Demand Response for High Performance Computing

Y. Qin[1], K. Fernsler[1], M. Jennings[1], G. Kurtzer[1], G. Jung[1]
V. Ganti[2], G. Ghatikar[2]

[1]High Performance Computing Services (HPCS)
[2]Demand Response Research Center (DRRC)
Lawrence Berkeley National Laboratory

Data Center Efficiency Summit @ SVLG, 2012

BERKELEY LAB
Lawrence Berkeley National Laboratory

The HPCS View

- HPC is a special type of DC service
- Load (Job) characteristics
  - Block of allocation
  - Non-interruptible
  - Non-migratable
- Scheduler plays a central role
- Dynamic load
- Dynamic response time
  - Best: instantaneous
  - Average: property of system
  - Worst: system-defined

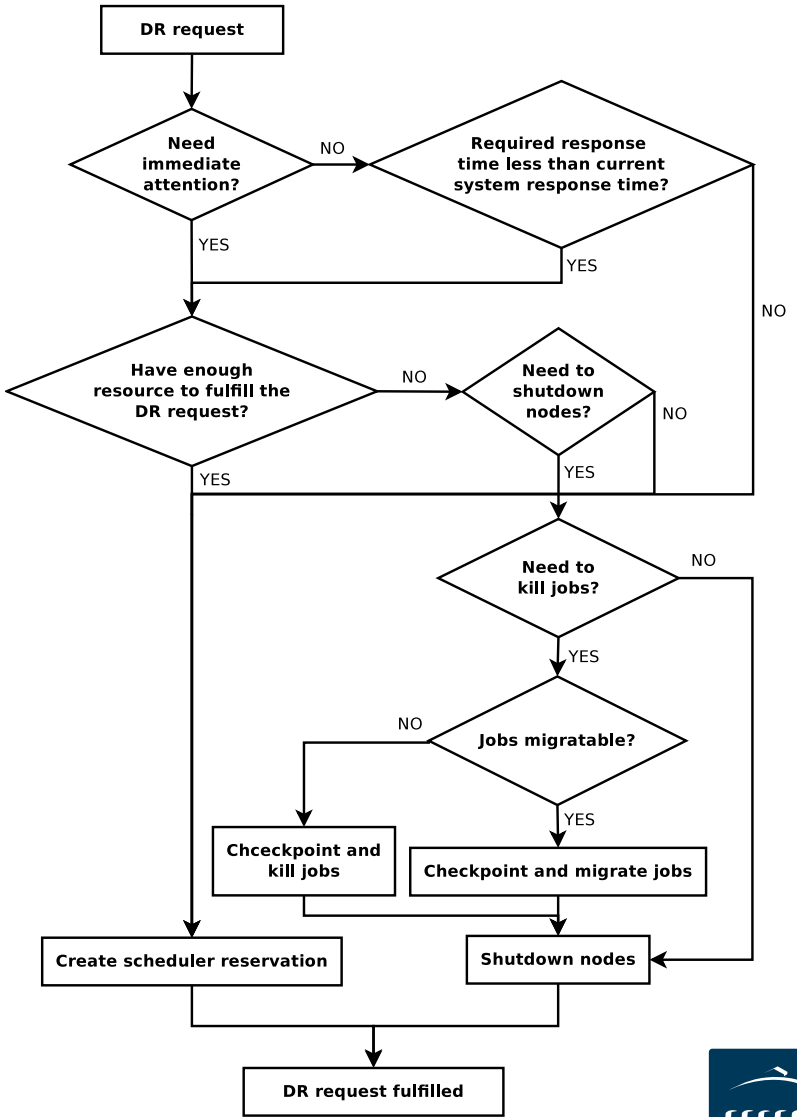

BERKELEY LAB
Lawrence Berkeley National Laboratory

Homogeneous Systems

- Mako: UCB
- Thresher: SDSC
- Load migration: Mako $\Leftrightarrow$ Thresher
- Two schedulers
- Average response time
  - Mako: 4 hours
  - Thresher: 6 hours
- Worst response time - 72 hours
- Test response time - 0 hours

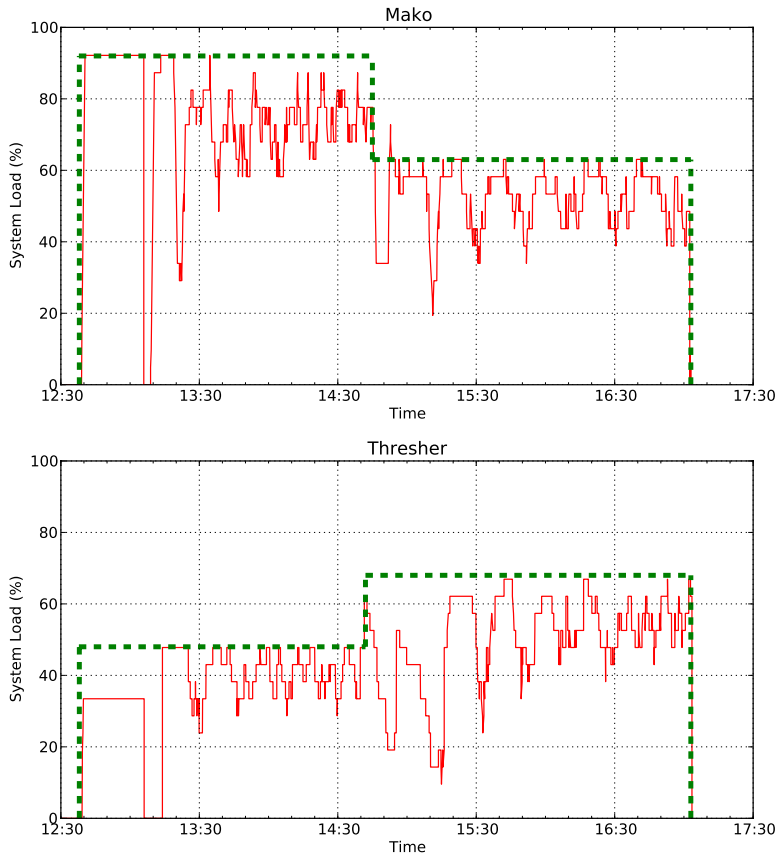

BERKELEY LAB
Lawrence Berkeley National Laboratory

Heterogeneous Systems

- LR1: SDSC
- LR2: LBNL
- Load migration: LR1 $\Rightarrow$ LR2
- One scheduler
- Average response time
  - LR1: 4 hours
  - LR2: 5 hours
- Worst response time - 72 hours
- Test response time - 4 hours
- Queue wait time on LR2 increased from $\sim$8 hours to $\sim$34 hours

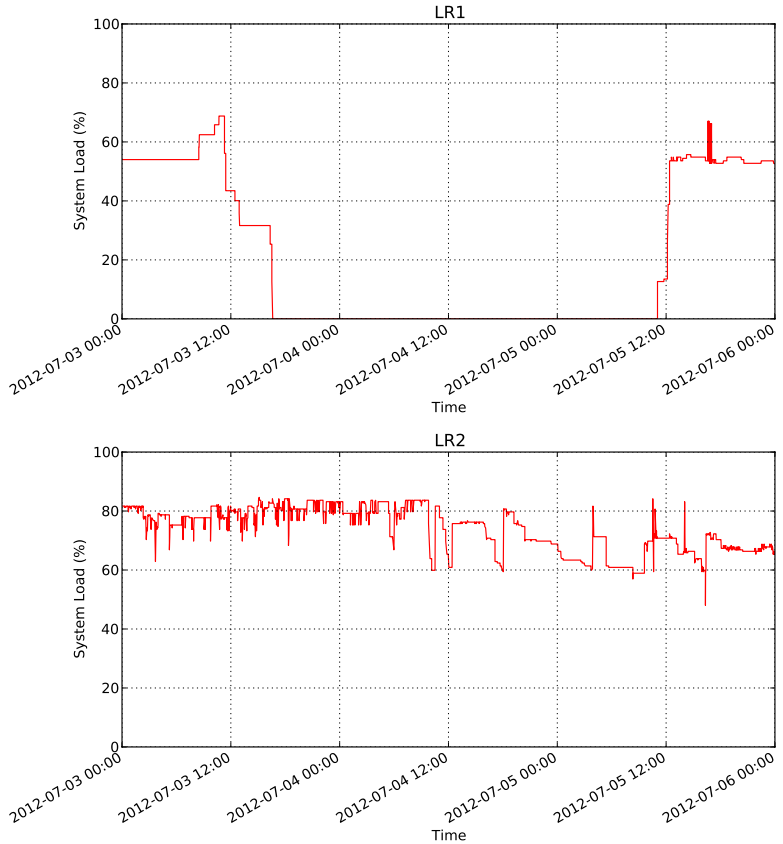

BERKELEY LAB
Lawrence Berkeley National Laboratory

Summary

- Enabling technologies
  - Scheduler with grid support - Moab
  - Provisioning system - Warewulf
  - Checkpoint/Restart
  - Power-aware scheduler
- Issues
  - Lack of system integration
  - Grid scheduler is still premature
  - Most jobs do not have checkpoint/restart feature built in
  - Automatic checkpoint/restart mechanism does not work well
  - Automation is difficult