



Leibniz Supercomputing Centre  
of the Bavarian Academy of Sciences and Humanities

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.

**Powering the Future:  
Tackling Energy Challenges in  
Supercomputing and AI**

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

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

28.11.24 | D. Kranzlmüller

## Examples

How much electrical power do we need for the digital transformation?



- **Smartphone:** iPhone 8

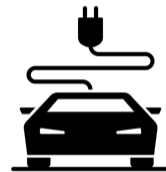
- Battery Capacity 1821 mAh
- Per charge (with 0,30 € per kWh) → approx. 0,0027 € per day
- With one charge per day → approx. 0,9855 € per year



Created by Thibin Matheir from Noun Project

- Personal computer:

- Current Intel/AMD multicore processor with dedicated GPU
- Per day (with 4 hours running) → 200 kWh per year
- approx. 60 € per year



Created by Symbolon from Noun Project

- Electric vehicle : Tesla

- Model S 60, 18.1 kWh/100 km → 5,43 € per 100 km
- Annual mileage approx. 15.000 km → 814,50 € per year

28.11.24 | D. Kranzlmüller

3

3

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

Partner for Digital Transformation of Science & Research

BA&W

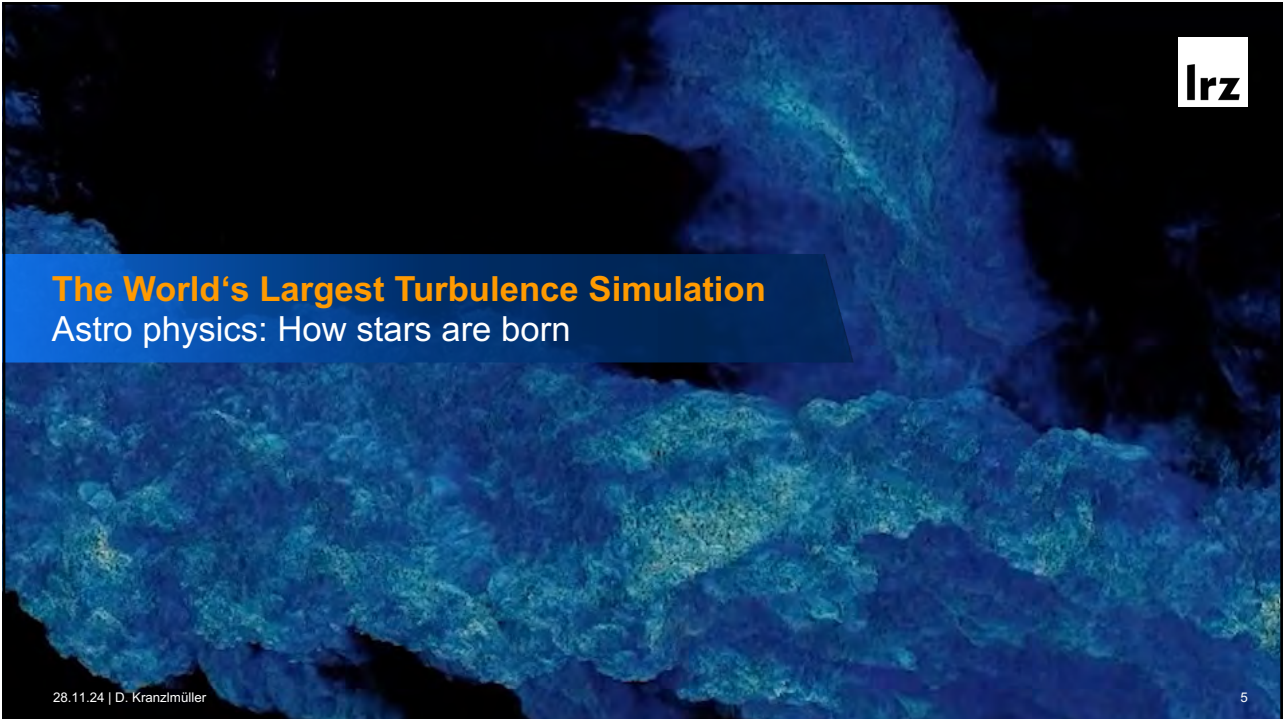
lrz

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

28.11.24 | D. Kranzlmüller



4

4

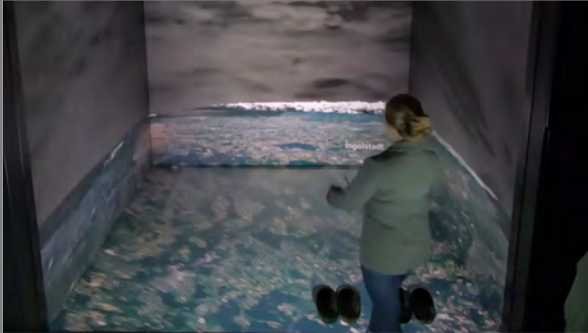


5


### Bavarian-Canadian Project ClimEx



Coupling HPC models for climate and hydrology, workflow and data management.<sup>1</sup>



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



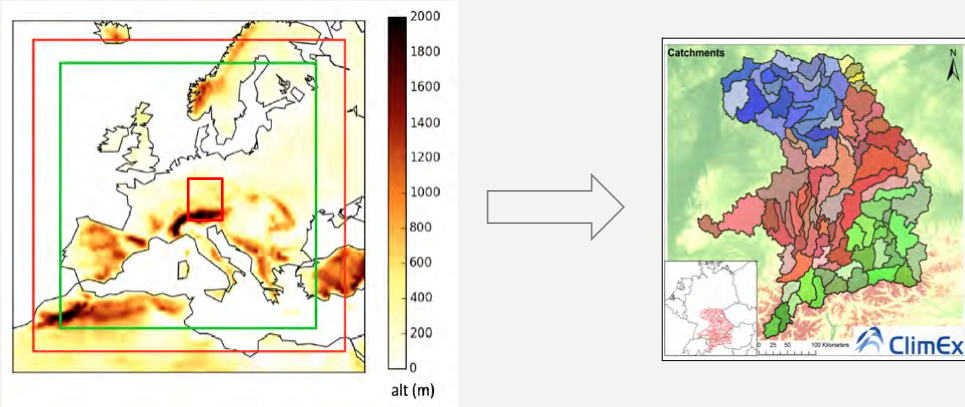
www.climex-project.org; 1 © LRZ / Ouranos; 2 Danube Flooding 2013 © sueddeutsche.de

6

6



# 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

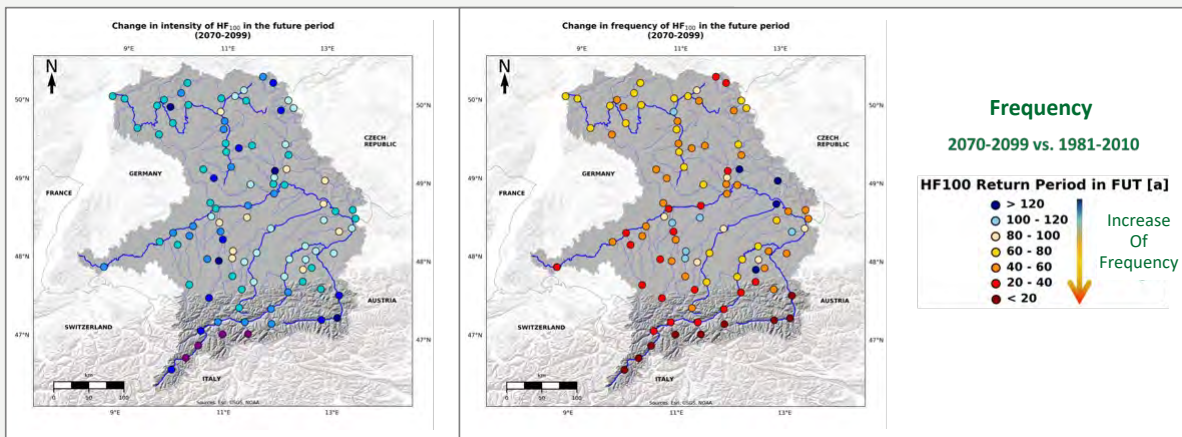
Slide courtesy Ralf Ludwig, LMU

7

# 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.)

Slide courtesy Ralf Ludwig, LMU

14

8

## Background...

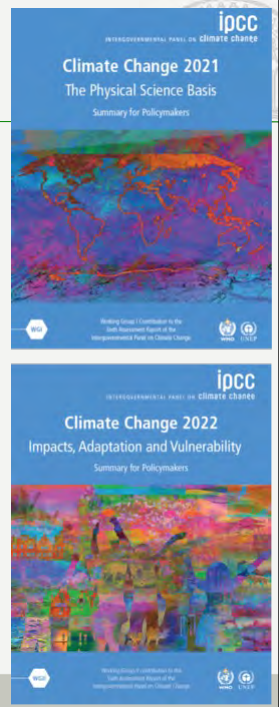
**From the 6th Assessment Report of the IPCC (AR6) (2021/2022):**  
**Working Group I: The Physical Science Basis**  
**Working Group II: Impacts, Adaptation and Vulnerability**

### **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





#	Site	Manufacturer	Computer	Country	Cores	Power (MW)	Power (Watt)
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 Daini 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	2.58
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 36C 3.3GHz, Intel Omni-Path	Germany	305,856	19.5	1.65
9	Oak Ridge National Laboratory	Cray	Titan Cray XK7, Xeon E5 12C 2.8GHz, Arles, NVIDIA Tesla K20	USA	560,640	17.6	8.2
10	Lawrence Livermore National Laboratory	IBM	Sierra IBM Power System, P9 22C 3.1GHz, Mellanox EDR, NVIDIA GV100	USA	1,572,864	17.2	7.8

11

**~9,6 Billion**  
Core Hours

**~2,4 Mio.**  
Jobs

**~515**  
Projects

**~1,400**  
Scientists

**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

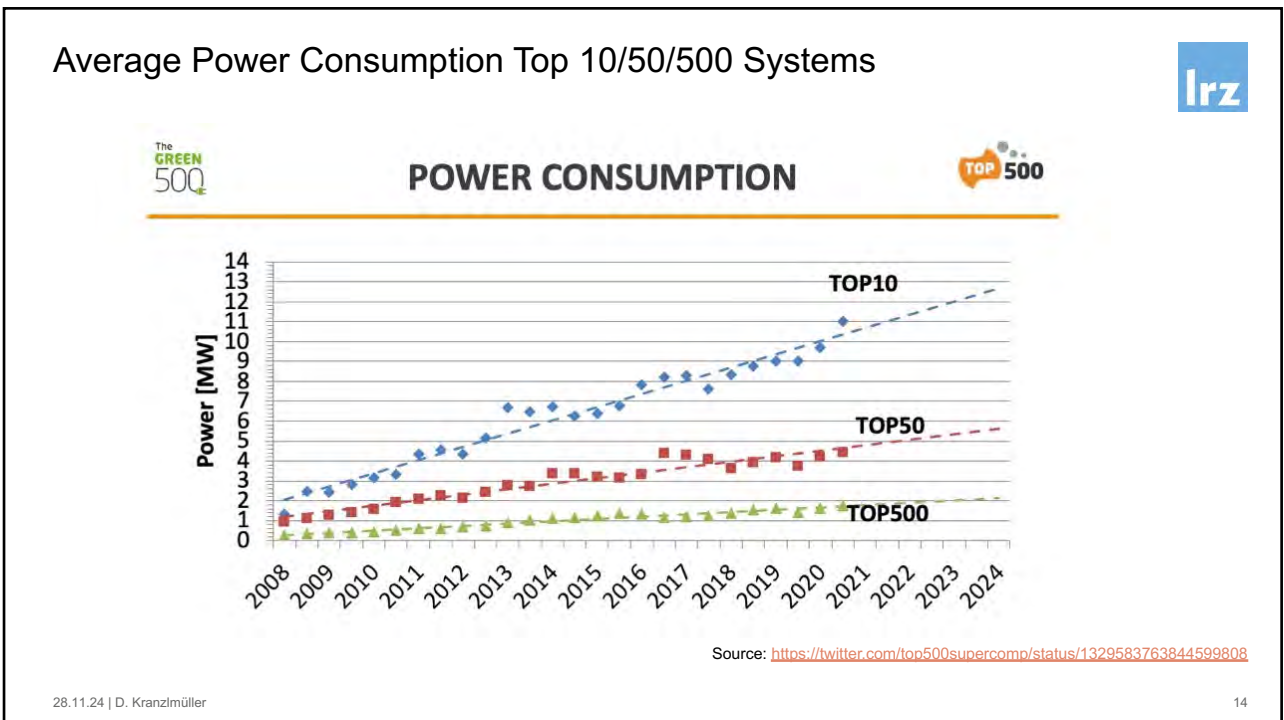
28.11.24 | D. Kranzmüller 12

12



**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

13



14

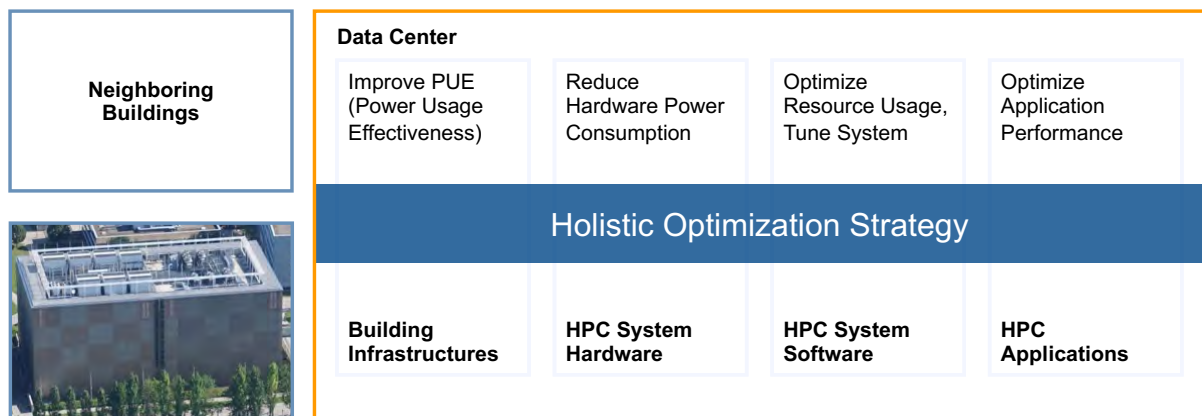
## LRZ HPC Systems and Energy Efficiency



- 2012: SuperMUC – first hot water cooled system (CPU+Memory)
- 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
- 2026/27: ExaMUC - ??? (procurement closed in October 2024)

15

## Energy Efficiency at LRZ 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>

16



## 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
- 2026/27: ExaMUC - ??? (procurement closed in October 2024)
- 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, ...

28.11.24 | D. Kranzlmüller

17

17

## Energy Efficiency at LRZ



Pioneer in  
direct-liquid  
hot-water cooling

Heat reuse  
(heating, adsorption  
chillers)

Energy-aware  
scheduling

Broad  
instrumentation  
of all components

DCDB:  
High frequency, high  
resolution monitoring  
and analysis

28.11.24 | D. Kranzlmüller

18

18

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.*

19

Some numbers from research online - Estimates



**ChatGPT-3 Energy consumption:**

- to train the GPT-3 model: 1064 MWh
- To operate search requests: 260 MWh/day
- A ChatGPT query: 6,79 Wh
- <https://medium.com/@zodhyatech/how-much-energy-does-chatgpt-consume-4cba1a7aef85#>

**Comparison:**

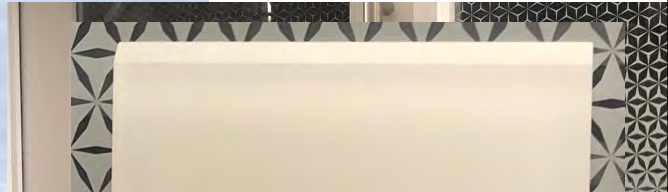
- With 1 MWh, an electric car can travel about 5000 km.
- The average per capita electricity consumption in Germany is about 1.5 MWh/a, per household about 4 MWh/a.
- 260 MWh/d correspond to the power consumption of a small town with 24000 households or 64000 inhabitants.

20

## NVIDIA Corporation HQ, Santa Clara, California



28.11.24 | D. Kranzlmüller



21

[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	Eos NVIDIA DGX SuperPOD - 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	Selene - 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	Sejong - 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	Yep1 - 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	BioHive-2 - NVIDIA DGX H100, Xeon Platinum 8480C 56C 2GHz, NVIDIA H100, Infiniband NDR400, Nvidia Recursion United States	40,320	23.32	34.22	

28.11.24 | D. Kranzlmüller

22

22





23

### 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

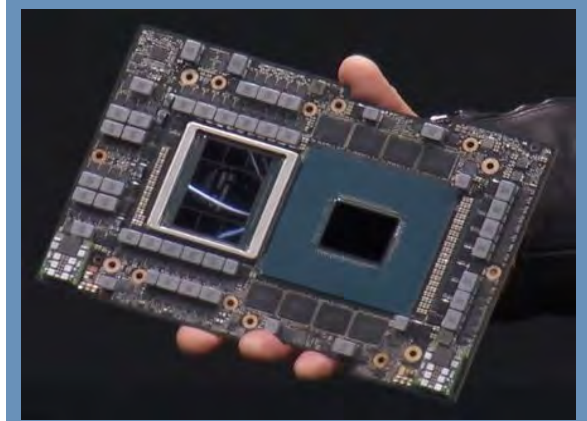
24

24

## 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>

28.11.24 | D. Kranzlmüller

2  
5

25

### 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

28.11.24 | D. Kranzlmüller

26

26

## 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

28.11.24 | D. Kranzlmüller 27

27

## Hot water cooling in the compute room





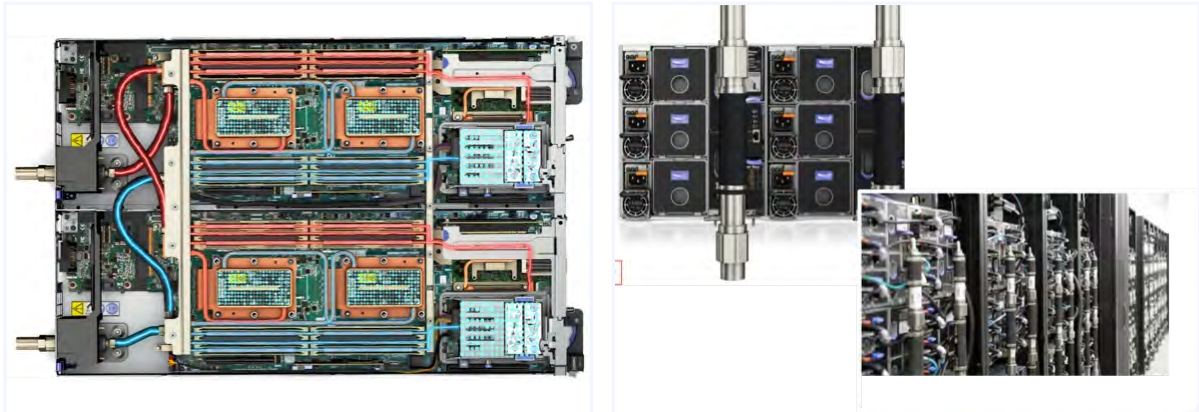
28.11.24 | D. Kranzlmüller 28

28



## SuperMUC-NG Node – with hot water cooling

lrz



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

29

29

## SuperMUC-NG Cool Manager

lrz

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%)



28.11.24 | D. Kranzlmüller

30

30

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

28.11.24 | D. Kranzlmüller

31

31



32



## 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.com33

33

### 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"

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

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




Foto: Reuters/Yves Herman

28.11.24 | D. Kranzlmüller34

34



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

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%.

28.11.24 | D. Kranzlmüller

**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

35

<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


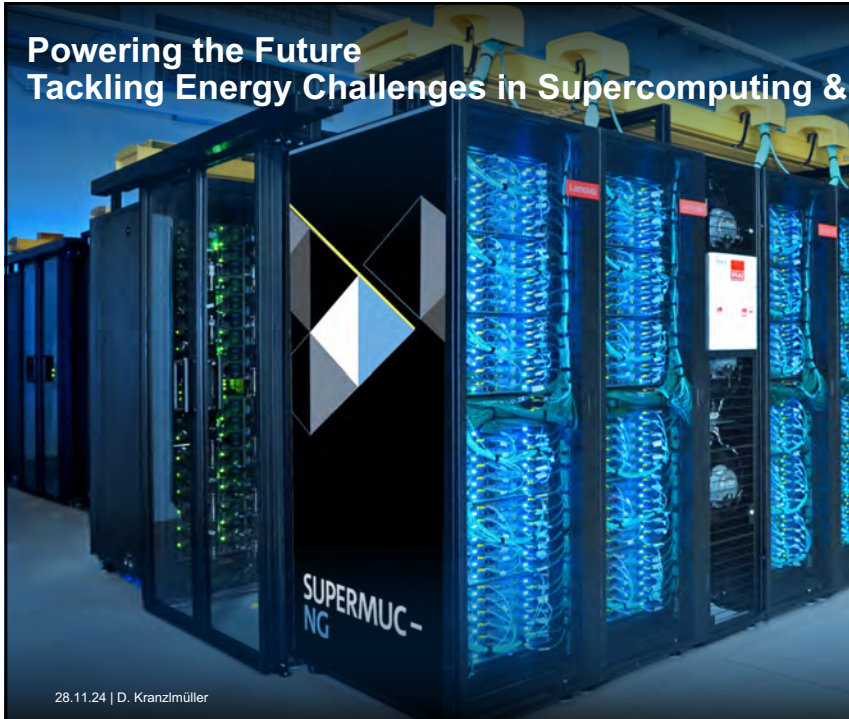
Foto von [Jeswin Thomas](#) auf [Unsplash](#)

28.11.24 | D. Kranzlmüller

36

36

### Powering the Future Tackling Energy Challenges in Supercomputing & AI






- **Energy efficiency** is essential for sustainability
- We need a **holistic approach** - more than GFlop/watt
- Aside from the **carbon footprint**, there is an important **financial aspect**

Contact: Prof. Dr. Dieter Kranzlmüller  
kranzlmuller@lrz.de | @Kranzlmuller

28.11.24 | D. Kranzlmüller

37

### Zusammenhang QC – Integrierte Supercomputer Architektur LRZ Quantencomputer Q-Exa – Deutschlands 1. Quantendemonstrator



28.11.24 | D. Kranzlmüller

38

38



39

German Federal Energy Efficiency Act - Increasing Energy Efficiency in Germany  
 Proposed German government bill **EnEfG 2023**



- 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.

28.11.24 | D. Kranzlmüller

40

40



[https://www.hugendubel.de/de/buch\\_kartoniert/kate\\_crawford-atlas\\_of\\_ai-41511599-produkt-details.html](https://www.hugendubel.de/de/buch_kartoniert/kate_crawford-atlas_of_ai-41511599-produkt-details.html)

## 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://indusmagazin.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>

