Cray XMT Resource Management

- Tim Carlson
- Pacific Northwest National Laboratory

Pacific Northwest
NATIONAL LABORATORY

Problem

- Cray XMT has limited resource management
  - mtarun assigns tasks to CPUs but has limited controls
  - Leads to over subscribed system
  - Users clashing with each other
  - Variable run times as resources are taxed
  - Mimimal process accounting.
    - Who is making use of the system?

Resource Managers

- SLURM
  - Simple Linux Utility for Resource Management
- Torque
  - Open source resource manager derived from PBS
- LSF
  - Load Sharing Facility
  - Commercial product from Platform

Dilemma

- XMT nodes not “traditional” compute nodes
- Resource managers usually required to know the state of the resources
- Difficult to get resource manager daemon processes running on XMT “nodes”

SLURM

- Provides a mechanism for managing resources on the XMT via “virtual” cluster
- Read the SLURM FAQ
  - https://computing.llnl.gov/linux/slurm/faq.html#multi_slurmd
- Define a number of virtual nodes equal to the number of processors
- Add Maui for sophisticated scheduling