ORACLE®

Eco-Efficiency Best Practices For Data Centers

John Kapson
Oracle Advanced Customer Services

Oracle ACS Systems Modernization

- The ACS Systems Modernization Practice is a business unit dedicated to helping our customers develop and implement IT strategies and plans that lead to efficient and sustainable data centers.
- The practice focuses on strategy, planning, design, implementation, and operations of large and complex information technology and data center facility infrastructures.
- Service offerings include:
  - Data center strategy
  - Data center design/planning/construction
  - Virtualization and consolidation
  - Data center power and cooling assessments
  - Implementation of power and cooling efficiency technologies and processes
  - Data Center / Operations Management

ORACLE

Agenda

- Today's IT Challenges
- Data Center Ecology and Economics
- The Power and Cooling Challenge
- Modularity in Data Center Design
- Next Steps Towards Eco-Efficiency

What is Data Center Efficiency?

- Managing the use of floor space, power and cooling through sustainable modular design and active management of the information technology and data center facility infrastructure.
- The goal is improved performance as measured by increases in the work unit of computing per watt of power consumed.
- Improvements start at the application level and continue through the entire facility infrastructure.


Eco – It's About
Economics,
Ecology
AND
the Ecosystem

Today's Global IT Challenges

- Under-Utilized, Inefficient Systems are Limiting Business Growth
- Disaster Recovery Planning is Non-Negotiable
- Power and Cooling Constraints are Very Real Issues
- Energy Costs are Draining the Bottom Line
- Sprawling IT Infrastructure is Increasingly Hard to Manage
- Ability to Deploy New Services is Critical to Remain Competitive

What's Driving Infrastructure Demand?

New Consumers.
New Content.

On the
Network...

Drives
Infrastructure
Demand.

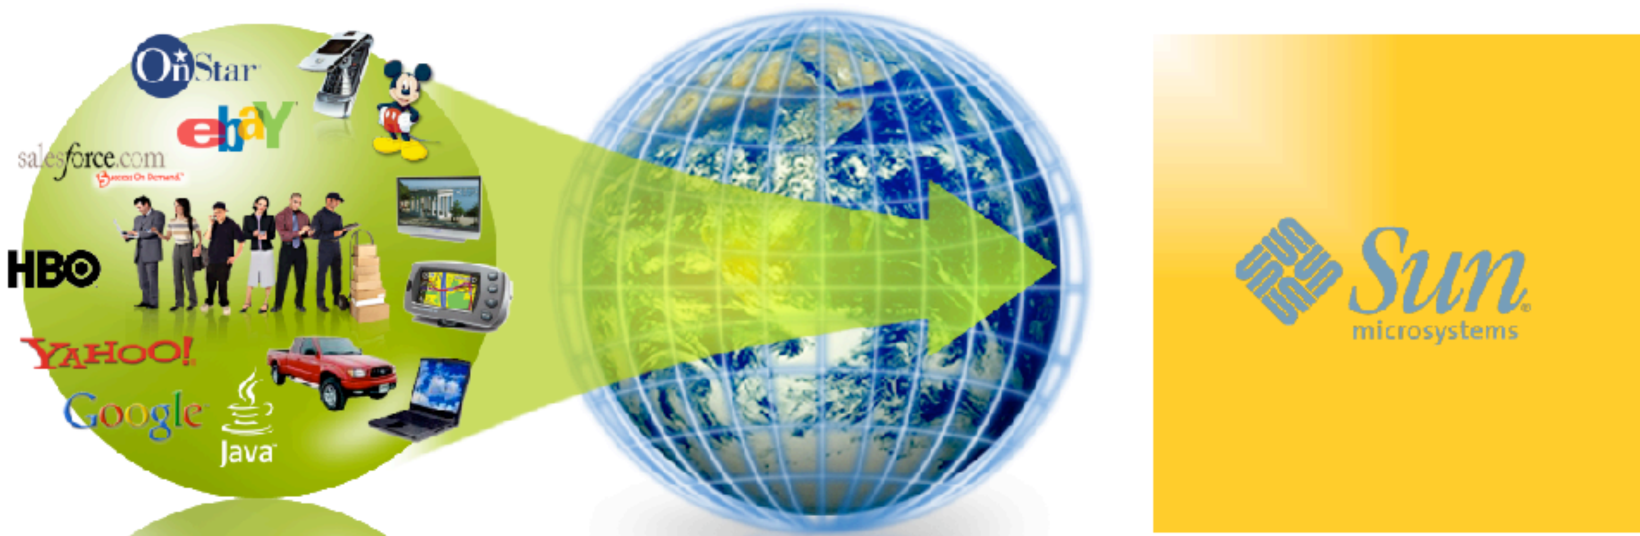

New Devices.
New Services.

Sun Network
Computing Infrastructure

Server Opex will soon Exceed Capex

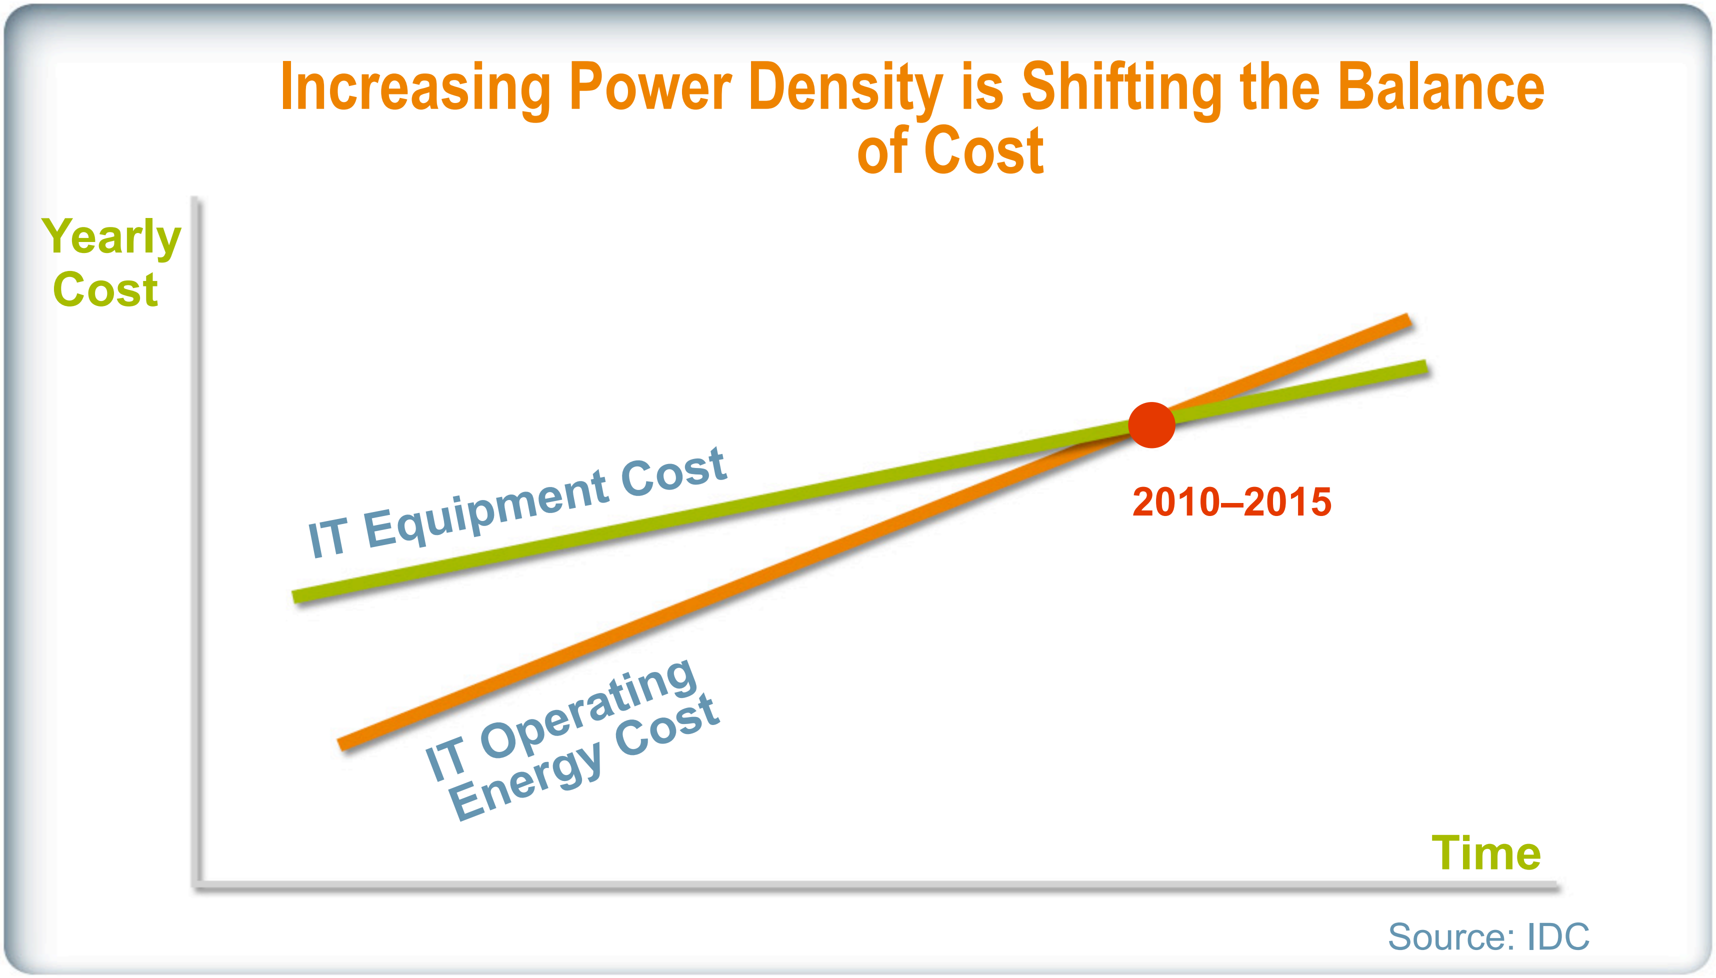

Rack Densities Continue to Increase…

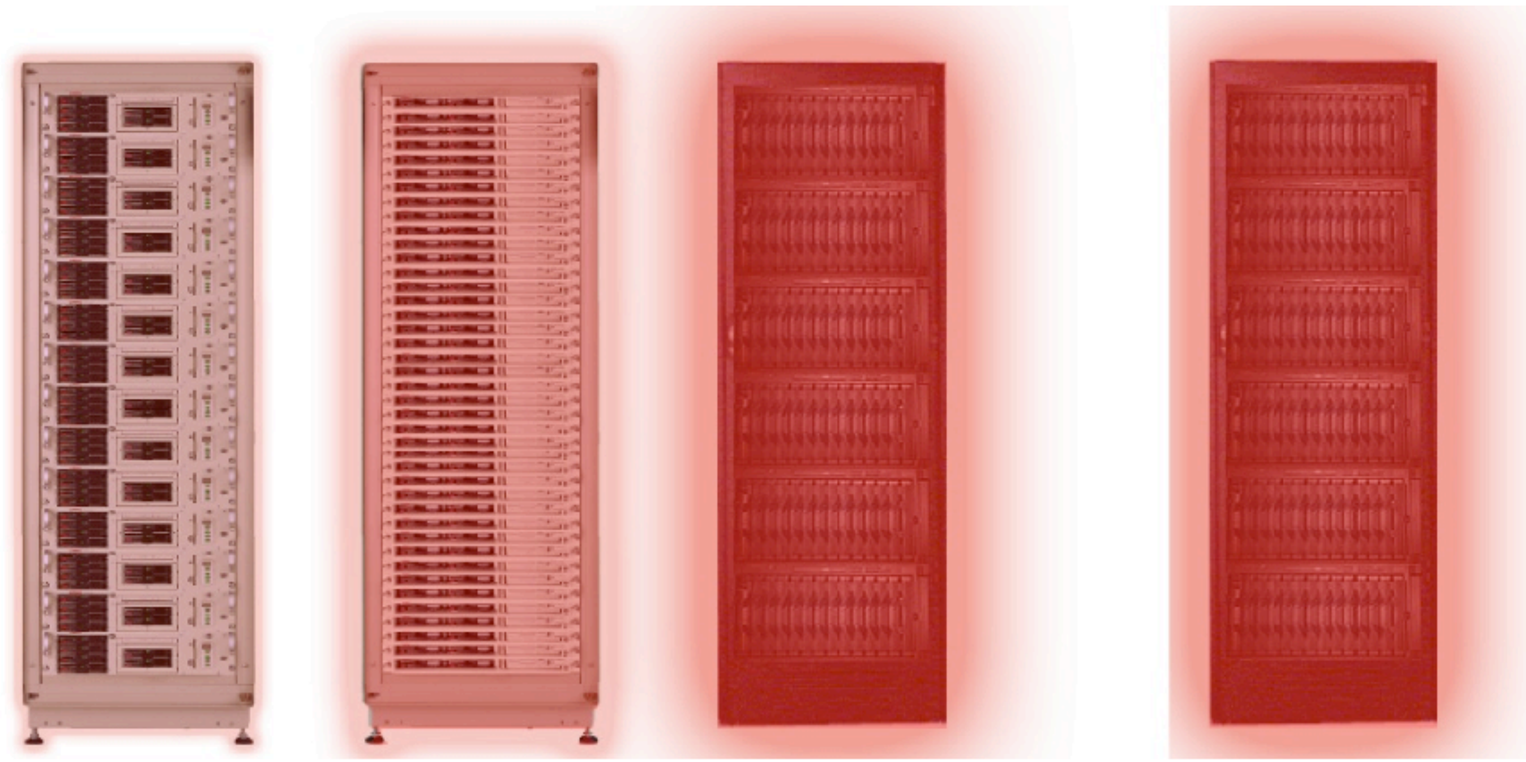

2000	2002	2005	2008
28x2U Servers	42x1U Servers	6 BladeCenters	6 BladeCenters
2kW Heat Load	6kW Heat Load	24kW Heat Load	30kW Heat Load

Source: Emerson Network Power / Liebert

Power Capacity Limits

Cooling Capacity Limits

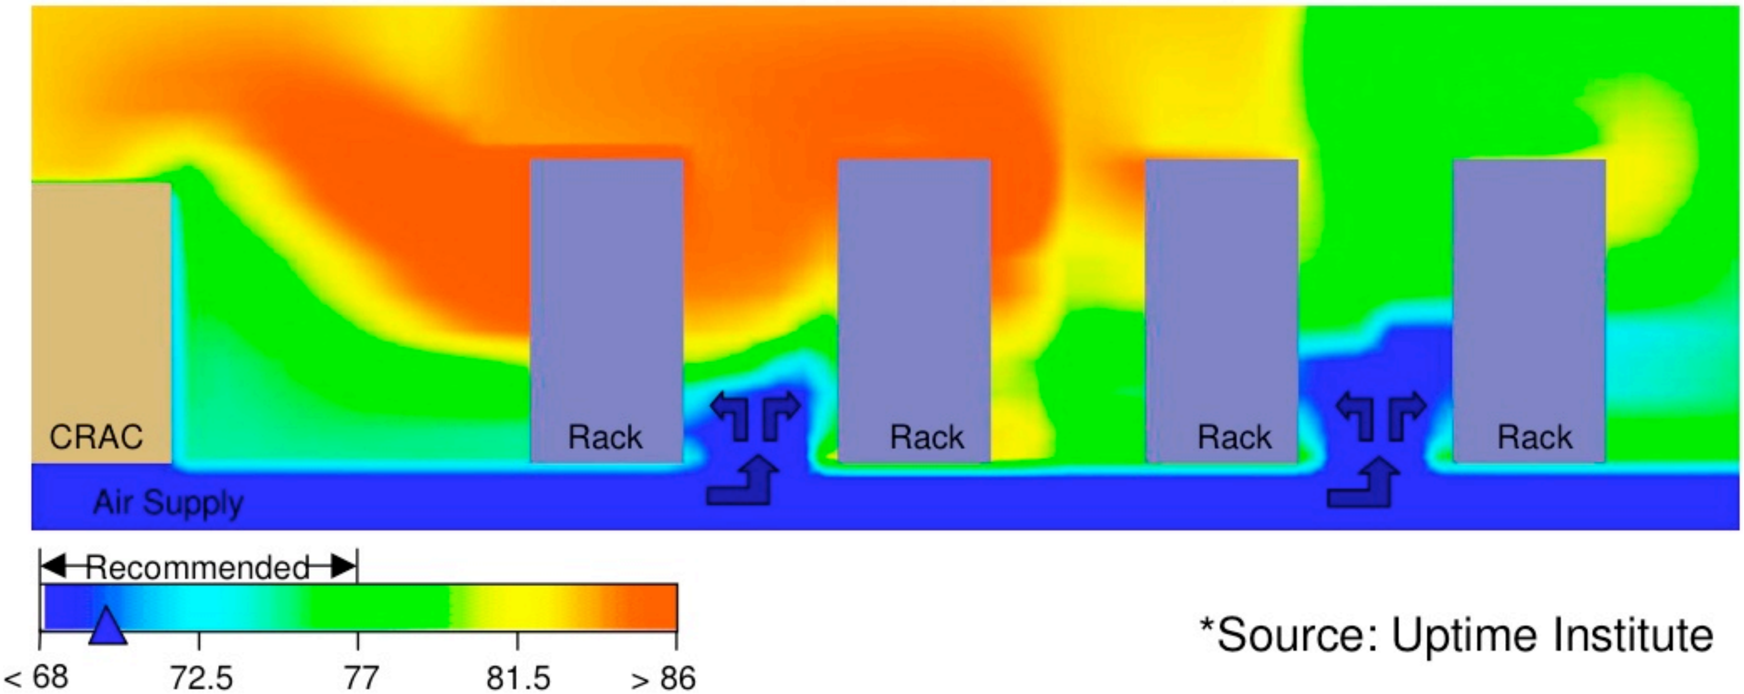

*Source: Uptime Institute

Impacts of the Challenges

- A “Set it and Forget it” IT hardware deployment strategy is no longer an option.
- Coherent design and planning must incorporate both the the IT and facilities infrastructure.
- Tools and processes must be used offset the increased complexity associated with running a denser and more efficient data center.
- Planning focus will be on the efficient use of all IT resources.
- Active monitoring and capacity management will apply to both the IT utilization and power/cooling.

“Would you tell me, please, which way I ought to go from here?

That depends a good deal on where you want to get to,” said the Cat.

-- Lewis Carroll

Ecology and Economics

Data Center Energy Productivity

- For each 1,000 Watts of electricity generated, about 450 Watts is delivered to the servers
- Of this, about 67 Watts (15%) is used to run server workloads - an overall energy productivity of about 7%

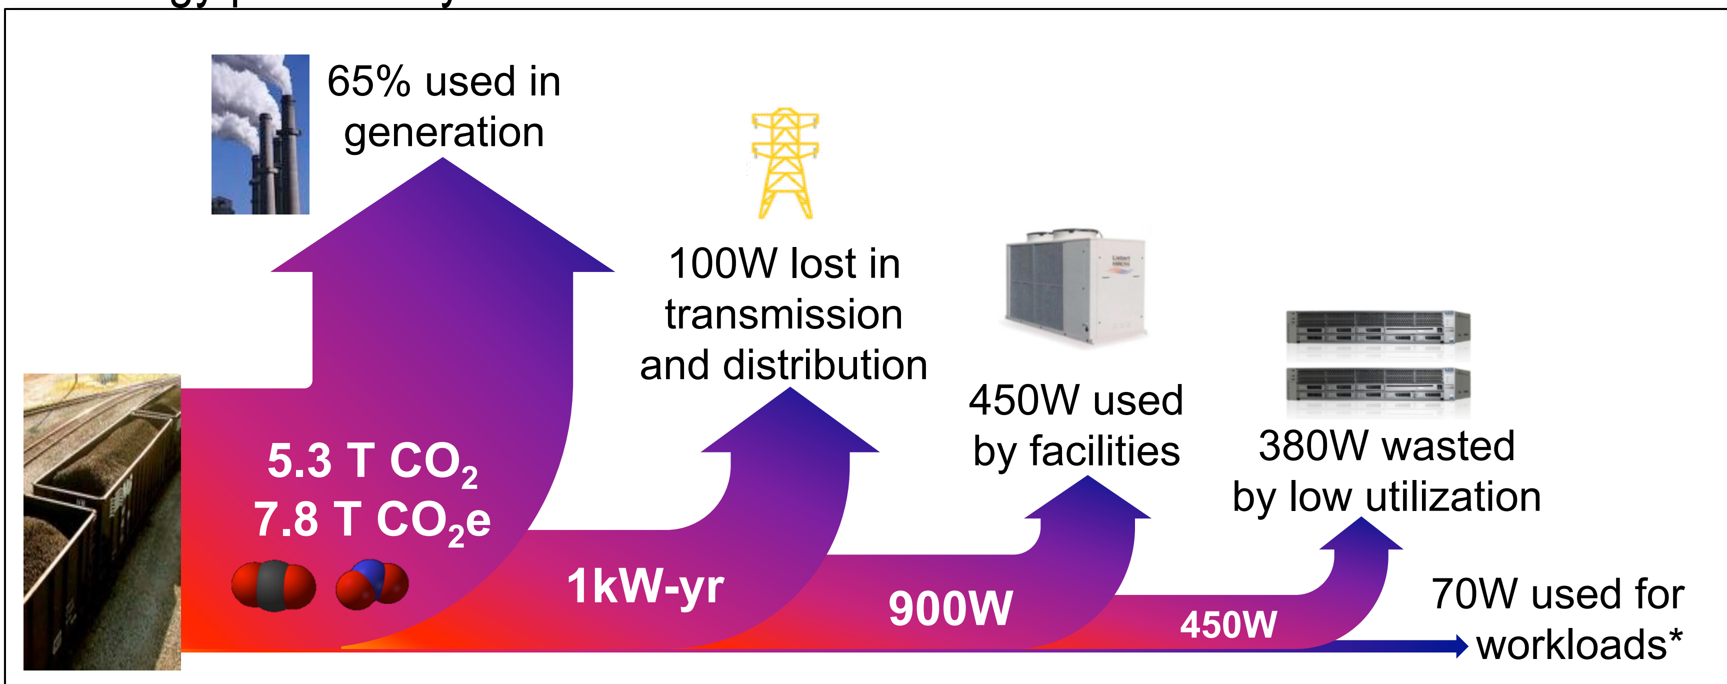

*No Power Management; EPA estimates for load levels (15%) and facilities power overhead of 50%

ORACLE

Greenhouse Gases and IT

Server CO_2 in Perspective

- A single server is responsible for about the same amount of CO_2 as a typical automobile driven for a year

Server	**Auto Travel**	**Air Travel**
440 Watt Server	Toyota Camry	Commercial Airliner
3,942 kWh/year	20,000 km/year	Vancouver-Toronto (6 trips)
5.3 Tonnes CO_2	4.4 Tonnes CO_2	4.4 Tonnes CO_2

The EPA estimate of 5.5 Tonnes includes all cars and light trucks, driven an average of 12,000 miles/year*

ORACLE

Data Center CO_2

- Even a small or medium sized Data Center can be responsible for huge amounts of CO_2

- 4,500 sq-ft @ 100W/sq-ft,
- 800 servers, plus storage and networking
- 450kW of IT load
- 1MW of electricity
- 5,300 Tonnes CO_2/year

* Assumes average of 450W per server

ORACLE

US Data Center CO_2 Emissions

- The cumulative impact of Data Centers is huge: 37 Million Metric Tonnes a year in the US alone
  - > Not including PCs, Laptops, Thin-Clients, Mobile Devices, etc..!

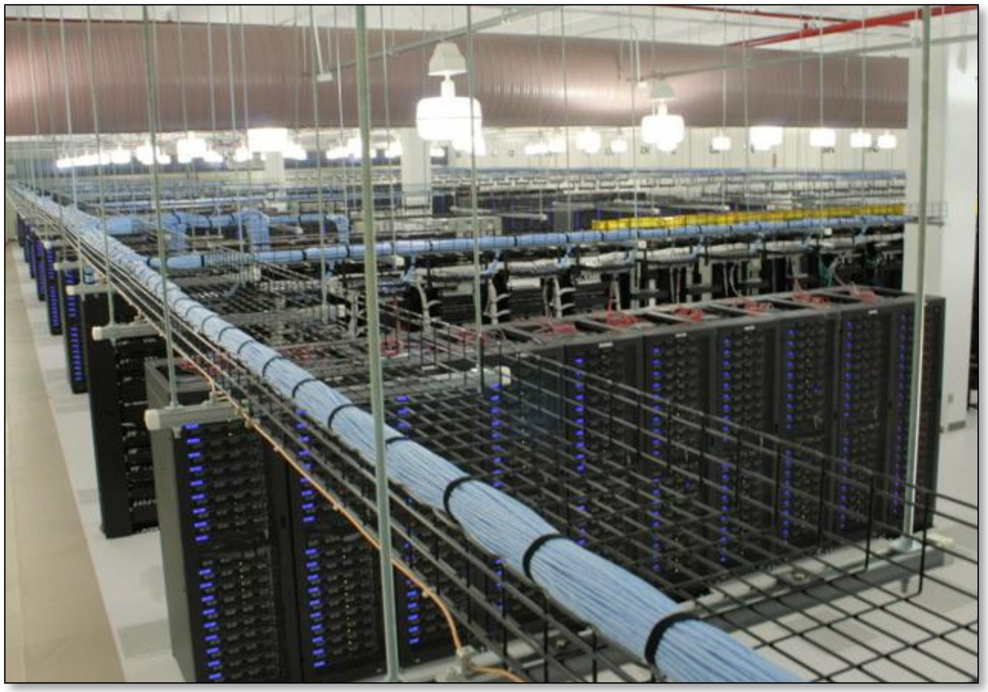

- 61.4 Billion kWh/yr
- 7 Million kW-yr of generation capacity
  - Equivalent to 7 Million 440W Servers, or 8.4 Million Autos
- 37,000,000 Metric Tonnes CO_2 per year

Data Center Growth Projections

- Data Center power usage is growing fast
  - > 42 Million Metric Tonnes more CO_2 – (about 8 million servers)
  - > $5.6 Billion in additional electricity cost (at 8 cents per kWh)
  - > $20 to $80 Billion in facilities costs (at $5K to $20K/kWh of IT load)

Figure 3-4. Comparison of Projected Electricity Use, All Scenarios, 2007 to 2011

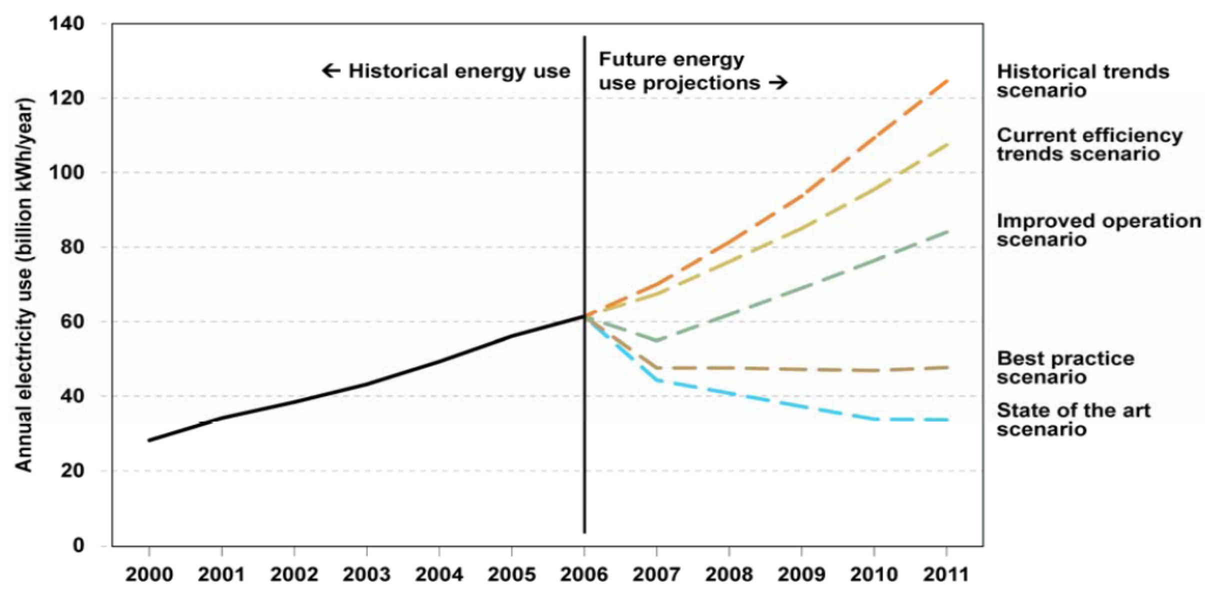

Reducing Data Center Energy Usage

- The EPA estimates that the forecasted energy usage can be cut by up to 70% using current technology

Figure 3-4. Comparison of Projected Electricity Use, All Scenarios, 2007 to 2011

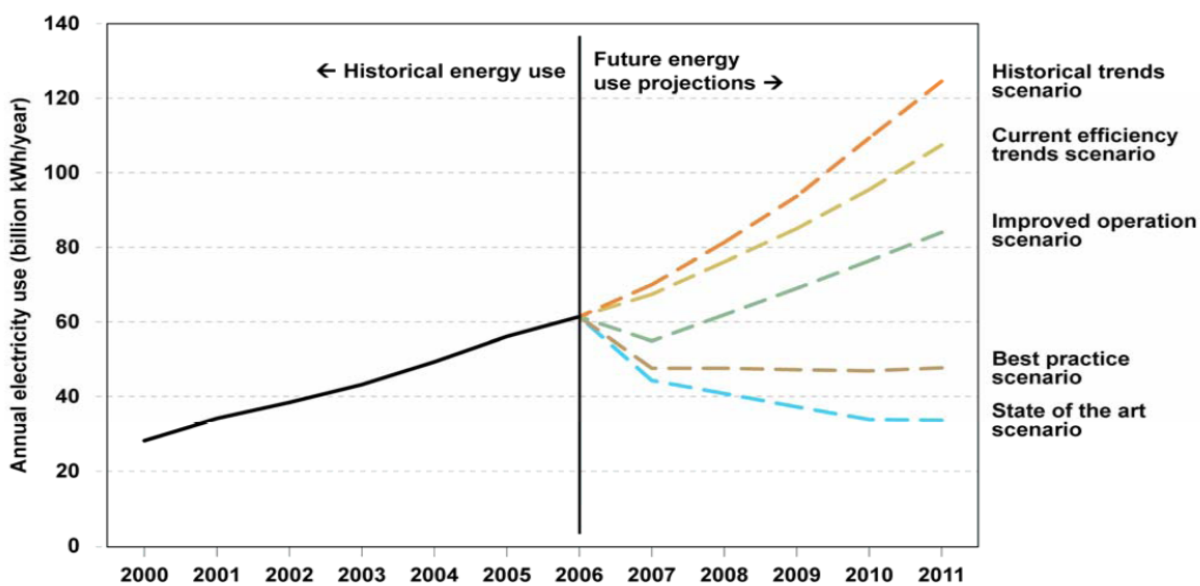

- ✓ 100% use of energy efficient servers
- ✓ All servers use power management
- ✓ Overall 5 : 1 server consolidation
- ✓ 5% of servers are retired
- ✓ Overall 2.4 : 1 storage consolidation
- ✓ Reduced facilities overhead by 55%

ORACLE

Impact of Higher Utilization

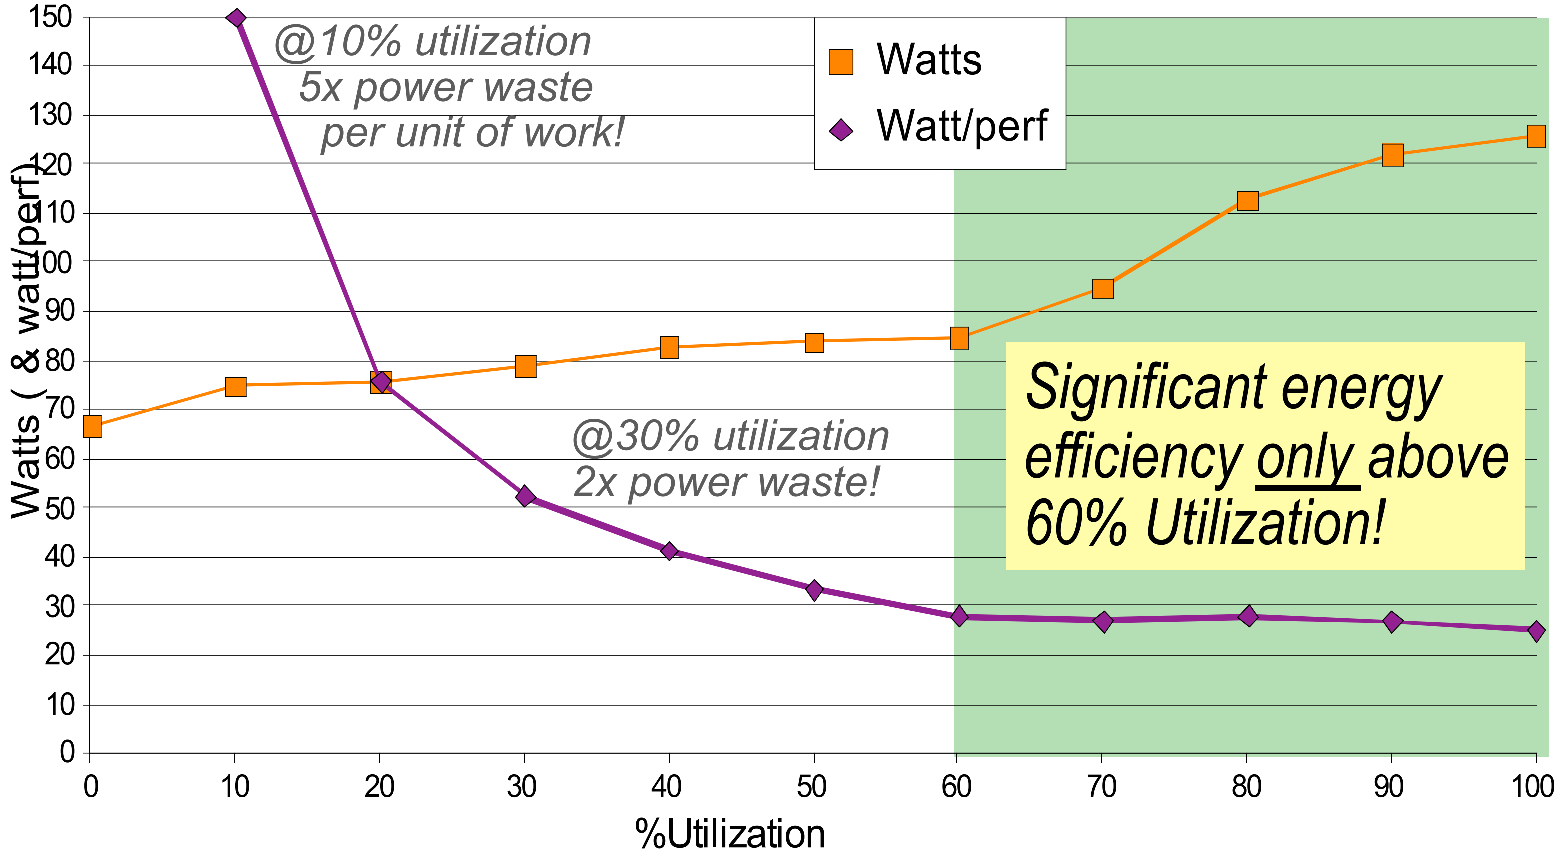

Increasing Utilization Saves More than Best Power Saving HW/SW

1 Server Watt saved equals 2.84 Watts

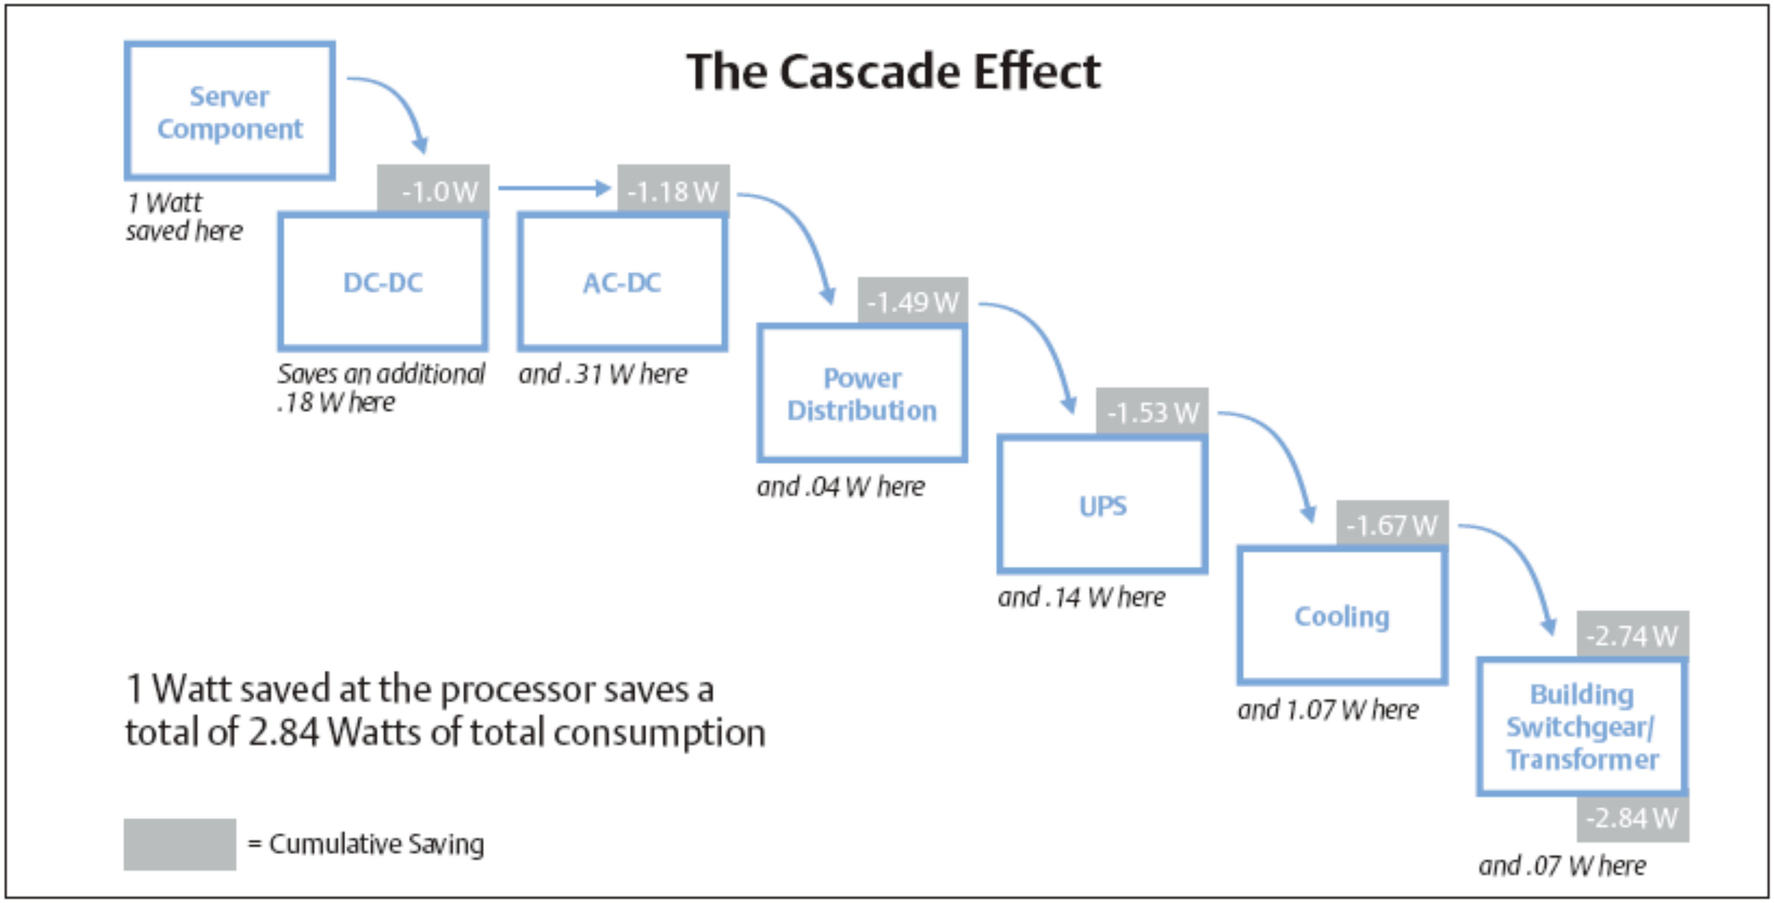

Source: EmersonLiebert, Energy Logic: Reducing Data Center Energy Consumption by Creating Savings that Cascade Across Systems, WP154-158-117 SL-24621

ORACLE

Eliminating Energy Inefficiencies

- Understand the economics of Data Center efficiency
- Obtain senior management support
- Develop models of your Data Center
- Measure and benchmark your current environment
- Identify and prioritize improvement opportunities
- Execute improvements
- Sustain improvements by building energy management into your capacity and financial management practices (e.g. ITIL)

ORACLE

The Power and Cooling Challenge

All power corrupts, but we need electricity.

Dan Galvin, TFTD-L@TAMU.EDU

ORACLE

More Power Equals More Cooling

- Power, heat, cooling, airflow, and temperature are all inter-related.

Tons/CFM

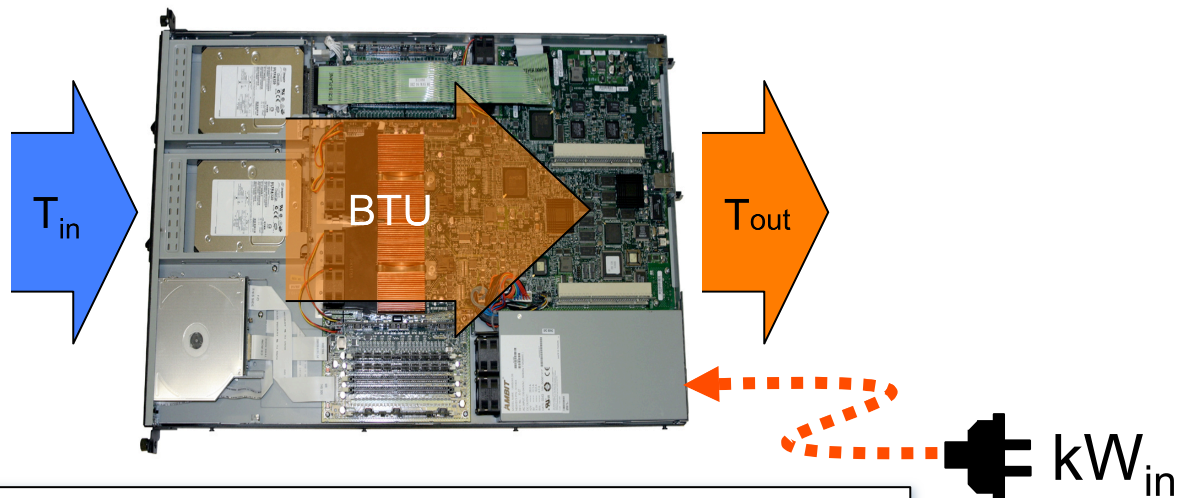

If Power Consumption = 1 kW/hr,
Heat Production = 3,410 BTU/hr,
AC Required = 0.283 Tons,
Airflow Required = 154 CFM,
for a ΔT of $20^{o}F$ @ 1 atmosphere

Note: Normal atmospheric pressure is defined as 1 atmosphere. 1 atm = 14.6956 psi = 760 torr.

Heat, Temperature and Airflow

- Under standard conditions (1 atm), heat, temperature, and airflow relate as follows:

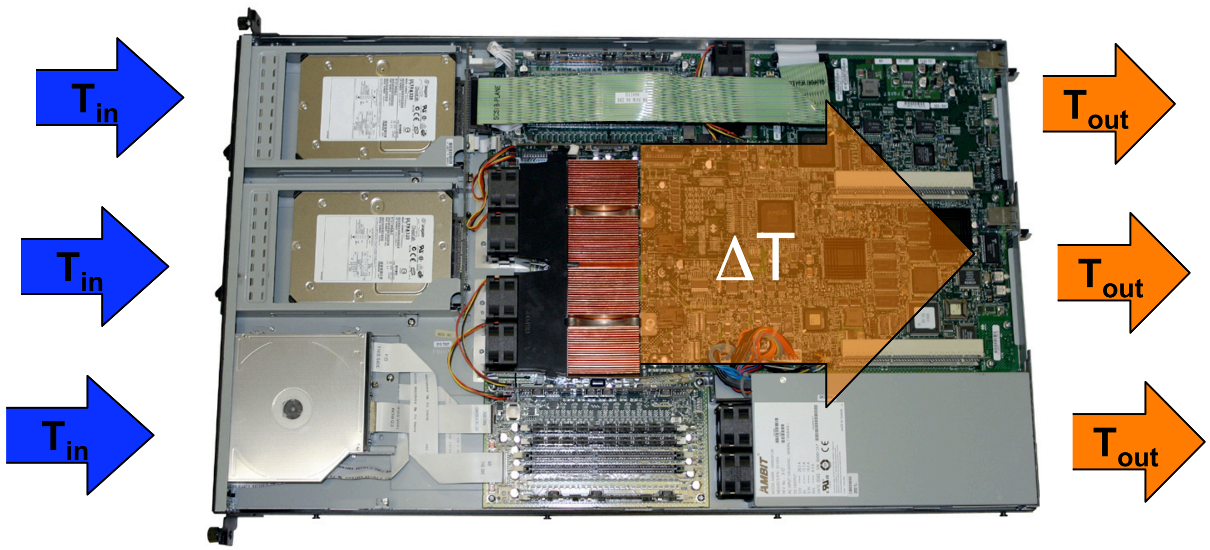

$$\Delta T(F) = \frac{3{,}412 \times kW}{1.085 \times CFM} \quad \text{or} \quad CFM = \frac{3{,}412 \times kW}{1.085 \times \Delta T(F)}$$

ORACLE

Data Center Cooling Rule Number 1

- All the heat generated by electronic equipment (server power) has to be removed from the room.
- Traditional raised floor cooling can typically handle up to 5 kW per rack. This assumes:
  - > raised floor is high enough – higher than 24”
  - > no obstructions – cables, trays, etc…
  - > hot aisle / cold aisle equipment layout – servers front to front, back to back

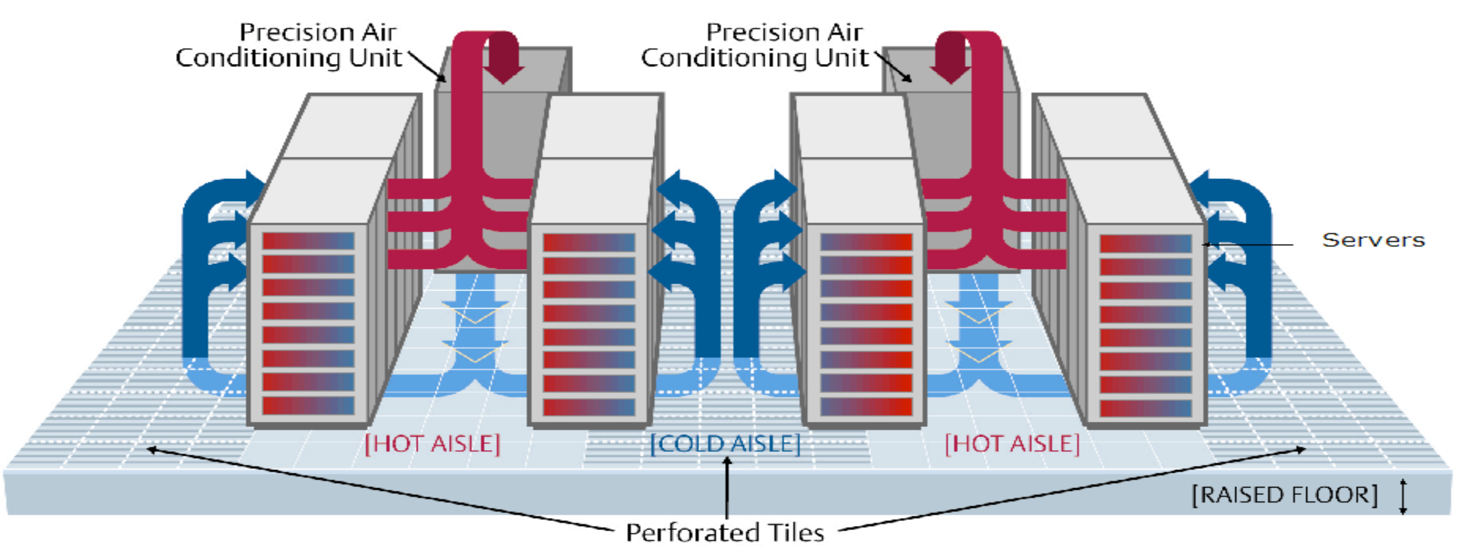

The Challenge is Managing Air

- Air is a fluid that is notoriously difficult to manage.
- Because data centers are constantly changing, thermal management becomes a problem.
- Traditional raised floor, forced air data centers have numerous challenges:
  - Leakage around floor tiles
  - Leakage around cable cutouts
  - Leakage to other rooms adjacent to the data center through access
  - Back pressure in CRACs from blocked filters and poorly engineered baffling

ORACLE

Raised Floor Cooling Limitations

- Standard tiles have a limited airflow delivery capability
  - Airflow above 500 CFM is considered 'extreme' and requires the use of special grate tiles (typically > 40% open area)

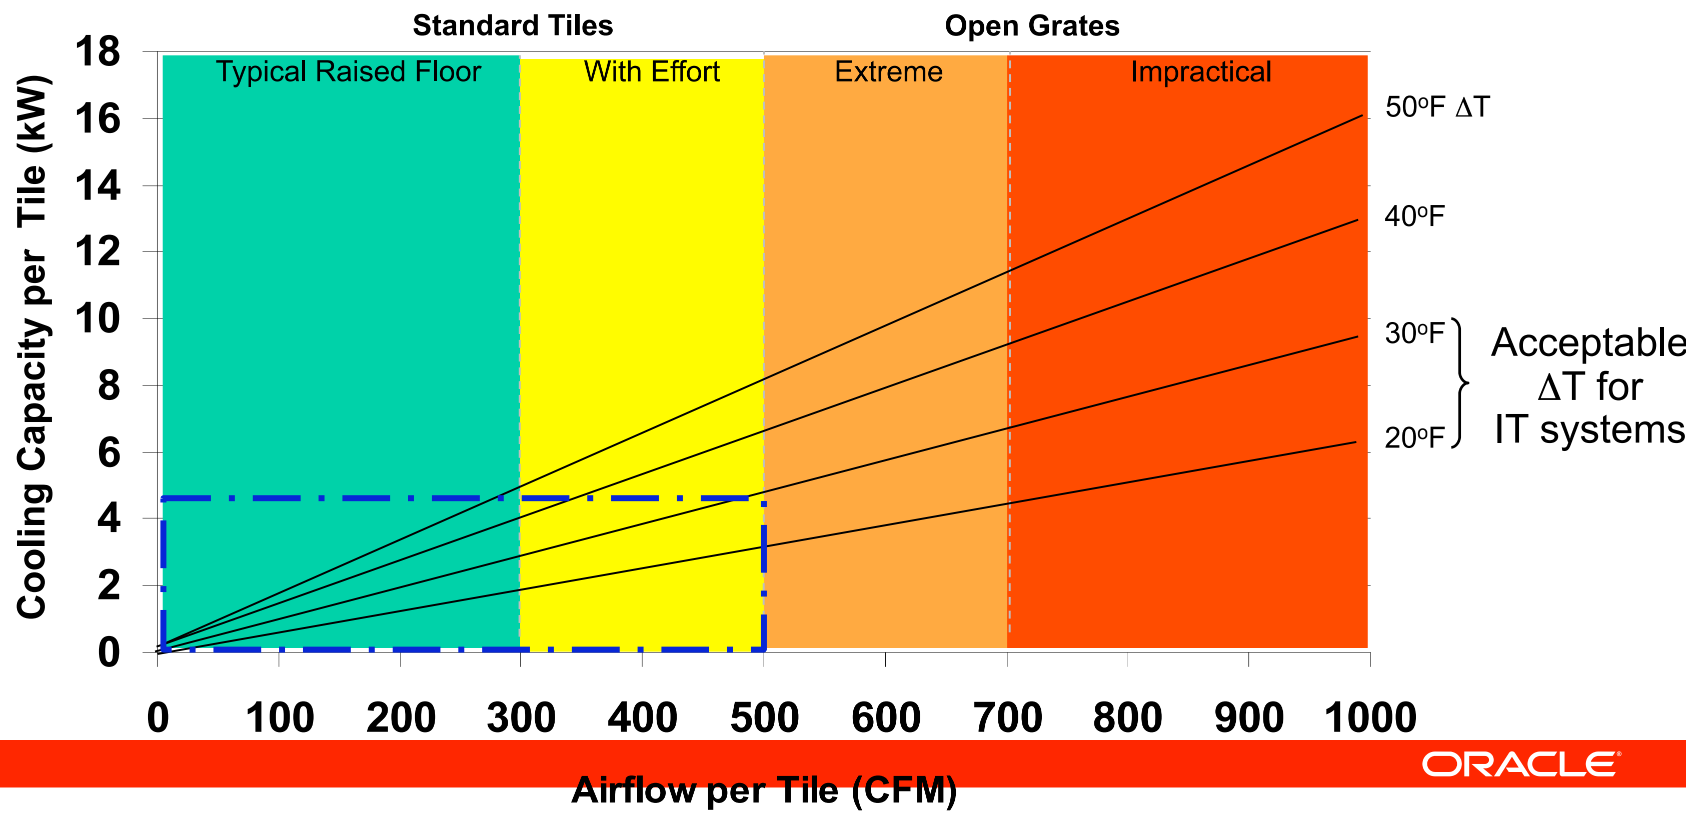

Limits to Raised Floors

- Floor load consideration
  - IT loads increasingly heavier – racks are beginning to weigh beyond 1 Metric Ton each – and designing raised floors will soon require calculation of point loads and rolling loads (but as a facility provider, do you know these at design phase?)
- Increasing Infrastructure costs
  - Increasing heat densities require much higher – and more expensive – raised floors height
- Decreasing energy efficiency
  - Raised floors can be very efficient at low heat densities but become much less efficient as air velocities and sub-floor pressures increase
- Increased design costs
  - High density raised floors require much more careful designs (i.e. CFD modeling)
- Fire suppression
  - Fire suppression is generally focus on isolation of smaller zones and release of a clean agent to extinguish the fire in that zone. With raised floor, you instantly double the number of zones you must monitor, and deploy fire suppression systems into.
- Cleanliness
  - Unless it was installed yesterday, all sort of dirt, dust, debris will accumulate and lurk beneath every raised floor in actual production.
  - Pollutants and contaminants in the air will lead to higher risk of failure.

ORACLE

Cooling Closer to Heat Source

Room → **Row** → **Rack** → **Chip**

- Blade Designs are increasing Density, Power, and Weight per Rack
- CRAC will handle 2 – 5 kW per rack. Beyond 5 kW, CRAC becomes increasingly less efficient.
- Current planning uses 9 kW per rack as the low end of the cooling requirement
- HPC installations are currently running at 25 to 28 kW per rack
- Managing air flow is now mandatory.
  - Capture Index for In-Row and Rack Cooling
    - Open Hot Aisle – 0.80
    - Hot Aisle Containment – 0.90
    - Rack Air Containment - 0.95

Efficiently Managing Heat

- Until we move cooling to the chip level, we will be using air to transfer heat from the IT source.
- The key to efficiency is managing the least amount of air and that is best done in a "contained" design
- A good design moves the initial heat transfer as close to the source as possible
- It then uses efficient and cost effective technology to move the heat off premises.

ORACLE

Weakness of Traditional DC Design

- Data centers are often described using average heat density in w/sq ft or w/sq m.
- In practice, data centers are not populated based on average heat density.
- Using CFD tools and trial and error, it is possible to run high density compute infrastructure in a traditional raised floor data center.
- Problem with this approach:
  - > Data centers change and the heat flux balance is impacted
  - > This works where a limited number of servers are high density. As average density increases, traditional designs start to fall apart.

Measuring Data Center Efficiency

- Data Center infrastructure Efficiency (DCiE)

$$\text{DCiE} = \frac{1}{\text{PUE}} \times 100\% \qquad \text{PUE} = \frac{\text{Total Facility Power}}{\text{IT Equipment Power}}$$

PUE = Power Usage Effectiveness

- A DCiE of 100% represents maximum (ideal) efficiency
- No comprehensive data in the industry related to DCiE, but early work suggests that many data centers are at 30% or less.
- With proper design, the Green Grid suggests that a DCiE of 62.5% is achievable
- The confirmed DCiE for SCA11-1500 (APC) is 78.57%

ORACLE

DCiE - A Small Change a Big Difference

- US Data Center energy use is currently about 61.4 Billion kWh/yr.
- Today's average data center operates at a DCiE of about 0.50 (50% efficient).
- This translates to 30.7 Billion kWh/yr for facilities
- Increasing DCiE to an average of 0.61 would result in:
  - 6.8 Billion kWh/year reduction in facilities overhead
  - Reduction in 862 MW of electricity generation
- This is the equivalent to retiring 862,000 x 440 Watt servers, or taking 1.0 Million Toyota Camarys off the road

ORACLE

Modularity In Data Center Design

The Sun Blueprint

Next Generation Data Center

- Modular
- Flexible
- Environmentally Conscious
- Designed in Concert with the Facilities
- Incorporate Broad Based Efficiency Improvements
  - Applications mapped to specific types of processors
  - Power supplies sized to server configuration
  - Electrical distribution at higher voltages
  - Integrated management and control
  - Cooling closer to the heat source

ORACLE

Sun Blueprints on Data Center Efficiency

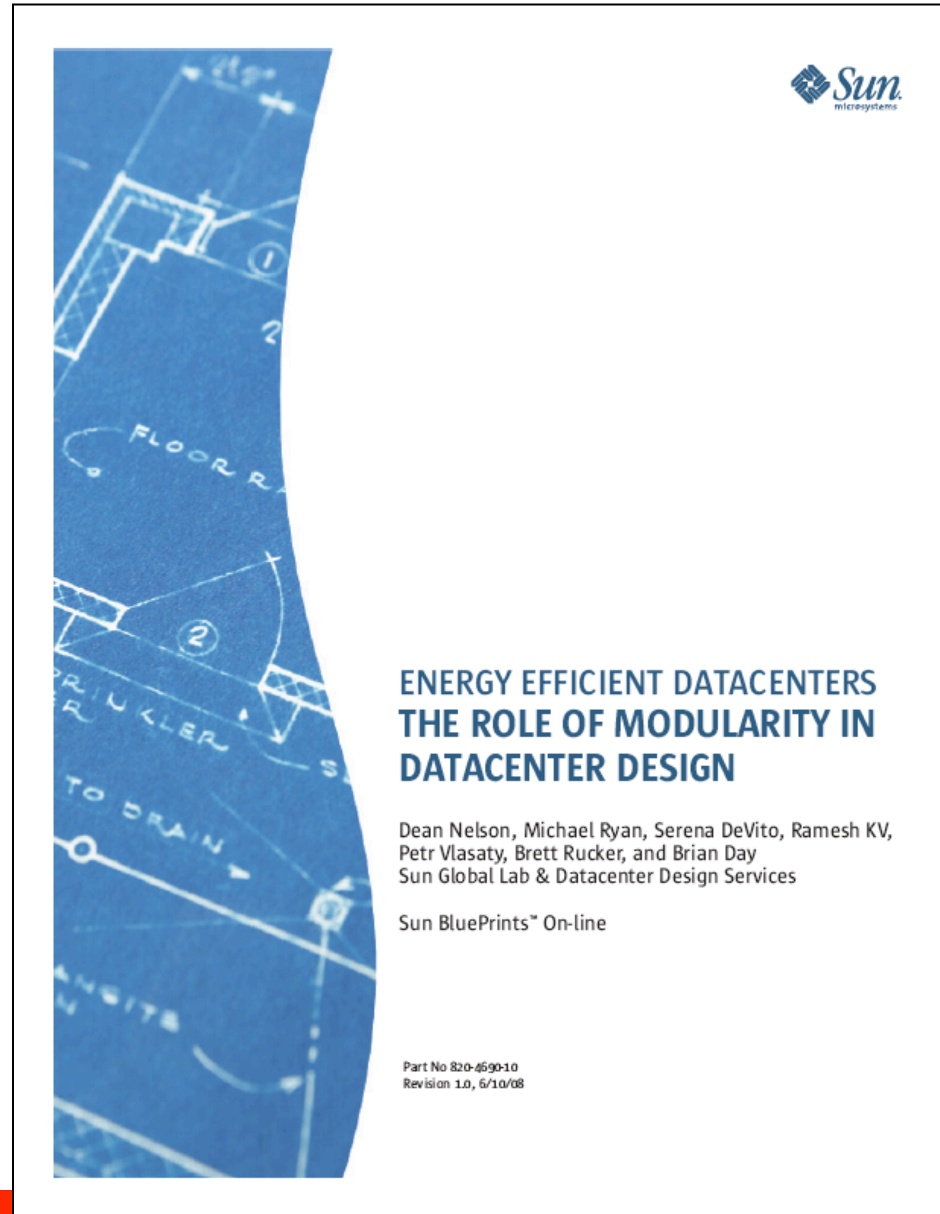

Sun microsystems

ENERGY EFFICIENT DATACENTERS
THE ROLE OF MODULARITY IN DATACENTER DESIGN

Dean Nelson, Michael Ryan, Serena DeVito, Ramesh KV,
Petr Vlasaty, Brett Rucker, and Brian Day
Sun Global Lab & Datacenter Design Services

Sun BluePrints™ On-line

Part No 820-4690-10
Revision 1.0, 6/10/08

Power and Cooling

- Building a data center with maximum power and cooling available from the date of commissioning is the wrong approach.
- Traditional forced air cooling works when a data center is operating at 40 W/sq. ft.
- Moving air is a large contributor to data center energy use.
- Delivering cool air to today's high density racks through raised floor is difficult and inefficient.

The Right Power Metric

- Watts per square foot is an architectural term created for office space.
- Sun has abandoned this measure for a more accurate metric of Watts per rack.
- 2008 industry average Watts per rack is between 4 - 6 kW per rack.
- However, averages are deceiving. Today's data center is heterogeneous and contains racks that range from less than 1kW to more than 30 kW.
- The minimum design point for a new data center should probably by 9 kW per rack or greater

ORACLE

Design to the Need

- Matching capacity to the temporal requirements increases efficiency.
  - Cyclical Workloads driven by time of day, month, or year
  - Frequency of Change - regularly moving servers and cabinets that have different power and cooling requirements
  - Rate of change - replacing older equipment with newer equipment; often driven by lease refresh.
- Physically different power and cooling requirements. Front to back cooling versus chimney.

Living within the Envelope

- Space
  - How much should you plan on using? (27 - 35 sq. ft. per rack)
    - Rack = 7 - 10 sq ft
    - Access = 5 - 10 sq ft
    - Mech/Elec/Other = 15 sq ft
- Efficiency
    - PUE - Total Facility Power divided by Power going to IT
    - A PUE of 2.5 means that only 40% of the power into the data center is going to the compute infrastructure.
    - The lower the PUE the shorter the investment pay back period
- Cooling Capacity
    - 2.5 mW IT load with a PUE of 2.0 requires 714 Tons of cooling plant.

ORACLE

Living within the Envelope

- Floor Loading
  - > Rack weight increased with compute density
  - > Today's blade racks can weigh as much as 2400 lbs and are rapidly moving into the 3000 lb range
  - > Floors often have to be reinforced.
- Raised Floor or Slab
  - > Raised floors tend to have complex and unpredictable airflow characteristics.
- Cooling
  - > Eliminate heat gain from people, solar, lighting
  - > Control humidity within the range of 30 - 60 percent
  - > Filter air
  - > Fresh air as required by code.
- Wider racks to accommodate increasing cable densities.

ORACLE

Sun's Approach to Data Center Design

- Scalable, repeatable, modular architecture
  - > Vendor independent
  - > Slab or raised floor
  - > Modular, right-sized power and cooling
  - > Simplified, flexible cabling and plumbing
- Real-time energy monitoring
- Scale cost with use
- Easy to build and expand

Entering a New Age of Engineered Data Centers

Modular Building Blocks

- POD Based Design
- Up to 36 Racks with a common hot/cold aisle
- Modular power, cooling, and cabling components
- Size is variable
- Tier of Service is variable
- Data center within a data center
- Eliminates most of the custom design cost
- Avoids cost of “sizing it wrong.”

ORACLE

The Sun POD Design

- Small, self-contained group of racks (generally 20 - 24) that optimize power, cooling, and cabling.
- Design at the pod level, rather than room level, simplifies the approach.
- The same architecture can be used in both small and large data centers.
- Integral to our building design in the pod architecture is a 10 - 15% larger MEP support infrastructure to future-proof the space.
- Minimizes both power and network cable lengths.

ORACLE

The Sun POD Design

- Physical layout is hot aisle/cold aisle with either:
  - Spot overhead cooling
  - In row hot aisle containment
- Power distribution uses a busway that runs either above the pod or below the floor.
- Equipment cooling is in-row or overhead as close to the heat source as possible. Room air conditioning is provided to meet code for habitable space, humidity, and air filtration.
- Connectivity allows short cable lengths with in-pod switching and patch panels either above or below each rack.

ORACLE

Pod Architecture using In-Row Cooling

Pod: A group of racks or benches with common hot or cold aisle used as a building block to simplify data center design for power, cooling and cabling

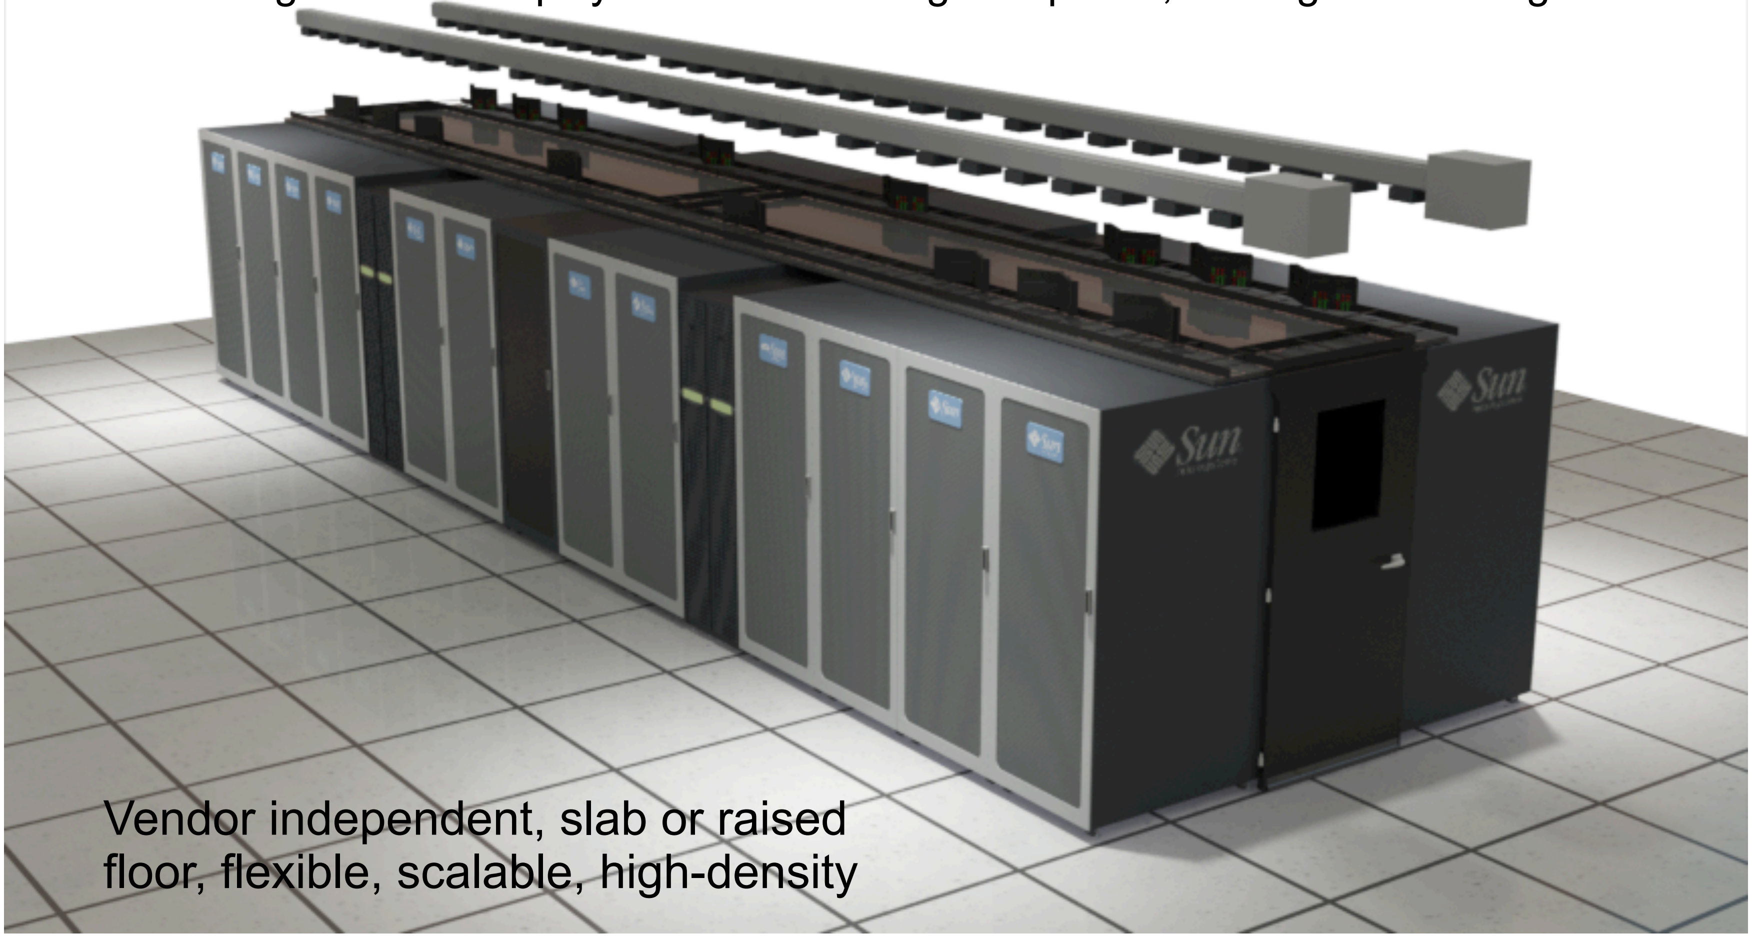

Vendor independent, slab or raised floor, flexible, scalable, high-density

Pod Architecture with Overhead Cooling

Hybrid Pod for Very High-Density

Bay Area Sun Ray Thin-Client Server Consolidation

Before
- 88 x Sun Fire V880
- 1320 ft.^{2} (44 racks)
- Power: 122kW load
- Utility Cost: $122k p.a.

After
- 58 x Sun Fire T2000
- 120 ft.^{2} (4 racks)
- Power: 11.6kW load
- Utility Cost: $14.6k p.a.

Results
- 4813 Users served with two-thirds the equipment
- 91% reduction in floor space (44 to 4 racks)
- 88% drop in annual utility costs (not including cooling)
- Cut carbon emissions by 210 metric tons

Santa Clara Lab Consolidation

Before

- 2177 Servers
- 202,000 ft.^{2} (550 racks)
- Power: 2.2MW
- Utility Cost: $1.89M p.a.

Bridging the gap between Facilities & IT

After

- 1015 Servers
- 76,000 ft.^{2} (65 racks)
- Power: 500kW
- Utility Cost: $739k p.a.

Results

- $9M Cost Avoidance
- Decommissioned >5000 servers, storage, and networking devices
- 61% drop in annual utility costs (not including cooling)
- Cut carbon emissions by 3227 metric tons
- Completed in three months with minimal downtime

ORACLE

Spotlight: Sun Engineering Data Center

- 76,000 ft^2 data center
- 9MW scalable to 21MW
- Largest Liebert/APC installs
- 15 Buildings to 2
- 152 data centers to 14
- Completed in 12 months
- $1.2M Utility rebates
- $250K Innovation Award
- 39% more efficient than ASHRAE standard (0.489 kW/ton)
- Reduced OPEX by 30+%

Delivered: Modular, scalable, future-proof and highly efficient in 63% less space. Under budget & on schedule.

Sun's DCiE Achievement

A DCiE of 78.57% translates into energy savings of $402,652 per year when compared to a more traditional data center built with a DCiE of 50% (PUE of 2.0).

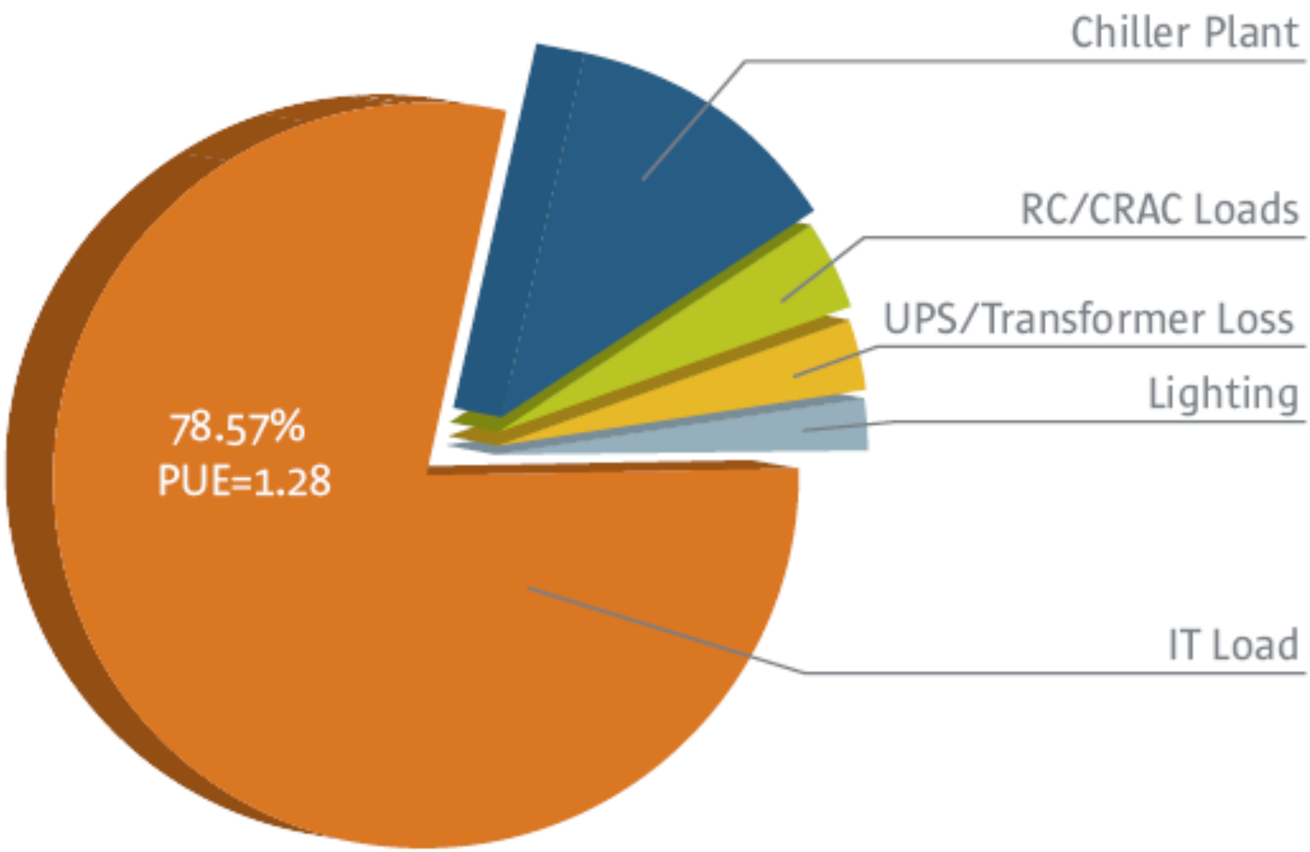


Pod Hot Aisle

APC in Row Cooling Units
Capacity of up to 30 kW Per Unit

Liebert XD

Liebert Pumping Unit and XDV
Capacity of 8 to 10 kW per Unit

StarLine Bus Power Distribution

Structured Cabling Solution

ORACLE

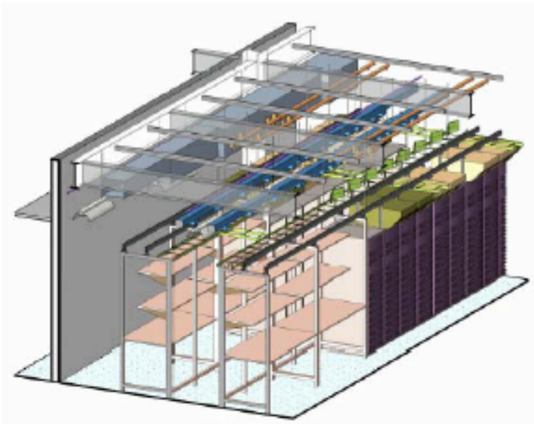

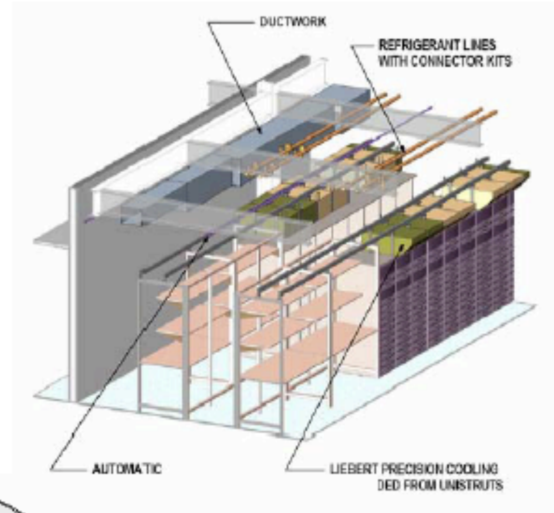

Isometric View of Computer Lab—
All services and components were coordinated in this three-dimensional model of the space.

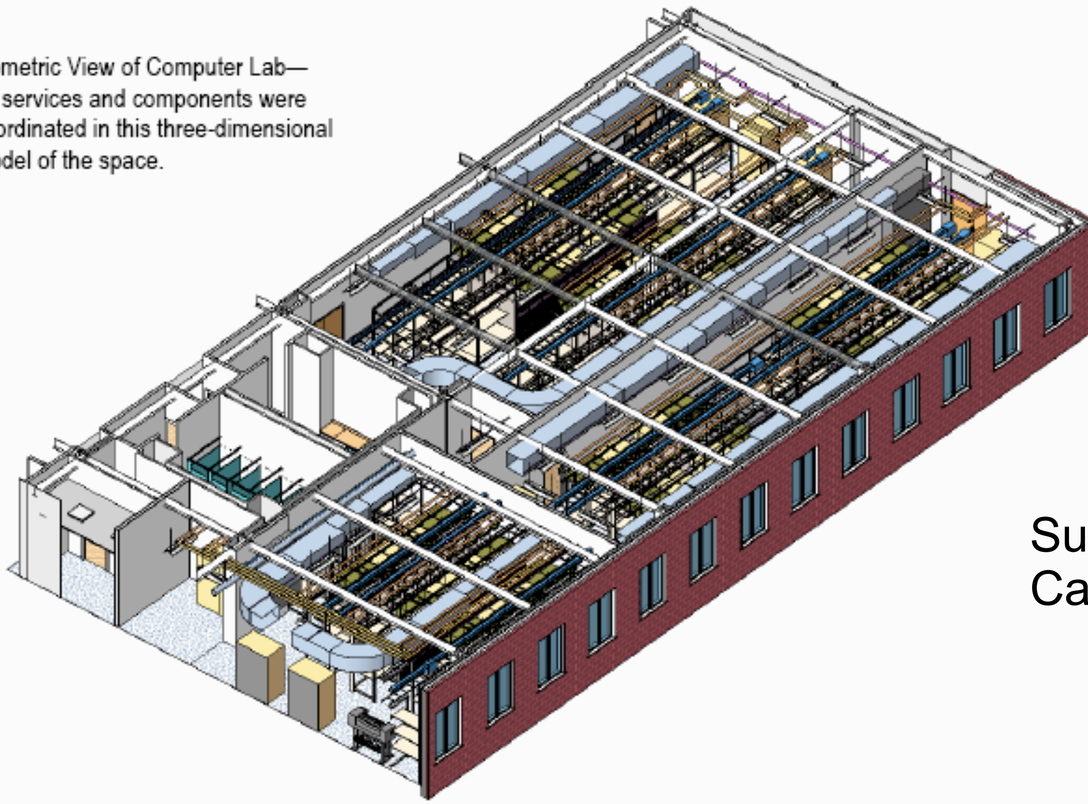

Sun Burlington Campus Lab

Conceptual Design
European High
Performance Computing
Center

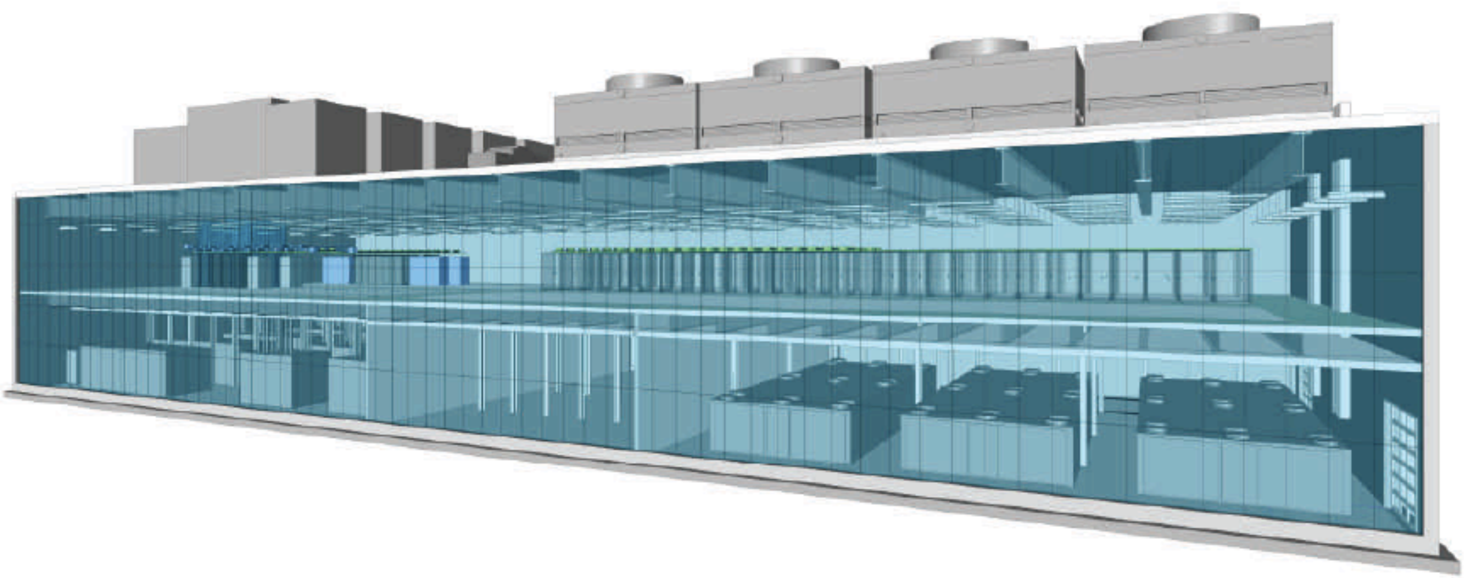

CLUMEQ "Colossus"

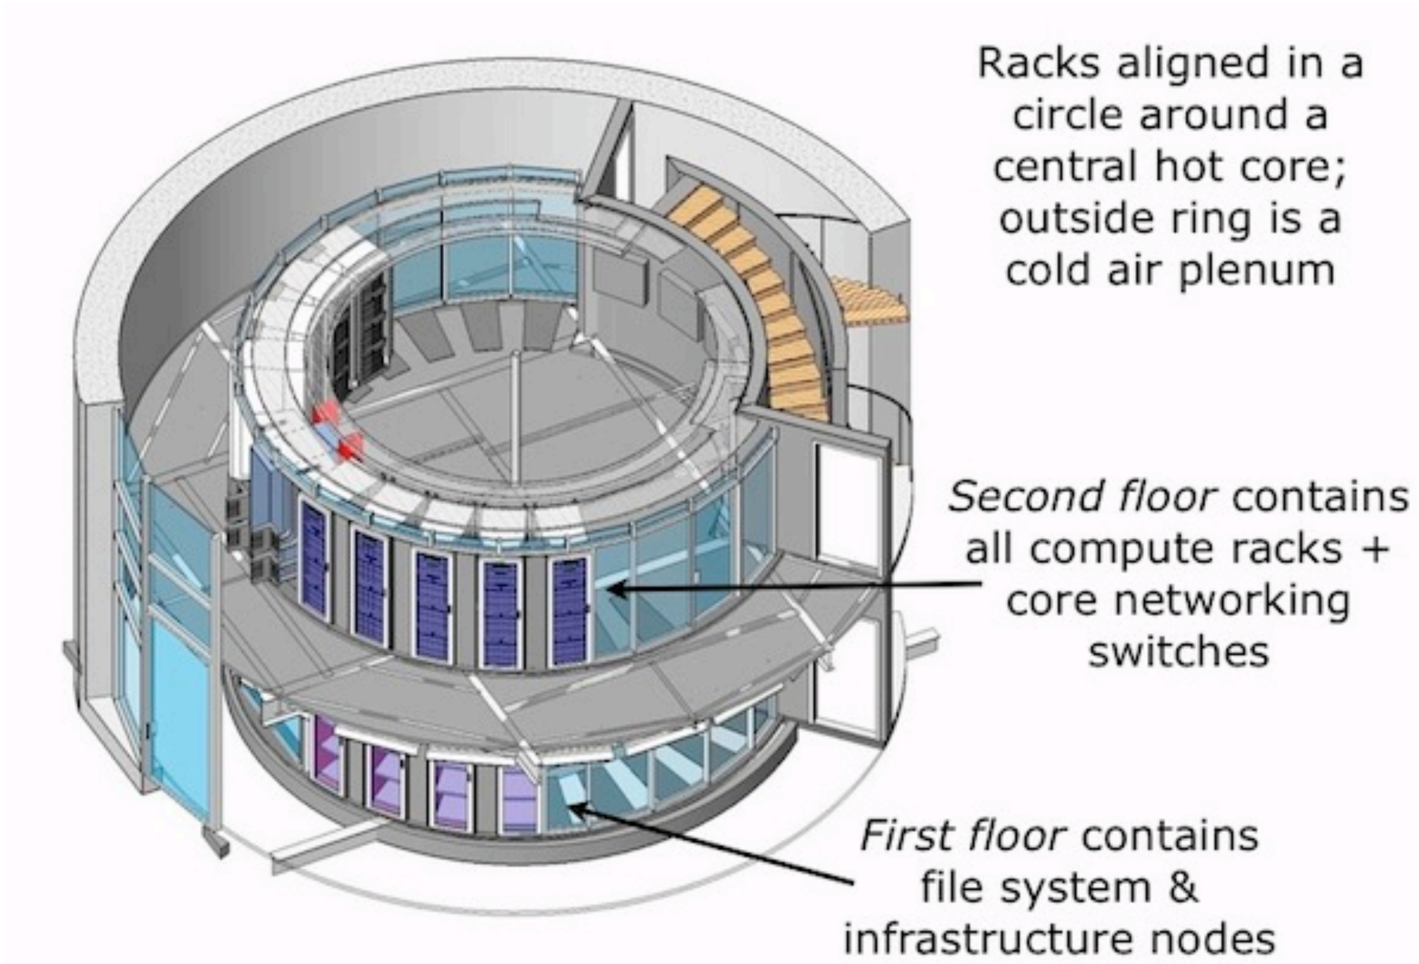

CLUMEQ

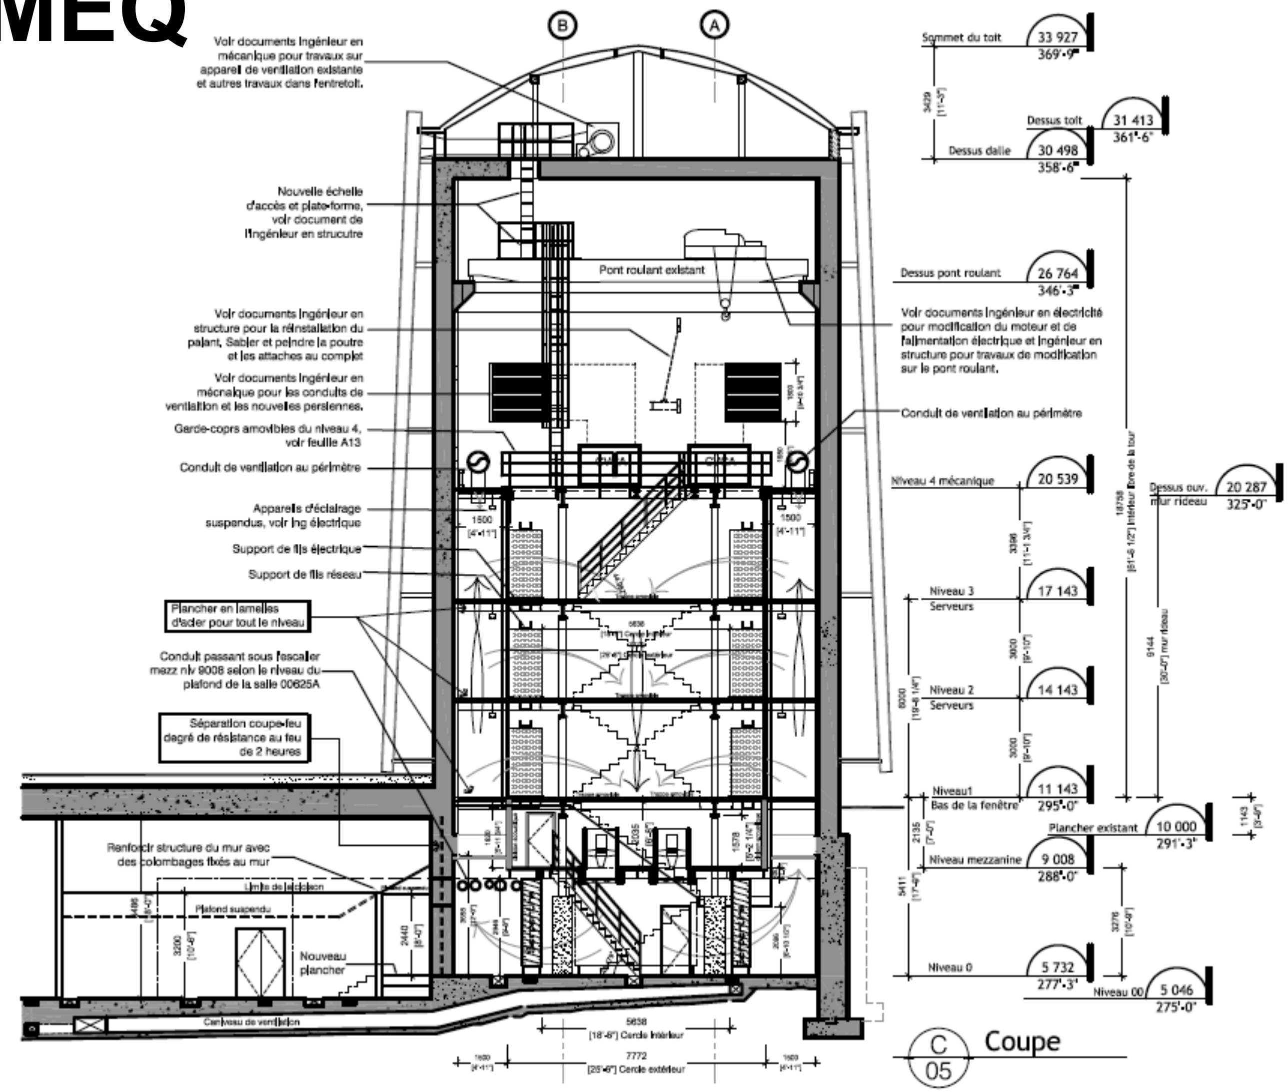

Voir documents Ingénieur en mécanique pour travaux sur appareil de ventilation existante et autres travaux dans l'entretoit.
Nouvelle échelle d'accès et plate-forme, voir document de l'ingénieur en strucutre
Pont roulant existant
Voir documents Ingénieur en structure pour la réinstallation du palant, Sabler et peindre la poutre et les attaches au complet
Voir documents Ingénieur en mécanique pour les conduits de ventilation et les nouvelles persiennes.
Garde-corps amovibles du niveau 4, voir feuille A13
Conduit de ventilation au périmètre
Appareils d'éclairage suspendus, voir ing électrique
Support de fils électrique
Support de fils réseau
Plancher en lamelles d'acier pour tout le niveau
Conduit passant sous l'escalier mezz niv 9008 selon le niveau du plafond de la salle 00625A
Séparation coupe-feu degré de résistance au feu de 2 heures
Renforcir structure du mur avec des colombages fixés au mur
Plafond suspendu
Nouveau plancher
Caniveau de ventilation
Sommet du toit
33 927
369'-9"
Dessus toit
31 413
361'-6"
Dessus dalle
30 498
358'-6"
Dessus pont roulant
26 764
346'-3"
Voir documents Ingénieur en électricité pour modification du moteur et de l'alimentation électrique et Ingénieur en structure pour travaux de modification sur le pont roulant.
Conduit de ventilation au périmètre
Niveau 4 mécanique
20 539
Dessus ouv. mur rideau
20 287
325'-0"
Niveau 3 Serveurs
17 143
Niveau 2 Serveurs
14 143
Niveau1 Bas de la fenêtre
11 143
295'-0"
Plancher existant
10 000
291'-3"
Niveau mezzanine
9 008
288'-0"
Niveau 0
5 732
277'-3"
Niveau 00
5 046
275'-0"
5638
[18'-6"] Cercle intérieur
7772
[25'-6"] Cercle extérieur
C
05
Coupe

CLUMEQ

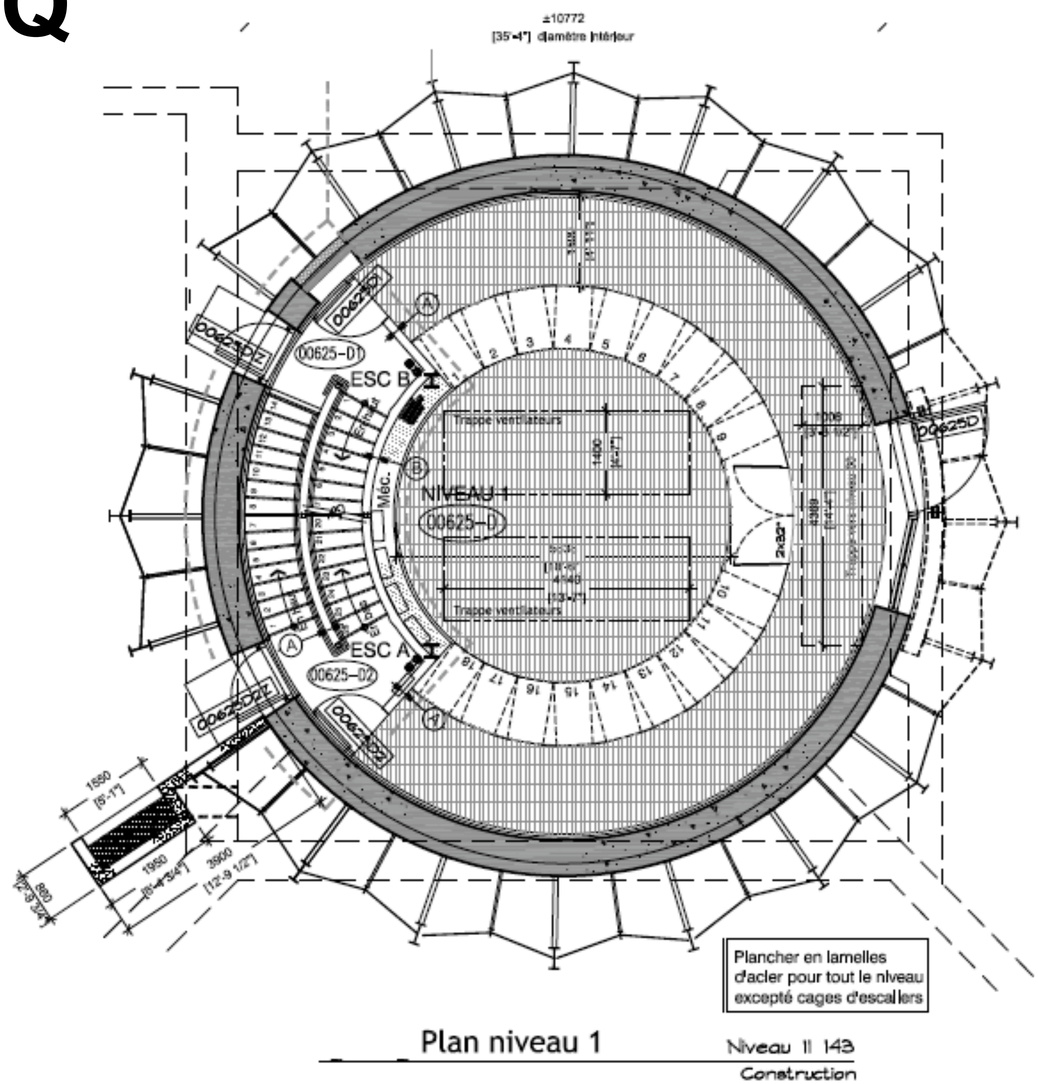
±10772
[35'-4"] diamètre intérieur
ESC B
ESC A
NIVEAU 1
Méc.
Trappe ventilateurs
Trappe ventilateurs
Plancher en lamelles d'acier pour tout le niveau excepté cages d'escaliers
Plan niveau 1
Niveau 1 1:43
Construction

Telco POD P.O.C.

Efficiency Strategies for Today’s Use

Top 10 Efficiency Strategies

Power savings can be addressed in many different ways depending on level of involvement by different categories of stakeholders.

Authority to Change

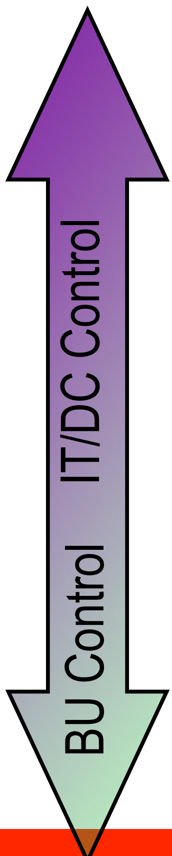

Top 10 Strategies

Potential Independent Energy Savings

Top 10 Strategies	*Point Projects*	*Entire DC*
Cooling Efficiency	30%-50%	5%-30%
Power Efficiency/Distribution	20%-40%	5%-25%
Storage Efficiency	20%-75%	5%-20%
Turn on CPU Power Management	5%-40%	5%-15%
Match Infrastructure to SLAs	20%-30%	10%-15%
Retire Unused Systems	50%-100%	5%-10%
Power Efficient Systems for New Deployments	40%-60%	10%-30%
Technology Refresh	50%-80%	30%-50%
Optimize System Utilization: Consolidation/Virtualization/ Workload Mgmt	50%-80%	30%-60%
More Efficient Applications	50%-95%	30%-70%

ORACLE

Ten Step Plan for Today

- Improve Solution Design
- Optimize System Utilization
- Enable Power Management
- Retire Unused Systems
- Use Power Efficient Systems
- Accelerate Technology Refreshes
- Improve Storage Efficiency
- Improve Cooling Efficiency
- Improve Power Distribution Efficiency
- Match Infrastructure to SLAs

ORACLE

In Closing…

- Externalities are changing the way we look at data centers
- Power, cooling, and space efficiency are critical
- Sun's modular, pod-based architecture raises efficiency, provides agility for balancing computational supply with business demand
- IT must integrate with business strategy
- Oracle ACS Systems Modernization can help:
  - IT Business Strategy
  - Data Center Services
  - Application and Infrastructure Modernization
  - Operational Management

ORACLE

Thank You!

john.kapson@oracle.com