Saving Energy with Free Cooling and the Cray XC30

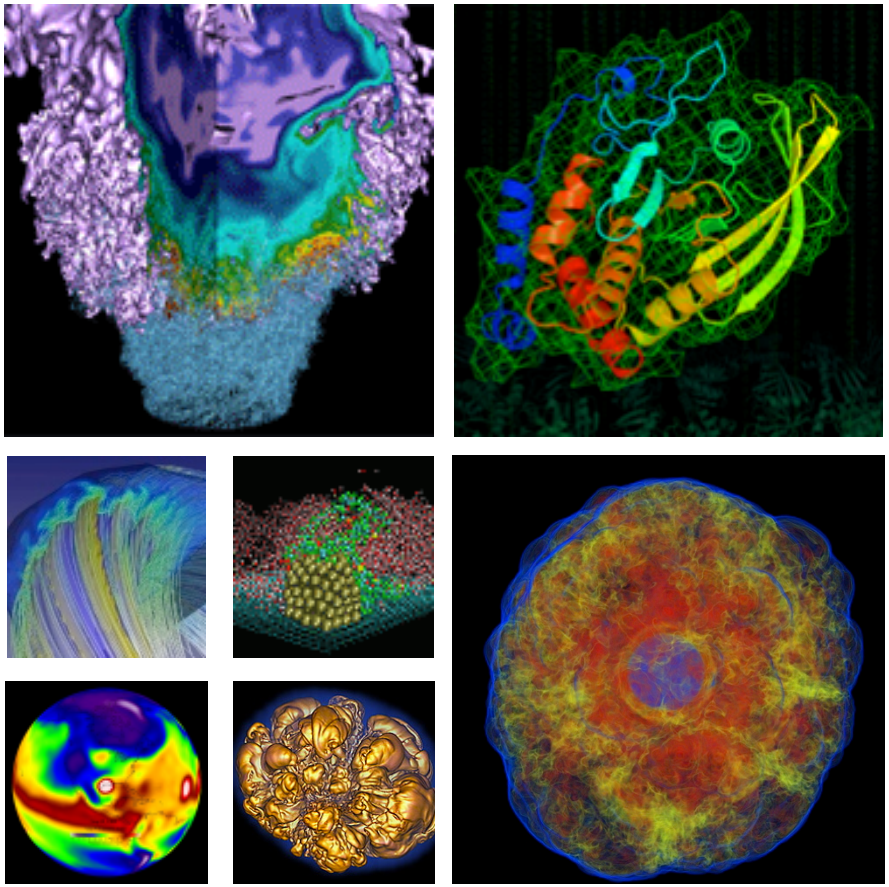

Brent Draney, Jeff Broughton,
Tina Declerck, John Hutchings

CUG 2013

What is "free" cooling?

- **Cooling without use of mechanical refrigeration**
- **Utilize environmental cooling source**
  - Lake water (CSCS)
  - Aquifer (iVEC)
  - Open air / evaporation (NERSC)
- **Not really "free"**
  - Power for pumps, fans, etc.
- **Still 2-3X less power**
- **Can be utilized full or part-time**

“The coldest winter I ever spent was a summer in San Francisco.”

Berkeley average temperatures

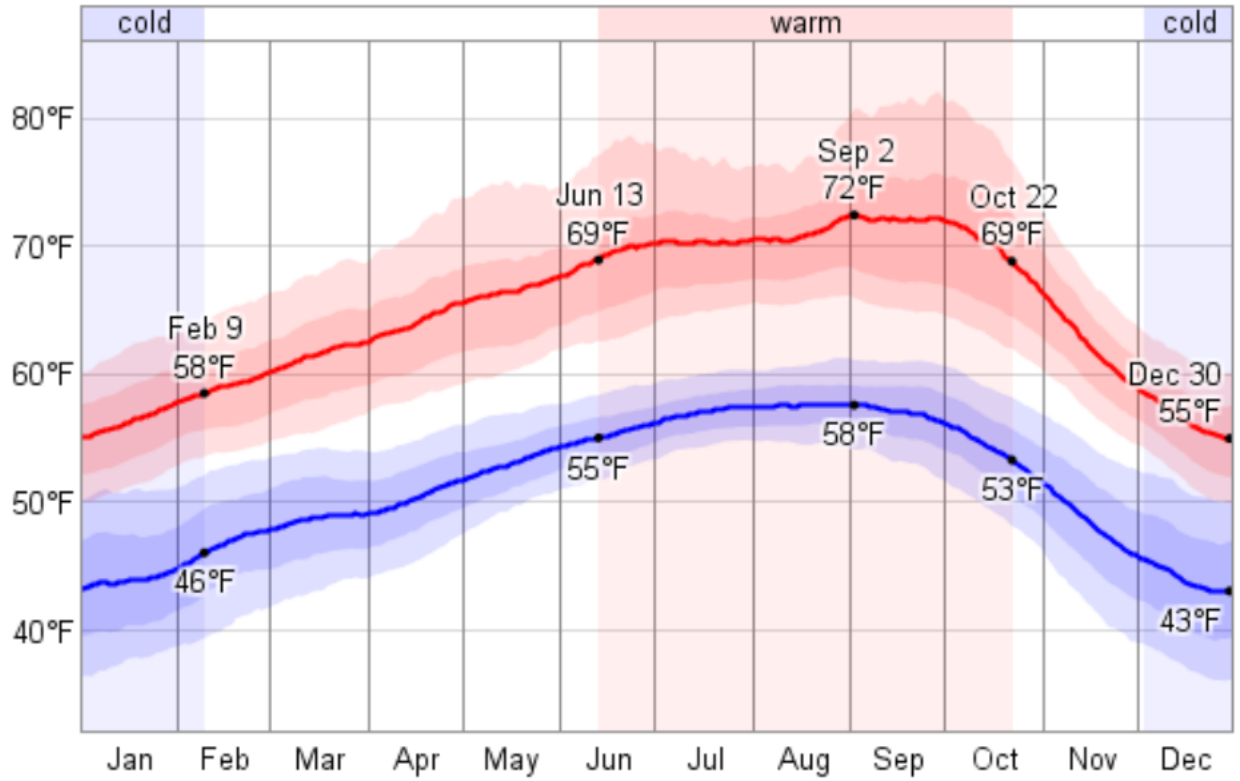

The daily average low (blue) and high (red) for Oakland/Berkeley with percentile bands (inner 25% to 75%, outer 10% to 90%). Data from weatherspark.com.

Edison (Cray XC30)

- First Cray XC30
- Intel processors
- Aries interconnect
- Water cooled
- Phase 1: 4 cabinets (installed)
- Phase 2: 24 cabinets (June)

- 5,200 compute nodes
- 333TB memory
- 6.4PB storage @ 140GB/s
- 12 login nodes, 6TB
- 1.6MW average, 2.0MW peak

Cray XC30 cooling mechanism

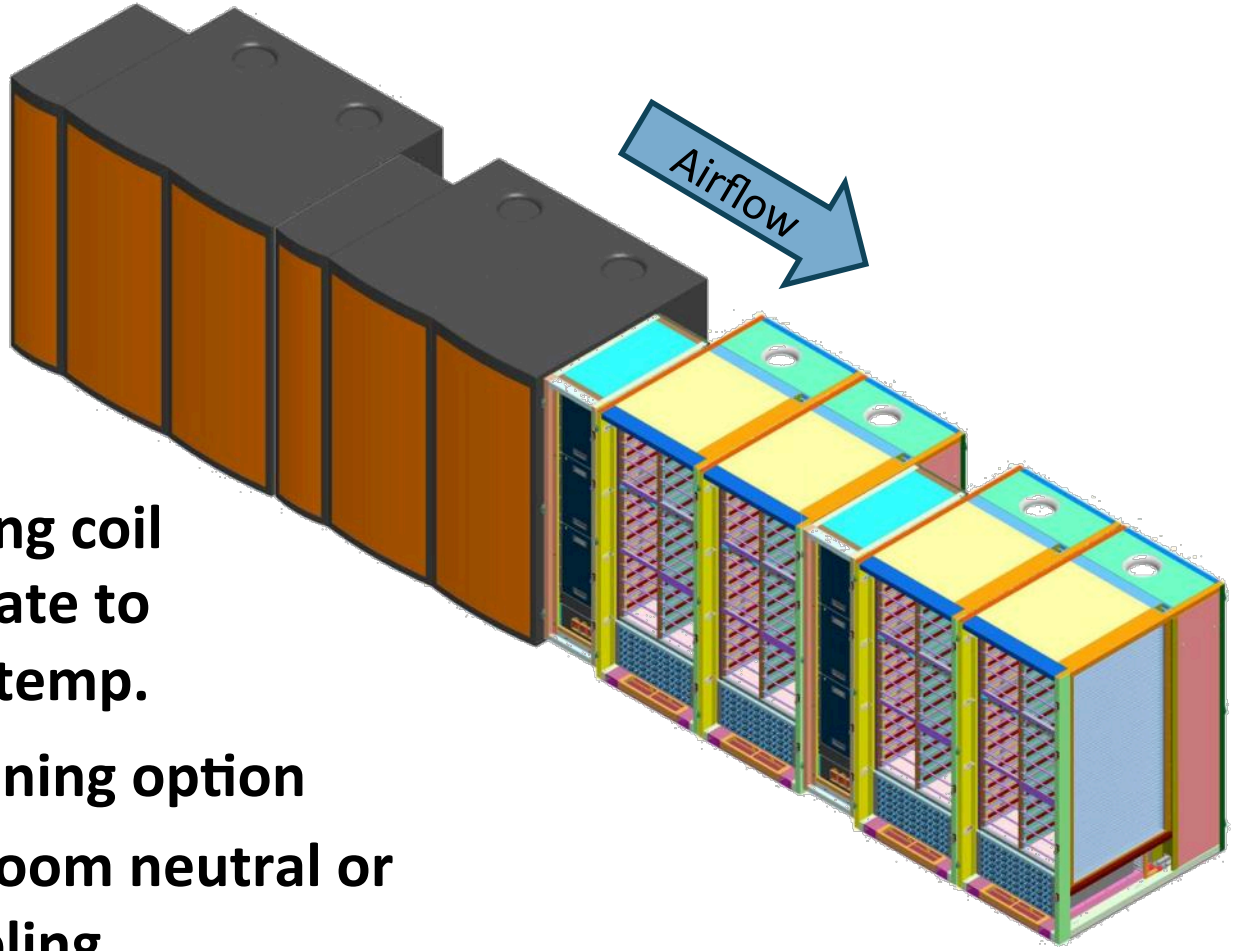

- One blower assembly for each cabinet pair (group)
- Compute rack cooling coil valve adjusts flow rate to maintain outlet air temp.
- Inlet air preconditioning option
- Exhaust air can be room neutral or require residual cooling

XC30 Cabinet Design

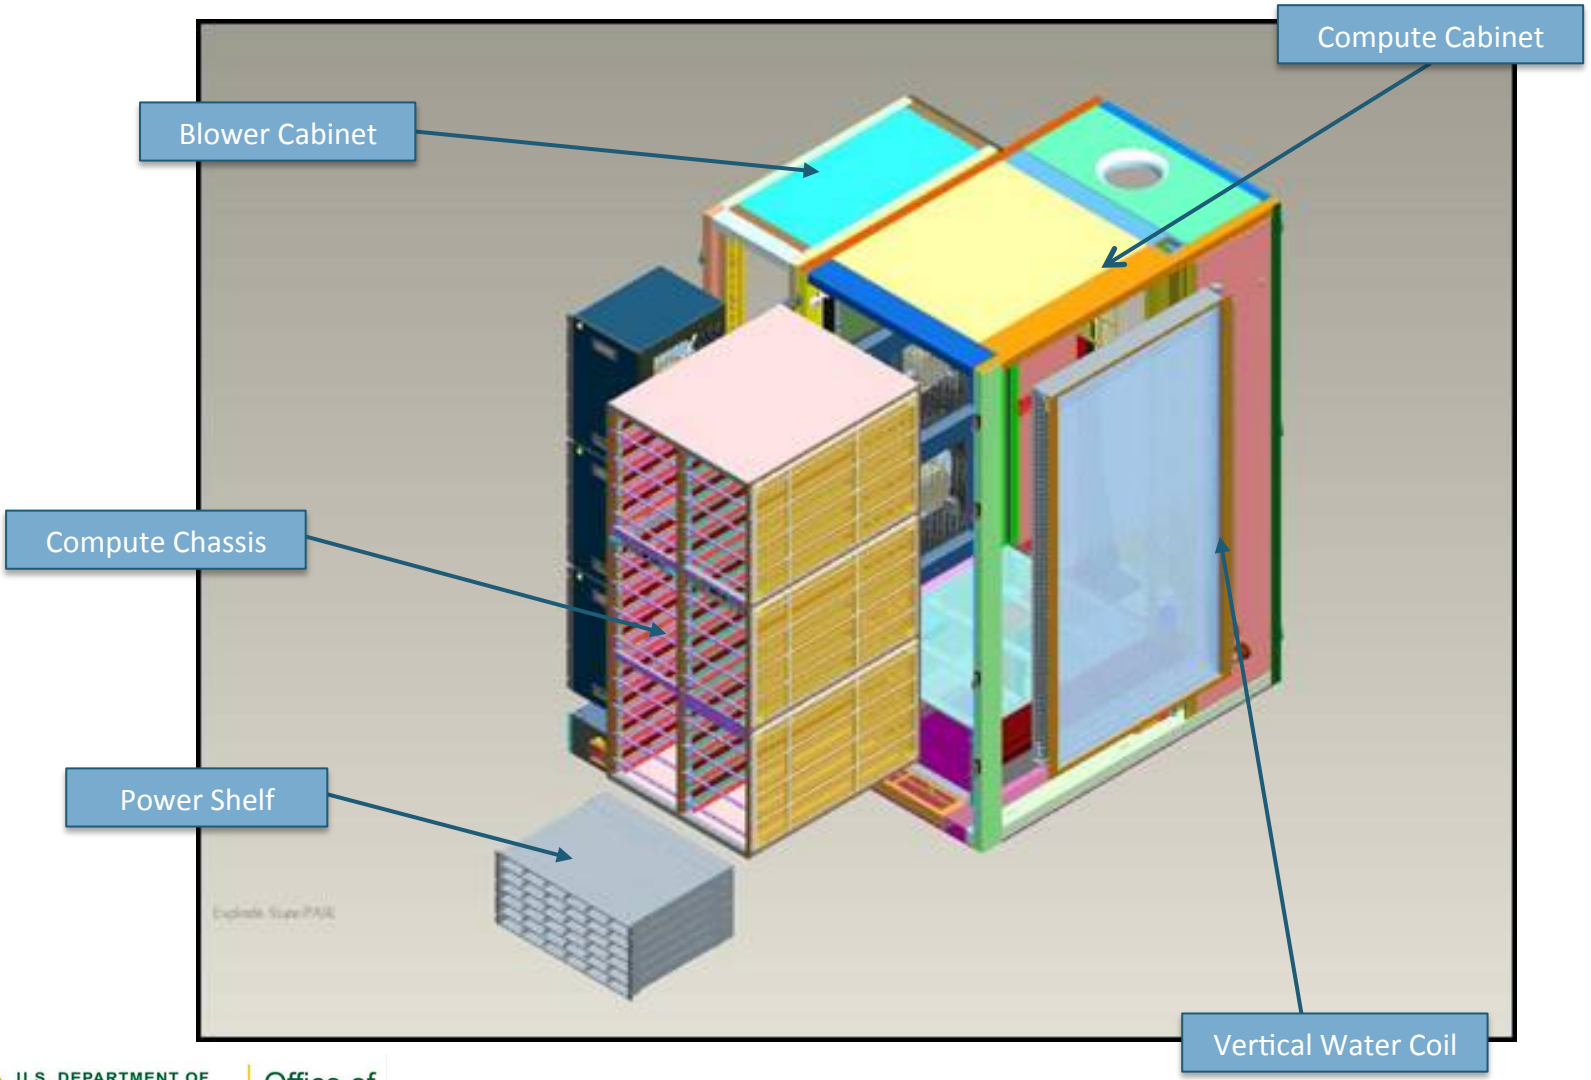

XC30 Blower Cabinet

- N+1 configurations
- Hot Swap Blower Assembly
- Low pressure and velocity of air
- Low noise (TN26-6 standard, 75 db/ cabinet)

- Heat generated by blowers is extracted by radiators in compute cabinets

Under-row Cooling Manifold

- Piping can get tight need dedicated rows for power and cooling
- 8” pipe is the largest you can get in a single row and that should work for 16 racks
- Consider feeding power and cooling from opposite ends to avoid conflicts

Parallel-Flow vs Counter-flow Heat Exchangers

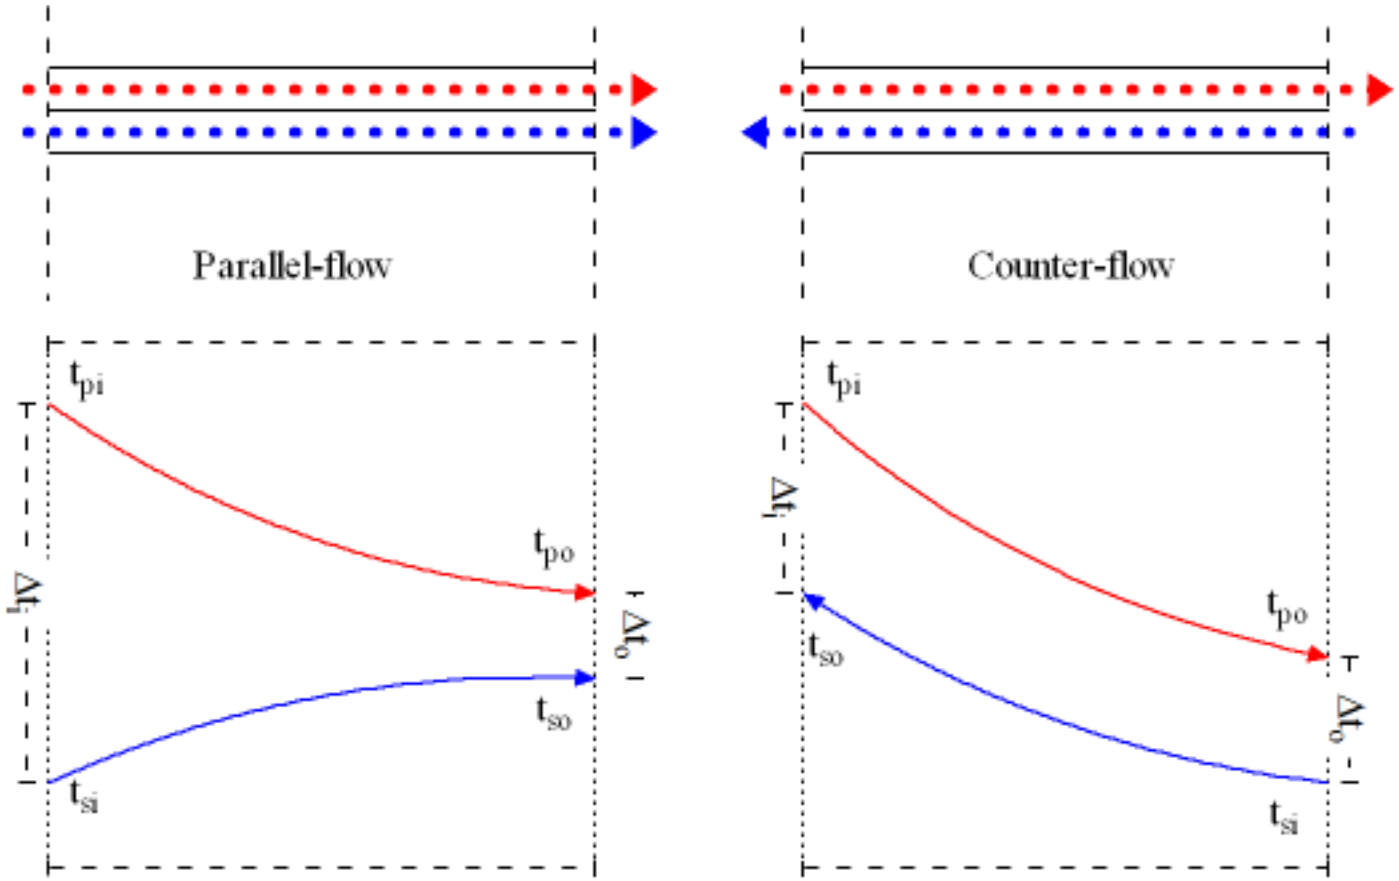

Things to consider for liquid cooling

- Water temperature
- Wet-bulb temperature
- "Approach"
- Dew-point
- Flow rate
- Differential pressure
- Processor type
- Altitude

115W, 8 Core SNB & 10 Core IVB 2.8 Ghz (0' elevation)

Water inlet C	Water outlet C	Flow US GPM	Flow m3/hr	Flow l/s	dP psi	Coil Outlet Air Temp C
10	30.5	12	2.78	0.76	3	27
11	30.5	13	3.04	0.83	3	27
12	30.5	14	3.29	0.9	3	27
13	30.5	15	3.54	0.97	4	27
14	30.5	17	3.8	1.04	4	27
15	30.5	19	4.3	1.18	4	27
16	30.5	21	4.55	1.25	5	27
17	30.5	22	5.06	1.39	5	27
18	30.5	24	5.57	1.53	5	27
19	30.5	26	6.07	1.67	6	27
20	30.5	30	6.33	1.74	6	27
21	30.5	32	7.34	2.01	7	27
22	30.5	36	8.6	2.36	8	27
23	30.5	42	10.12	2.78	10	27
24	30.5	73	14.17	3.89	23	27

Things to consider for liquid cooling

- **Water quality**

Parameter	Required
pH	9.5 to 10.5
Nitrite (NO2 ppm)	> 600
Molybdate (MoO4 ppm)	> 250
Conductivity (µS)	< 5,000
Combined Sulfate and Chloride (Na2SO4 + NaCl ppm)	< 500
Copper (ppm)	< 0.2
Aerobic bacteria (per ml)	< 1000
Anaerobic bacteria (per ml)	< 10
Total Hardness (ppm)	< 10
Iron (ppm)	< 1
Manganese (ppm)	< 0.1
Suspended solids (ppm)	< 20
Turbidity (NTU)	20

Bay area climate provides the ability to generate cool water all year round

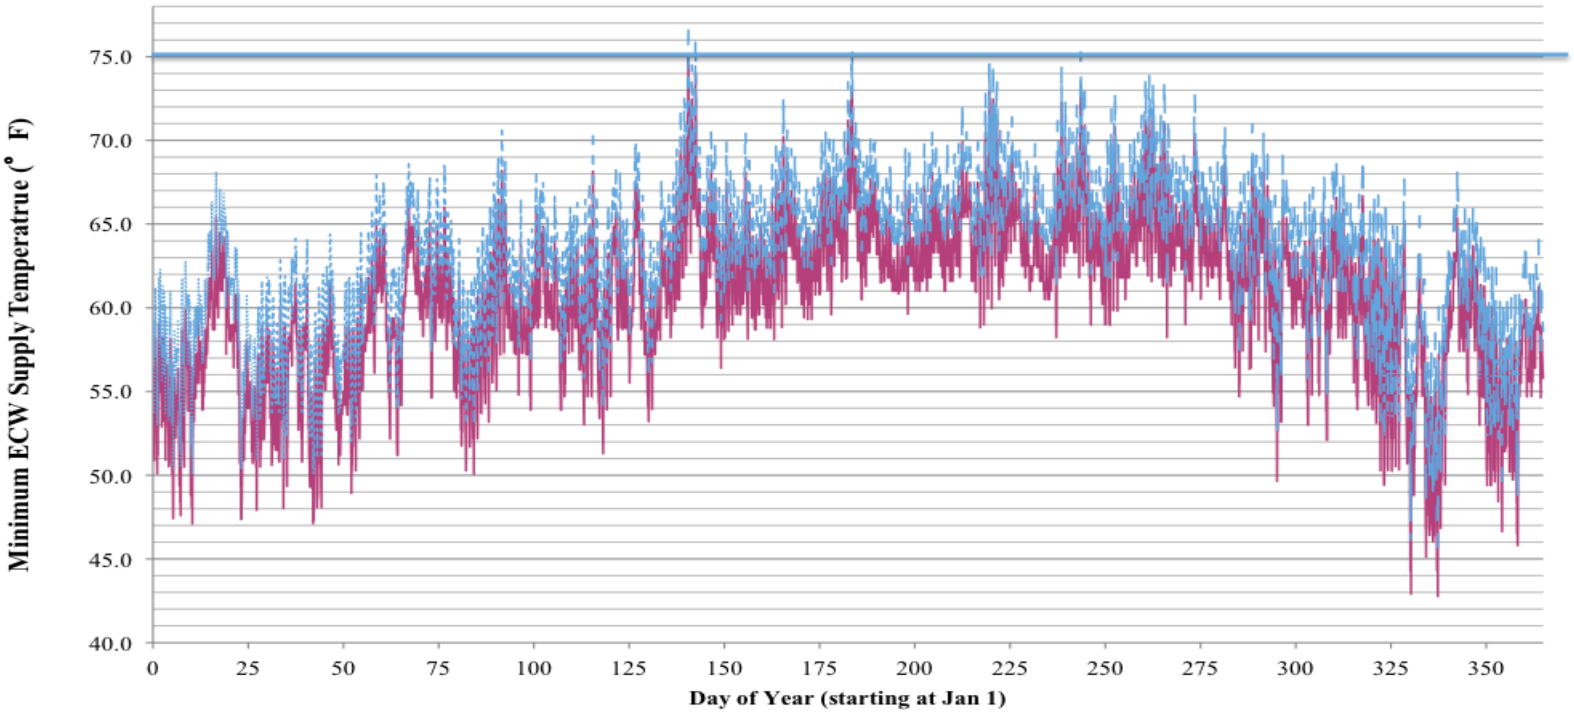

Supplied water temp (hours per year)	<70F	70-71F	71-72F	72-73F	73-74F	74-75F	>75F
One Cell	8258	261	101	64	45	19	12
Two Cell	8636	63	32	19	7	2	1

Power consumption with and without chillers

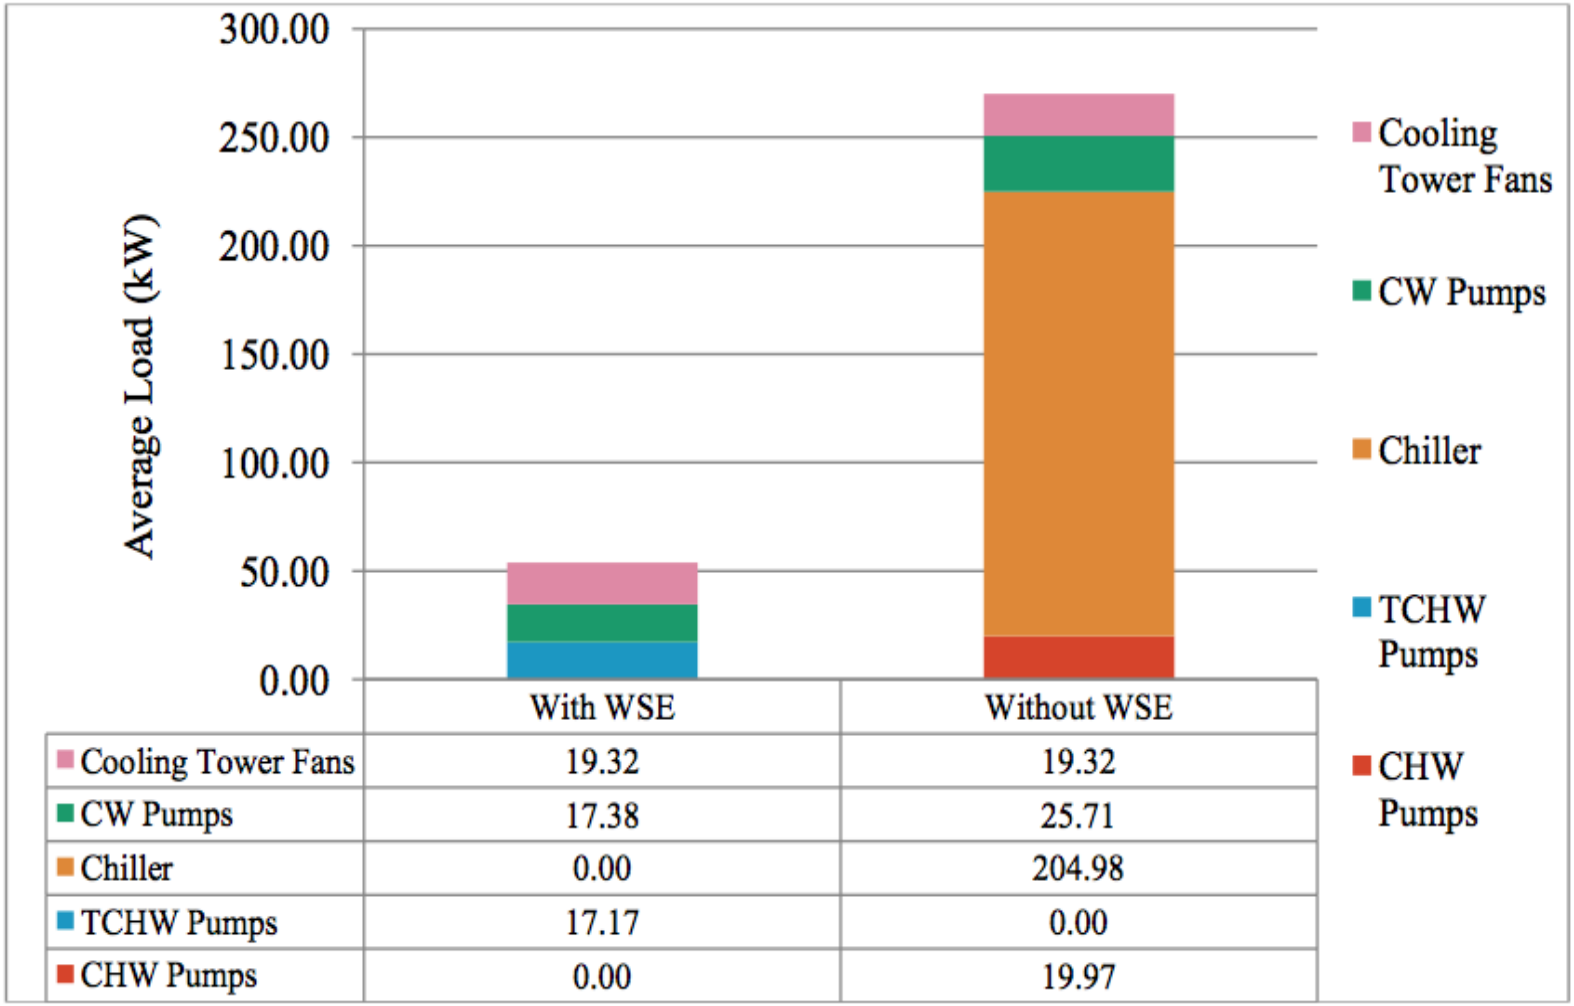

	With WSE	Without WSE
Cooling Tower Fans	19.32	19.32
CW Pumps	17.38	25.71
Chiller	0.00	204.98
TCHW Pumps	17.17	0.00
CHW Pumps	0.00	19.97

Power Savings and Payback Period

- **Annual average electrical consumption**
  - WSE (Free Cooling): 54 kW
  - Chiller Plant: 270 kW
- **Average Annual Load Savings of 216 kW (for WSE)**
  - 216 kW * 365 days * 24 hours/day = 1,892,000
- **Initial Cost of WSE (Free Cooling)**
  - Total rough estimate = $643,500
  - New construction of WSE system = $474,000
  - Demolition/Removal of packaged chiller plant = $169,500
- **Initial Cost of Chiller Plan**
  - Total rough estimate = $168,300
- **Annual savings = 1,892,000 kWh * $0.10/kWh = $189,200**
- **Cost Premium for Free Cooling = $643,500 - $168,300 = $475,200**
- **PG&E Rebate: $416,000**

$60,000 net install costs vs. $189,200 annual power savings!

Lessons Learned / Observations

- **Carefully calculate flow rate needs at the highest water temperatures, including pressure drop in distribution pipes**
- **Determine if "room neutral" or a "fixed set point" is right for your machine room**
- **Operational safety margins are built in, but automatic changes to control parameters are one way**
  - fixed in 7.0 UP03
- **Environmental control parameters are lost across a reboot**
  - Fixed in 7.0 UP02
- **Stainless steel fittings need to be handled carefully to avoid leaks**
- **Temperature stratification from the bottom to the top of the rack can still occur**

Future Work

- **Investigate feedback between XC30 and building control system**
  - Adjust pressure to meet XC30 needs and save on pumping costs
- **Control delays may introduce temperature and pressure optimizations**
  - XC30 modulates cooling by adjusting flow through cabinets
  - Effects may not be seen at cooling plant for several minutes due to time of travel of water through pipes
  - Tune plant controls to dampen oscillations
- **Optimize tradeoff between system power and energy for cooling**
  - Processor chips use less leakage current and perform more reliably (frequency scaling) at lower temperatures
  - Determine the best temperature to reduce total power not just improve PUE.

Seismic Isolation

- ISOBase "Ball and Cone" isolation platforms
- 8 in displacement
- Challenge is to route power, network and cooling through the platform
- XC30's flat bottom is particularly tricky

ISOBase In Action

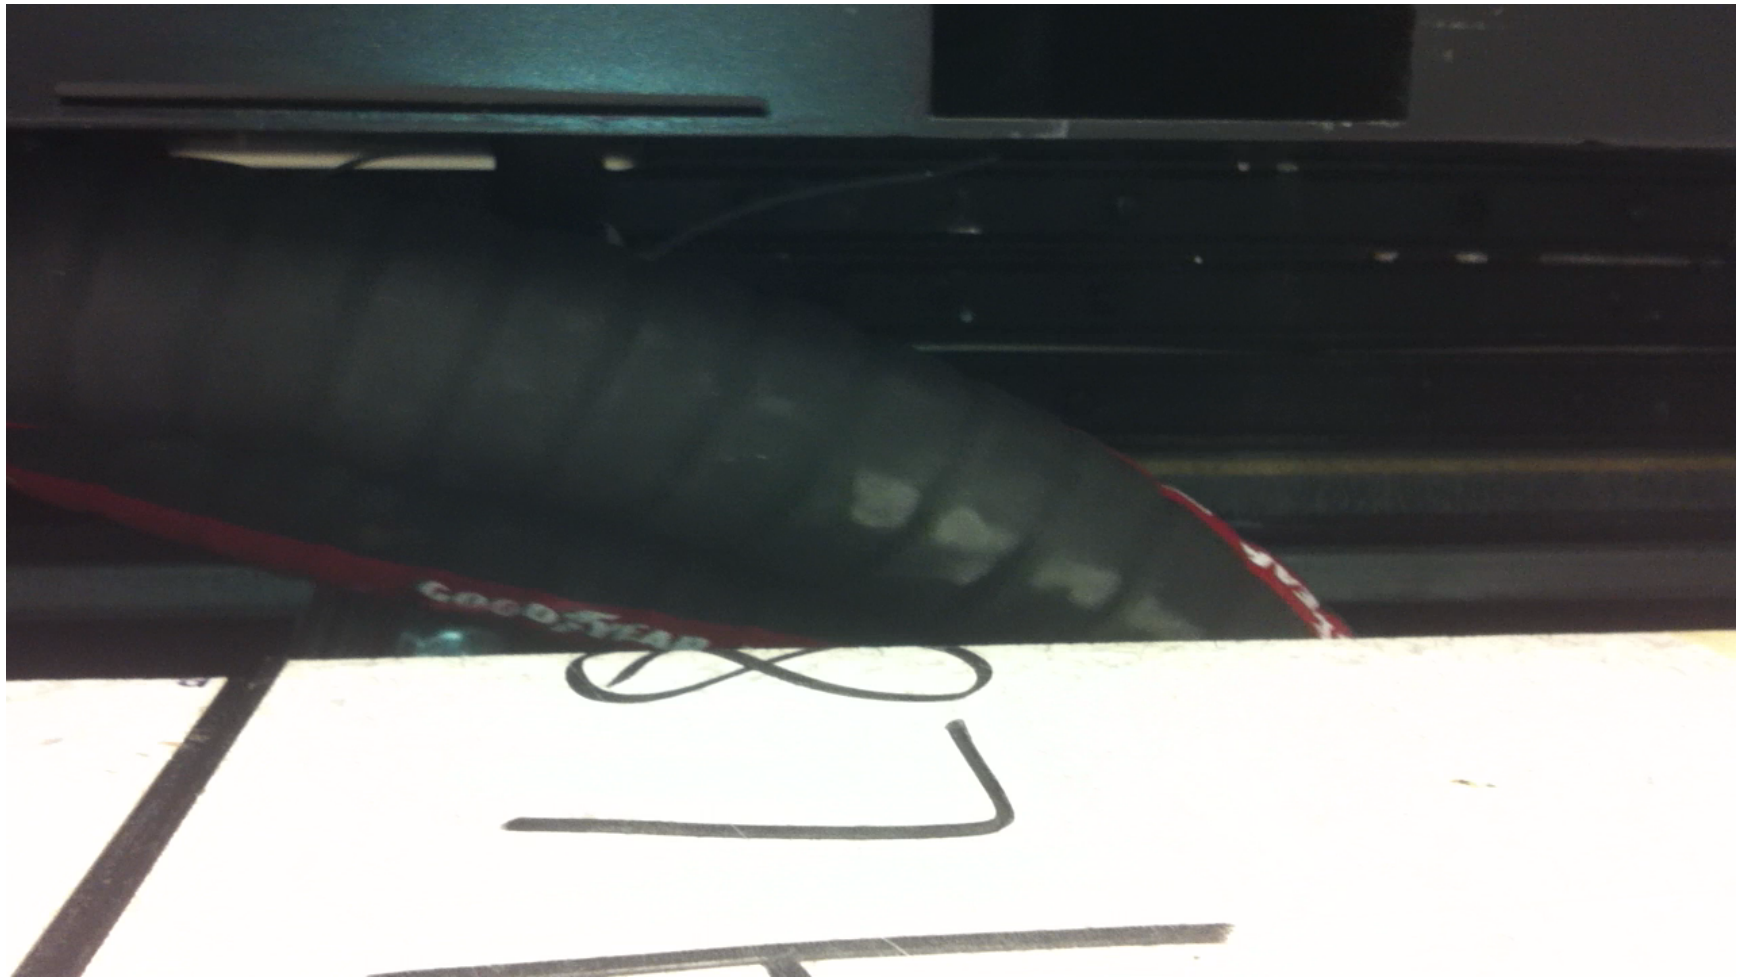

Edison Power Efficiency

	Edison Phase 1 (measured)	Edison Phase 2 (projected)	Hopper XE6 (measured)
System Cabinets	220 KW	1,540 KW	2,215 KW
Login Nodes and Storage	66 KW	69 KW	125 KW
Blowers	15 KW	90 KW	Incl. in system
Cooling Plant	38 KW	80 KW	538 KW
PUE (Blowers in IT load)	1.116	1.047	1.23
PUE (Blowers in mech. load)	1.175	1.106	N/A
Sustained Performance (SSP)	33 TF	236 TF	144 TF
Performance per Watt (including cooling)	99 GF/KW	132 GF/KW	50 GF/KW

NERSC will move to CRT in Fall 2014

- **Mixed office and data center building**
  - 300 offices on two floors
  - 20ksf HPC floor expandable to 28ksf
  - Mechanical space
  - 140ksf gross
- **Extremely energy efficient**
  - LEED gold design
  - "Free" cooling for air and water
  - Heat recovery
- **Expandable power**
  - 12.5MW at move-in
  - 42MW capacity to building

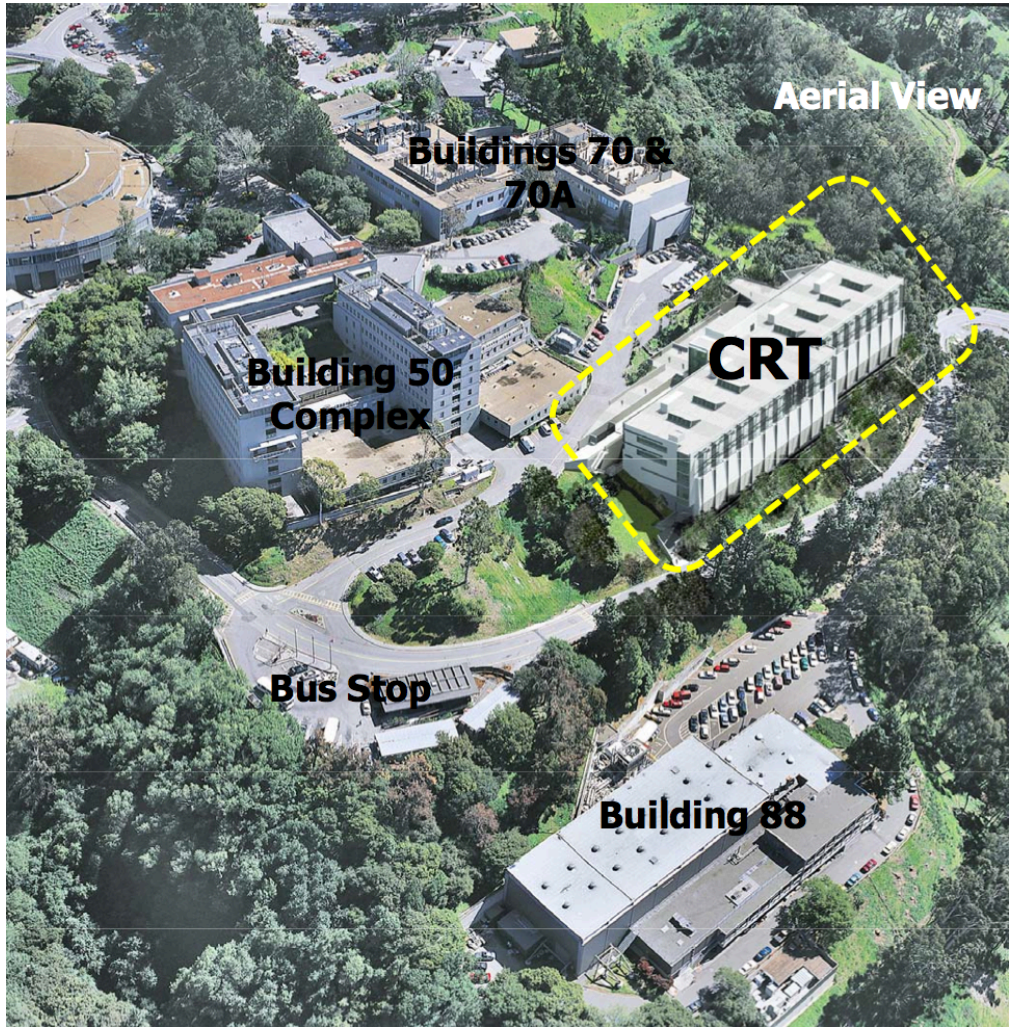

Local weather and hillside design enables energy savings

- **Air cooling**
  - 75F degree air year round without chillers
  - Computer room exhaust heat used to warm office floors
- **Liquid cooling**
  - 74F degree water year round without chillers
  - 65F degree water using chillers for 560 hours/year (6%)
- **Predicted Power Usage Effectiveness (PUE): < 1.1**

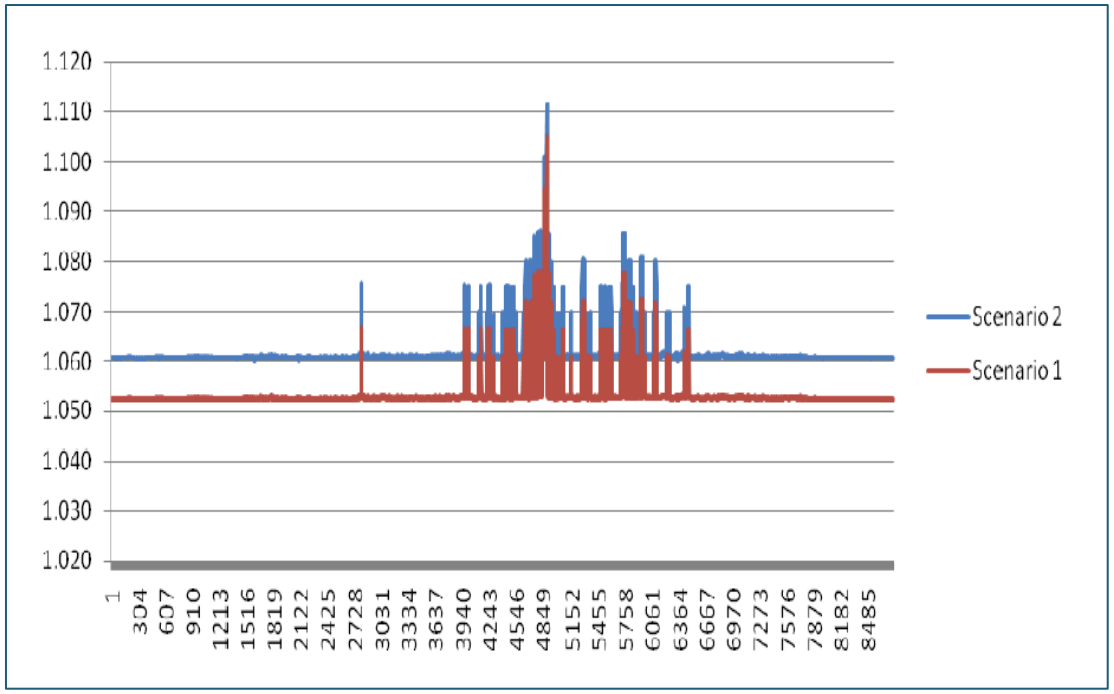

PUE simulation for two system configurations

National Energy Research Scientific Computing Center