

# Table of Contents

<b>ADCoSWIP .....</b>	<b>1</b>
<b>Requirements.....</b>	<b>2</b>
<b>CHECKLIST.....</b>	<b>3</b>
<b>Most Common Mistakes by Shifters.....</b>	<b>5</b>
<b>MC production.....</b>	<b>6</b>
General guide-line (Fast Troubleshooting).....	6
monitor.....	6
Job-states definitions in Panda.....	6
Task-states definitions in Panda.....	7
What to do when.....	7
A task is failing.....	7
A task is not assigned to site.....	8
A site is heavily failing.....	8
A site is not getting jobs.....	8
Details about errors.....	8
'Lost Heartbeat' error.....	8
'Job killed by signal 15'.....	9
'Exception caught in runJob'.....	9
'Pilot has decided to kill looping job'.....	9
'Get error : open connection to sename' or 'Get error: dccp/rfcp failed'.....	9
'Get error: with guid xxx not found at'.....	9
'GUID for xxxx not found in DQ2 !'.....	9
'Put error'.....	9
'/opt/lcg/bin/lcg-cr !'.....	9
'Transformation not installed in CE !'.....	9
'Transfer time out'.....	9
'No PFN found in catalogue for GUID'.....	10
ATLAS Software Releases problem.....	10
Missing input files on T2_PRODDISK.....	10
Missing files.....	10
checksum errors during dq2-get/lcg-cp.....	11
Waiting Jobs Procedure.....	11
How to know if the problems is task related or site related ?.....	12
Group production jobs.....	12
Group production jobs experience/hints.....	12
Task in pending state for long time.....	12
<b>DDM.....</b>	<b>13</b>
DDM dashboard shows timeout/SRMV2STAGER errors.....	13
What to fill into GGUS ticket subject (short description).....	13
Which Problems to report.....	14
Tier-0/Tier-1/Tier2 Data exportation.....	15
What to do when a site has no FREE disk space in space tokens?.....	15
What to do when subscriptions are not processing?.....	16
Checking blacklisted sites in DDM.....	16
Checking the deletion error rate per site.....	16
<b>Panda queues.....</b>	<b>18</b>

# Table of Contents

<b>Central services.....</b>	<b>19</b>
<b>Frontier.....</b>	<b>20</b>
<b>Miscelania.....</b>	<b>21</b>
Contacting the Cloud Support.....	21
Contacting the ticket portal support.....	21
Pilot Factories and methods.....	21
<b>Communication and Organization.....</b>	<b>22</b>
<b>Elog Management.....</b>	<b>23</b>
Choosing the right criticality in eLog.....	23
Replying to eLog entry.....	23
<b>Ticket Management.....</b>	<b>24</b>
Site naming convention - exceptions.....	24
General Rules.....	24
How to Submit GGUS Team-Tickets (direct routing to sites).....	24
Overview of Jira trackers which ADCoS shifter might need.....	25
How to find tickets in GGUS.....	25
Shift Reports and tickets.....	26
Ticket format.....	26
GGUS.....	26
Production or Validation Jira Bug Reporting.....	26
<b>Downtimes.....</b>	<b>28</b>
<b>Monitoring tools.....</b>	<b>29</b>
<b>Daily SHIFT report.....</b>	<b>30</b>
<b>Trainee evaluation report.....</b>	<b>31</b>
<b>ADC Virtual Control Room.....</b>	<b>32</b>
Troubleshooting.....	32
<b>GGUS ATLAS TEAM membership.....</b>	<b>33</b>
<b>Useful links.....</b>	<b>34</b>
<b>Tutorials.....</b>	<b>35</b>
<b>Shift Credits.....</b>	<b>36</b>
Expert Duties.....	36
<b>TEAM MEMBERS.....</b>	<b>37</b>

# ADCoSWIP

FOLLOW ME ON [twitter](#)

# Requirements

- ADCoS Shifter has to be an ATLAS member.
- All current ADCoS Shifter are registered in <https://e-groups.cern.ch/e-groups/Egroup.do?egroupName=atlas-project-adc-operations-shifts>
- **Grid Certificate requirements**
  - ◆ Every ADCoS Shifter on duty (Trainee, Senior, or Expert) is required to have a valid grid certificate registered to VO atlas at the time of the shift, see [WorkBookStartingGrid](#). When an ADCoS Shifter starts his/her shift without the valid grid certificate registered to VO atlas his/her shift booking will be cancelled and he/she will not get any OTP credit for the shift. Recurring certificate issues may result in discontinuation of possibility to sign up for ADCoS Shifts.
  - ◆ Valid grid certificate of the shifter must be in /atlas/team VOMS group, and that addition is done by ADCoS Coordinators at the Step 2 of Trainee Shifter setup procedure. If for some reason your certificate is not yet in /atlas/team VOMS group, ask **in advance** ADCoS Coordinators to add it. Particularly don't forget to do that when you get a completely new certificate.
  - ◆ Check whether you are able to submit ATLAS TEAM GGUS ticket, [https://ggus.eu/?mode=ticket\\_team](https://ggus.eu/?mode=ticket_team) (you should see TEAM option on you screen). Valid grid certificate (1) in your browser (2) with /atlas/team VOMS group (3) known to GGUS (4) is required for opening a TEAM ticket.
- **OTP requirements**
  - ◆ It is strictly forbidden to book more than 1 shift within 24 hours. Overcoming this rule may result in discontinuation of possibility to sign up for ADCoS Shifts.
  - ◆ ADCoS Shifter does book shifts in her/his name and does shift as the person who booked the shift. It is forbidden to book shift in one persons name on behalf of a different persons. The ADCoS Coordinator can book shifts on behalf of different persons, such a shift will be booked in name of the Shifter.
  - ◆ In case of an emergency situation leading into impossibility to take shift ADCoS Shifter immediately notifies ADCoS Coordinator, and the ADCoS Shift Captain of the shift timezone. ADCoS Coordinator then cancels shift booking in OTP. ADCoS Coordinator or ADCoS Shift Captain may announce a `Call for shifters` to find a replacement Shifter. Contact can be found on Team page
  - ◆ Booked shifts can only be cancelled 45 days in advance, after this time interval shifter has to find a replacement

# CHECKLIST

**NEW** For ADCoS Shifters [↗](#) page has useful CHECKLIST links on one page.

- Before coming to your shift, make sure that you fulfill shift requirements!
- Open Jabber and say hello in the Virtual control Room to let ADC/ADCoS community know you are on shift.
- **List of current shifters** [↗](#) to check who covers each shift.
- Check recent updates and additions to ADCoS procedures in this TWiki, as they will be marked with **NEW** stamp.
- If TEAM GGUS ticket does not work for you, please **DO NOT SUBMIT** non-TEAM tickets. Ask other shifter, ADCoS Coordination, or ADCoS experts to submit a TEAM ticket. When you are opening TEAM GGUS, please decide which priority you need, see [How\\_to\\_Submit\\_GGUS\\_Team\\_Tickets](#). There are rare cases of top priority problem for Tier-1s and especially for CERN-PROD.
- Open a browser with these tabs:
  - ◆ Production (BigPanda):
    - ◇ Error distribution page: <http://bigpanda.cern.ch/errors/?jobtype=production> [↗](#)
    - ◇ Job distribution in regions:
      - <http://bigpanda.cern.ch/dash/production/?cloudview=region> [↗](#)
        - Region in this context means home cloud of the sites. As many sites participate in multicloud production (see <http://bigpanda.cern.ch/sites/> [↗](#)), jobs from several clouds can run on given site. This view shows all jobs from all clouds which are running on the given site. In case it is necessary to check jobs running on given site for given cloud, there is cloud view (<http://bigpanda.cern.ch/dash/production/> [↗](#))
      - ◇ Job evolution in all clouds: set of ganglia plots [↗](#)
    - ◆ DDM Dashboard: <http://dashb-atlas-ddm.cern.ch/ddm2> [↗](#) (please check [DDMDashboardHowTo](#) if you're new to the Team)
    - ◆ ADC eLog: [ADC eLog](#) [↗](#)
    - ◆ GGUS interface (TEAM view): [GGUS Team- Tickets view](#) [↗](#) or [Open TEAM tickets view](#) [↗](#)
    - ◆ ADCoS Twiki
    - ◆ Central services machine status: [https://sls.cern.ch/sls/service.php?id=ADC\\_CS](https://sls.cern.ch/sls/service.php?id=ADC_CS) [↗](#) (see [#Central\\_services](#))
  - At the beginning of your shift please have a look at the **Known Problems** and [Daily\\_SHIFT\\_reports](#) of previous shifters
  - Check [ADC eLog](#) [↗](#)
  - Check the status of the TEAM tickets in GGUS and hand-over. See [TicketManagement](#)
  - Remember to report every action in eLog. (for 'new' entry, click on existing entry first. If you solve an issue, put [SOLVED] in the eLog subject)
  - Remember to check if sites are in Scheduled Downtime before opening bugs.
  - Check the deletion error rate and deletion backlog
  - Check recent changes in queue status: <http://bigpanda.cern.ch/dev/incidents/> [↗](#)
  - Check [http://bigpanda.cern.ch/jobs/?jobtype=production&display\\_limit=100&jobstatus=waiting](http://bigpanda.cern.ch/jobs/?jobtype=production&display_limit=100&jobstatus=waiting) [↗](#), follow [#Waiting\\_Jobs\\_Procedure](#)
  - Check jobs in transferring state [↗](#) for long time
  - Once per shift please check [Frontier](#) [↗](#) status and [Squids](#) [↗](#) ( Known exceptions ). See also [Frontier](#)
  - Once per shift please check for tasks that have cancelled jobs but not assigned to any site [http://bigpanda.cern.ch/jobs/?jobtype=production&computingsite=&display\\_limit=100&jobstatus=cancelled](http://bigpanda.cern.ch/jobs/?jobtype=production&computingsite=&display_limit=100&jobstatus=cancelled) [↗](#)
  - Once per shift please check for tasks long in submitted state (longer than one month) [http://bigpanda.cern.ch/tasks/?status=pending&statenotupdated=720&display\\_limit=100](http://bigpanda.cern.ch/tasks/?status=pending&statenotupdated=720&display_limit=100) [↗](#)

## ADCoSWIP < Main < TWiki

- At the end of shift Senior shifter is obliged to submit daily shift summary report. Trainee shifters do not submit daily report.
- Senior shifter is obliged to submit trainee evaluation report if trainee was present.

Some of these tasks might require for shifter to contact the ADCoS Expert. In case there is no ADCoS shifter on duty, shifter can ask ADCoS Coordinators.

# Most Common Mistakes by Shifters

- Opening regular GGUS ticket rather than GGUS Team ticket.
- Opening duplicate GGUS ticket . Please don't forget to check the list of open tickets before submitting a new GGUS.
- Reopening the closed GGUS ticket or updating the existing one, rather than opening a new GGUS ticket when the issue/problem is different from the one in the existing GGUS ticket. Please check with the expert shifter if you are not sure if it's a new problem or different manifestation of the one which has already been reported.
- Submitting a GGUS ticket on the site in downtime . Please always check the AGIS Downtime Calendar before opening a new ticket.
- Forgetting to write the site name in the subject of the GGUS ticket. It is strongly advised to start the subject with the site name. That will make browsing by ticket subjects (GGUS/ELOG) much easier. On the other hand, if the site name is at the end of a lengthy subject line, it may not show on the summary list of GGUS tickets.
- **NEW** Forgetting to add cloud support atlas-adc-cloud-[CLOUD]@cern.ch in CC field of the GGUS ticket, or adding an address atlas-support-cloud-[CLOUD]@cern.ch which GGUS can't process. Please always add the cloud support atlas-adc-cloud-[CLOUD]@cern.ch in GGUS CC.
- **NEW** Coming to shift with expired grid certificate, or with a new certificate not in /atlas/team VOMS group, hence having problem opening GGUS Team ticket. Before coming to your shift please verify that you are able to open GGUS Team ticket.
- Forgetting to put an ELOG entry after opening a new GGUS/Jira ticket. Major status updates, as well as closing the ticket, need an ELOG entry as well.
- Opening a new ELOG thread on the evolving issue, which already has entry(s) in ELOG. Instead please continue the existing thread.
- Forgetting to submit an evaluation report of the participant trainee shifter.
- Using email address from TWiki without editing it to exclude SPAMNOT . Remove the word SPAMNOT, otherwise the email will bounce back.
- **NEW** Submitting ticket about mcXX\_valid tasks to validation jira. As this twiki clearly states, only tasks starting with valid should be reported in validation jira. mcXX\_valid tasks should be reported in ADCSUPPORT jira

# MC production

## General guide-line (Fast Troubleshooting)

- Spot sites with major problems
  - ◆ Start chasing major failures in both interfaces: f.i. cloud failing 100% of data import/export.

### monitor

- Go to error distribution page [↗](#)
  - ◆ First look for sites with a lot of failing jobs
  - ◆ Then look for tasks with a lot of failing jobs
  - ◆ Compare them to know if problems are site related or task related
    - ◇ If task is found to be failing at several sites, probably it would be **task** problem, then file an ADCO-Support Jira bug report [↗](#) for non validation tasks or Validation Jira bug report [↗](#) for validation tasks.
    - ◇ When task is failing due to missing files, please follow the #Missing\_files procedure
    - ◇ If all jobs are failing at one single site probably it would be a site related problem, then file a GGUS team ticket.

## Job-states definitions in Panda

There are 10 values in Panda describing different possible states of the jobs, these are:

- defined : job-record inserted in PandaDB
- assigned : dispatchDBlock is subscribed to site
- waiting : input files are not ready
- activated: waiting for pilot requests
- sent : sent to a worker node
- running : running on a worker node
- holding : adding output files to DQ2 datasets
- transferring : output files are moving from T2 to BNL
- finished : completed successfully
- failed : failed due to errors

The normal sequence of job-states is the following:

```
defined -> assigned -> activated -> sent -> running -> holding -> transferring -> finished/failed
```

If input files are not available:

```
defined -> waiting
```

then, when files are ready

```
-> assigned -> activated
```

And the workflow is:

- defined -> assigned/waiting : automatic
- assigned -> activated : received a callback for the dispatchDBlock. If jobs don't have input files, they get activated without a callback.
- activated -> sent : sent the job to a pilot



- sent -> running : the pilot received the job
- waiting -> assigned : received a callback for the destinationDBlock of upstream jobs
- running -> holding : received the final status report from the pilot
- holding -> transferring : added the output files to destinationDBlocks
- transferring -> finished/failed : received callbacks for the destinationDBlocks

The job brokering for production is listed in PandaBrokerage#Special\_brokerage\_for\_production

The delay for job rebrokering is listed in PandaBrokerage#Rebrokerage\_policies\_for\_product

## Task-states definitions in Panda

- registered : the task information is inserted to the JEDI\_Tasks table
- defined : all task parameters are properly defined
- assigning : the task brokerage is assigning the task to a cloud
- ready : the task is ready to generate jobs
- pending : the task has a temporary problem
- scouting : the task is running scout jobs to collect job data
- scouted : all scout jobs were successfully finished
- running : the task is running jobs
- prepared : outputs are ready for post-processing
- done : all inputs of the task were successfully processed
- failed : all inputs of the task were failed
- finished : some inputs of the task were successfully processed but others were failed or not processed since the task was terminated
- aborting : the task is being killed
- aborted : the task is killed
- finishing : the task is forced to get finished
- topreprocess : preprocess job is ready for the task
- preprocessing : preprocess job is running for the task
- tobroken : the task is going to broken
- broken : the task is broken, e.g., the task definition is wrong
- toretry : the retry command was received for the task
- toincexec : the incexec command was received for the task
- rerefine : task parameters are going to be changed for incremental execution

For more details, see [https://twiki.cern.ch/twiki/bin/view/PanDA/PandaJEDI#Transition\\_of\\_task\\_status](https://twiki.cern.ch/twiki/bin/view/PanDA/PandaJEDI#Transition_of_task_status)

## What to do when

### A task is failing

- If the task is assigned to CERN cloud to a queue different from CERN-PROD, i.e. CERN-BUILDS, CERN-RELEASE, CERN-UNVALID, CERNVM, CERN\_8CORE, then forget about it. If the failing task is running in the CERN-PROD queue, then please follow standard procedure.
- Validation tasks. Submit ATLAS validation Jira ticket for **validation tasks** (those beginning with **valid**). There's no need to file a bug for those tasks with small number of failures or to report bugs that has been already reported before:
  - ◆ Make sure that the bug has not been reported before.
- Other tasks -those **not** beginning with **valid**: use ADCO-Support [🔗](#) for non validation tasks with high failure rate.

## A task is not assigned to site

- To find tasks that are not assigned to site and have cancelled jobs check [http://bigpanda.cern.ch/jobs/?jobtype=production&computingsite=&display\\_limit=100&jobstatus=cancelled](http://bigpanda.cern.ch/jobs/?jobtype=production&computingsite=&display_limit=100&jobstatus=cancelled), check the task then please follow standard procedure to report the problem. To check task one should click on task number on top of the page and then click on task name on newly opened page. Consider each task as a failing one and report properly.

## A site is heavily failing

- If the burst of errors are restricted to less than few hours and and there is no more error (from Panglia plot, the increase if failed jobs is sharp and flat since then), no action to take.
- If jobs are continuously heavily failing:
  - ◆ If the site is in downtime, the site should have already been automatically set to test (check [sites](#) and [incidents](#) pages)
  - ◆ Make sure is a true site issue (not Athena issue for example)
  - ◆ File a GGUS Team-Ticket as described in this section with cloud responsible in CC.

## A site is not getting jobs

- Remember if the site has no jobs assigned, there's no chance to run.
- Check that software versions requested by current are installed at the site ([monitoring](#)).
- Check if site/queues are online
  - ◆ If queue is offline and site not in downtime and if there's no incident ongoing, contact the cloud responsible, and file eLog
  - ◆ If queue is offline and there's no incident ongoing, contact the cloud responsible, ADCoS coordinators and file eLog.
- Check if pilots run in the last hours at the site
  - ◆ If you find problems, fill elog and contact the cloud responsible and the pilot factory responsible
- More elaborated checks to be done by squad

## Details about errors

Panda error codes (mapping of error code to diagnostics message) can be found here:  
<https://twiki.cern.ch/twiki/bin/view/AtlasComputing/PandaErrorCodes>

### 'Lost Heartbeat' error

An update by the pilot to the Panda server is sent to every 30 minutes. If there is no update within 6 hours, the job is declared 'lostheartbeat'. When the job is finished, the pilot will try to update the Panda Server 10 times separated by 2 minutes. PandaPilot contains all details about Pilots.

- Most common reason : the local batch system has killed the job because it used more than the accepted resources (CpuTime, WallTime, memory). The ATLAS requirements are published in the [VOid Card](#). By comparing similar jobs on different sites, try to identify which is the problematic variable. In this case, the number of failing jobs should be spread over time
- Site or batch system is broken : The failing jobs should be spread over a period of few minutes.
- CE has lost track of the job

If you ticket the site, provide an example to the site (jobID to be documented).

### **'Job killed by signal 15'**

The local batch system issued a warning to the pilot informing that the job would be killed soon. The job stopped himself to be able to report the log.

### **'Exception caught in runJob'**

No documentation yet

### **'Pilot has decided to kill looping job'**

No documentation yet

### **'Get error : open connection to sename' or 'Get error: dccp/rfcp failed'**

The pilot is not able to copy the input file from the SE. It means that the SE is totally or partially broken. Cross-check with DDM transfers.

### **'Get error: with guid xxx not found at'**

The dataset is supposed to be available at the site but the file catalog scan reveals that the file is not at the site. No action for the moment.

### **'GUID for xxxx not found in DQ2 '**

It usually means that a file was remove from dataset definition after the task defined the input files. In most case, the file was found as not available on SE with no other replica to recover. For the moment, there is no action from the shifter. The jobs will fail quickly and it is up to GDP to treat the task properly.

### **'Put error'**

No information yet

### **'/opt/lcg/bin/lcg-cr '**

The job is not able to copy the file on SE or register the file in LFC. The pilot tries the command twice with file deletion in between (not correct for the moment)

### **'Transformation not installed in CE '**

No documentation yet

### **'Transfer time out'**

The output was not transfered in time. If you manage to find the transfer on DDM dashboard, report it as any other failing transfer. If you cannot find it on the DDM dashboard, write Elog with all information you can get about the issue and then send email with the link and explanation to the cloud support to check the activity of the FTS channel.

## 'No PFN found in catalogue for GUID'

No information yet

## ATLAS Software Releases problem

- Jobs may fail if the needed software release is not present or badly installed. The procedure to follow is to notify ATLAS SW managers and ask for re-installation of the release at the site.
  - ◆ Shifters can check in the ATLAS SW installation system [the status of the release at the site to cross-check](#)
- We will discuss about asking GGUS team to add special tag when submitting TEAM Tickets to address this specific problem to the SW responsible, but the interim solution to follow is:
  - ◆ Open normal GGUS team ticket and after that, set the "type of problem" to "VO Specific Software", and select "VO Specific"="yes", providing involved ATLAS release, site name and CE.  
and put atlas-grid-install@cernNOSPAMPLEASE.ch in CC

## Missing input files on T2\_PRODDISK

- In case of missing files in T2\_PRODDISK: check if there is a deletion backlog in the cloud (<http://bourricot.cern.ch/dq2/deletion/#period=1>).
  - ◆ If yes, **do nothing (do not file GGUS to site, do not file DDM ops Jira)**, only **do create eLog** [about your observation](#).
  - ◆ If no, post a DDM ops Jira [ticket to cloud support](#).

## Missing files

- For missing files as inputs for MC jobs taking inputs from T2 PRODDISK, please follow instructions at Missing input files on T2\_PRODDISK.
- Otherwise, follow these instructions:
- Sometimes files look missing from the site but they were actually never registered or the task is misconfigured. Two simple steps that could help to understand what is happening (in case the errors are ambiguous):

In case you have difficulty in the procedure, please contact the expert either at the #ADC\_Virtual\_Control\_Room, or via the ADCoS team ML (be aware: if there is now answer after 15 min in the chat, send an email to the list)

Follow the procedure (slides [slides](#))

1. From the panda job page where you found the missing file; **click the file name**
2. Then on the page opened, you will find a SURL ( `srms://...` ) or a list of them
  - ◆ If a replica for the site investigated is not listed then it is not a site problem
  - ◆ If a replica is listed you should check the file with `lcg-ls -l`

```
lcg-ls -l SURL
```

    - ◆ if it is ONLINE or AVAILABLE, then try to download it using `lcg-cp`

```
lcg-cp --vo atlas <SURL> <LOCALFILE>
```

      - ◆ If it is accessible it was a transitorial problem
      - ◆ If it is unaccessible it is a site problem
3. If it is a site problem
  1. Open GGUS Team ticket [to the site](#) and report the files that have been lost.

2. Notify the cloud contact about the missing files by adding the following address in cc filed of GGUS ticket
  - ◇ atlas-adc-cloud-[CLOUD]@cern.ch [where [CLOUD] stands for CA, CERN, DE, ES, FR, IT, ND, NL, RU, TW, UK, US] (see #ContactingCloudSupport)
4. In any case (either a site problem or not a site problem)
  1. Open DDM ops Jira [🔗](#) ticket to notify data management team about the lost files with
    - ◇ the lost files information (SURL, file name and the associated dataset)
    - ◇ the panda job information (a link to the panda job and TaskID)
    - ◇ a link to the GGUS ticket
    - ◇ in Mail Notification Carbon-Copy List add cloud support and task owner
    - ◇ **Important!** If the reported issue is urgent, please explicitly state this in eLog. ADCoS Expert. ADCoS Expert will then escalate issue to AMOD.
  2. When opening DDM ops Jira [🔗](#) please follow the pattern
    - ◇ Task type Task xxxxx: task status in XX cloud
    - eg
      - MC production Task xxxxx: waiting in XX cloud
      - Reprocessing Task yyyyy: input RAW file corrupted at SITE
      - Group production Task zzzzz: waiting for input in ZZ cloud
    - ◇ If you have additional information, eg. the input dataset has been deleted, add this information to Jira:
      - Dataset xxxxx deleted from SITE
    - ◇ It is important for ddm-ops experts to know category of the problem(task type, dataset deleted) and the location (cloud, site), so that they can start do some real work without navigating through panda pages to collect this information. TaskID is necessary for ProdSys.

## checksum errors during dq2-get/lcg-cp

DDM tools can check file (on Storage) consistency with LFC/DDM catalogs (filled when the file is registered in DDM and before any replication) on file by file basis. As soon as you have a doubt about a file, follow the procedure:

1. Check the file consistency:
  - ◆ If the file belongs to a tid dataset: run a script to check the consistency of the file on the Storage at the source T1: [link](#)
  - ◆ If the file does not belong to a tid dataset : use dq2-get which will copy the file locally, compute the checksum and report inconsistency
2. If the file is not corrupted on the source T1 Storage:
  - ◆ Run dq2-get to check the file consistency on the SE used by your application. If the file is correct, go to next point . If the file is not correct, fill a Jira ticket to DDM Ops [🔗](#) providing the dataset name and the file name. Somebody with special priviledges will do the cleaning (not automatic yet). Consider the file as lost.
  - ◆ Check within your application. For example, it is possible that the file was not copied on the scratch disk associated to the CPU because it was full or the copy time-out occured before the file was completly copied.
3. If the file is corrupted at the source T1, it needs to be deleted from the Storage and DDM. Fill a Jira ticket to DDM Ops [🔗](#) providing the dataset name and the file name. Somebody with special priviledges will do the cleaning (not automatic yet). Consider the file as lost.

## Waiting Jobs Procedure

this is being discussed with the Production operation team

## How to know if the problems is task related or site related ?

- Check to see if any jobs are done (e.g. scouts) by entering the task number in the task field on the bottom of the panda browser. If the scout have gone through and there are a lot of failures at 1-2 sites, the sites are more suspect than the task.
- If the same task is found to be failing at several sites, probably it is a **task** related problem.
- If jobs from a task are failing at one single site and running OK at other clusters, probably it is site related problem.

### **NEW** Group production jobs

- Group production is now being handled by the new DPD Production Team (contact: atlas-phys-dpd-coordination@cernNOSPAMPLEASE.ch.)
- Monitoring [↗](#) for group production tasks
- **NEW** Twiki with usefull info for group production reporting
- Please report Group production jobs which did not finish within 1 day.
- Group production task should run less than 1 week.
- Problematic tasks are to be reported to ADCo Support [↗](#) Jira:
  - ◆ put "TASK" string and `task ID(s)` to the Subject
  - ◆ put task owner to CC
- Site issue is to be reported to GGUS.
- DaTrI requests [↗](#) for Group production datasets - please check at the beginning of your shift.
- Group contacts
- More info for Groups: DPDPProductionTeam

### Group production jobs experience/hints

- When job fails with "ATH\_FAILURE - Athena non-zero exit", please make sure that it has **inputs defined** first. If the task has no input defined, mention in in JIRA. In such case you don't have to check athena logs further and you don't have to put long excerpt of log into Jira ticket. In such case task has to be redefined by the task owner (whome you put into CC of the ticket).

### **NEW** Task in pending state for long time

- Look once at [http://bigpanda.cern.ch/tasks/?status=pending&statenotupdated=720&display\\_limit=100](http://bigpanda.cern.ch/tasks/?status=pending&statenotupdated=720&display_limit=100) [↗](#)
- If the task is failing, try to figure out the reason and report it according to Task Failing Section

# DDM

- DDMGlobalOverview
- Spot most problematic clouds in DDM dashboard: (begin with those in RED, then YELLOW and then BLUE):
  - ◆ Click on the Tier-1 name to get a breakdown for the sites. Chase the site(s) that is causing the low efficiency at the cloud by clicking on the error number (breakdown for errors).
    - ◇ Understand if the problem is site-related (DESTINATION error):

```
.FTS State [Failed] FTS Retries [3] Reason [DESTINATION error during PREPARATION phase: [CONNECTION] failed to contact on remote SRM...
```
    - ◇ or if the problem is outside of this site (SOURCE error):

```
.FTS State [Failed] FTS Retries [3] Reason [SOURCE error during PREPARATION phase: [CONNECTION] failed to contact on remote SRM
```
    - ◇ If the error message is : \* SOURCE error during TRANSFER\_PREPARATION phase: [REQUEST\_TIMEOUT] failed to prepare source file in 180 seconds
      1. click the number on the right. You will see the list of files with  
FAILED\_TRANSFER
      2. click some of the files to see the history of the file transfer
      3. if you see the error is persistent (many errors for more than 1 day), the problem should be reported explicitly mentioning the error is persistent.
      4. otherwise (if only a few errors, or errors within 1 day), no need to report
    - ◇ DDM is intrinsically linked as downtime on a site can cause collateral effects to all sites pulling or pushing data to it.
  - ◆ For problematic sites check Services column: DQ/Grid Status and report in case it is not OK (by this time only DQ is monitored)
- If you're new to the Team, please check DDMDashboardHowTo
- Please follow guidelines on what information to fill in DDM-specific ticket. This information is valid for both GGUS tickets to the site, and to DDM ops Jira tickets.

## DDM dashboard shows timeout/SRMV2STAGER errors

- When DDM dashboard shows timeout errors or SRMV2STAGER errors, you should wait and see if the error re-occurs and persists before you submit ticket to a site.
  - ◆ The aim of waiting is to make sure that the issue is still there, and to prevent us from sending false alarms to sites.
  - ◆ Duration of waiting period is shown on the error message, it can be from tens of minutes to day(s). In any case please file an eLog entry about the timeout/SRMV2STAGER issue, and mention this issue in your daily report.
- **NEW** Staging Statistics [↗](#) and Staging Errors [↗](#) logs. Use these 2 pages to get summary information about the staging failures. If the failure rate is too high, please consult with your fellow ADCoS Expert shifter whether a GGUS ticket should be filed.

## What to fill into GGUS ticket subject (short description)

- Site name or spacetoken name
- Short description of observed issue:
  - ◆ If the FTS error transfer contains 'locality is unavailable' : put `locality is unavailable` into GGUS ticket subject
  - ◇ Example ERROR MSG:

```
ERROR MSG: [FTS] FTS State [Failed] FTS Retries [1] Reason [SOURCE error duri
```

```

ource file [srm://ccsrm.in2p3.fr/pnfs/in2p3.fr/data/atlas/atlasdatadisk/step0
54L.physics_C.recon.ESD.closed/step09.20201041000054L.physics_C.recon.ESD.clo
ocality is UNAVAILABLE]

```

- ◆ If the FTS error contains 'gridftp\_copy\_wait: Connection timed out ' : put

```
gridftp_copy_wait: Connection timed out
```

- ◇ Example ERROR MSG:

```

[FTS] FTS State [Failed] FTS Retries [1] Reason [DESTINATION error during TRA
TION_ERROR] failed to contact on remote SRM [http://grid05.lal.in2p3.fr:8446
3 tries]

```

- ◆ Otherwise: try to describe problem within up to 4 words, please try to avoid phrases like "many transfer errors" only when better error description is provided in ERROR MSG.
  - ◇ If FTS error message states SOURCE error, put SITE\_X cannot export data (SITE\_X is name of the SOURCE site)
  - ◇ If FTS error message states DESTINATION error, put SITE\_X cannot receive data (SITE\_X is name of the DESTINATION site)
  - ◇ If FTS error message states TRANSFER error, put Transfer issues between SITE\_X and SITE\_Y

## Which Problems to report

- Tier-1s: No transfer reported at dashboard level during few hours (cross-check if the site is in Downtime before).
- Report DDM errors only if :
  - ◆ the source site is problematic (as reported by FTS). Probably site is down or it is lost file
  - ◆ T1/T0 <-> T1/T0
  - ◆ T1<->T2 within same cloud
  - ◆ T1/T0 <-> T2\_PRODDISK (affects production ) T2\_GROUPDISK (group datasets are not aggregated at final destination) (cross cloud or not)
  - ◆ Do **NOT** report issues with T2-T2 transfers.
  - ◆ **NEW** When ERROR code in DDM dashboard is [**DDM Site Services internal**] , just report the error following DDM-specific ticket
  - ◆ If FTS error means that it is a problem at source (pattern SOURCE in FTS error log)
- Dashboard: If there is no transfers shown in the DDM dashboard, please notify **immediately** the dashboard team: dashboard-support@cernNOSPAMPLEASE.ch. If you get no response after one hour and status is the same contact directly the ADC Expert.
  - ◆ DDM dashboard: Please report to atlas-adc-expert@cernNOSPAMPLEASE.ch and dashboard-support@cernNOSPAMPLEASE.ch. Please wait until the issue is resolved. Please monitor SAM SRMv2 tests in the meantime <http://tinyurl.com/ATLAS-SRM-last48> and report the most recent issues to sites. Do not report to atlas-dq2-support at cern.ch.
  - ◆ That could be a side problem of various things:
    - ◇ Dashboard agents are not working
    - ◇ Site-services are not working
    - ◇ No data is transferred to the sites
    - ◇ Everything fine, but no data at all -very rarely seen, as there is always traffic either in the Tier-0 or at the production dashboard-
- Please follow guidelines on what information to fill in DDM-specific ticket. This information is valid for both GGUS tickets to the site, and to DDM ops Jira tickets.
- When a ticket is solved by site, and the issue disappears from our monitoring tools (will not occur the new issue if the same kind within 1 hr from the ticket solution), consider the issue to be solved. When the same issue reoccurs after 1 hr after the old ticket was solved, please open a new ticket.



## Tier-0/Tier-1/Tier2 Data exportation

TBD

### What to do when a site has no FREE disk space in space tokens?

- In most cases when there is no free disk space in particular spacetoken this spacetoken should be blacklisted for writing automatically
- If you see error DESTINATION error [NO\_SPACE\_LEFT]

```
DESTINATION error during TRANSFER_PREPARATION phase: [NO_SPACE_LEFT] at Thu Jun 04 20:11:4
```

- ◆ this is not the site problem, but an atlas issue in usage of given resources, do not send ticket to ggus ticket to the site
  - ◇ check the free space at [http://bourricot.cern.ch/dq2/accounting/global\\_view/30/](http://bourricot.cern.ch/dq2/accounting/global_view/30/)
  - ◇ check whether the spacetoken is blacklisted for writing [http://bourricot.cern.ch/blacklisted\\_production.html](http://bourricot.cern.ch/blacklisted_production.html)
- ◆ If there is indeed no space left and spacetoken is not blacklisted submit an DDM ops [Jira](#) with the cloud support in CC so that they can take an action (increase the space, reduce the share, etc...)
- ◆ **EXCEPTIONS**
  - ◇ If you see error [NO\_SPACE\_LEFT] for DATAPE or MCTAPE, it is a site issue. Send email to [atlas-adc-expert\(at\)cern.ch](mailto:atlas-adc-expert(at)cern.ch).
- The following error messages means that the log area for the FTS server is full. In this case, submit a GGUS ticket to the site which hosts the FTS server

```
[FTS] FTS State [Failed] FTS Retries [1] Reason [AGENT error during TRANSFER_SERVICE phase cannot create archive repository: No space left on device]
```

or

```
Reason [AGENT error during TRANSFER_SERVICE phase: [INTERNAL_ERROR] error creating file for memmap /var/tmp/glite-url-copy-edguser/BNL-NDGF__2010-01-16-0659_m left on device]
```

or

```
[FTS] FTS State [Failed] FTS Retries [1] Reason [TRANSFER error during TRANSFER phase: [GR _ftp_client: the server responded with an error 500 Command failed. : write error: No space
```

The DDM endpoints \*\_LOCALGROUPDISK are not managed centrally or by the site admins. If a DDM endpoint is full, DDM automatically blacklists the site as destination. ADCoS shifter should submit a Jira ticket within ADCo Support [Jira](#) for information. The cloud squad should acknowledge and close the ticket. The squad is responsible to inform the local users.

The actions are similar for ATLASGROUPDISK (DDM endpoints mainly called PERF-\* or PHYS-\*). The Jira ticket should be submitted to ADCo Support [Jira](#) and assigned to **Group Production** and the DPD contact person should be put in CC (list in `DPDProductionTeam#Group_DPD_contact_persons`). The ticket should be acknowledged by the Group production responsible and closed.

For the other space tokens, an automatic cleaning algorithm is defined and running for all space tokens. If the site is full, it means that the cleaning procedure is not perfect. The cleaning monitoring can be found at

- [DATADISK/MCDISK/PRODDISK/SCRATCHDISK](#)
- More detailed information about the whole procedure (involving AMODs) is at [ADCOpsSiteExclusion](#).

## What to do when subscriptions are not processing?

- If Site Services are under suspect, follow: Cental Services procedure
- If the problems is related to data loss or catalogue inconsistencies:
  - ◆ Place as DDM ops [Jira bug](#)
    - ◇ Report the error message and associated link.
    - ◇ Report the dataset not transfered and when it was done.
    - ◇ Please follow guidelines on what information to fill in DDM-specific ticket. This information is valid for both GGUS tickets to the site, and to DDM ops Jira tickets.
- Hardware status
  - ◆ Find the cloud's machine in the [DistributedComputingMachines](#) page
  - ◆ Squeeze the lemon monitors [for info](#)

## Checking blacklisted sites in DDM

- When a site is heavily failing and bypassing certain error threshold (#errors/time), the site is removed from site services so no further transfers requests happens for the site. This is done by ADCoS Expert Shifters for T2 and T3 sites, and by AMOD for T0 and T1 sites. There could be several reasons for doing this: Long scheduled/unscheduled downtime, persistent storage problems, fts issues, etc.
  - ◆ Sites on downtime (GOCDB/AGIS) are excluded automatically.
  - ◆ **NEW** Sites failing SAAB [nagios tests](#) are excluded automatically. SAAB blacklistings can be monitored [here](#). Currently only put tests are active, so in case of the site problem only write/upload (w/u) part is blacklisted by SAAB. If the site has also a problem as a source or at deletions (r/f/d), shifter must treat them as regular transfer/deletion failure case. If site was blacklisted by SAAB, then the failing test issue should be followed up in GGUS ticket to the site. See more information in SAAB TWiki.
  - ◆ Sites which are not excluded automatically have to be excluded manually.
- Shifters should check at the end of the shift if the blacklisted sites are still in troubles or if they have solved the problems so the site can be set online again, this has to be mainly done following these steps:
  - ◆ Check which sites are blacklisted centrally in DDM: [http://bourricot.cern.ch/blacklisted\\_production.html](http://bourricot.cern.ch/blacklisted_production.html)
  - ◆ Check eLog and GGUS Team tickets and look for updates.
    - ◇ If site problems has been solved, the site should be brought online again.
- Instructions for ADCoS Expert shifters to exclude/re-include spacetoken from DDM: [ADCoSExpert#DDM\\_spacetoken\\_exclusion](#)

## Checking the deletion error rate per site

- Go to DDM dashboard [for info](#)
- For each cloud in the table
  - ◆ if a site has more than 1000 errors over the last 4 hours:
    - ◇ Check if the error rate is constant over these 4 hours
  - ◆ Report to ADCoS expert who will check if it is worthwhile to contact the site and fill GGUS ticket if necessary

- ◇ file a GGUS ticket to that site with CC to the corresponding cloud support
  - Ticket subject has to contain site name
    - e.g. Site CA-VICTORIA-WESTGRID-T2 has more than 18k deletion errors in last 4 hrs
  - Ticket details has to contain list of problematic spacetokens and examples of error extracted from the error table. The URL can be provided so that the site can check himself that, after correction, the error rate has decreased
- ◇ file an elog
  - reference created GGUS ticket in that eLog

# Panda queues

If a site or a queue at a site is in downtime or is heavily failing, the site should be set to test so that jobs are not directed anymore until the problem is solved. Currently, sites are permanently tested and manipulated by Hammercloud. The current status of queues can be found at [sites](#) page. Recent changes in general and also for given queue can be found at [incidents](#) page. Test jobs can be checked here [http://bigpanda.cern.ch/jobs/?jobtype=test&display\\_limit=100&prodsourcelabel=prod\\_test](http://bigpanda.cern.ch/jobs/?jobtype=test&display_limit=100&prodsourcelabel=prod_test).

---

# Central services

**Central services** (hosted at CERN) to be monitored [https://sls.cern.ch/sls/service.php?id=ADC\\_CS](https://sls.cern.ch/sls/service.php?id=ADC_CS) (need NICE login)

- See also [CompAtP1Shift#ATLAS\\_Central\\_and\\_Grid\\_Services](#) and [CompAtP1Shift#Database\\_Monitoring](#)
- If any kind of degradation is observed, log incident into ADC eLog. Some services need more actions to be taken (listed below):
- The escalation procedure for Site Services and Central Catalogues would be to send an email to the ADC expert at CERN, he will check the mail in daily basis and will decide if the problem need to be addressed to the IT piquet service at CERN.
  - ◆ [AMI ATLAS-AMI](#)
    - ◇ When degraded, contact the AMI developers via [atlas-tagcollector@lpscNOSPAMPLEASE.in2p3.fr](mailto:atlas-tagcollector@lpscNOSPAMPLEASE.in2p3.fr) with CC to [atlas-adc-expert@cernNOSPAMPLEASE.ch](mailto:atlas-adc-expert@cernNOSPAMPLEASE.ch)
  - ◆ [PanDA machines: Panda](#)
    - ◇ When degraded, contact the ADC experts, [atlas-dba@cernNOSPAMPLEASE.ch](mailto:atlas-dba@cernNOSPAMPLEASE.ch), and [atlas-adc-central-services@cernNOSPAMPLEASE.ch](mailto:atlas-adc-central-services@cernNOSPAMPLEASE.ch)
  - ◆ [Pilot Factories: PilotFactories](#)
    - ◇ When degraded, contact the ADC experts and [atlas-project-adc-operations-pilot-factory@cernNOSPAMPLEASE.ch](mailto:atlas-project-adc-operations-pilot-factory@cernNOSPAMPLEASE.ch)

# Frontier

The Frontier service provides access to the conditions data stored in the 3D databases which is streamed from CERN to several Tier1 sites. Conditions data accessed from Frontier is primarily used in user analysis jobs. Because conditions data changes relatively slowly a lot of requests are the same and so a series of squid caches have been set up to reduce the load on the Oracle databases. When a job requires conditions data it will first try and get it from a local site squid. If the required data is not in the squid, the squid should connect to the designated Frontier server which will connect to the Oracle database if it doesn't have the data cached already. The system is setup so that if a site squid or Frontier server fails then the request will try other Frontier / squid combinations in order to get their data. Problems with a site squid or Frontier server should therefore not cause jobs to fail, although this will cause additional load elsewhere. If this is allowed to build up the whole service could eventually fail.

Periodically (2-3 times per shift) check: <http://sls.cern.ch/sls/service.php?id=ATLAS-Frontier>

If this is not at 100% for any of the sites for more than an hour check:

- Is the site in downtime?
- Does the Frontier service appear down in MRTG? You can check MRTG here: <http://wlcg-squid-monitor.cern.ch/snmpstats/indexatlas.html> the Frontier machines all begin with Lpad

If the site is not in a downtime and it is down in both SLS and MRTG then submit an urgent GGUS ticket to the site and cc in [atlas-frontier-support@cernNOSPAMPLEASE.ch](mailto:atlas-frontier-support@cernNOSPAMPLEASE.ch). If in doubt email [atlas-frontier-support@cernNOSPAMPLEASE.ch](mailto:atlas-frontier-support@cernNOSPAMPLEASE.ch) and copy in the expert shifter.

Once per shift check:

[http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview#currentView=Frontier\\_Squid&highlight=false](http://dashb-atlas-ssb.cern.ch/dashboard/request.py/siteview#currentView=Frontier_Squid&highlight=false)

If the site is red, click on the link. This will take you to the MRTG monitoring page which will show you when the squid stopped working. Check if the site is in downtime and if it isn't and the squid has not been responding for more than 4 hours (no "had been up for" line) submit a less urgent GGUS ticket to the site and cc in [atlas-frontier-support@cernNOSPAMPLEASE.ch](mailto:atlas-frontier-support@cernNOSPAMPLEASE.ch). Exception sites are mentioned on Known Problem page.

# Miscelania

## Contacting the Cloud Support

- Carbon Copy always the Tier-1 expert list when submitting a GGUS
- Carbon Copy when an action is performed affecting the sites inside the cloud
  - ◆ atlas-adc-cloud-<CLOUD>@cern.ch
  - ◇ Where <CLOUD> stands for CA, CERN, DE, ES, FR, IT, ND, NL, RU, TW, UK, US

## **NEW** Contacting the ticket portal support

In case GGUS or Jira tracker pages are not available, you should try to

- clear the cache of your browser and try again
- ask other shifters or colleagues if they are affected by this problem or it is only you
- if you cannot find a cause of the problem on your side, send email to portal's support
  - ◆ GGUS
  - ◆ Jira
  - ◆ eLog

## Pilot Factories and methods

- Pilot factory monitoring: <http://apfmon.lancs.ac.uk/>
- Every factory has a contact person and his email address written on the top of the page.
- Put atlas-adc-expert@cernNOSPAMPLEASE.ch and cloud support to CC of every email related to pilot factory, since AMOD can help tuning the issue.

# Communication and Organization

- Before coming to your shift, make sure that you fulfill shift requirements!
- ADCoS Mailing list: atlas-project-adc-operations-shifts@cernNOSPAMPLEASE.ch
  - ◆ ADCoS Experts list: atlas-project-adc-operations-shifts-experts@cernNOSPAMPLEASE.ch
- Weekly meetings (CERN phone conference): INDICO INDICO [↗](#)
- Shifter Schedule is available in the OTP [↗](#).
  - ◆ Who is on shift? Click on **Report** in the top menu, select Task schedule report, fill in Task ID 529222, edit Dates From and To and confirm.
  - ◆ Booking shift? Click on **Home**, click on Book My Shifts, click on ID of the desired task, click on + button in front of the task ID and name, table with shift slots appears. Book your shift as a selection of red slots. Save.
- ADC Virtual Control Room



# Elog Management

## Choosing the right criticality in eLog

- 1) top priority: data export from cern, problem at the tier0, problems of the central services, central catalog etc.
- 2) very urgent: problems at the tier1s, like no acceptance data from the tier0, FTS down
- 3) urgent: problems that affect the cloud
- 4) less urgent: others

## Replying to eLog entry

- When replying to eLog entry modify the subject elog entry, for example :
  - ◆ If you updating information on the problem which has been already reported, put [update] in the subject
  - ◆ If the problem is solved, please put [SOLVED] in the subject
  - ◆ If the subject is saying that queues were set OFFLINE and you are setting them to TEST/ONLINE, reflect this in the subject: queues were set in TEST mode/ONLINE
  - ◆ If you can see that elog subject do not briefly report the problem (for example site name is missing in the subject) please modify elog subject (add a site name, if appropriate)

# Ticket Management

## Site naming convention - exceptions

- See GridSites

## General Rules

1. Check GGUS Atlas tickets (see How to find tickets section)
  - ◆ Now shifters can follow all the TEAM tickets on the GGUS interface
2. DO NOT open duplicate tickets.
  - ◆ If a ticket is already open about the SAME problem follow up on that ticket.
3. Open only TEAM tickets so that other shifters can find them and follow up
4. When you open a ticket: Write in the ADCoS eLog [eLog](#) the reason it was opened and put a link to the opened GGUS ticket
  - ◆ Tickets for the US can now be opened in GGUS so they can also be treated in the same way.
  - ◆ Some of T3's cannot be found in the list of available sites. Please use TPM option instead of direc route to A site in this case. It is very important to put site name in the description line ( for example, SITET3: transfers are failing because certificate has expired).
5. When you close a ticket: Write the solution ADCoS eLog [eLog](#)
6. Write everything that is in between in the ticket. The ticket is now the reference for what happens in between opening and closing.
  - ◆ In the ADCoS eLog [eLog](#) should only go the reason the ticket has been opened, the link to the ticket and the solution when the ticket is closed.
7. When updating the ticket do not change ticket status to "waiting to reply"; this status is reserved for sites. In this case when shifter check for tickets which needed to be updated, it is easy to see "waiting for reply" ticket.
8. Don't open tickets for sites in downtime.
9. Do not try re-open ALARM ticket. If ALARM ticket is solved, the problem re-appeared and you can't contact ADC expert, open new TEAM ticket.

## How to Submit GGUS Team-Tickets (direct routing to sites)

- Notice that when open the submit new ticket I/F a label appears on the top: Open TEAM ticket
- Clicking it you are in the I/F for our special tickets that gets routed directly to the site
  - ◆ Set the ticket priority:
    1. **top priority** : Problem at CERN Services (affecting exports to **every** site) should be marked as "uop priority". This includes LFC at CERN, FTS at CERN, SRM(CASTOR) at CERN
    2. **very urgent** : Problem at services at Tier-1s (affecting exports to the given Tier-1 and within the Tier-1 cloud) or services at calibration Tier-2s should be marked as "very urgent". This includes LFC at Tier-1s, FTS at Tier-1s, SRM at Tier-1s, SRM at Calibration Tier-2s.
    3. **urgent** : Any other problem should be marked as "urgent"
    4. **less urgent** : Informational entries should be marked as "less urgent"
  - ◆ Select type of the problem
  - ◆ Select MoU Area
  - ◆ Select site affected
  - ◆ Put cloud support in the CC field - atlas-adc-cloud-[CLOUD]@cern.ch [where [CLOUD] stands for CA, CERN, DE, ES, FR, IT, ND, NL, RU, TW, UK, US] (see

## #ContactingCloudSupport)

**Submit TEAM ticket**

[▶ Open GGUS ticket](#)
[▶ Open CMS ticket](#)

**User information**

Name: Michal.Svatos      E-Mail:

CC to: ?       VO: atlas

Notification mode ?  
 on every change  
 on solution

**Problem information**

Date / Time of Problem: 2014 / 10 / 01 / 14 : 08 UTC

\*Subject ?

Describe your problem providing the information listed here ?  
 characters left

Change priority:

Ticket category:

Type of problem:

MoU Area:

Notify SITE ?

Routing type:  TPM  SITE/ROC

- All people from ADCoS team with the correct certificate permissions in GGUS machinery can track and follow the tickets opened by anyone in our team.
  - ◆ Please notify GGUS support in case you find problems accessing to it.
- **NEW EXCEPTION** : For south-african sites (ZA-\*), the sites are not registered in GGUS. Submit the ticket to ROC NGL\_ZA (14 May 2012 and should be solved in the coming weeks)

## Overview of Jira trackers which ADCoS shifter might need

- ADCSUPPORT [↗](#) and ATLPHYSVVAL [↗](#) for tasks
- ATLDDMOPS [↗](#) for DDM issues
- ADCSITEEXC [↗](#) for situations which might require site exclusion
- several Jiras for monitoring tools

## How to find tickets in GGUS

- All GGUS TEAM tickets [↗](#)
- Open GGUS TEAM tickets [↗](#)
- GGUS Alarm Tickets [↗](#)

## Shift Reports and tickets

To help expert shifters to compile the weekly report and also help your fellow shifters to get oriented when their start their shifts:

- Write the ticket numbers of the the tickets you have opened and closed in the shifter report.
  - ◆ There is now a dedicated field for it.

## Ticket format

### GGUS [↗](#)

- Provide meaningful subject (non-ATLAS people should understand it), including the site name.
- Provide time information: When the failure(s) start to happen ?
- Extract of the error message. Shifter should understand the error before reporting and provide translation from "ATLAS" language to general "language".
  - ◆ When possible provide detailed info of the command failed (if the failure is reproducible)
- Provide link to the log(s) file(s) (panda/production dashboard for MC or DDM dashboard for data ditribution)
- Approximate number of failures (related to the problem reported):
  - ◆ Last 12h for Monte Carlo (default Panda monitoring view)
  - ◆ Last 4h/24h for Data distribution (possible views in DDM dashboard)
- Monte Carlo Specific:
  - ◆ Node(s) affected
    - ◇ Sometimes Worker Nodes act as black hole. High number of failures related to the same processing host could be an evidence.
  - ◆ Provide local batch system job ID
  - ◆ When providing link to Panda monitoring, please provide link to one particular failing job, "last12h" aggregation ling might be invalid at the time site tries to address the ticket
- DDM specific:
  - ◆ Cross check if the error is related to SOURCE or DESTINATION and open the ticket to the affected site
    - ◇ Pretty clear info provided in DDM dashboard inside "ERROR MSG" field
    - ◇ Provide also details of 1 failed transfer file (click on date in the Placement time column, detail will appear). Detail information contains transfer placement time, file name, GUID, tool id, number of attempts, SRC SURL, DEST SURL, transfer ID, channel info, and ERROR MSG
    - ◇ Provide time locked URL (file details contain red square saying LOCKED) of the page with this example of failed transfer in DDM dashboard.
      - Example URL:  
[http://dashb-atlas-ddm.cern.ch/ddm2/#d.error\\_code=166&d.src.cloud=%22CERN%22](http://dashb-atlas-ddm.cern.ch/ddm2/#d.error_code=166&d.src.cloud=%22CERN%22)
        - Do not provide URL only to the homepage of DDM dashboard  
<http://dashb-atlas-data.cern.ch/ddm2/> [↗](#)
  - ◆ Guidelines for GGUS subjects available in section What to fill into GGUS ticket subject (short description)

## Production [↗](#) or Validation [↗](#) Jira Bug Reporting

- Task:
  - ◆ Task ID

- ◆ Task name
- ◆ Task Progress (Done, ToBeDone, Running, Pending)
- ◆ Task efficiency
- ◆ Task details (release, trf\_version, DBrelease)
- Errors:
  - ◆ Error summary. The content of the job log file is accessible from its panda monitoring page (and the panda page is linked from the dashboard [↗](#)). Try the *Find and view log files* link. If it doesn't work, click on the job log file name (in the table above, file type *log*). At the bottom of the new page, you find the SURL(s) for the log file and you can download them directly in a shell.
  - ◆ Link to the Log files (Panda/ProdSys dashboard)
- Info flow:
  - ◆ Start the ticket body with the line mentioning the task owner's name (Task Owner: name). To put the correct name write @ symbol and start writing the name of the task owner. A drop-down menu will appear to help you find the right person and complete the name.
  - ◆ Add task owner to the watchers (see below).
  - ◆ Remember to eLog.

**NEW** Note! ATLAS Distributed Computing groups have moved away from Savannah and are using now the Jira Issue Tracking Service [↗](#) for bug reporting. Jira is quite easy and intuitive to use. In Jira to see the list of issues (tickets) click on "Issues" on the left side menu. On the "Issues" view one can select to look "All Issues" or only the ones belonging to certain category (Unresolved, Added recently, Resolved recently, etc.). To open a new ticket click on "Create Issue" button on upper menu and fill the form. In Jira there is no direct CC option, but it can be done by adding Watchers. To do so, find the label called "Watchers:" on the right side menu while inside a particular issue (ticket), and click on the colored circle with a number inside (the number indicates how many watchers that issue already has). You will be prompted to add a watcher. To do so simply start typing the first letters of the name of the person you want to CC (task owner for example). Then Jira will open a matching list from which you can select the desired name. Then Jira will add that person as a Watcher and send email notification every time the ticket is updated.

# Downtimes

- Check the AGIS downtime calendar [🔗](#) for downtimes.
- Check **all** ongoing entries for site in question. Do care also about downtimes marked as NO\_RISK\_FOR\_ATLAS.

# Monitoring tools

- Please report problems and feature requests for **monitoring tools** of ATLAS Distributed Computing (not for the actual athena, DDM, site, etc problems which the monitoring reveals!) to <https://its.cern.ch/jira/browse/ADCMONITOR>
- Problems with a particular monitoring tool should be reported as usual:
  - ◆ Dashboard: <https://its.cern.ch/jira/browse/DASHB>
  - ◆ Panda Monitoring: <https://its.cern.ch/jira/browse/ATLASPANDA>

# Daily SHIFT report

- Submit your daily shift summary report using the Interface located at: [Shift report elog form](#). This triggers an automated shift report that is sent to the ADCoS mailing list and also a disk copy stored in the elog: [Shift summaries elog](#)



# Trainee evaluation report

- If trainee shifter participated in the shift, send e-mail to : 1) ADCOS coordinators : atlas-adc-adcos-coordinators@cernNOSPAMPLEASE.ch, as well as 2) current ADCoS Expert, where name of current Expert shifter can be found from query (PDF) in the top of Checklist. The e-mail subject should be in the form : "Trainee evaluation of ShifterName (shift Number), date timezone", where example of the date and timezone is "10/10/2012 EU" and total number of trainee shifts taken so far should be reported, like "(shift 3)"
- Please, report the following (use copy/paste) :
  - ◆ Active presence: how much one is present and how much is proactive,
  - ◆ Monitoring tools understanding,
  - ◆ Errors understanding (at least as far as those explained in the twiki),
  - ◆ Ticket handling (learning how not to open duplicates and not to write a single error line... etc etc).
  - ◆ Evaluation grade: 0-3 range (1: new shifter, still learning; 2: quite experienced, but not yet ready for promotion; 3: ready to be promoted to senior shifter; 0: very quiet shift, not enough information to evaluate) [NEW](#)
- More detailed description of evaluation grades
  - ◆ 0: not enough information to evaluate - for example, no ticket submitted or updated during the shift, not enough interactions and discussions to evaluate the shifter experience.
  - ◆ 1: Shifter lacks understanding of ADCoS shifter duties and is in process of learning them.
  - ◆ 2: Shifter has basic understanding of ADCoS duties but makes mistakes while performing them. Examples could be submitting GGUS ticket to the wrong site or with incomplete information, submitting JIRA ticket to the wrong tracker or with incomplete information or making mistakes mentioned in Most Common Mistakes by Shifters section.
  - ◆ 3: Shifter knows how to perform all the duties mentioned in the Checklist. In case of transfer issues, (s)he is able to submit GGUS ticket to correct site following the rules for GGUS ticket content. In case of task issues, (s)he is able to submit JIRA ticket to correct tracker following the rules for JIRA ticket content

# ADC Virtual Control Room

- Jabber VCR info
  - ◆ Room: **adcvcr**
  - ◆ Server: **conference.chat.uio.no**
  - ◆ Password: will provided on demand by  
atlas-adc-adcos-coordinators@cernNOSPAMPLEASE.ch
- More detailed information on dedicated TWiki: [ADCVirtualControlRoom](#)

## Troubleshooting

- Currently, it looks like Gmail jabber accounts face strange behaviour when you join the chat: you'll see very ancient chat log, but you will not see the most recent one when you log in. In the meantime, please try to use other jabber account than Gmail, e.g. try out your jabber account at CERN ([jabber.cern.ch](#)).
- Jabber server jabbim.\* is a privately-held jabber server. If it stops working there is nothing ADC experts can do about it. In that case please check [jabbim.\\* servers monitoring](#) and wait for the servers to get back. If getting jabbim back takes too long, please use your CERN jabber account.
- If you are disconnected with a "Conflict" error, please reconnect again with a different nickname (handle).
- If you are disconnected during the jabbim.com server downtime, please use your CERN jabber account to reconnect.

# GGUS ATLAS TEAM membership

- If for some reason your grid certificate is not yet in /atlas/team VOMS group, ask ADCoS Coordinators to add it. Particularly don't forget to do that when you get a completely new certificate. GGUS receives this list in daily basis and updates the membership accordingly.

# Useful links

- ATLAS Site Status Board: Shifter view [with some introduction in SSB shifter instructions.](#)
- ADC Monitoring homepage: <http://adc-monitoring.cern.ch/>
- DDM Dashboard 2.0: <http://dashb-atlas-data.cern.ch/ddm2/>
  - ◆ ~~DDM Dashboard~~: <http://dashb-atlas-data.cern.ch/dashboard/request.py/site>
- T0 dashboard: [http://dashb-atlas-data.cern.ch/ddm2/#activity=\(1\)](http://dashb-atlas-data.cern.ch/ddm2/#activity=(1))
- Dashboard (DDM and ProdSys) Jira: <https://its.cern.ch/jira/browse/DASHB>
- ADC eLog: <https://atlas-logbook.cern.ch/elog/ATLAS+Computer+Operations+Logbook/>
- ADCoS Twiki: <https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/ADCoS>
- PanDA monitoring: <http://bigpanda.cern.ch/>
- PanDA Jira: <https://its.cern.ch/jira/browse/ATLASPANDA>
- PanDA Shift Guide: <https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/PandaShiftGuide>
- Task monitoring: <http://dashb-atlas-task-prod.cern.ch/templates/task-prod/>
- Monitoring tools Jira - general requests: <https://its.cern.ch/jira/browse/ADCMONITOR>
- AGIS Calendar: <http://atlas-agis.cern.ch/agis/calendar/?print>
  - ◆ GOCDB: <https://goc.egi.eu/portal/>
- FTS Monitoring for ADC@P1 shifters:  
[https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/ADCPoint1Shift#FTS\\_monitoring](https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/ADCPoint1Shift#FTS_monitoring)
- SAM test results <http://dashb-atlas-sum.cern.ch/dashboard/request.py/latestresultssmry-sum>
- ATLAS storage space monitor: [https://sls.cern.ch/sls/service.php?id=storage\\_space](https://sls.cern.ch/sls/service.php?id=storage_space)
- DDM Twiki: <https://twiki.cern.ch/twiki/bin/viewauth/AtlasComputing/DistributedDataManagement>
- FT monitor plots: [http://atladcops.cern.ch:8000/drmon/ftmon\\_tier2s.html](http://atladcops.cern.ch:8000/drmon/ftmon_tier2s.html)
- ATLAS SW installation: [https://atlas-install.roma1.infn.it/atlas\\_install/](https://atlas-install.roma1.infn.it/atlas_install/)
- FCR: <https://lcg-fcr.cern.ch:8443/fcr/fcr.cgi>
- TiersOfAtlas: <http://atlas-agis-api.cern.ch/request/toacache/TiersOfATLASCACHE.py>
- Ganga Robot Results: <http://gangarobot.cern.ch/>
- Gstat site-info: <http://goc.grid.sinica.edu.tw/gstat/>
- ATLAS off-line SVN repository: <https://svnweb.cern.ch/cern/wsvn/atlasoff>
- DDM Operations Jira: <https://its.cern.ch/jira/browse/ATLDDMOPS>
- ATLAS Validation Jira: <https://its.cern.ch/jira/browse/ATLPHYSVAL>
- ADC operations support Jira: <https://its.cern.ch/jira/browse/ADCSUPPORT>
- ADC Site Status Jira: <https://its.cern.ch/jira/browse/ADCSITEEXC>
- GGUS portal: <https://ggus.eu/>
- CERN IT Services: <https://sls.cern.ch/sls/index.php?view=services>
- GRIDVIEW: <https://gridview.cern.ch/GRIDVIEW/>
- ATLAS Acronym Glossary:  
<http://www.hep.man.ac.uk/atlas/ReadOut/.grsthist:AARG.html:448D57A5:0079B:117CE:=2FC=3DUK=2FO>
- ATLAS Distributed Analysis Shifters Team: AtlasDAST
- Comp@P1 shifters: CompAtP1Shift
- Shift booking in OTP [with some introduction in SSB shifter instructions.](#)
- SquadHowTo

# Tutorials

- Shifters Jamboree: ADC Tools Tutorial and ADC Shifts Tutorial, during the ATLAS Distributed Computing Facilities Jamboree and Shifters Jamboree (3-5 December 2014)  
<https://indico.cern.ch/event/276502/>
- Tutorial session for ATLAS Distributed Computing shifts: ADCoS, and Comp@P1 during ATLAS S&C workshop (14 July 2010) Agenda (slides & video)
- Tutorial session for ATLAS Distributed Computing shifts: ADCoS, Tier-0 and Point-1 during ATLAS week (24 February 2010) Agenda (slides & video)
- Tutorial in ATLAS Software & Computing Workshop (02 September 2009) Agenda
- World-Wide EVO tutorial (23th of July 2009): Agenda
- ADC Shifters Jamboree, Januray 2009: <http://indico.cern.ch/conferenceDisplay.py?confId=45394>
- World-Wide EVO Monday 30 June 2008: <http://indico.cern.ch/conferenceDisplay.py?confId=35652>
- Tutorial held at CERN during the Software Week (25th February 2008):  
<http://indico.cern.ch/conferenceDisplay.py?confId=22132>
- Link to the ADCoS workshop held at CERN during 21st-22nd. january 2008 Agenda
- Operation Tutorial: EGEE Production (slides + video), Oct 25, 2007
- Spacetokens availability (free space) page:  
[http://bourricot.cern.ch/dq2/accounting/datadisk\\_view/30/](http://bourricot.cern.ch/dq2/accounting/datadisk_view/30/) ,  
[http://bourricot.cern.ch/dq2/accounting/scratchdisk\\_view/30/](http://bourricot.cern.ch/dq2/accounting/scratchdisk_view/30/)

# Shift Credits

- Before coming to your shift, make sure that you fulfill shift requirements!
- ADCoS is a Class 2 shift. All ADCoS shift have the same weight within ATLAS OTP.
- **NEW** Each shifter is required to take at least 6 shifts every 4 months!!!
- We have 3 flavours of shifters: ADCoS Expert shifters, ADCoS Senior shifters, ADCoS Trainee shifters.
  - ◆ **Senior** shifter: 8 hours of shift, 2-days blocks (Mon+Tue, Wed+Thu) and Friday credited with 78% (scaled from 100%), 2-days block (Sat+Sun) bonus credit 155% (scaled from 100%).
    - ◇ No upper limit on number of shifts. Please take at least few shifts a month. ADCoS training should be repeated if no shifts were taken within a year.
  - ◆ **Trainee** shifter: 8 hours of shift, shifts slots available Mon-Sat, 0% shift credit (scaled from 100%), no Sunday shift.
    - ◇ Please book your Trainee shift slot only if it is "red" on the OTP calendar. Please do not overallocate shift slots.
    - ◇ Trainee period takes 10 shifts, however, the final number of Trainee shifts strictly depends on Trainee shifter's performance, it can be significantly lower or higher than 10.
    - ◇ There is a time limit of 3 months for trainee shifters to finish training. If trainee shifter did not take shifts within last 1 month this shifter would be automatically excluded from the list. Each shift will be evaluated by Senior shifter.
    - ◇ After promotion to Senior it is required to take first Senior shift as soon as possible (within a month).
  - ◆ **Expert** shifter: 9 hours of shift, 7-days shift Wed-Tue, credited with 100%, no weekend bonus.
- We provide 24/7 operations shift in three timezones (defined in CERN time)
  - ◆ 00:00 - 08:00 ASIA/PACIFIC (AP) - Shift Captain: Hiroshi Sakamoto
  - ◆ 08:00 - 16:00 EUROPE (EU) - Shift Captain: Alexei Sedov
  - ◆ 16:00 - 24:00 AMERICAS (US) - Shift Captain: Armen Vartapetian
- Shifts are booked on first-come-first-served basis in ATLAS OTP [↗](#).
- ADCoS tasks in OTP:
  - ◆ 529221 - ADCoS Expert shifts
  - ◆ 529222 - ADCoS Senior shifts
  - ◆ 529223 - ADCoS Trainee shifts
  - ◆ 86 - ADCoS Coordination Shifts
- Generally, more information about ATLAS shifts available at ATLAS OtpShiftClasses TWiki page.
- In case of questions please contact ADCoS coordinators ([atlas-adc-adcos-coordinators@cernNOSPAMPLEASE.ch](mailto:atlas-adc-adcos-coordinators@cern.ch)).

## Expert Duties

- See ADCoSExpert page

# TEAM MEMBERS

- See ADCoSTeam page

---

## Major updates:

-- XavierEspinal - 30 Jul 2008 -- JaroslavaSchovancova - 2010-2011 -- MichalSvatos - 2014

%RESPONSIBLE% AlexeySedov

%REVIEW% **Never reviewed**

-- MichalSvatos - 25 Jun 2014

---

This topic: Main > ADCoSWIP

Topic revision: r11 - 2015-01-13 - MichalSvatos



Copyright &© 2008-2019 by the contributing authors. All material on this collaboration platform is the property of the contributing authors. Ideas, requests, problems regarding TWiki? Send feedback