



Leibniz Supercomputing Centre  
of the Bavarian Academy of Sciences and Humanities

1

A photograph of a modern, multi-story building with a glass facade, partially obscured by a semi-transparent blue overlay. The LRZ logo is in the top right corner of the image.

**HPC and AI in Supercomputing Centres: Green Today, Greener Tomorrow**

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

With contributions from  
Michael Ott, Axel Auweter,  
Josef Weidendorfer,  
Nicolay Hammer, Max Höb,  
and many others ...

10.03.24 | D. Kranzlmüller

2

2

Leibniz Supercomputing Centre (LRZ)  
of the Bavarian Academy of Sciences and Humanities  
Partner for Digital Transformation of Science & Research

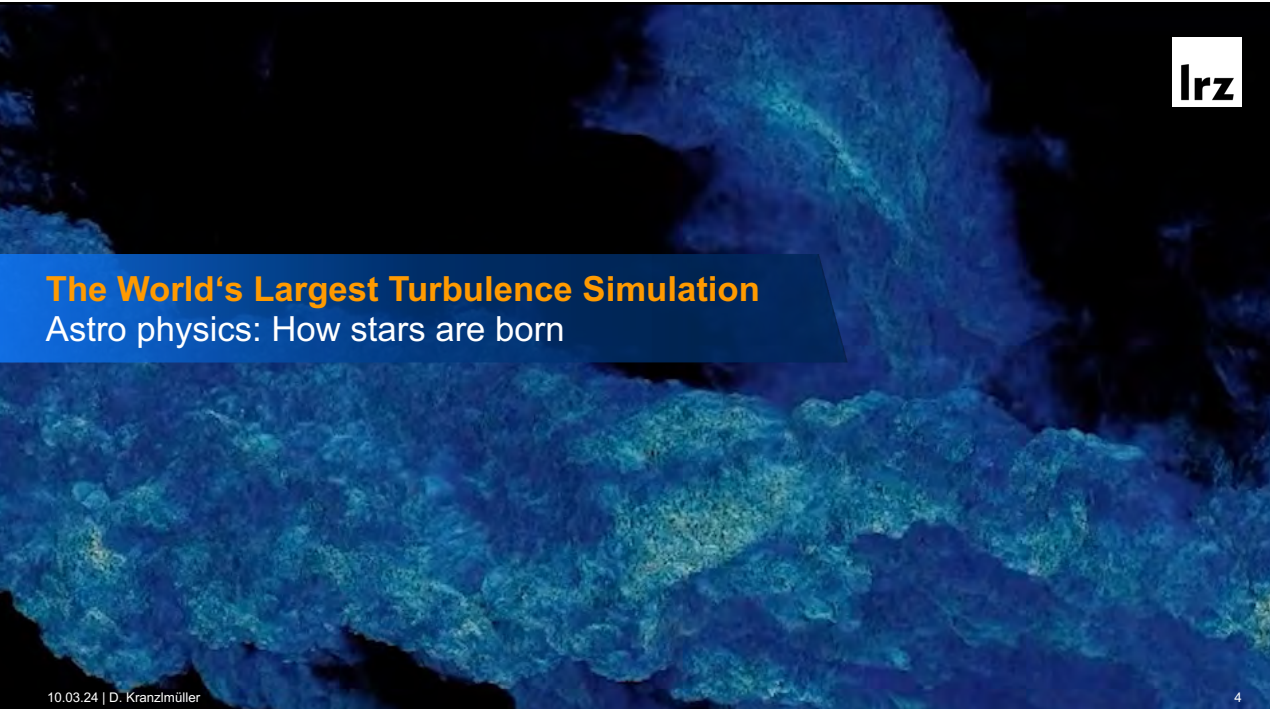


**BA&W** **lrz**

**IT-Services for Science since 1962**  
Munich Universities, Bavaria, Germany, ...

10.03.24 | D. Kranzlmüller 3

3



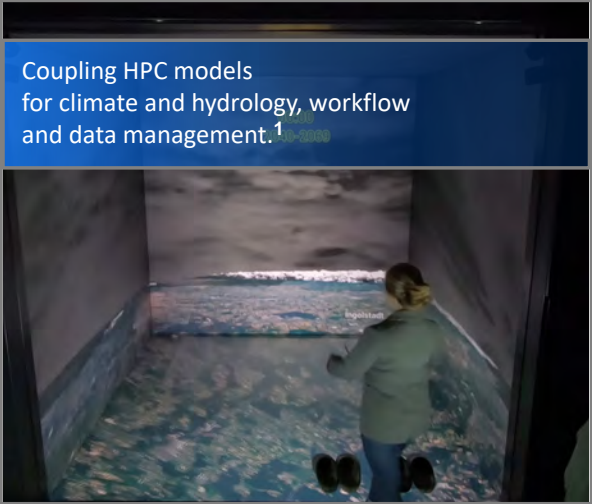
**lrz**

**The World's Largest Turbulence Simulation**  
Astro physics: How stars are born

10.03.24 | D. Kranzlmüller 4

4

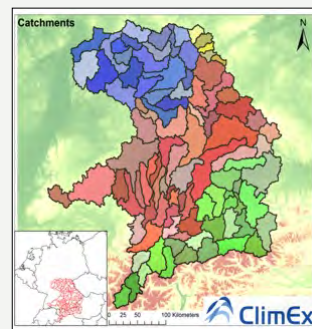
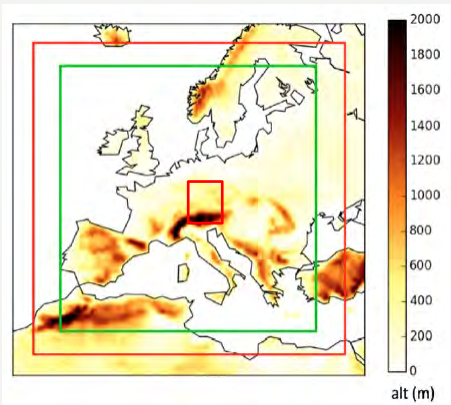
# Bavarian-Canadian Project ClimEx



www.climex-project.org; 1 © LRZ / Ouranos; 2 Danube Flooding 2013 © sueddeutsche.de

5

# ClimEx – Scales of application





- Scope:** Assess the CC related dynamics of extremes events from the continental to the river basin scale
- Goal:**
- a) Improve process understanding of non-linear hydro-meteorological extreme events
  - b) Provide adaptation and management options for stakeholders to reduce related risks

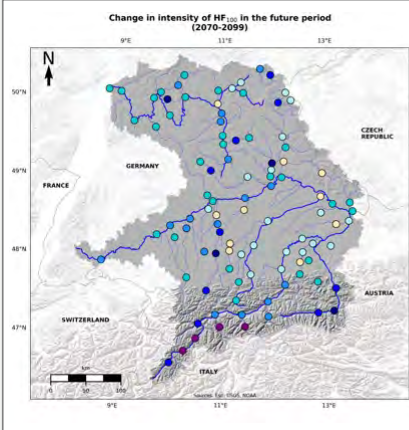
Slide courtesy Ralf Ludwig, LMU

6

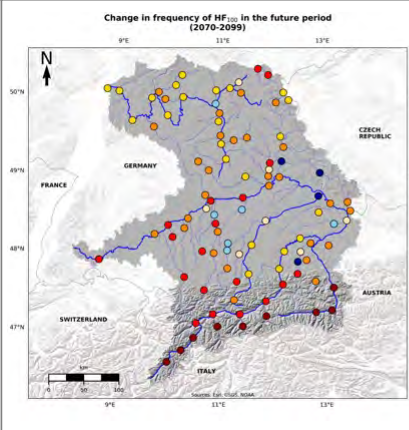
## Will extreme flows/floods (e.g. HQ100) be more severe?

### Changes in Frequency and Intensity of HQ100 in Bavarian river basins



Change in intensity of HF<sub>100</sub> in the future period (2070-2099)



Change in frequency of HF<sub>100</sub> in the future period (2070-2099)


Frequency

2070-2099 vs. 1981-2010

HF100 Return Period in FUT [a]

- > 120
- 100 - 120
- 80 - 100
- 60 - 80
- 40 - 60
- 20 - 40
- < 20

Increase Of Frequency



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

Slide courtesy Ralf Ludwig, LMU
14

7

## Background...

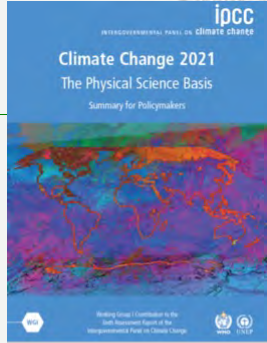
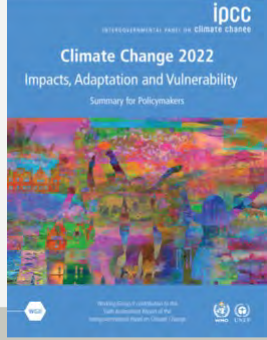
From the 6th Assessment Report of the IPCC (AR6) (2021/2022):  
Working Group I: The Physical Science Basis  
Working Group II: Impacts, Adaptation and Vulnerability

*WG1.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.*

Slide courtesy Ralf Ludwig, LMU

8



9

#	Site	Manufacturer	Computer	Country	Cores	Power (MW)	Power (MW)
1	Oak Ridge National Laboratory	IBM	Summit IBM Power System, P9 22C 3.97GHz, Mellanox EDR, NVIDIA GV100	USA	2,397,824	143.5	9.8
2	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA GV100	USA	1,572,480	94.6	7.4
3	National Supercomputing Center in Wuxi	NRCPC	Sunway TaihuLight NRCPC Sunway SW26010, 260C 1.45GHz	China	10,649,600	93.0	15.4
4	National University of Defense Technology	NUDT	Tianhe-2A ANUDT TH-IVB-FEP, Xeon 12C 2.2GHz, Matrix-2000	China	4,981,760	61.4	18.5
5	Swiss National Supercomputing Centre (CSCS)	Cray	Piz Daint Cray XC50, Xeon E5 12C 2.8GHz, Arles, NVIDIA Tesla P100	Switzerland	387,872	21.23	2.38
6	Los Alamos NL / Sandia NL	Cray	Trinity Cray XC40, Intel Xeon Phi 7250 68C 1.4GHz, Arles	USA	979,072	20.16	7.58
7	National Institute of Advanced Industrial Science and Technology	Fujitsu	AI Bridging Cloud Infrastructure (ABCI) PRIMERGY CX2550 M4, Xeon Gold 20C 2.4GHz, IB-EDR, NVIDIA V100	Japan	391,680	19.9	1.65
8	Leibniz Rechenzentrum	Lenovo	SuperMUC-NG ThinkSystem SD530, Xeon Platinum 9174 34C 3.1GHz, Intel Omni-Path	Germany	305,856	19.5	
9	Oak Ridge National Laboratory	Cray	Titan Cray XK7, Xeon E5 12C 2.8GHz, Arles, NVIDIA Tesla K20	USA	560,640	17.6	8.20
10	Lawrence Livermore National Laboratory	IBM			1,572,864	17.2	7.80

10



**~9,6 Billion**  
Core Hours

**~2,4 Mio.**  
Jobs

**~515**  
Projects

**~1.400**  
Scientists

**SuperMUC-NG**


10.03.24 | D. Kranzlmüller

**lrz**

**SuperMUC-NG**  
Lenovo Intel (2019)  
**311,040 cores**  
Intel Xeon Skylake  
**26.9 PetaFlops** Peak  
**19.5 PetaFlops** Linpack\*  
**719 TeraByte** Main Memory  
**70 PetaByte** Disk

11

11



**SUPERMUC NG PHASE 2**

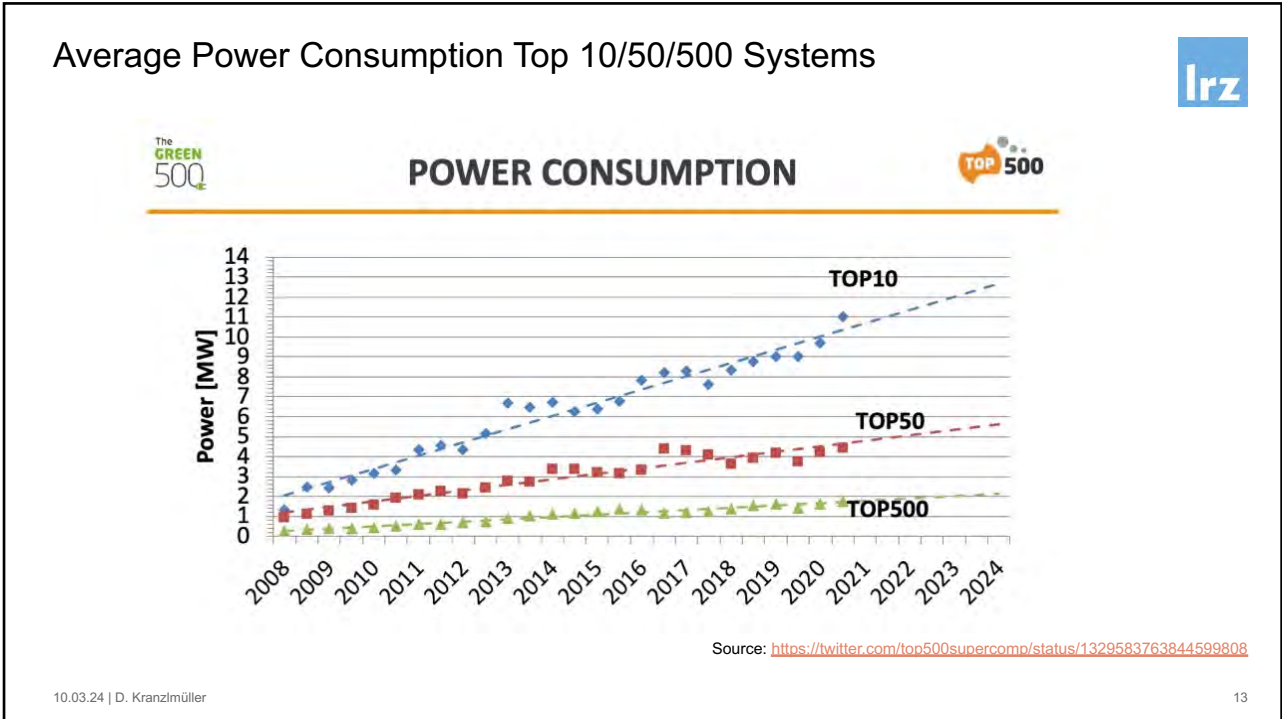
**HPC MEETS AI**

**LRZ@GCS**

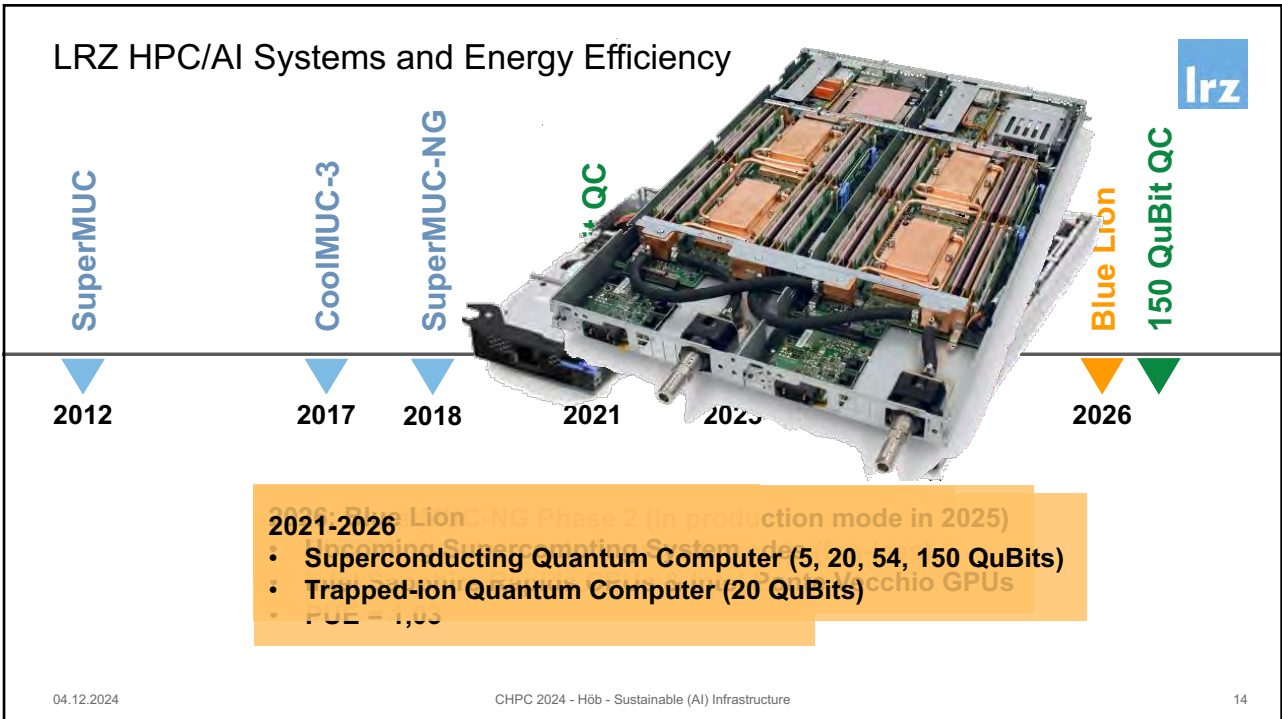
**SuperMUC-NG Phase 2**  
Lenovo Intel  
**240**  
Direct hot water cooled compute nodes (CPU:Sapphire Rapids + GPU:Ponte Vecchio)  
**SD650-I v3**  
Lenovo platform  
**1 PB**  
DAOS storage system mit Intel Xeon Scalable CPUs and 3rd Gen Intel® Optane™ persistent memory

12

12



13



14

### Energy Efficiency at LRZ The Four Pillar Model

**Neighboring Buildings**

**Data Center**

- Improve PUE (Power Usage Effectiveness)
- Reduce Hardware Power Consumption
- Optimize Resource Usage, Tune System
- Optimize Application Performance

**Holistic Optimization Strategy**

**Building Infrastructures**

**HPC System Hardware**

**HPC System Software**

**HPC Applications**

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>

10.03.24 | D. Kranzlmüller 15

15

### Lessons learned towards Sustainability

**SuperMUC** (2012)

**CoolMUC-3** (2017)

**SuperMUC-NG** (2018)

**5 QuBit QC** (2021)

**terrabyte** (2023)

**SuperMUC-NG Phase 2** (2024)

**20 QuBit QC** (2024)

**Blue Lion** (2026)

**150 QuBit QC** (2026)

- DVFS - Dynamic Voltage Frequency Scaling (on average: 68% of nominal frequency)
- Control of power capping / sleep modes
- 100% renewable energies by contract
- Cooling infrastructure optimizations - Glycol free cooling
- Waste heat usage
- Adsorption cooling machines
- Include job-power usage in job protocols
- Future computing technologies – Accelerators, Quantum, ...

04.12.2024 CHPC 2024 - Hüb - Sustainable (AI) Infrastructure 16

16



What about ChatGPT:  <https://chat.openai.com/>



Question: **Is ChatGPT energy efficient?**

Answer:

- ChatGPT's energy efficiency depends on factors like **hardware**, **model size**, and **deployment**.
- Large models can be **energy-intensive**.
- Researchers are working on techniques to improve efficiency.
- The actual efficiency varies based on **implementation** and **infrastructure**.
- Improvements are expected over time.

Question: **What is the energy consumption for training ChatGPT?**

Answer:

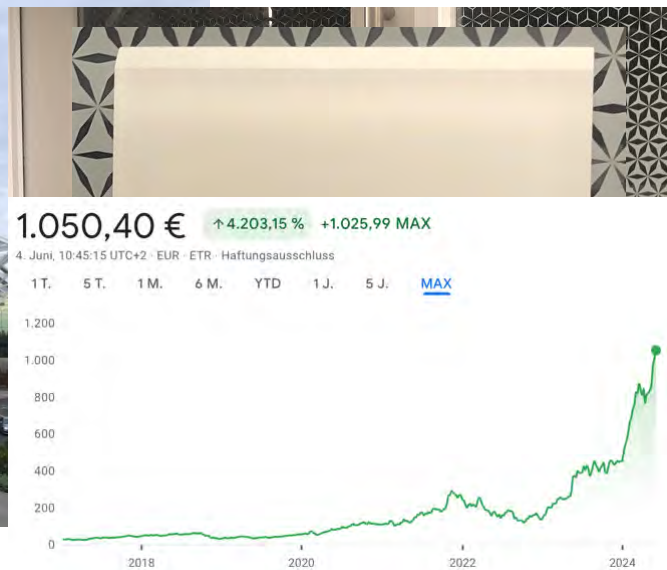
- To provide a rough estimate, training GPT-3, which has 175 billion parameters, could consume hundreds of megawatt-hours (MWh) of electricity during the training process.

10.03.24 | D. Kranzlmüller

17

17

NVIDIA Corporation HQ, Santa Clara, California



10.03.24 | D. Kranzlmüller

19

[www.top500.org](http://www.top500.org)

## Top 500 Supercomputers

- List: June 2024
- Vendor: NVIDIA
- Top 5 out of 22 entries

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
10	<b>Eos NVIDIA DGX SuperPOD</b> - NVIDIA DGX H100, Xeon Platinum 8480C 56C 3.8GHz, NVIDIA H100, Infiniband NDR400, Nvidia NVIDIA Corporation United States	485,888	121.40	188.65	
15	<b>Selene</b> - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	555,520	63.46	79.22	2,646
25	<b>Sejong</b> - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Infiniband HDR, Nvidia NAVER Corp South Korea	277,760	32.97	40.77	
34	<b>Yep1</b> - NVIDIA DGX H100, Xeon Platinum 8480C 56C 3.8GHz, NVIDIA H100 5XM5 80 GB, Infiniband NDR400, Nvidia Ahrefs US DC United States	73,584	23.32	34.58	
35	<b>BioHive-2</b> - NVIDIA DGX H100, Xeon Platinum 8480C 56C 2GHz, NVIDIA H100, Infiniband NDR400, Nvidia Recursion United States	40,320	23.32	34.22	

10.03.24 | D. Kranzlmüller

20

20



10.03.24 | D. Kranzlmüller

21

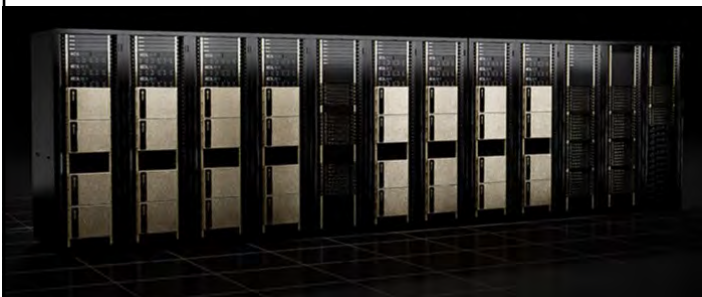
21

## Top 500 Supercomputers for HPC and Artificial Intelligence



### NVIDIA Technology:

- DGX A100  
<https://www.nvidia.com/en-us/data-center/dgx-a100/>
- NVIDIA DGX H100  
<https://www.nvidia.com/en-us/data-center/dgx-h100/>
- NVIDIA DGX SuperPod  
<https://www.nvidia.com/en-us/data-center/dgx-superpod/>



- All current systems are air cooled
- Water cooling is only a supplement
- Overhead for air cooling:
  - +30% Fans
  - +30% Cooling Infrastructure

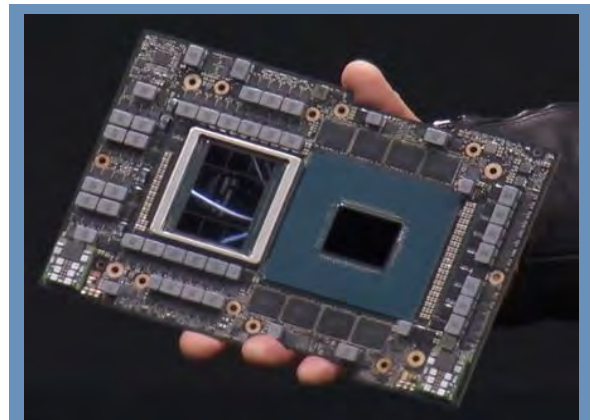
22

22

## NVIDIA H100 Tensor-Core-GPU



- Max. energy consumption (without cooling):  
700 W → per year: **3741 kWh**  
(at 61% utilization)
- Compare: average household in Germany per year (2020):  
approx. **3200 kWh**
- NVIDIA estimates sales for 2024 of approx. 2 Mio. H100
- Total power: 1.4 GW (without cooling)



NVIDIA Grace Hopper Chip  
<https://www.anandtech.com/show/18877/nvidia-grace-hopper-has-entered-full-production-announcing-dgx-gh200-ai-supercomputer>


10.03.24 | D. Kranzlmüller

2  
3

23

Example AI 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**

**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

10.03.24 | D. Kranzlmüller
24

24

## Air vs. Water

Name	Unit	Air	Water	Factor
Thermal Conductivity	W/(mK)	0.026	0.598	23x
Heat Capacity	J/(gK)	1.006	4.185	4x
Volumetric Heat Capacity	kJ/(m <sup>3</sup> K)	1.196	4178	3493x
Thermal Inertia	J/(m <sup>2</sup> Ks <sup>1/2</sup> )	5.563	1581	284x



10.03.24 | D. Kranzlmüller
25

25

### Hot water cooling in the compute room

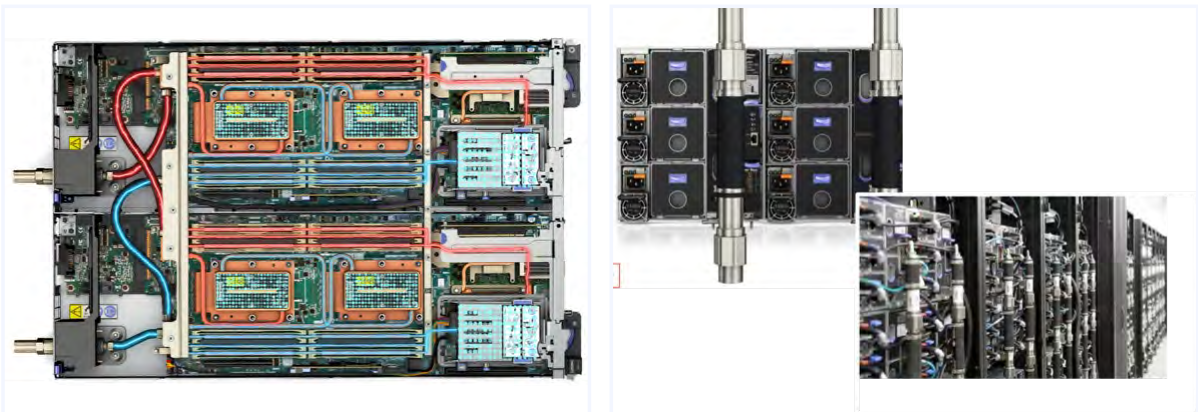


10.03.24 | D. Kranzlmüller

26

26

### SuperMUC-NG Node – with hot water cooling



Source: Photos Lenovo – Produkt: <https://www.lenovo.com/de/de/data-center/servers/high-density/ThinkSystem-SD650/p/77XX7DSSD65>  
10.03.24 | D. Kranzlmüller

27

27

## SuperMUC-NG Cool Manager



3.000 l/h pro 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%)



10.03.24 | D. Kranzlmüller

28

28

Example AI 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:**

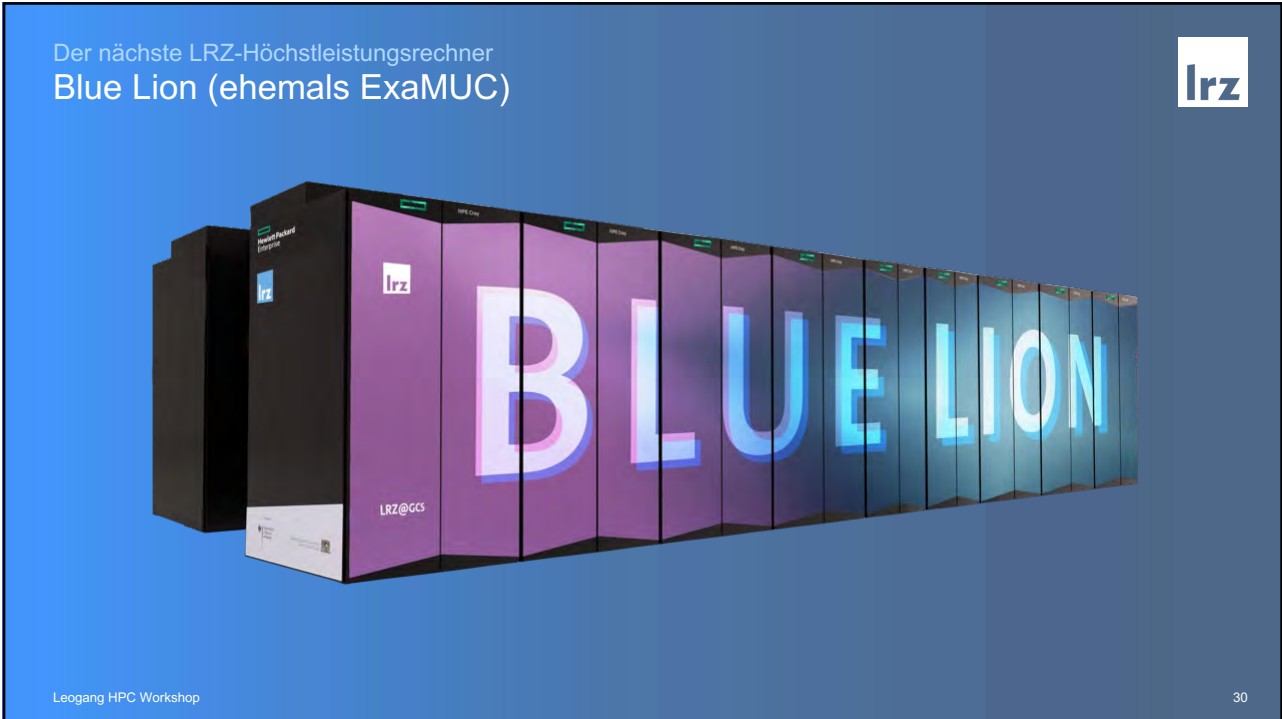
- Up to 144 GPUs per Rack
- Lenovo HGX-based nodes  
(4x A100)
- Direct hot-water cooled  
(allows free cooling)
  - PUE: 1,03-1,05



10.03.24 | D. Kranzlmüller

29

29



30



31

## Savings potentials through direct water cooling

Infrastructure		IT	
<b>Investment</b>	↓	Fewer chillers requires	↗ IT slightly more expensive (depending on quantity)
<b>Operations costs</b>	↓	Power consumption for refrigeration is significantly reduced	↓ Power consumption of water-cooled IT is lower (fewer fans)

www.megware.com32

32

### Example

## Google's Water Demand in Luxembourg

- New Google data center in Luxembourg - €1.2 billion
- 12% of Luxembourg's electricity demand
- Cooling water demand: 10 million l/day
- Google refuses to provide information on the grounds that it "could reveal details of the DC and the agreement with the Luxembourg government to competitors"



Foto: Connie Zhou/Google/Zuma <https://www.nature.com/articles/d41586-018-06610-y>

Foto: Reuters/Yves Herman

<https://orf.at/stories/3184229/>

10.03.24 | D. Kranzlmüller33

33



Nature News Feature - <https://www.nature.com/articles/d41586-018-06610-y>

## How to stop data centres from gobbling up the world's electricity

**ENERGY FORECAST**  
Widely cited forecasts suggest that the total electricity demand of information and communications technology (ICT) will accelerate in the 2020s, and that data centres will take a larger slice.

- Networks (wireless and wired)
- Production of ICT
- Consumer devices (televisions, computers, mobile phones)
- Data centres

20.9% of projected electricity demand

**Global electricity demand**

2015  
Best case 2030  
Expected 2030

Other demand

40,000 TWh

**INTERNET EXPLOSION**  
Internet traffic\* is growing exponentially, and reached more than a zettabyte (ZB,  $1 \times 10^{21}$  bytes) in 2017.

1987 2 TB†  
1997 60 PB  
2007 50 EB  
2017 1.1 ZB

\*Traffic to and from data centres.  
†TB, terabyte ( $10^{12}$  bytes); PB, petabyte ( $10^{15}$  bytes); EB, exabyte ( $10^{18}$  bytes).

©nature

The chart above is an 'expected case' projection from Anders Andrae, a specialist in sustainable ICT. In his 'best case' scenario, ICT grows to only 8% of total electricity demand by 2030, rather than to 21%.

10.03.24 | D. Kranzlmüller

34

<https://www.sueddeutsche.de/wirtschaft/microsoft-stargate-ki-openai-1.6509900>

## Microsoft Project „Stargate“

- 5 Steps Project to build AI Data Centers
- 2026 First Data Center jointly with OpenAI
- In Mt. Pleasant, Wisconsin
- Costs: 10 Billion USD
- 2030 Completion of Supercomputer „Stargate“
- Total: 115 Billion USD
- Electricity: 5 GWatt

**Three Mile Island is reopening and selling its power to Microsoft**

<https://edition.cnn.com/2024/09/20/energy/three-mile-island-microsoft-ai/index.html>

By Jordan Valinsky, CNN  
3 minute read | Updated 12:57 PM EDT, Fri September 20, 2024

Foto von [Jeswin Thomas](#) auf [Unsplash](#)

10.03.24 | D. Kranzlmüller

35

Zusammenhang QC – Integrierte Supercomputer Architektur

LRZ Quantencomputer Q-Exa – Deutschlands 1. Quantendemonstrator



10.03.24 | D. Kranzmüller

36

36

LRZ Quantencomputer von Alpine Quantum Technologies (AQT) auf Basis von Ionenfallen



Photo: MQV – Munich Quantum Valley

37



- Section 5 - Energy efficiency in data centers
- § 23 Energy efficiency and waste heat requirements for data centers
- (1) Data centers that begin operations on or after January 1, 2025, must, for the first two years from the start of operations
  - a planned power usage effectiveness (PUE) **smaller or equal to 1,3**
  - a planned percentage of **reused energy of at least 30 percent**; by January 1, 2027, of at least 40 percent.
- ...
- (3) For data centers that begin operations on or after January 1, 2024, the **minimum inlet temperature** for air cooling of information technology is **27 degrees Celsius**; a lower inlet temperature is only permitted if it can be achieved without the use of a refrigeration system.

## How sustainable is AI? Kate Crawford, „Atlas of AI“

- In addition to the electricity requirement:
  - **Rare earths** for chip production
  - **Water** for cooling
  - ...
- From Large Language Models to „**Smart**“ **Language Models**:
  - *Sepp Hochreiter*, JKU Linz, founder of NXAI, „*Möchte Open AI vom Markt fegen*“ <https://industriemagazin.at/news/sepp-hochreiter-ich-moechte-open-ai-vom-markt-fegen/>
  - *Stephan Günemann*, Director MDSI at TUM, founder of Pruna AI „*Smash your AI model*“ - <https://www.pruna.ai/>
  - *Björn Ommer*, LMU and Stable Diffusion, „*Kämpft gegen den Größenwahn der Branche*“ - <https://www.sueddeutsche.de/wirtschaft/kuenstliche-intelligenz-ki-neurips-konferenz-bjoern-ommer-1.6318009>

