

# **Campus Grid Technology**

Derek Weitzel

University of Nebraska – Lincoln  
Holland Computing Center (HCC)  
**Home of the 2012 OSG AHM!**

# Motivation for a Campus Grid

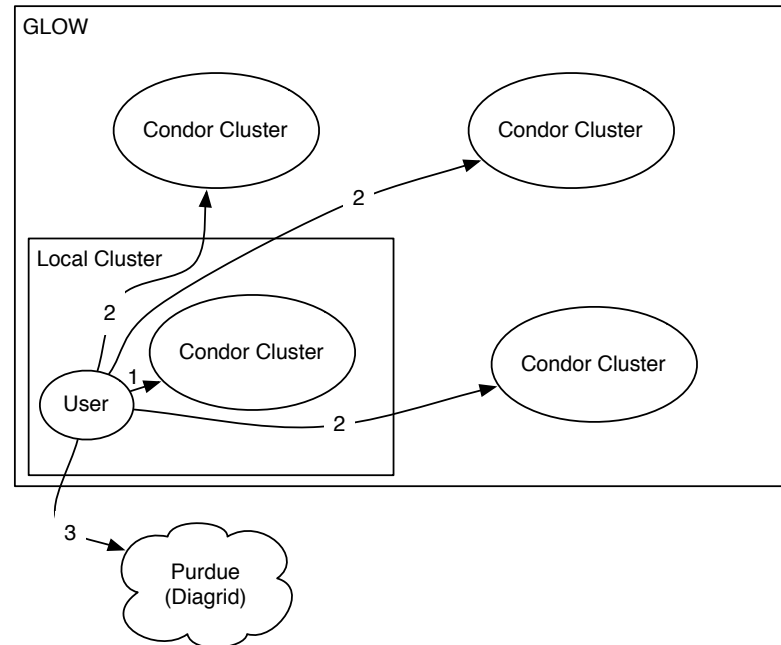
---

- Increase available resources by distributing jobs around the campus.
- Offload single core jobs to idle resources, making room for specialized (MPI) jobs.
- Independence from resource failures.

# Campus Grids you may know

- GLOW

- Condor is Everywhere!
- AFS for data



- Purdue

- Condor & PBS on nodes
- Shared Home directories, Large clusters with share file systems, Condor file transfer between departmental clusters.

# What changes do we want?

---

- Prefer not installing condor on every worker node. Less intrusive for sysadmins.
- PBS and Condor should coordinate job scheduling. We don't want PBS to kill condor jobs if it doesn't have to.

# Campus Grid Technology

---

- Condor + Glideins
- Condor will handle job submission, job execution, and file transfer
- Introducing Campus Grid Factory to handle on-demand glidein submissions.



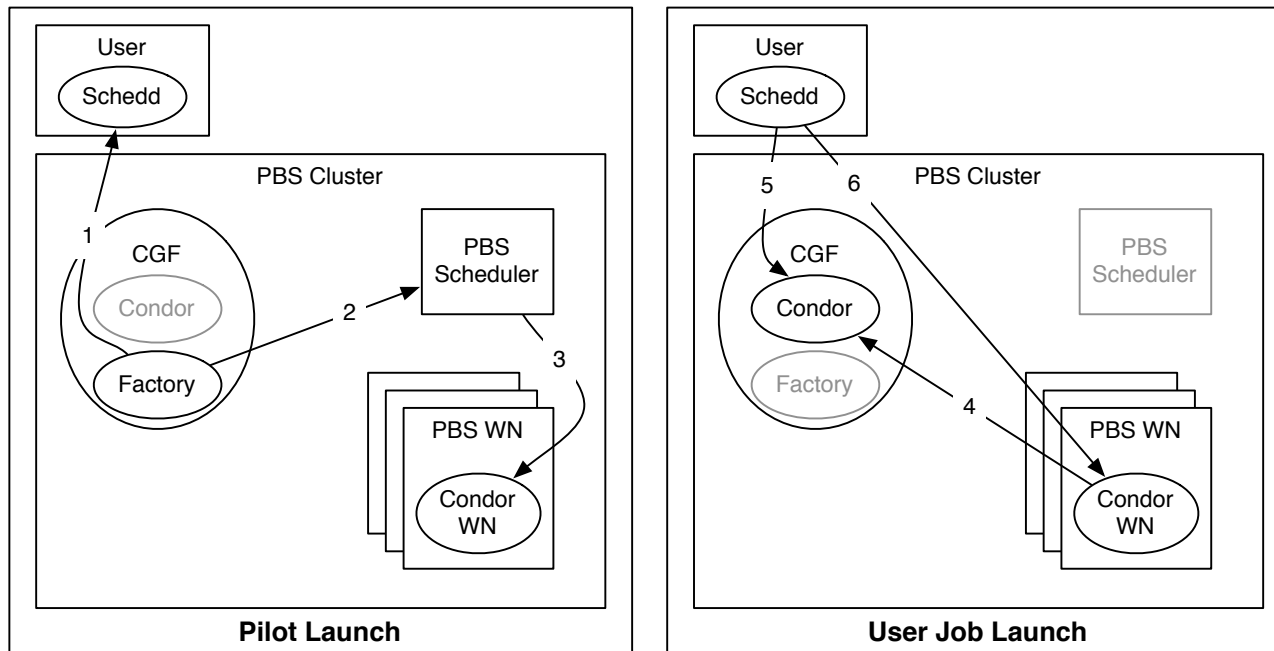
# Campus Grid Factory

---

- Submit Condor worker nodes as jobs to PBS
- Uses GlideinWMS libraries to manage glideins on the cluster.
- Condor + GlideinWMS + BLAHP + Custom Glue
- Provides virtual Condor pool for outside clients for Flocking.

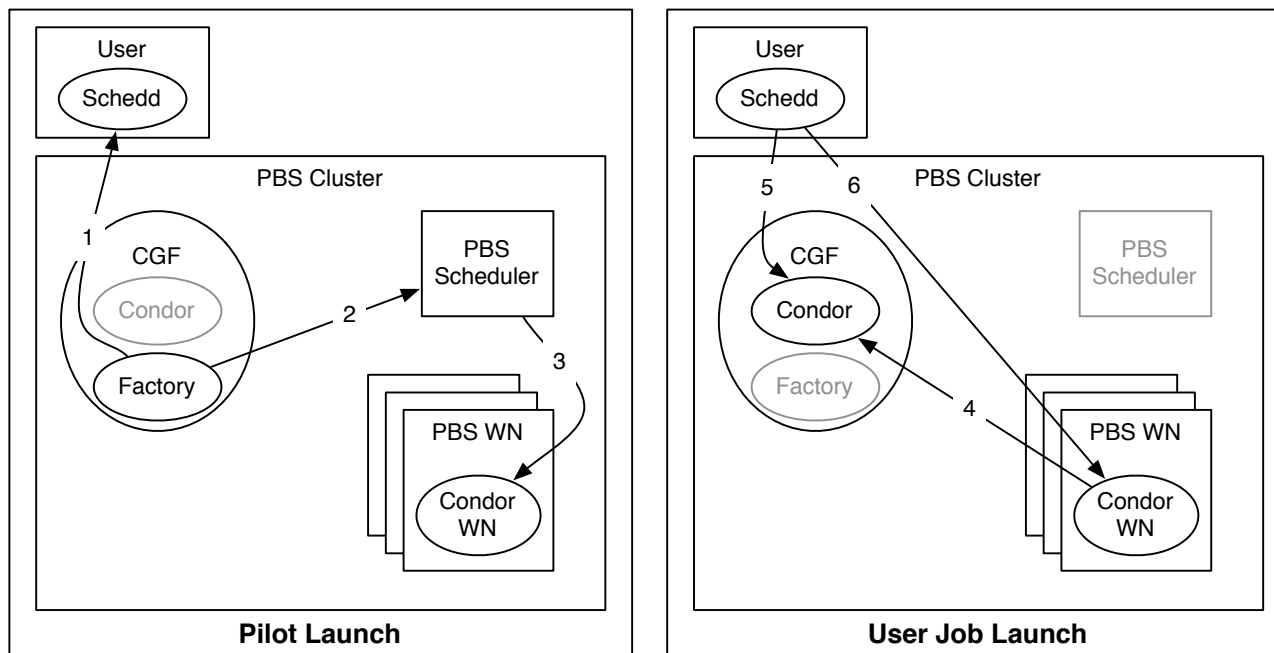
# Campus Factory Operation

1. Factory queries user queue
2. Submit to PBS (through Condor BLAHP)
3. PBS Starts the Glidein job.



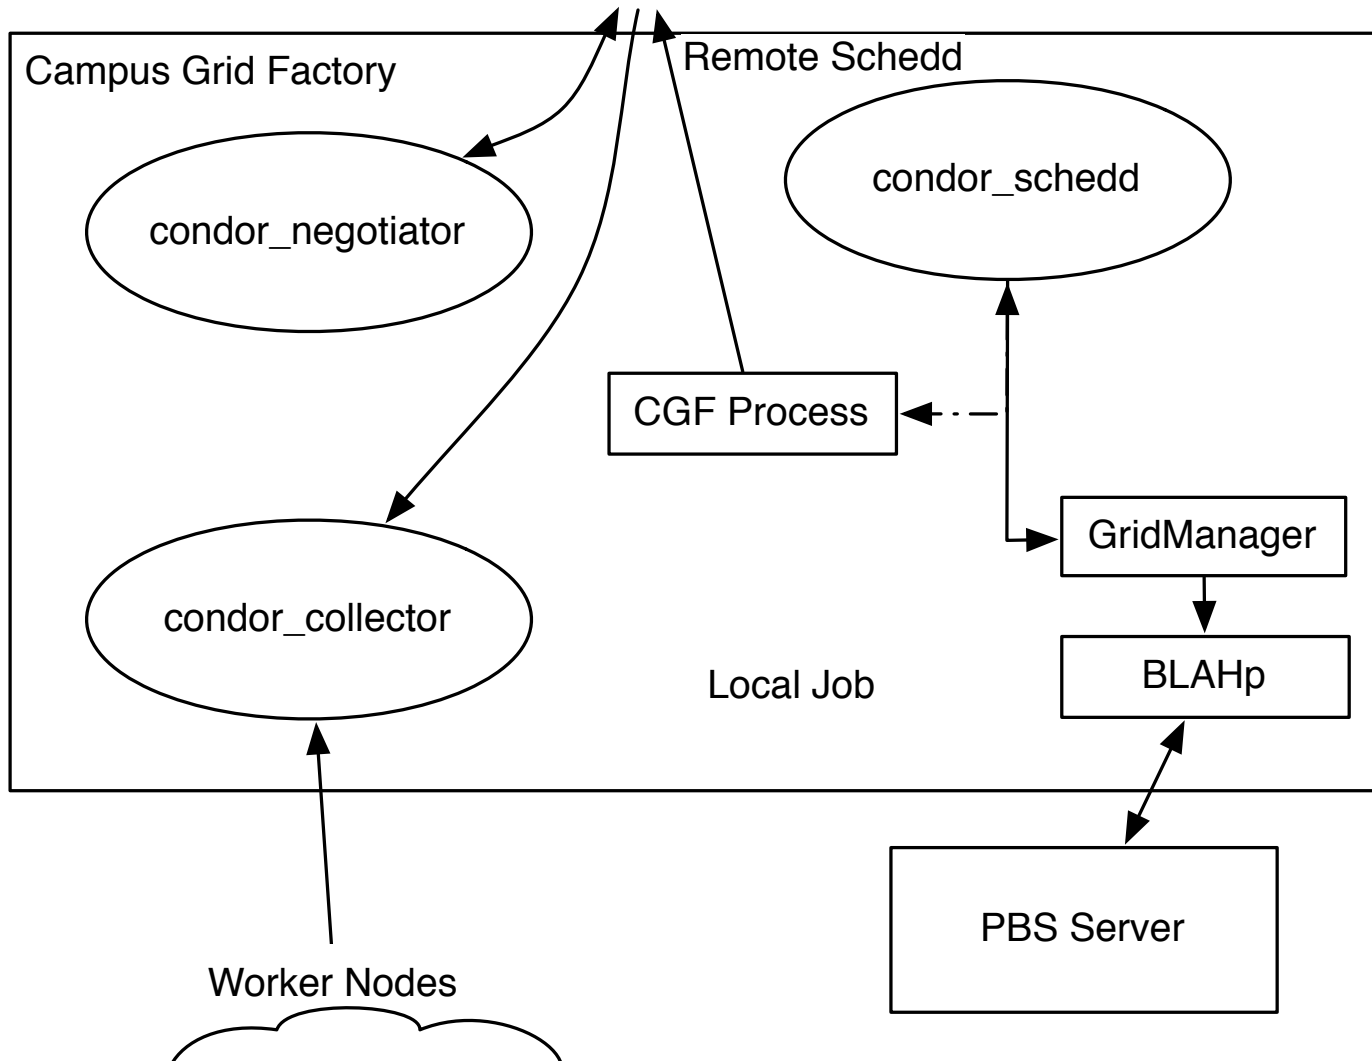
# Campus Factory Operation

4. Glidein reports to the factory's collector
5. User queue requests idle slots from factory
6. Condor negotiates and starts the job.





# Campus Grid Internals



# Advantages of the Campus Factory

---

- User is presented with an uniform Condor interface to resources.
- Can create overlay network on any resource Condor (BLAHP) can submit to (PBS, LSF,...)
- Uses well established technologies, Condor, BLAHP, and ideas, Condor, flocking.

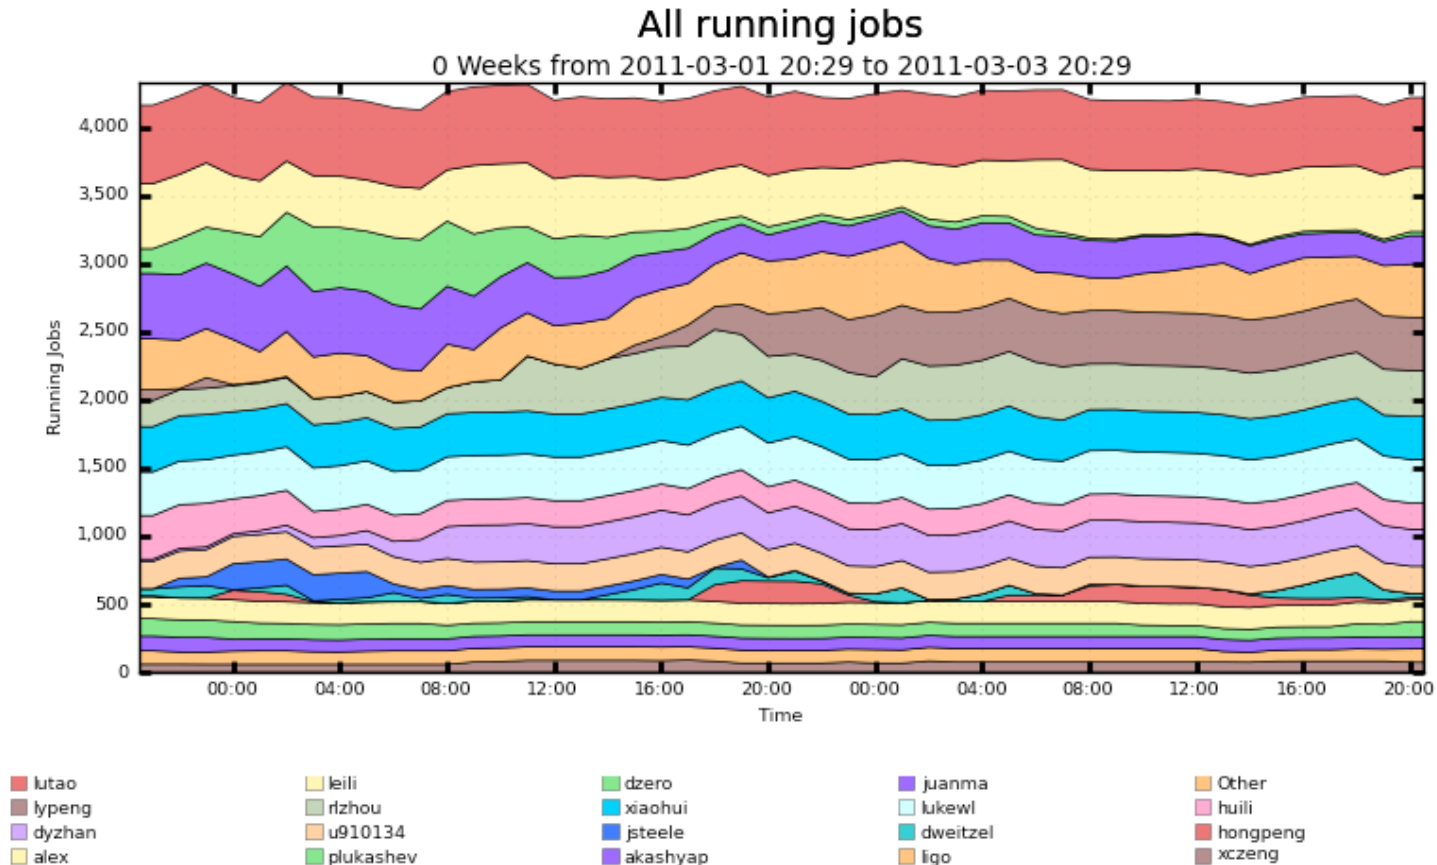
# Advantages for the Local Scheduler

---

- Allows PBS to know and account for outside jobs.
- Can co-schedule with local user priorities.
- PBS can preempt grid jobs for local jobs.

# Getting outside the Campus

- Campus resources are busy:



Maximum: 4,326 , Minimum: 0.00 , Average: 4,138 , Current: 4,216

Running jobs on Firefly. ~4300 cores

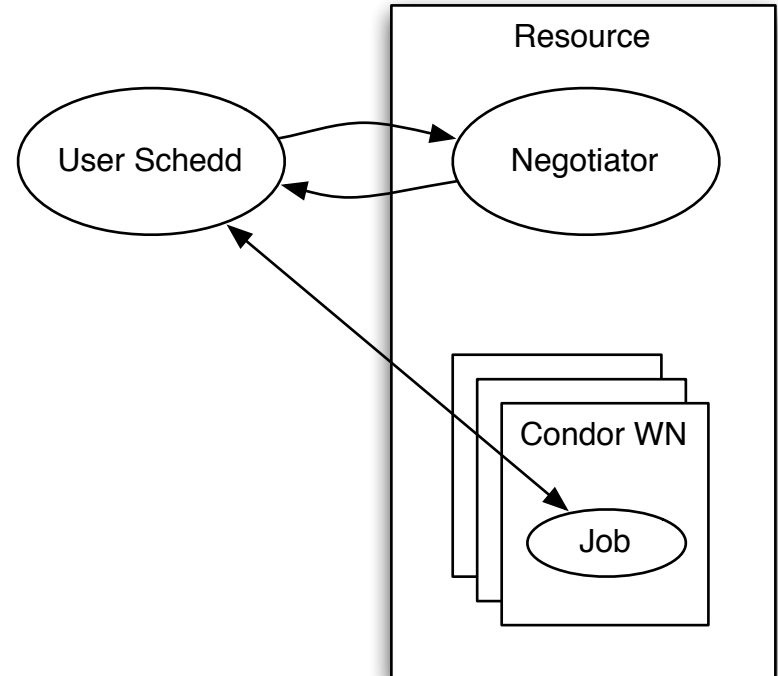
# Getting outside the Campus

---

- Users will use whatever resources are available, usually to the max.
  - Have you had a user submit 50,000 - 1 core jobs to PBS? We have!
- There are resources available out there, just need to get there.
- Users want a smooth transition to other resources.

# Getting outside the Campus: Flocking

- User Schedd advertises idle jobs to resource.
- Resource matches idle jobs with job slots
- User schedd transfers data directly to worker node.



# Getting outside the Campus: Flocking

---

- Flocking
  - Need to have Condor on the other side.
  - Need to have mutual trust relationships. Not necessarily x509.
  - Nebraska is part of Diagrid, large Condor flock led by Purdue.



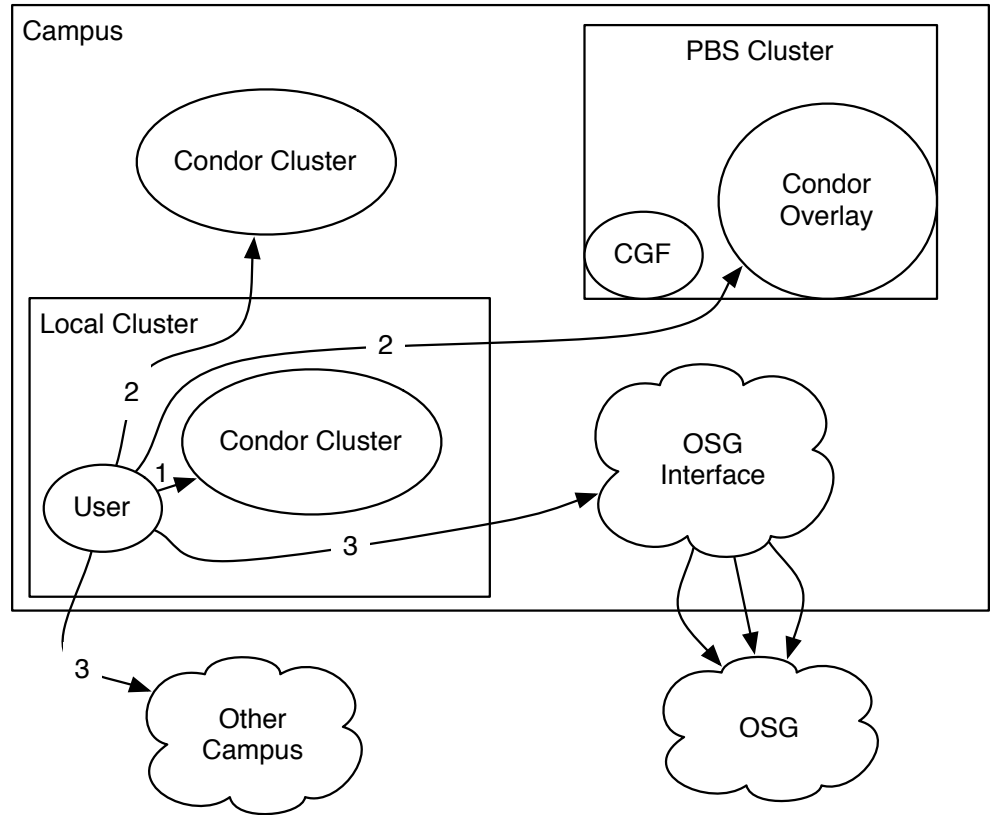
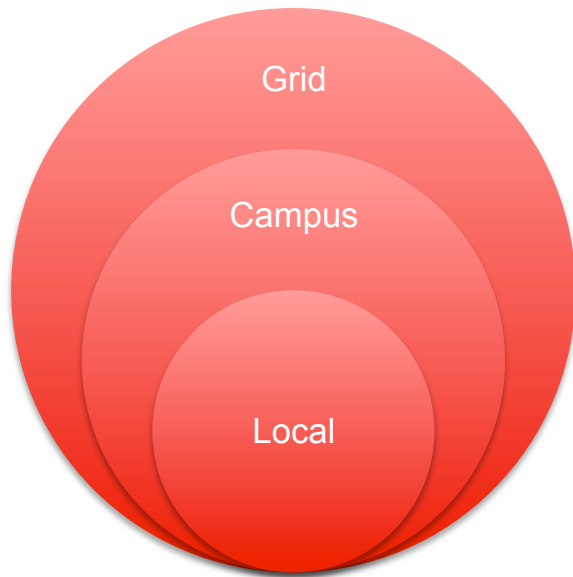
# Getting outside the Campus: OSG Interface

---

- Expand further with OSG Production Grid
- GlideinWMS
  - GlideinWMS Frontend provides virtual Condor pool of grid worker nodes. Campus submit hosts can Flock to the GlideinWMS Frontend.
  - Submission at the Frontend can also flock to campus resources.
  - We don't have to submit globus jobs; all Condor vanilla universe.
    - Factory is at UCSD submit globus jobs for us.




# Ever-increasing circle of resources



Me, my friends and everyone else

# Campus Grid at Nebraska

- Prairiefire PBS/ Condor (Like Purdue)
- Firefly – Only PBS
- GlideinWMS interface to OSG
- Flock to Purdue






# Usage of the Campus Grid

- Bridging to Purdue provides a lot of time
- OSG resources (through GlideinWMS)

Hours by resource type  
28 Days from 2011-02-04 to 2011-03-04





# Bonus: HTPC on Campus

---

- Specifying the right requirements, exactly the same as regular jobs.
- Flocking HTPC to Campus factory exactly the same, just specify requirements.

# Acknowledgements

---

- Brian Bockelman (Technical consult)
- Dan Fraser (Campus Grids Lead)
- David Swanson (Advisor)
- Igor Sfiligoi (GlideinWMS)
- Jeff Dost (UCSD Factory Support)