

Tier3 Support Working Group: Summary of Recommendations

Alessandro Di Girolamo – CERN IT-ES

Dan van der Ster – CERN IT-ES

on behalf of the Tier 3 Support Working Group

US ATLAS Tier 3 Workshop

June 9, 2010

- Earlier this year, the ATLAS T3 Support Working Group was formed to:
 - Evaluate the impact (especially the effort) in adapting HammerCloud to monitor the Tier3 along the same way of what we do for Tier2s (routine tests, stress tests, publication of results, etc...)
 - Using the experience of the ATLAS Distributed Analysis Support Team, evaluate the support load and new issues compared with the present situation (Tier2 only)
 - Evaluate/give direction to produce (using existing material eventually integrated with grid/tier3 specific material) an updated analysis cookbook
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/AtlasTier3SupportWG>
- This WG has generally wrapped up, and what follows is a summary of the draft recommendations.
 - We fully expect these recommendations to evolve as we get more practical experiences of Tier3s deployments.

- User Support
 - How do the users get started? Documentation and Links
 - How do the users get support when running on a Tier 3?
- Site Support:
 - How do site administrators get support?
 - How to run stress tests on your site?
 - How to run continuous validation jobs?

Tier 3 User Support

- Usage of PanDA/Ganga/Pathena for Tier 3 Analysis means that most of the workflows that worked on Tier 2 analysis should also work on Tier 3.
- User Documentation for the Tier 3 will be integrated to the existing entry-points for Computing and Analysis:
 - **ATLAS Computing Workbook** is the starting point for users:
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/ComputingWorkbook>
 - Links to Ganga/Pathena documentation
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/FullGangaAtlasTutorial>
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/DAonPanda>
- Here are some Tier3-specific docs:
 - <https://atlaswww.hep.anl.gov/twiki/bin/view/UsAtlasTier3/Tier3gUsersGuide>
 - <https://twiki.cern.ch/twiki/bin/view/Atlas/PandaTier3>

- Users have three general types of questions when running their analyses:
 - “Knowledge” problem. E.g. a question or misunderstanding about how to use the grid or one of the tools
 - Problem in the analysis or grid software: e.g. software bug, service outage, etc...
 - Site problem: e.g. software installation problem, storage inaccessible, etc...
- When users have a problem, they proceed through a few levels of support:
 - Ask a colleague
 - Ask a local or regional support team
 - Ask the central ATLAS support
- We expect that the above will continue to apply to Tier 3's

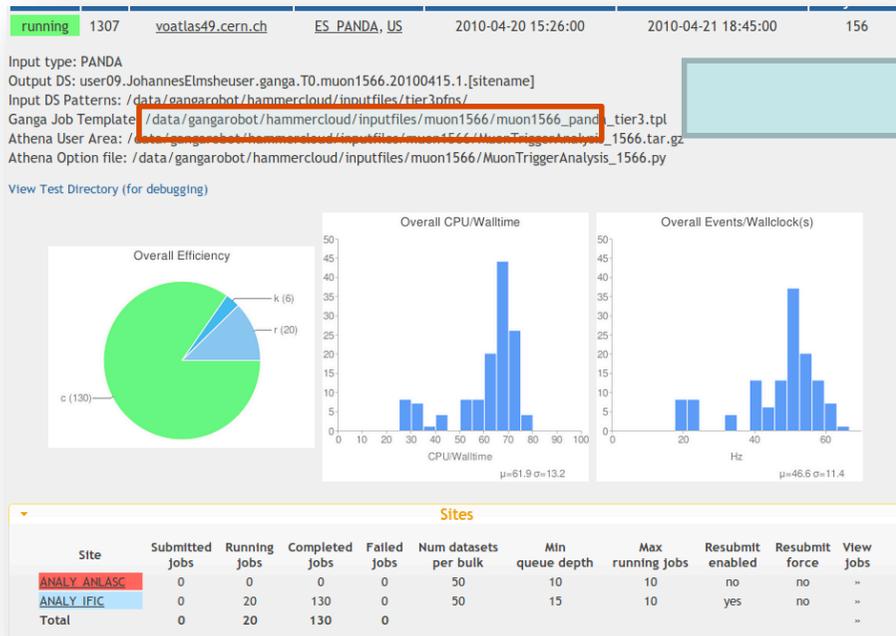
- When users need central support, the existing Tier 2 support infrastructure will apply
- Users use the aforementioned documentation, and when they still have a question they use the **DA help eGroup**:
 - hn-atlas-dist-analysis-help@cern.ch
 - This is a place for user-to-user discussions and support, but is also supported by a team of expert shifters: the Distributed Analysis Support Team (DAST)
- In the case of a site problem, the DAST shifter will need to contact the Tier 3 site managers:
 - We predict that most Tier 3 users will have close contacts with the site manager, so the user could contact the local support directly.
 - Otherwise, a site ticket or mail will be sent (more on this later).

Tier 3 Site Support

- In the case that a central support shifter detects a problem with a Tier 3 site, the shifter will:
 - Submit a GGUS or RT ticket to the site, if possible, and
 - Send a mail to the relevant atlas-support-cloud eGroup:
 - Relevant cloud support persons should be connected to this list so that issues can be routed to the Tier 3
 - (For the US the list is atlas-support-cloud-us@cern.ch, and at the moment it has one member: racf-wlcg-announce-1@lists.bnl.gov)
- Why this model?
 - We polled all ATLAS clouds about local plans for support and found:
 - 8/10 clouds have Tier 3's associated with a Tier 2, so the current procedures for ticketing a site via GGUS and using the atlas-support-cloud-xy eGroup will apply to Tier3's.
 - 2 clouds have a dedicated FTE to support Tier 3s.
 - Some Tier 3's may not be connected with a cloud. In this case, the site should provide a contact address.

- Tier 3 system managers should join the global eGroups list atlas-adc-tier3-managers@cern.ch, which will be the primary communication channel *to* and *between* Tier 3 sites:
 - This list has two main goals:
 - To act as the primary contact point for all the ATLAS Tier3 system managers
 - To facilitate discussions within the Tier 3 community regarding software, configurations, or other deployment and maintenance issues
 - “System managers” can mean local system administrators and/or local ATLAS contact persons; the list is not intended for users or user support.

- Background: Tier 2 sites are frequently tested with regular analysis tests:
 - Distributed Analysis Functional Tests (GangaRobot)
 - Distributed Analysis Stress Tests (HammerCloud)
- These tools have three main benefits:
 - Stress tests are useful to *get first-hand experience* with high loads of user jobs
 - E.g. during site commissioning or to evaluate different site storage/network configurations
 - Functional tests are useful to *continually validate* the functionality of the site
 - Participation in these tests is an effective way to *get wide-visibility* for your site
 - E.g. compare against similar sites to improve performance
- We are working to make the same tools available to Tier 3's



```
[atlldm07] /data/gangarobot/hammercloud/inputfiles/tier3pfns > ls
ANALY_ANLASC ANALY_IFIC
[atlldm07] /data/gangarobot/hammercloud/inputfiles/tier3pfns > ls -l
total 88
-rw-r--r-- 1 gangarbt cg 846 Apr 20 15:20 ANALY_ANLASC
-rw-r--r-- 1 gangarbt cg 78000 Apr 16 14:12 ANALY_IFIC
[atlldm07] /data/gangarobot/hammercloud/inputfiles/tier3pfns > cat ANALY_ANLASC
root://ascvmxrdr.hep.anl.gov/xrootd/mc08.105597.Pythia_Zprime_tt3000.recon.AOD.
e435_s462_s520_r808_tid091860/AOD.091860_000003.pool.root.1
root://ascvmxrdr.hep.anl.gov/xrootd/mc08.105597.Pythia_Zprime_tt3000.recon.AOD.
e435_s462_s520_r808_tid091860/AOD.091860_000120.pool.root.1
root://ascvmxrdr.hep.anl.gov/xrootd/mc08.105597.Pythia_Zprime_tt3000.recon.AOD.
e435_s462_s520_r808_tid091860/AOD.091860_000159.pool.root.1
root://ascvmxrdr.hep.anl.gov/xrootd/mc08.105597.Pythia_Zprime_tt3000.recon.AOD.
e435_s462_s520_r808_tid091860/AOD.091860_000045.pool.root.1
root://ascvmxrdr.hep.anl.gov/xrootd/mc08.105597.Pythia_Zprime_tt3000.recon.AOD.
e435_s462_s520_r808_tid091860/AOD.091860_000025.pool.root.1
root://ascvmxrdr.hep.anl.gov/xrootd/mc08.105597.Pythia_Zprime_tt3000.recon.AOD.
e435_s462_s520_r808_tid091860/AOD.091860_000146.pool.root.1
[atlldm07] /data/gangarobot/hammercloud/inputfiles/tier3pfns >
```

To run these tests, your site needs to be remotely accessible via:

- a Panda queue (ANALY_yoursite)
- a grid CE (this is still in development)

- HammerCloud determines the performance (e.g. Events/Second) of your site under stress
- To run a HammerCloud test you need to:
 - Download some (10's of GB's) data to your site which is useful to our test analysis (typically mc data with muons, e.g. mc09*mu*AOD*)
 - Send us a list of physical-file-names for this data
- Then we arrange a test time and we will submit the jobs.

- The requirement for sending a PFN-list could be removed if we used the following “trivial file catalog” method:
 - Store a site-specific ***pfnPrefix*** centrally or in a site environment variable
 - E.g. “/data/atlas/”
 - dq2-ls would provide a new option “--relativePfnList” which returns a list of relative file names, e.g.
 - ./data09_7TeV/AOD/x/y/file1
 - ./data09_7TeV/AOD/x/y/file2
 - Sites would store files with this convention in the ***pfnPrefix*** directory
 - /data/atlas/data09_7TeV/AOD/x/y/file1
 - /data/atlas/data09_7TeV/AOD/x/y/file2
- pathena/ganga/HammerCloud could then trivially generate the pfnList given just a dataset name
 - users could then submit as they do with a Tier 2, e.g. “pathena --inDS myDS --outDS y --site ANALY_MYTIER3 myOptions.py”

- Tier 3 User and Site Support is being modeled to be very similar to Tier 2 Support
- Key points:
 - To get central help, Tier 3 users will use the Computing Workbook and DA Help eGroup
 - When a site contact is needed, ATLAS shifters will send a GGUS/RT ticket and mail the atlas-support-cloud-xy list
 - Site managers should be connected to the atlas-adc-tier3-managers list to participate in ATLAS-wide Tier 3 discussions
 - Stress testing with HammerCloud is available to help commission your site and to get visibility