



Cluster2005

Daniel S. Katz

JPL

LSU (from mid-May)

d.katz@ieee.org



Cluster Questions

- How is a cluster distinguished from other computers?
 - Hardware?
 - Software/Operating System?
 - Network?

- Is a cluster a grid?
 - No, they are completely different - cluster can be element of grid...
 - Maybe, it depends on the middleware and how it's used



A Sample Application

- Montage: an astronomical image mosaic service for the National Virtual Observatory (<http://montage.ipac.caltech.edu/>)
- Core team at JPL and Caltech, grid architecture developed in collaboration with ISI

- Until the final SOS10 program came out, I was planning to present: D. S. Katz, G. B. Berriman, E. Deelman, J. Good, J. C. Jacob, C. Kesselman, A. C. Laity, T. A. Prince, G. Singh, and M.-H. Su, "A Comparison of Two Methods for Building Astronomical Image Mosaics on a Grid," *Proceedings of the 7th Workshop on High Performance Scientific and Engineering Computing (HPSEC-05)*, 2005.



Montage: Cluster vs. Grid Conclusions

- Montage is a custom astronomical image mosaicking service that emphasizes astrometric and photometric accuracy
- Final public release, Montage version 3betaX, available for download at the Montage website (<http://montage.ipac.caltech.edu/>) allows serial/MPI/grid processing
- A prototype Montage service has been deployed on the TeraGrid
 - It ties together distributed services at JPL, Caltech IPAC, and ISI
- MPI version of Montage:
 - Best performance
 - Requires a set of processors with a shared file system
- Pegasus (<http://pegasus.isi.edu/>) & Condor DAGman version of Montage:
 - Similar performance for large problems
 - Built-in fault tolerance
 - Can use multiple sets of processors



Cluster2005 and Grid2005

Daniel S. Katz

JPL
& LSU (from mid-May)

d.katz@ieee.org



Logistics

- **ClusterXY**

- 2005 meeting held in Boston for 4 days last September w/ ~200 attendees, see <http://cluster2005.org/>
- 2006 meeting to be held in Barcelona for 4 days this September, the week after EuroPVM/MPI, see <http://cluster2006.org/>

- **GridXY**

- 2005 meeting for 1.5 days at SC|05 was the 6th Grid workshop (5 of which were part of SC'XY) w/ ~450 attendees, see: <http://pat.jpl.nasa.gov/public/grid2005/>
- 2006 meeting will be 2-day conference co-located with Cluster2006 in Barcelona, see: <http://grid2006.org/>
- Last day of Cluster (workshops) will start after Grid keynote, Cluster-only attendees invited to Grid keynote

- **Think about attending both, to see how clusters and grids are related**



Cluster2005: Technical

- Program Chair, Thomas Stricker
- Acceptance Rate, 37%
- 1 Keynote & 1 Endnote:
 - On the Evolution of Communication Hardware and Parallel Programming Models, Mark Snir
 - Improving the Usability of Clusters, Bill Gropp
- 3 Plenary Papers (Best Papers):
 - Experiment Management and Analysis with perfbase, Joachim Worringer
 - Service Augmentation for High End Interactive Data Services, Matthew Wolf, Hasan Abbasi, Benjamin Collins, David Spain, Karsten Schwan
 - Distributed Out-of-Core Preprocessing of Very Large Microscopy Images for Efficient Querying, Benjamin Rutt, Vijay S. Kumar, Tony Pan, Tahsin Kurc, Umit Catalyurek, Yujun Wang, Joel Saltz



Cluster2005: Technical (2)

- 42 papers in two-track format, in sessions:
 - Optimizing with MPI
 - Scheduling
 - Parallel I/O and Efficient Communications
 - Fault Tolerance
 - Collective Communication Optimization
 - Networking
 - Security
 - Reliability
 - Systems Analysis and Performance
 - Distributed App.'s and OS's
 - Visualization and Simulation
 - Messaging



Cluster2005: Technical (3)

- 26 posters (really less, but I can't remember how many didn't show up)
- 3 tutorials
 - Parallel I/O in Practice
 - State of Infiniband in Designing Next Gen. ...
 - HA-OSCAR
- 2 workshops
 - RAIT (Remote Direct Memory Access (RDMA): Applications, Implementations, and Technologies)
 - HeteroPar
- Papers should be on web site (Thomas?)



Grid2005: Technical

- Program Chair, Daniel S. Katz (with help from 5 great vice-program chairs)
- Acceptance Rate, 18%
- Keynote:
 - Research Challenges for the Next Generation of Grid Systems, Dennis Gannon
- 6 Plenary Papers (Best Papers):
 - Wide Area Data Replication for Scientific Collaborations
 - Protecting Grid Data Transfer Services with Active Network Interfaces
 - Authorization and Account Management in the Open Science Grid
 - On the Creation & Discovery of Topics in Distributed Publish/Subscribe Systems
 - Grid-Enabling a Vibroacoustic Analysis Application
 - Collective Operations for Wide-Area Message Passing Systems Using Adaptive Spanning Trees



2006 Co-located Conferences

- Cluster2006 - <http://cluster2006.org/>
 - September 25-28, Barcelona
- Grid2006 - <http://grid2006.org/>
 - September 28-29, Barcelona
- The prior week:
 - EuroPVM/MPI (Bonn)
- Huge street festival “La Mercè”
 - Starts middle of the prior week, and builds to Sept. 24th
 - Plenty of street activities
 - Great fun



“There is no way I’ll miss this festival the weekend before the conference” - Mark Baker
 “The best night life ever! My students partied till the first metro opened in the morning” - Thomas Stricker
 “My wife said: If my boss does not allow me to go to Barcelona, I’ll quit my job!” - Bill Gropp

- (2007 - Austin)

