# La frugalité dans le HPC

20 juin 2022



**Grégoire SOLEIL**

QARNOT

qarnot.com

# Sommaire

# Qarnot

60 employees (⅔ of Engineers)
Entreprise créée en 2010
Montrouge, France

# Consommations IT


Switch, Nevada, 531 MW

Google, Arizona, 1.7M gal. water / day

Meta, Sweden, 120 MW

# Mutualiser pour économiser de l'énergie

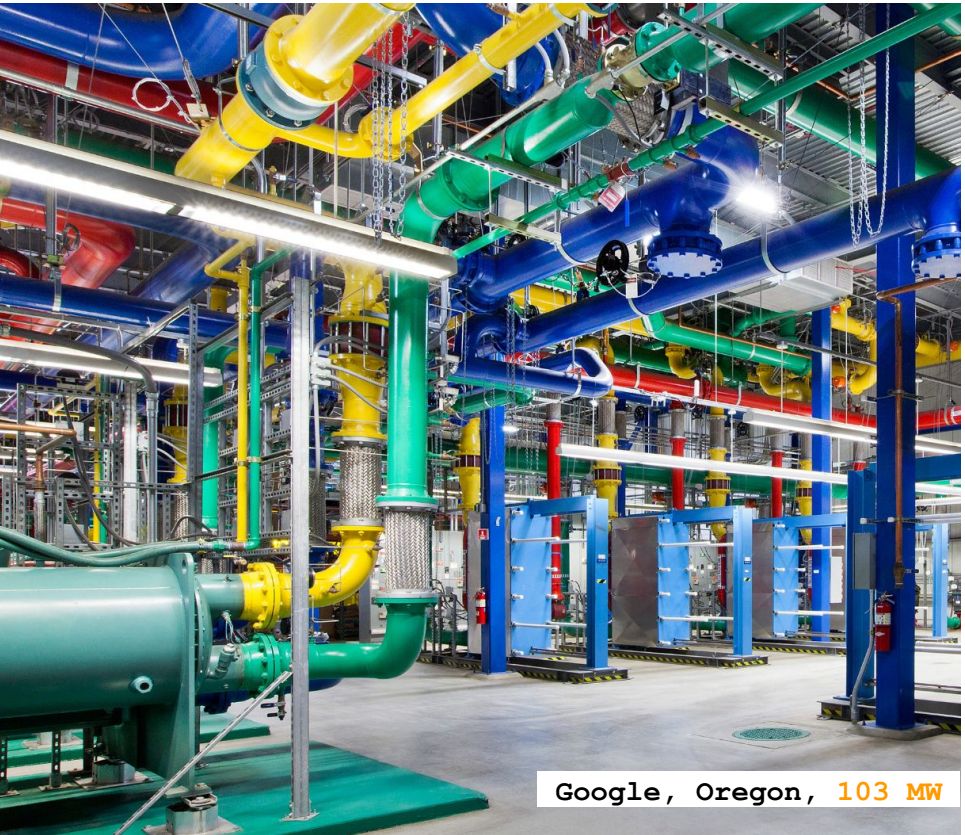

Google, Oregon, 103 MW

District heating plant,Sweden, 30 MW

# Un modèle économique bi-face

En valorisant un déchet dans l'informatique, la chaleur, pour en faire une ressource dans le bâtiment, Qarnot propose une solution d'économie circulaire numérique.

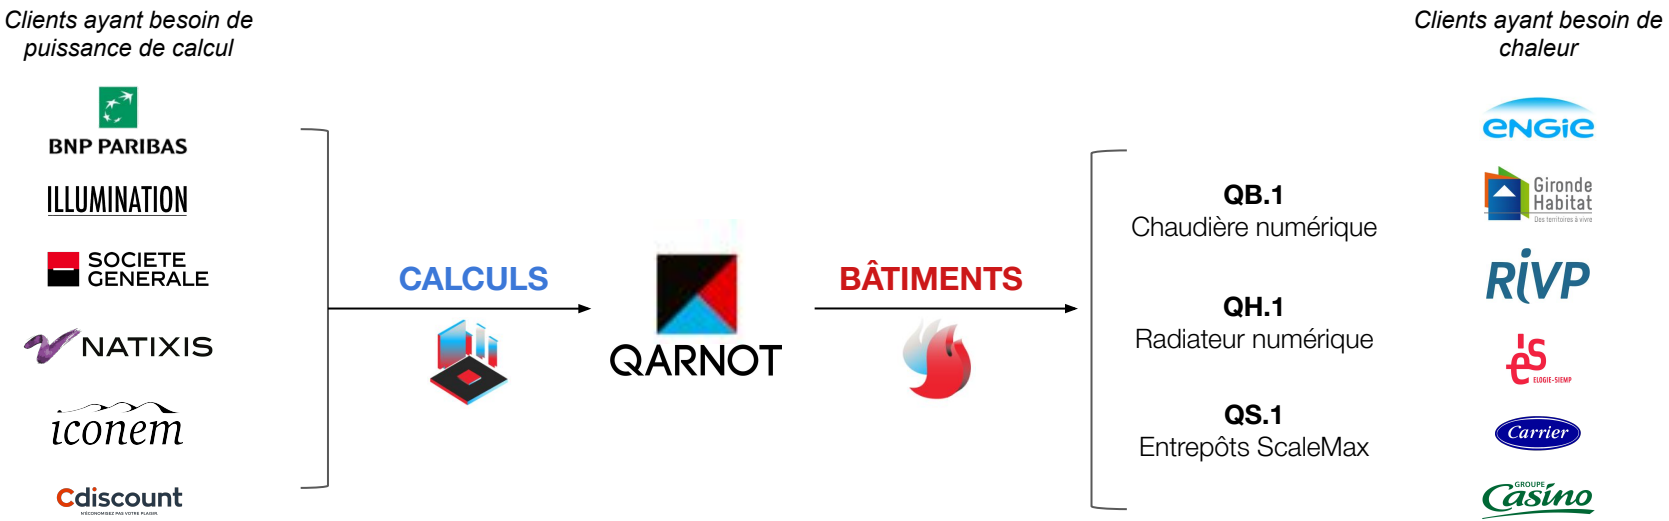

# La chaudière numérique QB · 1

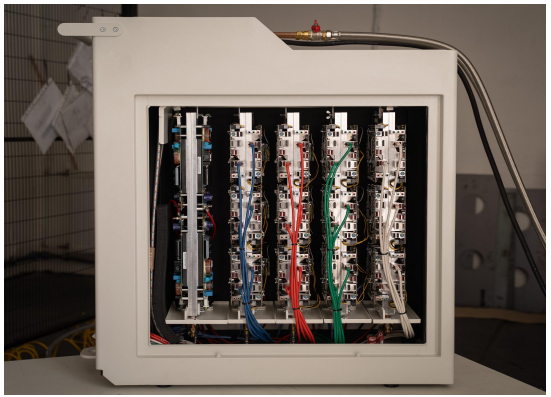

- Chaque module embarque 12 à 24 serveurs informatiques
- Récupération de 96 % de la chaleur
- En sortie, une eau chaude à 65 °C
- Conçue, fabriquée et assemblée en France

# Autres infrastructures

- Radiateurs numériques dans les logements
- Récupération de 100 % de la chaleur
- Fonctionnalités smart building

- Entrepôts ScaleMax (JV avec Casino)
- Récupération de 70 % de la chaleur
- Chaleur écologique vendue à l'entrepôt

# Distributed HPC infrastructure

QARNOT

*Qarnot's Q.Ware software suite optimally drives most bare-metal infrastructures (proprietary & third-party) with a partnership added-in to scale capacity*

Q•ware

**HOUSING**

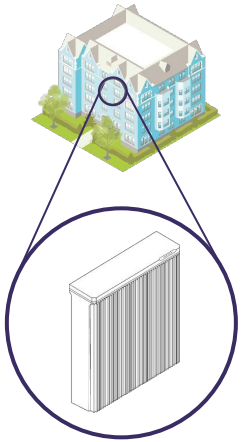

**QH.1**
computing heaters

**RETAIL WAREHOUSES**

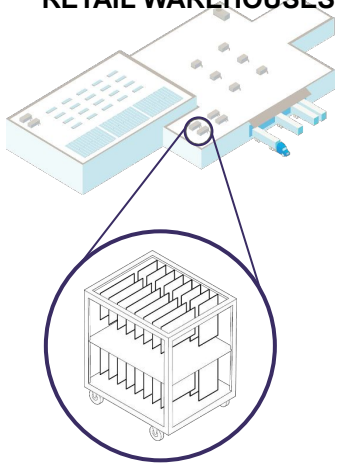

**QS.1**
low-cost clusters

**OFFICES/HOTELS**

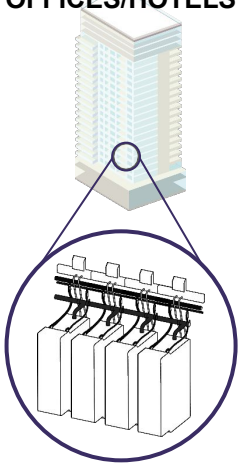

**QB.1**
computing boilers

**PROXIMITY HUBS**

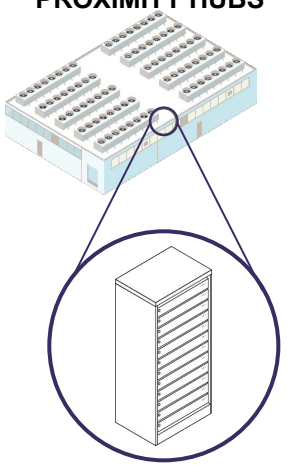

**VANILLA**
server racks

# Distributed HPC infrastructure

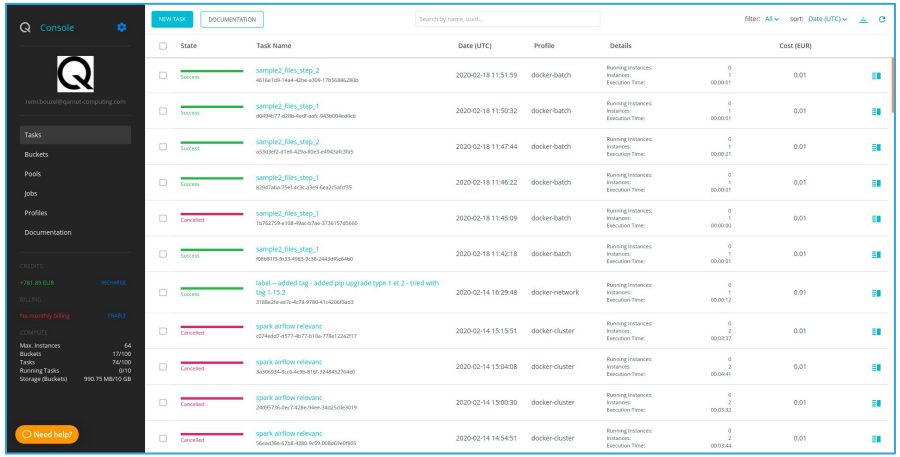


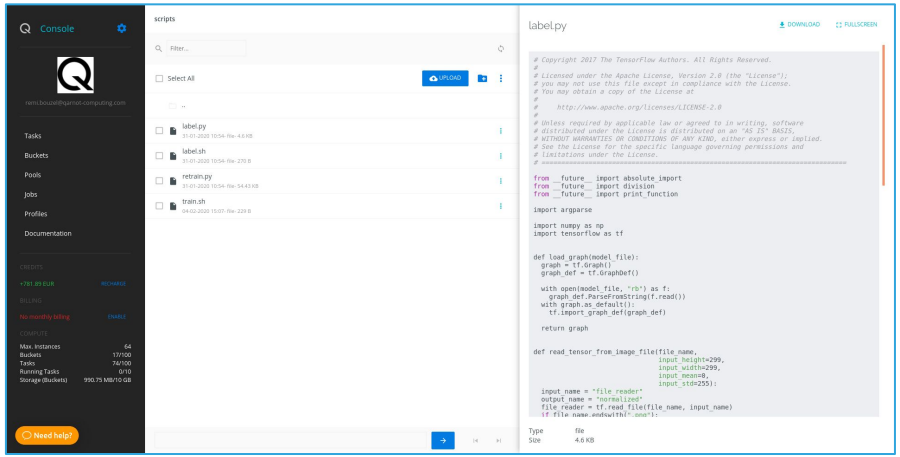


- Le Q.ware, un “platform-as-a-service”
  - Conception pour des secteurs verticaux avec des applications à forte intensité de calcul : animation 3D, banque, fonds spéculatifs...
- Applications flexibles
  - Les utilisateurs peuvent utiliser des conteneurs Docker publics ou personnalisés, ou des applications optimisées.
- Caractéristiques
  - API simple et intuitive, spécialisée pour le calcul distribué, SDK Python puissant, console user-friendly pour surveiller et gérer les tâches, les ressources et la facturation.

# Notre architecture

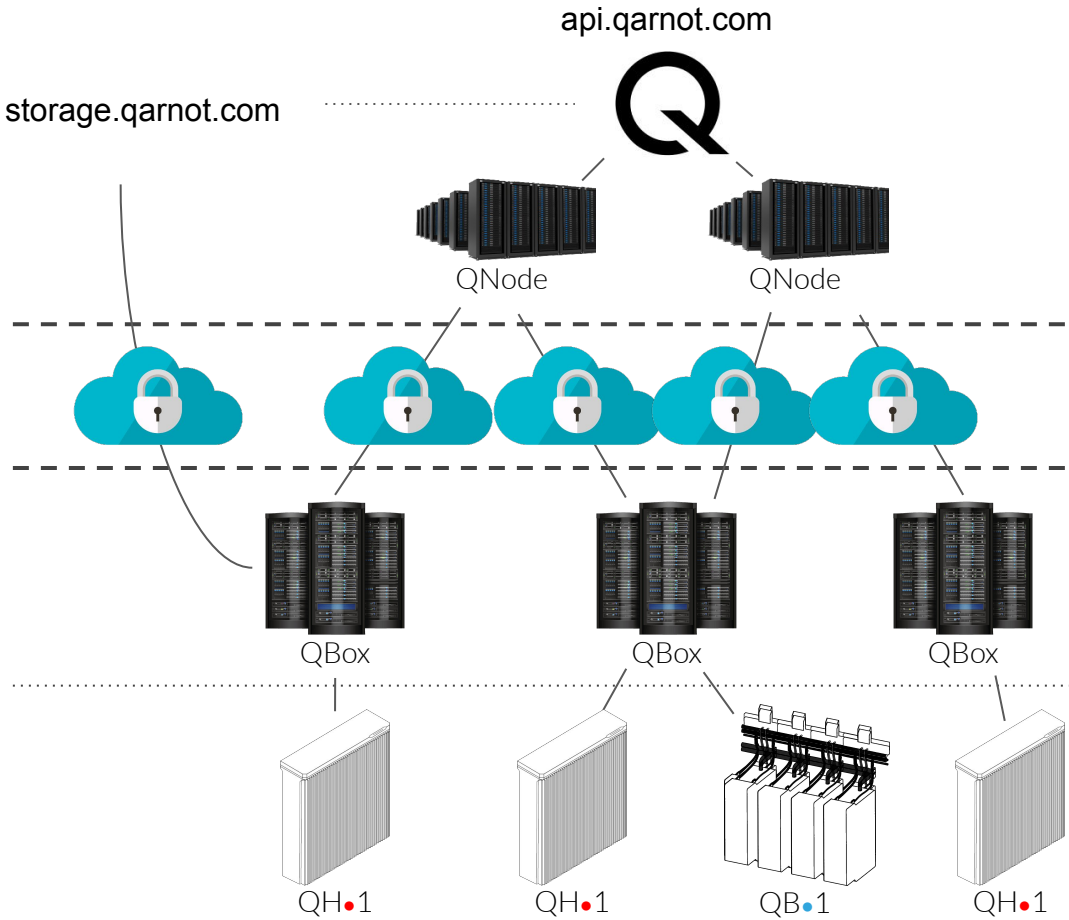


Datacenter

Internet

Dedicated room

Home and offices

Heated room

# Réduire les coûts et l'empreinte carbone


## Tarifs compétitifs

*Jusqu'à 4 fois moins chers que d'autres fournisseurs de Cloud*

## Haute sécurité

*Infrastructure bare-metal, chiffrement des données, accès sécurisé aux sites*

*Certifications ISO-27001, HDS*

## Calcul écologique

*Electricité verte et valorisation de la chaleur*

## Cloud de confiance

*Sites en France et en Europe*

## Q.WARE

*Couche applicative propriétaire et intuitive*

## Offres scalables

*De la start-up aux grands groupes, pour les besoins de calculs ponctuels ou récurrents*

## Carbon Facts

| | |
|---|---|
| Name | QB·1 3700X |
| Duration | 50000 h |
| Saved footprint Carbon (éq. $CO_2$) | 2.50 T |
| Reduced Emissions | 88,81 % |

**Energy**

| | |
|---|---|
| Total consumed | 5245 kWh |
| Average power | 105 W |
| Reused energy | 4786 kWh |
| Energy Reuse Effectiveness (ERE) | 0.09 |
| Avg Power Usage Effectivness (PUE) | 1.002 |
| Avg Energy Reuse Factor (ERF) | 0.91 |

**Carbon**

| | |
|---|---|
| Total carbon footprint (éq. $CO_2$) | 0.31 T |
| Carbon emission avoided by reused heat (éq. $CO_2$) | 1.09 T |
| Saved carbon footprint (éq. $CO_2$) | 2.50 T |
| Emmision reduced | 88.81 % |

# Cas d'usage HPC

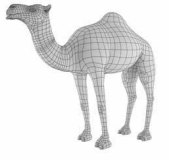

3D ANIMATION

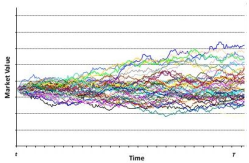


FINANCES

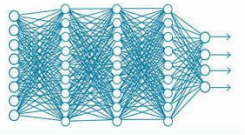

MACHINE LEARNING

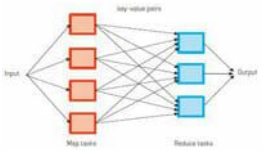


BIG DATA

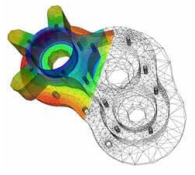

MECHANICS

ELECTRONICS

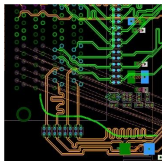

OPERATIONS RESEARCH

BIOTECH

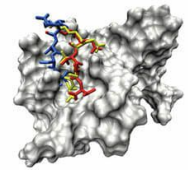

FLUID DYNAMICS

WEATHER FORECAST

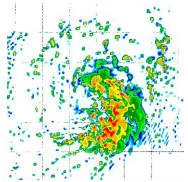

# Sommaire

# Qu’est ce que l’impact environnemental ?

**Les ressources**

- Epuisement des ressources minérales
- Epuisement des ressources fossiles
- Utilisation des sols

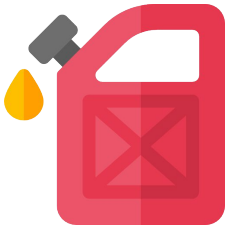

**La santé humaine**

- Epuisement de la couche d’ozone
- Toxicité humaine
- Émissions de particules
- Radiations ionisantes
- Formation d’ozone photochimique

**Les ecosystèmes**

- Acidification
- Eutrophisation terrestre
- Eutrophisation d’eau douce
- Eutrophisation marine
- Ecotoxicité aquatique

**Le changement climatique**

- Dérèglement climatique (potentiel de réchauffement global)

**L’eau**

- Utilisation de l’eau

# L'impact du numérique mondial

| CONSOMMATION D'ÉNERGIE PRIMAIRE (EP) : | EMISSIONS DE GAZ À EFFET DE SERRE (GES) : | CONSOMMATION D'EAU (EAU) : | CONSOMMATION D'ÉLECTRICITÉ (ELEC.)* : |
|---|---|---|---|
| 4,2 % | 3,8 % | 0,2 % | 5,5 % |

2 à 3 fois l'impact de la France

| ÉLECTRICITÉ | ÉNERGIE | GES |
|---|---|---|
| x 2,7 | x 2,9 | x 3,1 |
| entre 2010 et 2025 | entre 2010 et 2025 | entre 2010 et 2025 |

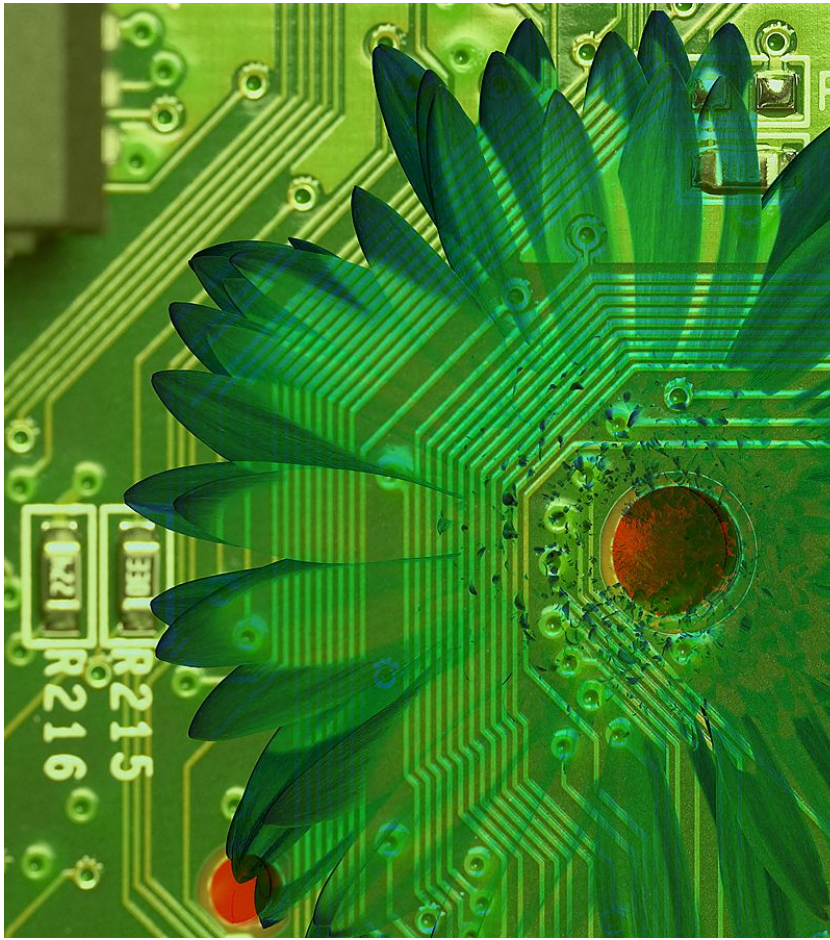

Empreinte environnementale du numérique mondial, GreenIT, 2019

# L’impact du numérique

**37%** des émissions de GES soutenables

**35%** de la consommation de matière premières soutenable

**32%** de la consommation d’eau douce soutenable

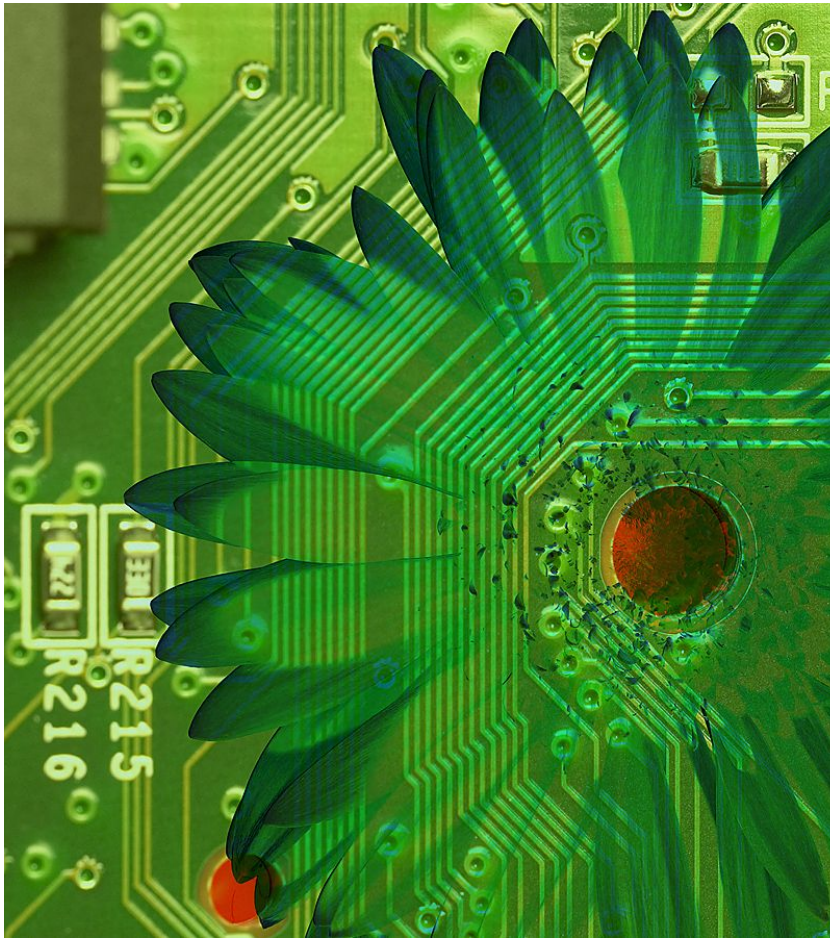

Empreinte environnementale du numérique mondial, GreenIT, 2019

# La mesure d'impact chez Qarnot

1 kWh IT

**Qarnot Boiler**

**Reference Scenario**

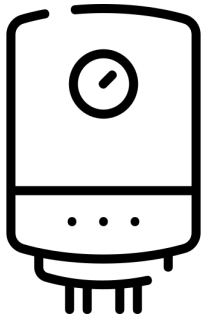

0.95 kWh gaz

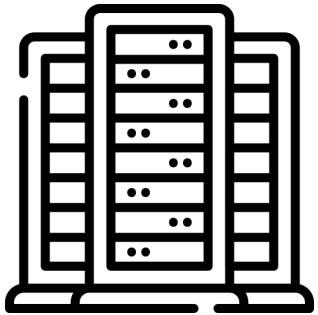

1 kWh IT

0,2 kWh cooling

# Une économie d'énergie

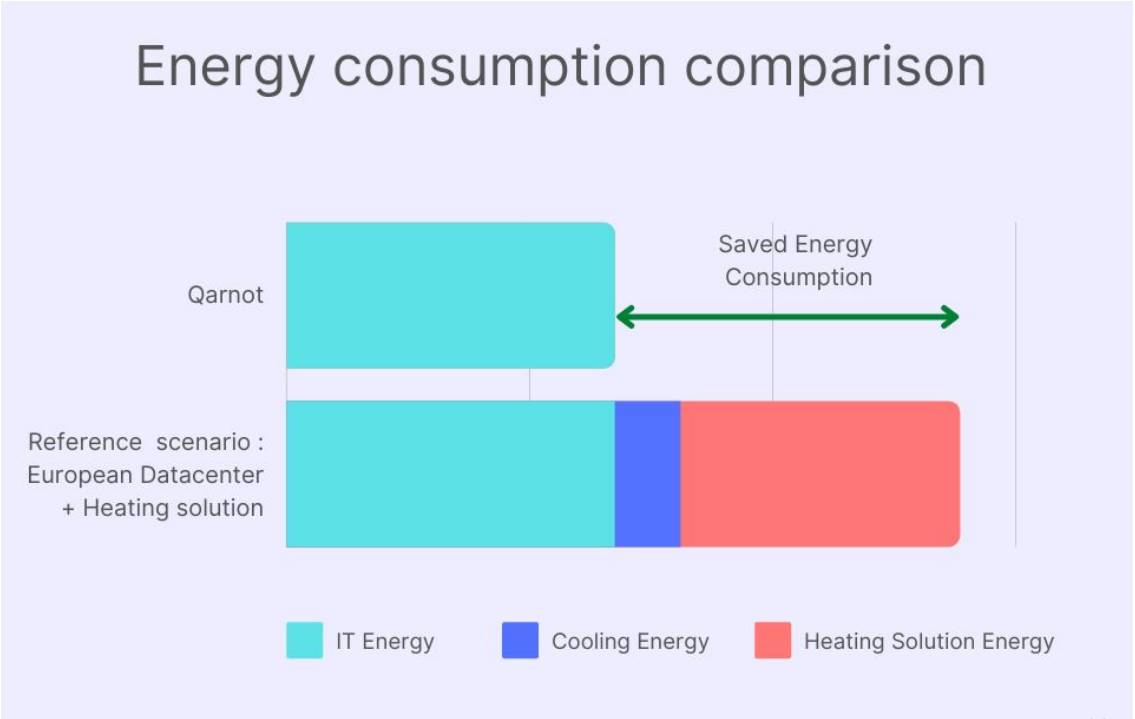

# Une économie de carbone

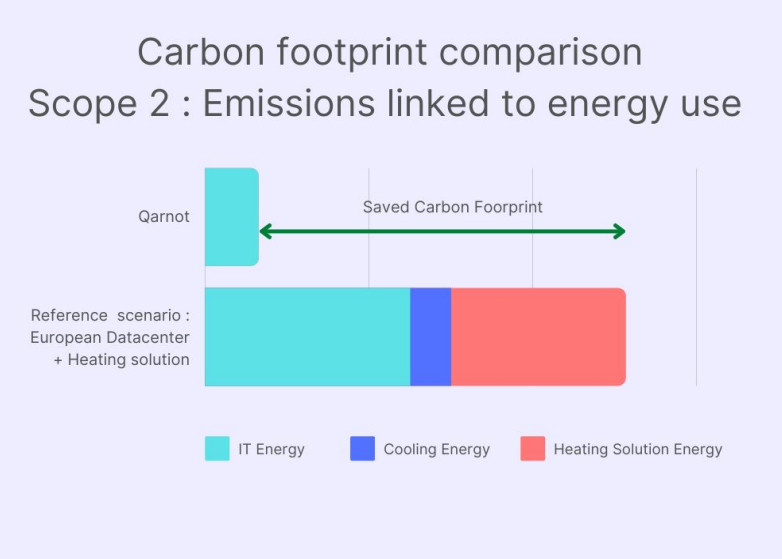


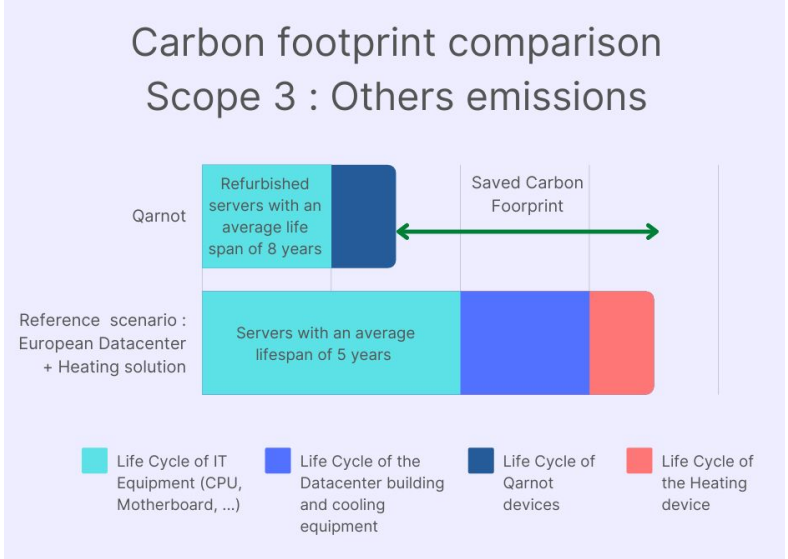

# Etiquette Carbon Facts

- Granularité à l'échelle du calcul
- Prendre en compte les 3 scopes
- ACV globale du calcul
- Valoriser les émissions évitées par la réutilisation de chaleur

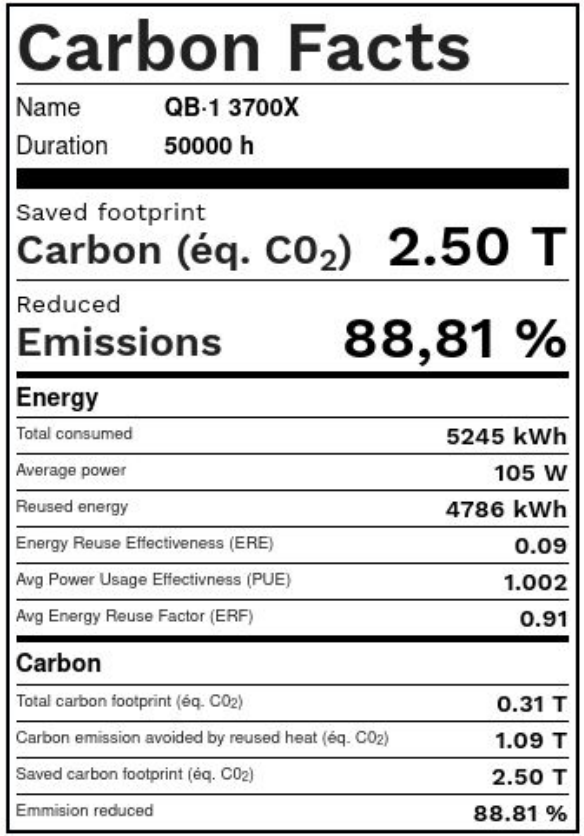

## Carbon Facts

Name QB·1 3700X

Duration 50000 h

Saved footprint

**Carbon (éq. $CO_2$) 2.50 T**

Reduced

**Emissions 88,81 %**

**Energy**

| | |
|---|---|
| Total consumed | 5245 kWh |
| Average power | 105 W |
| Reused energy | 4786 kWh |
| Energy Reuse Effectiveness (ERE) | 0.09 |
| Avg Power Usage Effectivness (PUE) | 1.002 |
| Avg Energy Reuse Factor (ERF) | 0.91 |

**Carbon**

| | |
|---|---|
| Total carbon footprint (éq. $CO_2$) | 0.31 T |
| Carbon emission avoided by reused heat (éq. $CO_2$) | 1.09 T |
| Saved carbon footprint (éq. $CO_2$) | 2.50 T |
| Emmision reduced | 88.81 % |

# Les autres impacts

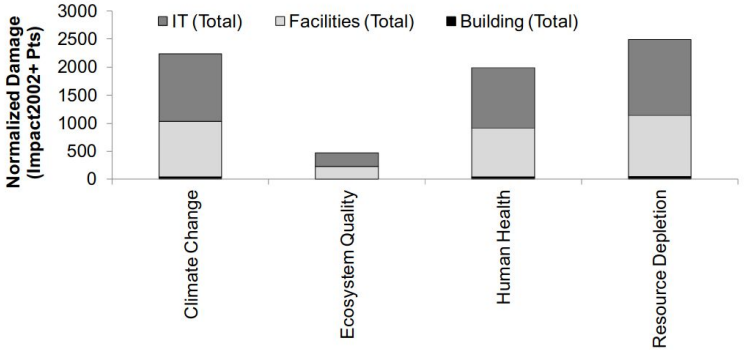


**Table 5(a). Relative index mapping influence of different dependent parameters on midpoint impact categories.**

| Parameter | Climate Change | Ecosystem Quality | | | | | Human Health | | | | | Resource Depletion | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | Climate Change | Aquatic Ecotoxicity | Land Occupation | Terrestrial Acidification & Nutrification | Aquatic Acidification* | Aquatic Eutrophication* | Human Toxicity | Ionising Radiation | Ozone Layer Depletion | Photochemical Oxidation | Respiratory Effects (Inorganics) | Mineral Extraction | Non-renewable Energy |
| Electricity Impact Factor (per kWh) | 0.83 | 0.60 | 0.66 | 0.08 | 0.83 | 0.65 | 0.77 | 0.82 | 0.78 | 0.74 | 0.81 | 0.07 | 0.83 |
| Buildling Construction & Decomissioning Impact Factor (per $ft^2$) | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| Power Delivery Infrastructure Manufacturing Impact Factor (per kW) | 0.01 | 0.23 | 0.13 | 0.95 | 0.01 | 0.02 | 0.04 | 0.01 | 0.01 | 0.02 | 0.02 | 0.47 | 0.01 |
| Cooling Equipment Manufacturing Impact Factor (per ton) | 0.01 | 0.23 | 0.13 | 0.95 | 0.01 | 0.02 | 0.04 | 0.01 | 0.01 | 0.02 | 0.02 | 0.47 | 0.01 |
| Compute Equipment Manufacturing Impact Factor (per W) | 0.03 | 0.08 | 0.10 | 0.00 | 0.02 | 0.21 | 0.07 | 0.04 | 0.08 | 0.12 | 0.04 | 0.47 | 0.03 |
| Storage Equipment Manufacturing Impact Factor (per W) | 0.03 | 0.09 | 0.11 | 0.00 | 0.03 | 0.23 | 0.07 | 0.04 | 0.09 | 0.13 | 0.04 | 0.50 | 0.03 |
| Networking Equipment Manufacturing Impact Factor (per W) | 0.03 | 0.09 | 0.11 | 0.00 | 0.03 | 0.23 | 0.07 | 0.04 | 0.09 | 0.13 | 0.04 | 0.50 | 0.03 |
| IT Equipment lifetime (yrs) | 0.03 | 0.09 | 0.11 | 0.00 | 0.03 | 0.22 | 0.07 | 0.04 | 0.08 | 0.12 | 0.04 | 0.49 | 0.03 |
| Facility Infrastructure lifetime (yrs) | 0.01 | 0.23 | 0.13 | 0.95 | 0.01 | 0.02 | 0.04 | 0.01 | 0.01 | 0.02 | 0.02 | 0.47 | 0.01 |
| Building Infrastructure Lifetime (yrs) | 0.02 | 0.01 | 0.02 | 0.00 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 | 0.00 | 0.02 |
| Fraction of IT equipment recycled (%) | 0.03 | 0.09 | 0.11 | 0.00 | 0.03 | 0.23 | 0.07 | 0.04 | 0.09 | 0.13 | 0.04 | 0.50 | 0.03 |
| Fraction of facility infrastructure recycled (%) | 0.01 | 0.23 | 0.13 | 0.95 | 0.01 | 0.02 | 0.04 | 0.01 | 0.01 | 0.02 | 0.02 | 0.47 | 0.01 |
| Fraction of building construction materials recycled (%) | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| Fraction of storage equipment in data center (%) | 0.03 | 0.09 | 0.11 | 0.00 | 0.03 | 0.23 | 0.07 | 0.04 | 0.09 | 0.13 | 0.04 | 0.50 | 0.03 |
| Fraction of networking equipment in data center (%) | 0.03 | 0.09 | 0.11 | 0.00 | 0.03 | 0.23 | 0.07 | 0.04 | 0.09 | 0.13 | 0.04 | 0.50 | 0.03 |
| Fraction of compute equipment in data center (%) | 0.03 | 0.09 | 0.11 | 0.00 | 0.03 | 0.23 | 0.07 | 0.04 | 0.09 | 0.13 | 0.04 | 0.50 | 0.03 |
| Level of Standby (Operational) Redundancy in facility (N + k) | 0.47 | 0.56 | 0.50 | 1.00 | 0.47 | 0.39 | 0.47 | 0.46 | 0.44 | 0.43 | 0.47 | 0.51 | 0.47 |
| Power Delivery Efficiency | 0.33 | 0.24 | 0.26 | 0.03 | 0.33 | 0.26 | 0.30 | 0.33 | 0.31 | 0.29 | 0.32 | 0.03 | 0.33 |
| Design Density of Active Floorspace ($W/ft^2$) | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 |
| Miscellaneous Runtime Building loads (W) | 0.02 | 0.01 | 0.02 | 0.00 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 | 0.02 | 0.00 | 0.02 |
| Fraction of Outside Air ("Free Cooling") usage (%) | 0.46 | 0.33 | 0.36 | 0.05 | 0.46 | 0.36 | 0.42 | 0.45 | 0.43 | 0.41 | 0.45 | 0.04 | 0.46 |
| Relative impact of changing facility supply air temperature (% per deg C) | 0.46 | 0.33 | 0.36 | 0.05 | 0.46 | 0.36 | 0.42 | 0.45 | 0.43 | 0.41 | 0.45 | 0.04 | 0.46 |
| Peak IT Equipment Power (W) | 0.56 | 0.68 | 0.65 | 0.97 | 0.56 | 0.65 | 0.59 | 0.56 | 0.58 | 0.61 | 0.57 | 0.98 | 0.56 |
| Idle IT Equipment Power (W) | 0.39 | 0.28 | 0.31 | 0.04 | 0.39 | 0.31 | 0.36 | 0.38 | 0.36 | 0.34 | 0.38 | 0.03 | 0.39 |
| Average IT equipment utilization (%) | 0.36 | 0.26 | 0.28 | 0.04 | 0.36 | 0.28 | 0.33 | 0.35 | 0.34 | 0.32 | 0.35 | 0.03 | 0.36 |
| Average Data Center Uptime (%) | 0.28 | 0.20 | 0.22 | 0.03 | 0.28 | 0.22 | 0.26 | 0.28 | 0.27 | 0.25 | 0.27 | 0.02 | 0.28 |
| Ratio of flow work to thermodynamic work in cooling infra. | 0.37 | 0.27 | 0.29 | 0.04 | 0.37 | 0.29 | 0.34 | 0.36 | 0.34 | 0.33 | 0.36 | 0.03 | 0.37 |
| Coefficient of Performance of chiller | 0.34 | 0.24 | 0.27 | 0.03 | 0.33 | 0.26 | 0.31 | 0.33 | 0.31 | 0.30 | 0.33 | 0.03 | 0.34 |

A. J. Shah, Y. Chen and C. E. Bash, "Sources of variability in data center lifecycle assessment," 2012 IEEE International Symposium on Sustainable Systems and Technology (ISSST), 2012, pp. 1-6, doi: 10.1109/ISSST.2012.6227975.

# Les autres impacts

De la vapeur d'eau s'élève au-dessus des tours de refroidissement du centre de données de The Dallas, en Oregon. Ces panaches de vapeur d'eau créent une brume paisible à la tombée de la nuit.

Transfert d'impact :

- Diminution du PUE
- Augmentation du WUE

WUE = 0.49, soit 0.49L par kwh IT :

0.49*2260KJ = 0,308 kWh/kWh IT

PUE = 1.2 + 0.308 = 1.508



https://www.google.com/about/datacenters/locations/berkeley-county/

# Sommaire

# Du HPC frugal ?

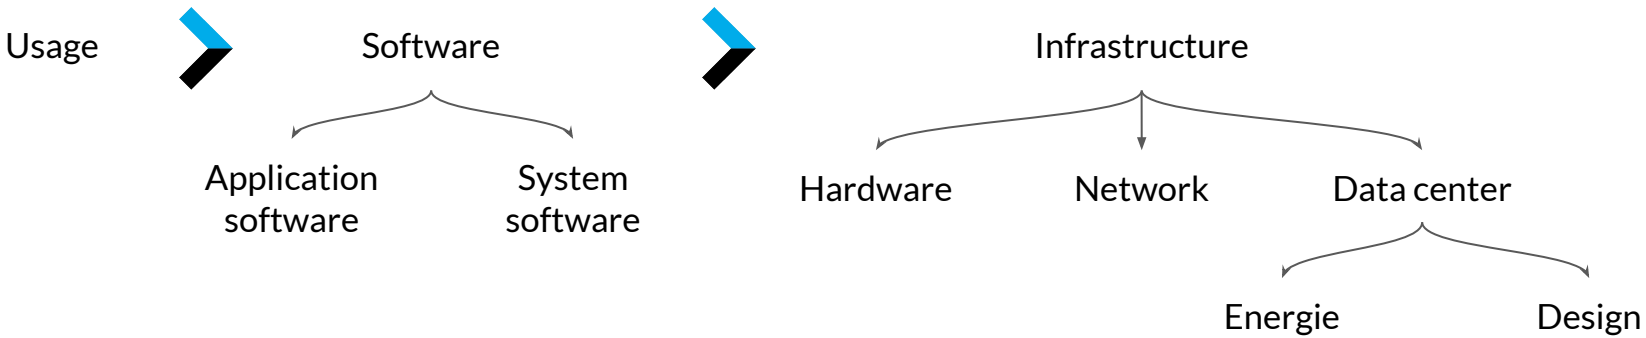

# Du HPC frugal ?

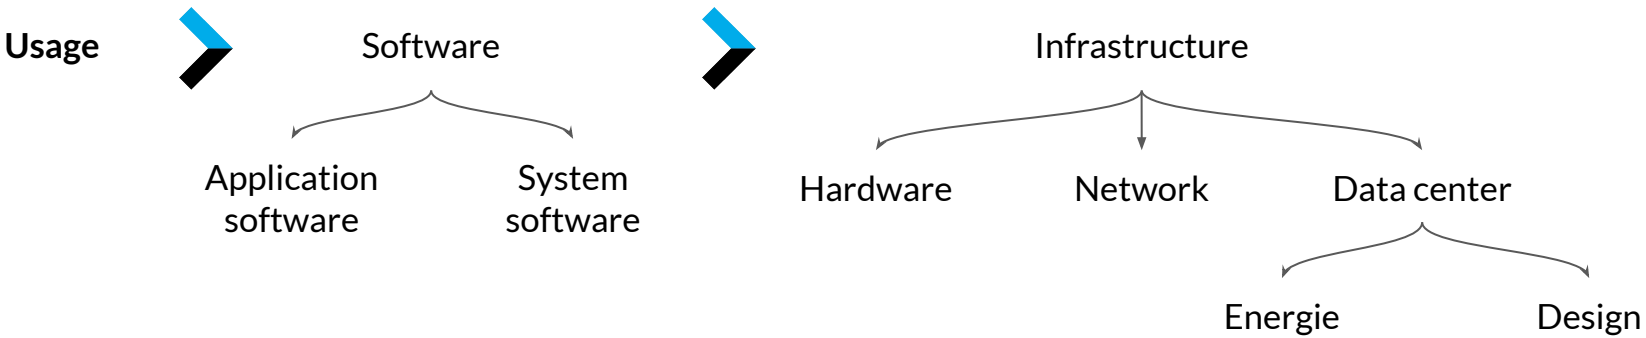

# L'Usage : l'exemple de l'IA

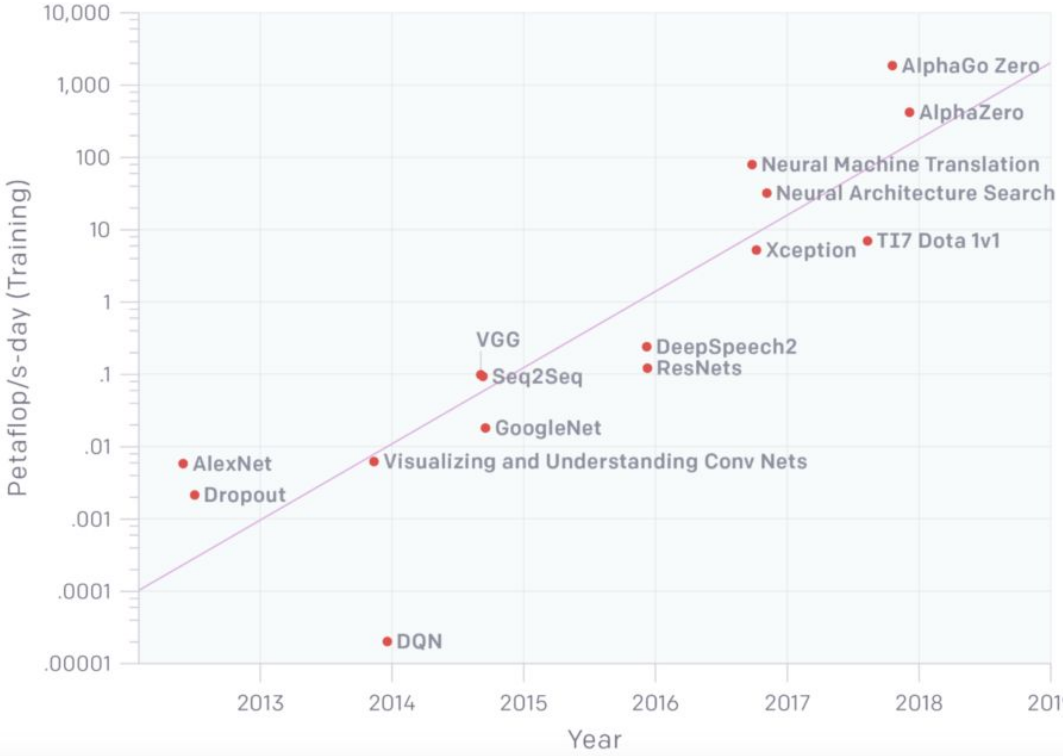


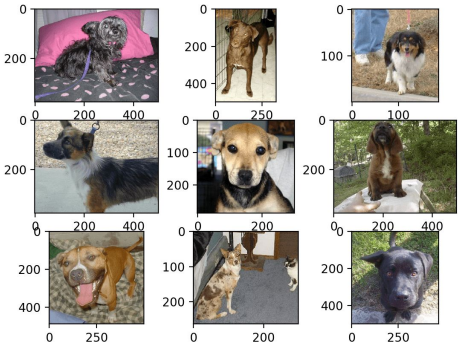

# Des usages durables

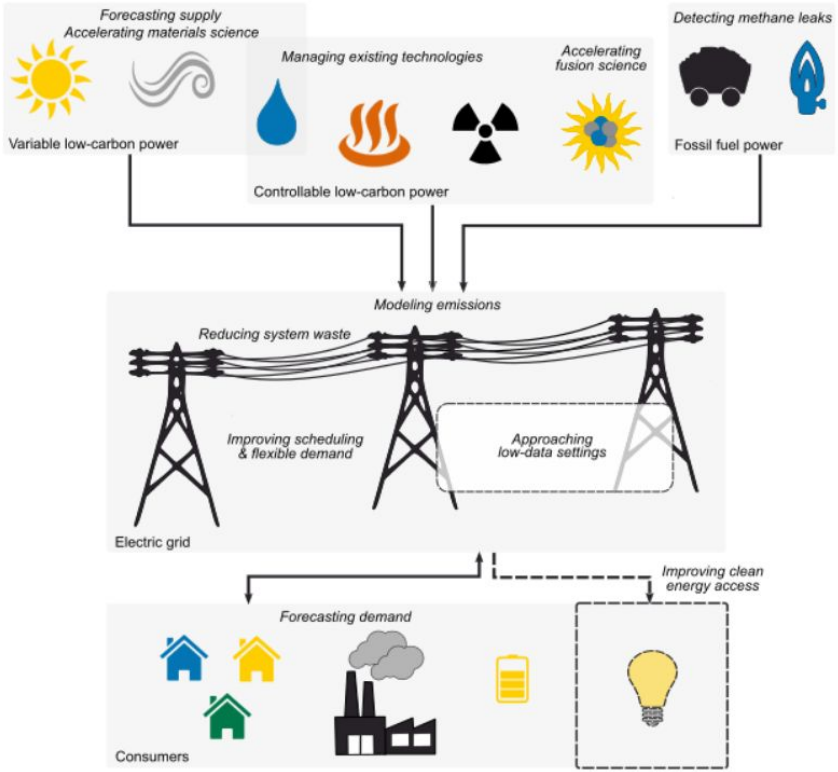


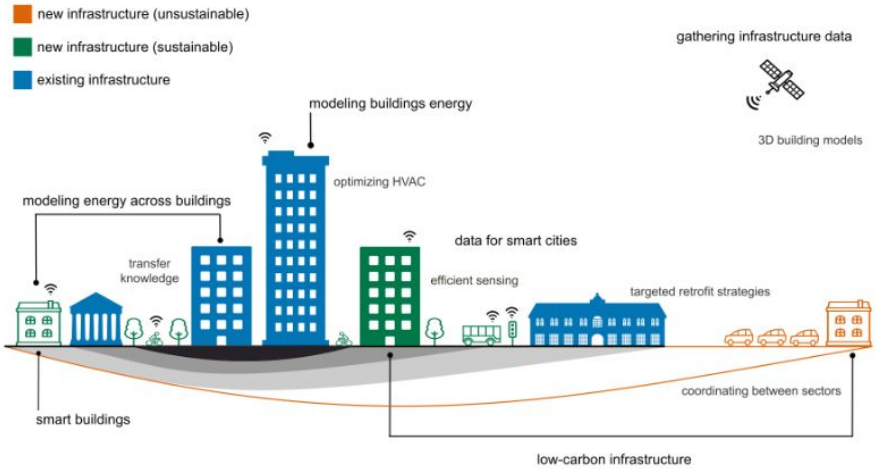


Figure 1: Selected opportunities to reduce GHG emissions from electricity systems using machine learning.

Rolnick, David, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, et al. « Tackling Climate Change with Machine Learning ». arXiv, 5 novembre 2019. https://doi.org/10.48550/arXiv.1906.05433.

# Green AI vs Red AI

Faire de l'efficacité des modèles d'IA un critère important

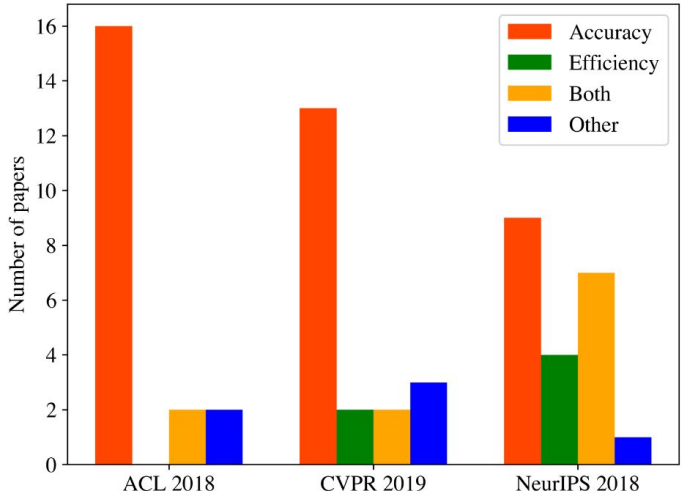


Des outils :

https://github.com/Swall0w/torchstat

https://github.com/Lyken17/pytorch-OpCounter

Schwartz, Roy, Jesse Dodge, Noah A. Smith, et Oren Etzioni.
« Green AI ». arXiv, 13 août 2019. https://doi.org/10.48550/arXiv.1907.10597.

# Du HPC frugal ?

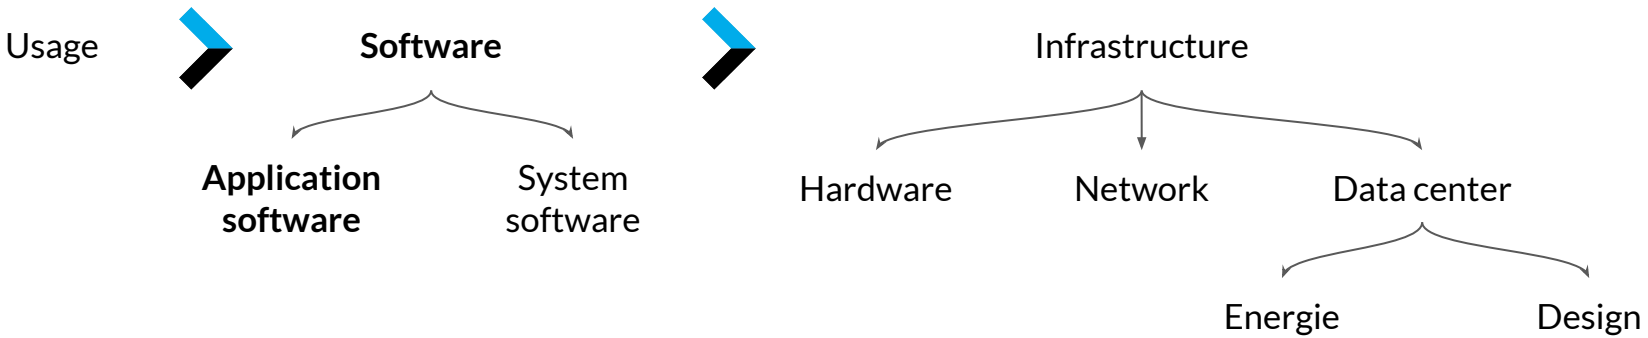

# Machine-learning Denoising

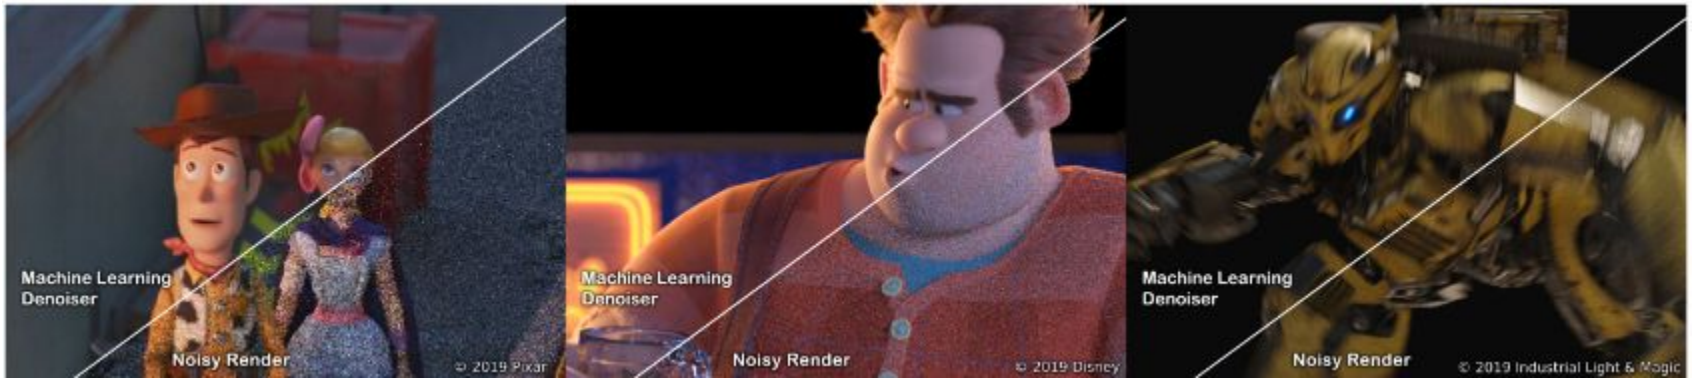


Dahlberg, Henrik, David Adler, et Jeremy Newlin. « Machine-learning denoising in feature film production ». In *ACM SIGGRAPH 2019 Talks*, 1-2. SIGGRAPH '19. New York, NY, USA: Association for Computing Machinery, 2019. https://doi.org/10.1145/3306307.3328150.

# Reduced order model

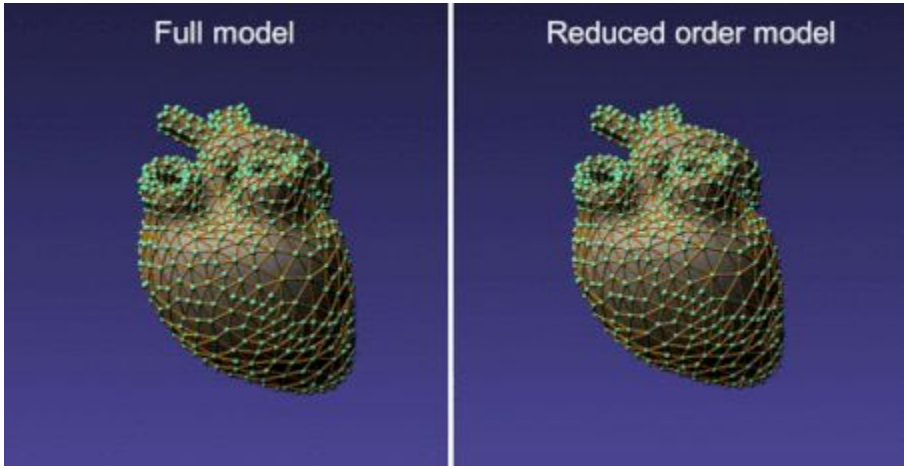

# Importance de la compilation

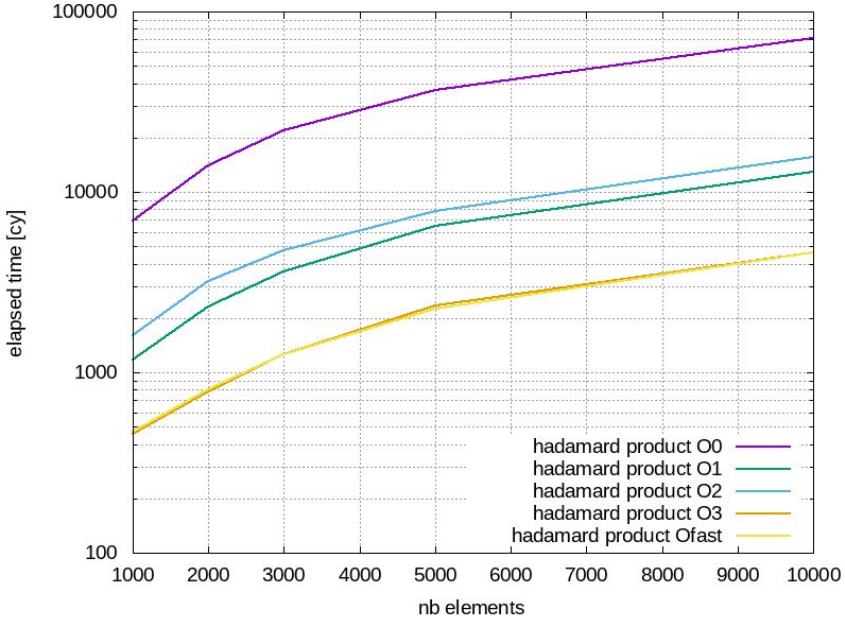


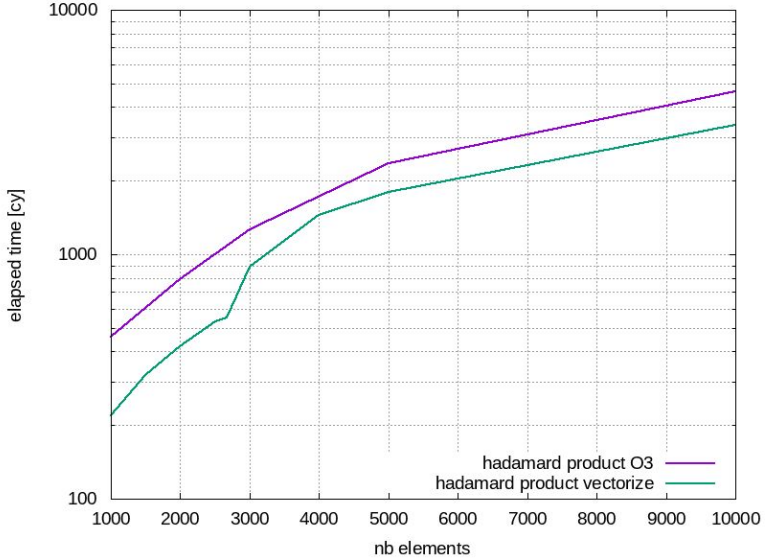


« How to optimize computation with HPC ».
https://lappweb.in2p3.fr/~paubert/ASTERICS_HPC/index.html.

# Du HPC frugal ?

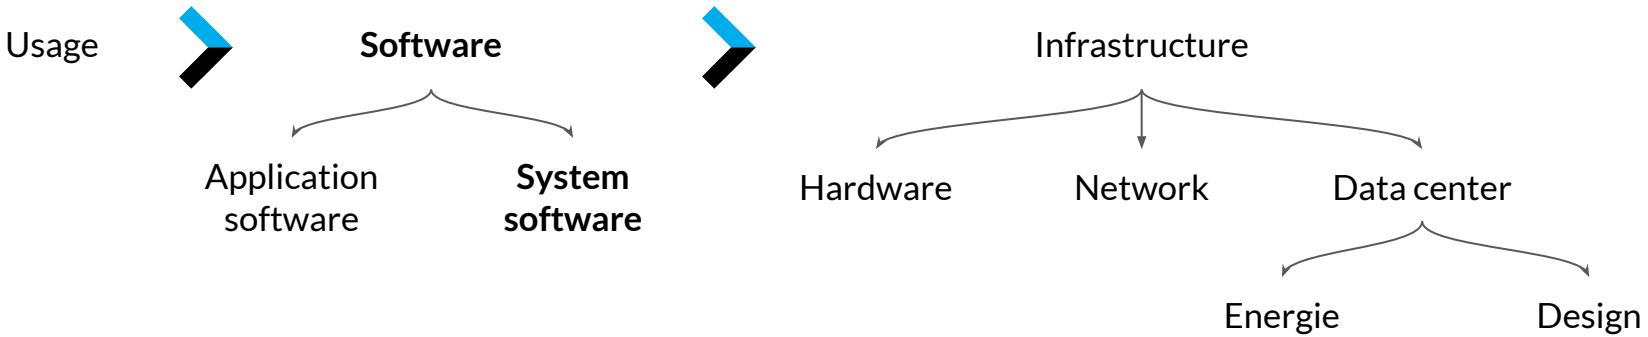

# System software

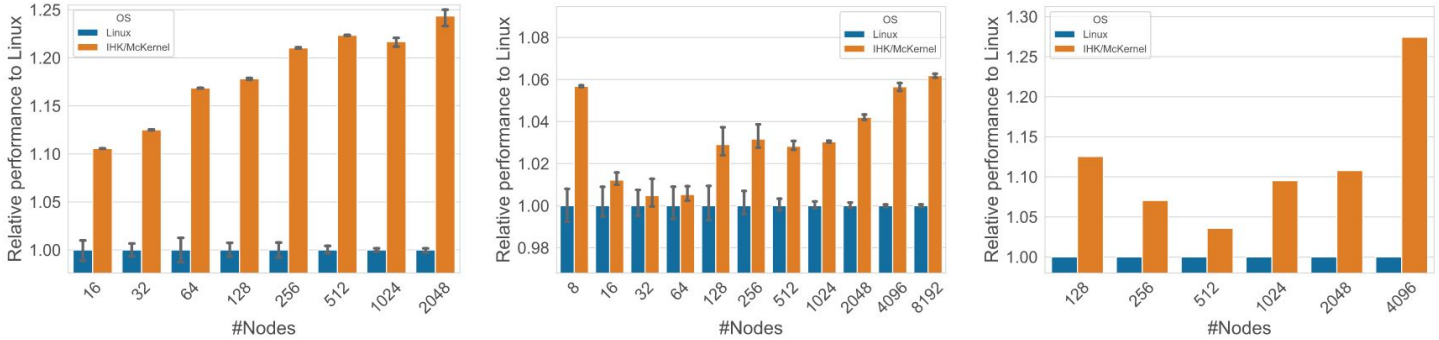


Gerofi, Balazs, Kohei Tarumizu, Lei Zhang, Takayuki Okamoto, Masamichi Takagi, Shinji Sumimoto, et Yutaka Ishikawa. « Linux vs. lightweight multi-kernels for high performance computing: experiences at pre-exascale ». In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 1-13. SC '21. New York, NY, USA: Association for Computing Machinery, 2021. https://doi.org/10.1145/3458817.3476162.

# Du HPC frugal ?

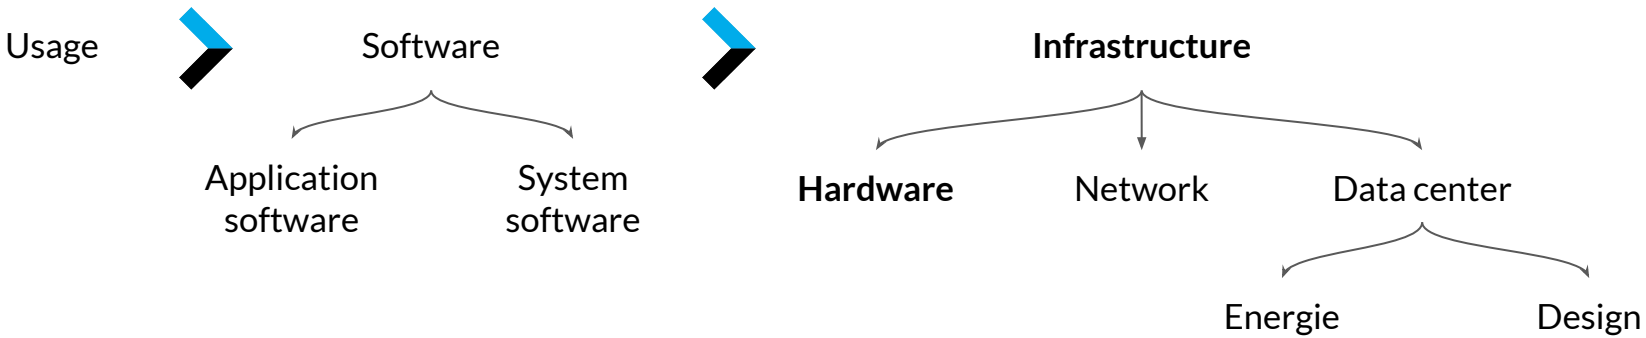

# Infrastructure : Hardware

## L’impact environnemental d’un serveur

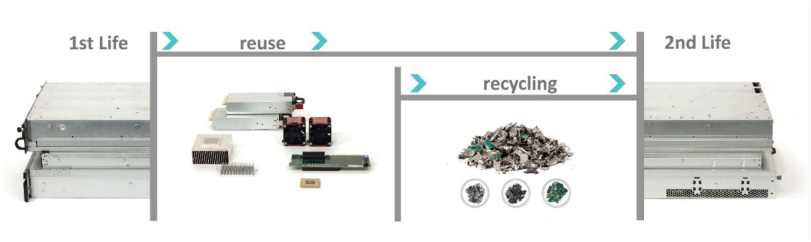


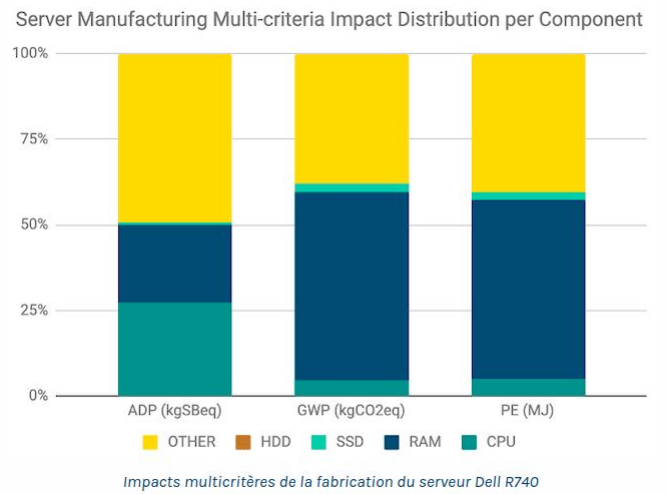


*Impacts multicritères de la fabrication du serveur Dell R740*

*Numérique et environnement : Comment évaluer l’empreinte de la fabrication d’un serveur, au-delà des émissions de gaz à effet de serre ?*, Boavizta, 2021

# Infrastructure : Hardware

**Des outils :**

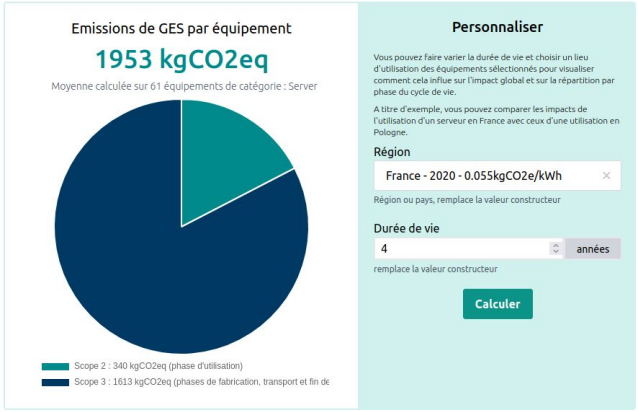


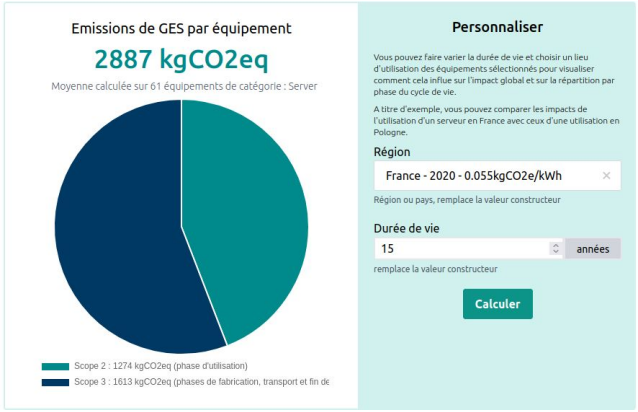


Des actions :

- Augmenter le taux d'utilisation
- Augmenter la durée de vie

Référentiel de données Boavizta :
https://github.com/Boavizta/environmental-footprint-data

API Boavizta :
https://doc.api.boavizta.org/

# Du HPC frugal ?

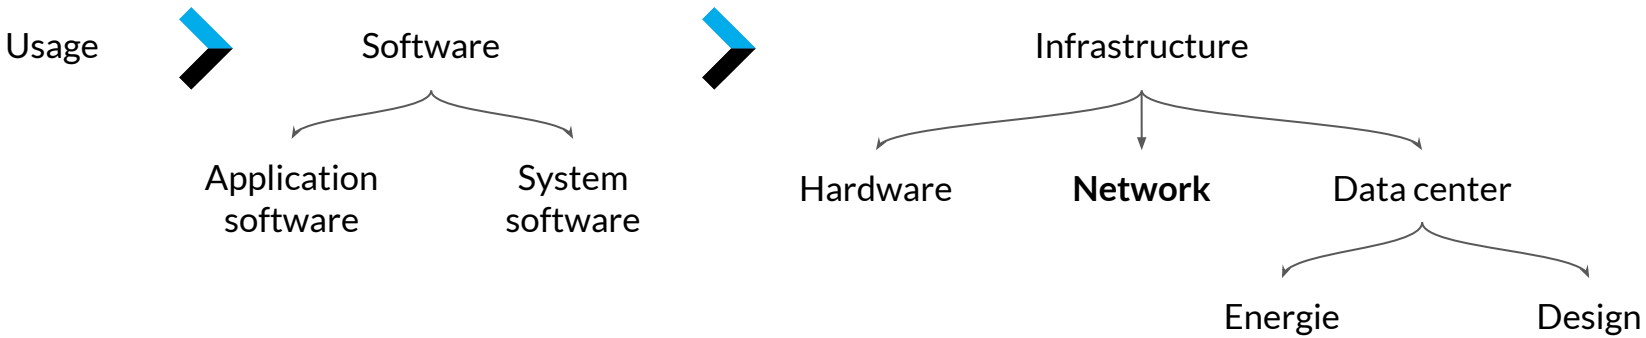

# Infrastructure : Network

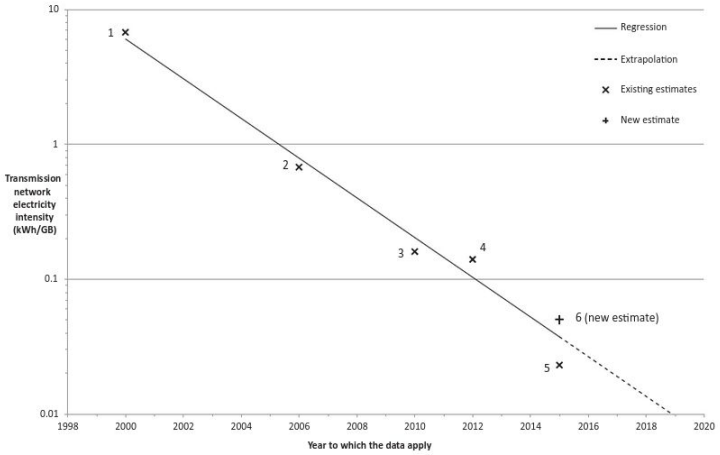


7,5 Wh/GB en 2021
Divisé /2 tout les 2 ans, mais …

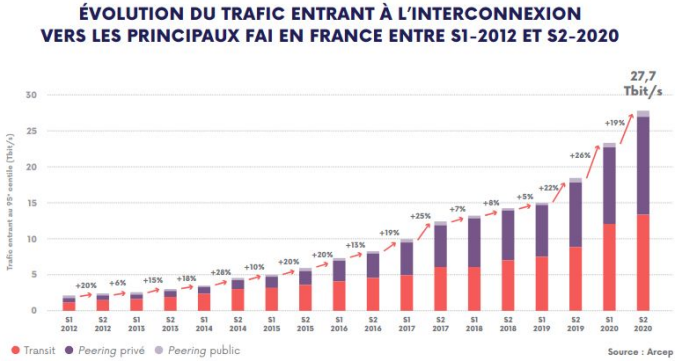


- Edge computing


Aslan, Joshua, et al. « Electricity Intensity of Internet Data Transmission: Untangling the Estimates: Electricity Intensity of Data Transmission ». *Journal of Industrial Ecology*, vol. 22, nº 4, août 2018, p. 785-98. *DOI.org (Crossref)*, doi:10.1111/jiec.12630.

Cao, Keyan, Yefan Liu, Gongjie Meng, et Qimeng Sun. « An Overview on Edge Computing Research ». *IEEE Access* 8 (2020): 85714-28. https://doi.org/10.1109/ACCESS.2020.2991734.

# Du HPC frugal ?

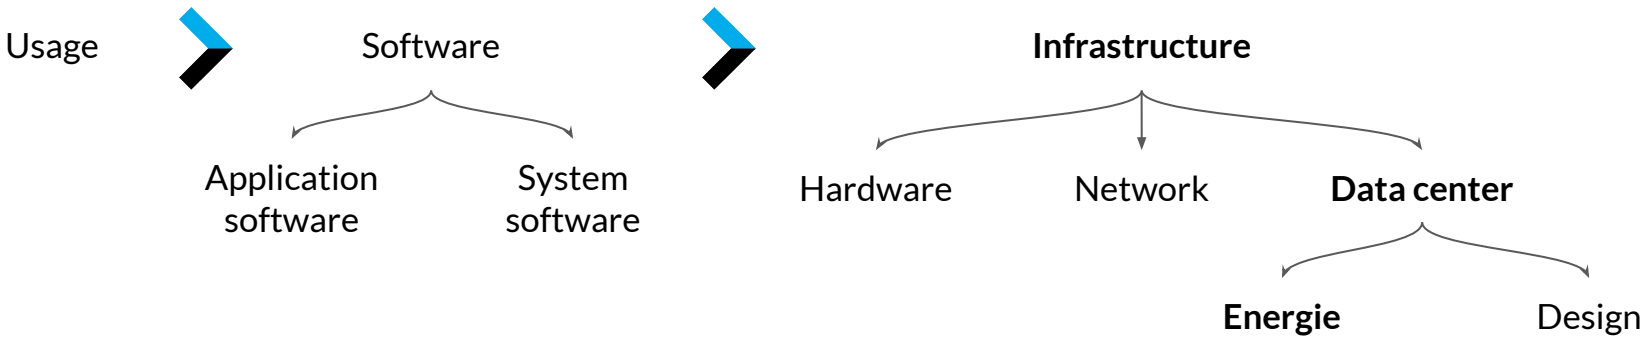

# Infrastructure : Energie

Dans quel pays faire ses calculs ?

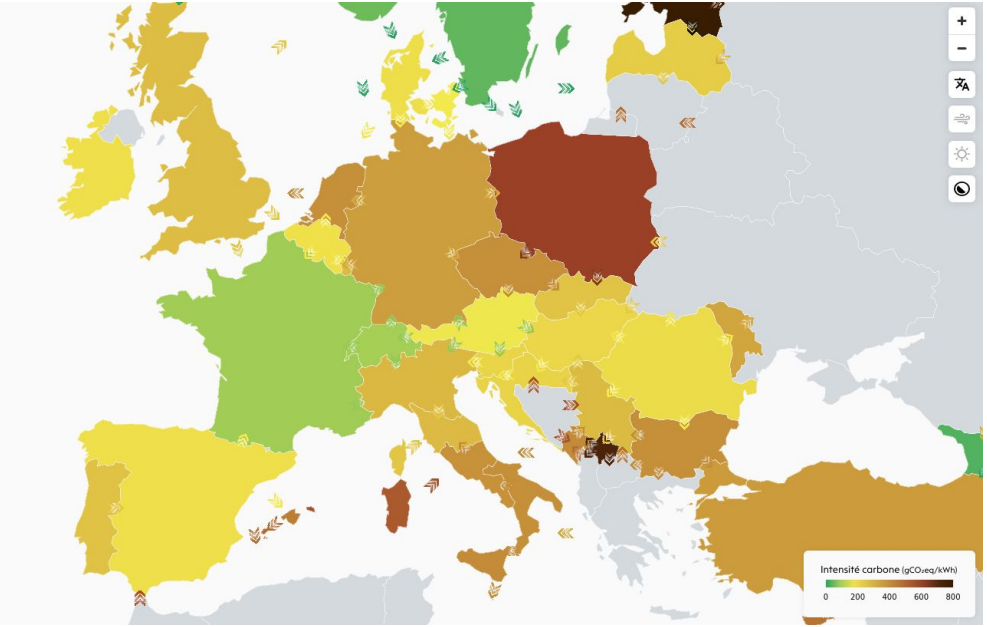


https://app.electricitymap.org/map

# Infrastructure : Energie

A quelle heure faire son calcul ?

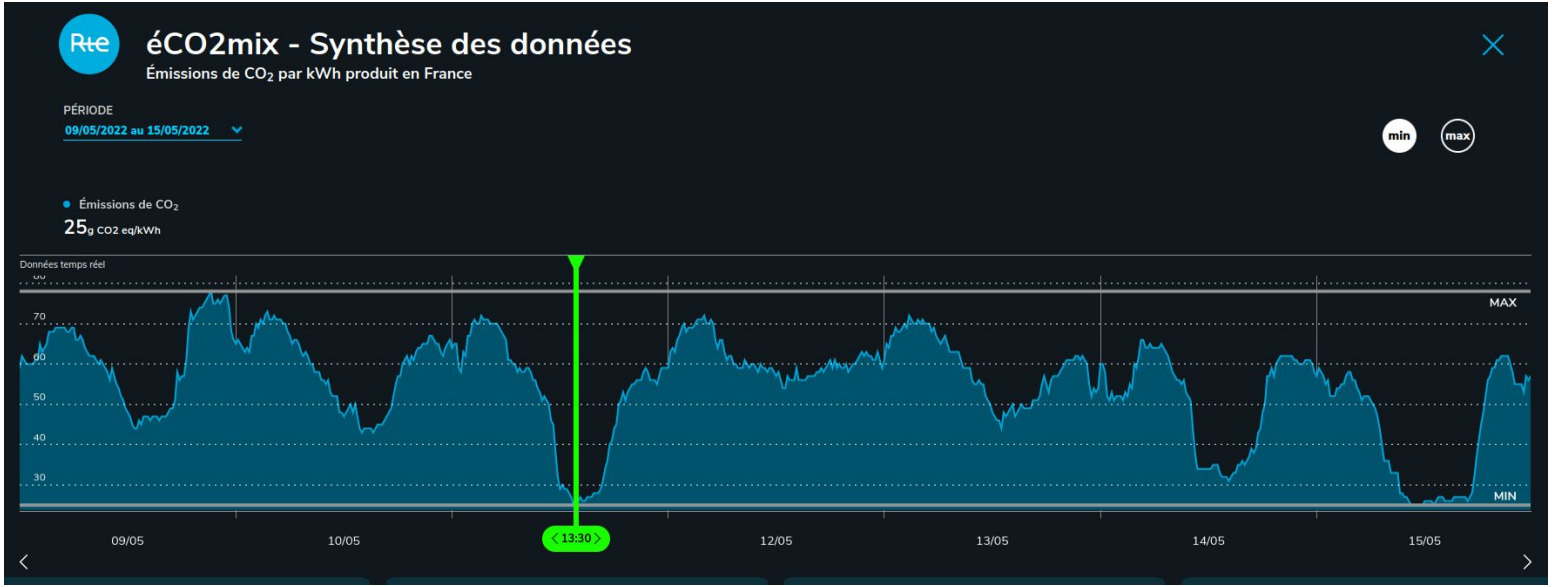


https://www.rte-france.com/eco2mix/synthese-des-donnees?type=co2#

# Du HPC frugal ?

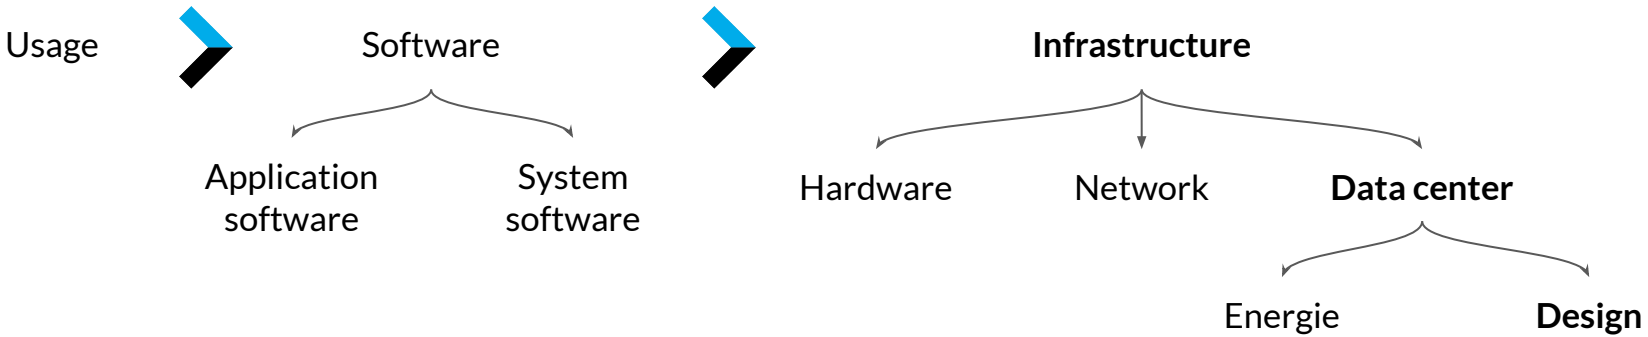

# Infrastructure : Design

Qarnot boiler

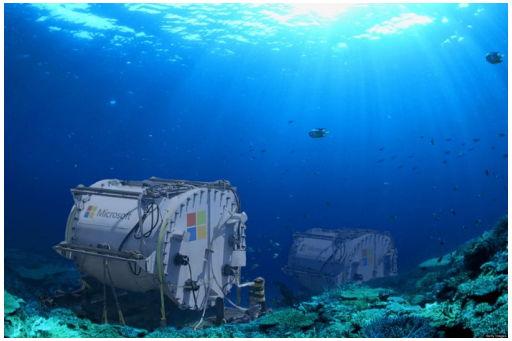

Microsoft underwater datacenter

Scandinavian Datacenter

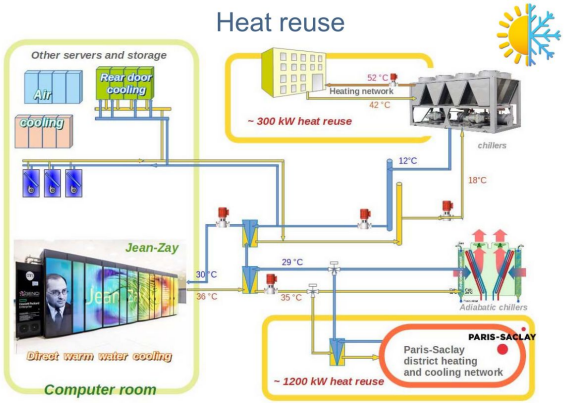

# Merci !

*Des questions?*


**Grégoire SOLEIL**

gregoire.soleil@qarnot.com