

Complementary Visualization of perfSONAR Network Performance Measurements

Andreas Hanemann
German Research Network (DFN)
c/o Leibniz Supercomputing Center
Barer Str. 21, 80333 Munich, Germany
hanemann@dfn.de

Olav Kvittem
Norwegian Research Network
(UNINETT)
Abels gate 5, 7465 Trondheim, Norway
olav.kvittem@uninett.no

Joe Metzger
Energy Sciences Network (ESnet)
Lawrence Berkeley National Laboratory
1 Cyclotron road, Berkeley, CA, USA
metzger@es.net

Vedrin Jeliakov
Institute for Parallel Processing
Bulgarian Academy of Sciences
Acad. G. Bonchev St 25-A, 1113 Sofia, Bulgaria
vedrin@acad.bg

Luís Marta
Fundação para a Computação
Científica Nacional (FCCN)
Av. do Brasil, 101, 1700-066 Lisboa, Portugal
lmarta@fccn.pt

Igor Velimirovic
Croatian Academic and Research Network
(CARNet)
J. Marohnica 5, 10 Zagreb, Croatia
igor.velimirovic@carnet.hr

Abstract

Research backbone networks are currently used by a variety of scientists and research projects interested in maximizing the benefit they receive in using the networks. Therefore, it is highly desirable for them to get access to network performance monitoring data which is prepared to serve their needs.

The perfSONAR measurement framework, which is jointly developed by the EU-funded GN2 JRA1 project, Internet2, and ESnet, has developed a prototype for measuring a variety of network performance metrics such as utilization, one-way delay, one-way delay variation, packet loss, available bandwidth, and others. Corresponding measurements are currently carried out in the contributing research networks in a prototypical manner and will be installed on a regular basis in the future.

For the visualization of measurements a set of tools is under development which is in the focus of this paper. The design of the tools is carried out to serve the interests of the broad user community and uses best practices gained from a survey of visualization possibilities.

1 Introduction

Many individual scientists and research projects make use of networking infrastructures which are provided by national or international research backbone networks and their local institutions. A local network administrator can make use of a variety of tools to monitor the LAN in order to ensure the availability and performance of the network. However, the situation gets much more complicated if information about the performance of a connection involving different administrative domains is requested. In addition to examples like transnational videoconferences, the monitoring of multi-domain connections is especially interesting for Grid projects which are currently being deployed across Europe and elsewhere. In general, a distinction has to be made between measurements on a regular basis (e.g. every five minutes) and on demand measurements for problem examination and performance improvement.

There is an increasing demand from users to get timely access to monitoring data which are preprocessed according to their needs. In addition, network providers would like to be able to customize views to control the information disclosure to users or for protecting the privacy of users against other users. These needs are addressed by user-oriented visualization tools which can be divided into two

major groups. *Centralized visualization tools* use a central server to collect data from different repositories. Usually these tools are easy to use and some refinement of information or functionality can be performed by the centralized server. Examples can be the caching of frequently demanded information, calculation of end-to-end metrics, central enforcement of information access policies. A disadvantage is that the central server is a single point of failure. In contrast, *decentralized visualization tools* allow the user to directly access basic measurement data. These tools allow the user to get a deeper insight into the basic information, but they require more understanding of the underlying measurements. Therefore, these tools are especially suitable for experienced users.

The rest of this paper is organized as follows. In Section 2 the perfSONAR measurement framework is presented which is in the process of measuring a set of metrics within several research backbone networks. Related work concerning measurement projects and their visualization tools is presented in Section 3. Building upon perfSONAR a set of visualization tools is offered to serve the needs of the different perfSONAR user groups. These tools are explained in detail in Section 4 focusing on their innovative features. Conclusion and future work are given in the last section.

2 perfSONAR Measurement Framework

The subproject *Joint Research Activity 1* [7] of the GN2 project aims at providing a framework for performing multi-domain measurements in the European Research Network (GÉANT2) and the connected National Research and Education Networks (NRENs). It is carried out in close cooperation with Internet2 and ESnet and provides a common system called *perfSONAR* (Performance focused Service Oriented Network monitoring ARchitecture) [13]. The name reflects the choice of a Service Oriented Architecture for the system implementation.

The general monitoring framework, which is explained in detail in the following, is depicted in Fig. 1.

The Measurement Points are the lowest layer in the system and are responsible for measuring and storing network metrics as well as for providing basic network information. The measurements can be carried out using active or passive monitoring techniques. The Measurement Point Layer of a domain consists of different monitoring components called Measurement Points where each of them provides information on a specific metric (e.g., one-way delay, jitter, loss, available bandwidth). Each network domain can, in principle, deploy Measurement Points of its choice.

The Service Layer [4] is the middle layer of the system and consists of administrative domains. It allows for the exchange of measurement data and management information between domains. In each domain, a set of entities (ser-

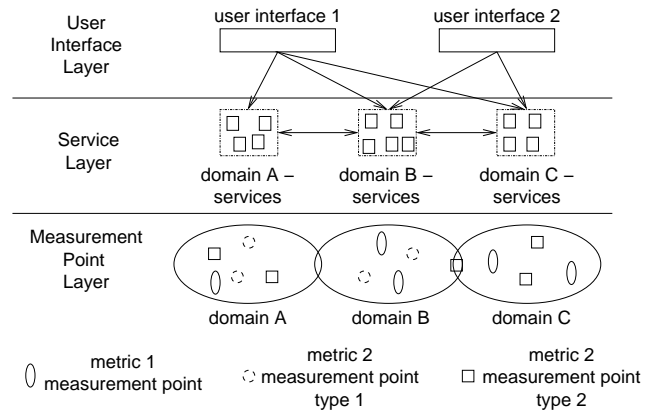


Figure 1. perfSONAR Measurement Framework

vices) is responsible for the domain control. Each of them is in charge of a specific functionality, like authentication and authorization, discovery of the other entities providing specific functionalities, resource management, or measurement of network traffic parameters. In particular, the Measurement Archive Service (MA) is designed as a repository for measurement results. The interaction of the entities inside a domain as well as the access to the Measurement Point Layer or other domains may not be visible to the end user. Some of the entities contain an interface which can be accessed by the User Interface Layer.

The User Interface Layer consists of visualization tools (user interfaces) which adapt the presentation of performance data to be appropriate for the needs of specific user groups. In addition, they may allow users to perform tests using the lower layers of the framework.

Different user groups are distinguished in the project.

End users are individual researchers running applications over the network or are interested in network research. They would like to have easy access to multi-domain network topologies and metrics.

Project members are interested to access data meaningful for their project (i.e. network topology of the project's network and metrics related to the project purpose).

NOC/PERT staff is interested in getting timely fault and performance data as well as in being able to perform on demand tests.

Network managers would like to get an overview about the network performance and to track long term trends useful for network planning and optimization.

3 Related Work

As the focus of this paper is on measurement visualization, related work concerning measurement projects and their visualization tools is presented in this section.

MonALISA The MonALISA (Monitoring Agents using a Large Integrated Services Architecture, [9]) system is designed as an ensemble of autonomous multi-threaded, self-describing agent-based subsystems which are registered as dynamic services, and are able to collaborate and cooperate in performing a wide range of information gathering and processing tasks.

MonALISA represents a variety of network performance metrics, such as round-trip time, bandwidth utilization, packet loss, and delay variation (jitter). These metrics are represented through graphs that are customizable with respect to colors, type of graph (line, scatter, area or combination), value ranges and zoom.

A drawback of the MonALISA is that the system is not designed for real multi-domain environments, but pursues a centralized approach. In addition, the visualization does currently not allow for starting on demand tests.

MRTG The Multi Router Traffic Grapher (MRTG, [11]) is a tool that monitors and graphs the traffic load on network links. Since the start of its development in 1994 it has become the de-facto industry standard for visualizing SNMP interface counters. MRTG is easy to use, highly configurable, and portable. It provides the ability to monitor network device interfaces using SNMP, store the results in a Round Robin Database (RRD), and generate web pages containing complex graphics. MRTG has powerful mechanisms to manage and aggregate data over time. It simplifies and automates the problems of dealing with missing data and handling error conditions. MRTG has been used by thousands of domains across the Internet. Its database and visualization facilities are often used as components in other monitoring tools. Despite of its initial design for SNMP counters, it has been used to measure and visualize many things including all types of network characteristics including latency and bandwidth, computing system performance and even weather, tides, and traffic jams.

SmokePing This tool [12] can measure, store and display the latency distribution in combination with packet loss. An RRD is used for data storage and to draw graphs, giving up to the minute information on the state of each network connection. For every round of measurement, SmokePing sends several packets. It then sorts the different round trip times and selects the median which is drawn on the graph. The other values are drawn as successively lighter shades

of gray in the background (smoke). In addition, packet loss is represented by the color of the median line changing accordingly to the number of lost packets. It is a good representation of network health, making it easy and intuitive to see the fluctuation of RTT and packet loss of the ICMP echo pings that are sent within the network.

DFN IPPM DFN Erlangen developed an IPPM [5] measurement system that is capable of determining delay, jitter and packet loss. That prototype has been improved continuously, and is being used in the networks of DFN and GÉANT2. In active measurement tasks, groups of UDP packets with timestamps are sent throughout the network at regular time intervals, to test jitter and one-way-delay. Each group of packets has a certain configuration of ToS bits, packet size, group size and distance between packets. Comparing the number of packets that were sent with the number that arrived at the destination gives the value of the packet loss metric. The measured metrics are represented as graphs within a web browser.

However, there is a tight link between the visualization and the measurements in MRTG, SmokePing and DFN IPPM which limits the visualization to a set of metrics.

Commercial tools For the management of networks a variety of tools is offered by commercial vendors. For example, HP OpenView allows for the management of enterprise networks and Infovista is suitable for the monitoring of nodes and links including the automatic alerting of operation staff. Mercury Interactive's tools provide an aggregated view of the fulfillment of service level agreements. However, these tools are designed for a network provider-internal use and do not support the customization of views for user purposes.

Summary As a whole, additional efforts are needed for the visualization of multi-domain measurements as several limitations within today's tools have been identified. Limitations exist concerning the multi-domain support. Multiple domains are only supported if the domains agree to deploy common tools and measurements while perfSONAR aims to avoid such limitations. Suitable visualizations of certain metric sets have been identified, but the aim is to be able to flexibly add new metrics. In addition, the customization of views for user needs is often only possible in a quite restrictive manner.

4 Visualization Tools

In this section the visualization tools which are offered for perfSONAR users are presented. These comprise three centralized tools (CNM, Nemo and VisualperfSONAR) as

well as a decentralized tool (perfSONARUI). The tools make use of RRD MAs and network topologies which are currently available for GÉANT2, CARNet, SWITCH, SURFnet, UNINETT, and ESnet as well as additional RRD MAs provided for CESNET, ISTF, PIONEER, and Abilene. The RRD MAs currently provide bandwidth and utilization data. Furthermore, a prototypical version of an IPPM MA [5] exists providing delay, jitter, and packet loss metrics. Table 1 gives an overview of the functionality that the tools provide to the different user groups.

4.1 Customer Network Management

In general, the term Customer Network Management (CNM) means to provide a controlled access to a service provider's management information. This access allows customers to retrieve management information which has been preprocessed according to their needs and to change the service configuration in a predefined manner.

In 1997 DFN started the development of its CNM tool [2] to provide universities and research institutions in Germany information about the national research backbone. Since the start of the JRA1 project in September 2004 DFN has begun to provide and evolve the CNM for its application in perfSONAR. In a prototypical version the tool shows six networks for which topological data is available. The aim is to provide an overview to users about the current and past performance of networks which is in particular useful to decide whether a problem is located in the user's own network or in the backbone.

The tool uses Java applets on the client side which communicate via CORBA with the CNM C++ server. Topology and metrics are stored in a PostgreSQL database, while perl scripts are in place for data retrieval.

Functionality The tool shows the network topology using a set of hierarchical topology maps. The maps contain network nodes (routers or switches) and links including status information and metrics. In the prototypical version the link utilization is the first metric which is displayed, but delay, jitter and loss metrics will be added shortly. It is not only possible to see the current situation in the network, but also to go back to a situation in the past which is enabled by a time navigation function. In addition, graphs can be opened to show the course of metrics over time using statistics on a daily, weekly, or monthly basis.

The CNM tool has the capability to customize views and to limit the access to metrics. This is done using its authentication mechanism. This functionality is useful for projects which would like to use the CNM since it can offer a view onto the network connecting the partners. In addition, this mechanism supports storing user preferences (e.g. preferred maps, preferred metrics).

Specific Developments Besides the integration of additional network topologies and metrics, the tool development aims at the development of special overview capabilities for network management. This *dashboard* will show a set of aggregated metric values. For example, the network manager might want to know the current overall availability of the network, and how this metric has evolved over a past time period. Another interesting question is whether there are problems with the network delay in general. Therefore, it is desirable to calculate an overall delay value from the single delay values.

4.2 Network Monitor System

The Network Monitor System (Nemo) [10] is a collection of tools that can visualize a network with metrics on a geographical map. It has similarities with the CNM, but focuses on an easy-to-use GUI for end users. An important feature is the automated construction of maps from given geographical coordinates which are combined with background maps. Even though geographical information is not of interest for network management per se, people's customization to geographical maps allows for an easy orientation. The first web-based map of UNINETT was presented around 1997.

The main purpose of Nemo was to provide visualizations to the customers showing how they are connected to the network and to provide them with statistics on their QoS. A secondary goal is to have a visual representation of the inventory database for data quality inspection. The system is in use by UNINETT and SUNET in Sweden.

The Nemo architecture is a loosely coupled collection of programs in Sh, Perl and Java. A topology database feeds the map planning, which is used to extract maps from the map database. The combination of maps with URLs being attached to network elements results in clickable maps.

Functionality The system features self-zooming and flexible connections to external web-resources for network components and other maps. Each entity on a map like a node, link or reference to another map is clickable and has extendable menus. Information that can be related to an element name on the map, can be presented on the map using color or size of nodes and links. Each type of node can have its own icon.

The system imports statistical information from various network measurement systems and stores it with resolution from a few minutes to a week. Statistics are provided for any past time period using the current network topology.

There are multiple clients of the system featuring both client side processing with Java applications and server side processing with Web/Javascript presentation.

	End user	Project member	PERT/NOC staff	Network manager
CNM	Customized view of network topology and metrics	Customized view of project network and metrics	View of complete network topology and additional metrics	Customized view of network topology and metrics and planned dashboard
Nemo	Easy to use unrestricted topology and metric visualization	General topology and metric visualization	General topology and metric visualization including link to on demand testing	Intuitive network topology and performance overview
VisualperfSONAR	View of general traffic statistics for IP path and geographical path representation	View of detailed traffic statistics for IP path and geographical path representation	View of detailed traffic statistics for IP path and geographical path representation	View of general traffic statistics for IP path and geographical path representation
perfSONARUI	Seamless access to all perfSONAR services	Seamless access to all perfSONAR services	Seamless access to all perfSONAR services	Seamless access to all perfSONAR services

Table 1. Features of the visualization tools (focus concerning group interests highlighted in bold)

The system is also used to present clickable maps for web servers in the academic sector, measurement equipment locations and eduoam sites.

Specific Developments The workplan foresees the development of Nemo interfaces for perfSONAR access, starting with collecting measurements from MAs and the construction of maps using geographical coordinates.

4.3 VisualperfSONAR

Following the first prototypical traceroute visualization tool by ESnet, CARNet provided a first prototype of its VisualperfSONAR [15] tool in January 2006. The former ESnet prototype was capable of parsing through the output of a traceroute command and displaying a list of interface utilization graphs for some of the ESnet, Abilene and GÉANT2 router interfaces found in the output using the MAs for these networks. In contrast to CNM and Nemo this tool is focused on tracking the performance on single paths within the networks which is interesting to examine problems in more detail.

VisualperfSONAR is designed as an application which enables its user to view the statistics for all the interfaces along a given path (most often the output of a traceroute command) for which data is stored in one of the implemented MAs. Also, a graphical representation of the path over the geographical map is provided. The back end uses Web services to retrieve measurement data from multiple MA locations.

Functionality The tool is implemented as a web based application. It enables users to upload a specific IP path, in a form of a traceroute output. The application then parses the traceroute output and queries MAs for each interface (in this case IP address of the interface) which is identified in the output. The list of operational MAs and list of interfaces for which a specific MA contains the data are stored in tool configuration files.

The tool retrieves data using Web services to query the MAs, and presents the data to the user as a graph. Also, the tool maps each of the interfaces to the IP address and draws a virtual path over geographical maps. At the time of this writing, there is no standardized solution of retrieving geographical coordinates. Therefore, the tool currently chooses random geographical coordinates. Google's map service [3] is used for presenting a geographical background map.

Specific Developments In order to provide better functionality and user perspective, some modifications are going to be implemented in the following months. Users will be able to choose timescales and select different metrics for visualizations. They will also be able set several graphing preferences.

4.4 perfSONARUI

The development of perfSONARUI [14] began in Q4 of 2005, triggered by the explicit need to provide an open source, publicly accessible, easy to use and yet powerful stand alone graphical user interface client, capable of querying all available perfSONAR services and displaying the returned data in a concise manner. In contrast to the other

tools the use of perfSONARUI requires an understanding of the perfSONAR architecture as it is in particular useful for problem analysis and troubleshooting purposes. At the time of this writing perfSONARUI supports all ten deployed RRD MA services and is prepared to support the IPPM MA Service.

A Java runtime environment on the client system is required for perfSONARUI which makes it platform independent. perfSONARUI uses the Apache Axis SOAP implementation [1] to query perfSONAR services. Its visualization is based on JFreeChart [6], an open source Java library for generating charts.

Functionality At present perfSONARUI is able to query one or several RRD MAs and visualize the returned data. Users can apply filters (a list of IPv4/IPv6 interface addresses or arbitrary traceroute output) in order to search for a particular subset of interface utilization data. The visualization is done both in tabular and graphical form. Users can select different criteria for sorting the data. The tool provides a quick overview of the interface utilization vs. capacity in several user-selected domains simultaneously, as well as detailed views for any particular interface. Time intervals for summary and detailed views are selectable through radio buttons and sliders.

Specific Developments The perfSONARUI development team is keen to implement support for every newly developed, deployed, and publicly accessible perfSONAR service in a timely manner. Plans for the next few months especially foresee implementing perfSONARUI support for BWCTL MP services.

5 Conclusion and Future Work

Based on the multi-domain measurement framework perfSONAR a set of visualization tools is provided to serve the needs of different user groups being summarized in Table 1. For end users, easy to use network maps are provided by Nemo and CNM and a network path will be easily traceable by VisualperfSONAR. A specific offer for project members will be provided by the CNM since it allows for the customization of maps. For troubleshooting VisualperfSONAR and perfSONARUI are very useful to investigate the behavior of a data transfer along a path. Nemo is also addressing this functionality by linking its network elements with test tools which is easily extensible. For network managers the CNM will provide a dashboard to aggregate data within a network domain.

There are some developments expected in the next few months that will impact all of the perfSONAR visualization tools.

A Topology Service is going to be deployed which should serve as a common basis for topological information. It will include geographical coordinates which are needed in particular for Nemo and VisualperfSONAR. Therefore, interfaces will be written to retrieve these coordinates and other topological information from the Topology Services.

There is currently a JRA1 [7] and JRA5 [8] collaboration to design and deploy an Authentication and Authorization Service within perfSONAR. It will be used to provide handles to users according to their data access rights. Once this service is available the visualization tools will be extended to interact with it and to reflect the defined data access policies.

Lookup Services will be deployed which allow for the automated discovery of service locations and capabilities across domains. This will provide new capabilities to all the visualization tools. For example, the perfSONARUI will be able to dynamically find all active MAs.

Acknowledgement Portions of this work are supported by the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

References

- [1] Apache Axis. <http://ws.apache.org/axis>.
- [2] CNM project homepage. <http://www.cnm.dfn.de/>.
- [3] Google map service. <http://maps.google.com/>.
- [4] A. Hanemann, J. Boote, E. Boyd, J. Durand, L. Kudarimoti, R. Lapacz, M. Swany, S. Trocha, and J. Zurawski. Perfsonar: A service oriented architecture for multi-domain network monitoring. In *Service-Oriented Computing - ICSOC 2005, LNCS 3826, Springer Verlag*, pages 241–254, Amsterdam, The Netherlands, December 2005.
- [5] IPPM measurement data. <http://www.win-labor.dfn.de/cgi-bin/ipqos/select-new.pl?config=geant>.
- [6] Jfreechart. <http://www.jfree.org/jfreechart>.
- [7] Joint research activity 1 (jra1): Performance measurement and monitoring. <http://www.geant2.net/server/show/nav.754>.
- [8] Joint research activity 5 (jra5): Roaming and authorisation. <http://www.geant2.net/server/show/nav.758>.
- [9] MONitoring Agents using a Large Integrated Services Architecture (MonALISA), California Institute of Technology. <http://monalisa.caltech.edu/>.
- [10] Netmonitor (Nemo). <http://drift.uninett.no/kart/nemo>.
- [11] T. Oetiker. Multi Router Traffic Grapher (MRTG). <http://people.ee.ethz.ch/~oetiker/webtools/mrtg/>.
- [12] T. Oetiker. SmokePing homepage. <http://people.ee.ethz.ch/~oetiker/webtools/smokeping/>.
- [13] perfSONAR (Performance focused Service Oriented Network monitoring ARchitecture). <http://www.perfsonar.net/>.
- [14] perfSONARUI homepage. <http://www.perfsonar.net/jra1-wiki/index.php/PerfsonarUI>.
- [15] VisualperfSONAR web interface. https://noc-mon.srce.hr/visual_perf/.