

Software Infrastructure Overview: Hot Topics, Coordination, and Miscellaneous

Frederick Luehring
Indiana University
September 21, 2004

Introduction

- The first part of the talk is a list the hot topics and give an overview of the problems.
 - Full details are left to individual talks.
- The second part of the talks is about the organization of SIT work.
- I apologize in advance for my slides which a very dense and don't have any pictures.

Building Issues I

- The nightly build keeps growing:
 - It takes more and more space.
 - Since we keep multiple (currently 7) copies of the debug and optimized builds as well as branch and test nightlies and the build size keeps increasing, afs space is being consumed at a prodigious rate.
 - It takes more and more time.
 - This has been alleviated by incremental builds but will certainly remain a problem.
 - Prior to introducing the incremental builds, the nightlies were taking more than 24 hours. Is the large increase in building time understood?
- Local build and retrying have improved reliability.
 - Still not perfect.
- More platforms need to be supported.
 - SLC3 is nearly ready
 - Lots of work done on MacOSX

Building Issues II

- We urgently need more fast, big memory machines.
 - David Asbury has supplied such a machine for SLC3.
 - One more SLC3 machine is needed for branch releases.
 - Dual SLC3/RH 7.3 operation will continue for a while (at least to 2005). The SPMB should decide how long.
 - Large memory, fast machines for interactive debugging and valgrind are sorely needed.
- Don't use the lxbuild005, it's reserved for TDAQ.
- AlexU has put a 500 MB file size limit on the nightly builds to protect against a problem where a build test generated a multi GB file and killed the nightly.
 - The largest file currently generated by the build is 265 MB!

Building Issues III

- DavidQ has issued a call to reconsider what happens when the major ID of a package changes.
 - Our tags are of the form pkg-ii-jj-kk-[ll] and ii is the major id.
 - Policy is to put use statements like “use pkg-ii-*” in the requirements of dependent packages to catch changes.
 - Some packages simply bypass this warning method by changing the the use statement to be “use pkg-*”.
- However when things are coded in the official way sometimes packages fail even if they are changed.
 - David has proposed changing to the second format of the use statement or changing CMT to stop the problem.
 - People responding to David’s email definitely prefer changing CMT.

Distribution

- There has been considerable work to make the software distribution system more reliable.
 - Thanks go to Christian, Grigori, and Alessandro
- That said work still remains and this will be discussed at length tomorrow in the software distribution session.
 - There are 50 minutes for user feedback on distribution and 50 minutes for user feedback on release testing/validation.
 - Users feedback is critical to this process - so please attend. If you can not attend feel free to supply me with slides containing your comments.
- AlexU is working on running the kit validation in the nightlies so that the final kit is more likely to be good.

Testing & Validation

- Currently there are multiple efforts to test / validate the releases:
 - There is the AtlasTest package which is running on every nightly and developers have contributed many test to.
 - Usually more than half of these tests fail on any given nightly.
 - There is the Run Time Tester (RTT) under development.
 - There is the Kit Validation package for the distribution kits.
 - There are also traditional human based validation work done by the Physics Validation, Reconstruction, and other groups.
- It is my opinion that we need to coordinate these efforts better and improve automation so that we catch problems more consistently and before releases are built.
- I remind everyone of the discussion on this tomorrow during the Distribution Working Group session where user input is invited.

Debugging Issues

- FrankP has raised three issues with running the debugger on Athena:
 1. In many cases, the line numbers are not shown for a large part of the stack. This makes debugging the code difficult or impossible.
 2. We need some instructions (web pages and/or examples) on how to debug Athena.
 3. It would be nice to compile on module in debug mode while the rest of Athena was compiled as optimized.
- We will try to find someone to work on these issues.

Global Flags

- Work is underway to define the global control flags consistently throughout the software.
 - Doing this is very important to making the software easy to use.
 - It will require follow-up with many different developers before the flags are honored consistently throughout the software.
 - Things have been improving.

Release Strategy I

- Recently there has been considerable discussion about our release strategy.
- Our current strategy is:
 - Every ~3 weeks there is a development release (x.y.0)
 - Even though some validation is done after the release is built, bugs found in the release are usually not fixed. Problems are simply noted on a webpage though occasionally problem packages are patched.
 - Distribution kits are usually built for each release and worked on until the kit validation test works. For failed releases no kit is built.
 - Every ~6 months there is a production release (x.0.0) followed by a series of patches (x.0.z).
 - Intensive validation commences as soon as the release is built and bugs are either fixed by a rebuild of the release or a new patch release.
 - Distribution kits are always built and worked on until they install and pass the kit validation tests.

Release Strategy II

- Some users have asked that we do more testing on production releases after they are built before announcing their availability.
 - This would amount to a go - no go decision based on the test results. A no go would result in patching or rebuilding to fix the problems that are discovered.
 - The testing would also be applied to development releases but only used to characterize problems and not to force patching or rebuilds.
- We probably should discuss this at SPMB.
- Also it is not clear to me that the 3 week release cycle is what we want - we often end up with more like a four week cycle in practice.

Release Strategy III

- Zhongliang will probably mention it in the release coordinator's report but I would like to ask everyone (esp. the core developers) to respect the deadlines:
 - Core closes a week before release.
 - The tag collector closes two days before the release.
- I am aware of at least one proposed non-bug fix core change for 9.0.0 that was suggested to be put in today (the day before the release).
 - While there may be legitimate reasons for such a change, the potential to break the release is always present with a last minute change. Also it may require many last minute changes by the developers who are looking for stability.
- Last minute core requests happen most releases.

Documentation / Web / Wiki

- The Developer's Guide needs updating to reflect recent changes in the CMT user setup.
 - Need to review the How-To and Release Building docs too.
- The web pages are sorely in need of updating.
 - Most of the the SIT web pages were last updated by Claire.
 - I have promised to work with Di on updating the web pages but so far I have not had the time.
 - We need to at minimum remove incorrect web pages.
- There has been discussion about the use of Wikis and some groups have been using it.
 - We need to coordinate the use of Wikis with other help & documentation tools.
 - I see lots of partially supported tools with little coordination.

Help List / Savannah

- I have been monitoring the atlas-sw-help list and find that requests for help go unanswered.
 - I have been chasing people to respond but the results have been mixed.
 - Many people use the the developers list in the way that atlas-sw-help was intended to be used.
 - People also confuse atlas-sw-support with atlas-sw-help.
- The SIT Savannah bug group has been dormant however recently Emil has started monitoring it.
 - Emil will say more in his talk.
- Again lack of people is hurting us in responding to requests but still I think that we can do better.

checkreq

- I believe that we should endorse turning on the email messages associated with checkreq after 9.0.0 is complete.
 - 9.1.0 is probably too soon but 9.2.0 might be OK.
 - Jakob's participation is ending because he has taken a job in industry which starts in November. So we need to find someone to take this over as quickly as possible.

User CVS Repositories

- For some time now we have been promising to implement user CVS repositories.
 - Promised “after” CVS migration which is now done...
 - Any user who requested it, would get a user/<username> container created for them.
 - The user would have freedom to freely do what they want in container.
- Now that we have remembered this, we will try to get the scheme implemented.

SIT Meeting Organization

- Currently the SIT group meets every other week and generally uses a standard agenda.
 - We always struggle to get through the agenda.
 - Tried to go through recent problems in one meeting but we were even more pressed for time.
- My new plan is:
 - Rotate what is covered in the SIT meeting and cover only the hot topics are but not every topic every week.
 - Try to discuss coordination and planning issues in these meetings.
 - Relegate detailed discussion on specifically selected topics to the alternate weeks.
 - These meetings would only be attended by the people working directly on the specific topic and not the whole SIT.
 - Interested users could be invited to give feedback at these meetings.

WBS and Staffing Levels

- As usual I am ignoring a request by Tom LeCompte for input to his WBS about the SIT effort levels.
 - I will try to talk to each of you individually to revise the numbers that Tom has. I am certain that substantial revision is required. I apologize to Tom for how long it has taken for me to get started on this.
 - However I do not need a WBS to tell you that we are short of people and that the people that we have are less than ideally distributed over the urgent tasks.

Conclusion

- Coordination and lack of people remain the two biggest issues in my opinion.

Shameless Advertisement

- I am currently looking to hire someone to simplify working in the ATLAS software environment and doing distributed data analysis using local resources at remote sites. This person would be:
 - Providing tools to help the development cycle.
 - Generating and improving python scripts to run Athena
 - Improving the web page and documentation content
 - Easing the use of Athena on the grid to do analysis
 - Interacting with the users to identify problems
 - Helping track help requests about the atlas software
 - Tracking SIT Bugs
 - Answering emails to atlas-sw-help
- Details are in the Atlas web Job Vacancies section.
 - Review of applications begins on October 15