

Subject: FW: MPI on the PPA cluster 2013/10/31 meetings notes
From: "Gruber, Shirley" <shirley@slac.stanford.edu>
Date: 10/31/2013 12:22 PM
To: "Marshall, Stuart L." <marshall@slac.stanford.edu>

From: <Gruber>, Shirley Gruber <shirley@slac.stanford.edu>
Date: Thursday, October 31, 2013 12:14 PM
To: Shirley Gruber <shirley@slac.stanford.edu>
Cc: Adeyemi Adesanya <yemi@slac.stanford.edu>, Renata Dart <renata@slac.stanford.edu>
Subject: MPI on the PPA cluster 2013/10/31 meetings notes

MPI on the PPA cluster meeting 2013/10/31 meetings notes

Attendees: Yemi Adesanya, Renata Dart, Steffen Luitz, Tom Glanzman, Richard Mount, Wei Yang, Richard Dubois, Shirley Gruber

Richard M. ATLAS to use supercomputers per DOE.

Stuart – goal is everyone gets to run, including large MPI jobs. MPI jobs run best when they get an entire host. Possible options:

1. aggregate single slot jobs on fewest hosts
 - o Try JOB_ACCEPT_INTERVAL=0 on xlong queue
2. Limit number of running jobs per queue per host group
3. Priority adjustment for large jobs
4. Job migration for long ss-jobs when occupancy/host is low
5. Backfill
6. Cpu affinity & cgroups, etc.
7. Re-architect whole thing - move to something like multi-cluster (maybe on pinto cluster or orange cluster)

Issues:

- resource constraints
- No test environment – could use some bullets

Planned order of tasks: 1 5 2 3

1. 1 Try JOB_ACCEPT_INTERVAL=0 - apply to atlasq this afternoon
2. 5 backfill
3. 2 Limit number of running jobs per queue per host group
4. 3 Priority adjustment for large jobs

Richard Dubois: How do we measure success? When do we decide we have to take an alternative approach (partition...)?

Wei: throughput of MPI easy to measure, latency harder to measure

Stuart: people willing to wait a long time to get a 1024 core job to run, 64 core jobs want shorter time, perhaps backfill

Realistic date to set for significant improvement would be something like December 1.

Stuart: It will take a month to do just part of this plan.

Richard Dubois: take a portion the section added to add-users – give to MPI 20K fairshares.

- reduce all users fairshare by 10K

Stuart: new norm for LSF is run an instance of LSF and have multiple instances of LSF running jobs

- 6 person-months of time

IBM/Platform support:

- escalate to get SME for a longer period of time to help achieve the appropriate outcome

Can Tom take advantage of backfill by setting run limits?

- may not work well with jobs in Fermi pipeline