

PanDA codebase and infrastructure upkeep

E. Karavakis, F. Barreiro Megino on behalf of the PanDA team
8th PanDA Community Forum

Thursday 7 November 2024



@BrookhavenLab

Stream 1 Progress

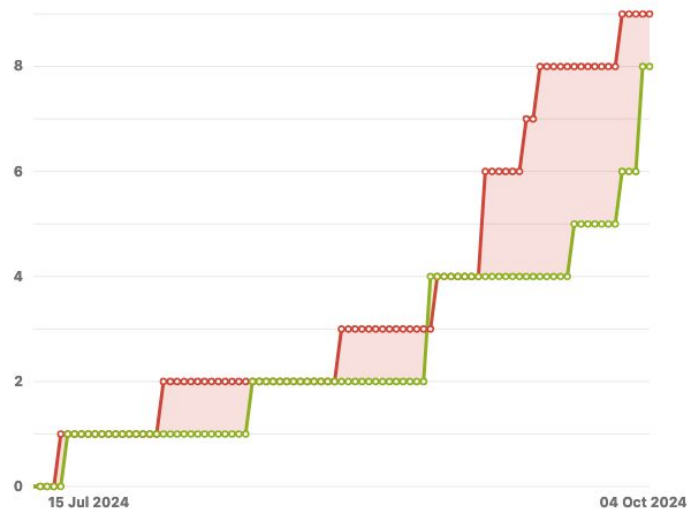
- Progress update on Stream 1 following up on the [presentation](#) during June's Community Forum
- PanDA server package started ~20 years ago
- ATLAS Distributed Computing has gone through many changes during that time
- Technical/infrastructure requirements have also evolved
- Several modules have become extremely complicated to understand/maintain



Stream 1

Assignee	OPEN	IN PROGRESS	RESOLVED	CLOSED	T:
	1	0	0	2	3
	1	2	0	6	9
	0	0	0	2	2
	2	1	1	4	8
	1	0	0	1	2
	0	0	0	2	2
	0	0	0	1	1
Unassigned	4	0	0	0	4
Total Unique Issues:	9	3	1	18	31

Grouped by: Status Showing 8 of 8 statistics.



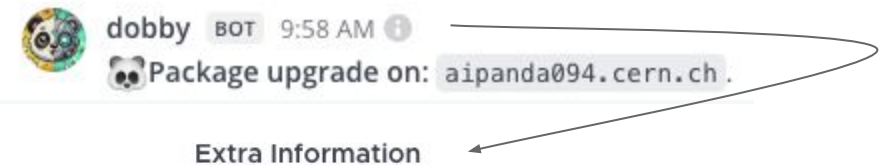
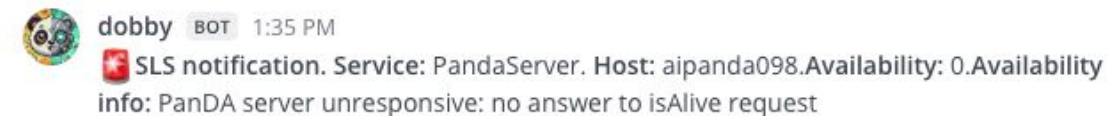
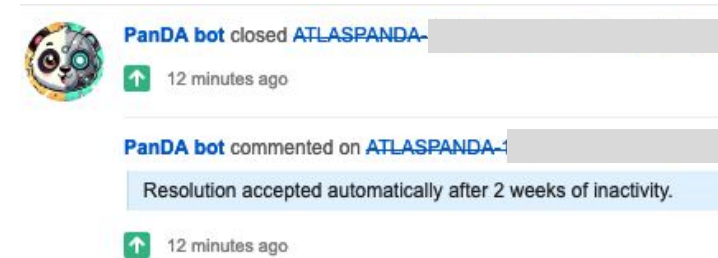
Issues in the last 90 days (grouped daily) [View in Issue Navigator](#)

- Created issues (9)
- Resolved issues (8)

- Improve and modernize our code and environment
- Overlooked area in the past, but with increased importance as we grow
- Some dedicated effort with participation across the team
- Issues range from small improvements to larger refactorings

Automating some of our daily activities

- A clean JIRA for a clear mind
 - Close tickets in resolved state after 2 weeks
 - Close tickets in any state after 6 months of inactivity
- Internal MM chatops with bot updates
 - Introductory task for Jammel Brooks (IU)
 - Optimize our infrastructure management
 - SLS notifications
 - Previously email notifications
 - ADC Ops might have noticed the improvement across all PanDA SLS categories in the last year
 - Server upgrade notifications: help synchronize our upgrades and increase awareness



Extra Information

Property	Value
Package	https://github.com/PanDAWMS/panda-server
Branch	prodsourceLabel=ddm
Commit	d24df2896e6a701a586371fc8e3ce968625192b0

Obsoleting CloudConfig and other old code/tables

- CloudConfig represented a static and obsoleted view of the ATLAS Grid
- Most concepts were obsoleted when moving to WORLD (nucleus+satellite) model
- CloudConfig defined default T1 queues: changing the name of the queue in CRIC has led to incidents in the past
- We deleted every reference to this table and simplified related code in multiple components/agents
- The effort led to multiple side tasks
 - ATLAS_PANDAMETA database schema has been greatly cleaned up: 5 tables left from >15
 - Obsolete code was removed for features that are no longer used/functioning (e.g. PanDA mover)
 - We are overall looking for and deleting unused code

DataService clean up volume II

- Significant effort already since beginning of January
- Code refactoring for the pandaserver/dataservice [directory](#), updating nine agents in total: [DynDataDistributor](#), [Activator](#), [Closer](#), [Finisher](#), [EventPicker](#), [EventLookupClientEI](#), [DDMHandler](#), [Setupper](#) and [Adder](#)
 - Roughly speaking, jobs go through UserIF → Setupper → Activator → JobDispatcher (→Watcher) → Adder → Closer → Finisher
- Removing obsolete code base and agents
- Modernizing the code base, refactoring and modularizing the code following PEP8 standards

DataService clean up volume II (cont.)

Refactoring of Setupper [ATLASPANDA-917] #335

Merged EdwardKaravakis merged 44 commits into master from setupper on Apr 30

Conversation 72 Commits 44 Checks 0 Files changed 13 +2,554 -2,616

EdwardKaravakis commented on Mar 29 • edited

PR of setupper for the usual refactoring: variable and method renaming for pep8 compliance and removal of some obsolete code, unused variables and etc.

PyLint report for setupper_atlas_plugin:
Your code has been rated at 8.73/10 - original code was 5.01/10

PyLint report for setupper:
Your code has been rated at 9.64/10 - original code was 3.65/10

Reviewers: fbarreir

Assignees: No one—assign yourself

Labels: None yet

Projects

Second mild refactoring attempt of Setupper #371

Merged EdwardKaravakis merged 16 commits into master from setupper on Jul 17

Conversation 10 Commits 16 Checks 0 Files changed 5 +280 -332

EdwardKaravakis commented on Jul 5

Second mild refactoring attempt of Setupper / Setupper ATLAS Plugin to address most of the comments received after the meeting with Tadashi and Fernando

- Removing obsolete code
- Correcting a bit the docs

This PR does **not** include splitting of the setup_source and setup_destination methods of the ATLAS plugin.

Your code has been rated at 9.13/10 (previously the first mild refactored code was 8.73/10)

Reviewers: fbarreir

Assignees: No one—assign yourself

Labels: None yet

Code gets reviewed and there are several iterations with Fernando and Tadashi if a code is obsolete / the functionality is still needed

DDMHandler & DDM refactor #321

Merged EdwardKaravakis merged 16 commits into master from ddmrefactor on Mar 27

Conversation 14 Commits 16 Checks 0 Files changed 14 +1,015 -818

EdwardKaravakis commented on Mar 22 • edited

Usual refactor, the DDMHandler file was in a very good state already though. DDM.py was refactored as well:

Your code has been rated at 9.18/10 (before it was 5.74).

Reviewers: fbarreir

Assignees: No one—assign yourself

Labels

First attempt to refactor Adder #410

Merged EdwardKaravakis merged 12 commits into master from adder last week

Conversation 5 Commits 12 Checks 0 Files changed 13 +1,647 -1,488

EdwardKaravakis commented 3 weeks ago • edited

First PR attempting to refactor Adder.

- PyLint and PEP8 compliant (renaming, changing code and etc).
- Removed unused variables.

I did not yet split big methods into smaller ones - it will come later on.

PyLint scores:
adder_atlas_plugin.py: Your code has been rated at 6.62/10 (2.79/10 before)
adder_dummy_plugin.py: Your code has been rated at 7.14/10 (1.43/10 before)
adder_gen.py: Your code has been rated at 9.13/10 (7.15/10 before)
adder_plugin_base.py: Your code has been rated at 8.75/10 (2.22/10 before)
adder_result.py: Your code has been rated at 10.00/10 (0.95/10 before)
adder_simple_plugin.py: Your code has been rated at 3.33/10 (0.25/10 before)

Reviewers: fbarreir

Assignees: No one—assign yourself

Labels: None yet

Projects: None yet

Refactor eventLookupClientEI.py [ATLASPANDA-919] #320

Merged EdwardKaravakis merged 4 commits into master from eventlookupclient on Mar 21

Conversation 2 Commits 4 Checks 0 Files changed 4 +81 -68

EdwardKaravakis commented on Mar 21

Usual refactor: pep8, docstrings, function annotations, variable/method rename and etc. Not sure why do we need the EventLookupClientEI class, since it is too small, it could be just another function in dyn_data_distributer I think.

Your code has been rated at 9.19/10 (original code was: 3.59).

Reviewers: fbarreir

Assignees: No one—assign yourself

Labels

Removing XML parsing #416

Merged EdwardKaravakis merged 4 commits into master from adder last week

Conversation 0 Commits 4 Checks 0 Files changed 1 +65 -147

EdwardKaravakis commented 2 weeks ago • edited

Removing XML parsing, adding a warning when JSON parsing fails

Reviewers: fbarreir

ATLAS Testbed for PanDA on Kubernetes

- Deployed a testing infrastructure within a Kubernetes environment to evaluate new PanDA features before they go into production [[ATLASPANDA-1009](#)]
- Modified PanDA Kubernetes helm charts and updated the Docker images to support connecting to an Oracle database instead of PostgreSQL to replicate production ATLAS environment
- Requested new INT8R accounts (thanks to Luca Canali) and installed the schemas for all components - WiP to also have ProdSys schema deployed
- Aim is to have a continuous, low-level task submission and job execution
- Based on experience, it could pave the way for future large-scale Kubernetes deployments for ATLAS, marking a transition from the current OpenStack VMs-based infrastructure

Conclusions

- Simplify PanDA: make it more **accessible** and **maintainable**
- Part of PanDA's codebase around for almost 20 years
 - While significant progress has been made to modernize the PanDA codebase and infrastructure: **continuous** effort is **essential**
- **Enhancing** development practices
 - Peer reviews for pull requests when possible/needed
 - Future-proof the codebase, ensuring it is maintainable and scalable
 - Introducing development tools to ensure code consistency and quality

Team is working across streams, doing support, operation, development, documentation and modernization at the same time

Thank you for your attention!

-

Questions?

