UC Davis
UC Davis Previously Published Works

Title

Big Data and Analysis of Data Transfers for International Research Networks Using NetSage

Authors

Gonzalez, Alberto

Leigh, Jason

Peisert, Sean

et al.

Peer reviewed

Big Data and Analysis of Data Transfers for
International Research Networks Using NetSage

Alberto Gonzalez[1], Jason Leigh[1], Sean Peisert[2,3], Brian Tierney[2,3], Edward Balas[4], Predrag Radulovic[4], and
Jennifer M. Schopf[4]

1. Laboratory for Advanced Visualization & Applications, University of Hawaii at Mānoa.
agon@hawaii.edu, leighj@hawaii.edu.
2. Lawrence Berkeley National Laboratory. sppeisert@lbl.gov, bltierney@es.net.
3. University of California, Davis
4. Indiana University. ebalas@iu.edu, pradulov@iu.edu, jmschopf@iu.edu.

Abstract—**Modern science is increasingly data-driven and collaborative in nature. Many scientific disciplines, including genomics, high-energy physics, astronomy, and atmospheric science, produce petabytes of data that must be shared with collaborators all over the world. The National Science Foundation-supported International Research Network Connection (IRNC) links have been essential to enabling this collaboration, but as data sharing has increased, so has the amount of information being collected to understand network performance. New capabilities to measure and analyze the performance of international wide-area networks are essential to ensure end-users are able to take full advantage of such infrastructure for their big data applications.**

NetSage is a project to develop a unified, open, privacy-aware network measurement, and visualization service to address the needs of monitoring today's high-speed international research networks. NetSage collects data on both backbone links and exchange points, which can be as much as 1Tb per month. This puts a significant strain on hardware, not only in terms storage needs to hold multi-year historical data, but also in terms of processor and memory needs to analyze the data to understand network behaviors.

This paper addresses the basic NetSage architecture, its current data collection and archiving approach, and details the constraints of dealing with this big data problem of handling vast amounts of monitoring data, while providing useful, extensible visualization to end users.

IRNC Measurement, Infrastructure, International, Tools, Visualization, Analytics.

I. Introduction

Much of modern science is now driven by data from scientific instruments, some producing even petabytes of data, which is then shared and analyzed by thousands or tens of thousands of scientists all over the world. Some of the "big data" challenges in this space are related to the growing archive of network data available for analysis, the diversity of available data sets, and the need for long-term storage in order to do trend analysis. New methods to monitor, analyze, and understand the performance of big data transfers are needed to assure that end-users are able to take full advantage of networking infrastructure. This complex cyberinfrastructure is rapidly increasing our ability to produce, manage, and use data [1].

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

This paper describes the NetSage system. In Section 2 of this paper we provide an overview of NetSage and its use cases. In Section 3, we describe the NetSage monitoring architecture, with additional details given in Section 4 for the data sources. Section 5 describes the archive, and Section 6- the visualization. In each section, we detail the "big data" needs and challenges, and how we are approaching them.

II. NetSage Overview

NetSage was originally developed to better understand the behavior of the NSF-funded International Research Network Connections (IRNC) backbone networks and exchange points, shown in Figure 1. The IRNC-funded backbones include: TransPAC4 [7], AmLight ExP [8], PIREN (Pacific Islands Research and Education Networks) [9], ACE (America Connects to Europe) [10], AtlanticWave [11], StarLight [12], and Pacific Wave [13],

Figure 1. The current NSF IRNC-funded backbones and Exchange Points, showing link capacity for each network.
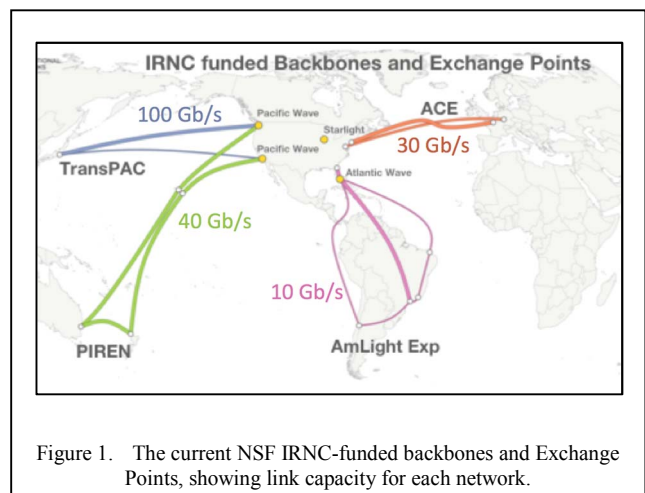
as well as their respective exchange points. Overall, these connections carry the majority of U.S. research traffic to and from their international collaborators. They operate 24 hours a day, 7 days a week, and combined, represent over 500Gb/s of network capacity.

NetSage is building and deploying advanced measurement services and an exploratory visualization platform to benefit science and engineering communities dealing with big data transfers in multiple ways by focusing on providing a better understanding of:

- current traffic patterns across IRNC links, and the ability to better understand growth trends for capacity-planning purposes;
- the main sources and sinks of "large, long fat network" (LFN) or "elephant flows," to know where to focus attention on outreach and training;
- 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;
- how the IRNC links are being used by the different research domains and institutions.

NetSage services will provide an unprecedented combination of passive measurements, including SNMP data [2] [3], flow data [4], and traffic header analysis [5], as well as active measurements [6], to create longitudinal network performance data visualizations.

For monitoring and analysis 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 Directors 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 Operations 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.
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 could 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.

TABLE I. DIRECTED QUESTIONS THAT ARE GUIDING NETSAGE DEVELOPMENT.

Goals	Insight Desired	Data Required
Understand traffic patterns for capacity planning	What is the max, min, average bandwidth used on links? What is the duration and are there any periodic patterns or peak periods? Which exchange points are congested?	Throughput data (SNMP) Packet Headers (Flow Data)
Understand the impact, location and cause of packet loss	Is loss due to the client, the link, or the exchange point? 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?	Packet loss and packet headers (tstat) Bandwidth (perfSONAR) Latency (perfSONAR)
Understand the nature of elephant flows that use the links	What is the max, min average duration of elephant flows? How many elephant flows tend to occur at the same time? Which exchange point or networks service the most elephant flows? Are the flows bursty or experiencing sustained traffic?	Packet Headers (Flow Data and tstat) Packet Loss (tstat)

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.

III. NETSAGE MONITORING ARCHITECTURE

The NetSage architecture design was primarily driven by the application/algorithm needs and the system/platform-level requirements as described in [14].

NetSage consists of a 3-layered architecture, shown in Figure 2, similar to many other monitoring approaches. The bottom layer consists of various data sources. The monitoring systems 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 combine multiple, different active and passive measurements, as detailed in Section 4.
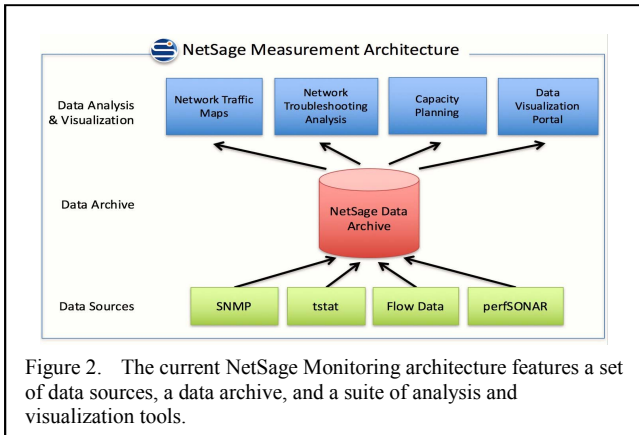

Figure 2. The current NetSage Monitoring architecture features a set of data sources, a data archive, and a suite of analysis and visualization tools.

The second layer consists of the NetSage archive service. In order to enable the ability to perform analysis in near-real time and to enhance the system's performance when coping with the inherently different data types (Time series data and aggregations of data) several approaches were explored. The first uses the Time Series Data Service (TSDS) [14] to implement a shared archive with the IRNC NOC (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 higher level services to integrate and query for these multisource time series data from the same repository. The second approach uses the "ELK stack" [16] to create aggregations summaries of individual flow records. (detailed in Section 5).

The third layer is a suite of visual analytics tools that address the questions outlined in Table 1. This is detailed in Section 6.

IV. NetSage Data Sources

The core of the NetSage infrastructure is the collection of data related to both the network links and exchange points. These are a combination of passive measurements, such as SNMP, flow data, and data from packet header inspection, and active measurements, currently perfSONAR.

The Simple Network Management Protocol (SNMP) is an application–layer protocol defined in RFC1157 [2] [3] 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
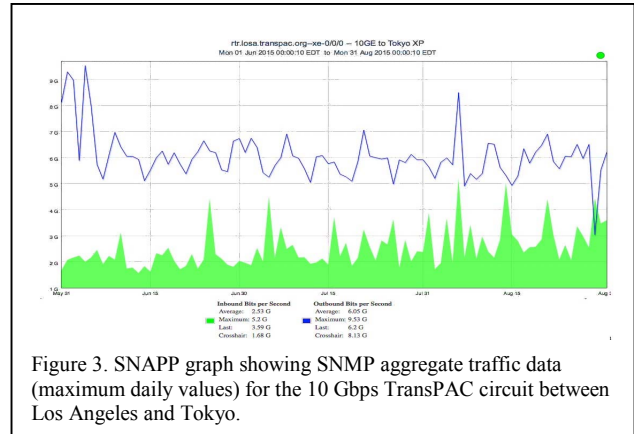

Figure 3. SNAPP graph showing SNMP aggregate traffic data (maximum daily values) for the 10 Gbps TransPAC circuit between Los Angeles and Tokyo.

example of SNMP data as shown in the open source tool SNAPP [17], maintained by the GlobalNOC [18]. The SNMP data is collected every 60 seconds for each of the links.

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 [19] 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 the "Tstat" tool [5] 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. We expect this data to grow to the order of at least about 1 TB per month per link, depending on what additional parameters we request.

Another source of passive data for networks is related to flow data collection. Depending on the hardware, this might be NetFlow [20], sFlow [21], or IPFIX [4]. Flow data will allow us 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. Currently, the TransPAC and ACE links generate 40-100GB of flow data per link per month. Flow data collected from exchange points is expected to be around 10 times bigger, up to 400GB - 1TB per exchange point per month.

Analysis of the flow data indicates that distribution of flows in R&E networks is heavily skewed towards large number of small flows (<100kB), and significant but comparatively smaller, number of very large flows. Flows with sizes of >100MB account for majority of the of the traffic, and that influenced decision to limit analysis of flows of that size or larger ("elephant flows"). Since such flows are much smaller in number, this eases strain on hardware for analysis and helps manage dataset sizes for multiple links and exchange points, from our preliminary tests we expect an average data reduction of 2 orders of magnitude.

The NetSage measurement services include active measurements as well, currently in the form of perfSONAR[4] [22] tests to gather latency and bandwidth information. PerfSONAR 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 our archive, TSDS. We have set up tests for each of the backbone links and exchange points across the full project. The current IRNC perfSONAR dashboard of tests is available at TransPAC dashboard [23]. PerfSONAR data is collected every 6 hours for bandwidth tests and every 60 seconds for losses. We are currently storing around a few GB of PerfSONAR monthly data.

V.   NETSAGE ARCHIVES

The second layer of the NetSage architecture is the data archive. NetSage uses a three-tier approach to flow data archives. With each tier having a specific set of capabilities and intended use

After flow data is de-identified these individual flow records are stored as raw unindexed messages in canonical repository. The canonical storage is not used for analysis directly, its purpose is to ensure we can regenerate later tiers if we need to redesign.

After raw storage, the flow records are inserted into an Elasticsearch database using Logstash. This tier provides flexible query access to indexed individual flow records. This data is kept for multiple weeks and is used to generate aggregated summaries of traffic activity and to perform adhoc / exploratory data analysis.

The final tier contains pre-generated aggregates statistics derived from the individual flow records stored in Elasticsearch. The types of summaries include: the distribution of traffic volume by protocol over time, the distribution of traffic by source, etc.

We are using the Time Series Data Service (TSDS) [17], 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 AS.

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.

The NetSage project is also using "E*lasticsearch*", part of the "ELK Stack" to provide a scalable yet flexible system for indexing structured data that could be used with flow records and other added contextual metadata. "Elasticsearch" is a distributed, JSON-based search and analytics engine designed for horizontal scalability, maximum reliability, and easy management. The NetSage project intends to leverage ELK's strengths by using "Elasticsearch" for internal analysis on individual flows and then create a set of queries to generate aggregate flow data which will be then stored in TSDS.

The NetSage project is using TSDS and ELK simultaneously to balance their different strengths and drawbacks. ELK stores individual records which are indexed directly. This is good for storing event logs and other types of points in time data, but for storing regularly interval measurements, like measurements done every minute, the large index size becomes problematic. TSDS excels at storing multiple measurements for a set of stable metadata, such as when it is required to measure the same groups of interfaces over a period of thousands of measurements. But this approach does not work with flow records due to the metadata variability, TSDS suffers losing efficiency due to metadata churn.

VI.   NETWORK DATA VISUALIZATION

The top layer of the architecture as depicted in Figure 2, contains the data analysis systems and visualization components that will query the underlying database. The NetSage visualization service will enable both near-real time monitoring and longitudinal analysis of the interconnected R&E networks that are necessary to address the inquiries in Table 1.

Large-Scale networks have become increasingly challenging to manage by system administrators and/or network managers. Network anomaly detection is difficult due to its vast data volume, large numbers of attributes, interconnectivity/causality and high dynamics (i.e., connections can be established or broken at any moment, because users may come and go) [30]. Even when data mining and machine learning have proven effective in detecting anomalies, it is often the case that patterns are not known ahead of time by network administrators and operators [31]. Visualization tools are needed to allow them
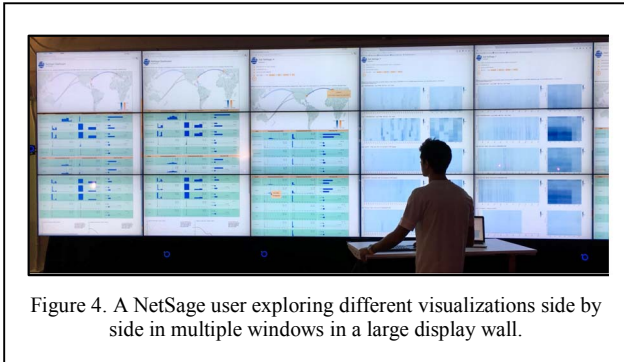

Figure 4. A NetSage user exploring different visualizations side by side in multiple windows in a large display wall.
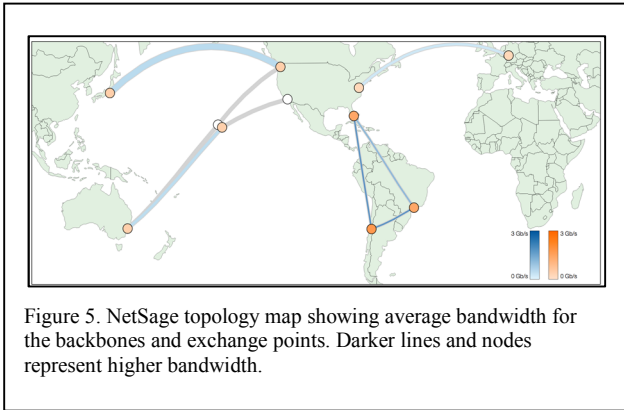

Figure 5. NetSage topology map showing average bandwidth for the backbones and exchange points. Darker lines and nodes represent higher bandwidth.

to find these patterns by examining past data. While numerous network management and visualization tools have existed in the past, few of them are lightweight enough or specifically designed towards anomaly detection in dynamic network traffic data. The area of visualizing anomalous events in particular is still actively been studied by the visualization research community [30] and is an area of interest in the NetSage project as well.
When visualizing large volumes of data, it is often difficult to see the bigger picture in the details. To improve big data exploration NetSage's approach is to enable visualizations from multiple queries to be juxtaposed simultaneously to provide a multi-faceted view of the network phenomenon being investigated. Prior research in ultra-high resolution display walls show that users are able to come to conclusions with greater creativity, speed, accuracy, comprehensiveness, and confidence using such

environments [27, 28]. To achieve this multi-faceted view NetSage leverages SAGE2 (the Scalable Amplified Group Environment) a widely used open source middleware system for supporting visualization on large ultra high resolution displays [29]. NetSage was designed so that each query can generate its own independent visualization. These visualizations can then be combined as needed in SAGE2 like jigsaw pieces in a puzzle to help the user answer complex questions about the networks, as shown in Figure 4.
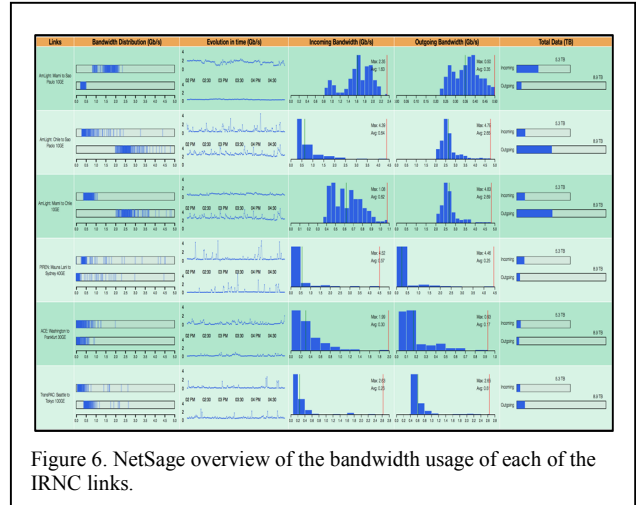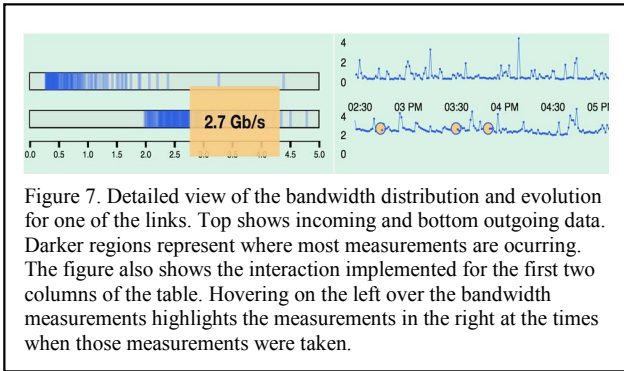

Figure 6. NetSage overview of the bandwidth usage of each of the IRNC links.

The NetSage web site shows a default dashboard that provides an overview of the current state of all IRNC links for the last three hours. The Dashboard, which updates every five minutes, begins with a map, represented in Figure 5 showing the state of the links and exchange points where the size of the links represents the relative capacity of the connections. Links are colored on a blue scale while nodes are colored on an orange scale. Grey links and white nodes have no data for the last three hours. While viewing the visualizations hovering over the data points shows the underlying data values.

The table below the map, as shown in Figure 6 gives more detailed information about the links where each column is intended to give a different perspective of the data. The first column of the table, shows the incoming and outgoing bandwidth distribution for each link in the past 3 hours where the horizontal axis represents Gb/s and each bandwidth measurement is represented as a blue vertical line. This allows NetSage users to visualize where most of the bandwidth measurements for each link are occurring, as can be seen by the slightly darker marks shown in Figure 6.

Figure 7. Detailed view of the bandwidth distribution and evolution for one of the links. Top shows incoming and bottom outgoing data. Darker regions represent where most measurements are occurring. The figure also shows the interaction implemented for the first two columns of the table. Hovering on the left over the bandwidth measurements highlights the measurements in the right at the times when those measurements were taken.

The second column shown in Figure 6 shows the incoming and outgoing bandwidth for each of the links in the last 3 hours. This column allows NetSage users to understand the evolution of the bandwidth for the last 3 hours. Column one and two of the table in Figure 7 are intended to be used simultaneously: When hovering over the bandwidth measurements of the first column, those same measurements are highlighted in the second column. Note that all the visualizations in this first and second column share the same axis to allow fair comparison across all the links.

The next two columns of the table show histograms of incoming and outgoing bandwidth for the connections as well as the maximum and average values, which are visualized as a red and green vertical lines respectively. The histograms show the distribution of the bandwidth over the link and the overall network.

The last column of the table shows incoming and outgoing data transferred per link relative to the total incoming and outgoing data transmitted by all the links in the IRNC network. This facilitates comparison of how much each link has contributed to the total data transmission of all links in the last three hours.

Clicking on the "Ask NetSage" button brings to view NetSage's query interface. The NetSage portal provides users with a query interface that directly follows the questions depicted in Table 1. It consists of a dropdown menu from which users can quickly customize, as shown in

Figure 8. Detailed view of the Ask NetSage query menu.

Figure 11. Detailed view of a heatmap showing latency data for one of the links from July 2015 till December 2015. Darker squares represent higher latencies for those hours. This is a simple visualization method that allows to spot patterns over long periods of time. We pretend to implement automated methods for detecting those patterns automatically in order to reduce the complexity to NetSage users.
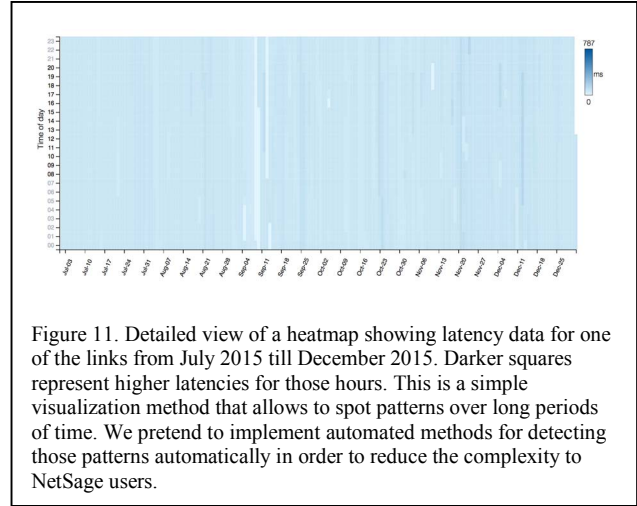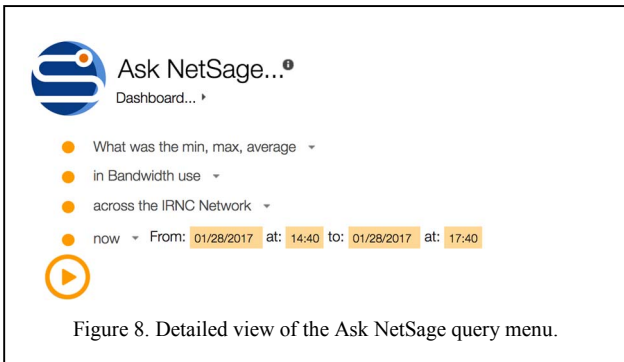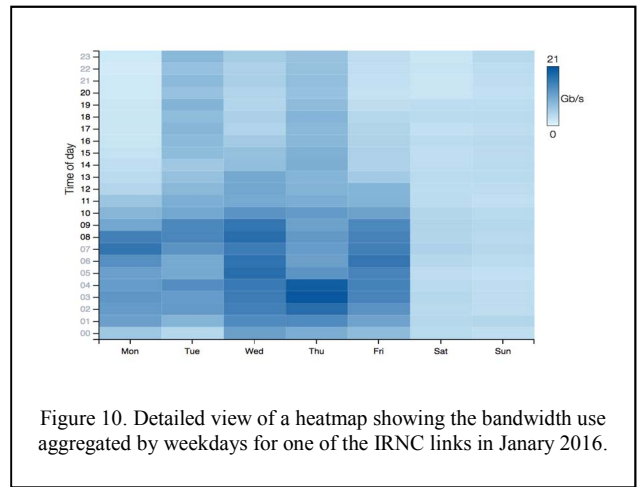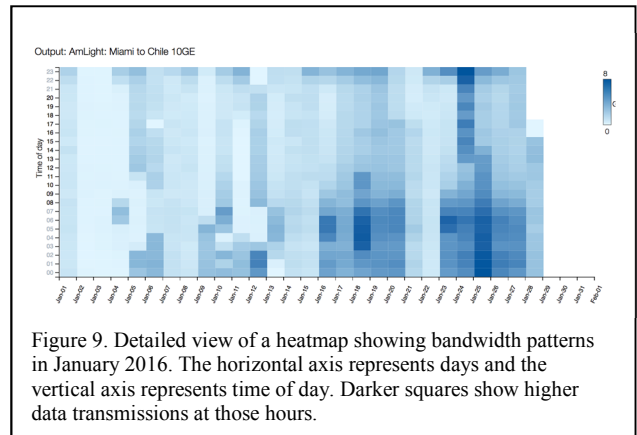
Figure 8 With this approach, it is easier to express complex queries and control the underlying data without having to be an expert in manipulating the query interface [24, 25].

As one example, the NetSage query: "What is the max, min, average in bandwidth use across the IRNC network in

Figure 9. Detailed view of a heatmap showing bandwidth patterns in January 2016. The horizontal axis represents days and the vertical axis represents time of day. Darker squares show higher data transmissions at those hours.

Figure 10. Detailed view of a heatmap showing the bandwidth use aggregated by weekdays for one of the IRNC links in Janary 2016.

the last 7 days?" automatically produces visualizations similar to the ones described earlier except that based on users customized query parameters.

As another example of how NetSage is leveraging the visualizations to address managing large scale data during long periods of time, the NetSage query: "What is the duration and are there any periodic patterns or peak periods in bandwidth use across the IRNC network in the last month?" automatically produces a series of heatmaps over the period selected. These heatmap visualizations allow to cluster data by day or weekday without losing resolution in the input data. As it can be seen in Figure 9, in the vertical axis we have the hour of day in the horizontal axis we have a progression of days. The darker the color in the heatmap the higher amount of data was transmitted in that given periods of time. The chart to the right, also shown in Figure 10, shows an aggregation of the same information over weekly periods. Hovering the mouse over a particular data point shows more detailed information. This same visualization approach, as shown in Figure 11, can effectively be applied to other measurements such as perfSONAR tests for network loss or latency.

As mentioned earlier, all new queries are created in new tabs to facilitate comparisons of the data across large displays or walls, as shown in Figure 4. Each of these queries generate unique URLs which allows charts to be easily shared between network users and administrators to help them troubleshoot problems together.

VII. CONCLUSIONS AND FUTURE WORK

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

Presently the fundamental components of NetSage have been deployed to enable us to begin to collect, archive and visualize data from the IRNC backbones and exchange points. Next steps include expanding the data sources, both in terms of types and coverage and working to include measurement data from common data sources. At the same time, we are working on creating a science registry which will allow us to map flows to science projects in order to understand how the IRNC networks are being used and to gain insight into new research questions that have not yet been addressed. For example: how do researchers worldwide in various disciplines share data with each other. In terms of analytics and visualization, the next steps include creating meaningful aggregations of flow data and creating tools to visualize those aggregations to gain insight on big data transmissions and elephant flows. We also intend to implement machine learning capabilities for predicting future use or predicted losses as well as tools to automatically detect patterns on the visualizations to reduce the cognitive effort of NetSage users.

While the big data challenge for NetSage's data gathering effort is in overcoming the sheer volume of the data, the challenge for visualization is helping network troubleshooters understand the complex relationships between network parameters that ultimately affect network performance.

The portal can currently be accessed from http://www.netsage.global and there is already growing interest by international network partners in adopting our methodology so that in the future an open standard for the collection, sharing, and analysis of network data can be created.

REFERENCES

[1] Atkins, D Revolutionazing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastucture, 2003. www.cise.nsf.gov/sci/reports/toc.cfm

[2] Case, J., Fedor, M., Schoffstall, M. Davin, J. Simple Network Management Protocol (SNMP). 1990

[3] Case, J., Fedor, M., Schoffstall, M. Davin, J. RFC1157. 1990

[4] Claise, B., Ed., Trammell, B., Ed., Aitken, P., Specification of the IP Flow Information Export (IPFIX) Protocol for the exchange of Flow information. 2013

[5] Marco M., Renato Lo C., Fabio N., Measuring IP and TCP behavior on edge nodes with Tstat, Computer Networks, Vol.47, No.1, pp.1-21, ISSN: 1389-1286. 2005

[6] Tierney, B., Metzger, J., Boote, J., Boyd, E., Brown, A., Carlson, R., Zekauskas, M., Zurawski, J., Sawny, M., Grigoriev, M. PerfSONAR toolkit. 2009. http://www.perfsonar.net

[7] TransPAC4, 2015 NSF founded project award: http://www.nsf.gov/awardsearch/showAward?AWD_ID=1450904.

[8] America's Lightpath (AmLight), 2015 NSF founded project award: http://nsf.gov/awardsearch/showAward?AWD_ID=1451018.

[9] PIREN, 2015 NSF founded project award: https://www.nsf.gov/awardsearch/showAward?AWD_ID=1451058.

[10] ACE, 2010 ACE NSF founded project award: http://www.nsf.gov/awardsearch/showAward?AWD_ID=0962973

[11] AtlanticWave 2006 AtlanticWave NSF founded project award: http://www.atlanticwave.net/index.html

[12] StarLight, 2010 StarLight NSF founded project award: http://www.nsf.gov/awardsearch/showAward?AWD_ID=1450871.

[13] Pacific Wave 2004 Pacific NSF founded project award: http://www.nsf.gov/awardsearch/showAward?AWD_ID=0441119

[14] Singh, D., Reddy, CK., A survey on platforms for big data analytics. Journal of Big Data. 2014.

[15] TSDS, 2016 TSDS home page available from: http://globalnoc.iu.edu/software/measurements/tsds.html

[16] ELK stack website: https://www.elastic.co

[17] SNAPP, 2016 SNMP Network analysis and Presentation Package. SNAPP home page is available at the following URL: http://globalnoc.iu.edu/grnoc-tools/snapp.html

[18] Global NOC website: http://globalnoc.iu.edu

[19] Paxson, V., Bro a System for detecting network intruders in real time, Computer Networks, 31(23): p. 2435-2463. 1999

[20] Claise, B., Ed., Cisco Systems NetFlow Services Export Version 9. 2004

[21] Phaal, P., Panchen, S., McKee, N., InMon Corporation's sFlow: A method for Monitoring Traffic in Switched and Routed Networks. 2001

[22] Tierney, B., Metzger, J., Boote, J., Boyd, E., Brown, A., Carlson, R., Zekauskas, M., Zurawski, J., Sawny, M., Grigoriev, M.,

perfSonar: Instantiating a global network measurement framework". Proceedings of the SOSP Wksp. Real overlays and Distrib. Sys. 2009

[23]   Transpac dashboard website is available at the following URL: http://data.ctc.transpac.org/maddash-webui/index.cgi?dashboard=Netsage%20Mesh

[24]   Yiwen S., Articulate: Creating meaningful visualizations from natural language. 2012

[25]   Aurisano, J., Kumar, A., Gonzalez, A., Leigh, J., DiEugenio, B., Johnson, A.,, Towards a Dialoge System that supports Rich visualizations of Data. The 17ᵗʰ annual SIGdial Meeting on discourse and Dialogue. SIGDIAL. 2012.

[26]   Marco Mellia, Renato Lo Cigno, Fabio Neri. Measuring IP and TCP behavior on edge nodes with Tstat. 2008.

[27]   Leigh, J., Johnson, A., Renambot, L., Peterka, T., Jeong, B., Sandin, D., Talandis, J., Jagodic, R., Nam, S., Hur, H., Sun, Y., Scalable Resolution Display Walls, Proceedings of the IEEE, Issue 99, pp. 1-15, May 1, 2012.

[28]   Reda, K., Offord, C., Johnson, A., Leigh, J., Expanding the Porthole: Leveraging Large, High-Resolution Displays in Exploratory Visual Analysis, Proceedings of CHI'14 EA on Human Factors on Computing Systems, Toronto, Ontario, Canada, pp. 2047-2052, ACM, April 26, 2014

[29]   Marrinan, T., Aurisano, J., Nishimoto, A., Bharadwaj, K., Mateevitsi, V., Renambot, L., Long, L., Johnson A., Leigh, J., SAGE2: A New Approach for Data Intensive Collaboration Using Scalable Resolution Shared Displays, In Proceedings of the 10th IEEE International Conference on Collaborative Computing: Networking, Applications and Worksharing, Miami, FL, October 22, 2014

[30]   Zhang, T., Liao, Q., Shi, L., Bridging the Gap of Network Management and Anomaly Detection through Interactive Visualization. IEEE Pacific Visualization Symposium. 2014

[31]   Effective Network management via dynamic network anomaly visualization. Liao, Q., and Li, T., International Journal of Network Management. 2016.