



An intelligent Data Delivery Service for and beyond the ATLAS Experiment

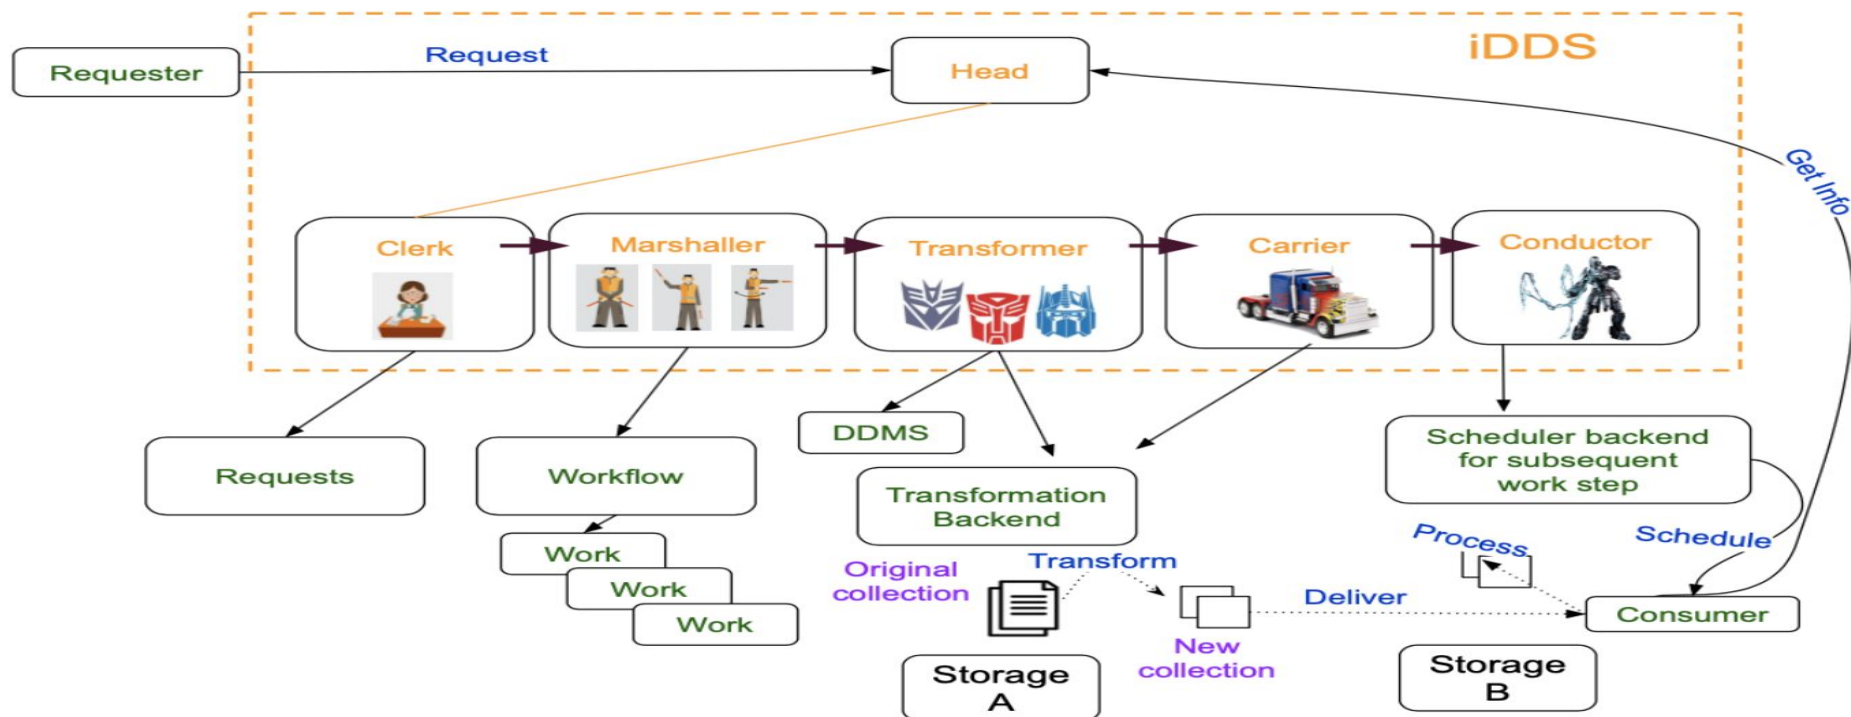
Wen Guan, Tadashi Maeno
Brian Bockelman, Torre Wenaus
Fernando Barreiro, Fahui Lin, Rui Zhang, Misha Borodin, Paul Nilsson

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

May 18, 2021

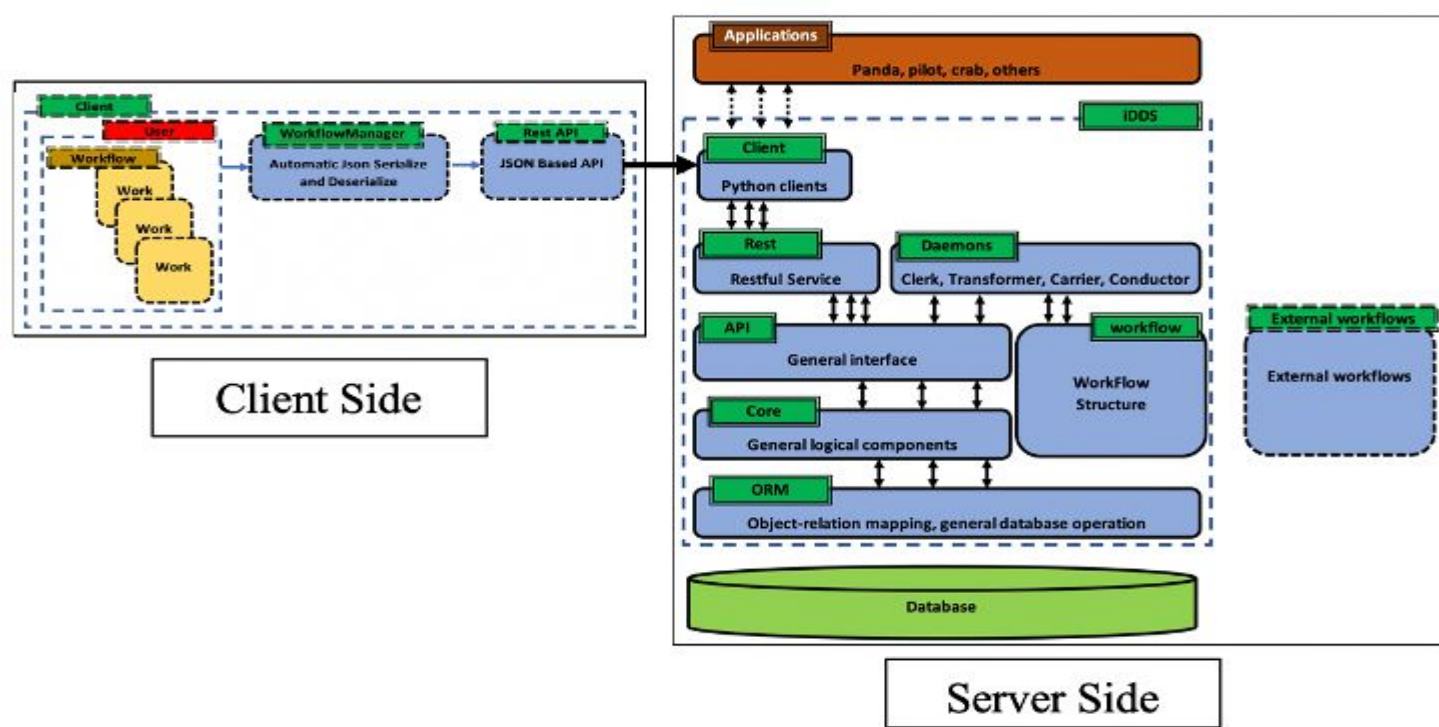
iDDS (a joint project with IRIS-HEP and ATLAS)

- An intelligent service to transform and deliver needed data to consumers, to orchestrate workflow and data management.
 - Experiment agnostic
 - Extraction and abstraction of functions for orchestration
 - Maintainability and extensibility with a 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

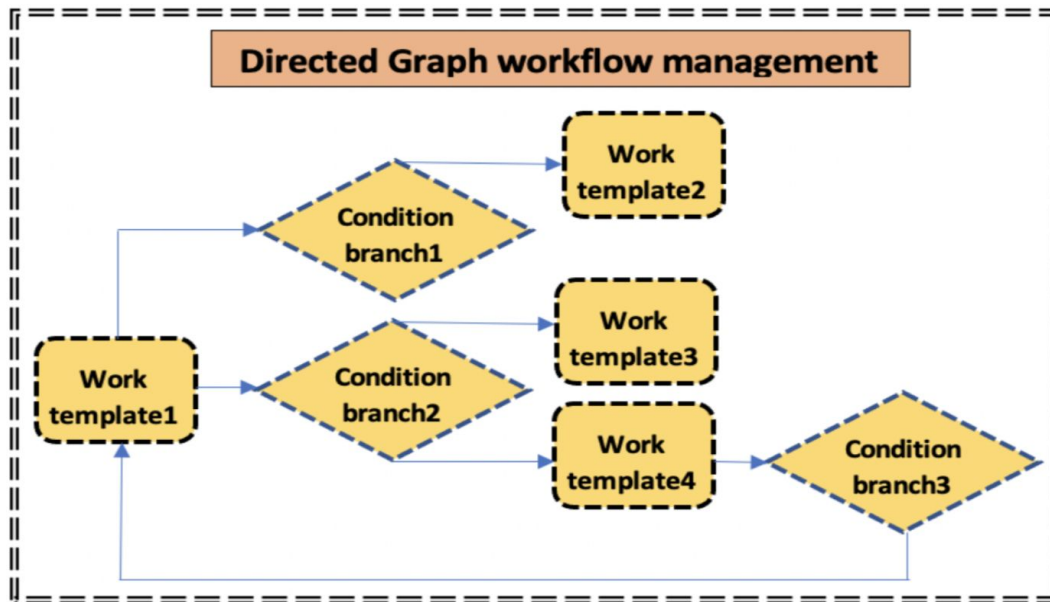
- **DAG workflow management**

- **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 workflow management system.
- When a job is terminated, dependent jobs are triggered.



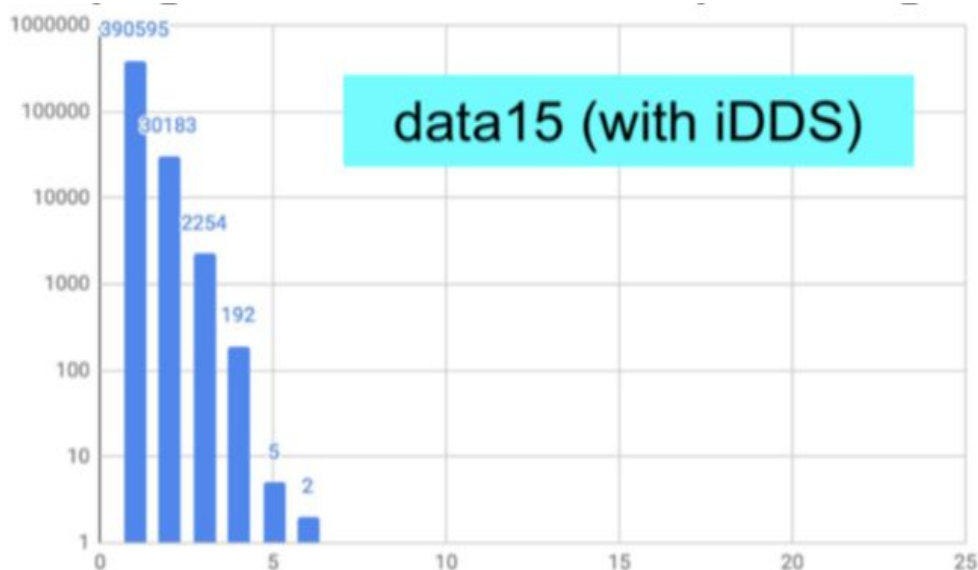
Task Level DAG



Job Level DAG

iDDS Data Carousel

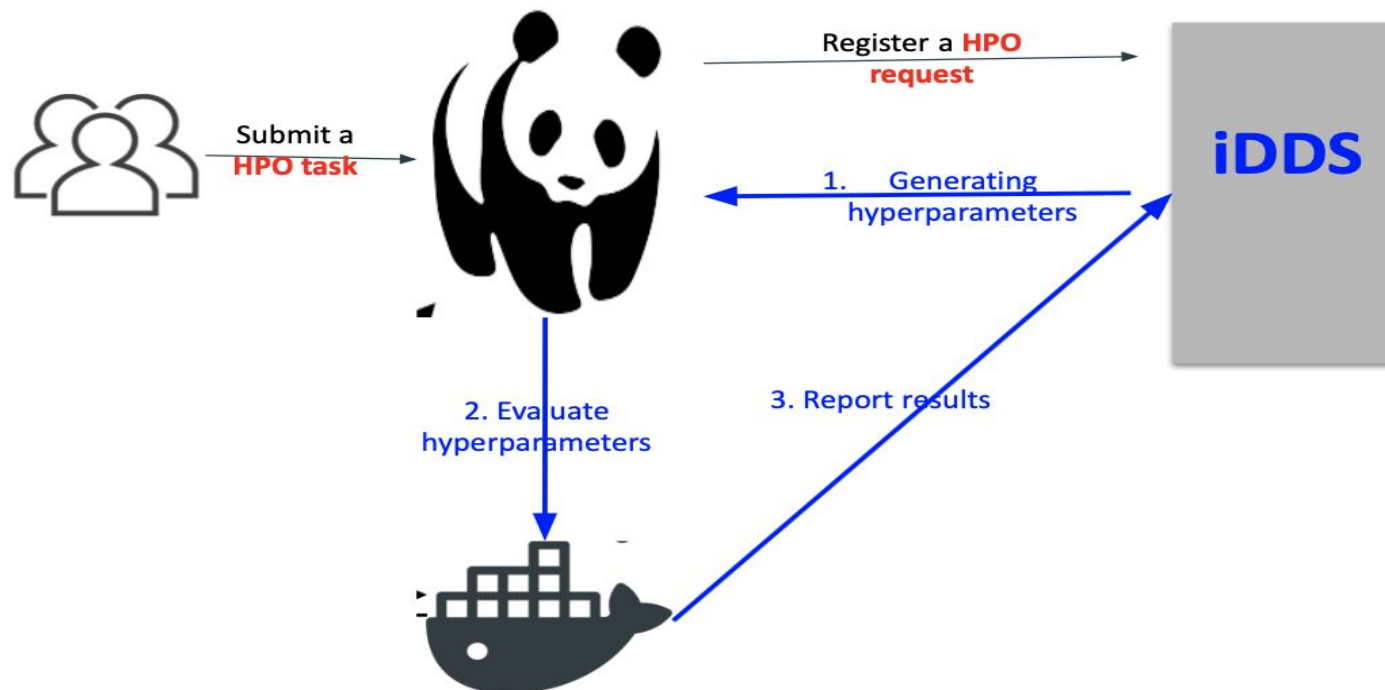
- Orchestrates Rucio to collect and digest file information, and lets JEDI/PanDA process only prestaged files with proper granularities and grouping.
- In ATLAS production since May 2020
- Has processed in total about 21 PB data.
- Reduced a lot of redundant job attempts.
 - With iDDS, the tail of jobs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



X: attempt times, Y: number of jobs

iDDS HPO(Hyper Parameter Optimization)

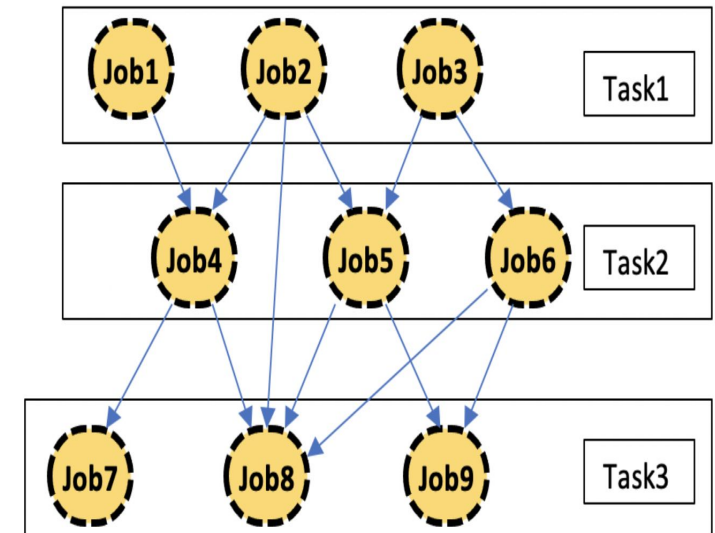
- To provide a fully-automated platform for hyper parameter optimization on top of geographically distributed GPU resources on the grid, HPC, and clouds.
- Used by ATLAS ML users, not specific to ATLAS.
- The usage is increasing. Different use cases are using the HPO framework to automate distributed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



iDDS Rubin Observatory (LSST)

- **LSST exercise**

- Use the experiment-agnostic DOMA PanDA instance to schedule jobs to Google Cloud.
- 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 30K DAG jobs, scaling test with 250K jobs is running

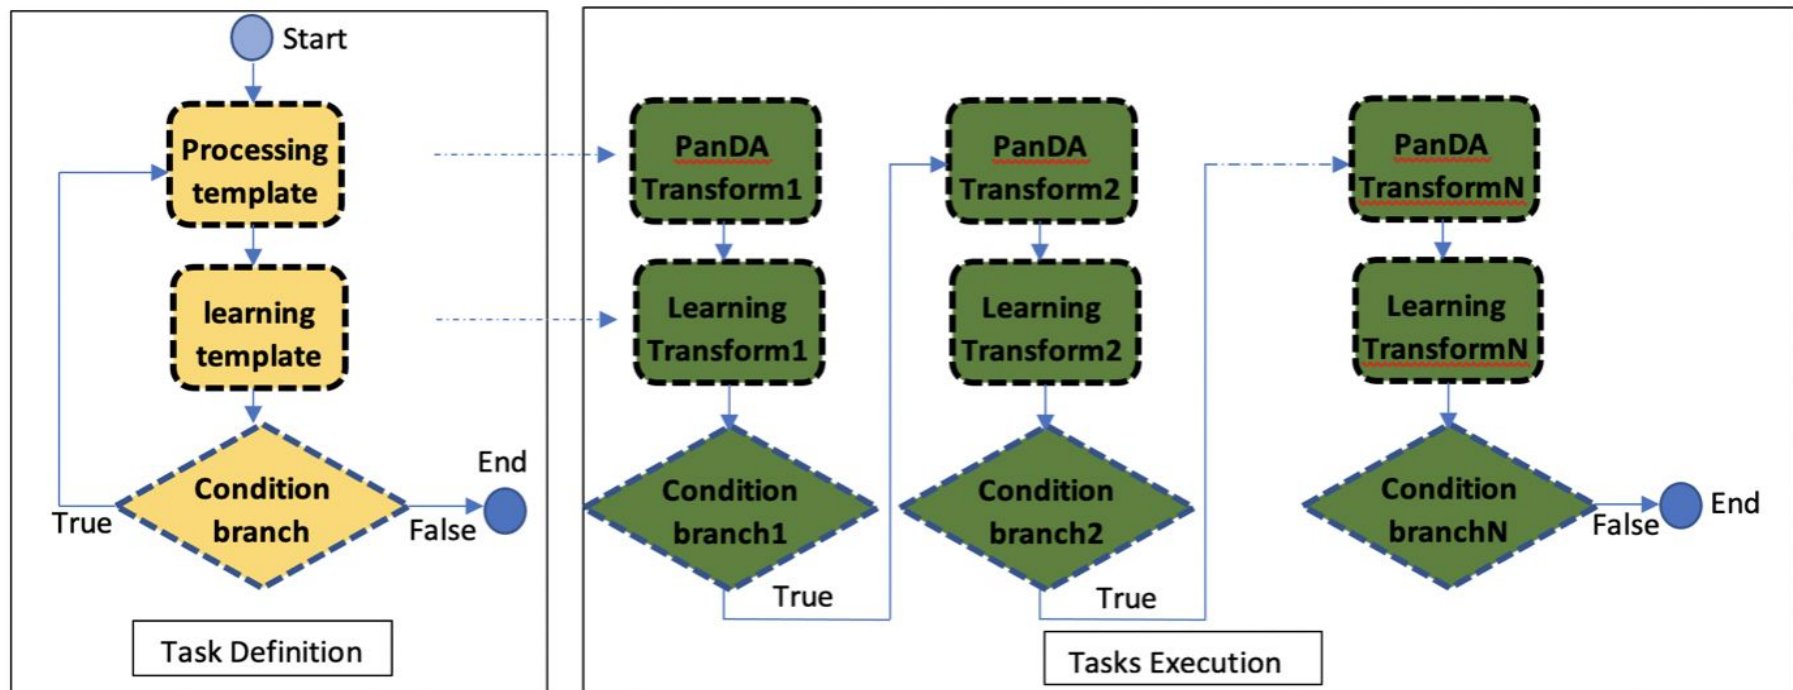


request_id	r_status	created_at	total_tasks	tasks	250K scaling tests	remaining_files	processed_files	processing_files	total_files
292	transforming	2021-05-01 09:19:12	24	Submitted(17) Running(5) Finished(2)		119133	119273	119133	238406
290	finished	2021-04-30 23:35:38	24	Finished(24)		0	143	0	143
272	finished	2021-04-14 16:59:02	15	Finished(15)	30K scaling tests	0	30987	0	30987
270	failed	2021-04-14 16:44:18	3	Failed(3)		6	8	6	14

iDDS ActiveLearning

- **ActiveLearning**

- A simple Task Level 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 analyzes 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

Backups

iDDS Rubin scaling test status

Show entries

Search:

request_id	r_status	created_at	total_tasks	tasks	remaining_files	processed_files	processing_files	total_files
292	transforming	2021-05-01 09:19:12	24	Submitted(17) Running(5) Finished(2)	119133	119273	119133	238406
290	finished	2021-04-30 23:35:38	24	Finished(24)	0	143	0	143
286	failed	2021-04-30 15:42:06	24	Finished(1) Failed(23)	205732	32674	205732	238406
284	transforming	2021-04-29 15:38:19	24	Submitted(23) Failed(1)	142	1	142	143
282	finished	2021-04-27 18:30:21	24	Finished(24)	0	143	0	143
280	transforming	2021-04-25 02:15:23	24	Submitted(7) Finished(16) Failed(1)	34	109	34	143
276	transforming	2021-04-24 20:07:22	24	Submitted(11) Finished(12) Failed(1)	40	103	40	143
274	finished	2021-04-24 16:01:44	3	Finished(3)	0	14	0	14
272	finished	2021-04-14 16:59:02	15	Finished(15)	0	30987	0	30987
270	failed	2021-04-14 16:44:18	3	Failed(3)	6	8	6	14

Showing 1 to 10 of 30 entries

Previous **1** 2 3 Next