

Tier2D Ops Issues/Communication

U.S. Facilities Workshop/SMU
October 12, 2011
Mark Sosebee

Outline

- Description of production shifts
- Tools used by shifters
- Contact options
- ATLAS Tier2D
- Example of how (not?) to handle problems
- Can we improve the situation?
- Summary

Shifter Perspective (I)

- Production shifts have evolved considerably over time
- In more US-centric days we called them “PanDA shifts”
- With widespread adoption of PanDA in the other clouds it became necessary to standardize the shifts
- Enter “ADCoS” (Atlas Distributed Computing Operations Shifts)

Shifter Perspective (II)

- ADCoS shifts are 8 hours x 3 time zones, including Sundays
- The full team includes:
 - **AMOD**
 - ADCoS “expert” shifter (US: Wensheng Deng, Yuri Smirnov, Douglas Smith)
 - ADCoS “senior” shifter (+ maybe a trainee)
 - **Point1 shifter**

Very Useful Twiki...

(Jarka Schovancova)

<https://twiki.cern.ch/twiki/bin/viewauth/Atlas/ADCoS>

File Edit View History Bookmarks Tools Help

ADCoS < Atlas < Twiki

https://twiki.cern.ch/twiki/bin/viewauth/Atlas/ADCoS

Google

Jump Search Atlas All webs

Atlas

Atlas Web > AtlasComputing > AtlasDistributedComputing > ADCoS (28-Sep-2011, JaroslavaSchovancova)

ADCoS FOLLOW ME ON [twitter](#)

- ↓ [Requirements](#)
- ↓ [CHECKLIST](#)
- ↓ [MC production](#)
 - ↓ [General guide-line \(Fast Troubleshooting\)](#)
 - ↓ [Panda monitor I/F](#)
 - ↓ [Job-states definitions in Panda](#)
 - ↓ [What to do when a task is failing?](#)
 - ↓ [How to know if the problems is task related or site related ?](#)
 - ↓ [Tips for ticketing the right site](#)
 - ↓ [Missing files](#)
 - ↓ [What to do when I get checksum errors during dq2-get/lcg-cp?](#)
 - ↓ [What to do when a site is not getting jobs](#)
 - ↓ [What to do when there is an LFC outage](#)
 - ↓ [Waiting Jobs Procedure](#)
 - ↓ [ATLAS Software Releases problem](#)
 - ↓ [Reprocessing jobs](#)
 - ↓ [Group production jobs](#)
 - ↓ [Group production data transfer requests in DaTri](#)
 - ↓ [Group production jobs experience/hints](#)
 - ↓ [Missing input files on T2_PRODDISK](#)
 - ↓ [Symptoms of input file missing at T2_PRODDISK](#)
 - ↓ [Checking the deletion backlog](#)
- ↓ [DDM](#)
 - ↓ [DDM dashboard shows timeout/SRMV2STAGER errors](#)
 - ↓ [What to fill into GGUS ticket subject \(short description\)](#)
 - ↓ [Which Problems to report](#)
 - ↓ [Tier-0/Tier-1/Tier2 Data exportation](#)
 - ↓ [What to do when a site is failing ?](#)
 - ↓ [What to do when a site has no FREE disk space in space tokens?](#)
 - ↓ [What to do when subscriptions are not processing?](#)
 - ↓ [Checking blacklisted sites in DDM](#)
 - ↓ [Checking transfers to CERN-PROD_PHYS-GENER](#)
- ↓ [Panda queues](#)
 - ↓ [Controlling Panda Queues](#)
 - ↓ [Cloud Control](#)
 - ↓ [HOWTO change queue status?](#)
 - ↓ [Cloud and queue status](#)
 - ↓ [How to track site status?](#)
 - ↓ [How to send test inhs \(procedure\)](#)

ATLAS Homepage
ATLAS Collaboration

ATLAS TWiki
Public Results
Physics
Detectors
Trigger
Computing
Data Preparation
Documentation Help
Help
Glossary

Create a LeftBar for this page

TWiki Search
MarkSosebee
SSO logout

Edit Attach PDF

Problem Reporting/Communicating Issues

- Three systems:
 - ggus (<https://ggus.eu/pages/home.php>)
 - eLog (<https://atlas-logbook.cern.ch/elog/ATLAS+Computer+Operations+Logbook>)
 - Savannah (<https://savannah.cern.ch/>)
 - Note: ggus tickets for U.S. sites should trigger a corresponding RT ticket
- Everything should be eLog'd – choice of ggus vs. Savannah depends on the problem
- Chat session always on-line for shift team (ADC VCR)

Shifter Tools

(might be useful for site admins?)

- Various panda monitor views:
 - Production summary
 - Error summary
 - Pilot information (see [clouds views](#))
- DDM dashboard (v2.0)
- Site downtimes (GOCDB for EGEE & NDGF, OIM for US sites)
- DDM deletions information
- SSB (site status board)
- Much more [here](#) (ADCoS TWiki)

Contacts

- Shifters: atlas-project-adc-operations-shifts@cern.ch
- ADCoS expert: atlas-adc-expert@cern.ch
- atlas-support-cloud-US@cern.ch (or substitute any other cloud name: CA, DE, ES, FR, IT, ND, NL, TW, UK)
- US-specific mailing list: pandashift@hepmail.uta.edu

How Are We Doing?

- Shifts have become more involved over time:
 - Good news: we're tracking more information
 - Not-so-good news: a lot to keep up with
 - Issues can slip through?
- More shifters on the team:
 - Good: distributes the work load
 - (Sometimes) bad: new shifters need time to come up to speed
- Overall, issues are well-covered and reported

ATLAS Tier2D*

- T2D's: “directly connected Tier2” (F.Barreiro, S.Campana, I.Ueda...)
- A Tier2 with direct connections to all Tier1's, Tier2D's and CERN
- Tier2D selection criteria:
 - Robustness
 - Network bandwidth and performance
- Goal: commission all Tier2's as Tier2D's

* Alexei Klimentov, NEC 2011, Varna

How Do We Handle Tier2D Issues?

- At present, possibly room for improvement...
- Multiple examples of slow file transfers from U.S.
=> other clouds – jobs fail due to transfer timeouts
- Is cloud support sufficient to manage these kinds of problems?
- Rob suggested a “Hiro-like” contact in the other clouds – does cloud support suffice?

A Current Issue...

- On 10/9 shifter created:
<https://savannah.cern.ch/bugs/index.php?87589>
- “MWT2_UC has many jobs transfer timeouts to TRIUMF”
- Discussion thread followed in eLog (so far no updates to the Savannah ticket, nor has a ggus ticket been created)
- See:
<https://atlaslogbook.cern.ch/elog/ATLAS+Computer+Operations+Logbook/30170>



LCG Savannah

Login Status:

Not Logged In

Login

New User

This Page

Clean Reload

Printer Version

Related Recipes:

- Markup Reminder
- Why log in?

Search

in **Projects**

Hosted Projects

- Register New Project
- Full List
- Contributors Wanted
- Statistics

Site Help

- User Docs: Cookbook
- User Docs: In Depth Guide
- Get Support
- Contact Us

Links

GNU/Savane



Atlas Distributed Computing Operation Support - Bugs: bug #87589, MWT2_UC has many jobs transfer...

- Group**
- Main
- Download
- Docs
- Support
- Mailing Lists
- Source Code
- Bugs**
- Tasks
- News

You are not allowed to post comments on this tracker with your current authentication level.

bug #87589: MWT2_UC has many jobs transfer timeout to TRIUMF

Submitted by: [Wen Guan <wguan>](#)

Submitted on: 2011-10-10 01:51

<u>Category:</u>	None	<u>Severity:</u>	3 - Normal
<u>Priority:</u>	5 - Normal	<u>Item Group:</u>	None
<u>Status:</u>	None	<u>Privacy:</u>	Public
<u>Assigned to:</u>	None	<u>Open/Closed:</u>	Open

(-) Discussion

2011-10-10 01:51, **original submission:**

Hi,

At MWT2 there are many jobs failed. The failed message is "taskBuffer: transfer timeout". The destination SE is TRIUMF.

(1)<http://panda.cern.ch/server/pandamo...>

(2)

http://panda.cern.ch/server/pandamo...=*&site=MWT2_UC&type=production&jobStatus=failed&hours=12

Wen Guan
<wguan>

(-) Attached Files

No files currently attached

List | New | Duplicate | Reply | Find | Move to CCRC'08 Logbook | Logout | Help

- Mon, Oct 10, 2011, 01:51 +0200, Wen Guan, wguan.icedew@gmail.com, Less Urgent, US, T2, No, MWT2_UC has many jobs transfer timeout to TRIUMF
- Mon, Oct 10, 2011, 04:00 +0200, Rob Gardner, rwg@hep.uchicago.edu, Less Urgent, US, T1, No, MWT2_UC has many jobs transfer timeout to TRIUMF
- Mon, Oct 10, 2011, 10:23 +0200, Rodney Walker, rodneywalker@lmu.de, Less Urgent, US, T1, No, MWT2_UC has many jobs transfer timeout to TRIUMF
- Mon, Oct 10, 2011, 13:03 +0200, Alastair Dewhurst, alastair.dewhurst@cern.ch, Less Urgent, US, T1, No, MWT2_UC has many jobs transfer timeout to TRIUMF
- **Mon, Oct 10, 2011, 13:18 +0200, Rodney Walker, rodney.walker@lmu.de, Less Urgent, US, T1, No, MWT2_UC has many jobs transfer timeout to TRIUMF**

Message ID: 30170 Entry time: Mon, Oct 10, 2011, 13:18 +0200 In reply to: 30169

Author:	Rodney Walker
Email:	rodney.walker@lmu.de
Severity:	Less Urgent
Cloud:	US
Level:	T1
Affect Users:	No
Subject:	MWT2_UC has many jobs transfer timeout to TRIUMF

Hi,
Maybe just invite them to upgrade the fts monitor and check the MWT2 transfer rate(without ticket). You are right, the rest is rather for ddm ops. For example, what transfer types are competing for the T2D-T1 channels.

Cheers,
Rod.

```
> Hi Rod
>
> I am not sure what should be in the ticket. Currently:
> Between [T2DS] and TRIUMF-LCG2
> 20 concurrent files on 15 streams each
>
> That seems a reasonable setting so I wouldn't have thought Triumph would want to increase that. As you say the quick fixes is to stop the production jobs at
> MWT2 going to CA. Can you please advise what if anything I should follow up with Triumph?
>
>
> Alastair
>
>
>
>> Hi,
>> A ticket to TRIUMF is the best bet. I have the feeling the T2D channels have become overloaded with user transfers, and maybe there is some current
>> consolidation too. This hurts the multicloud model - there is no separate shares for user/prod transfers. The quick fix is to remove CA from MWT2, and the
>> longer fix to do some HITS merging at T2.
>>
>> The TRIUMF fts monitor
>> http://gridinfo.triumf.ca/ftsmonitor/
>> seems to be an older, and less useful, version, so not much use to figure out what is happening.
>>
>> Cheers,
>> Rod.
>>
>>> We'll be interested to see how this one is resolved, since it happened last week as well. NOTE - this is not at Tier 2 issue, but a T1-T2 issue (but the
>>> actually the T1's FTS is probably in place, so its really a T1 issue). Who'll be following up?
```

Options

- Instruct shifters about identifying these errors as inter-cloud transfer problems
 - Appropriate ticket, cc list, etc.
- Escalate this issue to ADC ops
- Try involving cloud support
- Other possibilities?

Summary

- ADCoS shifts have mostly been a success
 - Coverage 24x7 across all time zones
 - Integration across all clouds (PanDA)
- With the generalization of some Tier2's => Tier2D's job file transfer problems are an intermittent problem
- At present we need a strategy to address this problem more systematically
- Even if shifters more precisely report this type of problem still need to understand/fix root causes

Go Rangers

