Measurement Infrastructures for Future Internet Testbeds

José Augusto Suruagy Monteiro

Núcleo de Pesquisas em Redes e Computação (NUPERC)
Universidade Salvador (UNIFACS)
R. Ponciano de Oliveira, 126 – Rio Vermelho – 41950-275 Salvador - Brazil

suruagy@unifacs.br

Abstract. *In this paper we review the proposed measurement infrastructures for Future Internet testbeds, including the GENI Instrumentation and Measurement System (GIMS), and PlanetLab Europe's TopHat measurement service. We conclude proposing an initial set of monitoring resources in order to open a discussion aiming at defining the measurement facility for our Future Internet subproject testbed.*

Resumo. *Neste artigo são revistas as infraestruturas de medições propostas para as redes experimentais de testes para a Futura Internet, incluindo o Sistema de Instrumentação e Medições da GENI (GIMS) e o serviço de medições TopHat do PlanetLab Europa. O artigo conclui propondo um conjunto inicial de recursos de monitoração de modo a abrir uma discussão visando definir os recursos de medções para a rede experimental do nosso subprojeto sobre a Futura Internet.*

1. Introduction

Even in today's technology, end to end performance may not be what users would expect from high speed backbones. In order to identify possible performance holes, it is desirable to instrument such networks, from the user's lab, through his campus, regional, national and international backbones with tools that can perform tests in order to obtain delay and throughput measurements.

In testbeds for future Internet experiments such as GENI, Federica, OneLab2, G-Lab, and AKARI, where each experiment runs on a given slice, it is even more important to instrument the testbed in order to collect all relevant data for each experiment, including not only network related measurements but also operational data which can help operations personnel on monitoring and troubleshooting the infrastructure itself. For more information on these testbeds, please look at a companion paper [Abelém et al., 2010].

GENI (Global Environment for Network Innovations) is a network research and engineering infrastructure which is currently being designed. GENI key concepts are the following: programmability, virtualization, federation, and slice-based experimentation [GENI 2009b]. The GENI development started in the so called Spiral 1, with 4 control frameworks and corresponding prototype clusters: PlanetLab, ProtoGENI, ORCA, and ORBIT. One of Spiral 2 objectives is to improve the integration of GENI prototypes as well as the design and prototyping of Instrumentation and Measurement.

In Europe the OneLab2 project, which is related to PlanetLab Europe, developed the so called TopHat measurement service [Borgeau et al. 2010].

As far as network measurement infrastructures are concerned, we emphasize perfSONAR, and ETOMIC.

perfSONAR [Hanemann et al. 2005] is a service oriented architecture for multi-domain network monitoring. perfSONAR have been deployed in several National Research and Education Networks (NRENs), including Internet2, ESNet, Géant, and RNP, through the MonIPÊ service [Sampaio et al. 2007]. perfSONAR is already being used in the ProtoGENI infrastructure, one of GENI's clusters, to monitor Internet2 backbone and other parts of it.

ETOMIC [Csabai 2010] is a network traffic measurement platform which was originally launched in 2004, and has been recently renewed as part of the OneLab project, and is being used by the TopHat measurement service, mentioned before.

This paper is inserted in the "Web Infrastructure" layer of the Brazilian Institute for Web Science Research. It describes preliminary work in the first activity for the 3rd Goal of the "Future Internet Architectures" which aims at designing and implementing a measurement infrastructure for the testbed proposed in the subproject first goal.

In this paper we review the proposed measurement infrastructures for Future Internet testbeds, including the GENI Instrumentation and Measurement System (GIMS) [Barford et al. 2009], and PlanetLab Europe's TopHat measurement service [Borgeau et al. 2010] and propose an initial set of monitoring resources in order to open a discussion aiming at defining the measurement facility for our project testbed.

This position paper is organized as follows. In Section 2 we present an overview of the FI testbeds evaluated, starting with the Slice-Based Facility Architecture (SFA) that was designed to federate them. In Section 3 we present GENI's Instrumentation and Measurement System (GIMS) as a general framework for Future Internet Testbeds Monitoring. In Section 4 we present an overview of relevant monitoring projects. Finally, in Section 5 we conclude proposing an initial set of monitoring resources for our project testbed.

2. Future Internet Testbeds

In this section we present an overview of the FI testbeds currently evaluated: OneLab, Emulab, Vini, Enterprise GENI (E-GENI), and Openflow. Starting with a subsection on the Slice-Based Facility Architecture (SFA) designed to allow the interoperation of a federation of sliced-based network substrates.

2.1. Slice-Based Facility Architecture (SFA)

Basically all Future Internet (FI) testbeds, rely on slice-based virtualization and for their integration it is interesting to achieve some sort of federation among them.

Slice-Based Facility Architecture, or simply SFA, aims at defining a minimal set of interfaces and data types that permit the interoperation of a federation of slice-based network substrates [Peterson et al. 2009]. The specification was designed to support federation among facilities such as PlanetLab, Emulab, VINI, and GENI.

SFA has two main abstractions: *components* and *slices*. A component is a resource such as an edge computer, a customizable router, or a programmable access point. While, a slice is a substrate-wide network of computing and communication resources capable of running an experiment or a wide-area network service.

SFA is composed by several building blocks. A Registry (R) is used to register users, slices, components, and authorities. A user interfaces a Slice Manager (SM) to create and control its slice. An Aggregate Manager (AM) is responsible for a set of components. Finally, a Component Manager (CM) controls a component.

2.2. OneLab

OneLab is a project based on PlanetLab Europe which provides an open, general purpose, shared experimental facility for experimentations on future Internet.

2.3. EMULAB

Emulab is a network testbed, giving researchers a wide range of environments in which to develop, debug, and evaluate their systems. The name Emulab refers both to a facility and to a software system.

2.4. VINI

VINI is a virtual network infrastructure that allows network researchers to evaluate their protocols and services in the wide area. VINI allows researchers to deploy and evaluate their ideas with real routing software, traffic loads, and network events. To provide researchers flexibility in designing their experiments, VINI supports simultaneous experiments with arbitrary network topologies on a shared physical infrastructure.

VINI is related to the PlanetLab project, although its testbed is distinct from the public PlaneLab. It uses MyPLC to manage the nodes. The so called PL-VINI [Bavier et al. 2006] forwards packets in the user space significantly limiting scalability.

Trellis [Bhatia et al. 2008] is a set of extensions to the PlanetLab kernel and tools which tries to balance speed, flexibility, and isolation using commodity hardware and software.

All PlanetLab users can now access the VINI testbed, and create virtual networks within their slices, via the SFA federation interface. I.e., VINI has an Aggregation Manager which controls its resources and interfaces with PlanetLab Slice Manager.

2.5. OpenFlow

OpenFlow [McKeown et al. 2008] is an open protocol to program the flow-table in different switches and routers. Using OF, a network administrator can divide traffic in production and research flows. Research flows can be controlled by a separate controller which adds and removes flow entries from the Flow Tables.

FlowVisor is a special purpose controller which slices the OpenFlow datapath allowing different flows be controlled by different controllers, i.e., it acts as an OF proxy [Sherwood et al. 2009].

SNMPVisor [Yap et al. 2009], in the other hand, is used to mediate device configuration access among experiments. It watches SNMP control messages and direct them to the correct datapath element, acting as a SNMP proxy.

2.6. Enterprise GENI

The Enterprise GENI (E-GENI)[1] is a project which aims at demonstrating how GENI experiments can be a part of local networks while simultaneously running production traffic. Its main entities are: an OF-based substrate, FlowVisor, and an Aggregate Manager.

The OpenFlow-based Substrate is composed by the switches that communicate using the OpenFlow protocol to a controller application. In E-GENI the OF-based substrate communicates with FlowVisor that isolates control traffic for the production traffic from the control of individual experiments traffic. Finally, the Aggregate Manager is an application which mediates between the Clearinghouse[2] and the E-GENI FlowVisor, over a SOAP-based SFA protocol.

3. Monitoring

In this section we present GENI Instrumentation and Measurement System (GIMS) as a general framework for Future Internet Testbeds Monitoring. In Section 3.1 we present GIMS General Requirements, followed by GIMS General Specification, not yet implemented, in Section 3.2, and the first definitions for its architecture, in Section 3.3.

3.1. GIMS General Requirements

As we mentioned before, GIMS General Requirements [Barford et al. 2009] is good set of requirements which should be met by any FI instrumentation and Measurement facilities. They are:

- Ubiquitous deployment,
- No (or at least measurable) impact on experiments,
- Extensibility (i.e., the ability to add new instrumentation and/or new measurement synthesis capability),
- High availability (at least as available as GENI systems on which experiments are conducted),
- Large capacity (i.e., the ability to support a diverse set of simultaneous activities from a large number of experiments),
- The ability to measure detailed activity with high accuracy and precision from physical layer through application layer (includes the ability to calibrate measurements),
- The ability to specify required measurements for an experiment in a slice (using either standard measurement types from a library or defining user specific

[1] http://www.openflowswitch.org/wk/index.php/E-GENI

[2] GENI Clearinghouse includes the SM and R functionalities.

measurements) and then having these measurements initialized in the infrastructure when an experiment is activated,

- Access control (i.e., the ability to specify what data is available from a particular device or collection of devices, to whom, and for how long),
- A large, secure central repository in which collected data can be anonymized and made available to users,
- A “data analysis warehouse” where tools for visualizing, interpreting and reporting measurement data can be developed and made openly available.

3.2. GIMS Specification

GIMS specification is organized into three general concepts: instrumentation, measurement synthesis, and analysis and archiving with security and access control.

Instrumentation is composed by physical taps in GENI systems. They basically identified three types of sensors: link sensors, node sensors, and time sensors. Specific link sensors depend on link technology: it could be an optical tap, or a switch port mirroring. Node sensors can be implemented through the /proc file system that is available in UNIX or through SNMP MIBs. Finally, Time sensors can be implemented using GPS receivers.

Synthesis corresponds to the set of systems that collect and synthesize the signals provided by the sensors into meaningful data. It includes the ability to frame and capture packets (e.g., tcpdump, Endace DAG cards). It can also Flow-export measurements for specified packet streams. They considered as a data gathering protocol either SNMP or the Simple Common Sensor Interface for PlanetLab that has a data query protocol based on HTTP. Besides it includes the task of applying timestamps to experiments.

The **Data archive** can be decomposed in two primary components: a repository and a data catalog. The repository should reside on high capacity storage systems, while the data catalog indexes the data in the repository. Their objectives are to provide a secure and reliable repository to users measurement data, and to provide a standard interface for data extraction, so that users can analyze their data at their own sites.

The general requirements for **security and access control** are:

- Sensors, collection/synthesis and archival systems must only be accessible by authorized users,
- Authorized users must only be allowed access to a specified proportion of the resources on the measurement systems,
- Different views of the same data will be required depending on the approved level of authorization.
- All of the measurement systems themselves must be secured against attack,
- No aspects of the data that is collected will compromise or violate the privacy of individuals or organizations that source/sink the data.

3.3. GENI Instrumentation and Measurement Architecture

The GENI Instrumentation and Measurement Architecture is still in an early draft stage which contains basically the document skeleton as well as the definition of a number of agreed upon key services [GENI 2010]. They are: Measurement Orchestration (MO) service, Measurement Point (MP) service, Measurement Collection (MC) service, Measurement Analysis and Presentation (MAP) service, and Measurement Data Archive (MDA) service. Besides these services, the document identifies the need for interfaces, protocols, and schemas for measurement data, as well as interfaces, protocols, and API's for using the I&M services.

4. Monitoring Projects

In this Section we present an overview of some monitoring projects included in GENI's Spiral 2, followed by OneLab's TopHat measurement component, and OpenFlow Measurements.

4.1. GENI Spiral 2 Projects

In GENI Spiral 2 there are a number of projects related to the Instrumentation and Measurement Design and Prototyping[3]. In this paper we will concentrate in three of them which were originally proposed in the scope of the ProtoGENI Cluster: Measurement System, LAMP, and INSTOOLS.

4.1.1. Measurement System

This project which is described in [GENI 2009a] is a first step on the prototyping of an instrumentation and measurement system for GENI. It focus on *passive packet capture*, and in particular in IPv4 packets, with support for no-loss packet capture at a line rate of 1Gbps, and the support for 256 simultaneous active experiments per node.

For the experiments they use a programmable network interface card (NIC) to capture packets from an optical tap. For the data repository they are proposing to use the flexible high capacity storage services offered by Amazon (S3 service) and others.

The system will provide the following transformation functionalities: sampling, filtering, anonymization, and aggregation through flow records, and packet counts.

The GIMS will include an Aggregate Manager in order to control the instrumentation plane for its users.

4.1.2. LAMP

LAMP stands for "Leveraging and Abstracting Measurements with perfSONAR"[4]. perfSONAR [Hanemann et al. 2005] is a multi-domain performance monitoring framework, which defines a set of protocols standards for sharing data between measurement and monitoring systems.

---

[3] http://groups.geni.net/geni/wiki/SpiralTwo

[4] http://groups.geni.net/geni/wiki/LAMP

perfSONAR is a middleware which interfaces with data collection tools and provides a number of web services to analysis and visualization tools. Its services are grouped into data services and infrastructure services. The former includes Measurement Points, Measurement Archives (MA), and Transformations. While the later includes a group of Information Services: Service Lookup, Topology, and Service Configuration; as well as Authentication and Authorization Services.

Leveraging on these services, the LAMP objectives include: collaborate on defining a common but extensible format for data storage and exchange (using OGF's Network Measurement WG schema used by perfSONAR as a starting point); develop a representation of GENI topology to be used to describe measurements and experiment configuration; collaborate in the design of a common GENI I&M architecture.

Due to its multi-domain nature, the federated GENI is a perfect match to perfSONAR. Furthermore, existing perfSONAR SNMP MA can be used to query ProtoGENI switches on Internet2. To support experiments perfSONAR will have to evolve, but this is not a problem due to its modular and extensible architecture.

4.1.3. INSTOOLS

INSTOOLS is the Instrumentation Tools project[5] at the University of Kentucky. It automatically deploys, configures, and enables slice-specific measurement infrastructure, and makes measurement results available via an easy to use interface.

Its architectural components can be mapped to Capture, Collection, Storage, Control, Processing, Access, and Presentation functionalities, facilitating the mapping to GENI's I&M architecture.

Furthermore, they have proposed to begin moving towards the perfSONAR formats.

4.2. TopHat

TopHat is PlanetLab Europe's active measurement component and provides PL applications with a topology monitoring service for the entire lifecycle of an experiment [Bourgeau et al. 2010]. Therefore, extending MySlice/Sys services which provide useful historical information about PL nodes. Using both services, users can better assign their slices to the nodes that better fit the experiment requirements, including connectivity and delays.

TopHat get information from several services such as DIMES[6], ETOMIC [Csabai et al. 2010], SONoMA[7], and Team Cymru[8], and is considering also to access perfSONAR's data. However, it also uses a dedicated measurement infrastructure, named TDMI (TopHat Dedicated Measurement Infrastructure), which runs active measurements including Paris Traceroute.

---

[5] http://groups.geni.net/geni/wiki/InstrumentationTools

[6] http://www.netdimes.org/ Studies the structure and the topology of the Internet.

[7] http://complex.elte.hu/sonoma/

[8] Provides an IP to AS mapping service. http://www.team-cymru.org/

Regarding the integration of TopHat and perfSONAR, while TopHat focus on experiment support, perfSONAR original focus was on troubleshooting across network boundaries. However, with projects like LAMP, it is being considered to focus also on slice monitoring. On the other hand, as stated in [Bourgeau et al. 2010] their interest on perfSONAR is to have access to router information such as queue lengths and packet drop rates.

4.3. OpenFlow Measurements

OpenFlow measurements can be achieved basically in two ways: through SNMP queries or via event logs exported to a controller. For instance, with NOX's events logs it is possible to record several network events including topology changes and link state.

5. Proposition

Based on this revision of related work, we propose the following technologies and measurement facility for our project testbed.

First of all, we advocate that currently the best combination for our testbed is a private instance of PlanetLab for node virtualization, and the OpenFlow based E-GENI for virtualizing the network resources.

For this testbed monitoring facility we propose to monitor the OpenFlow substrate via SNMP with the use of the SNMPVisor, and trough the access of the controller event logs.

Furthermore, we propose to closely follow the LAMP project, and to study and evaluate the detailing of GENI's Instrumentation and Measurement Architecture.

References

Abelém, A., Machado, I., Stanton, M., and Carvalho, T. (2010) "Design of a testbed for R&D in network architectures". Submitted to the First Workshop of the Brazilian Institute for Web Science Research. Rio de Janeiro, 2010.

Barford, P., Blodgett, M., Crovella, M., and Sommers, J. (2009) "Requirements and Specifications for the Instrumentation and Measurement Systems for GENI: Version 0.2", May 2009. http://groups.geni.net/geni/attachment/wiki/MeasurementSystem/MeasurementSysSpec.pdf.

Bavier, A. et al. (2006), "In VINI Veritas: Realistic and controlled network experimentation", In: Proc. ACM SIGCOMM, September 2006.

Bhatia, S., Motiwala, M., Muhlbauer, W., Mundada, Y., Valancius, V., Bavier, A., Feamster, N., Peterson, L., and Rexford, J. (2008) "Trellis: A Platform for Building Flexible, Fast Virtual Networks on Commodity Hardware", In: Proc. 3rd International Workshop on Real Overlays & Distributed Systems (ROADS'08), December 2008.

Bourgeau, T., Augé, J., Friedman, T. (2010) "TopHat: supporting experiments through measurement infrastructure federation". In Proceedings of the 6th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom 2010).

Csabai, I., Fekete, A., Hága, P., Hullár, B., Kurucz, G., Laki, S., Mátray, P., Stéger, J., Vattay, G., Espina, F., Garcia-Jimenez, S., Izal, M., Magaña, E., Morató, D., Aracil, J., Gómez, F., Gonzalez, I, López-Buedo, S., Moreno, V., Ramos, J. (2010) "ETOMIC Advanced Network Monitoring System for Future Internet Experimentation". In Proceedings of the 6th International ICST Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom 2010).

GENI (2006). "GENI Instrumentation and Measurement Systems (GIMS) Specification". Document ID: GDD-06-12. Available at http://groups.geni.net/geni/attachment/wiki/OldGPGDesignDocuments/GDD-06-12.pdf

GENI (2009a). "Requirements and Specifications for the Instrumentation and Measurement Systems for GENI". Available at http://groups.geni.net/geni/attachment/wiki/MeasurementSystem/MeasurementSysSpec.pdf

GENI (2009b). "System Requirements Document". Document ID: GENI-SE-SY-RW-02.0. Available at http://groups.geni.net/geni/attachment/wiki/SysReqDoc/GENI-SE-SY-RQ-02.0.pdf

GENI (2010). "GENI Instrumentation and Measurement Architecture". Document ID: GENI-SE-IM-ARCH-0.1. March 8, 2010. Available at http://groups.geni.net/geni/attachment/wiki/GeniInstrumentationandMeasurementsArchitecture/030810%20%20GENI-SE-IM-ARCH-0.1.pdf

Hanemann, A., Boote, J. W., Boyd, E., Durand, J., Kudarimoti, L., Lapacz, R., Swany, M., Trocha, S., and Zurawski , J. (2005). "Perfsonar: A service oriented architecture for multi-domain network monitoring". In Proceedings of the Third International Conference on Service Oriented Computing (ICSOC 2005). ACM Sigsoft and Sigweb.

McKeown, N. et al. (2008), "OpenFlow: Enabling Innovation in Campus Networks", In: ACM SIGCOMM Computer Communication Review, Volume 38, Number 2, April 2008, p. 69-74.

Peterson, L., Sevinc, S., Lepreau, J., Ricci, R., Wroclawski, K., Faber, T., Schwab, S., Baker, S. (2009). "Slice-Based Facility Architecture. Draft Version 1.04". April 7, 2009

Sampaio, L., Koga, I., Costa, R., Monteiro, H., Monteiro, J. A. S., Vetter, F., Fernandes, G., Vetter, M. (2007). "Implementing and Deploying Network Monitoring Service Oriented Architectures - Brazilian National Education and Research Network Measurement Experiments (Invited Paper)". In: 5th Latin American Network Operations and Management Symposium (LANOMS), Petrópolis, RJ, p. 28-37.

Sherwood, R., Gibb, G., Yap, K., Appenzeller, G., Casado, M., McKeown, N., Parulkar, G. (2009). "FlowVisor: A Network Virtualization Layer". OpenFlow Technical Report. Available at http://www.openflowswitch.org/downloads/technicalreports/openflow-tr-2009-1-flowvisor.pdf

Yap, K., Sherwood, R., Kobayashi, M., Handigol, N., Huang, T., Chan, M., McKeown, N., Parulkar, G. (2009). "Blueprint for Introducing Innovation into the Wireless Networks we use every day". OpenFlow Technical Report. Available at http://openflowswitch.org/downloads/technicalreports/openflow-tr-2009-3-openflow-wireless.pdf