

Thursday Seminar - HPC in the Cloud with Penguin Computing on Demand ("POD")

*

Penguin Computing on Demand (or POD) is Penguin's public cloud HPC resource. POD has been in commercial service for 18-months and has run over 8 million jobs for customers located throughout the world. Users range from individual graduate students to departments in Fortune 1000 companies, and who run a wide range of private codes, open source applications and commercial products. This seminar includes a general introduction of POD, presents how to access and use POD, reviews several case studies of current and previous users, and concludes with a hands-on demonstration of staging data and submitting jobs using "PODShell", a new tool developed in conjunction with Caltech.

*Introduction to POD

*

- What is POD
- How POD compares/contrasts to Amazon's EC2
- The new EDU POD
- Setting up your account
- Transferring data to/from POD
- Setting up your environment
- Running jobs
- GPU computing on POD
- Setting up a VPN
- Running commercial software
- Checking your usage and storage (PODReport)
- POD's fee structure

Case Studies

- Case Study #1: Graduate Student
- Case Study #2: Local Web Data Processing Company (shipping data from multiple international locations)
- Case Study #3: Fortune 1000 Company - Genetic Sequencing
- Case Study #4: TBA

POD Shell (live demonstrations)

- What is POD Shell
- Submitting Jobs
 - Switching between TORQUE and SGE
- Data Staging
- Scripting / workflow
- Security
- Report examples