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

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

ICSC Italian Research Center on High-Performance Computing, Big Data and Quantum Computing
Missione 4 • Istruzione e Ricerca

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	>=2
KPI2.2.5.3	Technical notes (in experiments and ICSC)	>=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.

The plan as presented at the previous report

What to do in the last months of 2023

- Push main (LHC) experiments to get access and start setting basic workflows
- Discussing the right policy on the node with CNAF farming support (who can have super user rights to install packages) and collecting requirements
- Asking for a condor/grid queue
- Goals
  - Having some workflows able to run, at least in interactive mode, by the end of the year
  - Setting condor/grid submission
  - Converging on some internal monthly(?) meetings

Updates from infrastructure (17/11 CNAF CdG)

- Interactive access provided in one of the two nodes (wn-arm-02)
- HTCondor submission enables (CMS is testing) on wn-arm-01
- 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

Internal meeting (16/11/2023)

We had a discussion last week with people of all the 4 LHC experiments.

Considering the available manpower the main use cases are related to ALICE, ATLAS and CMS experiments which are already running tests and configuring the node to be able to run different workflows (DATA, MC, ML, …).

In the next slides I tried to summarize the progresses done so far.

CMS

HTCondor

- Integration in the production scenario
  - CE configured with the possibility of routing on arm nodes
  - CMS ready to send pilot job

Workflow release validation

- CNAF is one of the test bed (with Glasgow) for the validation on ARM
- Some interactive tests already run on wn-arm-02 (pilot jobs will be sent to wn-arm-01)
- Using container from cvmfs

ALICE

- Validation: access to alien, pythia MC jobs with a tag from Dec22 from cvmfs
- Issues under investigation:
  - problems in building new tags (internal Jira ticket opened to follow up)
  - some missing packages for panda (same error on lxplus8-arm) affecting python scripts (MC)

ATLAS

- software ARM available on cmvfs
- Using container from cvmfs
- Full detector simulation -> validated
- NN training -> in setup
- next step: job via batch system

Points under discussion

- 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

backup

Current resources at CNAF

There are 2 ARM nodes (one already available)

- 256 cores
- 512 GB ram

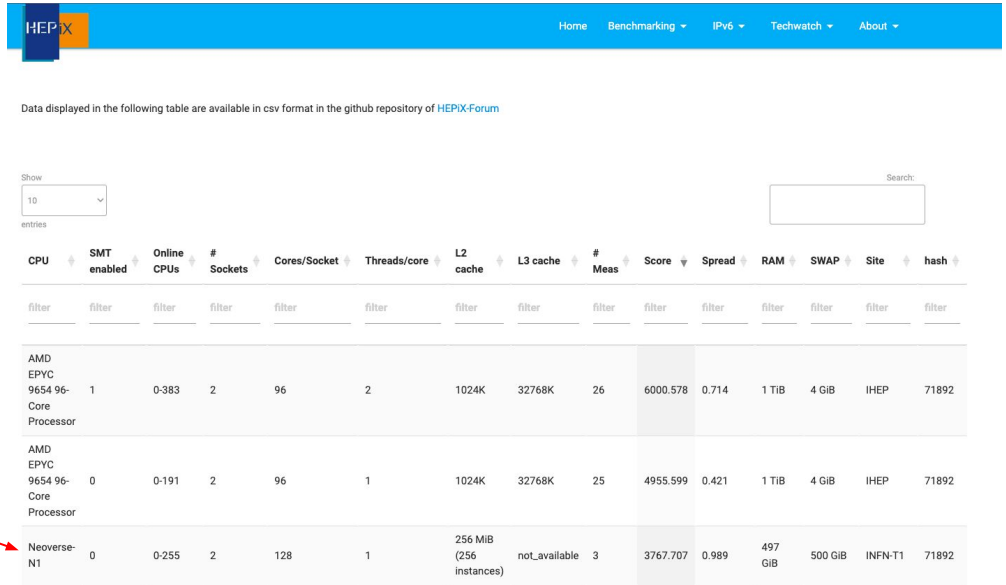

CPU	SMT enabled	Online CPUs	# Sockets	Cores/Socket	Threads/core	L2 cache	L3 cache	# Meas	Score	Spread	RAM	SWAP	Site	hash
filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter	filter
AMD EPYC 9654 96-Core Processor	1	0-383	2	96	2	1024K	32768K	26	6000.578	0.714	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

Current setting (still work in progress)

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

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.