



# WP2 report. Extending the scope

Sergey Panitkin



# CVMFS on Titan

- ◆ Cvmfs on parrot can be run on login nodes on Titan
- ◆ Titan admins copied ALICE cvmfs to the common file system on Titan so the it's visible on worker nodes
- ◆ Copy of the ATLAS cvmfs started three weeks ago. Moved to DT nodes because of general slowness. Still in progress.
- ◆ Will continue with tests of ATLAS software on worker nodes after cvmfs is available

The background image shows the ALICE detector components, including a large circular structure on the left and a complex network of pipes and machinery on the right. The word "ALICE" is overlaid in large white letters in the center.

# ALICE

- ◆ Discussed with Ken Read possible ALICE workloads for tests on Titan
- ◆ MPI wrapper is developed for GEANT4v10 (multithreaded) and tested on Titan
- ◆ Can run multi-node GEANT4 simulations with every GEANT4 job running in multithreaded mode
  - ◆ Year 3 deliverable (multi core/node batch slots, inter-core coordination)
- ◆ Ran two realistic/real workloads from ALICE, (simulation, digitization reconstruction chain in one job)
  - ◆ MPI wrappers with GEANT seed management
  - ◆ First ran in stand alone mode
  - ◆ Then ran with backfill capability on. Danila will comment on lessons learned.
- ◆ ALICE is interested in using PanDA for larger scale tests on Titan.
- ◆ Want more control in task definition.
- ◆ Suggested that they use PanDa server on EC2.



# PanDA Server on EC2

- ◆ Development server set up by Jarka
- ◆ Used for development and LSST tests
- ◆ Titan queue already defined in the server on EC2
- ◆ Jarka needs user's DN string from a Grid certificate to allow access
  - ◆ Me, Danila and Ken submitted DN's already
- ◆ I was able to submit a job for Titan to the server
- ◆ Monitoring is rudimentary but probably sufficient for current needs
- ◆ We will test the job submission chain via EC2
- ◆ The idea is that code and MPI wrappers will be developed on Titan
  - ◆ ALICE people, using PanDa user tools will submit job descriptions to EC2 server
  - ◆ Our pilot on Titan will contact PanDA@EC2 and run the jobs (using backfill option or not, see Danila's talk)
  - ◆ Output will remain local and moved by ALICE if needed.



# Backfill Tests

- ◆ Ran more backfill tests on Titan with ALICE workload (ppbench)
  - ◆ Realistic ~20 min of running per job
  - ◆ Ran with MPI wrapper , 1 rank per core
  - ◆ In some tests used several hundred worker nodes (each with 16 cores!)
- ◆ Se pretty much the same behavior that we saw with Hello World tests
- ◆ Wait time is small
- ◆ Scheduler seems to prefer jobs with space-time parameters that match currently available free resources
- ◆ We didn't see penalty for continuous running

# Some Backfill Test Results

JobID	Submitted	Nodes	Wait	Runtime	State	Completed	Comment
1931570	Apr, 29 17:43	19	0.02	0.25	Completed	Apr, 29 17:59	
1931605	Apr, 29 18:21	1.132	71.44		Removed	Apr, 29 18:31	Execution manually removed, due part of allocation was 'stolen' by another user, and job was rescheduled
1931611	Apr, 29 18:32	872	0.03	2.15	Completed	Apr, 29 20:43	Jumbo output. Pilots execution time few hours due to massive disk operations.
1932104	Apr, 30 13:41	1	0.05	0.20	Completed	Apr, 30 13:56	
1932133	Apr, 30 14:07	20	0.34	0.26	Completed	Apr, 30 14:43	
1932187	Apr, 30 14:56	77	0.01	0.30	Completed	Apr, 30 15:15	
1932205	Apr, 30 15:21	32	0.00	0.27	Completed	Apr, 30 15:37	
1932218	Apr, 30 15:40	34	0.01	0.27	Completed	Apr, 30 15:57	
1932238	Apr, 30 16:00	38	0.20	0.26	Completed	Apr, 30 16:28	
1932257	Apr, 30 16:30	103	49.29		Removed	Apr, 30 16:38	Execution manually removed, due part of allocation was 'stolen' by another user, and job was rescheduled
1932262	Apr, 30 16:40	93	0.47	0.37	Completed	Apr, 30 17:31	
1932380	Apr, 30 17:35	44	0.01	0.29	Completed	Apr, 30 17:53	
1932877	May, 01 11:38	103	0.00	0.39	Completed	May, 01 12:01	
1932910	May, 01 12:35	43	0.02	0.27	Completed	May, 01 12:52	
1932953	May, 01 13:14	2	0.00	0.20	Completed	May, 01 13:26	
1932994	May, 01 13:47	281	0.00	0.96	Completed	May, 01 14:45	
1933061	May, 01 15:16	732	0.04	1.87	Completed	May, 01 17:11	
1933236	May, 01 17:42	837	0.06	0.49	Completed	May, 01 18:15	Output contain about 13K directories. Was not able to serve by Pilot for the few hours.
1933343	May, 01 22:02	11	0.01	0.36	Completed	May, 01 22:24	
1933560	May, 02 09:09	7	0.02	0.36	Completed	May, 02 09:32	
1933595	May, 02 09:52	1.234	7.92		Removed	May, 02 10:12	Execution manually removed, due debug reasons
1933667	May, 02 11:04	50	0.00	0.48	Completed	May, 02 11:33	
1933685	May, 02 11:42	50	0.00	0.48	Completed	May, 02 12:11	
1933733	May, 02 12:15	7	0.03	0.38	Completed	May, 02 12:40	
1933762	May, 02 12:42	50	0.00	0.48	Completed	May, 02 13:11	
1933787	May, 02 13:15	50	0.00	0.49	Completed	May, 02 13:45	
1933826	May, 02 13:48	50	0.00	0.49	Completed	May, 02 14:17	
1933899	May, 02 14:21	50	0.00	0.46	Completed	May, 02 14:49	
1933991	May, 02 14:53	50	0.00	0.40	Completed	May, 02 15:17	Job exceed walltime, and was interrupted by sheduler. Intermediate (temporary) data was pcking by pilot. Output 35Gb
1934149	May, 02 17:02	3	0.08	0.19	Completed	May, 02 17:18	
1934156	May, 02 17:20	4	0.02	0.20	Completed	May, 02 17:34	



# eRHIC

- ◆ Talked to Alexander Kiselev from eRHIC group at BNL
  - ◆ He has expressed interest in tests on Titan
  - ◆ I attempted FairRoot installation on Titan.
    - ◆ Problems with Geant 3 installation
  - ◆ Agreed to work on this with Alexander this week



# Miscellaneous

- ◆ Our project on Titan was extended for another year. 500k CPU hours allocation.
- ◆ Got access to the newest machine at ORNL - EOS (Cray XC-30)
  - ◆ Deployed recently
  - ◆ Similar to Edison at NERSC
  - ◆ Somewhat different software setup than on Titan
- ◆ Technical writer from ORNL wrote a story about PanDA on Titan.
  - ◆ Comments from BNL, ORNL, UTA