



Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

Sustainable HPC and AI

Prof. Dr. Dieter Kranzlmüller
Leibniz Supercomputing Centre (LRZ) &
Ludwig-Maximilians-Universität (LMU) Munich

Leibniz Supercomputing Centre (LRZ)
of the Bavarian Academy of Sciences and Humanities

Partner for Digital Transformation of Science & Research

BAdW

lrz

IT-Services for Science since 1962

Munich Universities, Bavaria, Germany, Europe, ...



High speed networking



Dienstleistungs- und
Gebührenkatalog
2025



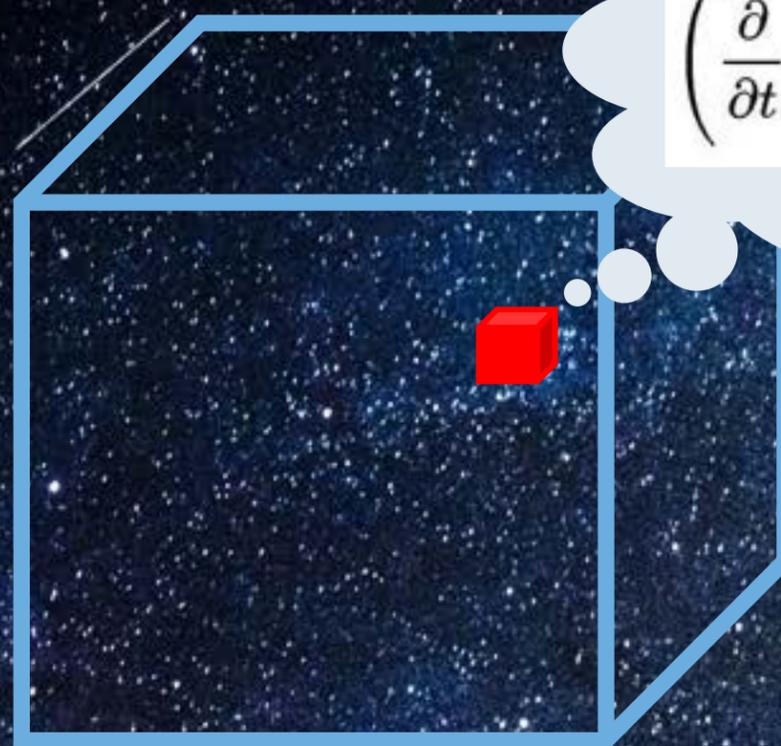
High Performance Computing



Visualization



Research & Development



$$\left(\frac{\partial}{\partial t} + \vec{v} \cdot \nabla_{\vec{x}} + \frac{\vec{F}}{m} \cdot \nabla_{\vec{v}} \right) f(\vec{x}, \vec{v}, t) = \frac{\partial f}{\partial t} \Big|_{\text{Stoss}}$$

<https://de.wikipedia.org/wiki/Boltzmann-Gleichung>

The background of the slide is a high-resolution simulation of a turbulent medium, likely a star-forming region. It features complex, swirling patterns of gas and dust, rendered in a color palette of deep blues, cyan, and teal. The structures are highly irregular and interconnected, showing the characteristic chaotic nature of turbulence.

The World's Largest Turbulence Simulation

Astro physics: How stars are born

Bavarian-Canadian Project ClimEx



Coupling HPC models for climate and hydrology, workflow and data management.¹

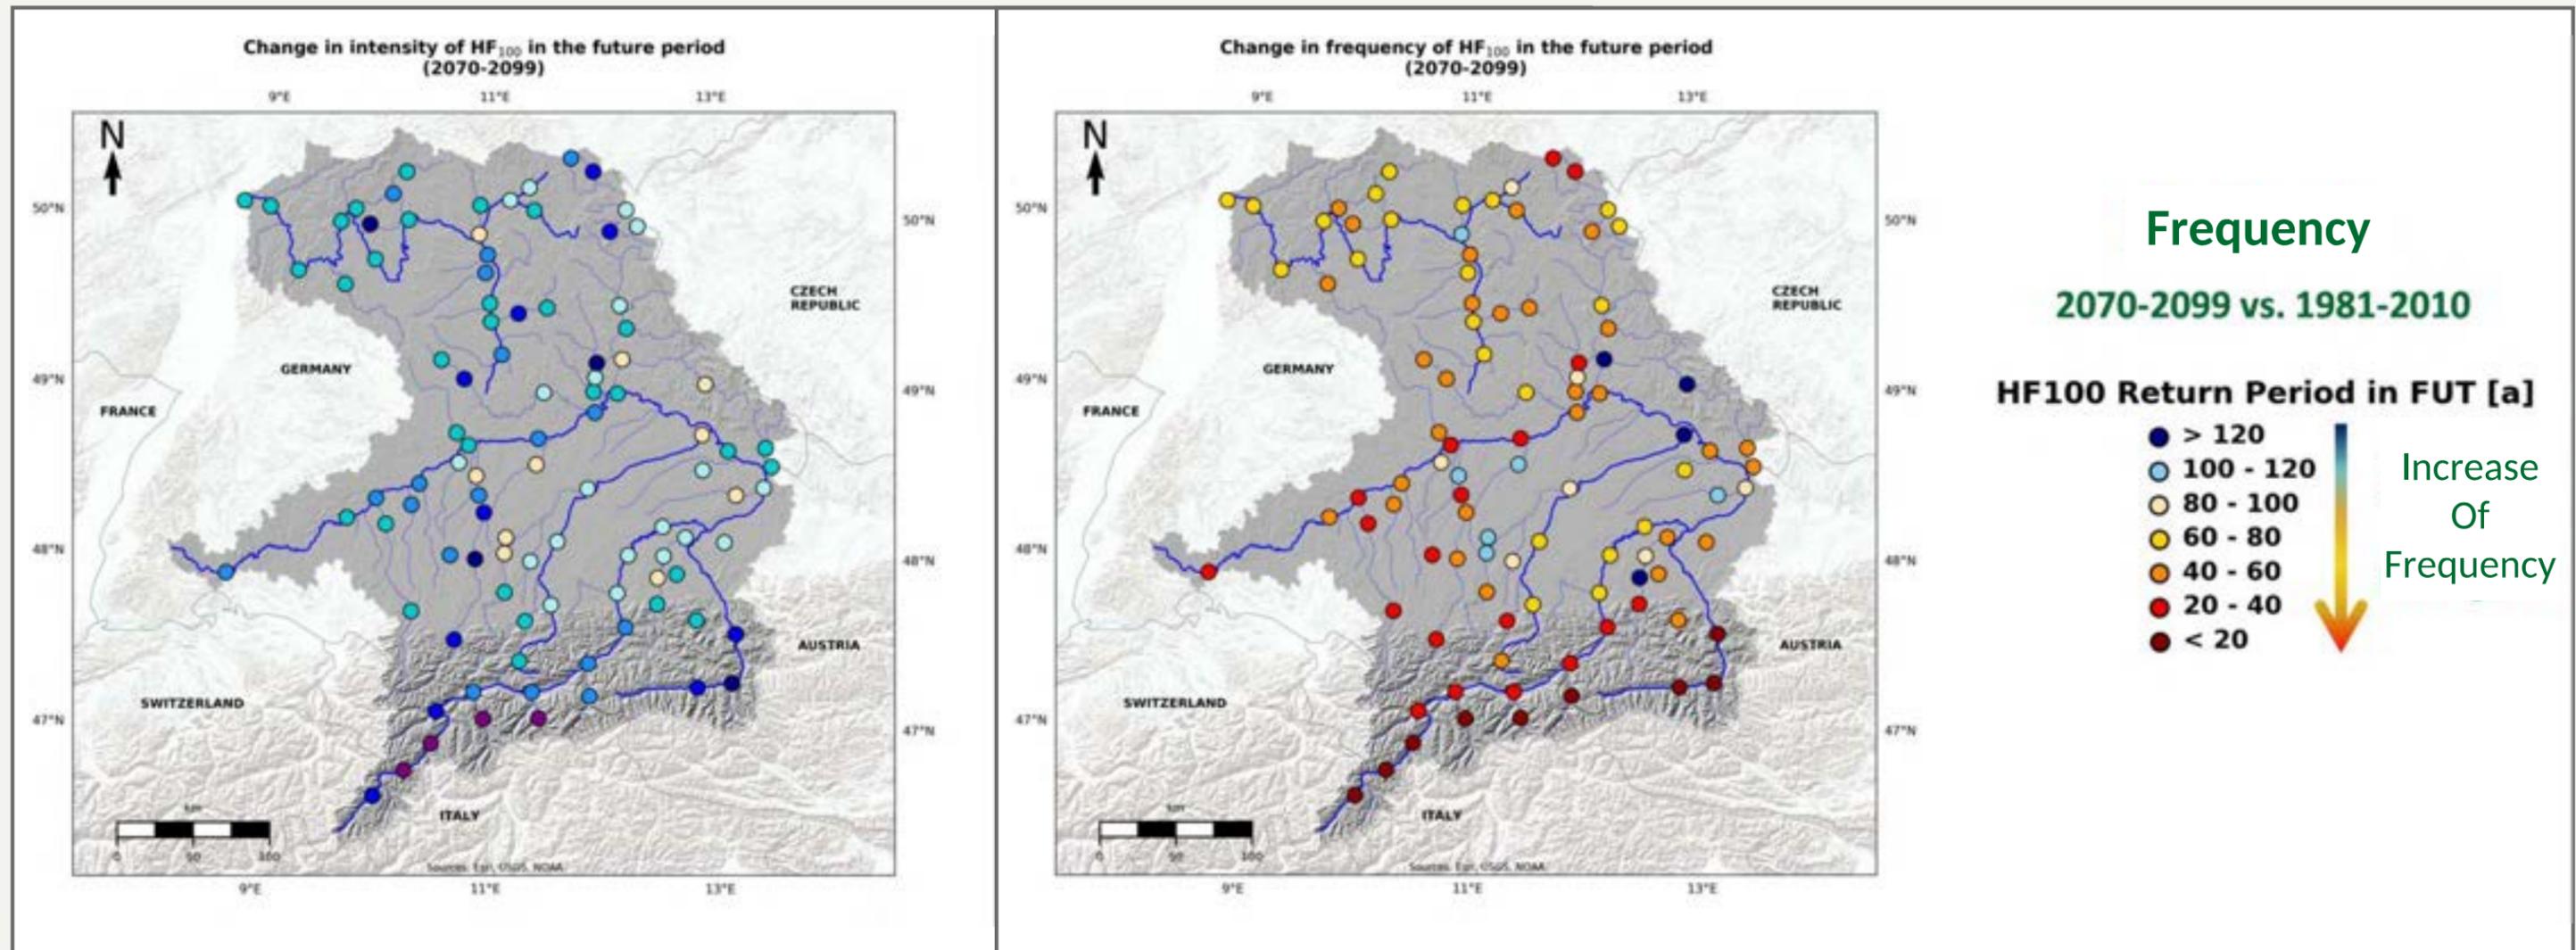


Impact of the climate crisis on extreme hydrological events such as floods and droughts.²

Will extreme flows/floods (e.g. HQ100) be more severe?



Changes in Frequency and Intensity of HQ100 in Bavarian river basins



(Willkofer, Ludwig et al., in prep.)

Background...

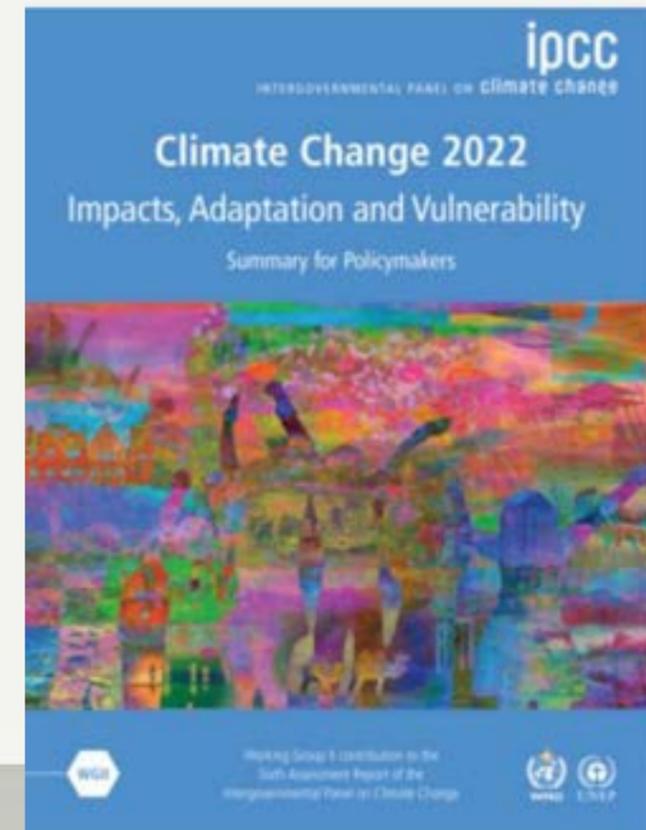
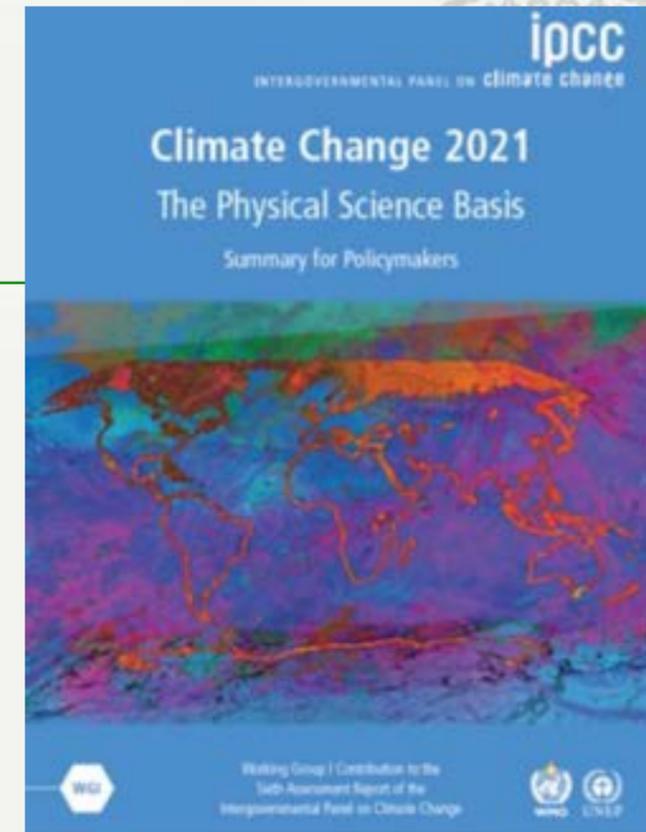
From the 6th Assessment Report of the IPCC (AR6) (2021/2022):
Working Group I: The Physical Science Basis
Working Group II: Impacts, Adaptation and Vulnerability

WGI.B – Possible Climate Futures

B.2 Many changes in the climate system become larger in direct relation to increasing global warming. They include **increases in the frequency and intensity of hot extremes**, marine heatwaves, and heavy precipitation, agricultural and ecological droughts in some regions [...]

WGII.C – Future Adaptation Options and their Feasibility

C.2 [...] **The effectiveness of adaptation** to reduce climate risk is documented for specific contexts, sectors and regions (high confidence) and **will decrease with increasing warming**. Integrated, multi-sectoral solutions [...] increase the feasibility and effectiveness of adaptation in multiple sectors.



Supercomputer SuperMUC-NG Phase 1 + 2



Bild: V. Hohenegger (LRZ)

Dieter Kranzlmüller | 02.10.2025

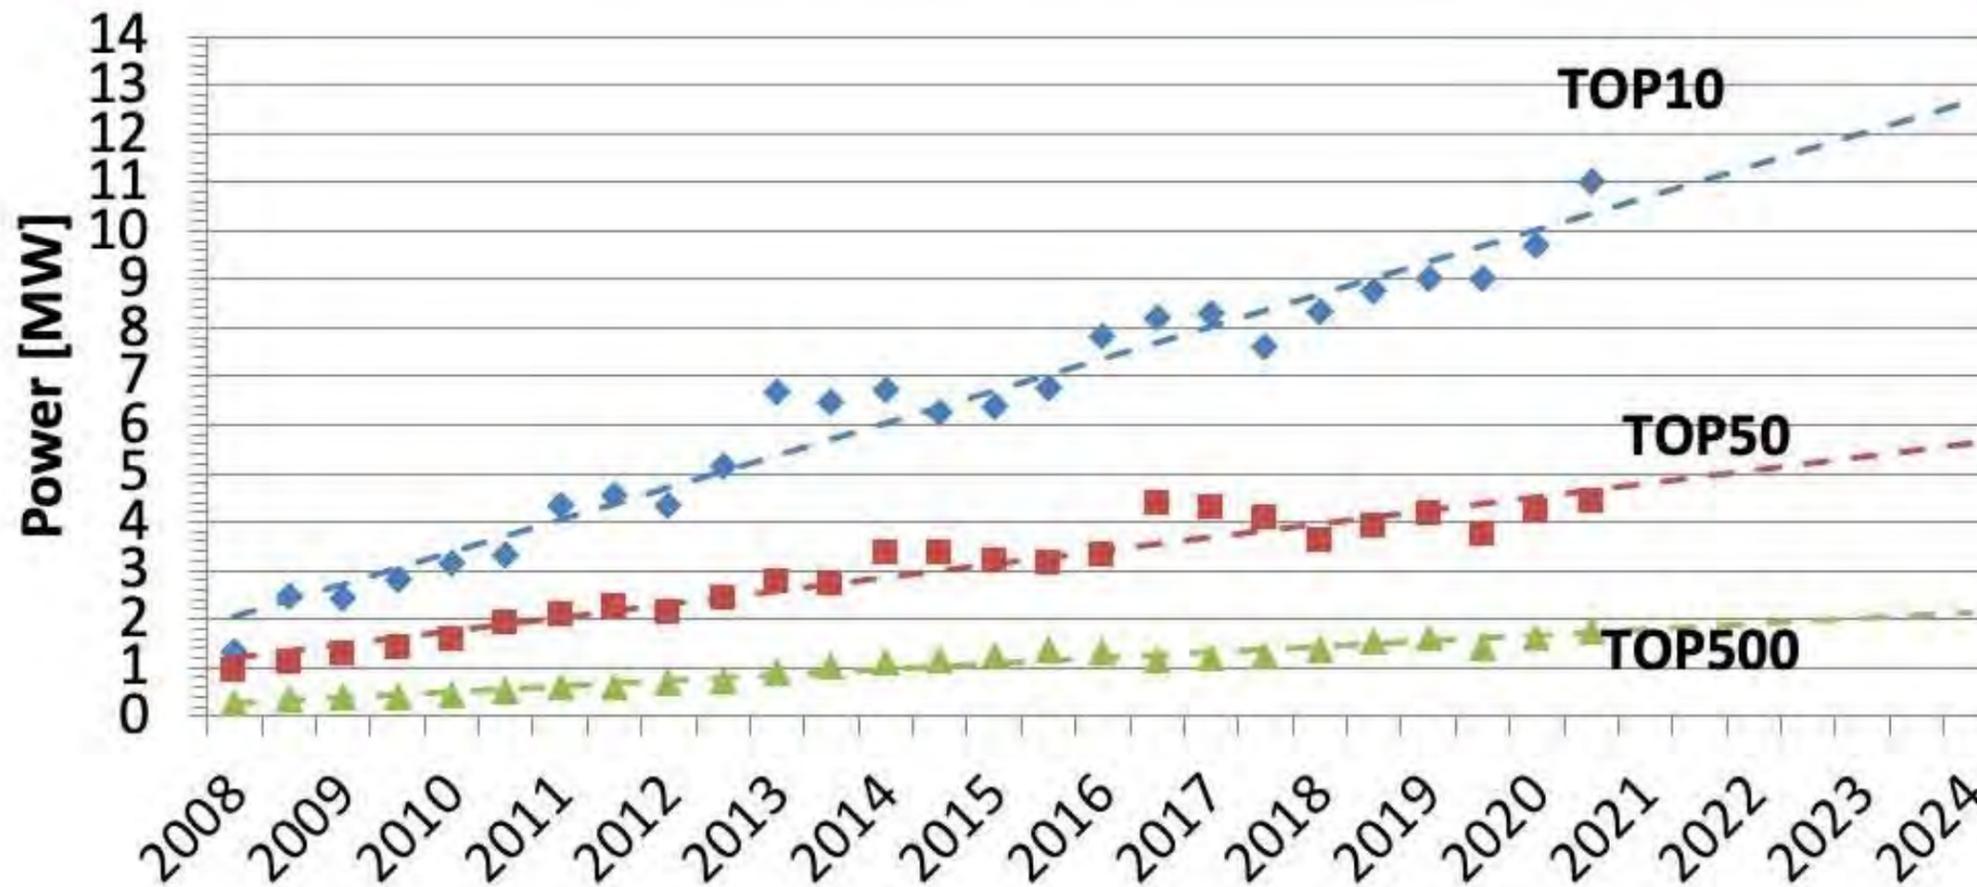
The next national supercomputer @ LRZ Blue Lion



Average Power Consumption Top 10/50/500 Systems



POWER CONSUMPTION

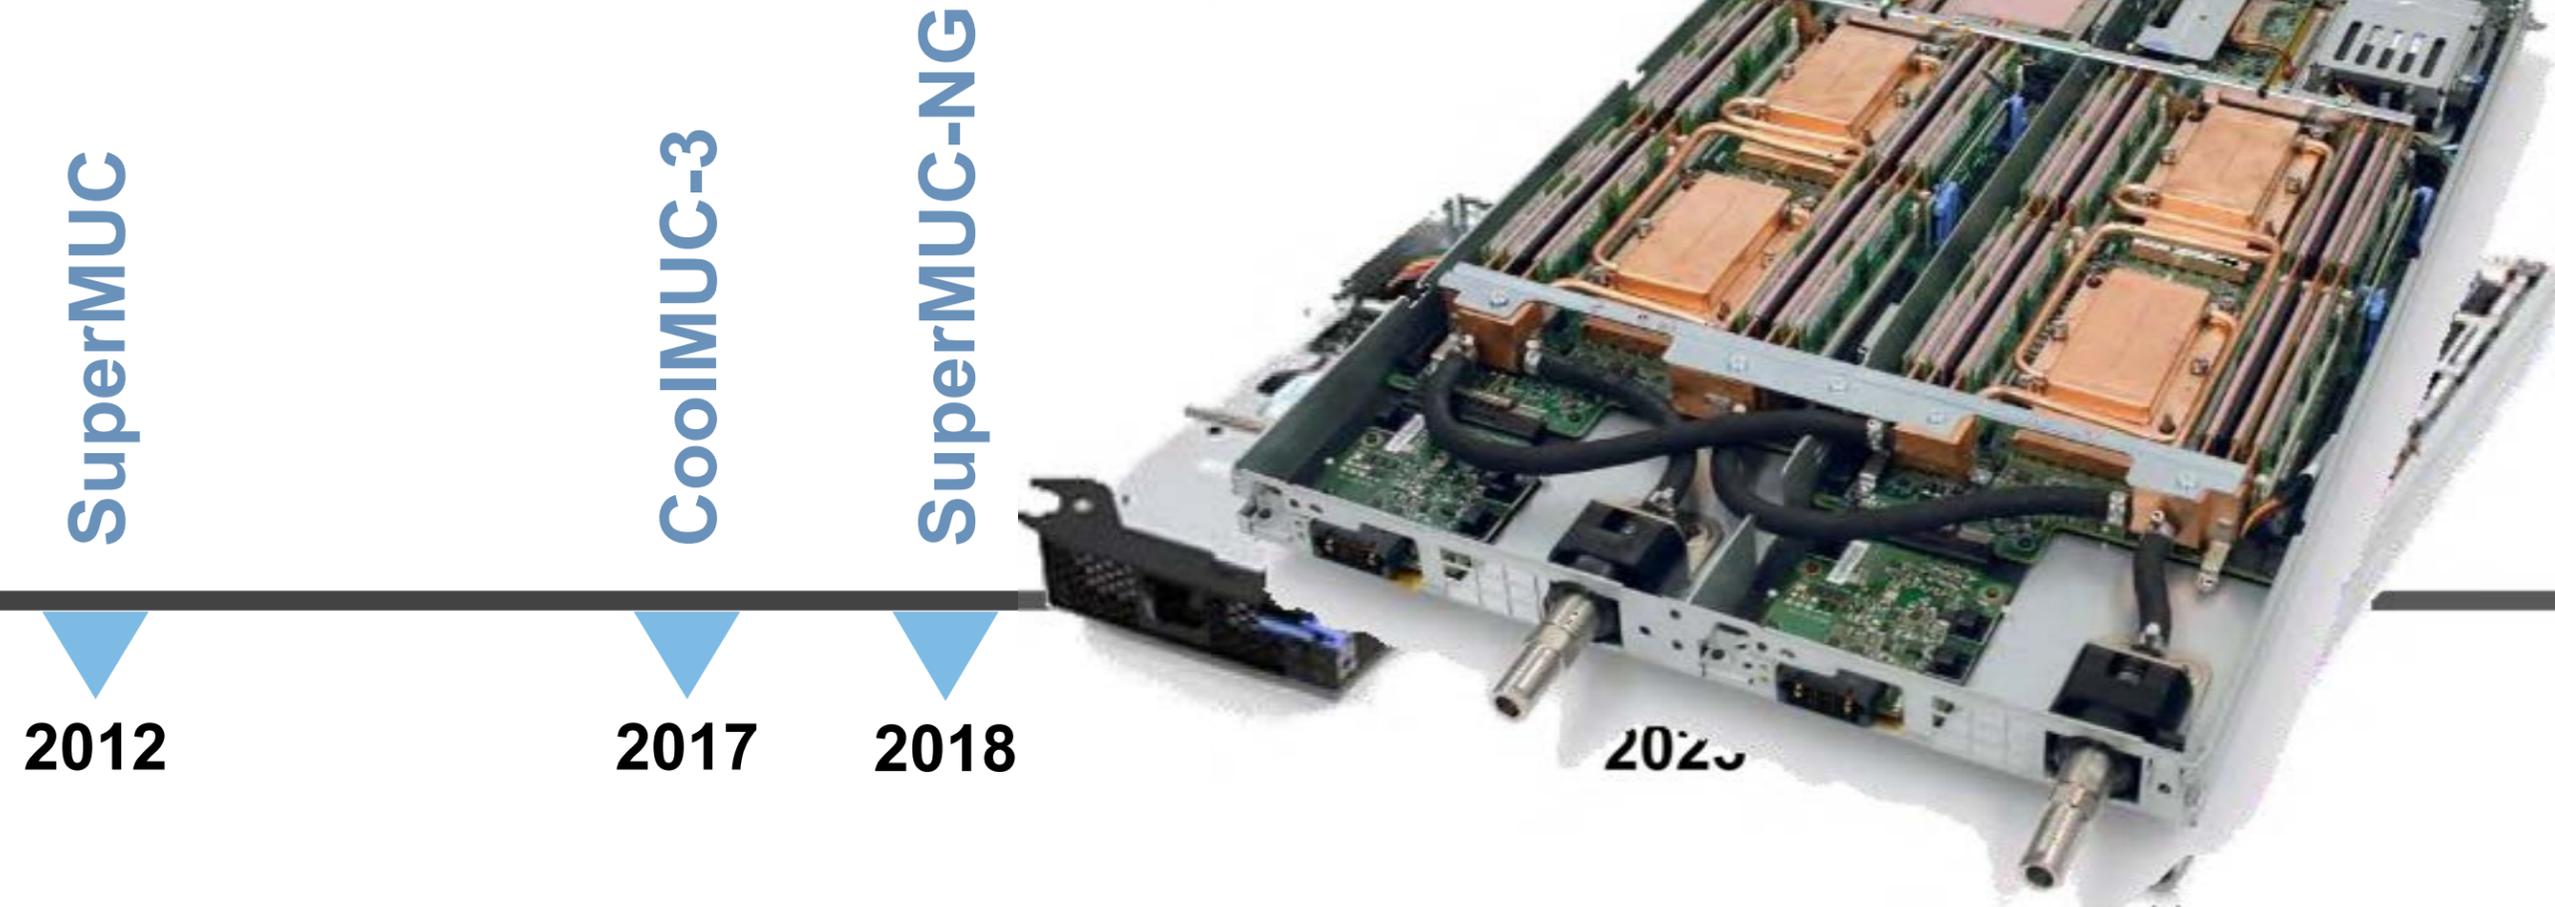


Source: <https://twitter.com/top500supercomp/status/1329583763844599808>

Hot water cooling in the compute room



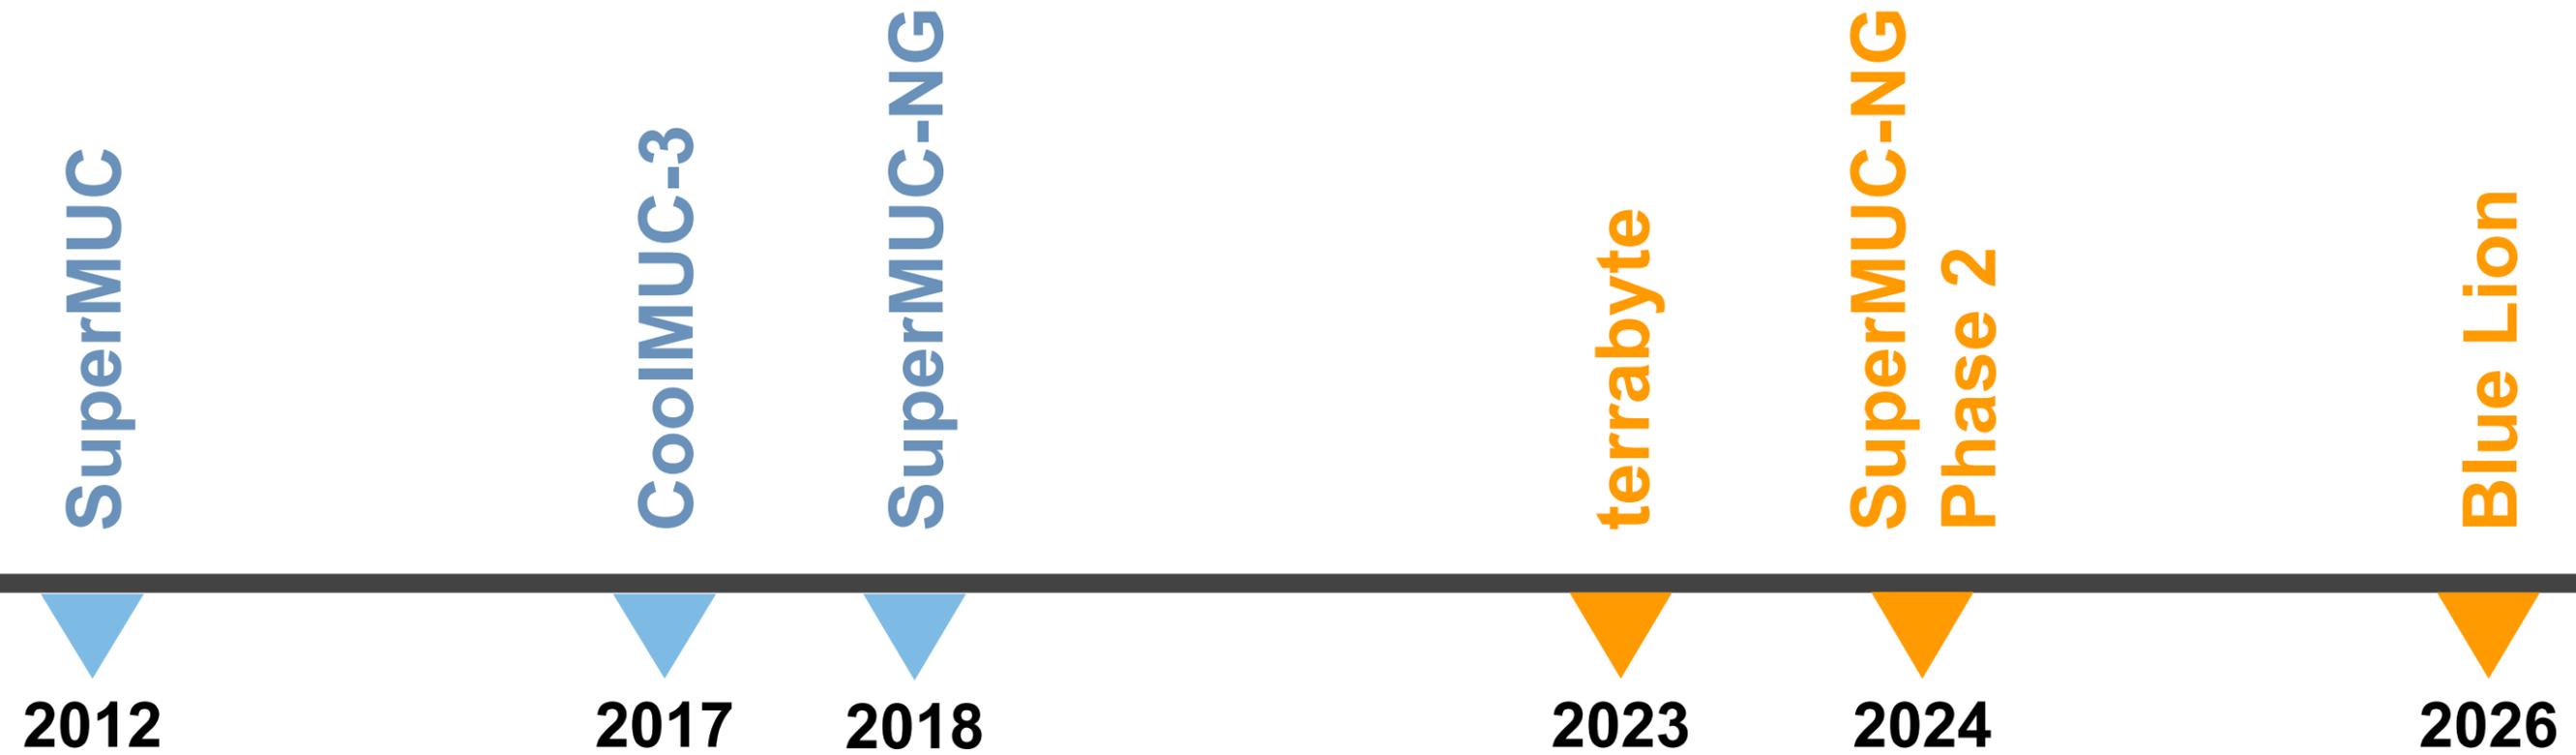
LRZ HPC/AI Systems and Energy Efficiency



2024: SuperMUC-NG Phase 2

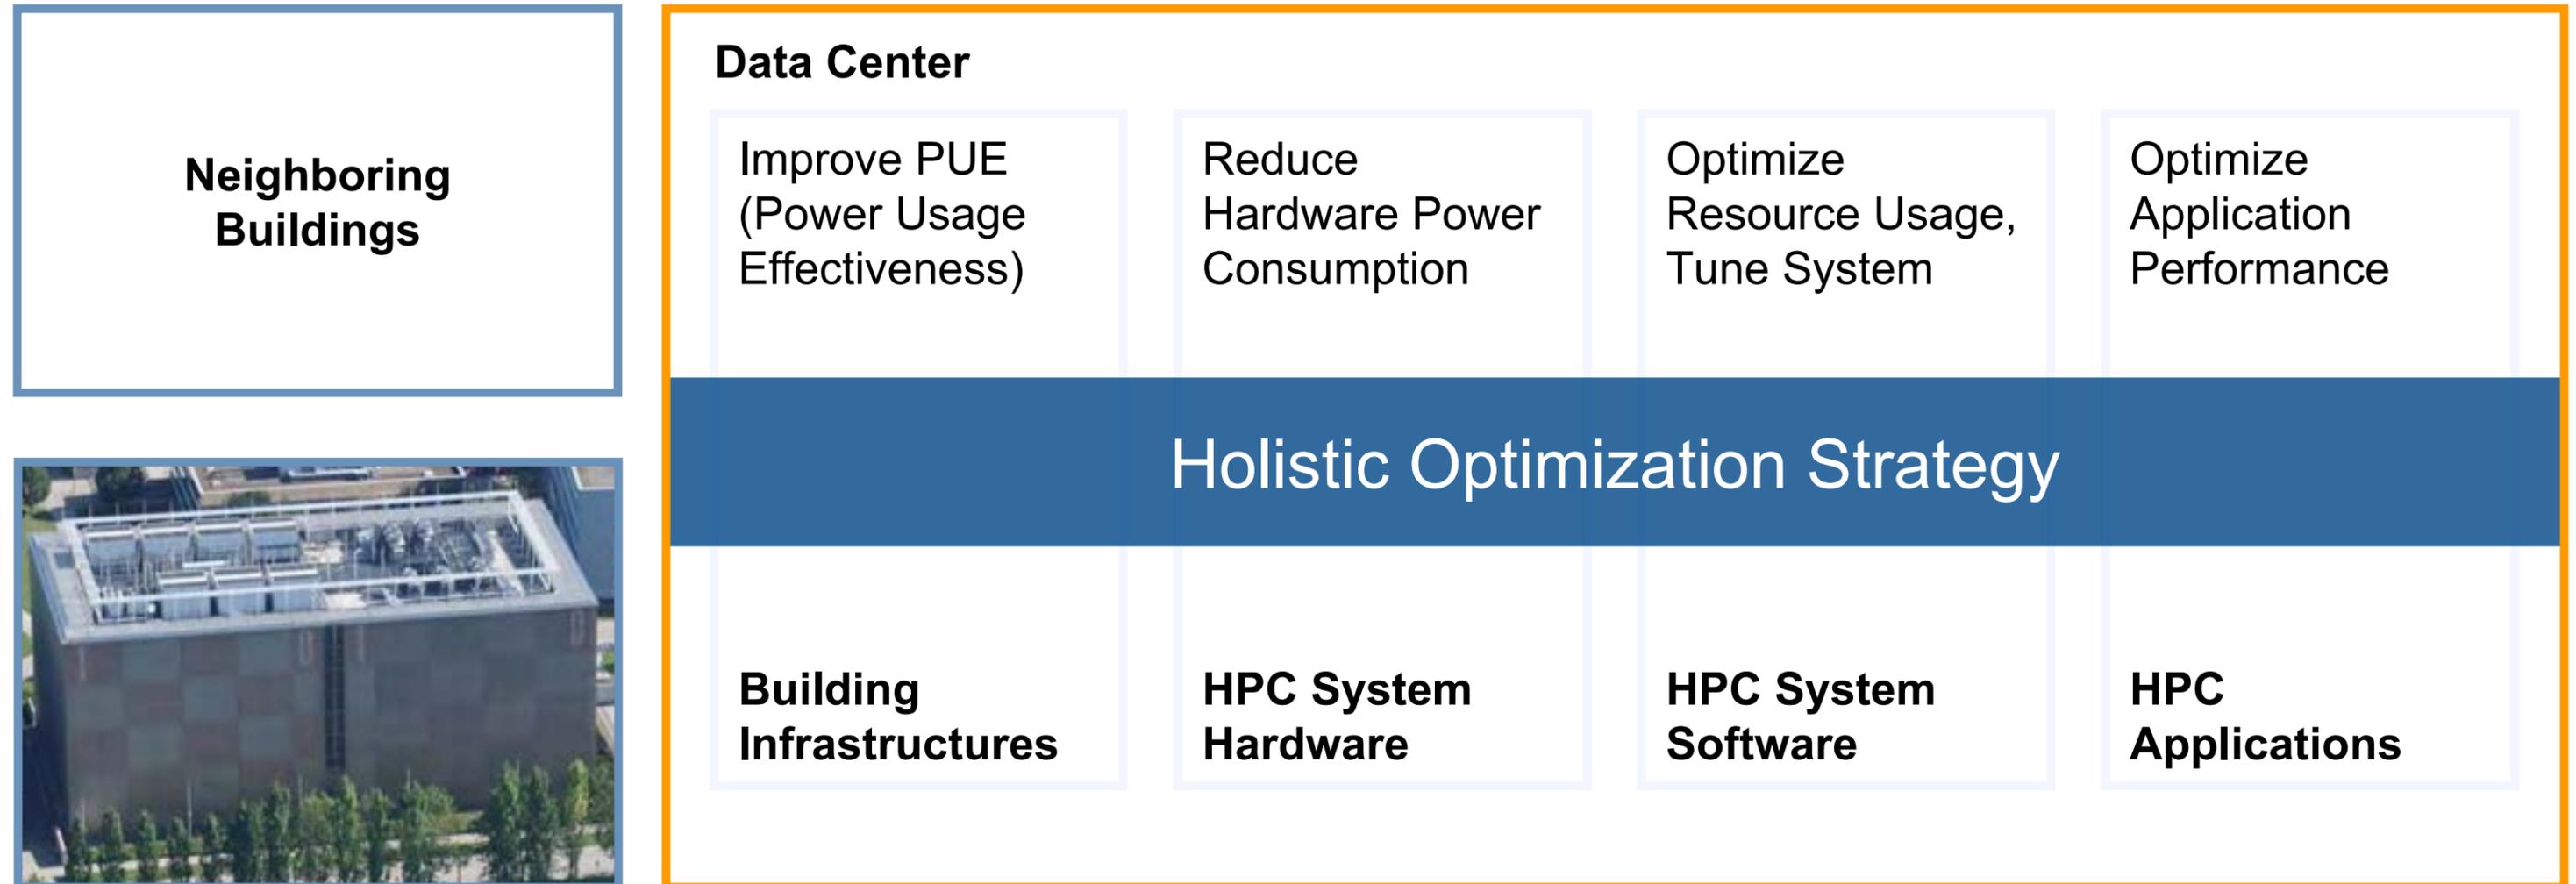
- Direct hot-water-cooled compute nodes
- Intel CPUs & Intel PVC GPUs
- PUE = 1,03

LRZ HPC/AI Systems and Energy Efficiency



2026: Blue Lion
• **Upcoming Supercomputing System**

The Four Pillar Model



Torsten Wilde, Axel Auweter, Hayk Shoukourian
„The 4 Pillar Framework for energy efficient HPC data centers“,
Comput Sci Res Dev (2014) 29: 241. <https://doi.org/10.1007/s00450-013-0244-6>

LRZ HPC Systems and Energy Efficiency



- 2012: SuperMUC – first hot water cooled system (CPU+Speicher)
- 2017: CoolMUC 3 - first 100% hot water cooled system (fan-less)
- 2018: SuperMUC-NG - higher temperatures (45 degree Celsius), 100% free cooling throughout the year, PUE = 1,03
- 2024/25: ExaMUC - ??? (procurement ongoing innovation partnership)
- DVFS - Dynamic Voltage Frequency Scaling (on average: 68% of nominal frequency)
- Control of power capping / sleep modes

SuperMUC-NG Cool Manager

3.000 l/h per rack

Reduced server power consumption

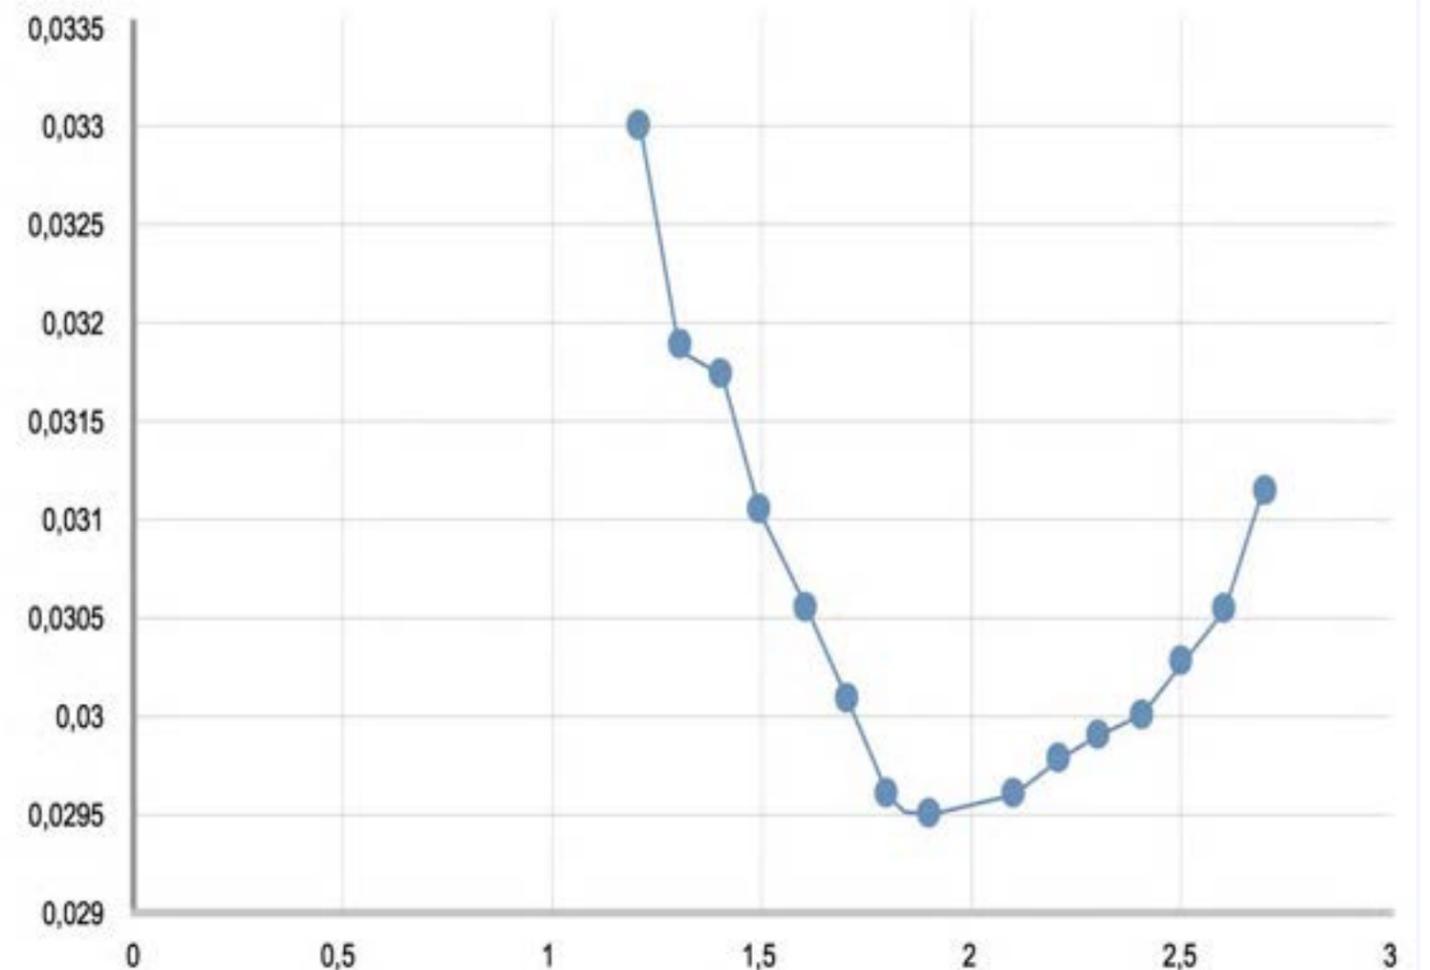
- Lower processor power consumption (~ 5%)
- No fan per node (~ 4%)

Reduce cooling power consumption

- At 40°C free cooling all year long (> 15%)

Energy Aware Scheduling

- Only CPU bound jobs get max frequency (~ 5%)



LRZ HPC Systems and Energy Efficiency



- 2012: SuperMUC – first hot water cooled system (CPU+Speicher)
- 2017: CoolMUC 3 - first 100% hot water cooled system (fan-less)
- 2018: SuperMUC-NG - higher temperatures (45 degree Celsius), 100% free cooling throughout the year, PUE = 1,03
- 2024/25: ExaMUC - ??? (procurement ongoing innovation partnership)
- DVFS - Dynamic Voltage Frequency Scaling (on average: 68% of nominal frequency)
- Control of power capping / sleep modes
- 100% renewable energies
- Cooling infrastructure optimizations - Glycol free cooling
- Waste heat usage
- Adsorption cooling machines
- Future (unconventional) computing technologies – Accelerators, Quantum, ...

Energy Efficiency Optimizing the PUE



Energy Efficiency at LRZ 2017-2022

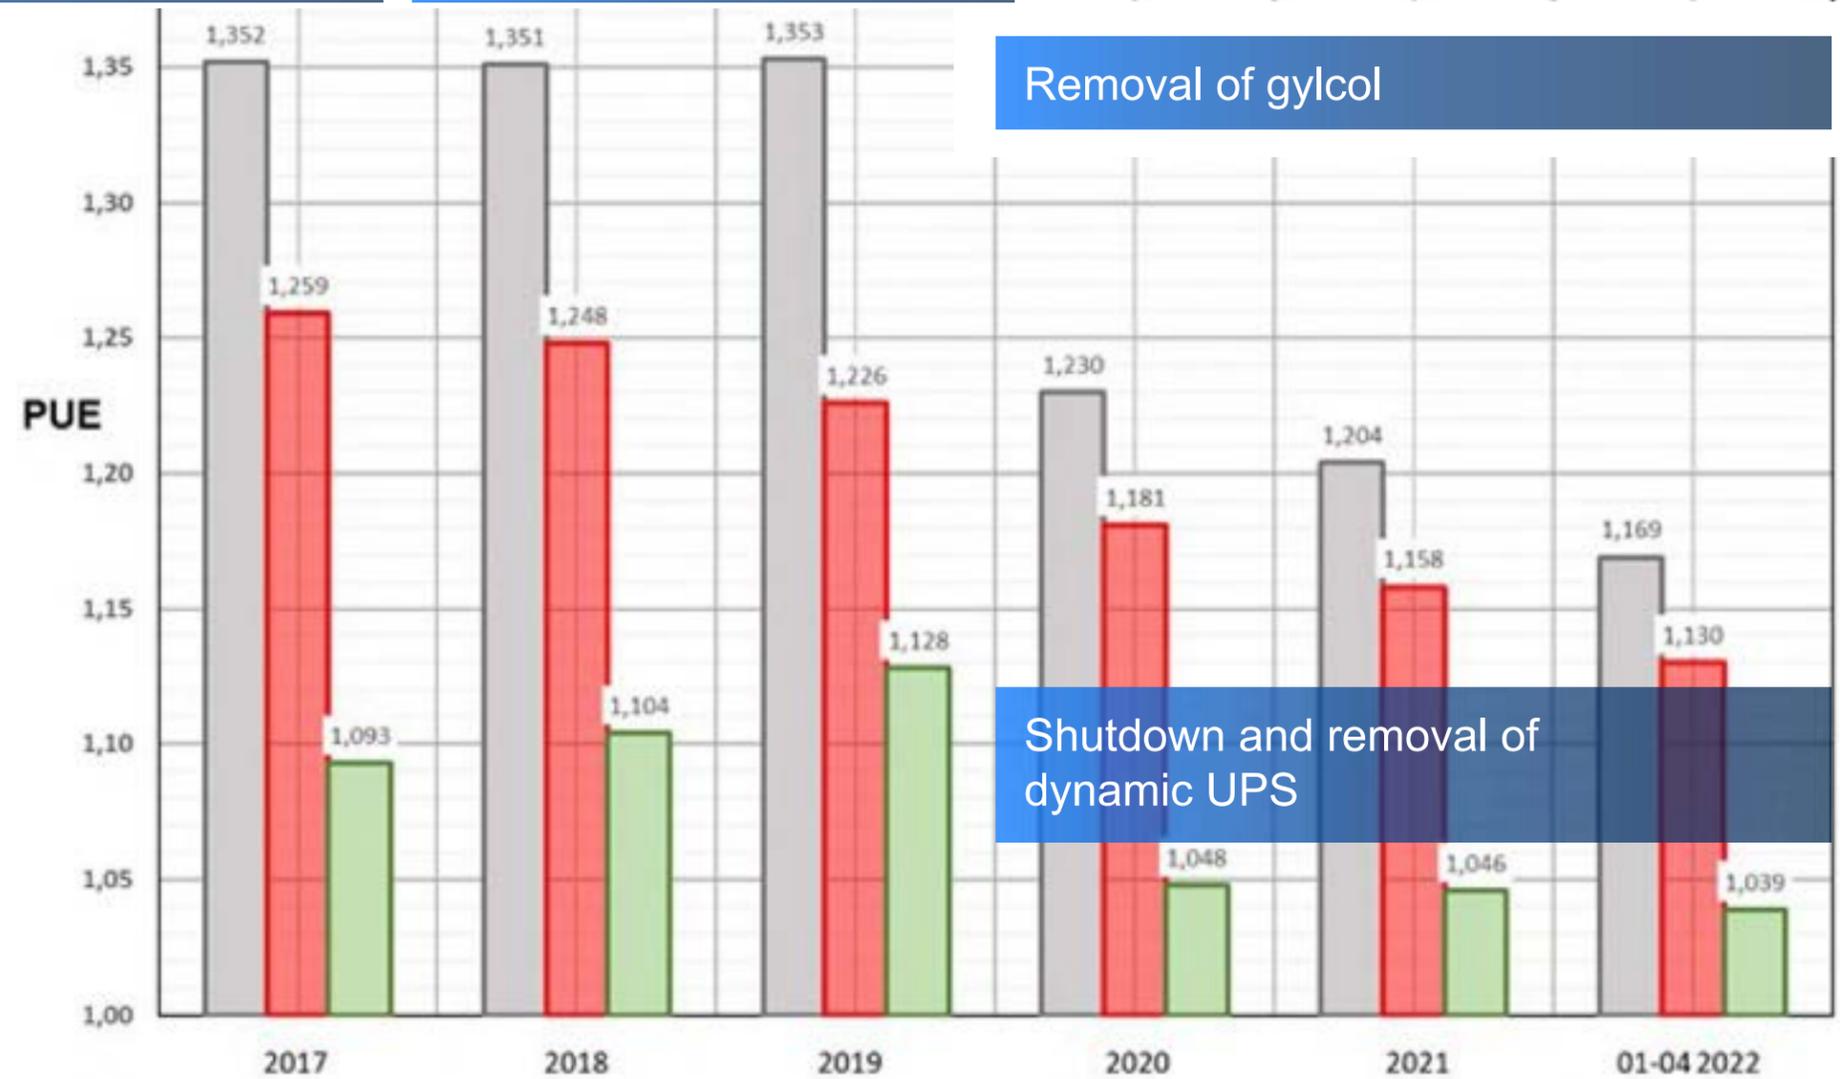
Removal of hydraulic switches

Optimization of controllers

Removal of glycol

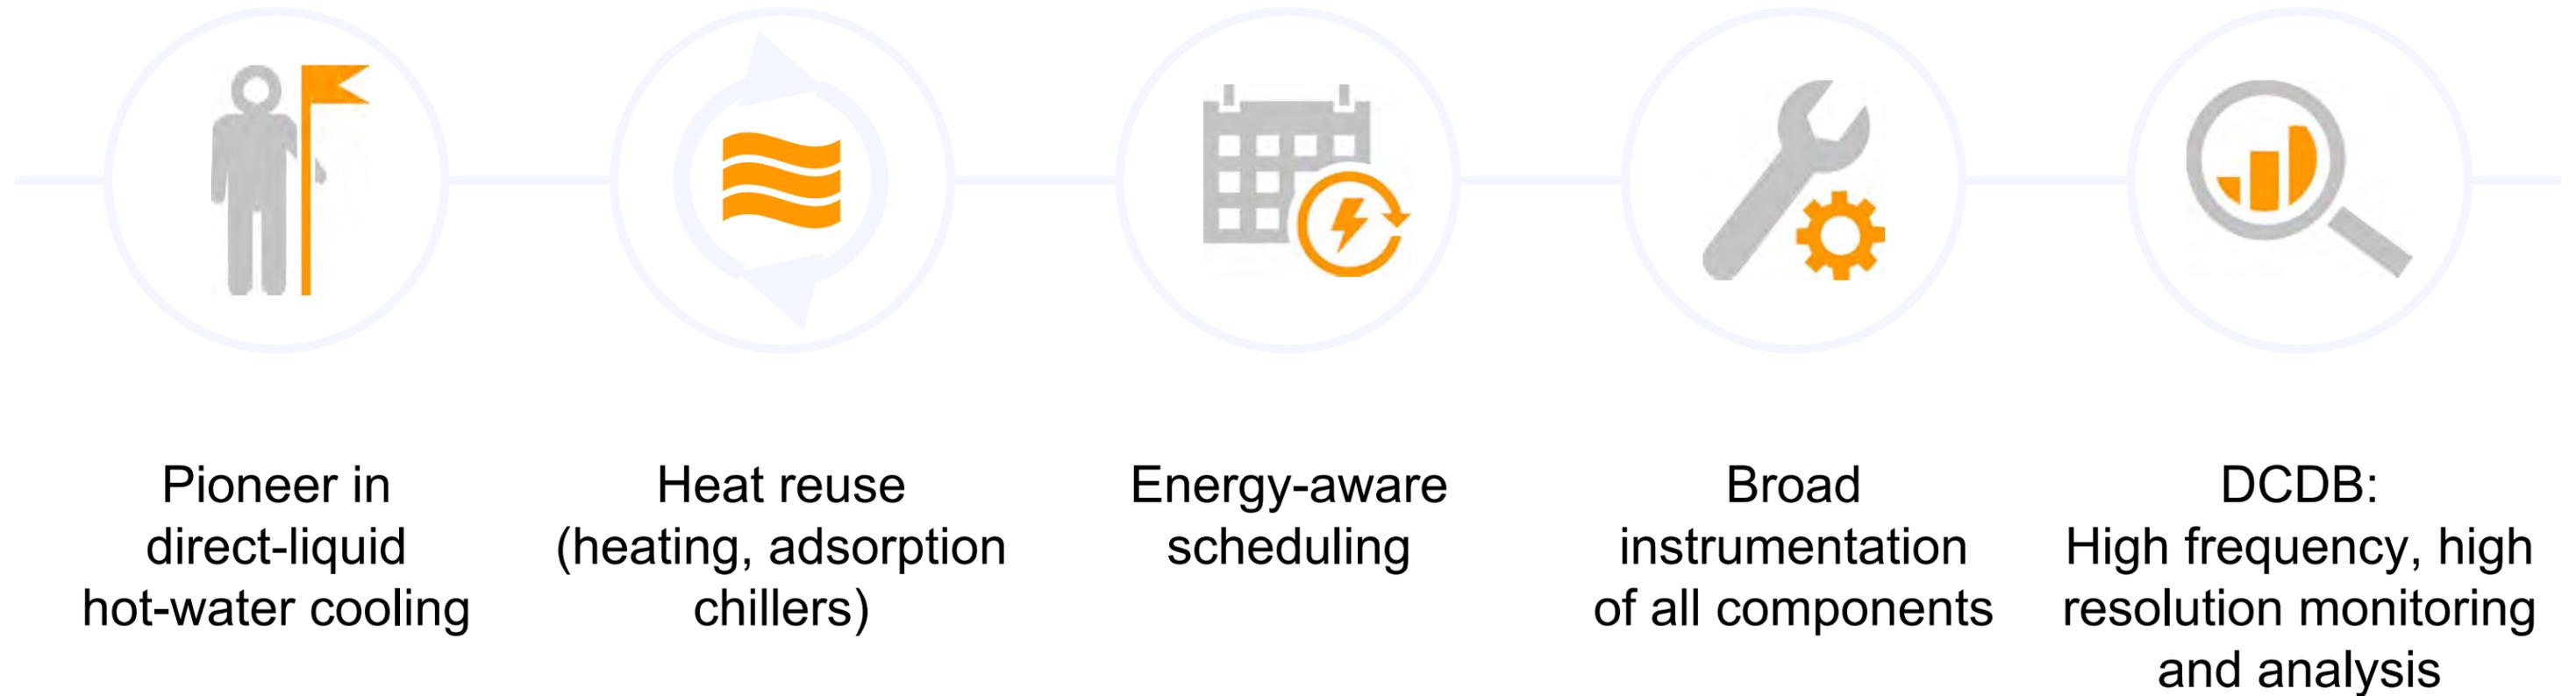
Shutdown and removal of
dynamic UPS

$$PUE = \frac{\sum P(i)}{P(IT)}$$



■ Total
 ■ Cooling
 ■ Power
 1 Power Usage Effectiveness

Energy Efficiency at LRZ



The Munich Center for Machine Learning (MCML)



Left:

- 2-3 NVIDIA DGX A100
- =16-24 GPUs per Rack
- Air-cooled
(in a water-cooled rack)
- PUE: 1,65-1,80

Power Usage Effectiveness:

the **total amount of power** entering a data center divided by the power used to run the **IT equipment**

Example:

PUE 1,65 = 1 für IT
+ 65% für Infrastructure

The Munich Center for Machine Learning (MCML)



Left:

- 2-3 NVIDIA DGX A100
- =16-24 GPUs per Rack
- Air-cooled
(in a water-cooled rack)
- PUE: 1,65-1,80

Right:

- Up to 144 GPUs per Rack
- Lenovo HGX-based nodes
(4x A100)
 - Direct hot-water cooled
(allows free cooling)
 - PUE: 1,03-1,05





Cerebras Wafer-Scale Engine (WSE-2)



NextGen AI System

**Cerebras CS-2 System
w/ HPE Superdome Flex**

850.000 Cores

2,6 Bio. Transistors

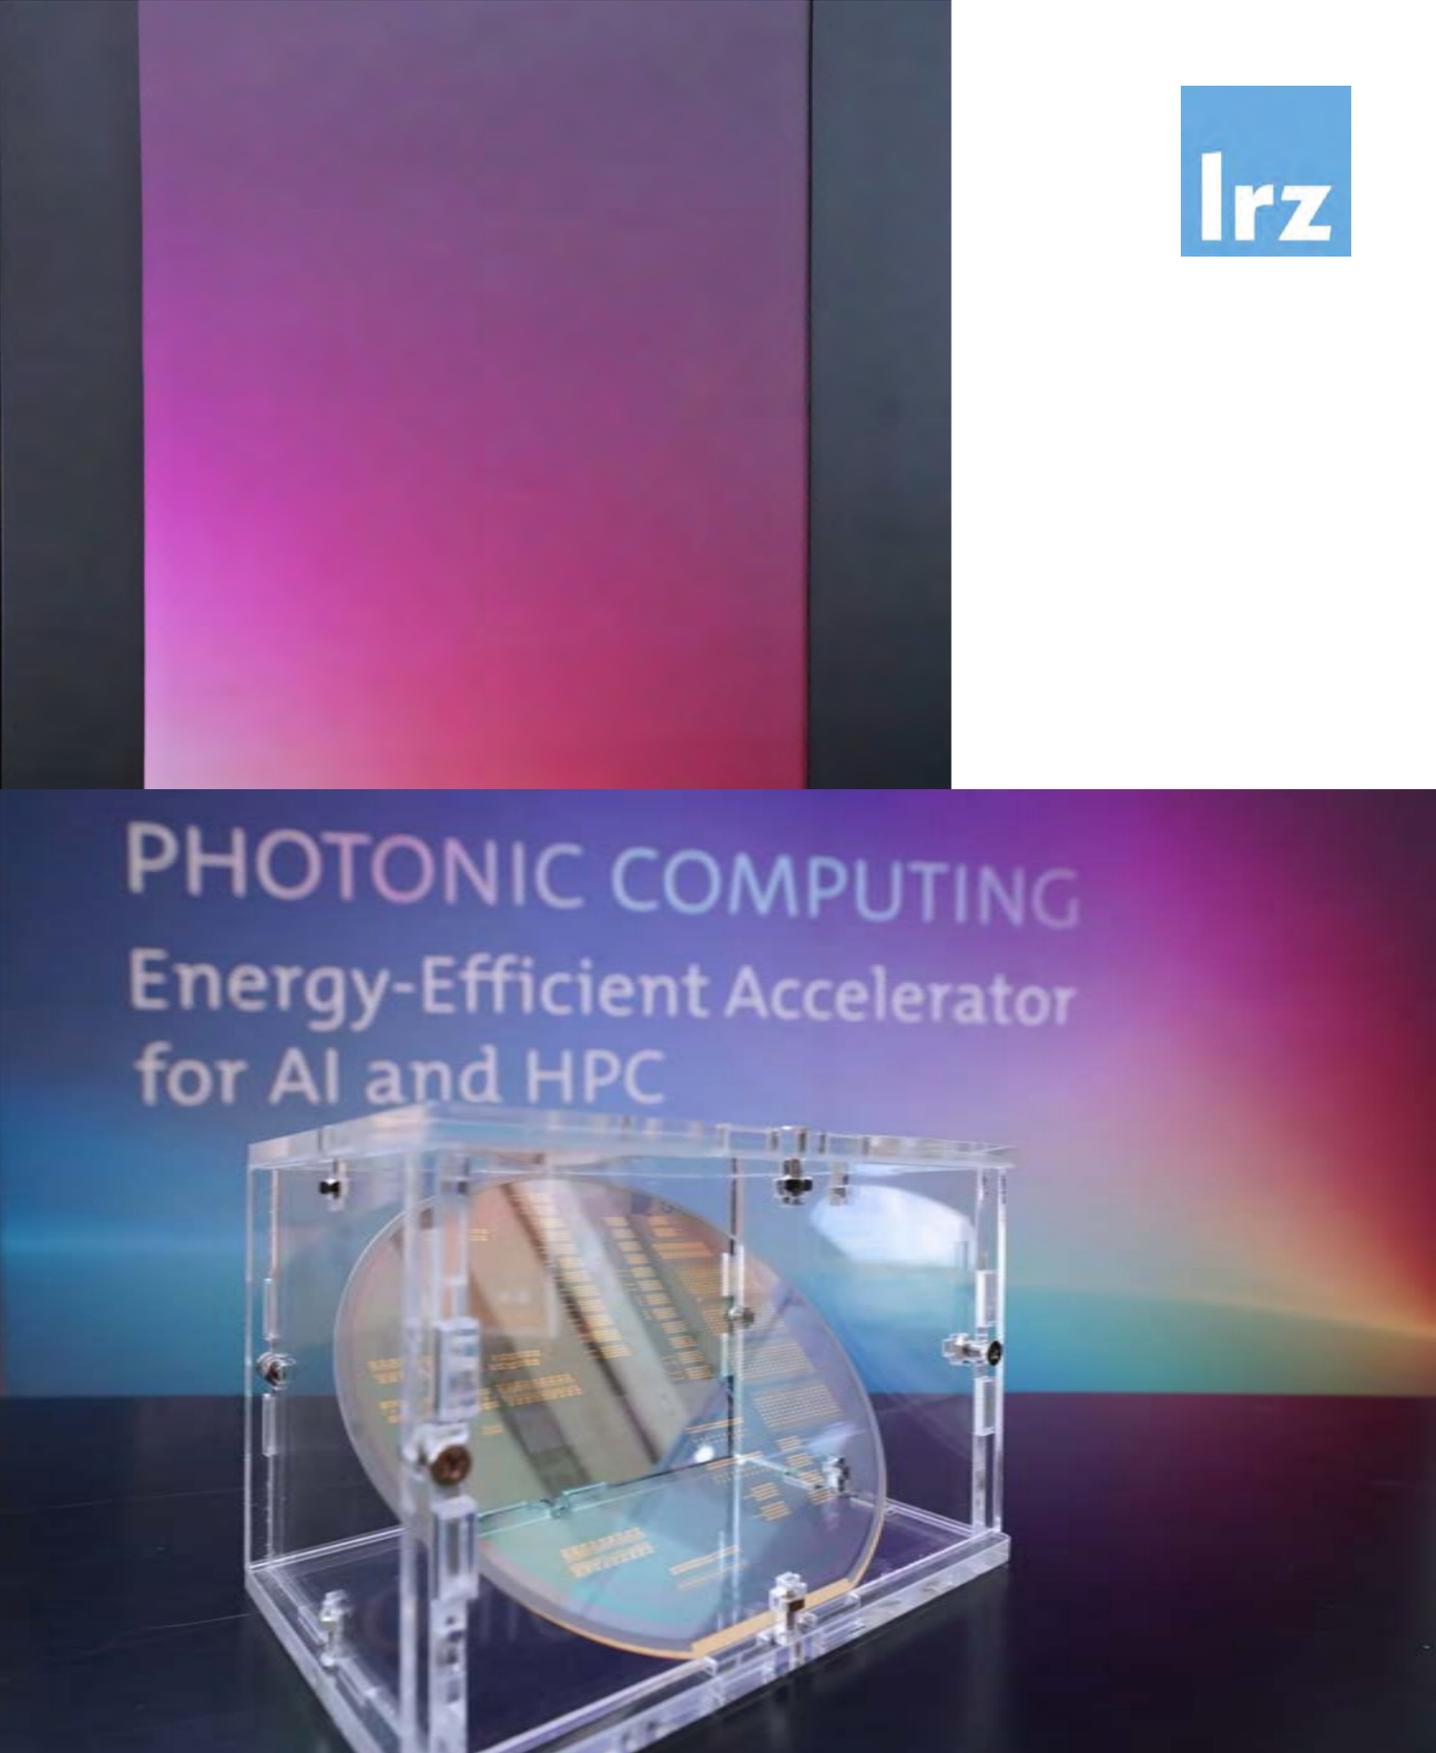
40 GB Memory

220 Petabyte/s Data Transfer

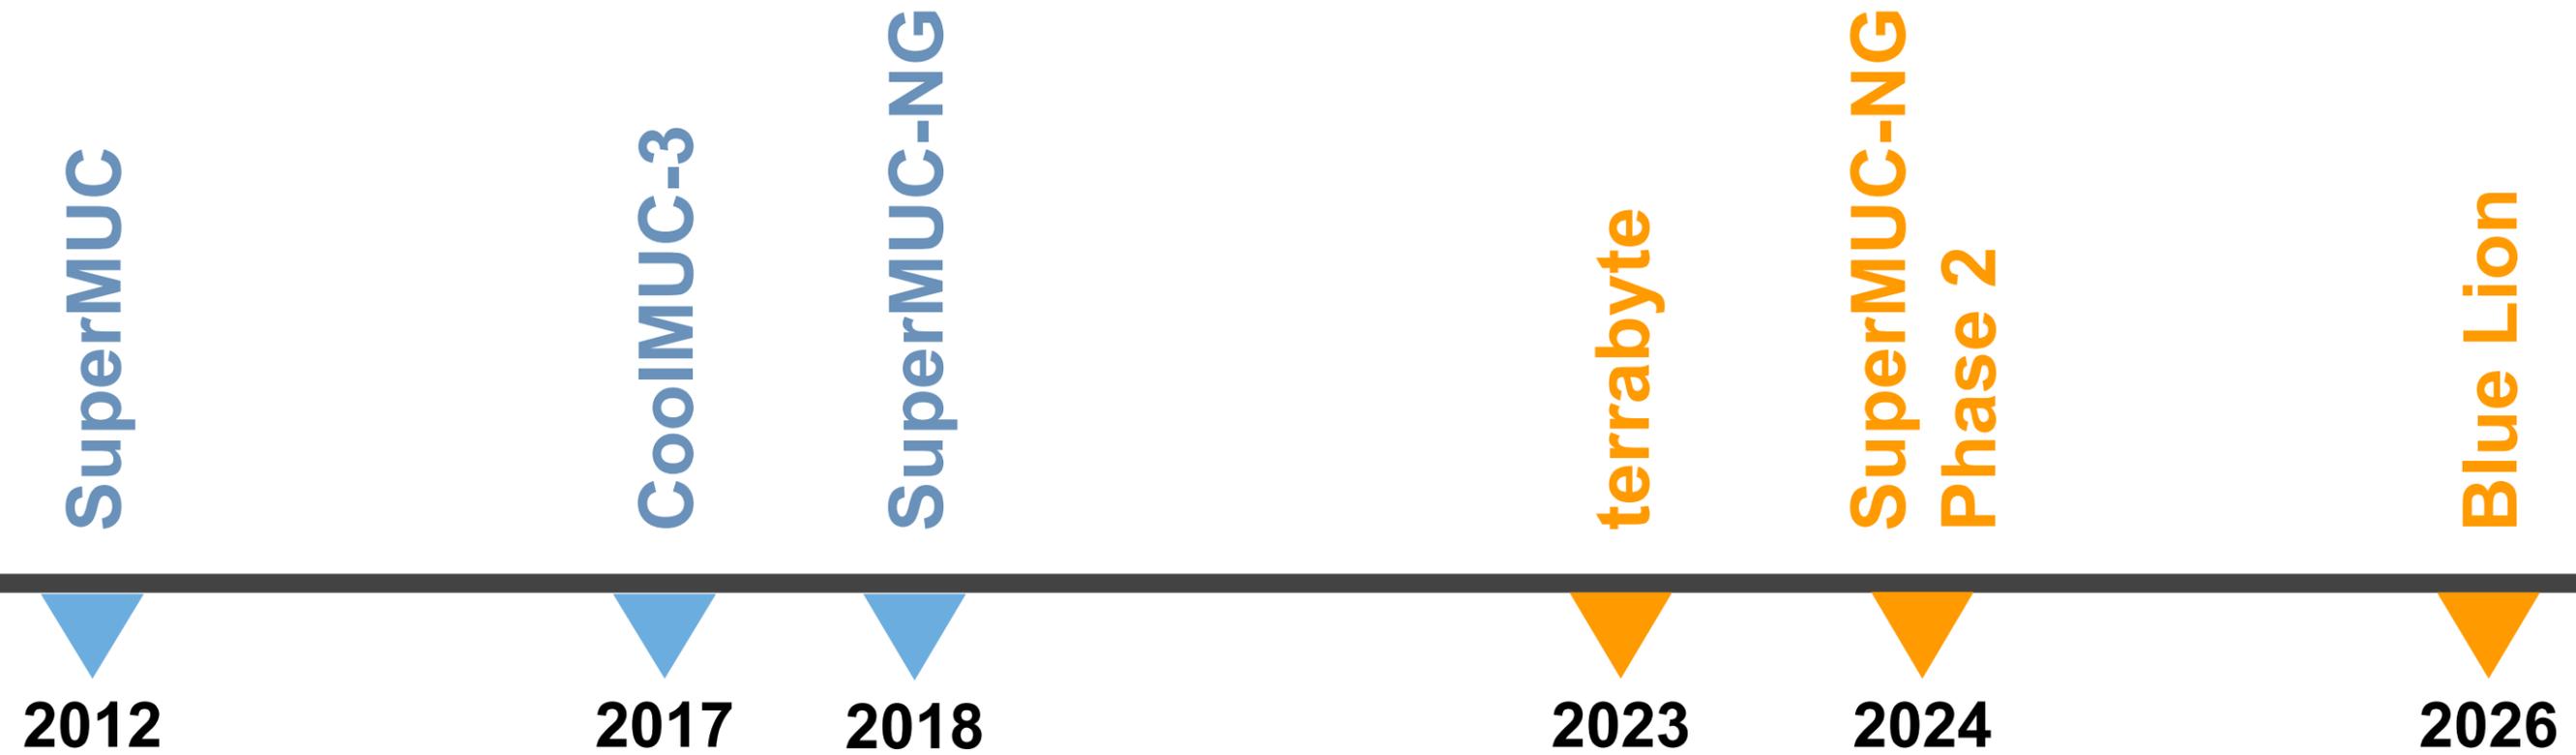
Germany's First Quantum Demonstrator Q-Exa



Photonic Computing @ LRZ



LRZ HPC/AI Systems and Energy Efficiency



2026: Blue Lion
• **Upcoming Supercomputing System**

Sustainable HPC and AI



- **Energy efficiency** is crucial for HPC and AI
- We need a **holistic approach** - instead GFlop/watt
- Future technologies such as quantum, photonics, and others must **be integrated**

Contact: Prof. Dr. Dieter Kranzlmüller
kranzlmue@lrz.de | @Kranzlmue

60 Years of LRZ

Towards the Future



Bildnachweis: StMWK – Axel König

HPC Roundtable 2025