Towards understanding and automating data center operations

Greg Ganger

PDL Director

The Data Center Observatory (DCO)

A unique endeavor, simultaneously:

- An operational data center with real users
- A research facility used to:
  - Study real data center operations
  - Evaluate innovations in real deployment

Operational costs out of control

Power and cooling

- Now on par with purchase costs
- Trends making it worse every year
  - Power/heat go up with speed
  - Cluster sizes increase due to commodity pricing

Projected Trend in Rack Power

Timeframe	Avg. Watts / Sq Ft	Avg. kW / rack
*2003	40	2
**2005	80	4
**2007	240	15
**2010	500	30

* Per surveys by 7x24 Exchange, Lawrence Berkley Labs, Uptime Institute

** Estimates based on gradual implementation of new server technology that is either already announced or currently on the market

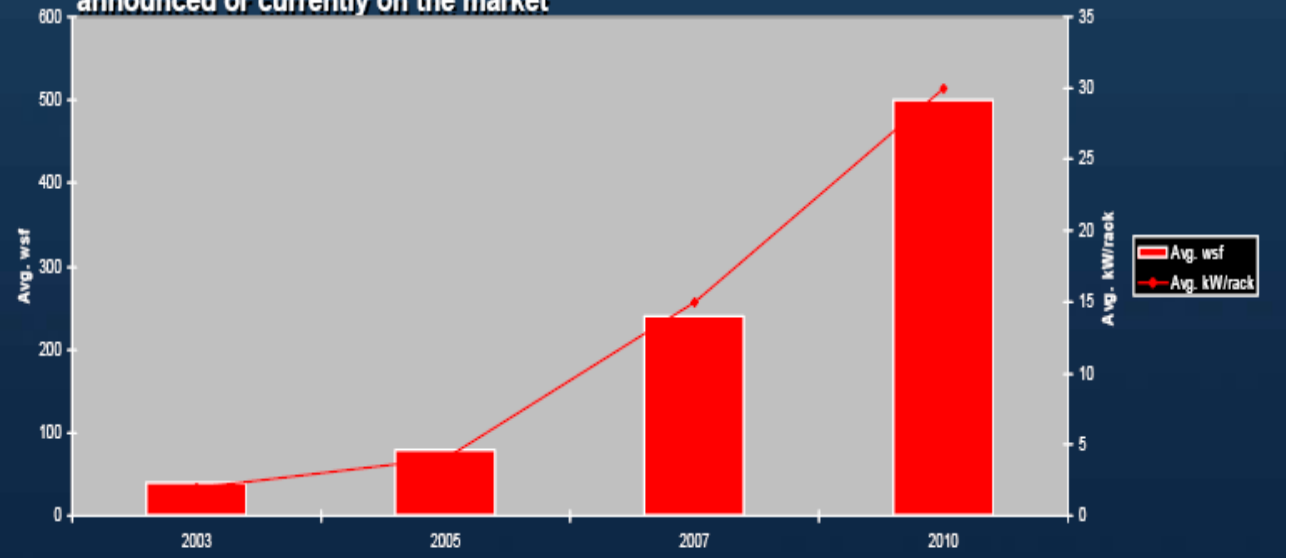

Source: "Data Centers of the Past, Present and Future"
California Data Center Design Group

THE WALL STREET JOURNAL.

DOWJONES MONDAY, NOVEMBER 14, 2005 - VOL. CCXLVI NO. 104 - ★★★★ $1.00 WSJ.com

Power-Hungry Computers Put Data Centers in Bind

Newer Hardware Guzzles Electricity and Spews Heat, Requiring Costly Alterations

By Don Clark

The University at Buffalo installed a $2.3 million Dell Inc. supercomputer last summer, hoping to bolster its image as a research institution. Instead, the big machine came to symbolize an increasingly vexing problem for data centers world-wide.

Once the machine was delivered, university officials discovered they had only enough electrical power to switch on two-thirds of the system. They have temporarily responded by throttling back use of an older supercomputer, but a $20,000 electrical-system upgrade will be needed to run both systems at full capacity.

"The calculations that were done fell slightly short," says Bruce Holm, a senior vice provost at the school, which is part of the State University of New York. "The bottom line was that they missed."

More such misses are likely. That's because, in its long-running race to boost performance, the computer industry has hit a major hurdle: The newest hardware—particularly the servers that run most business programs and Web sites—draws too much electricity and generates too much heat.

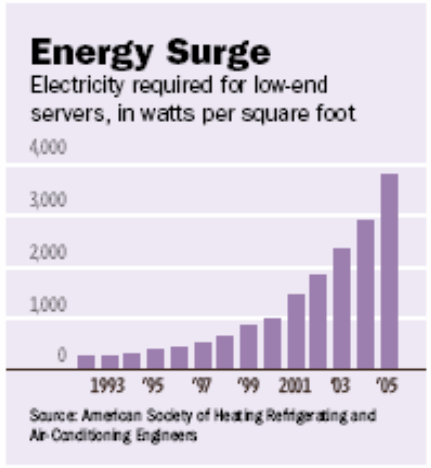

The power-hungry machines, along with rising energy prices, are generating enormous utility bills and forcing changes on Silicon Valley technology suppliers that are akin to Detroit's struggle to improve gas mileage. (See related article on page B4.)

Though more-energy-efficient computers are on the way, it could be years before companies replace the systems they have already purchased.

In the meantime, bringing in more electricity and cooling is expensive and difficult in some data-center buildings. Organizations face unpleasant choices that include building new facilities, putting off server purchases or leaving costly space in computer rooms unoccupied to avoid overwhelming their air-conditioning systems.

Facilities planners at the University at Buffalo, for example, originally erred because they thought an older supercomputer would no longer be needed by the time their new machine arrived, Mr. Holm says. The need for both systems caused the university to consider spending as much as $150,000 to upgrade the current data center's air conditioning, just as the university was on the verge of moving the systems to a more modern computer room. "It's that kind of juggling act," Mr. Holm says.

If planners miscalculate, servers overheat, damaging circuitry or causing shutdowns that disrupt operations. The Uptime Institute, an organization that represents data-center managers, predicts

Please Turn to Page A16, Column 1

Operational costs out of control

Power and cooling

- Now on par with purchase costs
- Trends making it worse every year
  - Power/heat go up with speed
  - Cluster sizes increase due to commodity pricing

Administration costs

- Often reported at 4-7X capital expenditures
- Trends making it worse every year
  - Complexity goes up with features, expectations and cluster size
  - Salaries go up while equipment costs go down

Where does administration effort go?

Shockingly little data available

- From IT organizations or system providers

Tough to solve an undefined problem

- Need to understand the challenges
  - and their relative sizes
- Need a way to evaluate solutions
  - to produce convincing quantitative evidence

As an observatory, the DCO provides

Views into a real data center's operations

- Sensors throughout physical environment
- Instrumented software systems
- Records of all administrator time

A testbed for new technologies

- Evaluate effectiveness of tools directly
  - on real situations in a real data center environment
- Evaluate effectiveness of deployed systems

As a data center, the DCO can

Enable researchers to focus on their work

- Instead of management of private infrastructures

Help Carnegie Mellon control its IT costs

- Aggregation reduces waste
- Power and cooling costs are a problem here too
- Administration is a problem here too

DCO on lobby level of the CIC

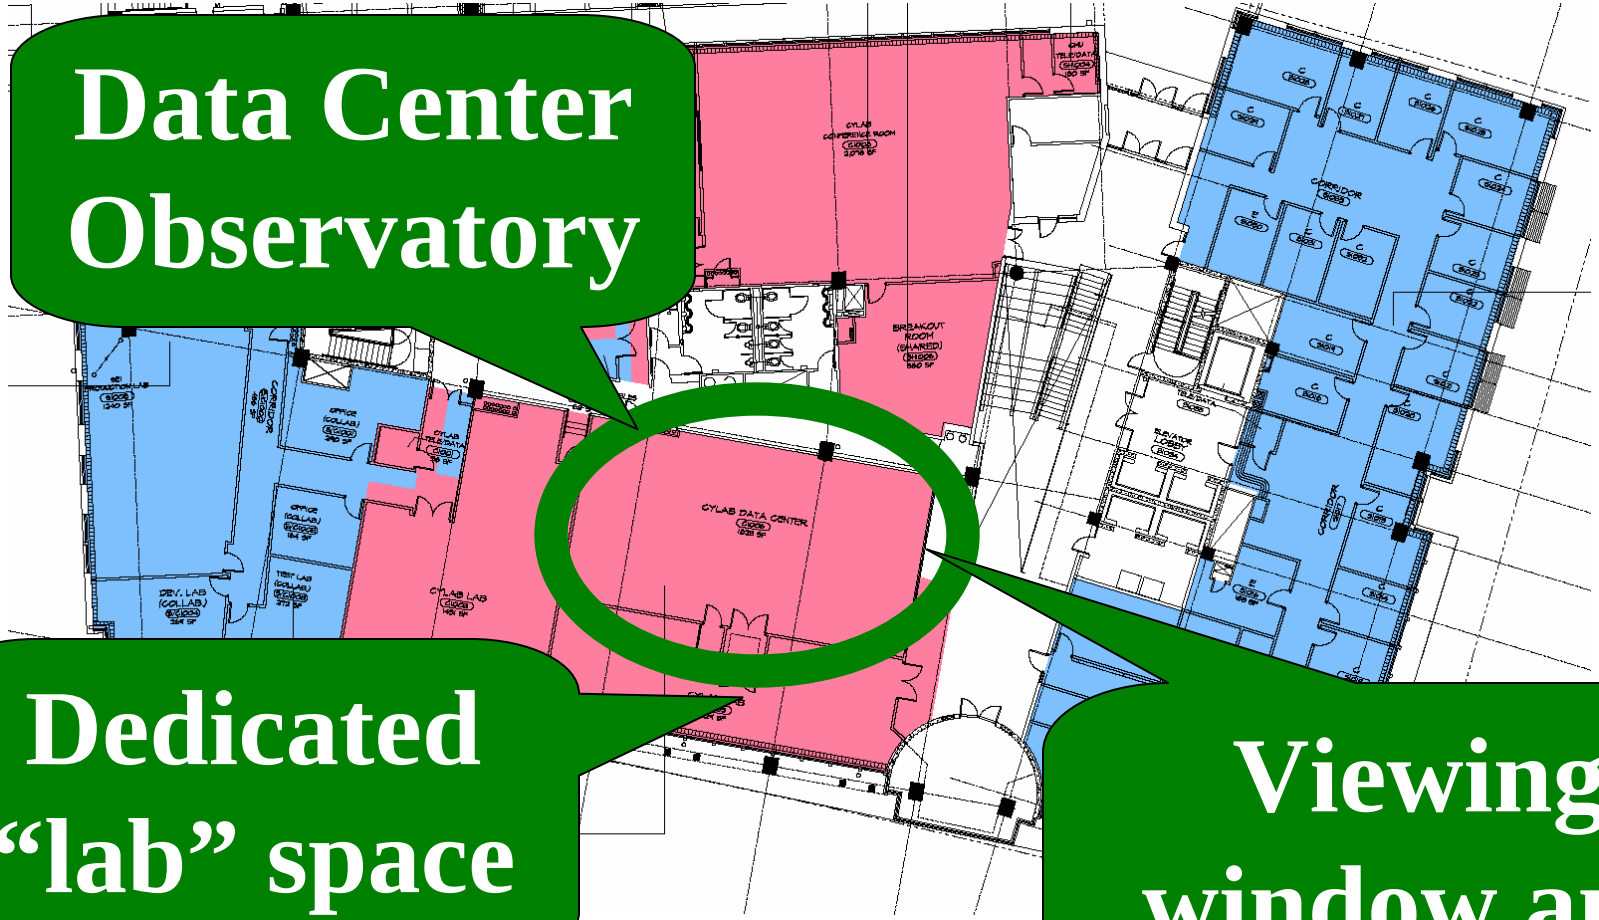

Partnering with APC for

Design and show-casing of forward-looking cooling technology

- Based on their novel hot aisle containment and “pay as you grow” concepts

Data collection and research

- Documenting power/cooling costs and developing novel approaches to reducing both
- Automating dynamic power and thermal management

First “zone” in place now

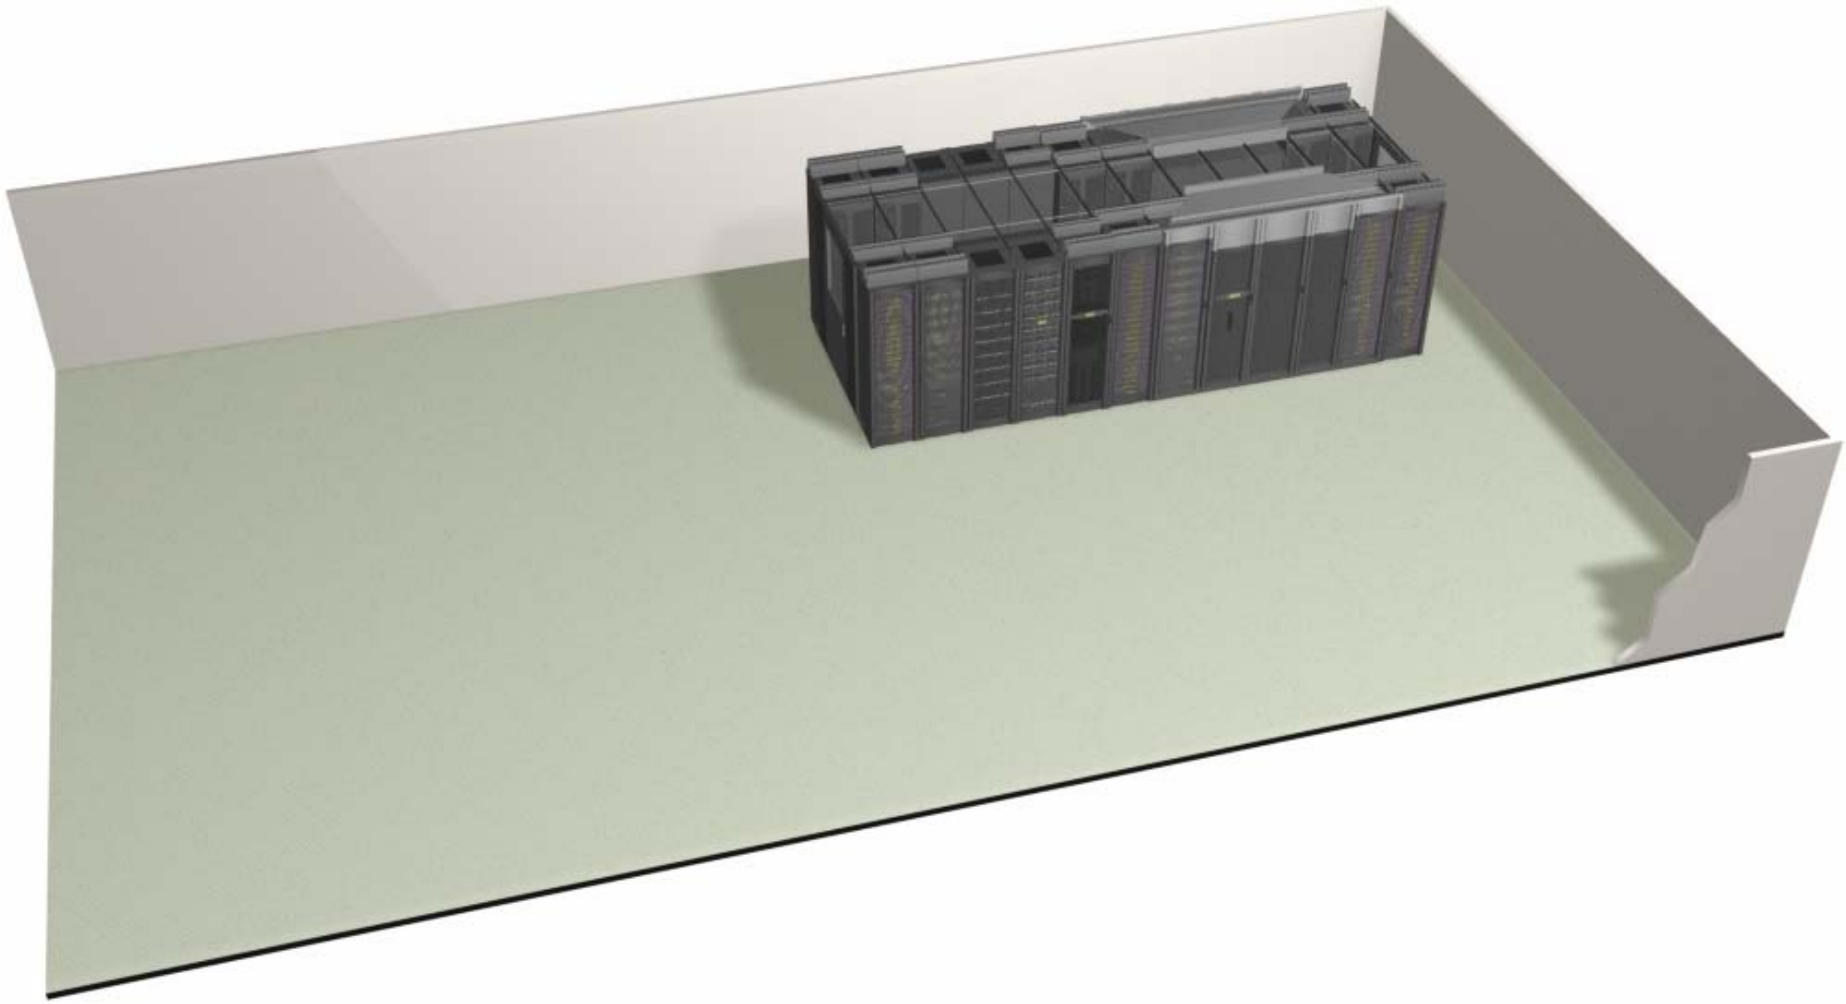

Will grow by adding additional zones

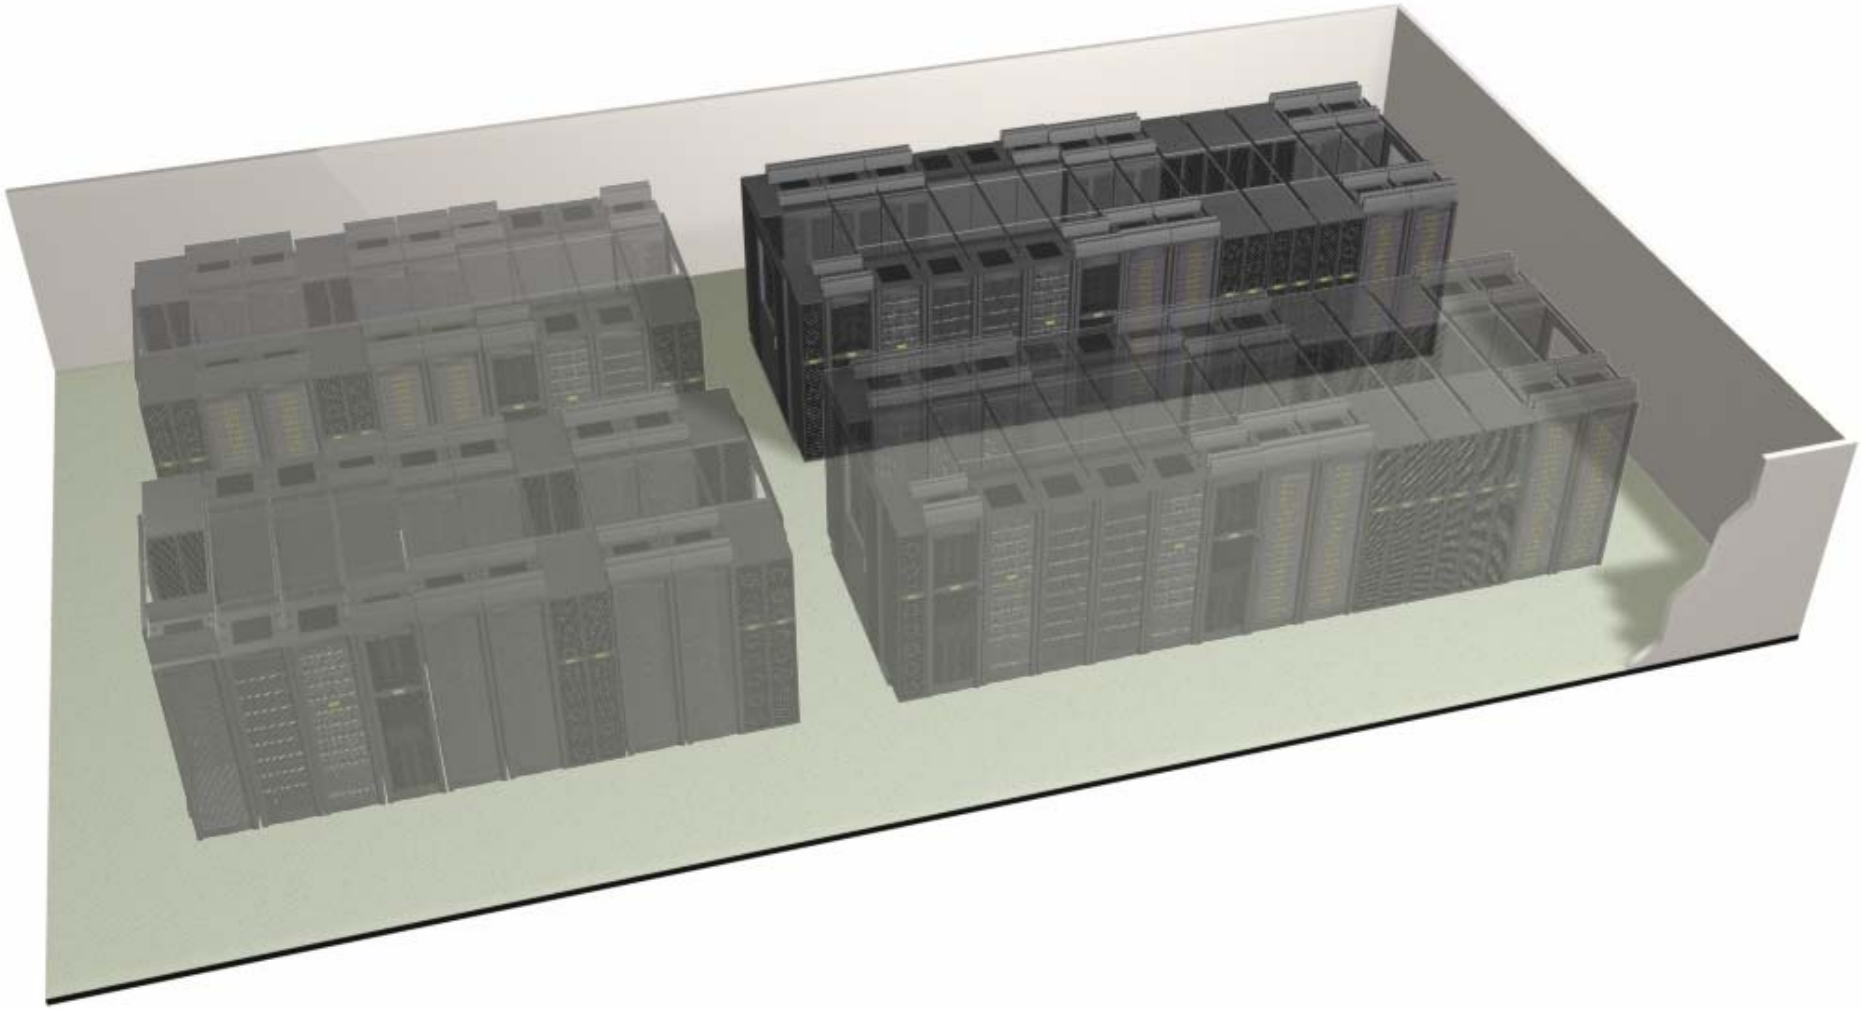

Some DCO details

Compute and storage will grow with need

- Target: ~40 racks of servers @ 20kw/rack
  - over a petabyte of storage
  - over 4,000 GHz of CPU

Machine room is approximately 2,000 ft^2

- Adjoining lab is approximately 1,100 ft^2

Highly-visible showcase

- Windowed wall on high-traffic hallway
- LCD display for telling story and showing stats

Support on many fronts

University

- Administration and facilities management

Funding for research

- NSF, DARPA, AFOSR, and ARO
- PDL and CyLab companies

Early computer equipment donations

- IBM, Intel, Seagate, Sun

Summary: Data Center Observatory

A utility and a research endeavor

- An administered utility
  - computation and storage for CMU researchers
- An observatory and testbed
  - measuring costs/difficulties in a real data center
  - deploying new technologies for reducing them

A major, long-term endeavor

- Construction finished in early April 2006
- First equipment now up and running
- Attack on operational costs now underway!