

Green Coding

programmier.bar meetup

 **GREEN CODING;**

Who are we

Green Coding Solutions GmbH

- LLC from Germany - 2 founders - 4 ppl. team
- Founders: Arne / Didi - Electrical Engineer / Software-Dev
- Main fields:
 - R&D - Academia / Governments / Public Institutions / Open-source
 - Company sustainability transformations - Consulting and integrations



Agenda

04.04.2024

- **Background - CO2 and energy consumption of ICT**
- **Definition: What is Green Coding**
- **Moving parts: What do we need to make code use less CO2?**
- **Overview of some tools**
 - Green Metrics Tool
 - Cloud Energy / Eco-CI
 - Power Hog
 - Software Lifecycle Assessment
- **Potentials and examples**
- **Optional Addons / Teaser**
 - How do we measure energy? Hardware / Software / OS (Windows, macOS, Linux)
 - How do we estimate in the cloud?
 - Organizations and Software certifications available (Blue Angel / SCI / ISO 14001)

Background

Some fundamental data and GHG concepts

Background

Why is digital sustainability relevant? - ICT sector prognosis




Figure : Two charts comparing (left) the greenhouse gas emissions of the aviation industry in blue with the ICT sector in green, and (right) projected estimates for emissions from the ICT sector by 2050 if nothing changes. Data is from ACM's 2021 Technology Policy Council report. (Image from KDE published under a [CC-BY-SA-4.0](#) license. Airplane icon by Simon Child and IT icon by Sari Braga licensed under a [CC-BY](#) license. Design by Lana Lutz.)

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 ...*

Background info

Where do ICT emissions come from exactly?

- Component drilldown

- Servers
- End-User-Devices
 - Computers
 - Mobile
 - Screens
 - ...
- Datacenters
- Network infrastructure

- Domain drilldown

- Energy consumption (Usage phase)
 - and thus CO2 emissions through energy generation
- Embodied Emissions
 - CO2
 - Water use
 - Land use
 - Toxic metals
 - ...
- End-of-Life - Recycling / Electronic waste

What is Green Coding?

Many organisations have different definitions

Definition from academia

Junger et. al - 2023

- **Green By IT:** [...] (e. g. smart grids, building automation, environmental management information systems (EMIS) like life-cycle assessment software (LCA), etc.) [...]
- **Green IT** deals with making ICT itself more sustainable (e. g. follow-the-sun load shifting, backward compatible software, designing energy- and resource-efficient hardware, etc.)
- **Green Coding:** Considered part of 'Green IT', concerned with the sustainability of the software itself, focusing on methods and tools that can be leveraged during development to impact software sustainability in all of its life-cycle phases (development, usage, and disposal phase).

Green Coding Solutions

Green Coding definition

- Making Energy and CO2 a functional KPI
- Increase transparency
- Efficiency through performance engineering and AI
- Reduce resource consumption / Increase utilization
- Prolong hardware change cycles
- Software developers contribution to climate change and sustainability
- **Summary:** Many parts are already known. However their composition is becoming a separate domain. Similar to DevOps which also distinct responsible domain.

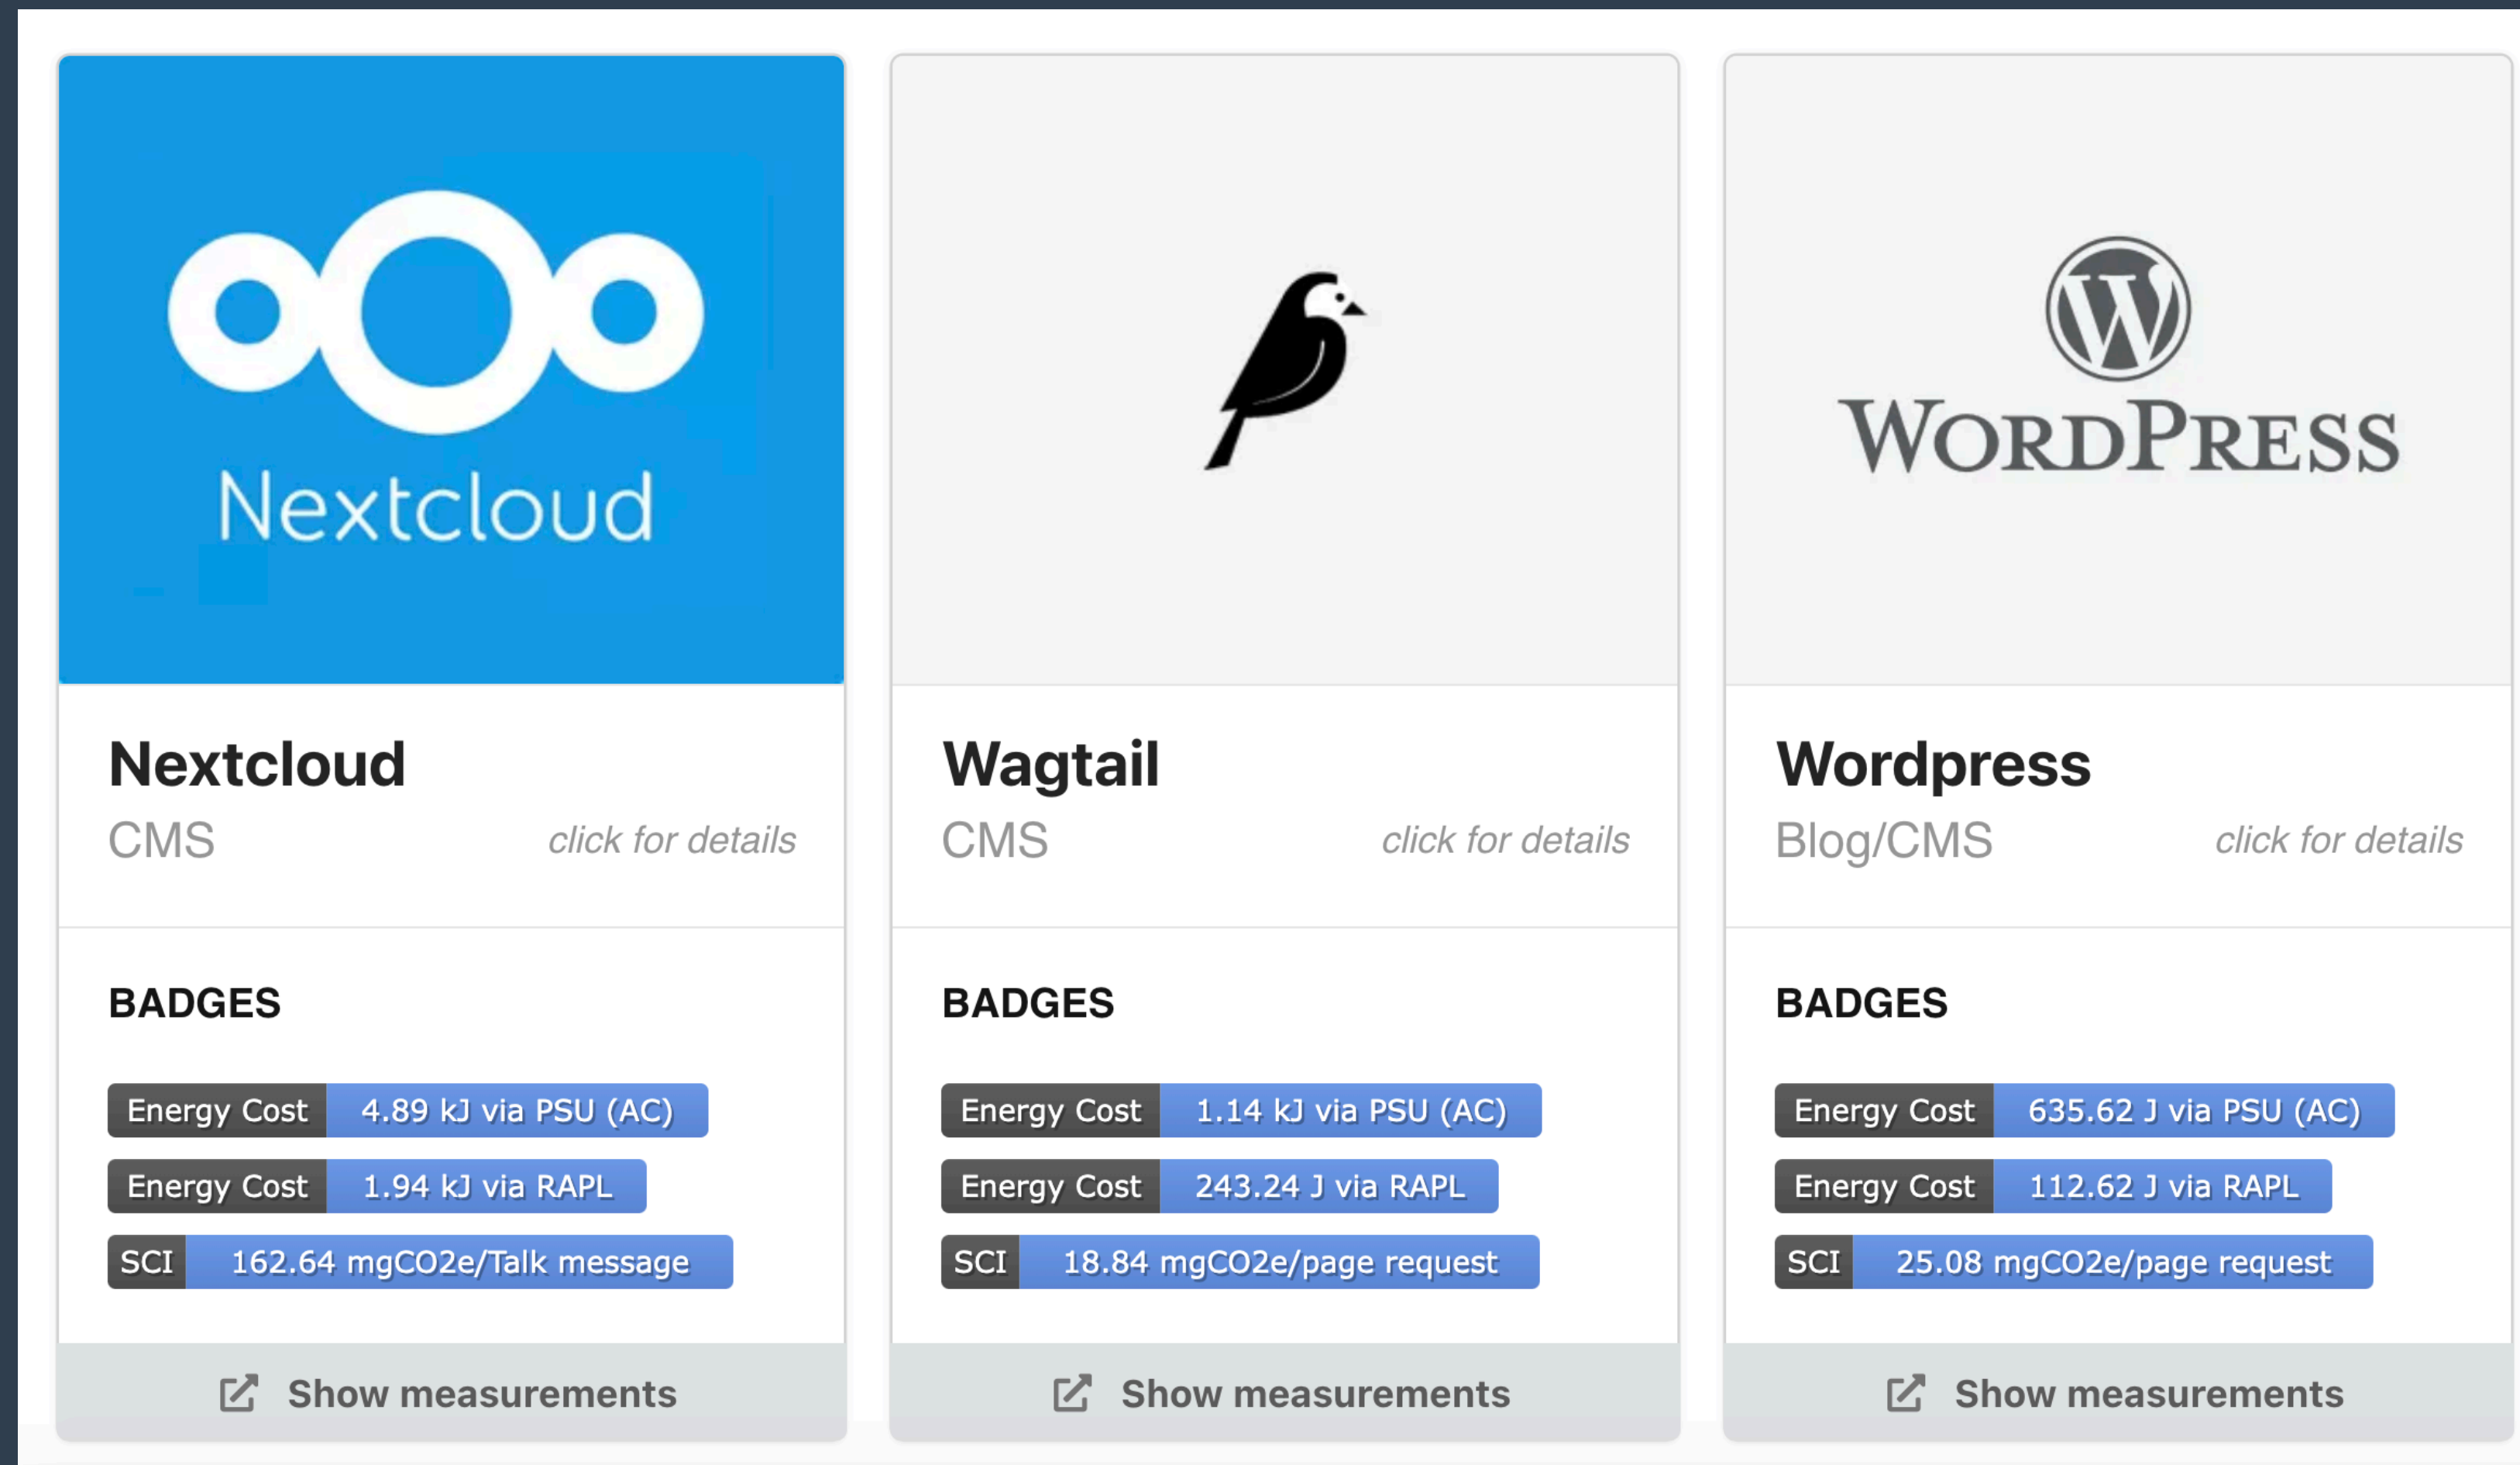


Moving parts

What do we need in order to quantify and optimize software for less carbon consumption?

This is where we want to be

Making a claim in CO2 for a software and it's use case



SCI by the Green Software Foundation

ISO Norm (tbd 2024) - Only includes the runtime

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

- (E) - Energy consumption (kilowatt hours) for different components:
 - Ex. CPU/GPUs, Data storage, Memory, Network
- (I) - Emissions factors
- (M) - Embodied emissions
 - Ex. data for servers, mobile devices and laptops
- (R) - Unit of work / use case

How to get embodied emissions

Using Life-Cycle-Assessment databases

- **Boavizta**

<https://dataviz.boavizta.org/manufacturereadata>

- **Microsoft**

<https://tco.exploresurface.com/sustainability/calculator>

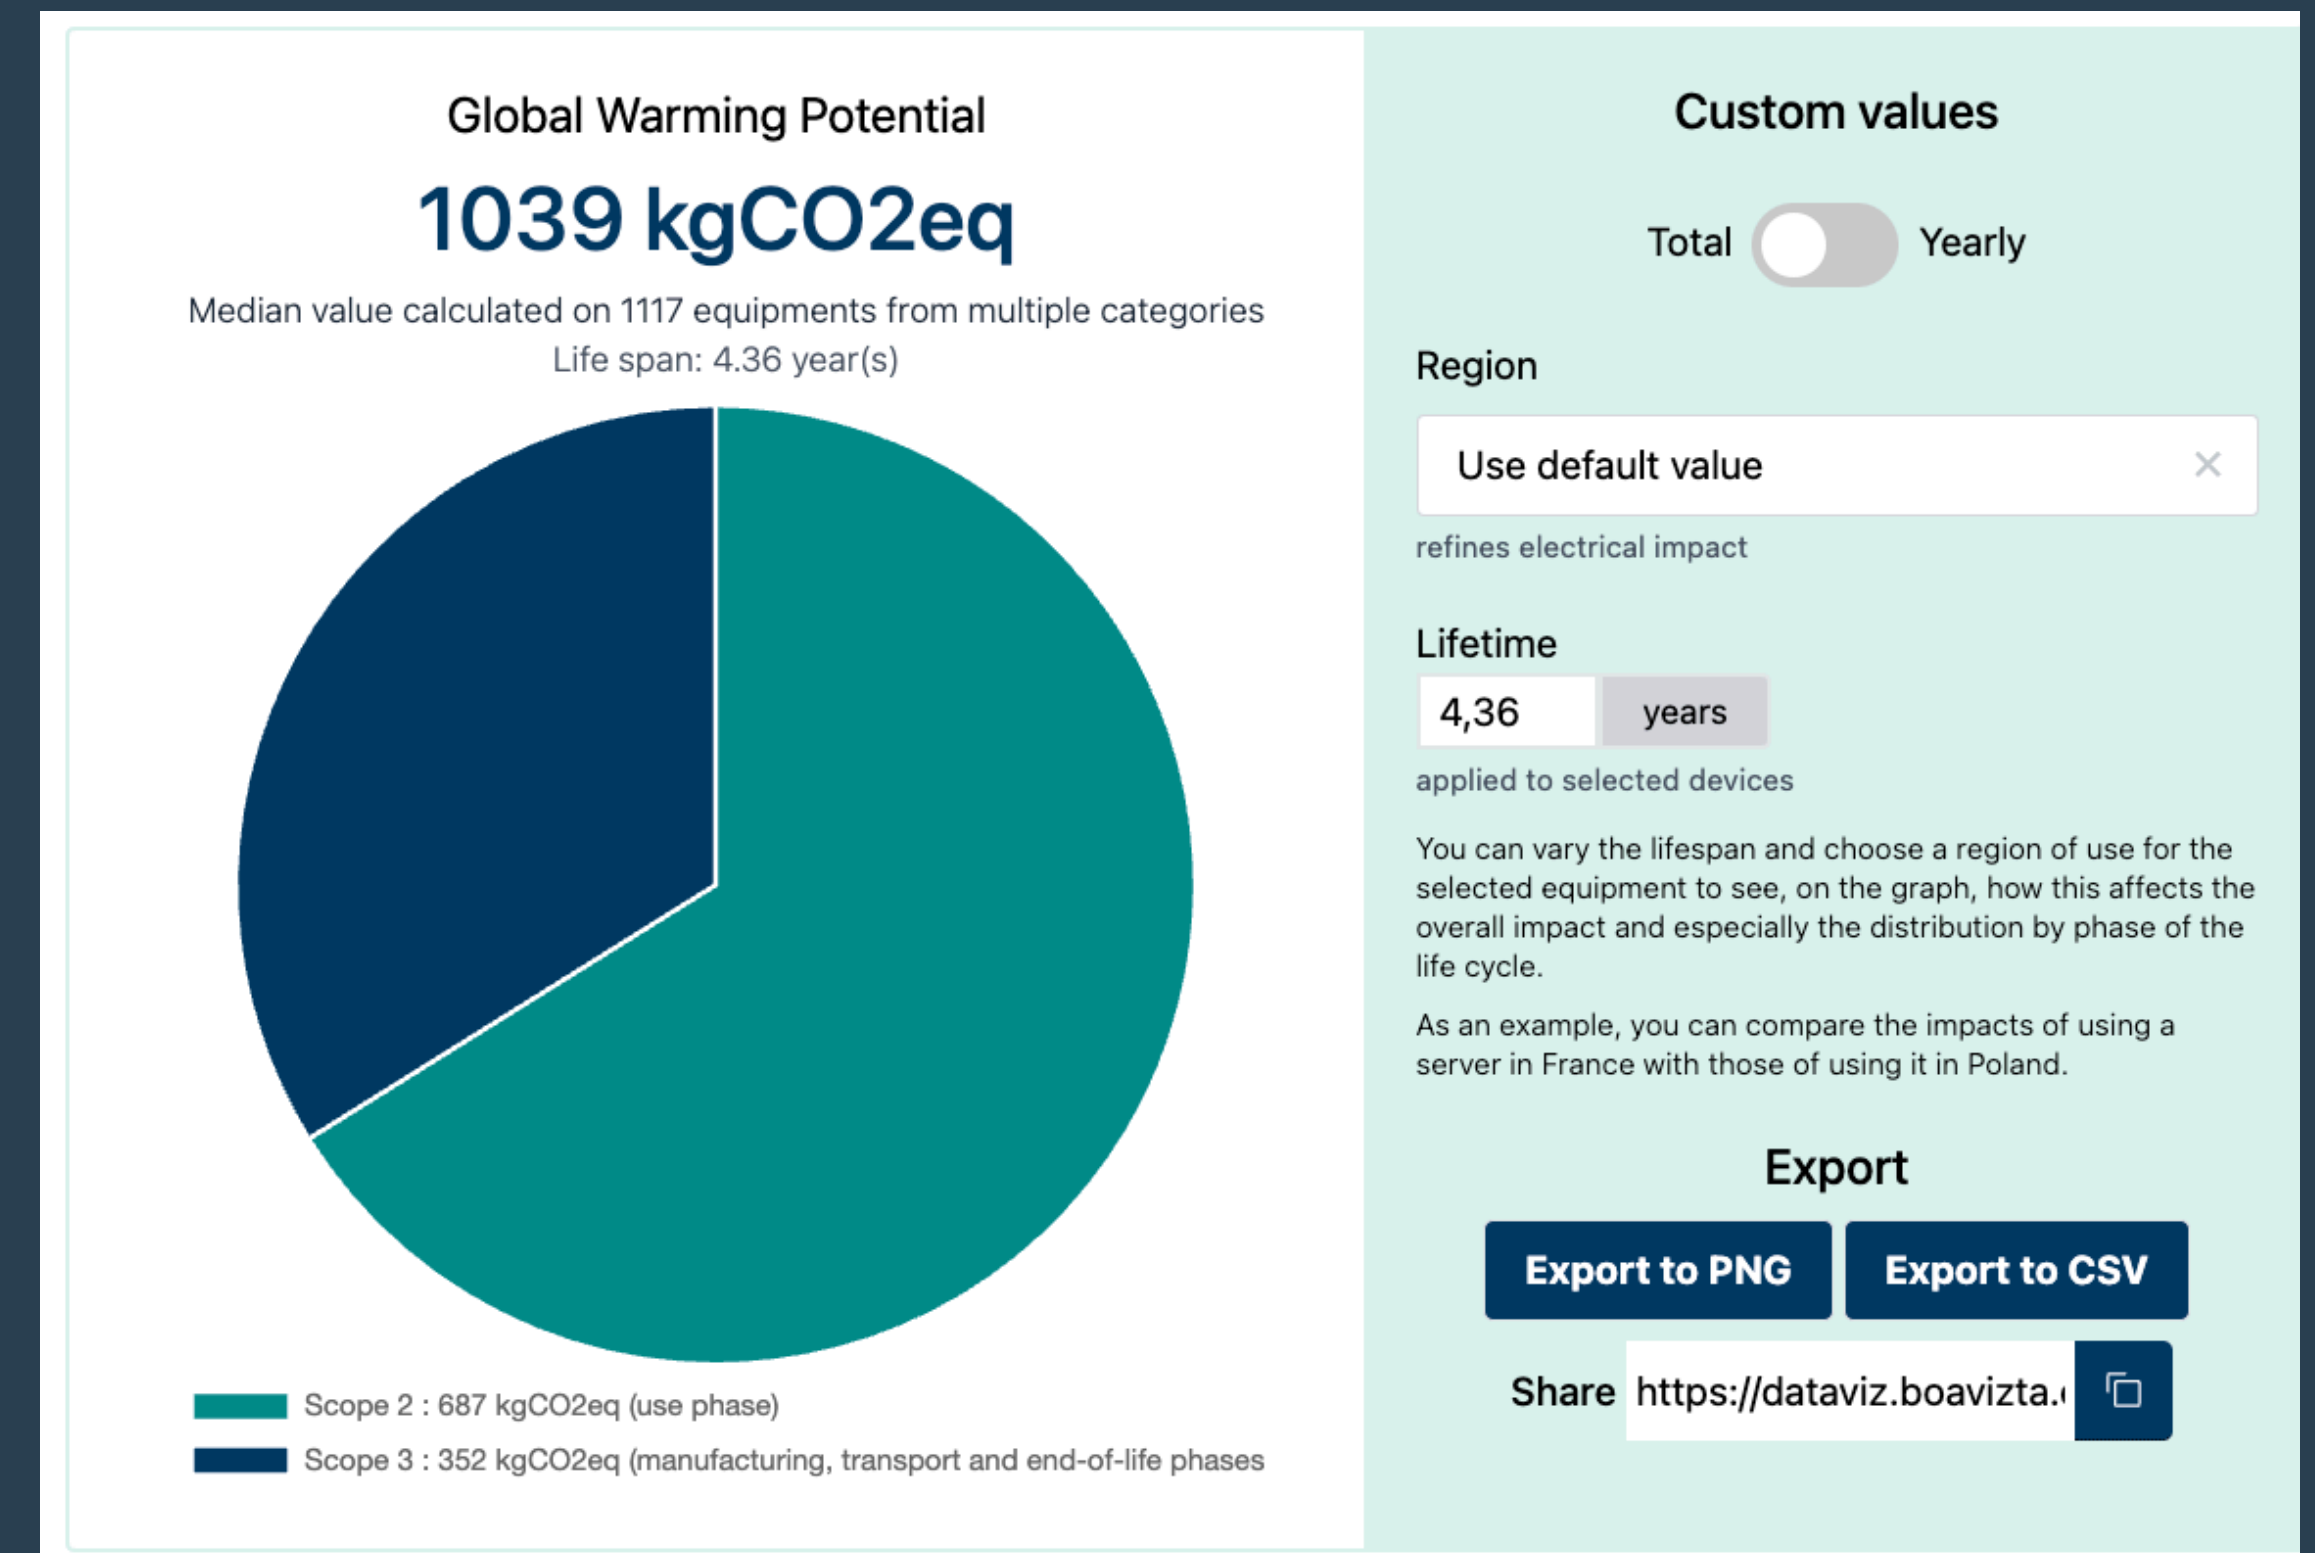
- **Dell**

Example: https://www.delltechnologies.com/asset/en-us/products/servers/technical-support/Full_LCA_Dell_R740.pdf

... many more

- **Important:**

They contain more than only CO2, also water usage, Toxic metals etc.

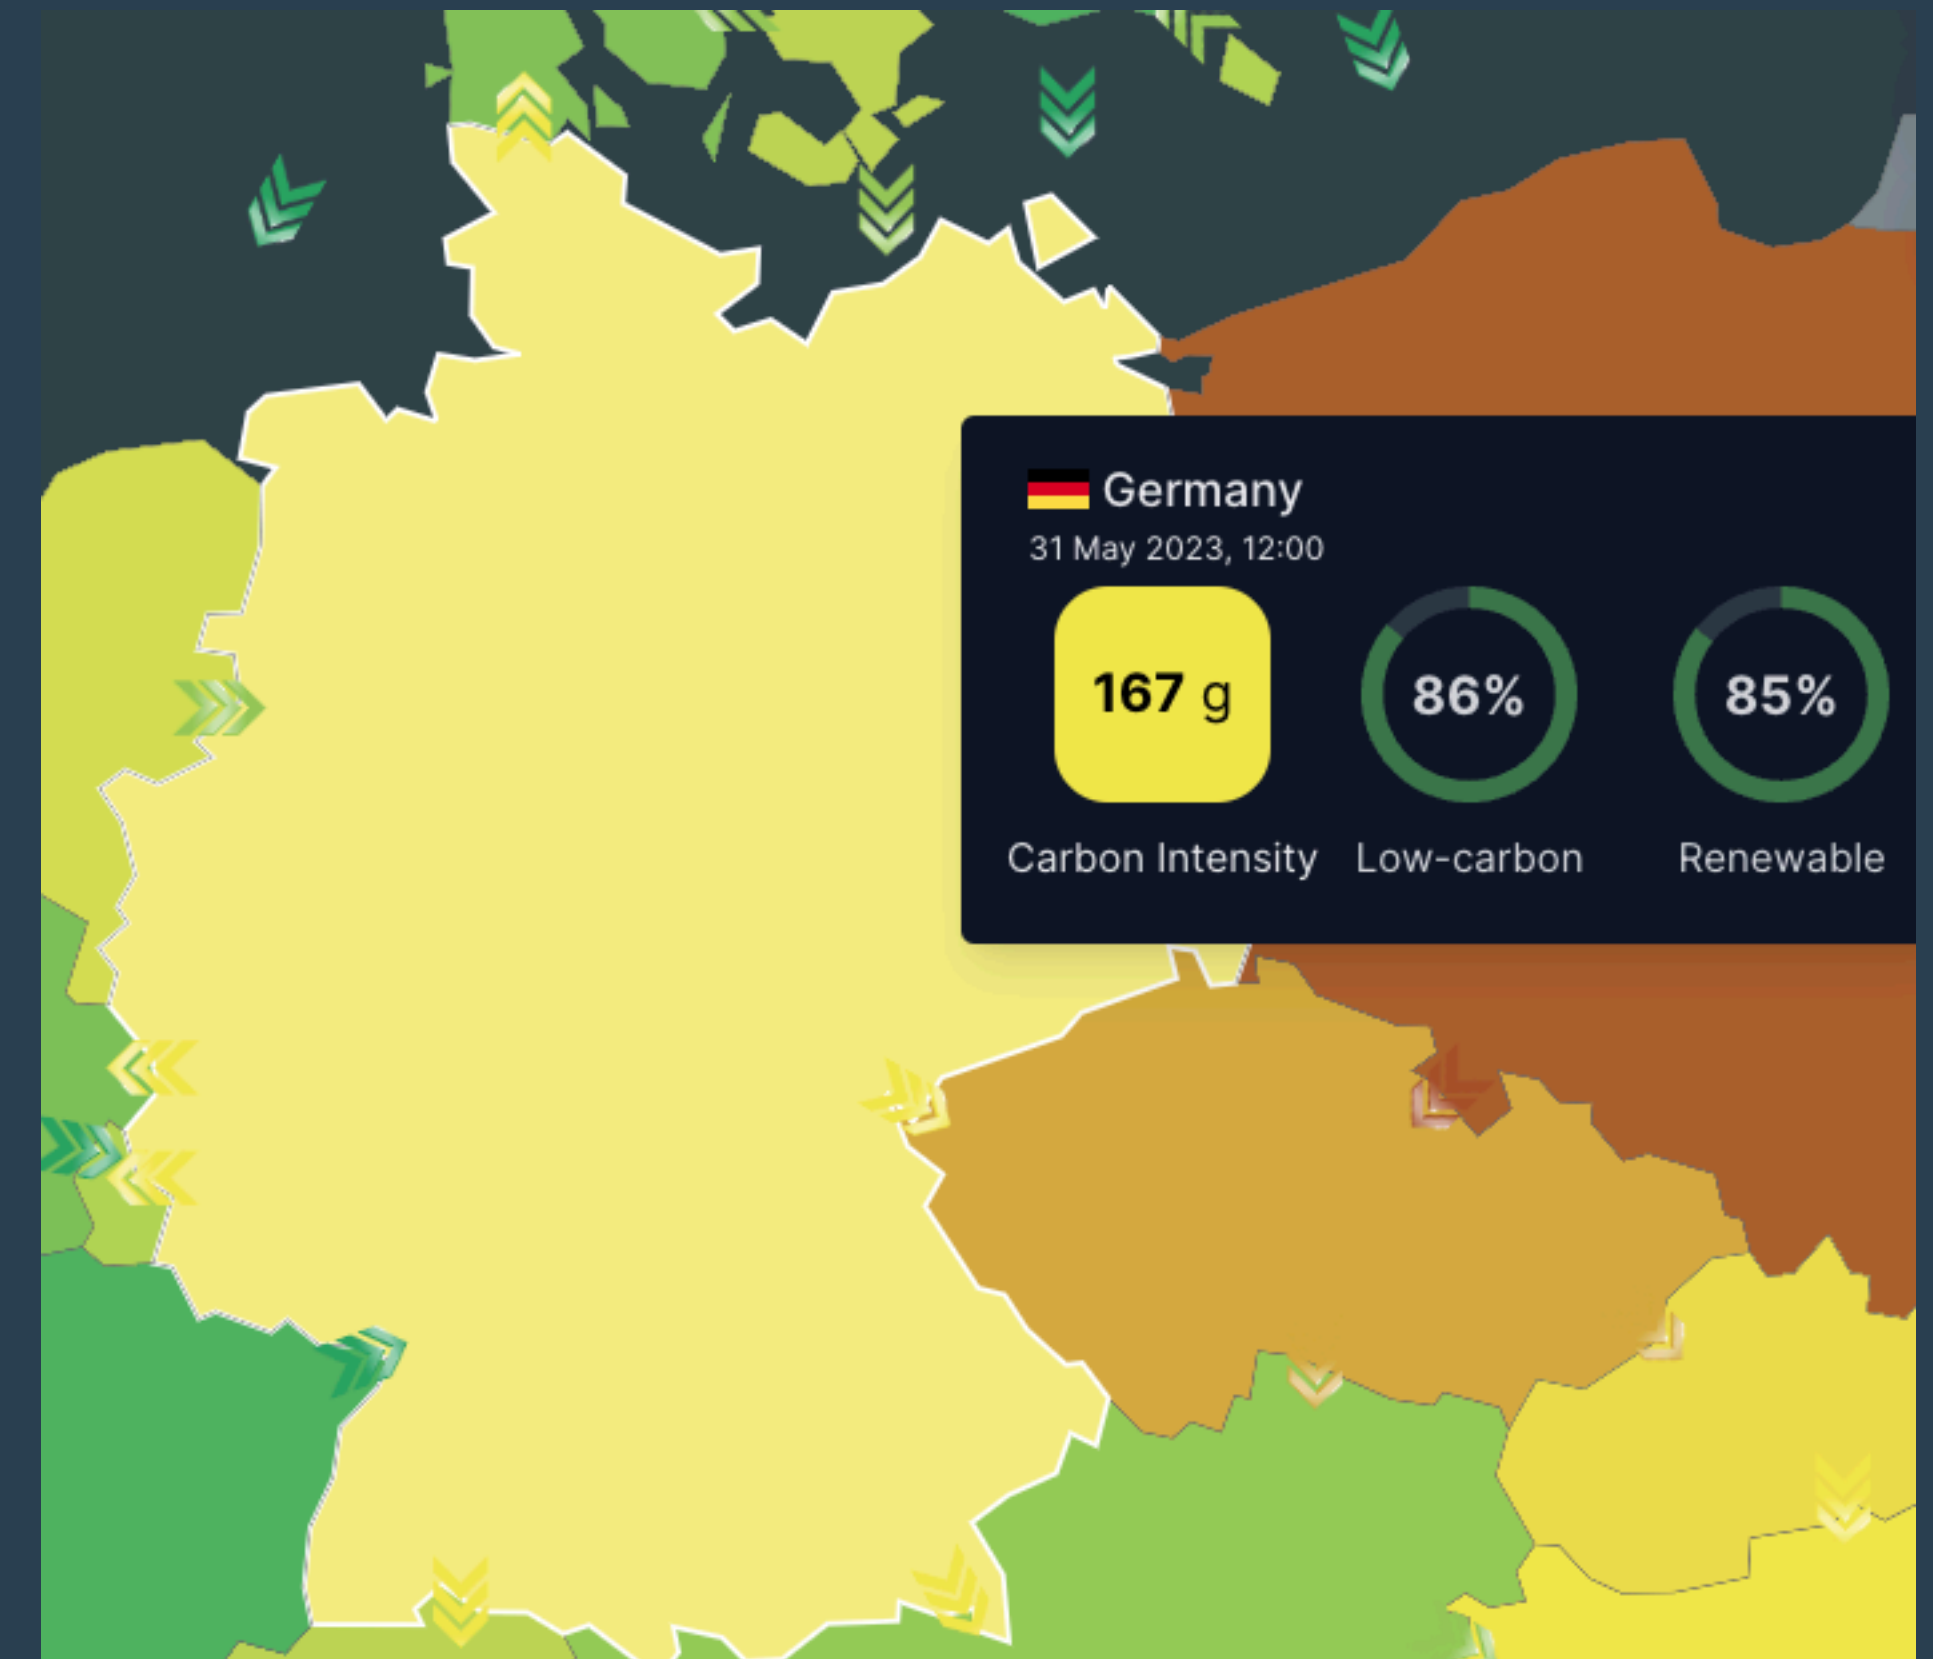


Source: <https://dataviz.boavizta.org/manufacturereadata>

How to get emission factors

Using open data sources

- **Electricitymaps**
<https://www.electricitymaps.com/>
- **Carbon-Aware-SDK**
<https://github.com/Green-Software-Foundation/carbon-aware-sdk>
- **Bundesnetzagentur (DE)**
<https://www.smard.de/home>
- **Watttime (US)**
<https://www.watttime.org/>
- ... many more

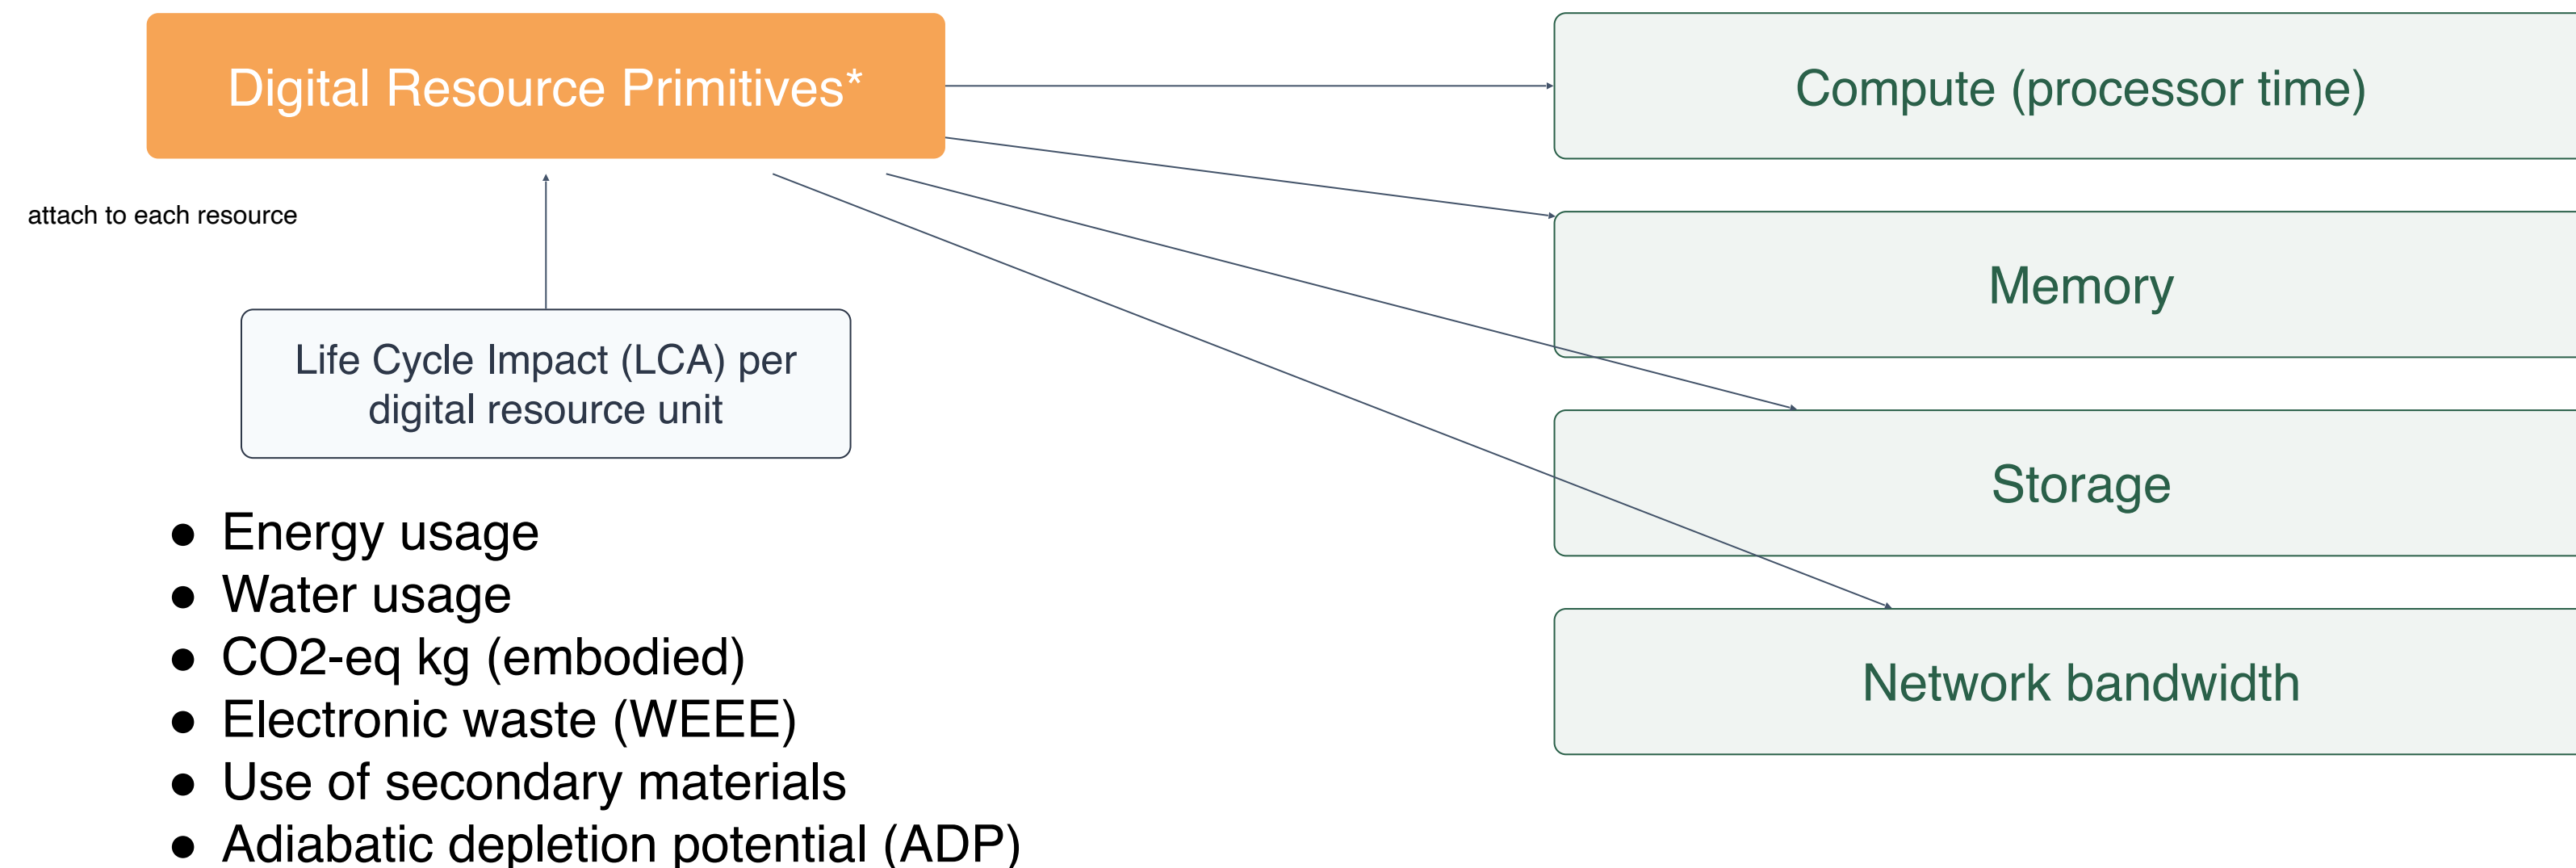


Source: <https://app.electricitymaps.com/zone/DE>

Sobald wir die Werte pro Einheit haben

können wir diese in abstrahierte Komponenten (Cloud Ressourcen) zerlegen

SDIA Modell der "Digital Resource Primitives"



Nicht nur CO2 zählt zu den Emissionen der Komponenten

Tools

Alrighty, let's look at some tools!

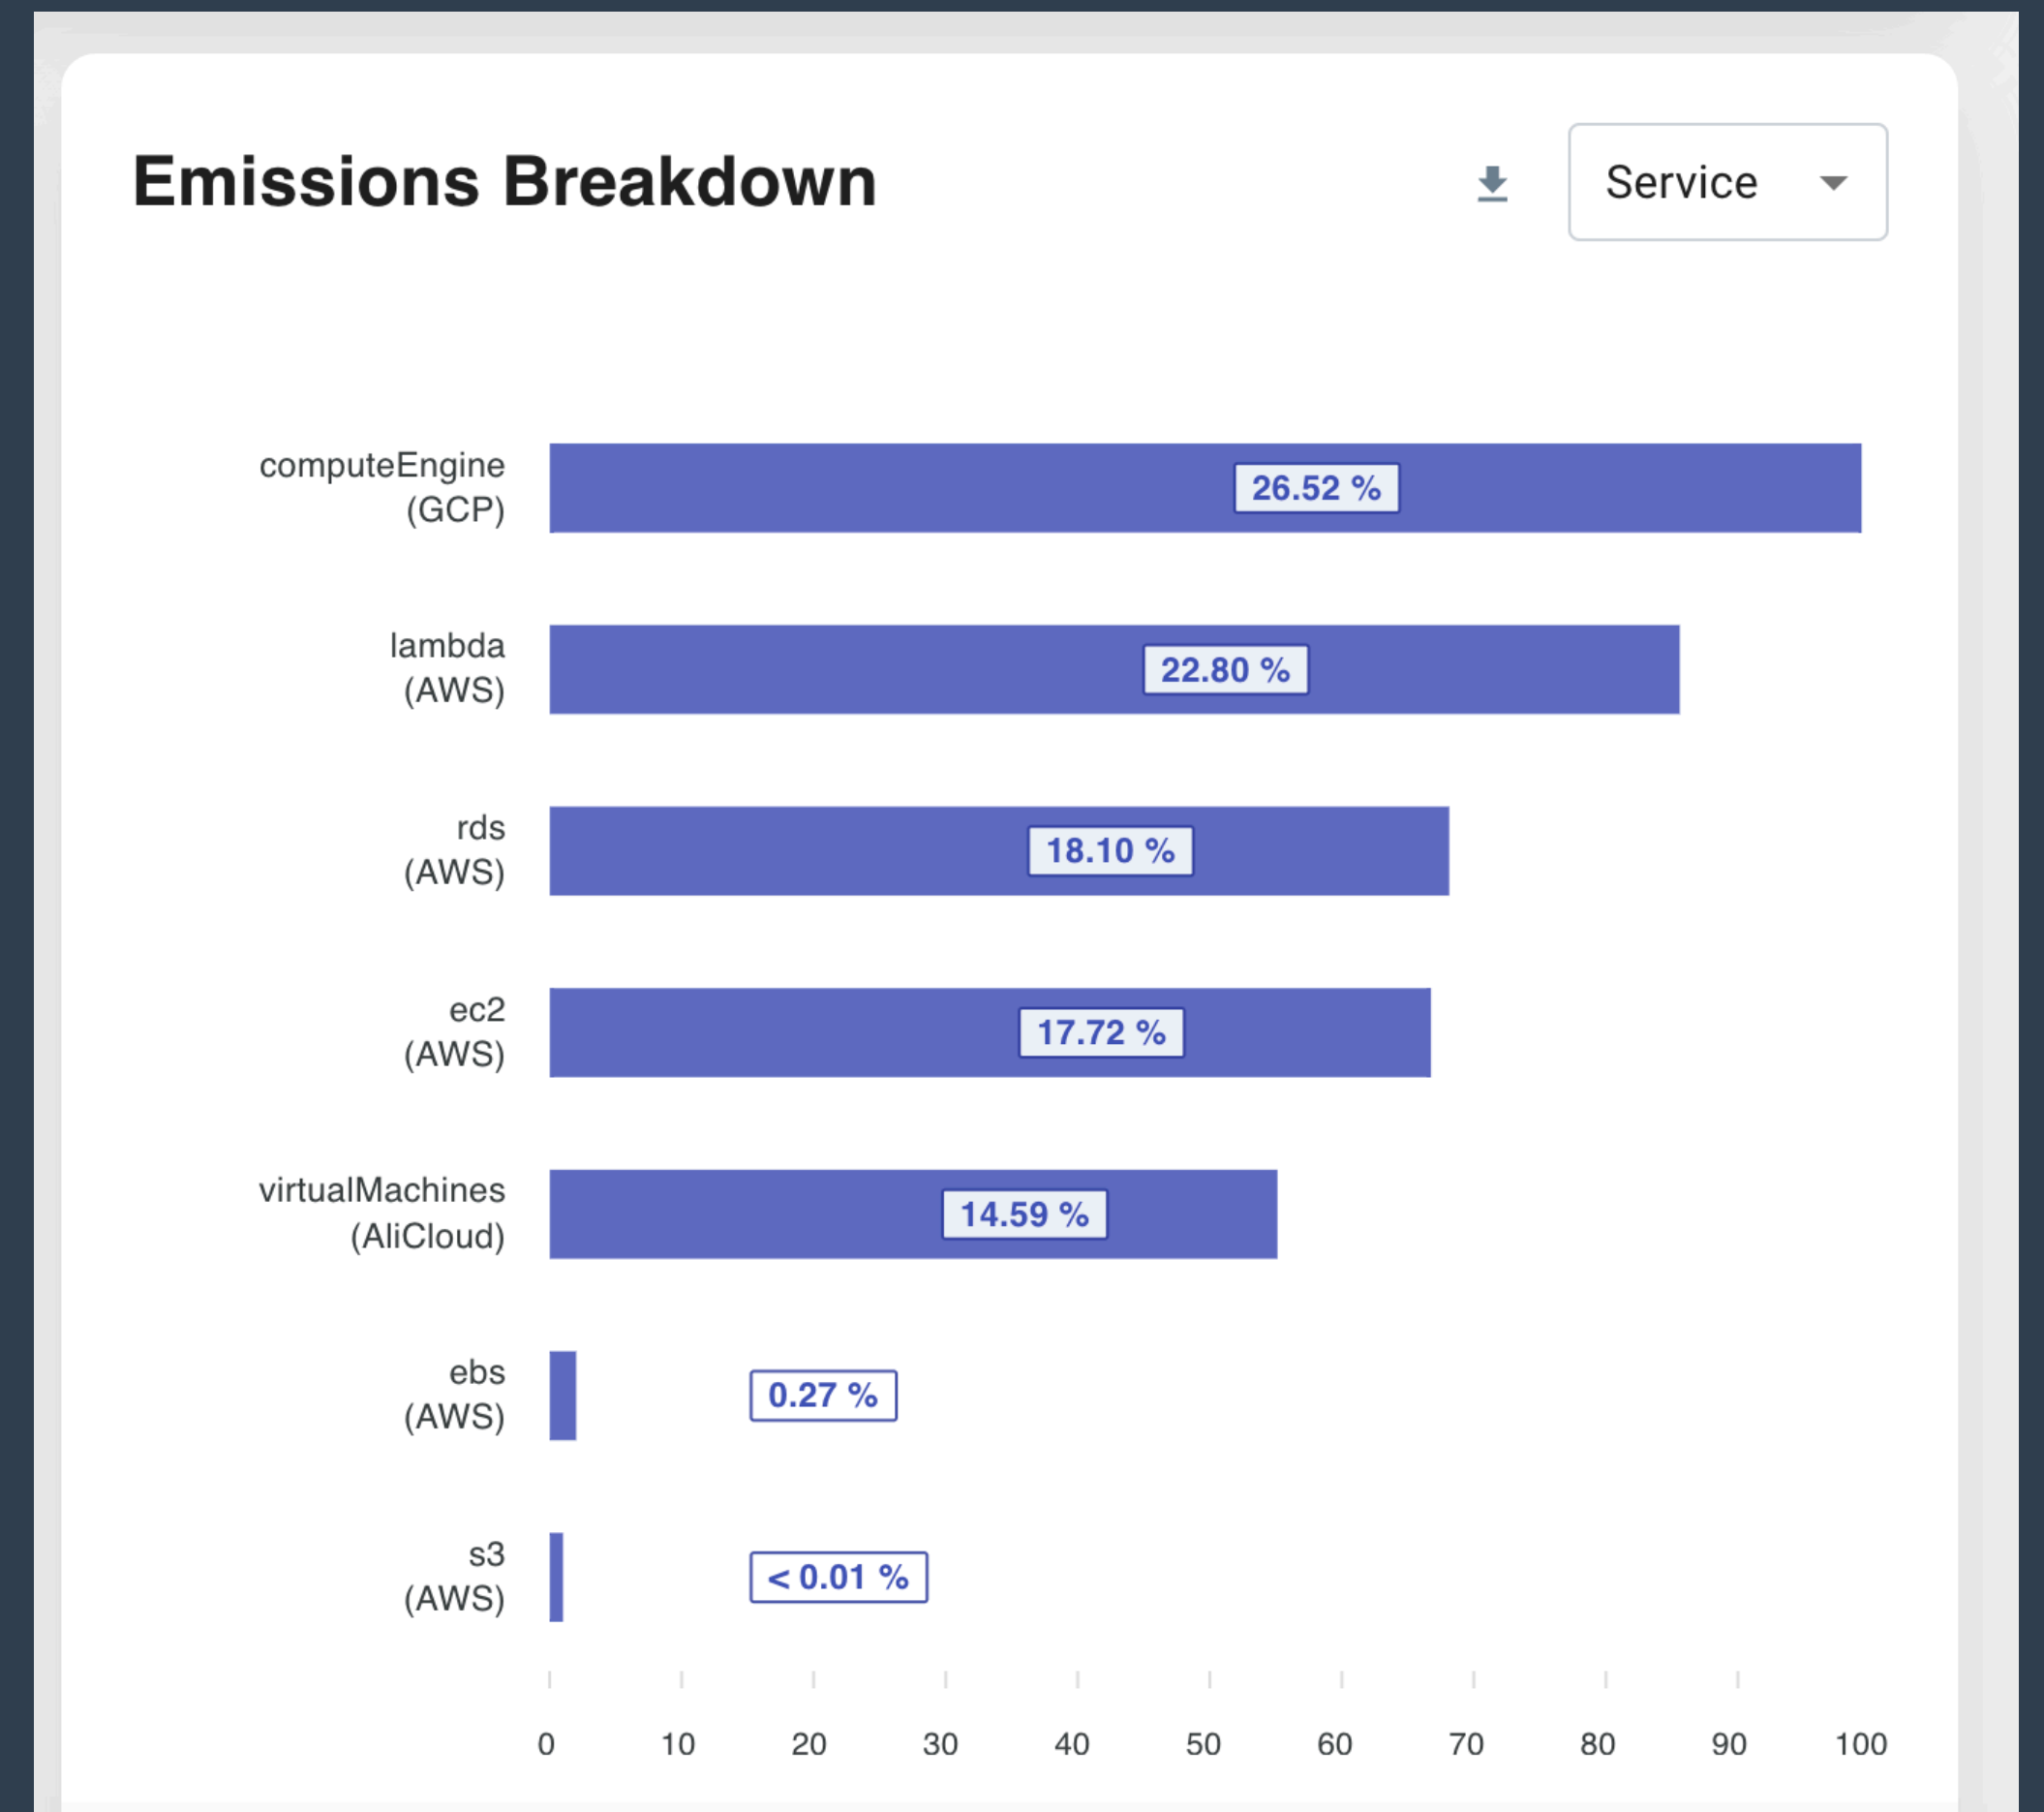
Estimation of IT CO2 emissions

If you do not know the energy directly - Cloud Carbon Footprint

- Uses billing data
- Based on this, a static calculation model is created [1] [2]
- Most cloud native services not covered
- IAM, DynamoDB, Logstash, etc.
- Downside: Only gives you the company total operation

[1] <https://www.cloudcarbonfootprint.org/docs/methodology/>

[2] Average Watts = Min Watts +
Avg vCPU Utilization * (Max Watts - Min Watts)
DRAM, hard disk etc. via coefficients



Anwendungsfall - pro Komponente

perf_events - CPU, DRAM, GPU (onboard), SMC

- Initial als Performance Engineering Tool / Profiler entwickelt
- Hervorragendes Tool, sofern neuer Kernel verfügbar [und man sich gut damit auskennt :)]
- Moderne Alternative für *manche* Fälle: eBPF

```
Performance counter stats for 'system wide':  
  
    76,49 Joules power/energy-pkg/  
    4,92 Joules power/energy-ram/  
    0,50 Joules power/energy-gpu/  
  
6,692703695 seconds time elapsed
```

Anwendungsfall - pro Prozess

Scaphandre - open-source RAPL based CLI tool

- Nettes Feature: Schlüsselte nach Prozess auf basierend auf CPU-%

```
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"
```

- Wie gut klappt das in der Praxis? -> Wir sehen uns im Workshop :)

Anwendungsfall - Programmcode

Beispiel: codecarbon.io



- Python
- RAPL-based
- NVIDIA GPU support

```
1 import tensorflow as tf
2
3 from codecarbon import Emission
4
5 mnist = tf.keras.dataloaders.MNIST
6
7 (x_train, y_train), (x_test, y_test) = mnist.load_data()
8 x_train, x_test = x_train / 255.0, x_test / 255.0
9
10
11 model = tf.keras.models.Sequential(
12     [
13         tf.keras.layers.Flatten(input_shape=(28, 28)),
```

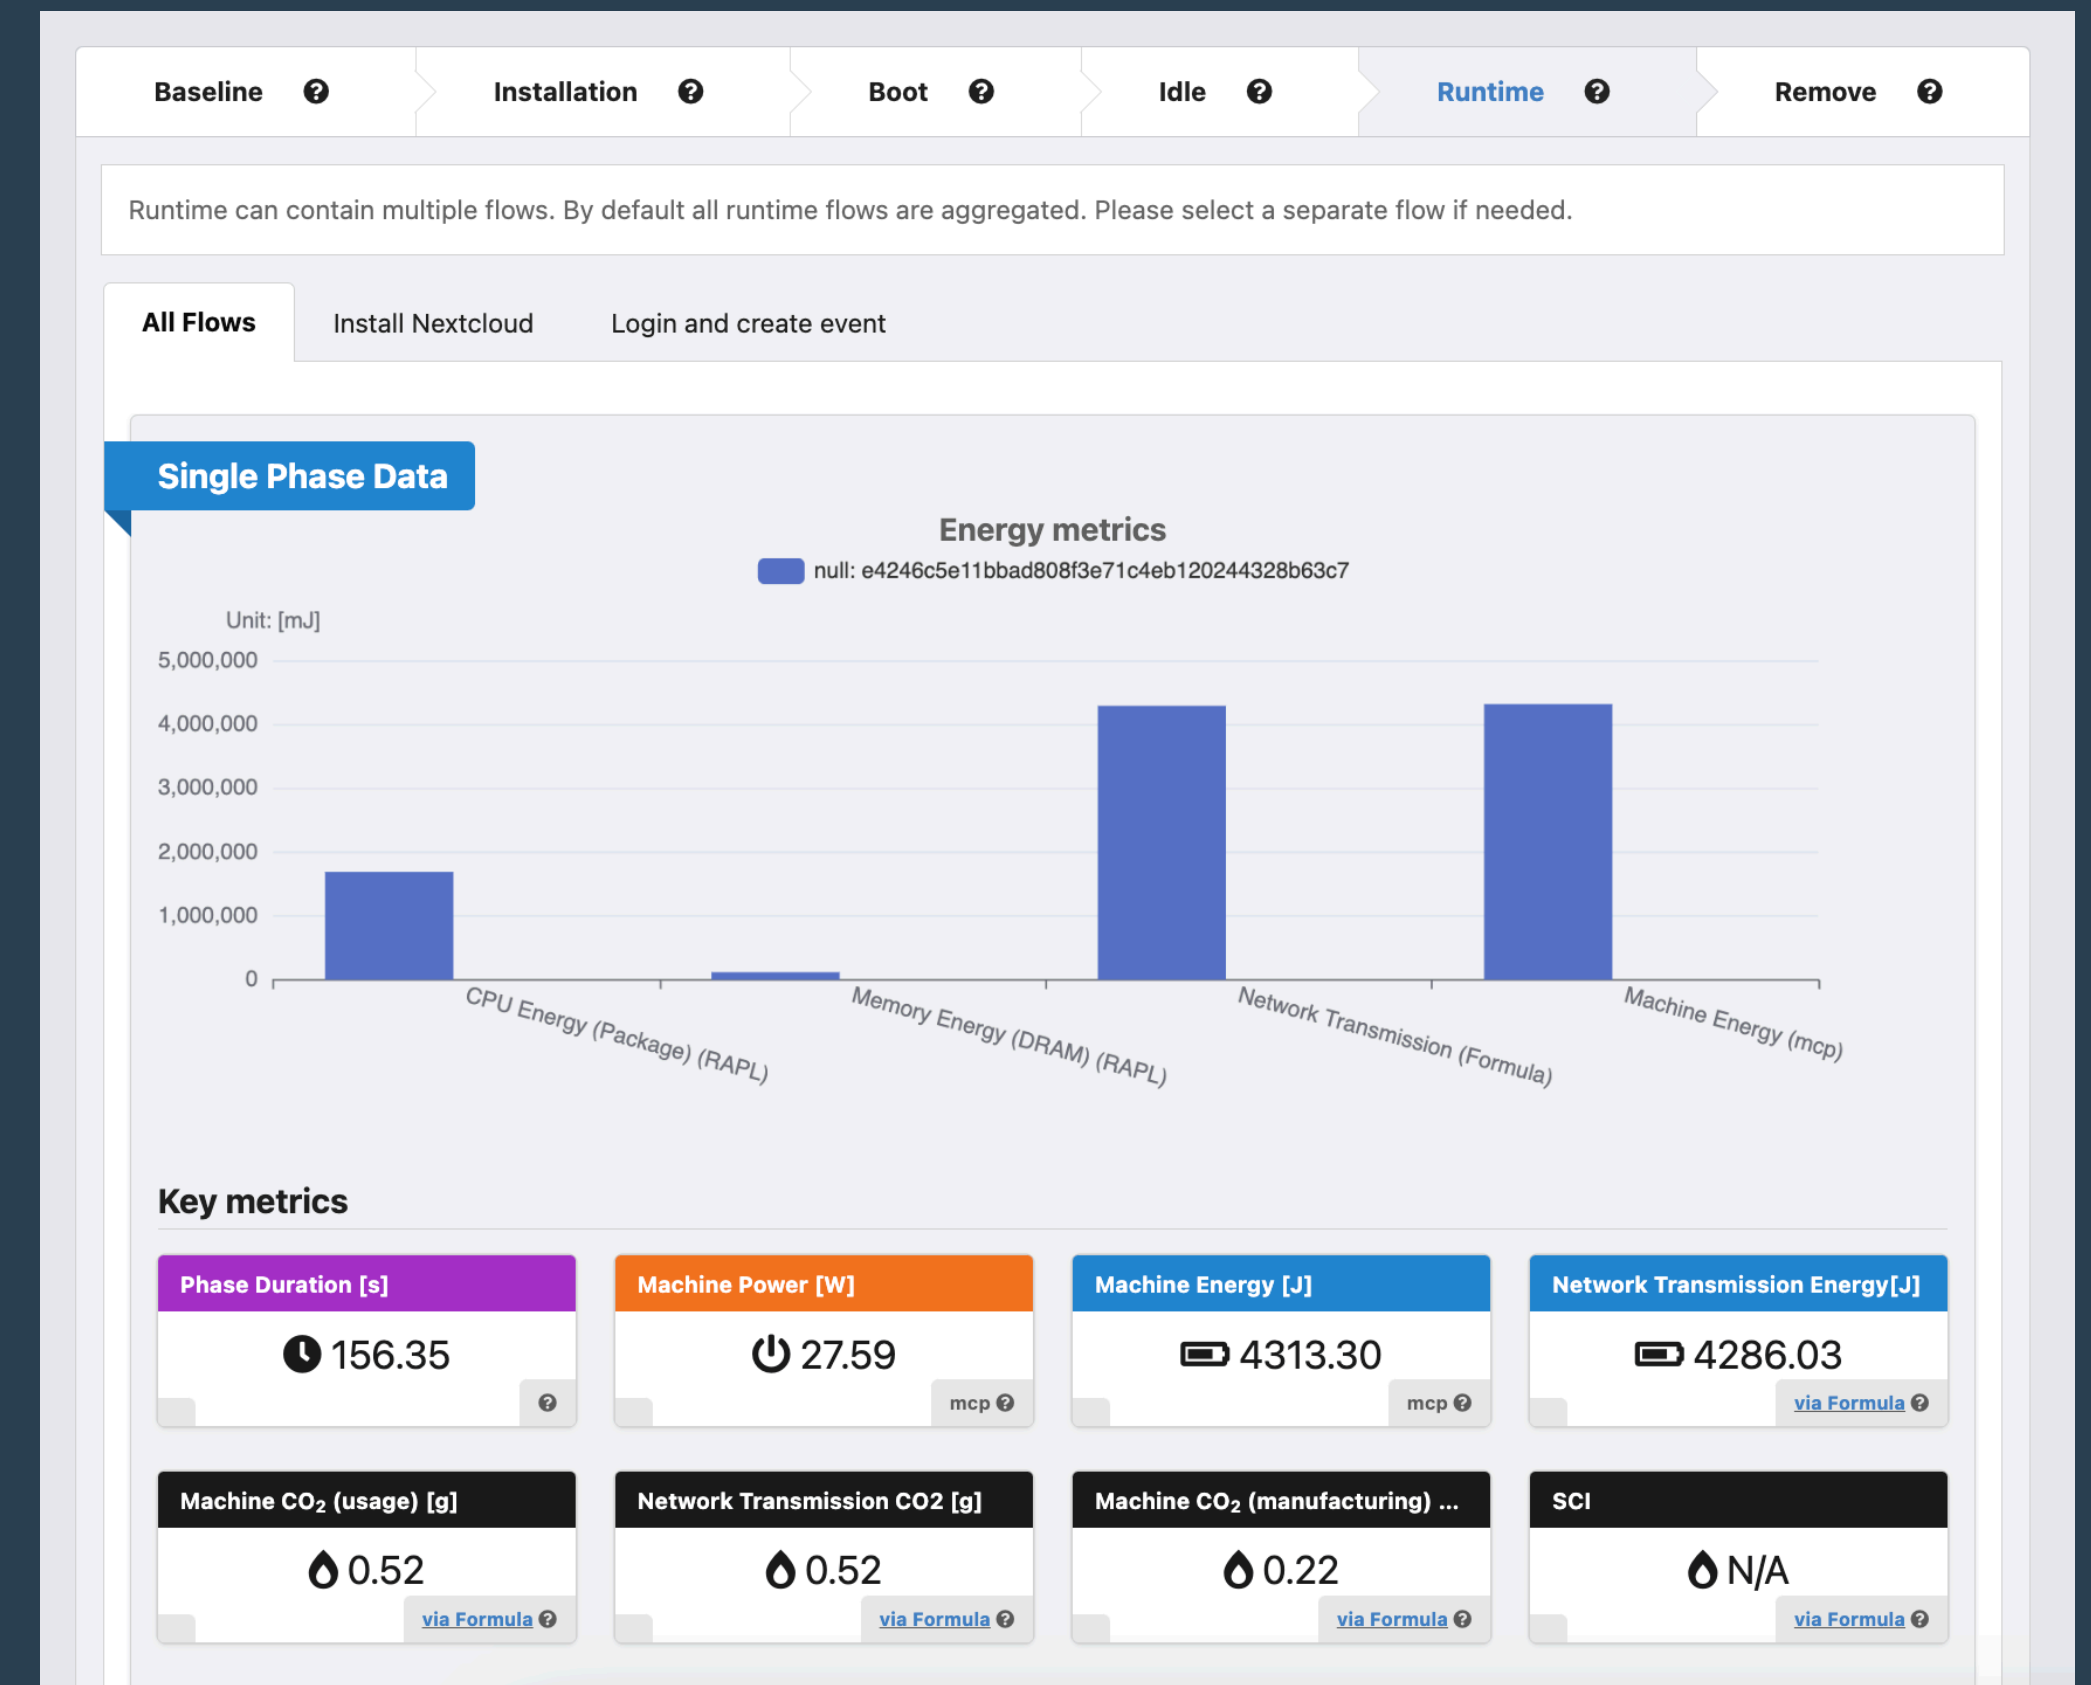
- Bessere Alternativen sind jedoch direkt die Runtime zu patchen / instrumentieren

Quantifying a product

Example: Green Metrics Tool

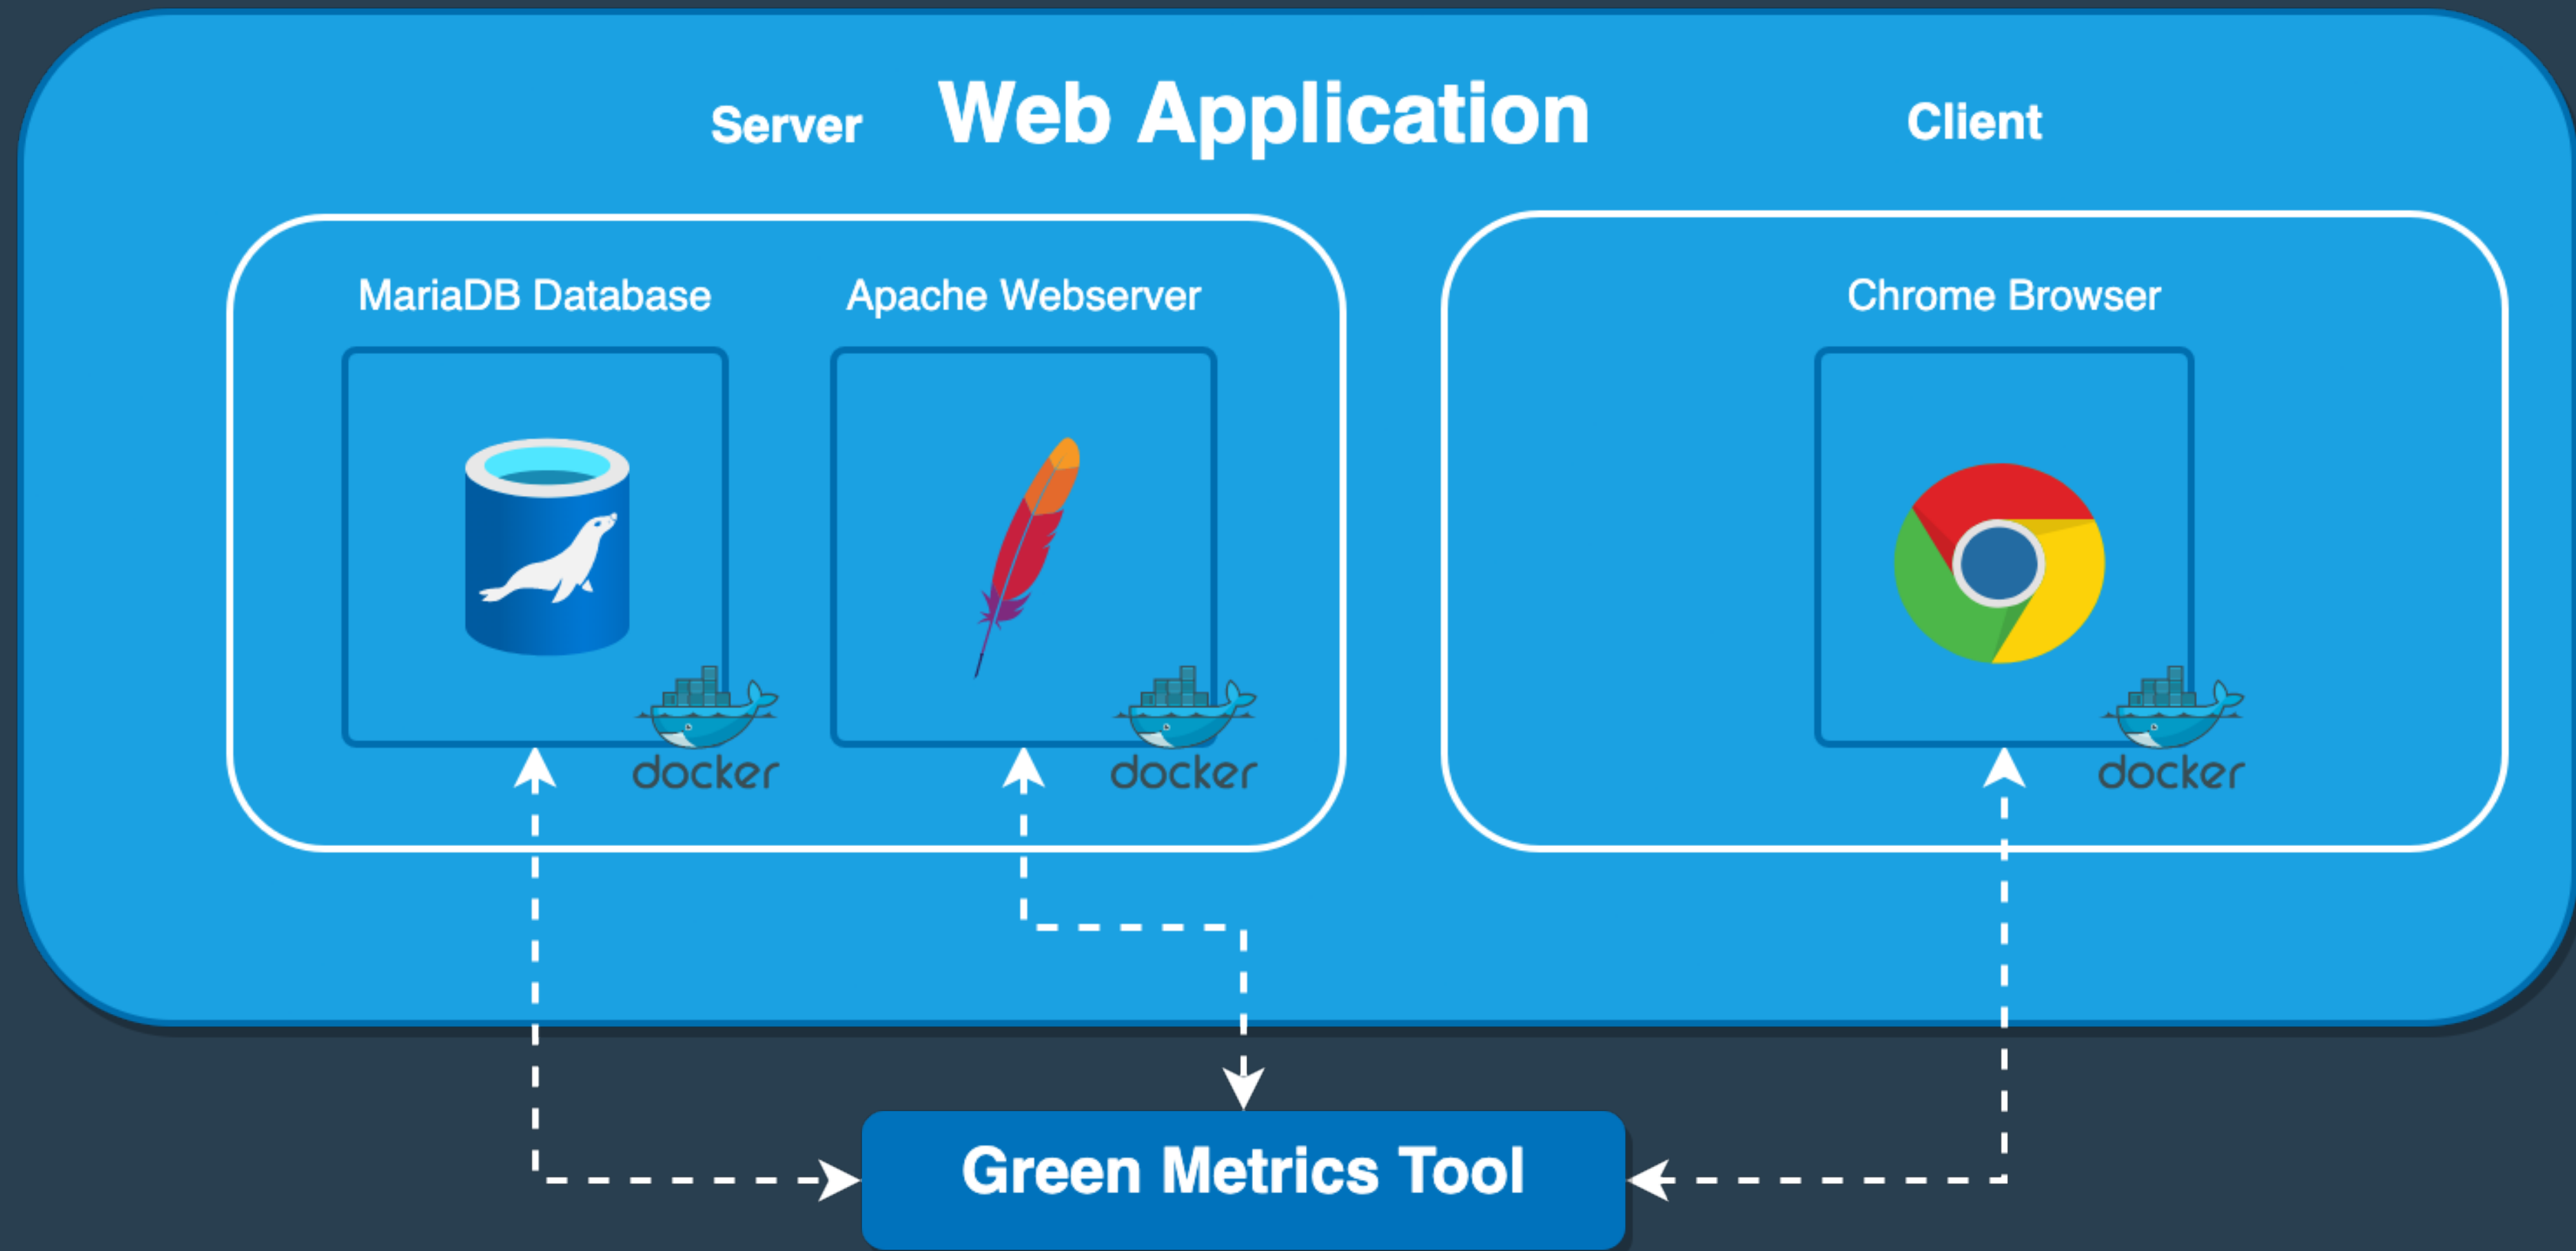
- Benchmarking via Standard-Usage-Scenarios (academia based via UCB / Öko-Insitut e.V.)
- Look at software in all phases of execution Building, Idle, Runtime etc.
- Include embodied carbon and modular metrics reporters
- SCI (ISO-Norm tbd 2024) / Blauer Engel compatible (German Eco Label)

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



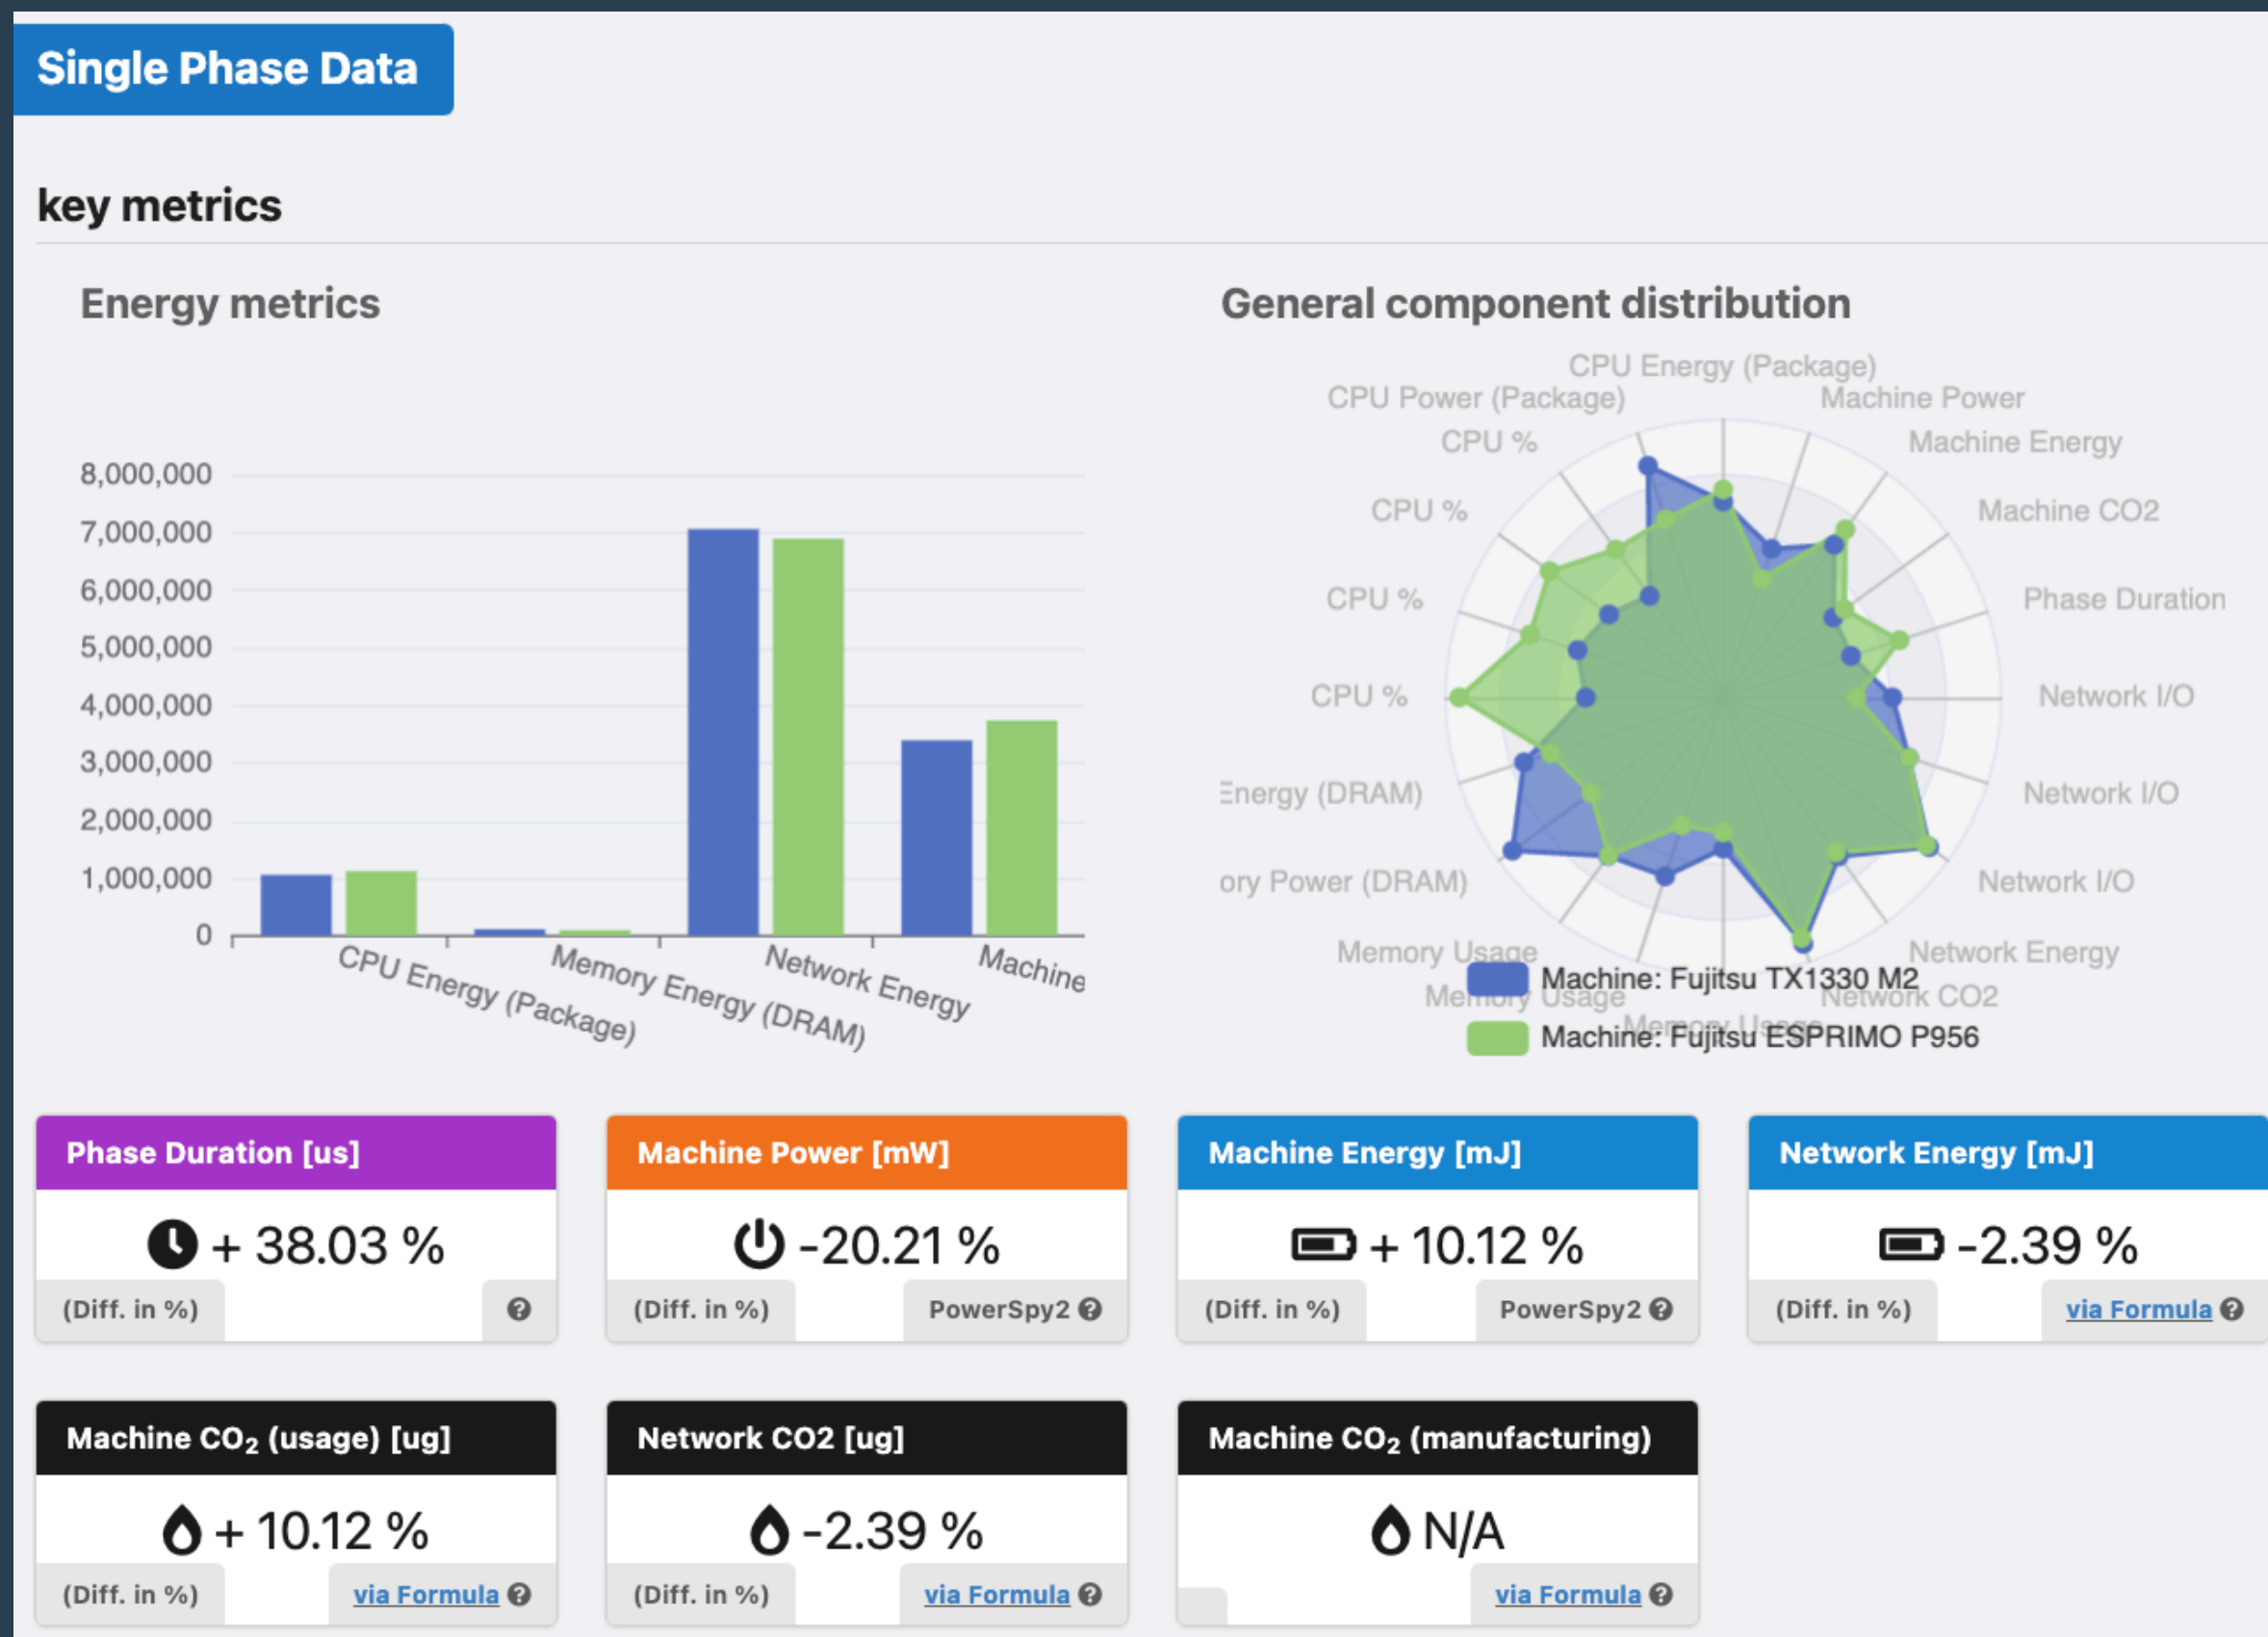
Green Metrics Tool

Aufbau einer Applikation für ein Usage Szenario - Client / Server

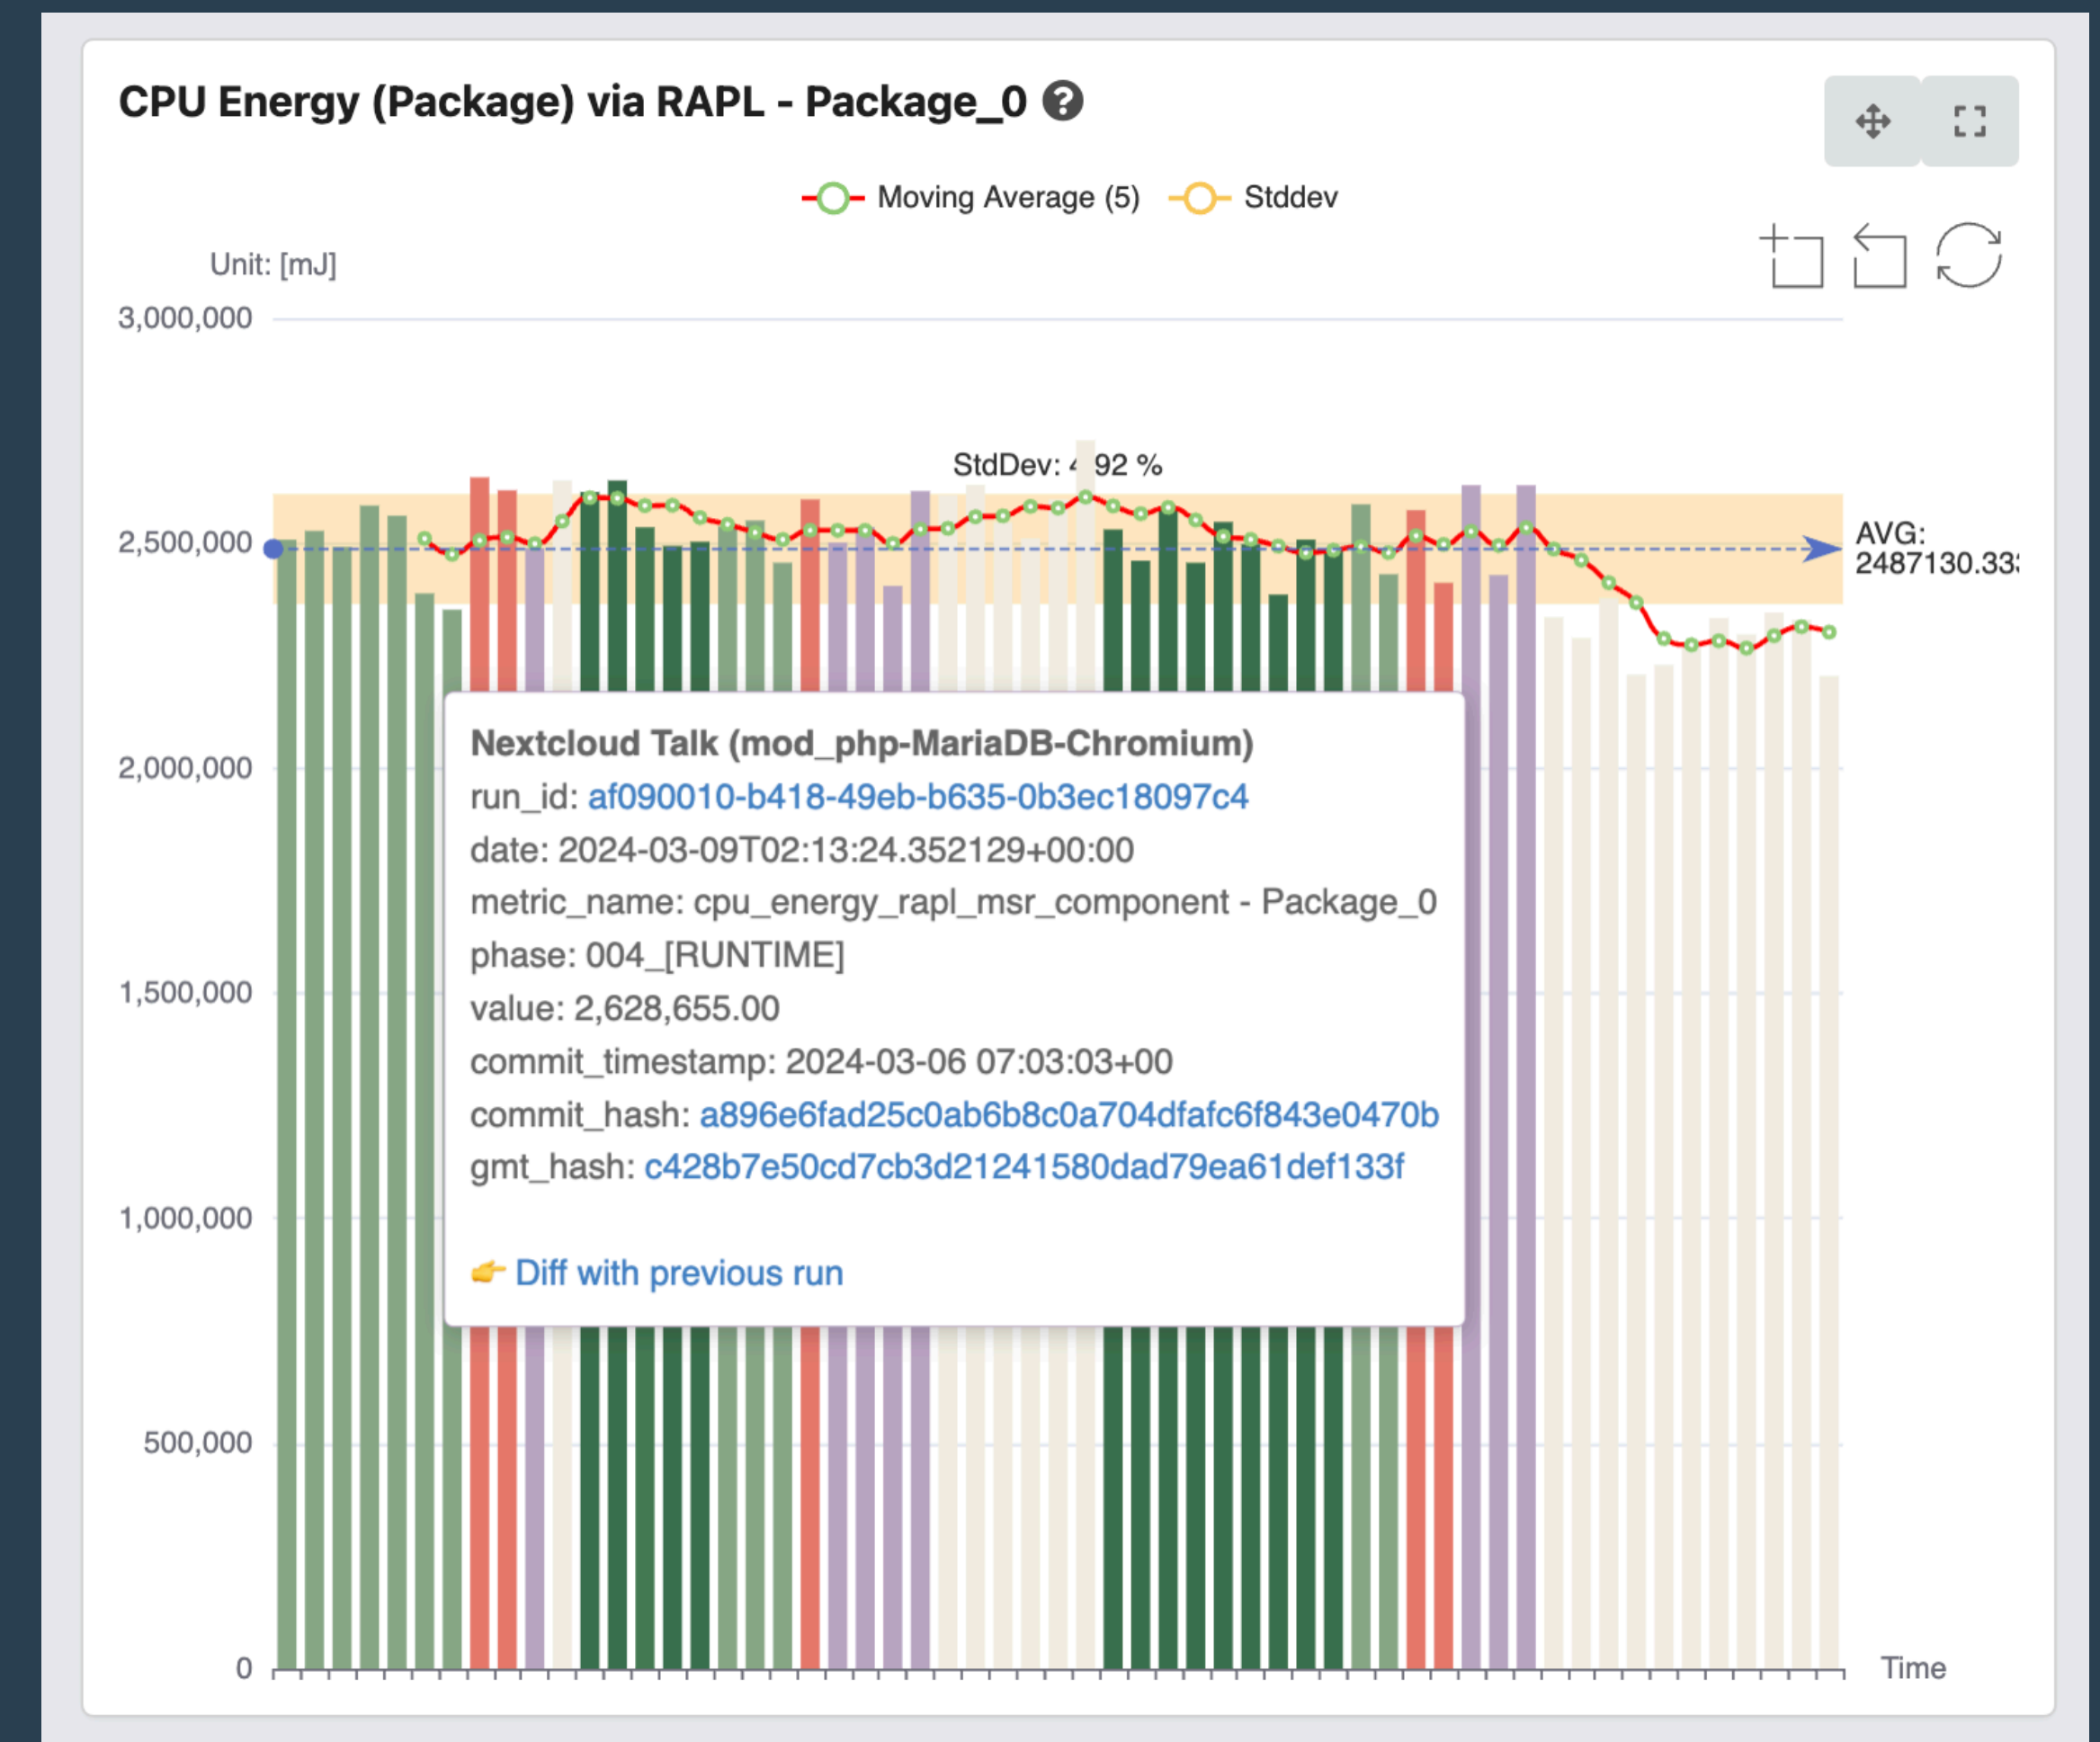


Green Metrics Tool

More Features



Comparisons between Hardware / Repository / Branch etc.



Measurement per Commit - Timeline View

Green Metrics Tool

More Features

- Optimization features
- Based on resource limits and performance metrics
- Includes runtime, build and Idle phases
- Currently in beta

Run Data

General Measurement Badges Machine Usage Scenario Logs Network **Optimizations 17**


This is a list of potential optimizations we have detected based off your run. We are constantly adding new features so please check our [documentation](#) for more details. If you think a warning does not apply to you, you can add an ignore into your `usage_scenario.yml` file.

- Docker container boot time** (Docker) `docker-boot-time`
The container takes very long to become usable ~ 456s. Containers are meant to be started and killed quite fast.
- Cpu container resource allocation** (Docker) `docker-cpu-allocation`
Container 'app' is maybe overprovisioned. CPU utilization was '2%'. Max was '31%'.
- Low IPC count** (cpu) `low-ipc-count`
The ipc count is 1. This is a sign that you are not using all features of the processor.
- High page faults** (memory) `high-page-faults`
Your system is experiencing a lot of page faults. See documentation for a detailed description on how to maybe fix this.
- Code with high optimization potential found** (Ai) `ai-low-code-quality`
This piece of code was rated a 18/100 by our Green Coding AI. We think this can be optimized
`object_specification.py:473: return [line.strip().split(',') for line in open(file_path, 'r')]`
- Energy hotspots** (Benchmark) `energy-hotspots`
We have detected energy hotspots in the flame graph. See documentation on how to run your code with full profiling enabled.

Anwendungsfall Web: Statische Analyse

Eco-Code / CAST etc.

- Basierend auf Empfehlung aus Wissenschaft und französischem Positionspapier Numérique Responsable



- Kein nachgewiesener Effekt in größeren Software-Projekten
- Empfehlungen wie i++ statt ++i sogar im Assembly Code / Byte-Code oft identisch
- Overhead der Tools unklar

Deep dive into one label: websitecarbon.com

Tool to get the carbon footprint estimate of a website


- Very visual, good usability, but also very basic
- Entering the URL is sufficient
- Unfortunately, there are many calculators and every one provides very different values
- Most only use network traffic (not packets)
- Some add constants for end device screen time or server costs
- Good for communicating that you care. Not so much for comparisons / reporting

Carbon results for
green-coding.org Share     

This page was last tested on 30 Jan, 2022. [Test again](#)

 **Hurrah! This web page is cleaner than **87%** of web pages tested**

 Only **0.18g of CO2** is produced every time someone visits this web page.

 **This web page appears to be running on **sustainable energy****

Eco-CI

Estimation in CI / CD Pipelines

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

github-actions bot commented 1 minute ago

Eco-CI Output:

Label	avg. CPU utilization [%]	Total Energy [Joules]	avg. Power [Watts]	Duration [Seconds]
Total Run	6.83923	24.7645	1.7689	27
Measurement #1	5.99	12.3823	1.7689	8
Sleep #2	5.72625	12.3823	1.7689	8

Energy graph:

CO2 Data:
City: San Francisco, Lat: 37.783, Lon: -122.417
Carbon Intensity for this location: 208
CO2eq emitted for this job: 0.005151


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

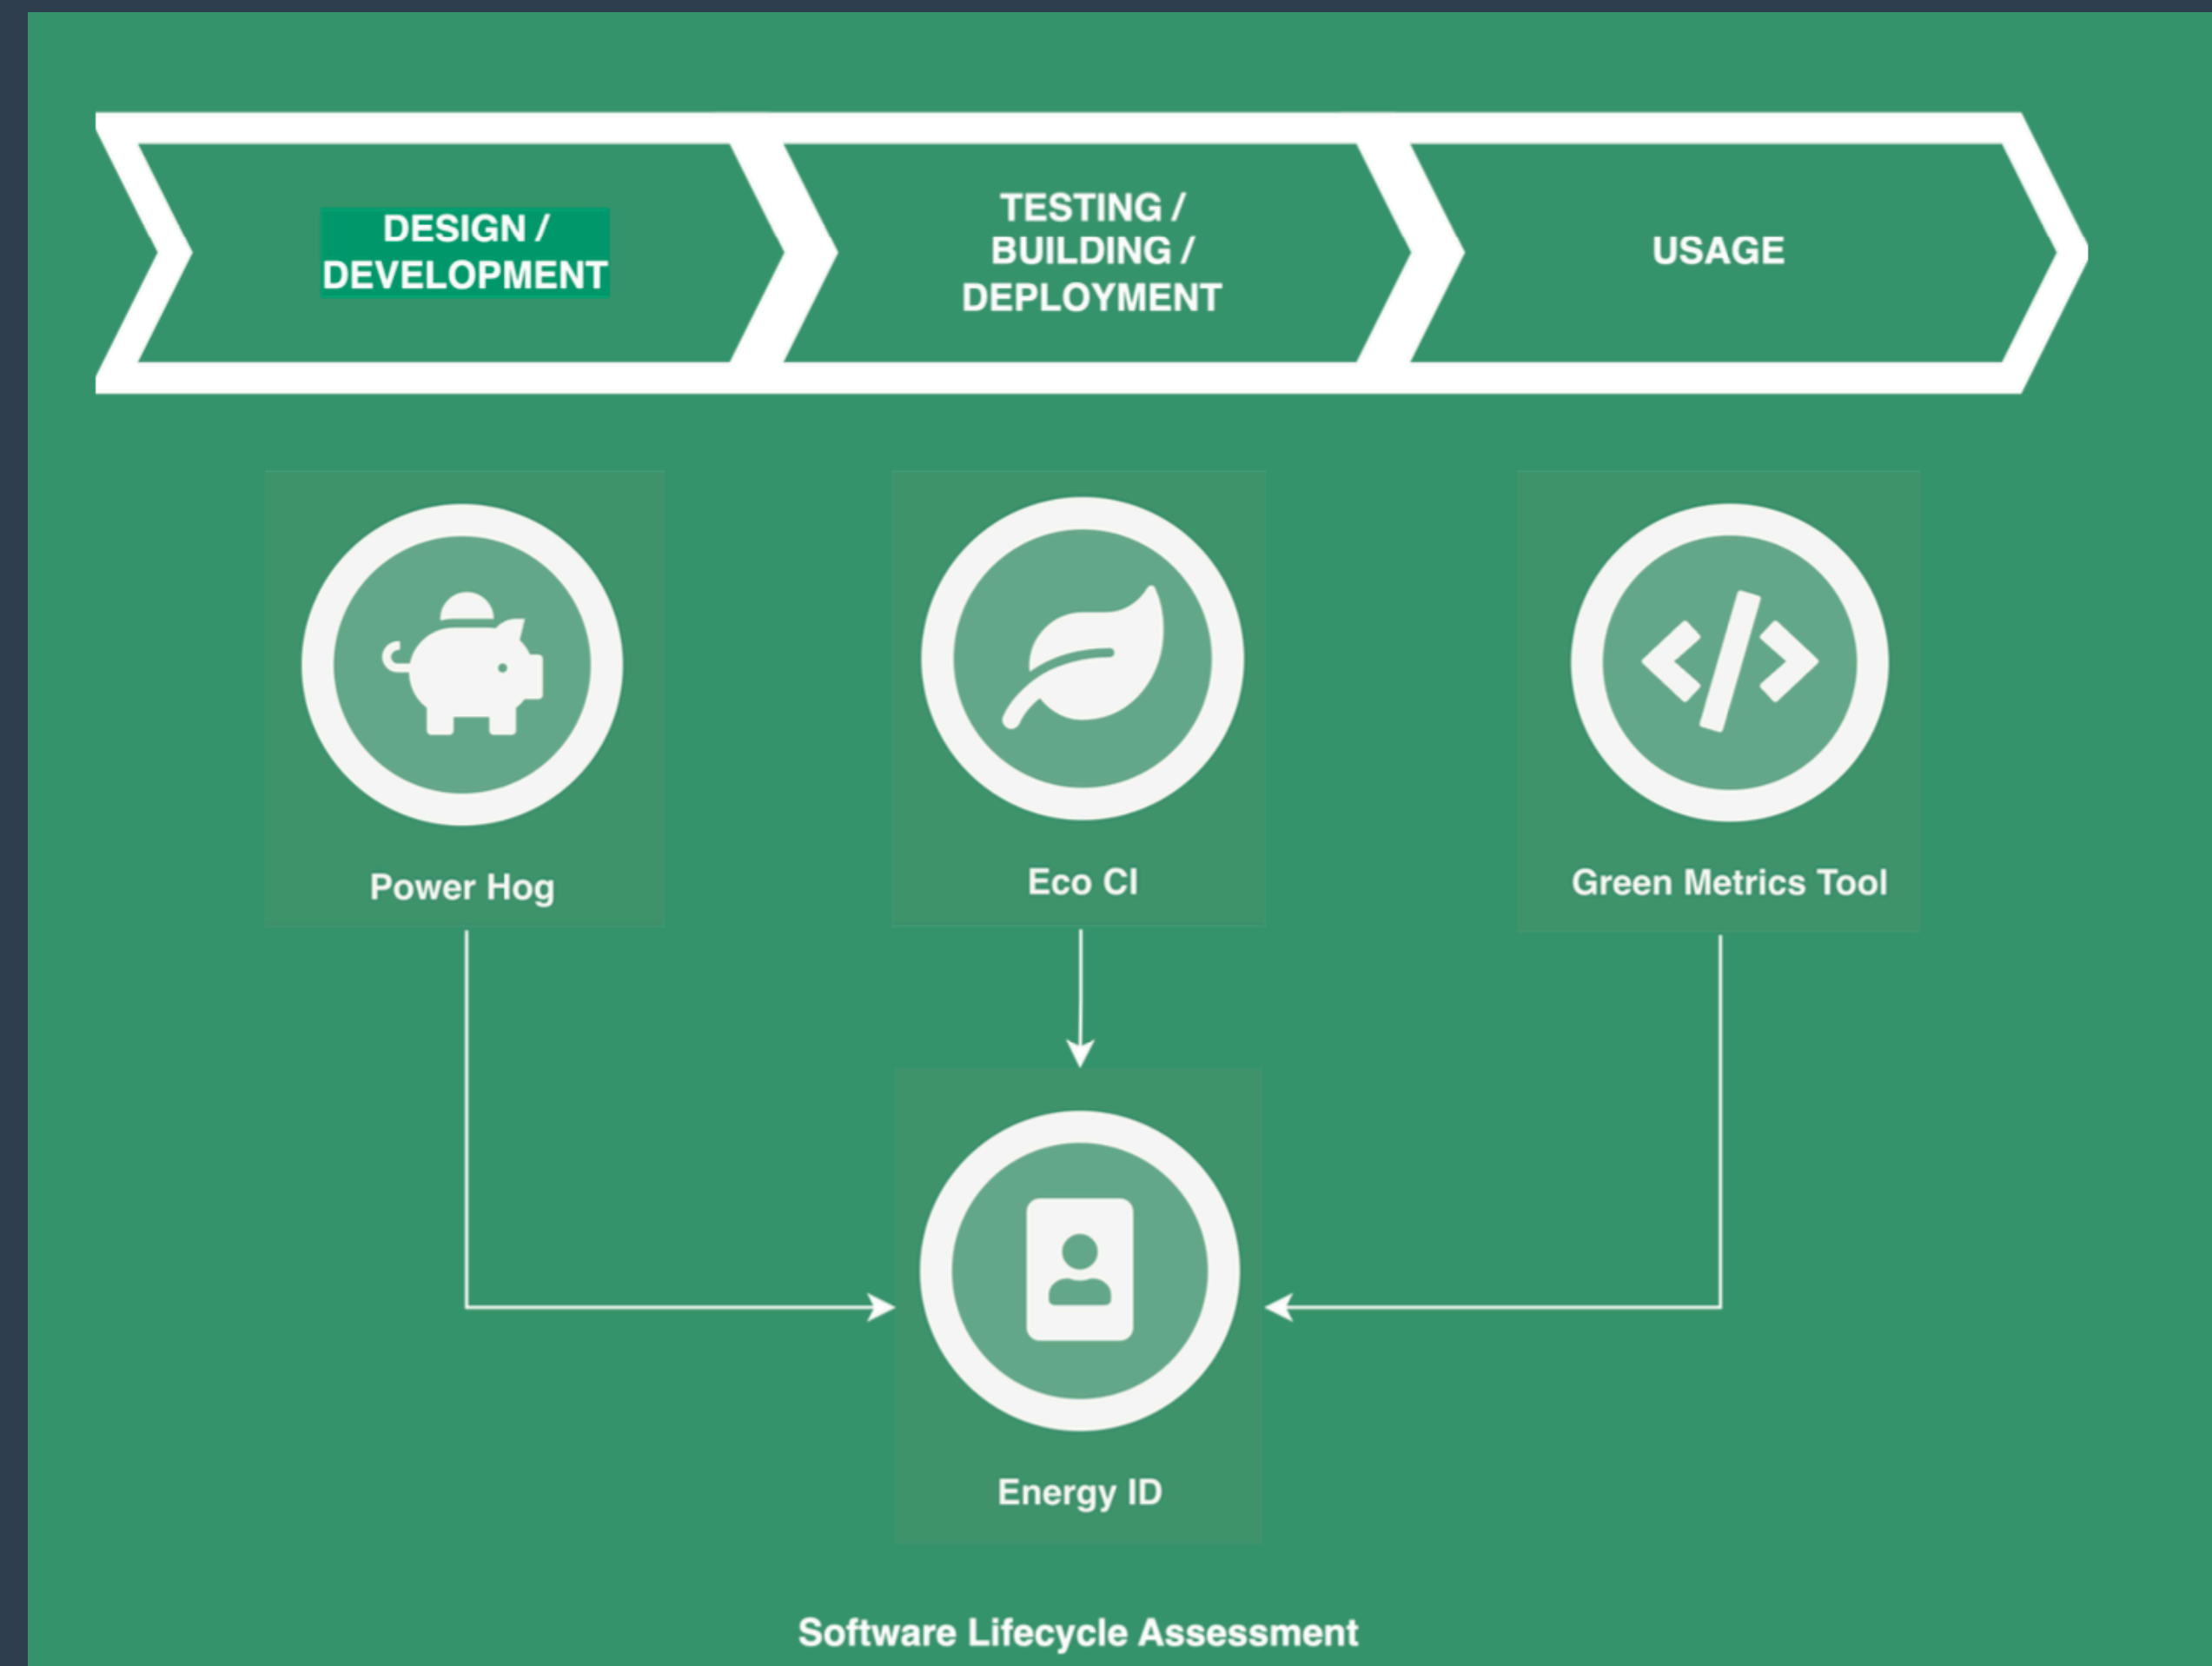
Development Cost measurement

- 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



Software Lifecycle Assessment

Combines all parts to a holistic perspective



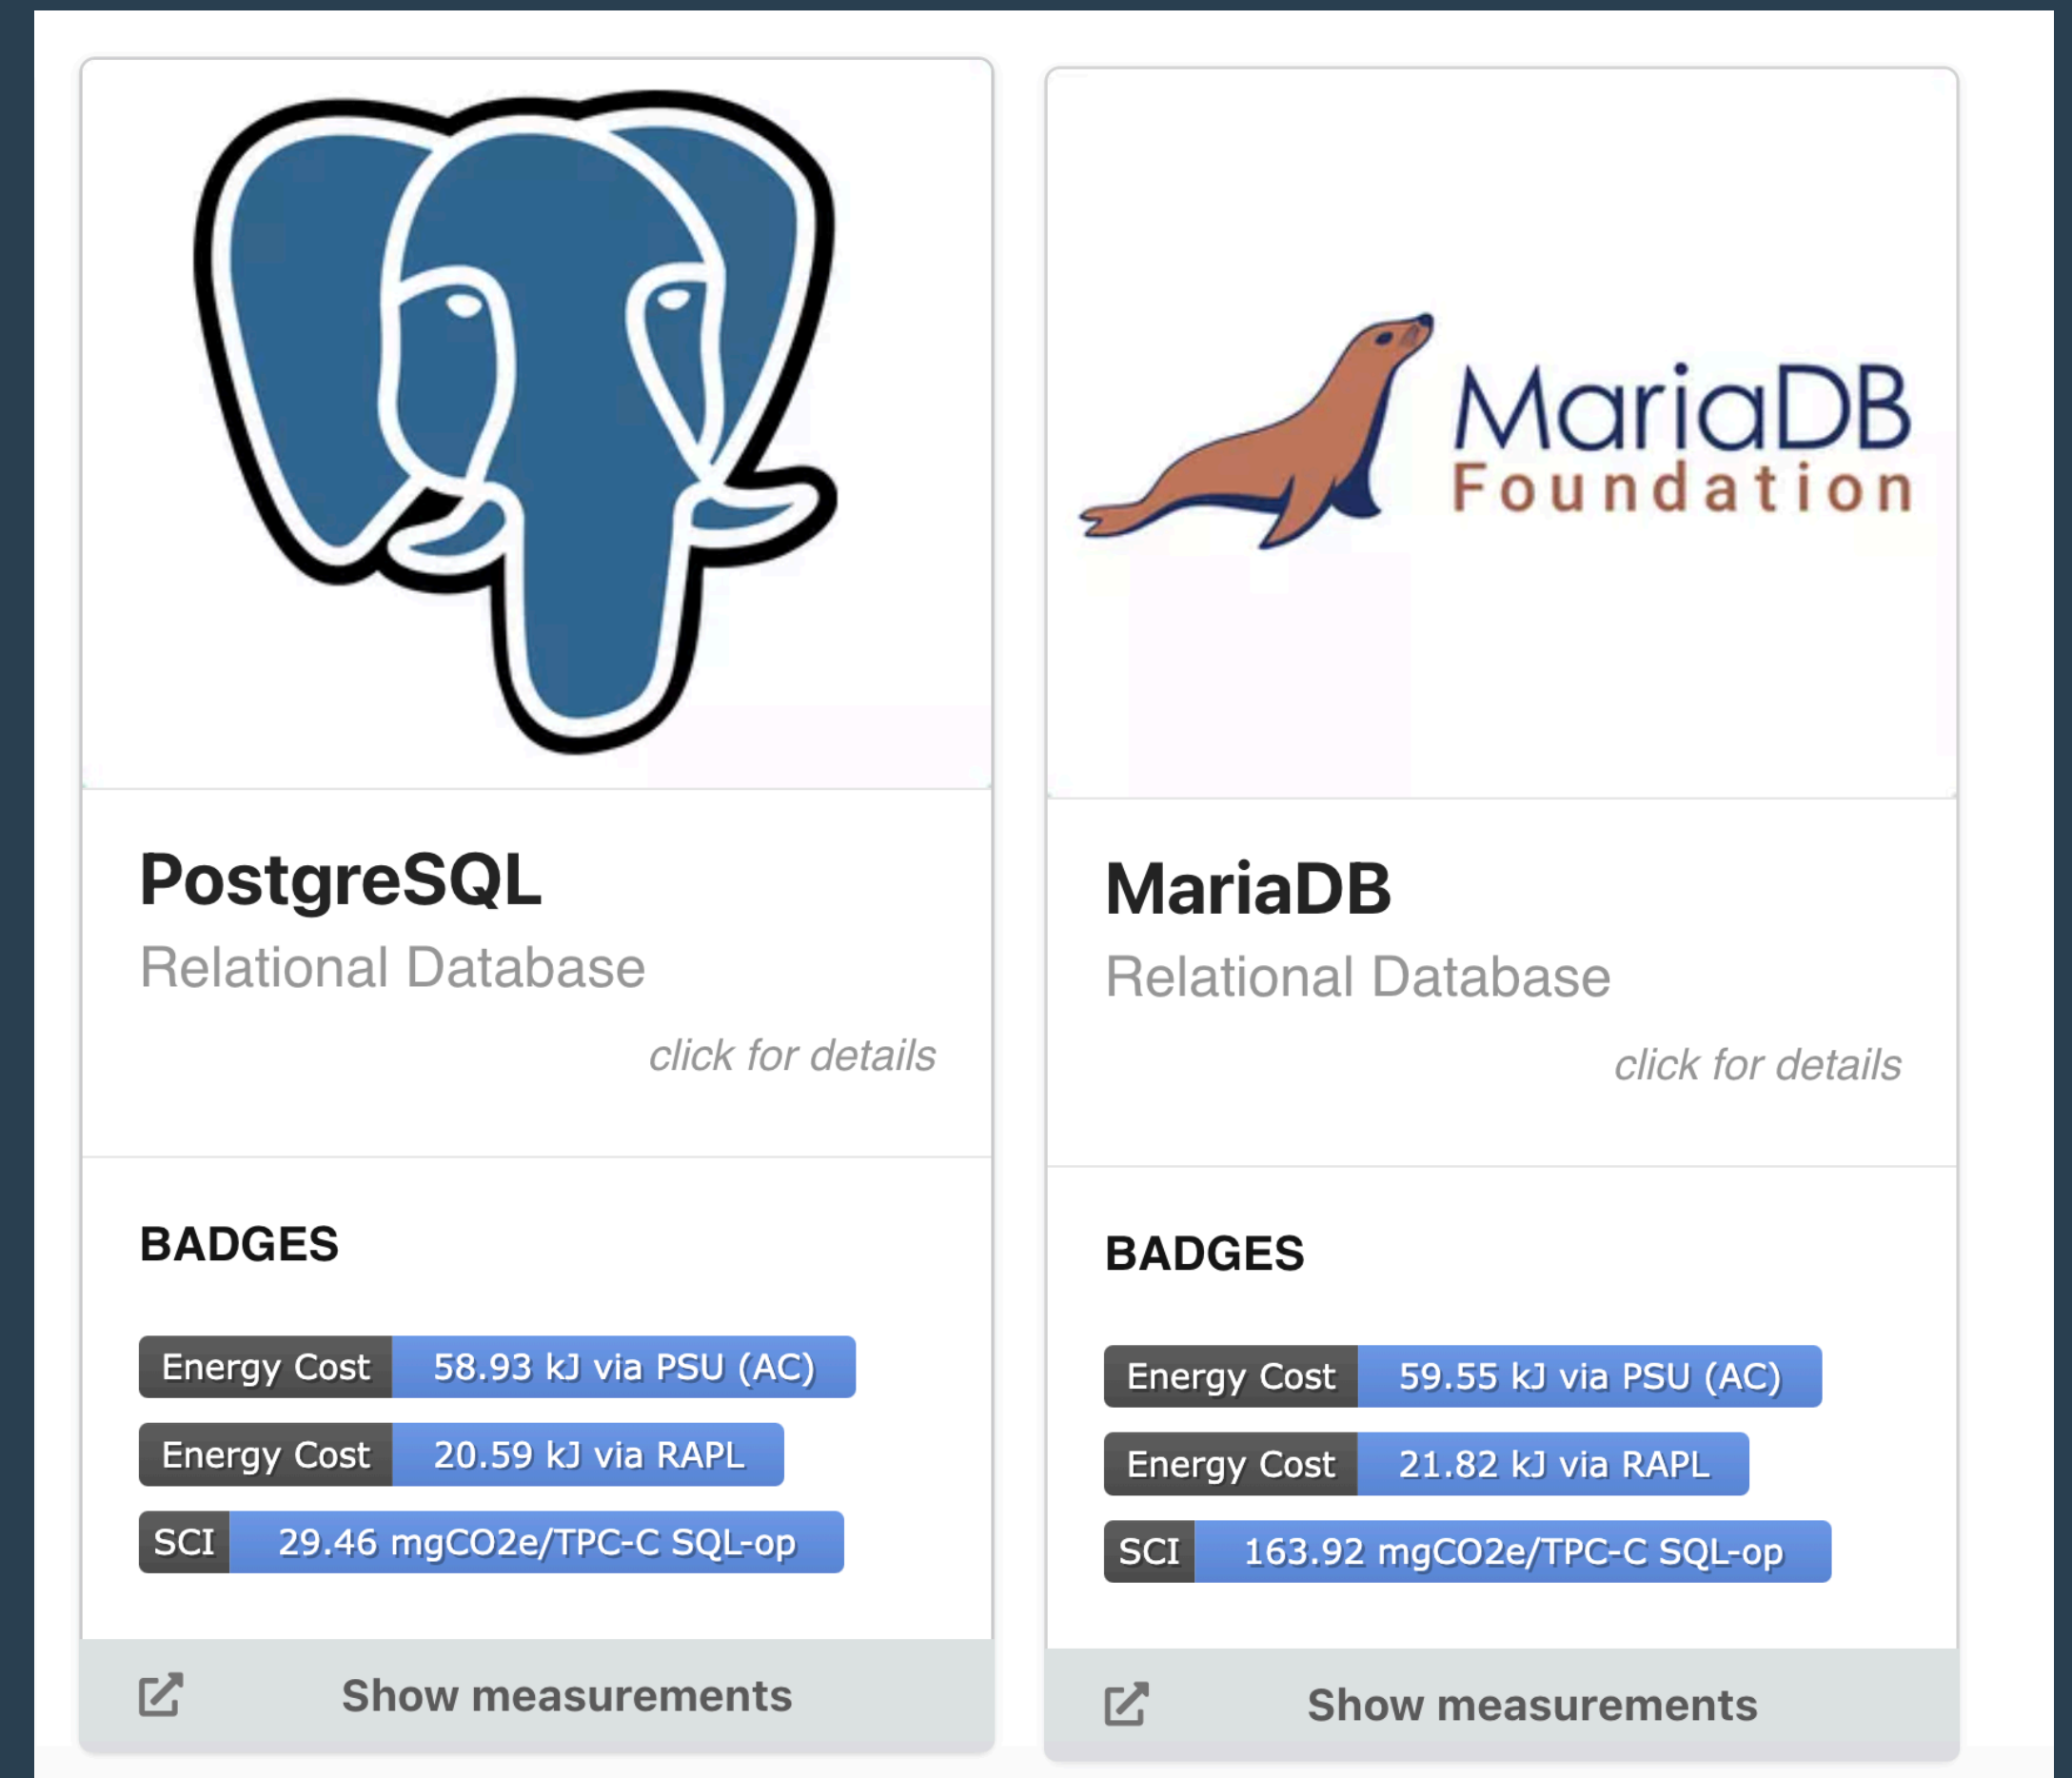
Potentials and levers

In saving and mitigation

Potentials in dependencies / components

Quantifying dependencies and toolchains in terms of carbon


- Software carbon and energy transparency
- Different software products for same use-case can have vastly different energy and carbon profiles
- Levers:
 - Library / Dependency selection
 - Default software recommendations



Potentials in Architectures

My website is already loading below 1s and very efficient

- Fast \neq efficient
 - Although a website might be fine in Google Pagespeed saving potentials are still enormous.
- Case Study: **Wordpress vs. HUGO**
 - Saving is 5x per Request
 - You can rebuild the page 5x for EVERY VISITOR and still would only break even.
 - On a small startup website (100.000 Page Views / month) this saving is already 2.4 kg!
- **Take away:** Perceived effectiveness is not carbon effectiveness



<https://www.green-coding.berlin/case-studies/wordpress-vs-hugo-cloudflare/>


Programming language potentials

- 75x in worst case!

Table 4. Normalized global results for Energy, Time, and Memory

Total					
	Energy		Time		Mb
(c) C	1.00	(c) C	1.00	(c) Pascal	1.00
(c) Rust	1.03	(c) Rust	1.04	(c) Go	1.05
(c) C++	1.34	(c) C++	1.56	(c) C	1.17
(c) Ada	1.70	(c) Ada	1.85	(c) Fortran	1.24
(v) Java	1.98	(v) Java	1.89	(c) C++	1.34
(c) Pascal	2.14	(c) Chapel	2.14	(c) Ada	1.47
(c) Chapel	2.18	(c) Go	2.83	(c) Rust	1.54
(v) Lisp	2.27	(c) Pascal	3.02	(v) Lisp	1.92
(c) Ocaml	2.40	(c) Ocaml	3.09	(c) Haskell	2.45
(c) Fortran	2.52	(v) C#	3.14	(i) PHP	2.57
(c) Swift	2.79	(v) Lisp	3.40	(c) Swift	2.71
(c) Haskell	3.10	(c) Haskell	3.55	(i) Python	2.80
(v) C#	3.14	(c) Swift	4.20	(c) Ocaml	2.82
(c) Go	3.23	(c) Fortran	4.20	(v) C#	2.85
(i) Dart	3.83	(v) F#	6.30	(i) Hack	3.34
(v) F#	4.13	(i) JavaScript	6.52	(v) Racket	3.52
(i) JavaScript	4.45	(i) Dart	6.67	(i) Ruby	3.97
(v) Racket	7.91	(v) Racket	11.27	(c) Chapel	4.00
(i) TypeScript	21.50	(i) Hack	26.99	(v) F#	4.25
(i) Hack	24.02	(i) PHP	27.64	(i) JavaScript	4.59
(i) PHP	29.30	(v) Erlang	36.71	(i) TypeScript	4.69
(v) Erlang	42.23	(i) Jruby	43.44	(v) Java	6.01
(i) Lua	45.98	(i) TypeScript	46.20	(i) Perl	6.62
(i) Jruby	46.54	(i) Ruby	59.34	(i) Lua	6.72
(i) Ruby	69.91	(i) Perl	65.79	(v) Erlang	7.20
(i) Python	75.88	(i) Python	71.90	(i) Dart	8.64
(i) Perl	79.58	(i) Lua	82.91	(i) Jruby	19.84

[17] Greenlab Study

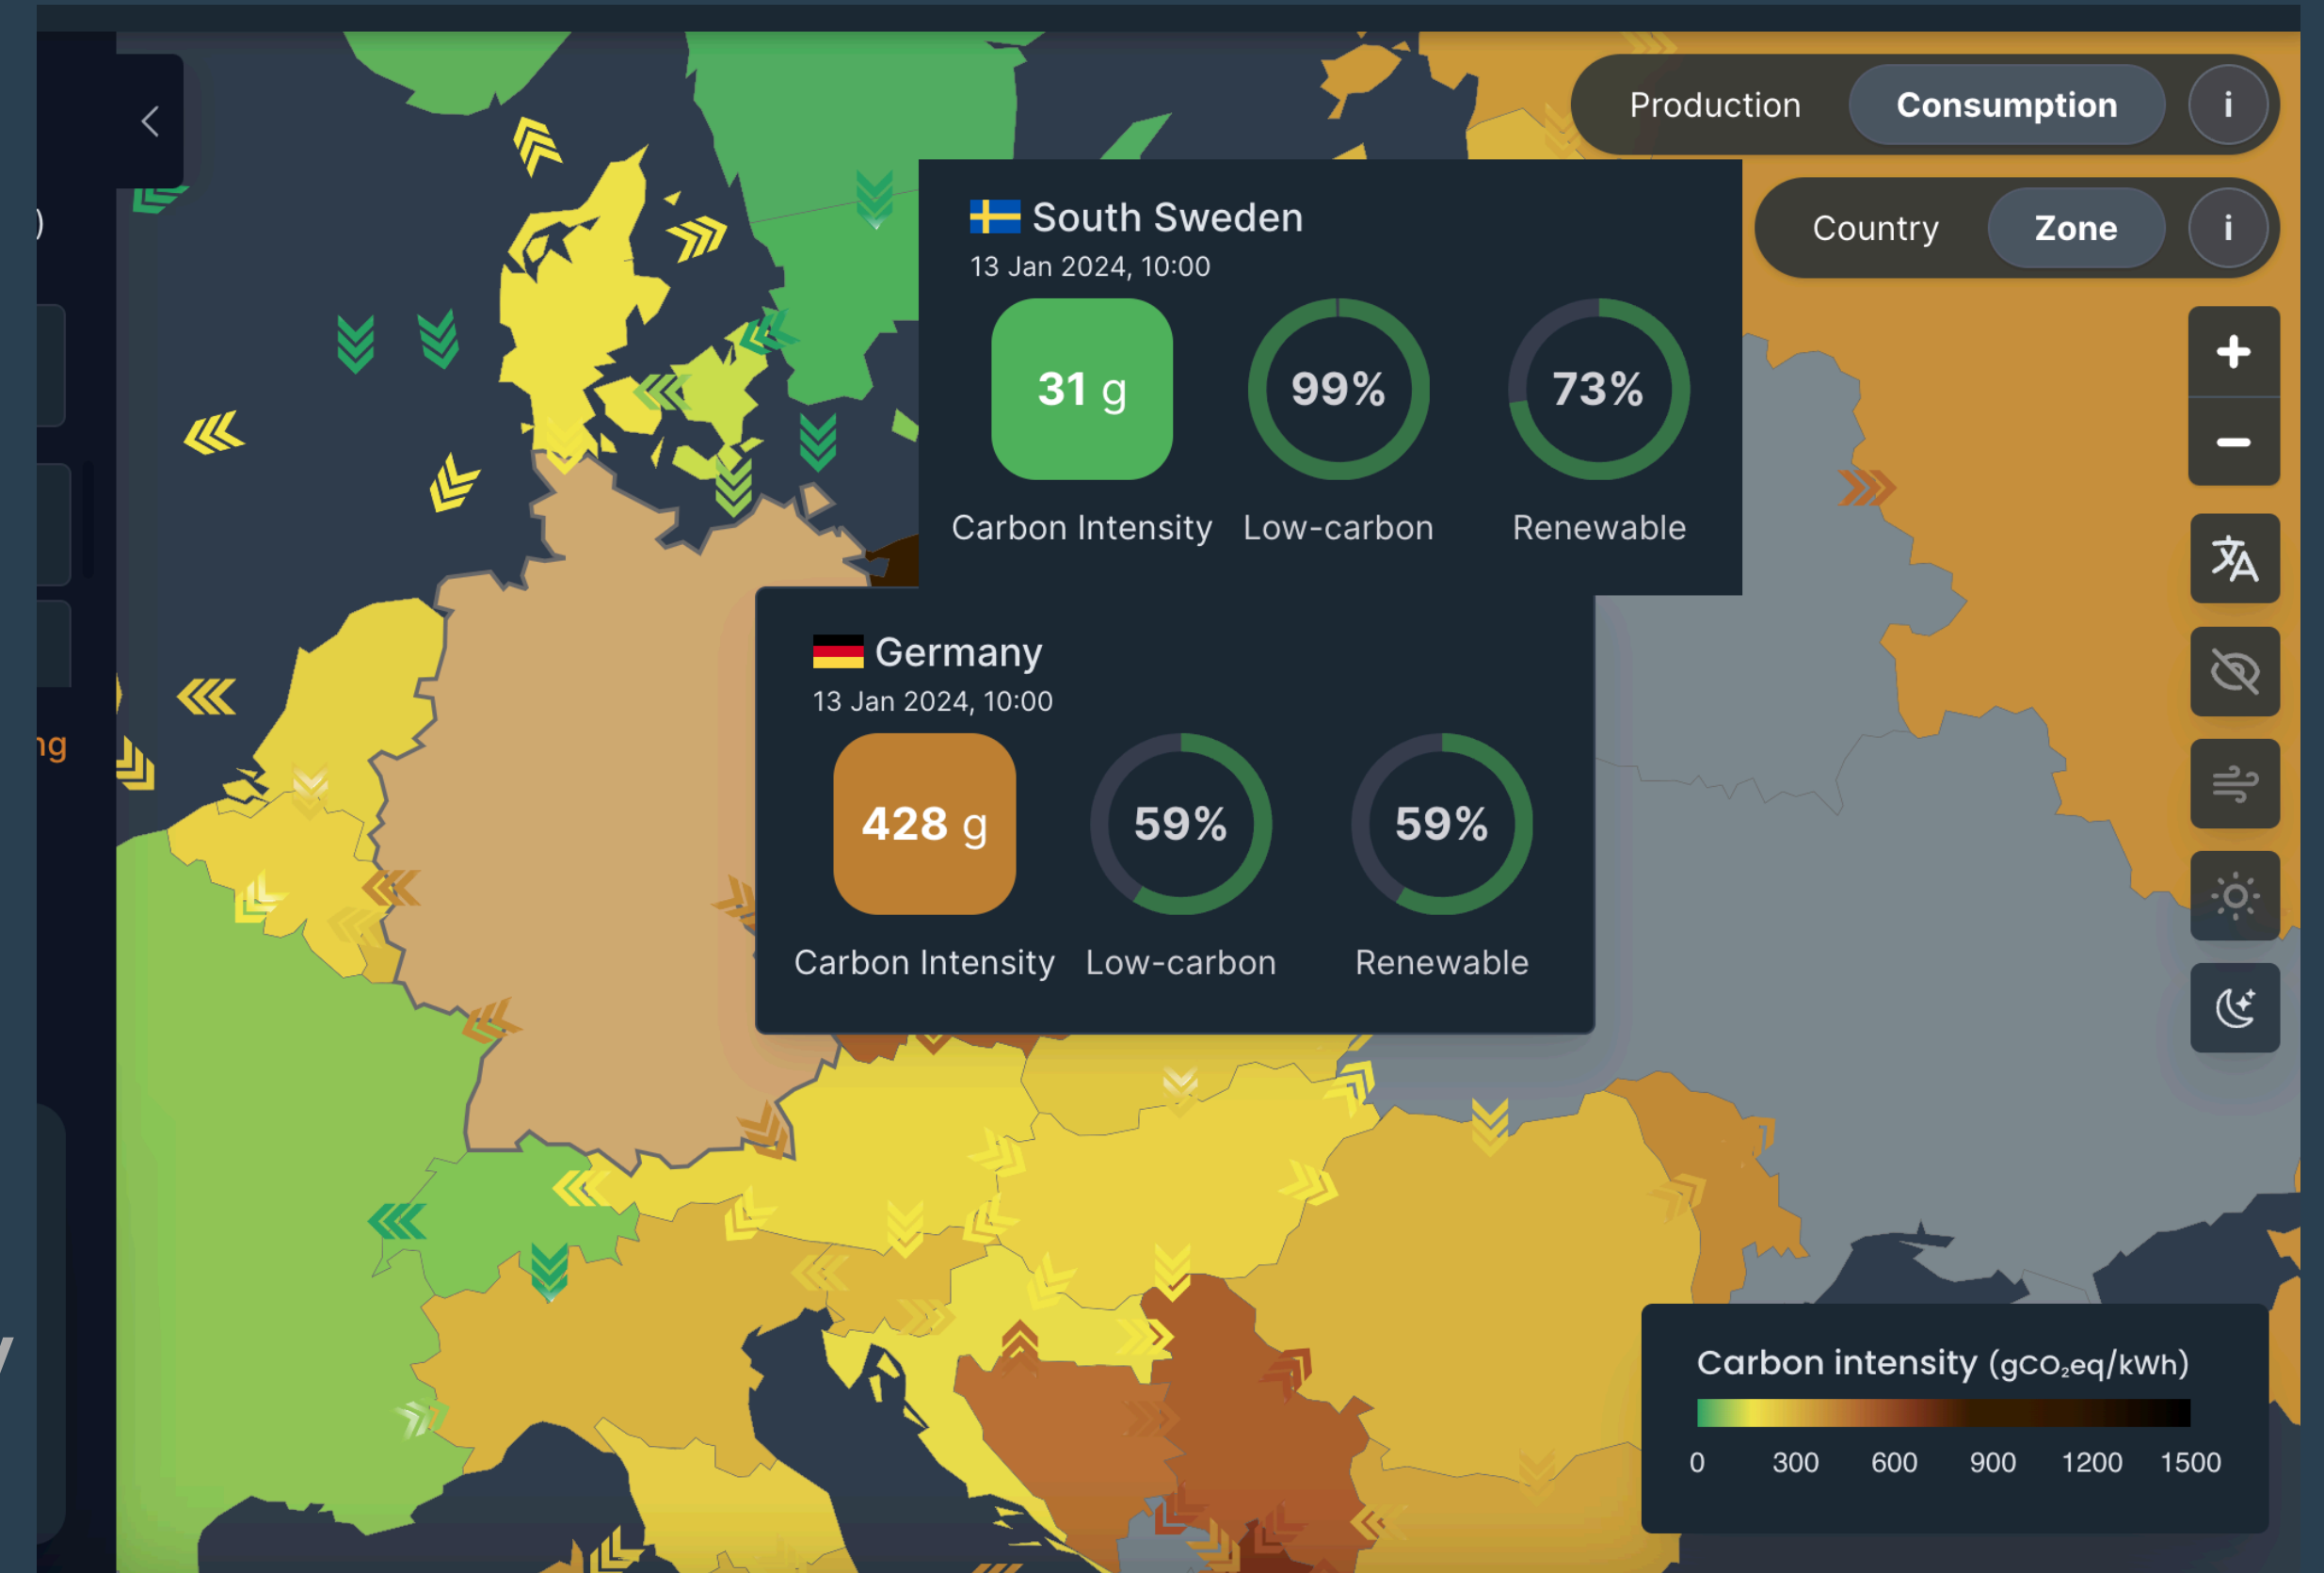


[18] Microsoft article on energy use of programming languages
1.8x - 3x difference

Potentials of Location Shifting

Example: Europe


- Moving a server to a low carbon region can save CO2
- Sweden / Germany is > 10x at times
- Moving workloads inside europe / US very easy due to low compliance risk
- Moving workload internationally tricky



Potentials of different hardware configuration


Power behaviour of hardware and how to leverage non-linearity. Cloud and non-cloud!

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)

Summary

Green Coding is a software technique FOR software

- Efficiency gains through performance engineering
- Making hardware last longer through software optimizations
- Using special CO2 reduction techniques through using green energy wherever possible
 - load-shifting
 - time-shifting
 - Demand-shaping
- Reduction of compute / data storage / network / complexity

Thank you

Want to know more

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

Thank you

Want to know more

- Website / Blog / Newsletter: <https://www.green-coding.io>
- Demo Open Data Repository: <https://metrics.green-coding.io>

- Unsere Projekte: <https://www.green-coding.io/#projects>
- Unsere Case-Studies: <https://www.green-coding.io/case-studies>

- Meetup Gruppe (Berlin): <https://www.meetup.com/green-coding>

- <https://www.linkedin.com/in/arne-tarara> / arne@green-coding.io

- Interested in more content and networking -> Join our conference!

Thank you

Want to know more



EcoCompute 2024

25-26 April - Munich

The first engineering conference on sustainability in hardware & software.

www.eco-compute.io

Thank you

Want to know more

Line-up



Hagen Pfeifer



Arne Tarara



Sasche Böhme




Didi Hoffmann



Yelle Lieder



David Kopp



Max Körbacher



Anna Zagorski




Max Schulze



Verena Majuntke



Andreas Brunnert



Petteri Kivimäki



Anita Schüttler



Ghazal Aakel & Eric Jochum



Backup

Potentials


For industries

Distribution of scope emissions

Example GoCardless

FIGURE 8

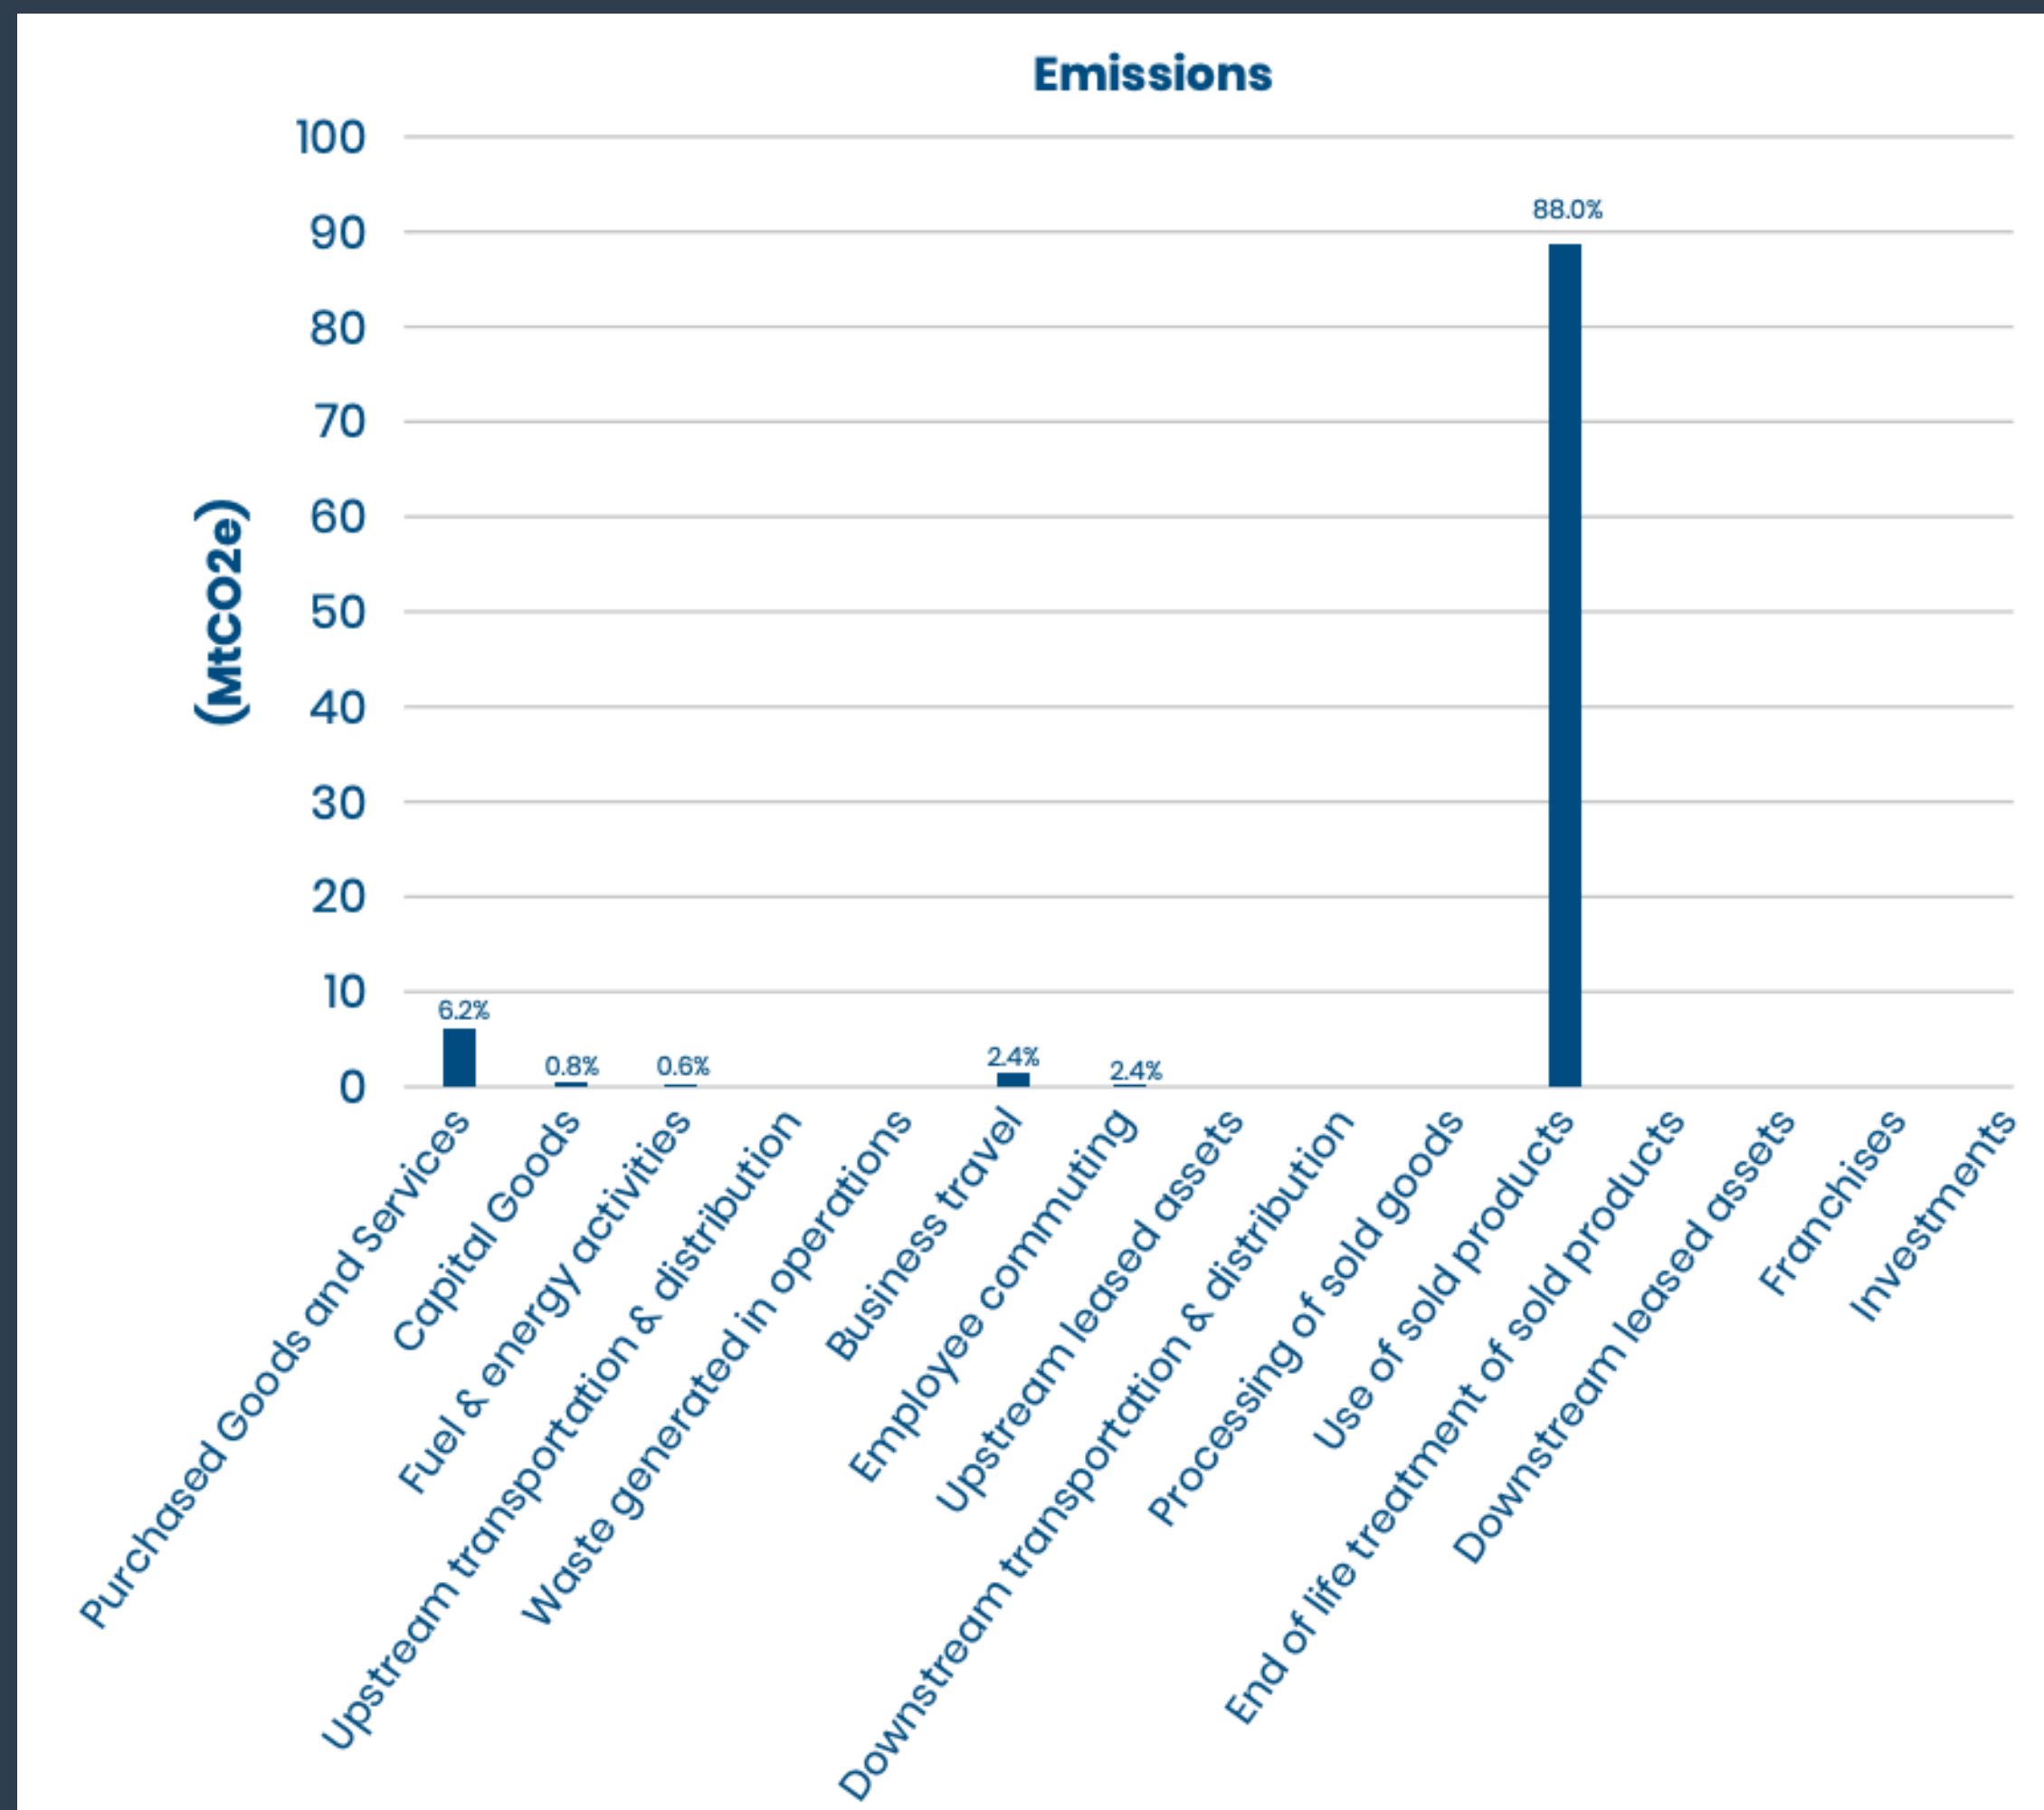
Scope 3 emissions reported by two software companies, Gocardless (2021) and AVEVA (2022)



- GoCardless is an online payment processor
- If you are a SaaS company your emissions are mostly Scope 3 upstream
- They originate from buying cloud services

Distribution of scope emissions

Example AVEVA



- AVEVA is an industrial software company.
- If you are an software integrator your emissions are mostly Scope 3 downstream
- They originate through clients using your software in production

ABN Amro - Datacenter emissions

Bank from the Netherlands - 11.45 Billion EUR market cap.

- Datacenters and SaaS account for 11.89 kTon CO2
- ~30% of Scope 3
- ~25% of total CO2e

3.1 Summary of carbon emissions within ABN AMRO scope			
Reported kton GHG emissions by region			2022
Scope 1 Tank-to-Wheel (TTW)	NL	RoW*	Total
Energy (natural gas + solar PV + other)	-	0,52	0,52
Mobility (lease cars - internal combustion engine)	2,18	-	2,18
Total scope 1	2,18	0,52	2,70
Scope 2 Tank-to-Wheel (TTW)**			
Energy (electricity, heating and cooling)	0,65	2,28	2,93
Mobility (lease - electric vehicles)	0	-	-
Total scope 2	0,65	2,28	2,93
Total scopes 1 + 2	2,83	2,80	5,63
Note: Decrease in scope 1 & 2 GHG emissions is mainly due to Covid-19 effects.			
			2022
Scope 3 Tank-to-Wheel (TTW) Own Operations	NL	RoW*	Total
Air travel, international business rail travel and hotel visits	1,93	1,37	3,30
Mobility (commuting in private vehicles and business travel)	5,28	-	5,28
Public transport	0,55	-	0,55
Home workplace	9,00	1,74	10,74
Off-premise datacenters + Software-as-a-Service	11,89	-	11,89
Total scope 3 TTW	28,65	3,11	31,77
Scope 3 Well-to-Tank (WTT) Own Operations	NL	RoW*	Total
Energy (electricity, gas, heating and cooling)	0,67	-	0,67
Home office (Netherlands)	2,64	-	2,64
Mobility (lease cars - internal combustion engine)	0,66	-	0,66
Mobility (lease - electric vehicles)	2,94	-	2,94
Private vehicles	1,73	-	1,73
Public transport	0,21	-	0,21
Total scope 3 WTT	8,85	-	8,85
Total Scope 3 emissions - Own operations	37,51	3,11	40,62

Source: <https://www.abnamro.com/en/information/sustainability-reporting-and-publications?selectedTabs=Reporting>

Allianz

Insurer - Germany - 93.73 Billion EUR market cap.

- Datacenters account for 325,000 TJ
- => 90 GWh
- => 45 kTons CO2e
- => 16% of total energy consumption
- Total CO2 not disclosed

Table ENV-3

Energy consumption

As of December 31		2022
Energy consumption from our office buildings	GJ	1,578,675
Energy consumption from our data centers	GJ	325,212
Total energy consumption¹	GJ	1,965,294


Source: https://www.allianz.com/content/dam/onemarketing/azcom/Allianz_com/sustainability/documents/Allianz_Group_Sustainability_Report_2022-web.pdf

- ~ 4 times the CO2 / 8-times the market cap

Potentials of different application classification

By looking at software via usage scenario

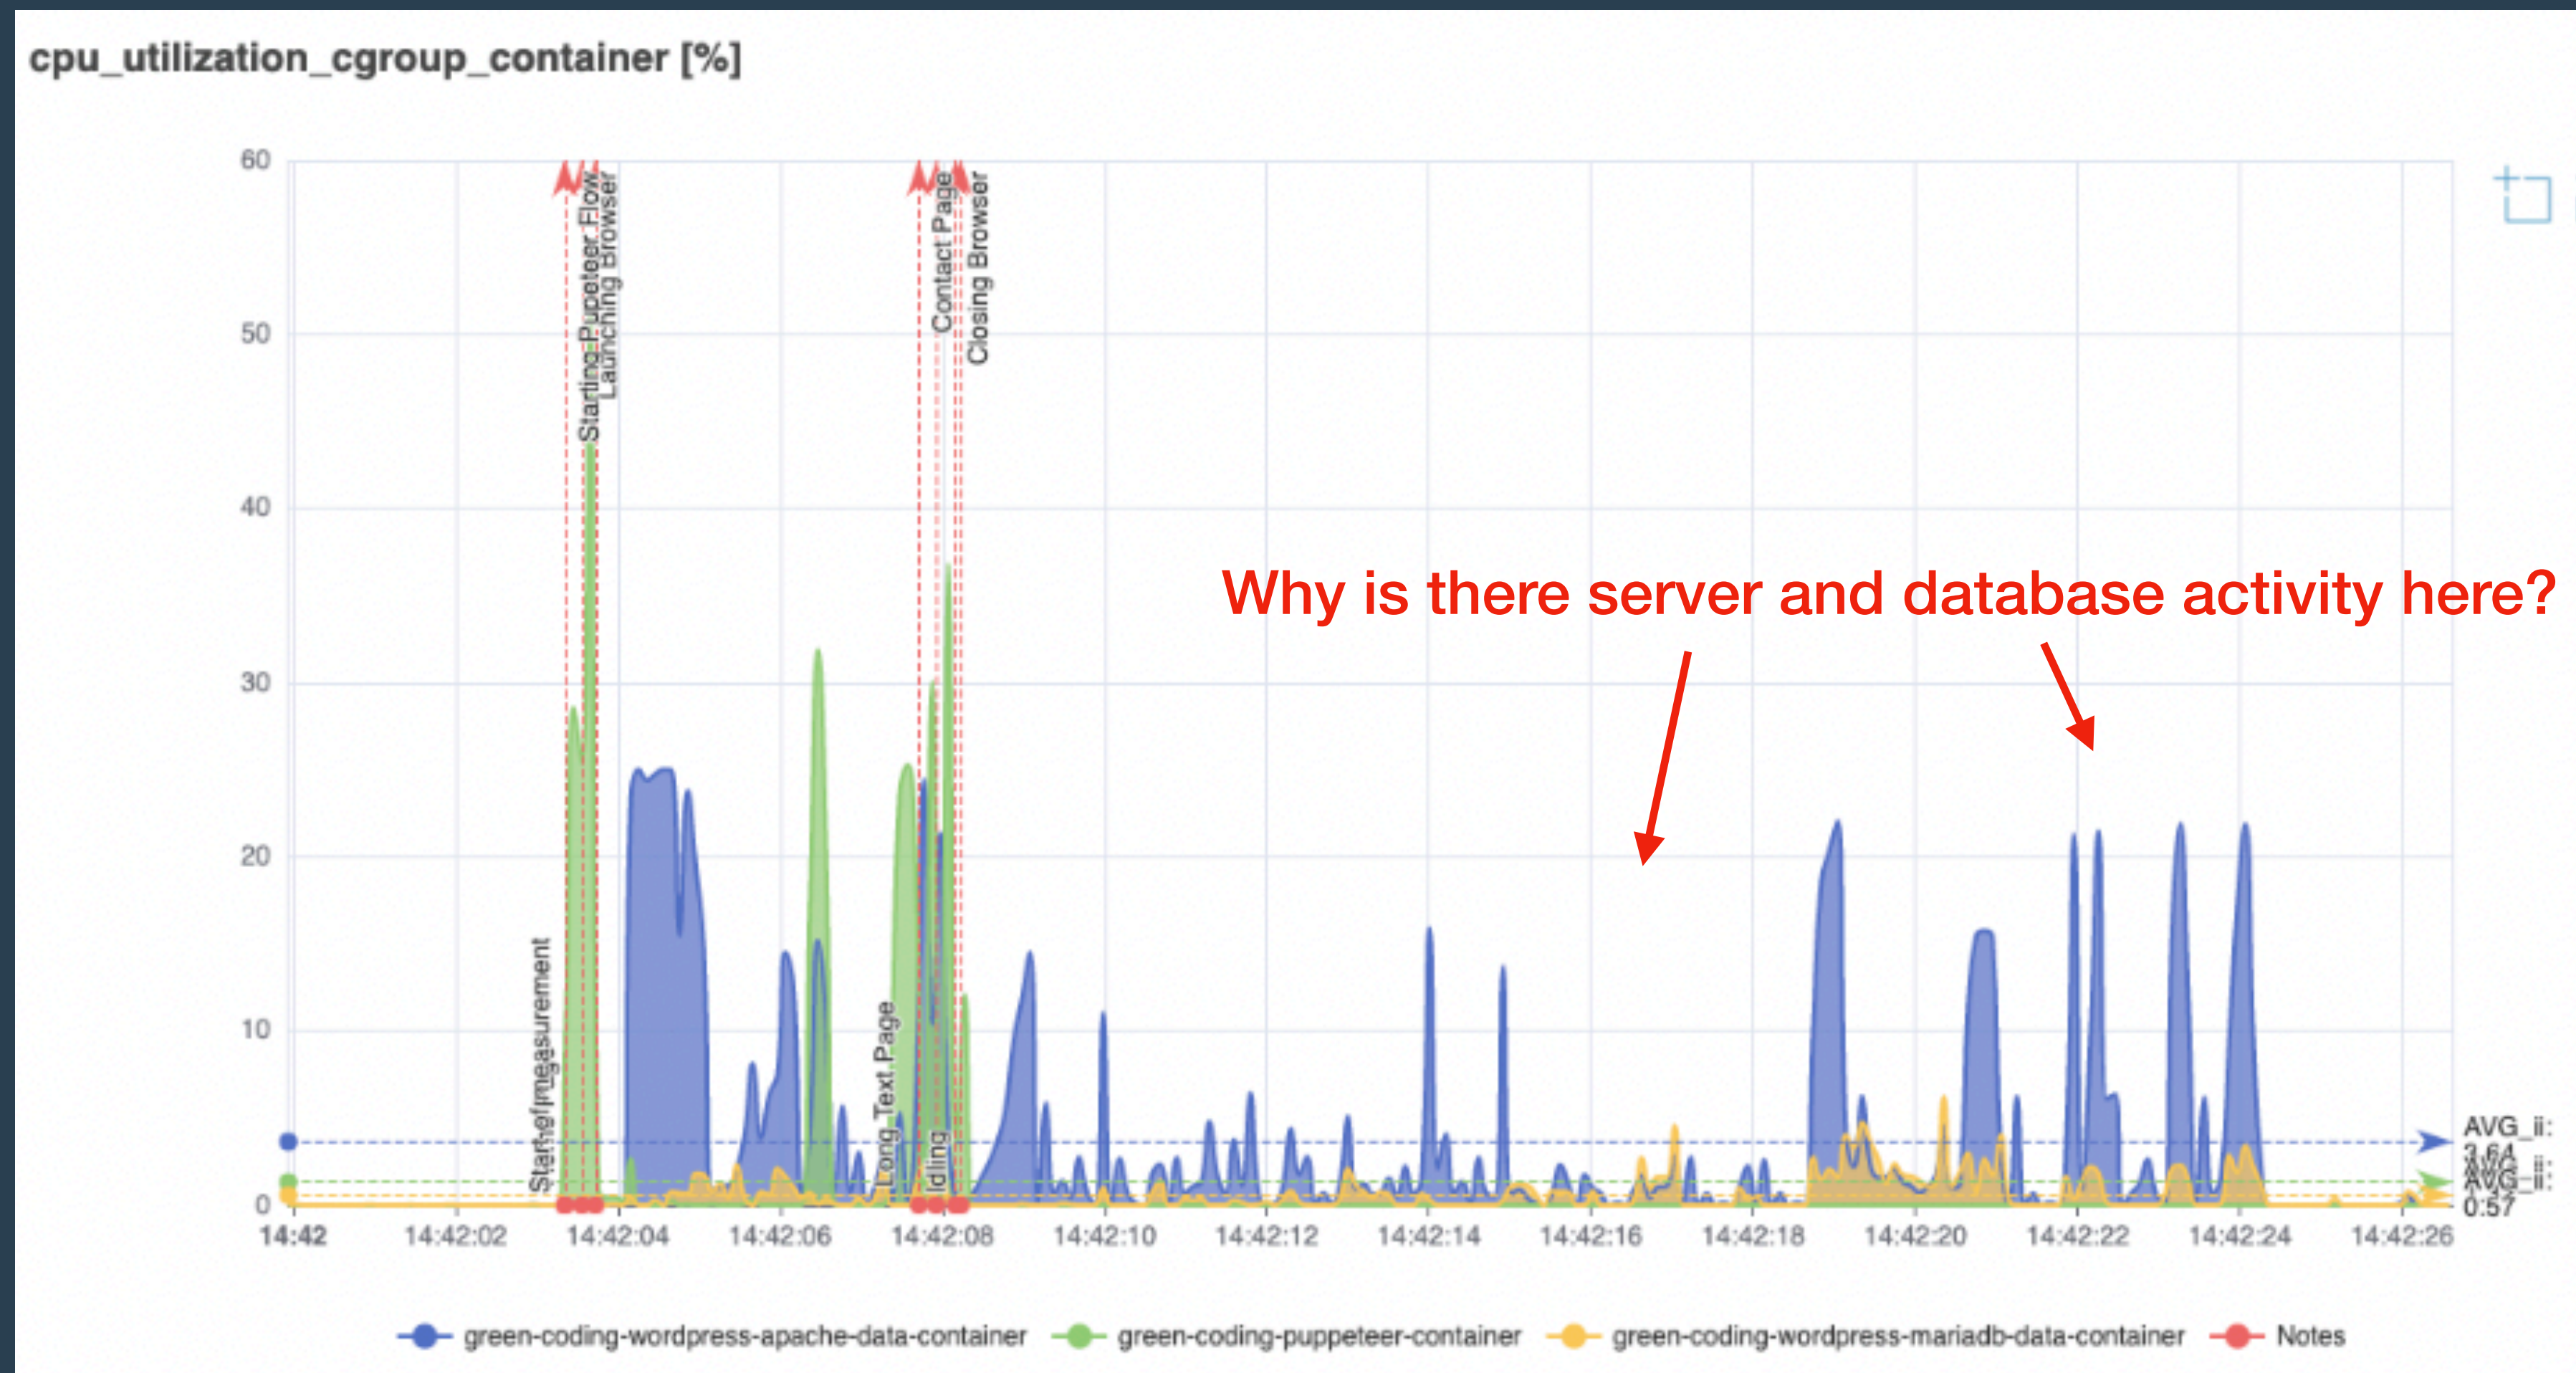
No actual data! Concept picture!



Energy consumption of Whatsapp vs. Telegram - per usage scenario

Potentials of simple energy bugs

In background application measurement data



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
- Main topics:
 - [climate justice](#)
 - [green energy](#)



Green Software Foundation

NGO for the sustainability of the digital sector



- UK founded NGO
- The goal is to reduce the CO2 emissions of software
- Main work is standardization and reference implementations
 - Software Carbon Intensity (ISO Standard)
 - Carbon Aware SDK
 - SCER - Software Carbon Efficiency Rating

SDIA

NGO for sustainability in the digital sector

- Association for stakeholders in the digital industry
 - datacenters / software vendors /
 - hardware manufacturers etc.
- Main topics:
 - Digital carbon footprint
 - Consume renewable energy
 - Increase utilization of hardware
 - Leaner Architecture of software
 - Extend hardware lifecycle



Open Source potentials

Due to scaling effects can be massive even with minimal effort

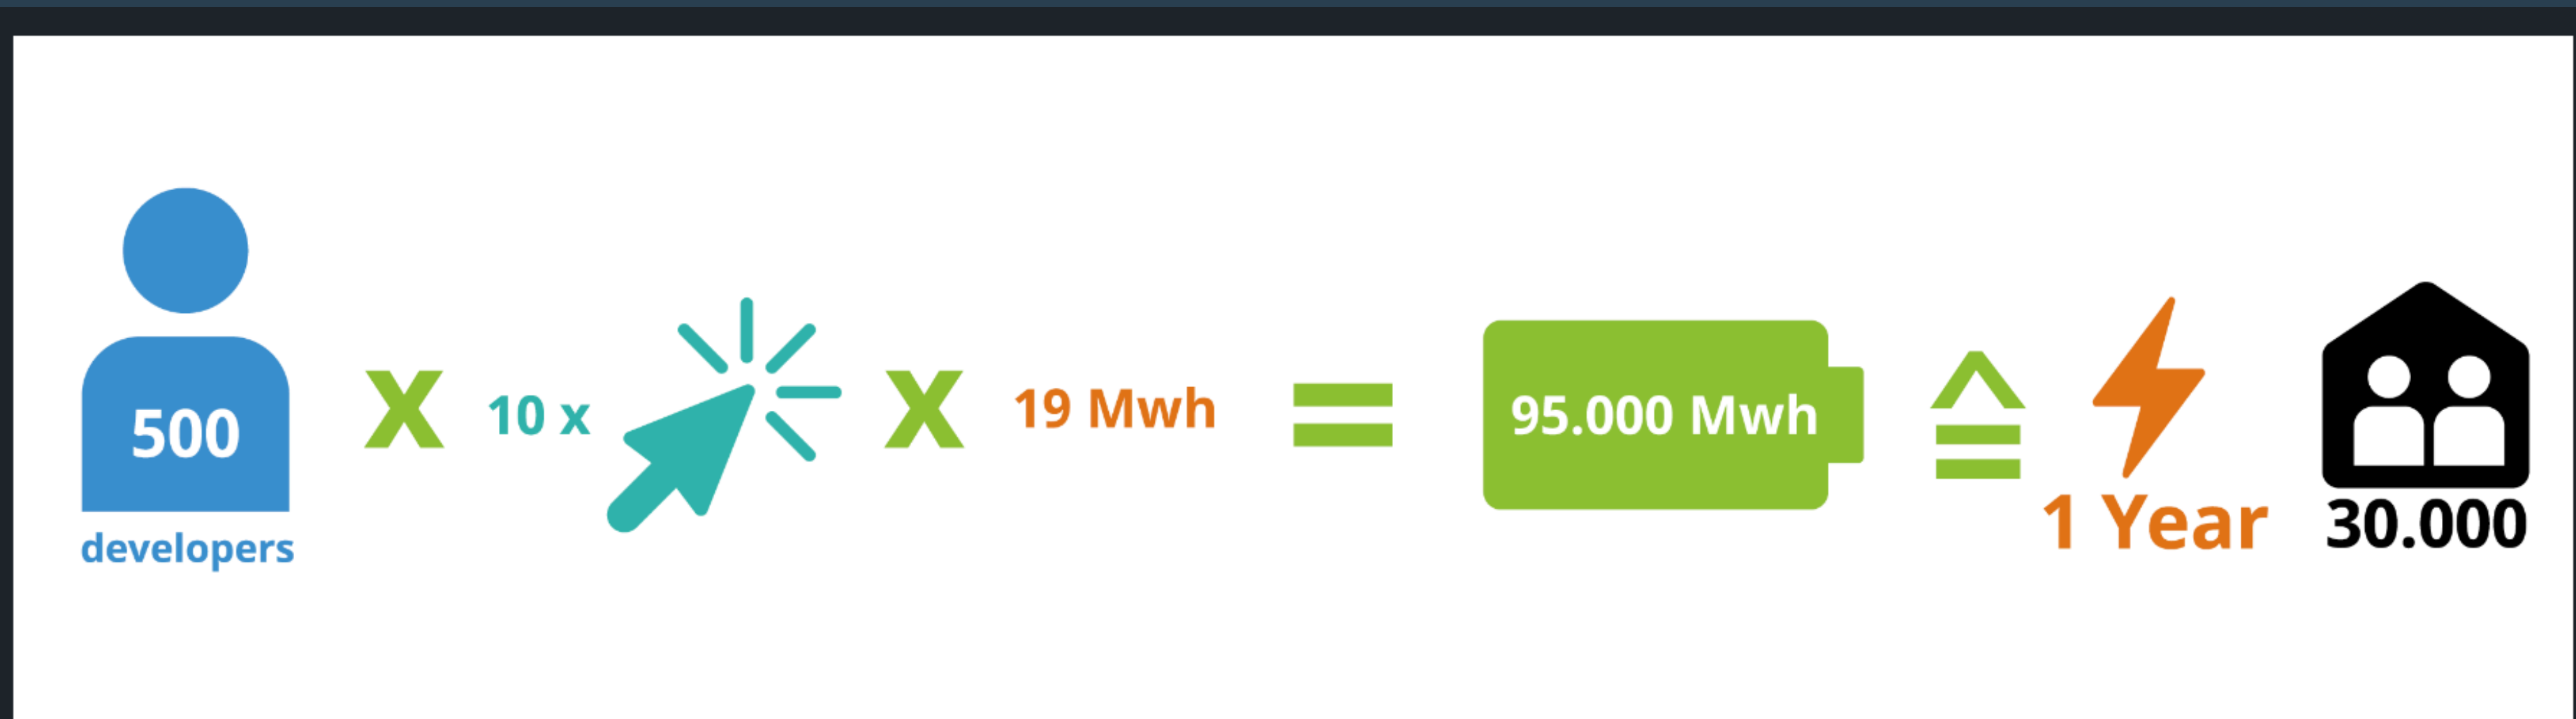
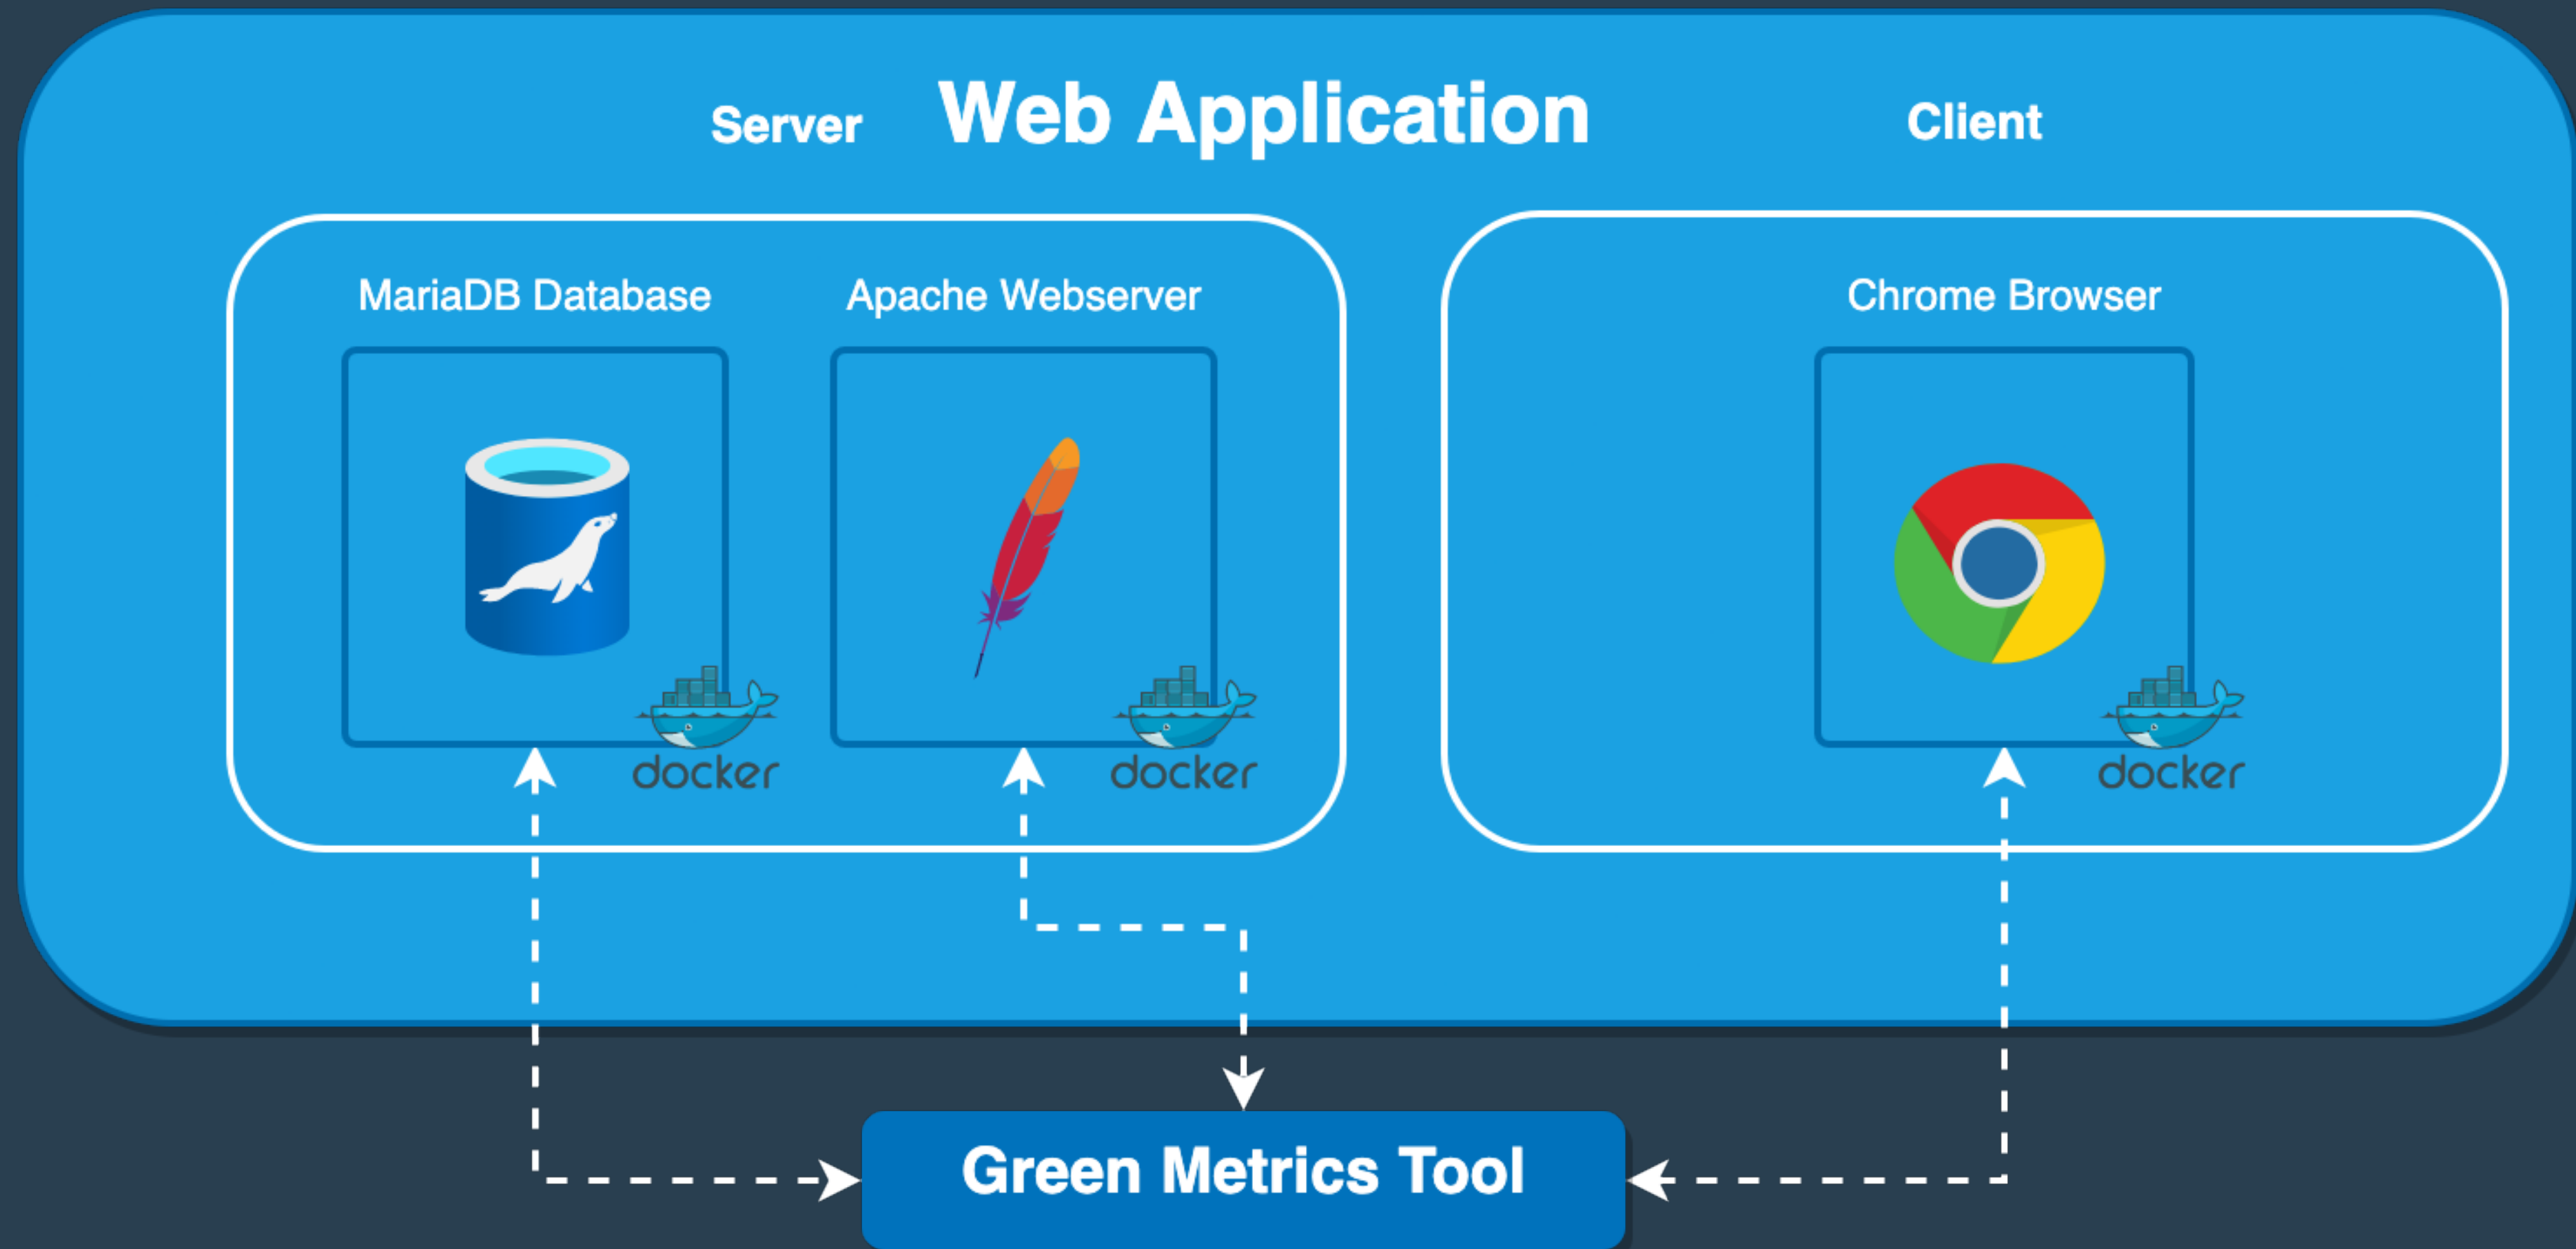


Figure : If 500 developers make 10 one CPU-second reductions, this is equal to 95 thousand megawatt-hours of savings, or the energy consumption of 30-thousand two-person households over one year. (Image from KDE published under a [CC-BY-SA-4.0](#) license. Cursor icon by Alice-vector licensed under a [CC-BY](#) license. Example from Detlef Thoms. Design by Lana Lutz.)

Green Metrics Tool

Aufbau einer Applikation für ein Usage Szenario - Client / Server



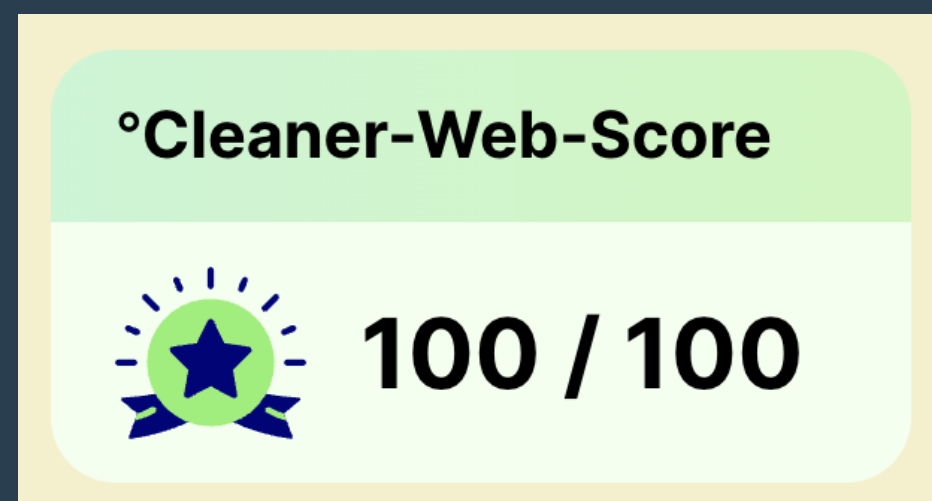
Labels and certifications

An excerpt

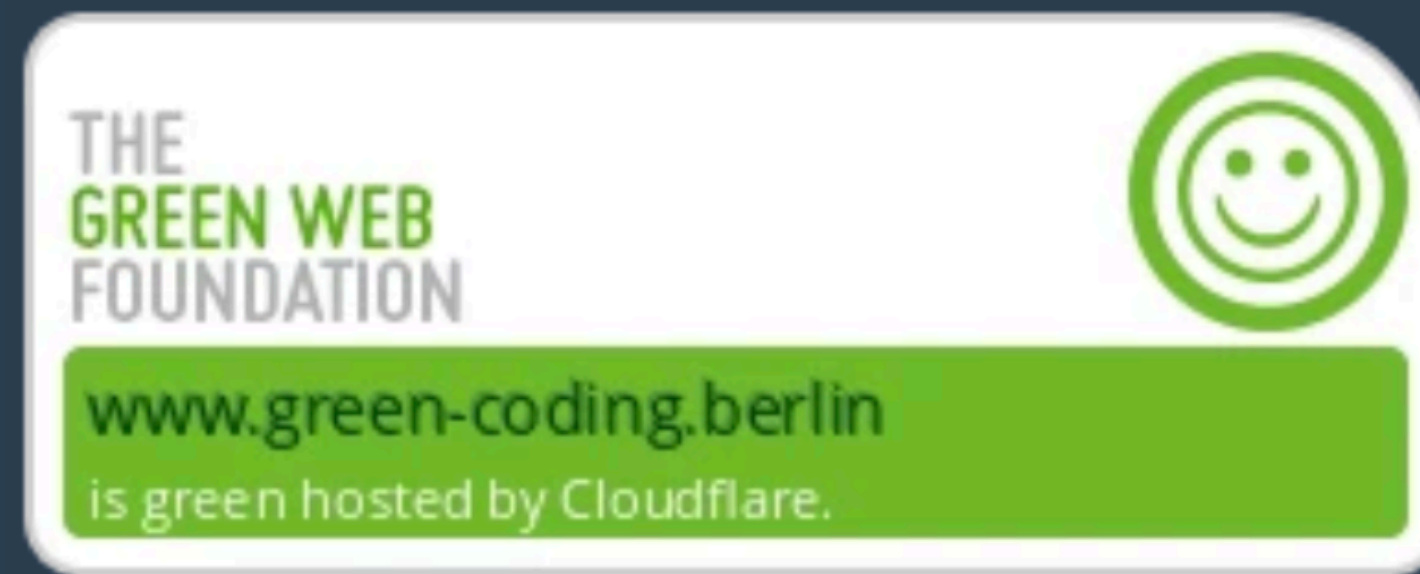
Labels for Green Software

Sparse ... only some really high profile / applicable



Website



Hosting




Application



Blue Angel - German Eco Label

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

- Reduce hardware obsolescence
- Increase transparency about software energy consumption
- Increase user autonomy
- Increase modularity and software re-use
- No CO2 concerns



Measurement

How do we measure?

Wie messen wir Energie?

An der Steckdose natürlich, ist doch klar!



Die Amazon Version

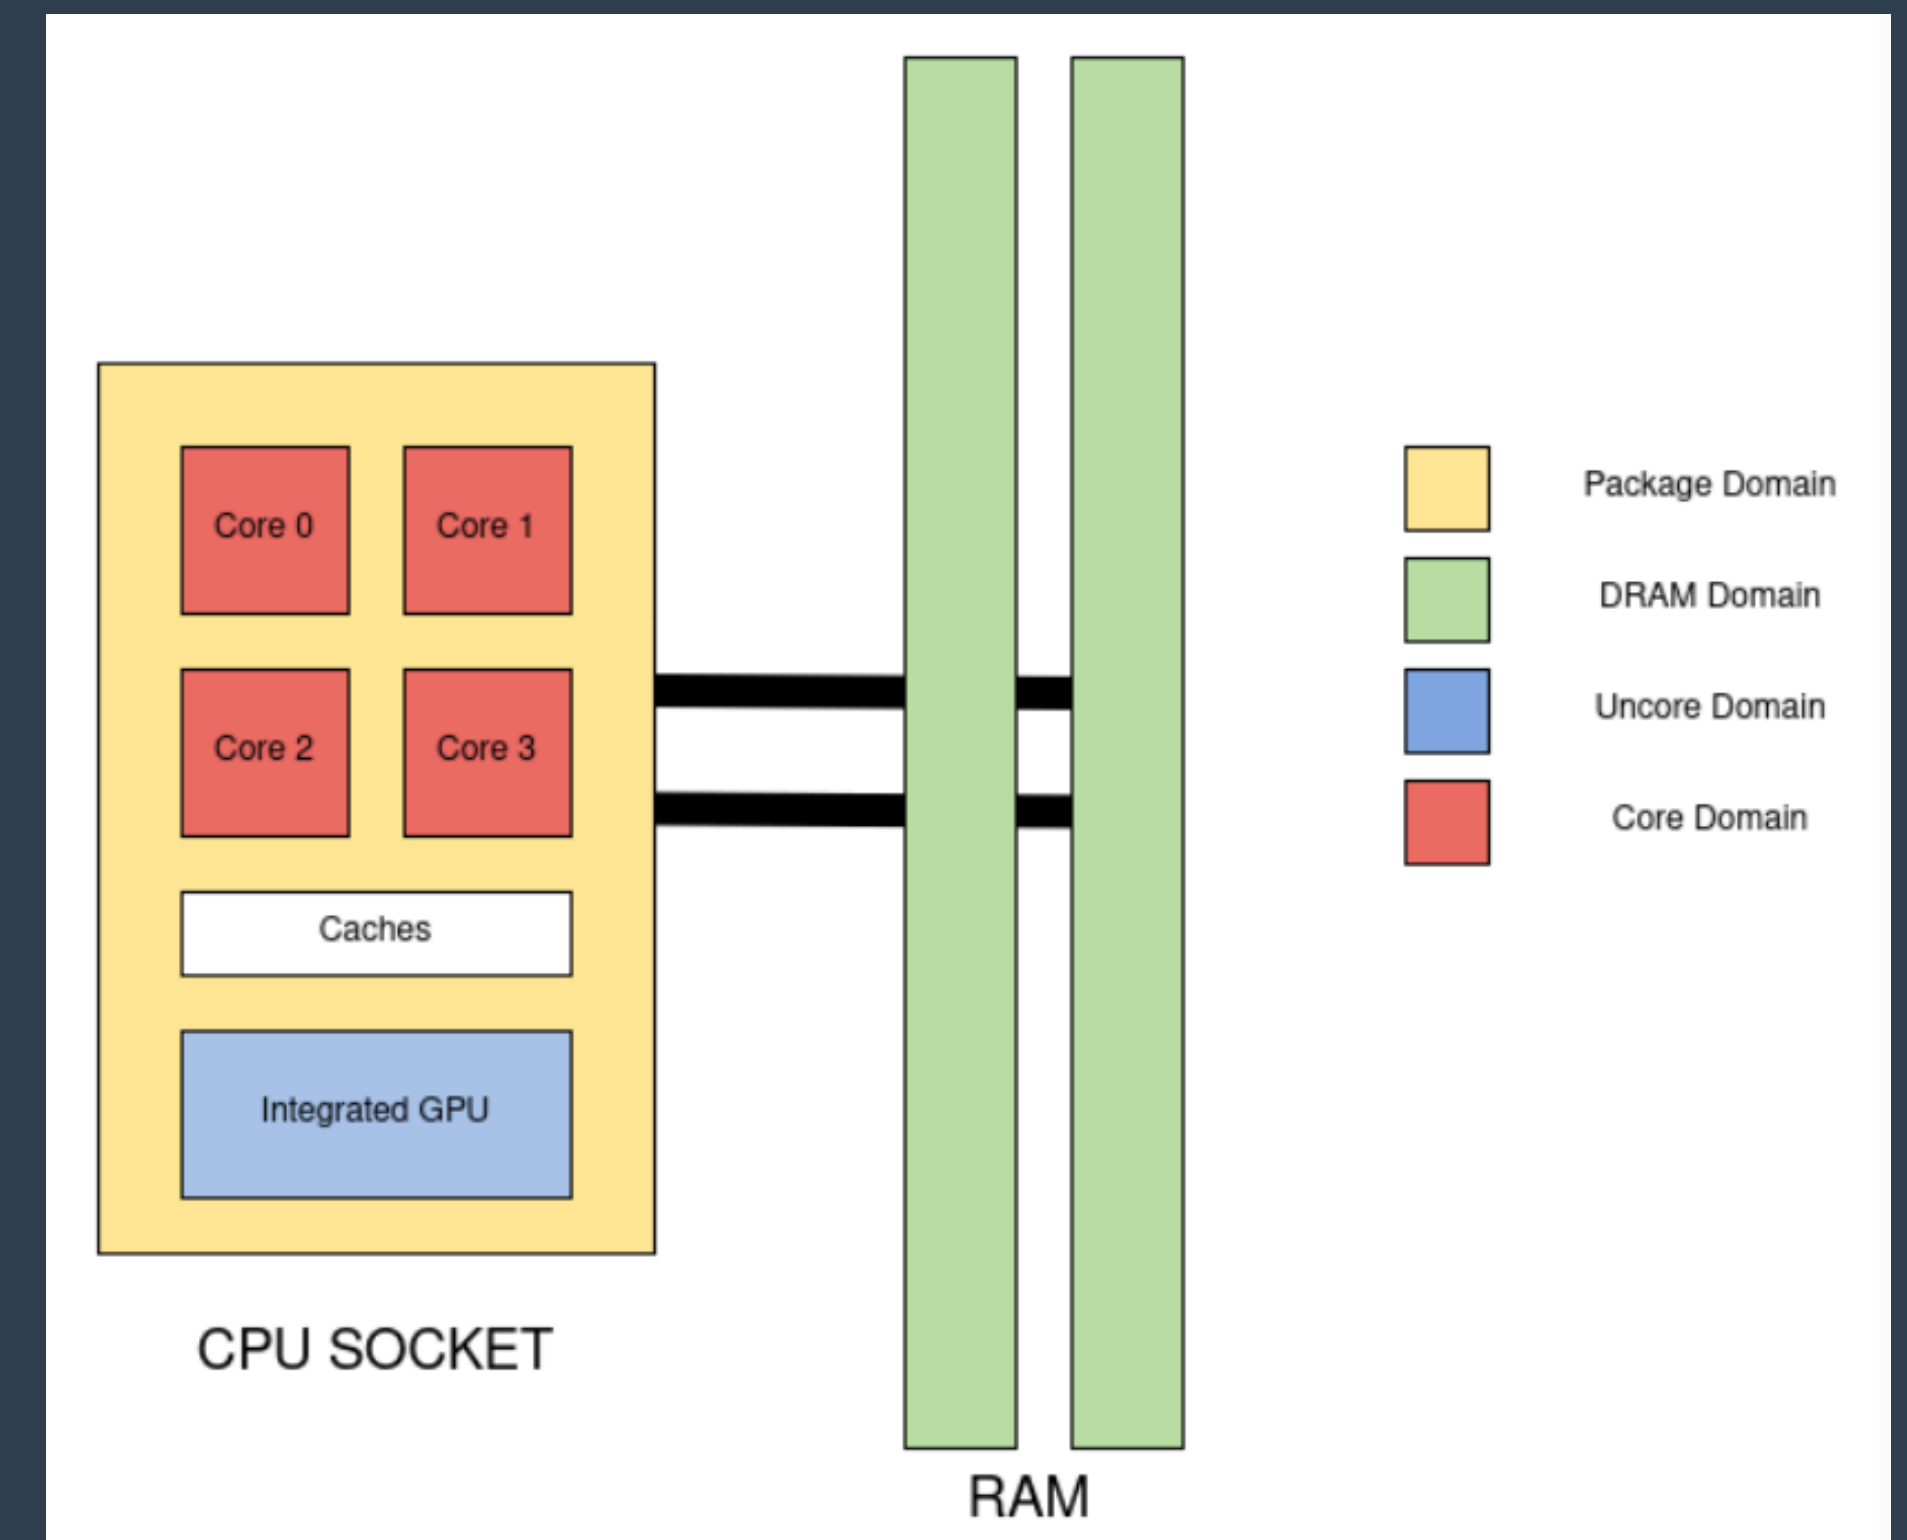


Die Laborversion

Wie messen wir Energie?

On-Board Hardware Interfaces


- Energiemessfunktionen für die meisten modernen Intel / AMD Prozessoren
- Messen:
 - CPU-Energie pro Kern/Paket
 - RAM
 - Integrierte GPU
- Software-Modell (proprietär - Spannungsregler etc.), aber mehrfach akademisch validiert
- Auflösung 1ms / 15,3 microJoules
- Bereitgestellt als Linux Kernel-Modul o. MSR




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

Wie messen wir Energie

Wie weit möchte man gehen?



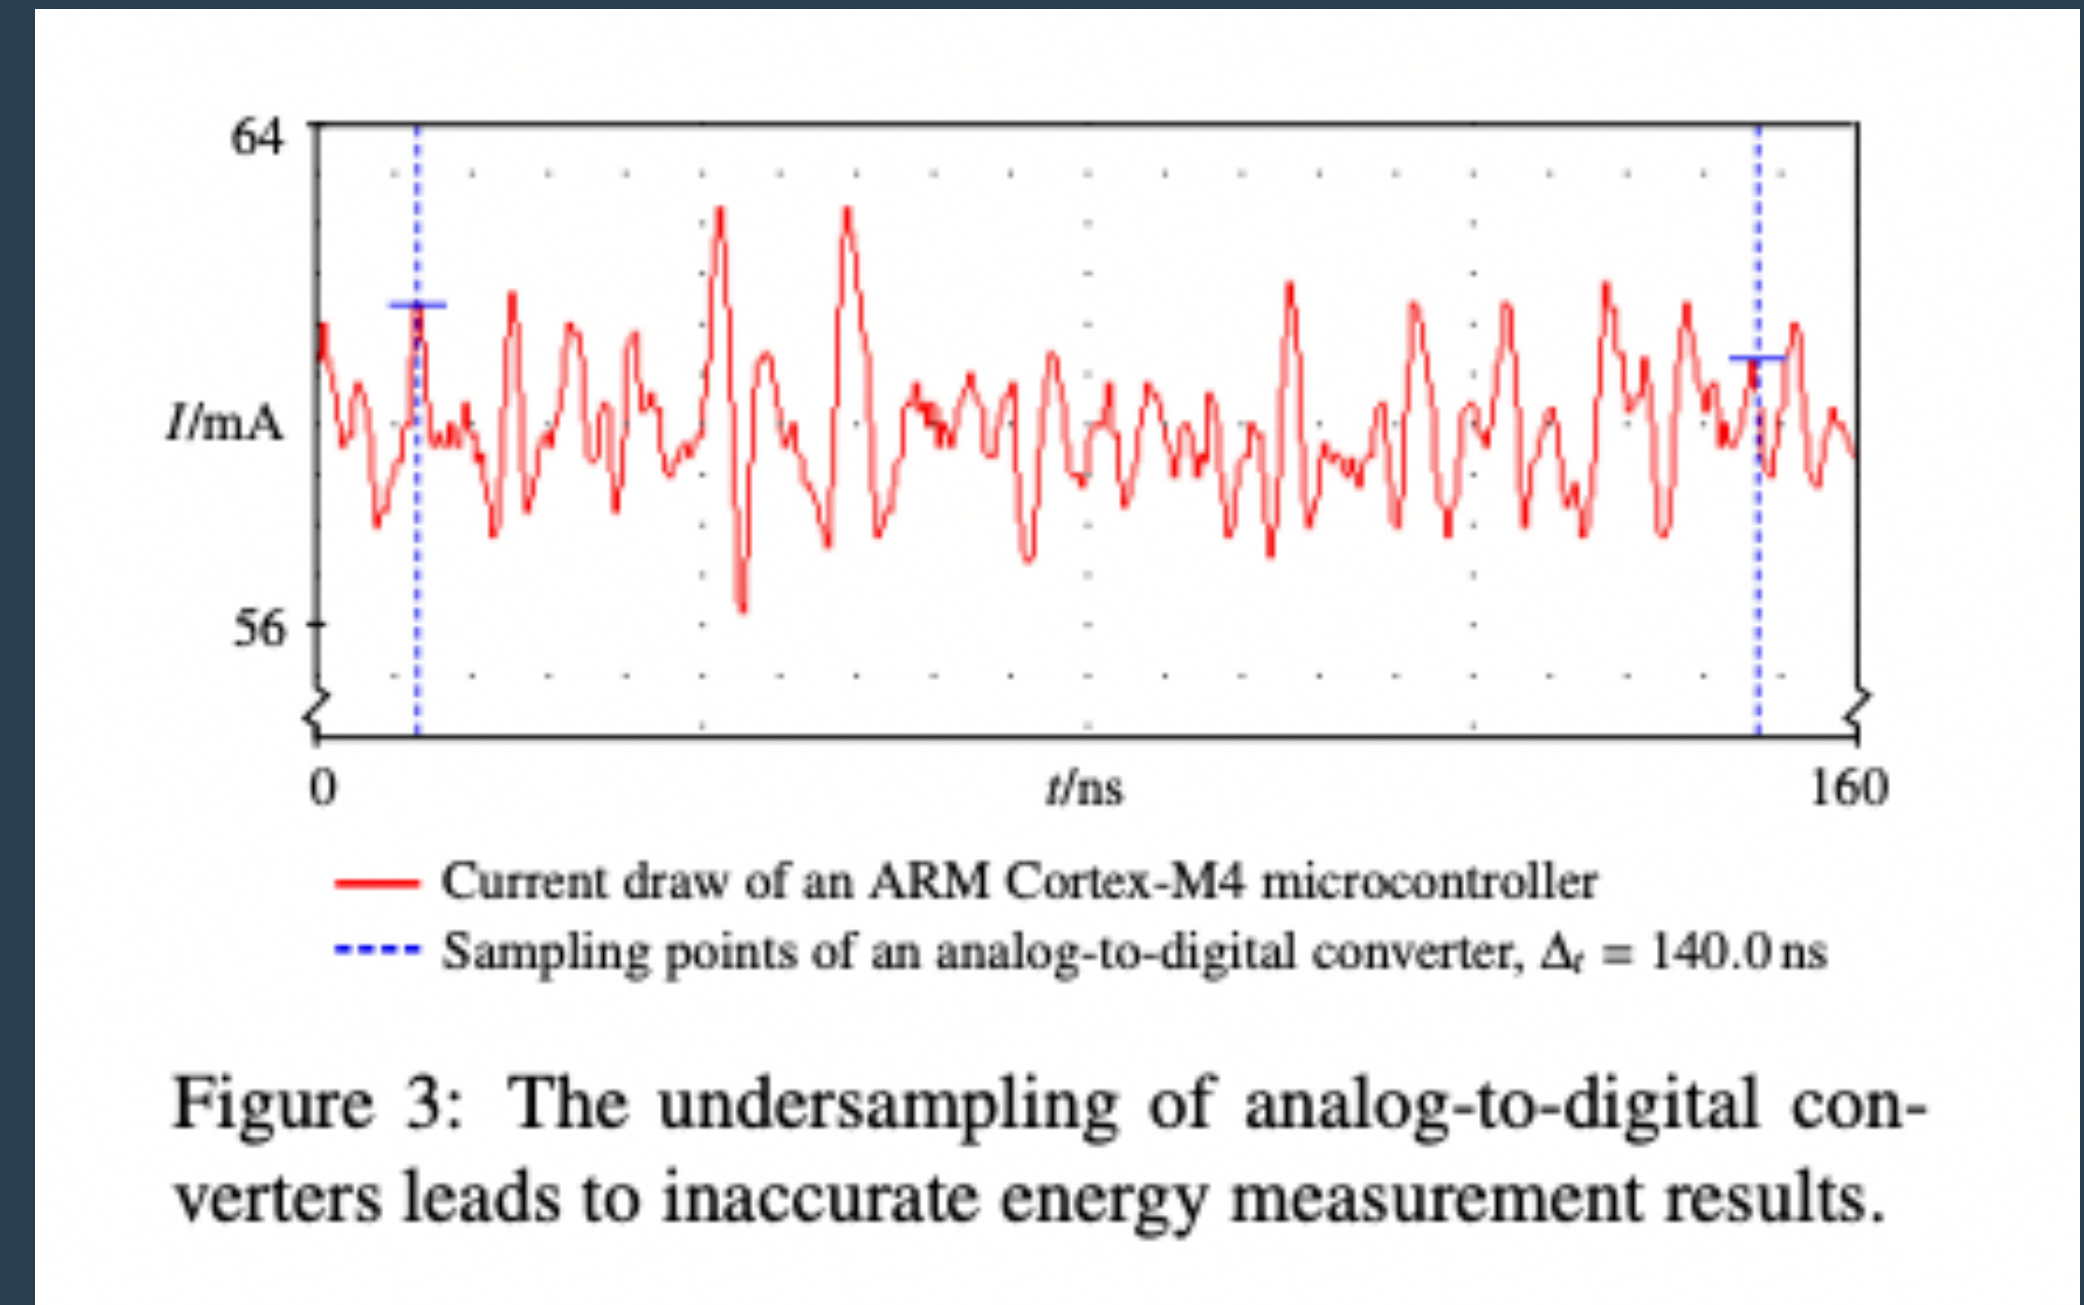
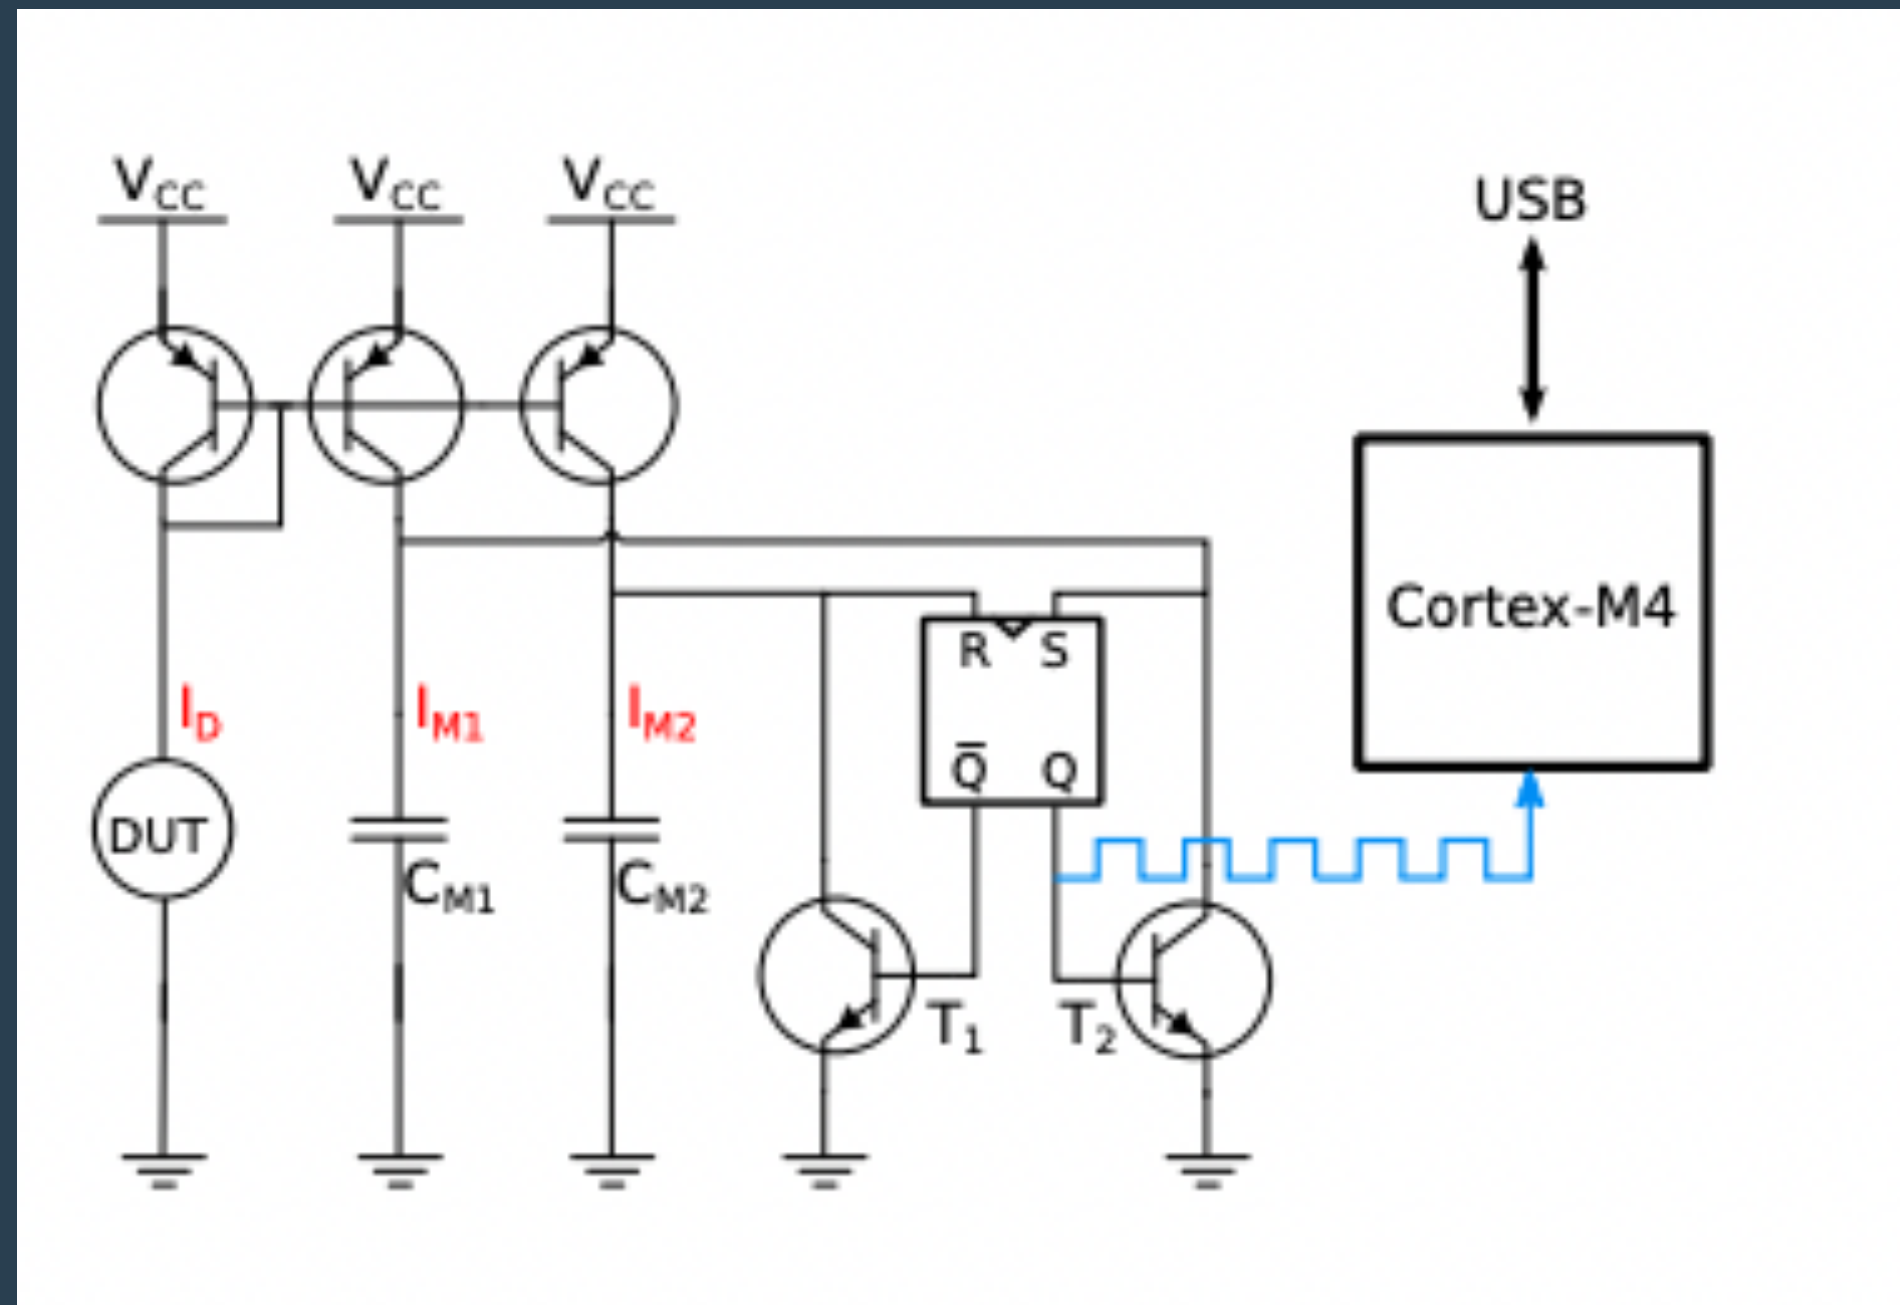
Auftrennen der ATX-Channels für DC Eingang ins Mainboard [Green Coding Berlin]



Direktanschluss an Versorgungsleitung im Mainboard
[https://web.eece.maine.edu/~vweaver/projects/rapl/2016_memsys_rapl.pdf]

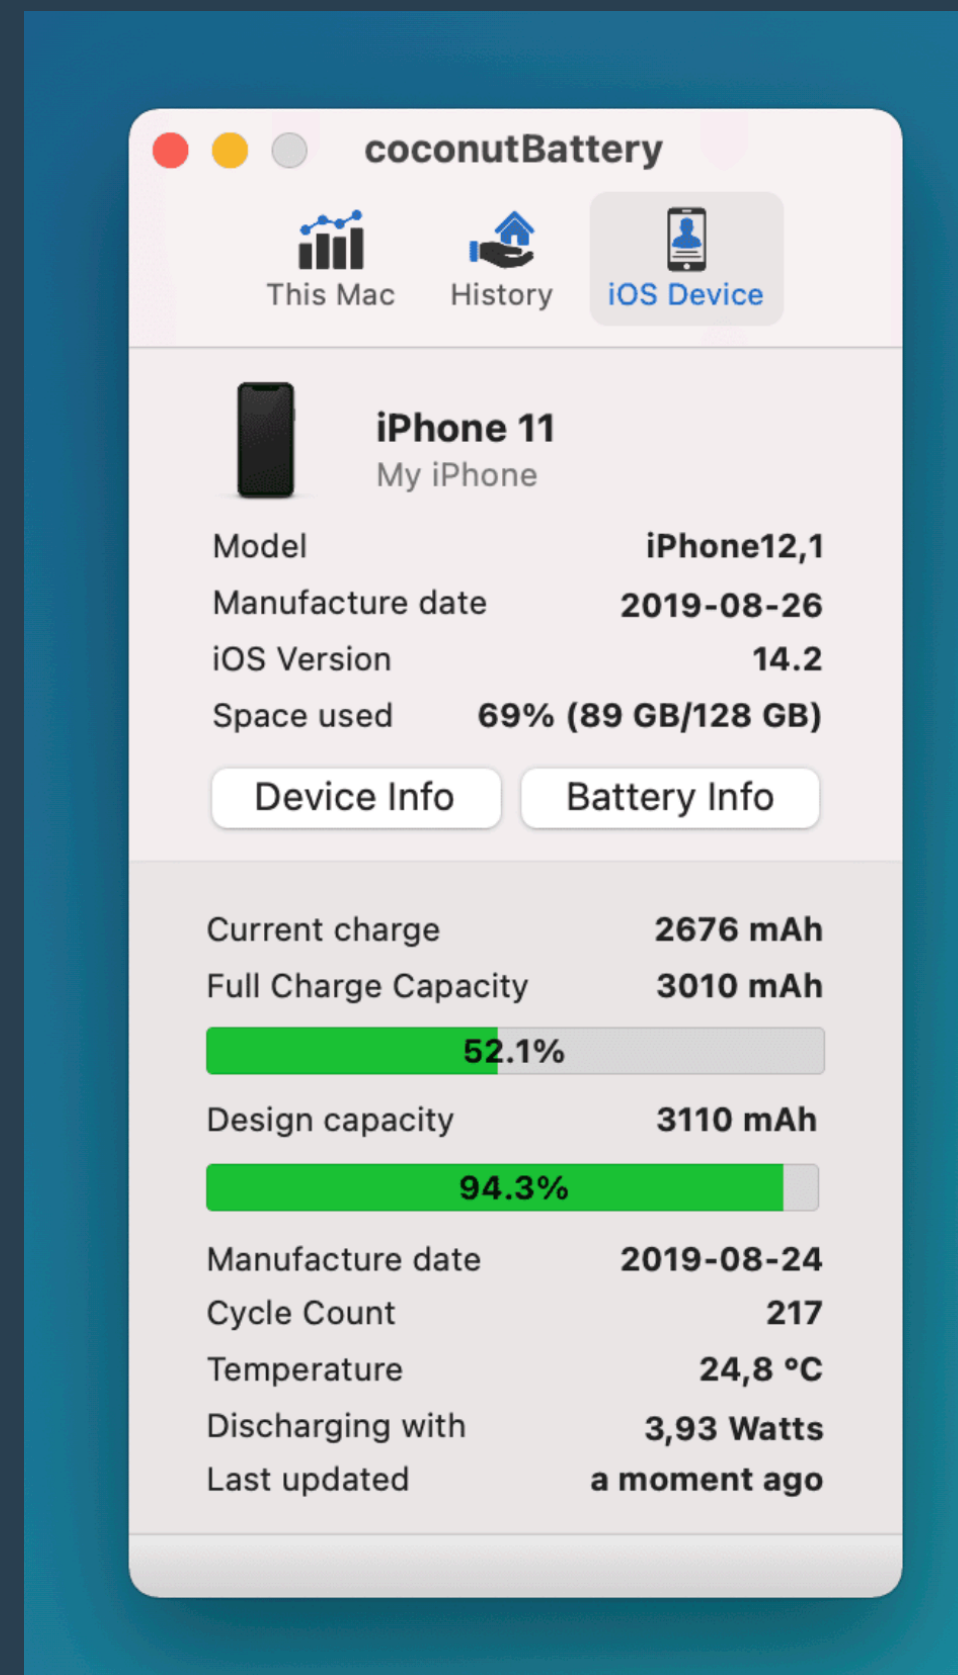
DC Energie-Messungen

Was wäre der bestmögliche Ansatz?

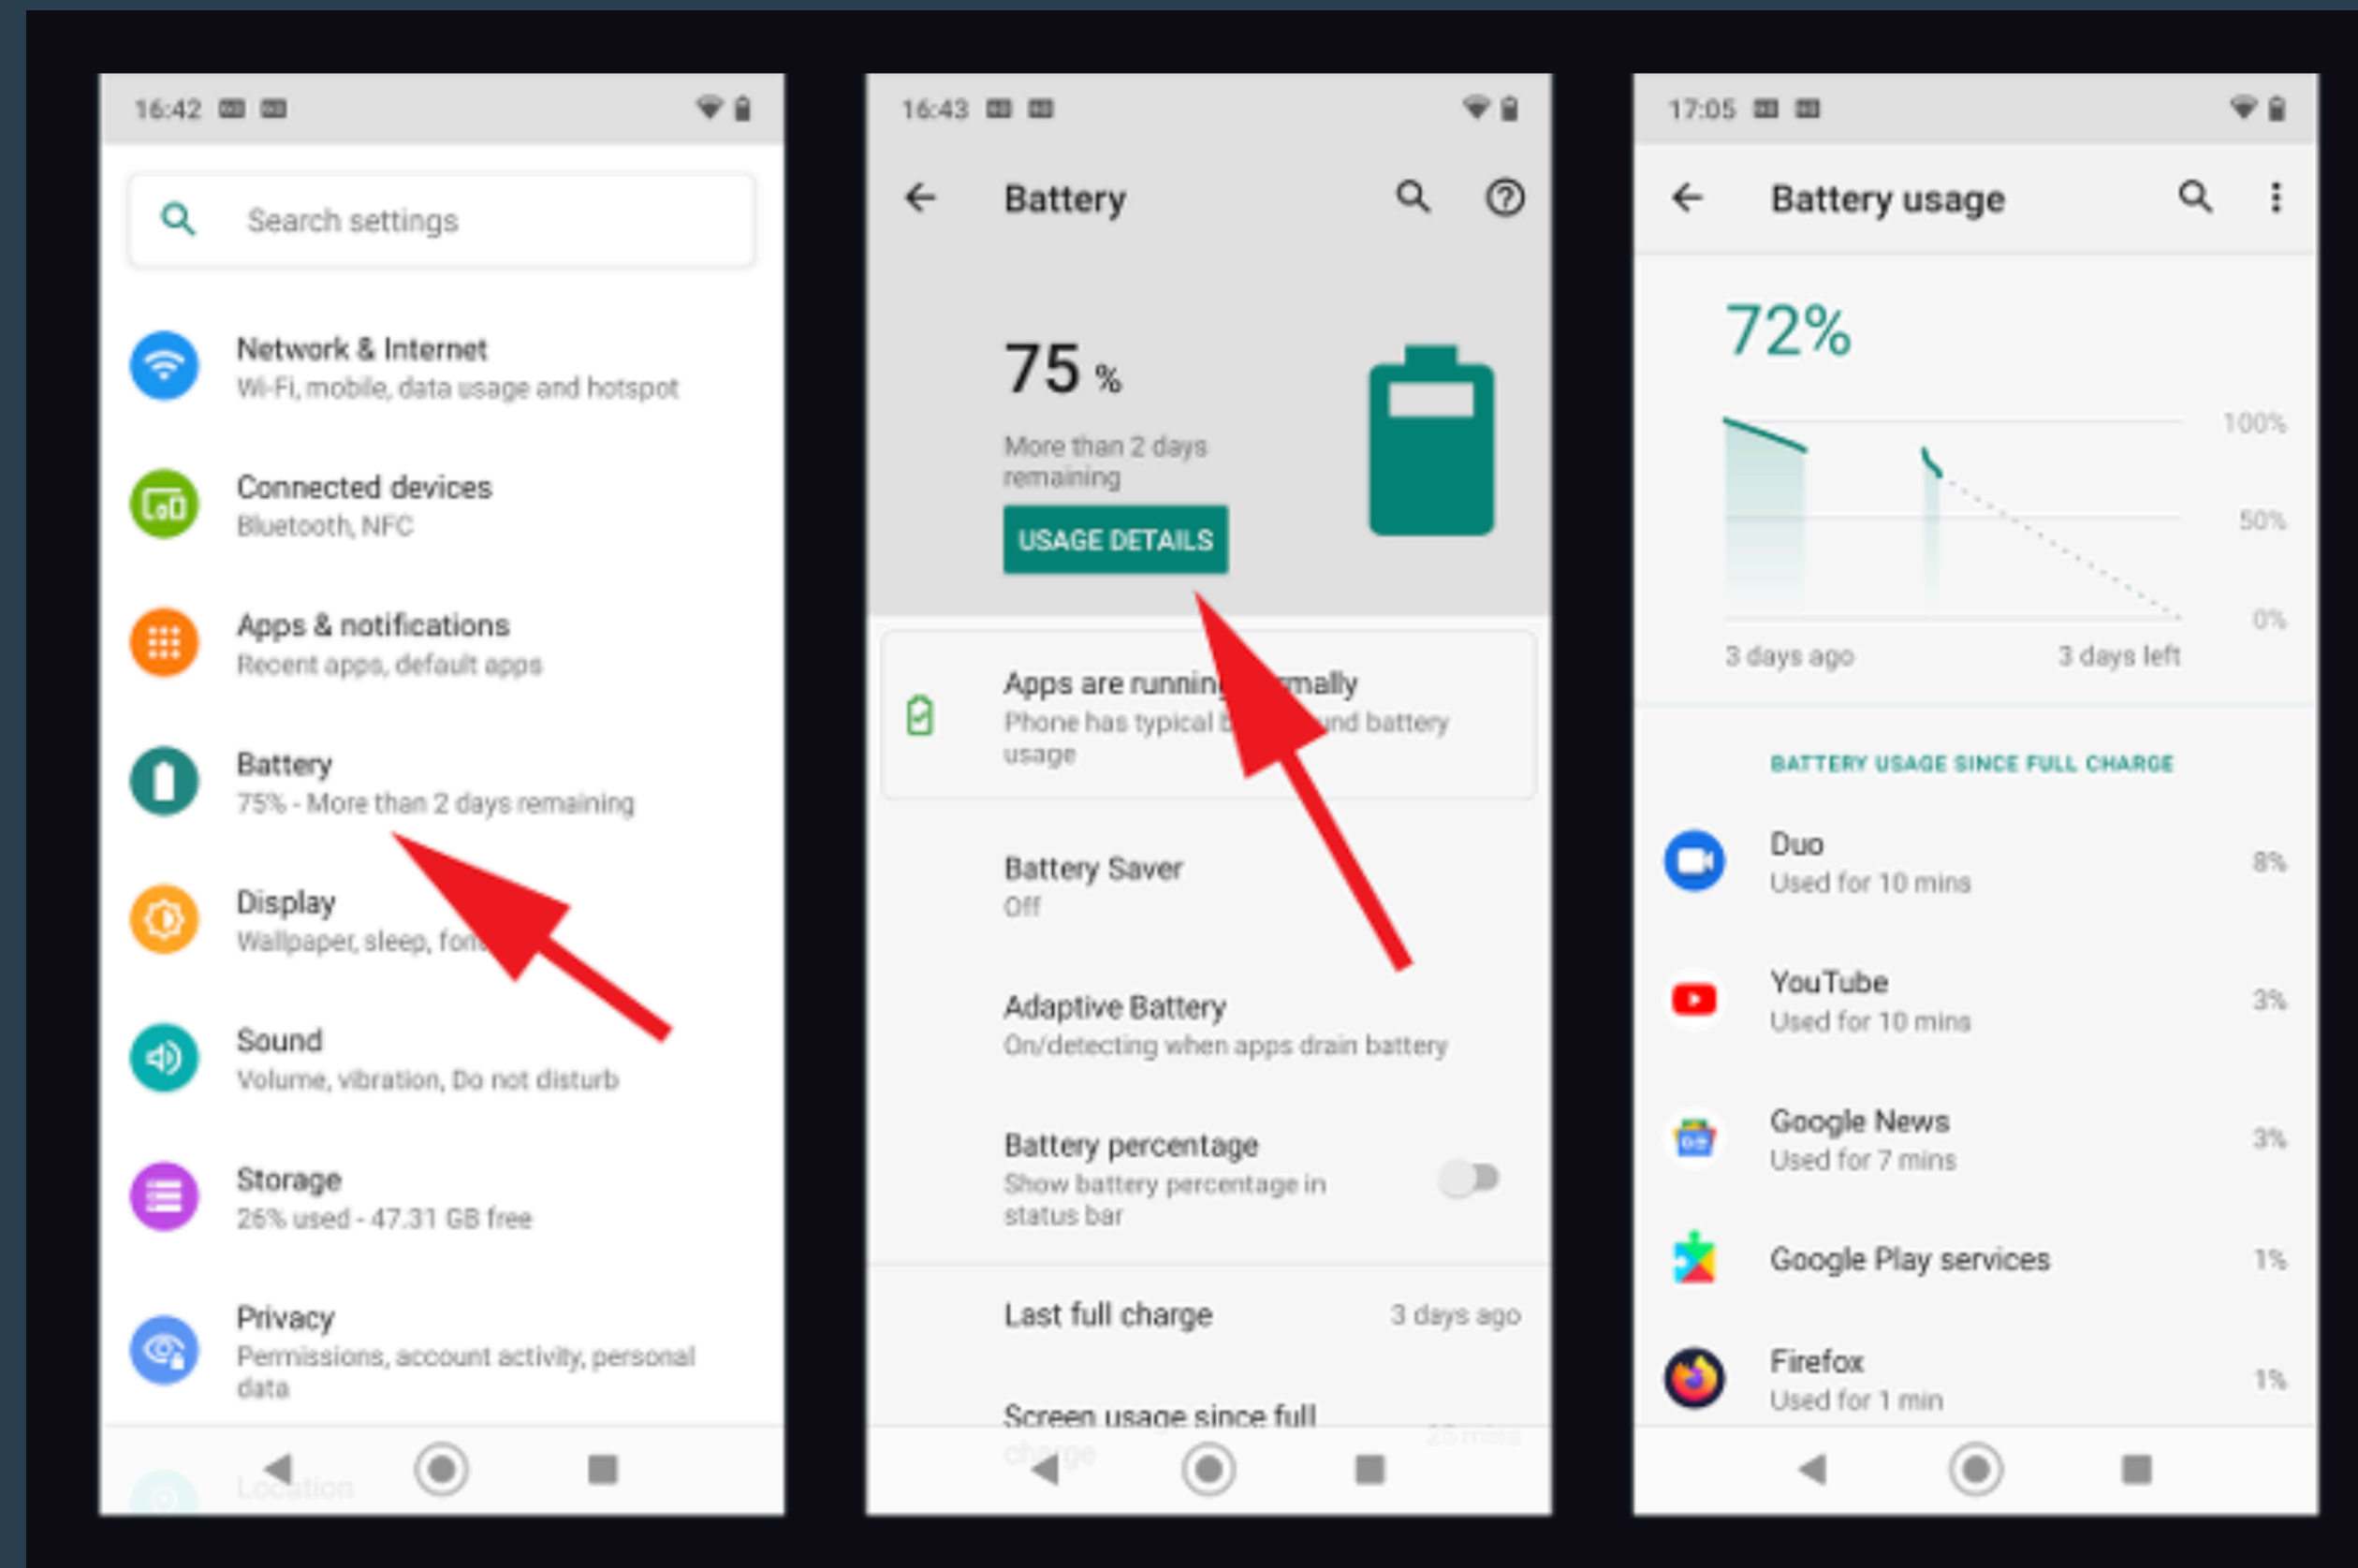


Wie messen wir Energie

Eine mögliche Alternative: Durch Batterieentladung bei mobilen Geräten



Example: Coconut Battery for macOS / iOS

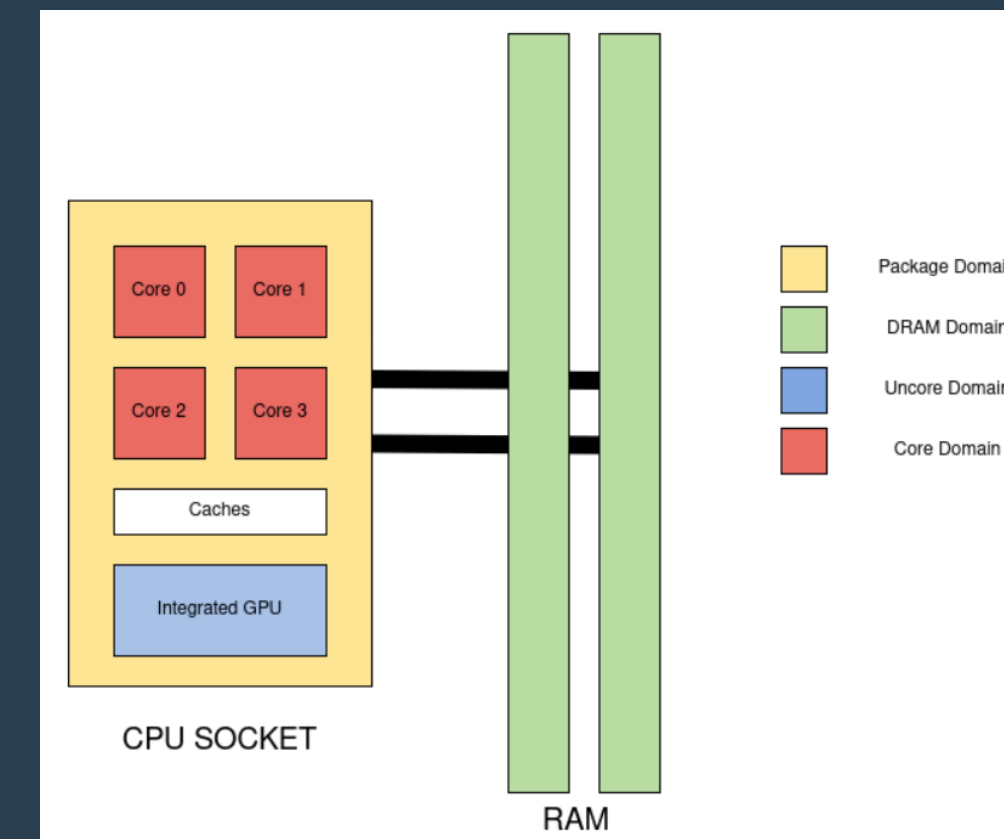


Android Battery usage (model)

Green Metrics Tool

Mess-Hardware / Cluster

- **Desktop:** Blue Angel Compatible Machine
- **Server:**
 - Fujitsu TX1330 M3 / M2 (2x)
 - Quanta Leopard - Multi-Tenant Server
SoftAWERE compatible
- **Developer:** MacBook Pro 2015 13"



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



MCP39F511N / AMD00706



NVIDIA SMI



Cloud

How can we estimate in the cloud?

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



Estimation of IT CO2 emissions

Static approach: Back of the envelope - to get first snapshot

- Assumptions:
 - \$2.5 - 5\$ per hour for full machine (bare metal) [1]
 - Typical machine draws ~ 600 W [2]
 - Worldwide grid intensity is 442 gCO2e/kWh [3]
- ~6,500 - 13,000 EUR in cloud spends equals 1 ton of CO2 emissions
- Downside: Only gives you the company total operation

[1] <https://aws.amazon.com/ec2/pricing/on-demand/>


[2] <https://datavizta.boavizta.org/cloudimpact>

[3] <https://www.green-coding.io/co2-formulas/>

Potentials of different hardware configuration


Power behaviour of hardware and how to leverage non-linearity

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)