

PanDA FOR COMPASS AT JINR

*A. Sh. Petrosyan*¹

Joint Institute for Nuclear Research, Dubna

PanDA (Production and Distributed Analysis system) is a workload management system, widely used for data processing in experiments at Large Hadron Collider and others. COMPASS is a high-energy physics experiment at the Super Proton Synchrotron. Data processing for COMPASS runs locally at CERN, on lxbatch, the data itself stored in CASTOR. In 2014 an idea to start running COMPASS production through PanDA arose. Such a transformation in experiment's data processing will allow COMPASS community to use not only CERN resources, but also Grid resources worldwide. During the spring and summer of 2015 installation, validation and migration work was performed at JINR. Details and results of this process are presented in this paper.

PACS: 89.20.Ff

COMPASS

COMPASS (Common Muon Proton Apparatus for Structure and Spectroscopy) is a high-energy physics experiment at a Super Proton Synchrotron at CERN [1]. The purpose of the experiment is to study hadron structure and spectroscopy with high-intensity muon and hadron beams. First data taking run of COMPASS started in summer 2002 and the sessions continue, each data taking session contains from 1.5 to 3 PB of data. More than 200 physicists from 13 countries and 24 institutes are the analysis user community of COMPASS.

COMPASS data taking started when mechanisms of distributed computing, such as Grid [2], were not available yet. That is why computing model of the experiment is very conservative, centralized and relies on CERN IT services, such as lxplus, Castor, lxbatch [3–5]. Later, EOS was added to the data flow chain [6].

Below there is a short and simplified description of implemented data flow, also presented in Fig. 1:

- all data are stored on Castor;
- data are being requested to be copied from tapes to discs before processing and then, once the data are ready, file moves to lxbatch for processing;
- the processing is done by the specific software [7];
- after processing the results are being transferred either to EOS for merging or short-term storage or directly to Castor for long-term storage.

¹E-mail: artem.petrosyan@jinr.ru

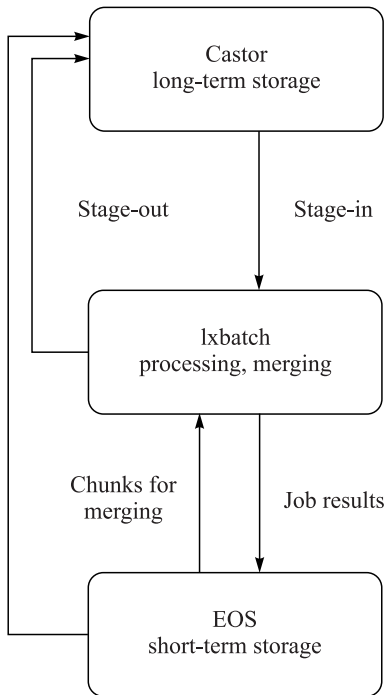


Fig. 1. COMPASS data flow

Current implementation of data-processing data flow has several limitations:

- data management is done by a set of scripts, deployed under production account on AFS, and not organized as a commonly run system;
- execution of user analysis jobs and production jobs are separated and managed by different sets of software;
- number of jobs which can be executed by collaboration at lxbatch is limited;
- available space on home of COMPASS production user at lxbatch and Castor is limited and strictly managed;
- Castor has a storage system which stores data on tapes and is not definitely designed for random access reading from many users simultaneously;
- although COMPASS data flow has conditions for distributed computing, due to historical premises it is implemented as a single-site processing which uses only one computing facility;
- absence of monitoring does not allow one to track how and by whom the data are being used.

CERN IT constantly decreases access to LSF and Castor and forces experiments to use more intensively available Grid infrastructure built to serve experiments at LHC. In the situation of limitations shown above, a

decision was made to evaluate one of the existing distributed analysis tools. The natural choice of such a tool is PanDA [8], production and workload management system, originally implemented to handle needs of the ATLAS collaboration.

PANDA

The PanDA (Production and Distributed Analysis System) has been developed by ATLAS to meet requirements of data-driven workload management system for production and distributed analysis processing capable at LHC data processing scale. PanDA manages both user analysis and production jobs via the same interface with processing rate from 250 to 300 K jobs on more than 170 sites every day. The PanDA ATLAS analysis user community numbers over 1400 physicists. PanDA data flow is presented in Fig. 2.

PanDA started in 2005 from World Wide LHC Computing Grid and ATLAS. In 2012, BigPanDA project was started. The main goal of the project was to bring PanDA outside ATLAS. The efforts during the BigPanDA project were concentrated on making PanDA modules distributable, adopted to work with several backends, on HPCs, on different cloud platforms, to allow several virtual organizations to use one physical instance, better use of network infrastructure, and to be instrumented by a better monitoring capabilities. Currently, PanDA supports various resources such as Amazon EC2 and Google cloud platform, Hyper Power Computers like world's number two TITAN, MIRA, and many others. Process of adopting PanDA to serve the distributed computing of various scientific projects is ongoing.

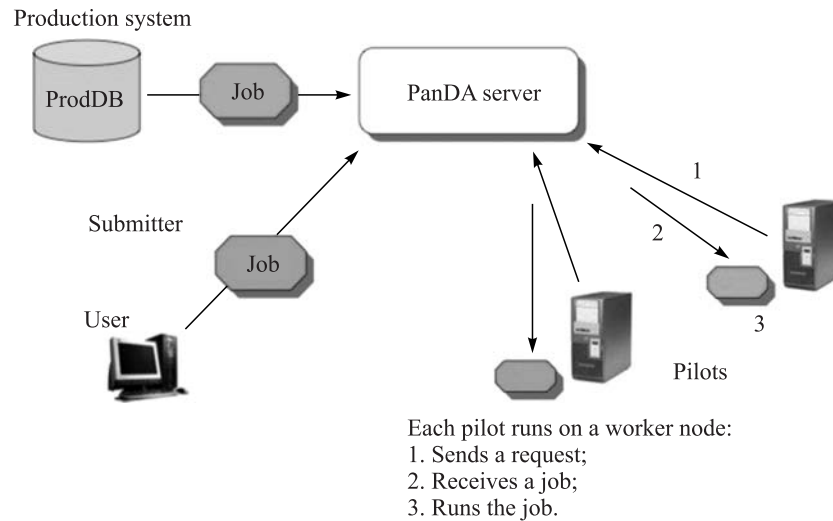


Fig. 2. PanDA simplified data flow

PANDA INSTANCE AT JINR

The goal of the project was to evaluate the possibility of COMPASS jobs to be executed by any COMPASS user via PanDA server installed at JINR on local cloud service. So, the main goal was split into the following tasks:

- evaluate PanDA for MySQL instance installation;
- evaluate PanDA monitoring application installation;
- extend PanDA server&pilot by modules to introduce COMPASS logic;
- evaluate whether COMPASS software is portable and can be executed outside CERN;
- test that COMPASS payload submission via PanDA;
- prepare and publish a wrapper to define COMPASS job via PanDA client;
- make sure that an ordinary COMPASS user can submit jobs via PanDA.

All these tasks described above were finished successfully. Three PanDA queues were defined: at JINR on local cloud service, at CERN on lxplus, and at CERN on lxbatch. Several jobs were sent to each queue. During the evaluation more than a thousand real COMPASS jobs were defined and processed by PanDA. Successful processing of the jobs was confirmed by COMPASS software coordinators. User's howto documentation was prepared and evaluated by production manager [9]. Several jobs were defined and sent by production manager of the experiment. Details of every job were presented on PanDA monitoring instance, installed at JINR [10].

CONCLUSIONS

Evaluation of PanDA for COMPASS was finished successfully. PanDA confirmed itself as an extendable system which may be adopted to run COMPASS payload. PanDA instance setup was evaluated.

The project is not finished, general direction of the development is concentrated on preparation of the production environment and on migration from the currently existing data flow management to PanDA. The following steps are being performed:

- install and configure Auto Pilot Factory (APF) to enable an automatic Pilot submission to several sites;
- implement a distributed data management plug-ins to open Grid for COMPASS: a common stage-in, stage-out and management of data on several sites through Grid infrastructure;
- migrate from cloud service to real machines, to avoid possible scalability and reliability problems which may arise in the future. Both JINR and CERN have all necessary capabilities to host PanDA instance;
- connect JINR as a computing site for COMPASS through PanDA, i.e., allocate the space and define PanDA queue at JINR's computing infrastructure;
- define PanDA queues in other participating in COMPASS institutes and make their data analysis distributed.

REFERENCES

1. COMPASS Experiment. <http://wwwcompass.cern.ch/>.
2. Worldwide LHC Computing Grid. <http://wlcg.web.cern.ch/>.
3. CERN AFS Service. <http://information-technology.web.cern.ch/services/afs-service>.
4. CERN Batch Service. <http://information-technology.web.cern.ch/services/batch>.
5. CERN Advanced STORage Manager (Castor). <http://castor.web.cern.ch/>.
6. CERN EOS Service. <http://information-technology.web.cern.ch/services/eos-service>.
7. COMPASS Software.
<https://wwwcompass.cern.ch/compass/software/offline/welcome.html#software>.
8. PanDA&BigPanDA. <https://twiki.cern.ch/twiki/bin/view/PanDA/PanDA/>
9. COMPASS User HowTo Twiki.
<https://twiki.cern.ch/twiki/bin/view/PanDA/COMPASSUserHowTo>.
10. Monitoring Link to PanDA COMPASS Queues at JINR.
http://vm127.jinr.ru/bigpandamon/jobjobs/?computingsite=ANALY_CERN_COMPASS_PROD.