



# Quels leviers pour maximiser le ratio performance / l'impact environnemental ?

Cyril Deblois

Senior Sustainability BDM, AWS

# Agenda



- Introduction
- Impact environmental du cloud : les ordres de grandeurs
- Qu'est-ce qu'un cloud frugal ?
- Conclusion

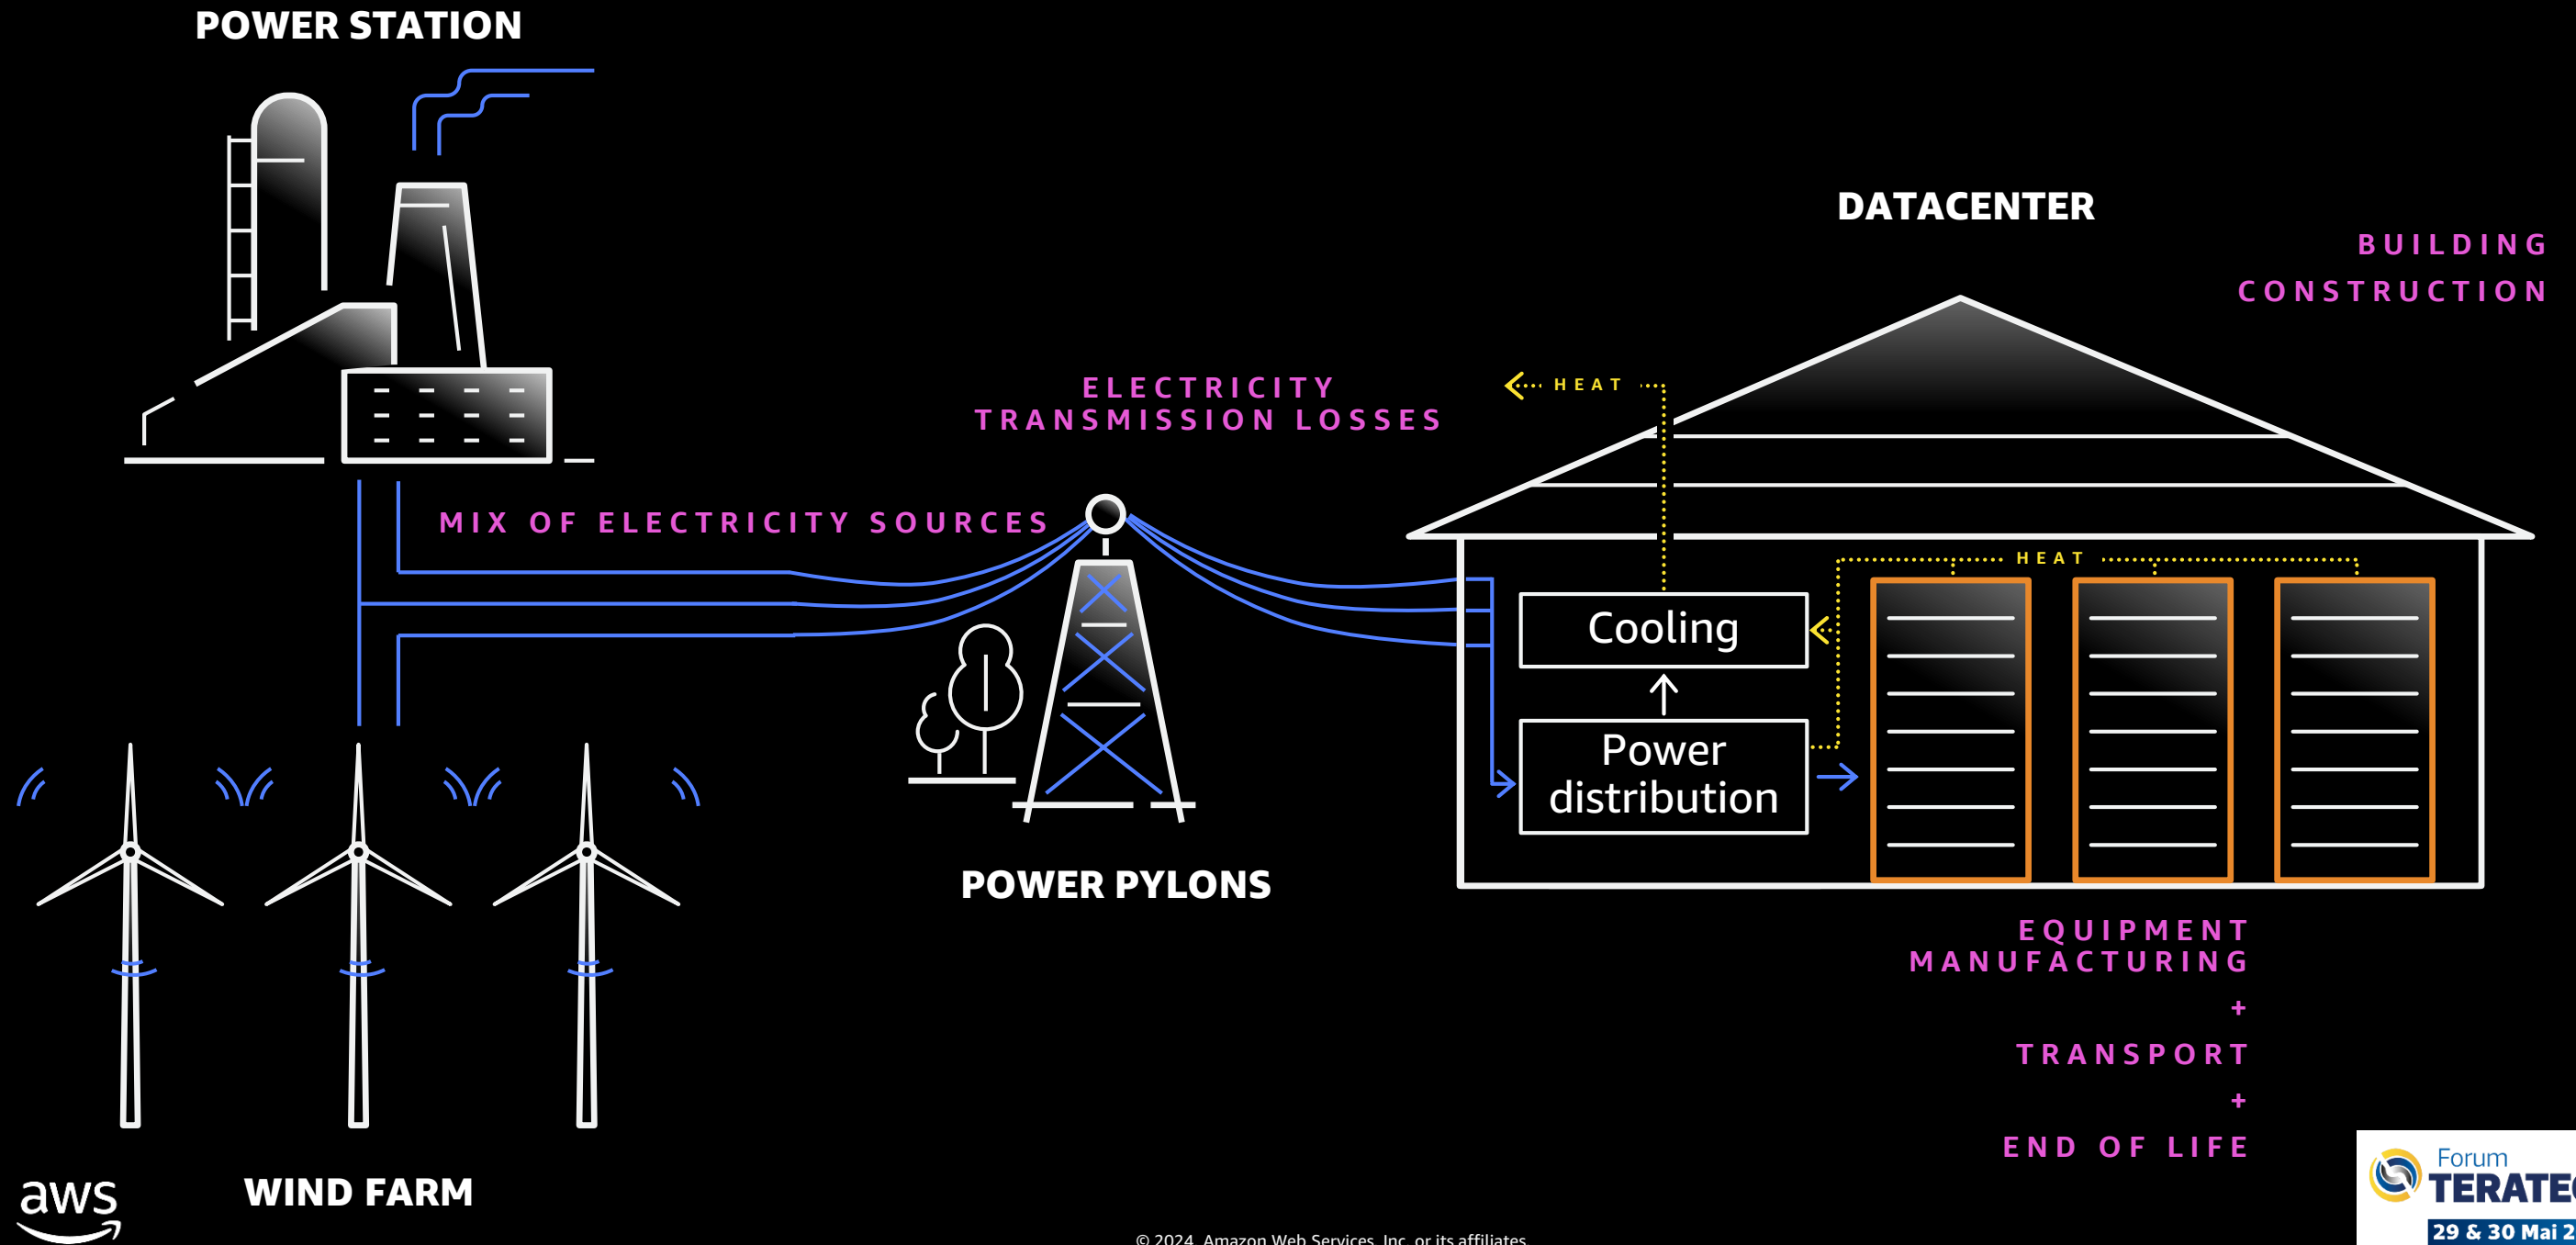


# Introduction



# Impact environnemental du cloud : les ordres de grandeurs

# Les charges de travail et les émission de GES



# Impacts environnementaux des Data Centres dans l'EU28

Manufacturing  
Scope 3  
27.1%

Use  
Scope 2  
72.2%

Scope 1  
0.7%

	Resource use, minerals and metals	Resource use, fossils	Acidification	Ecotoxicity, freshwater	Climate change	Ionising radiation, human health	Particulate matter	Photo-chemical ozone formation - human health	Raw materials	Waste production	Primary energy consumption	Final energy consumption (use)
Computing equipment	67.0%	12.8%	20.0%	25.1%	20.8%	9.4%	15.8%	20.8%	24.8%	34.1%	11.4%	0.0%
Storage equipment	12.8%	2.1%	3.4%	3.8%	3.6%	0.5%	2.8%	3.5%	4.0%	7.5%	1.9%	0.0%
Network equipment	2.8%	0.2%	0.2%	0.2%	0.2%	0.1%	0.1%	0.2%	0.4%	1.0%	0.1%	0.0%
Architecture and technical support equipment	16.7%	2.4%	3.0%	7.5%	2.5%	6.8%	2.8%	4.9%	10.9%	18.3%	2.1%	0.0%
Computing, storage and network energy consumption	0.4%	47.7%	42.5%	36.7%	41.7%	48.1%	45.3%	40.8%	34.6%	22.6%	48.8%	57.8%
Technical support equipment energy consumption	0.3%	34.8%	31.0%	26.8%	30.5%	35.1%	33.1%	29.8%	25.3%	16.5%	35.6%	42.2%
Refrigerant leaks	0.0%	0.0%	0.0%	0.0%	0.7%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%

PUE: 1.7

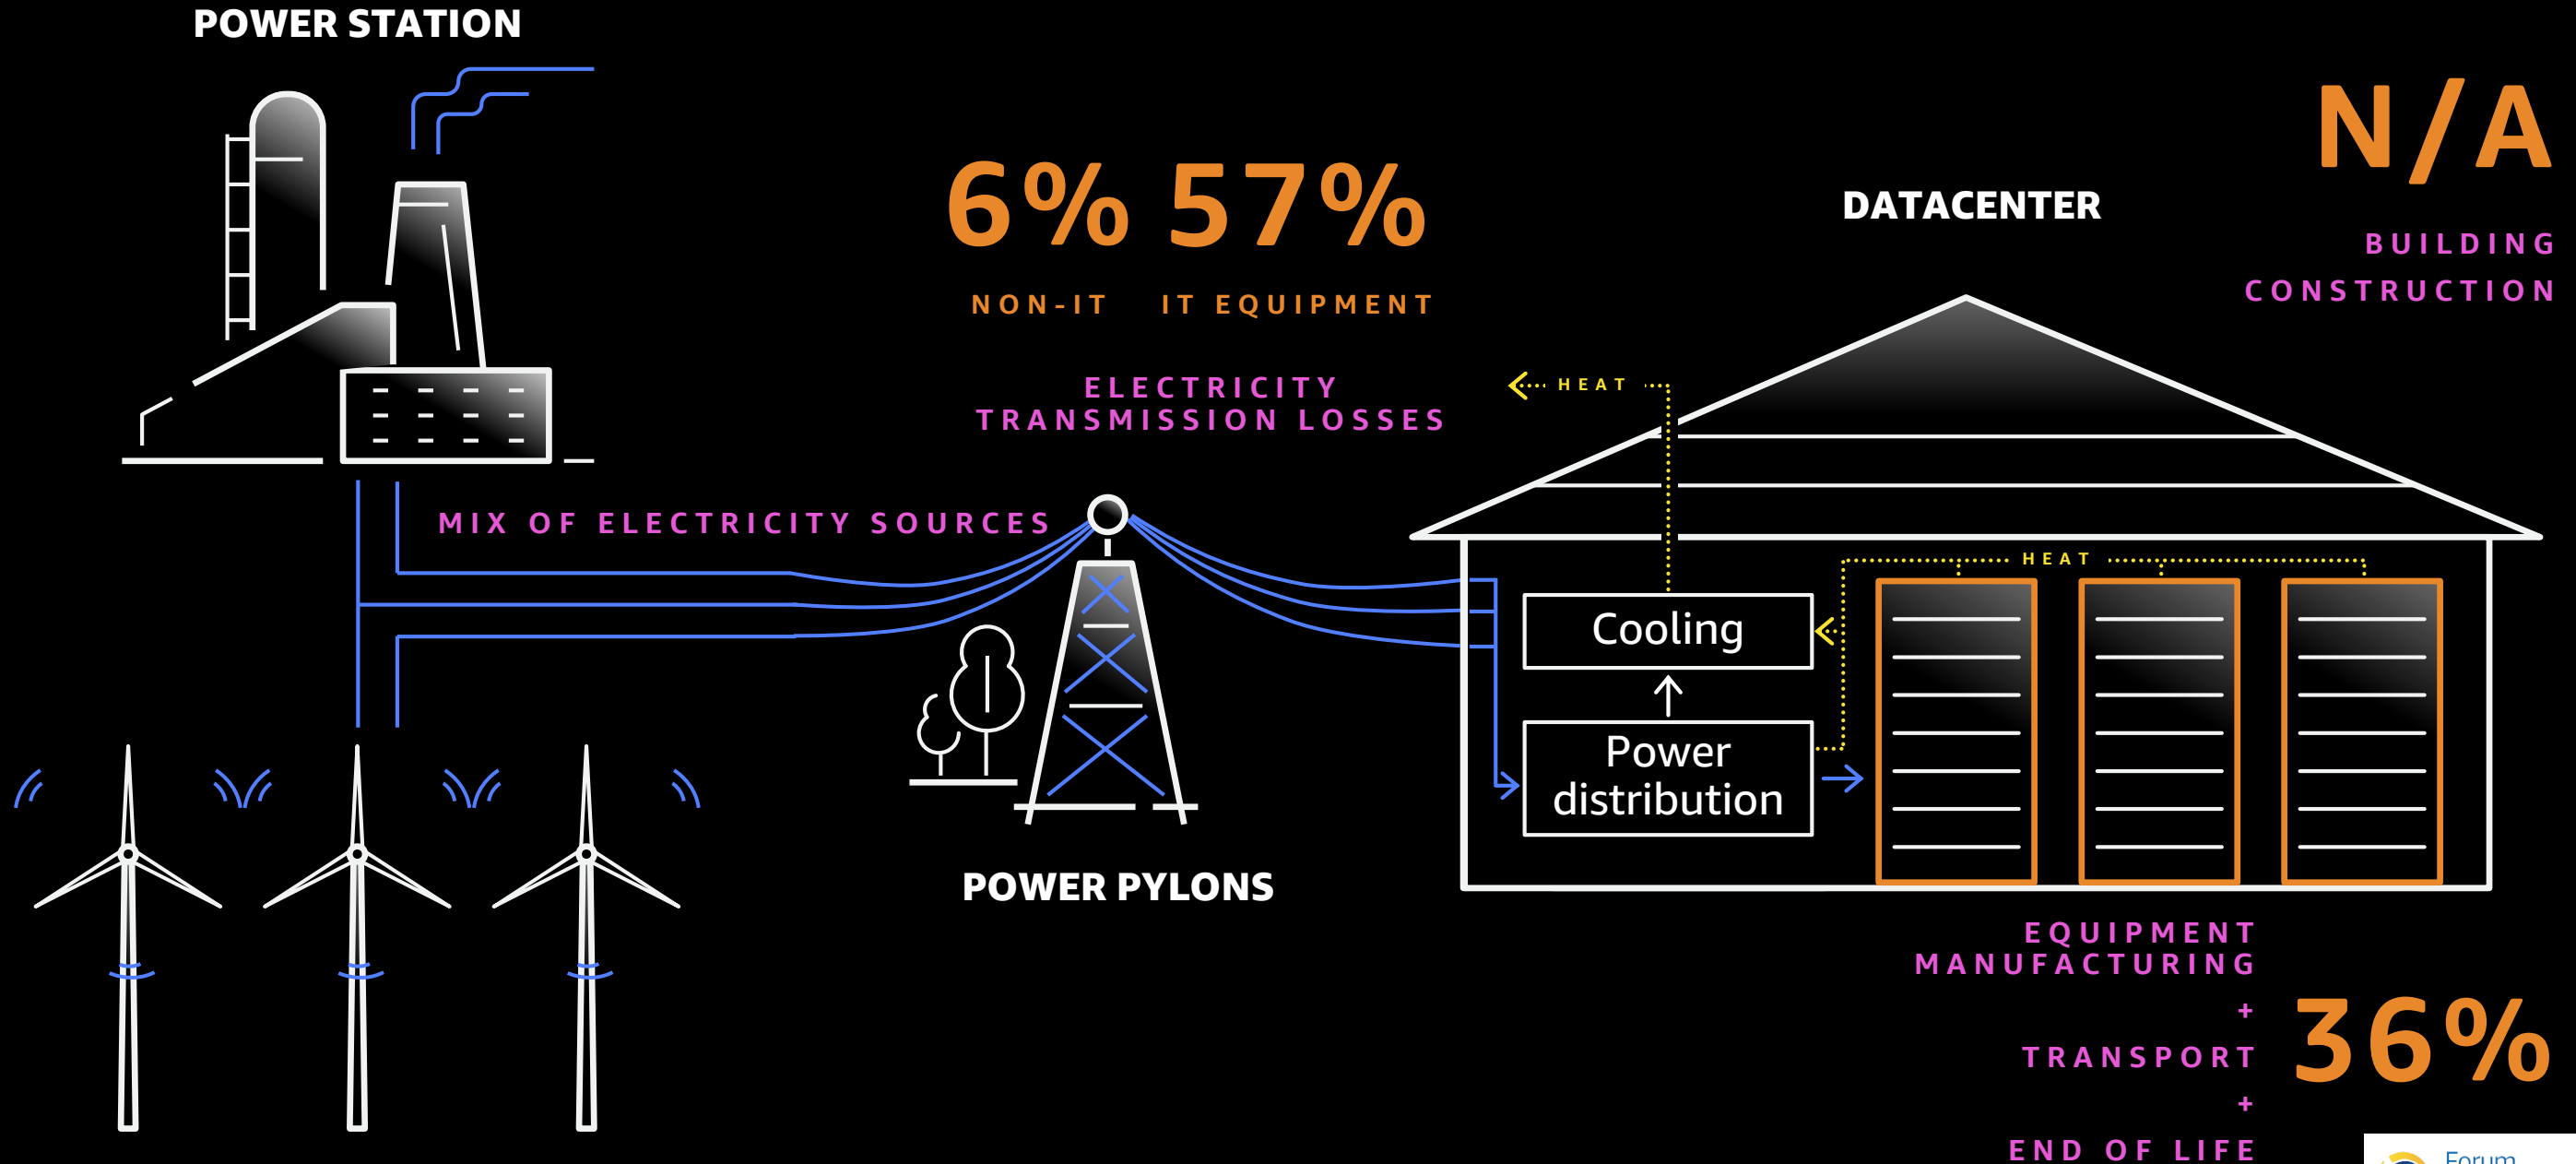
Source: "Digital Technologies in Europe" commissioned by the European Parliamentary group of the Greens/EFA.



# Impacts environnementaux des Data Centres AWS (EU28)



# Les charges de travail et les émissions de GES





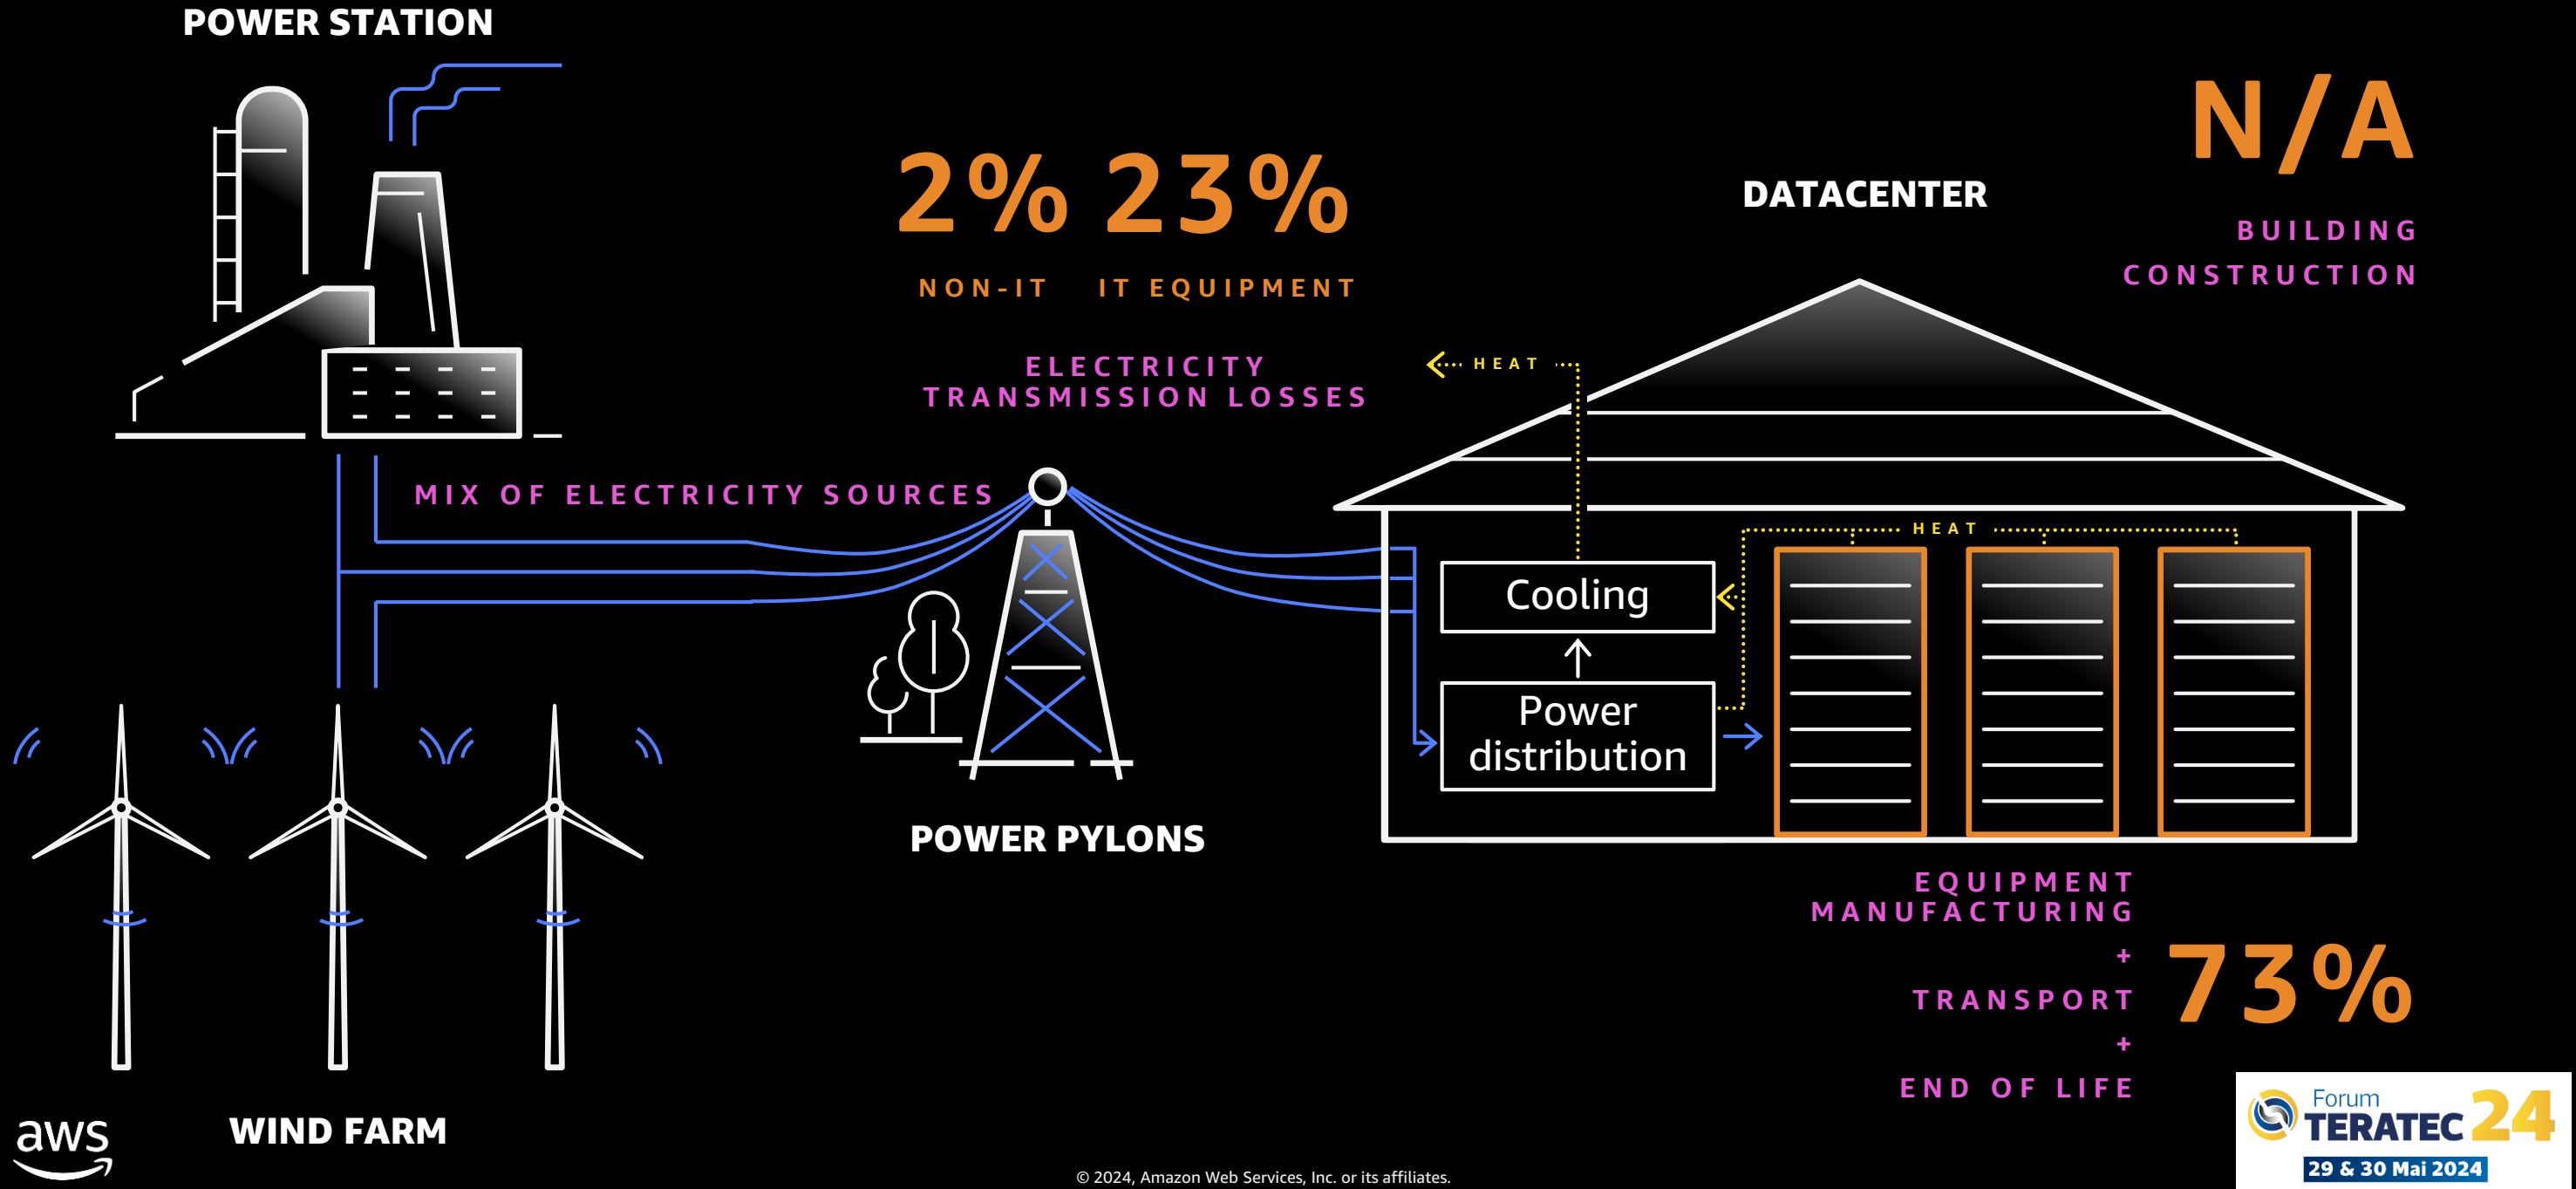
# Impact du type d'énergie dans les Regions AWS

Region	gCO2e/ kWh	Scope 1	Scope 2	Scope 3	Scope 1	Scope 2	Scope 3	TOTAL
EU-28	270	0.9%	63.3%	35.9%	0.9	63.3	35.9	100.1
US East (Virginia)	396	0.7%	71.6%	27.7%	0.9	92.8	35.9	129.6
EU West (Ireland)	371	0.7%	70.3%	29.0%	0.9	87.0	35.9	123.8
EU Cen. (Frankfurt)	400	0.7%	71.8%	27.5%	0.9	93.8	35.9	130.6
US West (Oregon)	206	1.1%	56.8%	42.2%	0.9	48.3	35.9	85.1
US East (Ohio)	396	0.7%	71.6%	27.7%	0.9	92.8	35.9	129.6
EU West (London)	199	1.1%	55.9%	43.0%	0.9	46.7	35.9	83.5
APAC NE (Tokyo)	536	0.6%	77.3%	22.1%	0.9	125.7	35.9	162.5
APAC SE (Sydney)	545	0.5%	77.6%	21.8%	0.9	127.8	35.9	164.6
APAC SE (Singapore)	488	0.6%	75.7%	23.7%	0.9	114.4	35.9	151.2
US West (California)	262	0.9%	62.5%	36.5%	0.9	61.4	35.9	98.2
APAC N. (Seoul)	444	0.6%	73.9%	25.5%	0.9	104.1	35.9	140.9
CA Cen. (Montreal)	31	2.0%	16.5%	81.5%	0.9	7.3	35.9	44.1
APAC S. (Mumbai)	747	0.4%	82.6%	16.9%	0.9	175.1	35.9	211.9
EU West (Paris)	53	1.8%	25.2%	72.9%	0.9	12.4	35.9	49.2
Gov Cloud (UGE1)	396	0.7%	71.6%	27.7%	0.9	92.8	35.9	129.6
SA East (Sao Paulo)	90	1.6%	36.4%	62.0%	0.9	21.1	35.9	57.9

\*Year 2018 from International Energy Agency: <https://www.iea.org/reports/european-union-2020>

\*\* From electricitymap.org in 2023 over the last 12 months, electricity consumed

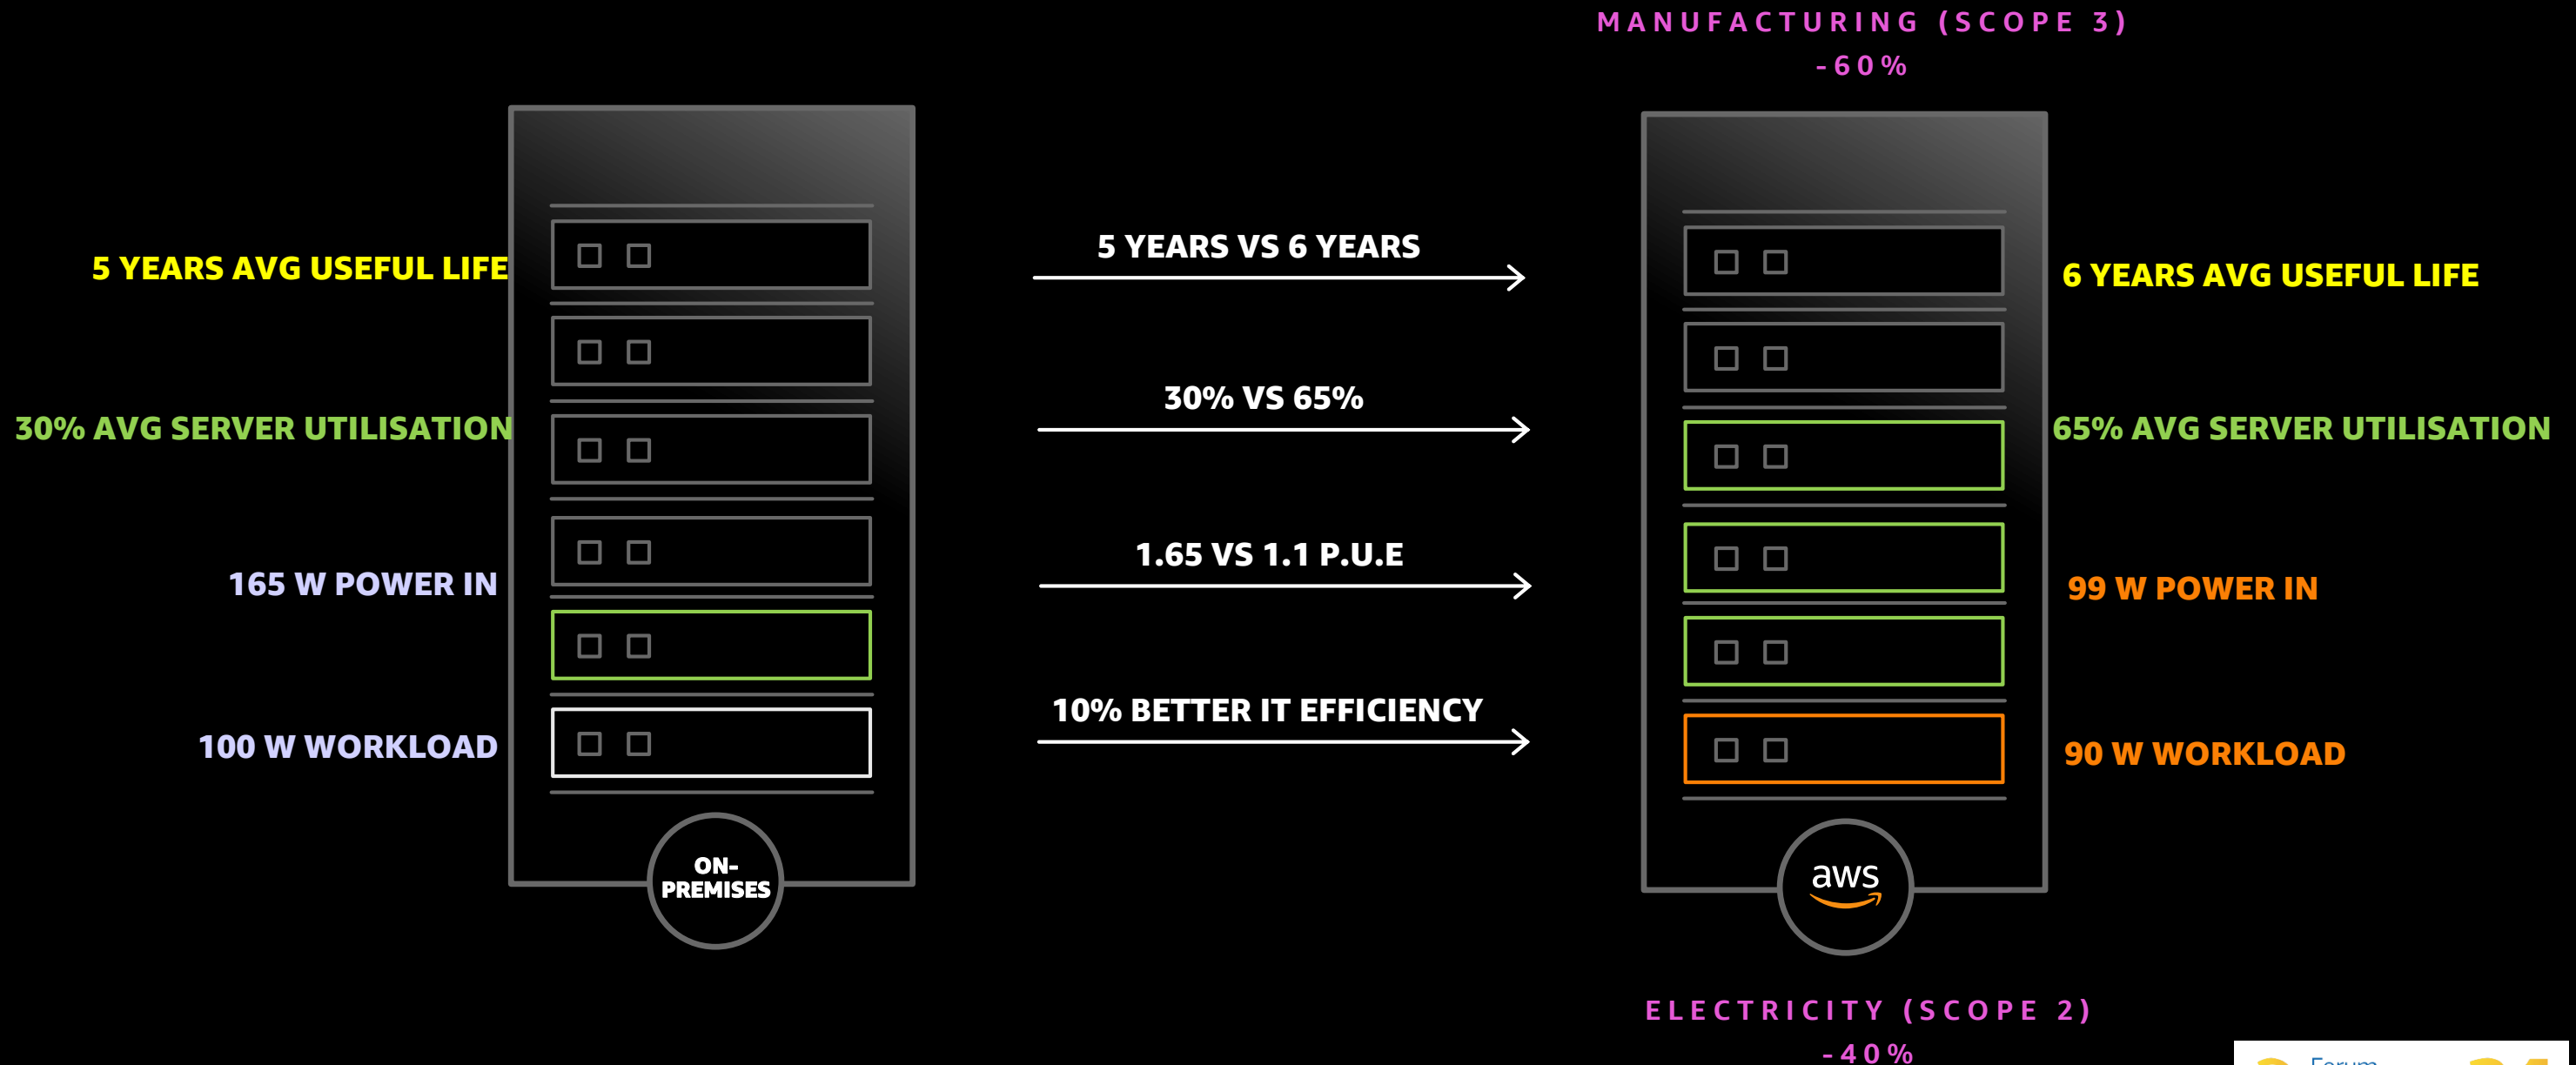
# Charges de travail et émissions de GES AWS en France



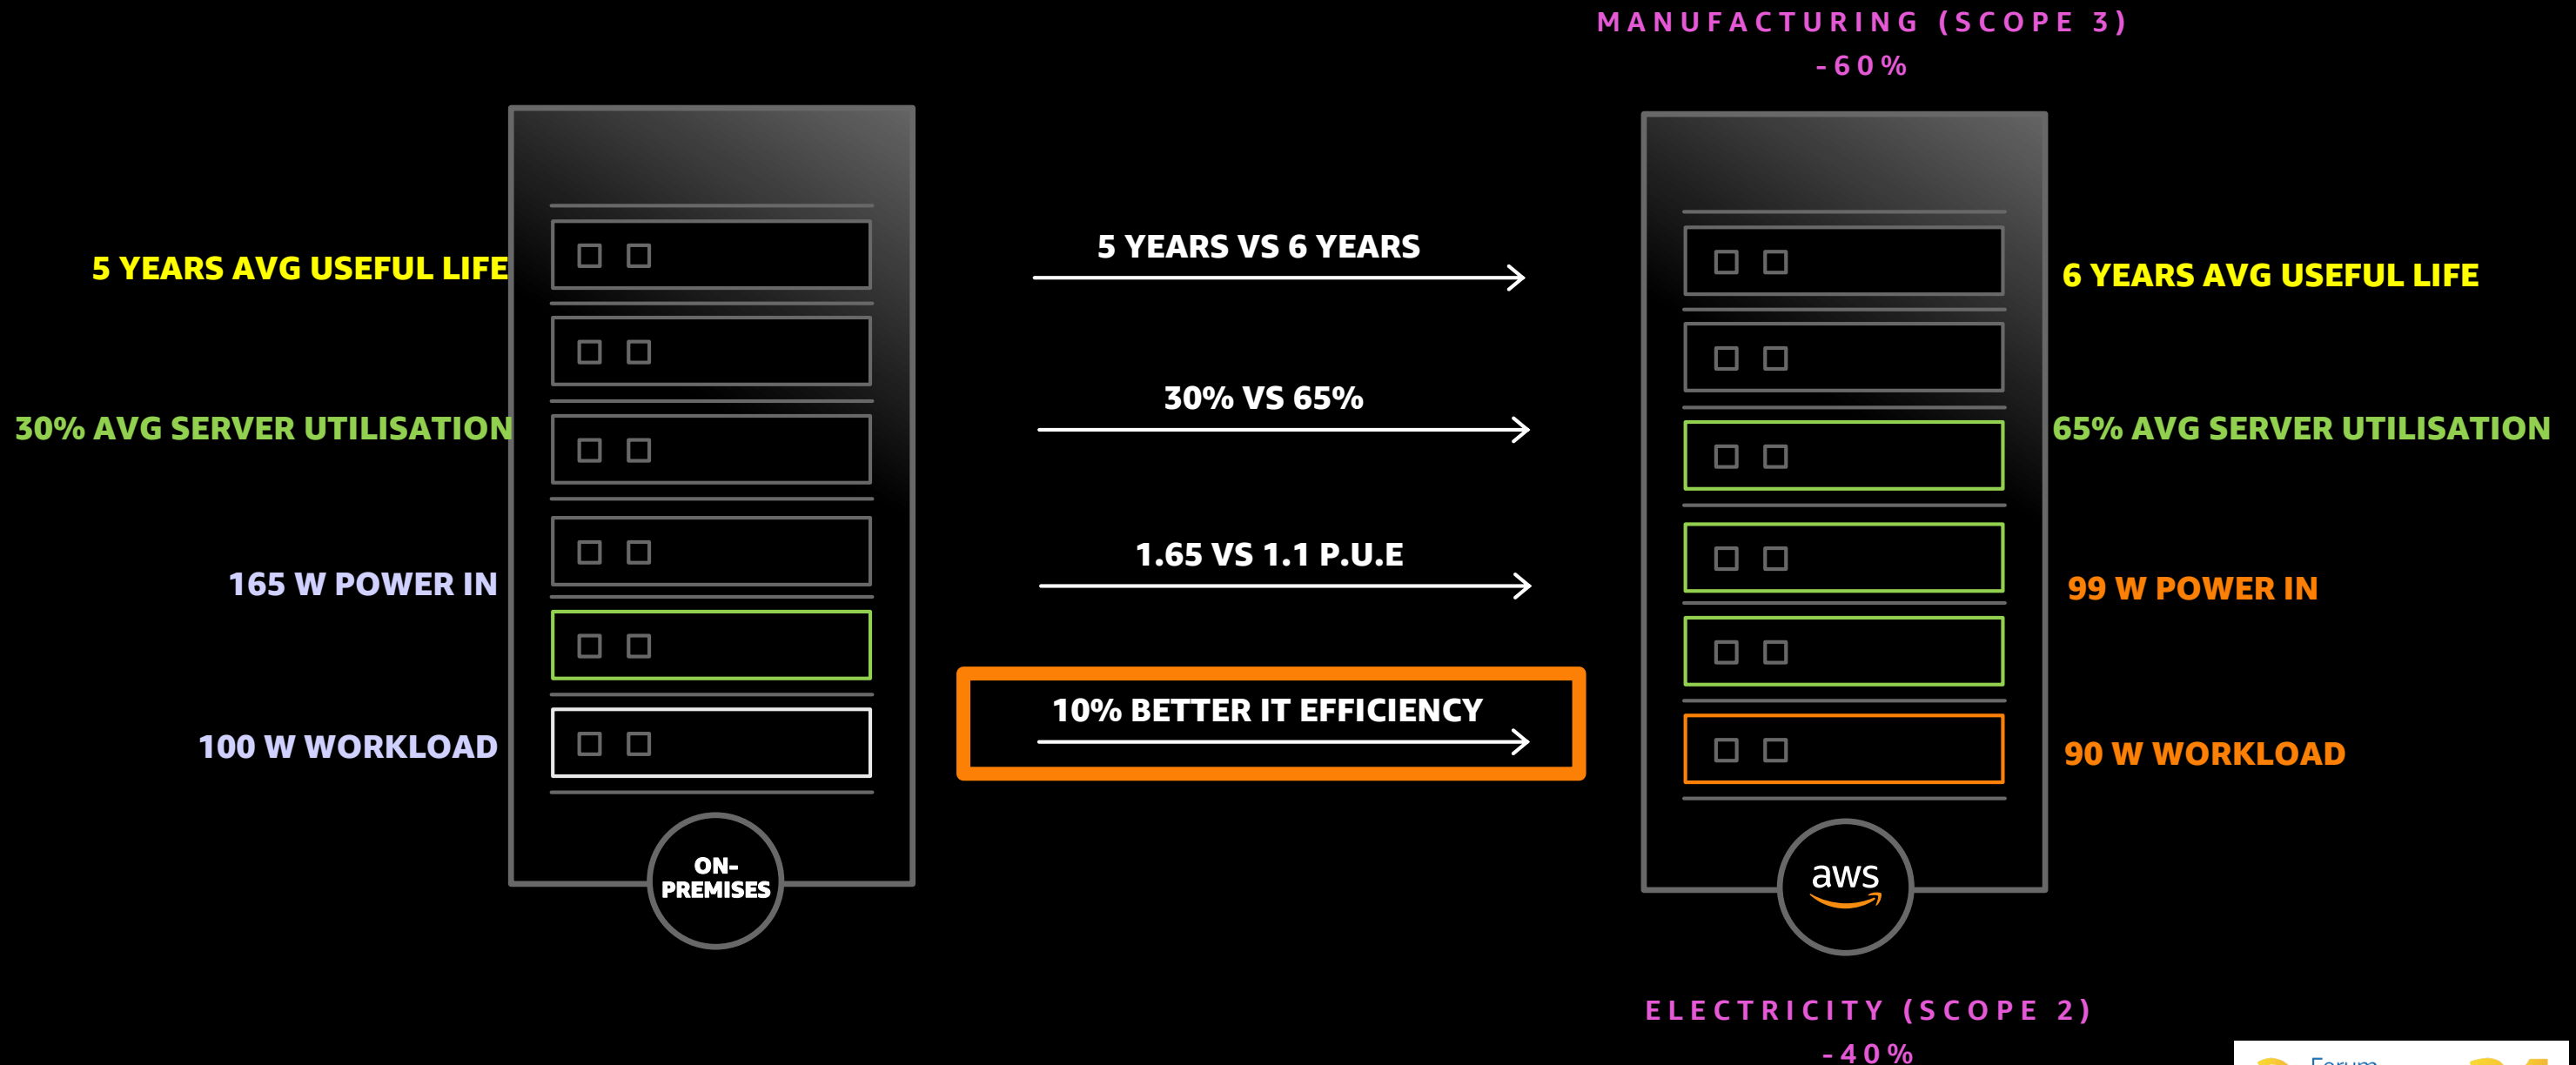


# Qu'est-ce qu'un cloud frugal ?

# Cas pratique sur une charge de travail



# Le cas de l'efficacité de l'IT



# AWS avec Graviton, Inferentia, et Trainium


Graviton3-based Elastic Compute Cloud instances use up to **60% less energy** than comparable Amazon EC2 instances

---

Models built on Trainium result in **energy-consumption reductions of up to 25%** versus comparable instances


---

Our Inferentia2 machine learning accelerator delivers up to **50% higher performance per watt**




# NVIDIA avec les GH200 ...

CPU Data Center  
1500 x86 CPU



GH200 Data Center  
72 GH200



## EQUIVALENT PERFORMANCE

6X Lower TCO

22X Fewer Server Nodes

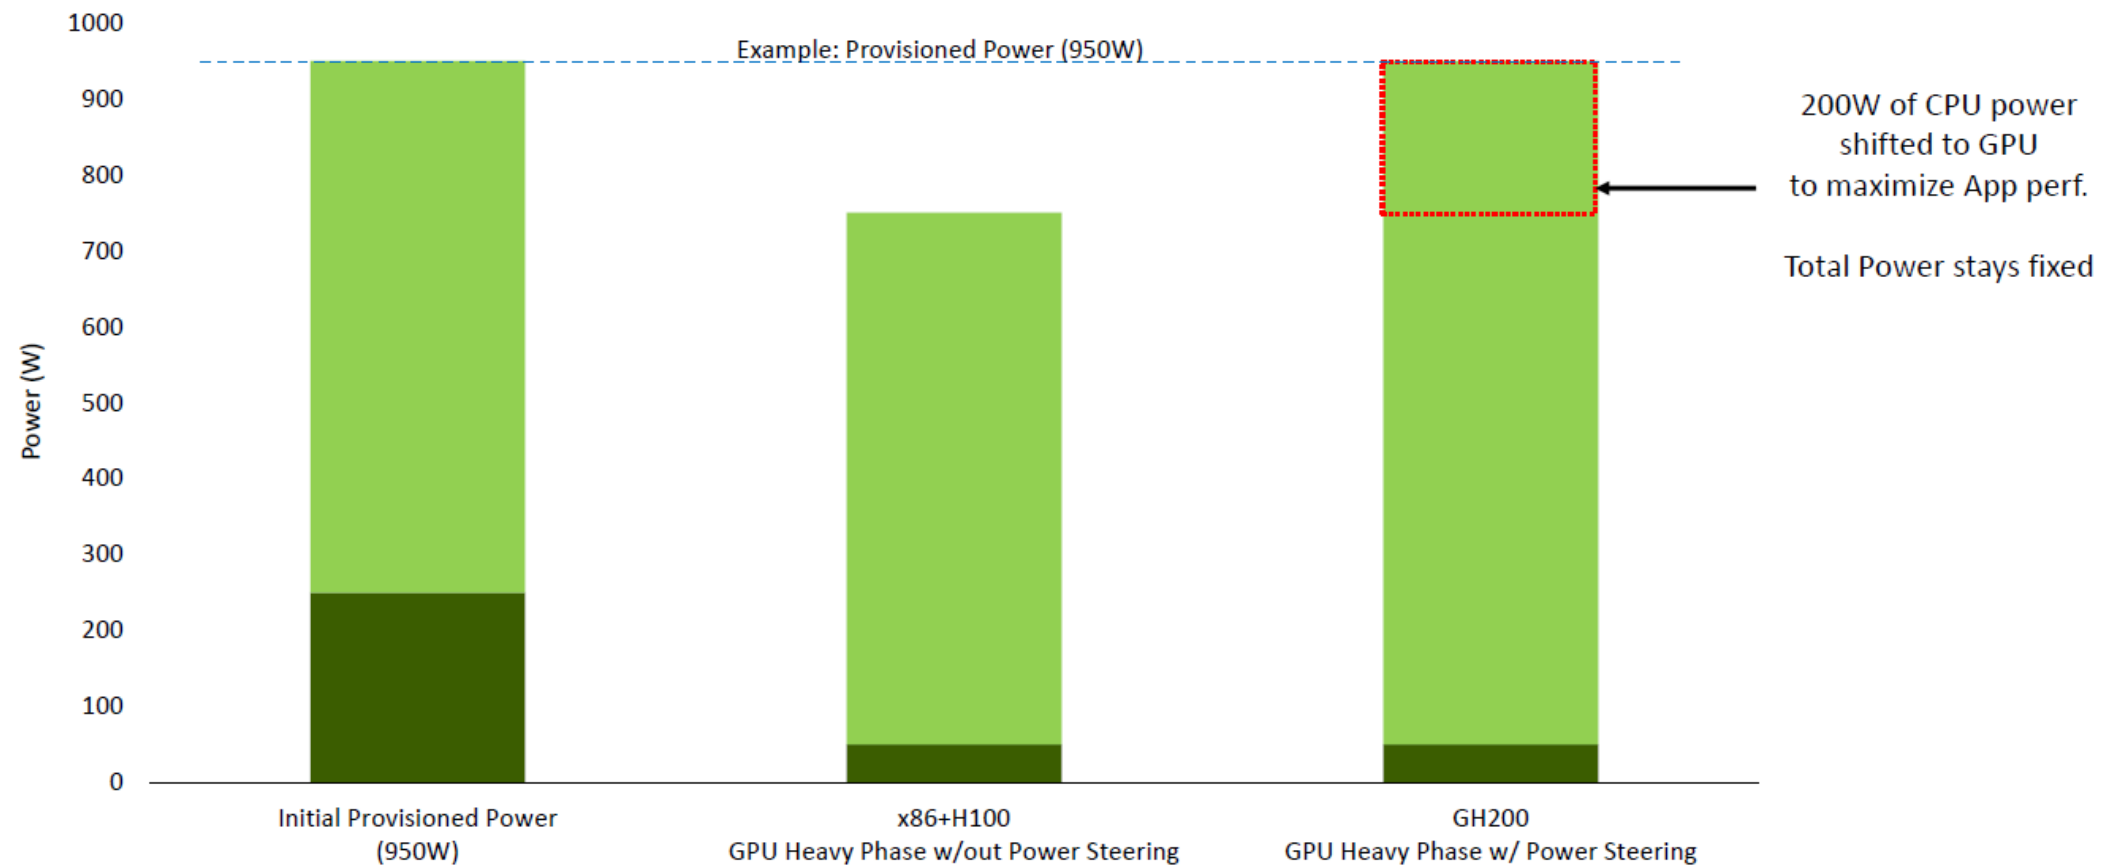
12X More Energy Efficient

# ... et le Power Steering

## Optimizing Performance Through Power Steering

Getting the Most Out of provisioned power

■ CPU ■ GPU





# La largeur et la profondeur des services est critique




Fig. 1. The tasks examined in our study and the average quantity of carbon emissions they produced (in g of CO<sub>2</sub>) for 1,000 queries. N.B. The y axis is in logarithmic scale.

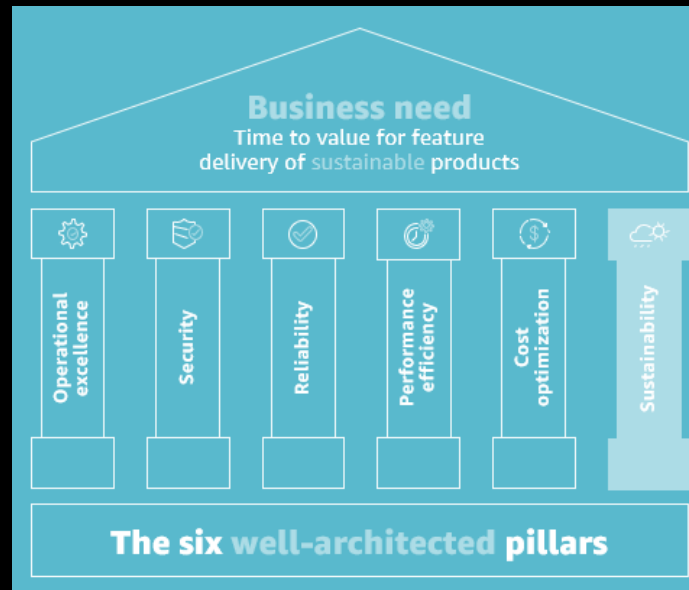
“Power Hungry Processing: Watts Driving the Cost of AI Deployment?” Alexandra Sasha Luccioni and Yacine Jernite, Hugging Face, Canada/USA  
Emma Strubell, Carnegie Mellon University, Allen Institute for AI, USA, 28 Nov. 2023

# AWS et l'accompagnement GreenOps

Well-Architected Sustainability Review (led by an AWS SA or an accredited AWS Partner)

Well-Architected Sustainability Labs. Self-paced labs to put in practice some of the 118 Best Practices of the Well-Architected Sustainability pillar

Well-Architected Sustainability Tool. To self-assess your architecture against the 118 Best Practices of the Well-Architected Sustainability pillar





# Conclusion



The logo for Forum TERATEC 24. It features a circular icon on the left with a blue and yellow swirl. To the right, the word "Forum" is in blue, "TERATEC" is in large blue letters, and "24" is in large yellow letters. Below this, a blue rectangular box contains the text "29 &amp; 30 Mai 2024" in white.

# Q&A



**Merci !**