



An intelligent Data Delivery Service for and beyond the ATLAS Experiment

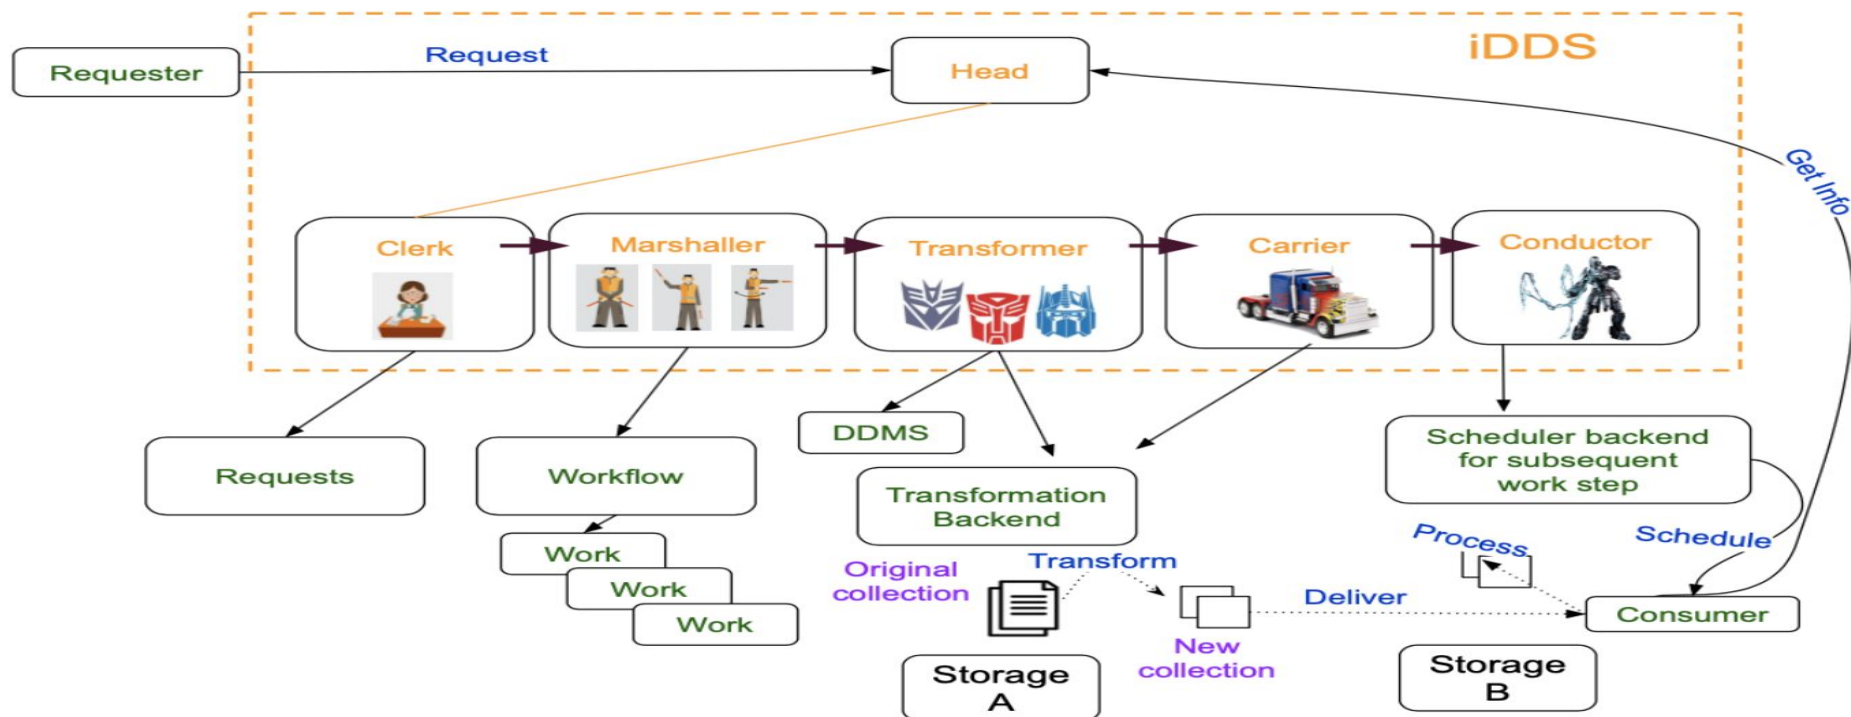
Wen Guan, Tadashi Maeno
Brian Bockelman, Torre Wenaus
Fernando Barreiro, Fahui Lin, Rui Zhang, Misha Borodin, Paul Nilsson

25th International Conference on Computing
in High-Energy and Nuclear Physics

Apr 26, 2021

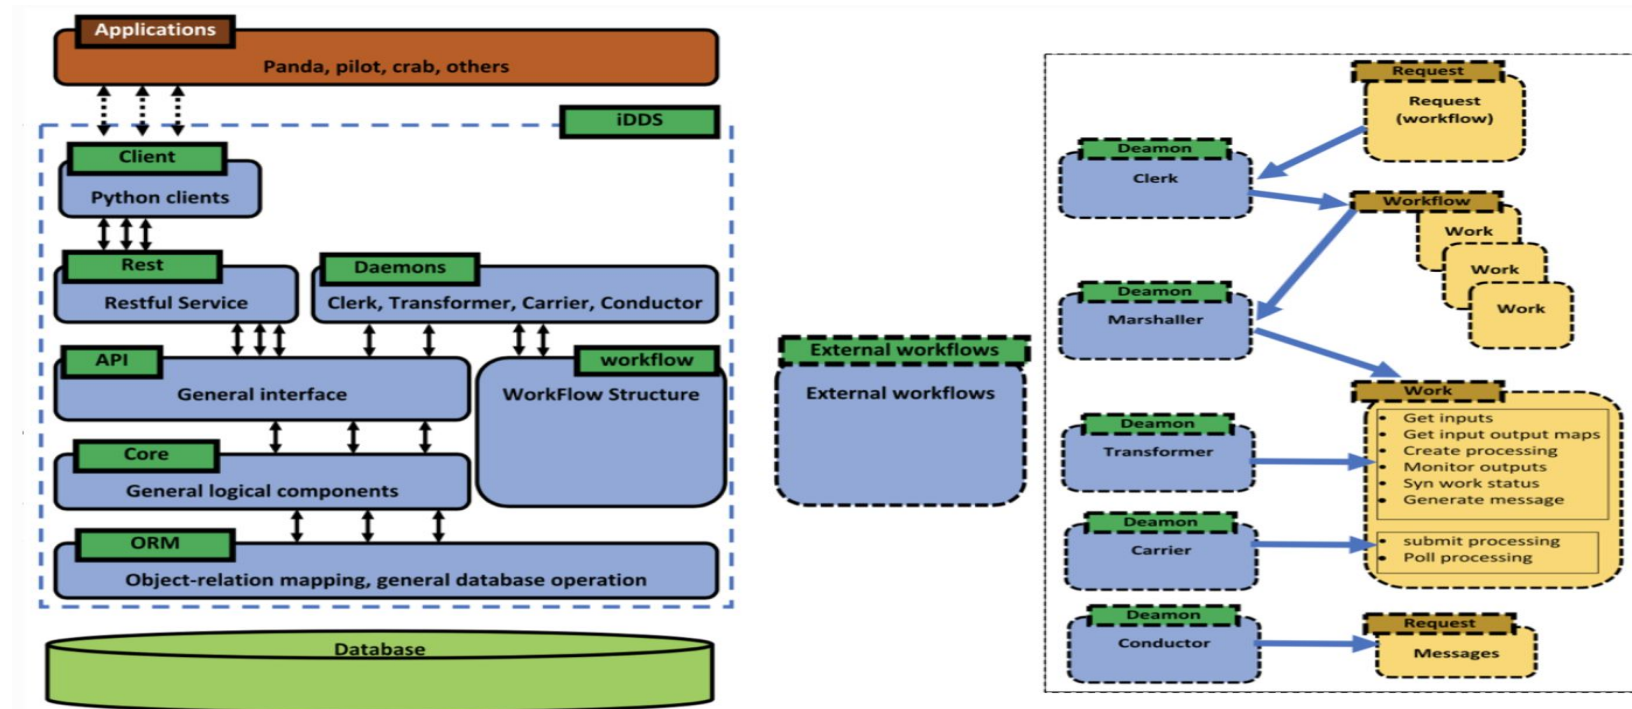
iDDS (a joint project with IRIS-HEP and ATLAS)

- An intelligent service to transform and deliver needed data to consumers, to orchestrate WFMS and DDMS with generalized workflows
 - Experiment agnostic
 - Extraction and abstraction of functions for orchestration
 - Maintainability and extensibility with plugin architecture



iDDS

- **Client/Server**
 - REST server to handle requests from clients.
 - Employ different backends for transformation: PanDA, Rucio, Condor and so on.
- **Layered architecture of the server**
 - Every layer abstracts a group of functions, hiding the complexity of different logics on different layers.



iDDS

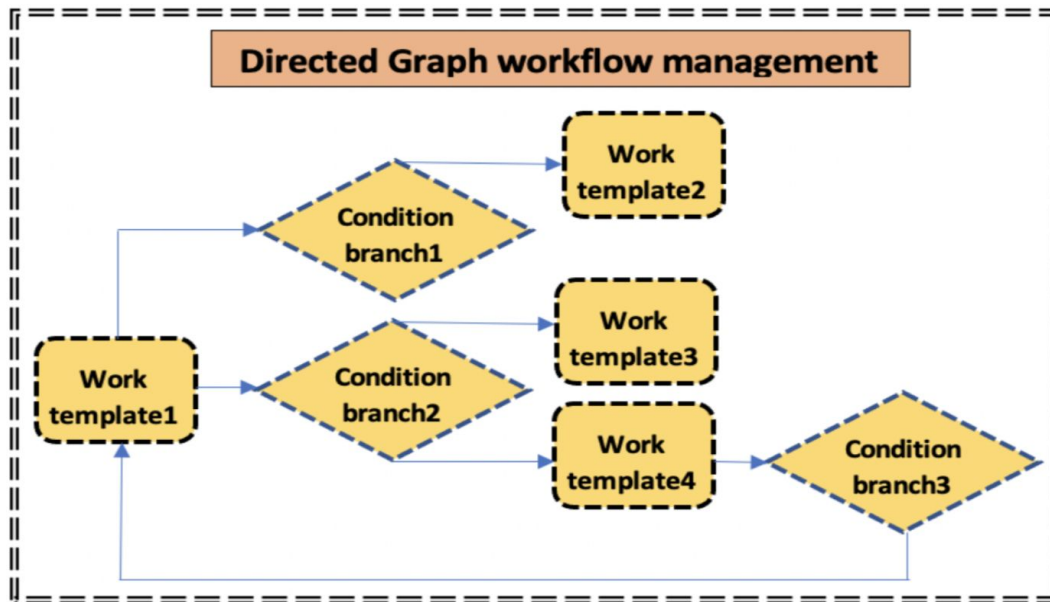
- **DAG workflow management**

- **Task Level DAG**


- Relation description is among tasks.
- When there are new outputs in a task, new jobs are generated for the dependent tasks.
- When a task is terminated, dependent tasks are triggered.

- **Job Level DAG**

- Relation description is among pre-defined jobs.
- Job grouping based on WFMS.
- When a job is terminated, dependent jobs are triggered.



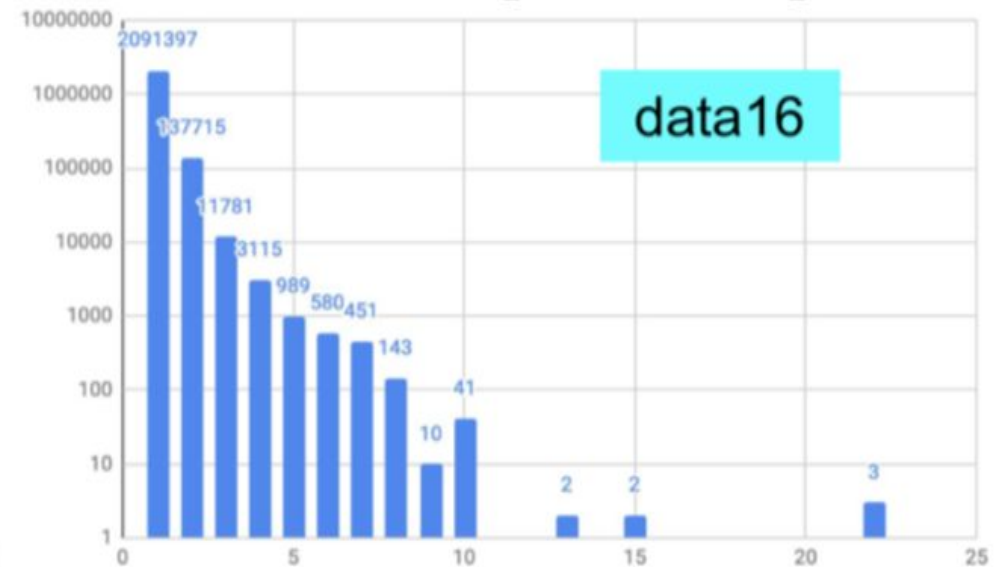
Task Level DAG



Job Level DAG

iDDS Data Carousel

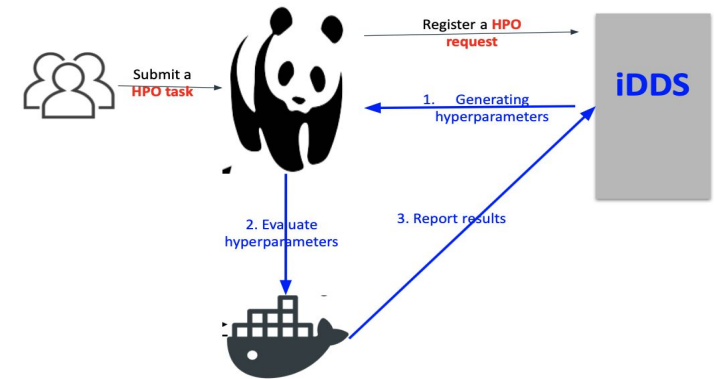
- Orchestrates Rucio to collect and digest file information, and lets JEDI/PanDA process only prestaged files with proper granularities and grouping.
- In ATLAS production since May 2020
- Has processed in total about 21 PB data.
- Reduced a lot of redundant job attempts.
 - With iDDS, the tail of jobs with a lot of attempts becomes shorter.



X: attempt times, Y: number of jobs

iDDS HPO(HyperParameterOptimization)

- To provide a fully-automated platform for hyper parameter optimization on top of geographically distributed GPU resources on the grid, HPC, and clouds.
- Used by ATLAS ML users, not specific to ATLAS.
- The usage is increasing.
 - FastCaloGAN
 - ToyMC




Requests:

Show 10 entries

Search: HPO

request_id	scope	name	status	transform_status	in_status	in_total_files	in_processed_files	out_status	out_total_files
37223	hpo	hpo.24509485.2021_03_10_19_04_57_473740655	Transforming	Transforming	Open	2	2	Open	2
37225	hpo	hpo.24509487.2021_03_10_19_04_57_829887699	Transforming	Transforming	Open	2	2	Open	2
37227	hpo	hpo.24509488.2021_03_10_19_05_11_831247870	Transforming	Transforming	Open	2	2	Open	2
37229	hpo	hpo.24509491.2021_03_10_19_05_44_755961431	Transforming	Transforming	Open	2	2	Open	2
37231	hpo	hpo.24509492.2021_03_10_19_05_45_00621075	Transforming	Transforming	Open	2	2	Open	2
37233	hpo	hpo.24509493.2021_03_10_19_05_48_072992706	Transforming	Transforming	Open	2	2	Open	2
37235	hpo	hpo.24509494.2021_03_10_19_06_18_694058984	Transforming	Transforming	Open	2	2	Open	2
37237	hpo	hpo.24509495.2021_03_10_19_06_18_859385180	Transforming	Transforming	Open	2	2	Open	2
37239	hpo	hpo.24509500.2021_03_10_19_06_39_335869695	Transforming	Transforming	Open	2	2	Open	2
37241	hpo	hpo.24509509.2021_03_10_19_06_47_514355545	Transforming	Transforming	Open	2	2	Open	2


Showing 911 to 920 of 922 entries (filtered from 25,895 total entries)



iDDS Rubin Observatory (LSST)

● LSST exercise

- Use the experiment-agnostic DOMA PanDA instance.
- Job Level DAG: cascade of chains for multiple-step processing.
- iDDS manages the dependencies and triggers to release jobs incrementally when all dependencies are ready, instead of blocking tasks until all previous tasks finish, to avoid long waiting.
- Scaling tested with 50K DAG jobs, will test with 150K jobs.



1072	shared_pipecheck_20210324T001908Z_detection test last atpilo1 RequestID: 1072 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 5	20% 1		2021-03-28 11:53:09	2021-03-28 11:53:09	900		US
1071	shared_pipecheck_20210324T001908Z_assembleCoadd test last atpilo1 RequestID: 1071 Errors	finished 5	20% 80% 1 4		2021-03-26 09:44:16	2021-03-26 09:44:16	900		US
1070	shared_pipecheck_20210324T001908Z_consolidateSourceTable test last atpilo1 RequestID: 1070 Errors	finished 49	73% 26% 36 13		2021-03-25 09:20:21	2021-03-25 09:20:21	900		US
1069	shared_pipecheck_20210324T001908Z_makeWarp test last atpilo1 RequestID: 1069 Errors	finished 49	83% 16% 41 8		2021-03-26 00:08:11	2021-03-26 00:08:11	900		US
1068	shared_pipecheck_20210324T001908Z_transformSourceTable test last atpilo1 RequestID: 1068 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 254	95% 224		2021-03-28 11:53:09	2021-03-28 11:53:09	900		US
1067	shared_pipecheck_20210324T001908Z_copyData test last atpilo1 RequestID: 1067 Errors	finished 49	73% 26% 40 9		2021-03-26 07:00:04	2021-03-26 07:00:04	900		US
1066	shared_pipecheck_20210324T001908Z_writeSourceTable test last atpilo1 RequestID: 1066 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 234	95% 0% 224 2		2021-03-28 11:53:09	2021-03-28 11:53:09	900		US
1065	shared_pipecheck_20210324T001908Z_characterizeImage test last atpilo1 RequestID: 1065 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 234	96% 0% 226 1		2021-03-28 11:53:09	2021-03-28 11:53:09	900		US
1064	shared_pipecheck_20210324T001908Z_calibrate test last atpilo1 RequestID: 1064 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 234	95% 0% 224 2		2021-03-28 11:53:09	2021-03-28 11:53:09	900		US
1063	shared_pipecheck_20210324T001908Z_isr test last atpilo1 RequestID: 1063 Errors	finished 234	97% 2% 227 7		2021-03-24 23:49:50	2021-03-24 23:49:50	900		US

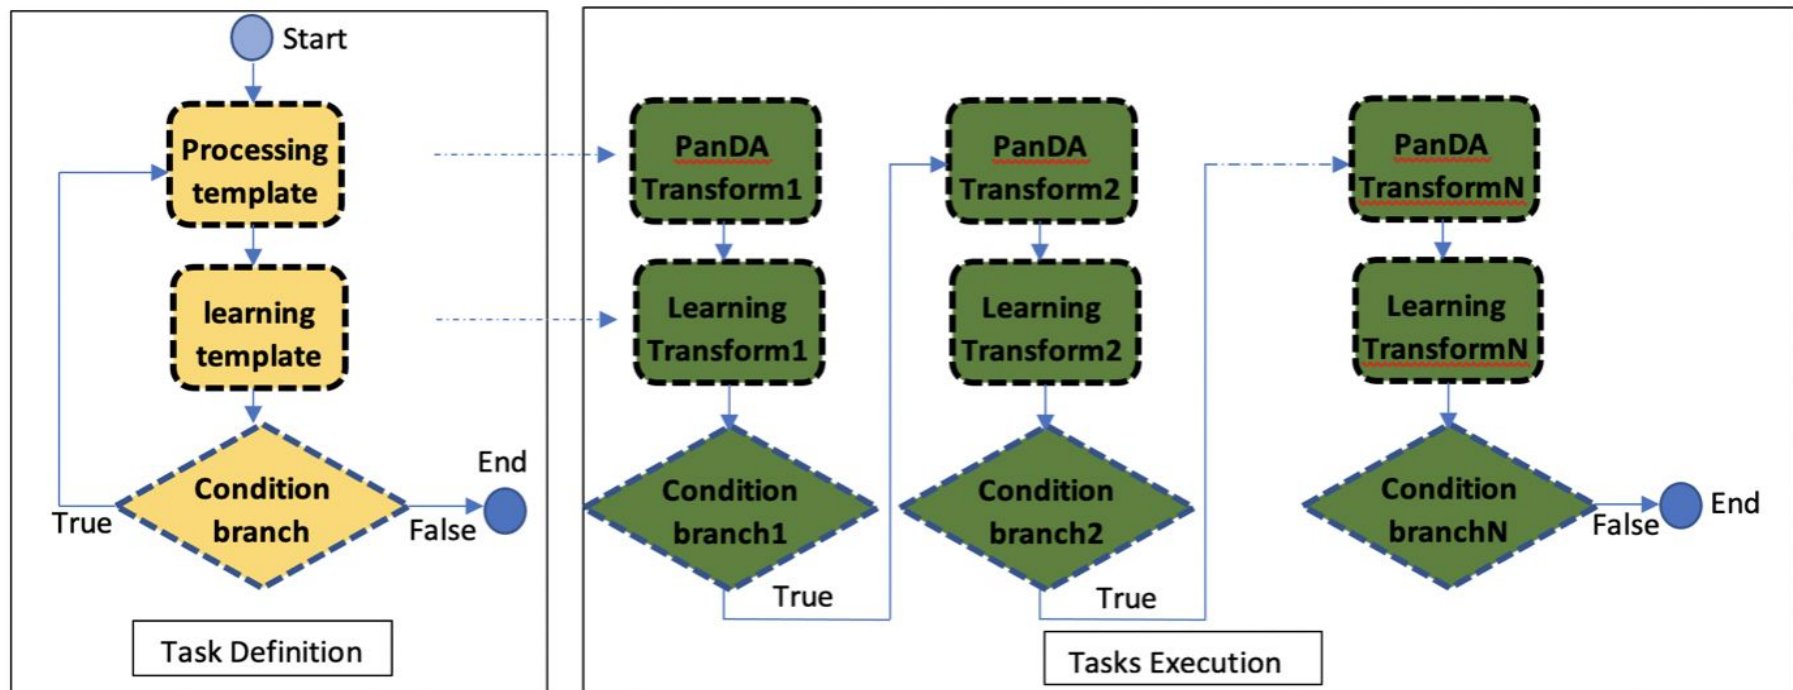
**A workflow with cascade of jobs:
Jobs are grouped to tasks**

1039	shared_pipecheck_20210323T160801Z_calibrate test last atpilo1 RequestID: 1039 Errors	finished 1	100% 1		2021-03-23 18:58:42	2021-03-23 18:58:42	900		US
1038	shared_pipecheck_20210323T160801Z_isr test last atpilo1 RequestID: 1038 Errors	finished 1	100% 1		2021-03-23 17:26:30	2021-03-23 17:26:30	900		US
1037	shared_pipecheck_20210323T160801Z_characterizeImage test last atpilo1 RequestID: 1037 Errors	finished 1	100% 1		2021-03-23 18:30:38	2021-03-23 18:30:38	900		US
1036	shared_pipecheck_20210323T160801Z_transformSourceTable test last atpilo1 RequestID: 1036 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 252	74% 18% 188		2021-03-28 11:53:08	2021-03-28 11:53:08	900		US
1035	shared_pipecheck_20210323T160801Z_pipetaskinit test last atpilo1 RequestID: 1035 Errors	finished 1	100% 1		2021-03-23 11:26:32	2021-03-23 11:26:32	900		US
1034	shared_pipecheck_20210323T160801Z_consolidateSourceTable test last atpilo1 RequestID: 1034 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 49	65% 10% 32 5		2021-03-28 11:53:08	2021-03-28 11:53:08	900		US
1033	shared_pipecheck_20210323T160801Z_characterizeImage test last atpilo1 RequestID: 1033 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 252	75% 1% 190 5		2021-03-28 11:53:08	2021-03-28 11:53:08	900		US
1032	shared_pipecheck_20210323T160801Z_isr test last atpilo1 RequestID: 1032 Errors	finished 252	77% 22% 195 57		2021-03-23 20:02:49	2021-03-23 20:02:49	900		US
1031	shared_pipecheck_20210323T160801Z_writeSourceTable test last atpilo1 RequestID: 1031 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 252	74% 0% 188 1		2021-03-28 11:53:08	2021-03-28 11:53:08	900		US
1030	shared_pipecheck_20210323T160801Z_calibrate test last atpilo1 RequestID: 1030 Errors insufficient inputs are ready, 0 files available, 1*1 files required	running 252	75% 0% 189 1		2021-03-28 11:53:08	2021-03-28 11:53:08	900		US

iDDS ActiveLearning

- **ActiveLearning**

- A simple DAG use case to chain processing and learning tasks.
- To define the subsequent processing task based on the decision making in the learning task which analyze the results of the previous processing task.
- Task templates to generate concrete tasks, and condition branches to control the workflow.
- Under integration with PanDA.



Summary: iDDS Current Status

❖ Main architecture

- iDDS database, core, REST API
- Plugins
- Agents
- Watchdogs

❖ Documents & monitors

- Home page: <https://ids.cern.ch>
- Codes: <https://github.com/HSF/iDDS>
- Documents: <https://ids.readthedocs.io>
- ATLAS monitor: <https://bigpanda.cern.ch/ids/>
- Different monitors are being enriched.

❖ Instances in production

- ATLAS, DOMA

❖ Future developments

- Main Structure improvements
- Monitor improvements
- New use cases
 - Dynamic transformation and placement on demand, for example Derivation on Demand
 - Fine-grained data transformation and delivery, such as Event Streaming Service