

Scalable and Extensible Network Monitoring



Sonia Fahmy, Sriharsha Gangam (Purdue University)
Puneet Sharma, Prakash K. Malligemane (HP)



<http://networking.hpl.hp.com/>
<http://www.cs.purdue.edu/homes/fahmy>

Motivation

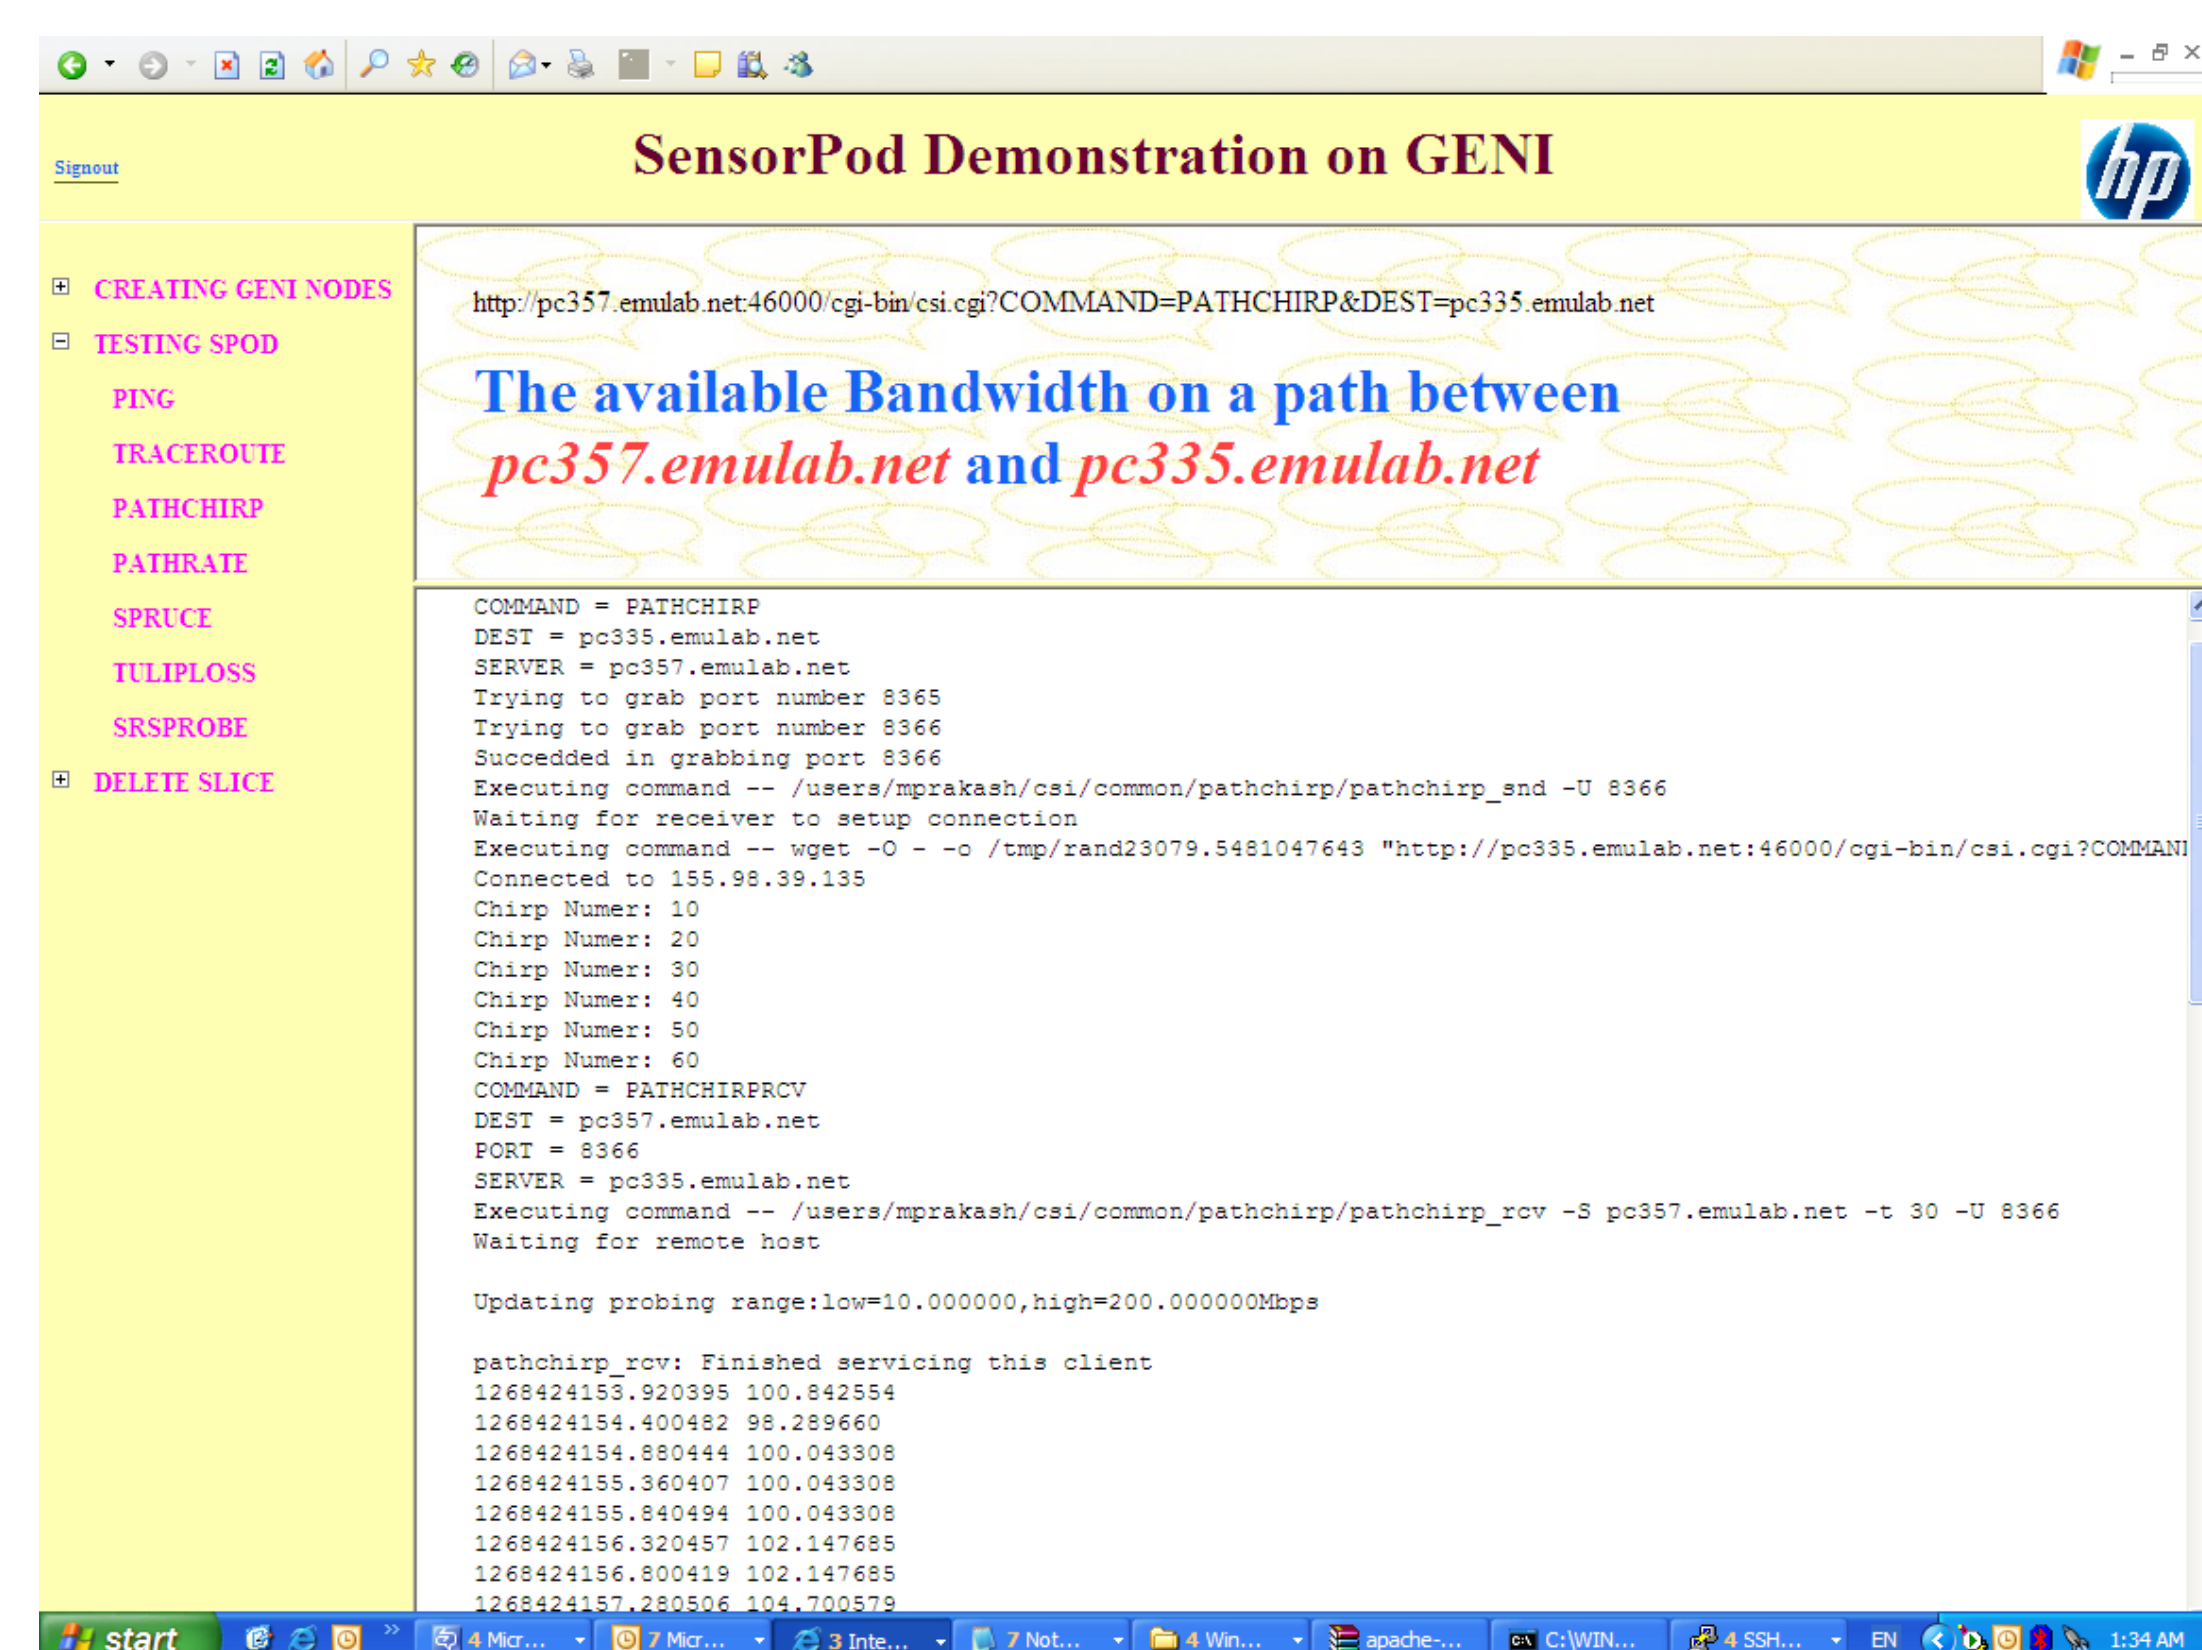
- Provide monitoring information (especially network state information) to ProtoGENI system administrators **and** experimenters

Goals

- Provide ProtoGENI system state *in real-time*
- Obtain network (and maybe node) state
 - Active and passive
 - E2E or leverages network element information 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
- Scalable, secure, and reliable

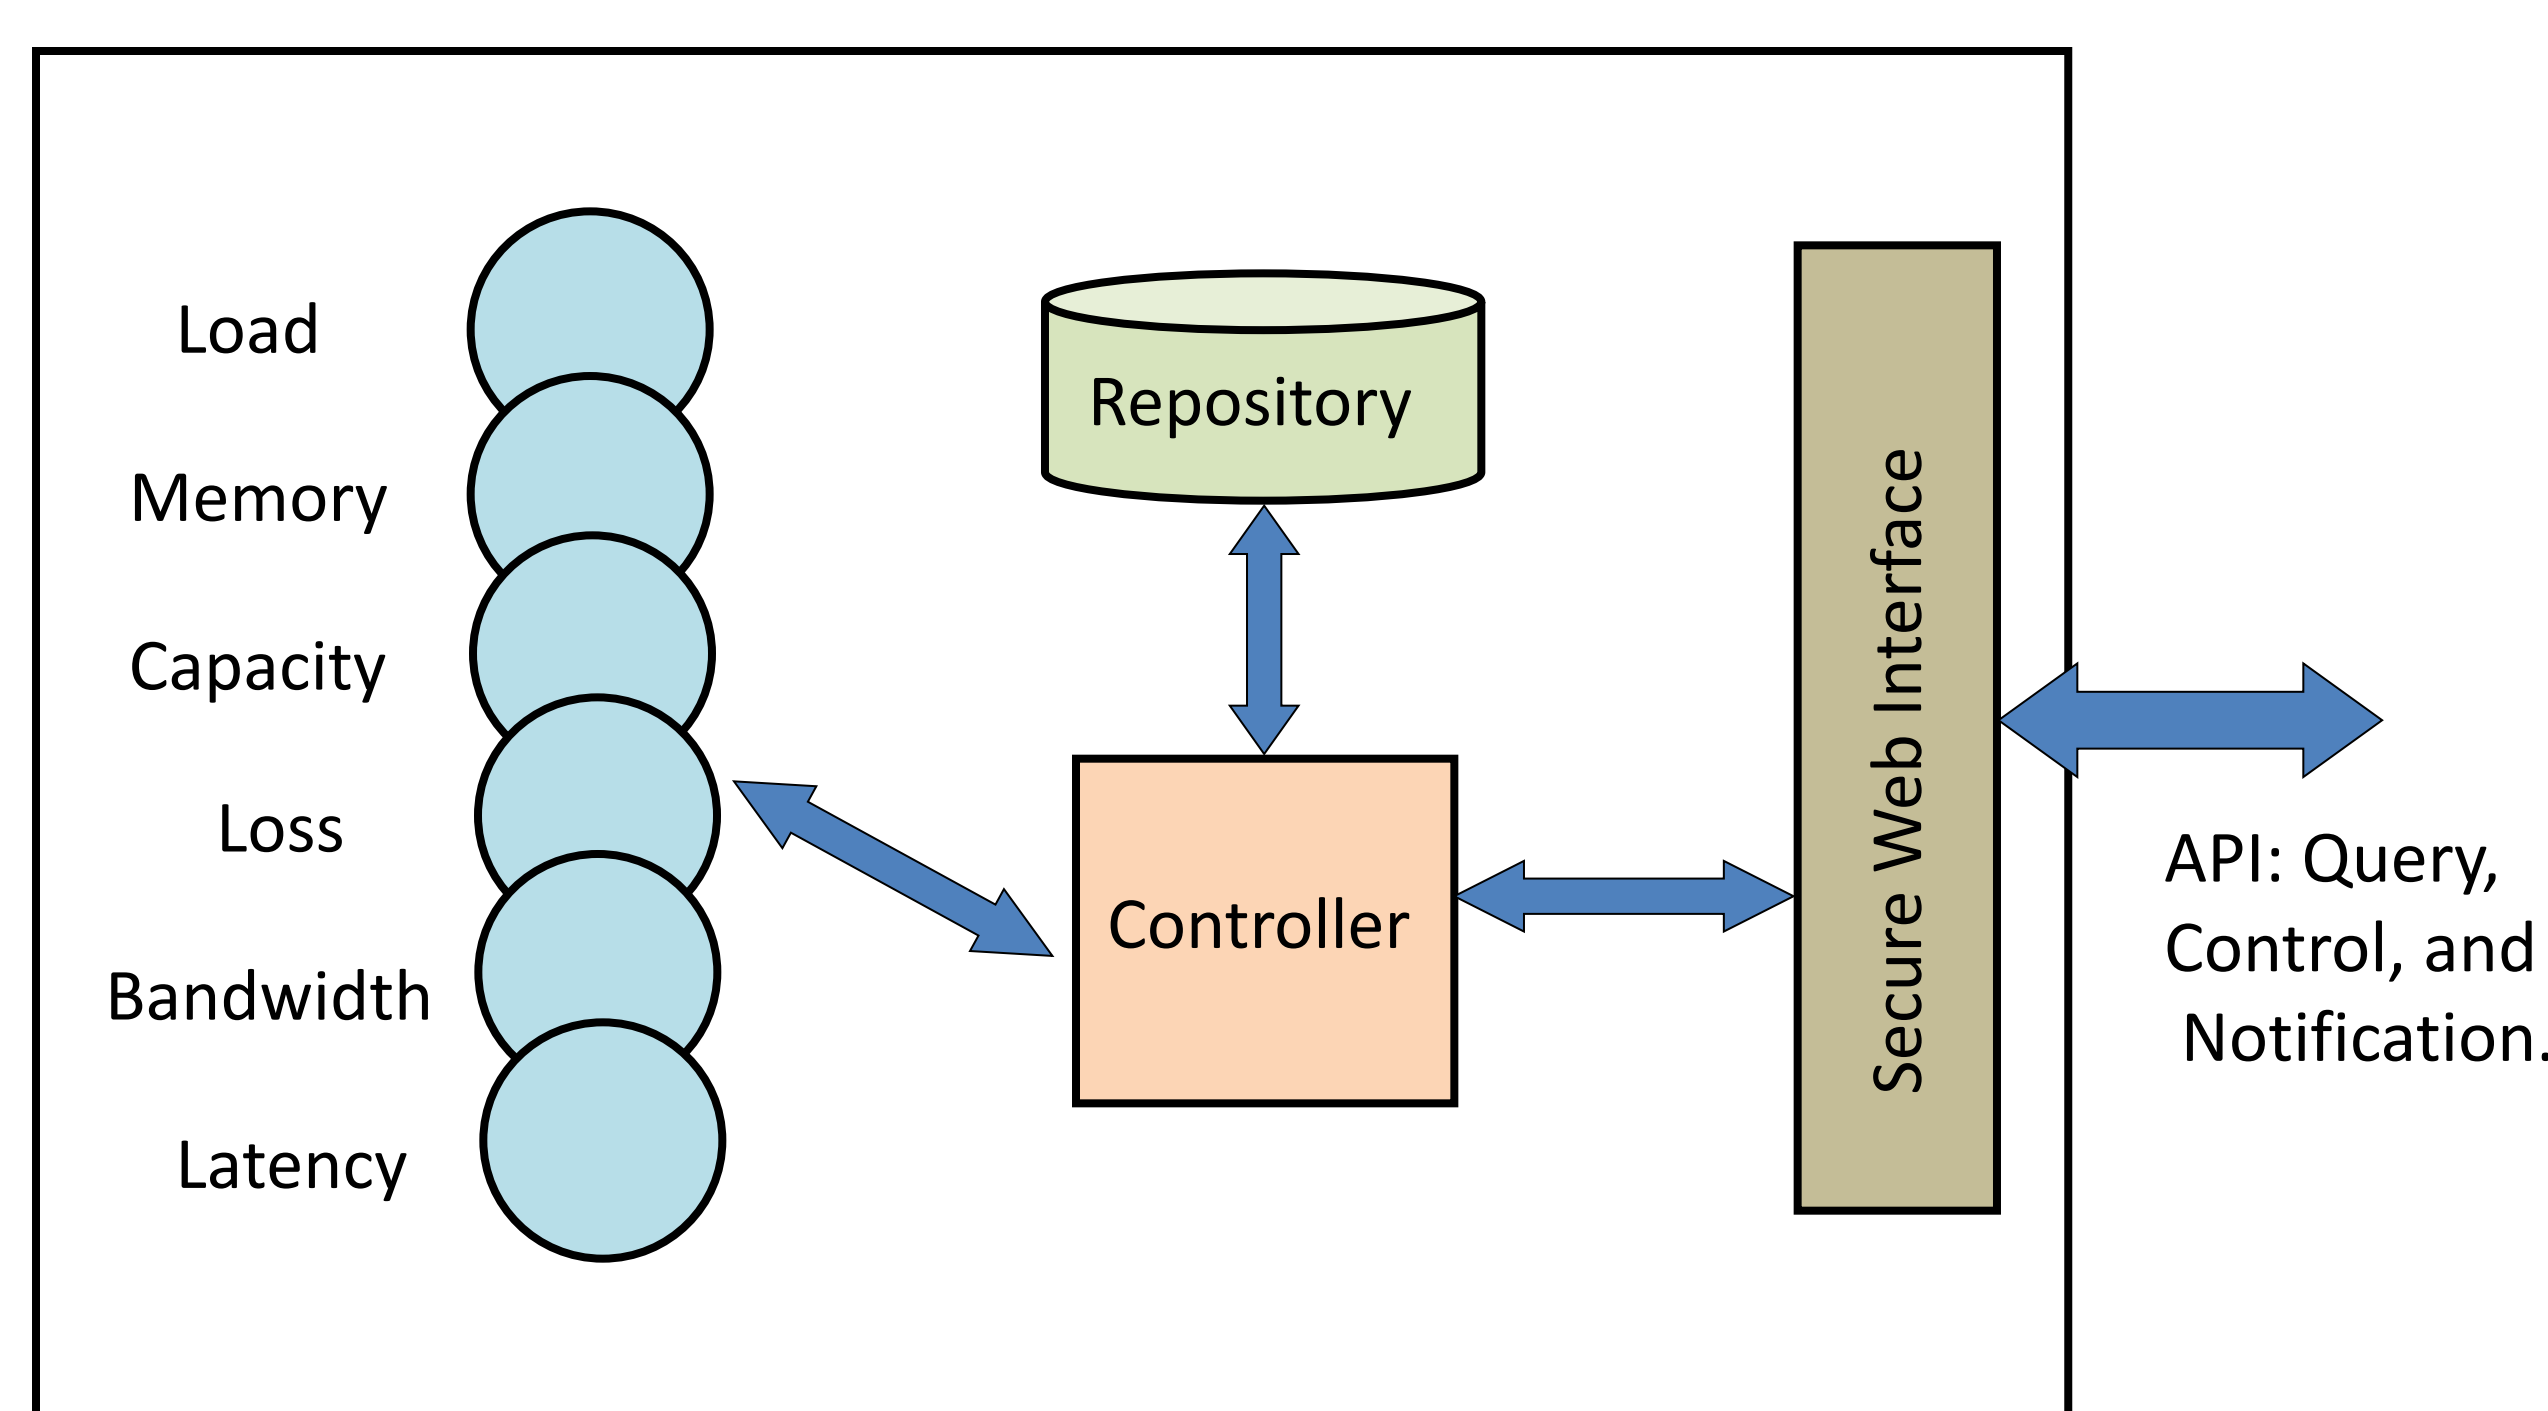
Challenges

- Active measurement tools previously tested only in point-to-point configurations
- Deployment in a large scale setting exposes 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

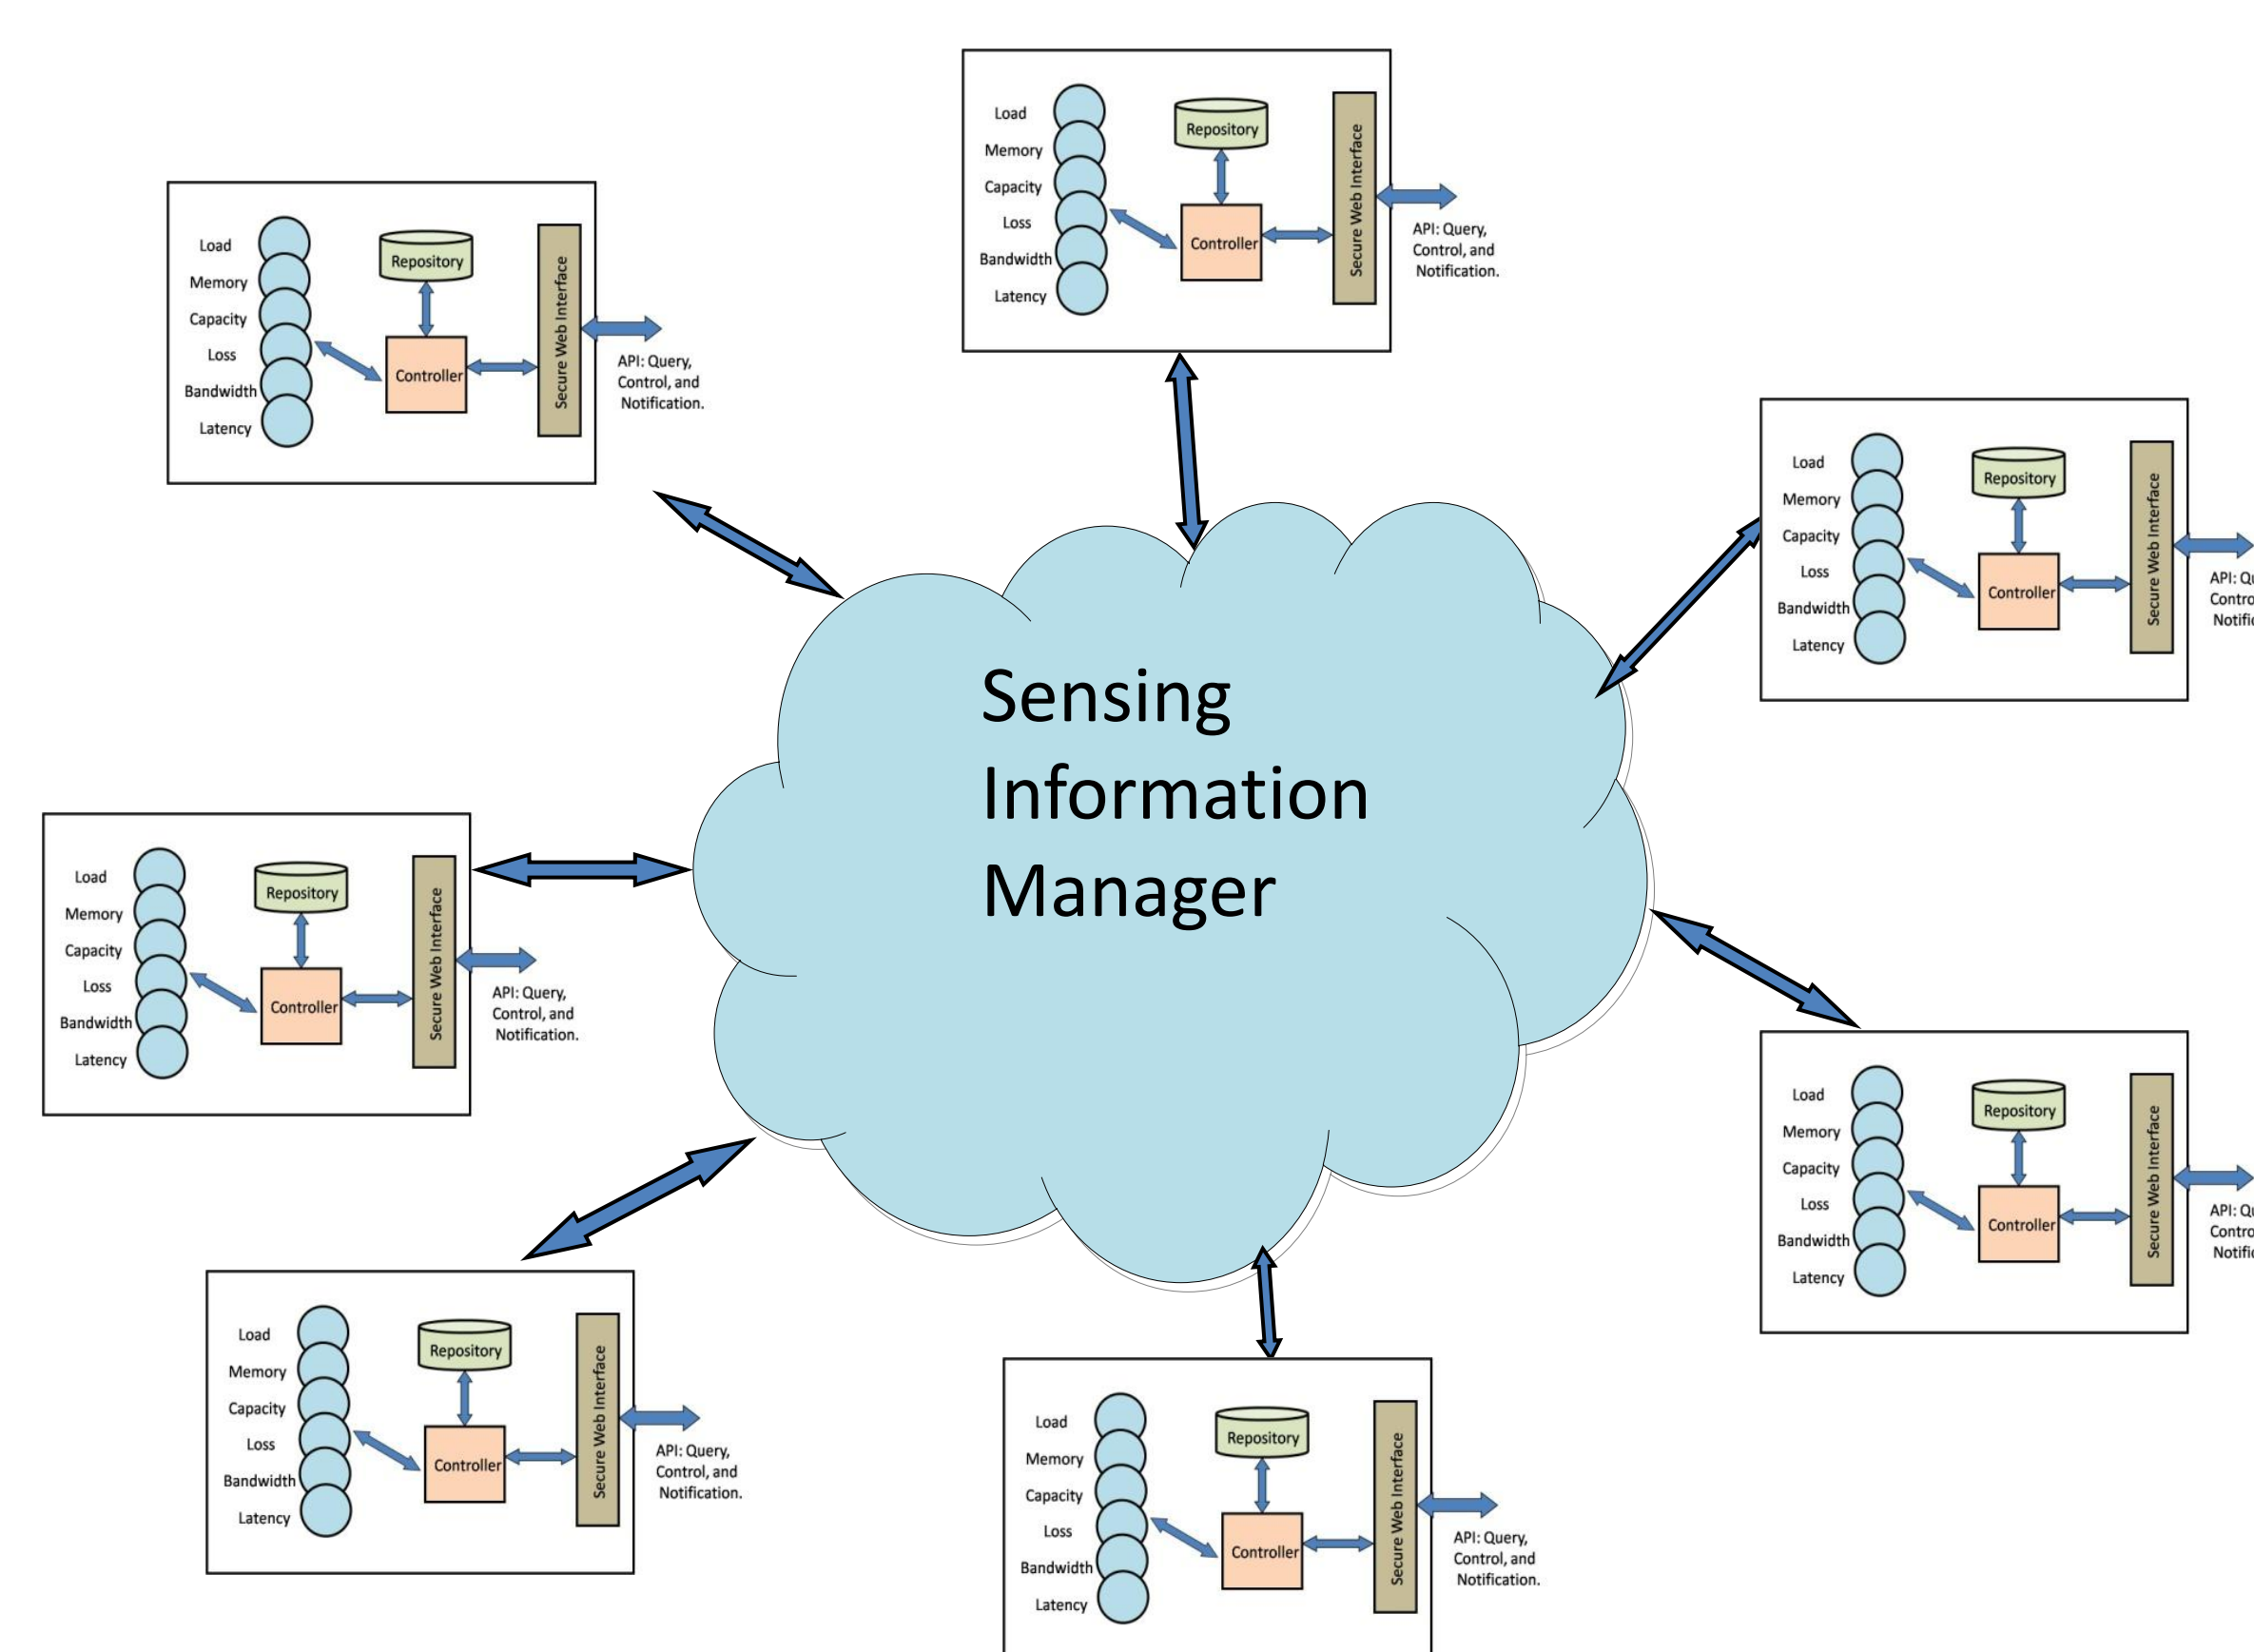


Scalable Sensing Service (S³)

- Sensor pods
 - Measure system state from a node perspective
 - Web-service enabled collection of sensors



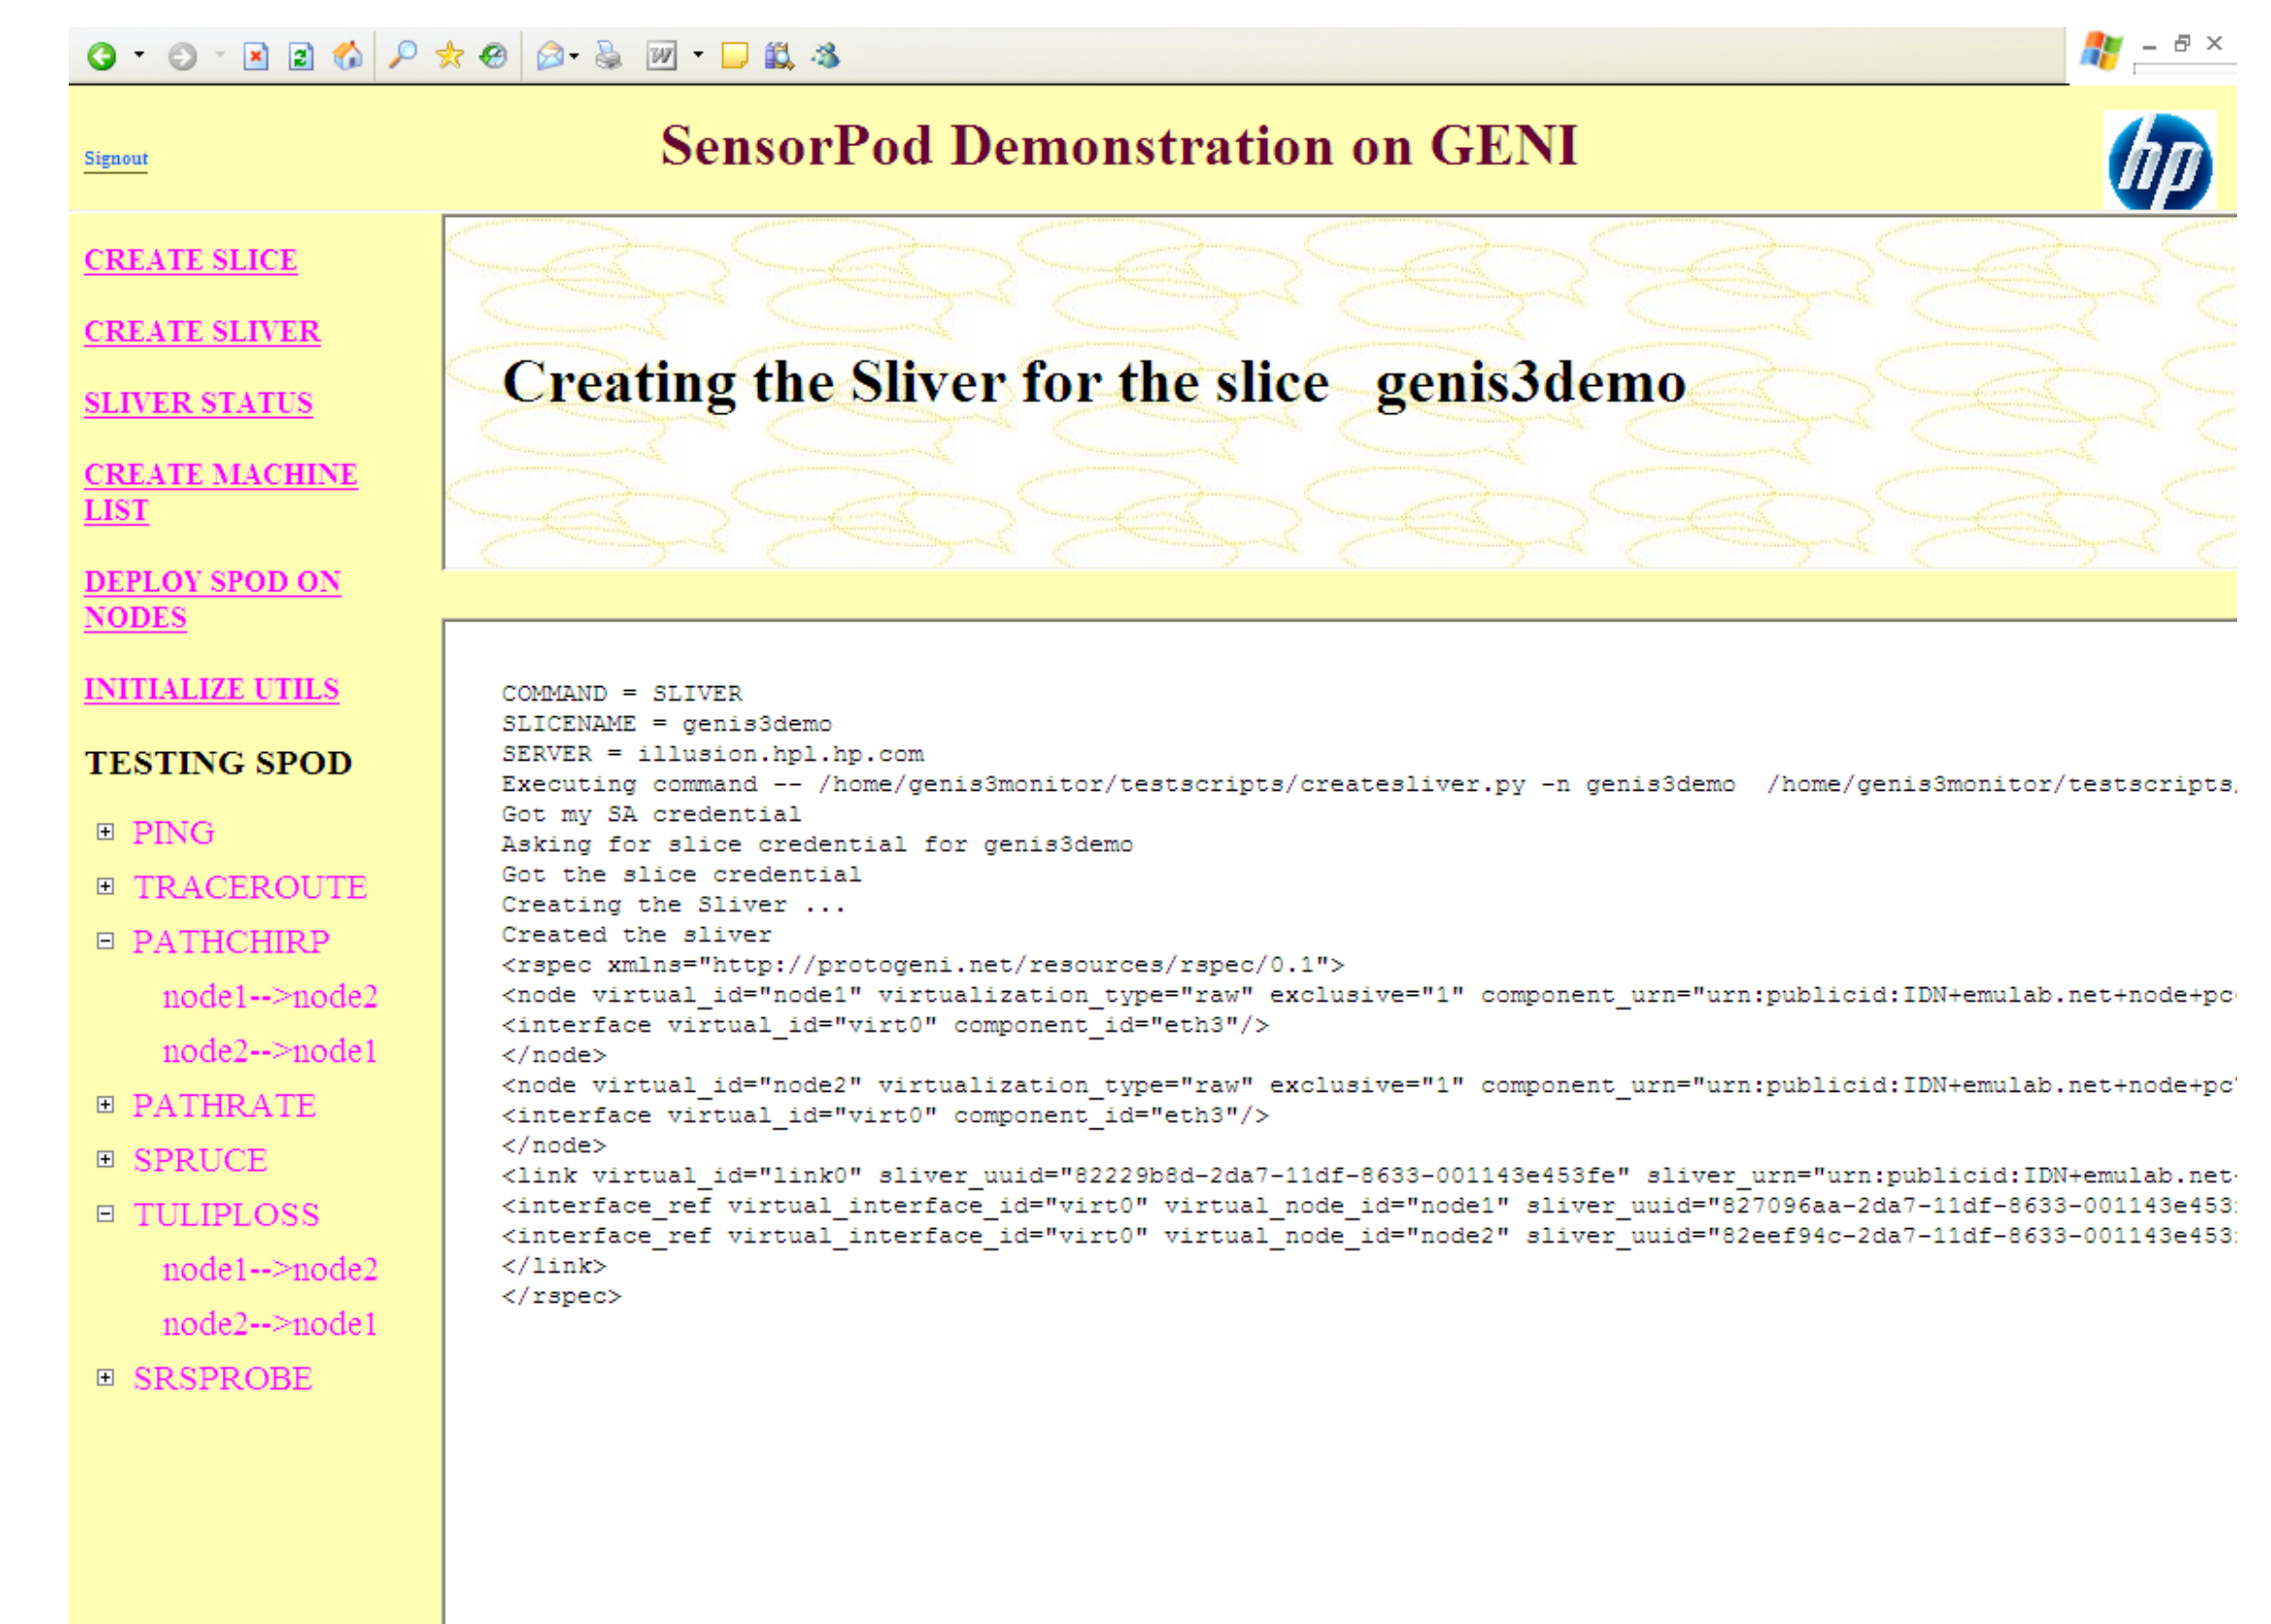
- Sensing information manager
 - Controls pods and aggregates measurements
 - A portal to request and invoke measurements
 - Answer research queries



- Inference engines
 - Infer $O(n^2)$ E2E path info by measuring a few paths
 - Dynamically schedule measurements on sensor pods

Deployment Strategies

- Deploy an instance of S³ per sliver
 - Expensive to deploy at each node
 - Measurements represent the network resources as seen by a user
 - Redundancy in service
- Deploy a single S³ instance for all users
 - Provides only an estimate of the measurements requested by the users
 - Needs sliver information to provide additional accuracy
 - Less expensive



Planned Work

- On-demand measurements at user defined times, frequencies, and tolerance to error/staleness
 - Policy-based scheduling and priorities
 - Estimation of load introduced by measurement probes
- Dynamic invocation of inference mechanisms based on measurement request workload
- Information manager for aggregation and query response
- Integration with ProtoGENI RSPEC

This project is funded by BBN Technologies as GENI project office (GPO).

