CERN-IT Sustainability

markus.schulz@cern.ch

CERN IT-GOV-GRC
Governance, Risk and Compliance Section

CERN's approach to the environment

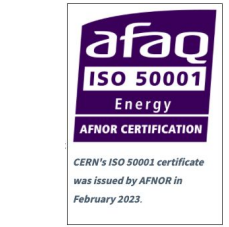

Energy Certification 2023

- CERN wide initiative to minimise energy use piloted by the CERN Energy Coordinator
  - **ISO 50001** certification obtained 2023

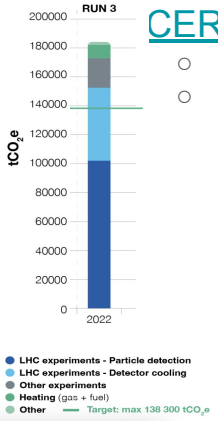

CERN and the Environment Town Hall is a good starting point!

- Scope 1,2 and 3 emissions and plans for improvements have been covered in detail
- Current goals for 2021-2025 and the revised goals for 2025-2030
  - **High Priority Goals for Horizon 2030**
    - **Scope 1:** Direct CO_2e emissions
      - -50% from 2018 (currently ~ 140-184000 tCO_2e)
    - **Scope 2:** Electricity
      - 1.5 TWh/year (+14% for 5-7.5 times the collisions)
      - Now: ~ 1.3 TWh/year → ~ 65000 tCO_2e
    - Water: 3651 ML (+5% 2018)
  - **Scope 3:** travel, commuting, catering, waste,....., procurement
    - 9000 tCO_2e + 105000 tCO_2e
  - **Total: ~320-364000 tCO_2e**

Mostly fluorinated gases (cooling and detectors)

French power mix:
53g CO_2e/kWh

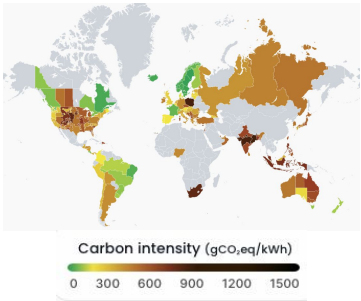

- Detailed Environment Reports compiled by HSE
  - Every 2 years updated
- IT's contribution is about 4%

IT's impact on the environment

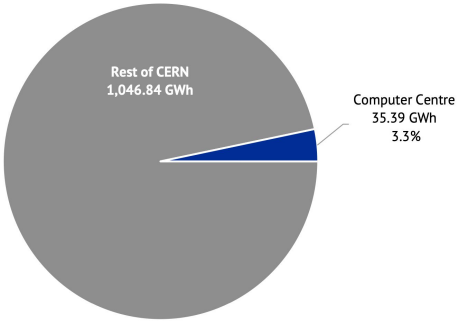

- Scope 1: Almost no direct emissions
- Scope 2: Electricity
  - We receive every year a "Virtual Electricity Invoice"
  - 2023: **35.39 GWh (2.71 MEuro)** up from 1.16MEuro in 2022
  - Very stable: 2.6-3.3 GWh/month
    - **Very high, constant workload**
      - **Almost all resources!**
  - **LHC-ON ~3.5% LHC-OFF ~9%**
  - Corresponding CO_2e: **2275 tCO_2e/year**
- Fine grain tracking of electricity use:
  - openDCIM Data Center Inventory (cern.ch)
  - Web Energy 2.0 (cern.ch)
  - Data Center Infrastructure Monitoring - HomePage - Dashboa
  - Zones - Dashboards - Grafana (cern.ch)
  - PDC (0775) Power - PDC (0775) - Dashboards - Grafana (cern.c

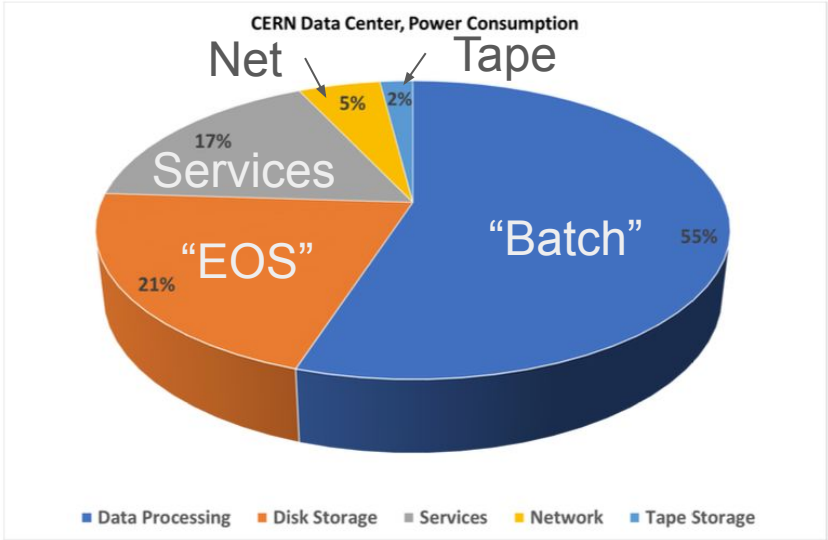

IT’s impact on the environment

- Scope 3: Purchases, travel, commute, waste and…..
  - Difficult to give a reliable estimate for IT equipment
    - Buying new IT hardware (CPU, Disk, Network, Storage)
      - There are numbers, but precision isn’t great Boavizta
      - Time of usage varies (+ we donate hardware for a second life)
    - Building new computer centres + reworking existing infrastructure
      - decades of amortisation…not taken into account
      - Transformers, UPS, cooling…
      - Detailed plan for transforming the Meyrin Data Centre
        - Will also improve PUE and reduce water usage!
  - Bernd Panzer did a rough estimate, based on a **5 year life** cycle of equipment
    - Based on this: 3300-4400 tCO_2e/year
    - Travel, commute etc. CERN has also goals to reduce these
      - ~300 people in IT, 300 twice a year to NYC ~540tCO_2e/year (high upper bound)

IT’s impact: Summing it up

- Scope 1: 0 , Scope 2: 2275 tCO_2e/year , Scope 3: 3570 - 4670 tCO_2e/year
- Conservative estimate: **6000 - 7000 tCO_2e/year**
- CERN Total: **~320000-364000 tCO_2e**
- IT contributes **less than 2.2% to CERN’s total impact**
- However…
  - The data LHC produces is responsible for far more CO_2
  - Every part of CERN has to contribute to minimise the environmental impact
- Then there is WLCG with a significant impact
  - **The T0 is only 20% of WLCG’s computing**
  - WLCG started some time ago to look at this
  - A holistic study of WLCG energy needs for the LHC scientific program
  - Plus many contributions by members to conferences and workshops

What does IT do?

- We build a new more efficient data centre: Prevessin Data Center (PDC)
  - PUE 1.1 vs 1.46 , 4 MW, later 8 MW
  - Hosting energy hungry services like Batch
  - From 2027 on: reusing heat
- Meyrin Data Centre (MDC) 3.8 MW
  - With 100% UPS + diesel
  - Will focus on Storage and critical services
- Planning the future of the MDC
  - Detailed document by the IT Fabric Group (no link :-/)
  - Improvements in cooling, UPS, transformers…
    - Less energy used
    - Better PUE
    - While extending capacity for HL-LHC needs

- Information on CO2 is part of the purchase process
- Extending the lifetime of equipment by longer use and donations
  - After about 5 years the CO_2 e from usage is equivalent of the CO_{2e} needed for production
- Offering additional (energy efficient) architectures (ARM ~2000 cores, GPUs)
- ML platform: ml.cern.ch for efficient training of models

What does IT do?

- Monitoring and understanding the loads on our resources, examples:
  - CERN wide energy monitoring: Energy Tracking
  - Monitoring energy consumption of IT equipment openDCIM Data Center Inventory
  - Web Energy 2.0 (cern.ch)
  - CO2 monitoring
  - PDC (0775) Power
- Understanding utilisation of our services
  - Fine grained utilisation and efficiency of the batch systems (about 80% efficiency)
  - Monitoring that looks at the energy efficiency via HEPScore23
    - Including looking at efficiency differences on WLCG via HammerCloud
    - Extending this work towards GPU based workloads
      - CMS use of GPU based MadGraph → CHEP24 Talk
      - New CMS tracking code based on the former Patatrack
      - FlashSim (CMS ML based fast simulation) → CHEP24 Talk
- There is no lack of data that is being collected
    - More analysis is needed
    - Actions: Overcommitting in batch, LHC@home on services, BEER (Batch on Storage)

Understanding HEP Workloads

- Detailed studies of power consumption of different machines and conditions
  - Domenic Giordano, Natalia Szczepanek, Kacper Kozik
  - Using HEPScore with additional functionality
  - HEP Benchmark Suite:..
  - Power Consumption Time Series:...
  - ……

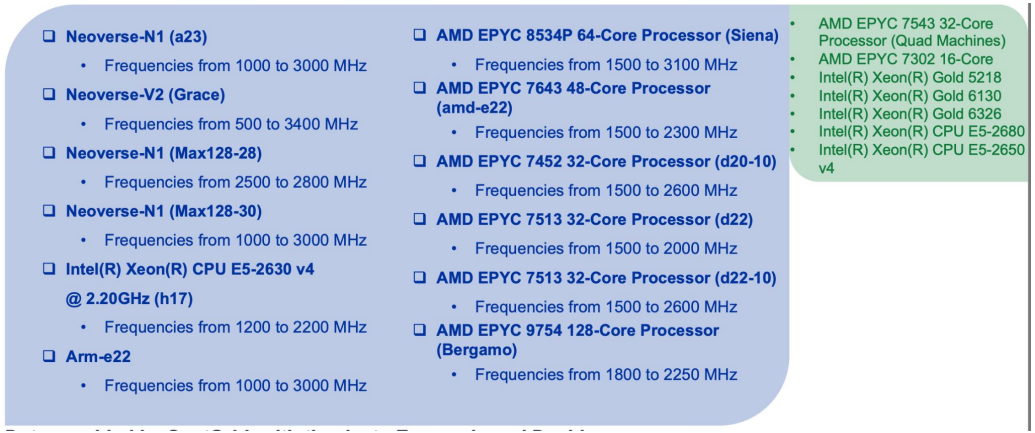

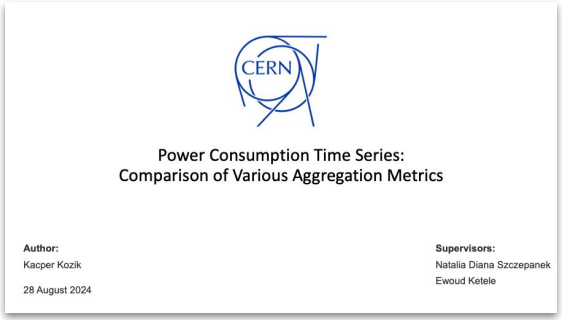

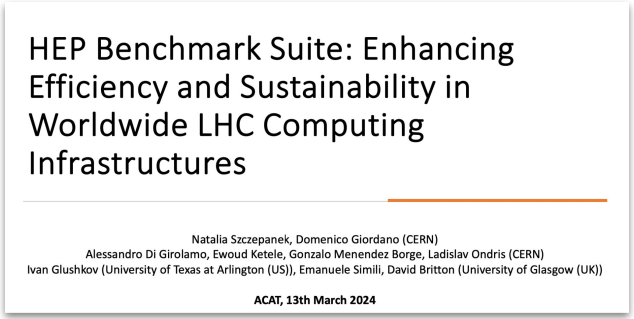

What does IT do? Projects

- Contributing to porting MadGraph5 to GPUs and vectorising the code for CPUs
  - First release for LO, NLO work is ongoing (more details tomorrow)
  - Developing new profiling tools
- Working with CERN SFT on optimising I/O efficiency for RNtuple
- Many Projects linked to openlab
  - NextGen Heterogeneous Computing, Event Generators, IT-ML-Service..
    - Software efficiency, energy efficient hardware, scheduling based on task…
  - KEPLER (Kubernetes Efficient Power Level Exporter)
  - Proposal to the CERN Environmental Steering Board Committee
- Started to link with the team providing electricity at CERN
  - What can IT do during the 15 critical days to help?
- In contact with several external projects and actors:
  - RF2.0 Includes activities related to computing, power …(more later)
  - Looking at activities that assess and certify the level of maturity in matters of sustainability
    - GreenDiSC
    - GreenDIGIT (EGI is a participant)
    - Swiss Datacenter Efficiency Association SDEA
  - Other communities, like HPC should follow soon
- Raising awareness within the teams (Climate Fresk)

My first impressions (active since October 24)

- The department is well aware of the importance of becoming more energy efficient
  - For sustainability and cost reasons
- We know our impact and efficiencies quite well (monitoring, assessments..)
  - Some updates needed (Bernd's detailed paper is from 2022)
  - Some aggregation of information needed (monitoring)
- There are scores of activities targeted at improvements
  - Not well linked, often not clear whether they are IT specific or community oriented
  - Do we focus on reducing the impact or on getting more compute for the same?
- We probably can benefit from the experience of other communities like HPC
  - Have to be aware of the differences (continuous upgrades vs from scratch)
- WLCG (Sites and Experiments) are key for developing a strategy