Digital Emissions Methodologies

Our approach to measuring carbon emissions from the sites we host and support

Authors

Thibaud Colas

Tom Usher

Version

1.0, 11 Sept 2023

Torchbox

Contents

Introduction

We use different methodologies to calculate digital emissions within different system boundaries:

- Sustainable Web Design (SWD)
- Hosting infrastructure
- GreenFrame
- Display power consumption

Sustainable Web Design (SWD)

This is the model used by a lot of online calculators, such as Website Carbon. This methodology has the widest system boundaries, including usage and embodied emissions from devices of website users. It's based on the well-established Sustainable Web Design model which uses page weight as a metric of energy efficiency, and page views as a metric of site utilisation. This model has clear known limitations, but is nonetheless ideal to provide high-level figures for a wide range of websites or pages.

Hosting infrastructure

This methodology focuses on emissions generated by the physical servers used to host websites, calculated based on our understanding of the computing resources involved.

It provides more precise numbers than Sustainable Web Design, but only accounts for infrastructure for which there is data available.

Greenframe

GreenFrame is what we use to determine energy usage for a specific utilisation scenario on a project. It provides very precise and repeatable measurements for those simulated scenarios, allowing before-after comparisons. It does not provide evaluations of embodied emissions.

Display power consumption

For the specific task of optimising the energy usage of a website's designs on OLED monitors – we measure energy usage with a watt-meter.

Choosing the appropriate methodology

Here are the methods we recommend for different purposes:

Carbon accounting & carbon footprint reporting

For inventory of emissions according to the GHG Protocol Scope 3 standard, we recommend either the Sustainable Web Design or Hosting infrastructure methodologies.

- Sustainable Web Design (SWD) uses the largest system boundary, in particular including product/website usage emissions generated by and embodied in user devices (home router, laptop, smartphone).
- Hosting infrastructure focuses on computing resources necessary to make the product/website available on the web.

Both approaches can be valid depending on the organisation's reporting goals. For compliance with UK SECR, EU CSRD, or other regulations, we recommend further research based on the specifics of the products/websites to be reported on.

Carbon reduction initiatives

To quantify the impact of carbon reduction initiatives, here are our recommendations in different scenarios.

Initiative	SWD	Hosting infrastructure	GreenFrame	Display power cons.
Infrastructure reduction or relocation	Avoid	**Recommended**	Inapplicable	Inapplicable
Back-end / database / server code optimisations	Avoid	Consider case-by-case	**Recommended**	Inapplicable
Front-end code optimisations	**Recommended**	Inapplicable	Consider case-by-case	Inapplicable
Dark theme	Inapplicable	Inapplicable	Inapplicable	**Recommended**
Design / product changes	**Recommended**	Inapplicable	Consider case-by-case	Inapplicable
Carbon reduction KPIs	Consider case-by-case	Consider case-by-case	Consider case-by-case	Consider case-by-case

Sustainable Web Design Methodology

Here are the system boundaries as defined by this methodology:

Datacentre usage

15% of total figure

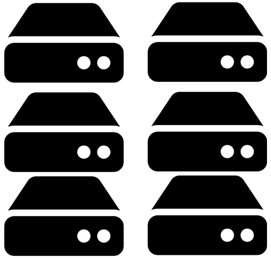

Carbon emissions from electricity to power servers, and keep them cool.

Network transfer

14% of total figure

Carbon emissions from electricity to power core networks, mobile networks, on-premise wifi and wired routers.

End user device usage

52% of total figure

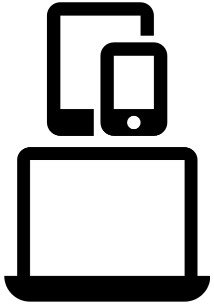

Carbon emissions from electricity to power end user devices - laptops, smartphones, tablets etc.

Production

19% of total figure

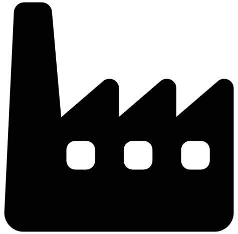

Carbon emissions from making the devices

Based on an image from the Green Web Foundation, CO2.js

GreenFrame is what we use to determine energy usage for a specific utilisation scenario on a project. It provides very precise and repeatable measurements for those simulated scenarios, allowing before-after comparisons. It does not provide evaluations of embodied emissions.

Display power consumption

For the specific task of optimising the energy usage of a website's designs on OLED monitors – we measure energy usage with a watt-meter.

based on the well-established Sustainable Web Design model which uses page weight as a metric of energy efficiency, and page views as a metric of site utilisation. This model has clear known limitations, but is nonetheless ideal to provide high-level figures for a wide range of websites or pages.

Hosting infrastructure

This methodology focuses on emissions generated by the physical servers used to host websites, calculated based on our understanding of the computing resources involved.

It provides more precise numbers than Sustainable Web Design, but only accounts for infrastructure for which there is data available.

Estimation approach

Refer to SWD's Calculating Digital Emissions for exact formulas and important predetermined factors.

To measure page weight, Torchbox relies on:

- HTTP Archive data where available, for accounting purposes over large numbers of websites.
- DebugBear data where available, for more precise accounting over large numbers of web pages.
- Google Chrome HAR file exports or Lighthouse measurements for specific pages / versions of pages.

To measure page views, we rely on:

- Google Analytics data exports where available, for precise page views per page URL.
- Chrome UX Report site popularity data where available, for accounting purposes over large numbers of websites.

To determine region-specific carbon factors, we rely on:

- Google Analytics data exports where available, for precise allocations of traffic by country.
- Websites ccTLD to country mappings, for accounting purposes over large numbers of websites.

Website-specific estimates

With the above data provided, we can provide estimates for any website, with increased precision based on the data source.

Without any data, here are indicative estimates for Torchbox websites, based on high-level calculations over a sample of 44 Torchbox-built websites:

- 1.5 MB average site page weight
- 6M average site page views per year
- 319 gCO2e/kWH median grid carbon intensity
- 0.32 gCO2e/pageview average emissions per page view
- 1.5 tCO2e/year average total site emissions per year

Hosting infrastructure methodology

Our reference client hosting configuration involves components from three primary providers:

- Heroku, a Platform-as-a-Service (PaaS) providing hosting for the applications we build
- AWS S3, a file storage solution for storing big files like documents and images.
- Cloudflare, a Content Delivery Network (CDN) which copies parts of our sites to locations closer to users so that they can be delivered faster.

While this calculation would be a lot less involved if these providers would share emissions numbers with their users, they do not, so we have to find ways to make the most educated estimates we can.

Thankfully, we can build on the work of tools such as the open-source [Cloud Carbon Footprint] (CCF) from Thoughtworks. This tool provides formulae and values to calculate emissions from a number of cloud hosting platforms, which we can use and adapt as required.

To establish a full picture of the emissions from hosting, we follow the same approach as CCF, calculating total emissions as:

Operational emissions + embodied emissions where

- Operational emissions are those from the active usage of power to run servers.
- Embodied emissions are those generated from the manufacture and provision of those servers by the cloud provider.

Heroku

Heroku's platform hosts sites on what they refer to as 'dynos' - their way to share resources between lots of applications running on the same physical server. While we don't have a way to determine exactly what that underlying server looks like, or what percentage of it any of our sites occupy, we:

- Know that Heroku runs on AWS
- Have emissions calculations for AWS regions from CCF
- Can define an equivalent AWS 'virtual server' based on the resources Heroku gives us.

For example, a client site running a 'Standard-1X' dyno with 512MB of RAM and 1x CPU Share (ref. 3) might be the equivalent of an AWS EC2 't2.nano', which also offers 512MB of RAM and 1 CPU.

Given that equivalence, we can use values provided by CCF as a baseline for our calculations.

First, we break down the server in to components for which we can calculate separate emissions values:

- Compute (CPUs)
- Memory
- Storage

Then, based on the Cloud Jewel's methodology established by Etsy we determine coefficients for each of those components by which we can multiply our usage to determine a kWh figure.

Each of these coefficients are calculated following figures provided by CCF.

Calculating coefficients

Compute

Compute's coefficient is calculated based on the average wattage of the appropriate CPU microarchitecture.

Following the CCF methodology, we can calculate the average wattage as:

*Minimum Watts + (Average CPU Utilisation (%) / 100) * (Maximum Watts - Minimum Watts)*

In the case of our Heroku equivalent, we consider the microarchitecture to be of Intel's 'Cascade Lake' generation, which following CCFs figures, has:

- A minimum wattage of 0.64
- A maximum wattage of 3.97

Considering an estimated average CPU utilisation of 50%, our CPU coefficient is calculated as:

*0.64 + (50 / 100) * (3.97 - 0.64) = 2.305 Watt Hours / vCPU Hour*

Memory

CCCF provides a coefficient of 0.392 Watt Hours / Gigabyte Hour for memory, based on an average of estimations provided by memory manufacturers.

The above compute wattages include an energy estimation for memory, so CCF only applies an additional calculation if the bare metal server in use is estimated to have more memory than those accounted for by the compute estimations.

CCF calculates an average of 98.12 GB per physical CPU chip.

The T2 instances have a maximum (as close to bare metal as we know) memory of 288 GB and 1 CPU chip.

This leaves us with an excess of 189.88GB of memory which is not accounted for by the compute figures. Using CCFs coefficient, the total Watt Hours for a bare metal T2 server would be:

189.88 x 0.392 = 74.432 Watt Hours.

To determine a coefficient for 1GB, we can calculate a Watt Hours / Gigabyte figure as

74.432 / 189.88 = 0.3920 Watt Hours / Gigabyte Hour

Storage

CCFs storage coefficient calculations are based on average capacity and average wattage of HDDs (spinning platter drives) and SSDs (solid state drives) projected from a 2016 U.S. Data Center Usage Report.

These calculations result in coefficients of:

- 0.65 Watt Hours / Terabyte Hour for HDDs
- 1.2 Watt Hours / Terabyte Hour for SSDs

Carbon estimates

Given these coefficients, we can continue with CCFs methodology to calculate CO2e estimates for these components.

To do this, we first multiply each value by a Power Usage Effectiveness (PUE) score; a value which indicates a data center's energy efficiency (where a value of 1 indicates 100% of energy goes towards powering servers). CCFs estimated PUE for AWS is 1.135.

This gives us the following values for each component:

- Compute: 2.616 Watt Hours / vCPU Hour
- Memory: 0.4449 Watt Hours / Gigabyte Hour
- HDD: 0.7378 Watt Hours / Terabyte Hour
- SSD: 1.361 Watt Hours / Terabyte Hour

The next step is to multiply these values by an emissions factor for the region in which our servers run. For most of our applications, this is AWS' 'eu-west-1' region for which CCF has listed an emissions factor of 0.0002786 metric tons/kWh (from data provided by the European Environment Agency).

This leaves us with the following figures (converting the above to kWh):

- Compute: (2.616 / 1000) x 0.0002786 = 0.0000007288 tCO2e/vCPUh
- Memory: (0.4449 / 1000) x 0.0002786 = 0.0000001239 tCO2e/GBh
- HDD: (0.7378 / 1000) x 0.0002786 = 0.0000002056 tCO2e/TBh
- SSD: (1.3619 / 1000) x 0.0002786 = 0.0000003794 tCO2e/TBh

Embodied emissions

To estimate embodied emissions, we first need to determine the total emissions generated during the creation of the server hardware throughout its lifetime. Then, we divide this total by the proportion of the server hardware utilised by our application.

To do this, we use CCFs interpretation of the Green Software Foundation's [Software Carbon Intensity methodology](), which calculates embodied emissions as:

Embodied Emissions = Total Emissions x (Time Reserved / Expected Lifespan) x (Resources Reserved / Total Resource)

For our calculations:

- *Total Emissions* are from [CCFs collected data](), in the case of our reference t2.nano, this value is 1.4772 tCO2eq.
- *Time Reserved* will be the entire time period we are calculating over, as our sites operate 24/7. For consistency with our other figures, this will be 1 hour.
- *Expected Lifespan* is fixed at 4 years, or 4 x 8760 hours, 35,040 hours.
- *Resources Reserved* is the number of vCPUs used by our t2.nano; 1 vCPU
- *Total Resources* is the total number of vCPUs on the largest 't2' server (as close to bare metal as we can calculate), 8 vCPUs in the case of a t2.2xlarge.

This leaves us with a calculation of:

1.4772 x (1 / 35040) x (1 / 8) = 0.000005270 tCO2eq/vCPUh

Application hosting totals

Given these values we can now determine a total emissions estimate for our applications over time.

Our hosted applications are typically made up of four key resources:

- The ‘dyno’ hosting the application
- A virtual Heroku-managed server hosting the database service
- A virtual Heroku-managed server hosting the Redis cache service
- Storage of large files (images and documents) in AWS’ S3 object store

Given Heroku’s CPU, memory and storage figures for each of these, we know that:

- A ‘standard-1x’ application dyno has 1 vCPU, 0.5 GB memory and roughly 512GB (0.5TB) of SSD storage.
- The ‘standard-0’ database addon has 2 vCPU, 4 GB of memory and 64GB (0.064TB) of SSD storage.
- The ‘premium-0’ tier of the Redis cache addon has 1 vCPU, 0.5GB of memory and 50MB (0.005TB) of SSD storage.

For the object store we consider compute and memory to be negligible and take a pessimistic estimate of 1TB of storage per site. We also assume the worst case in that all S3 data is stored on the SSDs.

Therefore, see the table on the following page for the detailed per hour breakdown. Over a longer period of time the totals are:

- **0.018404 tCO2eq/~month**
- **0.220844 tCO2eq/year**

Component	Resource	Scale	tCO2eq/h
Application Dyno	Compute	1 vCPU	0.000000728866355
	Memory	0.5 GB	0.000000061977356
	Storage	0.5 TB	0.0000001897266
	Embodied	1 vCPU	0.00000527
Database Addon	Compute	2 vCPU	0.00000145773271
	Memory	4 GB	0.000000495818848
	Storage	0.064 TB	0.0000000242850048
	Embodied	2 vCPU	0.00001054
Cache Addon	Compute	1 vCPU	0.000000728866355
	Memory	0.5 TB	0.000000061977356
	Storage	0.005TB	0.000000001897266
	Embodied	1 vCPU	0.00000527
Object Store	Storage	1 TB	0.0000003794532
Total			**0.00002521060105**

Supporting components

While Heroku and object storage are likely to be the most intensive parts of our hosting stack, there are a number of supporting components including but not limited to:

- Network traffic between Heroku and Cloudflare (our most commonly used CDN)
- Network traffic leaving Cloudflare
- Scheduled tasks in Heroku scheduler
- Processing and storage of backups
- Development processes such as Continuous Integration (CI) tasks.
- Supporting services such as log aggregation tools, monitoring, application error track, application performance monitoring and transactional email services.

While our figures don't currently take these into account, we intend to evolve this methodology to include as many of these components as possible over time.

Greenframe methodology

The GreenFrame methodology depends entirely on GreenFrame, a dedicated carbon measurement software. The GreenFrame model explains how energy usage is calculated.

To use this methodology on a project, there needs to be a test harness set up with pre-agreed representative usage scenarios, created via automated testing scripts.

Here is an example of configuring a demonstration Wagtail website to measure different scenarios with GreenFrame: bakerydemo-gold-benchmark.

Based on this model and our results for a representative website, we estimate the energy consumption of a single user session on a Wagtail website to be within a range of 0.2 to 2 Wh / session. This depends on the complexity of the website, specific scenario of the session, and total duration spent on the site. Based on the local electricity grid's carbon intensity at the user's location, this represents carbon-equivalent emissions between 0.01 and 2 gCO2e/session.

Display Power consumption methodology

This methodology is very specific to the energy usage of OLED monitors, which varies according to the luminance of the website's visual designs – with the OLED panel using no power when displaying pure black.

To use this methodology, we recommend connecting a watt-meter on an OLED TV's power socket (or any other OLED monitor):

- Measure power consumption when the TV panel is off (displaying pure black or software-level "off" switch for the panel only)
- Measure power consumption displaying a light variant of the designs.
- And a dark variant.

The three numbers can then be compared to understand what kind of improvements a dark theme (or color changes) could bring.