

# Compute node partitions

The compute nodes of Lise in Berlin (blogin.hlrn.de) and Emmy in Göttingen (glogin.hlrn.de) are organized via the following SLURM partitions:

## Lise (Berlin)

Partition (number holds cores per node)	Node name	Max. walltime	Nodes	Max. nodes per job	Max. jobs (running/queued) per user	Usable memory MB per node	CPU	Shared	Charged core-hours per node	Remark
standard96	bcn#	12:00:00	1204	512	16 / 500	362 000	Cascade 9242		96	default partition
standard96:test	bcn#	1:00:00	32 dedicated +128 on demand	16	1 / 500	362 000	Cascade 9242		96	test nodes with higher priority but lower walltime
large96	bfm#	12:00:00	28	8	16 / 500	747 000	Cascade 9242		144	fat memory nodes
large96:test	bfm#	1:00:00	2 dedicated +2 on demand	2	1 / 500	747 000	Cascade 9242		144	fat memory test nodes with higher priority but lower walltime
large96:shared	bfm#	48:00:00	2 dedicated	1	16 / 500	747 000	Cascade 9242		144	fat memory nodes for data pre- and postprocessing
huge96	bsn#	24:00:00	2	1	16 / 500	1522 000	Cascade 9242		192	very fat memory nodes for data pre- and postprocessing

12 hours are too short? See [here](#) how to pass the 12h walltime limit with job dependencies.

## Emmy (Göttingen)

Partition (number holds cores per node)	Node name	Max. walltime	Nodes	Max. nodes per job	Max. jobs per user	Usable memory MB per node	CPU, GPU type	Shared	NPL per node hour	Remark
standard96	gcn#	12:00:00	924	256	unlimited	362 000	Cascade 9242		96	default partition
standard96:test	gcn#	1:00:00	16 dedicated +48 on demand	16	unlimited	362 000	Cascade 9242		96	test nodes with higher priority but lower walltime
large96	gfn#	12:00:00	12	2	unlimited	747 000	Cascade 9242		144	fat memory nodes
large96:test	gfn#	1:00:00	2 dedicated +2 on demand	2	unlimited	747 000	Cascade 9242		144	fat memory test nodes with higher priority but lower walltime
large96:shared	gfn#	48:00:00	2 dedicated +2 on demand	1	unlimited	747 000	Cascade 9242		144	fat memory nodes for data pre- and postprocessing
huge96	gsn#	24:00:00	2	1	unlimited	1522 000	Cascade 9242		192	very fat memory nodes for data pre- and postprocessing
medium40	gcn#	48:00:00	368	128	unlimited	181 000	Skylake 6148		40	
medium40:test	gcn#	1:00:00	32 dedicated +96 on demand	8	unlimited	181 000	Skylake 6148		40	test nodes with higher priority but lower walltime
large40	gfn#	48:00:00	11	4	unlimited	764 000	Skylake 6148		80	fat memory nodes
large40:test	gfn#	1:00:00	3	2	unlimited	764 000	Skylake 6148		80	fat memory test nodes with higher priority but lower walltime
large40:shared	gfn#	48:00:00	2	1	unlimited	764 000	Skylake 6148		80	fat memory nodes for data pre- and postprocessing
gpu	ggpu#	12:00:00	3	3	unlimited	764MB per node (32GB per gpu)	Skylake 6148 + 4 Tesla V100		600	
gpu40					unlimited	(40GB per gpu)	?????? + 4 NVidia A100		600	see <a href="#">GPU Usage</a>

Which partition to choose?

If you do not request a partition, your job will be placed in the default partition, which is **standard96**.

The default partitions are suitable for most calculations. The `:test` partitions are, as the name suggests, intended for shorter and smaller test runs. These have a higher priority and a few dedicated nodes, but are limited in time and number of nodes. Shared nodes are suitable for pre- and postprocessing. A job running on a shared node is only accounted for its core fraction (cores of job / all cores per node). All non-shared nodes are exclusive to one job, which implies that full **NPL** are paid.

Details about the CPU/GPU types can be found below.  
The network topology is described **here**.

The available home/local-ssd/work/perm storages are discussed in [File Systems](#).

An overview of all partitions and node statuses is provided by: `sinfo -r`  
To see detailed information about a nodes type: `scontrol show node <nodename>`

## List of CPUs and GPUs at HLRN

Short name	Link to manufacturer specifications	Where to find	Units per node	Cores per unit	Clock speed [GHz]
Cascade 9242	<a href="#">Intel Cascade Lake Platinum 9242 (CLX-AP)</a>	Lise and Emmy compute partitions	2	48	2.3
Cascade 4210	<a href="#">Intel Cascade Lake Silver 4210 (CLX)</a>	blogin[1-8], glogin[3-8]	2	10	2.2
Skylake 6148	<a href="#">Intel Skylake Gold 6148</a>	Emmy compute partitions	2	20	2.4
Skylake 4110	<a href="#">Intel Skylake Silver 4110</a>	glogin[1-2]	2	8	2.1
Tesla V100	<a href="#">NVIDIA Tesla V100 32GB</a>	Emmy gpu partition	4	640/5120*	

\*Tensor Cores / CUDA Cores