

Best Practices in Green IT

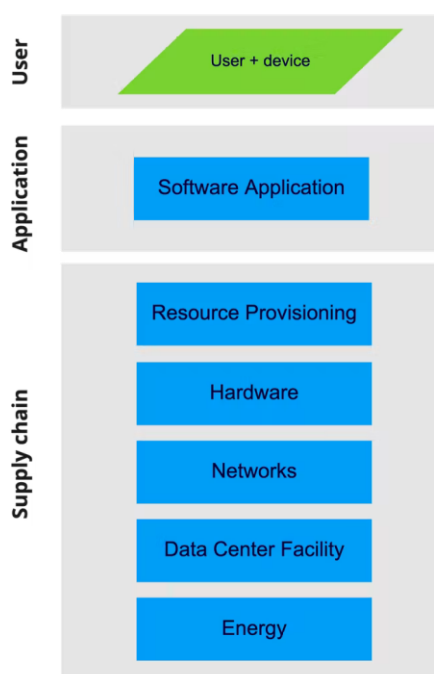
An overview of the current state of Green IT, responsibilities, and best practices for each actor.

Green IT explained	2
Responsibilities across the value chain	2
Best Practices for Users	4
Turn off applications or functionalities that are not needed	4
Ask for transparency from your application vendors	4
Best Practices for Software Application	5
Provide Life Cycle Information transparently to users	5
Minimize the number of resources being used when the application is running	5
Measure resource usage during the development of the application	6
Enable start-stop capabilities for an application	6
Classify the workload by priority	6
Best Practices for Resource Provisioning	8
Turn off or hibernate software applications when the resources they occupy could be used otherwise	8
Support switching off the actual physical infrastructure to save electricity	8
Best Practices for Digital Infrastructure	9
Recover heat (energy) or adapt power consumption to match physical green electricity	9
Support the removal of diesel generators by enabling signaling to the provisioning system or use alternative fuels	10

Green IT explained

Information technology (IT) itself can be divided into the application (what the user is seeing) and the infrastructure that is required to deliver the application to the user.

IT applications have a supply chain, it's what we call digital infrastructure in this document and it encompasses both the physical infrastructure (data center, electricity, etc.) as well as the IT infrastructure (e.g. hardware, networking, and virtualization software). Together, this digital infrastructure generates digital resources, which are used to power the IT application and provide connectivity, enabling users to connect to an application.



Users can access an application using a device, e.g. a smartphone or computer with a browser.

The application requires [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](#).

The [Radboud Centre for Green Information Technology](#) definition of Green IT: “energy-efficient hardware (greening of IT) [and] smart software solutions that contribute to the sustainable development goals (greening by IT).”

In this context, and using the taxonomy below, we clarify the definition of Green IT (focussing on the “greening of IT” aspect, as the practice of creating environmentally sustainable IT applications and digital infrastructure.

As Green IT can not exist without green digital resources (“the fuel of IT applications”), we cover the digital infrastructure aspect as part of this overview as well.

Responsibilities across the value chain

Each participant in the value-creation process has responsibilities when it comes to the 'greening' of their part of the value chain as well as the sustainable consumption of inputs and outputs.

	Role	Responsibility	Example
User	Consumer of digital applications & services, provides inputs & uses outputs	Practice sustainable & responsible consumption, e.g. turning off applications not in use, or deleting data not needed anymore	A user should turn off applications & functionalities that are not needed
Software Application	Turn user inputs into useful output, by utilizing digital resources & programming	Transparent resource consumption, minimizing resource use, use sustainable digital resources	A software developer should ensure that an application can be 'switched-off' when its not being used
Resource Provisioning	Allocating digital resources to applications	Allocating resources as efficiently as possible, recommendations on optimization towards application	A virtualization platform should re-allocate resources that are not being utilized
Digital Infrastructure	Producing digital resources for allocation towards applications	Producing digital resources that are sustainable	A data center should re-use generated heat or consume physical renewable energy sources

Participants who are active in any of these parts of the value chain can utilize the best practices below to improve the sustainability of their operations. Each measure has a direct impact on the GHG emission footprint of the organization.

Best Practices for Users

Turn off applications or functionalities that are not needed

Just like a light bulb in a room you don't use, applications or functionalities that you don't use can be turned off. They will still use natural resources¹, even though you might not notice or see it. Furthermore, if possible, try to disable the functionalities of an application, e.g. in a complex enterprise application like a CRM, that you do not need or are not using. Don't leave applications and or computers running when you don't need them.

- Turn off functionalities you don't need
- Turn off computing devices (phones, TVs, computers, servers) when you are not using them
- Uninstall or disable applications you are not using

Ask for transparency from your application vendors

You are now probably wondering how many natural resources¹ your app is using. Unfortunately, most applications do not share this information as of today. The data to provide this information is available, however, it's currently not shared by software vendors, hence you should ask for it whenever you make a purchase decision (especially as a business). Remember that also your IT application portfolio contributes to your total emissions as a company.

When you buy a car, you know how much fuel it uses per kilometer (on average). This way, we can also calculate the general environmental footprint of driving the car from A to B. For applications, we cannot do this, because the data is not available. Therefore, it's important that we start to ask for this information to be available, so we know the impact of our usage and how we can start to improve our usage behavior of these products and services.

- Request environmental footprint information when procuring IT applications (e.g. as a business or enterprise)
- Release the environmental footprint information of usage of the application to users in an organization (e.g. as a business IT operations in a company)

¹ Natural Resource: In the context of this document we define a natural resource, as a raw material, a form of energy or the emission absorption capability of an environmental medium.

Best Practices for Software Application

Provide Life Cycle Information transparently to users

To measure is to know. Users should be empowered to know what the impact of their use is. Therefore, the data about a software application should be open to the users or purchasers.

Applying methodologies like [Life Cycle Assessments](#) (LCA) can give developers and users great insights into the environmental impact of products and services. We believe that to measure is to know: the data used in an LCA can give us what we need to report on our environmental impact, and to improve our products, services, and therefore our businesses. LCA also allows for comparison between different products and services, and can, for example, help us choose the most sustainable option.

Releasing this information also enables users or purchasers to calculate the total environmental impact that is created from their usage.

- Release environmental impact data for the software application to users & procurers using a standardized framework such as the Life Cycle Assessment Methodology (for digital products & services you can find an [adapted methodology here](#))
- Enable users & procurers to understand the environmental impact created by their usage (e.g. with a calculation tool or by visualizing the impact within the user interface for each interaction)

Minimize the number of resources being used when the application is running

Especially as a software developer or vendor, there is a lot of control over the number of resources an application is using. Many IT practices suggest that digital resources are infinite; minimizing resources is not a principle that is found in many IT architecture guidelines or education.

However, digital resources are also finite and developers & software vendors are responsible for minimizing the number of resources used in the applications they publish or sell.

- Implement a resource minimization strategy & practice in your development practice
- Monitoring & measure the resource usage during the development process
- Ensure backward-compatibility of your software and minimize the amount of hardware-specific configurations you need (to ensure that it's not the software you publish or sell that drives the replacement of hardware)

Measure resource usage during the development of the application

Just as any other product, an application needs to be built/developed. This is also part of the supply chain of the product, and should in the end be reported to the customer. Therefore, resource usage during the development process should be measured, so it can be added to the environmental footprint of the usage of the application.

Further it can serve as a tool to monitor and make improvements during the process of developing the software. Similar to resource use, it is important to monitor resource efficiency during the development to identify 'hotspots' where resource efficiency is low. Based on this data, you will be able to improve your resource efficiency throughout the development process. This can also contribute to the sustainability of the final product.

- Measure and monitor the resource use and efficiency of the software application during the development process
- Identify and verify that the software application does not create hotspots when its running and that it does not contain major inefficiencies that drain resources when the application is running

Enable start-stop capabilities for an application

In order for resource provisioning platforms to be able to perform optimization, e.g. moving workloads across different environments or locations or scaling an application up or down, at the minimum, the application needs to be capable of being started and stopped without intervention.

Many legacy applications require developers or system operators to be there to verify that an application started correctly, however, this is outdated and it is possible to containerize applications so that they can be started/stopped/migrated efficiently.

It's the responsibility of the software vendor and developers to ensure that applications can be restarted without supervision to enable optimization.

- Make the application capable of restarting, e.g. when it's inside a virtual container that it can be turned off, migrated, and turned on again, without supervision
- Migrate applications into containers that contain the instructions on how they can be provisioned, what resources they require and how they can be restarted

Classify the workload by priority

In addition to providing the ability to start-stop an application, it's helpful if vendors or developers of software applications classify the requirements of their application towards the resource provisioning platform.

This way, more optimization strategies can be applied. A great example of this is to classify applications based on their availability requirements. A banking system or payment system for example might require 'always on' infrastructure, whereas a monthly payroll processing system only needs to be run on-demand (e.g. at the end of the month). Further, some applications can tolerate an outage (think of your personal blog or website) whereas others must be always available (think about the information system of an airport). ABN Amro has made application classification a tenant of their [Sustainable IT Strategy](#).

- Classify applications based on their availability criticality and make the classification information available to the resource provisioning platform for optimization
- Explicitly mark applications that can be switched off easily and are optional to run or are interruptible as well as applications that need to be highly available
- Avoid running applications on high-available infrastructure that does not require it

Best Practices for Resource Provisioning

Resource provisioning platforms are the connecting tissue between the digital resource production (e.g. from digital infrastructure: compute, memory, storage, network capacity) and consumption (e.g. IT applications).

Turn off or hibernate software applications when the resources they occupy could be used otherwise

It's the responsibility of the resource provisioning platform to make sure all resources are used as efficiently as possible. This means that resources should be re-allocated whenever possible, either by taking resources from an application away and assigning them to another one or by providing signals to applications to inform them that the requested resources are not being utilized and should be freed up.

Most resource provisioning systems already support reallocation of resources as well as overprovisioning to increase the overall utilization & efficiency of resource consumption.

- Enable signals towards applications when resources that are requested or blocked are not being utilized
- Reallocate resources that are not being utilized by an application automatically to maximize the efficiency of resource-usage
- Implement and support the workload classification and automated optimization strategies

Support switching off the actual physical infrastructure to save electricity

One of the biggest challenges in the past has been to actually turn off physical servers that are not being utilized. This had to do with mechanical components, such as spinning hard disk drives that could get damaged from many power on/off operations.

In today's server generations, most of these mechanical components have been removed or replaced with silicon/chip-based systems. For example, spinning hard-disk drives have been replaced by solid-state drives which can easily be switched on and off.

Resource provisioning platforms

Best Practices for Digital Infrastructure

[Read the guides](#) on best practices in saving energy in data centers.

Recover heat (energy) or adapt power consumption to match physical green electricity

Green electricity will become a scarce resource in our future society, especially given its importance in every industry - from transportation to heavy industry. Together with electrification itself, it's a key component in the decarbonization of our economy.

Data centers are increasingly purchasing those green electrons - some of the largest actors within the digital economy have also become the largest purchasers of green electricity (through Power Purchase Agreements, PPAs). In the northern hemisphere, which includes the Netherlands, most of the green electrons however are essentially wasted; data centers consume them, then transform them into heat energy, which is then given into the environment through cooling infrastructure. This represents a significant waste, as the energy of the green electron remains mostly unspent. Especially in the Netherlands, potential off-takers of residual heat are plentiful: from greenhouses to communities & district heating networks that are looking for carbon-free sources of heat energy.

The second phenomenon is the over-purchasing of green electrons to cover the constant power consumption of the data center ('baseload'). This is another form of wasteful behavior as it takes this scarce commodity away from other users and industries that also need access to green electricity. Instead of covering the baseload of the data center, the power consumption curve should be adjusted to match the production of renewable energy (instead of vice-versa). This can be done through thermal storage, battery systems, or through the deeper integration of resource provisioning platforms and the data center infrastructure to enable optimization strategies such as load-shifting.

- Adapt the power consumption profile of data center infrastructure to match the generation profile of nearby or national renewable energy production profiles
- Respond to signals from the grid & energy markets to adapt power consumption, instead of covering constant consumption with over-purchasing of green electrons
- As an alternative, recover & store the thermal energy that comes from the conversion of the green electrons within the data center
- [Publish a price sheet](#) for heat produced by the data center facility to facilitate third parties to 'come and get the heat'

Support the removal of diesel generators by enabling signaling to the provisioning system or use alternative fuels

Most physical redundancy has already been solved within the software application, networking and resource provisioning layer. Most modern IT applications do not rely on a 'single point of failure', e.g. a single data center facility. They use load balancers to avoid a single outage of a data center to impact the entire application. Hence many of the redundancy measures put in place within data centers, are outdated and are the result of similarly outdated standards & certifications that have been driven by financial institutions and are not applicable anymore to most modern applications.

Datacenter operators should challenge the need to install & maintain expensive and polluting diesel generators or at least replace the fuel with alternatives that are nearly carbon-neutral and pollution-free, such as SynFuels or Biofuels.

- Support signals towards the resource provisioning platforms when outages are likely and to be expected to support the further movement of redundancy and fail-over mechanisms into software
- Reduce test-runs of diesel generators to a minimum and evaluate other fuel options
- Challenge existing standards and norms and discuss with clients if diesel generators are really needed or can be removed