

The Hessian Competence Center for High Performance Computing **HPC Hessen**

Brainware for Science

The Hessian Competence Center for High Performance Computing (HPC Hessen) supports scientists striving for the efficient and sustainable use of modern HPC systems. HPC systems are in use at all universities in Hesse and accessible statewide.

To this end, HPC Hessen offers work-shops, consulting, and HPC user support. HPC experts support users locally at each university, but also coordinate their efforts.

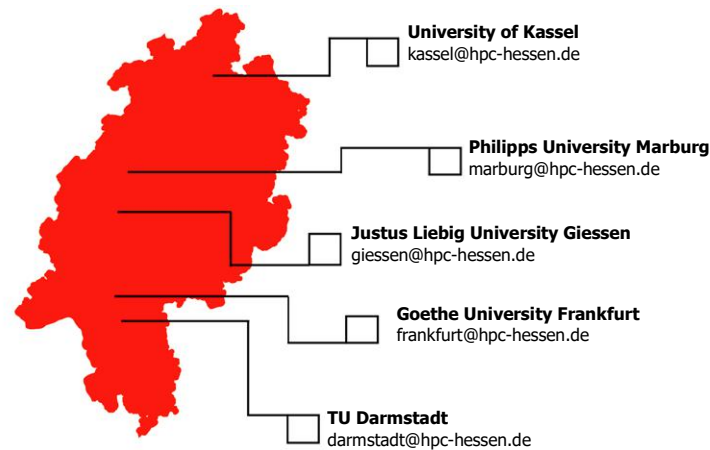
Our services includes the installation of a central monitoring infrastructure, cooperative support in using performance analysis tools, preparation and execution of work-shops, and monthly internal workshops to discuss and learn from new local experiences.

HPC Hessen was founded in 2014 by the universities of Darmstadt, Frankfurt, Gießen, Kassel, and Marburg with initial funding by the Hessian State Ministry of Higher Education, Research and the Arts.

www.hpc-hessen.de



Local Support:



Office:
Hessisches Kompetenzzentrum für
Hochleistungsrechnen
Alexanderstraße 2
D-64283 Darmstadt

E-mail: office@hpc-hessen.de
Phone: +49 6151 16-76161
URL: www.hpc-hessen.de



Which computer is right for you?

To find out which computer matches your research best, an overview of the main characteristics of the Hessian high performance computers is provided here.

HPC Hessen experts can advise you on the suitability of these systems for your needs and how to access them.

The Hessian Competence Center for High Performance Computing HPC Hessen (www.hpc-hessen.de)



Lichtenberg Cluster TU Darmstadt	Goethe-HLR Goethe University Frankfurt	Skylla Cluster JLU Giessen	Linux Cluster University of Kassel	MaRC2 Philipps University Marburg
-------------------------------------	-------------------------------------------	-------------------------------	---------------------------------------	-----------------------------------------

Typical Node Parameters	Note type A		Note type B		Node type A	Node type B	GPU Node			AMD Node		Intel Node			
	Typical	Max	Typical	Max	-	-	-	Typical	Max	Typical	Max	Typical	Max	Typical	Max
Cores (sockets x cores/socket)	2x8	8x8	2x12	4x15	2x40	2x40	2x6	2x6	4x8	2x16		2x12		4x16	
Memory	32 GB	1 TB	64 GB	1 TB	192 GB	768 GB	128 GB	32 GB	64 GB	64 GB	256 GB	256 GB		256 GB	
FLOPS/Core (DP, theor. peak)	20.8 GFlop/s		20.0 GFlop/s		76.8 GFlop/s	76.8 GFlop/s	27.8 GFlop/s	9.6 GFlop/s		18.38 GFlop/s		28.83 GFlop/s		9.2 GFlop/s	
CPU Type	Intel Xeon E5-2670		Intel Xeon E5-2680v3		Intel Xeon Skylake 6148 (2.4 GHz)	Intel Xeon Skylake 6148 (2.4 GHz)	Intel Xeon E5-2630v2 (2.6GHz)	AMD Opteron 2431 (Istanbul, 2.4 GHz)		AMD Opteron 6267 (2.3 GHz)		Intel Xeon E5-2650v4 (2.2 GHz)		AMD Opteron (Interlagos/AbuDhabi, 2.3 GHz)	
MPI Communication (pt2pt)	Intranode	Internode	Intranode	Internode	Intranode	Intranode	-	Intranode	Internode	Intranode	Internode	Intranode	Internode	Intranode	Internode
- Bandwidth	4.6 GB/s	4.7 GB/s	4.1 GB/s	6.5 GB/s	6.7 GB/s	6.7 GB/s	-	2.5 GB/s	1.6 GB/s	2.7 GB/s	4.2 GB/s	2.7 GB/s	4.2 GB/s	2.5 GB/s	1.9 GB/s
- Latency (64 bytes)	0.73 µs	1.34 µs	0.71 µs	1.29 µs	0.39 µs	0.39 µs	-	0.8 µs	3.8 µs	0.8 µs	2.5 µs	0.8 µs	2.5 µs	1.0 µs	3.8 µs
Memory Bandwidth	40 GB/s per socket		60 GB/s per socket		192 GB/s	192 GB/s	-	15.5 GB/s		51 GB/s		76 GB/s		71.5 GB/s	
Accelerators	2x NVIDIA K20Xm 1.3 TFlop/s, 6 GB each		2x NVIDIA K40m 1.4 TFlop/s, 12 GB each		-	-	2 x AMD FirePro S10000 1.5 Tflop/s (DP) each		-	-		-		-	
Local Temporary Storage	100 GB				-	-	-	400 GB		1.8 TB		2 TB		2 TB	
Node Allocation	shared and exclusive				-	-	-	shared and exclusive		shared and exclusive		shared		shared	

Global Cluster Parameters					
Processors (CPU, DP, peak)	771 TFlop/s		1,450 TFlop/s		9.3 TFlop/s
Accelerators (GPU, DP, peak)	180 TFlop/s		150 TFlop/s		-
Compute Cores (CPU)	27,928		19,360		992
Permanent Storage	0.5 PB		4 PB		8 TB
Scratch Storage	1.5 PB		2 PB		24 TB
Job Manager	Slurm Workload Manager		Slurm Workload Manager		Sun Grid Engine
Other Job Constrains	runtime: 24h, max 7d		runtime: max 21d		runtime: 5d/unlimited max 128 (192) cores/user
					runtime: 24h, max 48h
					runtime: 10d max 1000 jobs/user