

A Large Ion Collider Experiment

---



# TITAN INTEGRATION. STATUS REPORT

Pavlo Svirin (National Academy of Sciences of Ukraine (UA))

Andrey Kondratyev (Joint Inst. for Nuclear Research (RU))

---

# TITAN GENERAL INFORMATION



## Current projects running on Titan (six critical codes):

- Molecular dynamics (LAMMPS), molecular physics (S3D)
- Simulations for nuclear reactions (Denovo)
- Global atmosphere model (CAM-SE)
- Astrophysics (NRDF)
- Thermodynamics (WL-LSMS)



## TITAN GENERAL INFORMATION

<b>Architecture</b>	18,688 AMD Opteron 6274 16-core CPUs, 18,688 Nvidia Tesla K20X GPUs
<b>Operating system</b>	Cray Linux Environment
<b>Memory</b>	693.5 TiB (584 TiB CPU and 109.5 TiB GPU)
<b>Storage</b>	40 PB, 1.4 TB/s IO Lustre filesystem
<b>Speed</b>	17.59 petaFLOPS (LINPACK)
<b>Ranking</b>	TOP500: #2, June 2013



## PROJECT TARGETS

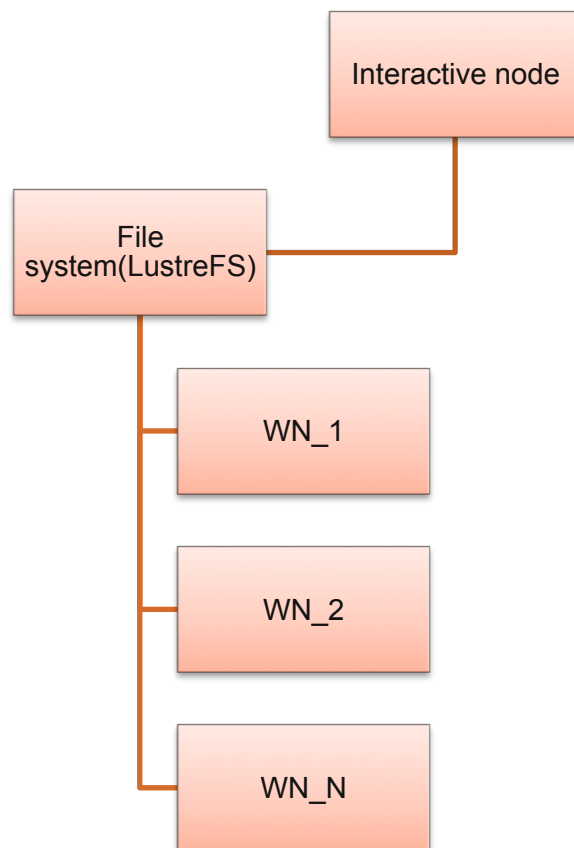
- Integrate ORNL Titan into AliEn Grid environment
- Update our Grid software according to specifics of Titan
- Recompile binary modules taking into account the specifics of operating system on Titan (Dario knows more...)
- Run Monte-Carlo jobs



## GRID SIDE UPDATES FOR TITAN

- Installed and managed to run the Vobox on Titan's interactive node
- Solved the problem with Grid certificate (thanks to Miguel)
- Implemented the Vobox module for Titan and managed to pass some jobs to the worker nodes which started successfully
- However ...

# TITAN ARCHITECTURE FOR ALICE



- Interactive node (IN) is the only one having internet access
- Communication between IN and worker nodes (WN) is done via a file system
- A static copy of CVMFS shared across the worker nodes



## JALIEN UPDATE

- JobAgent
- Sandboxing routines
- Communication with Central Services  
via filesystem

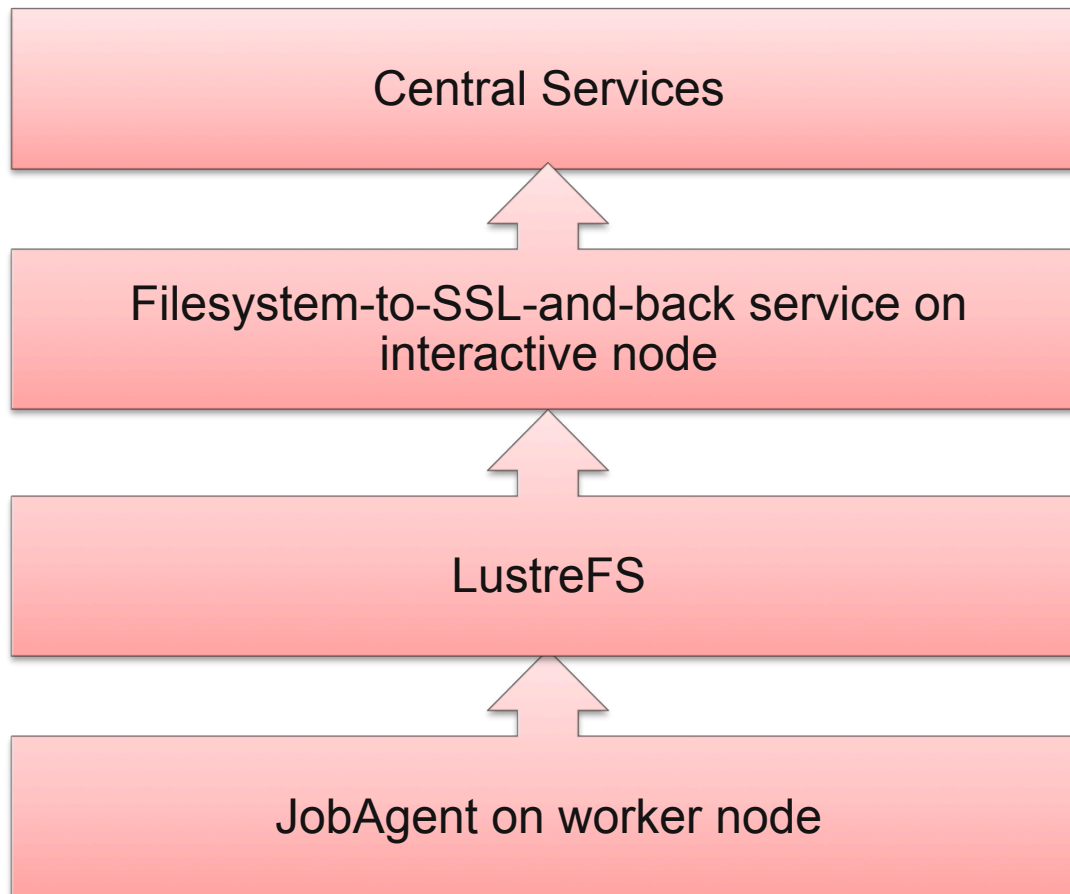


## A MINI MODEL OF TITAN

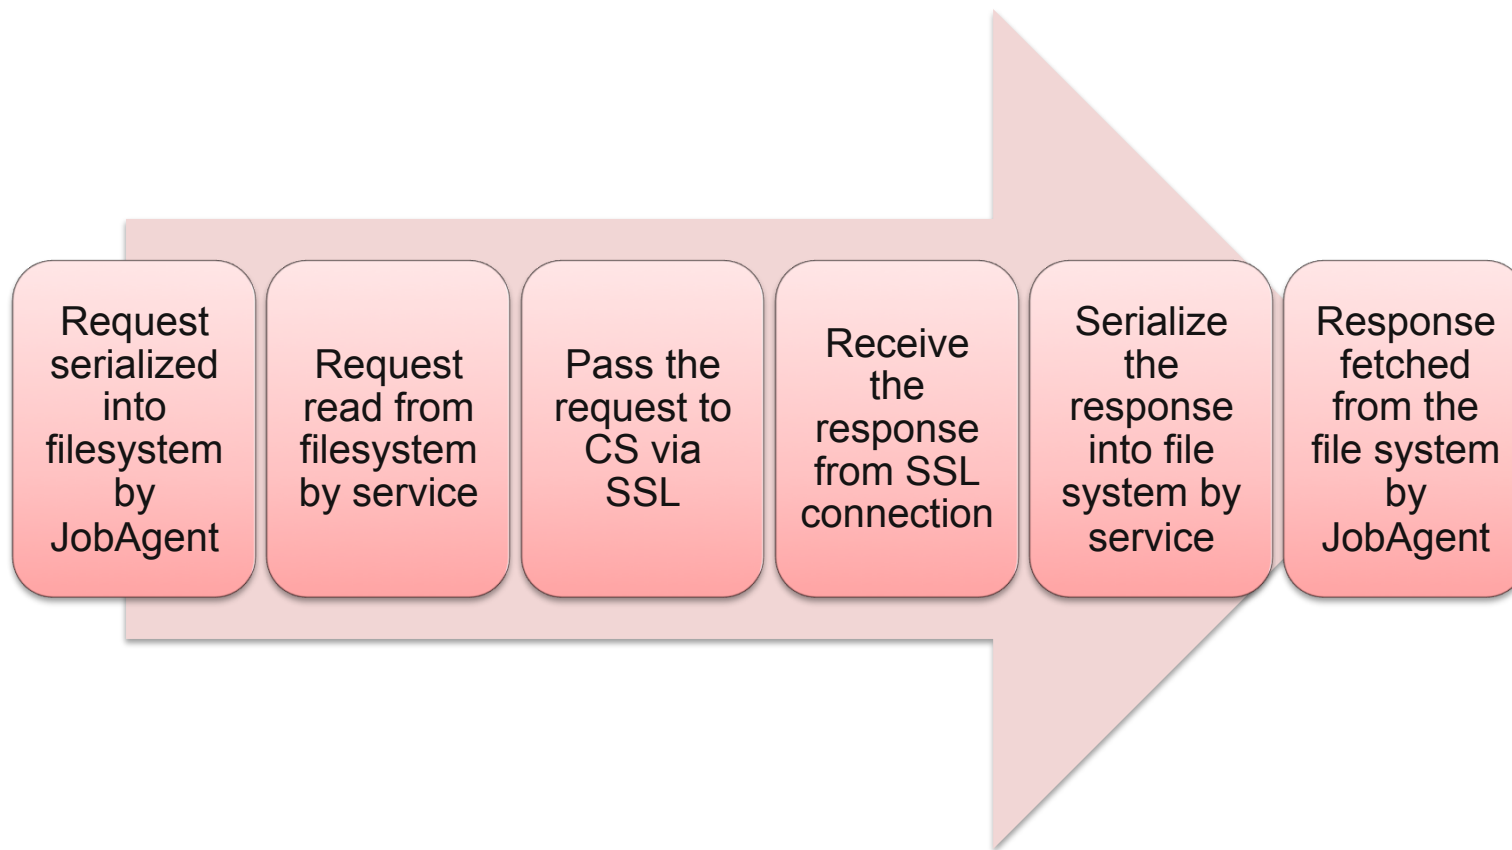
- Based on virtual machine run by Virtualbox (thanks to Costin)
- Has no network
- Emulates LustreFS through a shared Virtualbox folder
- Has an rsync'ed copy of CVMFS
- My desktop – served as an interactive node



# GENERAL INTERACTION FOR REQUESTS



# REQUEST DISPATCHING





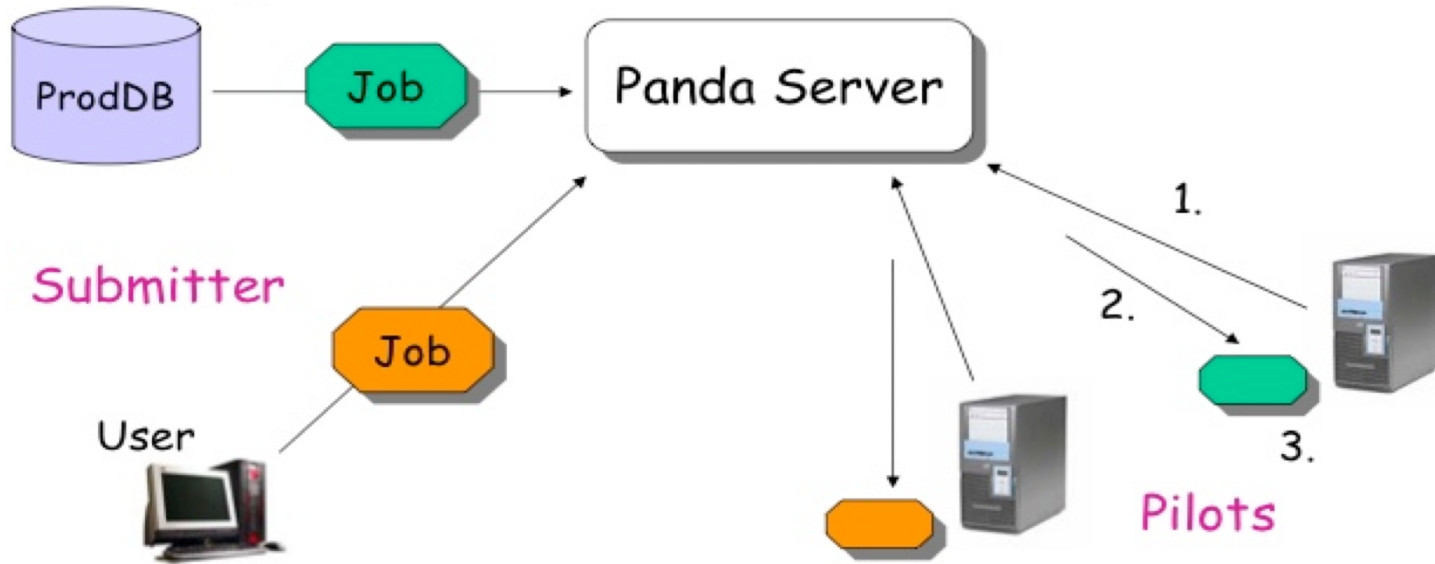
## PANDA WMS INTEGRATION

PanDA (Production and Distributed Analysis) system was originally developed for the ATLAS experiment at CERN's Large Hadron Collider (LHC) by teams at Brookhaven National Laboratory and the University of Texas at Arlington.

While PanDA continues to be developed for ATLAS, where it is used for the massively scaled distributed data-intensive production and analysis processing of the experiment at over 100 sites globally (150k concurrent jobs around the clock, a million jobs a day, analyzing a data set currently 150 petabytes in size), it is also being generalized, extended and packaged for use by other scientific communities through the BigPanDA project supported by the US Department of Energy.

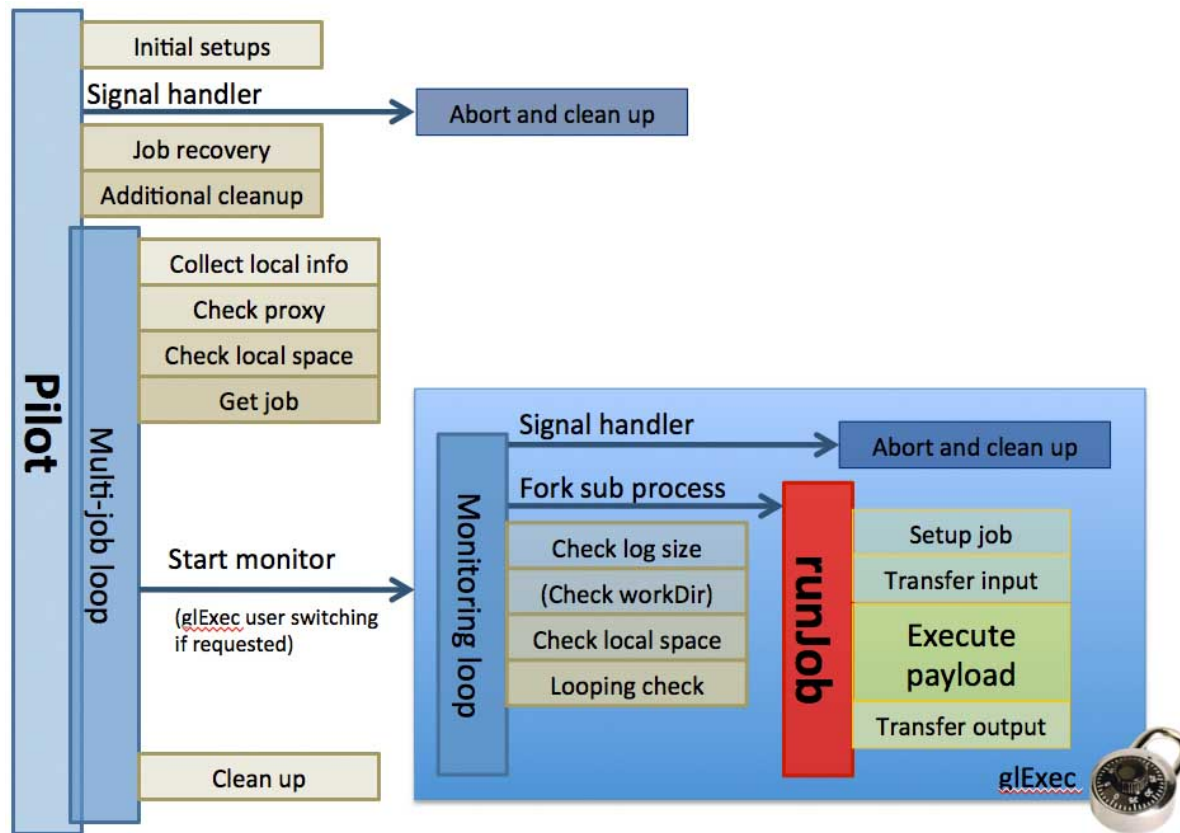
# PILOT FACTORY

Production system



Each pilot runs on a worker node  
1. send a request  
2. receives a job  
3. runs the job

# PILOT WORKFLOW

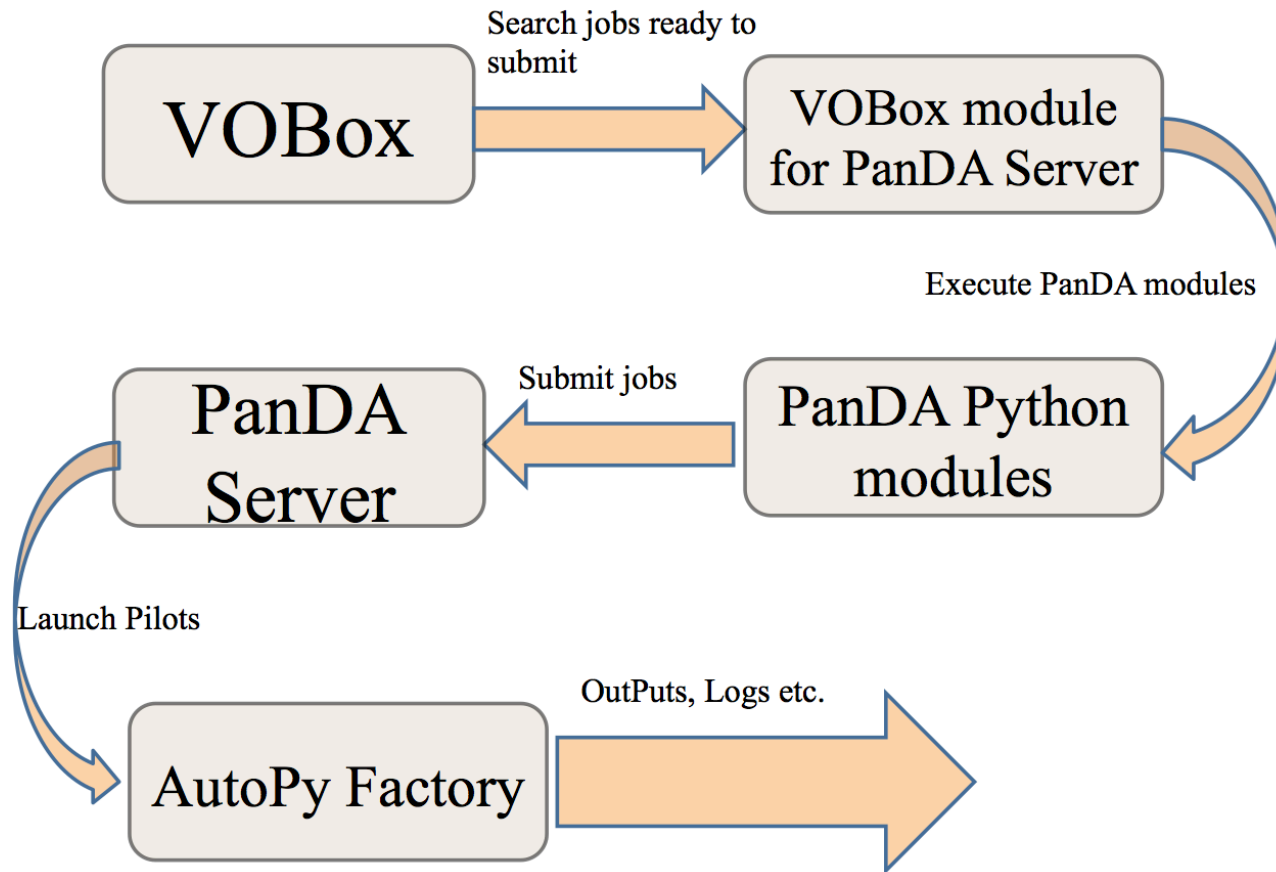




## PANDA INTEGRATION: COMPLETED STEPS

- Vobox module which operates with Panda tools (thanks to Andrey)
- Job submission to PanDA Server hosted in Amazon Cloud and uses Titan
- Implemented (need a further testing) feature for killing PanDA jobs
- A solution for fetching logs discovered (some steps have to be completed by Titan team to use curl approach)

# ALIEN-PANDA INTERACTION





# SUBMITTED JOBS TO TITAN

PanDA jobs

pandawms.org/dev/jobs/?jobstatus=finished

ВКонтакте mail.ru почта LostFilm.TV. Mail Services JINR Webmail RuTracker iTop Login Battle Сериал Теория PanDA PanDA BigPanDAinstance1

**Job list** Sort by PandaID, ascending mod time, descending mod time, priority, attemptnr, duration

PanDA ID Attempt#	Owner / VO Group	Request Task ID	Transformation	Status	Created	Time to start d:h:m:s	Duration d:h:m:s	Mod	Site	Priority	Job info
3108 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	finished	2015-10-29 19:41	23:7:46:22	0:0:02:52	11-22 03:35	ANALY_ORNL_Titan	2000	
	Job name: 87f5e5bd-e624-43c1-bc3f-0396588e7477 #0										
	Datasets: Out: panda.destDB.278cbadb-8643-41ae-9d1d-b9308c37883e										
3107 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	finished	2015-10-29 19:40	23:7:41:23	0:0:04:55	11-22 03:30	ANALY_ORNL_Titan	2000	
	Job name: c370216f-e0f0-47a6-9edb-a81997867ac5 #0										
	Datasets: Out: panda.destDB.f87b65e6-cf10-4e01-9db7-bb7c2c89b28c										
3106 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	finished	2015-10-29 19:39	23:7:33:25	0:0:04:56	11-22 03:25	ANALY_ORNL_Titan	2000	
	Job name: 1cad0f40-deed-4f38-80f1-284e1d4ce7b7 #0										
	Datasets: Out: panda.destDB.61a1b94e-1193-46f0-bc2f-4376ce75d0e5										
3105 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	finished	2015-10-29 19:38	23:7:20:28	0:0:04:54	11-22 03:10	ANALY_ORNL_Titan	2000	
	Job name: 821fa14c-5d9b-43b7-a012-21797ce61c42 #0										
	Datasets: Out: panda.destDB.623e5c24-7fdf-486d-b870-bdab01d4959d										
3104 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	finished	2015-10-29 19:37	23:7:16:20	0:0:02:54	11-22 03:05	ANALY_ORNL_Titan	2000	
	Job name: 7b1e5e5d-43b7-a012-21797ce61c42 #0										
	Datasets: Out: panda.destDB.623e5c24-7fdf-486d-b870-bdab01d4959d										





# SUBMITTED JOBS TO TITAN

PanDA jobs

pandawms.org/dev/jobs/?jobstatus=failed

vo (1) | alice (1100)

**Job list** Sort by PandalID, ascending mod time, descending mod time, priority, attemptnr, duration

PanDA ID Attempt#	Owner / VO Group	Request Task ID	Transformation	Status	Created	Time to start d:h:m:s	Duration d:h:m:s	Mod	Site	Priority	Job info
3298 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	failed	2015-10-29 22:58	26:8:44:42	0:3:01:41	11-25 10:50	ANALY_ORNL_Titan	2000	Trf exit code 244. trans: Unknown transexitcode error code 244
	Job name: 953464d0-d2a3-4777-879b-e4a4c2d931de #0										
	Datasets: Out: panda.destDB.fd4c9d05-cb58-460b-ada7-d6af5bd0ed0e										
3297 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	failed	2015-10-29 22:57	26:5:45:28	0:2:59:03	11-25 07:45	ANALY_ORNL_Titan	2000	Trf exit code 244. trans: Unknown transexitcode error code 244
	Job name: 31f49487-d30f-4885-a163-332885f29a1b #0										
	Datasets: Out: panda.destDB.4fd6e274-0ab4-42de-865b-02691ad9e245										
3296 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	failed	2015-10-29 22:56	26:4:46:27	0:0:58:44	11-25 04:45	ANALY_ORNL_Titan	2000	Trf exit code 244. trans: Unknown transexitcode error code 244
	Job name: 821f04f3-8794-4ea3-9056-4a38e7388211 #0										
	Datasets: Out: panda.destDB.386b8a9f-4022-4ce7-89fc-8a2dfc30f137										
3295 Attempt 0	Andrey Kondratyev / alice	1	mpi_wrapper_alice_A01alicegeo.py	failed	2015-10-29 22:55	26:4:19:31	0:0:27:08	11-25 03:45	ANALY_ORNL_Titan	2000	Trf exit code 244. trans: Unknown transexitcode error code 244
	Job name: 50422ad7-2771-443f-bb0e-3d0545550209 #0										



## CONCLUSIONS AND FUTURE WORK

- Complete testing for AliEn software running in a networkless environment
- Complete testing for PANDA module
- Sign OLCF agreement and get the access to Titan
- Finish binary software adjustment according to CrayOS specifics (by build masters)