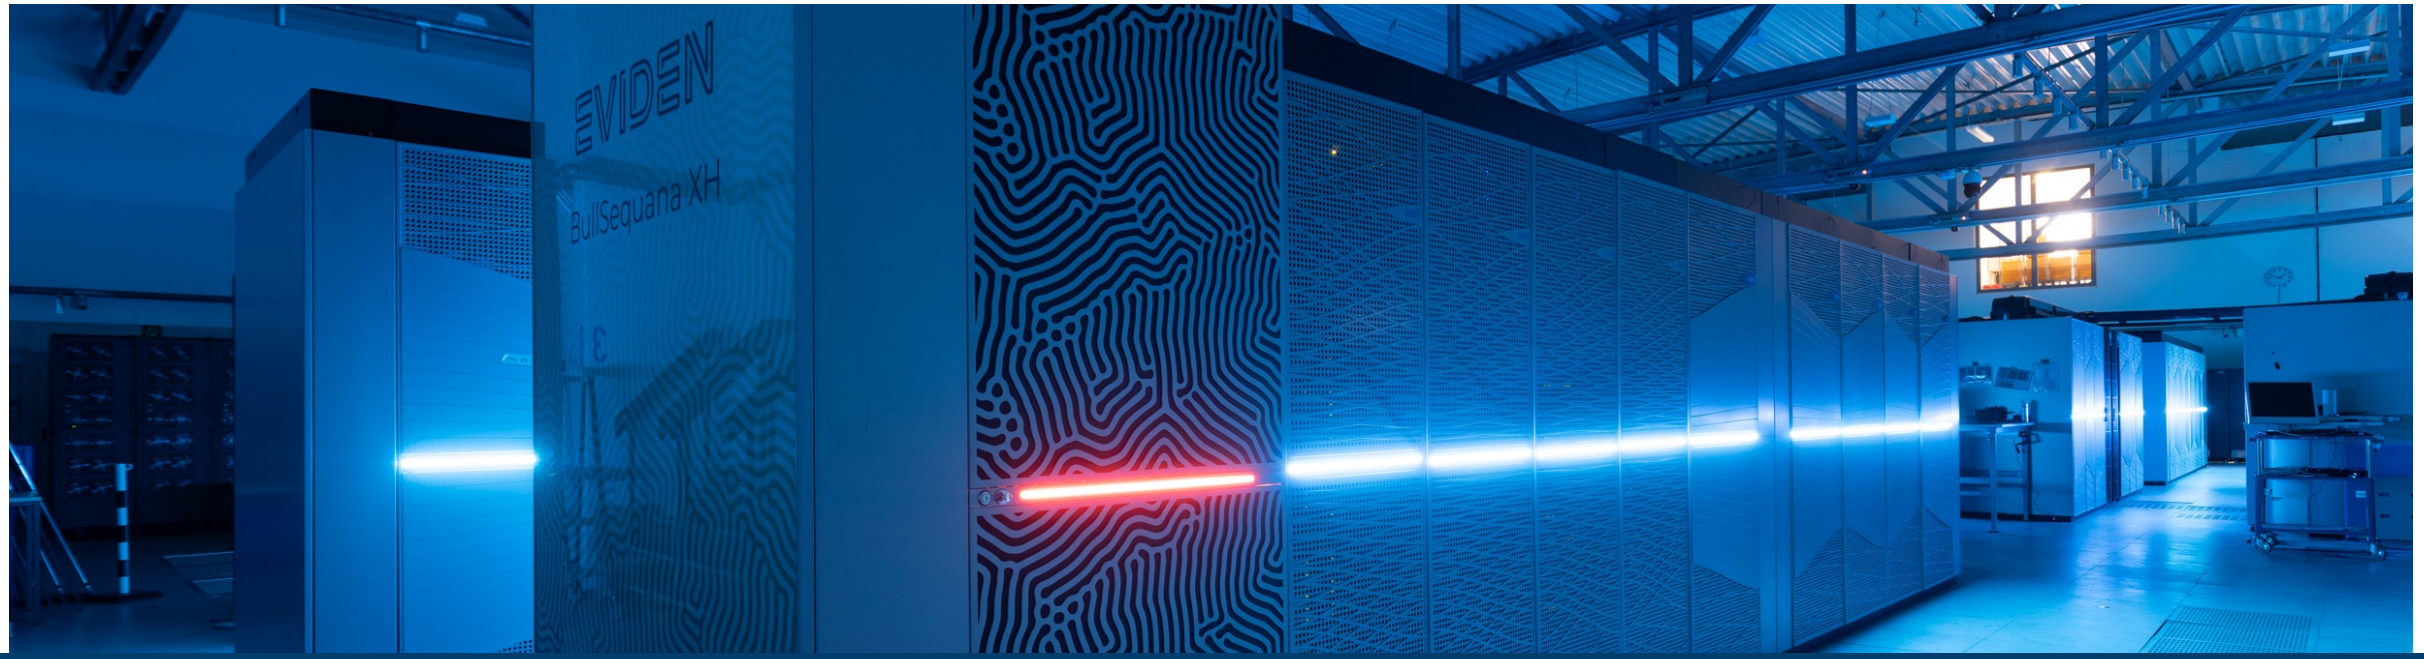

JEDI

The JUPITER Exascale Development Instrument

ISC24 - HAMBURG - GREEN500 BOF

EuroHPC Joint Undertaking

Member of the Helmholtz Association

Bundesministerium für Bildung und Forschung

Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen

GCS Gauss Centre for Supercomputing

JÜLICH Forschungszentrum

Shaping Change

JUPITER EXASCALE DEVELOPMENT INSTRUMENT

EuroHPC / Forschungszentrum Jülich

- Eviden BullSequana XH3000
  - 24x Compute nodes (12x Blades)
    - NVIDIA quad-GH200 96GB Grace Hopper Superchip
    - Memory: 480GB on CPUs + 384GB on GPUs
    - NVIDIA quad-rail InfiniBand NDR200
  - 1x Network switch:
    - NVIDIA Quantum-2 NDR InfiniBand switch
  - All components are Direct Liquid Cooled

EVIDEN BULLSEQUANA XH3000

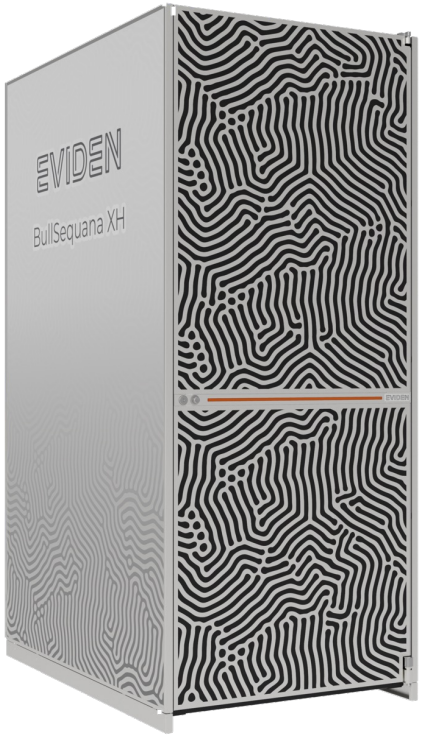

Fully integrated rack
(compute, network, cooling, management)

One blade hosts
two compute nodes

ENVIRONMENT

Performance Baseline

- NVIDIA recipe:
  - NVIDIA hpc-benchmarks 24.03 container
  - NVIDIA HPC-X (with NCCL)
- Identified ‘best’ performing HPL parameters (24 nodes)
  - Especially N, NB and PxQ

RANKING

List	Rank	System	Vendor	Total Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
06/2024	189	BullSequana XH3000, Grace Hopper Superchip 72C 3GHz, NVIDIA GH200 Superchip, Quad-Rail NVIDIA InfiniBand NDR200	ParTec/EVIDEN	19,584	4.50	5.13	67.31

POWER MEASUREMENT

Sources of Power Data

- XH3000 internal sensors: PMC
  - Embedded in rack management
  - Accuracy: 5%
  - Valid for Green500 L1 measurement

- Power meter at rack input level: PQBox:
  - Requires specific setup (don’t do that at home)
  - Accuracy: 2%
  - Valid for Green500 L1 measurement

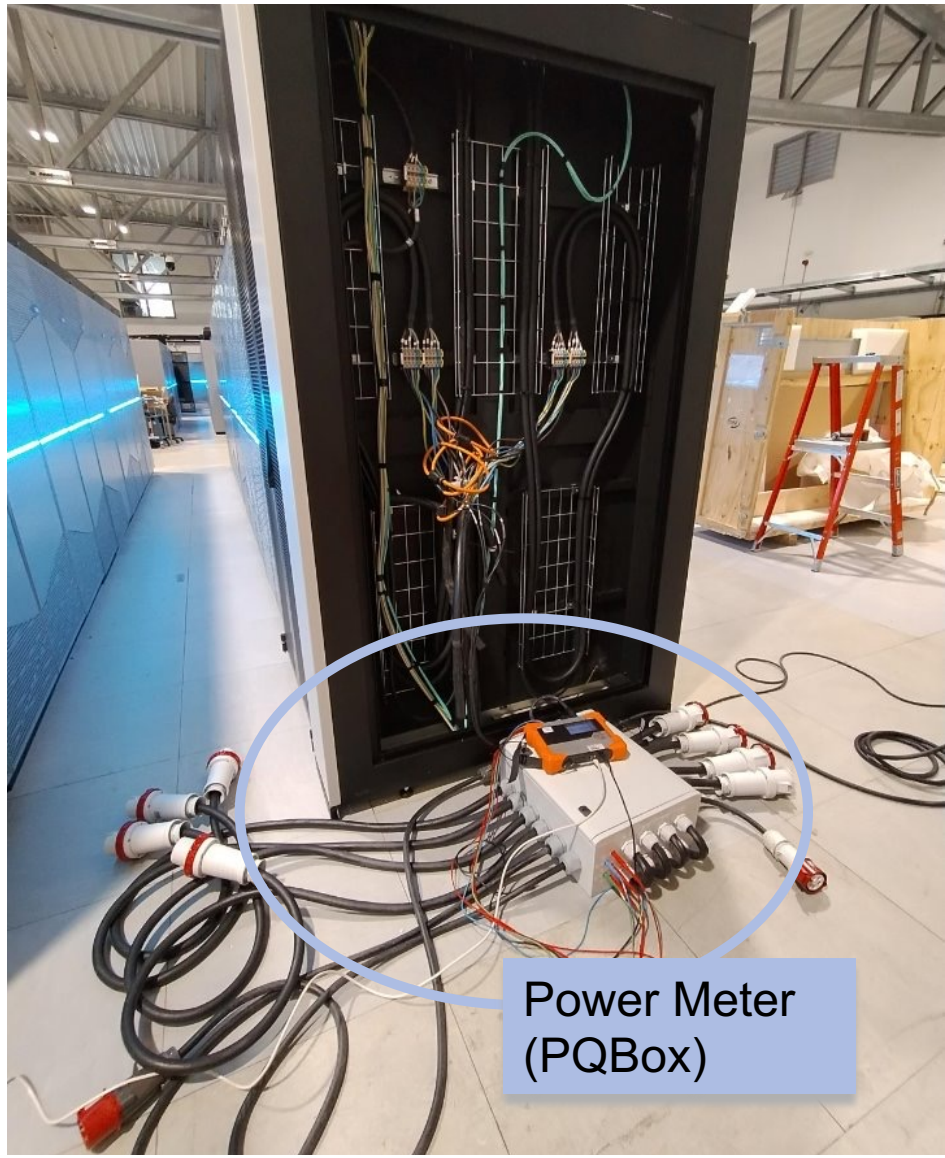

POWER MEASUREMENT

Choice of Source

- XH3000 rack sensors vs Power meter
  - Similar measurements
  - In all tests, variation between 0.5 and 1.3%
  - 73.17 GF/W vs 72.73 GF/W
- Choice to use data from power meter:
  - More conservative measurement results

Sensor	Measurement	Difference
Rack sensors	51838.4 W	0.0%
Power meter	52154.3 W	+0.6%

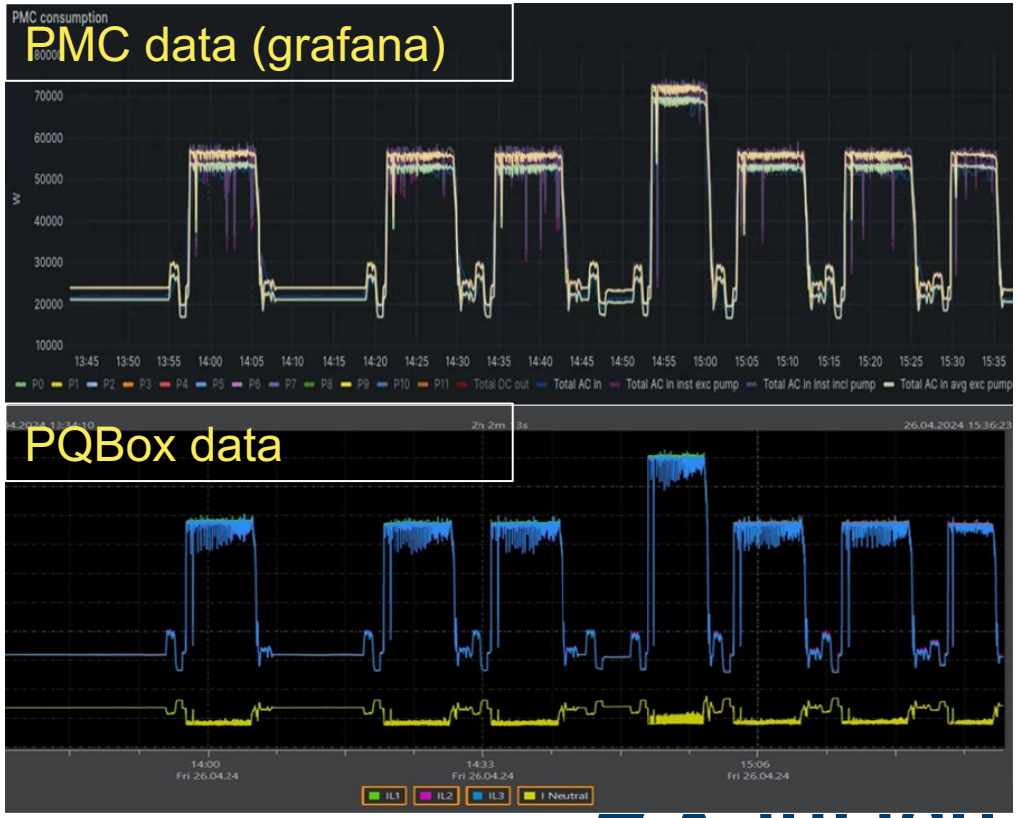

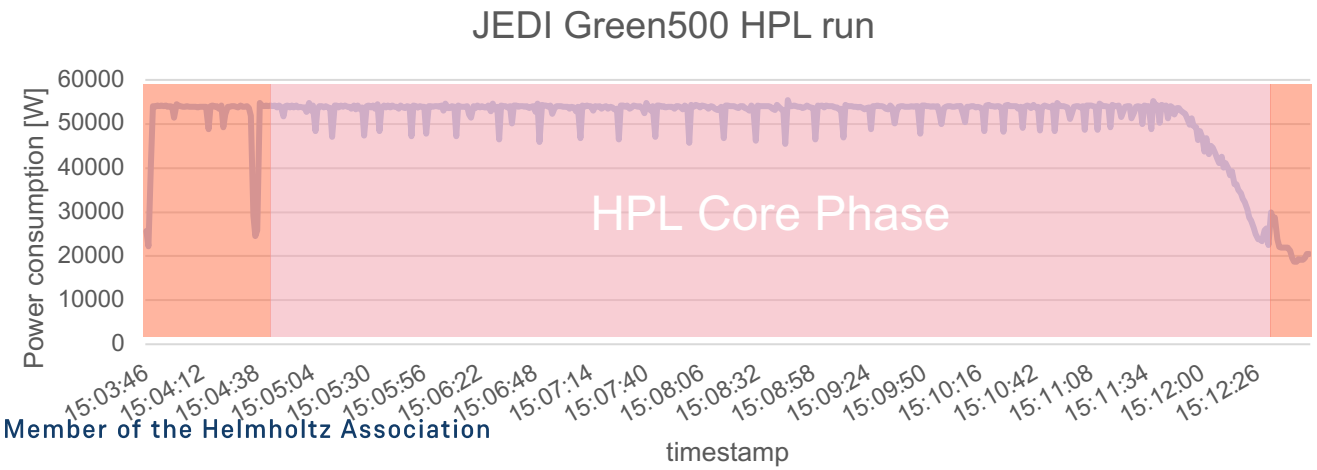

DIFFICULTIES ENCOUNTERED AND CHOICE OF LEVEL 1

- Difficulties to fill “Green500” form ☺
  - A « trap »:
    - HPL performance run is provided as TFLops
    - HPL optimized run is provided as GFlops
    - (HPCG is provided in PFlops)
  - No clear description of what is expected as evidence
    - “Papers” would take a long time, incompatible with submission deadline
  - Missing “APU” type of processor

- Quality level 1:
  - What about level 2 or level 3 ?
  - No clear benefit to claim higher quality level
  - Consistency with planned measurements at larger scale:
    - Scale > 200x
    - Intent to rely on BullSequana rack sensors with 5% accuracy (incompatible > L1)

GREEN500 – JUNE 2024

Green500 Data

Rank	TOP500 Rank	System	Cores	Rmax (PFlop/s)	Power (kW)	Energy Efficiency (GFlops/watts)
1	189	**JEDI** - BullSequana XH3000, Grace Hopper Superchip 72C 3GHz, NVIDIA GH200 Superchip, Quad-Rail NVIDIA InfiniBand NDR200, ParTec/EVIDEN EuroHPC/FZJ Germany	19,584	4.50	67	72.733

JOINING FORCES

Ministerium für
Kultur und Wissenschaft
des Landes Nordrhein-Westfalen

Thank you for the heavy lifting!

fz-juelich.de/jupiter