Datacenter Sustainability: Measured, not Guessed!

Babak Falsafi

EcoCloud

ecocloud.ch
sdea.ch

DATA CENTER
EFFICIENCY

OUR DIGITAL UNIVERSE

Fueled by:

- Data volume
- Data growth rate
- Monetization of data
- Data's impact on GDP
- ….now AI

DATACENTERS: THE BACKBONE OF OUR DIGITAL UNIVERSE

- 100s of thousands of commodity or home-brewed servers
  - Consuming 10s to 100s MW
- Centralized to exploit economies of scale
- Network fabric w/ µ-second connectivity
- Often limited by ingress
  - Electricity
  - Network
  - Cooling

Boydton DC, 300MW

CLOUDS AT VARIOUS SCALES

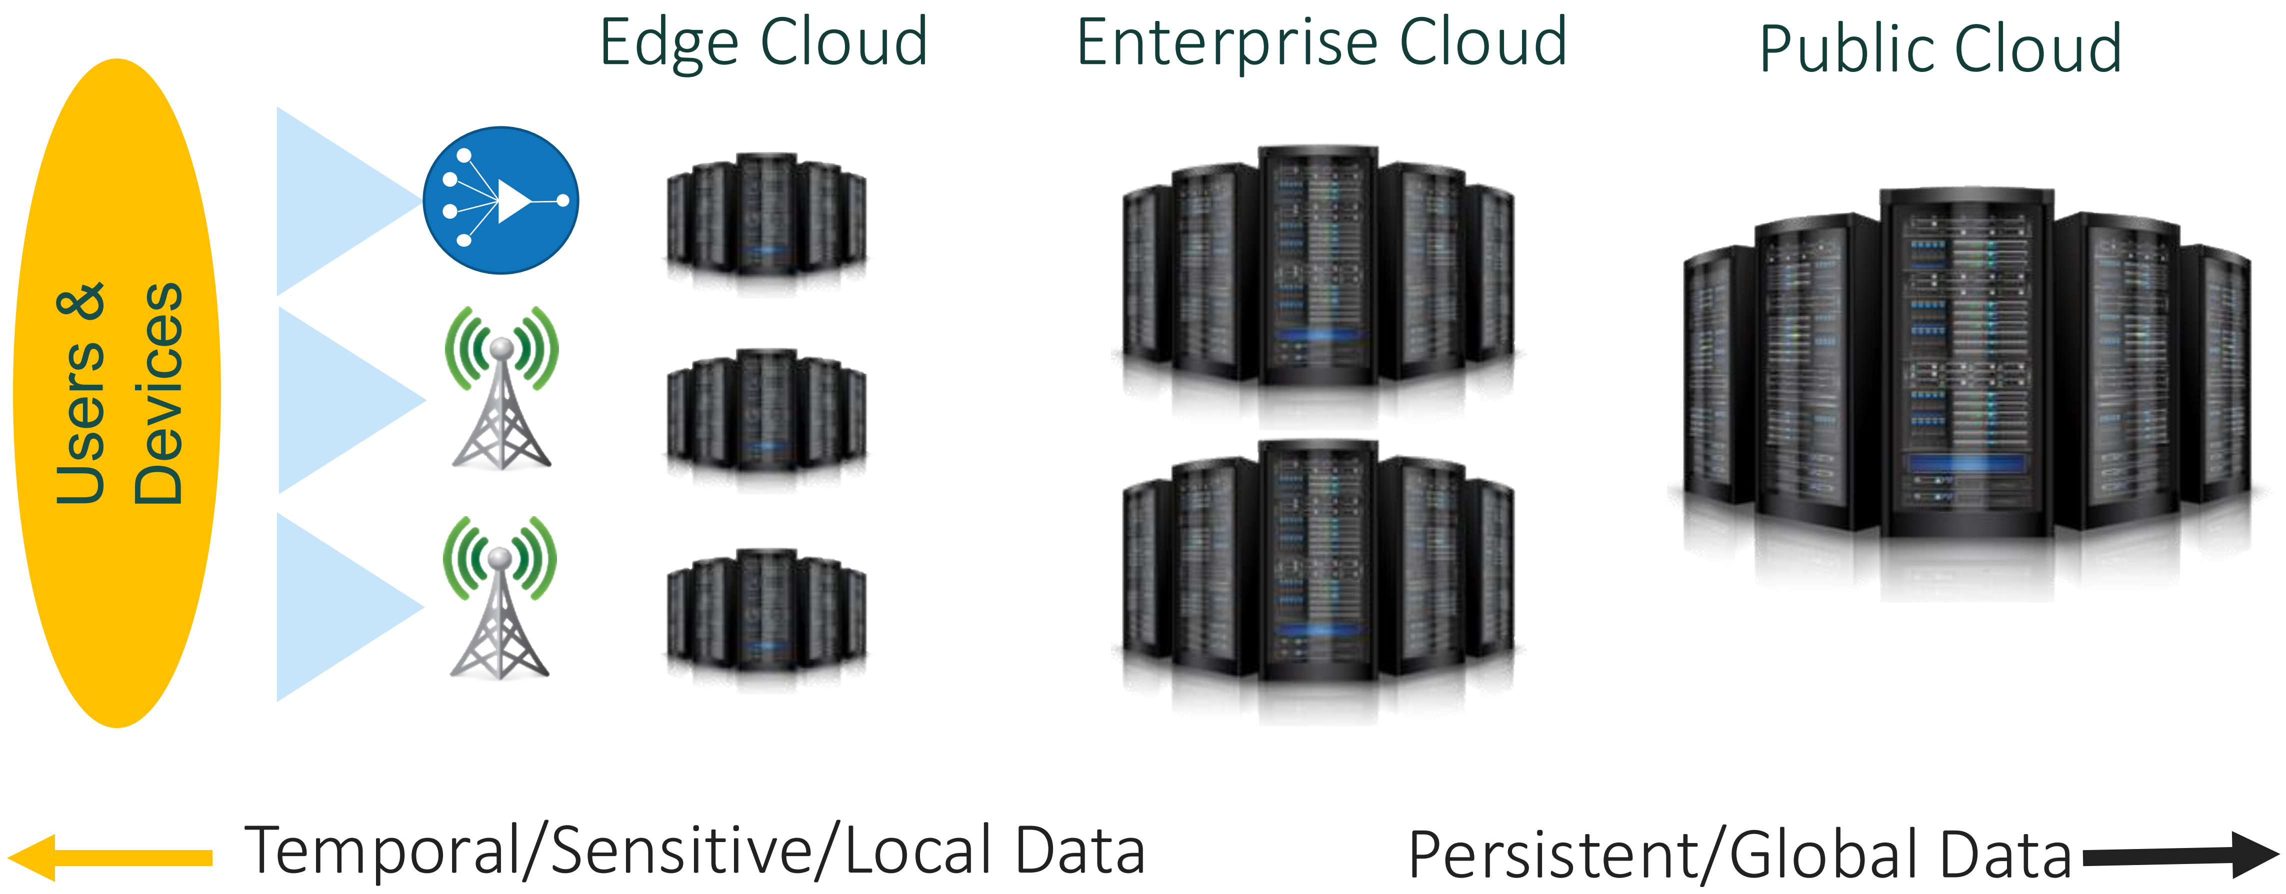

UNIVERSE MADE POSSIBLE BY MOORE'S LAW

1971

Intel 4004

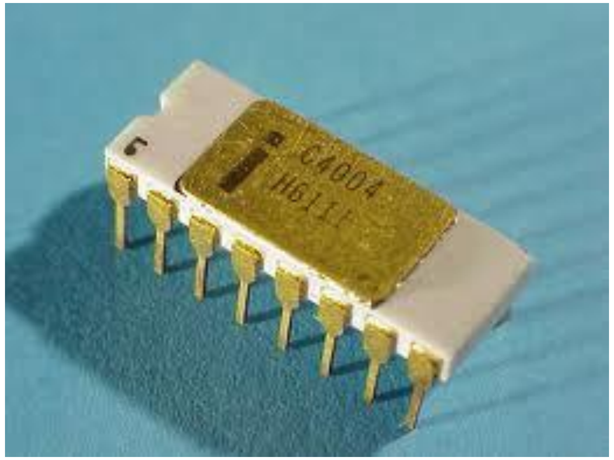

92,000 ops/s

1 Watt

2021

Intel Ice Lake

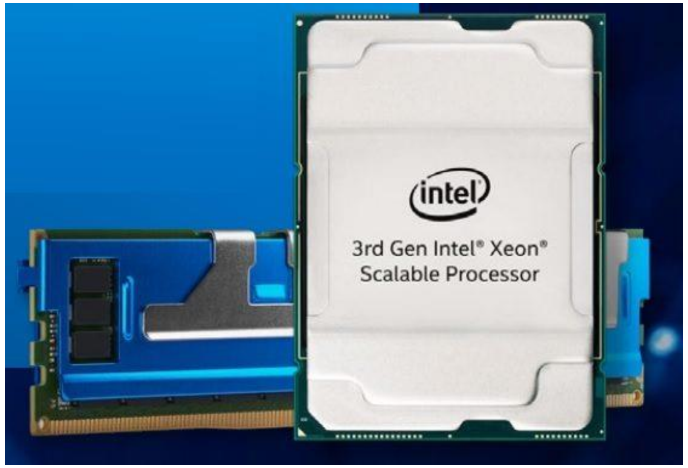

1,200,000,000,000 ops/s

270 Watts

MOORE'S LAW: EXPONENTIAL DENSITY & EFFICIENCY

1971

Intel 4004

2021

Intel Ice Lake

In 50 years:

13 million times faster

48 thousand times more efficient

92,000 ops/s
1 Watt

1,200,000,000,000 ops/s
270 Watts

LONG LIVE MOORE'S LAW

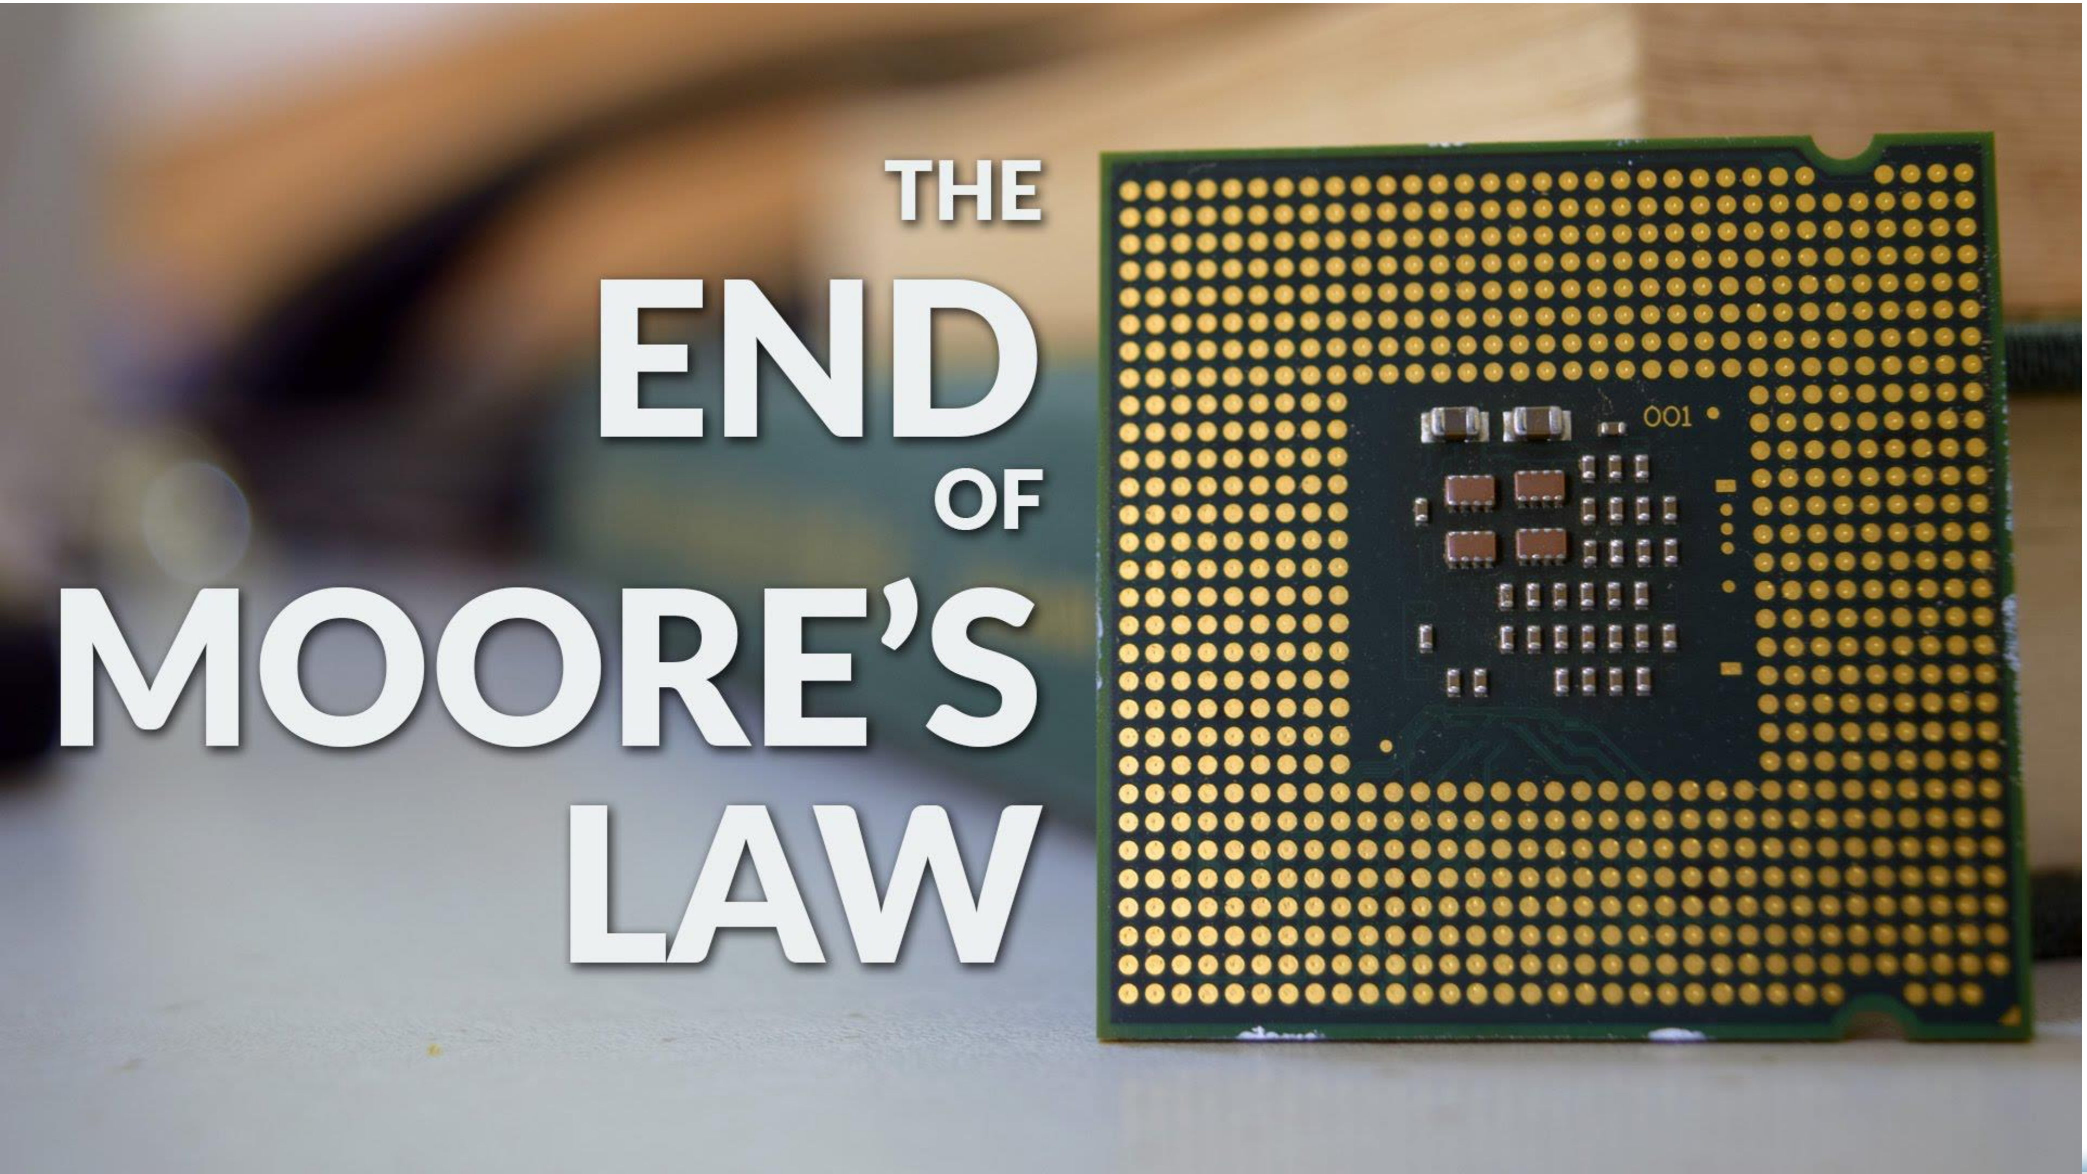

EcoCloud

BUILDING BIGGER & FASTER CHIPS

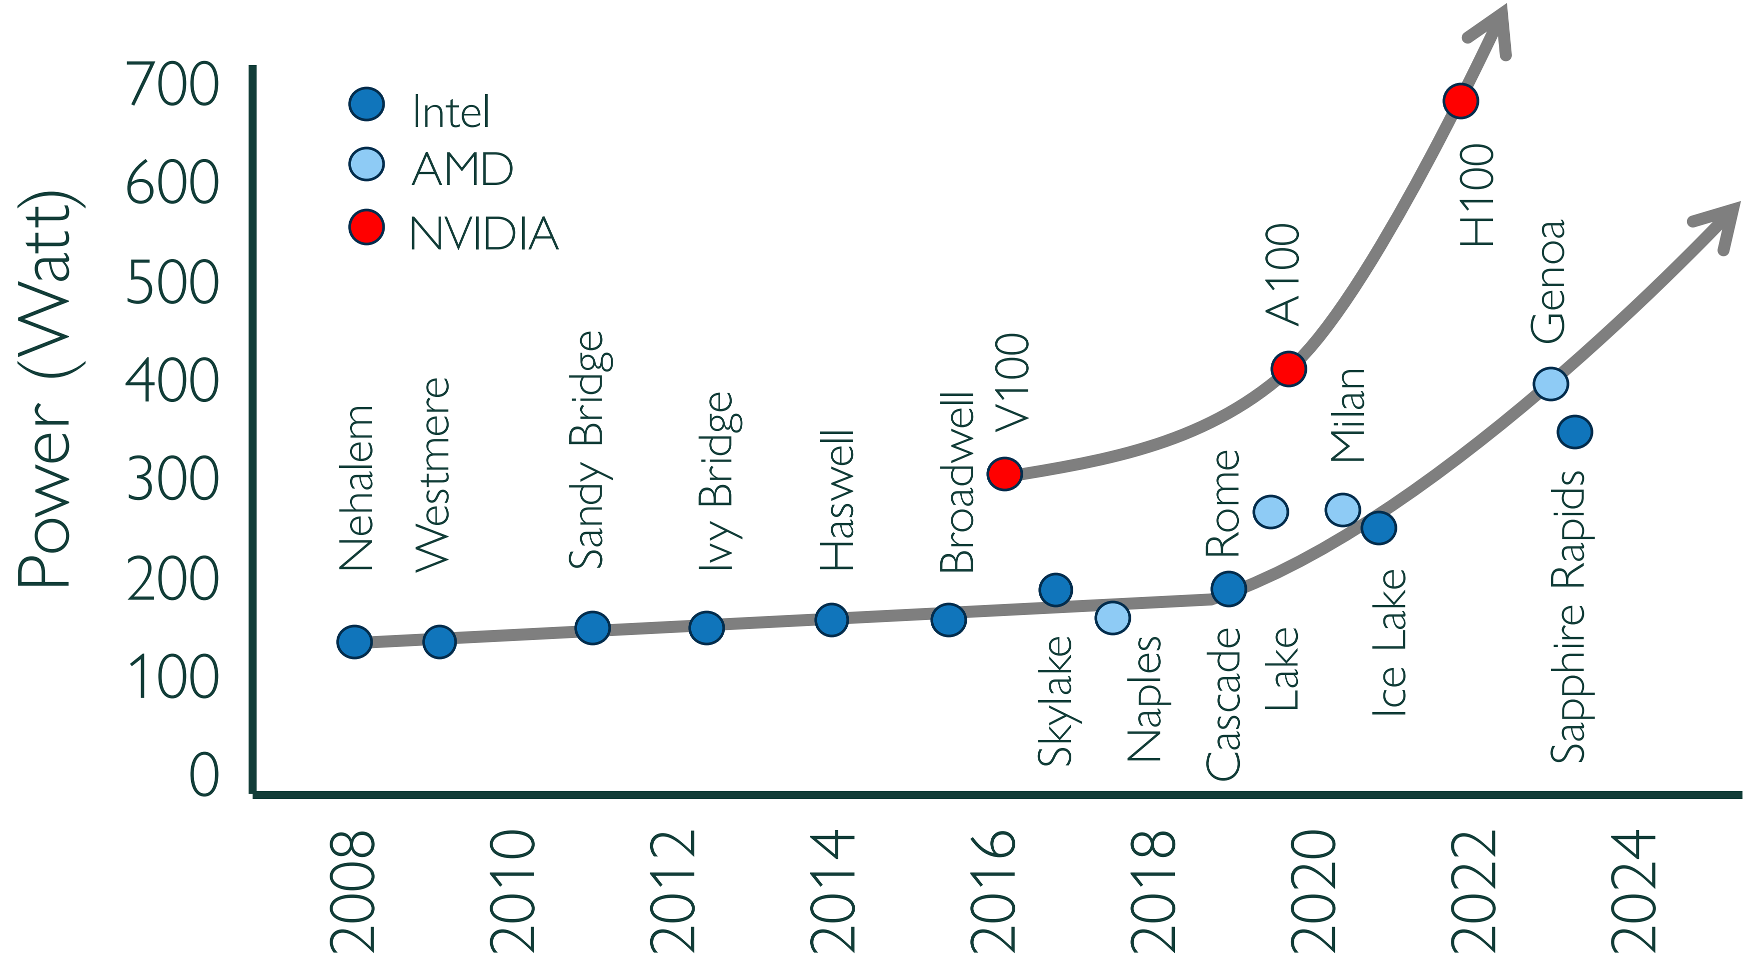

BUILDING BIGGER & FASTER CHIPS

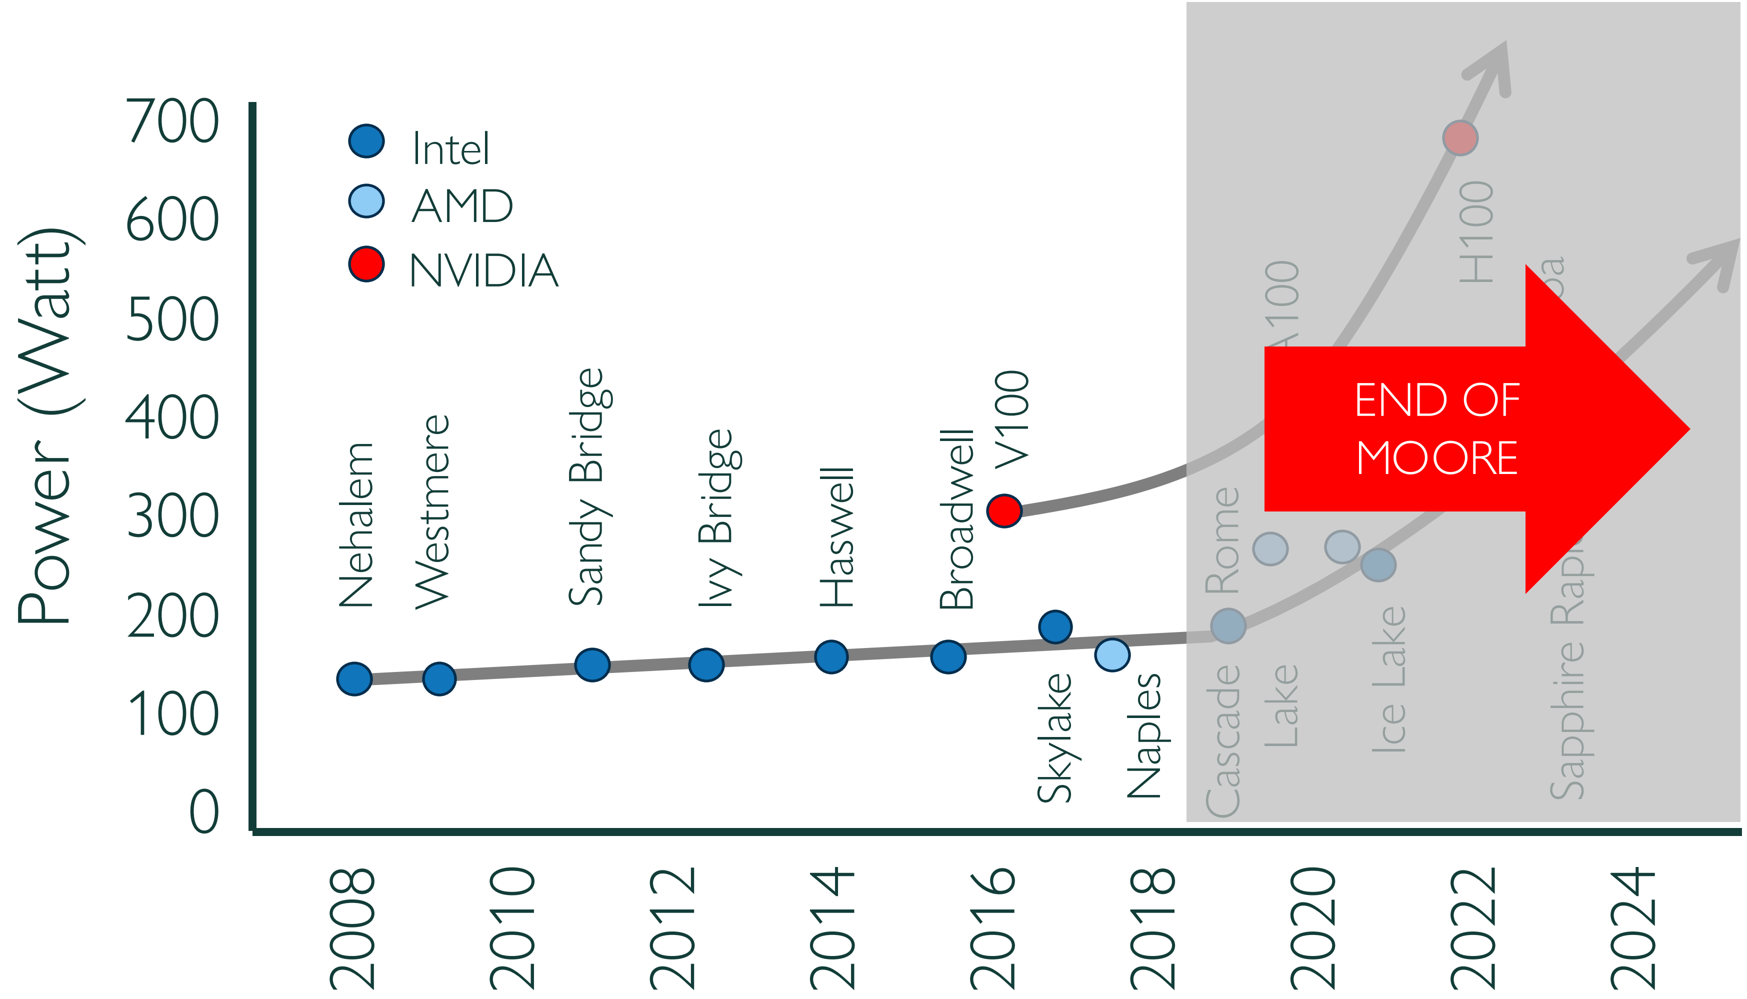

LLMS' GROWTH

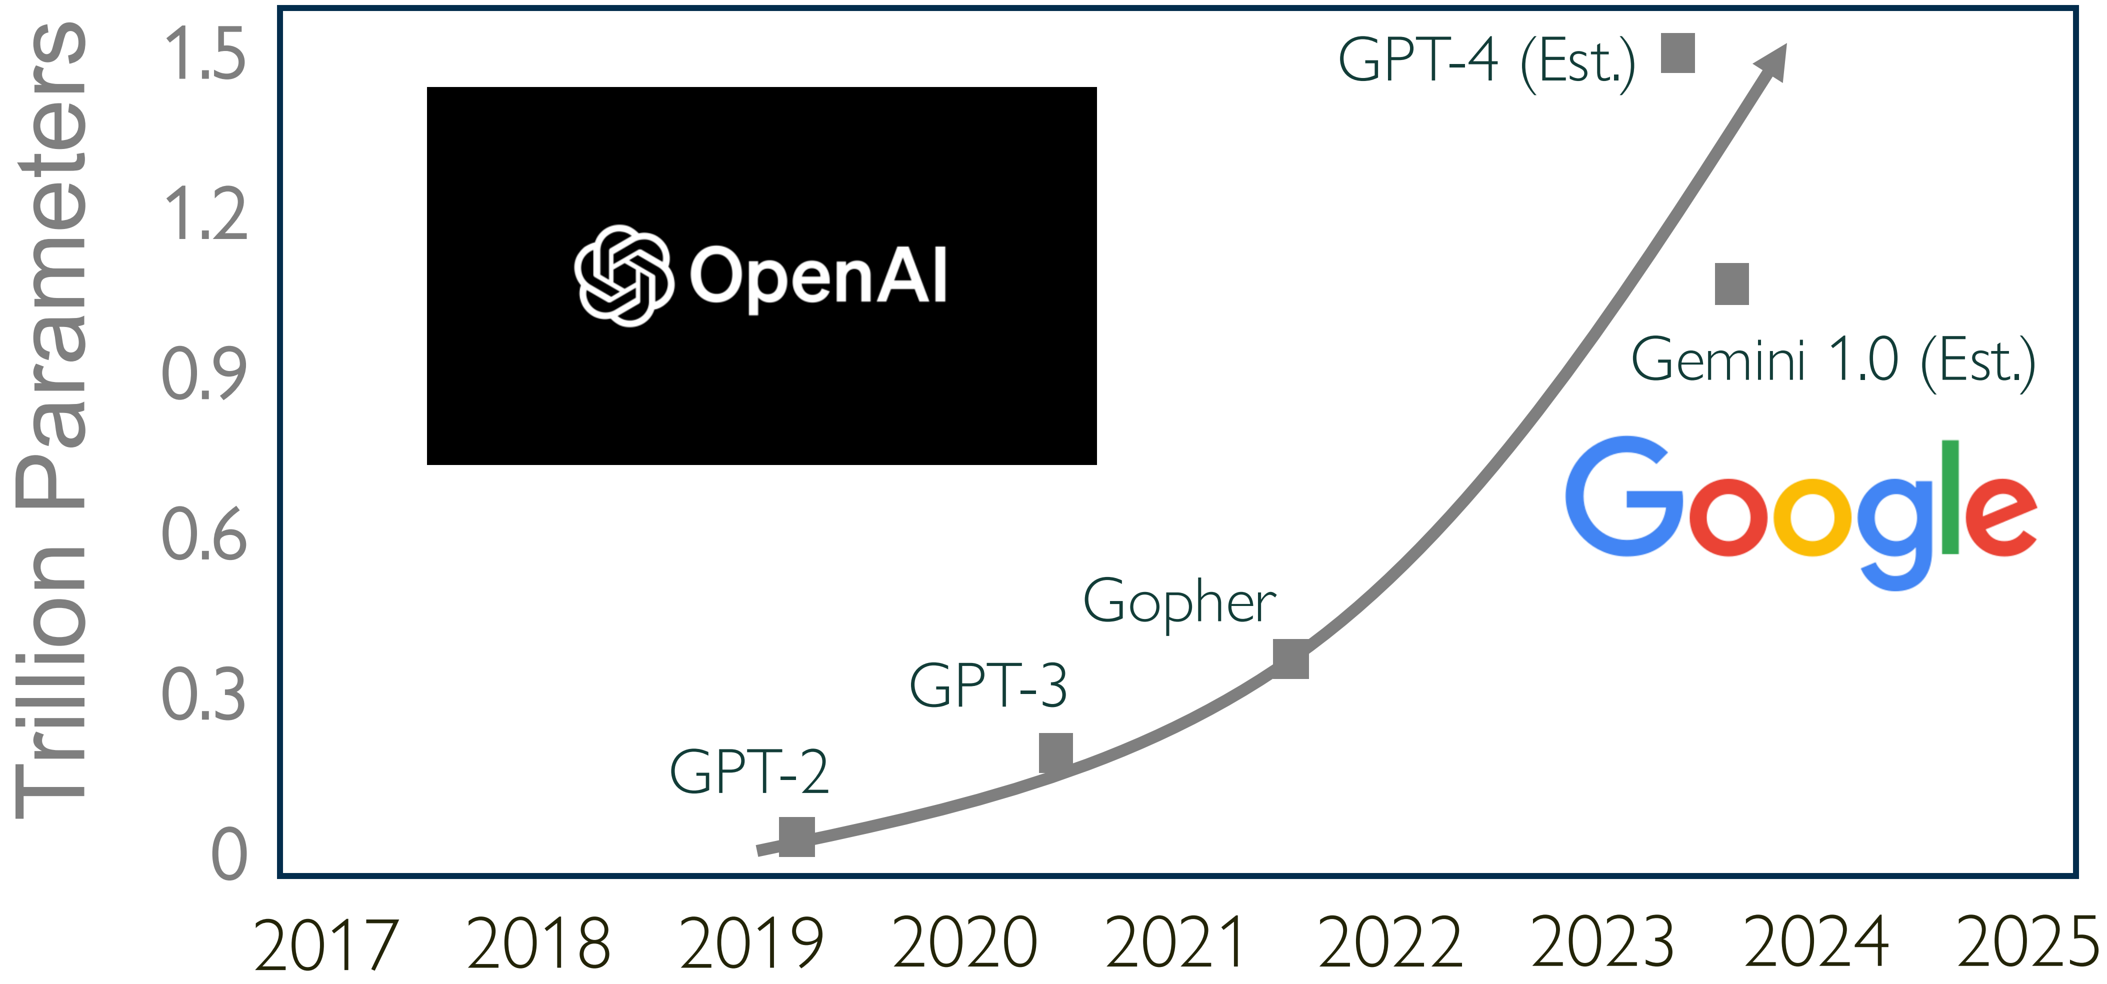
Trillion Parameters
1.5
1.2
0.9
0.6
0.3
0
OpenAI
GPT-4 (Est.)
Gemini 1.0 (Est.)
Google
Gopher
GPT-3
GPT-2
2017
2018
2019
2020
2021
2022
2023
2024
2025

NVIDIA CHIP EFFICIENCY

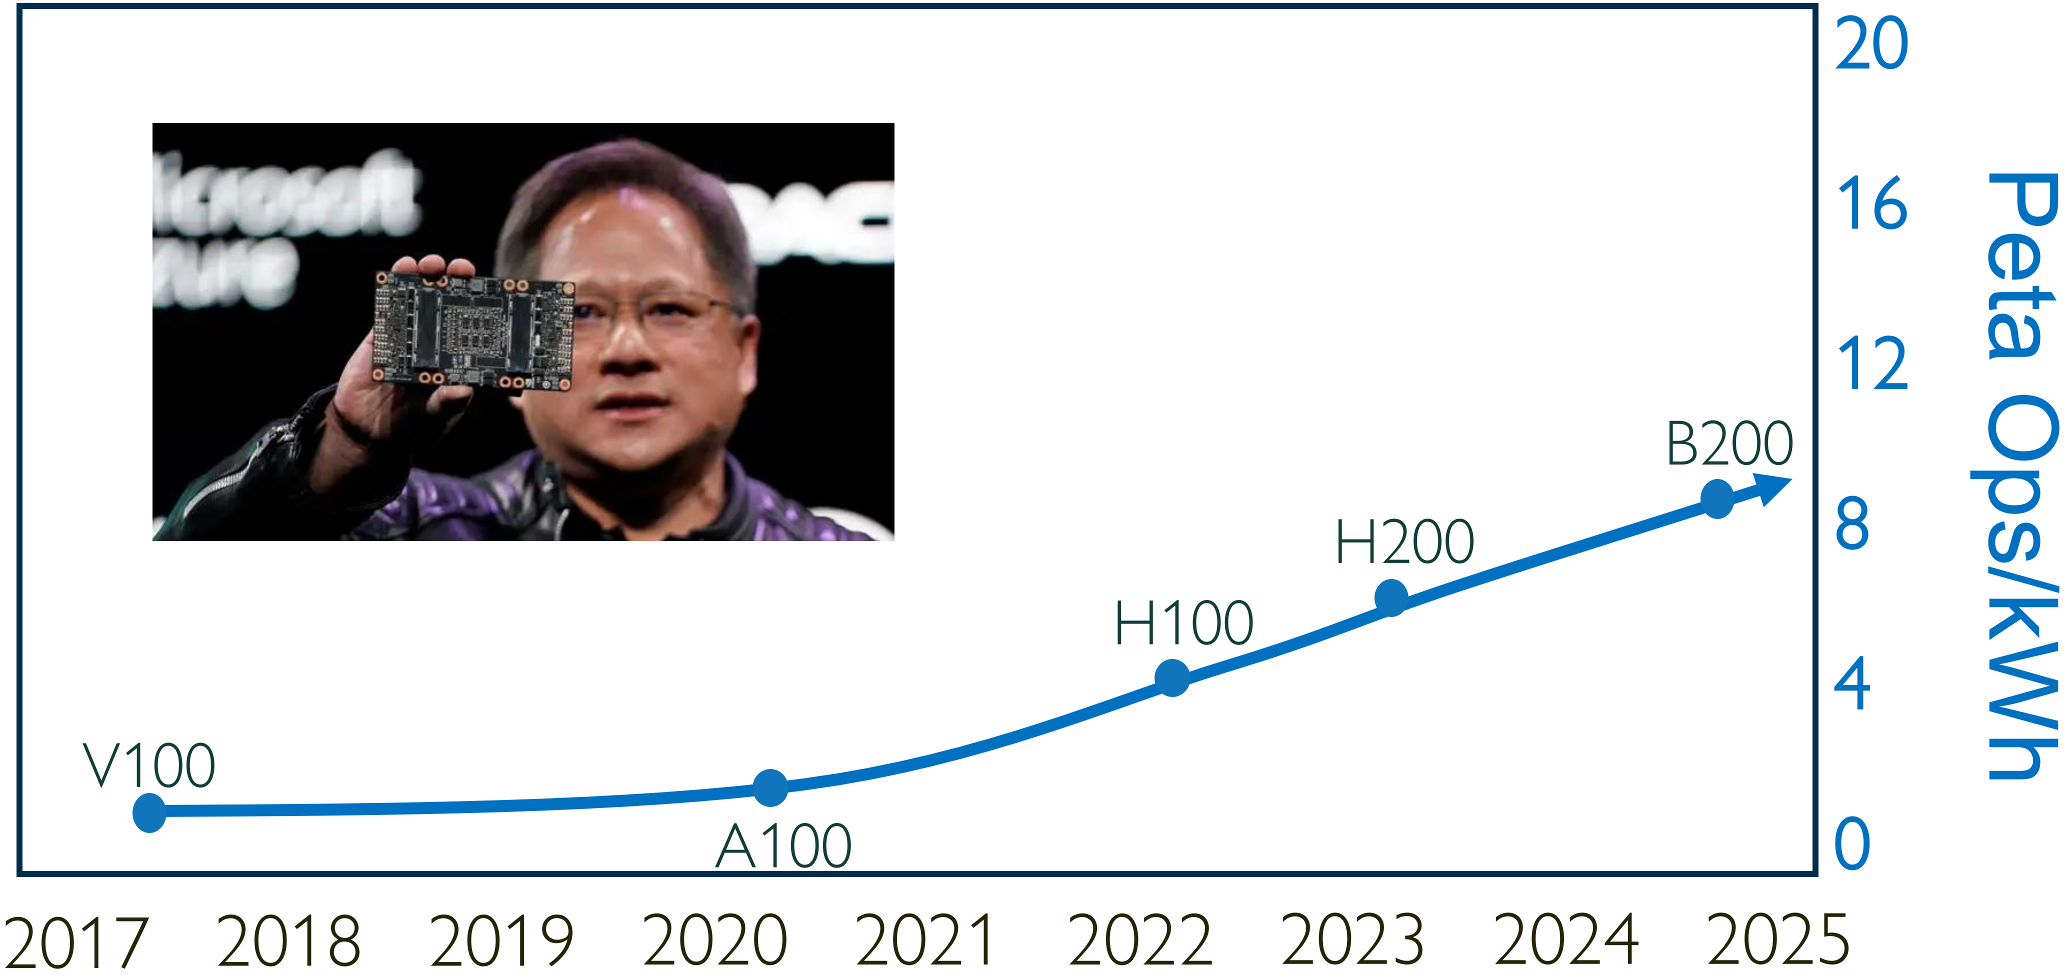

CATCH ME IF YOU CAN!

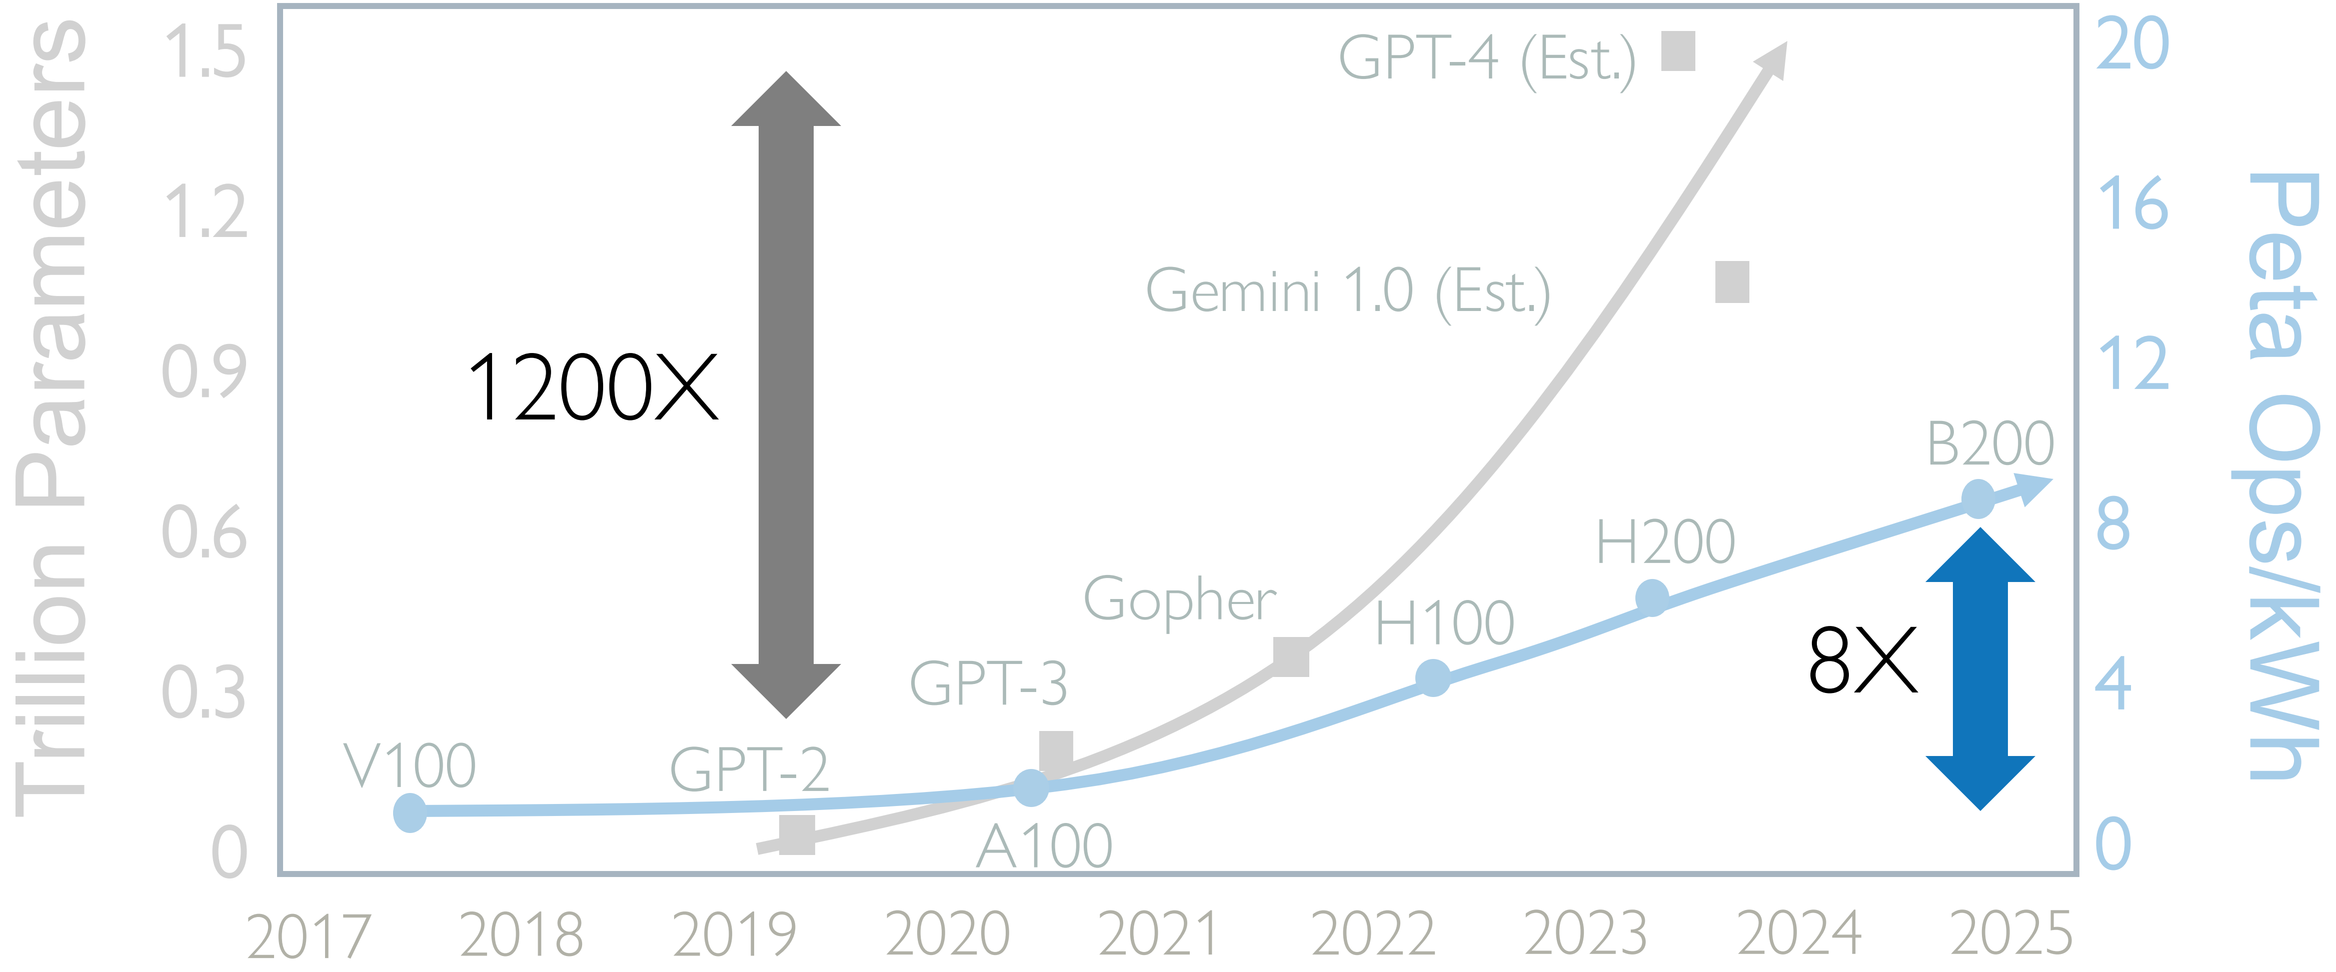

GROWTH IN DATACENTER ENERGY

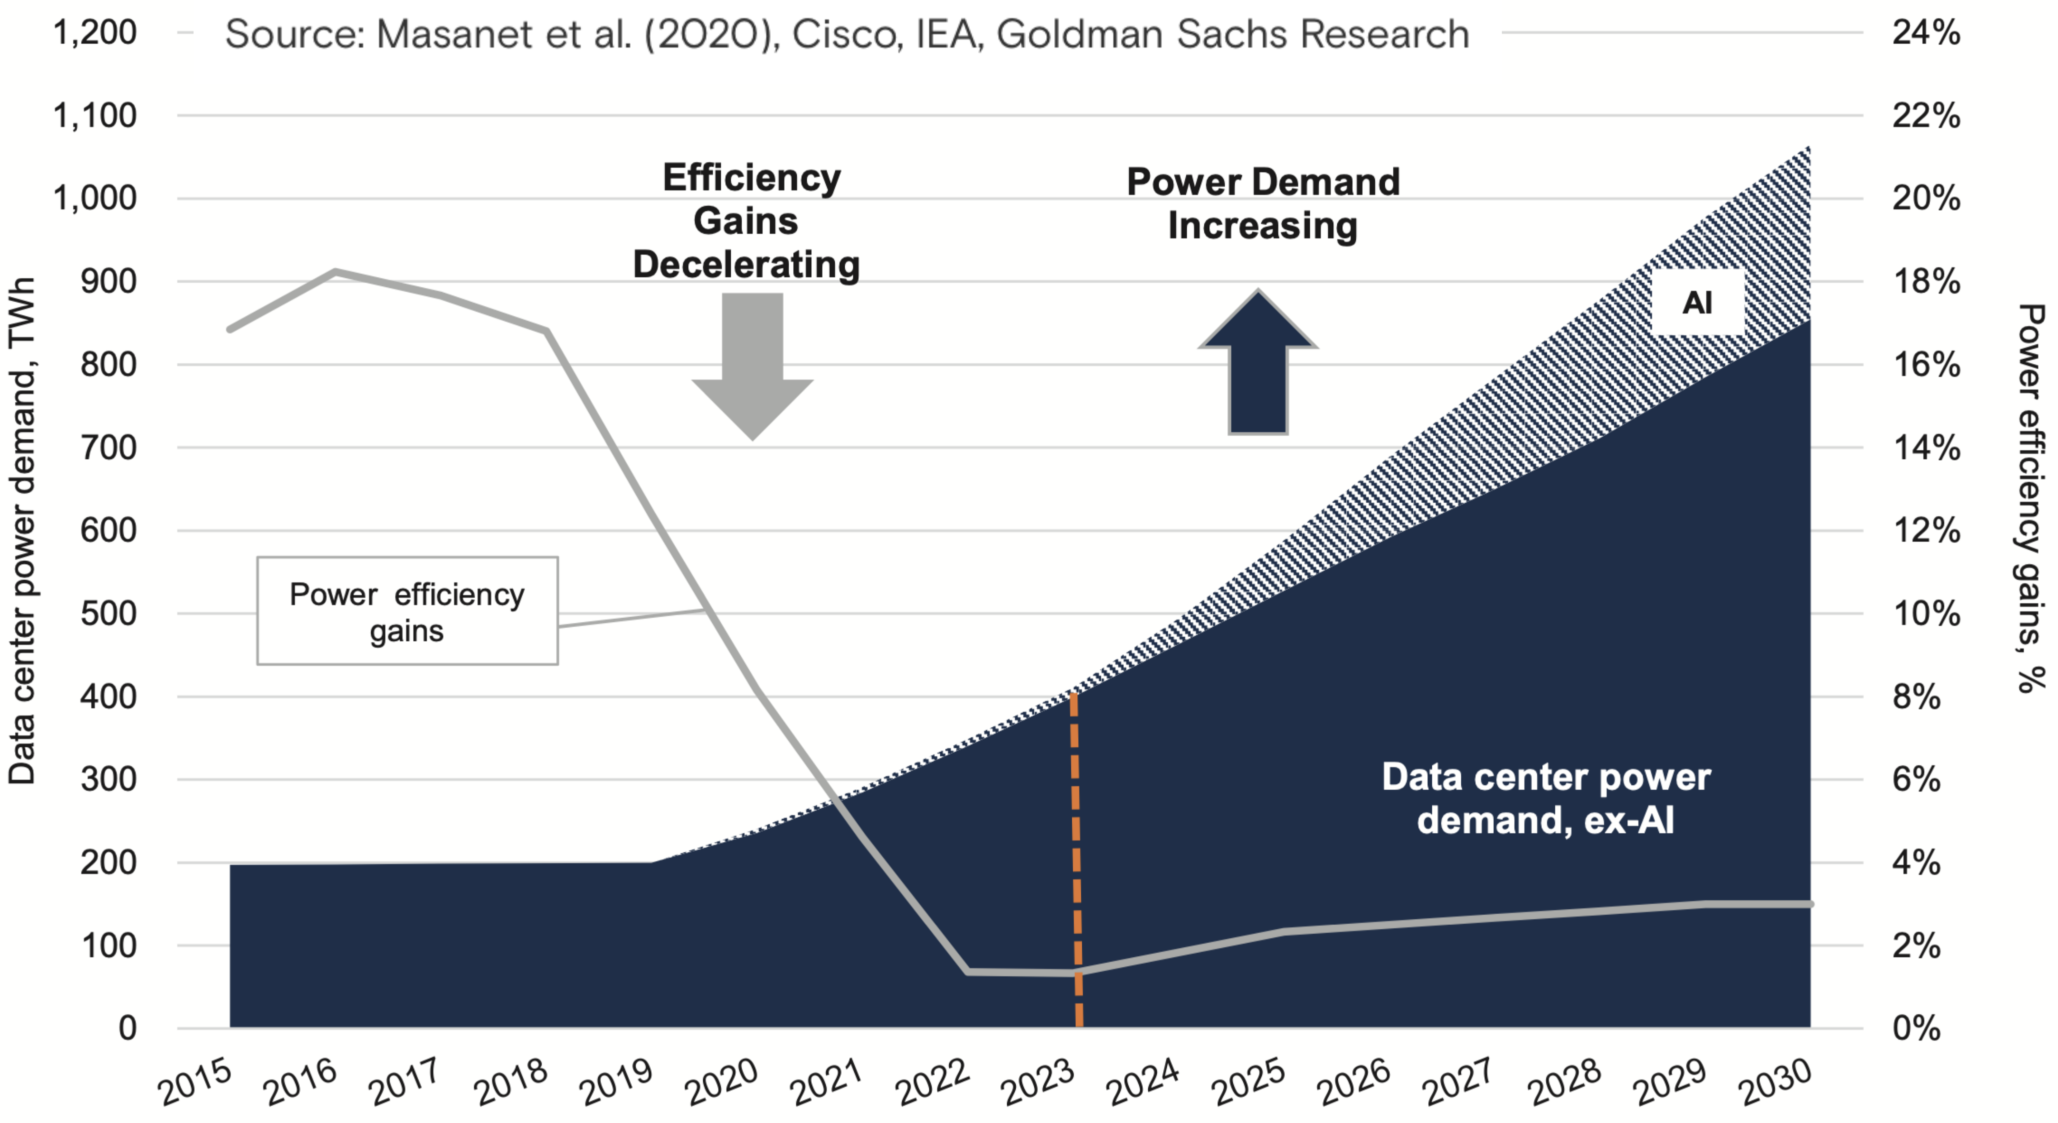

OPERATIONAL VS. EMBODIED EMISSIONS

EcoCloud

“The **use stage** GHG emissions in 2020 relating to electricity use account for the majority of **total GHG emissions**.”

- *Malmodin et al. (2020)*

OPERATIONAL EMISSIONS
Scope 1 & Scope 2

95 million tons CO_2

EMBODIED EMISSIONS
Scope 3

31 million tons CO_2

76%

24%

ENERGY VS. CARBON

What is our output?

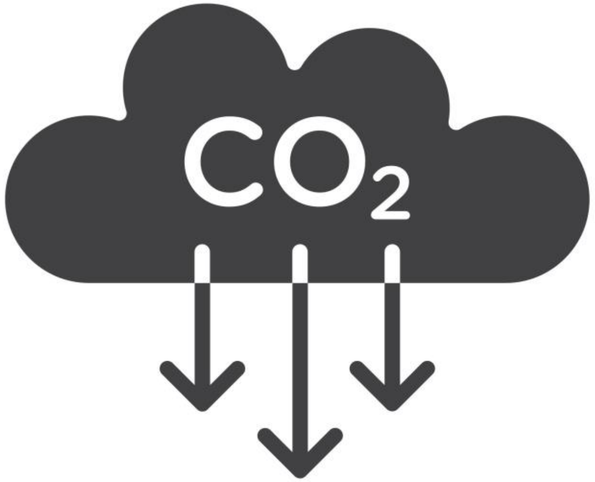

What is our environmental impact?

SUSTAINABILITY IN DATACENTERS

Metrics

Design

Methodologies

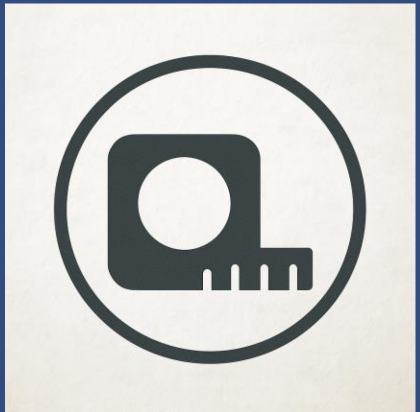

SUSTAINABILITY IN DATACENTERS

Metrics

Design

Methodologies

Today’s efficiency metric

power usage efficiency

$$PUE = \frac{\text{Total DC Power}}{\text{IT Power}}$$

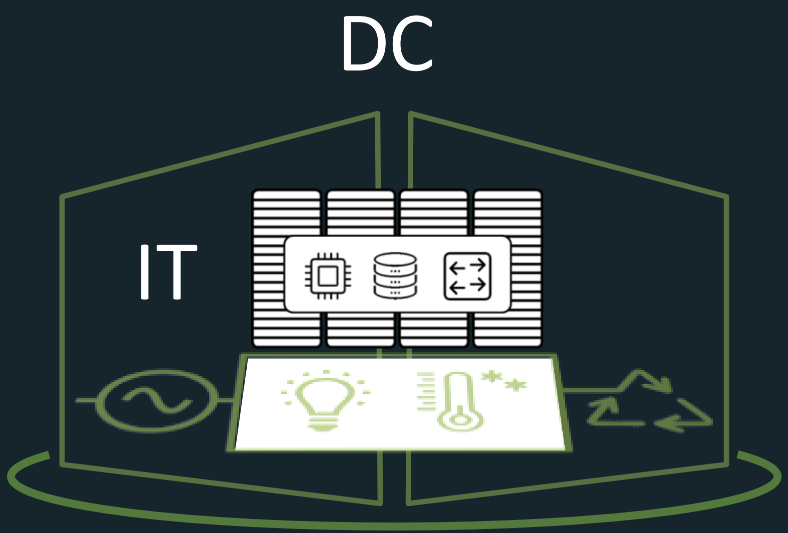

PUE has been around for two decades

INDUSTRY STANDARD: PUE

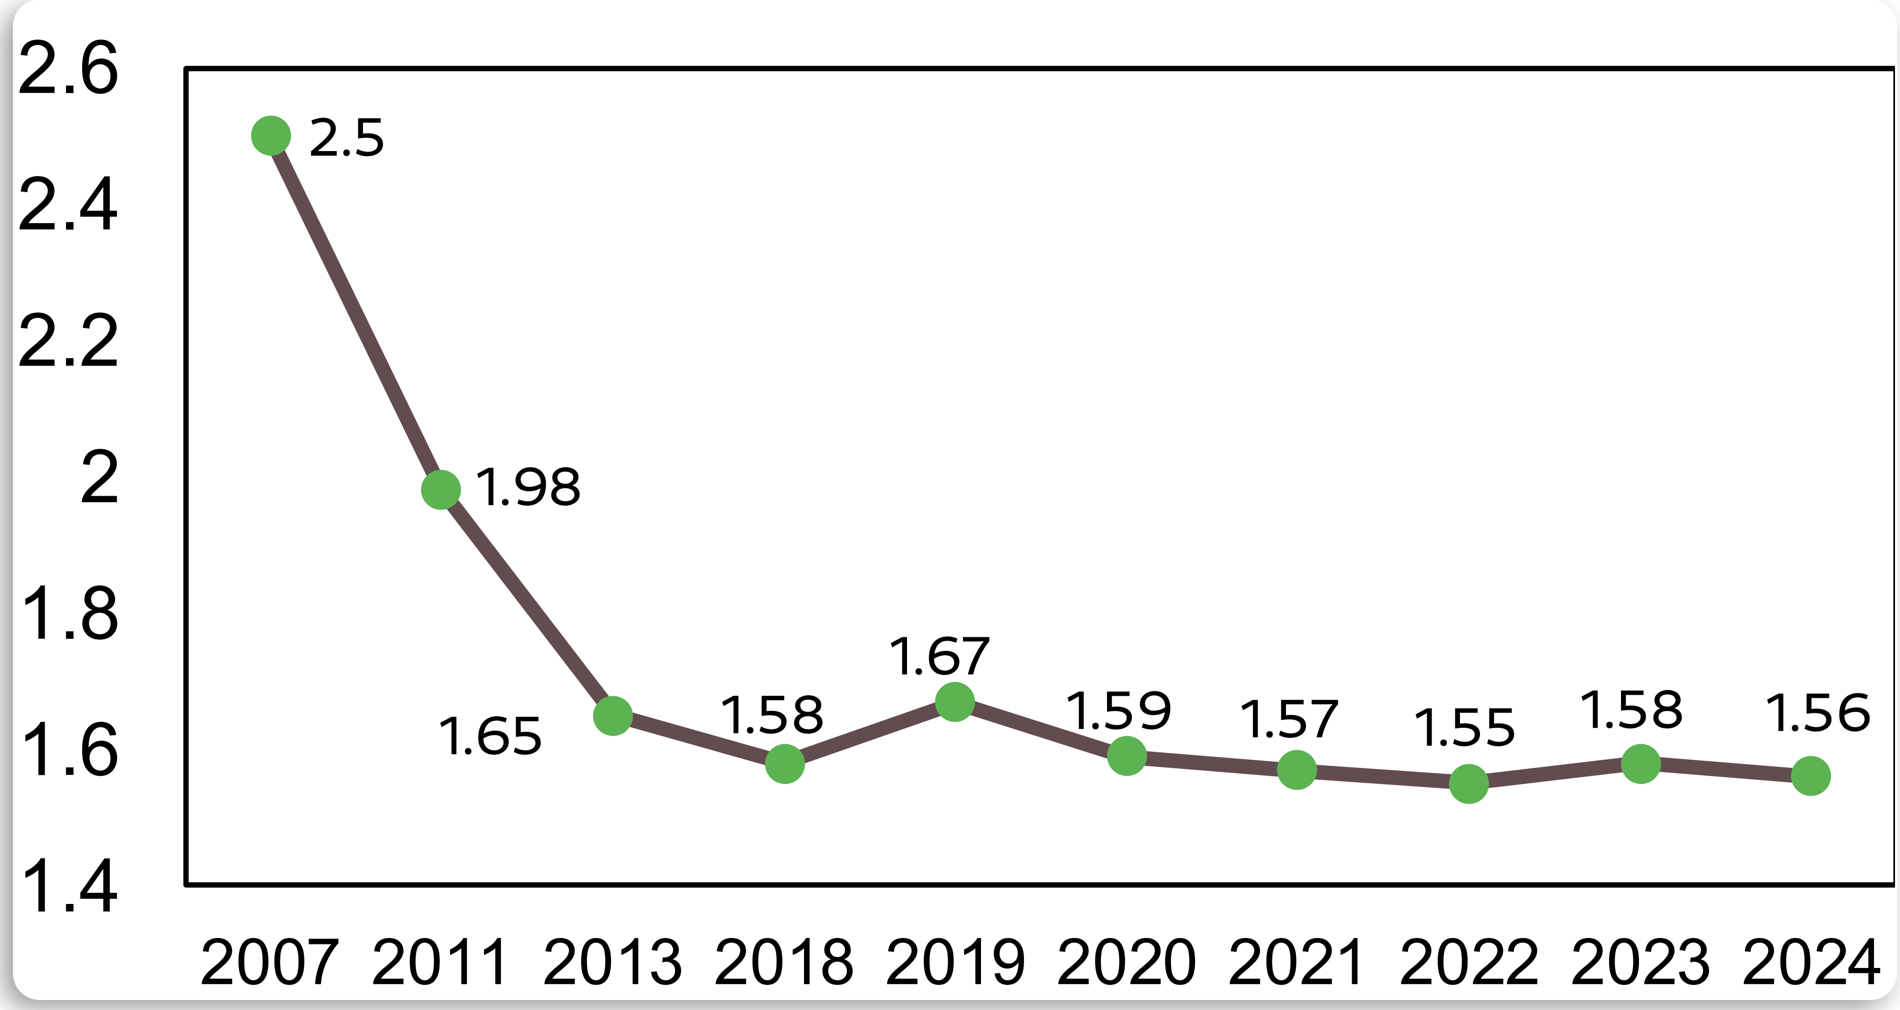

Global Average (2024): 1.56 *(= 64% of the electricity flows into IT)*

statista

LIMITS OF PUE

PUE IGNORES IT EFFICIENCY

INEFFICIENT OR UNDERUTILIZED SERVERS MAKE THE PUE LOOK GOOD

PUE: 1.2

PUE: 1.5

LIMITS OF PUE

PUE IGNORES IT EFFICIENCY

INEFFICIENT OR UNDERUTILIZED SERVERS MAKE THE PUE LOOK GOOD

PUE: 1.2

Av. Server Utilization: 15%

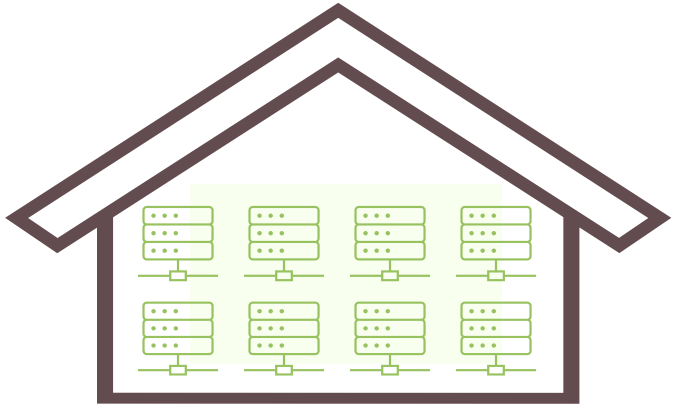

PUE: 1.5

LIMITS OF PUE

PUE IGNORES IT EFFICIENCY

INEFFICIENT OR UNDERUTILIZED SERVERS MAKE THE PUE LOOK GOOD

PUE: 1.2

Av. Server Utilization: 15%

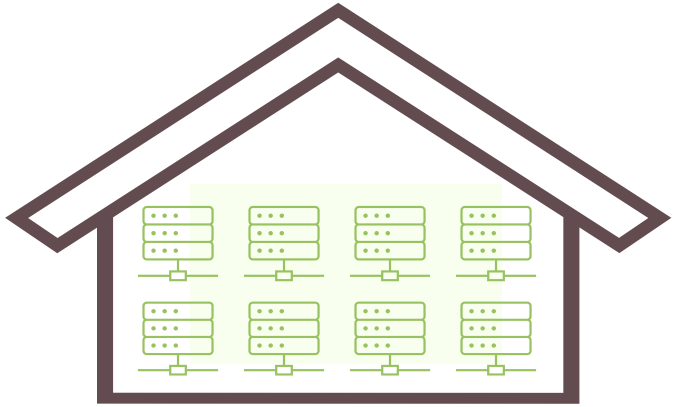

PUE: 1.5

Av. Server Utilization: 60%

LIMITS OF PUE

PUE IGNORES END-TO-END ENERGY FLOW

LIMITS OF PUE

PUE IGNORES END-TO-END ENERGY FLOW

PUE IGNORES HEAT RECYCLING OR ON-PREMISE RENEWABLE GENERATION

LIMITS OF PUE

PUE IGNORES END-TO-END ENERGY FLOW

PUE IGNORES HEAT RECYCLING OR ON-PREMISE RENEWABLE GENERATION

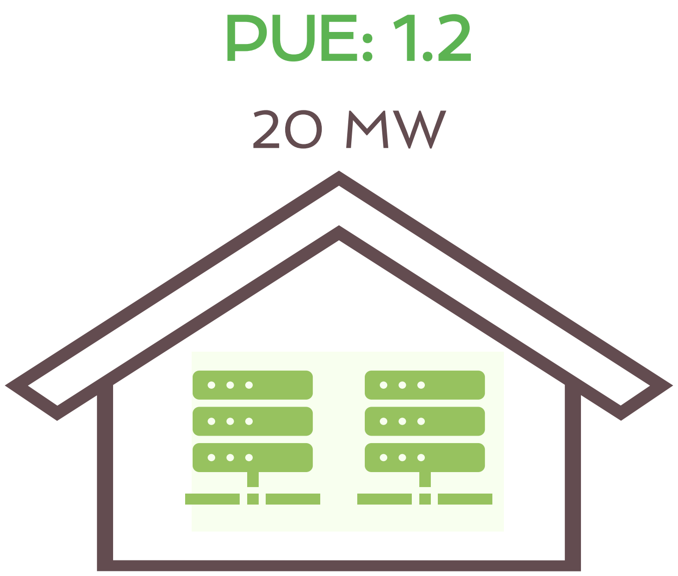

LIMITS OF PUE

PUE IGNORES END-TO-END ENERGY FLOW

PUE IGNORES HEAT RECYCLING OR ON-PREMISE RENEWABLE GENERATION

PUE: 1.2

20 MW

PUE: 1.5

20 MW

LIMITS OF PUE

PUE IGNORES END-TO-END ENERGY FLOW

PUE IGNORES HEAT RECYCLING OR ON-PREMISE RENEWABLE GENERATION

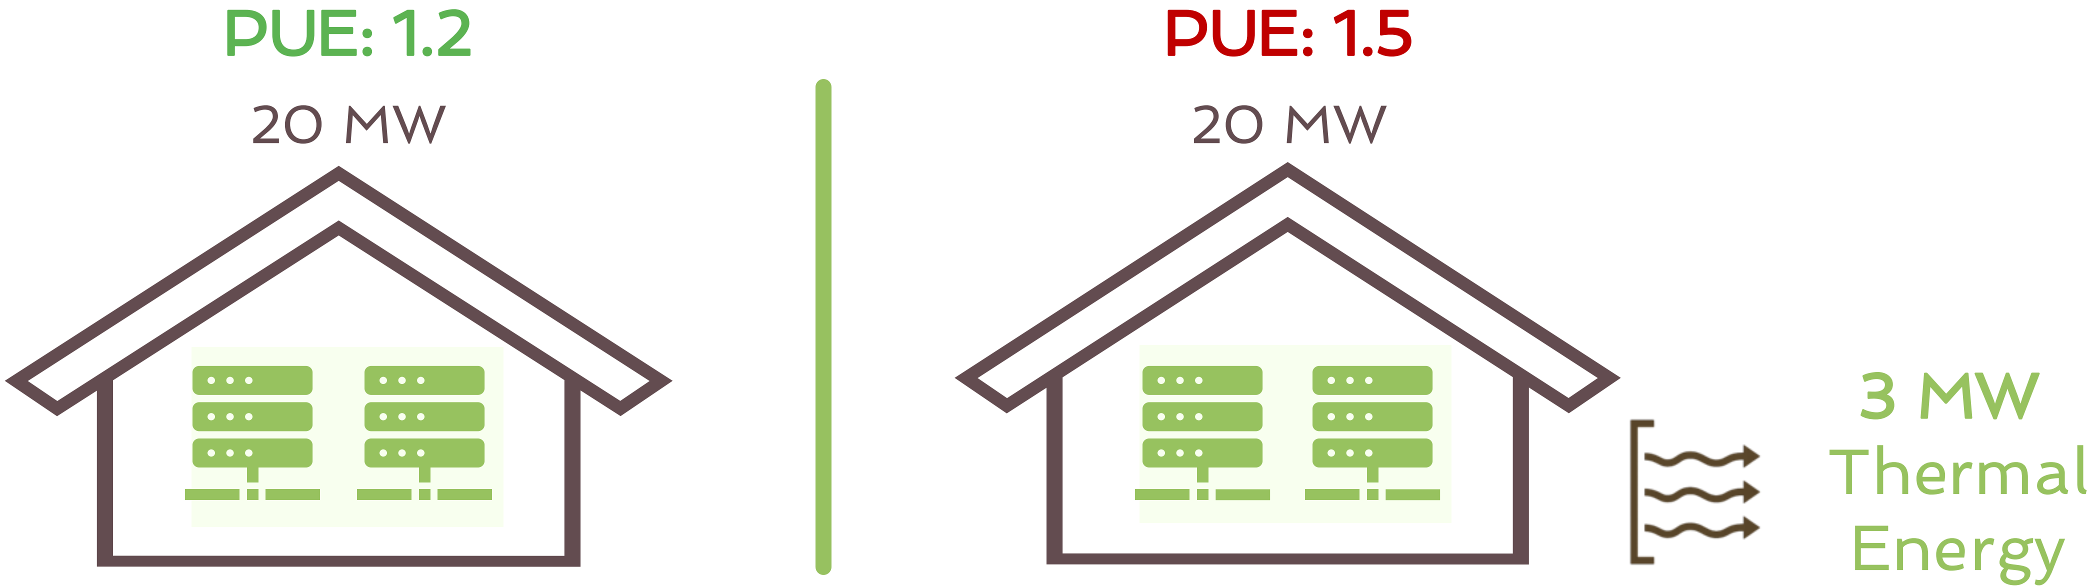

LIMITS OF PUE

PUE SAYS NOTHING ABOUT CO_2 EMISSIONS

LIMITS OF PUE

PUE SAYS NOTHING ABOUT CO_2 EMISSIONS

DATA CENTER EFFICIENCY

⚠ PUE IGNORES THE SOURCE OF ELECTRICITY

LIMITS OF PUE

PUE SAYS NOTHING ABOUT CO_2 EMISSIONS

PUE IGNORES THE SOURCE OF ELECTRICITY

PUE: 1.2

20 MW

PUE: 1.5

20 MW

LIMITS OF PUE

PUE SAYS NOTHING ABOUT CO_2 EMISSIONS

DATA CENTER EFFICIENCY

PUE IGNORES THE SOURCE OF ELECTRICITY

PUE: 1.2

20 MW

100% Coal Power

PUE: 1.5

20 MW

LIMITS OF PUE

PUE SAYS NOTHING ABOUT CO_2 EMISSIONS

PUE IGNORES THE SOURCE OF ELECTRICITY

PUE: 1.2

20 MW

100% Coal Power

140 k

tons CO_2/y

PUE: 1.5

20 MW

LIMITS OF PUE

PUE SAYS NOTHING ABOUT CO_2 EMISSIONS

DATA CENTER EFFICIENCY

PUE IGNORES THE SOURCE OF ELECTRICITY

PUE: 1.2

20 MW

100% Coal Power

140 k

tons CO_2/y

PUE: 1.5

20 MW

100% Renewables

LIMITS OF PUE

PUE SAYS NOTHING ABOUT CO_2 EMISSIONS

PUE IGNORES THE SOURCE OF ELECTRICITY

PUE: 1.2

20 MW

100% Coal Power

140 k

tons CO_2/y

PUE: 1.5

20 MW

100% Renewables

14 k

tons CO_2/y

IT SUSTAINABILITY METRICS

Need metrics to answer our 20 kW goes

- For a given output (e.g., LLM prediction)
  - How much electricity to we need?
  - How much silicon do we need?

- In practice
  - Hyperscalars build w/ commodity parts
  - Maximize utilization
  - (Other) IT operators don't know

RACK-LEVEL OPERATIONAL ENERGY

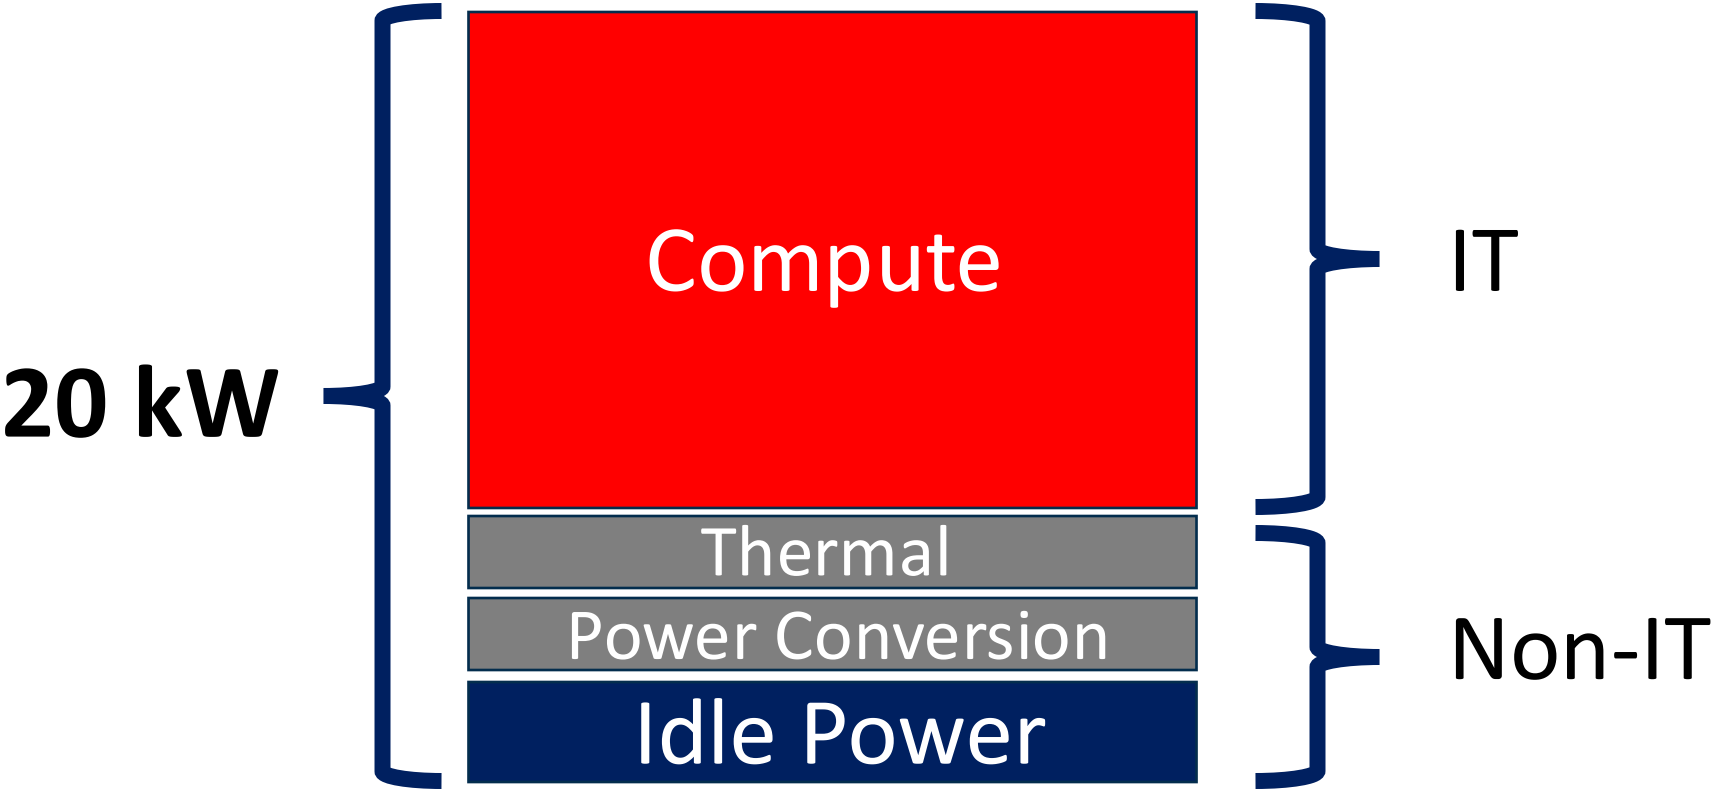

DC Efficiency Metrics Workstream, Open Compute Project (OCP)
EMEA Summit, April 29, 2025

COMPUTE EFFICIENCY METRICS

- Throughput/W (operational energy)
  - Logic dominates power
  - Air-cooled chips are power-bound (e.g., 0.4 W/mm^2)
  - We pack chips with (not useful) dark silicon (mostly SRAM)

- Throughput/mm^2 (embodied energy)
  - Liquid-cooled chips become area-bound (e.g., 1.2 W/mm^2)
  - What are the metrics to provision SRAM for area-bound chips?

- What other computational metrics?
  - Compute, memory, network, storage

OTHER METRICS

- Recycled heat
- Renewable energy
- Input/output water
- Carbon metrics
- ….

SUSTAINABILITY IN DATACENTERS

Metrics

Design

Methodologies

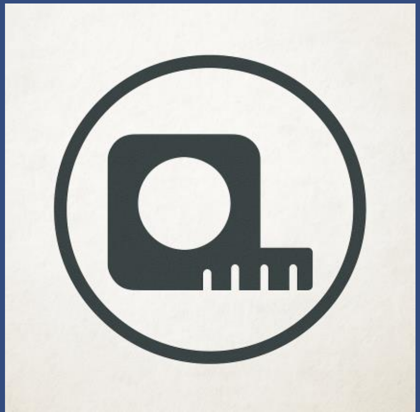

SUSTAINABILITY IN DATACENTERS

Metrics

Design

Methodologies

POST-MOORE CLOUD RACKS

Rack as an SKU with "ISA"

- Integration
  - reduce data movement
- Specialization
  - cut resources to analyze data
- Approximation
  - compress data & computation

From algorithms to infrastructure

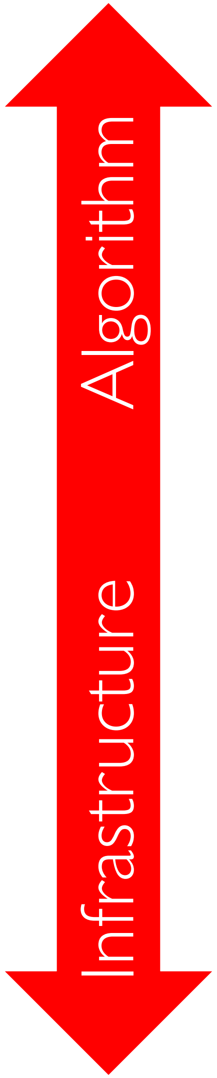

CENTER @ EPFL SINCE 2011

Mission
- Sustainable computing
- IT for sustainability
- Best practices, metrics & methodologies

Impact
- Server-grade ARM CPU
- Cloud-native network/database stacks
- Liquid-cooling from chip to rack

ecocloud.ch

TODAY’S SERVER = 90’S DESKTOP PC

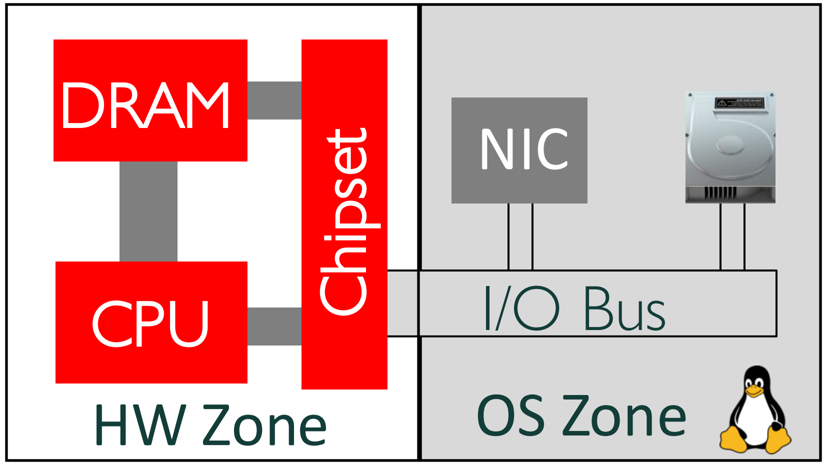

90’s Desktop PC

TODAY'S SERVER = 90'S DESKTOP PC

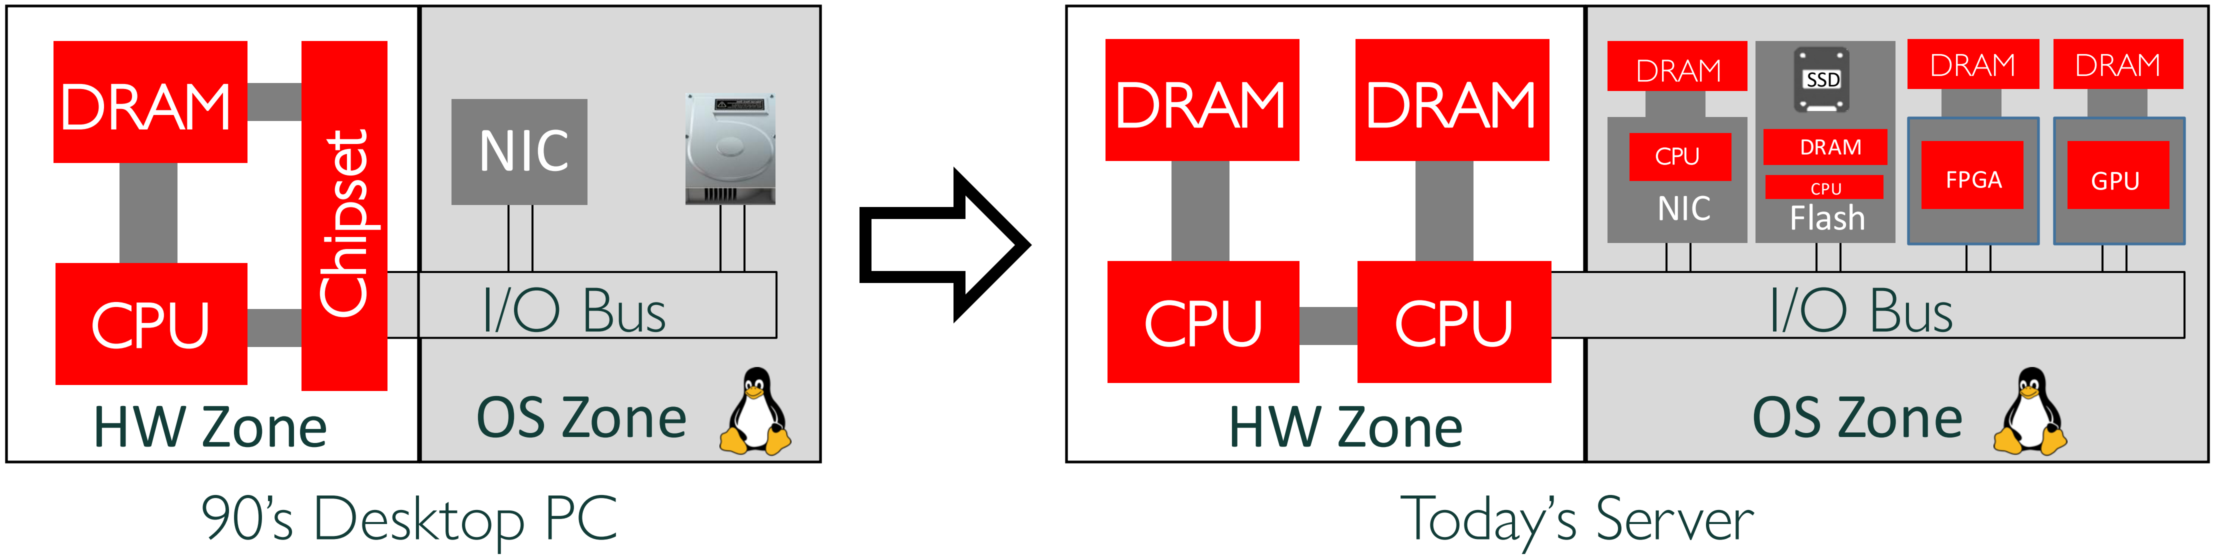

TODAY'S SERVER = 90'S DESKTOP PC

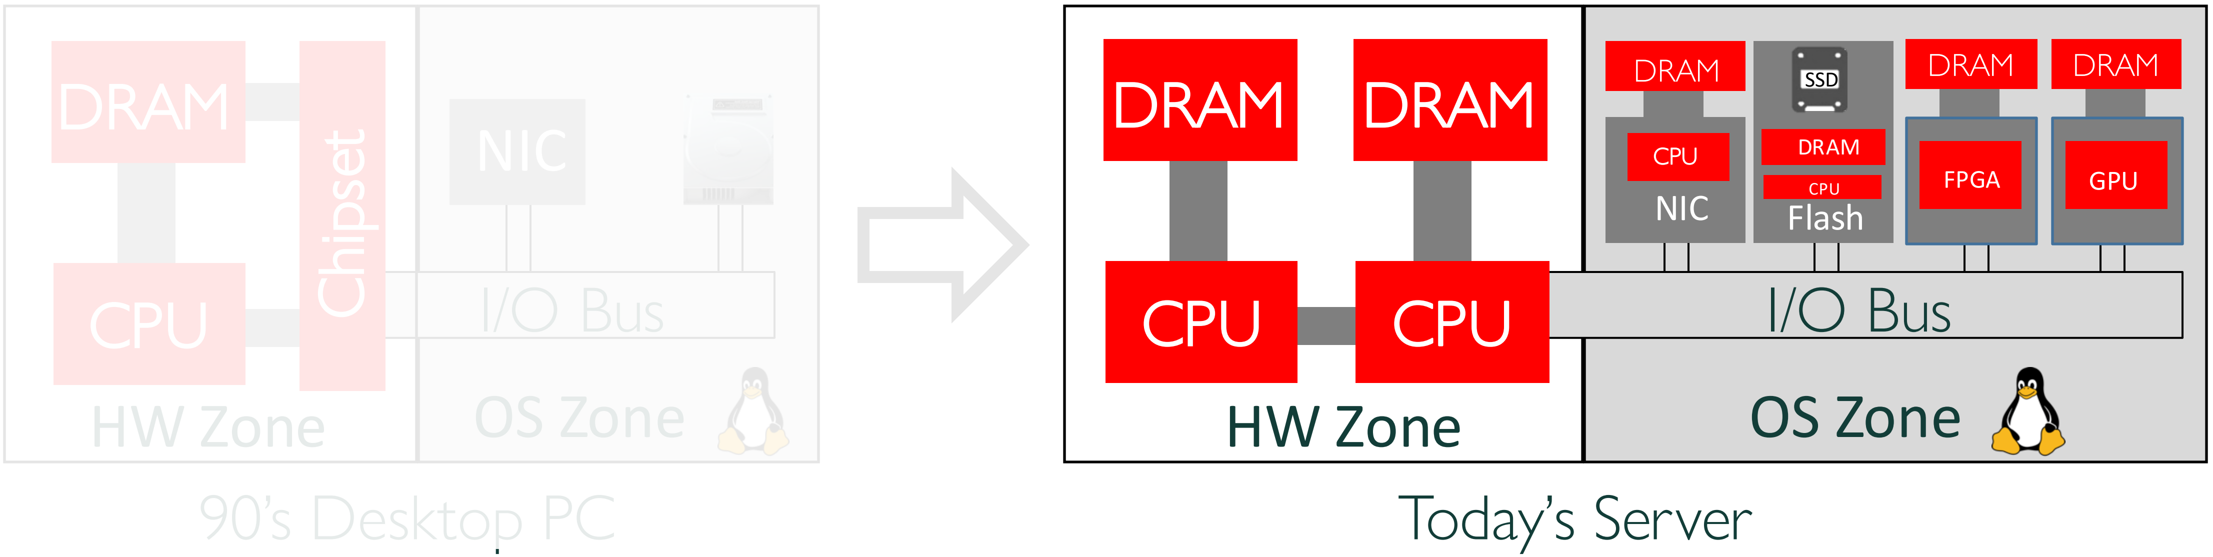

- Focused on minimizing cost (Google c.a. 2000)
- CPU, memory = nanosecond timescale, OS, I/O = millisecond timescale
- OS follows legacy interfaces (PCIe) and abstractions (POSIX)
- Silicon fragmented across legacy interfaces

EFFICIENCY PROBLEMS IN IT STACKS

Hardware/workload mismatch (EPFL, Meta, Google)

Datacenter tax ~ 20% (Google)

- 20,000 threads running per CPU
- Virtualization/containerization/FaaS using POSIX
- RPC

Memory wasted (Microsoft)

- 50% of containers do not use their memory
- 20% of memory is stranded

GPU utilization for deep learning < 50% (Microsoft)

AIR-COOLED POST-MOORE CPUS

[ISCA'12, ISCA50 Retrospective, IEEE Micro'24]

Today's server CPUs

- ✗ Designed for single-core performance
- ✗ Power-bound → ½ big cores + ½ memory
- ✗ Run at high frequency (power cubic with f ~ cv^2f)

AMD Zen 3
4.0 mm²
3.7W @ 3 GHz

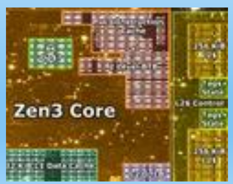

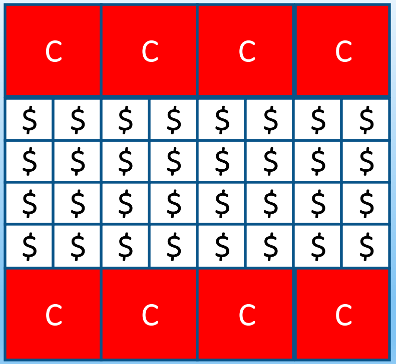

Cloud-native CPUs

- ✓ Low-width cores
- ✓ Recover power cubically w/ lower frequency
- ✓ Maximize throughput/W & number of cores
- ✓ Need only memory for per-core working set

ARM N1
1.4 mm²
0.7W @ 2 GHz

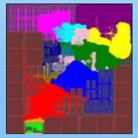

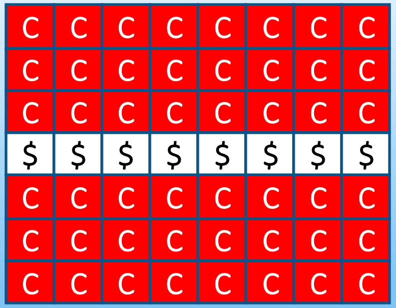

10X higher throughput with SLO!

SUSTAINABILITY IN DATACENTERS

Metrics

Design

Methodologies

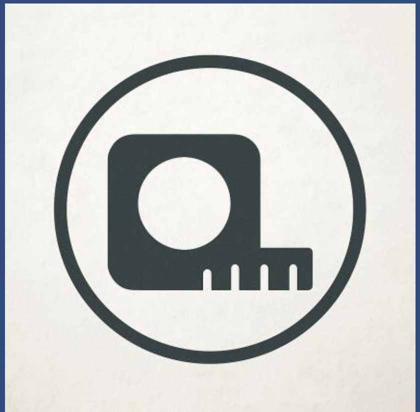

SUSTAINABILITY IN DATACENTERS

Metrics

Design

Methodologies

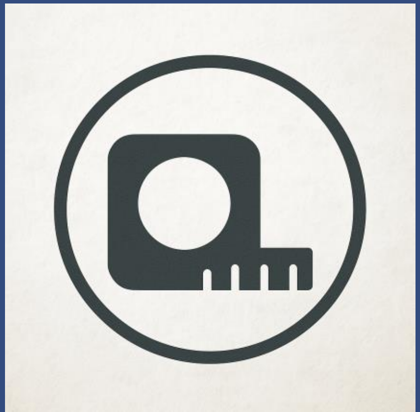

ADVERTISED VS. MEASURED

- Uptime reports [Feb. 2024]
  - Hyperscalars at PUE = 1.2
  - Average DC worldwide at PUE ~ 1.6
  - Reducing PUE from 1.5 to 1.3 is much easier than 1.3 to 1.1

Need proper measurement:

  - PUE highly varies over 12 months
  - Most builders/operators report “design” PUE

MEASURE FULL-STACK EFFICIENCY

DC EFFICIENCY

- electricity w/ renewables, cooling, heat recycling

IT EFFICIENCY

- compute, storage, network and workloads

CARBON FOOTPRINT

- emissions from input electricity sources

DATA CENTER

EFFICIENCY

sdea.ch

EPFL EcoCloud Hewlett Packard Enterprise asut Schneider Electric SWISS DATACENTER ASSOCIATION energieschweiz HOCHSCHULE LUZERN

SDEA NAVIGATOR

navigator.sdea.ch

PUE+	ITIE	CO_2
DC INFRASTRUC.	IT INFRASTRUC.	CO_2 FOOTPRINT

Digital Realty

100 MW
Zurich DC

awarded

DATA CENTER
EFFICIENCY
GOLD Plus

DATA CENTER
EFFICIENCY
SILVER Plus

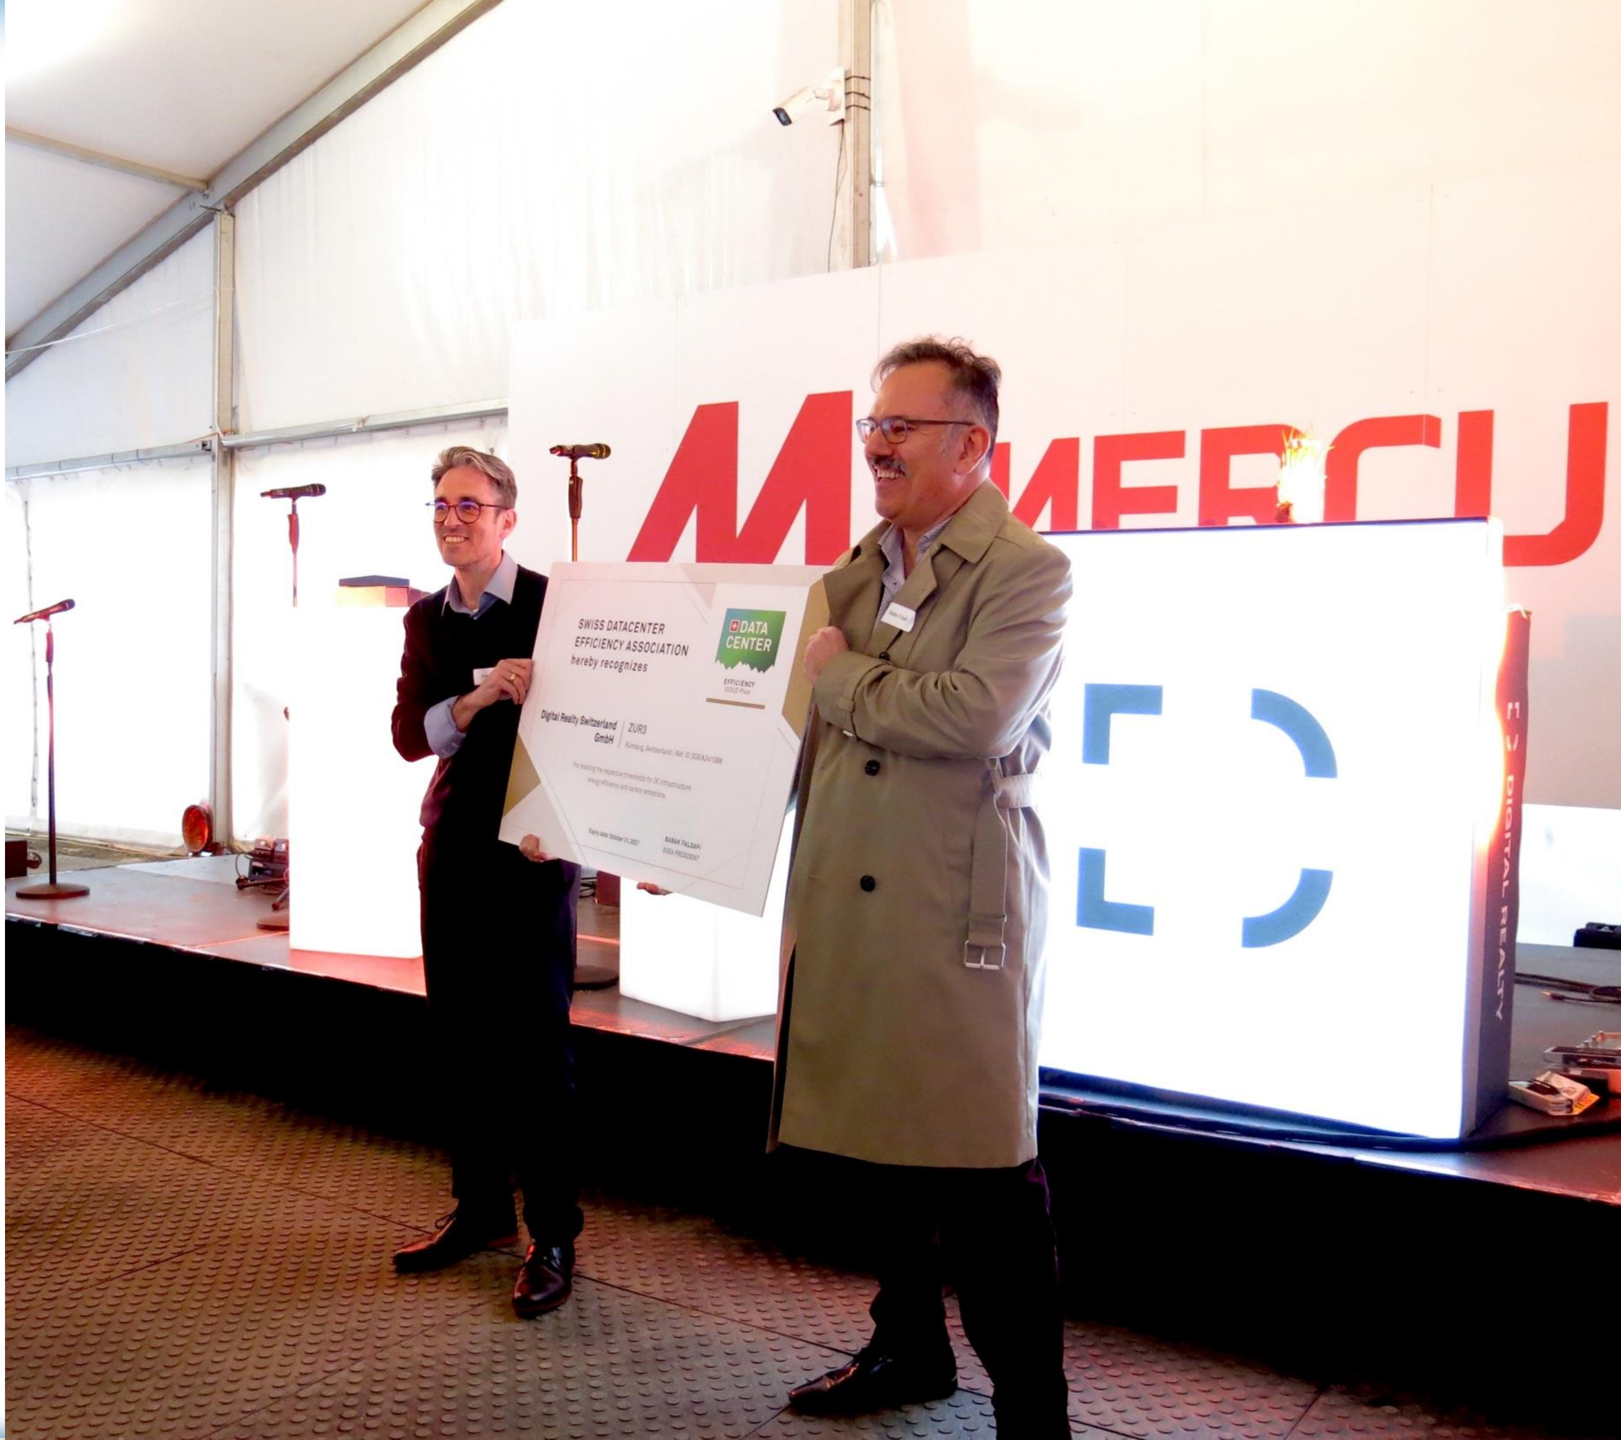

AI SUSTAINABILITY CLASSIFICATION

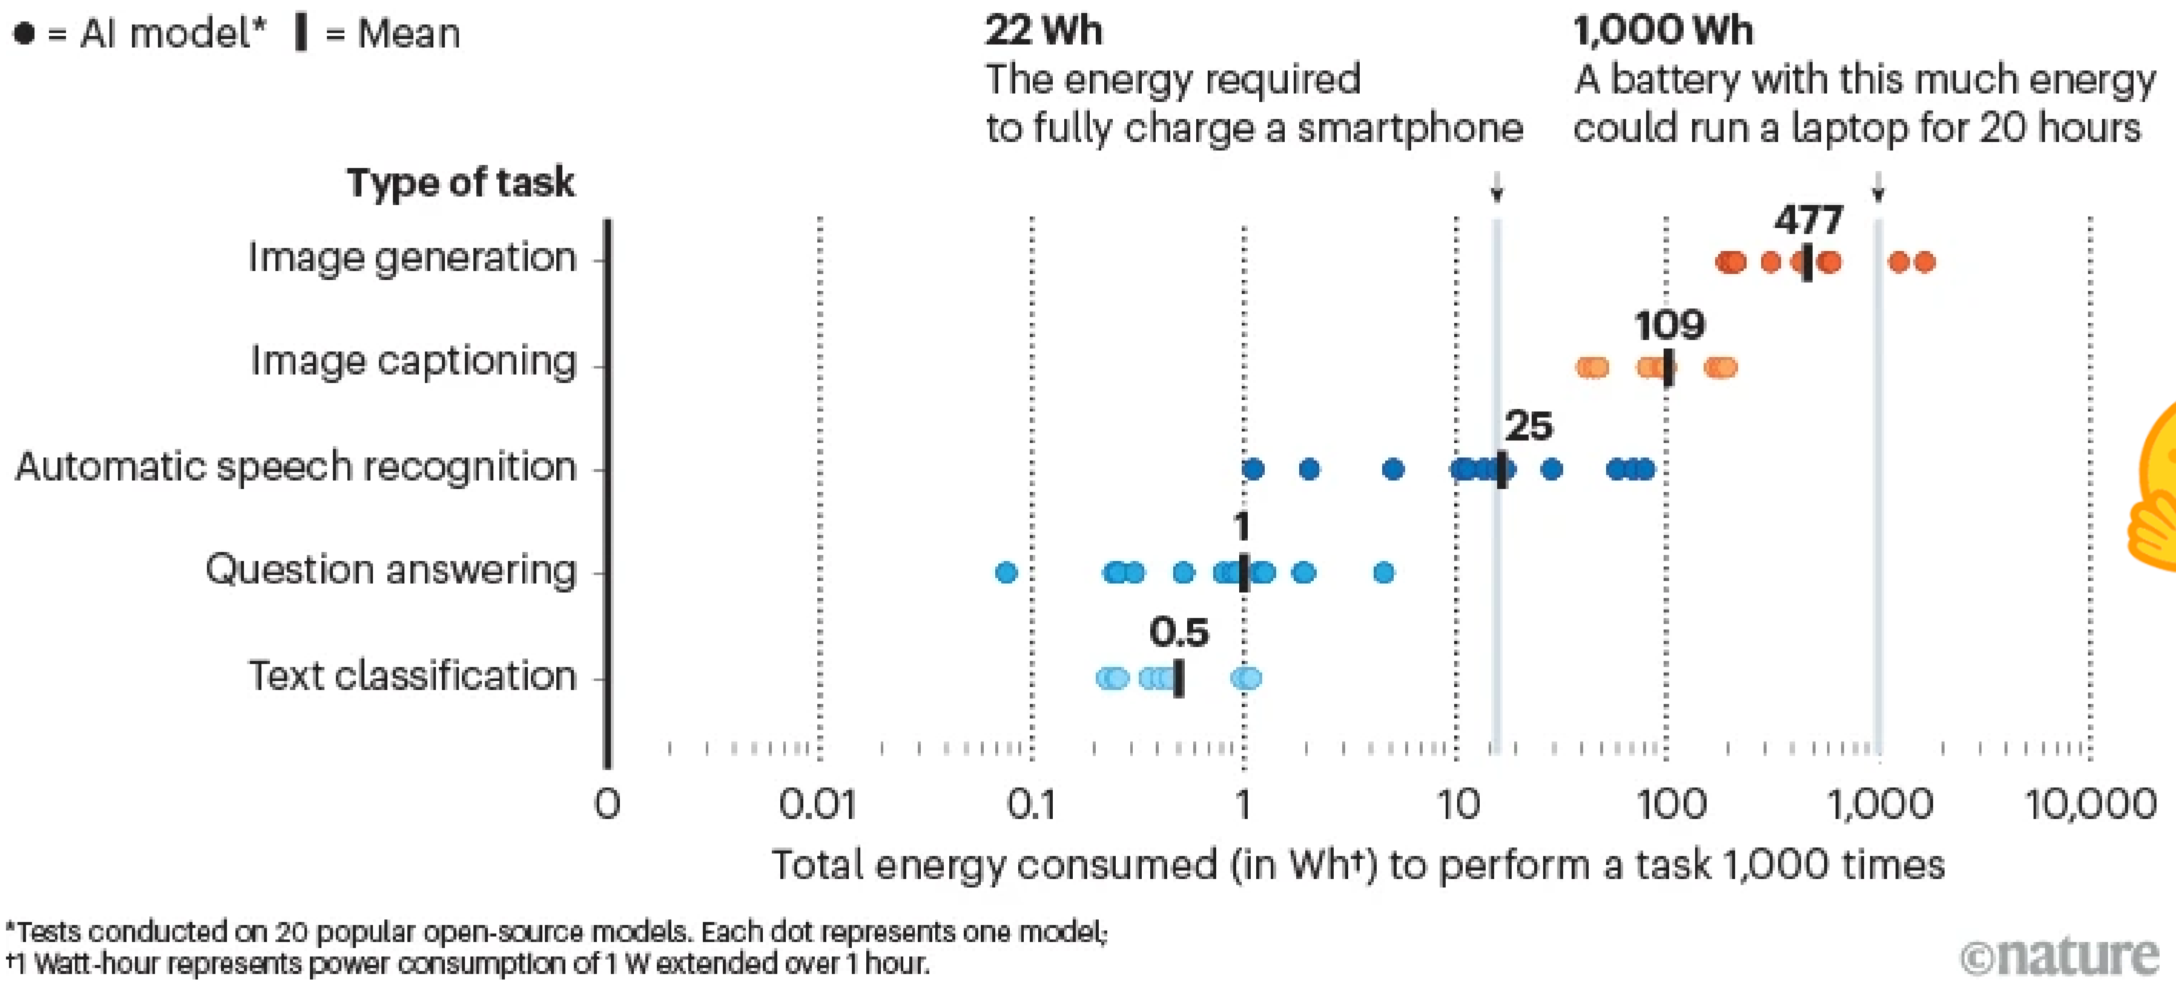

SUMMARY

DC energy consumption is growing at 16%/year

Moore's Law of silicon is dead

Need metrics & methodologies for efficiency

Need post-Moore DC design w/ "ISA":

- Integration + Specialization + Approximation

THANK YOU!

For more information, please visit us at

ecocloud.ch
sdea.ch