

Infrastructure

Tier2 clusters of Ghent University

The Stevin computing infrastructure consists of several Tier2 clusters which are hosted in the S10 datacenter of Ghent University.

This infrastructure is co-financed by FWO and Department of Economy, Science and Innovation (EWI).

Tier-2 login nodes

Log in to the HPC-UGent Tier-2 infrastructure via <https://login.hpc.ugent.be> or using SSH via `login.hpc.ugent.be`.

Read more info on [using the web portal](#), and [about making a connection with SSH](#).

Tier-2 compute clusters

CPU clusters

The HPC-UGent Tier-2 infrastructure currently included several standard CPU-only clusters, of different generations (listed from old to new).

For basic information on using these clusters, see our [documentation](#).

<i>cluster name</i>	<i># nodes</i>	<i>Processor architecture</i>	<i>Usable memory/ node</i>	<i>Local disk space/ node</i>	<i>Interconnect</i>
doduo (default cluster)	128	2x 48-core AMD EPYC 7552 (Rome @ 2.2 GHz)	250 GiB	180GB SSD	HDR-100 InfiniBand
gallade	16	2x 64-core AMD EPYC 7773X	940 GiB	1.5 TB NVME	HDR-100 InfiniBand

<i>cluster name</i>	<i># nodes</i>	<i>Processor architecture</i>	<i>Usable memory/ node</i>	<i>Local disk space/ node</i>	<i>Interconnect</i>
		(Milan-X @ 2.2 GHz)			
shinx	48	2x 96-core AMD EPYC 9654 (Genoa @ 2.4 GHz)	370 GiB	500GB NVME	NDR-200 InfiniBand

Interactive debug cluster

A special-purpose interactive debug cluster is available, where you should always be able to get a job running quickly, **without waiting in the queue**.

Intended usage is mainly for interactive work, either via an interactive job or using the [HPC-UGent web portal](#).

This cluster is heavily over-provisioned, so jobs may run slower if the cluster is used more heavily.

Strict limits are in place per user:

- max. 5 jobs in queue
- max. 3 jobs running
- max. of 8 cores and 27GB of memory in total for running jobs

For more information, see our [documentation](#).

<i>cluster name</i>	<i># nodes</i>	<i>Processor architecture</i>	<i>Usable memory/ node</i>	<i>Local disk space/ node</i>	<i>Interconnect</i>
donphan (*)	16	2 x 18-core Intel Xeon Gold 6240 (Cascade Lake @ 2.6 GHz) + 1x shared	738 GiB	1.6 TB NVME	HDR-100 Infiniband

<i>cluster name</i>	<i># nodes</i>	<i>Processor architecture</i>	<i>Usable memory/node</i>	<i>Local diskspace/node</i>	<i>Interconnect</i>
		NVIDIA Ampere A2 GPU (16GB GPU memory)			

(*) also see this [extra information](#)

GPU clusters

GPU clusters are available in the HPC-UGent Tier-2 infrastructure, with different generations of NVIDIA GPUs.

These are well suited for specific workloads, with software that can leverage the GPU resources (like TensorFlow, PyTorch, GROMACS, AlphaFold, etc.).

For more information on using these clusters, see our documentation.

<i>cluster name</i>	<i># nodes</i>	<i>Processor architecture & GPUs</i>	<i>Usable memory/node</i>	<i>Local diskspace/node</i>	<i>Interconnect</i>
joltik	10	2x 16-core Intel Xeon Gold 6242 (Cascade Lake @ 2.8 GHz) + 4x NVIDIA Volta V100 GPUs (32GB GPU memory)	256 GiB	800GB SSD	double EDR Infiniband
accelgor	9	2x 24-core AMD EPYC 7413 (Milan @ 2.2 GHz) + 4x NVIDIA	500 GiB	180GB SSD	HDR InfiniBand

<i>cluster name</i>	<i># nodes</i>	<i>Processor architecture & GPUs</i>	<i>Usable memory/ node</i>	<i>Local diskspace/ node</i>	<i>Interconnect</i>
		Ampere A100 GPUs (80GB GPU memory)			
litleo	8	1x 48 core AMD EPYC 9454P (Genoa @ 2.75 GHz) + 2x NVIDIA H100 NVL (96GB GPU memory)	315 GiB	1.4TB SSD	NDR-200 Infiniband

Tier-2 shared storage

<i>Filesystem name</i>	<i>Intended usage</i>	<i>Total storage space</i>	<i>Personal storage space</i>	<i>VO storage space (^)</i>
\$VSC_HOME	Home directory, entry point to the system	90 TB	3GB (fixed)	(none)
\$VSC_DATA	Long-term storage of large data files	1.9 PB	25GB (fixed)	250GB
\$VSC_SCRATCH	Temporary fast storage of 'live' data for calculations	1.7 PB	25GB (fixed)	250GB
\$VSC_SCRATCH_ARCANINE	Temporary very fast storage of 'live'	70 TB NVME	(none)	upon request

<i>Filesystem name</i>	<i>Intended usage</i>	<i>Total storage space</i>	<i>Personal storage space</i>	<i>VO storage space (^)</i>
	data for calculations (recommended for very I/O-intensive jobs)			

^ Storage space for a group of users (Virtual Organisation or VO for short) can be increased significantly on request. For more information, see our [documentation](#).

Infrastructure status

[Check the system status](#)