

EuroCC


Providing access to compute, courses, consulting

Speaker: Simeon Harrison
VSC Research Center and EuroCC Austria

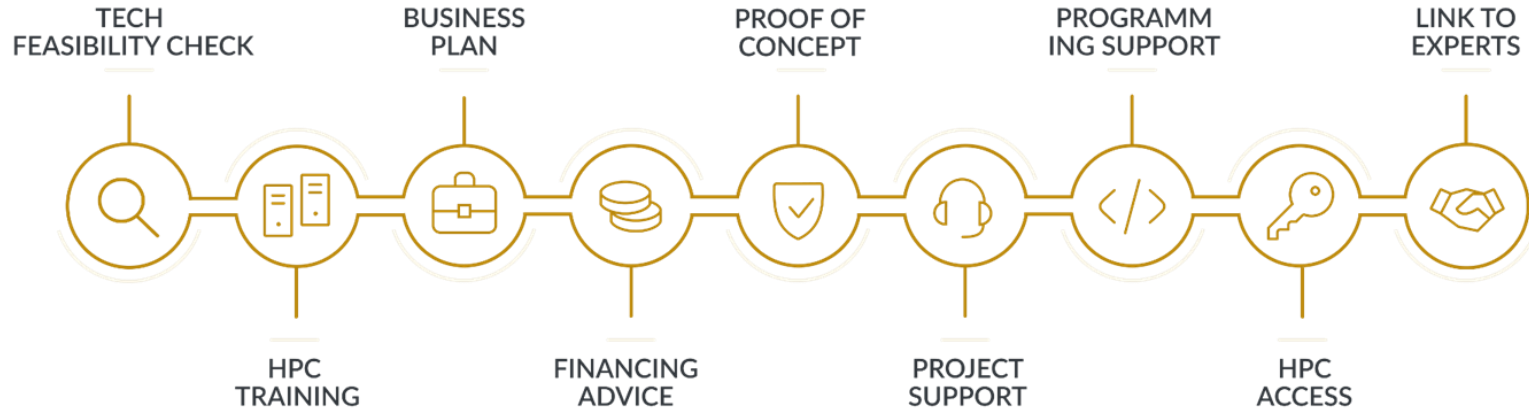
EuroCC

Fully funded EU project

- EuroCC is EU-funded international initiative aimed to support the uptake of AI and High-Performance Computing (HPC) in Europe
- Set up of 32 National Competence Centres (NCCs) across Europe
- EuroCC Austria is one of them
- Service Provider for AI, HPC and HPDA



EuroCC Austria's Services




INFRASTRUCTURE – TRAINING – CONSULTING

The Vienna Scientific Cluster

VSC-4 (2019)

790 CPU nodes

- 2x Intel Skylake Platinum CPUs
- 2x 24 cores per CPU
- 96 GB of memory per node



VSC-5 (2022)

770 CPU nodes

- 2x AMD EPYC Milan
- 2x 64 cores per CPU
- 512 GB of memory per node

60 GPU nodes 2x NVIDIA A100,

- 40 GB memory per GPU

40 GPU nodes 2x NVIDIA A40

- 48 GB memory per GPU

Need More Compute Power?

LUMI

- #5 in Top500
- Linpack: 380 PFlop/s
- AMD EPYC CPUs
- AMD Instinct MI250X GPUs (128 GB)

<https://www.lumi-supercomputer.eu/>

Leonardo

- #7 in Top500
- Linpack: 240 PFlos/p
- Intel Xeon CPUs
- NVIDIA A100 GPUs (64GB)

<https://leonardo-supercomputer.cineca.eu/>

Supercomputers in Europe

EuroHPC JU systems

Apply for access to EuroHPC supercomputers

Different access modes:

- Benchmark Access
- Development Access
- ...
- Extreme Scale Access

<https://eurohpc-ju.europa.eu/>




Supercomputers in Europe

- VSC5 (60 nodes with 2 Nvidia A100 40 GB)
- LEONARDO (3456 nodes with 4 Nvidia A100 64 GB)
- LUMI-G (2978 nodes with 4 AMD MI250x 128 GB)
- MUSICA (~200 nodes with 4 Nvidia H100 96 GB)

https://eurohpc-ju.europa.eu/supercomputers/our-supercomputers_en

Typical Setup of a Supercomputer



Access via:

- SSH (Terminal, IDE)
- JupyterHub
- NoMachine (GUI)

SLURM: Job Scheduler

VSC5 job script (2 nodes)

```
#!/bin/bash
#SBATCH --partition=zen3_0512_a100x2
#SBATCH --qos=zen3_0512_a100x2
#SBATCH --nodes=2                # Number of nodes
#SBATCH --ntasks-per-node=1     # Number of `srun` tasks executed per node
#SBATCH --gres=gpu:2            # Number of GPUS (1 or 2 on VSC5)

#SBATCH --time=3-00:00:00       # Time limit. Format: Days-hours:minutes:seconds

module purge                    # Start in a clean environment
module load miniconda3          # Load conda
eval "$(conda shell.bash hook)" # Initialize conda

srun bash -c "conda run -n conda_env_name --no-capture-output python script.py"
```

Course Offer

Series: Foundations of LLM Mastery

- 22 January 2025 [Fine-tuning on one GPU](#)
- 26 February 2025 [Fine-tuning on multi GPUs](#)
- 26 March 2025 [Prompt Engineering Essentials](#)
- 30 April 2025 [Retrieval Augmented Generation](#)
- And many other topics: <https://training.vsc.ac.at>



Text generation

STAY IN TOUCH



EuroCC Austria



@eurocc_austria



eurocc-austria.at

THANK YOU



This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 101101903. The JU receives support from the Digital Europe Programme and Germany, Bulgaria, Austria, Croatia, Cyprus, Czech Republic, Denmark, Estonia, Finland, Greece, Hungary, Ireland, Italy, Lithuania, Latvia, Poland, Portugal, Romania, Slovenia, Spain, Sweden, France, Netherlands, Belgium, Luxembourg, Slovakia, Norway, Türkiye, Republic of North Macedonia, Iceland, Montenegro, Serbia