

# General informations

The hardware structure of the DIFA-OPH computing cluster is summarised in the table below, listing all the compute nodes currently available with their individual hostnames, their specific resources (number of cores and available RAM), and their access policies.

In particular, bldNN nodes are part of the **BladeRunner island**, mt>NN nodes are part of the **Matrix island**, and gpuNN are part of the **GPU island**.

**The nodes associated with the OPH project are open to all users** while nodes associated with other individual projects may be subject to **access restrictions**. Such restrictions are indicated in the last column of the table.

- **Shared** nodes, the access is available for all users
- **Teaching** nodes, reserved for teaching purposes: can be accessed only by students of specific DIFA courses for their laboratory activities
- **Reserved nodes**, the access is restricted to users explicitly authorised by the corresponding project PI till the given **expiration date**.

Nodes	vCPUs/RAM	GPUs	Project (PI)	Access type (exp)
bld[01-02]	24 / 64G	-	OPH	Shared
bld[03-04]	32 / 64G	-	OPH	Shared
bld05	32 / 128G	-	OPH	Shared
bld[15-16]	16 / 24G	-	OPH	Teaching
bld[17-18]	32 / 64G	-	OPH	Shared
mtx[00-15]	56 / 256G	-	OPH	Shared
mtx[16-19]	112 / 512G	-	ERC-Astero (Miglio)	Reserved (2024-09) Shared
mtx20	112 / 1T	-	OPH (Di Sabatino)	Shared
mtx[21-22]	192 / 1T	-	SLIDE (Righi)	Reserved (2025-10)
mtx[23-25]	112 / 512G	-	OPH (Marinacci)	Shared
mtx26	112 / 512G	-	CAN (Bellini)	Reserved (2026-04)
mtx27	112 / 512G	-	FFHiggsTop (Peraro)	Reserved (2026-04)
mtx[28-29]	112 / 1.5T	-		
mtx30	64 / 1T	-	Trigger (Di Sabatino)	Reserved (2027-02)
mtx[31-32]	64 / 1T	-	EcoGal (Testi)	Reserved (2027-08)
mtx[33-34]	192 / 1.5T	-		
mtx[35-36]	192 / 512G	-	RED-CARDINAL (Belli)	Reserved (2028-04)
mtx[37-40]	192 / 512G	-	ELSA (Talìa)	Reserved (2026-12)
gpu00	64 / 1T	2xA100	VEO (Remondini)	Reserved (2026-04)
gpu[01-02]	112 / 1T	4xH100	EcoGal (Testi)	Reserved (2027-08)
gpu03	112 / 1T	4xH100	ELSA (Talìa)	Reserved (2026-12)

## Computing Resources

Resources are nodes, CPUs, GPUs<sup>1)</sup>, RAM and time. You'll have to select the resources you need for the job. Do not overestimate too much or you'll be "billed" too much. But don't underestimate or your

job won't be able to complete. **When a job completes, you receive a mail with seff output: this should help a lot to optimize future requests.**

Nodes are grouped by partitions. **DO NOT** specify neither nodenames nor partitions unless directed to do so by tech staff.

## Selecting nodes

To select a (set of) node(s) suitable for your job, **use constraints**. These include:

- **blade**: older nodes, usually for smaller/sequential jobs, quite heterogeneous
- **matrix**: newer nodes, for bigger parallel jobs; **allocated by “half node” units!**
- **ib**: require IB-equipped nodes (**all** nodes in matrix are IB-equipped ⇒ no need to specify)
- **filetransfer**: ask for a node with fast access to outside network to quickly transfer big files
- **intel**: require an Intel CPU
- **amd**: require an AMD CPU
- **avx**: require that the CPU supports AVX instructions
- **dev**: require that the node can be used to compile (deprecated: all nodes host build tools)
- **dida**: require nodes used for lessons (obsolete)
- **gpu**: require a GPU-equipped node

## Reserved nodes

Some nodes are reserved for specific projects (see table above). To be able to use 'em you have to be explicitly allowed by project manager (= added to the project group via DSA interface). Once you're in the allowed group (check with `id`) you can submit jobs specifying `-reservation=prj -....`

Project	Manager	AD group (DSA)	OPH group (id)	Reservation to use
CAN	Bellini	Str04109.13664-OPH-CAN	OPH-res-CAN	prj-can
ECOGAL	Testi	Str04109.13664-OPH-ECOGAL	OPH-res-ECOGAL	prj-ecogal
ELSA	Talia	Str04109.13664-OPH-ELSA	OPH-res-ELSA	prj-elsa
FFHiggsTop	Peraro	Str04109.13664-OPH-FFHiggsTop	OPH-res-FFHiggsTop	prj-ffhiggstop
RedCardinal	Belli	Str04109.13664-OPH-RedCardinal	OPH-res-RedCardinal	prj-redcardinal
SLIDE	Righi	Str04109.13664-OPH-SLIDE	OPH-res-SLIDE	prj-slide
Trigger	DiSabatino	Str04109.13664-OPH-Trigger	OPH-res-Trigger	prj-trigger
VEO	Remondini	Str04109.13664-OPH-VEO	OPH-res-VEO	prj-veo

## QualityOfService

What other clusters call “queue”, Slurm calls QoS.

By default all jobs are queued as “-qos=normal”.

Each QoS offers different features:

QoS	Max runtime	Priority	Note
normal	24h	standard	Default
debug	15'	high	Max 2 nodes, 1 job per user, <b>not billed</b>
long	72h	low	

## Storage Resources

Storage is detailed [on its own page](#).

It's important to select the correct storage for the use you're going to do.

1)

Only on GPU nodes

From:

<https://apps.difa.unibo.it/wiki/> - **Wiki DIFA**

Permanent link:

<https://apps.difa.unibo.it/wiki/oph:cluster:resources>

Last update: **2025/02/04 05:56**

