

Slurm at Rensselaer

Tim Wickberg

Slurm Overview

Originally "Simple Linux Utility for Resource Management" - SLURM

Has grown in to a vastly more capable workload scheduler

Plugin architecture keeps core scheduling logic efficient + simple
While no longer "simple", code is still easily accessible

Open source (GPL v2)

Scales from small lab clusters, up to the largest systems in the world
Linux machines through Blue Gene/L,P,Q, Cray

RPI Background

CCNI launched in 2007 with a 16-rack Blue Gene/L at #7 on the Top500.

BG/L complemented by an x86 cluster

Now includes Blue Gene/Q

Over 100 projects, 800 users from academic, government and industry

Our shift to Slurm

CCNI originally built using a certain vendor's proprietary workload manager

We switched, and have no zero regrets

Center has grown alongside Slurm, adding new features as they are released

Now have a support contract through SchedMD if we get stuck

How we schedule

Separate clusters for BG/L, BG/Q, Opteron systems

But, single job accounting system!

Multifactor fair-share scheduling

Backfill scheduling

Research Platform

Data-Staging research platform - the "RAMDISK Storage Accelerator"

8TB RAM dynamically carved up and served to compute nodes as a filesystem

Research involved in both using this, and effectively allocating resources to jobs