Analysis of conversion factors (corepower values) by using HS23

Natalia Szczepanek, Domenico Giordano (IT-TC-LCG)

WLCG Operations Coordination

3rd October 2024

Outlook

- **Corepower** value of a server is the HS06 score of a server normalized to single core
- For most sites corepower is still HS06/core (Transition from HS06 to HS23 in April 2023)
- Comparison of corepower declared by sites in ATLAS-CRIC with *runtime* corepower based on HEPScore23, measured via the job submission infrastructure of the ALTAS experiment
- Analysis of corepower from different data sources (ATLAS-CRIC, BDII: GLUE1/GLUE2, APEL*)
    - Visible discrepancies between data sources
- **Objective: Analysis of corepower**

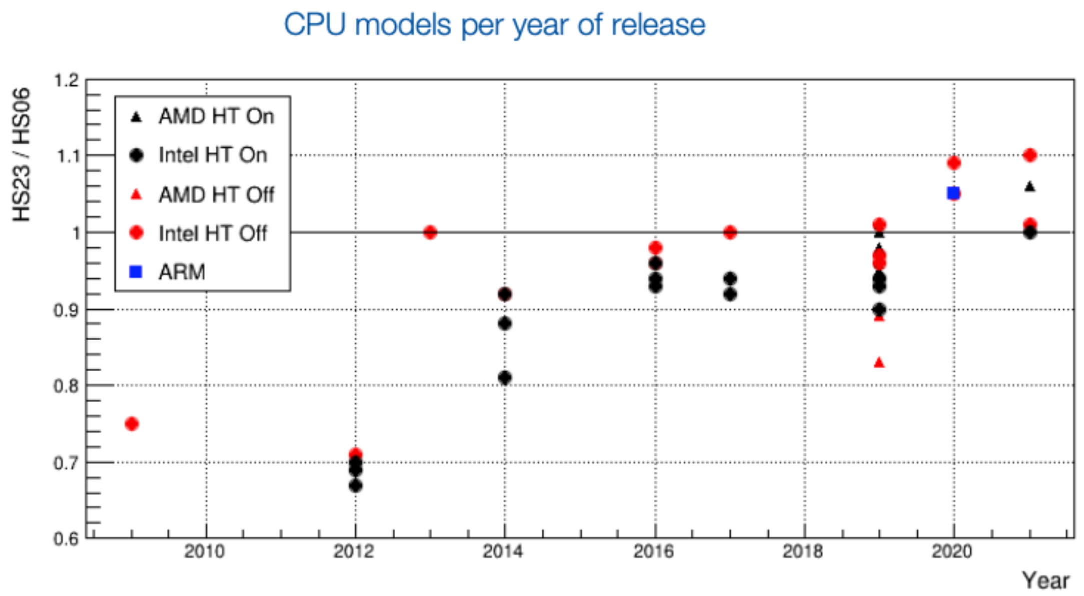

Submission Infrastructure

Automated submission of HS23 via HammerCloud

Infrastructure:

- PanDA, HammerCloud, Rucio, ActiveMQ, Elasticsearch, Grafana, Kibana…

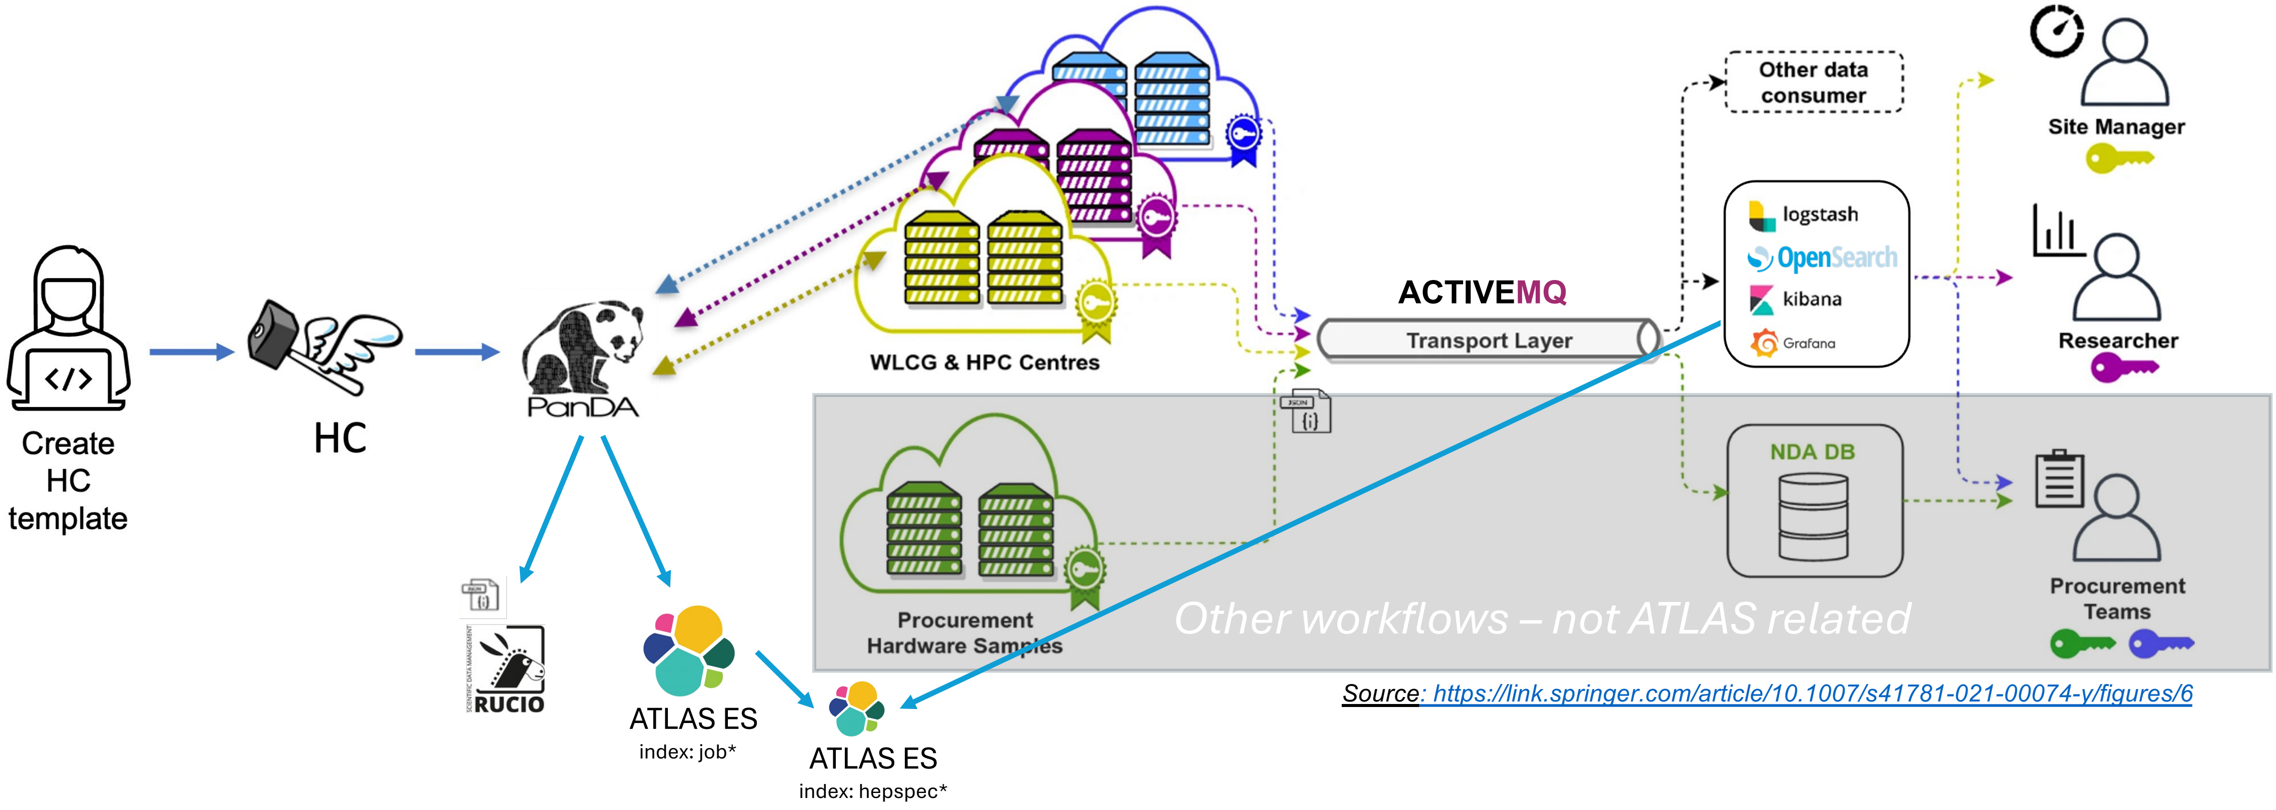

Source: https://link.springer.com/article/10.1007/s41781-021-00074-y/figures/6

Probing job slots

- Each site has servers with a variety of CPU models and number of cores (256, 128, 64...)
- We are running the benchmark injecting the HEP Suite script as a normal experiment job running inside the PILOT Apptainer
- We probe multi-core job slots (8 cores)

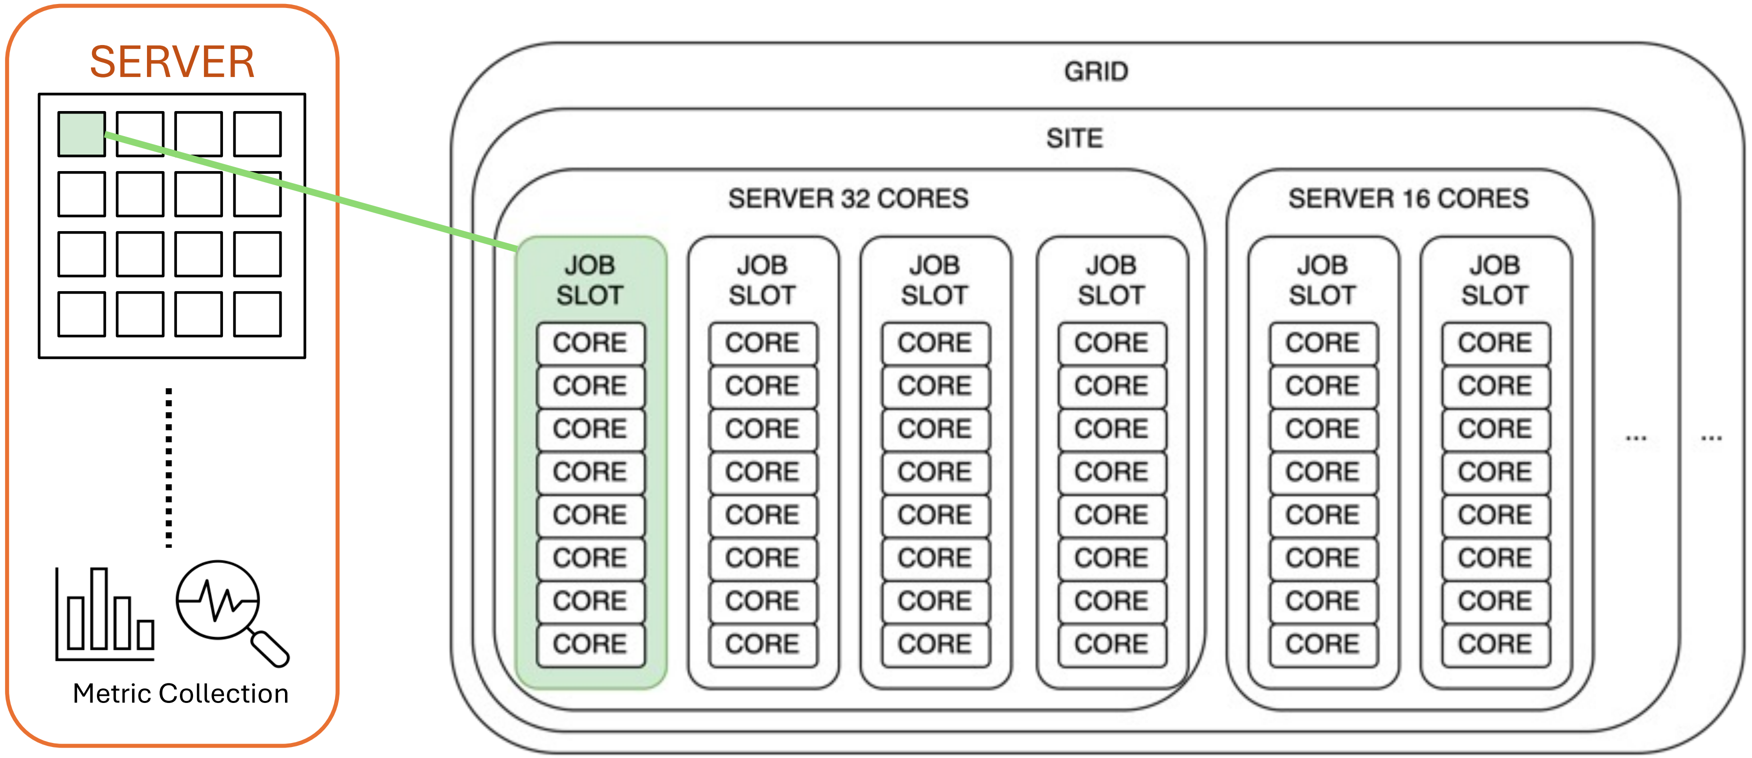

Declared corepower

- Corepower data for ATLAS is taken from ATLAS-CRIC
  - Essential metric to understand the computing capabilities based on the specific hardware
  - In ATLAS-CRIC each ATLAS site can have multiple Panda Queues (resources)
    - **Challenge: Mapping between Panda Queues and ATLAS Sites**
- Corepower reported by sites is the **weighted average** of different corepowers of given CPU models available at the site
  - **Challenge: Find the correct weights per site using finished jobs (next slide)**
- Compute the corepower of a site as weighted averages of the runtime HS23 values

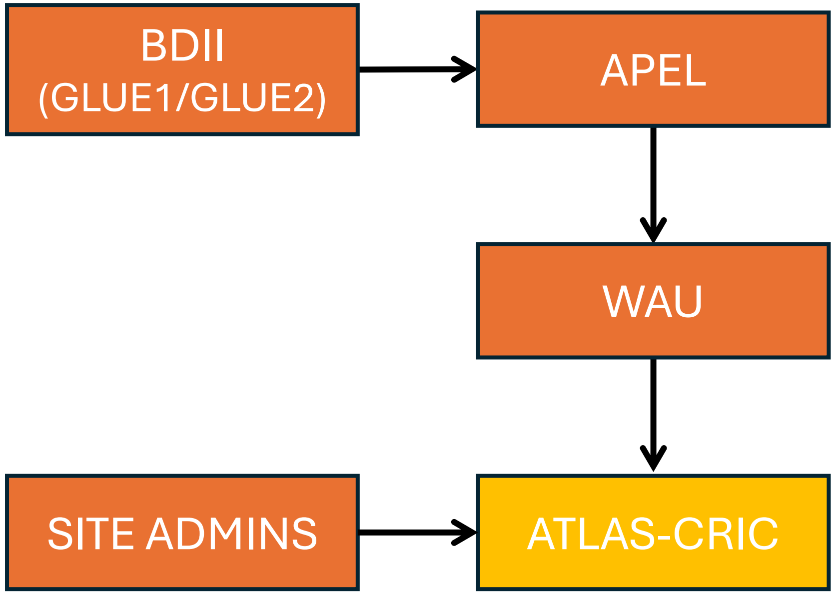

Based on : link and internal communication

Declared corepower

- Corepower data for ATLAS is taken from ATLAS-CRIC
  - Essential metric to understand the computing capabilities based on the specific hardware
  - In ATLAS-CRIC each ATLAS site can have multiple Panda Queues (resources)
    - **Challenge: Mapping between Panda Queues and ATLAS Sites**
- Corepower reported by sites is the **weighted average** of different corepowers of given CPU models available at the site
  - **Challenge: Find the correct weights per site using finished jobs (next slide)**
- Compute the corepower of a site as weighted averages of the runtime HS23 values

PanDA Queue	State	type	Cloud	Tier	Final status	core power
pic_MareNostrum4	ACTIVE	production	ES	T1	ONLINE	27.18 27.18
IFIC_MareNostrum4	ACTIVE	production	ES	T2D	ONLINE	27.18 27.18
UAM_MareNostrum4	ACTIVE	production	ES	T2D	ONLINE	27.18 27.18
HPC2N	ACTIVE	unified	ND	T1	online	25.45 25.45
SWT2_GOOGLE_VHIMEM	ACTIVE	production	US	T2D	online	24.6 24.6
praguelcg2_Barbora_MCORE	ACTIVE	production	DE	T2D	online	24.5 24.5
UNIBE-LHEP-UBELIX	ACTIVE	unified	ND	T2	online	21.7 21.7
SLAC	ACTIVE	production	US	T3D	online	20.05 20.05
ANALY_SLAC_GPU	ACTIVE	analysis	US	T3D	ONLINE	20 20
UNIGE-BAOBAB	ACTIVE	unified	ND	T2	online	19.98 19.98
CA-IAAS-T3	ACTIVE	production	CA	T3	online	19 19
SWT2_GOOGLE_ARM	ACTIVE	unified	US	T2D	online	18.77 18.77
DCSC	ACTIVE	unified	ND	T1	online	18.07 18.07
INFN-CNAF_ARM	ACTIVE	unified	IT	T1	online	17.9 17.9
UNI-FREIBURG_NHR	ACTIVE	unified	DE	T2D	online	16 16

https://atlas-cric.cern.ch/atlas/pandaqueue/list/

Runtime corepower

- corepower_runtime per site:
  - For each CPU model on each site calculate the weight as:

$$w_x = \frac{\sum_i \text{walltime\_x\_core}_i}{\sum_j \sum_k \text{walltime\_x\_core}_{jk}}$$

  - For each site calculate the weighted average (using available benchmarking CPU Models):

$$\text{corepower\_runtime}_s = \frac{\sum_x w_x \cdot \text{corepower\_runtime}_x}{\sum_x w_x}$$

- Relative change:

$$Relative\ change = \frac{\text{corepower\_runtime}_s}{\text{corepower\_declared}_s} - 1$$

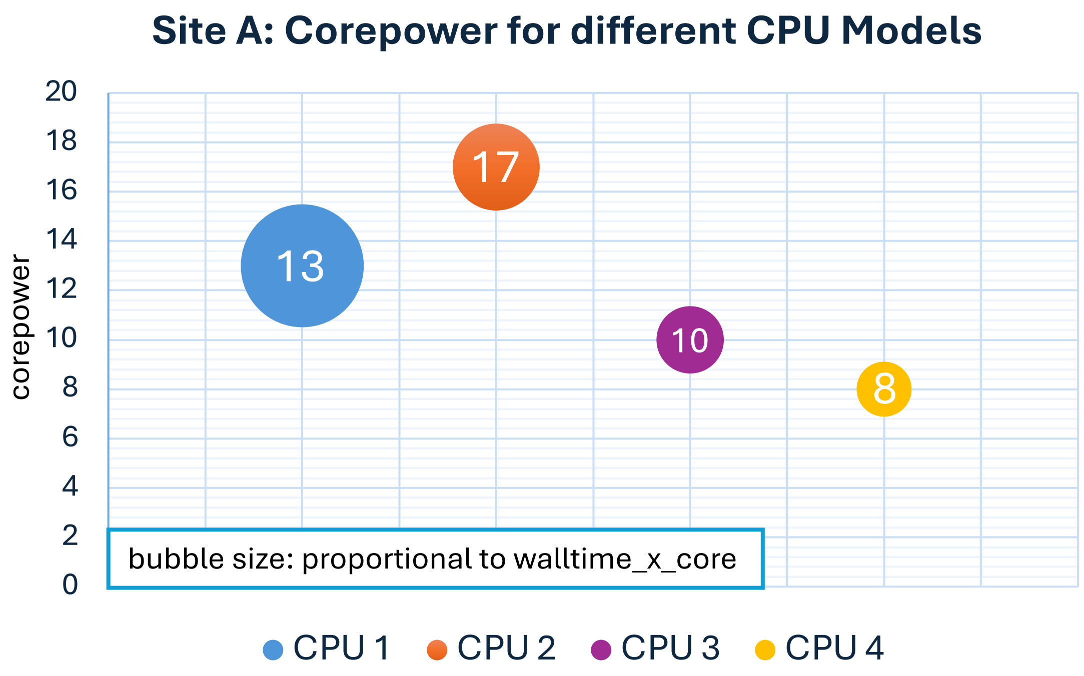

Runtime corepower

- corepower_runtime per site:
  - For each CPU model on each site calculate the weight as:

$$w_x = \frac{\sum_i \text{walltime\_x\_core}_i}{\sum_j \sum_k \text{walltime\_x\_core}_{jk}}$$

  - For each site calculate the weighted average (using available benchmarking CPU Models):

$$\text{corepower\_runtime}_s = \frac{\sum_x w_x \cdot \text{corepower\_runtime}_x}{\sum_x w_x}$$

- Relative change:

$$Relative\ change = \frac{\text{corepower\_runtime}_s}{\text{corepower\_declared}_s} - 1$$

Site A

weights derived from walltime_x_core

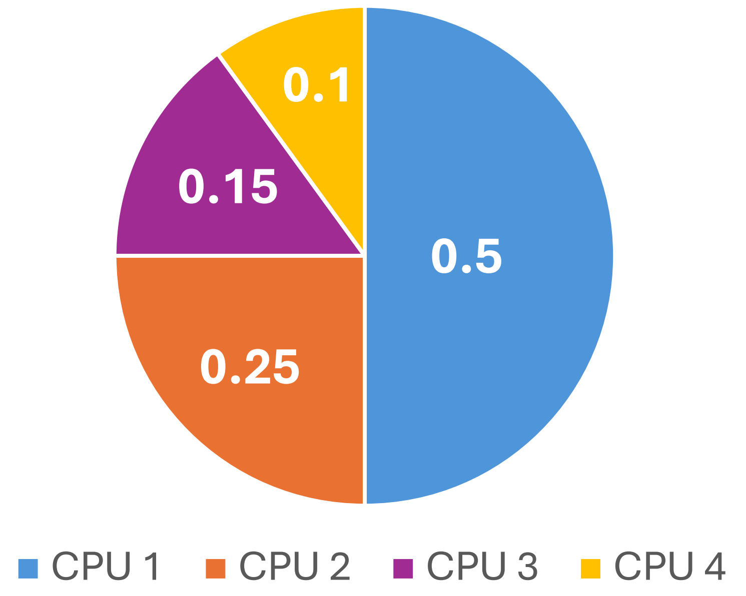

Runtime corepower

- corepower_runtime per site:
  - For each CPU model on each site calculate the weight as:

$$w_x = \frac{\sum_i \text{walltime\_x\_core}_i}{\sum_j \sum_k \text{walltime\_x\_core}_{jk}}$$

  - For each site calculate the weighted average (using available benchmarking CPU Models):

$$\text{corepower\_runtime}_s = \frac{\sum_x w_x \cdot \text{corepower\_runtime}_x}{\sum_x w_x}$$

- Relative change:

$$Relative\ change = \frac{\text{corepower\_runtime}_s}{\text{corepower\_declared}_s} - 1$$

Site A

weights derived from walltime_x_core

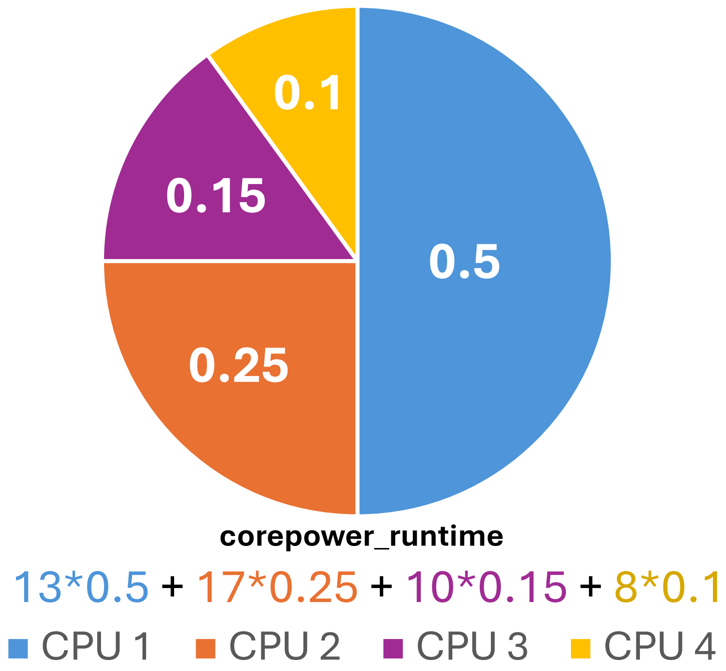

Runtime corepower

- corepower_runtime per site:
  - For each CPU model on each site calculate the weight as:

$$w_x = \frac{\sum_i \text{walltime\_x\_core}_i}{\sum_j \sum_k \text{walltime\_x\_core}_{jk}}$$

  - For each site calculate the weighted average (using available benchmarking CPU Models):

$$\text{corepower\_runtime}_s = \frac{\sum_x w_x \cdot \text{corepower\_runtime}_x}{\sum_x w_x}$$

- Relative change:

$$Relative\ change = \frac{\text{corepower\_runtime}_s}{\text{corepower\_declared}_s} - 1$$

Site A corepower_runtime

13.05

Results

Relative change for different ATLAS sites

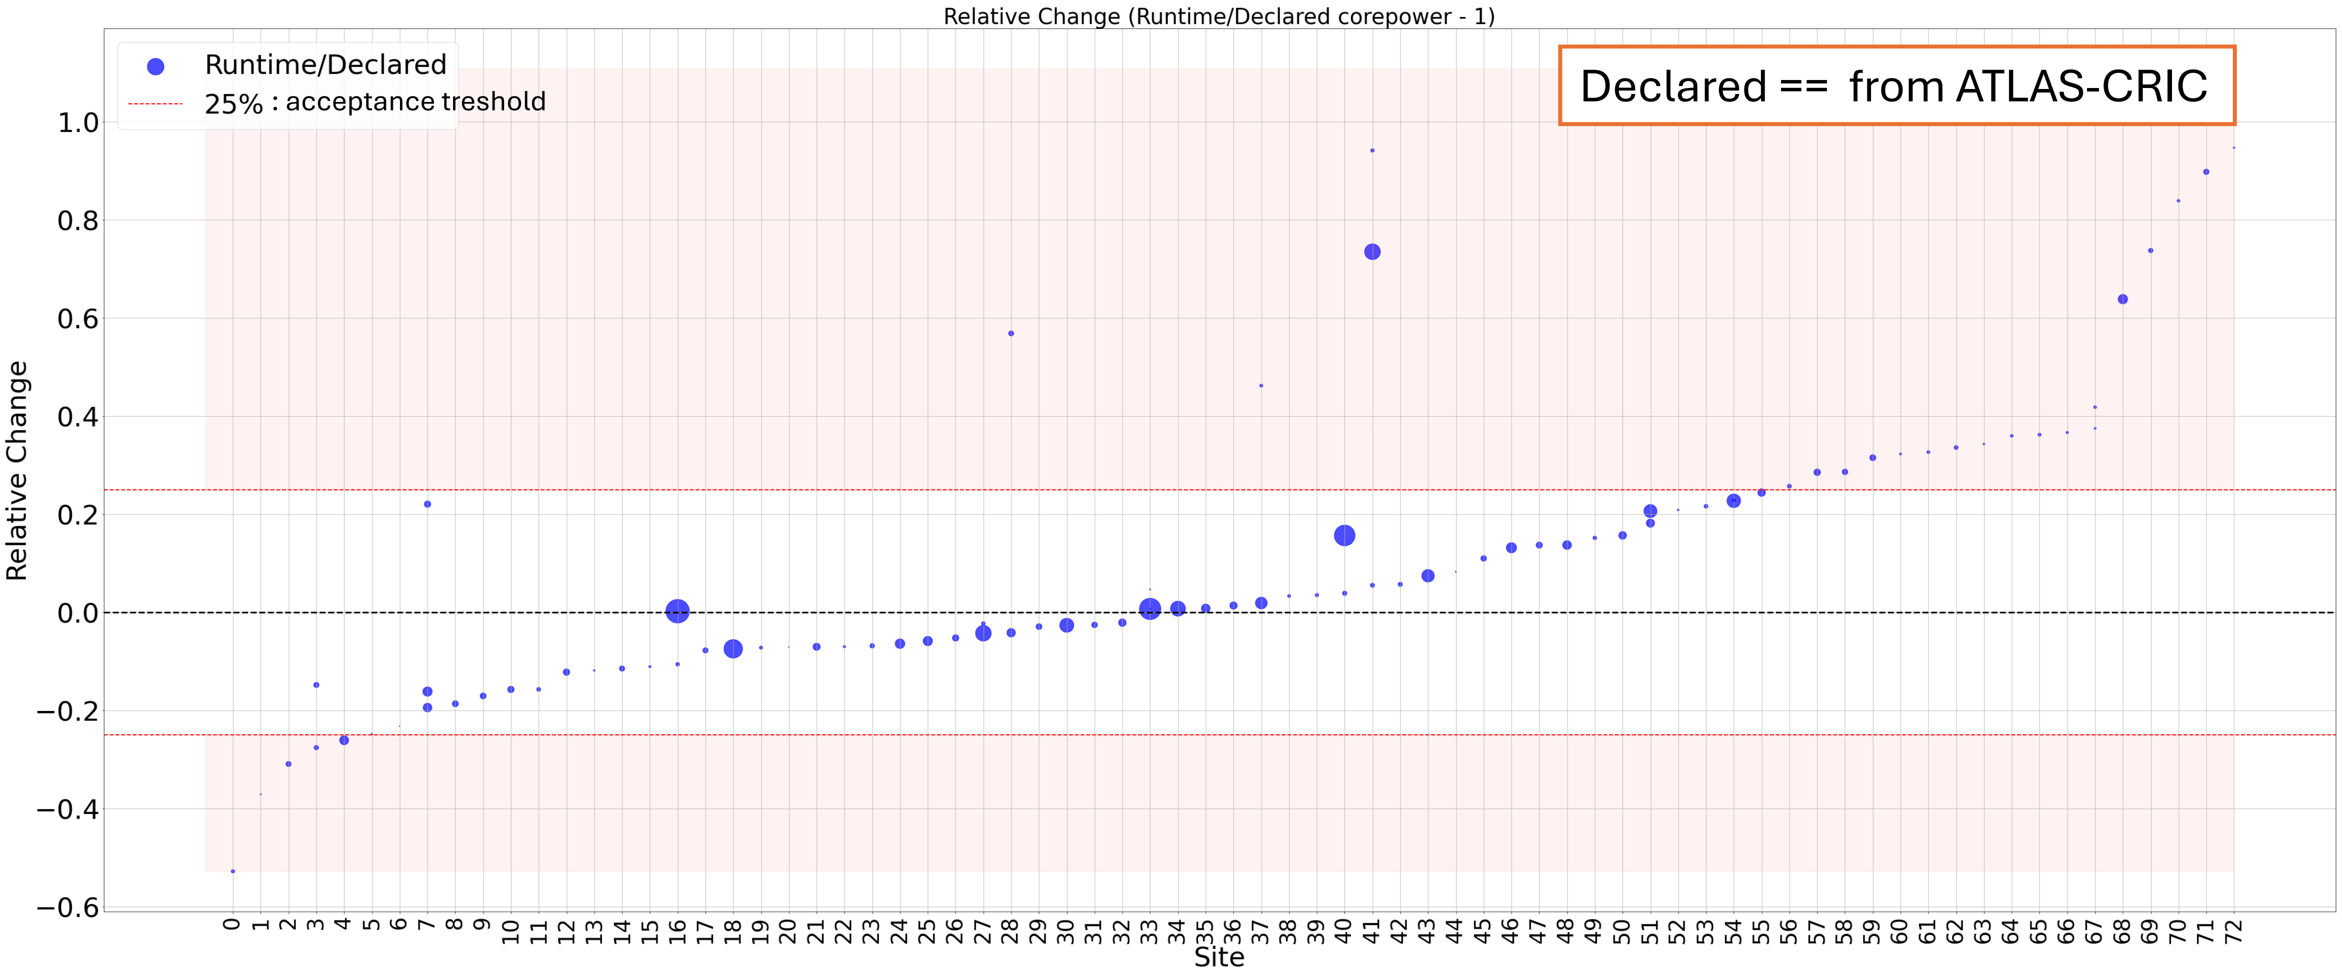

Relative change for different ATLAS sites

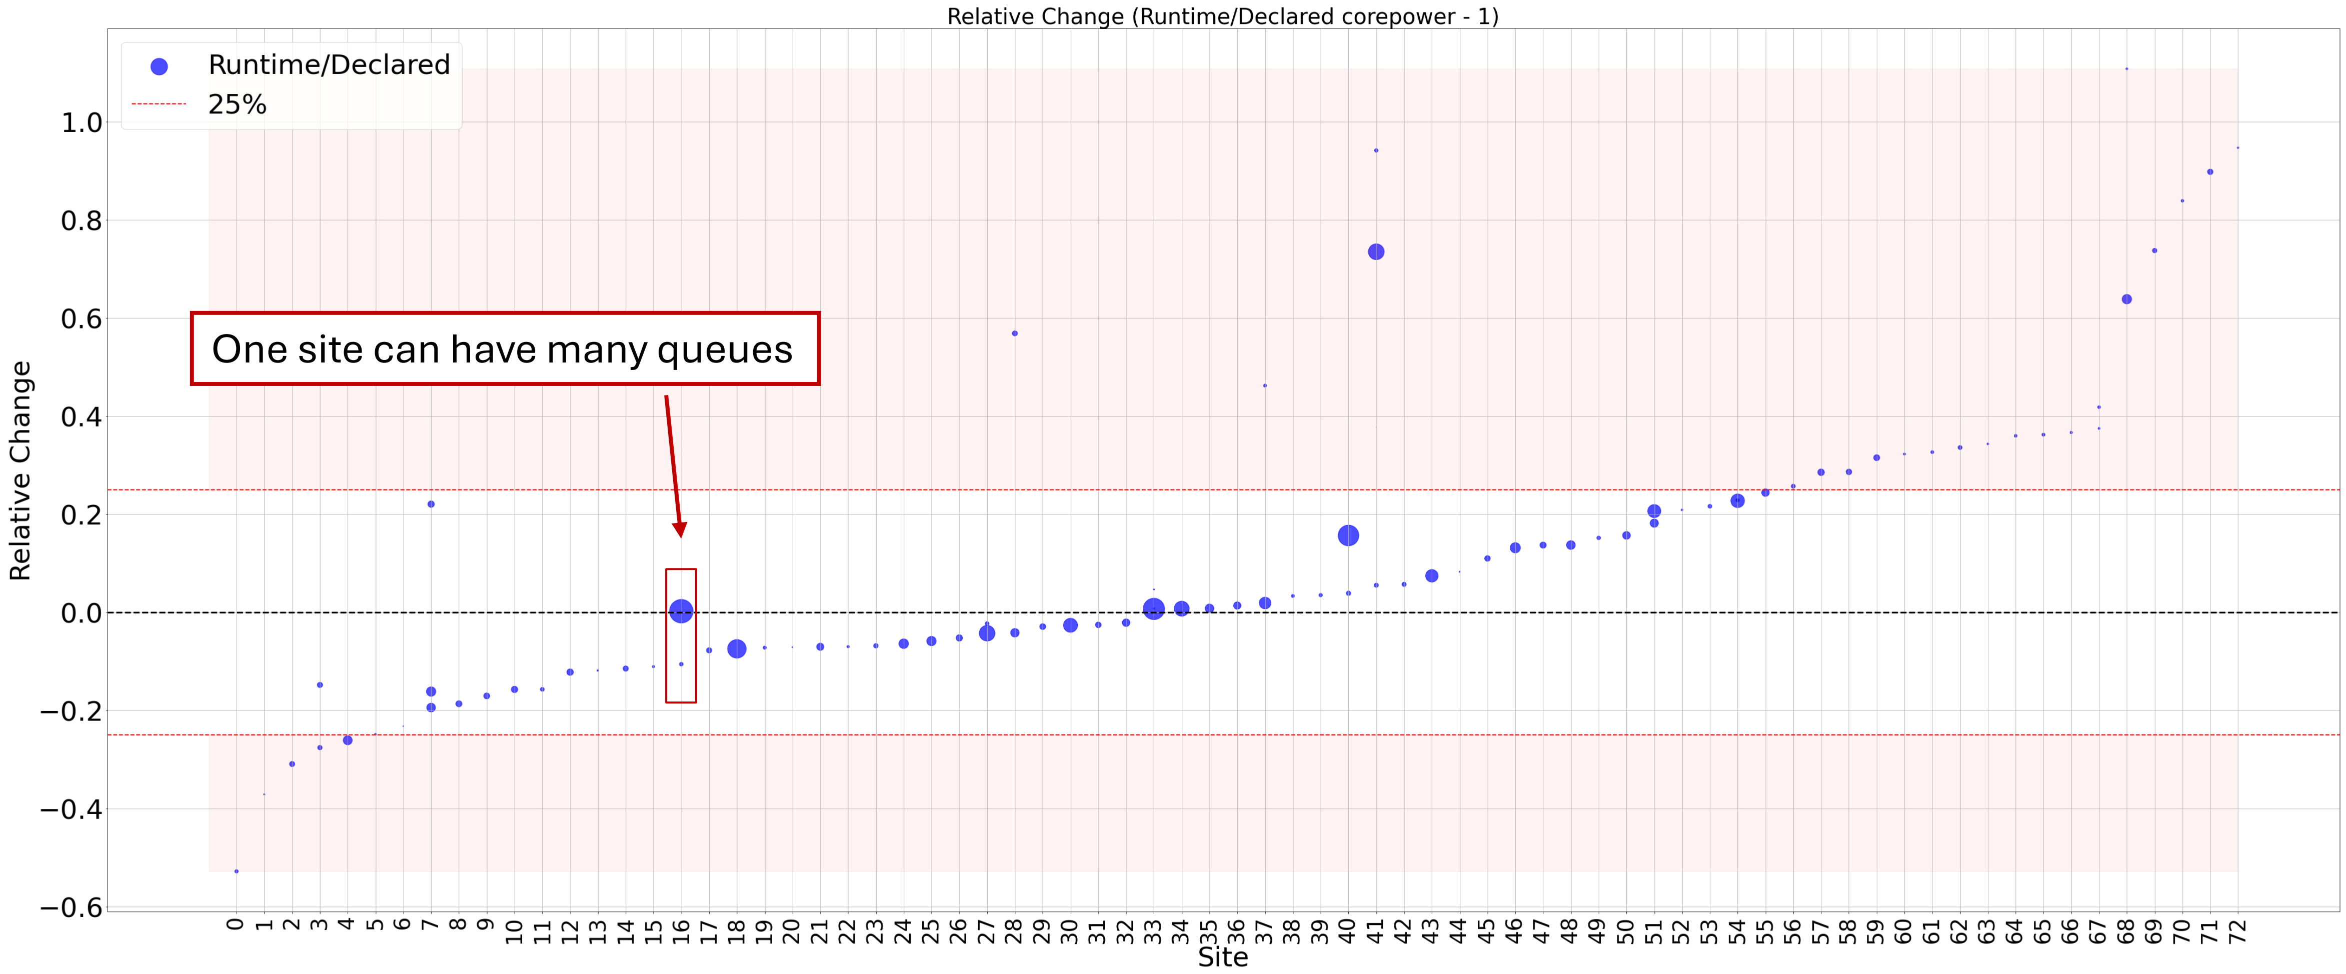

Relative change for different ATLAS sites

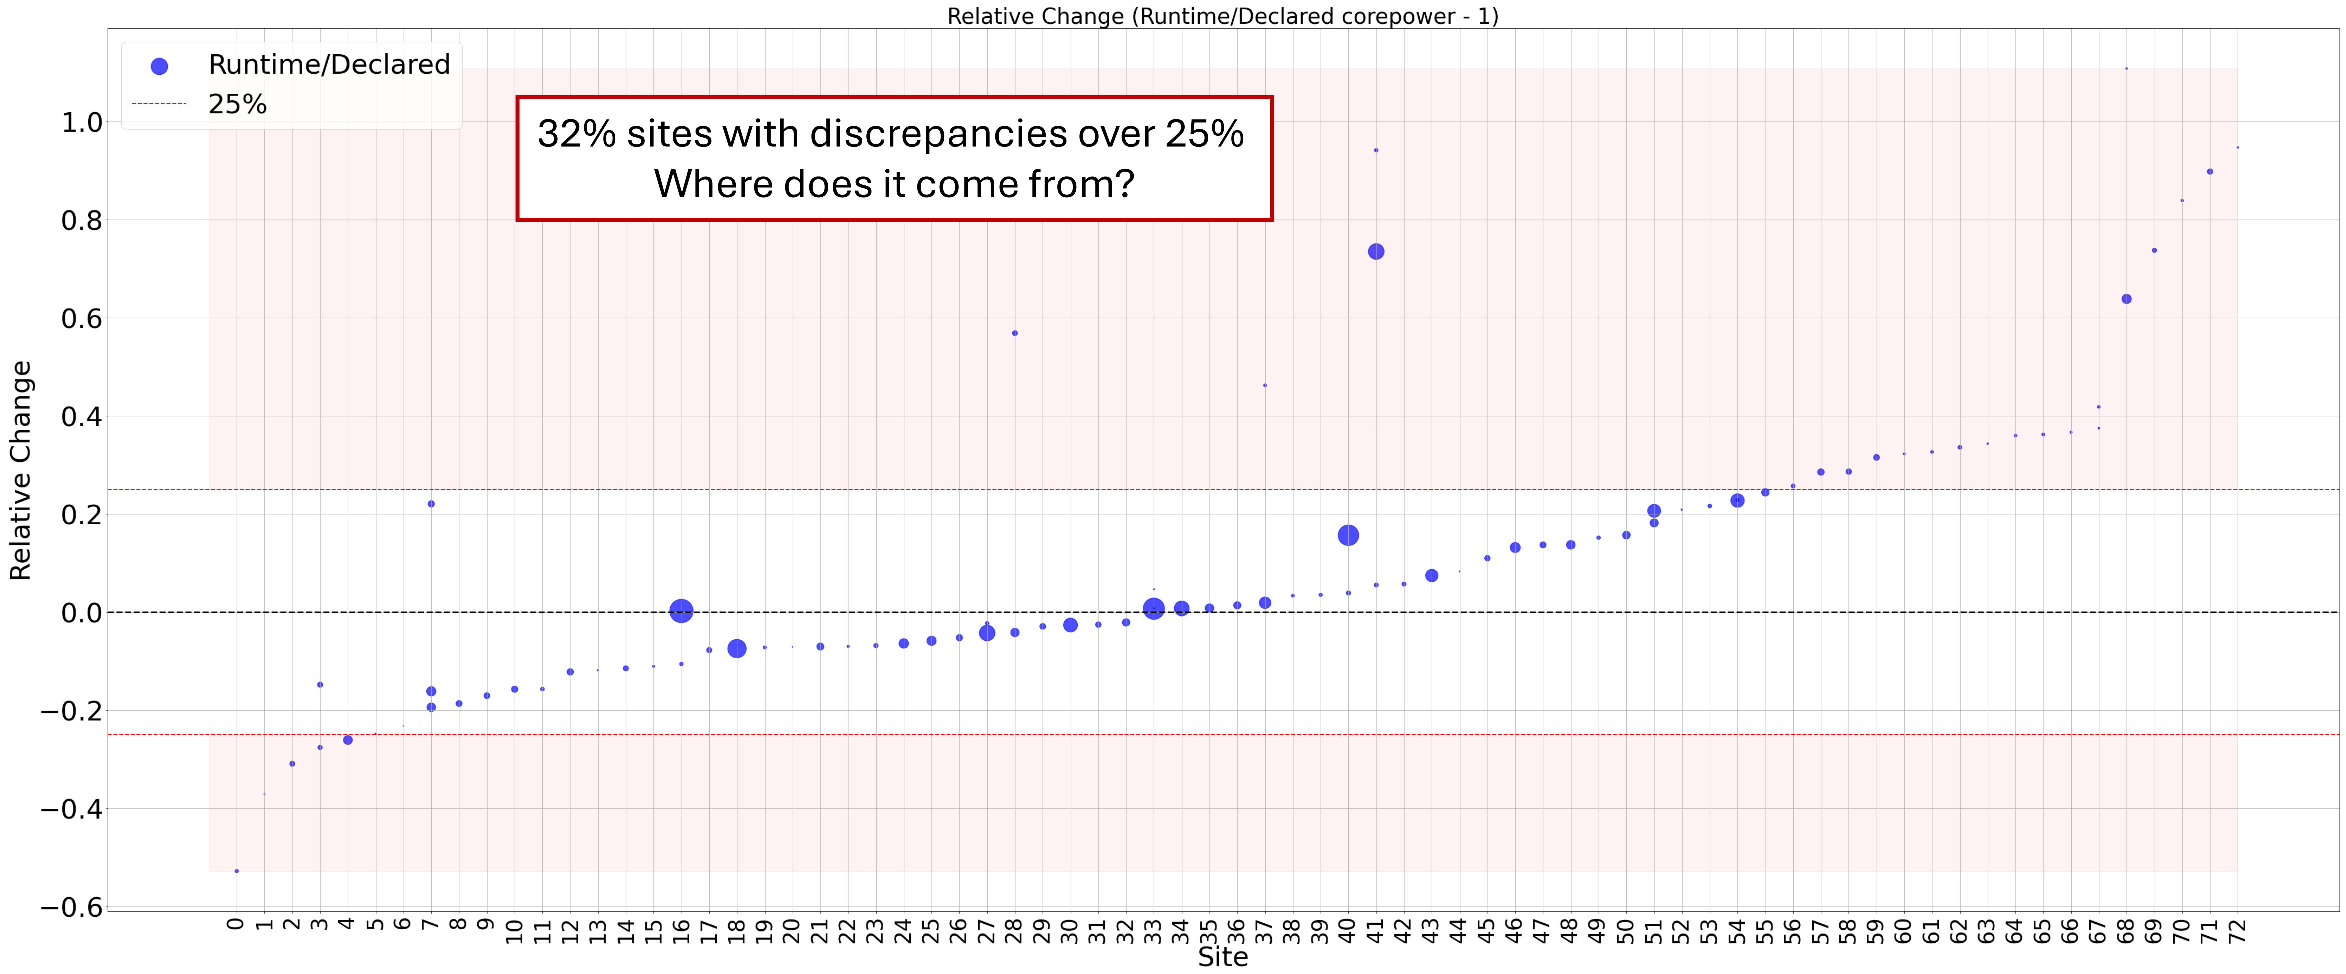

Similar studies – similar results

- Similar studies have been presented in ATLAS and at CHEP
    - *“A comparison of HEPSPEC benchmark performance on ATLAS Grid-Sites versus ideal conditions“* (2022)
- Same type of discrepancies was found

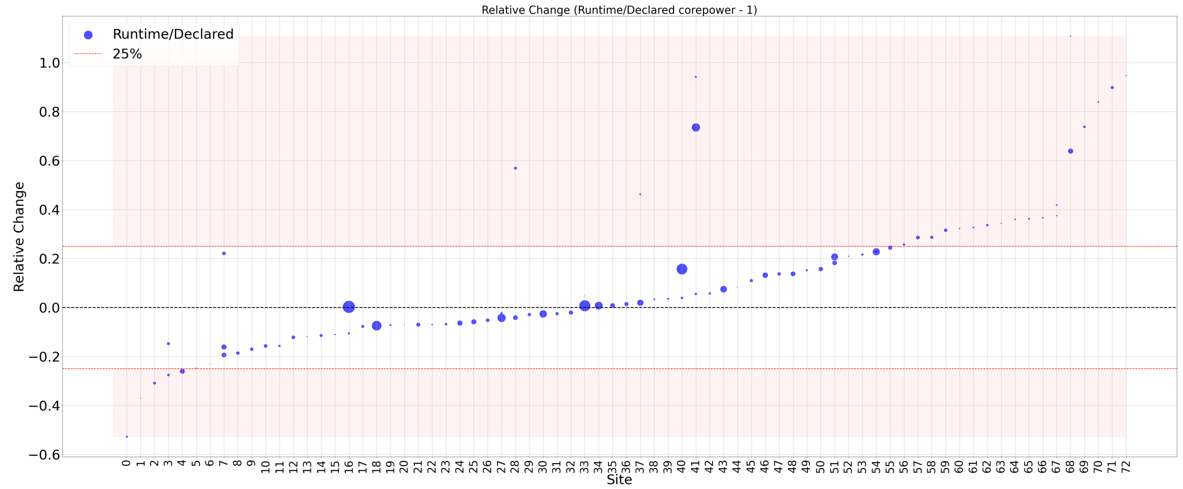

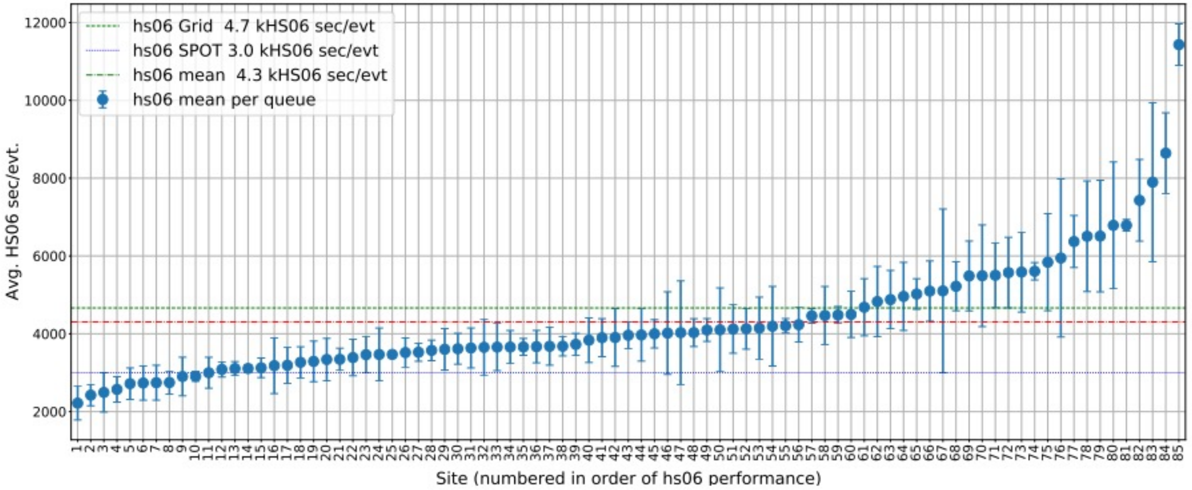

A comparison of HEPSPEC benchmark performance on ATLAS Grid-Sites versus ideal conditions

by Michael Boehler

Analysis of corepower values

- The corepower values from ATLAS-CRIC largely differ from runtime corepower for 35% of sites
- The corepower values from ATLAS-CRIC seem outdated
- It is worth to check whether those discrepancies are just a matter of lack of updates in ATLAS- CRIC from site admins, or does it occur also in other data sources, as BDII or APEL*
- We were able to collect conversion factors data from GLUE1 and GLUE2 and calculate corepower values from it

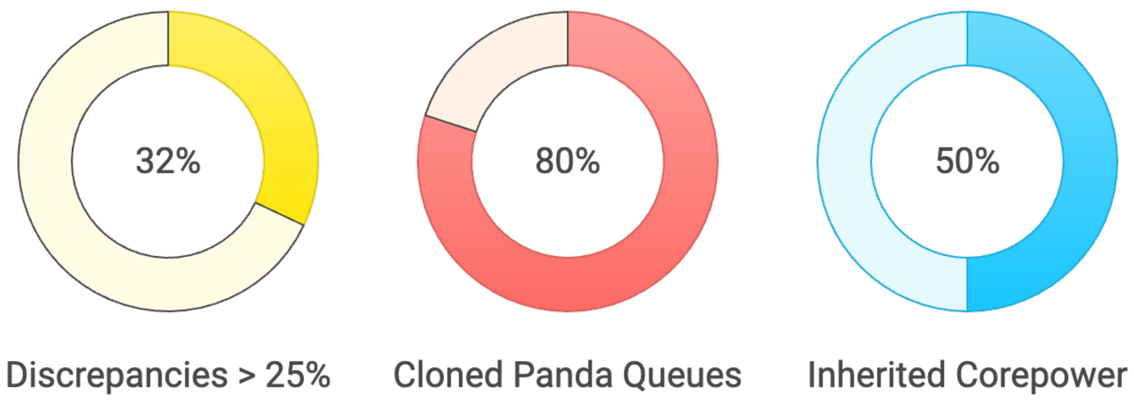


```
ldapsearch -x -LLL -h lcg-bdii.egi.eu -p
2170 -b o=grid \
objectClass=GlueCE GlueCECapability
```

```
ldapsearch -x -LLL -H ldap://egee-
bdii.cnaf.infn.it:2170 -b
"GLUE2GroupID=grid,o=glue"
'objectClass=GLUE2Benchmark' GLUE2BenchmarkT
ype GLUE2BenchmarkValue
```

https://twiki.cern.ch/twiki/bin/view/LCG/DataNormalization#Data_Normalization_as_of_25_09_2

Comparison between GLUE1 and GLUE2 values

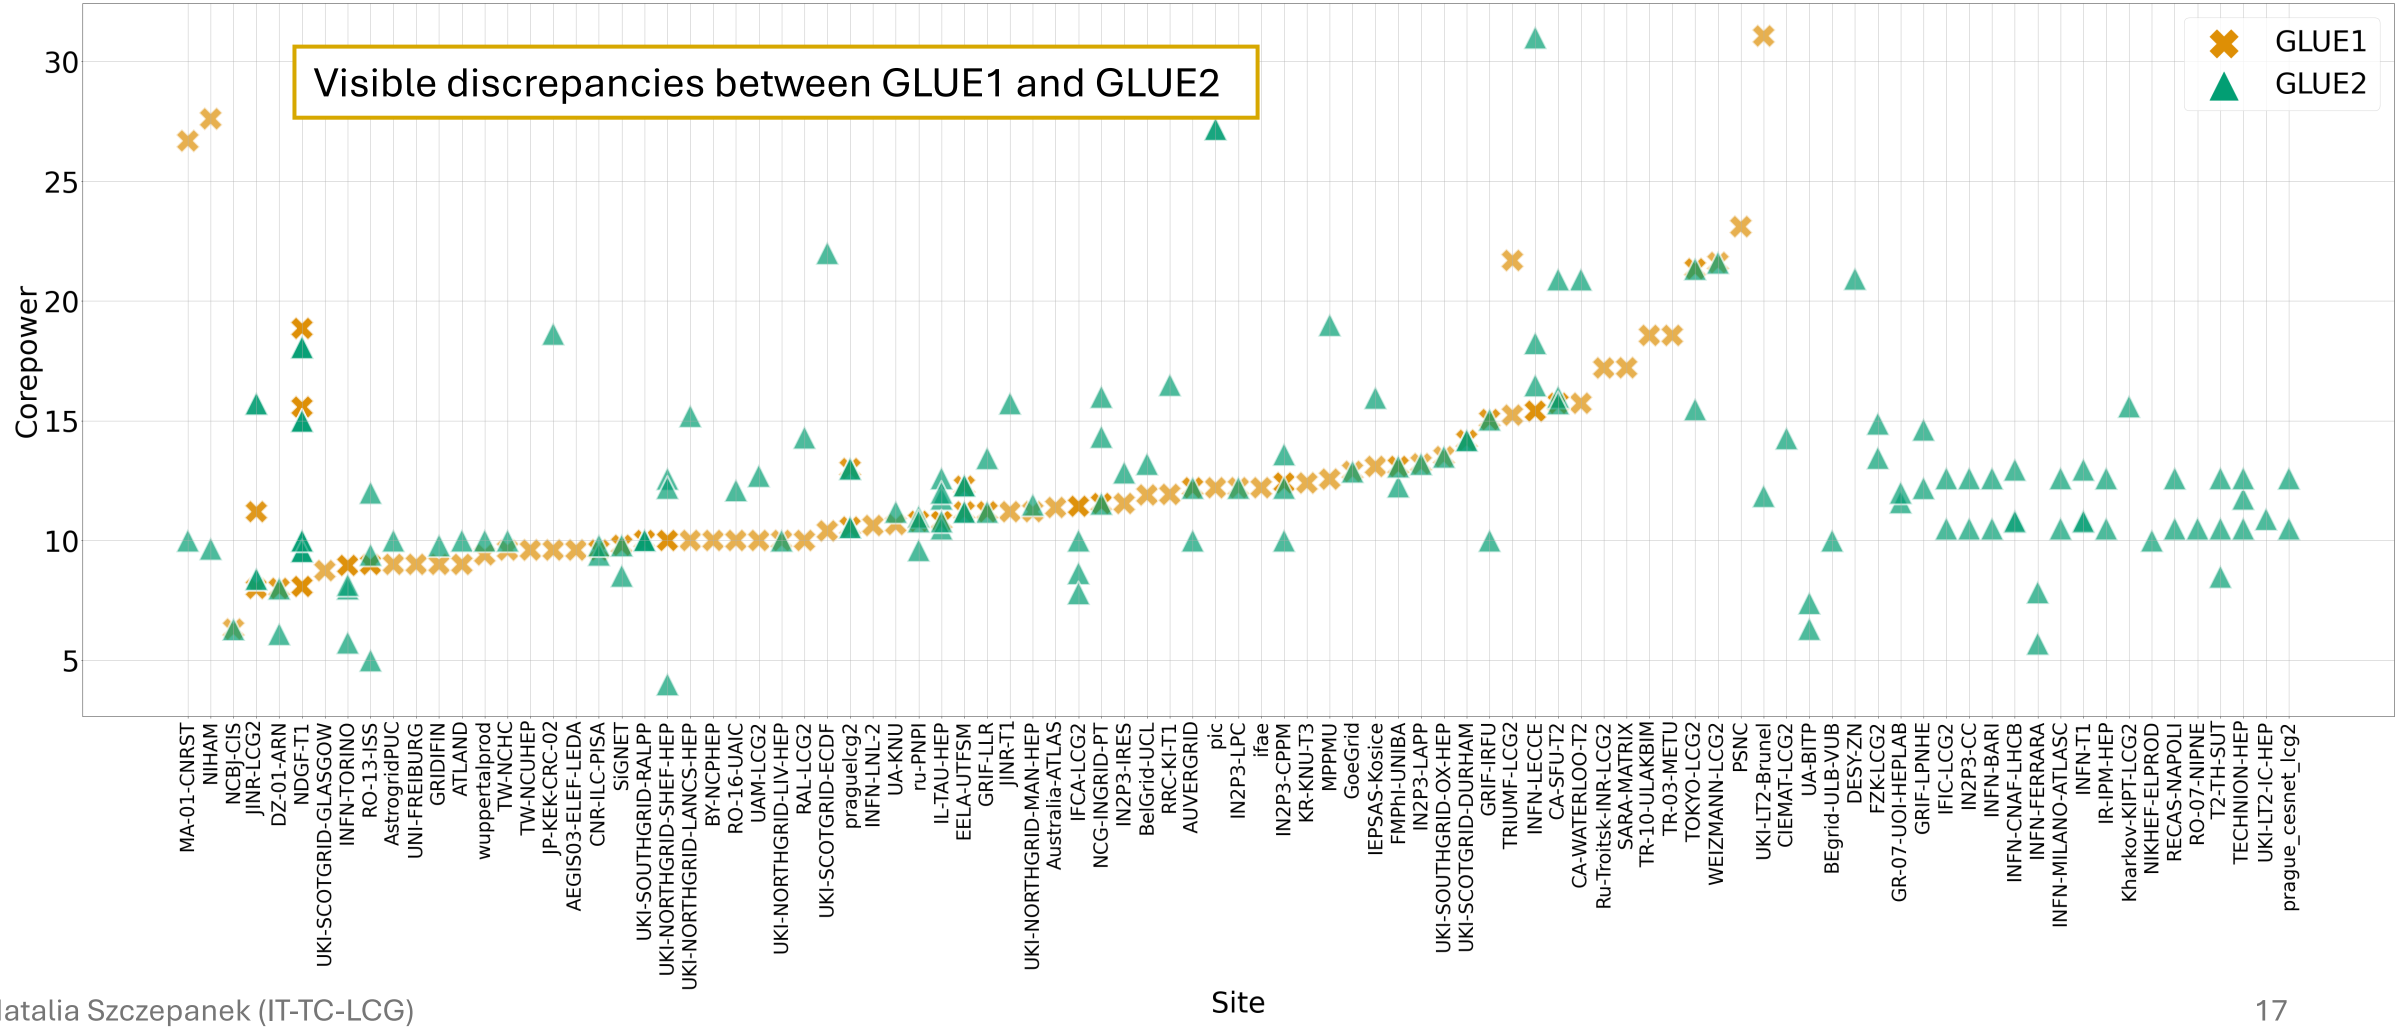

How does it look comparing to ATLAS-CRIC values?

Comparison between CRIC, GLUE1 and GLUE2 values

- Missing matches between CRIC/GLUE1/GLUE2
- Visible discrepancies between different data sources
- Are these discrepancies known and acceptable?

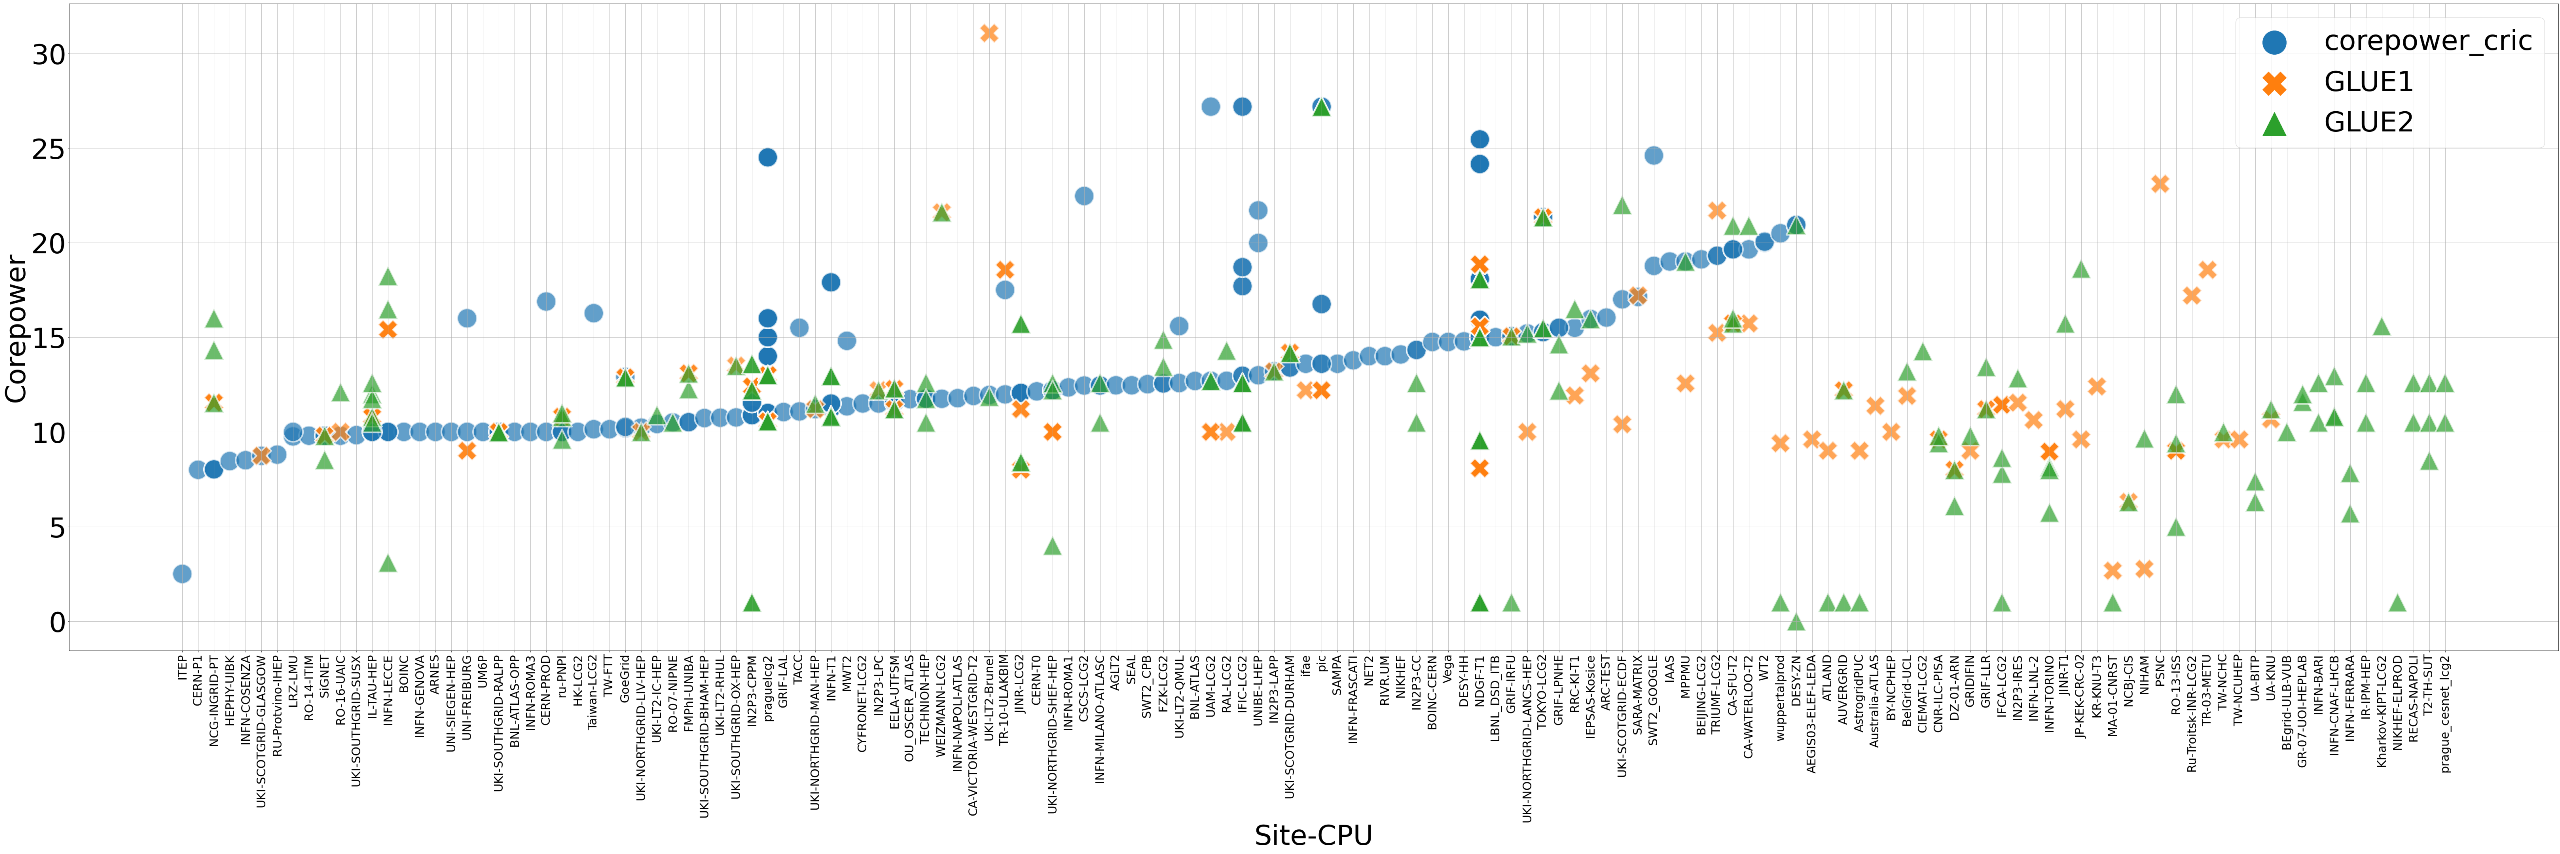

How does this relate to HS23 runtime corepower?

Relative change of our measurements and different data sources

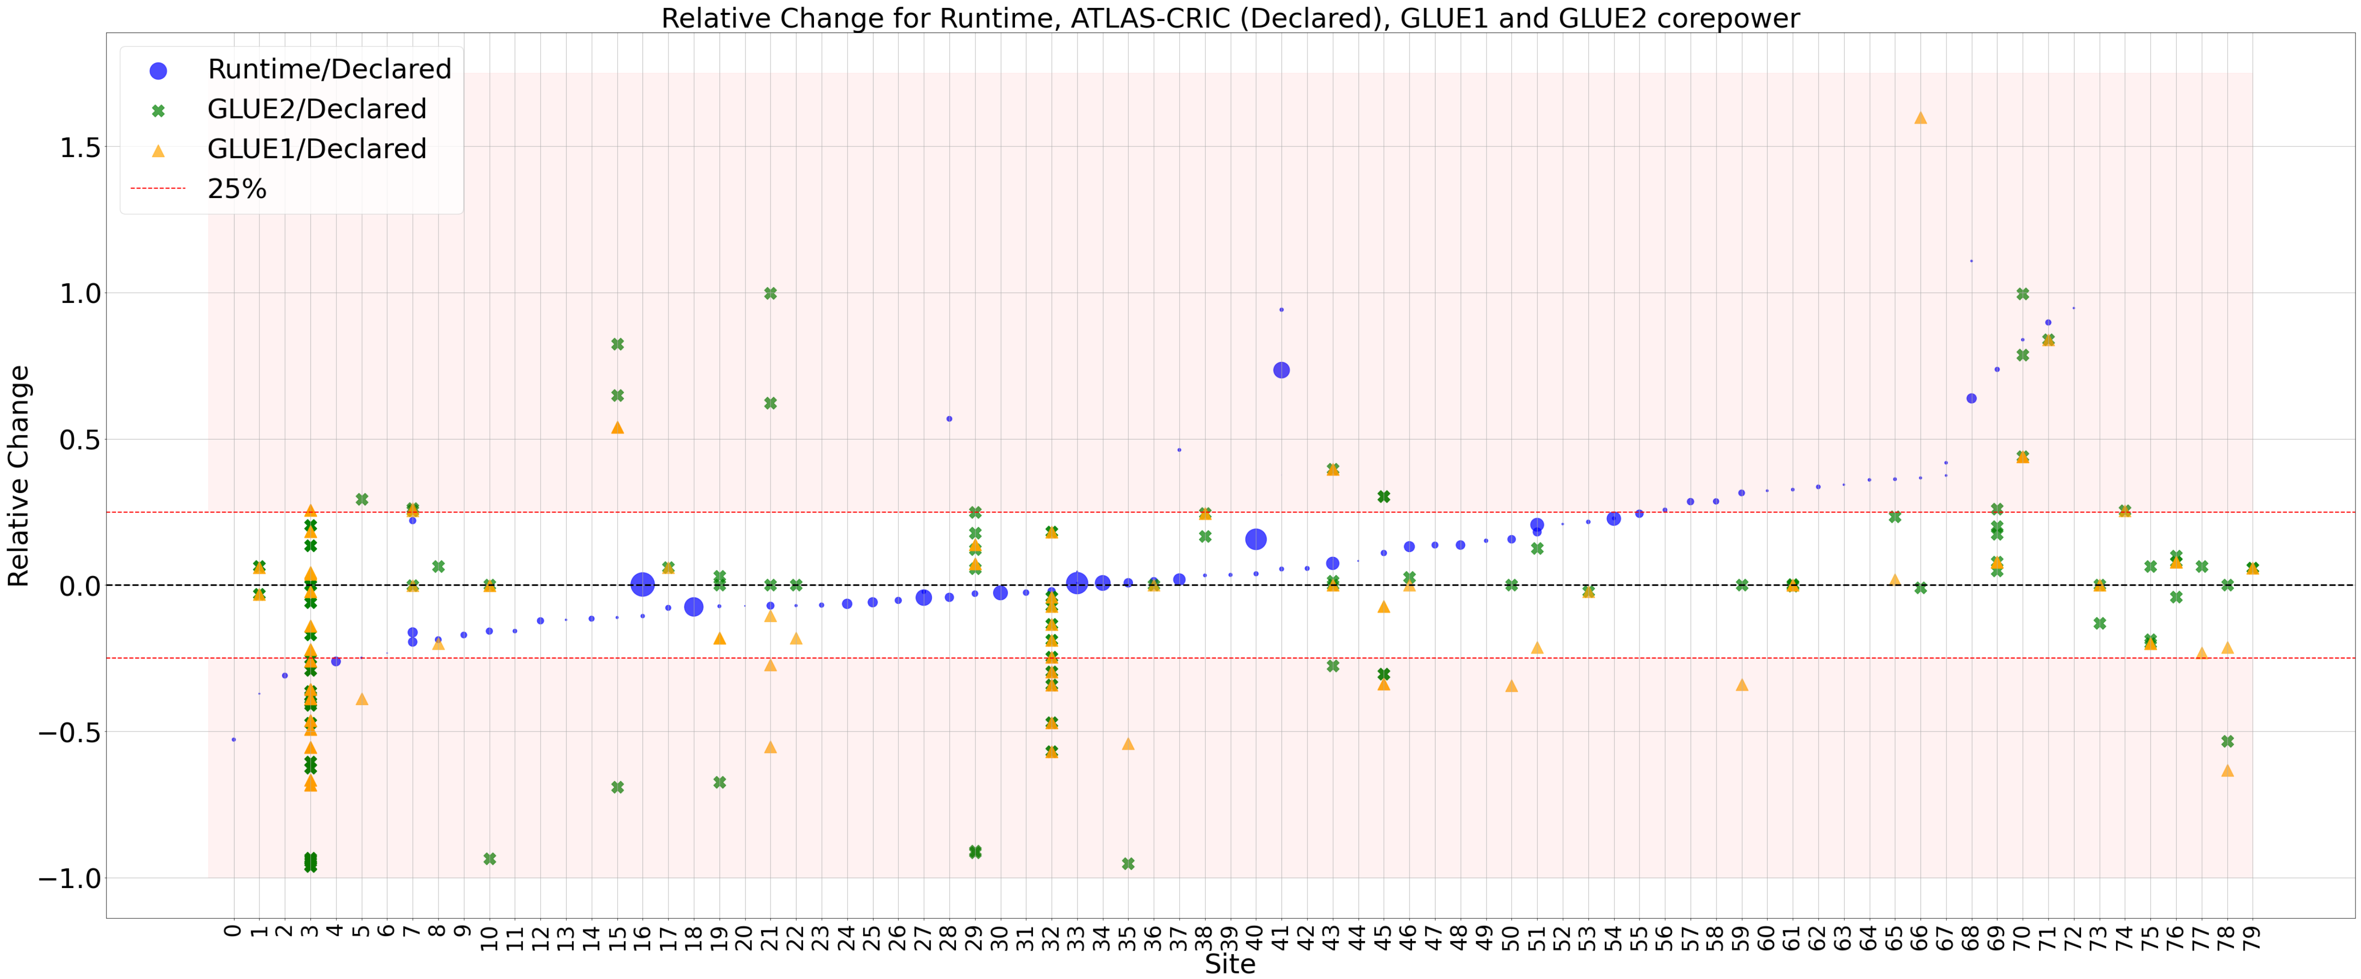

- Multiple entries for a given site, due to multiple CE (queues) available
- Discrepancies between measured corepower vs different data sources

Summary

- ATLAS-CRIC corepower values are different comparing to the measured runtime corepower values
  - Should be fixed
- There are discrepancies beetween ATLAS-CRIC, GLUE1 and GLUE2 data sources
  - Should be fixed
- It would be beneficial to update entries in the different information systems and possibly unify them