

---

# Meeting ATLAS-IT calcolo : highlights from S&C week

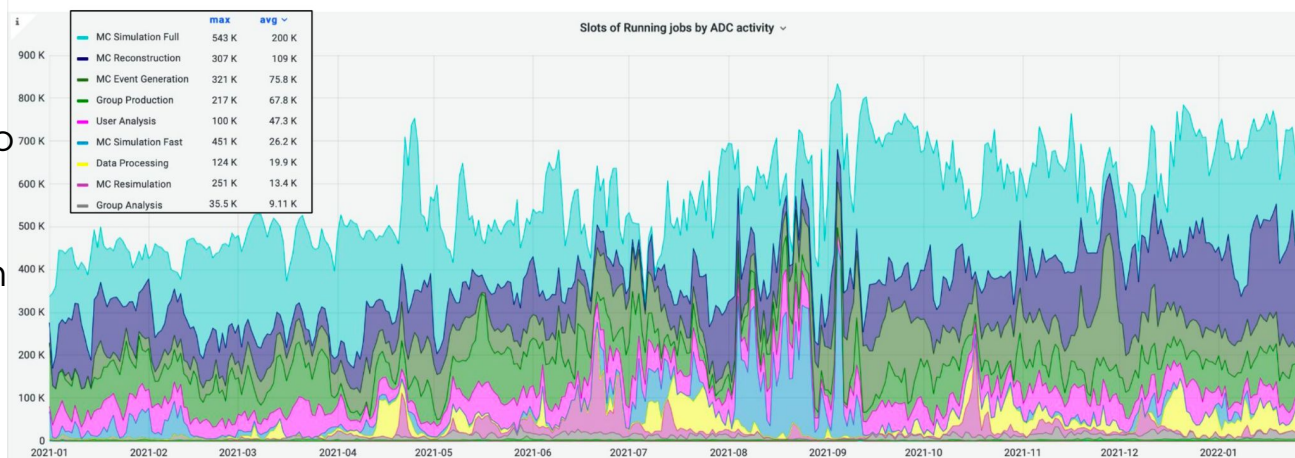
(A. Doria, L. Carminati)

## Highlights from last S&C: Computing

Outstanding performance of our distributed computing infrastructure : could run up to 800k jobs per day ( in many different workflows ).

- ❑ Run 2 data and MC reprocessing expected to finish this month

- ❑ Run 3 MC reconstruction for Combined Performance samples now underway

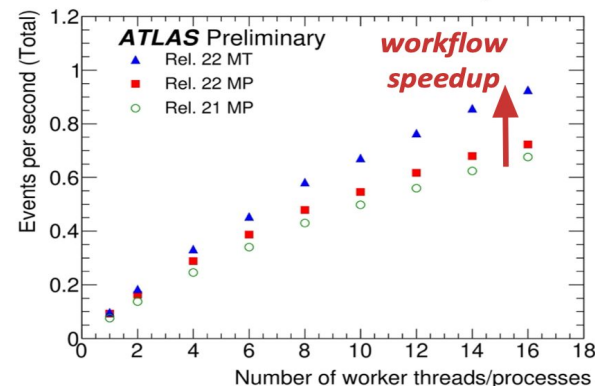
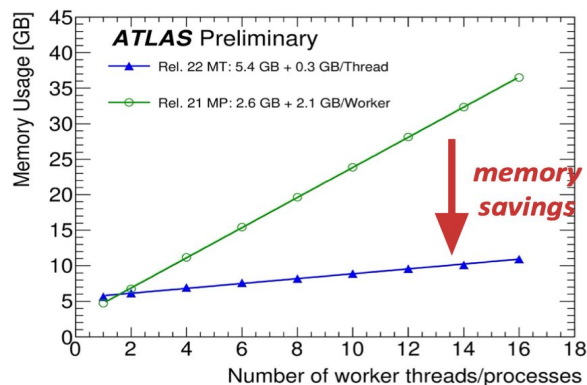


- ❑ HPCs continue to deliver at a feverish pace : ATLAS was able to run almost all data-flows
  - ❑ in the last few months HPC contributed up to 100% of the pledged computing resources
- ❑ Disk space is currently a critical item, basically reached to pledge, no contingency
- ❑ Data carousel mode running at full speed (both reco and derivation) : maximal usage of the tape resources

# Highlights from last S&C : Software

Continuous effort to improve the software performance. Several important achievements :

- ❑ First production relying on multi-threaded reconstruction (AthenaMT) : full reprocessing running at full speed
  - ❑ 3.5x memory savings
  - ❑ 20% workflow speed-up



- ❑ Impressive results from G4 optimisation task force : simulation up to 35% faster when including several optimisations, more to come
- ❑ New Analysis Model well advanced : two centrally produced formats
  - ❑ DAOD\_PHYS : ~30 kB/ev validated on Run2, ready for Run3
  - ❑ DAOD\_PHYSLITE : ~12 kB/ev currently being validated, crucial for HL-LHC



# Current Status of ATLAS (1)

[https://indico.cern.ch/event/1115413/contributions/4708335/attachments/2384202/4074331/Token\\_S%26C\\_20203-1.pdf](https://indico.cern.ch/event/1115413/contributions/4708335/attachments/2384202/4074331/Token_S%26C_20203-1.pdf)

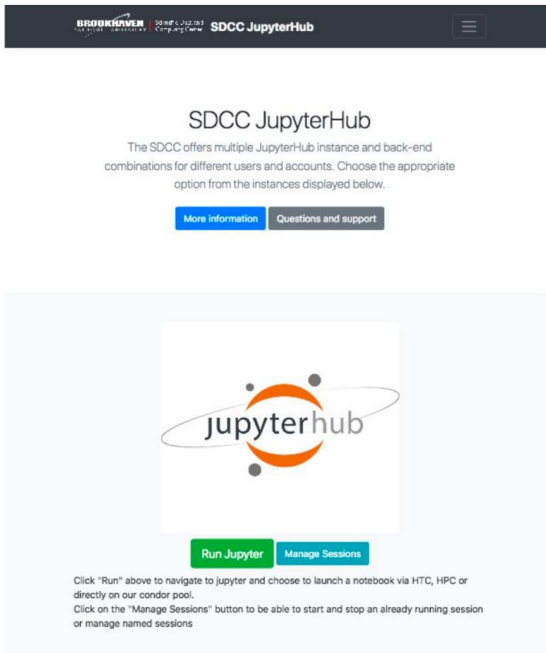
- Central Harvesters have been submitting with both voms proxy and ATLAS IAM token
  - Since 1 Nov. 2021: Met WLCG Token Transition Timeline M4 by end of Oct. 2021
  - “Central Harvesters”: Harvester\_ID = **CERN\_central\_A** , **CERN\_central\_B**
  - Central Condor instances upgraded to v9.0.6 (still supports GSI)
  - Harvester is configured to update tokens for all CEs on CRIC every 30 minutes
    - Token lifetime was originally 1 hr, while upcoming ATLAS IAM change will allow 6 hr
  - Harvester picks up the corresponding token file according to the CE chosen to submit to
  - If the CE (HTCondor-CE) supports token, it will authenticate with the token
    - Otherwise, will fallback to voms proxy
    - Harmless to ARC CEs or HTCondor-CEs without token setup
  - Currently, tokens with different IAM IDs are used to distinguish production vs analysis
  - Reference: [ATLAS talk in OSG Token Transition workshop in Nov. 2021](#)

# ATLAS Site Updates in 2022

<https://indico.cern.ch/event/1115413/contributions/4708346/attachments/2384056/4074544/ATLAS%20S%26C%20Week%20%2371%20-%20Timeline%20for%20grid%20site%20middleware%20updates.pdf>

- HTCondor-CE with [token support](#) – by June 2022
  - OSG 3.6 HTCondor-CE with default mapping to local usatlas{1,2,3} user
    - OSG have their own timeline given by OSG 3.5 EOL
  - Non-OSG HTCondor-CE 5.x/6.x(?) + [user mapping configuration](#)
- ARC-CE with REST support – by June 2022
  - 70% already OK, the [rest not reachable](#) (firewall?)
- [ATLAS VOMS updates](#)
  - Add ["vomses" for new ATLAS IAM VOMS](#) – January 2022
  - Next steps (removing "legacy" VOMS) – Autumn 2022
- Storage
  - Support for GridFTP protocol become optional in March 2022
    - Be aware we may still need / use SRM for space reporting at some sites
  - Updates to get production quality WLCG JWT token support
    - We may come with recommendation already at the end of 2022
    - The goal is to be ready for WLCG Data Challenge in 2023

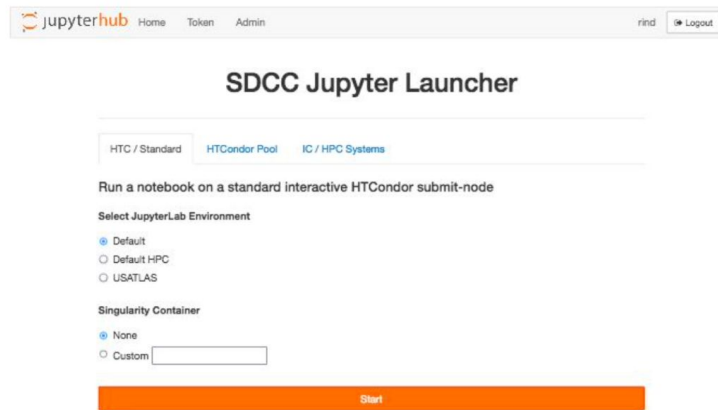
## BNL SDCC Jupyterhub Portal



The screenshot shows the BNL SDCC JupyterHub portal. At the top, there's a dark header with the BNL logo and 'SDCC JupyterHub'. Below the header, the text 'SDCC JupyterHub' is displayed, followed by a paragraph: 'The SDCC offers multiple JupyterHub instance and back-end combinations for different users and accounts. Choose the appropriate option from the instances displayed below.' There are two buttons: 'More information' and 'Questions and support'. Below this, there's a large white box containing the JupyterHub logo. At the bottom of this box, there are two buttons: 'Run Jupyter' and 'Manage Sessions'. Below the box, there's a paragraph: 'Click "Run" above to navigate to jupyter and choose to launch a notebook via HTC, HPC or directly on our condor pool. Click on the "Manage Sessions" button to be able to start and stop an already running session or manage named sessions'.

Custom Jupyterhub interface to multiple resources backends with account-based access control:

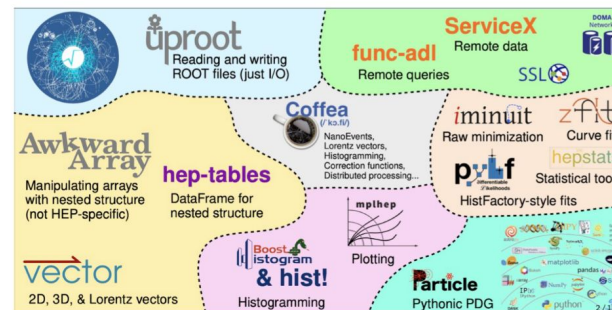
- Dedicated set of Jupyterhub nodes
  - Dask scheduling onto HTCondor pools
- Notebook spawning onto HTCondor pools
- Notebook spawning onto HPC cluster via Slurm



The screenshot shows the SDCC Jupyter Launcher interface. At the top, there's a header with the JupyterHub logo and navigation links: 'Home', 'Token', 'Admin'. There's a 'Logout' button. Below the header, the title 'SDCC Jupyter Launcher' is displayed. There are three tabs: 'HTC / Standard', 'HTCondor Pool', and 'IC / HPC Systems'. Below the tabs, the text 'Run a notebook on a standard interactive HTCondor submit-node' is displayed. There's a section 'Select JupyterLab Environment' with three radio buttons: 'Default' (selected), 'Default HPC', and 'USATLAS'. There's a section 'Singularity Container' with two radio buttons: 'None' (selected) and 'Custom' (with a text input field). At the bottom, there's a large orange button labeled 'Start'.

## Advanced Tools

- [Coffea-Casa](#) - low latency columnar analysis based on Dask and Jupyter
- [ServiceX](#) - data extraction and delivery
- [REANA](#) - platform for reusable analysis
- [RedHat OKD cluster](#) - k8s platform for current Reana and ServiceX deployments on test cluster at BNL
  - Chosen for use at DOE labs (BNL, FNAL) due to stronger security policies than k8s, but requires additional development effort



Close collaboration with IRIS-HEP!

<https://indico.cern.ch/event/1121552/contributions/4708724/attachments/2383561/4072980/AMGPerspective.pdf>

### Tape challenge '22 (1/4)

- We are planning tape challenge '22, again with multi-VOs
- The dates ?
  - Coordinate with detector commissioning
    - SFO available in the 3rd week of March for the tape challenge
    - So either the 2nd+3rd week or the 3rd+4th week of March
  - Coordinate with other experiments
    - To be discussed at the joint meeting with other experiments
- Again to cover both modes : data-taking (DT) mode and after-data-taking (A-DT) mode, one after the other

# Highlights from last S&C : dismissione DPM

Distributed Computing				
	MID	DID	Description	Due
	DC-1		Transition to tokens	Q4 2023
		1.1	Submission from Harvester to all HTCondor CEs with tokens	Q4 2022
		1.2	All users move from VOMS to IAM for X509	Q4 2022
		1.3	All job submission and data transfers use tokens	Q4 2022
	DC-2		Storage evolution	Q4 2025
		2.1	No GridFTP transfers at any site	Q1 2022
		2.2	SPM loss access to tape	Q4 2025
		2.3	Recommended transition plan from DPM completed	Q4 2021
		2.4	Transition plan from all DPM sites	Q4 2022
		2.5	All sites moved away from DPM	Q2 2021
	DC-3		Next operating system version	Q4 2022
		3.1	Ability to run on "future OS" on grid sites	Q4 2021
		3.2	Central services moved to "future OS"	Q1 2023
		3.3	(CentOS 7/8 EOL)	Q2 2024
	DC-4		Network infrastructure ready for Run 4	Q4 2027
		4.1	Network challenge at 10% expected rate	Q4 2021
		4.2	Network challenge at 30% expected rate	Q4 2023
		4.3	Network challenge at 60% expected rate	Q4 2025
		4.4	Network challenge at 100% expected rate	Q4 2027

WLCG sites should plan migrations before Centos 7 EOL (mid 2024)

○ From Centos 7 to 9 seems

promising during the period 2H

2022-2023 subject to middleware

availability

■ CentOS Stream 9 already available, RHEL/ELC 9s coming in the next few months

■ [ADC: would prefer sites to do such upgrades in winter shutdowns]

○ Centos 8 may also be used already as needed for hardware compatibility or functionality

WLCG sites should plan migrations before Centos 7 EOL (mid 2024):

- ❑ From Centos 7 to 9 seems promising during the period 2H 2022-2023 subject to middleware availability
  - ❑ CentOS Stream 9 already available, RHEL/ELC 9s coming in the next few months
  - ❑ [ADC: would prefer sites to do such upgrades in winter shutdowns]
- ❑ Centos 8 may also be used already as needed for hardware compatibility or functionality