

BLUE WATERS

BREAKING THROUGH THE LIMITS

Experiences with Blue Waters

Michelle Butler

GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTING



IBM



What is Blue Waters?

- National Science Foundation Track 1 system
 - 208M sustained PF machine in production 6/2011

System Attribute	Abe	Blue Waters
Vendor	Dell	IBM
Processor	Intel Xeon 5300	Power 7
Peak Performance	90TF	-----
Sustained Performance	5TF	>= 1PF
# of cores/chip	4	-----
# of processor cores	9600	>200,000
Amount of memory	14.4TB	>800TB
Amount of Disk	½ PB	> 10PB
Amount of Archival Storage	5PB	500PB
External Bandwidth	40Gbps	100-400Gbps

Current issues?

- **GPFS and GHI tightly coupled with HPSS as archive server**
 - Users access GPFS only, HPSS is in the background as tape only
- **File systems sizing**
 - Most performance while ease of use, and data flow to and from archive is achieved with performance considerations
- **What file systems are managed?**
 - Scratch? Home? Others? Blocking size?
- **Do we manage every file?**
 - Or just files that have been there for 1 week or more?
 - Dials that we can change to alter the course of the data ...

Issues continued...

- **Backups?**
 - How do we backup these large file systems?
 - NCSA is writing RAIT for HPSS
- **Import/Export in and out of machine?**
 - NCSA is writing a facility to not only move data in and out of the machine (internal and external), but to move data within the machine for jobs (to and from HPSS)
- **Schedule jobs based on data as a resource to be managed too like a memory resource**

BLUE WATERS

BREAKING THROUGH THE LIMITS



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

Questions?

-