Computing Newsletter 200912



Computing Newsletter

Disk Space

We continue to add about 5TB of new data a week bringing the total size of LAT data at SLAC to 600TB (400TB Level 1 output, 180TB MC and 20TB reprocessed data). 4 new file 30TB servers arrived at SLAC on Monday and hopefully by the time you are reading this the first one will have been installed averting the need to store older data only on tape. Although the LAT data is all kept on raid arrays with multiple redundant drives we have also acquired an additional 250TB of tapes so that we can continue to keep all data backed up in case of unanticipated disk problems.



Existing fileservers (wain017 - wain021) holding ~150 TB

Instructions for replacing disk or fan

Sun 'Thor' fileserver (~30 TB xroot c

New user disk space (too successful?)

The new user disk space, /afs/slac.stanford.edu/g/glast/users, is gaining users and usage. Currently, 3 TB of disk space is allocated (of which about 1/2 is actually used) amongst 167 users. In addition, there are nine science groups with allocations totaling 345 GB. From this perspective, the new user space is a success.

However, a usage pattern for user space has emerged that is stressing the server. Submitting 100s of simultaneous batch jobs can cause the server to become non-responsive which, in turn, causes batch jobs to stall and eventually fail. In addition, interactive users attempting to access this space will be unsuccessful. The SLAC Computing Division has been alerted of this issue with hopes for improved service. In the meantime, please be aware of the possibility that one can overload that server and affect other users. Batch jobs should be limited to prevent such overloading conditions. This can be done by dribbling in batch jobs a few at a time while monitoring the server. When the CPU utilization exceeds ~50%, you are entering the danger zone. General guidelines for using the SLAC batch system in any substantial way (>50 simultaneous jobs) should also be read to avoid making common mistakes which can unnecessarily burden the file servers.

Pass 7.2 reprocessing

The first 14 months of survey data was reprocessed in October and November with the new Pass 7.2 event classification. The data sample extends from run 239557414 (2008-08-04 15:43:34 UT) through run 277596392 (2009-10-18 22:06:32 UT), spanning 6581 runs and including over 14 billion events. The C&A group is currently evaluating this reprocessed data and, depending on their findings, there may be another reprocessing cycle early next year. See the C&A pages for additional details: Pass7.2 planning, and Analysis User Forum.

Computer Security Training (SLAC)

An email went out earlier this month to all Fermi LAT collaborators who had not yet completed the required computer security training. We were informed by the SLAC Cyber Security Team that beginning in January 2010, all non-SLAC employees who had not completed this training would have their SLAC computer accounts disabled. (The deadline for SLAC employees was July 2009.) These accounts are used for access to the glast-ground web site including data access, interactive logins (Linux), email access, and access to a variety of other web-based services. Don't get stuck!

It is already the case (since Oct 2009) that users who need their passwords reset by an administrator must have first completed this training.

For more information on the course, or if you have questions, contact Marilyn Cariola in SLAC Computer Security at 650-926-2820 (email mcariola@slac. stanford.edu).

LAT Workbook developments

The LAT Workbook continues to evolve, expand and be updated. Some highlights since the last newsletter include: LAT GRBanalysis (new); User and group disk space; Using the SLAC batch farm; pylikelihood analysis (updated); Science Tools environment setup (update, including new SCons section); new astroserver examples; ASP Data Viewer help (updated). View a full chronicle of the updates, or, better yet, just browse through the Workbook

Good-bye and good luck to Navid!

An email from Richard last month

A bittersweet tale: Navid will be leaving us at the end of November to take up a position in Earth Sciences at Goddard. He did his Masters with them and will carry on with a PhD.

Navid has been a key player in our success to date and we can barely thank him enough for all he has done. Wish him the best...

In the meantime, we will be trying to figure out how to fill his shoes. No easy task given the breadth and depth of contributions. He is doing a mind meld with confluence now. At least he will be around to answer questions!

Sniff, Richard

SCons On the Way

(Joanne Bogart)

SAS is in the process of changing build systems from CMT to SCons. This has been going on a long time (first investigation was over two years ago!) and isn't done yet, but the end is nearly in sight. ScienceTools builds are already being generated with the new SCons Release Manager as well as the CMT Release Manager; GlastRelease will take a few months longer. Should you care? That depends on how you use SAS software. End users who use prebuilt binaries will see at most small differences in set up procedures. Those who build from sources, e.g. because their platform is unsupported, will need to know something about how SCons works and how we've chosen to use it. Developers need the same kind of understanding of SCons machinery - the equivalent of requirements files, GlastPolicy package, and so forth - as they currently have for CMT, and may also benefit by learning about SCons analogs of MRvcmt and the CMT Release Manager web pages. Look for more information in a future Newsletter as this transition nears production and documentation becomes available.

Possible topics for this newsletter...

Level 1 developments

New astroserver release

Possible topics for the next newsletter (April 2010?)...

HPSS update Pass 8 reprocessing