

GRID COMPUTING IN THE ATLAS EXPERIMENT AT THE LHC

Johannes Elmsheuser

Ludwig-Maximilians-Universität München

11 December 2009/Göttingen



CHALLENGES IN DATA ANALYSIS



Data volumes

- ATLAS produces and stores several PetaBytes/year

CPUs

- Event complexity (large number of channels) and number of users demands: at least 100000 fast CPUs based on computing model

Software


- ATLAS has complex software environment and framework

Connectivity

- Data should be available 24/7 at a high bandwidth

AVERAGE ANALYSIS AT LHC I

Higgs-Search: $H \rightarrow WW^{(*)} \rightarrow \mu^+ \nu_{\mu} \mu^- \bar{\nu}_{\mu}$ für 1 fb^{-1}




Monte Carlo events needed :

- 4 mass points: $m_H = 130 - 190 \text{ GeV}$: 100k + 500k Systematic studies
- Background: Z/γ^* : 2M, $t\bar{t}$: 500k, $WW+WZ+ZZ$: 200k, $W+\text{jets}$: 1M
- Total: 4.3M
- Time for simulation: 200h @ 10000 CPUs with 0.5h/event (no overhead)

Data:

- 10^9 Events/year
- $\approx 50\text{d}$ time for reconstruction @ 10000 CPUs with 45s/event



Analysis:

- 10^6 data events from trigger and skim pre-selection
- Estimated time:
 - 1 week MC+data at 1 CPU with 10Hz
 - 4h MC+data at 1000 CPUs (Tier2-share)
 - Optimization of analysis demands much more time

Scaling up:

- Assume 2000 physicist with same analysis
- Time: 3h at 100000 CPUs
- Shown analysis is not the most time consuming
- Analysis with jets need much more CPU-time
- All given time: without additional overhead



Data

- Centrally organized data distribution by DDM/DQ2 according to computing model

Athena distribution kits

- Centrally organized installation on EGEE, OSG and NG

User jobs

- Model: „Job to Data”
- Tools for job management: Panda clients and Ganga

User Output:

- Store on site *SCRATCHDISK* or transfer to remote *LOCALGROUPDISK*
- Retrieval with dq2 command line tools


ATLAS GRID INFRASTRUCTURE

- Heterogeneous grid environment based on 3 grid infrastructures:




- Three major ATLAS Grid areas:
 - Production System (Panda): centralized MC simulation and Data reconstruction
 - Distributed Data Management (DDM/DQ2): centralized data movement
 - Distributed User Analysis: de-centralized individual analysis


COMPUTING CENTER AND TIER STRUCTURE



✓Left: All sites, ✓Right: Organized Tier Structure



ATLAS PRODUCTION SYSTEM JOBS - LAST YEAR



DDM/DQ2 TRANSFERS - 24H

[Data: All Activities](#)[Jobs: Production](#)[Jobs: Analysis](#)[Panda: Production](#)[SLS: Central Services](#)[Overview](#)[Dataset Info](#)[Page Help](#)[User Guide](#)[Feedback](#)

OVERVIEW

Activity Period

[Activity in Last Hour](#)[Activity in Last 4 Hours](#)[Activity in Last 24 Hours](#)[Activity in Last 7 Days](#)[Activity in Last 30 Days](#)[Activity in ...](#)


Selected Activities

[Production](#)[T0 Export](#)[Functional Test](#)[User Subscriptions](#)[Staging](#)[Data Consolidation](#)


Selected Cloud

[ASGC Cloud](#)[BNL Cloud](#)[CERN Cloud](#)[CNAF Cloud](#)[FZK Cloud](#)[LYON Cloud](#)[NDGF Cloud](#)[PIC Cloud](#)


Throughput (MB/s)




Data Transferred (GBytes)



Completed File Transfers



Total Number Transfer Errors



Activity Summary ('2009-12-08 20:00' to '2009-12-09 20:00' UTC)

DDM/DQ2 TRANSFERS - 30D

[Data: All Activities](#)[Jobs: Production](#)[Jobs: Analysis](#)[Panda: Production](#)[SLS: Central Services](#)[Overview](#)[Dataset Info](#)[Page Help](#)[User Guide](#)[Feedback](#)

OVERVIEW

Activity Period

[Activity in Last Hour](#)[Activity in Last 4 Hours](#)[Activity in Last 24 Hours](#)[Activity in Last 7 Days](#)[Activity in Last 30 Days](#)[Activity in ...](#)


Selected Activities

[Production](#)[T0 Export](#)[Functional Test](#)[User Subscriptions](#)[Staging](#)[Data Consolidation](#)


Selected Cloud

[ASGC Cloud](#)[BNL Cloud](#)[CERN Cloud](#)[CNAF Cloud](#)[FZK Cloud](#)[LYON Cloud](#)[NDGF Cloud](#)


Throughput (MB/s)




Data Transferred (GBytes)



Completed File Transfers



Total Number Transfer Errors



Activity Summary ('2009-11-09 00:00' to '2009-12-09 00:00' UTC)

DDM/DQ2 TRANSFERS - TIER0 - 30D

[Data: All Activities](#)[Jobs: Production](#)[Jobs: Analysis](#)[Panda: Production](#)[SLS: Central Services](#)[Overview](#)[Dataset Info](#)[Page Help](#)[User Guide](#)[Feedback](#)

OVERVIEW

Activity Period

- [Activity in Last Hour](#)
- [Activity in Last 4 Hours](#)
- [Activity in Last 24 Hours](#)
- [Activity in Last 7 Days](#)
- [Activity in Last 30 Days](#)
- [Activity in ...](#)


Selected Activities

- [Production](#)
- [T0 Export](#)
- [Functional Test](#)
- [User Subscriptions](#)
- [Staging](#)
- [Data Consolidation](#)


Selected Cloud

- [ASGC Cloud](#)
- [BNL Cloud](#)
- [CERN Cloud](#)
- [CNAF Cloud](#)
- [FZK Cloud](#)
- [LYON Cloud](#)
- [NDGF Cloud](#)
- [PIC Cloud](#)
- [SARA Cloud](#)


Throughput (MB/s)




Data Transferred (GBytes)



Completed File Transfers




Total Number Transfer Errors



Activity Summary ('2009-11-10 00:00' to '2009-12-10 00:00' UTC)

[Click on the cloud name to view list of sites](#)

DATA DISTRIBUTION: ATLAS




≈ 80 Tier1/2/3 sites managed by DQ2 right now

EVENT DATA MODEL: ATLAS

Missing here: dESD and dAOD (Distribution similar to DPD/AOD)

Refining the data by: Add higher level info, Skin, Thin, Slim






Naive assumption: Grid $\approx$ large batch system

- Provide complicated job configuration for Workload Management System
- Find suitable experiment software, installed in the Grid (100 CEs, 30 Software versions)
- Locate the data on different storage elements
- Job splitting, monitoring and book-keeping
- etc.

⇒ Need for automation and integration of various different components

Several ways lead into the Grid !

ATLAS DISTRIBUTED ANALYSIS




Data is centrally being distributed by DQ2 - Jobs go to data


Long ADC review: current model is fine, Ganga/Panda as a focus

DISTRIBUTED ANALYSIS: GANGA

How to combine all different components: **Job scheduler/manager:**
GANGA




JOB SCHEDULING



Job Push mode


- Dependent on information system and site status
- Decentralized
- Better control of site policies
- Ganga: LCG and NG backend



Job Pull mode

- Workarounds for some Grid problems
- Data pre-staging
- Panda clients or Ganga Panda backend


NUMBER OF ATLAS ANALYSIS JOBS I



- Daily 10-20k Jobs on LCG and Panda
- Exclude HammerCloud tests (JE certificate)

- Compare with daily ~ 50 - 100 k productions jobs
- Since beginning of the year increasing number of users - but many more expected !

NUMBER OF ATLAS ANALYSIS JOBS II



Example: Number of Analysis jobs on Panda
all peaks are due to HammerCloud tests - many analysis slots available

ATLAS is testing since several month all sites with very high automatic generated analysis load: HammerCloud

<http://gangarobot.cern.ch/hc>



Differences Analysis vs. MC Production:

- „unorganized” user analysis vs. „organized” MC production
- User Analysis puts much higher load on SE compared to CPU dominated simulation


Tests of different work-flows:

- Sequential AOD analysis of MC data
- Sequential cosmics analysis with DB/Frontier/Squid access

Some highlights:

- Analysis tools generally stable and reliable
- Some weak spots detected in site infrastructures, especially in input file access mode lots of tuning potential

HAMMERCLOUD WORKFLOW



HammerCloud uses MySQL DB - Web pages generated with Django

SOME TEST RESULT EXAMPLES

Different input file access modes:



- Ganga/WMS, DQ2_LOCAL: dcap/gsidcap, rfio, file, root during athena event loop
- Ganga/WMS, FileStager: copy-to-scratch (lcp-cp or rfcg/dccp) during athena event loop
- Panda on OSG, EGEE usually uses copy from local SE to worker node scratch disk in advance of athena event loop
- ARC offers local disk cache buffer
- FileStager and file protocol showed good event rates and CPU/Walltime in most HammerCloud tests
- dcap/gsidcap usually suffered from bad read-ahead buffer tuning, have now improved read-ahead buffer algorithm, ROOT and EDM team are actively working on improving file storage and input

Why is input file streaming important ?



- Conditions data pool files are accessed with dcap/rfio
- TAGs
- Don't access file twice, Multi-core worker nodes

HAMMERCLOUD RESULTS LRZ-LMU, GOEGRID



Default: HC 768, dcap



With libdcap.so patch,
HC 798





With libdcap.so patch,
HC 798





HAMMERCLOUD RESULTS DE AND LRZ-LMU

libdcap.so patch HC 798, all DE
(some DPM site, heavy loaded FZK)



FileStager, HC 748, only LRZ-LMU



ATLAS USER ANALYSIS TESTS (UAT) I

Took place end of October

Two main topics:

- Replicate huge UAT container dataset into all clouds in advance
- Experienced users submitting their analysis to the ATLAS Grid infrastructure

Goal:

- Identify bottlenecks in the infrastructure, tools, work-flows
- User feedback

Observations:

- Generally worked well
- Post mortem meeting on Thursday
- Some additional HammerCloud tests as base line and additional massive test

ATLAS USER ANALYSIS TESTS (UAT) II

Massive Test Results


- After UAT send 2 analysis to all sites using Panda
- Test duration 36h, ramped up to ~ 8000 simultaneous running jobs
- Some bias due to remnant user jobs in some clouds
- STEP09 (2 weeks): 970k jobs, 2.9m files, 20 bio events, 7.1 PB

	<u>Num Jobs</u>	<u>Frac. Failed</u>	<u>Num Events</u>
DE	25871	0.08	1309996110
US	23909	0.09	1208817132
FR	19967	0.07	961222848
UK	15970	0.08	872697795
NL	7136	0.15	229160559
ES	3968	0.00	181877439
ND	2118	0.13	27737499
TW	224	0.57	12526419
IT	1642	0.80	11356259
CA	104	0.67	1541610
TO	1473	0.00	0
Sum	102382		4816933670
AGLT2	8621	0.00	428664530
Glasgow	6466	0.06	384738138
MWT2	6449	0.24	378889164
LRZ-LMU	6010	0.00	347971614
MPPMU	6389	0.00	326289071
NET2	4873	0.01	267089587
DESY-ZN	4886	0.00	234530348




Grid Jobs


- **Ganga/WMS:**
 - 32 different people from 20 different institutes ran Ganga jobs on the Grid, mostly with ESD, AOD or DESD_COLLICAND as input
 - Top scorers are physics_MinBias.recon.ESD
 - Several 100s jobs/dataset
- **Ganga/NG:**
 - A few 100 jobs, mostly on DESD_COLLICAND datasets
- **pAthena:**
 - Most popular data are physics_MinBias.merge.DESD_COLLICAND
 - 9-15 users/dataset launching 100-300 jobs



Panda sites



Panda users



CURRENT USER PROBLEMS AND SUPPORT

Frequently asked questions or problems:

- Problems with conditions data access
- There is a problem with my special code configuration
- The job had problems with accessing the input data files
- The ratio of CPU and Wall-time is largely varying btw. 10% - 100% and depends on the site and user



Support:

- ATLAS wide user support mailing list for DA
- Shifters in EU and US time zone
- Hoping for more user2user support
- Has developed to one of the busiest mailing lists in ATLAS



CURRENT ATLAS DA ISSUES

- Trying to improve user output work-flows and output consolidation
- Trying to improve panda job scheduling
- Improve input file access modes, esp. for direct access (HammerCloud)
- Improve site overall analysis stability ! more active black-listing (Daily GangaRobot)
- Some problems with Conditions data access fixed with recent athena package updates
- Real Data: InputFilePeeker problematic for DPM rfio
- DAST is handling the user support load so far, but load will increase because of new responsibility to support DDM issues and maintain blacklists of analysis sites.
- DAST Recruitment and training is proceeding well; a plan to double the number of shifts during busy weeks is feasible.

CONCLUSIONS AND SUMMARY

What is working well:

- MC Production
- Automatic data distribution
- Analysis:
 - At most of the sites
 - Small scale MC production
 - Automatic Standard Job Configurations

What works, but needs improvement:

- Analysis:
 - 'Blind' job submission
 - Exotic use cases
- Overall site availability and input file access

For the distributed analysis it is vital to have:

- Easy interface that does not scare off physicists
- A reliable and robust service of many components