



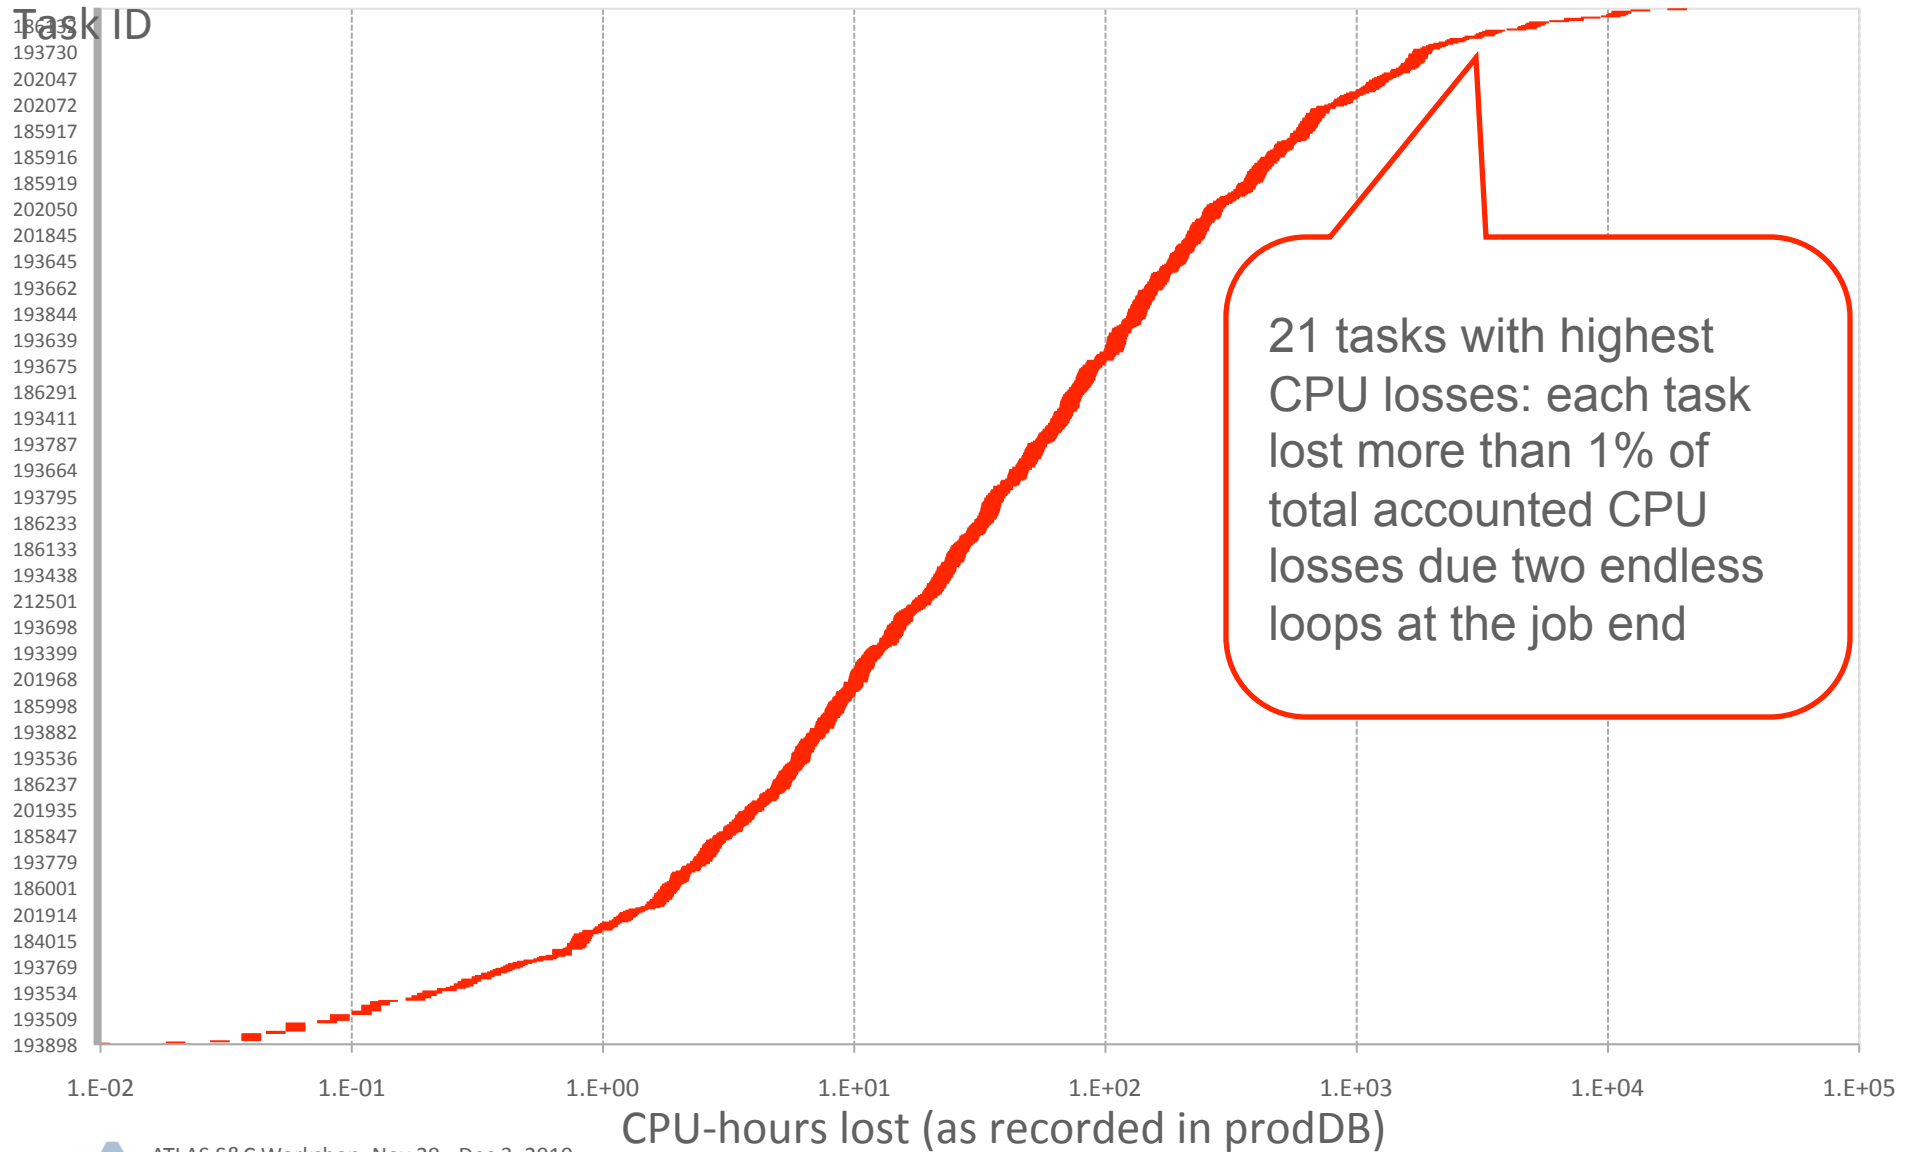
Distributed Computing: Error Analysis

Alexandre Vaniachine
ATLAS Software & Computing Workshop
November 29 – December 3, 2010
CERN, Geneva, Switzerland

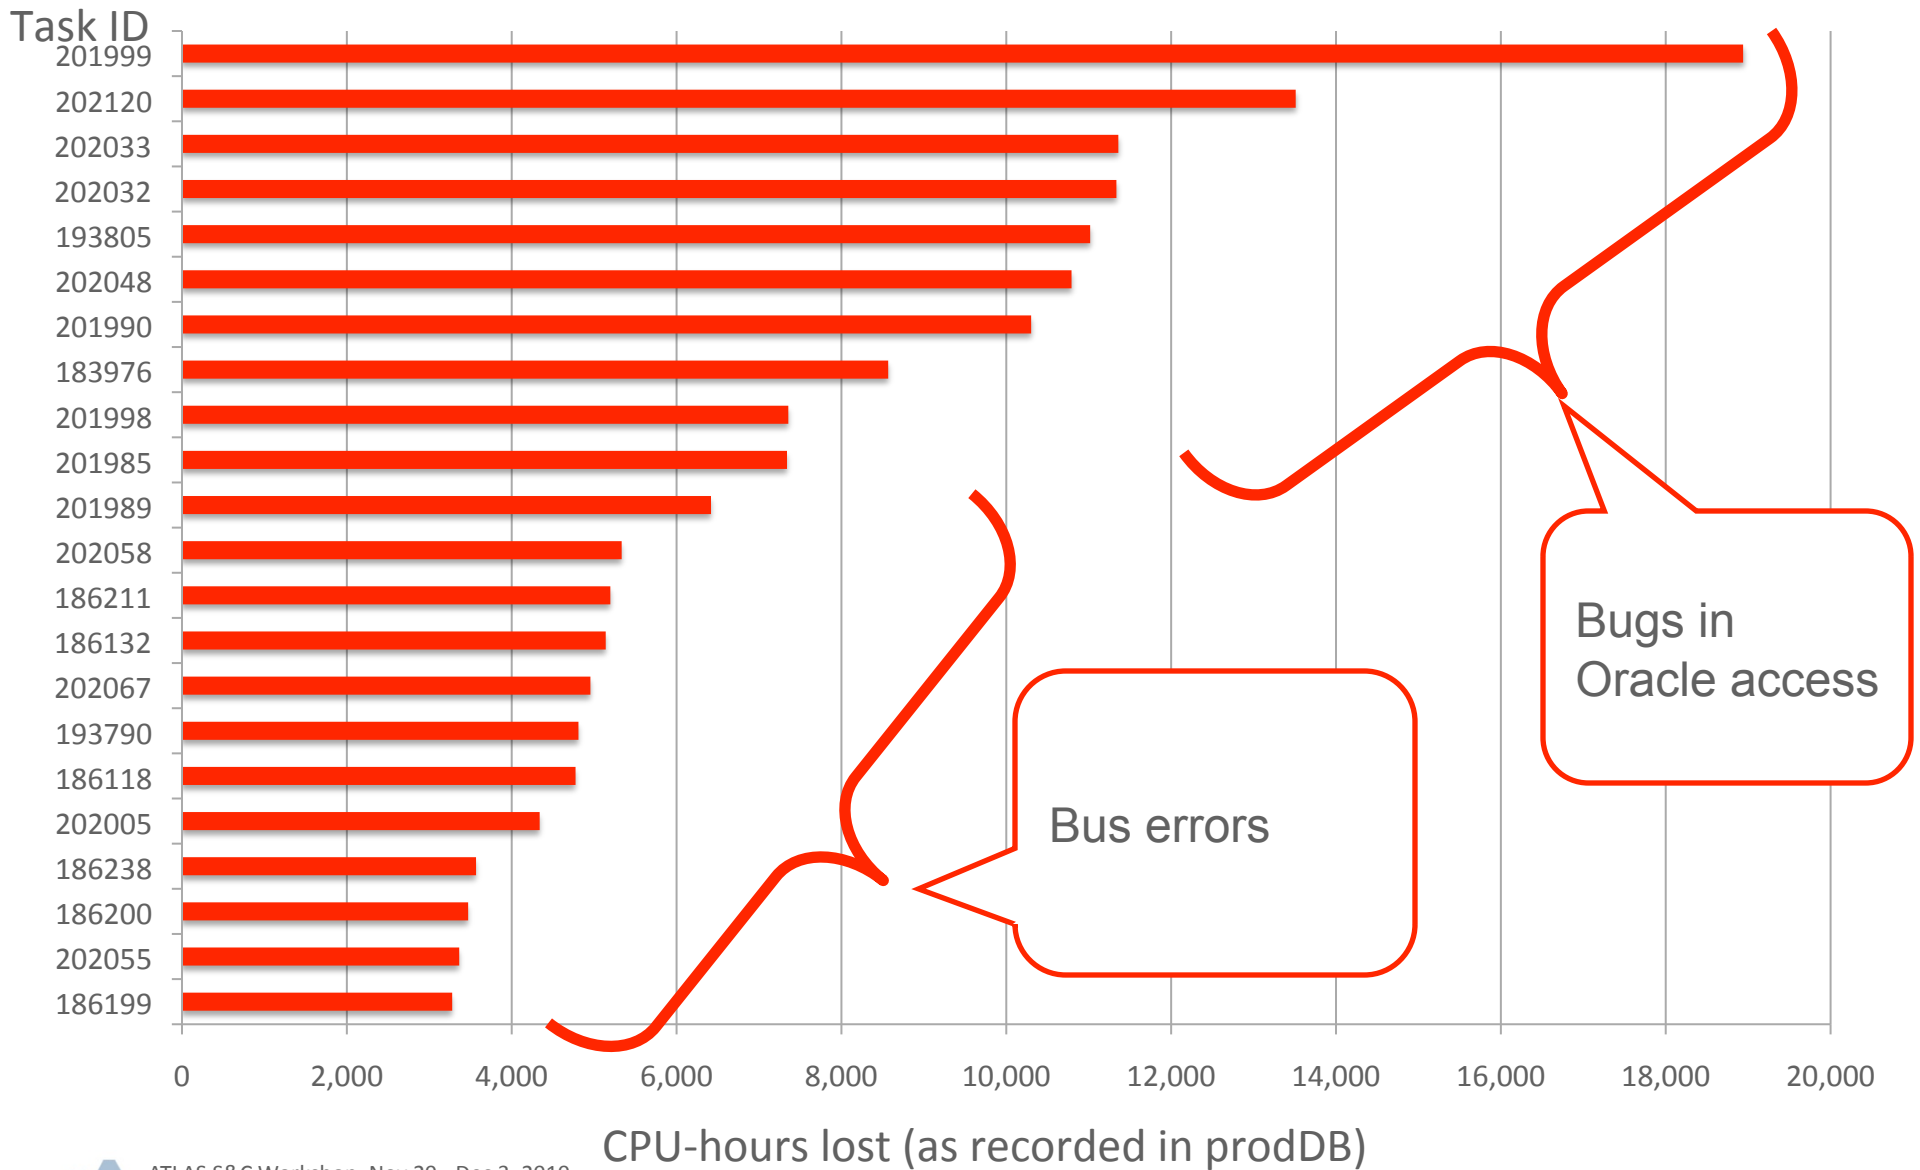
Preliminary Results

- Results are preliminary
 - Reprocessing of more than 1 PB of data on the Grid finished just yesterday
 - There are considerable “dark” CPU losses, which are unaccounted for yet
 - Exacerbated by the PanDA server failure due to Oracle server problem on November 19
- As a result, the analysis is limited to select unaffected cases with highest accounted CPU loss
 - Were we must improve the software and grid efficiency

Petascale Data Processing on the Grid: the S-curve



Two Reasons for Highest Accounted CPU Losses



Example from the Very Last r1774-Tag Jobs

Jobs:									
PandaID, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	Cloud/Site, Type	Priority	
1153997601 strandbe@umich.edu	data10_7TeV.00166658.debugrec_hltrej.recon.g1_r1774_tid210104_000005.job 7	finished	2010-11-29 02:57	0:03:42	9:15:19	11-29 12:16	US/BNL_ATLAS_2, production	954	
In: data10_7TeV.00166658.debugrec_hltrej.merge.RAW.g1 Out: data10_7TeV.00166658.debugrec_hltrej.recon.HIST.g1_r1774_tid210104_00									
1153427086 strandbe@umich.edu	data10_7TeV.00166658.debugrec_hltrej.recon.g1_r1774_tid210104_000005.job 6	cancelled	2010-11-28 05:56	0:02:54	20:46:32	11-29 02:45	US/BNL_ATLAS_1, production	945	
Error details: taskBuffer: killed by /C=UK/O=eScience/OU=Glasgow/L=CompServ/CN=graeme.stewart/CN=proxy									
In: data10_7TeV.00166658.debugrec_hltrej.merge.RAW.g1 Out: data10_7TeV.00166658.debugrec_hltrej.recon.HIST.g1_r1774_tid210104_00									
1152866580 strandbe@umich.edu	data10_7TeV.00166658.debugrec_hltrej.recon.g1_r1774_tid210104_000005.job 5	cancelled	2010-11-27 03:51	0:04:47	21:30:26	11-28 01:27	US/BNL_ATLAS_1, production	936	
Error details: taskBuffer: killed by /C=UK/O=eScience/OU=Glasgow/L=CompServ/CN=graeme.stewart/CN=proxy									
In: data10_7TeV.00166658.debugrec_hltrej.merge.RAW.g1 Out: data10_7TeV.00166658.debugrec_hltrej.recon.HIST.g1_r1774_tid210104_00									
1152559885 strandbe@umich.edu	data10_7TeV.00166658.debugrec_hltrej.recon.g1_r1774_tid210104_000005.job 4	cancelled	2010-11-26 16:31	0:05:34	10:54:54	11-27 03:31	UK/RAL-LCG2, production	927	
Error details: taskBuffer: killed by /C=UK/O=eScience/OU=Glasgow/L=CompServ/CN=graeme.stewart/CN=proxy									
In: data10_7TeV.00166658.debugrec_hltrej.merge.RAW.g1 Out: data10_7TeV.00166658.debugrec_hltrej.recon.HIST.g1_r1774_tid210104_00									
1151857153 strandbe@umich.edu	data10_7TeV.00166658.debugrec_hltrej.recon.g1_r1774_tid210104_000005.job 3	cancelled	2010-11-25 15:27	0:06:43	1 day, 0:42:26	11-26 16:17	UK/RAL-LCG2, production	918	
Error details: taskBuffer: killed by /C=UK/O=eScience/OU=Glasgow/L=CompServ/CN=graeme.stewart/CN=proxy									
In: data10_7TeV.00166658.debugrec_hltrej.merge.RAW.g1 Out: data10_7TeV.00166658.debugrec_hltrej.recon.HIST.g1_r1774_tid210104_00									
1150574877 strandbe@umich.edu	data10_7TeV.00166658.debugrec_hltrej.recon.g1_r1774_tid210104_000005.job 2	cancelled	2010-11-24 02:54	0:07:49	1 day, 11:53:26	11-25 14:55	UK/RAL-LCG2, production	909	
Error details: taskBuffer: killed by /C=UK/O=eScience/OU=Glasgow/L=CompServ/CN=graeme.stewart/CN=proxy									
In: data10_7TeV.00166658.debugrec_hltrej.merge.RAW.g1 Out: data10_7TeV.00166658.debugrec_hltrej.recon.HIST.g1_r1774_tid210104_00									
1150342655 strandbe@umich.edu	data10_7TeV.00166658.debugrec_hltrej.recon.g1_r1774_tid210104_000005.job 1	failed	2010-11-23 17:41	0:15:08	8:27:56	11-24 02:24	UK/RAL-LCG2, production	900	
Error details: pilot: Job killed by signal 15: Signal handler has set job result to FAILED, ec = 1201 (no time to send log)									
In: data10_7TeV.00166658.debugrec_hltrej.merge.RAW.g1 Out: data10_7TeV.00166658.debugrec_hltrej.recon.HIST.g1_r1774_tid210104_00									

- Several known bugs in our software for Oracle access led to the endless loop to get the direct Oracle connection at the ESDtoAOD stage

Bugs in Oracle Access

- Highest CPU losses were due to a combination of severe bugs in our software for Oracle access (#75596, #75508) plus a yet unknown Oracle problem at one site
- A. Valassi:
 - We will follow that up in CORAL bug #75596 and CORAL bug #58522 (the ORA-24327 is indeed our current highest priority and we are doing a lot of tests and prototyping for this).

POOL and CORAL - Bugs: bug #58522, "ORA-24327: need explicit...

Group Main Homepage Docs Support **Bugs** Tasks

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

bug #58522: "ORA-24327: need explicit attach" while re-connecting

Submitted by: **Andrea Valassi <valassi>** [Submit Changes](#)
Submitted on: 2009-11-09 18:09

<i>Category:</i>	CORAL (generic)	<i>Severity:</i>	3 - Normal
<i>Priority:</i>	7 - High	<i>Status:</i>	None
<i>Privacy:</i>	Public	<i>Assigned to:</i>	None
<i>Open/Closed:</i>	Open	<i>Fixed Release:</i>	None
<i>Effort:</i>	0.00	<i>Planned Release:</i>	

(-) **Discussion**

[Jump to the original submission](#)

2010-11-22 16:08, comment #19:
Affected ATLAS reprocessing see [bug #75573](#)

Alexandre Vaniachine <vanyash>

> 1 year



The Bus Error Bug



ATLAS Reprocessing - Bugs: bug #75613, Task 201956: Athena hangs instead...

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 #75613: Task 201956: Athena hangs instead of libSISconePlugin Bus error exit

Submitted by: Alexandre Vaniachine <vanyash>
Submitted on: 2010-11-23 10:55

Submit Changes

Category: Other software problem
Priority: 9 - Immediate
Status: Duplicate
Assigned to: None

Severity: 5 - Blocker
Item Group: Software-related
Privacy: Public
Open/Closed: Open

(-) Discussion

2010-11-23 10:55, **original submission:**

This looks like a duplicate of [bug #74015](#) but with a different behavior - instead of the Bus error exit Athena hangs for days.
Given the tight reprocessing schedule, such behavior results in lost events.

Alexandre Vaniachine
<vanyash>



- Despite this bug no events were lost in reprocessing
 - Just the CPU- years were lost



Heavy Ion Processing - Same Problems

- Job
 - <http://panda.cern.ch/server/pandamon/query?job=1151758756>
- Reconstruction goes well to the end of file, but gets a bus error at the termination phase
- This is understandable as it needs some previously opened shared libraries which disappeared when release was re-installed
- The next phase is the “crime” which we are desperately fighting for months:
 - Athena tries to produce gdb backtrace and of course fails and enters an infinite loop
- As a result, we are wasting thousands CPU-days

Reprocessing Requirements

- We need a working Athena error reporting system
 - Requested at the previous reprocessing post-mortem almost two years ago
- Since reprocessing will be more lengthy in duration, we need a continuous error analysis service
 - based upon new technologies like “streaming data mining”
- There should be no endless loops at the end of the job
 - Instead, we should be able to turn on the gdb backtrace for the looping grid jobs on-demand, e.g. upon resubmission of the job after the failed first attempt
- To reduce data corruption there should be a working (= robust and fast) validation of each output file