




# Batch Scheduling on XT3

Chad Vizino <[vizino@psc.edu](mailto:vizino@psc.edu)>  
Pittsburgh Supercomputing Center  
CUG 2005

# Overview

- Simon Scheduler
  - Design
  - Features
- XT3 Scheduling at PSC
  - Past
  - Present
  - Future



Back...

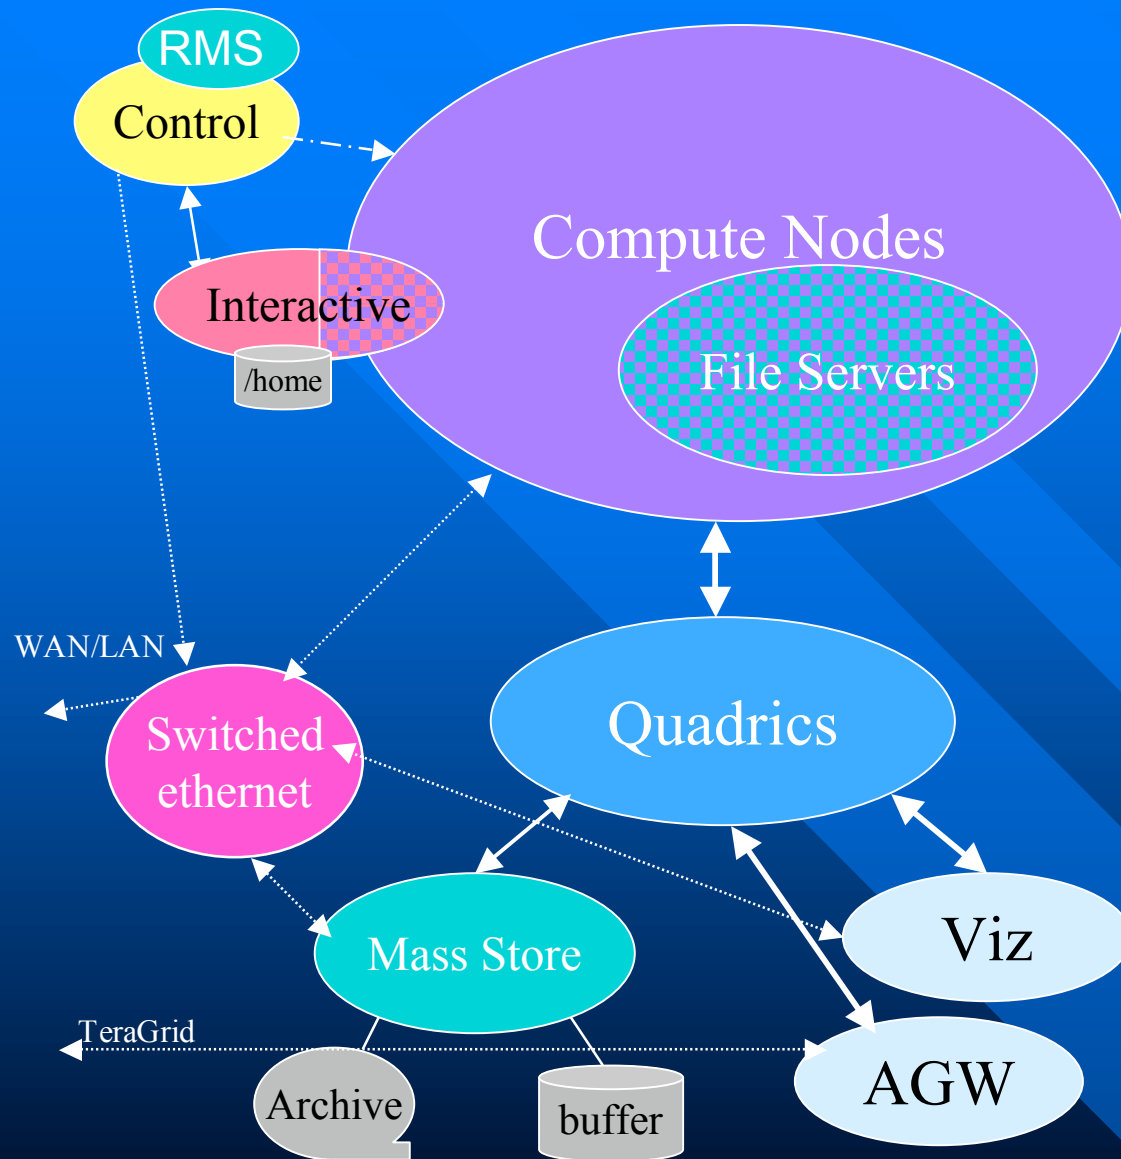
...to the Future!

# Scheduler Design Goals



- Support PSC Scheduling Goals
  - Encourage Large Jobs
  - Foster Parallel Development
  - Fair
- Maintain good utilization

# Terascale Computing System (LeMieux)

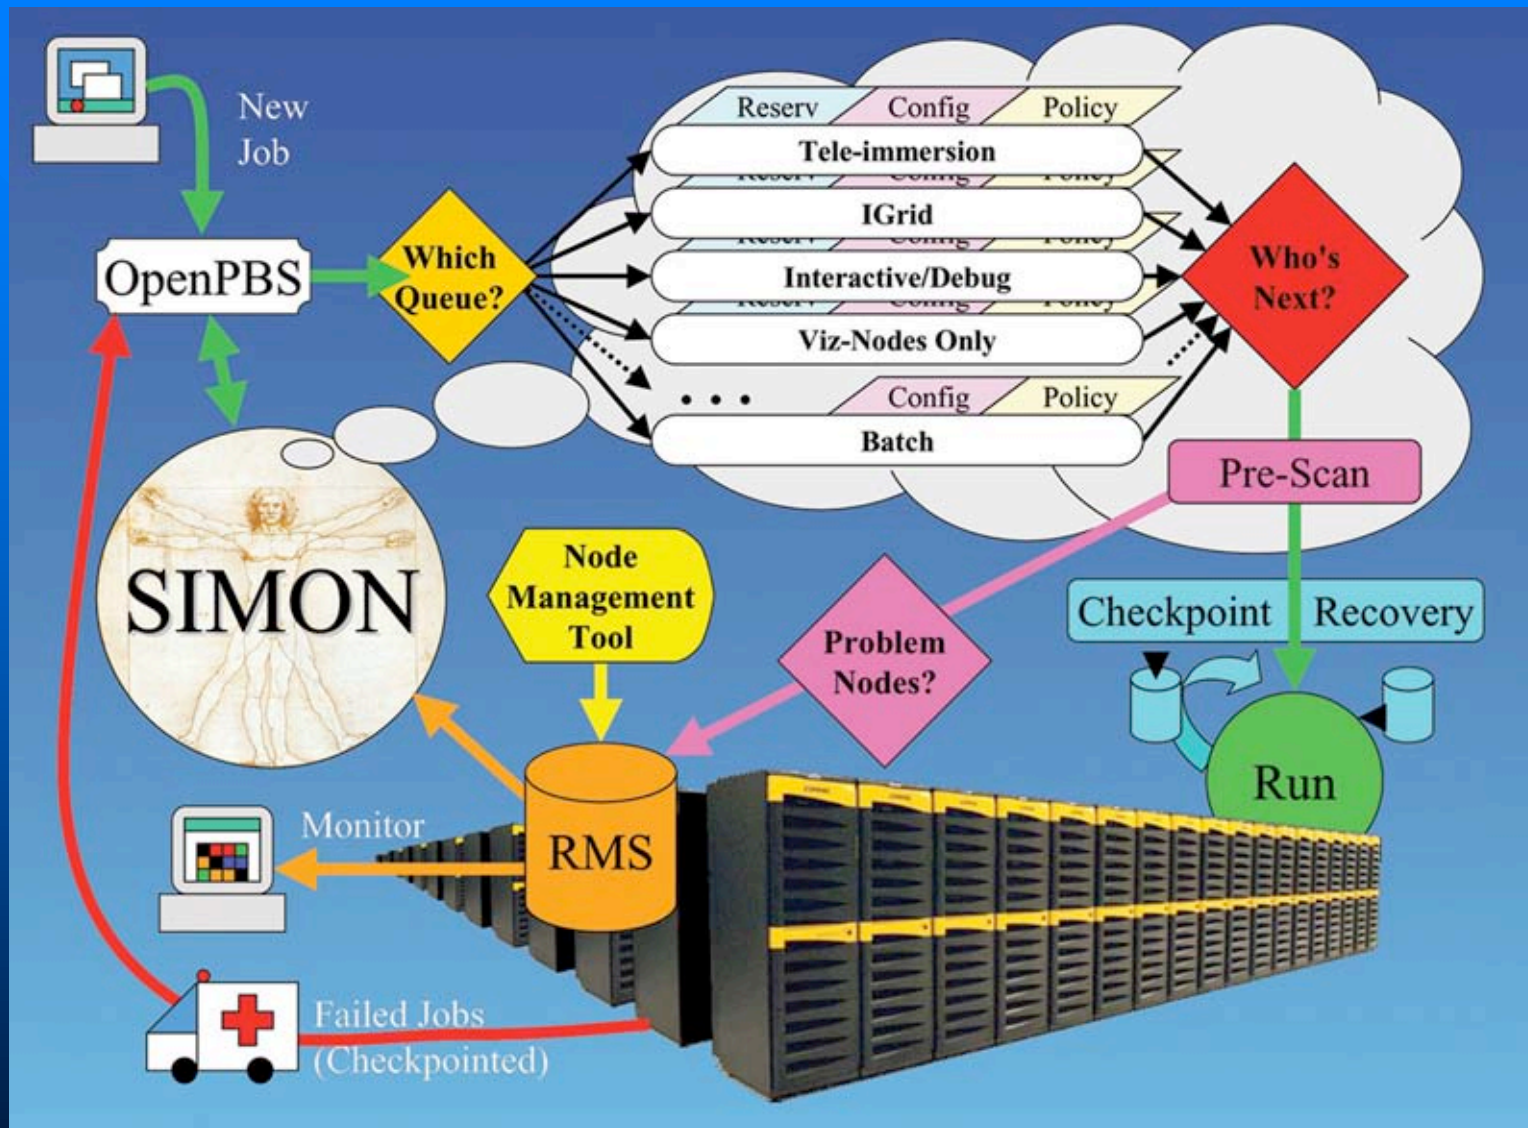


## Summary

- 750 Compute Nodes
- 3000 EV68 processors
- 6 Tf (peak, est >4Tf on LSMS)
- 3. TB memory
- 27 TB local disk
- Multi-rail fat-tree network
- Redundant monitor/ctrl
- WAN/LAN accessible
- Parallel visualization
- File servers: 30TB, ~32 GB/s
- Resource Management System (RMS)

PITTSBURGH SUPERCOMPUTING CENTER

CUG 2005



# *Simon Batch Scheduler*

- OpenPBS
- 6,500 lines of TCL
  - Rapid prototyping in TCL (`pbs_tclsh`)
  - Easy to add new modules
  - Only scripting language supported by OpenPBS
- Over four years in production
- Backfill
- Reservations
- Node Maintenance

# *Simon Batch Scheduler*

- Co-scheduling
  - Visualization cluster
  - Application GateWay nodes
- Pre-job scanning
- Ties to Usage Accounting
  - Queue statistics
  - Reservation charges





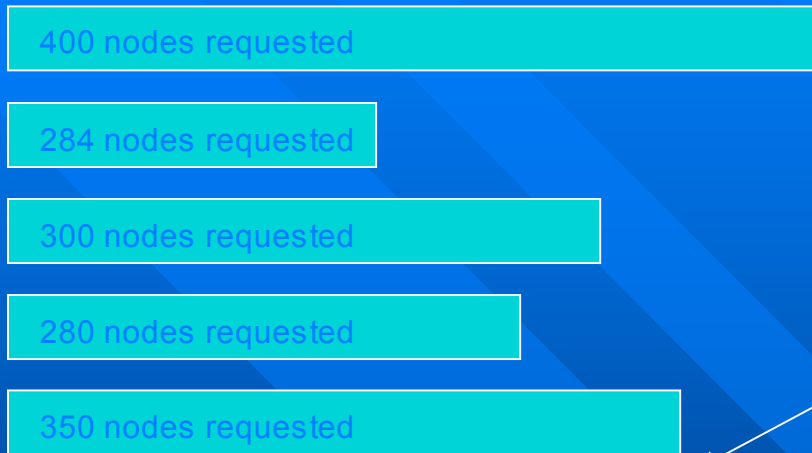
# Simon Namesake

- "Simon" is named for Dr. Herbert Simon (1916-2001), University Professor of Computer Science and Psychology at Carnegie Mellon University, and winner of the 1978 Nobel Prize in Economics.
- Argued that inevitable limits on knowledge and analytical ability force people to choose the first option that "satisfices" or is good enough for them.
- Scheduling a large computing system often requires making choices with limited knowledge.

# PSC Scheduling

## Batch Queue

upper  
bin



upper bin jobs are ordered  
in a “first in – first out” basis.

275 nodes requested

256 nodes

lower  
bin



lower bin jobs are ordered  
by largest number of nodes  
requested.

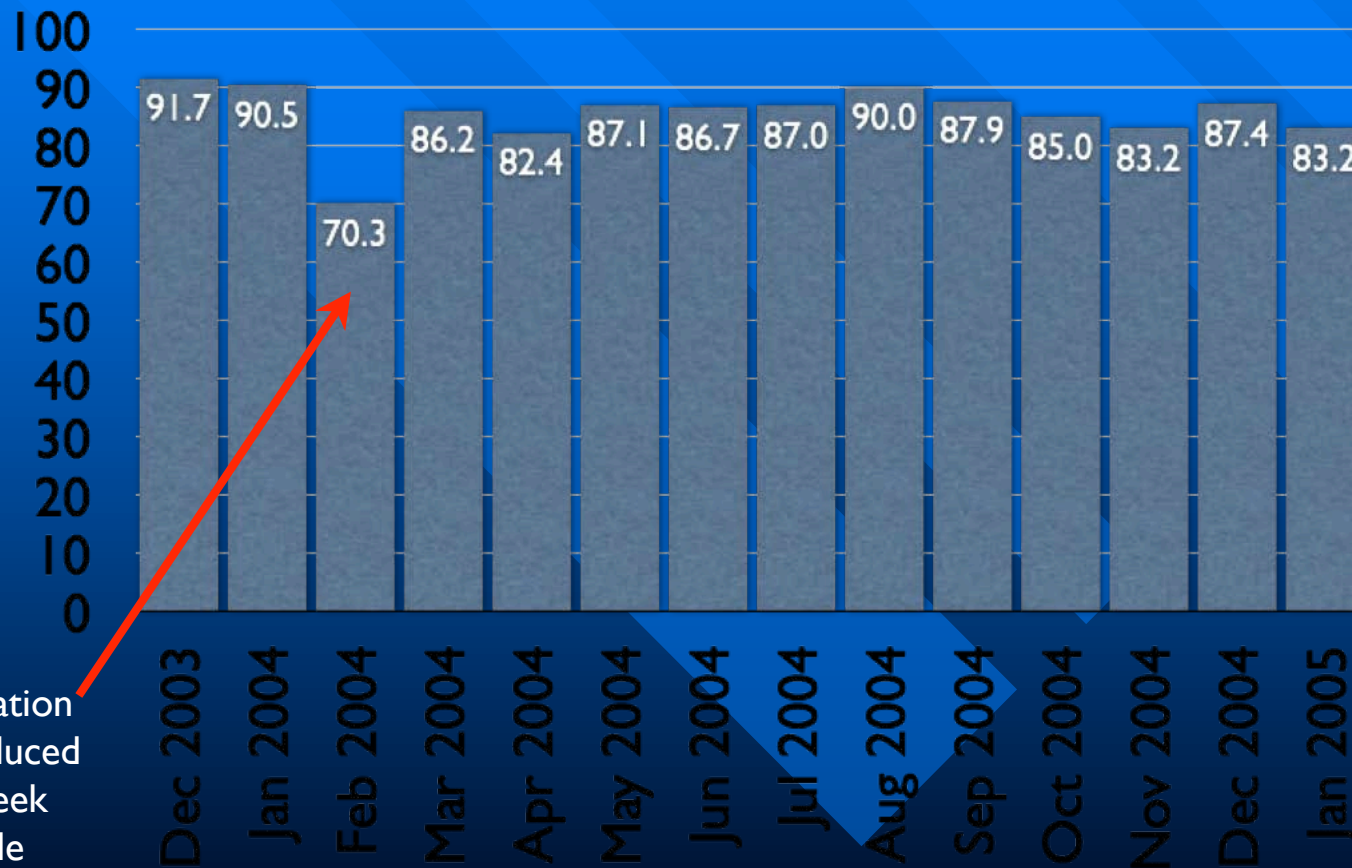
95 nodes requested

# LeMieux $\geq$ 1024 Processor Usage

Year	% Processor*Hours
2002	6.5
2003	28.6
2004	46.7

# There are no spare cycles...

Average Daily % Utilization of TCS (lemieux.psc.edu)  
750x4p Alpha EV68




Average utilization for month reduced due to one week system upgrade downtime

PITTSBURGH SUPERCOMPUTING CENTER


CUG 2005

# PBS-RMS Relationship



# Scheduling Support for Job Success

- Pre-Scan
  - File Systems
  - Orphan/busy processes
  - CPU availability/accuracy
  - Interconnect availability and performance
  - Failure removes the node from scheduling
- Job Profiling
  - Disk Performance
  - Live network connections
  - Application Fingerprints






Move to the recent past...

...the XT3

PITTSBURGH SUPERCOMPUTING CENTER

CUG 2005


# Single Cabinet System



Single cabinet  
October 2004  
Demonstrated Running  
Applications at SC04




# XT3 First Row Installed



Row 1


December 2004

# XT3 Second Row Installed



Row 2  
February 2005

# Original XT3 Design



PITTSBURGH SUPERCOMPUTING CENTER

CUG 2005

# Challenges

- System ran dev harness
  - No SDB
  - No CPA
  - No functional batch system
    - » Major setback for efficient applications work
  - Boot login nodes to clear bad nids
  - Yod –list 10..20,30,50,90..100 ....
- Lack of unique, writable file system on login nodes
- System partitioned into separate dev harness systems

# Solutions

- Use Torque (OpenPBS)
  - Builds on 64-bit platforms
  - Open source
  - Free
  - Have lots of experience with it
  - Can write custom scheduler
- Use RAM file system and load /var/spool/torque loaded from /usr/users (NFS mounted)
- Let flat files act as the SDB
- Let scheduler be the CPA
- Batch and interactive can be handled

# SDB/CPA Replacement

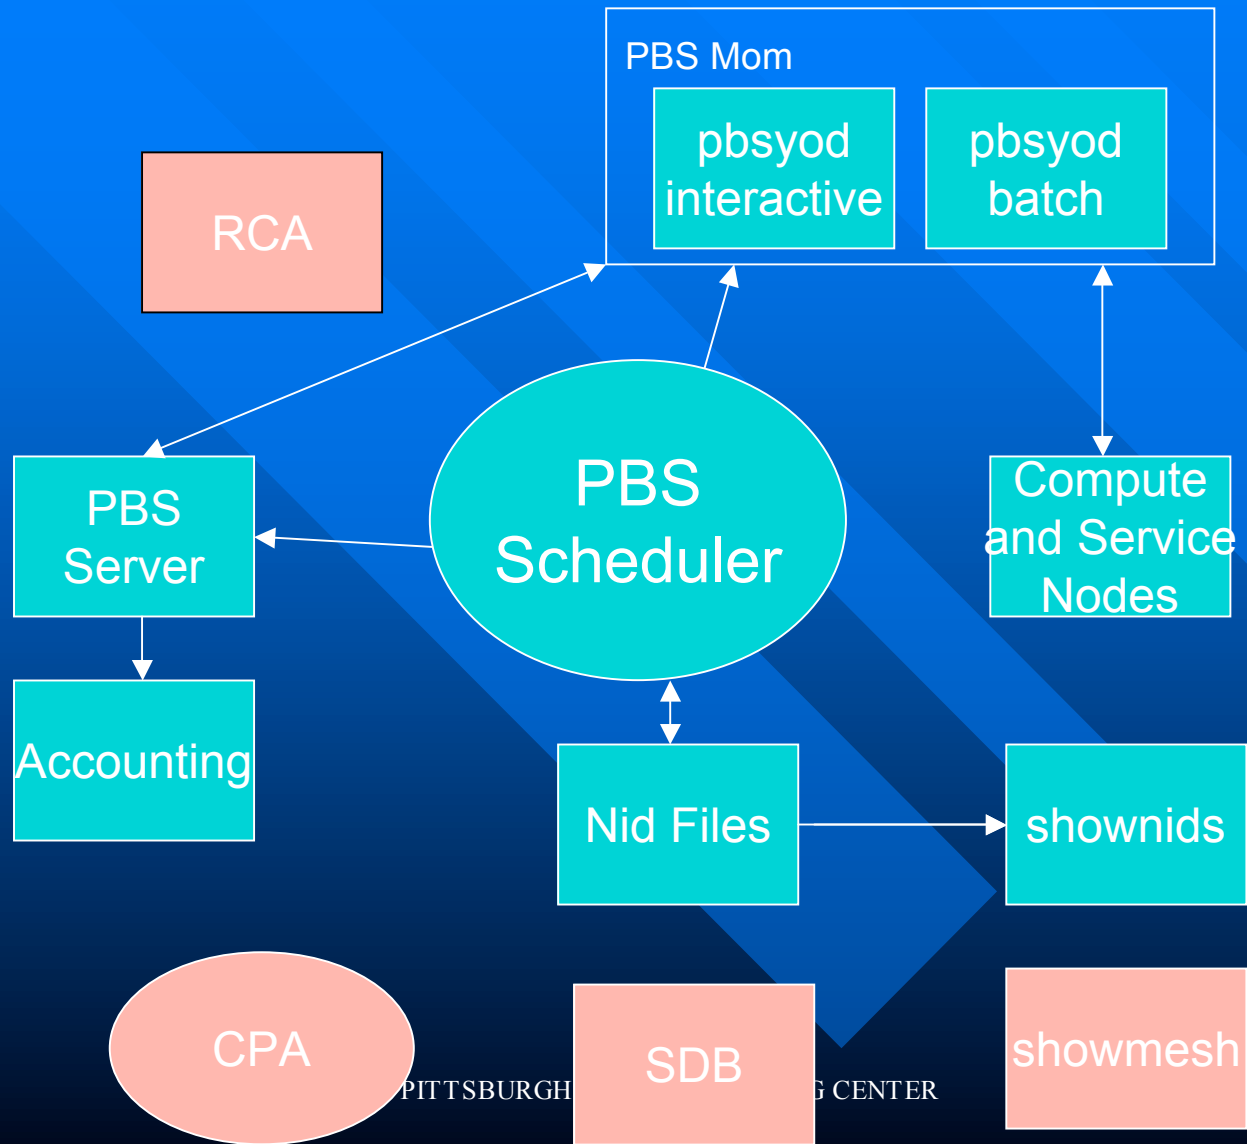
## ■ SDB processor table

- /opt/harness/default/ssconfig/sysN/node\_list
  - » Cray managed (HW list)
- /usr/users/torque/nids\_list\_loginN
  - » Scheduler managed
  - » Holds state of nids (enabled/disabled/allocated)
  - » Query via shownids

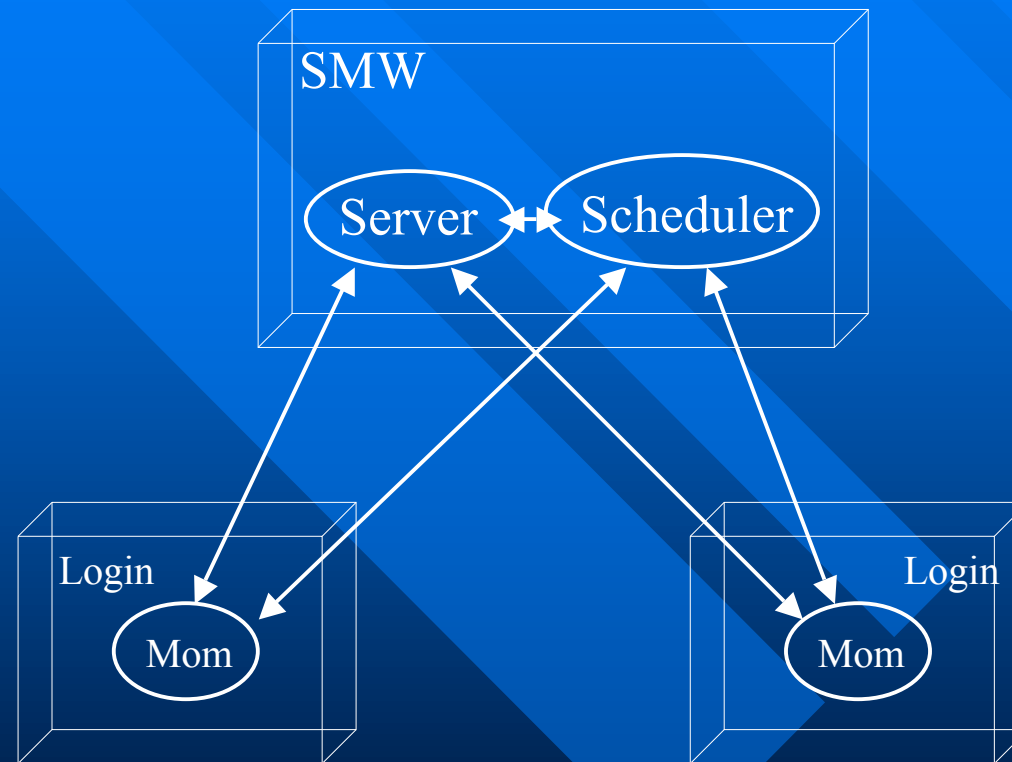
## ■ Yod

- pbsyod reads YOD\_NIDLIST from scheduler
- -size, -base options to stack in job

# Harness Layout



# Harness PBS Configuration





# Early Scheduler Features

- Adaptation to harness
  - Schedule to multi-cab arrangement
- Backfilling
  - System drains
  - Use aggressive backfilling (EASY)
    - » Switchable (can just use First Fit)
  - Top job (largest) gets reservation
    - » Use FIFO ordering
  - Other jobs can run as long as they don't delay the start of the top job (backfill)
- Pre-scan

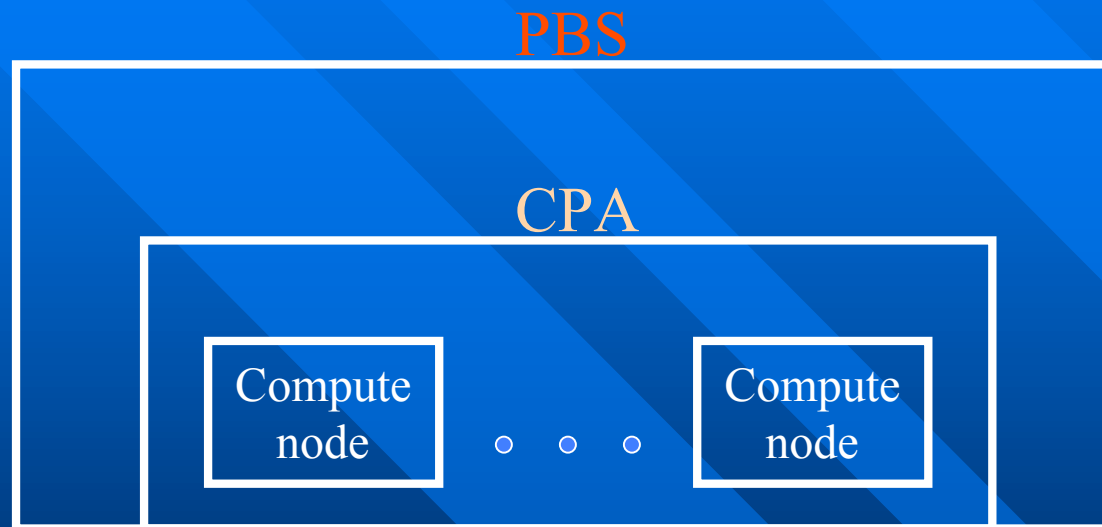
# Defensive Mechanisms

- Maintain lists of node categories
  - Checking
  - Bad
- Started calling ping\_node
  - After job finished
  - Slow
- Call ping\_list before job starts
  - Fail nodes
- Develop new, faster ping\_node
  - ping\_list -l 12..15,17..95,200..295
  - Returns good and bad lists
- Automated reboots
  - Let scheduler control

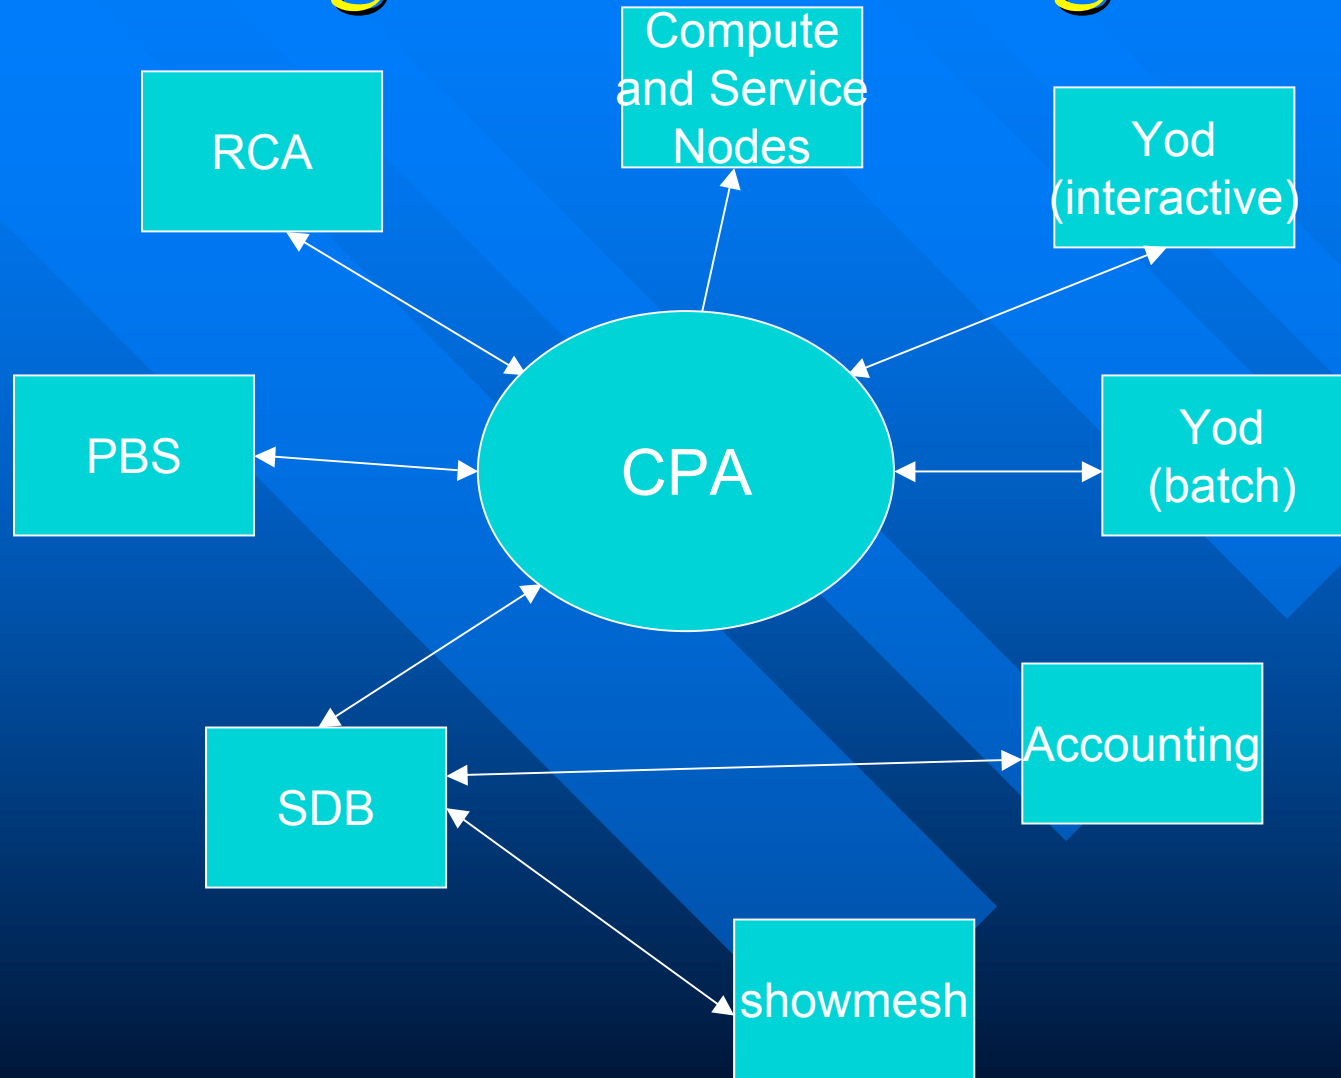
# On to CRMS...

- Installed late April
- Needed to integrate harness scheduling setup with CRMS
- CPA present
- SDB present
- Booting handled differently

# PBS-CPA Relationship



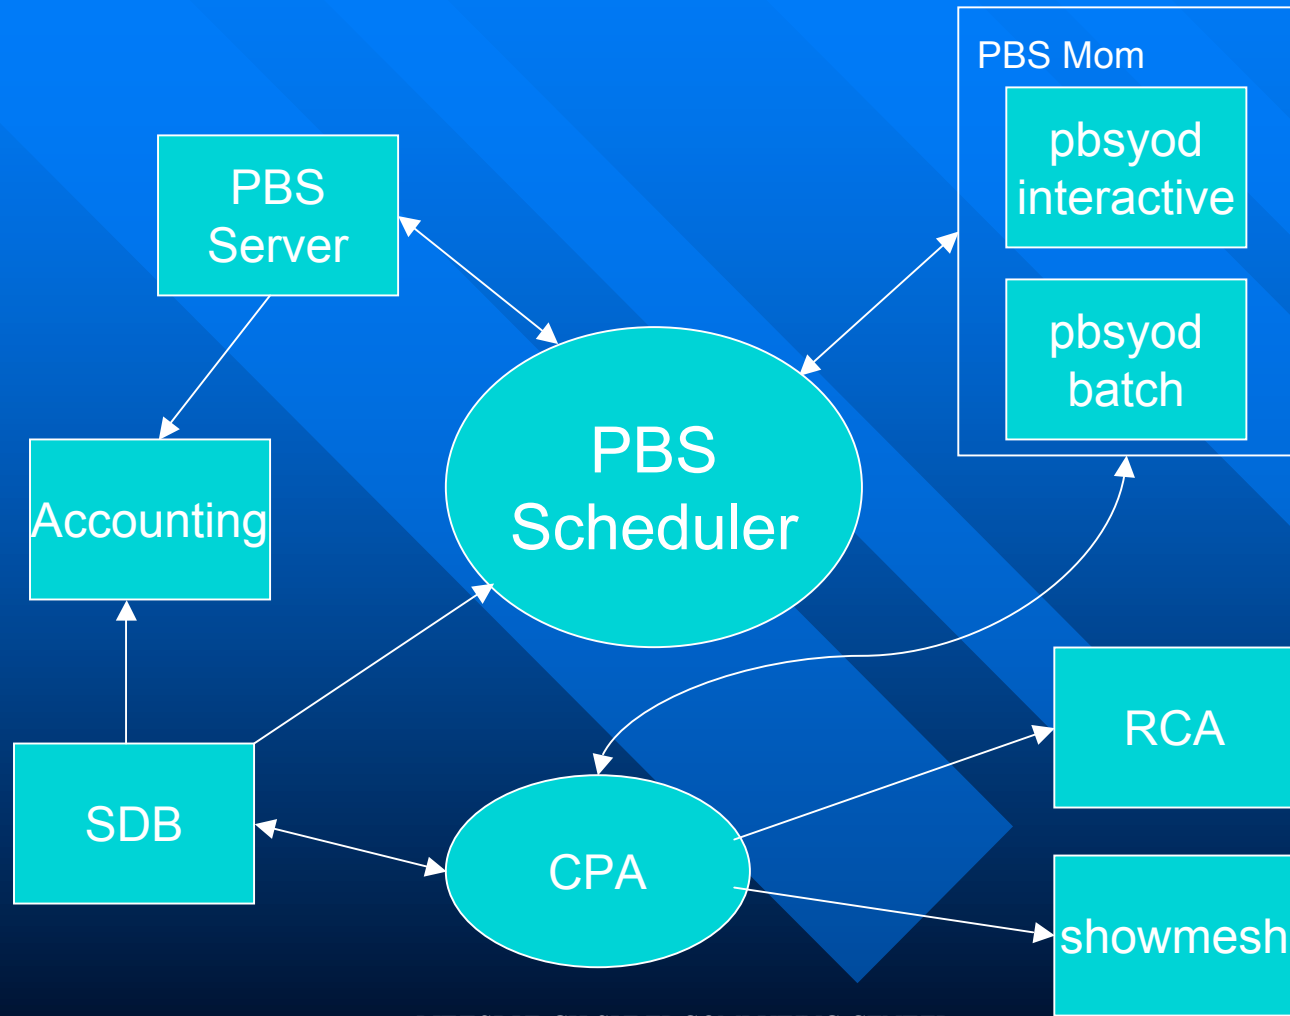
# Original XT3 Design



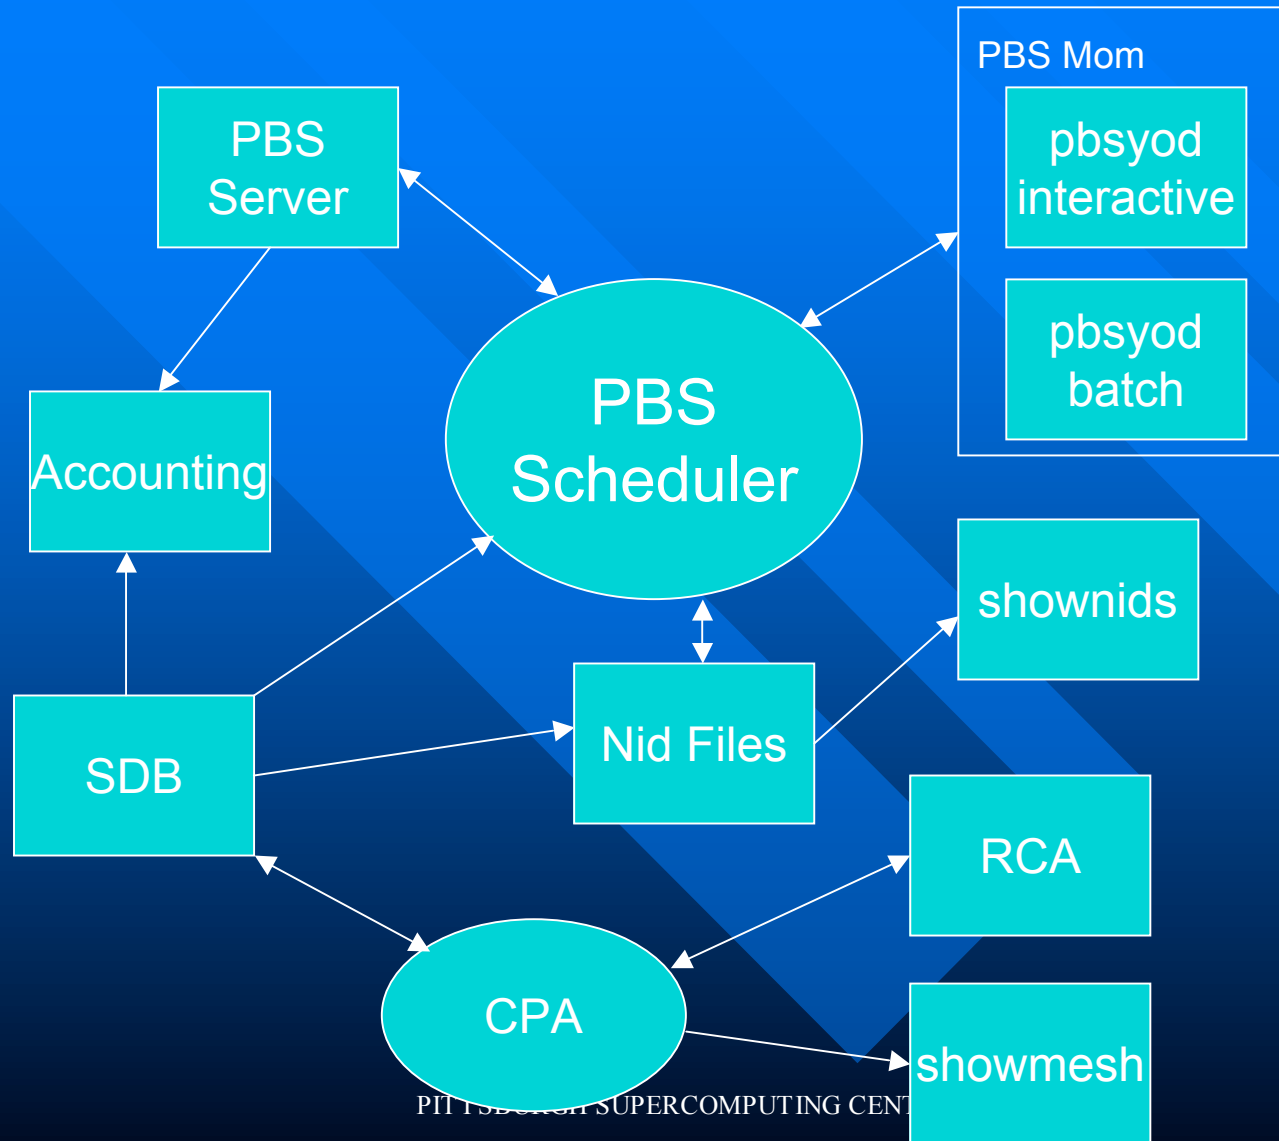
PITTSBURGH SUPERCOMPUTING CENTER

CUG 2005

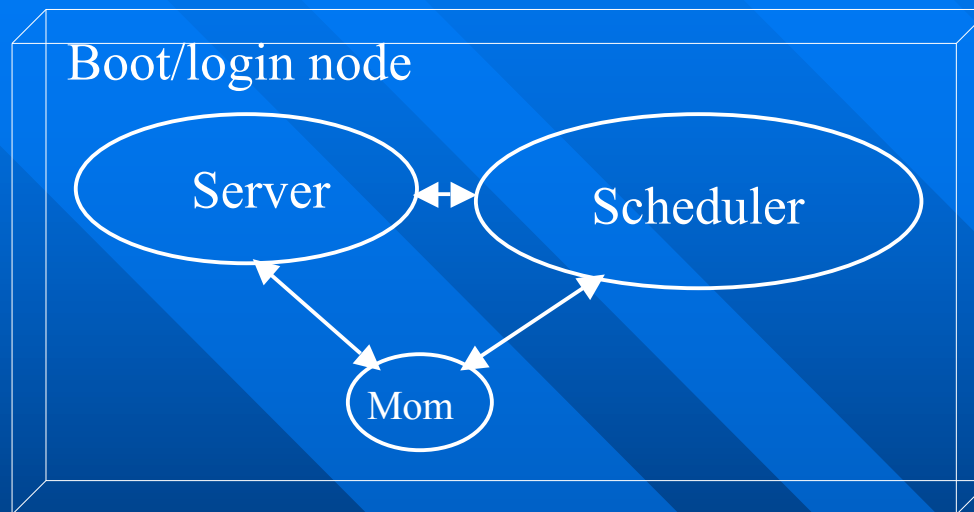
# Goal



# CRMS Phase 1



# Initial CRMS PBS Configuration

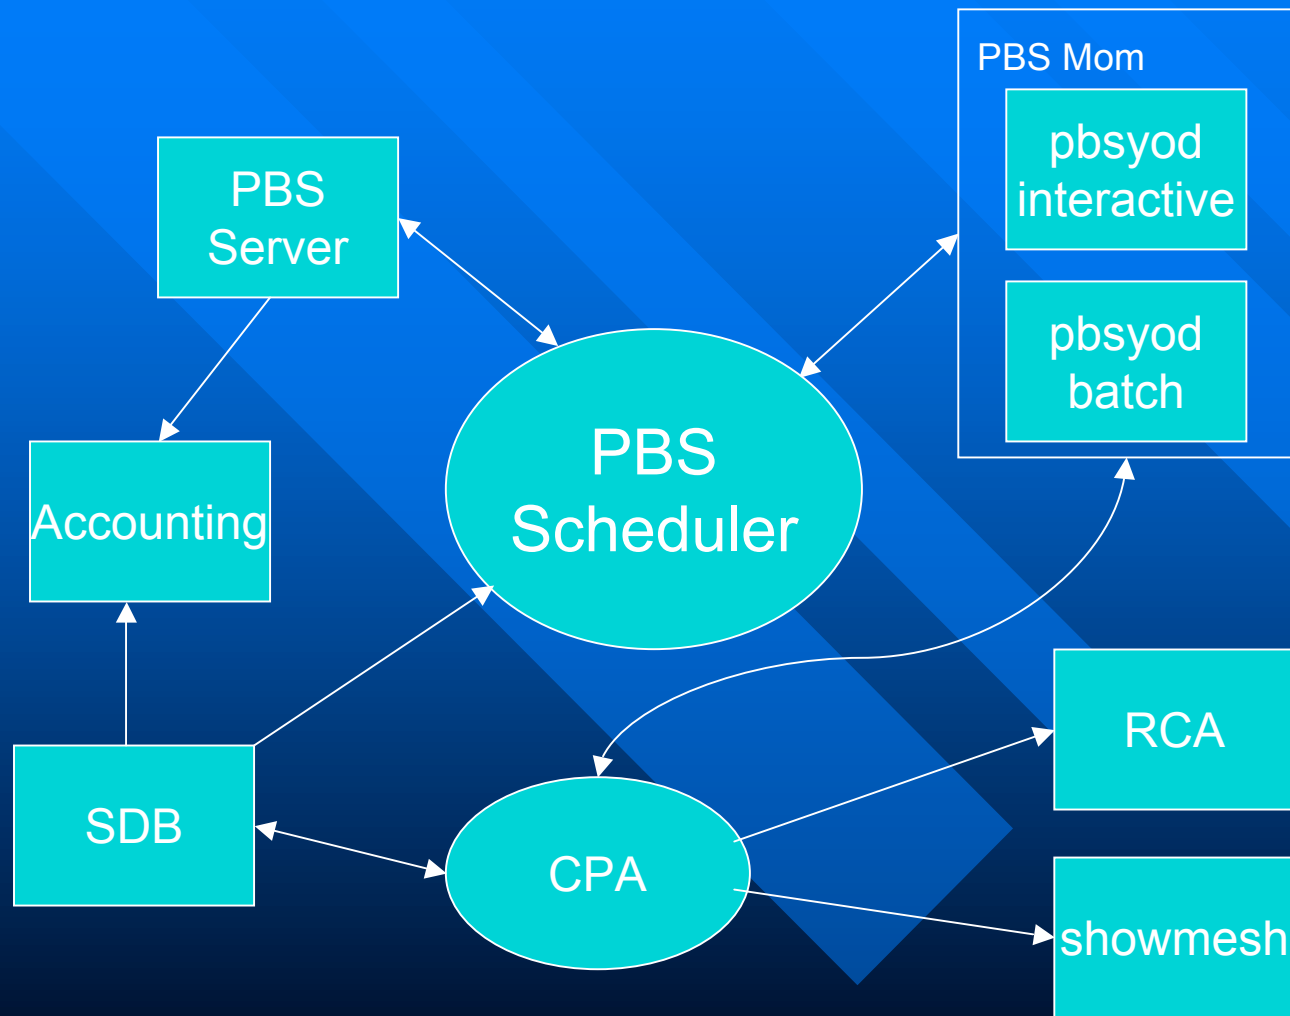




# Integration with CRMS

- At phase one with Torque
- Use “interactive” mode for processors
  - processor table in SDB
  - YOD\_STANDALONE to bypass CPA
    - » pbsyod
- Be gentle with SDB reads
  - Scheduler synchronizes itself with processor table every 5 minutes
- Logging tools

# CRMS Phase 2



PITTSBURGH SUPERCOMPUTING CENTER

CUG 2005

# Move to Phase 2 (in test)

- Use PBS Pro (instead of Torque)
  - Build pbs\_sched with TCL interpreter
  - New resources
  - Mom adaptation
    - » CPA calls
- Adapt scheduler to query SDB directly
  - Most of the other code will stay the same

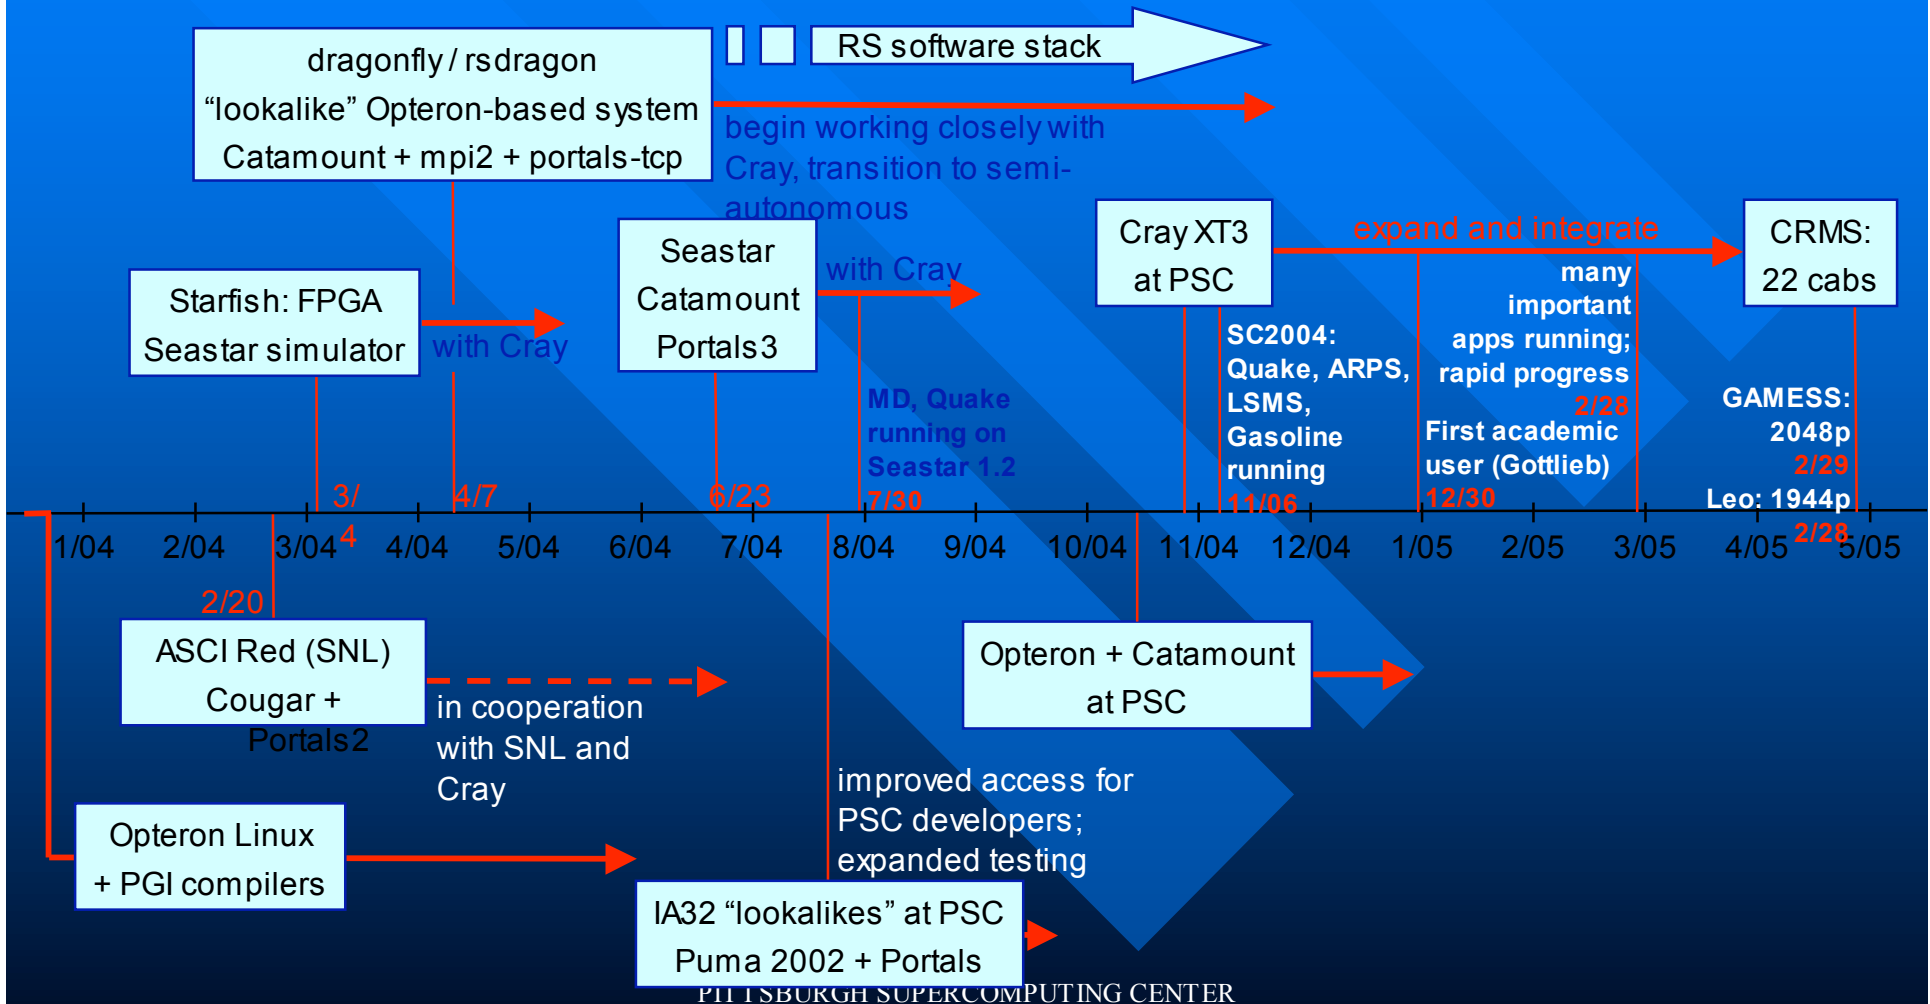
# Progress on Phase 2

- Build PRO pbs\_sched with TCL interpreter - Done
- New resources to PBS Pro - Done
  - nid\_list
  - linux\_nid\_list
  - Nidmask
- Changes to pbs\_mom - Done
  - Read nid\_list assigned by scheduler
  - Pass to CPA
- Changes to harness scheduler – In test

# Future

- More defensive node health checking
- Add more moms and schedule to these
- Integrate more features from Simon
- Investigate node allocation algorithms
  - No longer in CPA
  - Scheduler modules
  - Application specific
- Recode in Python
- Co-schedule service nodes
  - Viz
  - Data, etc.

# Successes facilitated by Batch System



PITTSBURGH SUPERCOMPUTING CENTER

# Other PSC XT3 Talks

- *Early Applications Experience on the Cray XT3*
  - Nick Nystrom, 2:30pm Today, Taos
- *Integrating External Storage Servers with the XT3*
  - Jason Sommerfield, 2pm Thursday, Taos