



Treffen der HPC NutzervertreterInnen

16.12.2022, 10:00 – 11:00 Uhr



TOPs

- Status HPC Cluster
 - aktuell: bis einschließlich Montag keine Stromgeneratoren, im Notfall müsste HPC Cluster abgeschaltet werden
 - \$WORK ist immer noch ziemlich voll (>95%)
- nächster HPC Cluster
 - Übersicht zum neuen System
- Sonstiges
 - ZKI Arbeitskreis Supercomputing Herbsttagung in OL (u.a. ein wiss. Fachvortrag)
 - HPC-Support durch KI
 - nächster Termin 31.03.2023

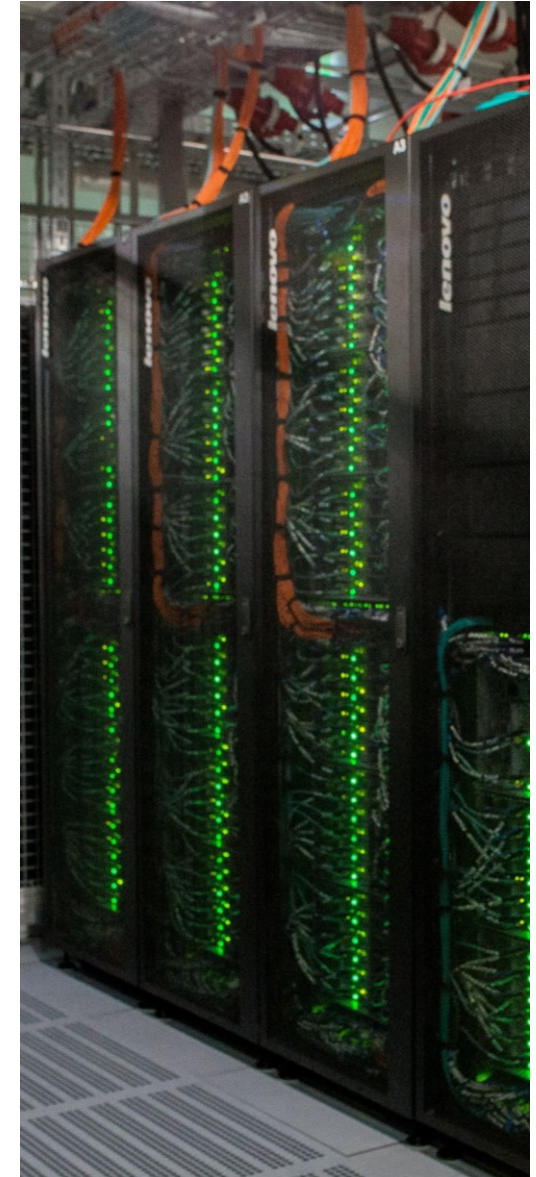
Overview HPC-Clusters

- CARL-2
 - multi-purpose compute (MPC) cluster as a basic computing resource
 - funded by the University/MWK and the DFG under grant number INST 184/225-1 FUGG (Forschungsgroßgerät nach Art. 91b GG)
 - total funding: 2.400.000€ (+460.000€ HIFMB)
 - responsible PIs: Thorsten Klüner, Jörg Lücke, Stefan Harfst
- RE-HPC and MOUSE (ForWind)
 - CFD cluster for wind energy research
 - funded by REACT-EU under grant number ZW7-95186744 and BMWK under grant number 03EE3067A
 - total funding: 1.667.000€ (REACT-EU, including 167.000€ Uni) and 478.956,21€ (BMWK, HPC-Cluster only, project total is 1.628.058,80€)
 - responsible PI: Laura Lukassen
- used as a shared HPC cluster
 - common infrastructure is shared (e.g. file systems, network)
 - shared administration

HPC-Cluster @ University Oldenburg

- shared HPC-cluster
 - close to 160 compute nodes
 - 4 login and 2 administration nodes
 - Infiniband HDR interconnect for parallel computing
 - 25/1GE network for management
 - parallel file system (GPFS) with >4 PiB capacity
 - NFS mounted central storage

 - Linux (RHEL) as OS
 - many scientific applications and libraries available
 - Job Scheduler (SLURM)



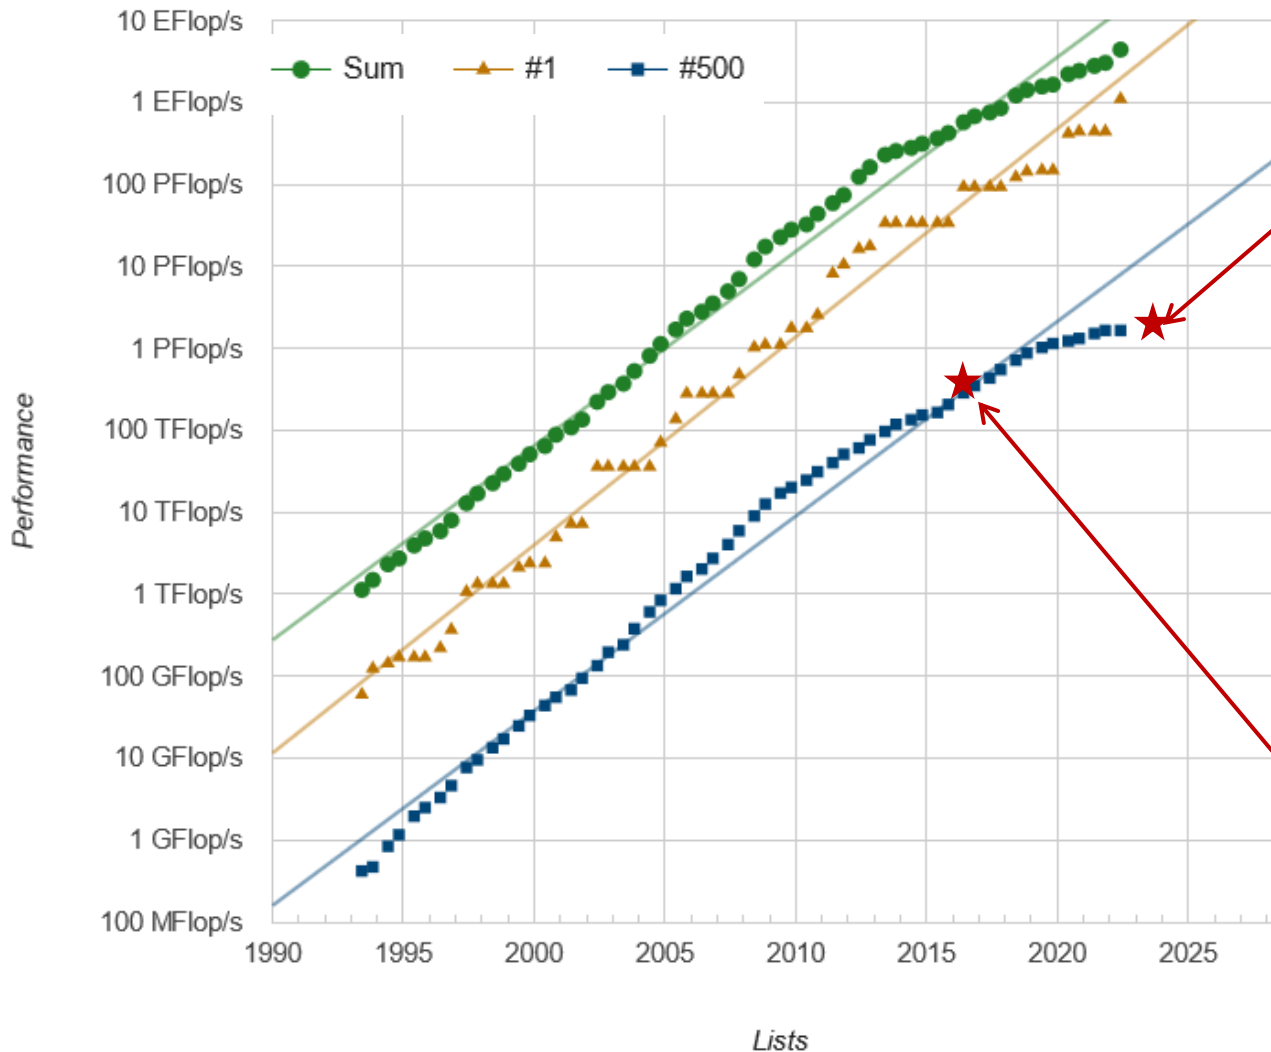
Summary MPC & CFD

Feature	MPC	CFD	Total
Nodes	91	70	159
CPU-Cores	11.648	8.960	20.608
RAM	91 TiB	54 TiB	145 TiB
GPUs	24x H100 94GB	8x A100 80GB 4x A100 94GB	8x A100 80GB 28x H100 94GB
GPFS	>2 PiB	>2 PiB	>4 PiB
Burst Buffer	92 TB	-	92 TB
Rmax (CPU)	~546 TFlop/s	~414 TFlop/s	960 TFlop/s
Rpeak (GPU)	~1.440 TFlop/s	~320 TFlop/s	1.760 TFlop/s

Node Configurations and Project Assignment

	Node Type	Node Cnt	CPU / GPU	RAM	Assignment
MPC	MPC-STD	42	2x AMD Genoa EPYC 9554, 64C, 3.1 GHz, 360W	768 GB	DFG
		14			HIFMB
	MPC-BIG	24		1124 GB	DFG
		2			AG EST
	MPC-PP	2		4096 GB	DFG
		1			HIFMB
MPC-GPU	6	CPU identical to MPC-STD, 4x Nvidia H100 94GB SXM	1024 GB	DFG	
CFD	CFD-STD	52	2x AMD Genoa EPYC 9554, 64C, 3.1 GHz, 360W	768 GB	RE-HPC
		16			MOUSE
	CFD-GPU	1	identical to MPC-GPU	1024 GB	MOUSE
		1	2x AMD Milan EPYC 7713, 64C, 2.0 GHz, 225W, 8x Nvidia A100 80GB SXM		RE-HPC

Top500 Performance Development



MPC/CFD

- not ranked
- Rpeak 2,720 TFlop/s (CPU+GPU)
- 161 nodes, 20,608 cores
- 145 TiB RAM



4.7x Rpeak
0.28x nodes
1.53x cores
1.48x RAM

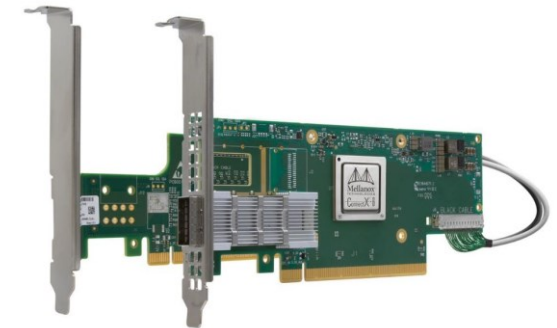
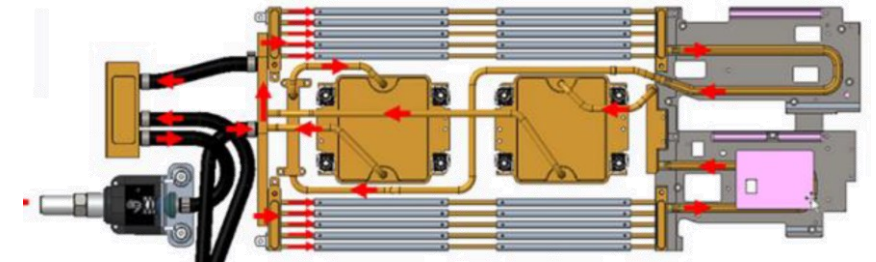
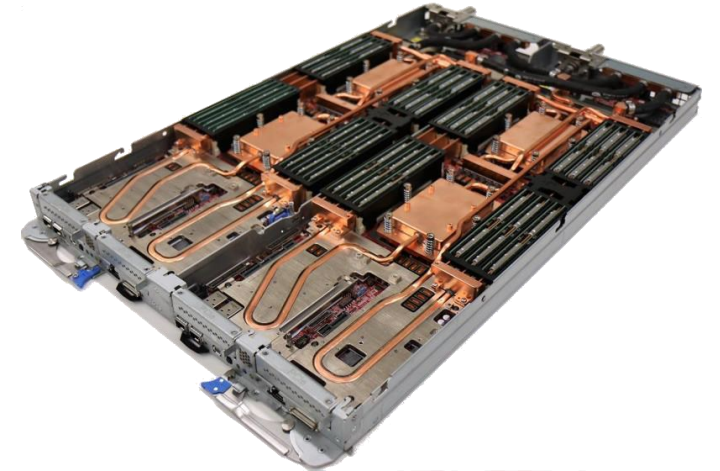


CARL/EDDY

- Rank 363
- Rmax (CPU) 457.2 TFlop/s
- Rpeak 577 Tflop/s (CPU+GPU)
- 571 nodes, 13,500 cores
- 98 TiB RAM

Details on MPC/CFD-STD Compute Nodes

- compute nodes equipped with
 - CPU: 2x AMD „Genoa“ EPYC 9554, 64C @ 3.1GHz, 360W TDP
 - RAM: 768 GB (24x 32GB DDR5 @ 4800MHz)
 - Infiniband: ConnectX-6 HDR/200GbE (SharedIO, see below)
 - direct water-cooling (DWC) for CPUs, RDIMMs, and more
- two nodes on a single tray
 - shared DWC
 - shared IB-connector



TOPs

- Status HPC Cluster
 - aktuell: bis einschließlich Montag keine Stromgeneratoren, im Notfall müsste HPC Cluster abgeschaltet werden
 - \$WORK ist immer noch ziemlich voll (>95%)
- nächster HPC Cluster
 - Übersicht zum neuen System
- Sonstiges
 - ZKI Arbeitskreis Supercomputing Herbsttagung in OL (u.a. ein wiss. Fachvortrag)
 - HPC-Support durch KI
 - nächster Termin 31.03.2023

Asking ChatGPT for HPC support

sh

How can I write a job script for a supercomputer using slurm if I want to run the program OpenFOAM on 256 core with 4GB per core for 12 hours?



Here is an example of a job script that you can use to run the program OpenFOAM on a supercomputer using the slurm workload manager. The script requests 256 cores with 4 GB of memory per core and runs the program for 12 hours:



Copy code

```
#!/bin/bash

#SBATCH --nodes=1
#SBATCH --ntasks=256
#SBATCH --mem-per-cpu=4096
#SBATCH --time=12:00:00

# Load the OpenFOAM module
module load OpenFOAM

# Run the OpenFOAM simulation
foamJob
```

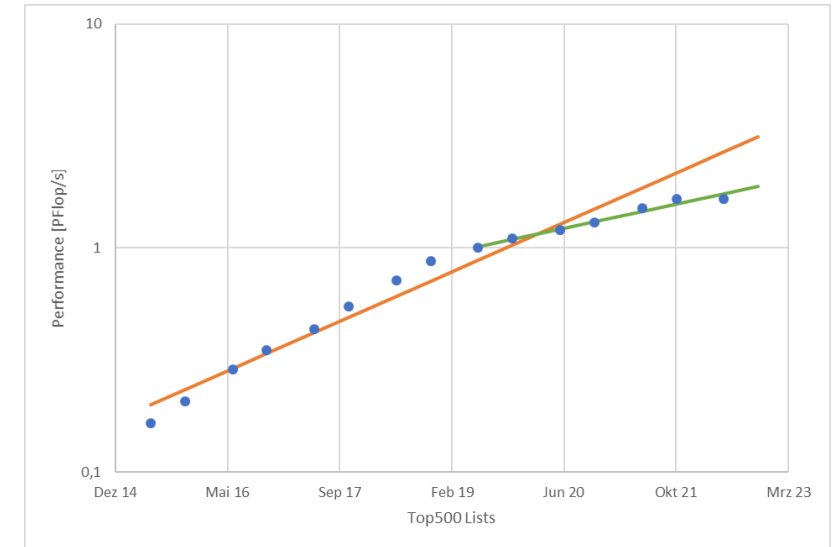
Try again



Merry Christmas &
a Happy New Year

How to reach the Top500

- based on the current development
 $R_{\max} \geq 2.15$ PFlop/s is needed for entry in Top500 in November 2023



- CPU Solution
 - a single compute node can reach up to 6 TFlop/s, the available 148 compute nodes can reach up to 888 TFlop/s
 - approximately 210 additional compute nodes are needed, estimated costs ~5 M€
- GPU Solution
 - a single H100 GPU can reach up to 60 TFlop/s(?), the available 28 H100 GPUs can reach 1.680 TFlop/s
 - approximately 2 additional GPU nodes with 4x H100 are needed, estimated costs 260 k€ (still no guarantees)