

BLUE WATERS

SUSTAINED PETASCALE COMPUTING

Resiliency

National Center for Supercomputing Applications, Department of Chemistry, Department of Computer Science, and Department of Electrical & Computer Engineering



GREAT LAKES CONSORTIUM
FOR PETASCALE COMPUTATION

Collaboration Ideas

- Analysis of “Fail in Place” Plans and Results
 - Modeling performance degradation
 - Modeling Failures over time - what happens when you use up your spares?
 - Actions - see below detailed actions
 - Rare Event Analysis
 - Possible collaborator - Pierre Del Moral
- Result Certification
 - INRIA Resources - COQ, Saclay, Rennes, Bordeaux
 - Action – Email exchange to define what might be done – Claude, Bill,

Collaboration Ideas

- Failure Modeling and Analysis
 - HELO - Ana
 - Tsubame Logs – Leonardo
 - ISC – Joshi
 - Investigate use of common architecture
 - Actions
 - Visit UIUC in April – Leonardo
- New Fault Tolerant Protocols
 - Clustering of events - Thomas, Estebane, Sebastien
 - Action – Submit a paper (ICS – Jan 2011, SC11 – Apr 2011)

Collaboration Ideas

- Job Failure & Correct Error Code Propagation to Applications
 - Catalogue errors that should be presented
 - Test Error code propagation in existing systems
 - Test Error code propagation on BW when HW arrives
 - Actions
 - Analyze existing job failure logs and “refund logs” – Galen
 - Analyze accuracy of error messages on existing systems – BP and BD
- Error rates and resource uses compared to energy consumption
 - Thomas
- Analysis of Tape Subsystem
 - Actions
 - Identify ‘projected’ device error rates – Mathias
 - Provide ‘projected’ device error rates – Michelle Butler/Jason
 - Record and analyze real data – 2H2011-2H2012
 - Logs from Tsubame - Leonardo

Collaboration Ideas

- Benchmark tests for Scotch on different systems
- C/R in memory
 - Action
 - Is there dev/HSM in Blue Waters? If so does that work? – Jason B
 - Presentation to PRAC meeting in March about this option – Leonardo
- How to reschedule work once a job failure occurs?
 - How to coordinate failure management and batch scheduler.
 - Scheduler has to be able to add resources, adjustment time limits, etc. if there is a failure
 - Actions -
 - Analyze LL functionality and provide recommendations for LL design – Olivier
 - Pointers to documentation for LL – Bill K
 - Access to run on BD and BP for INRIA – Bill K, Olivier
 - Proposal for an INRIA student to work on it

Collaboration Ideas

- C/R interval -
 - What is the best combination of C/R times compared to performance
 - Characterize error logs to see if it is a weibull distribution
 - Actions
 - Share pointer for error logs- Ana
 - Share expected numbers for BW – Time to C/R, memory rates, etc. – Bill K
 - Create and share model – Fredric V
- Concurrently optimized storage – how does it apply to C/R
 - Postdoc at UIUC – Jan 2011 – Jan 2012 –
 - Action – Study IBM's implementation of C/R I/O - Bogden & Bill