Prévessin Data Centre Powers Up

24 October 2024

CERN IT – Max DUPUIS

Outline

CERN Data centre in 2020

LHC Needs

Increase CERN Data centre capacity

PDC: Design

PDC: Trial Operation

PDC: Operations & Maintenance

PDC: 1st year of operation

PDC: Milestones

CERN Data centre in 2020

Available capacity 5 MW:

- 5M HEPScore [1] CPU processing
- 900 PB Tape Storage
- 320 PB Disk Storage

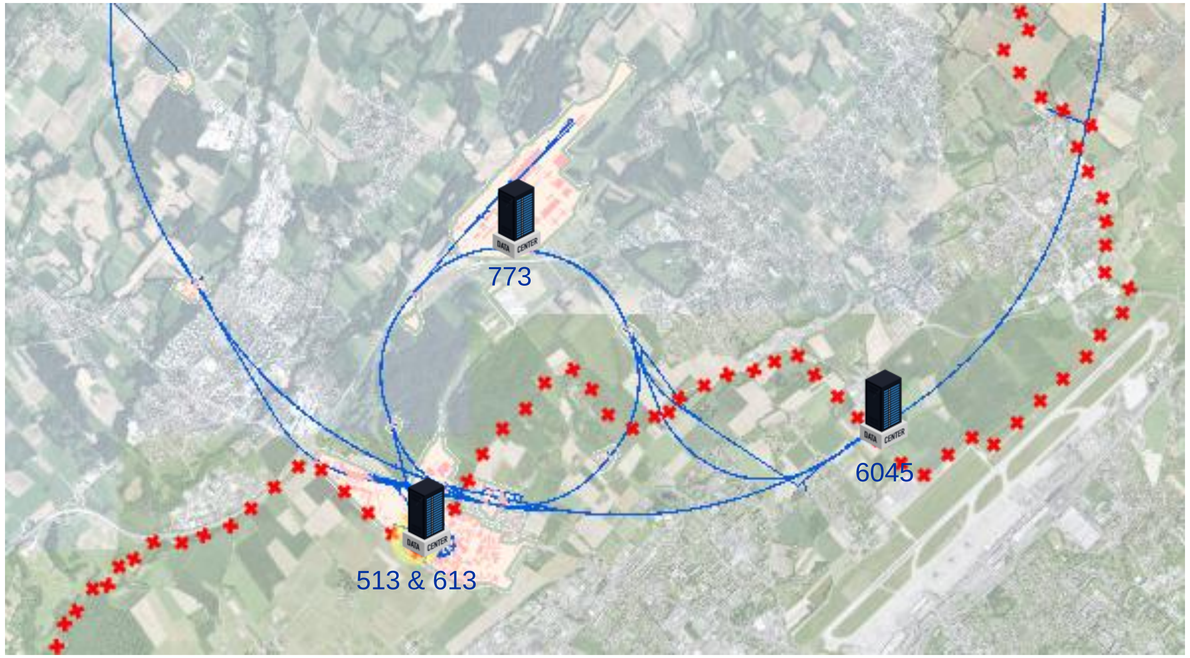

Meyrin Data Centre (MDC) – 1973

- 3.8MW for computing in 3 rooms, PUE ~1.5
- full UPS and partial diesel coverage, tape robots

B613 – 2001

- 100 kW for Backup tape robots
- full UPS and diesel coverage, tape robots

2nd Network Hub – 2016

- 100kW site network redundancy, PUE ~1.1
- full UPS and diesel coverage

LHCb containers – 2018

- 2 containers of 500kW for computing, PUE ~1.1
- No UPS or diesel coverage

LHC needs

The computing needs of CERN are largely driven by the LHC experiments (~85%)

Requirements evolve continuously

Each upgrade of the LHC and/or experiments leads to higher data rates and hence higher computing needs (RUN3 => RUN4 ~ factor of 10)

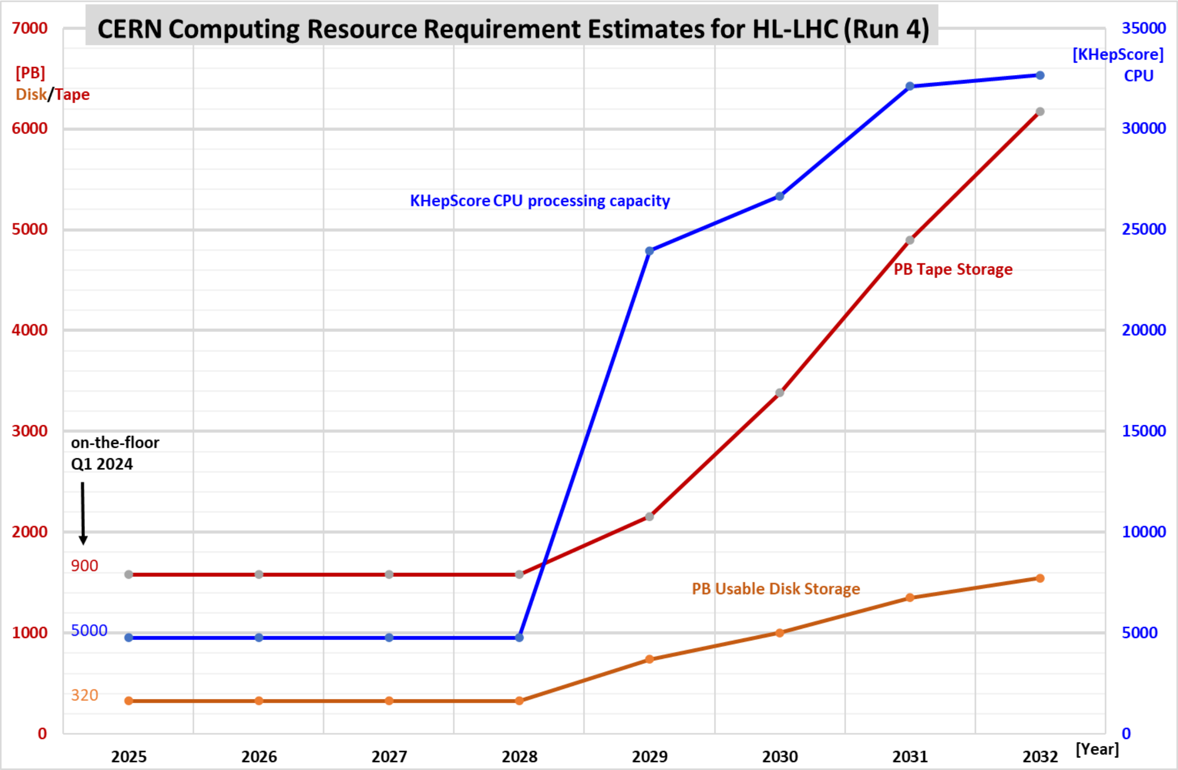

Increase CERN Data centre capacity

After several attempts the project has been approved at the FC in 2020. Mid 2021 contract has been signed with a consortium led by EQUANS.

Tender Solution Logic:

- Maintained the approach for the 2019 tender
- DC design is specialised and has been progressing quickly over the past years
- We wanted to benefit from industry expertise, therefore, a functional specification in the tender with minimal design prescriptions
- 10-year O&M to ensure good design

PDC: Design

Highly efficient data centre with a Power Usage Effectiveness (PUE) [2] ≤ 1,15

Construction in 1 phase

Installation will be in 3 phases over time for a total of 12MW

- **1st Phase**:   4MW - 2nd floor
- **2nd Phase**: +4MW - 1st floor
- **3rd Phase**: +4MW - ground floor

PDC: Design

Footprint 2250m2 (70mx37mx19m)

6400 m2 of which 2100 m2 for IT rooms

6 IT rooms of 2MW each, two per floor

Transformers and Low Voltage distribution, but High Voltage equipment’s are provided by CERN

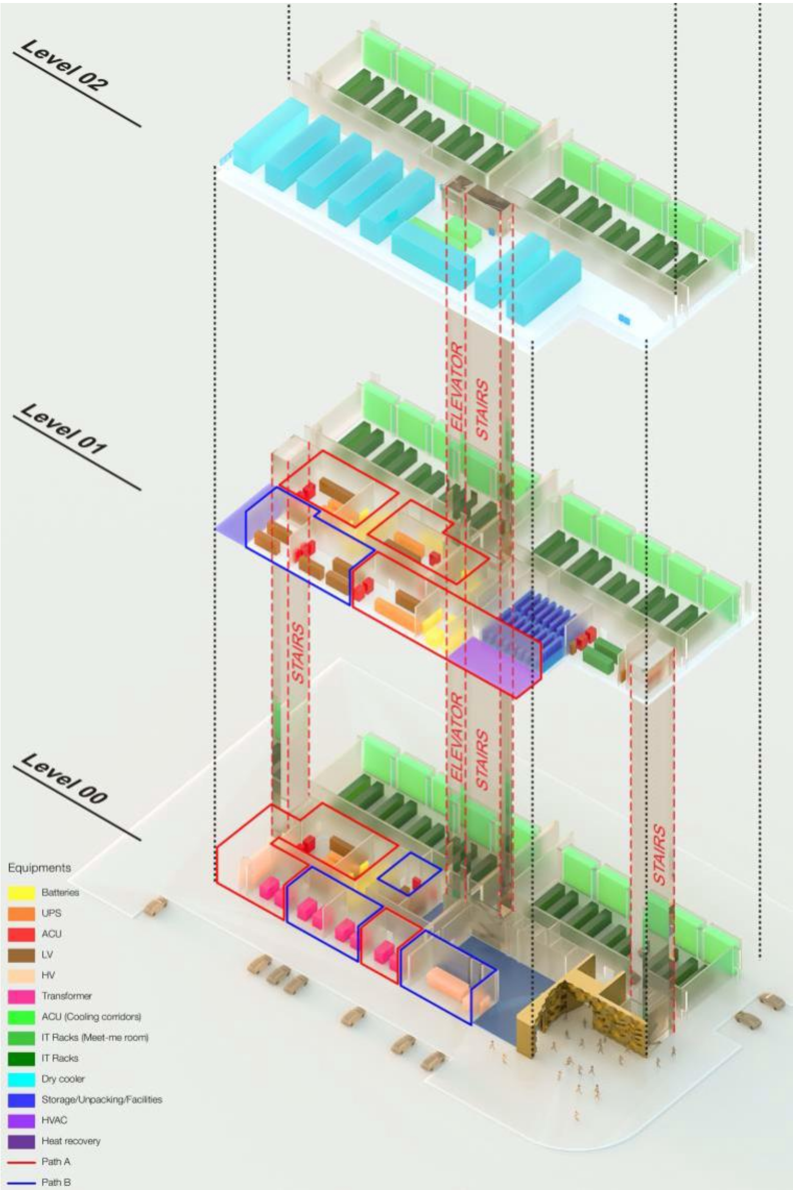

PDC: Design

4x 2MW Jaegi dry coolers in 1st Phase (2x2MW per additional phase) - N+1 redundancy plus 1 for technical areas

Adiabatic cooling from 20 degrees outside temperature with water recycling to improve Water Usage Effectiveness(WUE)[3]

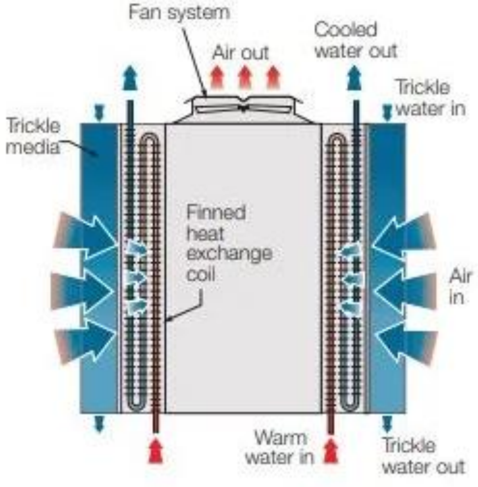

Hybrid cooler with trickle onto media

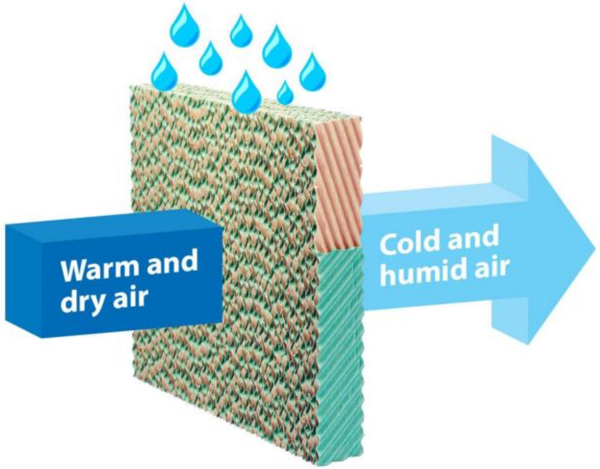

PDC: Design

5 x 500kW fan walls per room (Schneider largest in Europe), N+1 redundancy

ASHRAE [4] class A1, temp. environment up to 32 degrees

Heat recovery of 3 MW in 1st phase and 4 MW in 2nd phase

PDC: Design

90 Racks per room

- 54x 800mm wide, Non-UPS racks up to 25kW (physics data processing)
- 24x 600mm wide, UPS racks up to 12kW (business continuity)
- 12x 800mm wide, 50% UPS racks (network pod)

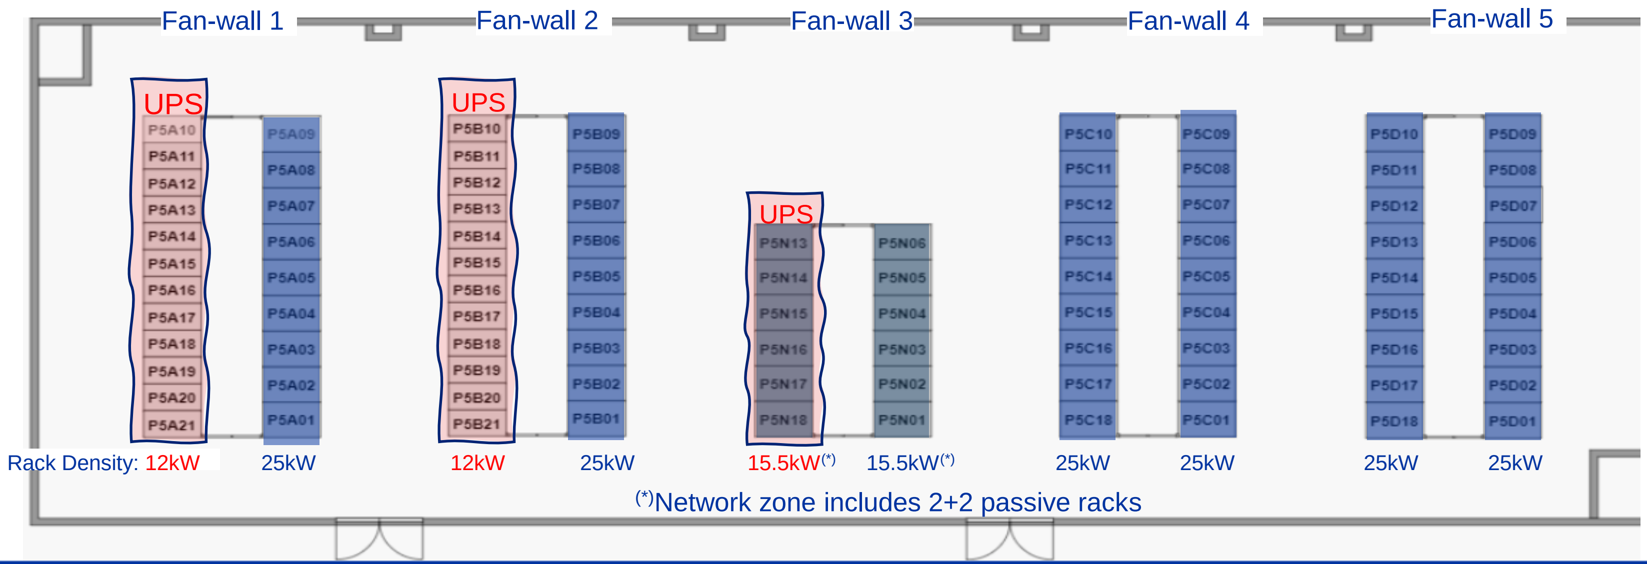

PDC: Design

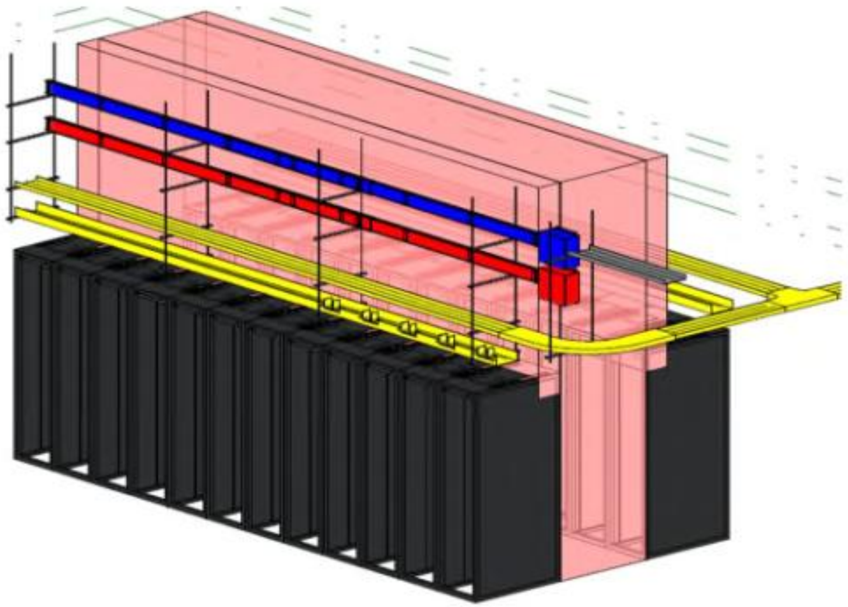

Air-cooled racks with hot-aisle containment

Two redundant power feeds: red and blue

- Red feed: 20% UPS coverage (Business Continuity)

Two type of Power Distribution Unit (PDUs):

- 63Amps (physics data processing)
- 32 Amps (business continuity)

PDC: Design

Computational Fluid Dynamics (CFD) has been used to determine the optimal IT room layout and to evaluate various failure scenarios.

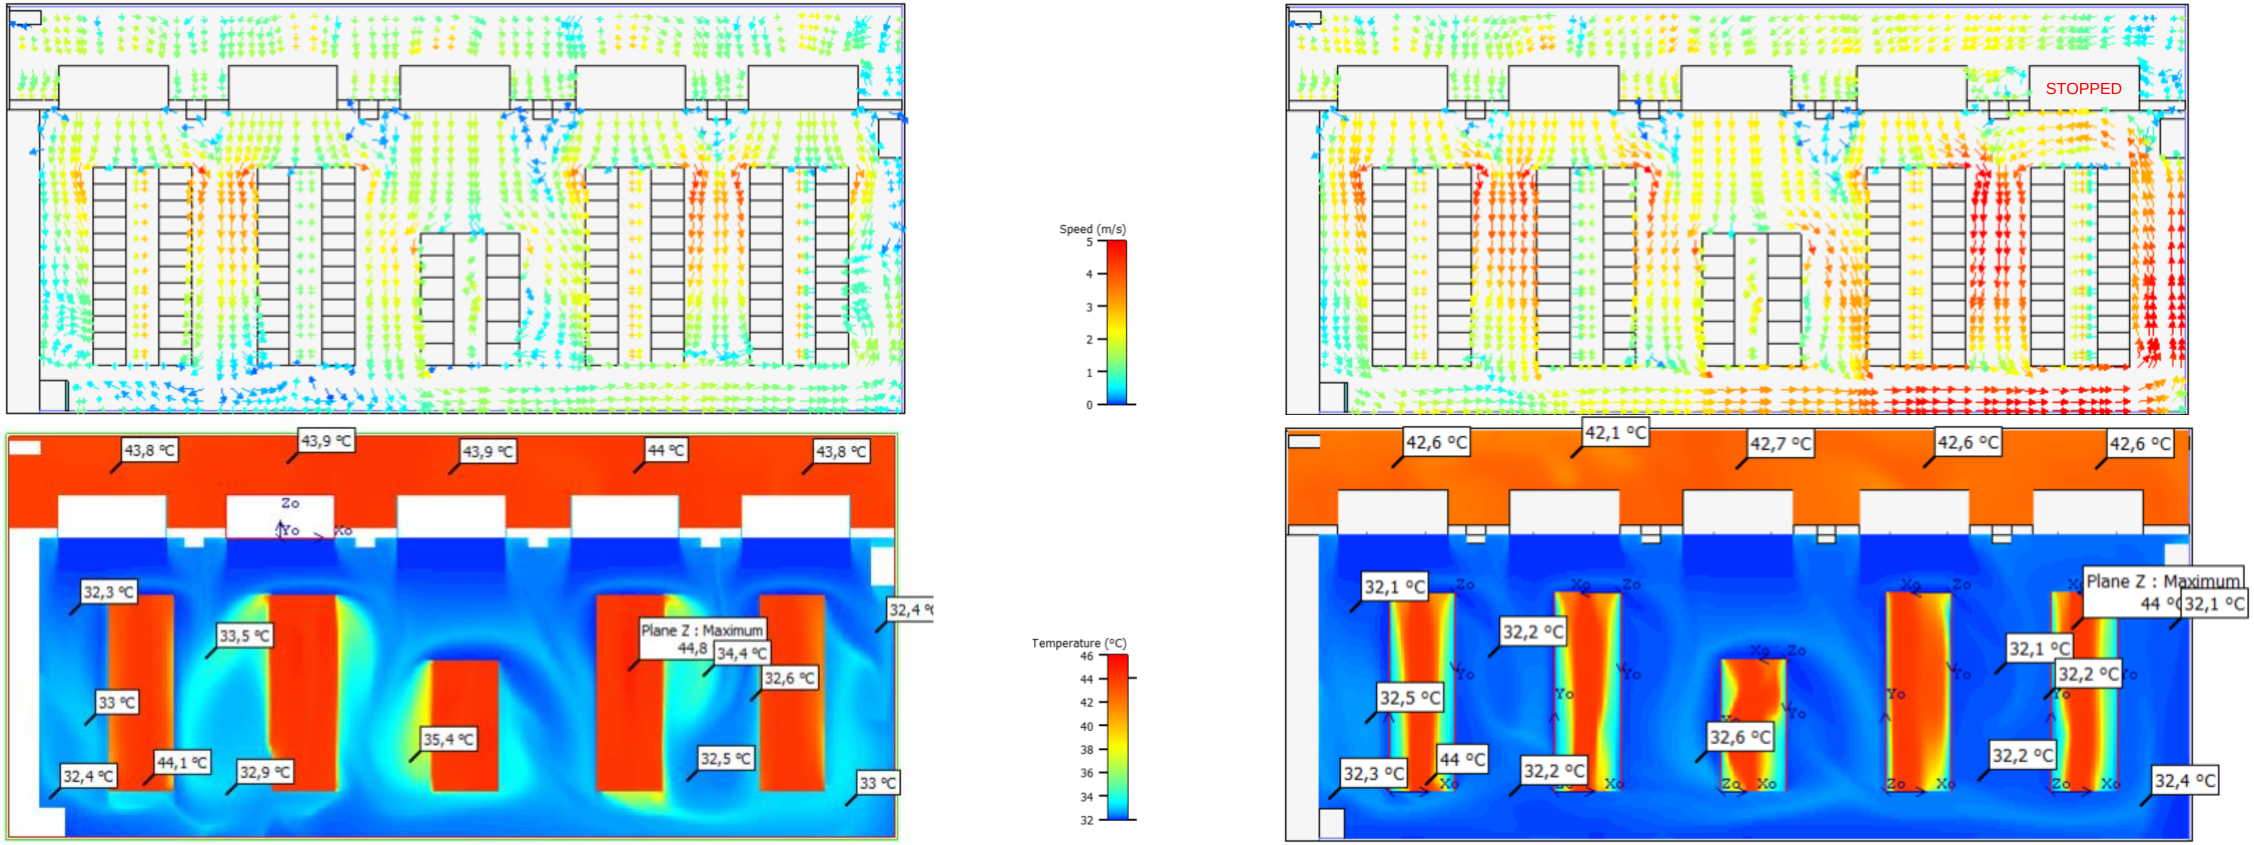

PDC: Trial Operation

At the end of the commissioning of the 1st phase, we run a test with servers with the following technical scope:

- Build a full non-UPS POD with batch servers:
  - 576 nodes (144 Quads) for a total of 432 kW
- All other racks sealed with front-panels
- Run standard burn-in for one week

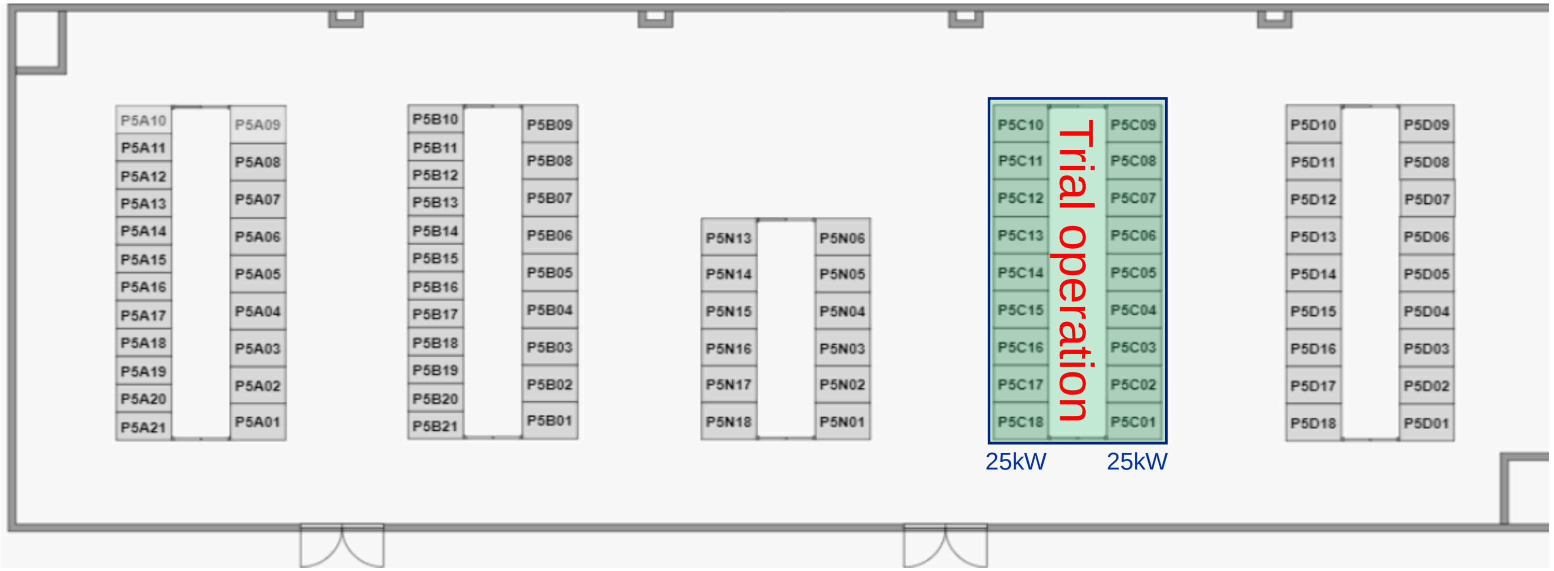

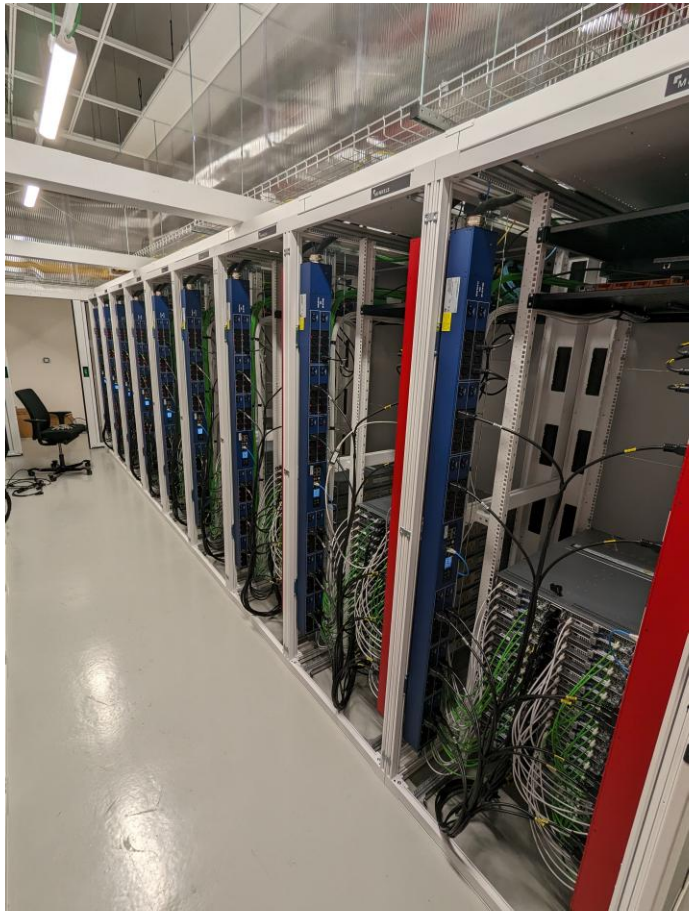

PDC: Trial Operation

Several scenarios have been tested:

- Verify operation in nominal mode
- Verify redundancy in power topology - cut a feed and verify
- Verify redundancy in cooling - cut 1 fanwall and verify
- Verify redundancy in cooling - cut 2 pumps of dry cooler and verify
- Verify operation on UPS – cut all fanwalls and let servers running for 5 min

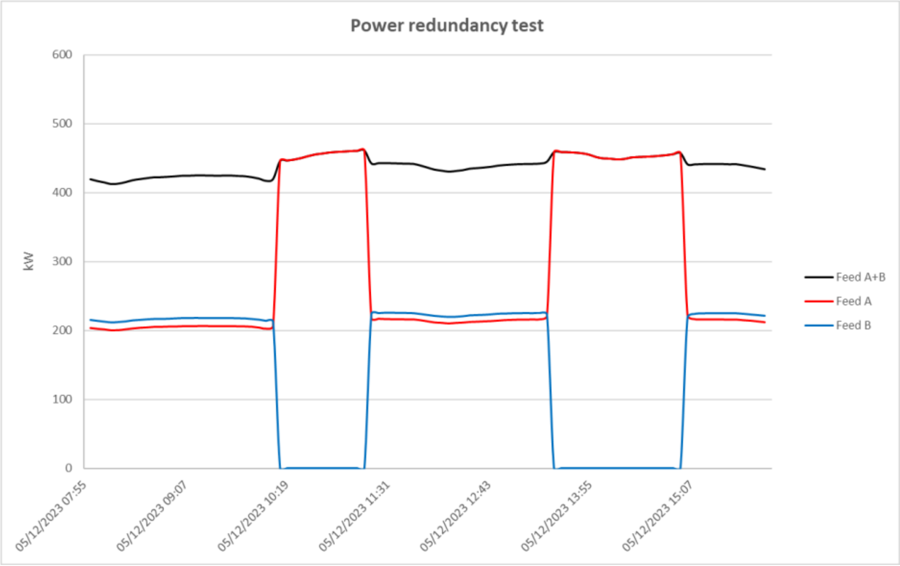

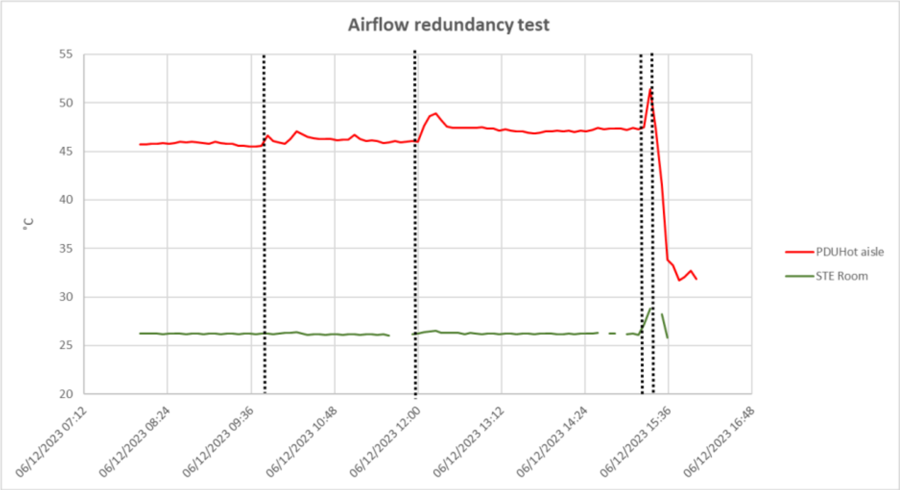

PDC: Operations & Maintenance

O&M managed by EQUANS has started once the commissioning + trial operation completed successfully in February 2024.

Current contract runs until 2034.

From CERN perspective, the contractor:

- Fulfills roles of CERN technical groups for the infrastructure
- Acts on PDUs failures triggered by us (PDUs are monitored + managed by CERN IT using standard tools)
- Provides access to the Building Management System (PCVue) for detailed internal views of whole infrastructure

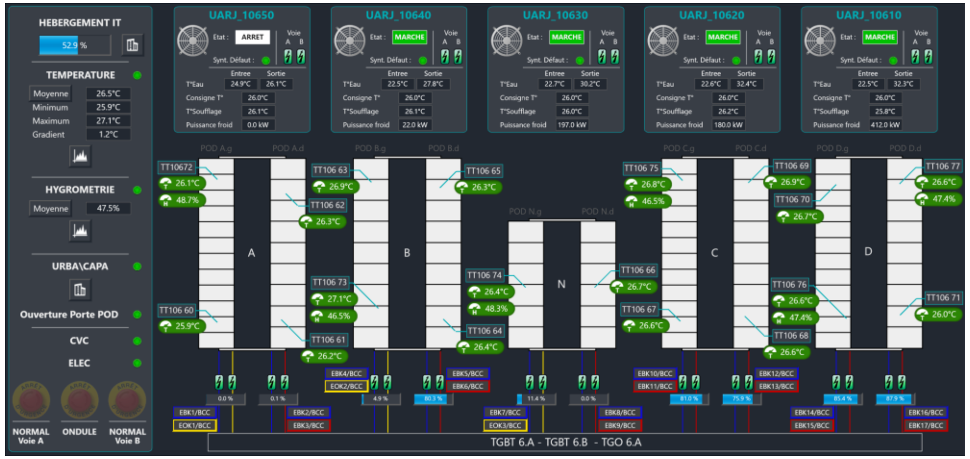

Monthly contract meetings

- Review issues and performance (PUE, WUE)

PDC: Operations & Maintenance

For O&M at rack level and Electrical distribution, we are using the same model as Meyrin Data Centre (MDC):

- Since Q1 2024, CERN IT is operating remote computing facilities effectively without the CC Operator providing on site corrective action
- DC resilience at a rack level has reached a point that doesn't require immediate CC Operator intervention to maintain service continuity
- Platform and Applications are designed with a high degree of fault tolerance and can withstand a single point of failure
- Critical Function has been reassigned to CERN Control Room (SSB entry following DC infrastructure incident, call IT piquet, monitor and handle network alarms)

PDC: 1st year of operation

Since the beginning of 2024, the two rooms of the 1st Phase have been operational. In 6 months, we successfully commissioned the PDC at 50% capacity

NEXT: In preparation to Run4, we will need to initiate the installation of the 2nd Phase of the PDC and begin procuring additional servers

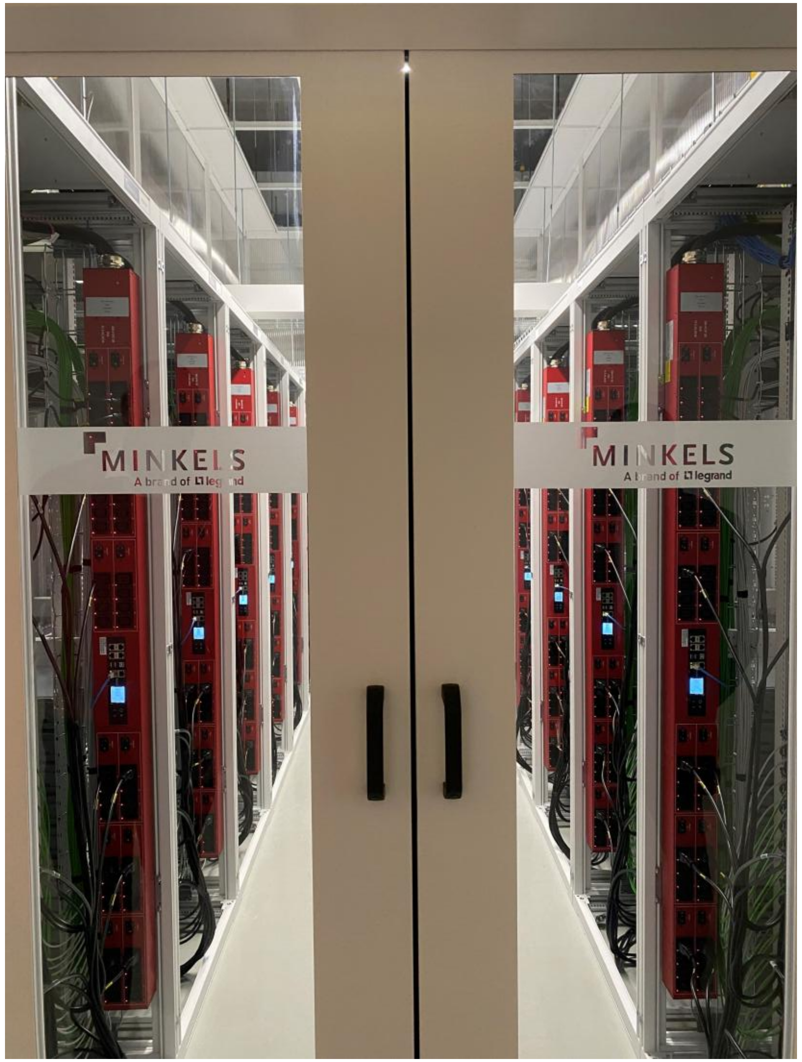

PDC: Milestones

- **2008** – 1st Tender with 4 concept design
- **2009**
- **2010**
- **2011** – 2nd Tender Remote hosting capacity
- **2012**
- **2013** – Wigner (HU) - 2.5MW – 7Y Ops
- **2014**
- **2015**
- **2016**
- **2017** – 3rd Tender
- **2018** – LHCb Containers - 1MW – 7Y Ops
- **2019** – 4th Tender
- **2020**
  - Q4 - Approved at the FC
- **2021**
  - Q2 - Contract Signed
  - Q3 - Building Permit
- **2022**
  - Q2 - Groundbreaking Ceremony
- **2023**
  - Q4 - Commissioning of the Data Centre
- **2024**
  - Q1 - O&M Began
  - Q4 – 50% Capacity reached

Thank you!

Questions?

home.cern

References

[1] **HEPScore** is a new CPU benchmark created to replace the HEPSPEC06 benchmark that is currently used by the WLCG for procurement, computing resource pledges and performance studies. HEPScore A new CPU benchmark for the WLCG

[2] **Power Usage Effectiveness** (PUE) metric, a tool designed to help boost energy efficiency in data center operations. PUE A Powering Change Across the ICT Industry Infographic

[3] **Water Usage Effectiveness** (WUE) metric to address water usage in data centers, which is emerging as extremely important in the design, location, and operation of data centers in the future. WUE A Green Grid Data Center Sustainability Metric

[4] **The American Society of Heating, Refrigerating and Air-Conditioning Engineers** (ASHRAE) is an American professional association seeking to advance heating, ventilation, air conditioning and refrigeration (HVAC&R) systems design and construction. Equipment Thermal Guidelines for Data Processing Environments