Hewlett Packard
Enterprise

GREEN VS. EXASCALE HPC : CARBON-NEUTRAL SITE OPERATIONS, ENERGY EFFICIENCY AND OVERALL SUSTAINABILITY

Utz-Uwe Haus
HPE HPC/AI EMEA Research Lab

2020-07-12 ISPDC22, Basel

ENERGY EFFICIENCY

- Energy efficient computing
- green data centers
- carbon neutrality
- prioritizing **energy-to-solution** over time-to-solution

- If ~1.5 Exaflops (10^{18}) costs ~30 MW
- How much does a Zettaflop (10^{21}) cost?

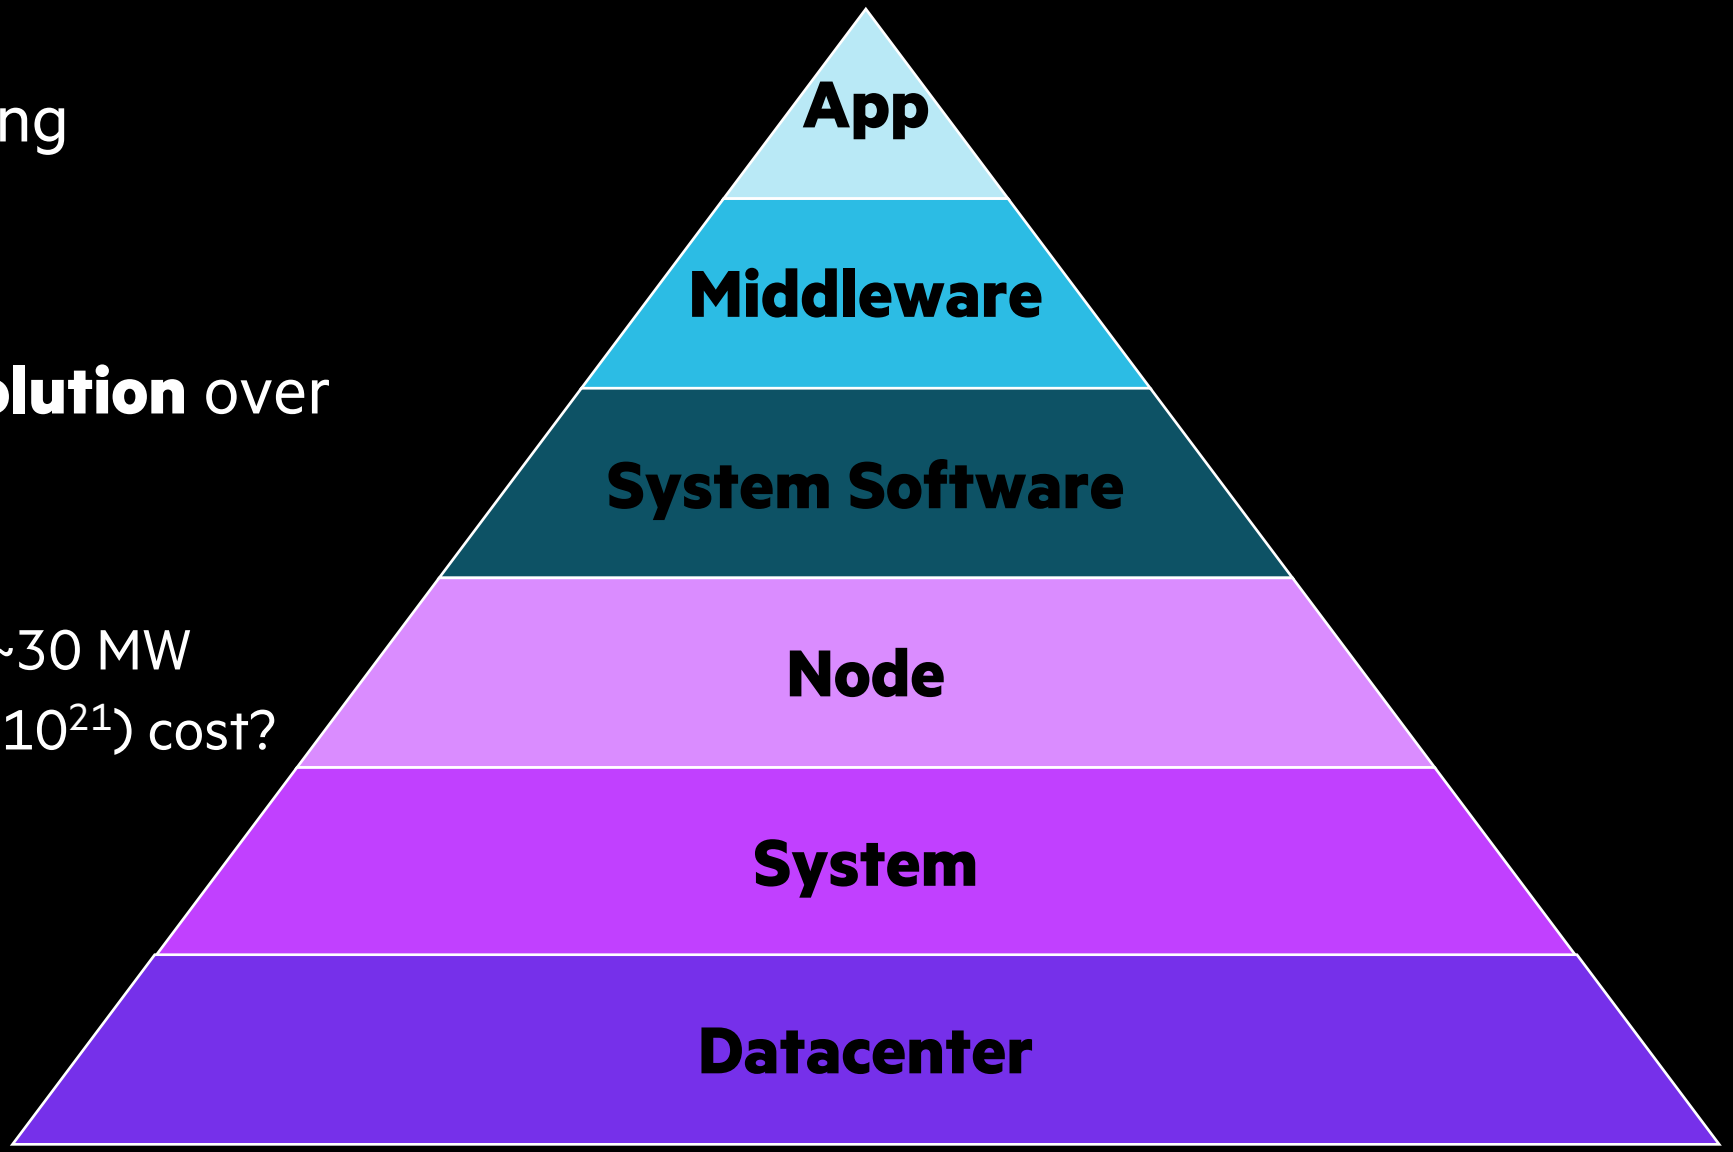

LUMI DATACENTER IN KAJAANI

High end:

- **Biggest European HPC system**
- **10 partner countries**
- **#3 on June22**
  - **Top500 HPL**
  - **Top500 HPCG**
  - **Green500**

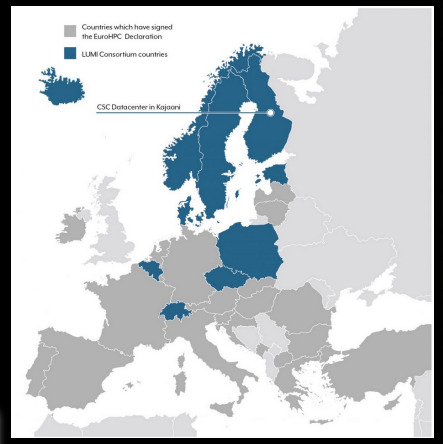

100% hydroelectric energy up to 200 MW

Very reliable power grid: Only one 2 min outage in 38 years

100% free cooling available, PUE 1.03

Waste heat reuse: effective energy price 35 €/MWh, negative CO_2 footprint: 13500 tons reduced every year

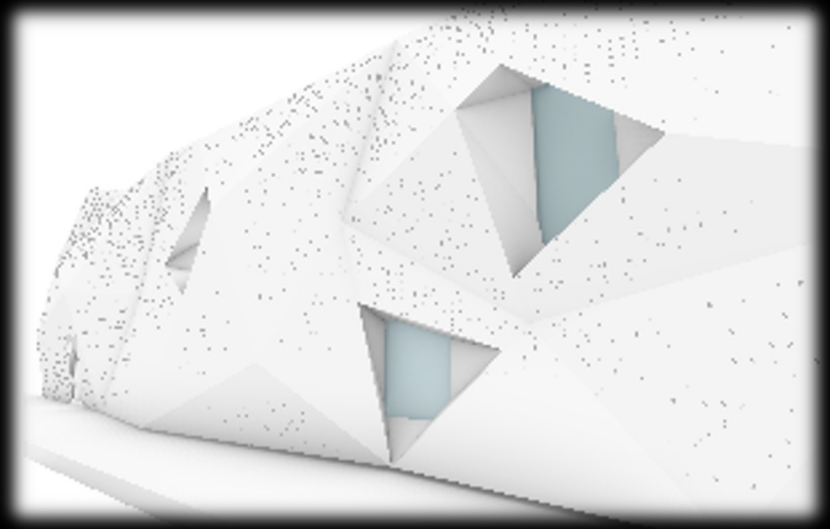

https://www.lumi-supercomputer.eu/

FEDERATED COMPUTING ENABLES GREENER IT
Really?

Gaia-X and Dataspaces

- Current USPs:
  - Compute near data
  - Vendor agnostic cloud-like computing
  - Trusted computing
  - Composable services
- Green aspects
  - Choose compute resources by ecological criteria
  - Decentralize compute by geographic opportunities
  - Attest Green IT aspects

Worldwide Datacenter Power

- 200 TW (= 4x Switzerland)

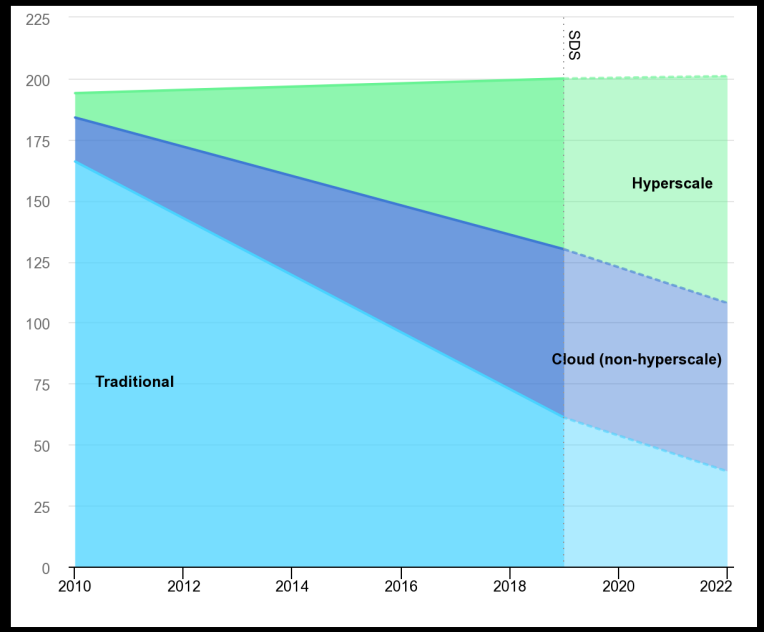

Usage growth by category

- normalized to DC energy use

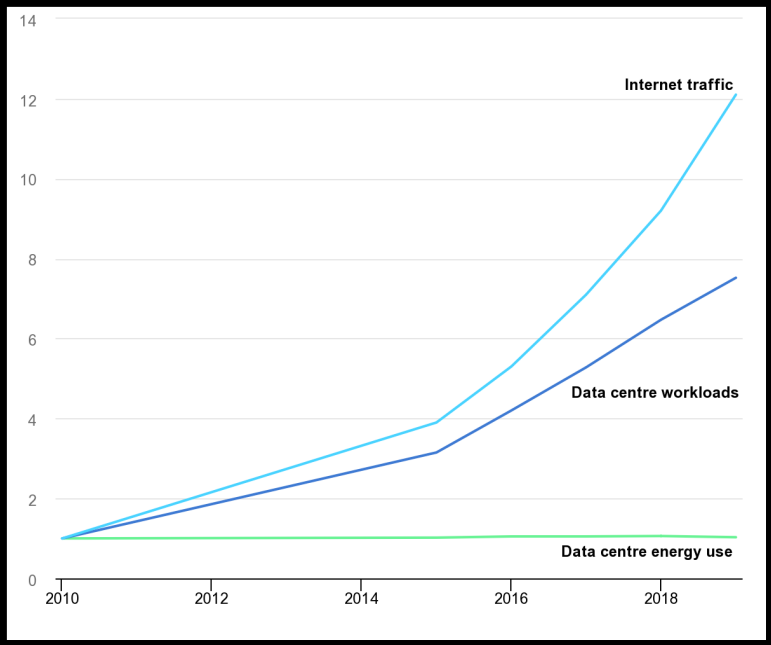

”How much CO_2 do we spend to compute a solution to stop global warming?”

IEA study, 2019/20, last update 2022-03-25 http://iea.org/

GREENHPC IS NOT JUST ABOUT TCO

- Refuse, Reduce, Reuse, Repurpose, Recycle
- Design for sustainability
- Minimize transport … incl. for recycling
- Grid-interactivity
- Energy consumption

SUSTAINABILITY BEYOND ENERGY USAGE

Refuse, Reduce, Reuse, Repurpose, Recycle

- Sustainability is becoming part of TCO calculation
  - Can become serious part of design space
- Extend hardware lifetime
  - By refurbishing
  - By smarter middleware that can handle degrading components
  - Heterogeneous compute architectures (may) help extend lifetimes of installations
- Use spare cycles
  - This is the original rationale of AWS
  - Edge-to-Cloud paradigm, HPE Greenlake

Key figures:

- 80% of environmental impact influened during design
- 30% of large DC servers are unused
- 73k tons of IT equipment recycled by HPE 2018-2020

Optimized Energy Network Operations

- Optimization models for short-term (operational), mid-term (planning), long-term (investment) models of multi-energy systems
- Evaluation of how to achieve stability when integration of renewables and power-to-gas happens : Value of Flexibilities

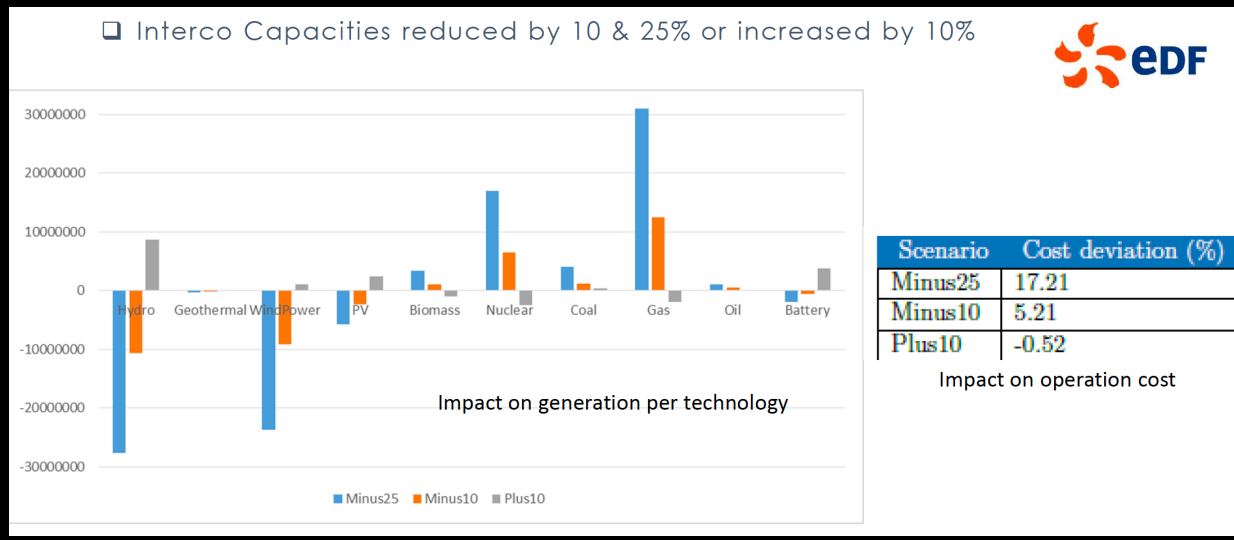

Scenario	Cost deviation (%)
Minus25	17.21
Minus10	5.21
Plus10	-0.52

Impact on operation cost

https://www.plan4res.eu/

plan4res

HOLISTIC DATACENTER & OPERATIONAL EFFICIENCY - 4 PILLAR FRAMEWORK

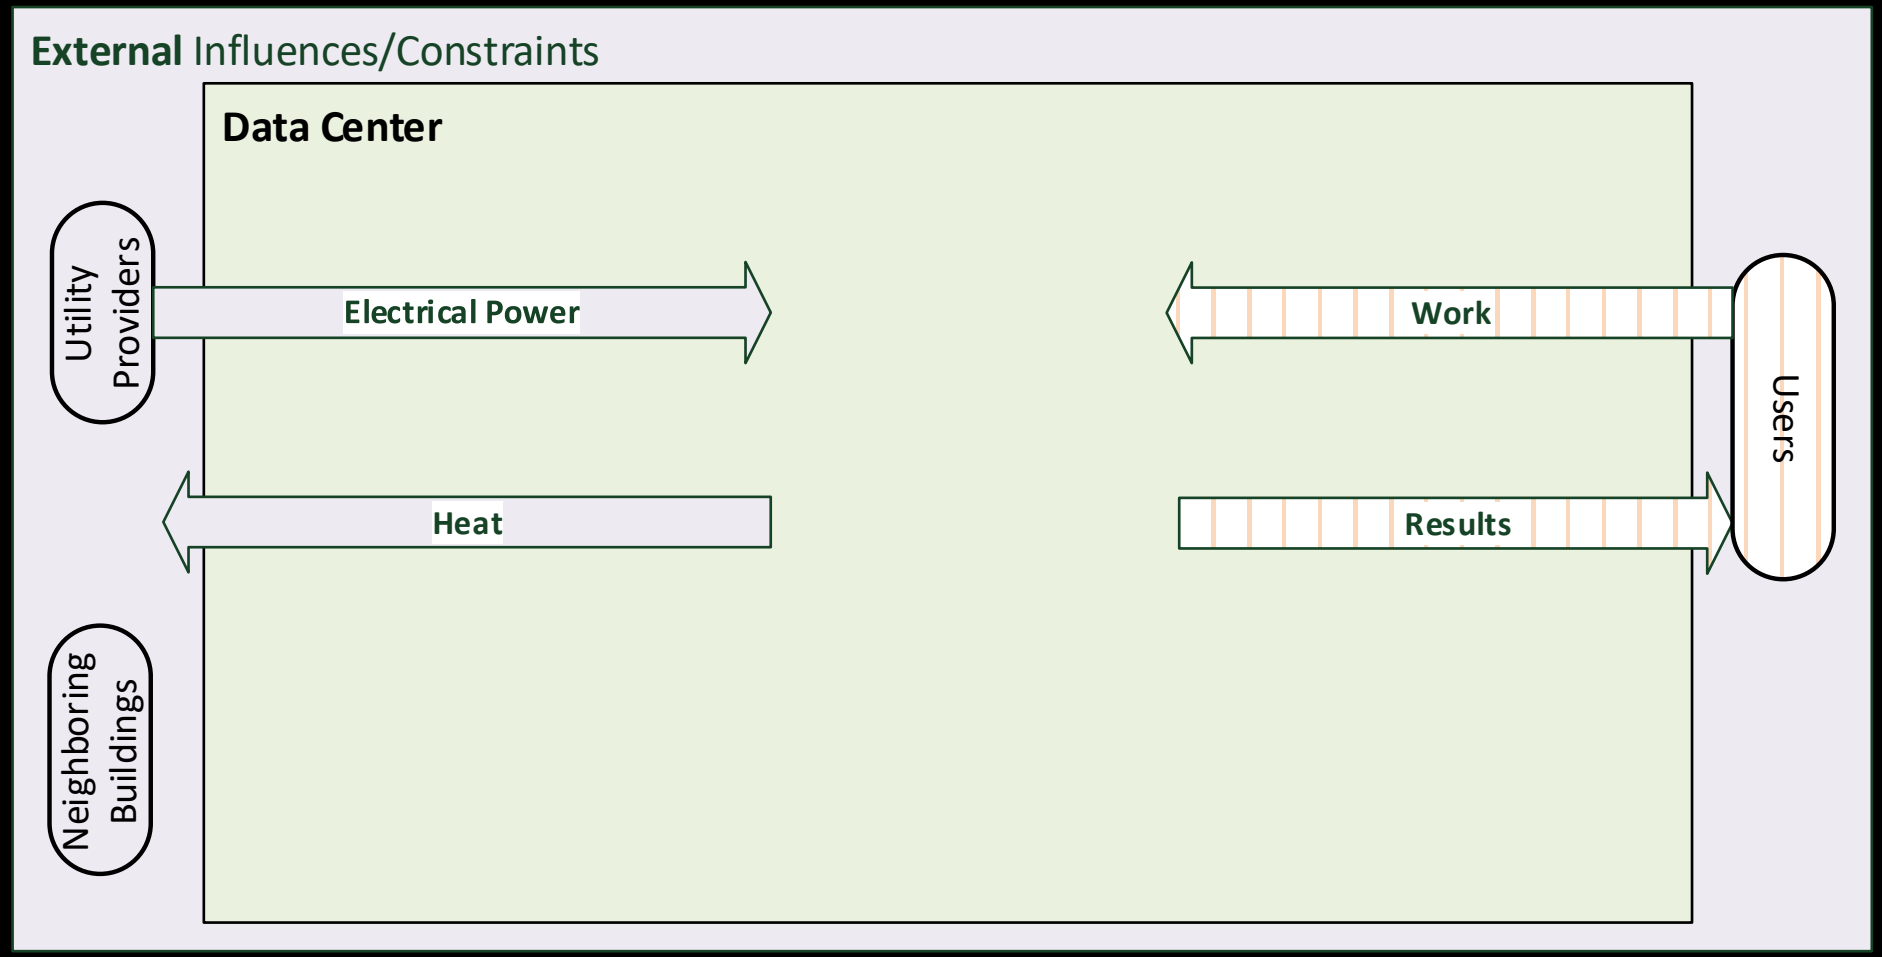

Torsten Wilde: Assessing the Energy Efficiency of High Performance Computing (HPC) Data Centers, Dissertation, Technical University Munich (TUM), 2018, http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss-20180326-1399734-1-7

HOLISTIC DATACENTER & OPERATIONAL EFFICIENCY - 4 PILLAR FRAMEWORK

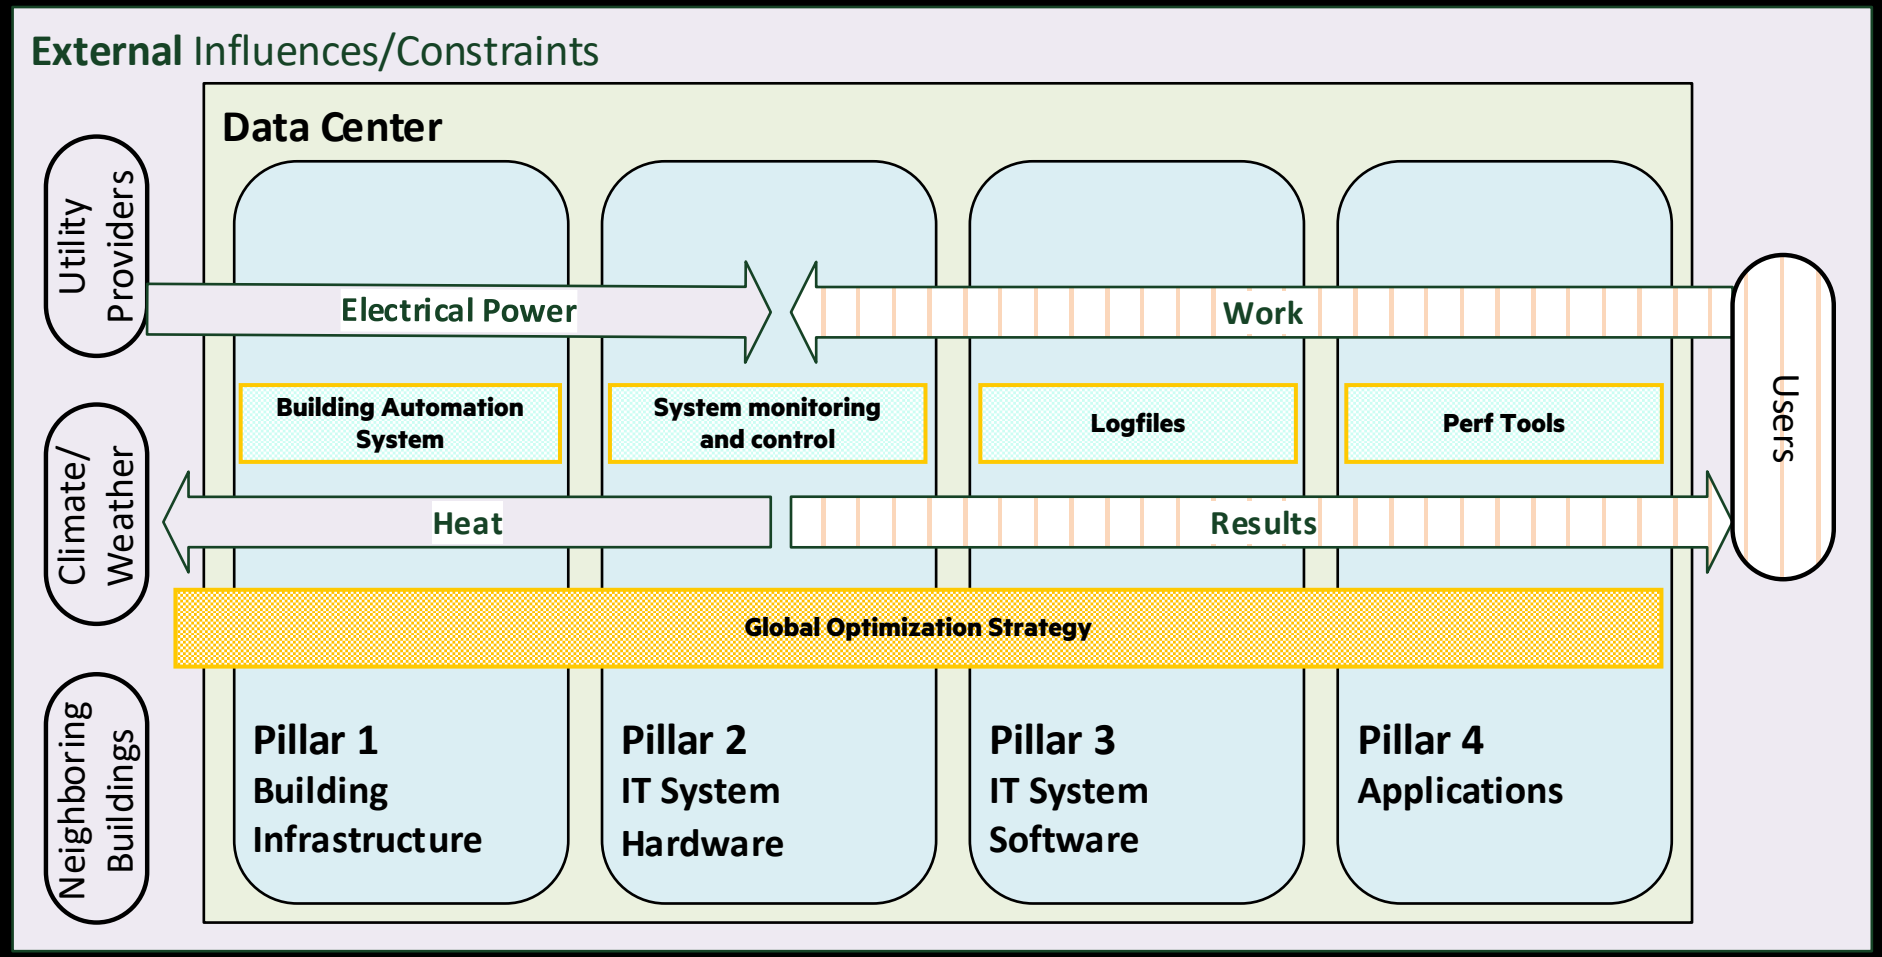

HOLISTIC DATACENTER & OPERATIONAL EFFICIENCY - 4 PILLAR FRAMEWORK

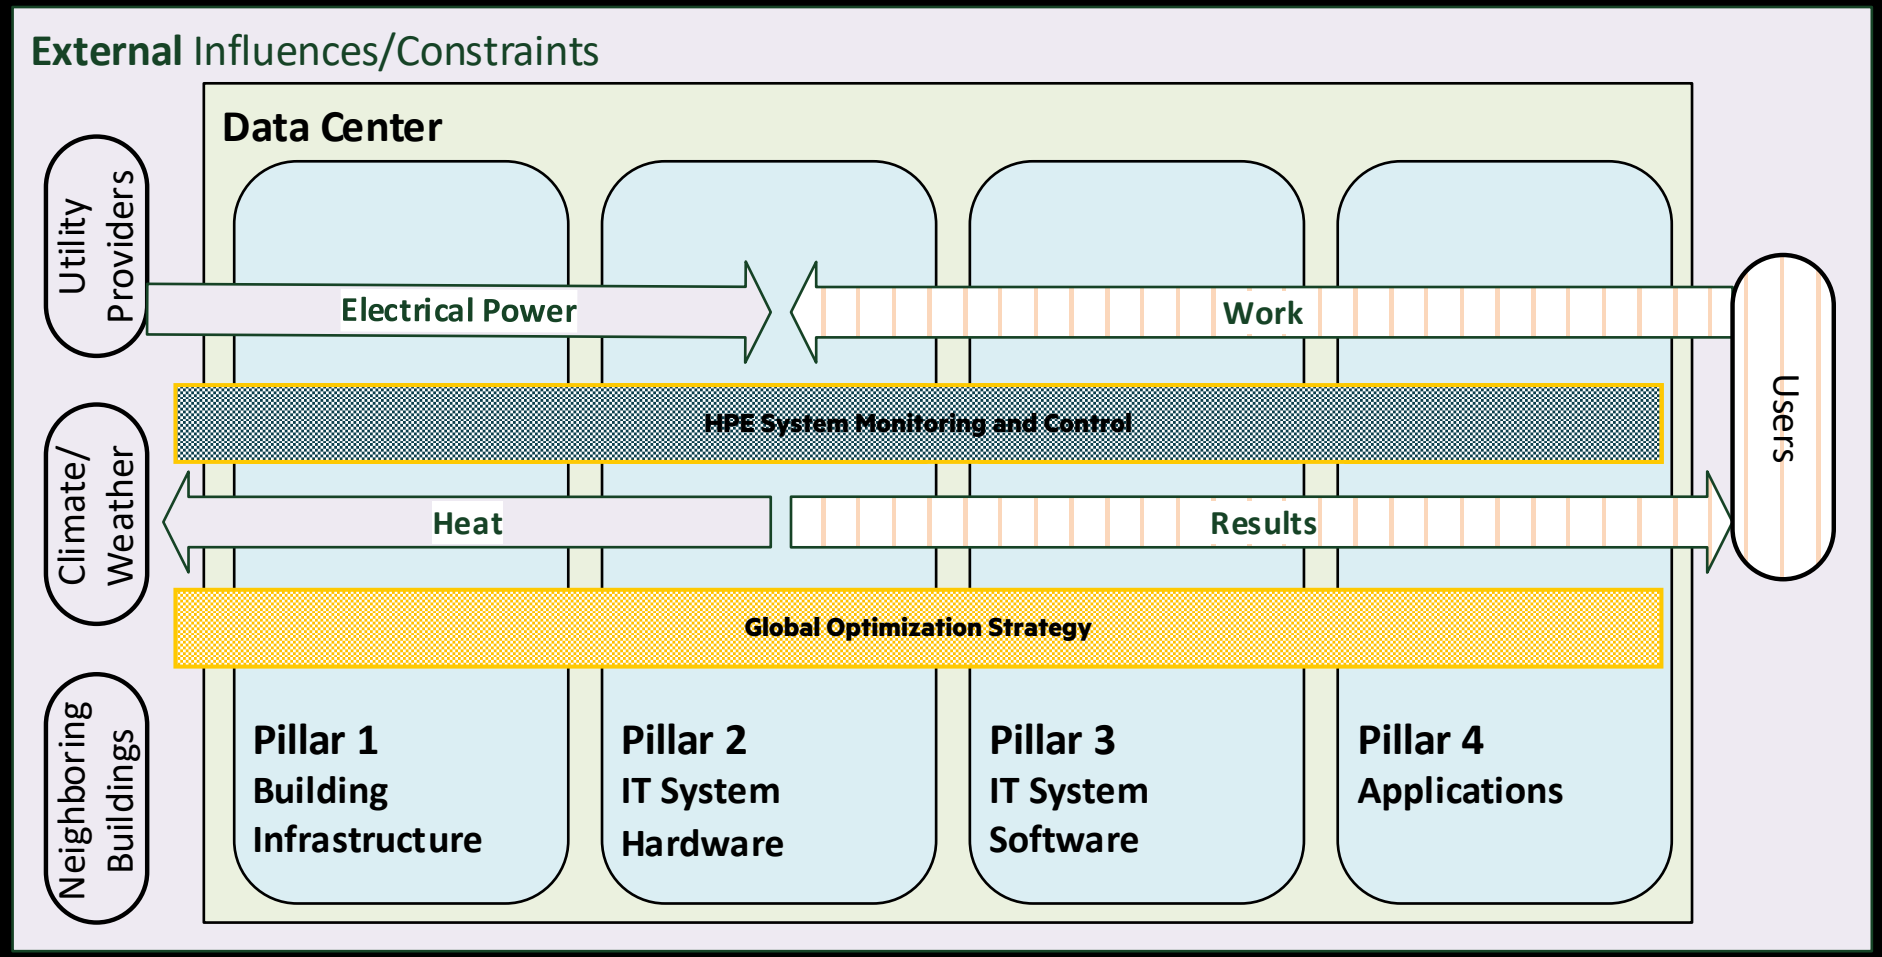

NODE HISTOGRAM POWER VARIABILITY SUPERMUC PHASE1 (INTEL SANDY BRIDGE-EP XEON E5-2680 8C) – ONE ISLAND (514 NODES)

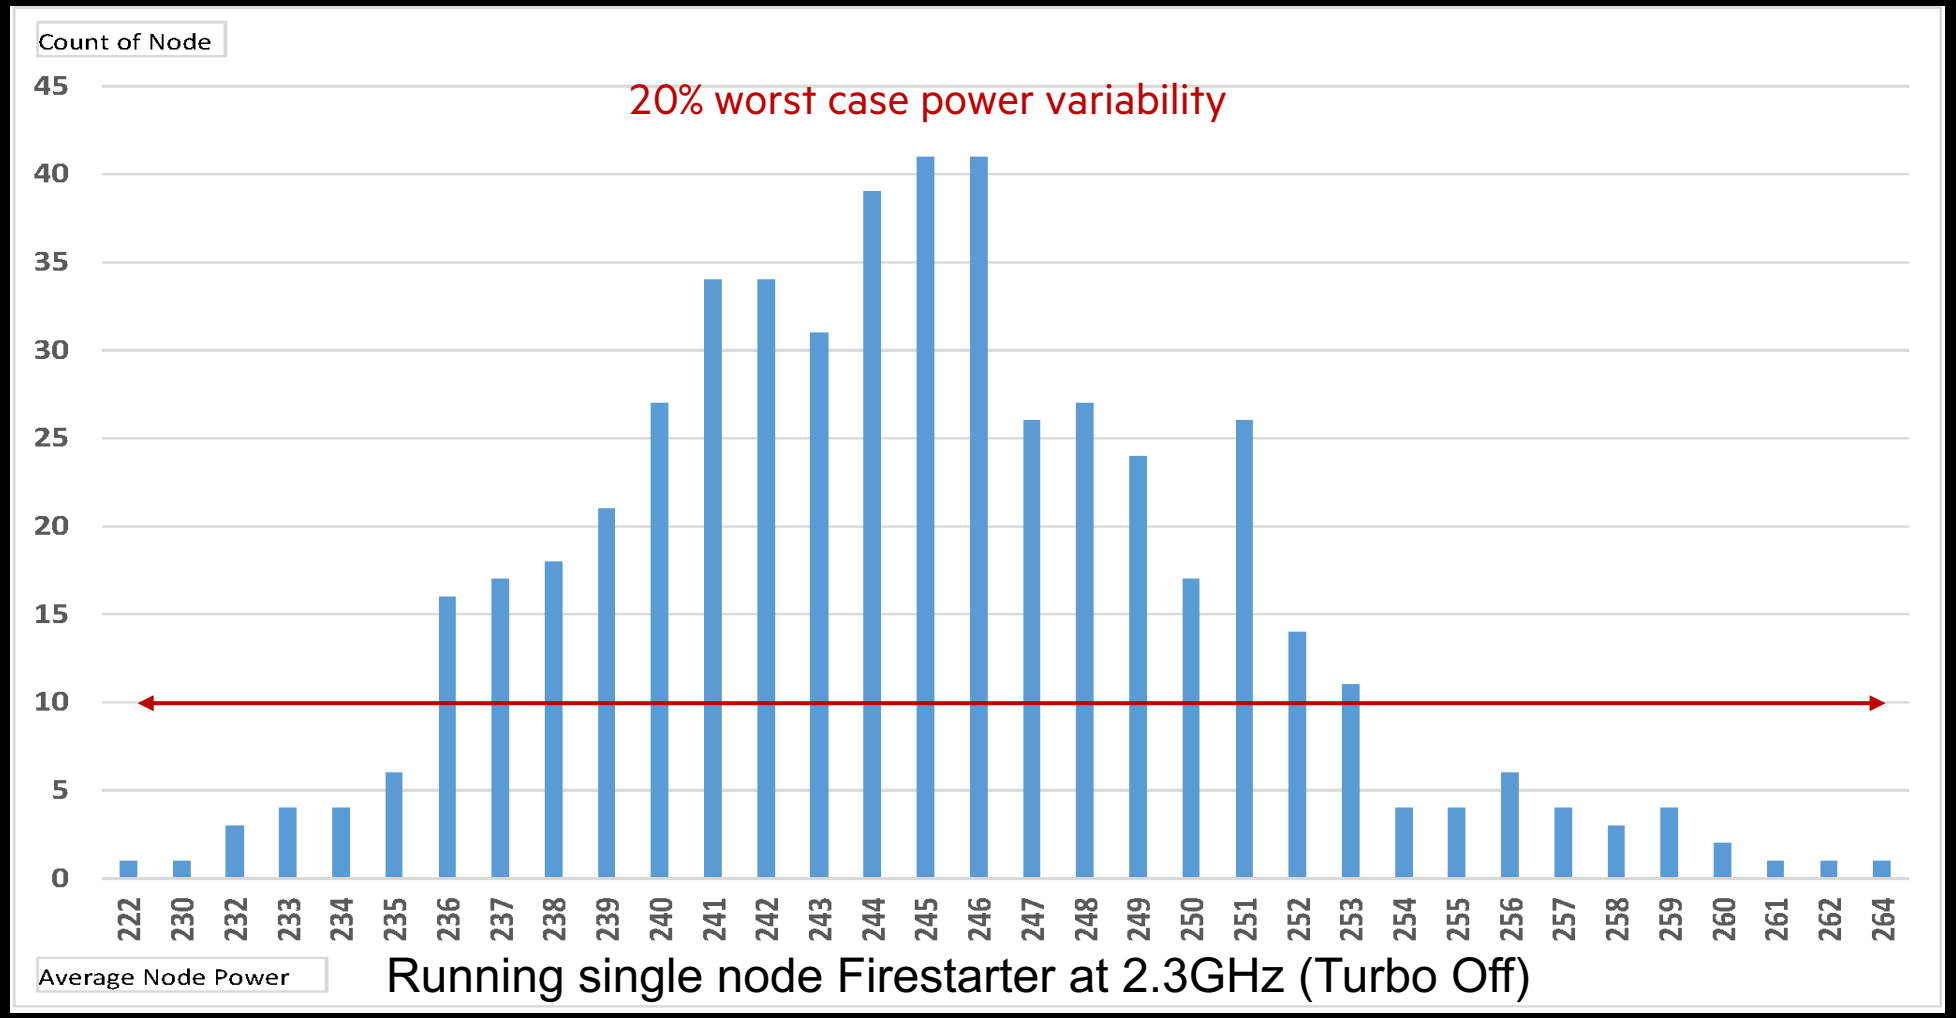

Wilde T., Auweter A., Shoukourian H., Bode A. (2015) Taking Advantage of Node Power Variation in Homogenous HPC Systems to Save Energy. In: High Performance Computing. ISC High Performance 2015. Lecture Notes in Computer Science, vol 9137. Springer, Cham

RUNNING NAS PARALLEL BENCHMARKS WITH DIFFERENT POWER CAPS ON AMD EPYC 7702

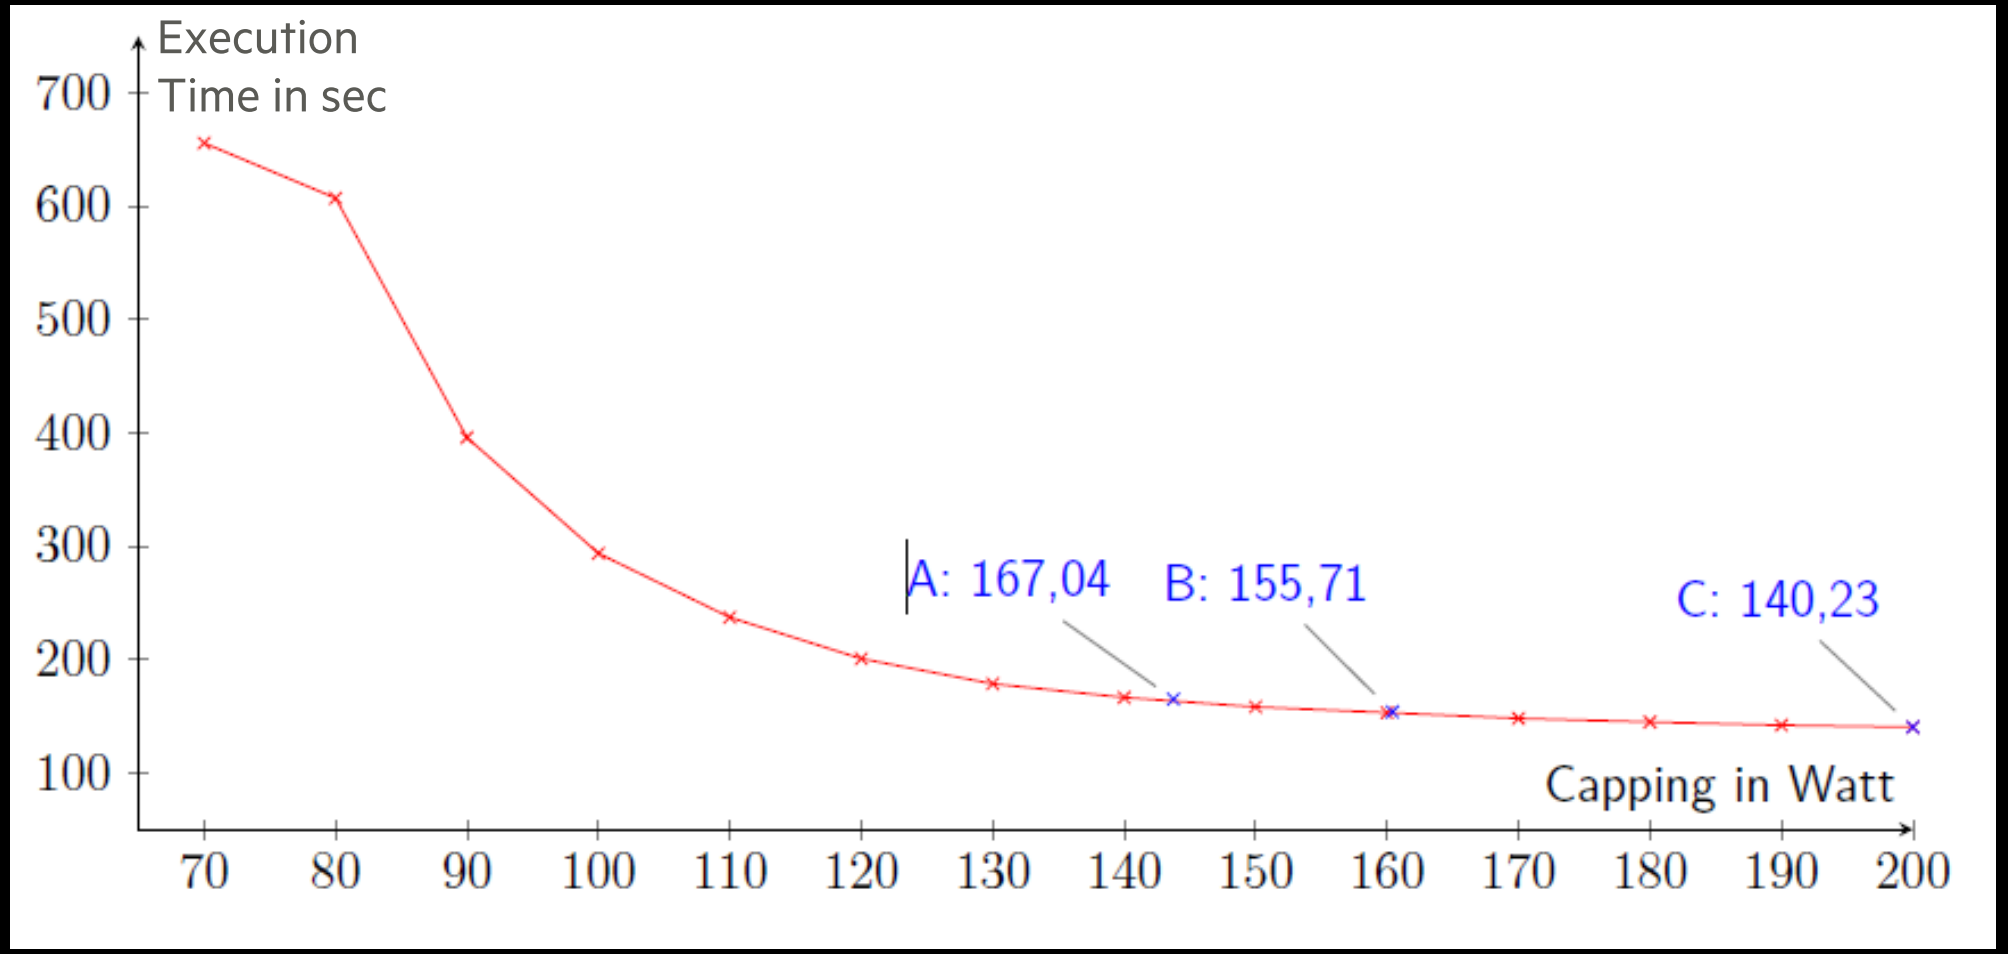

ORNL SUMMIT AVERAGE POWER BEHAVIOR*

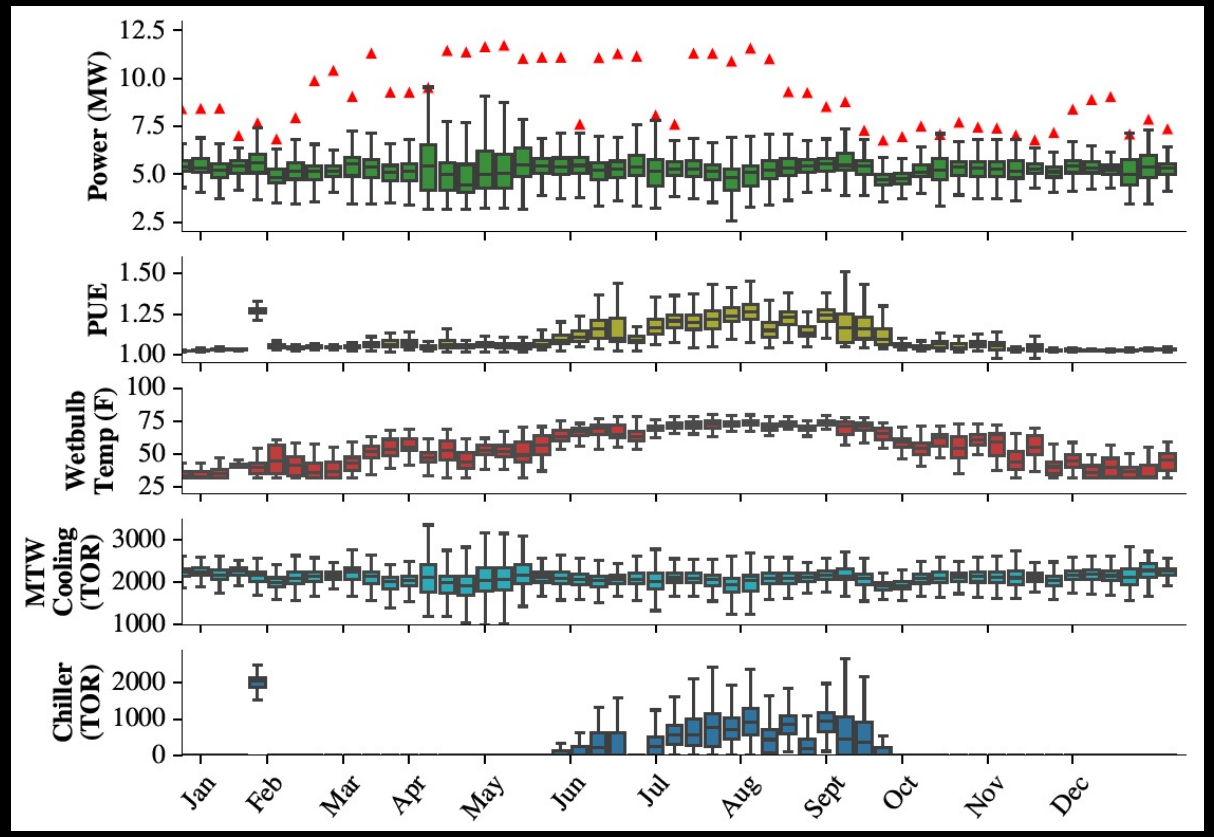

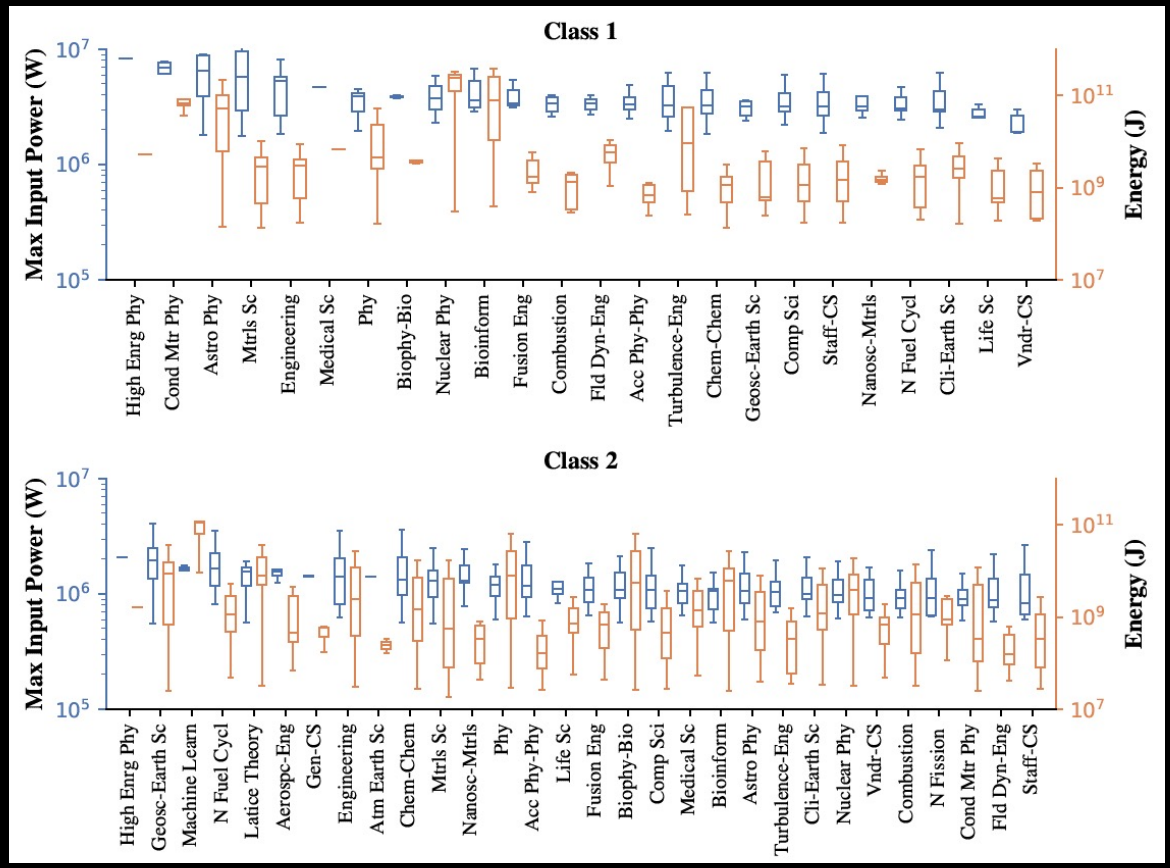

Average power consumption was between 5MW and 6MW with a constant small percentage of extremes that touches both the system idle (2.5MW) and peak (13MW) power consumption throughout the year."

*Woong Shin, Vladyslav Oles, Ahmad M. Karimi, J. Austin Ellis, Feiyi Wang, "Revealing Power, Energy and Thermal Dynamics of a 200PF Pre-Exascale Supercomputer", SC'21, best paper

DATA MOVEMENT MIDDLEWARES

Data movement is expensive

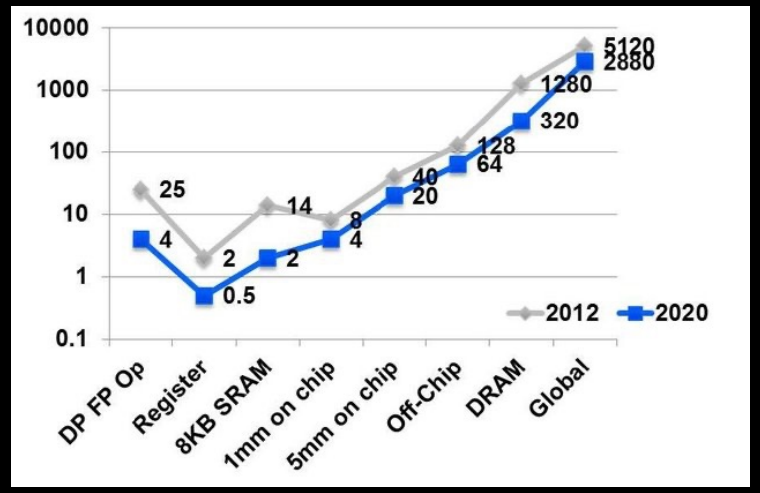

Energy cost of data movemen:

pJ per 64-bit FP op

Leland et al, SAND2016-9583

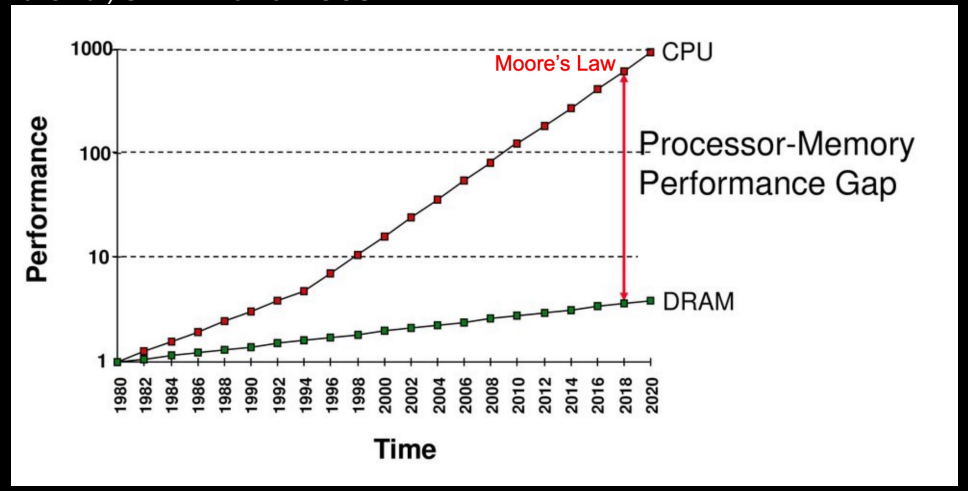

Patterson, UC Berkeley

Need data object abstraction

- Across applications, across memory tiers, across compute nodes

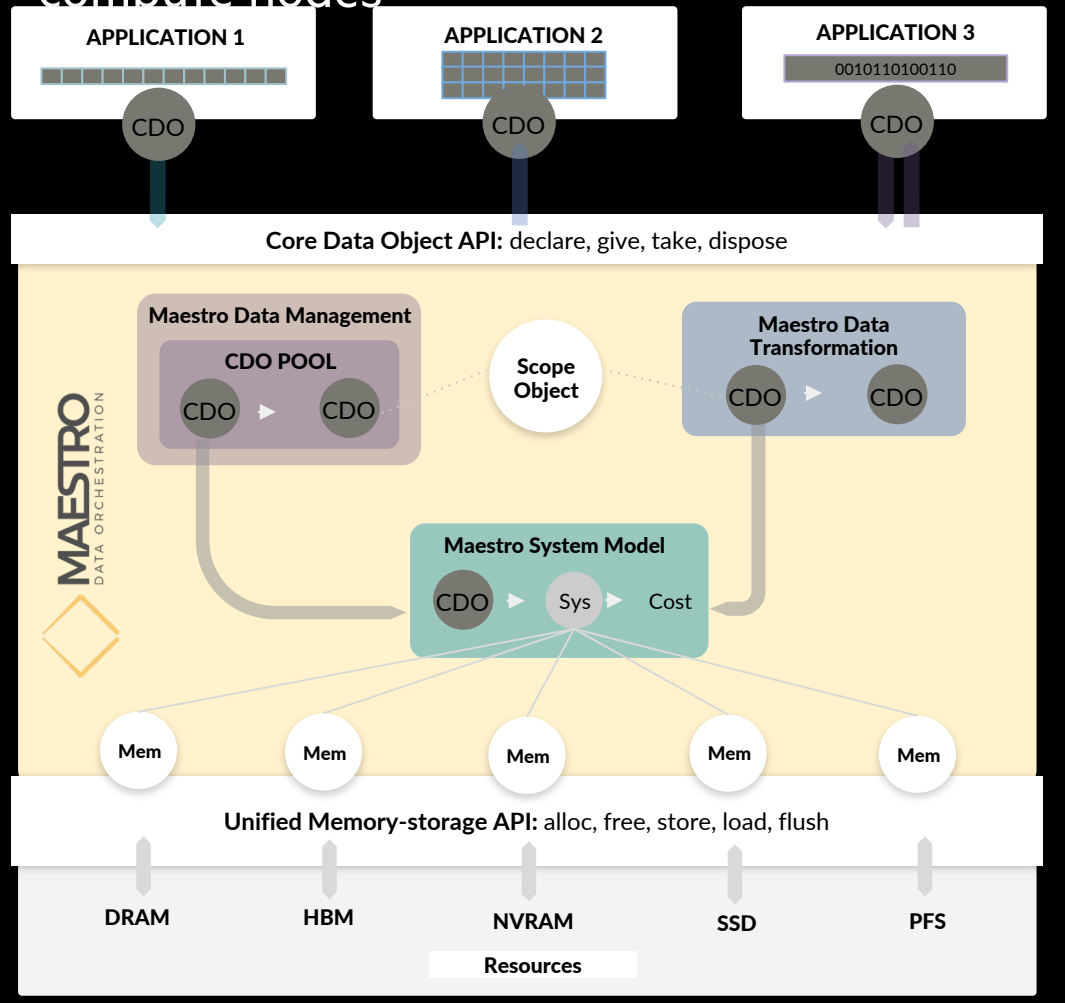

https://maestro-data.eu/

4 FACETS OF ENERGY EFFICIENCY

Optimized Job Performance under resource constraints

- Customers want to run hardware over-provisioned systems with for better overall system performance
- Need: Balance between available power/cooling and workload performance

Data Center Sustainability

- Governments (US, EU) are developing mandates for our customers to address sustainability aspects
- Need: bring down current energy usage and carbon footprint, optimize system operation according to data center TCO (balance facility efficiency with system operation)

Minimize Energy Consumption

- Customers want to reduce OPEX due to increased power needs of new technologies and increased energy prices
- Need: reduce energy consumption of workloads according to a TtS / EtS tradeoff metric

Maximize Resource Utilization

- Customers need to optimize power and cooling needs to support sustainable HPC efforts
- Need: optimized use of available resources, minimize stranded capabilities (power, cooling) in datacenter and HPC system

A WORD FROM YOUR SPONSOR: ERL IN 1 SLIDE

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 and 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

A research team in the CTO office of the HPC, AI & Labs business unit.

THANK YOU

uhaus@hpe.com