Datacenter cooling

ELECTRONICS COOLING AT MACRO SCALE

JONAS GUSTAFSSON

Data Centers in a Nutshell

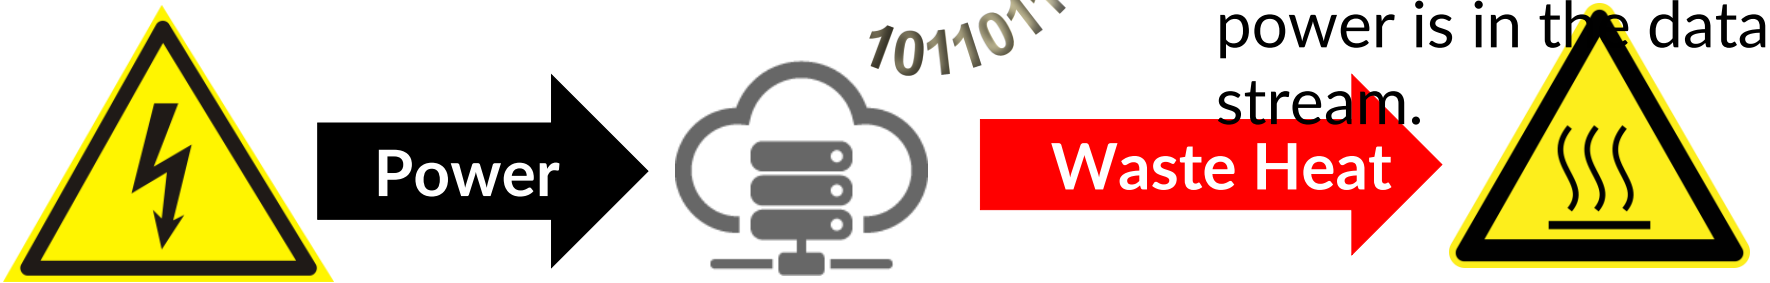

Today 0.03% of input power is in the data stream.

Today 99.97% of input power is in the thermal stream.

Based on:
Rolf Landauer, “Irreversibility and Heat Generation in the Computing Process,” **IBM J Res. Dev. 5, 183 (1961)**.
http://dx.doi.org/10.1147/rd.53.0183

RI. SE

Heat generation in data centers
- an increasing challange (opportunity)

- Moores law is slowing down… (debateable)
- Dennards scaling stopped around 2010
  - in every technology generation transistor density would double, circuit becomes 40% faster, while power consumption (with twice the number of transistors) stays the same.
- Multicore and virtualization has been the solution for the last ~10 years. But can it continue to match the ever increasing demand for digital services?

Energy Use Breakdown in Data Centers

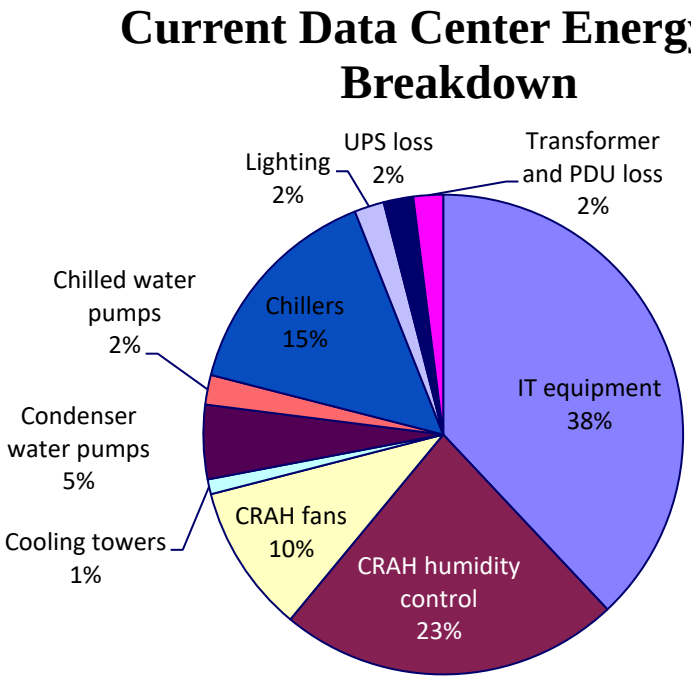

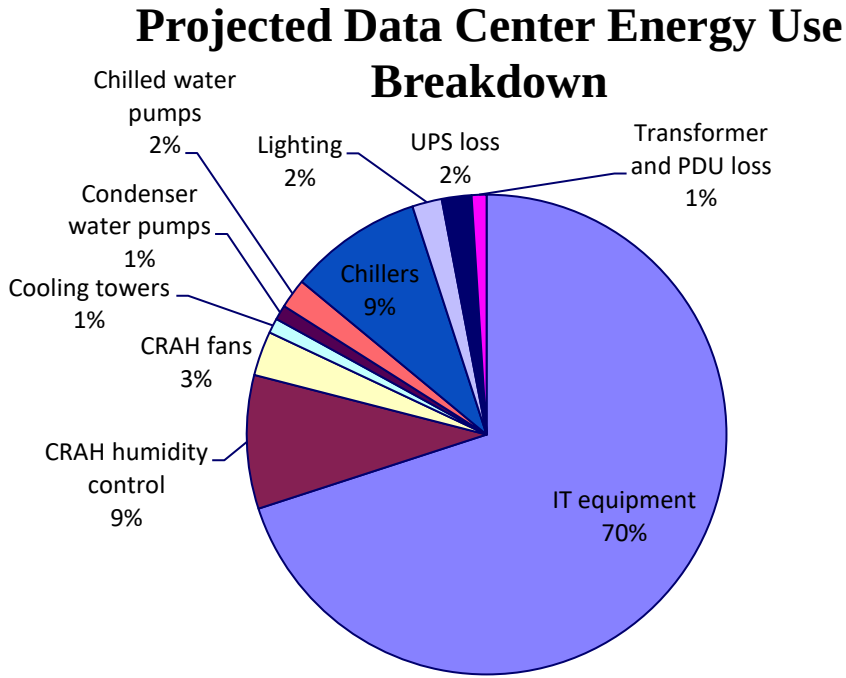

https://datacenters.lbl.gov/resources

Depends on temperature of operation and geographical ambient.

Cooling is a waste of energy and money!

Power Usage Effectiveness, PUE

$$PUE = \frac{Total\ Facility\ Power}{IT\ Equipment\ Power}$$

ICS › 35 › 35.020

ISO/IEC 30134-2:2016

Information technology — Data centres — Key performance indicators — Part 2: Power usage effectiveness (PUE)

Getting cooling where it is needed is a problem of scales.

- Microprocessors are small

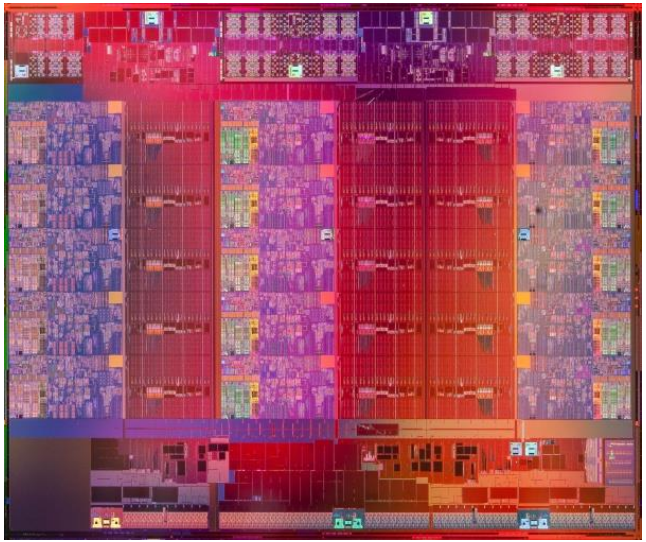

170mm^2

Intel Haswell EX (15 core)
approx 160W

- Data centers are large

13,823,530 fold increase in size!

2,350,000,000mm^2

2MW facility with approx 400 cabinets.

Datacenter cooling principles
Heat generation to heat rejection (or recovery)

Air cooling

- Server cooling
  - Remove air from CPU and other heat generating components
  - TIM, Heat sink design, server internal fan control
  - Maintain CPU temperature below specified threshold.
  - Server manufacturers domain.
  - Open Compute Project has stired things up!

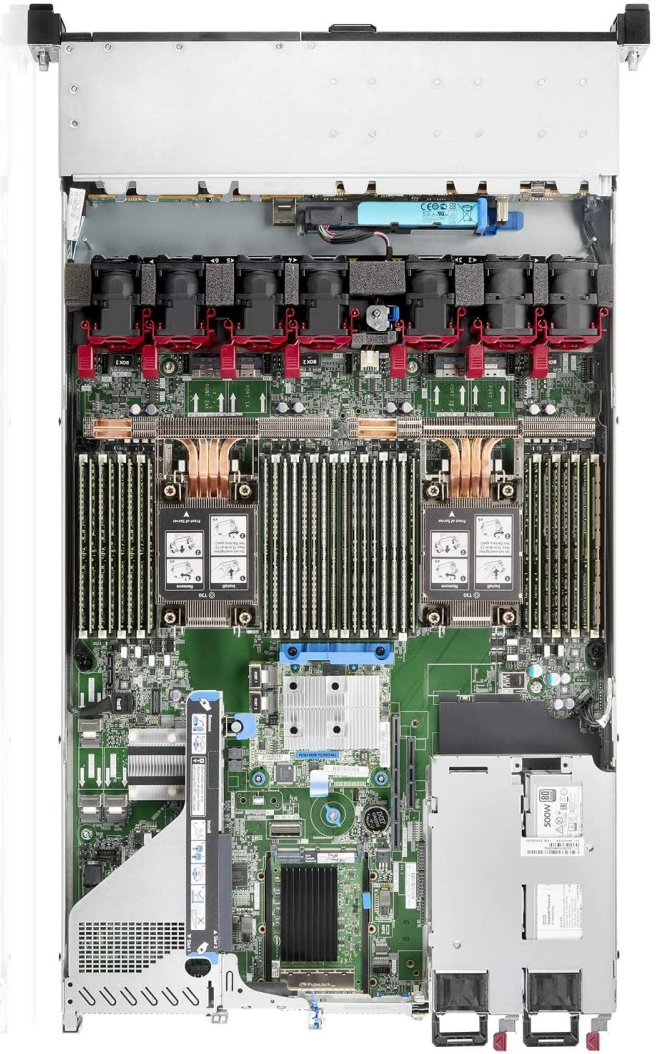

Air cooling

- Space cooling
  - Servers in racks, and racks in rows
  - Supply cold air through raised floor (often)

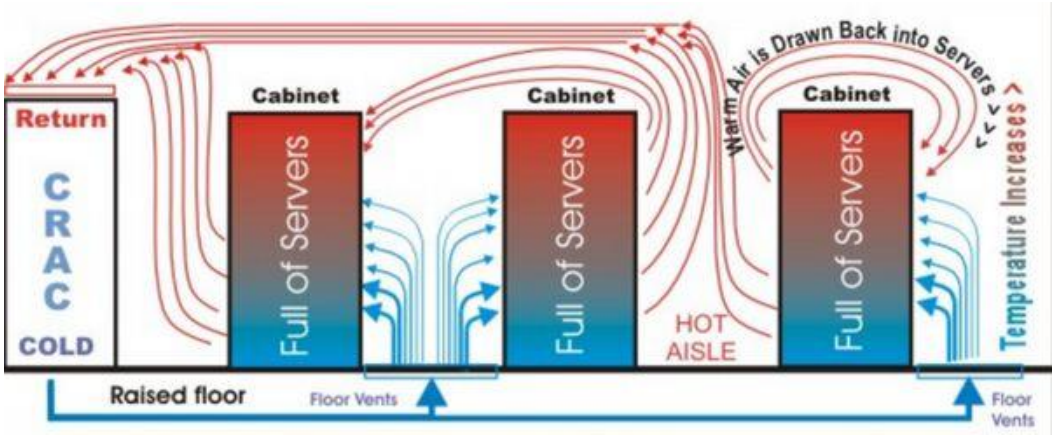

Air cooling

- Space cooling
  - Containment to reduce recirculation and increase efficiency

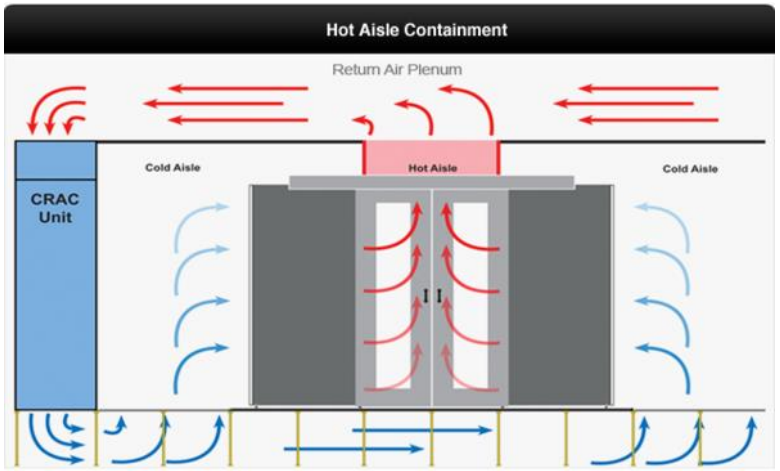

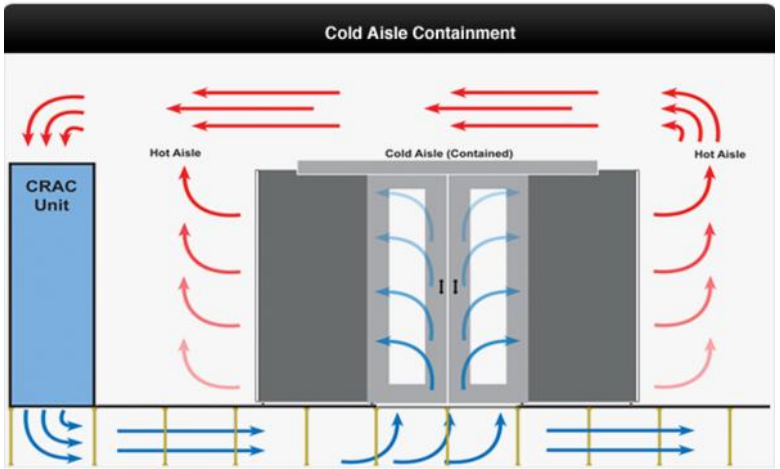

Air cooling

- Heat rejection
  - if (waste heat >> warmer than outside air); reject it (free cooling)
  - But! Incoming air also needs to be ”cold enough” (<~28°C)
    - With some themal losses, it often becomes challangeing to run free cooling all year around at most places.
    - Warmer server conditions also increases losses in CPUs and other chips. (good for PUE, but bad for power consumption!)
  - Indirect vs. Direct (fresh air) cooling.
  - Chillers are often needed to bring down the temperature of incoming air.
    - Datacenters are working intensly to remove compressors.

Liquid cooling

Pros.

- Higher IT densities
- Higher waste heat temperature
  - Enables more free cooling, and/or better heat recovery opportunities
- Easier to pump liquid than air.
- Fever moving parts

Higher specific heat capacity c_p, and higher density

Conc.

- System complexity
- Risk of leaks
- (Concurrent) maintenance
- Service friendlyness

Liquid cooling

Ways of implementing liquid cooling

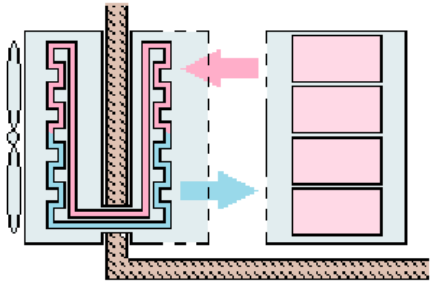

Traditional ”air cooling” (CRAH/CRAC)

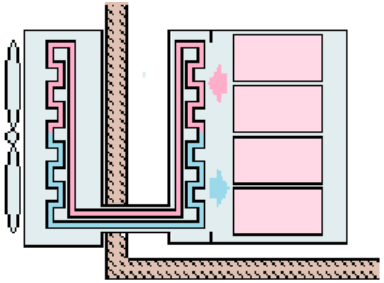

”Rear-door” cooling

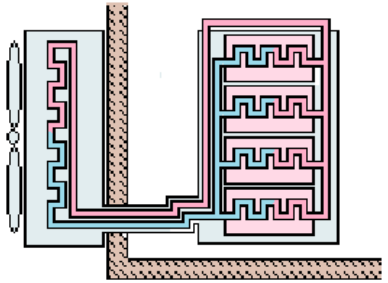

On-chip cooling

Source: www.deltapowersolutions.com

Source: nvent.com

Source: asatek.com

Liquid cooling

Ways of implementing liquid cooling – immersion systems

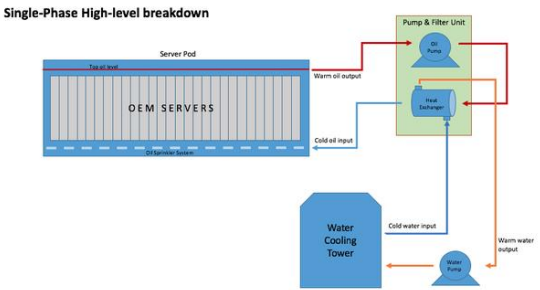

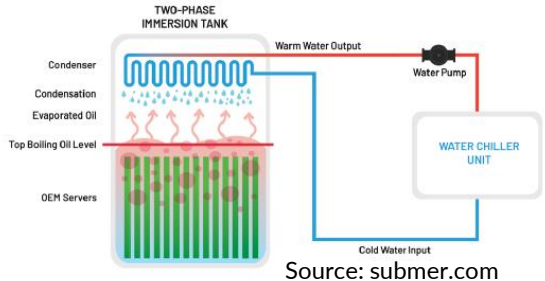

Source: submer.com

Single-phase immersion

Two-phase immersion

Other types…

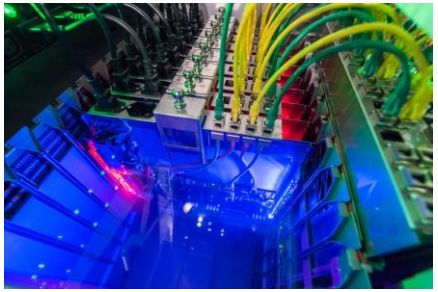

Source: asperitas.com

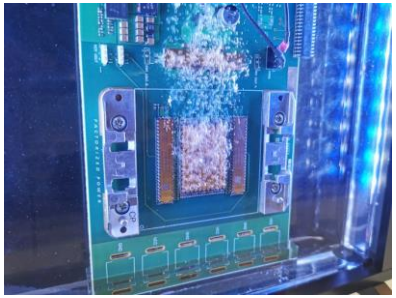

Source: https://www.anandtech.com/

Source: iceotope.com

Facility vs IT

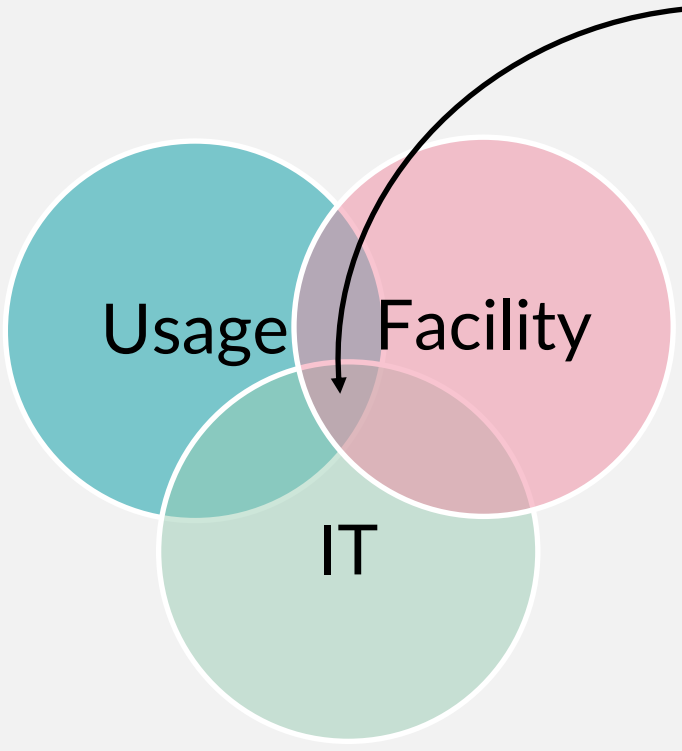

Holistic cooling

- Philosofy: Holistic
  - *”characterized by the belief that the parts of something are intimately interconnected and explicable only by reference to the whole.”*
- Achieve maxim (cooling) efficiency independently of IT-load
  - Why should CPU temperature vary?

Heat recovery examples

- District heating
  - Huge opportunities
  - Difficult to meet temperature demands.
- Greenhouses
  - Promising
  - Close proximity of DC
- Industrial processes

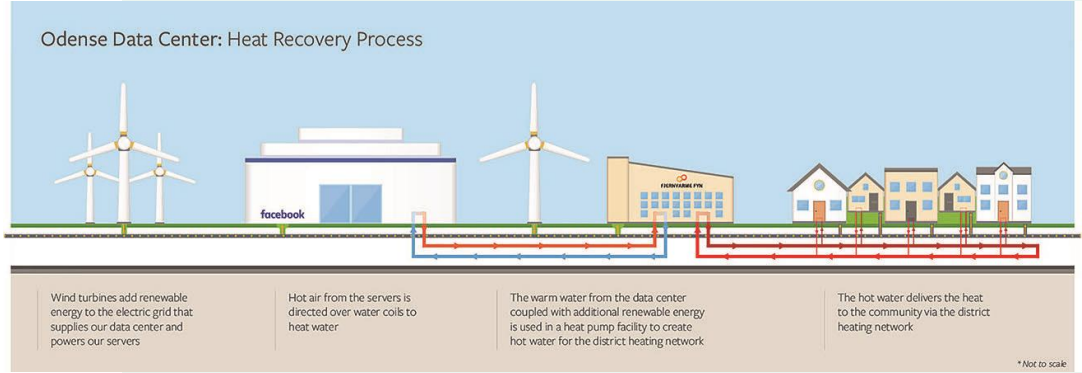

Research challanges, examples

- Optimize heat transfer, from source (CPU) to rejection or recovery.
  - Thermal interface material
  - Heat exchangers
  - Heat recovery application
  - …
- Further integrate the different ”parts” of the data center
  - Facility, IT and Software

Maintain as high temperature as possible as far as possible in the heat rejection (recovery) process

Cooling of electronics does not always end at the heat-sink

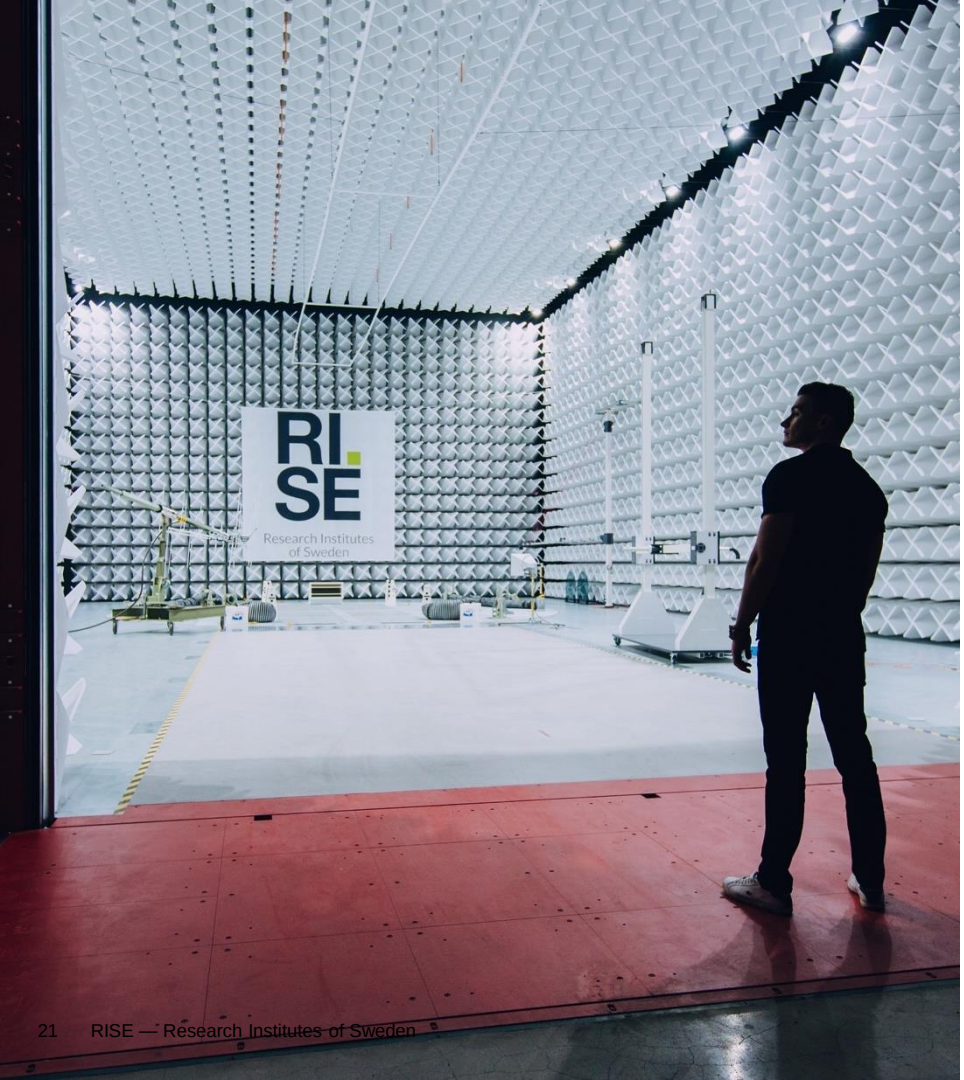

We at RISE ICE Datacenter

- Run experimental datacenter
- Test-rigs for data center equipment
- Direct projects for companies
- Together with companies find R&D&I funding for ”cool” ideas
- New partner program for fast and easy access of our results.

Partnerin program - Introduction

Description

- RISE ICE is offering a program for partners to engage in our research and innovation. The objective is to give **easy and early access** to results, data and resources for research and innovation projects, with the provision to **deep access** to additional results, data and resources.
- The partnership can extended with additional direct funded projects or with public funded projects with co-funding from the partner. See last slides.
- Legal framework to cover partnership available. NDA for the partner set-up to enable project discussion. Partner program agreement for the project year.

Three levels of yearly memberships

- Platinum Member - €100 000
- Gold Member - €50 000
- Silver Member - €25 000

	Access	Platinum	Gold	Silver
Webinars	Early and easy	2 exclusive +1 common	1 exclusive +1 common	1 common
Expert access	Deep	X	x	
Follow-up	Early	Every other month	Quarterly	Half year
Package of output	Easy	X	X	X
Junior research work	Deep & innovation	4 months	2 months	1 month
Data access	Easy	8 weeks	4 weeks	2 weeks
Thesis project	Easy & innovation	Master thesis student	Master thesis student	
Logos	Visibility	A front place	Next to front	Next to front

RI. SE

Jonas Gustafsson

Jonas.gustafsson@ri.se
070-333 00 35