

Software and Programming Support for Computational Research @ Princeton University

Ian A. Cosden, Ph.D.

Manager, HPC Software Engineering and Performance Tuning
Research Computing, OIT, Princeton University

icosden@princeton.edu

May 1, 2017

Overview

- What is the Research Computing department at Princeton?
- Who uses our services?
- Software and programming support
 - Current approach
 - Future direction


What is Research Computing?

- One team (24 people) - joint effort
 - Office information Technology (OIT)
 - Princeton Institute for Computational Science and Engineering (PICSciE)
- Work closely with Academic Departments, CDH, CSML
- Goal: Enable and accelerate faculty-driven computational research
- Hardware
 - 5 clusters ~28k cores, 320 GPUs
 - 4.5+ PB parallel filesystem
 - Many smaller systems and servers
- Groups/Focus
 - Systems & Storage
 - Distributed Computing Support
 - Advanced Networking
 - *Software and Programming*

<http://www.princeton.edu/researchcomputing>

Software and Programming Group

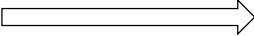
- Programmers, Software Engineers, Visualization Specialists
 - 9 people (7.5 FTE) and growing
- “Customer” Demographics
 - Over **750** *active* users on main HPC clusters from **21** departments
 - Constantly adding new users
 - All levels of experience

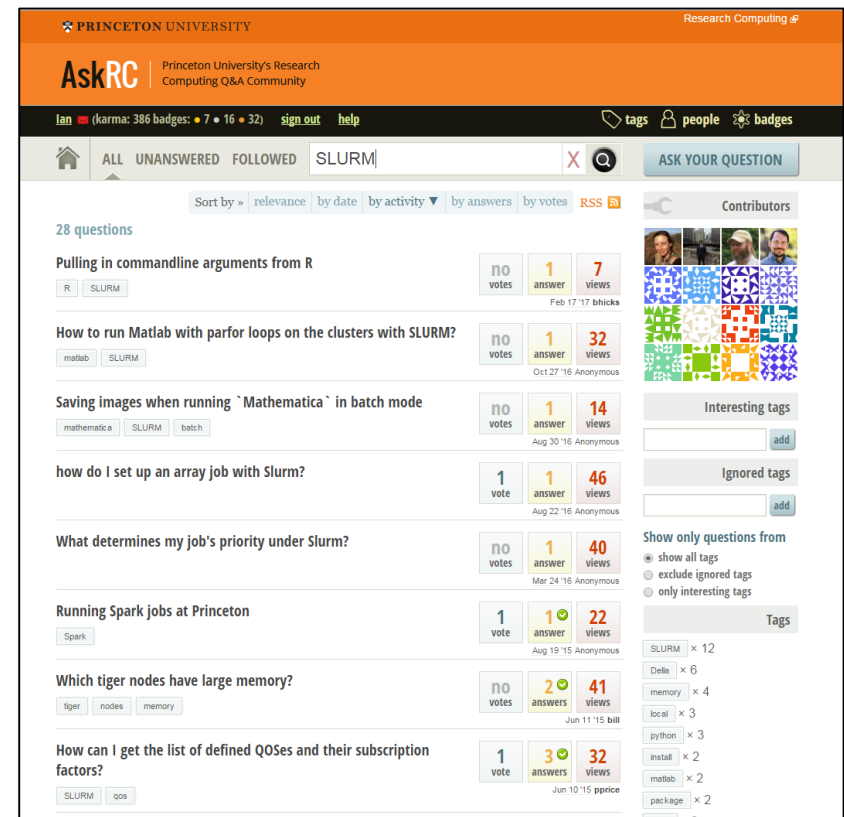


Focus of Software and Programming Support

1. Getting Started & Problem Solving
 - Direct user interaction & Documentation
2. Improving Productivity
 - User interaction, products/services
3. Training/Education
 - Classroom setting
4. Software Engineering
 - Solving research problems

1. Getting Started

- New users come every: year, semester, month, week
- Linux and HPC basics
- Building and compiling code
- Programming basics
- Version control
- Email help (cses@princeton.edu)
- Help sessions (twice a week)
- Website & FAQs
- AskRC.princeton.edu 
- Coming soon: RC blog



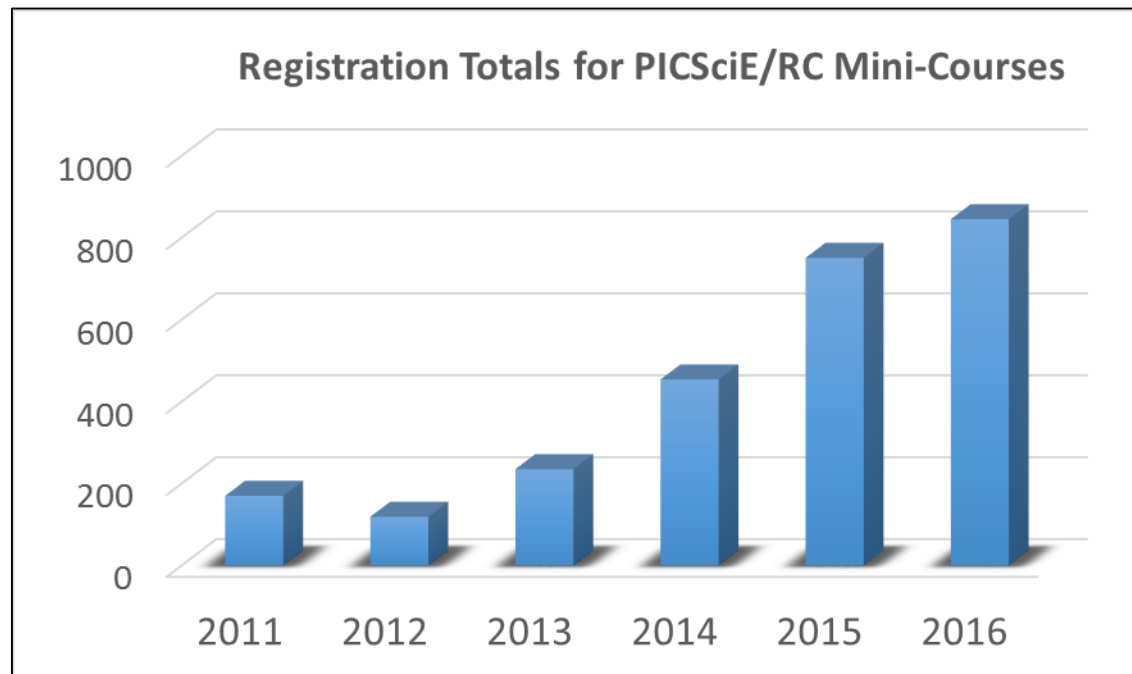
2. Improving Productivity

- Debugging
- Continuous Integration
- Performance Tuning & Optimization
- Parallelization
- Big Data & Data Science
- Visualization




3. Training/Education

- 11 Geographic and Information Systems (GIS) workshops
- Currently 11 unique programming mini-courses/workshops
 - Developed and taught by technical staff
 - Offered each semester
 - From 2 hours to 2 days



Academic Research Software Challenges

- Increasingly difficult for academic researchers to stay current
 - Hardware
 - Software



- Growing gap between skills and expertise needed to produce quality code
 - Grad students and postdocs = constant turnover
- Traditional academic reward structure doesn't prioritize software development

4. Research Software Engineers (RSE)

- Newest endeavor
 - Create permanent staff RSEs
- Combine:
 - Domain knowledge (PhD level)
 - Computational expertise
 - Software development
- Awarded based on proposal
 - Geosciences/Physics
 - Neuroscience
 - Astrophysics
 - Statistics and Machine Learning

Future Directions and Open Questions

- User base is expanding
 - Increased demand for:
 - Data science
 - Training
 - Define:
 - Role of RSE and programmers in academic research
 - Career path
 - Establish meaningful collaborations & develop community
-
- More info & contact Research Computing:
 - <http://www.princeton.edu/researchcomputing>
 - <http://www.princeton.edu/AskRC>
 - cses@princeton.edu