intelligent Data Delivery Service (iDDS)

Wen Guan

supported by IRIS-HEP and Wisconsin
University of Wisconsin-Madison
HL-LHC R&D topics
Mar 29, 2021

Contents

- iDDS introduction
  - Main Architecture
  - DAG
  - Monitors
- iDDS achievements
  - Data Carousel
  - HPO(HyperParameterOptimization)
  - LSST exercise
  - ActiveLearning
- Summary

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

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

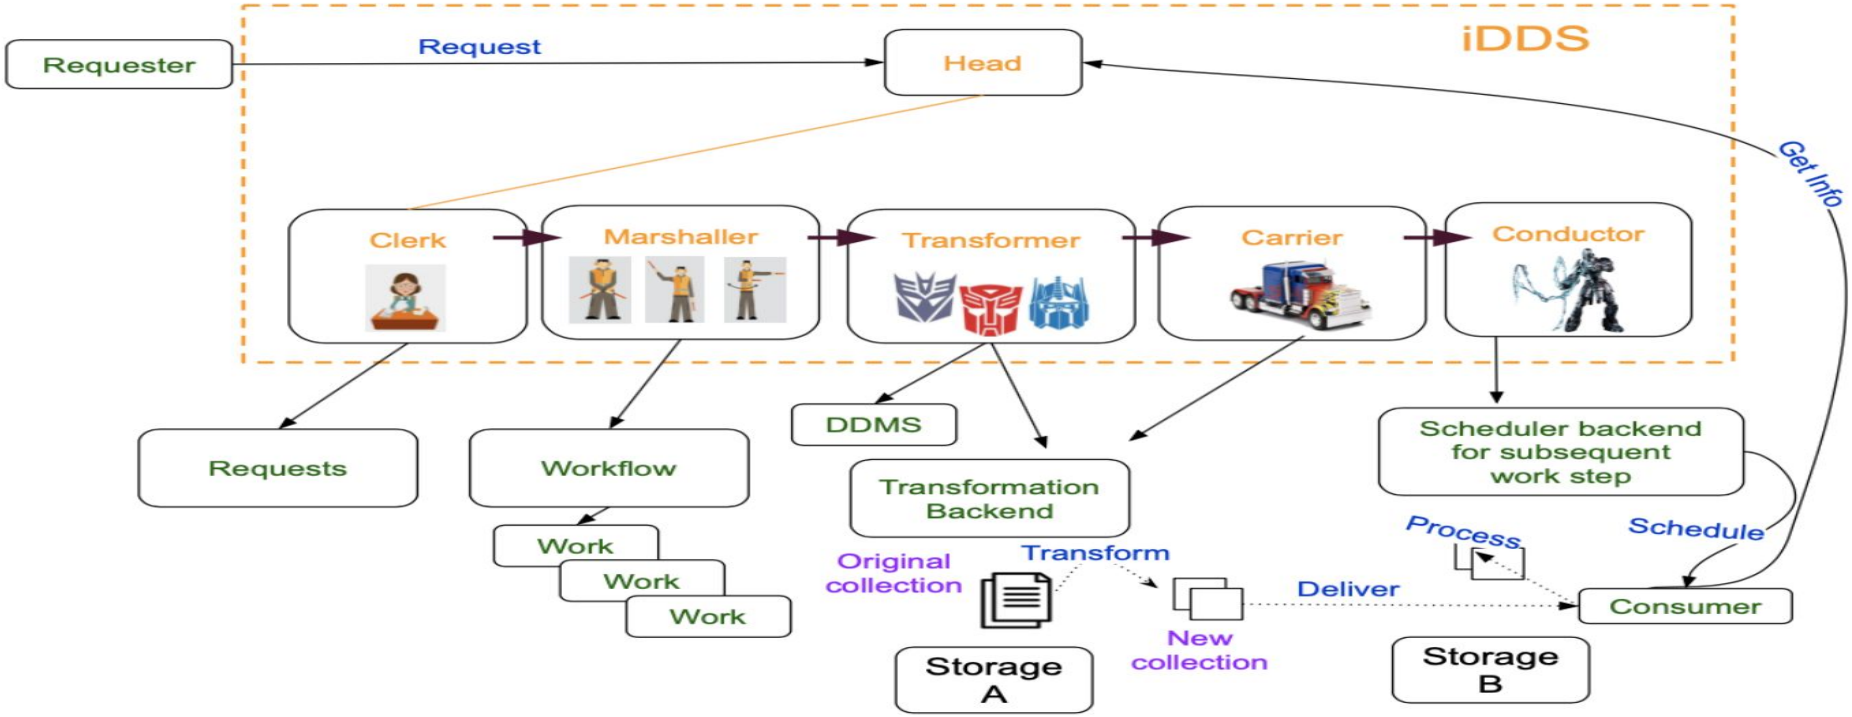

iDDS

- **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, hidden the complexity of different logics on different layers.
- **Workflow-based task management**
  - A workflow is a group of work and their relationship.
  - Work is a transform task: For different experiment or use cases, different work can be defined or developed.

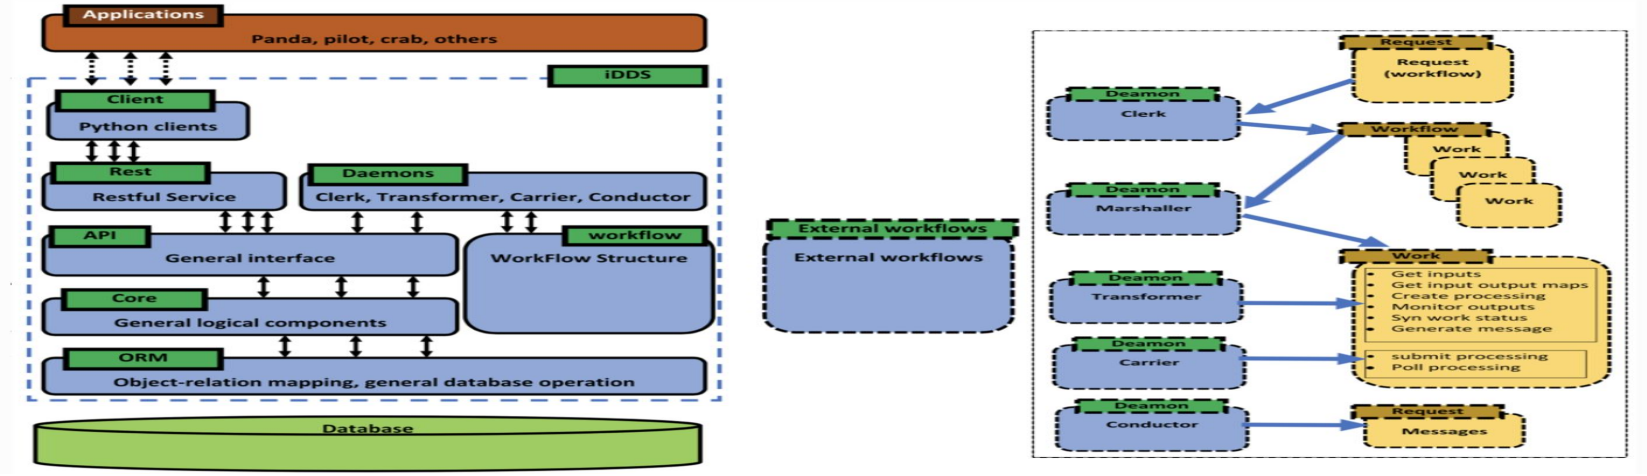

iDDS

- **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.

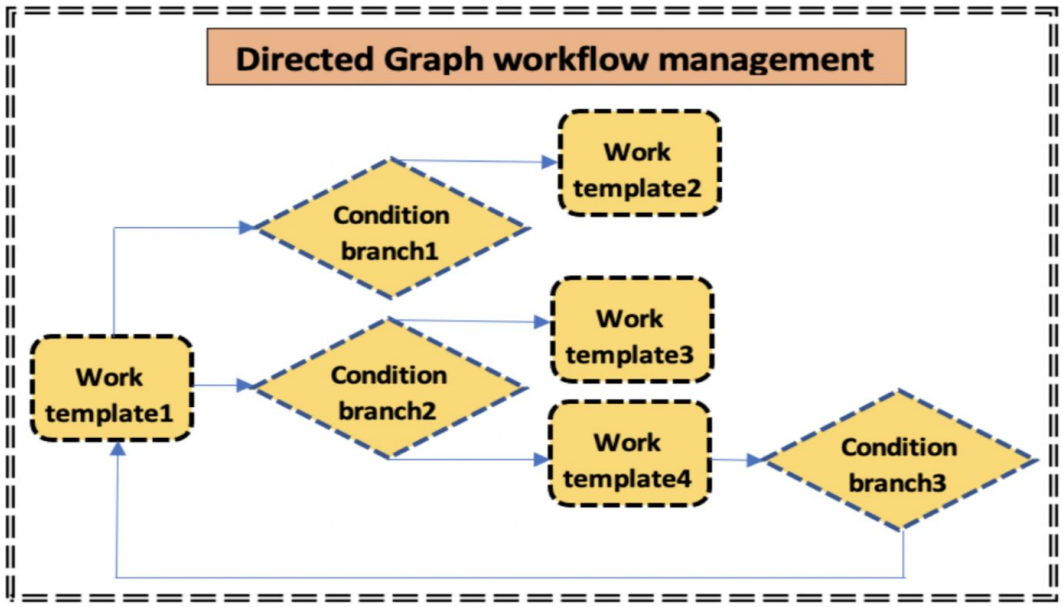

Task Level DAG

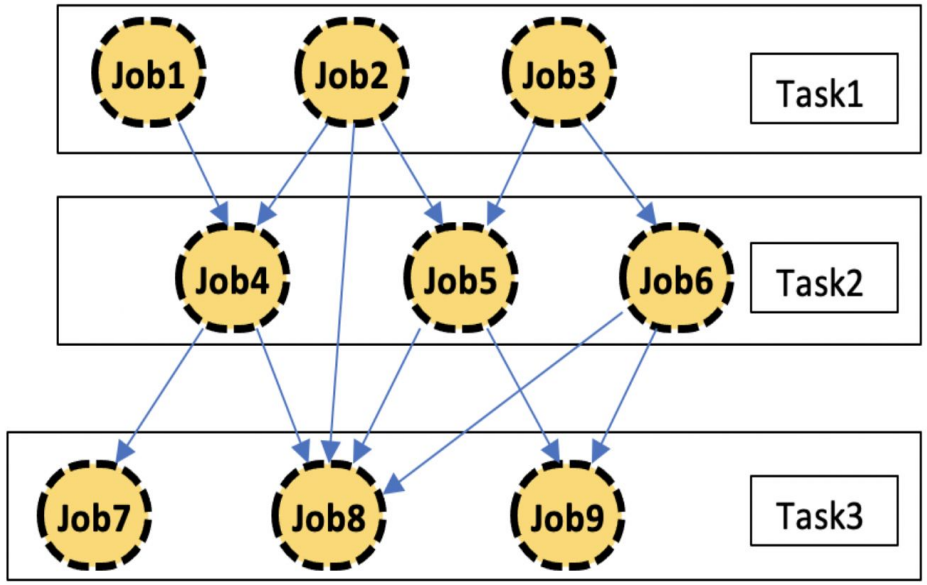

Job Level DAG

iDDS

- **Monitors**
  - iDDS publishes messages and many monitors are built based on these messages.
  - Processing status monitor:
    - Monitor task status
    - bigpanda, Grafana, Elastic Search
  - Service health monitor
    - Monitor whether iDDS is running in a good health
    - Elastic Search
- **ToBeDone:**
  - To construct an experiment-agnostic monitoring.

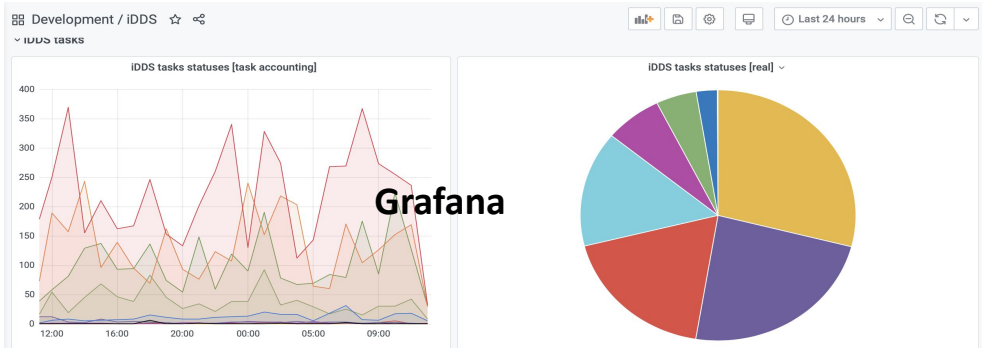

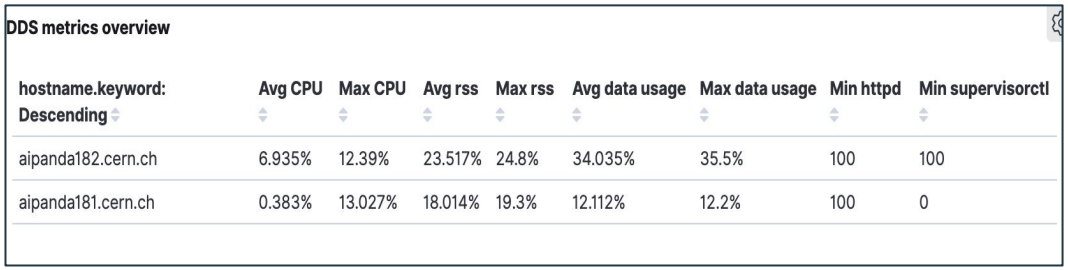

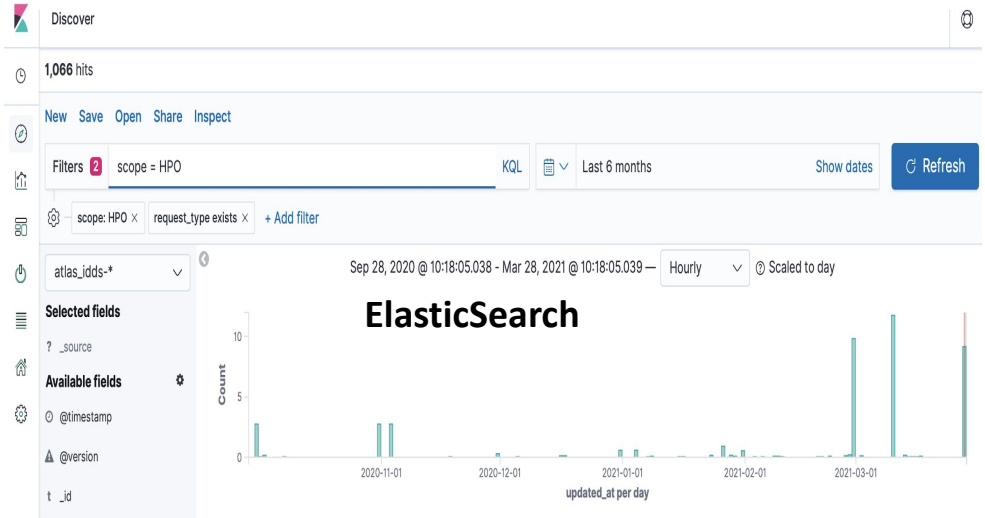

iDDS Achievements

- **Data Carousel**
  - In production since May, 2020
  - Solved the issues with the delayed start of processing data on tape
- **HPO (Hyper Parameter Optimization)**
  - To provide a fully-automated platform for hyper-parameter optimization on top of geographically distributed GPU resources on the grid, HPC, and clouds
  - Advertised to ATLAS ML users, not specific to ATLAS
- **JoB level DAG based workflow management**
  - Using new DOMA PanDA instance used for Rubin Observatory (LSST) exercise
    - Cascade of chains for multi-step processing with thousands of jobs per step
    - Release jobs incrementally for different steps to avoid long waiting
- **Task level DAG based workflow management**
  - High-level workflows specified by DAGs driving workload scheduling.
  - Active Learning for ATLAS dynamic task management (New tasks are generated based on the analysis result of previous tasks).

iDDS Data Carousel

- **iDDS Fine-grained data carousel.**
  - **Orchestrate Rucio to collect and digest file information, and lets JEDI/PanDA process only prestaged files with proper granualities and grouping, instead of processing with dataset-level granualities(not process data until the whole dataset is ready).**
  - **Trigger to process data based on messages from iDDS.**

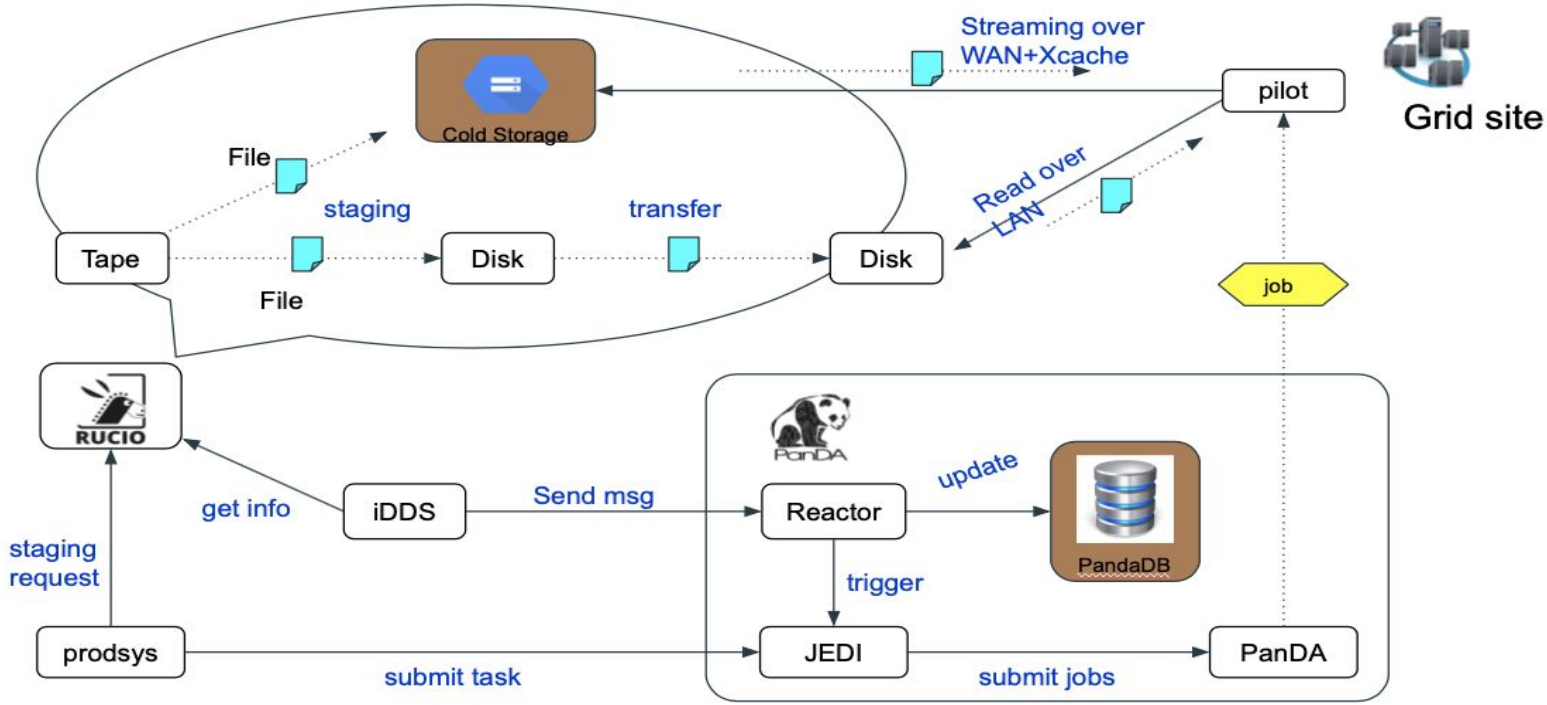

iDDS Data Carousel

- For ATLAS production since May, 2020
- Totally has processed about 21 PB data (Trillion).
- Reduced a lot of redundant job attempts.

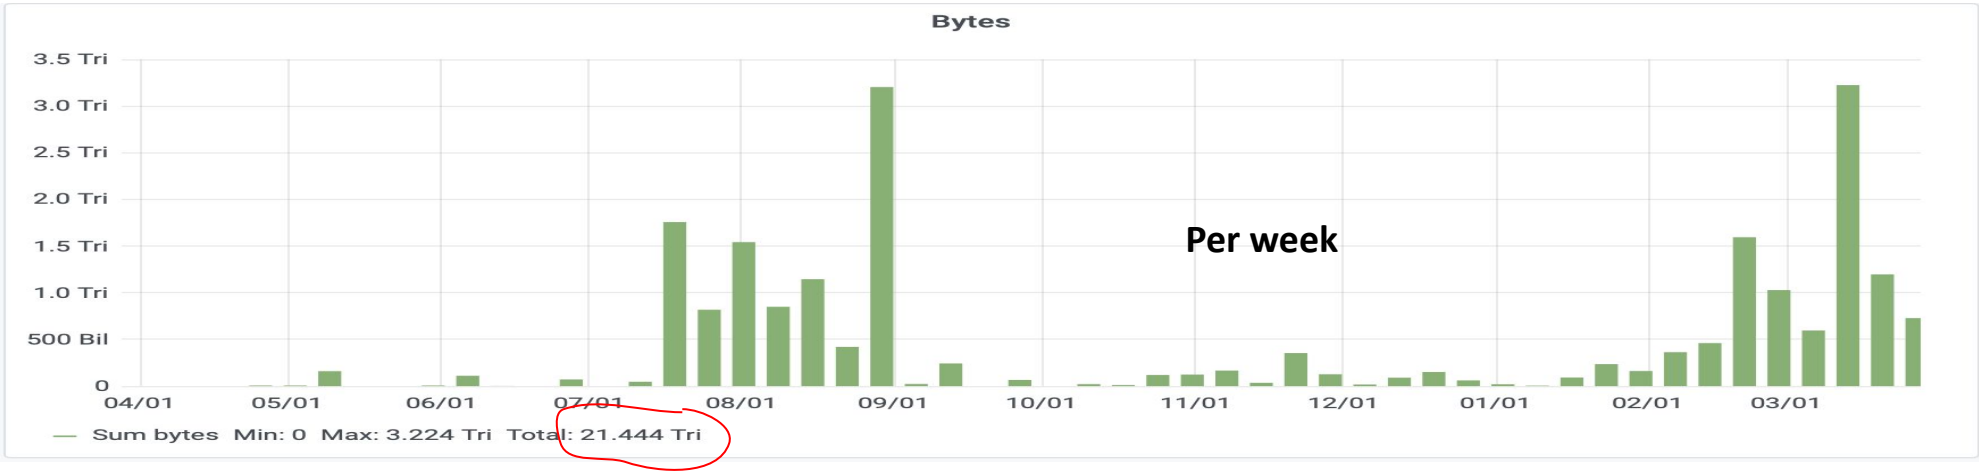

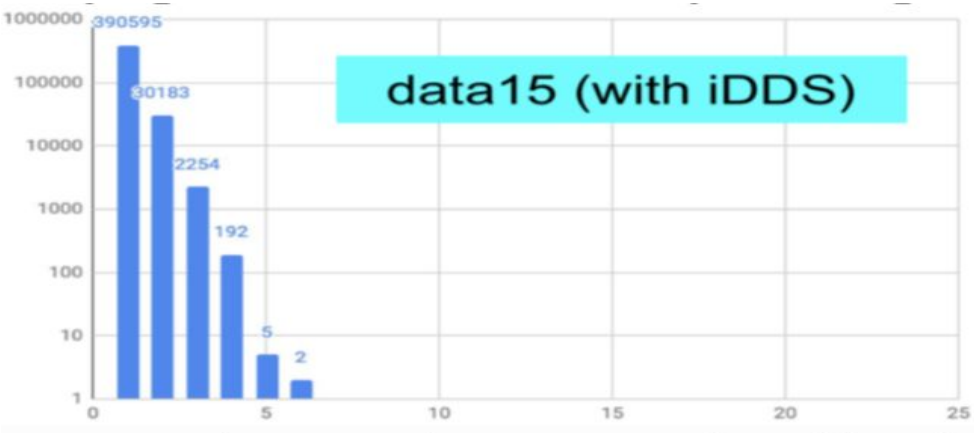

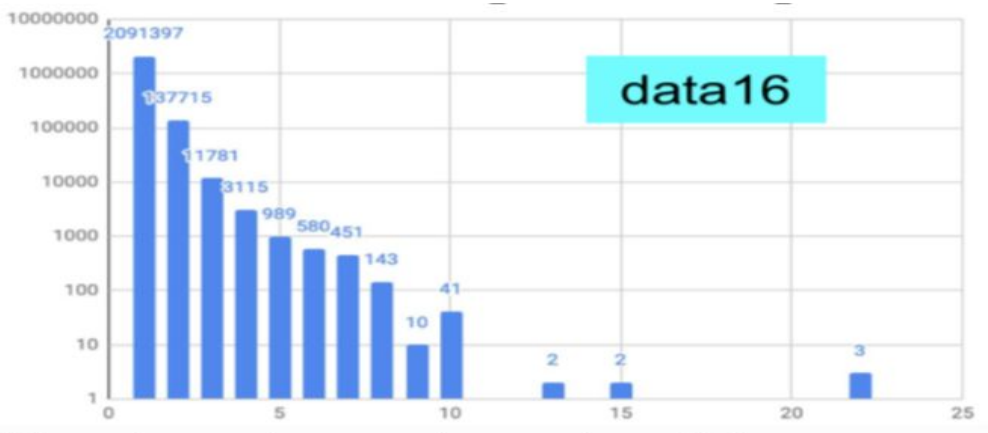

X: attempt times, Y: number of jobs

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
- iDDS generates hyperparameters by iterations and collects results from hyperparameter evaluation.
- Leveraging scalability and resources integration PanDA to evaluate hyperparameters.
- New hyperparameters are generated based on the results of previous iterations.

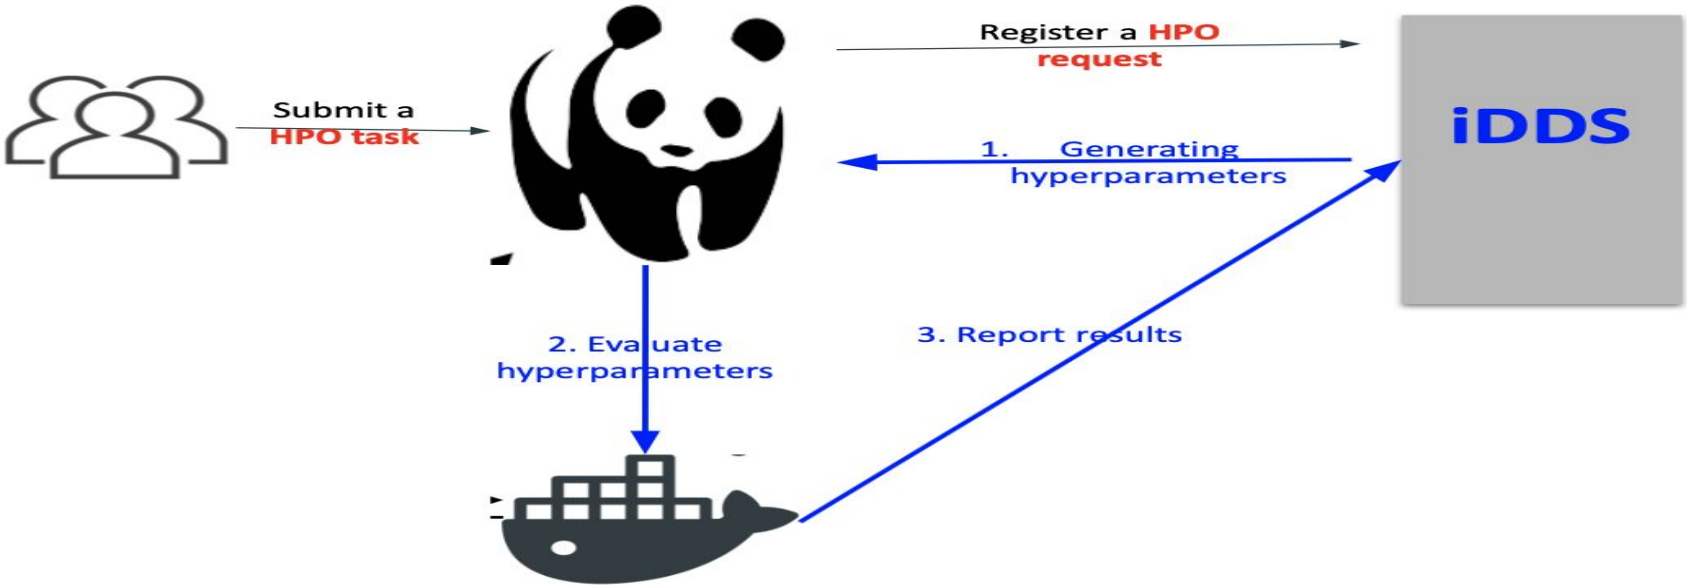

iDDS HPO(HyperParameterOptimization)

- Advertised to ATLAS ML users, not specific to ATLAS
- The usage is increasing.
- Advanced use cases
  - Segmented HPO, distributed training (See Rui Zhang's presentation on Mar 03, https://indico.cern.ch/event/1004145/)

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) Previous 1 ... 89 90 91 92 93 Next

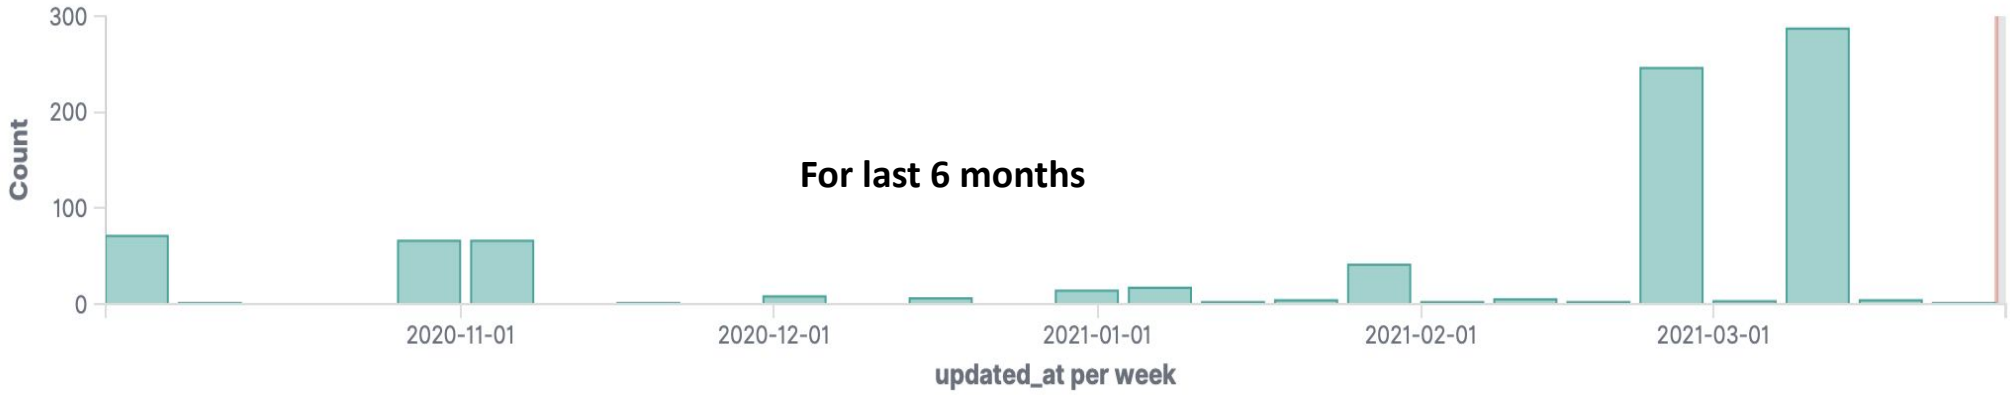

iDDS Rubin Observatory (LSST)

- **LSST exercise (Job level DAG)**
  - Thousands of jobs and their dependencies are defined: cascade of chains for multiple-step processing.
  - Group jobs to tasks based on their dependencies, to adapt PanDA task management.
  - 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.

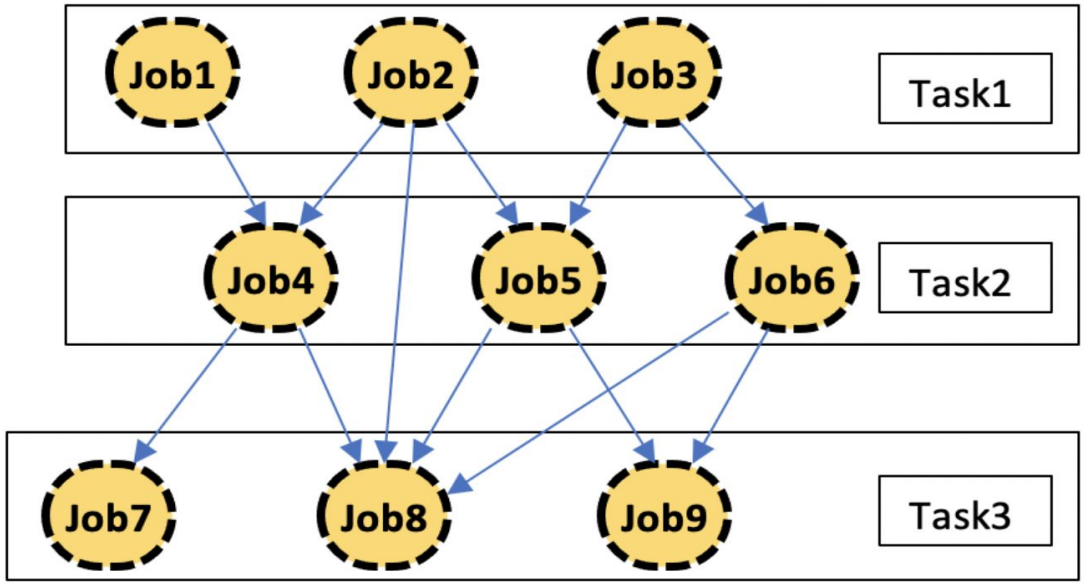

iDDS Rubin Observatory (LSST)

- **LSST exercise**
  - Use the experiment-agnostic DOMA PanDA instance.
  - A new DOMA iDDS instance is deployed, to work with DOMA PanDA instance.
  - Various workflows with cascade of jobs are submitted.
  - Scale tests are ongoing.

1072	shared_pipecheck_20210324T001908Z_detection test lsst atlpilo1 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 lsst atlpilo1 RequestID: 1071 Errors	finished 5	20% 80% 1 4		2021-03-26 00:44:16	2021-03-26 00:44:16	900		US
1070	shared_pipecheck_20210324T001908Z_consolidateSourceTable test lsst atlpilo1 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 lsst atlpilo1 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 lsst atlpilo1 RequestID: 1068 Errors insufficient inputs are ready. 0 files available, 1*1 files required	running 234	95% 224		2021-03-28 11:53:09	2021-03-28 11:53:09	900		US
1067	shared_pipecheck_20210324T001908Z_consolidateVisitSummary test lsst atlpilo1 RequestID: 1067 Errors	finished 49	81% 18% 40 9		2021-03-25 07:00:04	2021-03-25 07:00:04	900		US
1066	shared_pipecheck_20210324T001908Z_writeSourceTable test lsst atlpilo1 RequestID: 1066 Errors insufficient inputs are ready. 0 files available, 1*1 files required	running 234	95% 224		2021-03-28 11:53:09	2021-03-28 11:53:09	900		US
1065	shared_pipecheck_20210324T001908Z_characterizeImage test lsst atlpilo1 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 lsst atlpilo1 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 lsst atlpilo1 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 lsst atlpilo1 RequestID: 1039 Errors	done 1	100% 1		2021-03-23 18:58:42	2021-03-23 18:58:42	900		US
1038	shared_pipecheck_20210323T160801Z_isr test lsst atlpilo1 RequestID: 1038 Errors	done 1	100% 1		2021-03-23 17:26:30	2021-03-23 17:26:30	900		US
1037	shared_pipecheck_20210323T160801Z_characterizeImage test lsst atlpilo1 RequestID: 1037 Errors	done 1	100% 1		2021-03-23 18:30:38	2021-03-23 18:30:38	900		US
1036	shared_pipecheck_20210323T100331Z_transformSourceTable test lsst atlpilo1 RequestID: 1036 Errors insufficient inputs are ready. 0 files available, 1*1 files required	running 252	74% 188		2021-03-28 11:53:08	2021-03-28 11:53:08	900		US
1035	shared_pipecheck_20210323T100331Z_pipetaskInit test lsst atlpilo1 RequestID: 1035 Errors	done 1	100% 1		2021-03-23 11:26:32	2021-03-23 11:26:32	900		US
1034	shared_pipecheck_20210323T100331Z_consolidateSourceTable test lsst atlpilo1 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_20210323T100331Z_characterizeImage test lsst atlpilo1 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_20210323T100331Z_isr test lsst atlpilo1 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_20210323T100331Z_writeSourceTable test lsst atlpilo1 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_20210323T100331Z_calibrate test lsst atlpilo1 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.**

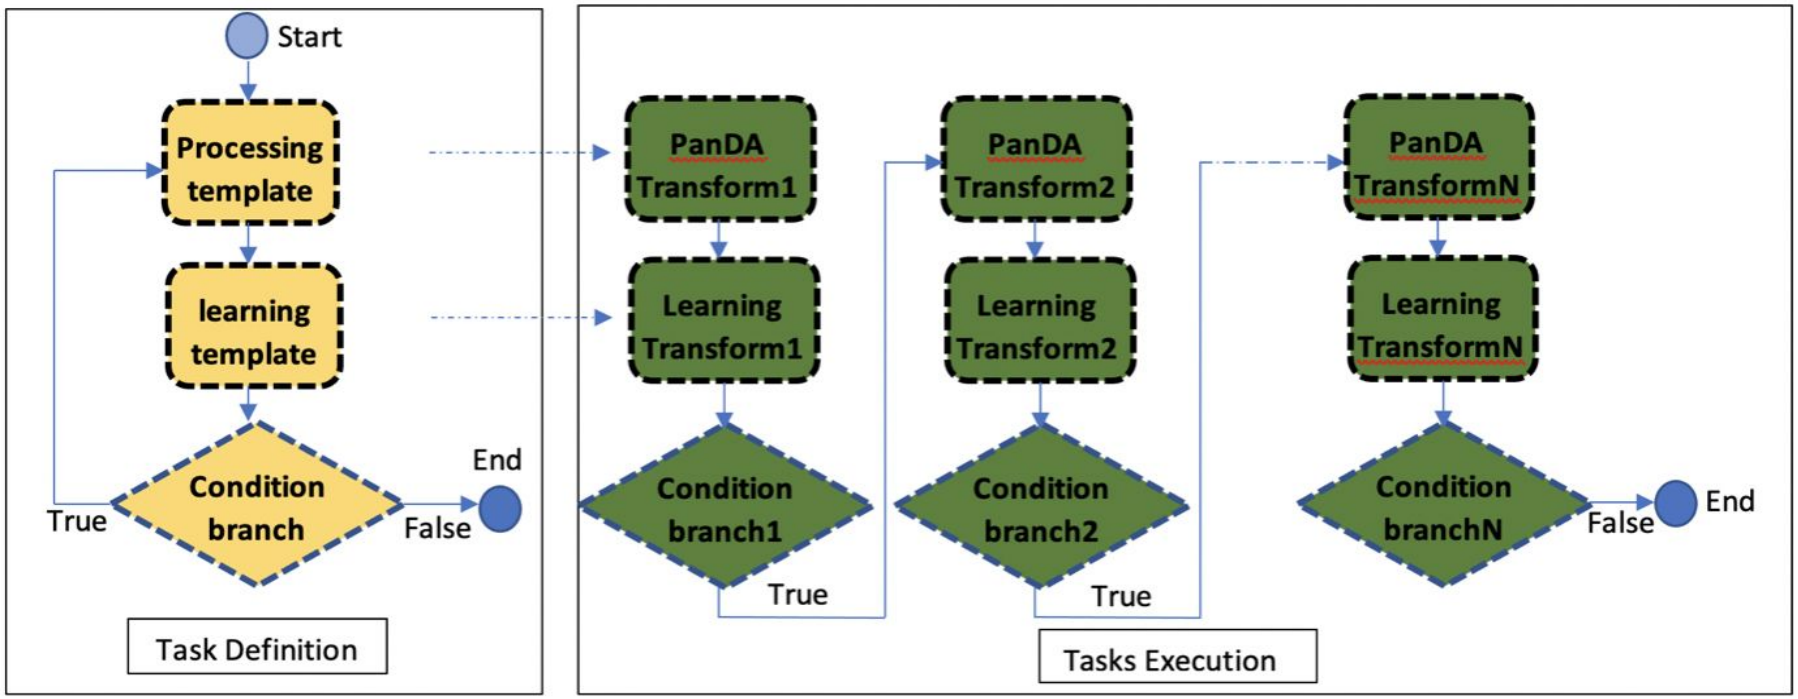

Summary: iDDS Current Status

- **Main architecture**
  - iDDS database, core, REST API
  - Plugins
  - Agents
  - Watchdogs
- **Documents & monitors**
  - Home page: https://idds.cern.ch
  - Codes: https://github.com/HSF/iDDS
  - Documents: https://idds.readthedocs.io (dev)
  - ATLAS monitor: https://bigpanda.cern.ch/idds/
  - Different monitors are being enriched.
- **Instances in production**
  - ATLAS, DOMA
- **Instance for dev**
  - Development and integration
  - New or not well-defined use cases.

Summary: iDDS Current Status

- **Use cases**
  - Fine-grained data carousel
  - Hyper Parameter Optimization (HPO)
  - LSST exercise
  - Decision making for active learning
- future developments
  - Main Structure improvements
  - Monitor improvements
  - New use cases
    - ToyMC (Based on HPO, similar structure)
    - Dynamic transformation and placement on demand, for example Derivation on Demand
    - Fine-grained data transformation and delivery, such as Event Streaming Service