G E N I

Global Environment for Network Innovations

Milestone 1
GENI Requirements for Real-Time Measurements

Document ID: GENI-MS1-ERM-March09-v1.1

February 27, 2009

Prepared by:
F. Fidler, C.P. Lai, and K. Bergman
Dept. of Electrical Engineering, Columbia University New York,
500 W. 120th Street ,
New York City, NY 10027
under Project Nr. 1631
“Embedding real-time measurements for cross-layer communications”

TABLE OF CONTENTS

1 Document Scope

This section describes this document's purpose, its context within the overall GENI project, the set of related documents, and this document's revision history.

1.1 Executive Summary

This technical note presents the results obtained in work package "Milestone 1: GENI requirements for real-time measurements" of Project Nr. 1631, "Embedding real-time substrate measurements for cross-layer communications".

The purpose of this milestone is to assess and evaluate GENI requirements for real-time user access to measurement data across a diverse set of technologies. In Section 2 we assess the capabilities of GENI's (future) infrastructure with respect to real-time measurements based on the current version (01/05/2009) of GENI's "Spiral 1 substrate catalog" [geni09_1], on personal communication with other working groups, as well as on information from equipment vendors and the internet. Section 3 deals with the requirements for real-time user access to measurement data across a diverse set of technologies. It is found that a unified measurement framework interfacing the substrate, the control plane, and the experimenter would be advantageous to limit the overhead and complexity both in hardware and software.

1.2 Related Documents

The following documents are related to this document, and provide background information, requirements, etc., that are important for this document.

1.2.1 GENI Documents

Document ID	Document Title and Issue Date
GENI_QR_ERM_Dec08	4Q08 Status Report
GENI-INF-PRO-S1-CAT-01.3	GENI Infrastructure Substrate Catalogue

1.3 Document Revision History

The following table provides the revision history for this document, summarizing the date at which it was revised, who revised it, and a brief summary of the changes. This list is maintained in chronological order so the earliest version comes first in the list.

Revision	Date	Revised By	Summary of Changes
1.0	27 February 09	F. Fidler	Initial draft
1.1	03 March 09	J. Jacob, F. Fidler	Incorporation of editorial notes from J. Jacob

2 GENI Infrastructure for Embedded (Real-Time) Measurements

To evaluate the GENI requirements for real-time user access to measurement data, we assessed the capabilities of GENI's (future) infrastructure with respect to real-time measurements. Our survey is based on the current version of GENI's "Spiral 1 substrate catalog" [geni09_1], on personal communication with different working groups, as well as on information from equipment vendors and the internet.

While information about the main equipment available in each GENI node was derived from [geni09_1], information about the equipment's real-time measurement capabilities had to be gained directly from the equipment vendors, or to some extent from the GENI working groups which control the nodes. General information about some measurement capabilities could be readily derived from publicly available product datasheets, vendor webpages, and personal communication with different working groups. We found that for integration purposes, it is absolutely necessary to work directly with several vendors to gain more specific information about how physical layer performance indicators are accessible on individual node equipment such as routers or switches.

Table 1 shows an overview of the equipment available in the nodes of networks used by different clusters. Based on this information, a survey was performed to reveal the real-time measurement capabilities of different nodes (cf. Table 2).

	Cluster A	*Cluster B*			*Cluster C*	*Cluster D*	
	TIED	*MAX*	*GpENI*	*SSP Overlay*	*Proto-GENI*	*DOME*	*BEN*
Project Nr.	*1609*	*1658*	*1595*	*1578*	*1579*	*1599*	*1582*
HP ProCurve 5400 switch					**X**		
Adva Optical Networking switch		**X**	**X**				
Ekinops 360 transport platform			**X**				
GPS device						**X**	
Juniper EX3200 switch							**X**
NetFPGA				**X**	**X**		
Polatis Fiber Switch							**X**
Ciena CN4200 switch			**X**				
Ciena CoreDirector switch			**X**				
Cisco 6509-E switch	**X**						**X**
Cisco 6513 switch	**X**						
Infinera DTN ROADM					**X**		**X**

Table 1: Node equipment in candidate GENI networks as given in [geni08_1, geni09_1].

Infinera DTN ROADM	*Adva Optical Networking transponder*	*Ekinops transport platform*	*Polatis fiber switch*	*Ciena CN4200 switch*	*GPS device*	*Juniper EX3200 switch*	*NetFPGA*
Bit error counter (FEC)	Bit error counter (FEC), optical power monitor	Bit error counter (FEC)	Optical power monitor, (optical power control)	Bit error counter (FEC), (optical power control)	Interpreted and raw position data	Detection of link brakes	Freely programmable FPGA
[inf08_1, geni09_1]	[adv09_1, geni09_1]	[eki09_1]	[polatis08_1, geni09_1]	[ciena08_1, geni09_1]	[geni09_1]	[jun09_1, geni09_1]	[netf08_1, geni09_1]

Table 2: Real-time measurement capabilities of GENI node equipment according to vendor's webpages, datasheets and information provided in [geni08_1, geni09_1].

From Table 2, the most promising candidates to be applied in real-time measurement based cross-layer experiments are:

- **Infinera DTN:** The *Infinera DTN* is a remotely configurable optical add/drop multiplexer (ROADM) employed in a majority of nodes within the GENI network infrastructure (BEN, Proto-GENI, Internet2, ...). One key advantage is its ability to provide bit-transparent services with digital performance monitoring [inf08_1]. The DTN's PIC technology allows cost-effective O-E-O conversion of the optical signals as they transit the network, allowing bit error ratio (BER) monitoring from forward error correction (FEC) and performance monitoring of SONET/SDH and G.709 overhead at every digital node. The *Infinera Digital Network Administrator* (DNA), a comprehensive integrated element and network management system, provides users with a graphical interface with performance management capabilities to access the BER information. Performance monitoring information can also be exported by means of *.csv data files.
- **Adva Optical Networking:** The *Adva Optical Networking layer 1 equipment* (FSP3000) provides pre-FEC bit error counters for its 10G transponders [morkel08_1, geni09_1]. A Reed Solomon (255,239) code is used for forward error correction yielding a 6.2 dB net coding gain at 10^{-15} BER. To achieve further performance enhancement, an adaptive receiver threshold algorithm is implemented in the module. Feedback from the integrated FEC ASIC to adjust the receiver threshold decision level is used to minimize total pre-FEC errors with improvements up to two decades in pre-FEC error performance. Pre-FEC bit error ratio can be measured.
- **Polatis 24port fiber switch:** The *Polatis switch* is employed within BEN [geni09_1] and offers the possibility to remotely monitor the optical power on a per port basis [polatis08_1]. It also can be used as a variable optical attenuator, which could be employed to compensate for power transients in a cross-layer based fashion. The switches come with a programmable GUI control interface and/or with a visual touch screen display to observe real-time optical power across all ports/selected ports, set thresholds for visual alarming, or activate automatic switchover functionality based upon preconditioned thresholds. For remote power monitoring, the interfacing can be done via Ethernet, serial interface (RS232), or GPIB.

- **Ciena CN4200 switch:** The installation of the *Ciena CN4200* is planned within networks used by Cluster B [geni09_1]. Transponder modules (e.g. the F10-T_10G or FC4-T module) already incorporating FEC devices with the possibility of accessing pre-FEC BER are available [ciena08_1] but not yet installed. A ROADM module available for the CN4200 platform allows for automatic power equalization (power level control), which could be used to compensate for power transients in a cross-layer based fashion.
- **Ekinops 360:** In the *Ekinops 360 transport platform* [eki09_1], multiple client inputs are encapsulated together with a digital wrapper that can optionally contain FEC. This functionality is provided by Ekinops' 10G long-haul transponders (PM 1008, PM 1001LH/RR) and aggregation modules (PM 1004); this could then be used to deduce pre-FEC BER. The management card is accessible via two Ethernet ports and a command line interface allows for (remote) management and performance monitoring.
- **NetFPGA:** The *NetFPGA* is a low-cost platform, primarily designed as a tool for teaching networking hardware and router design [netf08_1]. It has also proved to be a useful tool for networking researchers. Currently, it is implemented in networks used by Cluster B and Cluster C. At a high level, the board contains four 1 Gigabit/second Ethernet (GigE) interfaces, a user programmable Field Programmable Gate Array (FPGA), and four banks of locally-attached static and dynamic Random Access Memory (SRAM and DRAM). It has a standard PCI interface allowing it to be connected to a desktop PC or server. The use of an FPGA allows implementing customized performance monitoring functionalities from scratch. The 1 Gb/s interfaces are rather slow but definitely suitable for exchanging measurement information. The next version of NetFPGAs is planned to incorporate 10Gb/s serial interfaces [netfpga08_2].

Disclaimer: The summary of measurement equipment available in GENI's substrate infrastructure given in this section does not raise the claim of completeness. GENI working group members are invited for further discussions. Any information about additionally, current or future, available (real-time) measurement equipment usable for physical layer performance metrics should be sent to the authors (cf. contact information [geni08_1]) to be included in future revisions of this document.

3 Requirements for User Access to Measurement Data

Interfacing the various available performance monitors (PMONs) identified within Section 2 to the control framework and to the access point of the GENI researcher in a straightforward way may lead to the following obstacles:

- Depending on the number of available PMONs, the number of required interfaces might become very large, leading to a vast overhead in design (different interface specifications, huge number of individual and potentially vendor-specific proprietary implementations, data formats, …), hardware (physical interfaces, access terminals, …) and software (many different drivers, exchange formats, …).
- Designing the interfaces requires detailed knowledge about the control mechanisms of the PMONs and of how measurement data is exported. Because of the very nature of the substrates, which mainly employ commercial equipment, this information might not be publicly available to every researcher (or GENI partner) to the necessary extent. Some abstraction is highly desired.
- A certain amount of editing and preparation of the measured data prior to the delivery to the researcher could be desirable, e.g. if not all information is required, if only part of the measured data is requested, if control signals should be derived, or if the researcher is only allowed to view part of the available measurements.
- Extensibility of the measurement framework and manageability, e.g. the slivering by the control framework, of all (or only a subset of) PMONs is difficult if each PMON has to be controlled individually.

Therefore, rather than interfacing every performance monitoring device within the substrate directly with the control framework and the experimenter, we recommend the design and use of an unified measurement framework (UMF). The UMF represents a universal measurement platform which can be accessed by the control framework and the researcher via a limited number of well defined interfaces. The main tasks and functionalities of the UMF are (cf. Figure 3-1):

- Acquiring measurement data from the various performance monitoring devices within the optical substrate,
- abstraction of measurement capabilities and equipment from several manufacturers,
- provide a single point of access,
- basic processing of the measurement data, e.g., to extrapolate signaling data for the researcher,
- provisioning of some storage capacity for non-time-sensitive measurements,
- interfacing to the researcher via a unified measurement interface which allows requesting and controlling certain measurements and delivers the measured data, and
- interfacing with the control framework so that the measurement framework (or a subset of PMONs) can be allocated to a requesting GENI researcher and reconfiguration of the slice is made possible.

As such, the UMF will present a uniform view of measurement capabilities within the substrate, abstracting the measurement capabilities of the substrate and make them accessible/sliceable to a control framework.

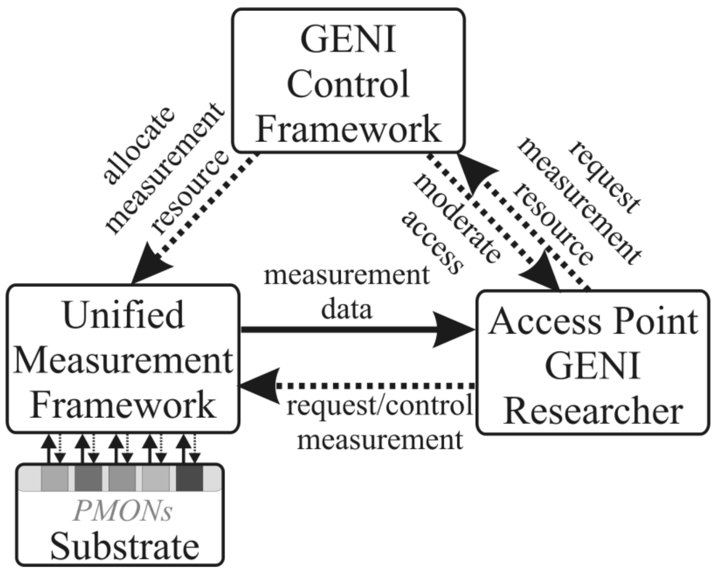

Figure 3-1: Schematic of interaction between performance monitors (PMONs) of substrate, unified measurement framework (UMF), control plane, and GENI researcher

4 Bibliography

[1] [geni08_1] Global Environment for Network Innovations – Wikipedia, "Embedded Real-Time Measurements (Project Nr. 1631)" (2009, January) [Online]. Available: http://groups.geni.net/geni/wiki/Embedded%20Real-Time%20Measurements

[2] [geni09_1] GENI Project Office, "Spiral 1 substrate catalog", 05. January 2009

[3] [inf08_1] Infinera, "The DTN in detail" (2008, December) [Online]. Available: http://www.infinera.com/products/dtn_detail.html

[4] [eki09_1] Ekinops, "Ekinops 360 Platform" (2009, February) [Online]. Available: http://www.ekinops.net/products-01.html

[5] [adv09_1] Adva Optical Networking, "The Scalable Optical Transport Product Family" (2009, February) [Online]. Available: http://www.advaoptical.com

[6] [polatis08_1] Polatis, "VOA and Power Meter Switch Tray VST" (2008, December) [Online]. Available: www.polatis.com

[7] [ciena08_1] Ciena, "CN 4200 FlexSelect Advanced Services Platform Family" (2008, December) [Online]. Available: http://www.ciena.com/products/products_cn4200_overview.htm

[8] [netf08_1] (2008, December) [Online]. Available: http://www.netfpga.org

[9] [netfpga08_2] J.W. Lockwood, "NetFPGA - A quick update for the GEC3 conference" (2008, October) [Online]. Available: *http://g*roups.geni.net/geni/attachment/wiki/presentations/NetFPGA_Status_Update_at_GEC3.ppt?format=raw

[10] [jun09_1] Juniper, "Juniper Networks EX-Series" (2009, February) [Online]. Available: http://www.juniper.net/products_and_services/ex_series/index.html

[11] [morkel08_1] P.R. Morkel et. al., "Integrated IP-Optical Networks. Demonstration of DWDM Router-to-Router IP Transport Over 574km SMF Fiber Link Using 11.1Gbit/s OTN Pluggable Interface with Integrated G.709 and FEC," OFC 2009, Paper NME4