

Software Infrastructure Hot Items

Frederick Luehring
Indiana University
Friday Plenary SW Week
September 15, 2006

Introduction

- Hot areas in Software Infrastructure:
 - Documentation coordination
 - Information protection
 - Inclusion of database release in distribution kits
 - Kit production in nightlies
 - Nightly build performance improvements is a separate talk.
 - Support for a new platforms and compiler versions
 - This is a separate talk.
 - Current RTT status
 - RTT Review is a separate talk.

New SW Documentation Coordinator

- Last February's Documentation Review recommended that ATLAS management appoint a Documentation Coordinator to facilitate development of comprehensive and coherent documentation.
 - The selection was made with input from Dario, Fabiola, and DavidQ.
- Stephen Haywood is the new ATLAS software documentation coordinator.
 - He is responsible for managing (NOT creating!) the ATLAS software documentation.
 - The Documentation Coordinator sits on the CMB
 - Maria Smizanska is the documentation coordinator for the rest of ATLAS (e.g. the Physics Working Groups web pages)
 - Please work with Stephen - he needs everyone's help.

Who am I ?



ATLAS for 13 years:

- **ID Performance** for TDRs - User
- **e/ γ Working Group** – User
- **Computing Reviews** (~7 years ago), inc. Architecture TaskForce
- **SCT End-cap Engineering**

ALEPH: Software, esp. Mini-DST

Little exposure to current Software/Computing:

Pros & Cons

Lots to learn ... looking forward to it



Plans

Sooner

Resource loaded plan in conjunction with Fred (SIT)

Get on top of TWiki

Review Top Level Pages

Identify contacts for all areas

Review Doxygen documentation for key areas – coming out of 2005

Reviews

Chase “TBD” pages

Encourage use of Hypernews

Later

Ensure safe and robust access to documentation 24/7

Identify further “advanced” Workbooks

Information Protection

- ATLAS Management is very concerned about protecting ATLAS data and physics results from unauthorized access and premature release.
 - It is the responsibility of every ATLAS member to protect our data and physics results.
 - An Information Protection Officer will be appointed.
- Traudl Hansl-Kozanecka has spent some time writing up a comprehensive survey of where we are currently exposing ATLAS information. Here is the URL:
http://cern.ch/atlas-computing/temp/2_securityissues/informationProtection.html
- Please send comments to Traudl by Monday.

Overview

- What to protect
 - Data analysis and related communication (most important)
 - Collaboration internal information
 - Not considered here: protection of data on GRID
- Applications and repositories for information exchange
 - Their present security status
 - Comments on individual passwords, keywords for groups
- Plans of IT/CERN
 - Management and maintenance of group lists
 - Comment on VOMS
- Issues concerning SIT
 - Protection of code and secure distribution of code
 - Document available on atlas-computing, but not yet officially approved
 - Feedback welcome

Protect ...

- Physics analyses during the premature analysis stages against non-authorized access
 - includes protection of preparatory documents and graphs, of presentations and discussions in meetings, and of discussions in emails, or in other web-based media like HyperNews or the TWiki
- Data: always write-protected
 - also protected against unauthorized read access
- Software: only accessible to ATLAS members
 - repository for the “Kits”, protect by individual passwords?
 - Note: CVSView (allows download), LXR, Doxygen TreeView
- Protect s/w and data against "malicious" attacks
 - which would be a nuisance rather than a real threat, because copies exist

For real data, software and data can be independently protected. This is not the case for ATLFast, a "weakness" which was already exploited by theoreticians



DB Release

- starting from 12.0.1, DB Release is an integral part of the distribution kit
- uses for setup the `cmtsite/requirements` file generated by the distribution kit at install time

Nightly Kits

- Grigori Rybkine and Alex Undrus have setup the nightly so that it now builds and installs a kit.
 - With help from Alessandro De Salvo, we are now running kit validation on the installed nightly kit.
- The motivations for creating a nightly kit:
 - Creating and validating the kit checks the whole chain used to install the software at the production sites.
 - Eventually we will install the nightly kit in place of the current copy from a local repository used to put the software onto AFS. This will make CERN identical to all of the outside sites using the kit.
 - The RTT validation will need to use the kit when it runs on the grid to test the nightly build.
 - Debug nightly kits will allow users to debug at their local sites.



Nightly Kit (1/5)

Aim:

ensure that once a release has been built, a fully functional distribution kit for it can be built immediately

- build a distribution kit of each nightly build to detect and solve possible problems with its packaging
- install the kit as a way of basic testing
- make the kit (both packaged and installed) available for further testing

Nightly Kit Builds

- Nightlies for project with one package: PackDist (author Grigori Rybkine)
- NICOS page http://atlas.web.cern.ch/Atlas/GROUPS/SOFTWARE/OO/dist/nightlies/global_kits
- Kits created with PackDist for 12.0.3 nightlies in [/afs.cern.ch/atlas/software/builds/kitrel/nightlies/bugfix](http://afs.cern.ch/atlas/software/builds/kitrel/nightlies/bugfix)
- Kits are downloaded with PackDist on a local disk
 - Some problems after AtlasProduction & AtlasOffline swap
- Kits are validated with KV suite (thanks Alessandro for help)
 - The basic tests are OK while DC3 tests fail (due to intrinsic release problems?)

Stability of the nightlies

- Fair unrecoverable failures rate (~4 per month, but spike in July)
- Failures reasons:
 - Machine problems, power cuts (~2 per month)
 - AFS problems, inc. lack of space (~1.2 per month)
 - Atlas.Release is now notified when AFS is full
 - TC (~0.6 per month)
 - Abnormal slowness, CVS shutdown, error in nightly configuration (<0.5 per month)

Runtime Tester (RTT) Status

- The RTT team is busy responding to a set of milestones proposed by a review committee (Laurent Vacavant will talk on this later).
- The RTT team also busy with the ongoing task keeping the RTT running every night.
 - The validation supplied by the RTT is essential when deciding to build a release.
 - Keeping the RTT running reliably to meet this need puts significant load on the RTT team.
- Given the recent growth in demand for RTT testing, there is an immediate need for additional hardware resources to run the RTT on.
 - New resource are being sought at Lancaster and CERN.



Production status

- Latest News


- 5 new packages added to the production runs since last software week
- First run ATN tests from RTT (parsing the ATN part of the unified configuration file)
- Local running fixed
- Pilot jobs introduced (run athena on one event and submitted to the short queue)



Production Status(1)

- Towards 12.0.3
 - Running only N.0.X branch of the nightlies (requested by Fred and Ian)
 - RTT producing results daily (~130 jobs with ~80 running OK)
 - RTT had problems closing down on a few runs
 - An escalating issue is the insufficiency and reliability of available hardware
 - The hardware situation needs considerable attention as demand on the RTT increases

Hardware Resource Issues(1)



- How well is the UCL farm doing?
 - Running to full capacity
 - It is a shared resource, throughput not guaranteed
 - Scheduled/unscheduled power cuts
(next one: Fri pm-Mon am)
 - Recent disk failures
 - Insufficient power of the launch nodes (cannot launch more than one nightly build on same node)
 - RTT has outgrown its environment at UCL

Electronic Calendar

- Luc Goossens has requested that there be a calendar system that will allow us to advertise when members of the production and release building teams are unavailable (on holiday or sick).
- After a brief investigation Luc found the Google calendars seems to work nicely for this purpose.
 - It is my suggestion that people in key positions sign up for Google calendars and start to advertise when they are away.
 - We still need to provide a web page with links to various peoples Google calendars.
 - Before proceeding with this plan, it needs to be endorsed by the SPMB and/or computing management.

Summary

- There is lots of activity in the software infrastructure. Here are a couple of highlights:
 - Feel free to give Stephen Haywood your comments on the software infrastructure.
 - We will also soon need nominees for the Information Protection Officer position.
 - There should be an official request for nominees out soon.
- Much of the new activity will be done by four new / returning people working on software infrastructure at CERN.
 - Igor Kachaev
 - Krzysztof Ciba
 - Jose Antonio Coarasa Perez
 - Vasily Kabachenko