# NSF Workshop on Designing Tools and Curricula for Undergraduate Courses in Distributed Systems

## Boston, Massachusetts
## July 8, 2012

# Welcome!

- Thanks for coming!
- Acknowledgements
  - Organizing Committee
    - Jay Akait and Mark Berman
  - Funding and support
    - NSF
    - Jeff Forbes, Bryan Lyles, and Keith Marzullo

# Logistics

- For reimbursement, please mail receipts to:

  Jeannie Albrecht

  47 Lab Campus Dr

  Williamstown, MA 01267

  by **AUGUST 1, 2012**.
- If you plan to stay for GEC, you must register
- Workshop webpage:
  http://www.cs.williams.edu/~jeannie/nsf-workshop/index.html

# Wireless Access

- SSID: GlobalSuiteMeeting
- Meeting Room: Wireless
- Group/Company Name: GENI2012
- Password: GENI33

# Why Distributed Systems?

- Distributed systems - collections of networked computers that function as single systems
- When designed correctly, distributed systems improve scalability, fault tolerance, response time of Internet applications
- Distributed systems are widely used by many major companies

# Challenges

- Developing, debugging, deploying distributed systems introduces new challenges
  - Can be overwhelming to new developers
- Students benefit from the opportunity to design, implement, deploy, and evaluate ***real*** systems in ***real*** environments
  - Learn techniques for coping with common challenges

- Unfortunately undergraduate curriculums rarely offer Distributed Systems courses
  - Students are not fully prepared for jobs/grad school

# Overcoming Limitations

- Until recently, many colleges and universities did not have the local computing resources required for large-scale experimentation

- Options were mainly emulation and simulation

# Advancements in Technology

- New options have appeared in last decade
- There are now many public testbeds available
  - Developers can "rent" Amazon, Google, and Microsoft resources
  - Shared platforms like GENI provide a variety of deployment options
- Role of MOOCs (Massively Open Online Courses)

# Workshop Goals

- (Re)Define undergraduate distributed systems educational goals
- Integrate new technologies with classic concepts
- Leverage availability of public testbeds to give students hands-on experience with the development of distributed systems

# Workshop Deliverables

- Written report
- Collection and aggregation of course-related resources for educators
- Community building

# Draft Agenda

- 8:00 - 8:30: Continental Breakfast, Registration
- 8:30 - 9:00: Opening Remarks
- 9:00 - 10:30: Session 1 – Platforms
- 10:30 - 10:45: Coffee/Tea Break
- 10:45- 12:15: Session 2 – Educator Experiences
- 12:15 - 1:30: Lunch
- 1:30 - 3:00: Session 3 – Educator Experiences
- 3:00 - 4:00: Integrating classical and modern topics
- 4:00 - 4:15: Coffee/Tea Break
- 4:15 - 5:15: Panel/Round-table discussion and wrap-up

# Opening Remarks

- Keith Marzullo
  - Division Director for the Computer and Network Systems Division (CNS) in the Computer and Information Science and Engineering Directorate (CISE) of the NSF
- Jeff Forbes
  - Program Director for the Education and Workforce Program for the CNS Division in the CISE Directorate of the NSF

# Session 1 – Platforms

- Getting access, managing student accounts, how to use in classroom, etc.
  - Justin Cappos (NYU Poly) – Seattle
  - Gary Wong (Utah) – ProtoGENI
  - Armando Fox (UC Berkeley) – Amazon EC2 and other cloud-based educational technology
  - Jeff Chase (Duke) – ORCA

# Session 2 – Educator Experiences

- Assignments, textbooks, etc., primarily used in small colleges
  - Jeannie Albrecht (Williams College) – Undergrad Distributed Systems
  - Joel Sommers (Colgate University) – Undergrad Computer Networks
  - Tia Newhall (Swarthmore College) – Undergrad Parallel and Distributed Computing
  - Zongming Fei (Kentucky) – Undergrad Networks and Distributed Operating Systems / Calvin College Emulab

# Session 3 – Educator Experiences

- Assignments, textbooks, etc., primarily used in large universities
  - Armando Fox (UC Berkeley) – Massively open online courses (MOOCs)
  - Sonia Fahmy / Ethan Blanton (Purdue) – GENI-based classroom exercises
  - Anish Arora (The Ohio State) – Projects designed for local and remote testbeds
  - Mark Berman (BBN/GENI) – Sample assignments designed for GENI

# Common Themes

- Breadth over depth; introduce students to a variety of technologies
- "Learn by doing"
- Debugging in wide-area is challenging
- Read/write papers in addition to code
- Expose students to low level details (i.e., sockets)
- Open-ended/Independent final project
- "Short"/well-defined labs early in semester
- Experimentation and analysis of systems

# Common Themes

- Specific skills/concepts:
  - Makefiles, revision control, debugging tools, scripting, latex, gnuplot, using threads, sockets
- Combination of textbook and research papers
- External resources critically valuable (for small and large universities)
- Give students experience with “real” resources/environments

# Discussion

- Any remaining unanswered questions?
- What is the best way to disseminate the material we've discussed today?
- Best way to share materials more generally?
- Final thoughts?