Lessons Learned in Achieving Data Center Efficiency Goals – 3 Case Studies

Rich Reyher, Director, PayPal Core Platform & Infrastructure: Global Data Center, PayPal, Inc.
David Breland, Manager, Enterprise Data Center Facilities, Intuit, Inc.
Dale Sartor, Staff Scientist/Engineer, Lawrence Berkeley National Laboratory

Better Buildings®
U.S. DEPARTMENT OF ENERGY

Developing Innovative, Replicable Solutions with Market Leaders

- Better Buildings Challenge
- Better Buildings Alliance
- Better Buildings, Better Plants
- Better Buildings Accelerators
- Better Buildings Residential
- Strategic Energy Management

State, Local, and Federal Governments Leading by Example

- Better Communities Alliance
- Performance Contracting

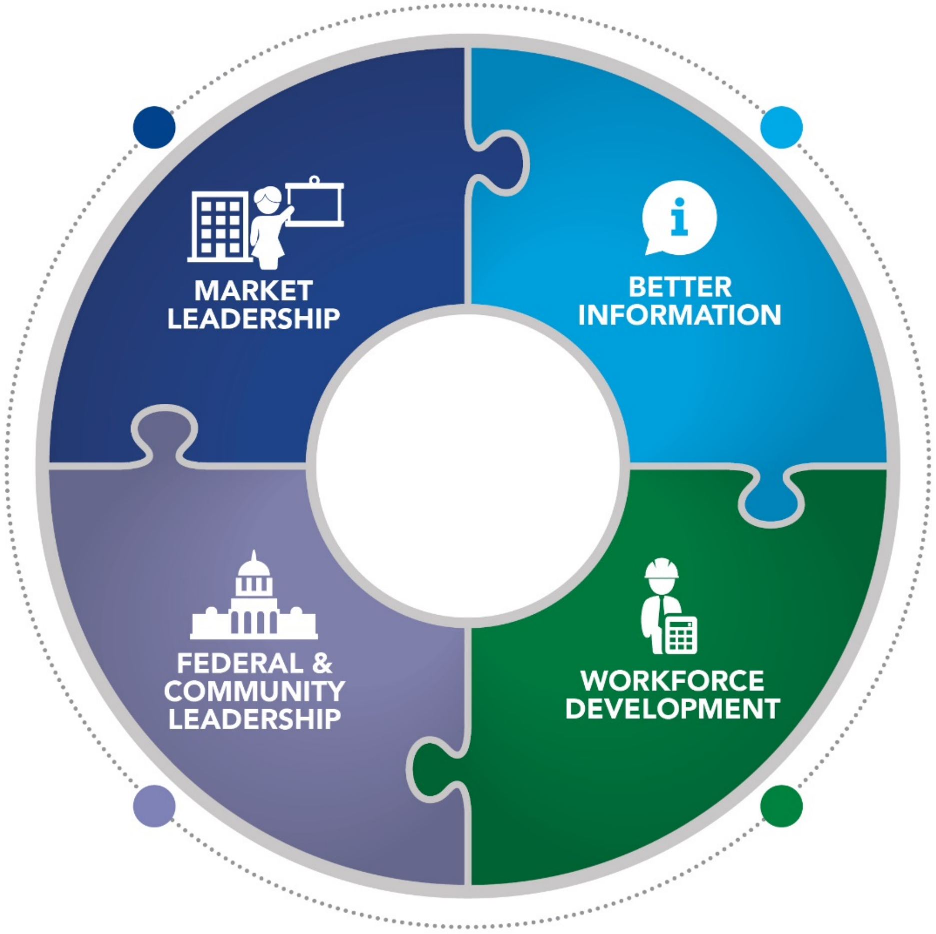

Making Energy Efficiency Investment Easier

- Better Buildings Solution Center
- Financing Navigator
- Improved Data Consistency and Access
- Tools to Assess the Efficiency of Buildings/Homes
- Tools for Energy Management

Expanding the Workforce

- Better Buildings Workforce Guidelines
- Industrial Energy Management Workforce

Better Buildings and Data Centers

Challenge partners set 10-year, 20% energy-savings targets across portfolio of buildings and share results

Better Buildings Challenge
CenturyLink
Digital Realty Trust
eBay Inc.
Facebook
Intel Corporation
Intuit
IO Data Centers
Iron Mountain Data Centers
Michigan State University
Sabey Data Center Partners
Staples

Accelerator partners set 5-year, 25% energy-savings targets across one or more data centers and provide regular updates on progress

Better Buildings Accelerator	
ANL	Defense Health Agency
EMSL	Georgia Tech
Indiana University	LBNL
LLNL	LANL
NERSC	NREL
ORNL	Stanford
State of Michigan	Home Depot
DOD-DISA	DOJ-DEA
VA	EPA
NASA	SSA
University of Colorado	University of Iowa
Virtustream	Waste Management

Data Center Path to Efficiency and Automation

October 2017

Data Center Path to Efficiency and Automation

PayPal's data center landscape

PayPal has grown rapidly over the last decade and our data center landscape has become increasingly complex. A key challenge was to maximize efficiency to ensure available capacity.

16 Countries

33 Data Centers

~2,500 Racks

27 Megawatts

3,000+ DC Infrastructure Assets

1M+ Monitoring Endpoints

3 primary DC partners

13 Tools & Platforms

~130 People

900+ DC Procedures

Reduced carbon footprint

Phoenix

- PayPal Owned and Operated
- Future PayPal Primary
- Tier IV & Tier II Design

Bluffdale

- Colocation Provider
- Tier III Design
- Home of DCG 13

Las Vegas

- Colocation Provider
- Tier IV Design
- Home of Analytics

Salt Lake City

- Colocation Provider
- Tier IV & Tier II Design
- Home of CCG & DCG

The efficiency challenge

We captured an initial baseline to determine our energy usage and identify several sources of inefficiencies

- Always-on cooling
- Lack of proper containment
- No temperature control zones
- Inefficient CRAH motors

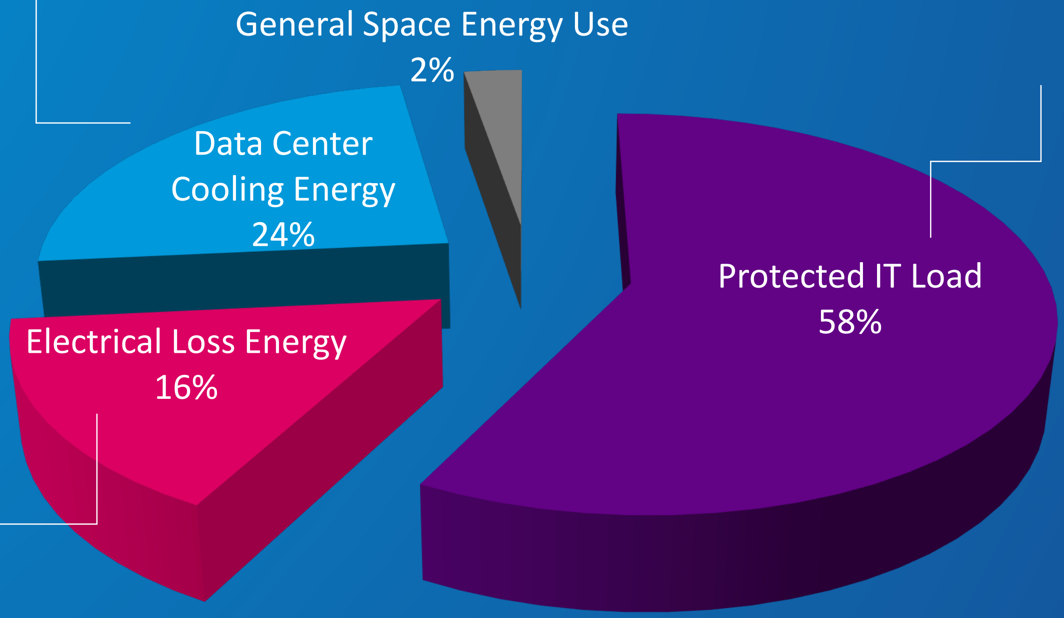

- Older transformer-based UPSs

- Lack of transparency on equipment ownership
- No mechanism for clear utilization showback
- Require a large buffer in capacity to ensure we meet business demand
- Over-allocation for failure scenarios
- Significant amount of unused hardware (i.e. zombie assets). Estimated utilization is ~40–50%

Our ultimate goal is to reach 70% true utilization for space and power in the data center and this requires tackling efficiency on multiple fronts

Video

Our journey so far

We tackled efficiency improvements in multiple phases moving from physical space and power efficiency to people, process and tools.

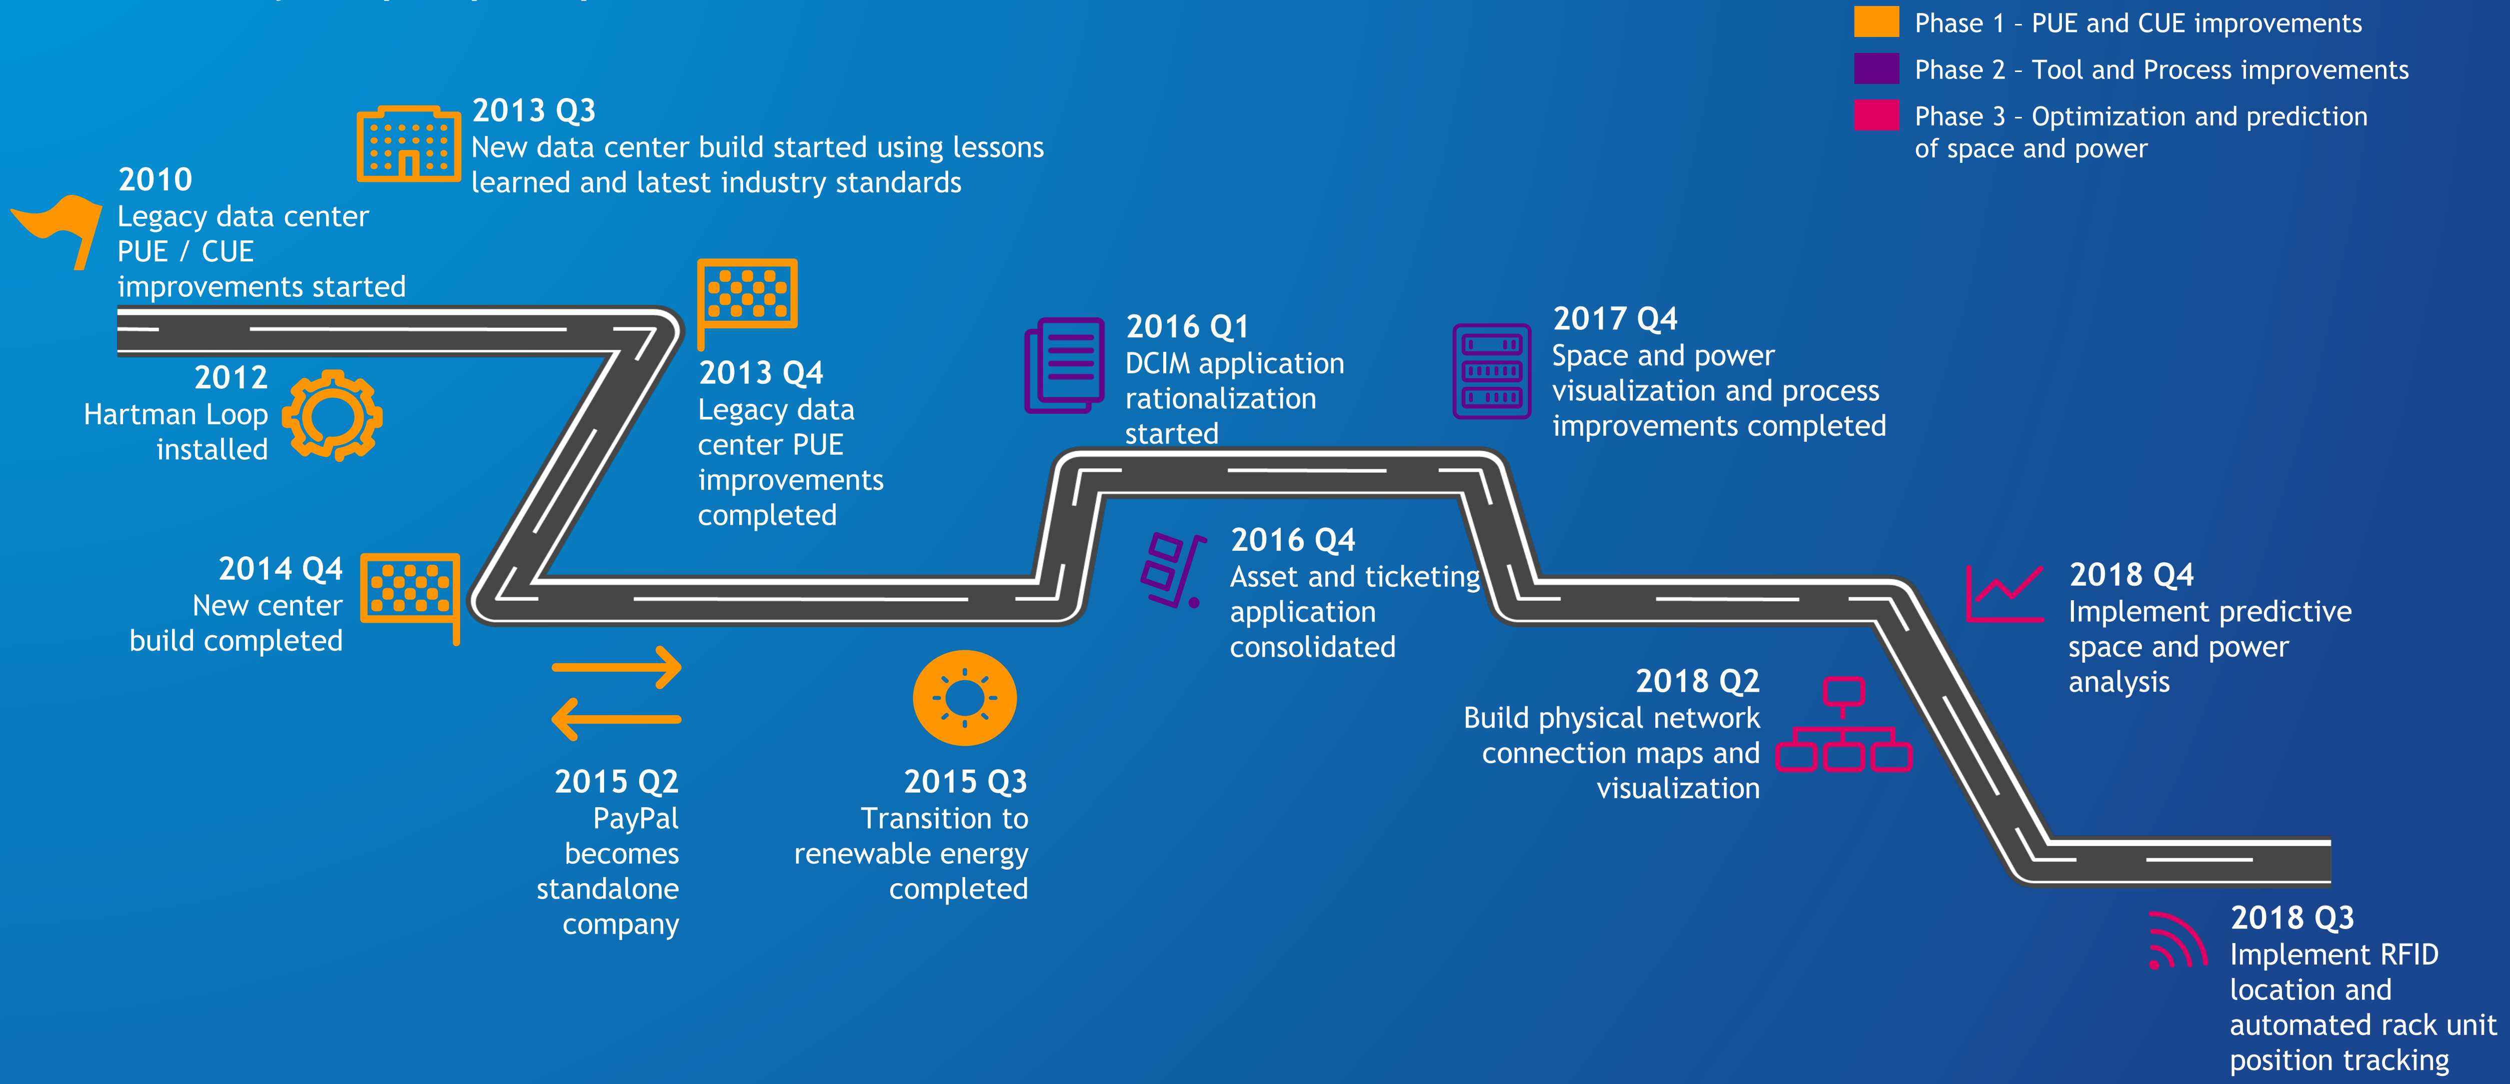

PayPal

Phase 1 – Improving space and power efficiency

We began improving efficiency by tackling quick wins in our legacy Phoenix data center

Current Situation
Reducing Data Center cooling energy

- The data center was operated as a standard tier IV facility
- The electrical losses and IT load required long lead times to address

Solution
Three year energy management plan

- Cold Aisle Containment
- Hartman loop chiller plant optimization
- Variable Speed Drives and new high efficiency motors on the CRAH systems
- Harmonic Mitigating Transformers in the CRAH electrical distribution system

Phase 1 – Impact on data center efficiency

The net result of the efficiency projects was an improvement of PUE from 1.85 to 1.49 (19%) across owned data centers.

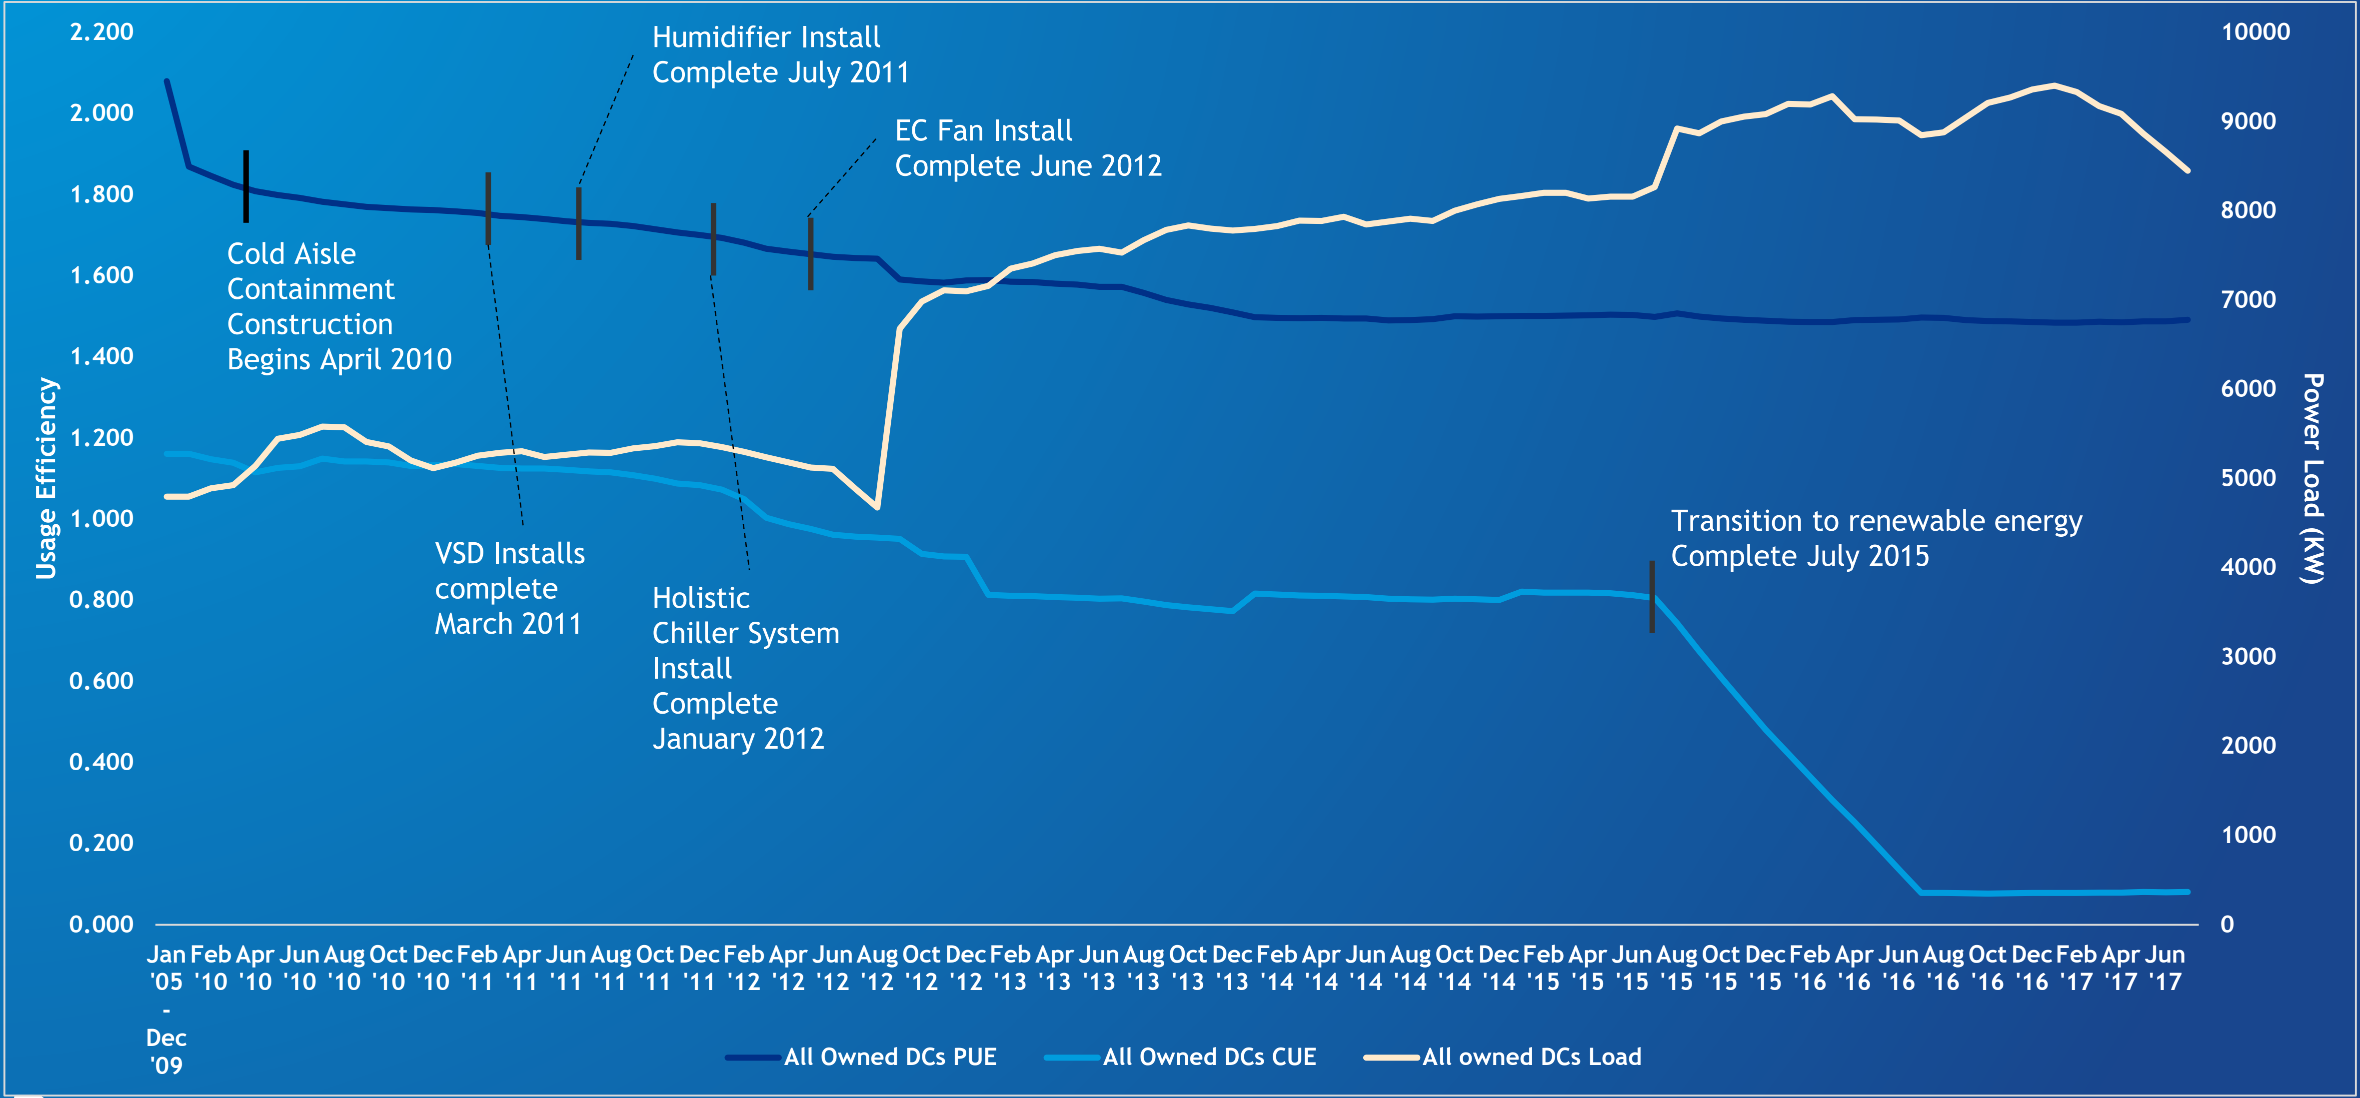

Phase 2 - How we have traditionally managed our data centers

Ticketing, monitoring, alarming, and data have historically been fragmented across multiple platforms

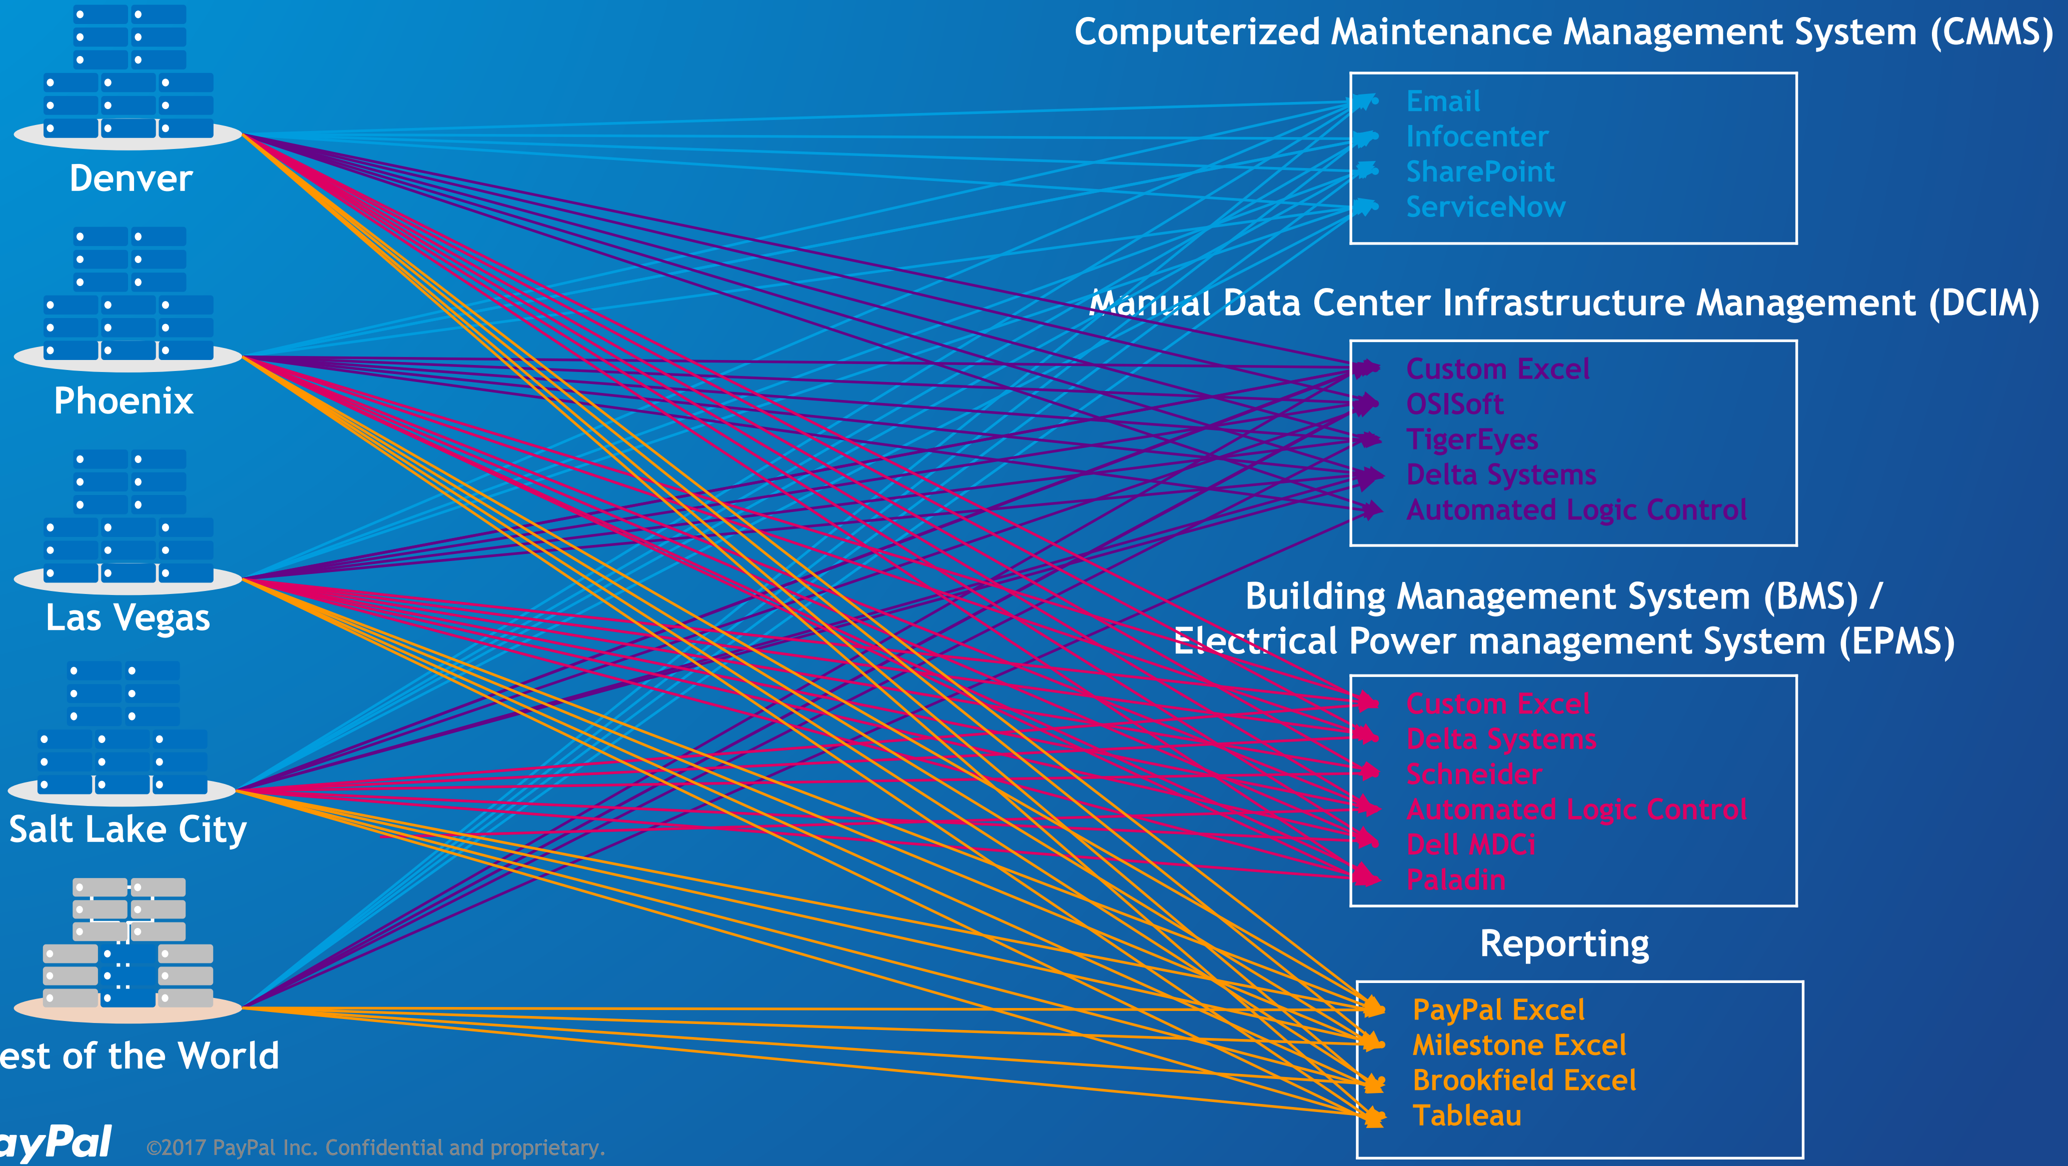

Phase 2 - How our data centers will be managed moving forward

Ticketing, monitoring, alarming, and data will be consolidated into a single pane of glass

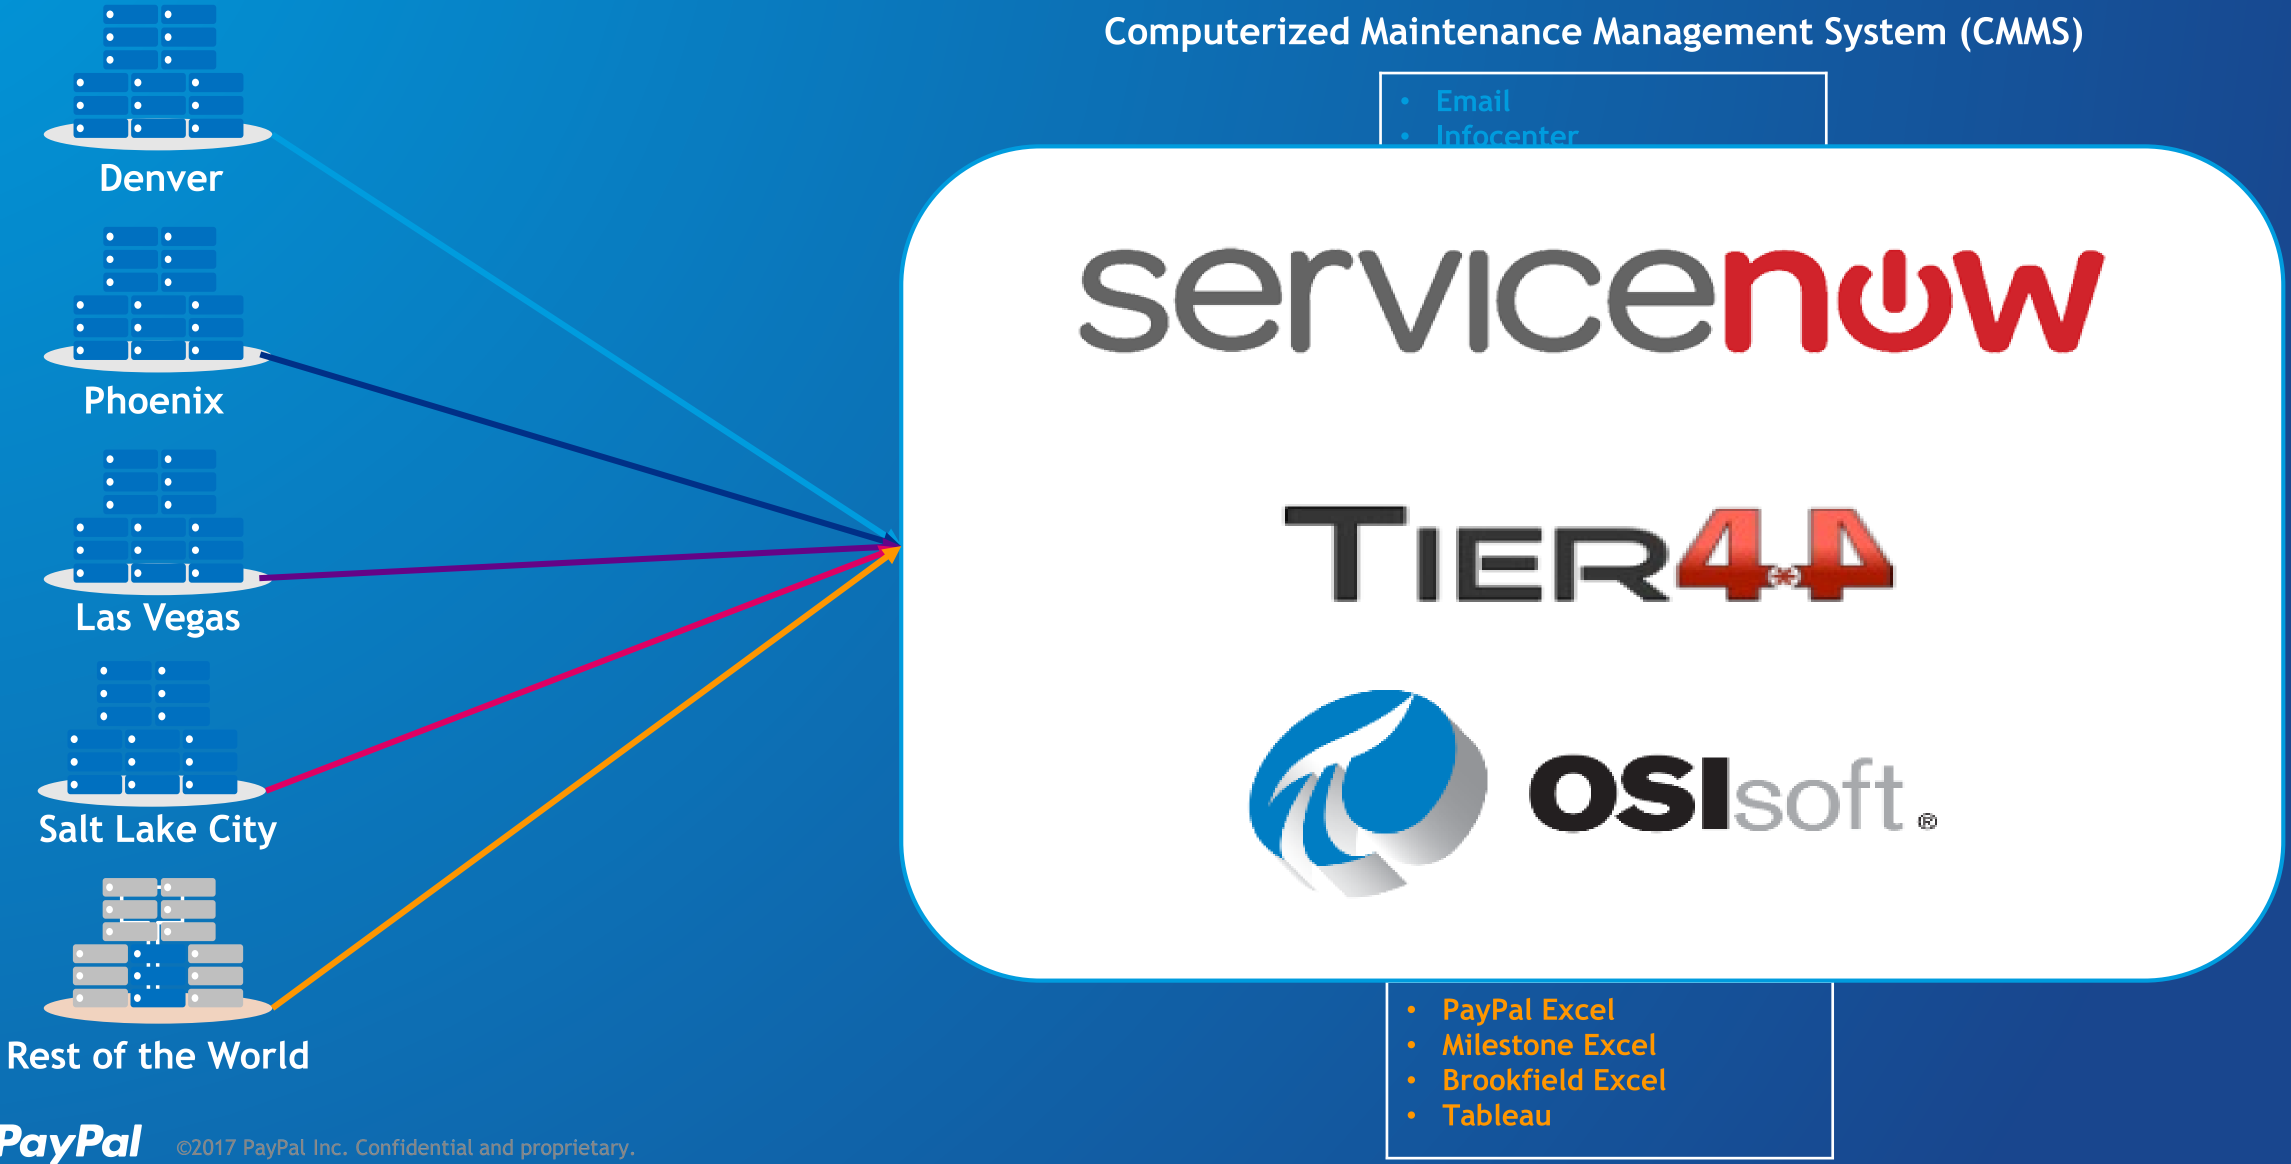

Phase 3 – Planned Future Improvements - "What's next?"

With the physical and application improvements in place we can begin to further automate our systems and predict future demand

Current Situation

Significant ongoing effort needed to continually audit assets to validate their location and position with racks

Lead times for provisioning can be weeks if there is no space and power available at the time of a request

Solution

Power Strip Tagging
Provides automated asset unit position within the rack saving time for audits and provides rack usage indicators.

RFID Tagging
Provides automated building level position and will be used to automatically receive assets and check them in/out of stock

Predictive analysis of space and power

1. Objective is to use data to more accurately project expected space and power usage
2. Projections would be used to automatically issue procurement or work tickets to stay ahead of the demand, reducing provisioning times from weeks to days

PayPal

HOW HOT CAN I GO?

Intuit's Journey in Quincy Washington

Dave Breland
Manager, Enterprise Data Center Facilities

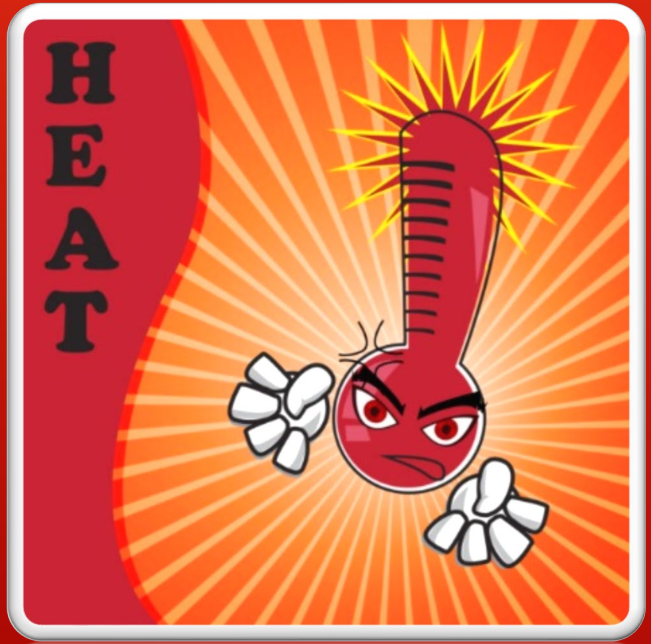

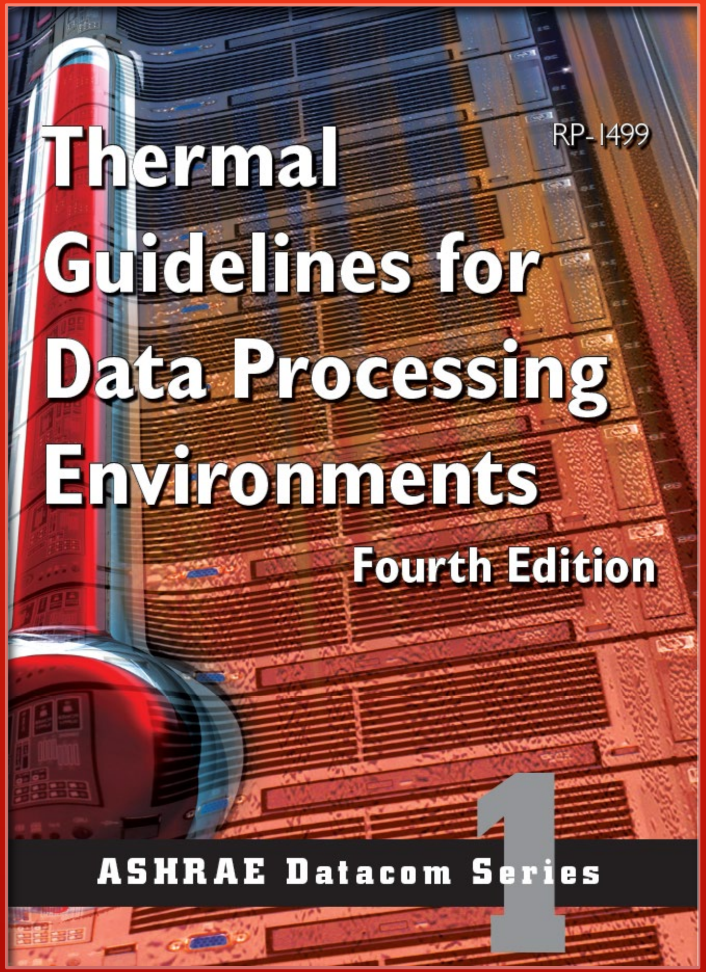

FIRST- I NEEDED A GUIDE

ASHRAE Datacom Series formerly TC9.9

INSIDE ASHRAE'S DATACOM GUIDE

Thermal Guidelines for Data Processing Environments, Fourth Edition | 13

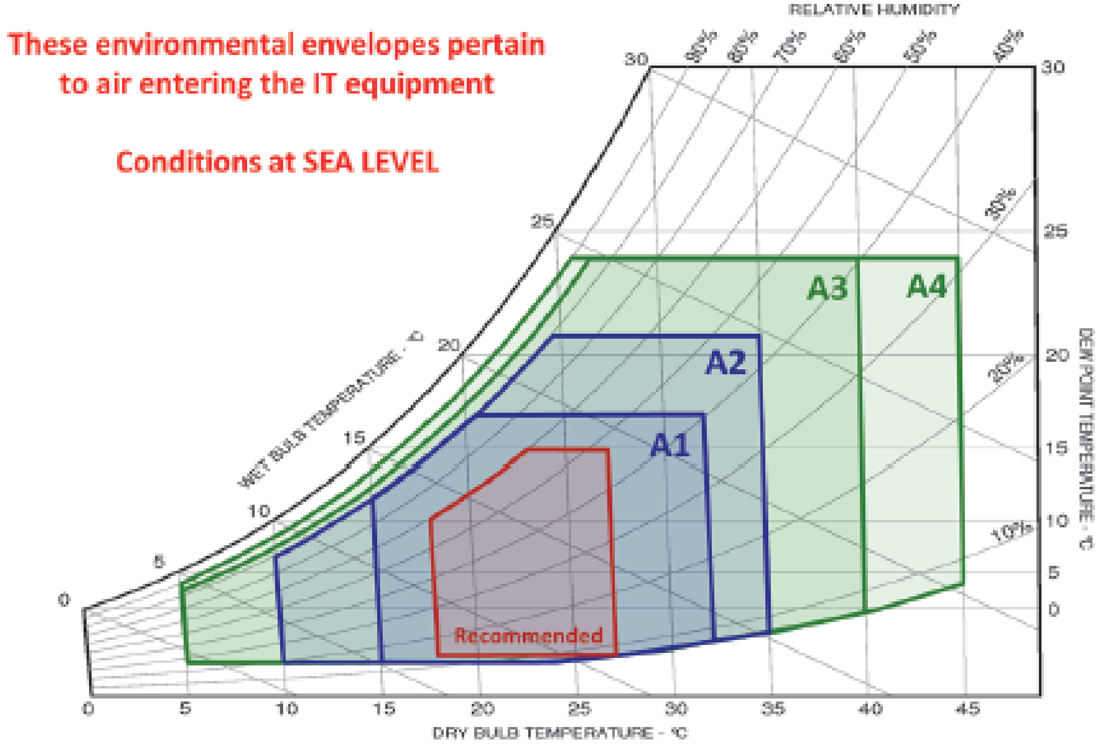

Figure 2.3 2015 recommended and allowable envelopes for ASHRAE Classes A1, A2, A3, and A4.

2.2.1 Environmental Class Definitions for Air-Cooled Equipment

Compliance with a particular environmental class requires full operation of the equipment over the entire allowable environmental range, based on nonfailure conditions.

Class A1: Typically a data center with tightly controlled environmental parameters (dew point, temperature, and RH) and mission-critical operations; types of products typically designed for this environment are enterprise servers and storage products.

IN ENGLISH

Recommended vs. Allowable

Table B.1 2015 Thermal Guidelines—I-P Version (SI Version in Table 2.1)

Equipment Environment Specifications for Air Cooling							
	Product Operation[b,c]					Product Power Off[c,d]	
Class[a]	Dry-Bulb Temperature[e,g], °F	Humidity Range, Noncondensing[h,i,k,l]	Maximum Dew Point,[k] °F	Maximum Elevation[e,j,m], ft	Maximum Rate of Change[f], °F/h	Dry-Bulb Temperature, °F	Relative Humidity,[k] %
Recommended (Suitable for all four classes; explore data center metrics in this book for conditions outside this range.)							
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 and 8% rh to 62.6°F DP and 80% rh	62.6	10,000	9/36	41 to 113	8 to 80
A2	50 to 95	10.4°F DP and 8% rh to 69.8°F DP and 80% rh	69.8	10,000	9/36	41 to 113	8 to 80
A3	41 to 104	10.4°F DP and 8% rh to 75.2°F DP and 85% rh	75.2	10,000	9/36	41 to 113	8 to 80
A4	41 to 113	10.4°F DP and 8% rh to 75.2°F DP and 90% rh	75.2	10,000	9/36	41 to 113	8 to 80
B	41 to 95	8% to 82.4°F DP and 80% rh	82.4	10,000	N/A	41 to 113	8 to 80
C	41 to 104	8% to 82.4°F DP and 80% rh	82.4	10,000	N/A	41 to 113	8 to 80

[a] For potentially greater energy savings, refer to the section "Detailed Flowchart for the Use and Application of the ASHRAE Data Center Classes" in Appendix C for the process needed to account for multiple server metrics that impact overall TCO.

HOW DO I KNOW?

Is there some way for me to know for sure?

Dell states 10°C – 35°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within Dell Requirements

Table 33. Environmental specifications

Fresh Air: temperature, humidity, altitude de-rating	
Continuous operation	10°C to 35°C (50°F to 95°F) at 10% to 80% relative humidity with 26°C (78.8°F) maximum dew point (maximum wet bulb temperature). De-rate maximum allowable dry bulb temperature at 1°C per 300m above 950m (1°F per 547 ft above 3117 ft).

Dell states 10°C – 35°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within Dell Requirements

Table 27. Environmental specifications

Fresh Air: temperature, humidity, altitude de-rating	
Continuous operation	10°C to 35°C (50°F to 95°F) at 10% to 80% relative humidity with 26°C (78.8°F) maximum dew point (maximum wet bulb temperature). De-rate maximum allowable dry bulb temperature at 1°C per 300m above 950m (1°F per 547 ft above 3117 ft).

HP states 10°C – 35°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within HP Requirements

HP ProLiant DL170e G6 Server - Overview

Product description

The DL170e G6 is optimized for efficiency, density and flexibility and can be serviced individually without impacting the operation of other server nodes sharing the same chassis. The DL2000 consists of up to 4 independent DL170e G6 servers in the 2U HP ProLiant e2000 G6 Chassis. The servers share power supplies and fans, providing greater power and cooling efficiencies.

Environmental specifications

Item	Description
Thermal output (maximum operating)	392 W/hr
System inlet temperature, operating	50° to 95° F (10° to 35° C) at sea level with an altitude derating of 1.8°F per every 1000 ft (1.0°C per every 305 m) above sea level to a maximum of 10,000 ft (3050 m), no direct sustained sunlight. Maximum rate of change is 18°F/hr (10°C/hr). The upper limit may be limited by the type and number of options installed. System performance may be reduced if operating with a fan fault or above 86°F (30°C).

HP states 10°C – 35°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within HP Requirements

System Inlet Temperature	Operating

50° to 95° F (10° to 35° C) at sea level with an altitude derating of 1.8°F per every 1000 ft (1.0°C per every 305 m) above sea level to a maximum of 10,000 ft (3050 m), no direct sustained sunlight. Maximum rate of change is 18°F/hr (10°C/hr). The upper limit may be limited by the type and number of options installed.
System performance may be reduced if operating with a fan fault or above 86°F (30°C).

HP ProLiant DL2000 server technologies

Cisco Nexus 9396PX

Cisco states 0°C – 40°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within Cisco Requirements

Environmental Specifications

Environment		Specification
Temperature	Ambient operating temperature	32 to 104° F (0 to 40° C)
	Ambient nonoperating	-40 to 158° F (-40 to 70° C)
Relative humidity	Ambient (noncondensing)	5 to 95%
Altitude	Operating	0 to 13,123 feet (0 to 4,000 meters)

Cisco states 0°C – 40°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within Cisco Requirements

Climatic Environment

Table 10 summarizes the climatic environment for the Cisco Nexus 7000 4-, 9-, 10-, and 18-Slot chassis.

Table 10. **Climatic Environment**

Item	Description			
	Cisco Nexus 7000 4-Slot Chassis	Cisco Nexus 7000 9-Slot Chassis	Cisco Nexus 7000 10-Slot Chassis	Cisco Nexus 7000 18-Slot Chassis
Temperature	• Operating: 32 to 104°F (0 to 40°C) • Short term: 23 to 131°F (-5 to 55°C)* • Nonoperating: -40 to 158°F (-40 to 70°C)* Note: • Chassis external thermal requirements are defined in the GR-63-CORE Network Equipment Building Standards (NEBS) specification published by Telcordia Technologies in Section 4.1.2, Operating Temperature and Humidity Criteria. • Short term refers to a period of not more than 96 consecutive hours and a total of not more than 15 days in 1 year (a total of 360 hours in any given year, but no more than 15 occurrences during that 1-year period). * 131°F (55°C) and 25% relative humidity for equipment shelves that fill less than half the frame (length)			
Humidity	• Relative humidity (nonoperating): 5 to 95%, noncondensing • Relative humidity (operating): 5 to 90%, noncondensing Note: An ambient relative humidity between 45 and 50% is suggested to reduce corrosive problems, to provide an operating time buffer in the event of failures, and to reduce interference from static discharges.			

NetApp Disk Shelves and Storage Media Technical Specifications

NetApp states 0°C – 40°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within NetApp Requirements

Environmental Specifications

Environment		Specification
Temperature	Ambient operating temperature	32 to 104° F (0 to 40° C)
	Ambient nonoperating	-40 to 158° F (-40 to 70° C)
Relative humidity	Ambient (noncondensing)	5 to 95%
Altitude	Operating	0 to 13,123 feet (0 to 4,000 meters)

Apple states 0°C – 35°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within Apple Requirements

Electrical and operating requirements

- Line voltage: 100–240V AC
- Frequency: 50Hz to 60Hz, single phase
- Maximum continuous power: 450W
- Operating temperature: 50° to 95° F (10° to 35° C)
- Relative humidity: 5% to 95% noncondensing
- Maximum altitude: 16,400 feet (5000 meters)
- Typical acoustical performance, sound pressure level (operator position): 12 dBA at idle

Mac

Mac Pro

Pure Storage states 5°C – 35°C while the ASHRAE Guideline allows 15°C – 32°C.
ASHRAE Guidelines are well within Pure Storage Requirements

ENVIRONMENTAL

//M10, //M20, //M50, //M70	RANGE	NOTES
Operating Temperature	5° to 35° C	Derate 1 C per 300 m above 950 m
Non-operating Temperature	0° to 60° C	
Operating Humidity	10 – 80%	Non-condensing
Non-operating Humidity	5 – 95%	At a maximum temperature of 33° C. Non-condensing, web bulb 33° C.

Appendix E
OSHA and Personnel Working in High Air Temperatures

Although OSHA does not have a particular regulation or standard that covers high-temperature environments, the General Duty Clause, Section 5(a)(1) of the Occupational Safety and Health Act of 1970 (OSHA 2012b), requires each employer to "furnish to each of his employees employment and a place of employment which are free from recognized hazards that are causing or are likely to cause death or serious physical harm." OSHA has interpreted this rule such that employers shall provide means and methods that will reduce the likelihood of worker heat stress. These means or methods may include issuing personal protective equipment (PPE), minimizing exposure through frequent breaks, frequent hydration, and developing a heat stress program. There are various manufacturers that produce PPE for hot working environments.

Table E.1 Permissible Heat Exposure Threshold Limit Value (TLV) (ACGIH 1992)

Work/Rest Regimen	Work Load[1]		
	Light	Moderate	Heavy
Continuous work	30.0°C (86°F)	26.7°C (80°F)	25.0°C (77°F)
75% work, 25% rest, each hour	30.6°C (87°F)	28.0°C (82°F)	25.9°C (78°F)
50% work, 50% rest, each hour	31.4°C (89°F)	29.4°C (85°F)	27.9°C (82°F)
25% work, 75% rest, each hour	32.2°C (90°F)	31.1°C (88°F)	30.0°C (86°F)

1. Values are in °C and °F (wet-bulb globe temperature [WBGT]).

INTUIT'S JOURNEY IN QUINCY WASHINGTON

THEN

- Original Design 72°f top of Rack
- CWS Temp 47°f
- Economizer Start OAWBT 45°f with a CWST of <47.5°f.
- Economizer Stop 48.5°f CWST
- Averaged 48% of year on economizer at current load

NOW

- Current Control 78.5°f top of Rack Warn 80°f Alarm
- CWS Temp 61°f
- Economizer Start OAWBT 58°f with a CWST of <61.5°f.
- Economizer Stop 64°f CWST
- Now average 65% on water side economizer

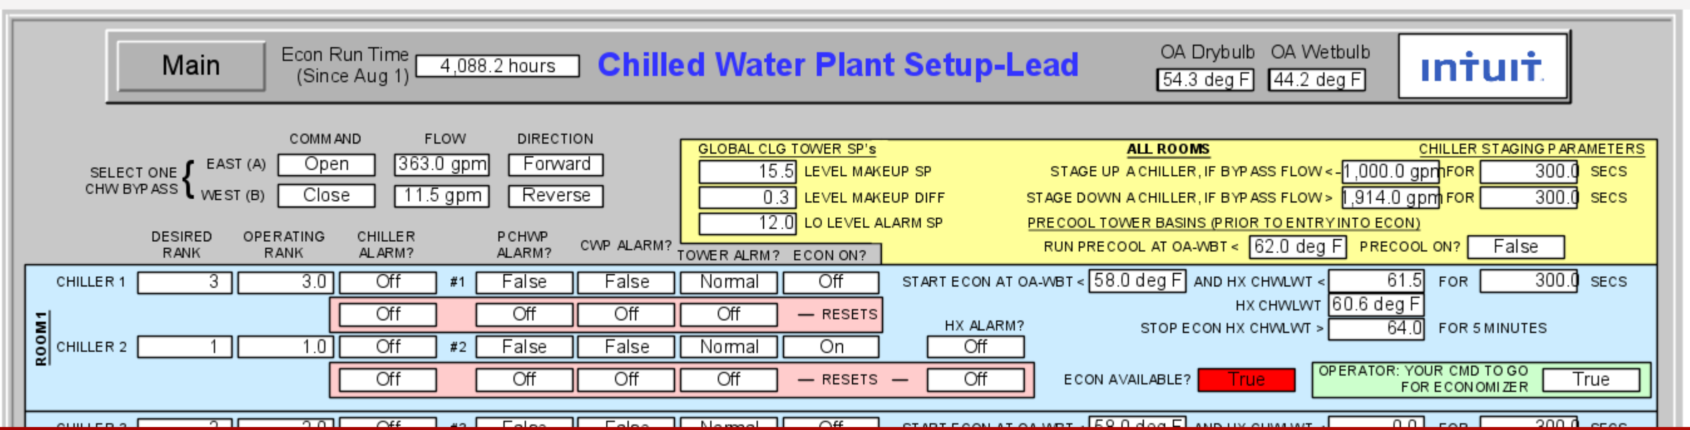

One of Our Cooling Plants

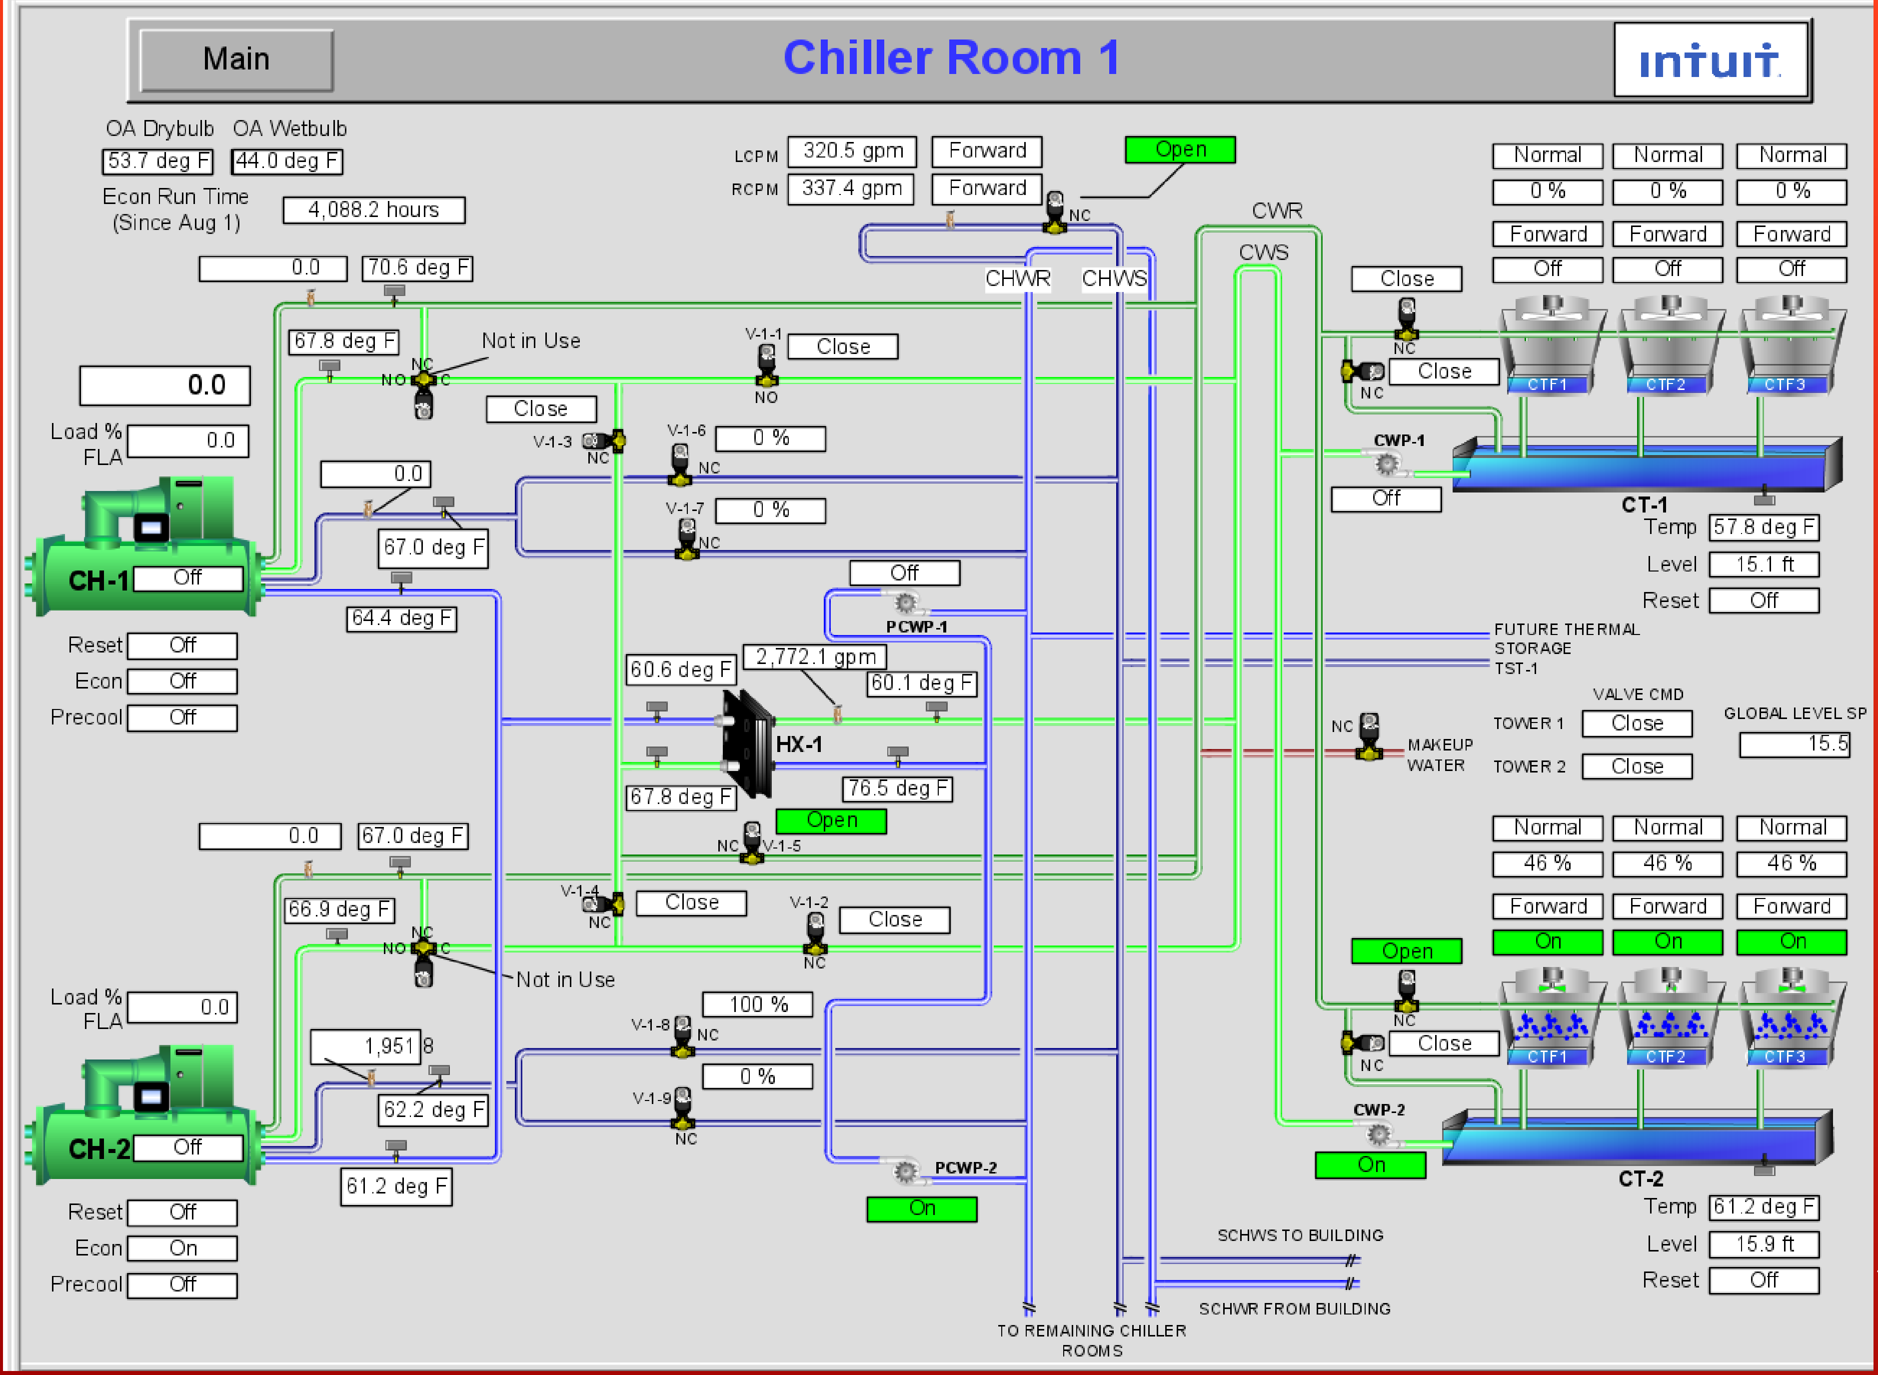

- Balance data halls using Purkay instruments and Tileflow Software, to ensure cold aisle temps ≤ 78.5°f. at top of rack.
- Raised CWST to get all of the CRAH control valves at their mid range.
- Operators monitor weather patterns and redundant cooling tower storage to extend economizer hours when possible.
- My biggest bang is staying on my economizer as long as possible.
- With my water side economizer, it's as far as I can go. Most later version data centers around me are using 100% outside air or probably half have gone to the Munters Evaporative Cooling units.

SO THIS IS HOW WE HAVE DETERMINED HOW HOT WE CAN GO

National Energy Research Scientific Computing Center (NERSC)

Shyh Wang Hall

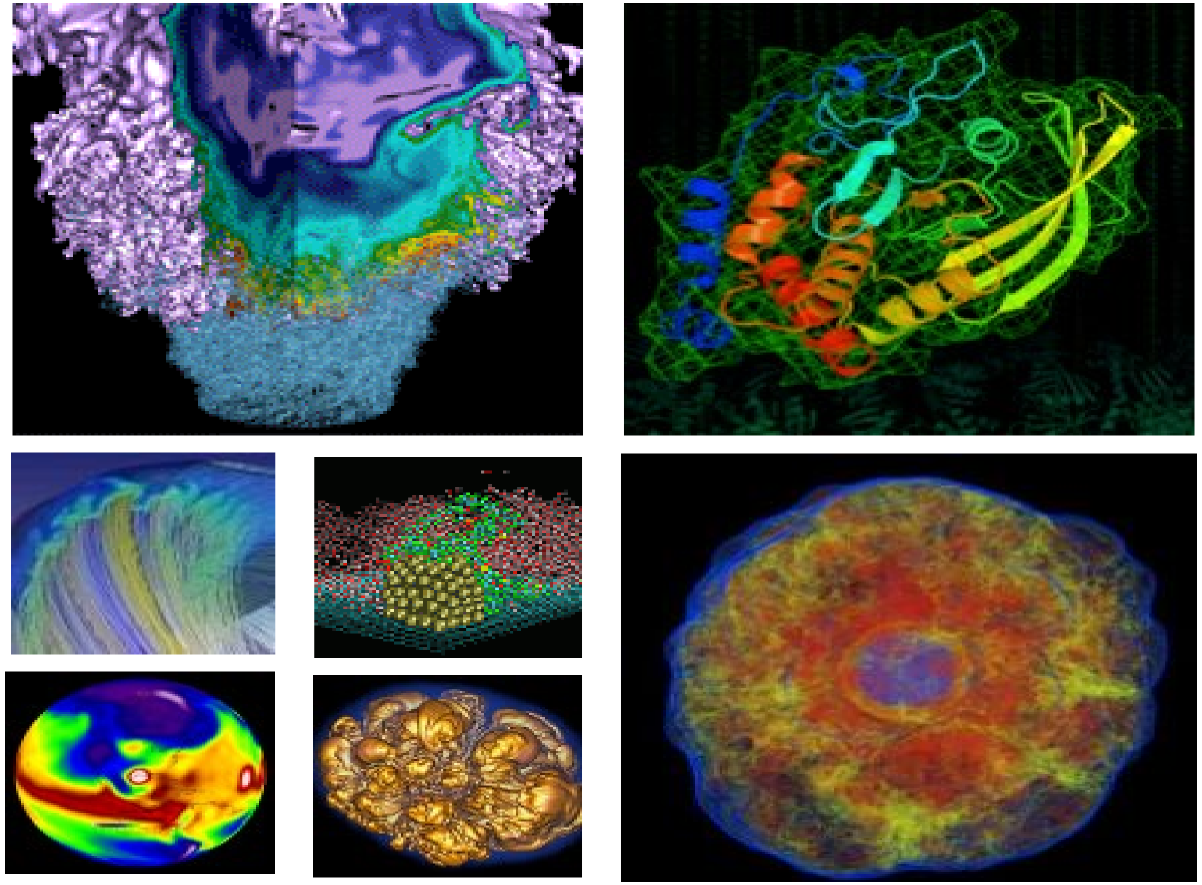

7x24 – Shaping The Future
Phoenix, Arizona
October 17, 2017
Dale Sartor, PE

NERSC

U.S. DEPARTMENT OF ENERGY | Office of Science

Version: 090817

BERKELEY LAB
Lawrence Berkeley National Laboratory

Shyh Wang Hall, the New Home for NERSC

UC's Shyh Wang Hall

- 142,000 square feet total
- 7 MW IT load to start, then up to 17, then ???
- IT load will dominate building
- 4 large AHUs for air-cooled loads
- 4 cooling towers for water cooled loads
- Water-cooled supercomputers
- Air and water side economizers
- IT loads cooled without compressors
- Air-side heat recovery for heating offices

High and Low Density Computing

NERSC hosts Cori
(#5 Top 500, Nov 2016)

File Systems and Air Cooled Computers

Free Cooling - Air System Design

- 2 MW Air Cooling (17 MW max)
- Annual PUE less than 1.1
- Air-side economizer
- Direct evaporative cooling for humidification & cooling

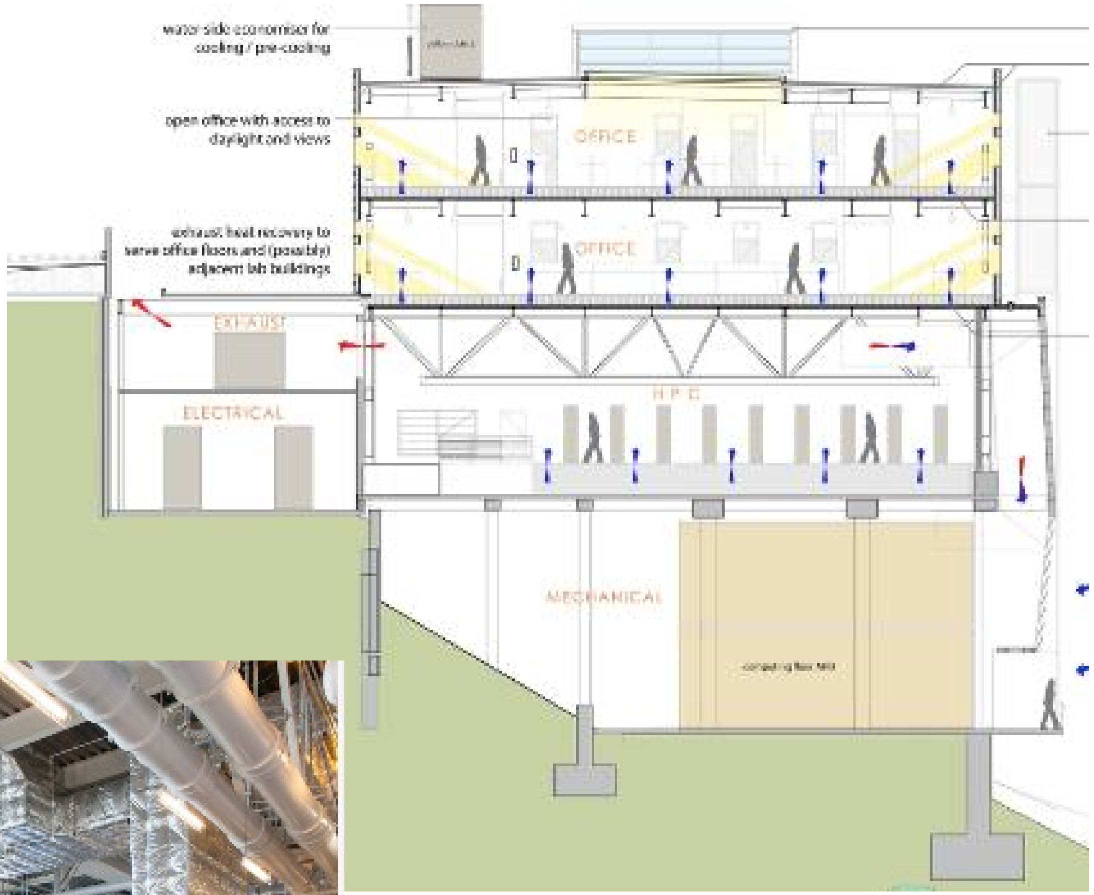

Outside Air Based Design

1. Blue = original design supply
2. Green can become blue mixing return and outdoor air
3. Most of the conditions below and right of blue can be satisfied w/ evaporative cooling
4. Hot and humid hours will enter the "allowable" range

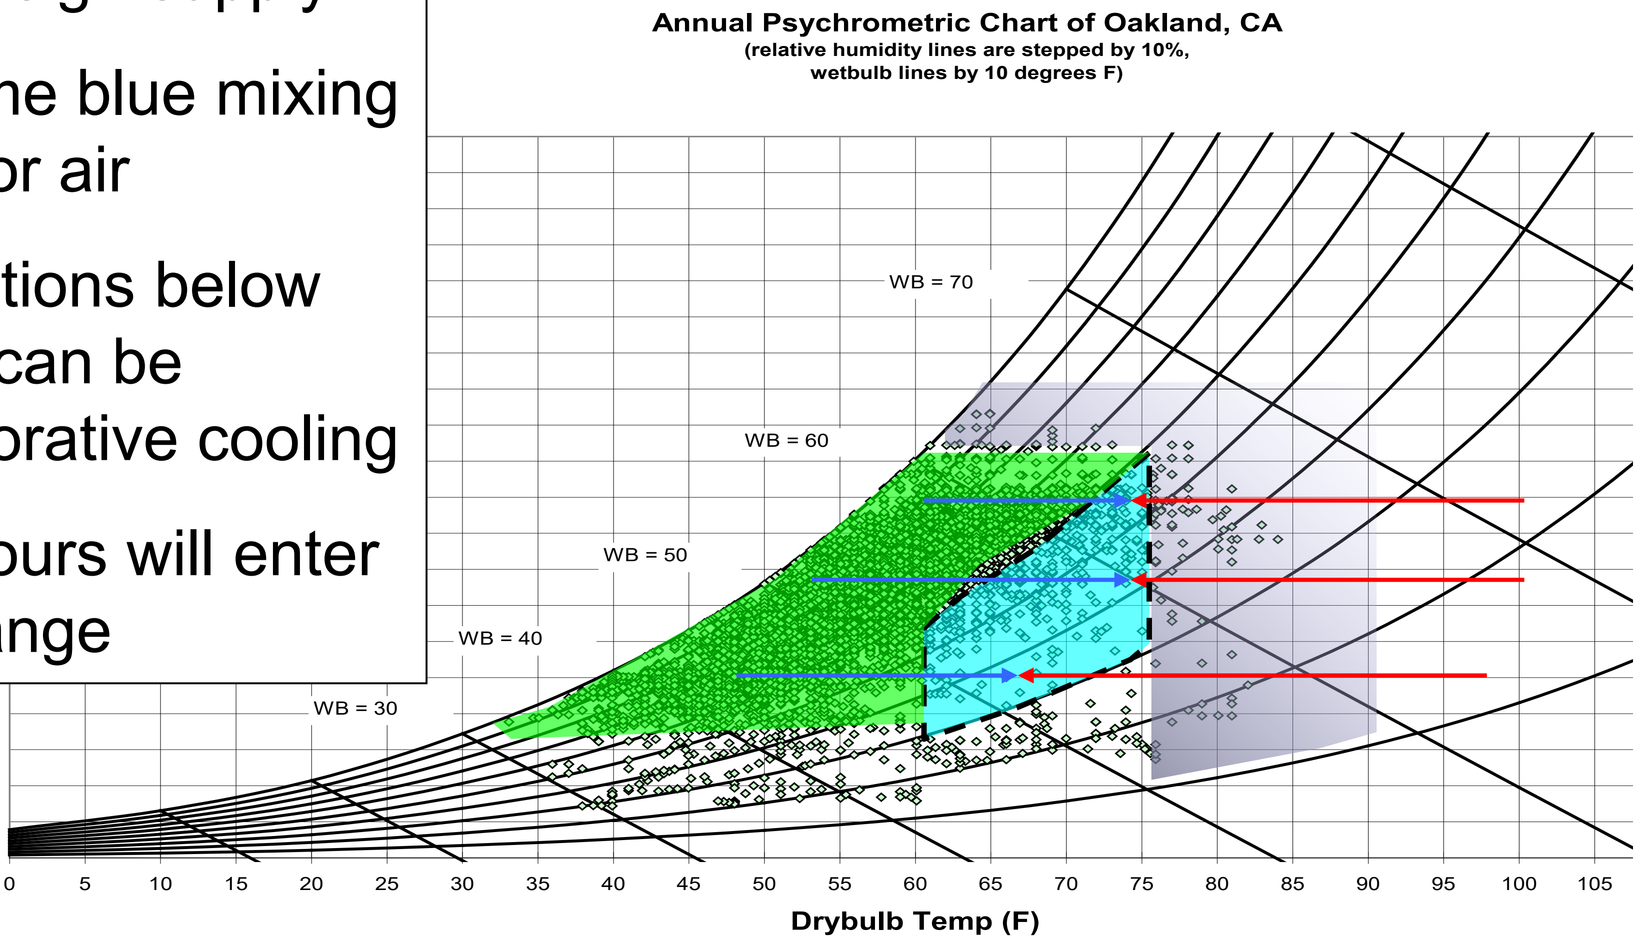

Free Coolng - Water Based Design

- 10 MW liquid cooling (20 MW max)
- Annual PUE less than 1.1
- Closed-loop treated cooling water from cooling towers
- Headers, valves and caps for modularity and flexibility

Real Life Results

- Performance
- Opportunities for improvement/optimization
- Lessons learned

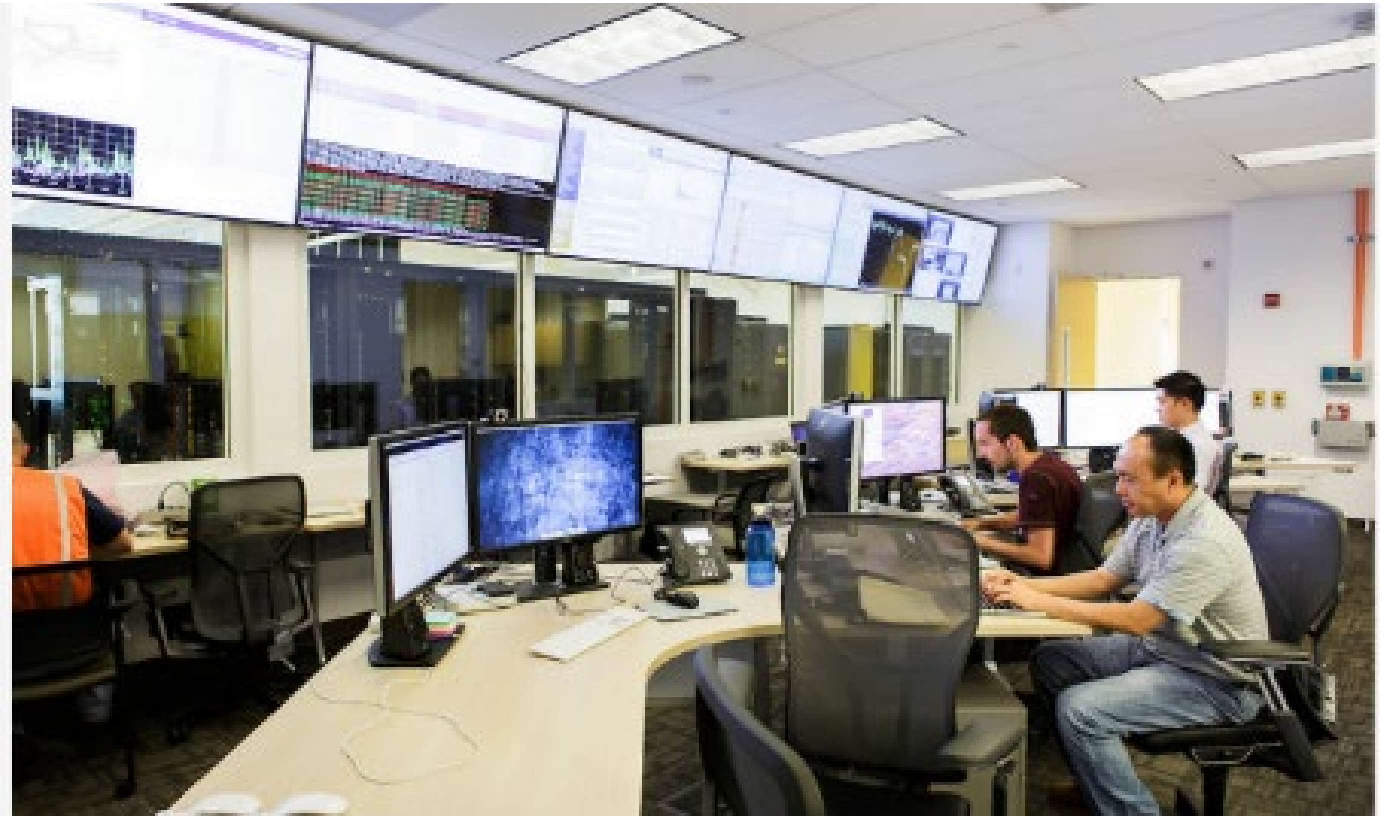

Liquid Cooling Performance Baseline

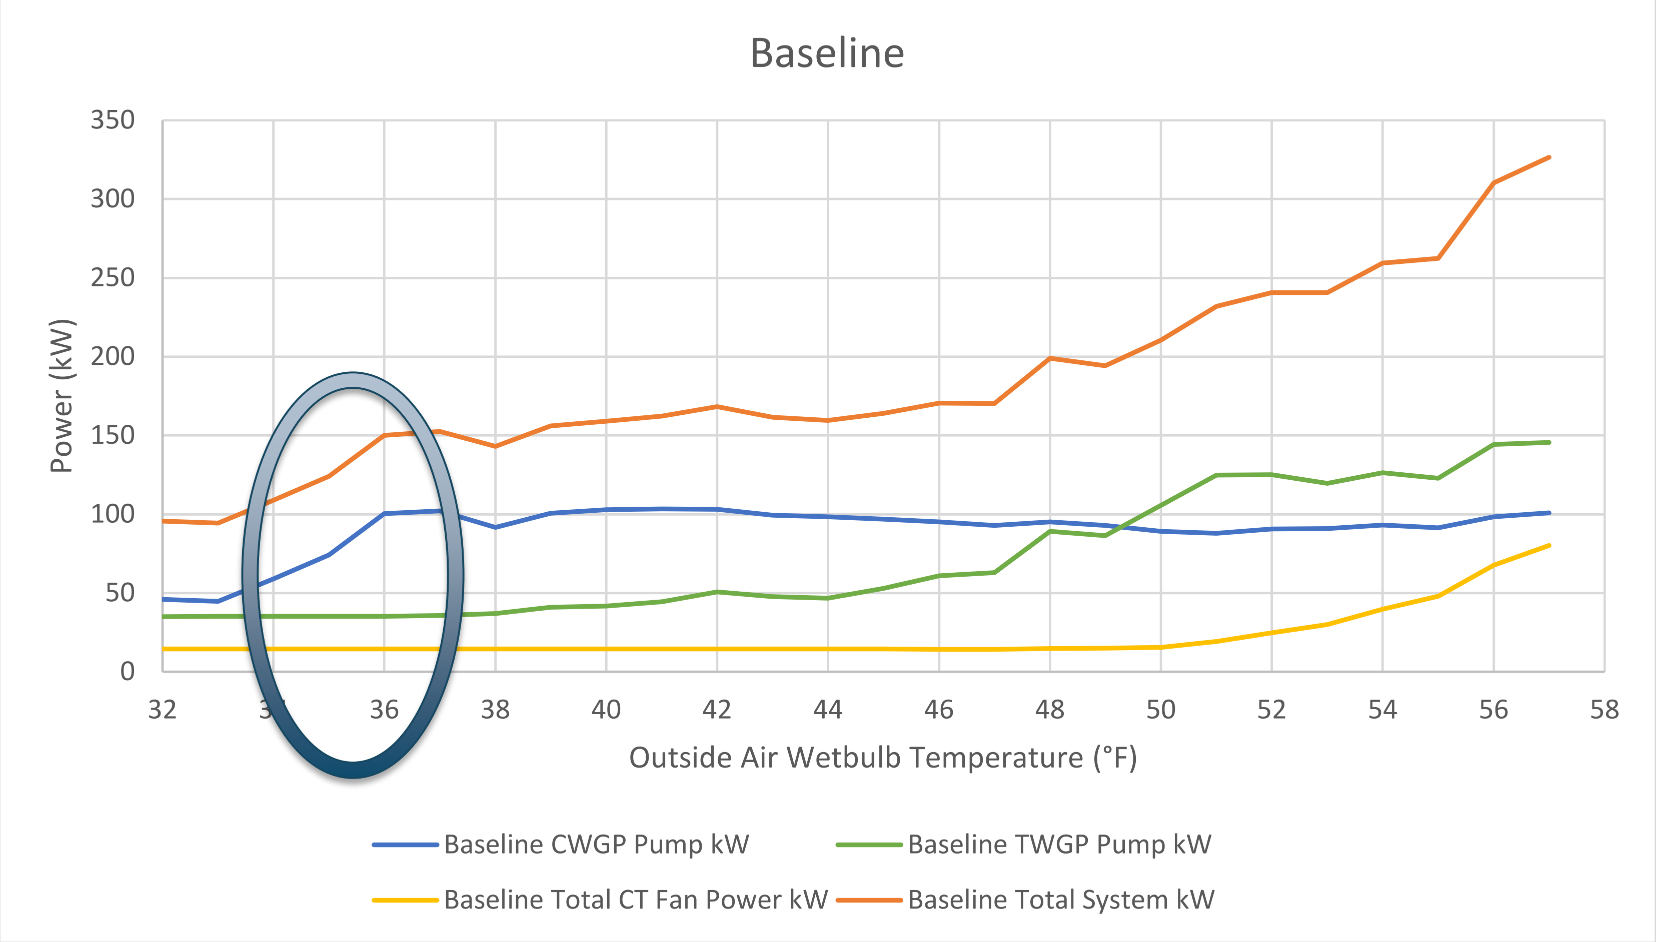

Liquid Cooling Performance Balanced

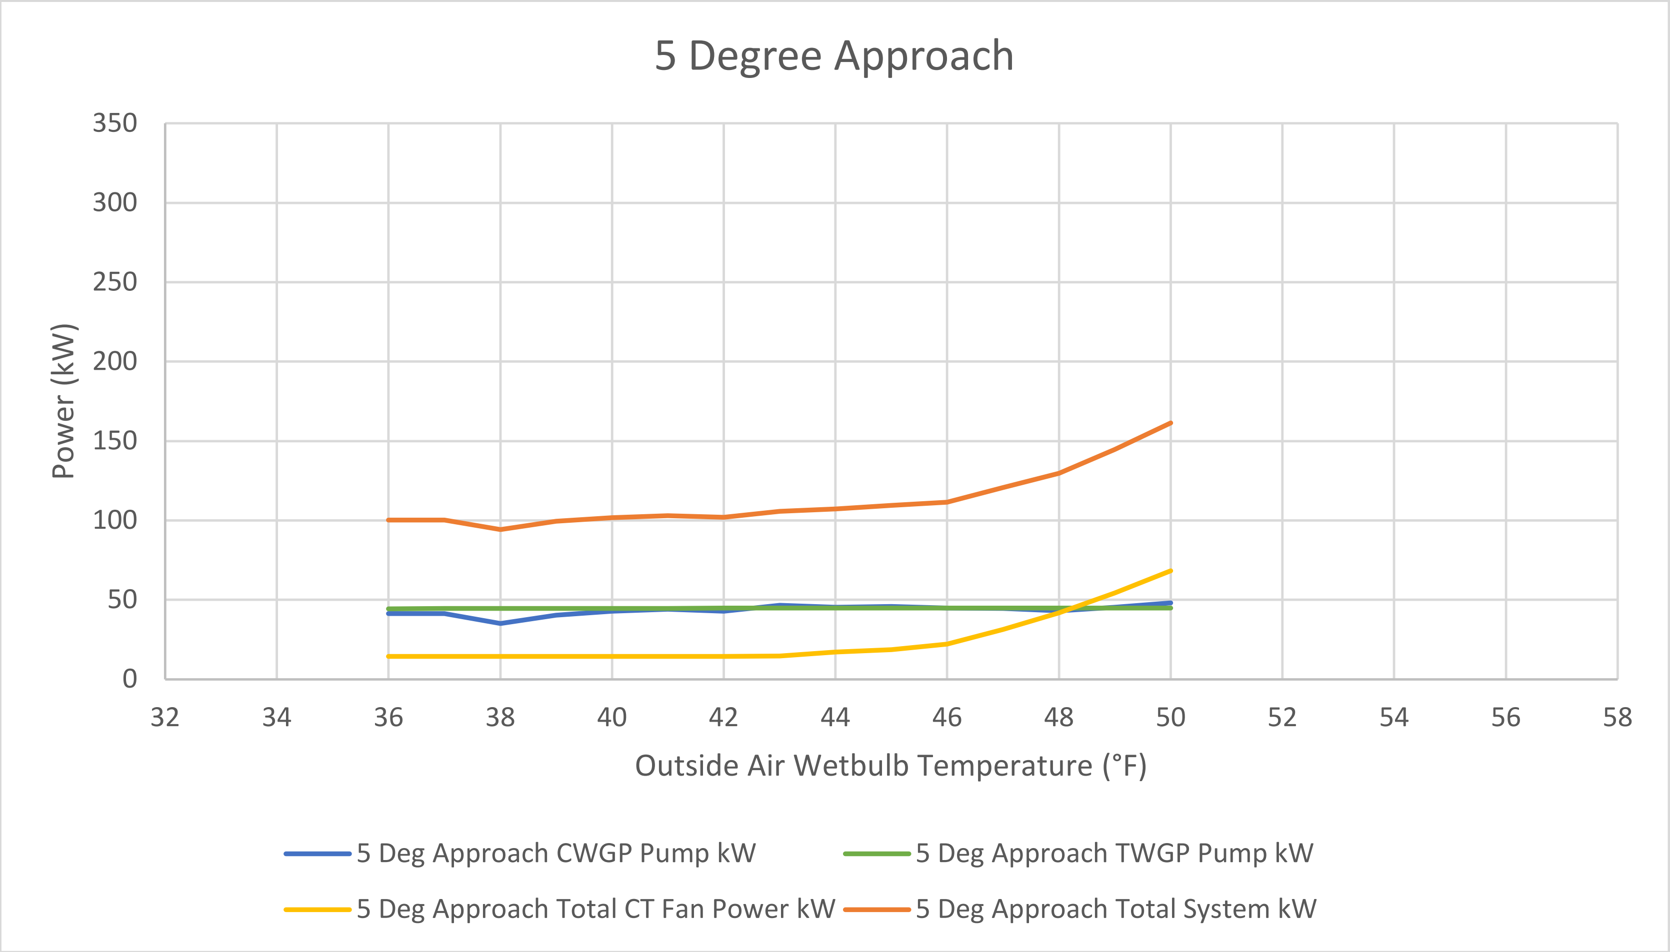

Taking Full Advantage of Recommended Ranges

- Allowable ranges have expanded
- Trade off between high humidity and high temperature
  - ➢ We chose high humidity
- Allowing dryer air would save water

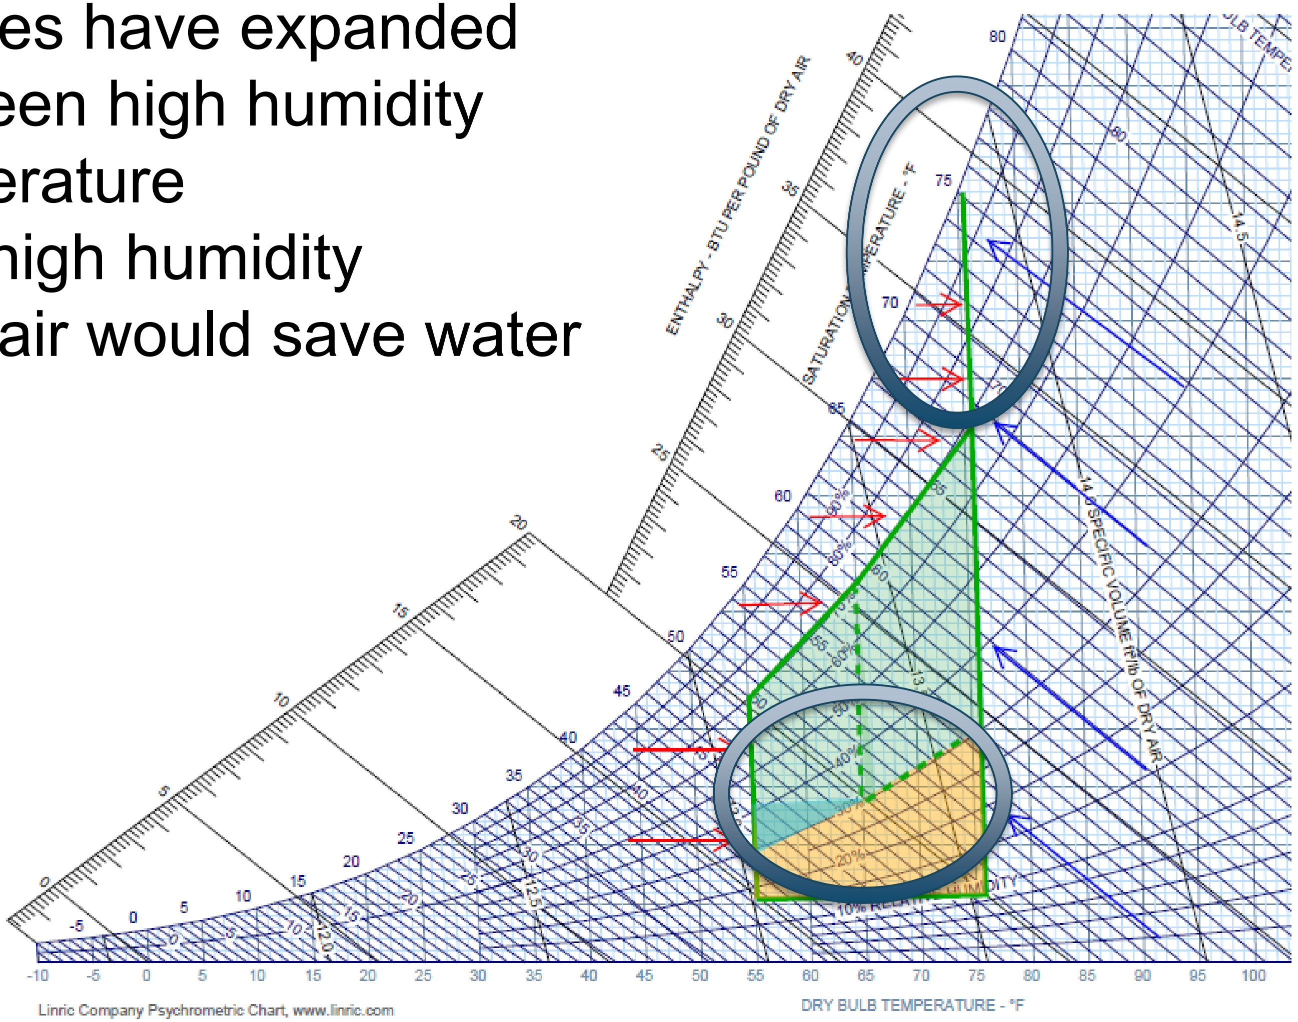

Cooling System Optimization

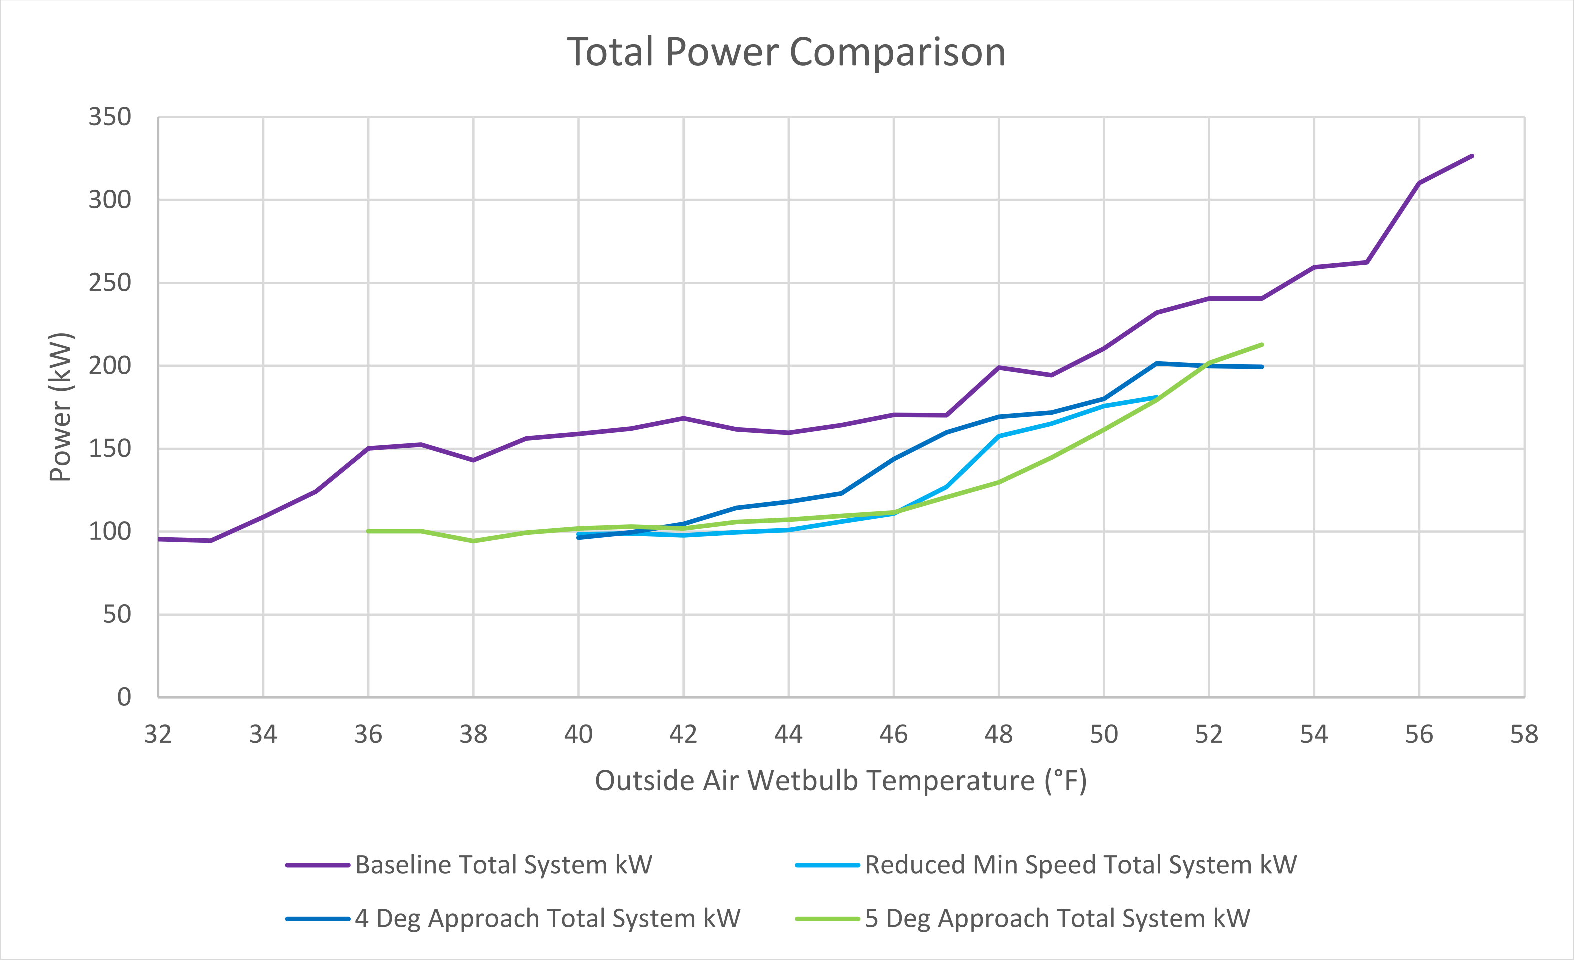

UPS Double Conversion Losses

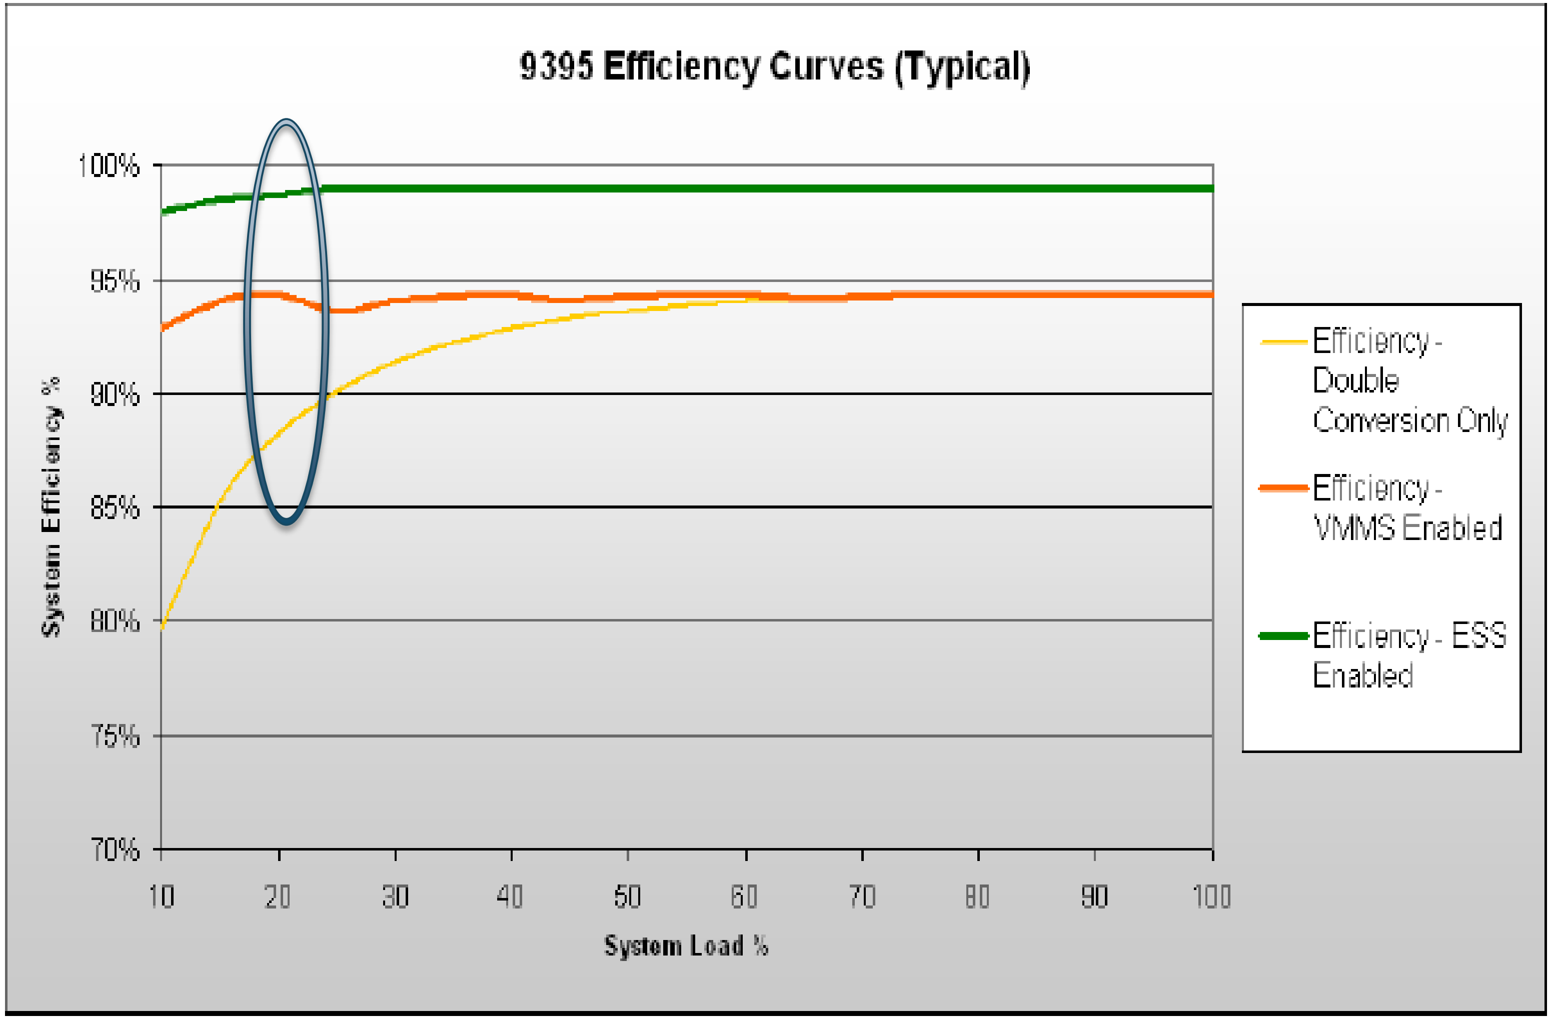

Electricity and Water Savings

	Measure Title	Energy Savings (kWh) Estimated	Energy Savings (kWh) Verified	Water Savings Gallons	Cost Savings $	PUE Reduction
Controls						
1	Optimize Cooling Tower Fan and Pump Controls	-	360,000	100,000	$ 20,880	0.007
2	Optimize Closed Loop Pump Control	240,000	-	110,000	$ 13,920	0.005
3	Optimize AHU SAT and Flow Control	300,000	-	-	$ 17,400	0.006
4	Reset Cooling Water Supply Temperature	600,000	-	220,000	$ 34,800	-
5	Install Firmware to Enable ESS Mode for UPSs	190,000	-	65,000	$ 11,020	0.004
Physical Projects					$ -	
6	Cold Aisle Partial Containment	100,000	-	-	$ 5,800	0.002
	Total	**1,400,000**	**400,000**	**500,000**	**$ 100,000**	**0.025**

IT kWh	48,200,000	*Extrapolated based on typical operation*
Total Non-IT kWh	3,200,000	*Does not include CRAY fans*
PUE	1.07	
Estimated Post-Case PUE	1.04	
Savings as a Fraction of Cooling System kWh	56%	

Our Improvement Team Approach

1. Identify opportunities
2. Identify critical control boundaries
3. Try it out!
4. Verify using trend data
5. Repeat until optimized

All parties work closely throughout the project

Takeaways

- Involve full team from measure development through implementation
- Question everything
- Incremental, iterative process can be most effective
- Don't underestimate interconnectivity of systems
- A sub 1.1 PUE data center can still be improved!

Our Identified Savings are Equivalent to a 1 MW Solar Array

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/