

# Environmental Sustainability in Digital Research Infrastructures - the GreenDIGIT Project

**Catalin Condurache** (*EGI Foundation*)

**Gergely Sipos** (*EGI Foundation*)

**Yuri Demchenko** (*University of Amsterdam*)



# Outline

- GreenDIGIT project scope and goals
  - Shared responsibility model for sustainability
- Landscape survey
  - Objectives, targets
  - Responses, findings
  - Opportunities, recommendations
- Environmental sustainability metrics
- 'Green' optimised frameworks
- Conclusions








# GreenDIGIT Project - Founding Digital RIs (ESFRI)

- **EBRAINS** - An open research infrastructure that gathers data, tools and computing facilities for brain-related research
- **EGI** - International federation delivering e-Infrastructure and open solutions for advanced computing and data analytics in research and innovation
- **SLICES** - Scientific Large-scale Infrastructure for Computing and Communication Experimental Studies
- **SoBigData** - Distributed, pan-European, multi-disciplinary research infrastructure aimed at using social mining and Big Data to understand the complexity of our contemporary, globally interconnected society



<https://greendigit-project.eu>







GreenDIGIT

# GreenDIGIT Project (2024-2027) - Objectives

- **O1: Assess status and trends** of low impact computing within 4 DIGIT RIs (EGI, SLICES, SoBigData, EBRAINS) and in the broader digital service provider community of ESFRIs, to produce **recommendations and roadmaps** for providers during and beyond the project.
- **O2: Provide reference architecture and design principles**, as well as an actionable model for RIs about environmental impact assessment and monitoring, reflecting on the **whole RI lifecycle** and including the digital infrastructure components and their interaction with the broader environment.
- **O3: Develop and validate new and innovative technologies, methods, and tools** for digital service providers within European Research Infrastructures through which they can reduce their energy consumptions and overall environmental impact.
- **O4: Develop and provide for researcher technical tools** that assist them in the design, execution and sharing of environmental impact aware digital applications with reproducibility, Open Science and FAIR data management considerations.
- **O5: Educate and support digital service providers and researchers** in the RI communities about good practices on environmental impact conscious lifecycle management and operation of infrastructures and services.







# Sustainability Aspects: Energy Efficiency – Decarbonisation – Environmental Impact

- **Energy Efficiency in Digital Infrastructures:**
  - **Definition:** This refers to optimizing digital infrastructures to consume as little energy as possible for a given workload or service. It's about achieving more computational or storage results with less energy input.
- **Decarbonization of Digital Infrastructures:**
  - **Definition:** This specifically targets the reduction of carbon emissions associated with the operation and maintenance of digital infrastructures.
- **Reducing Environmental Impact of Digital Infrastructures:**
  - **Definition:** This is a more comprehensive consideration of the various ways digital infrastructures might affect the environment, going beyond just energy consumption and carbon emissions.







# Sustainability Aspects: Energy Efficiency – Decarbonisation – Environmental Impact

## ● Energy Efficiency in Digital Infrastructures:

- **Definition:** This refers to optimizing digital infrastructures to consume as little energy as possible for a given workload or service. It's about achieving more computational or storage results with less energy input.

Architecture,  
Design,  
Recommendations

## ● Decarbonization of Digital Infrastructures:

- **Definition:** This specifically targets the reduction of carbon emissions associated with the operation and maintenance of digital infrastructures.

Operation,  
Monitoring, KPIs

## ● Reducing Environmental Impact of Digital Infrastructures:

- **Definition:** This is a more comprehensive consideration of the various ways digital infrastructures might affect the environment, going beyond just energy consumption and carbon emissions.

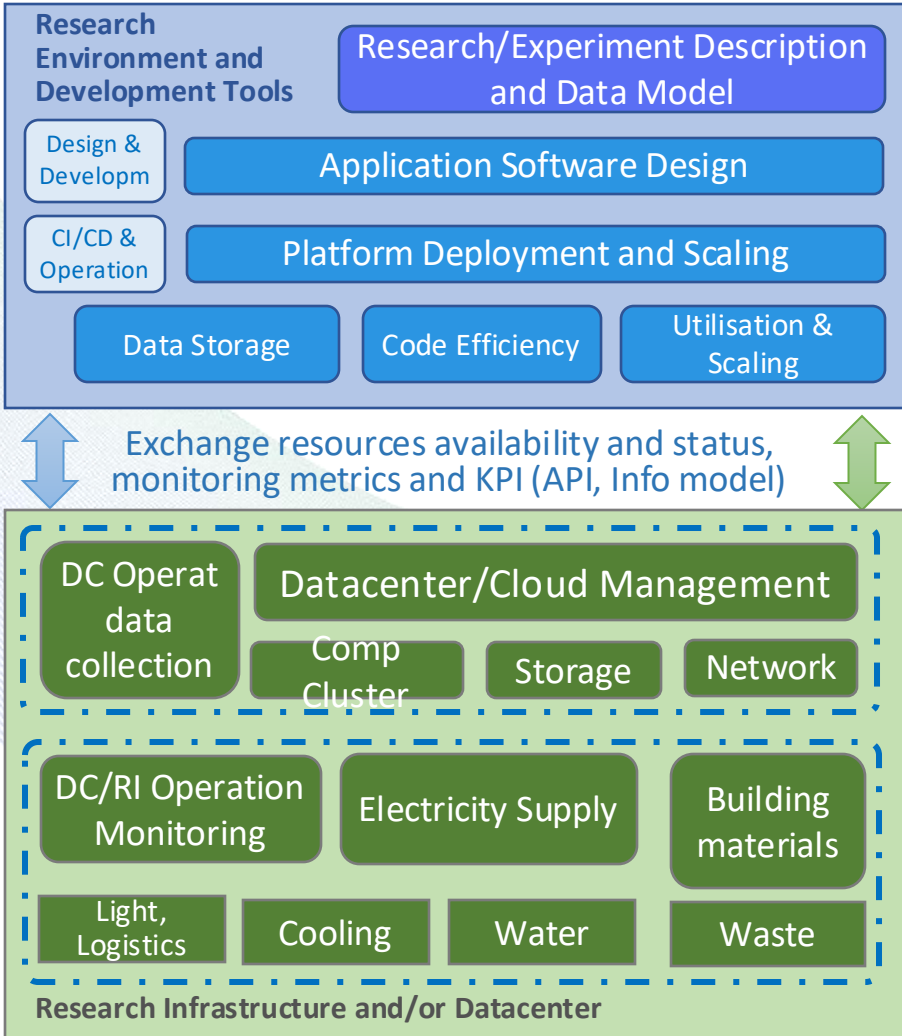
Lifecycle, Policy,  
Training







# Shared Responsibility in Sustainability



Four categories identified

- Manufacturers
- RI operators
- SW developers
- RI users

Each category must assume their share towards the common goal

GreenDIGIT brings various solutions to help last 3 groups and incentivise first group

**Project/Researcher Responsibility on RI**  
 Applications Development, Deployment, Operation, Energy usage and KPI monitoring

**Provider/Operator Responsibility of RI**  
 Research Infrastructure or Datacenter, Monitoring Energy and environmental impact metrics and KPI

**Standards and regulations**  
 Software Development Quality and Design Patterns

**Standards and regulations**  
 Datacenter and RI Building and Operation



# RIs Landscape Survey

- Objective

- To collect and analyse RIs practices, needs and gaps, and opportunities-barriers-good practices
- What did the RIs do so far? → Good practices
- Where are they going? → Initiatives
- Where could we help? → Opportunities

- Target

- GreenDIGIT consortium (EGI, SoBigData, EBRAINS, SLICES)
- Thematic RIs with **significant** digital infrastructures

- Methodology

- Preparation, survey and interviews, analysis of results and identification of key aspects







# RIs Landscape Survey - Responses

PREPARATORY PHASE

PRE-PRODUCTION PHASE

PRODUCTION PHASE



EBRAINS



ICT facilities at 14 sites  
(5G, Networks, Clouds)

**Data, Computing and Digital**

6 Computational  
25 Research facilities

3 HPC centers

~200 HTC sites  
~30 cloud sites



CERN (Particle acc., Intern.)  
ELI-Beams (Laser facility, CZ)

**Physical Sciences**

CLARIN (Language resources)

**Social Sciences & Humanities**

GÉANT (Networks, Intern.)  
Switch Cloud (CH)  
University of Freiburg (DE)

**Data, Computing and Digital**







# RIs Landscape Survey - Findings

- Included separated responses from EGI Federation sites
- Responses represent over 50 academic providers from 20+ countries:
  - 80% offering "Computation resources" as the primary resource
  - 60% offering "Networking / Communication resources"
  - 40% offering archives and scientific data
  - 30% offering also "Scientific equipment / instruments"
  - Each of the ~50 centers running between 50-10,000 physical core systems







# RIs Landscape Survey - Findings

RIs served (<https://ri-portfolio.esfri.eu/>)

- **Energy**
  - EUROfusion
- **Environment**
  - LifeWatch
- **Health & Food**
  - ELIXIR, French Bioinformatics, Instruct, METROFOOD
- **Physical Sciences & Engineering**
  - WLCG, CTA, HESS, SKAO, LOFAR, BELLE II, ESA, AGATA, FOCAL
- **Social Sciences & Humanities**
  - CLARIAH, CESSDA, ODISSEI (NL)







# RIs Landscape Survey - Findings

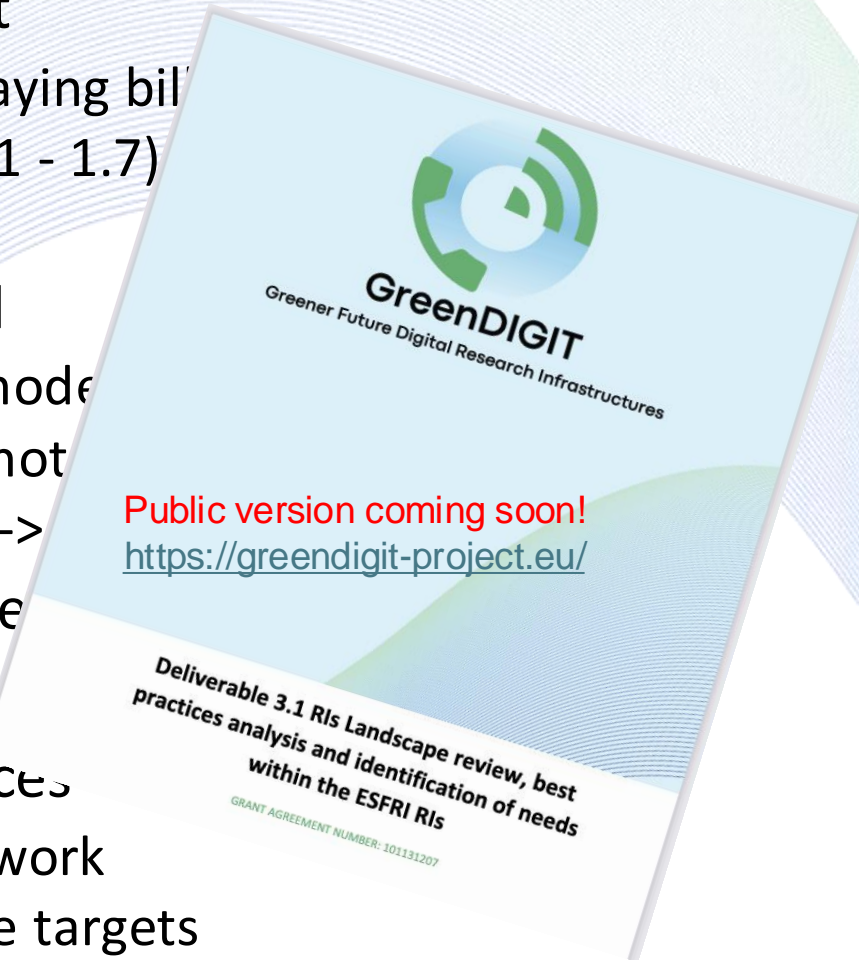
- **Energy** is primary, **water** is secondary environment impact
  - Tracking total energy consumption is nearly everywhere (paying bills)
  - Power Usage Efficiency (PUE) is most common (between 1.1 - 1.7).
  - Others: Watts/VM, Watts/GPU, Watts/vCPU, FLOPS/Watt
- Initiatives to **lower consumption/waste** exist, but isolated
  - Physical upgrade - Not trivial because of investment cost (modernise HW, A/C, building)
  - Software based optimisation at few places. Workflows are not designed for energy saving
  - Few practices for waste heat reuse (heating offices/homes -> significant investment)
- Green computing **policy** is rare. An example at IRES, France:  
<https://scigne.fr/resources/policies/green-computing-policy.pdf>
- Very few centres **encourage/educate** staff on good practices
  - Switching lights or equipment off, reduce travel, cycling to work
  - Both staff operating the IT equipment and the end users are targets





# RIs Landscape Survey - Findings

- **Energy** is primary, **water** is secondary environment impact
  - Tracking total energy consumption is nearly everywhere (paying bills)
  - Power Usage Efficiency (PUE) is most common (between 1.1 - 1.7)
  - Others: Watts/VM, Watts/GPU, Watts/vCPU, FLOPS/Watt
- Initiatives to **lower consumption/waste** exist, but isolated
  - Physical upgrade - Not trivial because of investment cost (modernisation)
  - Software based optimisation at few places. Workflows are not optimised
  - Few practices for waste heat reuse (heating offices/homes -> district heating)
- Green computing **policy** is rare. An example at IRES, France  
<https://scigne.fr/resources/policies/green-computing-policy.pdf>
- Very few centres **encourage/educate** staff on good practices
  - Switching lights or equipment off, reduce travel, cycling to work
  - Both staff operating the IT equipment and the end users are targets







# Survey -> Opportunities -> GreenDIGIT Plans

- **Framework** for managing the full digital RI lifecycle from the environmental impact perspective
  - For status assessment, Design, Implementation, Operation, Decommissioning
- **Software tools** for power usage efficiency
  - Metrics infrastructure, Schedulers, Reproducibility frameworks, Resource managers (VMs, containers, workflows, AI/ML models, Data, etc.)
- **Develop skills and collaborations**
  - Training resources, events, courses, forums





# Planned SW Tools for Power Usage Efficiency

PROVIDERS' SIDE

- **Metrics infrastructure**  
Gather relevant environmental sustainability metrics from providers  
Serve metrics to decision making systems (schedulers, frameworks, VREs)
- **Energy-efficiency aware schedulers**  
Reshuffle jobs based on its type (CPU bound or I/O bound)  
Delay execution based on electricity cost, regulations, carbon intensity

- **Developer tools for energy efficient software design**  
JupyterHub, Greenspector
- **Scientific reproducibility and resource management platform**  
JupyterHub, Binder, DataHub, SoBigData, etc.

USERS' SIDE





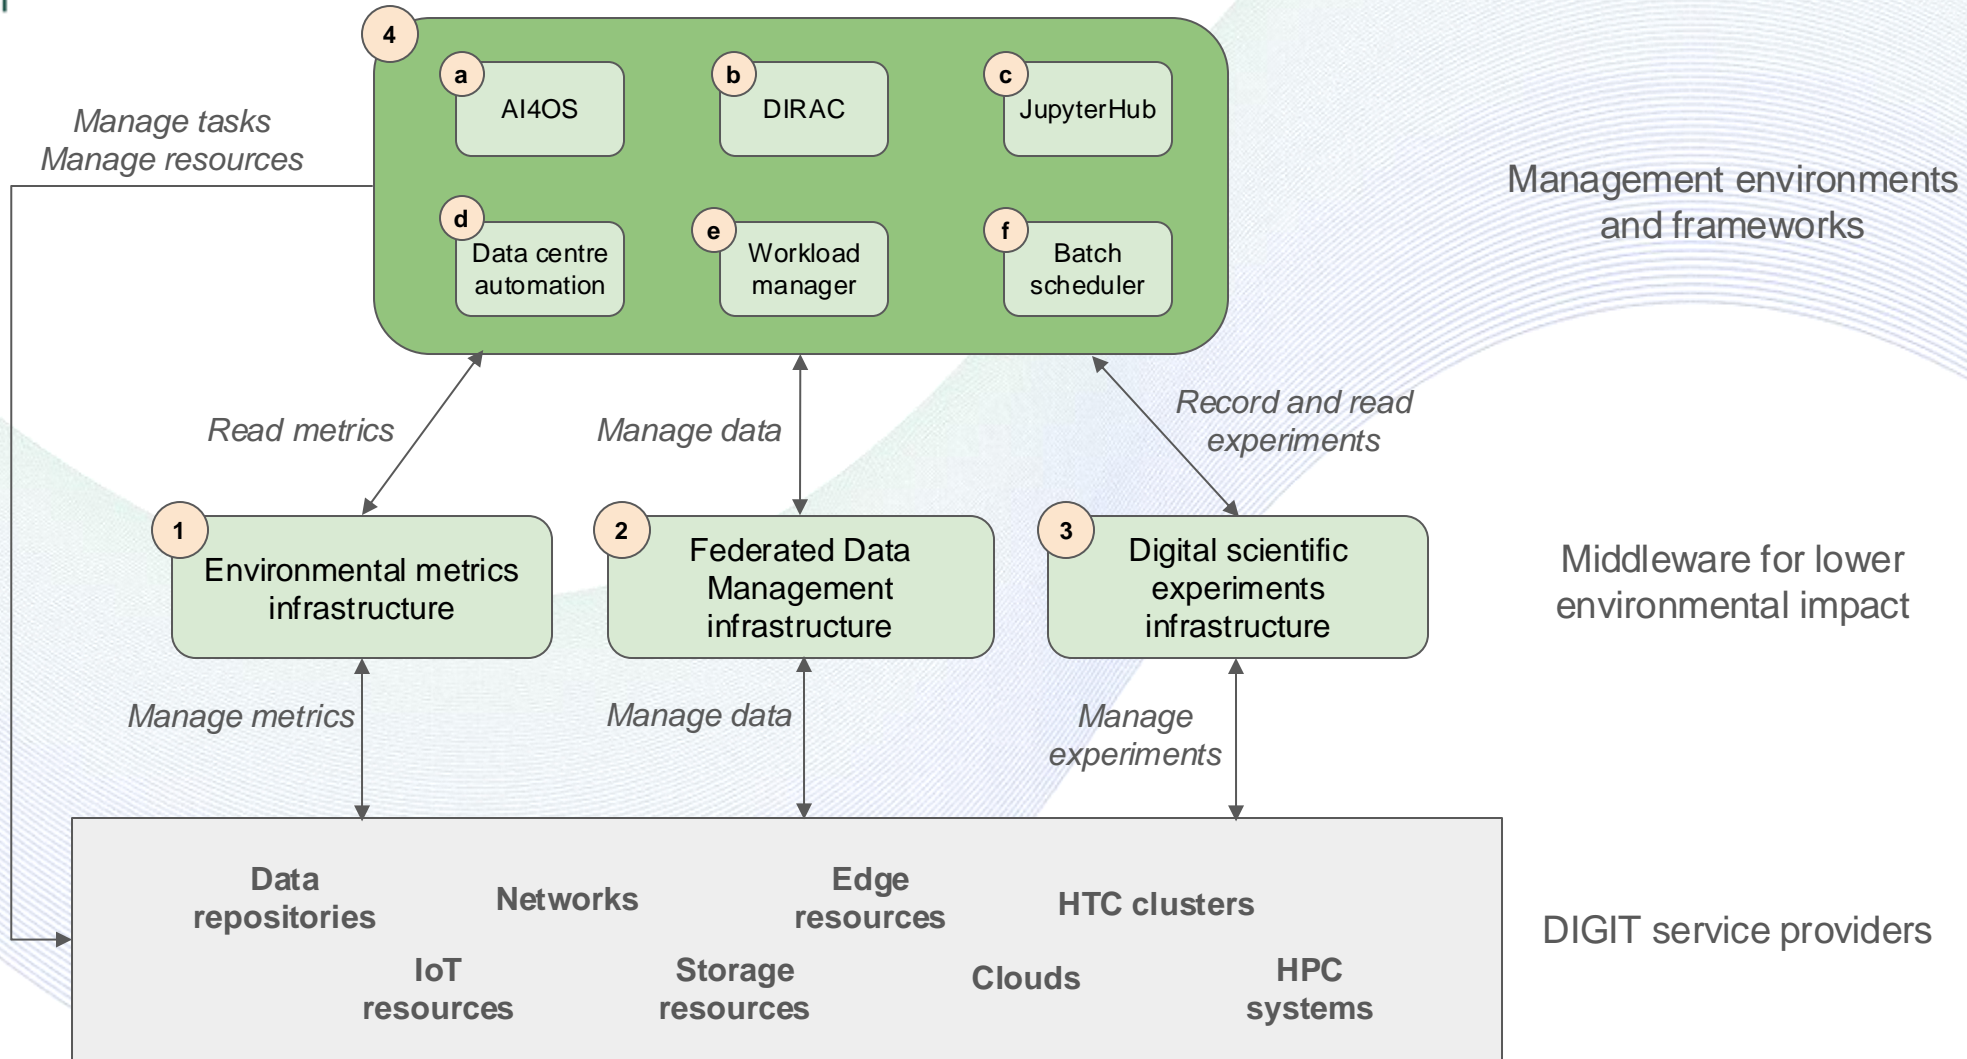
GreenDIGIT

# Initial Recommendations

Recommendation	How to achieve	How to measure
<b>Use less (per flop/Byte/bit)</b>	HW: Modern building; Modern hardware (e.g. flash instead of disks); Led lights SW: Optimised code (AI/workflow); Job monitoring (to avoid ill configured jobs on HPC nodes; Switching off unused nodes Physical: Less staff travel; Data centre in cold region	W / YEAR (energy) m <sup>3</sup> (water)
<b>Use for the right purpose</b>	Higher A/C temp; Hot/cold aisles; Turn off lights	PUE
<b>Use from renewable source</b>	Change supplier; Deploy solar panel; Delay workload to when energy is greener	Green %
<b>Share the infrastructure</b>	Invest into shared infrastructure (use) for better cost-value ratio, to eliminate 'downtime waste'	In-house - Cloud ratio
<b>Reuse your waste</b>	Repair; Waste heat reuse (office heating); Circular economy recycling (paper, plastic, ...)	Waste reuse %
<b>Train people, support new ideas</b>	Office training; Information dashboards	# training hours # new initiatives
<b>Measure and improve</b>	Regularly update the metrics	

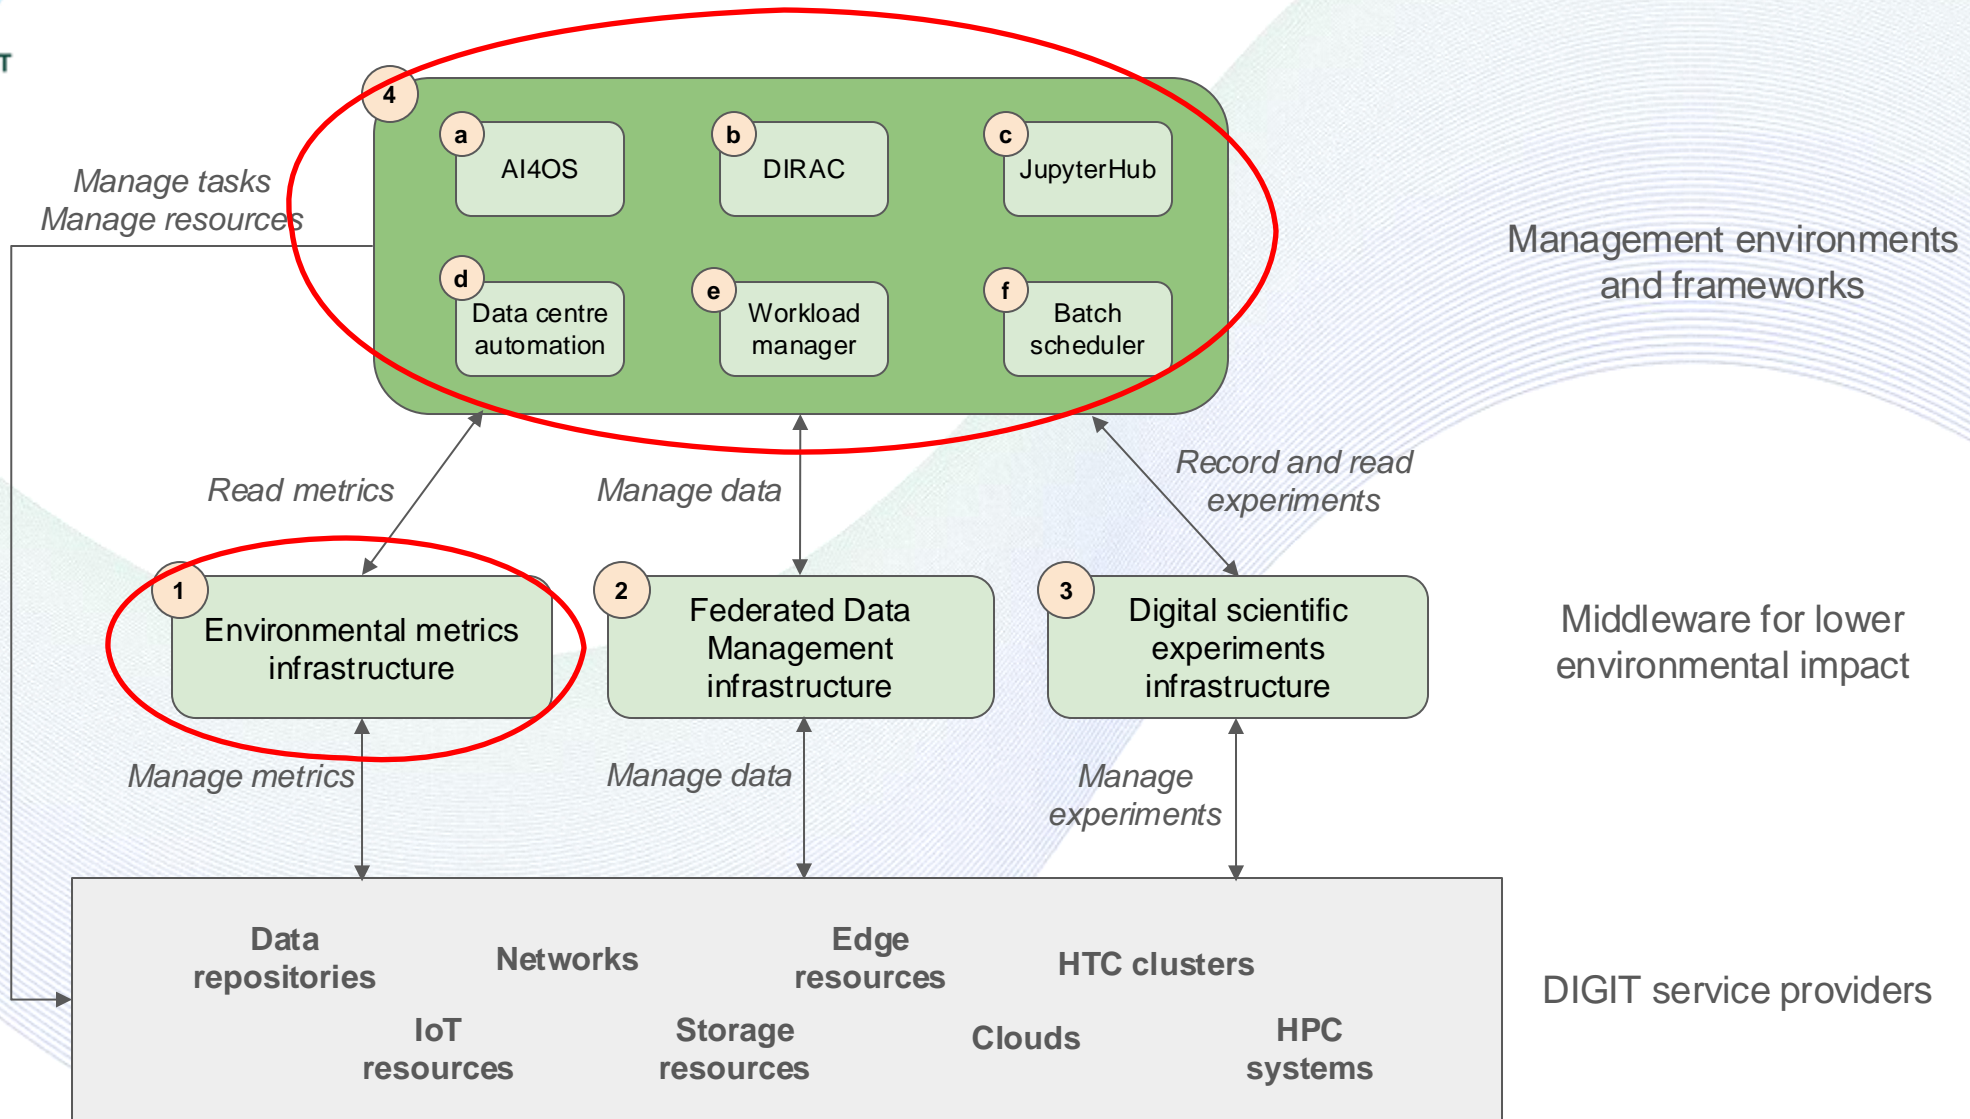


# GreenDIGIT Technical Architecture





# GreenDIGIT Technical Architecture



# Environmental Sustainability Metrics

(work in progress...)

EGI partners in consortium involved in choosing and proposing a use for them

- *Simplification 1* – Focus on **energy** as the main environment impact (water is secondary)
- Needs
  - Metrics that capture these aspects
  - An infrastructure that collects/aggregates the metrics
- Challenges
  - What can be measured at equipment level?
  - What are providers willing to share?
- *Simplification 2* – consider '**Operation**' phase only
  - ESFRI lifecycle – Concept development, Design, Preparation, Implementation, **Operation**, Termination
  - **Operation** – Frontier research results, services to scientific community, continuous upgrade of instrumentation and methods, outreach, political and financial support for long-term operation







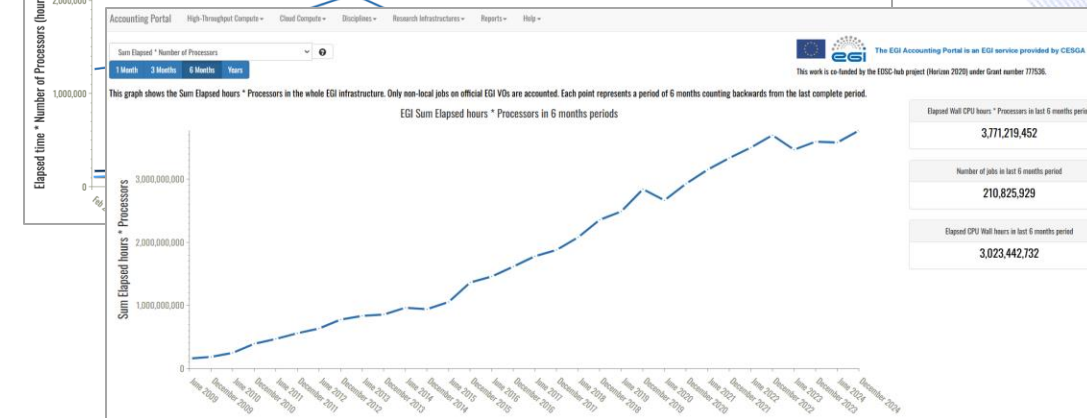
# Environmental Sustainability Metrics

(work in progress...)

- A very small set to start with from
  - PUE, WUE, ERF, REF, Carbon intensity, Energy consumption, ...
  - [EC Delegated Regulation \(EU\) 2024/1364](#) of 14 March 2024
  - Likely only 2-3 for the beginning
- EGI Accounting system ([accounting.egi.eu](https://accounting.egi.eu)) proposed to collect and manage them
  - Monthly aggregation (even if records are more frequent)
  - It already has usage data

Spain — Elapsed time \* Number of Processors (hours) by Resource Centre and Month (All VOs)

Resource Centre	Jun 2024	Jul 2024	Aug 2024	Sep 2024	Oct 2024	Nov 2024	Dec 2024	Jan 2025	Total	Percent
SIR	80	12,680	0	0	0	0	0	0	315,173	1.33%
ESGA-CLUST	00	181,456	187,500	192,277	186,417	203,702	198,878	204,254	2,278,516	8.97%
ESGA-GRID	22	165,120	111,840	114,580	111,181	114,576	110,885	114,520	1,144,520	4.52%
ESGA-LOG2	00	2,066,626	1,751,120	1,454,882	1,548,083	1,887,573	1,641,476	1,781,450	1,783,365	69.44%
Total	94	2,365,282	2,050,460	1,761,689	1,847,681	2,205,853	1,856,161	2,106,250	2,071,470	23,709,757
Percent	%	9.98%	8.62%	7.43%	7.9%	9.30%	8.2%	8.8%	8.4%	





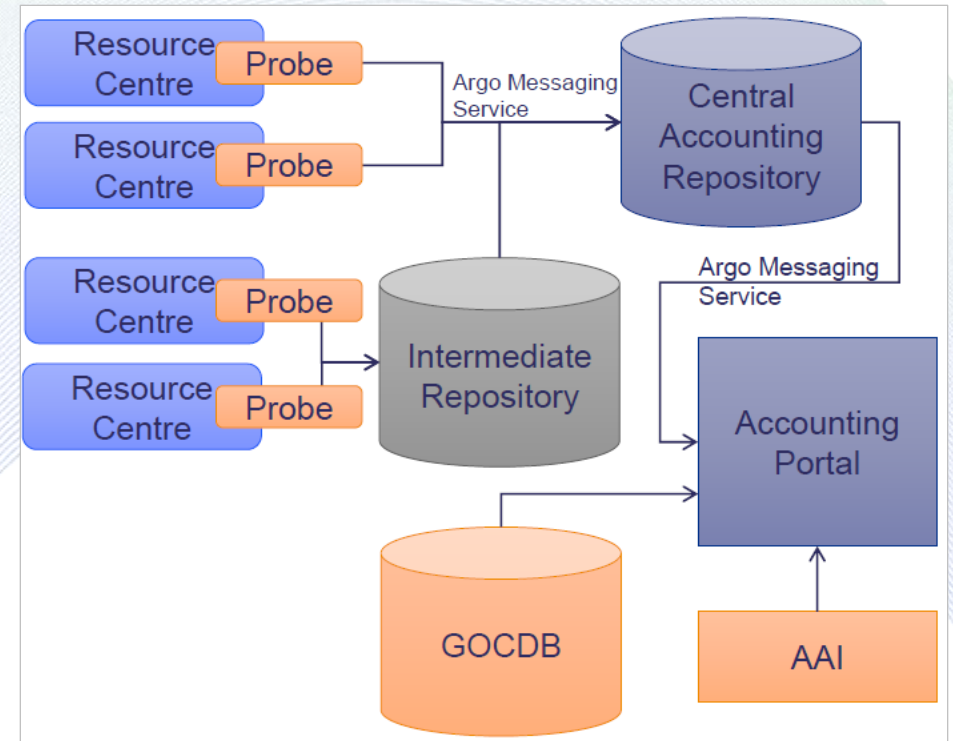


# Environmental Sustainability Metrics

(work in progress...)

## EGI Accounting system for metrics

- Features
  - scalable for O(100) providers
  - scalable for O(10) consumers
  - can store historical values for the stored metrics
  - support for time-series queries
  - available GUI and read-write API
  - support for cumulative reporting of consumption of long-running workflows







# Towards 'Green' Optimised Frameworks

4

GreenDIGIT will implement energy aware workflow/orchestrator optimisation algorithms within relevant frameworks

- DIRAC Workload Manager (@CNRS)
  - Perform validation through scientific user cases (WeNMR, Biomed) for HTC and Cloud
  - Possibilities to alternate production runs with/without 'green' metrics to validate algorithms
- AI4EOSC framework (@CSIC)
  - Validation through iMagine ML/AI models



# AI4EOSC Power Consumption & Environmental Impact

AI4EOSC aims to make users aware of the impact produced by training and deploying AI models because it can leverage significant computing power.

We need to instrument at different levels our platform in order to:

Produce metrics at **datacenter level**.

- PUE, energy quality, efficiency, hardware characteristics.
- Currently static, working to make data collection dynamic.

Produce metrics at **model level**.

- Impact on training phase.  
(i.e. how much does it *cost* to build a model)
- Impact on inference/prediction.  
(i.e. how much does it *cost* to use the model)



The approach is to measure the power consumption of the platform through monitoring the power consumption of the processes that runs the VMs where is deployed on. This has to be done from bare metal machines (It requires access to Infrastructure Level).

- Not trivial, as GPUs and other specialized hardware must be taken into account.
- Nowadays, it is not possible to automate the instrumentation of the whole OpenStack cloud stack without **cooperation from resource providers**.





# Towards 'Green' Optimised Frameworks

GreenDIGIT will implement energy aware workflow/orchestrator optimisation algorithms within relevant frameworks

- Terraform DevOps tool (@SZTAKI)
  - Validation via Environmental science use cases using Terraform containers
- Data Center Energy scheduler (@CESNET)
  - Power-grid aware computing algorithms
  - Control consumption based on
    - Electricity cost, power-grid regulation needs
    - Current carbon intensity of regional electricity
    - Reused heat demand







# GreenDIGIT, EGI and WLCG

- Following [WLCG Environmental Sustainability Workshop](#) in Dec 2024
  - **power consumption per CPU core** and **power consumption per HEPscore23** metrics are considered
  - Might need some HOWTO if available
- Sites and/or experiments are welcome to join
  - Either during validation process
  - Or during integration and acceptance within RIs
- EGI Green Computing Task Force
  - Re-activated following EGI-ACE project
  - The entity that links interested EGI stakeholders/sites to GreenDIGIT
  - <https://confluence.egi.eu/display/EGIBG/Task+Force%3A+Green+Computing>
  - Email list: [green-computing-TF@mailman.egi.eu](mailto:green-computing-TF@mailman.egi.eu)







# Conclusions

- GreenDIGIT works on an overarching framework for environmental impact reduction methodology in digital sciences
- The framework considers the RI lifecycle and shared responsibility models
- Technological solutions to be provided for compute centers and software developers
- Get involved via the EGI GC TF - [green-computing-tf@mailman.egi.eu](mailto:green-computing-tf@mailman.egi.eu)





**Thank you!**  
**Questions!**







GreenDIGIT

# Backup slides

**COMMISSION DELEGATED REGULATION (EU) 2024/1364 of 14 March 2024 on the first phase of the establishment of a common Union rating scheme for data centres - [https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L\\_202401364](https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=OJ:L_202401364) - result of [Directive \(EU\) 2023/1791](#)**

## Proposed metrics - KPIs (Annex II)

- **Total energy consumption** (“ $E_{DC}$ ”, in kWh) of the reporting data centre shall be measured as defined by, and by using the methodology in the CEN/CENELEC EN 50600-4-2 standard or equivalent. Total energy consumption includes the use of electricity, fuels and other energy sources used for cooling.
- **Total energy consumption of information technology equipment** (“ $E_{IT}$ ”, in kWh)
- **Total water input** (“ $W_{IN}$ ”, in cubic meters) shall be measured as defined by, and by using the methodology set out in the CEN/CENELEC EN 50600-4-9 standard WUE Category 2
- **Waste heat reused** (“ $E_{REUSE}$ ”, in kWh) shall be measured as defined by, and by using the methodology set out in, the CEN/CENELEC EN 50600-4-6 standard or equivalent
- **Average waste heat temperature** (“ $T_{WH}$ ”, in degree Celsius) shall be measured as the temperature of the fluid used to cool the information and communication technology equipment in the data centre computer room, averaged over the year, and across every measurement point.







GreenDIGIT

# Backup slides

- **Total renewable energy consumption** (“ $E_{RES-TOT}$ ”, in kWh)  
 $E_{RES-TOT}$  is the sum of  $E_{RES-GOO}$ ,  $E_{RES-PPA}$  and  $E_{RES-OS}$ , as defined below
- **Total renewable energy consumption from Guarantees of Origin** (“ $E_{RES-GOO}$ ”, in kWh) shall be determined as the sum of the Guarantees of Origin purchased and retired by the reporting data centre
- **Total renewable energy consumption from Power Purchasing Agreements** (“ $E_{RES-PPA}$ ”, in kWh) shall be determined as the amount of energy from Power Purchasing Agreements made by the reporting data centre
- **Total renewable energy consumption from on-site renewables** (“ $E_{RES-OS}$ ”, in kWh) shall be measured as the energy generated from on-site renewable energy sources within the data centre boundary.

**Proposed Data Centres sust. Indicators (Annex III)** - shall be calculated based on the information and key performance indicators communicated to the European database on data centres in accordance with Annexes I and II.

- **Power Usage Effectiveness (PUE)** =  $E_{DC} / E_{IT}$
- **Water Usage Effectiveness (WUE)** =  $W_{IN} / E_{IT}$   
 $E_{IT}$  in MWh
- **Energy Reuse Factor (ERF)** =  $E_{REUSE} / E_{DC}$
- **Renewable Energy Factor % (REF)** =  $E_{RES-TOT} / E_{DC}$   
Default values from EC Country Report

