

# BigPanDA monitoring system evolution in the ATLAS Experiment

Tatiana Korchuganova<sup>1,\*</sup>, Aleksandr Alekseev<sup>2</sup>, Alexei Klimentov<sup>3</sup>, Torre Wenaus<sup>3</sup>, and Zhaoyu Yang<sup>3</sup> on behalf of the ATLAS Computing Activity

<sup>1</sup>University of Pittsburgh, Pittsburgh, PA, USA

<sup>2</sup>University of Texas at Arlington, Arlington, TX, USA

<sup>3</sup>Brookhaven National Laboratory, Upton, NY, USA

**Abstract.** Monitoring services play a crucial role in the day-to-day operation of distributed computing systems. The ATLAS Experiment at LHC uses the Production and Distributed Analysis workload management system (PanDA WMS), which allows a million computational jobs to run daily at over 170 computing centers of the WLCG and opportunistic resources, utilizing 600k cores simultaneously on average. The BigPanDA monitor is an essential part of the monitoring infrastructure for the ATLAS Experiment that provides a wide range of views, from top-level summaries to a single computational job and its logs. Over the past few years of the PanDA WMS advancement in the ATLAS Experiment, several new components were developed, such as Harvester, iDDS, Data Carousel, and Global Shares. Due to its modular architecture, the BigPanDA monitor naturally grew into a platform where the relevant data from all PanDA WMS components and accompanying services are accumulated and displayed in the form of interactive charts and tables. Moreover the system has been adopted by other experiments beyond HEP. In this paper we describe the evolution of the BigPanDA monitor system, the development of new modules, and the integration process into other experiments.

## 1 Introduction

BigPanDA monitor is the next-generation monitoring system of the ATLAS PanDA Workload Management System (WMS) [1], which was developed in the ATLAS Experiment [2] and brought into production in 2014. It has replaced the technologically outdated PanDA monitor web application [3]. Since then, it has been continuously improved in line with the advancement of the PanDA WMS to meet new challenges brought by the ATLAS collaboration, in particular the integration of HPCs and commercial clouds, the central control of computing activity distributed among available resources, and orchestration of the data movement between tape storage and disks. Also, the success of the PanDA WMS brought interest from other experiments in HEP and beyond. This implied the need for changes in the system development to become experiment-agnostic.

---

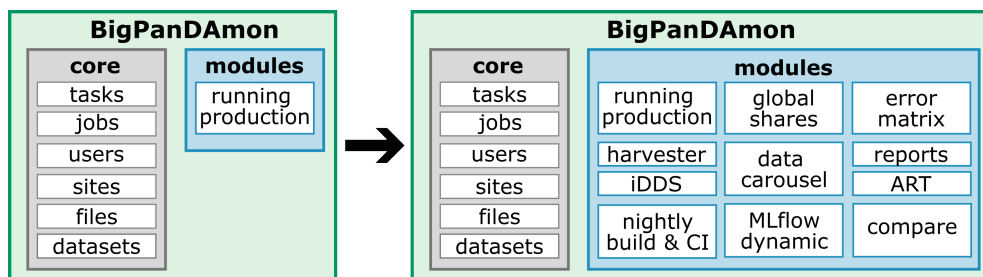
\*e-mail: tatiana.korchuganova@cern.ch



## 2 BigPanDA monitor evolution

BigPanDA monitor is a Django-based web application that collects data from various sources, such as a relational DB (Oracle, PostgreSQL, MySQL), Elasticsearch storage, APIs and prompt clients of other systems (CRIC information system [4], Rucio distributed data management system [5]), aggregates and represents it in the form of interactive tables and charts.

Initially, the project comprised a set of core views covering the basic needs of monitoring key aspects of the PanDA WMS, in particular jobs, tasks, files, datasets, computing sites, and users. With the advancement of the PanDA WMS [1] in the ATLAS Experiment, several new components have been developed, such as Harvester [6], iDDS [7], Data Carousel [8], and Global Shares [9]. All of them are closely related to key WMS objects and have to be monitored in an integrated way. Naturally, BigPanDA monitor started to grow and fulfil new monitoring needs. Also, the interest of other experiments brought a need for the ability to plug in/out components depending on the unique requirements of each experiment. This triggered the transformation of the BigPanDA monitor architecture into a modular platform. From 2017 to 2023, the number of implemented modules increased from 1 to 11 (Figure 1). To perform this transition, we generalized and separated common functions, so they could be



**Figure 1.** BigPanDA monitor structure evolution from 2017 to 2023.

reused across modules. The system was updated by:

- translating request GET/POST parameters into queries for Django’s built-in Object Relational Mapper (ORM) ;
- introducing caching of the prepared-to-be-rendered data, which allows adding user-specific data on top of it for subsequent requests;
- using a common base HTML template and nested blocks of the built-in Django template-rendering engine;
- improving configuration of the application to handle plug in/out of modules;
- and a testing mechanism by creating unit-tests for most of the views.

These improvements make it possible to create new modules and views efficiently and faster because developers only need to work on data querying, aggregation, and visualization with all tools at hand. The next section describes the implementation of several key modules.

## 3 Implementation of new modules

### 3.1 Harvester monitoring module

The Harvester [6] is a resource-facing service between the PanDA WMS and pilots for resource provisioning and workload shaping. It has two entities: a Harvester instance that can

run centrally with the required plugins activated for different resources or on edge nodes of HPC centres, and a Harvester worker that can be a pilot, VM or MPI job. The Harvester worker has a one-to-many relation with the PanDA job.

The requirements for monitoring the Harvester components were: to present a list of all Harvester instances and their states; to be able to search workers per Harvester instance, computing site or computing element; to find a worker for a PanDA job and, vice versa, PanDA job(s) for a worker. To fulfil the needs listed above, we developed four views:

- Harvester instance list;
- Harvester worker summary (Figure 2):
  - workers attribute summary which allows drilling down to workers of interest;
  - list of workers with links to batch logs;
  - list of associated PanDA jobs;
  - workers statistics;
  - diagnostic messages to identify the internal problems of a Harvester instance;
- Harvester worker info page accumulates all information related to the selected worker. It can be found by the worker ID or PanDA job ID;
- Harvester slots page represents the number of slots for special resources (HPCs) working in push mode;

We faced a performance issue during the development of the Harvester worker summary view. In the ATLAS Experiment, the average number of daily submitted workers is 600K and therefore the time to query and aggregate the data for the last day could reach one minute, which is unacceptable for most users. To mitigate this problem, it was decided to use Elasticsearch, Logstash and Kibana (ELK) stack for global monitoring of Harvester. The data is copied every 10 minutes from the PanDA DB to the central Elasticsearch storage provided by CERN IT [10]. The implemented views of the Harvester monitoring module are mostly used for debugging immediate problems related to a particular worker or PanDA job because BigPanDA monitor reads data directly from the PanDA database without additional delays.

### 3.2 Global Shares monitoring module

The Global Shares component [9] of the PanDA WMS is responsible for optimizing the distribution of workloads across the available computing resources, depending on current priorities of computing activities set by the ATLAS collaboration. The shares of computing resources may have a nested structure; for example, on the top level, it divides into production and analysis, then production divides into event generation, MC simulation, reconstruction, and derivation activities. The current setup has three levels. The main requirement for the monitoring of this component was the ability to see how close the actual distribution of computing resources is to the target for every share on any level. In addition, it must provide an insight into which type of resources (among the Grid, HPC, Cloud and GPU) contributes the most to each share.

We implemented a dedicated dashboard view to meet these requirements, which consists of three blocks:

- an overview with plots where the level of Global Shares and type of resources can be selected (Figure 3);
- a table with the list of global shares representing the target and actual values linked to the corresponding PanDA job list;

| Harvester worker attributes summary |   |
|-------------------------------------|---|
| <b>computingelement (257)</b>       | gate03.aglt2.org:9619 (1145) gridgk07.racl.bnl.gov:9619 (1370) calc2.t1.grid.kiae.ru:443 (140) gridgk02.racl.bnl.gov:9619 (1338) gridgk06.racl.bnl.gov:9619 (1549) gate01.aglt2.org:9619 (1804) ce02.tier2.hep.manchester.ac.uk (507) ce03.tier2.hep.manchester.ac.uk (273) lcge1.shef...<br><a href="#">more</a> |
| <b>computingsite (185)</b>          | AGLT2 (2780) BNL (8245) RRC-KI-T1 (350) UKI-NORTHGRID-MAN-HEP (1397) UKI-NORTHGRID-SHEF-HEP (129) UKI-SOUTHGRID-BHAM-HEP (57) SWT2_LANCIUM (114) GoeGrid_LODISK (37) GOOGLE_EUW1 (2780) MWT2_K8S_UCORE (86) RIVR.UM (108) ...<br><a href="#">more</a>   |
| <b>harvesterid (13)</b>             | CERN_central_A (39220) CERN_central_B (27801) test_aipanda185 (114) CERN_central_k8s (5535) CERN_central_ACTA (22654) CERN-dcv (183) SuperMUC (122) cern_cloud (908) CERN_central_1 (1546) CERN_central_0 (8546) tacc-frontera (9) CERN_central_dask (3) NERSC_perlmutter_test (1)                                |
| <b>jobtype (8)</b>                  | managed (55627) user (46111) prod_test (606) ANY (3061) panda (944) DUMMY (28) rc_alrb (142) rc_test2 (123)   |
| <b>nativestatus (15)</b>            | running (43002) completed (13555) idle (19712) removed (552) (13488) starting (4981) done (7152) unknown (64) donefailed (691) donecancelled (191) held (882) tovalidate (15) transferring (1) sent (2354) finished (2)   |
| <b>resourcetype (4)</b>             | MCORE (40253) SCORE_HIMEM (6749) SCORE (59632) MCORE_HIMEM (8)  |
| <b>status (7)</b>                   | running (48416) finished (22670) submitted (33889) cancelled (1192) missed (117) failed (279) idle (79)   |

| Workers                         | Jobs      | Worker Statistics      | Dialog Messages              |           |             |               |                           |                     |                         |
|---------------------------------|-----------|------------------------|------------------------------|-----------|-------------|---------------|---------------------------|---------------------|-------------------------|
| Number of workers               | 1000      | <a href="#">Reload</a> |                              |           |             |               |                           |                     |                         |
| Show                            | 10        | entries                | Search: <input type="text"/> |           |             |               |                           |                     |                         |
| Workers submitted by harvesters |           |                        |                              |           |             |               |                           |                     |                         |
| Instance                        | Worker ID | Jobs                   | Last update                  | Status    | Batch ID    | Computingsite | Computing element         | Submit time         | Start time              |
| CERN_central_A                  | 463520836 | 1                      | 2023-04-04 02:10:49          | running   | 15428467.34 | AGLT2         | gate03.aglt2.org:9619     | 2023-04-04 11:15:02 | 2023-04-04 11:50:23     |
| CERN_central_A                  | 463520835 | 1                      | 2023-04-04 02:10:49          | running   | 4356492.16  | AGLT2         | gate01.aglt2.org:9619     | 2023-04-04 11:15:02 | 2023-04-04 11:50:23     |
| CERN_central_A                  | 463520833 | 1                      | 2023-04-04 12:59:26          | cancelled | 13859329.0  | RRC-KI-T1     | calc2.t1.grid.kiae.ru:443 | 2023-04-04 11:15:02 | 2023-04-04 11:50:02 202 |
| CERN_central_A                  | 463520832 | 1                      | 2023-04-04 02:10:42          | running   | 4356492.15  | AGLT2         | gate01.aglt2.org:9619     | 2023-04-04 11:15:02 | 2023-04-04 11:50:22     |
| CERN_central_A                  | 463520831 | 1                      | 2023-04-04 02:00:37          | running   | 15428467.33 | AGLT2         | gate01.aglt2.org:9619     | 2023-04-04 11:15:02 | 2023-04-04 11:20:11     |
| CERN_central_A                  | 463520830 | 1                      | 2023-04-04 01:50:37          | finished  | 15428467.32 | AGLT2         | gate03.aglt2.org:9619     | 2023-04-04 11:15:02 | 2023-04-04 11:50:22 202 |
| CERN_central_A                  | 463520829 | 1                      | 2023-04-04 12:33:45          | finished  | 4356492.14  | AGLT2         | gate01.aglt2.org:9619     | 2023-04-04 11:15:02 | 2023-04-04 11:50:22 202 |

**Figure 2.** Harvester worker summary page for the most recent six-hour period.

- a set of tabs containing more detailed information, such as the global share distribution across computing sites, types of computing resources (Grid, HPC, Cloud etc.), resource types (single/multicore, ordinary/high memory), and PanDA job statuses.

### 3.3 Intelligent Data Delivery Service monitoring module

The intelligent Data Delivery Service (iDDS) [7] aims to orchestrate workflows and manage data pre-processing, delivery, and primary processing in large-scale workflows. It supports workflows defined via a Directed Acyclic Graph (DAG).

For monitoring workflows managed by the iDDS we created a workflow progress dashboard (Figure 4), which shows a list of workflows, their type and status, and the progress quantified in number of files. The iDDS brings flexibility to workflow structures, and connections between its internal steps. Therefore, one of the requirements for monitoring was to build a visualization of workflow complexity. Figure 5 shows an example of such a graph diagram, where each node is a computational task, and edges represent the flow of data. We are planning to improve this visualization by adding more information to it, in particular the status and progress of each task.




Figure 3. Global Shares dashboard.

Requests:

Show 10 entries

Search:

| request id | username | workflow status | graph | workflow name                              | created on (UTC)    | total tasks | tasks                       | transform type | total files | released files | unreleased files | finished files | failed files |
|------------|----------|-----------------|-------|--|---------------------|-------------|-----------------------------|----------------|-------------|----------------|------------------|----------------|--------------|
| 508223     | User 1   | SubFinished     | plot  | pseudo_input.2023_06_26_20_24_18_752267769 | 2023-06-26 20:24:19 | 10          | Finished(6)<br>Failed(2)    | Processing     | -           | 0              | 0                | 95.7%          | 4.3%         |
| 489653     | User 2   | SubFinished     | plot  | pseudo_input.2023_05_15_20_22_16_527614670 | 2023-05-15 20:22:17 | 10          | Failed(6)<br>SubFinished(4) | Processing     | -           | 0              | 0                | 50%            | 50%          |
| 489639     | User 3   | SubFinished     | plot  | pseudo_input.2023_05_14_14_56_16_550921349 | 2023-05-14 14:56:17 | 10          | Finished(4)<br>Failed(6)    | Processing     | -           | 0              | 0                | 57.8%          | 4.5%         |
| 486243     | User 4   | SubFinished     | plot  | pseudo_input.2023_05_01_01_21_18_02287980  | 2023-05-01 01:21:19 | 10          | Failed(6)<br>SubFinished(4) | Processing     | -           | 0              | 0                | 50%            | 50%          |

Figure 4. iDDS workflow progress view.




Figure 5. Graph plot of one of workflows defined via DAG. Each node is a computational task, which has a unique identification number.

### 3.4 Data Carousel monitoring module

The Data Carousel [8] orchestrates data processing between workload management, data management, and storage services with the bulk data resident on tape. The processing is executed by staging and promptly processing a sliding window of inputs onto faster buffer storage such that only a small percentage of input data are available at any one time.

The development of this module was particularly challenging due to its intermediate position between the workflow and distributed data management systems. The main source of information are the ATLAS Production System [11] DB tables, which store the datasets that need to be staged-in, the corresponding Rucio rules, and their progress. Detailed information regarding staging transfers comes from the central Elasticsearch storage filled by Rucio.

A dedicated dashboard has been developed for monitoring staging activities (Figure 6), which contains various visualizations of staging datasets/files/their volume and a table with staging requests and associated tasks that have input on tape. In addition, a notification

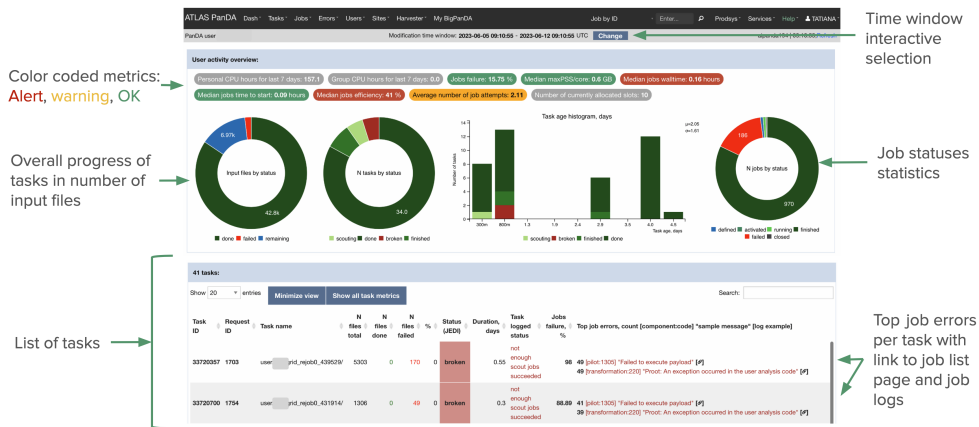


Figure 6. Data Carousel dashboard.

mechanism for stalled staging requests was implemented. It regularly searches for staging requests that have not progressed for  $N$  days (configurable, currently set to 10) and sends an email report to the list of relevant experts.

## 4 MyBigPanDA view

The development of the new modules described in Section 3 was initiated by requests from experts: WMS component developers, shifters, production managers, and system administrators of computing centres. This group of users is very active with almost half of requests to BigPanDA monitor despite the small number of people (around 100). On the other side, the number of physicists sending their analysis tasks to the system reaches 1500 people in peaks. Only 25% of improvement requests or bug reports come from analysis users, while the other 75% are from experts. Therefore, BigPanDA monitor naturally tends towards becoming a more complicated expert-focused system that is difficult to learn for experiment members performing their analysis. This necessitates the creation of a simple overview page that collects all analysis tasks submitted by a user, provides insight into task progress, and helps debugging problems. An example of this page is shown in figure 7. It consists of two blocks. The first is an overview showing statistics of user activity: the progress of the submitted tasks in the number of processed input files, the distribution of task statuses, a histogram of task age that allows slow tasks to be spotted, and colour-coded job metrics. These metrics can provide clues to possible problems. The second block is a table with a list of tasks. The



**Figure 7.** MyBigPanDA page. The red median job duration label highlights that jobs are very short (~ 10 min), which affects the overall efficiency of data processing, because typical Grid jobs are expected to run for a duration of order 24 hours. The error message in the last column of the table (“Failed to execute payload. An exception occurred in user analysis code”) clearly indicates the problem. In this case, the user can see the log with just one click.

last column of the table shows the most common errors across jobs for each task and the link to the log for one job. The MyBigPanDA view is actively used; during the last year, 1471 out of 2123 unique users visited it at least once. The number of daily visits by a single user varies from 0 to 212.

## 5 Adapting to other experiments

As of 2023 several other experiments are using the PanDA WMS for distributed data processing and BigPanDA monitor for monitoring, in particular COMPASS [12] at SPS, Vera Rubin observatory [13], and sPHENIX [14] at BNL. Each experiment has unique requirements and can choose which components of PanDA WMS to use. The main difficulty in adapting the monitoring system was the requirement to use different database vendors. ATLAS uses Oracle DB, while COMPASS requires MySQL. V. Rubin Observatory and sPHENIX use PostgreSQL. BigPanDA monitor views mostly use the Django built-in Object Relational Mapper (ORM) for querying data from databases. However, there are cases where the functionality is not enough, and for those we have added functions which change the syntax of queries depending on the database provider. Also, we have introduced new configuration variables to manage the content of views depending on the experiment. Another requirement was to run the BigPanDA monitor on a Kubernetes cluster, therefore we have created a Docker image and configured a GitHub action to build a container automatically for new releases.

## 6 Summary

The BigPanDA monitor system has significantly evolved since its first release in 2014. The architecture naturally transformed from an application into a multimodule platform which accumulates monitoring views for all components of the PanDA WMS. We are continuously working on improving system usability and transparency for ATLAS collaboration members doing their physics analysis.

## Acknowledgement

We appreciate the assistance of our colleagues from the PanDA team and ATLAS Distributed Computing group. We thank Siarhei Padolsky for a significant contribution to the development and supervising the project for many years.

## References

- [1] T. Maeno, K. De, T. Wenaus, P. Nilsson, R. Walker, A. Stradling, V. Fine, M. Potekhin, S. Panitkin, G. Compostella, *Journal of Physics: Conference Series* **396**, 032071 (2012)
- [2] ATLAS Collaboration, *J. Inst.* **3**, S08003 (2008)
- [3] A. Klimentov, P. Nevski, M. Potekhin, T. Wenaus, *Journal of Physics: Conference Series* **331**, 072058 (2011)
- [4] A. Anisenkov et al., *EPJ Web Conf.* **245**, 03032 (2020)
- [5] M. Barisits et al., *Comput Softw Big Sci* **3**, 11 (2019)
- [6] Maeno, Tadashi, Barreiro Megino, Fernando Harald, Benjamin, Doug, Cameron, David, Childers, John Taylor, De, Kaushik, De Salvo, Alessandro, Filipcic, Andrej, Hover, John, Lin, FaHui et al., *EPJ Web Conf.* **214**, 03030 (2019)
- [7] W. Guan et al., *EPJ Web Conf.* **251**, 02007 (2021)
- [8] M. Borodin et al., *EPJ Web Conf.* **251**, 02006 (2021)
- [9] Barreiro Megino, Fernando, Di Girolamo, Alessandro, De, Kaushik, Maeno, Tadashi, Walker, Rodney, *EPJ Web Conf.* **214**, 03025 (2019)
- [10] F.H.B. Megino, A. Alekseev, F. Berghaus, D. Cameron, K. De, A. Filipcic, I. Glushkov, F. Lin, T. Maeno, N. Magini, *EPJ Web Conf.* **245**, 03010 (2020)
- [11] F.H. Barreiro, M. Borodin, K. De, D. Golubkov, A. Klimentov, T. Maeno, R. Mashinistov, S. Padolski, T. Wenaus, on behalf of the ATLAS Collaboration, *Journal of Physics: Conference Series* **898**, 052016 (2017)
- [12] P. Abbon, E. Albrecht, V. Alexakhin, Y. Alexandrov, G. Alexeev, M. Alekseev, A. Amoroso, H. Angerer, V. Anosov, B. Badefek et al., *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **577**, 455 (2007)
- [13] Z. Ivezic et al., *Astrophys. J.*, **873**, 111 (2019)
- [14] A. Adare et al. (PHENIX) (2015), 1501.06197