Hewlett Packard Enterprise

Hewlett Packard Labs

Lifecycle Efficiency: Beyond Scope 1-3

Utz-Uwe Haus, Head of HPC/AI EMEA Research Lab

2025-03-10 at SC Asia 25 workshop “Advancing Energy and Resource Efficiency in Data Centers”

IT Lifecycle Efficiency

GreenHPC is not just about TCO

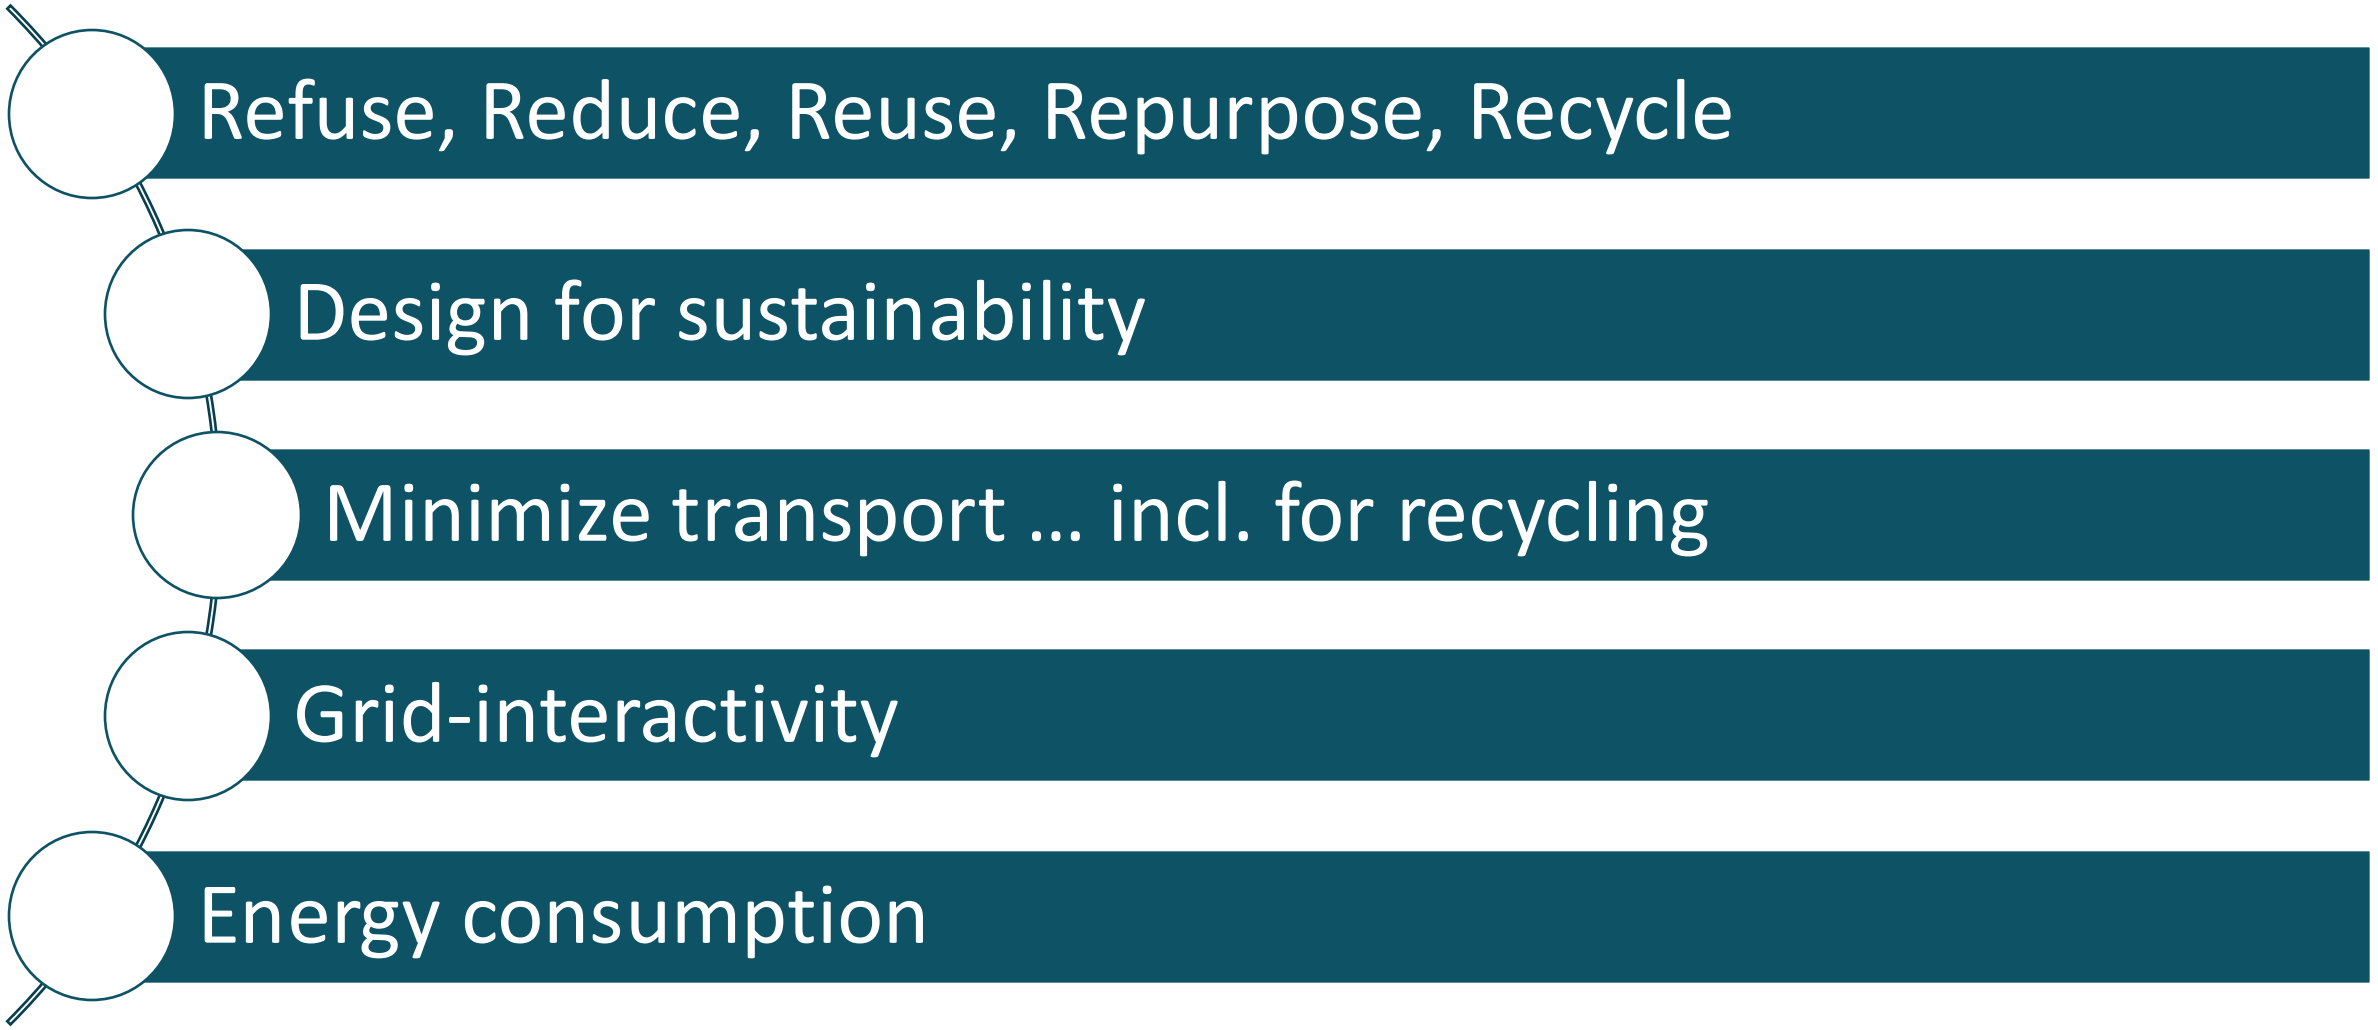

What’s wrong with Scope-based tracking?

- Product focused model
  - Scope 1: own contribution
  - Scope 2: upstream contribution
  - Scope 3: ‘rest of the product lifecycle’
    - Hard to measure, often based on proxy values or industry averages
    - Often the largest part of a company’s footprint
- Encourages shifting responsibility
- Lack of standardization
- Hard to quantify overlaps (double-counting)

HPE is striving to become a net-zero enterprise by 2040

We continue to align our goals to the latest climate science with a new suite of science-based targets. Beginning in 2022, executive compensation will be tied directly to climate metrics.

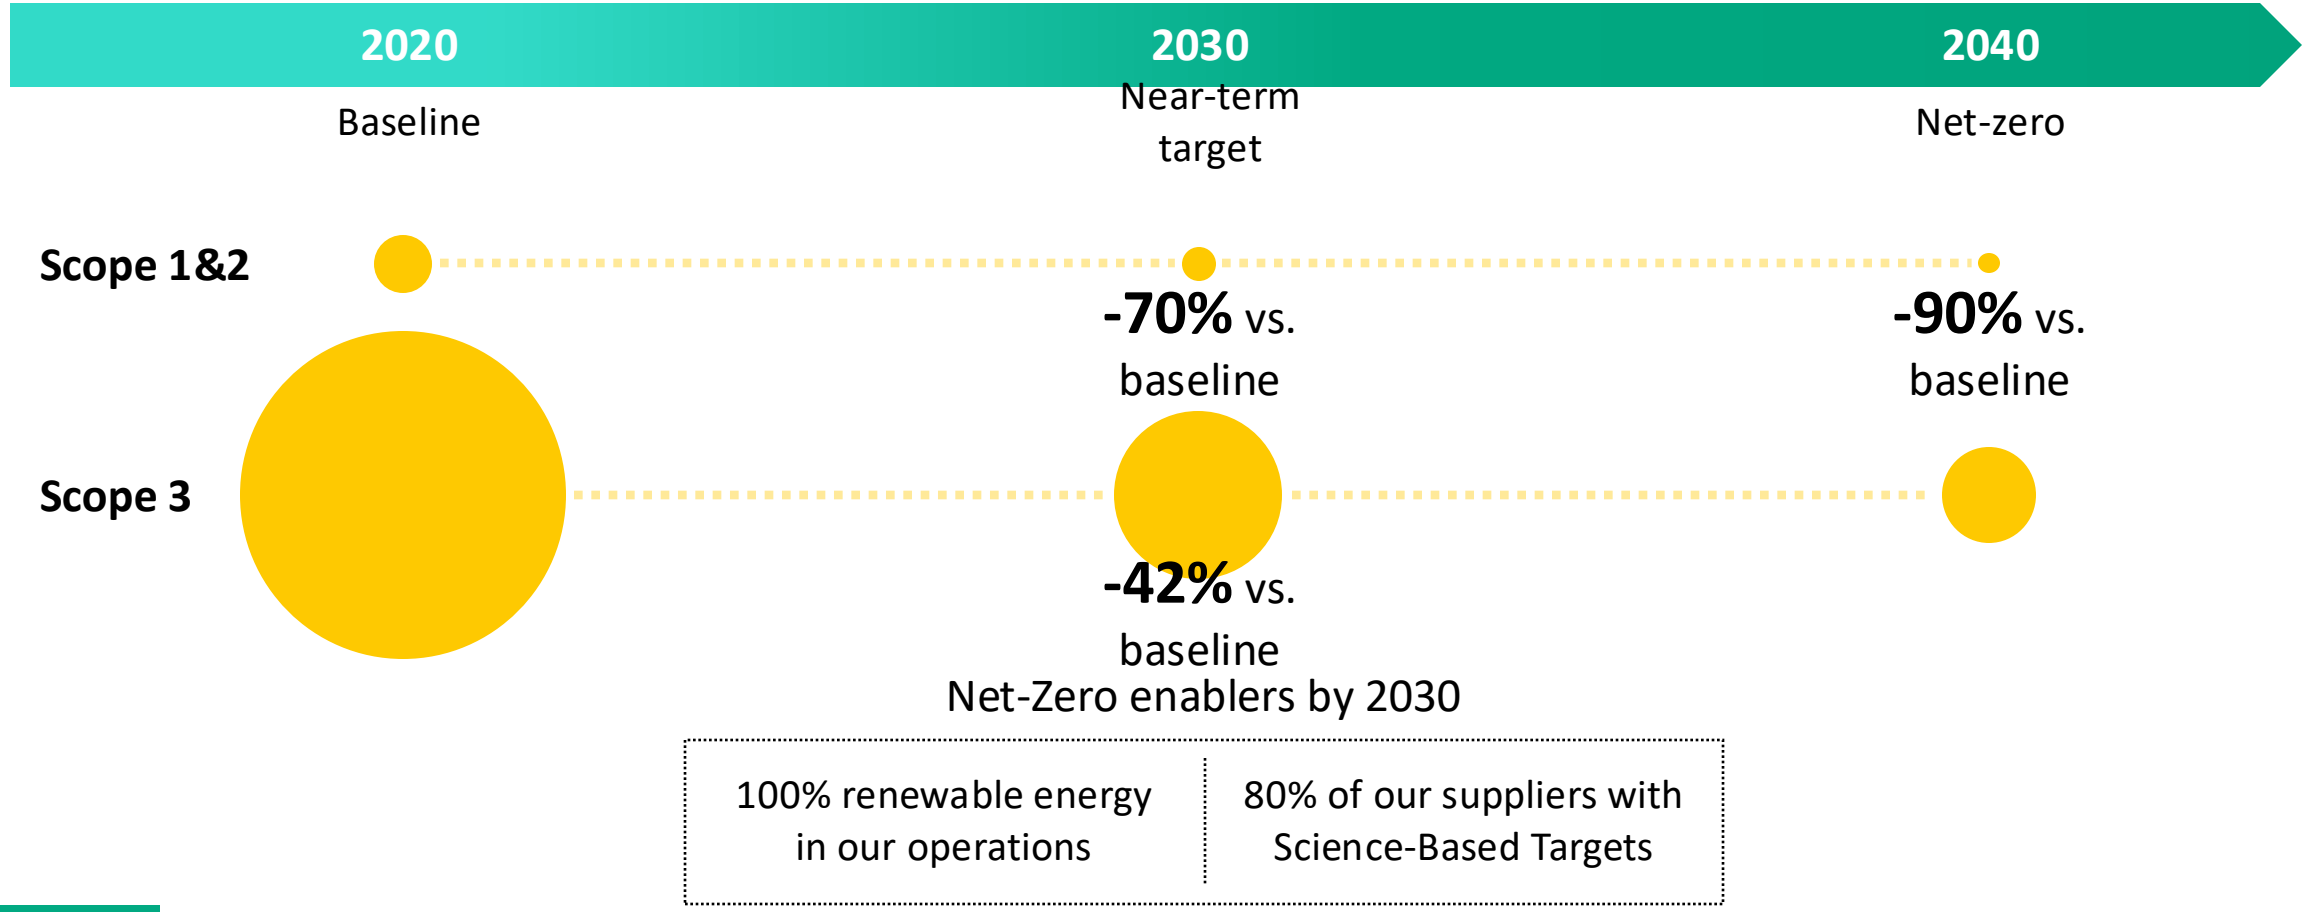

Net-Zero enablers by 2030

100% renewable energy in our operations	80% of our suppliers with Science-Based Targets

Energy efficiency and CO_2 footprint

- Energy efficient computing
- Green data centers
- Carbon neutrality
- Prioritizing **energy-to-solution** over time-to-solution

1.5 Exaflops cost ~30 MW

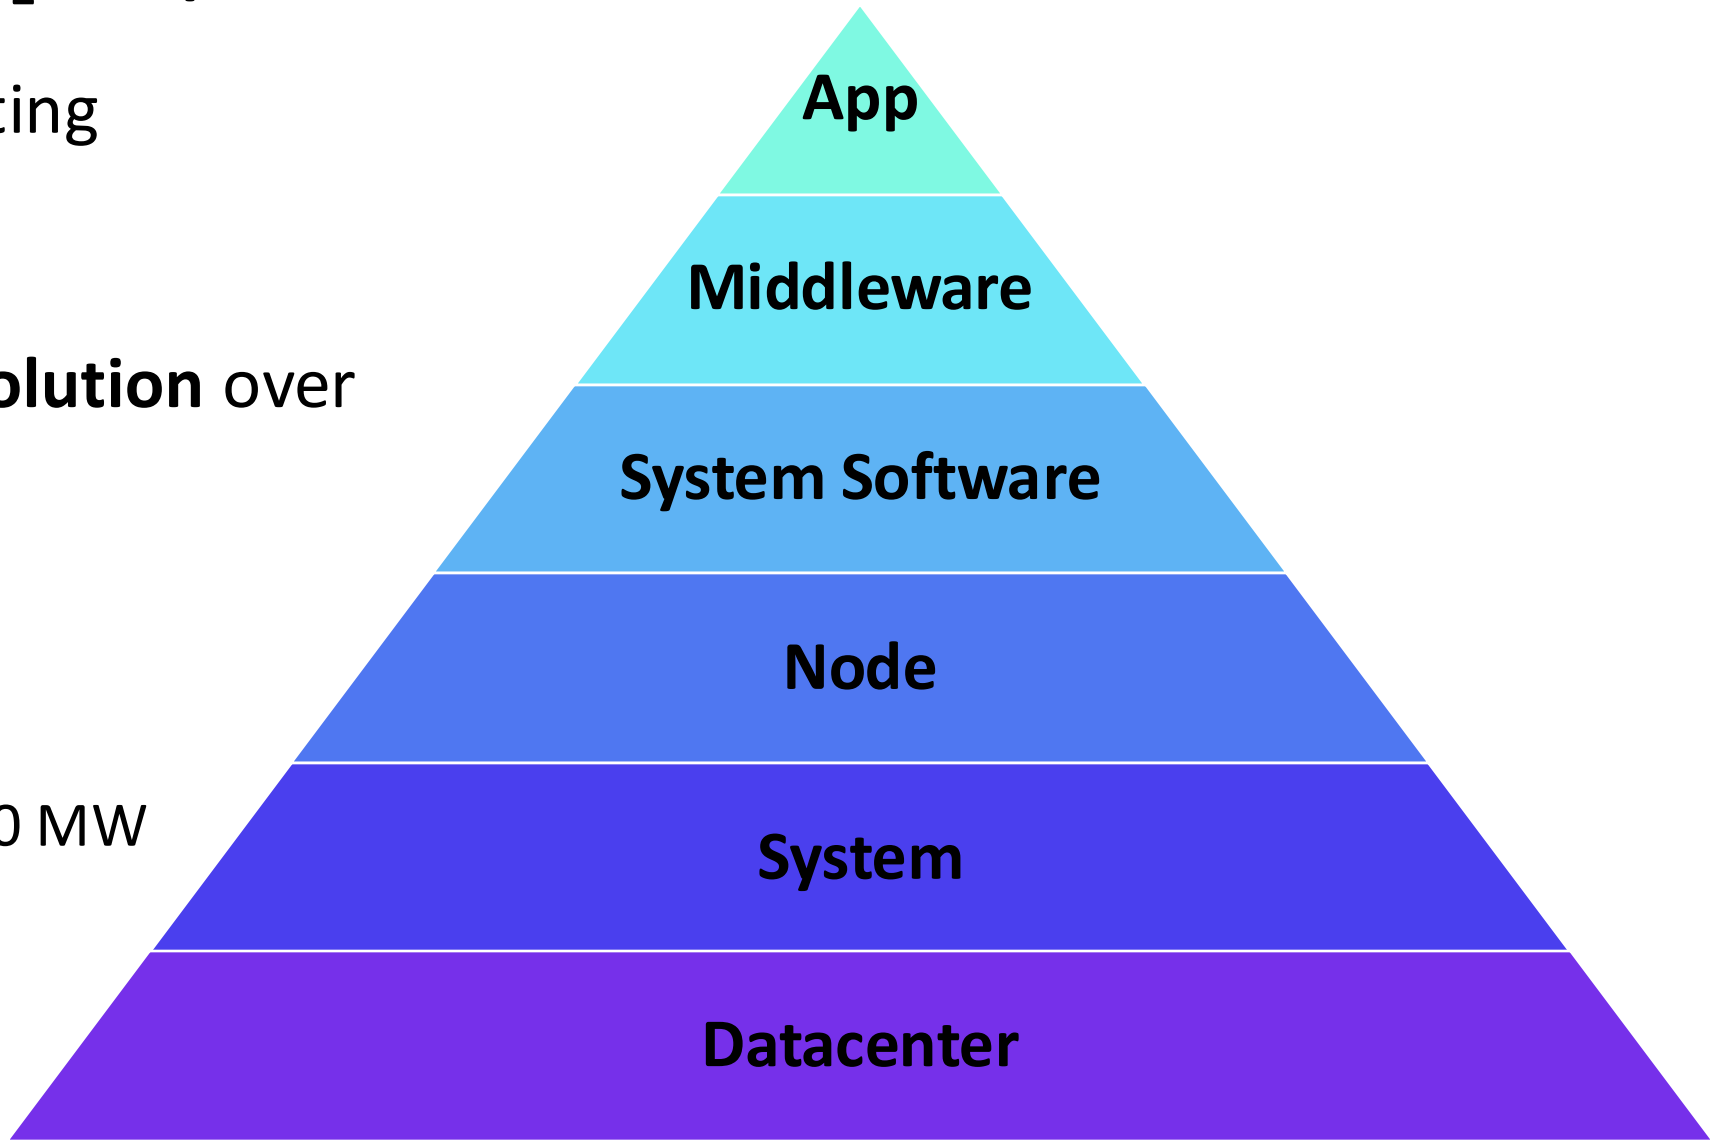

Data Movement Middlewares

Data movement is expensive

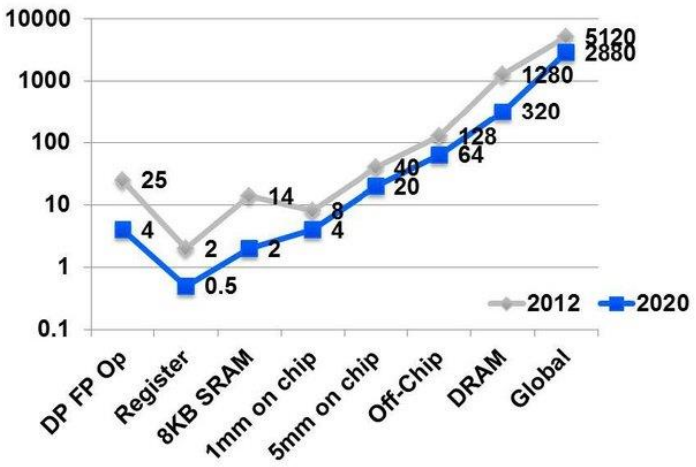

Energy cost of data movemen:

pJ per 64-bit FP op

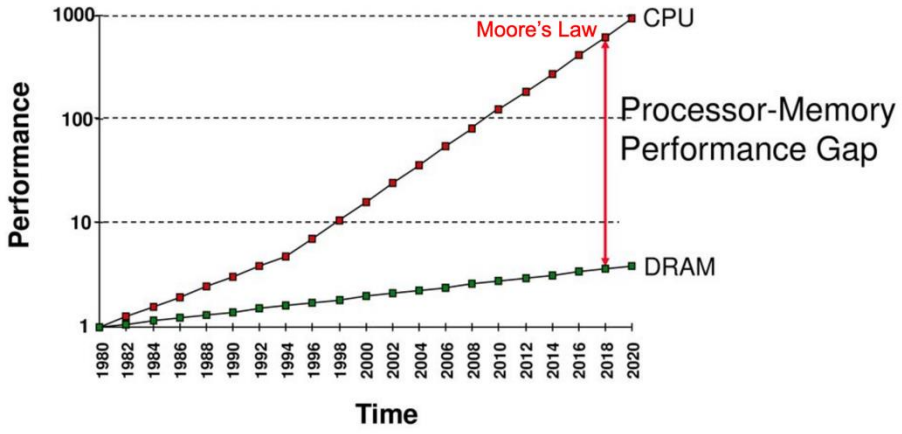

Patterson, UC Berkeley

Need data object abstraction+system model

- Across applications, across memory tiers, across compute nodes

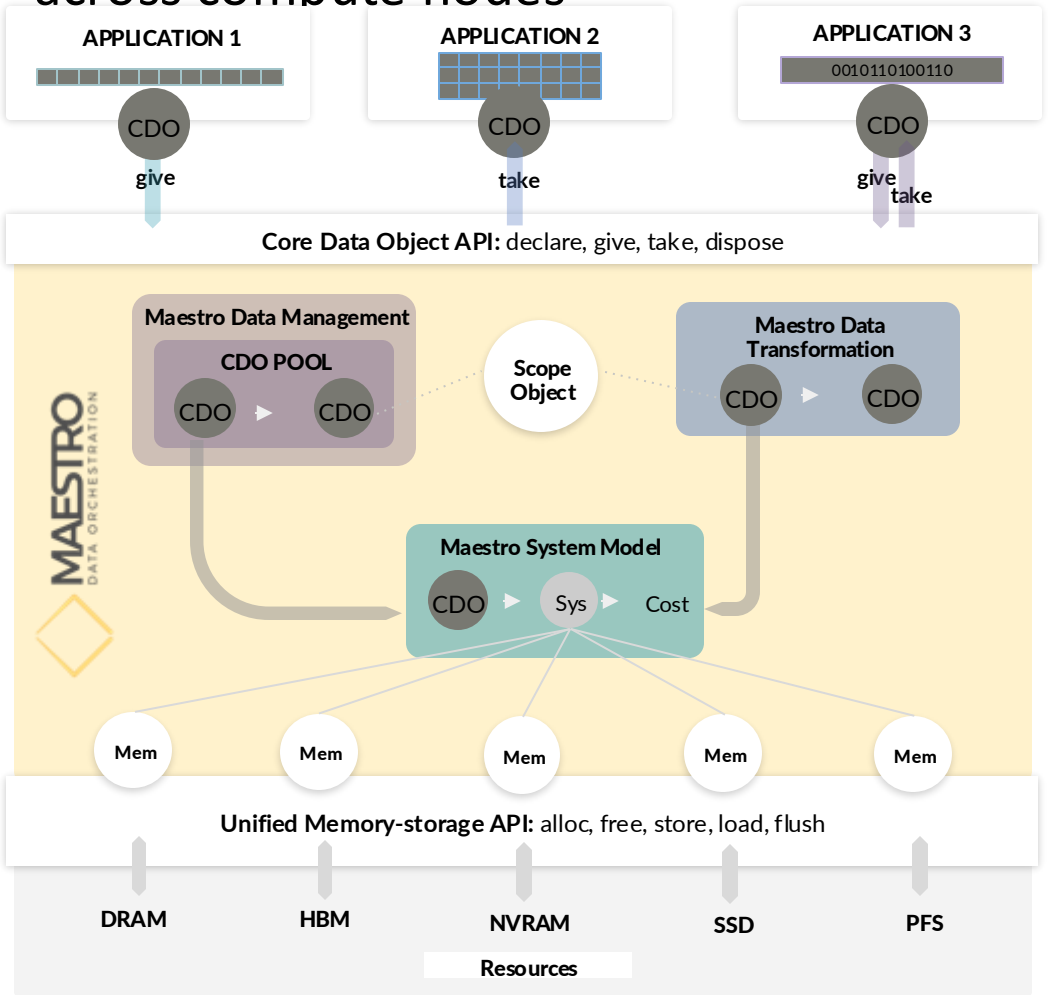

https://maestro-data.eu/

Global view: Data movement cost

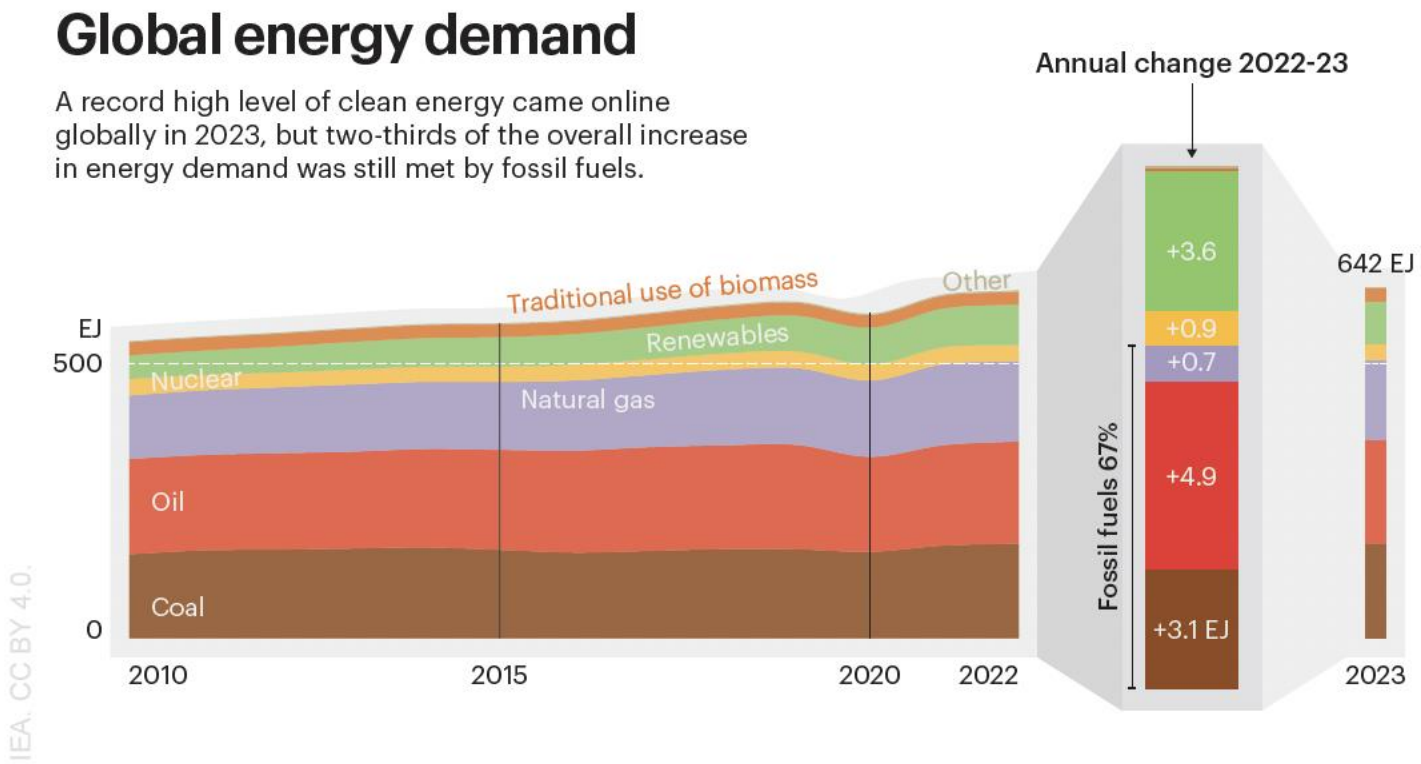

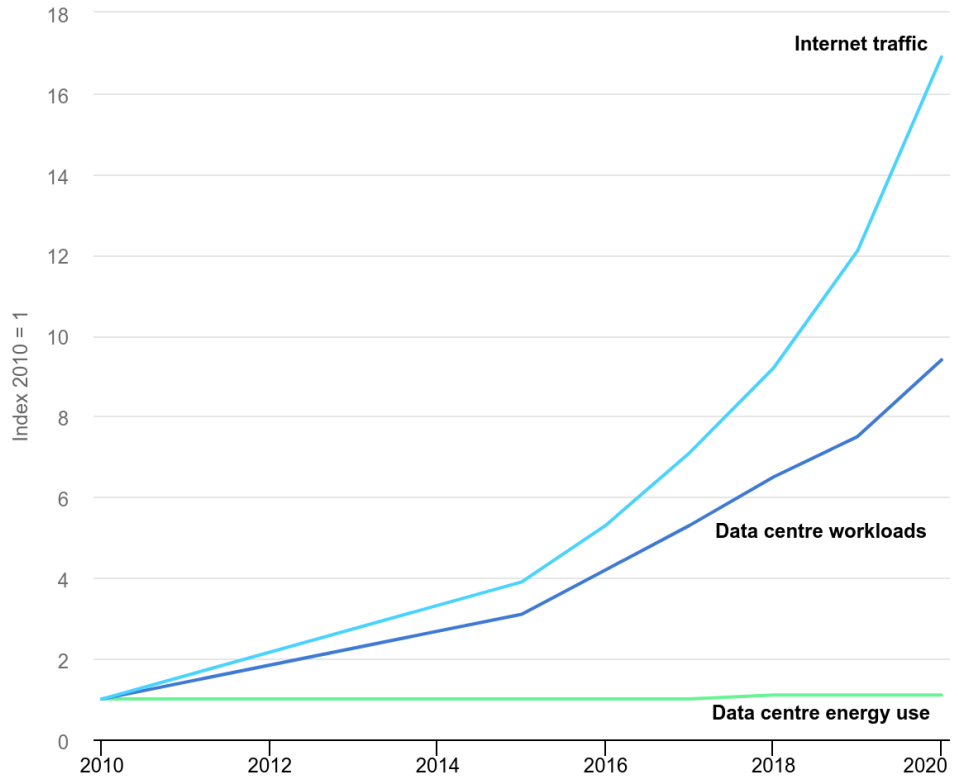

IEA World Energy Outlook 2024

https://www.iea.org/data-and-statistics/charts/global-trends-in-internet-traffic-data-centres-workloads-and-data-centre-energy-use-2010-2020


HPC/AI Power consumption

Power to Train Frontier Models

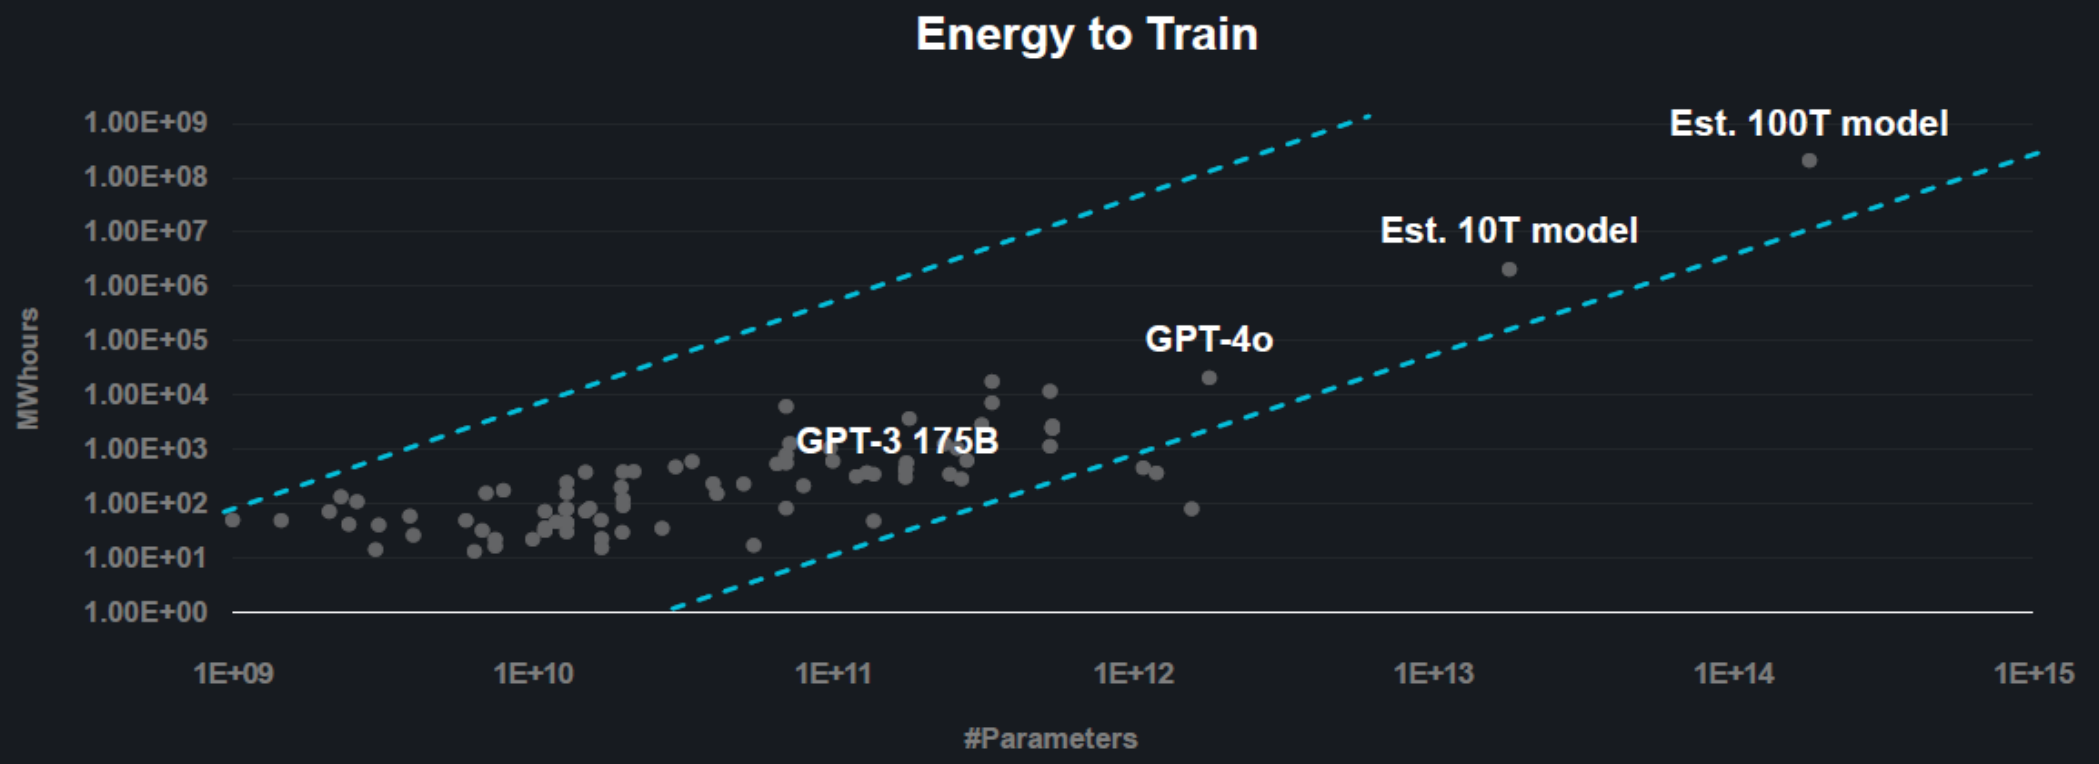

Exponential growth in model sizes drives massive increase in energy required for training

Total data center power capacity limits model size

Based on published parameter counts of leading training models; and AMD internal calculations approximating energy to train

Ref: Victor Peng, HotChips 2024

Utilized Power as % of Nameplate Power

Provisioned Power vs. Average and Peak Power

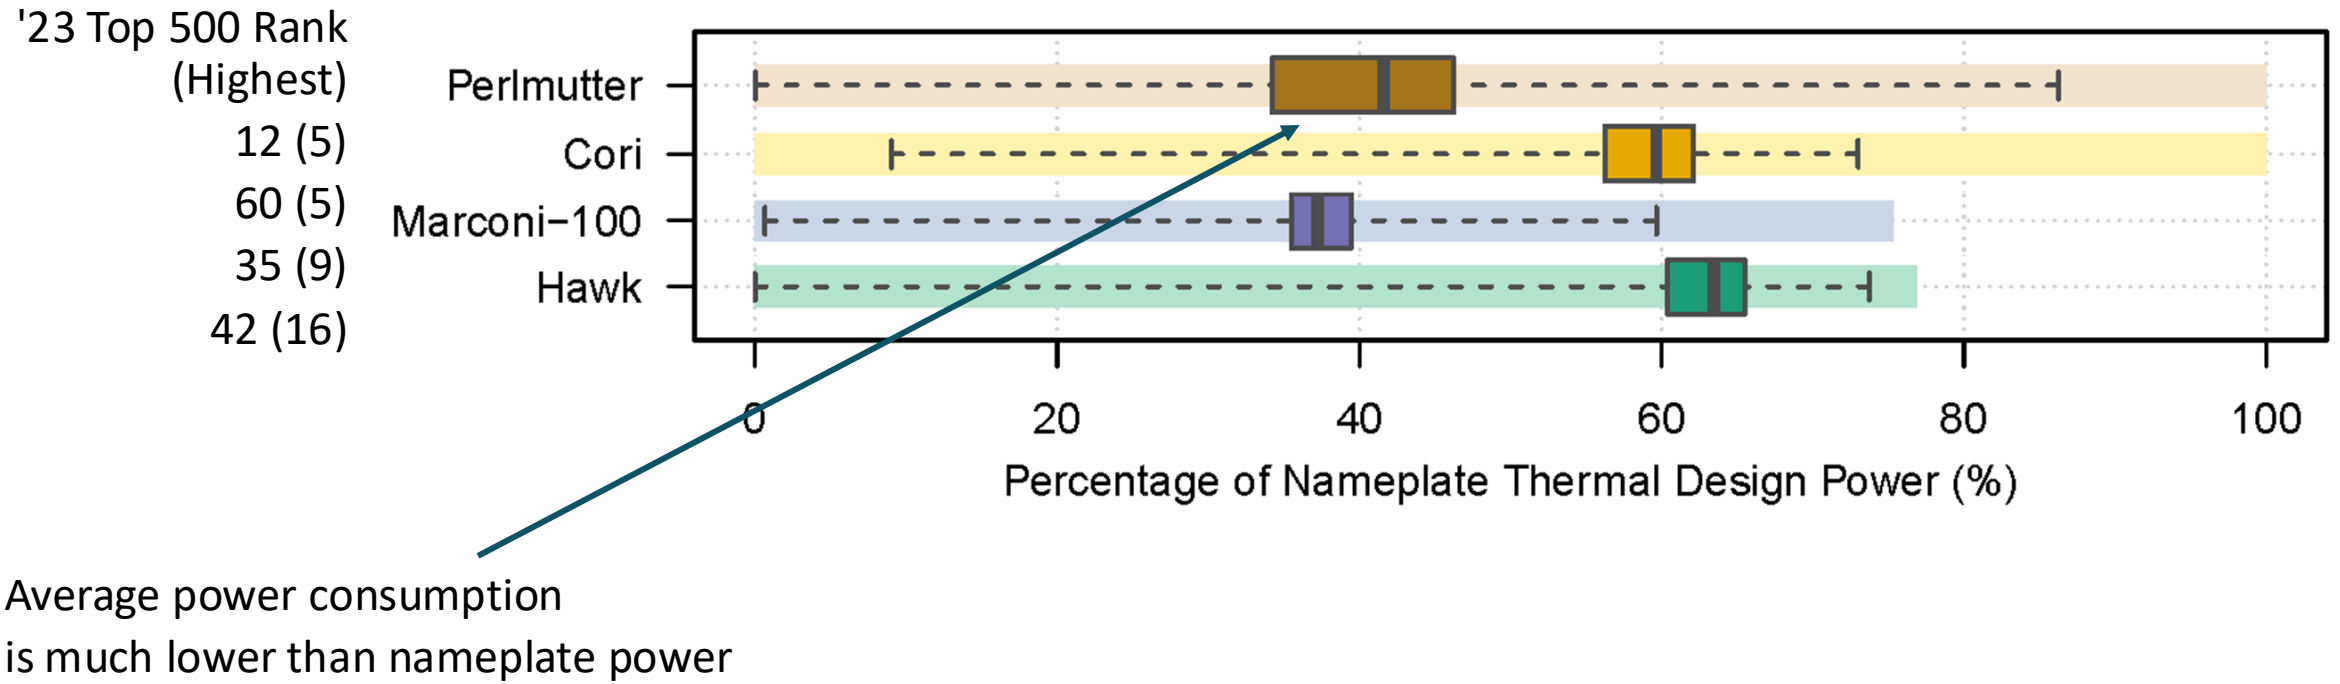

Average power consumption
is much lower than nameplate power

Customer HLRS setup of Hawk HPC system

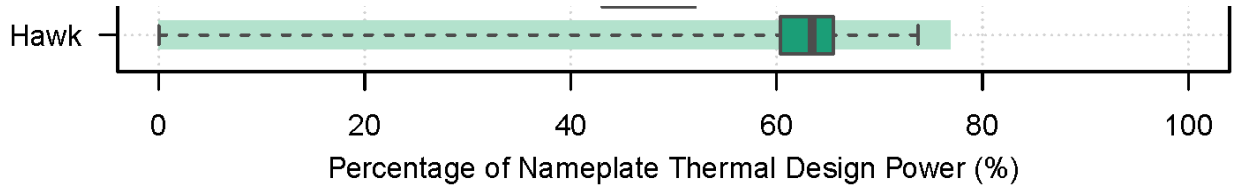

	Top500 Rank Nov. 2023 (Highest)	Year Installed	Nameplate TDP (MW)	Provisioned Power (MW)	**Provisioned Power used**
Hawk (HLRS)	#42 (#16)	2020	4.49	**3.45**	Avg 81% (max 97%)

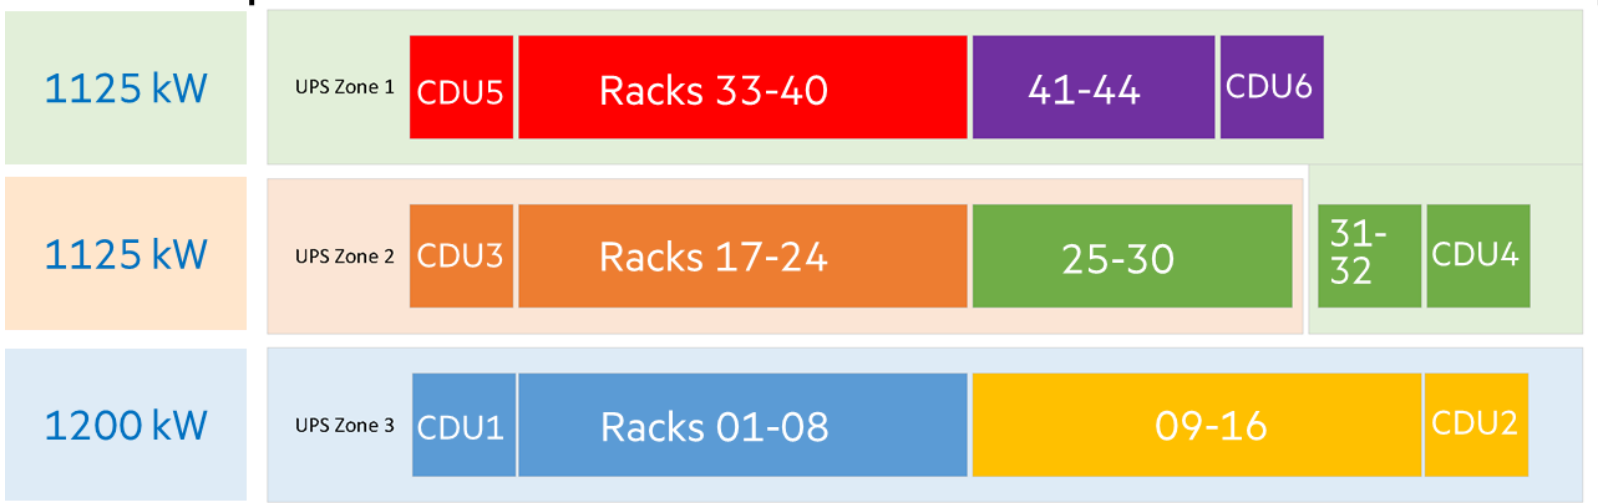

Uncapped (UC) vs Static Capping (SPC) vs. Dynamic Capping (DPC) on Hawk

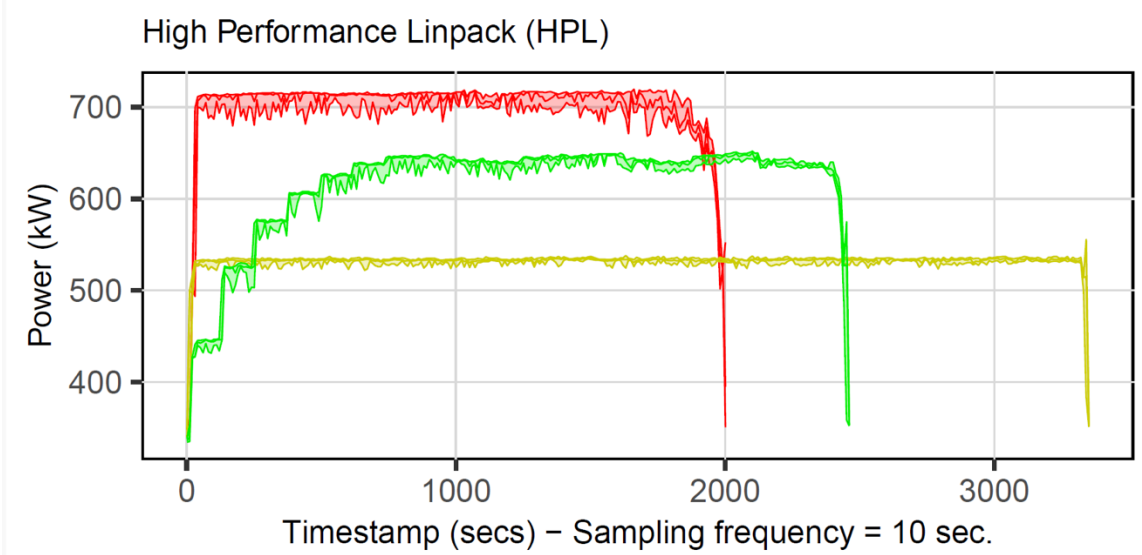

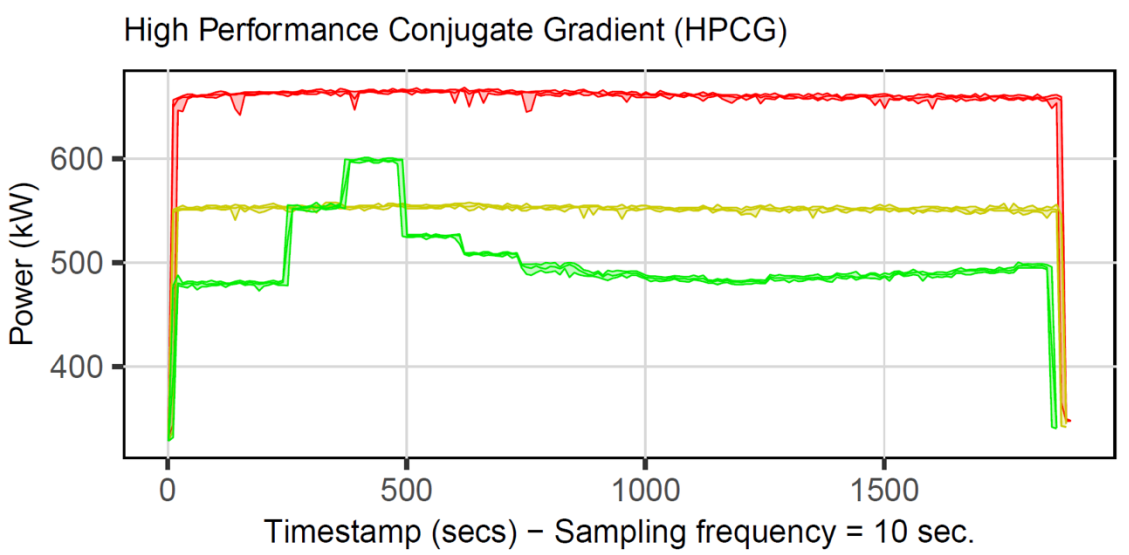

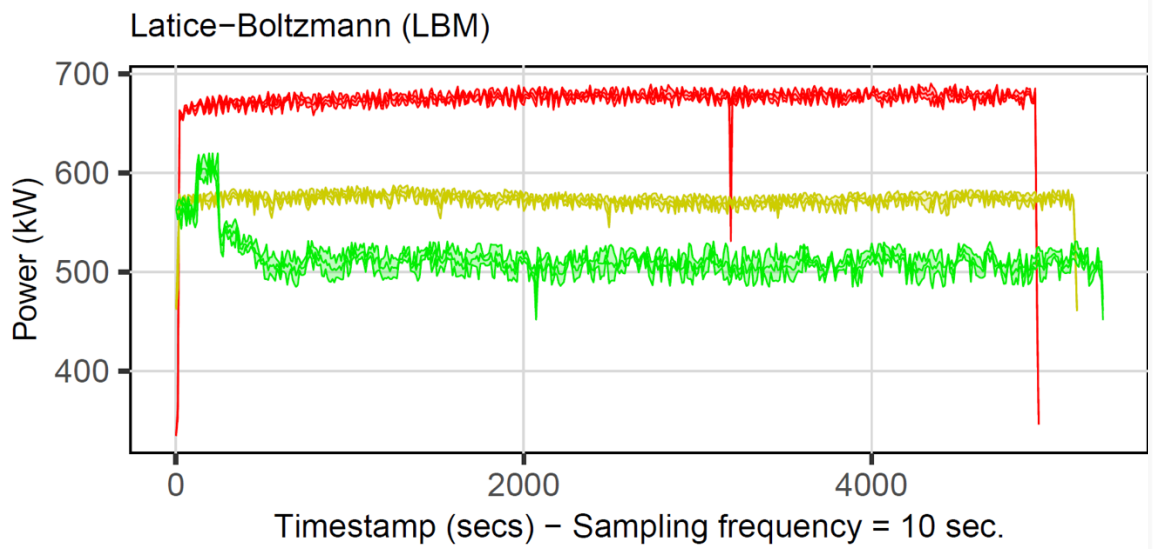

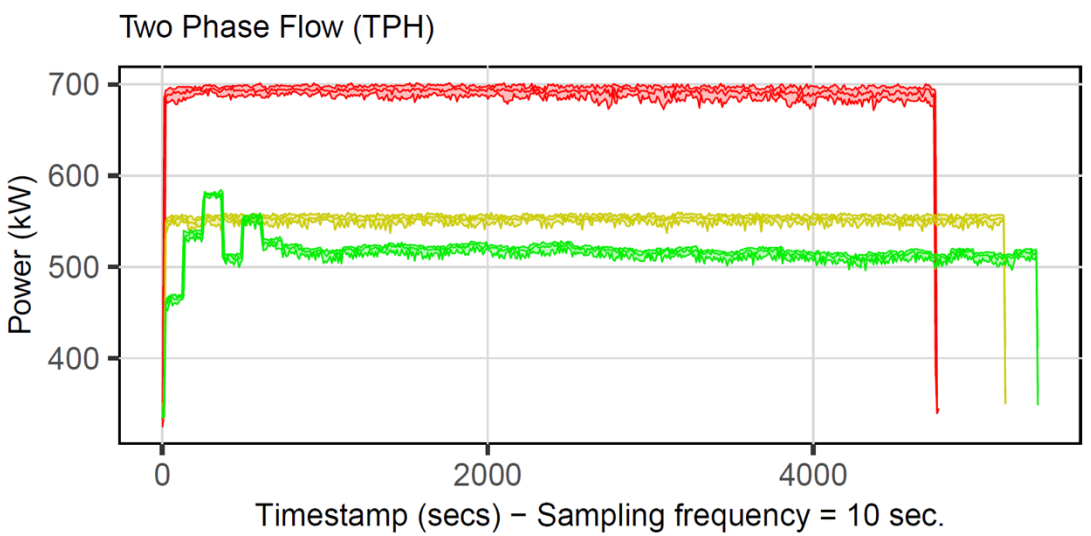

PowerSched* – A POC for hardware over-provisioning

Customer HLRS setup of Hawk HPC system

- Results show reduced performance impact against standard static uniform node power caps under the same power restriction
  - PowerSched gains over static capping
    - –up to 25% in performance for compute bound applications
    - –up to 14% in energy savings for memory bound applications

Test on one rack of Hawk	HPL (average node power draw in W)	HPCG (average node power draw in W)	TPH (average node power draw in W)	LBM (average node power draw in W)	App mix (average node power draw in W)
Static capping at 172W	536	552	553	574	2215
PowerSched (max 220W)	639	490	517	515	2161

- The PowerSched PoC is currently managing the production HPC system Hawk at HLRS

Scheduler level operational changes

Archer2 scheduling simulation…

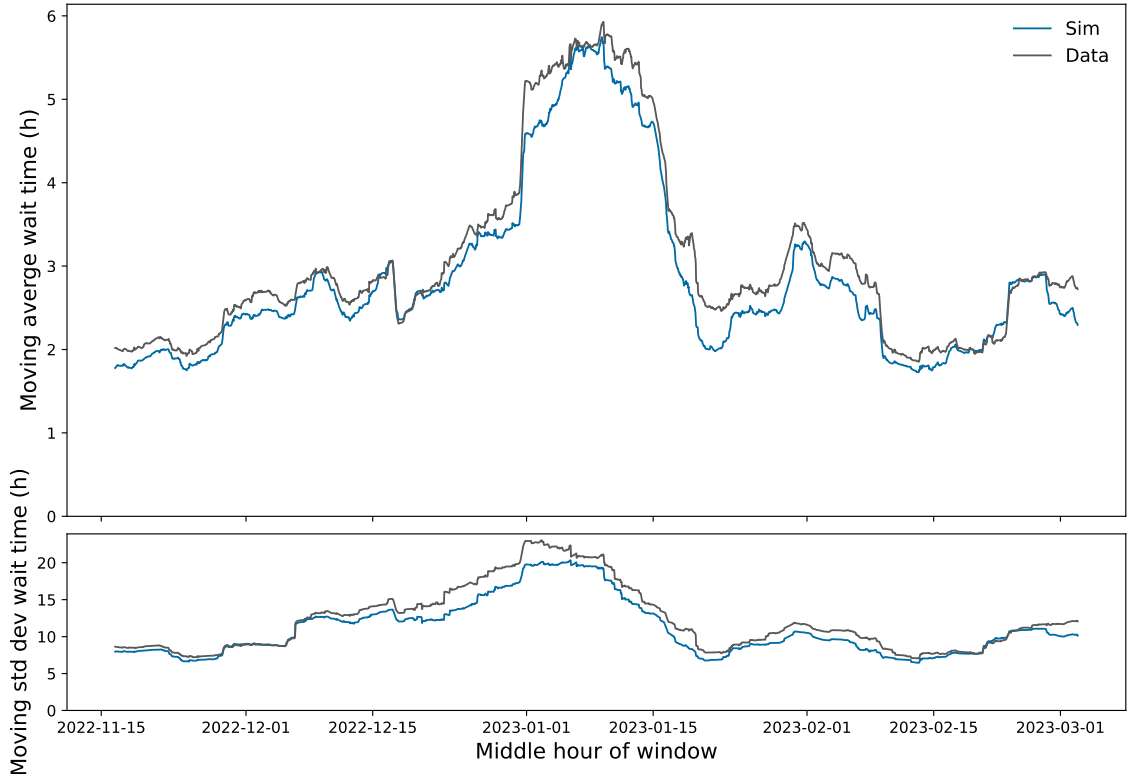

… across diffent QoS classes

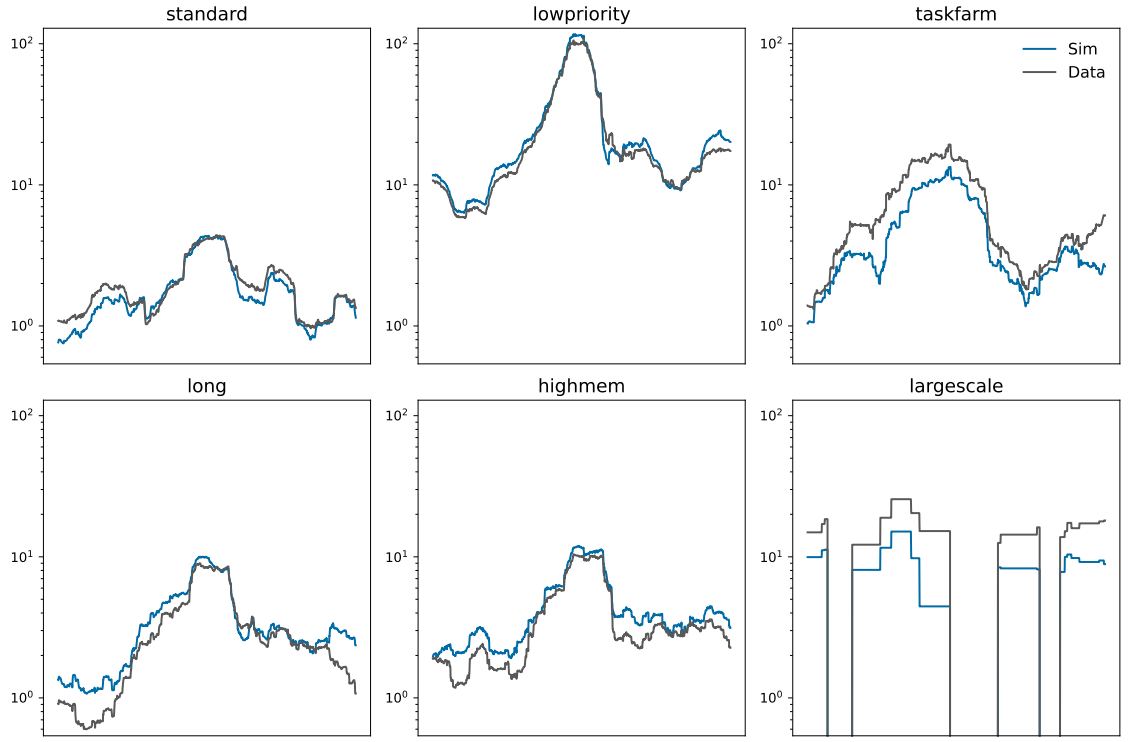

Alex Wilkinson et al, MODA4@ISC23

Scheduler (2)

Shifting power across the schedule…

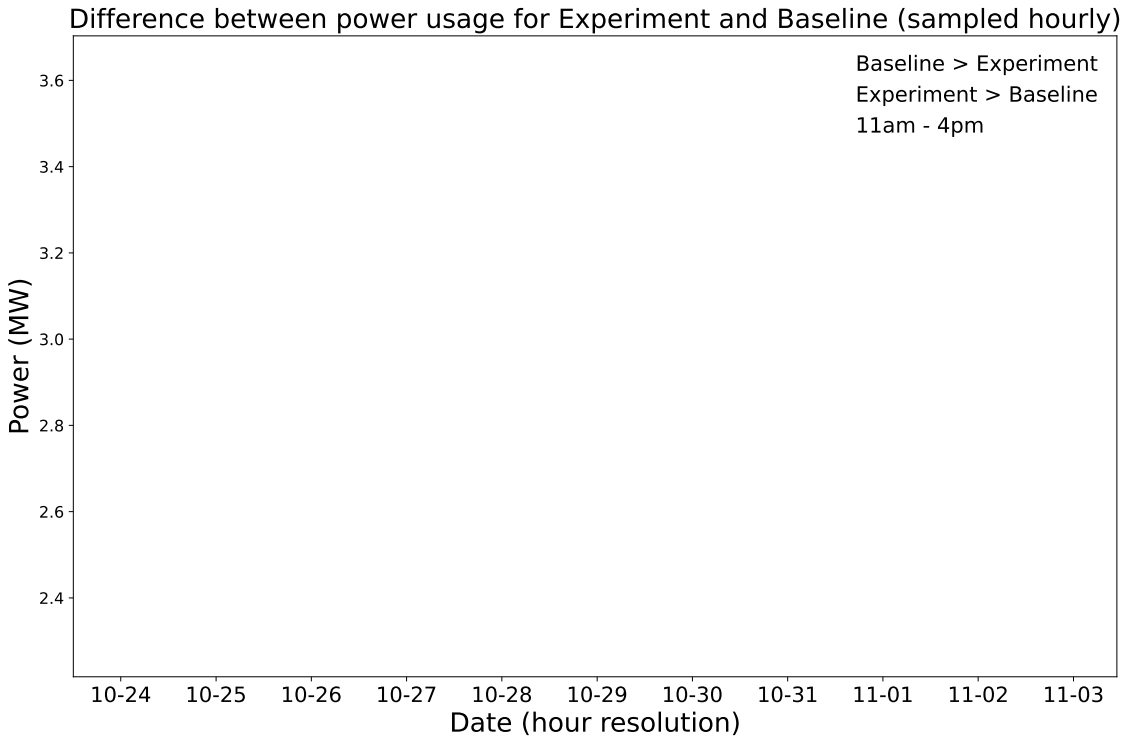

… with low impact on users

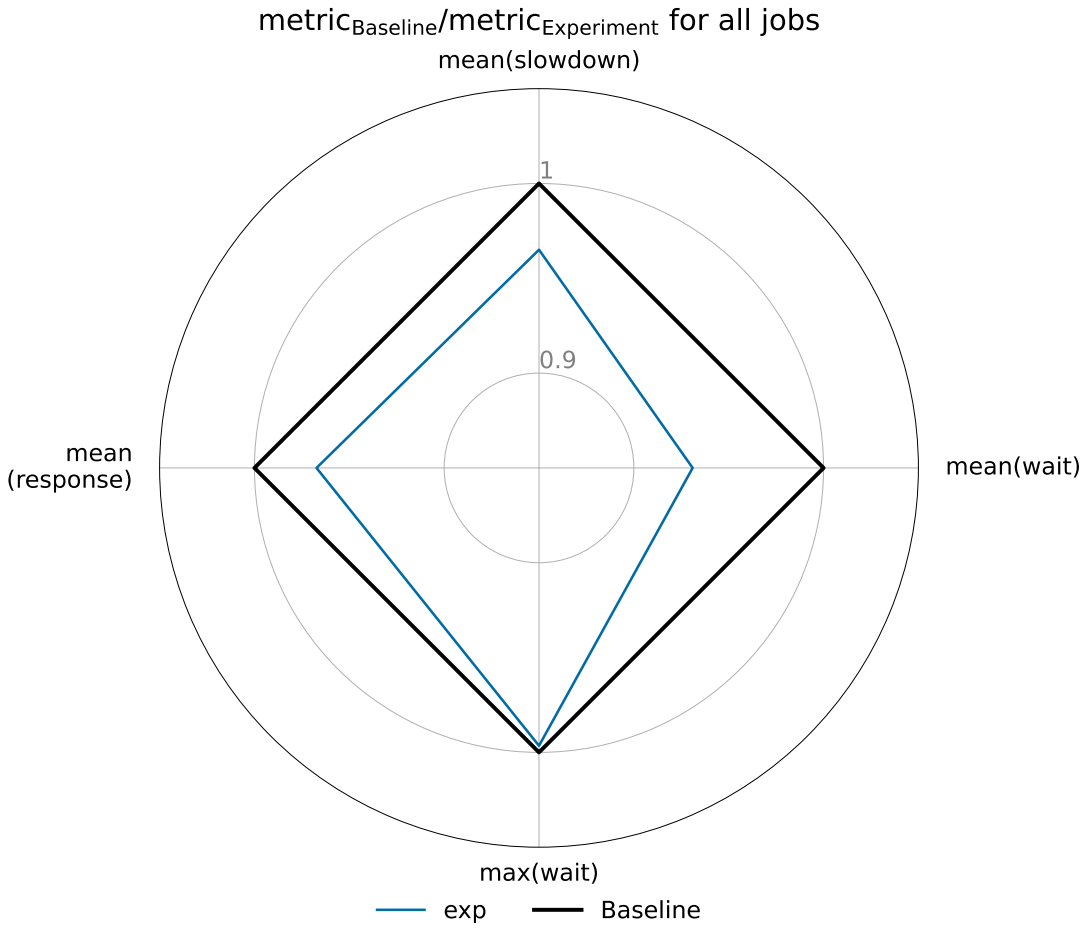

Scheduler-level toolkit

Components

- Power prediction
  - To allow controlling system power
- Scheduling simulation
  - To evaluate dynamic decisions near-realtime
  - To evaluate strategy changes before deployment
  - Avoiding full slurm simulations
- Application identification
  - to apply appropriate micro-models

Power series forecasting

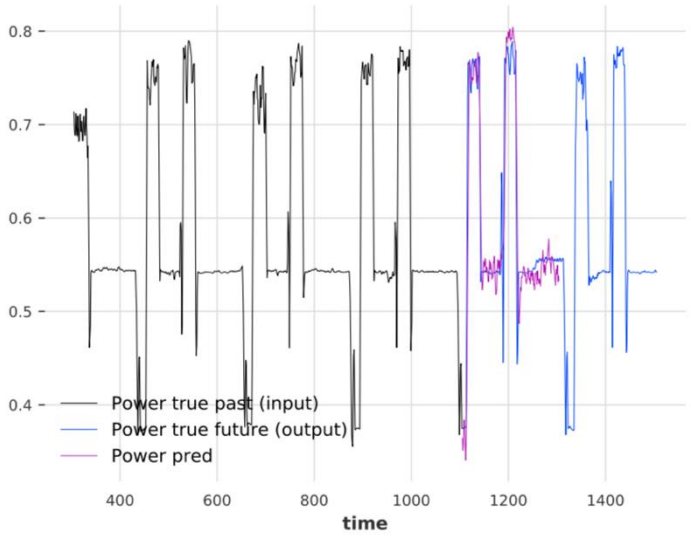

Frequency response tests

Use the scheduling simulation to quantify the performance to energy tradeoff for running jobs at low frequencies when queue not very full

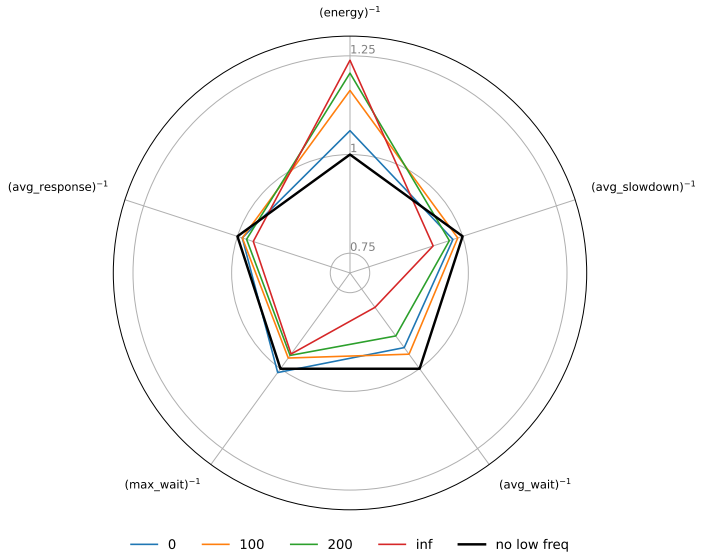

About ERL

HPE HPC/AI EMEA Research Lab Summary

Partnering with leading organizations in the EMEA region to advance supercomputing R&D

Our Role

- Deep technical collaboration with industry, academia, and public sector.
- Long term technical relationships surrounding research, co-design, and operational support.
- Focus on new technologies, driving HPE products.
- Create reusable PoCs & European IP

Research Interests

- HPC, Cloud, AI, Quantum
- Data movement, analysis, and workflows
- Heterogeneous computing and novel accelerators
- Programming languages and models
- Compilers and mathematical optimisation
- Performance portability, security, and containerisation
- Energy efficiency and sustainability

Engagement Models

- Centres of Excellence
- Advanced Collaboration Centres
- Value-add projects
- Joint-funded research projects
- Nationally/internationally funded research projects
- Ph.D. and Placements

The EMEA Team of Hewlett Packard Labs

Advanced Collaboration Centers In EMEA(A)

Coming soon:
BlueLion at LRZ

Archer2, UK	Shaheen 2&3, KAUST	GW4 Isambard 3 + UoB Isambard AI, UK	LUMI, EuroHPC JU	Setonix, Australia
• Focus is on delivering science • Revisiting LASSi - IO Monitoring and Analytics adding Power/Network • Application tuning and GPU testbed • MODA internship (energy and scheduling) • New focus on high-end users and capability runs (> 2000 nodes)	• Machine learning for Bioinformatics • Numerical linear algebra libraries • Combustion CFD • Coupled visualization workflows • Hybrid HPC+AI • Sustainability topics in datacenter	• ARM ecosystem development • Grace/Hopper + Slingshot tuning • Hackathons • Modular Data Center monitoring • Secure workload enclaves • Multitenant supercomputing	• European Pre-Exascale • Projects covering Containerisation, Easybuild and porting to LUMI-G (GPUS). • Extensive training program and Hackathons • Federated HPC • Holistic Energy-Awareness • Secure/Trustworthy HPC	• Australian national service • Software containers, AI workloads on HPC • Secure workflows • Data center digital twins • SKA workloads • Holistic power management
				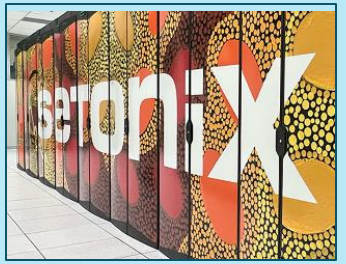

Thank you

uhaus@hpe.com