

Million-node GENI (Seattle): Year 1 Review and Outlook

Justin Cappos
Computer Science and Engineering
University of Washington

Progress in Year 1

Built:

Repy (VM), Node Manager, Seash (shell)

Prototyped:

SeattleGENI (component manager)

Collaborated:

PlanetLab, GpENI, ProtoGENI*, Digital Object Registry*,

Publications:

SIGCSE, PN-ASEE/WCERT, NW-DCSD workshops, 12 talks

Testbed:

1 to > 1000* nodes

Challenges in Year 1

Adoption

- PlanetLab, GpENI

- Universities

- Firewalls

- End users

Robustness

- SeattleGENI (component manager)

- Software updater

Portability

- sockets, file systems, resource accounting

Lots of undergrads, few grad students / senior personnel

ProtoGENI Clearinghouse Integration

Seattle running on ProtoGENI

- Automatic grader for instructors (proof of concept@GEC 5)

Obtaining Million-node GENI resources from ProtoGENI

- ProtoGENI brokers resources on Seattle (plan)
Implement the interface but not the semantics (?)

Storing node information in the Digital Object Registry

- Backup data store

Internet2 progress

Campus: June 1st, 1999

Million-node GENI nodes: N/A

Adding Internet2 nodes at UW (?)

Plans for Outside Experimentation

Educational use:

- Lots of docs for educators / undergraduates / developers
- Used in 6 classes
- Support from NW-DCSD
- **Want more adoption this Fall!**

Research / Development use:

- Adding support software
- Interest from many groups

Spiral 2 Plans

Increasing support for researchers / developers

16 mo: Installer / **end user interface** v0.1, example apps

- add NTP recovery, NAT traversal

20 mo: Repty v0.2, Node manager v0.2

- Performance / resource reassignment / measurement

- **Services**

23 mo: Specification for end-host Clearinghouse API

24 mo: Prototype of end-host Clearinghouse

- Mash-ups
- Identity management

Collaborate with O&M team and security team

Specific contribution to GENI outreach plan