Energy Exchange℠

Better Buildings®
SUMMIT

AUGUST 21-23, 2018 • CLEVELAND, OHIO

Liquid in the Rack: Liquid Cooling Your Data Center

Dale Sartor, PE

Lawrence Berkeley National Laboratory

Norman Bourassa

Lawrence Berkeley National Laboratory

Otto Van Geet, PE

National Renewable Energy Laboratory

Energy Exchange | Better Buildings SUMMIT

Contact Information

Dale Sartor, P.E.
Lawrence Berkeley National Laboratory
MS 90-3111
University of California
Berkeley, CA 94720

DASartor@LBL.gov
(510) 486-5988
http://datacenters.lbl.gov/

Agenda

- Introduction to Liquid Cooling
- Case studies
- Interactive Discussion

Liquid Cooling Solution

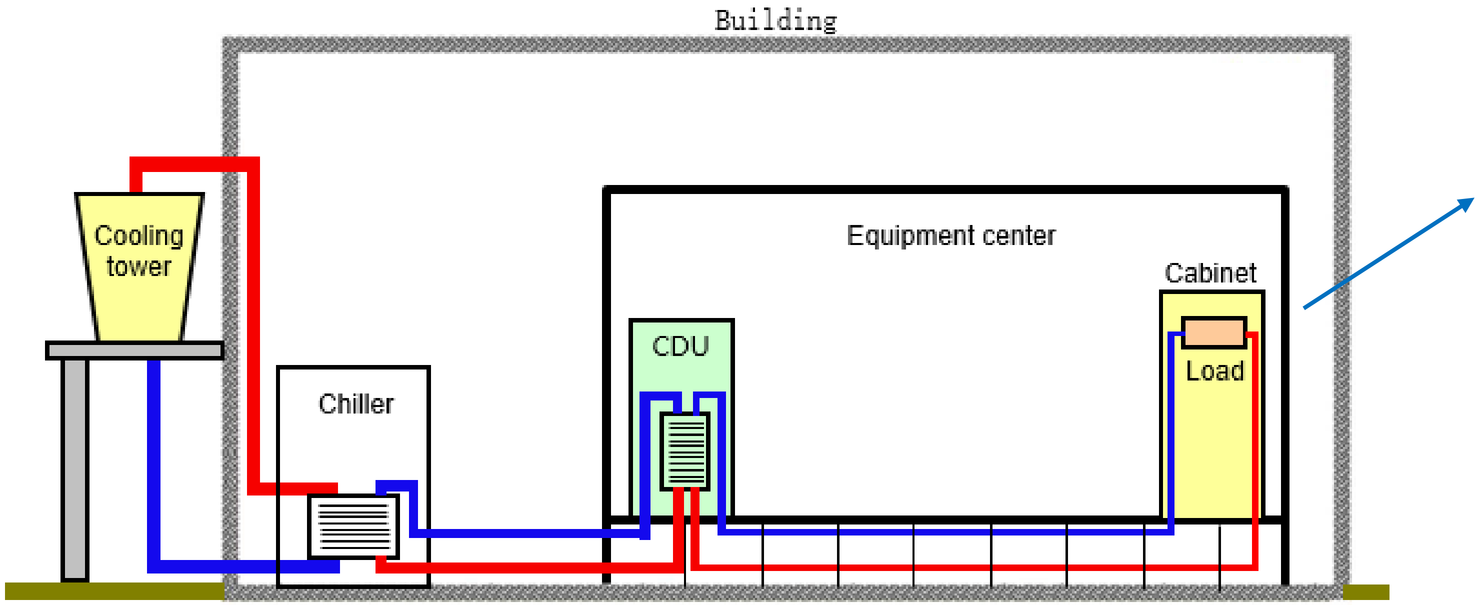

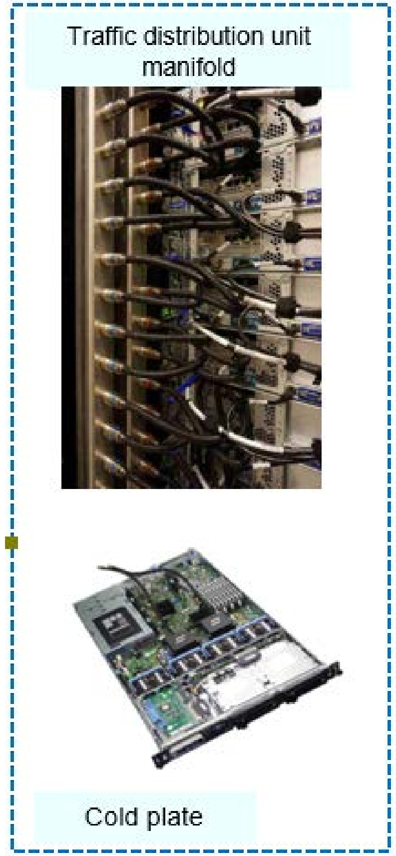

Typical liquid cooled equipment room, with external coolant distribution units (CDUs)

For most locations these data centers may be operated without chillers in a water-side economizer mode

Benefits of Liquid Cooling

- Higher compute densities
- Higher efficiency
  - Heat removal
  - Transport energy
  - cooling plant
  - Increased economizer hours
  - Potential use of waste heat

Moving (Back) to Liquid Cooling

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

(5,380 m^3)

(1.5 m^3)

Water

=

Air

Why Liquid Cooling?

- Higher compute densities
- 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

In-Row Liquid Cooling

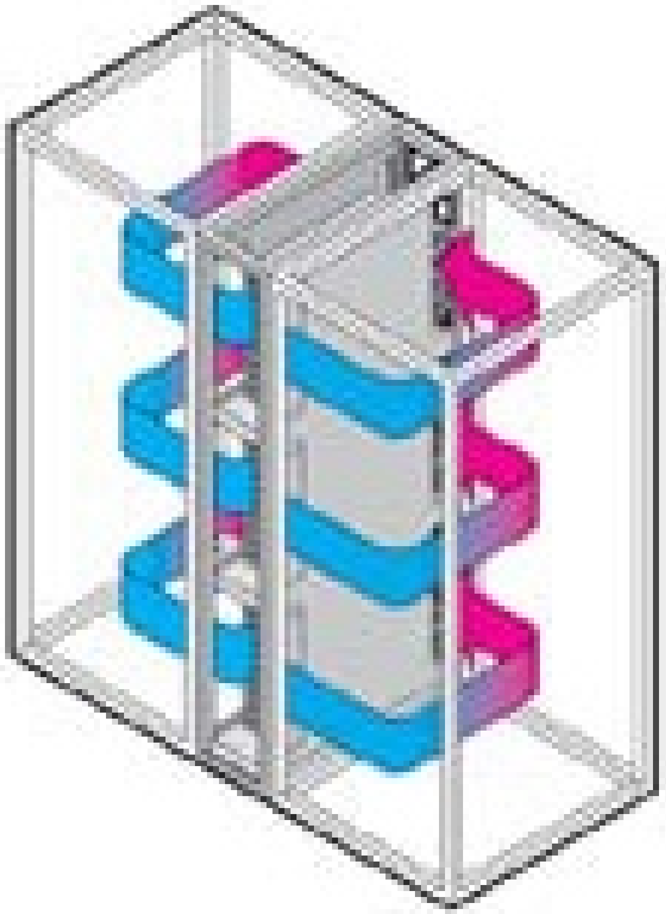

Graphics courtesy of Rittal

In-Rack Liquid Cooling

Racks with integral coils and full containment:

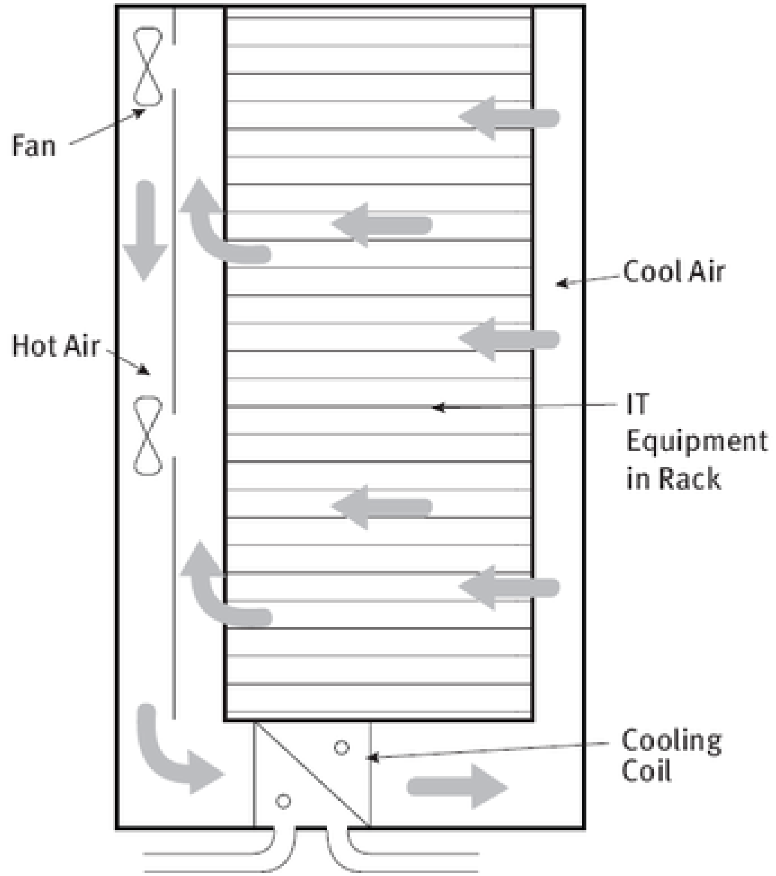

Rear-Door Heat Exchanger

- 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

- Clustered Systems design
- Conducting heat to a cold plate containing refrigerant

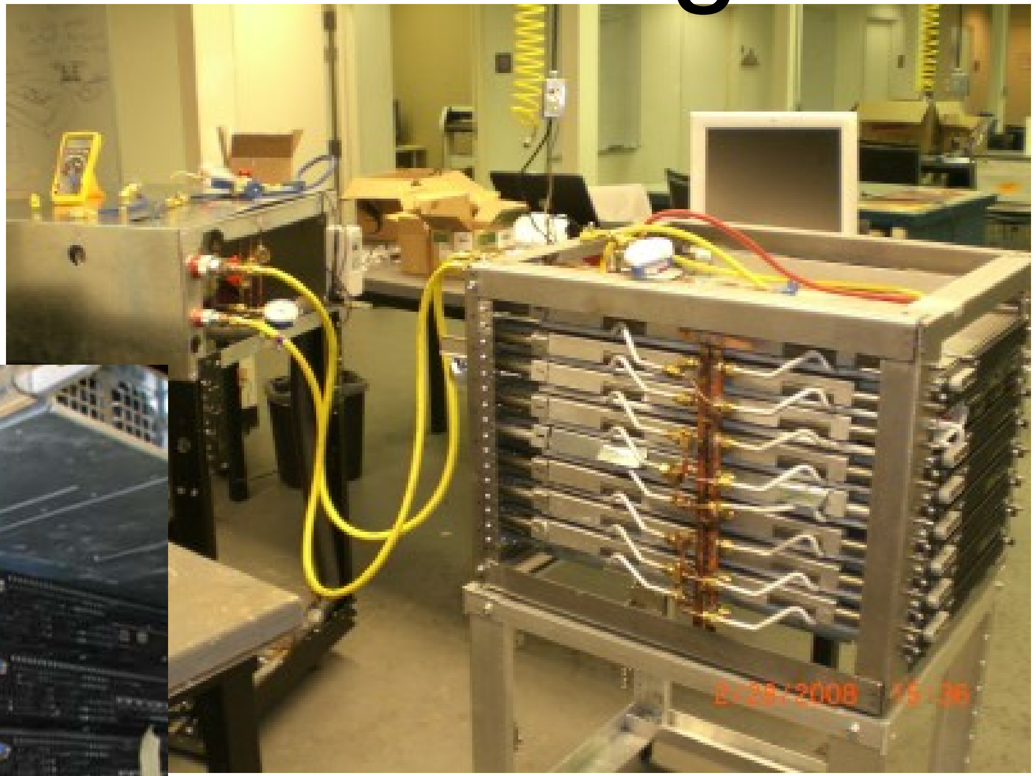

Liquid On-Board Cooling

- Server fans are removed
- Heat risers connect to the top plate, which has a micro channel heat exchanger

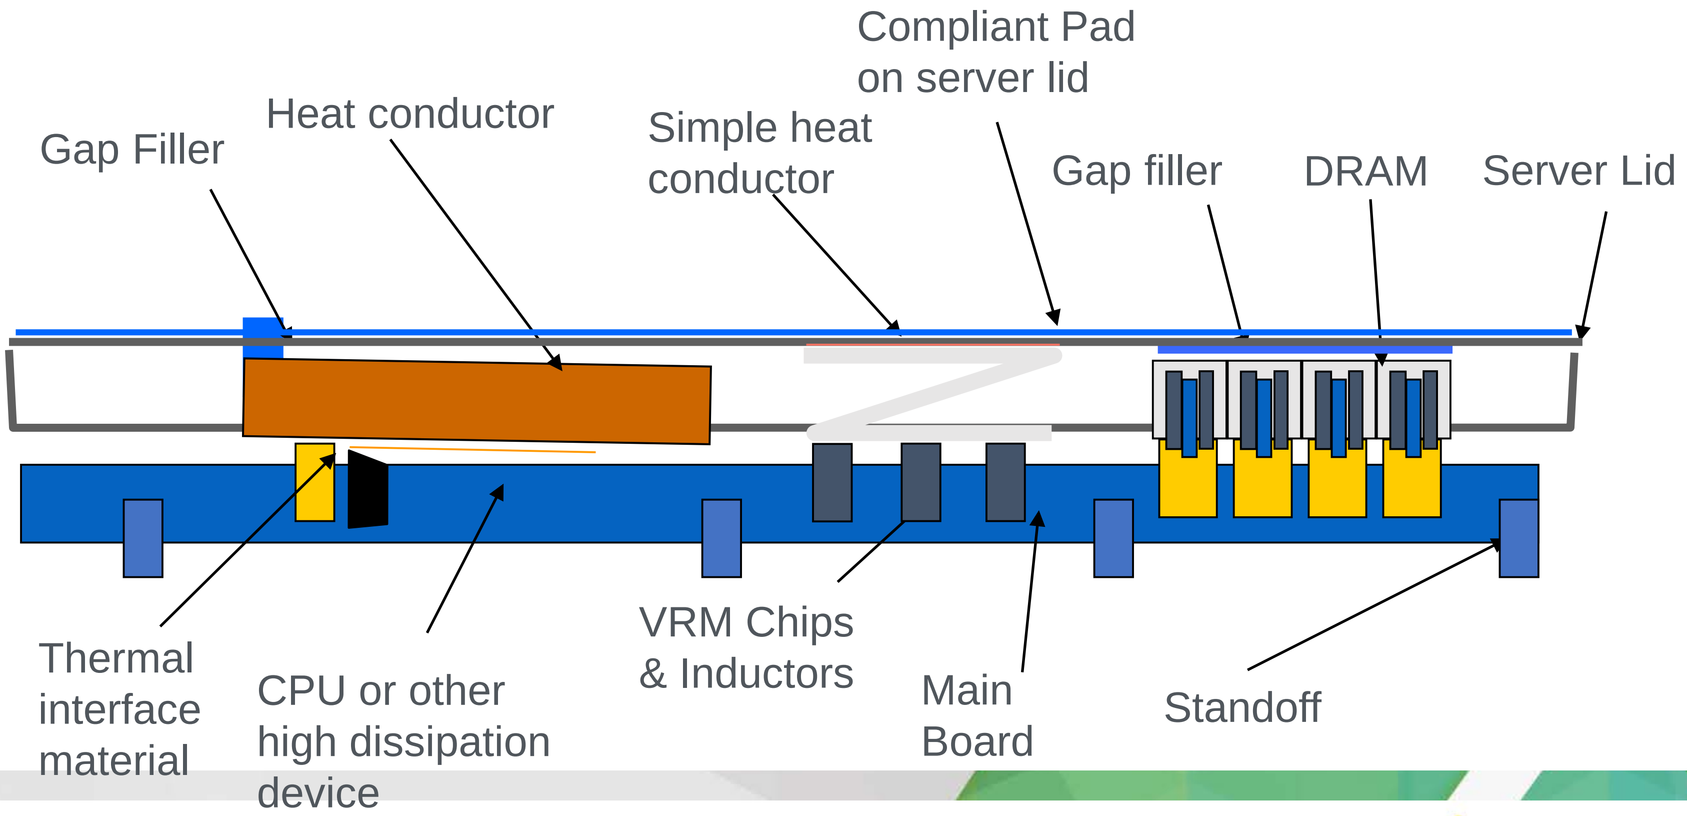

Liquid On-Board Cooling

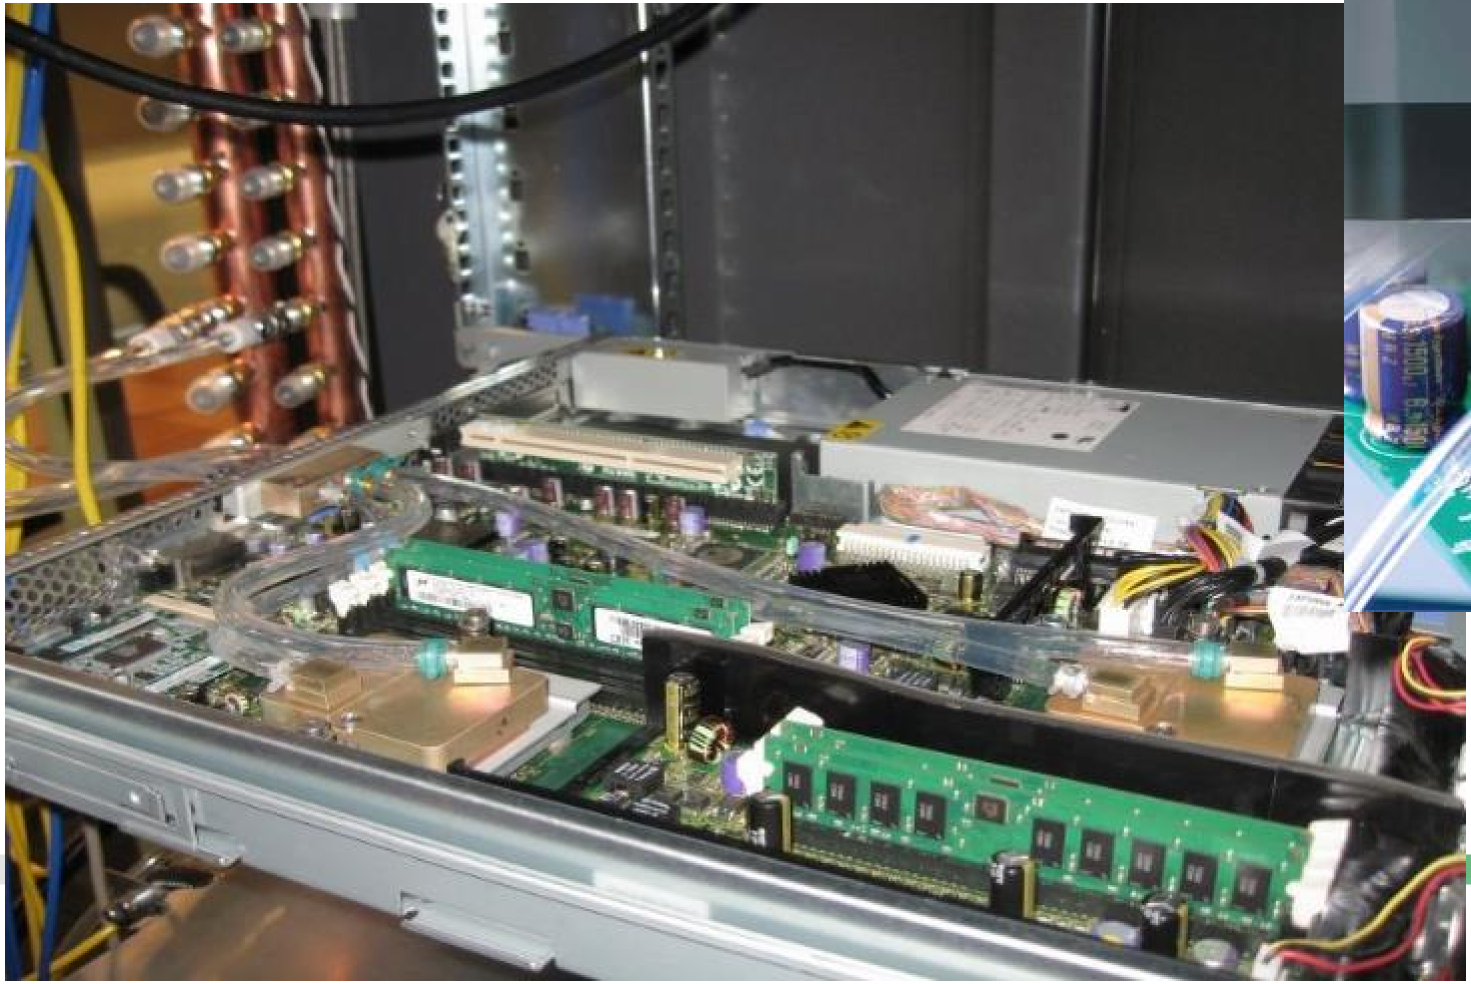

Liquid Immersion Cooling

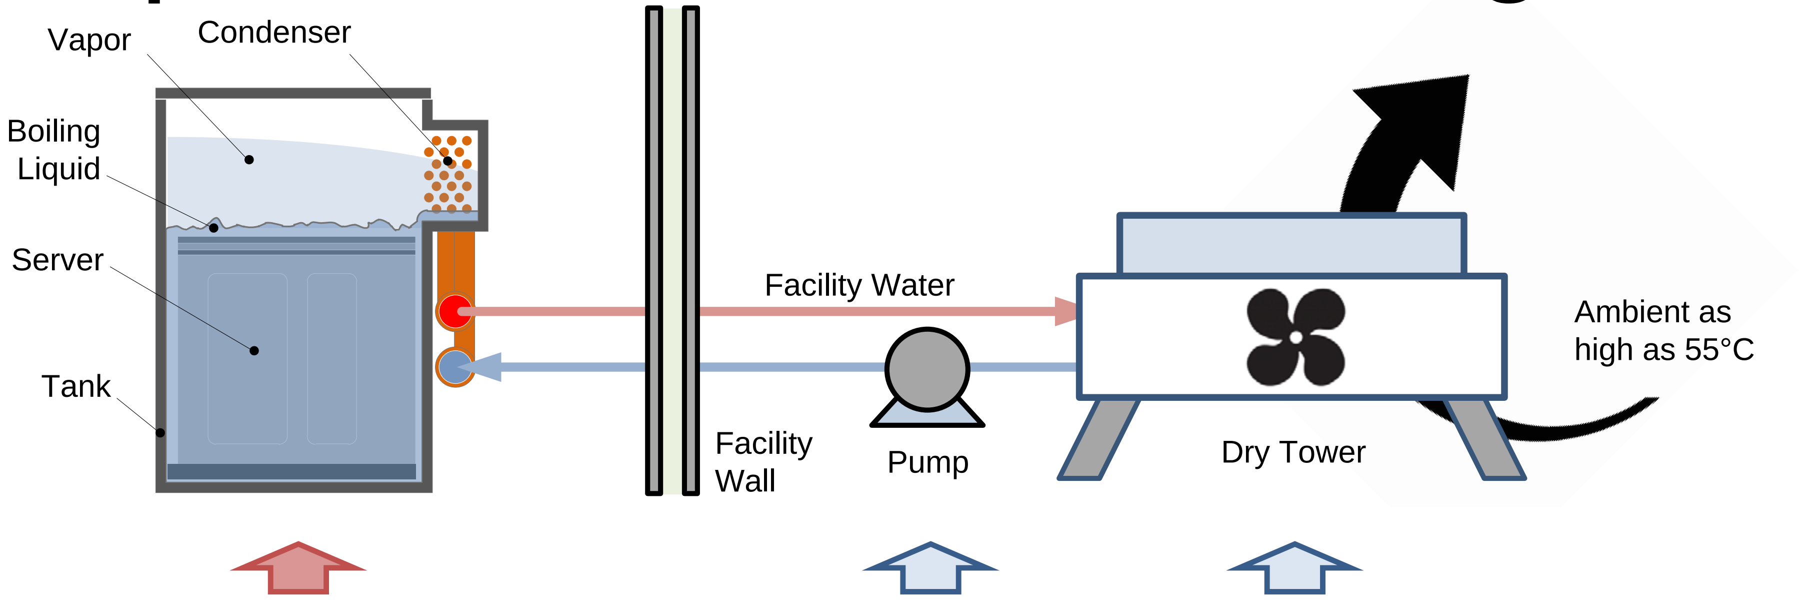

IT Power

Cooling Power = Pump + Fan

No longer requires:

- chillers
- cooling towers
- water use
- raised floors
- computer room air conditioners
- earplugs!

Phase Change of Dielectric Fluid Removes Heat Efficiently

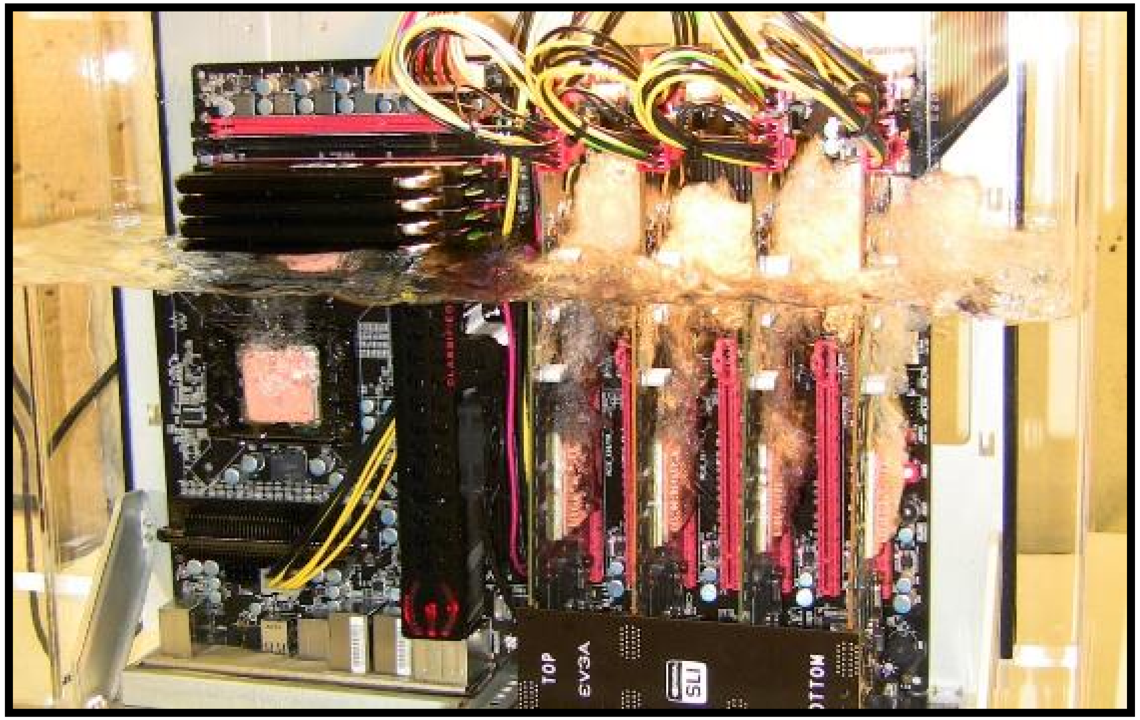

Computer in glass tank

3M Corp.
4 server system

"Chill-Off 2" Evaluation of Liquid Cooling Solutions

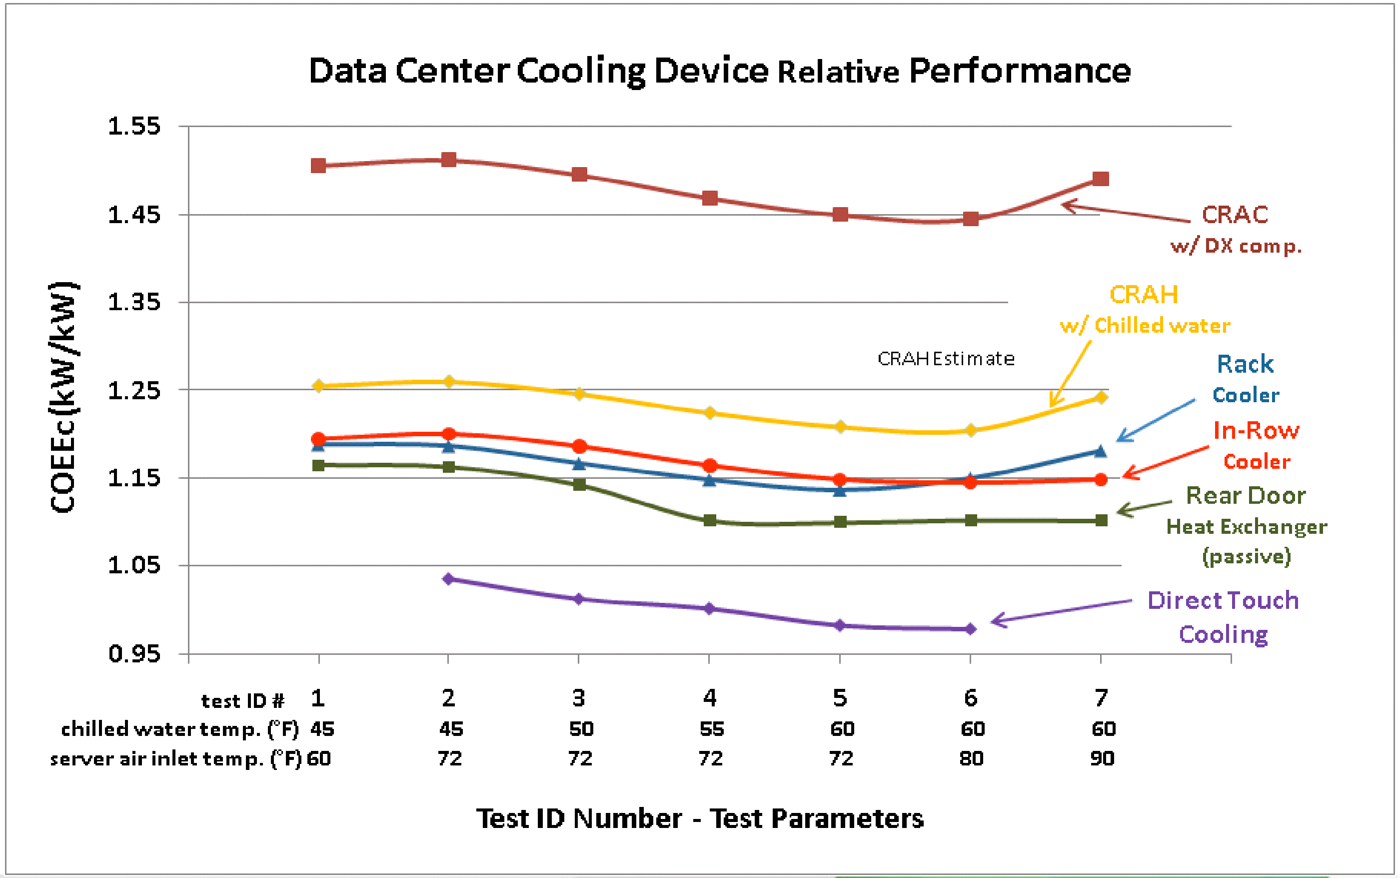

“Free” Cooling w/ Water-Side Economizers

- 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 compressor cooling

Cooling tower and HX = Water-side Economizer

Potential for Tower Cooling

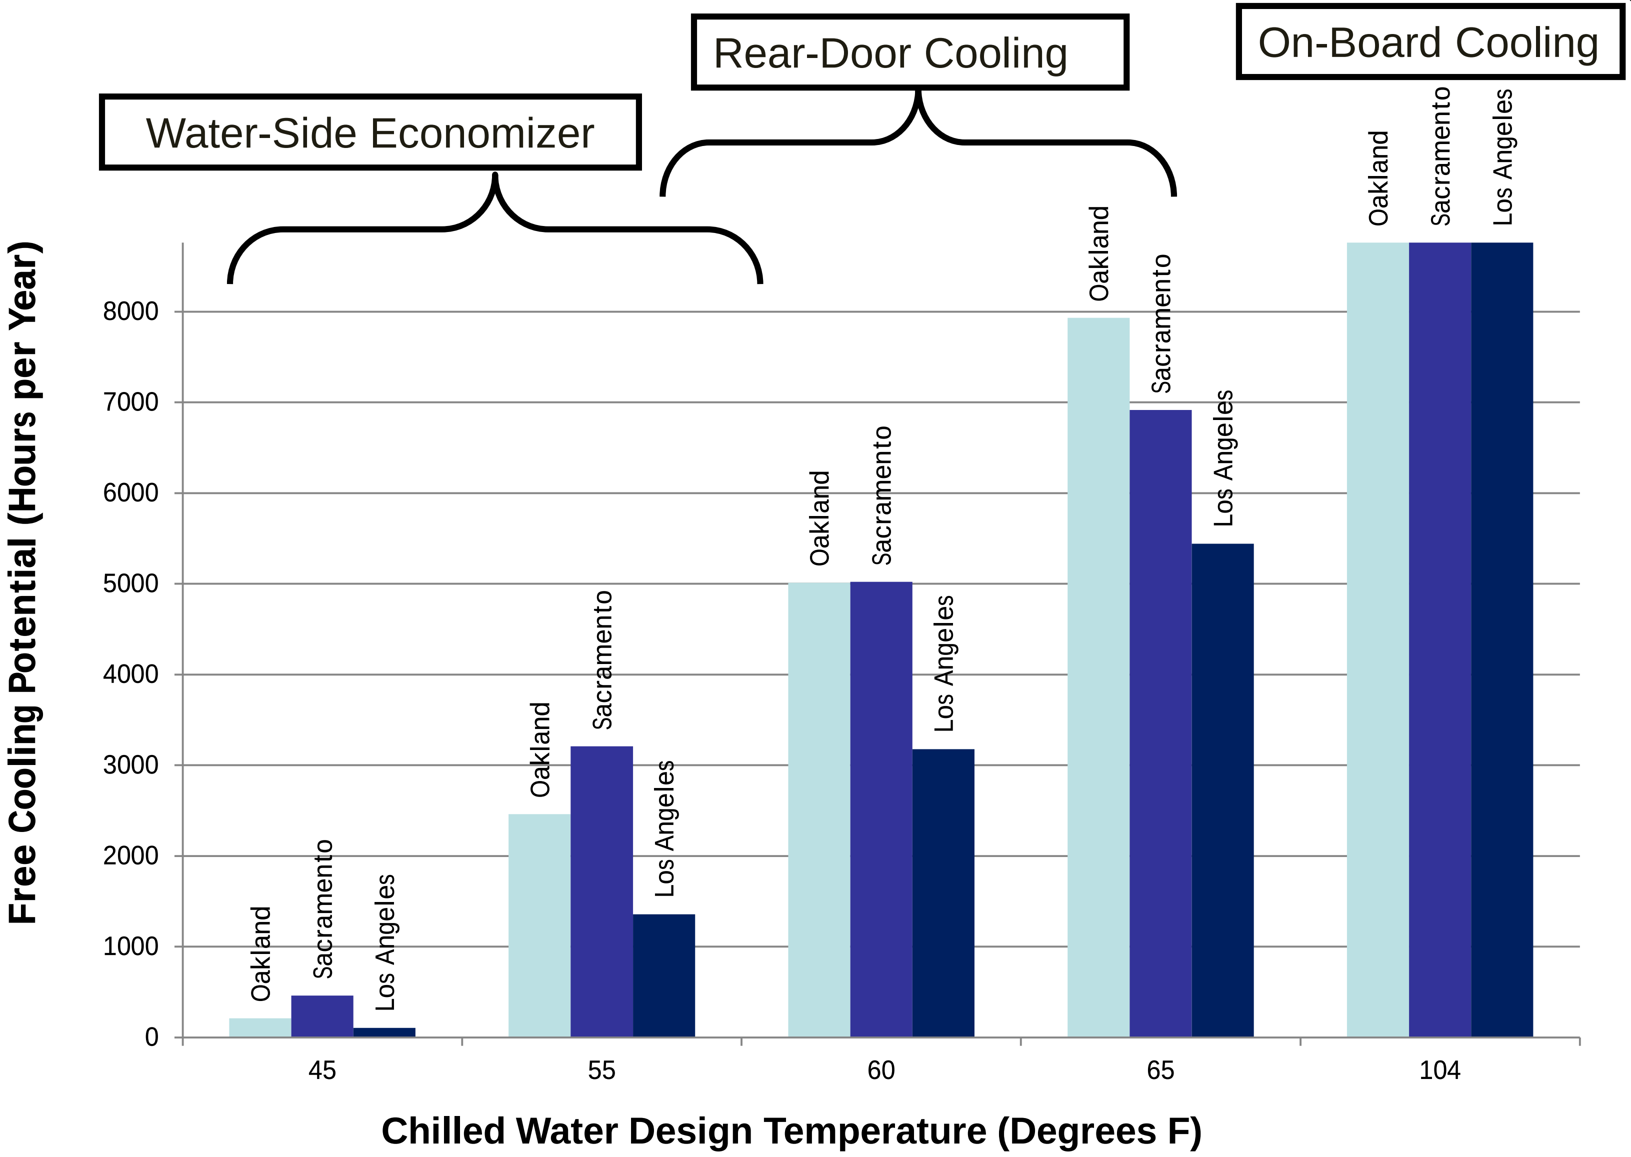

Re-Use of Waste Heat

- 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 feed 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 (in Europe)

ASHRAE Design Reference Conditions - 2015

Air side Conditions

Legacy equipment → A1

New racks with rear door ht exchangers → A3

Class	2015 ASHRAE Equipment Environmental Specifications for Air Cooling						
	Product Operations					Product Power Off	
	Dry-Bulb Temperature °F	Humidity Range, Non-Condensing	Maximum Dew Point °F	Maximum Elevation ft	Maximum Temperature Range [15 min/1 hr] °F	Dry-Bulb Temperature °F	Relative Humidity %
Recommended (Suitable for all 4 classes)							
A1 to A4	64.4 to 80.6	15.8°F DP to 59°F DP and 60%RH					
Allowable							
A1	59 to 89.6	10.4°F DP & 8% RH to 62.6°F and 80% RH	62.6	10,006	9/36	41 to 113	8 to 80
A2	50 to 95	10.4°F DP & 8% RH to 69.8°F and 80% RH	69.8	10,006	9/36	41 to 113	8 to 80
A3	41 to 104	10.4°F DP & 8% RH to 75.2°F and 85% RH	75.2	10,006	9/36	41 to 113	8 to 80
A4	41 to 113	10.4°F DP & 8% RH to 75.2°F and 90% RH	75.2	10,006	9/36	41 to 113	8 to 80
B	41 to 95	8% to 82.4°F DP and 80% RH	82.4	10,006	9/36	41 to 113	8 to 80
C	41 to 104	8% to 82.4°F DP and 80% RH	82.4	10,006	9/36	41 to 113	8 to 80

Water side Conditions

New Racks
(W2 Condition - with the exception of delivering 84 deg F water on peak WB days) → W2

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	Chiller/Cooling Tower	Water-side Economizer (With Drycooler or Cooling Tower)	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

Liquid Cooled Rack Specification

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

Goal for Liquid Cooled Rack Specifications

- Goal:
  - *A liquid cooled rack specification that could accommodate multiple vendors and provide a reusable infrastructure for multiple refresh cycles with a variety of liquid cooled servers/suppliers.*
- The working group has focused on the following specifications:
  - The wetted material list (all components must be compatible with this list)
  - Water based transfer fluid quality and treatment
  - A universal (multi-vendor) quick connect
  - General operating specifications such as pressure and temperature ranges
- https://datacenters.lbl.gov/technologies/liquid-cooling

International Open Data Center Specifications

- Target Organizations:
  - The Open Compute Project (OCP) – US
- Collaborators:
  - Facebook
  - Google
  - Intel
  - Microsoft
  - Baidu
  - Alibaba
  - Tencent
  - YOU

Resources

- Center of Expertise for Energy Efficiency in Data Centers
  - datacenters.lbl.gov/
  - https://datacenters.lbl.gov/technologies/liquid-cooling

Liquid Cooling at NERSC:

Good Design & Regional Climate make for Compression-Free Cooling

Norm Bourassa

August 22, 2018

Energy Exchange | Better Buildings Summit 2018

Liquid in the Rack: Liquid Cooling Your Data Center

What is NERSC?

- NERSC is the National Energy Research Scientific Computing Center
- Founded in 1974
- Focused on open science
- Located at Lawrence Berkeley National Laboratory
- Operated for the U.S. Department of Energy (DOE) Office of Science

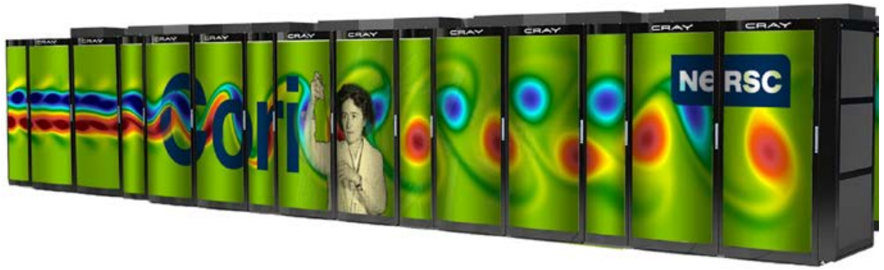

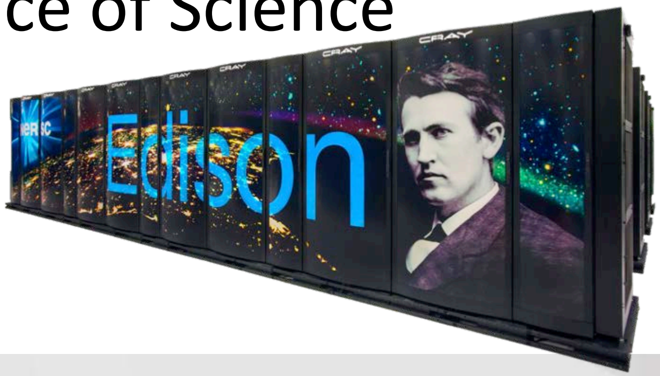

NERSC Provides HPC and Data Resources for DOE Office of Science Research

U.S. DEPARTMENT OF ENERGY | Office of Science

Largest funder of physical science research in U.S.

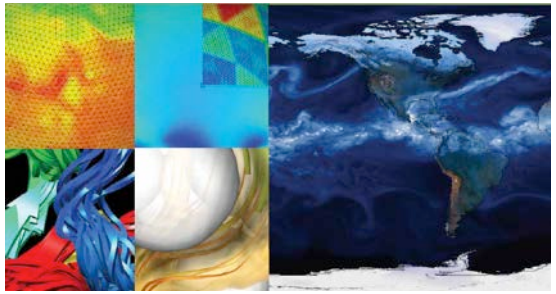

Biology, Environment

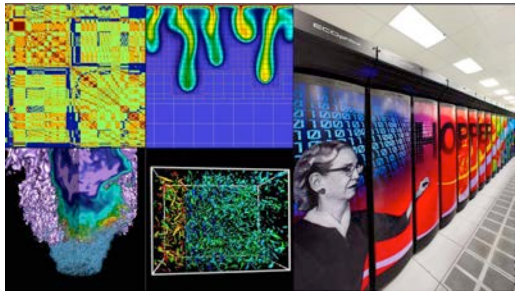

Computing

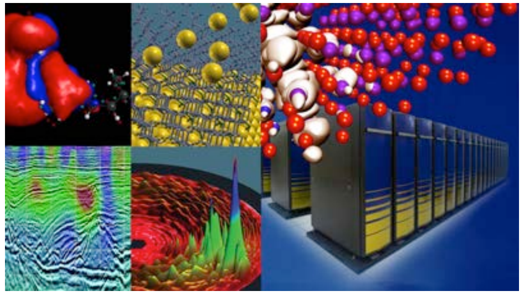

Materials, Chemistry, Geophysics

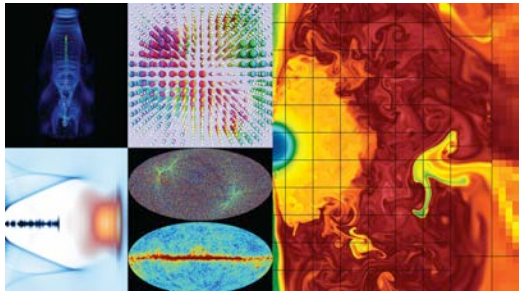

Particle Physics, Astrophysics

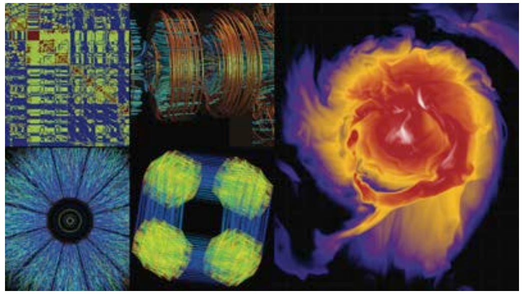

Nuclear Physics

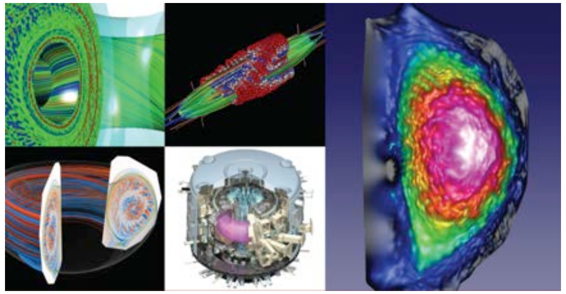

Fusion Energy, Plasma Physics

~7000 Users world wide

Moved from Oakland Site in 2015

- 11 MW power (6+5MW)
- 19,000 ft^2, 250 lbs/ft^2
- Chilled Air + Water
- **PUE ~ 1.3**
- Extensive environmental and energy-usage monitoring
- ESnet @ 10Gb/s & 100Gb/s
- Control room staffed 24 × 365
- 6.5 miles from LBNL main campus

Shyh Wang Hall (B59) completed June 2015

- Four story 150,000 GSF
  - Two 20ksf office floors, 300 offices
  - 20k -> 28ksf HPC floor
  - 12.5MW expandable to 20+MW
  - Current demand is ~7MW
- Energy Efficient
  - Year-round compressor free air and water cooling
  - LEED Gold
  - **PUE < 1.1**
- Extensively Instrumented
  - Substations, panels, PDUs, UPS
  - Cray Systems SEDC
  - One-wire Temp & RH sensors
  - BMS through ALC/BACNET
  - Indoor Particle counters
  - Weather station

Water Cooling

- Cooling towers
  - (3+1) x 3.375MW cooling towers
    - Expansion room to Qty. 7 total
  - Cooling towers operate every day
- 74°F water to racks and AHUs
  - 65°F water for 92% of the year
  - 65-70°F water for 5-6% of the year
  - 70-74°F water for 2-3% of the year
- Water Consumption
  - At 7.5MW 18M gallons/yr
  - At 17MW, estimated to be 40M gallons/yr

Climate provides the ability to generate cool air and water all year round

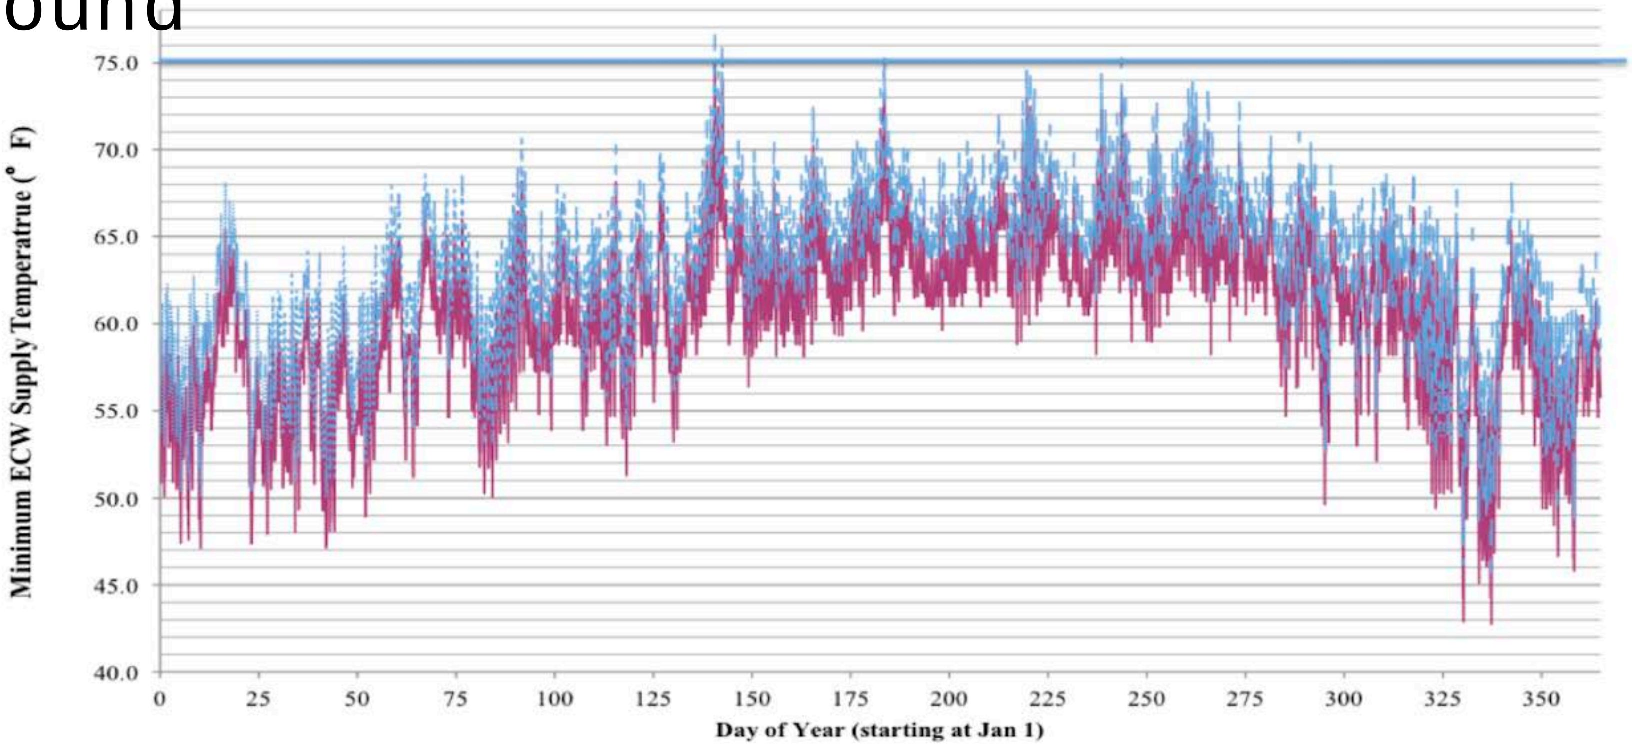

Supplied water temp (hours per year)	<70F	70-71F	71-72F	72-73F	73-74F	74-75F	>75F
One Cell	8258	261	101	64	45	19	12
Two Cell	8636	63	32	19	7	2	1

Liquid Cooling Performance Baseline

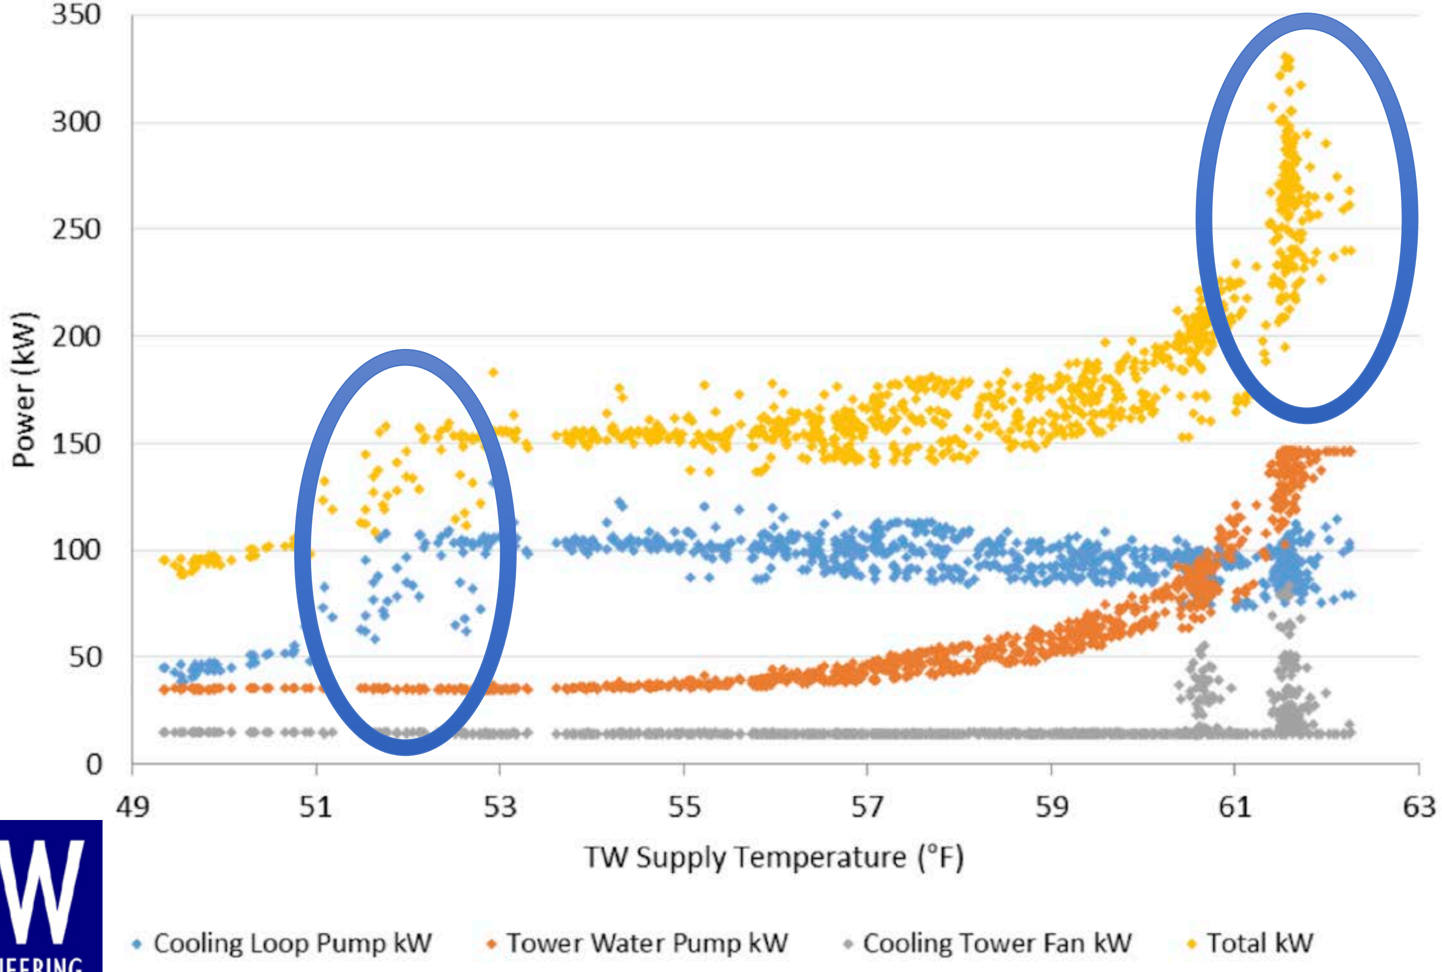

kW ENGINEERING

Energy Exchange

Better Buildings SUMMIT

Liquid Cooling Performance Balanced

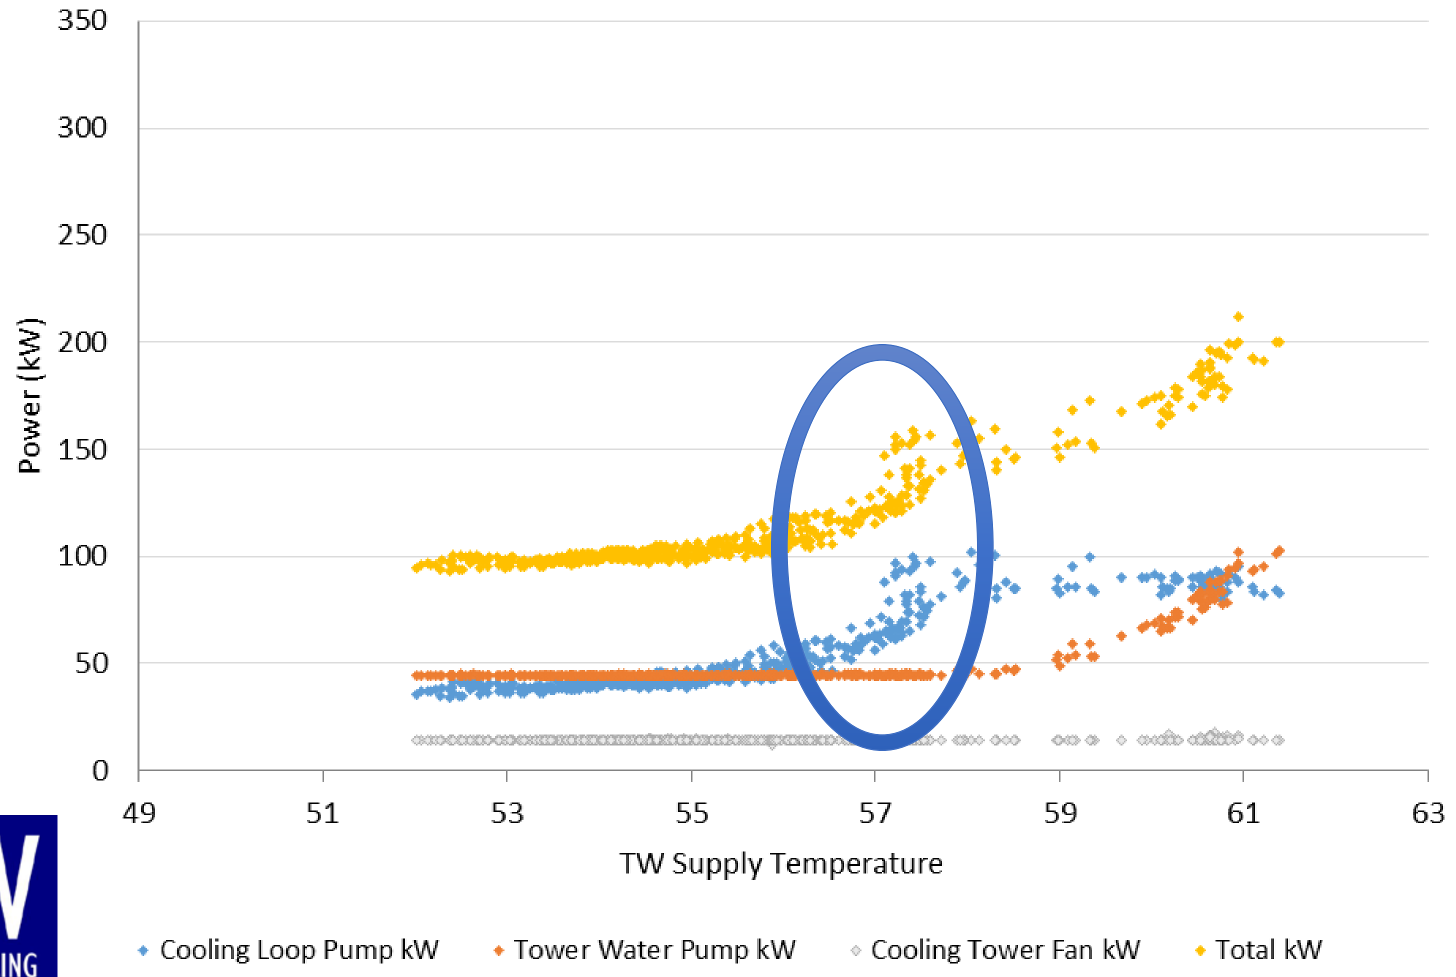

Cray XC Series - Liquid Cooled

- Edison and Cori are XC Series Liquid Cooled Systems
  - Essentially an IT fan powered “Rear-Door Coil” system
  - Each compute cabinet has a cooling coil extracting its generated heat from the air-stream before the next compute cabinet in row
    - Edison 30 compute cabinets (8/row)
    - Cori 64 compute cabinets (12/row)

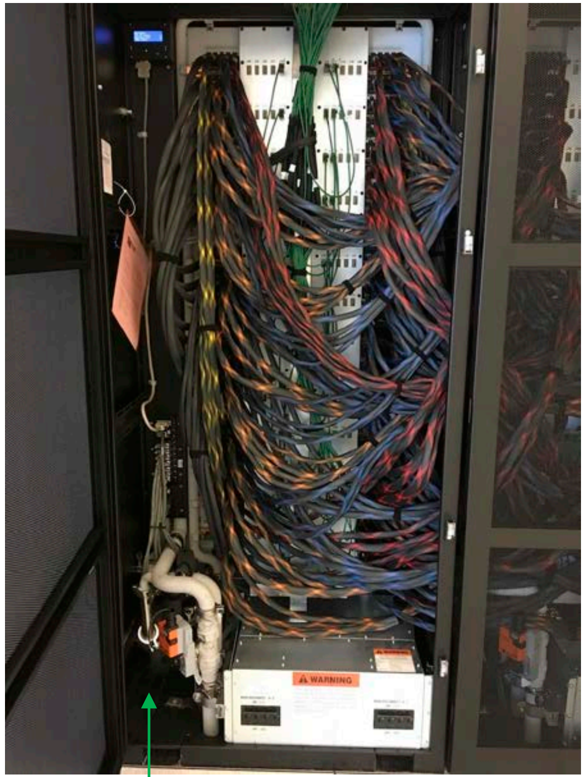

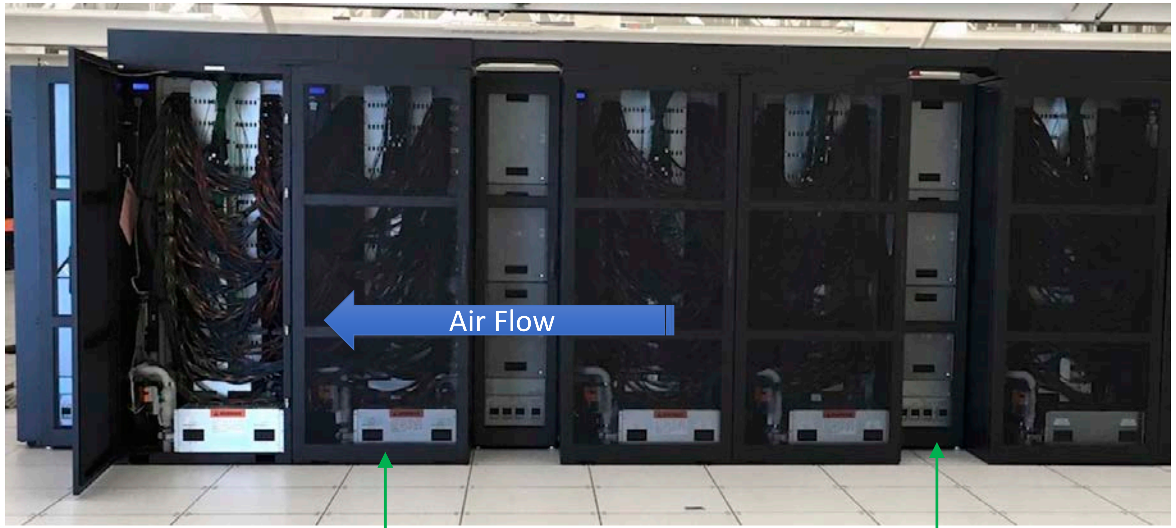

Controls to Save IT Fan Energy

- Blower Cabinets (6 Fans per)
  - Variable from 2500 to 4000 rpm
    - Dynamic Speed option saves energy sensing on HPC load with the processor temp.
  - Cooling water controlled to exiting air temp. (Typ. 69 – 71°F)
    - Can move cooling load more onto liquid when water cooling conditions favorable

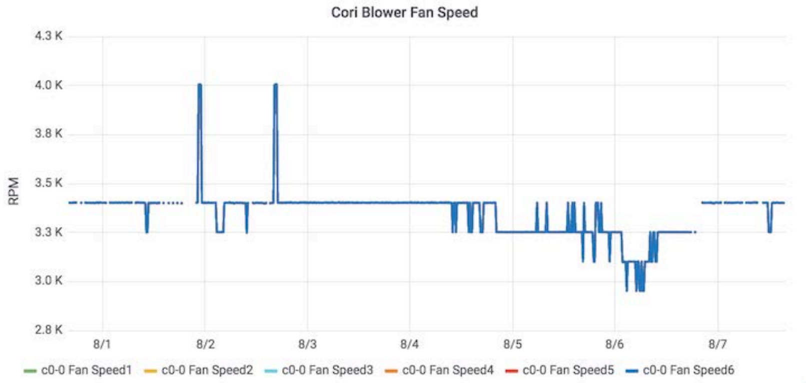

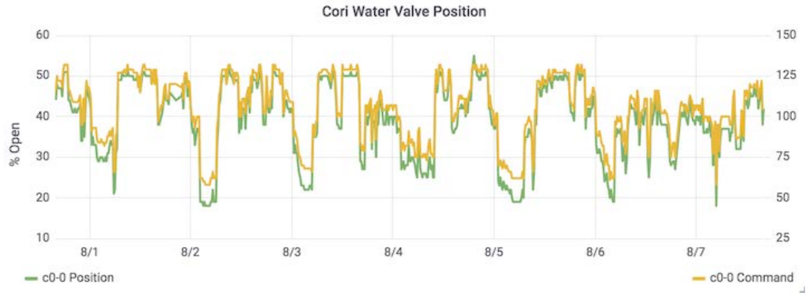

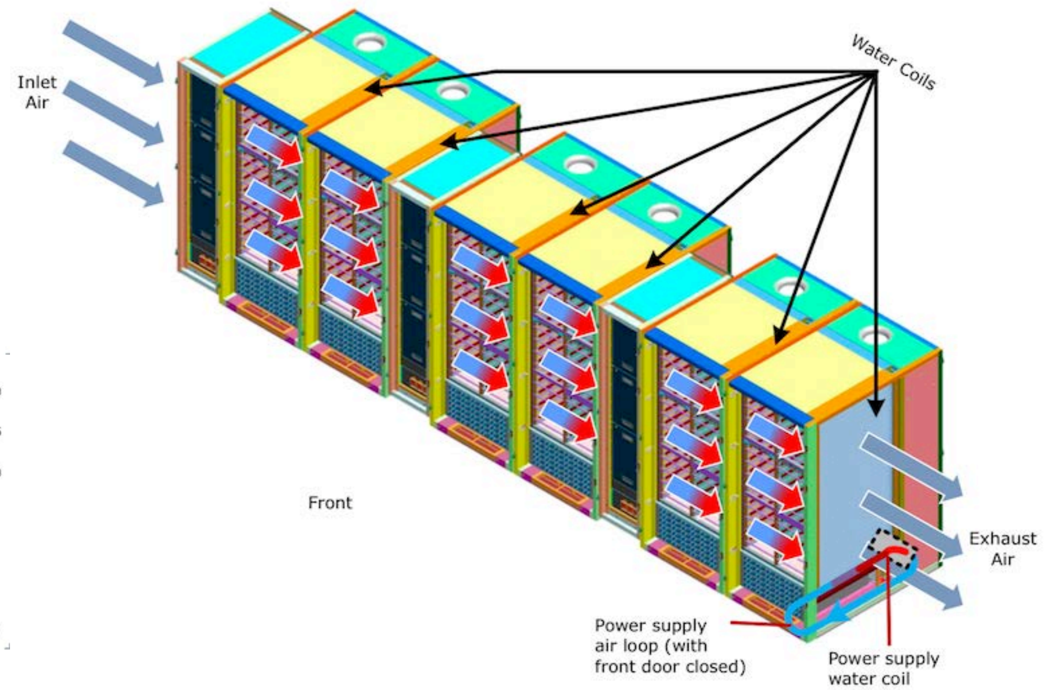

Record High Temperatures Test NERSC

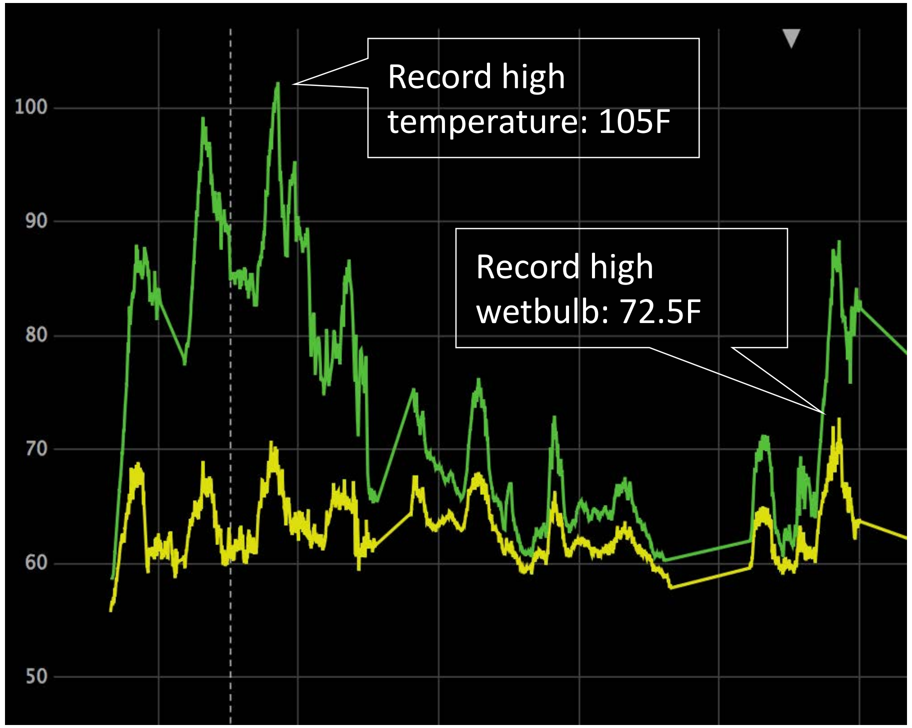

- Cooling towers alarmed on low water
- Cray's worked but required careful monitoring
  - Compute job curtailments were initiated as a precaution

Record Temps - Watching the Crays

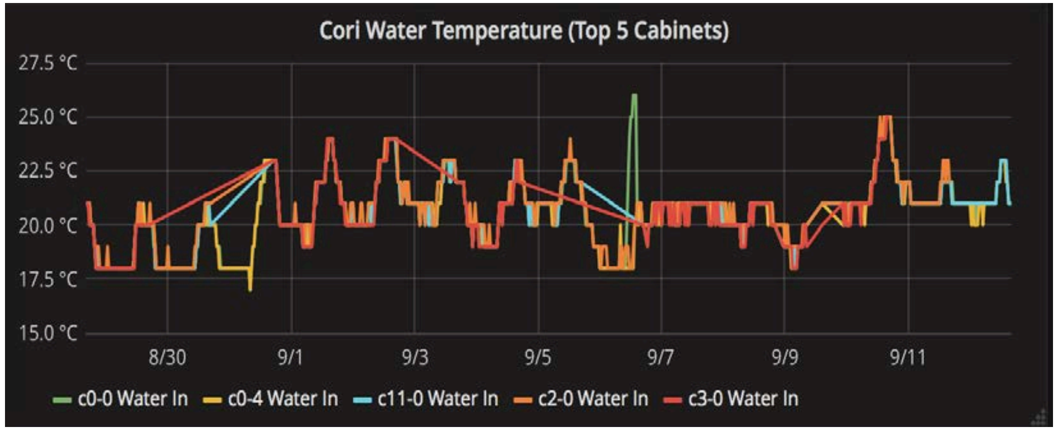

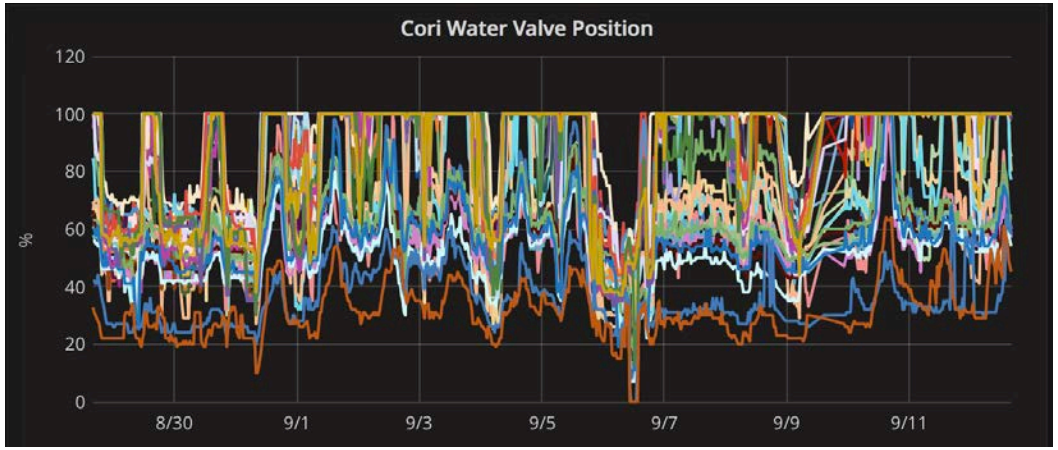

- Supply water temperature reached 78ºF (above design target of 74ºF)
- Cray SEDC provides extensive internal performance data
  - Only some cabinets had to open valves all the way
  - Almost no thermal throttling
- The Crays worked!

Cooling Tower Water Level Upgrade

- Accurate radar level sensor transmitters recently installed
  - Mid July, 2018

The Future – Moving more towards Liquid

- NERSC-9 Project Underway
  - 10MW power service upgrade (to 22.5MW building capacity)
  - Will be entirely liquid cooled
  - 3 more cooling towers with plant infrastructure to be installed
    - Also looking into new alternatives like thermosyphon
  - System installation forecast to begin Summer 2019, fully up in 2020

Lessons & Conclusions

- NERSC trend is even more towards liquid cooling
  - Hopper, (air cooled) and Edison (XC Series LC) retire in 2019
  - NERSC post 2019: Cori & NERSC 9 period of 2 systems, working towards a single liquid on-board cooled system
- Cooling tower water use is very high and becoming a more important savings measure focus
  - WUE calculation project pending – Current estimate is 1.2 to 1.5 liters/kWh (WUEsite)
  - Blowdown water saving measures are being implemented
  - Looking into State Point Liquid Cooling, new indirect evaporative system
- Thermosyphon is under consideration (NERSC 9)
- No space for Dry Air Coolers presently (Possibly NERSC 10)
- High resolution instrumentation of liquid systems provide beneficial information for operation improvements, but not without work overhead
  - Data analysis can be difficult – Good visualization tools a necessity
  - Keeping the system running requires considerable admin attention

National Energy Research Scientific Computing Center

Norm Bourassa, njbourassa@lbl.gov

NERSC Optimization Project - Ongoing

- Saved ~2,150,000 kWh & 150,000 gal water over the last year
  - Estimated 0.045 PUE improvement
  - 1,800,000 kWh Non-IT savings
  - 750,000 kWh IT savings

		Energy Savings (kWh)		Water Savings
	Measure Title	Estimated	Verified	Gallons
1	Optimize Cooling Tower Fan and Pump Controls	-	310,000	(300,000)
2	Optimize Closed Loop Pump Control	-	1,050,000	110,000
3	Optimize AHU SAT and Flow Control	300,000	-	-
4	Reset Cooling Water Supply Temperature and Optimize CRAY Controls	400,000	400,000	220,000
5	Install Firmware to Enable ESS Mode for UPSs	-	350,000	120,000
6	Replace Bypass Valves	-	40,000	-
7	Cold Aisle Partial Containment	100,000	-	-
	Total	**800,000**	**2,150,000**	**150,000**

Still more savings on the table: ~600,000 kWh/year

- supply air fan optimization
- Cray Dynamic Fan Speed Control
  - CW plant inter-tie
- Bypass valves for pumping
- HPC Water HX heat reclaim

Liquid in the Rack: Liquid Cooling Your Data Center

NREL ESIF Data Center

Otto Van Geet, PE - NREL

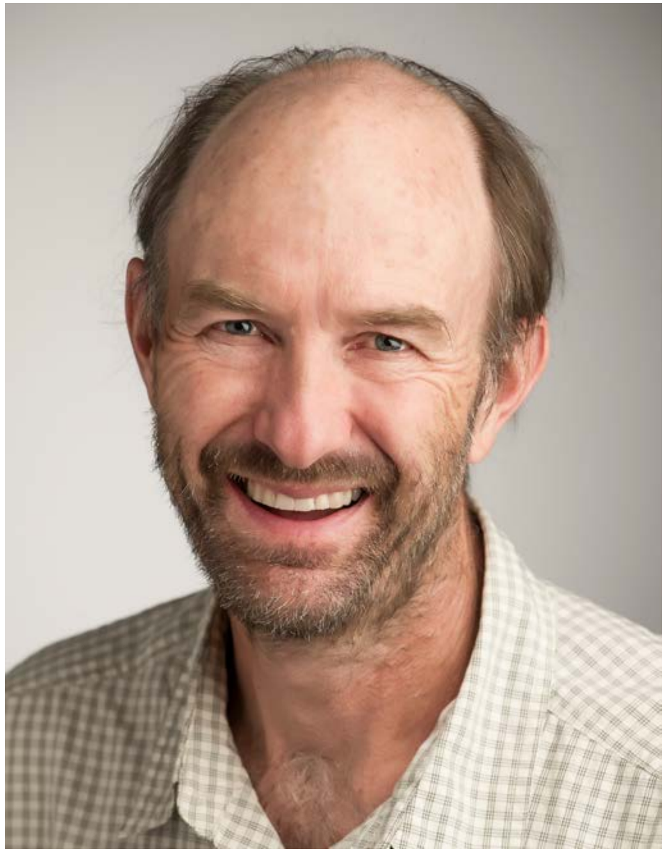

Otto Van Geet, PE

Principal Engineer, NREL

Otto.vangeet@nrel.gov

NREL's Dual Computing Mission

- Provide HPC and related systems expertise to advance NREL's mission, *and* push the leading edge for data center sustainability
- Demonstrate leadership in liquid cooling, waste heat capture, and re-use
- Holistic "chips-to-bricks" approaches to data center efficiency
- Showcase data center at NREL's Energy Systems Integration Facility (ESIF)
- **Critical Topics Include:**
  - Liquid cooling and energy efficiency
  - Water efficiency

Planning for a New Data Center

- Started planning for new data center in 2006
- Based on HPC industry/technology trends, committed to direct liquid cooling
- **Holistic approach:** integrate racks into the data center, data center into the facility, the facility into the NREL campus
- Capture and use data center waste heat: office and lab space (now) and export to campus (future)
- Incorporate high-power density racks—more than 60 kW per rack
- Implement liquid cooling at the rack, no mechanical chillers
- Use chilled beam for office/lab space heating. Low-grade waste heat use.
- Considered two critical temperatures:
  - Information technology (IT) cooling supply—could produce 24ºC (75ºF) on hottest day of the year, ASHRAE "W2" class
  - IT return water—required 35ºC (95ºF) to heat the facility on the coldest day of the year

Build the World's Most Energy Efficient Data Center

NREL Data Center

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*)

Utilize the bytes and the BTUs!

Data Center Features

- Direct, component-level liquid cooling, 24ºC (75ºF) cooling water supply
- 35-40ºC (95-104ºF) return water (waste heat) is 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 60- to 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

Metrics

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

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

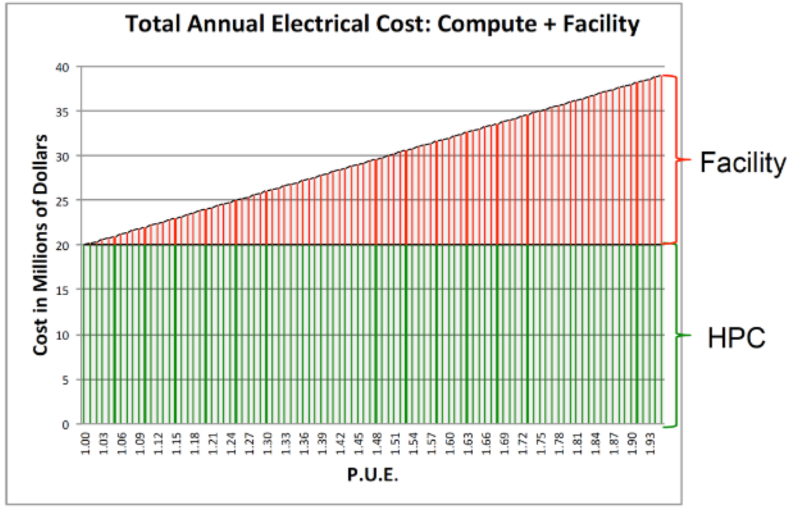

Assume ~20MW HPC system & $1M per MW year utility cost.

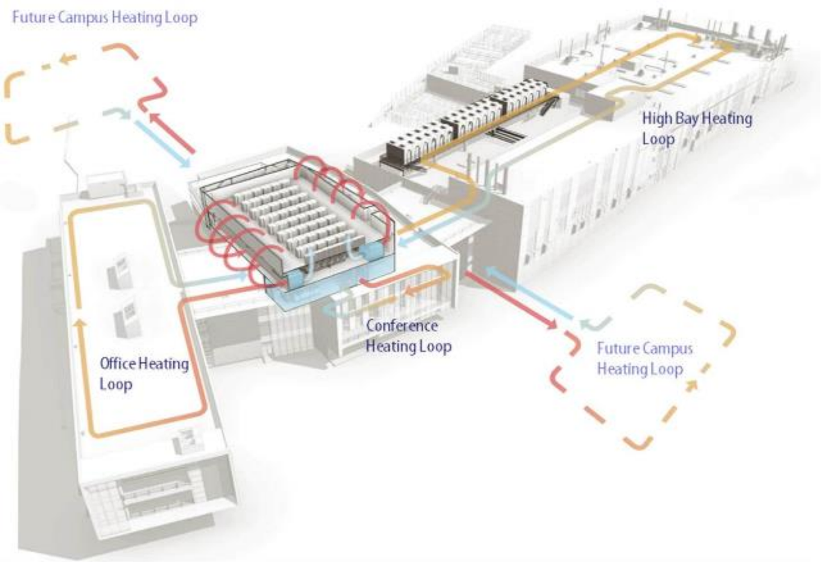

Metrics

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

the units of WUE are liters/kWh

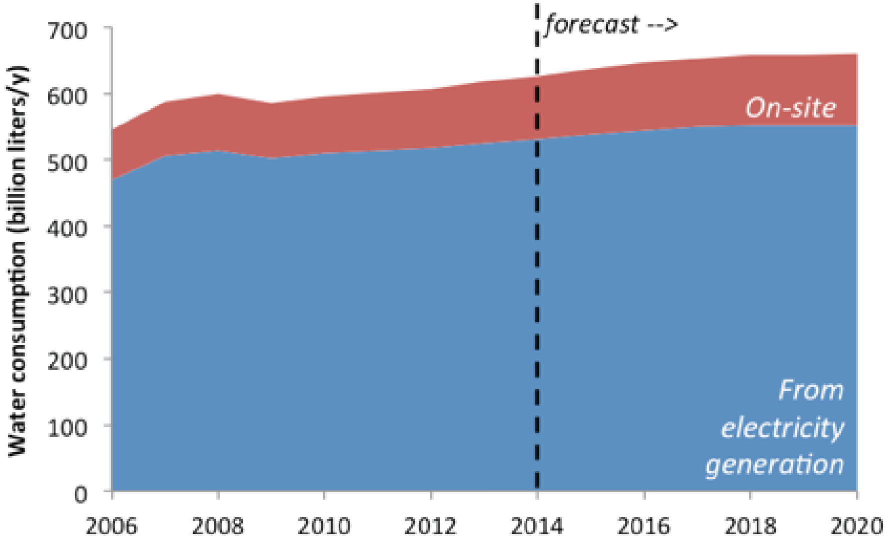

$$WUE_{SOURCE} = \frac{\text{"Annual Site Water Usage"} + \text{"Annual Source Energy Water Usage"}}{\text{"IT energy"}}$$

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

where EWIF is energy water intensity factor

Air-Cooled to Liquid-Cooled Racks

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

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

Liquid-Cooled Server Options

Peregrine HPC Tray

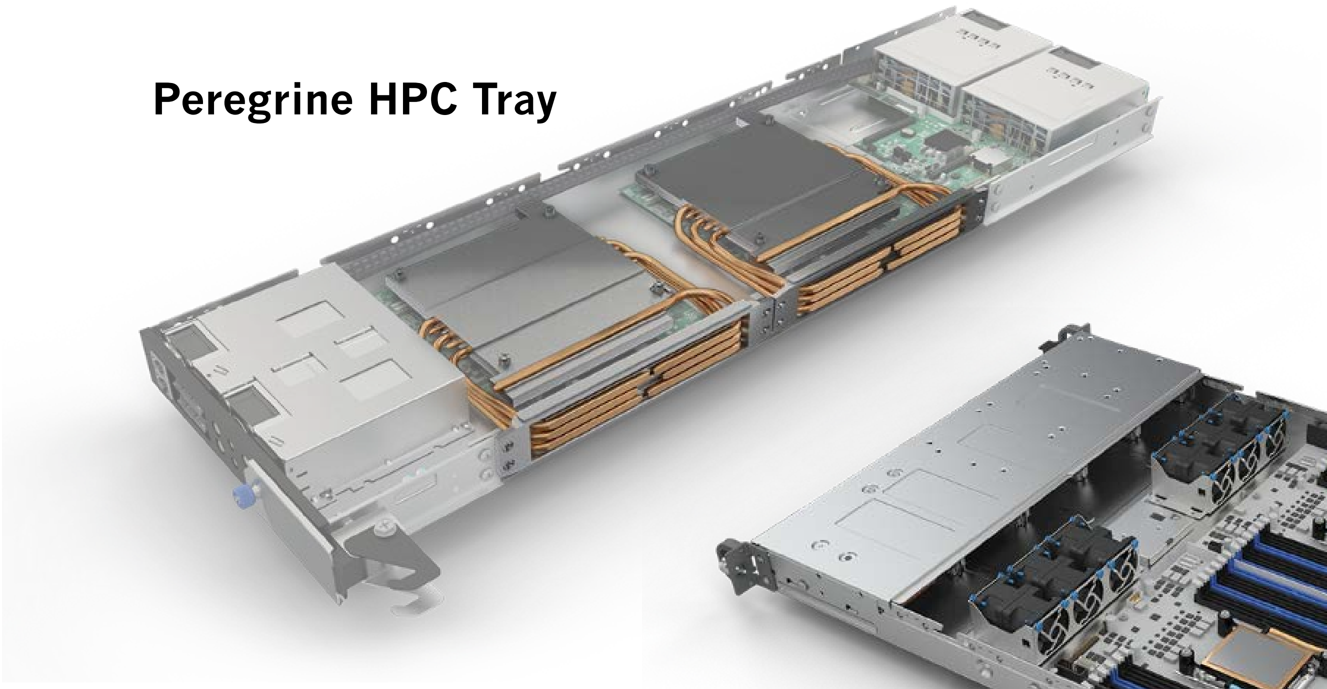

Cold Plates

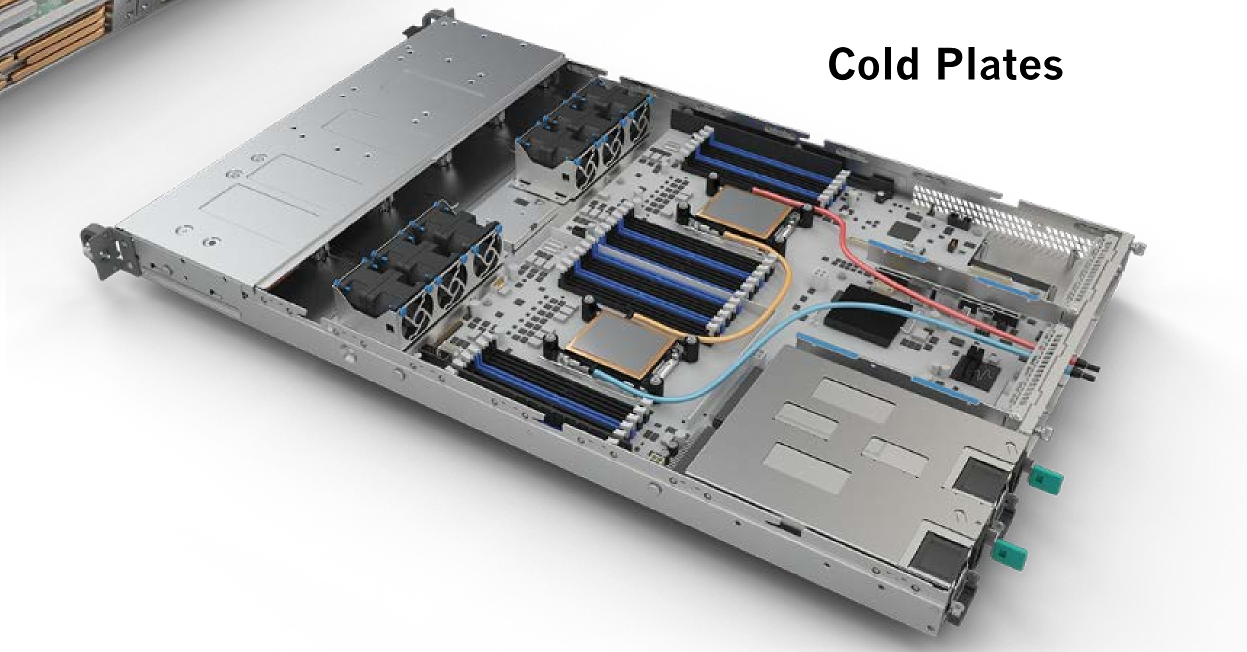

Fanless Liquid-Cooled Server Options

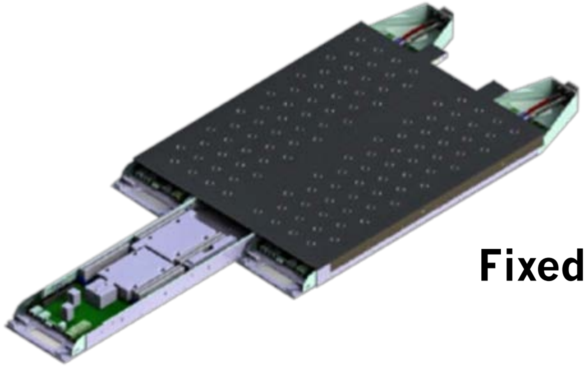

Fixed Cold Plate

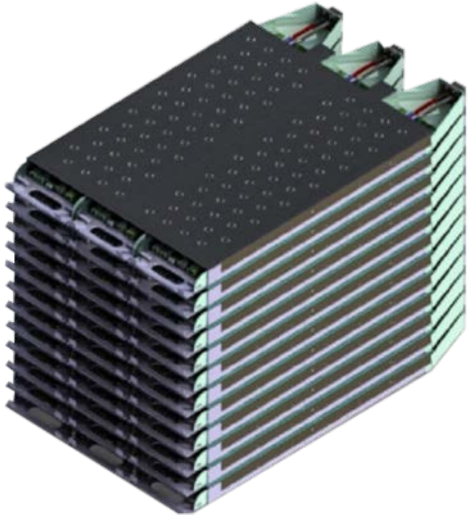

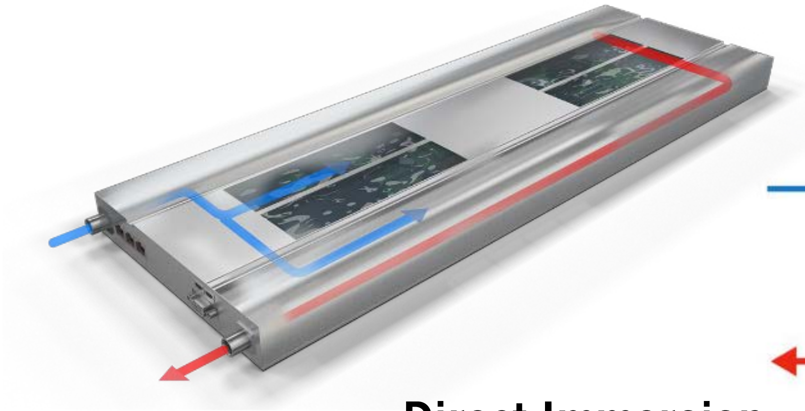

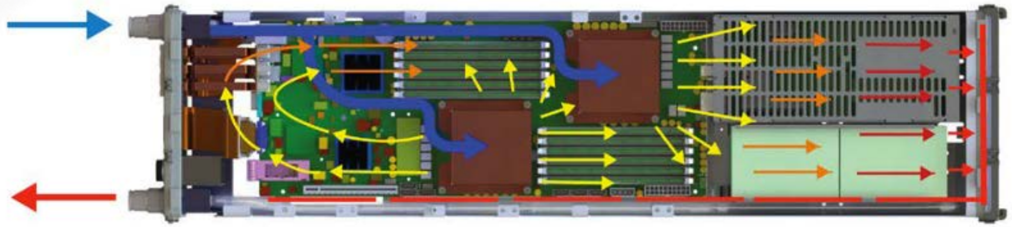

Direct Immersion

Data Center Water Distribution

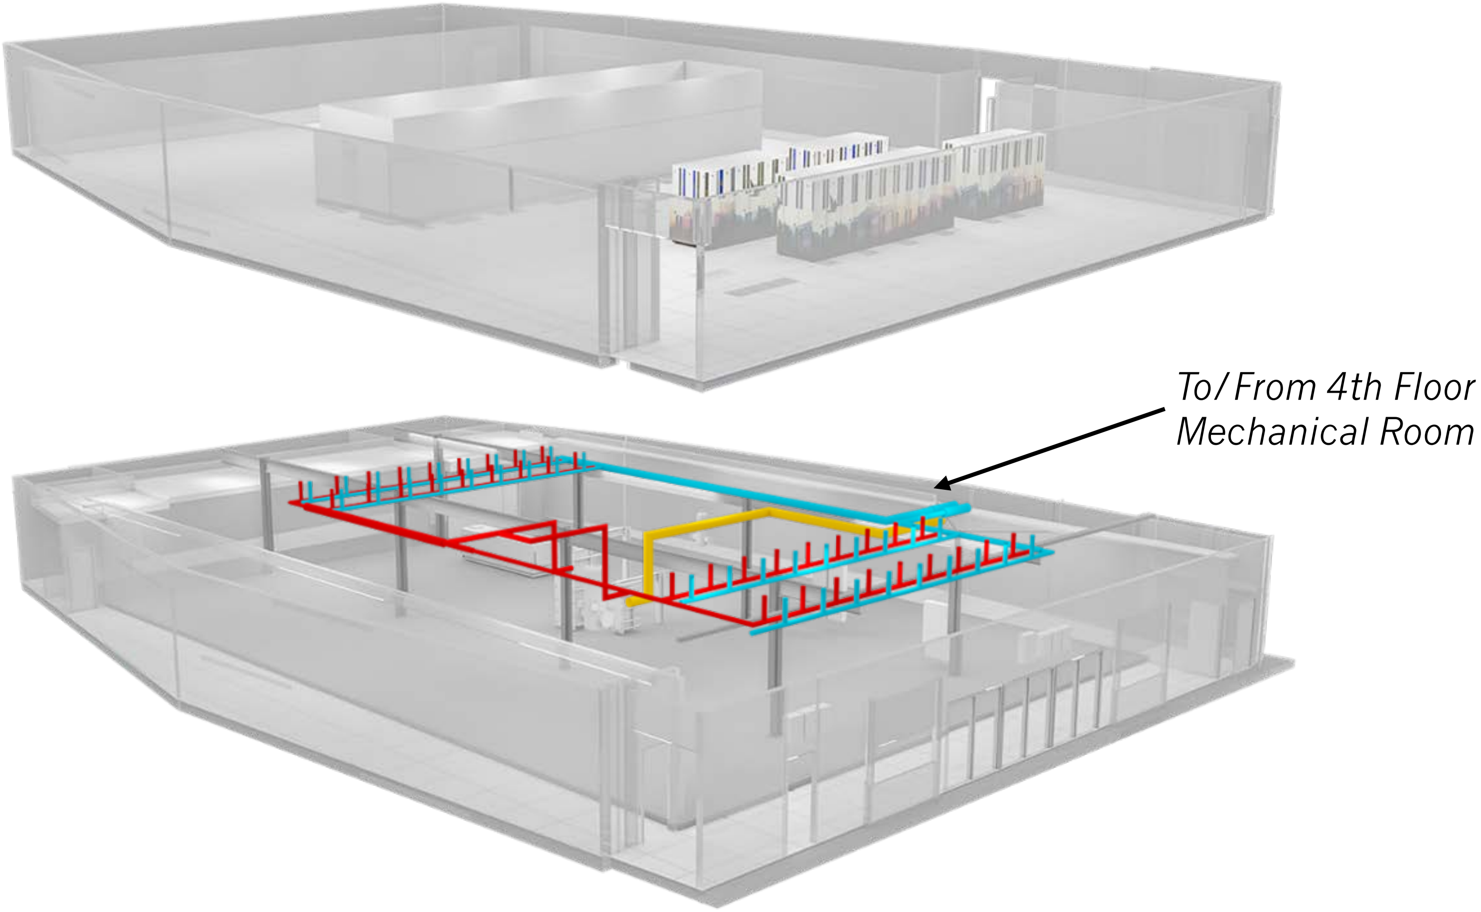

Liquid Cooling—Considerations

- Liquid cooling essential at high-power density
- Compatible metals and water chemistry is crucial
- Cooling distribution units (CDUs)
  - Efficient heat exchangers to separate facility and server liquids
  - Flow control to manage heat return
  - System filtration (with bypass) to ensure quality
- 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

Basic Hybrid System Concept

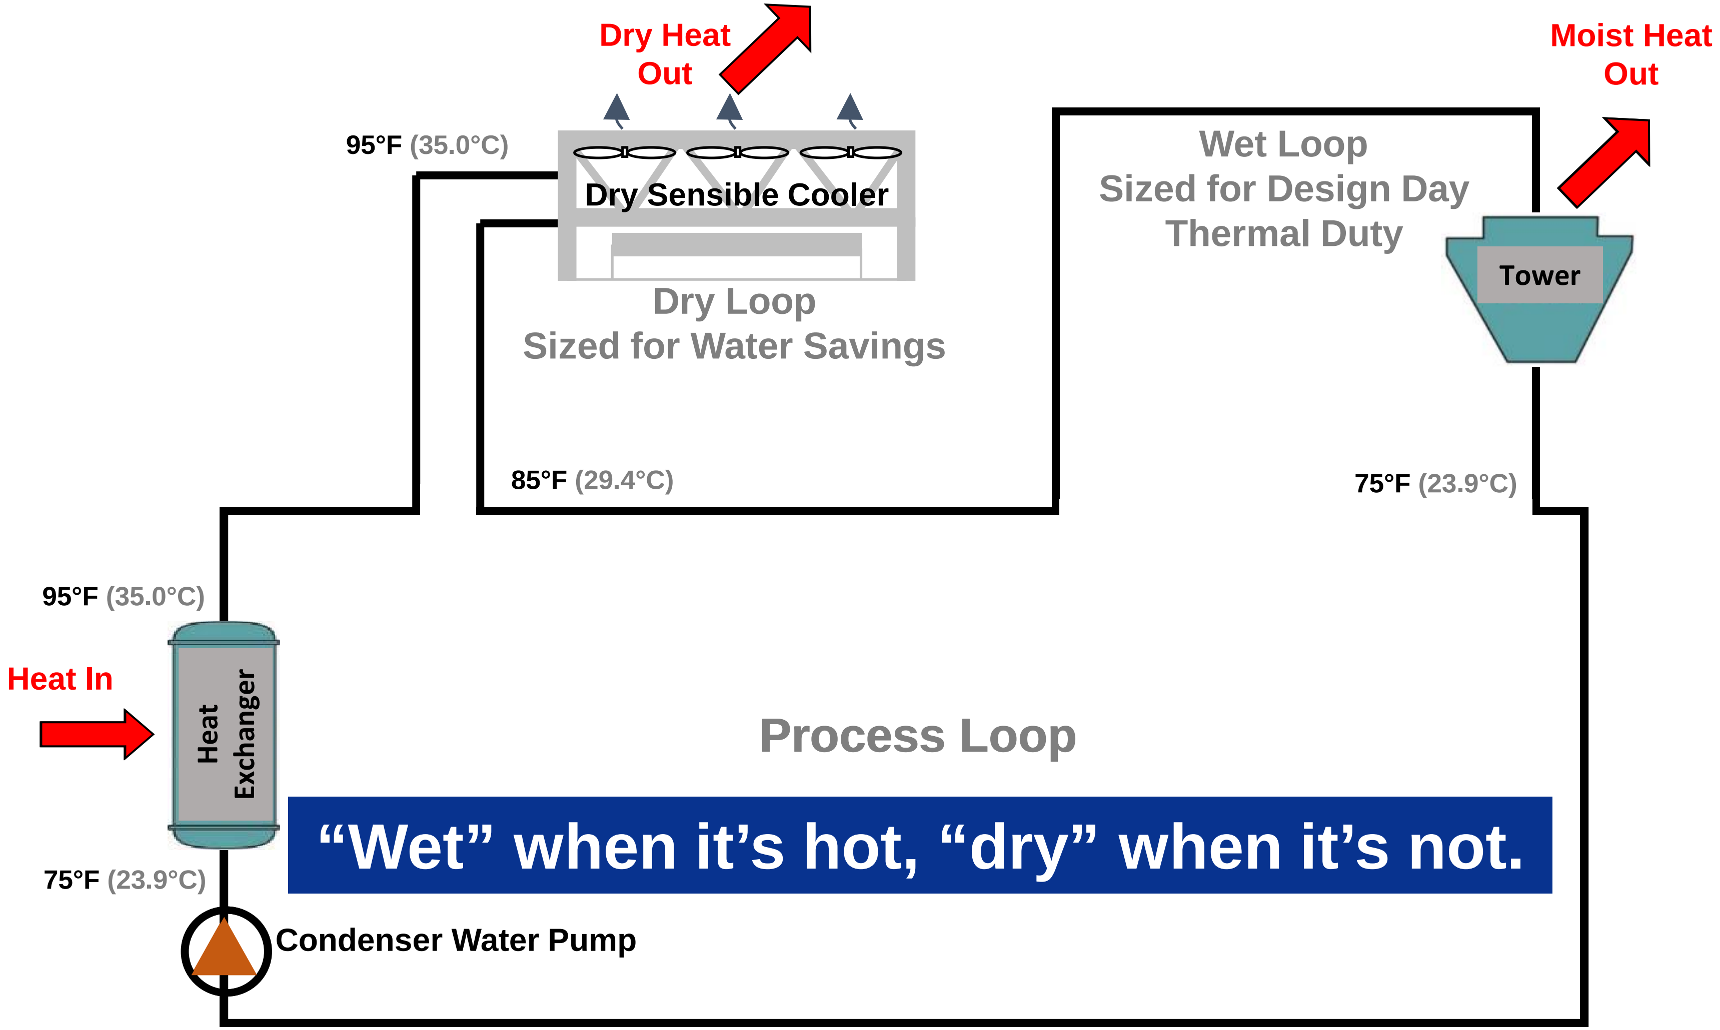

"Wet" when it's hot, "dry" when it's not.

Improved WUE—Thermosyphon

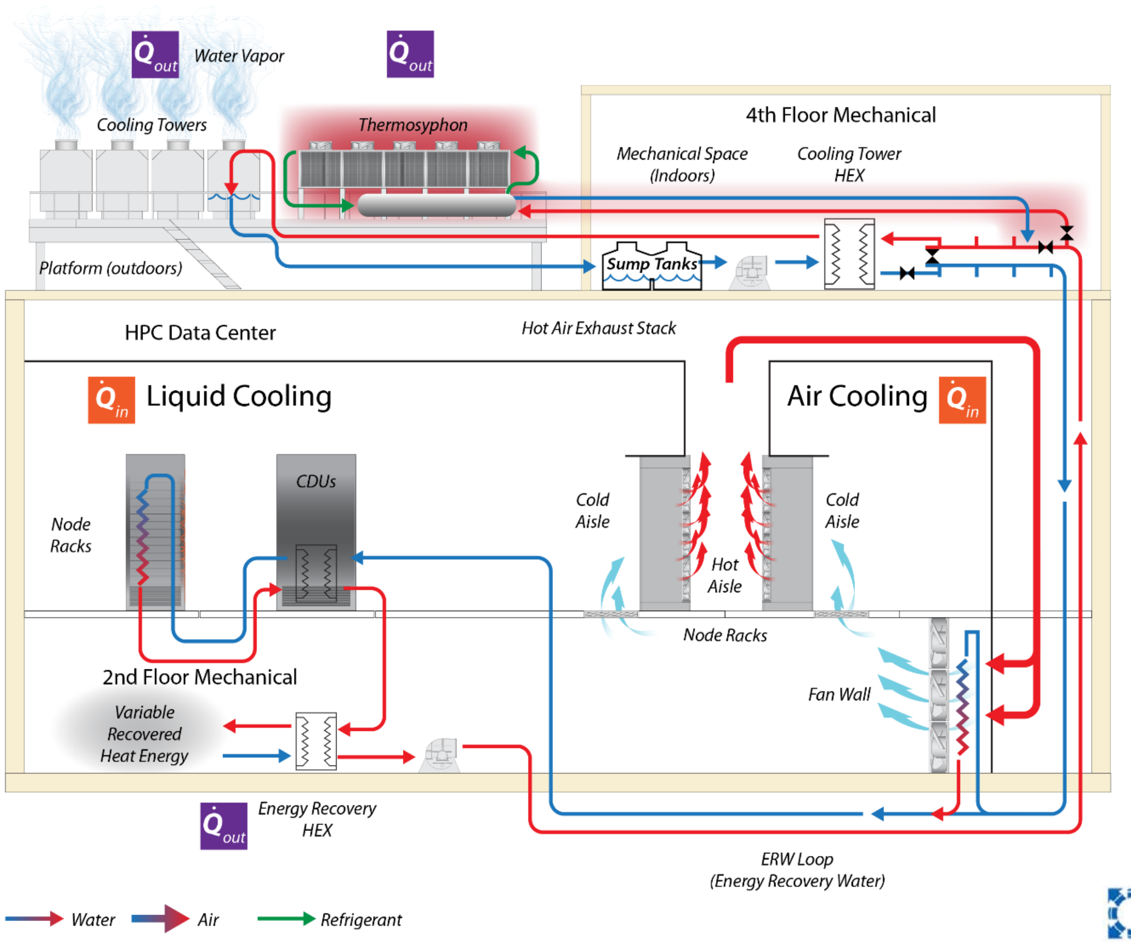

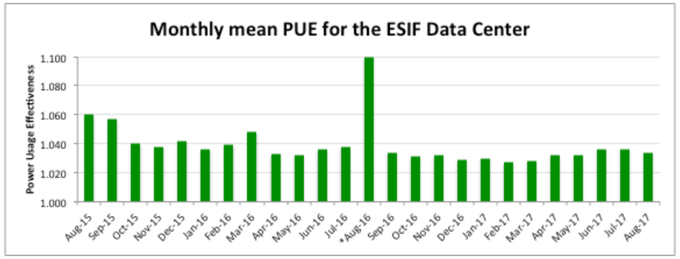

NREL NATIONAL RENEWABLE ENERGY LABORATORY
Johnson Controls
Sandia National Laboratories

ESIF Data Center Efficiency Dashboard

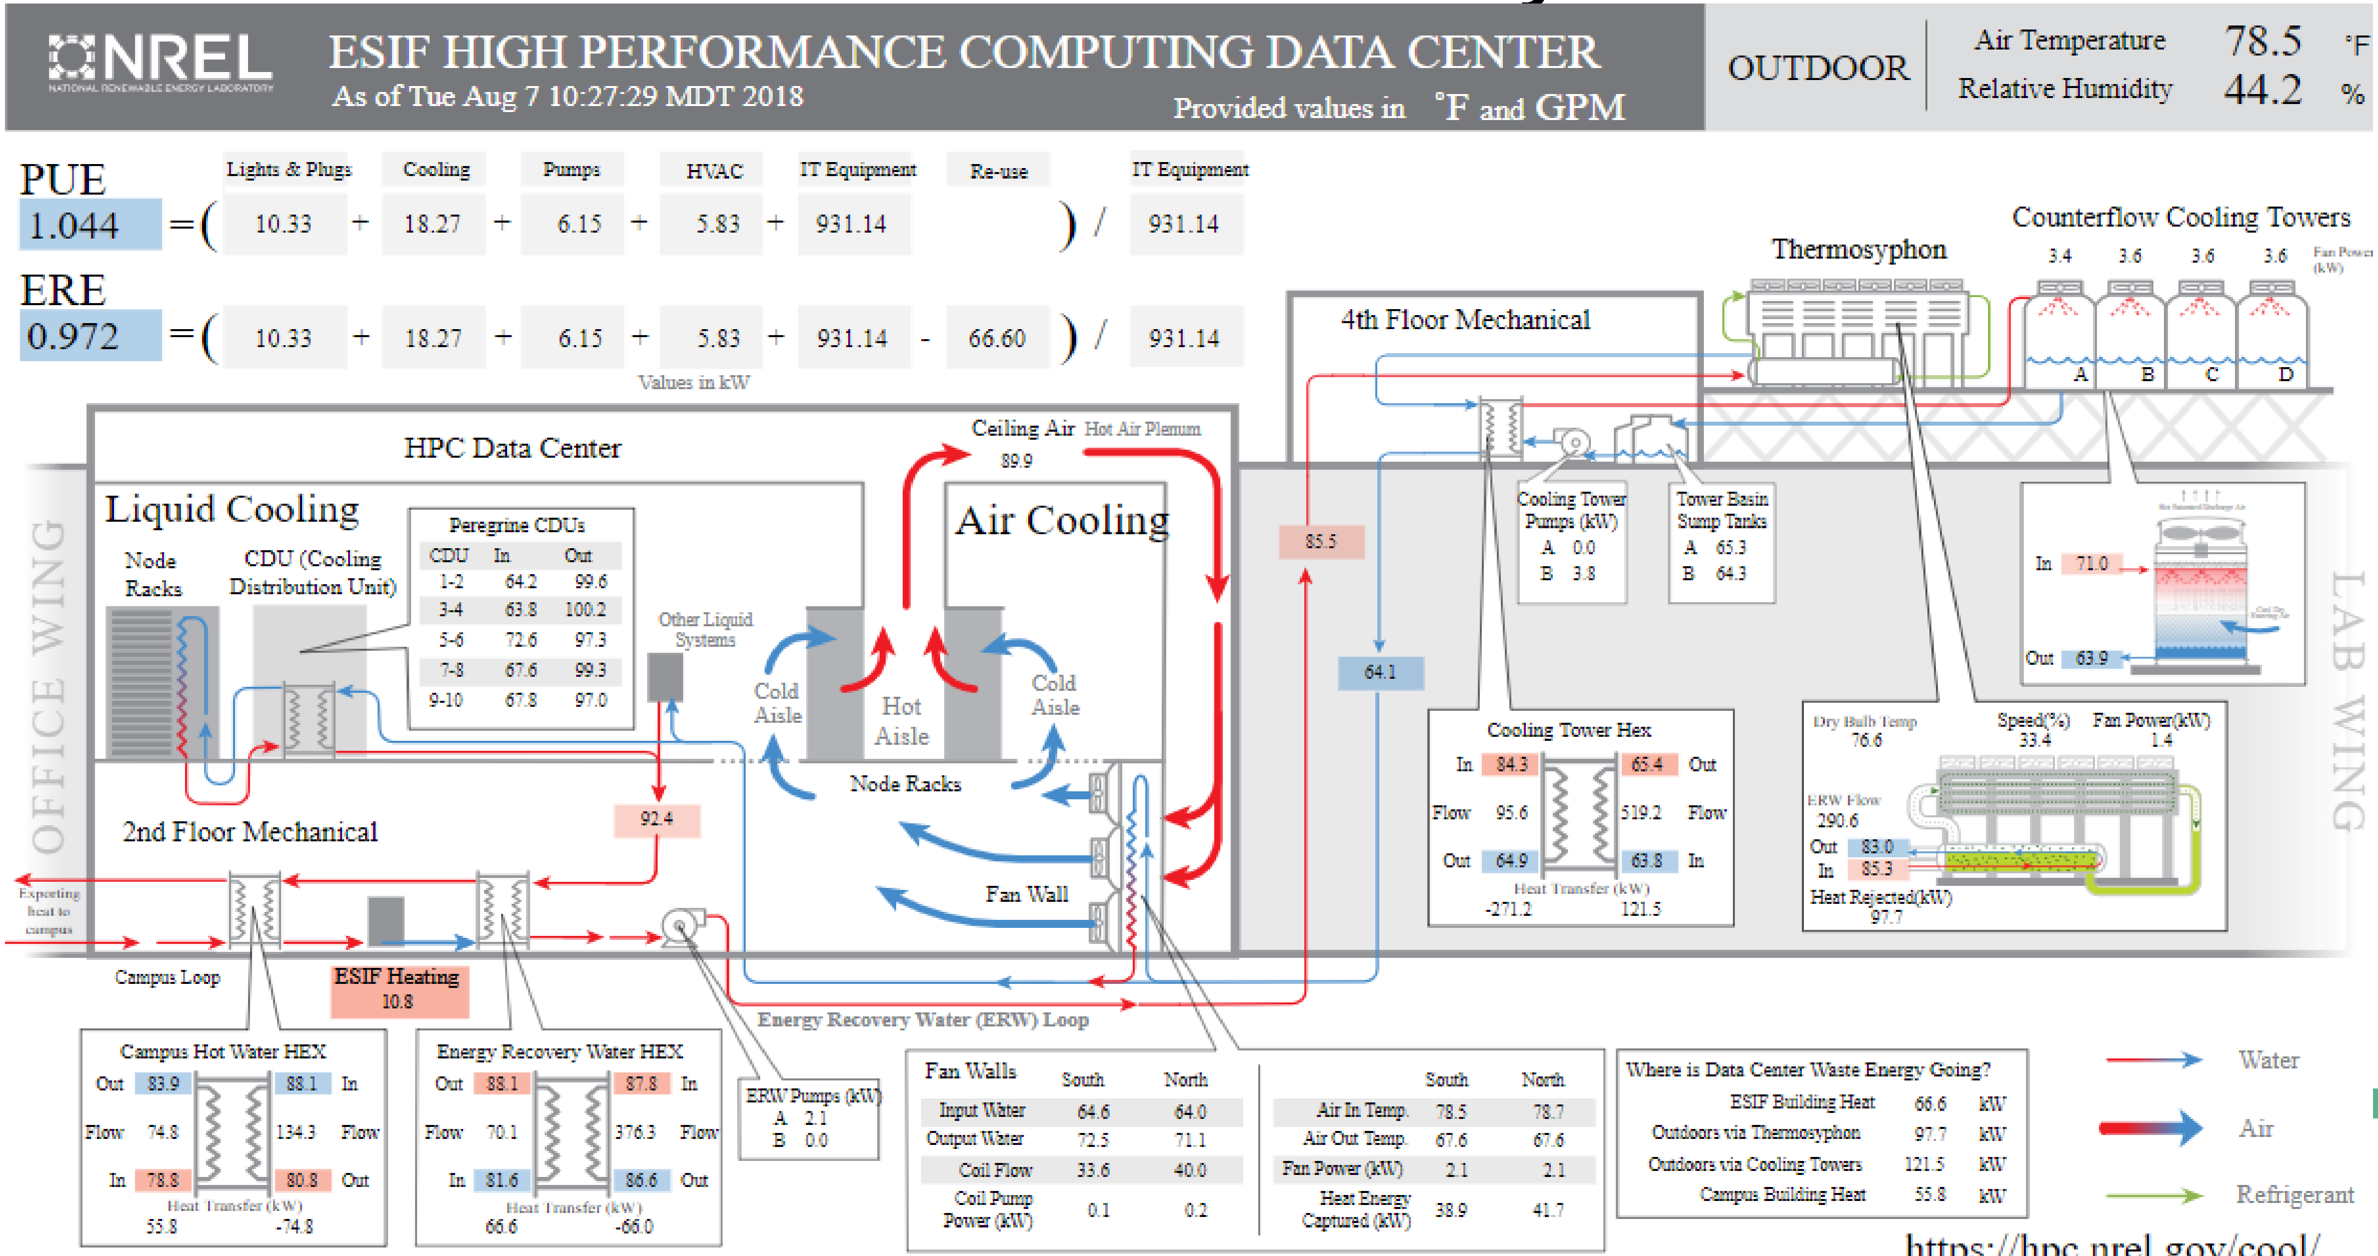

https://hpc.nrel.gov/cool/

Data Center Metrics

First year of TSC operation (9/1/2016–8/31/2017)

Hourly average IT Load = 888 kW

PUE = 1.034

ERE = 0.929

Annual Heat Rejection

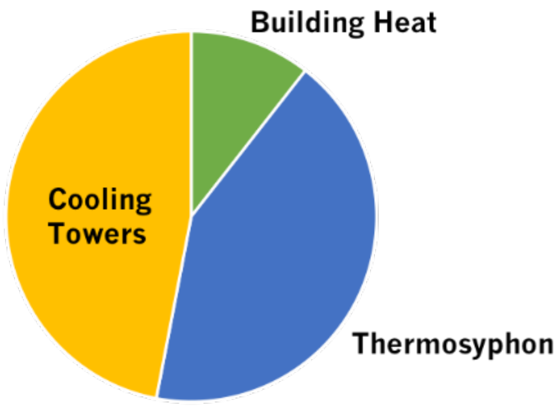

WUE = 0.7 liters/kWh

(with only cooling towers, WUE = 1.42 liters/kWh)

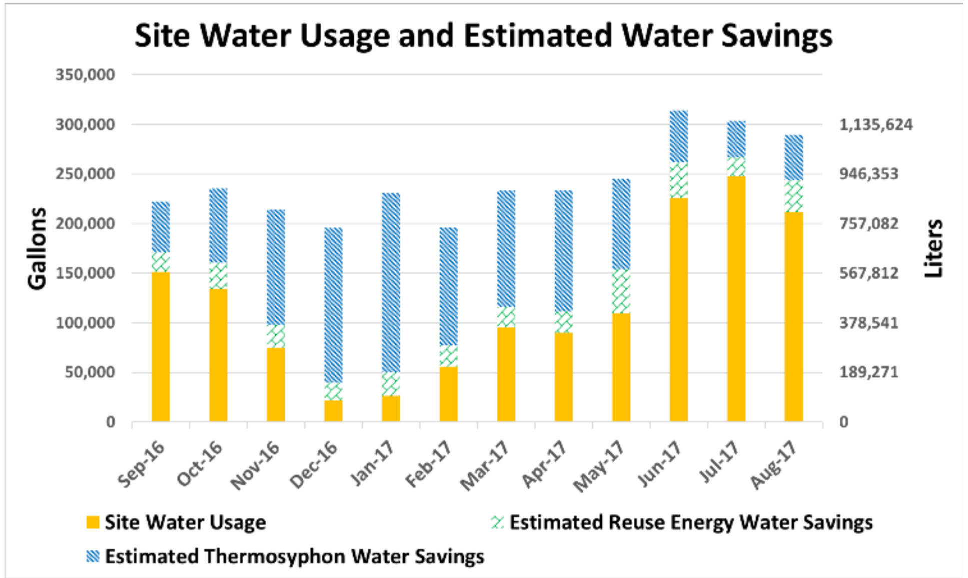

WUE_{SOURCE} = 5.4 liters/kWh
WUE_{SOURCE} = 4.9 liters/kWh if energy from 720 kW PV (10.5%) is included

using EWIF 4.542 liters/kWh for Colorado

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.