CMS - CdG T1

Daniele Spiga
INFN-PG

16.02.2024

Running cores

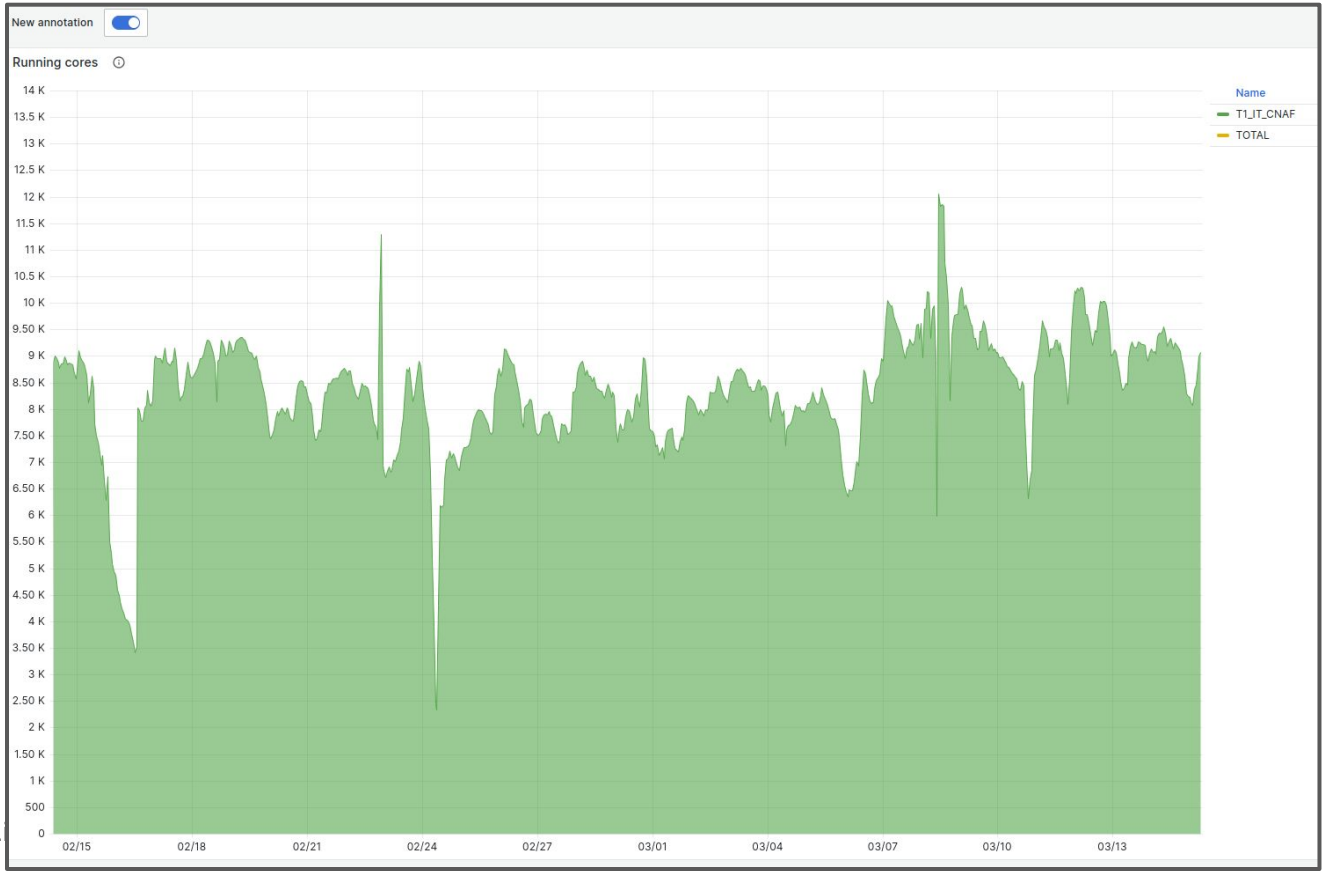

spiga@pg.infn.

Efficienza di CPU

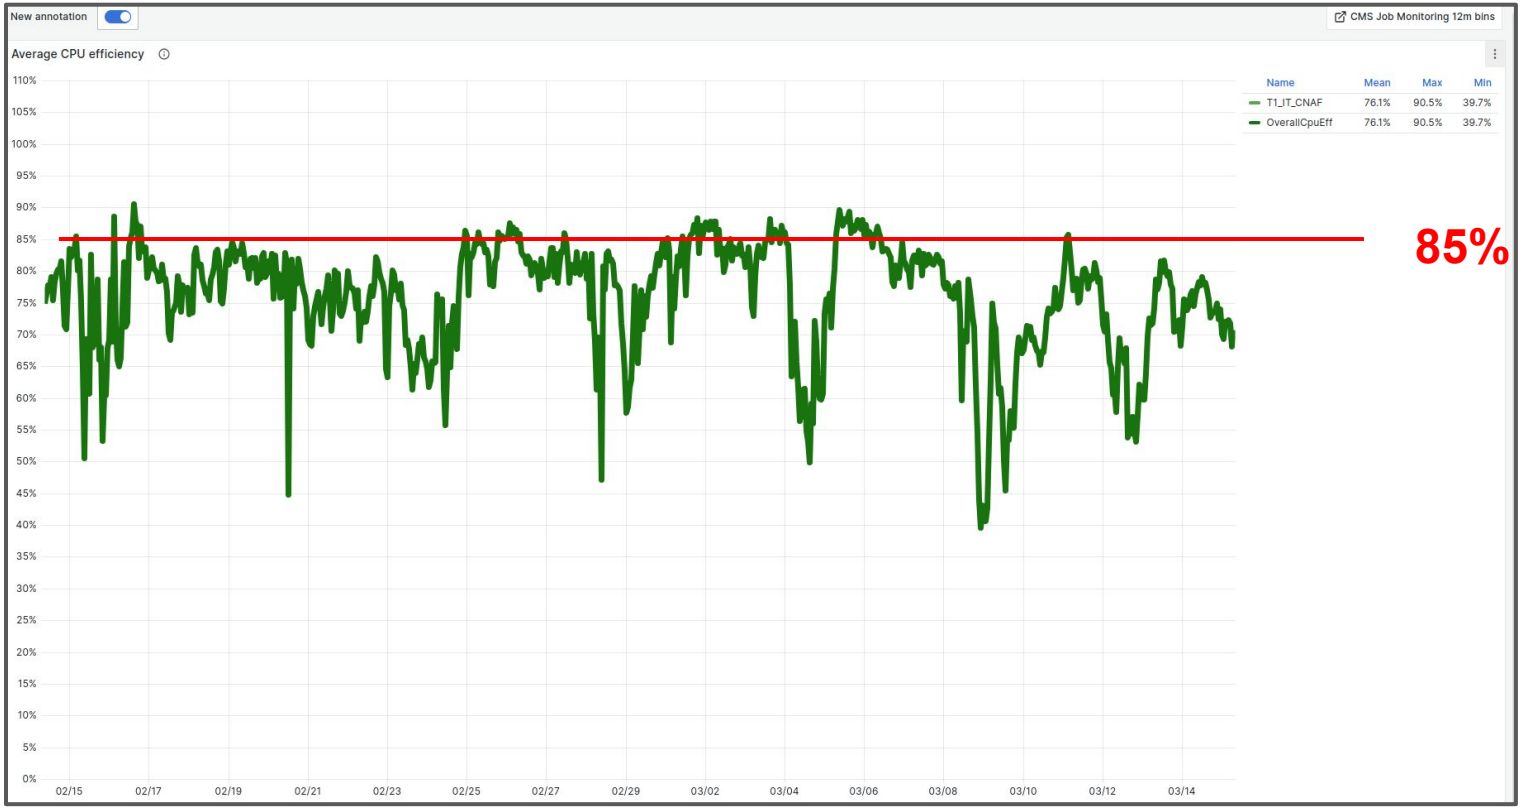

Overloading pilots

Ovvero come recuperare i cicli di CPU inutilizzati inserendo più payload nello stesso pilot job

- Un'idea semplice è quella di "ampliare" i nostri progetti pilota in modo che accettino più lavori di carico utile nelle stesse risorse acquisite dal sito
  - Esempio: aggiungendo il 25% di core CPU e memoria ai valori nominali del nostro pilota standard a 8 core

Partita prima una sperimentazione che ha dato risultati interessanti (Spagna e Germania e Usa). Poco tempo fa abbiamo abilitato i T2 italiani e negli ultimi giorni anche il T1

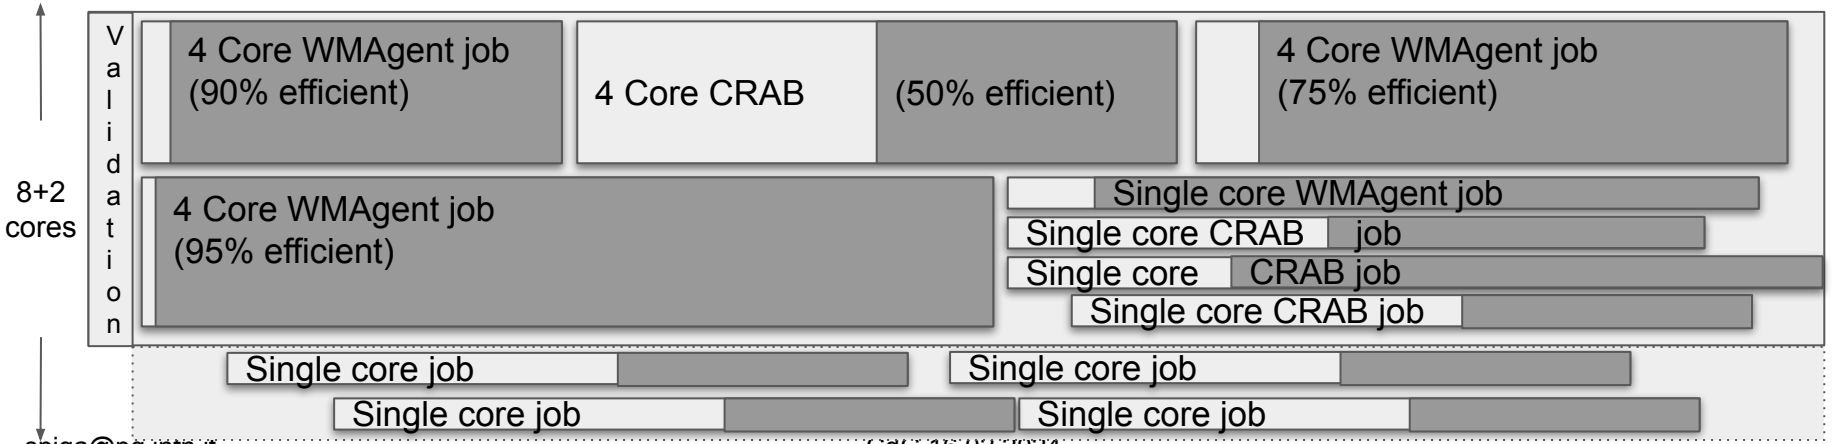

ARM

Validato Grace

Per CMS tutti i nodi sono in produzione

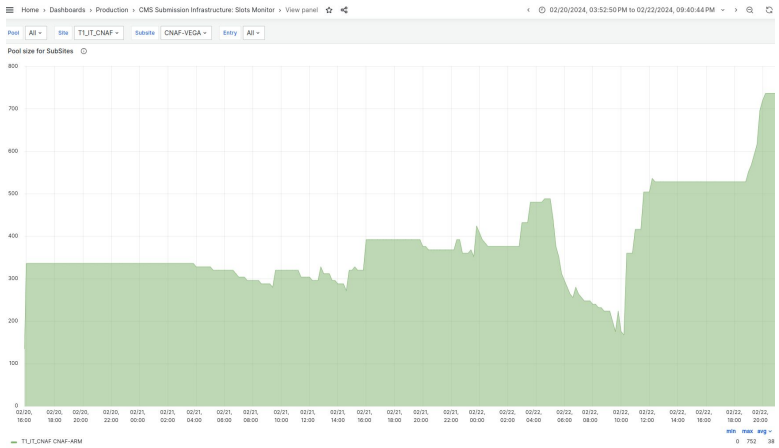

Rimane aperto IPV6 ma non sembra ne legato ad ARM ne CMS..

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.jlab.org/event/459/contributions/11636/
- https://indico.jlab.org/event/459/contributions/11553/

Risposta CMS Computing ai referee: CNAF come centro utilizzato per la convalida ARM insieme a Glasgow

Vega in Produzione

Vega EuroHPC sta girando in produzione come SubSite del CNAF

- Possibile impatto su Storage (e Xrootd) ma a questa scala (bassa) non abbiamo reali problemi

Faremo report a CRB prossima settimana

Maximise the impact of the grant
Goals and Roadmap

Use VEGA resources to contribute to current Alpaka (**GPU**) validation campaign needed for the HLT

- Approach agreed between C&O / DRP, PPD / PdmV (Adriano) and TSG (Andrea B.)
- Using the offline resources to execute online workflow

Demonstrate the ability to execute “any type of workflow”

- Task chain making access to both primary and secondary input file
- Standard MC workflow with remote read of the pre-mixed pileup

Technically VEGA was made a sub-site T1_IT_CNAF as a storage-less site relying on AAA + CNAF storage

- General concept already proven for other integrations
- Simplify the resource assignment process
- We rely on something we can easily keep under our control, simplify the technical integration process and allow to focus on the “new” aspect: using GPUs for a validation campaign

CRB, March 21th 2024

7

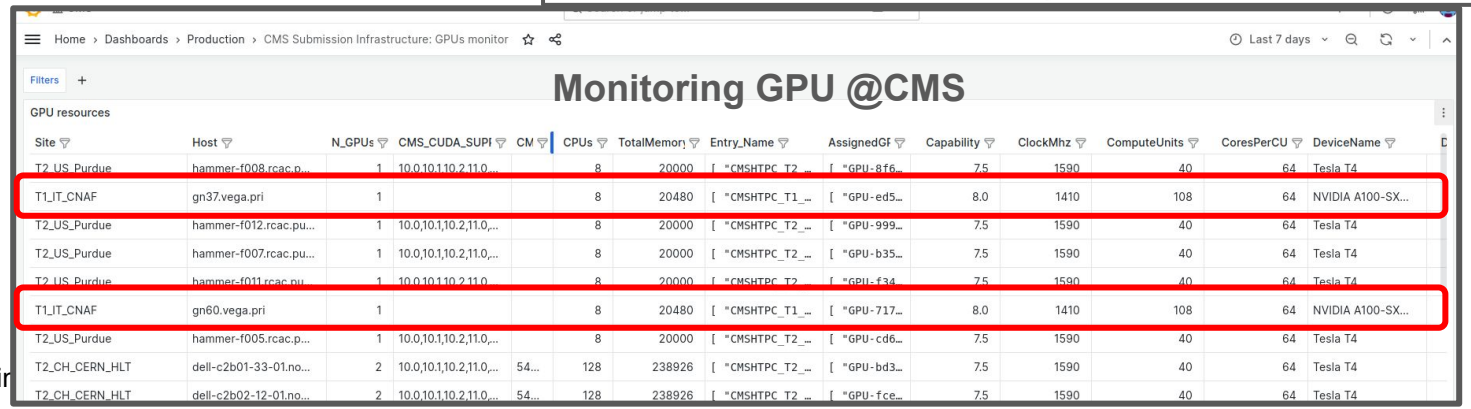

spiga@pg.in