



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# HEP software validation on ARM: status and next steps



F. Noferini, L. Anderlini, T. Boccali, C. Bozzi, L. Carminati,  
L. Rinaldi, D. Spiga, M. Veltri

ICSC Italian Research Center on High-Performance Computing,  
Big Data and Quantum Computing

Missione 4 • Istruzione e Ricerca

# Current status

## The plan in view of MS8 (end of phase I)

- First period (tentatively month 13-22 - aligned with MS8): procure and configure ARM machines in order to provide access to the experiment software and storage via a production infrastructure; select and document workflows to be benchmarked from the most representatives; **prepare a validation strategy agreed with the experiments.**




A specific document was provided with MS8 report

## ARM validation strategy (TAR2.15)

The goal of the "Validation of event reconstruction code on ARM" flagship use case is to provide the hardware and the software infrastructure to enable a validation of the software used by experiments to process data and Monte Carlo (MC), at least for two major LHC experiments, using resources from the ICSC datalake and in particular hosted at the INFN CNAF Tier-1 facility.

Describing 3 use cases  
for phase II:

- ALICE
- CMS
- ATLAS



# Current status (II)

For all the experiments considered there are two main areas where validation can run:

- Full simulation (MC)
- Data reconstruction

Both cases are explored in all the 3 cases

Where we are now...  $\frac{2}{3}$  experiments are already in an advance status:

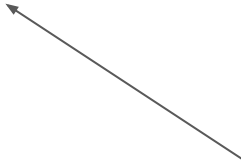
- ALICE (to be done)
- CMS (already advanced)
- ATLAS (ready to go)

# Recent presentations at workshops and conferences

2024 reports:

- INFN-CCR 2024 workshop (Palau), L. Rinaldi et al., [“Utilizzo di risorse con architetture ARM negli esperimenti di LHC”](#)
- Congresso SIF 2024 (Simposio supercalcolo, Bologna), F. Noferini et al., [“Uso di nodi ARM nella fisica delle alte energie”](#)
- CHEP 2024 (Kracow), D. Lattanzio et al., [“Heterogeneous computing at INFN-T1”](#) →23/10/2024

Part of this talk will focus on activities at CNAF with ARM



# Milestones

- First period (ICSC month 13-22 - aligned with MS8): procure and configure ARM machines in order to provide access to the experiment software and storage via a production infrastructure; select and document workflows to be benchmarked from the most representatives; prepare a validation strategy agreed with the experiments.
- Second Period (ICSC month 23-36 - aligned with MS10): test and validate the selected workflows (most probably from data reconstruction and simulation); validate the submission infrastructure and perform  $O(1 \text{ week})$  exclusive tests as needed by the validation strategy. Report the results to the experiments and in the ICSC documentation; disseminate the results at topical conferences.

**We are now in phase II**

# ATLAS

Physics validation ready (test @ Glasgow)


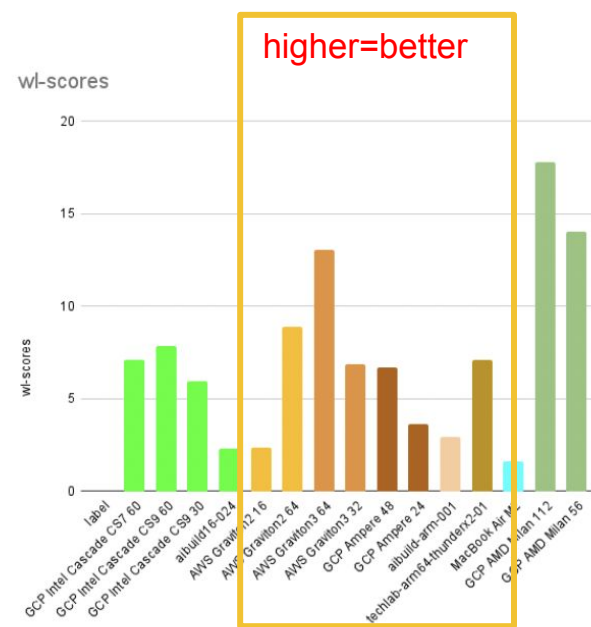
→ to be repeated at CNAF

How much (Tier-0 + other sites)?

- At peak: 2.6% world-wide pledge

@CNAF

- At peak: 12.5 kHS23
  - 9% INFN-T1 pledge



What (@CNAF)?

- MC FULL Simulation

How many events per HS23/hour?

- 3 events/hour/HS23

VERY PRELIMINARY!

- Results are comparable with other ATLAS measurements

# CMS

## The ARM nodes at T1\_IT\_CNAF have been integrated as a sub site of the regular Tier1 and thus accessed via GRID

- Minimize the effort on CMS and simplify the site admins life
- At the moment data are accessed via Xrootd protocol. The plan is to provide direct access also via GPFS.

## Technical validation fully done.

- Being the first allocation at a Tier site, the CNAF nodes where essential to finally validate that CMS Computing is multi-architectures enabled.

## Physics validation in progress:

- most of the subsystem reports green light (especially when looking at MC).
- While some discrepancies spotted on DATA require further analysis for a better understanding
- This step has been carried on both at CNAF and on Glasgow temporary allocation




RelVal	TICKETS	RELVALS	DASHBOARD	Logged in as Daniele Spiga
Preprod			<b>Workflows (jobs in ReqMgr2)</b>	
	<input type="checkbox"/>		<b>pdmlrexy_RVCMSSW_14_0_0_pre3Run3MET2020_CNAFARM_ReVal_202011_092819_7275</b> open in Status2 status: normal-archived	
			<b>FEVDEBAGHLT</b> completed: 159.27%, events: 1,296,791, type: <b>VALID</b>	
			<b>AGD</b> completed: 106.33%, events: 1,227,653, type: <b>VALID</b>	
			<b>MINIAD</b> completed: 106.33%, events: 1,227,653, type: <b>VALID</b>	
			<b>NANADQ</b> completed: 0.00%, events: 0, type: <b>VALID</b>	
			<b>DOMG</b> completed: 0.00%, events: 0, type: <b>VALID</b>	
			<b>FEVDEBAGHLT</b> completed: 123.17%, events: 788,261, type: <b>VALID</b>	
			<b>MINIAD</b> completed: 119.45%, events: 764,476, type: <b>VALID</b>	
			<b>NANADQ</b> completed: 120.25%, events: 769,595, type: <b>VALID</b>	
			<b>DOMG</b> completed: 0.00%, events: 0, type: <b>VALID</b>	
			<b>FEVDEBAGHLT</b> completed: 143.03%, events: 1,902,017, type: <b>VALID</b>	
			<b>AGD</b> completed: 135.98%, events: 1,809,262, type: <b>VALID</b>	
			<b>MINIAD</b> completed: 135.98%, events: 1,809,262, type: <b>VALID</b>	
			<b>NANADQ</b> completed: 82.25%, events: 671,245, type: <b>VALID</b>	
			<b>DOMG</b> completed: 0.00%, events: 0, type: <b>VALID</b>	
			<b>FEVDEBAGHLT</b> completed: 83.32%, events: 678,398, type: <b>VALID</b>	
			<b>AGD</b> completed: 82.25%, events: 671,245, type: <b>VALID</b>	
			<b>MINIAD</b> completed: 82.25%, events: 671,245, type: <b>VALID</b>	
			<b>NANADQ</b> completed: 82.25%, events: 671,245, type: <b>VALID</b>	
			<b>DOMG</b> completed: 0.00%, events: 0, type: <b>VALID</b>	

# ALICE

Preliminary tests showed good balancing in using resources but still some instability in ARM builds didn't allow to run a validation so far.

- Data reconstruction



- GRID setting tuned for 8 cores per node
- CPU efficiency consistent with what observed in the GRID node
- Physics validation on the output not yet done

Performance test in 2023 @E4 showed good performance both in data and simulation (full MC)

Recent builds on Almalinux9 presented some issues (mainly fixed) → tests are going to be resumed

GRID submission @CNAF (GRID submission + aptainer container) → validated



# Outlook and future plans (not requested in the validation strategy)

- An important piece is missing: measurement of **power consumption**
  - Define a strategy for the measurements and a good metric (e.i. evnt/sec/watt ?)
- The validation is on the way..
  - LHCb is performing test on Glasgow resources, CNAF will be used in next phase for DIRAC provisioning
- The use case would enable production workflows on ARM at the INFN CNAF center, a Tier-1 in the WLCG hierarchy and a node of the ICSC datalake
- On top of this, a successful validation would pave the way to more communities (and in particular smaller HEP and Astro experiments)
- In perspective, the path could enable the Italian Computing Infrastructure (INFN and ICSC, in this particular case) **to provision ARM machines in the near future**, with a sizable reduction of computing TCO.

backup

# People involved in this use case

## *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)

# KPI

## KPIs

KPI ID	Description	Acceptance threshold
KPI2.2.5.1	Software validation on ARM in the full GRID chain	50% (2/4 LHC experiments)
KPI2.2.5.2	Presentation at conferences	$\geq 2$
KPI2.2.5.3	Technical notes (in experiments and ICSC)	$\geq 2$

- **First period (tentatively month 13-22 - aligned with MS8): procure and configure ARM machines in order to provide access to the experiment software and storage via a production infrastructure; select and document workflows to be benchmarked from the most representatives; prepare a validation strategy agreed with the experiments.**
- Second Period (tentatively month 23-36 - aligned with MS10): test and validate the selected workflows (most probably from data reconstruction and simulation); validate the submission infrastructure and perform O(1 week) exclusive tests as needed by the validation strategy. Report the results to the experiments and in the ICSC documentation; disseminate the results at topical conferences.

# Available resources @ CNAF



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



- Resource procurement started in the second half of 2023 when CNAF acquired two ARM nodes allowing the experiments to logging on those machines and running preliminary tests. By the end of 2023 CNAF acquired two additional nodes and setup a GRID-HTCondor queue.
- 4x 2U Dual socket ARM – Ampere Altra Max
  - 2x Ampere AltraMax M128-30 128Core 2.8Ghz 250W
  - 1TB RAM
  - 2x NVMe U.3 3.84TB discs
  - 1x Dual port SFP28 ethernet

Current setting (still work in progress)

- Cvmfs available
- Network: access to external network
- Gpfs client -> now available on ARM but only with a very recent version, not yet in production at INFN-T1
- Condor/GRID -> in production

