Migrate and run HTCondor jobs to Slurm cluster via container

HTCondor Week 2018

Jingyan Shi
On behalf of scheduling group
of Computing Center, IHEP

Outline

1. Introduction and Motivation
2. Design and development
3. Status and next step
4. Summary

Introduction to IHEP Computing Platform

- IHEP: Institute of High Energy Physics, Chinese Academy of Sciences.
- IHEP Computing Center: network, computing and storage services provider to the HEP experiments.

BESIII (Beijing Spectrometer III at BEPCII)

DYB (Daya Bay Reactor Neutrino Experiment)

JUNO (Jiangmen Underground Neutrino Observatory)

YBJ (Tibet-ASgamma ARGO-YBJ Experiments)

Large High Altitude Air Shower Observatory

Hard X-Ray Moderate Telescope

CMS

ATLAS EXPERIMENT

CEPC

Batch Systems at IHEP

- Computing clusters
  - HTCondor Cluster for HTC computing
    - ~10,000 CPU cores
    - Avg. 100,000 jobs/day
  - Slurm Cluster for HPC Computing
    - ~2800 CPU core + 8 GPU cards
    - 122TFLOPs
- Login nodes
  - 32 login nodes shared by all users
- A unified job management tool for two clusters

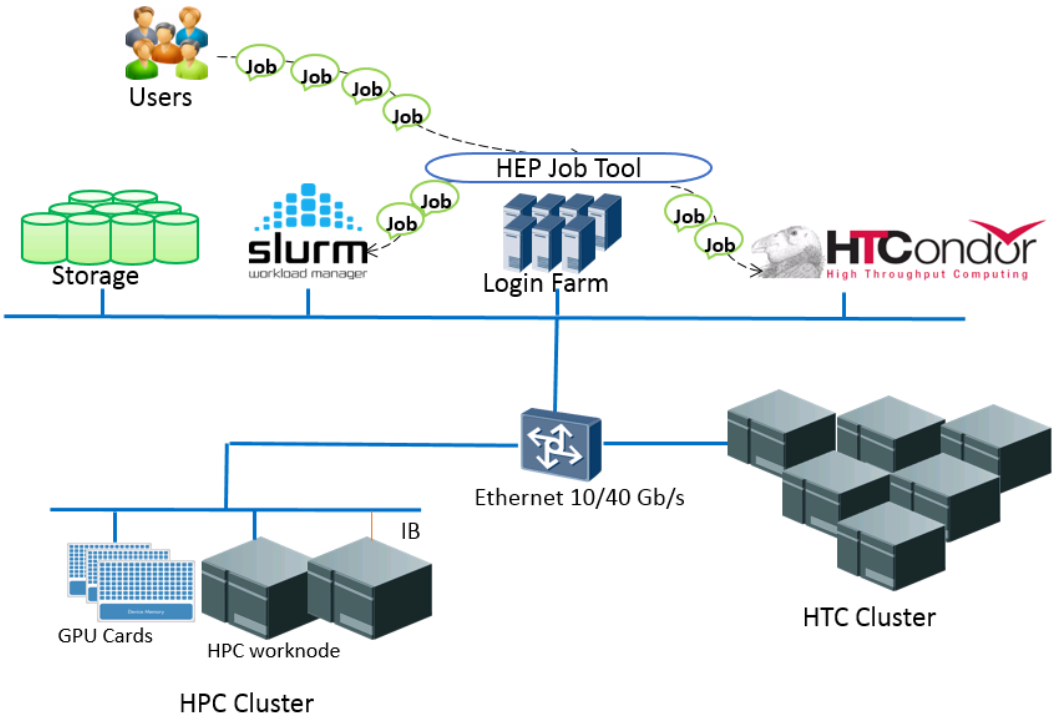

Computing -- HTCondor Cluster

- Resources
  - funded from different experiments
  - shared by all experiments
- HTCondor version: 8.4.11
- Job slots
  - Exclusive job slots: match jobs from the owner
  - Shared job slots: match jobs from all users
- Sharing strategy
  - Jobs are preferred to run on exclusive slots
  - Shared slots are kept for busy experiments
  - Group quota to each experiment, which can be exceeded if there are free shared jobs slots
  - Shared slots are matched according to the relative ratio of quota among the busy groups.
- Job Slots utilization is quite high -- > busy cluster
  - Job slots utilization: ~90%

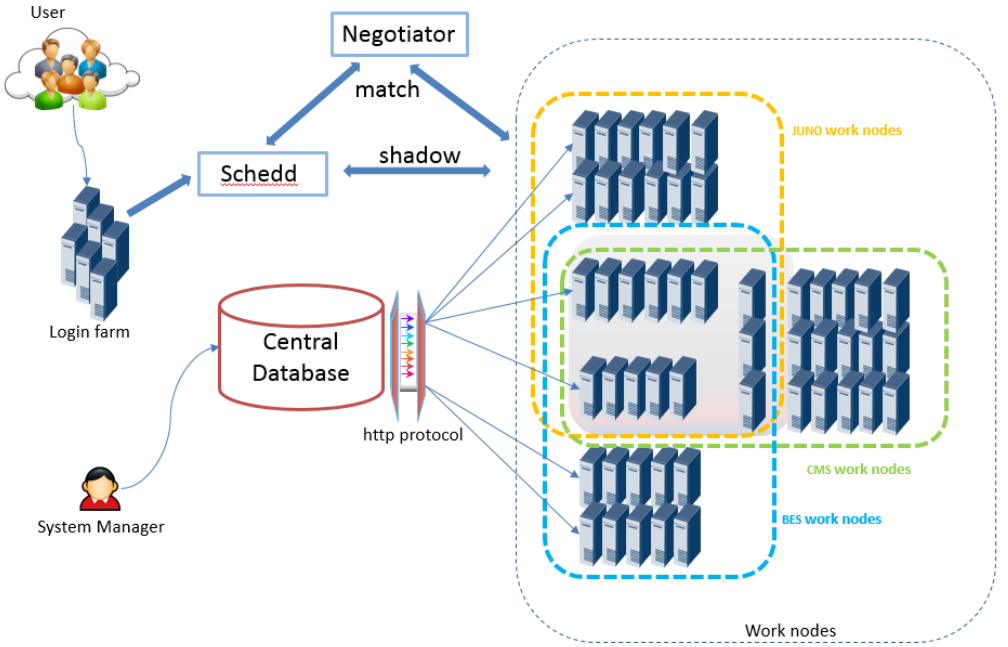

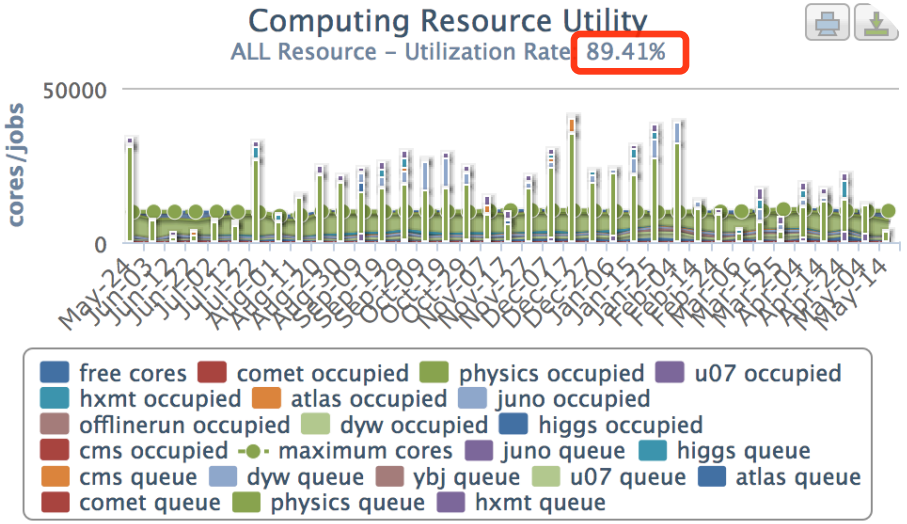

Computing -- Slurm Cluster

- Resources
  - 1 master node
  - 1 accounting & monitoring node
  - 125 work nodes: 2,808 CPU cores + 8 GPU cards
- Lower utilization: free job slots most of the time
  - Utilization: ~50%
  - Jobs (2018.1~2018.4)
    - Jobs: ~5300
    - CPU hours: ~3 million
- GPU servers procurement
  - NVIDIA Tesla V100, 1 PFLOPs (single precision)
  - Procurement in process, expected to be done in 2018.

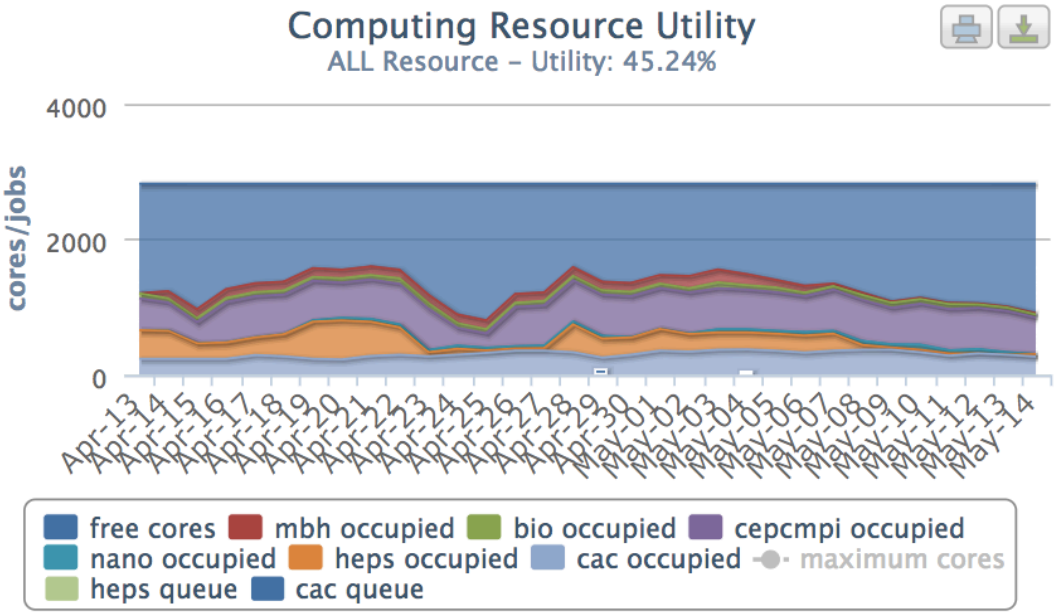

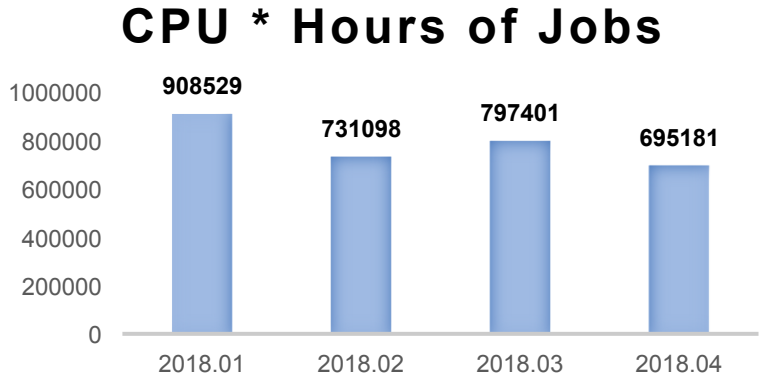

A Unified Job Toolkit -- HepJob

- Targets
  - To provide a unified method for users to submit and manage jobs
  - Simple user interfaces
  - To help administers achieving new scheduling policies: new experiment, Container, high priority job, etc.
- Implementation
  - Based on Python
  - Works with IHEP specific environments
    - Server names, group names …
    - Standard job templates for each group/ experiment
    - Unified user interfaces for two clusters

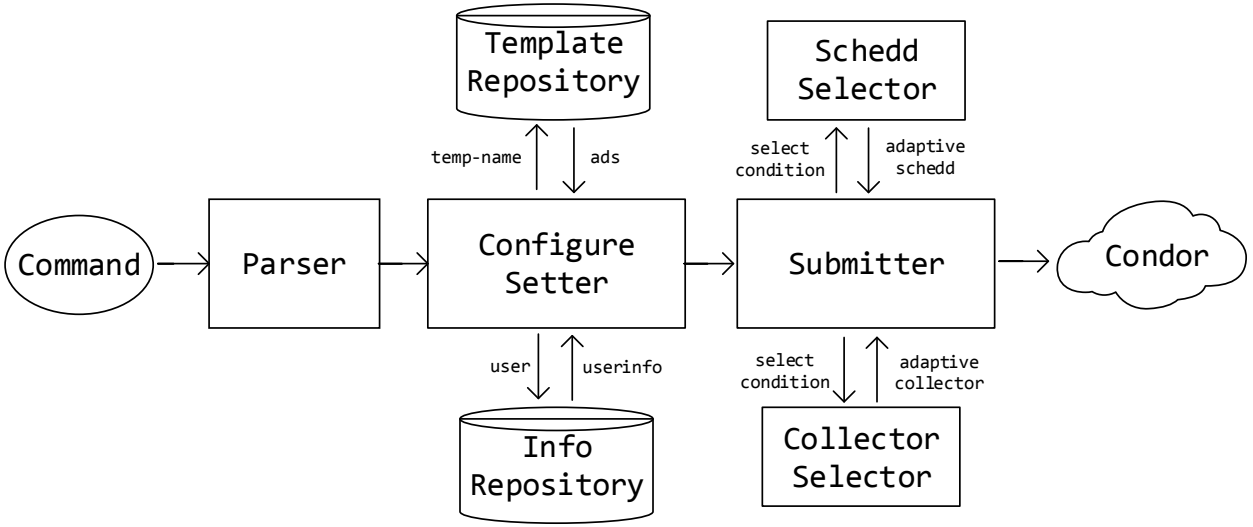

Modules of the the HepJob toolkit

Docker Jobs running on the HTCondor Cluster

- Container jobs
  - SL6 is the host OS running on physical machines
  - SL7 requirements from users
  - Some experiments don’t want to expose their files to others
- Docker images created to fit versatile requirements
  - Image saved under AFS with ACL accesses
- Jobs submitted with specific options, e.g.:
  - hep_sub −os sl7 −g juno juno_script.sh

HTCondor jobs running on the SLURM Cluster

- HTCondor busy queue vs. Slurm free slots
  - HTCondor jobs → Slurm job slots
- Which types of jobs to be migrated
  - Jobs queues at the end of long queue
  - Users agree to get migrated during job submission
    - Risk acknowledgement: jobs may get preempted and re-queued
- How
  - Add extra job attributes to queuing jobs
  - Start “startd” daemon on Slurm work nodes
  - “startd” is added to HTCondor resources
  - Jobs are scheduled by HTCondor to the startd slots at SLURM
- Status : under development

Design

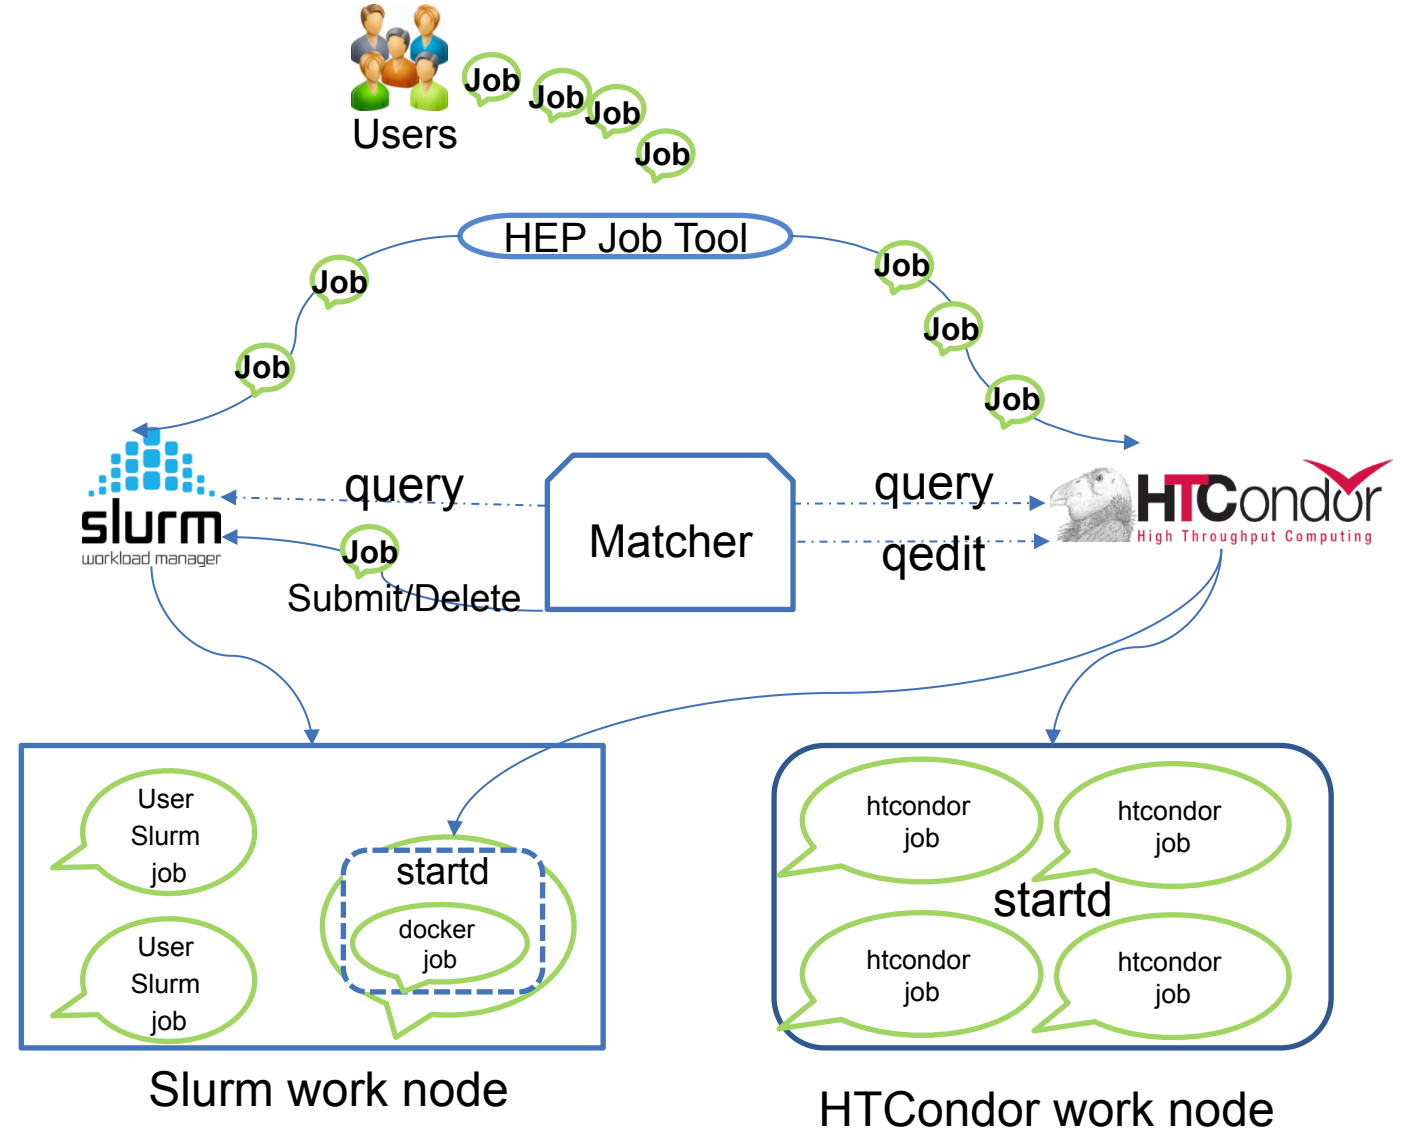
Users
Job
Job
Job
Job
HEP Job Tool
Job
Job
Job
Job
Job
slurm
workload manager
query
Job
Submit/Delete
Matcher
query
qedit
HTCondor
High Throughput Computing
User
Slurm
job
User
Slurm
job
startd
docker
job
Slurm work node
htcondor
job
htcondor
job
startd
htcondor
job
htcondor
job
HTCondor work node

Matcher

- Matcher:
  - Python implementation
  - HTCondor and Slurm client
    - Submit and delete Slurm job
    - Add attributes to HTCondor jobs
  - Function:
    - Query free slurm job slots and htcondor queuing job
    - Matching
    - Add extra attributes to the job to be migrated
    - Submit slurm jobs
    - Delete slurm jobs when preemption is necessary
  - Run with crontab

Others

- Migrated job selection
  - User agree the job could be migrated
    - hep_sub −hpc −g juno job.sh
  - Select from the end part of the queue
  - More selection policies would be added
- Docker Image
  - SL6 and SL7 docker images created
- Slurm job script
  - Start condor startd daemon under user “condor”
- Startd running in slurm slots: accept jobs with dedicated attributes
- HTCondor and Slurm scheduler: transparent to the schedulers

Preemption

- Step:
  - New Slurm job coming
  - Detected by the matcher
  - Delete Slurm jobs running startd
  - HTCondor jobs get re-queued
- Job deleted:
  - Latest jobs started would be deleted first

Test

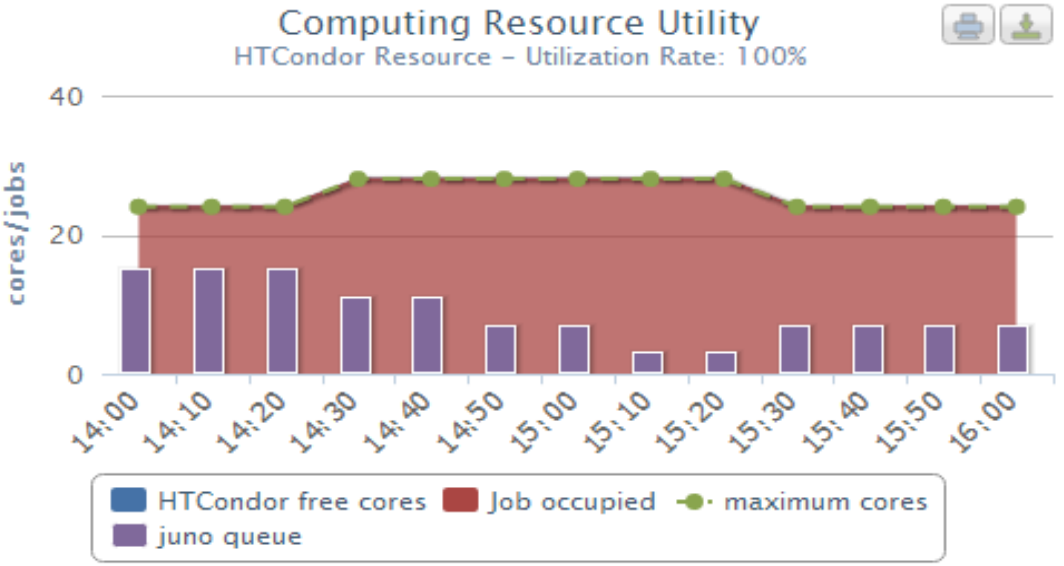

HTCondor cluster

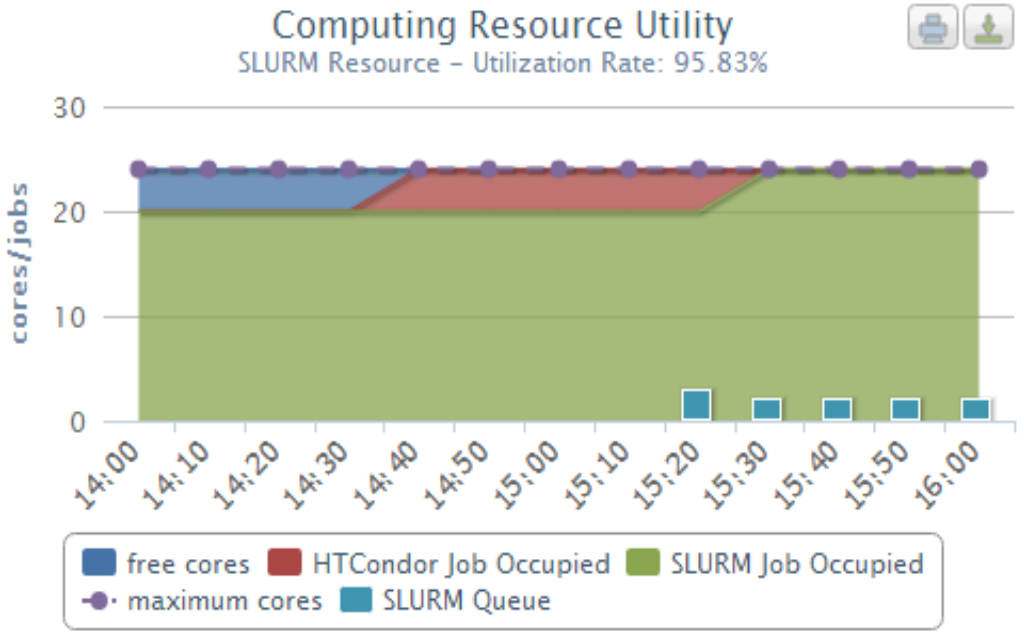

Slurm Cluster

Next Step

- More migration policies are coming.
- Prepare for the production systems.
- Consider to migrate HTCondor jobs to remote small sites.

Thanks &
Comments ?