

Team RACKlette

Swiss Federal Institute of Technology Zürich
Swiss National Supercomputing Centre



About Us



Our Institution. ETH Zürich, is one of the world's leading universities, renowned for its research and innovation in science and technology. Established in 1855, it has produced numerous Nobel laureates and remains a driving force behind global advancements in engineering, natural sciences, and computer science. Attracting students from around the world, it is currently home to over 20,000 students.

The whole Team. RACKlette (Yes, like the Cheese!) is a student team from ETH Zürich passionate about high-performance computing (HPC). Comprising 18 active members from diverse fields like Mathematics, Computer Science, Physics, and Computational Science & Engineering (CSE), the team collaborates on optimizing computations from hardware to algorithmic performance.



Swiss National Supercomputing Center. Our work is closely integrated with CSCS, which provides the infrastructure, ongoing maintenance for our clusters, and valuable workshops. The image depicts one of their clusters, ALPS, currently ranked 6th on the Top500 list.



Diversity of the Competing Team

Diversity of Genders:

- We have a female team captain and actively create a welcoming environment for both women and men throughout the recruitment process and our Team.

Diversity of Major/Skills:

- Our competing Team includes students from computer science, mathematics and computational science & engineering.
- CS students specialize in hardware and software optimization, while CSE students bring deep scientific insights and intuition, uncovering hidden optimization approaches.

Diversity of Origin:

- Our members come from a range of backgrounds, including German, Romanian, Chinese, Peruvian, and Kurdish origins.
- Many belong to groups underrepresented at ETH Zürich, in Switzerland, and across Europe.

Diversity of Experience:

- Half of our team has already competed in major competitions, namely ISC24, while the other half is new to SCCs.
- We focus on recruiting first- and second-year students based on their curiosity about HPC and motivation, rather than prior knowledge.

Competing Team



Hannah Osswald (Team Captain)
Major: Computer Science
Interests: Computer Systems, Tuning, Biology
Working on: Reproducibility, HPL



Leonard Schmidt
Major: Computer Science
Interests: ML, Numerics, Parallel Programming
Working on: ICON, MLPerf



Renas Sahin
Major: Computer Science
Interests: Algorithms, Mathematics
Working on: NAMD



Yi Zhu
Major: Computer Science
Interests: Automation, Profiling, GPU Programming
Working on: NAMD, ICON



Nick Schaller
Major: Computational Science & Engineering
Interests: Computer Architecture, Biology
Working on: Reproducibility



Bogdan Murgu
Major: Computer Science
Interests: Security, Computer Networking, Linux
Working on: ICON



Hussein Harake (Main Advisor)
Senior Systems Engineer at CSCS



Prof. Thorsten Hoefler (Advisor)
Chief ML-Architect at CSCS
Professor for Computer Science at ETHZ

Hardware & Software Configuration

Hardware

Our cluster, named *Sbrinz* after a traditional Swiss cheese variety, is designed for high performance and energy efficiency with its 4 nodes. Its architecture consists of the following components:

Nvidia GH200

- 8 x Nvidia Grace and 8 x H100 NVL (two per node) – Grace CPUs with 72 cores per CPU having excellent power efficiency due to its ARM architecture.
- Up to 600GB/s highspeed NVLink-Bridge between each pair of GPUs within a node – enables more efficient GPU utilization.
- 480GB of LPDDR5X memory dedicated to each Grace CPU, 94GB of HBM3 for each H100 NVL – ensuring ample memory bandwidth for both CPUs and GPUs.

Storage:

- 6TB NFS shared storage of NVMe drivers over RDMA – to handle data-intensive workloads.

InfiniBand:

- 400Gbit/s NDR InfiniBand Switch – ultra-fast, low-latency communication between nodes.

Ethernet Switch:

- 10Gbit/s Ethernet Switch – general network connectivity.

Software

OS:

- Rocky Linux – widely used in HPC, ensures compatibility with broad range of HPC applications.

Cluster Management:

- Bright Cluster Manager (setup: one master and compute node, four compute nodes) – streamlines cluster administration and workload, very little overhead, significantly improves efficiency in managing resources.

Package Manager:

- Spack – handles software dependencies and environment, allows us to flexibly manage compilers and libraries

Compilers:

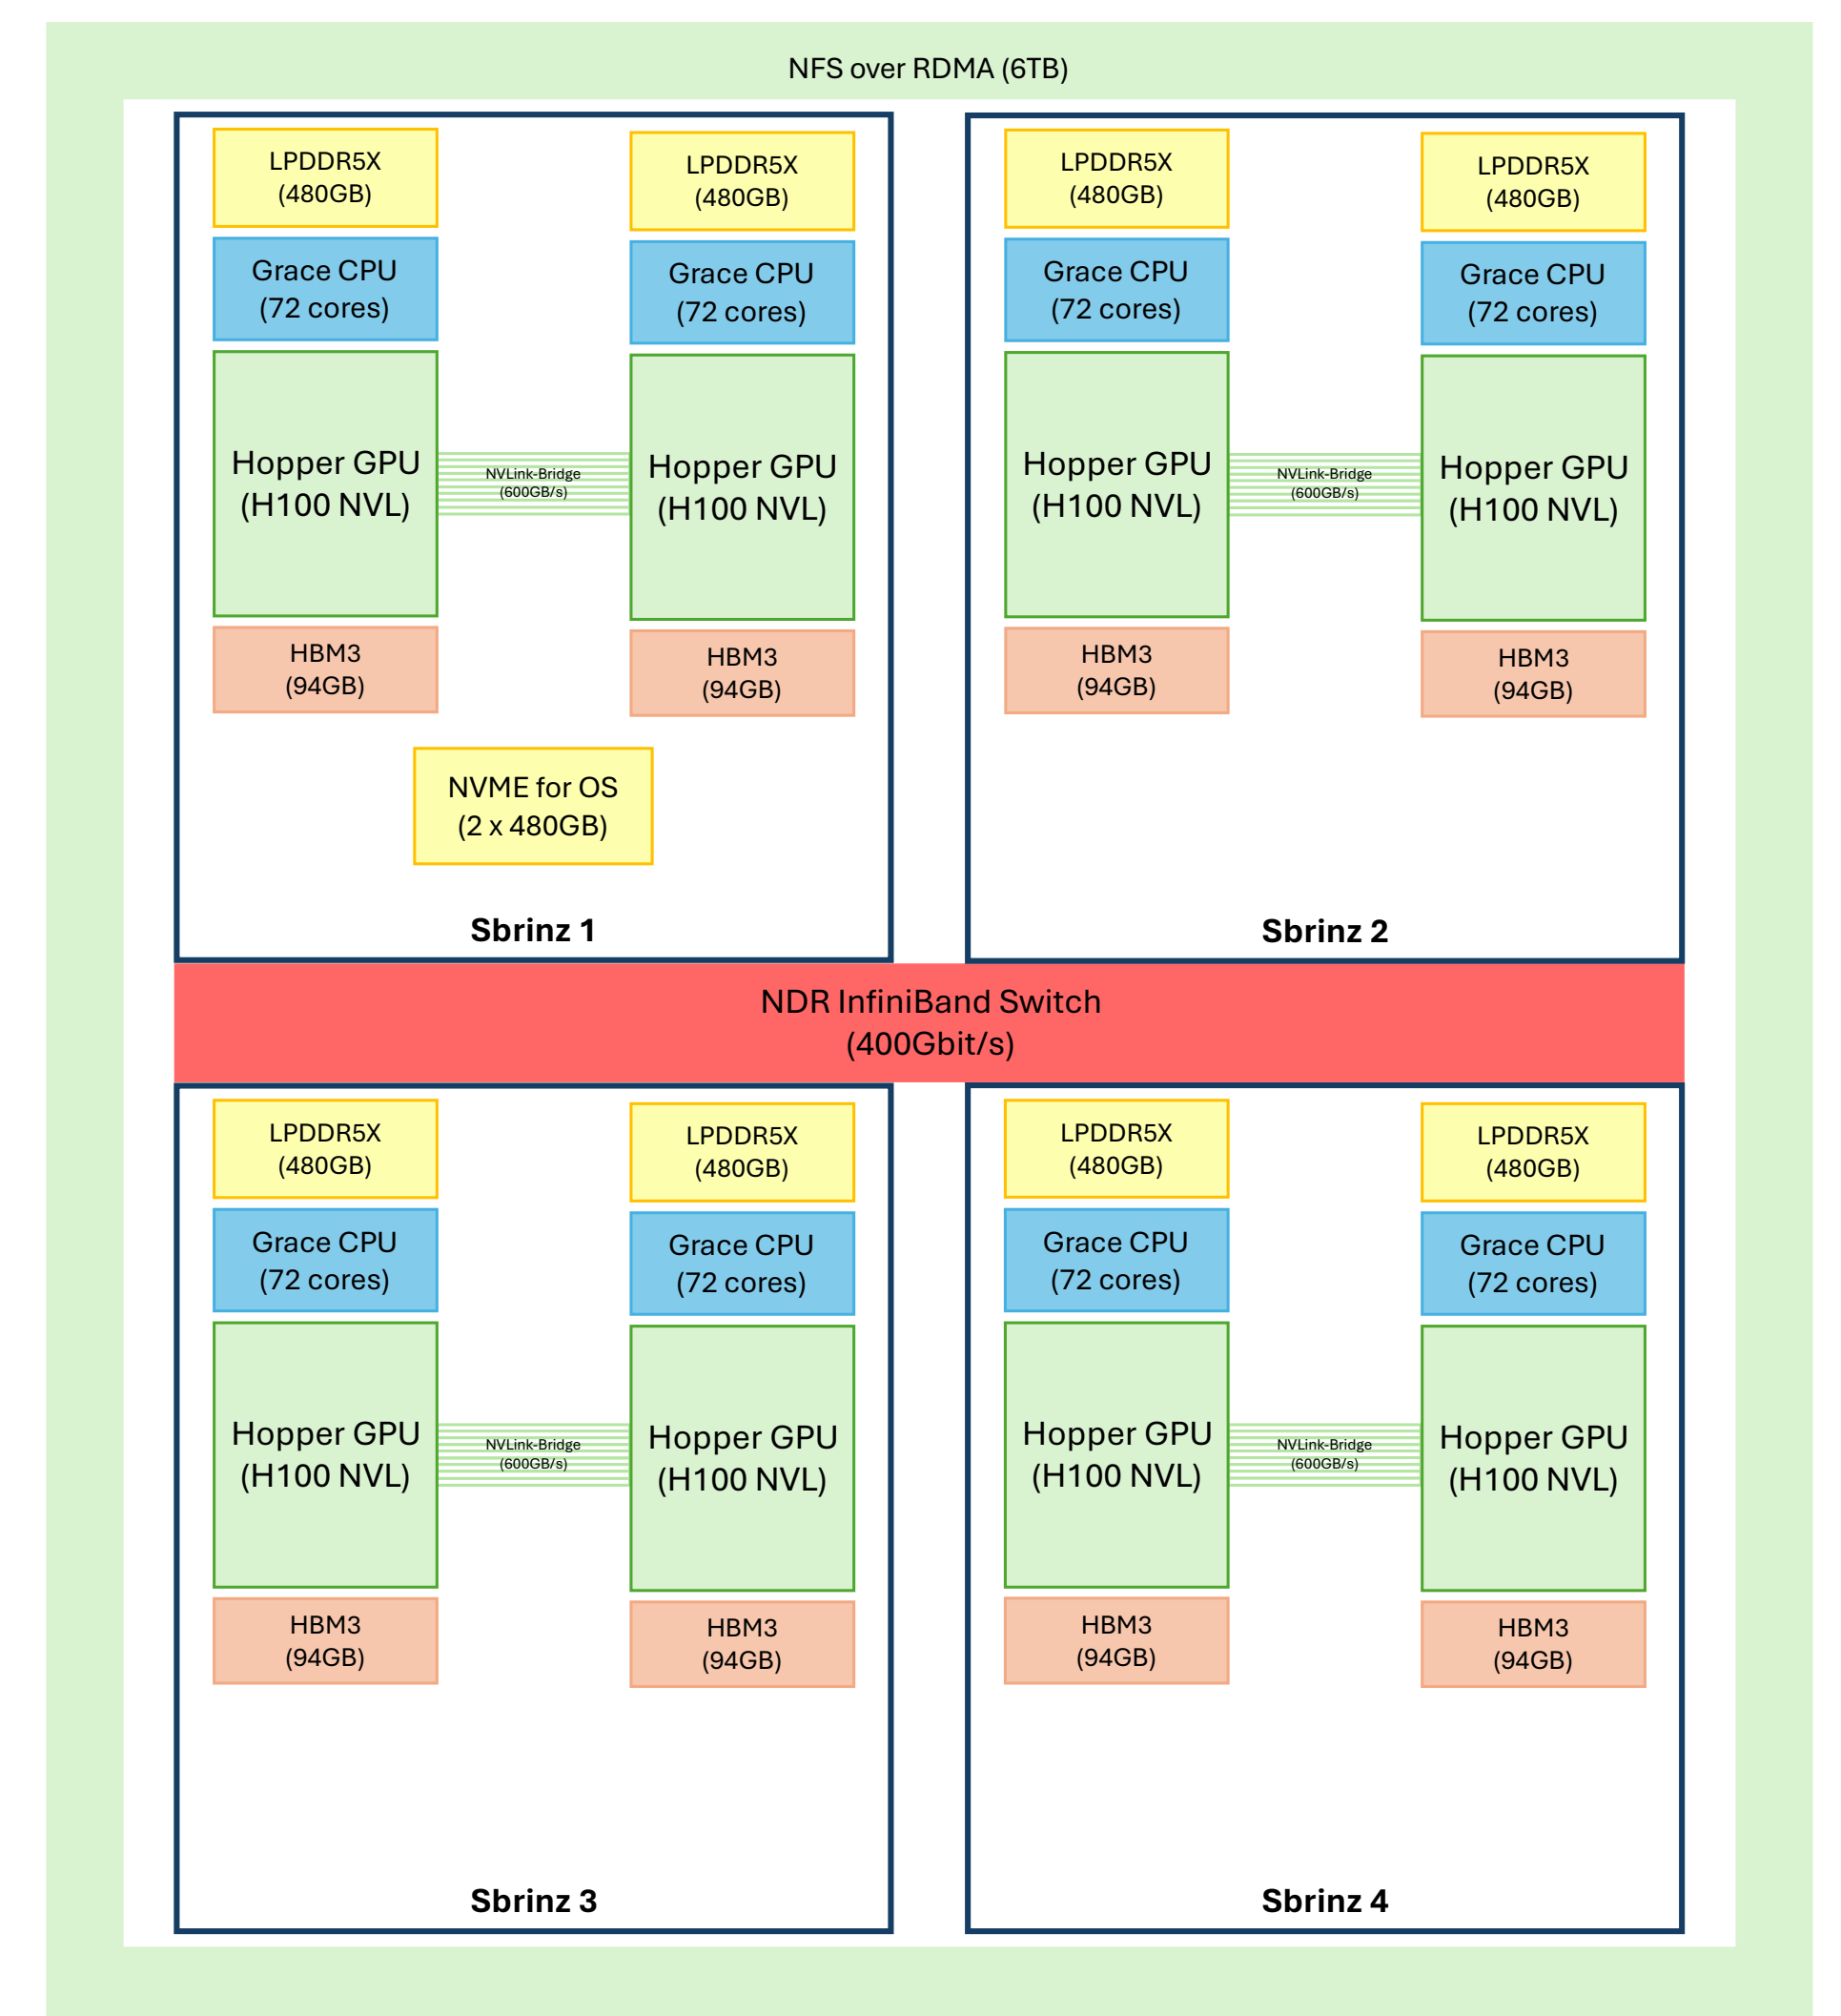
- Among others: GNU, NVHPC, ARM
- Experiment with many to compare performance on our architecture. Select the most suitable for each application.



NVIDIA H100 NVL



Sbrinz at ISC24



Sbrinz Hardware Topology

Preparation & Optimization

General Preparation

- Weekly self-organized meetings to discuss progress and address questions or issues
- Regular presentation from team members who have already attended competitions
- Professionally organized workshops by CSCS experts, focused on competition-specific applications
- Access to profiling & optimization tools developed by CSCS
- Teams of two, pairing new and experienced members, to foster diverse perspectives and ensure a valuable learning experience
- Computational Science and Engineering and Physics students provide scientific context to CS students

Staying under the Power and Noise Limit


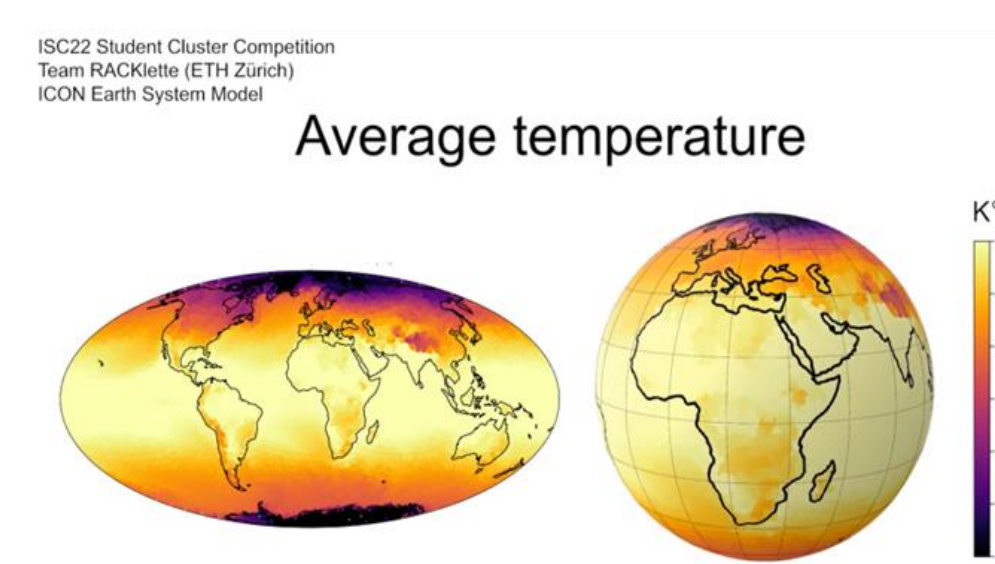
- Visualizing and monitoring power consumption to detect spikes early
- Enforce strict power limit (0.85 kW per Node)
- Sound-proof server cabinet for keeping fan noise down, reduces noise up to 40%

During the Competition

- Good preparation and understanding of every application – key to working efficiently
- Evaluate the difficulty of the tasks – assign the amount of time per task proportionally
- Working in teams of two – avoid getting stuck or making careless mistakes
- Taking turns working and resting – keeping energy and morale high during the 48h of competition

ICON

- Evaluate different compilers and libraries to identify which combination performs best
- Ensure dependency libraries are built with essential features for efficient I/O operations
- Utilize AGI profiling tool (developed by CSCS) to find performance bottlenecks
- Investigate methods to improve filesystem I/O efficiency, particularly for handling large datasets in HDF5 and NetCDF formats
- Explore optimizations for the I/O worker processes that manage data collection, distribution, and storage.




Ribosome in Violet – Rezvan Shahoei, 2014

NAMD

- Experiment with different Math libraries for fftw interface
- Try out different compilers and MPI versions
- Explore different communication setups of Charm++

Reproducibility Challenge

- Experiment with using Dataflow- Lifecycle-Graphs (DLF-Gs) and the methods described in the paper on different applications to gain a deeper understanding of the paper



Example Dataflow Lifecycle DAG presented in the paper

High-Performance Linpack (HPL)

- Use the latest pre-optimized NVIDIA container
- Choose problem size that fully utilizes available GPU memory
- Determine optimal blocking size so that matrix blocks fit within the cache, improving data locality
- Fine-tune these parameters for additional performance gains
- Leverage SLURMs built-in performance flags
- Optimize NUMA locality by applying CPU and GPU masking to minimize data transfer latencies
- Manually clear caches before running HPL to avoid cache trashing overhead

MLPerf Inference

- Explore different backends to compare the performance

Mystery Application and Mystery Benchmark

- Prepare different compilers and MPI implementations to quickly adapt to the given application



Team RACKlette benchmarking during SC23

Our Sponsors

