ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 5, Issue 4, April 2018

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

[1] Dr. A P Nirmala, [2] Siddhant Kumar Sahu
[1] Sr. Asst. Prof, [2] PG Student
[1][2]Dept of MCA, New Horizon College of Engineering

Abstract: NetSage is 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.

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].

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 NC 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 NC links are being used by the different research domains and institutions.

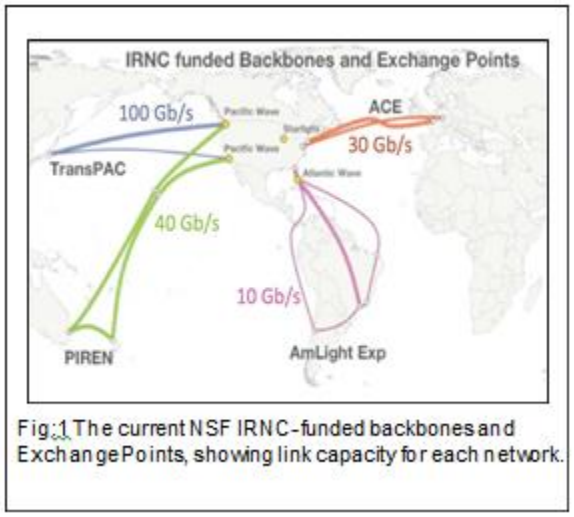

Fig:1 The current NSF IRNC-funded backbones and Exchange Points, showing link capacity for each network.

NETSAGE MONITORING ARCHITECTURE:

NetSage consists of a 3-layered architecture, shown in bellow Figure, 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.

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 5, Issue 4, April 2018

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) [11] 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" to create aggregations summaries of individual flow records.

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

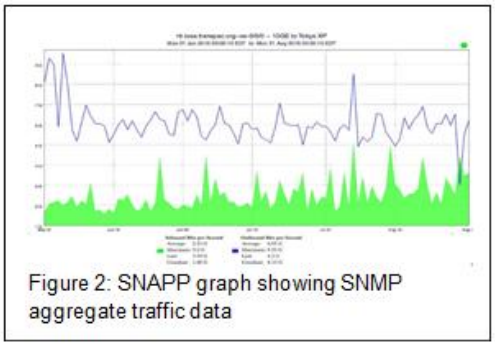

Figure 2: SNAPP graph showing SNMP aggregate traffic data

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.

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 [7] 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 Instead, we are currently deploying the "Tstat" tool [5] for our packet header inspection tool. Tstat is part of the EU Measurement Plane (mplane) FP7 project developed by Munafo and Mellia at Polytechnic 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 Net Flow , sFlow, 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] tests to gather latency and bandwidth information. PerfSONAR [4] is a network measurement toolkit designed to provide federated coverage of paths and help to establish end-to-end usage expectations. We

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 5, Issue 4, April 2018

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 [12]. 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.

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 Elastic search 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 exploratory data analysis.

The final tier contains pre-generated aggregates statistics derived from the individual flow records stored in Elastic search. 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) [12], 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 "Elastic search", 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.

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.

Large-Scale networks have become increasingly challenging to manage by system administrators and/or

Figure 3: A NetSage user exploring different visualizations side by side in multiple windows in a large

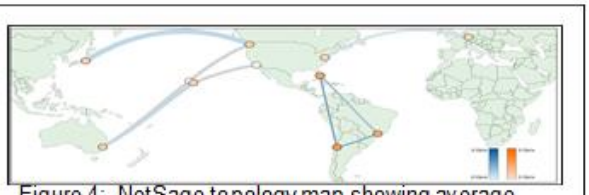

Figure 4: NetSage topology map showing average bandwidth for the backbones and exchange points. Darker lines and nodes represent higher bandwidth.

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. Even when data mining and machine learning have proven effective in detecting anomalies, it is often the case that patterns are

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 5, Issue 4, April 2018

not known ahead of time by network administrators and operators.

Visualization tools are needed to allow them 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 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 [7]. 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 [13]. 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 Figure5.

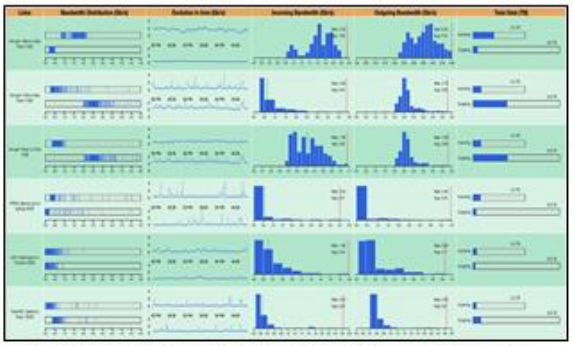

Figure 5: NetSage overview of the bandwidth usage of each of the IRNC links.

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 .

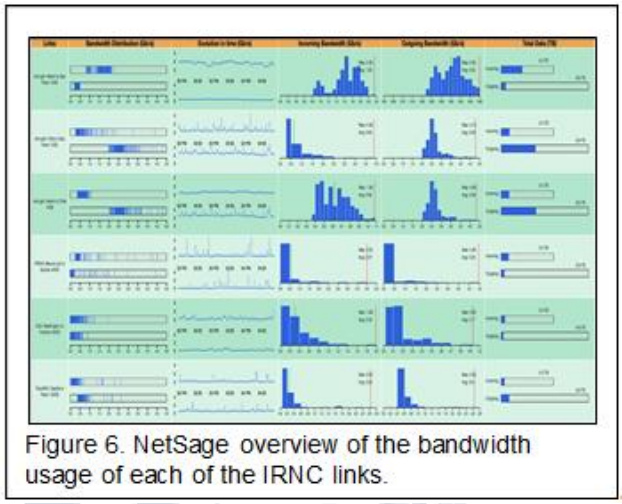

Figure 6. NetSage overview of the bandwidth usage of each of the IRNC links.

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.

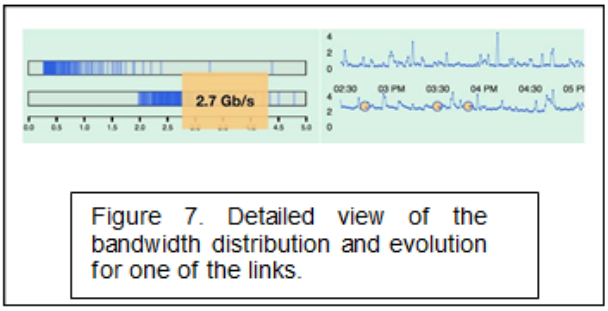

Figure 7. Detailed view of the bandwidth distribution and evolution for one of the links.

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

ISSN (Online) 2394-2320

International Journal of Engineering Research in Computer Science and Engineering (IJERCSE)
Vol 5, Issue 4, April 2018

the table. Hovering on the left over the bandwidth measurements highlights the measurements in the right at the times when those measurements were taken. 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. It consists of a dropdown menu from which users can quickly customize, as shown in 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.

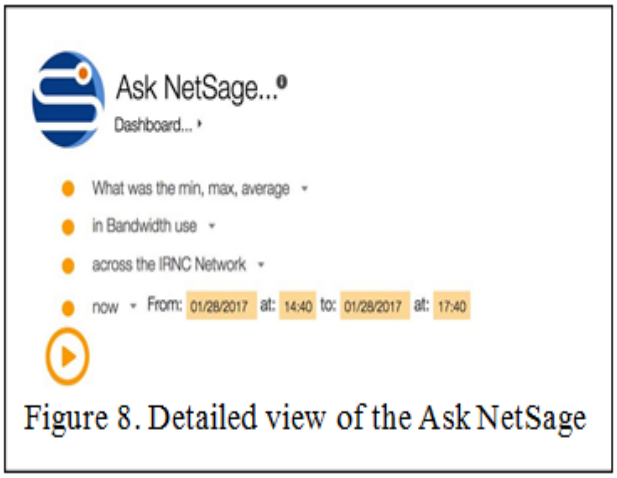

Figure 8. Detailed view of the Ask NetSage

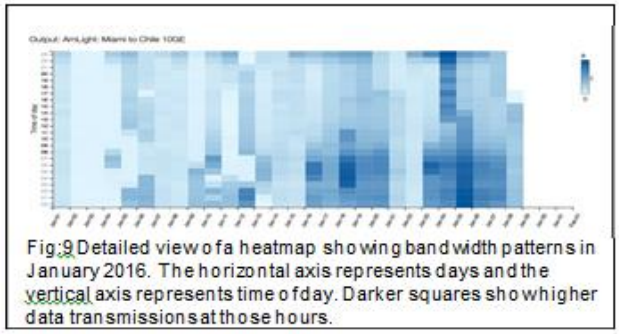
Fig:9 Detailed view of a heatmap showing bandwidth patterns in January 2016. The horizontal axis represents days and the vertical axis represents time of day. Darker squares show higher data transmissions at those hours.

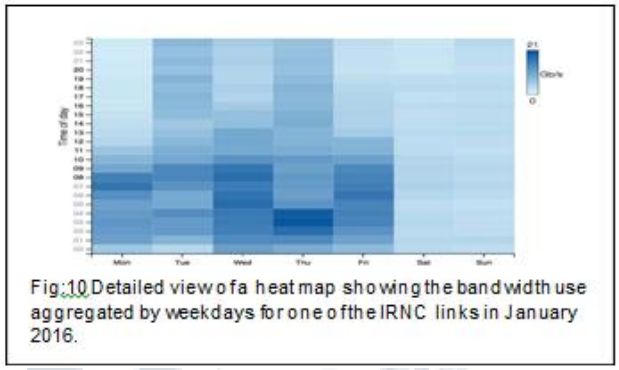
Fig:10 Detailed view of a heat map showing the bandwidth use aggregated by weekdays for one of the IRNC links in January 2016.

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 can effectively be applied to other measurements such as perfSONAR tests for network loss or latency.

CONCLUSIONS AND FUTURE SCOPE:

The NetSage is developing a framework for unified measurement and monitoring of the big data transfers over the research-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. 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.

REFERENCES:

1 Atkins, D Revolutionazing Science and Engineering Through Cyberinfrastructure: Report of the National Science Foundation Blue-Ribbon Advisory Panel on Cyberinfrastructure, 2003.

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.

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

8 Claise, B., Ed., Cisco Systems NetFlow Services Export Version 9. 2004.

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

10 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

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

12 Transpac dashboard,M., Schoffstall, M. Davin, J. RFC1157. 1990.

13 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 Work-sharing, Miami, FL, October 22, 2014.