Shared Measurement Services

Sonia Fahmy (Purdue University)
Puneet Sharma (HP Labs)

Leverages prior work joint with: Praveen Yalagandula, Sujata Banerjee, Sujoy Basu, SJ Lee (HP Labs), Ethan Blanton, Sriharsha Gangam, Greg N. Frederickson (Purdue University)

http://networking.hpl.hp.com

http://www.cs.purdue.edu/homes/fahmy/

PURDUE UNIVERSITY

hp invent

Motivation

- Experimenters, system administrators, and applications require real-time comprehensive network state
  - For fault diagnosis or increasing performance
- Current network state information
  - Centralized or fragmented ➔ not scalable; no E2E picture
  - Device-centric ➔ Not flow-centric
  - SNMP-based (typically >5-min frequency) ➔ not real-time
  - Limited exposure of only a few metrics

PURDUE UNIVERSITY

hp invent

Most Related GENI Projects

- Started spiral 1
  - U of Wisconsin
  - U of Kentucky
- Starting spiral 2
  - U of Delaware
  - OSU
  - Perhaps others…

PURDUE UNIVERSITY

hp invent

Goals

- Provide system state in real-time
  - Both network and node state
    - *Active* and passive
    - *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)
  - *Support complex queries*
    - To which node do I have the largest bandwidth?
    - Which game server is within 10ms latency?
- Share measurement info across applications
  - Eliminate redundant expensive measurements
- *Scalable*, secure, and reliable

PURDUE UNIVERSITY

hp invent

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 Sensing Service (S^3)

- Sensor pods
  - Measure system state from a node perspective
  - Web-Service enabled collection of sensors
- Backplane
  - Distributed programmable fabric
  - Connects pods, and aggregates measured system state
- Inference engines
  - Infer $O(n^2)$ E2E path info by measuring a few paths
  - Dynamically schedules measurements on pods
  - Aggregates data on backplane

PURDUE UNIVERSITY

hp invent

Sensor Pod

Web-Service (WS) enabled collection of sensors

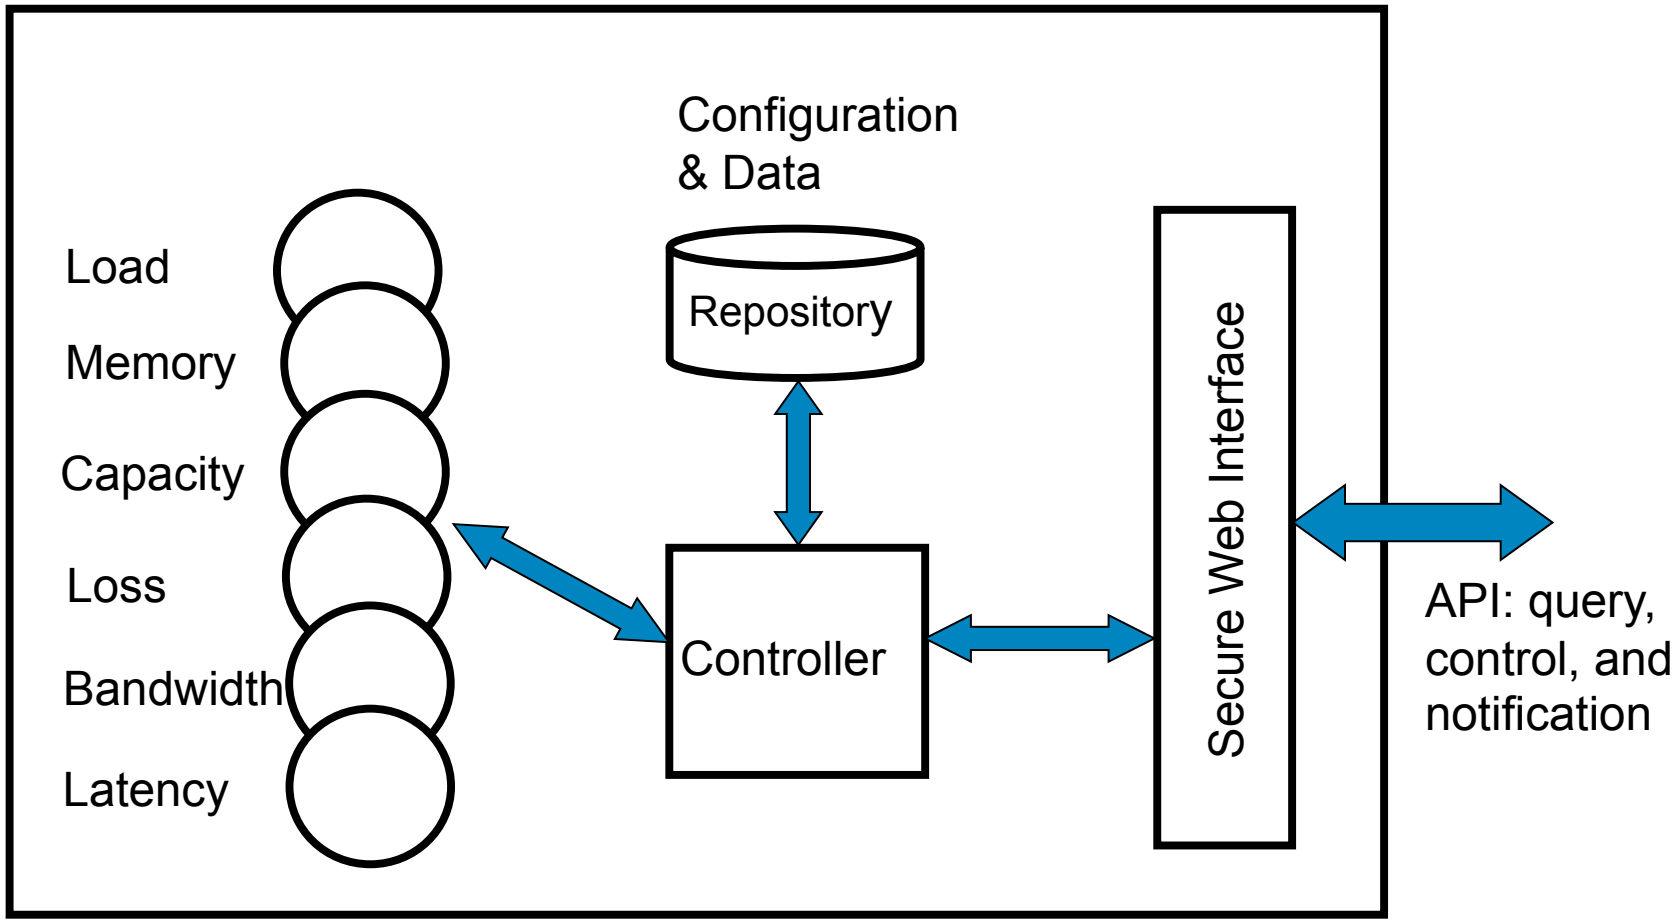

Sensing Backplane

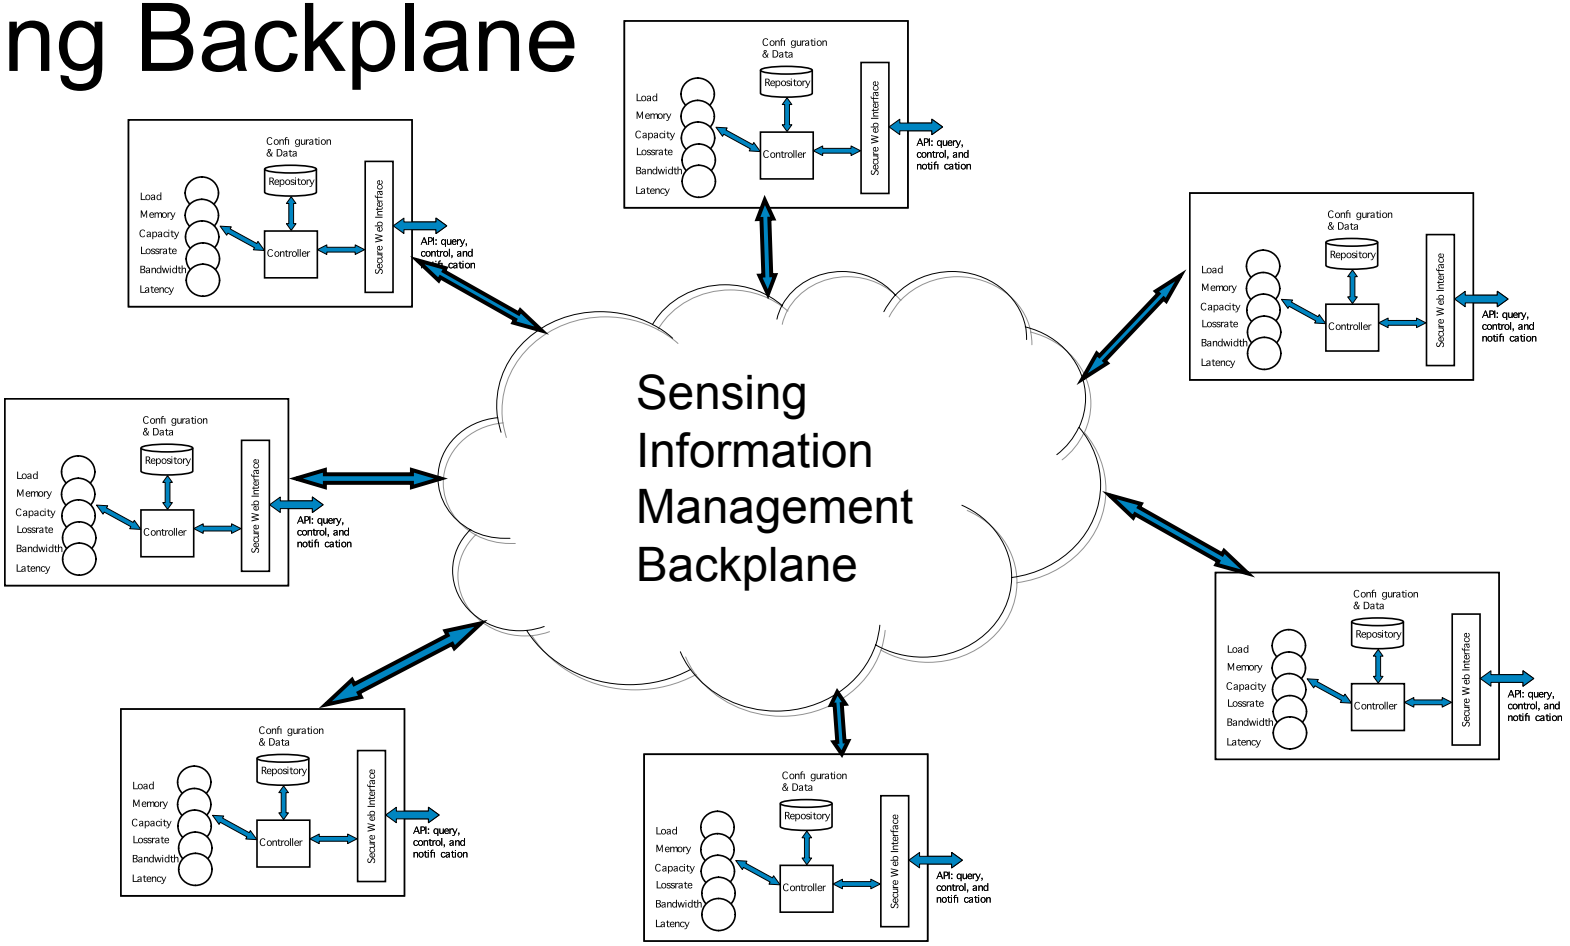

- Aggregate data from end-points
- Answer queries
- E.g., SDIMS [SIGCOMM 2004]

PURDUE UNIVERSITY

hp invent

Scalable Inference Engines

- Large overhead for probing and data exchange
  - $O(N^2)$ measurements in a network of N nodes
  - Dynamically changing ➔ 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]

Prototype Deployment on PlanetLab

- 700+ nodes scattered across 350+ sites
- Running since January 2006
- All pair network metrics: E2E latency, BW, Capacity, Loss
- Simple backplane: central server
  - Maintains pods, schedules measurements, collects and publishes data
- Stats:~14GB raw data every day, ~1GB compressed

S^3 Data Usage

- Web server stats (2006):
  - ~200 unique visitors/month
  - ~20GB download BW/month

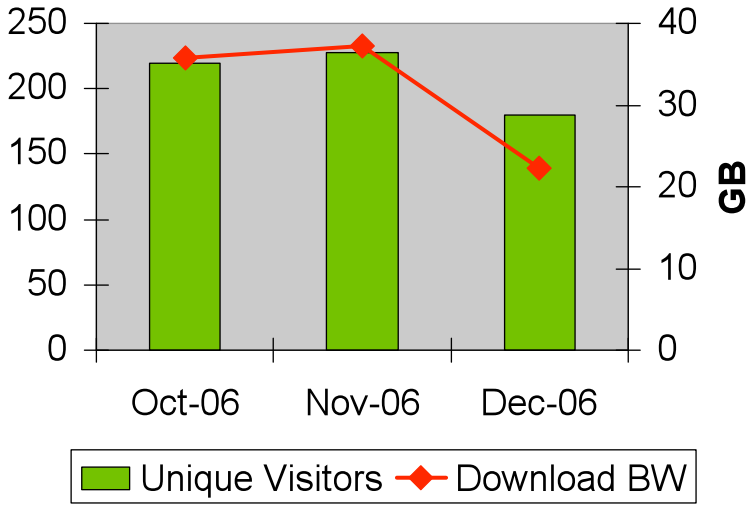

- Projects
  - HP Labs: Bandwidth inference, Resource-aware monitoring, semantic store
  - Others: Purdue University, MSR, U of Washington, Georgia Tech, Harvard, Princeton, Boston University, etc.

PURDUE UNIVERSITY

hp invent

Screenshot: Hop-by-hop Loss Sensor

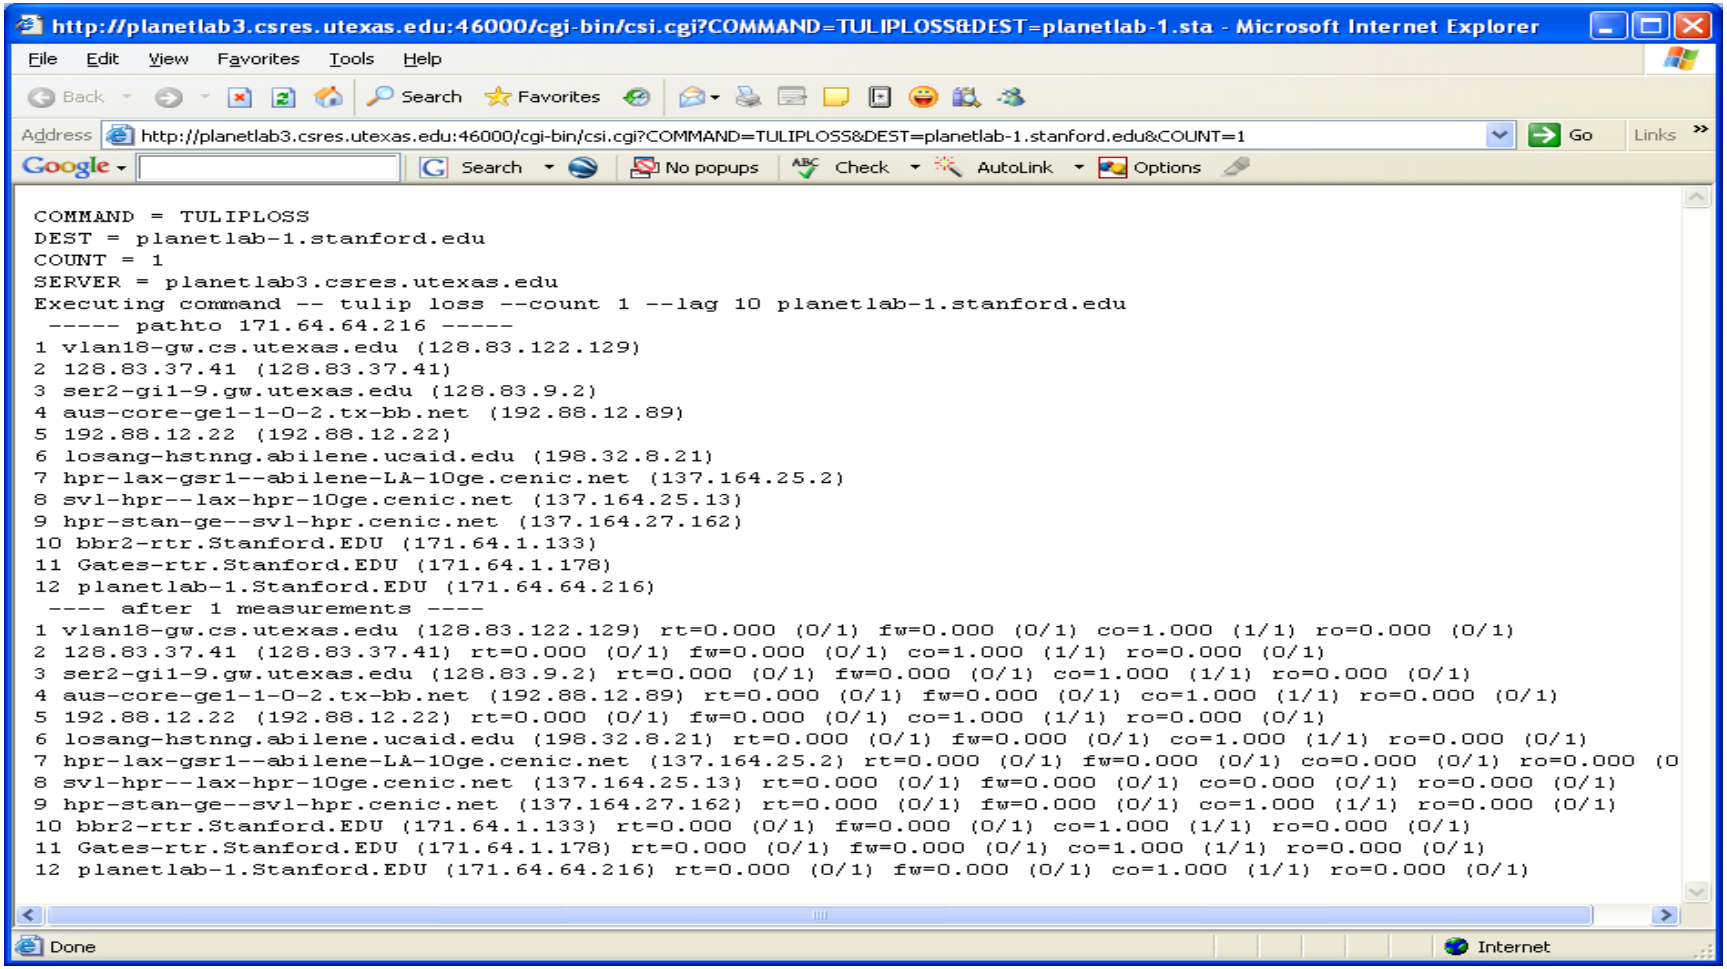

PURDUE UNIVERSITY

hp invent

S^3 Screenshot

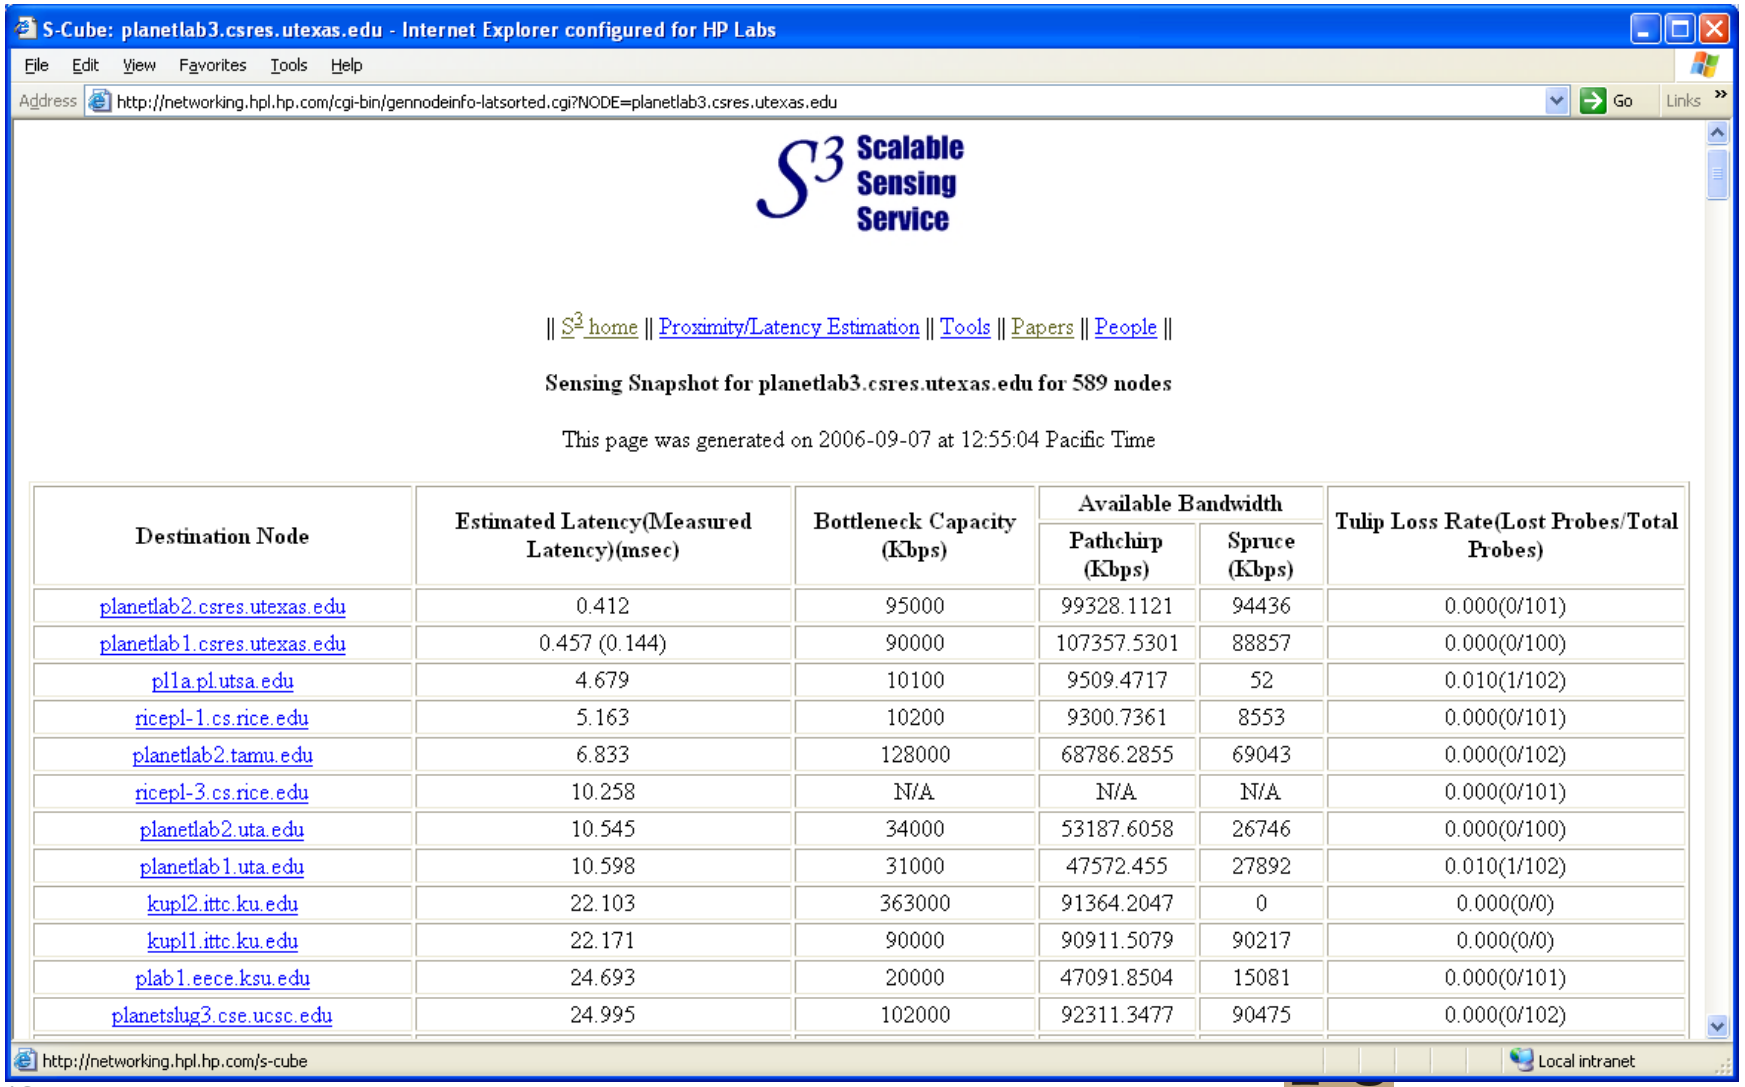

Destination Node	Estimated Latency(Measured Latency)(msec)	Bottleneck Capacity (Kbps)	Available Bandwidth		Tulip Loss Rate(Lost Probes/Total Probes)
			Pathchirp (Kbps)	Spruce (Kbps)	
planetlab2.csres.utexas.edu	0.412	95000	99328.1121	94436	0.000(0/101)
planetlab1.csres.utexas.edu	0.457 (0.144)	90000	107357.5301	88857	0.000(0/100)
pl1a.pl.utsa.edu	4.679	10100	9509.4717	52	0.010(1/102)
ricepl-1.cs.rice.edu	5.163	10200	9300.7361	8553	0.000(0/101)
planetlab2.tamu.edu	6.833	128000	68786.2855	69043	0.000(0/102)
ricepl-3.cs.rice.edu	10.258	N/A	N/A	N/A	0.000(0/101)
planetlab2.uta.edu	10.545	34000	53187.6058	26746	0.000(0/100)
planetlab1.uta.edu	10.598	31000	47572.455	27892	0.010(1/102)
kupl2.ittc.ku.edu	22.103	363000	91364.2047	0	0.000(0/0)
kupl1.ittc.ku.edu	22.171	90000	90911.5079	90217	0.000(0/0)
plab1.eece.ksu.edu	24.693	20000	47091.8504	15081	0.000(0/101)
planetslug3.cse.ucsc.edu	24.995	102000	92311.3477	90475	0.000(0/102)

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
- Backplane for aggregation and queries
- Integration

Needs

- Duration:
  - Continuously running
- Layering:
  - Uses legacy tools running on IP/UDP/TCP/HTTP
- Domain:
  - Currently storage, CPU
- Scale:
  - All
  - Measurement fidelity

Selected Publications

- 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 Measurement 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.

PURDUE UNIVERSITY

hp invent