



Finanziato  
dall'Unione europea  
NextGenerationEU



Ministero  
dell'Università  
e della Ricerca



Italiadomani  
PIANO NAZIONALE  
DI RIPRESA E RESILIENZA



# Validation of event reconstruction code on ARM Spoke 2 use case

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

# 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.

# MS7 Final Report

There is no explicit milestone in the project from MS7, hence we provide a status report.

The goal of this 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). Such a validation is expected to be done at least for two major LHC experiments, using resources from the ICSC datalake and in particular hosted at the INFN CNAF Tier-1 facility.

Resource procurement started in the second half of 2023 when CNAF acquired two ARM nodes allowing the experiments to login on those machines and running preliminary tests. By the end of 2023 CNAF acquired two additional nodes and setup a GRID-HTCondor queue.

CMS was the first LHC experiments which started to submit jobs following the standard GRID chain. At the end of January 2024 three of the LHC experiments were able to send standard GRID jobs, namely ALICE, ATLAS and CMS. See in the Figure 1 below an example taken from the CMS Dashboard (<https://cmsweb.cern.ch/wmstats/index.html>, access with experimental credentials required). CMS has also recognized the contribution from CNAF in a recent review at CERN (Figure 2).


At the moment data access is provided using Xrootd protocol but some work is ongoing to provide direct access also via GPFS. All the experiments mainly ran MC jobs but also release validation jobs on data can run in the same nodes, CMS already ran several tests.

Before moving to test extensively data jobs some optimizations are still needed to cope with the higher I/O required.

**Currently there is a discussion within the ARM flagship team to define, in view of MS8, the strategy for the validation of experiments' software.**

# MS7 Final Report (II)

## C-RSG Question #7: ARM Performance



We would really need to run real production workflows to answer this question. Fortunately, there are now ARM resources (~1K cores) newly available at CNAF. The core software team is being offered access. Thank you!

**[CMS-7] Page 13: Are the CPU efficiency of ARM based machines different from IA64 machines? And are there any notable learnings wrt. performance.**

CMS physics validation workflows, which were run on the ARM resources at Glasgow, are not optimized for CPU efficiency. We do not have any measurements running standard production workflows on ARM resources. We are, however, currently running a small validation campaign on ARM resources at the CNAF Tier-1 site. The results of the physics validation are not yet complete.

*We had an interesting discussion with the chief LHCC Referee about accessing ARM resources in **commercial clouds** as a way to perform validations on new architectures without waiting for our sites to purchase new machines. Are we (DRP?) interested in revisiting this capability? See two ATLAS CHEP presentations:*


- <https://indico.ilab.org/event/459/contributions/11636/>
- <https://indico.ilab.org/event/459/contributions/11553/>

Figure 2: CMS Computing responding to referees, and citing CNAF as the center used for ARM validation.

# CMS

## GRID

- All issues spotted in the previous reports **fixed** (xrootd access working fine after a fine tuning).
- In the last weeks technical validation of grace node was completed
  - → **ARM GRID instance declared as ready**



Up to ~500 cores used by concurrent jobs

# CMS

## GRID

- All issues spotted in the previous reports **fixed** (xrootd access working fine after a fine tuning).
- In the last weeks technical validation of grace node was completed
  - → **ARM GRID** instance declared as **ready**
- Ready to start a full validation (it will be done centrally)

RelVal	TICKETS	RELVALS	DASHBOARD	Logged in as Daniele Spiga
<input type="checkbox"/> Prepid				
			<b>Workflows (jobs in ReqMgr2)</b>	
			1. pdmvserv_RVCMSSW_14_0_0_pre3RunJetMET2023D_CNAFARM_ReVal_2023D_240215_092819_7275 open in: Stats2 status: normal-archived	
			◦  datar: FEVTDEBUHLT, completed: 100.72%, events: 1,265,761, type: VALID	
			◦  datar: AOD, completed: 106.33%, events: 1,227,653, type: VALID	
			◦  datar: MINIAOD, completed: 106.33%, events: 1,227,653, type: VALID	
			◦  datar: NANOAO, completed: 106.33%, events: 1,227,653, type: VALID	
			◦  datar: DQMIO, completed: 0.00%, events: 0, type: VALID	
<input type="checkbox"/> CMSSW_14_0_0_pre3_Data_2023_CNAFARM-RunJetMET2023D-00001			1. pdmvserv_RVCMSSW_14_0_0_pre3RunDisplacedJet2023D_CNAFARM_ReVal_2023D_240215_092843_9419 open in: Stats2 status: normal-archived	
			◦  datar: FEVTDEBUHLT, completed: 123.17%, events: 788,261, type: VALID	
			◦  datar: AOD, completed: 120.25%, events: 769,596, type: VALID	
			◦  datar: MINIAOD, completed: 119.45%, events: 764,476, type: VALID	
			◦  datar: NANOAO, completed: 120.25%, events: 769,596, type: VALID	
			◦  datar: DQMIO, completed: 0.00%, events: 0, type: VALID	
<input type="checkbox"/> CMSSW_14_0_0_pre3_Data_2023_CNAFARM-RunDisplacedJet2023D-00001			1. pdmvserv_RVCMSSW_14_0_0_pre3RunEGamma2023D_CNAFARM_ReVal_2023D_240215_092825_2415 open in: Stats2 status: normal-archived	
			◦  datar: FEVTDEBUHLT, completed: 143.03%, events: 1,903,017, type: VALID	
			◦  datar: AOD, completed: 135.98%, events: 1,809,262, type: VALID	
			◦  datar: MINIAOD, completed: 135.98%, events: 1,809,262, type: VALID	
			◦  datar: NANOAO, completed: 135.98%, events: 1,809,262, type: VALID	
			◦  datar: DQMIO, completed: 0.00%, events: 0, type: VALID	
<input type="checkbox"/> CMSSW_14_0_0_pre3_Data_2023_CNAFARM-RunEGamma2023D-00001			1. pdmvserv_RVCMSSW_14_0_0_pre3RunTau2023D_CNAFARM_ReVal_2023D_240215_092826_2548 open in: Stats2 status: normal-archived	
			◦  datar: FEVTDEBUHLT, completed: 83.12%, events: 678,398, type: VALID	
			◦  datar: AOD, completed: 82.25%, events: 671,245, type: VALID	
			◦  datar: MINIAOD, completed: 82.25%, events: 671,245, type: VALID	
			◦  datar: NANOAO, completed: 82.25%, events: 671,245, type: VALID	
			◦  datar: DQMIO, completed: 0.00%, events: 0, type: VALID	

Some official release validation workflows already running on the outputs of CMS jobs@ARM-CNAF


# ATLAS

## GRID

- In the last week reached a plateau with 8-core full-simulation
  - **12.6 kHS23 (at plateau)** → ~10% of ATLAS pledge at CNAF
- no major issues

## From previous report

- software ARM available on cmvfs
- Using container from cvmfs
- Full detector simulation -> validated
- NN training -> in setup





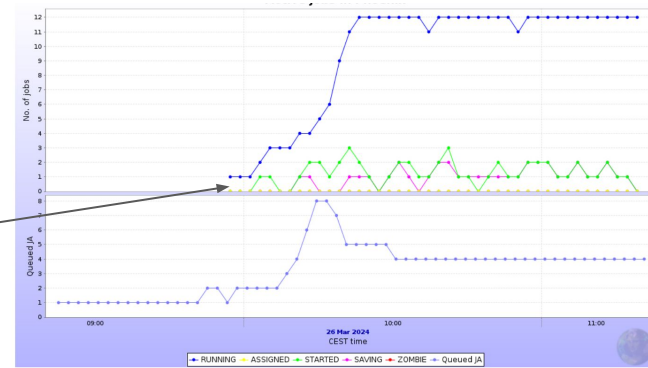
# ALICE

## GRID

- Issues under investigation:
  - **Software issue under investigation** (not reproducible at Glasgow) →ALICE MC expert is looking into that (interactive access requested)
  - ALICE queue was blocked due to a wrong path of one script (pilot job) → fixed this morning

## Workflow release validation

- **Setting of data reconstruction jobs → work in progress (queue unblocked this morning, we need to wait to finish the backlog of MC jobs)**



# Next plans

As announced we had an internal discussion about abstract submission to conferences:

- CCR workshop (Palau, 20-24 May)
  - Abstract submission deadline → 29th March
- Congresso SIF (Bologna, 9-13 September)
  - Abstract submission deadline → 22nd April
- CHEP (19-24 October)
  - Abstract submission deadline → 10th May (to be discussed with experiments ~one month in advance)
  - We thought it may work if site-oriented (e.g. it could be a part of a more general farming talk)

backup

# Current resources at CNAF

There are 2 ARM nodes (one already available)

- 256 cores
- 512 GB ram

Current setting (still work in progress)

- Cvmfs available
- Network: access to external network
- Gpfs client -> not available for ARM
- Condor -> not yet available

Data displayed in the following table are available in csv format in the github repository of [HEPIX-Forum](#)

Show: 10 entries

Search:

CPU	SMT enabled	Online CPUs	# Sockets	Cores/Socket	Threads/core	L2 cache	L3 cache	# Meas	Score	Spread	RAM	SWAP	Site	hash
filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter
AMD EPYC 9654 96-Core Processor	1	0-383	2	96	2	1024K	32768K	26	6000.578	0.714	1 TiB	4 GiB	IHEP	71892
AMD EPYC 9654 96-Core Processor	0	0-191	2	96	1	1024K	32768K	25	4955.599	0.421	1 TiB	4 GiB	IHEP	71892
Neoverse-N1	0	0-255	2	128	1	256 MiB (256 instances)	not_available	3	3767.707	0.989	497 GiB	500 GiB	INFN-T1	71892

It is in third position in the Hepscore23 DB

Opening access request was announced at the last CNAF CdG

Some experiments already got access and start to play with it.

# 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.