

How can we quantify ...
the "green-ness" of code?

 **GREEN CODING;**

Who am I

Arne Tarara / Green Coding Solutions GmbH



- Electrical Engineer by training
- Software-Dev 16+ years - CTO / CEO / Developer
- Founder and entrepreneur of GCS with 2 software company track record
- Current: **Green Coding Solutions GmbH**
 - R&D for sustainable software solutions



Cloud Energy

Estimation for Cloud and VMs

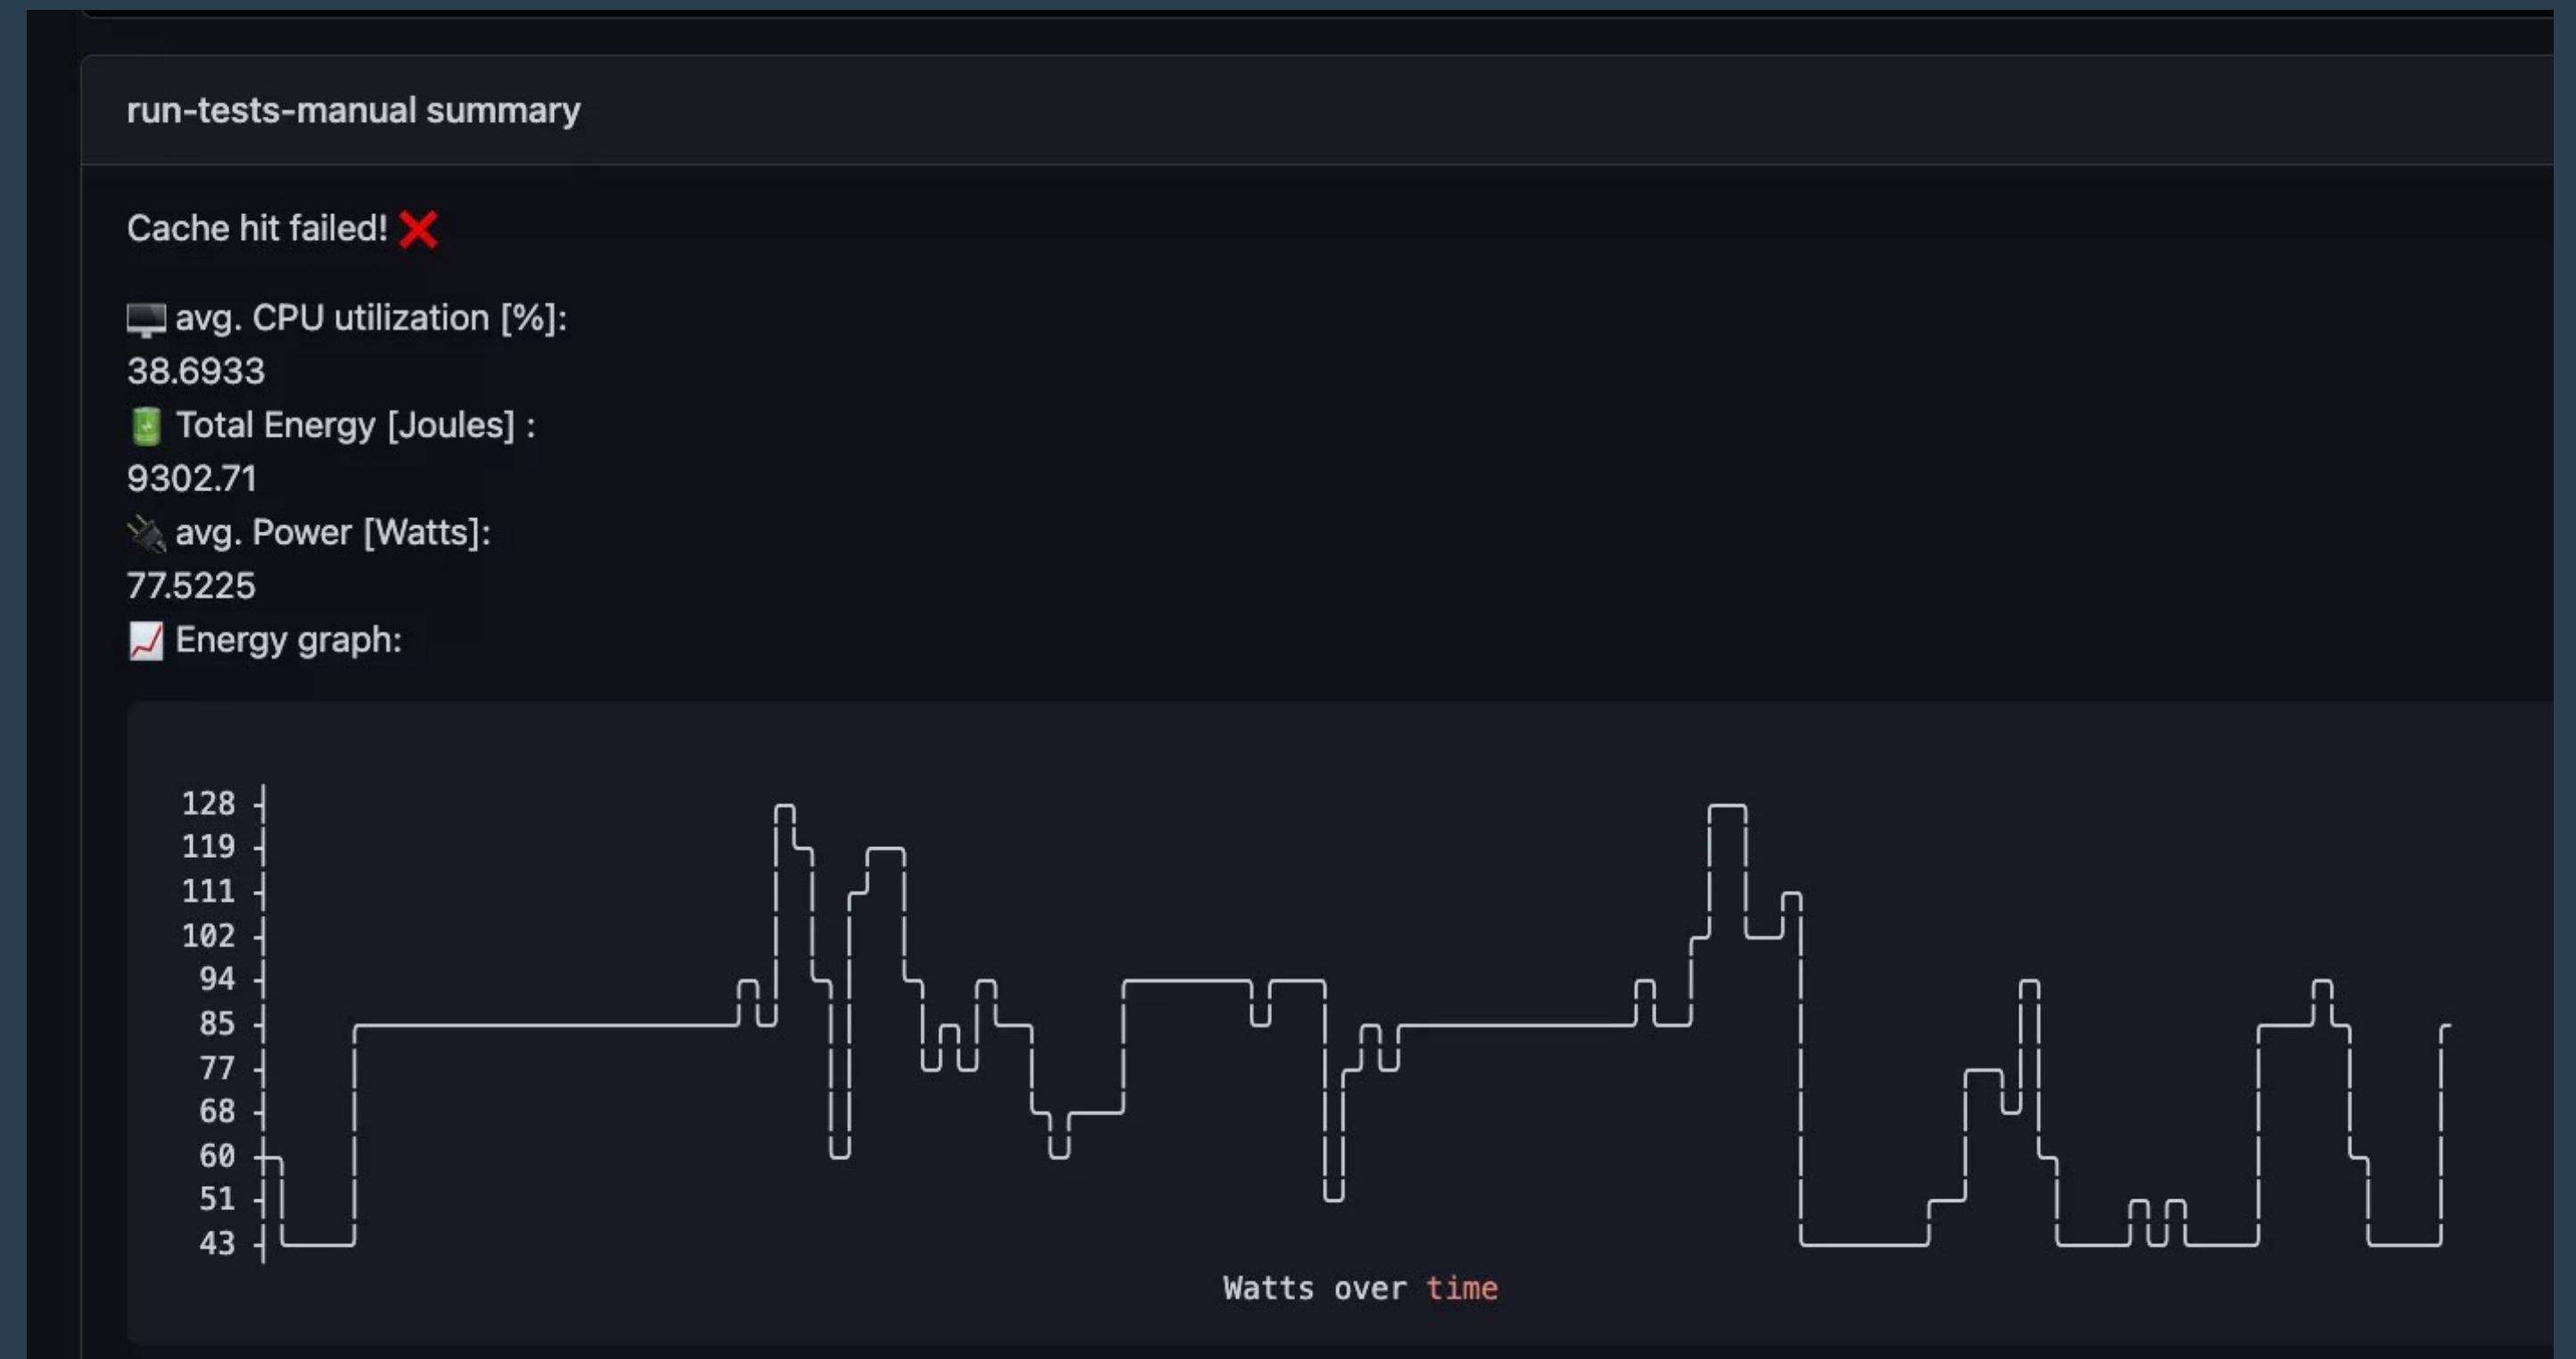
- The setup of the model is based on a research paper from Interact DC and the University of East London
- 90%+ Accuracy in/out-of sample
- Near 0% overhead
 - XGBoost + POSIX stream implementation



Eco-CI

Estimation in CI / CD Pipelines

- Integration into Github / Gitlab directly
- Export to central dashboard
- Statistical comparisons over time
- PR-triggered measurements



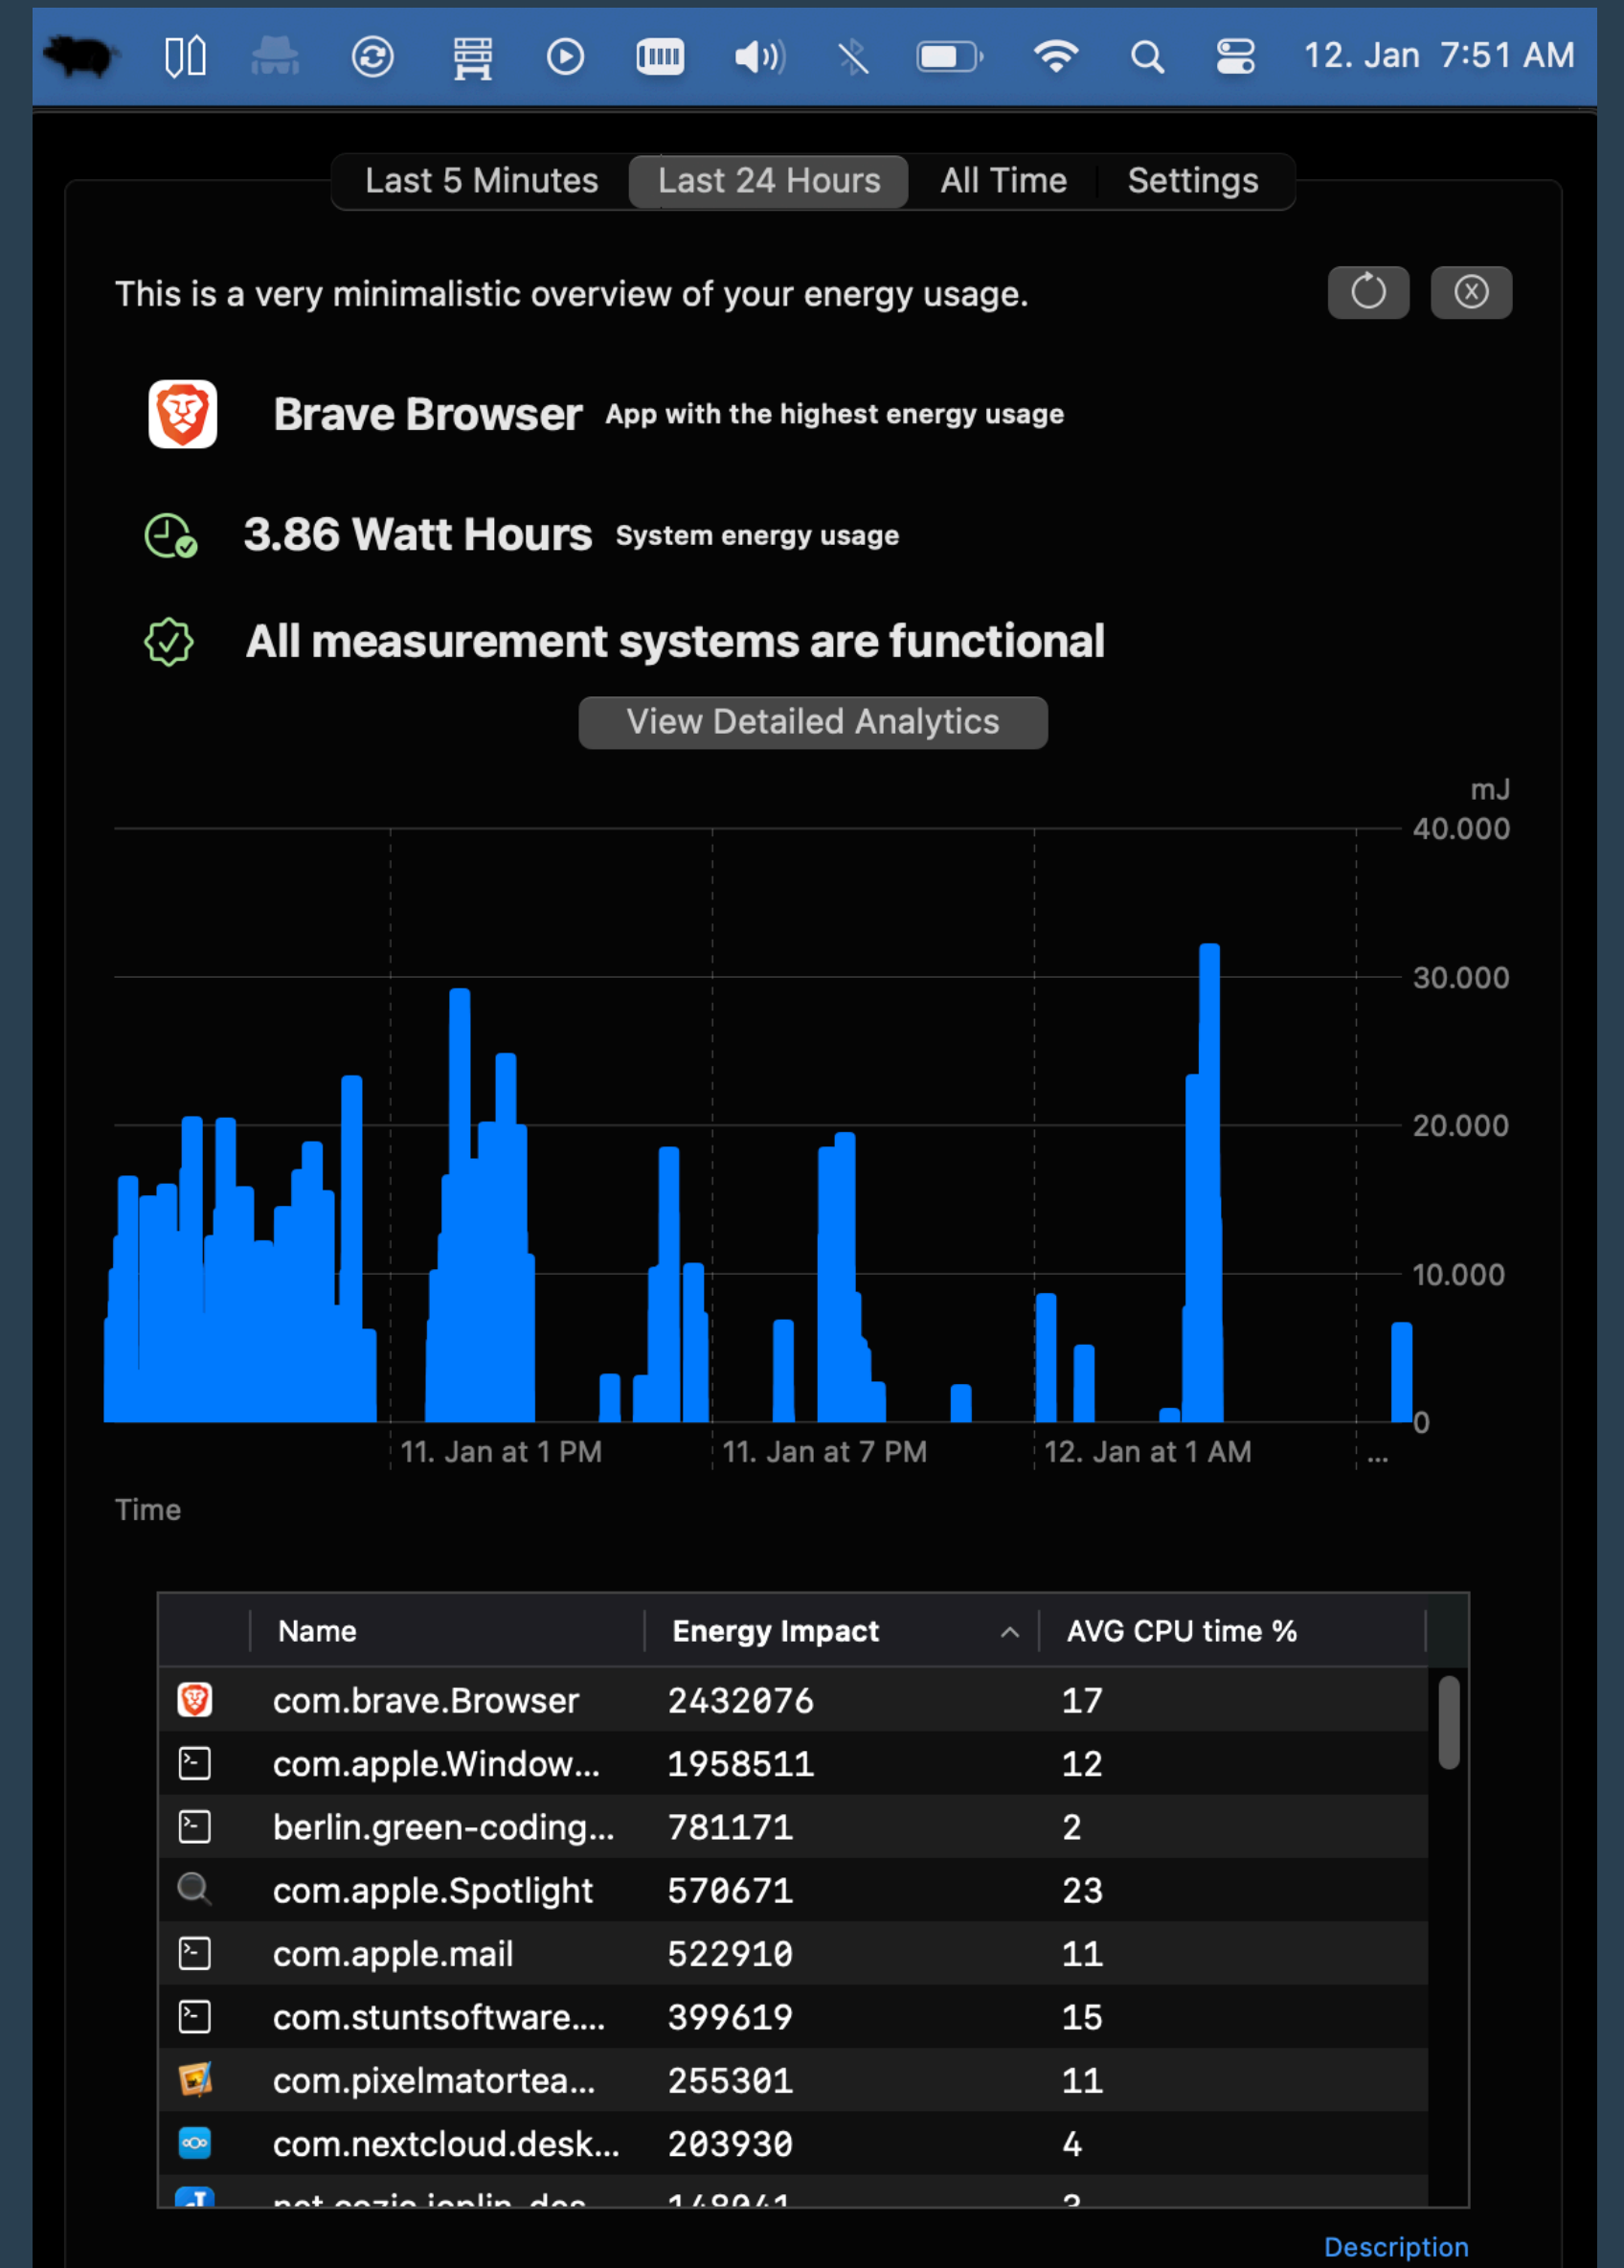
Example for Github PRs: <https://github.com/green-coding-berlin/green-metrics-tool/pull/653>

Example for Django Measurements over time: <https://metrics.green-coding.io/ci.html?repo=green-coding-berlin/django&branch=main&workflow=60545070>

Power Hog

Development Cost measurement

- Direct measurement of energy on the developer machine
- Granularity per process
- Aggregation per project
- API and drill-down option in dashboard
- Local database for quick analyses

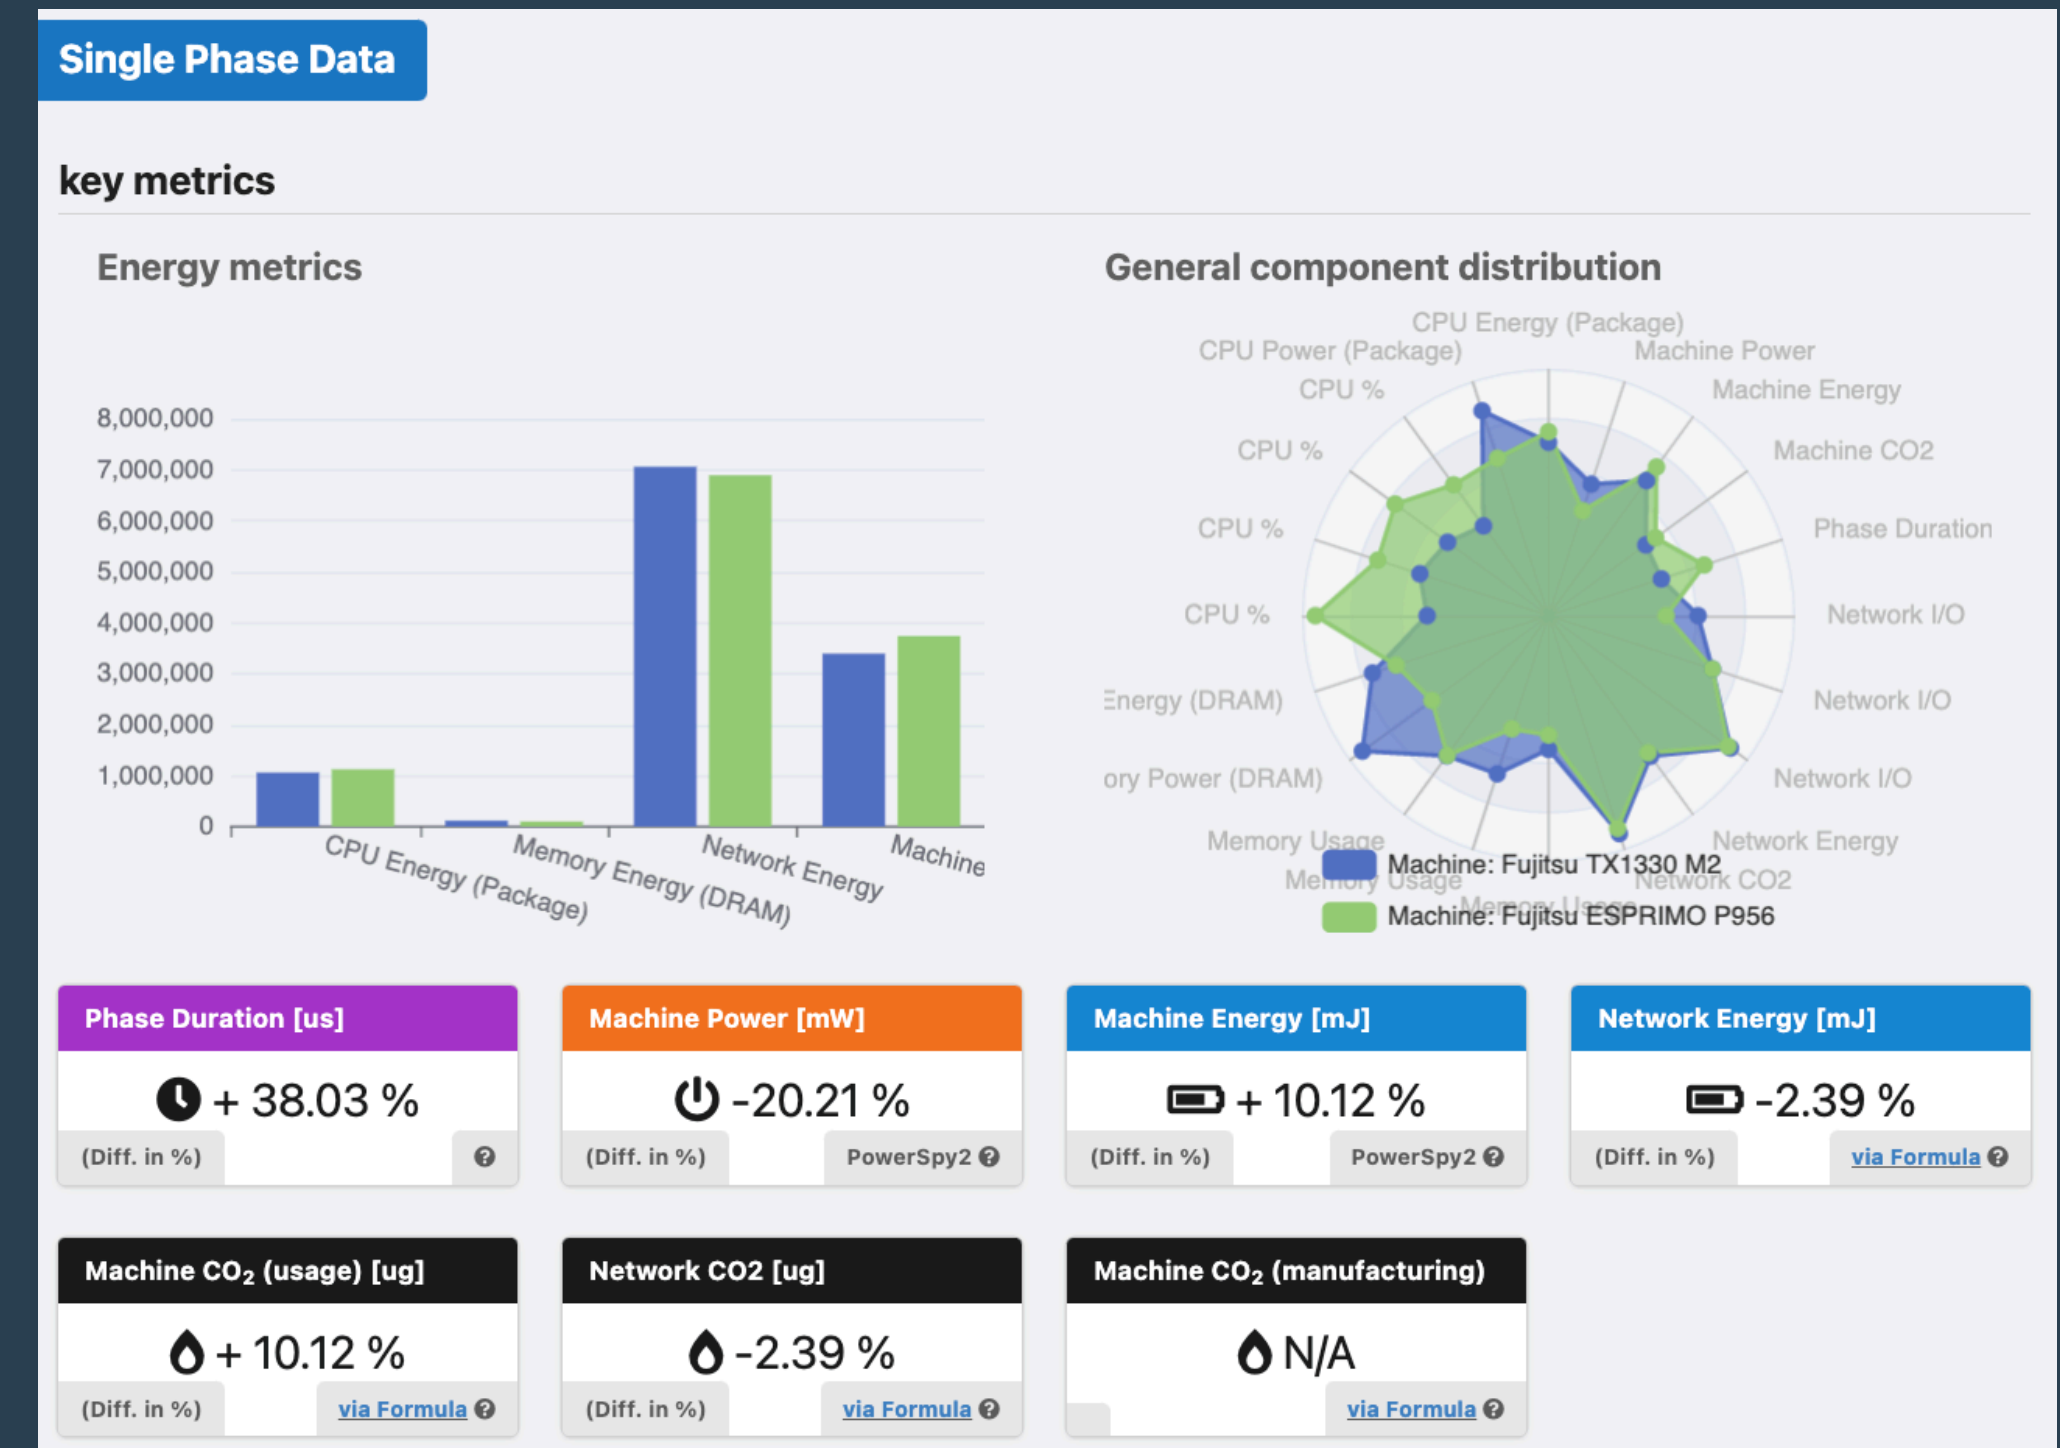


Green Metrics Tool

CO2- and Energy-Measurement

- Benchmarking via Standard-Usage-Scenarios (academia based via UCB / Öko-Insitut e.V.)
- Reproducibility / DevOps Integration via version control, orchestration and HW-Limits
- Dashboard with statistical comparison
- Modular / Extendable reporters for: Energy (Components, Machine), CO2 (Grid Intensity), Network, Memory, Temperature, Frequency, Embodied Carbon, VMs, distributed architectures ...
- SCI (ISO-Norm tbd 2024) / Blauer Engel compatible (German Eco Label)
- Cluster-Automation with Measurement-Accuracy-Control

Online Dashboard: <https://metrics.green-coding.io/index.html>




Agenda

18.01.2024

- **What quantifying tools do we have already?**
 - The basic tools Scaphandre, Kepler etc.
 - CloudCarbonFootprint, powertop
 - ecoCode / CAST
 - Green Metrics Tool, Eco-CI
 - The Green Web Foundation / Green Software Foundation
 - Did I miss anything?
- **Definition of sustainability as a proxy for green**
- **Let's try a greenfield approach**
- **Energy-ID project**
- **Summary and Outlook**

Before we try do a greenfield approach

Let's see what others are already doing



We have finally many tools That provide plain metrics output



```
Host: 13.1463 W
package core dram uncore
Socket0 13.1463 W | 10.879847 W 0.748591 W 0.071402 W

Top 5 consumers:
Power PID Exe
10.400553 W 16621 "stress"
2.08011 W 16610 "scaphandre"
0.166408 W 2786 "gnome-shell"
0.083204 W 3915 "Xwayland"
0.041602 W 4621 "guake"
```



powertop

List of hardware / os-settings tips

```
PowerTOP 2.12  Overview  Idle stats  Frequency stats  Device stats  Tunables  WakeUp
>> Bad      Enable Audio codec power management
Bad        Runtime PM for I2C Adapter i2c-2 (i915 gmbus panel)
Bad        Runtime PM for I2C Adapter i2c-3 (i915 gmbus dpc)
Bad        Runtime PM for I2C Adapter i2c-0 (i915 gmbus ssc)
Bad        Runtime PM for I2C Adapter i2c-1 (i915 gmbus vga)
Bad        Runtime PM for I2C Adapter i2c-9 (SMBus I801 adapter at efa0)
Bad        Runtime PM for I2C Adapter i2c-4 (i915 gmbus dpb)
Bad        Runtime PM for I2C Adapter i2c-5 (i915 gmbus dpd)
Bad        Runtime PM for port ata3 of PCI device: Intel Corporation 8 Series SATA Controller 1 [AHCI mode]
Bad        Runtime PM for PCI Device Intel Corporation 8 Series SATA Controller 1 [AHCI mode]
Bad        Runtime PM for port ata1 of PCI device: Intel Corporation 8 Series SATA Controller 1 [AHCI mode]
Bad        Runtime PM for disk sda
Bad        Runtime PM for disk sdb
Bad        Runtime PM for PCI Device Intel Corporation Wireless 7260
Bad        Runtime PM for PCI Device Intel Corporation 8 Series USB EHCI #1
Bad        Runtime PM for PCI Device Intel Corporation Haswell-ULT Integrated Graphics Controller
Bad        Runtime PM for PCI Device Intel Corporation 8 Series USB xHCI HC
Bad        Runtime PM for PCI Device Intel Corporation Haswell-ULT HD Audio Controller
Bad        Runtime PM for PCI Device Intel Corporation 8 Series LPC Controller
Bad        Runtime PM for PCI Device Realtek Semiconductor Co., Ltd. RTS5227 PCI Express Card Reader
Bad        Runtime PM for PCI Device Intel Corporation 8 Series HD Audio Controller
Bad        Runtime PM for PCI Device Intel Corporation 8 Series HECI #0
Bad        Runtime PM for PCI Device Intel Corporation Ethernet Connection I218-LM
Bad        Runtime PM for PCI Device Intel Corporation 8 Series SMBus Controller
Bad        Runtime PM for PCI Device Intel Corporation Haswell-ULT DRAM Controller
Bad        Runtime PM for port ata2 of PCI device: Intel Corporation 8 Series SATA Controller 1 [AHCI mode]
Good      Enable SATA link power management for host1
Good      VM writeback timeout
Good      Enable SATA link power management for host0
Good      Enable SATA link power management for host2
Good      NMI watchdog should be turned off
```

Cloud Carbon Footprint

Provides Recommendations in Demo - Do they work?

The screenshot displays the Cloud Carbon Footprint interface. At the top, there's a navigation bar with the logo and the text "Cloud Carbon Footprint". On the right side of the bar, the word "RECOMMENDATIONS" is underlined. Below the navigation bar, there are three filter boxes: "Cloud Providers: 4 of 4", "Accounts: 16 of 16", and "Regions: 9 of 9". Below these filters, there are two tabs for "CO2e Units": "metric tons" (which is selected and highlighted in blue) and "kilograms".

Cloud Provider	Account Name	Region	Recommendation Type
AWS	aws account 3	us-west-2	Delete
AWS	aws account 4	us-west-1	Delete
GCP	gcp account 0	us-west1	DELETE_IMAGE
GCP	gcp account 1	us-west1	SNAPSHOT_AND_DEL...
GCP	gcp account 2	us-west1	CHANGE_MACHINE_T...
GCP	gcp account 3	us-east1	DELETE_ADDRESS
GCP	gcp account 4	us-west2	DELETE_DISK
GCP	gcp account 2	us-east1	STOP_VM
GCP	gcp account 4	us-east1	STOP_VM

The right side of the interface shows a detailed view of a recommendation. The "Recommendation Details" section includes:

- Cloud Provider:** GCP
- Account Name:** gcp account 4
- Account ID:** gcp account 4
- Region:** us-east1
- Resource Name:** test-instance-9
- Resource ID:** 8928403120086348000

The "Recommendation Type" is "STOP_VM". The "Recommendation Detail" states: "Save cost by performing a STOP_VM for instance: test-instance-9." At the bottom of the details panel, there are three columns for savings:

Cost Savings (USD)	CO2e Savings (Metric Tons)	Energy Savings (kilowatt hours)
0	0	0

Anwendungsfall Web: Statische Analyse

Eco-Code / CAST etc.

- Based on recommendations from science and the French position paper Numérique Responsable
- No proven effect in larger software projects
- Recommendations such as `i++` instead of `++i` are often identical even in assembly code / byte code
- Overhead of the tools unclear

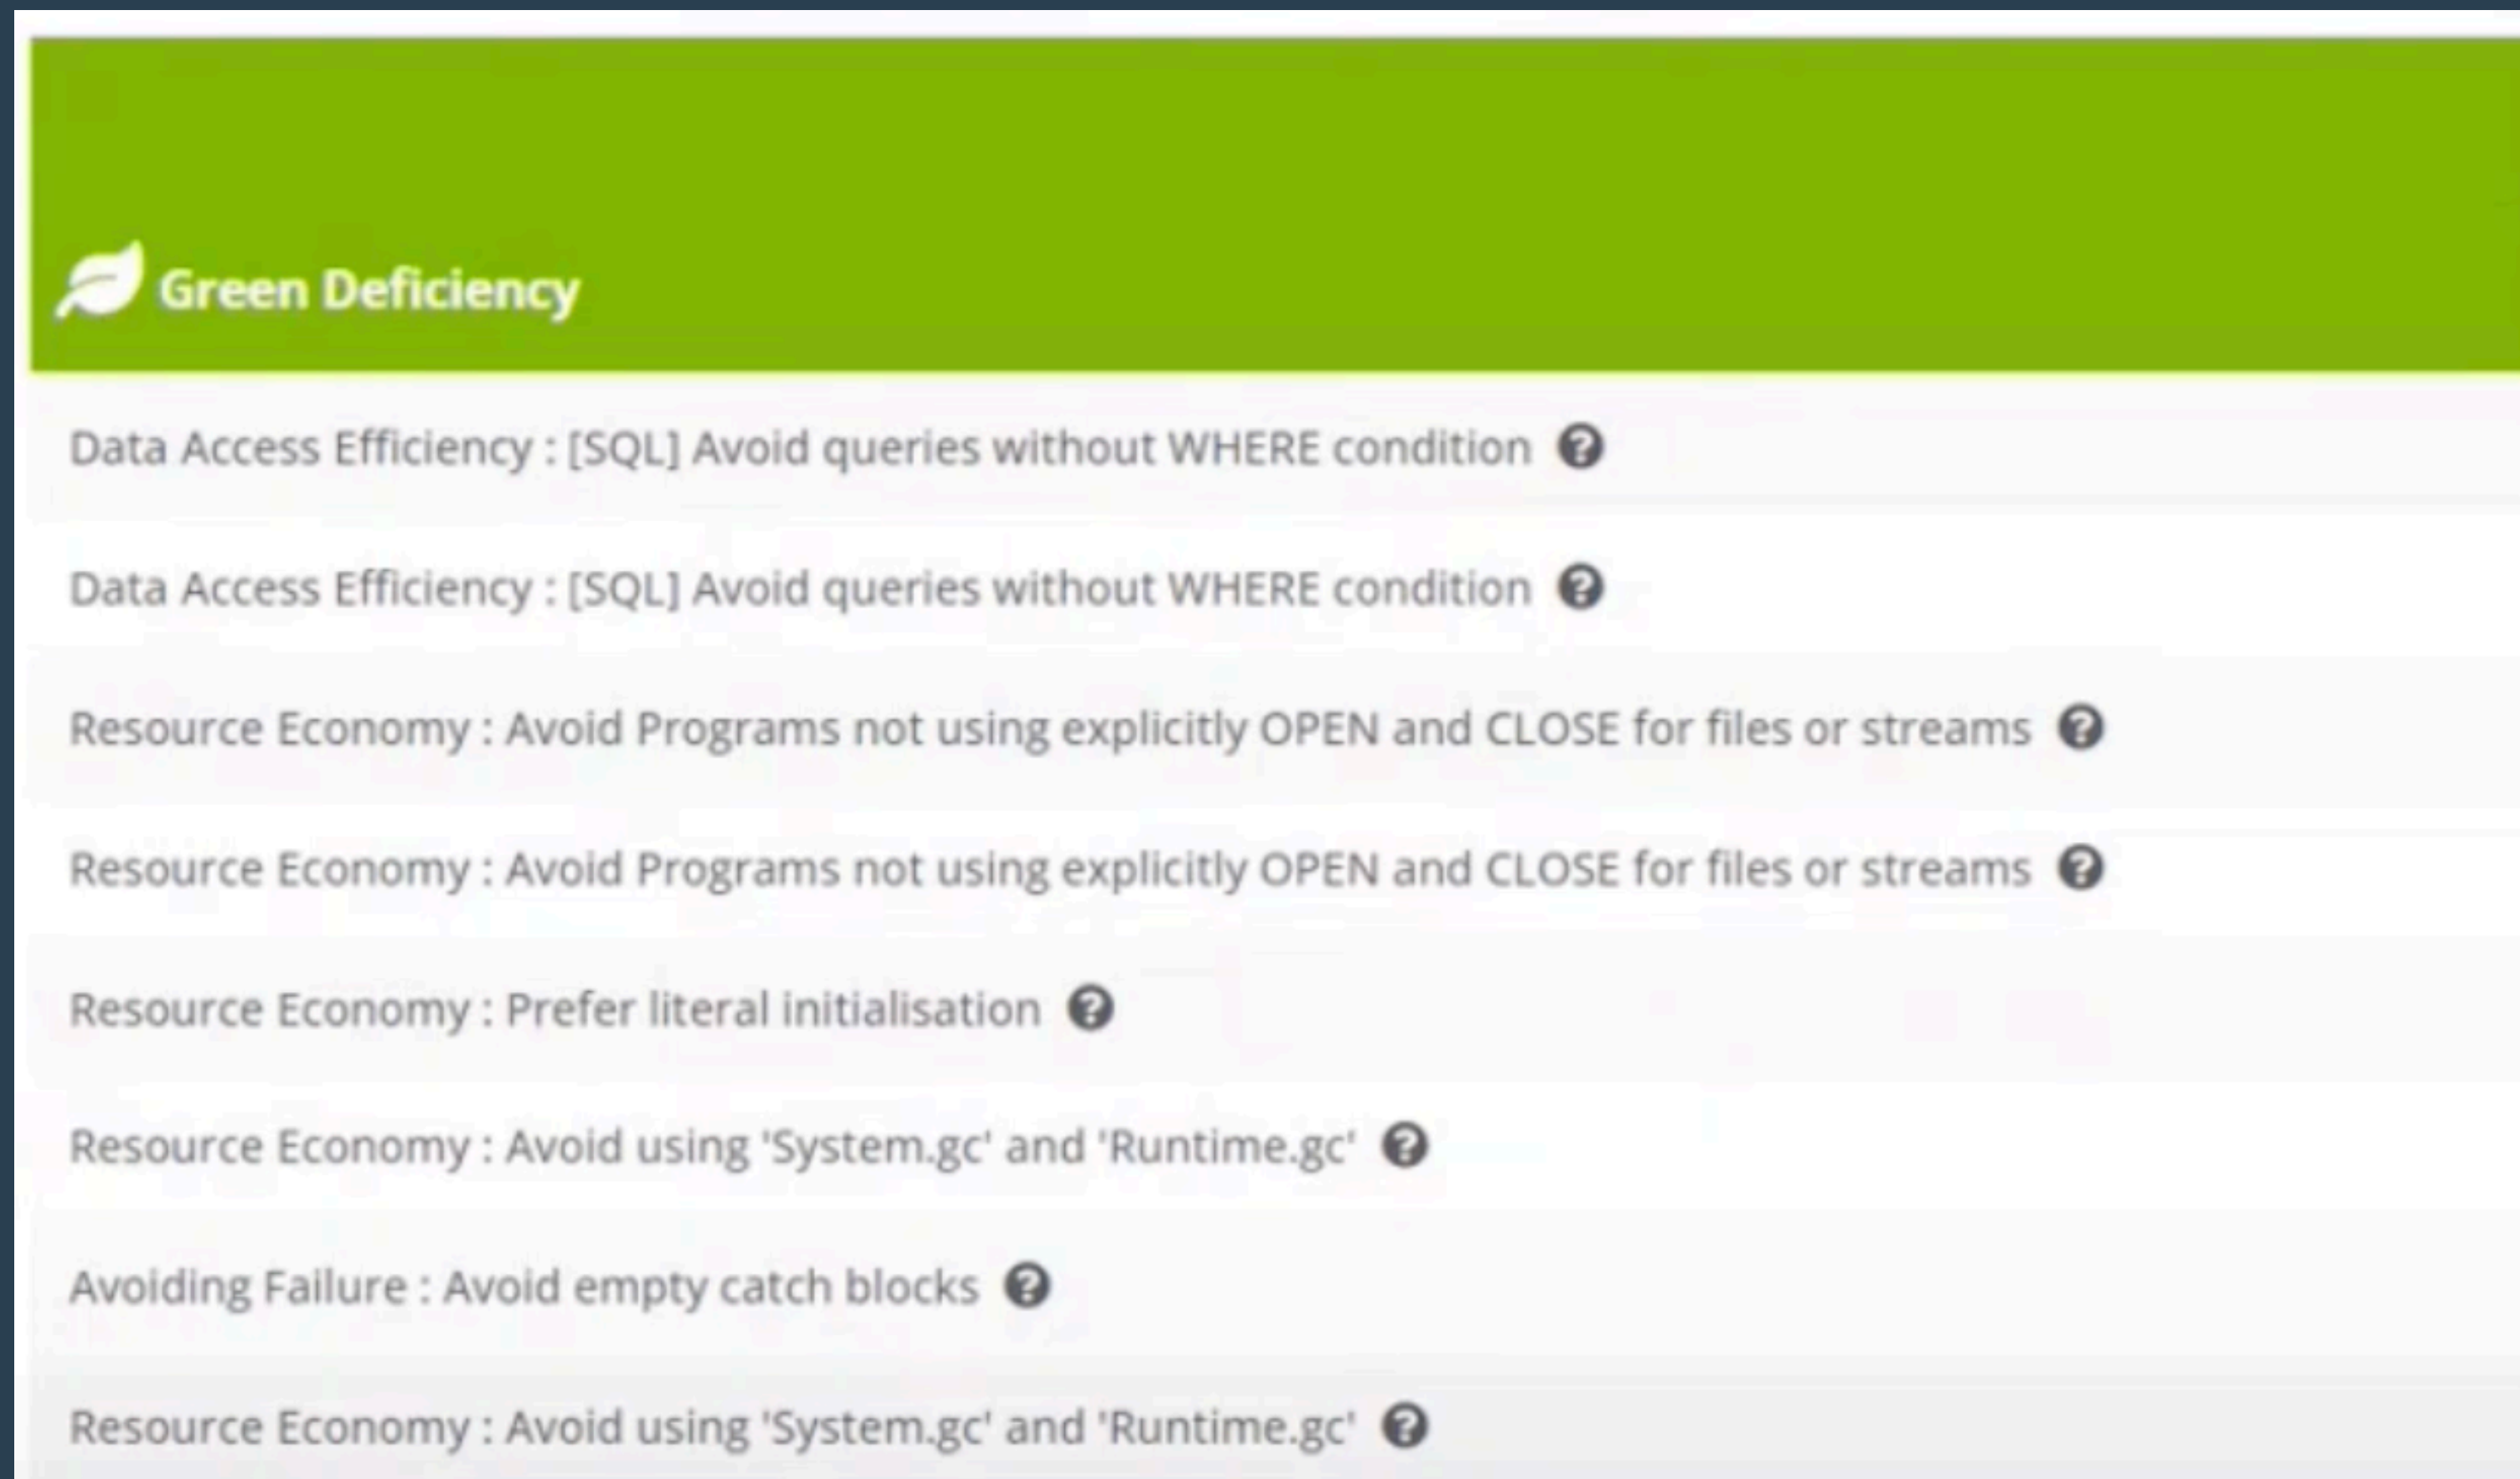


The screenshot shows the GitHub profile for the Green Code Initiative. The profile picture is a green square with a white leaf and blue code brackets. The name is "Green Code Initiative" with the tagline "Let's write low-carbon software". It has 84 followers, is located in France, and has a LinkedIn link to "company/green-code-initiative". Below the profile is a code snippet in JSON format:

```
1  {
2    "title": "Limit the number of returns for a SQL query",
3    "type": "CODE_SMELL",
4    "code": {
```

CAST

Static Code Scanner for Enterprise



Green Deficiency

- Data Access Efficiency : [SQL] Avoid queries without WHERE condition ?
- Data Access Efficiency : [SQL] Avoid queries without WHERE condition ?
- Resource Economy : Avoid Programs not using explicitly OPEN and CLOSE for files or streams ?
- Resource Economy : Avoid Programs not using explicitly OPEN and CLOSE for files or streams ?
- Resource Economy : Prefer literal initialisation ?
- Resource Economy : Avoid using 'System.gc' and 'Runtime.gc' ?
- Avoiding Failure : Avoid empty catch blocks ?
- Resource Economy : Avoid using 'System.gc' and 'Runtime.gc' ?



Green Impact
Programming practices and engineering principles that make software more environment-friendly

72.0

WORST	AVERAGE	BEST
36.2	72.0	88.5

More tools ... did we miss any?

Happy to talk about in Q&A !

- powertop
- powermetrics
- [GreenFrame.io](#)
- turbostat
- powerJoularX
- PAPI
- ...


Hmm, tools don't really help atm

Maybe labels, organizations and norms?

The Green Web Foundation

Green Energy Label for Datacenters / Servers

- The Green Web Foundation is pursuing and accelerating the transition to a fossil-free Internet
- Provides datasets/API for green hosting
- Label allows to show that your hosting is green
- If you are either connected to a green power plant or you have on-site generation or you buy PPAs



Green Software Foundation 1/2

SCI - An approach to quantify a software



$$SCI = (E * I) + M \text{ per } R$$

- (E) - Energy consumption (kWh) for various components:
 - Ex. CPU/GPUs, Data storage, Memory, Network
- (I) - Emission factor
- (M) - Embodied emissions
 - z.B. Data from servers, laptops, mobile devices, tablets etc.

Green Software Foundation 2/2

<https://patterns.greensoftware.foundation/>



- A collection of many tips from either common sense or performance engineering
- Many can also backfire
 - "Run AI models at the edge"
 - "Reduce transmitted data"
- Many provide no guidance
 - "Use energy efficient AI models"

Cache static data

Choose the region that is closest to users

Compress transmitted data

Delete unused storage resources

Encrypt what is necessary

Evaluate other CPU architectures

Terminate TLS at border gateway

Implement stateless design

Match your service level objectives to business needs

Match utilization requirements of virtual machines (VMs)

CNCF - TAG ENV

Cloud Native Computing Foundation eco group



- For instance, programs written in energy efficient languages or running on more optimized runtimes are generally “greener”.
 - So we use C and we are green?
- Automated scaling
- Only schedule pods when green energy is there

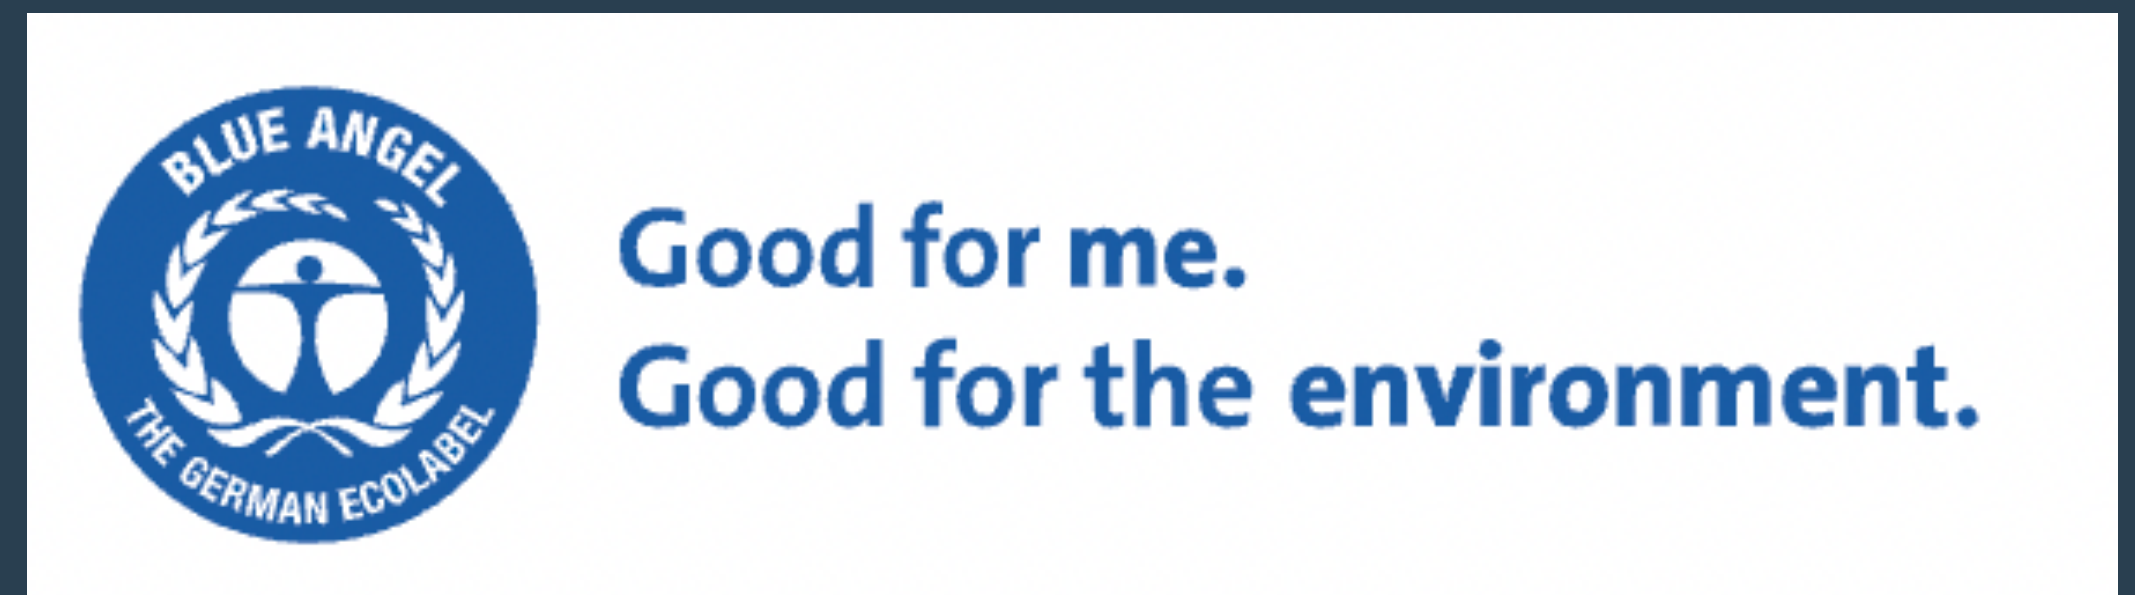
- **Green System Architecture**
- **Current Sustainable Cloud Computing Landscape**
 - **Data centers**
 - **Smart Data Centers**
 - **Cooling / BMC**
 - **Methodologies**
 - **Measurement Methodologies**
 - **Observability Methodologies**
 - **Observability Tooling**
 - **Infrastructure Tooling**
 - **Scheduling At The Cluster Level**
 - **Scaling**
 - **On-Node Power Management Tuning**

Source: <https://tag-env-sustainability.cncf.io/landscape/#green-system-architecture>

Blue Angel - German Eco Label 1/2

Ressource- and energy-efficient software products (DE-UZ 215)

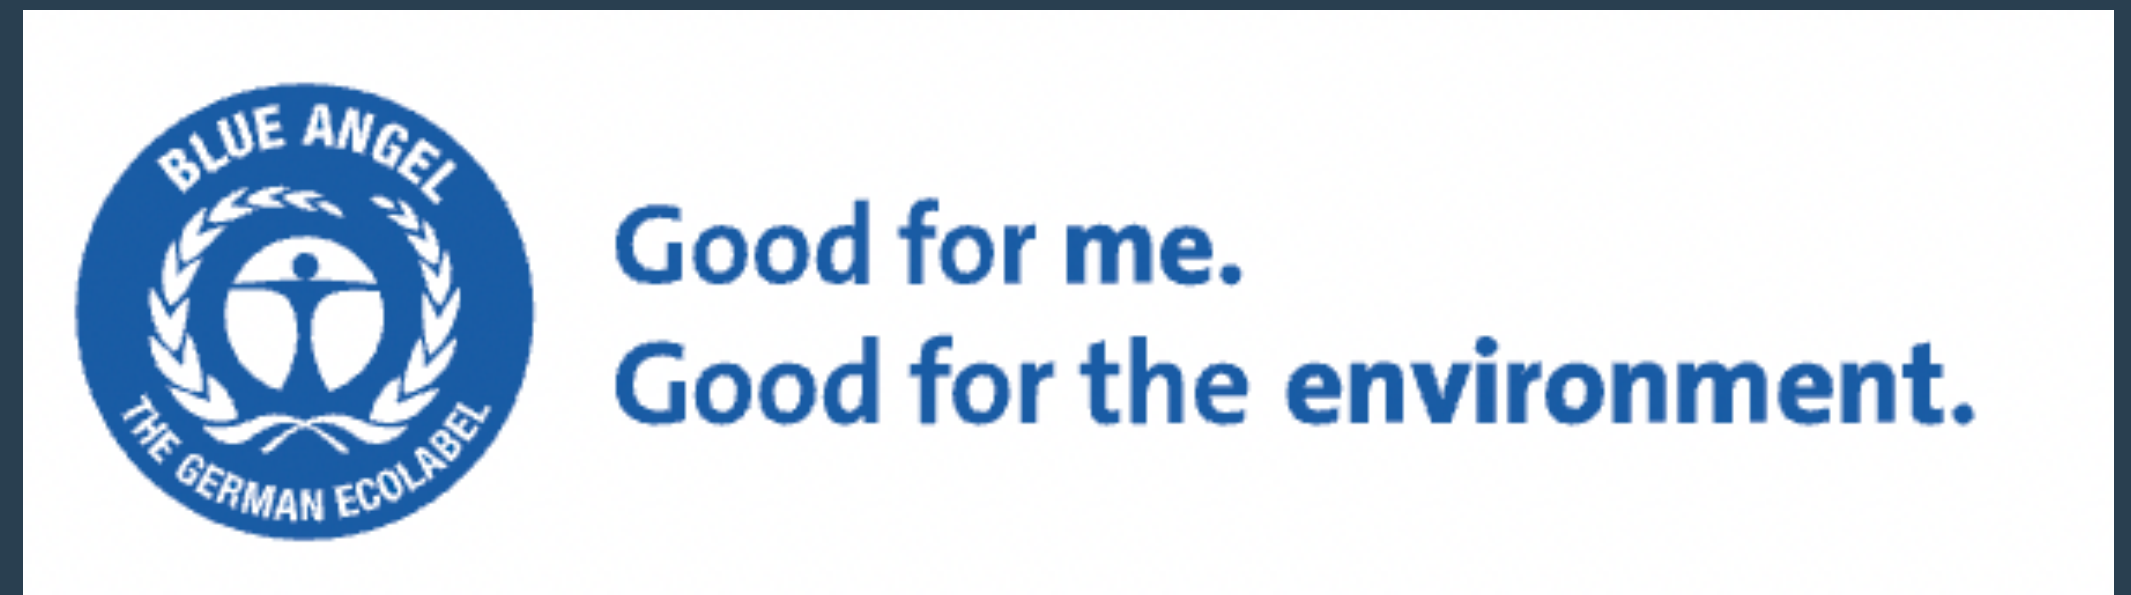
- Software certificate from the (quite ?) well-known German eco-label
- Developed in collaboration with academia and research institutions in Germany
- Currently only for desktop applications without online functionality 🤪
- Does require measurement, but no max resource consumption
- No technical reference implementation
- Criterias -> Page 2



Blue Angel - German Eco Label 2/2

Ressource- and energy-efficient software products (DE-UZ 215)

- Code should be reusable
- Code should run on 5 year old hardware
- Code should survive a Hibernate
- Code should be uninstallable without leaving files behind
- Code should not contain advertising
- Code should be documented
- ...



Normen

Mostly focus on life cycle assessment


- ISO 140001
 - The standard for a life cycle assessment.
But requires methodology reference for software
- GHG Protocol - Chapter 6
 - Concrete recommendations for software life cycle assessment. Even in VMs.
- So do we qualify as green if we just quantify?




And so much more ...

Did we miss anything?

Website



Application




- 1. Gibt es ein definiertes Vorgehensmodell bei der Softwareentwicklung?
- 1. Werden die Stories konsequent vom Project Owner priorisiert?
- 1. Ist ein fachlicher Domänenspezialist unter den Entwicklern?
- 1. Sind die Stories klar definiert und strukturiert?

Sadly we did not find the green-ness yet

Let's maybe take a step back and try to define green through sustainability as a proxy

What is the definition of sustainability ...

in a more general way



What actually is sustainability

The typical display via an ESG-like approach



Source: <https://www.anevis-solutions.com/>

What actually is sustainability

ESG sees sustainability as a "risk" to mitigate. Alternative view:

ESG vs Sustainability

Two distinctive perspectives

Earth Systems and Society: how to operate within a 'safe space for humanity' based upon environmental thresholds and societal foundations

Quantifying the status of threshold issues such as climate, inequality and nature

Quantifying the current performance of business and finance with respect to the safe operating space

Identifying the societal value at risk and non-financial impact of current financial performance

Sustainability:
'Inside-out'
focus on earth
and social systems

Business & Finance: how sustainability concerns impact enterprise value at risk

Quantifying impact of sustainability concerns on financial performance

Identifying financial investments needed to improve non-financial performance

Identifying forward risks of current business model with respect to safe operating space

ESG: 'Outside-in'
focus on enterprise
value

Source: <https://cisweb.lancaster.ac.uk/>

Sustainability - As used in this talk

Based on the definition of the UN-SDGs

- The Brundtland report from the **United Nations (UN)** defines sustainable development as the ability to :
 - *“meet the present needs without compromising the future generation abilities for their own needs“*
- General understanding often says:
 - *... the ability to refill itself at a quicker rate than it is consumed / damaged ...*

Let's do an intermediate summary

Of goods and bads so far

What is good and bad

of what we have seen so far

- Some tips mostly everybody would agree this is a green practice
 - Green Hosting. Deleting VMs you do not use. Deleting unused data ...
- Some tips require proper implementation and can backfire:
 - Location- / Time-Shifting; Switching to other algorithm; Running at edge
- Some have no data to back their claims:
 - Use `i++` instead of `++i` ...

Let's do a greenfield approach then

What would we want from green code?

Greenfield approach

What would be the most green we can think of?

- In a truly "sustainable" code it would be
 - code that consumes only so much resources as we can provide "renewably"
 - or software that we can provide for future generations in the same way
- This means we can only use so much hardware as we can also recycle
- And we can only use regenerative power
- Furthermore we would have to distribute the allowed amounts globally for every software product

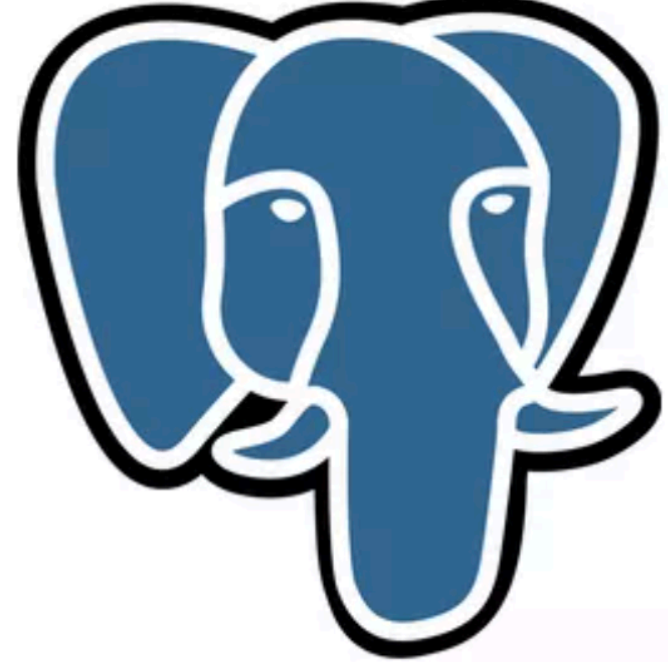
- Ok ... that is a "stretch goal" for now :)

But can we find a best in class

And create a gold standard?

Energy-ID project

Using standardized functional benchmarks for software



PostgreSQL


Relational Database

[click for details](#)

BADGES

- Energy Cost 58.93 kJ via PSU (AC)
- Energy Cost 20.59 kJ via RAPL
- SCI 29.46 mgCO2e/TPC-C SQL-op

[Show measurements](#)



MariaDB


Relational Database

[click for details](#)

BADGES

- Energy Cost 59.55 kJ via PSU (AC)
- Energy Cost 21.82 kJ via RAPL
- SCI 163.92 mgCO2e/TPC-C SQL-op

[Show measurements](#)



Wagtail


CMS

[click for details](#)

BADGES

- Energy Cost 1.14 kJ via PSU (AC)
- Energy Cost 243.24 J via RAPL
- SCI 18.84 mgCO2e/page request

[Show measurements](#)



Wordpress

Blog/CMS

[click for details](#)


BADGES

- Energy Cost 635.62 J via PSU (AC)
- Energy Cost 112.62 J via RAPL
- SCI 25.08 mgCO2e/page request

[Show measurements](#)

Energy-ID project

Using standardized functional benchmarks for software



Future works of Energy ID

Implementation into the GMT

- Use best practices as we have for running code
 - Use resource planning
 - Check how software reacts to carbon signals
 - Check with AI models for energy hotspots.
 - Identify LoC in different software iterations - Flamegraph
 - Apply AI test-optimizations and see if software gets better
- Monitor software over time if feature-addition impacts energy
- Goal: Identify what is the current most efficient code and create a **gold standard**

Thank you

Want to know more

- Website / Blog / Newsletter: <https://www.green-coding.io>
- Demo Open Data Repository: <https://metrics.green-coding.io>
- Our projects: <https://www.green-coding.io/#projects>
- Our case-studies: <https://www.green-coding.io/case-studies>
- Meetup Group (Berlin): <https://www.meetup.com/green-coding>
- <https://www.linkedin.com/in/arne-tarara> / arne@green-coding.io

Eco-Compute conference

The first engineering conference on sustainability in hardware & software

- We are looking for talks!
- Case studies with clients that save carbon / energy
- New tools and practices in Green Coding
- Hardware / Software interaction for efficiency
- ... <https://www.eco-compute.io/call-for-speakers/>

EcoCompute Conference 2024

The first engineering conference on
sustainability in hardware & software

April 25 - 26

Munich, Germany




<https://www.eco-compute.io/call-for-speakers/>

Backup Slides

Zusammenspiel von Hardware und Software


Datenbanken für Energiekennlinien durch Hardware und Konfiguration

High Idle, but almost linear




ASUSTeK Computer Inc. RS720Q-E9-RS8 (2019)

50% Power increase at 70% utilization



Hewlett Packard Enterprise ProLiant DL110 Gen10 Plus

Idle optimized

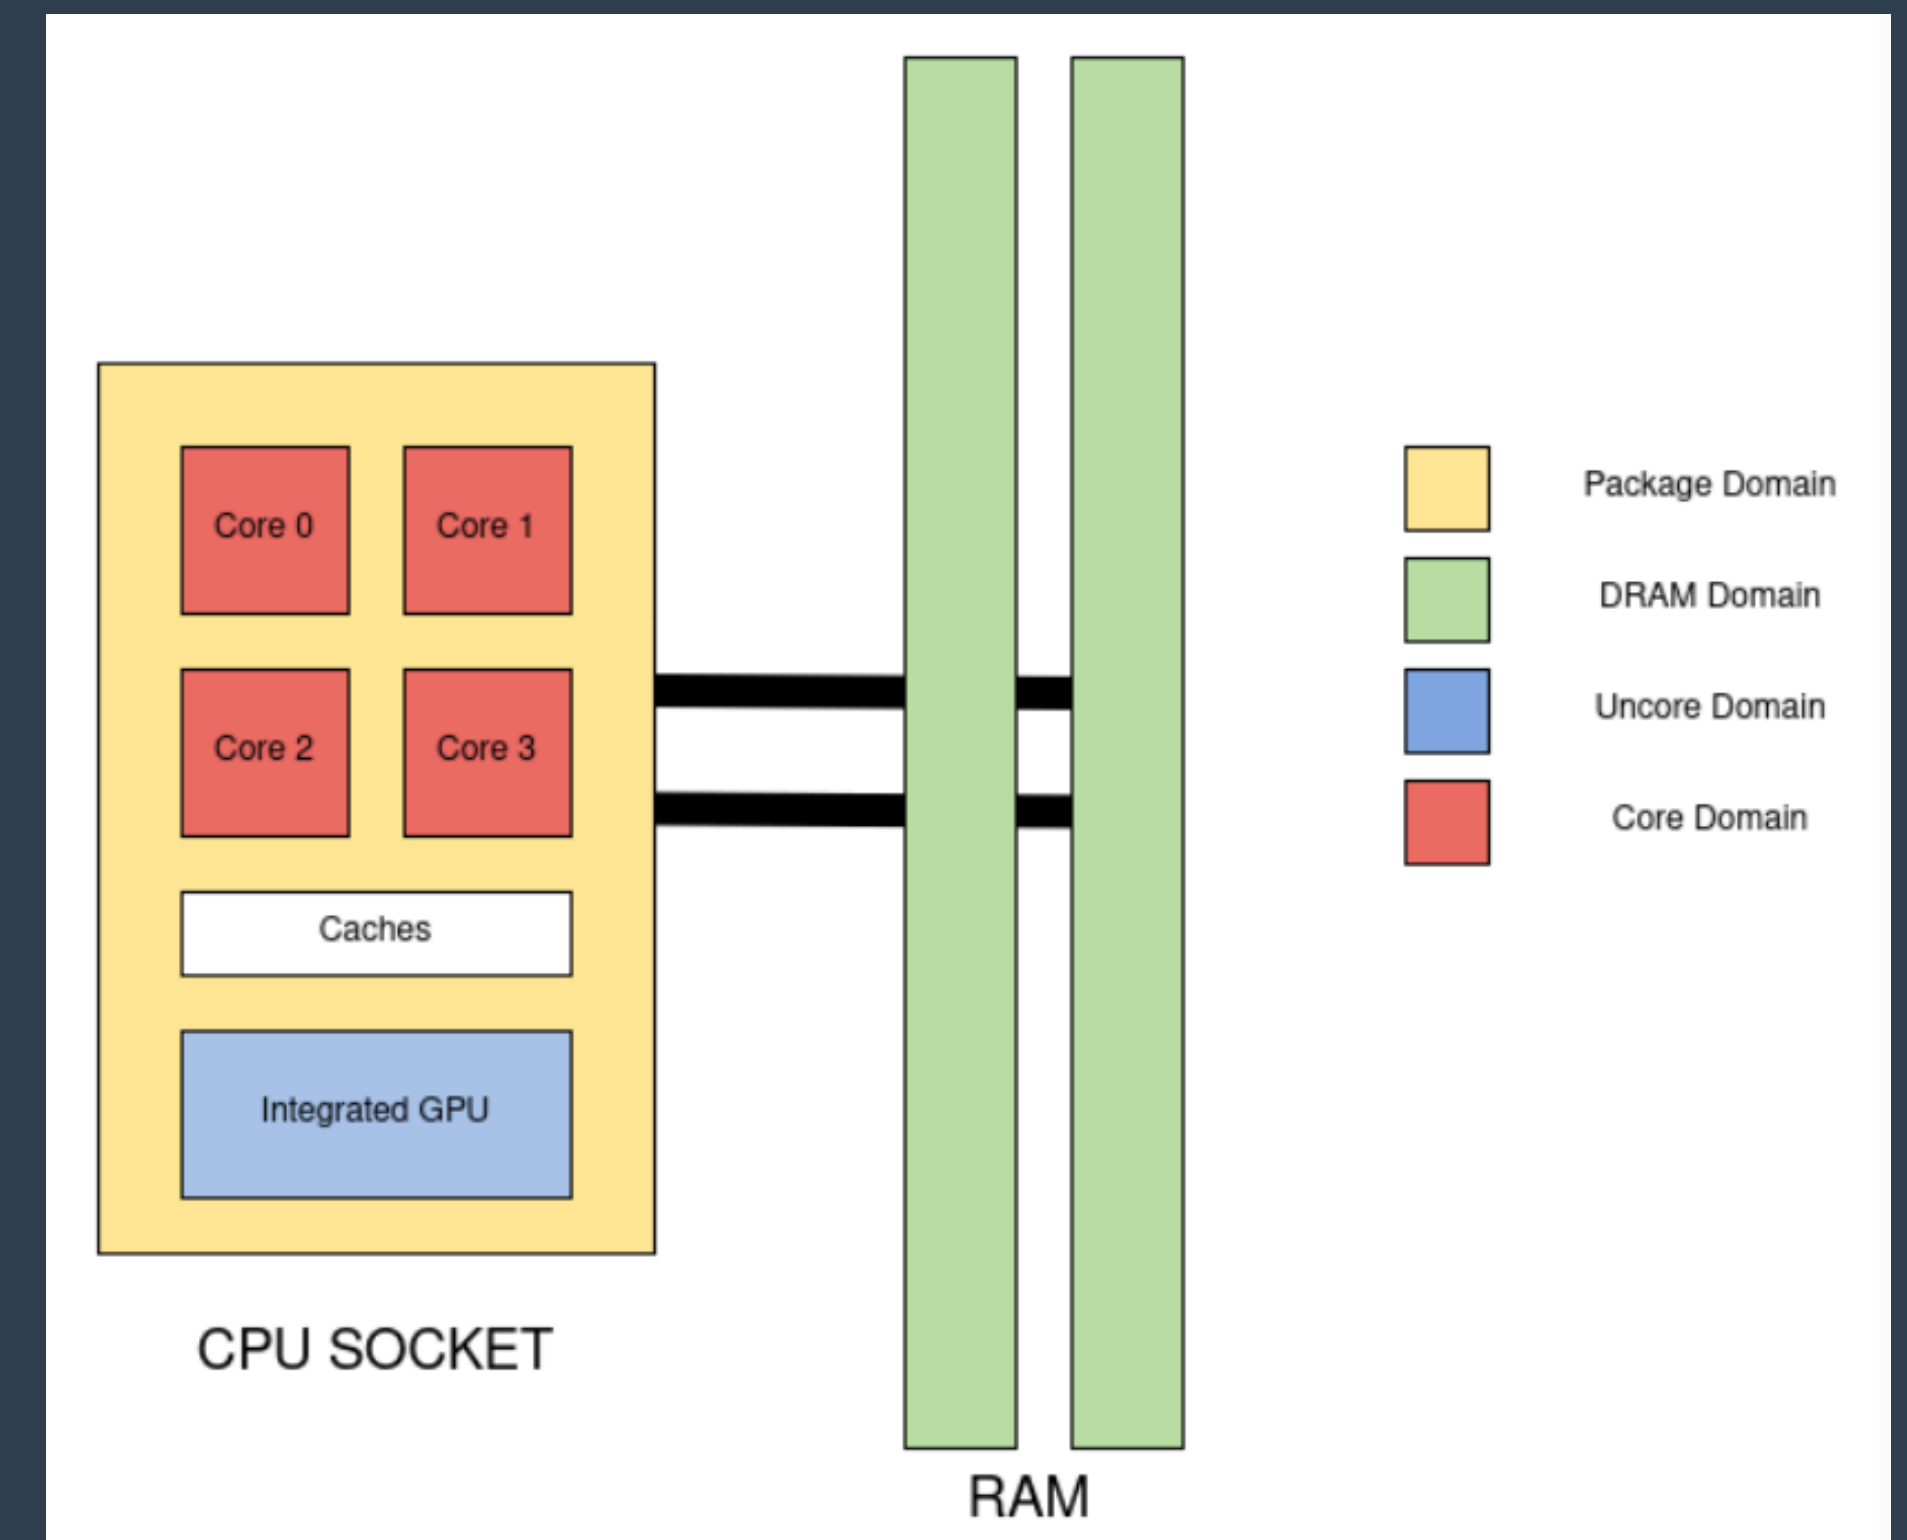


QuantaGrid D43K-1U (2022)

Details on RAPL

The most used technology atm

- Energy measurement capabilities on most modern Intel/AMD processors
- Measure:
 - CPU Energy per Core / Package
 - RAM
 - Integrated GPU
- Software model of capacitor readings on mainboard
 - Resolution 1ms / 15.3 microJoules
- Exposed in Linux kernel through device



Source: https://pyjoules.readthedocs.io/en/stable/devices/intel_cpu.html

GREEN CODING;