

GPU usage practices of the STARS team

Inria Center at Université Côte d'Azur

Please note: Everyone in or affiliated to the STARS team who uses any calculation means of Inria (NEF, Abaca, ...) automatically agrees to the rules listed in this document entirely. Otherwise, each violation will be discussed and consequences will be drawn individually.

We (the STARS team) have various computational nodes, each have a few GPUs, always the same type within a node.

Specifically, this is the pool of the most powerful nodes we currently¹ have:

- nefgpu39 - 4x RTX 8000, each 48GB of RAM (the RAM of a single GPU)
- nefgpu46 - 3x Tesla V100, each 32GB of RAM
- nefgpu56 - 3x A100, each 40GB of RAM
- Other GPUs added to the STARS team resources added after the day of ¹.

(More about the nodes can be found here:

<https://wiki.inria.fr/ClustersSophia/Hardware> (while searching for "STARS")).

Naming convention remark:

The nodes on NEF are currently¹ being migrated to Abaca. Therefore, their names are changing. Not all the names are known at the moment. However, this guideline concerns the STARS' GPUs in any state - shall they be currently placed on NEF, Abaca or somewhere else. Therefore, if someone abuses the usage of the GPUs e.g. referred above as nefgpu56, which at that moment will be placed under a different name on Abaca or such, it will still be considered as a violation and consequences will be drawn.

The usage guidelines (points 1-5) are as follows:

Short Version

(See the Precise Version below for details and clarifications. In case of ambiguities, the content of the Precise Version is privileged)

- 1) Long operations (training models on multiple GPUs) should be performed on GENCI or on public GPUs. NEF, Abaca is for inference and evaluations.
- 2) Occupy minimal resources. By default, take 1 GPU at a time. Set gpunum=1 in oarsub. Increase only when the job runs out of memory. Allocate low walltime.

¹ Status as of 14.05.2025

3) Monitor your jobs. Check on your jobs often and kill them if they give you wrong results. Each job should be able to stop without big consequences.

4) Be available. Everyone with submitted or running jobs should be available for contact, even during the weekend.

5) For non-critical jobs, use the Best Effort mode. Set -t besteffort in oarsub. You will have available nodes of other teams and your jobs could (but almost never will) be killed by a more important job.

Precise Version

Point 1.

NEF, Abaca and other local calculation resources at Inria are not designated for long operations, such as training models or using multiple GPUs (be it for training or for other operations). These local resources are intended for quick experiments (no training), debugging, inference, and evaluations. For training and other operations as well as for using multiple GPUs, [GENCI](#) computing platform must be used (requires an access request).

Please note: If you do not have access to GENCI, then you must either collaborate with someone else who does (e.g. use their account) or use lower/weaker/less powerful GPUs from our team (still, without extensive usage, blocking, etc.). GPU nodes referred on NEF as nefgpu39, nefgpu46 and nefgpu56 or newer (including the naming convention remark above) cannot absolutely be used for the GENCI-replacement purposes as stated above.

Point 2.

Occupy minimal resources. Set gpunum=1 in oarsub. It is strictly forbidden to use all the GPUs of any node from nefgpu39, nefgpu46, nefgpu56, or newer (including the naming convention remark above) per one person and/or project (in case of collaborations or joint work) at the same time.

Further, usage of more than 1 GPU at a time from the pool of GPUs from all the 3 powerful nodes must be well justified, requested and approved. Any objection or request from another side (another STARS team member) must result in immediate reduction to using only 1 GPU from the powerful pool by the user.

In case of the other STARS nodes, different than nefgpu39, nefgpu46, nefgpu56 or newer (including the naming convention remark above), max k-1 of GPUs can be used at the same time, where k indicates the number of GPUs at the given node.

Further, allocate low walltime (applicable to all STARS' nodes and GPUs). Max 8 hours in interactive mode. Same for the passive mode. For longer jobs, see Point 5.

Point 3.

Monitor your jobs. Check on your jobs often and kill them if they give you wrong results. Each job should be able to stop without big consequences.

Point 4.

Be available. Everyone with submitted or running jobs should be available for contact, even during the weekend, day off or a national holiday.

Point 5.

For non-critical jobs, use the Best Effort mode. Set -t besteffort in oarsub. You will have available nodes of other teams and your jobs might be killed (not too often though) by a more prioritized important job.

If you really need to perform longer operations on local calculation resources at Inria without stopping with the 8 hour intervals, you must do it with besteffort (especially with other team GPU's, but possibly also with ours).