

WLCG Environmental Sustainability Forum

Why do we use the term Environmental Sustainability?

- Outside our community sustainability is understood much more broadly
- The UN defined 17 goals for sustainable development
 - Only some are related to protecting the environment
- To avoid confusion we use “Environmental Sustainability”
 - Internally “Sustainability” is sufficient



Why do we need a Forum?

The WLCG [Environmental Sustainability Workshop](#) and [WLCG/HSF Plenary on Environmental Sustainability](#) covered or touched on:

- Efficiency of facilities
 - PUE, Hardware, scheduling, heat recovery....
 - Lifecycle management
 - Location
- Impact of different CPUs, GPUs, storage technology
- Impact of the efficiency of software
 - Experiments and packages (Generators, Detector Simulation...)
- Optimising workflows, reducing replicas, minimising the impact of errors, sharing results..
 - (Very) compact data formats
 - (Very)Fast Simulation....
- Impact of new approaches
 - Machine Learning, interactive analysis, Quantum Computing
- Liaising with other, related projects
-

Too many to fit into the existing WLCG events to follow up on the progress

Since then many student projects related to these topics and the open questions have been started

Primary Function:


Discuss and act on the three recommendations of the [WLCG Strategy 2024-2027](#)

- [ENV-1] WLCG to agree on **metrics** and provide a **framework to collect information** related to energy efficiency.
- [ENV-2] WLCG to facilitate the use of more energy-efficient hardware where possible, depending on the readiness of the **experiment software** and the **common libraries**.
- [ENV-3] WLCG to develop and promote a **sustainability plan** to improve energy efficiency and reduce carbon footprint, covering **software**, **computing models**, **facilities**, and **hardware technology** and **lifecycle**

This leads to the mandate for the forum

Mandate

Ensure that progress happens on action items concerning measuring and improving the energy and resource awareness of WLCG sites and jobs by:

- Providing a **public discussion** forum, attended by a **core set** of people
- Identifying areas where progress is needed
- Setting **clear goals**  Leading (hopefully) to activities by the stakeholders
 - following up on them
 - Communicating and disseminating the outcomes to:
 - WLCG Management Board, HEPIX, Open Technical Forum...
 - Receiving their feedback
- Liaise with the WLCG Technical Coordination Board
- Work **closely** with the LHC **experiments**
 - Follow up on their activities and get feedback

Leading (hopefully) to activities by the stakeholders

Mandate

There are many activities and projects related to sustainability in computing

- They are popping up like mushrooms after rain....
- Some are entangled with the HEP and LHC community, not all are relevant for us
- The Forum will establish **informal** links with (some) of them
 - Evaluate opportunities for collaboration
 - Serve as a point of contact
 - Based on our finding the WLCG Management Board will decide on formal collaboration
 - Some examples: [HEP Software Foundation](#), [GreenDIGIT](#), [HECAP+](#), [Research Facility 2.0](#), [ODISSEE](#)...
 - At the moment it isn't clear how we will deal with companies
 - Some companies, like Dell and HP, are doing very detailed studies on embedded carbon and product lifecycle, far beyond what we could do.
 - However, there is an increased risk to take part in **Green Washing™**



Organisation

- Coordinators
 - Caterina Doglioni, David Britton, Markus Schulz
 - These will work closely with **domain experts**
 - For benchmarking, evolution of hardware, software trends, lifecycle.....
- Contact:
 - Sign up to wlcg-sustainability-forum@cern.ch for announcements
 - Contact the coordinators via: wlcg-sustainability-forum-coordination@cern.ch.
 - Let us know whether a Mattermost channel is needed
- Meetings: First meeting → **October 1st 16:30**
 - Short, one hour, meetings, once a month on the first Wednesday at 16:30
 - This avoids collisions with HSF and the Compute and Accelerator Forum
 - **One topic, in-depth discussions**
 - **Suggestions for topics are welcome!!**



A (non-exhaustive) list of topics/projects we'd like to keep an eye on

- Estimation of environmental cost (and gains) of simulation chain
 - Examples:
 - Simulation steps on GPU / after specific optimisation
 - **MC generators** and parton shower programs on GPUs
 - Could also consider other kinds of simulation (e.g. fastSim / with ML)
- HEPscore23 with added power consumption monitoring capabilities, and its use for different sites
- Scripts to collect power consumption information for different sites
 - With the goal to integrate this into the accounting tools
- Sites with specific energy management needs (e.g. DESY in summer)
 - Digital Twins
- How to estimate / simulate power consumption of workflows and sites
- Embedded CO₂e of hardware and lifecycle management
 - Building a model, improving on Wadensteins work
- Experiments' efforts and discussions towards HL-LHC computing models
 - Can we escape Jevons Paradox? (higher efficiency is compensated by increased demand)
- Data challenges including power consumption

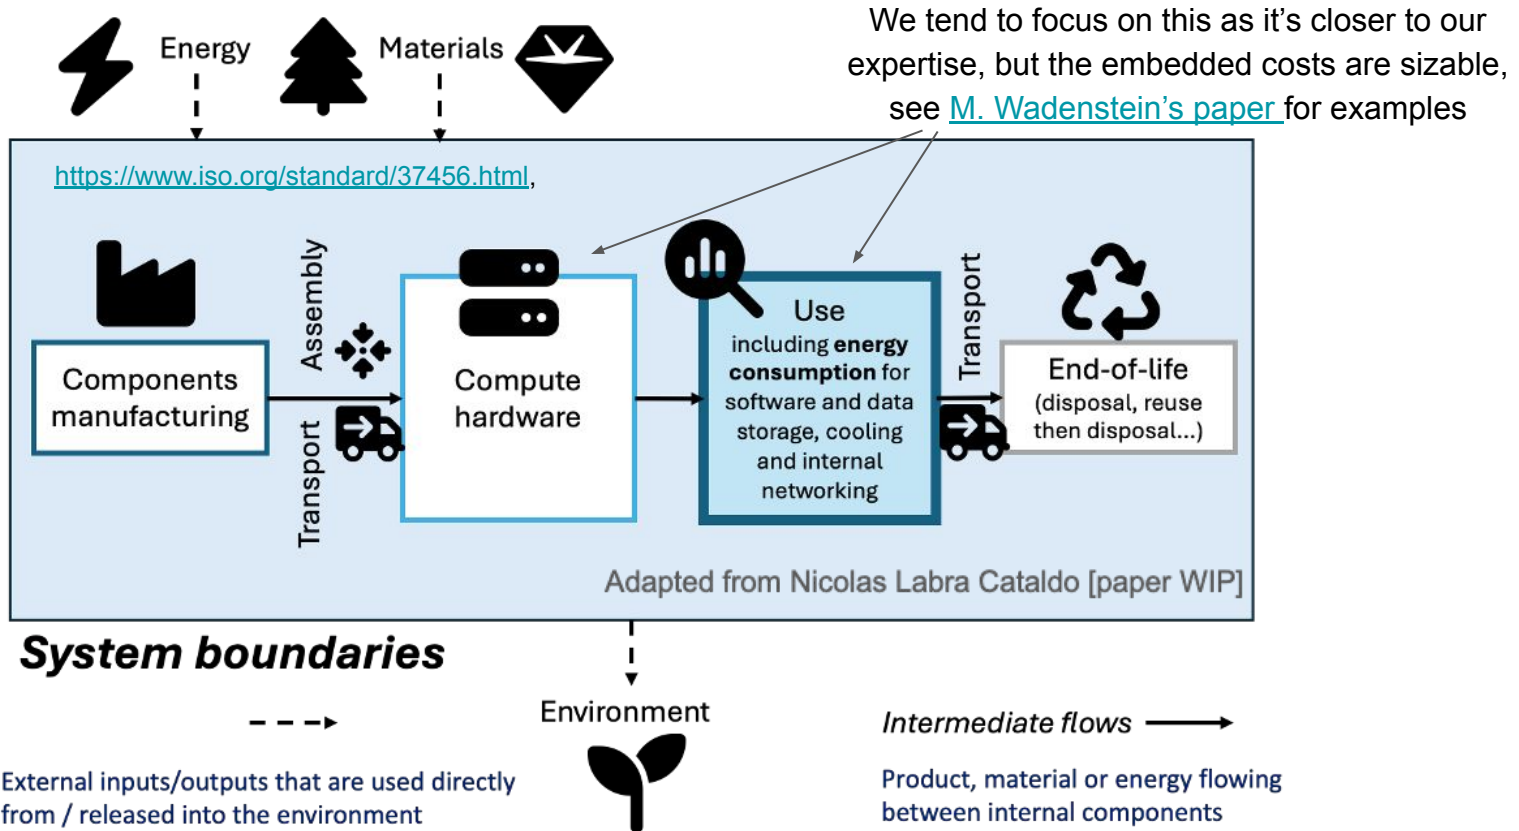
Next Talk

Discussion

- How can we best involve the WLCG sites in this Forum ? (bottom-up / top-down strategies)
- Anyone an idea for a visual identity (logo)?




Embedded carbon / lifecycle: setting the scene



A few words on lifecycle management

Link to the report

- Just to raise interest.. (Expert couldn't present, waiting for management consent)
- Why discuss lifecycle management?
 - Embedded CO2e is [complex](#) and is a significant part of the overall impact 
 - Depending on the gCO2/kWh in a region this can be the dominant factor
 - In places with “clean” electricity the break even point can be > 7 years
 - Technology changes
 - The rise of the SSD, GPUs are here to stay..
 - Different components of a system **have now different times of usefulness**
 - GPUs are outdated after 3-4 years
 - SSDs have limited lifetime, depending on usage ...
 - Newer CPUs/GPUs are more power efficient, with similar embedded CO2e
 - Replacement strategies have to take many aspects into account
 - Failure rates depend on age and component type
 - **CERN IT started to measure this, more soon**
 - Cost of: energy, hardware, maintenance, power and space envelope

Age distribution of the CERN IT hardware

