Creating a digital environmental footprint: a Life Cycle Assessment approach

Abstract

With the Digital Environmental Footprint, the first activity on the Roadmap of the Sustainable Digital Infrastructure Alliance (SDIA), we aim to bring together experts and make fruitful connections to develop an open methodology to assess the digital environmental footprint of digital products and services. With this methodology we aim to raise the bar and create transparency around the impact of software and its infrastructure on our environment. In this paper we document what we, together with our community, know today and chart a path towards a standardized methodology.

We have chosen to build on the Life Cycle Assessment (LCA) methodology, since it can give us great insights on the environmental impact of products and services. We believe that to measure is to know: information revealed through performing an LCA can give us what we need to report on our environmental impact, and to improve digital products and services, and therefore our businesses. By defining a clear taxonomy of the value chain, and an innovative approach of digital resource primitives as well converting those digital resources into environmental impact, we are a step closer to assessing the environmental footprint of digital products and services.

Table of Contents

About the Alliance

The Sustainable Digital Infrastructure Alliance (SDIA) is an independent alliance of stakeholders working across the digital sector. As an Alliance, we are committed to transparency, and the sustainable development of the digital sector, as exemplified in our 2030 Roadmap to Sustainable Digital Infrastructure. Together with our community we have sketched a vision for what a digital economy could look like that is sustainable, for people, environment and business alike.

1. A clear taxonomy and defining responsibilities

After attempting to break down and communicate about information technology (IT) we found out that if we look at every component of IT and its infrastructure, things quickly get complicated and problems become impossible to both identify and resolve. We had to come up with a solution for this, and created a simpler taxonomy to describe the complex value chain of digital and IT. This is a change of perspective that we think will drive us forward.

Our school of thought is simple: software applications have a supply chain. This supply chain is what we call digital infrastructure. Digital infrastructure generates digital resources used to power the application and provides connectivity, enabling users to connect to an application. The user can access an application using a device, e.g. a smartphone or computer with a browser.

While the application itself is not a physical product, it does require physical digital resources to run. These resources are provisioned by digital infrastructure. Digital infrastructure is the combination of IT hardware, networking and a data center facility. The main input in the supply chain is electricity. As digital infrastructure generates heat, we use the term energy to capture both. The main output are digital resources (Figure 1).

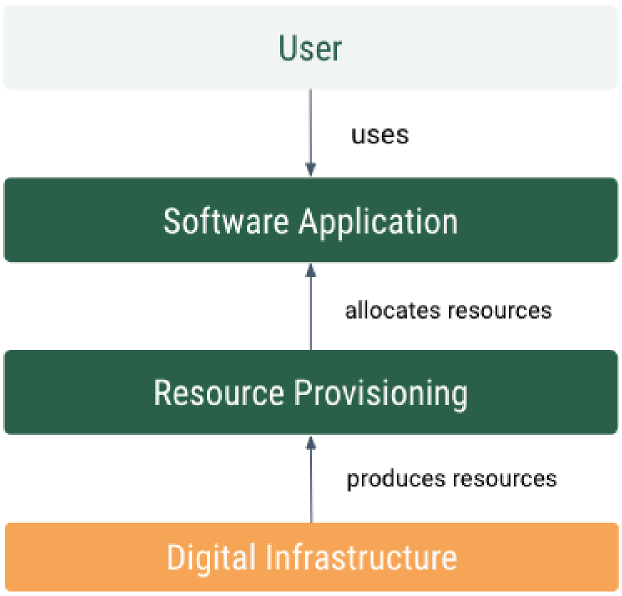

Figure 1: A simplified taxonomy

1.1. A framework for defining responsibility

Finding out the environmental footprint of a digital product is useful, but it creates new questions: which actor is creating the environmental impact within the value chain and who is responsible for reducing it?

This way of thinking also leads us to a simplified version of roles and responsibilities along the value chain. We learned that very few people are taking responsibility for the digital resource consumption of their products & services. Most of the attention is directed towards data centers, who represent only a small part of the value chain.

Our responsibility model, described in figure 2, shifts the majority of the responsibility to the creators and providers of digital product & services while giving the infrastructure actors a clear role to produce environmentally sustainable digital resources.

Role	Responsibility
User	• Sustainable use • Only when needed • Disable functionalities
Software Application	• Measure resource use & be transparent • Minimize resource usage • Use sustainably-made resources to operate
Resource Provisioning	• Reduce idling & unused resources (drive resource efficiency) • Enable access & verification of sustainable resources
Digital Infrastructure	• Generate sustainable resources • Make sufficient resources available • Label resources to differentiate

Figure 2: Roles and responsibilities across the value chain

1.2. The Missing Pieces: Measurement and Transparency

When considering these responsibilities, it becomes clear that the software application is responsible for the use of digital resources, hence it is important it can measure the resource usage and tie that resource consumption to environmental impact. This requires an assessment methodology that is tailored to the digital realm, which is what our methodology seeks to enable. **The idea behind this is that software does not exist in the physical world, but the resources it consumes do have a physical environmental impact.** Therefore, we can only measure how many digital resources the software is using and then convert those digital resources into environmental impact.

In conjunction with transparency and a label for purchasers or consumers of digital products and services, this enables more sustainable decisions, giving more control over picking the more environmentally sustainable product. Further, the developers, architects and software vendors in general gain insights on where the hotspots of environmental impact in their product are.

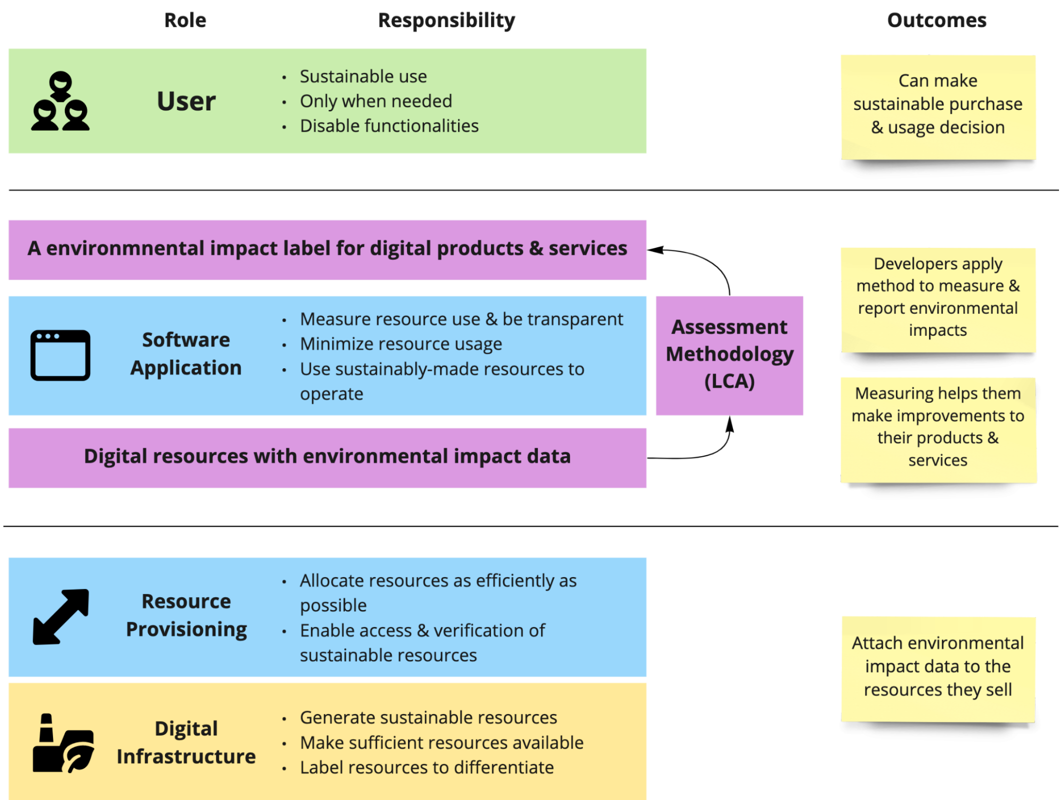

Figure 3: Roles and responsibilities across the value chain and the additional outcomes of adding an LCA methodology to assess the environmental impact of digital resources

Each actor along the value chain also has a responsibility to be transparent about their impacts on the environment. We want companies in the software industry to be enabled to do these LCA's themselves, and feel empowered to share this with their customers, improve their product, and create a more sustainable future together.

1.3. Defining digital resources as the key input to digital products

Digital Resource Primitives are defined as the low-level resources required for digital products & services to operate. They can be seen as the fuel that powers software applications. This is the key idea that we call Digital Resource Primitives which you can see in figure 4.

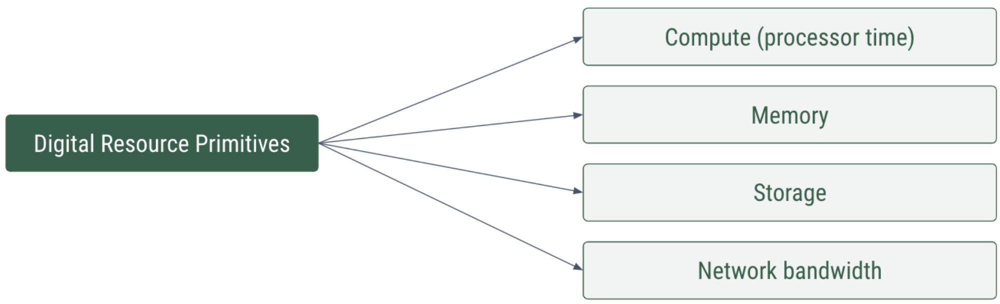

Figure 4: Digital Resource Primitives

Dividing this up in digital resources allows us to see the use of a specific resource usage across all layers of the value chain. Digital resources originate inside data centers and can be viewed as the 'output resource' of server, storage & network equipment. Therefore, we can assign an environmental impact per resource, by attributing the environmental impact of the chain of physical equipment (server, cables, rack, data center equipment, building) to the digital resource output they produce.

An example:

A generic server might have a digital resource production capacity of 48 vCPUs, 100 Gbit of network bandwidth, 128 GB of memory and 2 TB of storage. In order for this server to operate and be able to produce these resources, the value chain must be attributed to the server:

- Cabling
- Rack enclosure, power supply and distribution units
- Network switches & backbone network gear (for internet connectivity)
- Backup power systems (UPS, and backup generators)
- Cooling systems (to transport the generated heat away)
- Building & whitespace materials and components

- Energy used

Each of these has an associated environmental impact from manufacturing and use which can be allocated to the server, and through the server to each digital resource that it generates.

1.4. Attributing digital resources to an application & waste

At this point, we know the environmental impact of all digital resources that are produced by a specific server in a specific data center. On the other hand, we a digital application or process, or even virtual machine, for which we can measure the actual digital resource consumption (e.g. CPU time, memory usage, network bandwidth consumption, storage usage).

However, we can not directly convert all digital resources into equal environmental impact, as each type of digital resource represents a different fraction of the total environmental impact of a server. An example is the power consumption of the CPU that is much larger than any other component.

Hence we need to define an allocation weighting for each type of digital resource, which we have outlined in figure 5.

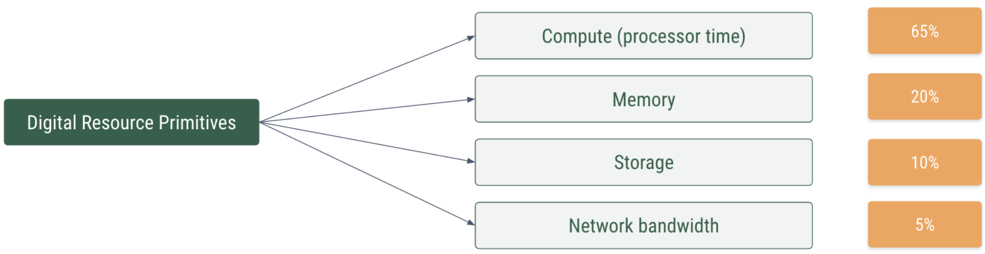

Figure 5: Digital resource impact weighting

The weighting of digital resource within servers is critical, and the following example will illustrate the challenge of doing so. Assuming there is a physical server that is virtualized and it has 3 virtual machines. Further resource allocation is the following:

Physical server: 128 GB Memory, 2 CPU Sockets (12 cores each), 24 cores total, 2 x 3.6 TB NVMe disks, 2 x 10 GbE Ethernet, total power consumption: 0.6 kWh

- Virtual machine 1: 12% memory, 24% CPU, 3% storage, 2% network capacity
- Virtual machine 2: 40% memory, 38% CPU, 12% storage, 18% network capacity
- Virtual machine 3: 33% memory, 3% CPU, 24% storage, 8% network capacity
- Idling resources (remainder of non-utilized digital resources)

- Overhead from virtualization/containerization itself
- Operating system overhead

As you can see, each virtual machine (but the same is true for an application, a process) consumes different amounts of the digital resource primitives, of which each creates an different environmental impact. With the proposed weighting ratio's, it becomes possible to allocate the total environmental impact of the underlying infrastructure to each type of primitive and in turn to attribute the environmental impact, using digital resource usage as a proxy to the application or virtual machine.

Existing research on attribution of energy to digital resources

There is limited existing research on this weighting, however, there are some studies that we can use to verify our assumptions.

Previous research by Hayri et al has found that *"the power used on servers is increasing and the two largest consumers of power are the processor and the memory"* (Hayri, 2016). In addition, research works have attempted to "*optimize systems to reduce DRAM power consumption*" (Hayri, 2016).

At this point in time, it seems that there is no reliable research on how to attribute 1 watt of input power into a server to its various components. In the research found, the CPU seems to still make up the majority of the power consumption as well as variance. However, most of the research found merely focussed on memory and CPU rather than including storage (SSD, HDD, NVMe) and GPUs in their comparisons.

Unfortunately, the database of Spec.org & the SERT results are not public which may have helped defined a more accurate weighting or could help create an API that allows the per-server calculation of these weightings.

Therefore, we settle on the following ratios:

Resource Utilization	**Allocation %**
CPU	65%
Memory	20%
Storage	10%
Network	5%
GPU	Exclude

Table 1: Assumed ratios for each resource.

Second Iteration

The ratio of allocating resource usage to power consumption are not the same across different systems. In the first iteration, we defined our default model. This represents our fallback. In order to improve the accuracy of the formula, it should be possible to parameterize it. This means that if the CPU, type of memory, storage, and network is known and we have a specific, verified measurement on the allocation, we can use that instead of using the default assumption. This would require that a tool exists, which given a standard usage scenario, can measure the resource allocation distribution on a computer or server system. Second it would require that a public database exists from which those values can be pulled. This is similar to spec.org (see example here), however it does not cover CPU power consumption, but all components of the system and their respective share of the total power consumption of the system in various usage scenarios.

Further, if we measure the amount of unused digital resources per server (e.g. idling), then we are also able to convert those 'wasted resources' into environmental impact as well.

Digital resources can not be stored, so once they are generated, they must be consumed. Measuring this waste is critical to improve the efficiency of resource provisioning. Further, it can lead to the improvement of the overall utilization of digital resources produced and and reduce overbuilding/overprovisioning of digital infrastructure.

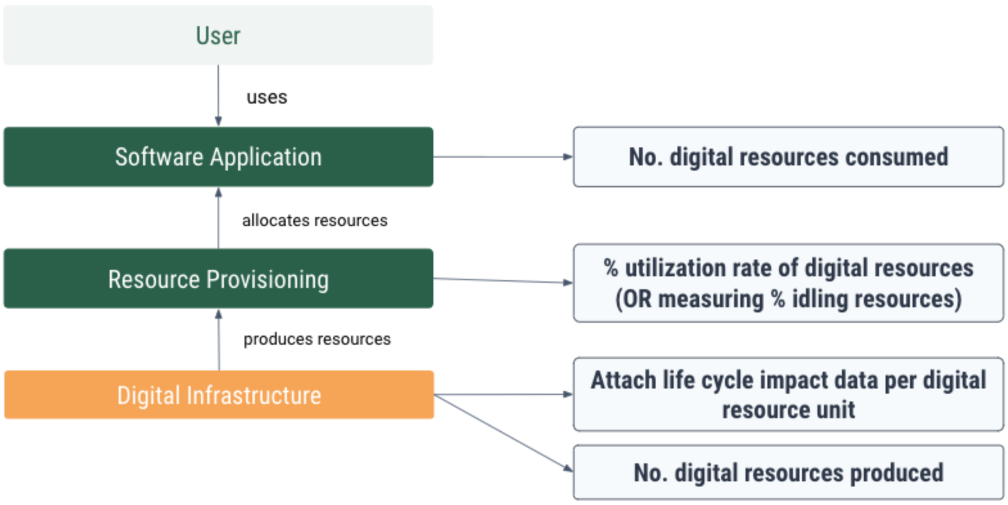

Figure 6: Model of how the environmental impact of a digital product's holistic measuring approach.

2. Methodology

2.1. Principles

To develop the methodology for a "digital environmental footprint", the SDIA and its community has defined clear principles for its development as an overarching, independent approach to quantifying the environmental impact.

- **Transparent**: Methodology, framework and calculations must be open and accessible to all stakeholders; digital products & services are enabled to create transparency on their own environmental footprint.
- **Trustworthy**: All information gathered and reported must be verifiable and scientifically-proven to ensure trustworthiness.
- **LCA Approach**: Utilize the existing LCA methodology and collaborate with academia to adapt existing frameworks to be applicable to digital products & services.
- **Open standard**: Create an open standard that is accessible to all stakeholders at no cost and is non-competitive.
- **End-to-end**: Make the entire environmental footprint across the supply chain and lifecycle of a digital product & service visible to all stakeholders.
- **Open-source**: The methodology, tools and frameworks should be part of the collective equity of society to support the reduction of the total environmental impact of the digital sector.
- **Impact focussed:** Focus on energy, resource & pollution reduction rather than merely efficiency, in order to combat the climate crisis through minimizing the environmental impact burden created by the digital economy.

The long-term vision for the environmental footprint is to capture all environmental impacts created across the value chain and attribute them to the server-side application, and eventually the user's usage of that application (Table 2).

2.2. Applying Life Cycle Assessment Methodology to digital products

LCA's are adopted broadly in the sustainability realm. For many sectors, ranging from the building sector to agriculture and energy sectors, widespread use of LCA's to identify environmental footprints of products can be observed. However, it has not been applied to the digital sector. This is one of the reasons why efforts to make the digital sector more sustainable are difficult: we don't have the right information and data on the environmental impact to improve the environmental performance of digital products. Applying methodologies like a Life Cycle Assessment can give us these insights on the environmental impact of digital products.

"LCA can support the pathway to a more environmentally sustainable digital transformation by pointing out environmental hotspots of ICT, by identifying environmental tradeoffs, by monitoring the environmental impacts of the digital transformation and by showing the sustainability potentials of substituting digital for non-digital technologies with high environmental impacts." [1]

LCA also allows for the comparison between different digital products and services, and can, for example, help us choose the most environmentally sustainable option. While it is a complex methodology, we believe it is also especially applicable to the digital sector because of this complexity. It is able to cover the many aspects of the value chain. Standards are already in place, which set the norms that define an LCA and its requirements, for example ISO-14040 and ISO-14044. Part of this is a review process, where the results are reviewed by an independent party.

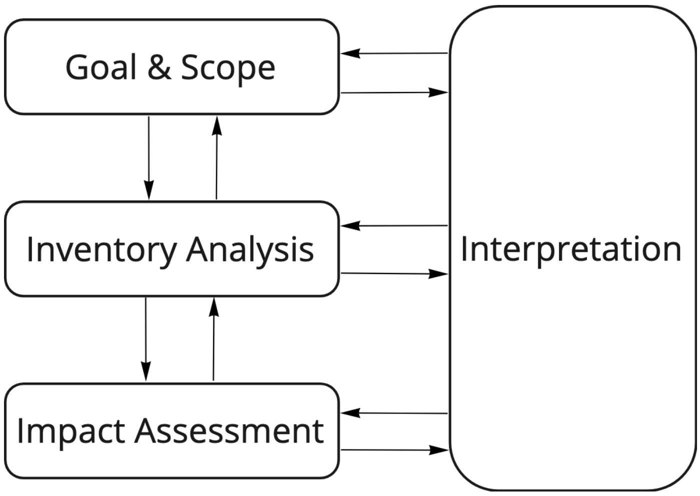

Figure 6: LCA Framework

We learned that software or a digital product do not actually exist physically, however they do consume physical resources. These resources do not fit the standardized LCA impact categories, because software does not consume energy or physical resources in the same way as a physical product, e.g. a paper notebook, does. Instead, it consumes computation (CPU cycles, GPU cycles, operations per second), memory, storage, and network bandwidth. Each of these does consume energy and has resources embodied from the manufacturing of the equipment needed to transform energy into computation, memory, storage and network bandwidth.

Therefore, we need to perform the transformation towards digital resources as described before. For each digital resource produced, the environmental impact can be measured and attributed to the resource. For each application, digital resource usage can be recorded, and environmental impacts can be allocated using digital resources as a proxy.

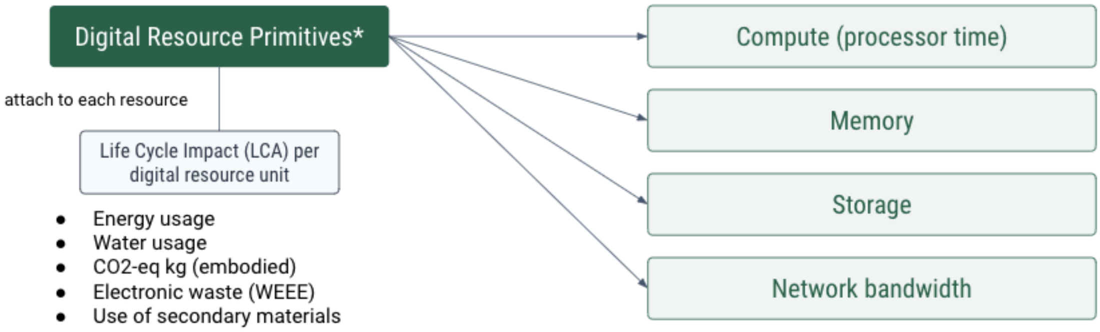

Figure 7: Digital Resource Primitives and the LCA impact categories that can be calculated per unit of digital resource.

2.3. Selecting the relevant LCA impact categories

The indicators that are recorded in the inventory analysis step of an LCA are comprehensive, as the following table illustrates.

Environmental	**Impact category / Indicator**	**Unit**
	Climate change – total, fossil, biogenic and land use	kg CO2-eq
	Ozone depletion	kg CFC-11-eq
	Acidification	kg mol H+
	Eutrophication – freshwater	kg PO4-eq
	Eutrophication – marine	Kg N-eq
	Eutrophication – terrestrial	mol N-eq
	Photochemical ozone formation	kg NMVOC-eq
	Depletion of abiotic resources – minerals and metals	kg Sb-eq
	Depletion of abiotic resources – fossil fuels	MJ, net calorific value
	Human toxicity – cancer, non-cancer	CTUh
	Eco-toxicity (freshwater)	CTUe
	Water use	m3 world eq. deprived
	Land use	Dimensionless
	Ionising radiation, human health	kBq U-235
	Particulate matter emissions	Disease incidence
Parameters that describe resources used	**Parameter**	**Unit**
	Primary renewable energy (materials)	MJ

	Primary renewable energy (energy)	MJ
	Primary renewable energy (total)	MJ
	Primary non-renewable energy (materials)	MJ
	Primary non-renewable energy (energy)	MJ
	Primary non-renewable energy (total)	MJ
	Use of secondary material	kg
	Use of fresh water	m3
	Use of renewable secondary fuels	MJ
	Use of non-renewable secondary fuels	MJ
Other environmental information: Waste type	**Indicator**	**Unit**
	Hazardous waste disposed	kg
	Non-hazardous waste disposed	kg
	Radioactive waste disposed	kg
Other environmental information: Output flows	**Indicator**	**Unit**
	Components for re-use	kg
	Materials for recycling	kg
	Materials for energy recovery	kg
	Energy production	MJ

Table 2: List of LCA impact categories (source: Ecochain)

From this list we selected the indicators for which it is possible today to either collect the required data from the manufacturers & producers of digital resources, the underlying infrastructure or equipment or that we are confident to make adequate estimates of. Throughout the document we will refer to these indicators as the 'environment impact [indicators]' of digital resources.

Environmental	**Impact category / Indicator**	**Unit**
	Climate change – total, fossil, biogenic and land use	kg CO2-eq
	Depletion of abiotic resources – minerals and metals	kg Sb-eq
	Water use	m3 world eq. deprived
	Land use	Dimensionless
Parameters that describe resources used	**Parameter**	**Unit**

	Primary renewable energy (energy)	MJ
	Primary renewable energy (total)	MJ
	Primary non-renewable energy (energy)	MJ
	Primary non-renewable energy (total)	MJ
	Use of fresh water	m3

Table 3: Selection of LCA impact categories for digital resources by the SDIA

2.4. Variable & fixed impact indicators

When it comes to digital resources, the most variable indicators of environmental impact are the energy- and water-use of the equipment and infrastructure producing digital resources.

Both of these change dynamically based on the utilization of both equipment and infrastructure. As an example, consider a server which is running at 5% utilization vs. 100%. In this case, the energy usage of the 5% utilization will be lower than in the 100% scenario. Further, the cooling infrastructure required to transport the generated heat away from the server will be at higher load in the 100% scenario (as more energy is converted to heat that needs to be removed).

Further, for both energy- and water-use there is an additional climate change impact from both. This needs to be determined based on the physical location of the digital resource production, as local environment, e.g. energy sources or water availability can have a significant impact on the climate change potential.

On the other side, the embodied impacts, such the depletion of abiotic resources and the climate change impact from manufacturing of infrastructure and equipment can be considered constant, however, they need to be attributed to each digital resource pro-rata.

As an example, consider a data center with 100 racks, containing 10 servers each, a total of 1000 servers. Let's assume the embodied CO2-eq (climate change impact) of the building is 100 tonnes. In this example, each rack is responsible for 1 tonnes of the building's embodied carbon, each server for 100 kg of CO2-eq respectively. These are constant values, however, the pro-rata attribution will change when need racks & servers are added to the facility, hence these indicators need to still be calculated dynamically.

Within the server, these fixed values need to be attributed to each digital resource produced, which requires the measurement of the total 'digital resource output' capacity, e.g. how much memory, computation, storage capacity and network bandwidth the server can produce at 100% utilization and then attributing the constant values to each unit. To that the variable indicators, water and energy need to also attributed to each digital resource, pro-rata based on the ratios specified in section 1.4.

The result is that for each digital resource produced by the specific server, in a specific location, we can determine the environmental impact indicators outlined in Table 3.

2.5. Attribution of environmental impacts in virtualized & containerized infrastructure

The majority of IT infrastructure environments are either virtualized or containerized, in which case it becomes more complex to connect the physical footprint of each virtualized container. In an ideal scenario, the virtualization software or container orchestration platform would expose the environmental impact indicators per digital resource reserved, used and unused. However in the majority of environments today, including cloud infrastruture, this information is not available.

So in these scenarios, the virtual machine or the container need to be able to make an adequate estimate of the footprint of the digital resources it is reserving on the host system, as well as consumed digital resources.

Previous research conducted on behalf of the German Environmental Protection Agency (2021 Gröger et al.) provides a general LCA model that can be used to calculate the environmental impact indicators and energy-use of a generic server, e.g. treating the virtual machine as a physical server. Furthermore, Etsy and Benjamin Davy have build estimates for the different AWS instance types that can be used. Both of these models & databases have been unified and are provided as an open API by Boavizta.

The data from these sources can be used to perform the same transformation as outlined in section 2.3, albeit on a virtual machine level.

However, it is of course much more desirable and accurate to expose the environmental impact indicators per unit of digital resource directly from the underlying infrastructure or the server. This is why the SDIA is advocating for cloud infrastructure, hosting and other digital infrastructure providers to make such information available per product unit they are responible for (e.g. per rack for a co-location space, or per digital resource for cloud infrastructure & hosting business).

3. Case study: a Life Cycle Assessment of a webpage (work in progress)

3.1. Goal and Scope

While developing our methodology, we have started to apply it to a case study. For simplicity, we will first be focusing on performing an LCA of one webpage (functional unit), namely our own homepage (www.sdialliance.org). To simplify the complexity of the LCA itself during the first application of the methodology, we have chosen to focus on two impact categories.

- Climate change – total, fossil, biogenic and land use (kg CO2-eq)
- Electricity consumption (kWh)

3.2. Data Quality Goals

Here, we "define needs for representativeness, including temporal, geographic, and technological aspects, and completeness" of our data used (LCI Data Quality Guidance).

- Temporal: We will collect data over 1 year (01-01-2021 to 31-12-2021).
- Geological: This is a tricky point within digital infrastructure, as the servers that are being used can be location all over the world.
- Technological: process design, operating conditions, material quality and process scale
- Completeness: The process completeness data quality goal details the system boundary and all flows entering, exiting and within the system boundary.

3.3. Inventory Analysis

We created a high-level and exploratory version of the lifecycle of a software product and the resources needed answering the questions: what are the inputs and outputs if we want to load the homepage of the SDIA website, and what happens if it's deleted? (Figure 8).

At the top you see the state of the product. In blue you see the different stages of a classic LCA and how it relates to the stages of the digital product. Under that you see the different components needed to create that product at each stage. We assume to build on the server, and the testing is also done on the build server, as is how it's done for the SDIA webpage.

Digital resources are the main resources that drive all steps in the chain. They are the main input of the development process, of the build process and the delivery of the website through a content delivery network.

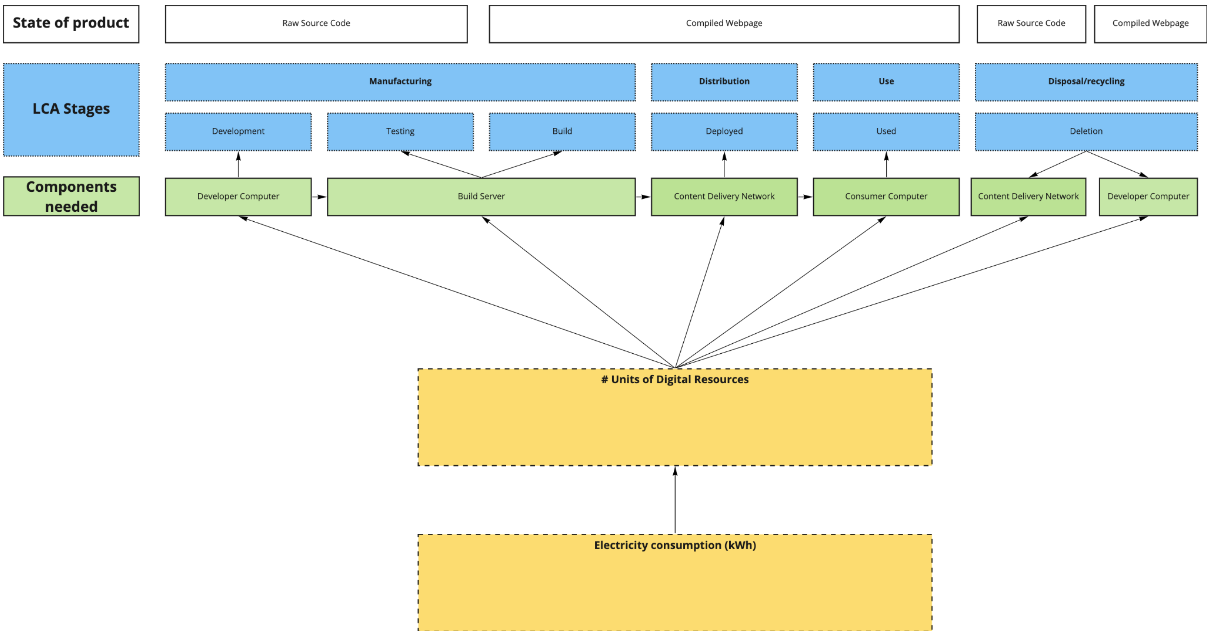

Figure 9: Lifecycle of a webpage.

Going one step deeper, we came to the following model (Figure 10) of the type of website construction we are applying our case-study to (a static website builder). In the manufacturing phase it shows the components used to generate the static web page and how they interact.

We will ultimately have to find out the environmental impact of all these different blocks shown in the model and the external data sources. The same counts for the delivery mechanism of the website (content delivery network).

The use phase is relatively simple compared to the other phases, as we will only look at the time the website is looked at by the user in the web browser.

For the end of the life of the website there are different options to consider. Either the developer ends the life of the product by deleting it, or the user closes the website and deletes it from their cache.

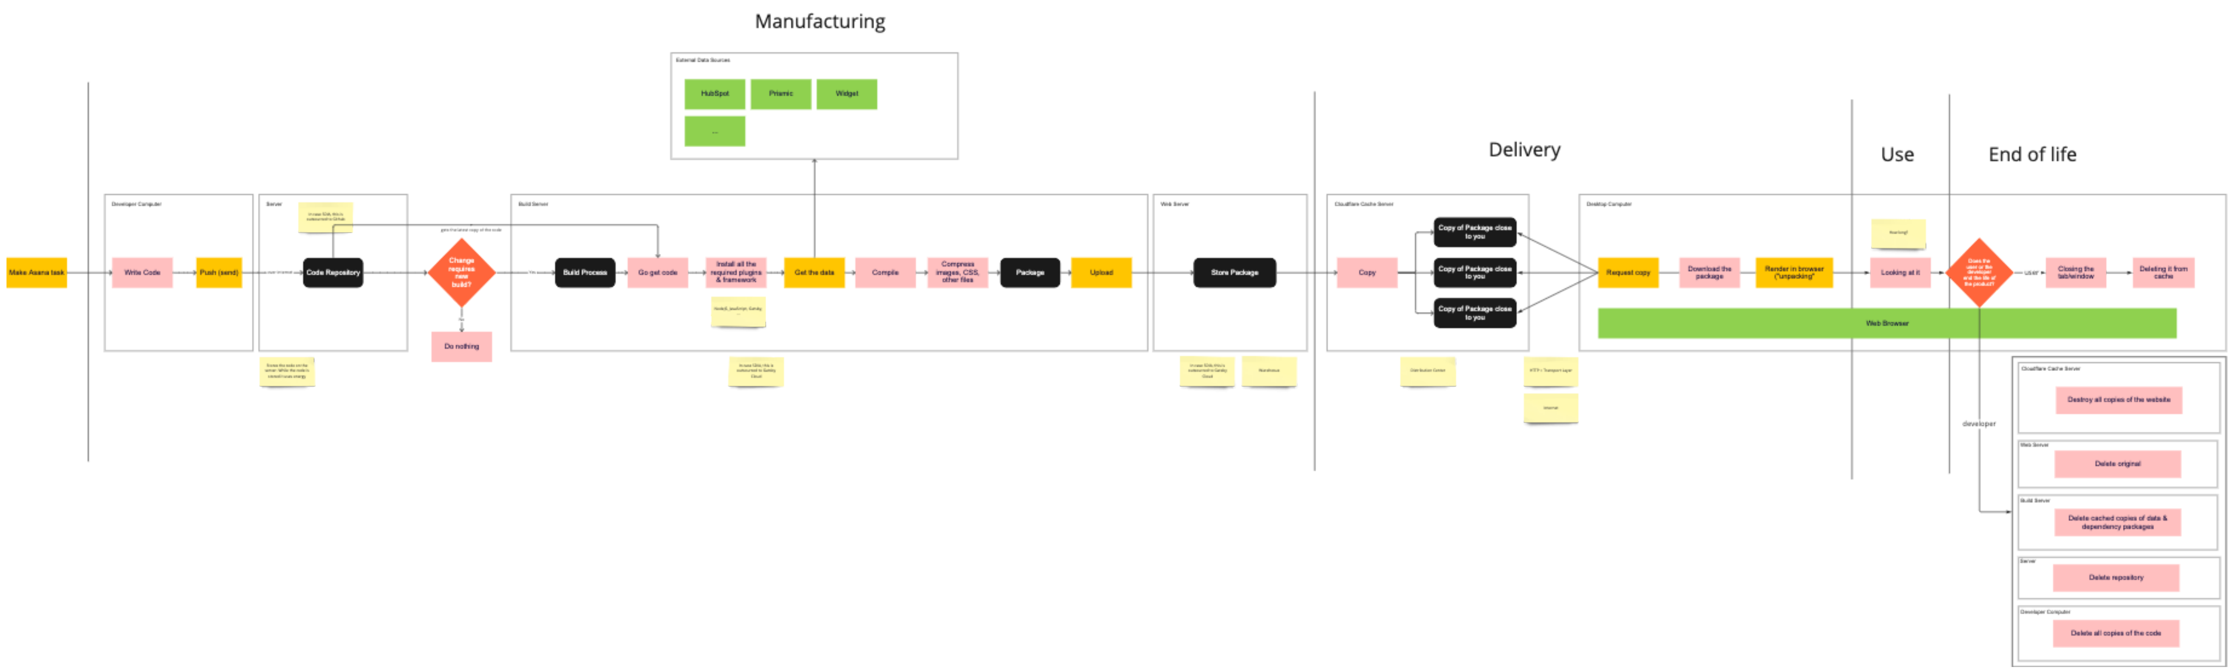

Figure 10: A more detailed view of the life cycle stages of one SDIA webpage.

4. Next steps

We believe that our methodology will open doors for sustainable development of the digital economy and digitalization, as we will get closer to measuring the actual impact of digital products and services. There are still many hurdles to overcome within our methodology and case study. Some that we have identified up till now are:

- We need **more data** from the entire digital value chain (servers, infrastructure, tools, external services, cloud infrastructure) to accurately measure the environmental footprint and do a full LCA.
- We learned that digital products can change 100 times a day through continuous delivery and integration. This is why it's key that we can **dynamically generate an LCA** for each version of the product or application.
- We need to **standardize** the conversion formulas of digital resources into environmental impacts.
- In the case of our own website, the packaged website will be distributed via a **content delivery network** will be located in different places. We need to find a way to account for this (deep dive of current case study) this type of distribution & networks.
- Getting a clear understanding of our scope. At this point in time, we will be keeping raw materials, end of life of developer and customer devices, and software obsolescence **out of scope**.

We will elaborate on a few of these next steps below.

4.1. Data transparency

A recurring issue throughout the sector is that we are missing the right data to make accurate calculations. A digital product or applications needs to make at least two things transparent: 1) the power consumption it is responsible for, 2) the embedded environmental impact in the underlying infrastructure — beginning with servers.

Right now, software does not have access to that data, as we have outlined throughout this document, because of the isolation from the resource provisioning layer. We urgently need more transparency from both resource provisioning (virtualization, container platfroms) and digital infrastructure operators on the environmental impact per-resource.

Without transparency and customer choice, there is no healthy competition focused on reducing environmental impact, only on the speed of innovation and growth. The negative externality of the resulting environmental impact is not accounted for by the market — therefore society bears the cost rather than the producer. It is the role of the regulator to step in and demand transparency from anyone selling digital products and services to customers and businesses.

4.2. Standardised methodology and conversion formulas to turn digital resources into environmental impact

We are currently working on standardised methodology. This paper and the ongoing case study are a start. We hope to continue to work on standardising this methodology with a group of experts in the field.

We do not aim for perfection of this methodology, but for something that can give an 80% accurate environmental footprint of a digital service so we can build from there. We believe that it is better to have something that has potential to create accurate and transparent results and is not greenwashing.

As for digital products the LCA has to be applied faster and should be automated based on the idea of digital resources, we see that regular LCA does not fully fit this product. Therefore, we propose to create a Digital Life Cycle Assessment (D-LCA).

We are currently working on creating formulas for outlined conversion of digital resource primitives into environmental impacts, so these can be applied to our case study as well as baseline conversion formulas that can be applied anywhere, without having data from the infrastructure or underlying server.

4.3. A label for digital products, services and applications

To accomplish the application of this model, firstly we are creating a standardized approach to make an LCA for digital products and services and secondly we will develop a label which displays the results of the LCA transparently, creating visibility for the improvements made by the vendor of the digital product, service or application.

In addition, we will also need a label for each digital resource, which contains the environmental impact information, that is produced by some cloud or hosting provider. The latter takes place in our group on A label & criteria for sustainable cloud infrastructure (for more information on this, please see our Knowledge Hub).

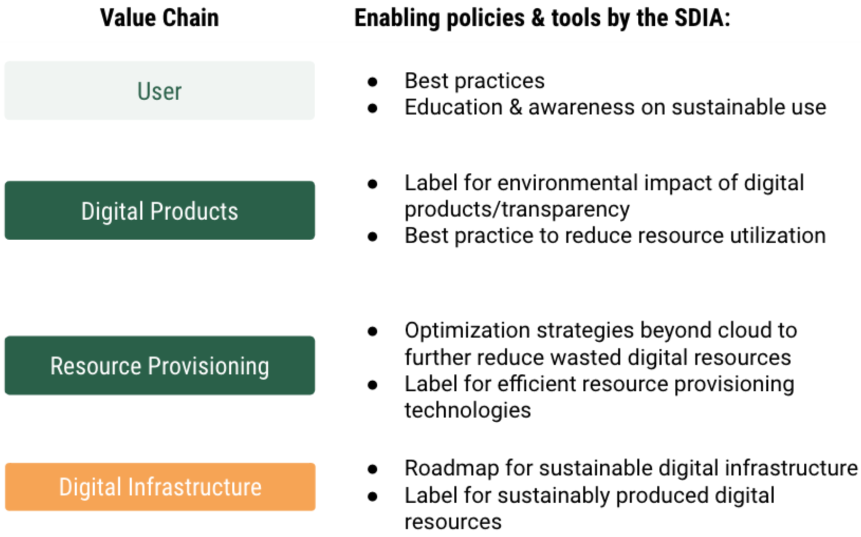

Figure 11: Plans for the labels to be created for Digital Products, Resource Provisioning, and Digital Infrastructure.

4.4. Creating a community to share & develop an open methodology

The SDIA has created a community of like-minded people and organisations. For this specific effort, we want to lean into the feedback-driven aspect, and believe it's important that we build a global community around our LCA approach, our digital resources, and conversion formulas.

We are aware of many other organisations and initiatives that address this issue. We have noticed that sometimes there is no connection between these organisations and initiatives or they are focused on creating a closed- or industry solution that is not centered on transparency & maximum reduction of environmental impact.

As the SDIA we orchestrate collaboration and knowledge sharing between all of the organizations while setting the bar high when it comes to transparency and the scope of environmental impacts that we are accounting for.

References

1. Itten, R., Hischier, R., Andrae, A. S., Bieser, J. C., Cabernard, L., Falke, A., ... & Stucki, M. (2020). Digital transformation—life cycle assessment of digital services, multifunctional devices and cloud computing. The International Journal of Life Cycle Assessment, 25(10), 2093-2098.
2. Hayri Acar, Gülfem Isiklar Alptekin, Jean-Patrick Gelas, Parisa Ghodous. Beyond CPU: Considering Memory Power Consumption of Software. Smartgreens 2016, Apr 2016, Rome, Italy. pp.417-424. Ffhal-01314070f
3. Karyakin, A., & Salem, K. (2017, May). An analysis of memory power consumption in database systems. In Proceedings of the 13th International Workshop on Data Management on New Hardware (pp. 1-9).
4. Jens Gröger, Ran Liu, Dr. Lutz Stobbe, Jan Druschke, Nikolai Richter (2021, June): https://www.umweltbundesamt.de/publikationen/green-cloud-computing