

*TIDC 2023 Annual Meeting
on Nov/25th, 2023*

COMPUTING REPORT

*Kai-Feng Chen
National Taiwan University*



OVERVIEW

- **Computing is an essential service for HEP** — no matter what kind of experiments or detectors, we always need sufficient computing power to process, to store, and to analyze data!
 - Resources and knowledge required for modern system operations are non-trivial — need to consider space/power/cooling/knowhow/*people*.
 - Optimal is to have a centralized system instead of clusters running individually, not to add lots of management loads to each group.
- Thanks to the supports from ASGC, now **a new Tier-3 cluster under the name of TIDC** has been initiated, mainly for user analyses.
 - Seeding cost allocated under TIDC project.
 - Computing equipments injected from individual grants (George's summit, grants from Ming & myself, etc).
- ATLAS grid running by ASGC & CMS Tier-2 running at NCHC are considered as separated projects so far.

Many thanks for all the supports!

COMPUTING SUPPORTS @ ASGC

- Currently ASGC is operating the **NSTCCore** services to support general HPC users & applications:
 - CPU: 2976 cores (*to be increased in the next years*), GPU: V100×32+A100×8.
 - Working space: Ceph filesystem (3 Tb / group).
 - Batch jobs: Slurm management system, with the entry UI: slurm-ui.twgrid.org.
 - Interactive jobs with web-based UI via dicos.grid.sinica.edu.tw

It is possible to start a Jupyter notebook with GPU supports.

The screenshot displays the DiCOS Jupyter Lab interface. At the top, there is a navigation bar with links for DiCOS, Resources, Policy, About, Documentation, Apps, Contact, and Live Chat. Below the navigation bar, the word "Jupyter" is prominently displayed. The main content area features a grid of six Jupyter Lab environment cards. Each card includes the Jupyter logo, the environment name, its version, the GPU/CPU configuration, and the resource usage percentage. A green progress bar indicates the resource usage, and a "Launch" button is provided for each environment.

Environment Name	Version	Resources	Usage
Jupyter Lab	Version: CPU with Tensorflow v1	Resources: 100%	100%
Jupyter Lab gpu 3090	Version: GPU with Tensorflow 3090	Resources: 58%	58%
Jupyter Lab GPU 1080ti	Version: GPU with Tensorflow	Resources: 14%	14%
Jupyter Lab GPU V100	Version: GPU with Tensorflow V100	Resources: 22%	22%
Jupyter Lab GPU A100	Version: GPU with Tensorflow A100	Resources: 75%	75%
Jupyter Lab Cryocare GPU	Version: GPU with 1080ti	Resources: 11%	11%

COMPUTING SUPPORTS @ ASGC (II)

- ASGC is welcoming new users — in order to use the resources, **group PI has to create a group account first** and then ask the users to register;
 - PI has the privilege to monitor the user activities (*and will receive the cost bill monthly.*)
- Group/user accounts creation:
 - <https://canew.twgrid.org/ApplyAccount/groupcreate.php>
 - <https://canew.twgrid.org/ApplyAccount/ApplyAccount.php>
- TIDC cluster is managed under the same dicos system — if you would like to use TIDC CPU/storage, you have to apply the same accounts too.
 - At this moment people affiliated to NTU and NCU should fill Kai-Feng Chen and Chia-Ming Kuo as PI.

TIDC CLUSTER OVERVIEW (II)

➤ Software access:

- CVMFS is available — so basically all the CERN related libraries/tools should be accessible directly.
- We have only tested CMS software (CMSSW) and analysis frameworks (RDataFrame/Coffea) so far, but the general analysis flow should be similar for other HEP projects too.


➤ Storage access:

- Grid EOS storage is accessible by xrootd:
<root://tidc-smstor1.grid.sinica.edu.tw/eos/>
 - ➔ Or via fuse mount under UI: `/eos`
- Private EOS area: `/eos/tidc/<group_name>` (50 Tb / group)
- Working space: `/ceph/work/<group_name>` (3 Tb / group).

ps. This is basically too small for modern analyses, PI can ask for more when needed.

HOW TO USE?

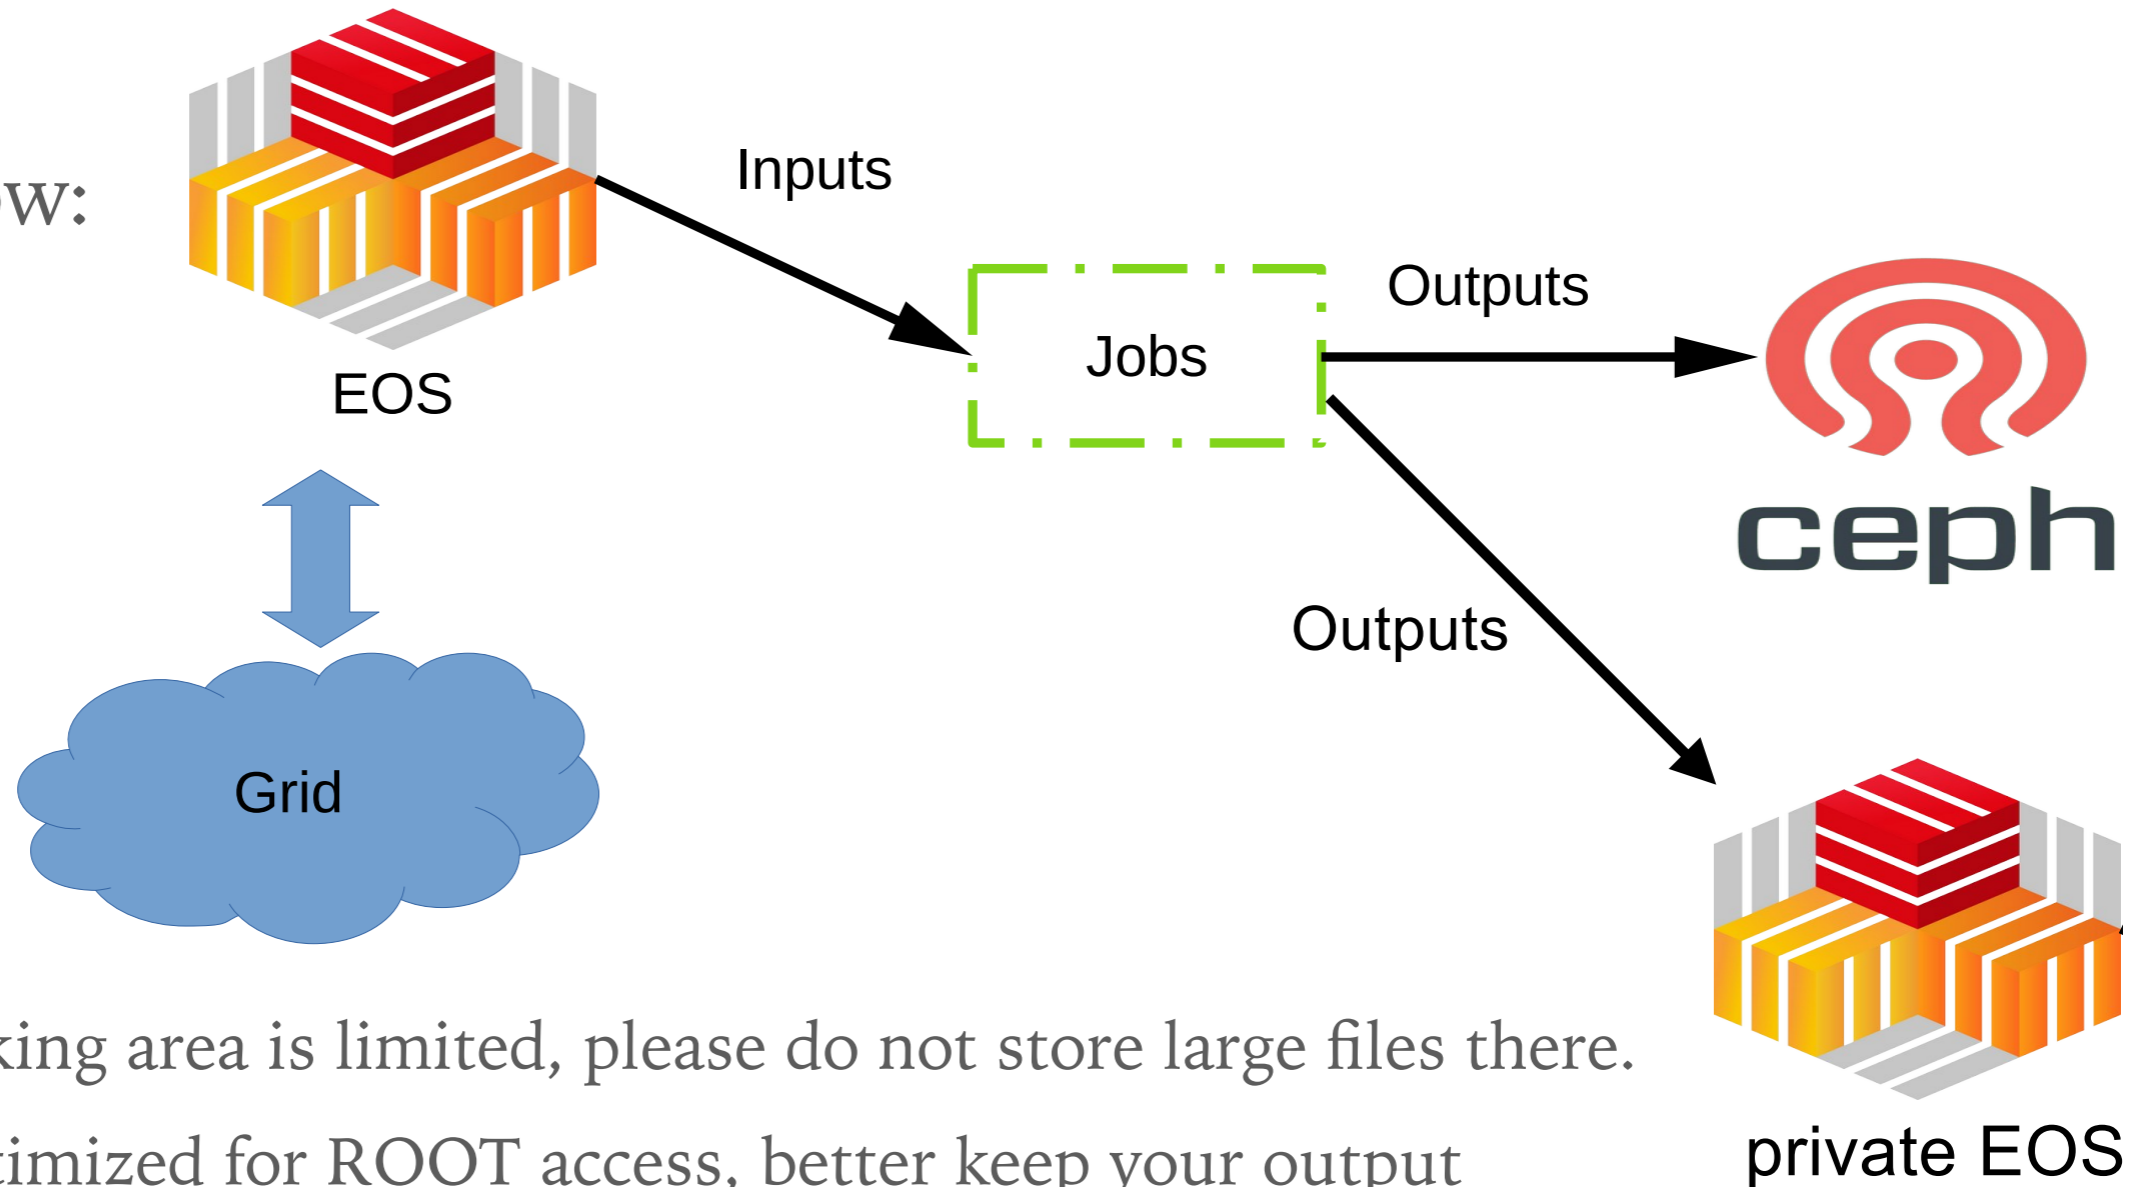
- ASGC is running hands-on tutorials every 3 months. Latest one was held at NTU (*many thanks to ASGC & TIDC again!*):
<https://indico4.twgrid.org/event/35/>
with a special session for TIDC cluster.
- Topics of interests:
 - Access to TIDC & Condor scheduler (by Felix Lee):
<https://indico4.twgrid.org/event/35/#b-595-hands-on-computing-servi>
 - General CMS software setup & condor jobs (by You-Ying Li):
<https://indico4.twgrid.org/event/35/#b-597-hands-on-analysis-framew>
 - Running analysis with RDataFrame (by Cheng-Han Wu) & Coffea (by Yu-Hsuan Chou):
<https://indico4.twgrid.org/event/35/#b-587-computing-service-for-he>



*Z → e⁺e⁻ peak generated
by Coffea running on
TIDC system, using
CMS open data.*

ANALYSIS WORKFLOW

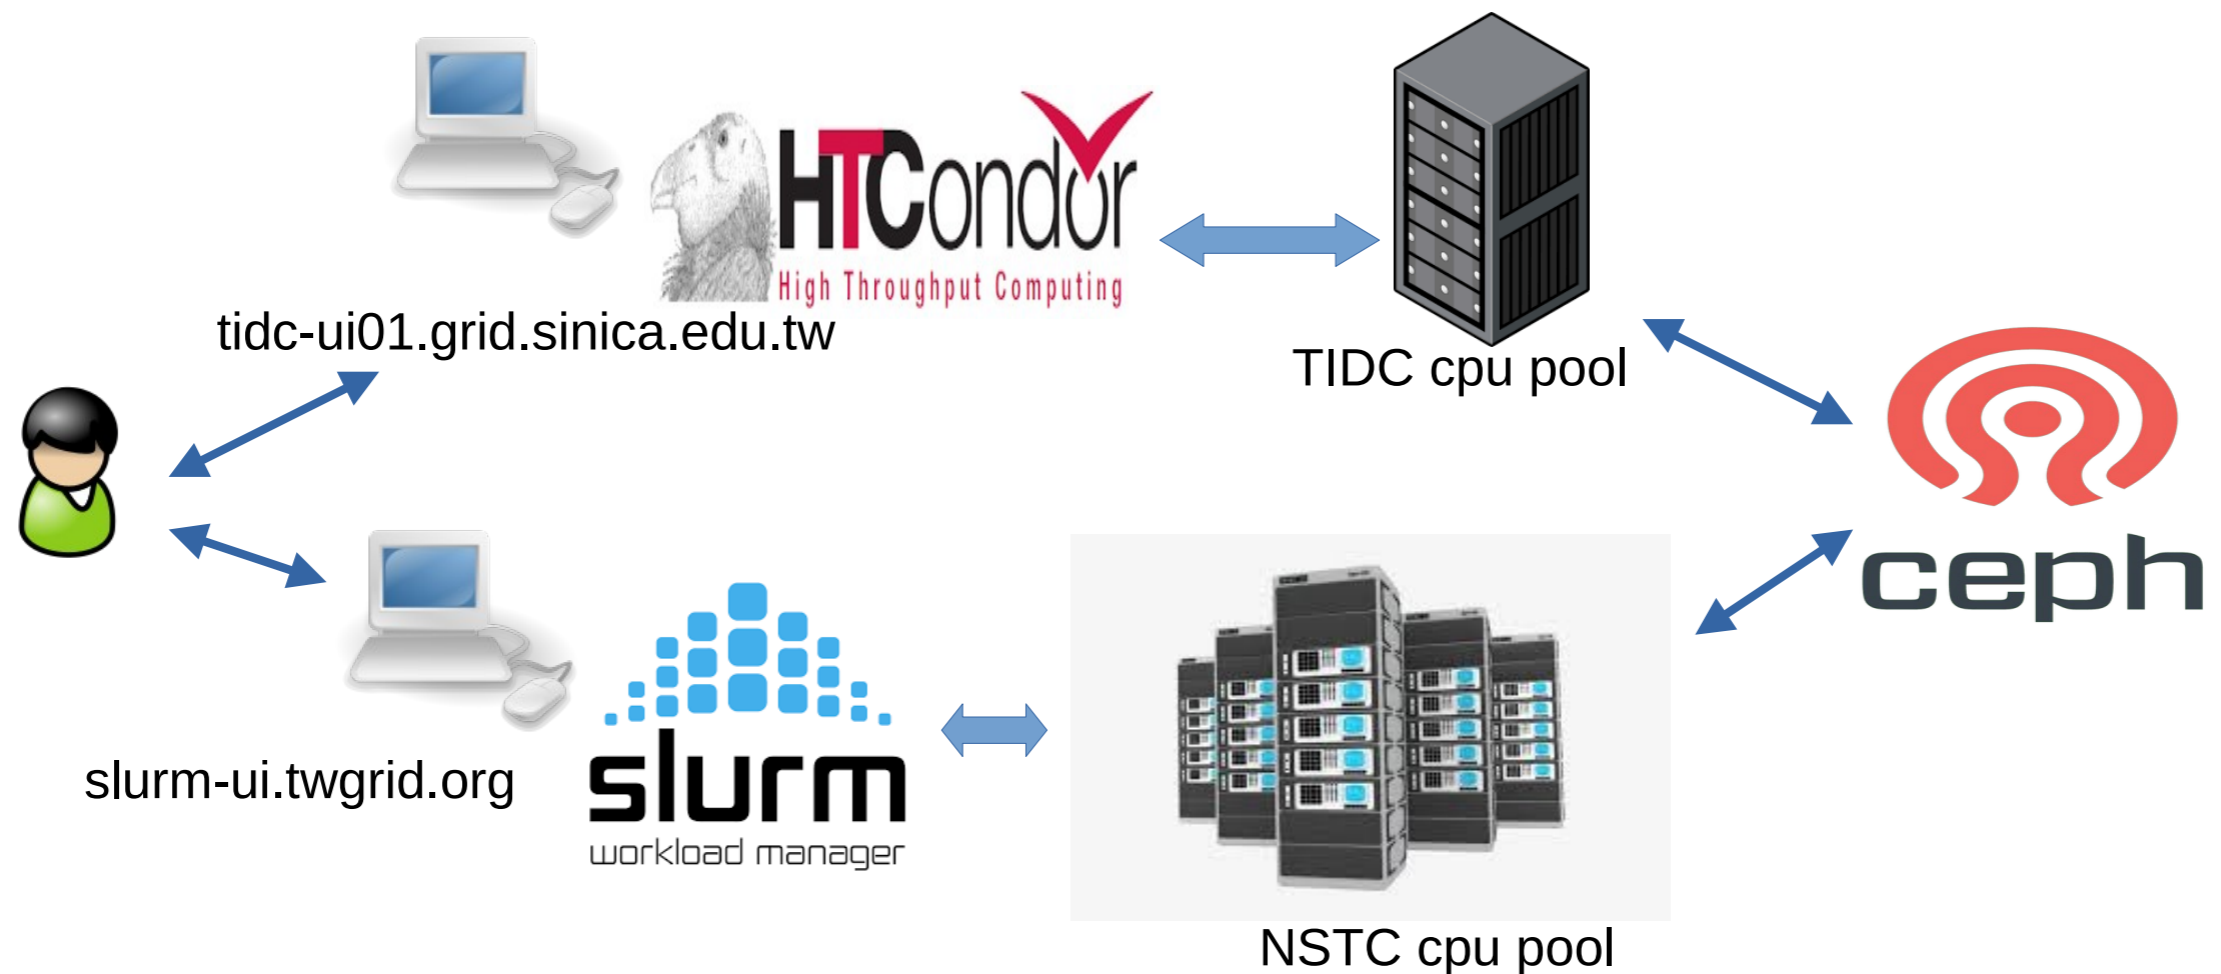
➤ Typical workflow:



- Ceph working area is limited, please do not store large files there.
- EOS is optimized for ROOT access, better keep your output ROOT files there if needed.
 - ➔ It is a bad idea to write your private files back to grid EOS (might be killed by global grid management), please use the private EOS area instead.

ANALYSIS WORKFLOW

- Integration with NSTCCore services:
 - It is possible to use the CPU/GPU allocated under NSTCCore too (as the Dicos user area & Ceph workspace are shared).
 - No direct access to experimental software but possible to run general analysis tasks (e.g. statistical studies, ML training, etc).



SUMMARY & PROSPECTS

- Many thanks to ASGC's support, now we have a **new TIDC cluster up and running**.
 - Please apply group/user accounts to access the resources.
 - Now it is CMS compatible, but it should be able to run the analysis jobs for other experiments as well (*CERN projects should be straightforward*) — further integrations can be discussed!
- **Only one UI at this moment — please be gentle / not to overload it!** (*more UIs to be built in the near future*)
- Expected to double the specs based on the summit grant; NCU resources to be integrated into the same pool.



**HAPPY RUNNING
ANALYSIS JOBS!**