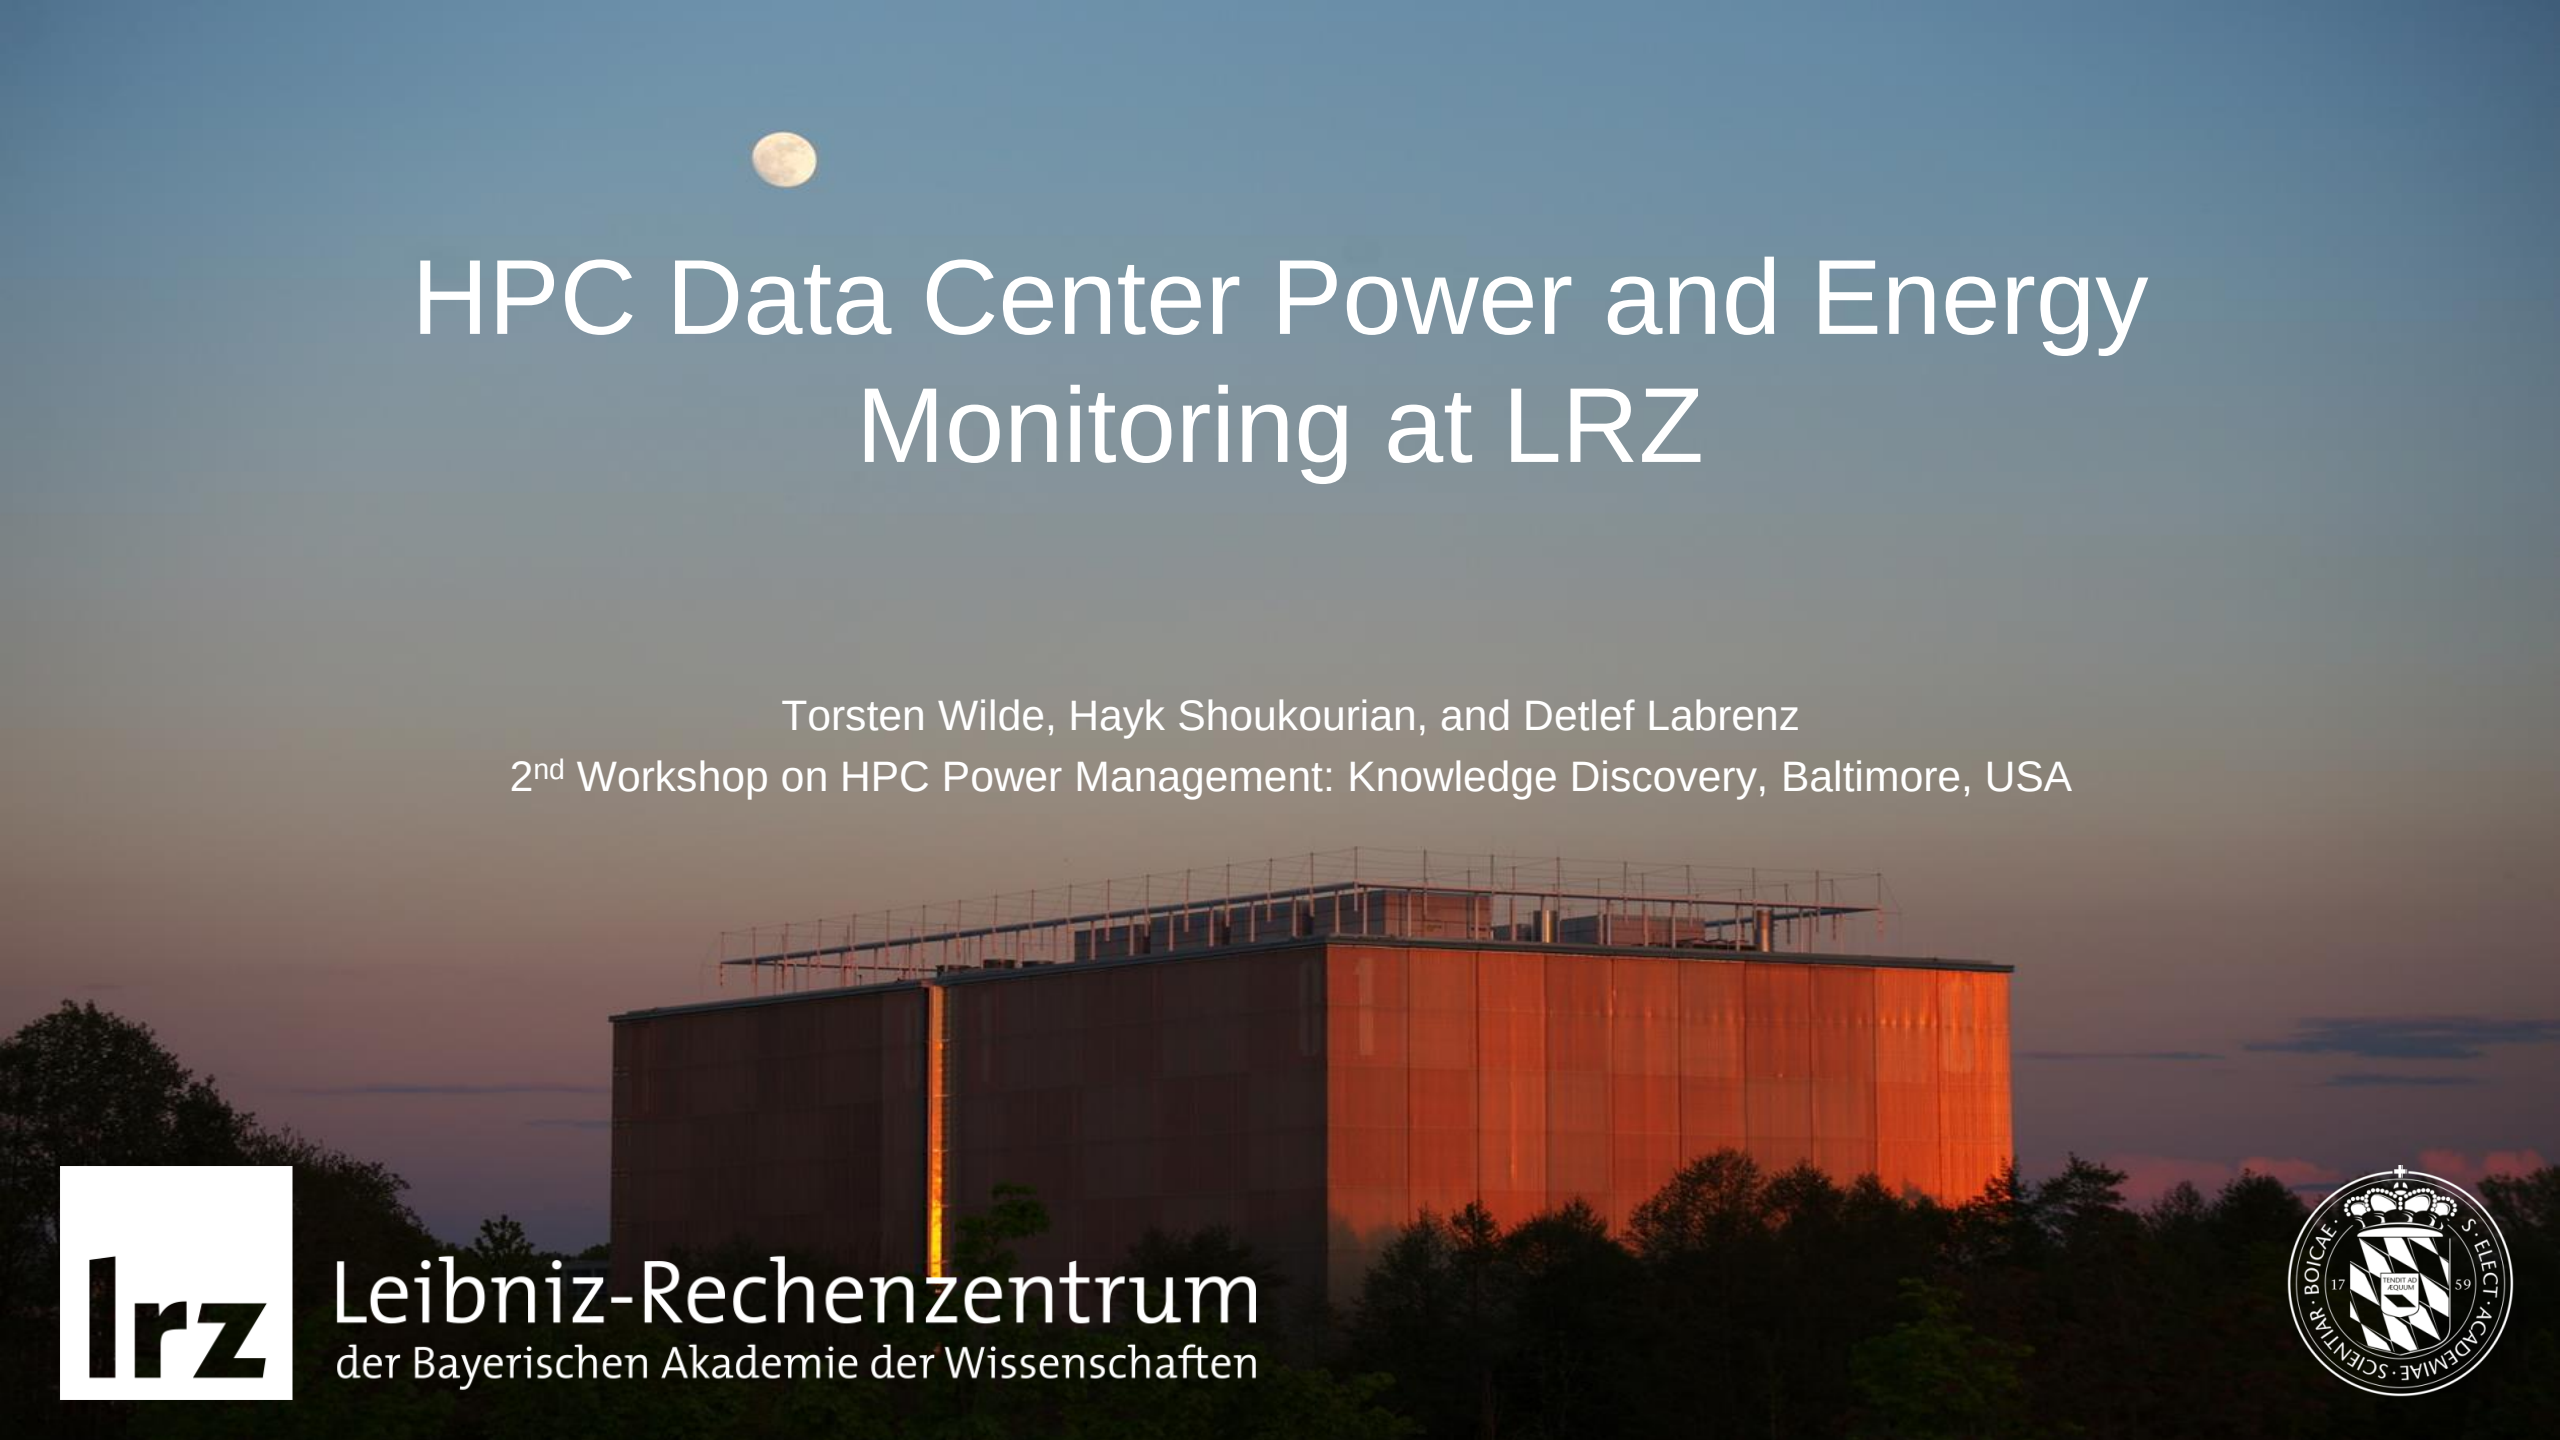

HPC Data Center Power and Energy Monitoring at LRZ

Torsten Wilde, Hayk Shoukourian, and Detlef Labrenz

2nd Workshop on HPC Power Management: Knowledge Discovery, Baltimore, USA

lrz Leibniz-Rechenzentrum
der Bayerischen Akademie der Wissenschaften

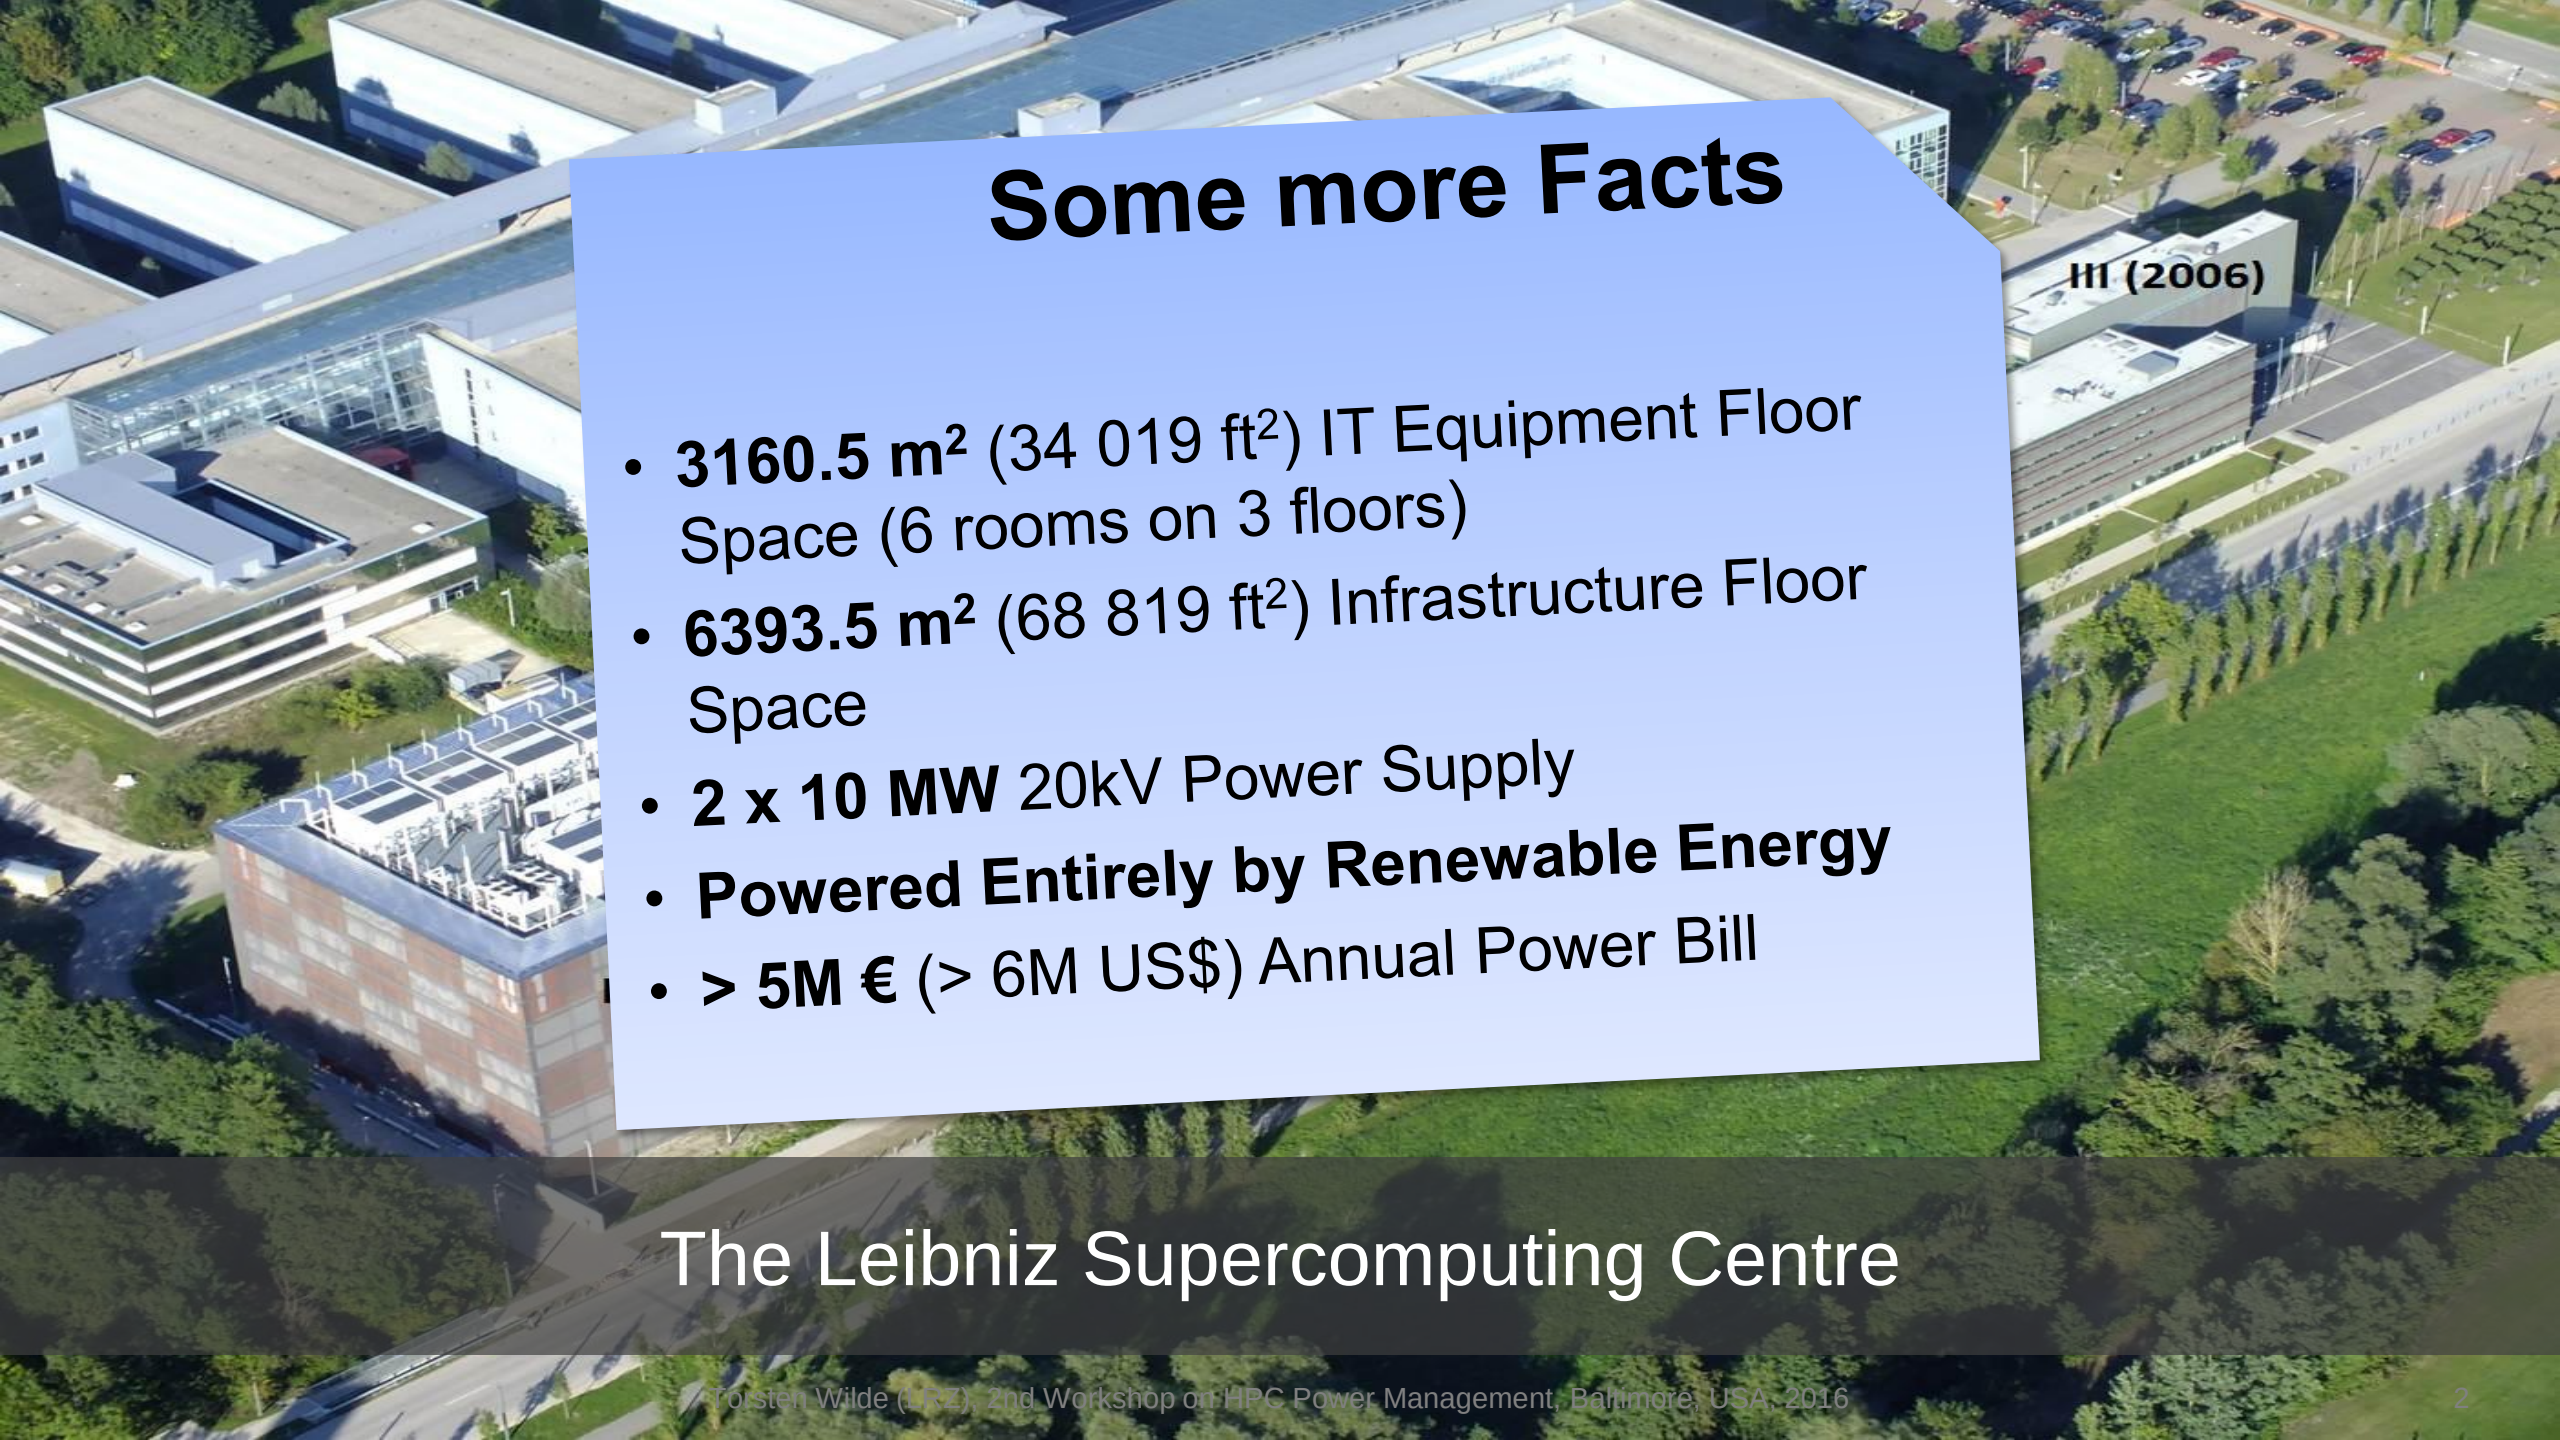

Some more Facts

- **3160.5 m²** (34 019 ft²) IT Equipment Floor Space (6 rooms on 3 floors)
- **6393.5 m²** (68 819 ft²) Infrastructure Floor Space
- **2 x 10 MW** 20kV Power Supply
- **Powered Entirely by Renewable Energy**
- **> 5M €** (> 6M US$) Annual Power Bill

The Leibniz Supercomputing Centre

Scientific Computing at LRZ

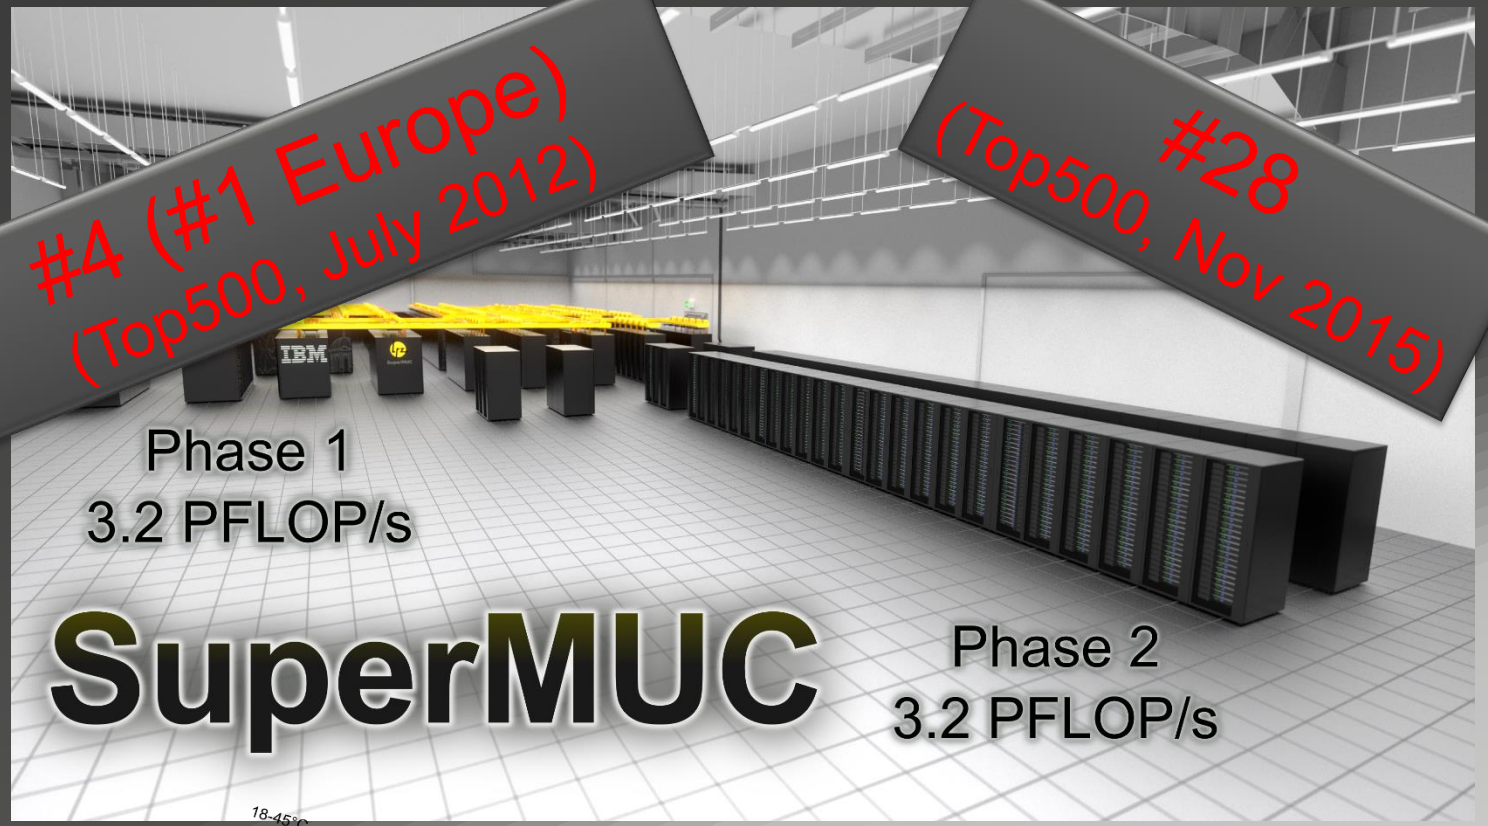

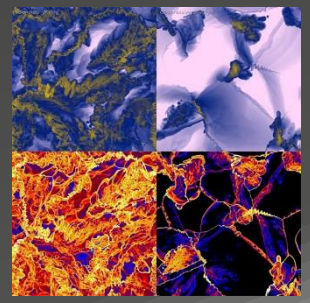

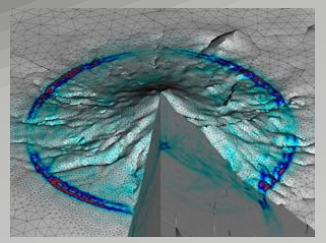

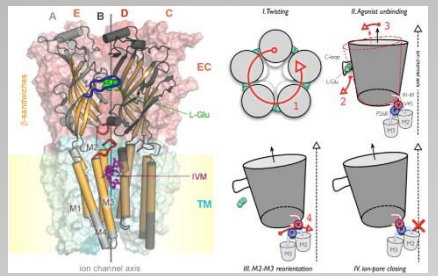

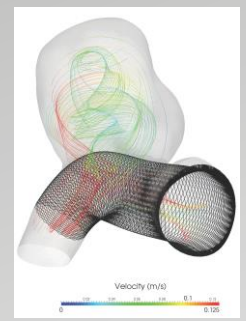

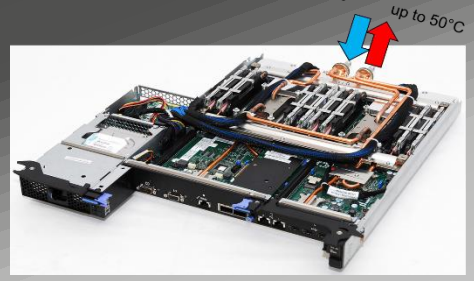

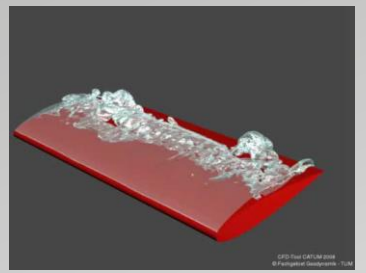

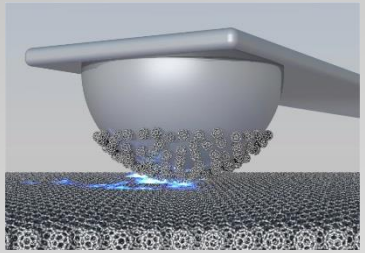

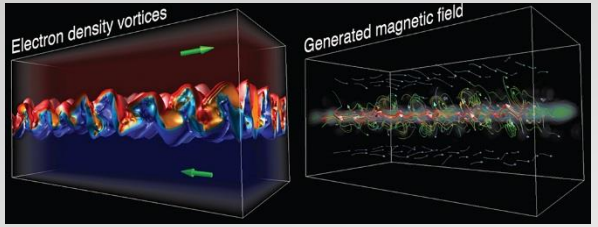

http://www.lrz.de/services/compute/supermuc/magazinesbooks/2014_SuperMUC-Results-Reports.pdf

http://www.gauss-centre.eu/gauss-centre/EN/Projects/projects_node.html

Motivation: Improve the Energy Efficiency of LRZ

LRZ Trend of Energy Costs and Energy Consumption from 2000 till 2015 with prediction for 2016 and 2017

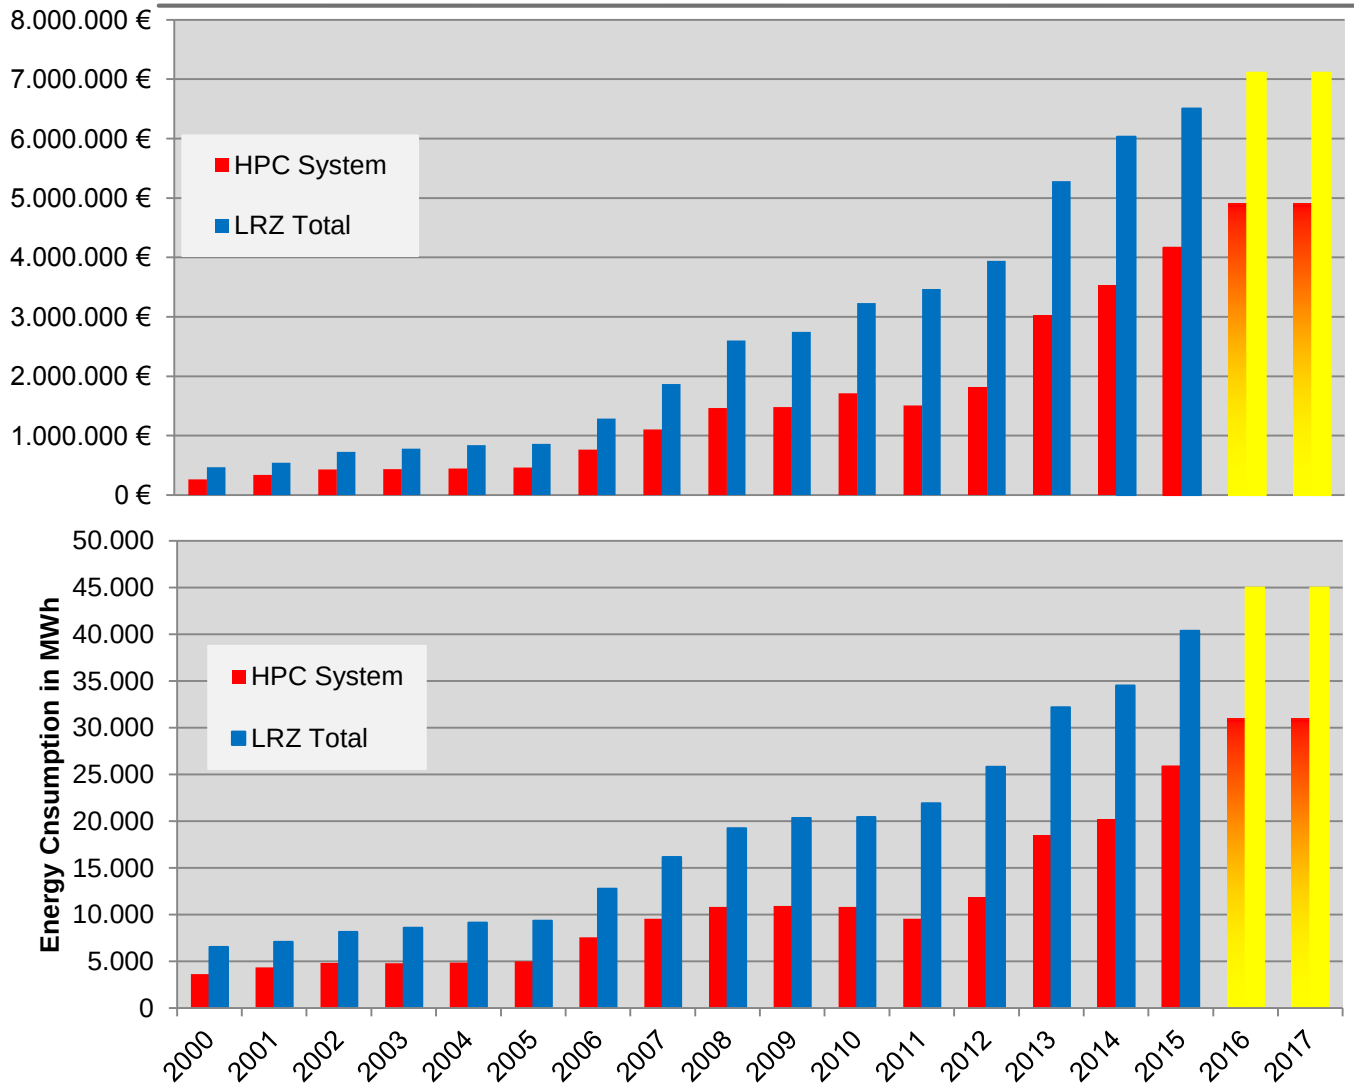

Costs for 1kWh: 2000 – 0.07€
2015 – 0.161€

The 4 Pillar Framework – HPC Data Center Work/Energy Flow

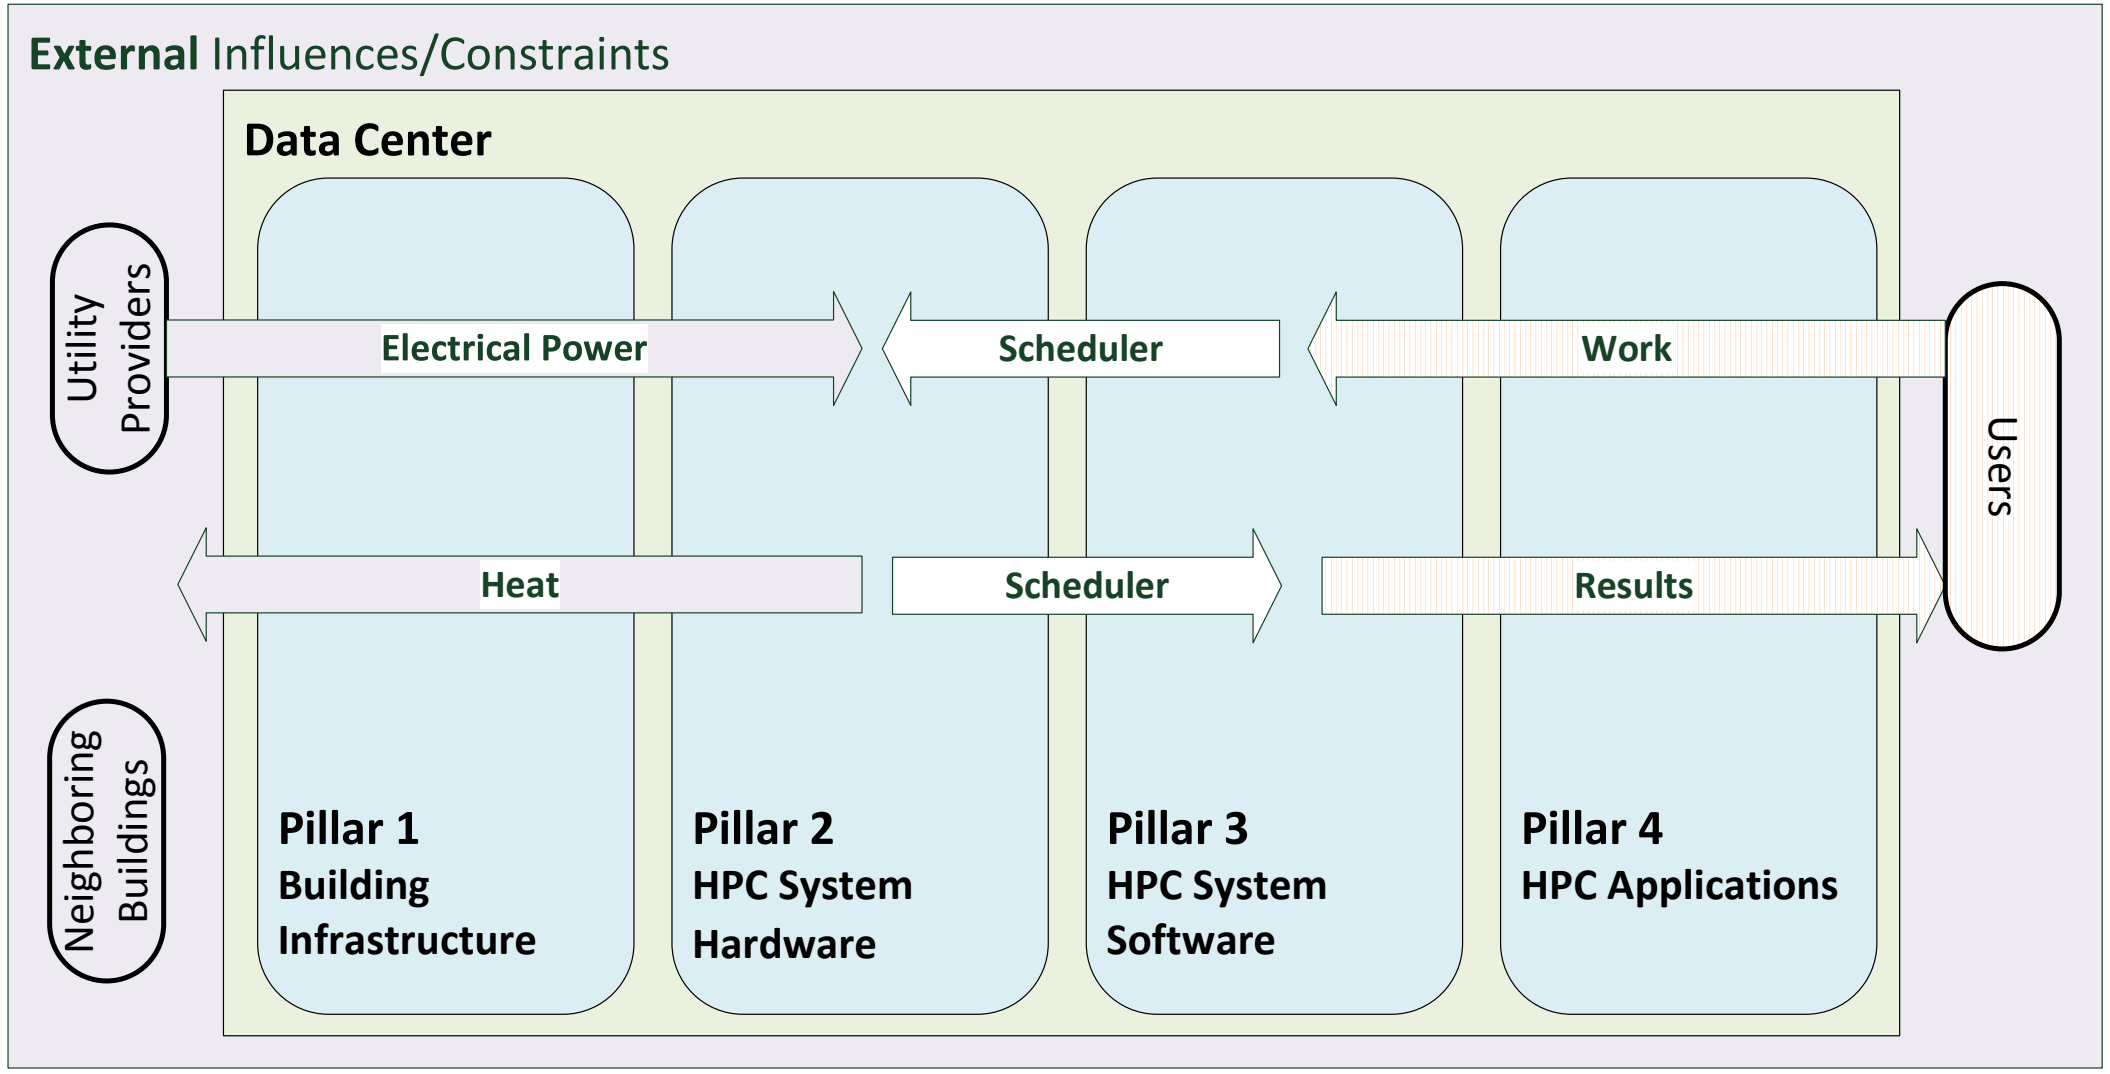

Where does Our Energy go?

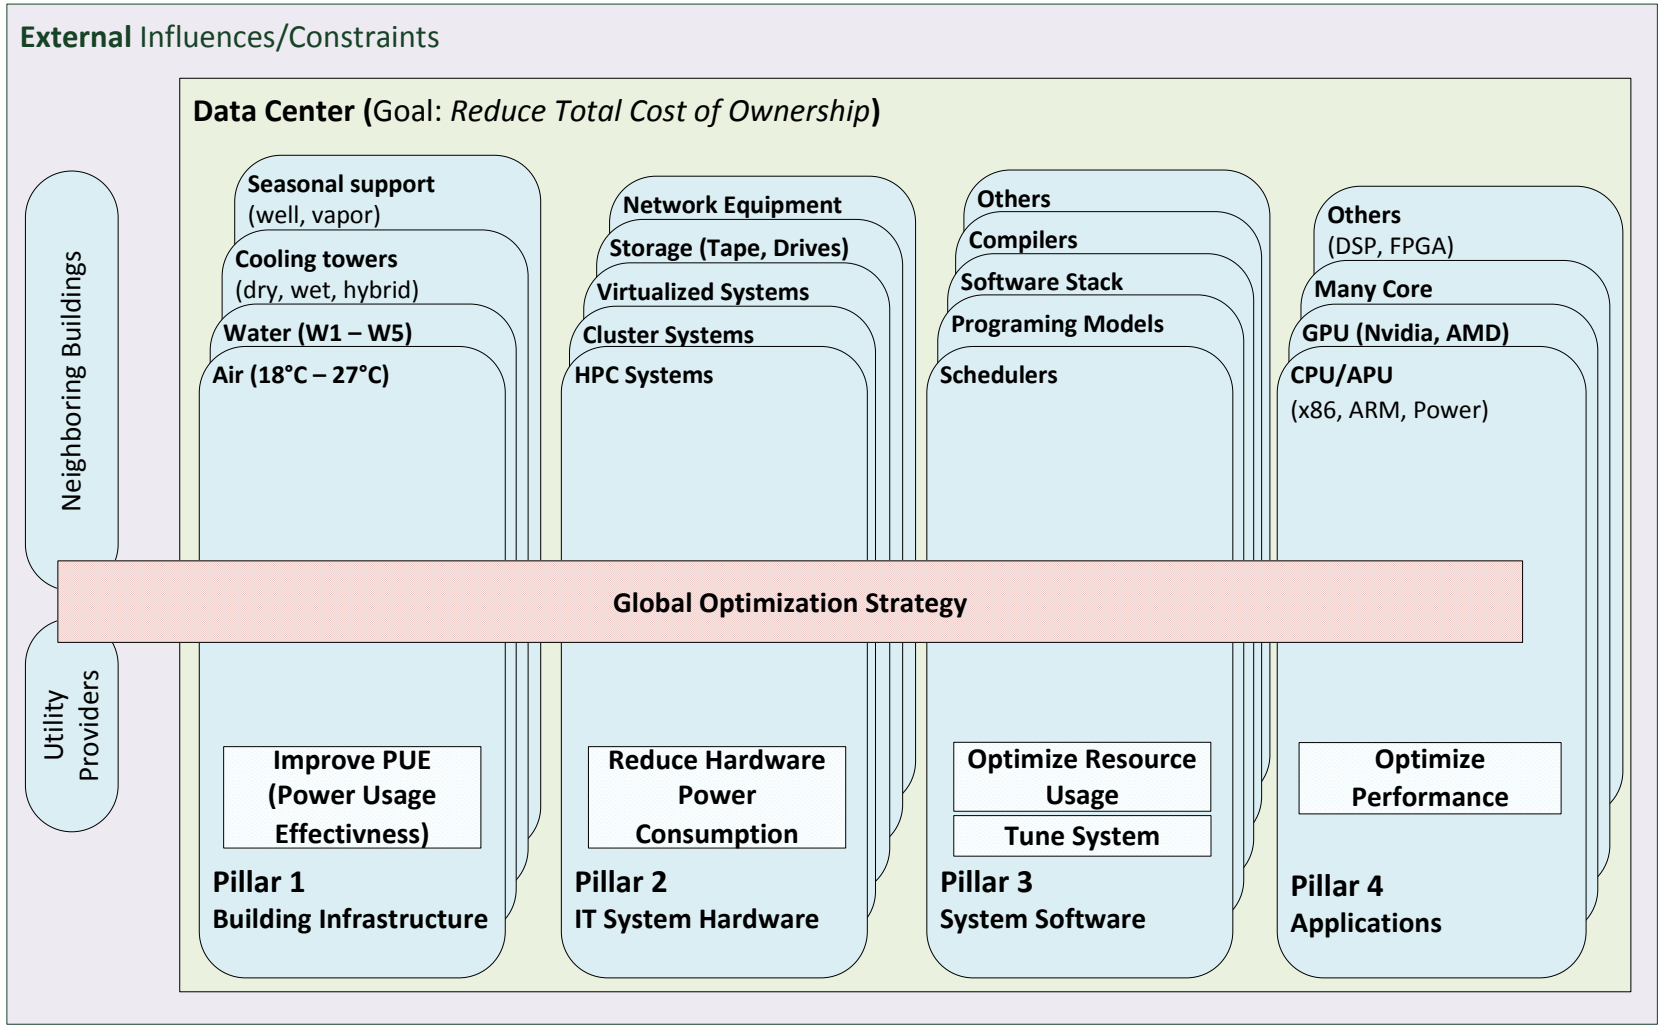

- Need to understand each pillar
- Optimize and measure (KPIs) for each
- Need global approach for optimal results
  - includes utility provider
  - define operating points
  - keep infrastructure efficiency constant over the whole operating range
  - measure and assess

Open Access 4 Pillar Framework Paper: http://www.springerlink.com/openurl.asp?genre=article&id=doi:10.1007/s00450-013-0244-6

Data Collection

Current Data Center Monitoring at LRZ

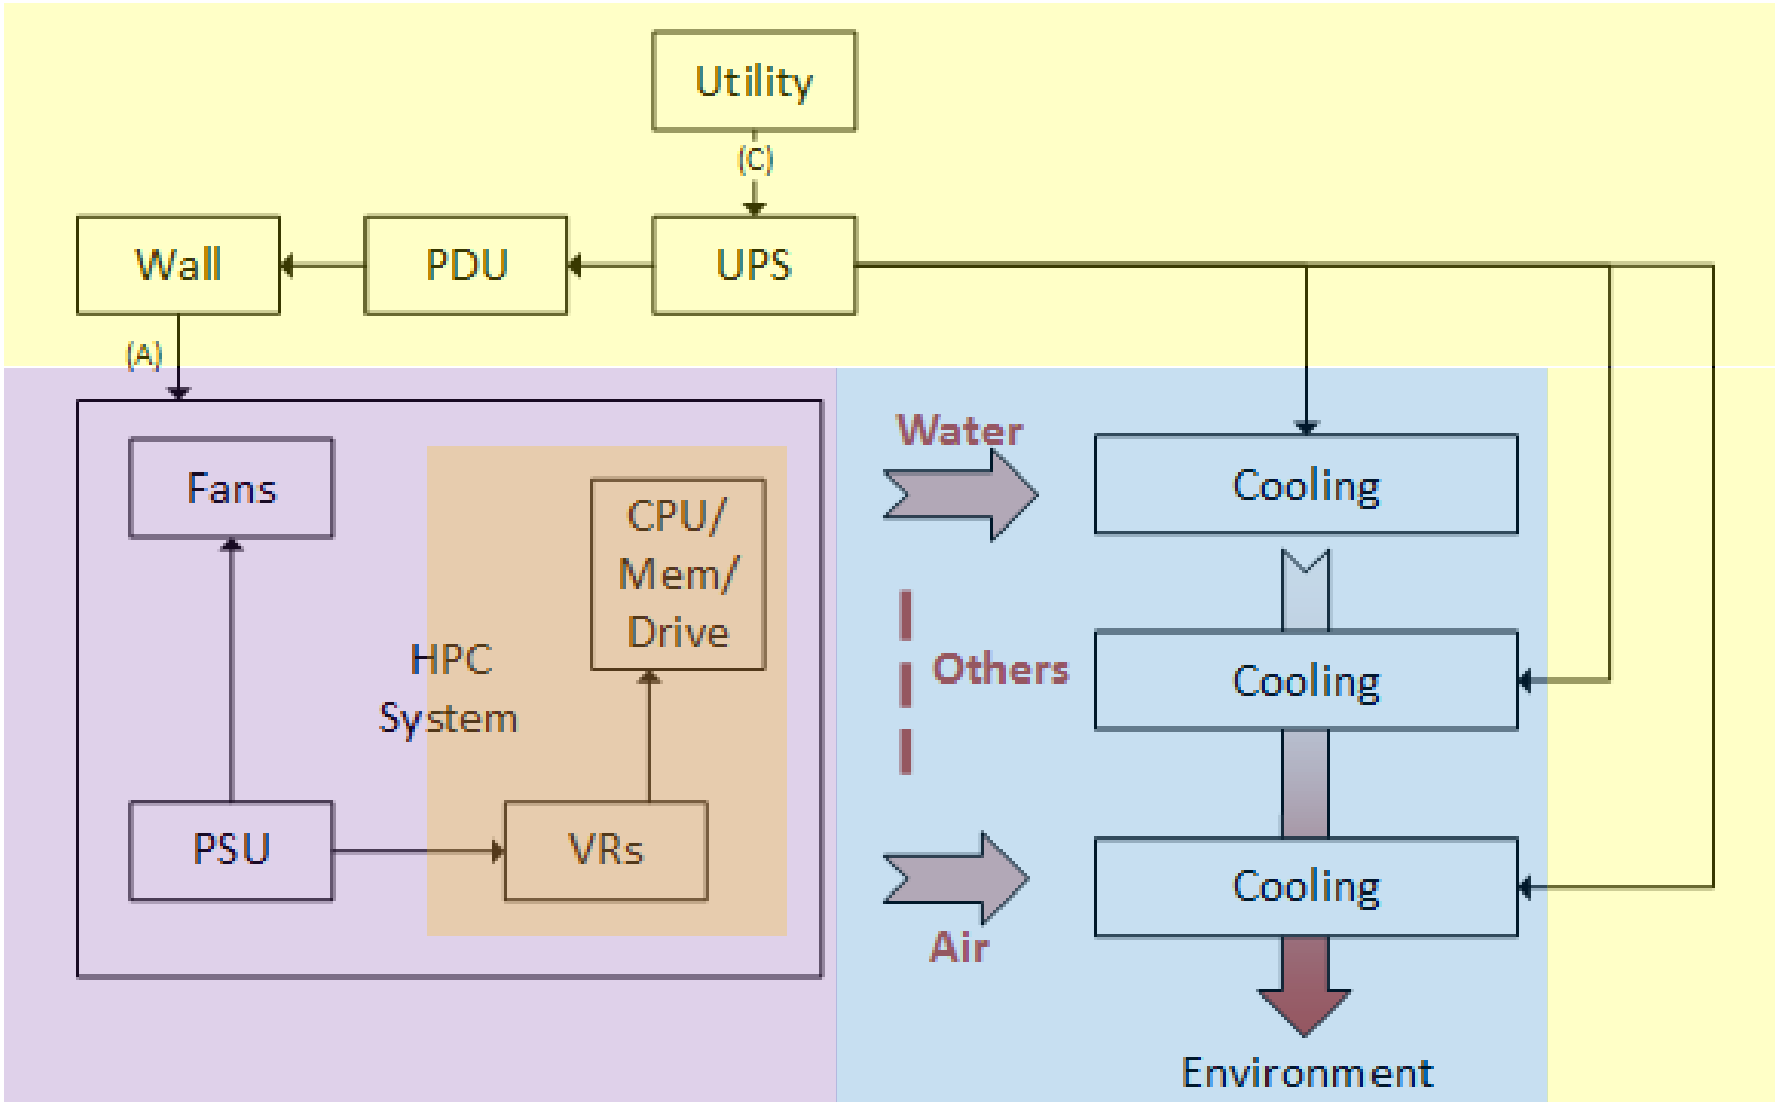

Data Consolidation

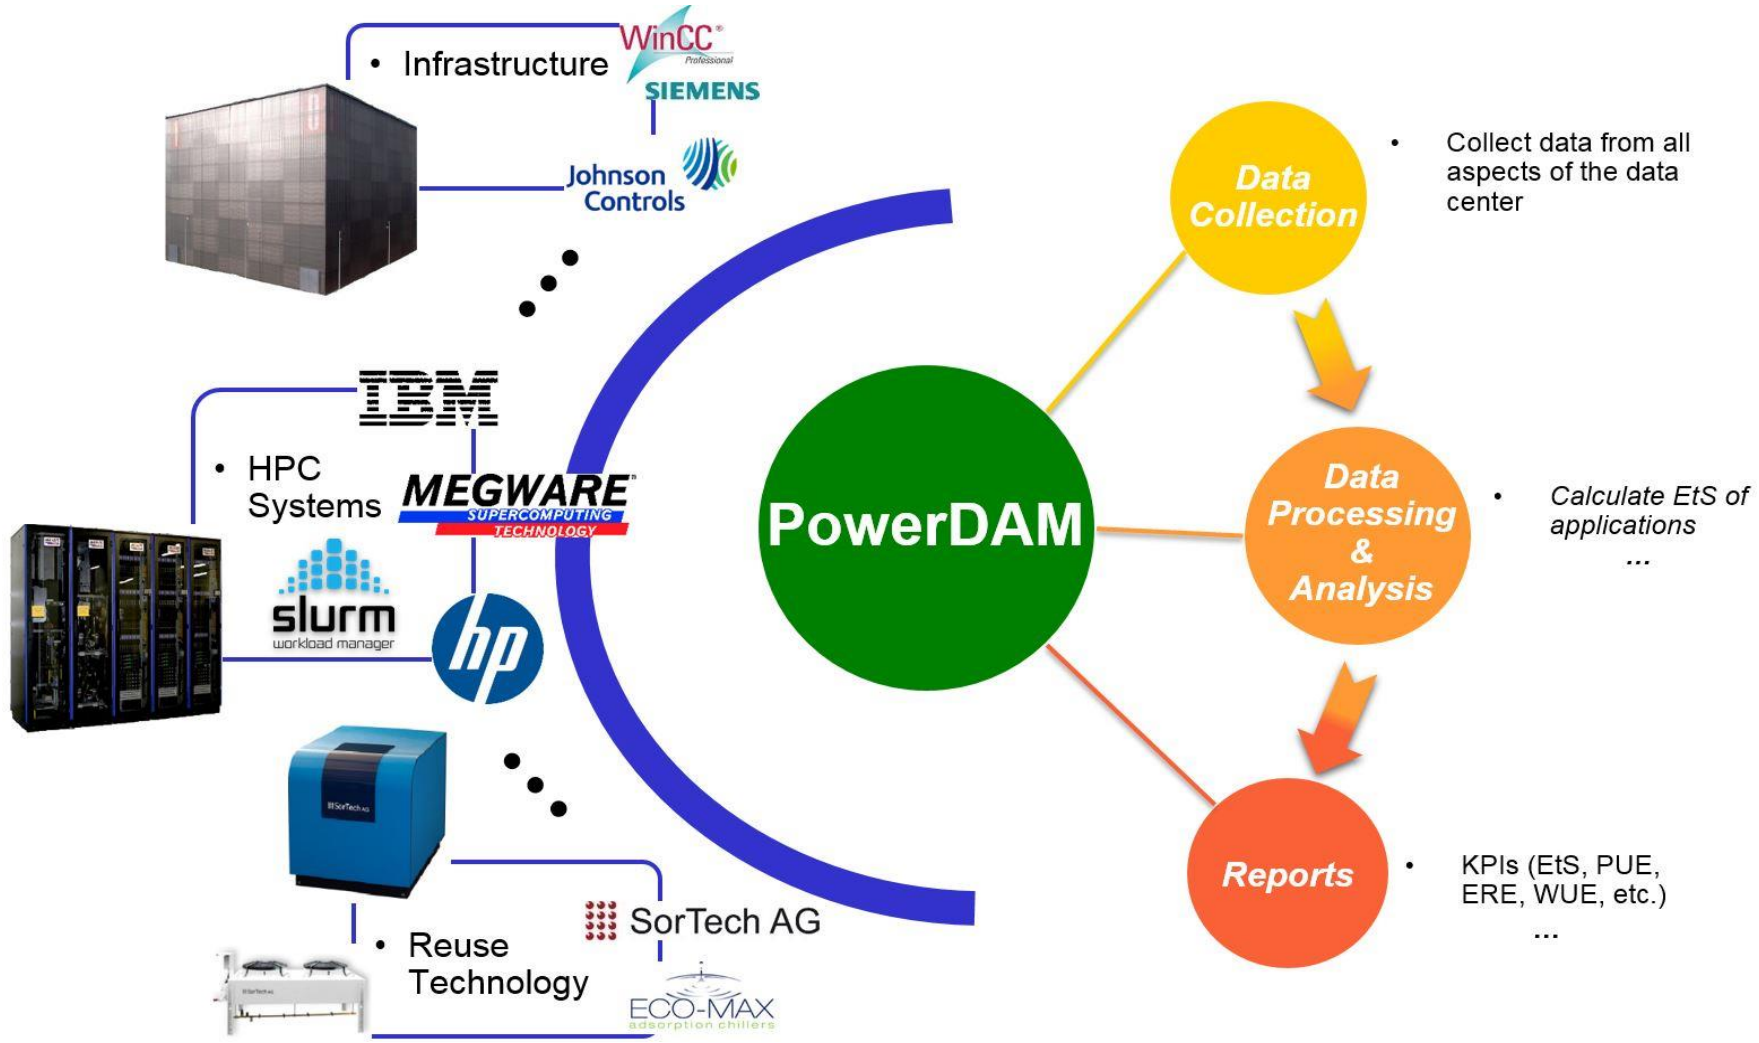

Hayk Shoukourian, Torsten Wilde, Axel Auweter, Arndt Bode: "Monitoring Power Data: A first step towards a unified energy efficiency evaluation toolset for HPC data centers" published in Environmental Modelling & Software (Thematic issue on Modelling and evaluating the sustainability of smart solutions), Volume 56, June 2014, Pages 13–26; DOI: http://dx.doi.org/10.1016/j.envsoft.2013.11.011

SuperMUC under a Power Bound

Power Contract Details

- Yearly payment consists of two parts:
  - Connection fee – Charged by the owner of the physical power lines, fixed charges per kWh
    - Can save 50% of it (≈ 250 k€)
  - Energy Costs – Charged by the power provider, price of used energy, depends on final (yearly) energy consumption (kWh)

- Biggest cost factor is Energy Consumption
  - LRZ Main Goal: **Improve Energy Efficiency**
    - Energy Aware Scheduling
  - LRZ Secondary Goal: Limit Power Peaks to 10% of Data Center Average Power Consumption in billing cycle (15min at LRZ)
    - Analyze past occurrences, avoid in the future (currently manually)

SuperMUC Phase1 and Phase2 Power Profile for 2015 (Turbo Off)

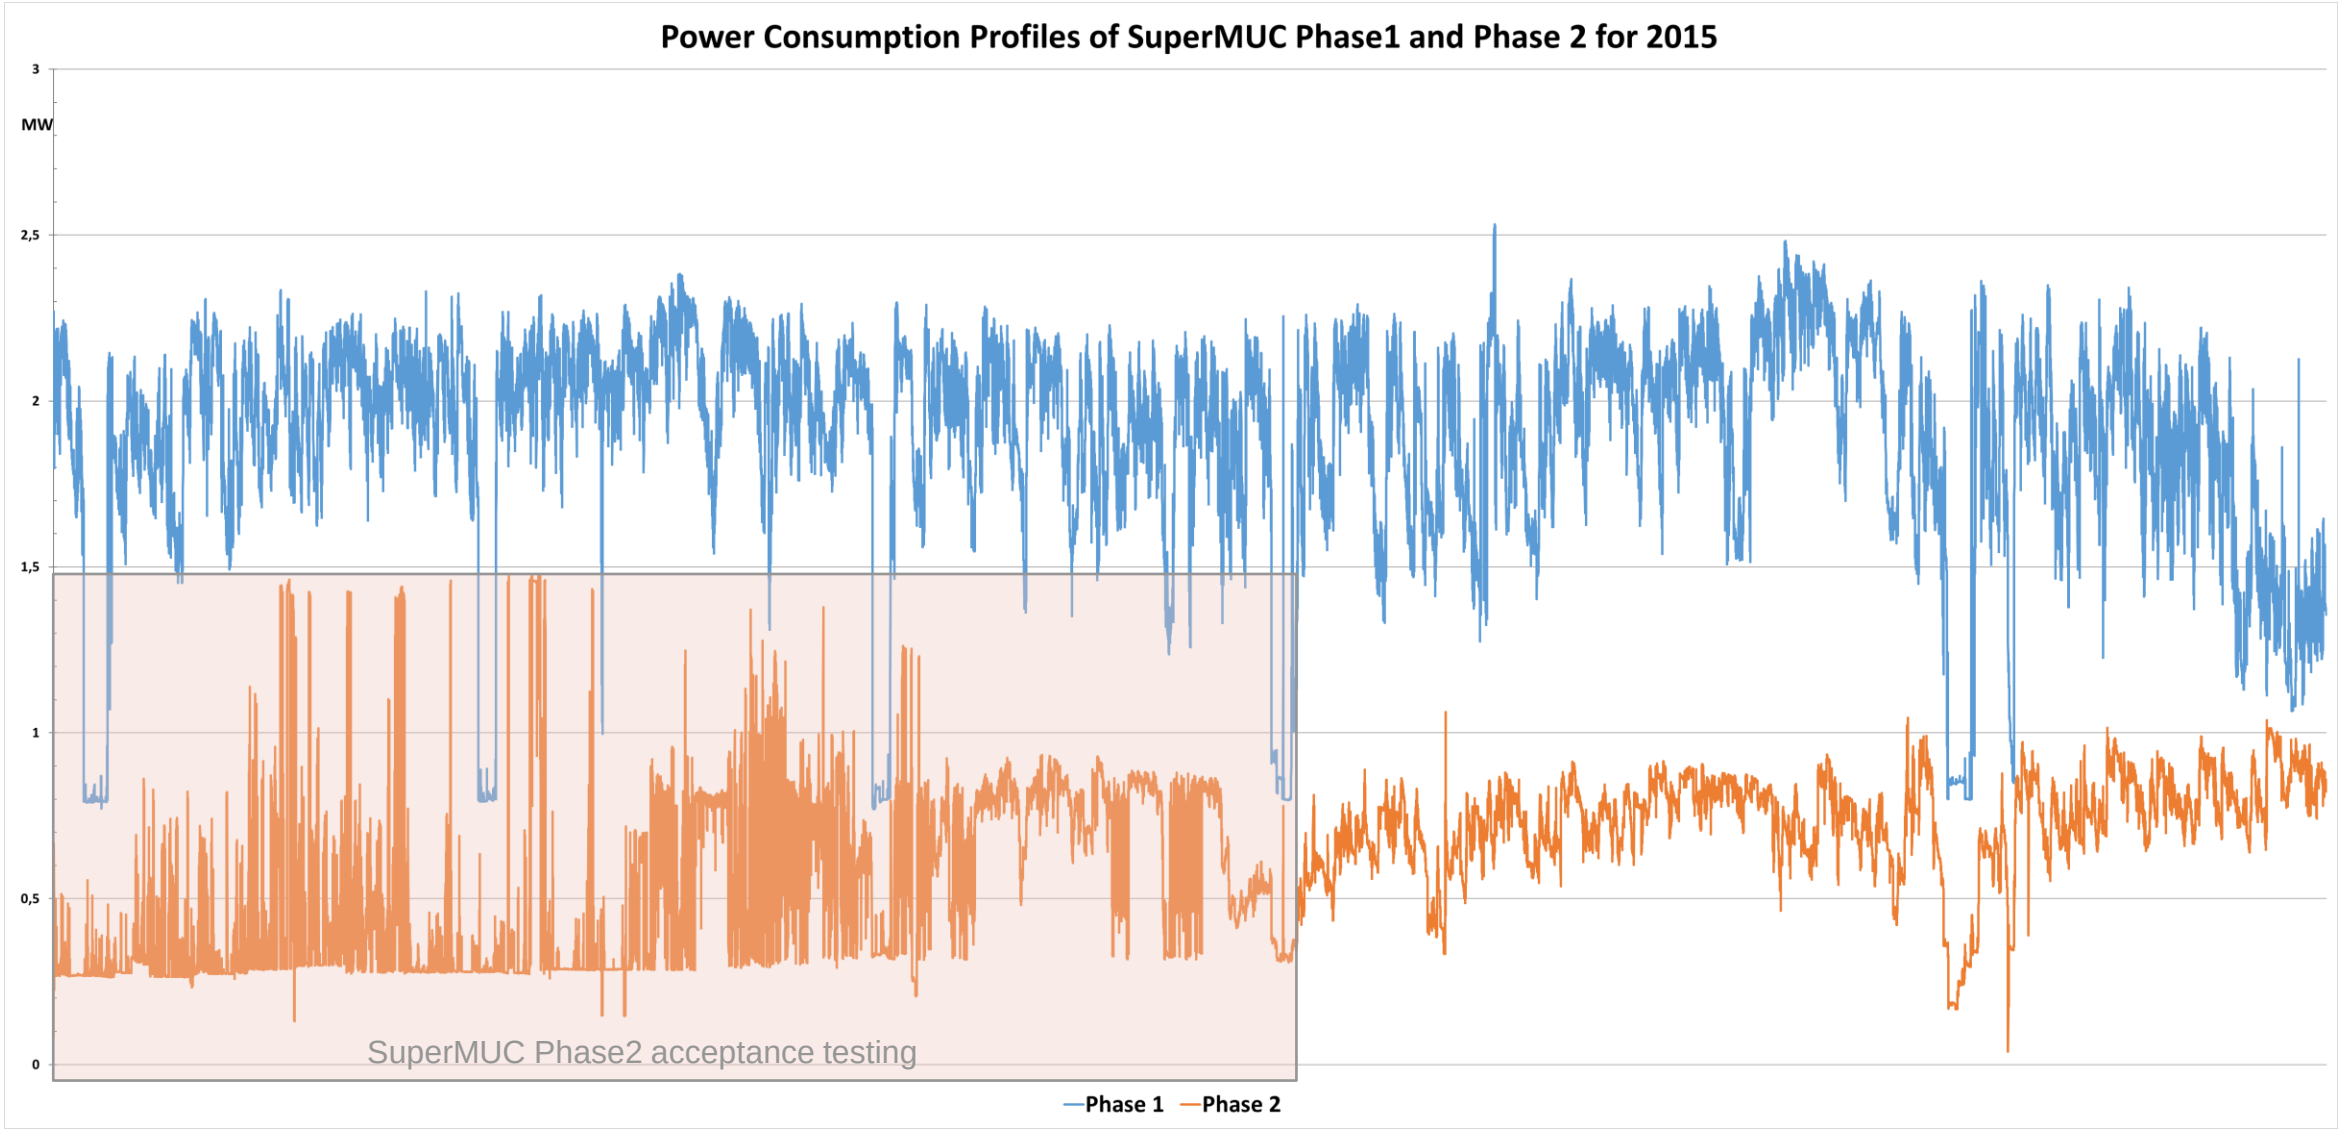

IT is not the only Power Driver in a Data Center

Outside Conditions and LRZ Cooling Efficiency

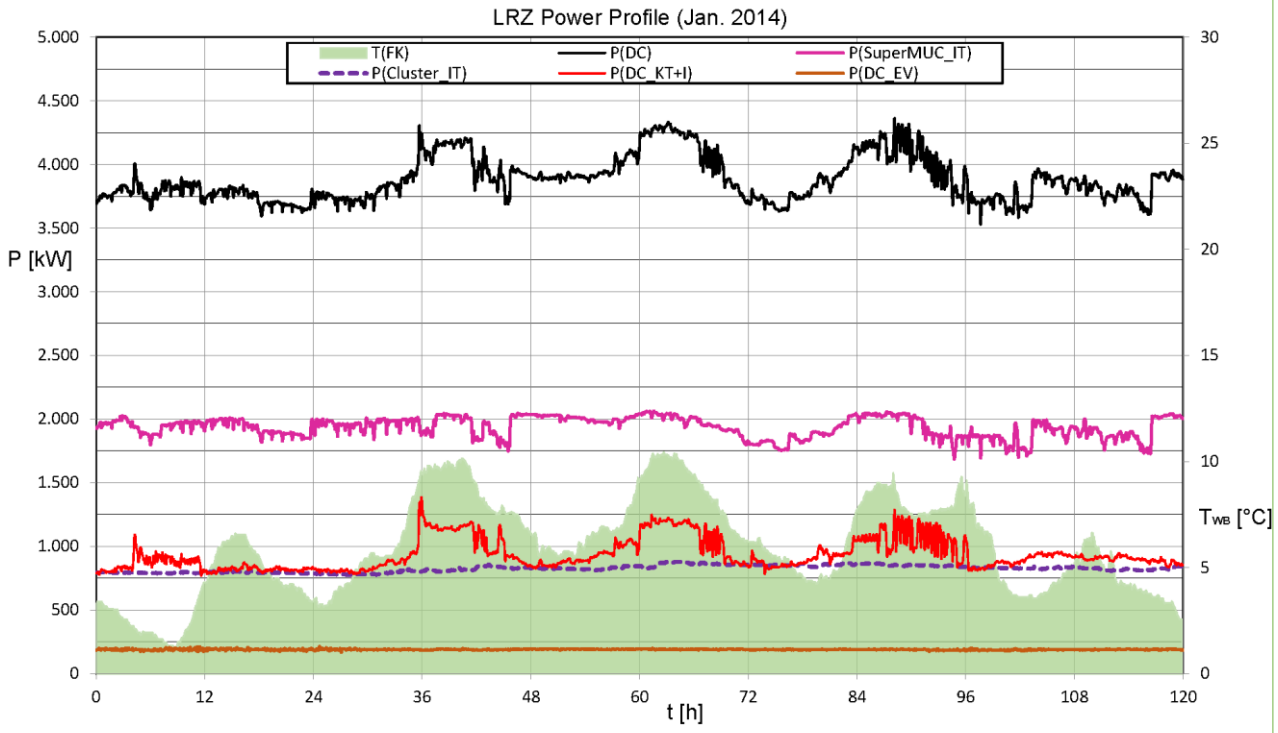

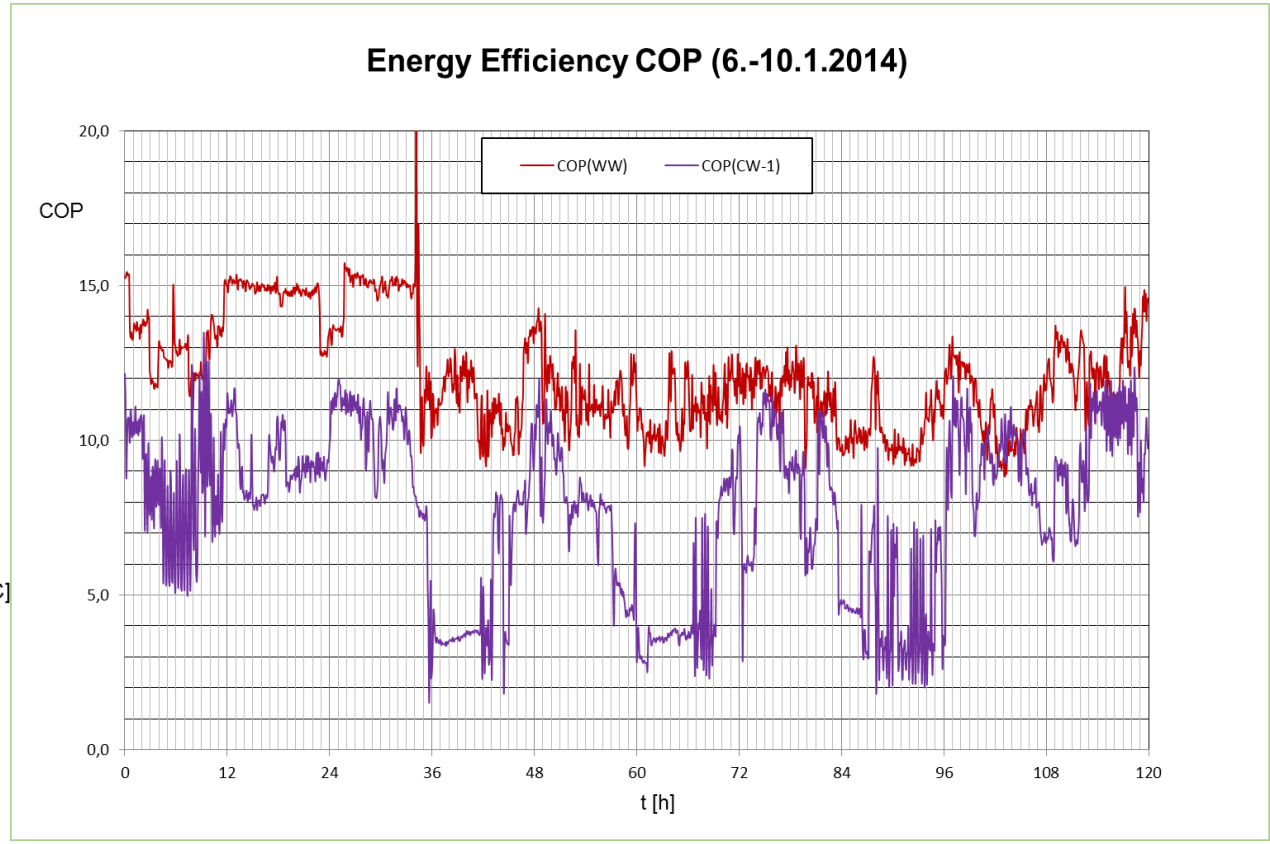

$$COP = \frac{Q}{P}$$

Power Bound Discussion

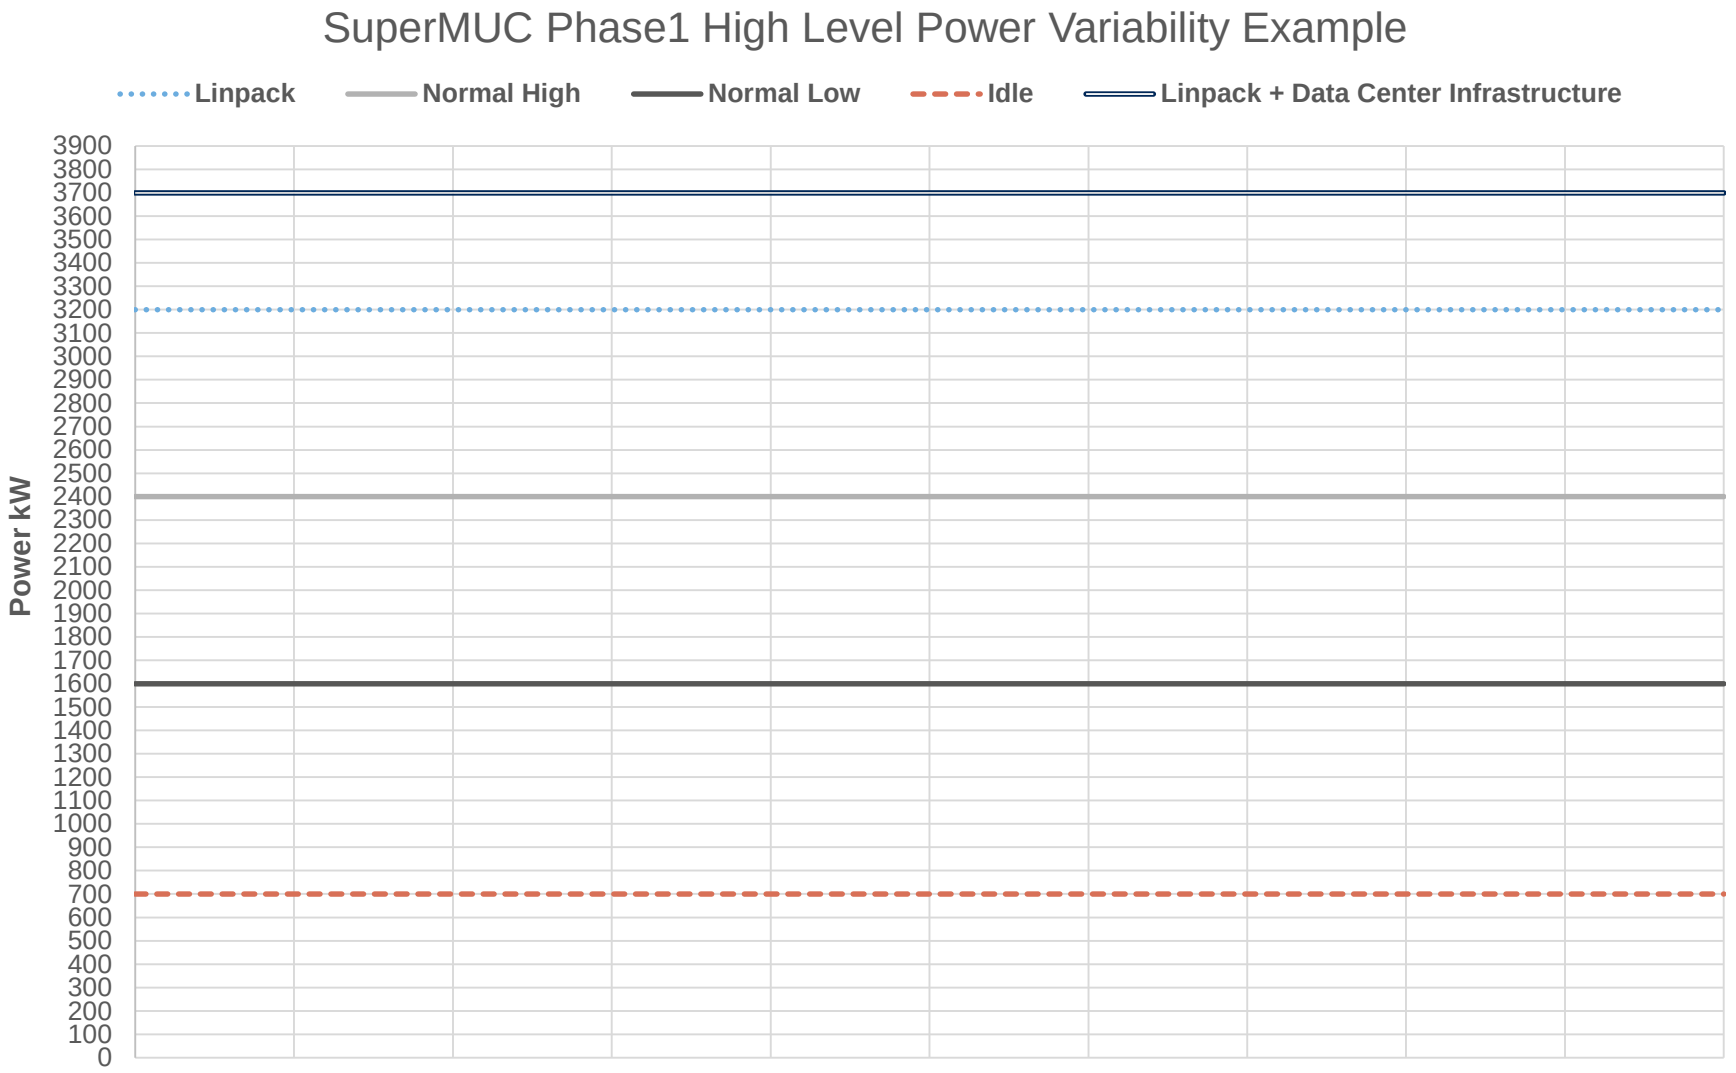

- Where is the Power Bound?
- What is the Power Bound?
- Why do I have a Power Bound?

- *Provision DC for Linpack*
  - Can install additional system for 800kW
  - Over-provisioning
- *Right size DC (provision for normal high)*
  - Simple power bound sufficient
  - Add portable power and cooling for Linpack

Analysis Example

Monitoring vs. Data Analysis

- Data collected by building automation system can be inaccurate for system analysis
  - Changes of +-X are ignored (not recorded) since value is used for control decision
  - Sensor readout frequency limited by connectivity. Same value is recorded for in-between timestamps.
- Gaps in data not considered critical
- Data quality and accuracy is not guaranteed

- Currently thinking about data verification and validation options

Machine Learning - COP Prediction (March 2015 learned, April predicted)

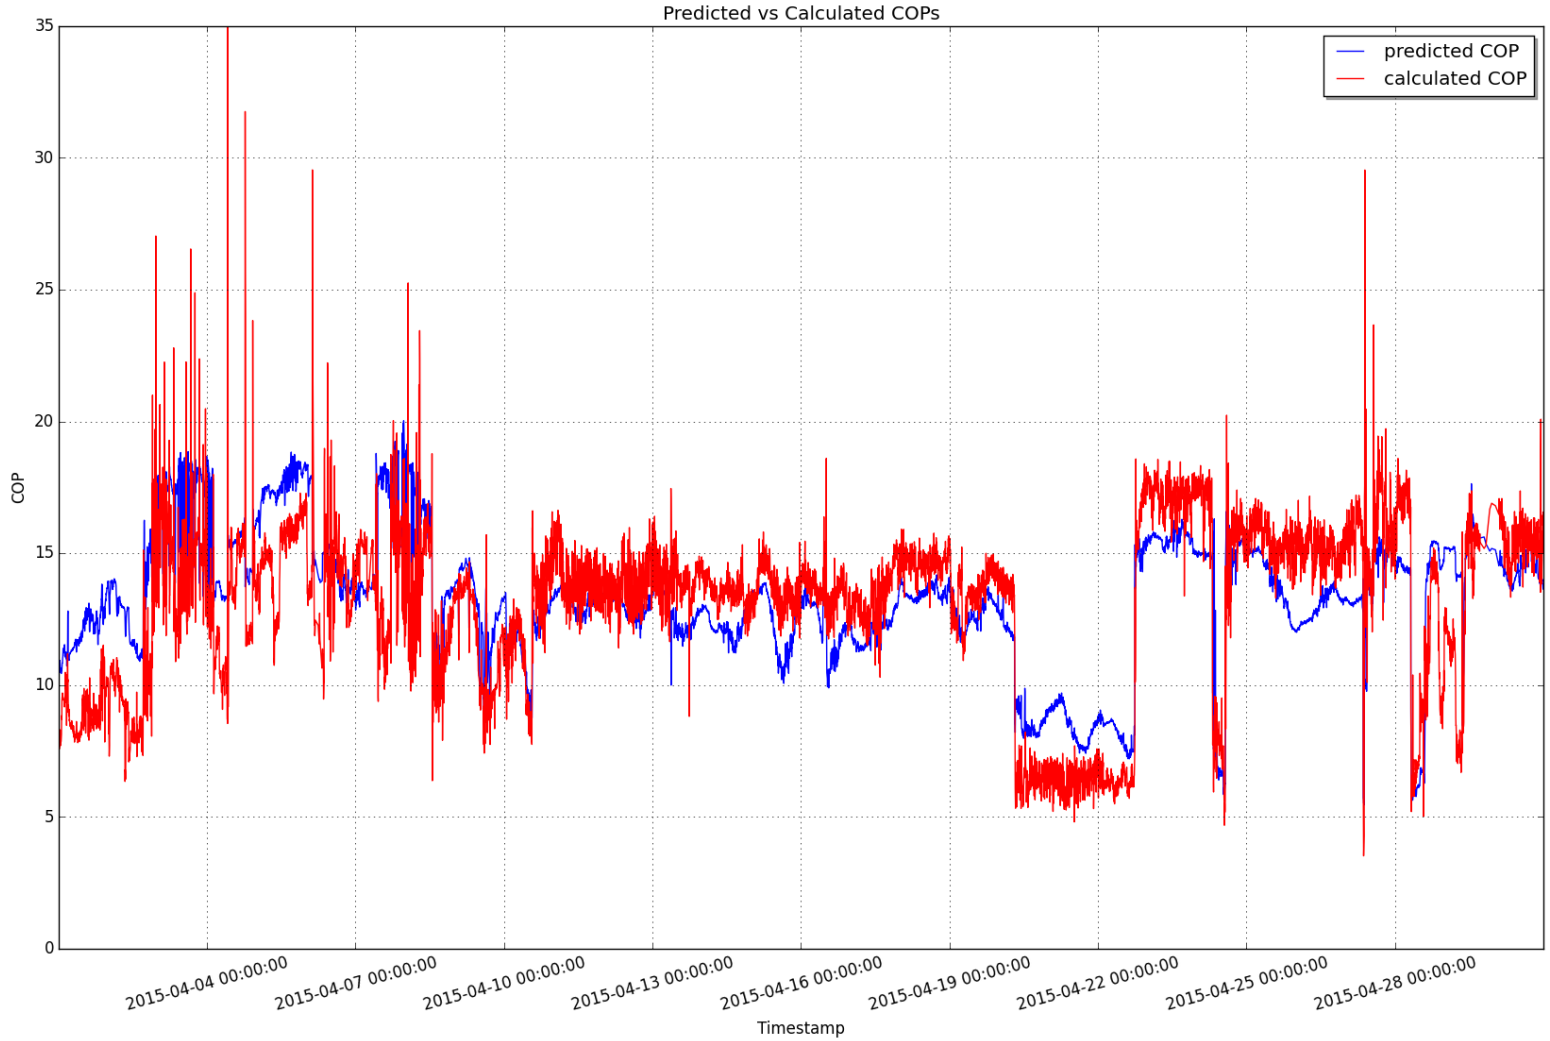

KLT14 (2/2015) Power Profile P(Q)

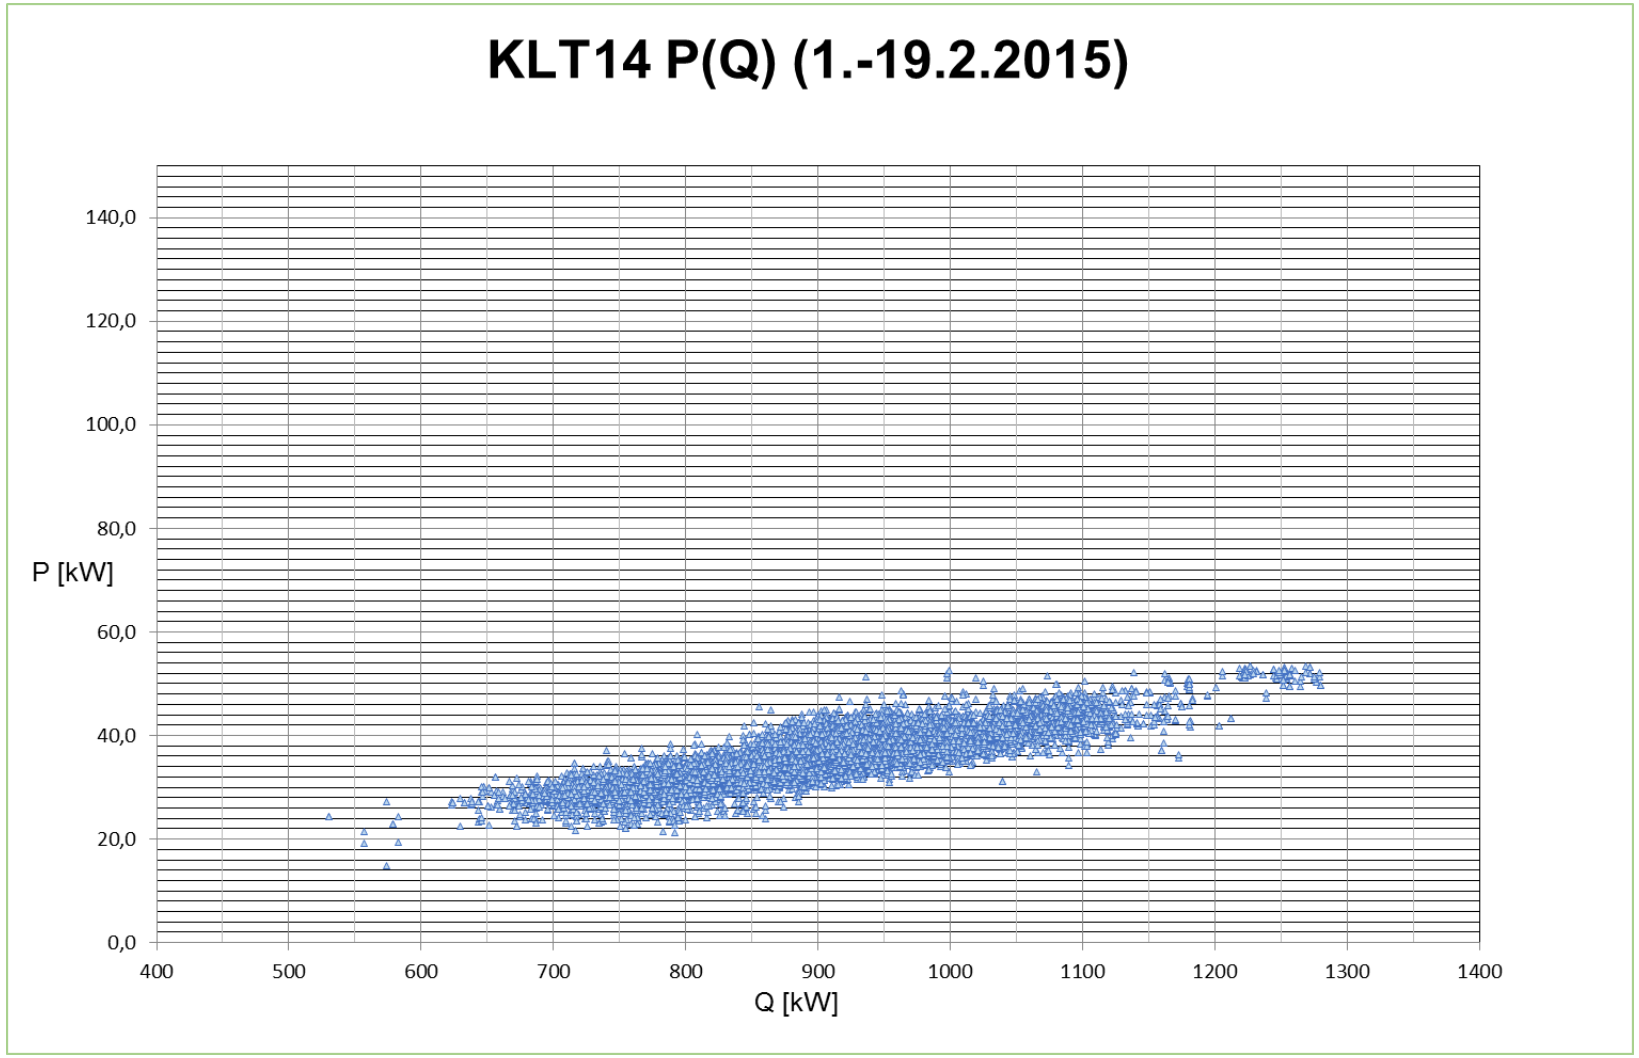

KLT14 (9/2015) Power Profile Q(P)

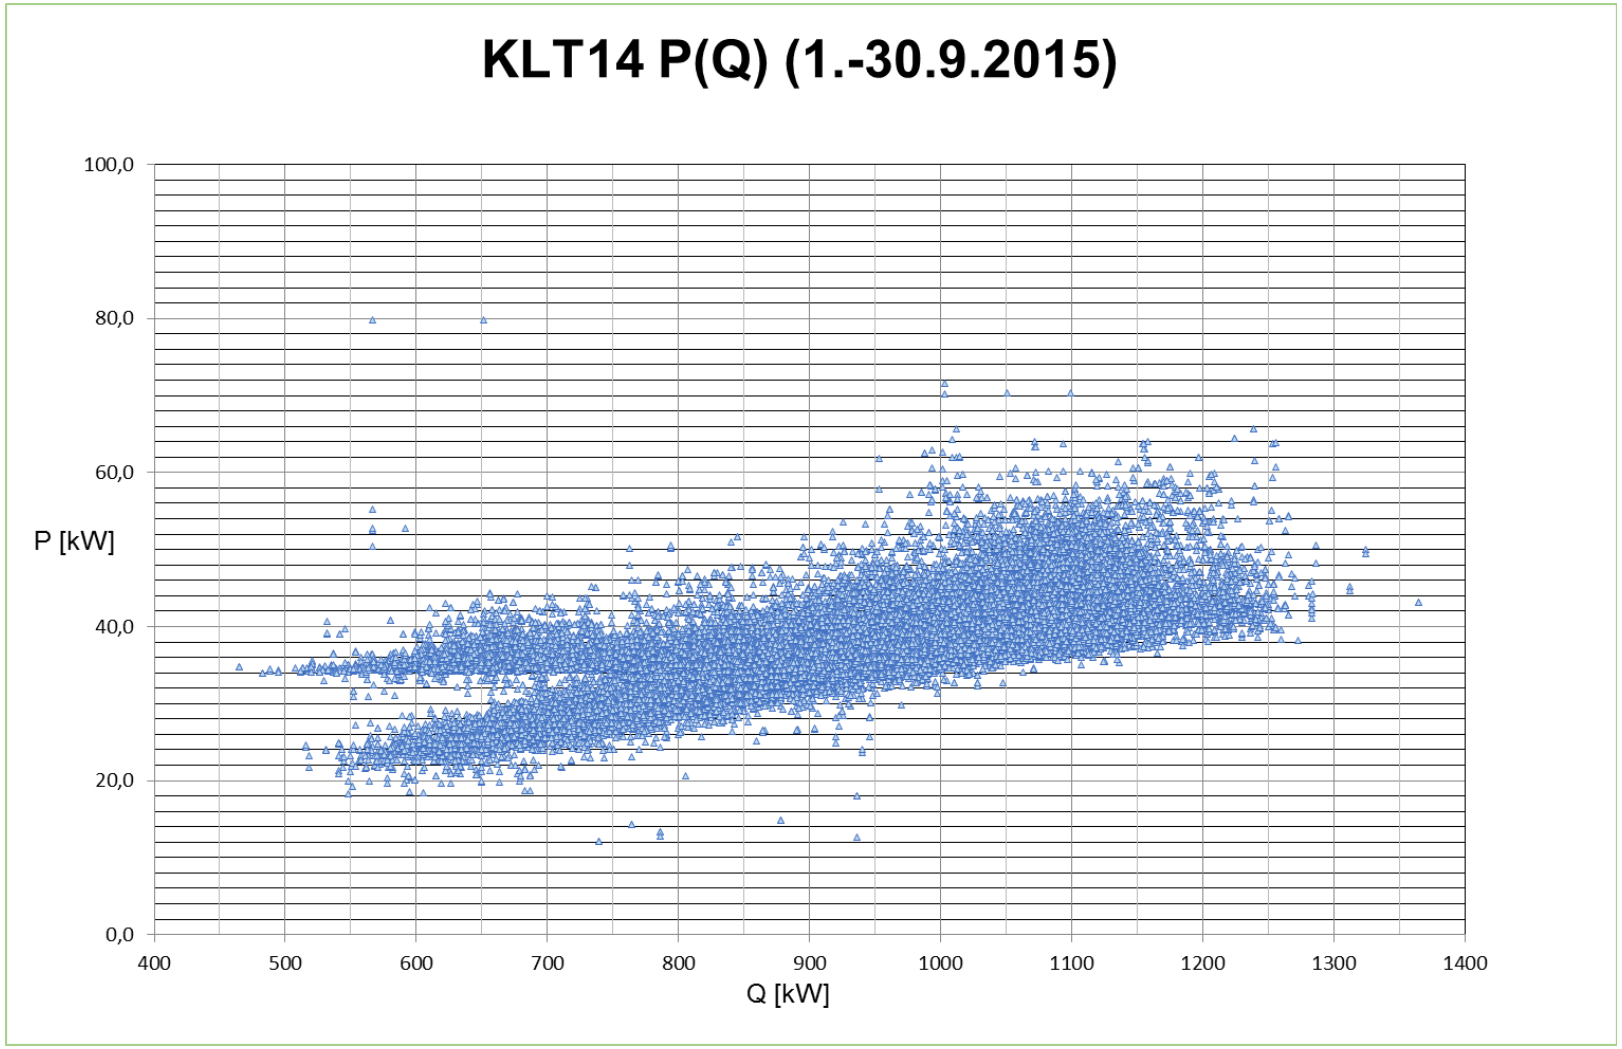

KLT14 (9/2015) Mode of Operations of Pumps

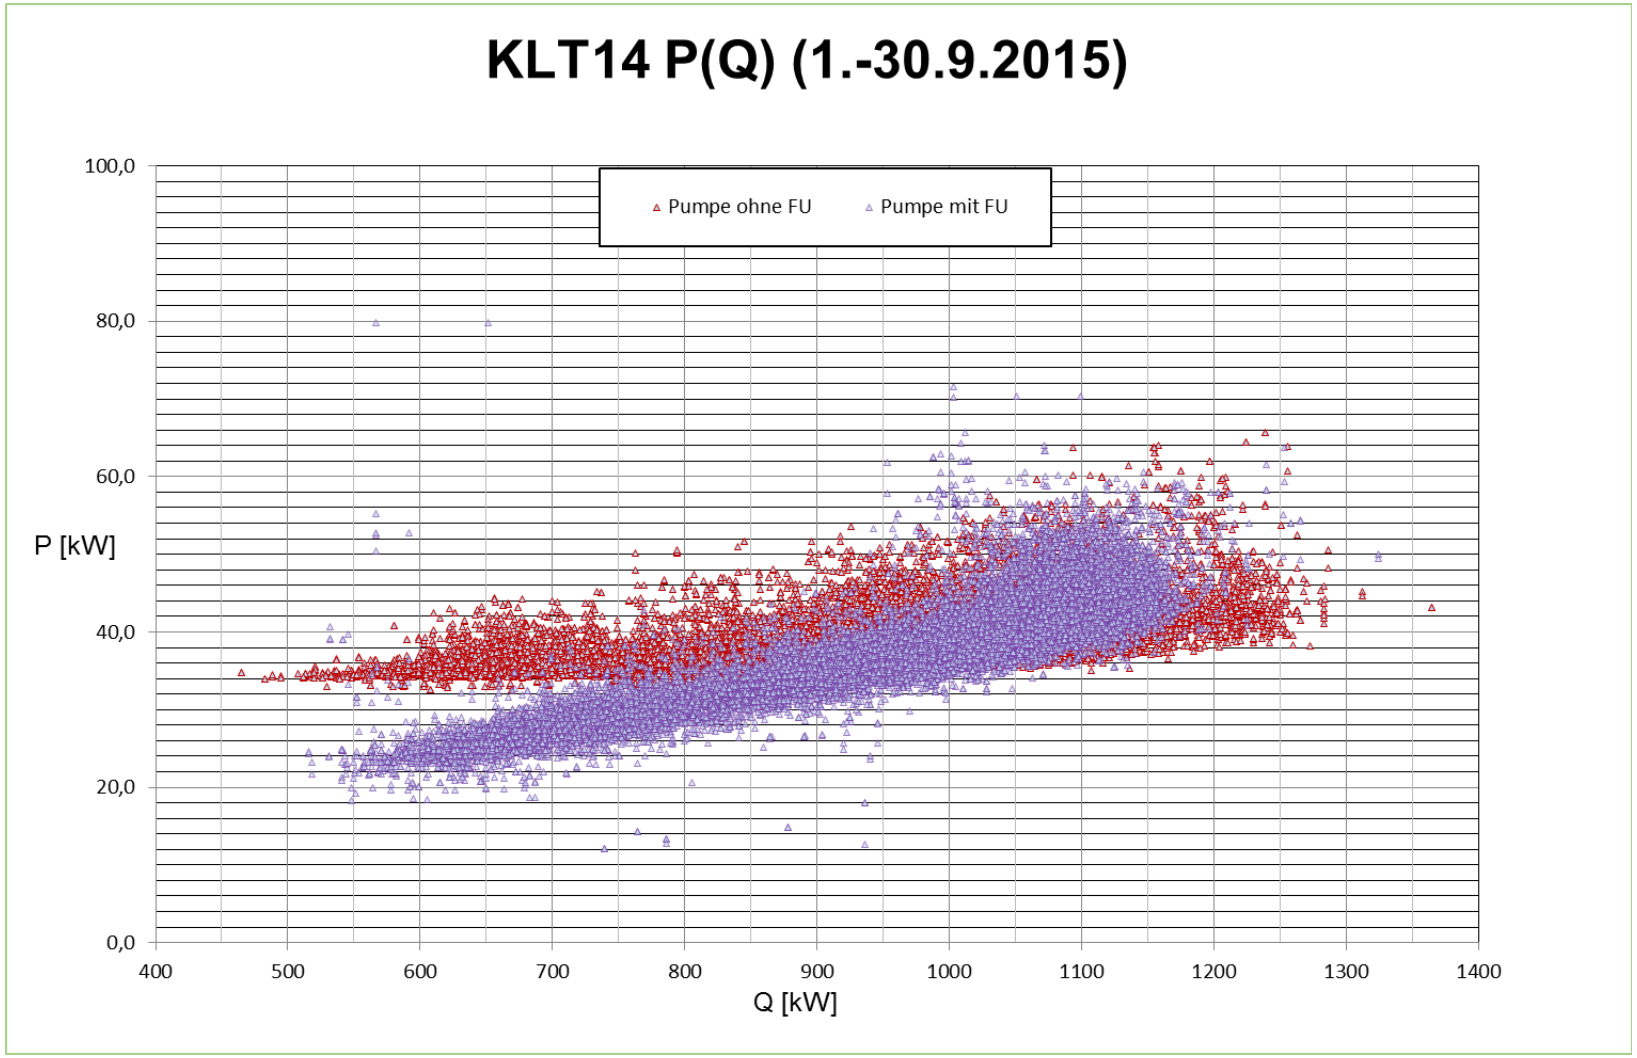

Cold generation vs. Power Consumption (Chillerless Cooling Circuit KLT14)

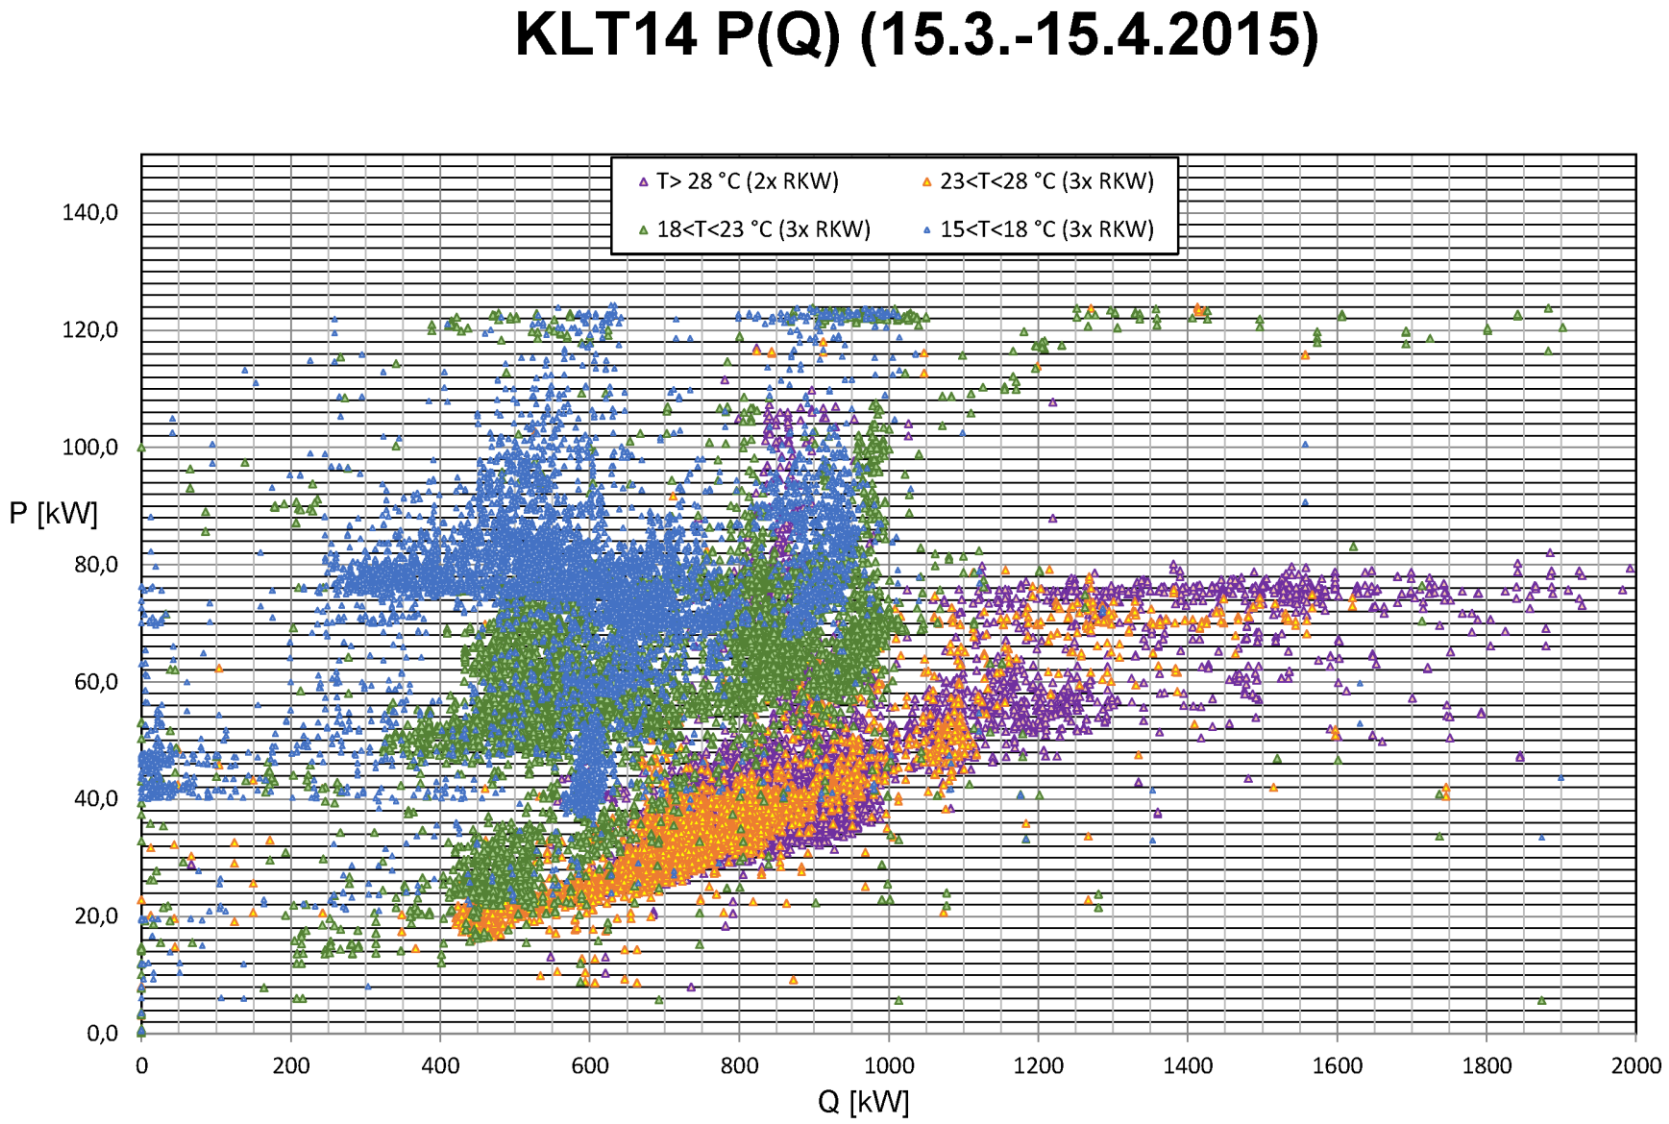

Wishlist

- Data Analytics Support
- Documented Access To All Data Center Monitoring and Management Systems
- Hooks into Data Center Management Systems

Acknowledgment

Part of the work presented here has been carried out within the SIMOPEK project, which has received funding from the German Federal Ministry for Education and Research under grant number 01IH13007A, at the Leibniz Supercomputing Centre (LRZ) with support of the State of Bavaria, Germany, and the Gauss Centre for Supercomputing (GCS).

Thank You!
Torsten.Wilde@lrz.de

Torsten Wilde

torsten.wilde@lrz.de
www.simopek.de

Leibniz-Rechenzentrum
der Bayerischen Akademie der Wissenschaften