LORENZO RINALDI

HPC & NEW ARCHITECTURES

OUTLINE

- Early step on LEONARDO
- HPC extension model
- ARM@CNAF
- Other activities

ATLAS EXPERIMENT

FIRST (PRE-)STEP INTO LEONARDO

It is now possible to request interactive access to Leonardo

(https://enccs.se/news/2023/09/how-to-login-to-leonardo-supercomputer/)

–> User that already have an account could ask to T. Boccali for project association

Two-factor authentication needed for ssh login

The big problems:

- **no** CVMFS and **isolated** nodes (cvmfs-exec doesn't work…)

- 2FA might be limiting for some automatic procedure

For the time being: try to do some test with stand-alone sw (CaloGAN training…)

ATLAS EXPERIMENT

HPC EXTENSION MODEL

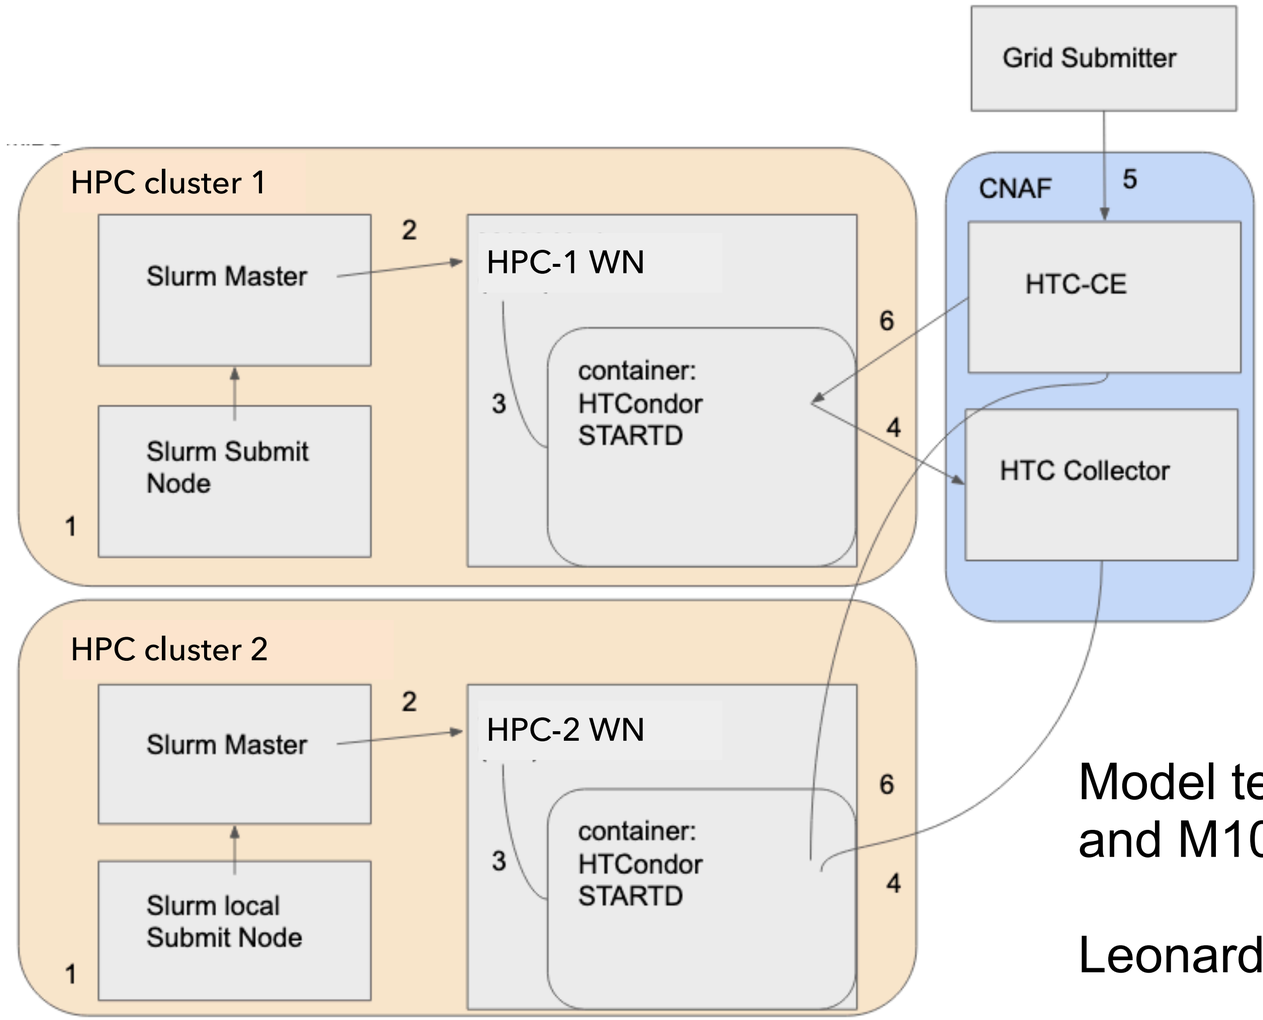

1. Owned script submit Slurm Jobs
2. The Slurm job launches a singularity container
3. The container starts a HTCondor STARTD
4. The STARTD has token credentials to join the HTC pool at CNAF
5. Jobs requiring HPC resources are properly routed and queued
6. for execution in the HPC node

Model tested with CINECA Marconi and M100, UniBO cluster.

Leonardo?

ATLAS EXPERIMENT

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
ICSC Centro Nazionale di Ricerca in HPC, Big Data and Quantum Computing

ARM@CNAF

Validation of event reconstruction code on ARM
Spoke 2 use case

Spoke	2
WP	2, 5
Use case short name	ARM_ON_DATALAKE
Use case ID	UC2.2.5
Expected Completion	31/8/2025

Participating Institutions

- Leader: INFN (Francesco Noferini, Daniele Spiga, Tommaso Boccali, Lucio Anderlini, Concezio Bozzi)
- Participants: INFN, UNIBO
- Experiments: ALICE (F. Noferini), CMS (D. Spiga, T. Boccali), ATLAS (L. Rinaldi, L. Carminati), LHCb (L. Anderlini, M. Veltri)

Available resources

There are 2 ARM nodes (Almalinux9)

- 256 cores
- 512 GB ram

Current setting (still work in progress)

- Interactive access provided in one of the two nodes (wn-arm-02)
- HTCondor submission enables on wn-arm-01
- Network: Cvmfs available and access to external network
- Gpfs client -> not yet available for ARM

- 2 more nodes acquired with 1TB ram (ram extension foreseen also for the two nodes already in production, now at 512 GB)
- 2 nodes NVIDIA acquired
  - Grace -> only CPU ARM (144 cores)
  - GraceHopper → CPU (72 cores) and GPU (Nvidia H100) on single chip

ATLAS EXPERIMENT

ARM@CNAF

End October 23 - Successfully tested some SIM jobs (interactive with centos7 container from cvmfs)

20/11 - Setup of a new Panda Queue **INFN-T1_ARM**: thanks to Pascolini (CNAF-farming) and Rod, Ivan and Asoka

- First HC jobs failed (missing --alma9 flag for picking up el9 container, now fixed)

- Rod forced a small full-chain task—> that was OK

ATLAS EXPERIMENT

ARM@CNAF

Number of Successful and Failed jobs

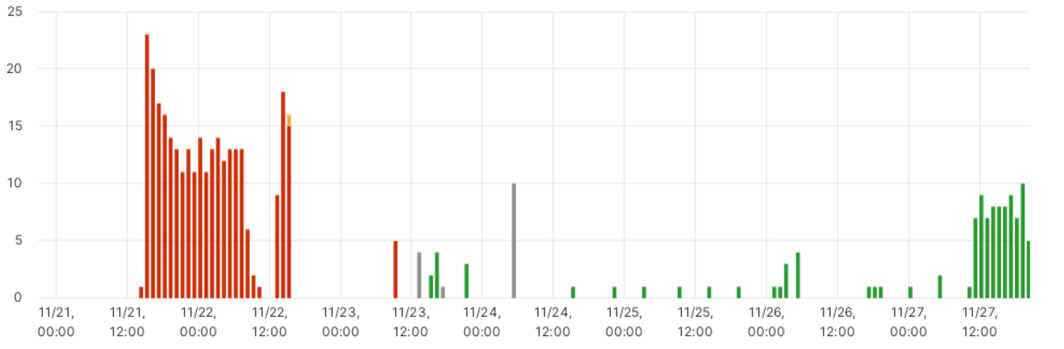

	min	max	avg	total
failed	0	23	1.77	298
finished	0	10	0.649	109
closed	0	10	0.0893	15
cancelled	0	1	0.00595	1

Slots of Running jobs

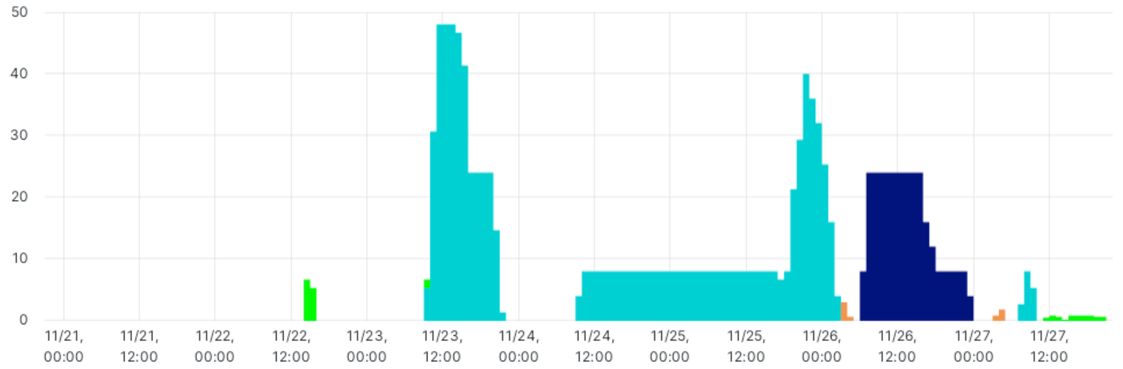

	min	max	avg
MC Simulation Full	0	48	5.13
MC Reconstruction	0	24.0	1.75
Testing	0	6.67	0.121
MC Merge	0	3	0.0375

Completed jobs

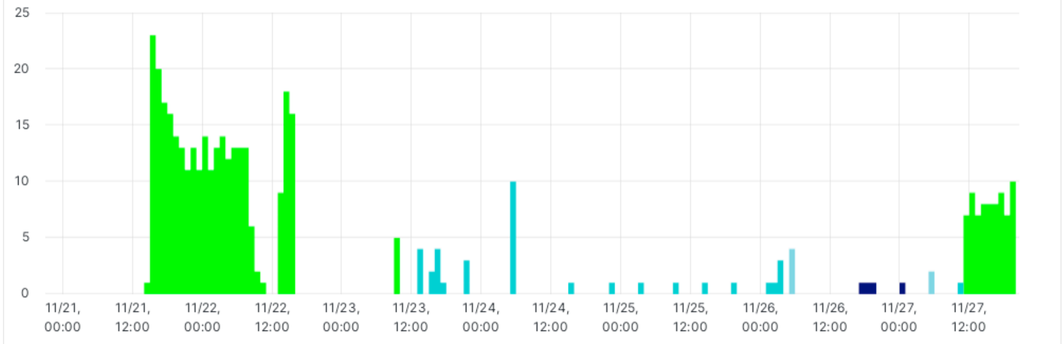

	min	max	avg	total
Testing	0	23	2.24	377
MC Simulation Full	0	10	0.214	36
MC Merge	0	4	0.0357	6
MC Reconstruction	0	1	0.0238	4

CPU/Wall Efficiency: Successful jobs

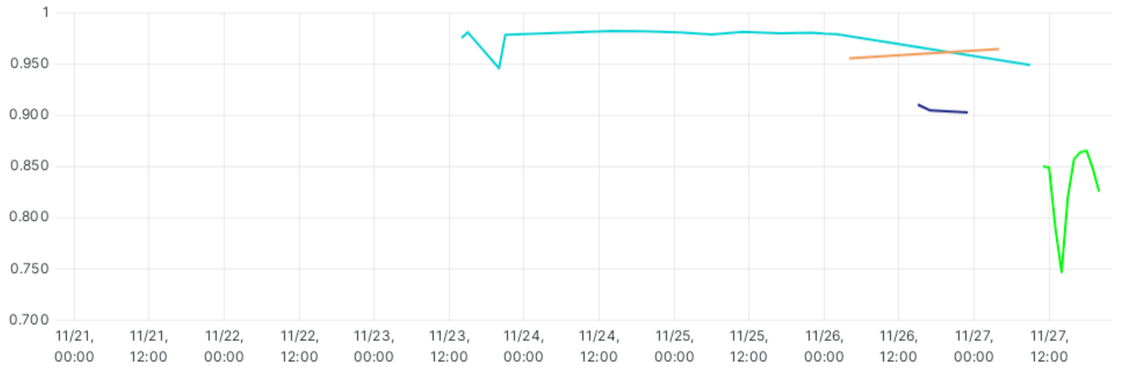

	min	max	avg	current
MC Simulation Full	0.946	0.982	0.975	0.949
MC Merge	0.955	0.964	0.960	0.964
MC Reconstruction	0.903	0.910	0.906	0.903
Testing	0.747	0.865	0.832	0.826

ARM@CNAF: NEXT STEP AND IDEAS

- ATLAS SW already validated on ARM for production workflow
- Coordinate activities with ATLAS and other exp (knowledge sharing)
  - Try to run different workflow (ML) and test new nodes with GPUs
  - Is there a way to quantify the performance in an ARM node with a direct/ indirect measurement of the **power consumption**? To be investigating with the farming team

ATLAS EXPERIMENT

OTHER ACTIVITIES

- FastCaloGAN (see Federico's talk) –> CN-HPC -spoke2- WP2 - FlashSim flagship
- PhD student in Data Science and Computation (Giacomo Levrini) interested in Big Data and Analytics platforms (now working at CNAF, ATLAS qualified, try to involve in ATLAS comp activities + period at CERN)
- Recent collaboration with Huawei for a PhD co-funding (PNRR)
  - New PhD student coming from India
  - Research topic: study and development of ML/AI algorithms for enhancing discovery potential in HEP (not-ATLAS, also technical part on data compression)

ATLAS EXPERIMENT