HPC4 ENERGYINNOVATION

HPC4 MATERIALS

HPC4 MANUFACTURING

National Laboratories Partner with U.S. Industry to Increase Innovation and Energy Efficiency

Aaron Fisher
Acting Director HPC4EnergyInnovation

U.S. DEPARTMENT OF ENERGY Energy Efficiency & Renewable Energy

U.S. DEPARTMENT OF ENERGY Office of Fossil Energy

U.S. DEPARTMENT OF ENERGY Office of Science

AMES LABORATORY · Argonne National Laboratory · Berkeley Lab · INL Idaho National Laboratory · Lawrence Livermore National Laboratory · Los Alamos National Laboratory · National Energy Technology Laboratory · NREL · Oak Ridge National Laboratory · Pacific Northwest National Laboratory · Sandia National Laboratories

HPC4EI connects the DOE HPC ecosystem to US industry

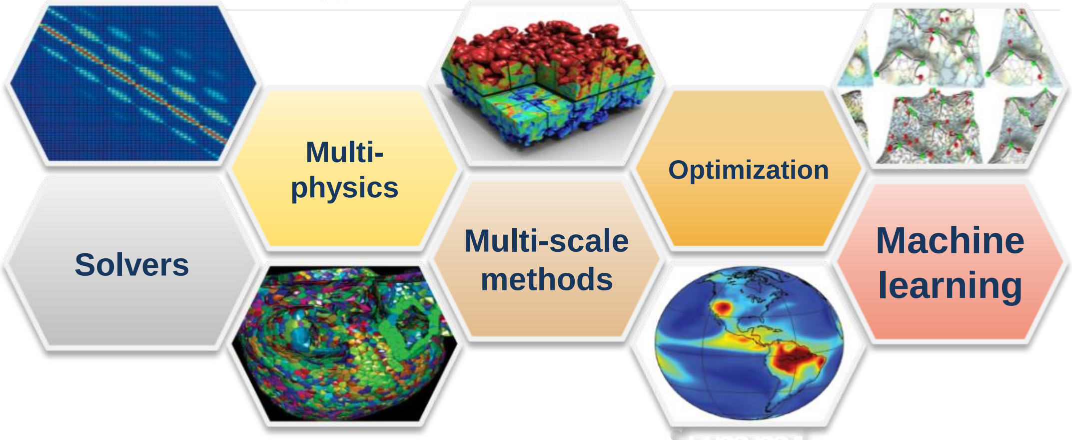

World-class supercomputers

Unique software

Subject Matter Experts

https://hpc4energyinnovation.llnl.gov/

HPC4EI has established an efficient process for DOE Labs to partner with industry


HPC4 MANUFACTURING

Using Machine Learning Tools for Fast, Accurate, and Real-Time Process Control in Plate Glass Manufacturing

Livermore is working with Vitro Glass Company to apply deep learning tools to glass furnace operation.

Manufacturing Glass Sheets Using Next-Generation Process Control

Used in the production of plate glass—a type commonly used in windows and doors—a glass melt furnace performs multiple functions such as melting and homogenizing raw materials, as well as stripping gases and impurities from the melt. Small variances in the process can shift the location in the furnace and how efficiently these functions occur, which impacts the quality of the final product.

The first step in modern process control for glass melt furnaces consists of setting the furnace temperatures based on complex computation fluid dynamics (CFD) models that simulate the process. When the measured furnace temperatures deviate from desired temperatures, the CFD code is run again to tell operators how to adjust the furnace set points—a process that currently takes up to two weeks. Engineers at Vitro Glass (formerly known as PPG Glass) are interested in reducing the time between measuring the deviation in the process control points and making the corrections to the parameters for operating furnaces. They posed this problem to researchers at Lawrence Livermore National Laboratory (LLNL) and asked for help to solve it.

To address the challenge, researchers at LLNL proposed to use VisIt—a premier, open source visualization and analysis tool for high performance computing (HPC) environments—along with machine learning methods to develop a reduced-order glass furnace model so plant engineers can make informed process adjustments in real time. The machine learning algorithm is much less computationally intensive than the original CFD model, and can be used as a real-time furnace control system running on a desktop computer instead of a larger supercomputer. The new control system based on machine

The VisIt tool produces advanced visualizations such as this one, which shows tracer flows in a glass furnace.

learning can instantly inform the furnace operator what new set points to use when temperatures deviate from normal. Through a project funded by the Department of Energy High Performance Computing for Manufacturing (HPC4Mfg) program—which aims to provide supercomputing resources to industry partners—LLNL is working with Vitro Glass to demonstrate a

Lawrence Livermore National Laboratory

- Two solicitations a year
  - AMMTO, IEDO, FECM
- $400k / Project
- ~$6M /yr Budget
- 11 participating National laboratories
- 150+ projects funded with 90+ companies over 7 years

Industry Concept Paper → Identify Lab PI → Submit Full Proposal → Technical Review → Award → PI Secures HPC Resources → Execute

Visit our website for solicitation details

FOA is expected soon!

Topics

…

- Improvements in semiconductor technologies that will result in operational energy efficiency improvements, supporting achieving the goals of EES2

…

Questions can be sent to hpc4ei@llnl.gov

Join the hpc4ei-info@llnl.gov distribution lists via the web to receive program announcements

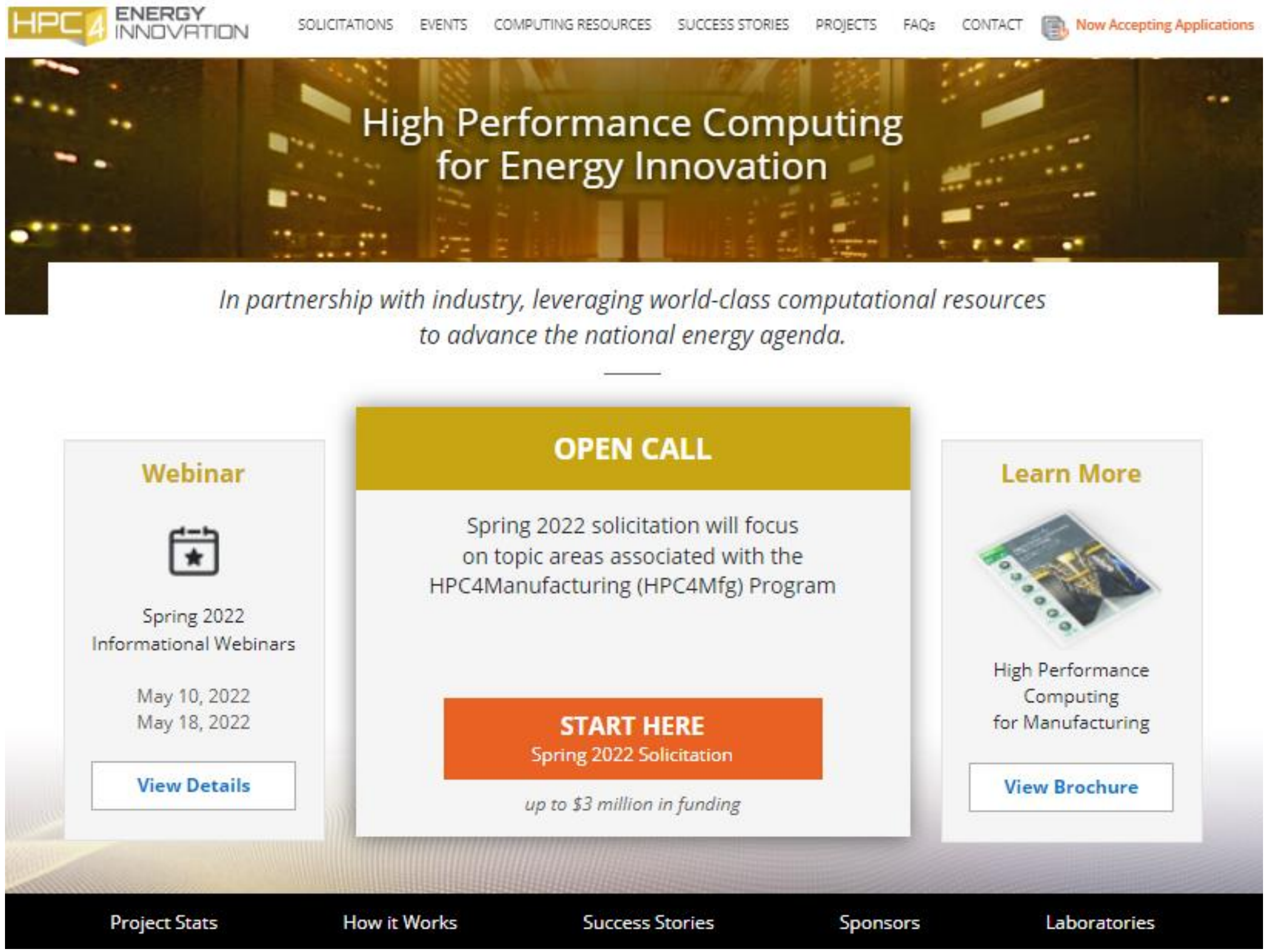

www.hpc4energyinnovation.org

Post-exascale System Architectures: Hardware overprovisioning and energy efficiency

EES2 Meeting

20 July 2023

Barry Rountree and Tapasya Patki
{rountree|patki1}@llnl.gov
Center for Applied Scientific Computing

LLNL-PRES-XXXXXX
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

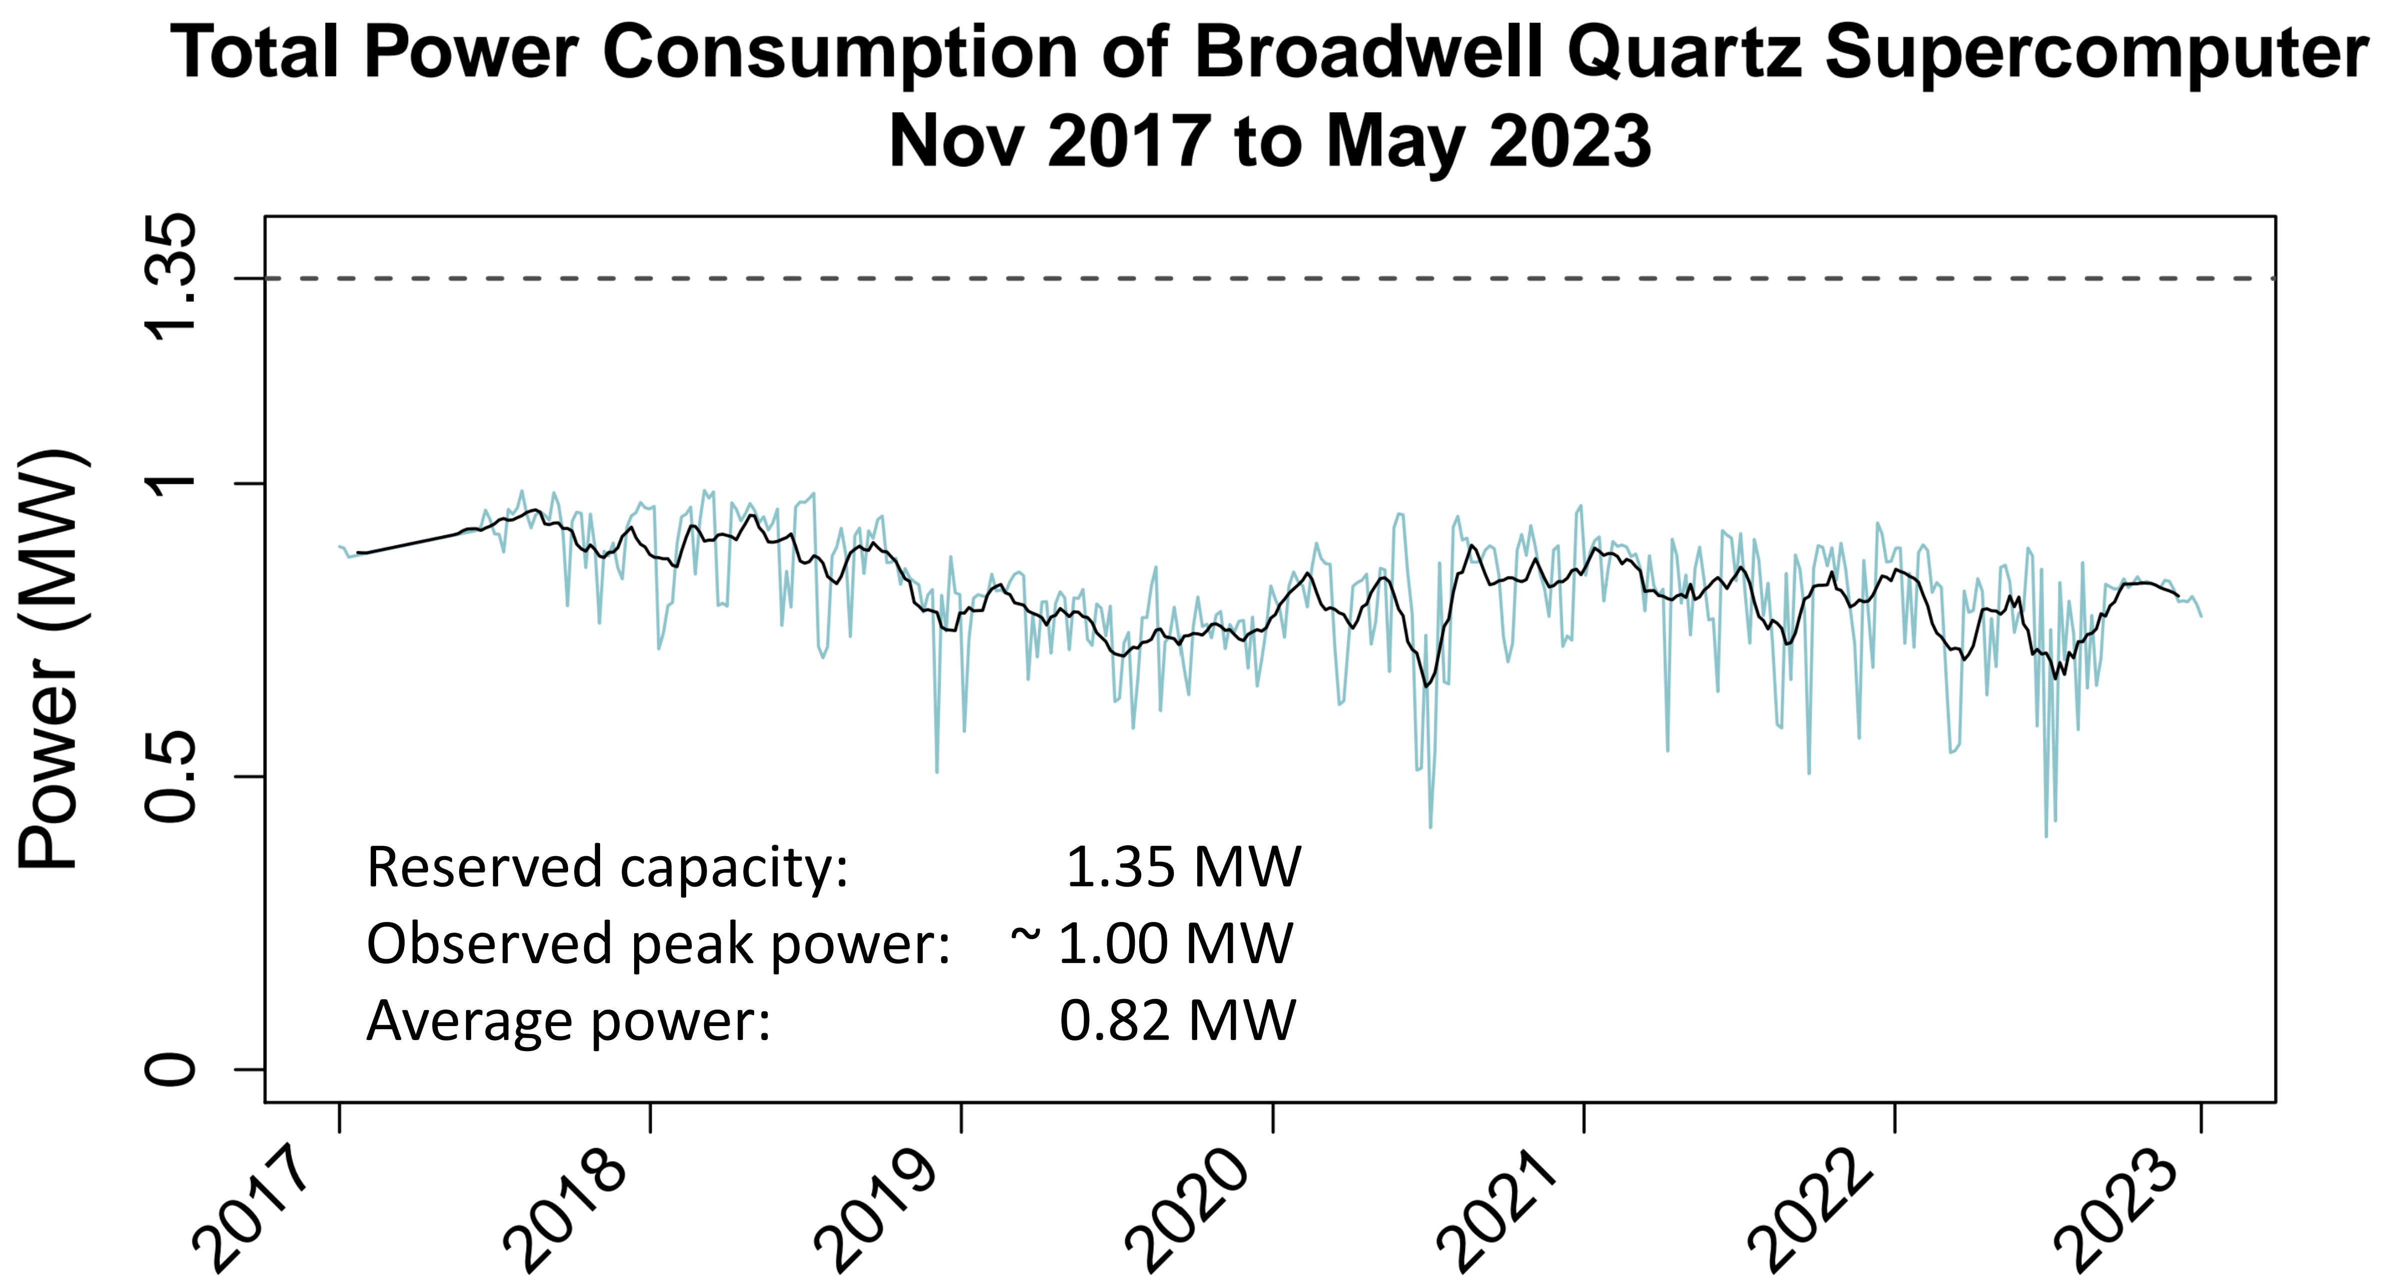

Total Power Consumption of Broadwell Quartz Supercomputer
Nov 2017 to May 2023
Power (MW)
1.35
1
0.5
0
Reserved capacity: 1.35 MW
Observed peak power: ~ 1.00 MW
Average power: 0.82 MW
2017
2018
2019
2020
2021
2022
2023

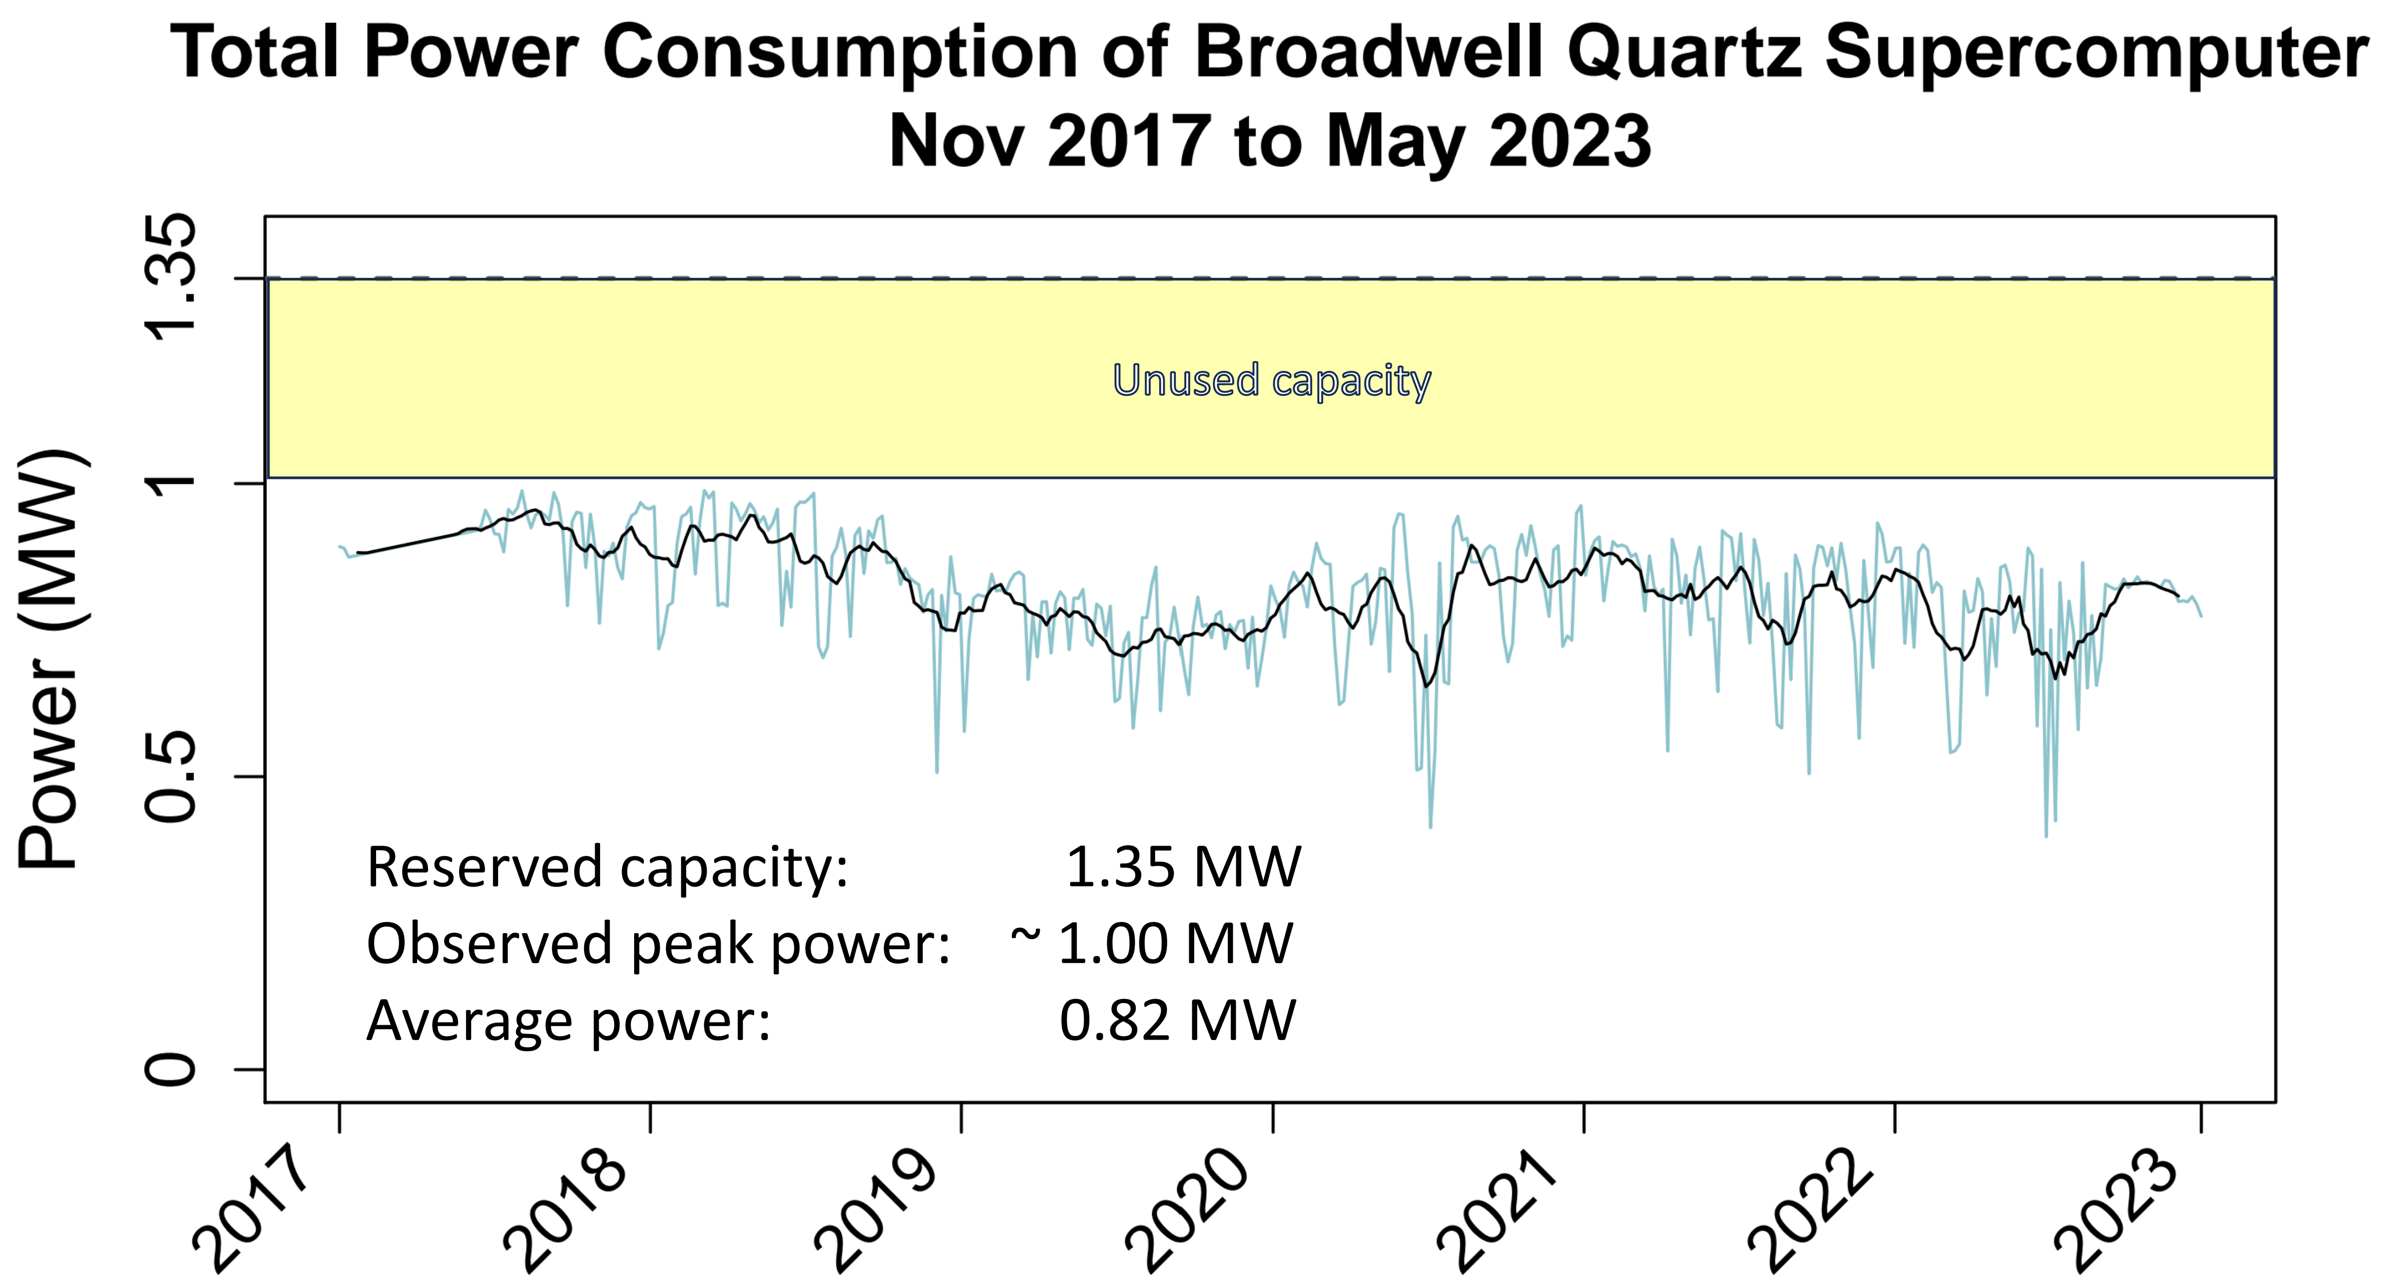
Total Power Consumption of Broadwell Quartz Supercomputer
Nov 2017 to May 2023
Power (MW)
1.35
1
0.5
0
Unused capacity
Reserved capacity: 1.35 MW
Observed peak power: ~ 1.00 MW
Average power: 0.82 MW
2017
2018
2019
2020
2021
2022
2023

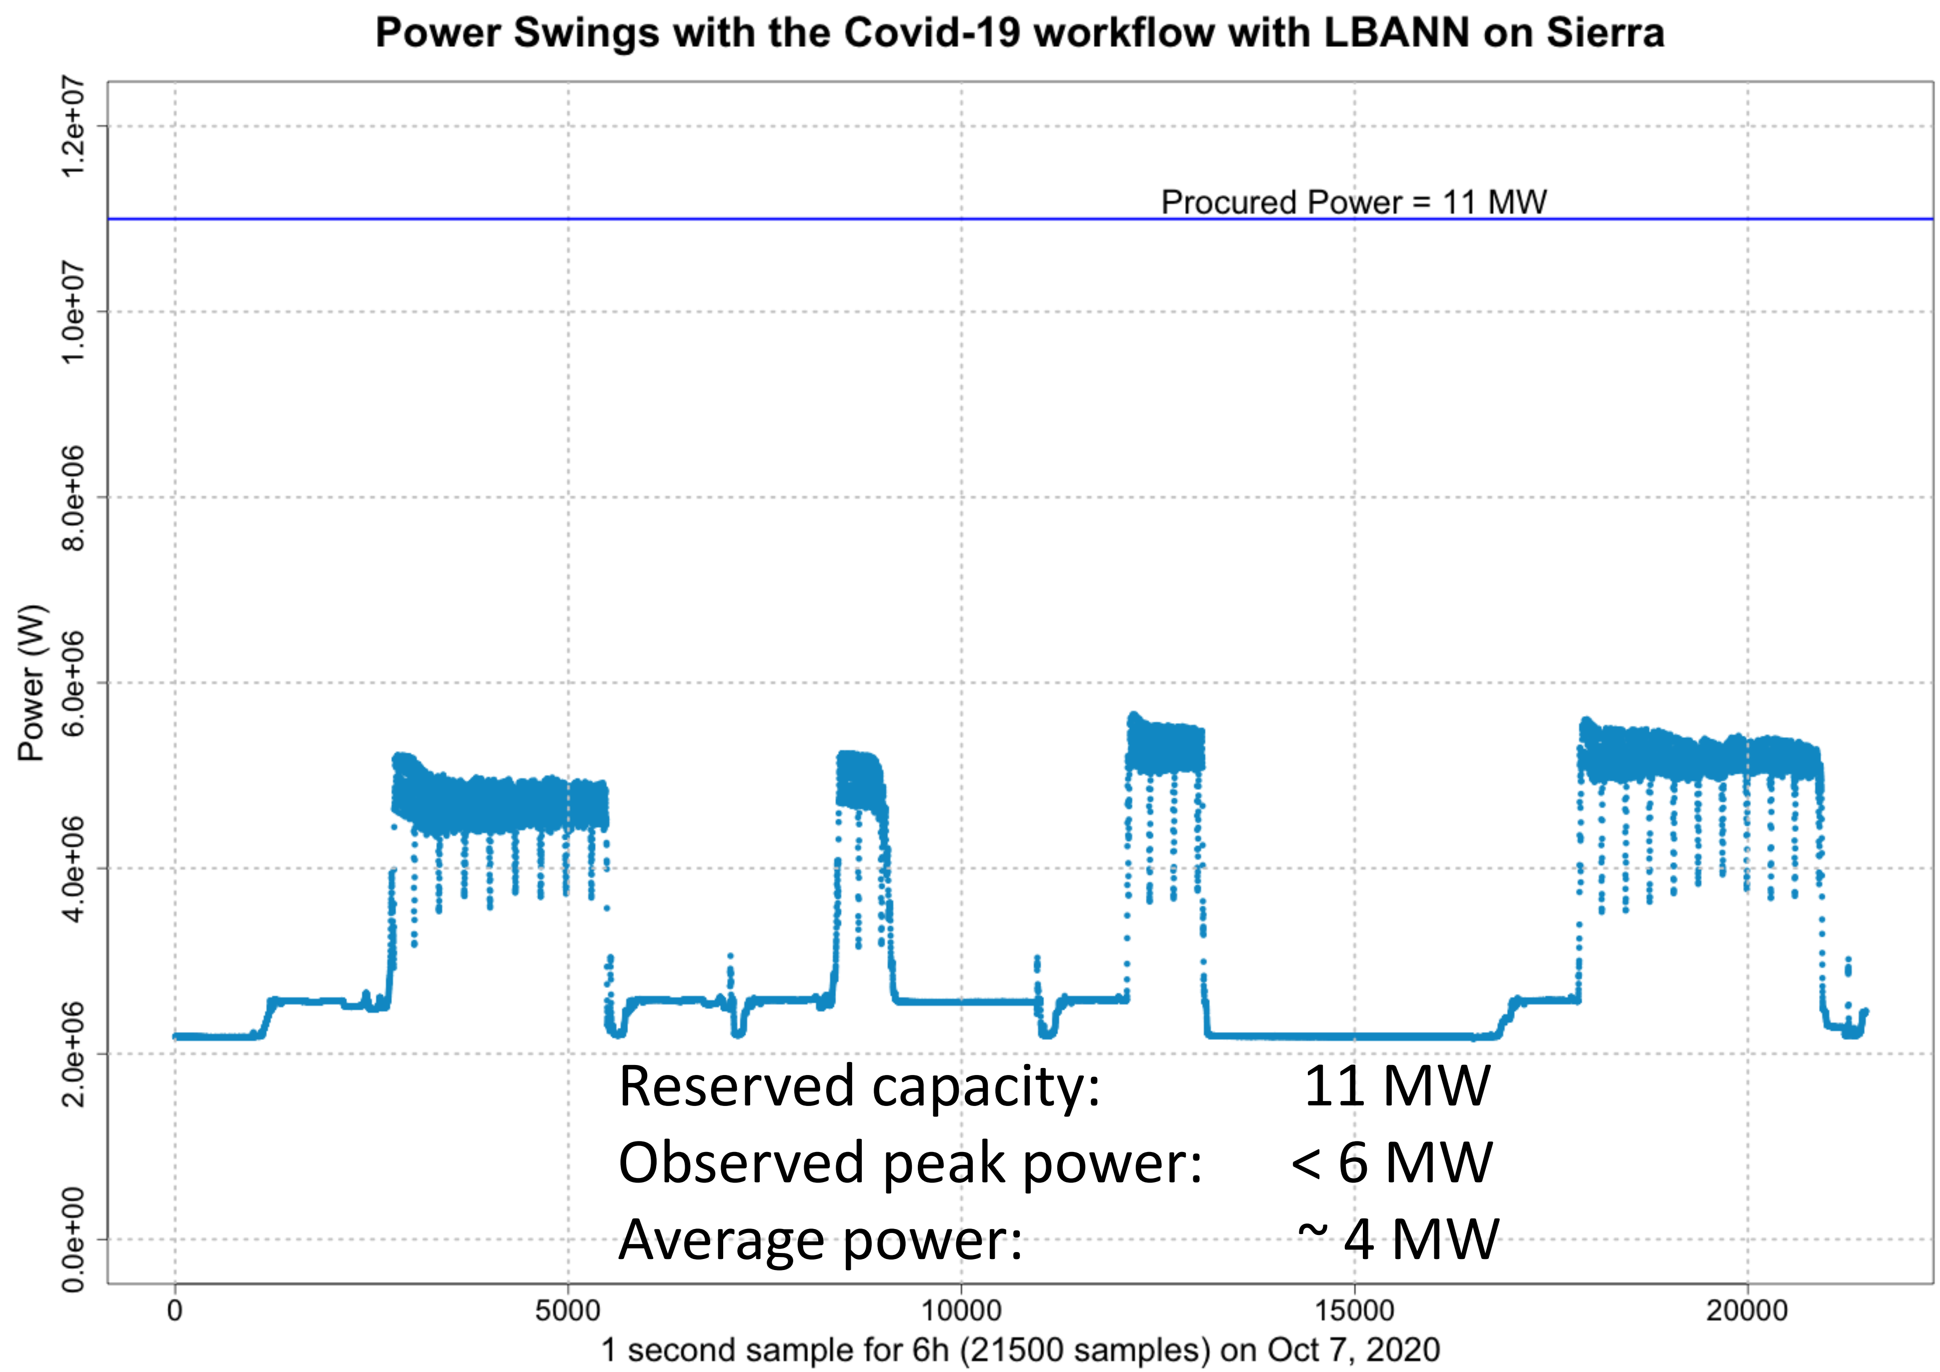
Power Swings with the Covid-19 workflow with LBANN on Sierra
Procured Power = 11 MW
Power (W)
0.0e+00
2.0e+06
4.0e+06
6.0e+06
8.0e+06
1.0e+07
1.2e+07
0
5000
10000
15000
20000
1 second sample for 6h (21500 samples) on Oct 7, 2020
Reserved capacity: 11 MW
Observed peak power: < 6 MW
Average power: ~ 4 MW

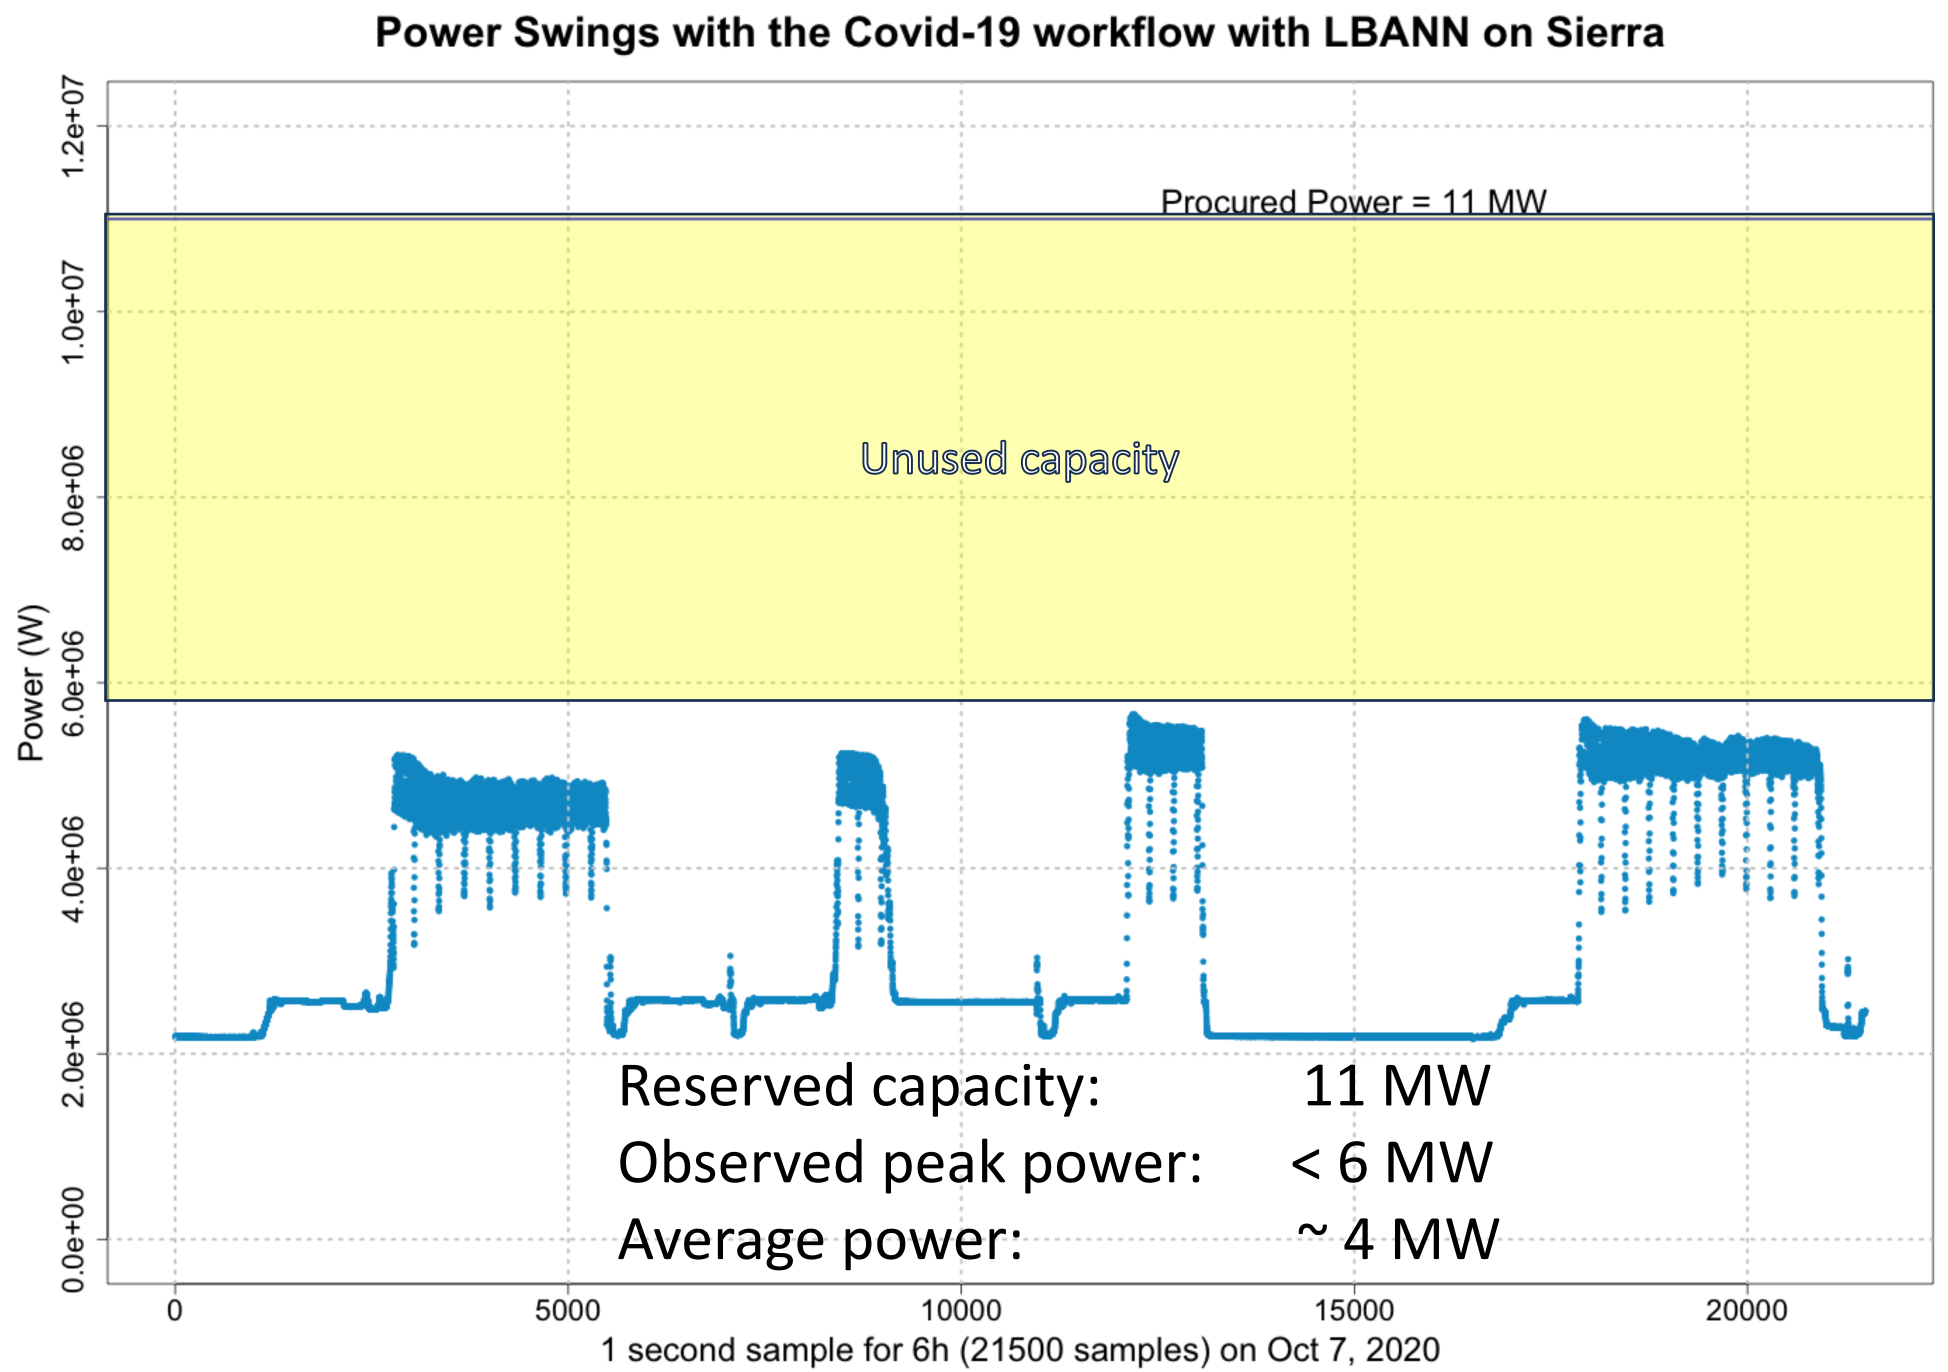
Power Swings with the Covid-19 workflow with LBANN on Sierra
Procured Power = 11 MW
Unused capacity
Power (W)
0.0e+00
2.0e+06
4.0e+06
6.0e+06
8.0e+06
1.0e+07
1.2e+07
0
5000
10000
15000
20000
1 second sample for 6h (21500 samples) on Oct 7, 2020
Reserved capacity: 11 MW
Observed peak power: < 6 MW
Average power: ~ 4 MW

Why we can't have more FLOPS:

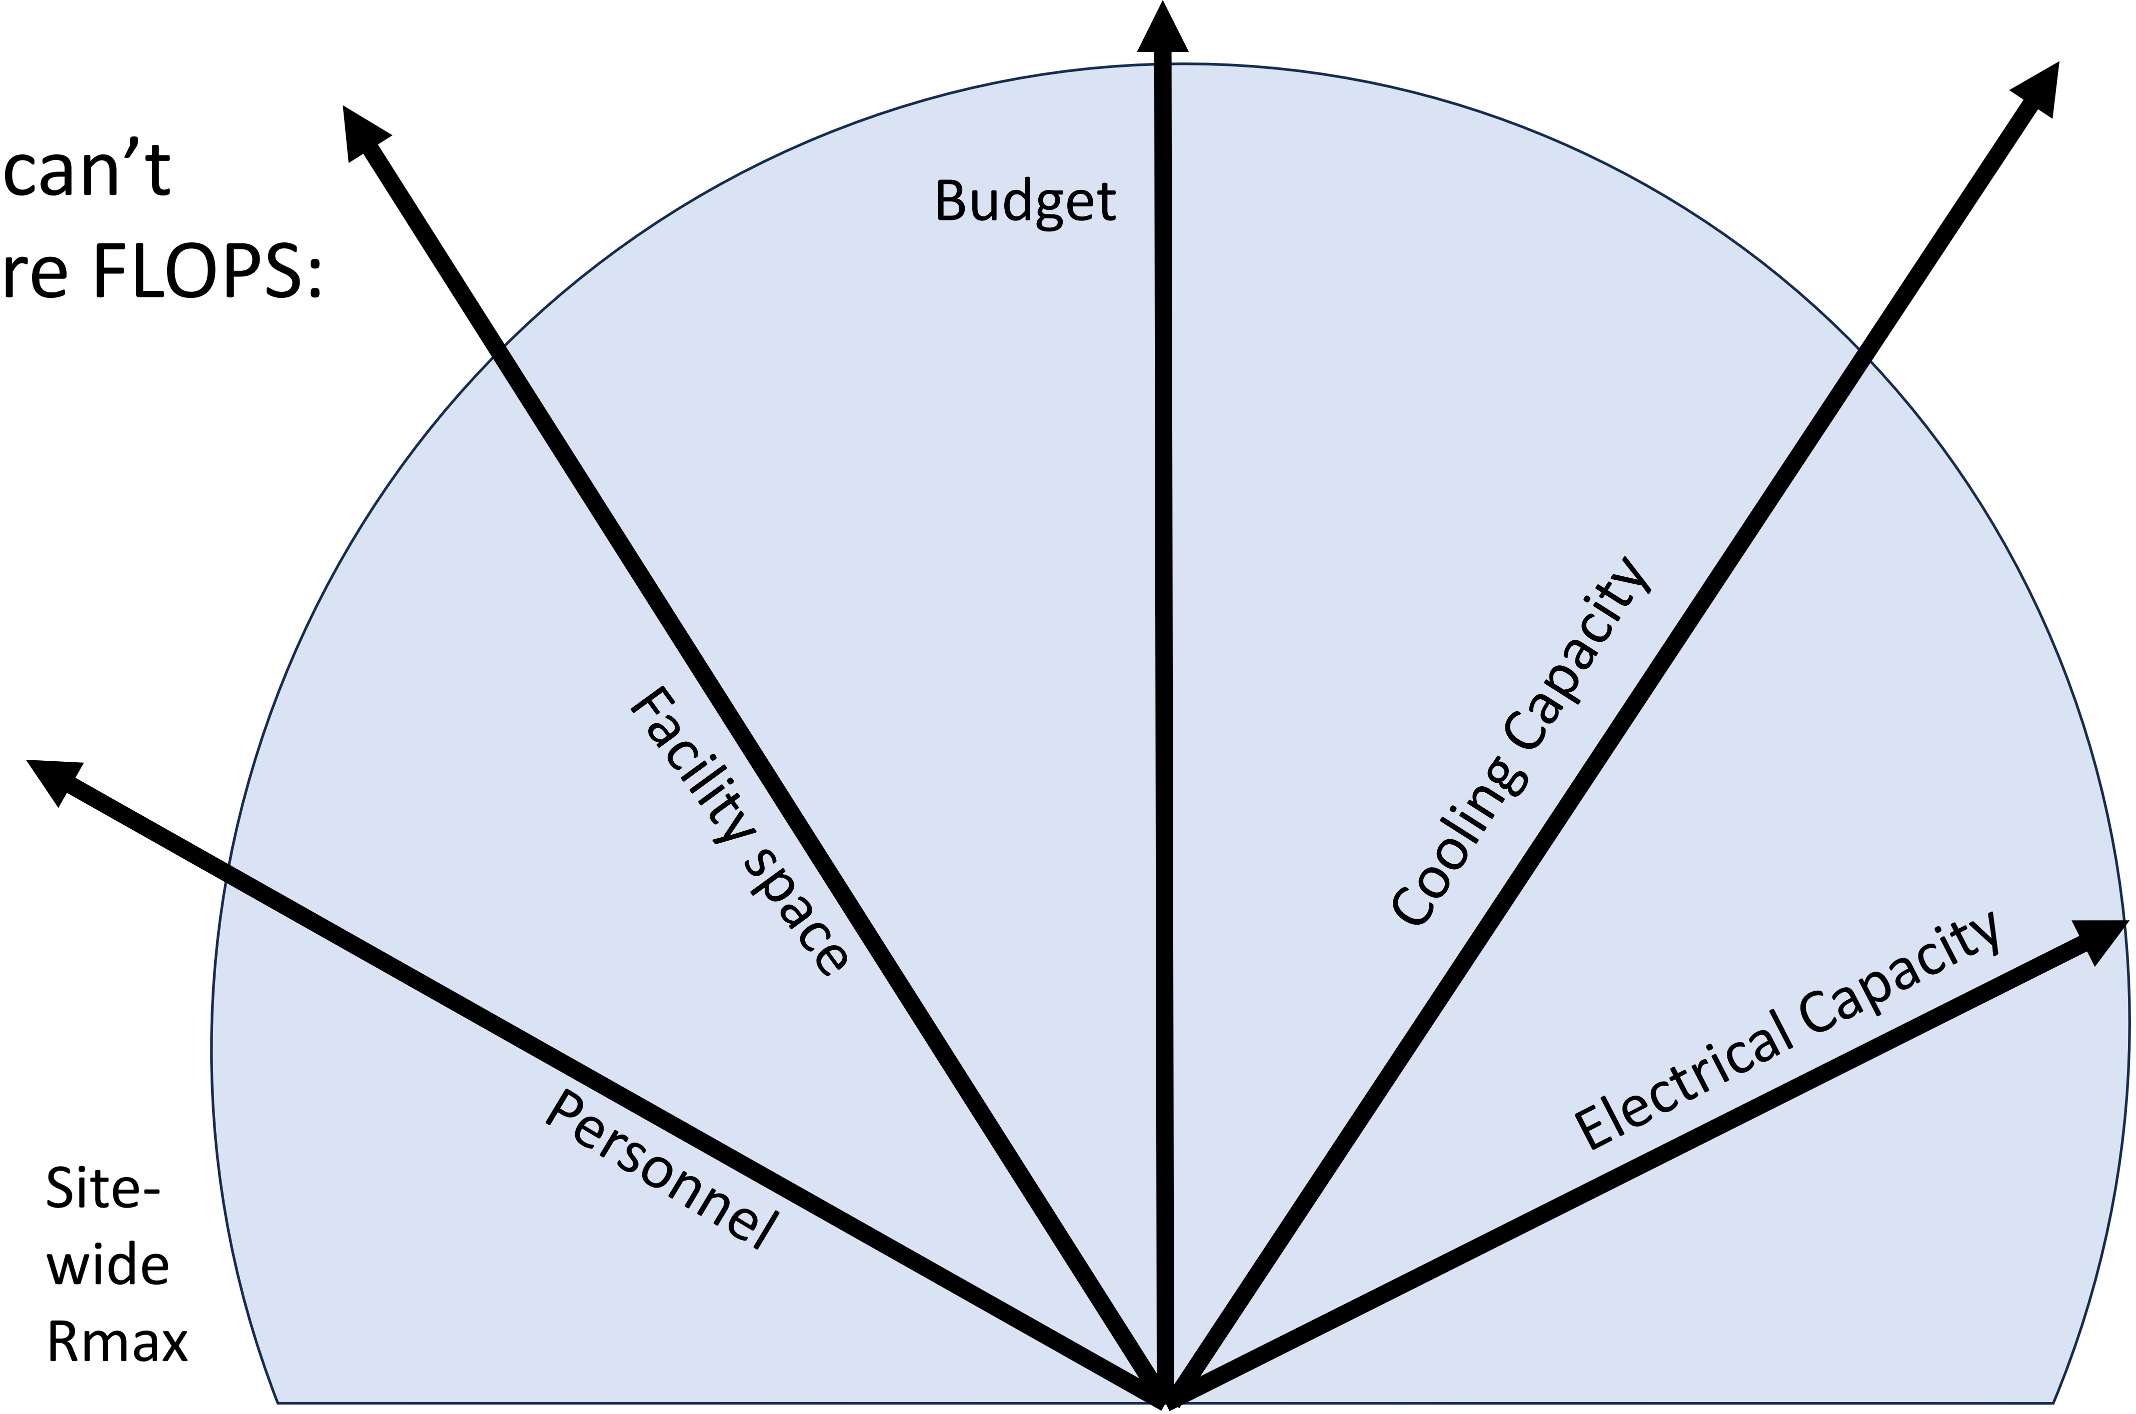

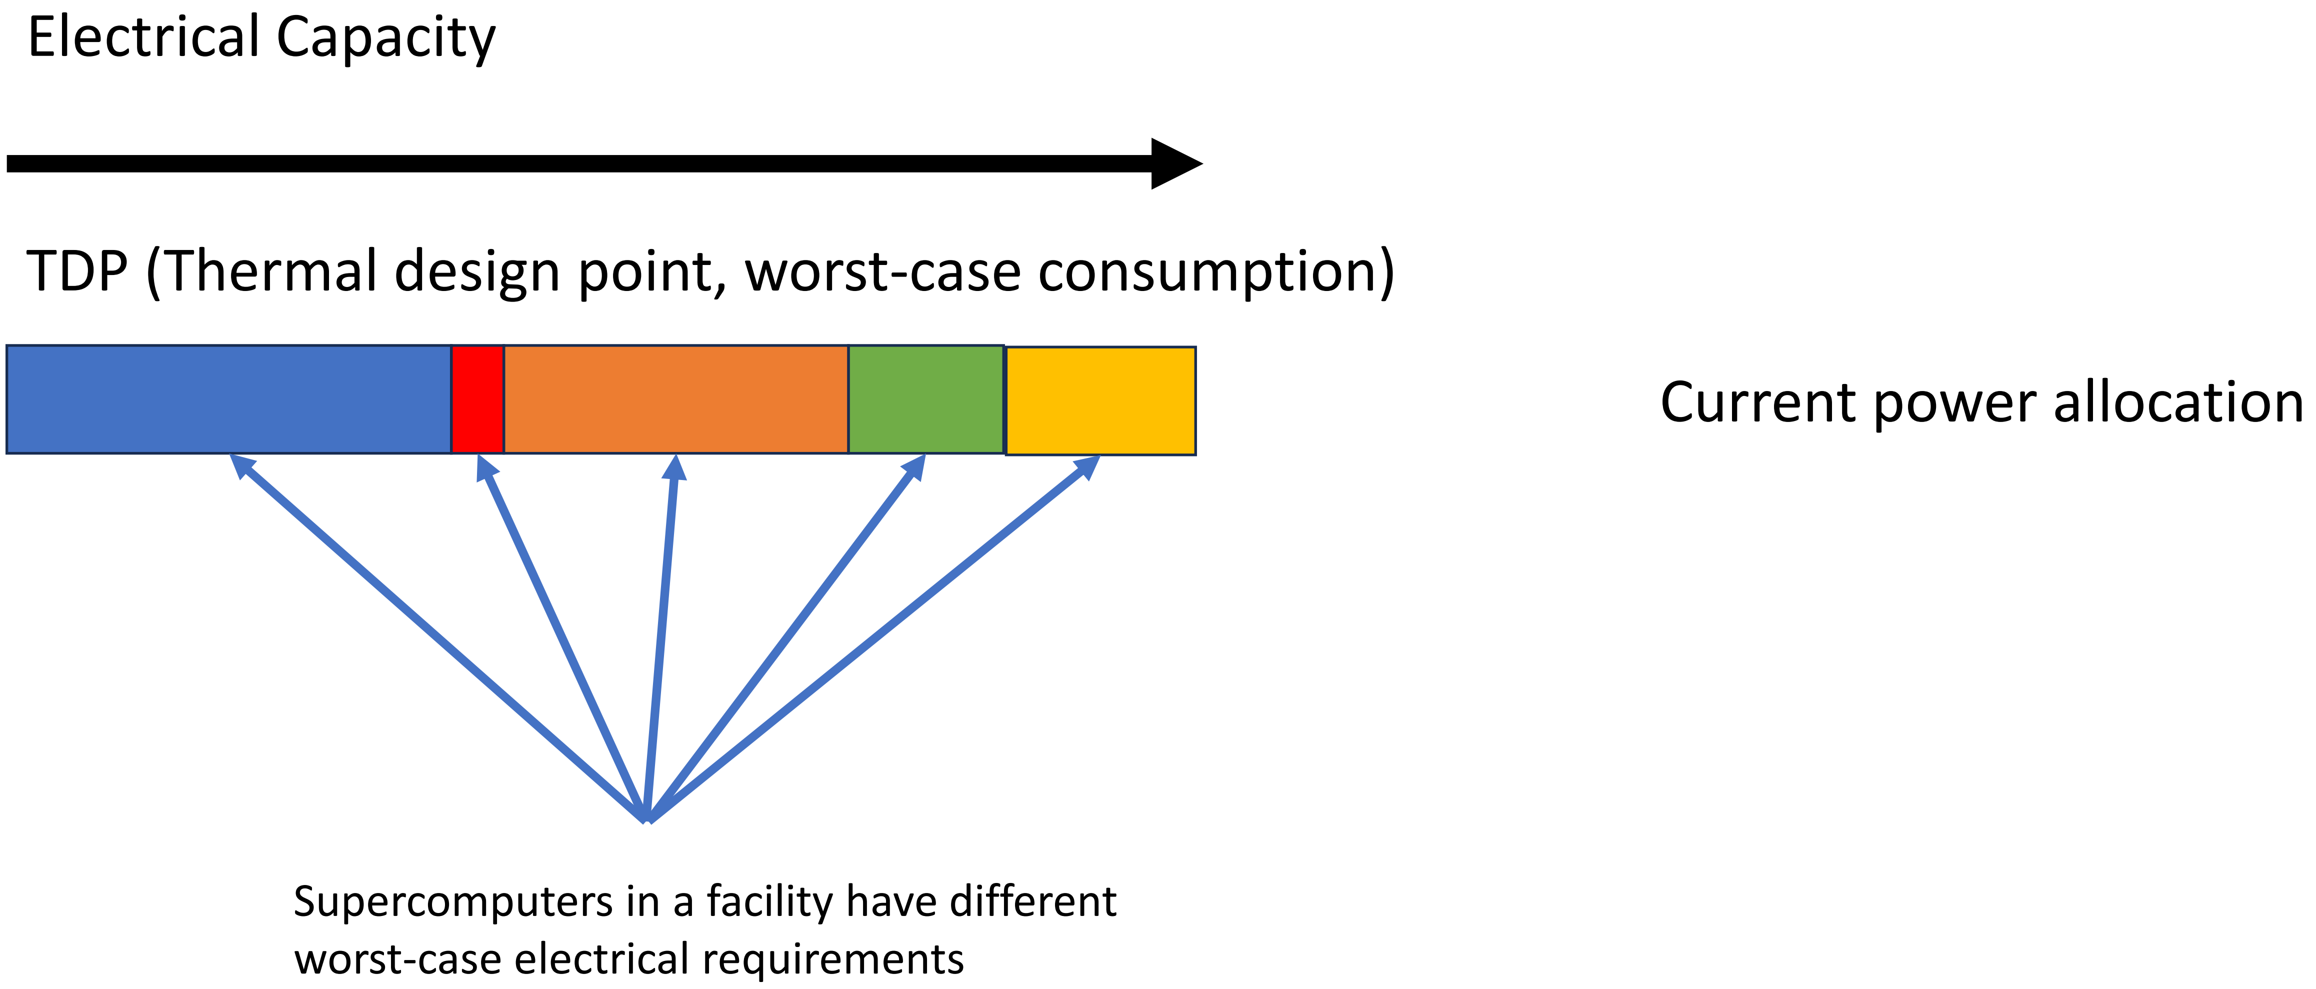
Electrical Capacity
TDP (Thermal design point, worst-case consumption)
Current power allocation
Supercomputers in a facility have different worst-case electrical requirements

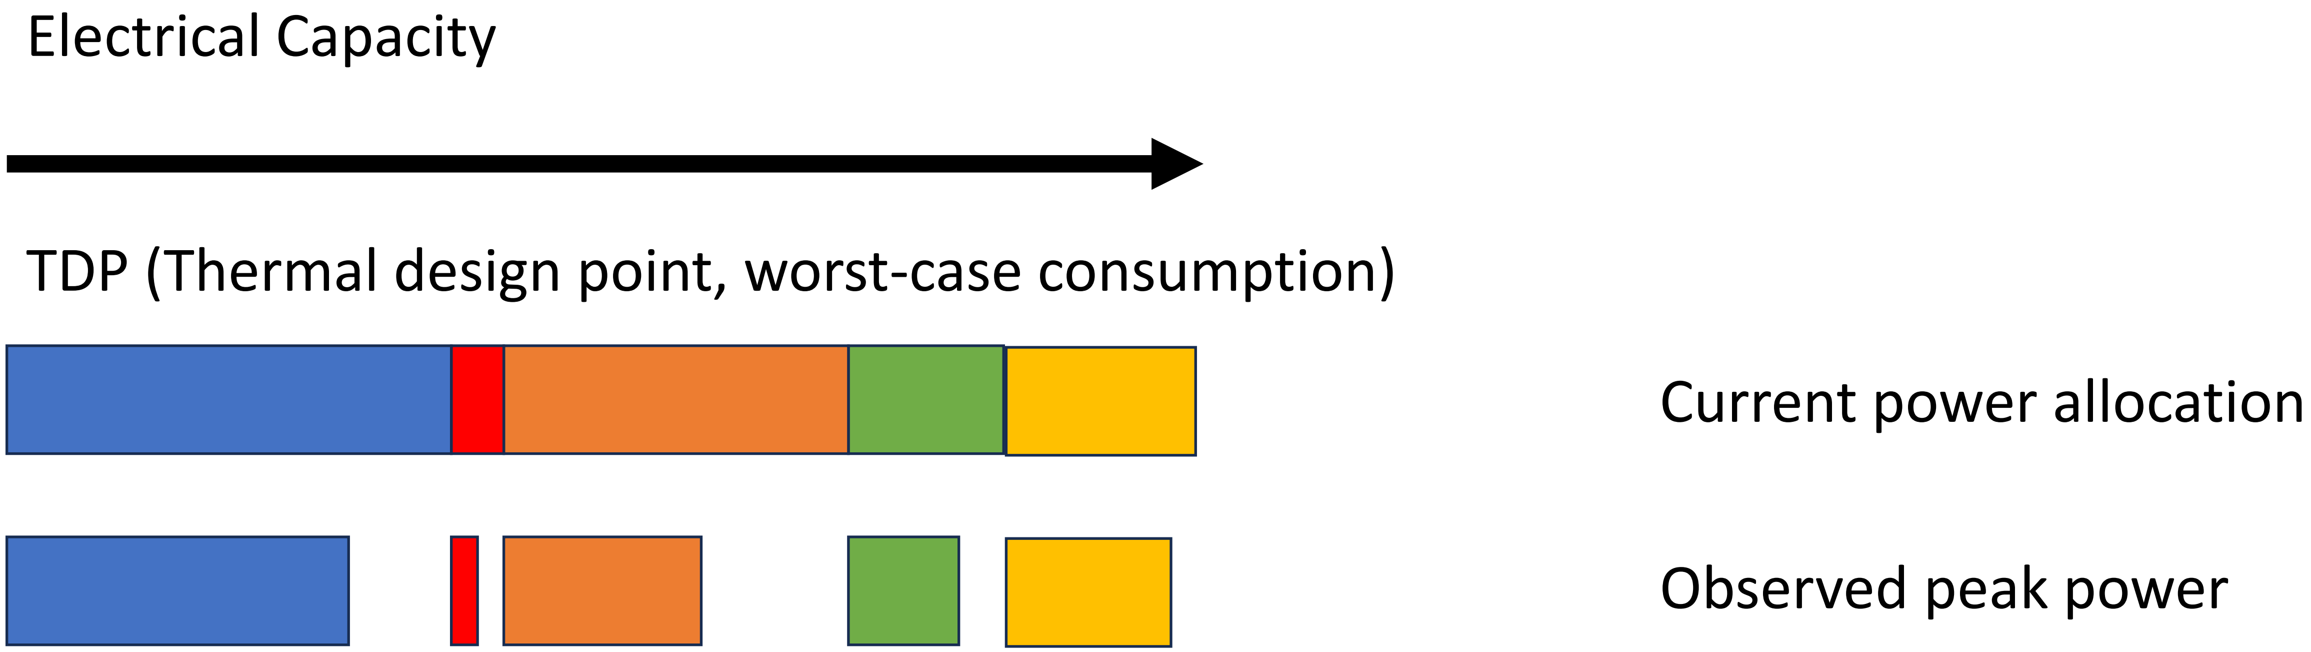
Electrical Capacity
TDP (Thermal design point, worst-case consumption)
Current power allocation
Observed peak power


Electrical Capacity
TDP (Thermal design point, worst-case consumption)
Current power allocation
Observed peak power
Observed average power

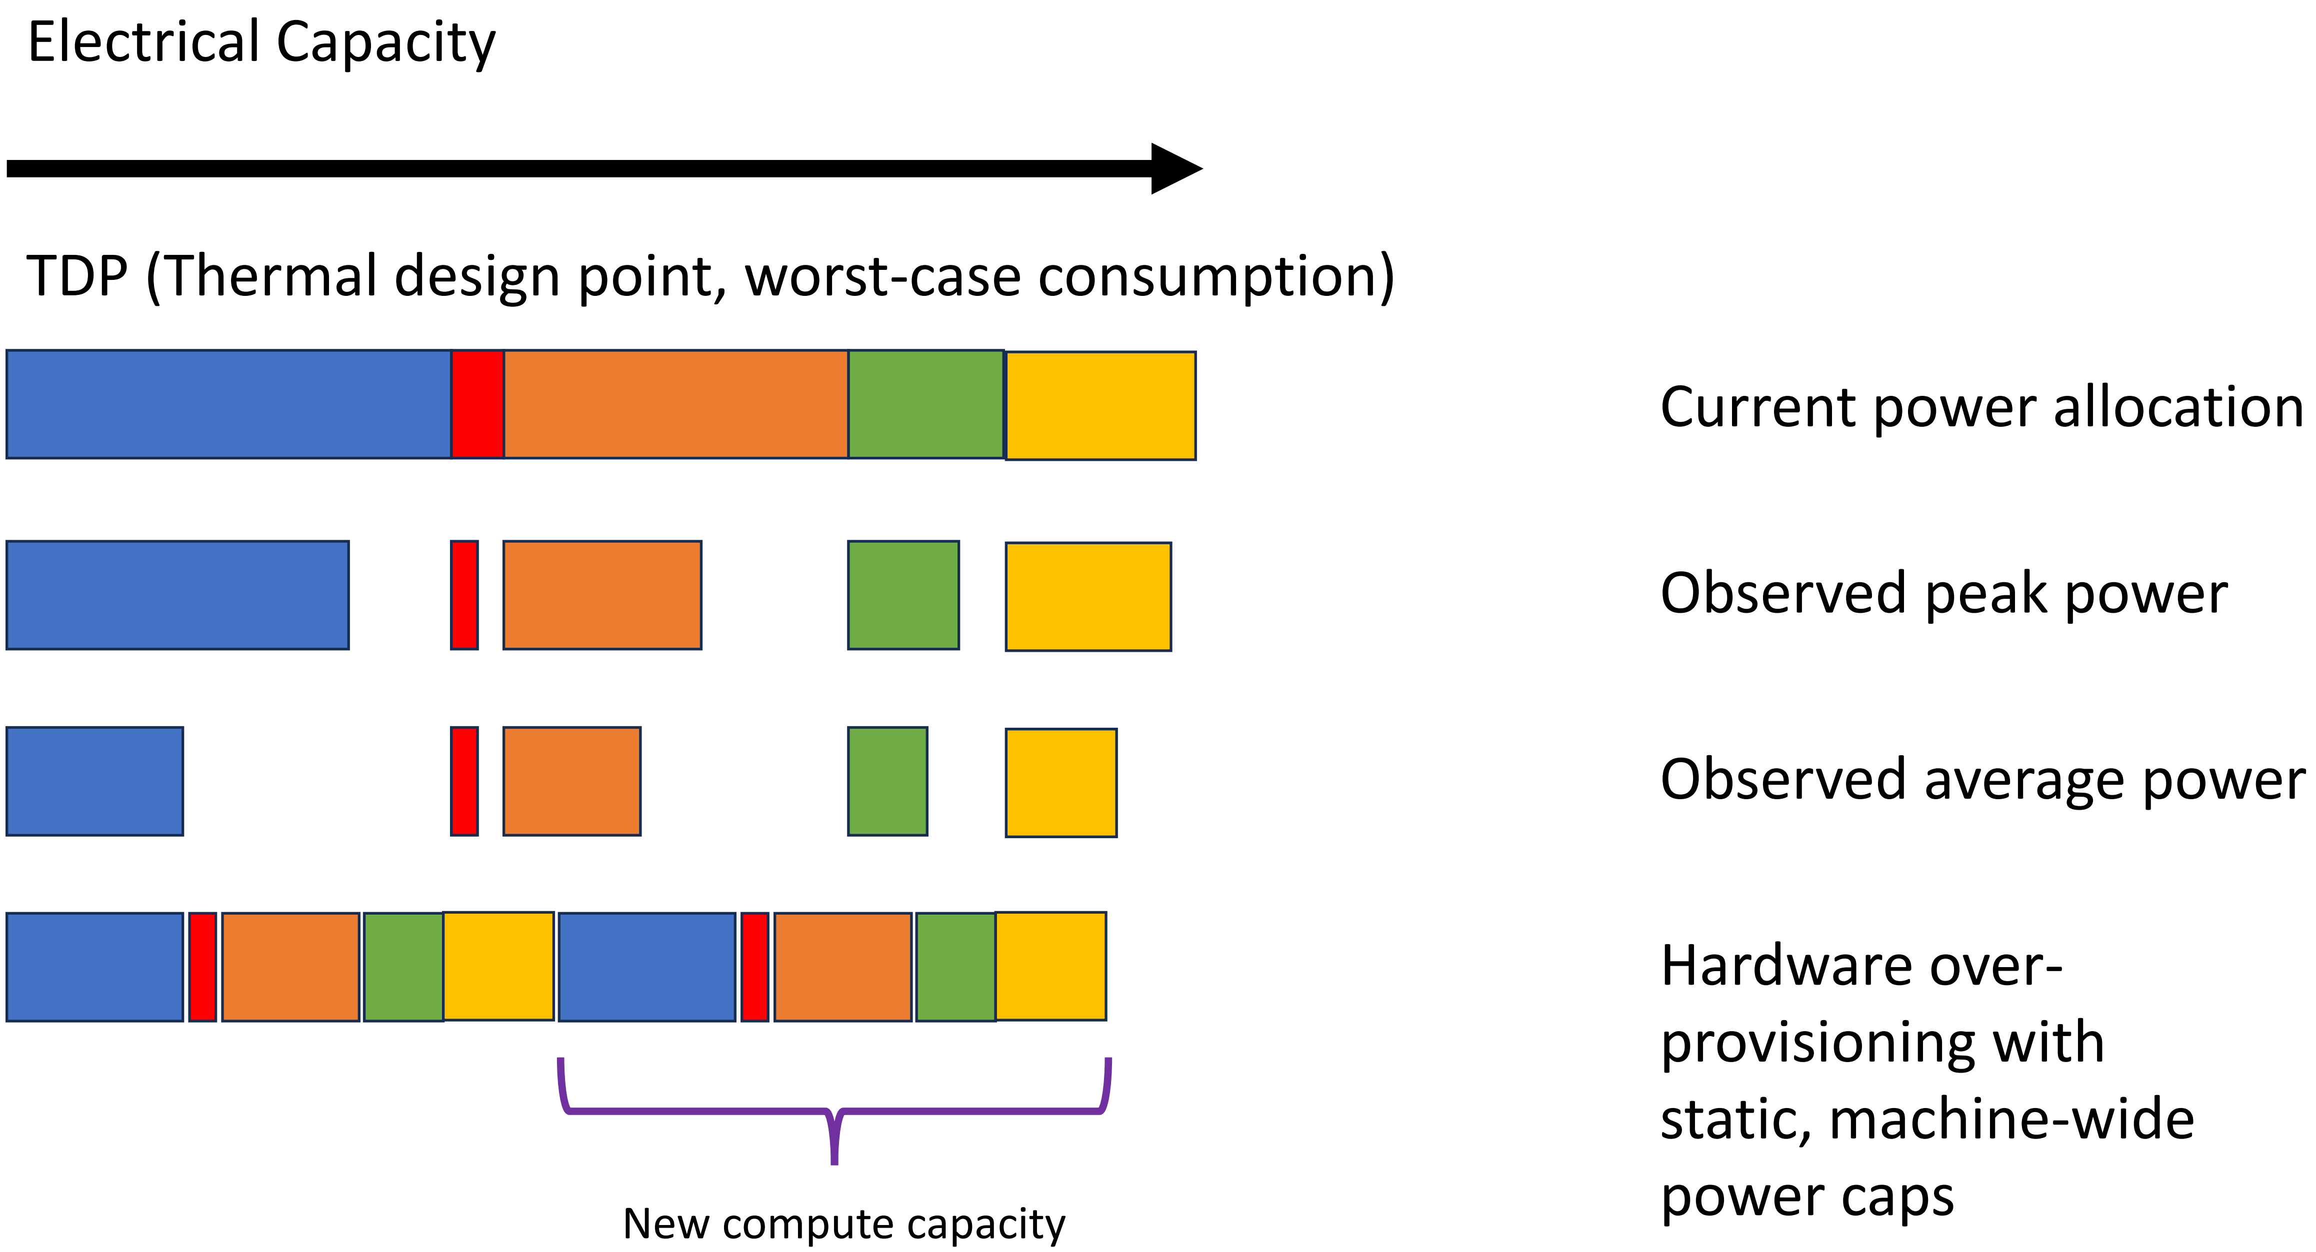
Electrical Capacity
TDP (Thermal design point, worst-case consumption)
Current power allocation
Observed peak power
Observed average power
Hardware over-provisioning with static, machine-wide power caps
New compute capacity

No energy savings

Modest efficiency increase

Electrical Capacity

TDP (Thermal design point, worst-case consumption)

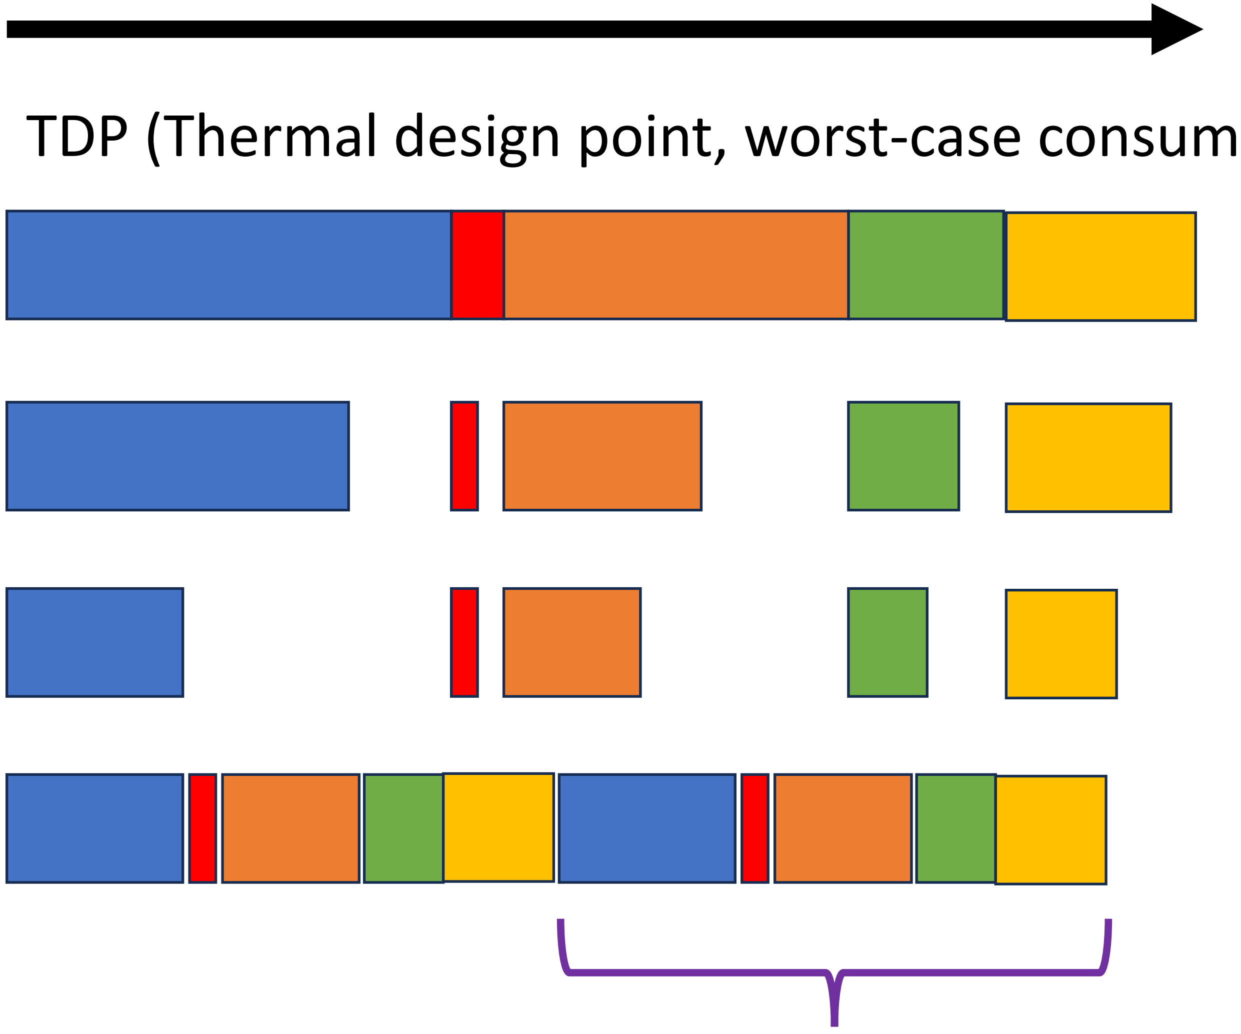

New compute capacity

Current power allocation

Observed peak power

Observed average power

Hardware over-provisioning with static, machine-wide power caps

Electrical Capacity

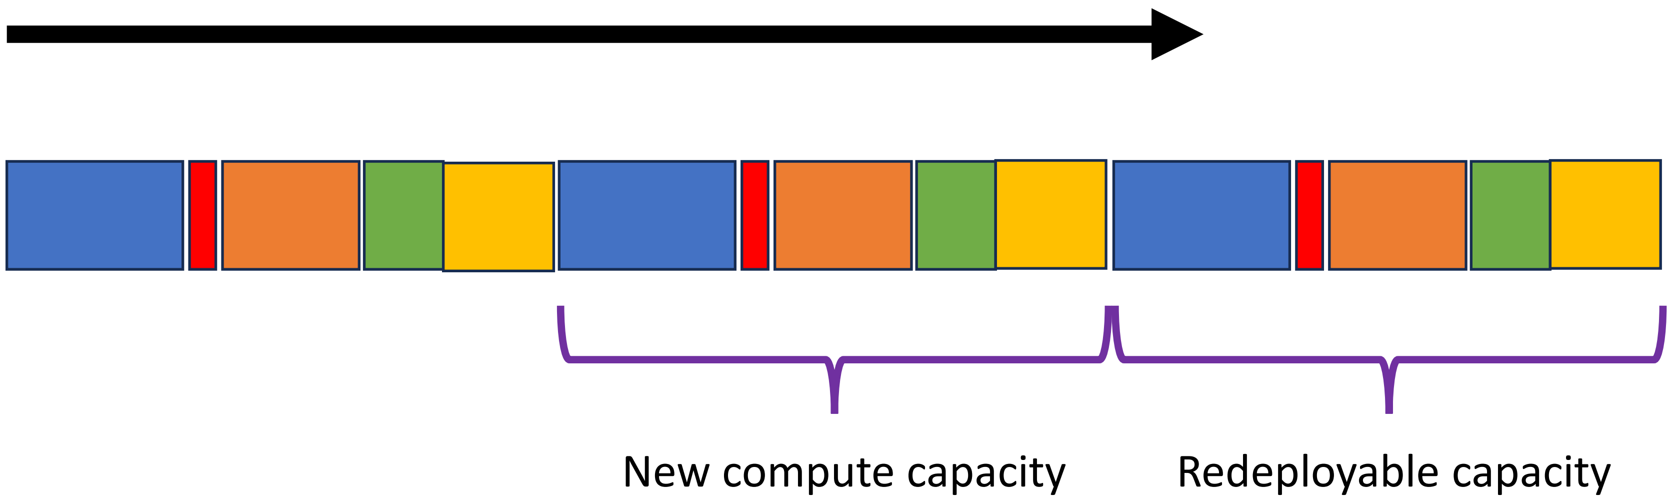

Dynamic hardware overprovisioning allows redeployable computing

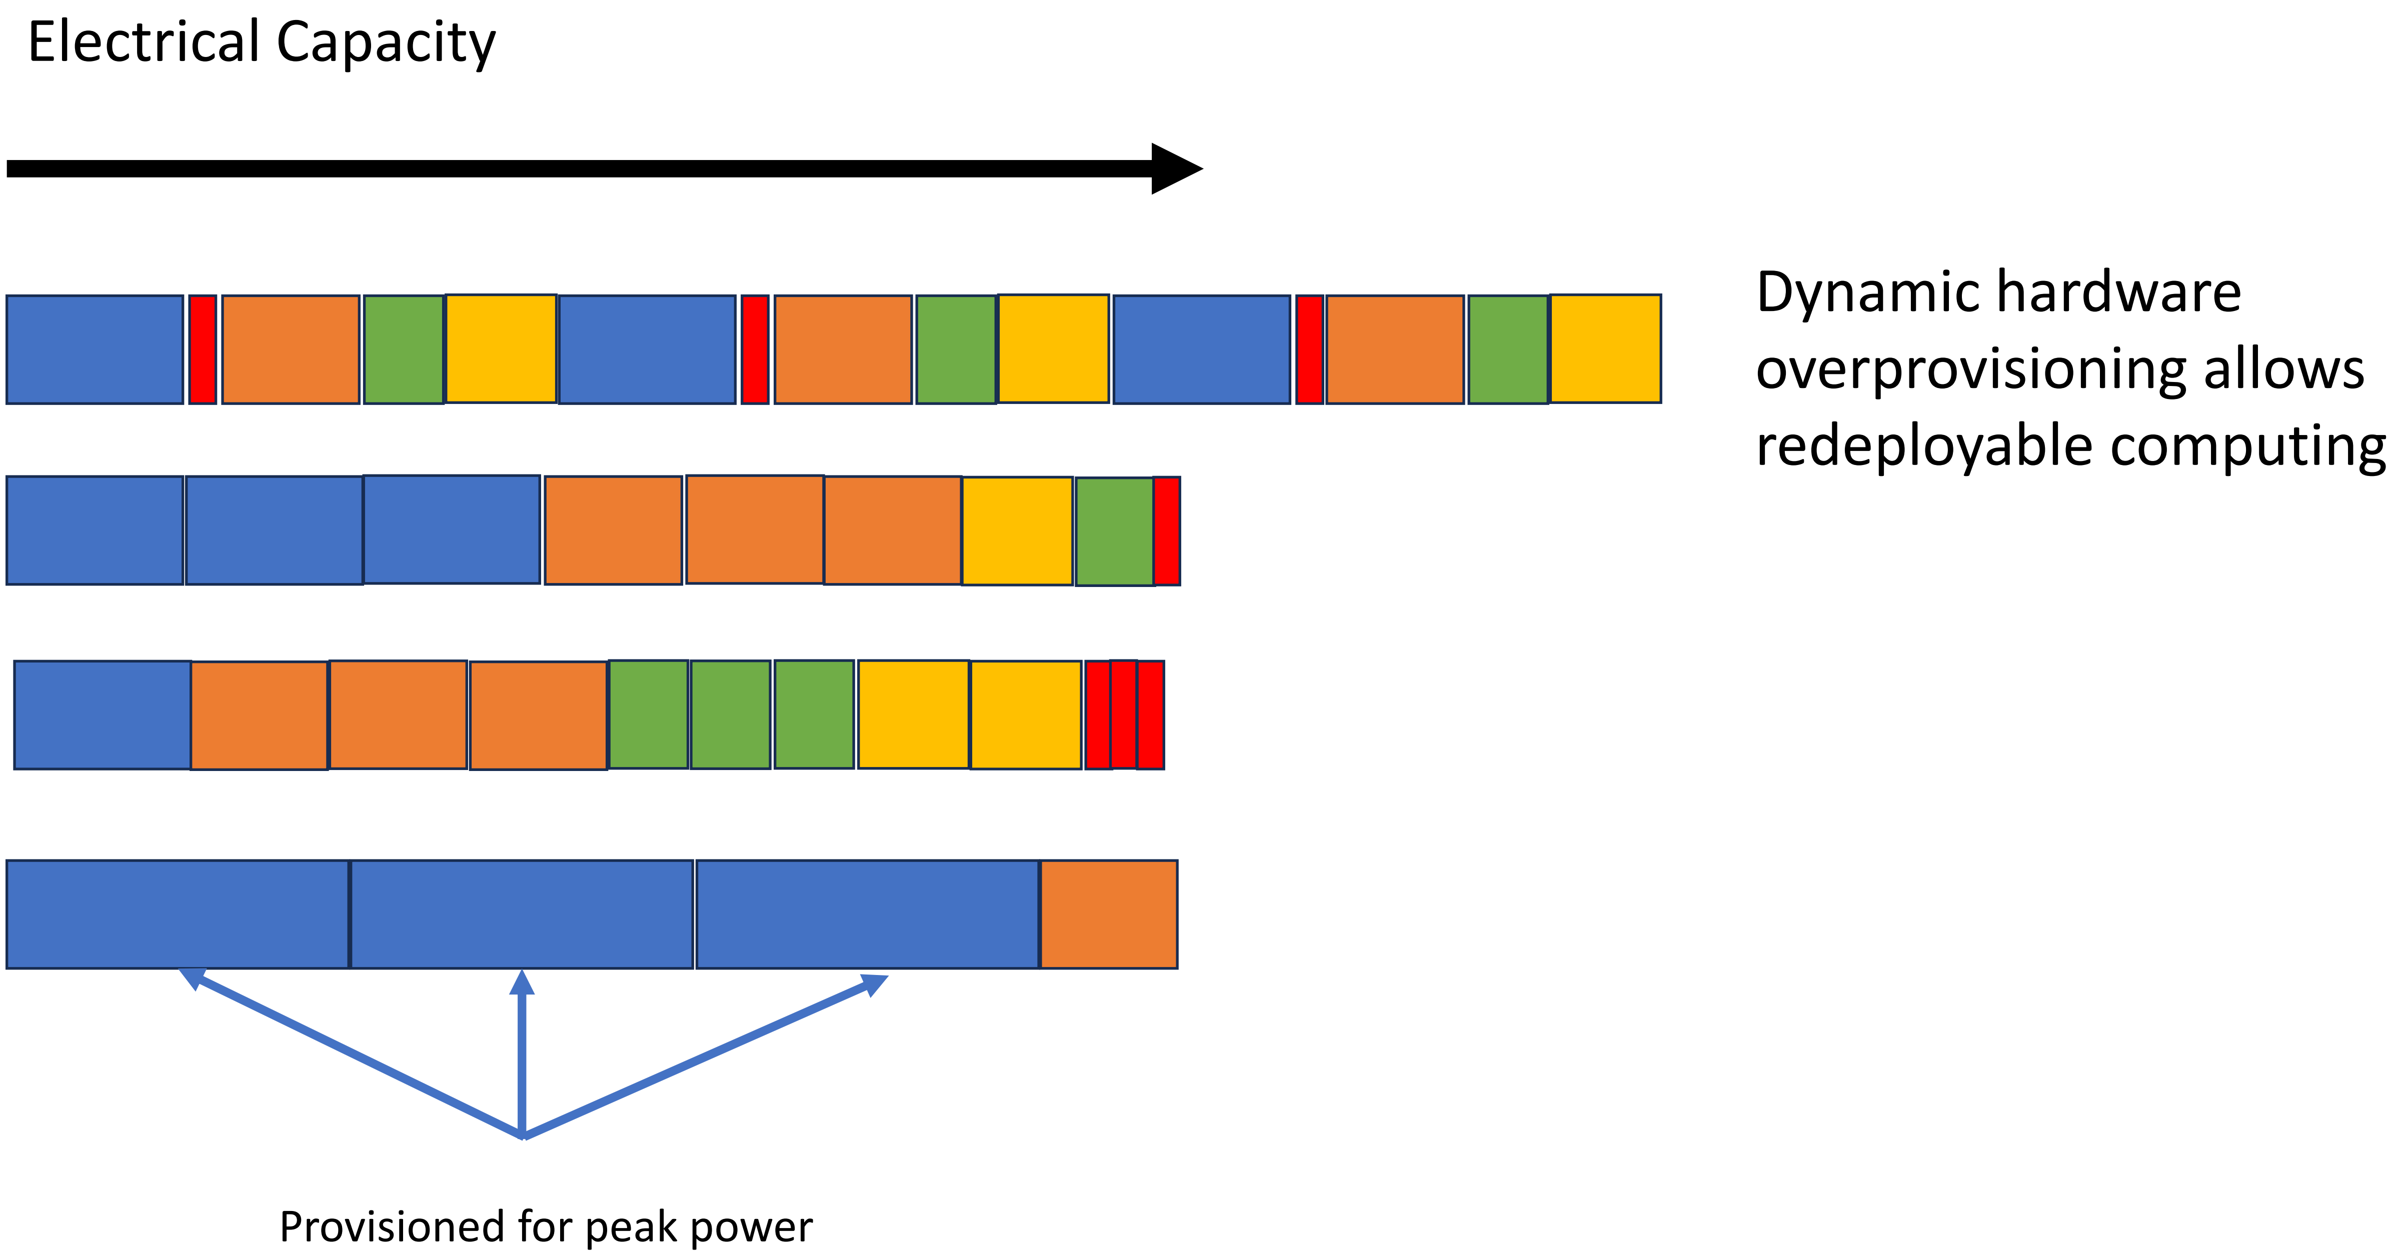
Electrical Capacity
Dynamic hardware overprovisioning allows redeployable computing
Provisioned for peak power

Electrical Capacity

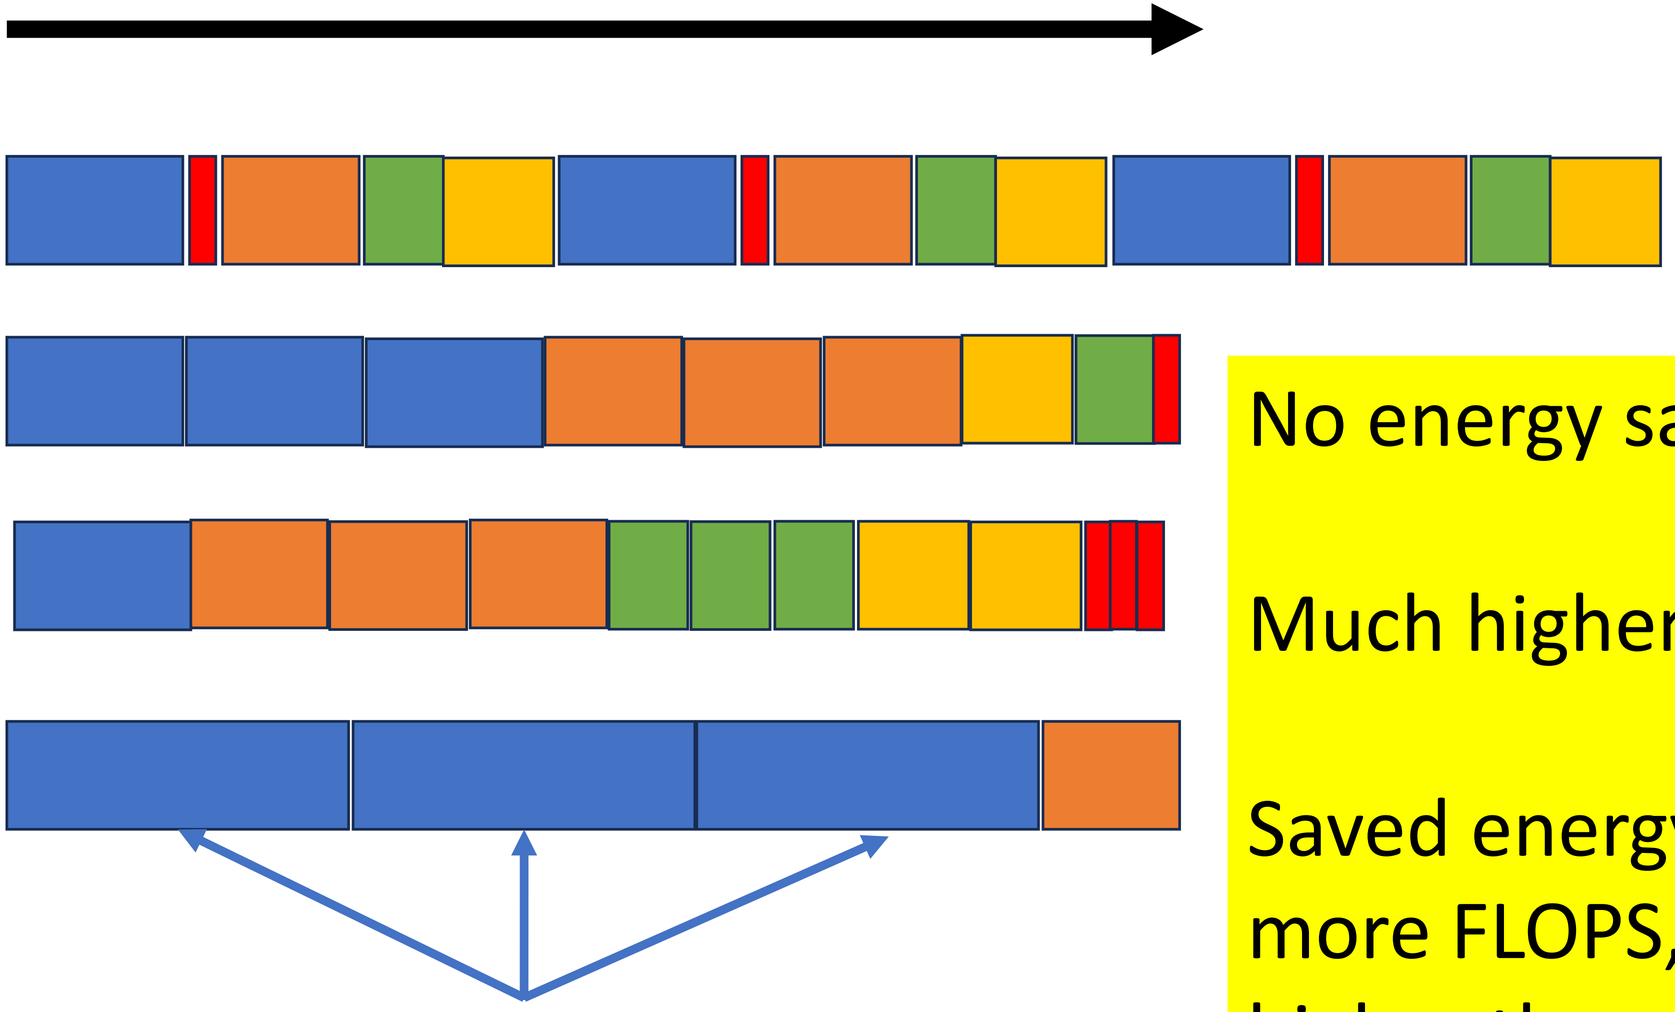

Provisioned for peak power

Dynamic hardware overprovisioning allows redeployable computing

No energy savings

Much higher efficiency

Saved energy translated into more FLOPS, faster turnaround, higher throughput

NABC

- Need
  - Increase compute capacity within the existing facility envelopes
- Approach
  - Apply existing hardware and system controls to cap power and waste heat
    - Level 1: Static, machine-wide caps
    - Level 2: Dynamic machine caps, static job caps   power-aware machine scheduling
    - Level 3: Dynamic job caps, uniform node caps   power-aware scheduling
    - Level 4: Dynamic node caps   power-aware runtime systems
- Benefit
  - More FLOPs, use the capacity we've already paid for
- Competition
  - Cloud computing approaches

LLNL Software Stack

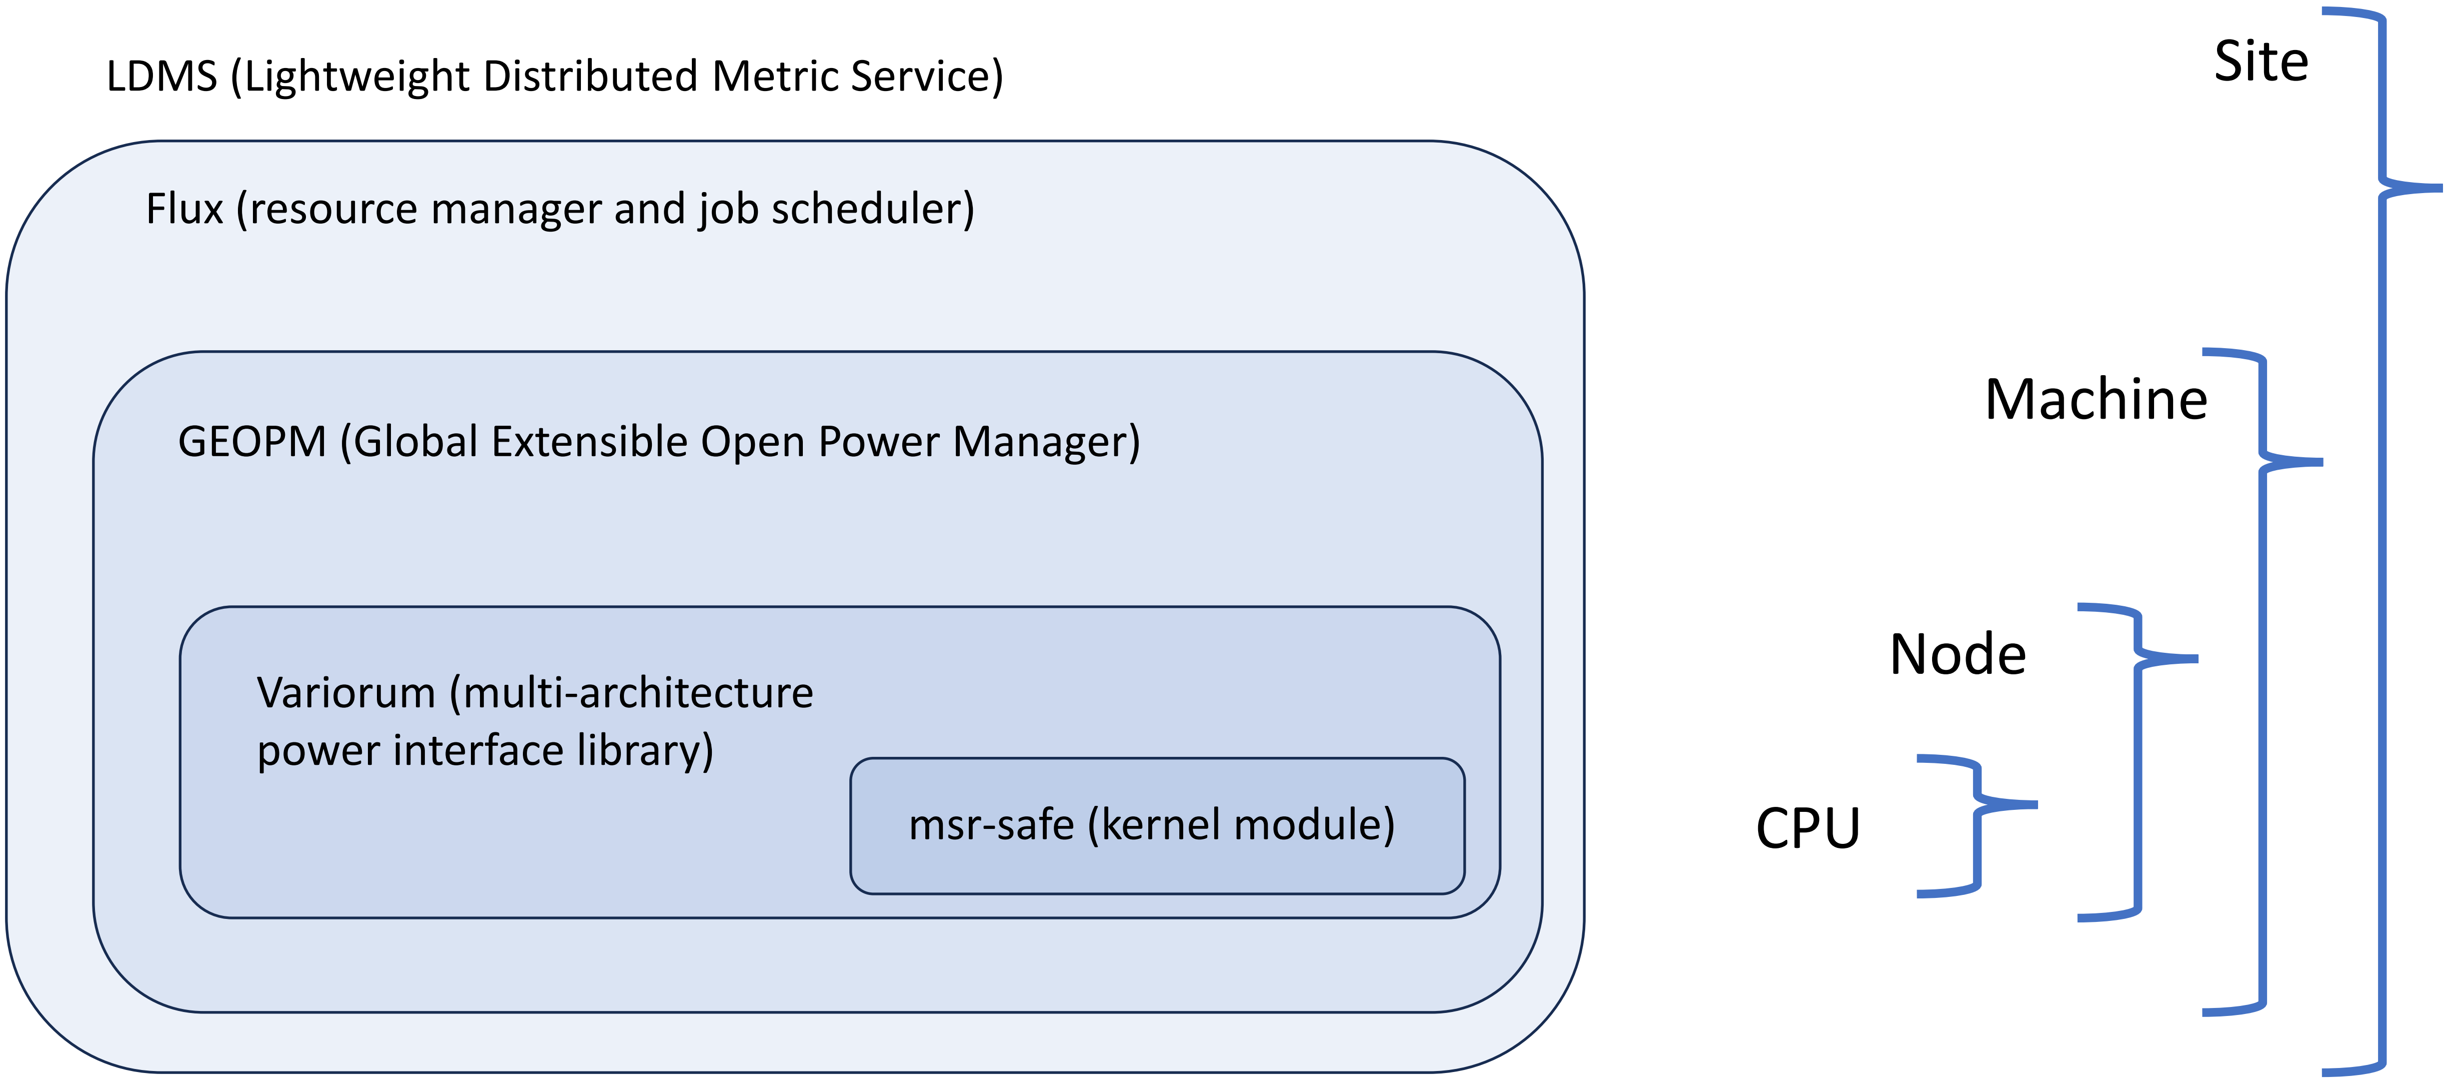

Lessons Learned

- DOE NNSA labs care about energy efficiency.... to a point
  - Efficient machine design is important
  - Once the machine is delivered, run it as fast as possible
- Measure total system energy, not just CPU/GPU energy
  - Faster algorithms are (almost) always more energy efficient
- Respect diversity of optimization goals
  - Speed, turnaround time, throughput need different strategies
- Energy efficiency matters most when energy is a constraint

Open Questions

- How to simultaneously schedule jobs, energy, hardware, and time?
- How much is performance reproducibility worth?
- What replaces “node hours” for user accounting and prioritization?
- How much complexity/choice should be revealed to users?
- How to we guarantee safe operation?