U.S. DEPARTMENT OF ENERGY | Energy Efficiency & Renewable Energy

New Horizons:
Getting (Back) to Liquid Cooling

7×24 Exchange International Fall Conference

Phoenix, AZ

October 29, 2019

Dale Sartor, P.E.
Lawrence Berkeley National Laboratory

Otto Van Geet, P.E.
National Renewable Energy Laboratory

Dave Martinez
Sandia National Laboratory

Version: 9/19/19

U.S. DEPARTMENT OF ENERGY
Office of ENERGY EFFICIENCY & RENEWABLE ENERGY

Liquid in the Rack: Liquid Cooling Your Data Centers

Dale Sartor, PE, Lawrence Berkeley National Laboratory

FEMP
Federal Energy Management Program

Benefits of Liquid Cooling

- **Higher compute densities**
- **Higher efficiency**

  ➢ Vision: Eliminate compressor based cooling and water consumption

Moving (Back) to Liquid Cooling

- As heat densities rise, liquid solutions become more attractive
- Volumetric heat capacity comparison:

(1.5 m^3)

Water

=

(5,380 m^3)

Air

Why Liquid Cooling?

- **Liquids can provide cooling at higher temperatures**
  – Improved cooling efficiency
  – Increased economizer hours
  – Potential use of waste heat
- **Reduced transport energy:**

Heat Transfer		Resultant Energy Requirements			
Rate	ΔT	Heat Transfer Medium	Fluid Flow Rate	Conduit Size	Theoretical Horsepower
10 Tons	12°F	Forced Air	9217 cfm	34" Ø	3.63 Hp
		Water	20 gpm	2" Ø	.25 Hp

Liquid Cooling Solution

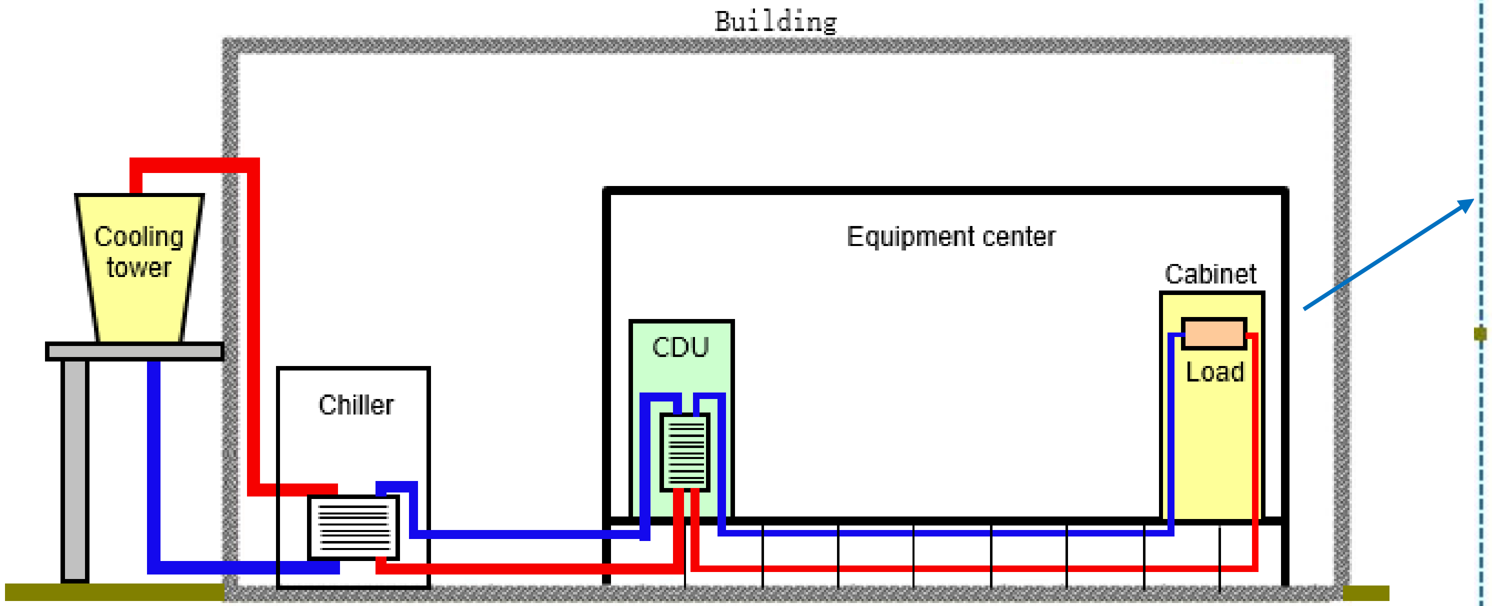

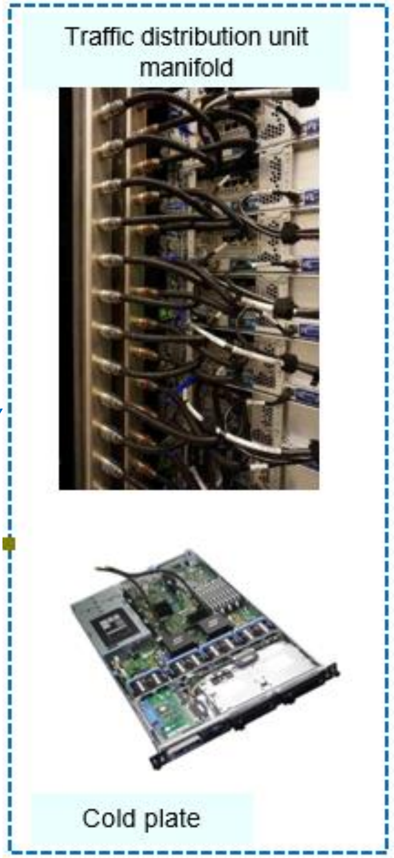

Typical liquid cooled equipment room, with external coolant distribution units (CDUs)

- For most locations data centers may be operated without chillers with a water-side economizer

In-Rack Liquid Cooling

- Racks with integral coils and full containment:

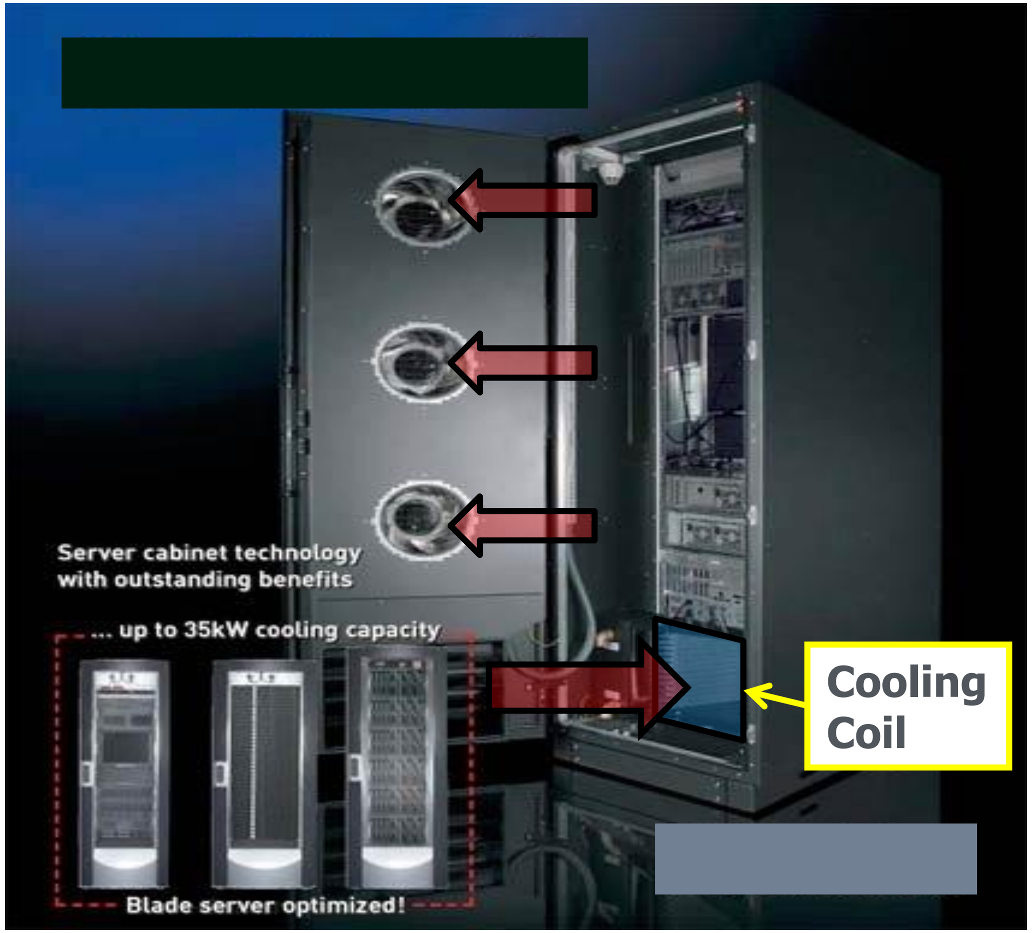

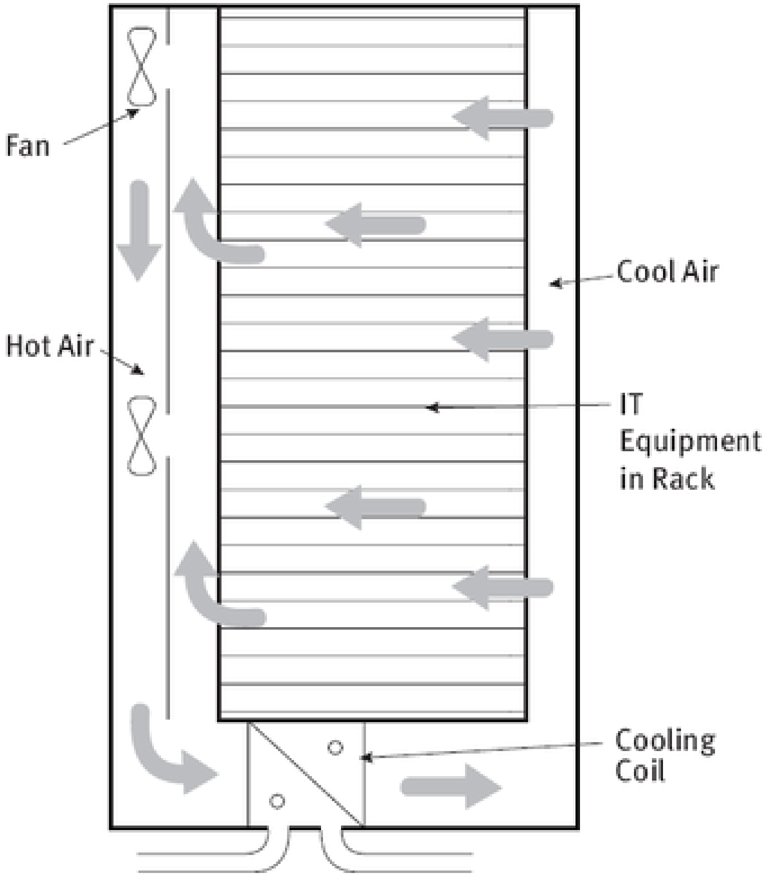

Rear-Door Heat Exchanger

- Passive technology:
  relies on server fans for airflow
- Active technology:
  supplements server fans with external fans in door
- Can use chilled or higher temperature water for cooling

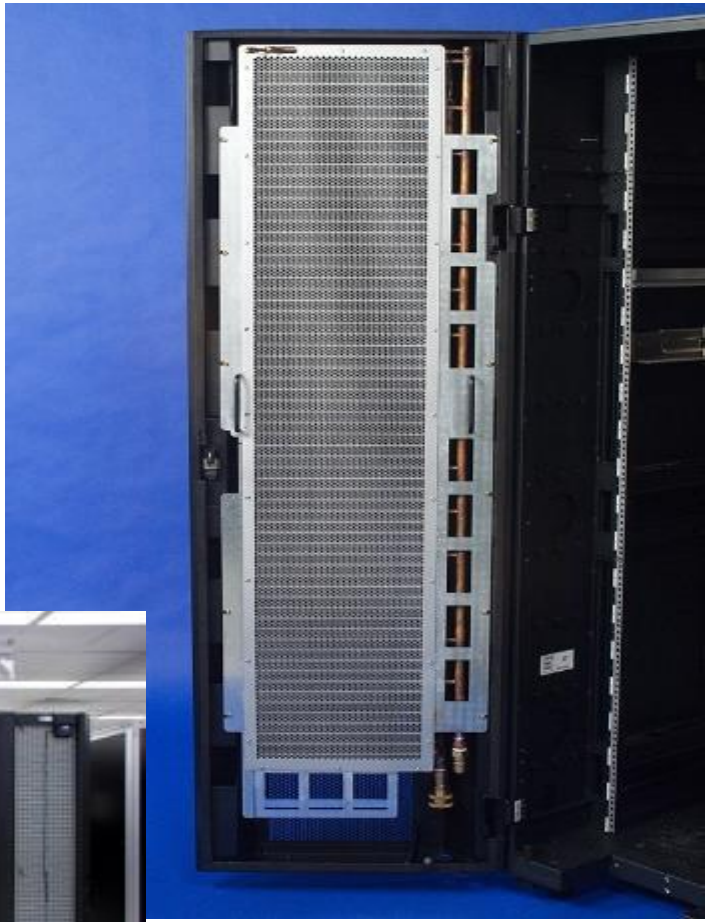

Photo courtesy of Vette

Liquid On-Board Cooling

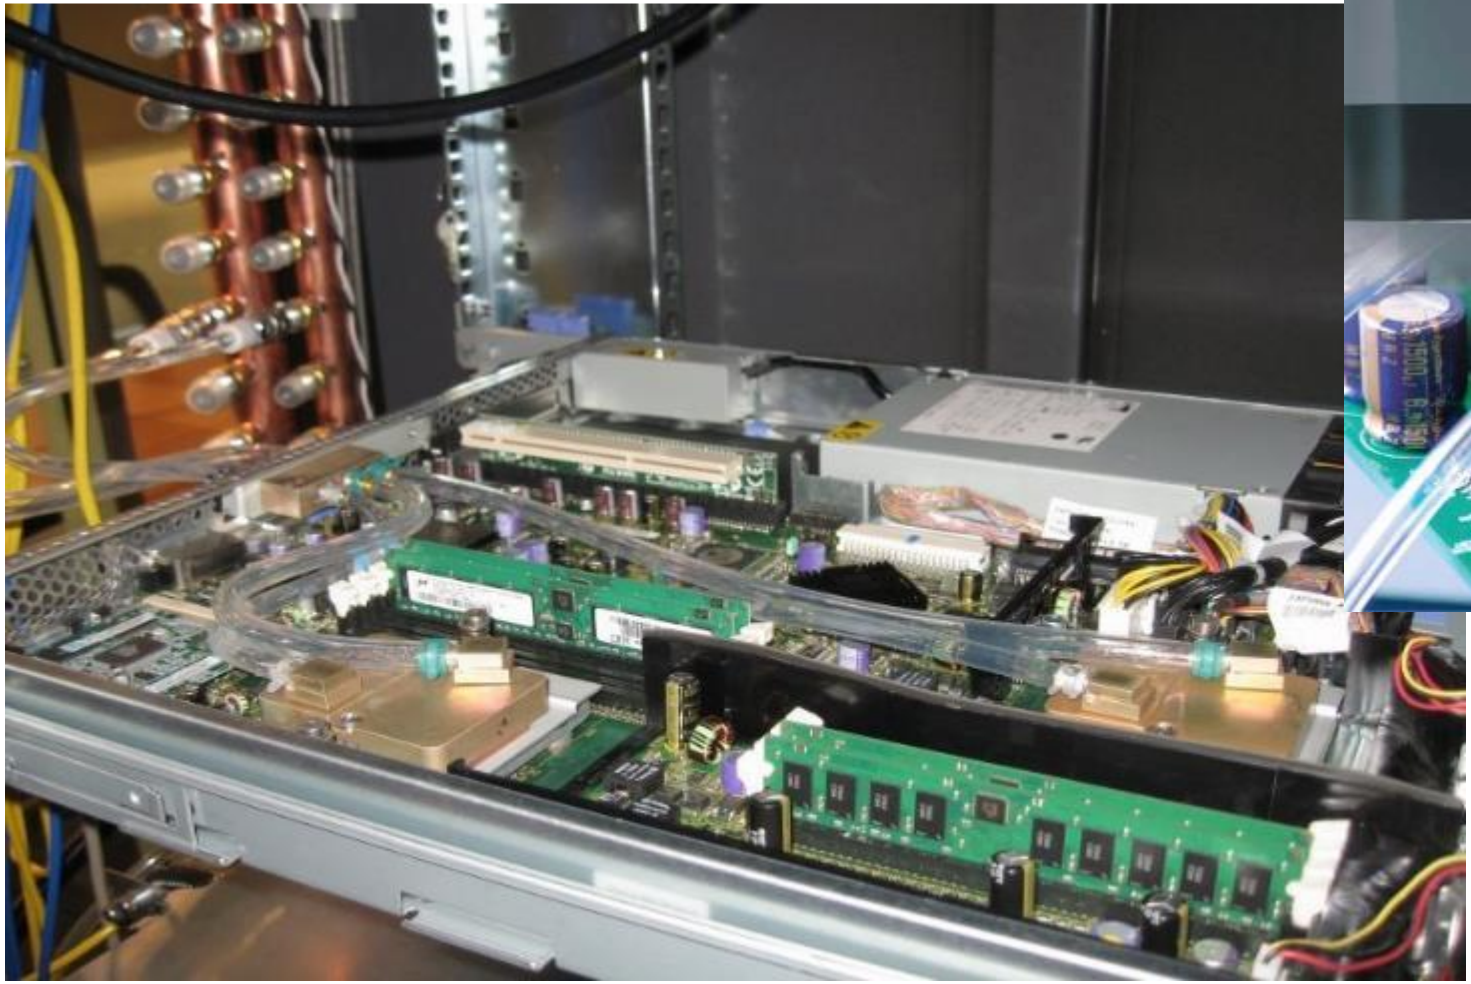

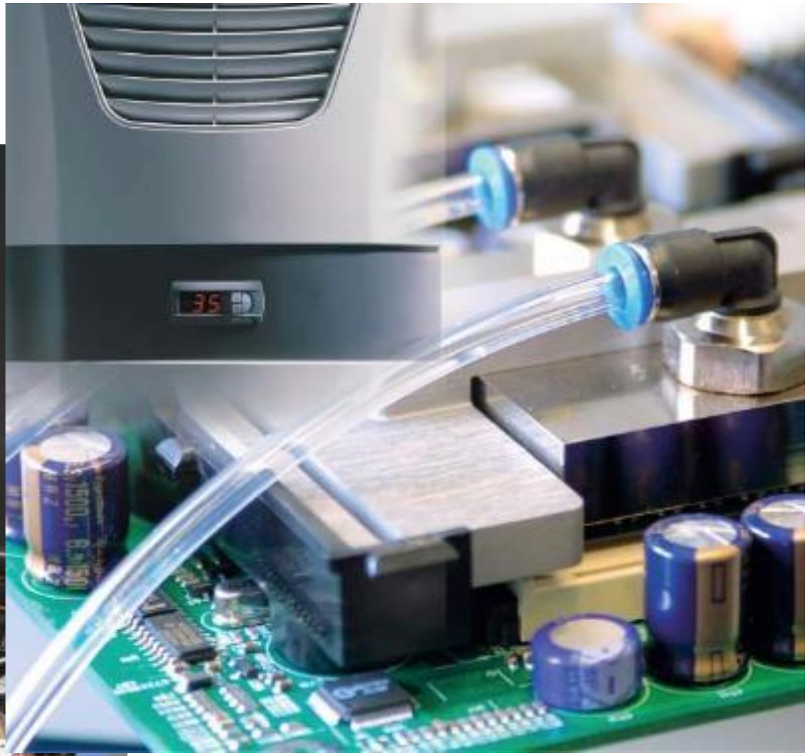

Example: Maui DOD HPC Center Warm Water Cooling

IBM System x iDataPlex

- 90% water cooled
- 10% air cooled
- Cooling water temperature as high as 44°C

Water inside

Dry Coolers, 10 kW each compared to 100 kW Chillers

Liquid Immersion Cooling

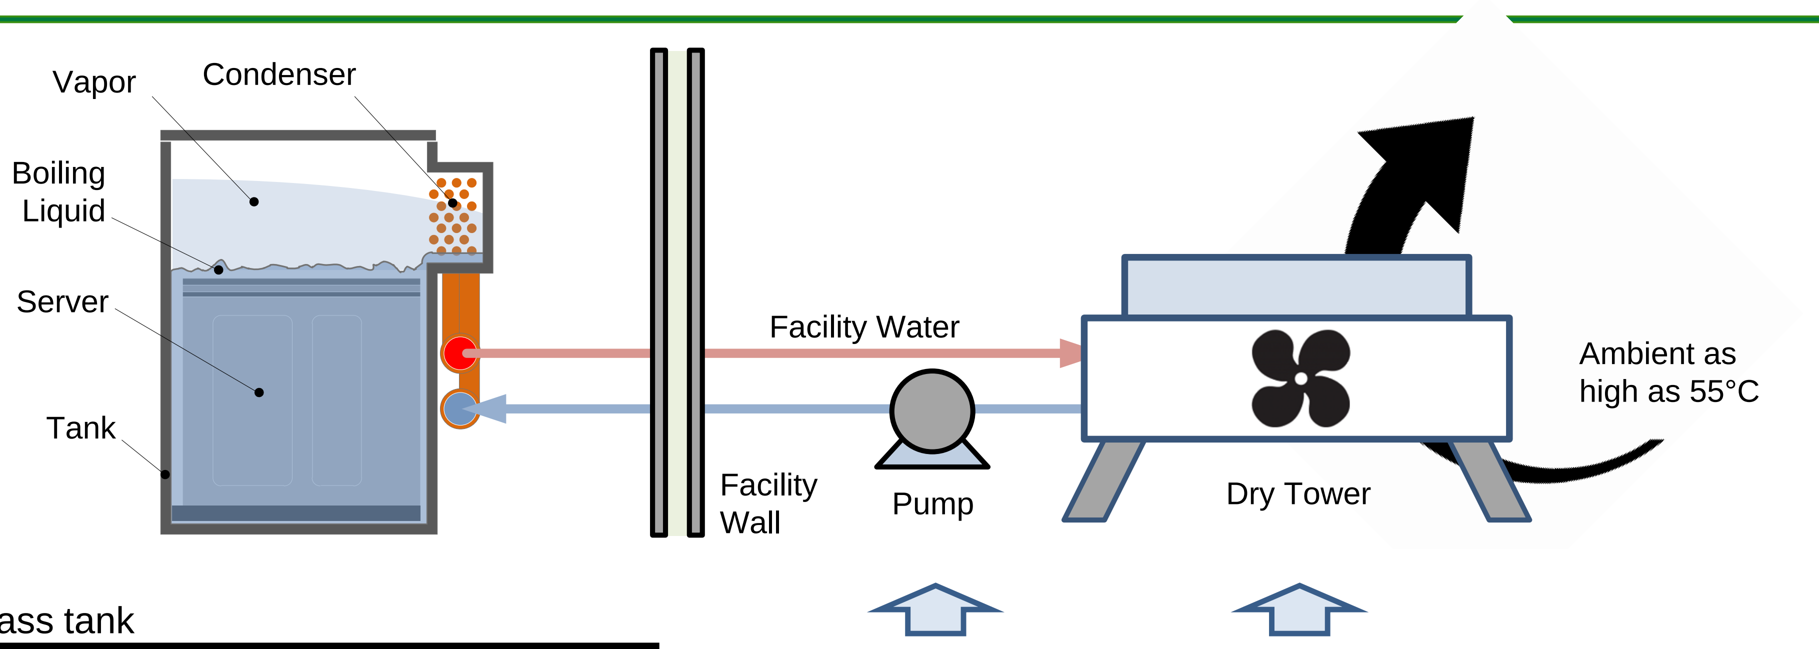

Cooling Power = Pump + Fan

Computers in glass tank

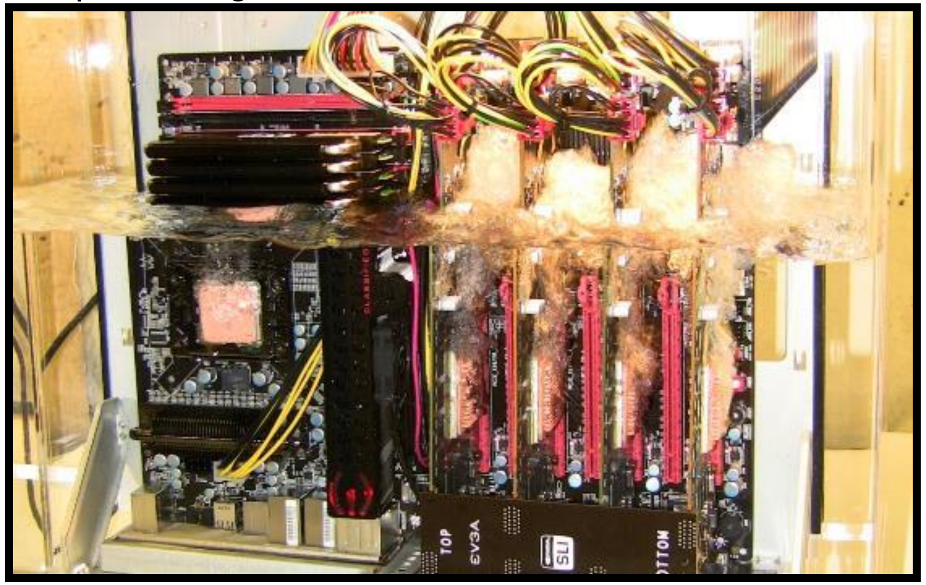

No longer requires:

- chillers
- cooling towers
- water use
- raised floors
- CRACs
- earplugs!

"Chill-Off 2" Evaluation of Liquid Cooling Solutions

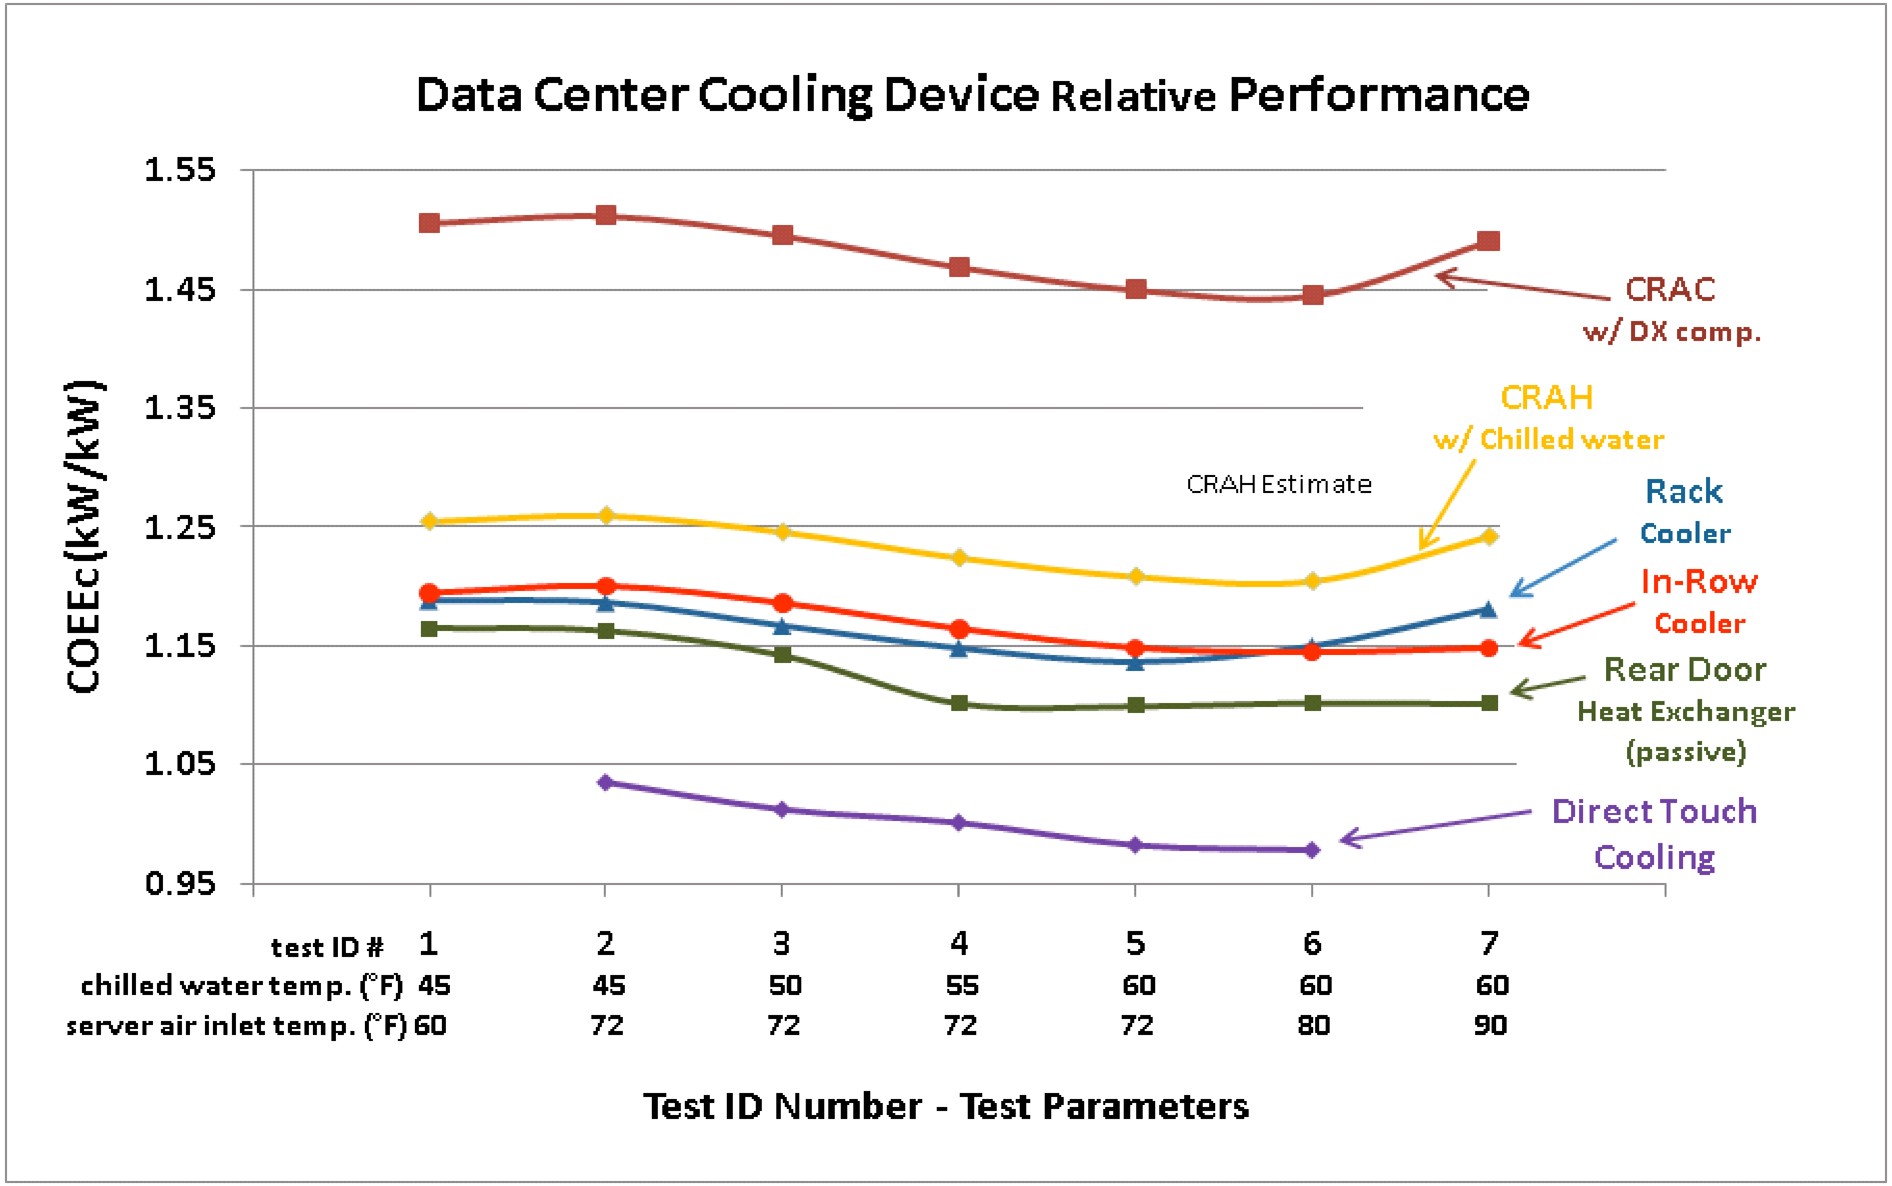

“Free Cooling” w/ Water-Side Economizers

- Cooling without Compressors
- Easier retrofit
- Added reliability (backup in case of chiller failure)
- No contamination issues
- Put in series with chiller
- Uses tower or dry cooler

No or minimum compress or cooling

Cooling tower and HX = Water-side Economizer

Re-Use of Waste Heat

- **Heat from a data center can be used for:**
  - Heating adjacent offices directly
  - Preheating make-up air (e.g., “run around coil” for adjacent laboratories)
- **Use a heat pump to elevate temperature**
  - Waste heat from LBNL ALS servers captured with rear door coolers feeds a heat pump that provides hot water for reheat coils
- **Warm-water cooled computers are used to heat:**
  - Greenhouses, swimming pools, and district heating systems

Open Specifications for Liquid Cooled Rack

- **While liquid cooling potential is understood, uptake is slow**
- **Most solutions are unique and proprietary**
- **Needed:**
  - Multi-source solutions
  - Reusable rack infrastructure
- **Users can drive faster technology development and adoption**

Open Compute Project (OCP) and EEHPCWG Initiative(s)

- **Goal:**
  - An open specification for a liquid cooled rack that would encourage a multiple vendor solution, and provide a lasting infrastructure for multiple refresh cycles with a variety of liquid cooled server suppliers.
- **The "cold plate" working groups are working on:**
  - Water based transfer fluid: quality, treatment and compatibility
  - Wetted material list (OK and not OK)
  - Universal (multi-vendor) quick connectors
  - Operating conditions (e.g. supply pressure, temperatures)
- **Challenges**
  - Proprietary nature of components, e.g. chemical compositions of transfer fluids and the quick connects

ASHRAE Design Reference Conditions - 2015

Liquid Cooling Classes	Typical Infrastructure Design		Facility Supply Water Temperature
	Main Heat Rejection Equipment	Supplemental Cooling Equipment	
W1	Chiller/Cooling Tower	Water-side Economizer (With Drycooler or Cooling Tower)	35.6°F to 62.6°F
W2			35.6°F to 80.6°F
W3	Cooling Tower	Chiller	35.6°F to 89.6°F
W4	Water-side Economizer (With Drycooler or Cooling Tower)	N/A	35.6°F to 113°F
W5	Building Heating System	Cooling Tower	>113°F

Resources: Center of Expertise for EE in Data Centers

- datacenters.lbl.gov
- datacenters.lbl.gov/technologies/liquid-cooling

Use CoE's Energy Efficiency Toolkit

Filter CoE's many resources by type and topic.

Choose from upcoming live webinars, pre-recorded trainings, and in-person Data Center Energy Practitioner (DCEP) trainings.

Search resources by topics of interest.

Explore the diverse activities that CoE is engaged in.

Follow us on Twitter @DataCenterCoE

Contact Info

Dale Sartor, P.E.

Lawrence Berkeley National Laboratory

MS 90-3111

University of California

Berkeley, CA 94720

BERKELEY LAB

DASartor@LBL.gov

(510) 486-5988

http://datacenters.lbl.gov/

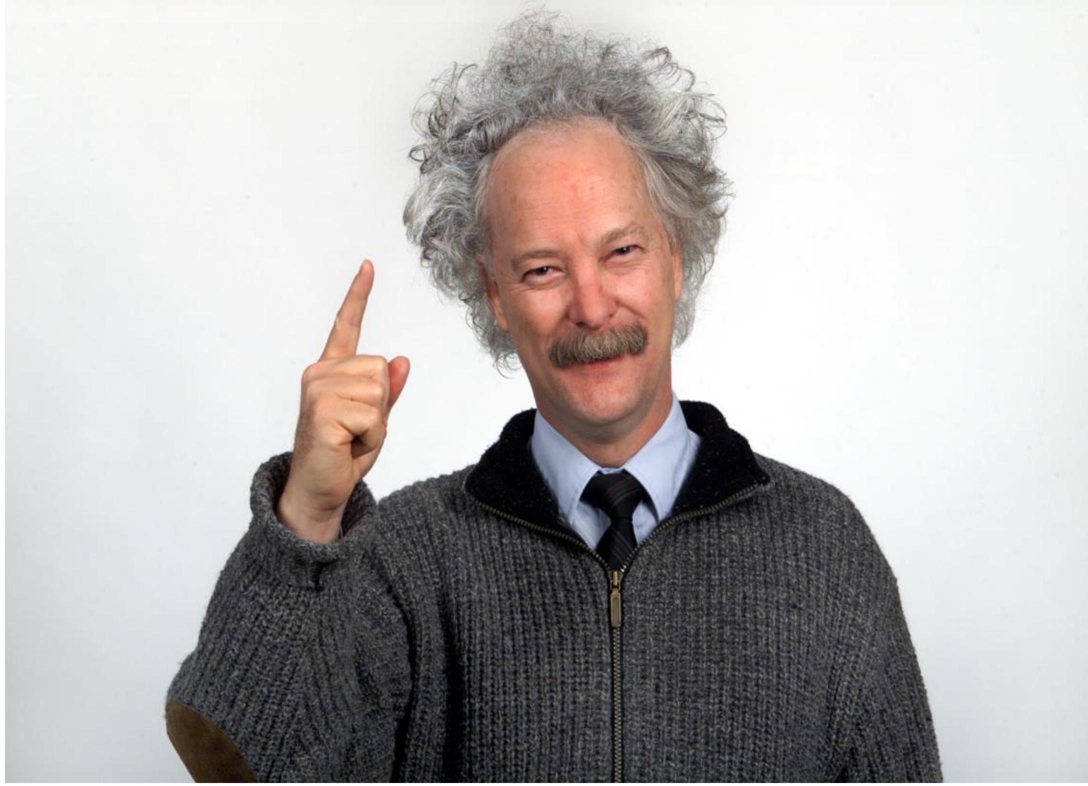

U.S. DEPARTMENT OF ENERGY
Office of ENERGY EFFICIENCY & RENEWABLE ENERGY

Data Center Optimization Strategies

Otto Van Geet, PE, NREL

October 2019

FEMP
Federal Energy Management Program

NREL Data Center Design

- **Showcase Facility**
  - ESIF 182,000 ft.^{2} research facility
  - 10,000 ft.^{2} data center
  - 10 MW at full buildout
  - LEED Platinum facility, PUE ≤ 1.06
  - No mechanical cooling (eliminates expensive and inefficient chillers)

- **Data Center Features**
  - Direct, component-level liquid cooling, 24°C (75°F) cooling water supply
  - 35°C–40°C (95°F–104°F) return water (waste heat), captured and used to heat offices and lab space
  - Pumps more efficient than fans
  - High voltage 480 VAC power distribution directly to high power density 60kW–80 kW compute racks
- **Compared to a Typical Data Center**
  - Lower CapEx—costs less to build
  - Lower OpEx—efficiencies save

Integrated "Chips to Bricks" Approach

Utilize the bytes and the BTUs!

Liquid Cooling – Considerations

- Liquid cooling essential at high-power density
- Compatible metals and water chemistry is crucial
- Redundancy in hydronic system (pumps, heat exchangers)
- Plan for hierarchy of systems
  - Cooling in series rather than parallel
  - Most sensitive systems get coolest liquid
- At least 95% of rack heat load captured directly to liquid

Air-Cooled to Liquid-Cooled Racks

Traditional **air-cooled** allow for rack power densities of 1kW–5kW

Liquid-cooled when rack power densities in 5kW–80kW range, have several options

Liquid-Cooled Server Options

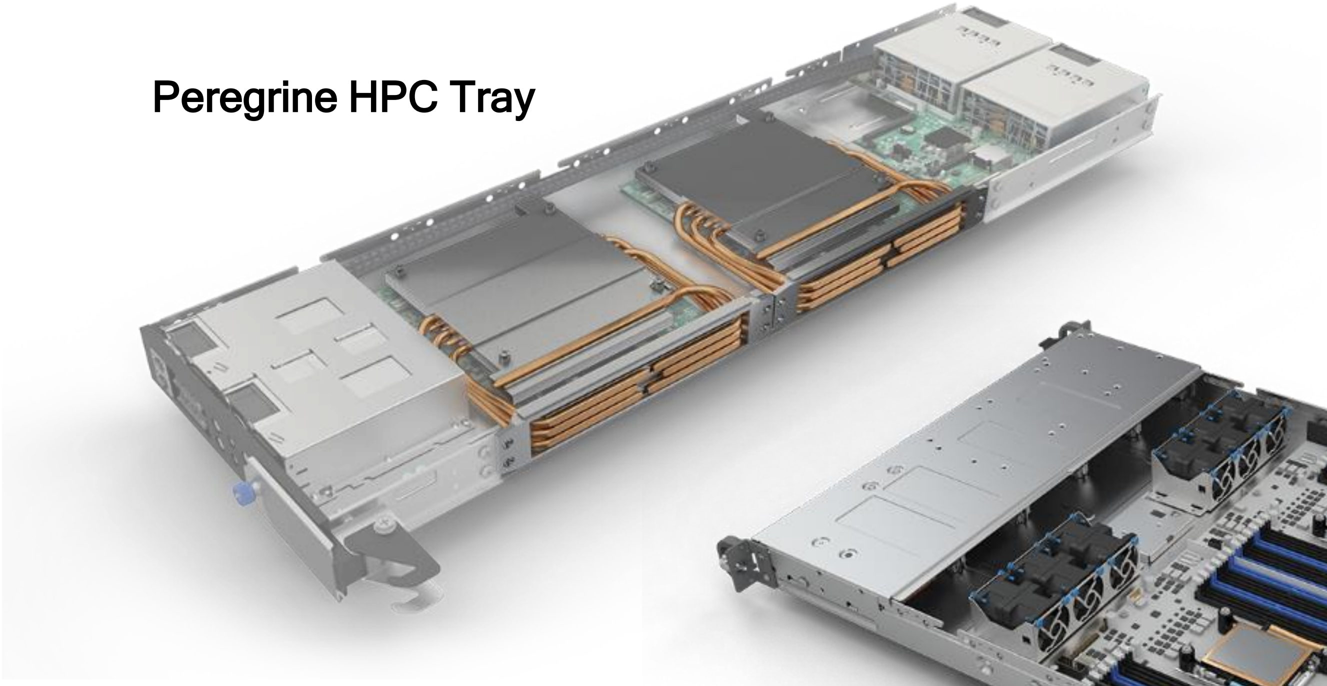

Peregrine HPC Tray

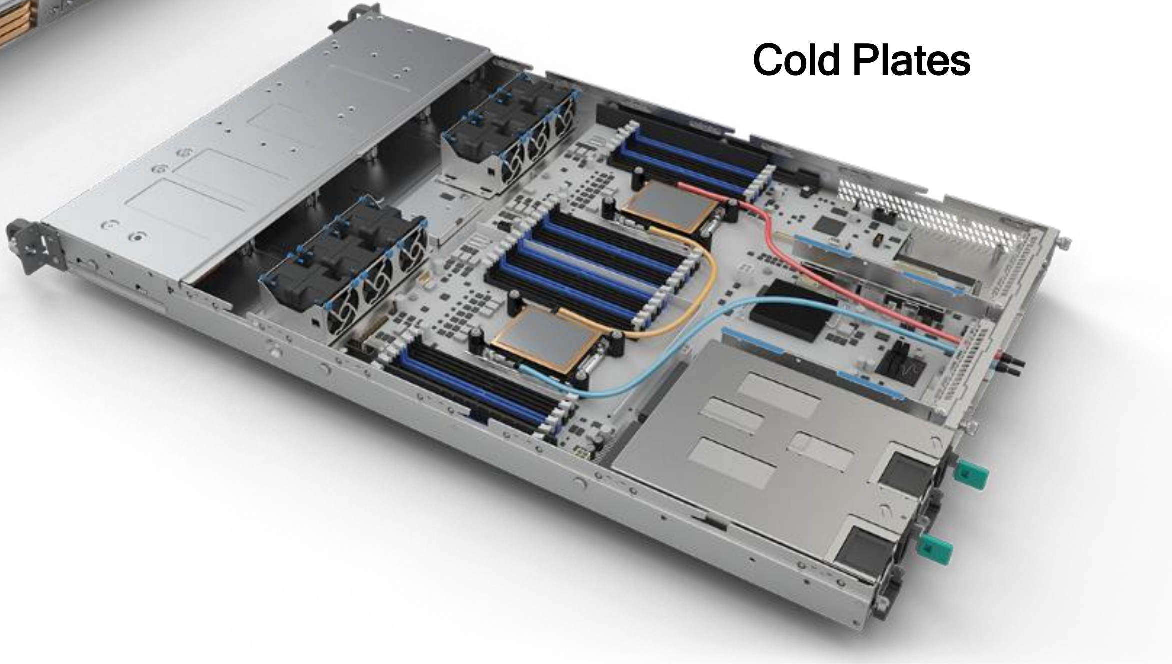

Cold Plates

Fanless Liquid-Cooled Server Options

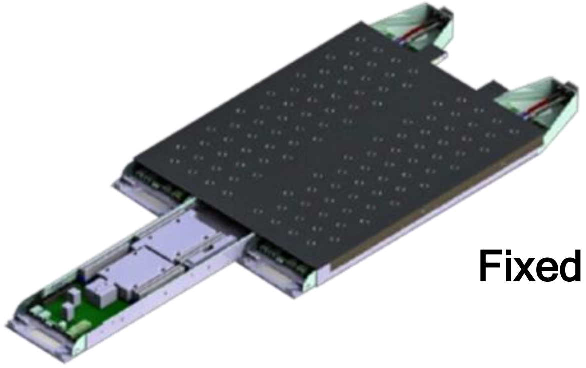

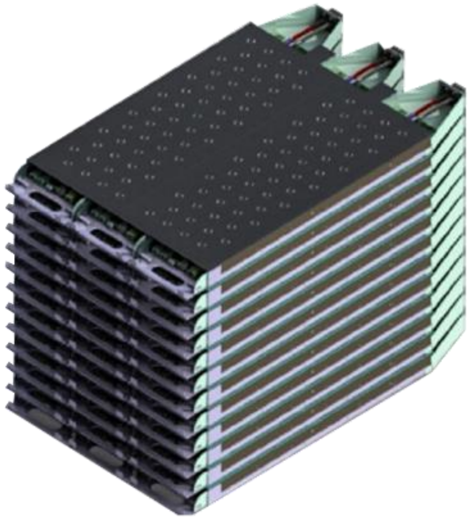

Fixed Plate

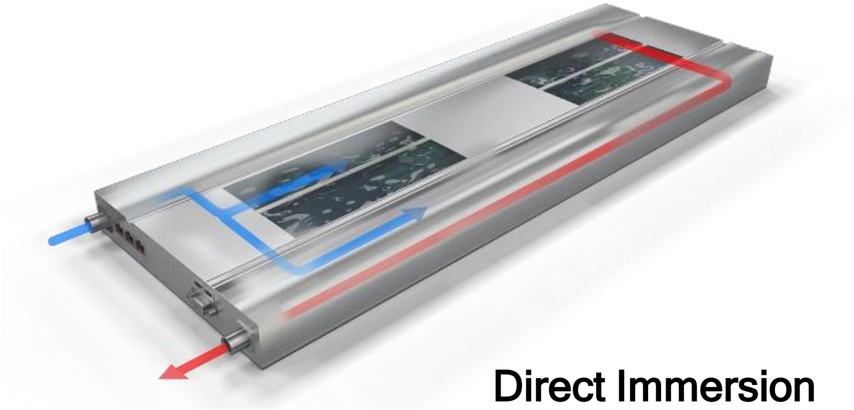

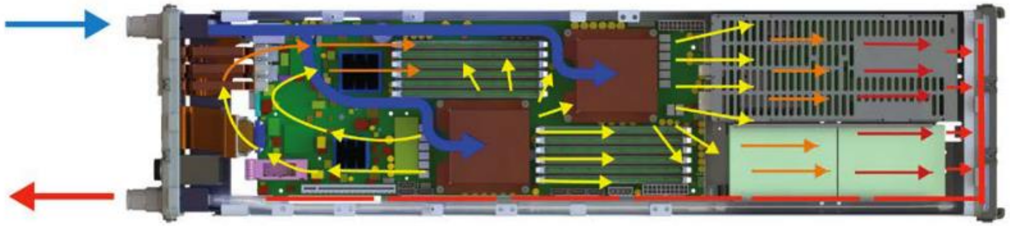

Direct Immersion

Metrics

$$PUE = \frac{\text{"Facility energy"} + \text{"IT energy"}}{\text{"IT energy"}}$$

$$ERE = \frac{\text{"Facility energy"} + \text{"IT energy"} - \text{"Reuse energy"}}{\text{"IT energy"}}$$

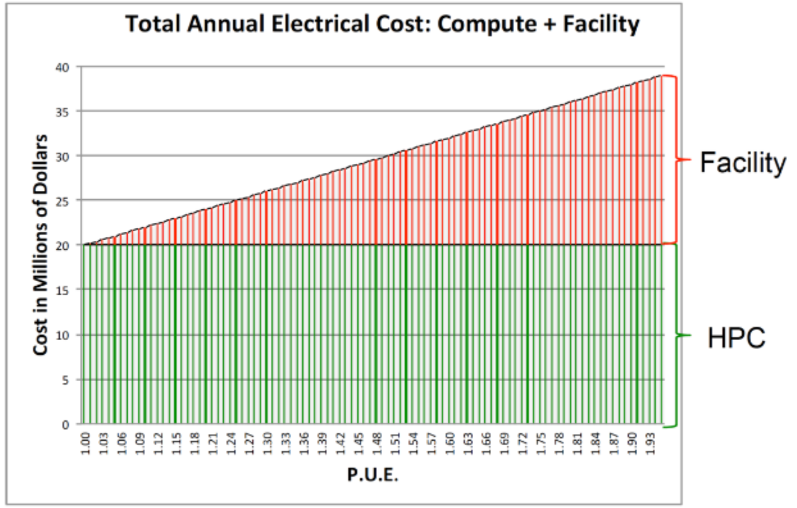

Assume ~20MW HPC system & $1M per MW year utility cost.

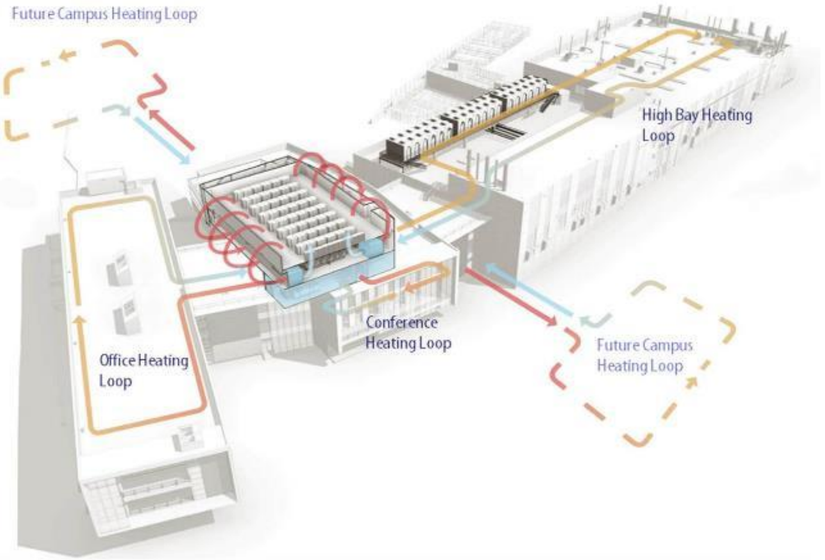

Metrics

$$WUE = \frac{\text{"Annual Site Water Usage"}}{\text{"IT energy"}}$$

the units of WUE are liters/kWh

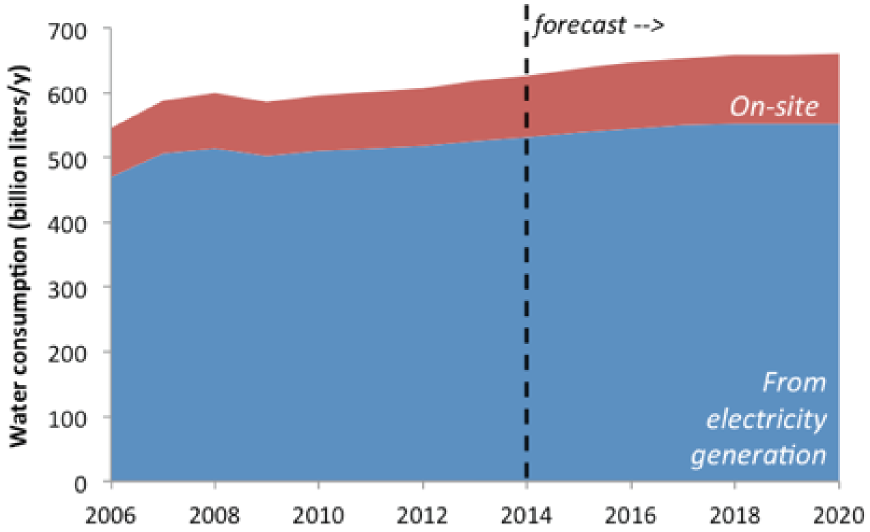

$$WUE_{SOURCE} = \frac{\text{"Annual Site Water Usage"} + \text{"Annual Source Energy Water Usage"}}{\text{"IT energy"}}$$

$$WUE_{SOURCE} = \frac{\text{"Annual Site Water Usage"}}{\text{"IT energy"}} + [EWIF \times PUE]$$

where EWIF is energy water intensity factor

Air- and Water-Cooled System Options

- **Air-Cooled System**
  - Operation is based on DRY BULB temperature
  - Consumes no water (no evaporative cooling)
  - Large footprint requires very large airflow rates
- **Water-Cooled System**
  - Operation is based on the lower WET BULB temperature
  - Evaporative cooling process uses water to improve cooling efficiency
    - 80% LESS AIRFLOW = lower fan energy
    - Lower cost and smaller footprint
  - Colder heat rejection temperatures improve system efficiency

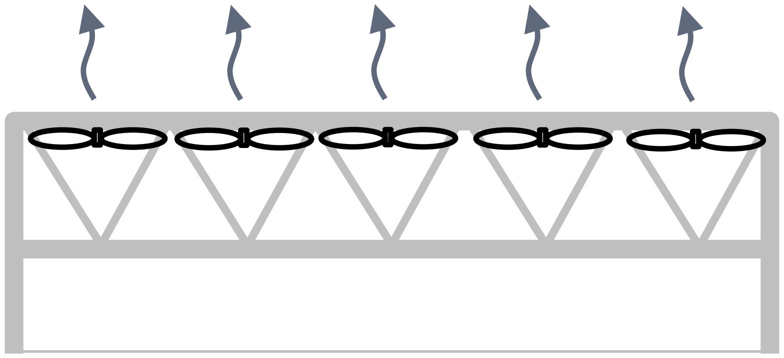

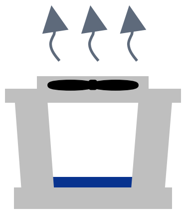

However, water-cooled systems depend on a reliable, continuous source of water.

Traditional Wet Cooling System

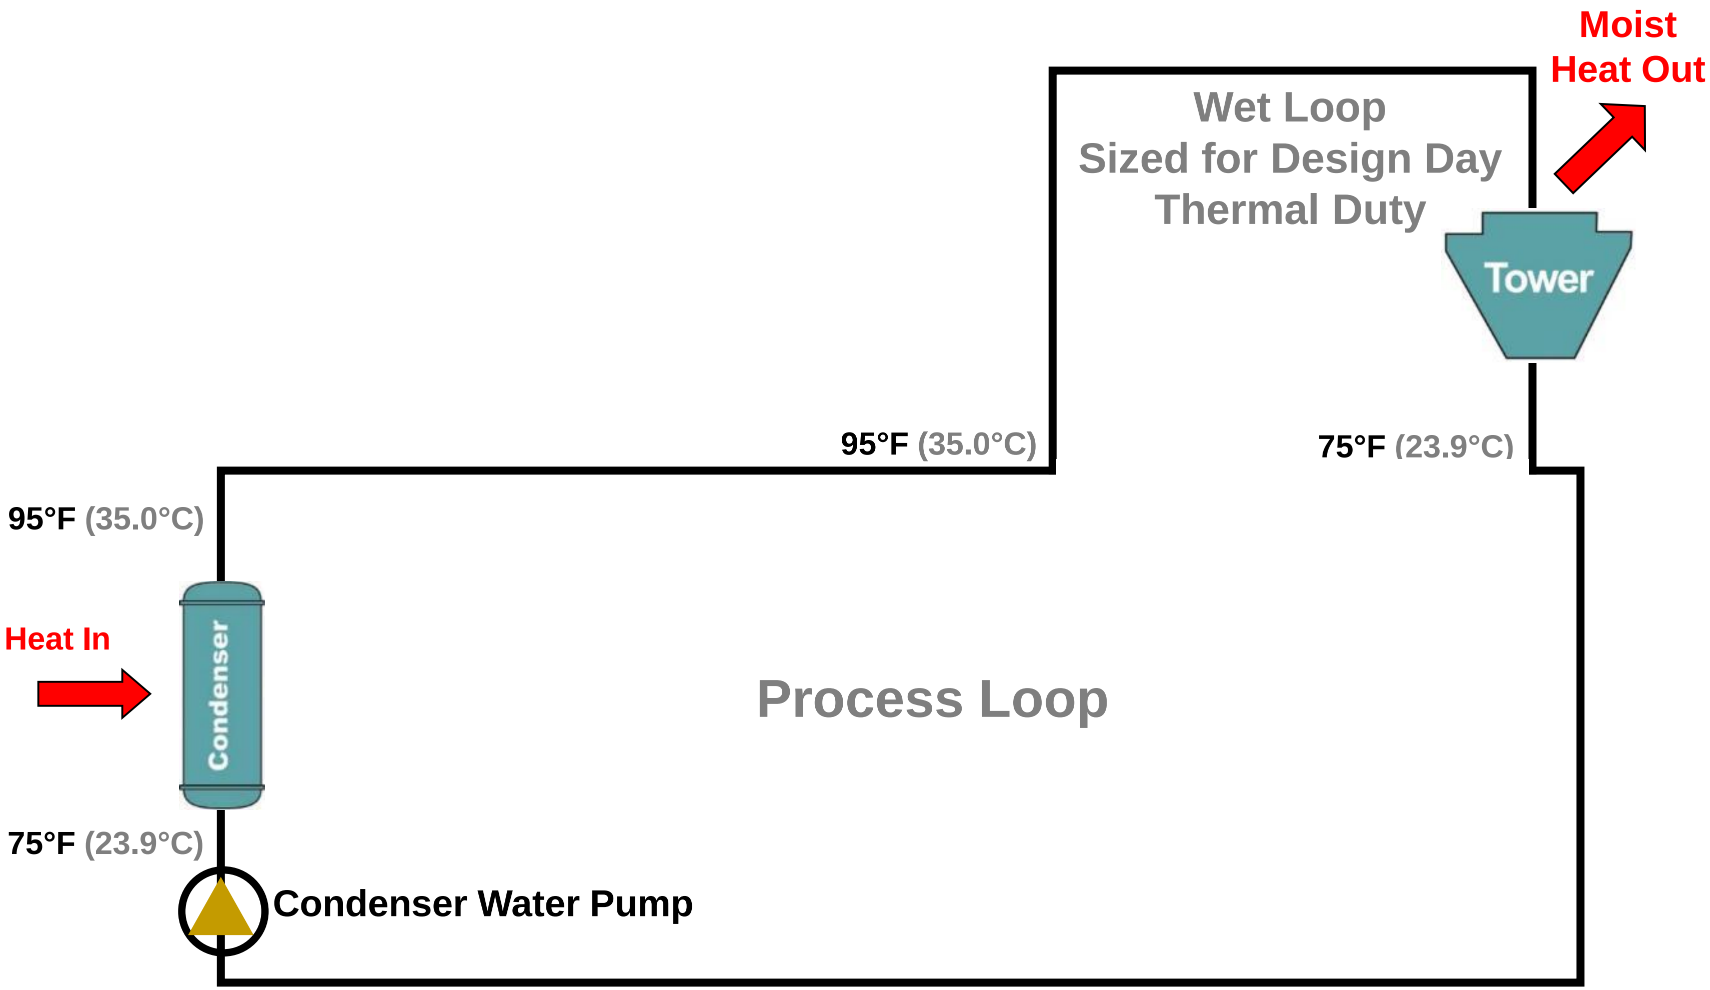

Basic Hybrid System Concept

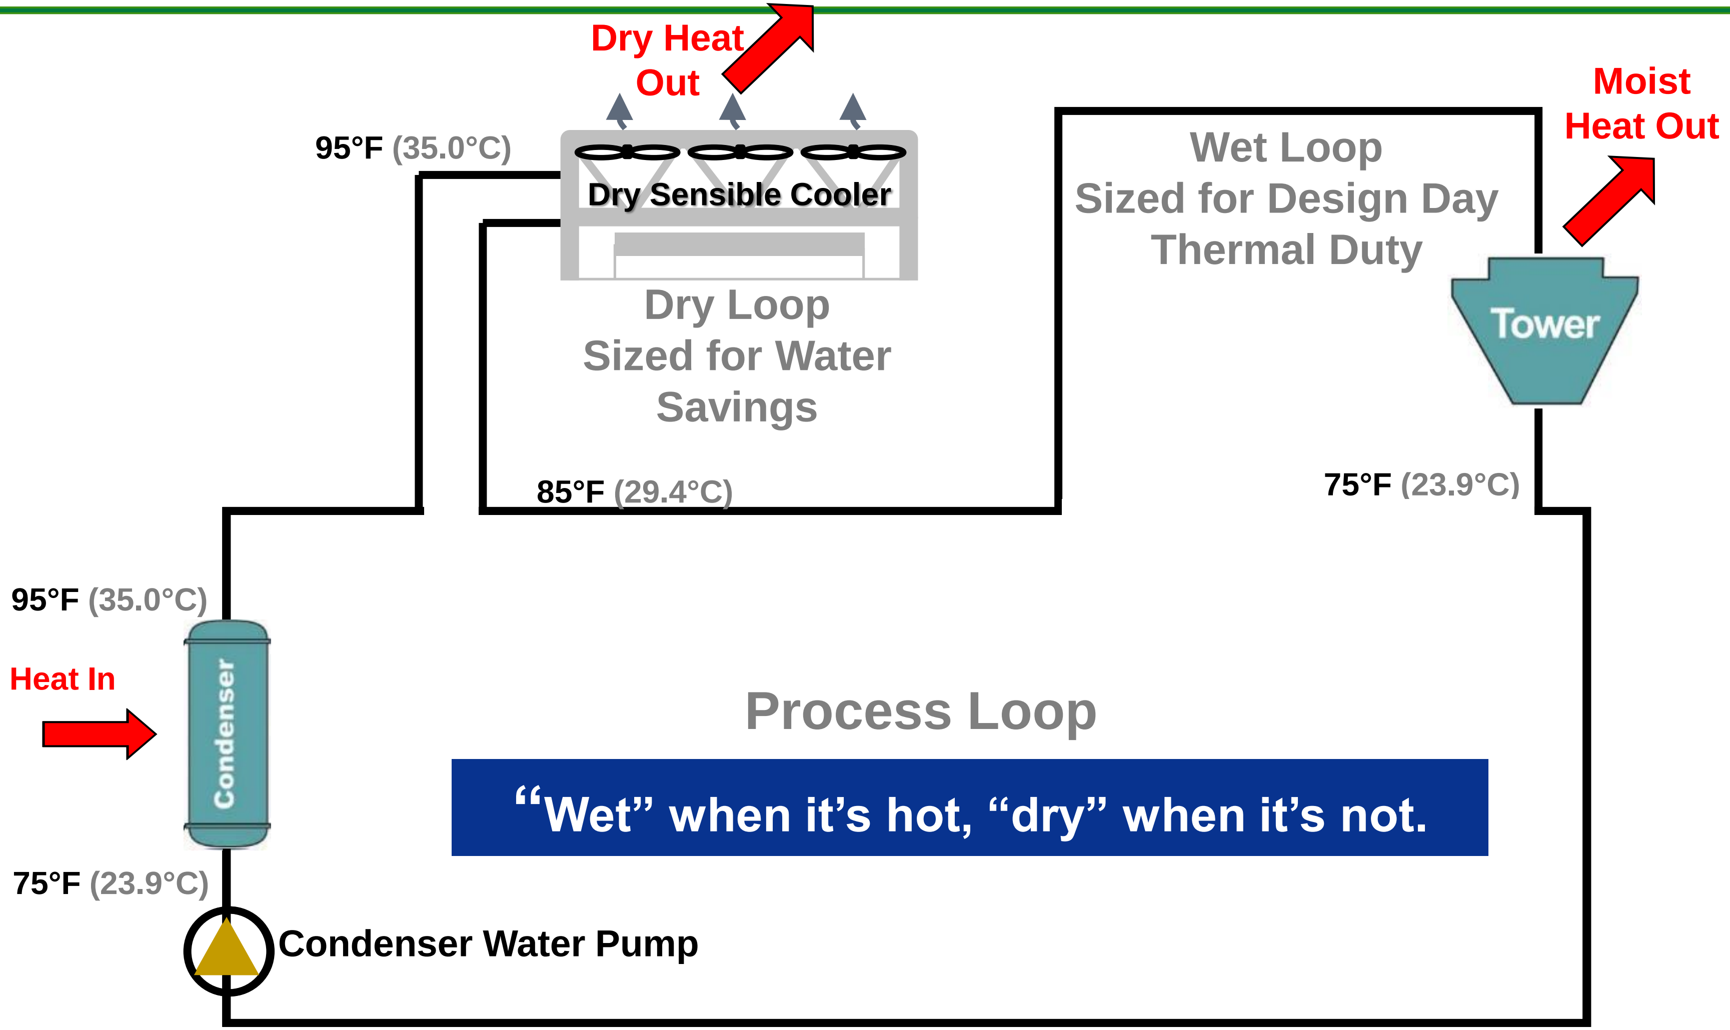

Improved WUE—Thermosyphon

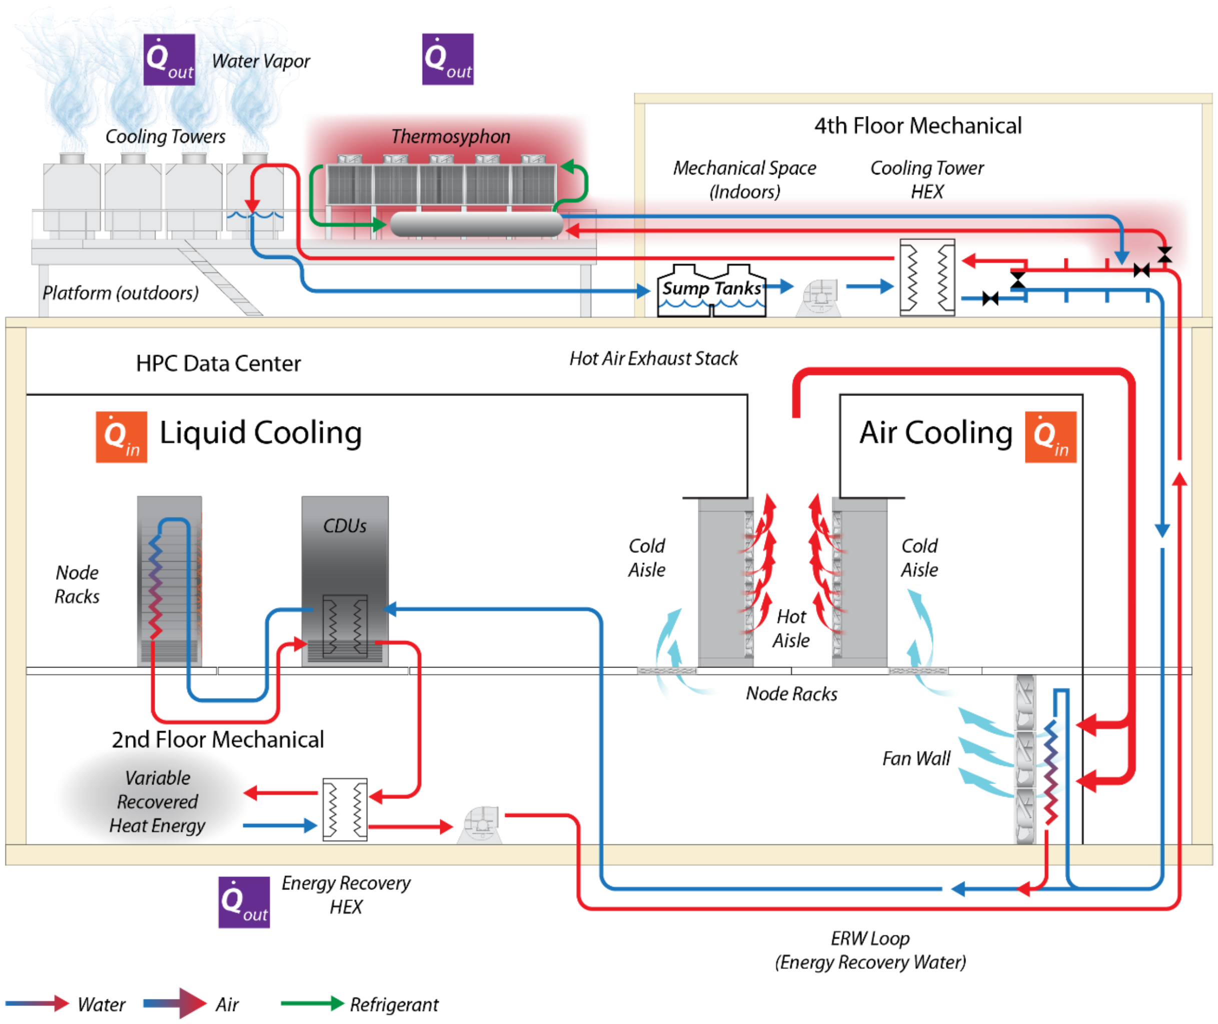

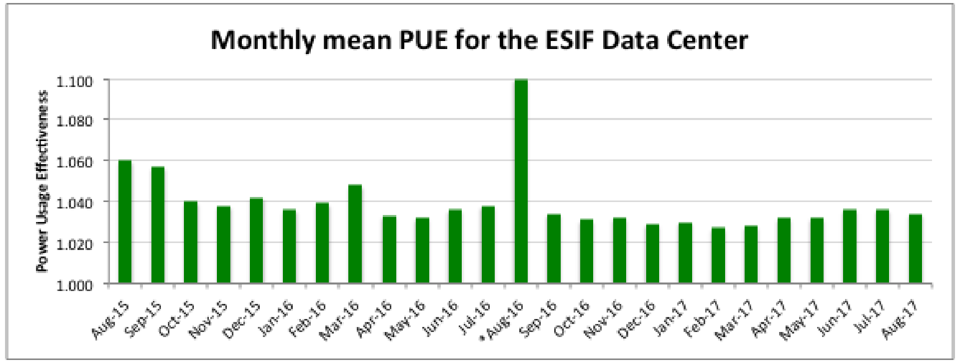

NREL NATIONAL RENEWABLE ENERGY LABORATORY

Johnson Controls

Sandia National Laboratories

ESIF Data Center Efficiency Dashboard

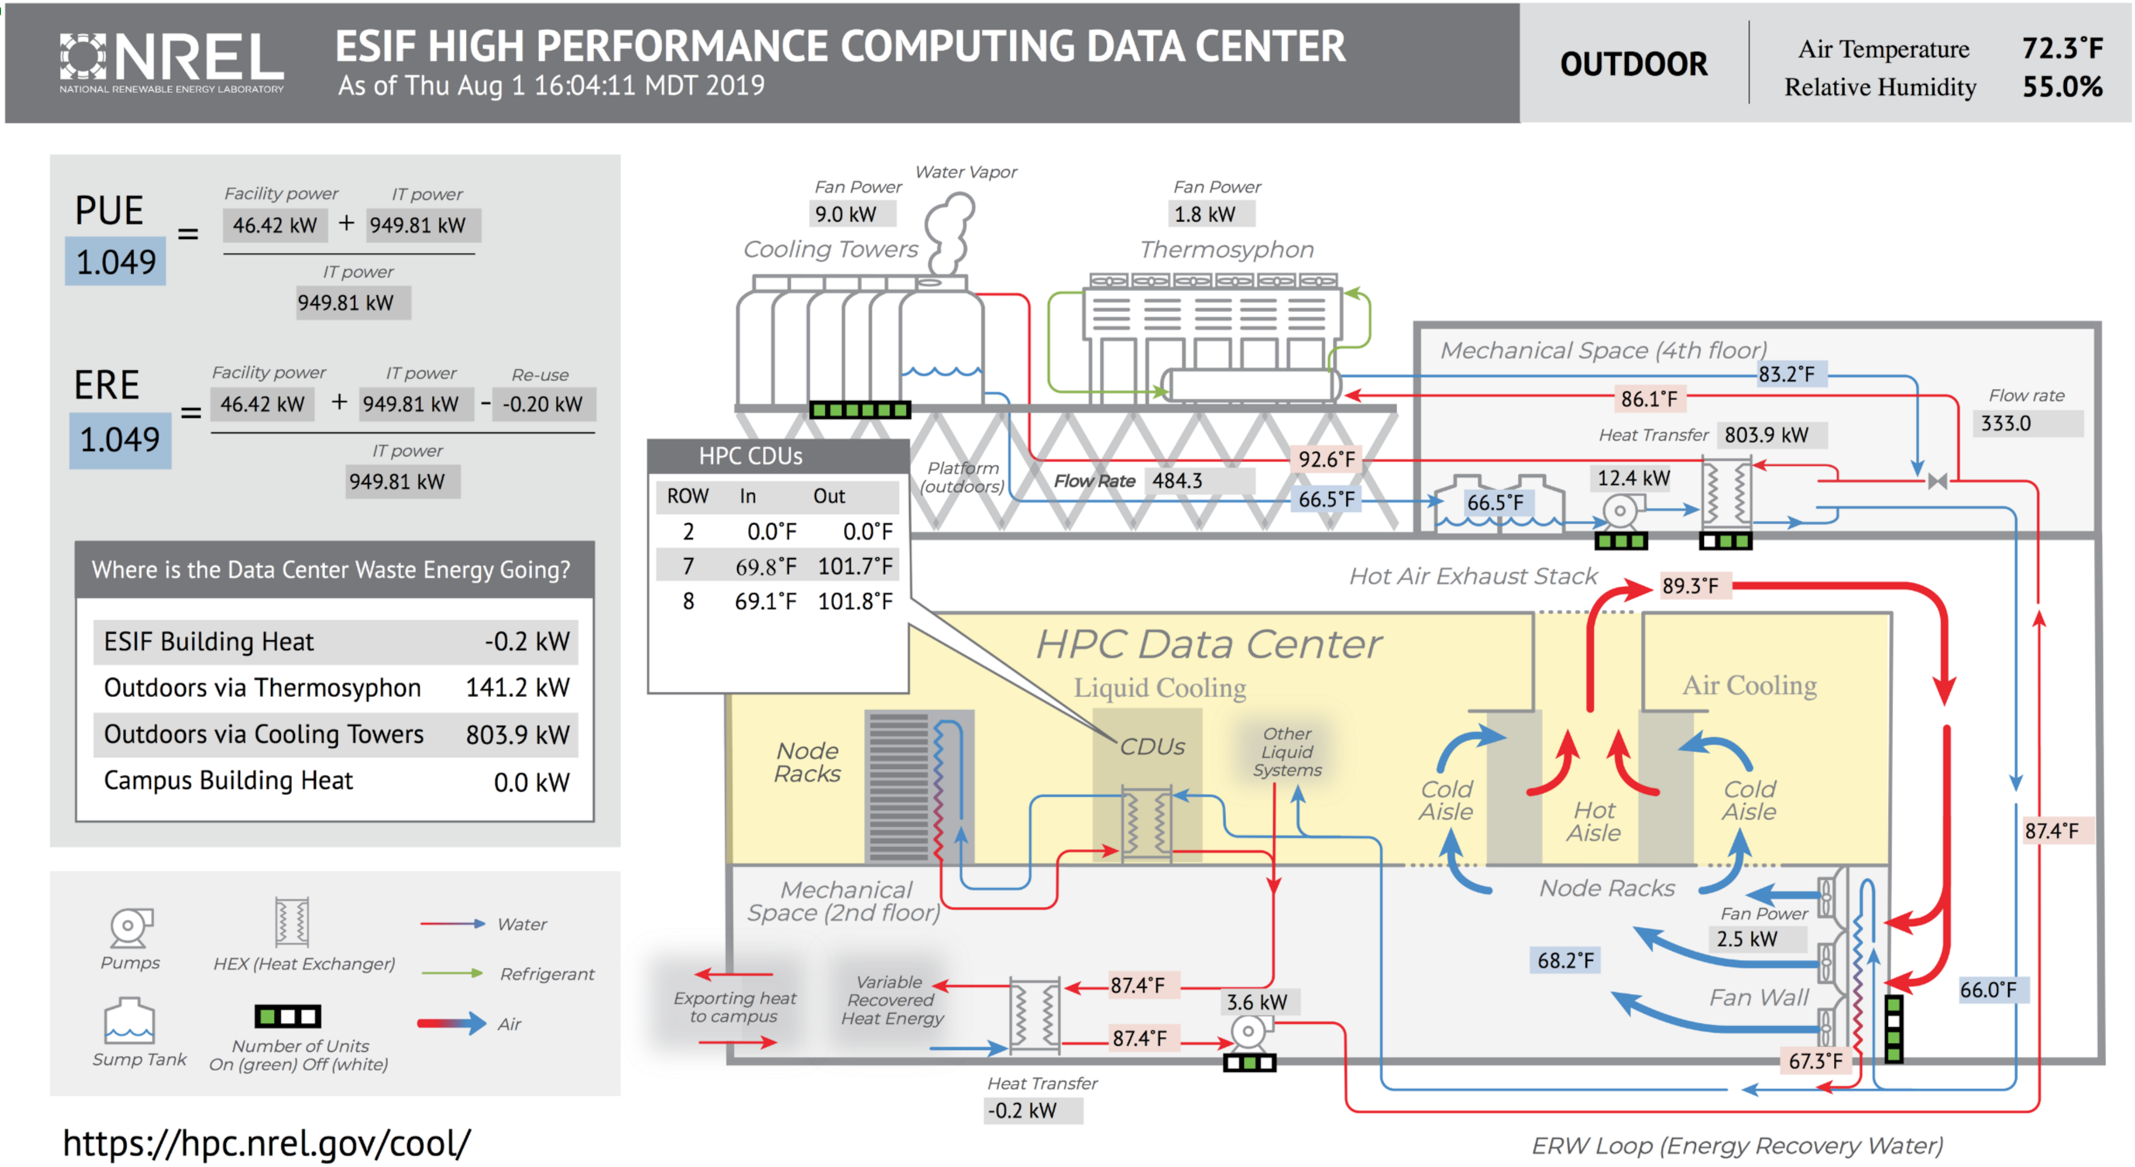

https://hpc.nrel.gov/cool/

Applications

Any application using an open cooling tower is a potential application for a hybrid cooling system, but certain characteristics will increase the potential for success.

- **Favorable Application Characteristics**
  - Year-round heat rejection load (24/7, 365 days is best)
  - Higher loop temperatures relative to average ambient temperatures
  - High water and wastewater rates or actual water restrictions
  - Owner's desire to mitigate risk of future lack of continuous water availability (water resiliency)
  - Owner's desire to reduce water footprint to meet water conservation targets

Sample Data: Typical Loads and Heat Sinks

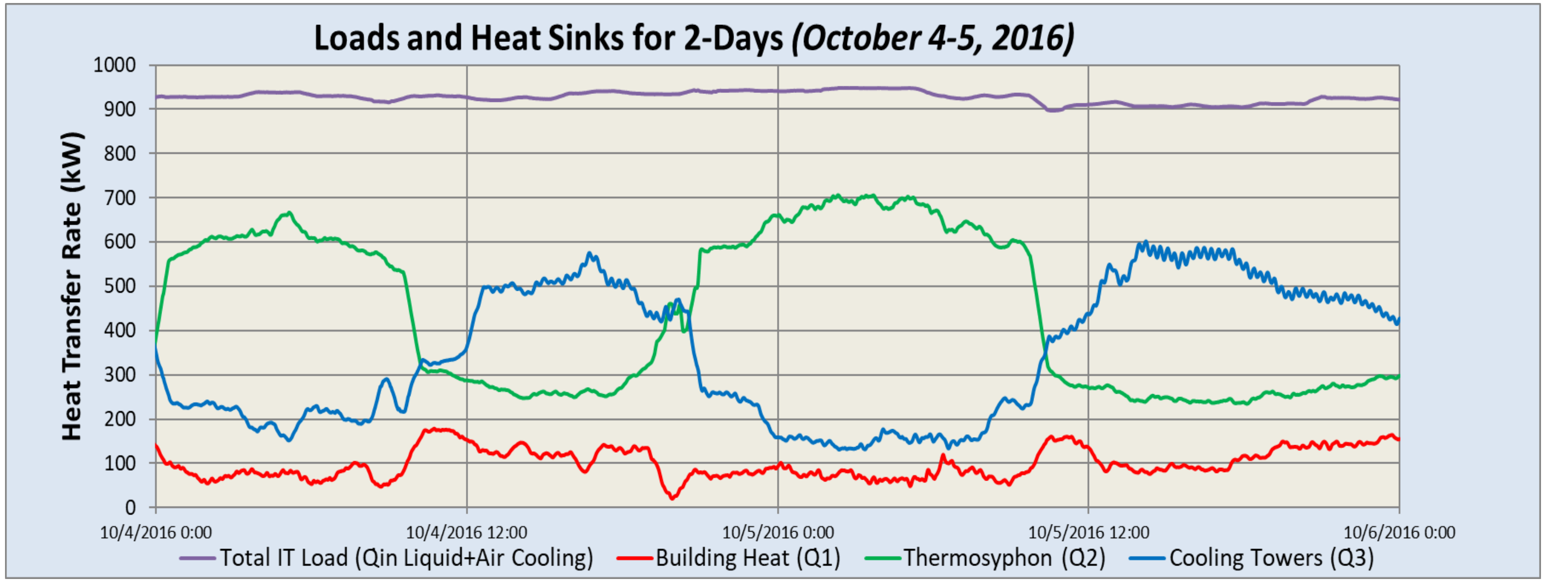

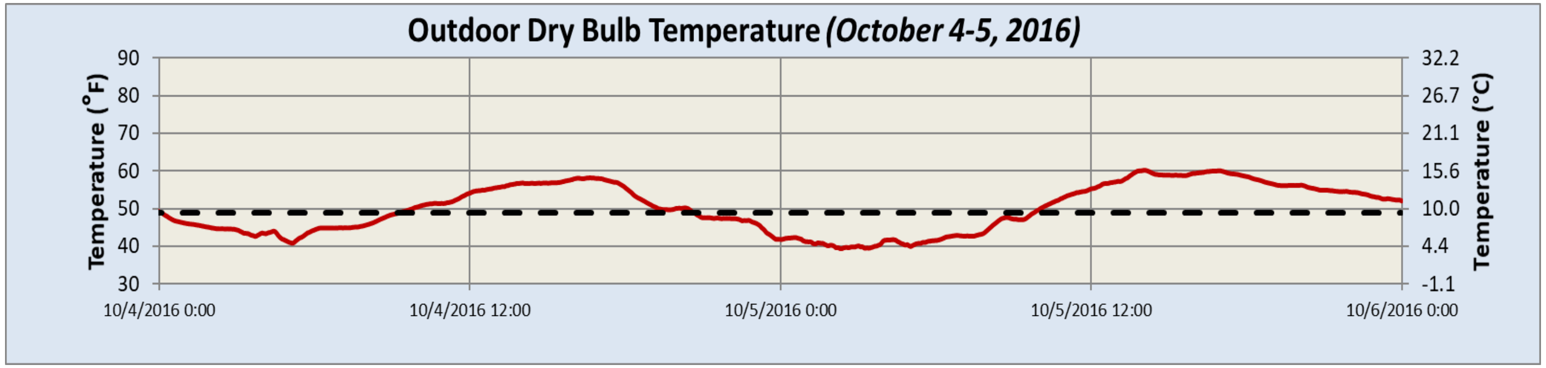

Data Center Metrics

First Year of TSC Operation (9/1/16–8/31/17)

Hourly average IT Load = 888 kW

PUE = 1.034

ERE = 0.929

Annual Heat Rejection

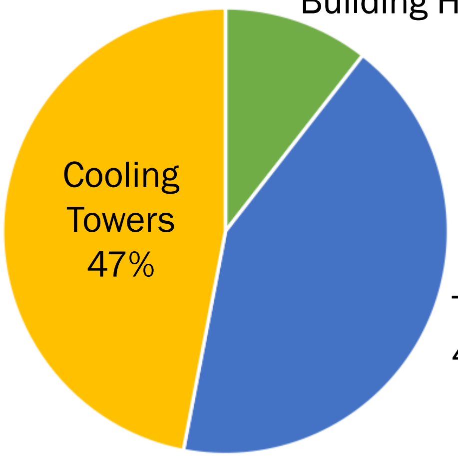

https://www.nrel.gov/docs/fy18osti/72196.pdf

WUE = 0.7 liters/kWh

with only cooling towers, WUE = 1.42 liters/kWh

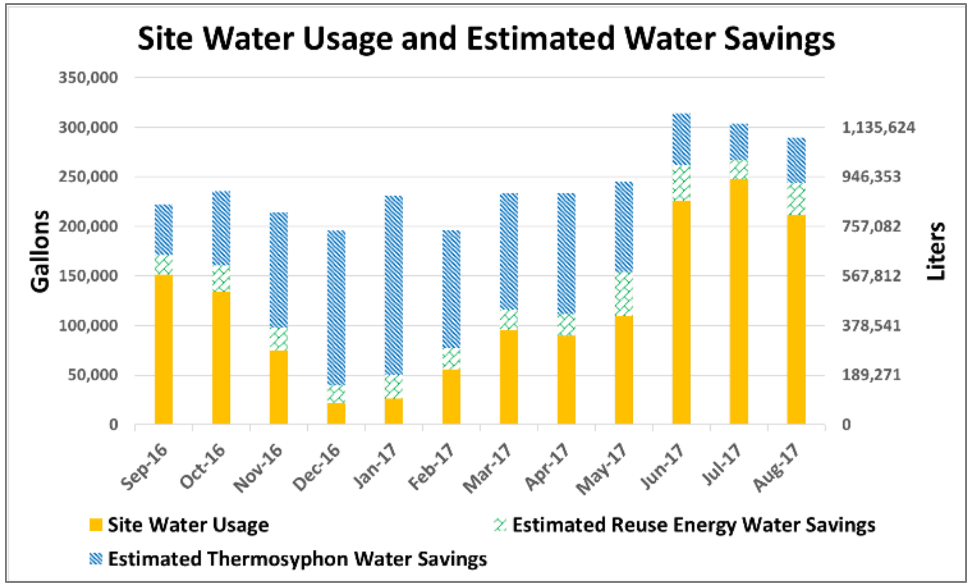

WUE_{SOURCE} = 5.4 liters/kWh
WUE_{SOURCE} = 4.9 liters/kWh if energy from 720 kW PV (10.5%) is included
using EWIF 4.542 liters/kWh for Colorado

Cost and Water Savings

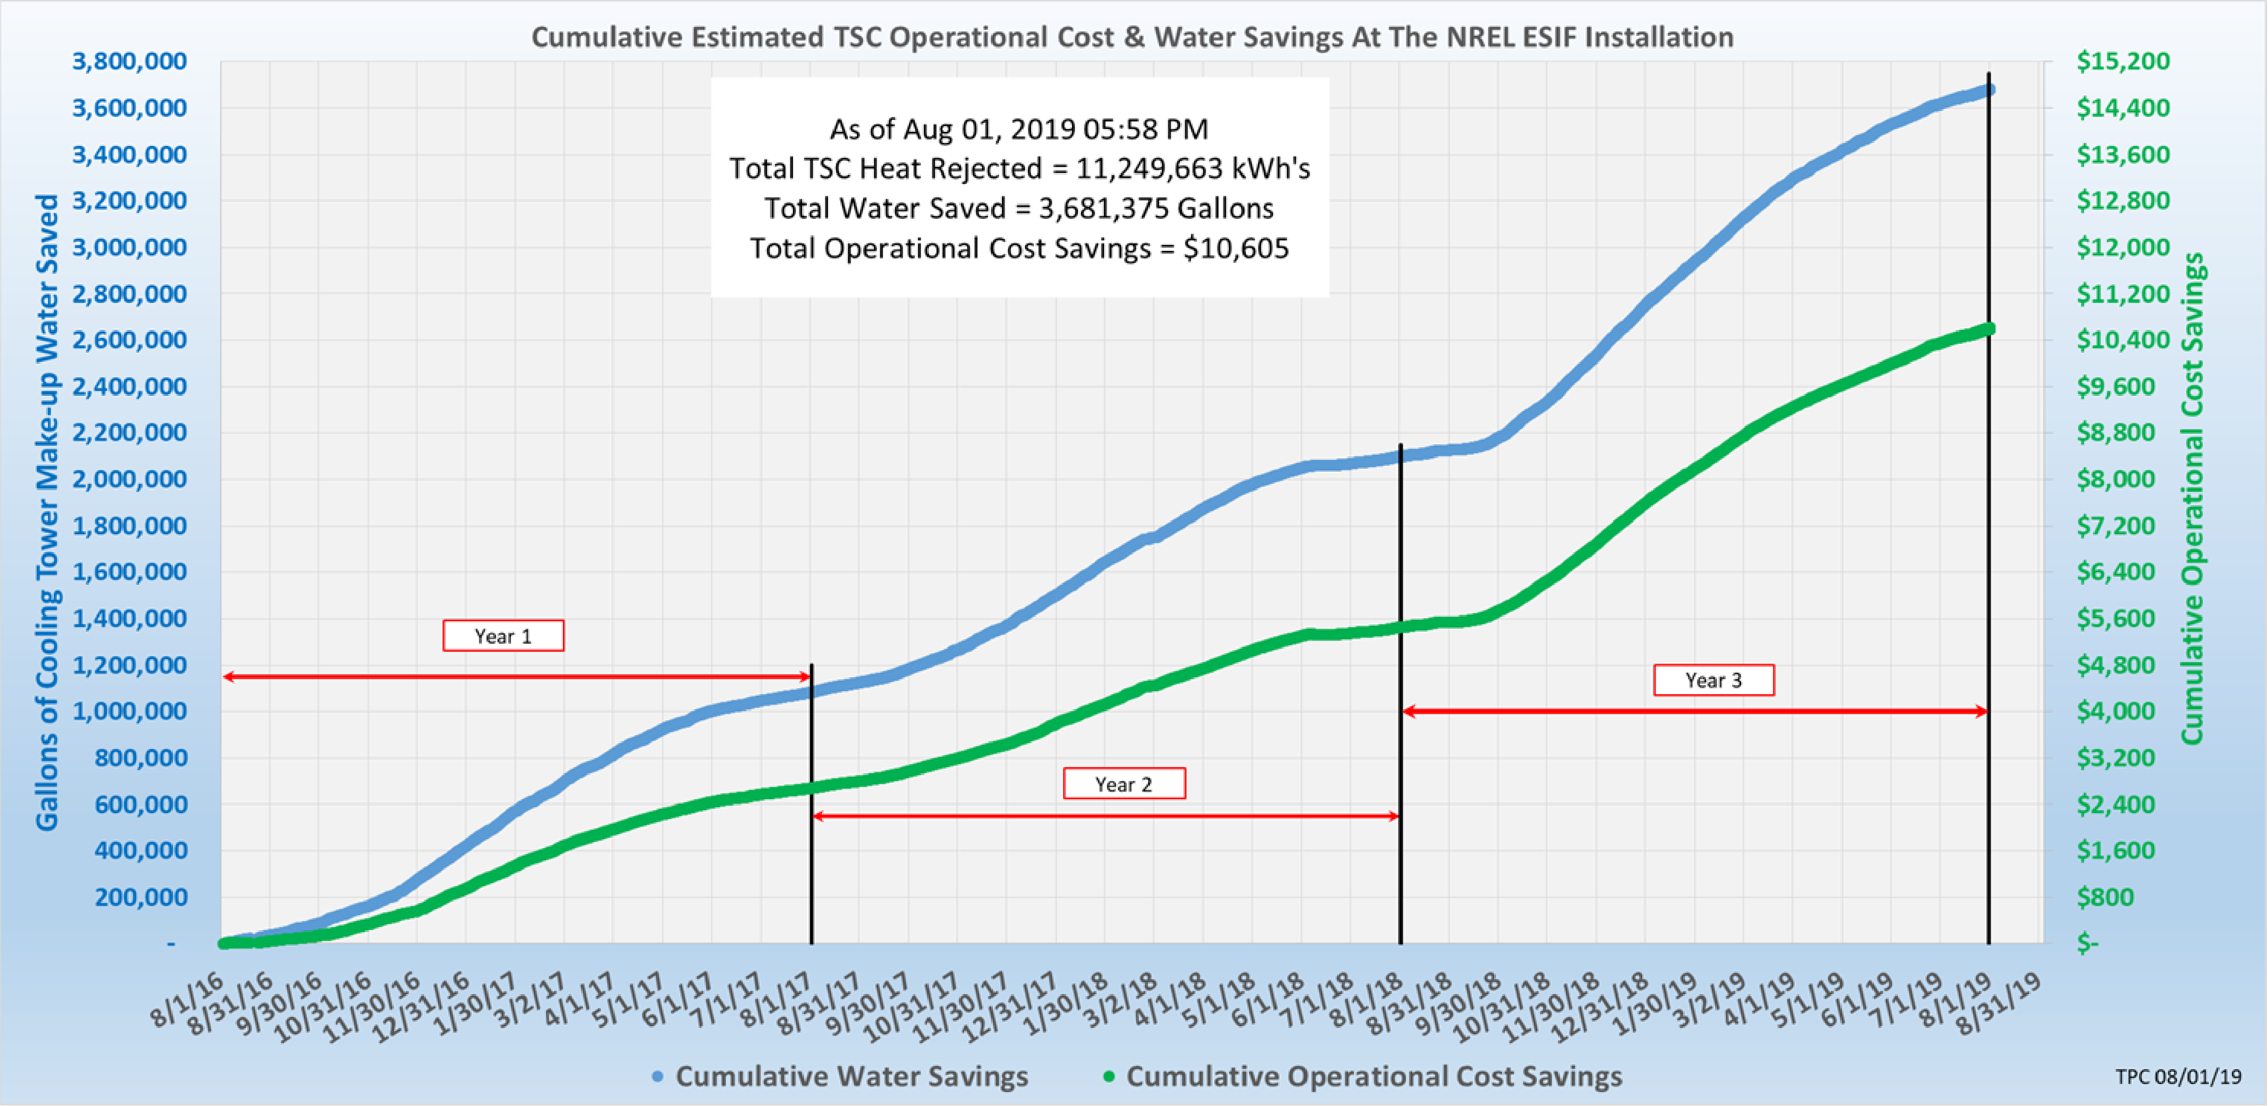

Contact

Otto VanGeet, 303-384-7369, otto.vangeet@nrel.gov

Notice

This research was performed using computational resources sponsored by the U.S. Department of Energy's (DOE) Office of Energy Efficiency and Renewable Energy and located at the National Renewable Energy Laboratory under Contract No. DE-AC36-08GO28308. Funding provided by the Federal Energy Management Program. The views expressed in the presentation do not necessarily represent the views of the DOE or the U.S. Government.

Sandia National Laboratories

7X24 Exchange International
Fall 2019

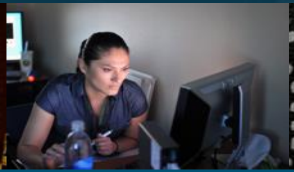

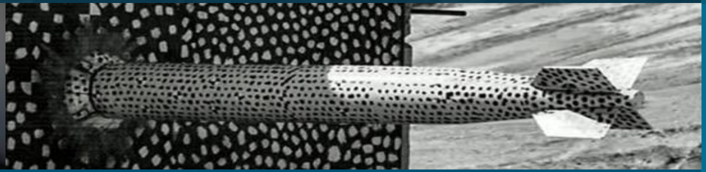

David J. Martinez

Sandia National Laboratories

Liquid Cooling Close Coupled Systems

1st Generation: Liquid indirectly on the chip positive pressure ~typical 65% heat captured to liquid

System 1: Fixed plate cooling technology indirect cooling with thermal interface 90% heat capture to liquid

System 2: Liquid indirectly on the chip works in a vacuum ~70% heat capture to liquid

Immersion cooling several vendors – dielectric fluid with total board submerged with ~100% heat capture liquid

All good solutions with warm water some are less maintenance and less initial infrastructure cost ,several advantages vs disadvantages all solutions are more effective in heat capture than current air solutions along with less OPEX.

System Innovative Solutions: Aquarius Warm Water Cooling "Yacumama"

- Aquarius 480 volt 15 kW half rack Standard Open Compute Project Rack (OCP)
- Emerson 12VDC Power rectifier
- 36 compute nodes (1/2 rack)
- Nodes never throttled back while testing (2 week test) water parameters were in a operating range between 75F - 95F entering at 10 GPM with a delta of 12 to 15 degrees
- Energy Performance Evaluation of Aquila's Aquarius Fixed Cold Plate Cooling System at NREL's High Performance Computing Center
  https://www.nrel.gov/docs/fy19osti/73356.pdf

Node and heat transfer material

Aquarius Inside Look

NREL testing station

Water connections 1" quick coupling

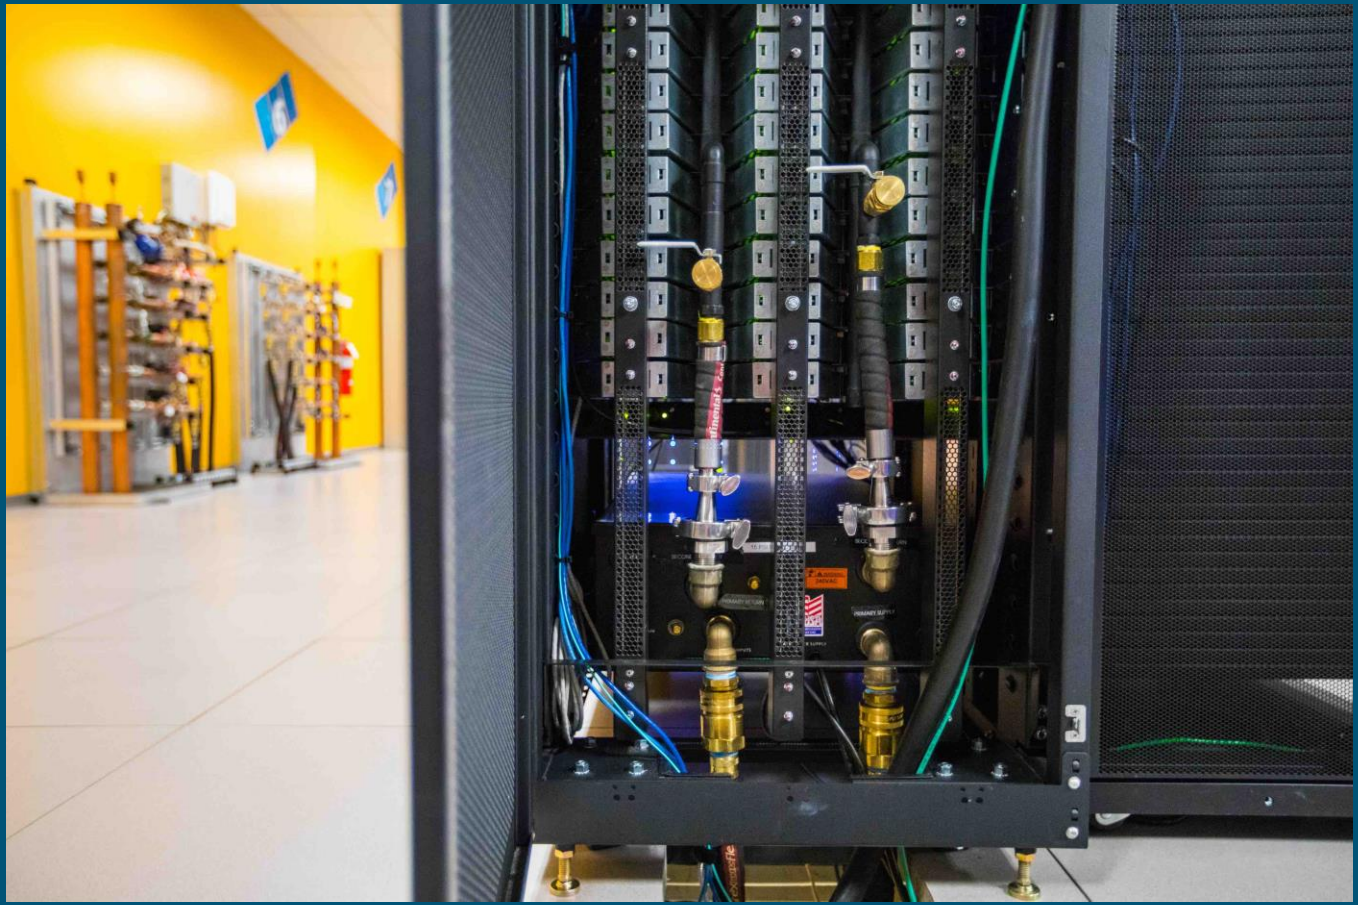

Plates before node install

System 2, ATTAWAY COMPUTER

1,488 compute nodes (8 scalable units)
- 1.93 PFLOPs peak
- Will be the fastest production cluster at Sandia

2.3 GHz processors
- Nodes have dual sockets with 18 cores each
- Intel Xeon Gold 6140 "Skylake"
- ~1.3 TFLOPs per node!

192 GB RAM per node (5.3 GB per core)

Intel Omni-Path high speed interconnect
- 2:1 oversubscription
- Mellanox ConnectX4 HCA's in IO nodes

Penguin Tundra™ hardware platform
- Open Compute (OC) 21" racks
- 3 nodes compute nodes slotted in 1 Open Unit (OU

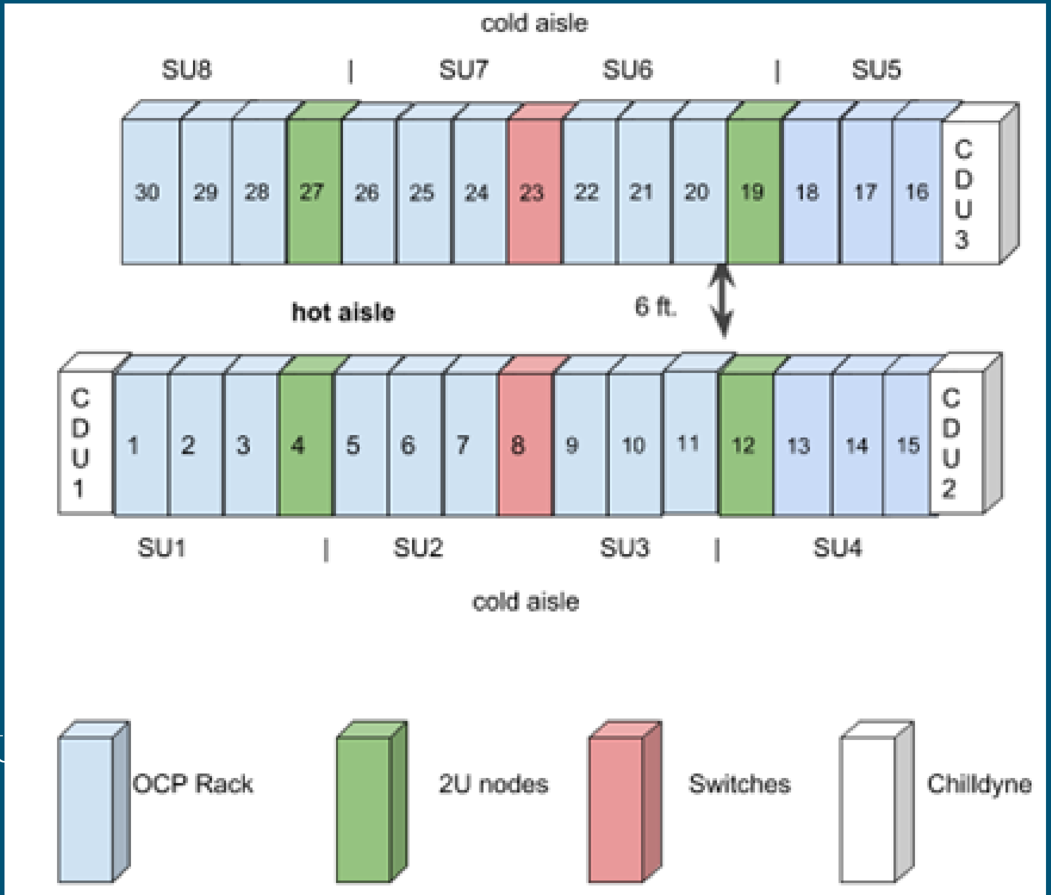

Next generation water cooling:
- Processors are directly cooled by cold plate
  - Fins provide additional air cooling/heat exchange
- Floor based CDU's provide negative pressure to prevent leaks and to improve system resilience.
- ~60-70% of node heat is removed via water

Delivery in August/September 2019

Attaway Computer System

- 24 OCP compute racks 6 sys administrative racks
- 480 volt 60 amp per rack
- ~30 kW per rack 12 volt DC (OCP Bus Bar)
- 70% liquid cooled medium temp supply water 70 F return water ~90 F, entering air temp 78 F
- Reverse Osmosis (RO) process water with rust inhibitor and biocides

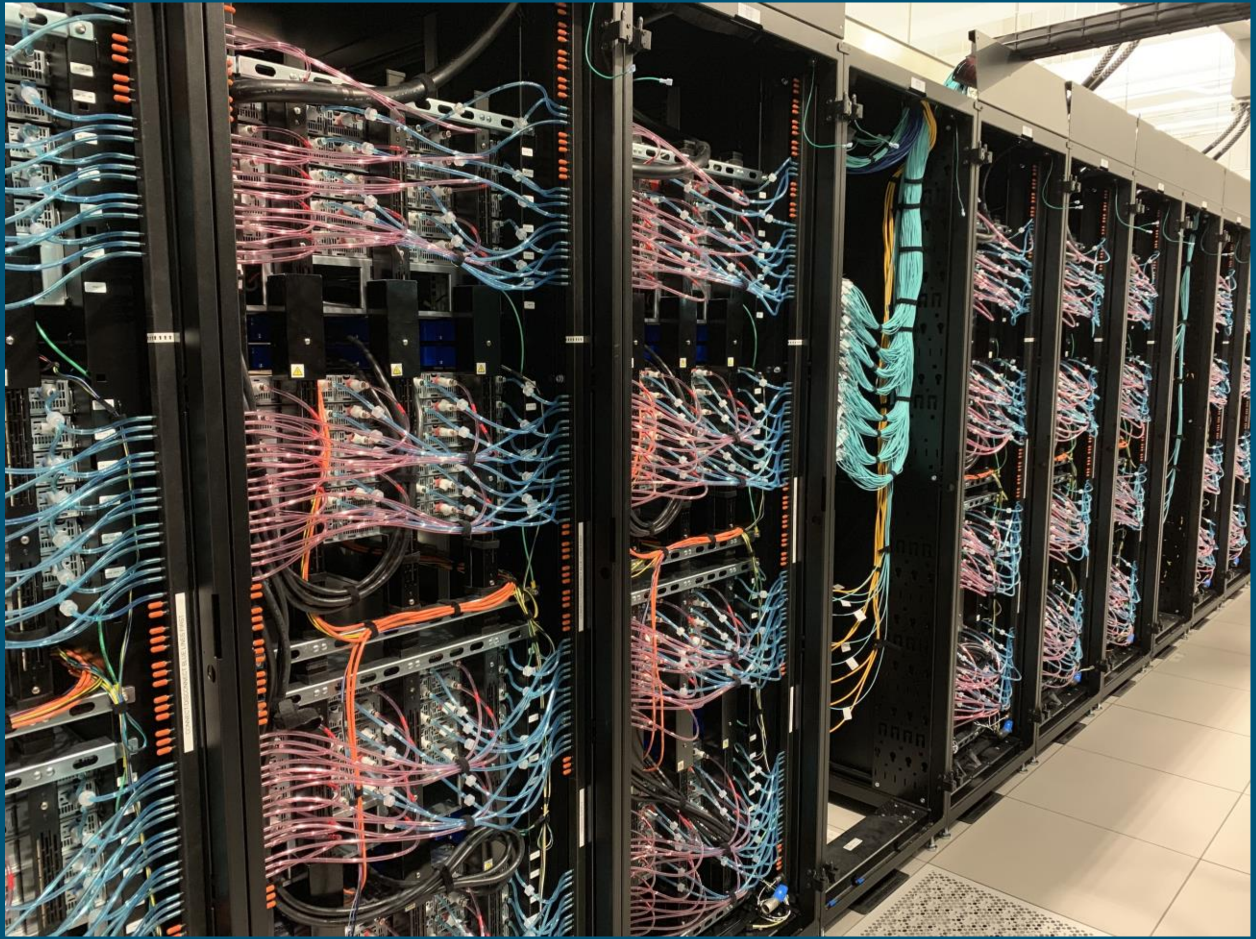

Node For New Attaway Computer

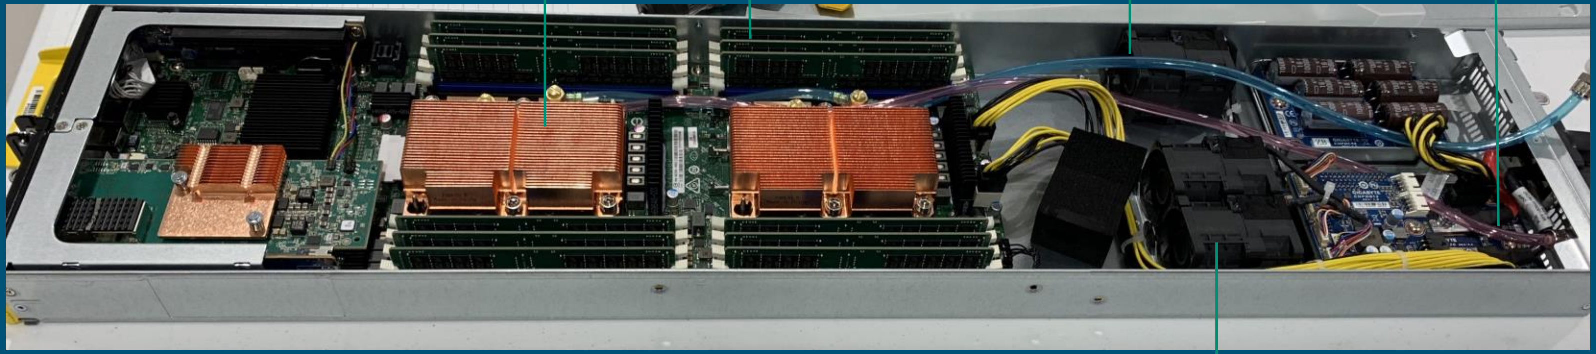

Cooling Distribution Unit (CDU)

1. Pump Chamber
The Chamber is where coolant is stored, supplied to the servers and received from the servers. The system cycles through the main and auxiliary chambers allowing for a steady flow

2. Heat Exchanger (2x)
Transfers the heat created by the servers to the cooling tower or chiller. The HX are connected in series to minimize the processor temperature on hot humid days with warm facility water.

3. Liquid Ring Pump (LRP)
LRP uses water as a seal to provide the required vacuum necessary to propel the coolant. The water seal does not wear out.

4. Microprocessor Control
The temperature in the fluid reservoir is controlled to maintain the coolant temperature above the dew point in the data center.

5. Water Quality Control
The water quality is monitored and controlled to maintain corrosion and bacterial protection. Automatic fill, drain, air purge and leak test are included and coolant additive is stored on board.

6. Coolant Handling Manifold
Standard 6 cooling loops exiting bottom of CDU. Optional single feed coolant supply for ICEdoorTM applications at either bottom or top.

Chilldyne Negative Pressure CDU

- Three CDU's share cooling load that two CDU's can handle
- Each CDU cools up to 300 KW
- Failover valves control which 2 CDU's do the cooling
- Closed loop in data center with 2x heat exchangers in each CDU removes heat to cooling tower
- 480V power + stainless steel chamber pump new for Sandia
- 65 GPM primary water

Chilldyne Cold Plate

Robust Cold Plate design

Includes fins for back up air cooling

Turbulator

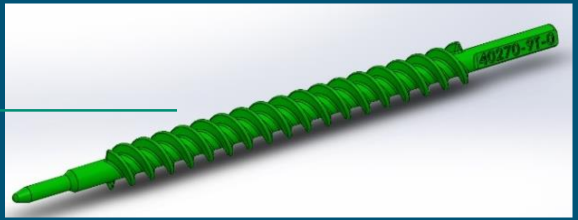

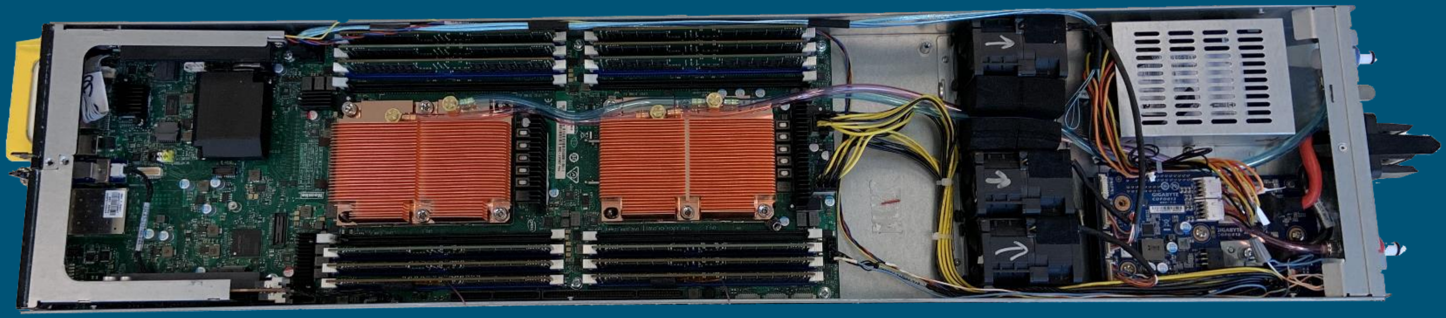

Sandia Rack and manifold design

- Manifold design for negative pressure
- Straight Barbs
- Temperature and pressure sensors in each rack
- Room for additional servers (up to 102 possible)

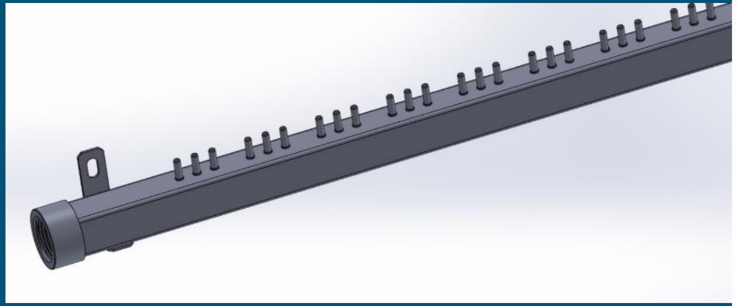

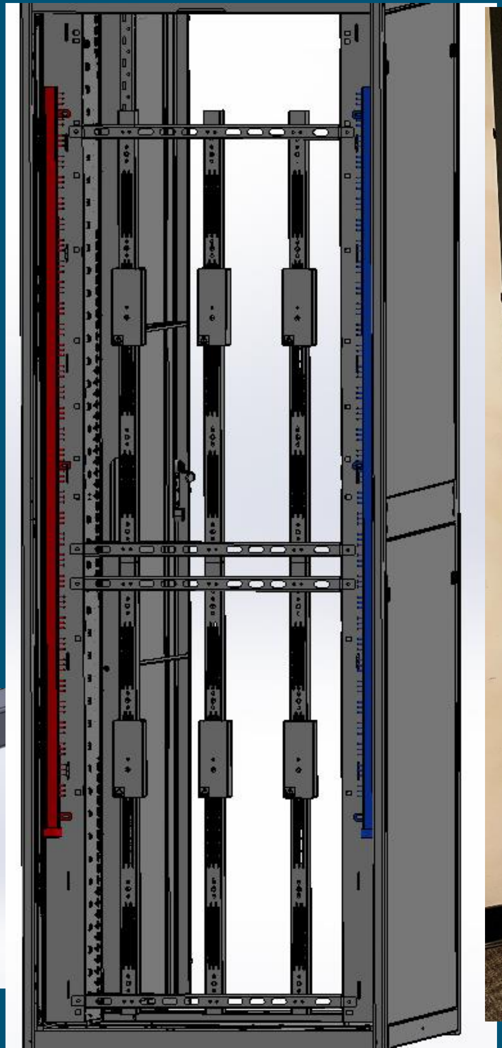

Conclusion (two different but like systems)

Each system have different characteristic's however both systems have quite a few common points:

- Both systems are noninvasive to the system administrator (can service nodes without having a mechanical engineering degree or plumber background)
- Both systems have great heat transfer and are flexible when adding additional systems
- Both systems can utilize warm water cooling with a possibility of capturing reheat use
- Both systems are designed with further enhancement capabilities
- Both systems are designed with ease of accessibility while servicing
- Both systems utilize simplistic chemistry in the water composition
- Both systems facility interface connections are standard to the industry

Contact Info

Dave Martinez

Engineering Program/Project Lead

Infrastructure Computing Services

PO BOX 5800 MS-0823

Office: 505-844-6531

Cell: 505-506-6865

davmart@sandia.gov