Towards sustainable HPC at the Jülich Supercomputing Centre

CIUK 2022 | December 1, 2022 | Thomas Eickermann

Research Centre Jülich by Numbers

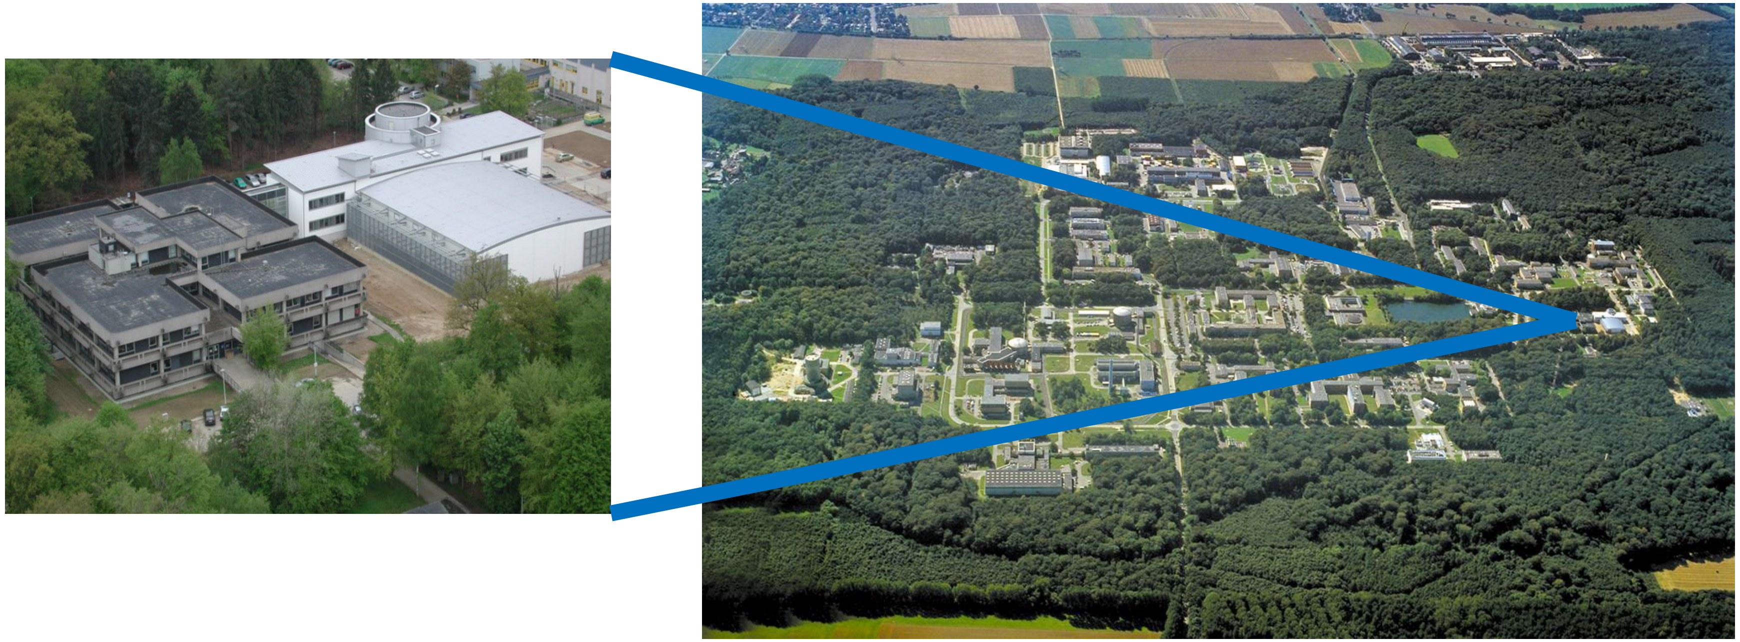

Research areas

- Information
- Energy
- Bioeconomy

- **Budget:** 861 Mio €,
  including 395 Mio € third party funding
  171 Horizon 2020 projects, 420 national projects
- **Employees:** 7.120
  incl. 2.626 scientists including PhD students
  934 guest scientists from 65 countries
- **Publications:** 3.081
  (source: fact sheet 2021)

Jülich Supercomputing Centre (JSC)

Facts and Figures

Staff:

220 Total (185 FTE)

160 Scientists

13 PhD Students (+13 external)

Budget:

30 Mio. € Institutional Funding (PoF)

15 Mio. € Third Party Funding

Jülich Supercomputing Centre at a Glance

- **Supercomputer operation for**
  - Centre – FZJ
  - Region – RWTH Aachen University
  - Germany – Gauss Centre for Supercomputing (GCS)
    John von Neumann Institute for Computing (NIC)
  - Europe – PRACE, EU projects, EuroHPC
- **Application support**
  - Unique support & research environment at JSC
  - Peer review support and coordination
- **R&D work**
  - Methods and algorithms, computational science, performance analysis and tools
  - Scientific Big Data Analytics with HPC
  - Computer architectures, Co-Design, Exascale Labs together with IBM, Intel, NVIDIA
- **Education and training**

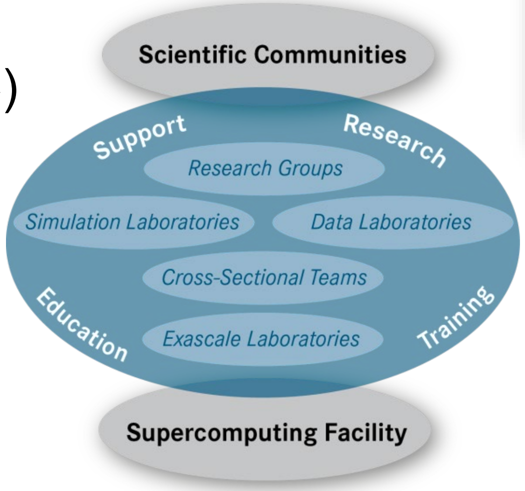

DEEP

Towards Sustainable HPC at JSC

Optimisation of Energy Usage

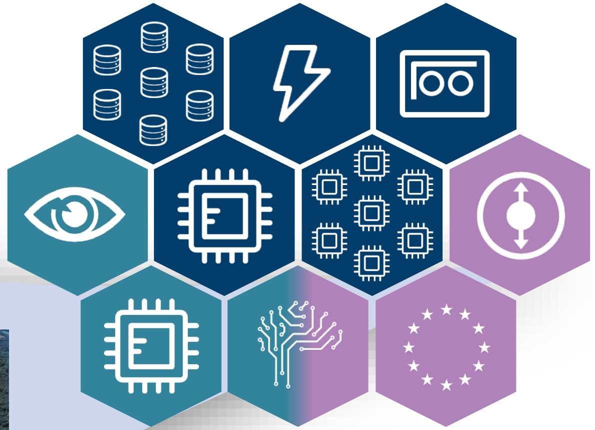

Future Plans

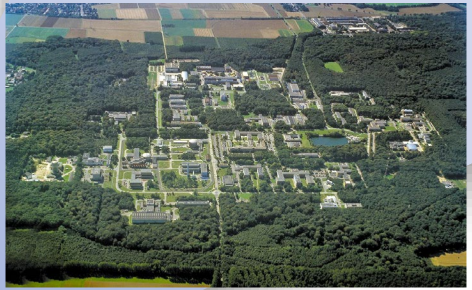

Campus Level

Data Centre Level

System Level

(DUAL) hardware strategy at JSC

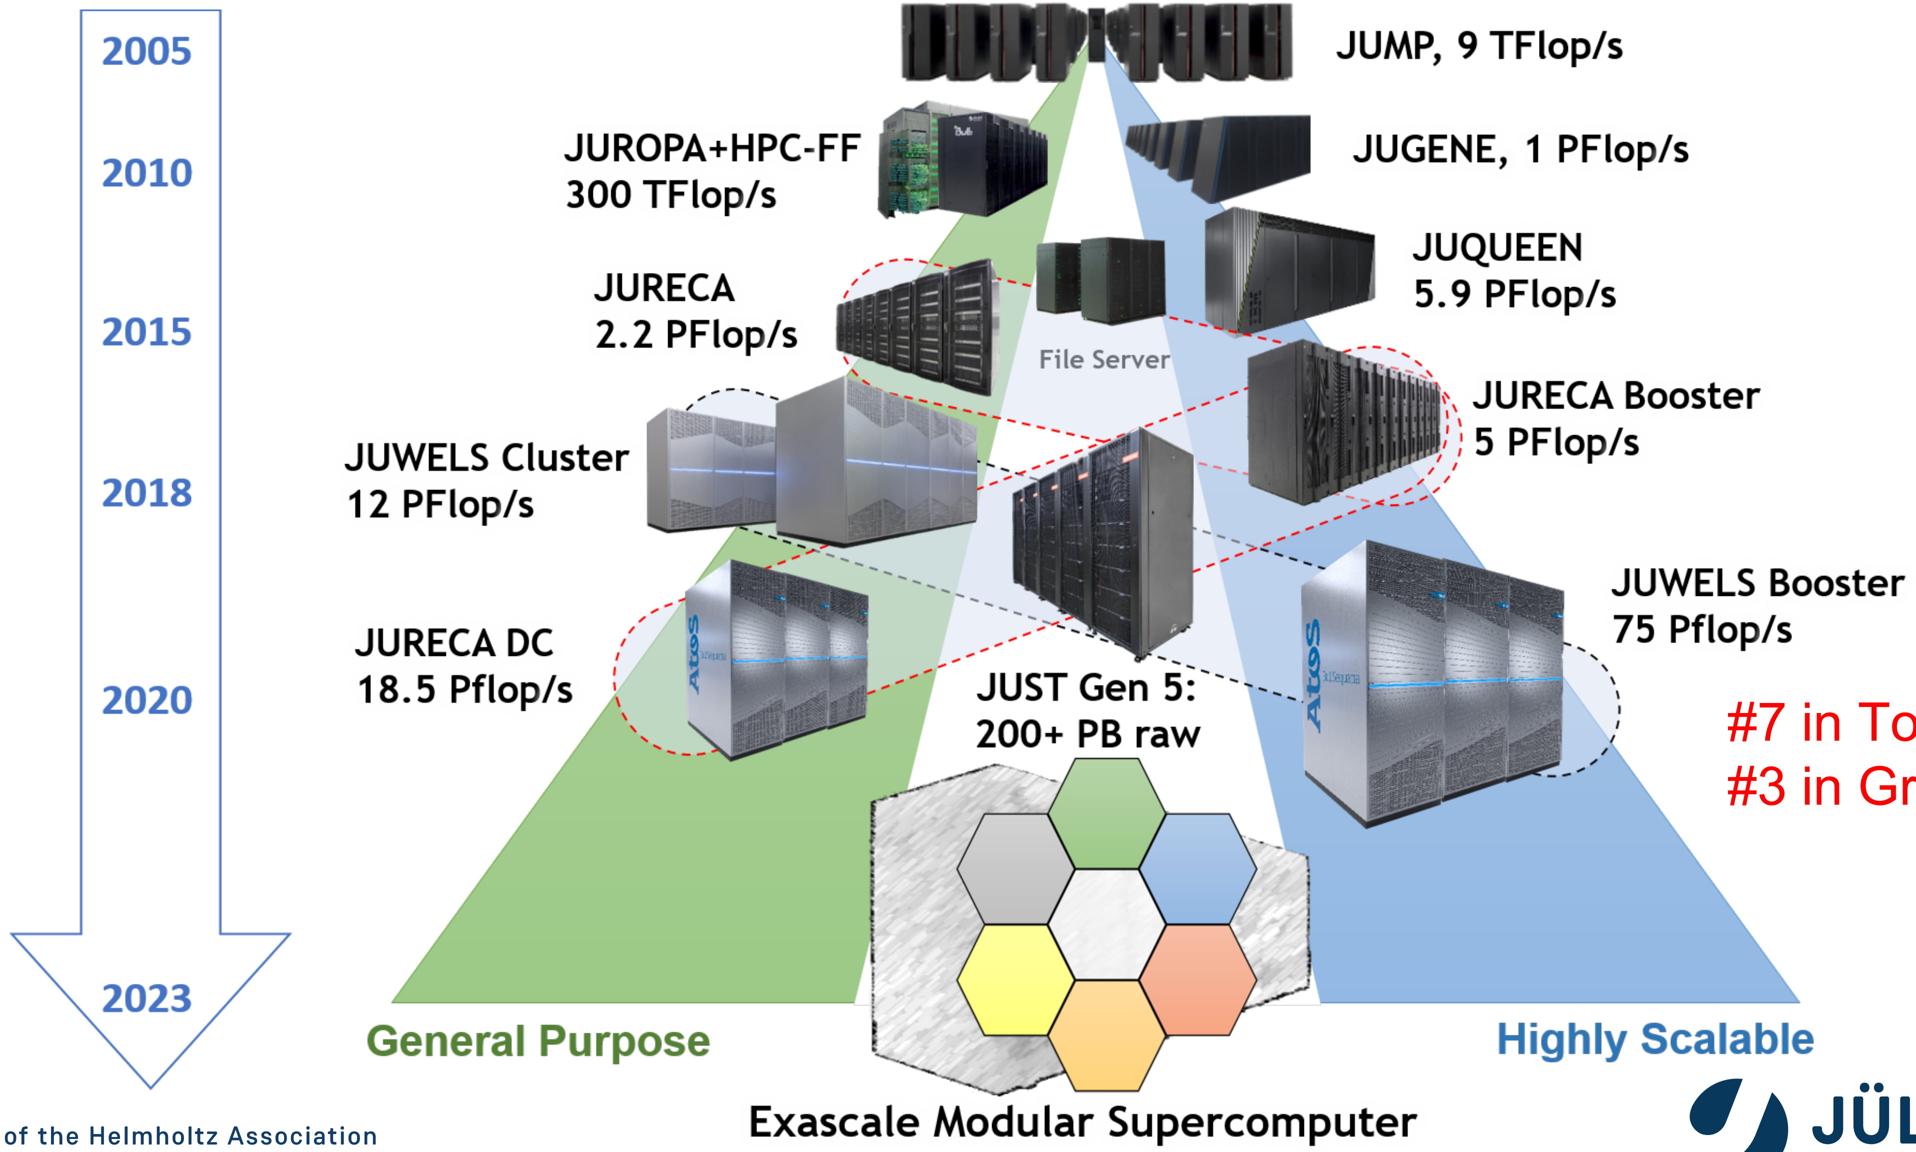

#7 in Top500 11/2020
#3 in Green500 11/2020

System Level

- **Energy-efficient compute nodes**
  - GPU accelerators boost Flops/W

- **Energy-efficient system architectures**
  - Many applications cannot benefit from GPUs (today)
  - Idle GPUs are not energy-efficient
  - Dual hardware strategy: General Purpose + Highly Scalable system for different demands and mixed workflows
  - 35% of JUWELS Booster projects have also allocations on the JUWELS Cluster
  - Modular Supercomputer Architecture: tight integration of heterogeneous resources

Modular supercomputing architecture

Composability of heterogeneous resources

- Cost-effective scaling
- Effective resource-sharing
- Match application diversity
- Large-scale, complex workflows

- **E. Suarez**, N. Eicker, Th. Lippert, "*Modular Supercomputing Architecture: from idea to production*", Chapter 9 in Contemporary High Performance Computing: from Petascale toward Exascale, Volume 3, p 223-251, CRC Press. (2019)
- **E. Suarez**, N. Eicker, and Th. Lippert, "Supercomputer Evolution at JSC", Proceedings of the 2018 NIC Symposium, Vol.49, p.1-12, (2018)

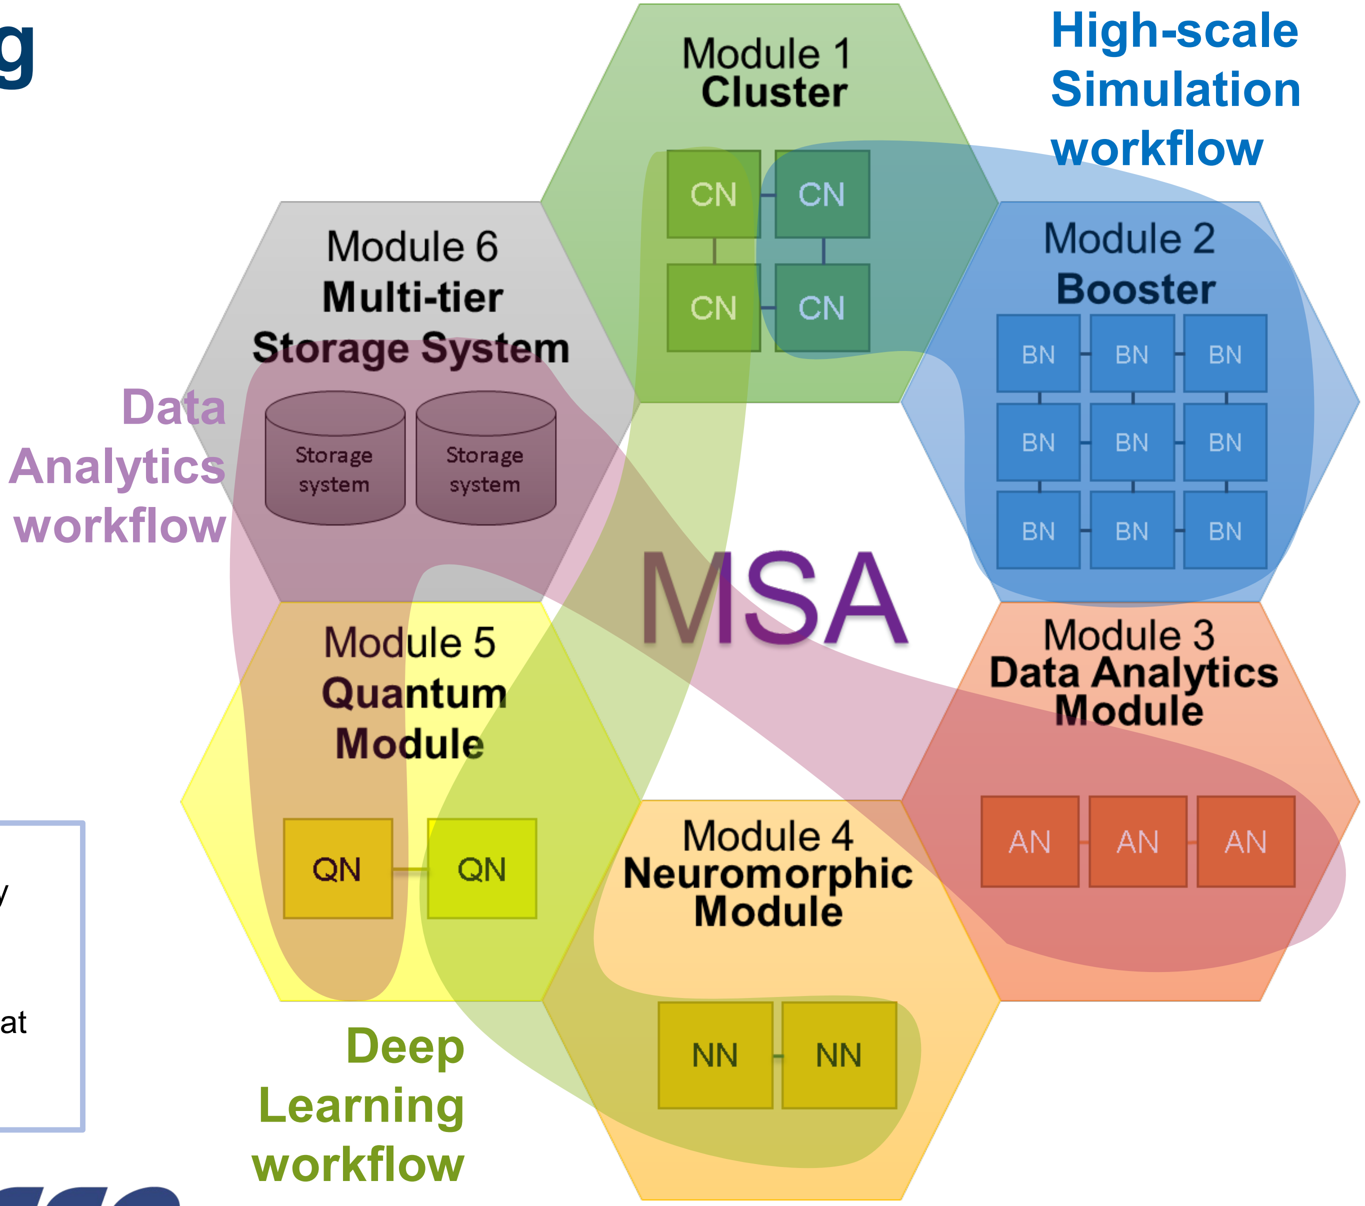

System Operation: Adaptation of GPU / CPU Frequencies

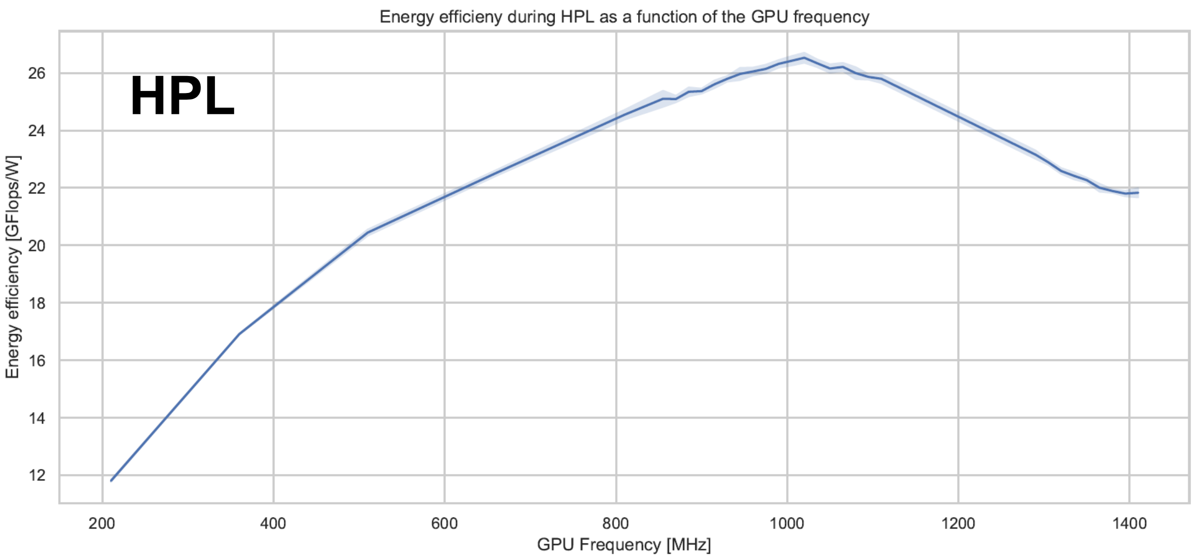

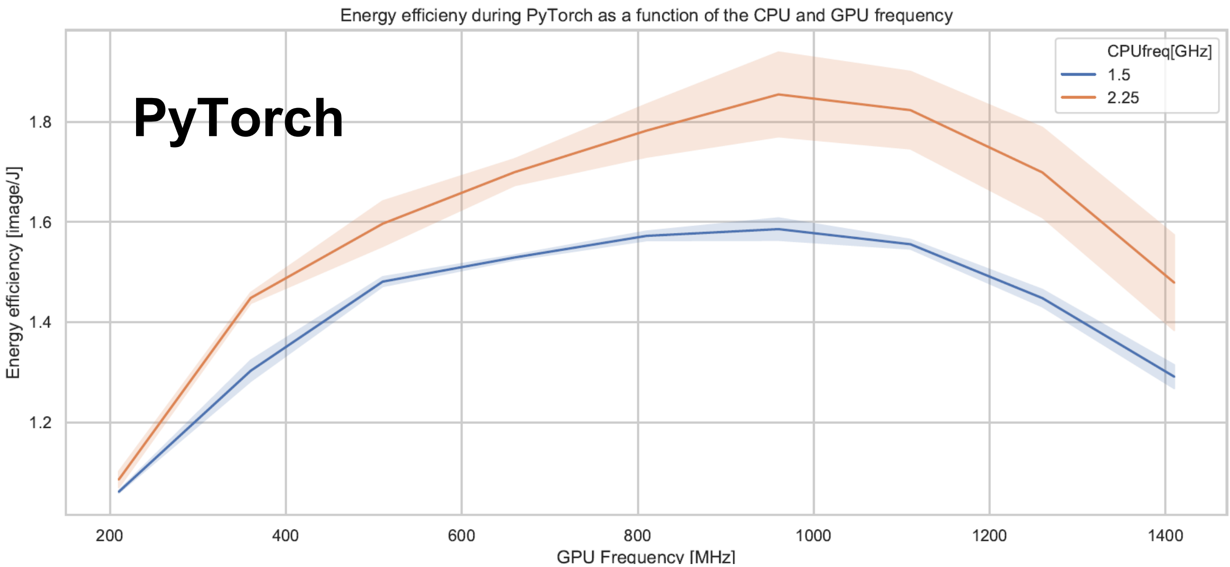

GPU freq. [MHz]	Time-to-solution	Energy to solution	TCO to solution
1410	-	-	-
1095	7.82 %	- 15.59%	- 6,2%
1029	11.33 %	- 17.71%	- 6,1%

GPU freq. [MHz]	Time-to-solution	Energy to solution	TCO to solution
1410	-	-	-
1110	3.32%	18.85%	-10,0%
960	9.01%	20.18%	-8,5%

Measurements on JURECA-DC: 2x AMD EPIC 7742, 4x NVIDIA A100-SMX4-40GB

by Sebastian Achilles (JSC)

System Operation Cont.

- **GPU Frequency adaptation**
  - Extended test opportunities provided for JUWELS users
  - No significant gain in energy-to-solution for many applications, 5-10% for some
- **Powering off idle nodes**
  - JUWELS is fully loaded, but …
    - Scheduling a mix of small and large node-count jobs leads to idle periods
  - Tested and put into production on smaller systems, incl. JURECA-DC
    - Reduced interconnect stability
    - Little impact on user experience

Data-Centre Level – Cooling

- **Until 2022**
  - JUWELS and JURECA-DC use direct liquid cooling
  - Chilled water is centrally supplied for the Jülich campus
  - Coefficient of Performance ~ 2.5
  - Supported by free cooling in winter

- **Since May 2022**
  - 1.8 MW Hybrid warm-water cooling:
    - inlet ~ 34 °C outlet ~ 42 °C
  - Free cooling + water evaporation in hot periods
  - PUE ~ 1.1
  - Extension to 3 MW is underway
  - Chilled water only for air cooled components: storage, network

Research Centre Level – Waste Heat Usage

- **Living Lab Energy Campus**
  - A project to develop and deploy an integrated campus-wide energy management
    - Renewable energy production and storage
    - Monitoring and predicting usage & steering energy production (e.g. gas-fired combined cooling, heating and power (CCHP) plant) and battery usage
- **Under Construction**
  - Low-temperature (~ 40 °C) district heating system powered by JSC waste heat
  - Temperature is sufficient to directly heat buildings fulfilling current German insulation standards
  - Heat pumps used to achieve temperature levels (~ 70 °C) required for older buildings – such as the JSC

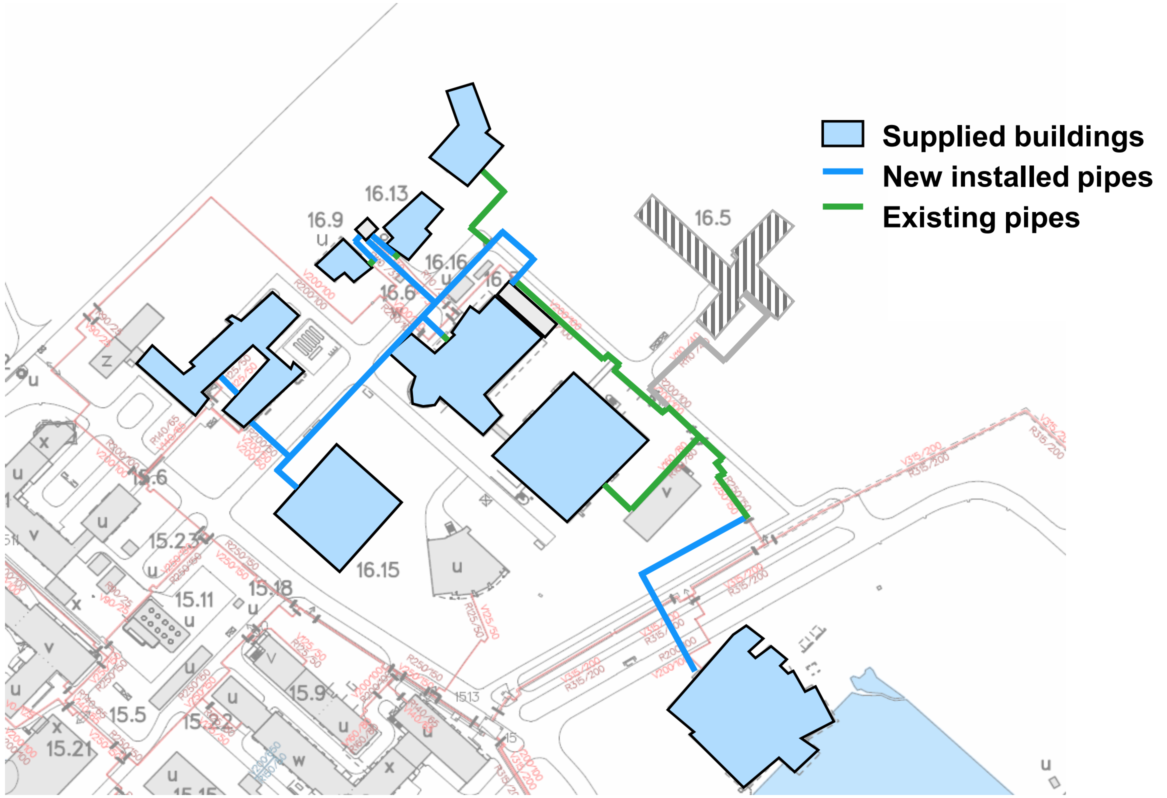

JUPITER – The 1st European Exascale System

EuroHPC Joint Undertaking

SPONSORED BY THE Federal Ministry of Education and Research

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

- **EuroHPC Joint Undertaking**
  - Joint undertaking between EU, member states, private partners
  - Took over funding of HPC related projects from EC
  - Co-funds Petascale, and owns Pre-Exascale, and Exascale systems

- **JUPITER - JU Pioneer for Innovative and Transformative Exascale Research**
  - Selected on June 14, 2022 as the 1st EuroHPC Exascale system
  - Installation in Jülich targeted for end of 2023
  - 500 Mio. € Total Costs, equally shared between EuroHPC and Germany (federal and state of North Rhine-Westfalia funding)

JUPITER – Modular Supercomputer

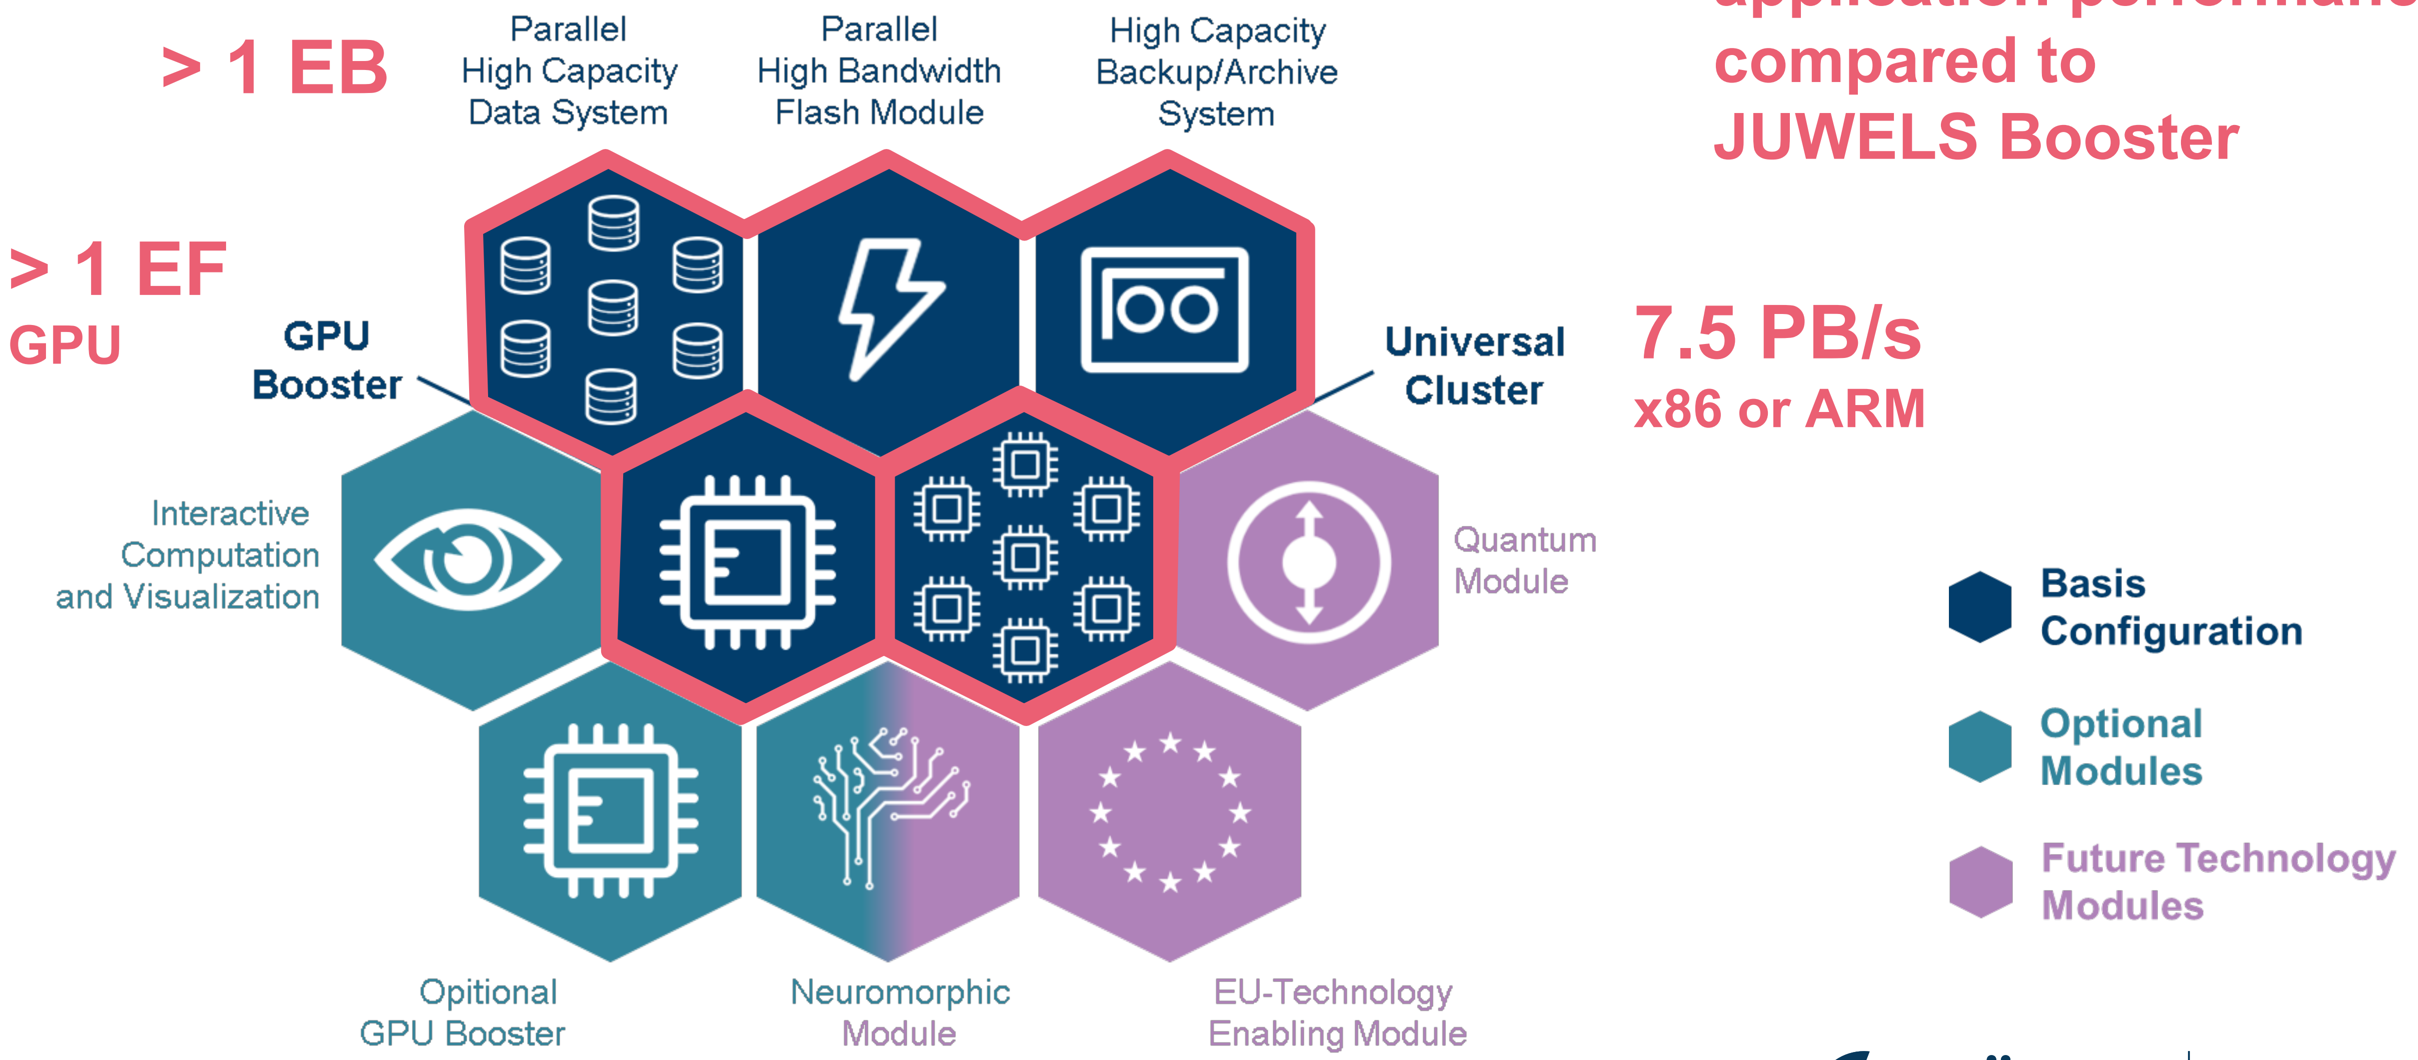

Target >20× application performance compared to JUWELS Booster

JUPITER - Towards Sustainability

- **JUPITER will leverage all of the above:**
  - Modular Supercomputer Architecture
  - GPU-based booster
  - Operated with green electricity
  - Direct warm-water cooling
  - Waste heat usage: funding secured for
    - a high-performance heat pump (> 1 MW)
    - Measures on the campus that enable broader utilization of JUPITER's waste heat
  - Optimisation of energy supply
    - From: 110 kV → 35 kV → 10 kV → 480 V
    - To: 110 kV → 35 kV → 480 V

Location of Exascale Data Centre

SPONSORED BY THE

Federal Ministry of Education and Research

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

Waste Heat Usage – Long-Term Vision

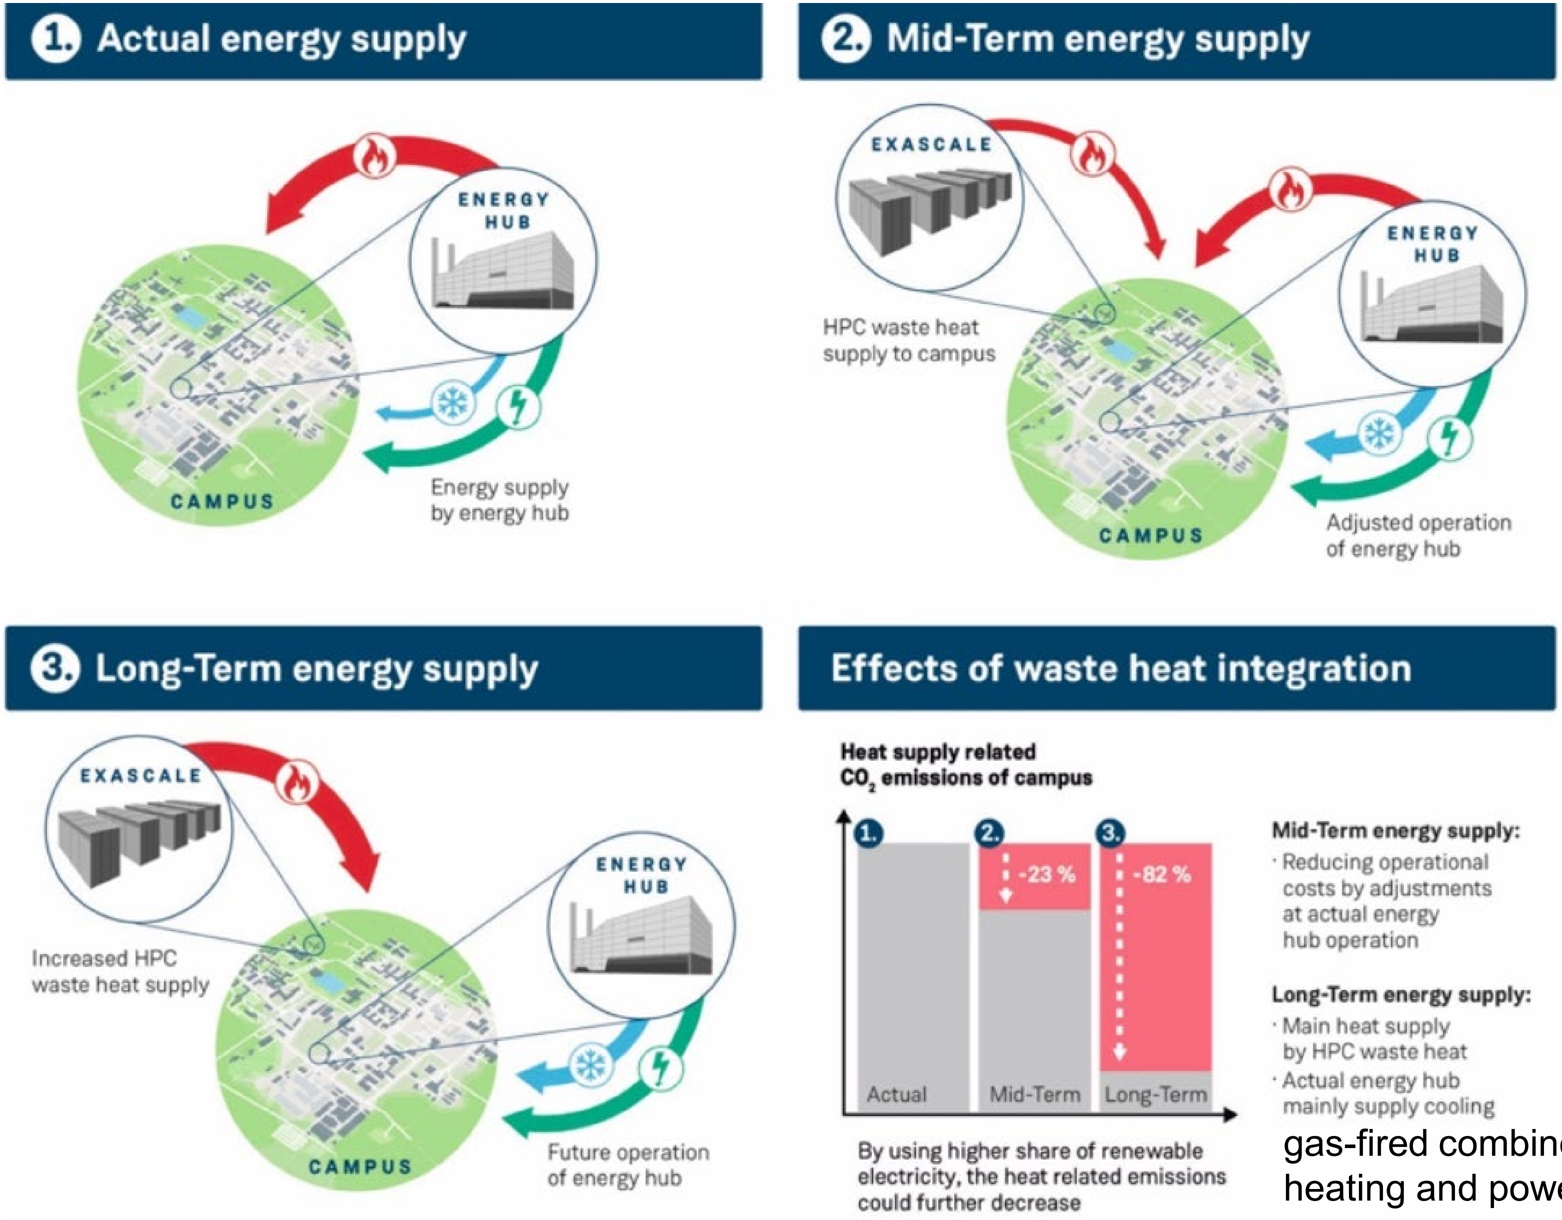

- JUPITER average power ~ 15 MW
- Campus heat demand ~7.5 MW

gas-fired combined cooling, heating and power (CCHP) plant

JÜLICH Forschungszentrum | JÜLICH SUPERCOMPUTING CENTRE

Towards Sustainable HPC at JSC

Optimisation of Energy Usage

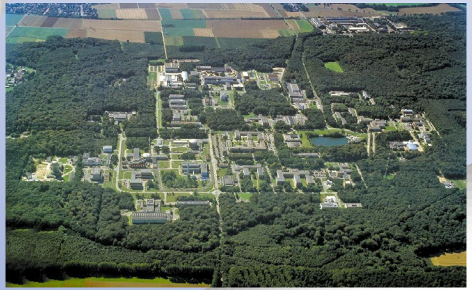

Future Plans
JUPITER

Campus Level:
waste heat usage

Data Centre Level: free cooling

System Level: GPUs, Modular Architecture