CINECA

How to Run Workflows and AI Jobs on Leonardo

AI-Friendly EuroHPC Systems

L. Morselli

CINECA: A GROWING CONSORTIUM

SUPPORTING THE ITALIAN ACADEMIC SYSTEM SINCE 1969

CINECA

117 MEMBERS

Ministry of University-Research and Ministry of Education, Universities, Public Research Organisations

5 LOCATIONS

Bologna, Milano, Roma, Napoli, Chieti

≃1100

Employees

LEONARDO HPC

THE 7TH MOST POWERFUL SUPERCOMPUTER IN THE WORLD

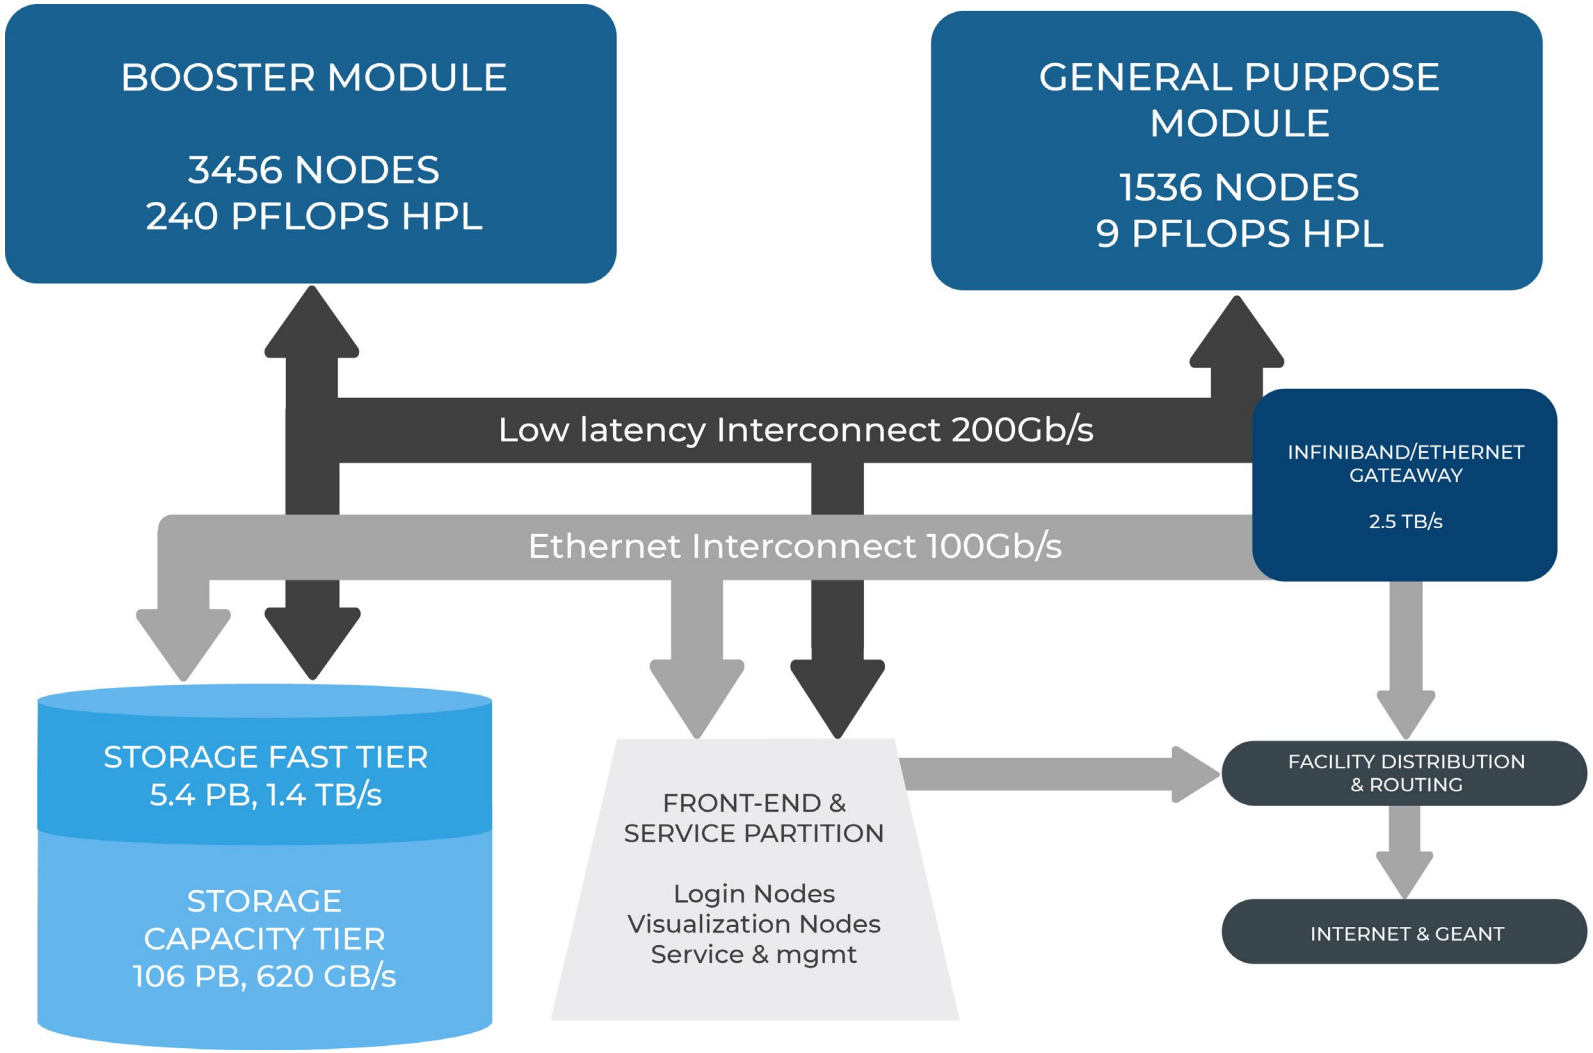

LEONARDO + LISA

TOWARDS AN AI OPTIMIZED INFRASTRUCTURE

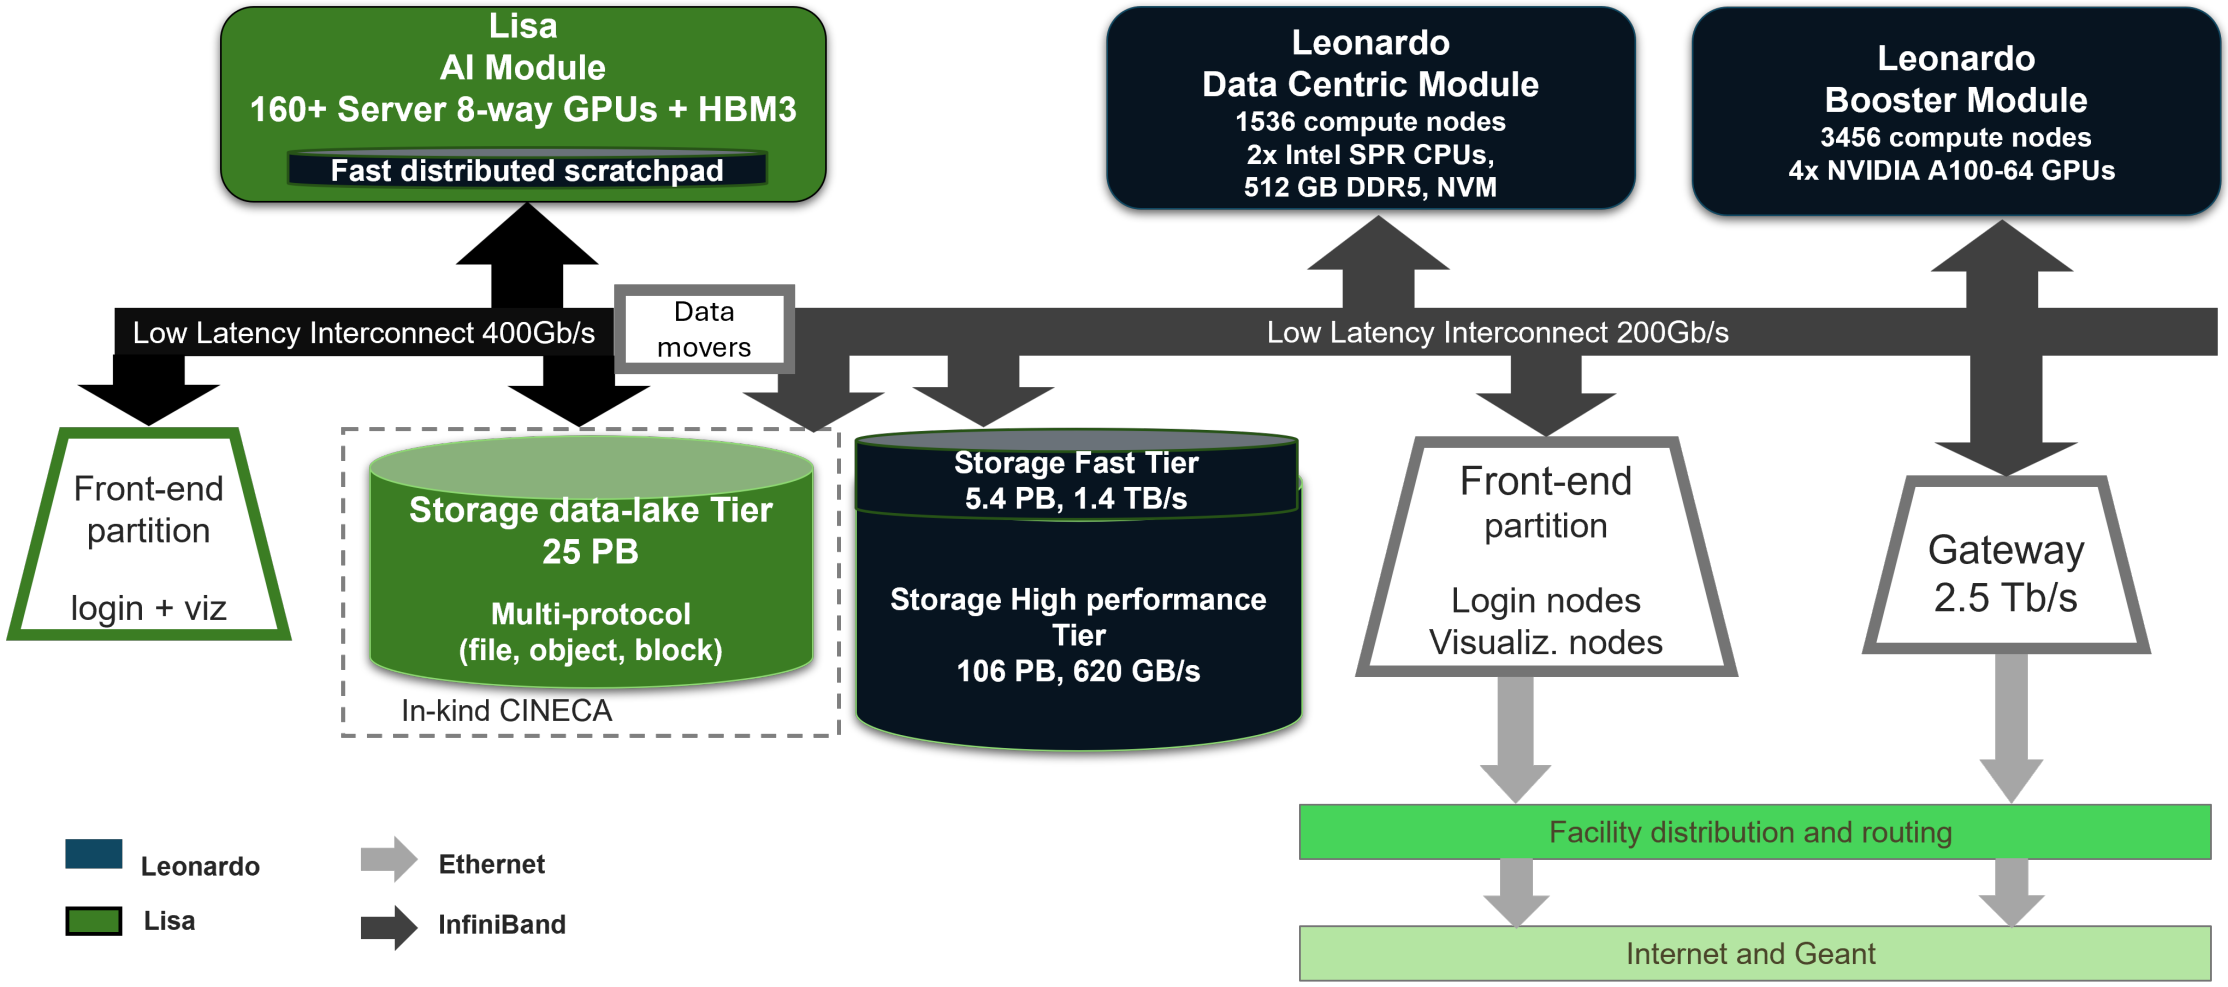

ACCESS TO COMPUTE

THE ITALIAN SUPERCOMPUTING RESOURCES ALLOCATION ORGANIZATION

- **Class B (Max 250'000 GPU hours on Leonardo Booster – 12 months)**
  Class B projects are received twice a year. They go under peer-review evaluation and a 5 months period is expected before access to HPC resources.

- **Class C (Max 10'000 GPU hours on Leonardo Booster – 9 months)**
  Class C projects are received through continuous submission and reviewed once per month. An average period of about 30 days is required for activating the project.

Projects' PIs need to be affiliated to an Italian research institution, while no restriction is applied for the Co-PI and collaborators. It is expected that the research will be performed at Italian institutions.

CINECA provides **high-level technical support** to each project through its User Support Group. Dedicated specialist support can be requested for the enabling and optimization of the applications necessary for the project.

ACCESS TO COMPUTE

THE EuroHPC ACCESS CALLS

- **Benchmark Access** → 3’500 node hours (on BOOSTER, 3 months)
- **Development Access** → 4’500 node hours (on BOOSTER, 1 year)
- **Regular Access** → 52’000-220’500 node hours (on BOOSTER, 1 year)
- **AI and Data-Intensive Applications** → 50’000 node hours (1 year)
- **Extreme Scale** → min 245’000 node hours (on BOOSTER, 1 year)

The CINECA AI Ecosystem

The cineca-ai modules: ready-to-use environments always up to date, optimised for the Leonardo GPUs and prepared by CINECA experts.
https://gitlab.hpc.cineca.it/cineca-ai/cineca-ai

Transformers · pandas · PyTorch · scikit learn · python · TensorFlow · jupyter · NumPy · dask · BeautifulSoup

Wiki Cineca-AI

CINECA-AI versions

- CINECA-AI 2.0.2
- CINECA-AI 3.0.0
- CINECA-AI 3.0.1
- CINECA-AI 4.0.0
- CINECA-AI 4.1.0
- CINECA-AI 4.1.1
- CINECA-AI 4.3.0

CINECA-AI 4.3.0

Package	Version
absl-py	2.1.0
accelerate	0.27.0
aiohttp	3.9.3
aiosignal	1.3.1
alabaster	0.7.16
annotated-types	0.6.0
anyio	4.3.0
apex	0.1
appdirs	1.4.4
argon2-cffi	23.1.0
argon2-cffi-bindings	21.2.0
arrow	1.3.0
asgiref	3.8.1
asttokens	2.4.1

Transfer hundreds of TBs of data!

Datamovers are dedicated nodes on each HPC cluster that are designed for transferring data FROM/TO a cluster. Available transfer protocols: **rsync, scp, sftp, wget, curl.**

Workflow and Data Management in progress

Pachyderm
version-controlled data science

mlflow

CINECA AI Accademy

HPC Courses Catalogue

23 HPC and general computing courses

python · julia · R · OpenACC · LEONARDO CINECA

15 AI, machine and deep learning focused courses

jupyter · PyTorch · TensorFlow · mlflow · dask

14 HPC, coding and domain-specific schools

GPU · BBS · blender

Interaction with the SLURM scheduler

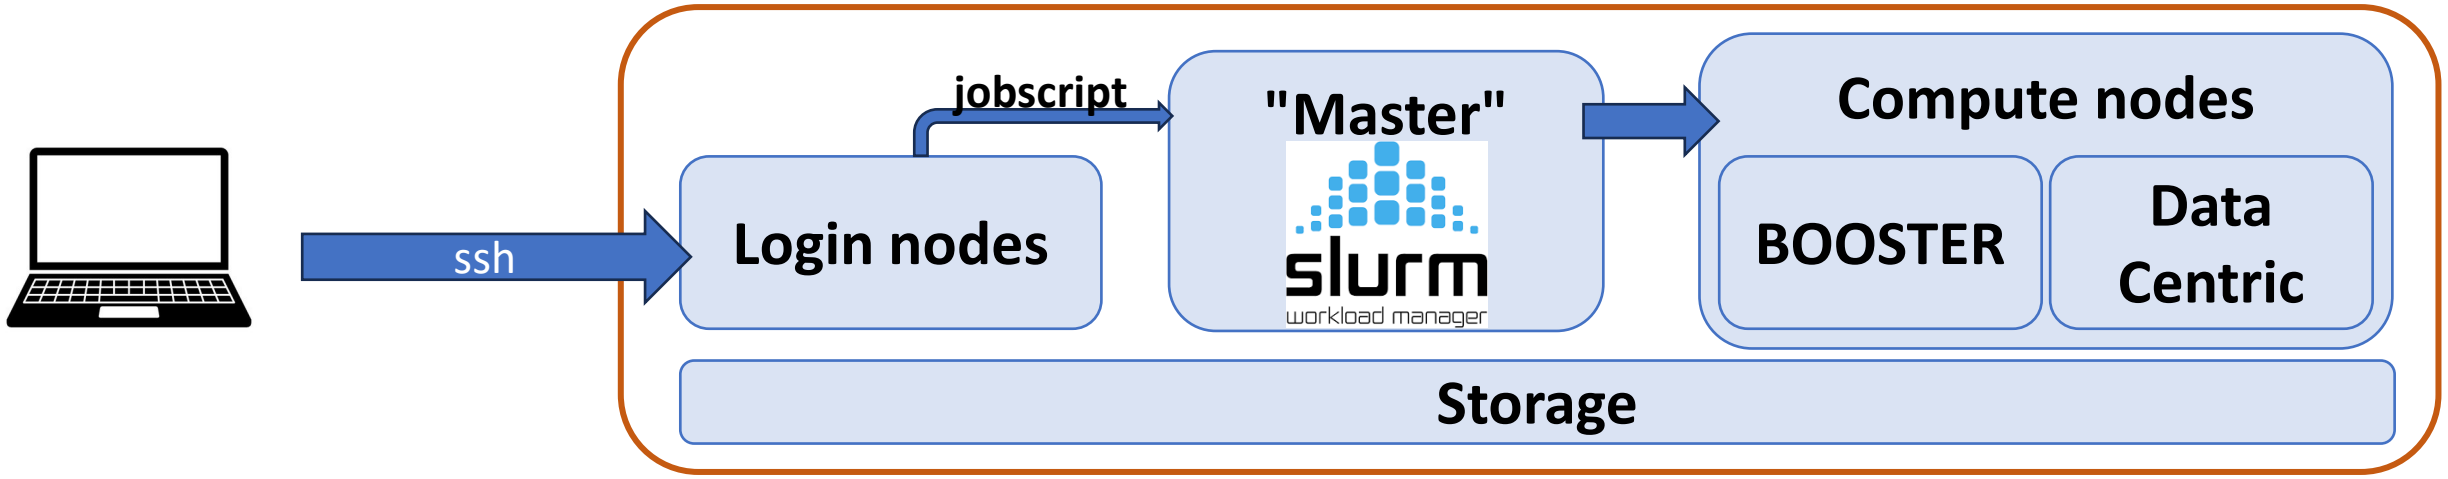


```
 GNU nano 2.9.8                    job_inference.sh                          Modified

#!/bin/bash
##SBATCH --out=%j.out
##SBATCH --err=%j.err
#SBATCH -N 1
#SBATCH --cpus-per-task=8
#SBATCH --exclusive
#SBATCH --gres=gpu:2
#SBATCH --partition=boost_usr_prod
#SBATCH -A cin_staff
#SBATCH --time=00:30:00
#SBATCH --qos=boost_qos_dbg


module purge
module load python
source env_iscra/bin/activate

#echo 'all modules loaded'

echo "path : $path"

torchrun --nproc_per_node 2  /leonardo_scratch/large/userinternal/lcavall1/llama/prova_def.py \
    --path $path \
    --ckpt_dir /leonardo_scratch/large/userinternal/lcavall1/llama/llama-2-13b-chat/ \
    --tokenizer_path  /leonardo_scratch/large/userinternal/lcavall1/llama/tokenizer.model\
    --max_seq_len 1536 --max_batch_size 50
```

A continuous improvement

Challenge	Solution
Slow I/O performance on HDD storage	Improved I/O on Full-Flash storage
AI jobs with long runs / high number of nodes may experience system instability	Sharing best practices on check points
Easy the uptake of the CINECA AI environment for the AI community	- Improve the usability of CINECA AI env - Promote the awareness of the HPC services to the AI communities
Readiness of the users: - Queues vs reserved resources - Lack of exposure of AI community to port AI environment on large scale AI supercomputers	- Improve skills the AI community on an efficient and sustainable use of the AI systems - HPC centres can support and train AI community

Support to AI users

General "HPC"- related support services, as for other communities

Dedicated services tailored for the AI scientific/industrial communities working on the challenges AI factory should address

- Level 0: Awareness of the need of AI community and of the HPC user support service
- Level 1: Support for porting AI applications and workflows on HPC infrastructures
- Level 2: Support for the use and mastering of AI libraries on HPC Architecture
- Level 2: Support for upscaling of workflows
- Level 3: Support to engineering Data pipelines
- Level 3: Support for the pre-training of large-scale AI models
- Level 3: Support for the specialization of various AI models

Efficient integration with existing and upcoming service-delivering European Initiatives (e.g. EPICURE, upcoming MINERVA)

CINECA

Thank you