MONITORING OF DISTRIBUTED DATA PRODUCTION AND ANALYSIS FOR HIGH ENERGY PHYSICS EXPERIMENTS

Tatiana Korchuganova

UNAB (Santiago, Chile), ISP RAS (Moscow, Russia)

06/12/2019, Ivannikov ISP RAS Open Conference

OUTLINE

ATLAS DISTRIBUTED COMPUTING

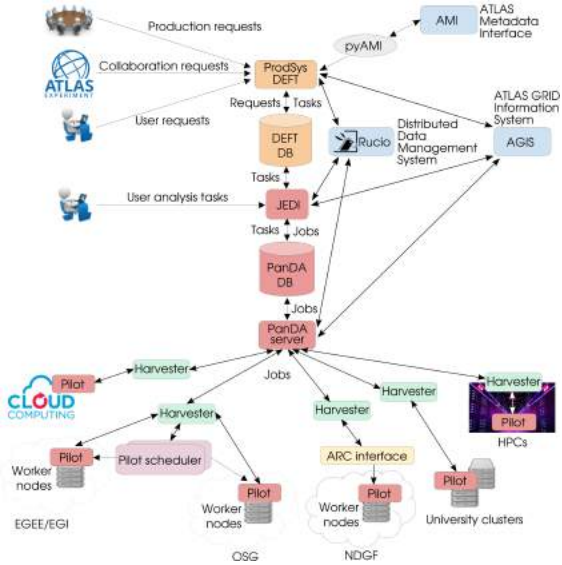

Figure 1 — ATLAS Workflow Management System scheme

- ProdSys - Production system;
- DEFT - Database Engine For Tasks;
- PanDA - Production and Distributed Analysis;
- JEDI - Job Execution and Definition Interface;
- PanDA server - the main component of the PanDA system that distributes and manages jobs among computing resources;
- Harvester - resource-facing service between the PanDA server and collection of pilots;
- AMI - ATLAS Metadata Interface;
- AGIS - ATLAS Grid Information System;

ATLAS DISTRIBUTED COMPUTING

Computing resources usage during Run2 (from 2015 to 2018)

Figure 2 — Number of running CPU core slots grouped by resource type: GRID sites, HPC with closed network infrastructure, Clouds, HPC with closed network infrastructure, BOINC - volunteered computing resources

ATLAS DISTRIBUTED COMPUTING

PanDA processing scales at ATLAS

Figure 3 — Objects of PanDA system

OUTLINE

MONITORING

1. **Monitoring of current system state, access to logs for troubleshooting, and providing a comprehensive and coherent view of the tasks and jobs executed by the system, from high level summaries to detailed drill-down job diagnostics for short-term period**;
2. A system state analytics in medium term and resource usage accounting.
3. Analytics of a system working in long term period by finding correlations, trends, anomaly detection and building models for prediction of system behaviour in future.

↓

A purpose of the BigPanDA monitoring system is providing a way for fast errors troubleshooting and tracking production and analysis tasks progress

BIGPANDA MONITORING

Requirements

- System state monitoring and troubleshooting by immediate access to logs;
- 24/7 accessibility;
- Ability to scale;
- Acceptable response time for different usage scenarios;
- Possibility to aggregate data from various sources:
  - DB (Oracle, MySQL, PostgreSQL);
  - Indexed data from ElasticSearch and Kibana visualizations;
  - log files in text or JSON format from Rucio;
  - plots from MONIT accounting service based on Grafana.
- Developing a common visualization mechanism to satisfy monitoring needs of different groups of users:
  - Physicists;
  - Operators or experts;
  - Computing site administrators;
  - Production managers;
  - Coordinators;
  - Software developers.
- Providing an API for programmatic analysis of system objects state.

Use-case Diagram for «Physisist»

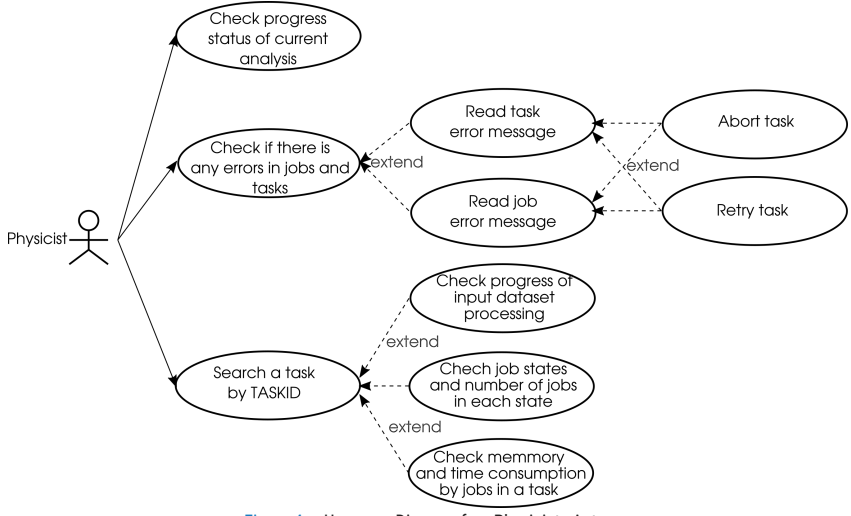

Figure 4 – Use-case Diagram for «Physisist» Actor

BIGPANDA MONITORING SYSTEM

Technology stack

1. Using Model-Template-View approach provided by Django framework;
2. Various DB backends (**Oracle**, MySQL, PostgreSQL);
3. Apache + WSGI;
4. NGINX server as load balancer;
5. Ceph as shared file storage between nodes;
6. OAuth2 for SSO provided by CERN, Google, and GitHub;
7. Angular и Ajax, jQuery for dynamic data delivery;
8. Responsive Web designing with ZURB Foundation;
9. DataTables plugin;
10. D3.js for advanced plots generation on client side;
11. Matplotlib for plots generation on server side;
12. ELK-stack for self-monitoring system (more details in a talk presented by Aleksandr Alekseev).

BIGPANDA MONITORING

Architecture

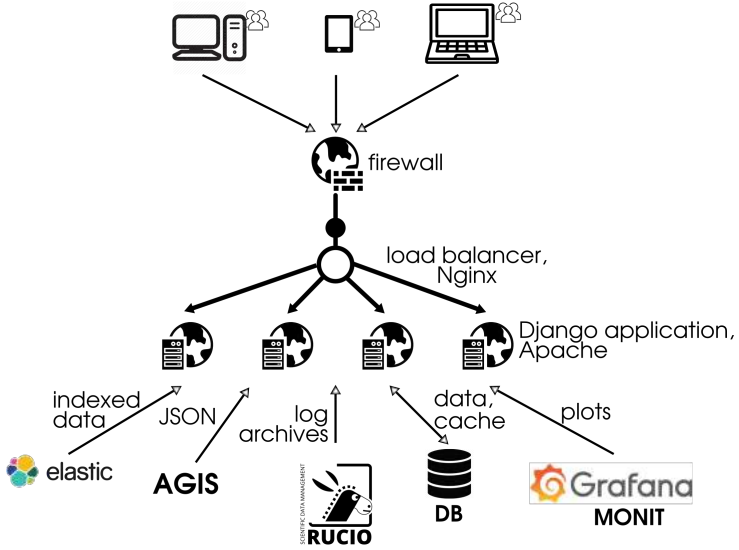

Figure 5 – Generalized scheme of BigPanDA monitoring system architecture

DATA FLOW IN BIGPANDA MONITORING SYSTEM

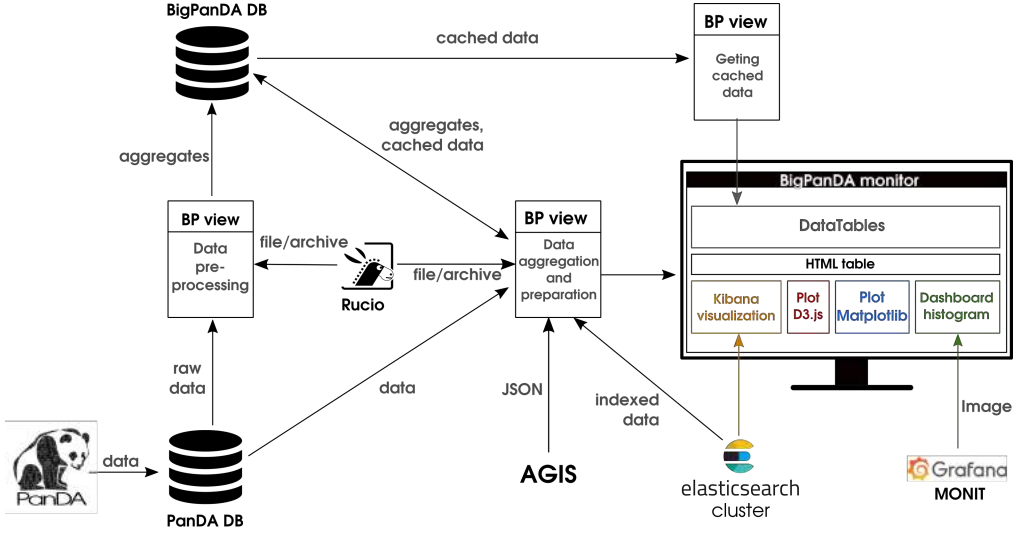

Figure 6 — Data-flow diagram of BigPanDA monitoring system

OUTLINE

USE-CASE DIAGRAM FOR «PHYSISIST»

Figure 7 — Use-case Diagram for «Physisist»

USER PAGE EXAMPLE (1/4)

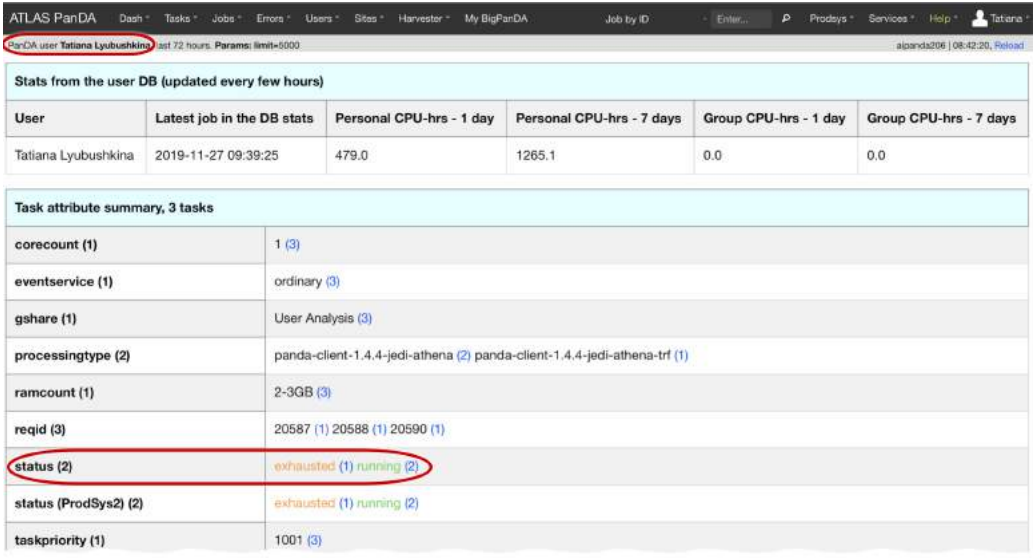

Figure 8 — Tasks attribute summary for last 3 days on the user page

USER PAGE EXAMPLE (2/4)

3 tasks, sorted by jeditaskid Show in task list page

ID Parent	Jobset	Task name TaskType/ProcessingType Campaign Group User Logged status	Task status Nfiles	Input files finish% fail% Nfinish Nfail	Started	Modified	State changed	Priority
19847334	20590	user.tlyubush.mc12_8TeV.Bcst_Bcm_JpsiMu_208491.NTUP.r42/ anal/panda-client-1.4.4-jedi-athena Tatiana Lyubushkina	running 16	6% 1	2019-11-22 15:30:08	2019-11-25 19:03:07	2019-11-22 15:31:12	1001
19847270	20588	user.tlyubush.mc12_8TeV.Bcst_Bcp_JpsiMu_208490.NTUP.r42/ anal/panda-client-1.4.4-jedi-athena Tatiana Lyubushkina	running 16		2019-11-22 15:27:00	2019-11-26 00:01:52	2019-11-22 15:27:48	1001
19820547	20587	user.tlyubush.MC16.999015.deriv.DAOD_BC3MU.e7289_a875_r10724_p3809_r700/ anal/panda-client-1.4.4-jedi-athena-trf Tatiana Lyubushkina log brokerage failed for 1 input datasets when trying 1 datasets. timeout while in pending since 2019/11/20 16:31:57	exhausted 20		2019-11-20 16:30:35	2019-11-21 17:36:39	2019-11-21 17:36:39	1001

Figure 9 — List of tasks for last 3 days on the user page

USE-CASE DIAGRAM FOR «PHYSISIST»

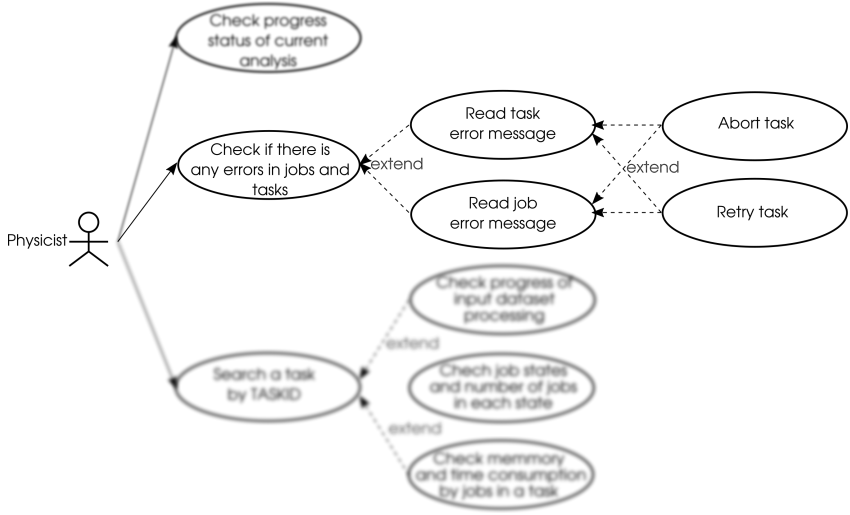

Figure 10 — Use-case Diagram for «Physisist»

TASK PAGE EXAMPLE

Figure 11 — Task page

USER PAGE EXAMPLE (3/4)

87 jobs

Job **modification times** in this listing range from **2019-11-24 08:42:16** to **2019-11-27 08:42:16**.

Job **current priorities** in this listing range from **994** to **2000**. See priorityrange in the job attribute summary to see how priorities are distributed.

Recent job attributes	
atlasrelease (1)	Atlas-17.2.14 (87)
attemptnr (7)	2 (16) 3 (16) 4 (3) 5 (19) 6 (19) 7 (2) 8 (2)
cloud (6)	CERN (16) DE (33) IT (15) ND (3) UK (2) US (18)
computingsite (7)	ANALY_ARNES_DIRECT (3) ANALY_CERN_T0 (16) ANALY_DESY-HH (16) ANALY_GOEGRID (17) ANALY_INFN-LECCE (15) ANALY_MANC_UCORE (2) ANALY_MWT2_UCORE (18)
durationmin (22)	0-0 (17) 1-180 (38) 180-359 (0) 359-538 (1) 538-717 (0) 717-896 (0) 896-1075 (0) 1075-1254 (0) 1254-1433 (4) 1433-1612 (10) 1612-1791 (0) 1791-1970 (0) 1970-2149 (2) 2149-2328 (3) 2328-2507 (5) 2507-2686 (0) 2686-2865 (2) 2865-3044 (0) 3044-3223 (3) 3223-3402 (0) 3402-3581 (1) 3581-3760 (1)
homepackage (1)	AnalysisTransforms-AtlasProduction_17.2.14.12 (87)
inputfileproject (1)	mc12_8TeV (77)
inputfiletype (1)	AOD (77)
jeditaskid (2)	19847270 (33) 19847334 (54)
jobsetid (2)	20588 (33) 20590 (54)
jobsetrange (1)	20500:20599 (87)
jobstatus (6)	assigned (7) cancelled (9) closed (1) failed (53) finished (5) running (12)

Figure 12 – Jobs attribute summary for last 3 days on the user page

USER PAGE EXAMPLE (4/4)

87 recent jobs Note that with JEDI, tasks are the best tool to monitor the progress of your submissions! See the help section below.

PanDA ID	Owner Group	Task ID	Transformation	Status	Created	Time to start d:h:m:s	Duration d:h:m:s	Mod	Cloud Site	Priority	Job info
	Tatiana Lyubushkina	19847270	runAthena-00-00-12	running	2019-11-26 00:01:51	0:7:14:32	1:1:25:57	2019-11-27 08:38:23	DE ANALY_GOEGRID	997	
4559292379	Job name: user.tlyubush.mc12_8TeV.Bcst_Bcp_JpsiMu_208490.NTUPr42/.4555880563										
	Datasets: **In:** mc12_8TeV:mc12_8TeV.208490.Pythia8B_AU2_CTEQ6L1_FakeBcst_Bcp_Jpsi_mu2p5mu2p5_mu2p5_nu.merge.AOD.e7479_a874_a270_a271_r4348_tid17988287_00 **Out:** user.tlyubush.mc12_8TeV.Bcst_Bcp_JpsiMu_208490.NTUPr42_EXT0.283040398										
	Tatiana Lyubushkina	19847270	runAthena-00-00-12	failed	2019-11-26 00:01:51	0:7:14:42	1:0:39:17	2019-11-27 08:01:23	DE ANALY_GOEGRID	996	**pilot, 1212:** Payload ran out of memory **trans, 40:** Athena crash - consult log file
4559292383	Job name: user.tlyubush.mc12_8TeV.Bcst_Bcp_JpsiMu_208490.NTUPr42/.4555880570										
	Datasets: **In:** mc12_8TeV:mc12_8TeV.208490.Pythia8B_AU2_CTEQ6L1_FakeBcst_Bcp_Jpsi_mu2p5mu2p5_mu2p5_nu.merge.AOD.e7479_a874_a270_a271_r4348_tid17988287_00 **Out:** user.tlyubush.mc12_8TeV.Bcst_Bcp_JpsiMu_208490.NTUPr42_EXT0.283040398										
	Tatiana Lyubushkina	19847270	runAthena-00-00-12	failed	2019-11-24 07:41:01	0:0:03:02	0:0:03:29	2019-11-26 00:00:51	US ANALY_MWT2_UCORE	998	**trans, 40:** Athena crash - consult log file
4557635926	Job name: user.tlyubush.mc12_8TeV.Bcst_Bcp_JpsiMu_208490.NTUPr42/.4555880586										

Figure 13 — List of jobs for last 3 days on the user page

JOB PAGE EXAMPLE

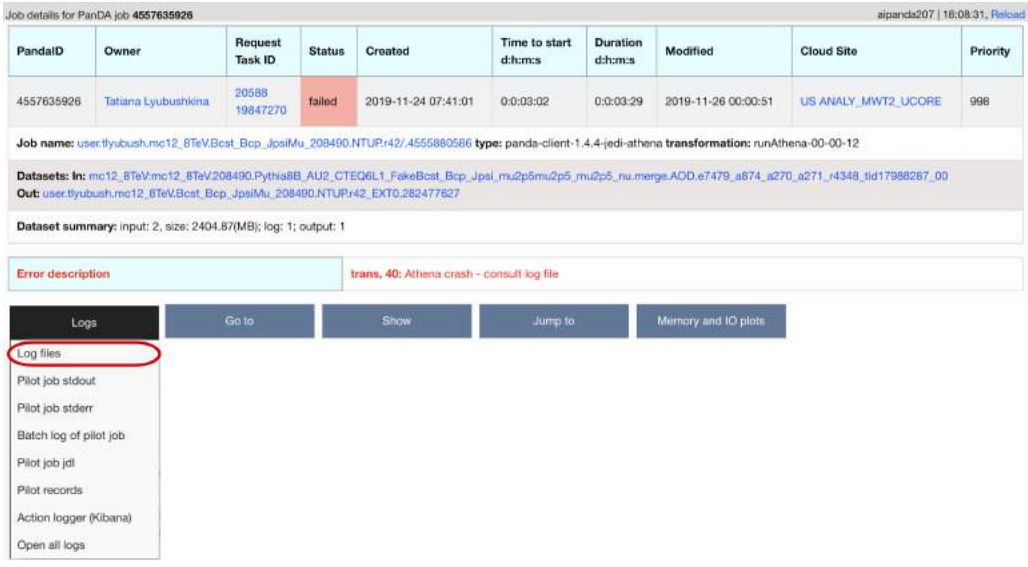

Figure 14 – Job page, a way to access job logs example

FILE BROWSER PAGE EXAMPLE

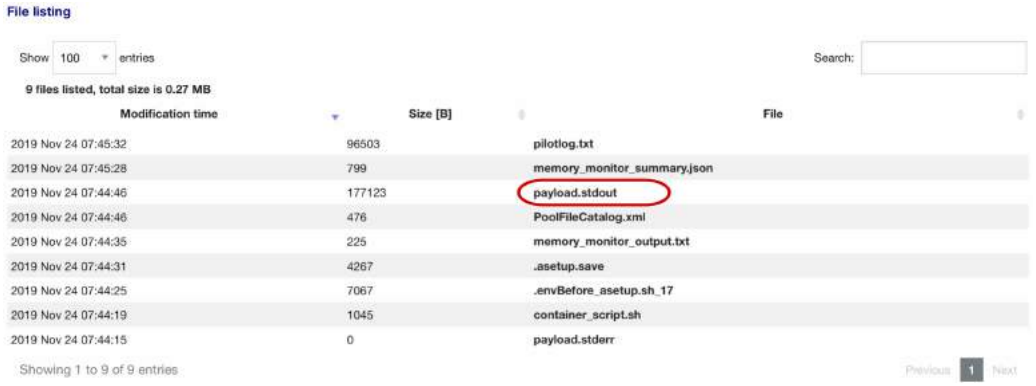

Figure 15 — File browser page example showing a content of logs archive file

LOG FILE BROWSING EXAMPLE

```
Py:InputFilePeeker   ERROR Unable to build inputFileSummary from any of the specified input files. There is probably a problem.
Py:AutoConfiguration   ERROR No RunNumber stored in InputFile!
Shortened traceback (most recent user call last):
  File "./BPhysAnalysisMasterAuto.Bc2S.MC.py", line 77, in <module>
    include ("RecExCommon/RecExCommon_topOptions.py")
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-
opt/17.2.14/AtlasReconstruction/17.2.14/InstallArea/jobOptions/RecExCommon/RecExCommon_topOptions.py", line 97, in <module>
    include ( "RecExCond/RecExCommon_flags.py" )
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-opt/17.2.14/AtlasEvent/17.2.14/InstallArea/jobOptions/RecExCond/RecExCommon_flags.py", line
235, in <module>
    ConfigureFromListOfKeys(rec.AutoConfiguration())
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-opt/17.2.14/AtlasEvent/17.2.14/InstallArea/python/RecExConfig/AutoConfiguration.py", line 909,
in ConfigureFromListOfKeys
    rec.projectName=GetProjectName()
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-opt/17.2.14/AtlasEvent/17.2.14/InstallArea/python/RecExConfig/AutoConfiguration.py", line 294,
in GetProjectName
    from RecExConfig.GetCool import cool
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-opt/17.2.14/AtlasEvent/17.2.14/InstallArea/python/RecExConfig/GetCool.py", line 32, in <module>
    cool=_setup()
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-opt/17.2.14/AtlasEvent/17.2.14/InstallArea/python/RecExConfig/GetCool.py", line 19, in _setup
    lb=GetLBNumber()
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-opt/17.2.14/AtlasEvent/17.2.14/InstallArea/python/RecExConfig/AutoConfiguration.py", line 69,
in GetLBNumber
    if inputFileSummary['nentries']==0:
KeyError: 'nentries'
Py:Athena            INFO leaving with code 8: "an unknown exception occurred"
Traceback (most recent call last):
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-opt/17.2.14/AtlasEvent/17.2.14/InstallArea/python/RecExConfig/InputFilePeeker.py", line 49, in
_setup
    fi = athFile.fopen(inFile)
  File "<string>", line 2, in fopen
  File "/cvmfs/atlas.cern.ch/repo/sw/software/i686-slc5-gcc43-opt/17.2.14/AtlasProduction/17.2.14.12/InstallArea/python/PyUtils/AthFile/__init__.py", line
33, in _update_cache
    res = fct(*args)
  File "<string>", line 2, in fopen
  File "/cvmfs/atlas.cern.ch/repo/tools/slc6/cmt/InstallArea/python/PyCmt/Decorators.py", line 109, in forking
    status, result = pickle.load(f)
EOFError
```

Figure 16 — Example of «payload.stdout» job log

USE-CASE DIAGRAM FOR «PHYSISIST»

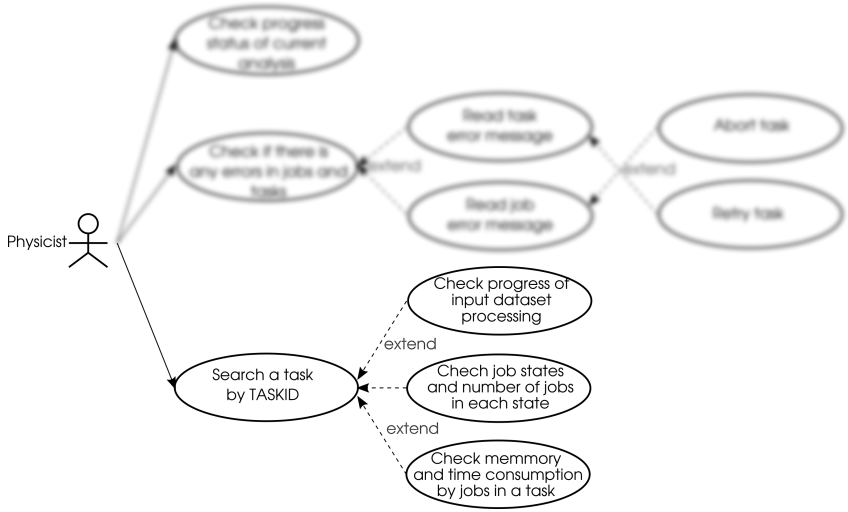

Figure 17 — Use-case Diagram for «Physisist»

TASK PAGE EXAMPLE

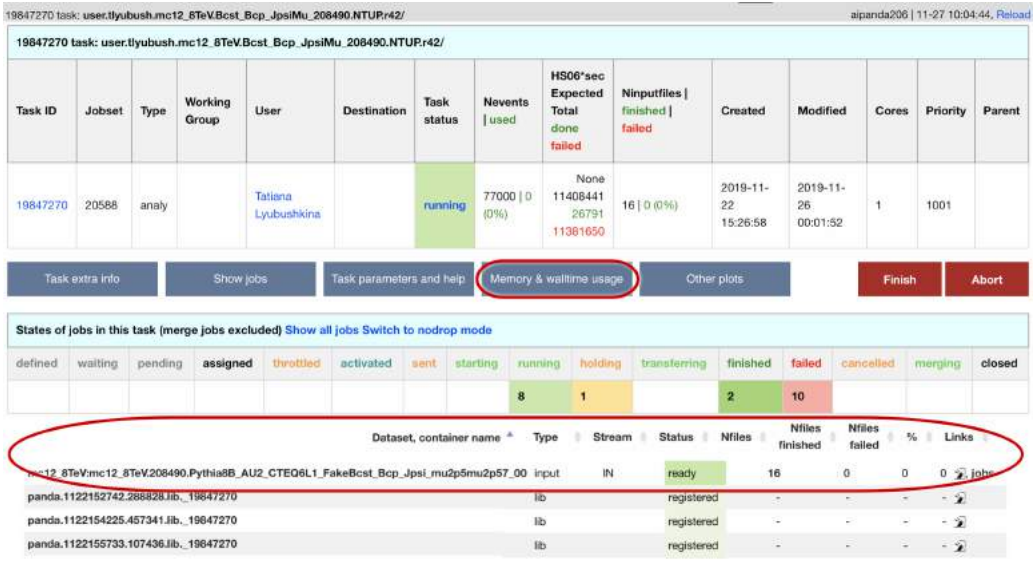

Figure 18 — Task page example showing a way to access information of jobs memory consumption

TASK PAGE EXAMPLE

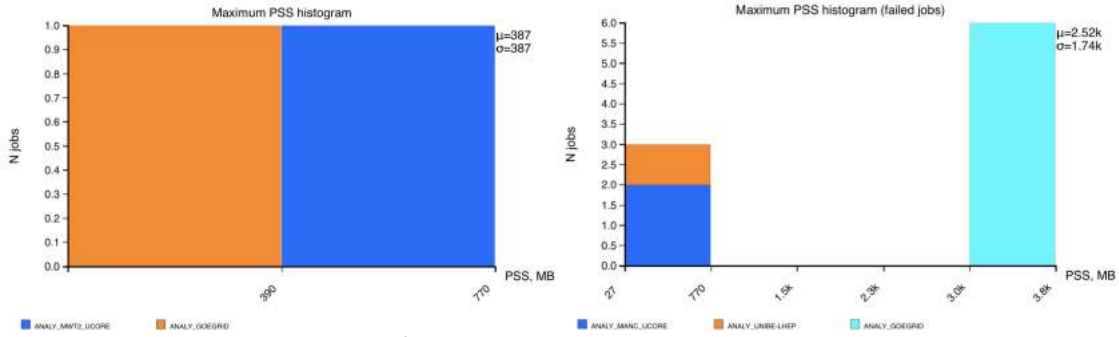

Figure 19 — Part of task page presenting jobs memory consumption histograms

JOB MEMORY CONSUMPTION EXAMPLE

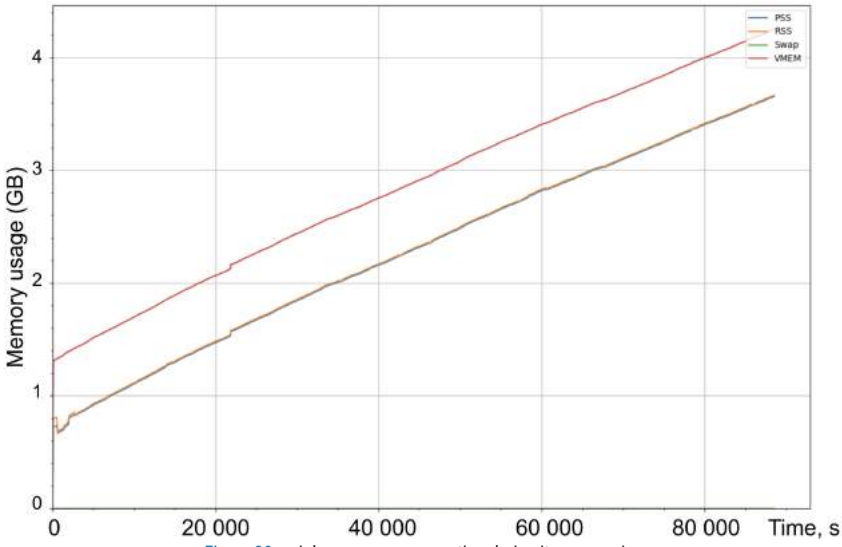

Figure 20 — Job memory consumption during its processing

OUTLINE

JOBS PAGE EXAMPLE

PanDA jobs , last 12 hours. **Params: jobtype**=production **hours**=12 **Total jobs found ~** 590000 aipanda203 | 08:10:23, Reload

576357 jobs in this selection

Job type: production

Job attribute summary, sort by count, alpha	
actualcorecount (18)	1 (378917) 2 (70) 3 (45) 4 (7825) 5 (37) 6 (6575) 7 (112) 8 (171174) 12 (1898) 16 (853) 32 (3038) 36 (2) 48 (235) 64 (127) 72 (792) 80 (56) ... more
atlasrelease (61)	Atlas-19.2.3 (1841) Atlas-19.2.4 (3) Atlas-19.2.5 (16125) Atlas-20.7.5 (21519) Atlas-20.7.7 (1677) Atlas-20.7.8 (1414) Atlas-20.7.9 (54776) Atlas-21.0.102 (5) Atlas-21.0.103 (49119) Atlas-21.0.15 (86478) Atlas-21.0.16 (2654) Atlas-21.0.19 (58) Atlas-21.0.20 (7368) Atlas-21.0.31 (70651) Atlas-21.0.40 ... more
attemptnr (29)	1 (409156) 2 (71991) 3 (38932) 4 (30058) 5 (13133) 6 (2665) 7 (839) 8 (434) 9 (350) 10 (280) 11 (105) 12 (123) 13 (78) 14 (30) 15 (33) 16 (93) ... more
gshare (14)	Data Derivations (50381) Event Index (298) Express (6208) HLT Reprocessing (1043) MC 16 (128838) MC 16 evgen (133626) MC 16 simul (102209) MC Derivations (38835) MC merge (13662) MC Other (17098) Reprocessing default (49073) Special (25015) Test (7879) Validation (2192)
harvesterinstance (9)	CERN-dev (137) CERN_central_0 (36048) CERN_central_1 (4018) CERN_central_A (133028) CERN_central_ACTA (141164) CERN_central_B (87430) cern_cloud (129) harvester_k8s (552) NERSC_test (4043)
jobstatus (15)	activated (116831) assigned (14641) cancelled (34) closed (40997) closed:toreassign (122) defined (38) failed (34516) finished (209777) holding (1755) merging (4749) pending (177) running (93134) sent (2257) starting (11521) transferring (45808)
priorityrange (13)	0:99 (23765) 100:199 (69682) 200:299 (99562) 300:399 (141882) 400:499 (69397) 500:599 (26680) 600:699 (861) 800:899 (114564) 900:999 (17363) 1000:1099 (1240) 1100:1199 (1045) 5000:5099 (2677) 10000:10099 (6529)

Figure 21 — Production jobs attribute summary for last 8 hours

RUNNING PRODUCTION TASKS PAGE EXAMPLE (1/2)

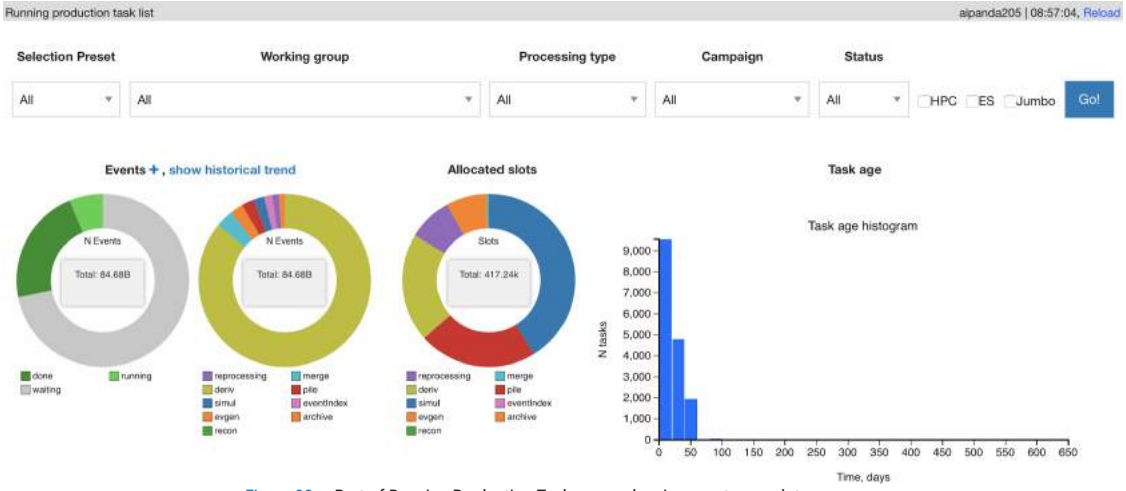

Figure 22 — Part of Running Production Tasks page showing events ans slots summary

RUNNING PRODUCTION TASKS PAGE EXAMPLE (2/2)

Show 20 entries

Search:

16403 tasks listed

Campaign	Req ID	Run	Task ID	RJobs	Status (PS2)	Status (JEDI)	Type	Evts total	Evts done	Evts running	Evts to be used	%	Failed files	Priority	Age
data18	20176	00363198	19843359	2	running	running	eventIndex	801,715,279	799,831,690	1,883,589	1,883,589	99.8	0	900	11
MC16e	26997	411287	19291746	2	running	running	evgen	210,000,000	207,000,000	1,000,000	3,000,000	98.6	400	900	61.4
MC16a	27851	411346	19772769	0	running	running	evgen	200,000,000	199,982,000	2,000	18,000	100.0	16	900	18.9
[illegible]	[illegible]	[illegible]	[illegible]	0	running	pending	evgen	199,861,000	65,038,000	2,100,000	134,812,000	32.5	1050	580	18.8
mc16	24224	366001	19917830	0	running	running	archive	145,297,000	85,421,000	3,239,000	59,876,000	58.8	0	888	1.9
mc16	24224	301000	19923238	2	running	running	archive	139,665,000	139,125,000	540,000	540,000	99.6	0	900	0.9
data16_13TeV	28235	00302872	19924654	0	running	running	deriv	131,571,245	549,102	35,687,960	131,022,143	0.4	0	880	0.7
mc16	24224	366001	19917828	32	running	running	archive	108,999,000	58,387,000	18,955,000	50,612,000	53.6	0	888	1.9
data16_13TeV	28235	00302393	19924644	0	running	running	deriv	106,733,895	761,555	7,202,583	105,972,340	0.7	0	880	0.7

Showing 1 to 20 of 16,403 entries

Previous 1 2 3 4 5 ... 821 Next

Figure 23 — Part of Running Production Tasks page showing list of tasks

OUTLINE

ADAPTATION OF BIGPANDA MONITORING SYSTEM FOR COMPASS

COMPASS experiment

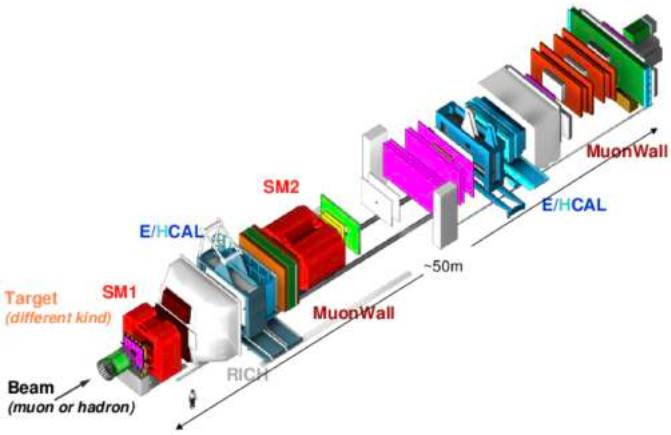

Figure 24 — COMPASS detector scheme

- **CO**mmon **M**uon **P**roton **A**pparatus for **S**tructure and **S**pectroscopy;
- The purpose of the experiment is the study of hadron structure and hadron spectroscopy with high intensity muon and hadron beams;
- First data taking run started in summer 2002 and sessions continue;
- Each data taking session containing from 1.5 to 3 PB of data;
- More than 200 physicists from 13 countries and 25 institutes are the analysis user community of COMPASS.

ADAPTATION OF BIGPANDA MONITORING SYSTEM FOR COMPASS

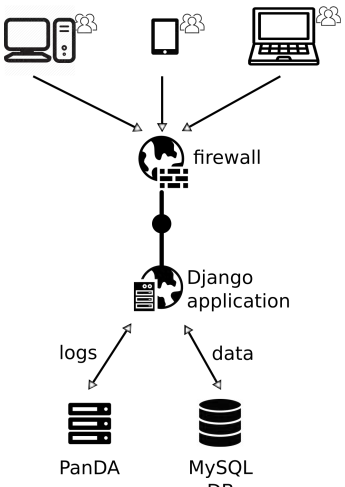

Figure 25 — Architecture of COMPASS PanDA monitoring system instance

- a new instance at JINR, Dubna;
- core views (jobs, files, dash, errors, sites) needed minor changes;
- tasks module was adapted to COMPASS production system;
- a completely new module was designed and implemented on the top of existed architecture;
- it is in production since November 2017 serving 300 requests/day.

Figure 26 — Tasks page example of COMPASS PanDA monitoring

PRODUCTION SUMMARY PAGE IN COMPASS PANDA MONITORING

Run	Number of events	Total number of chunks	Staged	Sent	Running	Failed	Finished	Status of mDST merging	X-checked mDST	mDST migration	Status of histogram merging	Histogram migration	Status of event dump merging	X-checked event dump	Event dump migration
280632	13086935	836	-	-	-	-	836	finished	yes	sent	finished	finished	finished	yes	finished
280630	12585657	841	-	-	-	-	841	finished	yes	finished	finished	finished	finished	yes	finished
280631	9549638	829	-	-	-	-	829	finished	yes	finished	finished	finished	finished	yes	finished
280633	2006458	124	-	-	-	-	124	finished	yes	sent	finished	finished	finished	yes	finished

Figure 27 — Production summary page example

OUTLINE

BIGPANDA MONITORING SYSTEM STRUCTURE

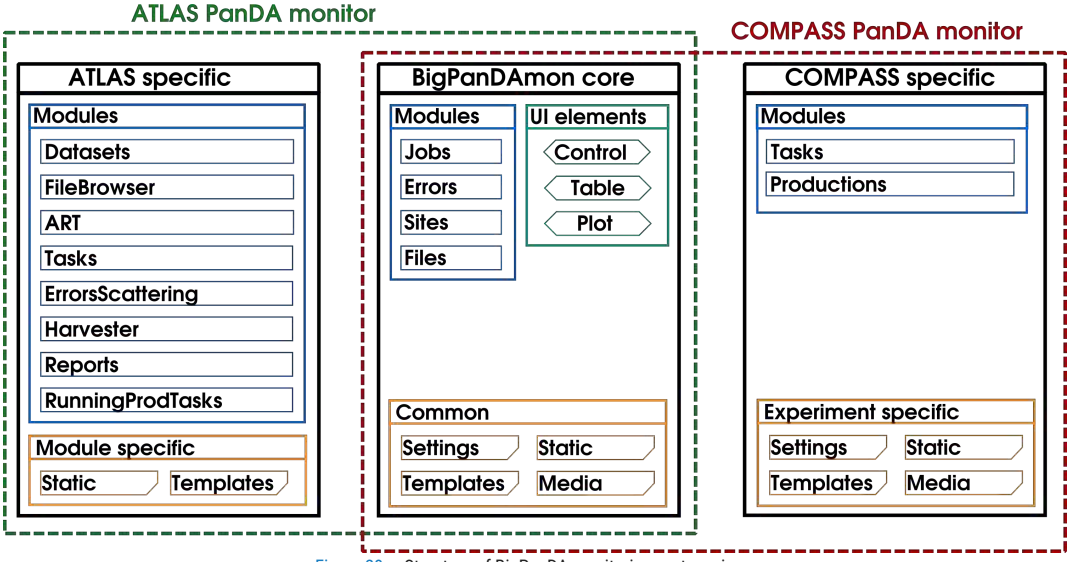

Figure 28 – Structure of BigPanDA monitoring system views

ATLAS BIGPANDA MONITORING SYSTEM USAGE STATISTICS

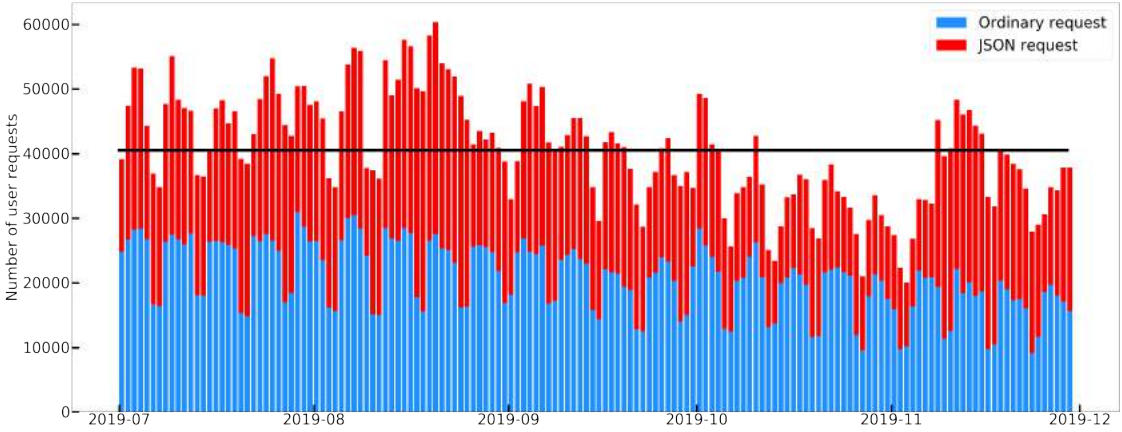

Figure 29 — Number of daily user requests, from July of 2019 to December of 2019

ATLAS BIGPANDA MONITORING SYSTEM USAGE STATISTICS

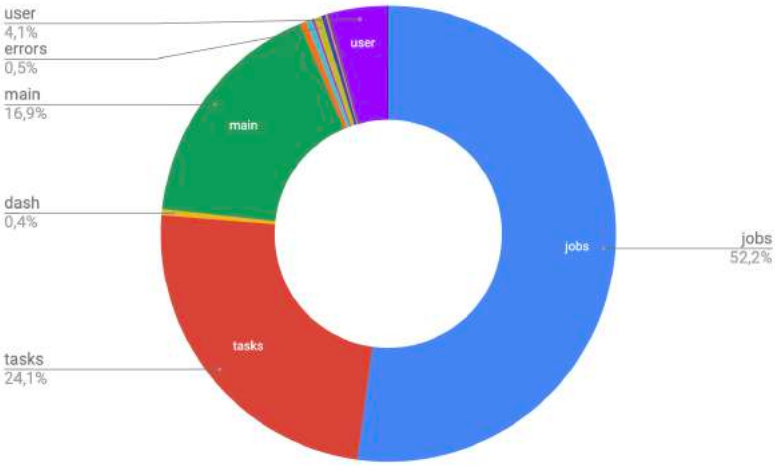

Figure 30 — Requests

- from 1 to 626 pages a day per user;
- 342 unique daily users;
- 1110 unique monthly active users.

RESULTS

1. ATLAS PanDA monitoring system serving needs of different groups of users to monitor and troubleshooting of up to $2 \cdot 10^6$ jobs distributed among 170 sites including Grid, Clouds, HPCs, and volunteer computing resources;
2. For 5 years in production, the system has only an average of 183 minutes of downtime per year (lasting more than 10 minutes) which proves a high availability — $99,965\%$;
3. Developed architecture allows to effectively serve $40 \cdot 10^3$ of daily requests in average;
4. The core of PanDA monitoring system has been successfully integrated into COMPASS experiment at SPS and provide various views for monitoring up to $1 \cdot 10^5$ jobs per day.

Acknowledgements:
This work was partially funded by the Russian Science Foundation under contract No.19-71-30008 (research is conducted in Plekhanov Russian University of Economics).

?

CERN

Large Hadron Collider (LHC)

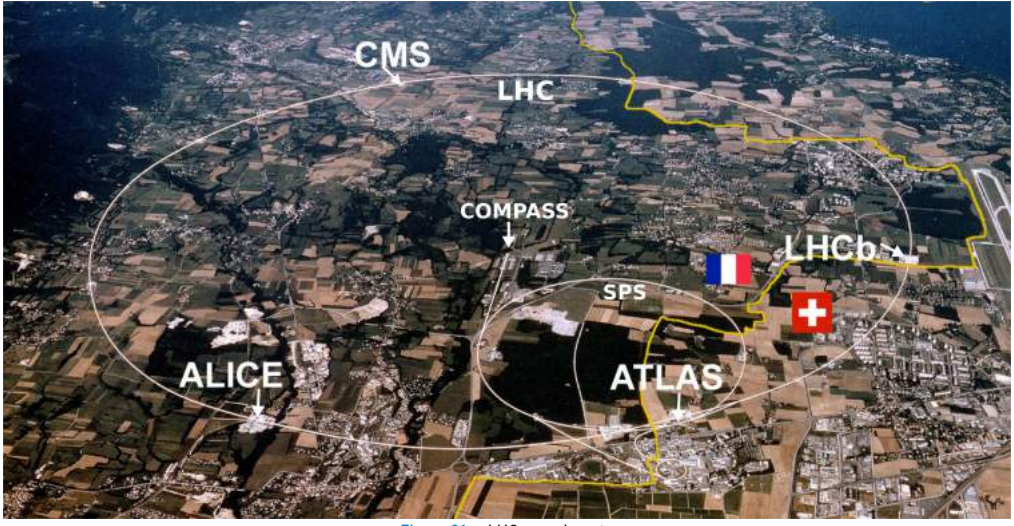

Figure 31 — LHC experiments

ATLAS EXPERIMENT

Collaboration

- >5 000 physicists, students, engineers, technicians
- >1 000 PhD students

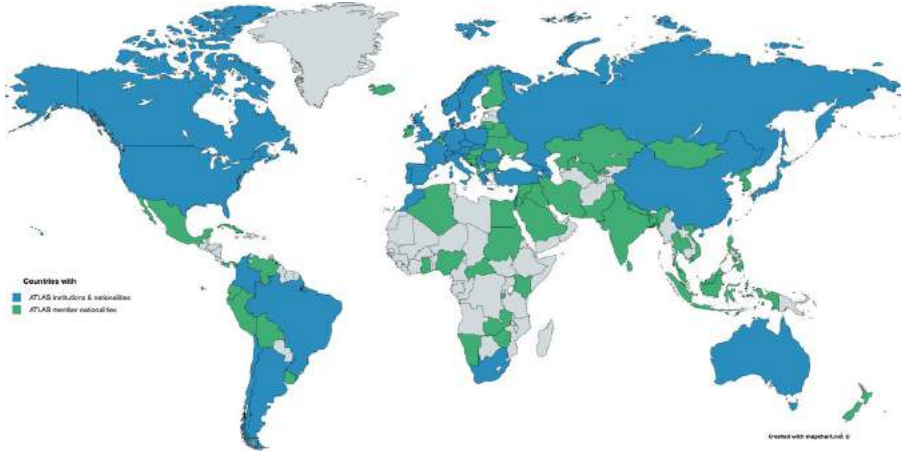

Figure 32 — ATLAS Collaboration members nationalities

ATLAS EXPERIMENT

Detector

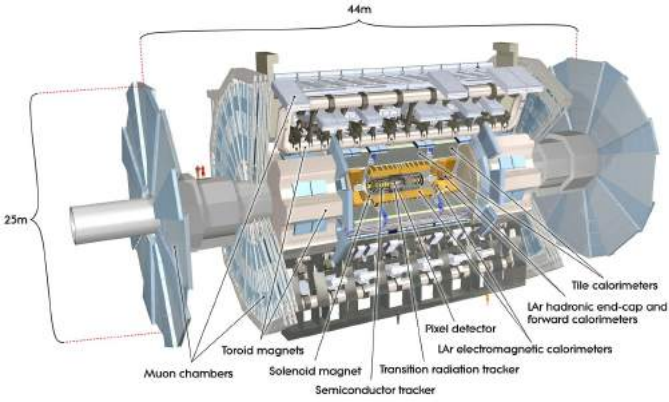

Figure 33 – Scheme of ATLAS detector

- 25×44 m., 7 000 tonnes
- $\sim$ 101 million of electronic channels
- $\sim$ 1 billion of charged particle bunch collisions per second
- trigger system filter the most interesting events in real time, which decreases data flow from $\sim$1PB/s to $\sim$1GB/s

DATA PRODUCTION IN ATLAS EXPERIMENT

GENERATION

SIMULATION

DIGITALIZATION

RECONSTRUCTION

DERIVATION

TRIGGER

RECONSTRUCTION

DERIVATION

ANALYSIS

Online

Tier0

GRID

Local

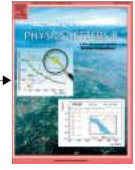

Figure 34 — Scheme of data production and analysis cycle in ATLAS experiment

WORLDWIDE LHC COMPUTING GRID (WLCG)

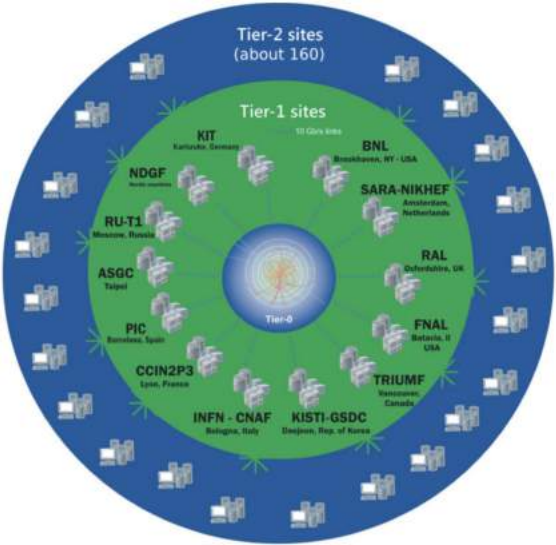

Figure 35 – WLCG scheme

WLCG: an International collaboration to distribute and analyse LHC data

Integrates computer centres worldwide that provide computing and storage resource into a single infrastructure accessible by all LHC physicists.

- Tier-0: data recording, reconstruction and distribution;
- Tier-1: permanent storage, re-processing, analysis;
- Tier-2: Simulation, end-user analysis;
- $\sim$ 170 sites, 42 countries;
- $\sim$ 1 M CPU cores;
- $\sim$ 1 EB of storage;
- > 2 million jobs/day;
- 10-100 Gb links.

WLCG INFRASTRUCTURE

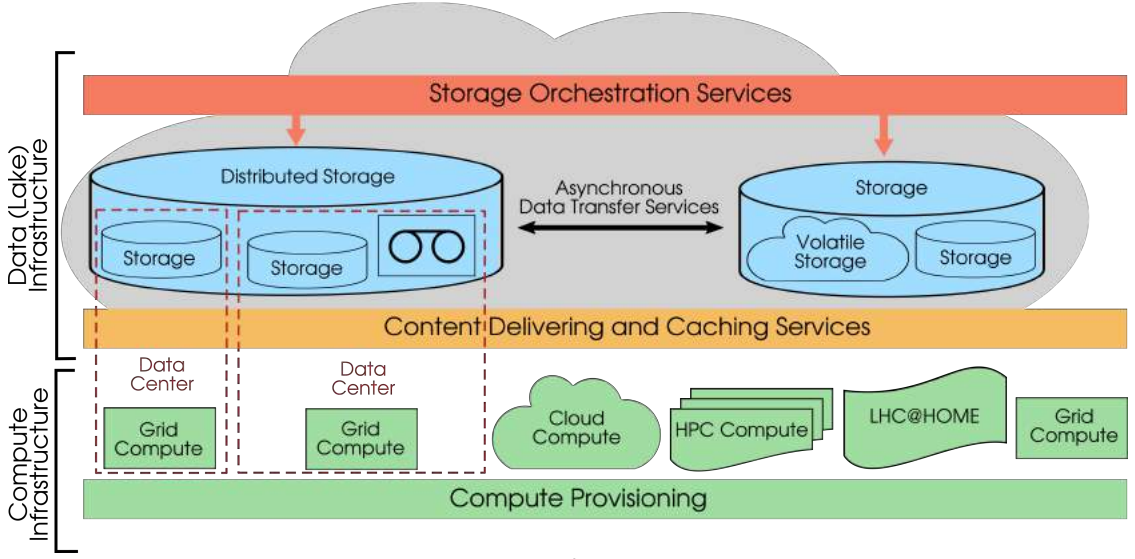

Figure 36 — WLCG Infrastructure scheme