September 2024

ensiie
Paris | Évry-Courcouronnes

QARNOT

Qarnot: an introduction through sustainable and technical challenges

Charles Anteunis
Software Engineer
Promo 2021

Jérémie Spiesser
Software Engineer
Promo 2022

- **How carbon intensive is cloud computing ?**
- What are the levers to reduce it?
- How is Qarnot helping?
- A few technical challenges

What is the cloud?

The cloud is composed of different services

Office 365

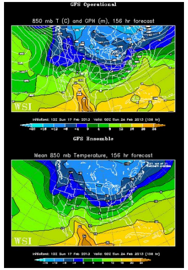

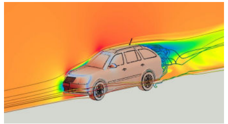

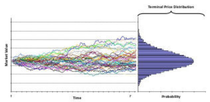

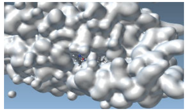

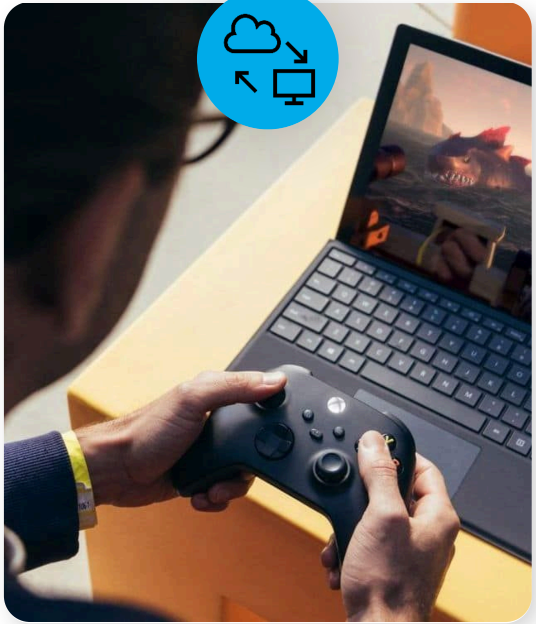

What is the cloud?

Are cloud infrastructures as light as this cloud?

QARNOT

What is the cloud?

Important infrastructures behind cloud services

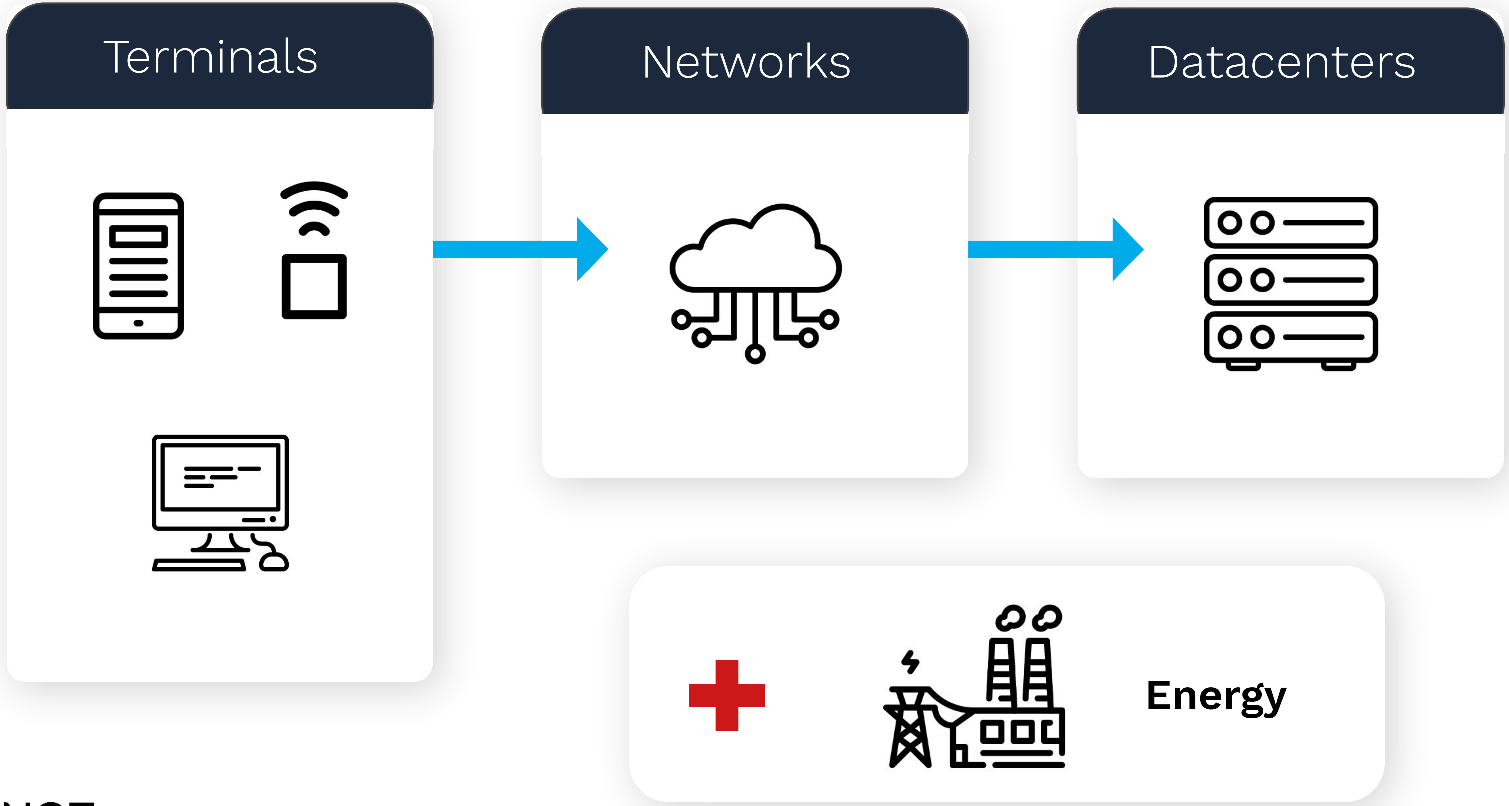

What is the cloud?

Important infrastructures behind cloud services


Internet cable laying vessel


Intel's Fab42 factory in Arizona, USA


Facebook datacenter in Texas, USA


Aluminium factory in Dunkerque, FR

Like every industry, digital has an impact

1,600 millions tons of CO_2 per year emitted by the Information Technology Industry (ICT)

4%+ of worldwide carbon footprint
That's more than the whole Aviation Industry

Growing **9%** per year

ICT also consumes scarce resources
(land, water,rare metals...)

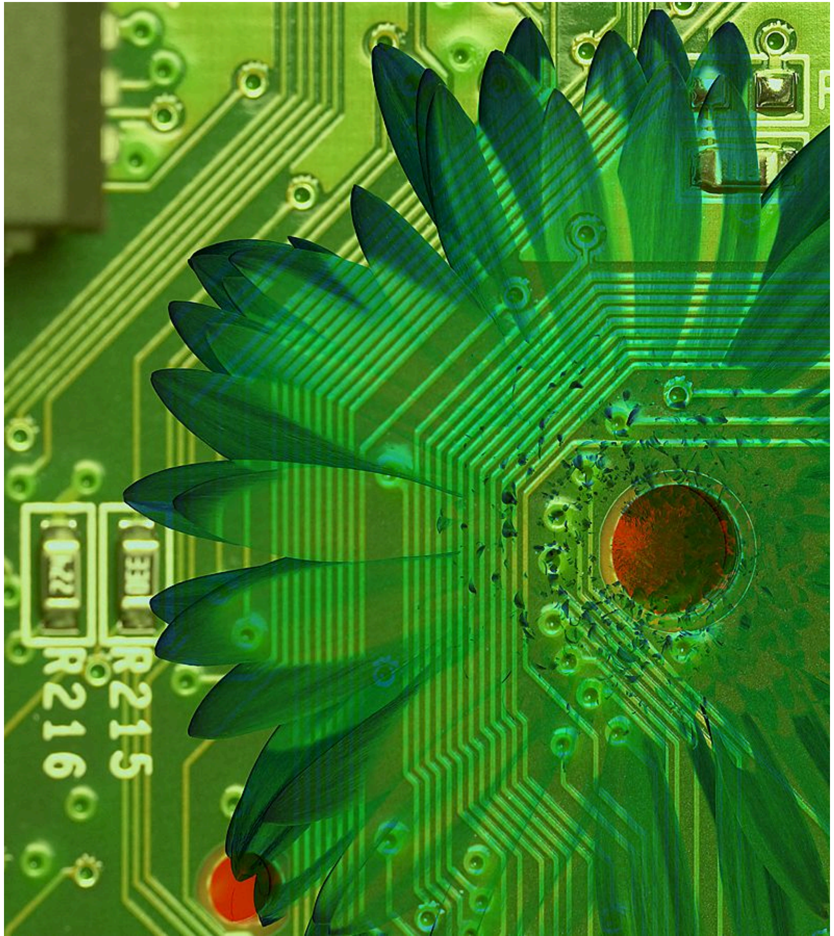

QARNOT

Source: [Lean ICT - Towards sobriety]() - The Shift Project - March 2019

Communication from hyperscalers

“Switching from on-prem to the cloud will save 90% of your carbon footprint”

Azure

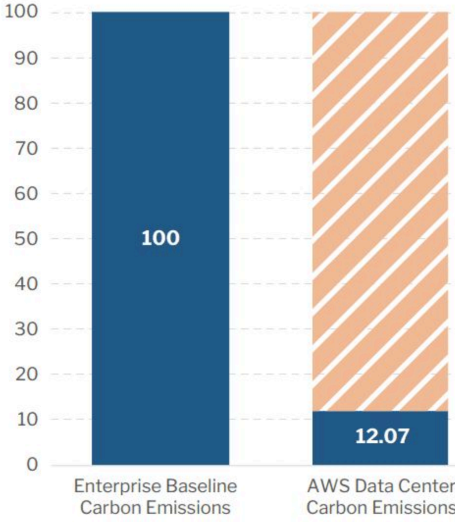

aws

Communication from hyperscalers

“The cloud is carbon neutral”

aws

Neutres d’ici 2040

« Nous avons pour objectif d’atteindre zéro émission nette de CO_2 d’ici 2040 »

Microsoft

Négatif d’ici 2030

« D’ici 2030, Microsoft aura une empreinte carbone négative »

Google

Neutre depuis 2007

« En 2007, nous étions la première grande entreprise à s'engager et à atteindre un bilan neutre en carbone »

- How carbon intensive is cloud computing ?
- **What are the levers to reduce it?**
- How is Qarnot helping?
- A few technical challenges

Carbon footprint of the cloud

Carbon footprint =

energy consumption x **electricity mix**
– energy double usage
+ manufacturing
+ end of life
– carbon offset
+ …

Carbon footprint of the cloud

Carbon footprint =

energy consumption x **electricity mix**

– energy double usage

+ manufacturing

+ end of life

– carbon offset

+ …

Reduce energy consumption

Increase utilization rate

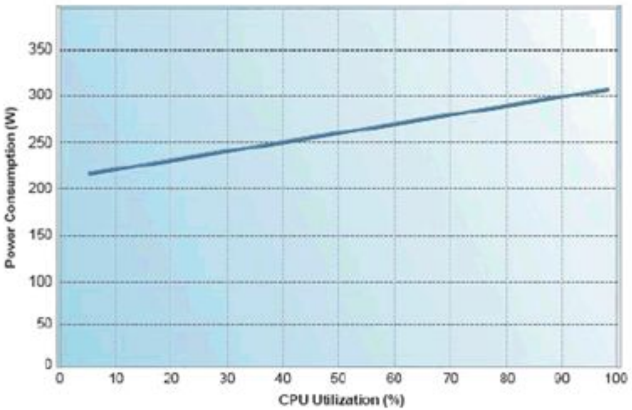

CPU Utilization and Power Consumption (Source: Blackburn 2008)

Servers

Option 1: traditional solution

Large number of servers with low utilization and (often) sub-optimal efficiency

Option 2: cloud-based solution

Smaller number of cloud servers with high utilization and efficiency

Reduce energy consumption

Evolution of the cooling technology

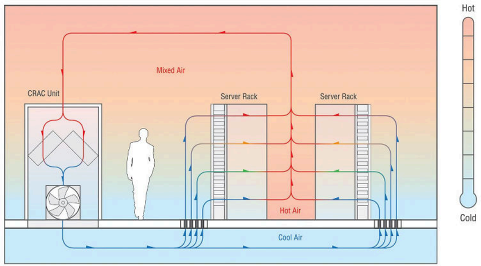

Traditional cooling

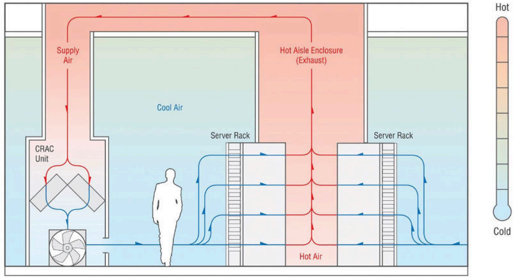

Hot aisle enclosure

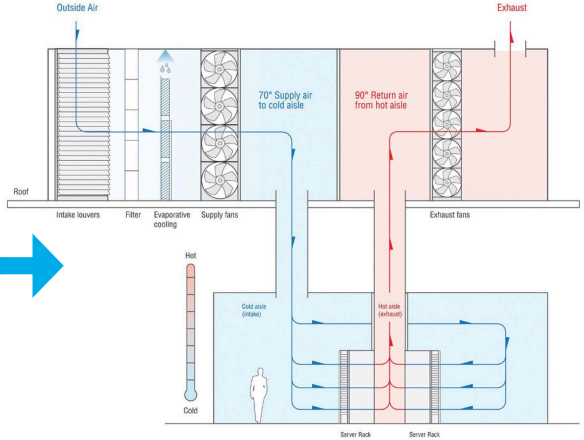

Free and evaporative cooling

Reduce energy consumption

Evolution of the cooling technology

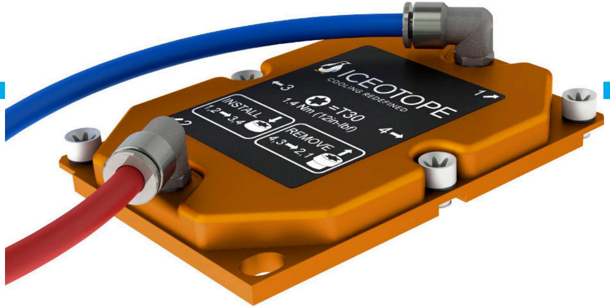

Direct Liquid Cooling

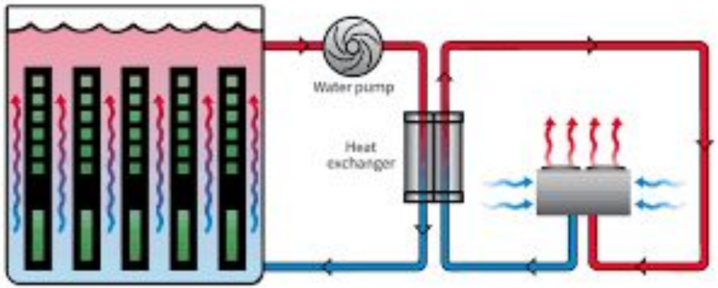

Immersion cooling

Reduce energy consumption

PUE: a key metric for the energy efficiency

PUE

Power Usage Effectiveness

$$\frac{EnergyIT + EnergyCooling + EnergyMisc}{EnergyIT}$$

PUE : impact de la métrique

DC efficiency gains have flattened out, we need new perspectives

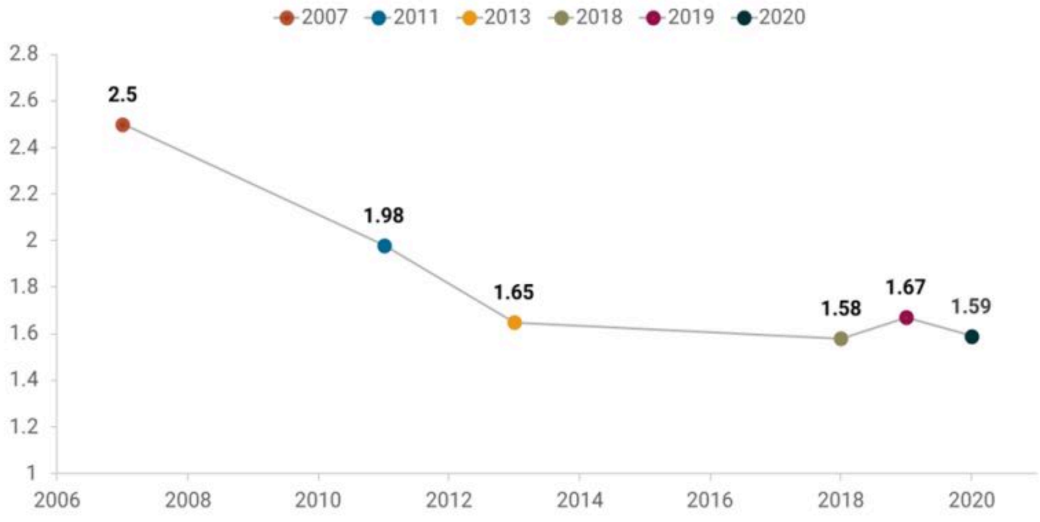

Source: Reported data center PUE figures in global Uptime Institute surveys from 2007 to 2020

UptimeInstitute | INTELLIGENCE

QARNOT

Carbon footprint of the cloud

Carbon footprint =

energy consumption x **electricity mix**

– energy double usage

+ manufacturing

+ end of life

– carbon offset

+ …

Electrical mix

Buying renewables certificates doesn't help reaching Paris agreement targets

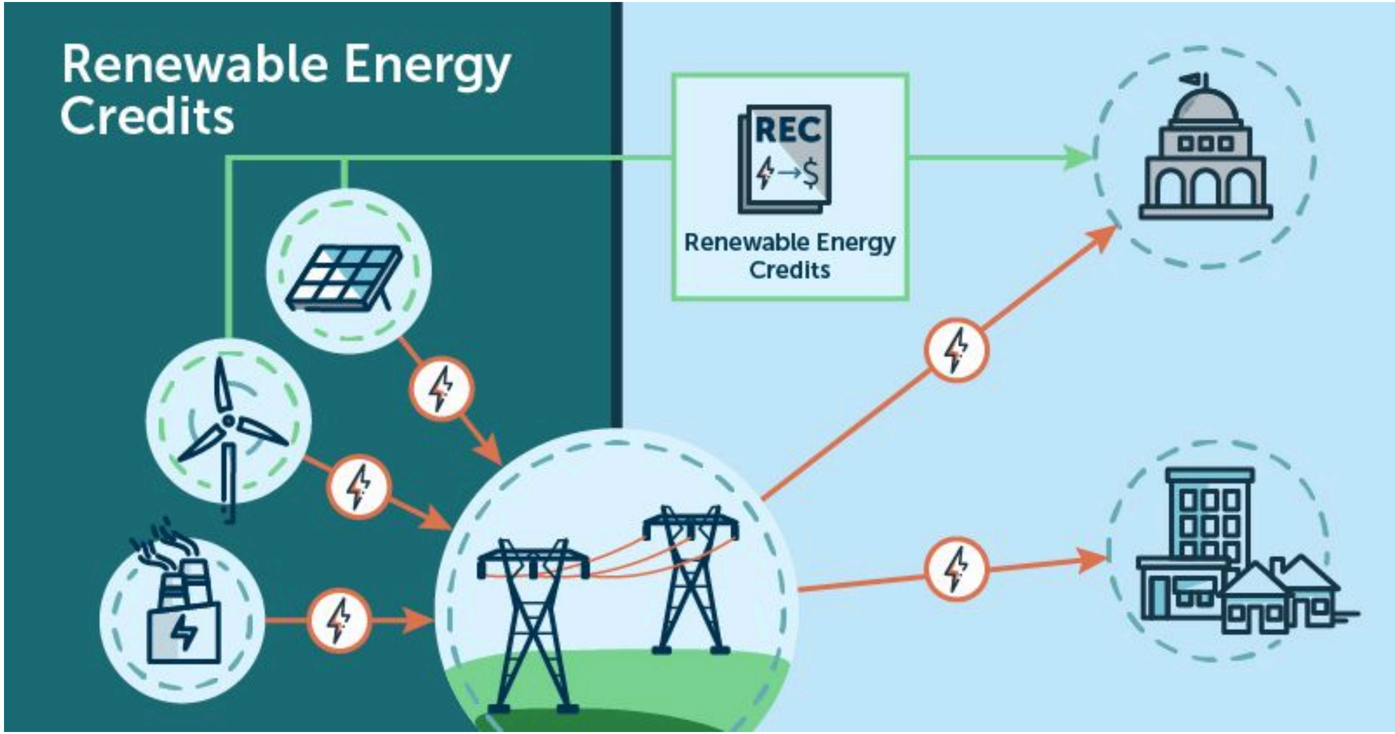

Carbon footprint of the cloud

Carbon footprint =

energy consumption x **electricity mix**

– energy double usage

+ manufacturing

+ end of life

– carbon offset

+ …

Carbon offsetting

Carbon offset is part of the solution but comes with a great deal of uncertainty

Today

In 30 years ???

Carbon footprint of the cloud

Carbon footprint =

energy consumption x **electricity mix**

– energy double usage

+ manufacturing

+ end of life

– carbon offset

+ …

Carbon footprint of the cloud

Consider the full Life Cycle Analysis

Manufacturing

Usage

Carbon footprint of the cloud

Consider the full Life Cycle Analysis

USA

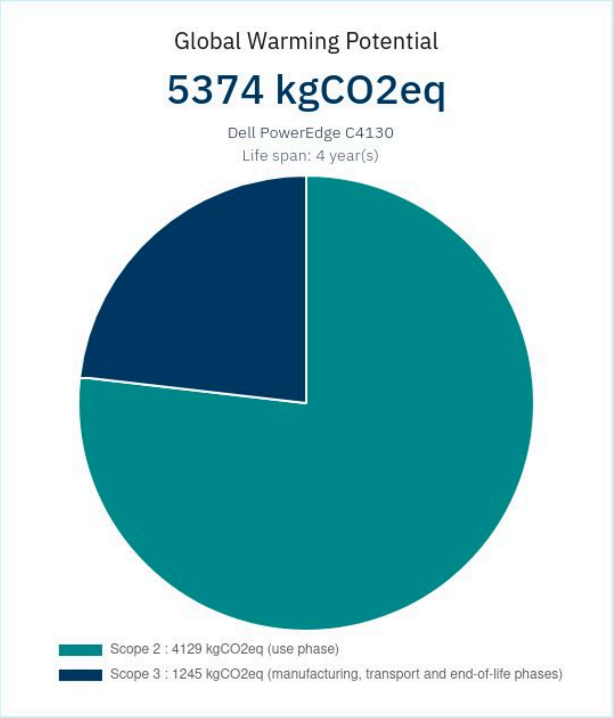

France

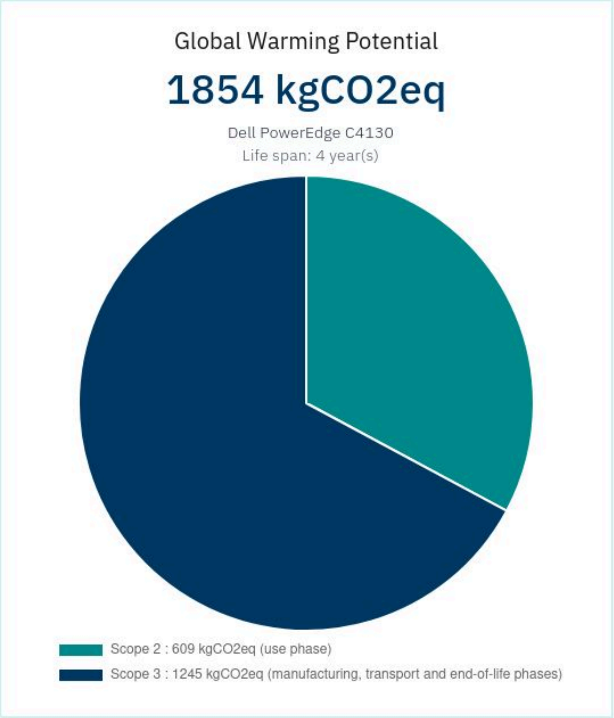

Carbon footprint of the cloud

Consider the full Life Cycle Analysis

Parts	Emissions (kgCO2eq)
8*3.84TB Solid State Drives	3379
12*32GB DIMMs Memory	533
Mainboard PWB	109
1*400GB Solid State Drive	64
Riser card 1 - Riser card 2 - Riser card 3 - Ethernet card - HDD Controller - Q-logic - Intel Ethernet X710	59
2*Xeon CPUs with housing	47
Chassis	34
PSU	30
Mainboard Connectors - Transport	20
Fans	13
Total	**4288**

Répartition des emissions liées à la fabrication du Dell R740

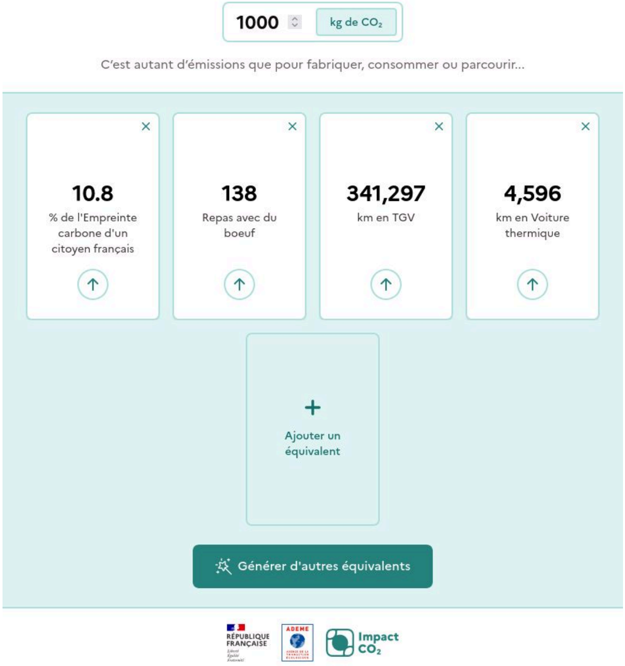

DC initiatives

Microsoft underwater datacenter

Scandinavian Datacenter

Jean Zay supercomputer
heating IDRIS CNRS labs

DC initiatives

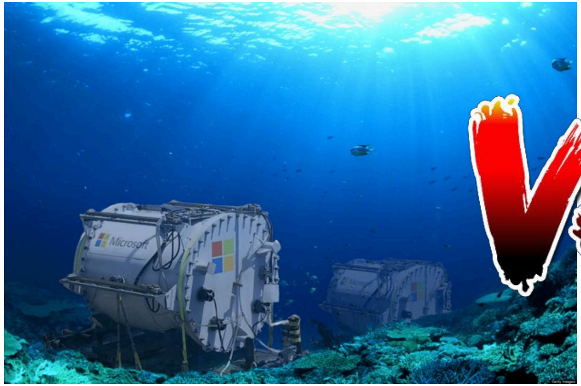

Microsoft underwater datacenter
Pure marketing

What about maintenance?

Scandinavian Datacenter
Better

Reduced cooling
Better electricity mix?

Jean Zay supercomputer heating IDRIS CNRS labs
Much better

Reduce used energy
Not reduced cooling

Sustainability is not just about carbon

Carbon footprint =
carbon footprint
+ water usage
+ abiotic depletion
+ …

DATACENTER ACTU

La Hollande soumise aux pénuries découvre que les DC de Microsoft ont consommé 84 millions de litres d'eau potable

YVES GRANDMONTAGNE 19 AOÛT 2022

Alors qu'elle a subi des pénuries d'eau sur son territoire, la Hollande découvre avec stupeur que les datacenters Middenmeer de Microsoft ont consommé 84 millions de litres d'eau potable en 2021, soit 4 à 7 fois plus que les engagements pris.

Water usage

Abiotic depletion

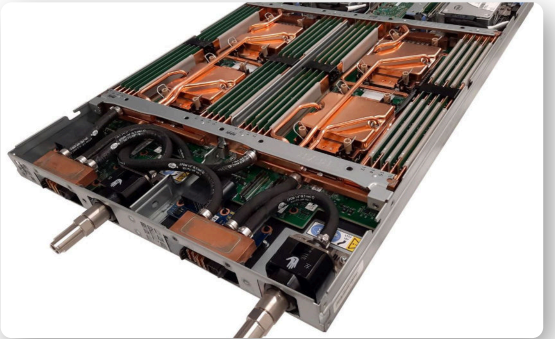

Source : https://datacenter-magazine.fr/la-hollande-soumise-aux-penuries-decouvre-que-les-dc-de-microsoft-ont-consomme-84-millions-de-litres-deau-potable/

September 2024

- How carbon intensive is cloud computing ?
- What are the levers to reduce it?
- **How is Qarnot helping?**
- A few technical details

QARNOT

A specific hardware solution : QB·1 and QB·X

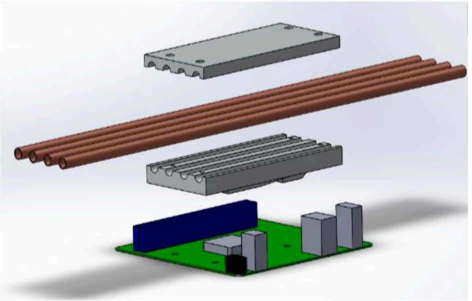

A specific hardware solution : QB·X

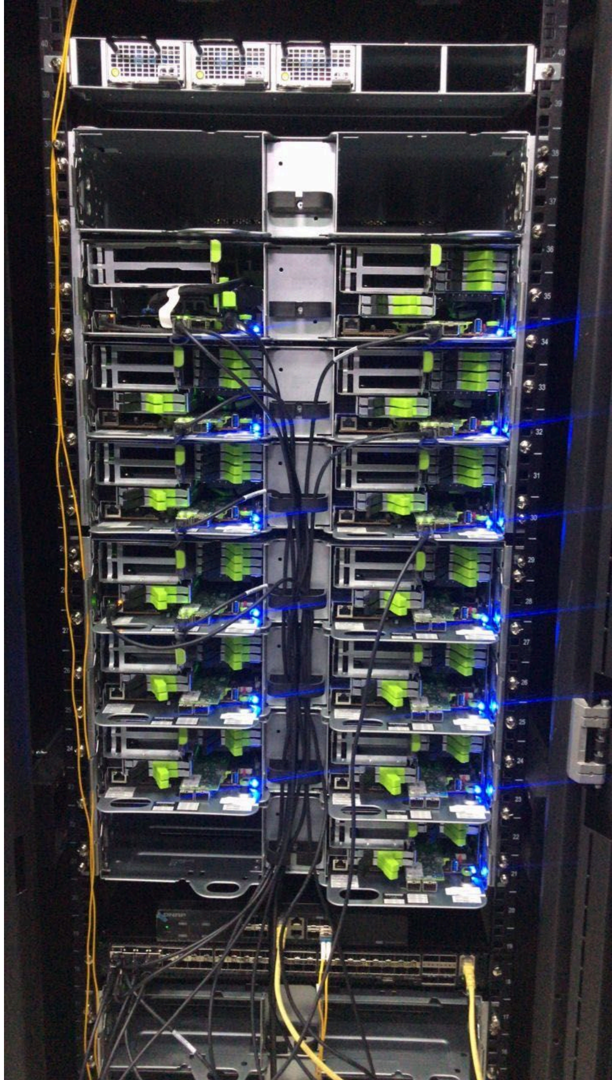

QARNOT

Installing our servers where the heat is needed

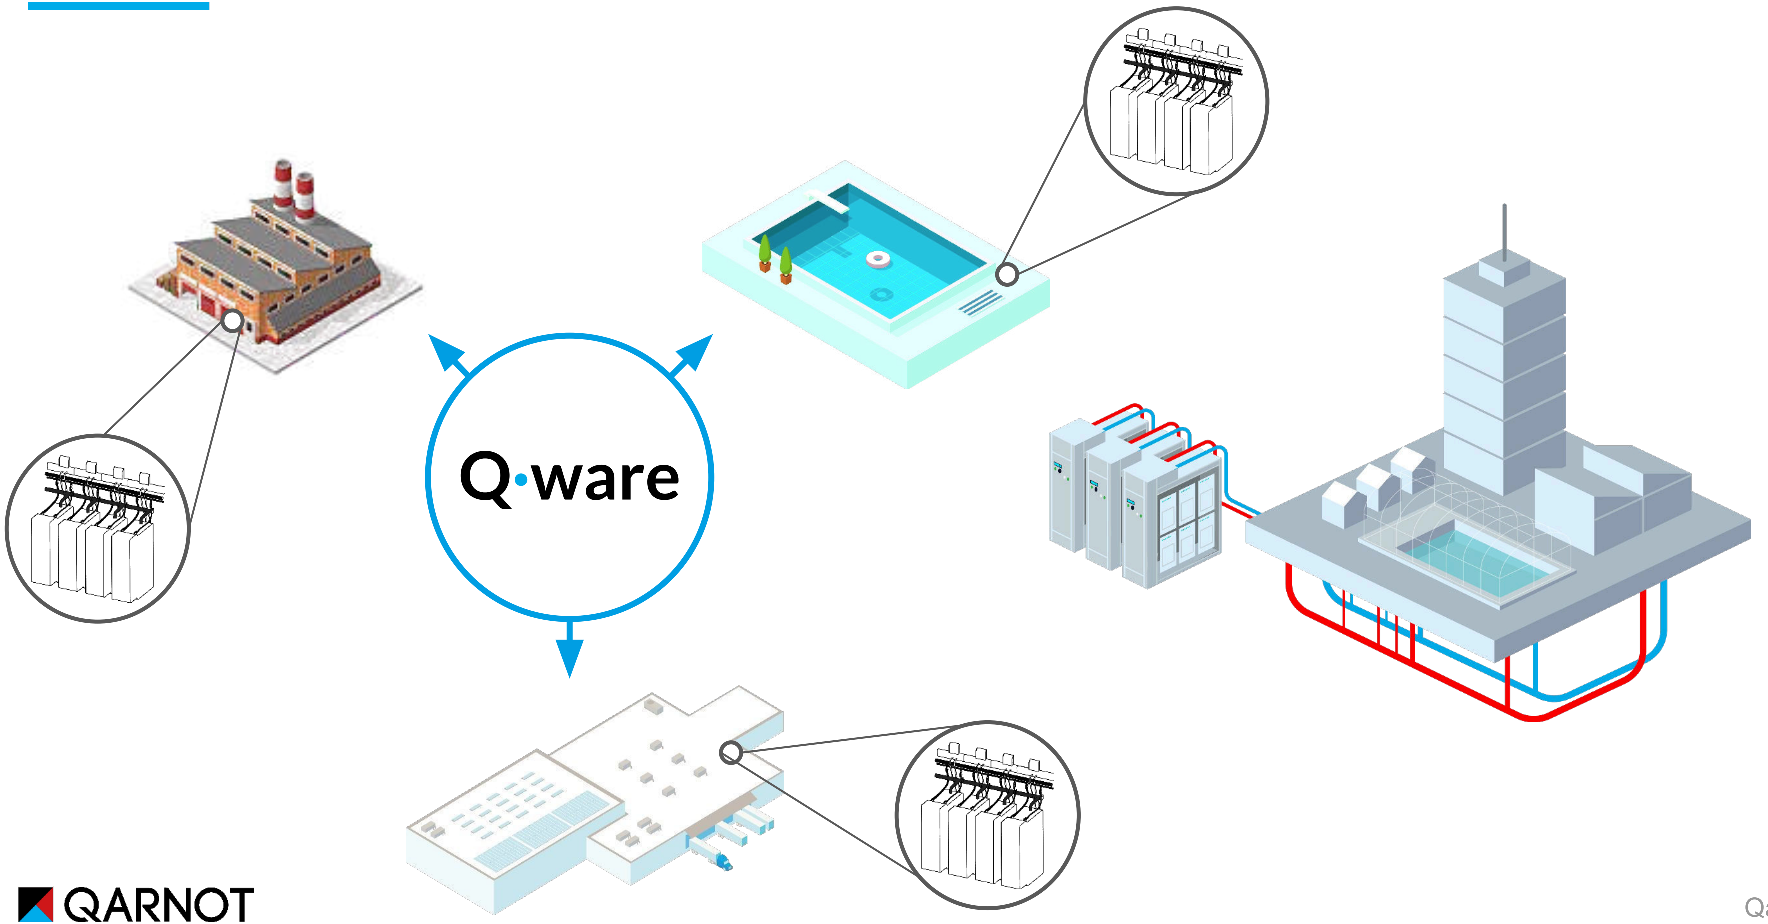

How do you manage in summer?

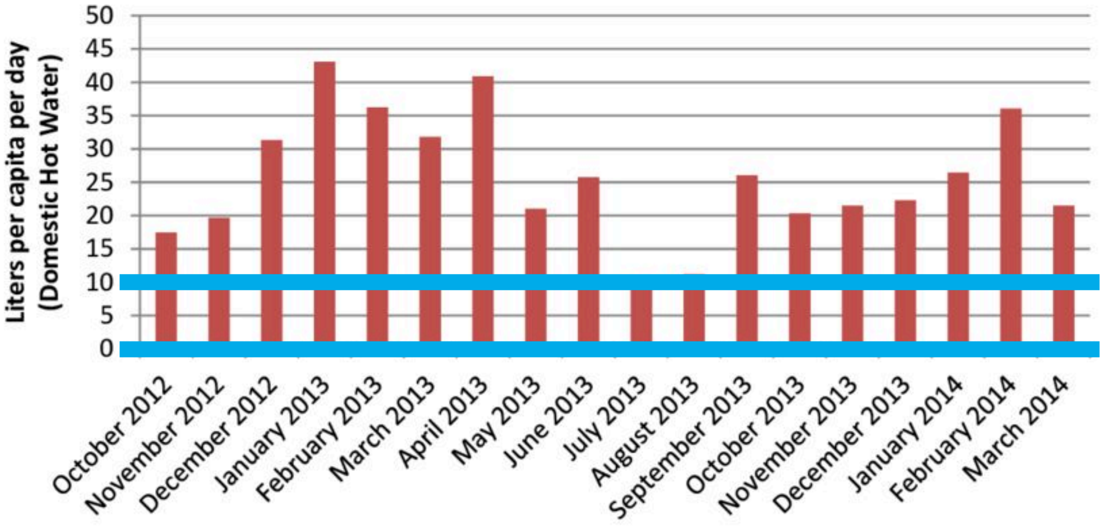

Carbon footprint of Qarnot

Carbon footprint =

energy consumption x **electricity mix**

– **energy double usage**

\+ **manufacturing**

\+ **end of life**

– **carbon offset**

\+ **…**

Carbon footprint of Qarnot QB.x

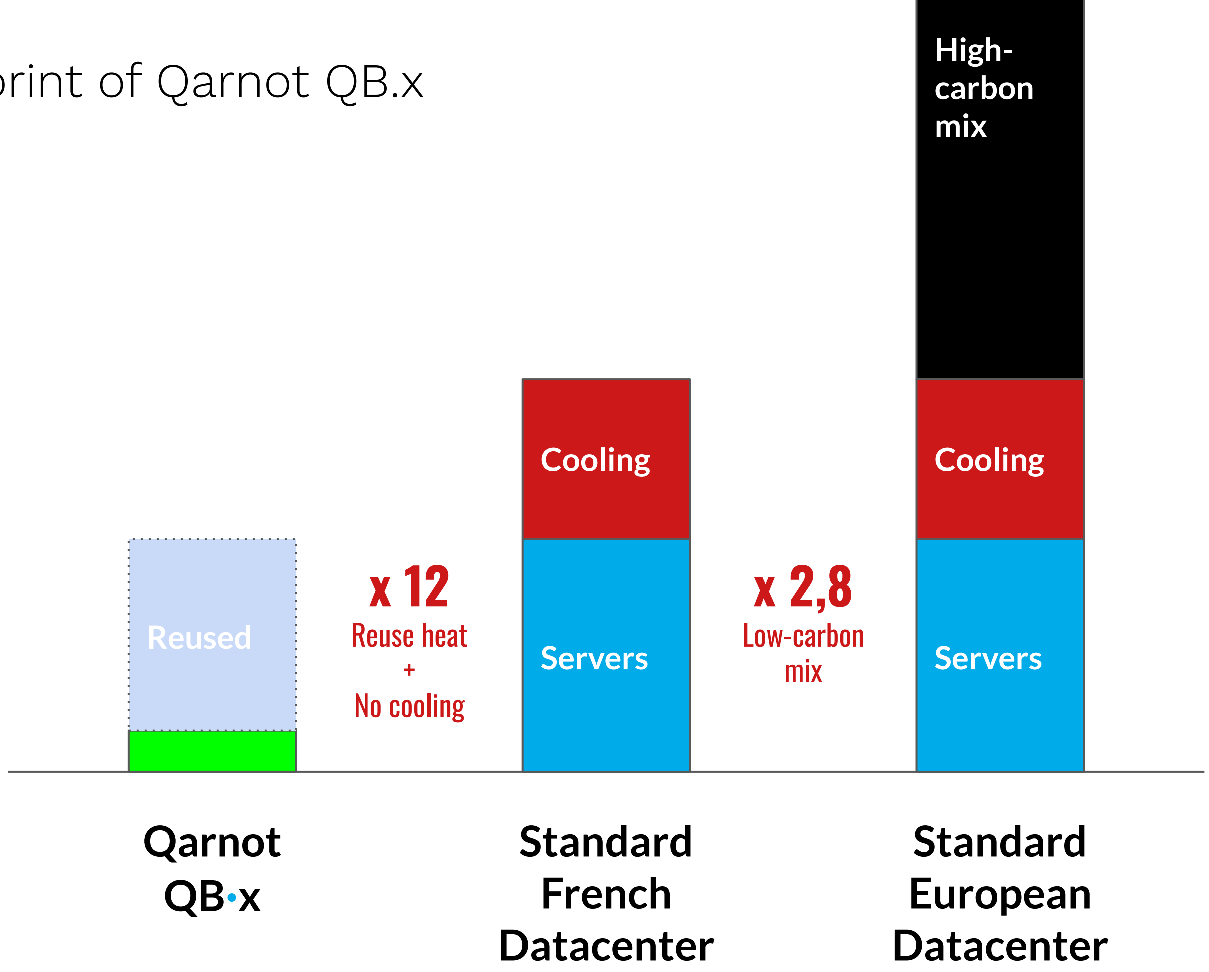

QARNOT

Carbon footprint of a computation on QWare

Carbon Facts

name	**Task on a mobo on a QBx OCP Leopard**
duration	**365d**

Saved Carbon Footprint **1.53** T CO_2e

Reduced Emissions **88.05** %

Energy

Total consumed	**3.47** MWh
Average power	**396** W
Reused energy	**3.10** MWh
Energy Reuse Effectiveness (ERE)	**0.11**
Avg Power Usage Effectivness (PUE)	**1.001**
Avg Energy Reuse Factor (ERF)	**89.2** %

Carbon

Qarnot carbon footprint	**0.21** T CO_2e
Carbon emission avoided by reused heat	**0.70** T CO_2e
Saved carbon footprint	**1.53** T CO_2e
Emmision reduced	**88.05** %

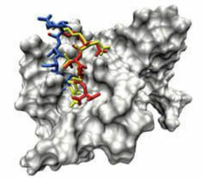

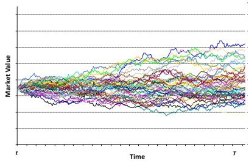

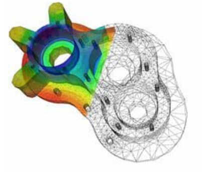

QARNOT

Computing
Clients who need computing power

QALWAY
QARNOT

Building
Clients who need heat

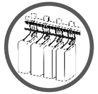

Computing boiler

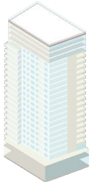

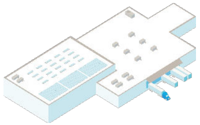

HPC use cases

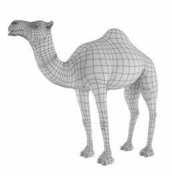

3D ANIMATION

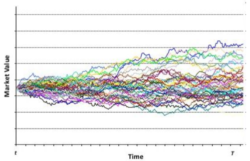

FINANCES

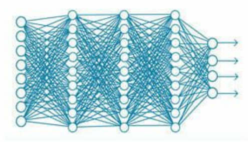

MACHINE LEARNING

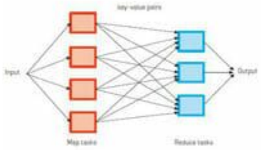

BIG DATA

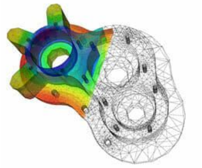

MECHANICS

ELECTRONICS

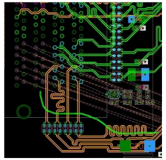

OPERATIONS RESEARCH

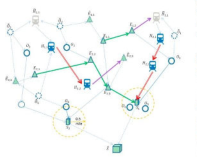

BIOTECH

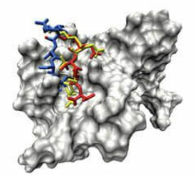

FLUID DYNAMICS

WEATHER FORECAST

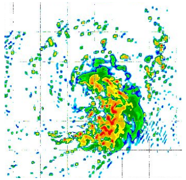

How to launch a computation?

API

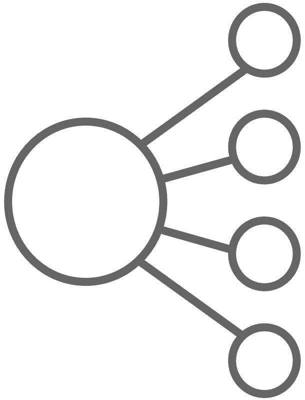

SDKs

Console

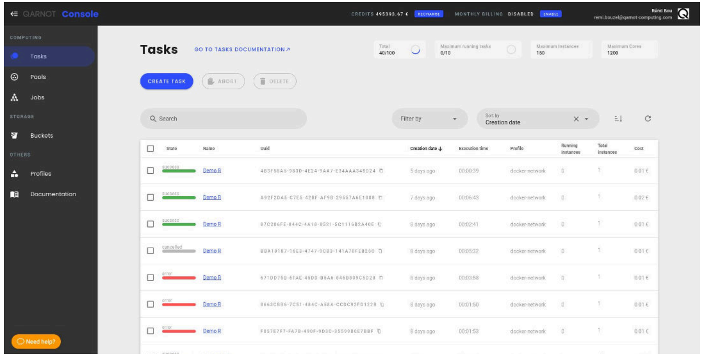

Software Architecture

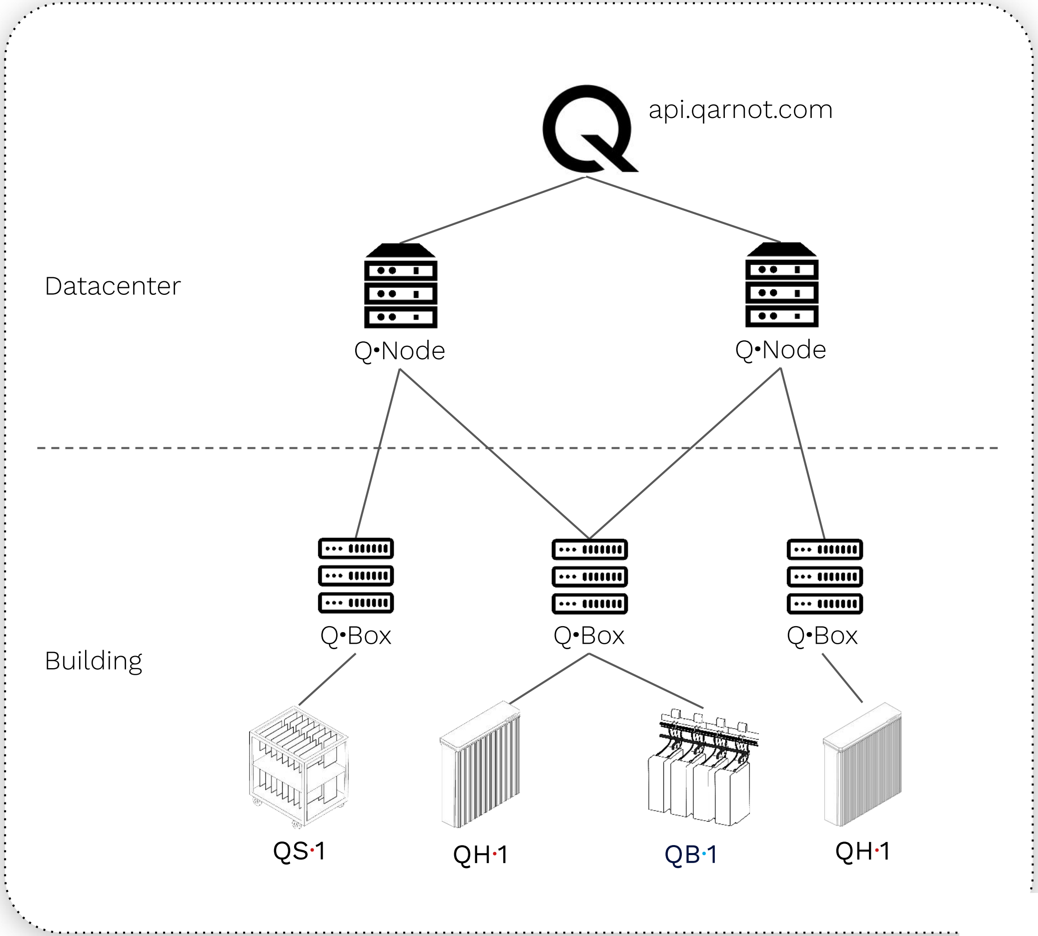

Software Architecture

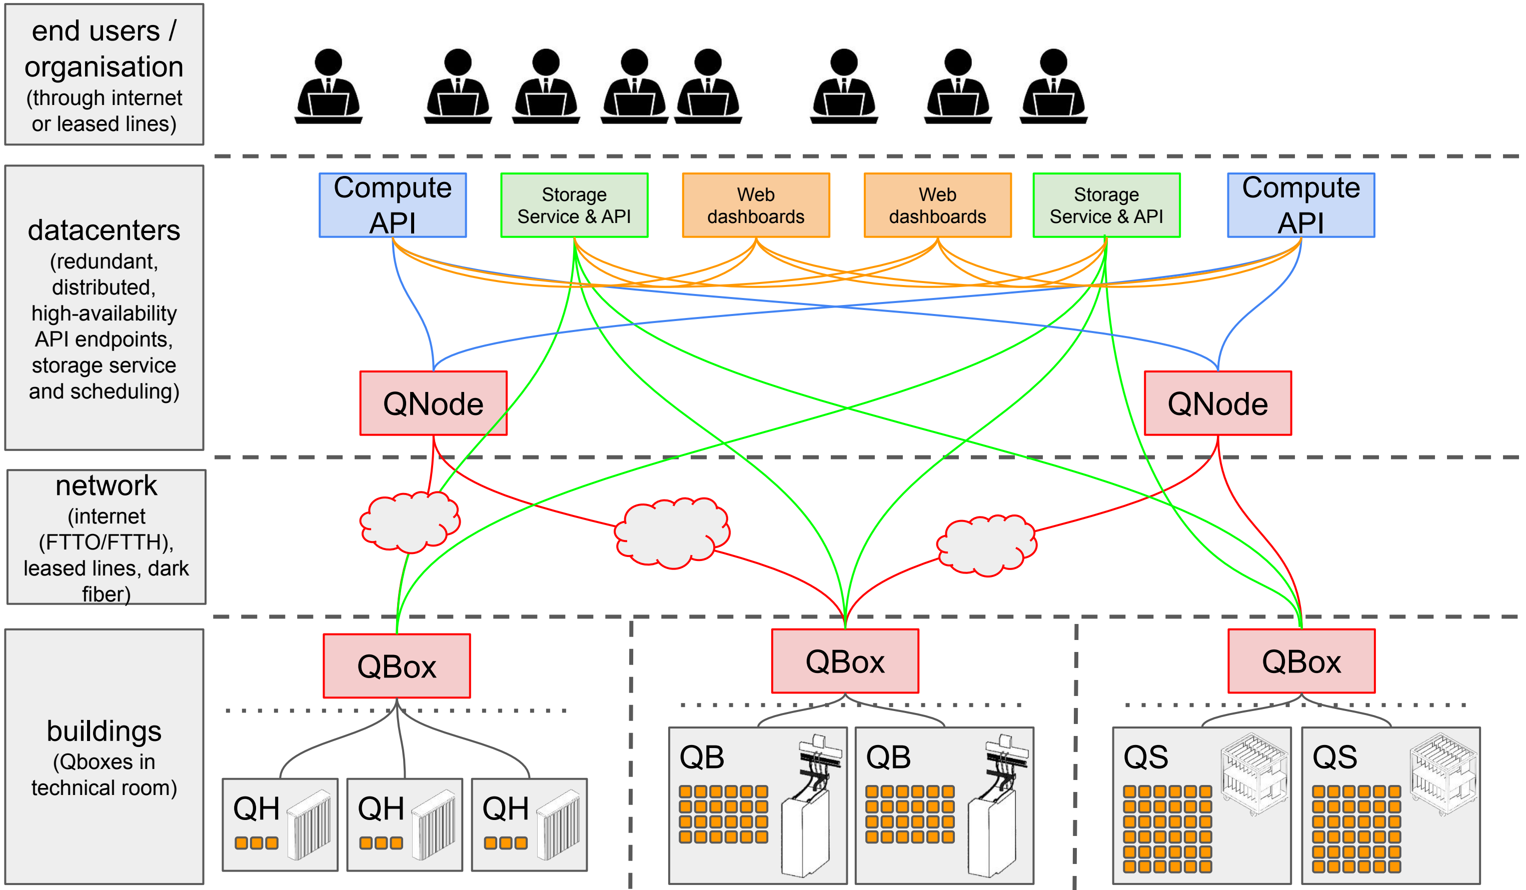

Qarnot

- How carbon intensive is cloud computing ?
- What are the levers to reduce it?
- How is Qarnot helping?
- **A few technical challenges**

1 - Security

Approche distribuée : sécurité hardware ?
=> kill switch
=> données client en RAM

Sécurité Software
=> TPM
=> config kernel de l’OS (module USB désactivés etc.)
“zero trust”

2 - Parallel computation

Communicate and coordinate nodes effectively

To do this, we use the MPI (message passing interface) standard which defines an interface to communicate between nodes.

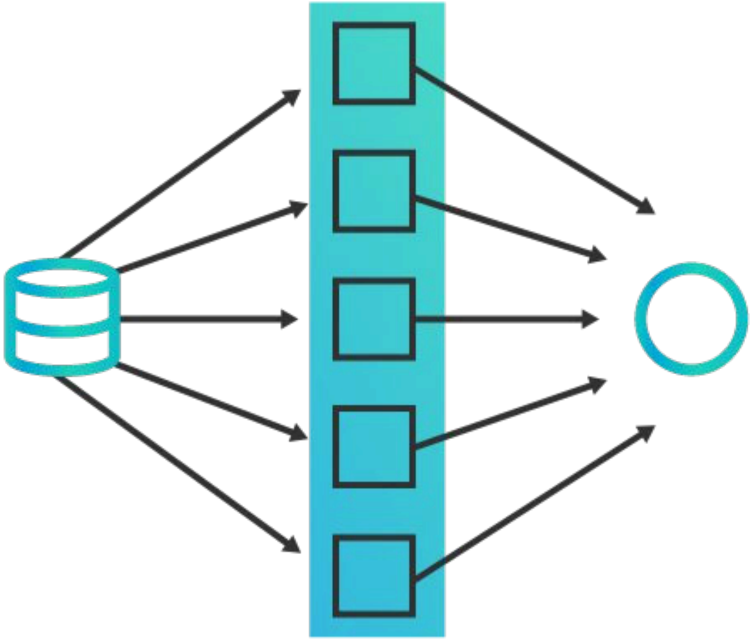

Speedup

$$\text{Speed up} = \frac{\text{Time single-process}}{\text{Time multi-process}}$$

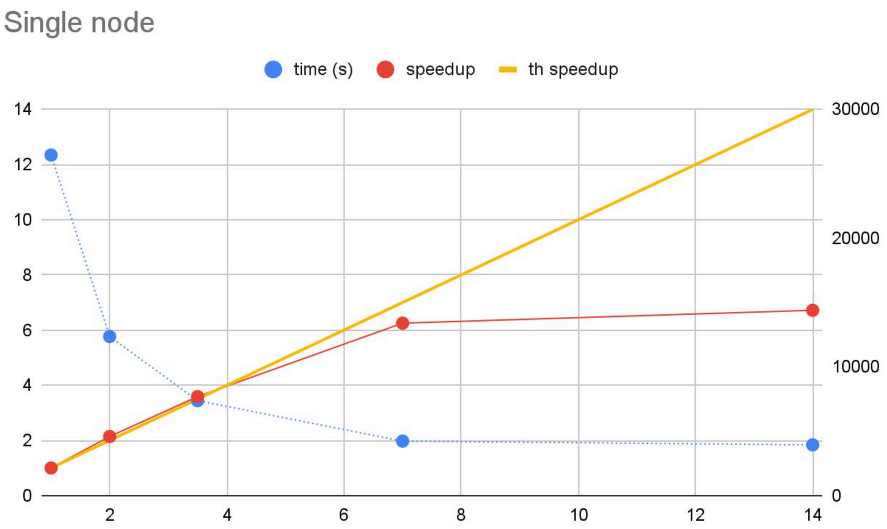

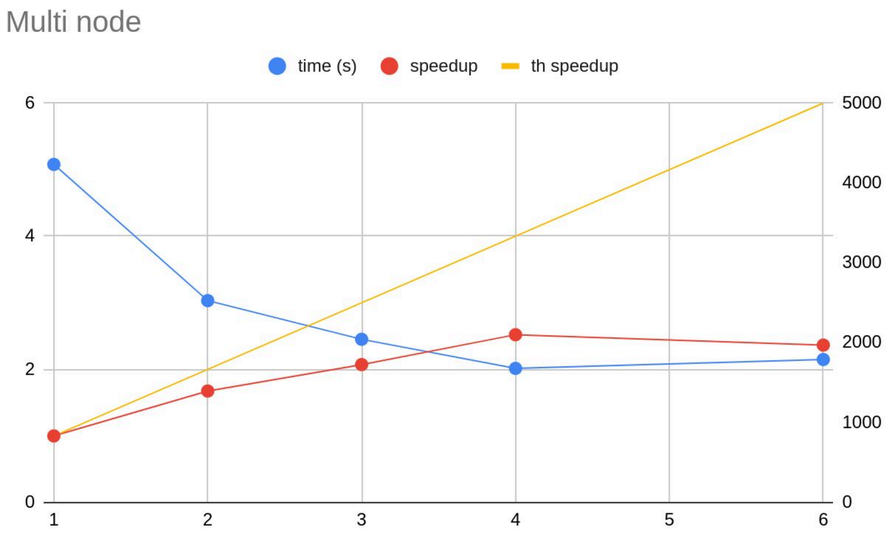

RDMA, Infiniband and RoCE

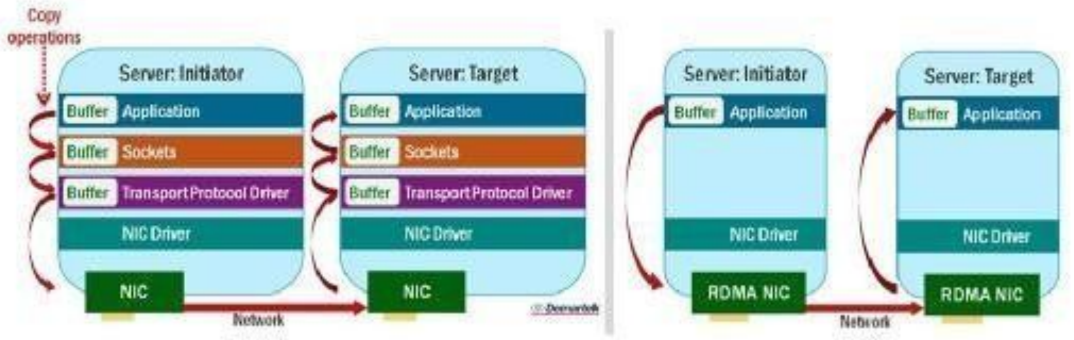

RDMA (Remote Direct Memory Access) :

Direct access to resources through the network without passing through the OS layer and without intermediate copy (zero-copy network).

Infiniband :

Specification of a protocol and network buses that natively support RDMA

RoCE (RDMA over Converged Ethernet) :

Protocol enabling RDMA by encapsulating Infiniband packets in UDP/IP frames (RoCEv2)

RoCE recipe

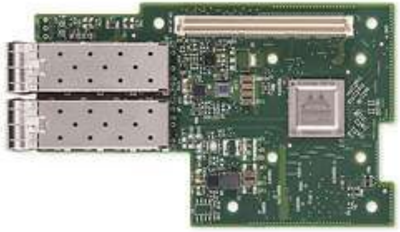

NIC handling RoCE

OCP ACCEPTED

switch handling RoCE

```
mpirun -np 112 [...] \
      -x UCX_NET_DEVICES=mlx5_0:1 \
      [...]
```

option MPI pour utiliser RoCE

```
0> 2022-04-05 11:53:22 || 'executing
0> # OSU MPI Latency Test v5.9
0> # Size          Latency (us)
0> 0                     14.42
0> 1                     13.19
0> 2                     13.19
0> 4                     13.15
0> 8                     13.18
```

Latency divided by 9 !

```
root@1c-34-da-7f-9b-b2:~# mpirun -
# OSU MPI Latency Test v5.6.2
# Size          Latency (us)
0                        1.54
1                        1.53
2                        1.52
4                        1.52
8                        1.52
```

Results

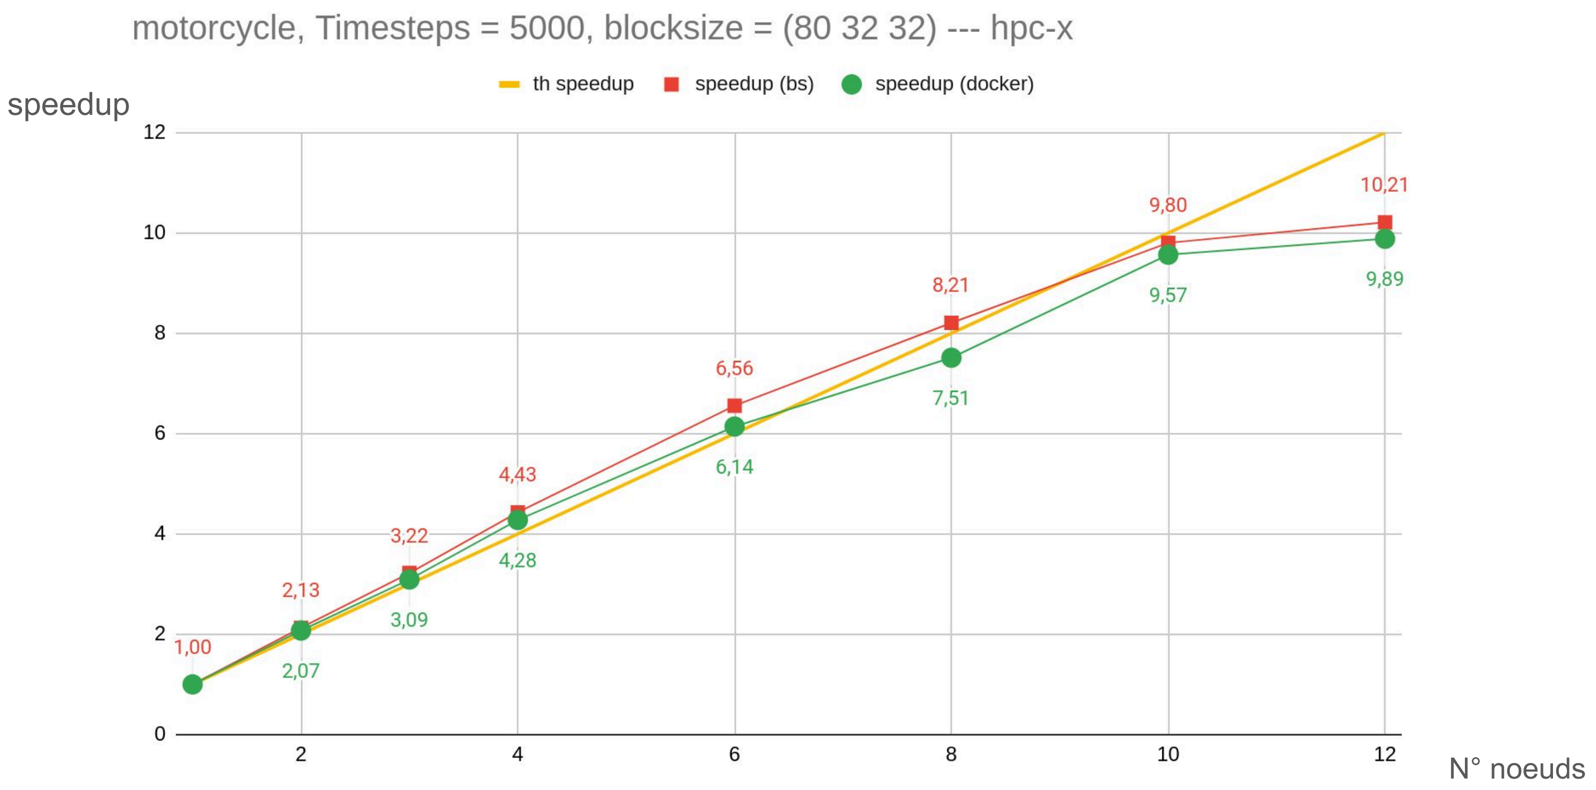

QARNOT

Thank you for your attention!

Questions?

charles.anteunis@qarnot-computing.com

jeremie.spiesser@qarnot-computing.com

jobs@qarnot-computing.com