

Minutes of the HiPCAT Meeting 24 Aug 2005

HiPCAT attendees:

- James Abbott (TTU)
- Kim Andrews (Rice)
- Jay Boisseau (UT)
- David Chaffin (TTU)
- Francis Dang (TAMU)
- Borries Demmeler (UTHSCSA)
- Nick Grishin (UTSW)
- Olin Johnson (UH)
- Andre Kerstens (UTEP)
- Josten Ma (UH)
- Rosalinda Minendez (UH)
- Jerry Perez (TTU)
- Paul Roberts (UH)
- Tommy Minyard (UT)
- Alan Sill (TTU)
- Phil Smith (TTU)
- Warren Smith (UT)
- David Steffen (BCM)
- Pat Teller (UTEP)

(Andre Kerstens and Pat Teller from UTEP joined late because of a time zone confusion)

- TIGRE update by Phil Smith:
 - Alan Sill hired by TTU.
 - Rice hired Mark Mazine to work on the TIGRE project.
 - TACC made a job offer, but the candidate declined.
 - Warren Smith is working on TIGRE.
 - TIGRE activity on HiPCAT website. Some data was initially lost, but the system admin at UH was able to recover the data. (Kudos!).
- THEGrid update by Alan Sill:
 - Visit at TTU by Texas A&M physics faculty Profs. Teruki Kaman and David Toback, and postdoc Sungwon Lee.
 - Discussion of technologies is ongoing (e.g. online electronic logbooks).
 - Most active in THEGrid :UT Arlington, UH, Rice , A&M, TTU.
 - Plan mid-fall workshop (date tba).

Institutional updates

- UT
 - UT received Teragrid Award from NSF starting Aug. 1 for 5 years
 - UT is looking for a candidate to chair the Distributed and Grid Computing group. Other positions already hired: a materials science researcher from Minnesota and a researcher from CMU. Still looking for computational biology person.

- TACC is expanding their machine room with cooling and power to grow the Rangler cluster from 256 processors to 656 processors, possibly to a 1000 processors.
- Received IBM second year funding for UTGrid: 2 grid technology people from IBM are working onsite at TACC and will be working on TIGRE, Teragrid and campus cyber-infrastructure.
- Teaching 'Introduction to Parallel Computing for Science and Engineering' starting Sept. 1. Making push to teach 'Visualization and Data Analysis for Science and Engineering' and 'Intro to Scientific Computing' in Spring. Technologies related to the intro course are: databases, python, xml and grid computing.
- TACC wants to become member in OpenScience Grid. Jay thinks this will be worked out soon.
- There will be a ceremony for a new building on Sept 29. HiPCAT members will be invited. The new building will house about a 100 staff and students. The machine room will double in size compared to current one, totaling approx. 10'000 sq. feet
- UTSW
 - An in-person meeting will be held at UTSW on Thursday Sept. 22. Agenda will be sent out within 2 weeks.
 - UTSW has bought new computer equipment that is housed in and maintained by TACC as part of a collaboration with TACC.
 - They are expanding their existing cluster with 34 dual-Opteron nodes. The cluster is housed in their campus machine room.
 - Stuart Johnson from TACC is helping UTSW setup their computational facilities (he's based at UTSW).
- Rice
 - Have purchased a 336 processor dual-core Opteron cluster which will be operational by mid-October.
 - They have an open postdoc position for the TIGRE project.
- BCM
 - The new computer room is up and running and currently holds a couple of hundred nodes (dual Opterons, 4 GB RAM) Question for David: what disk I/O is he using? He will inform on this later.
- UH
 - Combine TLC² under research division with IT department to support HPC and advanced networking services.
 - New 4000 sq. feet machine room. All clusters combined (approx. 3 TFlops) in one place.
 - They are actively recruiting HPC, network, and storage professionals.
 - Building the Reno network which is managed by the HPC group; Baylor and UT will be brought in soon.
 - UH is organizing the 2nd cyber-infrastructure workshop (December 6-8).
 - UH is organizing the 'Innovation through Collaboration' workshop sponsored by Johnson Space Center Houston.
- UTHSCSA

- Thanks from Borries to the group for inviting them to HiPCAT
- Installed a microway cluster: 20 dual Opteron 250 nodes with 4 GB RAM. A quad headnode with Opteron 850 processor, 8 GB RAM. Linked up by dual-channel bonded gigabit cards. Bringing an 8 TB storage device online. The cluster is mainly running bioinformatics apps like blast, sam, mass-spec, crystallography and an app to model biological micro-molecules.
- Working on porting biophysics apps to TIGRE grid (a meeting with TACC took place to look at this).
- UTHSCSA is planning to connect to the LEARN network later this year.
- They have a preliminary agreement with TACC that Jay will come down to UTHSCSA to help identify how to integrate better in HiPCAT and how to make best use of local resources and share these with other institutions.
- Question to Jay: Is steering committee looking at certain application areas. Phil says to postpone this to later.
- TAMU
 - Interviews are ongoing for a senior position at TAMU.
- UTEP
 - Delivery of a 24-processor p590 and 4-processor p550 to the academic server room of UTEP. An event is organized on Sept. 23 to celebrate this.
 - Luis Hernandez appointed HiPCAT representative from the IT department.
 - IT has funded a new position to take support and maintain the equipment in the academic server room.
 - Andre Kerstens was appointed HiPCAT secretary; The College of Engineering of UTEP is financially supporting his attendance to in-person meetings.
 - A gigabit Ethernet switch and a UPS were installed in the academic server room; environmental monitoring software (Netbotz) is under review.
 - 3 HPC-related papers have been accepted at HPC conferences.
 - Richard Zamudio, Michela Taufers student, is back from a 3 month visit to SDSC in San Diego. He has been working on the Topaz project. This has intensified the collaboration between UTEP and SDSC.
 - Andre Kerstens, Michela Taufer and Pat Teller are attending SC2005 and are currently working on materials for the booth (posters, trifolds, gadgets).
- TTU
 - Power has finally been installed for the new 256-processor Xeon cluster which has Infiniband connectivity. Dell will be coming on-site for 2 weeks to install this system.
 - Is HiPCAT executive committee restructuring needed because of the current growth of the organization?
 - Jerry Perez has just come back from Las Vegas where he presented a talk on TTU grid initiatives. Sponsored by Sybase. He had good and productive conversations with Sybase concerning academic collaborations.
 - Alan mentions a report on deploying grids in the commercial sector. Conclusion is that storage is the primary obstacle to deploying large-scale grids in companies.
 - TTU will teach an advanced MPI course this Fall. Received input from Jay Boisseau.

- HiPCAT website update by Josten Ma
 - TIGRE is using the website as it was intended to be. Josten is gratified to see this, and offers continued training and support.
 - The site recently migrated to a new version of Zope/Plone that has much more functionality than the previous version.
 - Webserving is done by a primary and secondary server which are synchronized every day during low traffic hours.
 - Josten is implementing critical file and directory backup of the servers.
 - He had to get website data back from a week old backup because of a server crash.
 - New job postings page has been added to the HiPCAT website.
 - The HiPCAT members-only area is only accessible for HiPCAT people. Question from Kim on page privileges.
 - Mailing list info added to the website. Contact Josten or Phil if there are problems with the lists (e.g. bounces)
 - Josten requests all members to update the HPC resources page and keep it up to date from now on.
 - According to Rosalinda, Josten is available to all members for website training and questions.

- Update from Jay on EPIC – Engaging People In Cyber-infrastructure (this is a funded proposal)
 - Approx. 20 people around the country are involved in this initiative.
 - This has much to do with training, methodologies, education outreach and materials to get people interested in advanced computing.
 - UT is currently preparing 4 classes in advanced computing and they will package these classes as part of the EPIC initiative.
 - Jay proposes that HiPCAT explore broadcasting technologies to send content out in real-time to interested institutions.
 - Is there interest from BCM, UTSW and UTHSCSA or their local communities in these materials if HiPCAT would deliver and broadcast? Interest is expressed from all of these parties.
 - Teaching ‘Introduction to Parallel Computing for Science and Engineering’ in the Fall, ‘Visualization and Data Analysis for Science and Engineering’ in Spring and ‘Distributed and Grid Computing for Science and Engineering’ in Summer.
 - There is a lot of interest in introductory classes in scientific application development, how to use scientific libraries, developing codes that perform well on numerical operations, how to manage scientific data. Put together intro courses for next Spring.
 - No recording/broadcasting of the Fall classes for now, because they are not ready preparing the infrastructure needed for this.
 - Send email to Jay if other institutes are interested in this kind of content.
 - Olin mentions that UH already has such content in the form of short courses.

- Jay likes to differentiate between short courses and full academic classes.
 - Olin mentions that we should get distant-learning coordinators on campuses involved to integrate these courses with local distant-learning courses.
 - Alan mentions that we should try to take viewpoint of students and see what they could get out of such broadcasted courses. There is a big difference between taking a 'real' class and a recorded class.
 - Interactive asynchronous online courses are doable, but this will need funding.
 - Nick mentions that courses should also be distributed asynchronously for students to listen to the lecture later in time.
 - Jay proposes to create a formal project out of parallel computing class content by HiPCAT members. Phil will make sure the executive committee votes on this later in time.
- Technologies for recording and distributing content (video and slides)
 - Jay mentions that MS has a team that developed Conference XP: free, but not cross-platform; it can be used for building Conference XP network across campuses. Tools can be built upon the open API. MS provides technology grants to groups to test-drive this.
 - Alan mentions that we should not go only the MS way, because they only want to extend their stronghold on the IT world. Technology should be cross-platform. Borries agrees on this completely, because his part of the campus runs mainly on Linux.
 - Jay mentions that distinctions have to be made between the node infrastructure to record the class and the format we distribute it as. The students should not care HOW the classes are recorded, so it shouldn't matter if it runs on windows or something else.
 - Borries likes the personal access grid nodes that can be used on a personal computer or laptop without the need to go to a conference room. Alan mentions that Insors has become much better and reliable lately and that it can be bridged to H.323 equipment. It will also be available cross-platform soon. Easy to use for the average computer user.
 - Jay mentions that Conference XP has a bridge for multicast.
 - CDLT should probably look at Conference XP. Jay will talk to MS to see if we can get resources that help us set such a system up.
- Phil adjourns the meeting and mentions that this was one of the best and productive meetings ever (although somebody complained about the dial-in cost ☺)