

Scalable and Extensible Network Monitoring For GENI

PIs: Sonia Fahmy (Purdue University), Puneet Sharma (HP Labs)

Contributors: Ethan Blanton, Sriharsha Gangam, Nabeel Butt
(Purdue University)

Sumit Kala, Nagarjuna Reddy Chevuru, Ranjan Kumar (HP)

Leverages prior work joint with:

Sujata Banerjee, Sujoy Basu, SJ Lee, Praveen Yalagandula (HP Labs),

Greg N. Frederickson (Purdue University)

<http://illusion.hpl.hp.com/genis3monitor>



RECAP

Goals

- Provide ProtoGENI system state in real-time
 - Primarily network (and maybe node) state
 - *Active* and passive measurements
 - *E2E* or leverages network element info when available
- Flexible and extensible
 - Easy to add new measurement tools to be developed!
 - Configurable time scales (start time, frequency, number)
- Share measurement info across applications/slices
 - Eliminate redundant expensive measurements
- *Scalable*, secure, and reliable

Leverages the following prior work

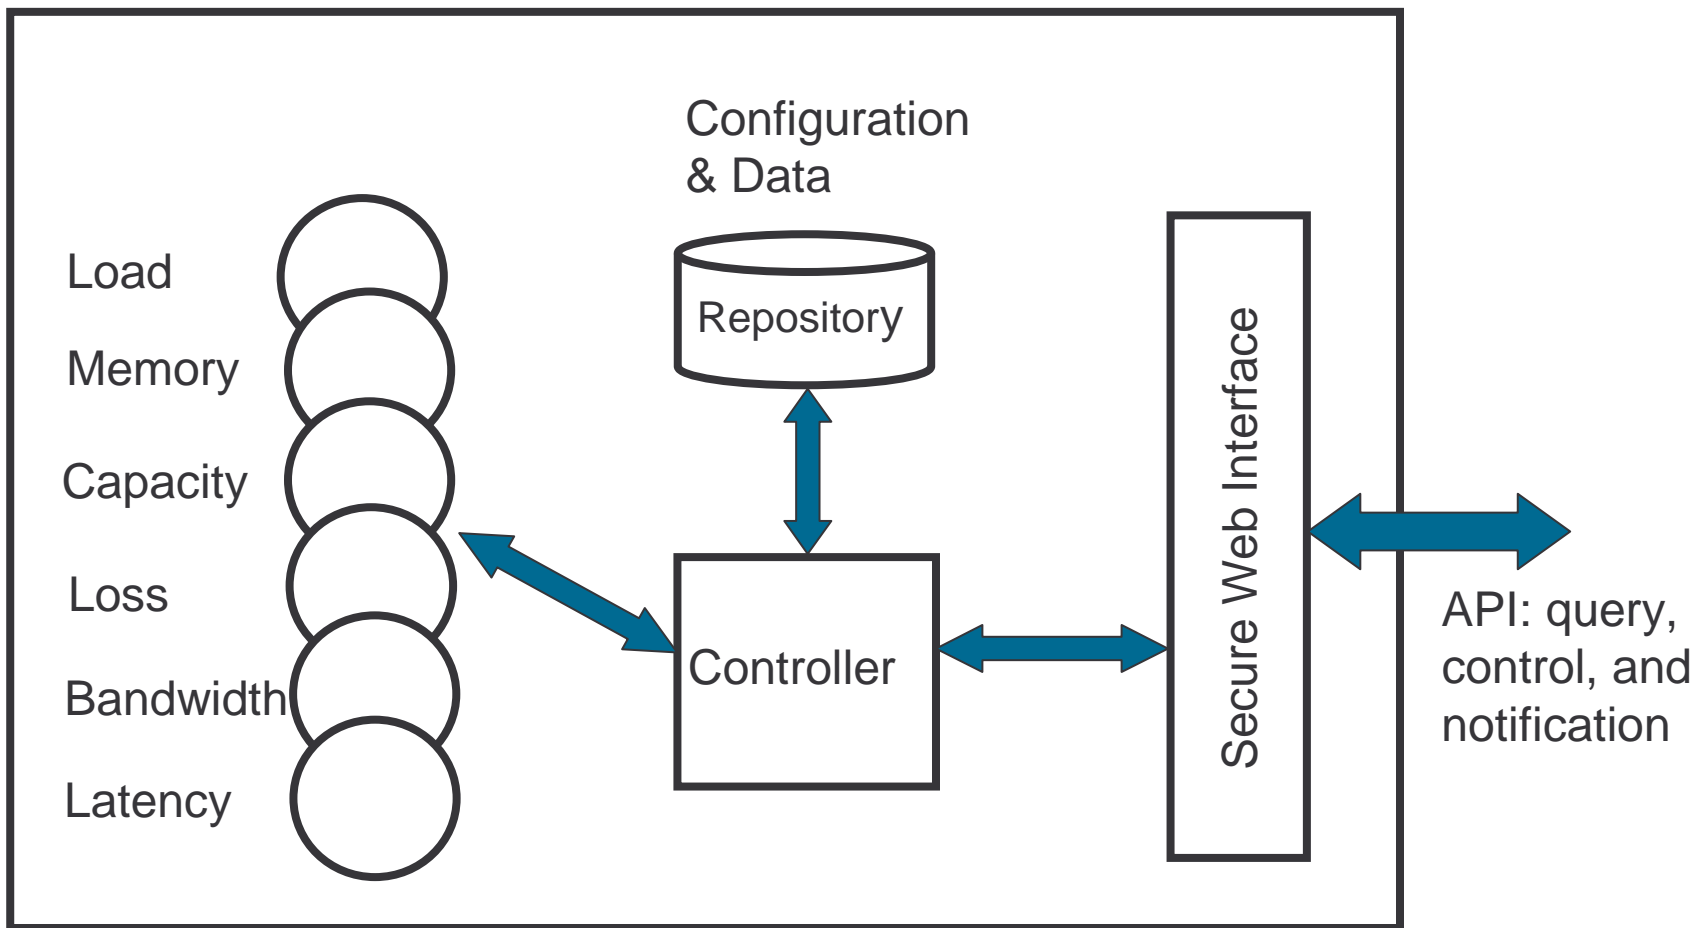
- <http://networking.hpl.hp.com/s-cube>
- Ethan Blanton, Sonia Fahmy, Greg N. Frederickson, "On the Utility of Inference Mechanisms," In Proceedings of IEEE International Conference on Distributed Computing Systems (ICDCS), 8 pp., June 2009.
- Ethan Blanton, Sonia Fahmy, Sujata Banerjee, "Resource Management in an Active Measurement Service," In Proceedings of the IEEE Global Internet Symposium, 6 pp., April 2008.
- P. Yalagandula, P. Sharma, S. Banerjee, S.-J.Lee, and S. Basu, [S3: A Scalable Sensing Service for Monitoring Large Networked Systems,](#) In *Proceedings of the Workshop on Internet Network Management 2006*, Pisa, Italy, September 2006.
- Praveen Yalagandula, Sung-Ju Lee, Puneet Sharma, and Sujata Banerjee, [Correlations in End-to-End Network Metrics: Impact on Large Scale Network Monitoring,](#) In *IEEE Global Internet Symposium*, Phoenix, AZ, April 2008.

Scalable Sensing Service (S³)

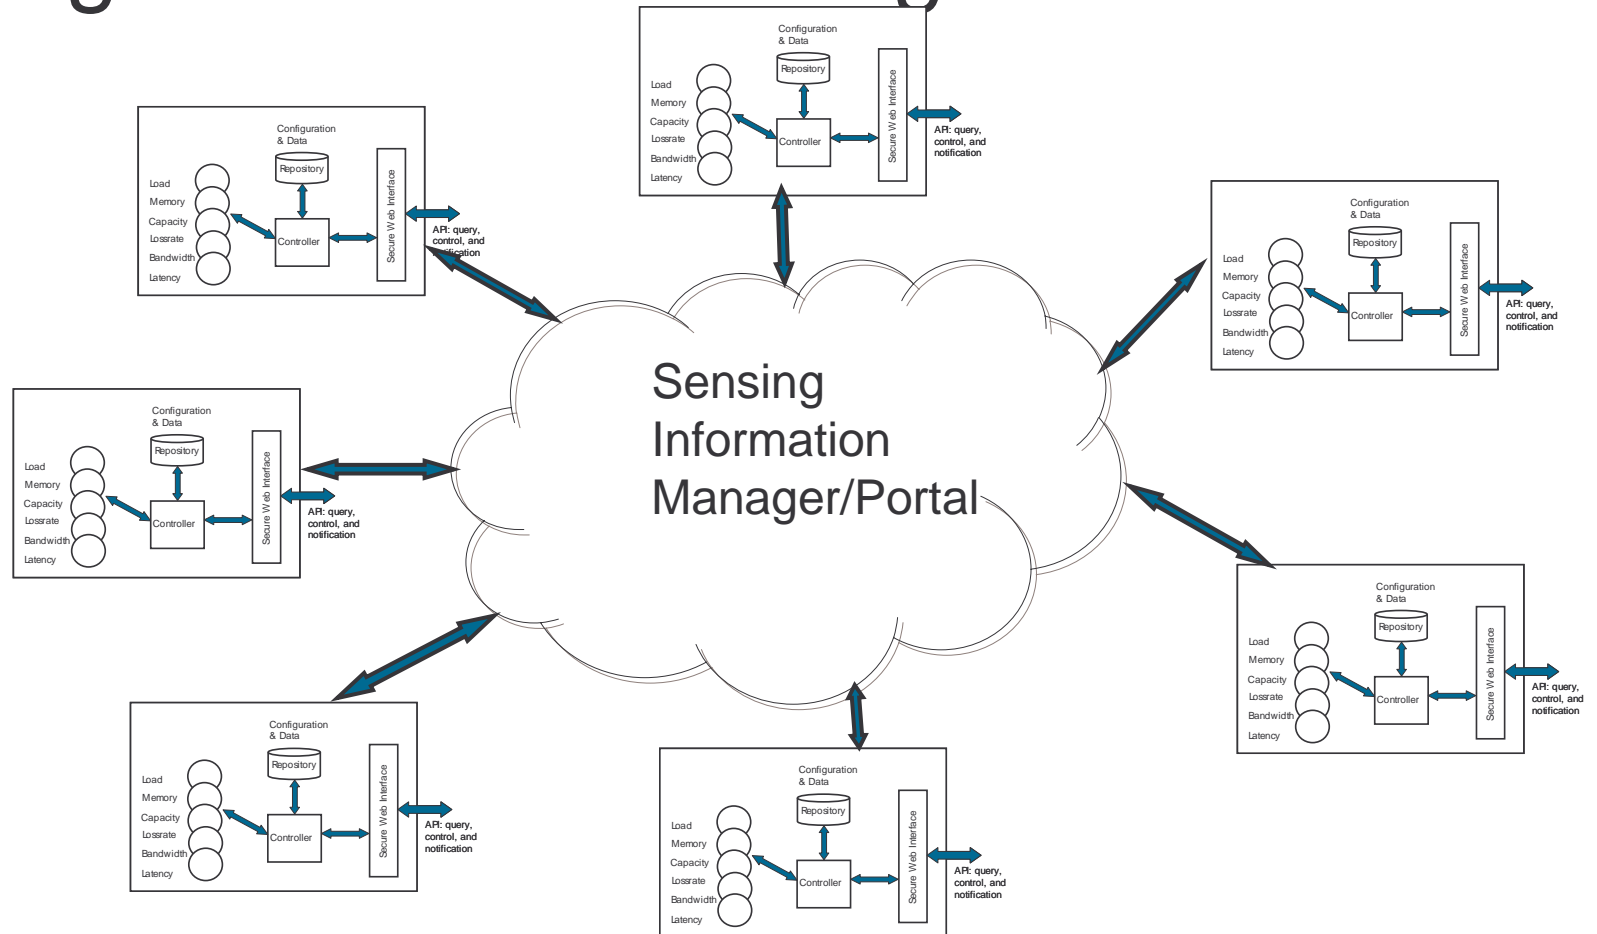
- Sensor pods
 - Measure system state from a node perspective
 - Web-Service enabled collection of sensors
- Sensing information manager
 - Controls pods, and aggregates measured system state
 - Portal to request and invoke measurements
- Inference engines
 - Infer $O(n^2)$ E2E path info by measuring a few paths
 - Dynamically schedules measurements on pods

Sensor Pod

Web-Service (WS) enabled collection of sensors



Sensing Information Manager

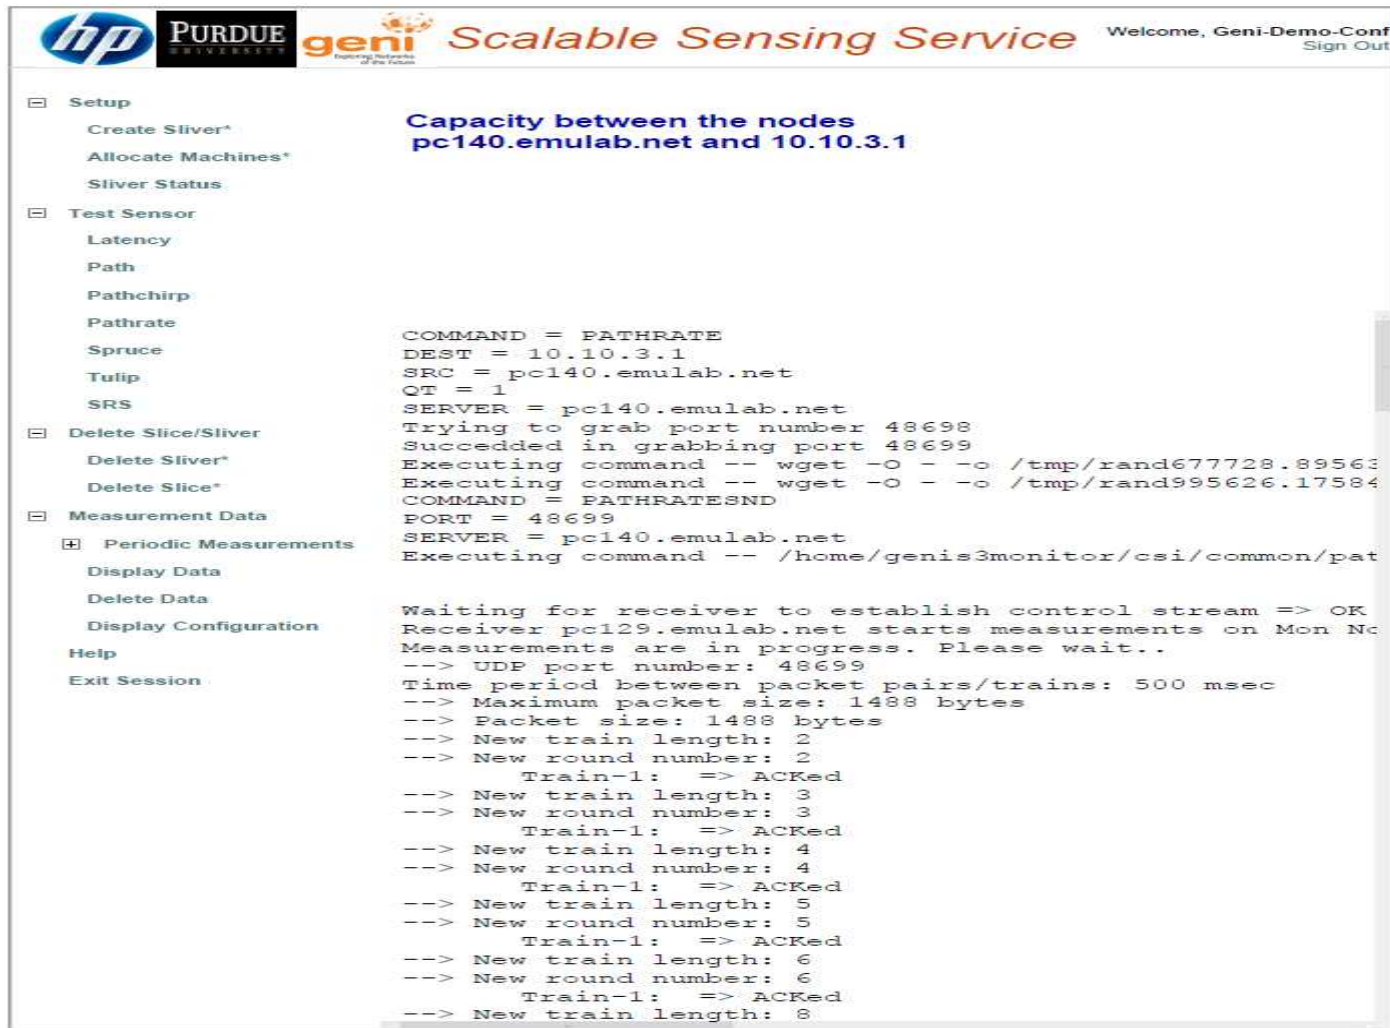


- Control the sensor-pods
- Aggregate data from sensor
- Answer researcher queries

YEAR 1 Accomplishments

- Integration with ProtoGENI API and mechanisms
- S3Monitor platform v1.0 available for deployment
 - Sensor pod
 - Sensing Information manager
- Support for on-demand and periodic measurements
- Archiving and querying capabilities for measurement data
- Continuous experimentation/demo
- Demo page up at:
 - <http://illusion.hpl.hp.com:8180/geni/demo.jsp>

Screenshot: On-demand Measurement



The screenshot displays the 'Scalable Sensing Service' web interface. The header includes logos for HP, PURDUE UNIVERSITY, and geni, along with the text 'Scalable Sensing Service' and 'Welcome, Geni-Demo-Conf Sign Out'. The main content area is a terminal window titled 'Capacity between the nodes pc140.emulab.net and 10.10.3.1'. The terminal output shows the execution of a 'PATHRATE' command, the successful establishment of a control stream, and the start of measurements. The results show a time period of 500 msec, a maximum packet size of 1488 bytes, and a series of train lengths (2, 3, 4, 5, 6, 8) and round numbers (2, 3, 4, 5, 6, 8) for 'Train-1'.

```
COMMAND = PATHRATE
DEST = 10.10.3.1
SRC = pc140.emulab.net
QT = 1
SERVER = pc140.emulab.net
Trying to grab port number 48698
Succeeded in grabbing port 48699
Executing command -- wget -O - -o /tmp/rand677728.89563
Executing command -- wget -O - -o /tmp/rand995626.17584
COMMAND = PATHRATESND
PORT = 48699
SERVER = pc140.emulab.net
Executing command -- /home/genis3monitor/csi/common/pat

Waiting for receiver to establish control stream => OK
Receiver pc129.emulab.net starts measurements on Mon No
Measurements are in progress. Please wait..
--> UDP port number: 48699
--> Maximum packet size: 1488 bytes
--> Packet size: 1488 bytes
--> New train length: 2
--> New round number: 2
    Train-1: => ACKed
--> New train length: 3
--> New round number: 3
    Train-1: => ACKed
--> New train length: 4
--> New round number: 4
    Train-1: => ACKed
--> New train length: 5
--> New round number: 5
    Train-1: => ACKed
--> New train length: 6
--> New round number: 6
    Train-1: => ACKed
--> New train length: 8
```

Screenshot: Query on Periodic Data Archive

hp PURDUE UNIVERSITY geni Scalable Sensing Service Welcome, Geni-Demo-Conf Sign Out

Setup
 Create Sliver*
 Allocate Machines*
 Sliver Status

Test Sensor
 Latency
 Path
 Pathchirp
 Pathrate
 Spruce
 Tulip
 SRS

Delete Slice/Sliver
 Delete Sliver*
 Delete Slice*

Measurement Data
 Periodic Measurements
 Display Data
 Delete Data
 Display Configuration

Help
 Exit Session

Displaying measurements

125 Record(s) have been found.

Time	Measurement Type	Value	Measurement Type	Value
10-28-2010 10:15:05	Min. Capacity Estimation (patharte) in Mbps	69.0	Max. Capacity Estimation (patharte) in Mbps	78.0
10-28-2010 10:15:05	Min. Capacity Estimation (patharte) in Mbps	69.0	Max. Capacity Estimation (patharte) in Mbps	78.0
10-28-2010 10:15:05	Min. Capacity Estimation (patharte) in Mbps	69.0	Max. Capacity Estimation (patharte) in Mbps	79.0
10-28-2010 10:15:05	Min. Capacity Estimation (patharte) in Mbps	69.0	Max. Capacity Estimation (patharte) in Mbps	78.0
10-28-2010 10:15:05	Min. Capacity Estimation (patharte) in Mbps	69.0	Max. Capacity Estimation (patharte) in Mbps	78.0
10-28-2010 10:15:05	Min. Capacity Estimation (patharte) in Mbps	69.0	Max. Capacity Estimation (patharte) in Mbps	79.0
10-28-2010 10:20:06	Min. Capacity Estimation (patharte) in Mbps	70.0	Max. Capacity Estimation (patharte) in Mbps	77.0
10-28-2010 10:20:06	Min. Capacity Estimation (patharte) in Mbps	70.0	Max. Capacity Estimation (patharte) in Mbps	77.0
10-28-2010 10:25:05	Min. Capacity Estimation (patharte) in Mbps	59.0	Max. Capacity Estimation (patharte) in Mbps	66.0
10-28-2010 10:30:05	Min. Capacity Estimation (patharte) in Mbps	68.0	Max. Capacity Estimation (patharte) in Mbps	75.0
10-28-2010 10:35:06	Min. Capacity Estimation (patharte) in Mbps	56.0	Max. Capacity Estimation (patharte) in Mbps	67.0
10-28-2010 10:40:06	Min. Capacity Estimation (patharte) in Mbps	56.0	Max. Capacity Estimation (patharte) in Mbps	67.0
10-28-2010 10:45:06	Min. Capacity Estimation (patharte) in Mbps	55.0	Max. Capacity Estimation (patharte) in Mbps	65.0
10-28-2010 10:45:06	Min. Capacity Estimation (patharte) in Mbps	55.0	Max. Capacity Estimation (patharte) in Mbps	66.0
10-28-2010 10:50:06	Min. Capacity Estimation (patharte) in Mbps	70.0	Max. Capacity Estimation (patharte) in Mbps	77.0
10-28-2010 10:55:06	Min. Capacity Estimation (patharte) in Mbps	69.0	Max. Capacity Estimation (patharte) in Mbps	76.0

YEAR 1: Other Accomplishments

- Instrumentation and Measurement WG
 - Collaboration with other researchers on GENI I&M specification document version 0.5
 - Discussions on similarities and differences with other measurement projects; participation in conference calls, etc.
- Service on TPCs of SIGMETRICS, INFOCOM, ICNP, COMSNETS, ...
- Presentation on INFOCOM panel
- Posters on the project at several events

YEAR 1: Publications

- "Distributed Partial Inference under Churn," Sriharsha Gangam, Sonia Fahmy, In Proceedings of the IEEE Global Internet Symposium, 6 pp., March 2010.
- "Leveraging Correlations Between Capacity and Available Bandwidth To Scale Network Monitoring," Praveen Yalagandula, Sung-Ju Lee, Puneet Sharma, Sujata Banerjee, In Proceedings of IEEE GlobeCom 2010, accepted for publication.
- "On the Cost of Network Inference Mechanisms," Ethan Blanton, Sonia Fahmy, Greg N. Frederickson, Sriharsha Gangam, IEEE Transactions on Parallel and Distributed Systems (TPDS), 14 pp., accepted for publication.

YEAR 2: Planned Work

- Extend deployment to PlanetLab and ORCA GENI Clusters
- Include measurement admission control
 - Careful admission control and scheduling of measurement requests prevents misuse
- Integrate S3Monitor data with GENI's I&M Measurement WG's data and meta-data formats
- Enable and support experimenters for using S3 platform on GENI clusters

Questions/Comments/Code

<http://illusion.hpl.hp.com/genis3monitor>

puneet.sharma@hp.com

fahmy@cs.purdue.edu

Backup

Challenges with Existing Tools

- Tools previously tested only in point-to-point configurations
- Deployment in a large scale setting exposed several issues
 - Hard-coded port numbers leading to port conflicts
 - Need to be started at source and destination simultaneously
 - Large resource requirements leading to end-node crashes
 - Long running times leading to web server timeouts

Scalable Inference Engines

- Large overhead for probing and data exchange
 - $O(N^2)$ measurements in a network of N nodes
 - Dynamically changing \mathcal{E} Need frequent probing
- Measurement/Monitoring failures
 - Failed or slow end machines
 - Measurement tool failures
- Inference based on incomplete information
 - Exploit properties such as triangular inequality
 - A coarse estimate may suffice for many applications
- Prediction based on archived information
- Tradeoff between accuracy and overhead
- When and where to use inference? [Blanton et al., ICDCS09]