


Future Needs of User Support (in ATLAS)

Dan van der Ster, CERN IT-GS & ATLAS


WLCG Workshop – Prague, Czech Republic

Sunday March 22, 2009

- In ATLAS, the model is **tutorials + help forum (+ validation)**
 - Physics Analysis Workbook “Running on Large Samples”
 - Offline Software Tutorials every ~6 weeks
 - hn-atlas-dist-analysis-help@cern.ch – catch all for both distributed analysis tools (Ganga + Pathena)
 - (Validation is behind-the-scenes automated Functional and Stress Testing)




- DA Support Team (DAST) formed:
 - To relieve developers of the support burden
 - To support Pathena & Ganga through a single forum
 - (the tools are working toward common source code)
 - To maintain documentation, enable users to help themselves
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/AtlasDAST>
- DAST was modeled after the ATLAS production shifts:
 - Reused their infrastructure (scheduling + calendar, some procedures)
- We asked the user community for volunteers to become expert shifters
 - Started Oct 2008 with 4 NA + 4 EU shifters
- Each week, we have 1 NA + 1 EU on shift:
 - Third time zone has no coverage ☹
 - Shifters are responsible for (a) directly helping users, (b) monitoring the analysis services, and (c) helping with user data management issues




- DAST is not a help desk:
 - Support is via an eGroups forum to enable user2user support
- Shifters need a shared interface to label, flag, and privately discuss the various threads/issues.
 - RT, Remedy, Savannah are not appropriate
 - We use a shared Gmail account

Labels	Issue	Status	Priority	Category	Assignee	Created
- Labels	Stephanie Jezequer	ACTIVE				Jan 20
ACTIVE (13)	Adam .. Hurng-Chun (18)	ACTIVE	WAITING	GLITE bulk subm		Jan 20
CLOSING	Cedric Serfon	ACTIVE		Re: Issue with SE at FZK - H		Jan 20
DDM Transfer Req	Ingo Reisinger	ACTIVE		Error in opening SUSY file fo		Jan 20
DQ2 Expert	Olivier Arnaez (2)	ACTIVE		File bad-unregistered at FZK-		Jan 18
DQ2 Savannah	Margar Simonyan	ACTIVE		compilation issues - Discuss		Jan 17
ESCALATED	Olivier .. Daniel (12)	ACTIVE	ESCALATED	GGUS Compi		Jan 17
EU investigating	Renaud, Daniel, Nurcan (5)	ACTIVE	NA please help	GangaPand		Jan 16
EU please help	Chris, Alden, Nurcan (9)	ACTIVE	NA investigating	pathena_uti		Jan 16
FIXED in next releas	helpdesk	ACTIVE		GGUS-Ticket-ID: #45276 TIC		Jan 15
Ganga Expert	ZHU, Alden (2)	ACTIVE	NA investigating	WAITING		Jan 14
Ganga Savannah	helpdesk (2)	ACTIVE	ESCALATED	GGUS WAITI		Jan 14
GGUS	Paolo, Daniel, Benjamin (3)	ACTIVE	ESCALATED	Ganga Expert		Jan 14
NA investigating	iacopo .. Margar, Daniel (4)	ACTIVE	WAITING	Troubles with gan		Jan 14
NA please help	Christopher .. Daniel (4)	ACTIVE	WAITING	Shallow retry cou		Jan 14
pathena Expert						


1. Usual DA issues:
 - Why did my job fail? My job ran yesterday but not today?
 2. User support is not just DA support
 - The user workflow is (a) look for input data, (b) run the jobs, (c) retrieve the output data
 - Need to support more than Ganga/pathena; (especially data management tools).
 3. Users aren't aware of the very nice monitoring:
 - Many users find it more convenient to ask why their job failed, rather than check what the monitoring is showing
 4. Users don't (and might never) know the policies:
 - i.e. where they can run, what inputs they can read, where they can store outputs, which storage locations are temporary/permanent, ...
 - Policies are dynamic and inconsistently implemented
- 3 & 4 above imply that the end-user tools need to
 - fully enforce the policies, and
 - be fully integrated with the monitoring, especially by being aware of site downtimes




- Part II: The Future Needs of User Support




- How does the ideal user get help with a problem?
 1. Read documentation (--help, tutorials, FAQs, mailing list archives)
 2. Ask office mates or other colleagues
 3. Post to mailing lists or discussion forums
 - *(the problem with #2 is that the wide community doesn't always benefit)*
- Today, we get many “stupid” questions on the mailing lists:
 - “My job/data-transfer failed!” at site in posted downtime
 - Question that is already answered in docs, FAQ, or yesterday.
- But are they really stupid? Two opinions:
 1. As users get experience, these questions will disappear.
 2. There are no stupid questions: they indicate weaknesses in the documentation or end-user tools, and we need to improve:
 - User-oriented monitoring
 - Organization and relevance of the documentation
 - User discussion forums
- **I think that we are not enabling users to act ideally today.**



- User-oriented monitoring:
 - AFAICT, users are not aware of and don't look at the monitoring
 - Users look at (a) the *job submission tool* and (b) the *job monitoring tool*. Sometimes these 2 are the same.
 - The relevant monitoring info must be obvious from the end-user tools:
 - E.g. A job listed as “activated” should say how long until it will run. If the selected site is now offline, this should be stated.
 - Current monitoring is too spread around; users need a single entry point.
- “Status Awareness” in the end-user tools:
 - Some tools already look at the service statuses: Ex: GangaRobot blacklisting in ATLAS.
 - Needs to be implemented in all end-user tools:
 - Ex: data retrieve tools should not try a transfer from a down site
 - The information must be updated frequently; sites need to be de-blacklisted ASAP.



- Users are asking questions that are already covered in documentation or other threads
 - Like the site/service monitoring, the documentation is spread around.
 - Tutorials, Twiki docs, FAQs, `<command> --help`
 - And some documentation does not exist:
 - HOWTO knowledgebase of working job configurations
- We also need to enable and encourage user2user support.
- What would Google do?
 - Google Help is a combination of Discussion Forums + Question & Answer
 - **Forum:** Same as eGroups, but we need to be able to “stick” (highlight) some important threads
 - **Q&A:** Users ask questions, and others post answers. Answers are voted on and the marked read only when the real answer is given.



- Will we always need expert user shifters?
 - Most (physicist) users will never become grid experts
 - There will always be only a few who understand the details which result in some job failures
 - Ex: we can never expect a user to recognize that an athena crash was due to an overloaded storage pool.
 - I think that expert user shifters will be needed for the visible future.
- In ATLAS, we estimate that we will need to at least double the manpower on DAST shifters
 - From 1+1 → 2 US + 2 EU (+2 ASIA)
- With more than 2 shifters on at once, communication between shifters (including production shifters) becomes critical:
 - Issue tracking: Gmail probably won't scale with >1 shifter on at once. Need another collaborative support tool.
 - Communication with other shifters: Experimenting with Virtual Control Room

