

# Centre Informatique National de l'Enseignement Supérieur

HPC

IT Hosting

Archiving

CENTRE INFORMATIQUE NATIONAL DE L'ENSEIGNEMENT SUPÉRIEUR



## Optimizing GPU Frequency for Sustainable HPC: Lessons Learned from a Year of Production on Adastra, an AMD GPU Supercomputer

CUG 25

May 6 2025

Gabriel Hautreux   Naïma Alaoui   Etienne Malaboeuf  
Head of HPC Dept.  
hautreux@cines.fr

# Adastra Supecomputer (installed in 2022)

## GPU Partition

356 GPU nodes :

- 8 AMD MI250X GCD, **64GB HBM2/GCD**
- Same tech as:
- **Frontier, LUMI , ENI**

## Partition CPU

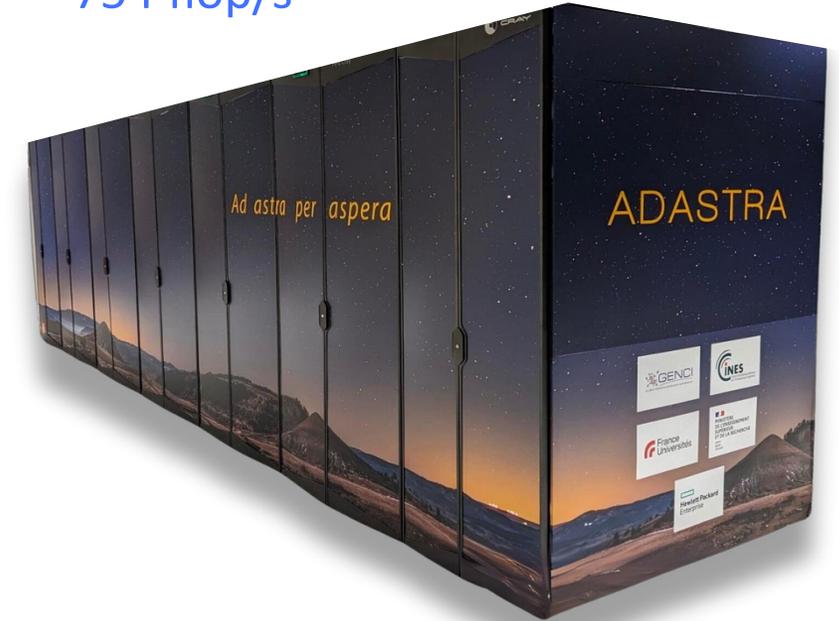
544 CPU nodes:

- 2 AMD Genoa EPYC 9654 96c@2.4 GHz,  
768G DDR5-5200 /node

## High bandwidth network and storage

- Slingshot 200Gb/s
- 2 ClusterStor 2Po SSD +12Po HDD

75 Pflop/s



## Rankings

	June 2022	Nov 2022	June 2023	Nov 2023	June 2024
	4	3	3	3	9
	10	11	12	17	20

# Adastra Supecomputer (installed in 2022)

## GPU Partition

356 GPU nodes :

- 8 AMD MI250X GCD, **64GB HBM2/GCD**
- Same tech as:
- **Frontier, LUMI , ENI**

Partition  
studied

## Partition CPU

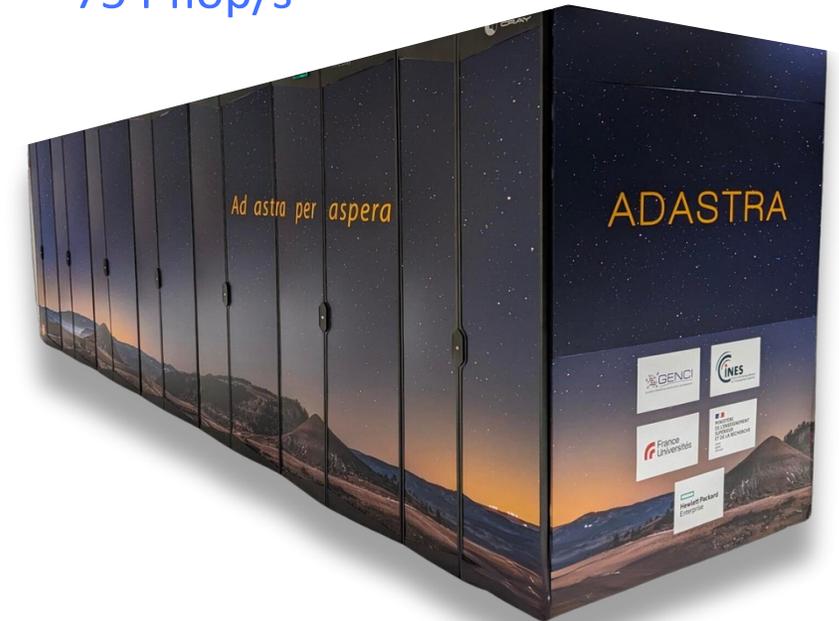
544 CPU nodes:

- 2 AMD Genoa EPYC 9654 96c@2.4 GHz,  
768G DDR5-5200 /node

## High bandwidth network and storage

- Slingshot 200Gb/s
- 2 ClusterStor 2Po SSD +12Po HDD

75 Pflop/s



## Rankings

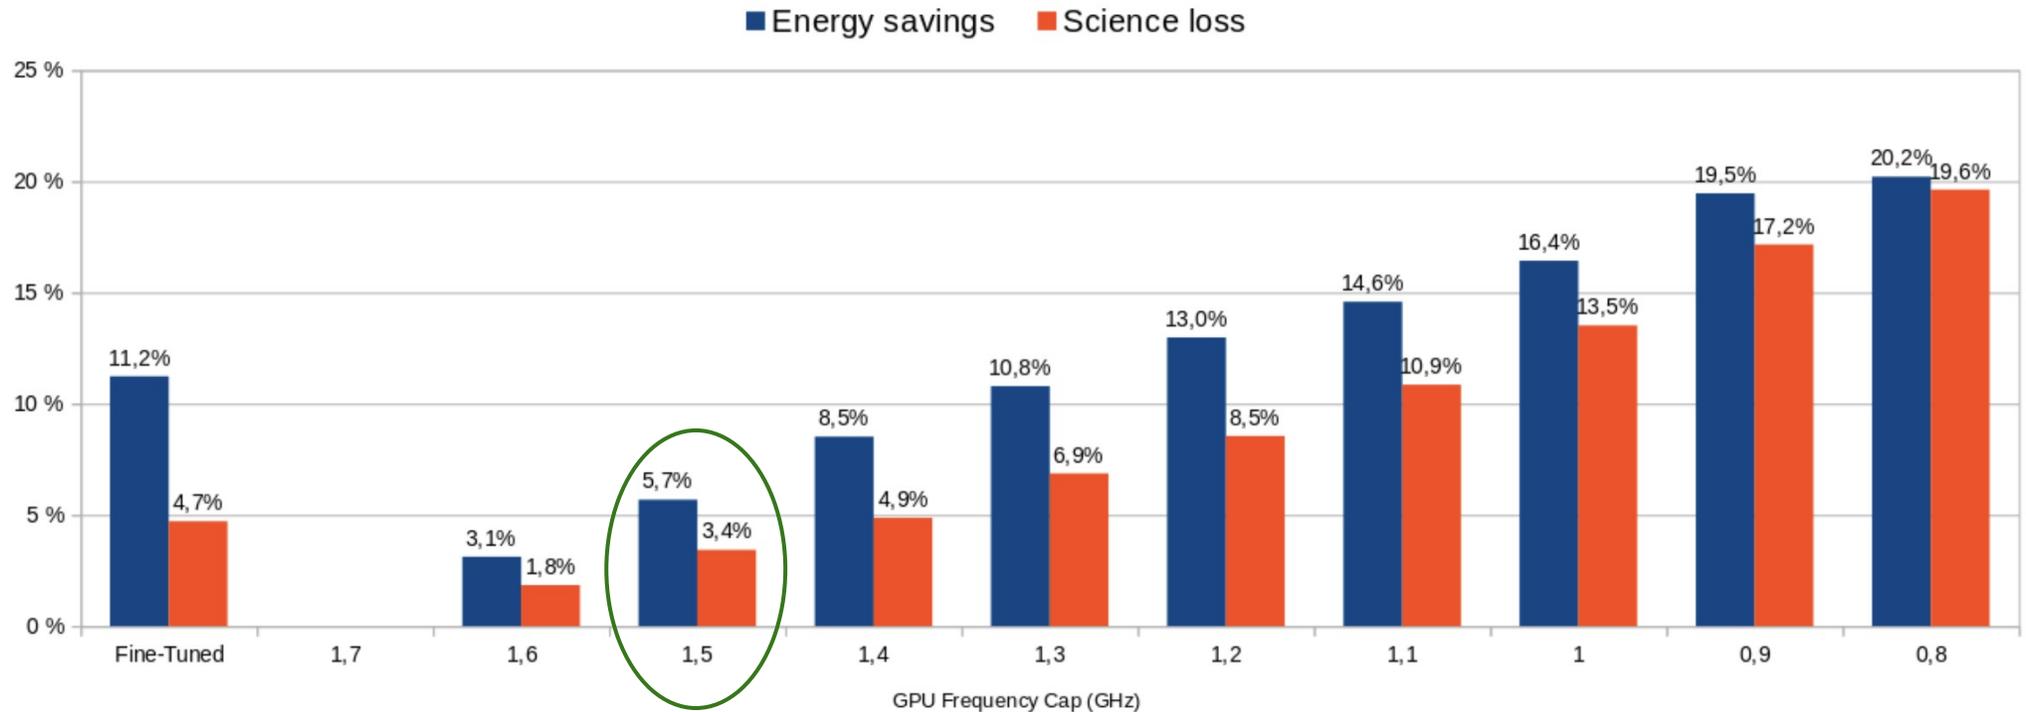


June 2022	Nov 2022	June 2023	Nov 2023	June 2024
4	3	3	3	9



June 2022	Nov 2022	June 2023	Nov 2023	June 2024
10	11	12	17	20

## Initial work (SC23)



<https://doi.org/10.1145/3624062.3624268>

### On our benchmark code base

- Moving from 1.7GHz as the default GPU frequency to 1.5GHz enabled ~6% energy savings for *only* ~3% science loss
- **Due to increasing operational costs, we decided to push it into production in July 2024.**

# User Base Evaluation

## Objective

- Validate user base consistency between the two semesters

## Data Collection

- We used SLURM to retrieve job data for each semester
- Counted common users and measured GPU hours consumed

## Analysis

- 145 active users in both S1 and S2 2024, representing 48% of the total user base.
- These users consumed 95.2% of total GPU hours.

**Consistent user base across semesters, crucial for assessing frequency change impact.**

	Number of users	M of GPU.hours consumed
2024	292	10.693
S1	223	4.794
S2	214	5.899
S1-S2 common	145 (48%)	10.181 (95.2%)

**Table 1: Users and GPU.hours consumption for MI250X Partition**

# Impact of GPU Frequency Capping on Jobs

## GPU@1.7GHz

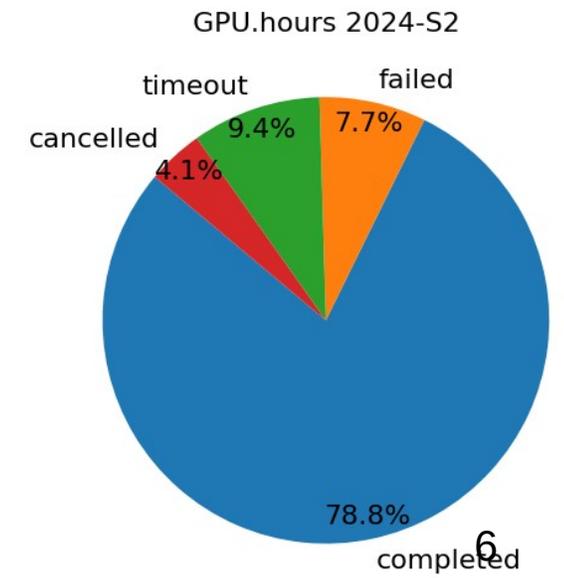
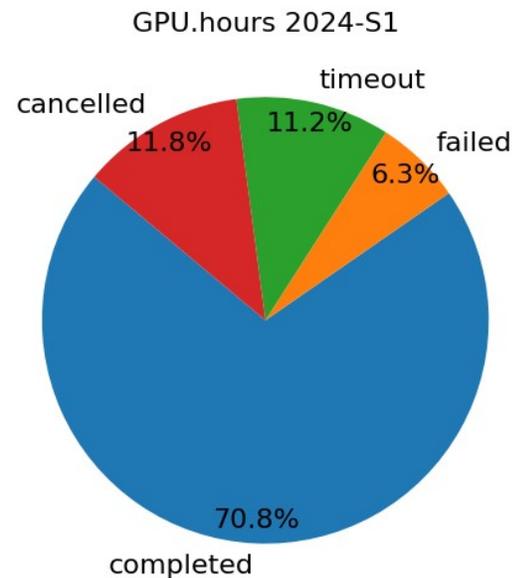
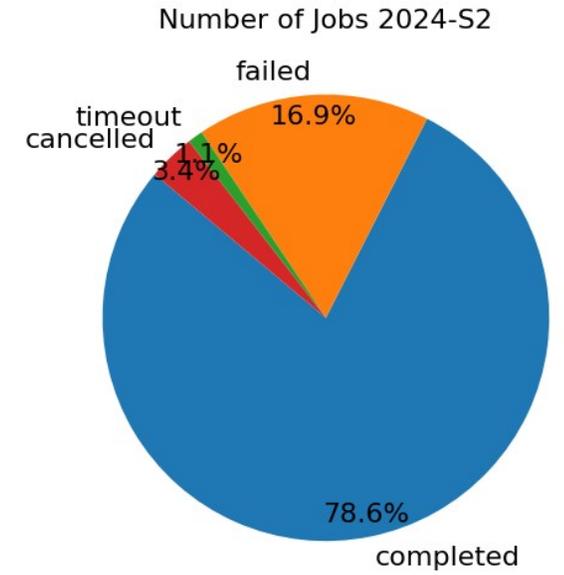
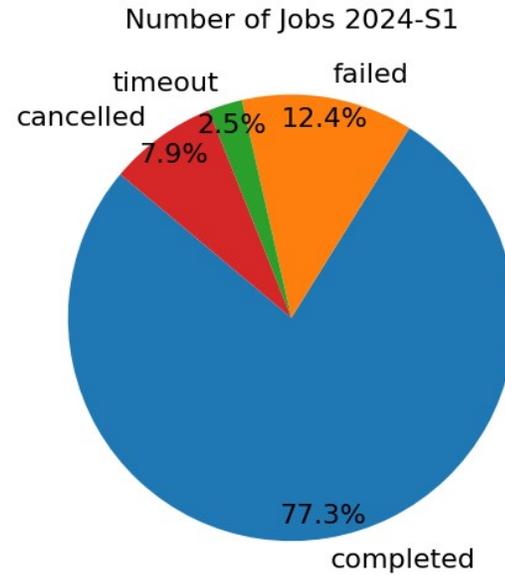
## GPU@1.5GHz

### Number of jobs

- % of completed jobs stays the same
- More failed jobs (+4 points)
- Less timeout jobs

### GPU.hours consumed

- +8 points in completed jobs
- Less cancelled
- Decrease of timeout volume



# Impact of GPU Frequency Capping on Jobs Power Draw

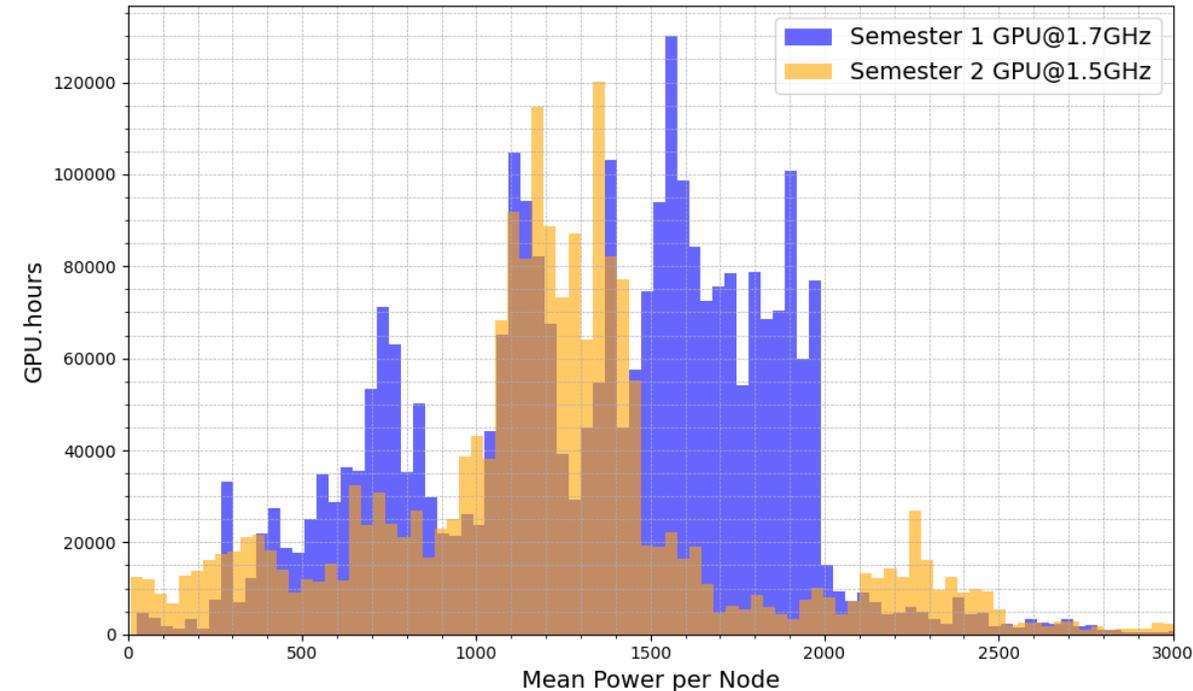
## Data Collection

- Retrieved job energy data from SLURM (using performance counters with a 30-second sampling rate).
  - Removed jobs with missing energy information.
  - Excluded short-duration jobs (less than 10 minutes).
- Calculation: Computed mean power draw per node for each job using the formula:

$$\text{Mean Power Draw per Node} = \frac{\text{Total Energy Consumption(J)}}{\text{Number of Nodes} \times \text{Time (s)}}$$

## Analysis

- Noted a shift in mean power draw from [1500-2000W] to [1000-1500W] between semesters.
- Identified an unexpected bump around [2200-2400W].

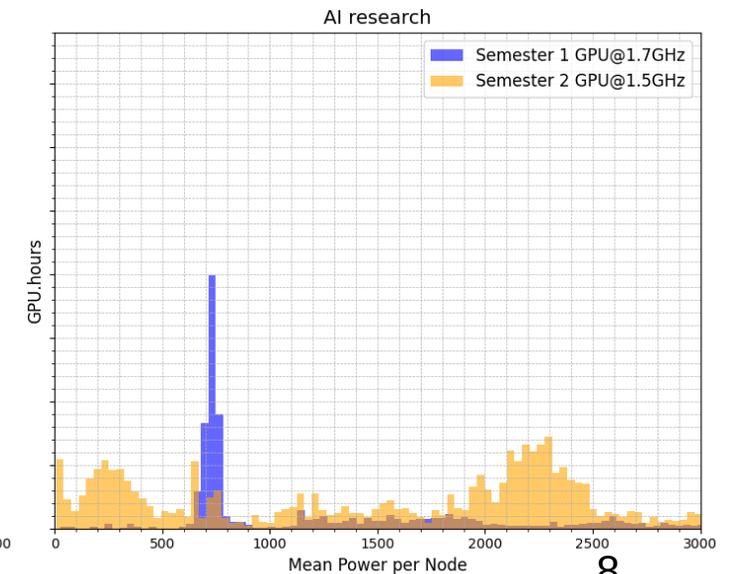
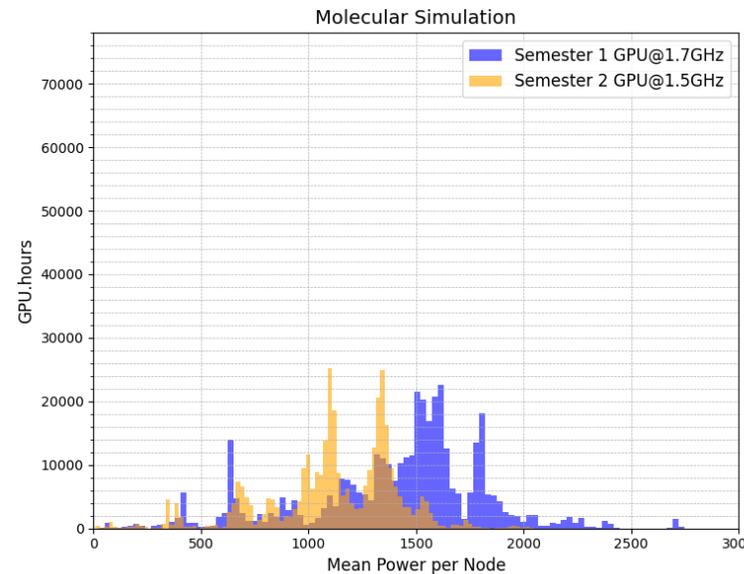
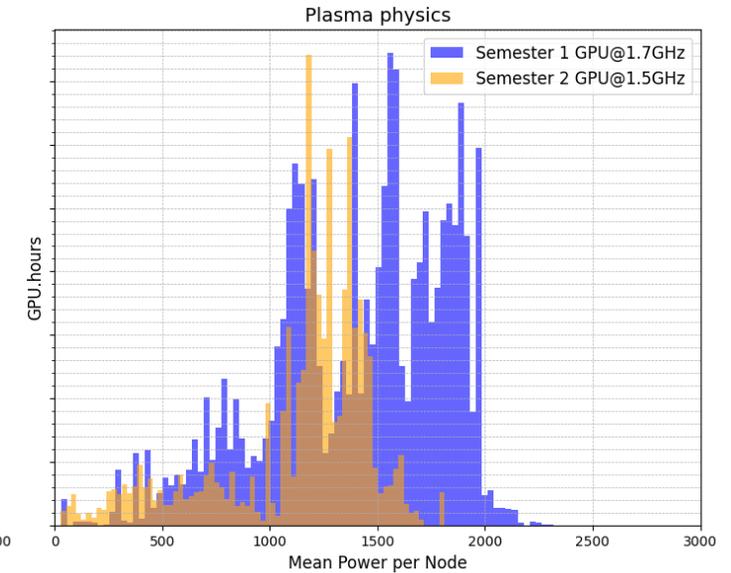
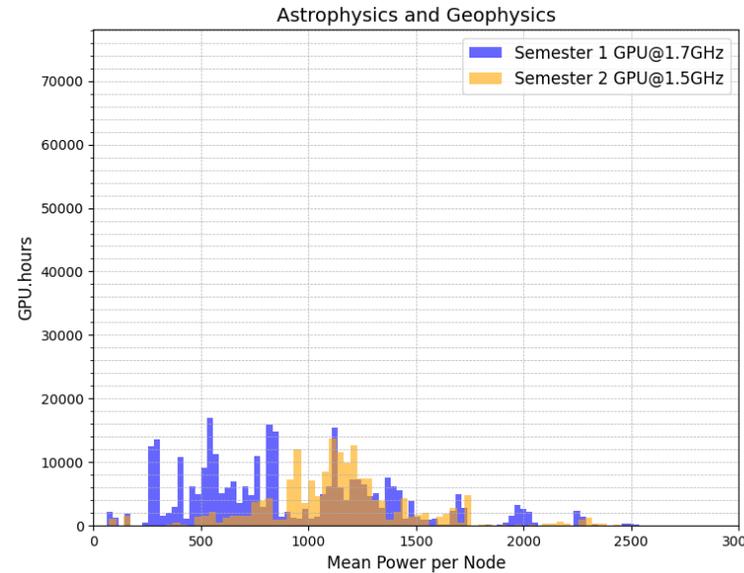


# Impact of GPU Frequency Capping on Jobs Power Draw

## Scientific Fields Analysis

- Focused on four fields:
  - Astrophysics and Geophysics
  - Plasma Physics
  - Molecular Simulation
  - AI Research.
- Observed significant shifts in power draw for Plasma Physics and Molecular Simulation.
- AI Research showed intense GPU usage in the second semester, with a mean power draw in the [2000-2500W] range, compared to other fields

Lack of data for AI workloads



# Impact of GPU Frequency Capping on AI Workloads

## Method

- Milabench used:  
<https://github.com/mila-iqia/milabench>
- Benchmark runs at 1.7GHz and 1.5GHz

## Analysis

- Max TTS increase ~4%, Mean 1.36%
- Max ETS decrease ~15%, Mean 2.11%

**This seems to validate the limited impact of the GPU frequency cap for commonly used AI workloads. Also explaining the still high power draw of AI workloads on our system.**

AI Benchmark	Time to solution	Energy to solution
bert-fp16	+1.13%	+3.27%
bert-fp32	-3.15%	-1.43%
bert-tf32	+1.17%	-1.76%
bert-tf32-fp16	+2.82%	+0.89%
dqn	+1.11%	-0.34%
lightning-gpus	+4.11%	-3.10%
pna	-0.86%	-4.35%
ppo	+4.35%	-15.07%
reformer	+0.28%	-4.57%
resnet152-ddp-gpus	-0.59%	+4.39%
t5	+1.70%	+2.66%
whisper	+4.21%	-5.93%
Average	+1.36%	-2.11%

**Table 2: Impact of GPU frequency change from 1.7GHz to 1.5GHz on AI workloads**

# Impact of frequency change on Aadastra MI250X partition

S1 : 1.7GHz 80% Activity

S2 : 1.5GHz 86 % Activity



## Key numbers:

+6% activity  
-7% mean power draw

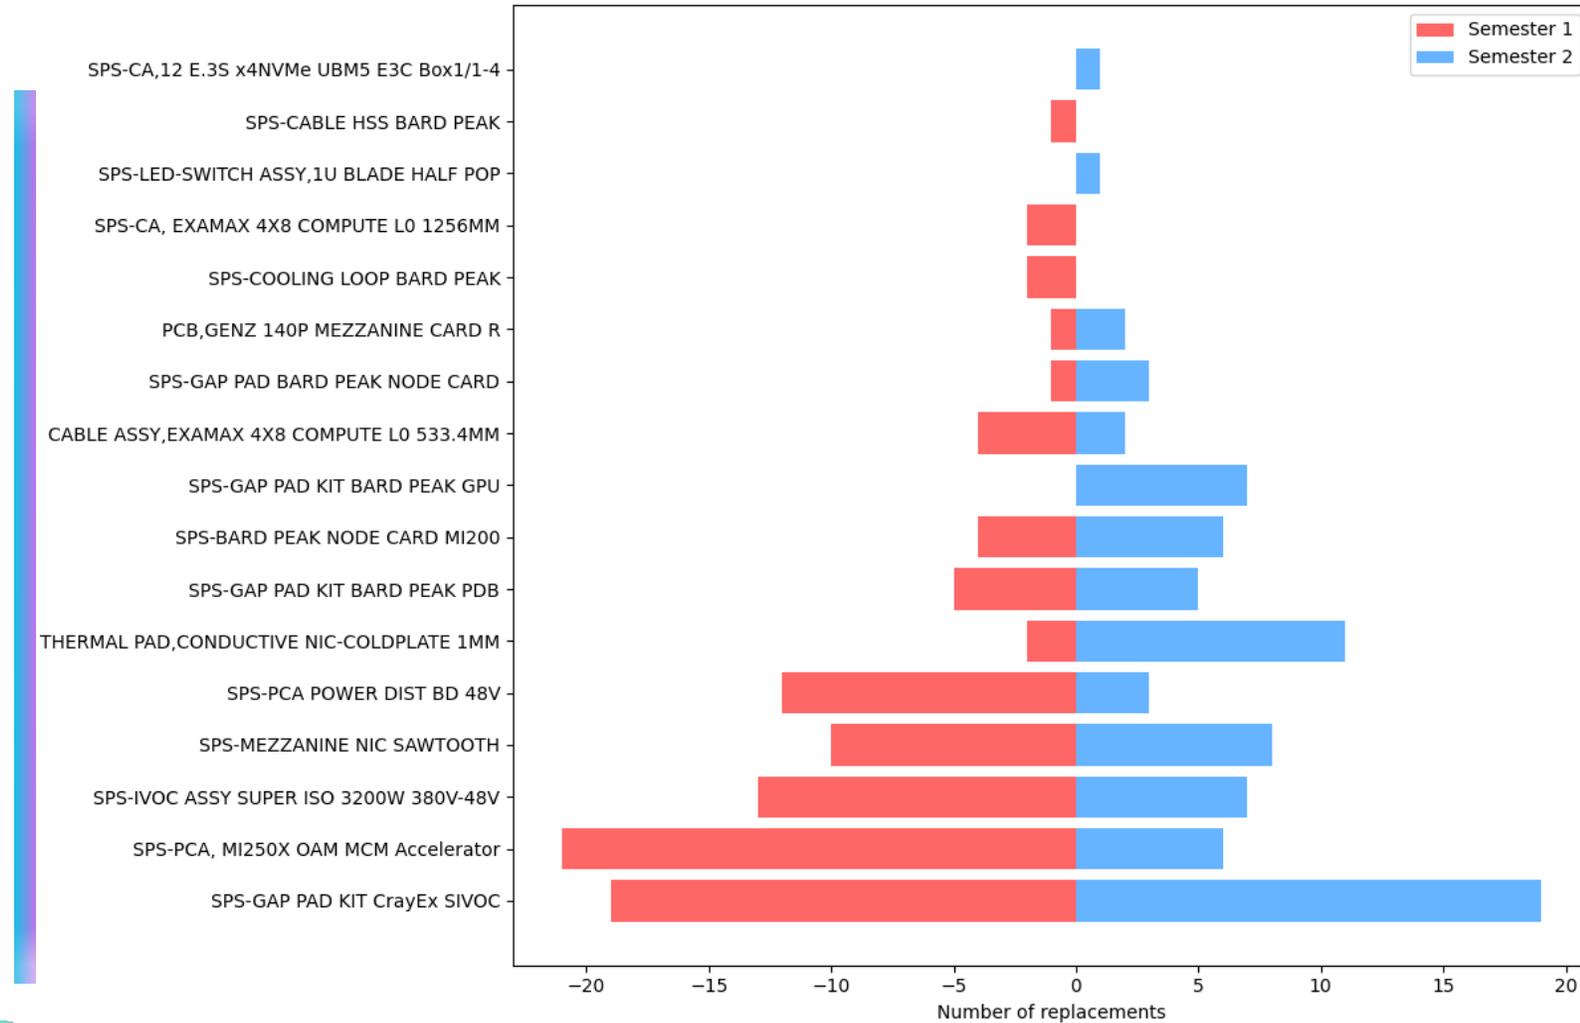


~14% less power per active node



~ 100k€ of energy savings

# Impact of GPU Frequency on Hardware Reliability



## Data Collection

- Infos sourced from HPE's intervention database

## Findings

- Reduction in MI250X OAM MCM Accelerator (GPU card) replacements
  - 21 GPUs replaced in the first semester, compared to 6 in the second semester
- Need to run the same assesment on multiple sites, we will continue to operate at 1.5GHz and can share the data

## User Satisfaction Impact

### Context

Users were unaware of the frequency change until the end of 2024.

### Method

- We analyzed user tickets database using simple classification and LLM inference.

### Findings

- No relevant information found regarding performance degradation.
- Cannot definitively conclude that no users were impacted, but that is the feeling of the user support team

## Conclusion

### Impact of frequency capping

- GPU frequency capping from 1.7GHz to 1.5GHz reduced energy consumption
- No substantial impact on computational performance observed
- Hardware reliability could be positively impacted by this policy

### Future work

- Further investigation needed for hardware stability and user satisfaction (multi-sites/HPE collaboration?)
- Continued monitoring of the Aadastra system for longer-term effects.

### Study expected impact

- Highlights the role of frequency management in HPC systems for energy and financial savings
- Discuss future strategies for sustainable and efficient computing.

MI250X Partition (356 nodes)	S1	S2	Gain
Active nodes	285.44	307.09	+7.6%
Mean Power Draw (kW)	552.75	513.16	-7.2%
Power Draw Standard Deviation (kW)	146.75	114.55	-22%
Mean Power per Node (kW)	1.94	1.67	-14%

**Table 3: Summary of GPU Frequency Change from 1.7GHz to 1.5GHz impact on Aadastra MI250X partition**

# Centre Informatique National de l'Enseignement Supérieur

HPC

IT Hosting

Archiving



MINISTÈRE  
DE L'ENSEIGNEMENT  
SUPÉRIEUR,  
DE LA RECHERCHE



Thank you

Questions?

---

Gabriel Hautreux   Naïma Alaoui   Etienne Malaboeuf  
Head of HPC Dept.  
hautreux@cines.fr