Green Data Center

Sabita Maharjan

Professor, University of Oslo, Norway

Feb. 23, 2023

Learning Objectives

Throughout this lecture, it is aimed for the students to be able to

- **Learn the basic concepts and architecture of data centers**
- **Learn about energy consumption and energy efficiency metrics for data centers**
- **Understand different techniques to reduce energy consumption in data centers with a vision to realize green data centers**

Industry Invited Talk Today

Speaker

Tor Kristian Gyland

Executive Consultant

Previously- CEO (Chief Executive Officer)

Green Mountain

Title

Green Mountain: Industry perspective in designing and operating green datacenter

Green Mountain

One of the largest operators in the data center sector in Norway with clients in Finance, IT, Government, Health, etc

DNB til Green Mountain

Banken utvider avtalen med datasenterselskapet og flytter store deler av sin it-plattform inn i fjellet. Verdien ligger på omlag 150 millioner kroner.

KENNETH CHRISTENSEN | itbransjen
kenneth.christensen@cw.no
Publisert: onsdag 15. april 2015, kl. 10:47 • Endret: 15. april 2015, kl. 10:48

Outline

Data Center Overview

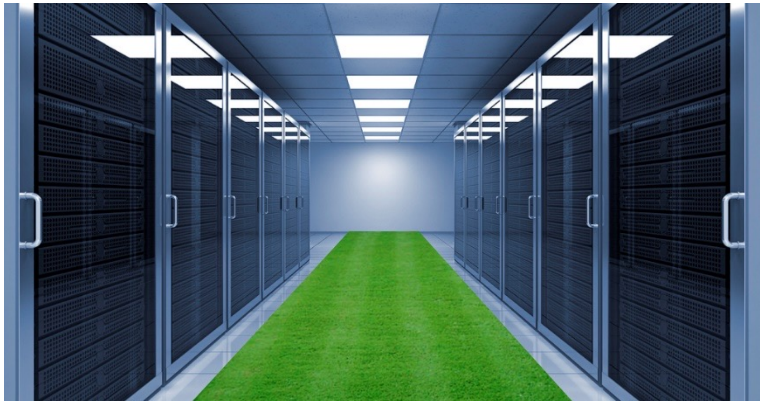

Green Data Center

DATA CENTER BASICS

Daily services supported by data centers

- Search
- Gmail
- Driving & Navigation
- Youtube
- Facebook, Whatsapp, Messenger
- Online classrooms

Google Data Centre at Mayes County, Oklahoma

Q: more?

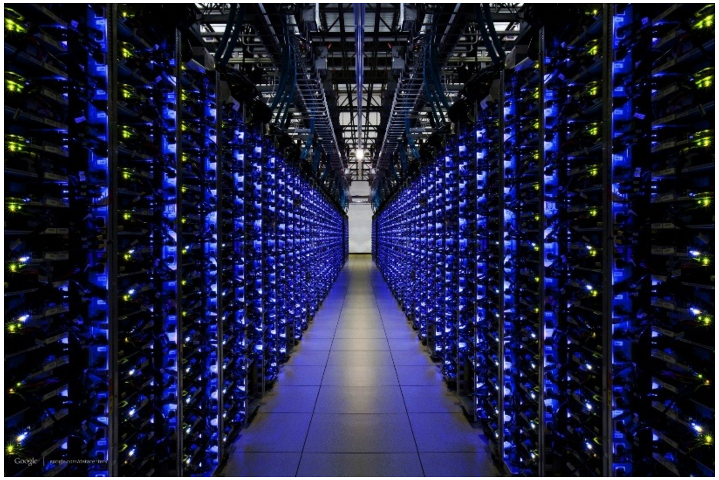

Inside view of a data center

Google data centers worldwide

Norway is very active in the data center business.

In 2018, the Norwegian government published a plan to boost Norway's data center industry and make the country a world-class player in the sector

Many data center players now:

Green Mountain

DigiPlex

LEFDAL MINE DATACENTER

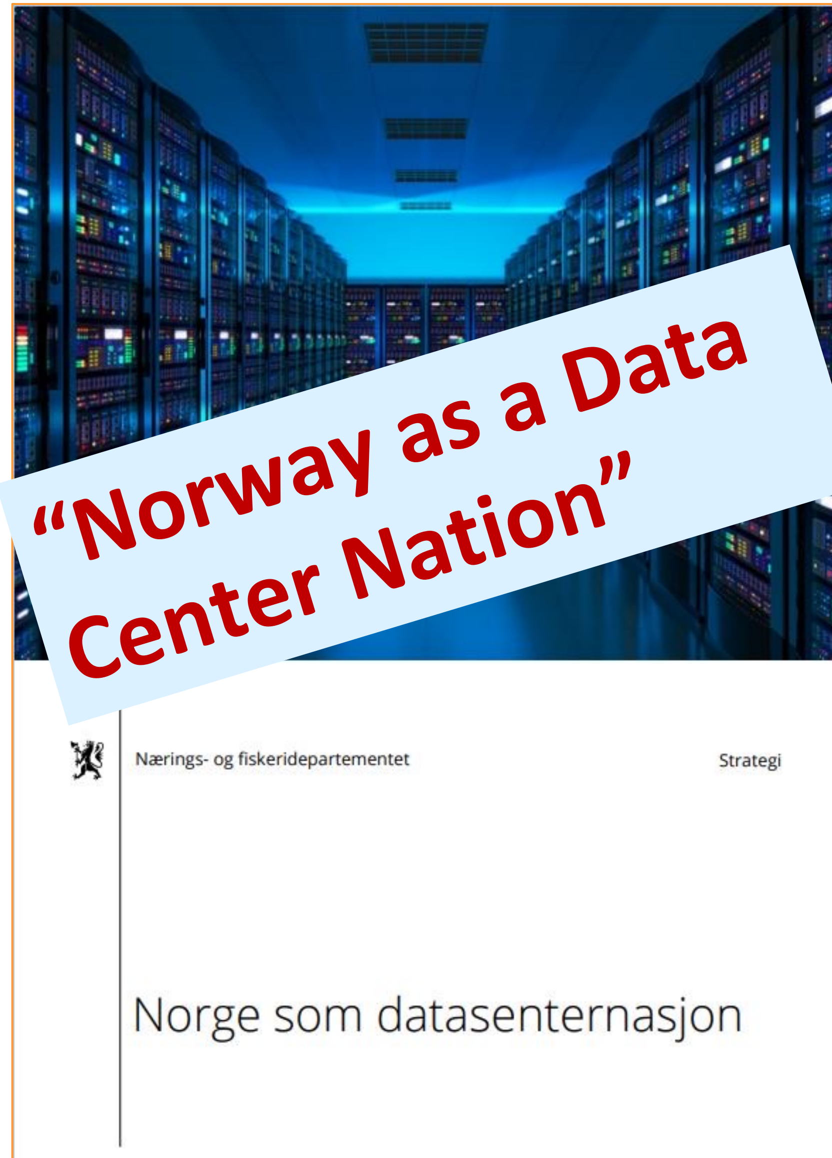

Data Center Definition

A data center is a building/container used to house computing equipment such as servers, along with associated components such as networking devices, storage systems, power distribution units, and cooling systems.

Common features

- **Power: equipped with a guaranteed power supply**
- **Communications: equipped with high bandwidth connectivity**
- **Reliability: redundant networks, power and IT infrastructure**
- **Environment: maintain a specified temperature and humidity**
- **Security: ensure that the facility and its data remain secure**

Data Center Physical Layout

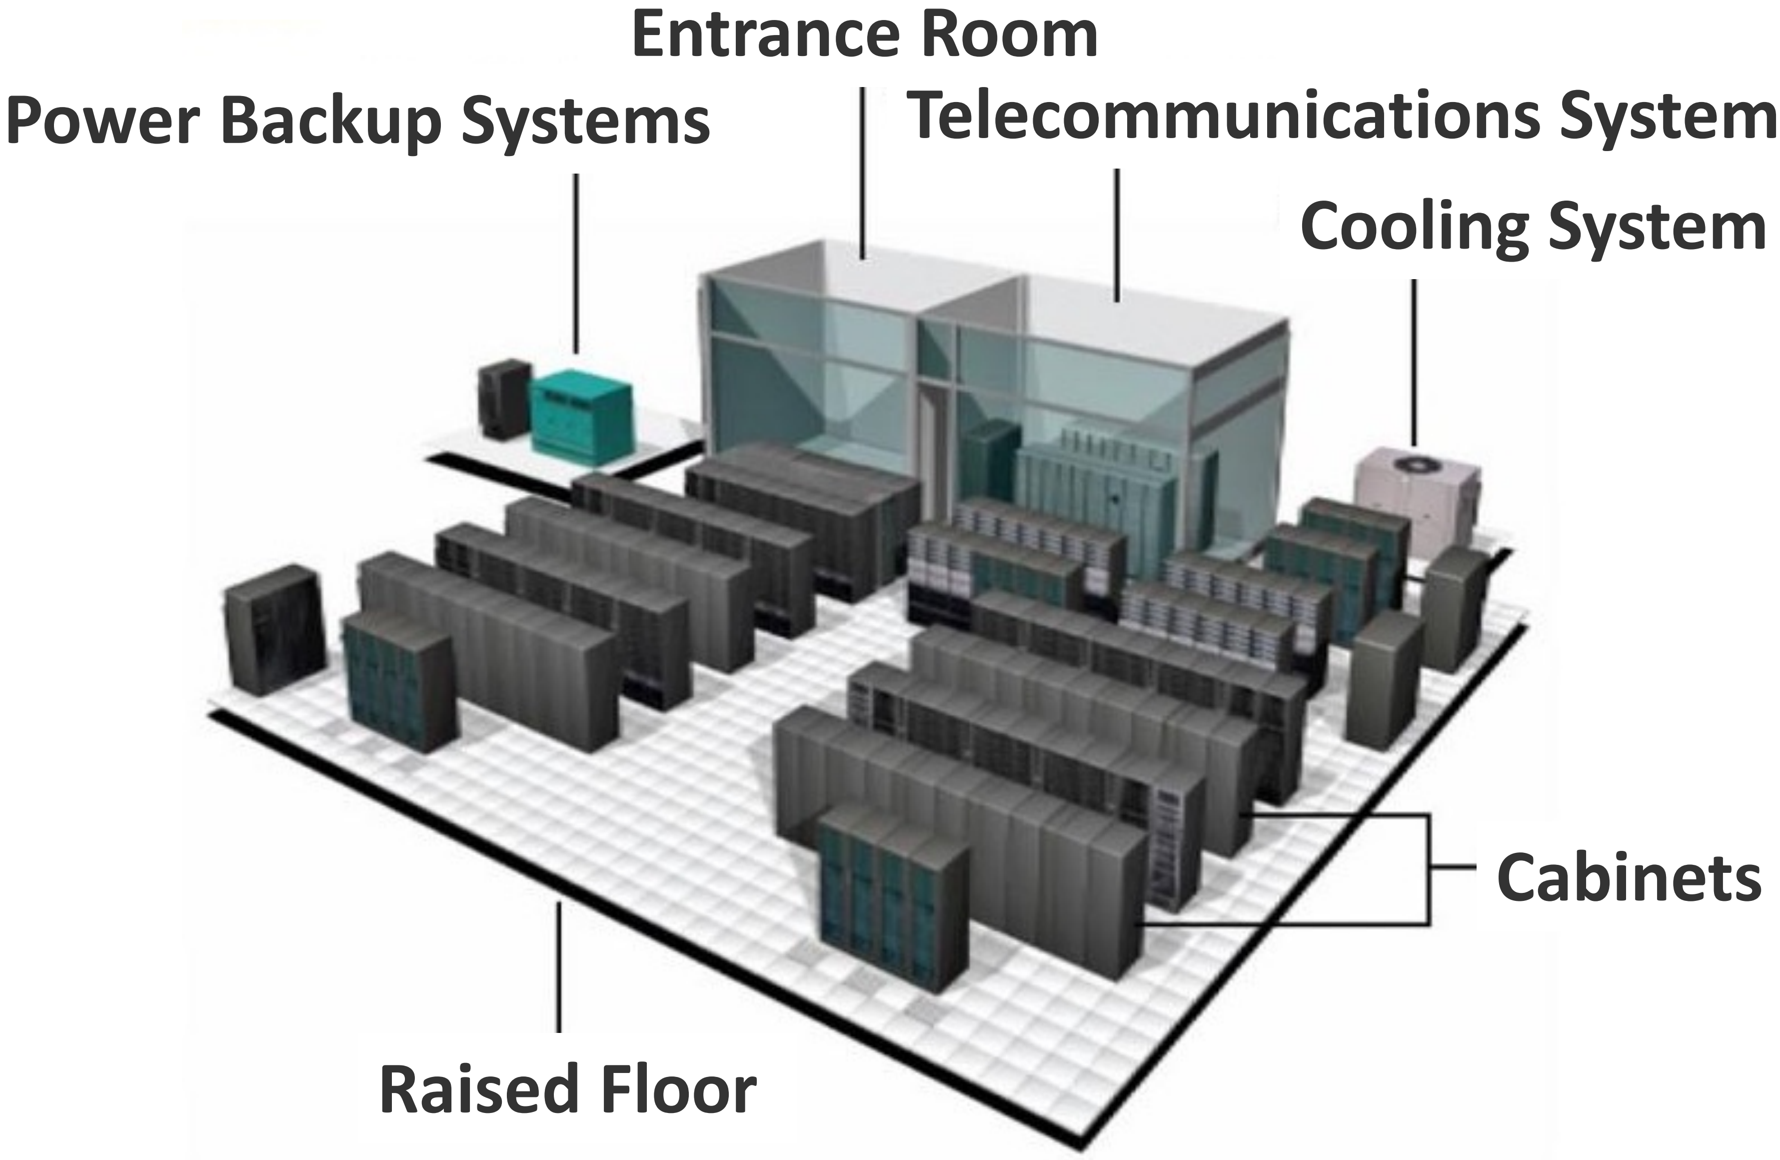

Data Centers – Hardware/Software Components

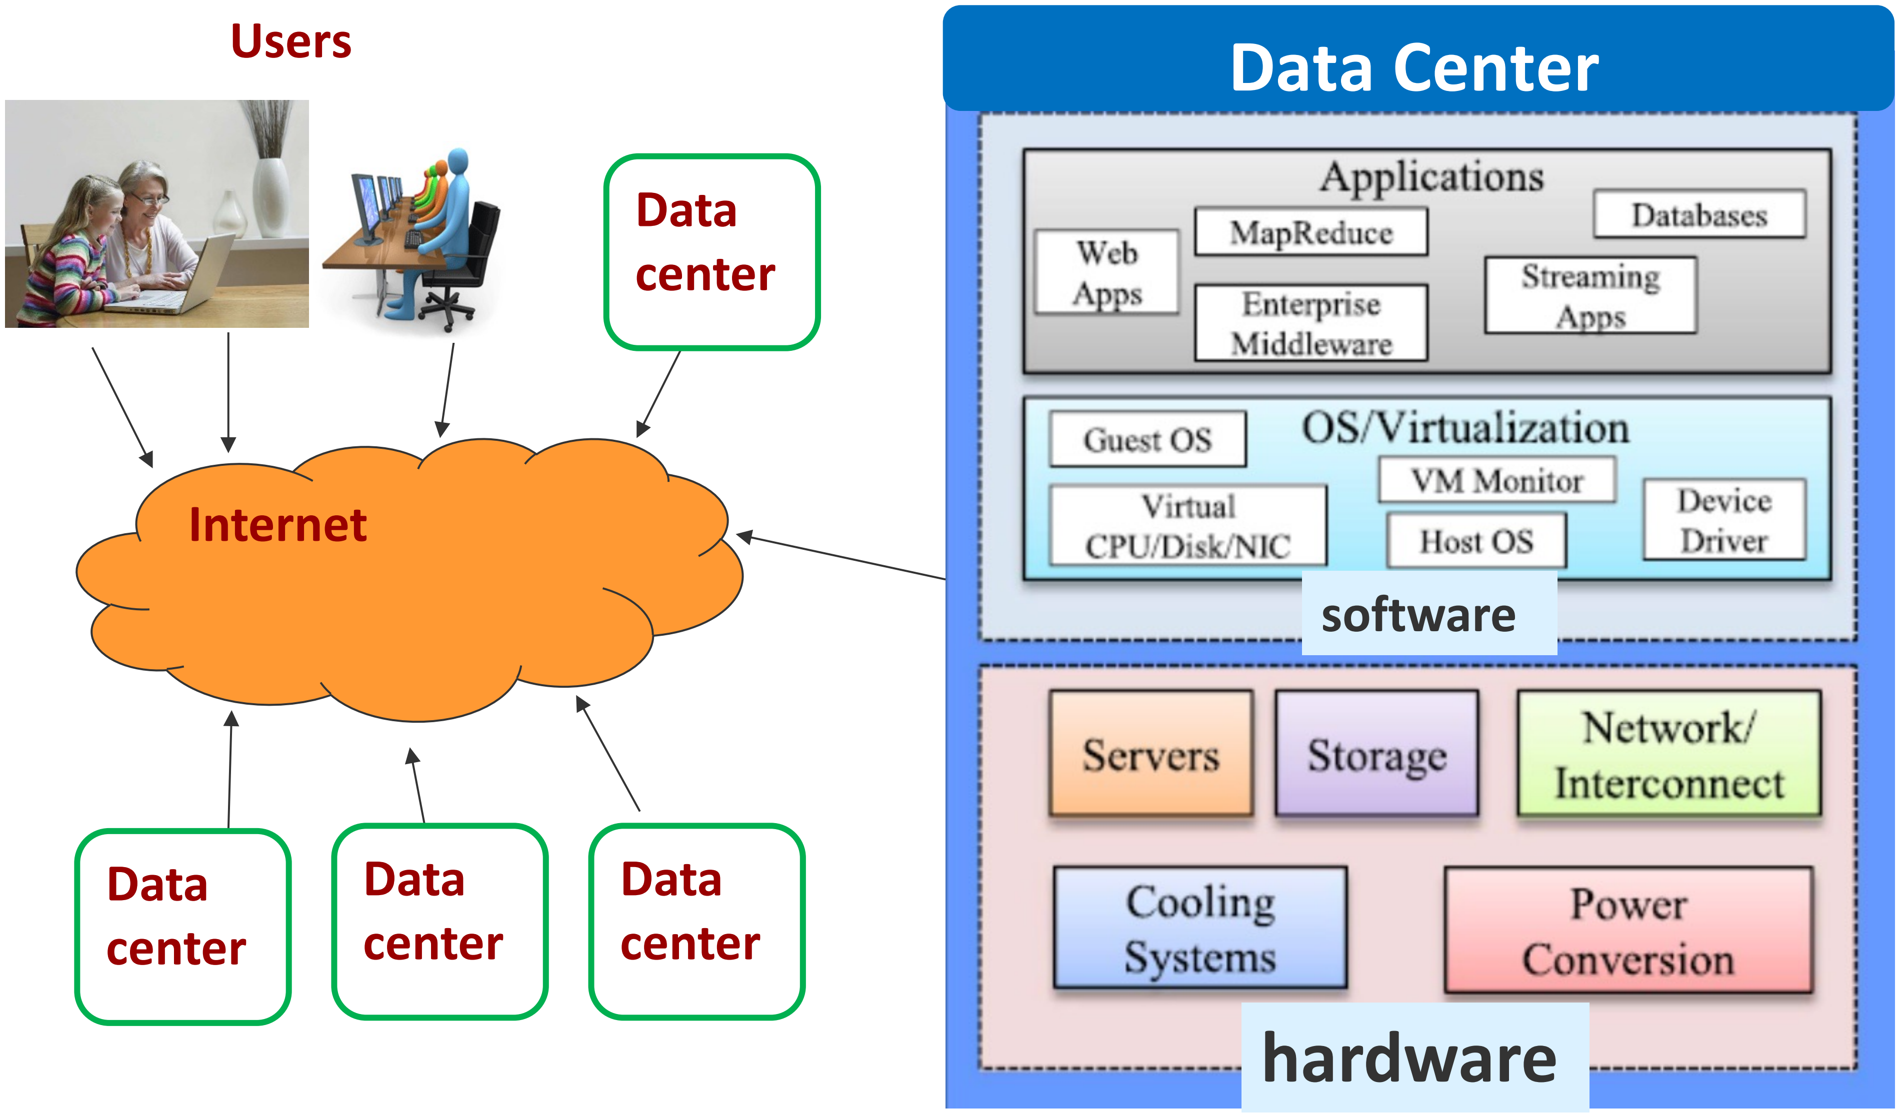

Servers organized as racks

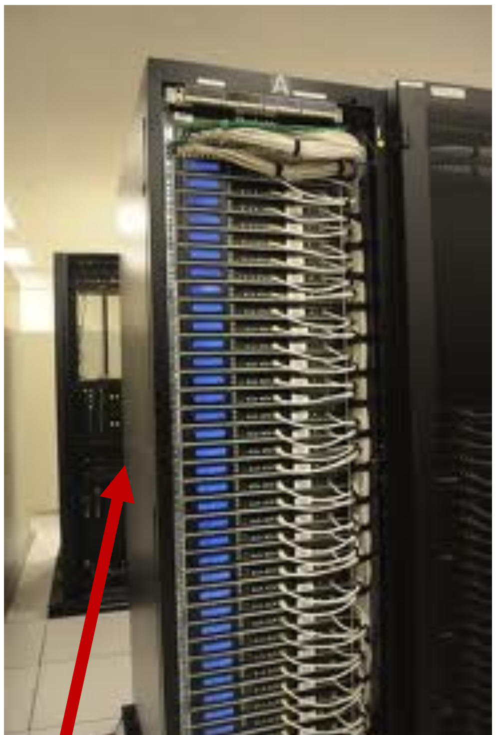

A rack

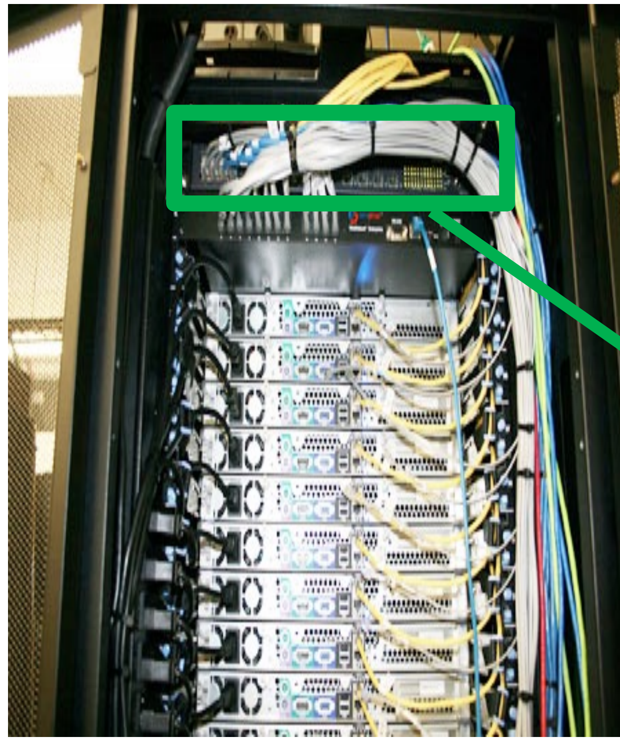

each rack has a "Top of Rack" (ToR) switch

ToR switch

Q: In practice a data center usually has multiple racks. How should those racks be connected?

“Top of Rack” (ToR) architecture connects many racks

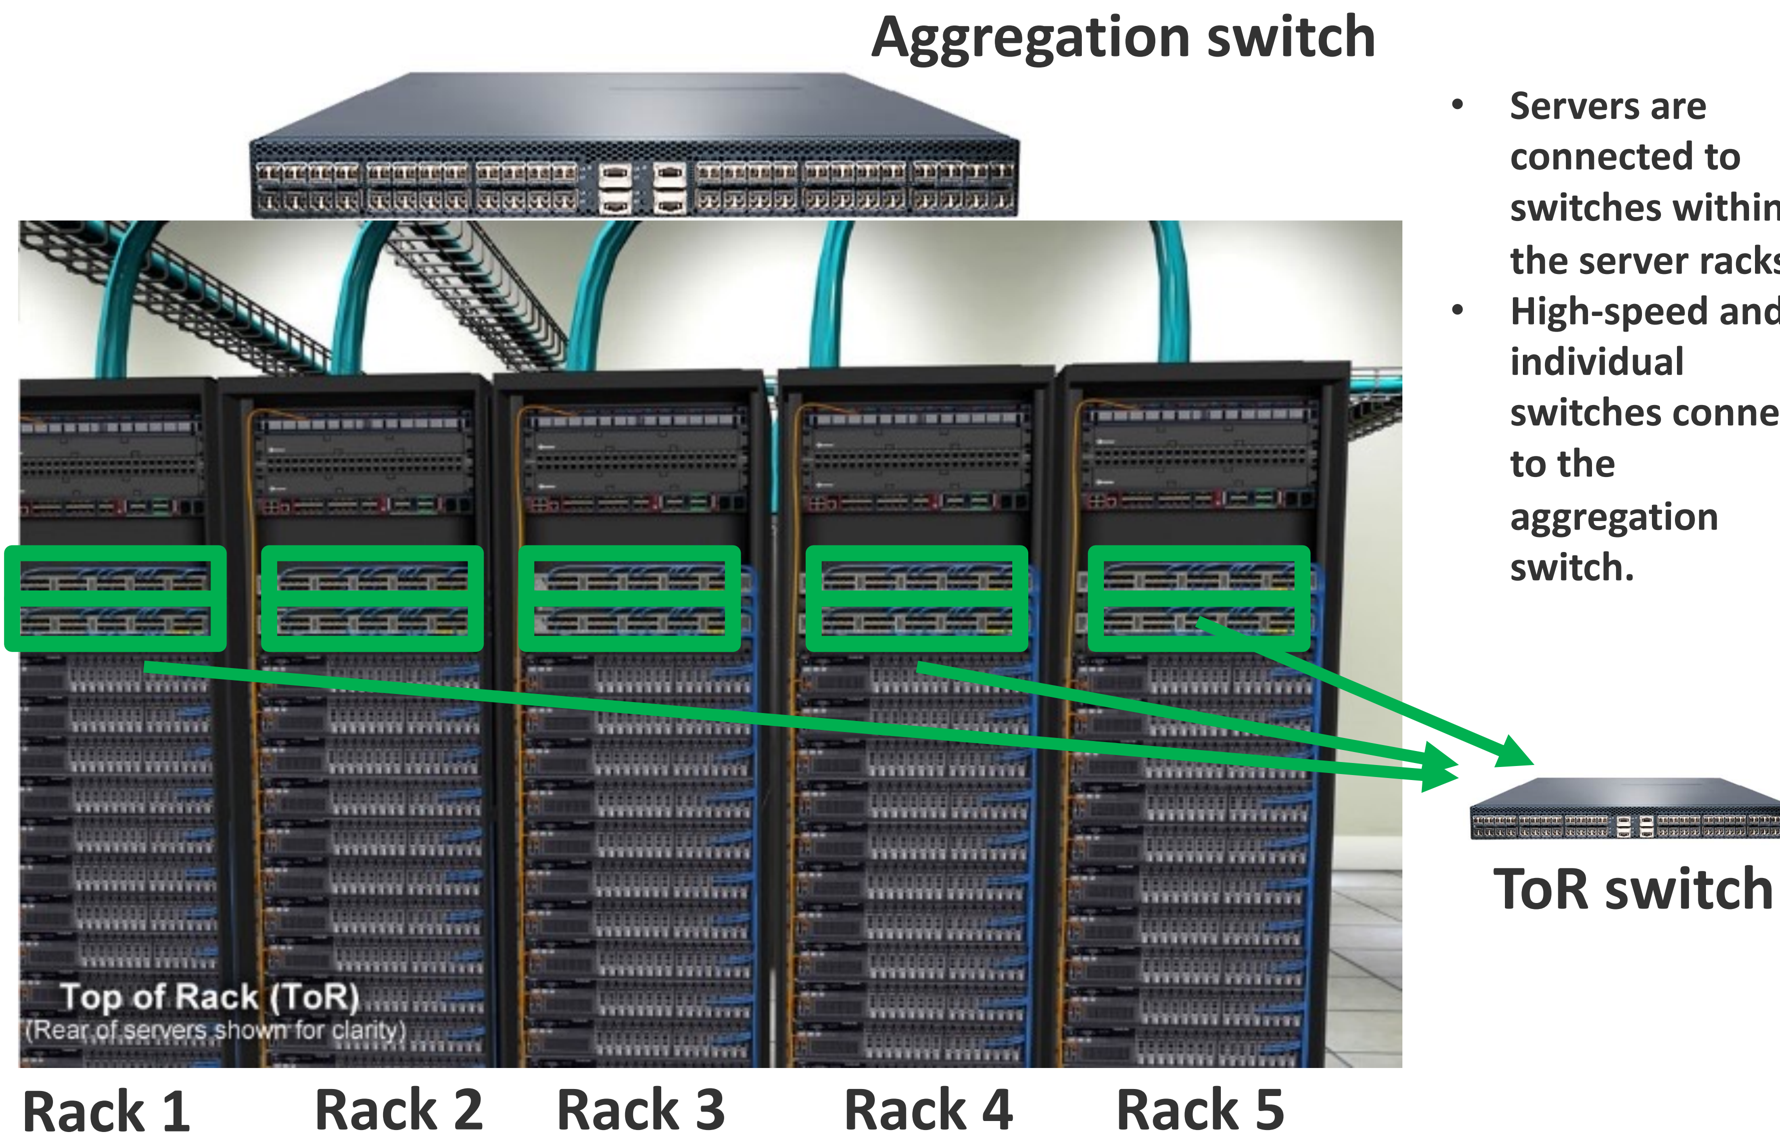

- **Servers are connected to switches within the server racks**
- **High-speed and individual switches connect to the aggregation switch.**

"Top of Rack" (ToR) architecture – abstract model

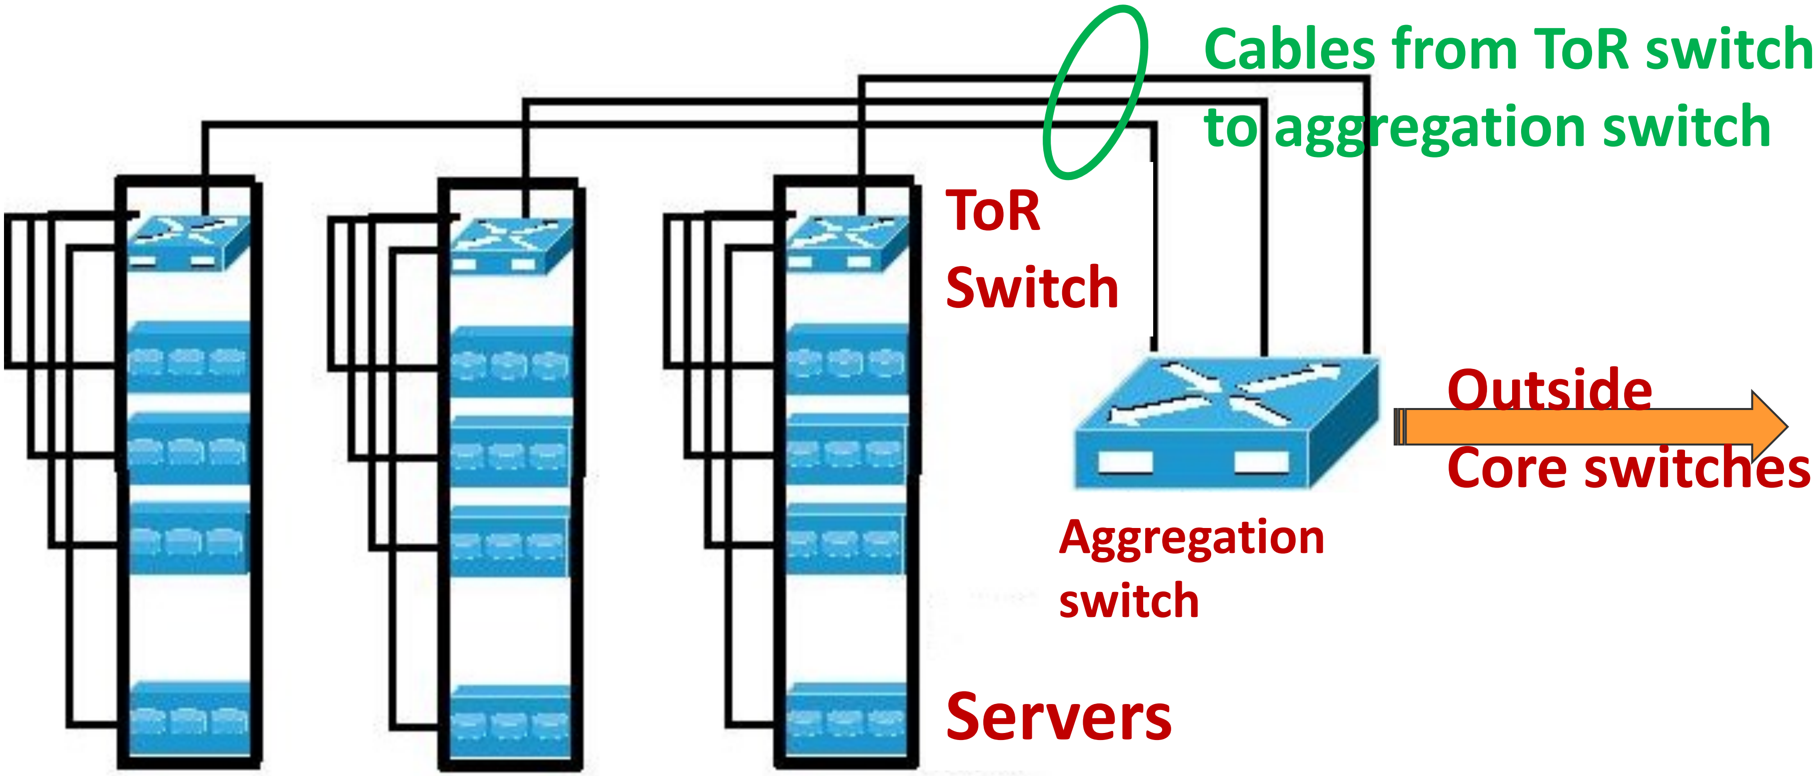

Source: http://www.excitingip.com/2802/data-center-network-top-of-rack-tor-vs-end-of-row-eor-design/

ToR architecture refers to the physical placement of the switch at the top of a server rack. Servers are directly linked to the switch. Each rack usually has two switches. (Q: why two switches?)

All the ToR switches are connected to the aggregation switch. Only a small amount of cables are needed to run from server rack to aggregation switch. Aggregation switch will be further connected to the outside core switches.

Major Components in a Data Center

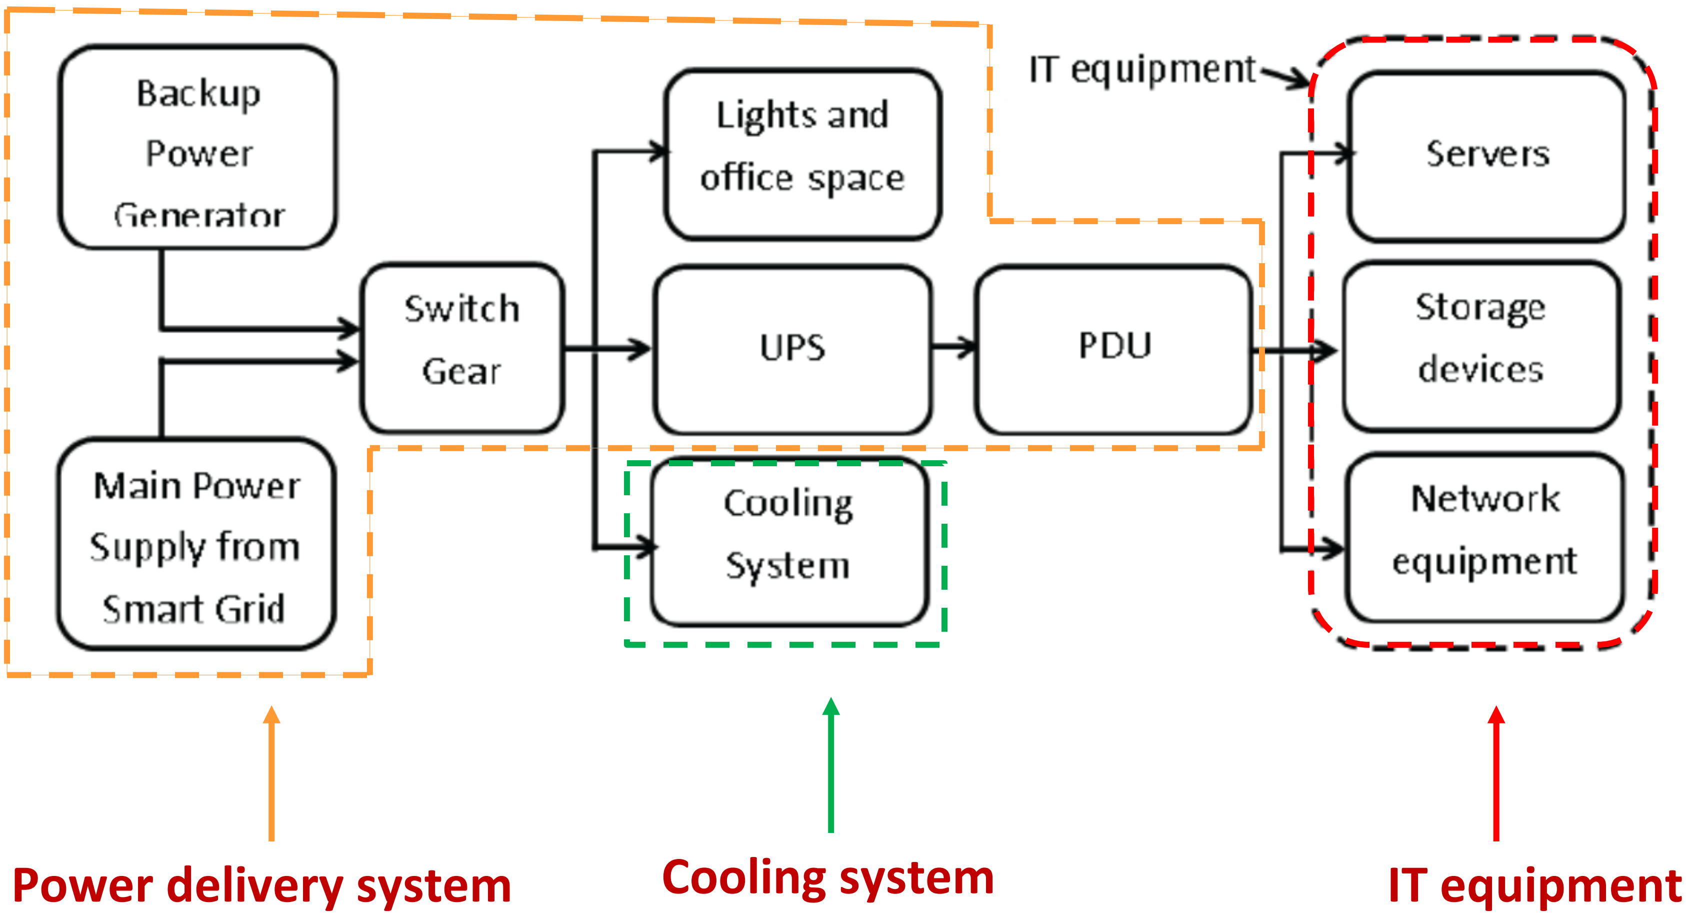

PDU: power distribution unit
UPS: Uninterruptible power supply

Data Center Major Components– IT equipments

Three kinds of IT equipments hosted in a typical data center

- **servers** for data processing
- **storage** equipment for data storage
- **network** equipment for data communications

Data Center Major Components - power delivery system

The power delivery system typically contains power conversion units, voltage regulators and backup equipment.

Power backup is often provided by Uninterruptible Power Supply (UPS) unit which prevents the IT equipment from experiencing power disruptions and possible serious business disruption or data loss.

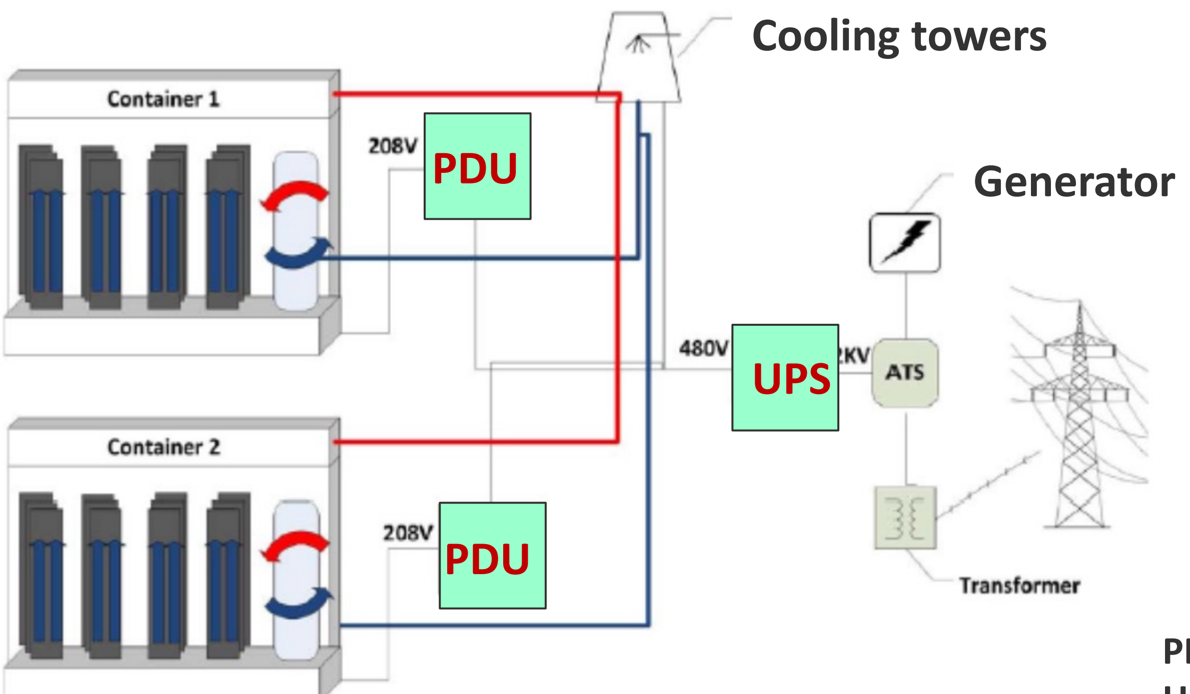

PDU: power distribution unit
UPS: Uninterruptible power supply

Data Center Major Components – cooling systems

Cooling in data centers is often provided by *Computer Room Air Conditioning (CRAC) units*.

Perforated floor tiles

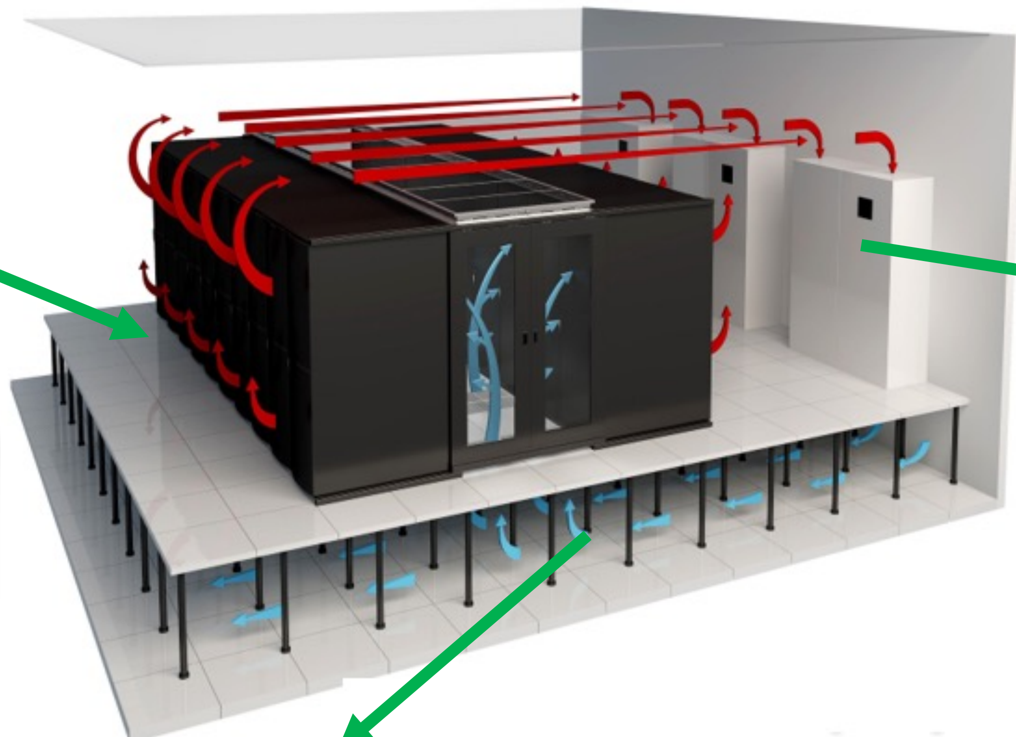

CRAC: Containing chilled water

Raised floor

Air Flow in Data Center

1. **Outside cool air enters the top of CRAC units**
2. **Air is conditioned by passing chilled water pumped from a chiller located outside the building.**
3. **The chilled air goes through a raised floor**
4. **The chilled air is then supplied to the servers**
5. **The cold air, while passing through perforated floor tiles, is pulled by the fans located inside the servers.**

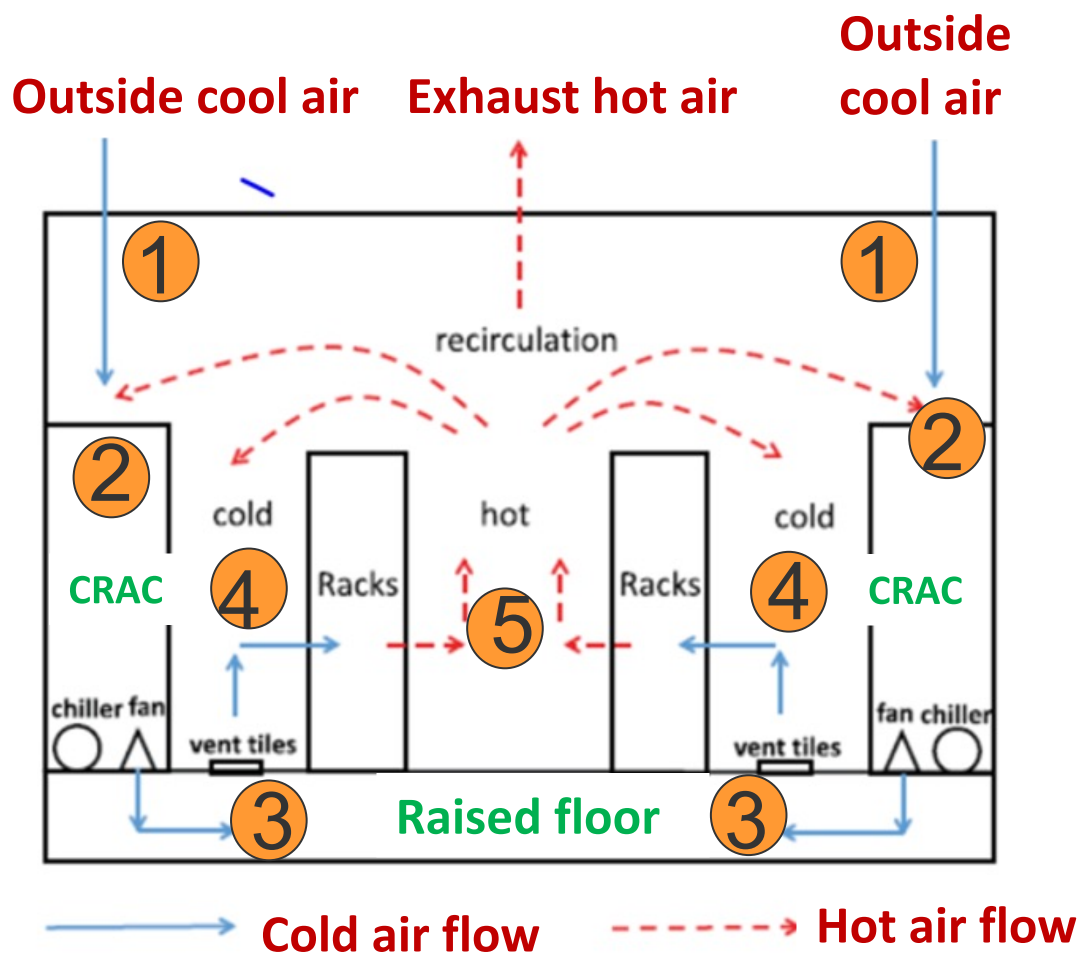

Still, data centers use much energy

Q: Why place data centers in Finland or undersea?

Data centers are huge energy consumers, in particular cooling systems, and

- **pay a lot for electricity bill**
- **make power grid instable during peak hours**

Cooling system uses sea water from the Bay of Finland and reduces energy use

Data Center Energy Consumption

INDEPENDENT News Voices Culture Lifestyle Tech Sport Daily Edition Charity Appeal

Environment

Global warming: Data centres to consume three times as much energy in next decade, experts warn

416.2 terawatt hours of electricity world's data centres us… consumption

Tom Bawden Environment Editor | @BawdenT…

"US data centers consumed about 70 billion kWh electricity in 2014"

- US Department of Energy

…uch Energy All US …enters Consume

BY YEVGENIY SVERDLIK ON JUNE 27, 2016 1 COMMENT

162 Like Tweet 597 Share 45 G+1

It's no secret that data centers, the massive but bland, unremarkable-looking buildings housing the powerful engines that pump blood through the arteries of global economy, consume a huge amount of energy. But while our reliance on this infrastructure and its ability to scale capacity grows at

Growth rate of total US data center energy use from 2000 until 2020.

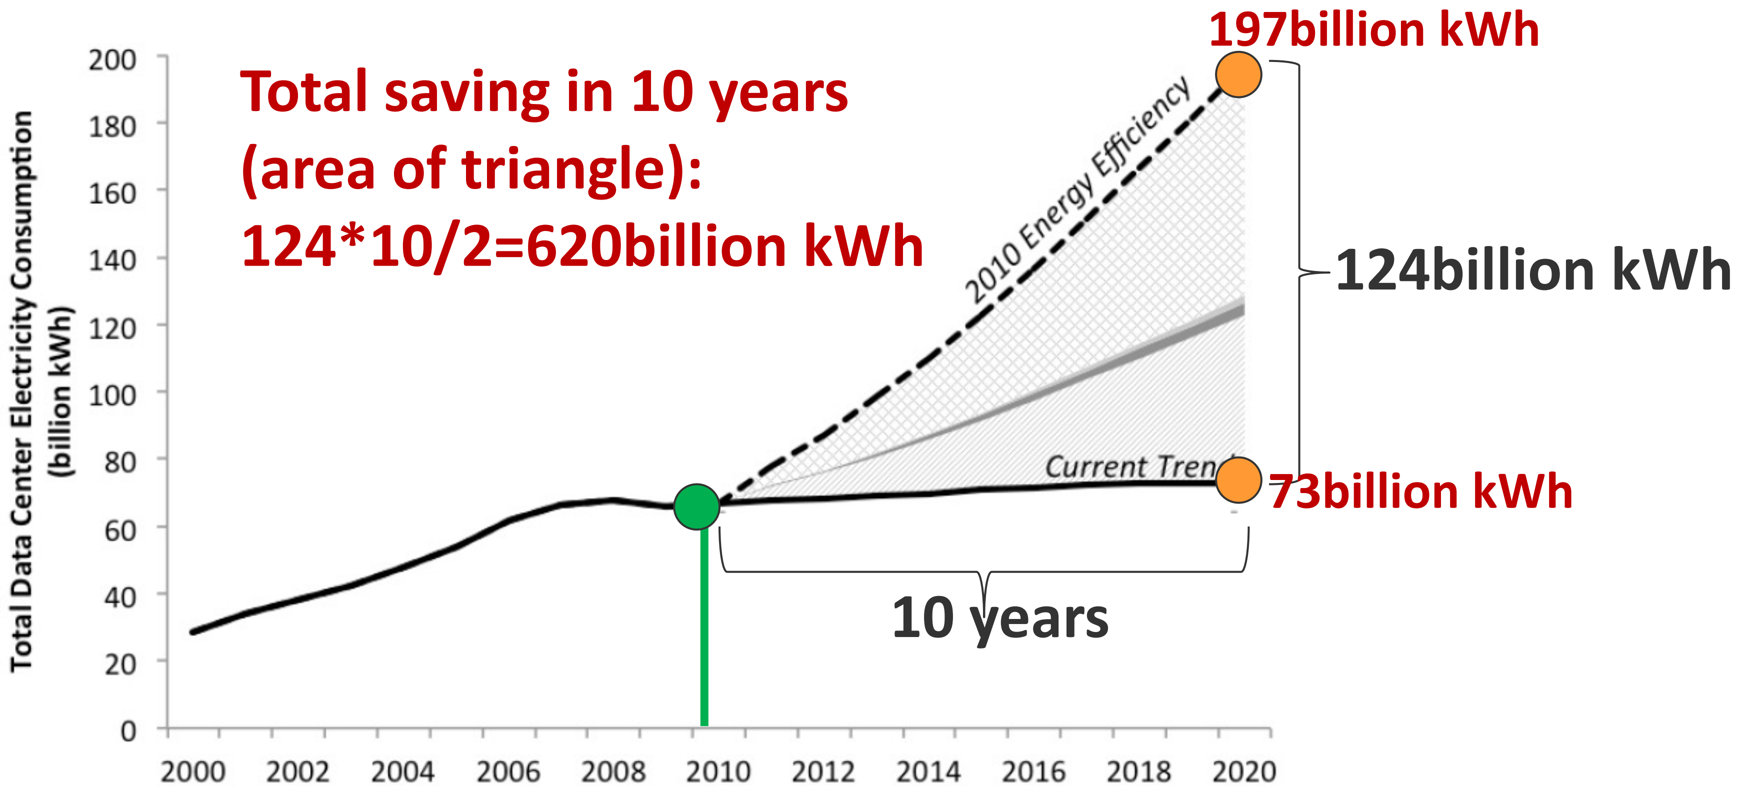

Industry has considered efficiency since 2010. The illustration shows how much faster data center energy use would grow if the industry, hypothetically, did not make any further efficiency improvements after 2010. (Source: US Department of Energy, Lawrence Berkeley National Laboratory**)**

Energy consumption in 2020 was estimated to reach about 73billion kWh

ENERGY CONSUMPTION IN DATA CENTER

Global electricity demand of data centers

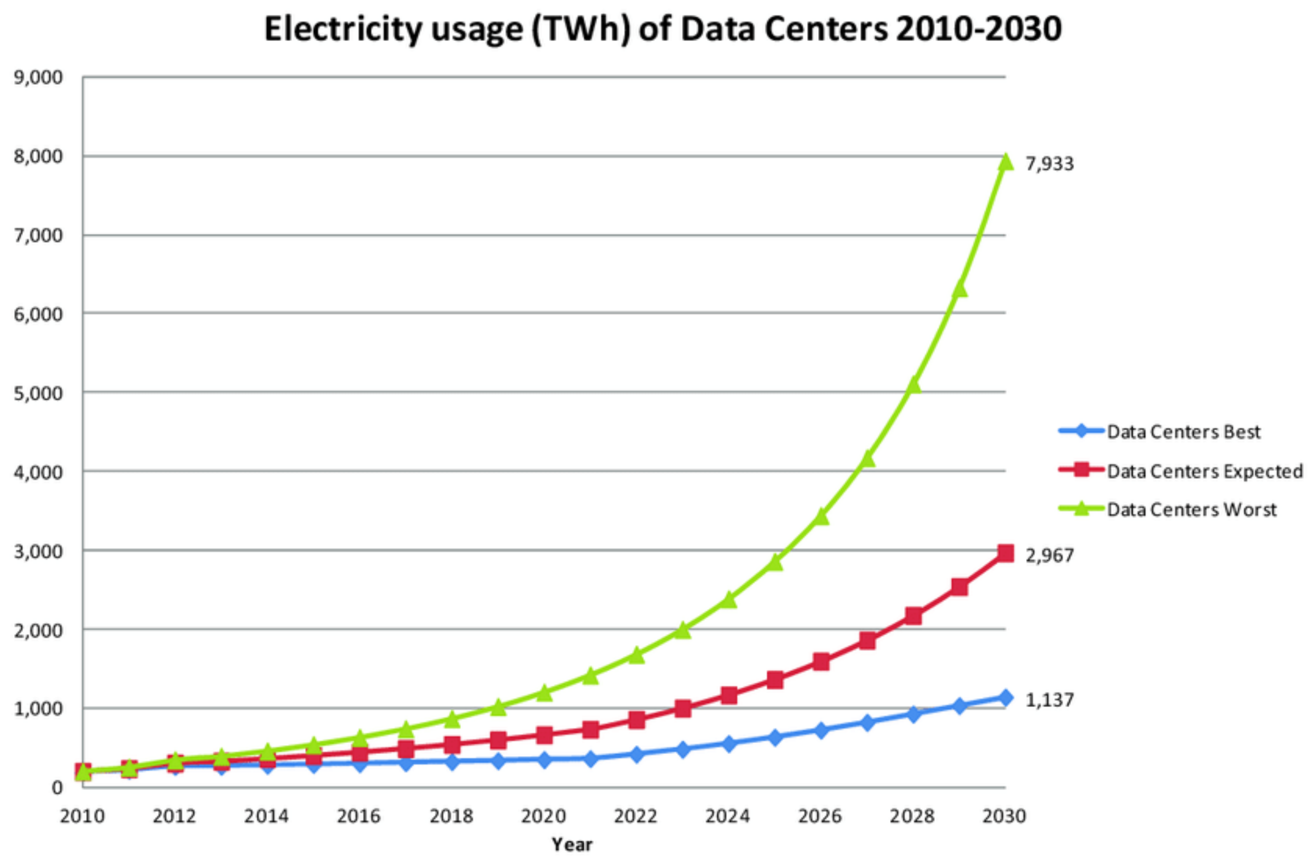

Fig. source
https://www.researchgate.net/figure/Global-electricity-demand-of-data-centers-2010-2030_fig2_275653947

Energy Consumption in Data Center

The **largest energy consumer** in a typical data center is the **cooling infrastructure (50%),** while servers and storage devices (26%) rank second in terms of energy consumption.

[according to the statistics published by the Infotech group] (**Note:** that these values might differ from data center to data center)

A breakdown of energy consumption by different components shows: the cooling infrastructure consumes a major portion of the data center energy followed by servers and storage, and other infrastructure elements.

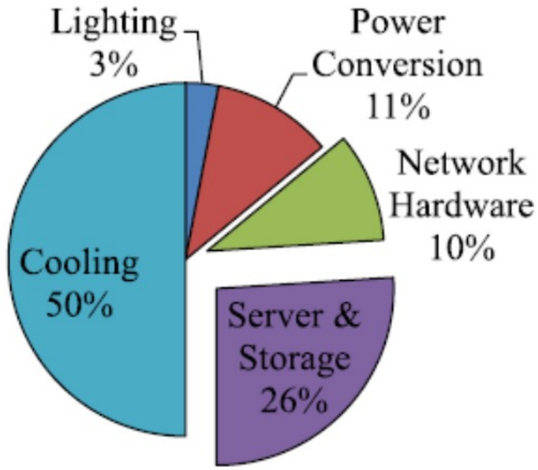

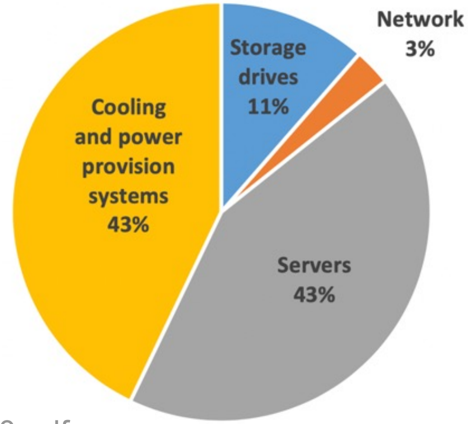

Figure 1. Fraction of U.S. data center electricity use in 2014, by end use. Source: Shehabi 2016.

Fig. src: http://static.infotech.com/downloads/samples/070411_premium_oo_greendc_top_10.pdf
https://energyinnovation.org/2020/03/17/how-much-energy-do-data-centers-really-use/

Energy Efficiency Metrics for Data Centers

In order to quantify the energy efficiency of data centers, several energy efficiency metrics have been proposed to help data center operators to improve the energy efficiency and reduce operation costs of data centers.

Two important energy efficiency metrics are

- **Power Usage Effectiveness (PUE)**
- **Data Center energy Productivity (DCeP)**

Power Usage Effectiveness (PUE)

PUE: the most commonly used metric to indicate the energy efficiency of a data center.

$$PUE = \frac{Total\ Power\ Consumption\ of\ a\ Data\ Center}{Total\ Power\ Consumption\ of\ IT\ Equipment}$$

$$= \frac{Total\ Power\ Consumption\ of\ IT\ Equipment + Total\ Power\ Consumption\ of\ nonIT\ Equipment}{Total\ Power\ Consumption\ of\ IT\ Equipment}$$

$$= 1 + \frac{Total\ Power\ Consumption\ of\ nonIT\ Equipment}{Total\ Power\ Consumption\ of\ IT\ Equipment} > 1$$

Power Usage Effectiveness (PUE): an example

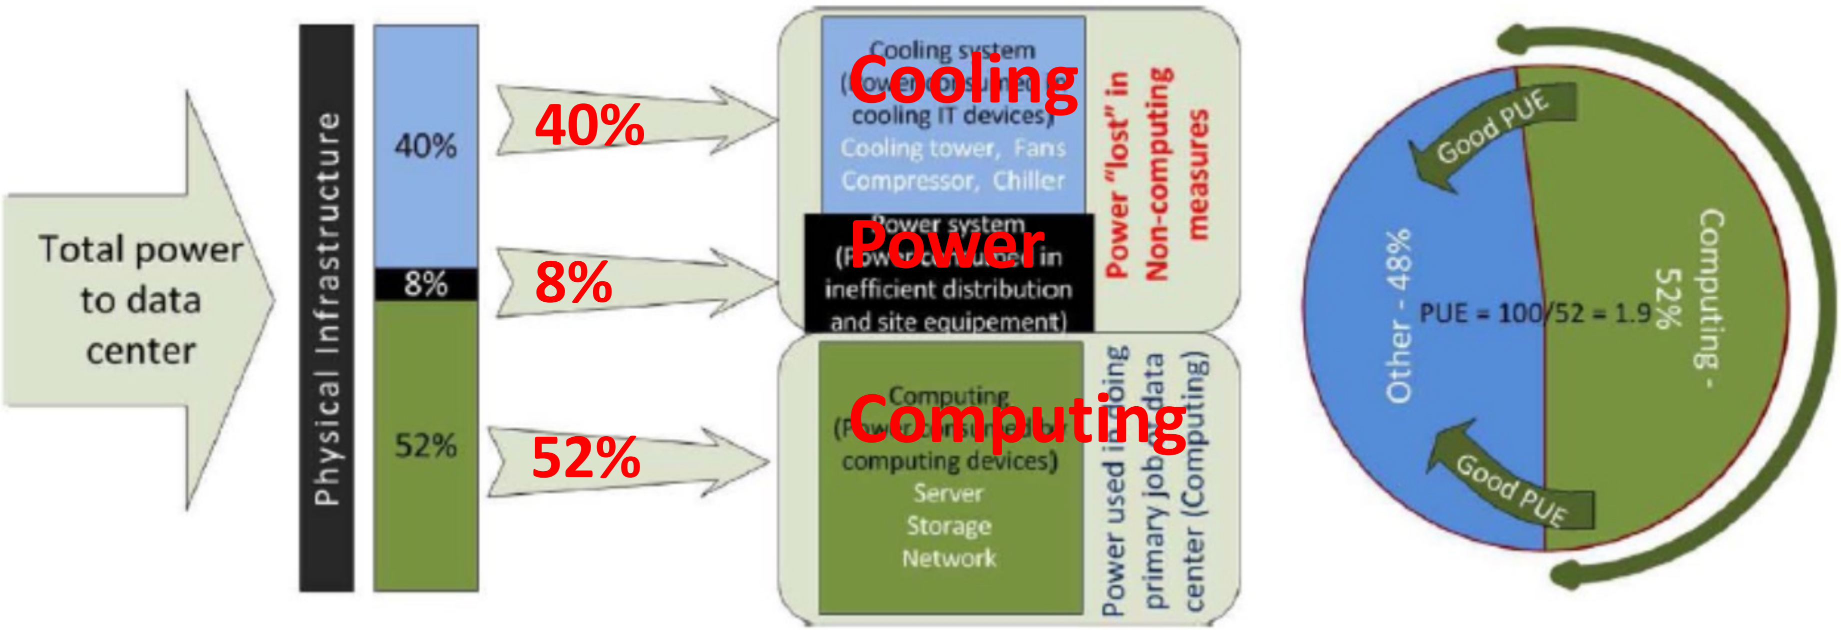

Source: J. Shuja, et al.,"Survey of Techniques and Architectures for Designing Energy-Efficient Data Centers," IEEE Systems Journal, vol. 10, no. 2, pp. 507-519, 2016

Distribution of data center power among cooling, power distribution, and computing units

In this example, 52% power goes to computing while 48% power goes to non-computing. Then, PUE = 100/52=1.9

Google's PUE

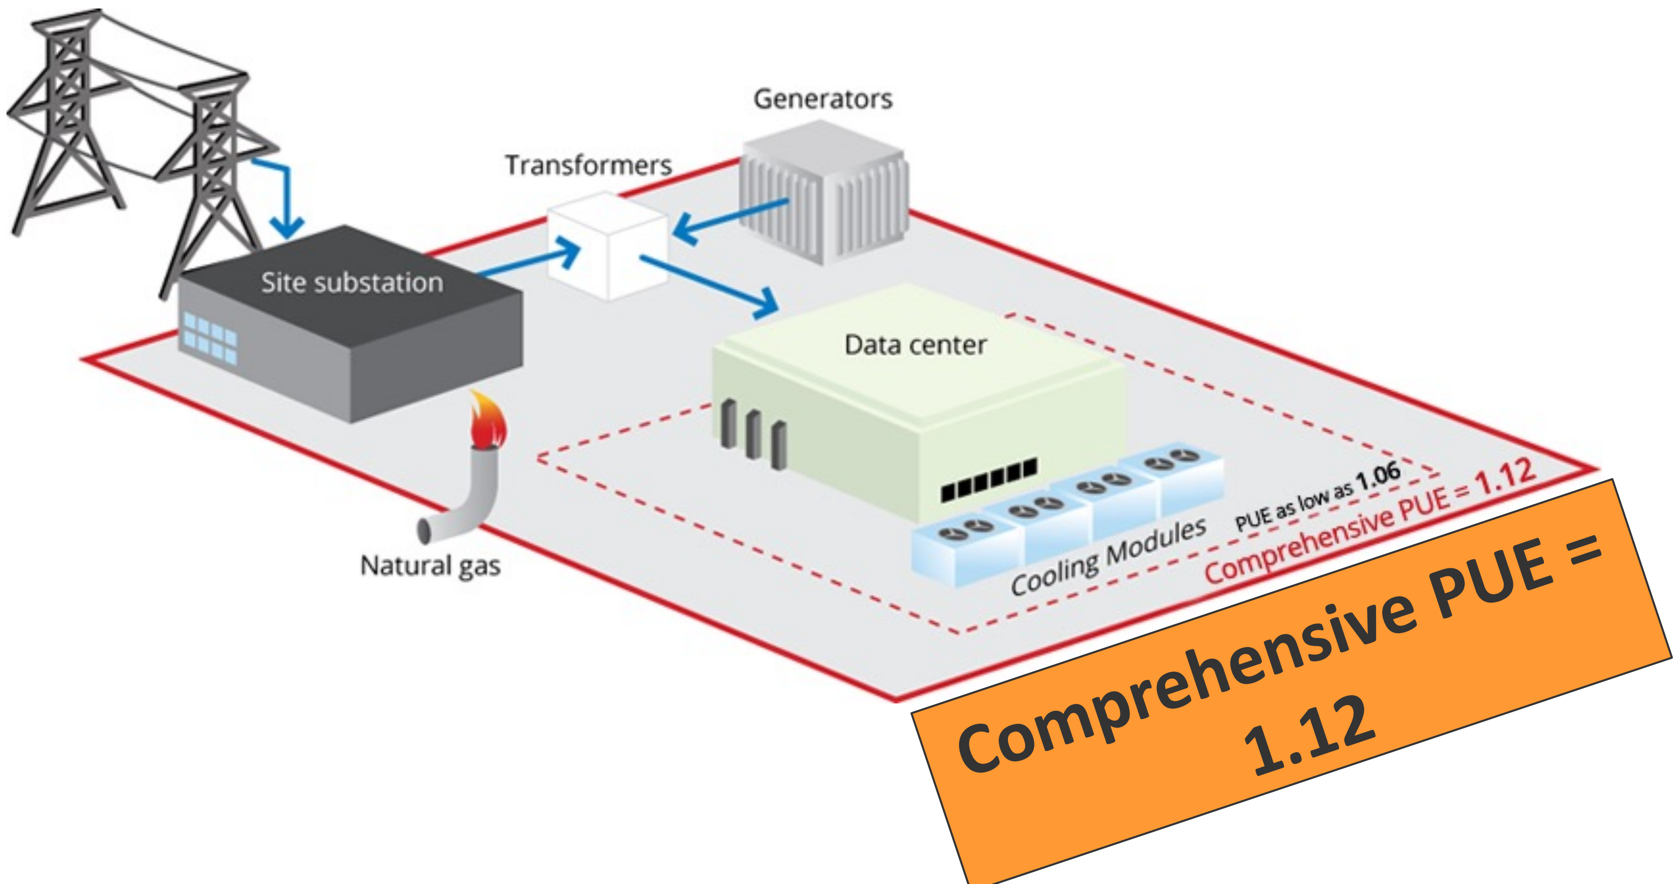

A PUE closer to 1.0 means nearly all of the energy is used for computing.

Google reports: the average PUE for all Google data centers is 1.12. Best sites can be less than 1.06.

More information: https://www.google.com/about/datacenters/efficiency/internal/

PUE Features

PUE measures the total power consumption overhead caused by the data center support equipment, including the cooling systems, power delivery, and other facility infrastructure like lighting.

PUE=1.0 implies: there is no power overhead and all power consumption of the data center goes to the IT equipment.

A higher PUE: a greater portion of the electricity coming to the data center spent on cooling and the rest of the infrastructure

Q: Can a good PUE value guarantee the global efficiency of a data center? Menti

A good PUE value may not be enough to guarantee the global efficiency of the data center. PUE metric does not consider the actual utilization (applications and workloads) of computational resources, some computation may not be necessary.

Data Center energy Productivity (DCeP)

Energy efficiency and energy productivity are closely related to each other.

- **Energy efficiency focuses on reducing unnecessary power consumption to produce a work output.**
- **Energy productivity measures the quantity of useful work done relative to the amount of power consumption of a data center in producing this work.**

DCeP measures the useful work performed by a data center relative to the energy consumed by the data center in performing the work

$$DCeP = \frac{Useful\ Work\ Produced}{Total\ Data\ Center\ Power\ Consumed\ for\ Producing\ this\ Work}$$

DCeP

- **DCeP metric allows the user to define the computational tasks, transactions, or jobs that are of interest, and then assign a measure of importance of economic value to each specific unit of work completed**
- **DCeP allows the continuous monitoring of the productivity of a data center as a function of power consumed by a data center.**
- **DCeP metric tracks the overall work product of a data center per unit of power consumption expended to produce this work.**

Q: Easy to use this metric?

– subjective and may not be easy to define *"Useful Work"*

GREEN DATA CENTER

Green Data Center

Green Data Centers
energy-aware, energy-efficient with minimum CO_2 emission designs, protocols, devices, infrastructures and algorithms for data centers

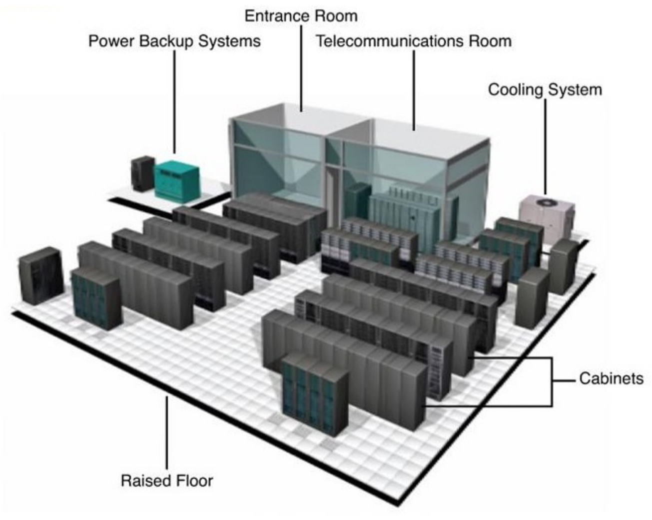

Reduce energy consumption at an individual data center

Reduce energy consumption across globally located data centers

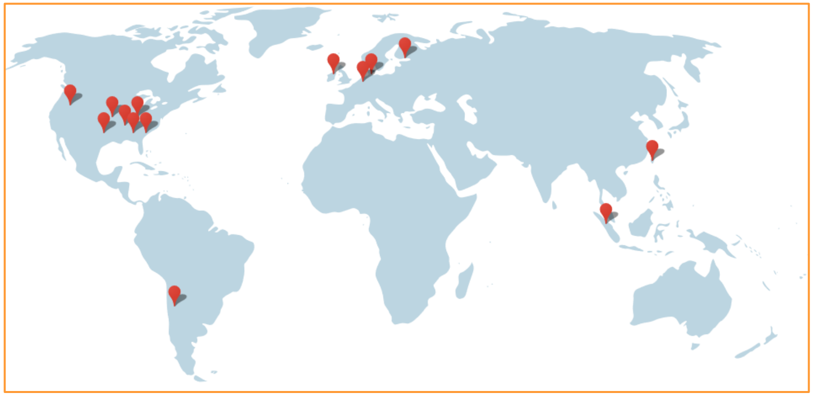

Techniques to Improve Energy Efficiency of Data Centers

Smart cooling and thermal management

Power management – Chip level

Power management – Server level

Power management – Data center level

Power management – Inter data center level

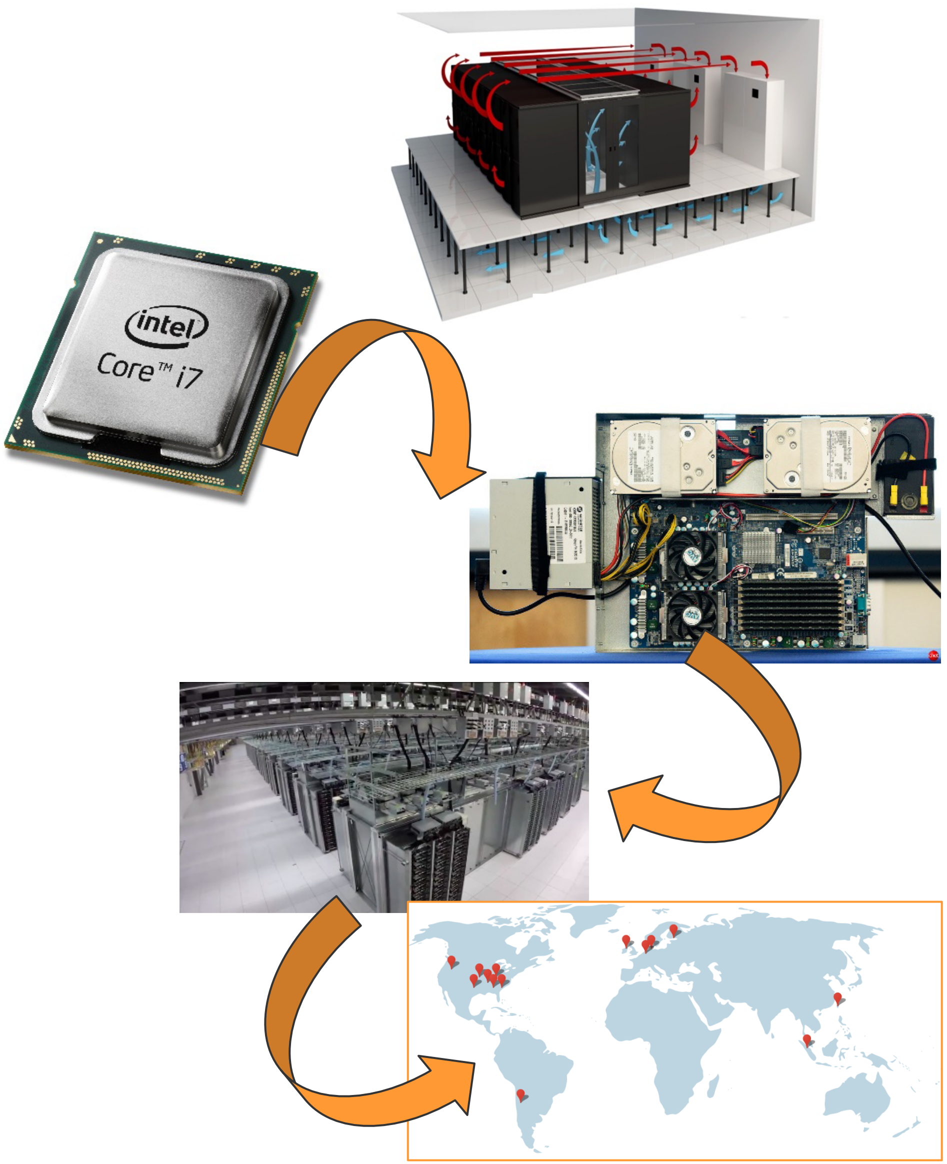

Chip Level Management

Dynamic voltage and frequency scaling (DVFS) is a commonly-used technique to save power at chip level power management

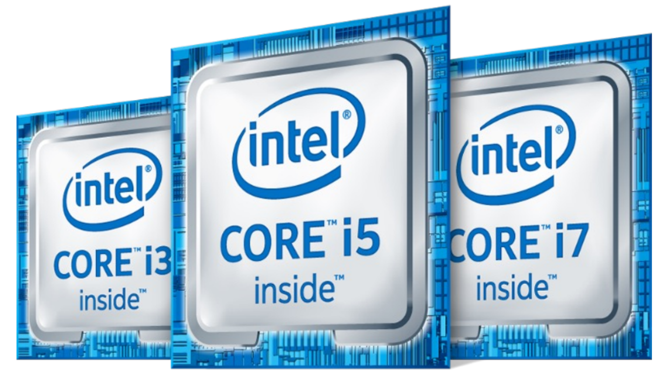

Main principle: DVFS reduces the power consumption of CPU by reducing the operating frequency or the supply voltage, as shown by

$$P = CfV^2 + P_{static}$$

physical capacitance → C

Operating frequency → f

Supply voltage → V

The voltage can be reduced as the frequency is reduced. This can yield a significant reduction in power consumption because of the V^2 relationship.

Q: what can be the consequence of reducing frequency? ⟷ **Delay!**

Server Level Management

A server has many components, including CPU, fans, disk, memory, etc.

DVFS can be highly effective in reducing CPU power. However, **CPUs account for a portion of total system power.**

Approach to manage power in the entire server, e.g., transition into a **low-power state or sleeping state** when there are no tasks, are necessary.

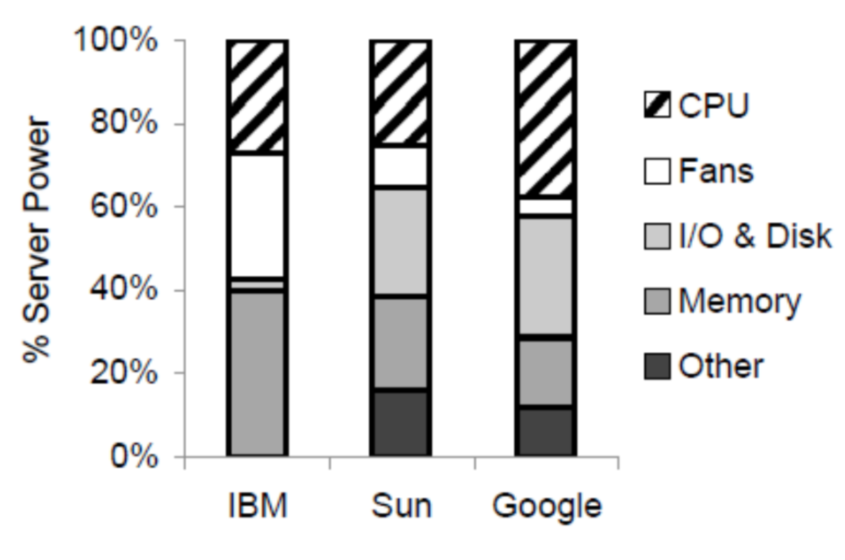

Typical server power breakdowns for the IBM p670, Sun UltraSparc T2000, and a generic server by Google.

PowerNap for a Server

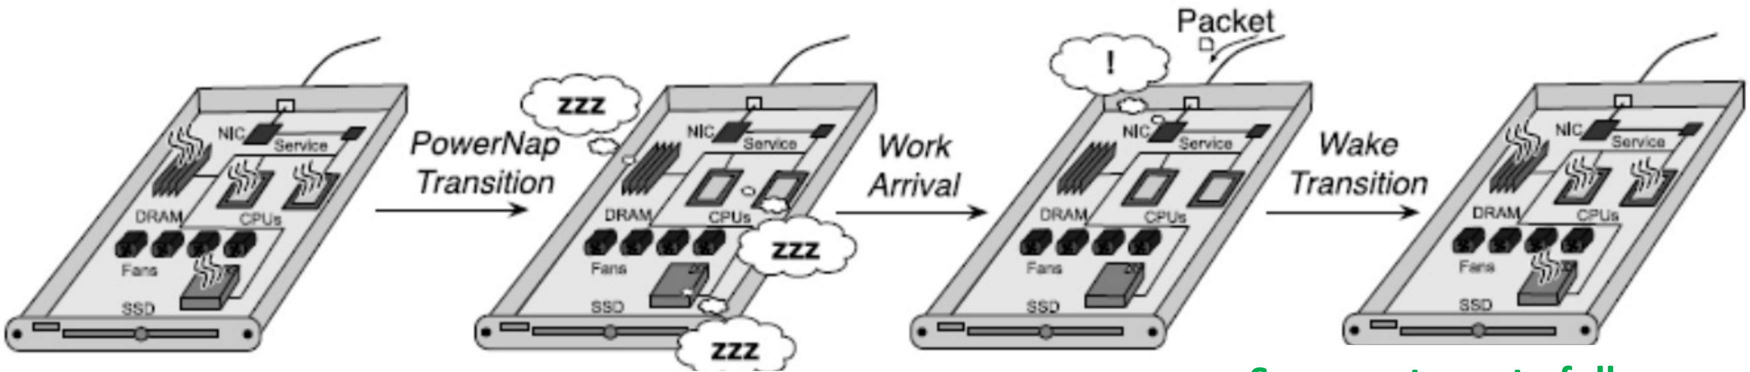

The server operates at full performance to complete existing tasks

System components nap while the server is idle

Task arrival is detected.

Server returns to full performance mode to execute the tasks

Each time the server completes current and all pending tasks, it transits to *the nap state*.

In the nap state: nearly all system components enter sleep mode. Power consumption is low, no processing can occur. **Components that signal the arrival of new tasks, or expiration of a software timer, remain partially powered**.

In the active state: When new tasks arrive, the system wakes and transitions back to the active state. When the work is complete, the system returns to the nap state again.

Data Center Level - Workload Migration

Observation: different workloads can result in different levels of power consumption.

Reason: too many computation tasks are running in Zone A while only few tasks in Zone B.

Q: how to solve this?

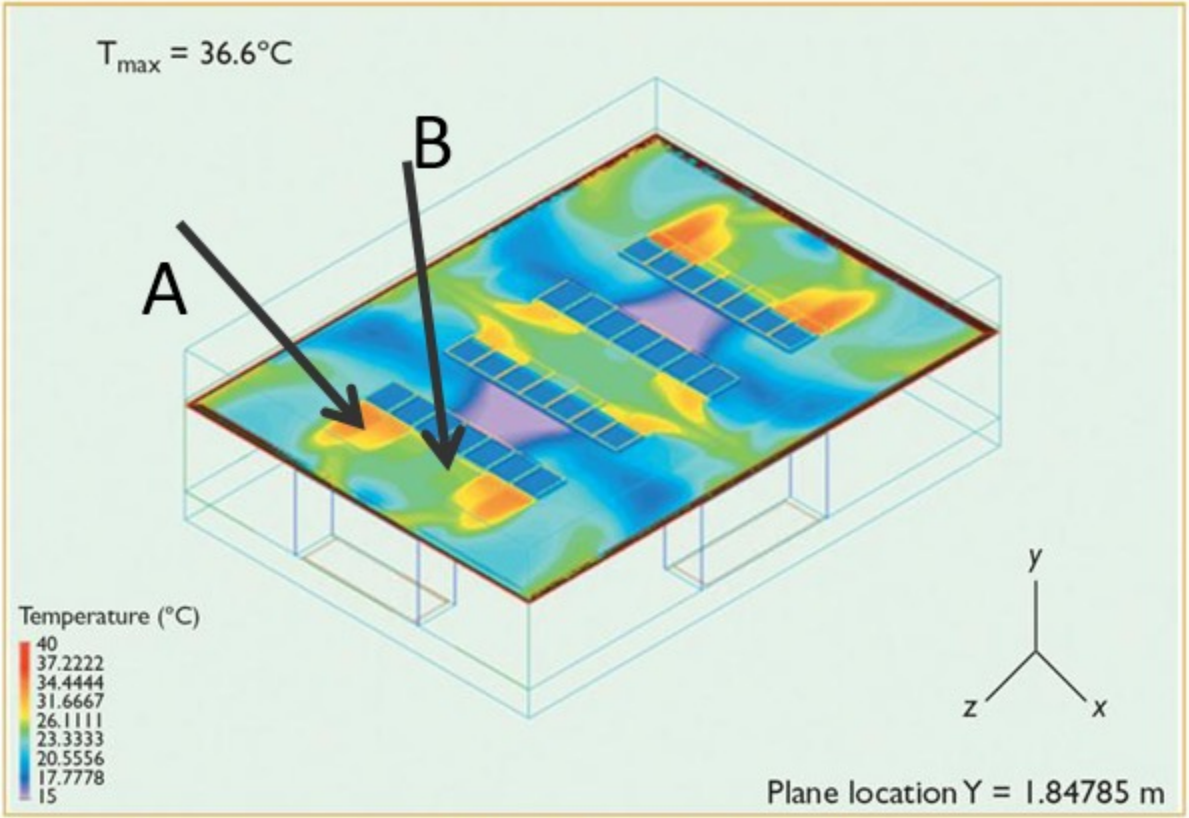

Temperature distribution in data center

Solution Approach: Move some computation tasks from Zone A to Zone B to reduce overall power consumption

Data Center Level – Dynamic Component Deactivation

World Cup

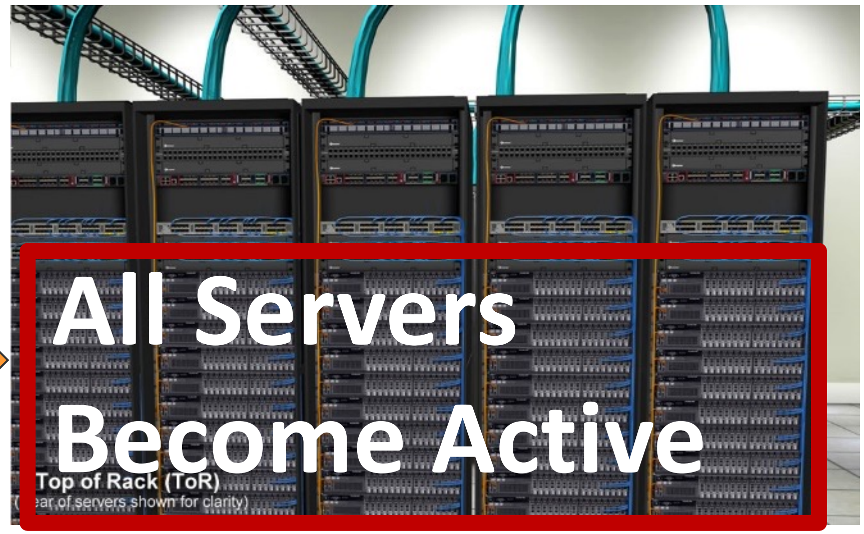

Main idea: Allow servers to be dynamically turned off to save energy and only turned on minimum servers to finish tasks.

Midnight

Inter Data Centers Level: geographical load balancing

Q: **what are the differences between these locations to consider when we need to handle the service requests at Oslo?**

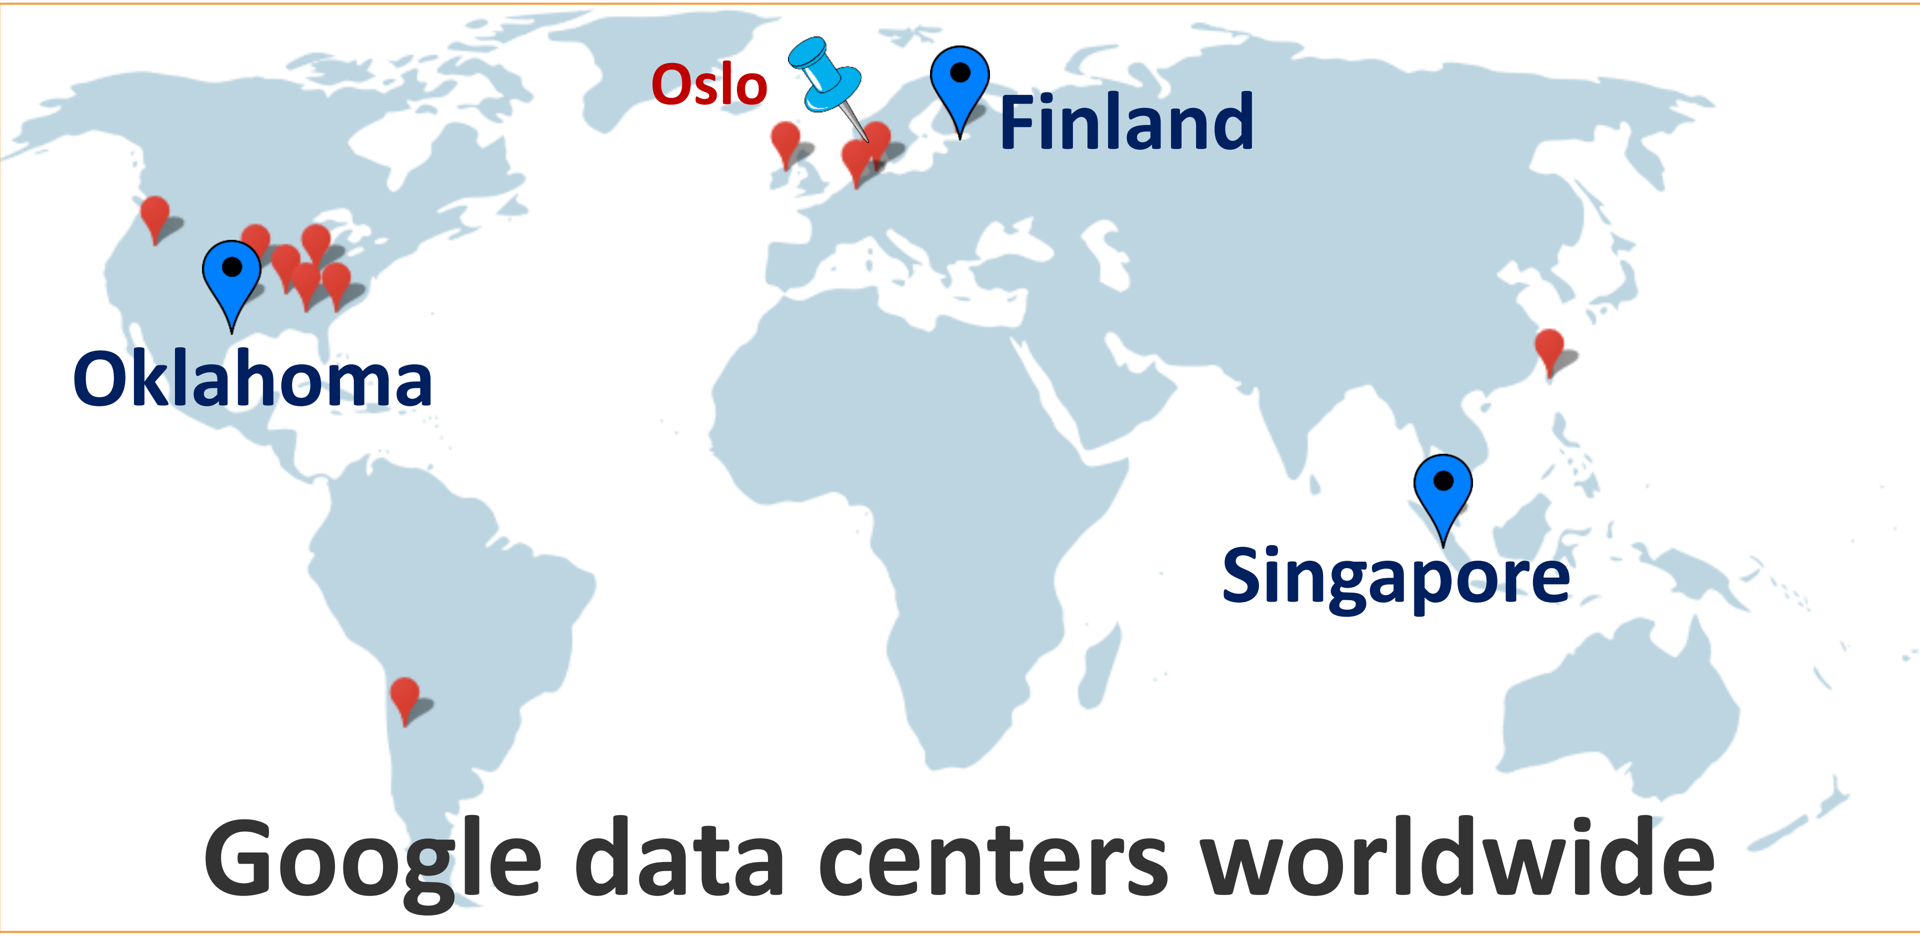

An Example: allocating tasks to minimize energy cost (similar as demand response concept in Lecture 5)

Three locations have different electricity prices. We need to allocate computation tasks to different locations to minimize the total energy cost

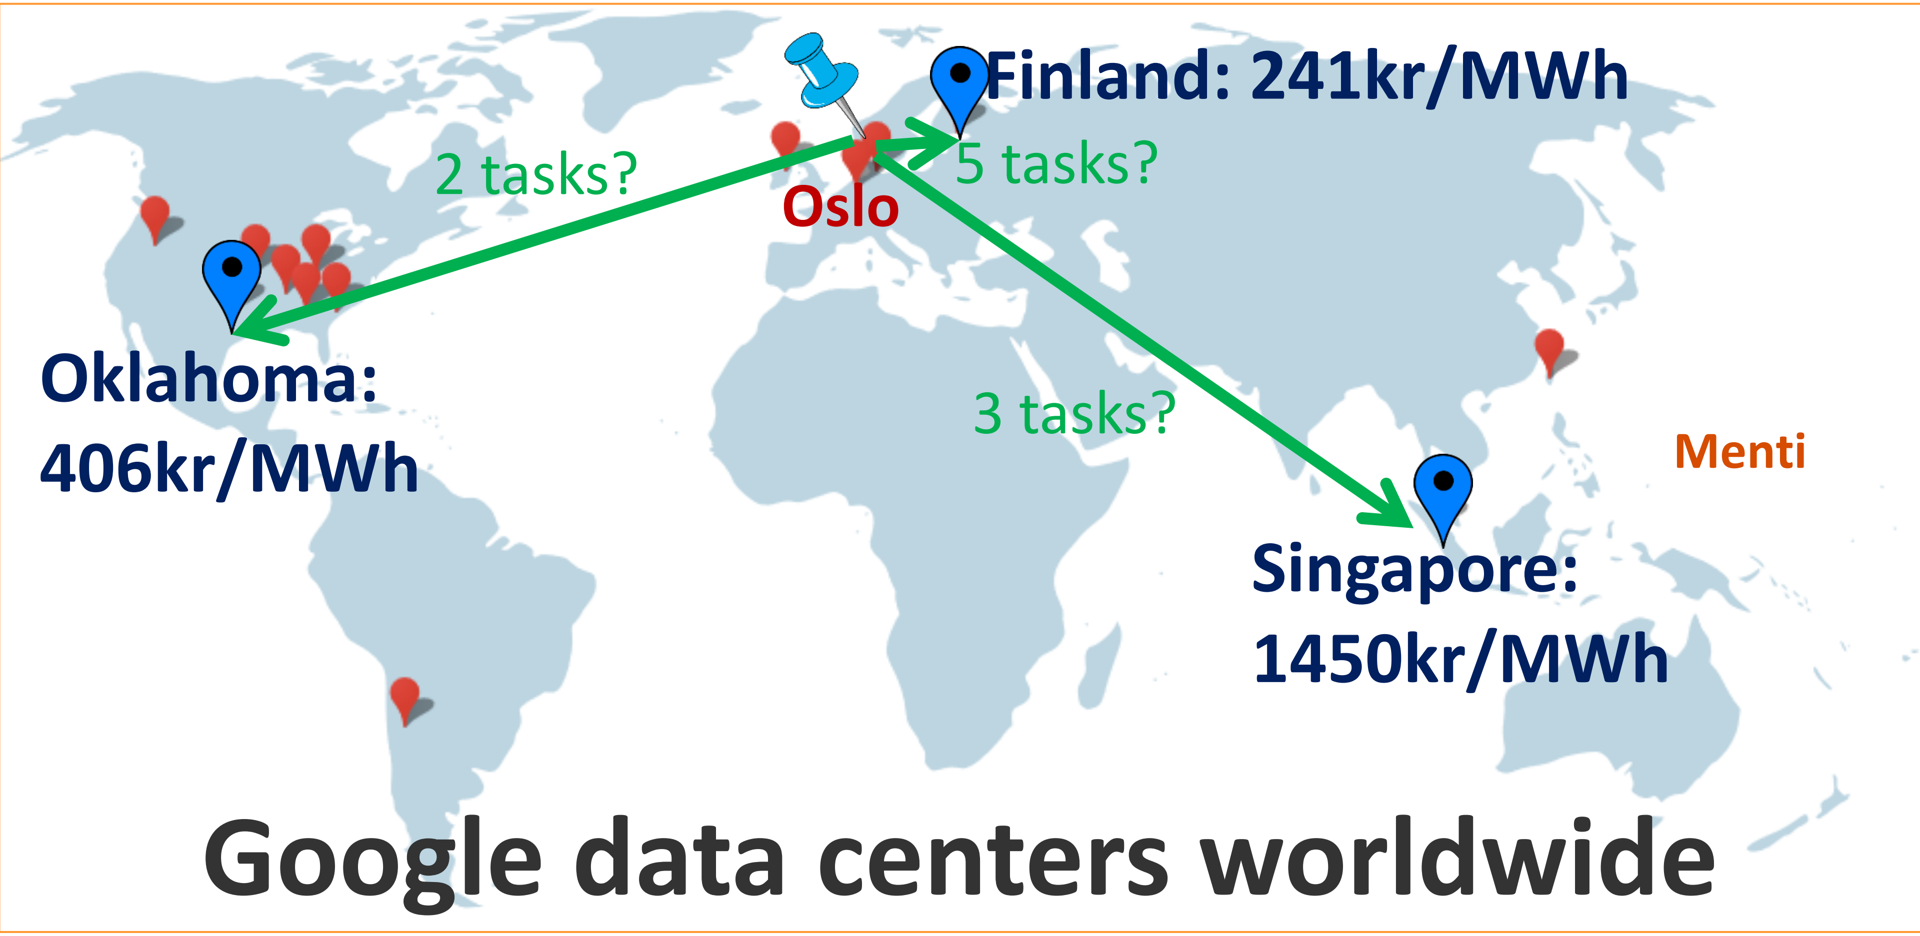

Inter Data Centers Level: exploiting location diversity

- **Different geographical distance from service requests**
  - **Your requests may be directed to the data center in Finland since you are geographically close to Finland**
- **Different working load**
  - **Your requests may be directed to the data center where there is low working load**
- **Different electricity price**
  - **Your requests may be directed to a data center with low electricity price to reduce energy cost**
- **Different renewal energy sources availability, greenness, and CO2 emission**
  - **Your requests may be directed to the data center where renewable energy is readily available**

An Example with N=3 Data Centers

Requests at a front-end server

A client request is first handled by a front-end server, then it is forwarded to one server at a specific location of the data center to be processed.

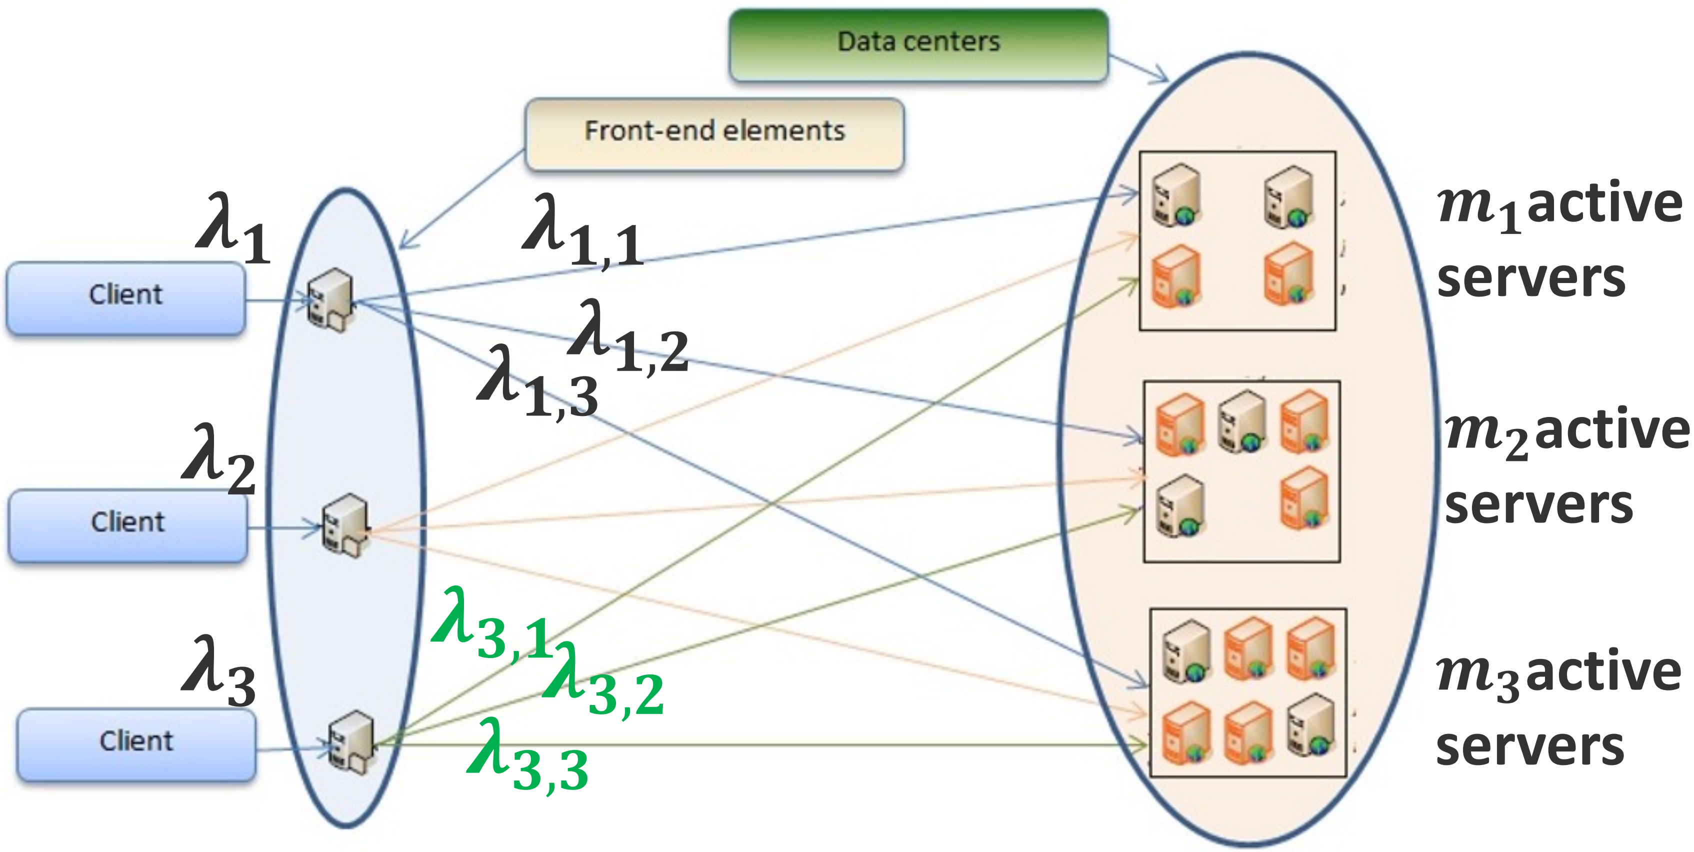

C=3 front-end servers (e.g., servers at UiO)

N=3 Data Centers (e.g., Google Map/Searching servers)

Energy Cost Problem – Workload Constraint

Requests at front-end server j

arrival rate of the tasks requests from front-end server *j* to data center *i*

$$\sum_{i=1}^{N} \lambda_{j,i} = \lambda_j; j = 1,2 \ldots C$$

$$\begin{cases} \lambda_1 = \lambda_{1,1} + \lambda_{1,2} + \lambda_{1,3} \\ \lambda_2 = \lambda_{2,1} + \lambda_{2,2} + \lambda_{2,3} \\ \lambda_3 = \lambda_{3,1} + \lambda_{3,2} + \lambda_{3,3} \end{cases}$$

Number of active servers

Total number of servers M_i at each data center i

$$m_i \leq M_i; i = 1,2 \ldots N$$

$$\begin{cases} m_1 \leq M_1 \\ m_2 \leq M_2 \\ m_3 \leq M_3 \end{cases}$$

Energy Cost Problem – Delay requirement

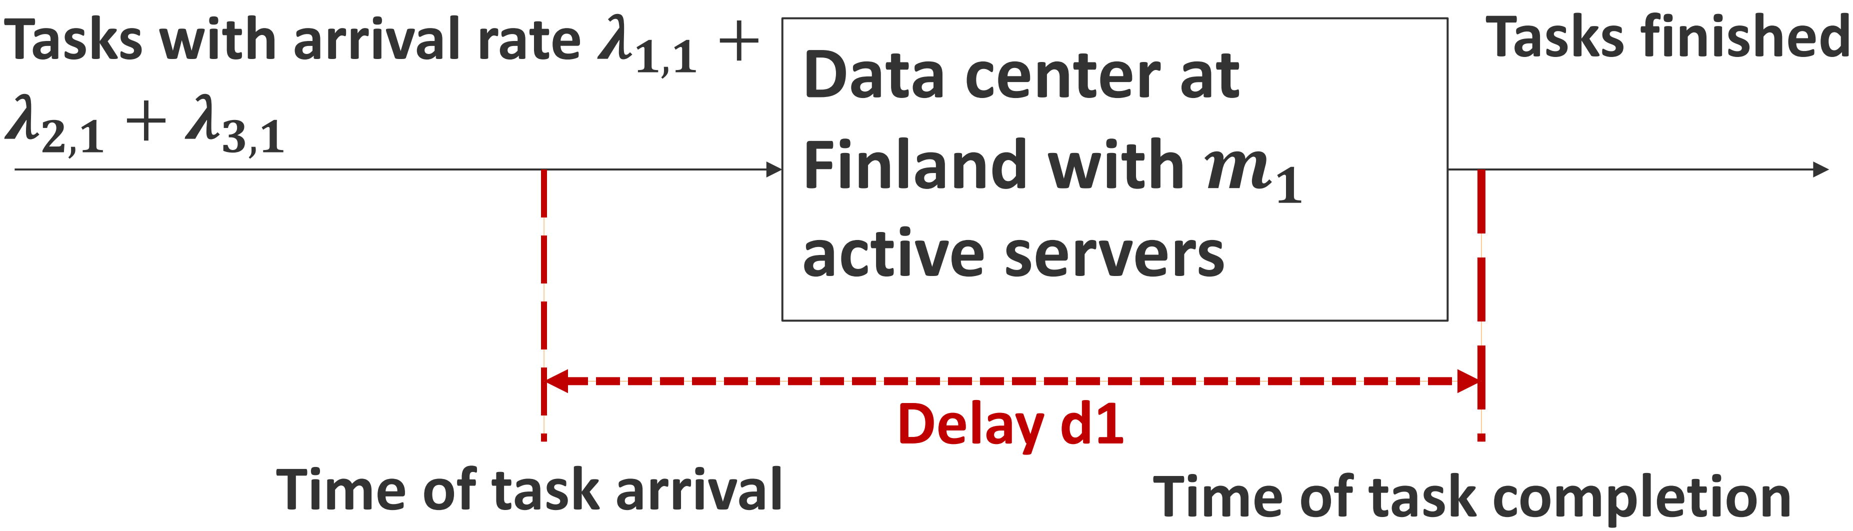

- m active servers in a data center will deal with tasks from all customers.
- When all requests go to Finland with the lowest electricity price, the energy cost may be very low. However, the customers may have to wait for a very long time to get response which is not favorable.
- Hence, there is a constraint related to the allowed latency, e.g., 1 sec. That is, the service delay d_1 should not be longer than the allowed latency D_1.

Energy Cost Minimization

Energy cost minimization:

Price of electricity at data center i at h

$$\min_{\lambda} \sum_{h=1}^{24} \left(m_1 P_1^h + m_2 P_2^h + m_3 P_3^h \right)$$

Cost from all data centers in hour *h*

Subject to

$$\sum_{i=1}^{N} \lambda_{j,i} = \lambda_j; j = 1,2 \ldots C$$

$$m_i \leq M_i; i = 1,2 \ldots N$$

$$d_i \leq D_i; i = 1,2 \ldots N$$

Q: Is this a linear programming optimization problem?

The delay constraint is not linear.

More interest? Need to read a queueing theory book and M/G/m queueing model

MORE CONSIDERATIONS…

Apple uses solar energy to power Data Centers

Apple plans to build a 200 MW solar array to power its data center in Reno, Nevada
Q: what are the main challenges data centers face in using solar power?

Reliability

- **Because of intermittency with power generation from solar panels, data centers must be prepared to store the power or supplement it.**
- **Centers typically use electricity from solar panels to power daytime operations, and apply any excess to charge UPS (uninterruptible power supply) systems.**
- **The real challenge is not to find alternatives to supplement solar power, but to have the intelligence to optimize the use of each power source.**

Quality: The second issue must be addressed in order to protect sensitive IT assets i.e., power quality to stabilize against fluctuations is important.

Google's DeepMind cut its Energy Bills by 40%

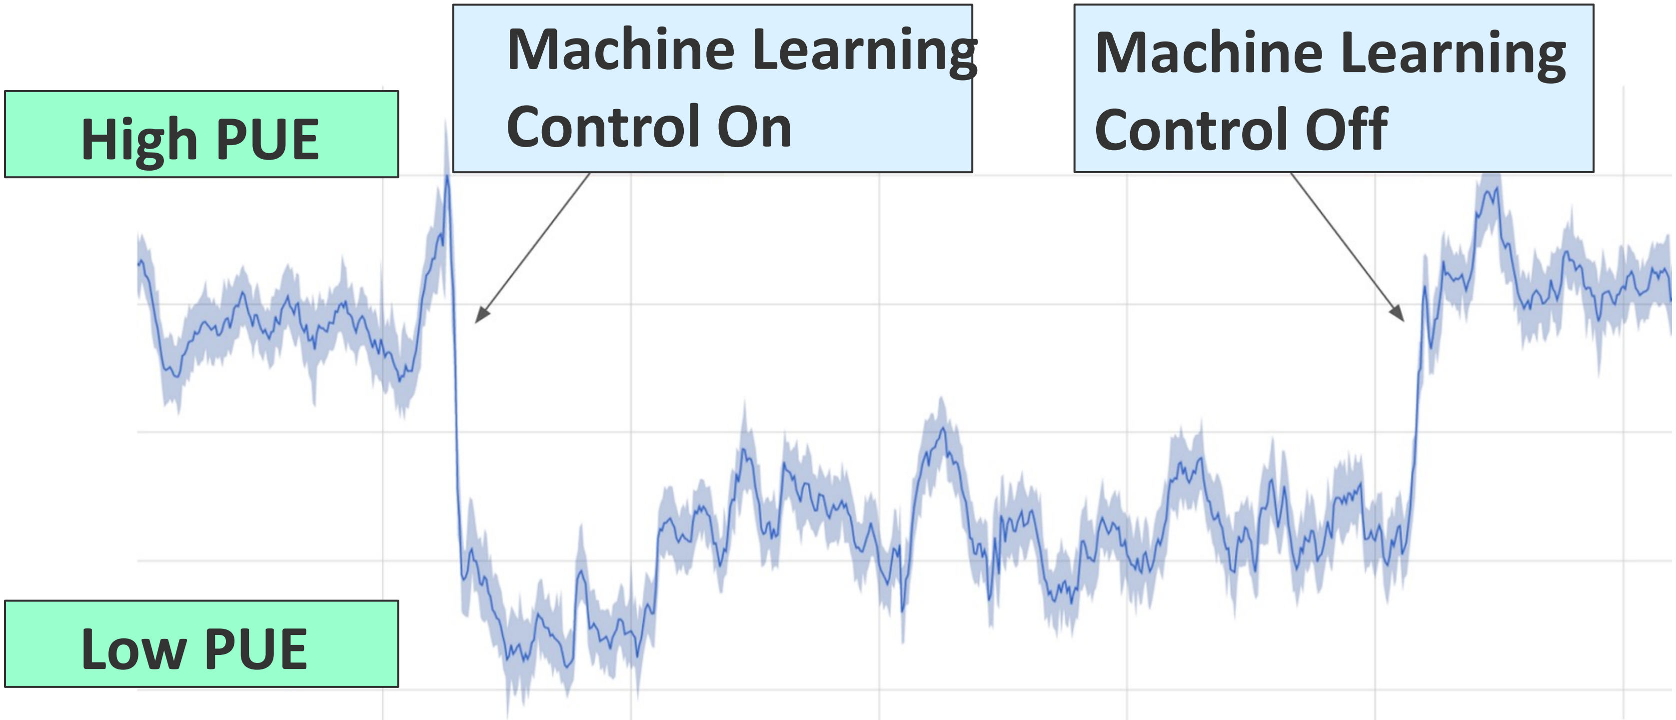

"…build a better predictive model that essentially uses less energy to power the cooling system by more accurately predicting when the incoming computation load is likely to land,"

The DeepMind team collected data for five years and created a prediction model for how much energy would be needed by the data center based on the amount of server usage that was likely.

Green mini-Data Center in Self-driving Cars

- Ford's self-driving car road test
- **Q:** why do we need mini-data center in self-driving car?

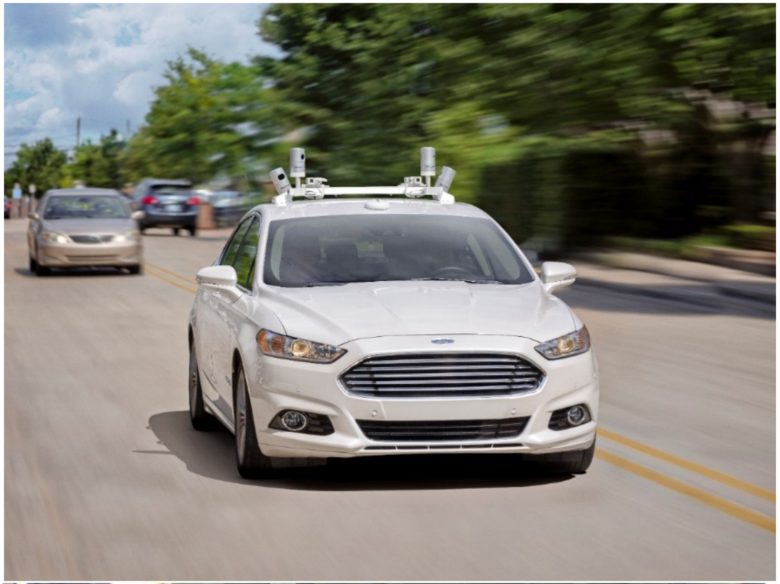

- A self-driving car generates **1 GB/second data** and it requires **real-time processing** to make correct decisions.
- Computation equipment in the car

NEWS FEATURE | 12 September 2018 | Correction 13 September 2018

How to stop data centres from gobbling up the world's electricity

The energy-efficiency drive at the information factories that serve us Facebook, Google and Bitcoin.

Nicola Jones

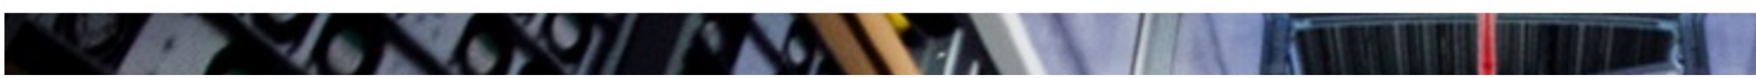

https://www-nature-com.ezproxy.uio.no/articles/d41586-018-06610-y

Data centres contribute around 0.3% to overall carbon emissions, whereas the information and communications technology (ICT) ecosystem as a whole — under a sweeping definition that encompasses personal digital devices, mobile-phone networks and televisions — accounts for more than 2% of global emissions.

That puts ICT's carbon footprint on a par with the aviation industry's emissions from fuel.

But one of the most worrying models predicts that electricity use by ICT could exceed 20% of the global total by the time a child born today reaches her teens, with data centres using more than one-third of that (see 'Energy forecast')[1].

If the computationally intensive cryptocurrency Bitcoin continues to grow, a sharp rise in energy demand could come sooner rather than later (see 'The Bitcoin bite')

References

- **“Er, what is a data center”, Intellect, TechUK.** **https://www.techuk.org/images/documents/Data_Centres_-_CCA/Note_03_Er_what_is_a_data_centre.pdf**

- **A. Rahman, X. Liu and F. Kong, "A Survey on Geographic Load Balancing based Data Center Power Management in the Smart Grid Environment", in IEEE Communications Surveys and Tutorials (COMST), vo.16, no.1, 2014.**

Thank you for your attention!

UiO : University of Oslo

Sabita Maharjan, PhD
Professor,
Department of Informatics,
University of Oslo
Email: sabita@ifi.uio.no; Tel. +4740054152