their aggregation group.

6. NetSage Archive

The second layer of the NetSage architecture is the data archive. We are using the Time Series Data System [TSDS, 2016], an Open Source software developed on commodity hardware, that provides a common archive shared with IRNC NOC. The system allows for well-structured and high performance storage and retrieval of time series data. TSDS is capable of tracking and reporting based on metadata, for example the system allows to view interface throughput from the viewpoint of a VLAN or BGP peer sessions from a particular ASN.

TSDS also provides the Time Series Query Language, which grants the possibility of easily generating reports about gathered data, including the ability to aggregate/summarize data over time, aggregate/summarize data based on one or more non-time dimensions, execute sub-queries to obtain incremental results, as well as the ability to perform a set of common aggregation functions for determining central tendency, frequency distribution etc. To give an example, once flow data is stored along with sufficient meta data in TSDS, a single query can be used to show the distribution of data transfer sizes between all known science facilities over the previous year, summarized by month and broken out by science domain.

7. Network Data Visualization

The top layer of the architecture (depicted in Figure 2) contains the data analysis systems and visualization components that will query the database to provide the desired services for end users. While much of the analysis work has yet to be started, the visualization team has produced several prototypes.

The NetSage visualization service will enable both near-real time and longitudinal monitoring of the interconnected R&E networks that are necessary to address the inquiries in Table 1. In line with this design approach, the NetSage portal will provide users with a query interface (depicted in Figure 4) that consists of a drop down menu of template natural language queries from which users can quickly customize. This approach will enable us to extend NetSage in the future to provide full natural language queries similar to Siri, for example. With this approach, it easier to express complex queries without having to be an expert in manipulating the query interface [Yiwen, S., 2012].




Figure 4: NetSage query selector.

As one example, the NetSage query: “What is the max, min, average bandwidth used between the IRNC links?” automatically produces the appropriate collection of charts needed to answer the query. First, a topology map is produced (as shown in Figure 5). The map shows color-coded links and exchange points depicting the average throughput. Second, the backbone link and exchange point throughput is represented on separate aligned histograms (shown in Figure 6) enable seamless visual comparison of individual links and exchange points relative to all the others. Similarly, the total data transmitted per element is also shown in relation to the total data transmitted.

The current prototype also allows for subsequent queries to be appended to facilitate custom report generation, such (as shown in Figure 7). This is a common workflow for end users understanding the data. Queries can also be saved as a URL that can be forwarded to provide NOC operators or engagement teams to enable a common view of the network when troubleshooting.



Figure 5: Topology map showing average bandwidth for the backbones and exchange points. Darker lines and nodes represent higher bandwidth.




Figure 6: Detail view of the aligned histograms for a set of the backbone links. These aligned histograms enable an easy comparison between links.

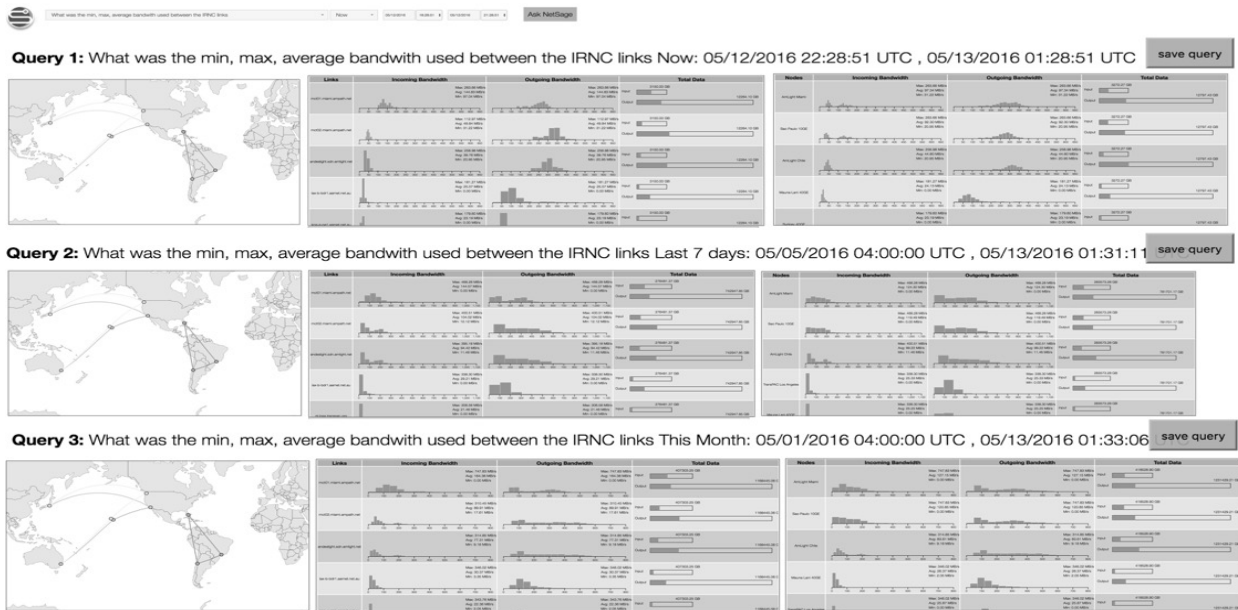


Figure 7: Current NetSage visualization prototype with three queries.

8. Conclusions and Future Work

The NetSage project is developing a framework for unified measurement and monitoring of the IRNC-funded backbones and exchange points, with an emphasis on open source software, privacy, analysis and visualization. The suite of tools being deployed enable end users to better understand network performance in a scalable, flexible way, using natural language queries.

The next steps of the project include expanding the data sources, both in terms of types and coverage and working to include measurement data from common data sources. In terms of visualization, the next steps include creating tools to visualize periodic patterns and elephant flows. We are also studying the possibility of implementing a Siri-like natural language querying interface. For more information, please visit our website, <http://www.netsage.global>.

References

- Dart, E., Rotman, L., Tierney, B., Hester, M., Zurawski, J., 2013 The Science DMZ: A network design pattern for data-intensive science. Proceedings of the IEEE/ACM Annual SuperComputing Conference (SC13).
- Case, J., Fedor, M., Schoffstall, M. Davin, J. 1990 Simple Network Management Protocol (SNMP).
- Tierney, B., Metzger, J., Boote, J., Boyd, E., Brown, A., Carlson, R., Zekauskas, M., Zurawski, J., Sawny, M., Grigoriev, M., 2009, perfSonar: Instantiating a global network measurement framework". Proceedings of the SOSP Wksp. Real overlays and Distrib. Sys.
- Tierney, B., Metzger, J., Boote, J., Boyd, E., Brown, A., Carlson, R., Zekauskas, M., Zurawski, J., Sawny, M., Grigoriev, M., 2009. perfSONAR toolkit. <http://www.perfsonar.net>
- Paxson, V., 1999 Bro a System for detecting network intruders in real time, Computer Networks, 31(23): p. 2435-2463
- Claise, B., Ed., 2004 Cisco Systems NetFlow Services Export Version 9.
- Phaal, P., Panchen, S., McKee, N., 2001 InMon Corporation's sFlow: A method for Monitoring Traffic in Switched and Routed Networks.
- Claise, B., Ed., Trammell, B., Ed., Aitken, P., 2013 Specification of the IP Flow Information Export (IPFIX) Protocol for the exchange of Flow information.
- TransPAC4, 2015 NSF funded project award: http://www.nsf.gov/awardsearch/showAward?AWD_ID=1450904
- PIREN, 2015 NSF funded project award: https://www.nsf.gov/awardsearch/showAward?AWD_ID=1451058
- America's Lightpath (AmLight), 2015 NSF funded project award: http://nsf.gov/awardsearch/showAward?AWD_ID=1451018
- SNAPP, 2016 SNMP Network analysis and Presentation Package. SNAPP home page available from: <http://globalnoc.iu.edu/gmcc-tools/snapp.html>
- ACE, 2010 ACE NSF funded project award: http://www.nsf.gov/awardsearch/showAward?AWD_ID=0962973
- StarLight, 2010 StarLight NSF funded project award: http://www.nsf.gov/awardsearch/showAward?AWD_ID=1450871
- AtlanticWave 2006 AtlanticWave NSF funded project award: <http://www.atlanticwave.net/index.html>
- Pacific Wave 2004 Pacific NSF funded project award: http://www.nsf.gov/awardsearch/showAward?AWD_ID=0441119
- Marco Mellia, Renato Lo Cigno, Fabio Neri, 2005 Measuring IP and TCP behavior on edge nodes with Tstat, Computer Networks, Vol.47, No.1, pp.1-21, ISSN: 1389-1286.
- Marco Mellia, Renato Lo Cigno, Fabio Neri. 2008 tstat home page available from: <http://tstat.polito.it/>
- Case, J., Fedor, M., Schoffstall, M. Davin, J. 1990 RFC1157.
- TSDS, 2016 TSDS home page available from: <http://globalnoc.iu.edu/software/measurements/tsds.html>
- Yiwen S., 2012 Articulate: Creating meaningful visualizations from natural language.

Biographies

Alberto Gonzalez is a PhD student at the University of Hawai'i at Mānoa. He works and develops his research at the LAVA Laboratory (<http://lava.manoa.hawaii.edu>) which focuses on visualization, analytics and intelligent systems.

Dr. Jason Leigh is the Director of LAVA: the Laboratory for Advanced Visualization & applications (<http://lava.manoa.hawaii.edu>), and Professor of Information & Computer Sciences at the University of Hawai'i at Mānoa. He is also Director Emeritus of the Electronic Visualization Lab and the Software Technologies Research Center at the University of Illinois at Chicago, where he maintains appointments in the Computer Science and Communications departments. His research expertise includes: Big data visualization; virtual reality; high performance networking; and video game design. He is co-inventor of the CAVE2 Hybrid Reality Environment, and SAGE: Scalable Adaptive Graphics Environment software, which has been licensed to Mechdyne Corporation & Vadiza Corporation, respectively.

Dr. Sean Peisert is jointly appointed as a staff scientist at Lawrence Berkeley National Laboratory, chief security strategist at CENIC, and associate adjunct professor of computer science at the University of California, Davis. His research interests cover a broad cross-section of computer security. Several recent projects include intrusion detection for control systems in power grids, and security in high-performance computing and networking environments. Dr. Peisert is vice chair of the IEEE Computer Society Technical Committee on Security & Privacy, an editorial board member of IEEE Security & Privacy magazine, and past general chair of the IEEE Symposium on Security and Privacy, the flagship conference for security research. He received his Ph.D., Master, and Bachelor degrees in Computer Science from UC San Diego.

Brian L. Tierney is a Staff Scientist and group leader of the ESnet Advanced Network Technologies Group at Lawrence Berkeley National Laboratory (LBNL), and is PI of ESnet's 100G Network Testbed Project. His research interests include high-performance networking and network protocols; distributed system performance monitoring and analysis; network tuning issues; and the application of distributed computing to problems in science and engineering. He has been the PI for several DOE research projects in network and Grid monitoring systems for data intensive distributed computing. Brian has been at LBNL since 1990.

Andrew Lee is the Principal Architect for the International networks group at Indiana University. In addition to his work with NetSage, he supports the TransPAC and America Connects to Europe (ACE) projects. Lee joined Indiana University in 2004, working within the Global Network Operations Center (GlobalNOC) on various research and education networks, including Internet2 and IU's Research Computing network. Prior to IU, he spent 10 years working for various commercial internet providers.

Dr. Jennifer M. Schopf is the director of International Networks at Indiana University, and the PI for the IRNC-funded NetSage, TransPAC, and ACE projects. Her group oversees a multi-million dollar per year system of networks and monitoring between the US, Asia, and Europe. Prior to IU, she was a program officer at the US National Science Foundation. She has co-edited a book, co-authored over 50 refereed publications, and given over 120 invited talks.