



International Research Networks Using Big data and Analysis of NetStage Data Transfers

G.Kavitha¹, Mrs.K.Bala², Dr.D.Rajinigrinath³

¹ Final year, CSE Department,

² Assistant Professor, CSE Department,

³ HOD, CSE Department,

Sri Muthukumaran Institute of Technology, Chikkarayapuram, Chennai, Tamil Nadu-600069, India.

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.

Keywords: 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 cyber infrastructure 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. NETSTAGE 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].

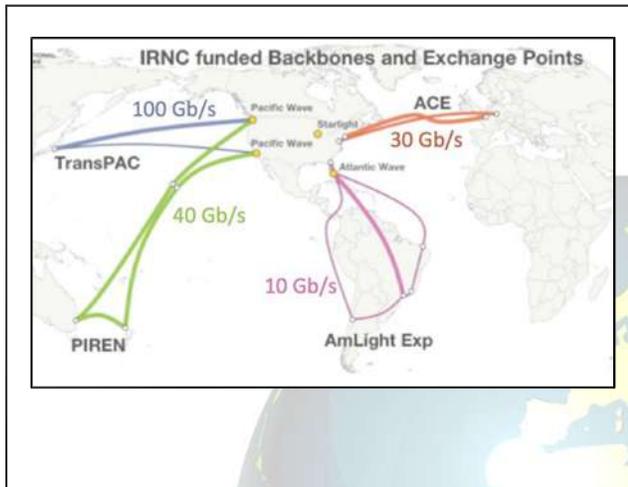


Figure 1. The current NSF IRNC-funded backbones and Exchange Points, showing link capacity for each network

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: 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 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

- Current traffic patterns across IRNC links, and the 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. NETSTAGE 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.

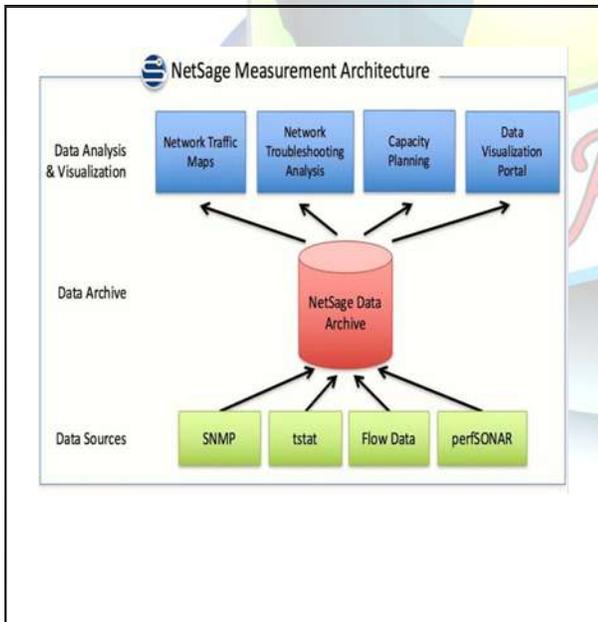


Figure 2. The current NetSage Monitoring architecture features a set of data sources, a data archive, and a suite of analysis and visualization tools

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.

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 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

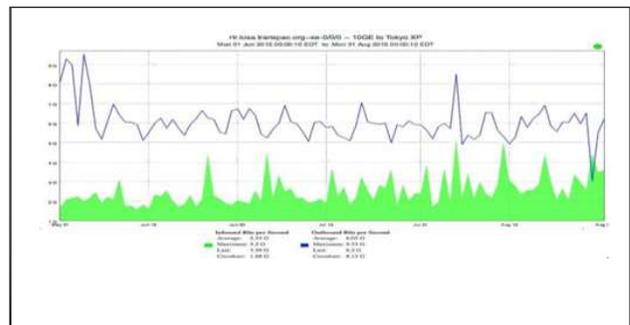




Figure 3. SNAPP graph showing SNMP aggregate traffic data (maximum daily values) for the 10 Gbps TransPAC circuit between Los Angeles and Tokyo

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.

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 “Elasticsearch”, 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. Project developed by Munafó and Mellia at 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.

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 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.

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. 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.

Another source of passive data for networks is related to flow data collection. Depending on the hardware, this might be NetFlow [20], 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.

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. 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 one example, the NetSage query: “What is the max, min, average in bandwidth use across the IRNC network in the last 7 days?” automatically produces visualizations similar to the ones described earlier except that based on users customized query parameters.

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.

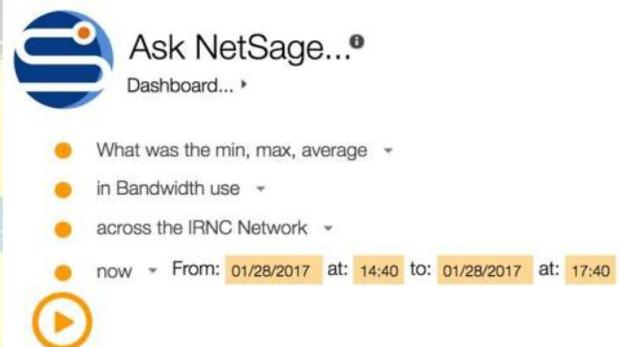


Fig 4: Asking netstage for task

VII. CONCLUSION

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.

REFERENCES

- [1]. Atkins, D Revolutionazing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure, 2003. www.cise.nsf.gov/sci/reports/toc.cfm
- [2]. Case, J., Fedor, M., Schoffstall, M. Davin, J. RFC1157. 1990
- [3]. Claise, B., Ed., Trammell, B., Ed., Aitken, P., Specification of the IP Flow Information Export (IPFIX) Protocol for the exchange of Flow information. 2013



- [4]. 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
- [5]. 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>.
- [6]. TransPAC4, 2015 NSF founded project award:http://www.nsf.gov/awardsearch/showAward?AWD_ID=1450904.
- [7]. America's Lightpath (AmLight), 2015 NSF founded project award:
http://nsf.gov/awardsearch/showAward?AWD_ID=1451018.
- [8]. PIREN, 2015 NSF founded project award:
https://www.nsf.gov/awardsearch/showAward?AWD_ID=1451058.
- [9]. ACE, 2010 ACE NSF founded project award:
http://www.nsf.gov/awardsearch/showAward?AWD_ID=0962973
- [10]. AtlanticWave 2006 AtlanticWave NSF founded project award:
<http://www.atlanticwave.net/index.html>
- [11]. StarLight, 2010 StarLight NSF founded project award:http://www.nsf.gov/awardsearch/showAward?AWD_ID=1450871.
- [12]. Pacific Wave 2004 Pacific NSF founded project award:
http://www.nsf.gov/awardsearch/showAward?AWD_ID=0441119
- [13]. Singh, D., Reddy, CK., A survey on platforms for big data analytics. Journal of Big Data. 2014.
- [14]. TSDS, 2016 TSDS home page available from:
<http://globalnoc.iu.edu/software/measurements/tsds.html>
- [15]. ELK stack website: <https://www.elastic.co>
- [16]. SNAPP, 2016 SNMP Network analysis and Presentation Package. SNAPP home page is available at the following URL: <http://globalnoc.iu.edu/gnoc-tools/snapp.html>
- [17]. Global NOC website: <http://globalnoc.iu.edu>
- [18]. Paxson, V., Bro a System for detecting network intruders in real time, Computer Networks, 31(23): p. 2435-2463. 1999
- [19]. Claise, B., Ed., Cisco Systems NetFlow Services Export Version 9. 2004
- [20]. Phaal, P., Panchen, S., McKee, N., InMon Corporation's sFlow: A method for Monitoring Traffic in Switched and Routed Networks. 2001
- [21]. 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
- [22]. NLK stack website: <https://www.elastic.co>

