Open. Together.

OCP
REGIONAL
SUMMIT

Case Studies

Building and Operating an OCP Data Center at Small Scale

Dr Jon Summers, Research Leader,
Research Institute of Sweden (RISE)

OPEN COMMUNITY®

Agenda

- What is the BodenTypeDC project?
- The design of an OCP style of data center at small scale.
- The opensource software monitoring chain.
- Operational results.
- Live and (hopefully) interactive demonstration of remote operation of the small scale OCP data center.



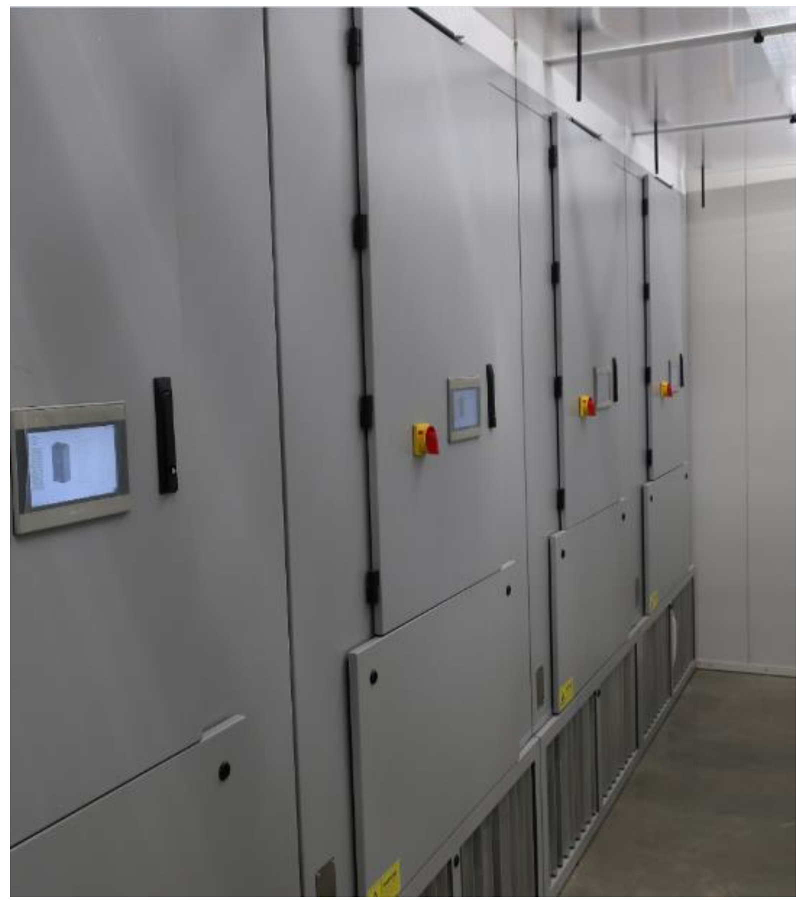
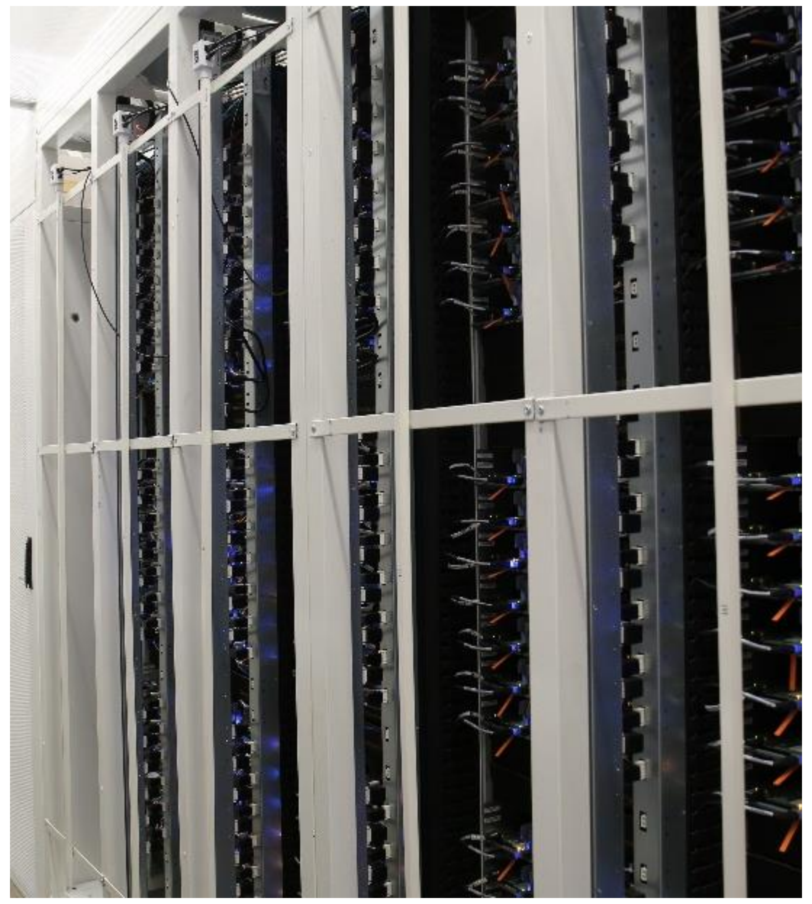

Case Studies

BodenType Data Center H2020 Project

- The BodenType Data Center (BTDC) project is funded by the EU
- To pilot a cost and energy efficient sub 1MW prototype data center.

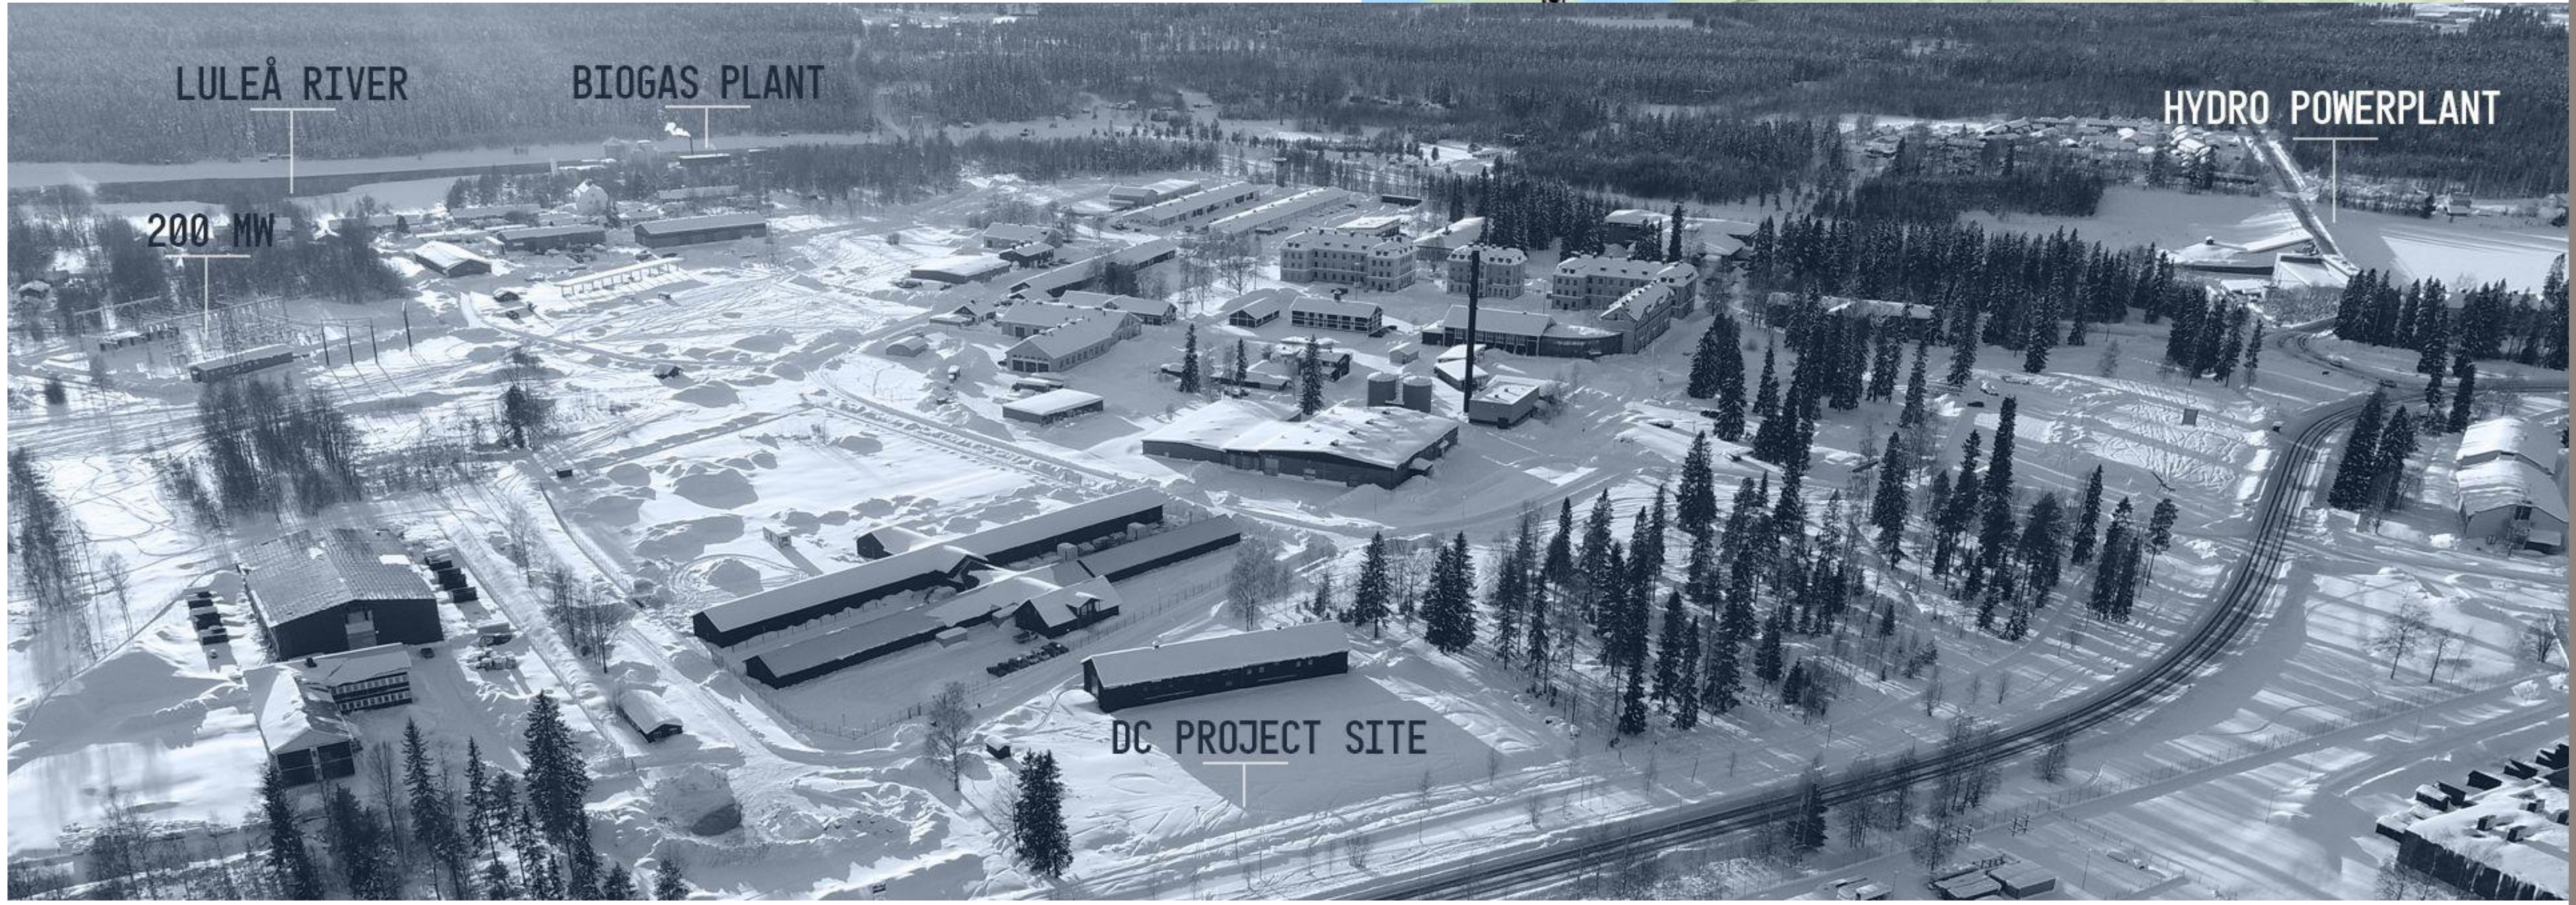

GA 768875

- Three Data Center PODs
- Designed for 500kW of Information Technology.
- Test POD 1 has legacy Open Compute Project IT – Windmill.
- POD 2 and 3 are designed for HPC/GPU test partners for a target 350kW.
- **This project has received funding from the European Union´s Horizon 2020 Research and Innovation programme under grant agreement No 768875.**

Visit https://bodentypeDC.eu

BodenType DC H2020 Project

- Rapid growth in data center energy consumption in Europe from 86TWh in 2013 to an anticipated 104TWh in 2020 (reference P.Bertoldi from the EU JRC).
- Aim of BodenType DC project to create a pilot data center that is cost and energy efficient with minimal environmental impact – to demonstrate this requires a comprehensive monitoring and measurement tool.
- Efficiency is focussed on reducing power losses – no UPS, reducing cooling power consumption – using direct air (measured comparison with other methods) and better utilization of the IT systems – workload deployment and management.

P. Bertoldi, The European Programme for Energy Efficiency in Data Centres: The Code of Conduct, EU DG JRC, Institute for Energy and Transport, 2016.

GA 768875

Visit https://bodentypeDC.eu

- Project partners are:

H1 SYSTEMS ONE STEP AHEAD

EcoCooling FRESH AIR COOLING

BODEN BUSINESS AGENCY

Fraunhofer IOSB

BodenType DC H2020 Project

- Rapid growth in data center energy consumptio Europe from 86TWh in 2013 to an anticipat 104TWh in 2020 (reference P.Bertoldi f JRC).
- Aim of BodenType DC project not data center that is cost and ene with minimal environmental impact strate this requires a comprehensive nd measurement tool.
- Efficiency is reducing power losses – no UPS, re g power consumption – using dire red comparison with other d better utilization of the IT systems – deployment and management.

EE-20-2017 - Bringing to market more energy efficient and integrated data centres

Bertoldi, The European Programme for Energy Efficiency in Data Centres: The Code of Conduct, EU DG JRC, Institute for Energy and Transport, 2016.

BODEN TYPE DC ONE

GA 768875

Visit https://bodentypeDC.eu

- Project partners are:

H1 SYSTEMS ONE STEP AHEAD

EcoCooling FRESH AIR COOLING

BODEN BUSINESS AGENCY

Fraunhofer IOSB

RI. SE

The complete BTDC One Building and plan view of POD1

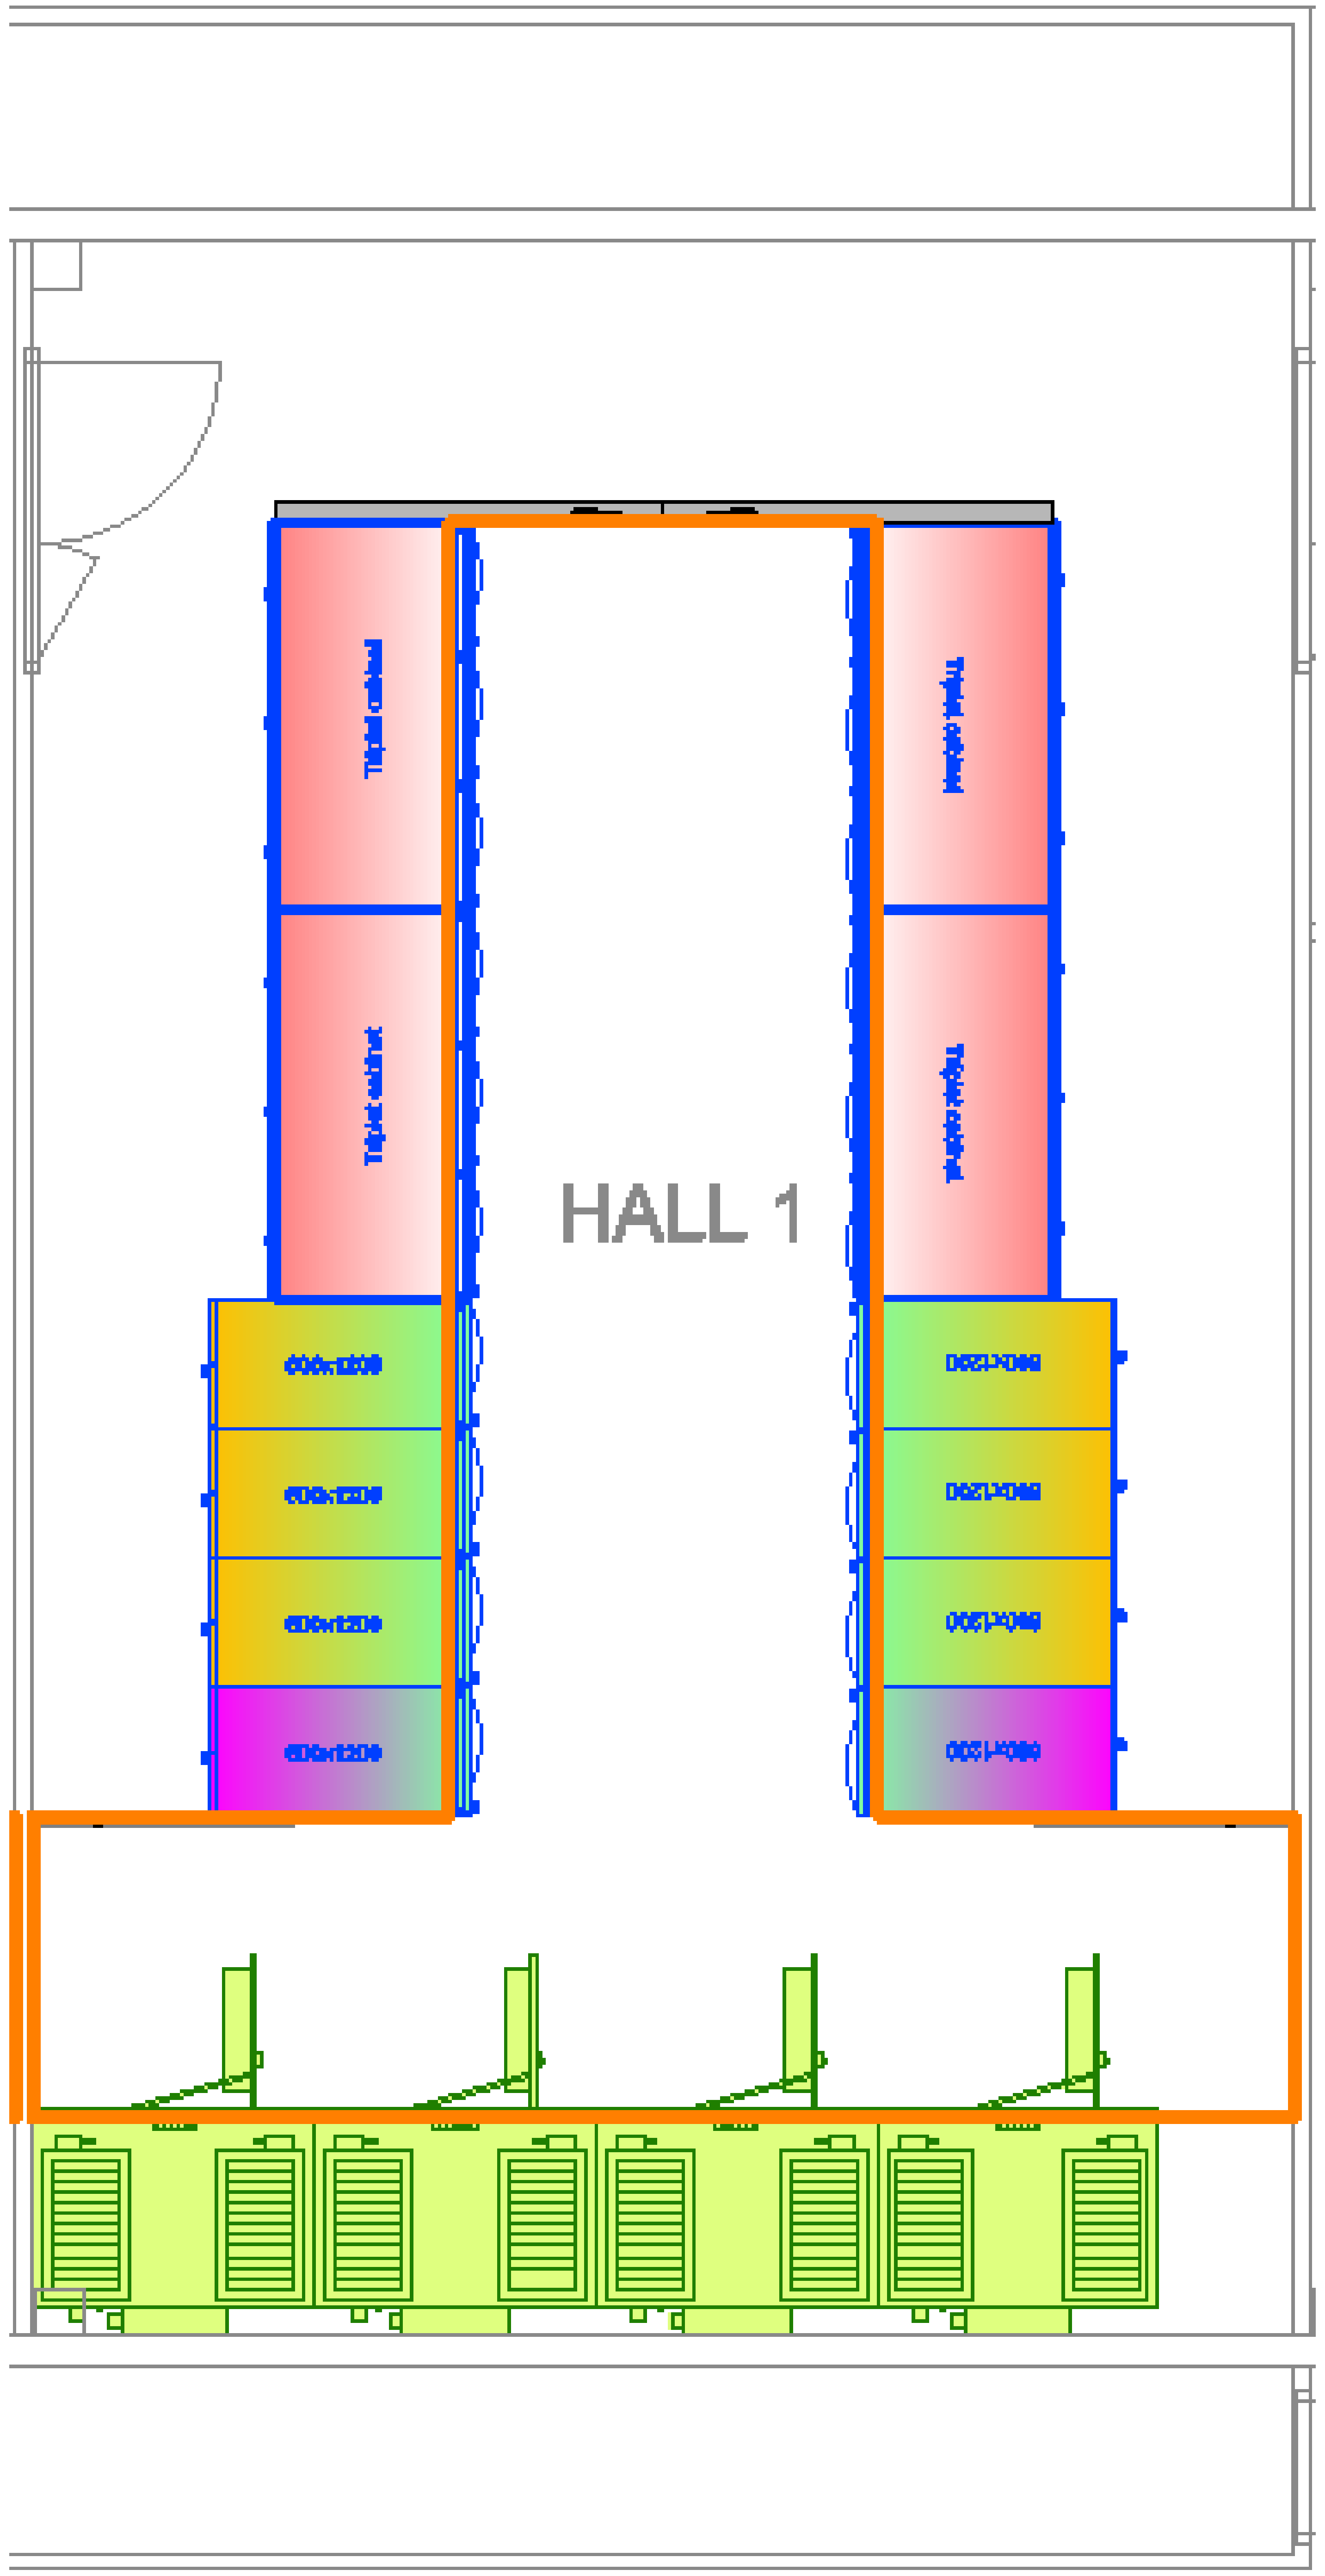

GA 768875

The complete BTDC One Building and plan view of POD1

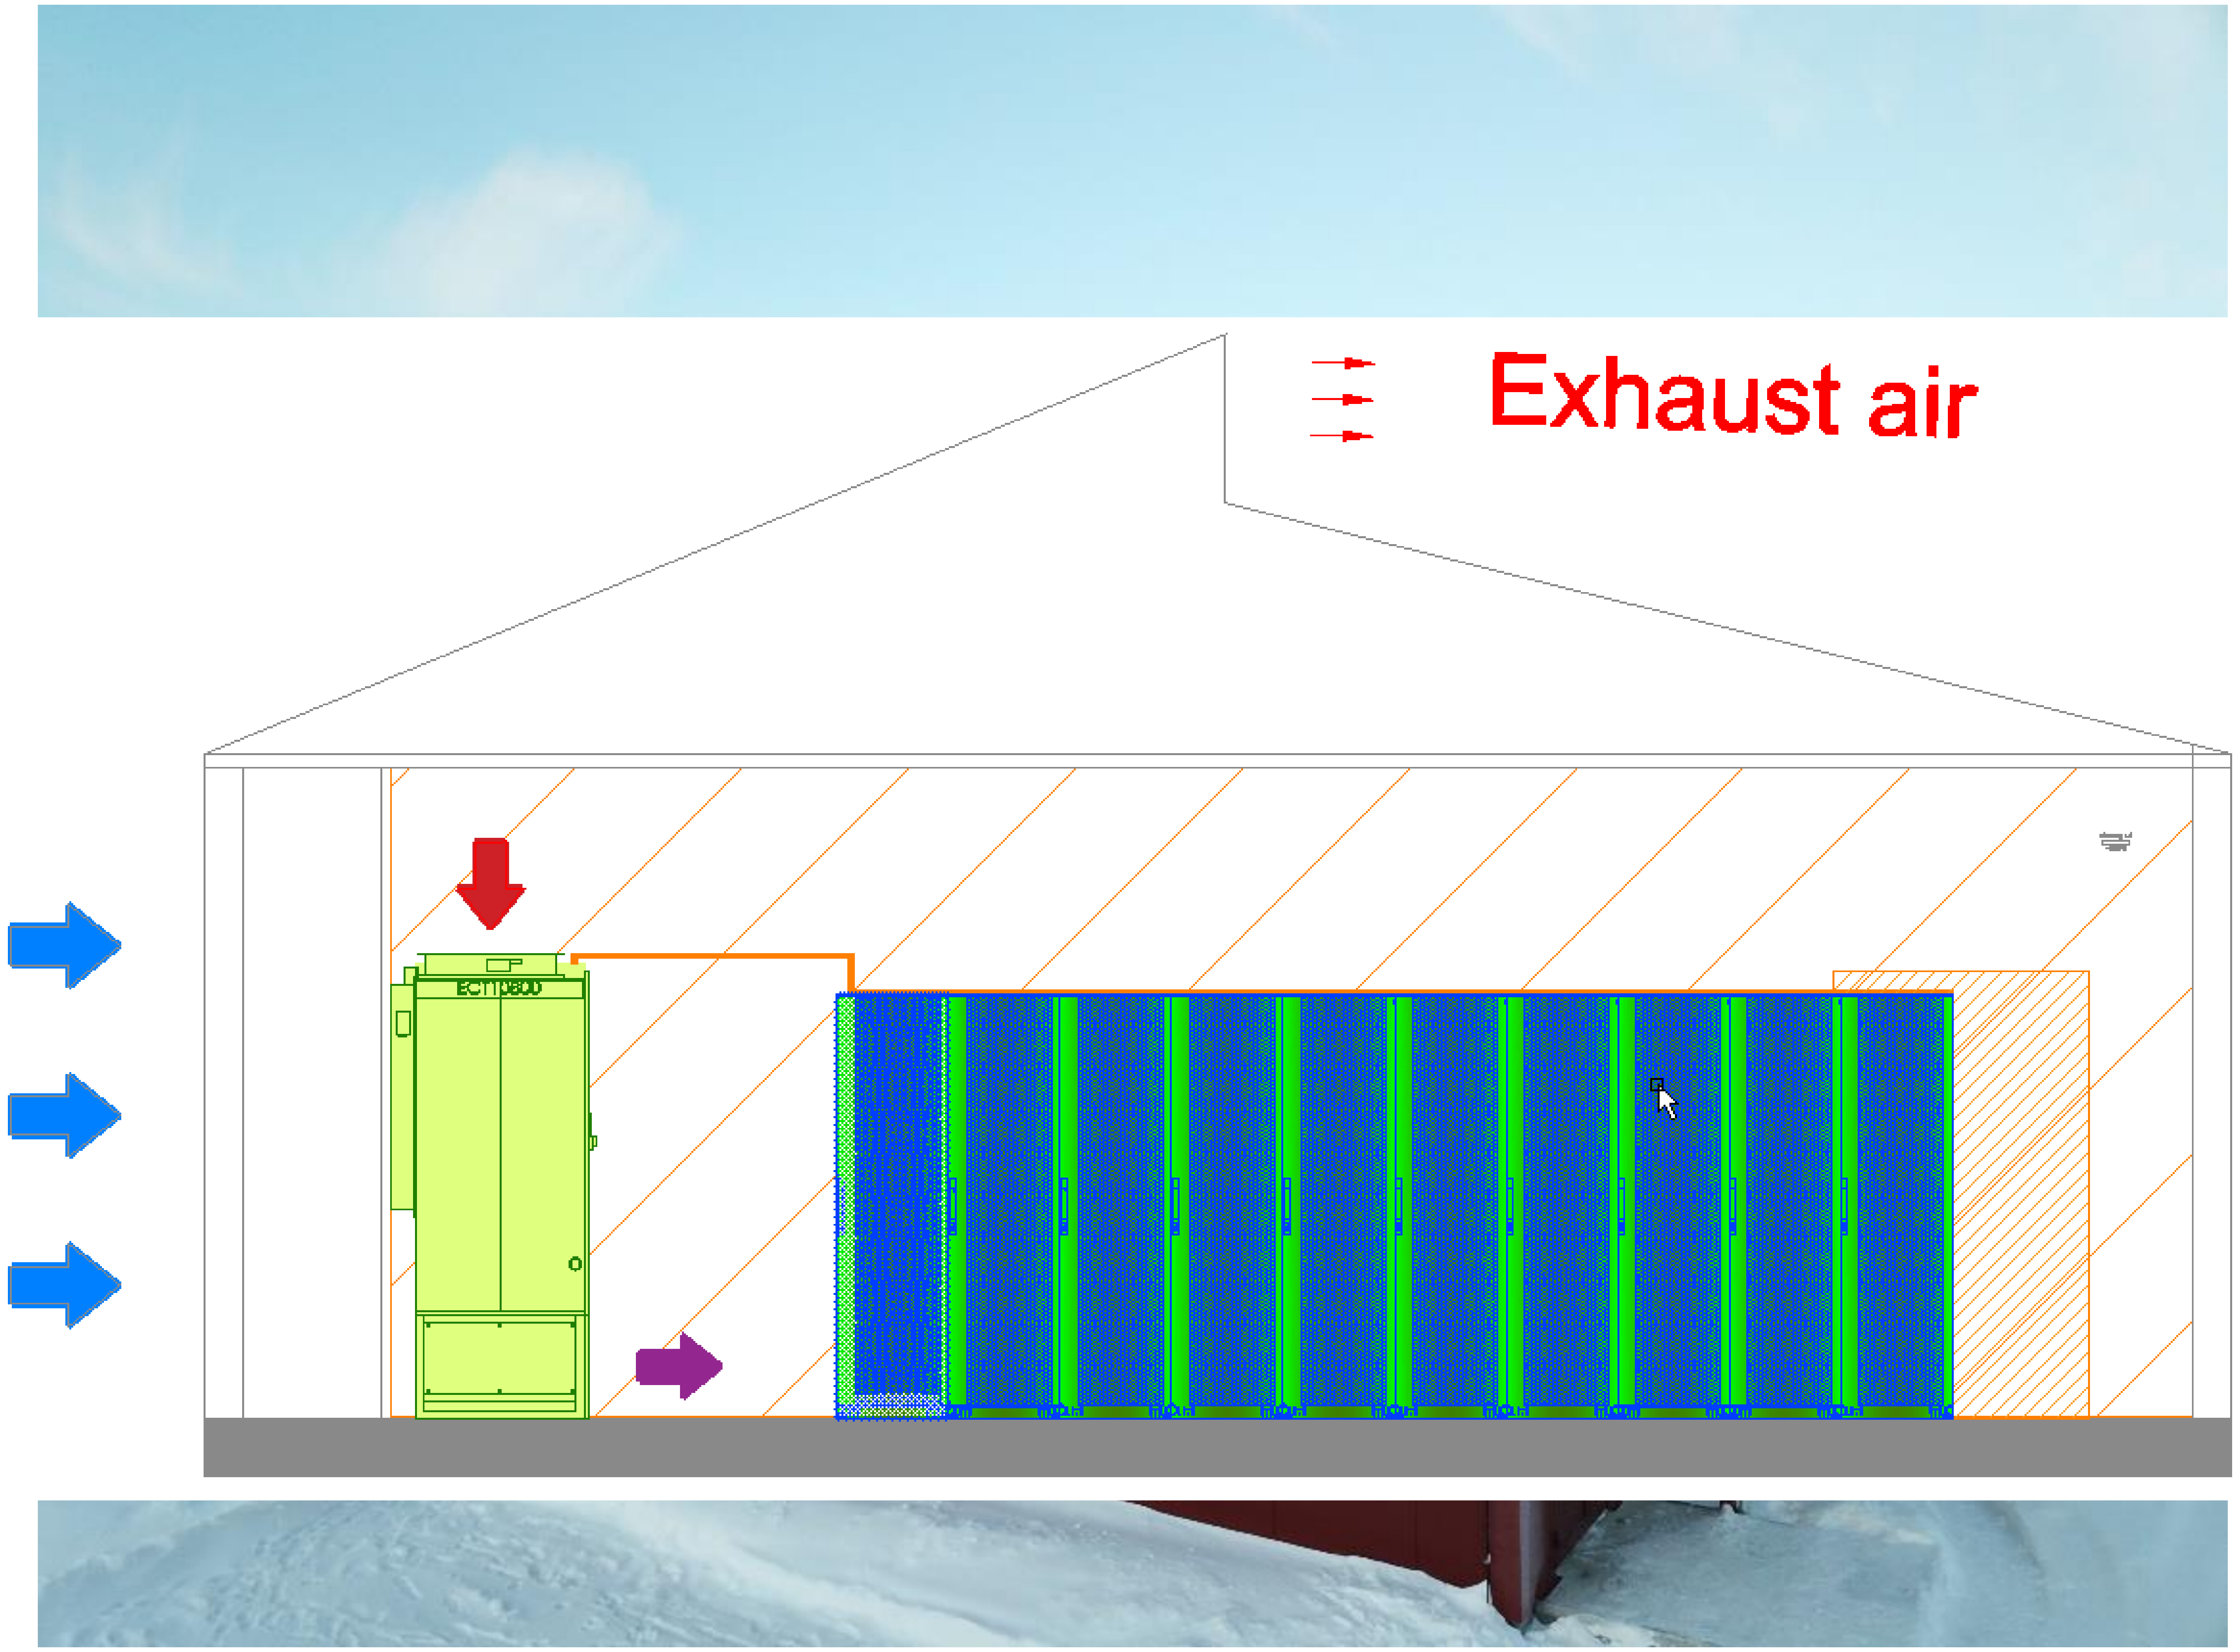

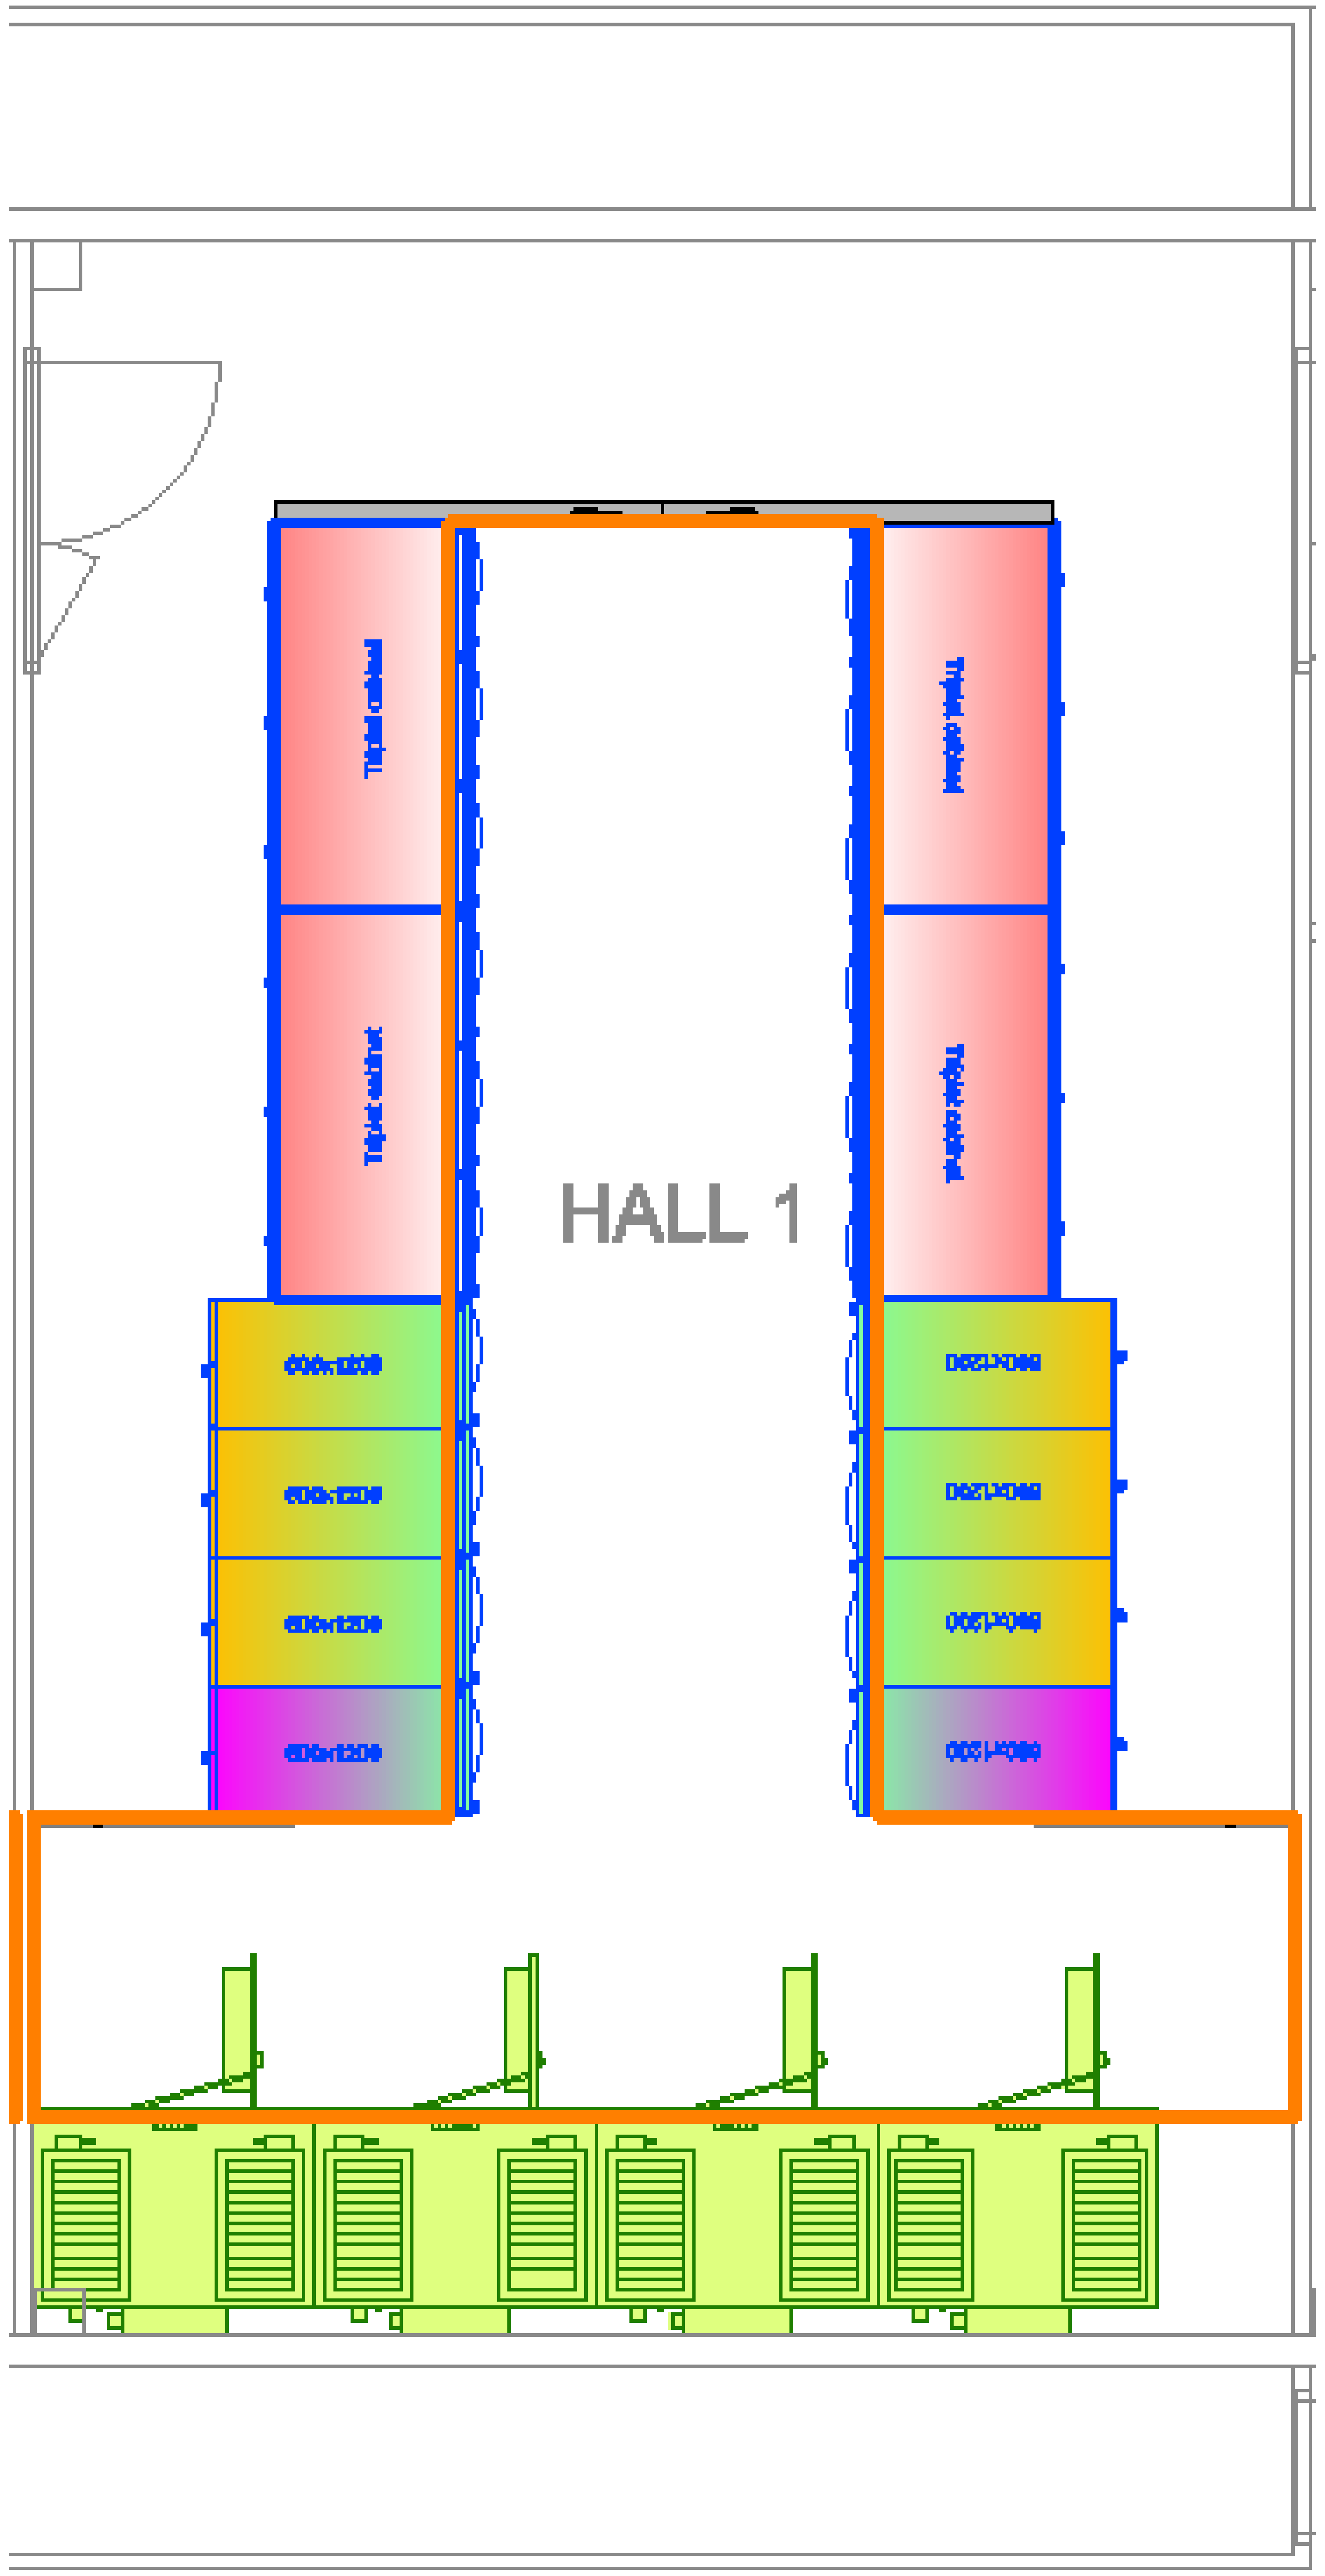

GA 768875

The complete BTDC One Building and plan view of POD1

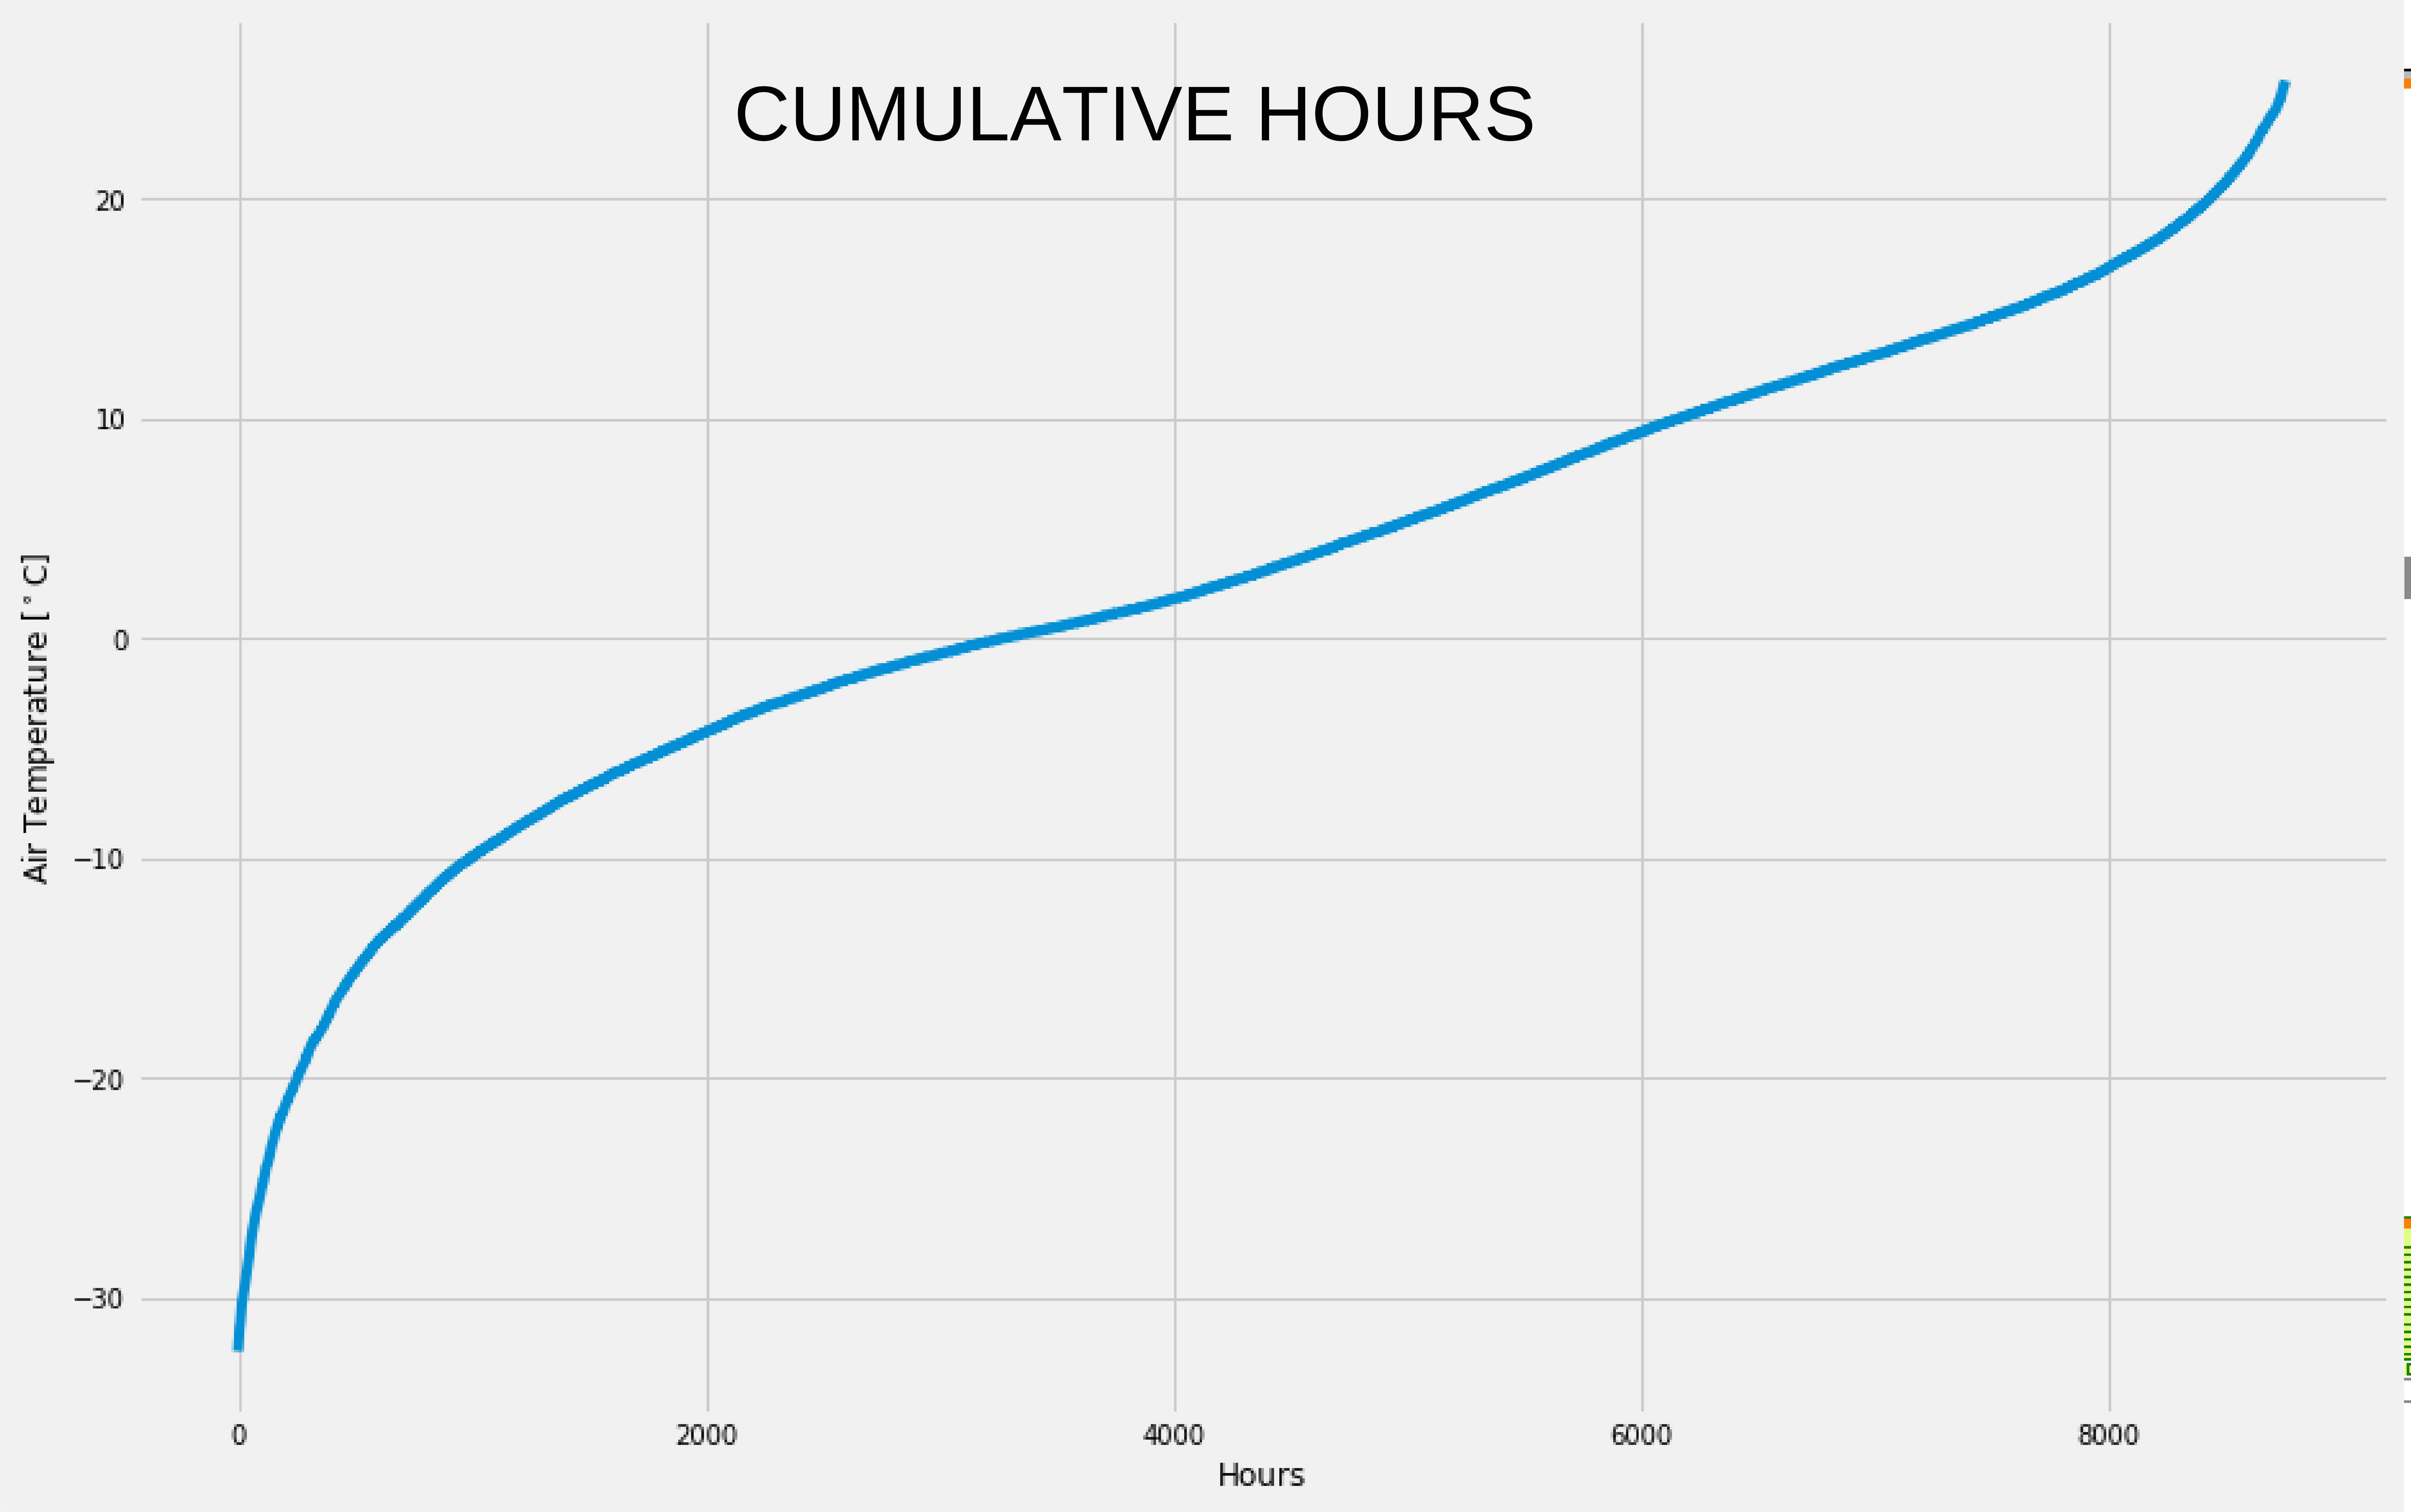

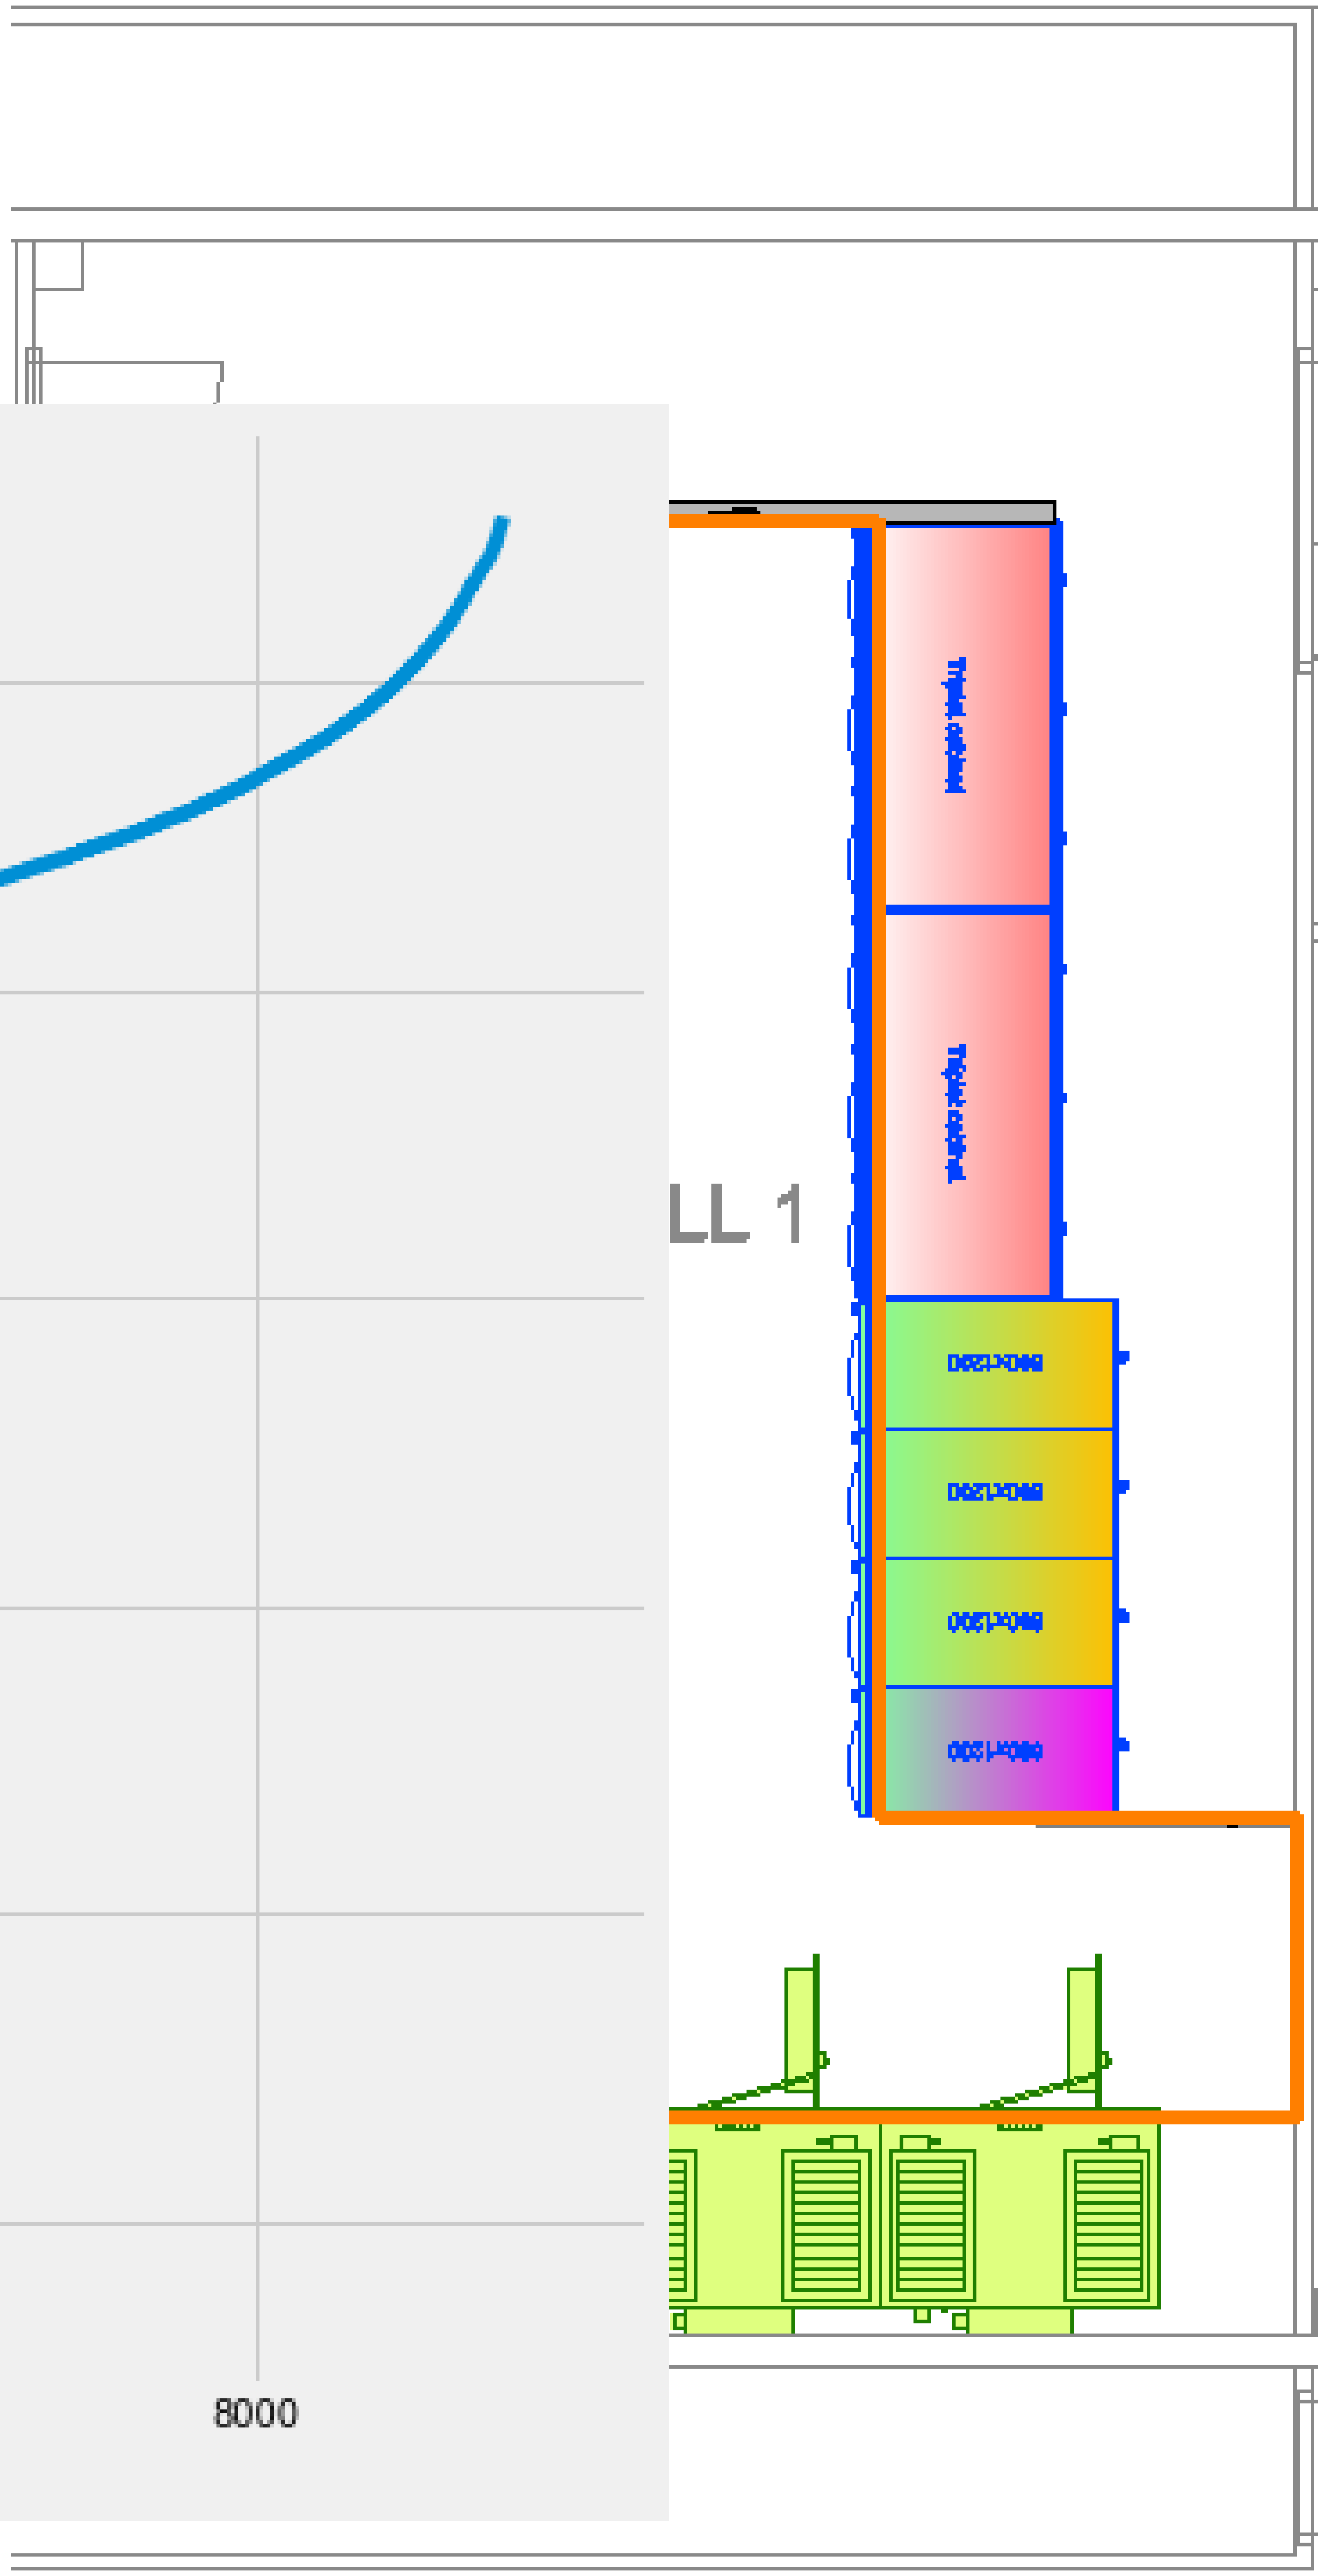

GA 768875

Open source monitoring of BTDC One

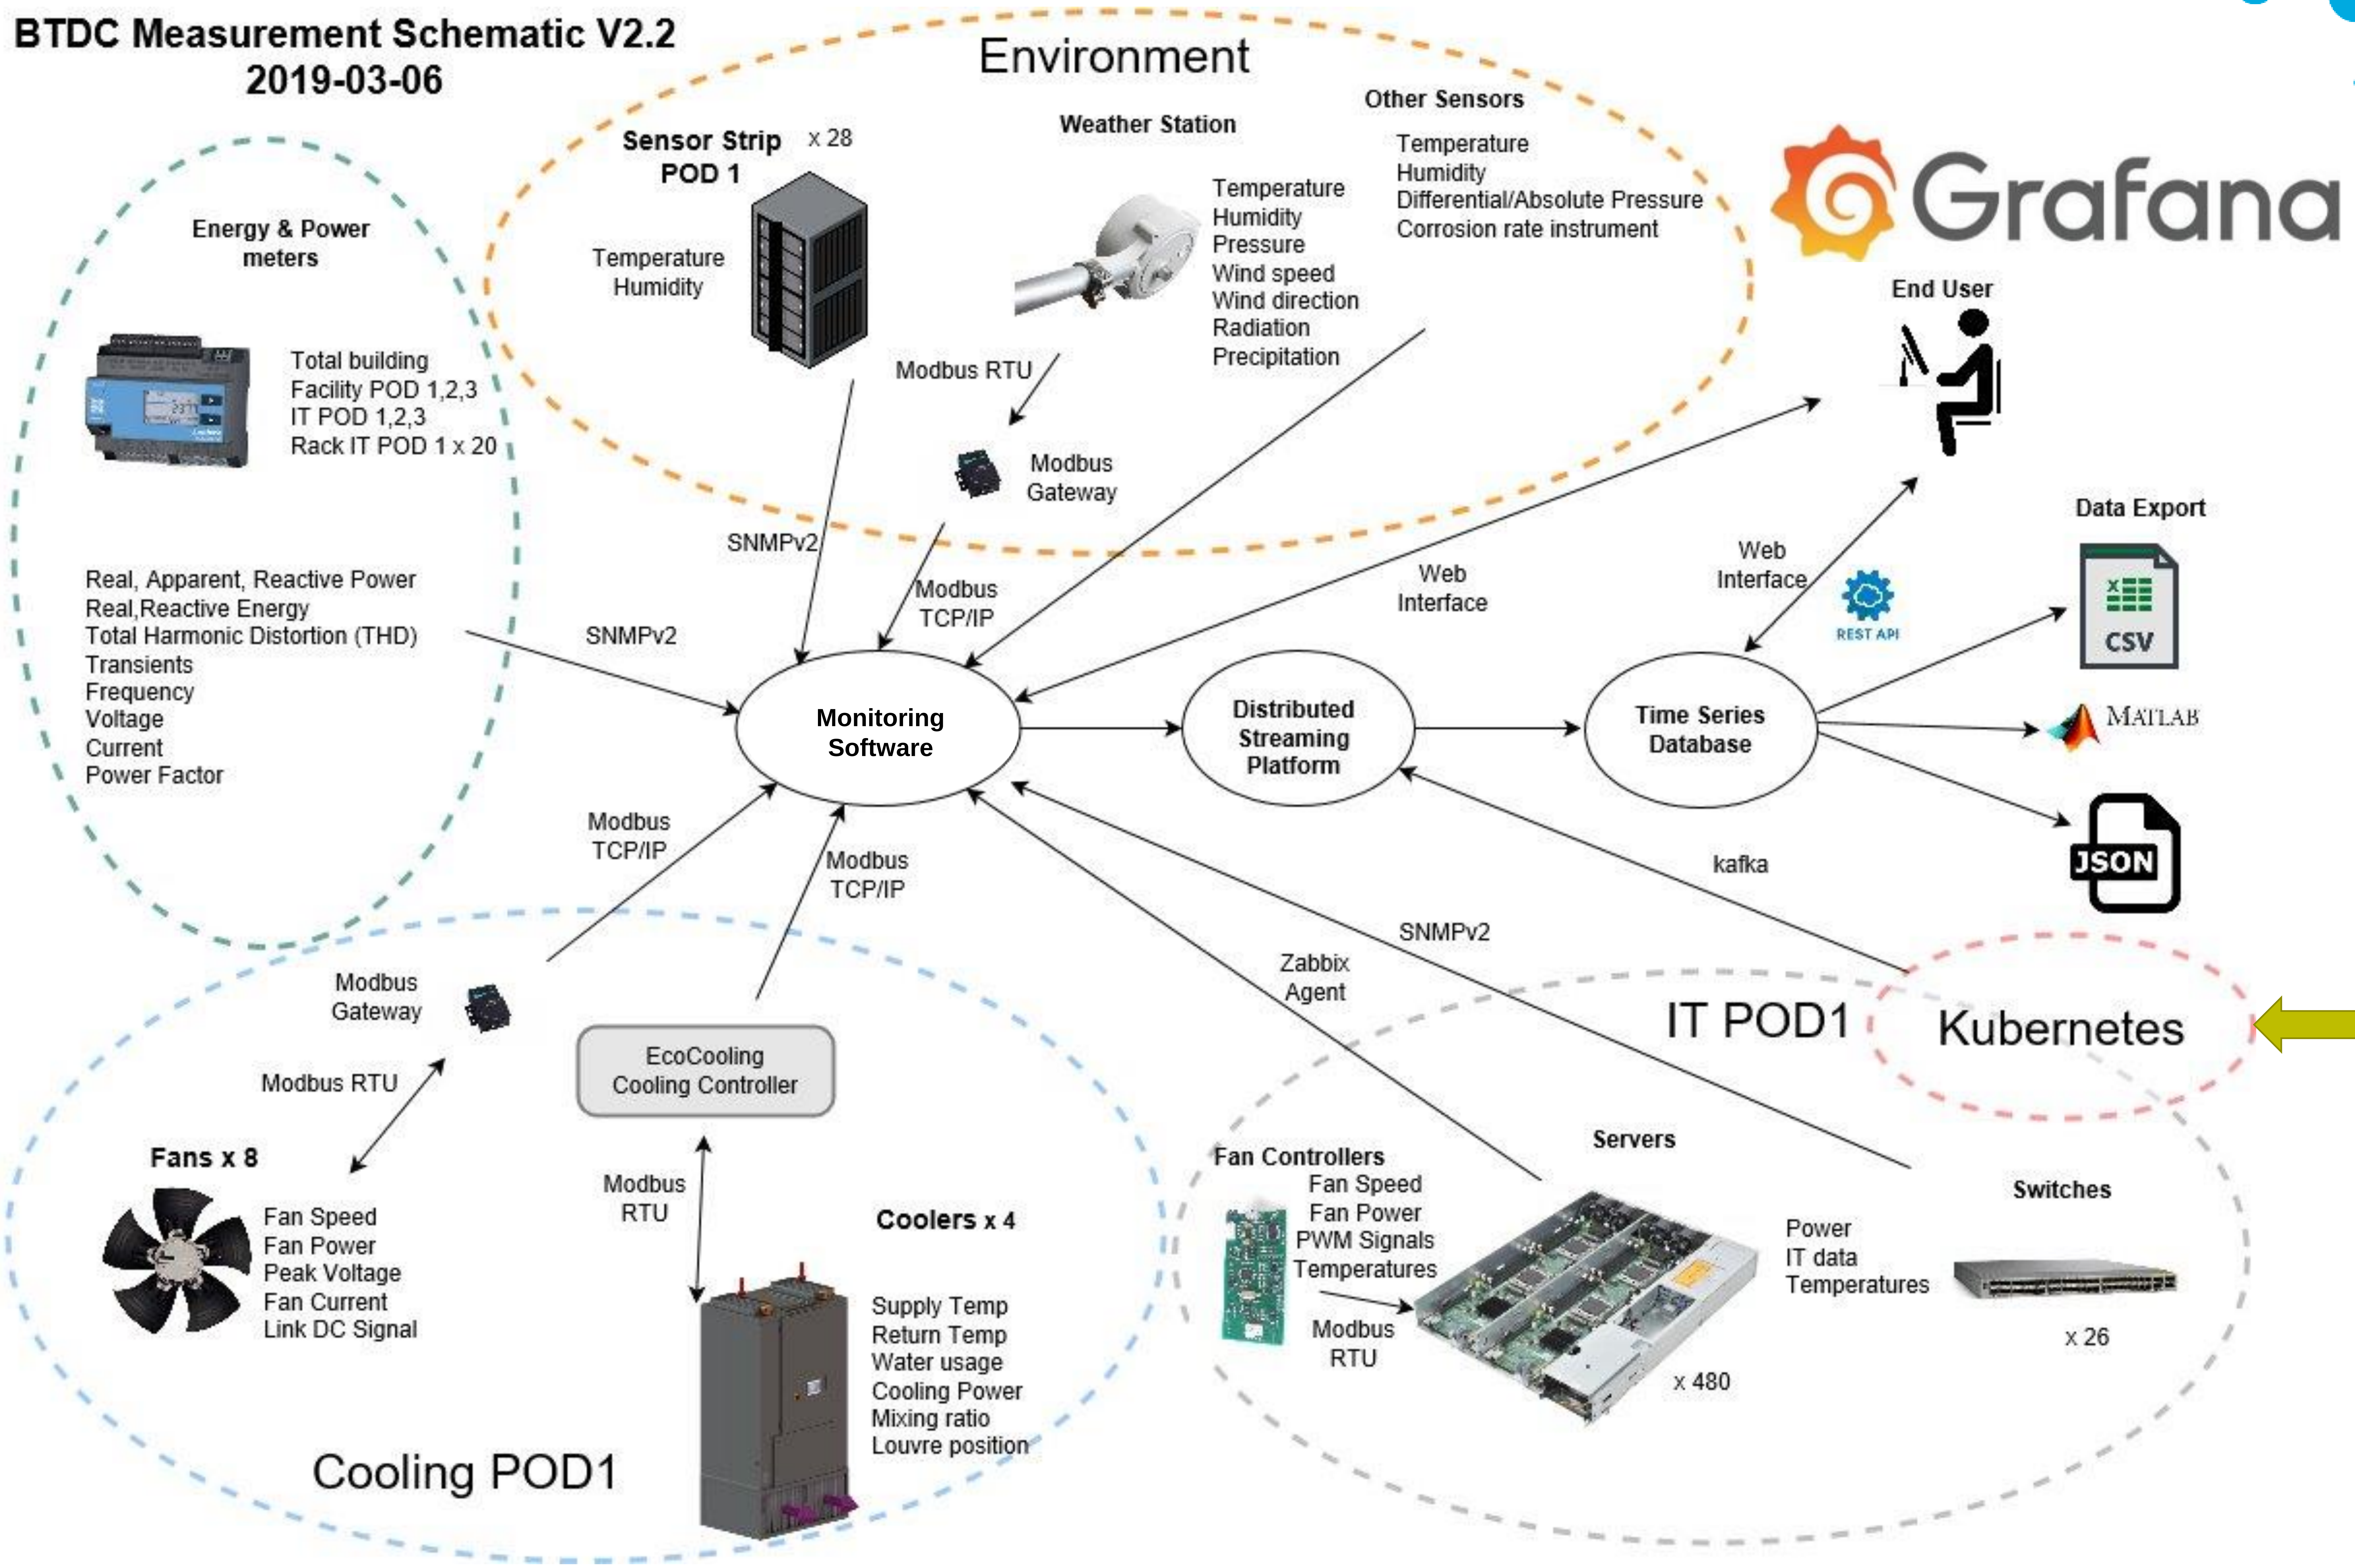

GA 768875

Server wind tunnel arrangement

GA 768875

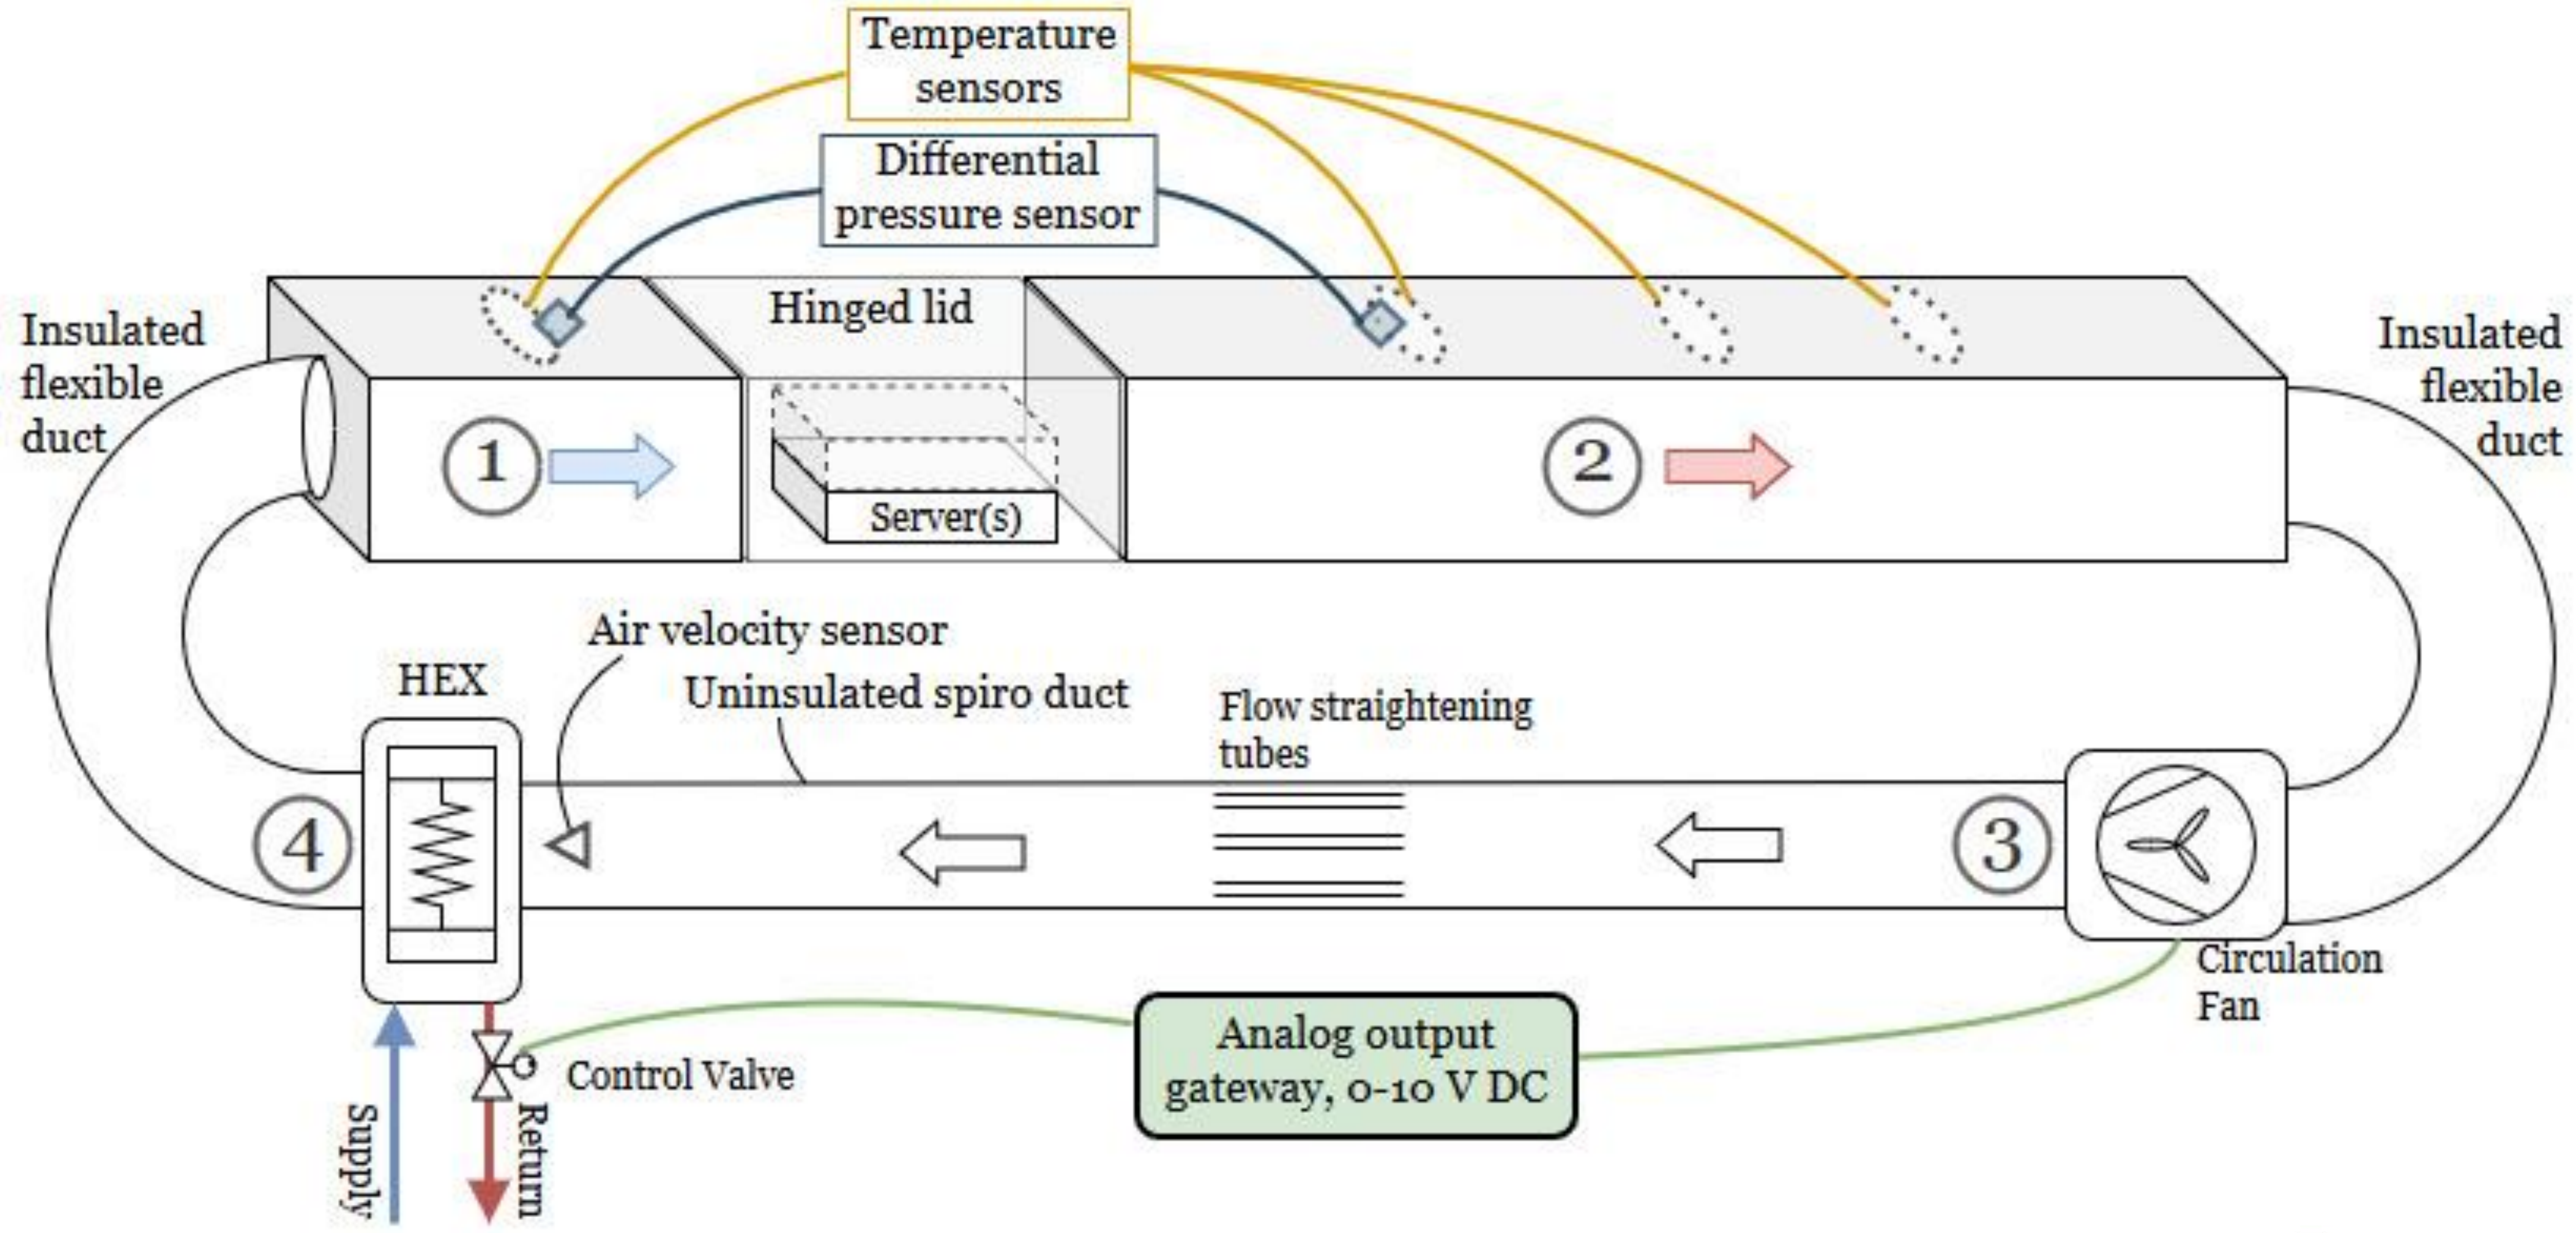

Server wind tunnel

- Set temperature with the HEX chilled water flowrate and wind tunnel FAN speed.
- Air flowrate with the wind tunnel FAN – can pressurize the cold aisle as would happen in the contained cold aisle in the DC.
- Upstream and downstream temperature measurements.
- Can support between 1U and 4U of Information Technology.
- Currently limited control over the air relative humidity.

Results with the OCP Windmill Server

- Wind tunnel results are captured using the same data center monitoring tools.

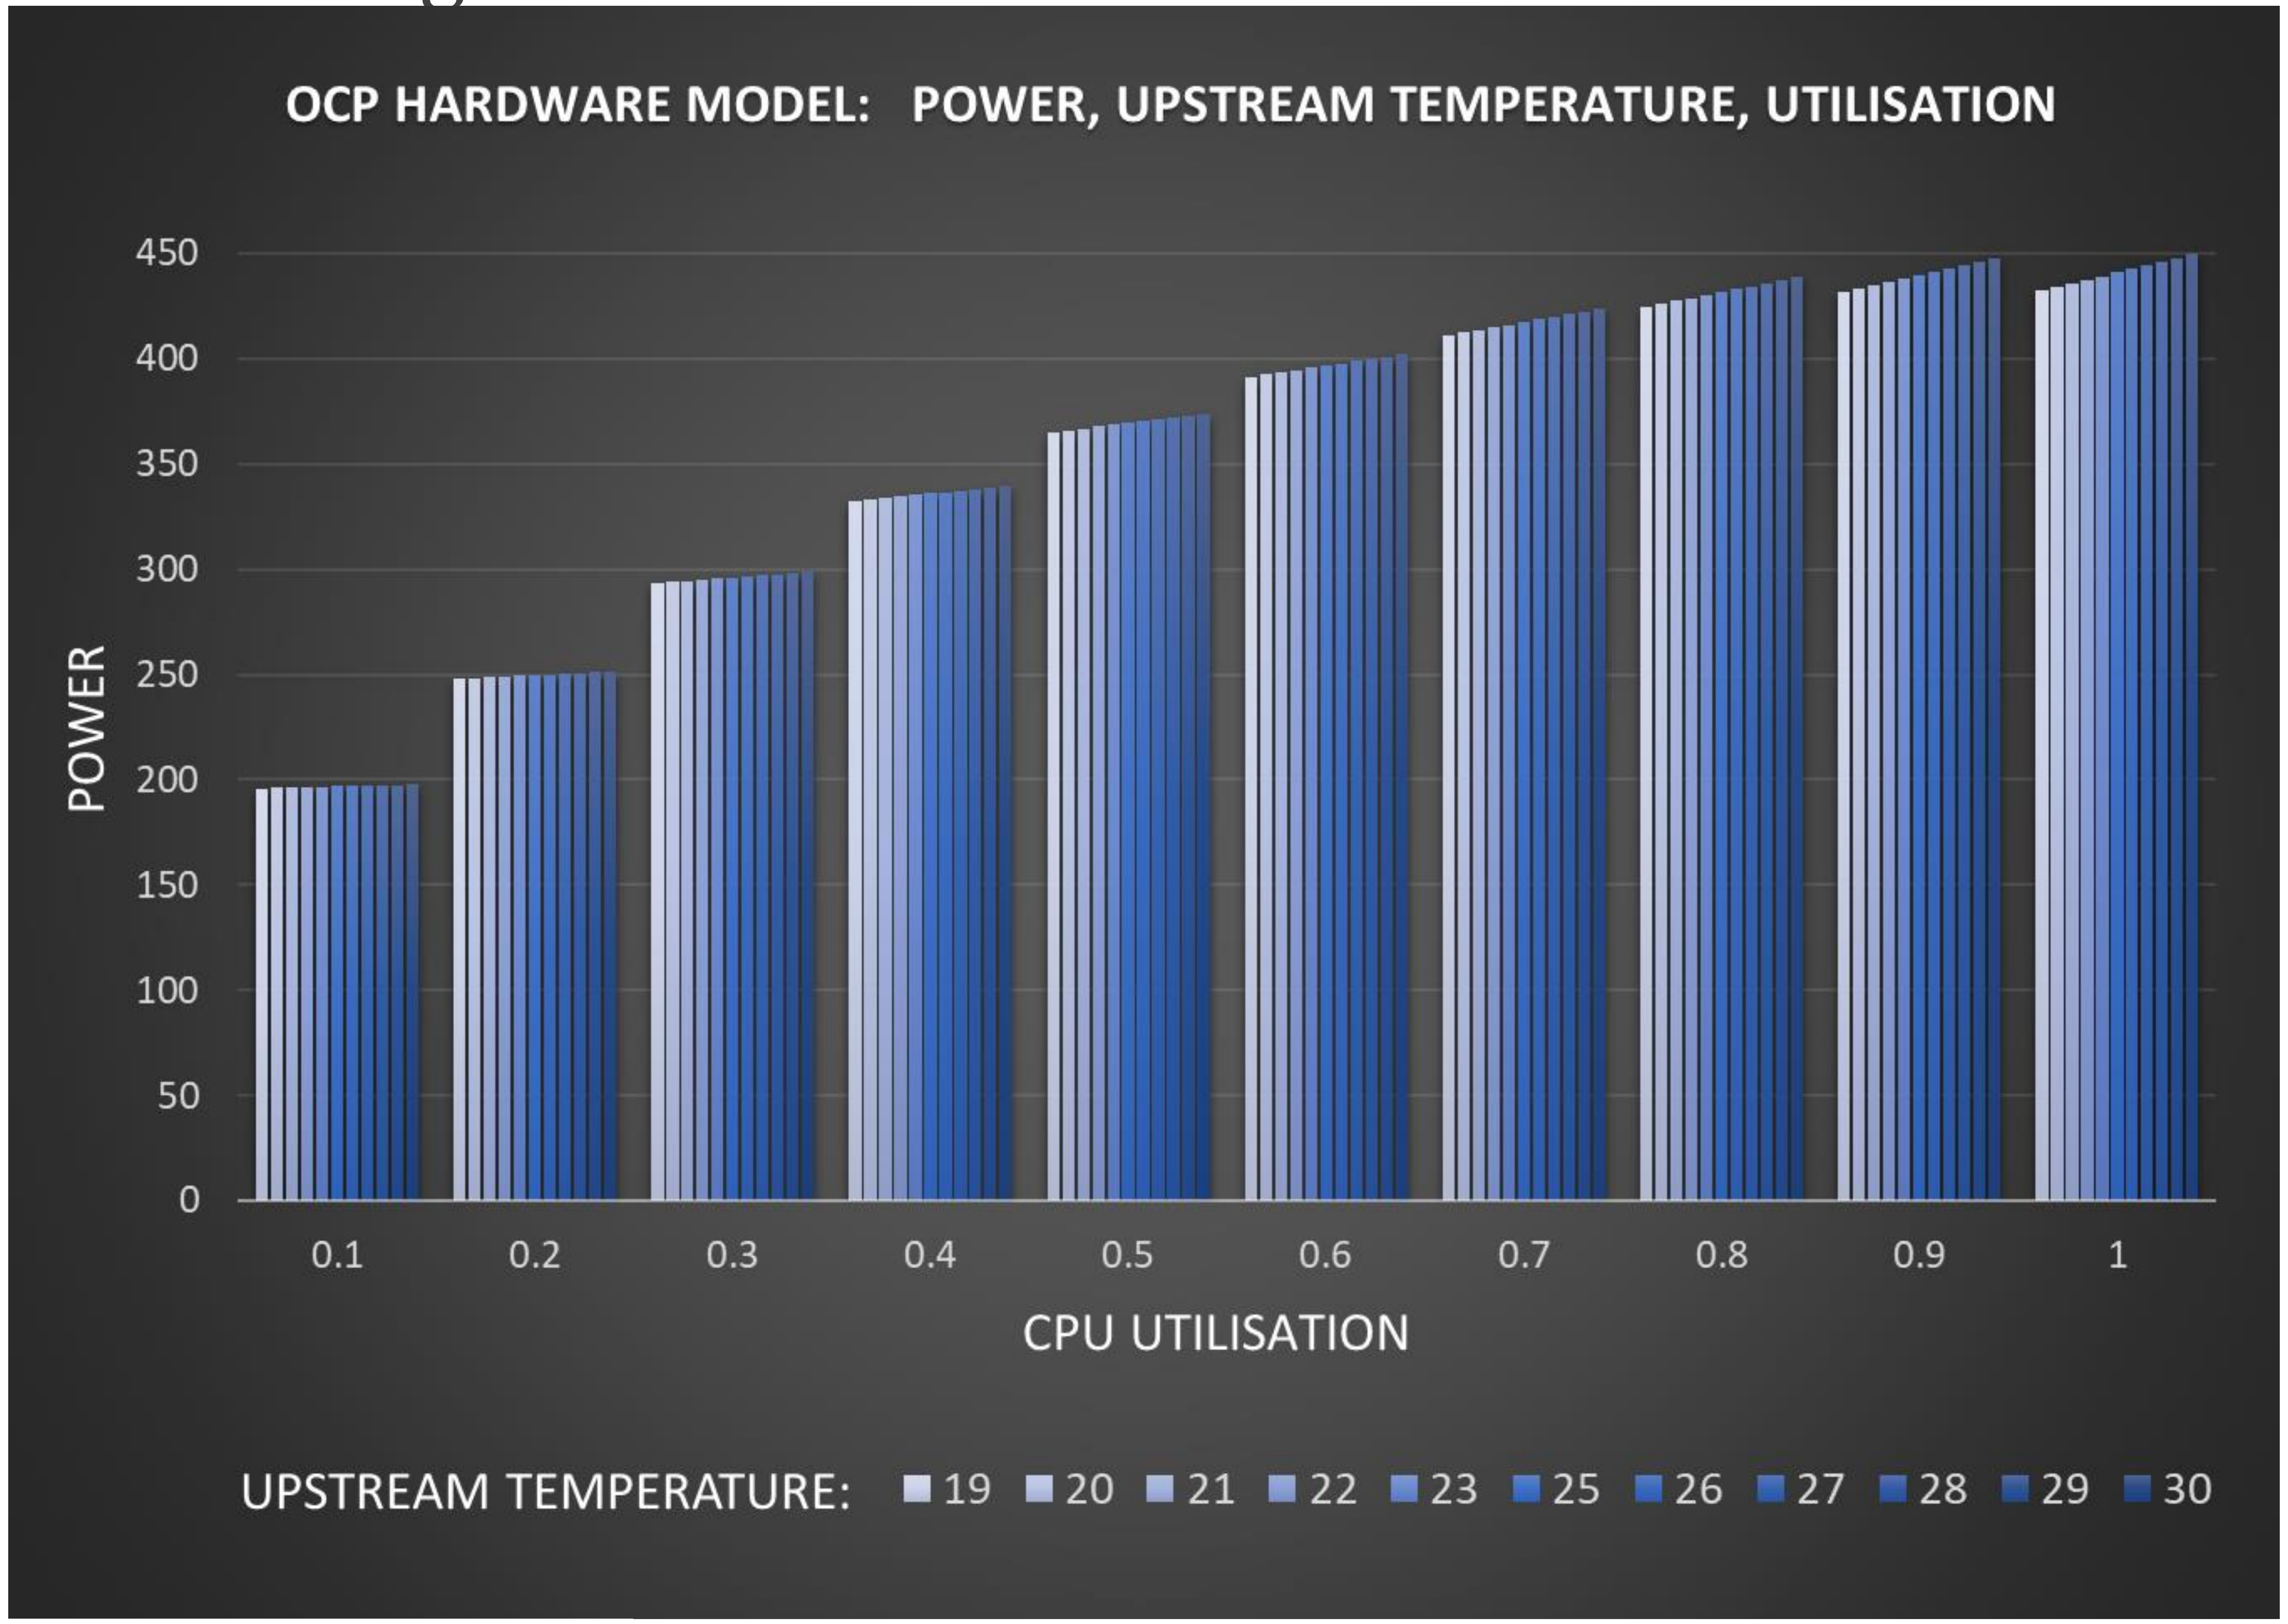

- Uses the same opensource monitoring system as the data center.
- The project uses many different workload generation approaches.
- Note the variation of power consumption with server supply temperature – combination of increased fan power and current leakages.
- Temperatures are in ºC
- Pressure drop was not kept constant for the different inlet temperatures.

Edgetic

RISE

Relationship between the delta T and power consumption of the IT.

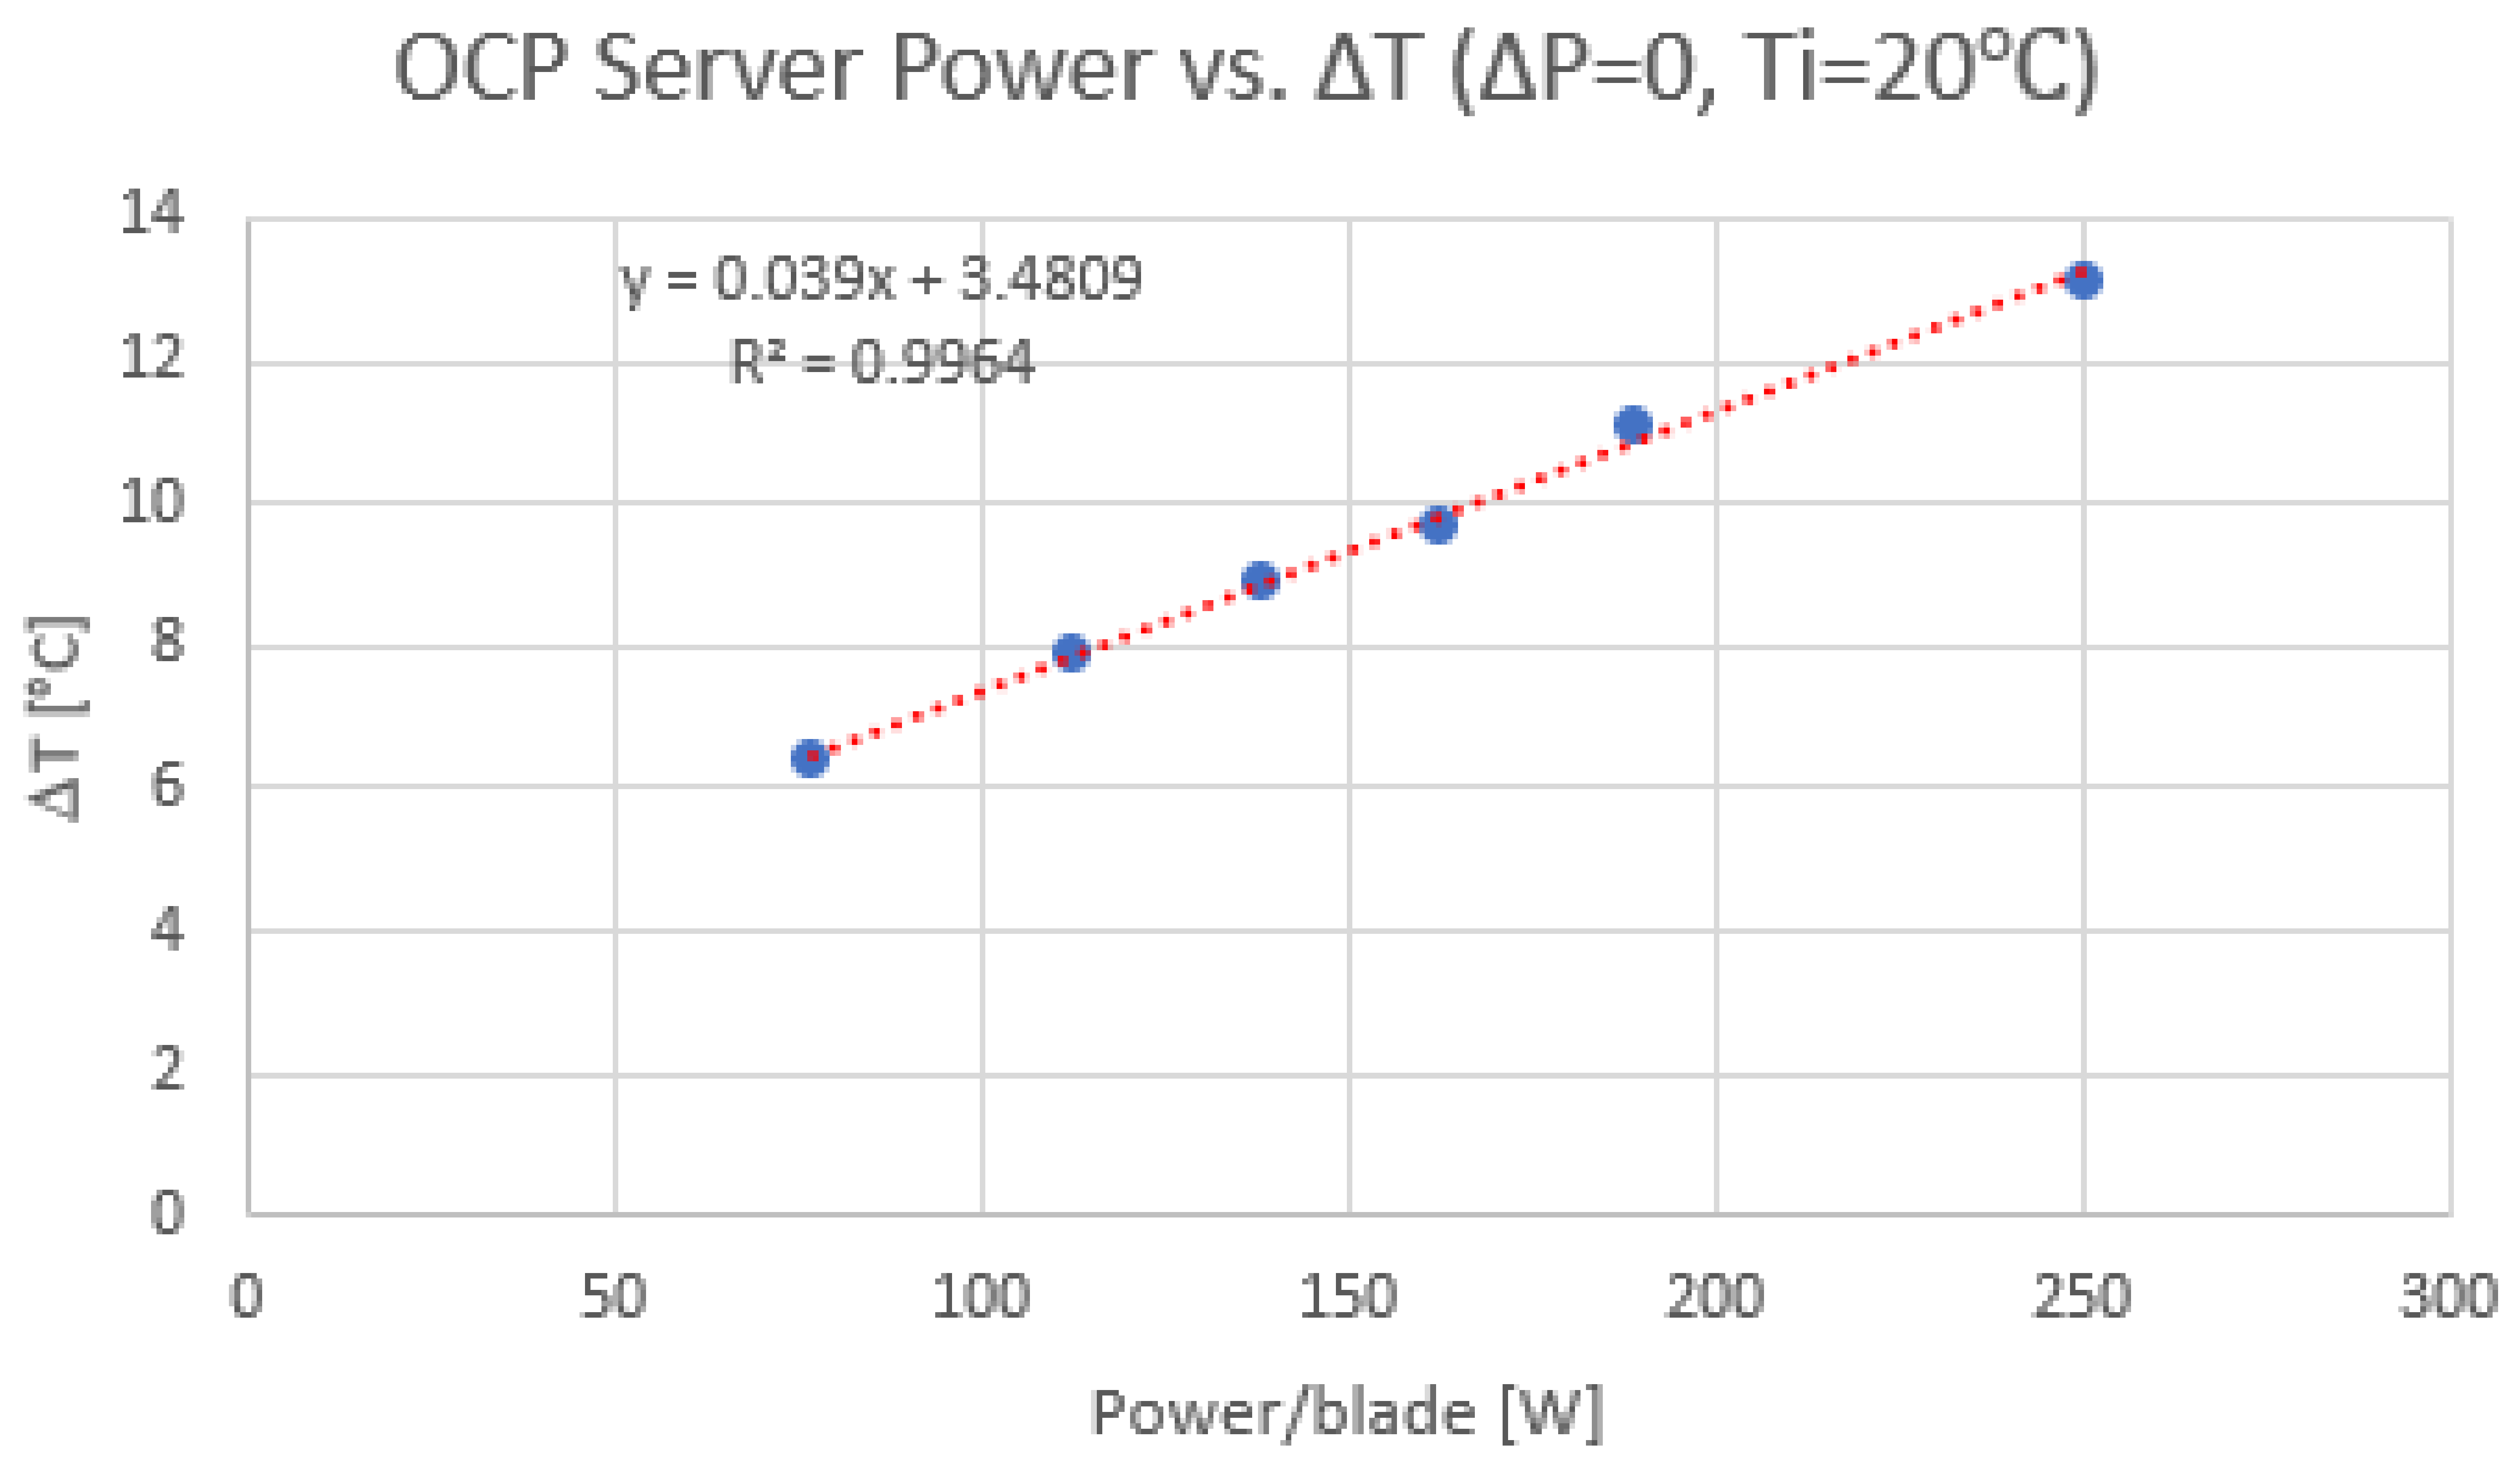

x 480 servers

- The cooling strategy is to set a delta T on the cooling systems to match idle power.
- When IT power increases the coolers will speed up the supply and extract fans to maintain the supply temperature and to try and maintain the return temperature based on the idle delta T.
- This will result in pressurisation of the contained cold aisle.
- A more efficient strategy is to allow the return temperature to match the exhaust temperature from the servers.
- Wind tunnel HEX kept inlet temperature constant and the FAN kept the server pressure drop to 0.

Innovation around thermal management

POD 1 – **Cooling Units set T_{supply} and T_{return}**

GA 768875

Traditional
Set supply
Temperature and
expected
delta T

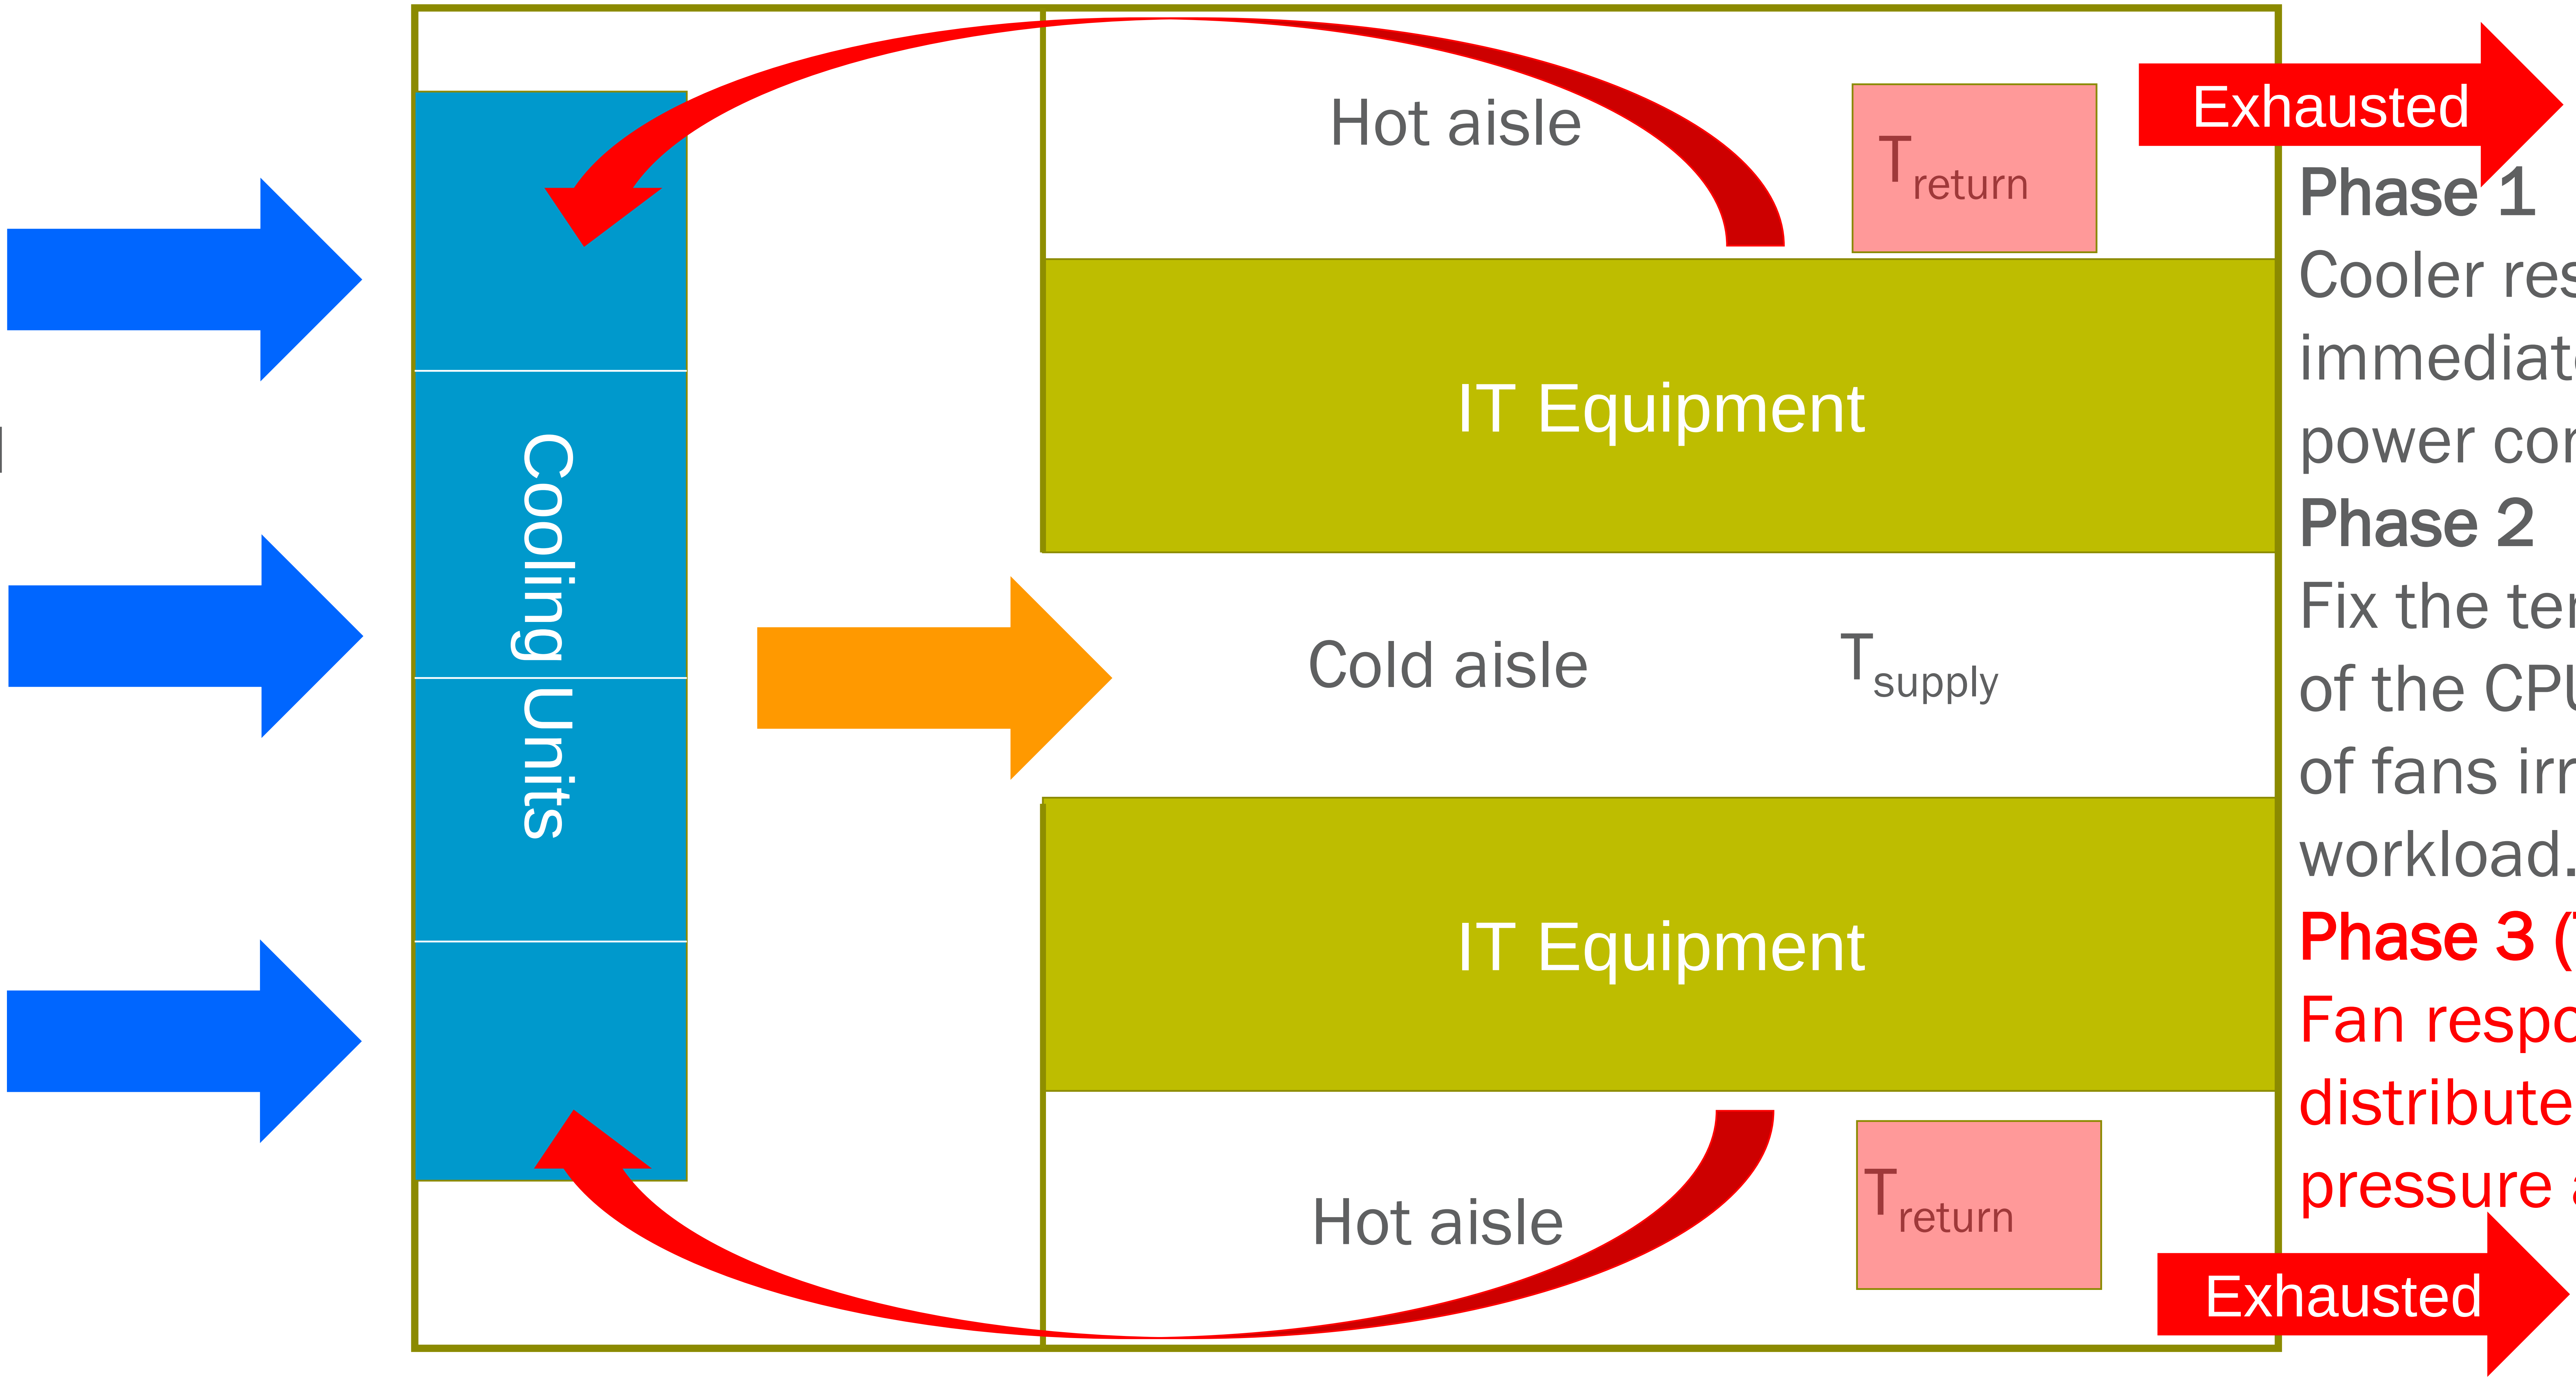

Phase 1
Cooler responds immediately to IT power consumption.

Phase 2
Fix the temperature of the CPUs via control of fans irrespective of workload.

Phase 3 (TODO)
Fan response to local distributed workload, pressure and more.

Wind tunnel to determine ΔT, fixed CPU temp, regulate air consumption on pressure.

Innovation around thermal management

POD 1 – **Cooling Units set T_{supply} and T_{return}**

GA 768875

Tr
Se
Te
ex
de

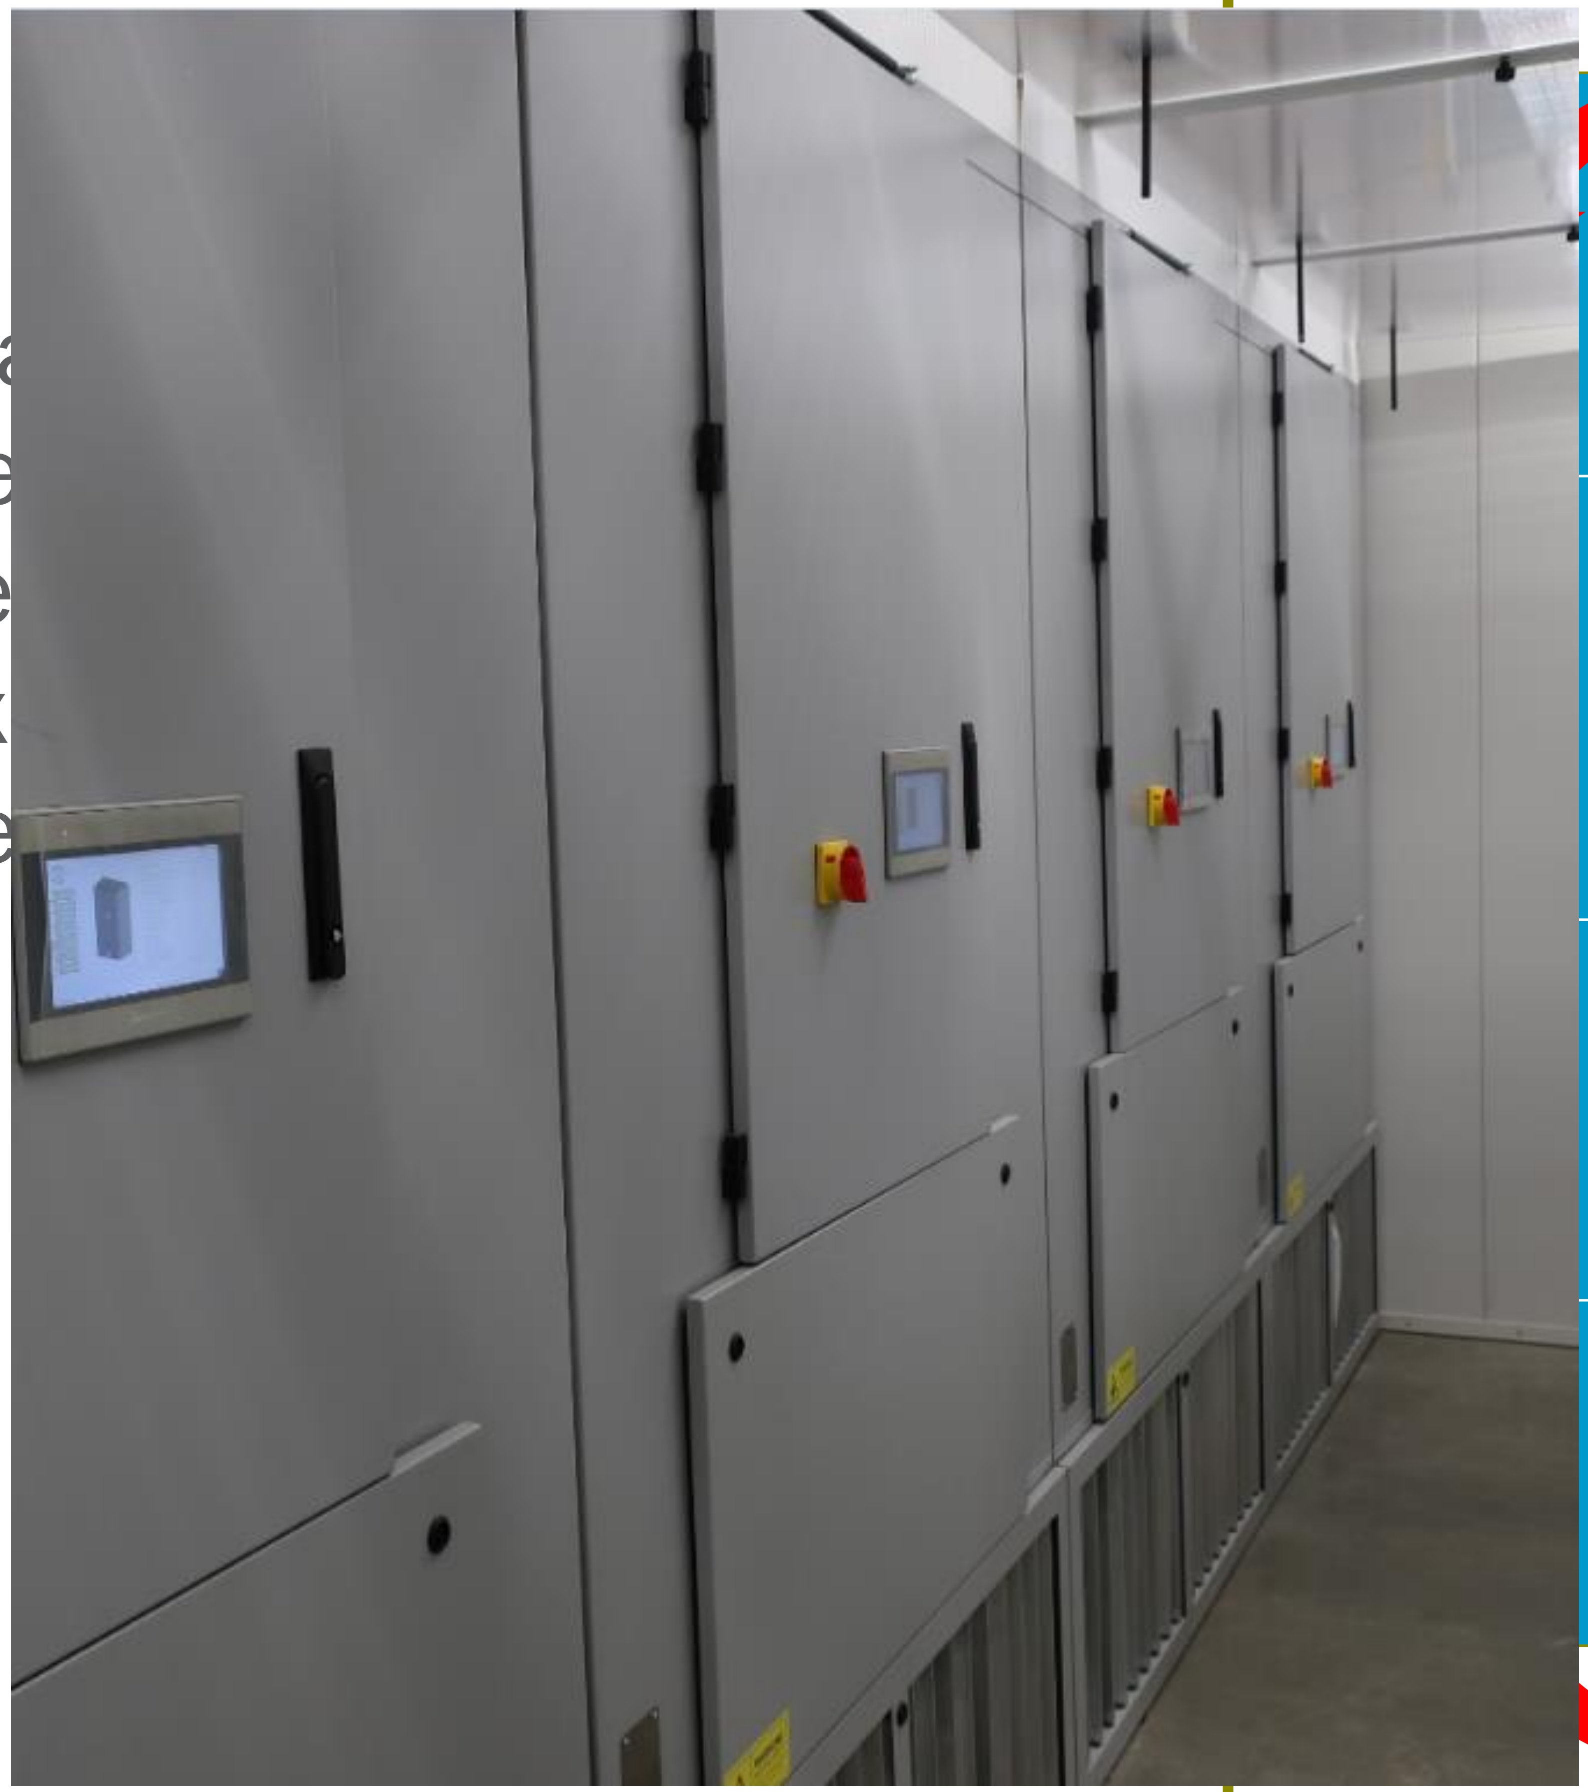

hausted

Phase 1

Cooler responds immediately to IT power consumption.

Phase 2

Fix the temperature of the CPUs via control of fans irrespective of workload.

Phase 3 (TODO)

Fan response to local distributed workload, pressure and more.

Exhausted

Wind tunnel to determine ΔT, fixed CPU temp, regulate air consumption on pressure.

Cooling system strategy

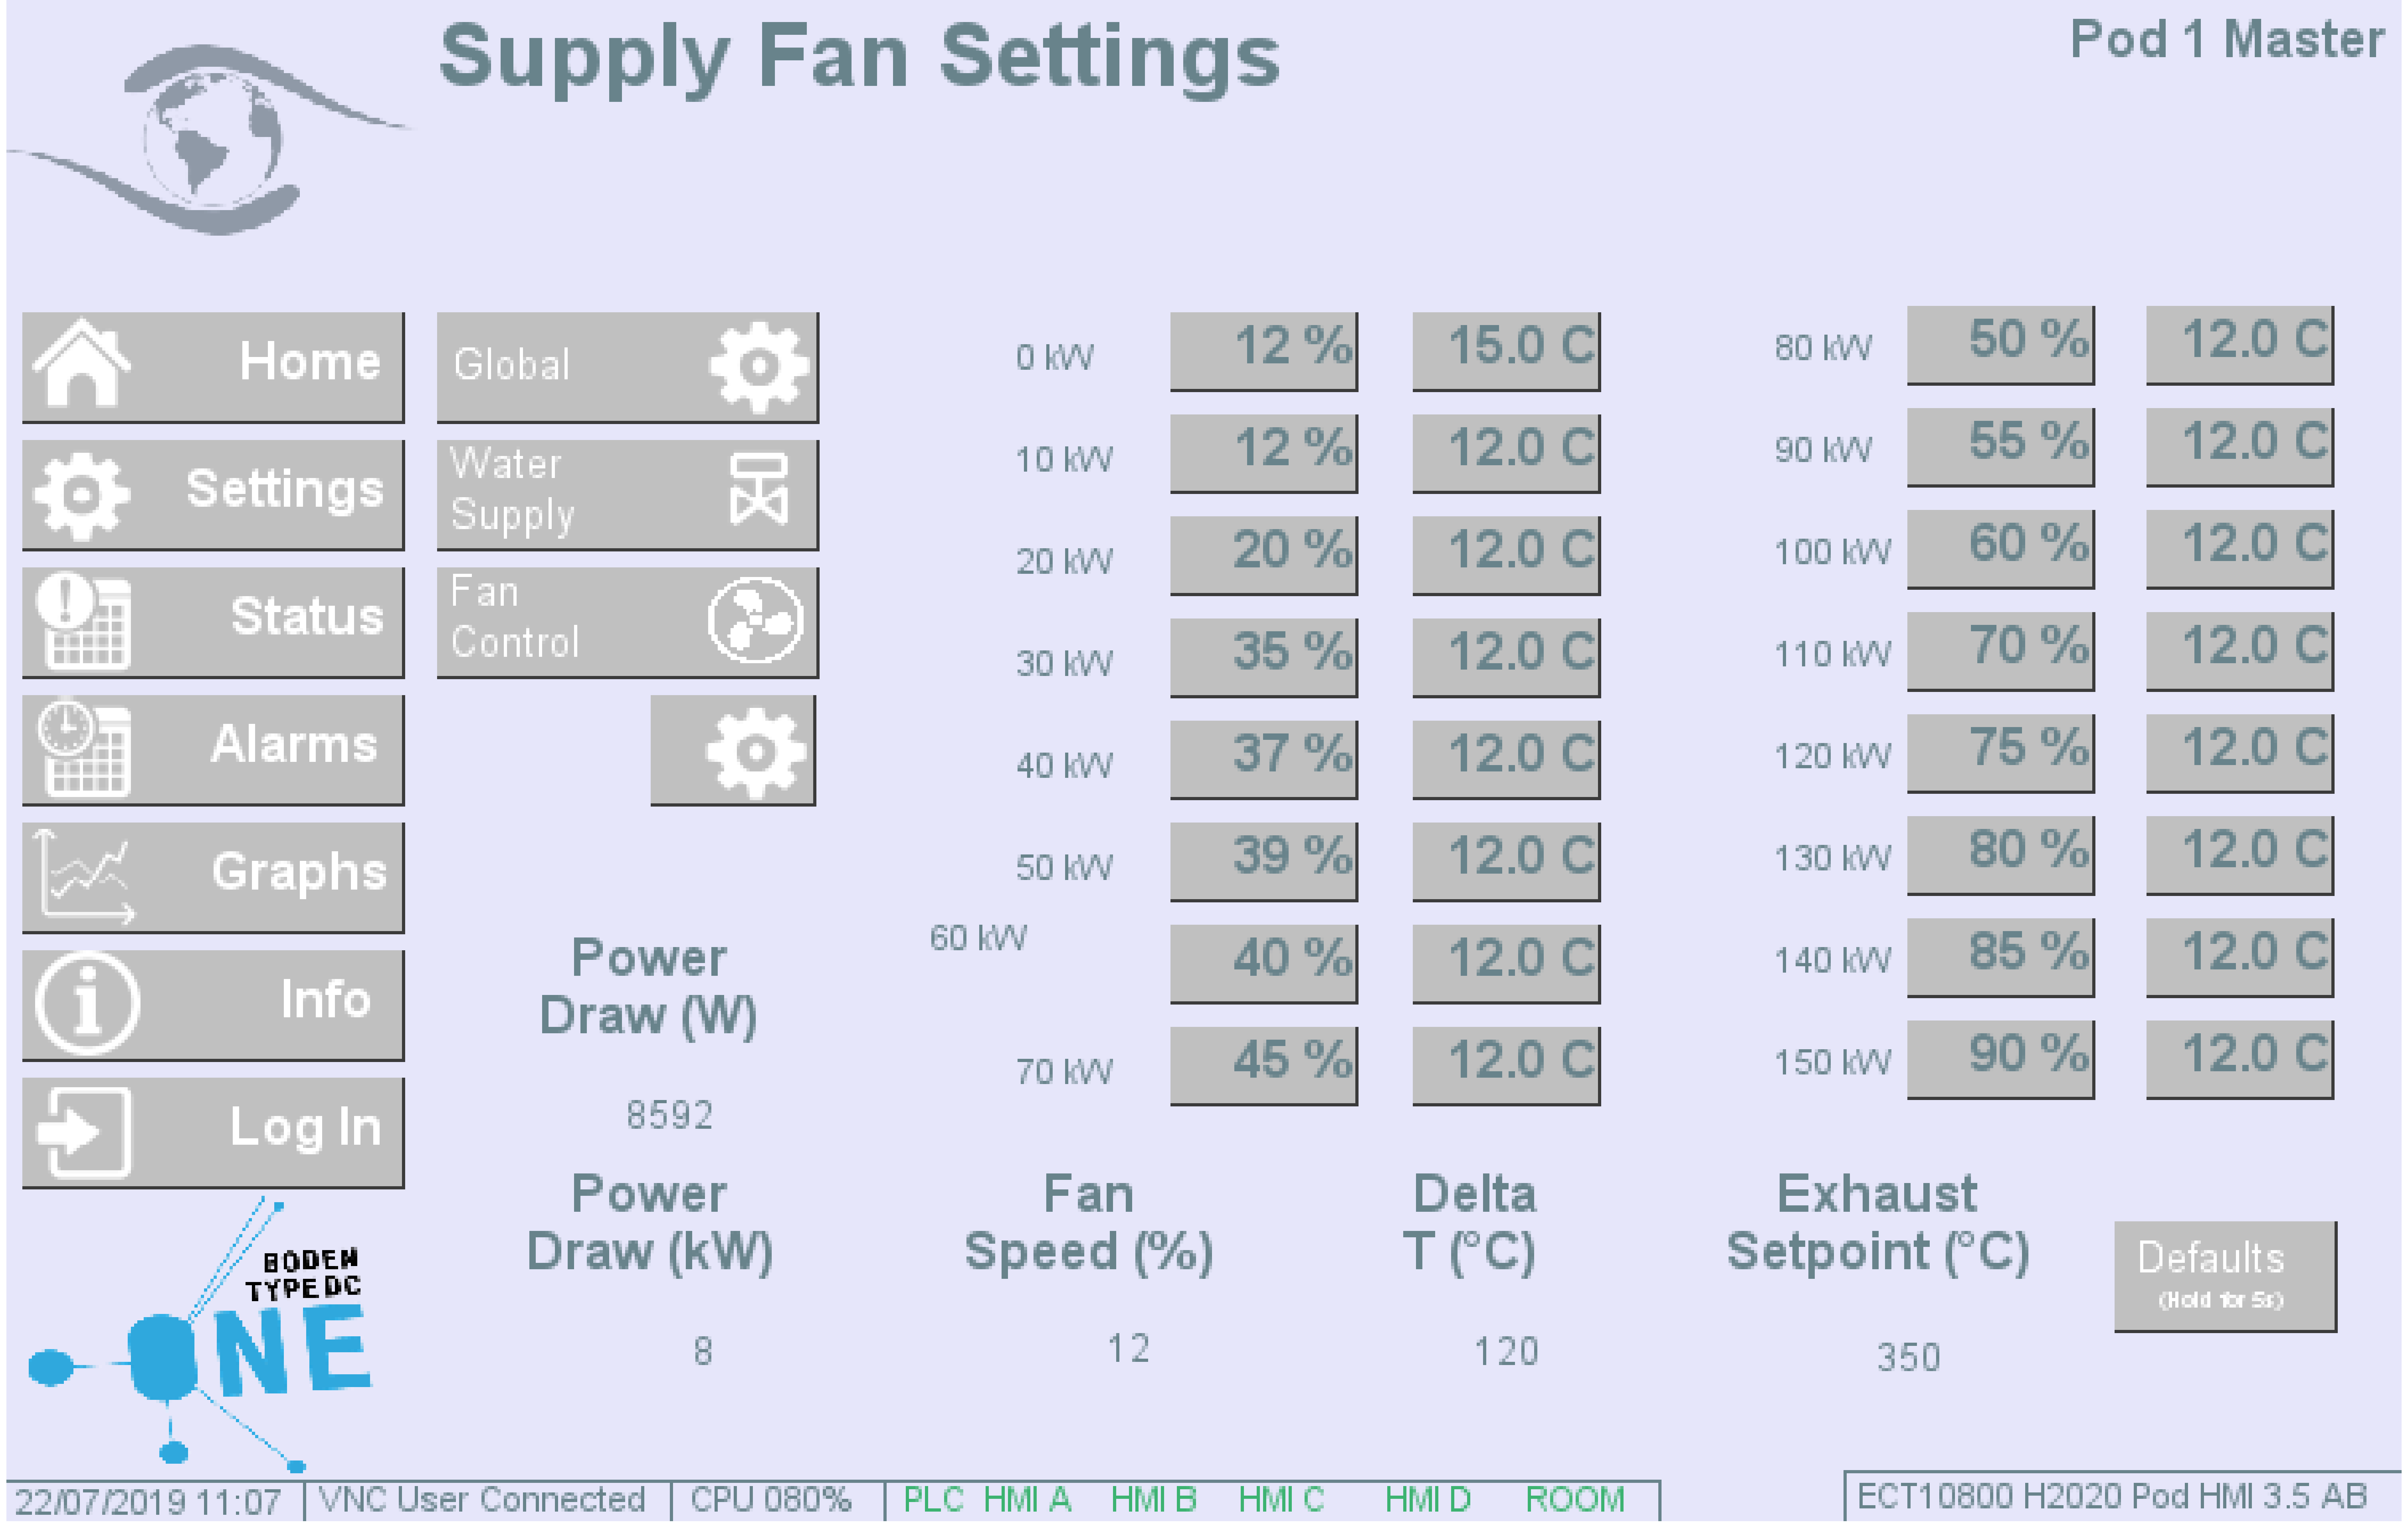

GA 768875

Cooling system strategy

GA 768875

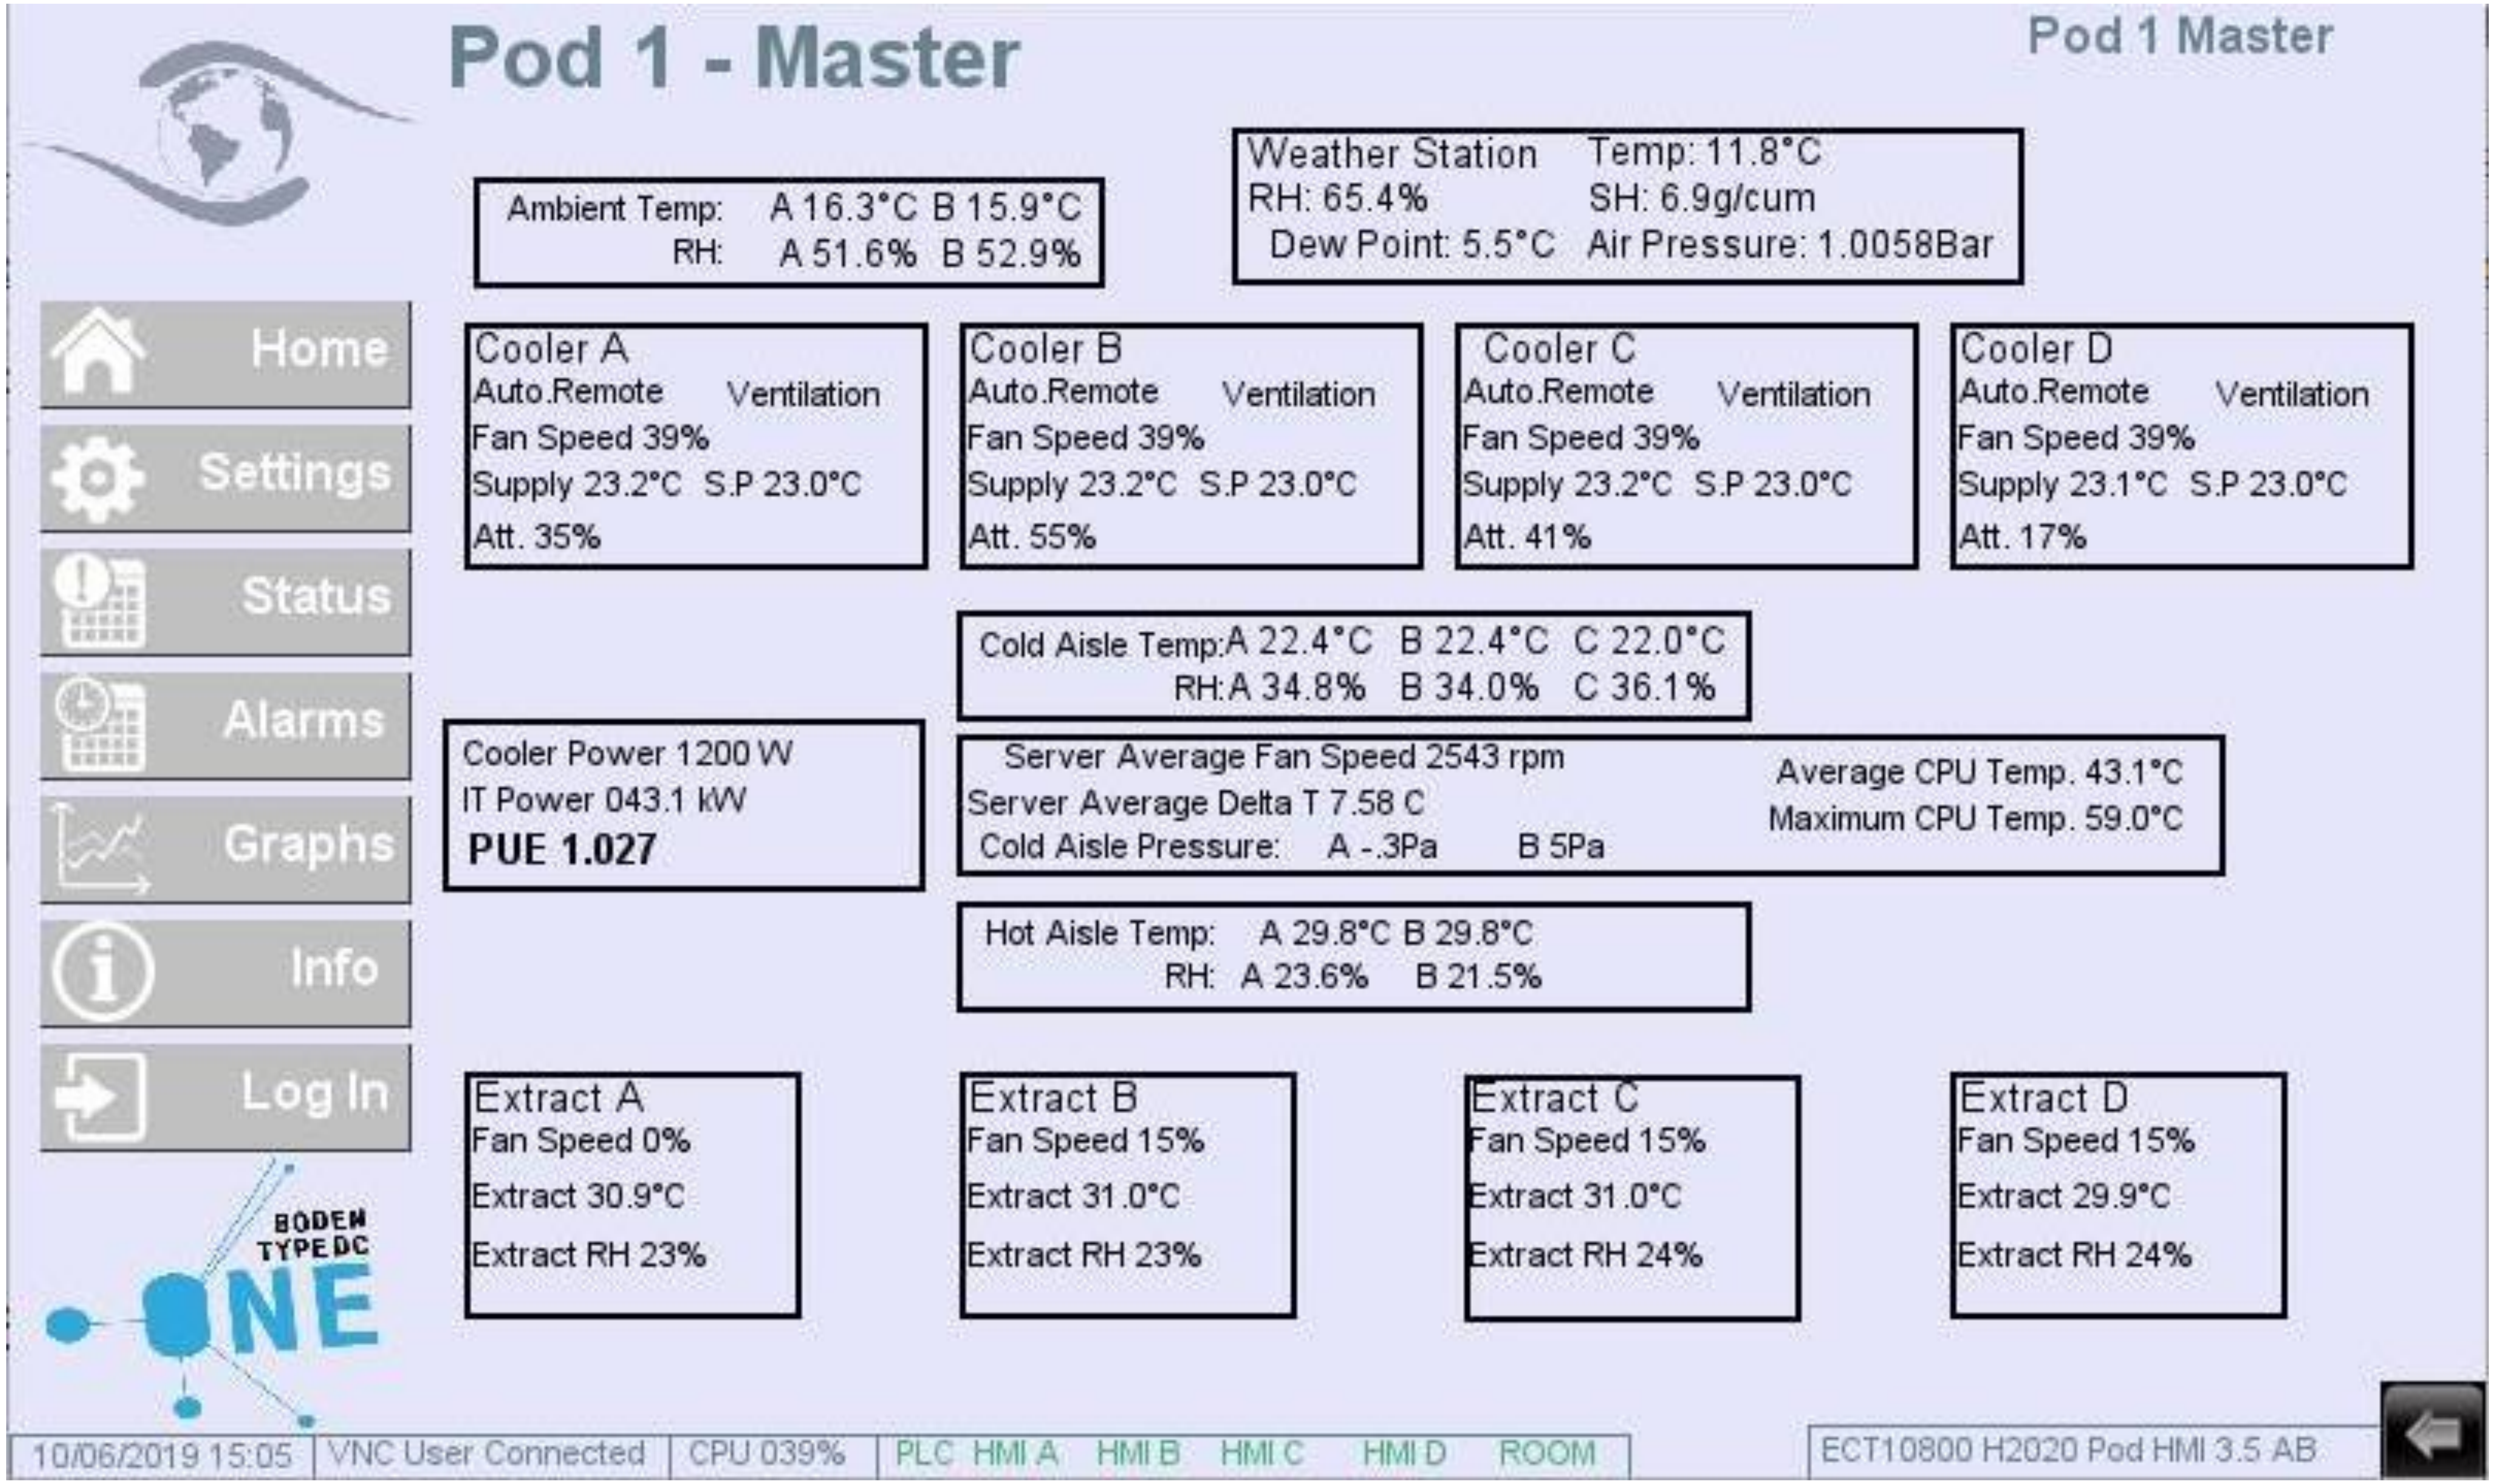

Workload deployment and orchestration

GA 768875

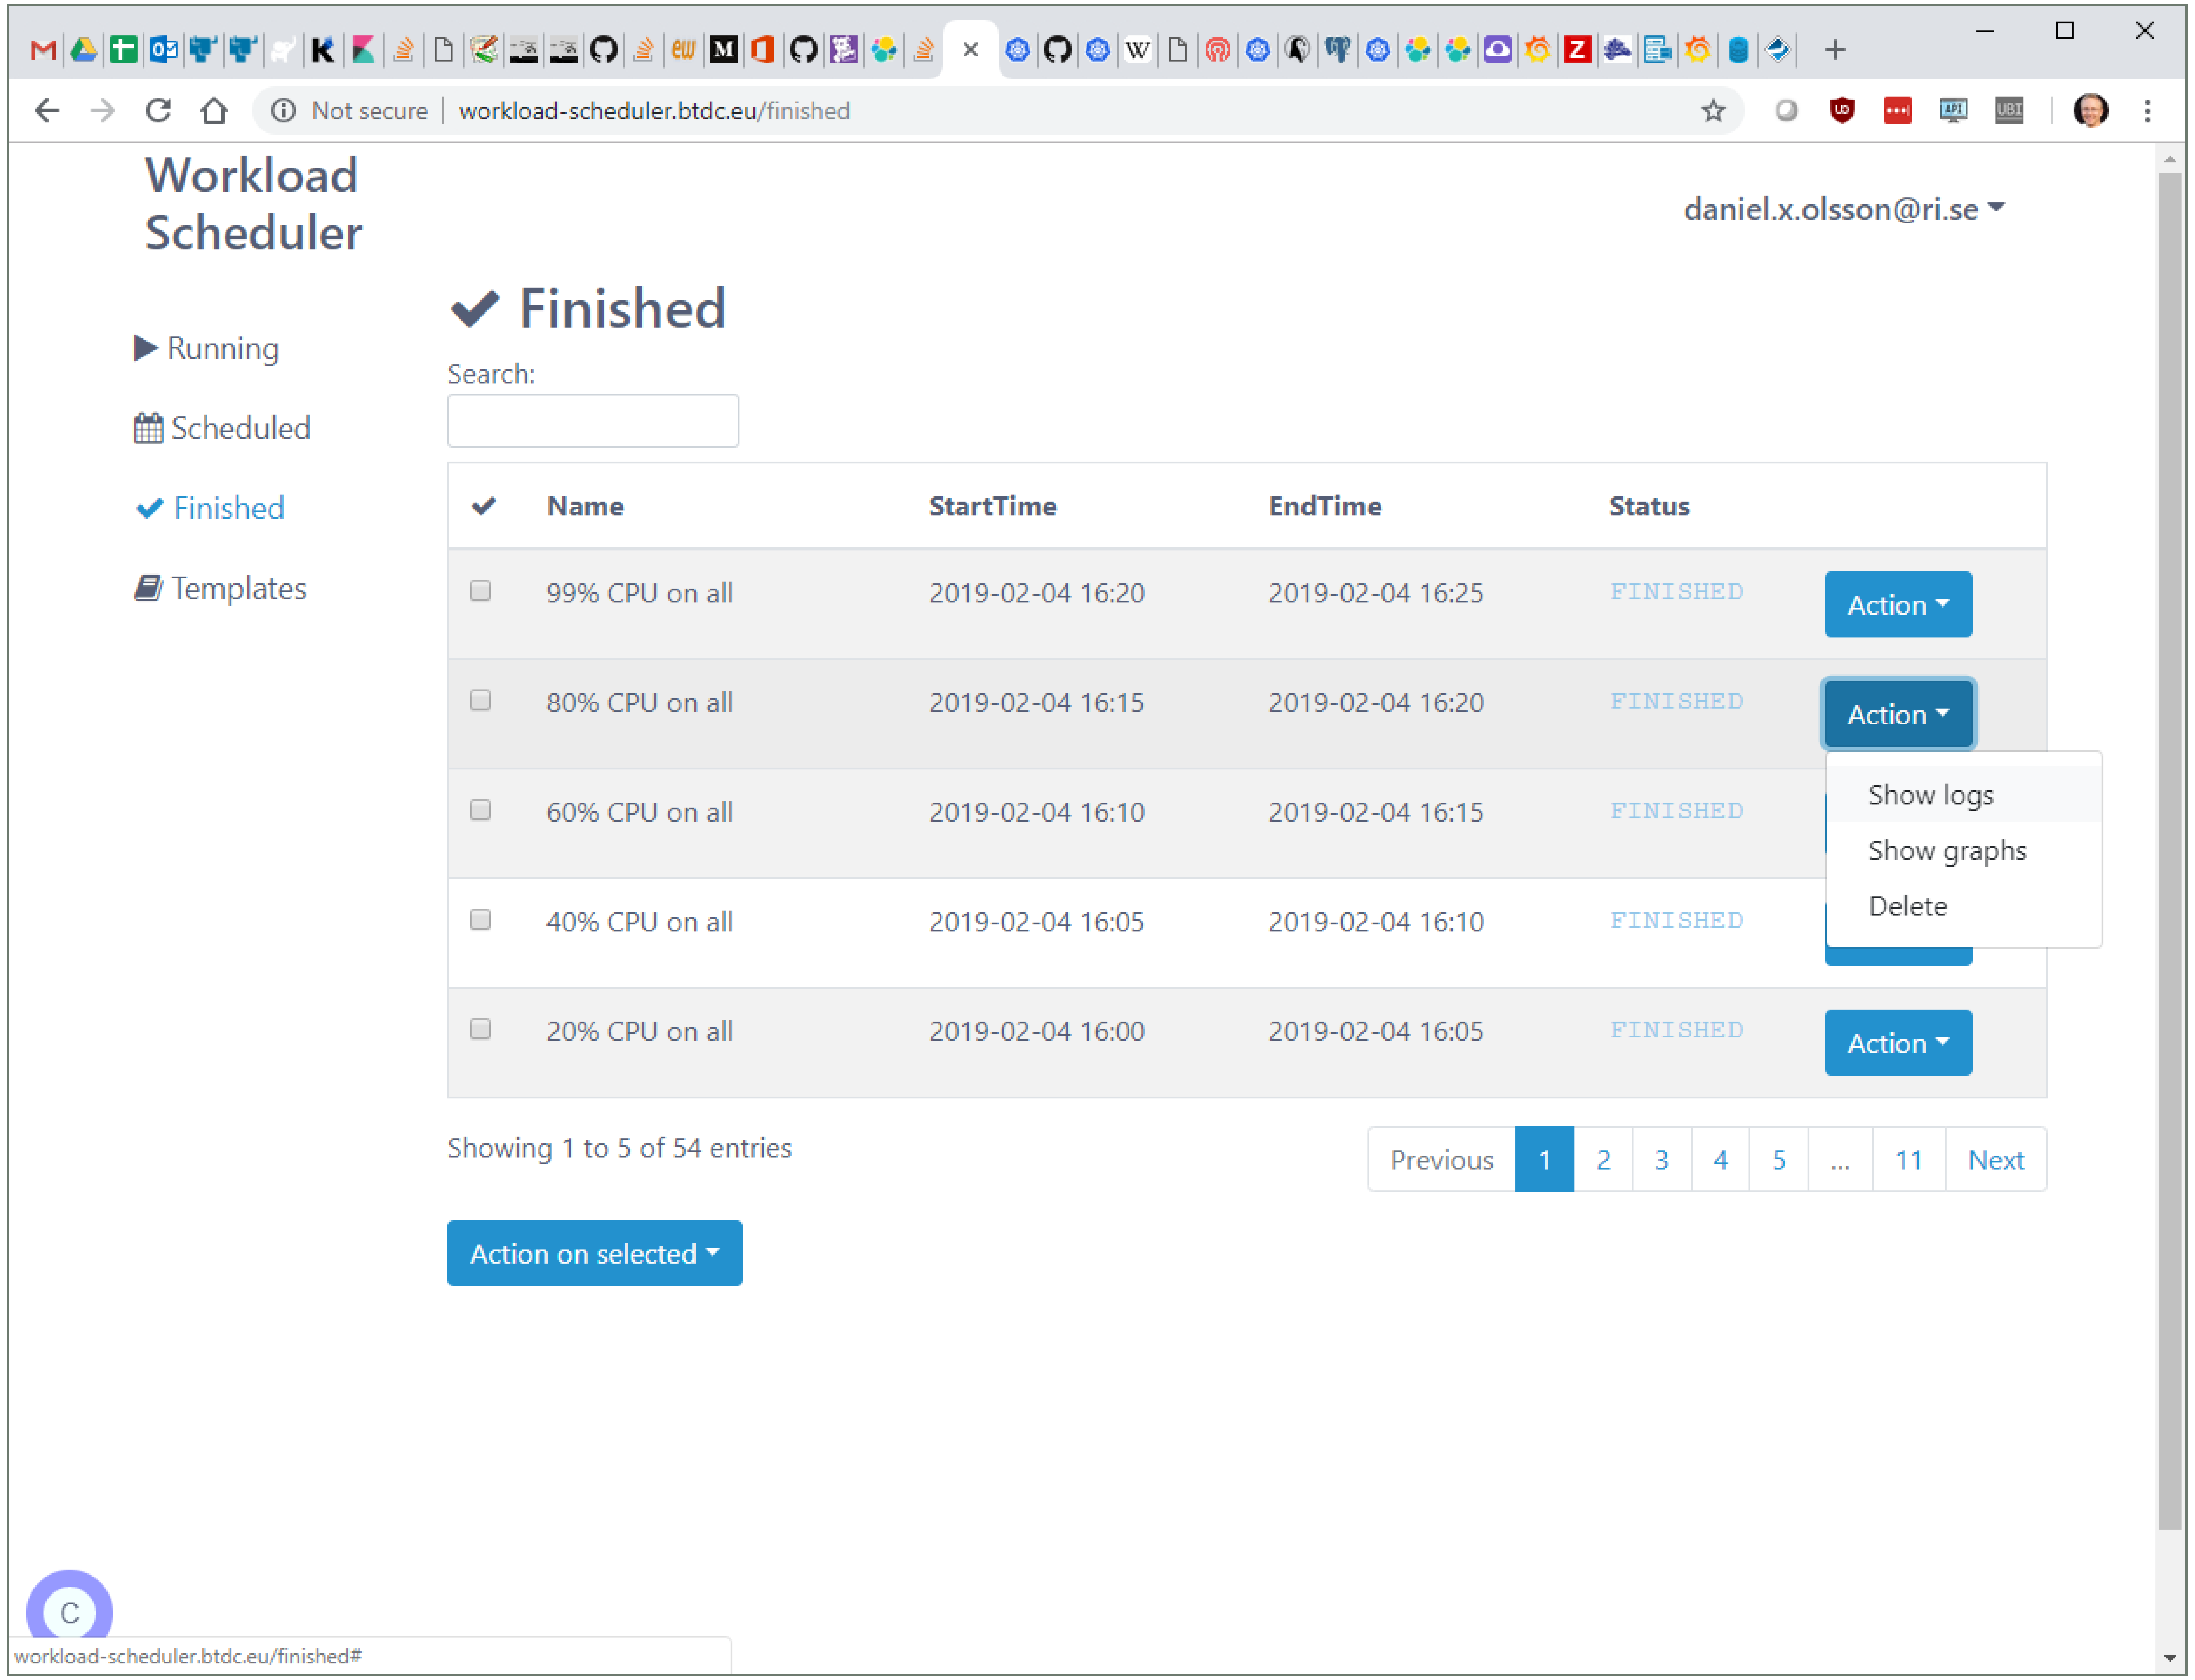

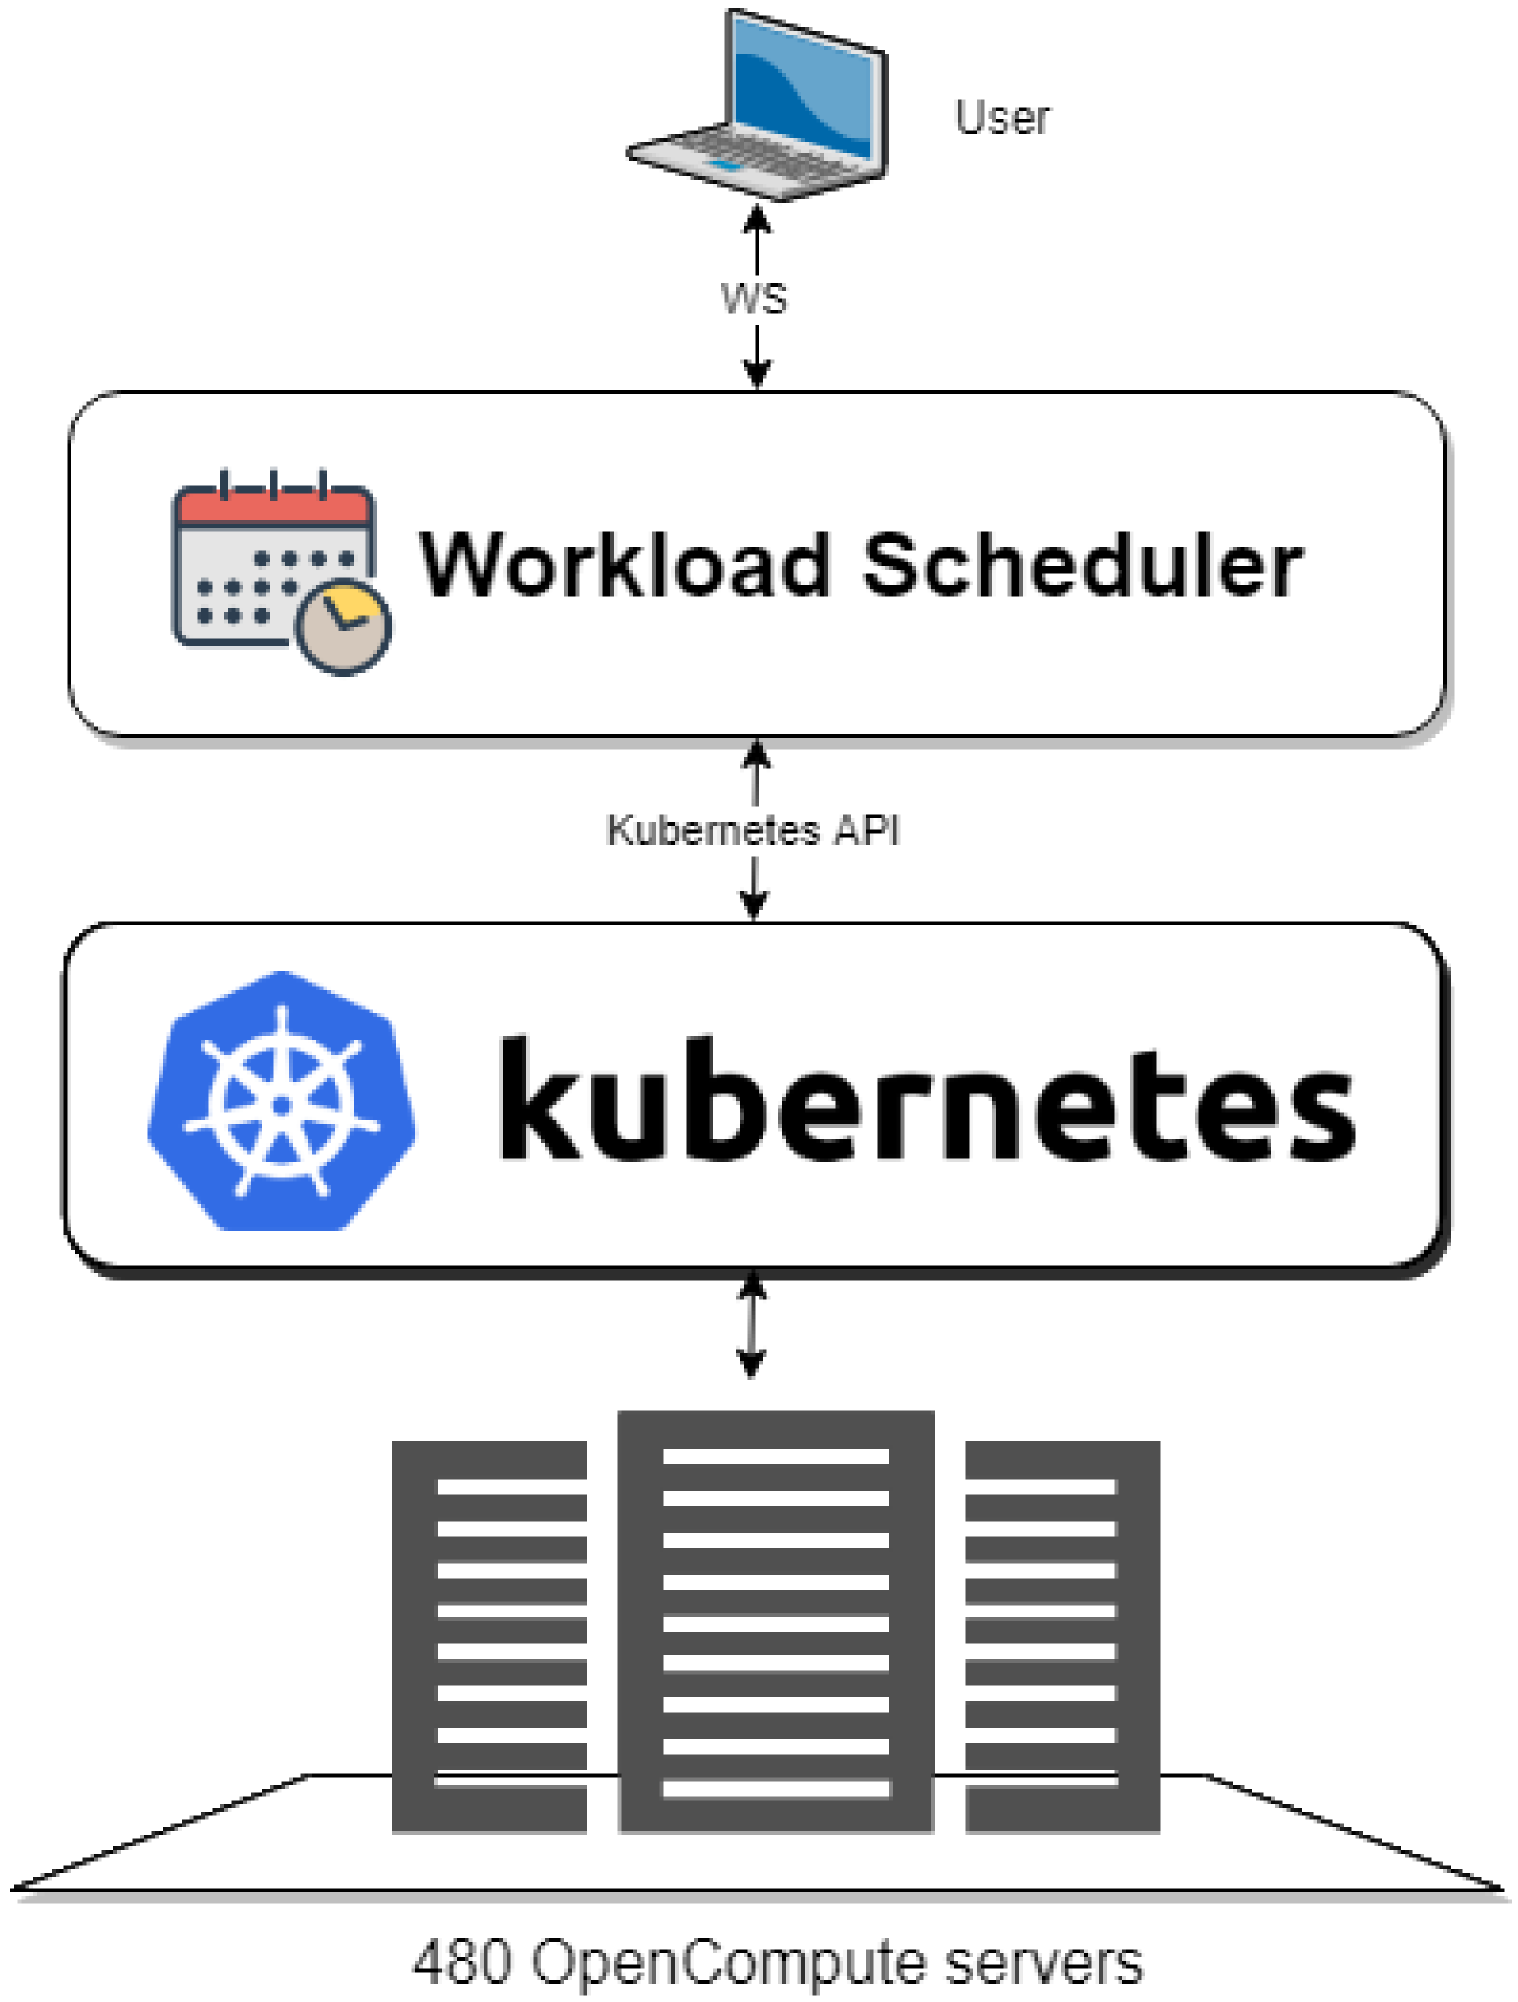

Varying CPU utilisation on servers using synthetic and application specific workloads.

GA 768875

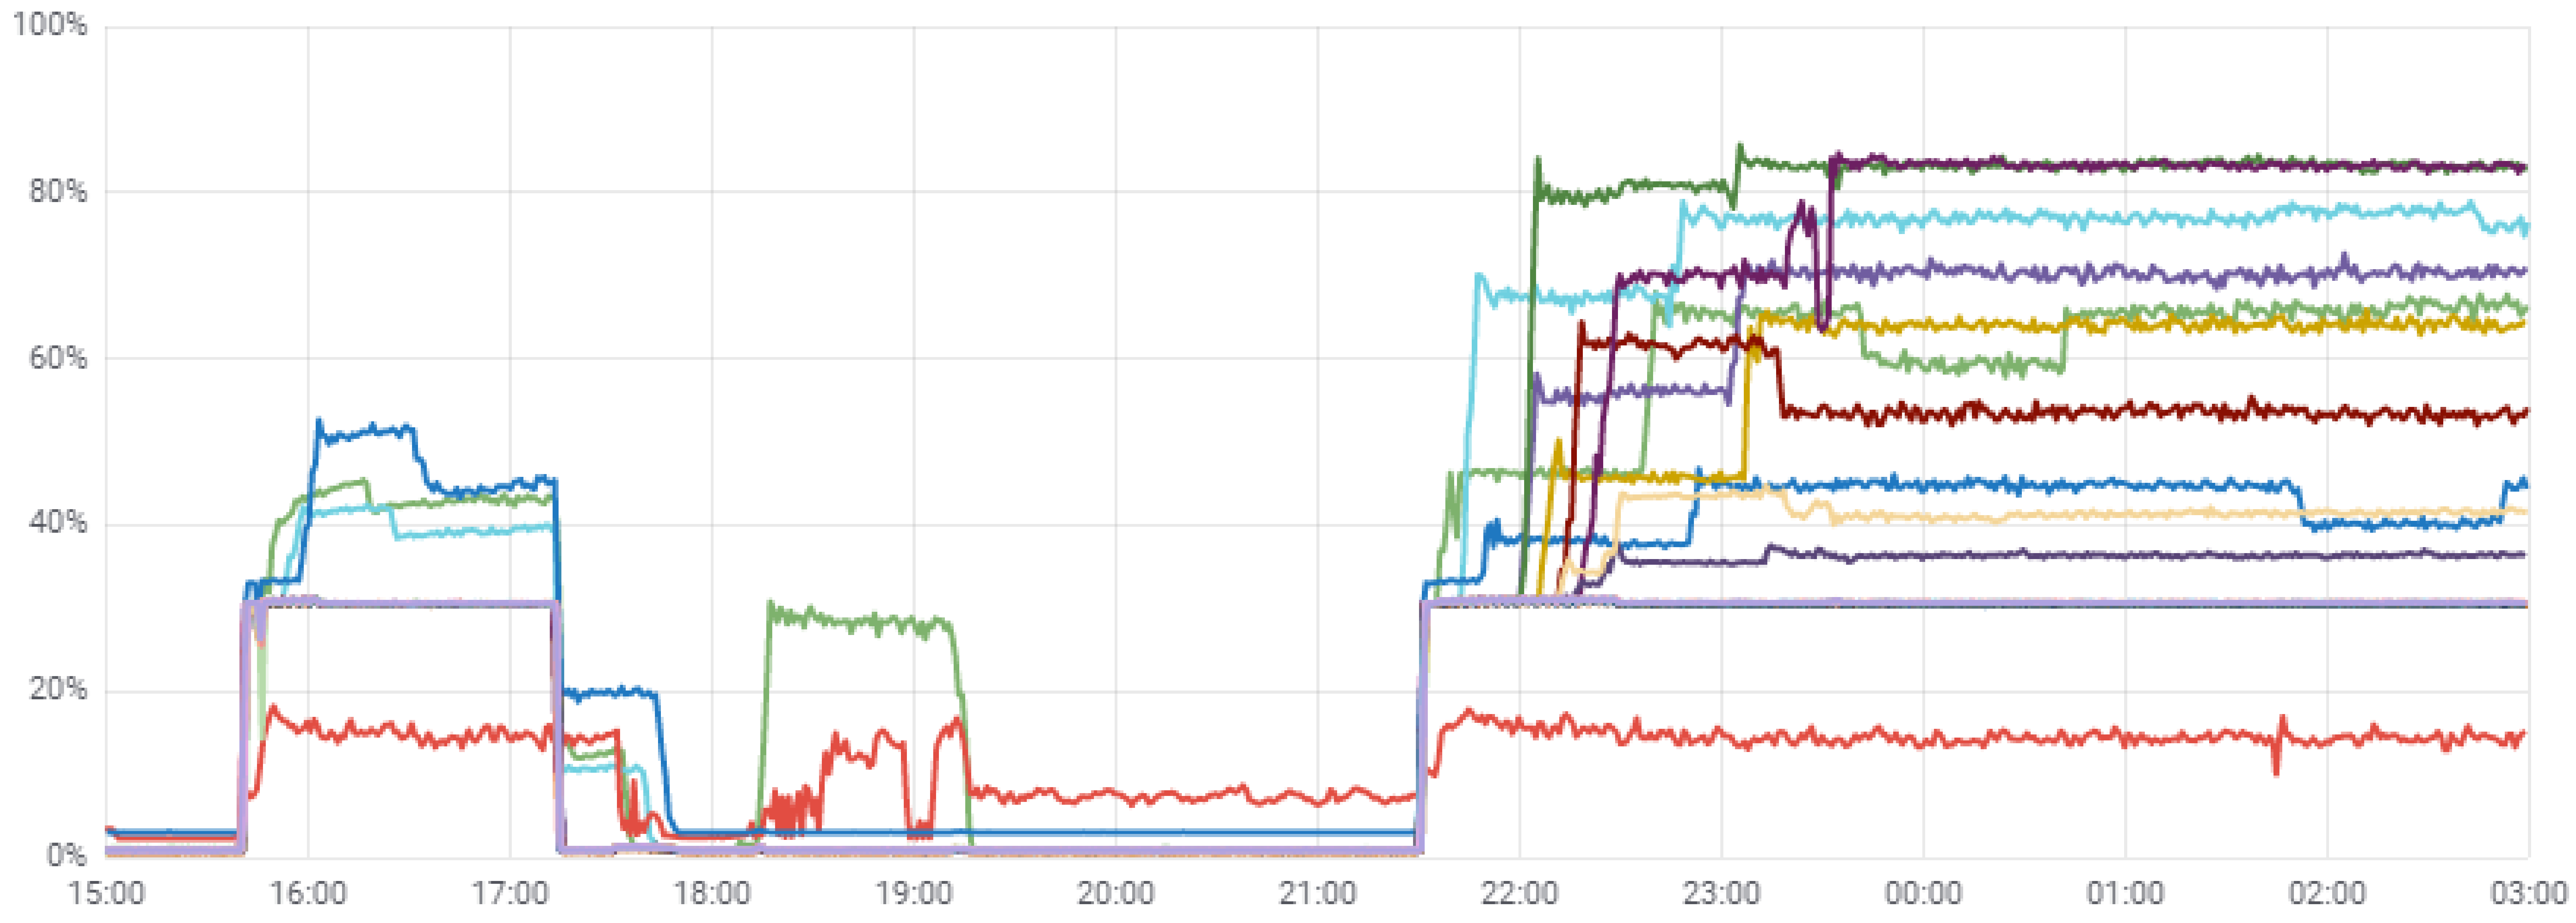

Results in the following profile of power consumption

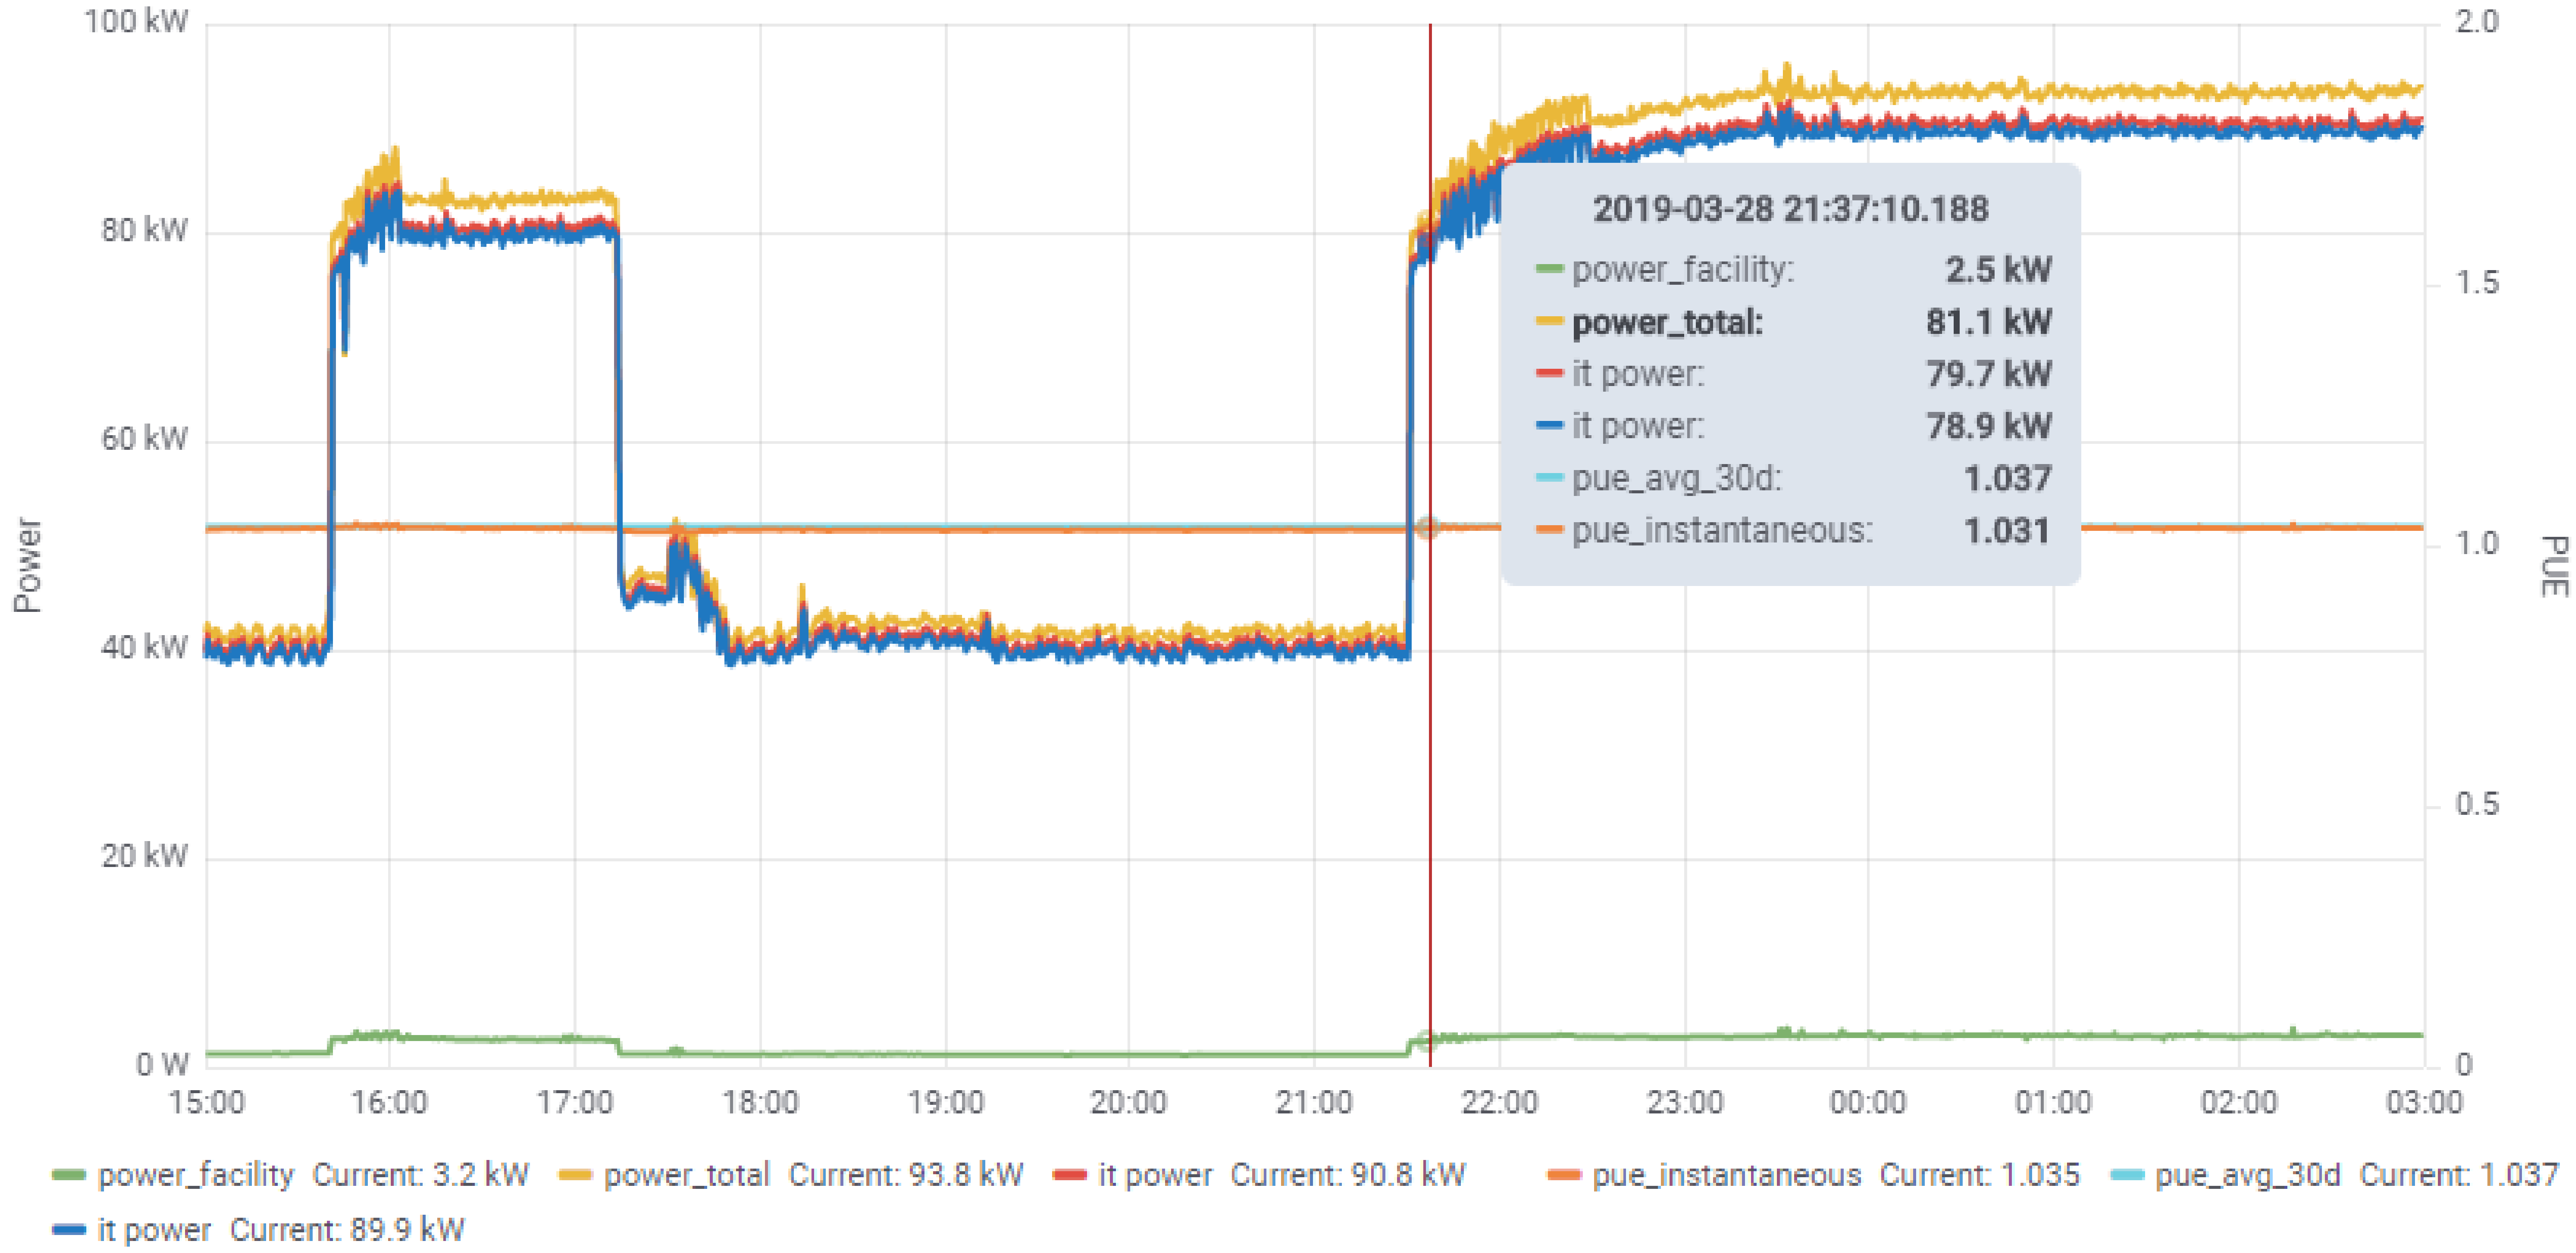

GA 768875

- Changes in IT workloads causes large variation in IT power consumption
- The facility power consumption shows only small changes
- Instantaneous PUE and 30 day average PUE remains fairly constant because of the communication between IT and cooling controller

Visit https://bodentypeDC.eu

Supply and return temperatures

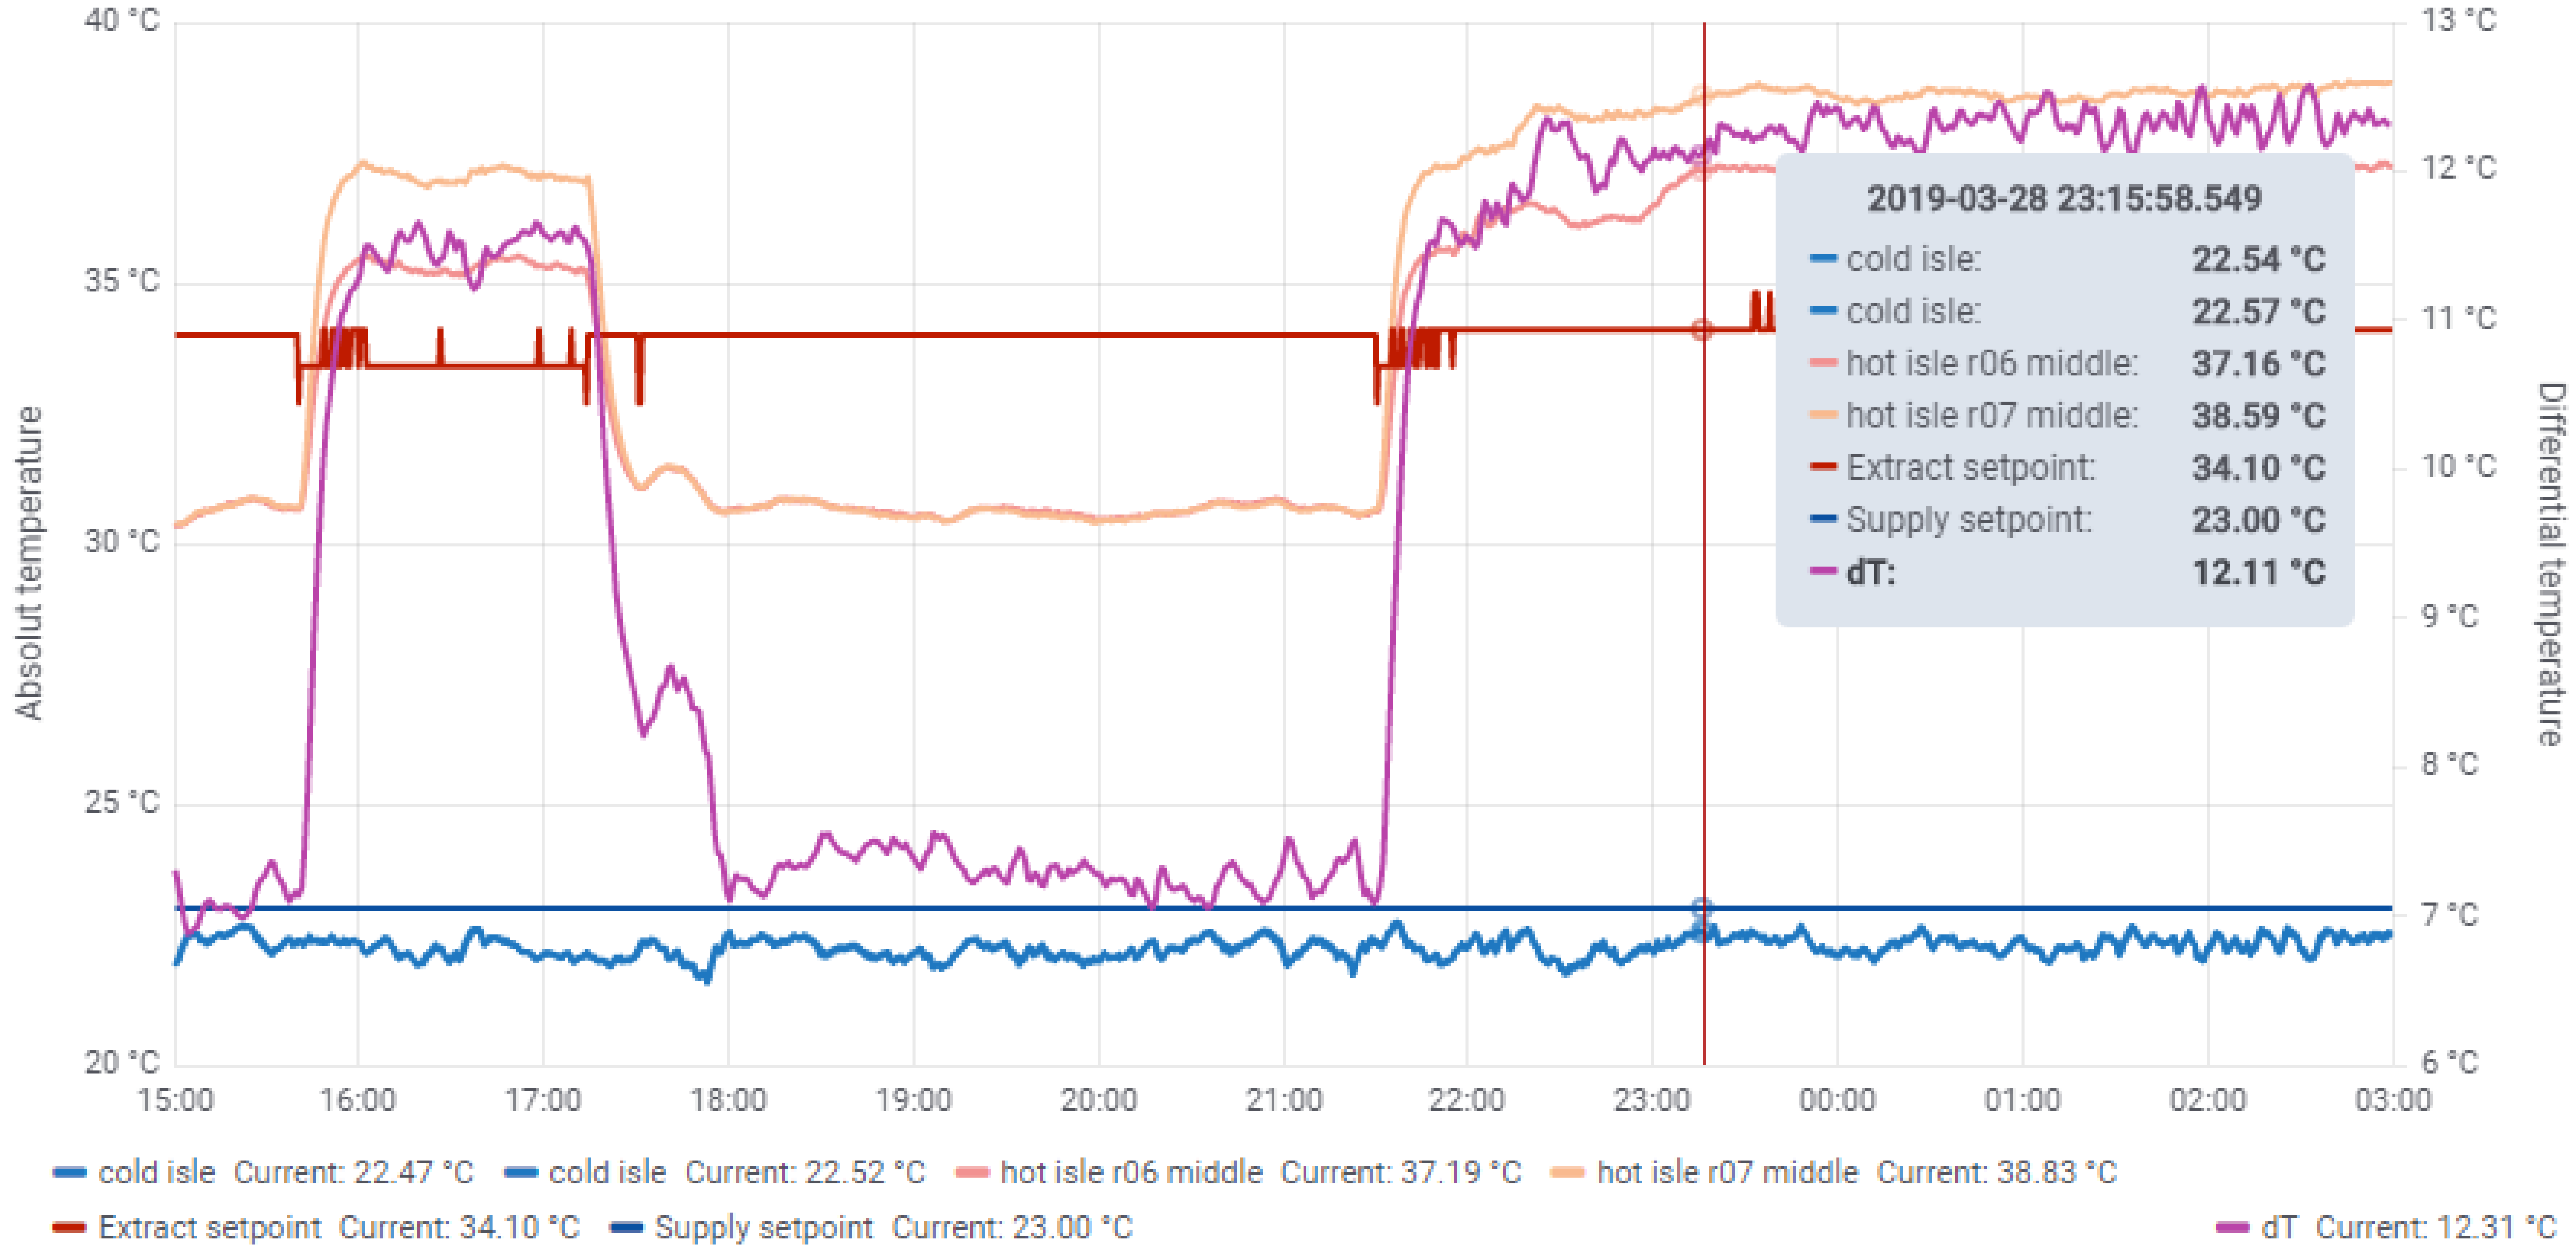

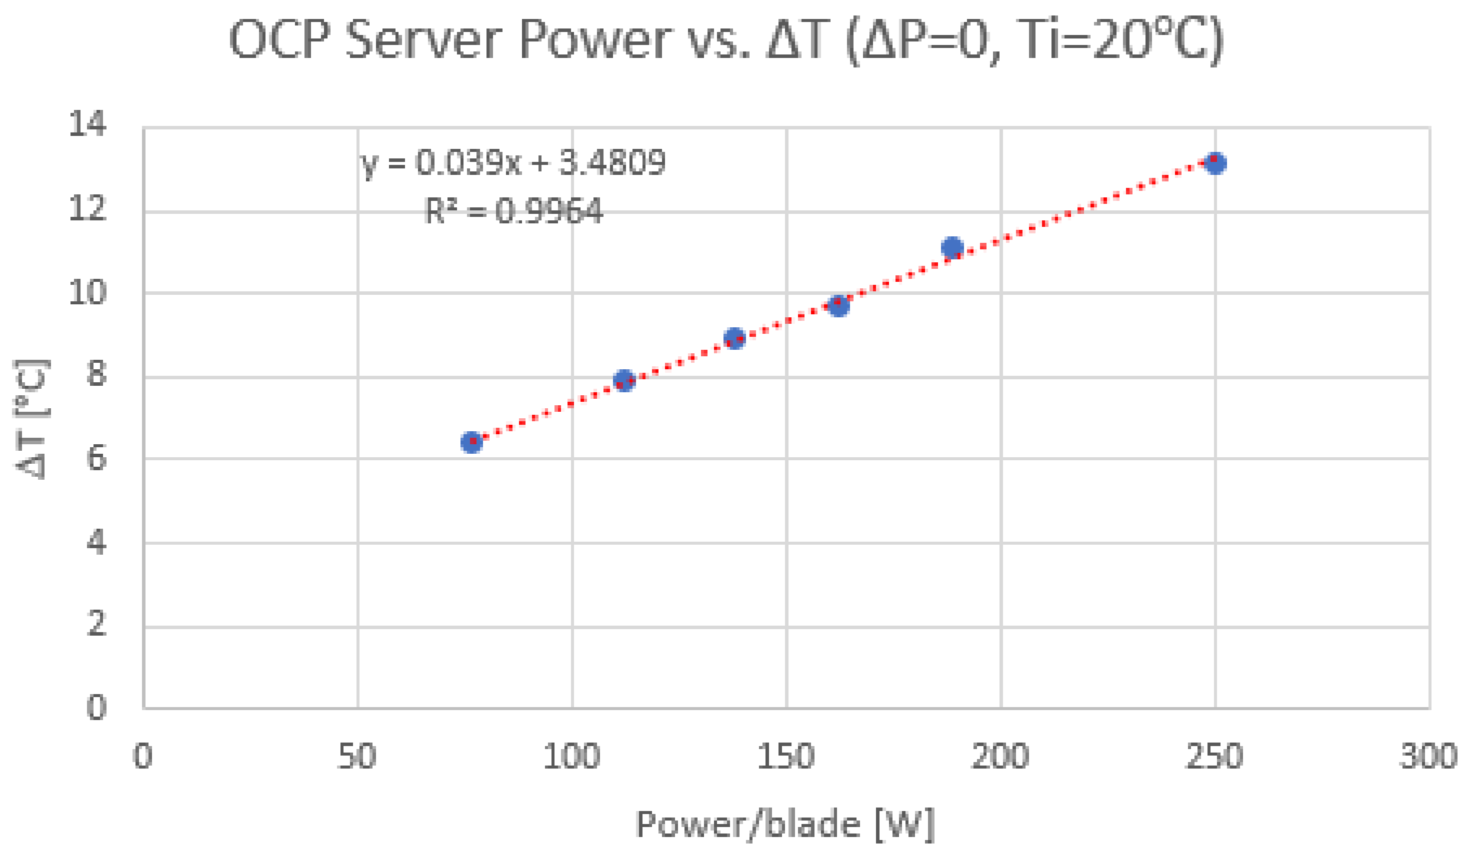

GA 768875

- The hot aisle temperatures varies with varying IT workload
- The supply temperature to the cold aisle remain at desired setpoint
- The delta T across the IT follows the expected values based on the relationship studied in the wind tunnel

Visit https://bodentypeDC.eu

Differential pressure and cooling units fan speeds

GA 768875

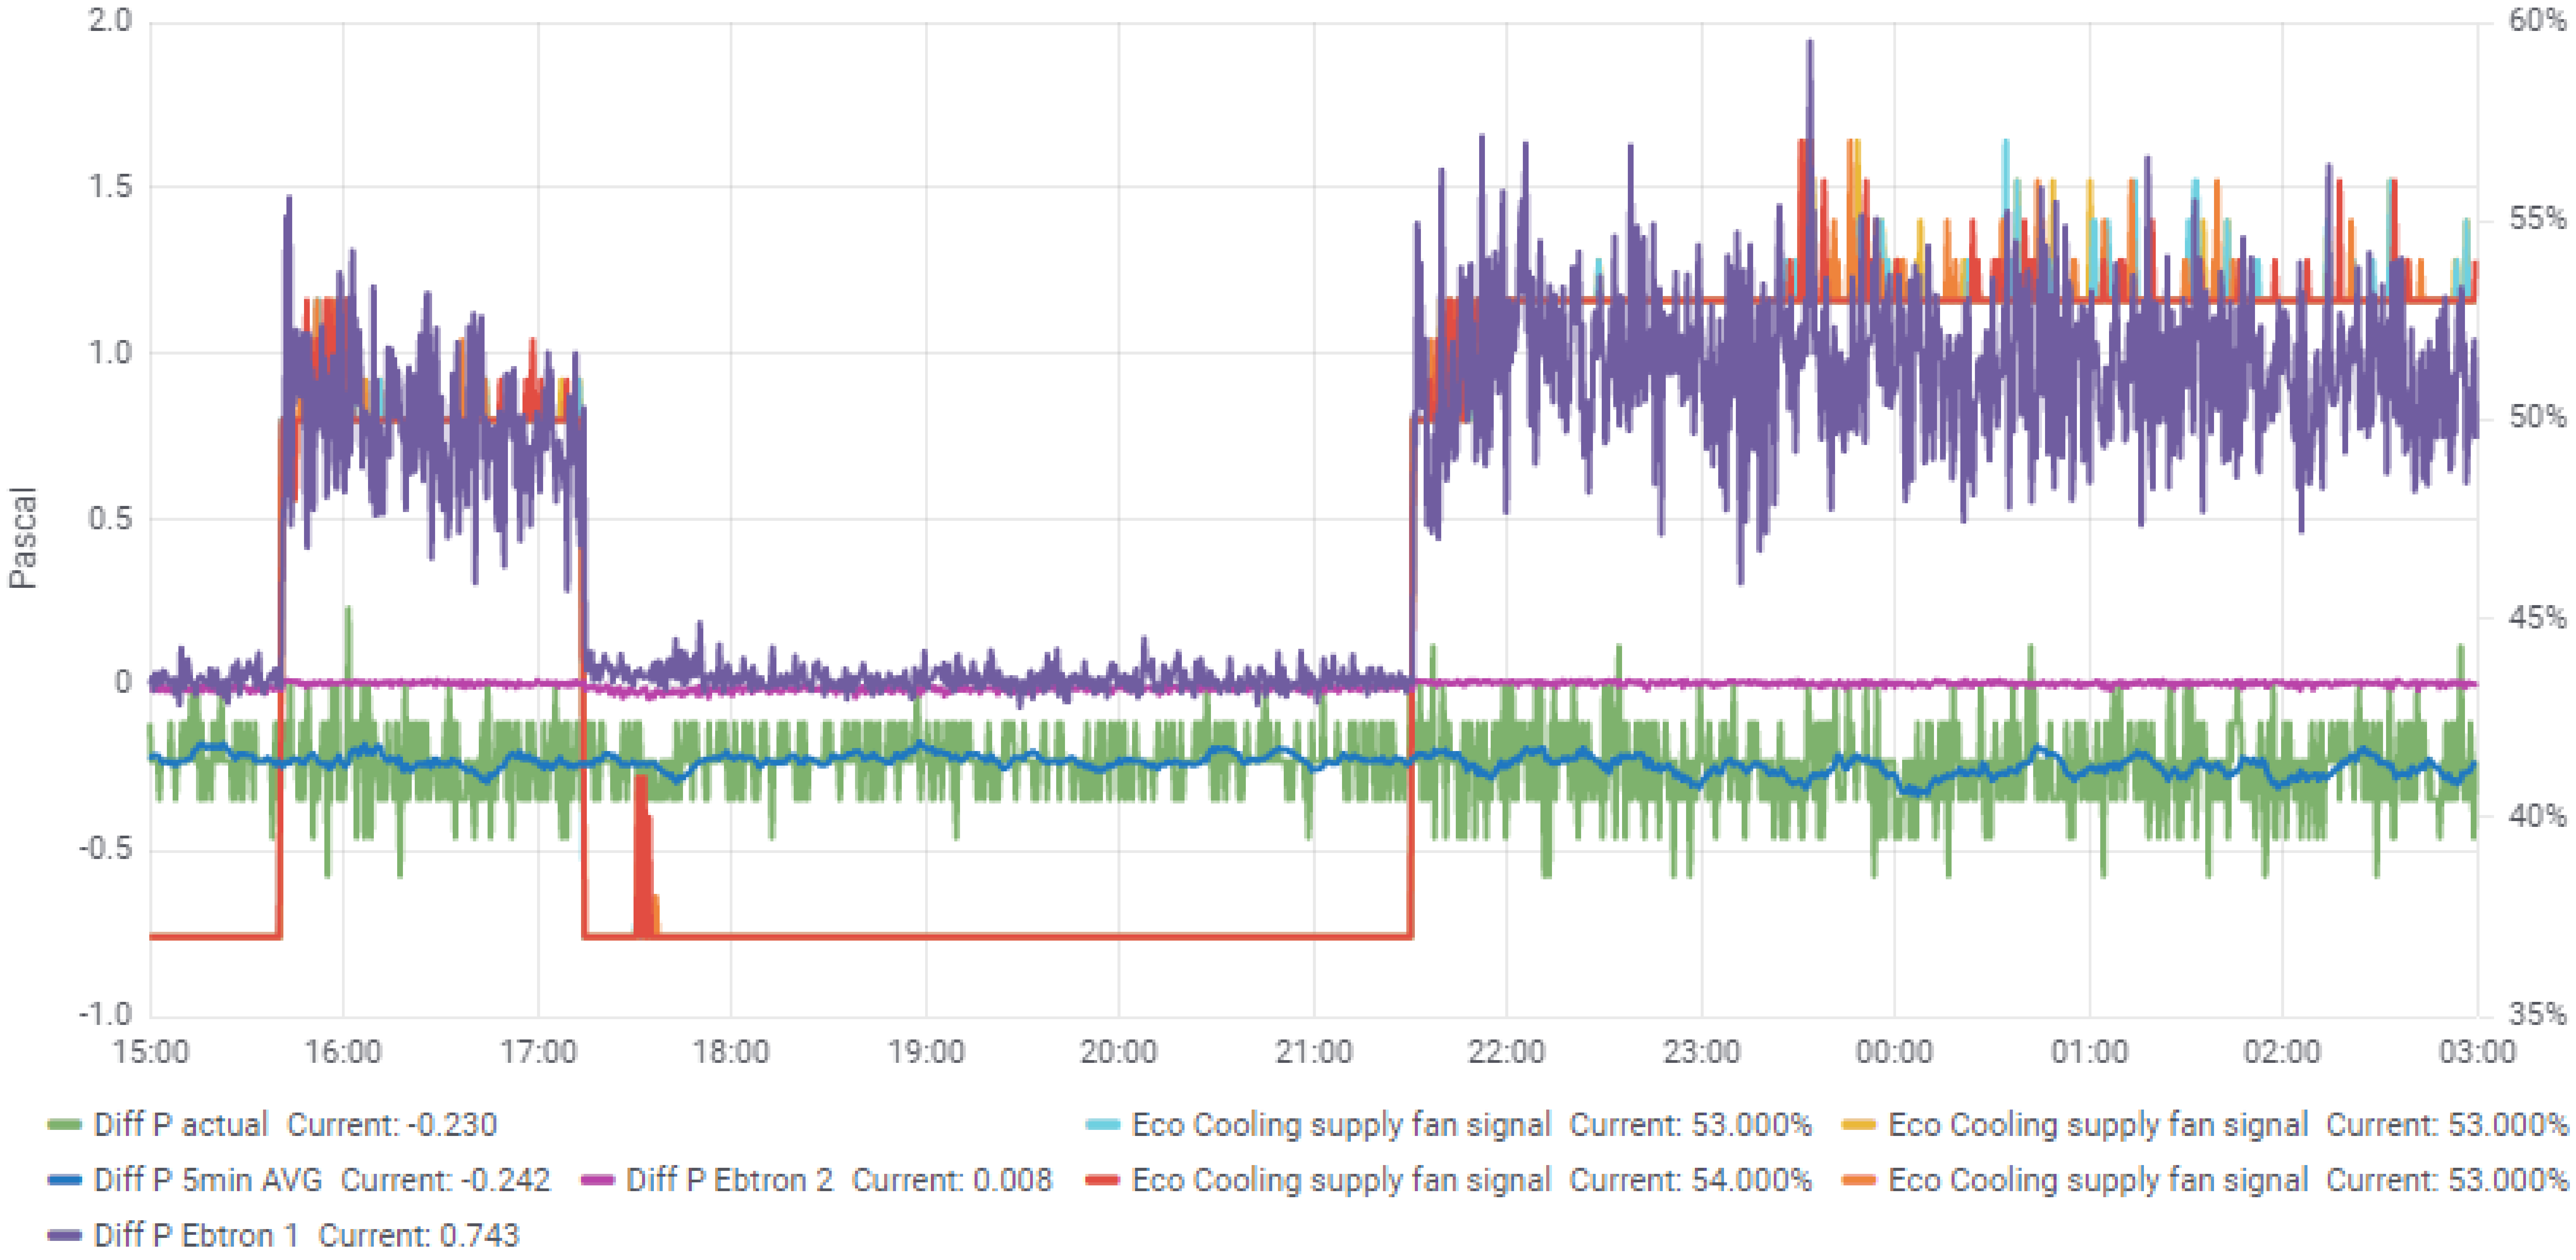

Heat sinks in OCP Servers

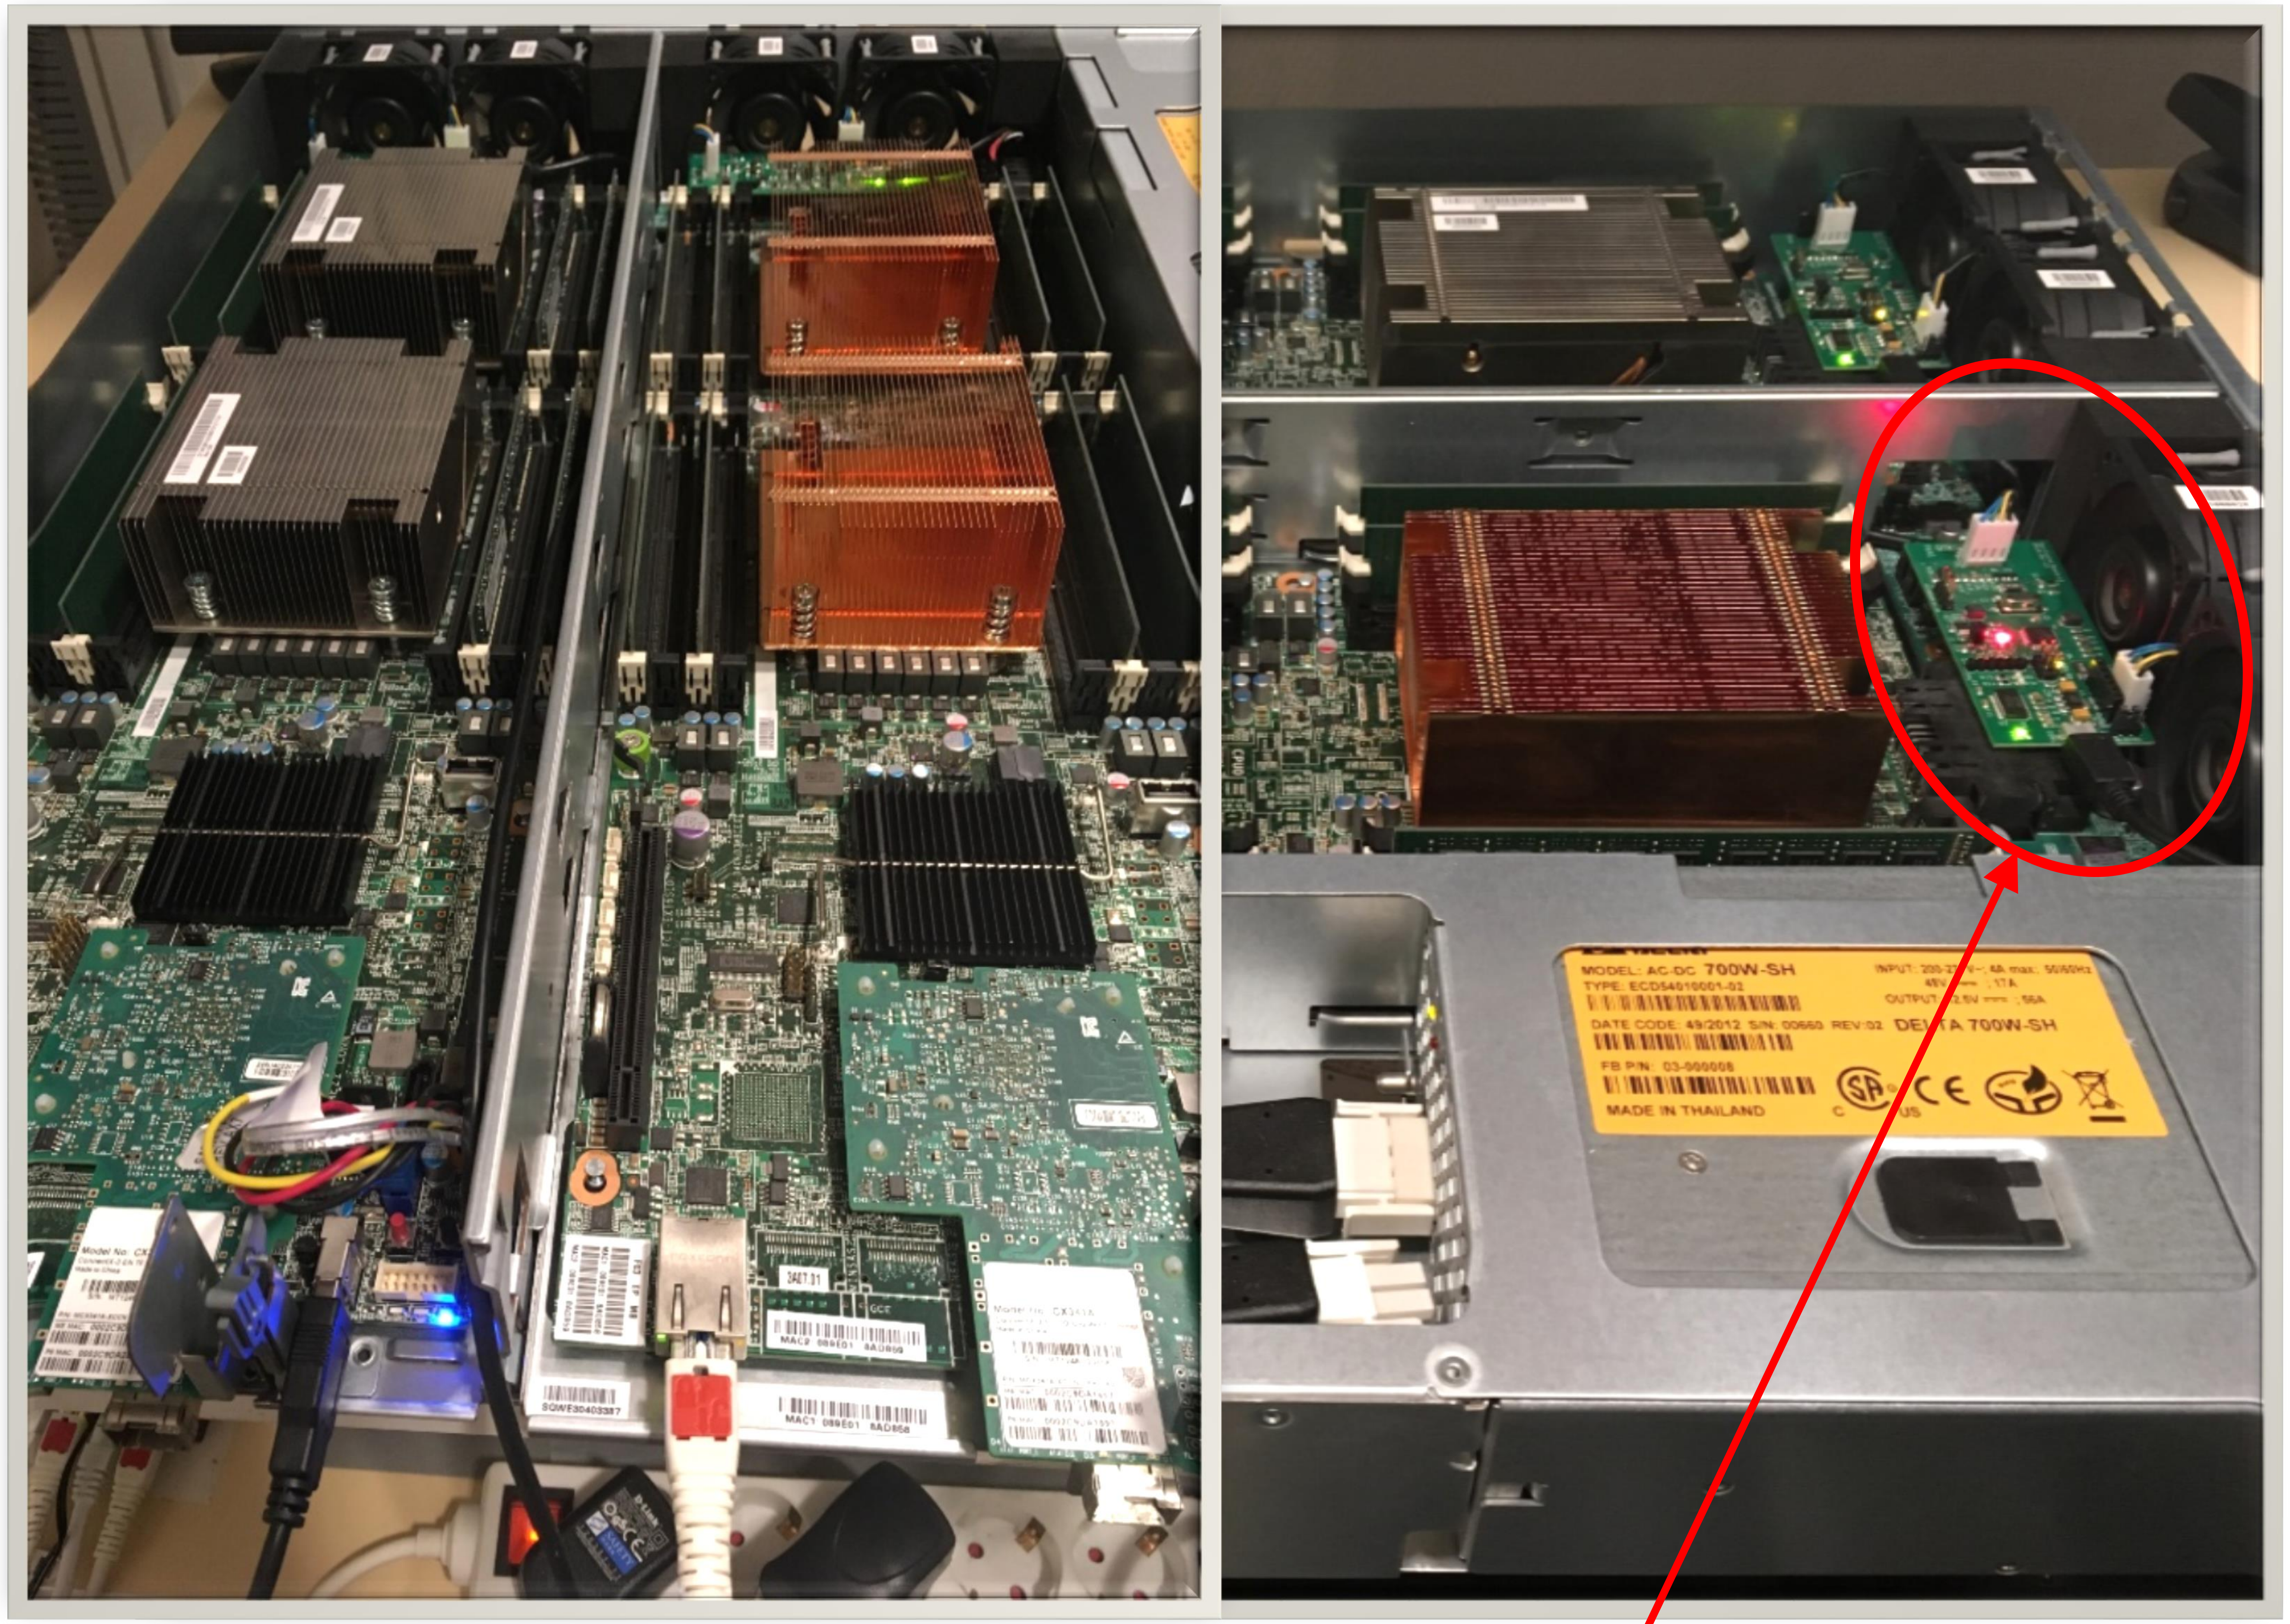

LOCALLY DEVELOPED FAN CONTROLLER

- Aim to achieve lower flow rates and higher delta temperature across the Windmill servers.
- Flowrate control is also achieved by the use of some locally developed fan controller cards.
- Run to a fixed CPU temperature on all servers despite workload or calculate flowrate through each server and provide this information to the fresh air cooling system.
- Replaced aluminium heat sinks containing embedded heat pipes with copper heat sinks.

Phase 2 of the holistic cooling – OCP servers

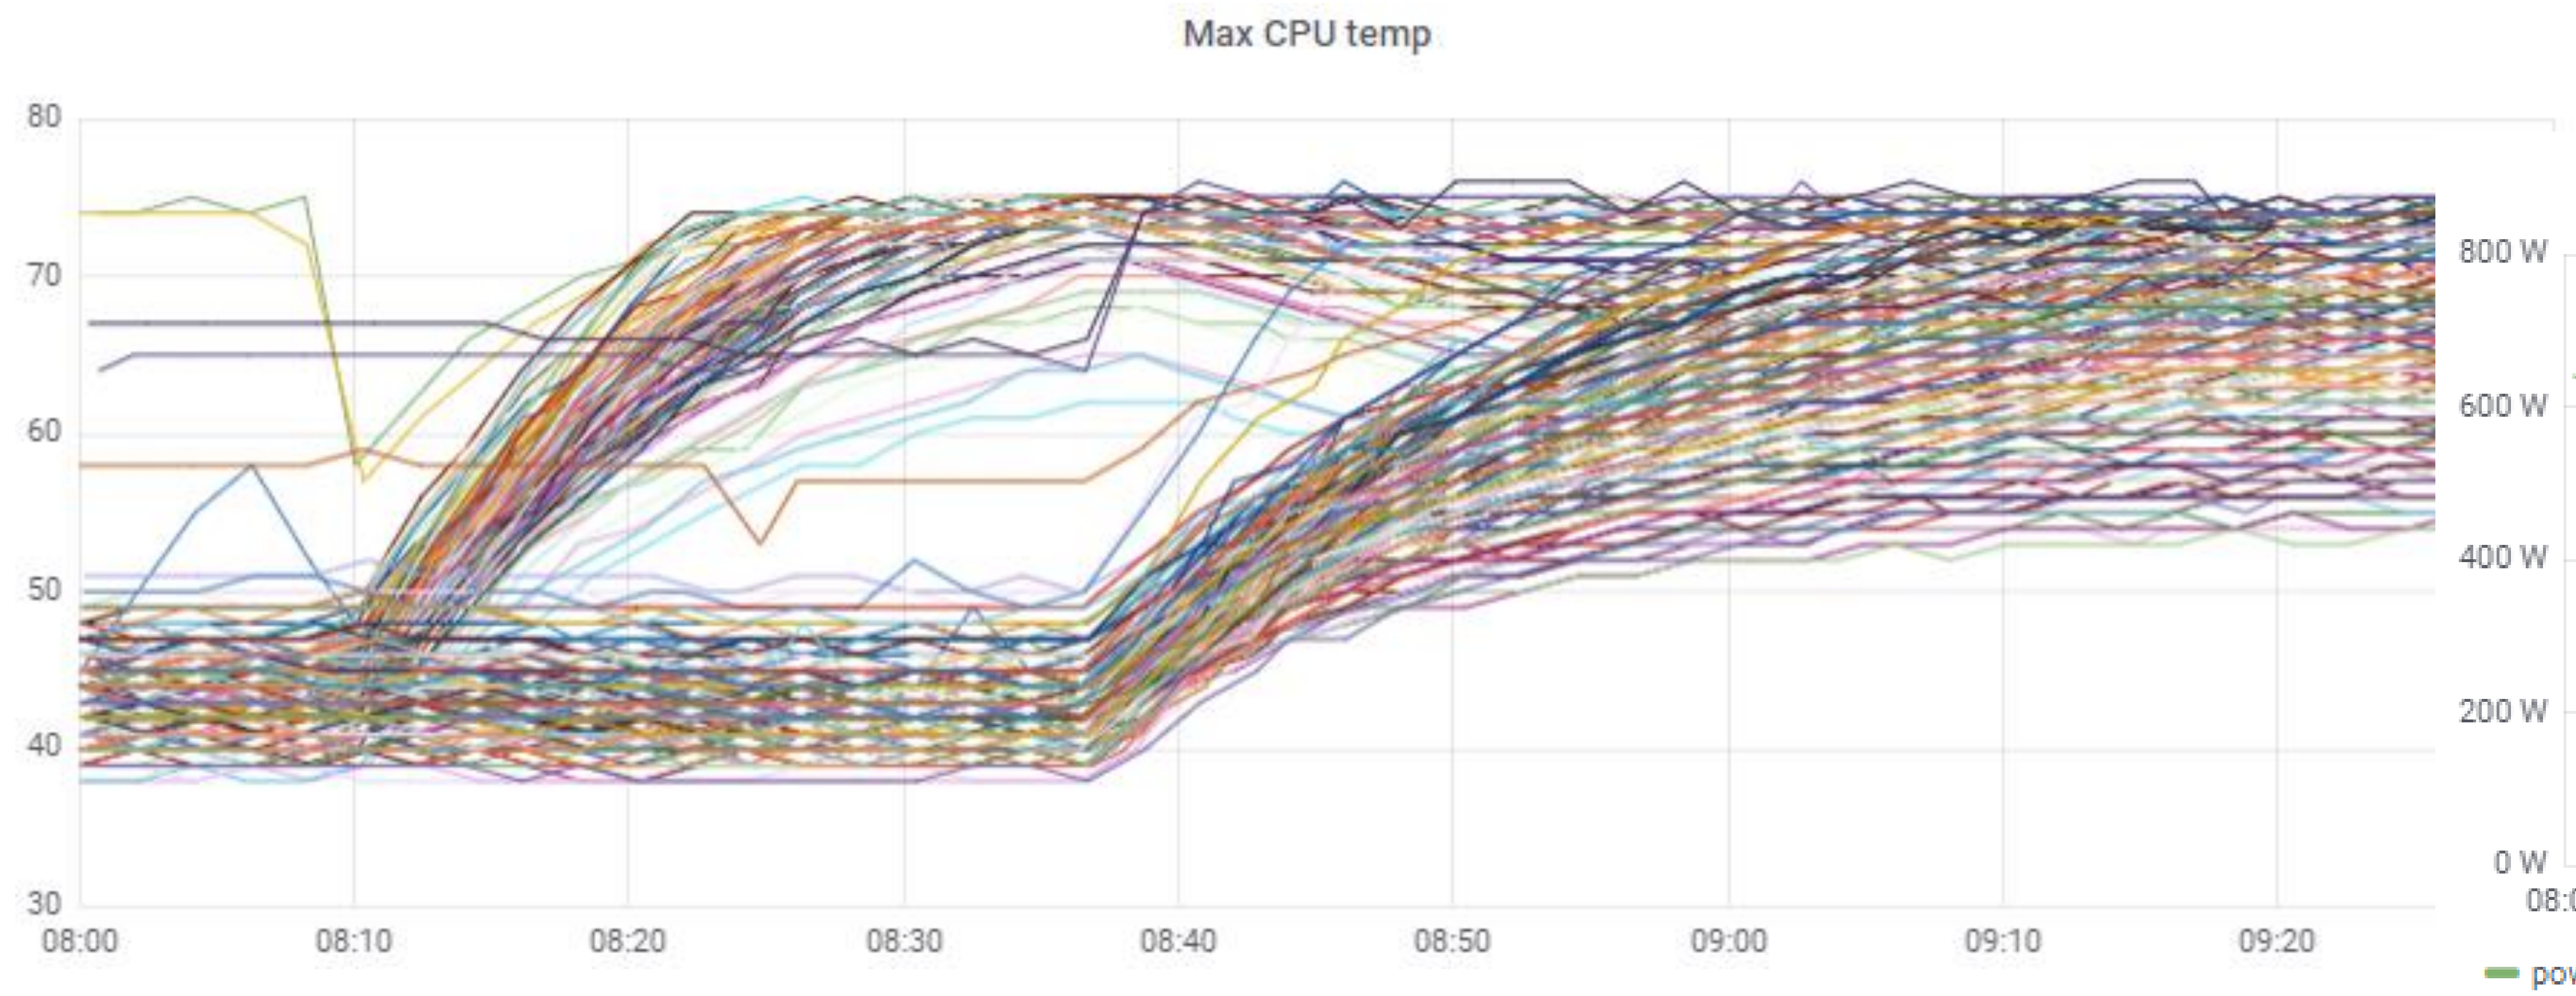

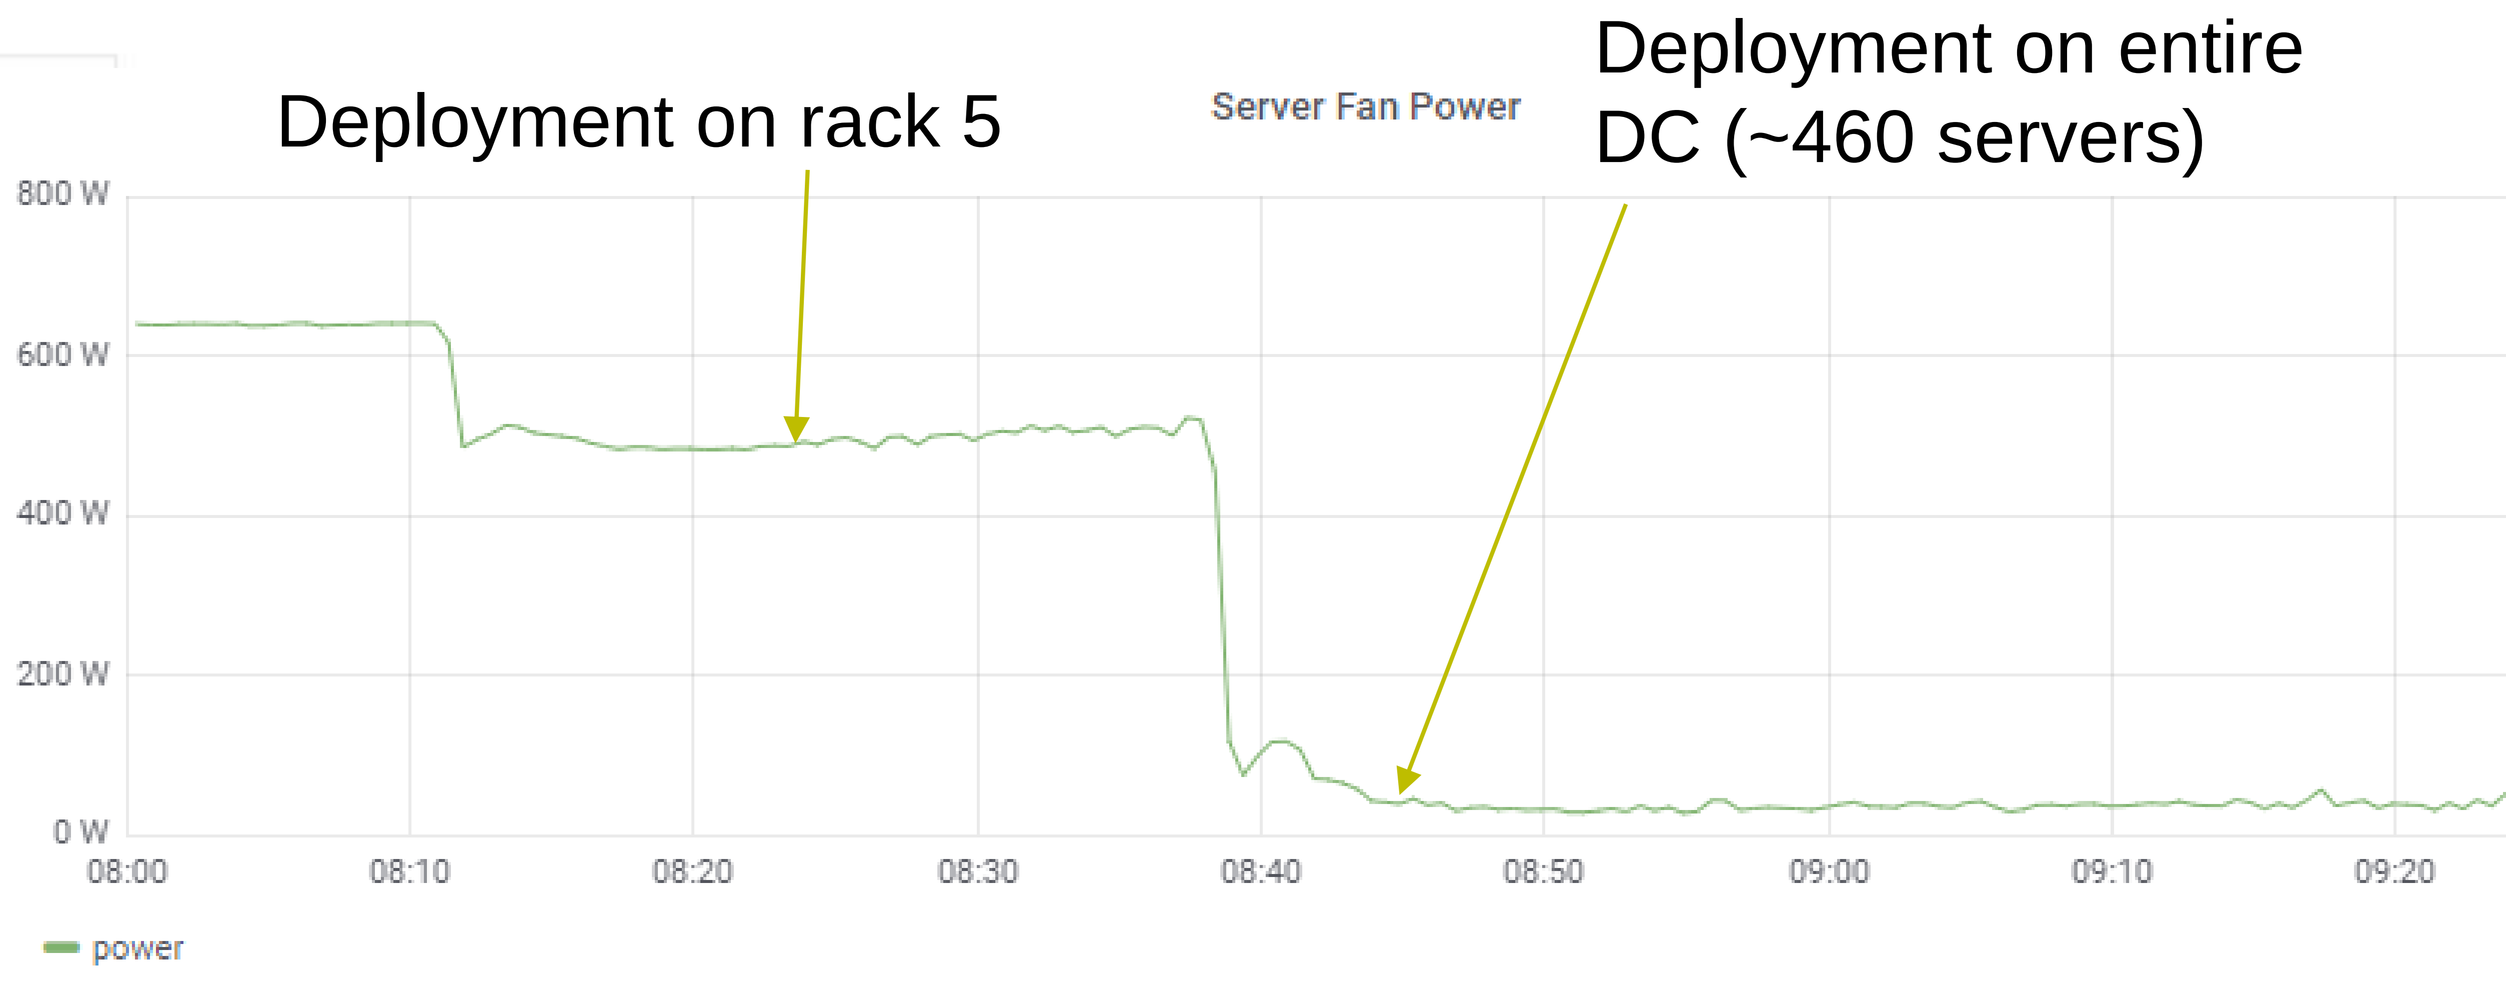

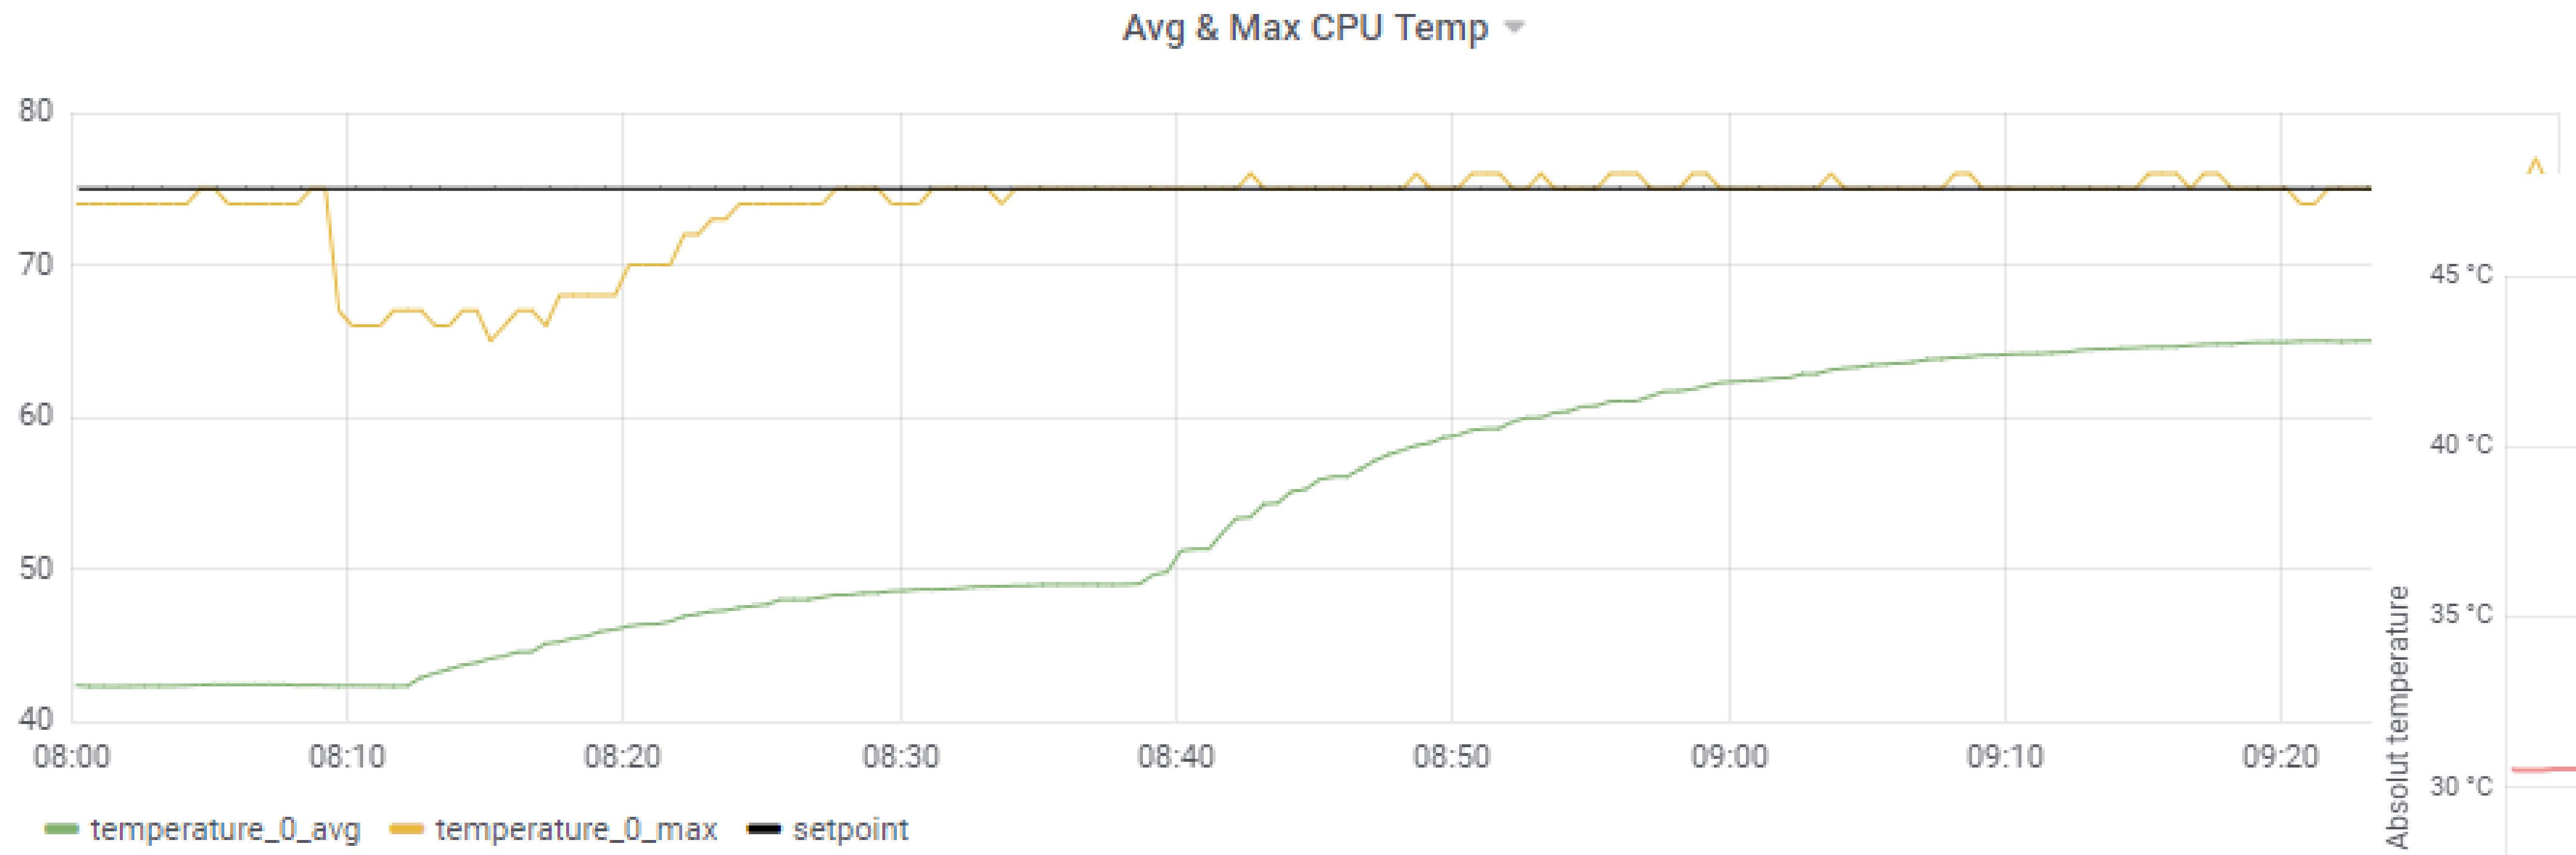

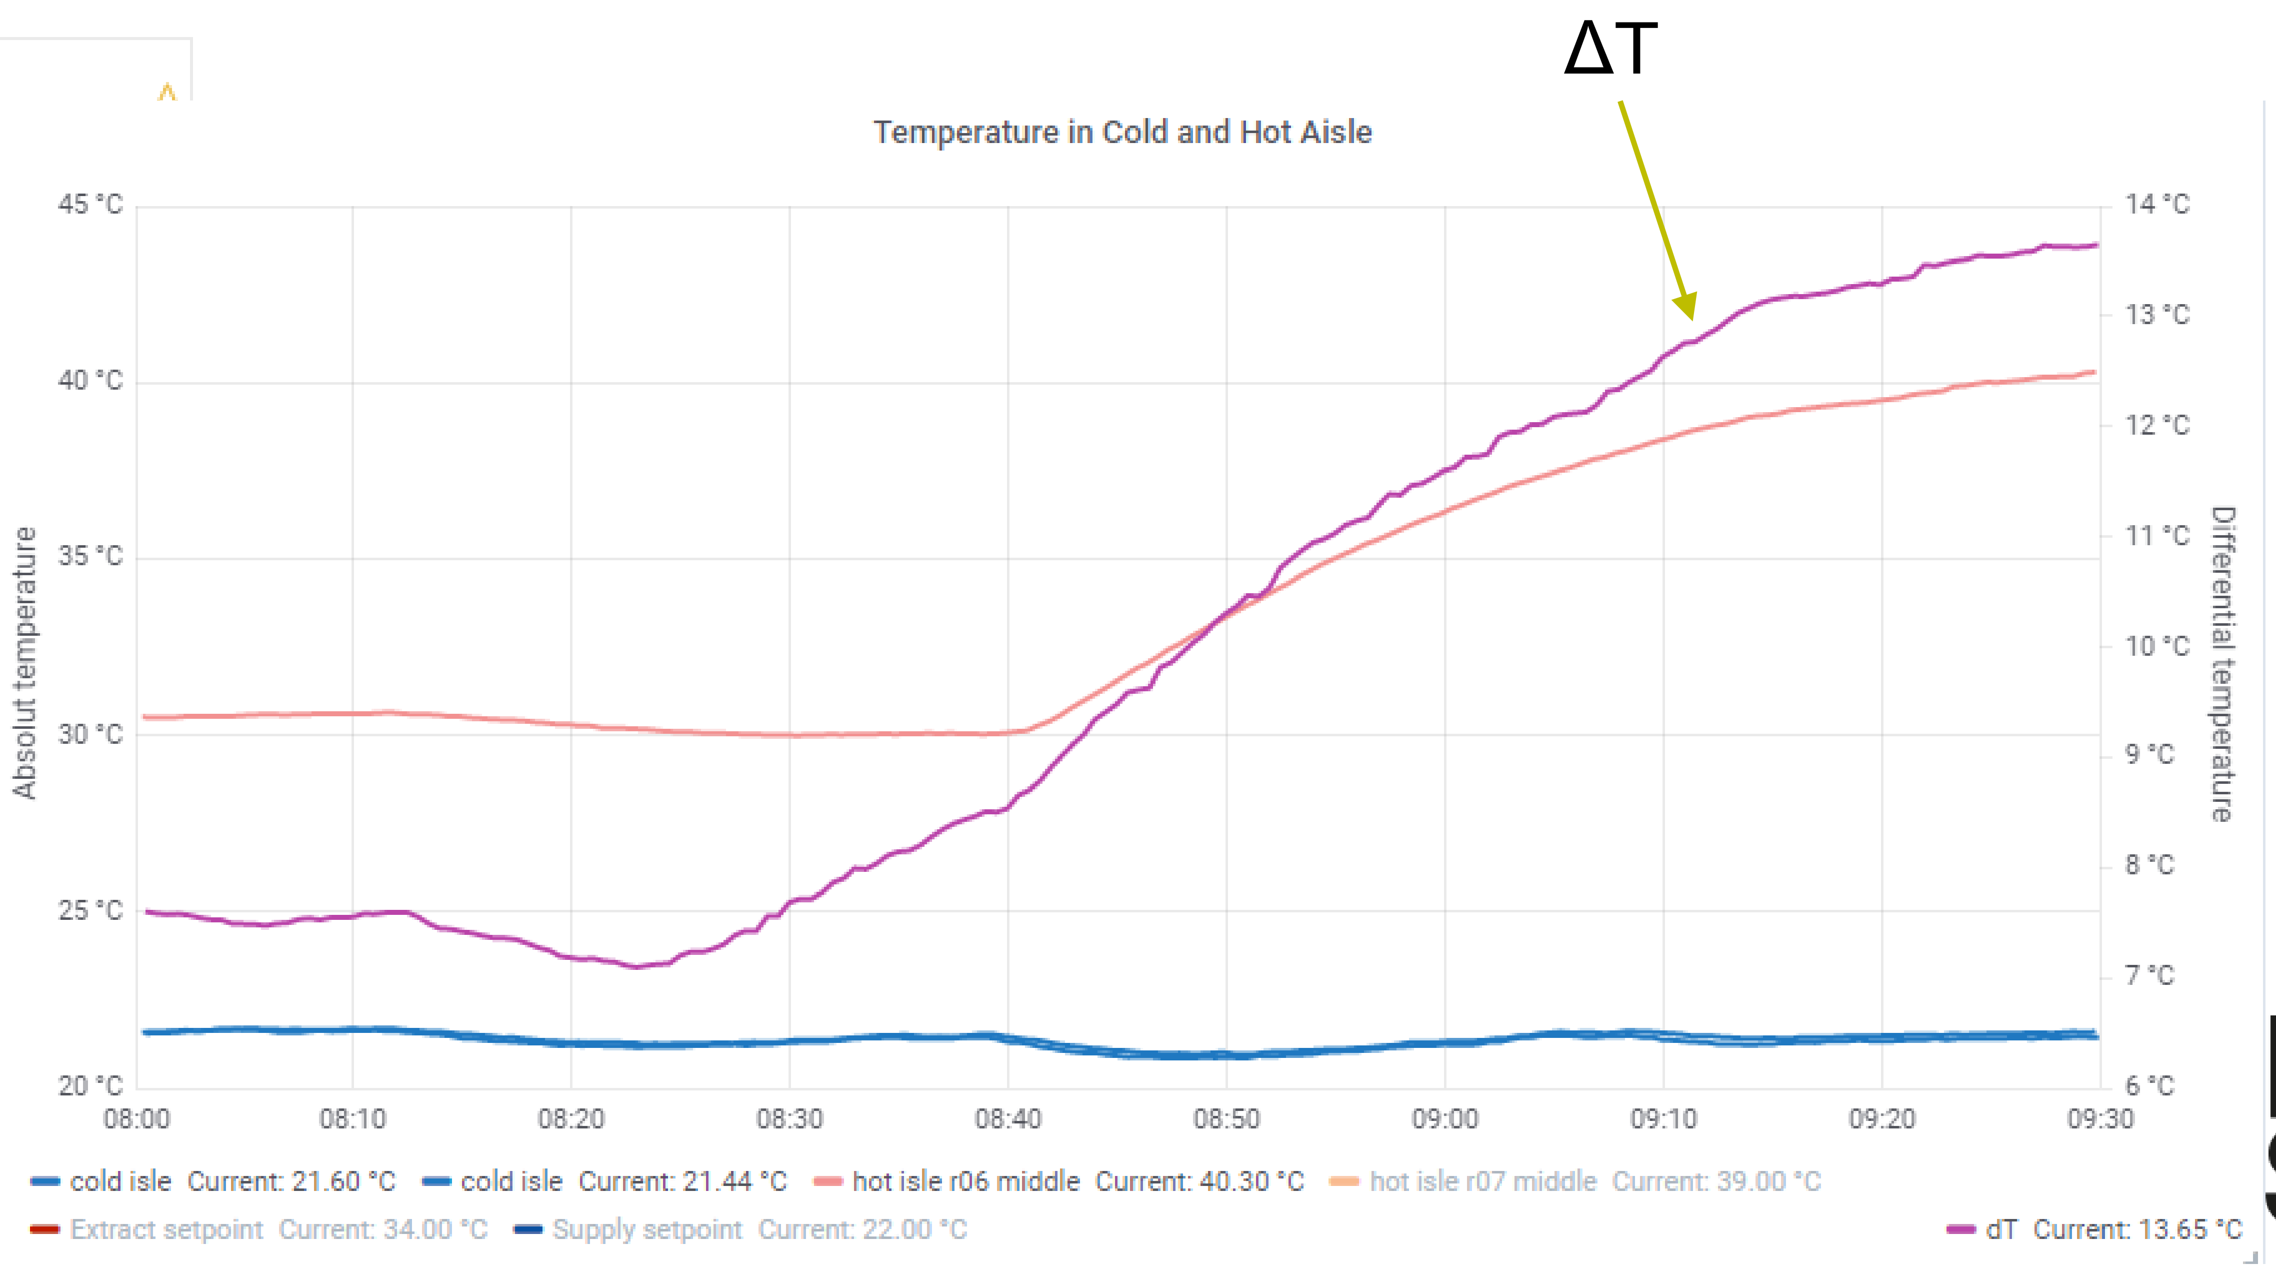

Phase 2 of the holistic cooling – OCP servers

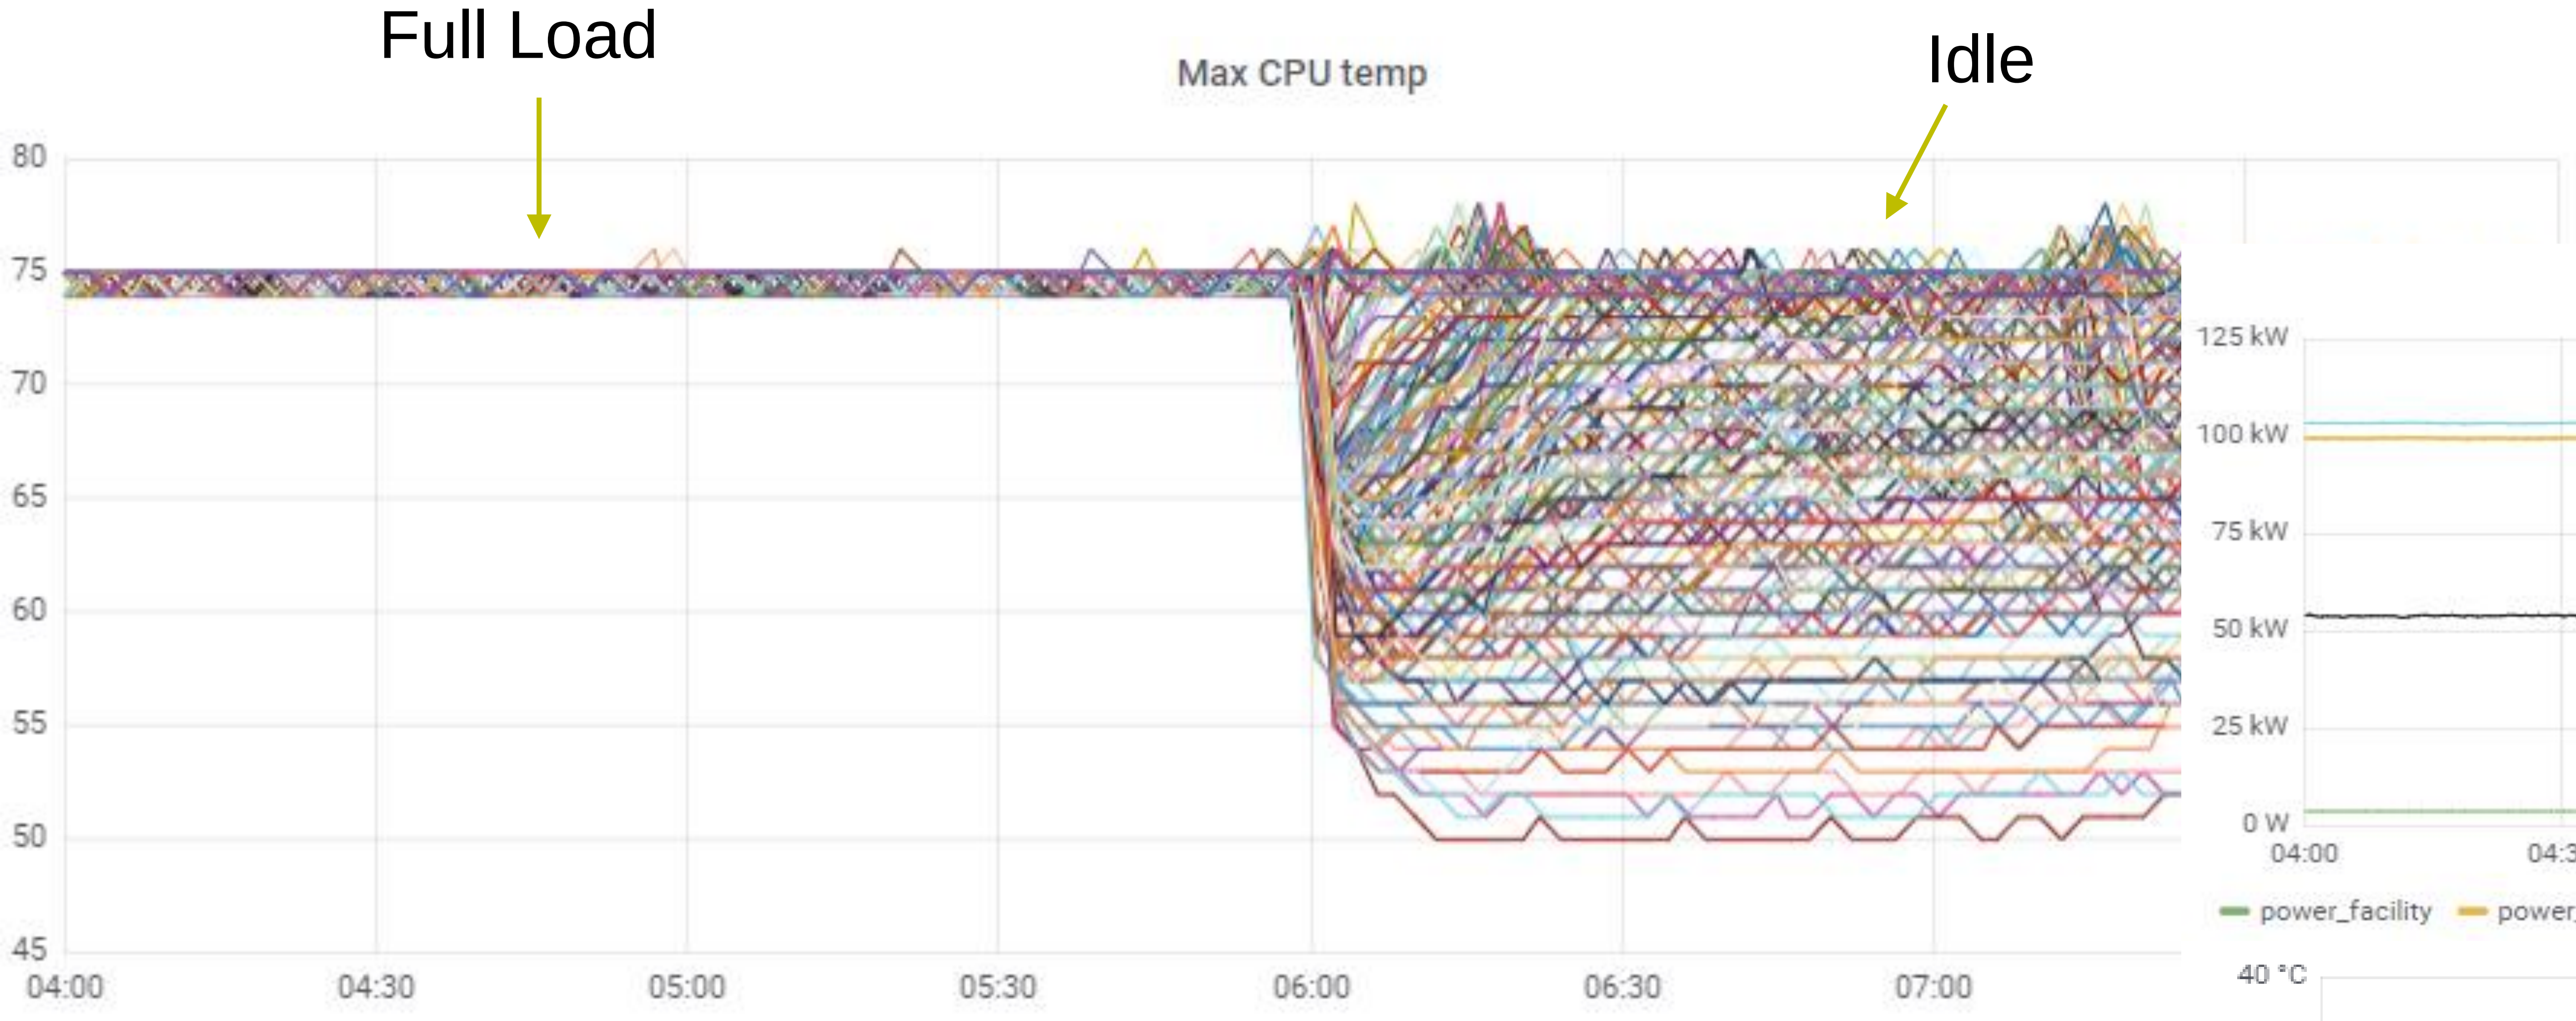

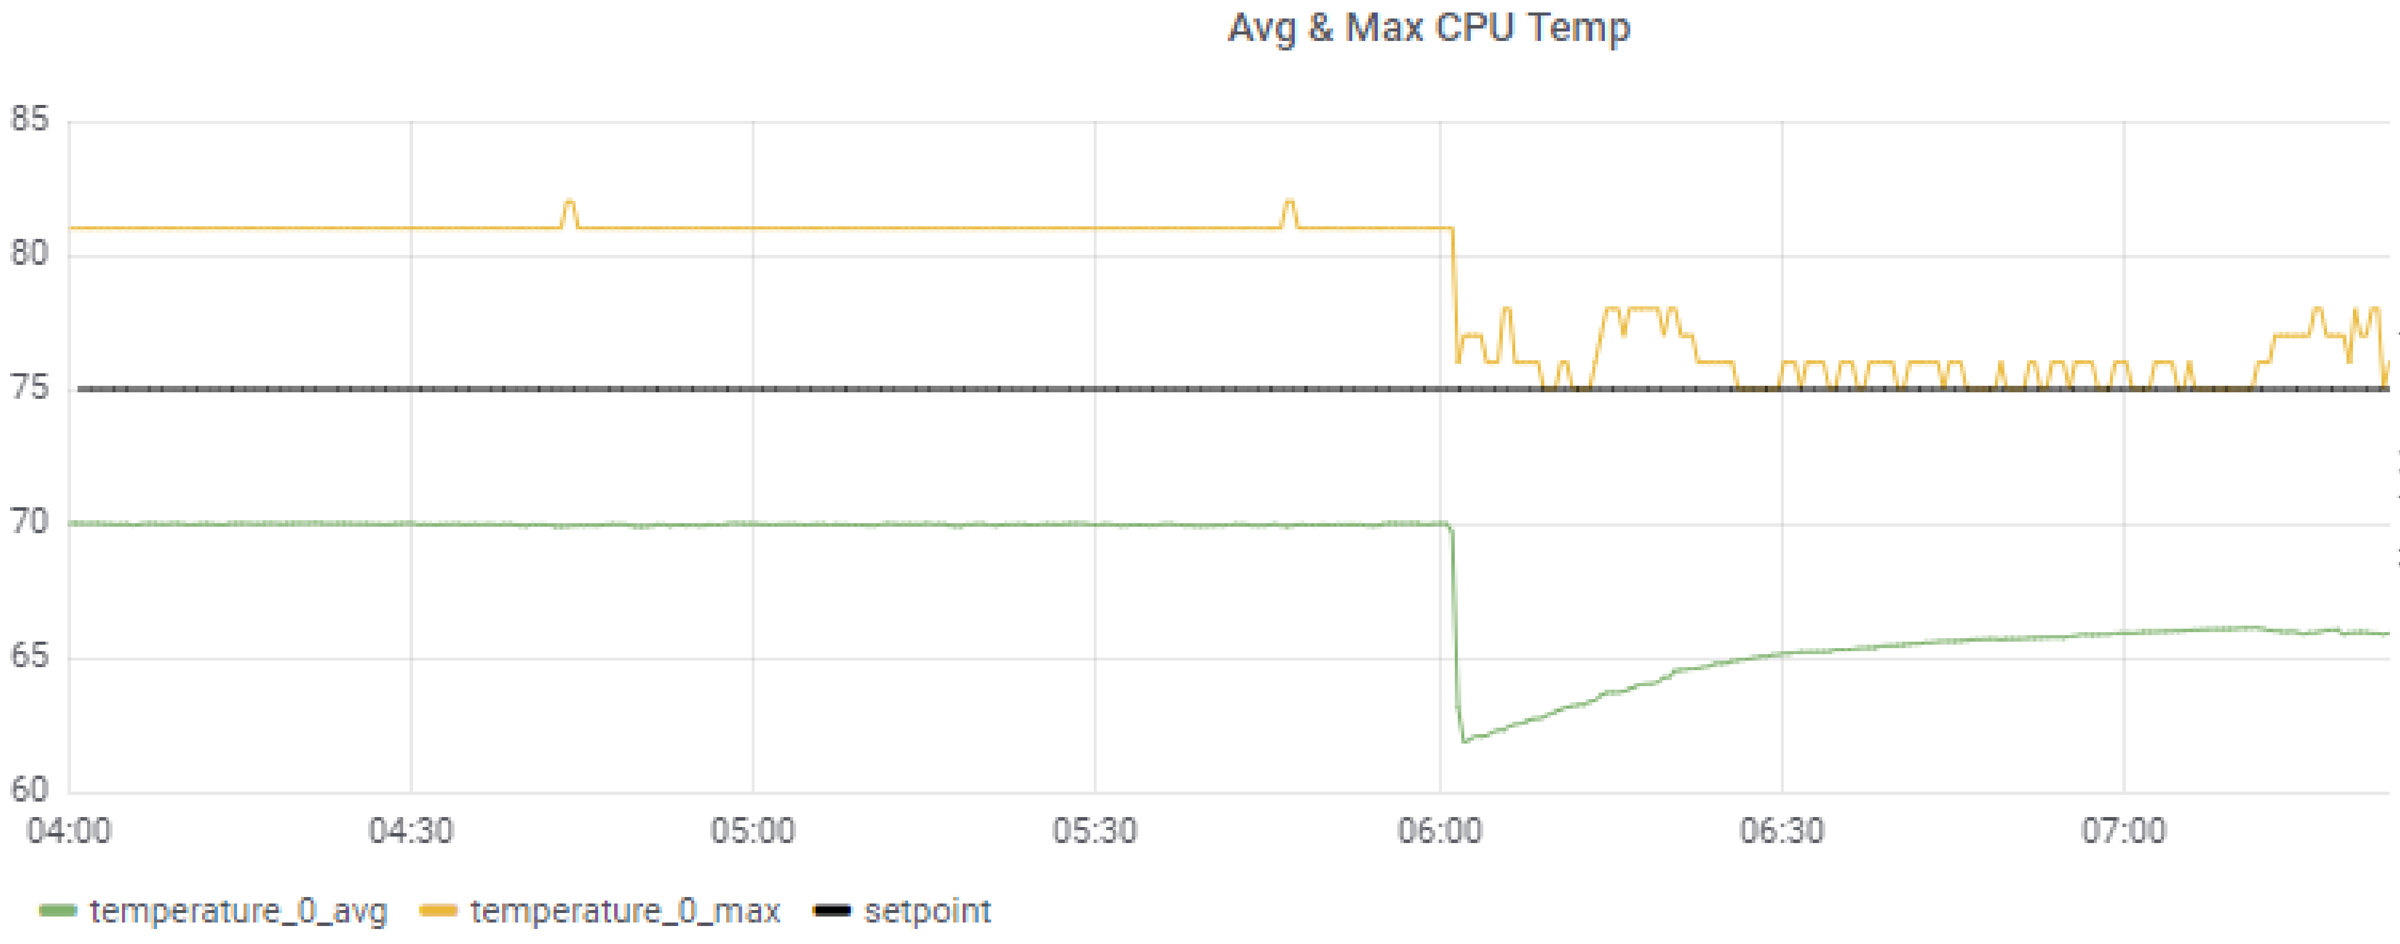

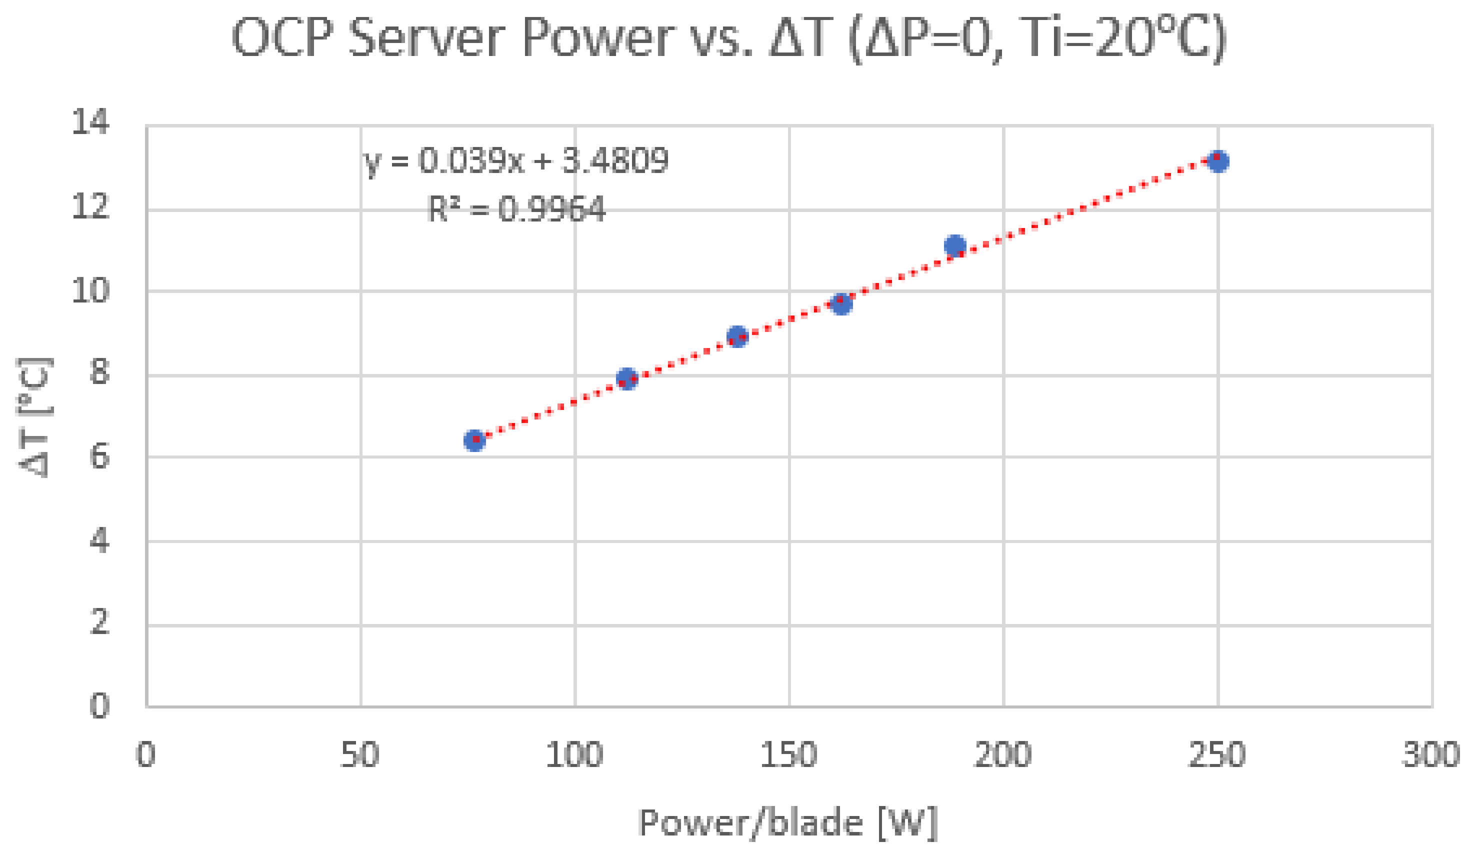

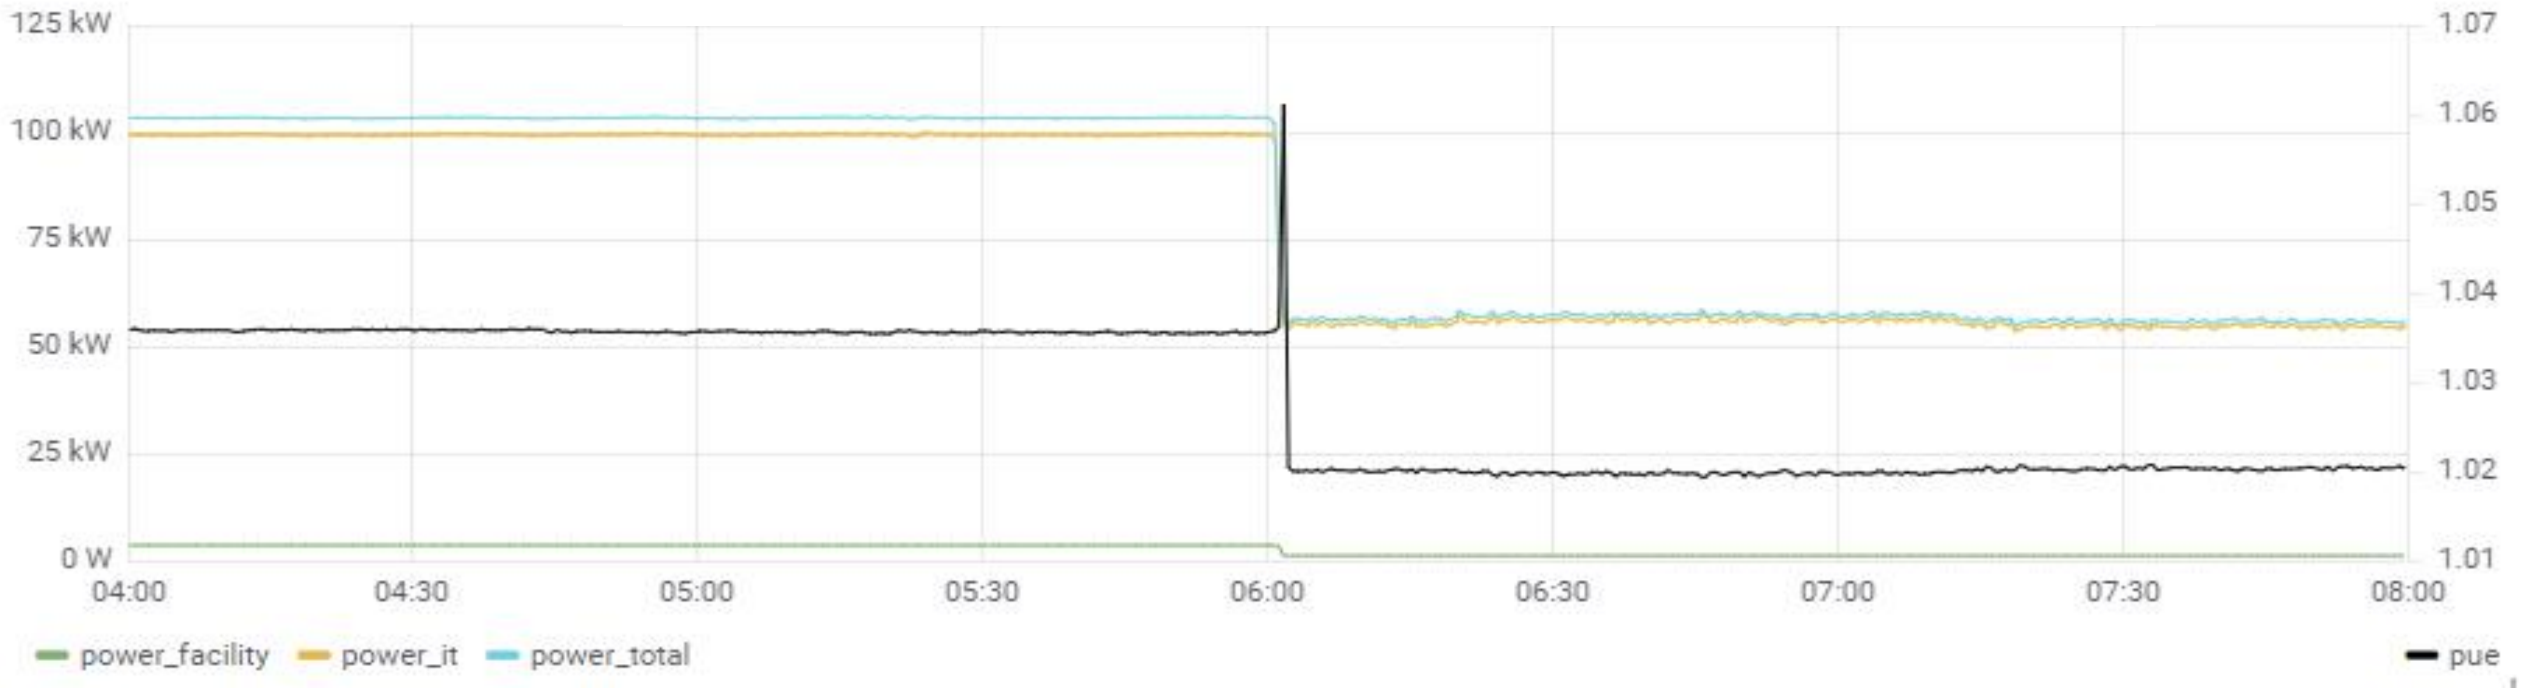

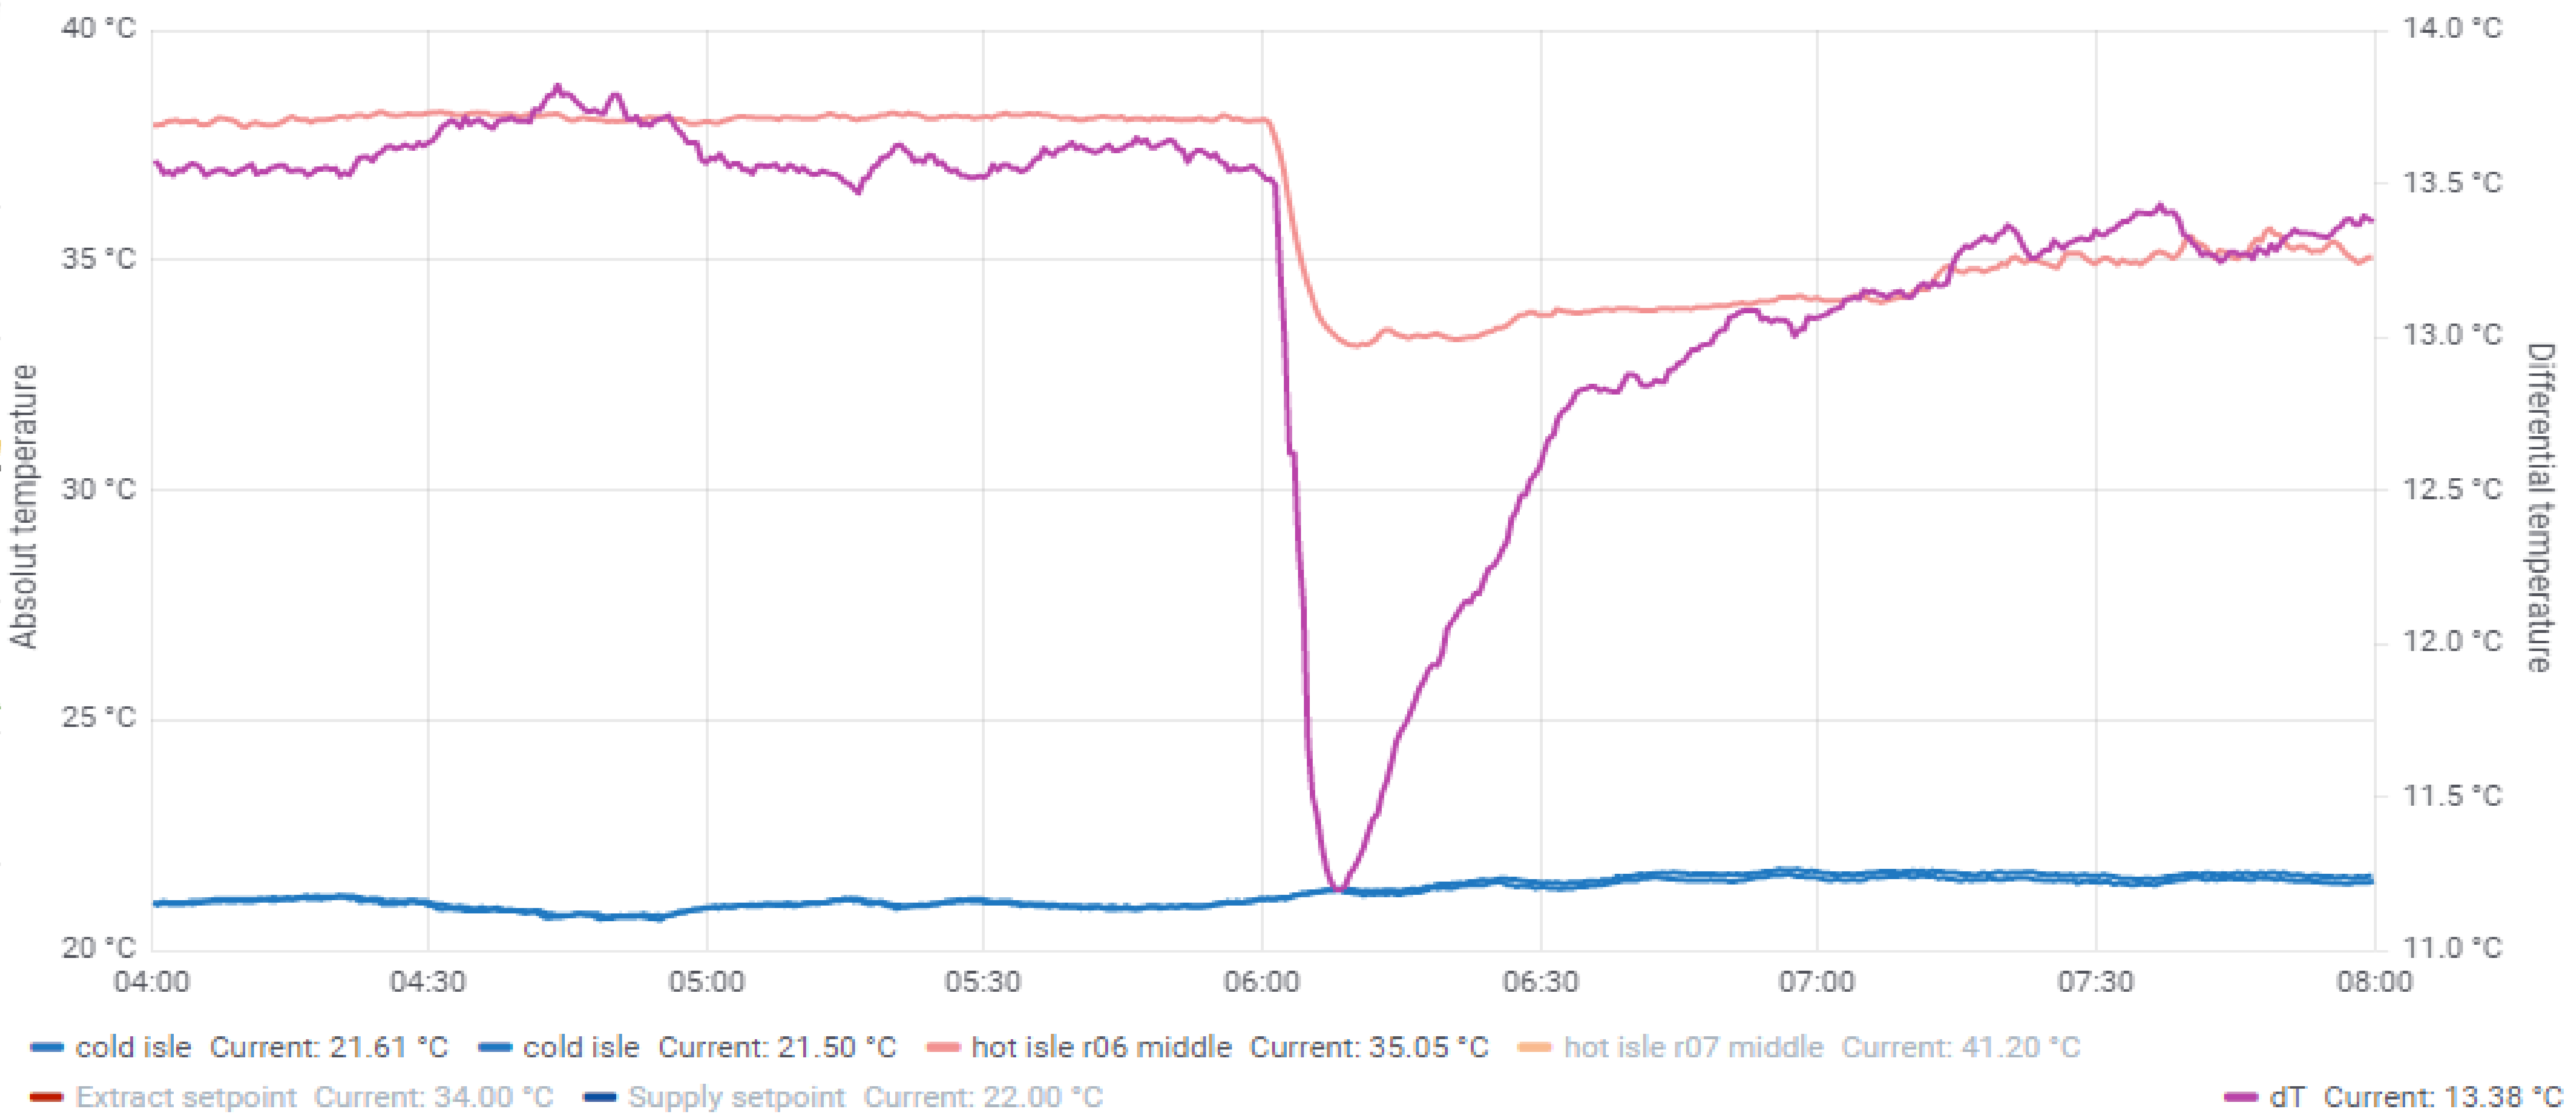

Phase 2 of the holistic cooling – cooling systems

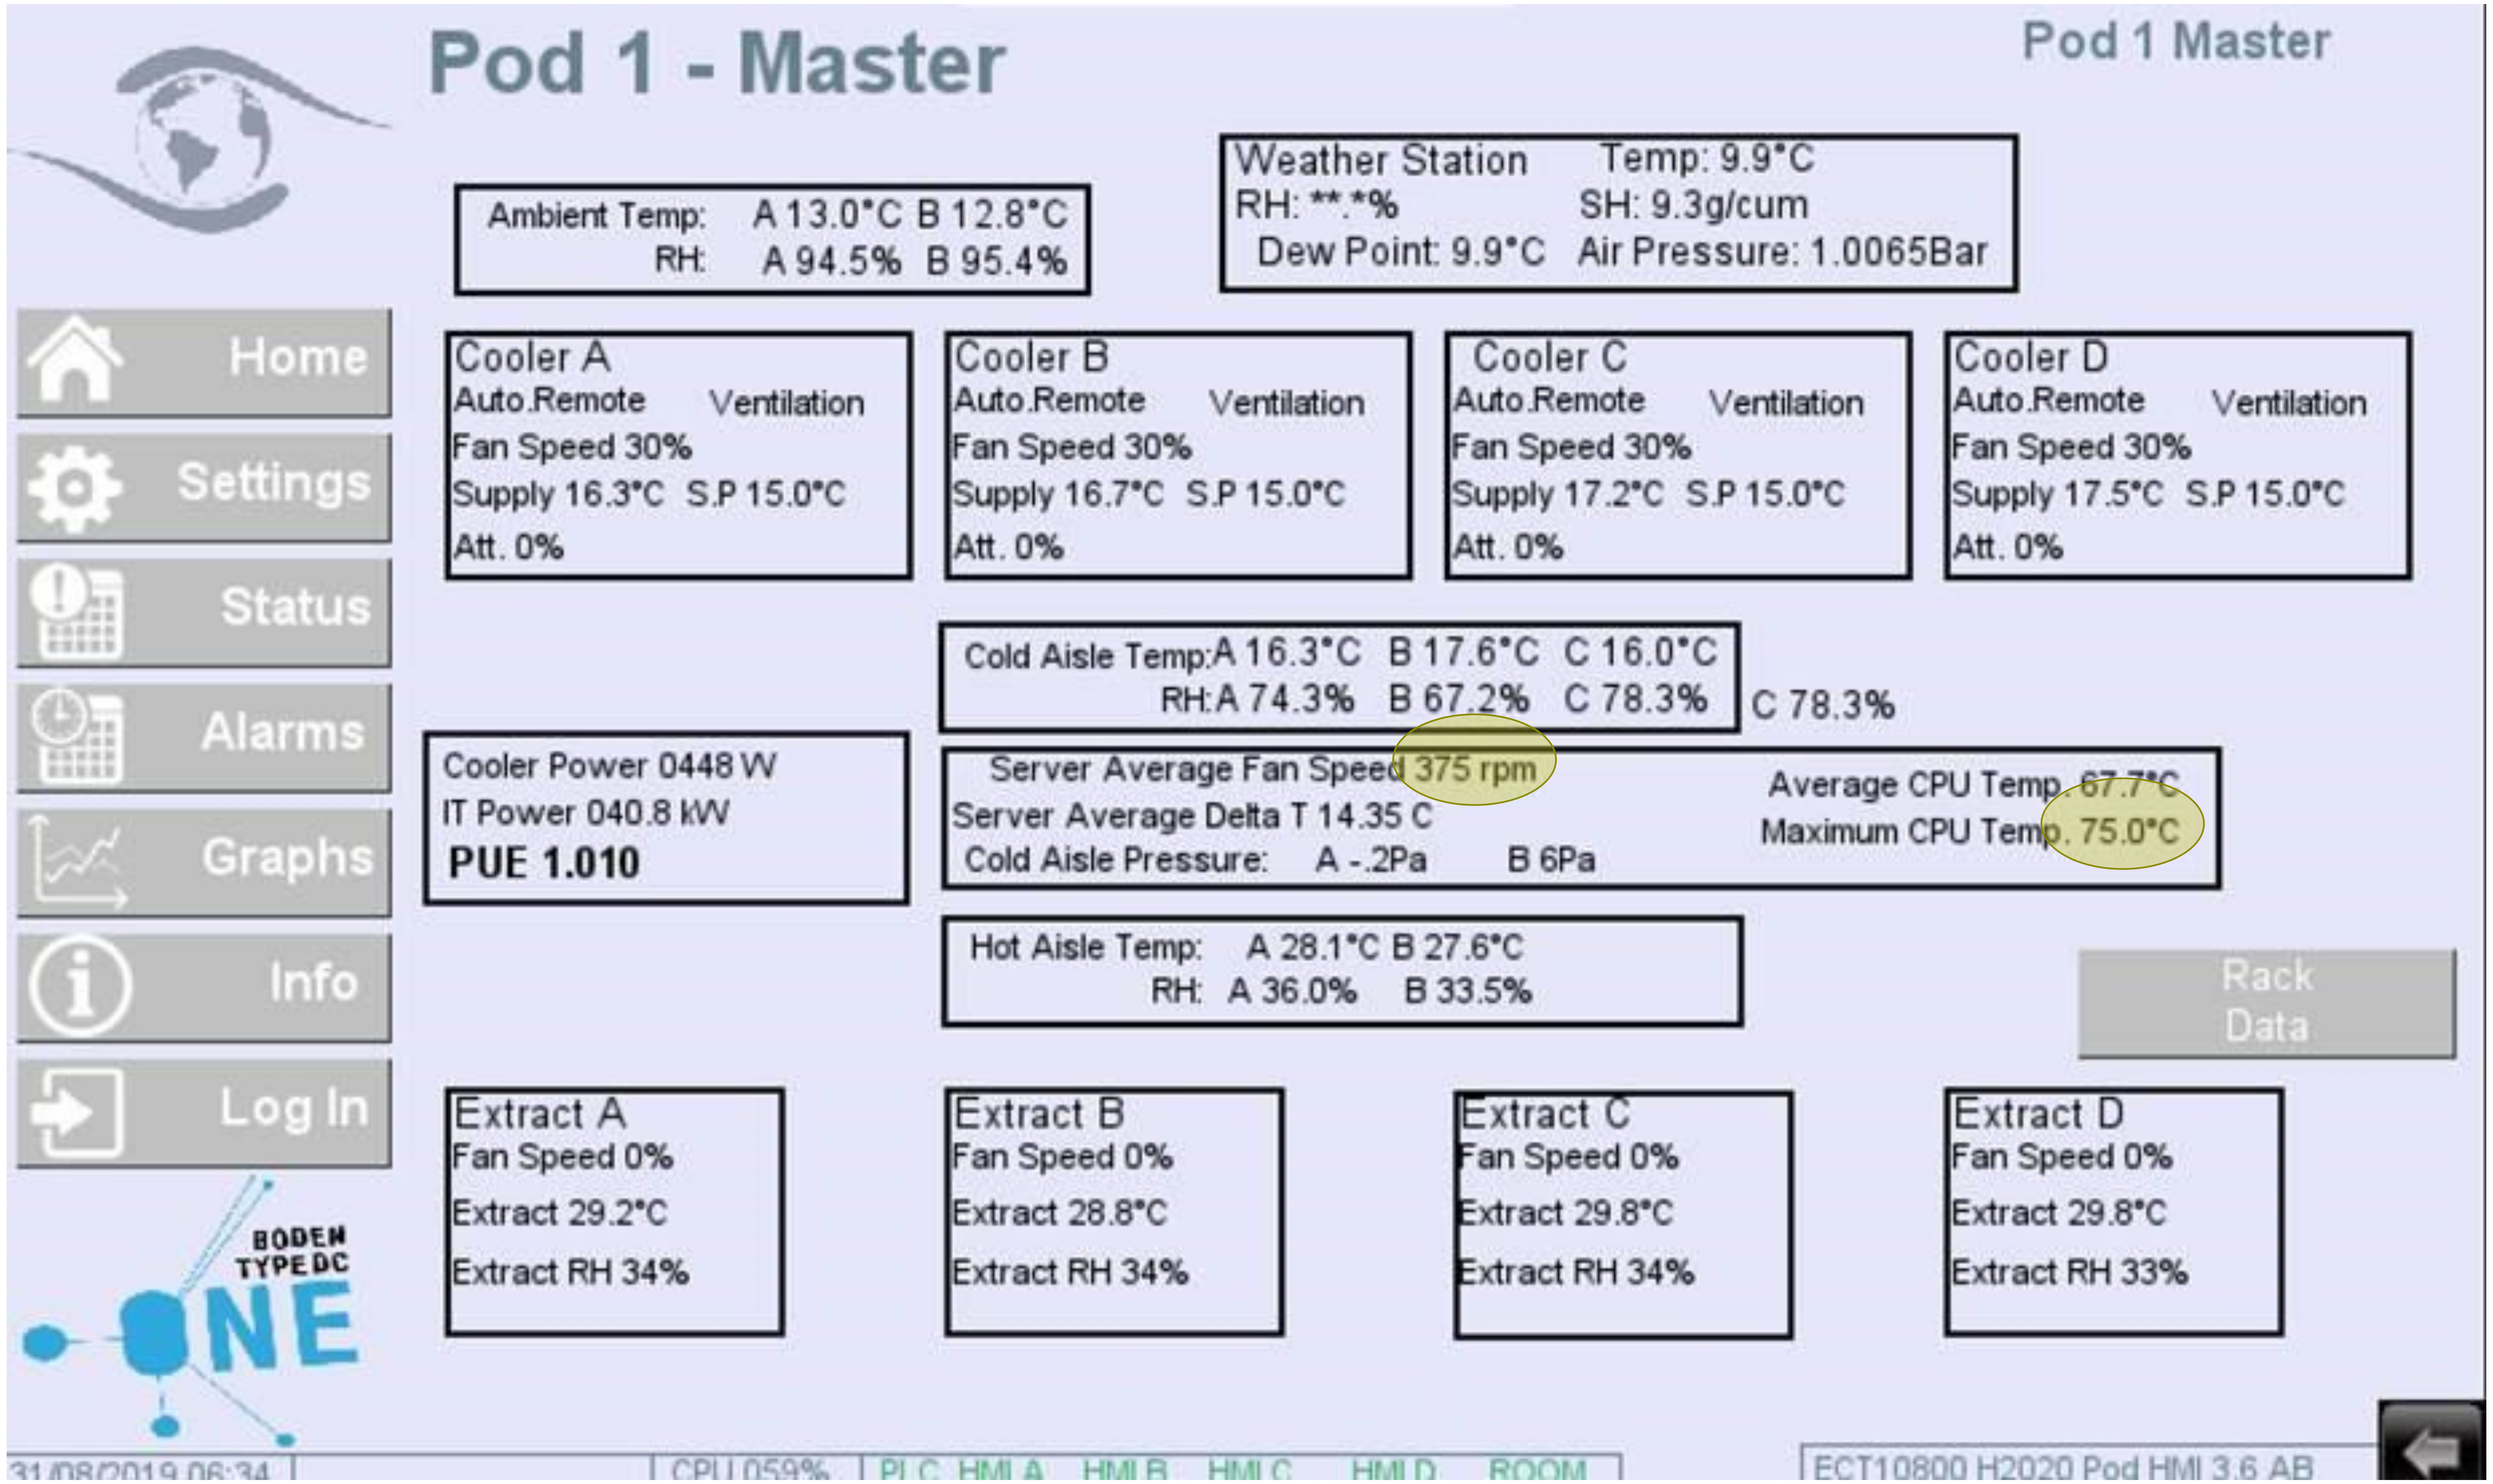

Phase 2 of the holistic cooling – cooling systems

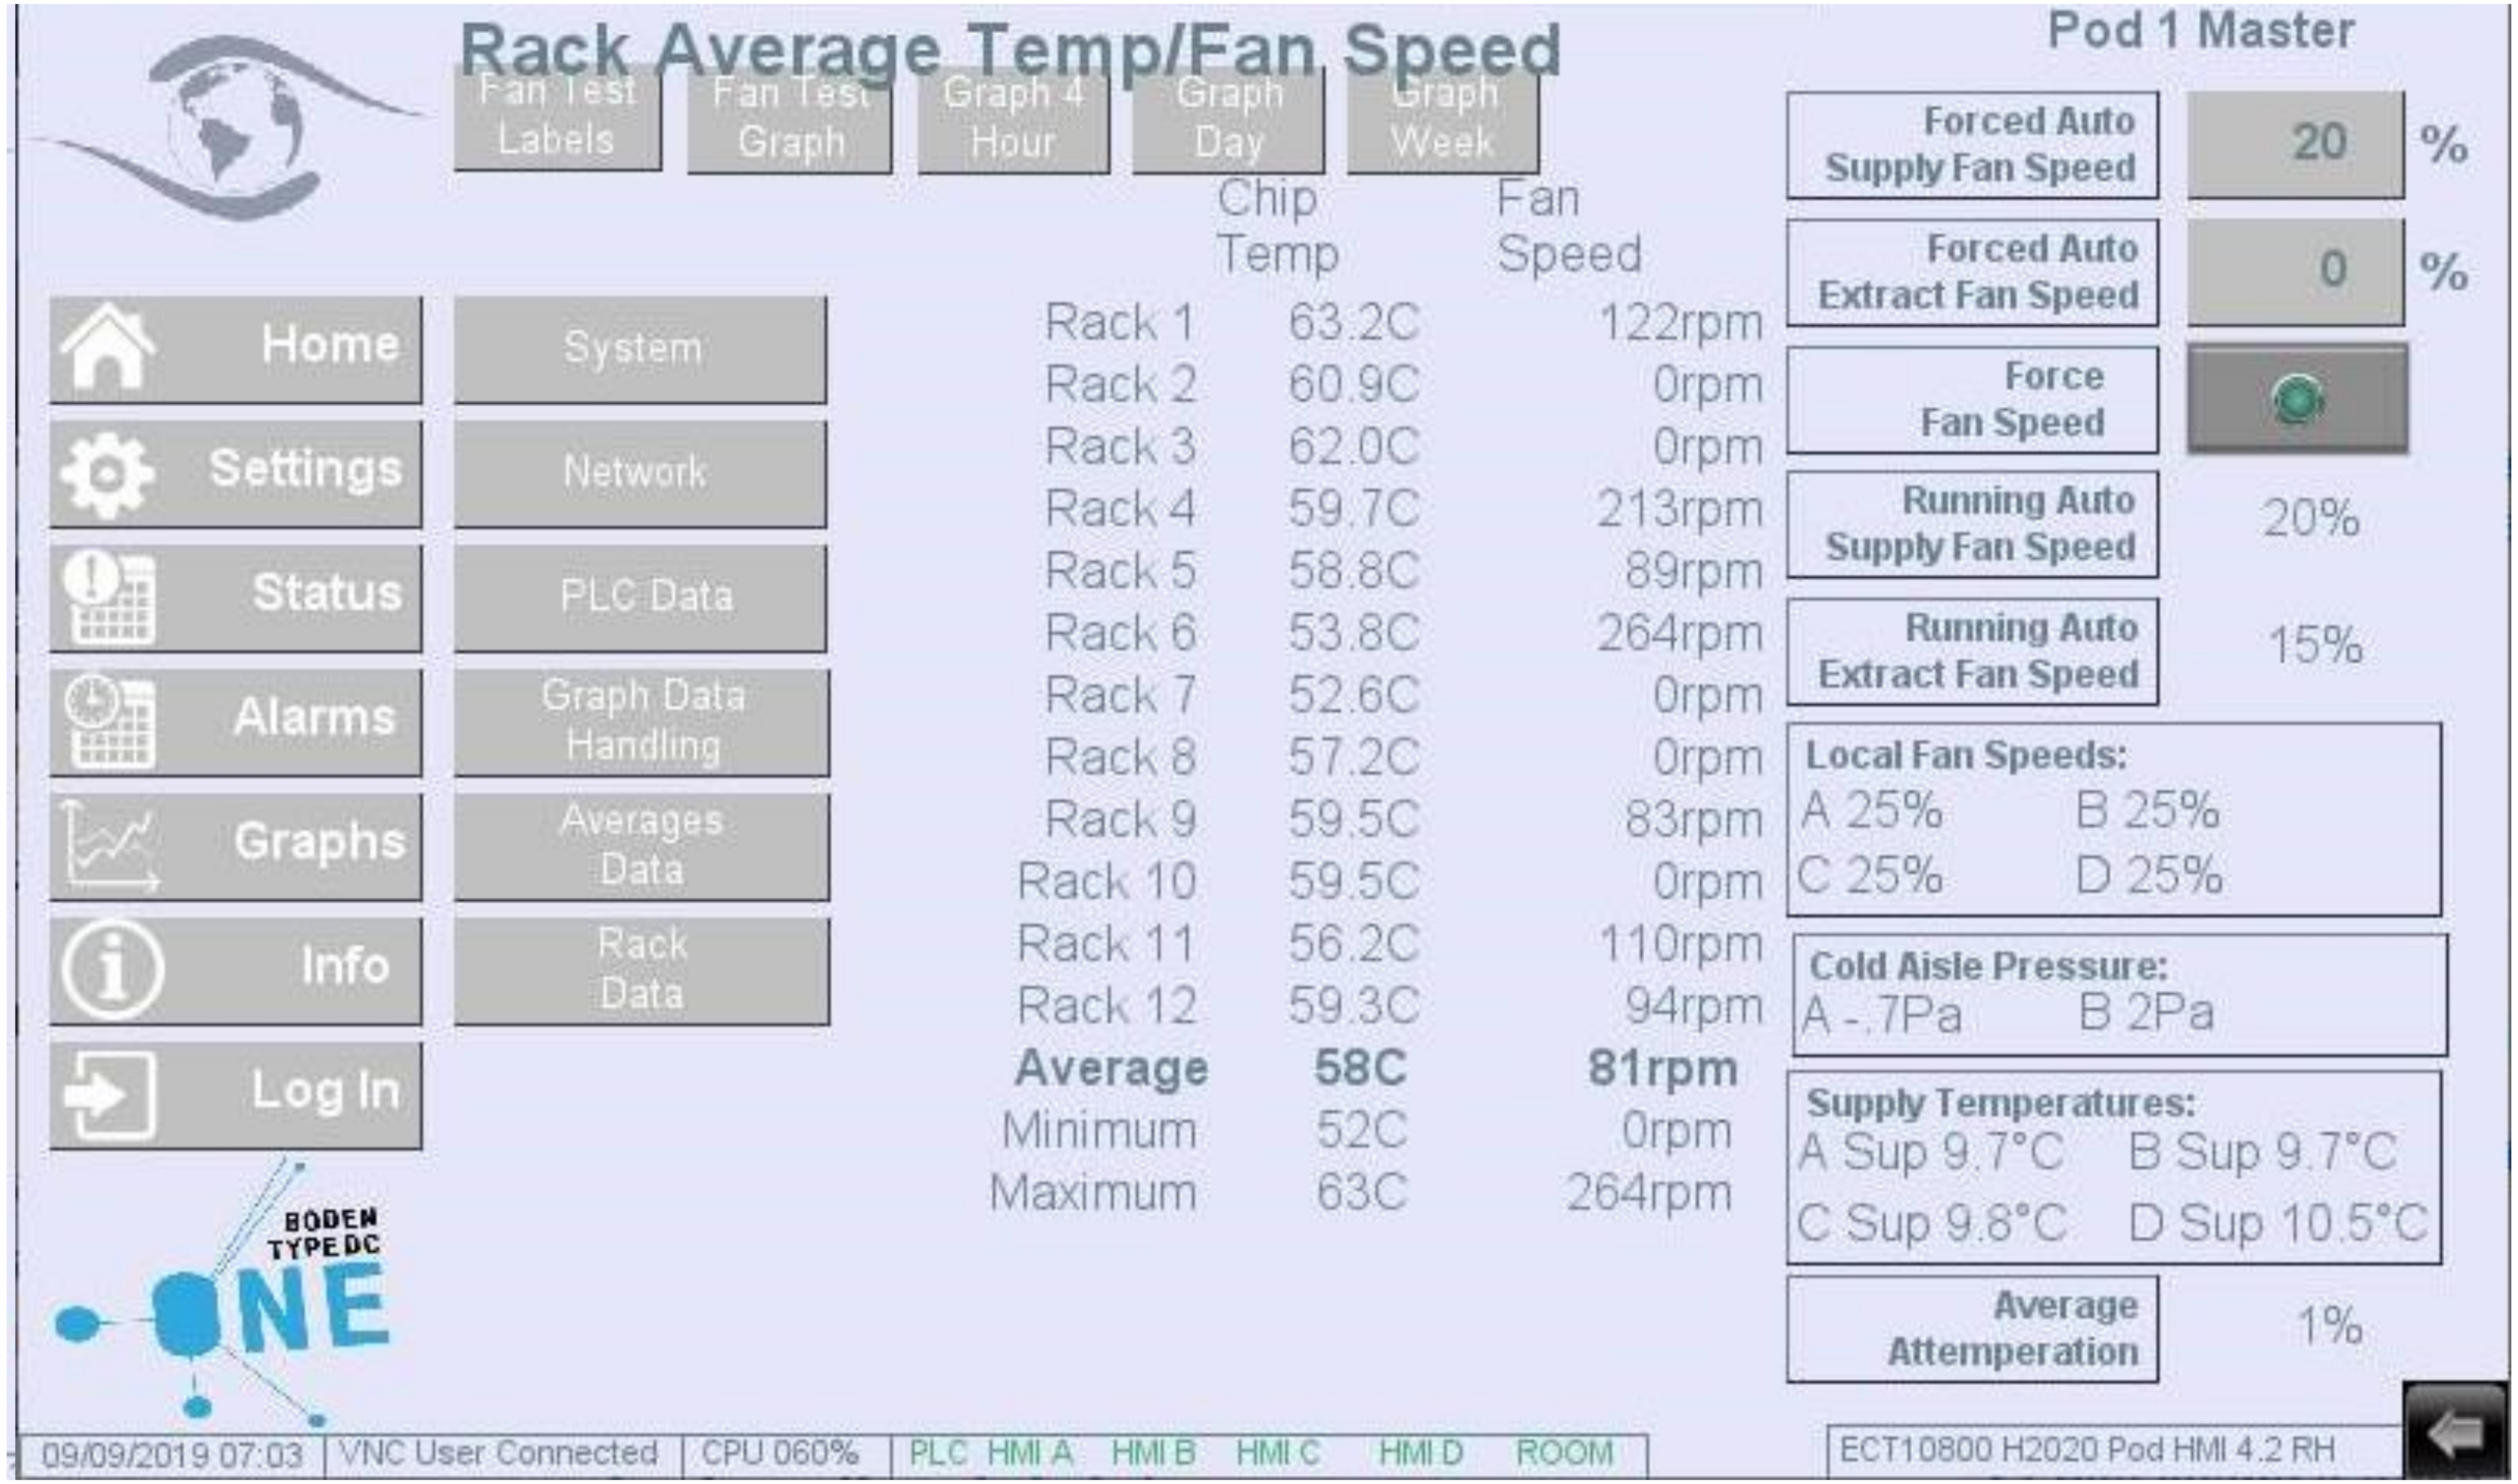

Phase 2 of the holistic cooling – cooling systems

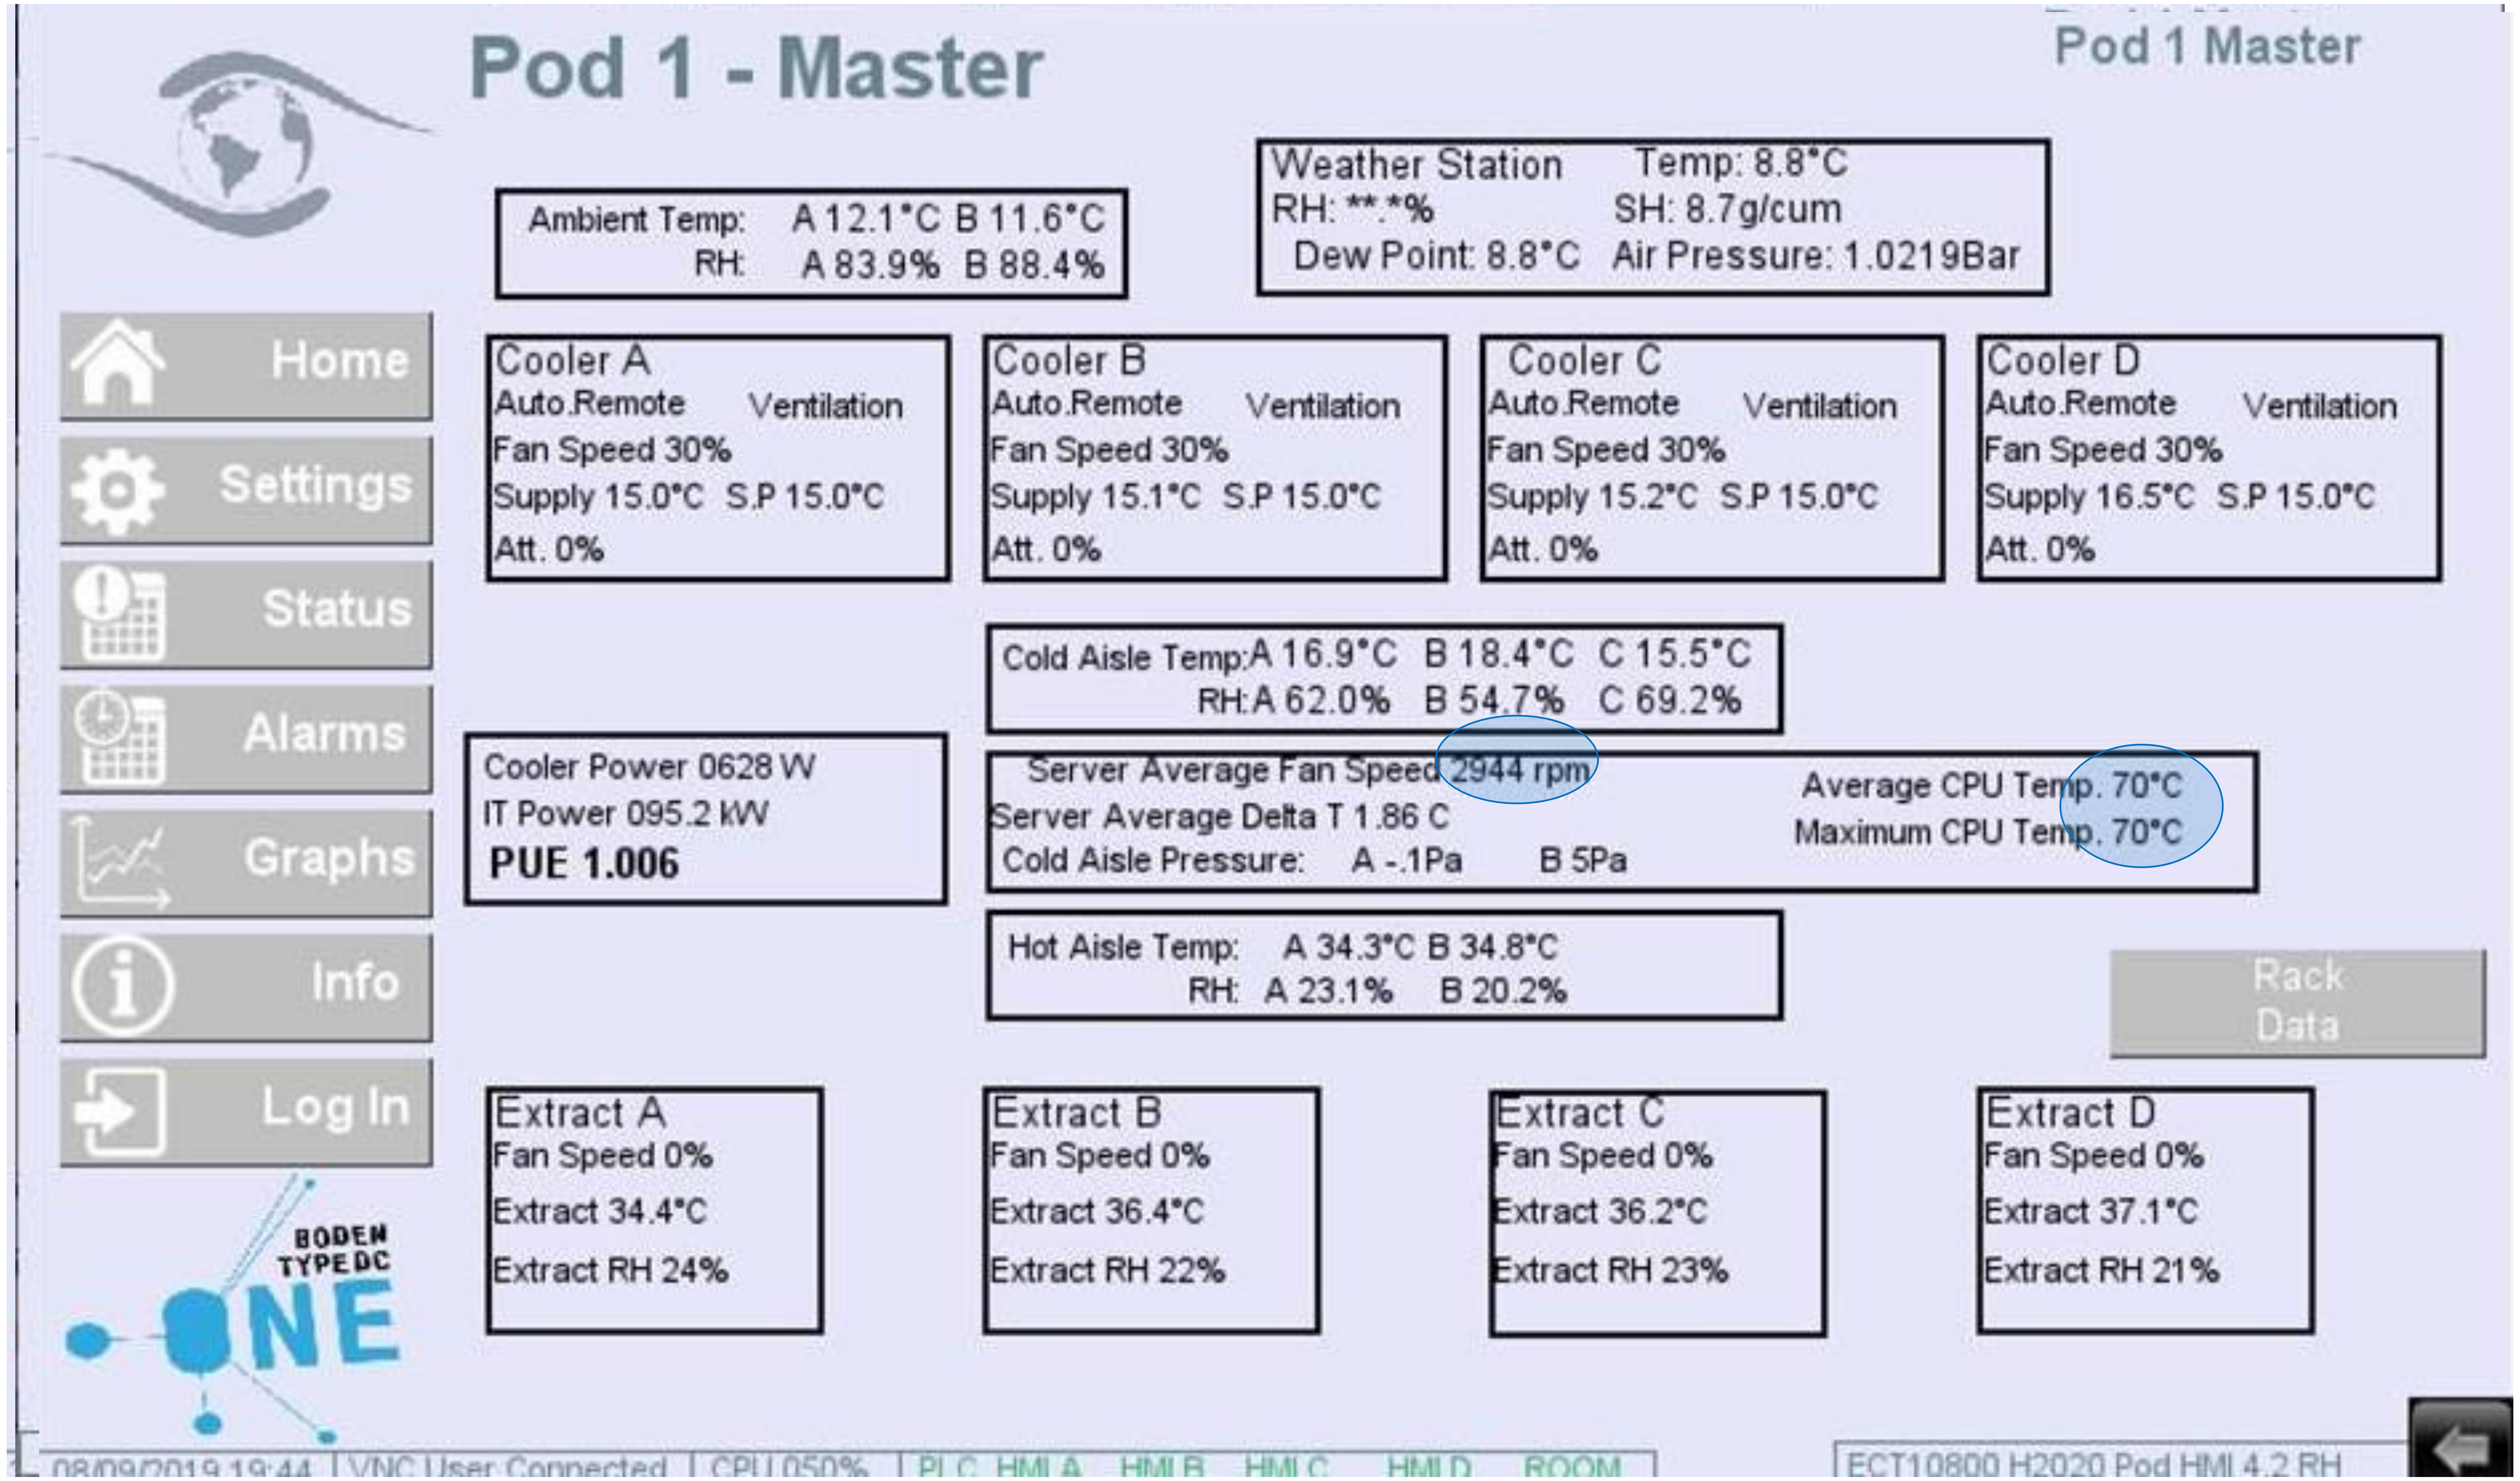

Live and interactive demonstration of the OCP POD1 operation

GA 768875

- Placeholder slide for around 5/10 minutes of demonstration. (Needs an interent connection!)
- Show different workload deployment scenarios – kubernetes versus loading OCP servers to their sweet spots for the environmental conditions at the time.
- Hope to show results of some Key Performance Metrics.
- Data will show power draw, temperature differences.

- RUN DEMO

Call to action / ask of the OCP community

- Should data center operators be trusted to take control of the server fans? YES/NO
- Could there be open access to extract server fan energy consumption as well as fan metrics and data.
- Could the Open System Firmware allow greater access for server fan data and include control for synchronised control of the airflows between the coolers and the servers.

SUMMARY

- The project operates and monitors an experimental, but realistic test data center.
- Deployed a chain of opensource software to monitor the data center, IT utilisation and deterministic control of workload.
- Created a 140kW IT footprint using legacy OCP Windmill servers, with power, network and environmental and workload control.
- Piloting a data center cooled by direct air with a low cost humidification and an ambitious linking of facility and server fans.
- RISE can characterise thermal and power envelopes of servers using a uniquely developed server wind tunnel.
- Open Research Data Pilot will be available from the end of March 2020.

Visit https://bodentypeDC.eu

- Thanks to RISE colleagues:
  Jonas Gustafsson
  Daniel Olsson
  Jeffrey Sarkinen
  Magnus Nilsson-Mäki
  Filip Blylod
  Sebastian Fredriksson
- This project has received funding from the European Union´s Horizon 2020 Research and Innovation programme under grant agreement No 768875.
  Thanks to the other BodenType DC Project partners
  H1 Systems
  Fraunhofer Institute of Optronics
  EcoCooling
  Boden Business Agency

GA 768875

- Thanks also to

Open. Together.

OCP Regional Summit

26–27, September, 2019