


Distributed Analysis User Support in ATLAS (and LHCb)


Dan van der Ster, CERN IT-GS & ATLAS
Contributions from Andrew Maier, CERN IT-GS & LHCb

WLCG Workshop – Prague, Czech Republic


Sunday March 22, 2009




- From A. Maier (LHCb):
 - In the last 3 months close to 200 unique users in LHCb
 - Stably increasing trend
 - $\frac{1}{3}$ of the entire collaboration
 - Probably close to the total number expected physicists involved in analysis
 - More than 60 physicists using Ganga on average per week
- For ATLAS, Ganga & Pathena has seen ~500 + ~500 unique users in the past six months:
 - ~125 + ~125 total unique per month
 - The number of users is still expected to increase
- For both experiments, the number of jobs will of course increase, so *we are not yet at the peak user support load.*



- In LHCb, the model is **tutorials + help forum (+ validation)**
 - Ganga introduction and hands-on tutorial is a part of the LHCb core software training (repeated every 2 months)
 - Additional external software training organised 2-3 per year
 - User support through mailing lists:
 - `lhcb-distributed-analysis@cern.ch` – covers all aspects of distributed analysis questions including Ganga
 - `lhcb-ganga@cern.ch` – covers specialised ganga support questions.
 - Validation with SAM tests
- 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



1. Provide help via the DA Help Forum:

- We see two basic problems:
- **How do I do x?**
 - We (create and) forward the user to documentation
- **My analysis doesn't submit / run / complete / output!**
 - View logfiles, check sites, try to reproduce problem; these take *time* to solve.
 - **Escalate to another expert:**

Issue	Escalate to	Email
Ganga Athena	Ganga Developers	project-ganga-developers@cernSPAMNOT.ch
Ganga AthenaMC	Ganga Developers	project-ganga-developers@cernSPAMNOT.ch
Ganga bugs	Ganga Savannah	Ganga Savannah
Glite WMS problems	WMS Support Ticket	wms.support@cernSPAMNOT.ch
DQ2	DQ2 Support	atlas-dq2-support@cernSPAMNOT.ch
User dataset deletion	DQ2 clients how-to	https://twiki.cern.ch/twiki/bin/view/Atlas/DQ2ClientsHowTo
Site or grid issues in general	ADC elog and shifter	atlas-project-adc-operations-shifts@cernSPAMNOT.ch
	GGUS	https://gus.fzk.de/pages/ticket.php
	RT	How to submit RT tickets
pathena bugs	Panda Savannah	Panda Savannah
Athena setup, cmt in the US	Physics Analysis Support	HN-PhysicsAnalysisSupport@bnlSPAMNOT.gov
Analysis tools	Physics Analysis Tools Help	hn-atlas-PATHelp@cernSPAMNOT.ch
Software related issues	General Offline Software Help	hn-atlas-offlineSWHelp@cernSPAMNOT.ch
Database issues, DBRelease files	Sasha Vanyashin	sacha.vaniachine@cernSPAMNOT.ch
If no help from GGUS or RT received	Cloud support	US cloud: Nurcan Ozturk, nurcan@utaSPAMNOT.edu

Oct 1 -> March 4:


621 “conversations”
~125 per month
~4 per day

February 2009:

155 conversations
~5.5 per day

2. Help with Dataset Replication Requests


- Users can request a dataset to be replicated to a new site
 - This can't be allowed freely (model is *jobs to data*, not *data to jobs*)
- If the request is for >10GB transfer and the data is already available within the destination cloud:
 - DAST intervenes and helps the user to process the data in the present location




- 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	Created
- Labels	Stephanie Jezequer	ACTIVE		Cleaning of xxx_OSERDISK	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 for	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_util	Jan 16
FIXED in next release	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					

- Gmail works for us, but it isn't perfect
- Our procedures for using gmail consistently:
 - Gmail Issue Tracking Procedure
 - 1. Open (unresolved) issues are to remain in the Inbox. Closed (resolved) issues are to be archived.
 - 2. Escalated issues are to be considered open until the problem is resolved.
 - 3. Open threads can be in one of 5 states:
 - 1. Requiring attention: these include any thread with no labels, or is starred, or that has an unread reply.
 - 2. Waiting for user response: label these WAITING, then ignore, and finally remove the label when the user has replied.
 - 3. Requiring urgent attention: label these URGENT, and act accordingly.
 - 4. Escalated: label these ESCALATED, and add a "to where" label (see (4)(1)); shifters should inform the user then close the issue when it is resolved. If an issue is escalated but still unresolved after a reasonable amount of time, we should contact "to where" for an update.
 - 5. Fixed in the next release: label these "FIXED in next release", and contact the user after the next release to remind/verify that the issue is fixed, and finally close. Use your judgment to decide if the thread can instead be immediately closed (i.e. if you feel that we don't need to follow up after the release).
 - 4. Labels other than those mentioned in (3) can be used for information purposes, including:
 - 1. to where an issue has been escalated, e.g. Ganga Expert, GGUS, DQ2 Savannah...
 - 2. temporary labels used to track common issues, e.g. mc08 dataset problem.
 - 3. other labels for arbitrary information/tracking purposes.
 - 5. To close an issue, remove any labels mentioned in (3), (4)(1), or (4)(2) and then archive the thread. Labels from (4)(3) can optionally remain on closed issues for reference later.
 - The above guidelines have the following implications:
 - 1. To find issues needing attention, just browse to the Inbox and look for unlabeled, starred, and unread threads.
 - 2. There will be many items in the Inbox that will not require attention. This is OK.
 - 3. Issues will stay in WAITING until the user responds. I don't think we need to contact users if they are too uninterested to reply. I suggest we close inactive WAITING threads after 7 days.
 - 4. Threads that we close will be automatically reopened if the user or anyone else replies to the thread. Thus, you can safely "close" a thread and it will reopen itself if the user doesn't agree with you. Perhaps this means that the WAITING state is redundant, but at least I find it useful to keep these obviously open threads in view, and thus in mind.
 - 5. If another user happens to resolve an issue without DAST intervention, just archive the thread and move on.
 - 6. If you find a thread in an inconsistent state, try to find out its real status and correct the labels.
 - 7. Feel free to create new labels under (4)(2) or (4)(3); please communicate their meanings to the other shifters if they are to persist, or otherwise delete them at the end of your shift week.



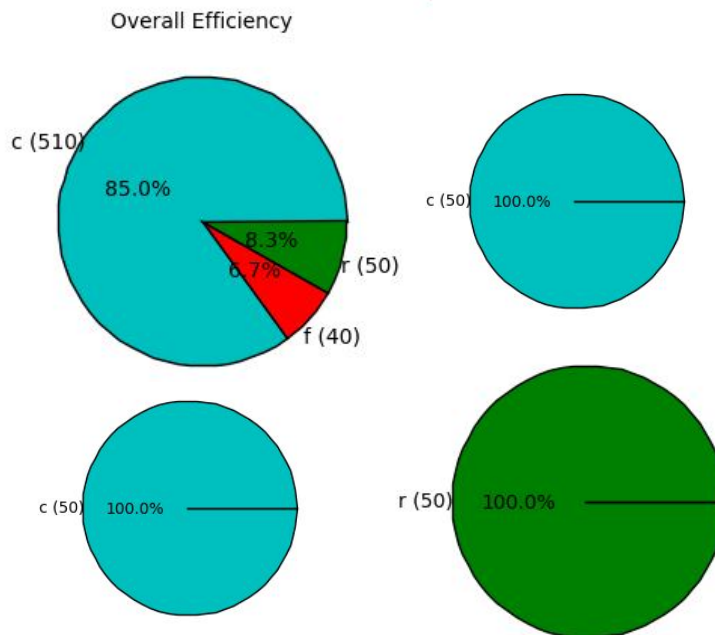
1. Usual DA issues:
 - Why did my job fail? My job ran yesterday but not today?
 2. User support is not just DA support
 - The user workflow is (a) look for input data, (b) run the jobs, (c) retrieve the output data
 - Need to support more than Ganga/pathena; (especially data management tools).
 3. Users 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




- Already saw the automated functional testing: GangaRobot
- Also running large automated stress tests with **HammerCloud**:

74 sites tested; top sites tested >25 times
>50000 jobs with average runtime of 2.2 hours.




>10.5 million files (>3 billion events)



% CPU Used



Events/second



Testing different data I/O configurations (e.g. posix I/O vs copy-and-process)
Also used to evaluate new or changed sites.