A comparison between DataCenter Liquid Cooling Solutions

Paolo Bianco
Sr Systems Engineer, Compute & Networking,
Italy HPC Lead
paolo.bianco@dell.com

intel

DELL EMC

Today Challenge: Economical

Average PUE of your largest data center:

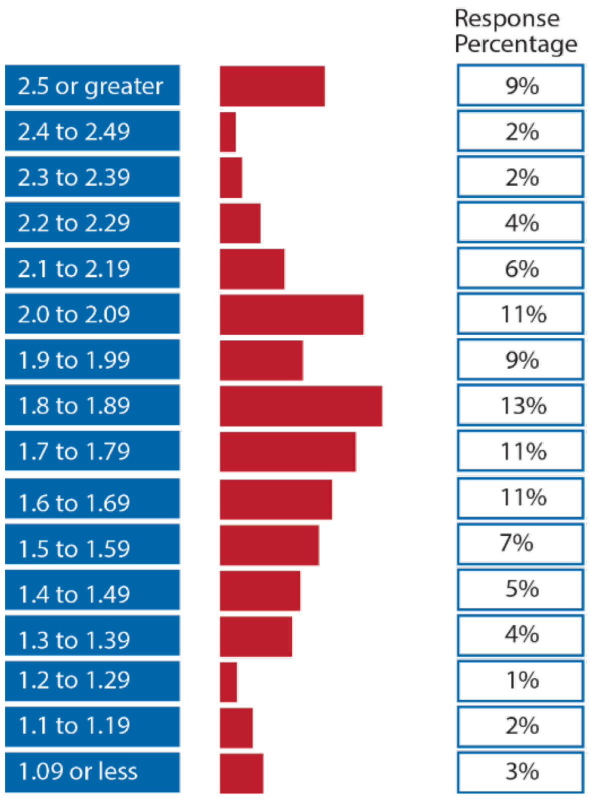

AVERAGE PUE 1.8 – 1.89

Source： Uptime Institute survey of over 1100 data centers

$$\text{Data center Power usage effectiveness(PUE)} = \frac{\text{Total data center power (kw)}}{\text{Total IT Power (kw)}}$$

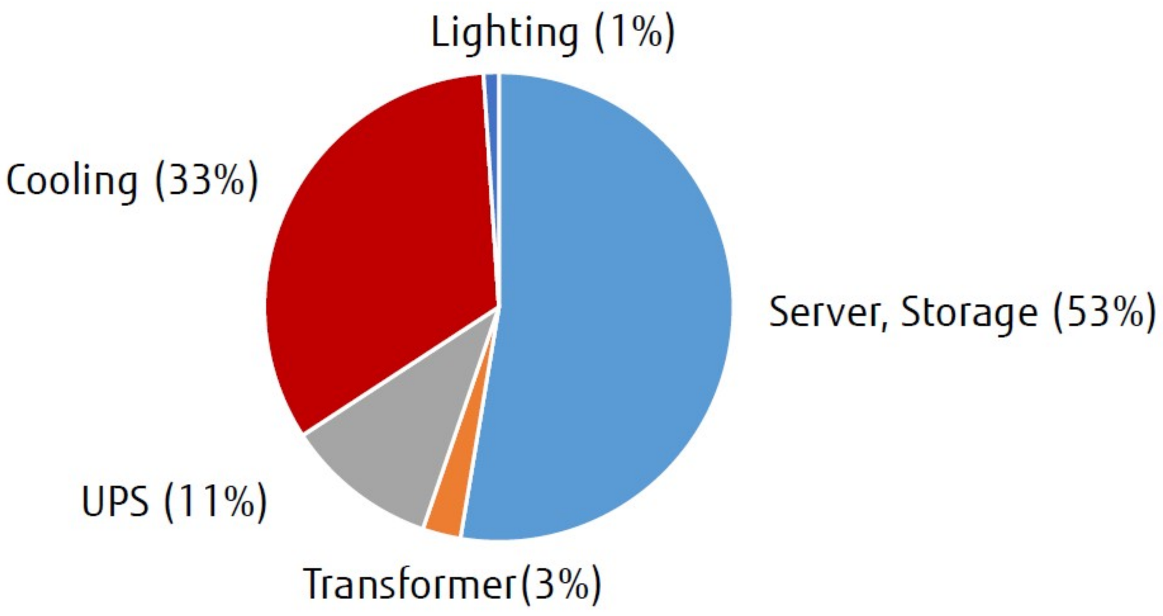

A Shehabi: LAWRENCE BERKELEY NATIONAL LABORATORY "United States Data Center Energy Usage Report", 2016.

- The energy consumption data center for cooling is quite large.
- How to achieve low PUE and optimized TCO? That has become a new challenge.

A new Challenge: Technical

- *Every 1 kiloWatt (kW) of rack power needs 1kW of cooling.*
- Common standard density compute Rack loadings are about 10-12kW per Rack
  - only 25U or so in use
- Common design max loadings are about 12-20kW per Rack
- To ensure mid-term future proofing then you need to be able to cope with 40-45kW per rack.
  - Today, a 42U Rack full of C6420 2S nodes (80 nodes) can draw up to 40kW !
- Full Air-Cooling is not going to meet the heat dissipation demand any more

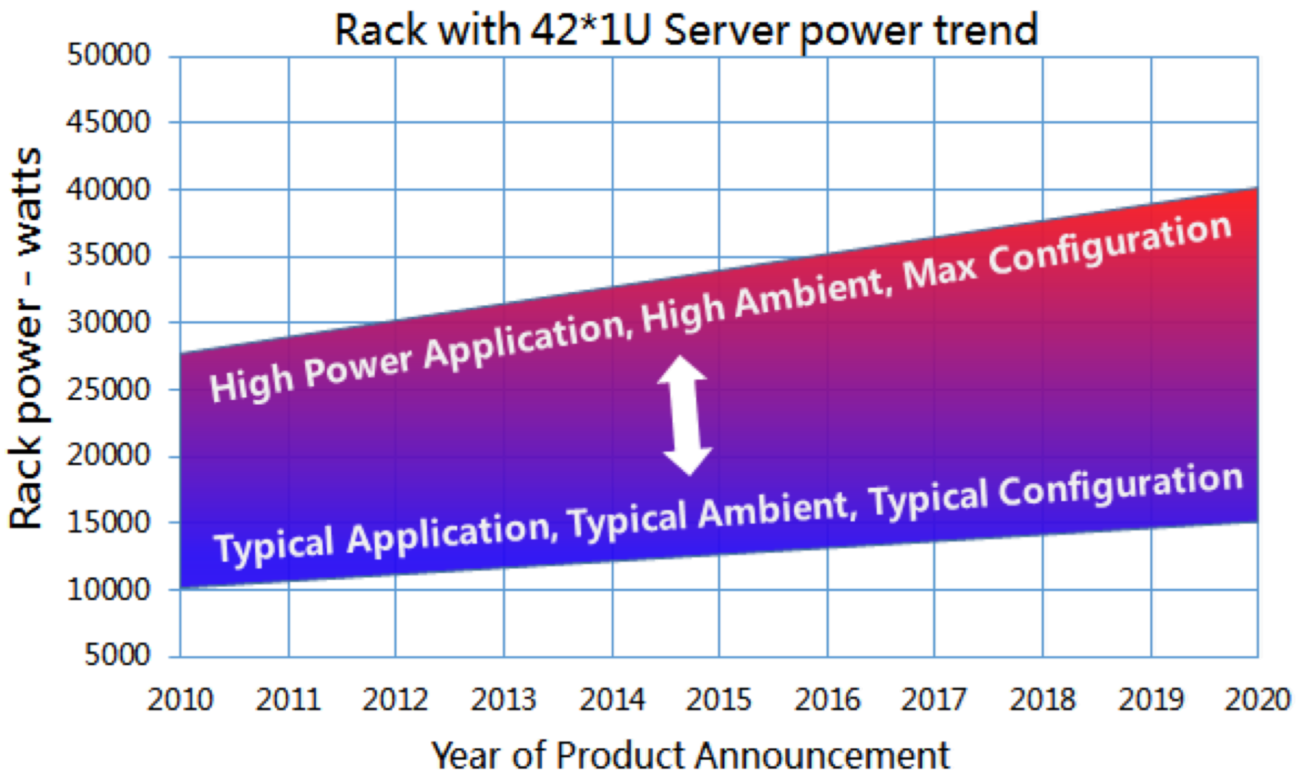

Datacom Equipment Power Trends and Cooling Applications --ASHRAE

TDP trends

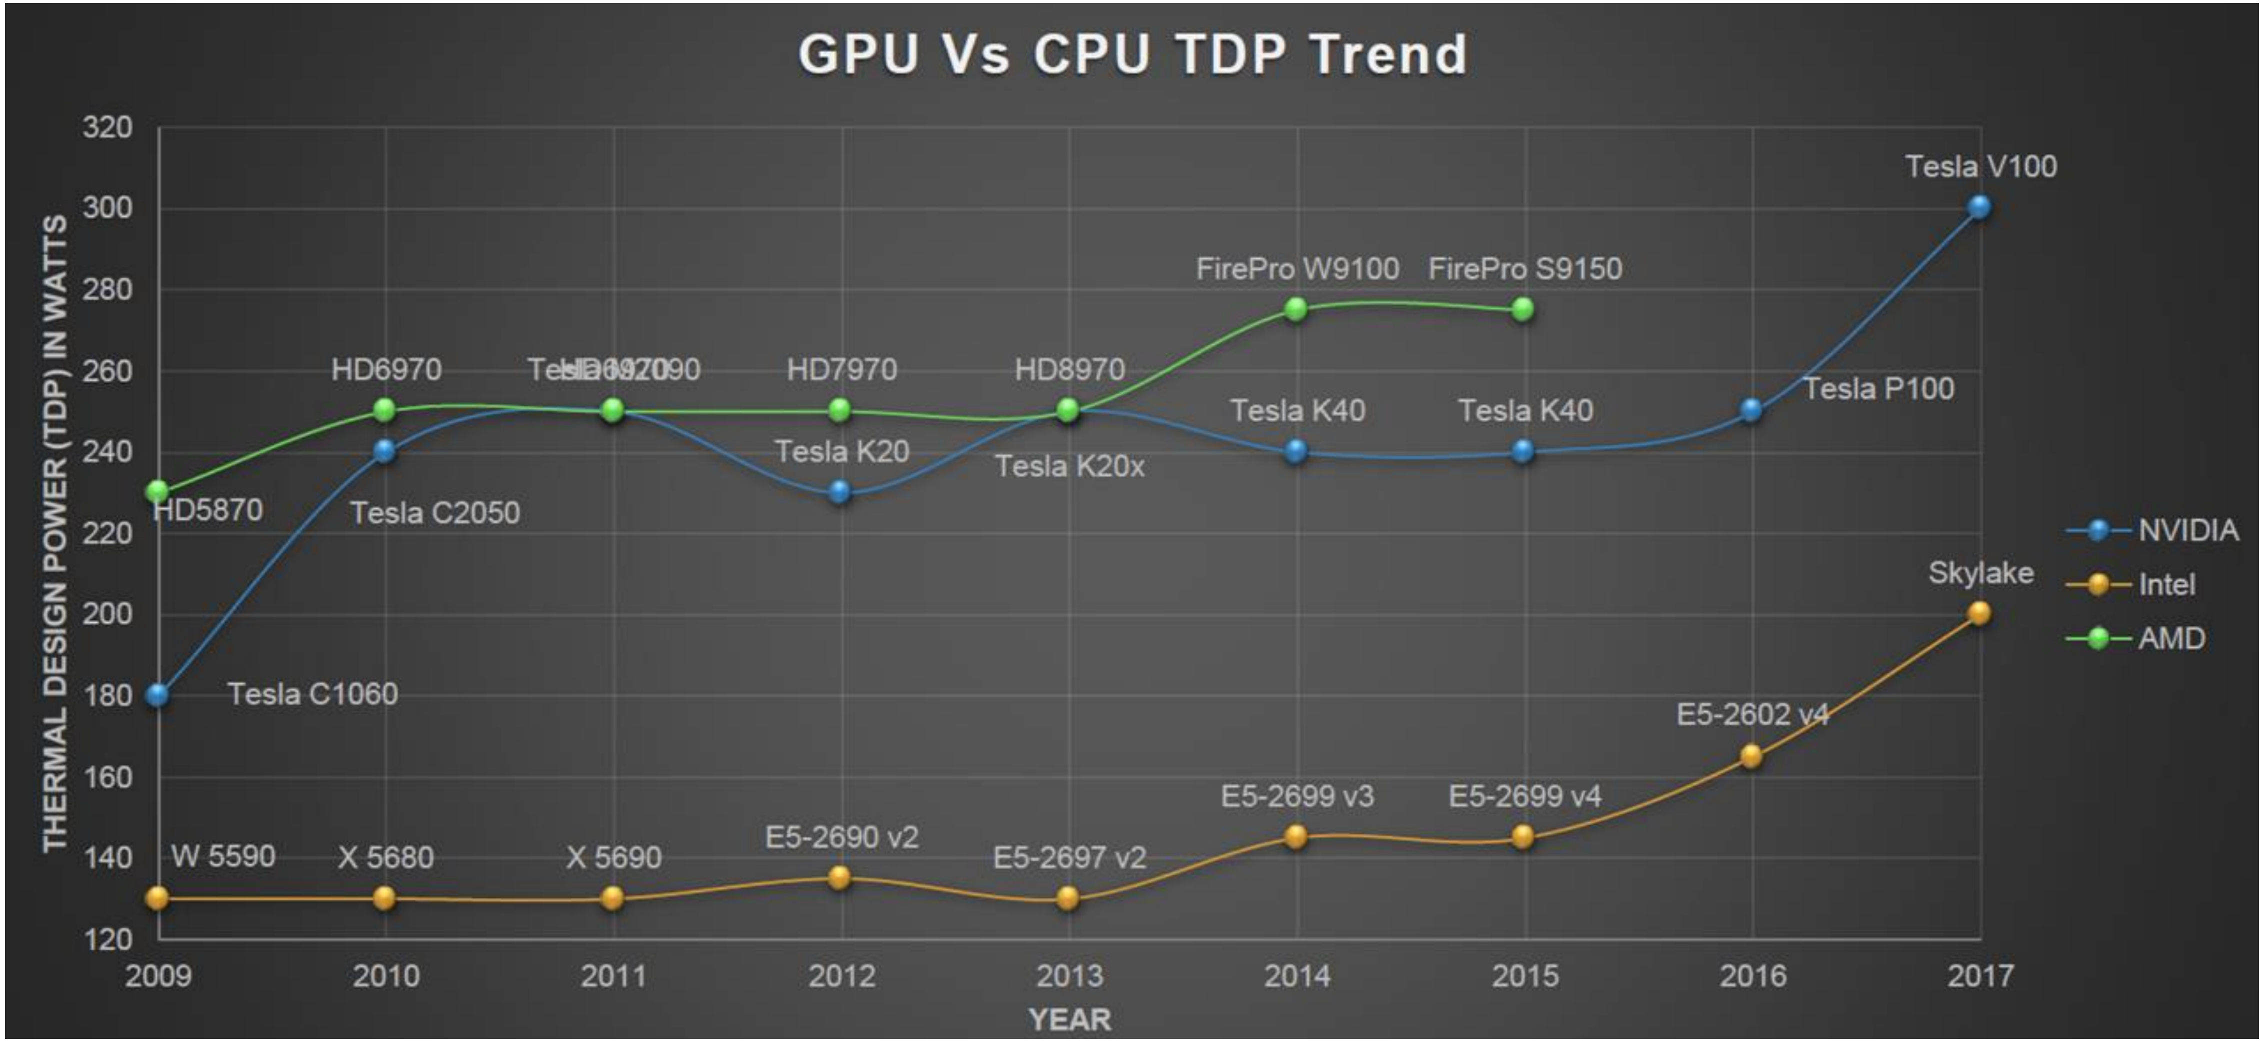

Source： Alibaba.com – Immersion cooling for Green Computing - OCP2018

DELL EMC

Technology trends

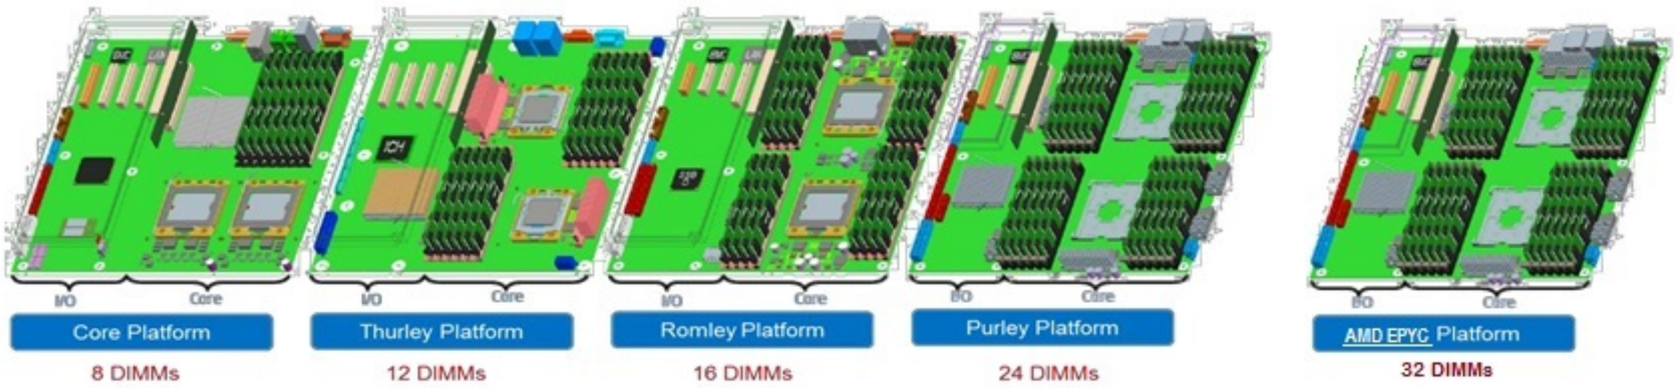

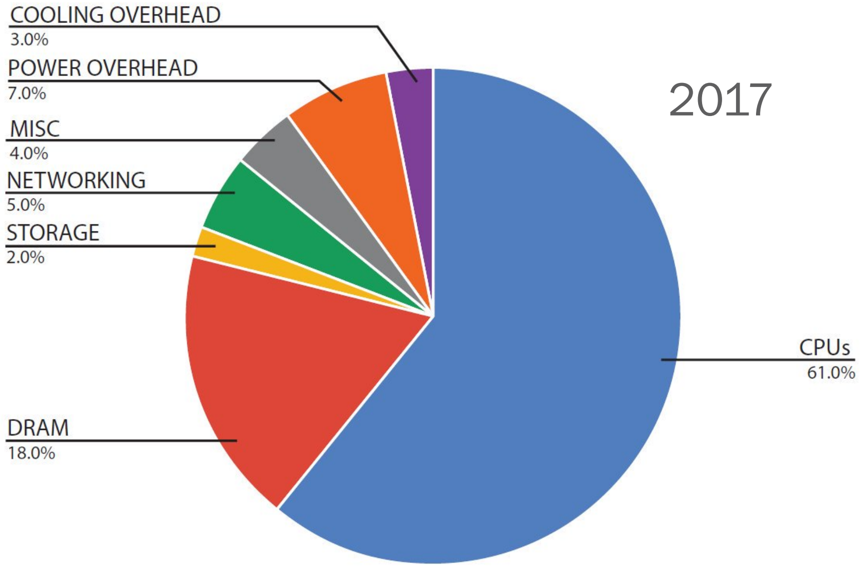

Figure 1.8: Approximate distribution of peak power usage by hardware subsystem in a modern data center using late 2017 generation servers. The figure assumes two-socket x86 servers and 12 DIMMs per server, and an average utilization of 80%.

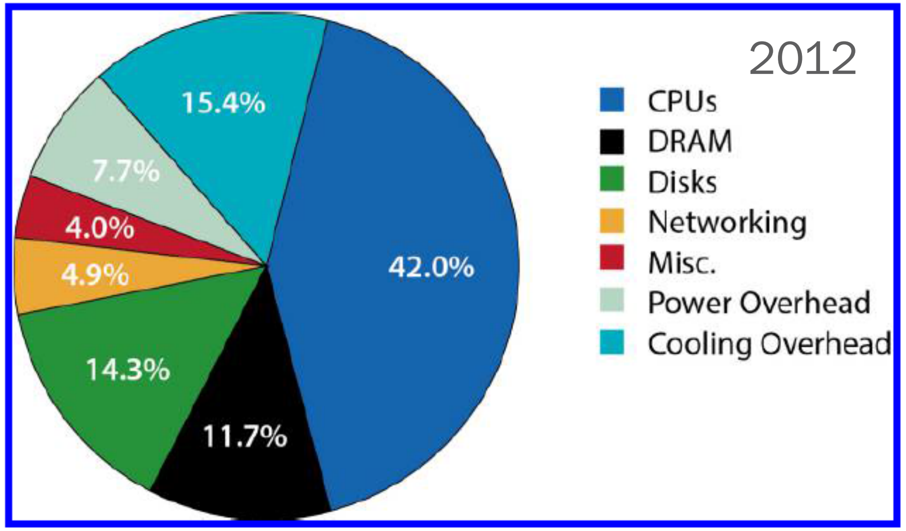

Example breakdown of peak power usage of a datacenter using 2012 generation servers. Assumes two-socket x86 servers, 16 DIMMs and 8 disk drives per server, and an average utilization of 80%

Chung-Ta King - Department of Computer Science - WAREHOUSE-SCALE COMPUTERS, National Tsing Hua University, Taiwan

Source : Barroso, Holzle, Ranghanathan – The Datacenter as a Computer – M&C

Increasing Efficiency

- Efficiency increase = Reduction in Total Cooling Power consumption. Where to act ?
  - Compressors/Heat Exchangers (better HW, ΔT)
  - Pumps (reducing speed/ reducing mass / eliminating)
  - Fans (better HW, reducing speed / eliminating)
  - Better coolant media (better heat carrying capacity / better fluid temperature)
  - Better air flow paths (compartimentazione)

- As heat density increases air becomes a less efficient coolant

- Liquid cooling provides a mean in which thermal resistance can be reduced dramatically
  - Principle: $Q_{LOAD} = mCp\Delta T = rVCp\Delta T$
  - Water has about 3.300 times higher heat carrying capacity than air!

- Rule of thumb: *the more the cooling liquid comes near the heat source, the better the efficiency is*

Liquid Cooling Landscape

- **Close Coupled Cooling**
  - In-Row®, In-Rack, Rear Door Exchangers

- **DLC (Cold Plates)**
  - Positive Pressure, Negative Pressure

- **Immersion Cooling**
  - Single Phase, Two-Phase

Close Coupled

- Air is still the only mean to cool board/chip
- Limited by existing chip maximum temps
- Requires additional fans

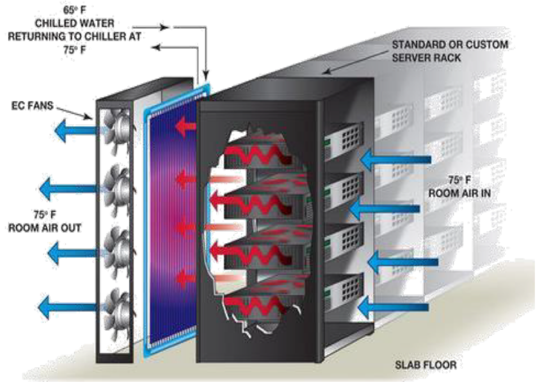

Cold Plates

- Individual Heatsinks
- Board-Specific
- Individual Chip Fluid-cooling

Immersion

- Specialty fluids ($$$)
- Can Require separate Cooling Coil
- Orientation sensitive

In-Row® Cooling

- Open-Loop Solution (air flows interact with the ambient room environment.)
- Row-Based Air Conditioning units are installed inside the rack rows.
- Heat removal of **up to 70 kilowatts per unit**
- Controlled in-row cooling
- Row air containment
- Modularity
- Typical **Inlet water temperature 12-23°C**
- Air flows generally follow short and linear paths, **reducing, in this way, the necessary power needed to rotate the fans** and increasing the energy efficiency.

Heat Exchanger and Fan Assembly | Front View | Rear View

Images: www.apc.com

In-Rack Cooling

- Closed-loop Solution (no interaction with the ambient room environment).
- The cooling system is joined to the server rack and both of them are completely sealed; the solid doors on the enclosure and In-Row Air Conditioners contain the air flow, directing cold air to the server inlet and exhaust air, by using fans, through the cooling coil.
- Heat removal of **up to 30-35 kilowatts per unit**
- One unit can serve up to two racks (up to 48U)
- Typical **Inlet water temperature 12-23°C**
- Right Equipment orientation (ie front-to-back) in the rack and proper cabling and positioning a must!

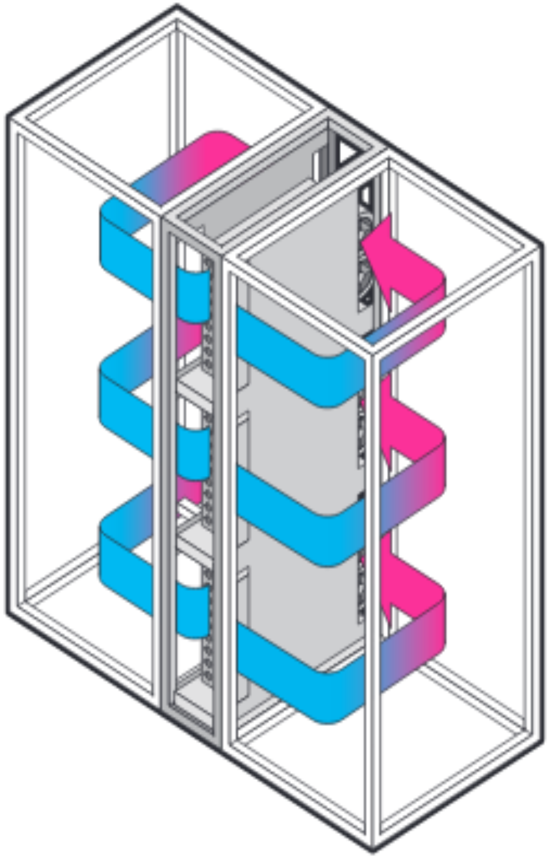

Rear-Door Heat Exchangers

- Open-Loop Solution (air flows interact with the ambient room environment.)
- This type of solution is based on the substitution of the rear door of an existing rack
- Heat removal of **up to 80 kilowatts per unit**
- One unit per rack (up to 48U)
- Rack adapters available for most standard racks
- Typical **Inlet water temperature 12-23**°C

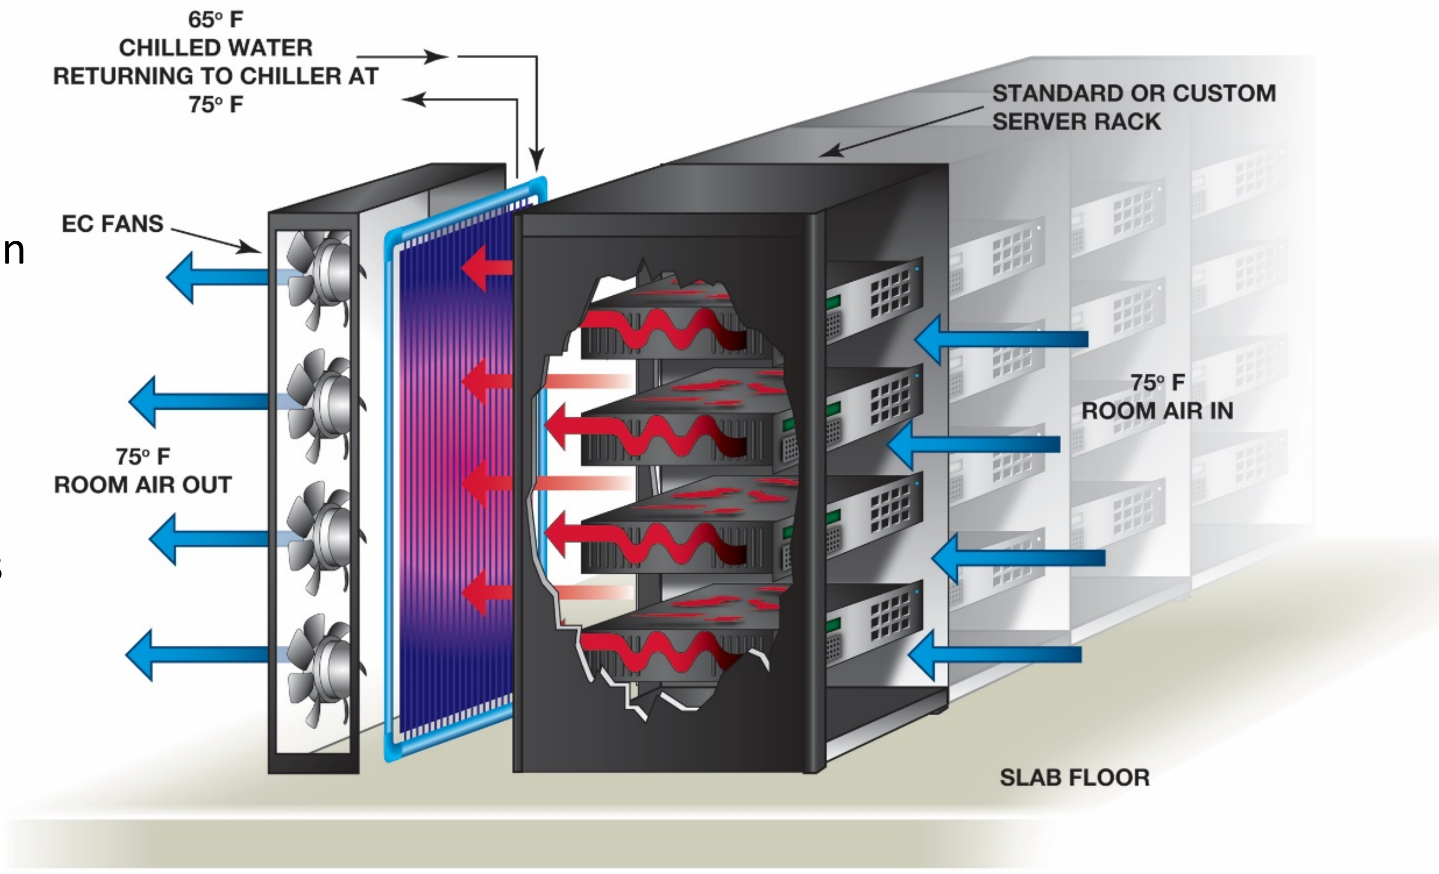

Exploded View of Active Rear Door Rack Cooling System

Why Direct Liquid Cooling?

- Heat density of servers is getting greater.
- Getting sufficient air flow in the confined space within servers is getting harder
- The latest Intel CPUs are > 200W & some cannot be air cooled
- New systems design do accept a higher thermal delta between Inlet and Outlet air temperatures of equipment (ASHRAE TC 9.9). This can lead to «normal» Tout of about 60°C (!)
  - Consequences for cables, optics and other components

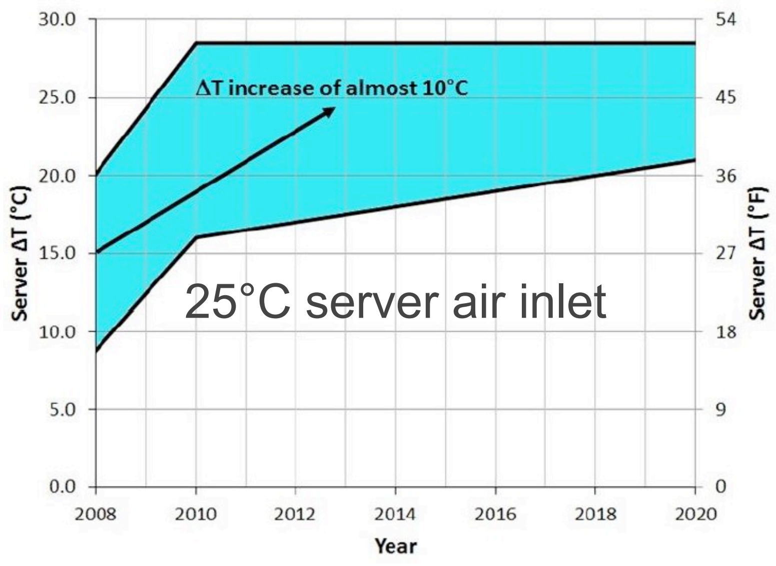

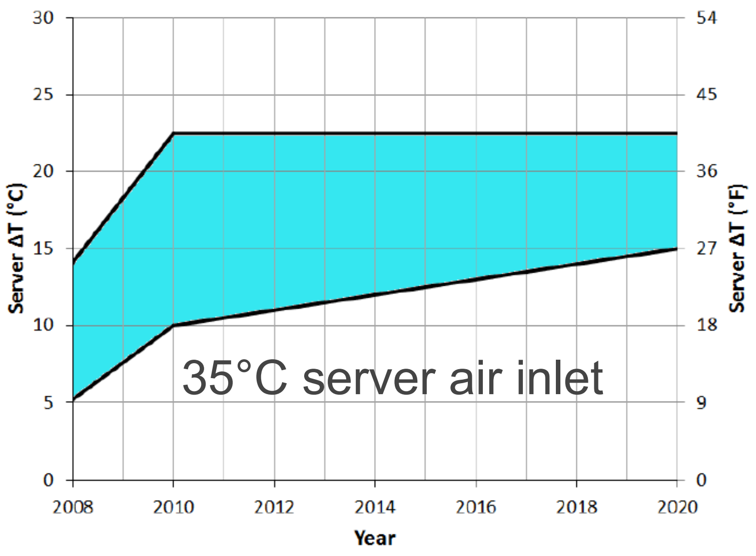

DELLEMC

DLC Components

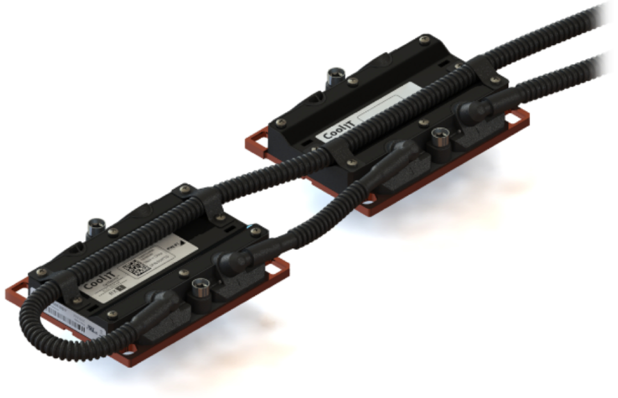

- DLC or Direct Liquid Cooling, also known as Cold Plate, uses a hollow metallic connection to the parts of a server, such as CPU or memory. Through this hollow metallic link, liquid flows taking the heat away within the server chassis. Coldplate assemblies replace heatsinks and are purpose-designed
- Connection to these heatsinks is via dry-break couplings, exiting the rear of the server, thereby permitting connection & disconnection with no leaks. A manifold within the rack manages & merges the water connections. Rack manifolds are made with reliable stainless steel and typically occupy a PDU mounting location at the rear of the rack.

- Heat Exchanger Modules connect to the rack manifold to pump coolant to the racks and exchange heat from the servers with facility water.
- Solution sold and supported by Dell on C6420
- Heat removal of **up to 80 kilowatts per rack (!)**.
- Typical **Inlet water temperature 27-45°C**

DLC: Pros

- Wider Cooling Margin allows higher water temperatures, which means less chiller infrastructure.
- Also allows heat extraction & reuse from the outlet water (i.e. through building heating system).

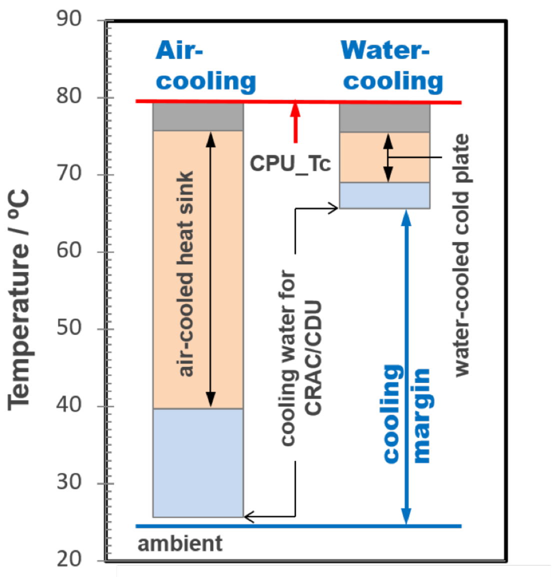

- High Temperature water allows for more hours of free cooling through low-maintenance dry coolers
- EER increasing with Inlet water temperature increase
- Enables lower TCO costs with more efficient cooling

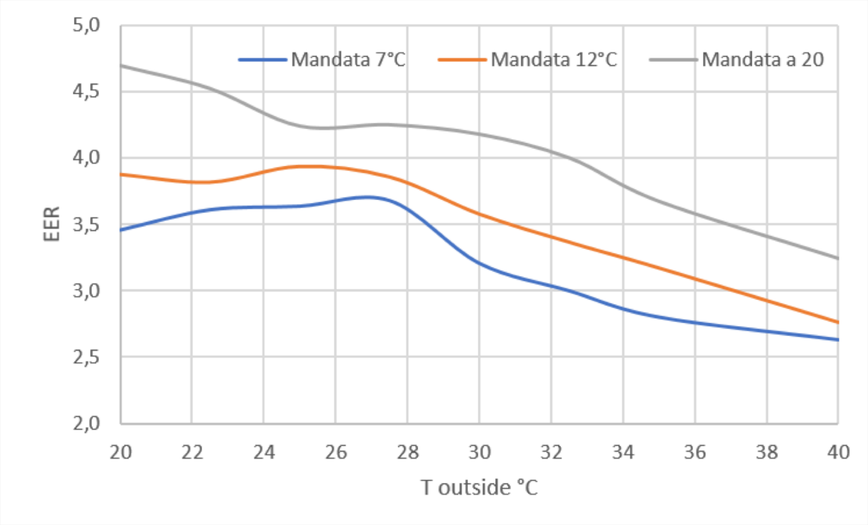

Source : R.Ricci – Infrastruttura: come cambia il clima in sala calcolo – INFN CCR Workshop 2018

DELL EMC

DLC: Cons

- Only a portion (up to 80%) of the total generated heat is removed with liquid, fans still required.
- A significant portion (about 20%) of the total DLC system investment is lost when servers are replaced (coldplate assemblies and tubing are purpose-designed)
- No warranty that different server vendors will adopt the same DLC system in the future (quite all enterprise servers manufacturer are however converging toward Staübli push-fit connectors, at today)
- High water temperatures (usually) require specific cooling circuits
- Proper Water Quality required (ASHRAE-D-90564)
- Using a single cooling system for traditional cooling and DLC may not make sense, as the cooling fluids would usually be at different temperatures.
  - But output water from Rear Door Heat Exchangers could be reused as inlet water for DLC (if DC cooling system can support this)

Parameter	Recommended Limits
pH	7 to 9
Corrosion inhibitor	Required
Sulfides	<10 ppm
Sulfate	<100 ppm
Chloride	<50 ppm
Bacteria	<1000 CFU/mL
Total hardness (as $CaCO_3$)	<200 ppm
Residue after evaporation	<500 ppm
Turbidity	<20 NTU (nephelometric)

Example: TACC Frontera Room Neutral System

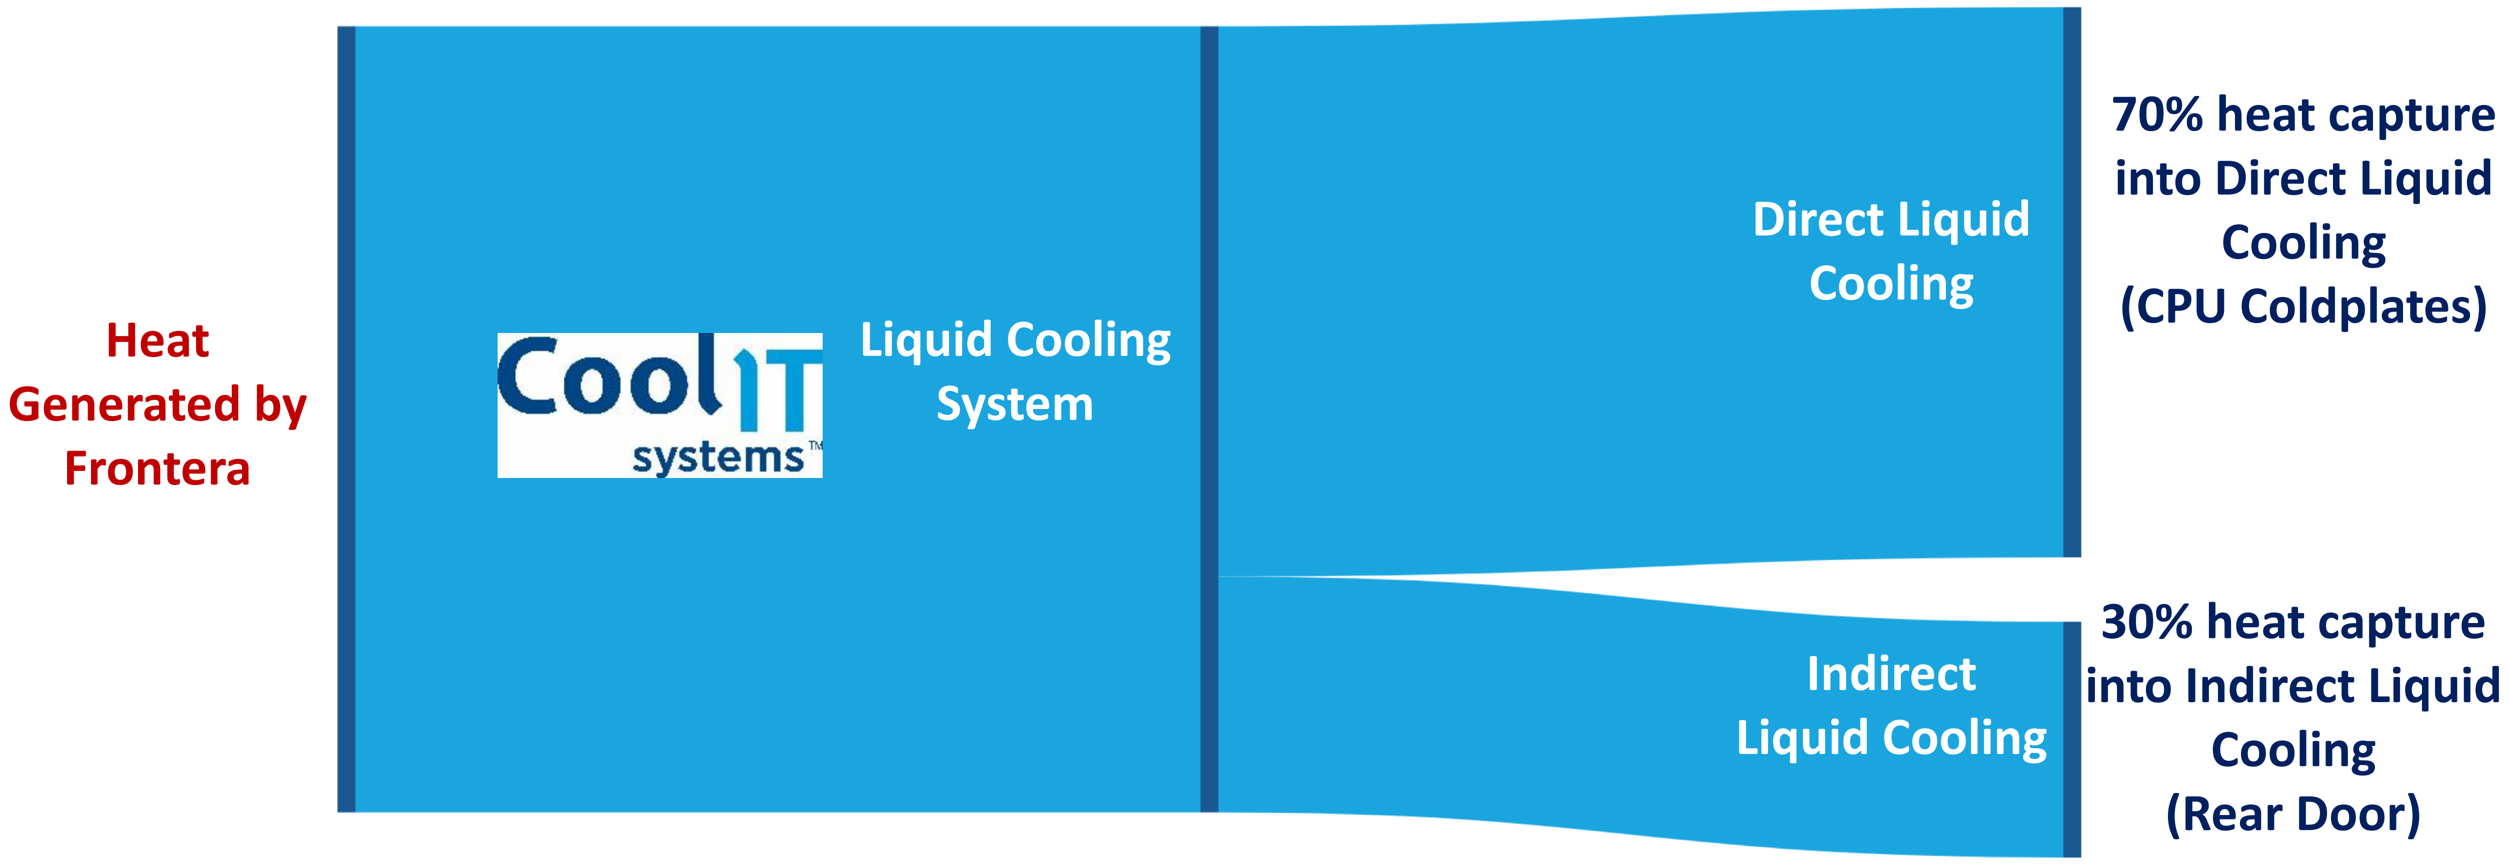

Immersion Cooling

- Immersion cooling is the practice of submerging computer components or systems in a thermally, but not electrically, conductive liquid (dielectric coolant) such as mineral oil, fluorocarbon-based fluids, synthetic fluids.
- In traditional cooling systems, heat must transfer through several levels of thermal interface materials, air, heat exchangers and working fluids. Immersion cooling simplifies thermal design and increases heat transfer efficiency.
- Depending on the properties of the coolant fluids, we can classify the Immersion Cooling techniques in:
  - Single-Phase Immersion Cooling
  - Two-Phase Immersion Cooling

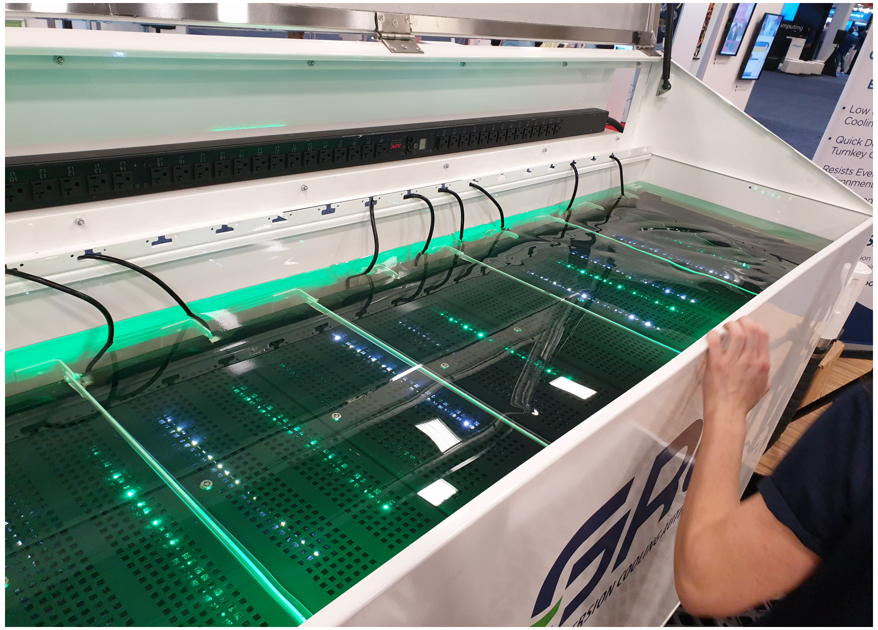

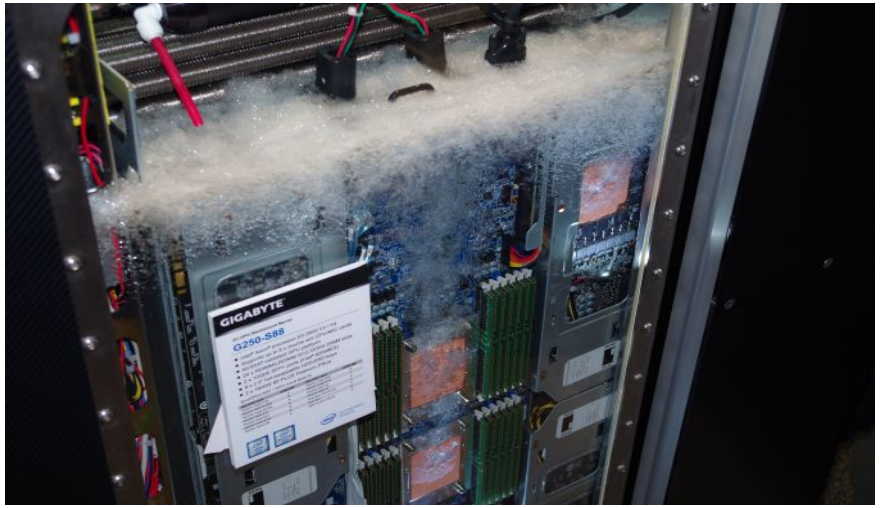

Single-Phase

- In single-phase immersion cooling the fluid has a higher boiling point and remains in its liquid phase throughout the process.
- Electronic components are submerged in a non-conductive liquid bath in an accessible, sealed enclosure
- The heat from the chip is transferred to the fluid
- Pumps are often used to flow the heated fluid to a heat exchanger, where it is cooled and cycled back into the bath
- Heat removal of up to **hundreds of KW per unit**
- Environment-Friendly Fluids (even cooking oil concepts! ☺)

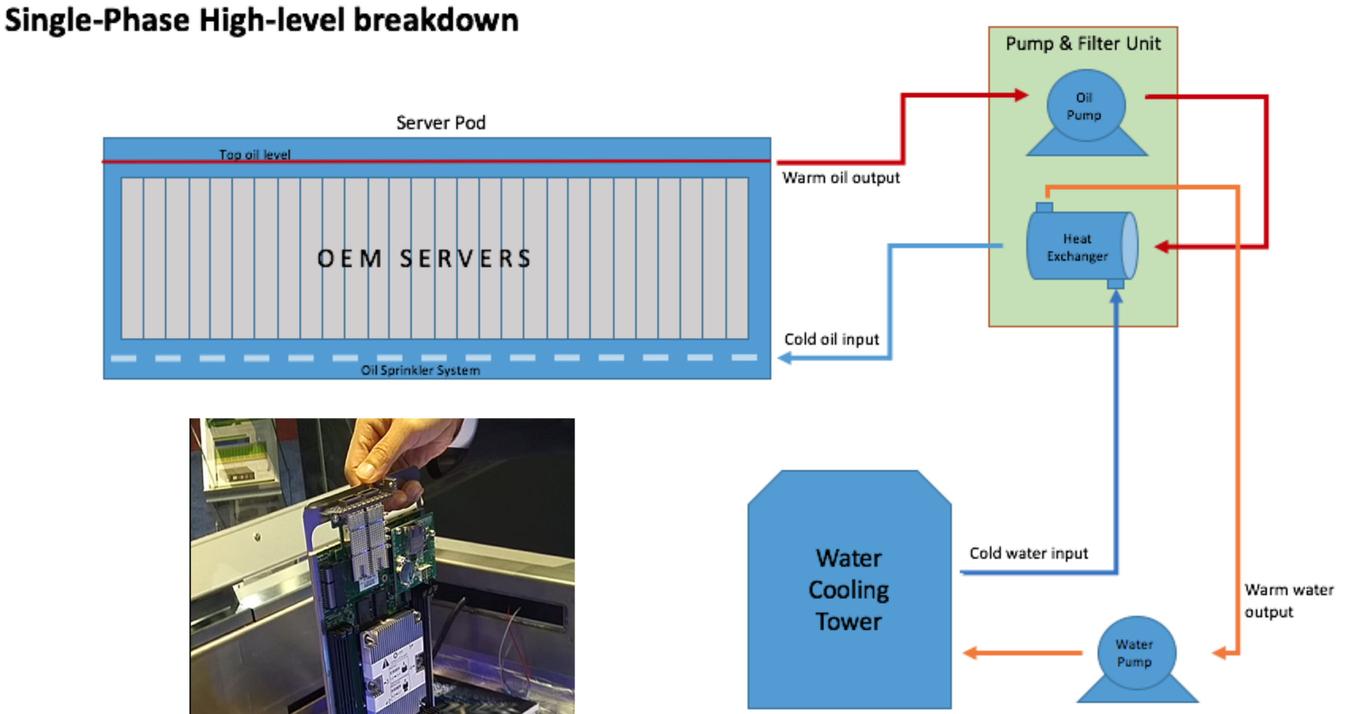

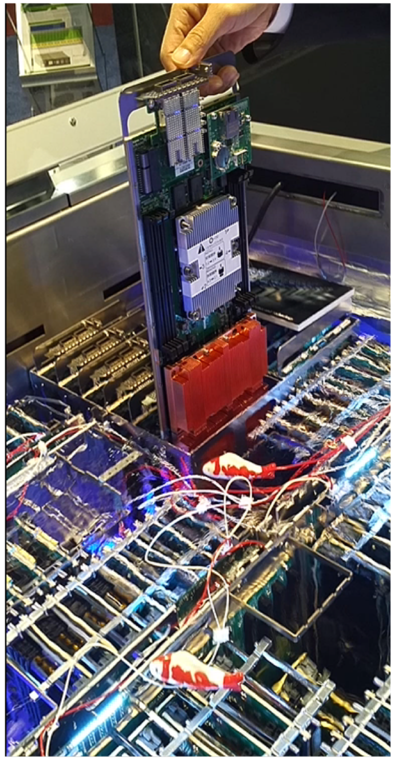

Two-Phase

- Two-phase passive immersion cooling exponentially increases heat transfer efficiency from boiling and condensation of cooling fluids.
- Electronic components are submerged in a non-conductive liquid bath in an accessible, sealed enclosure
- The generated heat boils fluid into a gaseous state and it rises to the top of the enclosure
- The gas condenses on the lid or a condensation coil and rains back into the bath to repeat the cycle
- Heat transfer efficiencies can exceed boiling water in a material that can safely come in direct contact with electronics (High latent heat in liquid vaporization)
- Heat removal of up to **hundreds of KW per unit**

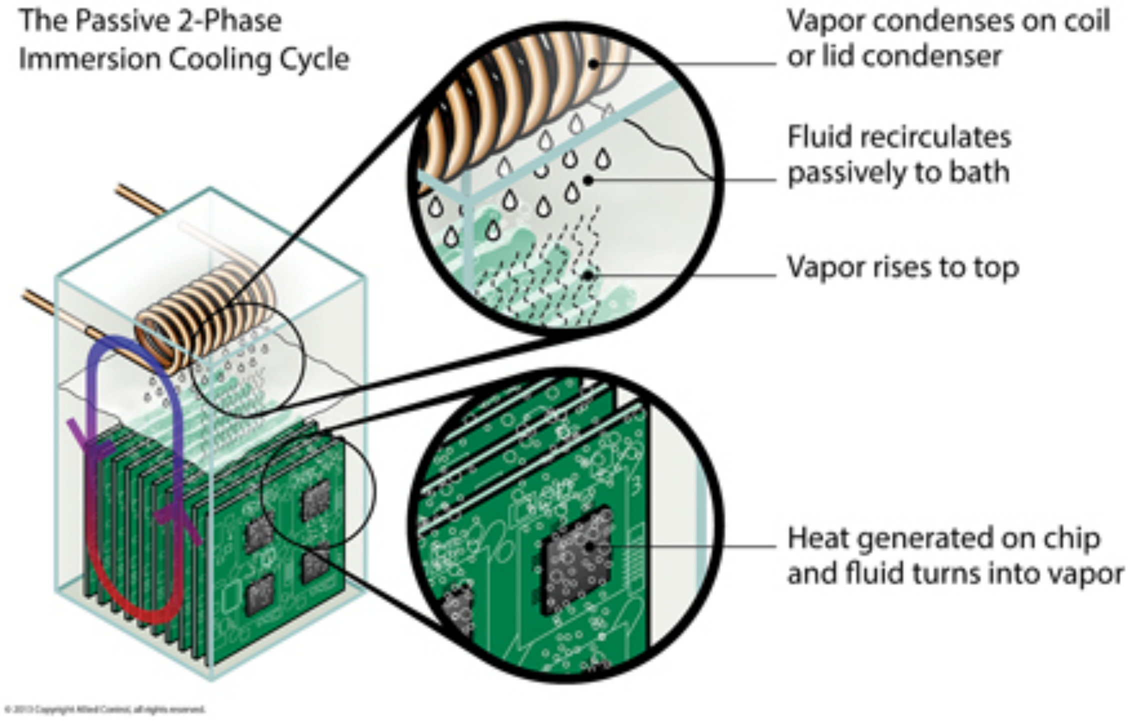

Immersion Cooling: Challenges

- Material Compatibility
  - Several studies demonstrated coolant fluid pollution due to dissolved chemicals coming from cables and HW/chips/plastic containers
- Signal Integrity
  - PCB strips designed for copper-in-the-air signal transmission (dielectric fluid affecting EM properties)
- Reliability
  - Micro-cavitation in Two-Phase; Long-Term effects on electronics not well known
- Optics and Optical devices
  - Refraction Index of coolant fluid very different from air!!
- IT Equipment Serviceability
- Environmental Footprint (Global Warming Potential)
- Thermal Resistance
  - DLC may be far better of immersion in cooling liquids! (We've found 0,05 °C/W vs 0,2 °C/W!)

Top 500: Cooling Methods

Cooling solutions for supercomputers

- Air cooling
- Water cooling
- One-phase immersion cooling
- Two-phase immersion cooling

Cooling solutions applied in Top50 Supercomputing systems [1]

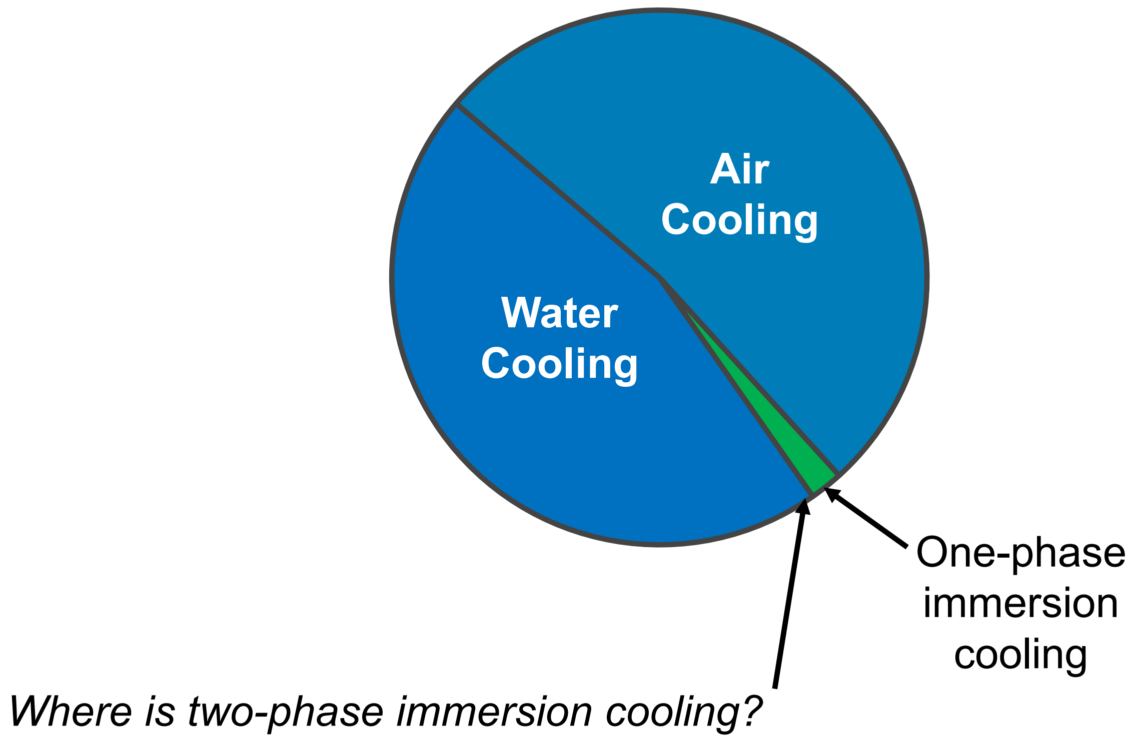

[1] "Top 500 supercomputers," Top 500 list as of November 2017.
www.top500.org

Conclusions

- Close Coupled Cooling is still the best solution, if Racks TDP falls into specifications
  - Full Investment recovery at IT HW replacement
  - Very long service life (independent from IT HW)
  - Very Low maintenance costs
  - Efficiency quite similar to Hybrid DLC solutions
  - TCO balanced by minor CAPEX investment at start
  - Rear Door Exchangers may remove the same heat qty as DLC
- DLC to be used only when Close Coupled not sufficient, not applicable, or in new systems
  - About 20% of the DLC solution investment is lost at IT HW replacement
  - Standard air cooling solution still required
  - Difficult implementation within existing cooling solutions (cooling fluids at different temperatures)
- Immersion cooling still a niche solution and, at today, practically unapplicable for standard DC IT
  - Higher PUE and kW/rack values, however:
    - Supportability, Environmental and Technological concerns
    - Needs specific support infrastructure (because of weight, exchangers etc)
    - Exciting PUE, however TCO reduction unclear
    - TCO influenced by replacement fluid costs.

Thank you