# DEF Formulas - Formula

## Purpose of this document

The purpose of this document is to explain a way to relate the total energy consumption of a server to the energy consumption of its components: CPU, memory, storage and networking. Especially in the context of the Cloud, it is often impossible to directly observe the energy used by a server, but often it is possible to gather metrics on the usage of these components.

Therefore, having a formula that provides a means to derive the total energy consumption based on metrics that we can observe and some necessary assumptions, would enable companies making use of Cloud-services to assess and report their Cloud-related energy consumption.

Additionally, these formulas serve to demonstrate the necessity for additional formulas relating easy-to-measure metrics to energy consumptions

# Mathematics

Given that the total energy consumed by a server can be written as the sum of the energy consumed by its component, multiplied by a overhead factor (PUE) and adding a baseline energy consumption to represent idling, we can write:

1) $$E_{tot} = [E_{cpu} + E_{mem} + E_{IO} + E_{net} + \mathbf{B}_{idle}] * PUE$$

Now, we can generally not measure the energy consumption for any of these components directly. However, we can split them out into a factor we can measure and a factor gathering the underlying complexity. For example, research has shown empirically that energy consumption of a CPU is not necessarily linear when more of a CPU or multiple CPUs are used. On the other hand, often the Thermal Design Power (TDP), which reflects the maximum amount of heat a CPU core can dissipate, is used as a linear proxy for the amount of energy a CPU consumes.

In any case, let's start off with a more general version before introducing more assumptions. The total energy used by a server, including cooling overhead is then defined as:

2) $$E_{tot} = [U_{cpu} f_{cpu}(U_{cpu}) + U_{mem} f_{mem}(U_{mem}) + U_{IOr} f_{IO}(U_{IOr}) + U_{net} f_{net}(U_{net}) + \mathbf{B}_{idle}] * PUE$$

Whereas:

- $E_{tot}$ : The total energy used by a server
- $PUE$ : Power Usage Effectiveness. The overhead, mostly due to cooling
- $U_{cpu}$ : Current CPU usage, measured in number of CPUs
- $f_{cpu}(U_{cpu})$ : Function of the CPU energy consumption per CPU usage in number of CPUs
- $U_{mem}$ : Current memory usage, measured in bits
- $f_{mem}(U_{mem})$ : Function of the memory energy consumption per memory usage in bits
- $U_{IO}$ : Current storage usage, measured in bits
- $f_{IO}(U_{IOr})$ : Function of the storage energy consumption per storage usage in bits
- $U_{net}$ : Current network usage, measured in bits
- $f_{net}(U_{net})$ : Function of the network energy consumption per network usage in bits
- $\mathbf{B}_{idle}$ : Energy consumption of the server when idling

Note that the functions denoted by $f_x(U_x)$ represent the energy allocation per unit of usage. Its unit is thus $E/U_x$. Also note that these are not necessarily linear, hence the functional depiction. I already assumed here the dependency on the actual Usage $U_x$, but perhaps there are other factors that are more relevant.

As mentioned, it is generally relatively straightforward to get measurements on $E_{tot}$, $U_{cpu}$, $U_{mem}$, $U_{IO}$ and $U_{net}$. Especially when working in a controlled environment where you can make use of voltmeters to directly observe $E_{tot}$. When any of these variables cannot be observed directly, we need to rely on more assumptions and approximations.

In contrast, the functions to relate energy consumption to usage of specific components are more difficult and generally depend on empirical hardware performance tests from the manufacturer or institutions such as SPEC (https://www.spec.org/).

## Empirical Estimations

Empirical estimations generally involve increasing the load while measuring the energy consumption. For CPU, you generally look at FLOPS (floating point operations per second) or SSJOPS (server-side java operations per second). See for example: https://www.spec.org/power_ssj2008/results/res2017q3/power_ssj2008-20170829-00780.txt

As far as we know, these are not readily available for all types of hardware and all types of components. Even if they are available, they often require quite a bit of domain knowledge to be properly understood.

Yet, a mathematical formulation linking energy consumption directly to usage would be the most accurate way to estimate energy consumption afterwards.

## Cloud

Looking at formula 2), a user of Cloud services is not able to perceive the energy consumption directly ($E_{tot}$) , and for various services it is not even possible to observe the usage of CPU, memory, storage and networking ($U_x$).

For the services that do expose these measures, we generally do not know the exact energy consumption given the level of usage ($f_x(U_x)$). Without these formulas, we can only get an estimation of the total energy used by a server, as well as its individual components, by using several assumptions.

## Best Effort Assumptions

So, let's provide a best-effort methodology, using the following assumptions:

- $\alpha_{CPU}$ of the energy consumed by a server when running on full load is due to CPU
- $\alpha_{mem}$ of the energy consumed by a server when running on full load is due to memory
- $\alpha_{IO}$ of the energy consumed by a server when running on full load is due to storage
- $\alpha_{net}$ of the energy consumed by a server when running on full load is due to networking
- The energy consumed by a server increases linearly relative to the increase in usage
- All energy of a server can be allocated to these 4 components:
  $\alpha_{CPU} + \alpha_{mem} + \alpha_{IO} + \alpha_{net} = 1$
- We can use the TDP as a proxy for energy consumption by a single CPU core when fully using the CPU,

Let $E_{tot-max}$ denote the total energy consumed by a server when running on full performance.

And:

3) $$E_{cpu-max} = U_{cpu-max} f_{cpu}(U_{cpu-max}) = TDP * N_{core}$$

Whereas:

- $N_{core}$ denotes the number of CPU cores

Then:

4) $$E_{tot-max} = [E_{cpu-max} * PUE]/\alpha_{CPU} = [TDP * N_{core} * PUE]/\alpha_{CPU}$$

Substituting $E_{cpu-max}$ by $U_{cpu-max} f_{cpu}(U_{cpu})$ and solving for $f_{cpu}(U_{cpu})$:

5) $$f_{cpu}(U_{cpu}) = (TDP * N_{core})/U_{cpu-max}$$

And finally for the energy related to the CPU:

6) $$E_{cpu} = U_{cpu} f_{cpu}(U_{cpu}) = (U_{cpu}/U_{cpu-max})([TDP * N_{core}]/\alpha_{CPU})$$

For the other components we get:

7) $$E_{tot-max} = [E_{mem-max} * PUE]/\alpha_{mem} = [U_{mem-max} f_{mem}(U_{mem-max}) * PUE]/\alpha_{mem}$$

Plugging in 4) for $E_{tot-max}$, and solving for $U_{mem} f_{mem}(U_{mem})$:

8) $$\alpha_{mem} [[TDP * N_{core}]/\alpha_{CPU}] = U_{mem} f_{mem}(U_{mem})$$

Or:

9) $$(\alpha_{mem}/\alpha_{CPU})[TDP * N_{core}] = U_{mem-max} f_{mem}(U_{mem-max})$$

Solving for $f_{mem}(U_{mem})$, and since $f_{mem}(U_{mem})$ does not actually depend on $U_{mem}$:

10) $$f_{mem}(U_{mem}) = (\alpha_{mem}/\alpha_{CPU})[TDP * N_{core}]/ U_{mem-max}$$

Meaning we end up with the total energy related to memory:

11) $$E_{mem} = (U_{mem}/ U_{mem-max})(\alpha_{mem}/\alpha_{CPU})[TDP * N_{core}]$$

The energy related to the other components can then be found in a similar fashion, due to symmetry.

Given that the allocation vector $\boldsymbol{\alpha}$ is a constant, the energy functions $f_x(U_x)$ become constants too. Which implies that the partial derivatives, the marginal energy cost related to usage of component x, can be simplified to that constant multiplied by the PUE:

12) $$\delta E_x / \delta U_x = f_x (U_x) PUE$$

# Example

| Name | Variable | Value | Unit |
|---|---|---|---|
| PUE | *PUE* | 180 | % |
| Total CPU sockets per server | $U_{me}$ | 2 | Sockets |
| Total cores per physical machine | $U_{cpu-max}$ | 96 | vCPUs |
| Total memory per physical machine | $U_{mem-max}$ | 384 | GB |
| Total IO per physical machine | $U_{IO-max}$ | ? | ? |
| Total network bandwidth per server | $U_{net-max}$ | 50000 | MBit s |
| Rated TDP per CPU | *E* | 240 | W |
| CPU power attribution | $\alpha_{CPU}$ | 65 | % |
| Memory power attribution | $\alpha_{mem}$ | 20 | % |
| Storage power attribution | $\alpha_{IO}$ | 10 | % |
| Network power attribution | $\alpha_{net}$ | 5 | % |

Energy per CPU:

$$E_{cpu} = (1/U_{cpu-max})\,[TDP\ *\ N_{core}] = (1/96)[240\ *\ 2] = 5$$

Energy per GB:

$$E_{mem} = (1/\,U_{mem-max})(\alpha_{mem}/\alpha_{CPU})[TDP\ *\ N_{core}\ *\ PUE\,]$$

$$E_{mem} = (1/384)(0.2/0.65)[240\ *\ 2]$$

$$E_{mem} = 0.384$$

Note that this only contains the direct Scope 1 energy consumption. If we also want to account for indirect Scope 3 consumption, we can get a rough estimate by multiplying these energies by the PUE:

$$E_{cpu-s3} = E_{cpu}\ *\ PUE = 5\ *\ 1.8 = 9$$

$$E_{mem-s3} = E_{mem}\ *\ PUE = 0.384\ *\ 1.8 = 0;692$$

By symmetry, the energy for the other components can be obtained by replacing $\alpha_{mem}$ and $U_{mem-max}$ to corresponding values.