Hardware specifications of cluster compute nodes

This page provides an overview of the compute cluster's hardware specifications, including node details, partitions, and configuration settings. We strive to keep this information up to date, but changes may occur. For the most accurate and current cluster details, please use the following command: `clusterinfo`

The LUIS computing cluster is a heterogeneous general purpose system designed for a variety of workloads. All nodes in a sub-cluster (“partition”) are interconnected using Mellanox Infiniband (at least QDR) non-blocking fat tree network. We use SLURM as the job scheduler.

By policy, the compute nodes cannot access the internet outside the computing cluster. Exceptions need to belong to the LUH network. If you need such an exception, contact cluster support stating IP address, port number(s), protocol(s) and account name(s) that should be allowed to use the exception as well as a contact person, the reason and duration of the exception. However, the compute nodes have access to cloud storage systems provides by LUIS. For detailed information please refer to the Rclone usage instructions.

You will notice that the columns “(useable) Memory/Node (MB)” and “Memory Total (GB)” differ slightly, which takes into account the difference of total physical memory per node vs. the memory configured in the batch scheduler SLURM avilable to jobs. The latter number is smaller since the operating system needs memory, too. If you want to autoritatively find out the maximum allocateable memory per node in SLURM, use the `clusterinfo -n` command on a login node.

Nodes running in the “FCH” service ("Forschungscluster-Housing", nodes owned by institutes that are integrated into the cluster) are too varied to be listed in these tables. They contribute significant additional power to the cluster, mostly during the night and over the weekend, but are usually reserved exclusively for institute accounts on week days. Your jobs have a chance of running in the night when they request less than 12 hours of walltime, or during weekends, for jobs that request less than 60 hours. You can find out more about nodes in this part of the cluster using the `clusterinfo` command on a login node.

Parallel Clusters (MPP)

Partition	Nodes	CPUs	Cores /Node	Cores Total	(useable) Memory /Node (MB)	Memory Total (GB)	Gflops /Core [1]	Local Disk /Node (GB)	Node Interconnect
mpp.share	27	2x AMD EPYC 9534	128	3456	500.000	13.500	80	800 (NVMe)	InfiniBand NDR, 200 Gbs
mpp.single	10	2x AMD EPYC 9534	128	1280	500.000	5.000	80	800 (NVMe)	InfiniBand NDR, 200 Gbs
amo	80	2x Intel Cascade Lake Xeon Gold 6230N	40	3200	180.000	15360	75	400 (SSD)	InfiniBand HDR, 100 Gbs
taurus	24	2x Intel Skylake Xeon Gold 6130	32	768	120.000	3072	67	500 (HDD)	InfiniBand EDR, 100 Gbs
haku	20	2x Intel Broadwell Xeon E5-2620 v4	16	320	60.000	1280	34	80 (SSD)	InfiniBand FDR, 40 Gbs
lena	80	2x Intel Haswell Xeon E5-2630 v3	16	1280	60.000	5120	38	180 (SSD)	InfiniBand QDR, 40 Gbs

Large Memory Servers (SMP)

Partition	Nodes	CPUs	Cores /Node	Cores Total	(useable) Memory /Node (MB)	Memory Total (GB)	Gflops /Core [2]	Local Disk /Node (GB)
smp	9	2x AMD EPYC 9534	128	1152	1.024.000	9216	40	800 (NVMe)
	2	2x AMD EPYC 9354	64	128	1.020.000	2048	52	3600 (NVMe)

GPU Servers

Partition	Nodes	CPUs	GPUs	Cores /Node	Cores Total	(useable) Memory /Node (MB)	Memory Total (GB)	Local Disk /Node (GB)
gpu	4	2x AMD EPYC 9555	4x NVIDIA H200 141 GB	128	512	1.150.000	4490	5900 (NVMe)
	4	2x Intel Xeon Gold 6230	2x NVIDIA Tesla V100 16 GB	160	40960	125.000	512	300 (SSD)
	3	2x Intel Xeon Gold 6342	2x NVIDIA A100 80 GB	48	288	1.025.000	3072	3500 (NVMe)

[1] [2]
,

Performance values are theoretical

From:
https://docs.cluster.uni-hannover.de/ - **Cluster Docs**

Permanent link:
https://docs.cluster.uni-hannover.de/doku.php/resources/computing_hardware

Last update: **2026/02/16 16:31**