



Recovering Stranded Capacity in the Data Center

Greg Bell

Driving e-Research Across the Pacific `09

Sydney: 11/13/09



AmazonGreen™



dlinkgreen



Green Power Switch®

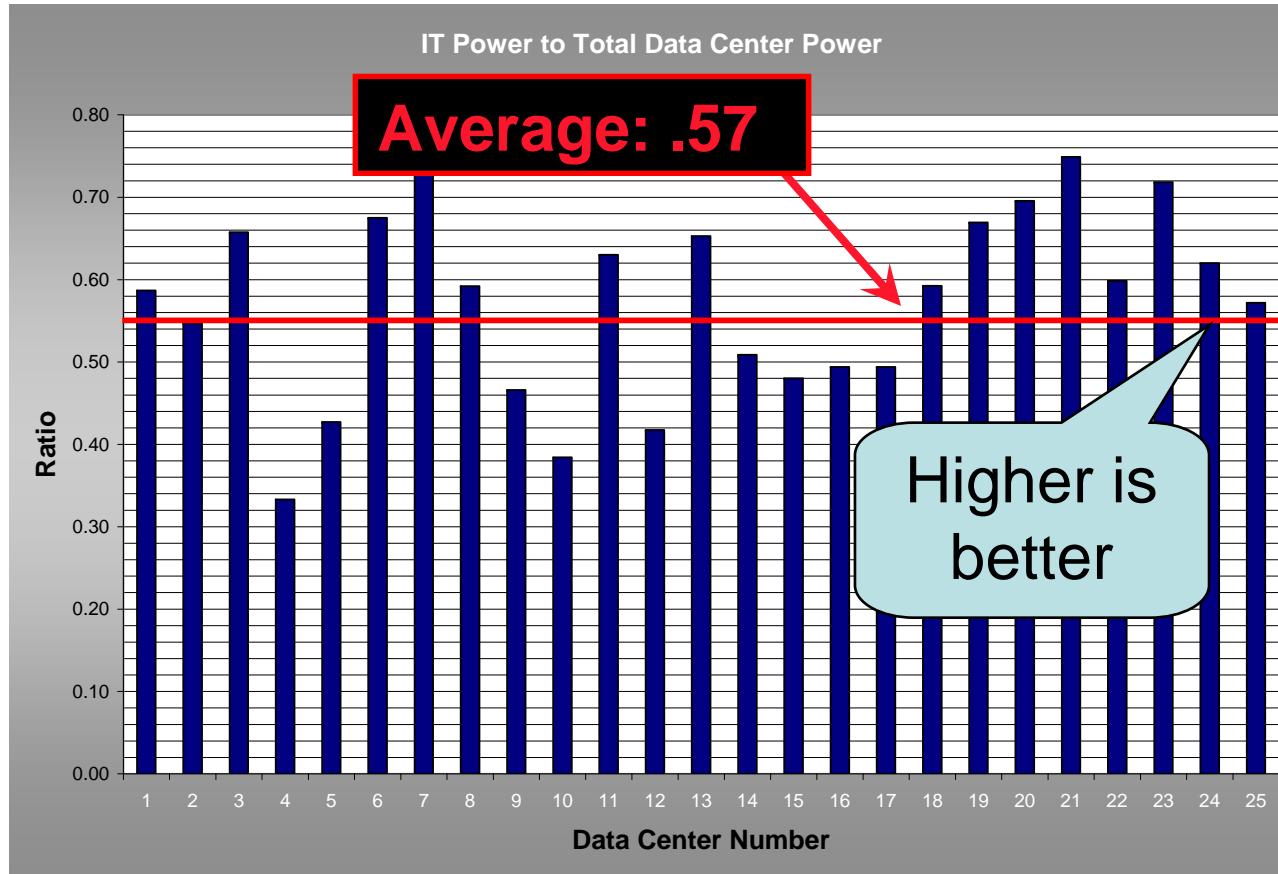


Recovering Stranded Data Center Capacity

- most data centers are overcooled & inefficient
 - waste of energy
 - waste of scientific capacity
- how to recover stranded capacity?
 - opportunities abound
 - we added 100kW load in a legacy data center, without new cooling infrastructure




Brief Digression: Benchmarking Data Centers



Source: LBNL Benchmarking / Dale Sartor

Brief Digression: Efficiency by Power Density



Source: LBNL Benchmarking / Dale Sartor




Smaller Rooms & Closets Grossly Inefficient

	Closet	Growing High Density	Mini Datacenter	Lights Out Cinderblock	High Density
PUE	2.14	1.78	2.94	1.70	1.27
DCiE (IT Load)	47%	56%	34%	59%	79%
UPS/PDU	No / No	No / No	12% / 12%	Redundant	30% / 70%
Normalized PUE	2.36	2.00	3.14	1.70	1.42
Normalized DCiE	41%	50%	32%	59%	70%

source data: Joyce Dickerson, Stanford University [Oct 09]

Strategy #1: Collect Environmental Data


- wireless sensor network
 - 700 monitoring points
 - temperature, humidity, pressure, current
- data center no longer a mysterious black box
- quantified understanding of environmental conditions
- we can perform experiments, monitor results, iterate and tune



\$10 Solutions (or, Data Center Truisms)

- what effect does adding a blanking panel have in a cluster?
 - takes 5 minutes
 - costs \$10
- input air temperatures can drop $\sim 20^{\circ}\text{F}$

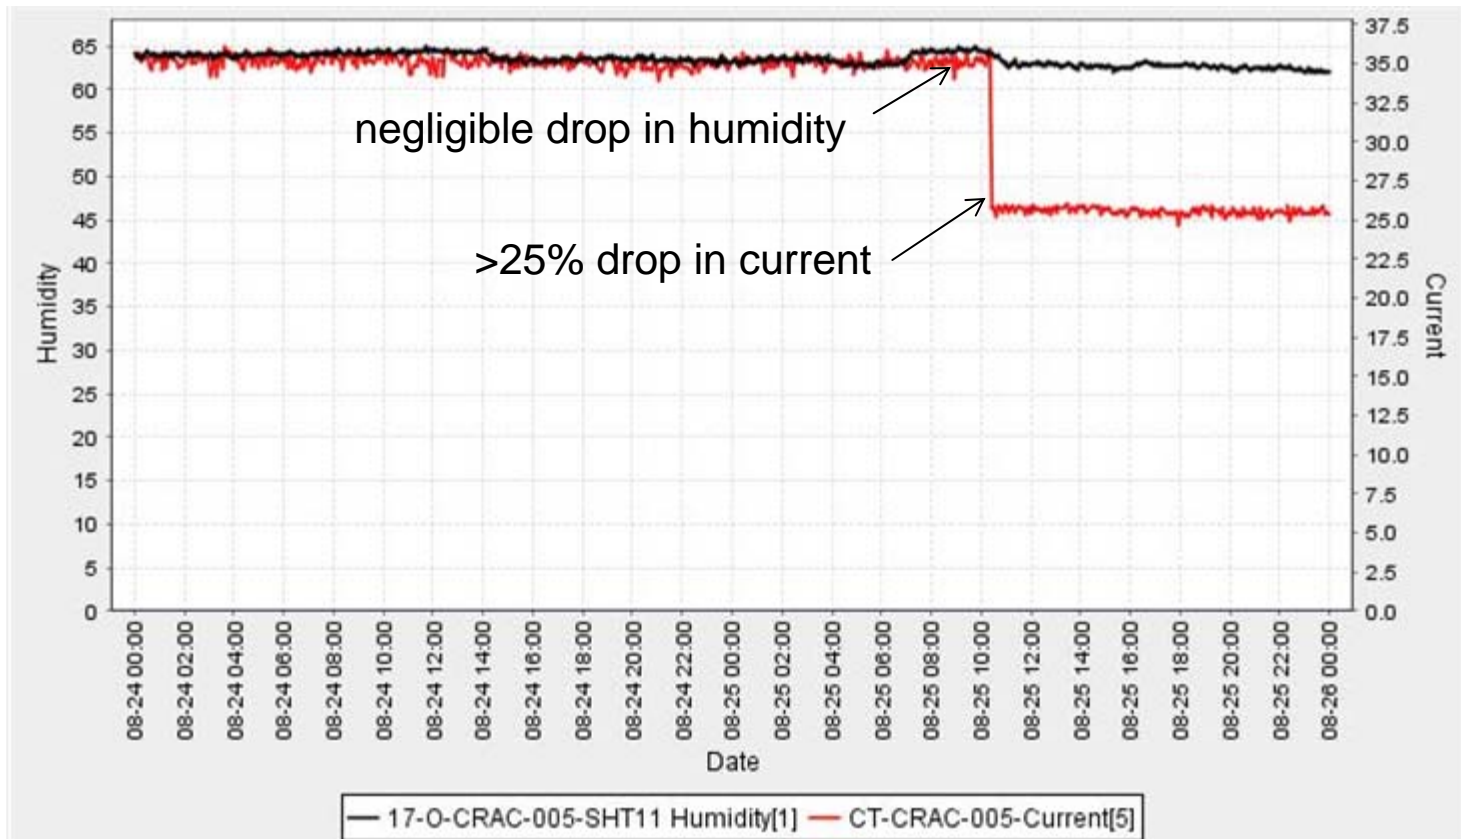
Top of rack



Middle of rack




Do We Need Humidification Control?



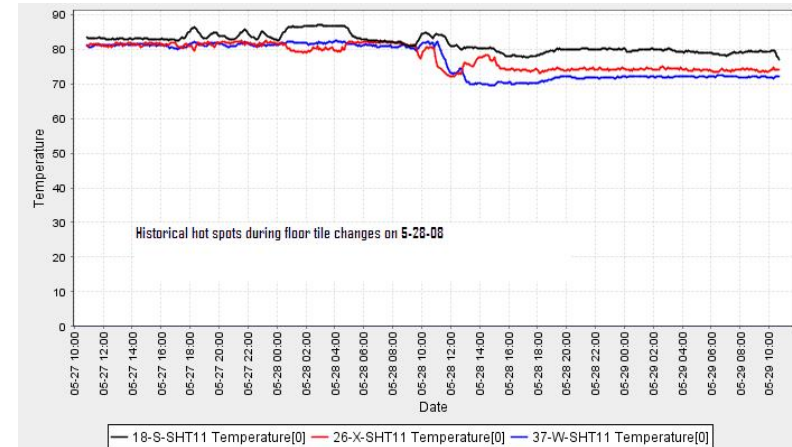
Tuning the Room

- we removed 28 permeable floor tiles
- when airflow is restricted
 - under-floor pressures go up
 - rack-top temps go down
- without monitoring & visualization, this process is guesswork
 - how many tiles to remove?
 - are we overheating our clusters?

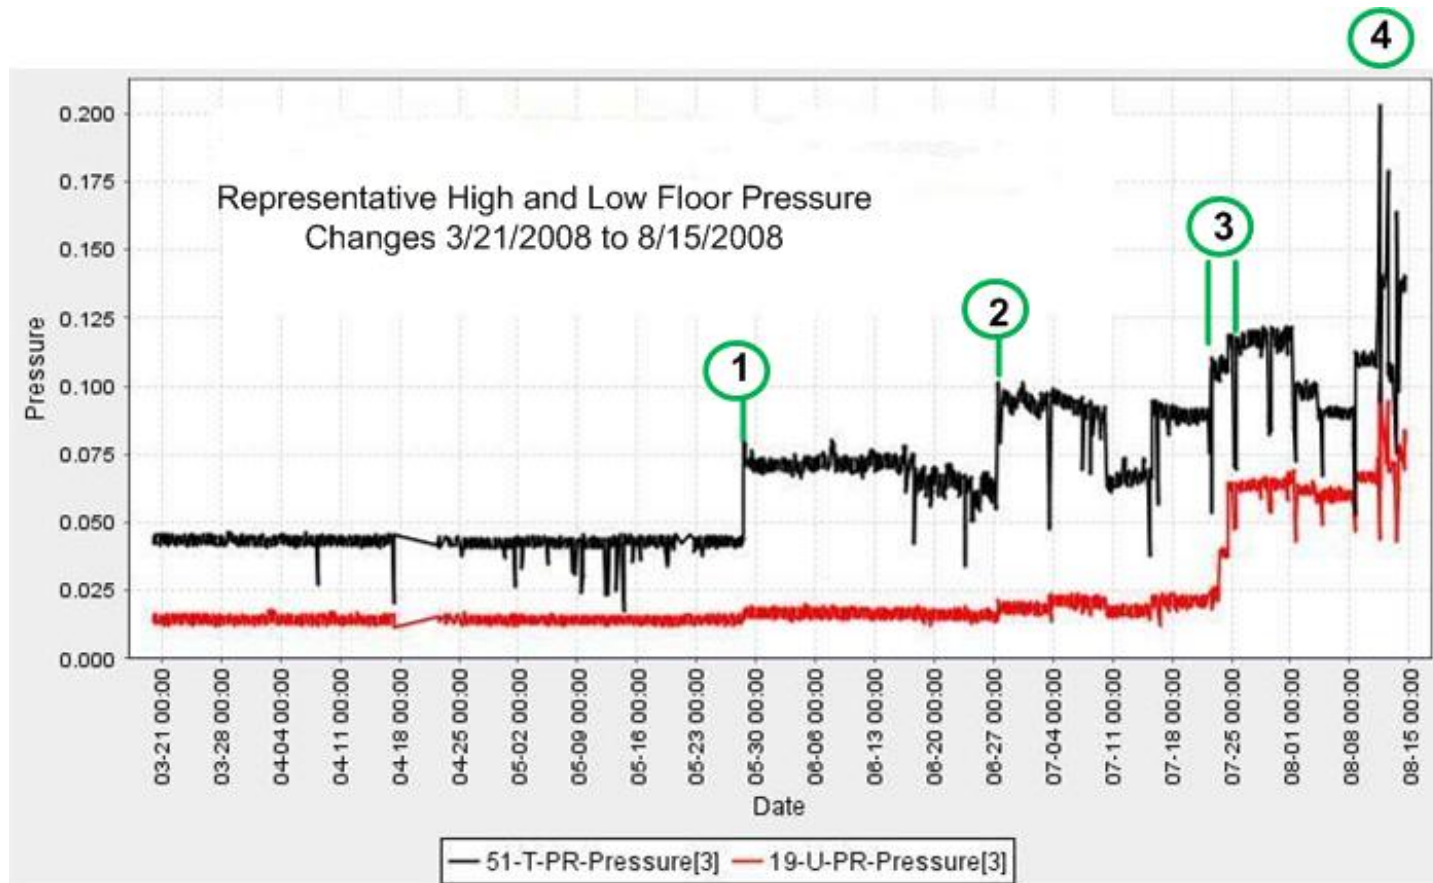
Under-Floor Pressure



Rack-Top Temperatures

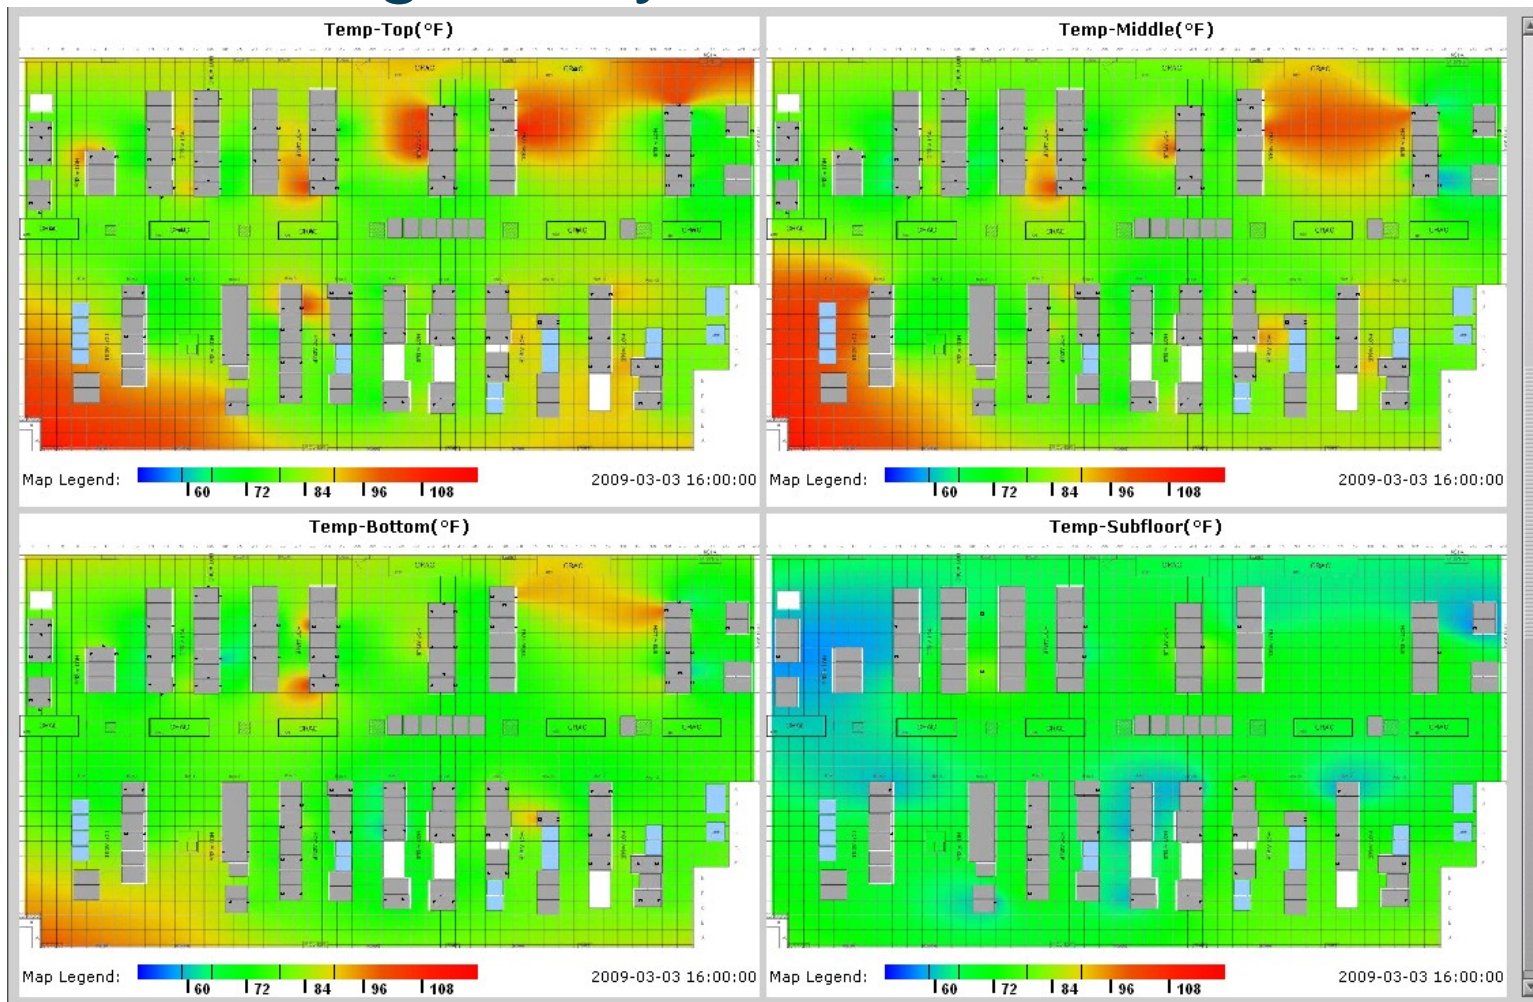


Comparing the Impact of Changes



1, 3 = floor tile removal; 2= plenum conversion;
 4 = CRAC maintenance [belt replacement; de-icing condenser pan]

Visualizing in Layers





Ultimate Goal: Efficiency Dashboard

- real-time calculation of DCiE
 - $DCiE = IT \text{ power} / \text{total power}$
 - integrated with our wireless monitoring system
 - much easier said than done
- our monitoring plan includes 45 measurement points
 - if you're building a new data center, include instrumentation for real-time efficiency calculation from the start



Strategy #2: Better Air Flow Management

- blanking panels & other cheap airflow gadgets
- transparent curtains – easy & flexible
- consider fire code first
- issues: smoke detection, gaseous fire suppression, behavior of curtains during fire



Curtains for Air Flow Management



Strategy #3: Liquid Cooling

- air is a poor, though convenient, conductor of heat
- water already in cooling loop
 - how to bring it closer to the CPUs?
- most solutions unfortunately require very cold water
 - on the horizon, servers that tolerate warmer air intake
 - opens up many possibilities: outside air for cooling, reuse of waste heat
- how to choose among liquid cooling products?
 - ‘Chill-Off’ competition (Sun/LBNL/SVLG)
 - in second year, with new results in March



Liquid Cooled Rack Doors in Place



plumbing for liquid-cooled doors

60kW Cluster with Liquid Cooled Doors




- water cooling without chillers: 74F intake water

Another Solution


- Clustered Systems
- modified servers (fans & heat sinks replaced with heat risers)
 - ‘net cost is negative’
- modified racks (pressure plates, water circulation)
 - design lifetime = 10 years
- claimed 2-3 yr payback
- Chill Off testing now underway



Heat Risers for CPU, Memory



Rack as Liquid-Cooled Heat Sink





Backing Up: Rhetorical Strategies for Motivating Such Work

- “we’ll save electricity”
 - may not be sufficient by itself
- “we’ll increase computing capacity while using existing mechanical infrastructure”
 - moderate investment x products capacity increase y
- “we’ll postpone a massive data center expense”
 - cost avoidance also persuasive





Summary

- most data centers & ‘server closets’ inefficient
- opportunities for improvement abound
- instrumentation helpful, not mandatory
- largest benefits to Berkeley Lab so far
 - cost avoidance
 - recovered capacity for science
- ...but the strategy doesn’t work forever
 - our next project: instrumented, chiller-less container
- please engage your infrastructure team on the topic of efficiency!

