Green Data Center? or Energy Efficient Data Center

27th -28th October 2009,
Congress Frankfurt

SNIA Europe
DATACENTER TECHNOLOGIES

Josh Krischer

Josh Krischer & Associates

Energy Consumption Challenges & Facts

- The Uptime Institute reports data center energy use doubled between 2000 and 2006 and predicts it will double again by 2012. 25% + of an IT budget is consumed by energy costs
- The “green DC” started with the introduction of high density “blade” servers
- Most of the existing data centers are design for equipment which consumes less than1KW / m^2
- Many data centers can not supply required power
- Blade servers require less energy than stand alone servers with the same computing power.
- The energy costs for cooling depend on cooling infrastructure effectiveness
- Mixing cold and hot air lowers cooling effectiveness and increase costs
- Millions of Euro/Dollars are being wasted on poorly planned and deployed cooling
- Poor-quality data centers will become increasingly uncompetitive and costly to run
- DC “Greenness” awareness impacts corporations’ image

Data Centers' Energy Consumption "Dirty Secrets"

- Most of the CIOs are not aware to their DC energy consumption and costs
  - Only 25% of respondents say that the IT department pays for its energy bill today *
  - Only 28% of respondents report they know accurately what the energy consumption of their data center is*
- Till recently IT has been deployed without consideration of energy requirements, costs and CO_2 emission
- A common practice is to overprovision in servers and storage capacities
- Poor communication between IT architects and facilities management

Carbon tax legislation currently before the U.S. Congress will significantly increase the cost of energy. A similar legislation is occurring around the world.

* IBM 2008 CEO Study

Reasons for Adopting Green Solutions

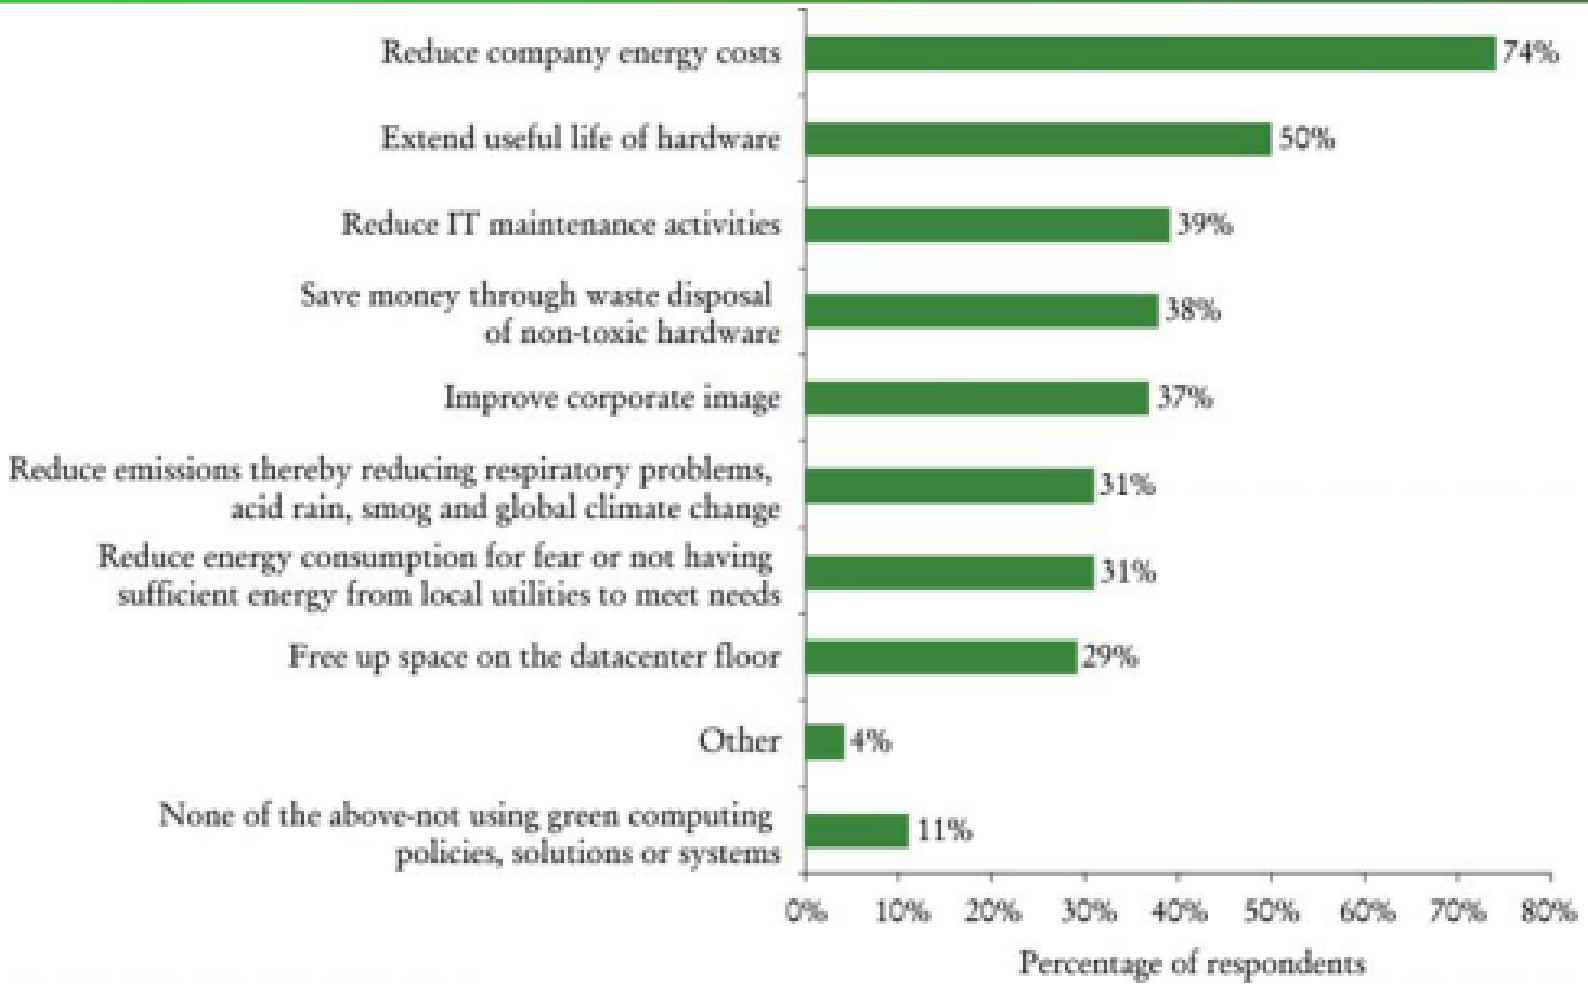

(Source: "Survey the green-tech landscape", Green-Tech Insights, an InfoWorld Research Report conducted by IDG Research Services Group, November 21, 2007, InfoWorld)

How to Adopt Green Solutions?

How to Address DC Energy Consumption Issues

- Short term – tactical actions to get immediate results
- Mid-term – consolidate, evaluate new technologies
- Long-term – design new or re-design you data center using fresh air, water cooling, solar or wind energy, etc.

How to Address DC Energy Phase III example

From deep underground, data center will help heat Helsinki homes

By Andrew Nusca | Nov 30, 2009 | 3 Comments

SHARE PRINT RECOMMEND 1 VOTES

The Finnish capital of **Helsinki** is preparing to house what may be the greenest data center on the planet.

Hidden deep within the bedrock of a massive cave underneath popular orthodox Christian landmark Uspenski Cathedral, the planned data center — which will be comprised of hundreds of computer servers — is expected to emit substantial amounts of heat.

That heat will then be captured and channeled into the city's district heating network, a system of water-heated pipes that are used to warm homes in

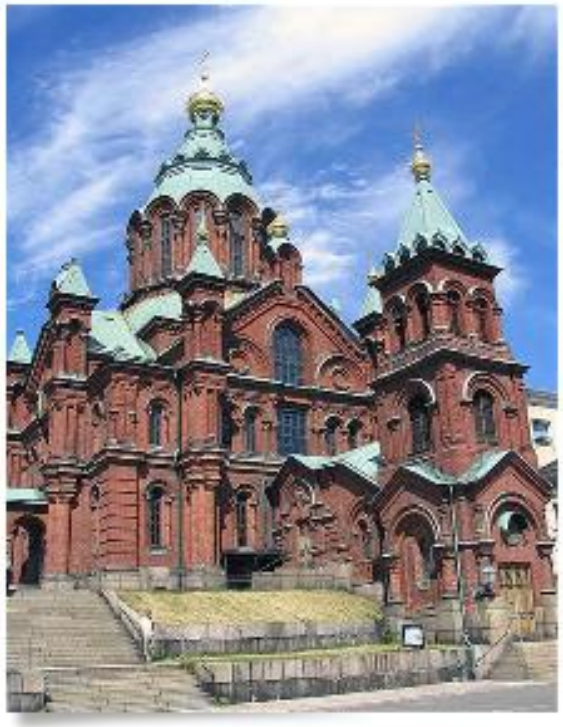

How to Address DC Energy another Phase III example

HP Opens First Ever Wind-cooled Data Center

PCWorld

Jeremy Kirk, IDG News Service – Wed Feb 10, 9:50 am ET

RELATED QUOTES

^DJUSS	421.98	-2.60
^IXIC	2,157.87	-19.54
^IXK	1,083.35	-11.17

From the outside, Hewlett-Packard's newest data center looks like a massive, well-secured loading dock, devoid of logos and surrounded by a robust barbed-wired fence in a nondescript industrial park.

The low-profile approach is intentional, as HP's Wynyard center is intended to hold the most valuable asset for many companies: their data. HP will use the data center to compete with companies such as IBM for IT services and management contracts, a growing source of revenue that requires secure data centers.

HP is hoping several of the environmentally friendly design features of the 360,000-square-foot Wynyard facility will push it ahead. It is HP's most energy efficient data center that it has built, said Maurice Julian, U.K. facilities project director. Half of the facility is now complete, comprising four data halls, with room to create four more data halls as demand dictates.

The data center was originally started by IT outsourcer EDS, which was then acquired by HP for US$13.9 billion in May 2008. The building sits in a blustery and chilly area about eight miles west of the North Sea in the northeast of England. It is entirely air-cooled: HP has built eight 2.1-meter stainless steel and plastic intake fans to draw cool air.

The air runs through a massive bank of modular filters to remove dust and other contaminants before it circulates in a massive cavity, called a plenum, below its data center halls.

Product Heat Density Trend Chart

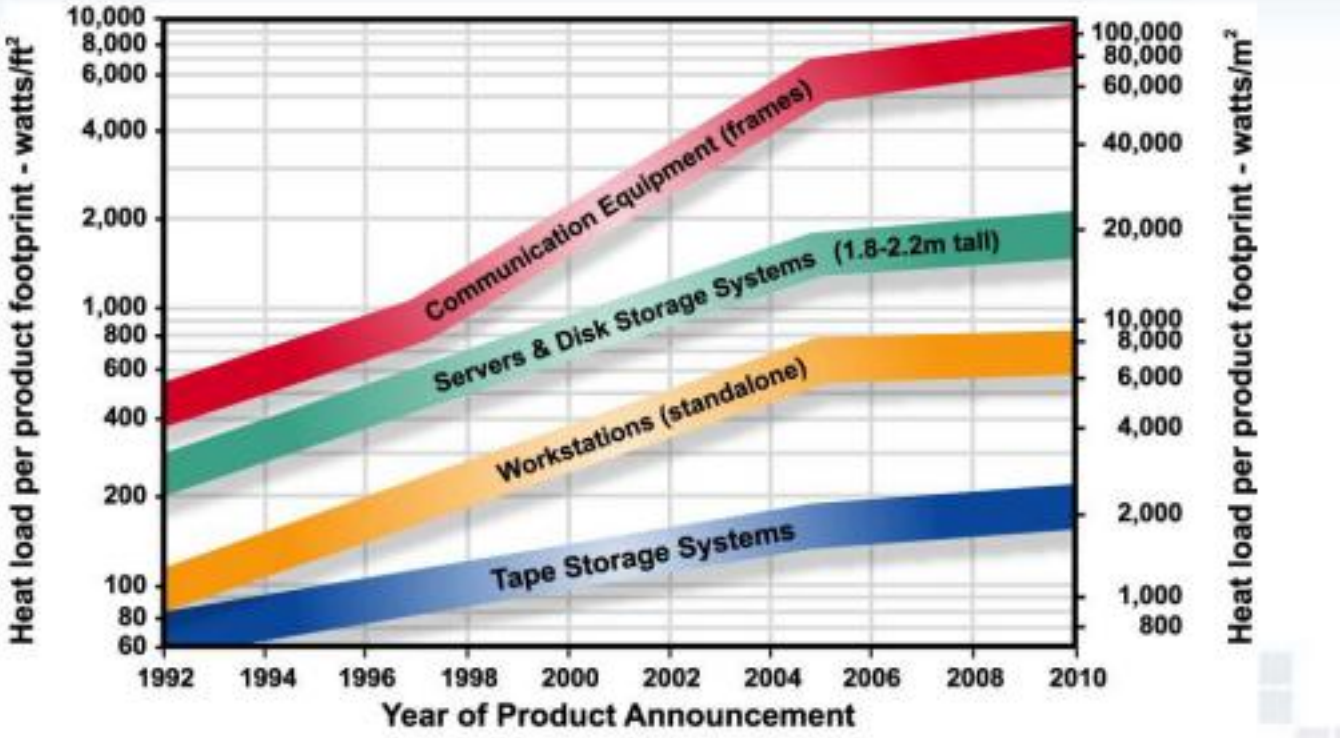

Created by members of the “Thermal Management Consortium on Data Center and Telecom Room”. Participants were from 15 major equipment manufacturers and service providers.

Shift in DC Economics

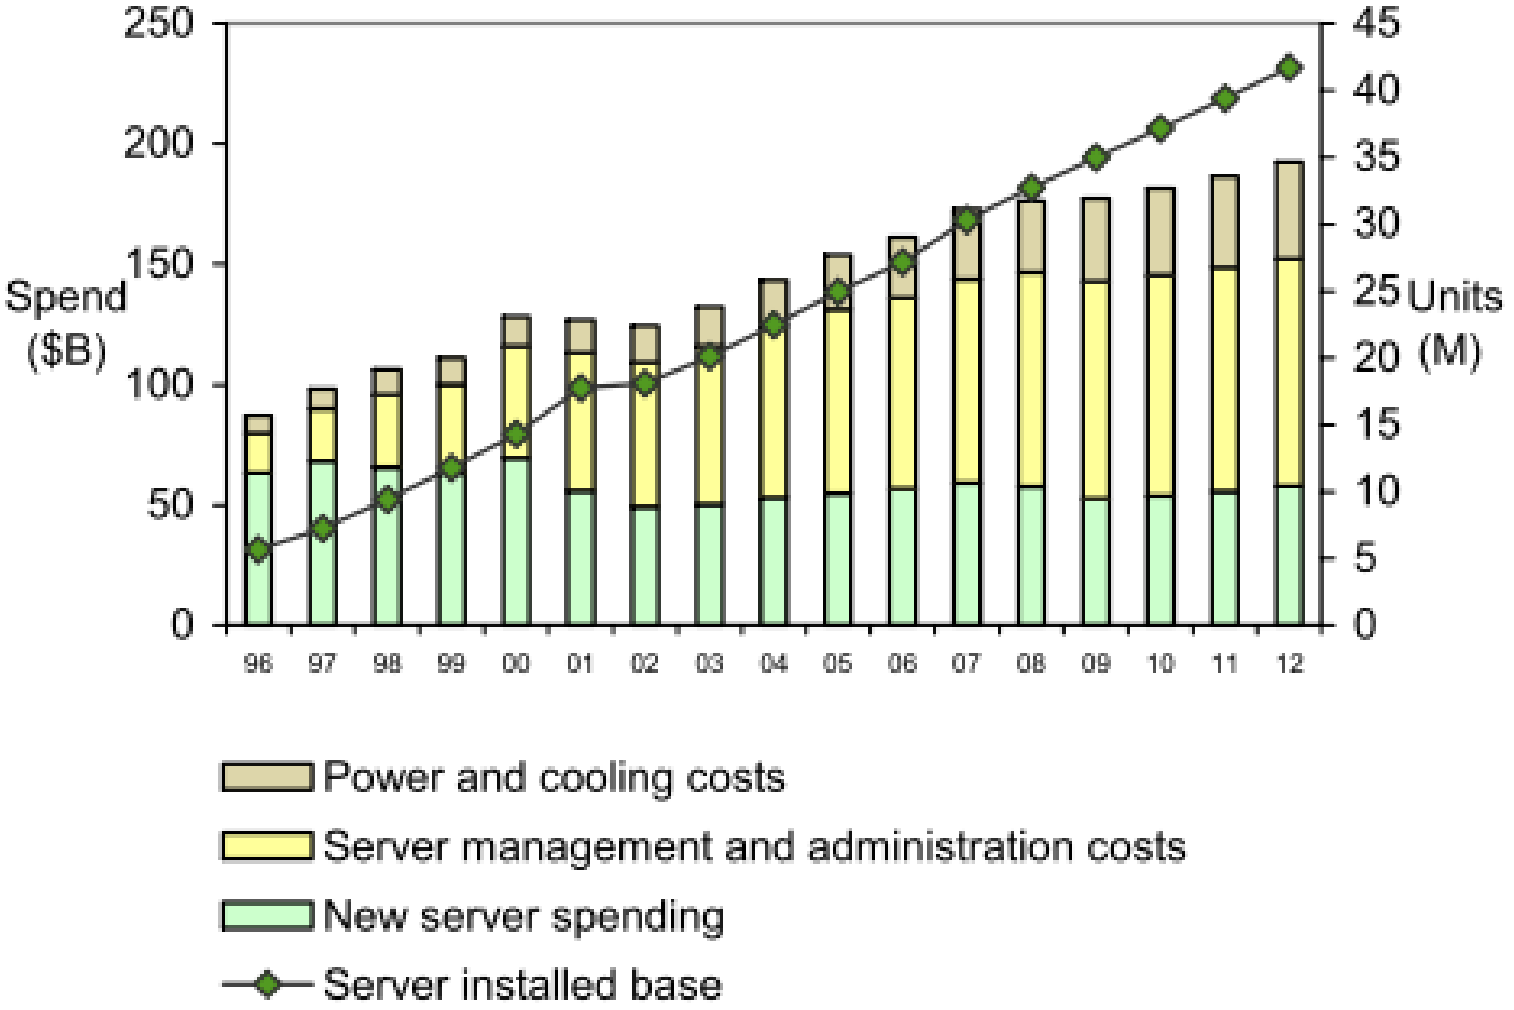

Source: IDC 2009

Breakdown of electricity consumption of a typical data center (source: APC)

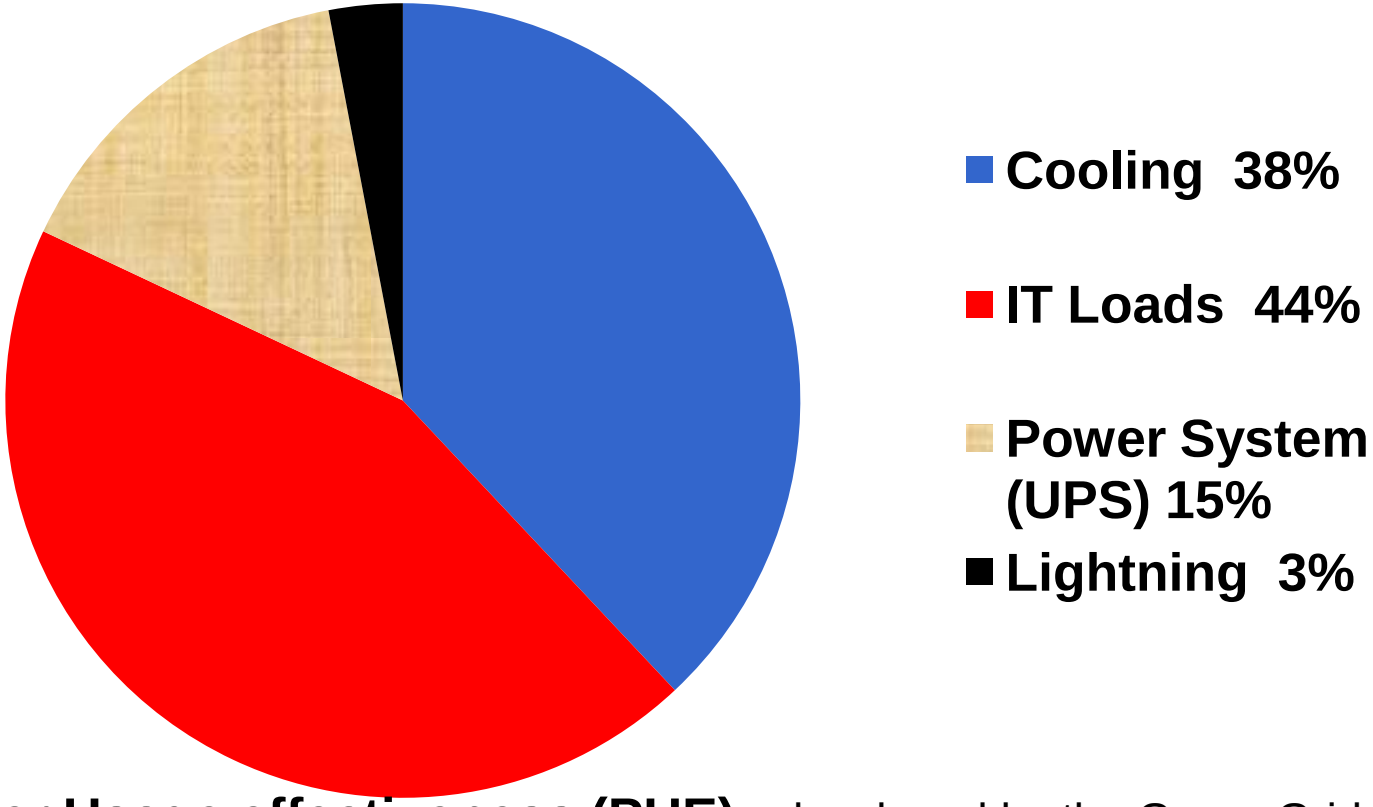

Power Usage effectiveness (PUE) - developed by the Green Grid

The PUE is calculated by dividing the total utility load (the power coming into your facility) by the total IT equipment load (the power consumed by switches, routers, servers and related gear), the lower the PUE, the better.

Running at a full load, HP has calculated that the Wynyard facility has a 1.2 PUE

Saving of electricity consumption

- Building Transformer
- UPS
- Power Distribution
- AC to DC conversion
- DC to DC conversion
- Cooling

1 Watt saved at the processor saves approximately

2.20 Watts of total consumption

Cooling

- Cabinets and racks should be arranged in an alternating pattern to create "hot" and "cold" aisles.
- Perforated tiles in the raised floor of the cold aisles allow cold air to be drawn into the face of the equipment. This cold air washes over the equipment and is expelled out the back into the hot aisle.
- In the hot aisle, there are no perforated tiles, which keeps the hot air from mingling with the cold
- Cooling air recirculation can increase required energy and costs by 10-25%
- Consider auxiliary cooling - rear door heat exchangers (water, Freon) 50-60% heat removal
- Consider to install Thermal Zone Mapping – large vendors service which uses heat sensors and mapping analysis software to pinpoint cooling problem areas

Cooling Issues/Common Mistakes

- The cool air is mixed with the hot air coming out of the equipment
  - Bad physical planning (spaces between racks, racks facing the same direction)
  - Missing blanking panels
  - Wrong placing of perforation tiles
  - Holes in tiles
  - Perforated tiles in places where there is no equipment
  - Large spaces between adjacent racks
- The cooling airflow is obstructed
  - Under floor cabling , junk
- Humidity set too high, different setting for different CRACs (Computer Room Air-conditioning Units)

Cooling II

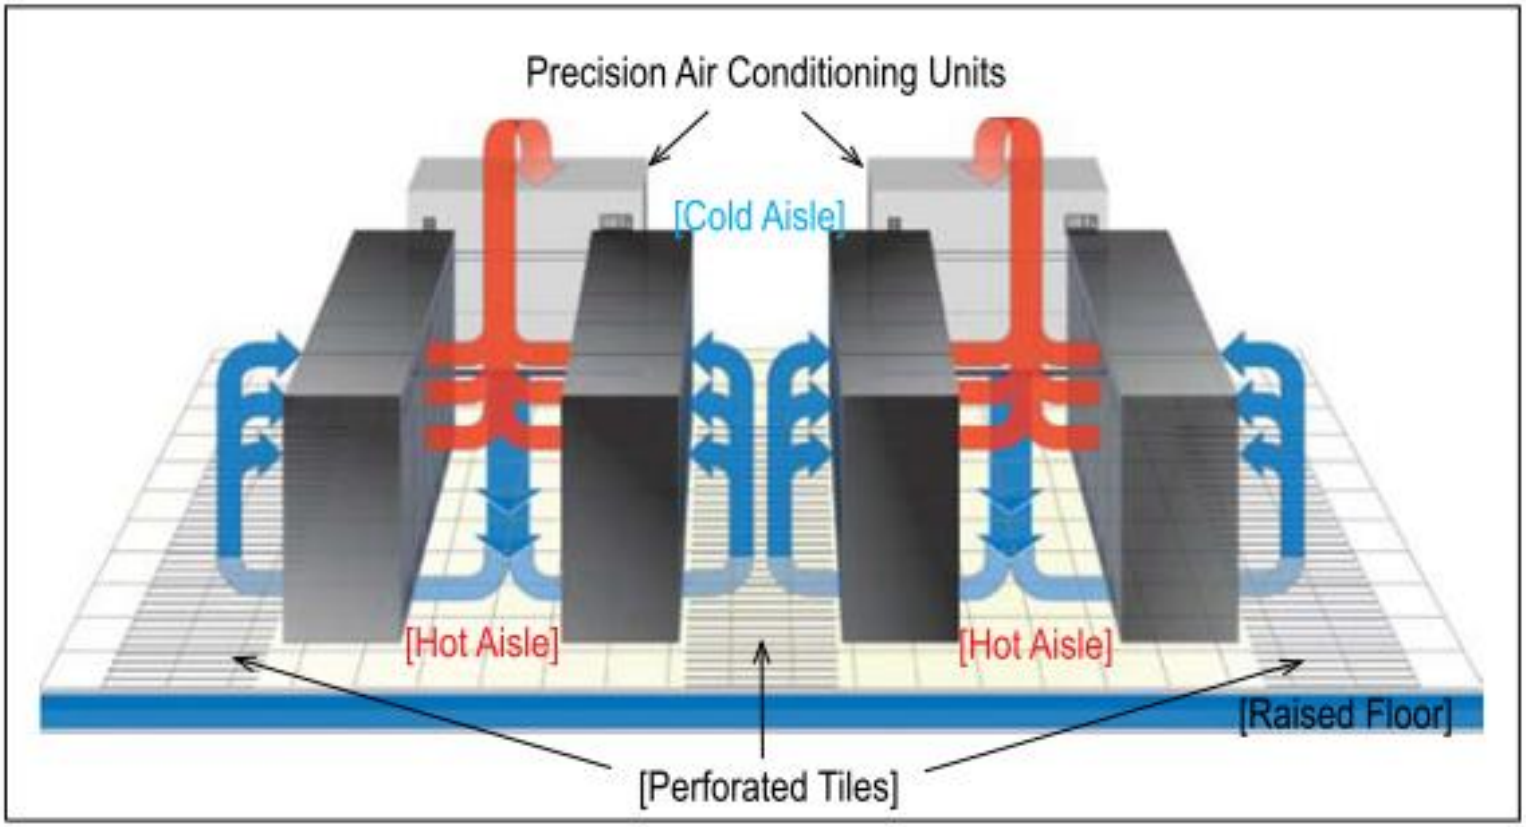

Source: Emerson Network Power

Putting racks in the same direction causes air recirculation, creates hot–spot problems, and increases the operating costs

If the racks and the floor grid do not align, airflow can be significantly compromised

Raised floor implementation using a dropped ceiling as a hot air return plenum

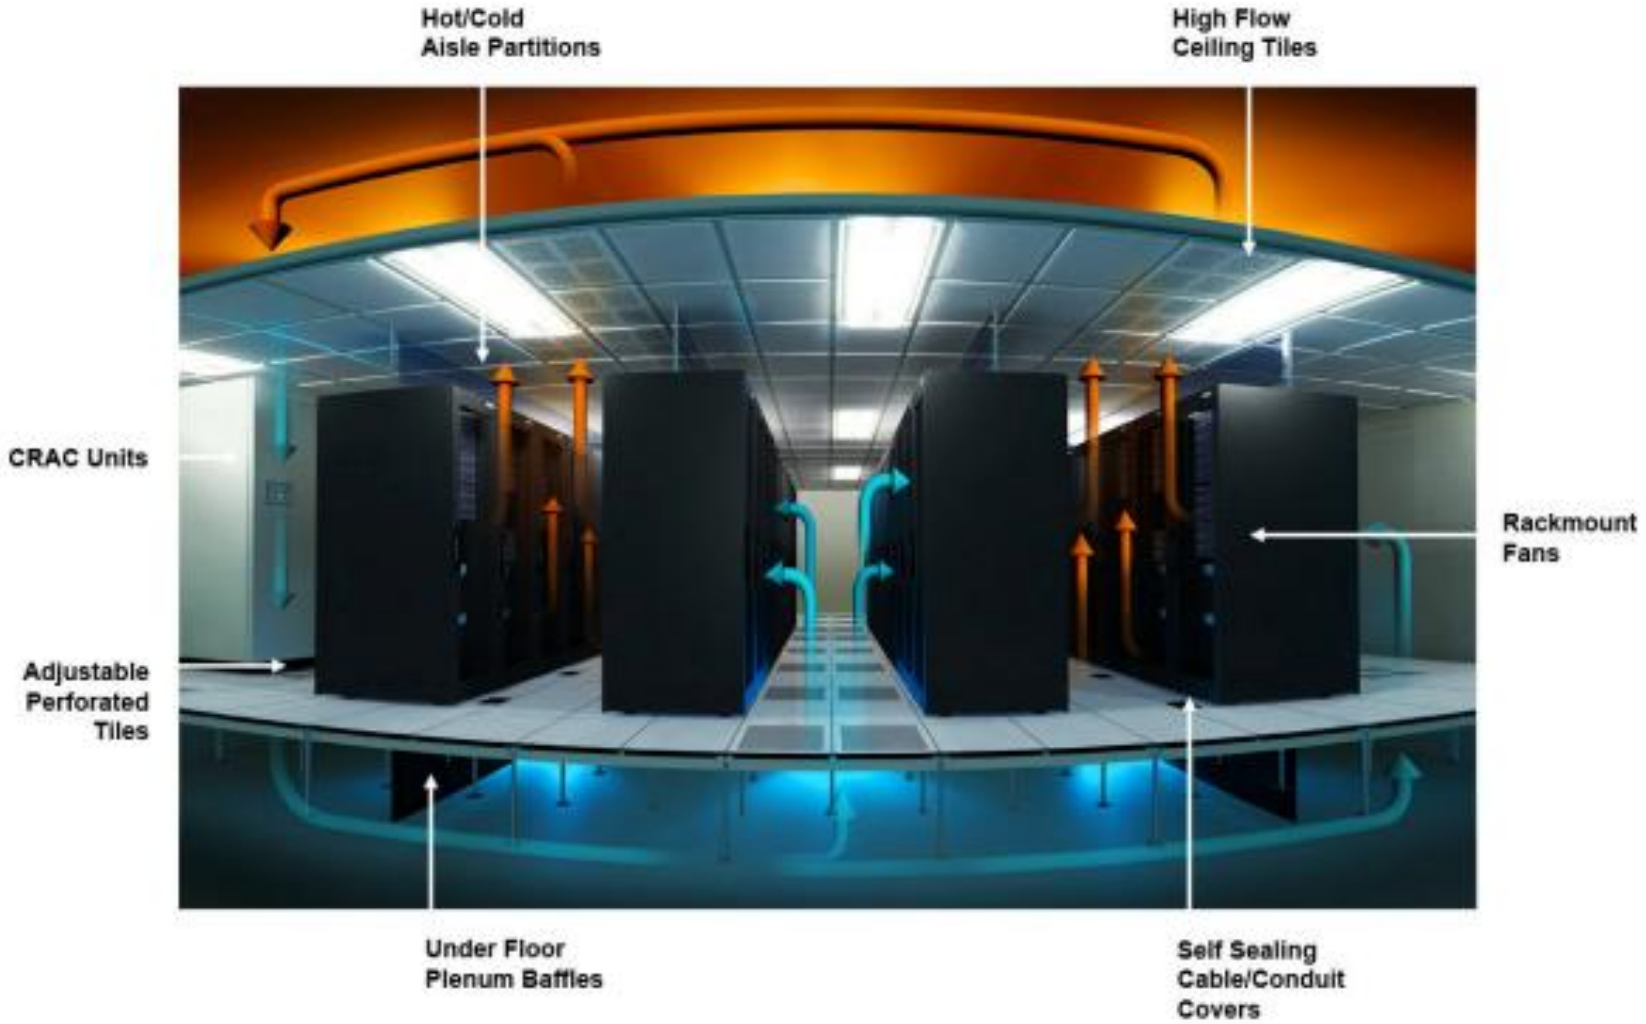

Source: Eaton Corporation

High or low power racks?

- 75% of data center TCO costs are driven by power and only 25% of costs are driven by area.
- TCO costs do not decline significantly with increasing power density, but actually increase after an optimum power density, which is on the order of 6 kW per rack.
- When the power density exceeds 10 kW per rack, the unpredictability of airflow becomes the dominant problem.
- The recommended raised floor in the low-density room is 18 inches, while the high-density room will need a 30-inch height
- Blade servers require ca. 120 cfm per each KW of power rating
- Dedicated high density areas of 1.1-4.3 kW / / m^2 or 3-12 kW per rack require supplemental cooling technologies (rack, row, back-panel, spot cooling, etc.) and not the typical raised floor design.
- Design the system to permit later installation of supplemental cooling devices
- Consider floor loading !

Physical Planning Recommendations

- Every room has two primary layout axes, or directions that the rows can be oriented. The axis selection is one of the most critical decisions in a data center plan and has a large impact on space usage and cooling.
- All rack rows should align with the same axis (i.e., all a rows are parallel, with no perpendicular rows)
- Align the floor grille with the racks
- For plain walls, a minimum of 2 tiles is an acceptable spacing for a row-end; larger data centers often prefer 3 tiles to provide better accessibility.
- Supporting columns should not cause large “open space” gaps.

Raised floor plan

The 42U Universal Server Rack is compatible with a wide variety of servers and rack mountable networking equipment, including Dell, HP / Compaq IBM, and Sun products

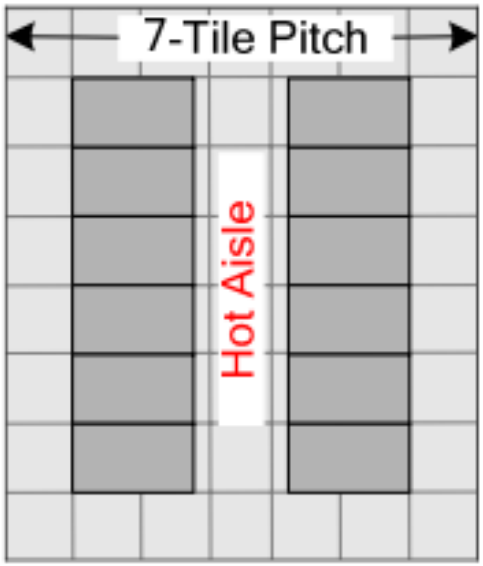

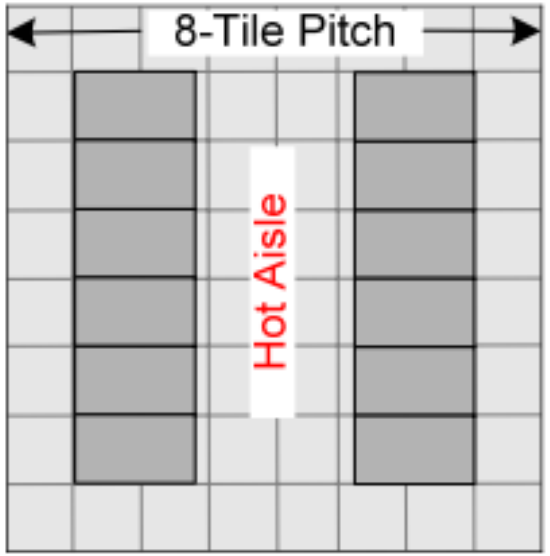

Low power racks

with raised floor air distribution

High power racks

Standard tile 60x60 cm,

42U Universal Rack Dimensions : Width: EIA Standard 19" Rack Rails, Ext. Width:60cm

Height: 78.74" – 2 m, Depths: 100 cm & 105 cm

Saving on Equipment

- Consolidation of DCs
- Deploy virtualization
- Don’t oversize
- Increase utilization percentage
- Analyse how many copies of data exists
- Set printing regulations
- Detect “ghost servers”
- Storage resource management (SRM) programs

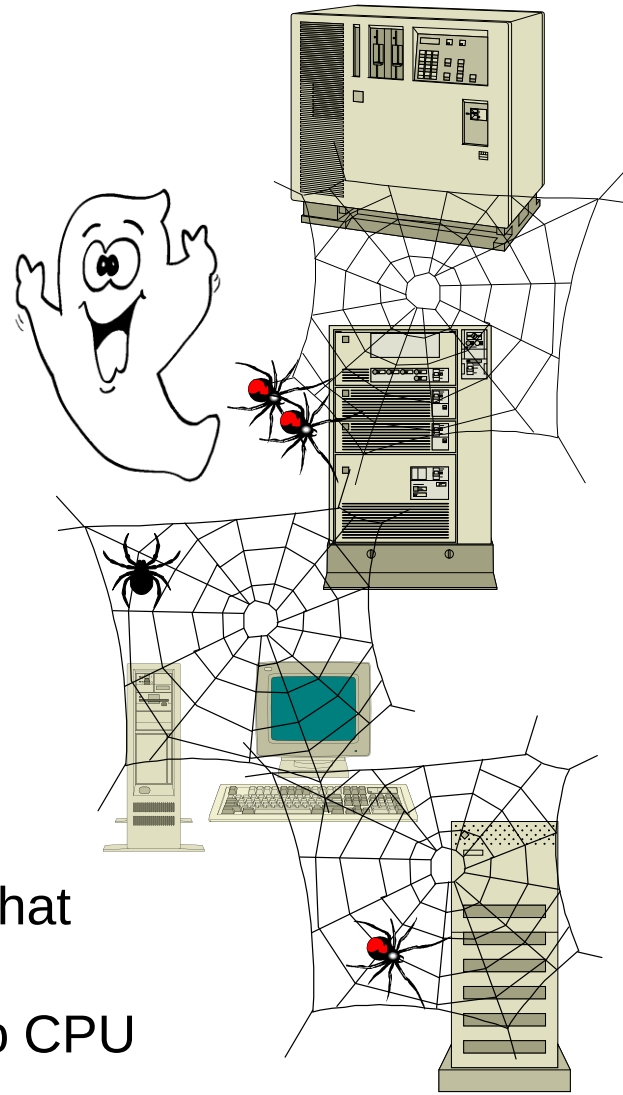

Large vendor making consolidation project found that about 150 of 1,800 servers had no current usable function, and an additional 200 machines had zero CPU utilization.

Storage Resource Management (SRM)

You can't manage what you can't measure !

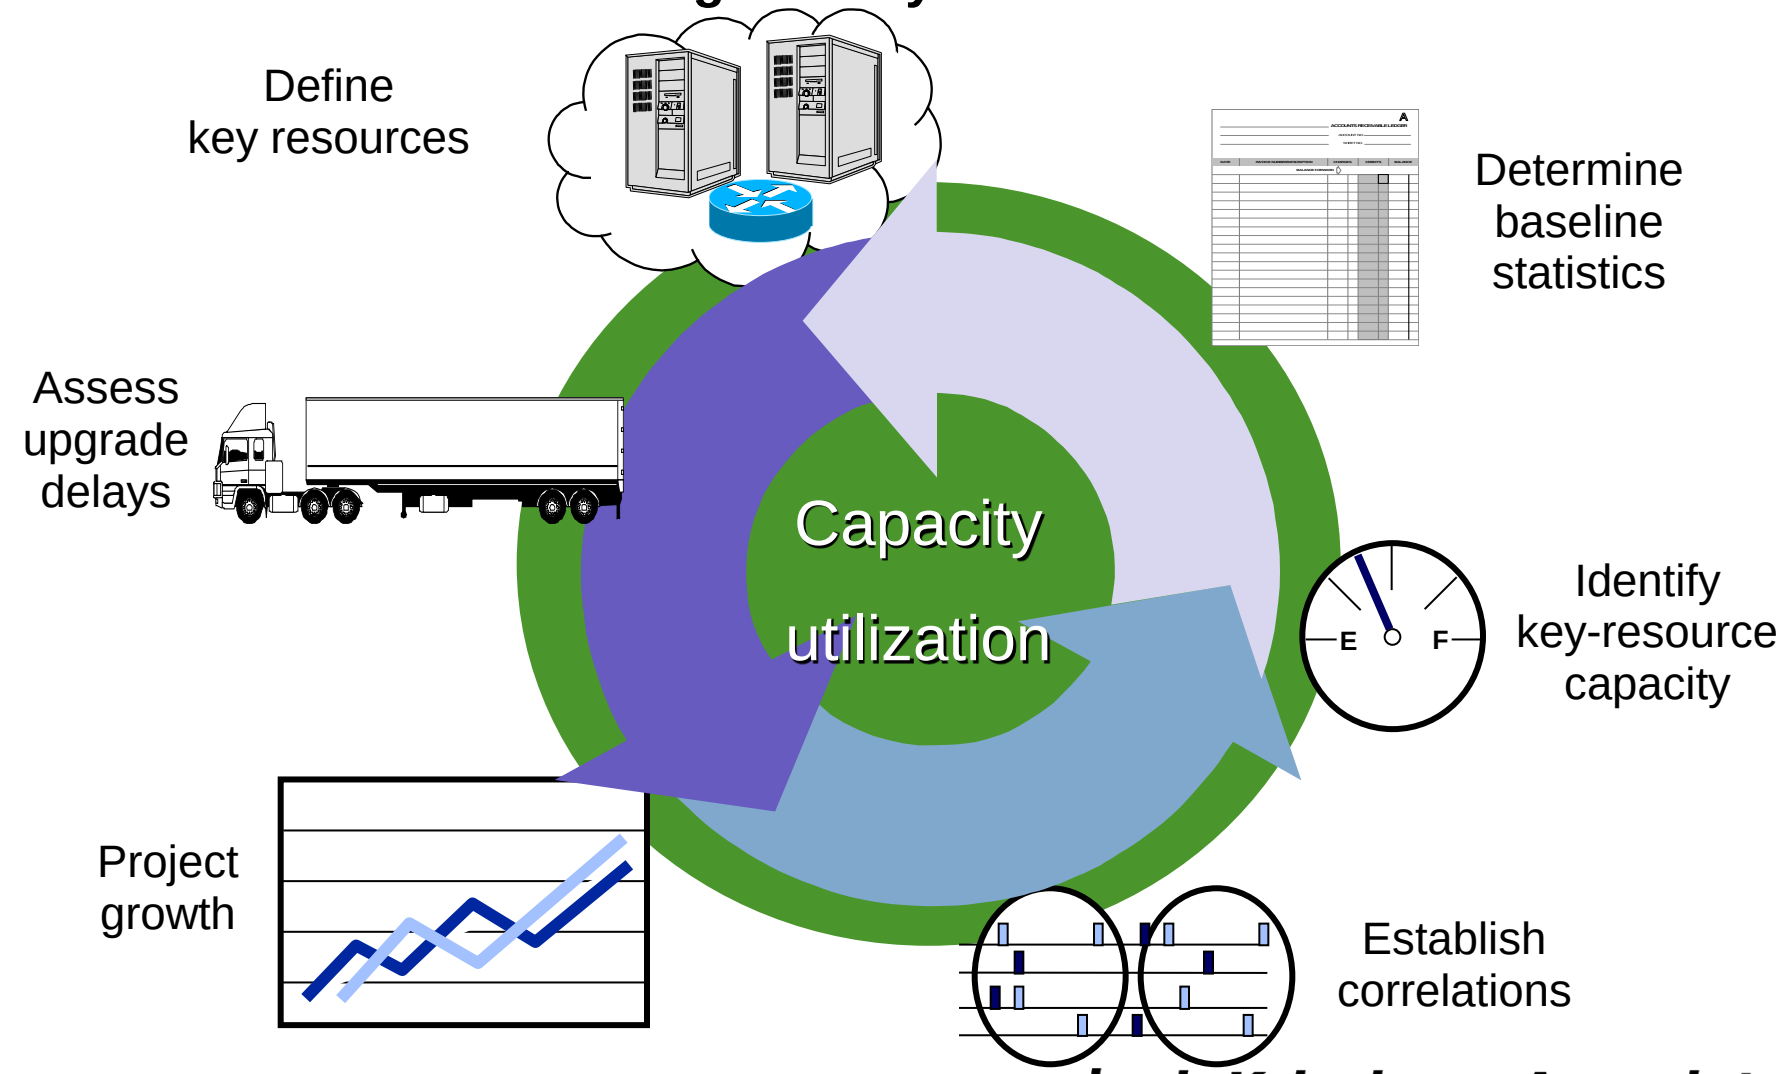

Saving on Servers' Energy consumption

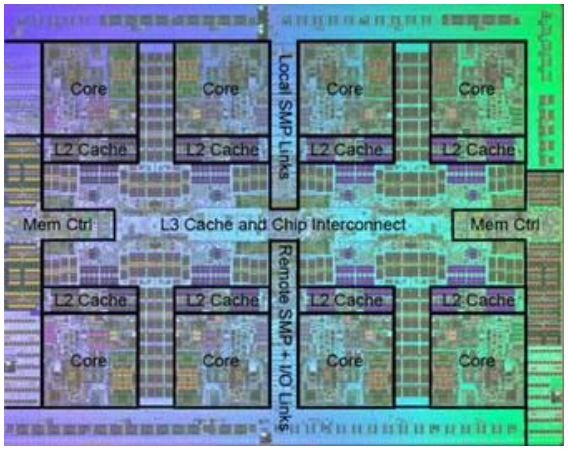

- Equipment refresh, multi core processors
- Switching-off unused processors
- Adaptive cooling: automation that adjusts blower and fan speeds based on system activity
- Dynamic Power Capping-limiting the amount of power consumed by servers

How to Lower the Storage TCO and help to save the Environment?

- Tiered storage
- Storage consolidation within a tier
- De- duplication, compression
- Thin provisioning
- Alternative technologies (tapes, VTLs)
- Virtualization
- New HDDs (larger capacity, smaller forms)

Low hanging fruits

Data storage devices had the highest power consumption growth rate (191%) and the highest overall power consumption.

Tiered Storage

- Not all the data is equal
- In energy vs. capacity, SATA disks are significantly more effective than the FC based disks
- An example: A Seagate Cheetah 300GB 15K FC RPM disk uses **17.6W,** Seagate Barracuda 1TB SATA - **11.6W**.
- 1 TB FC uses 58.7 Watts, which gives ratio 1:5
- 500 GB STEC Zeus uses **8.4 W**

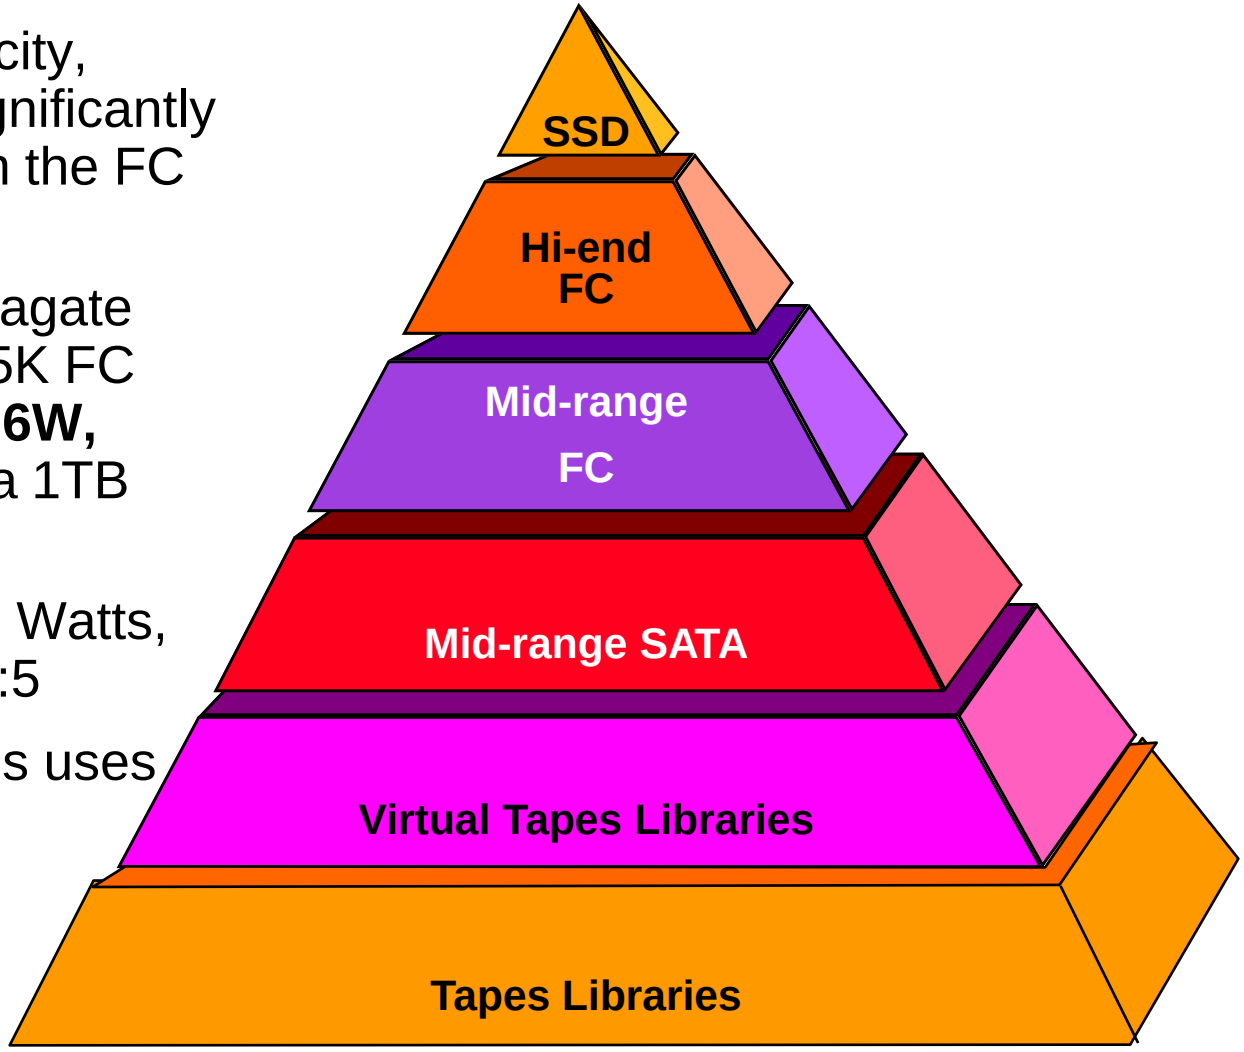

Consolidated Disk Storage

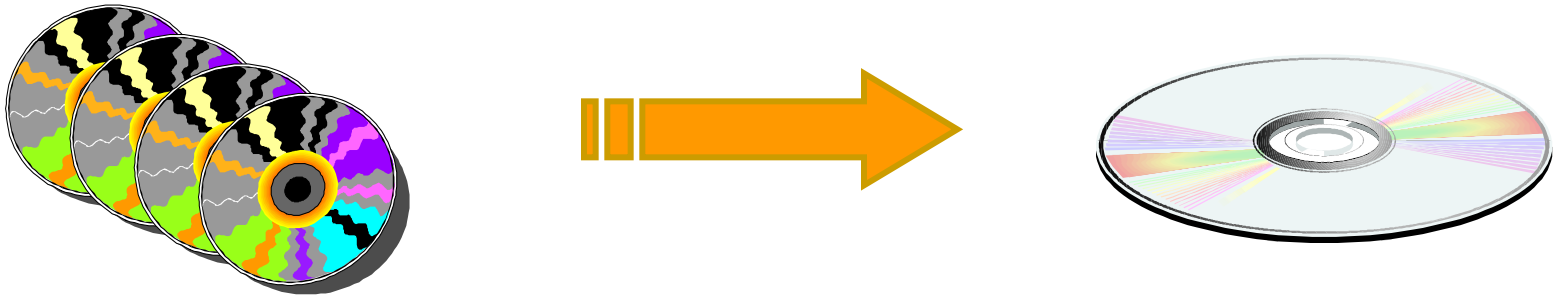

SAN, NAS or High-End Enterprise Storage

- Larger investment in procurement, but
- Better use of “spare capacity,” flexible LUN sizes with dynamic reconfiguration
- Improved availability, security and disaster recovery
- Ability to consolidate and automate backups, tape libraries
- Multiplatform data transfer
- More efficient access, sharing and distribution of information throughout the enterprise
- **Lower storage management costs**

Storage Designed for Archiving

- Disk Archiving subsystems
- Designed specially for „persistent“ data
- Low cost (usually SATA) HDDs
- Modular built, scale well,
  cluster connected
- FC or LAN attached
- Contain server with fast search and
  retrieval software
- Metadata on storage
- May, or not include
  imbedded CAS mechanism
- EMC‘s Centera, IBM‘s DR550, Sun‘s
  „honeycomb“,
  Hitachi Content Archive Platform (HCAP)
  Hp‘s ExDS 9100

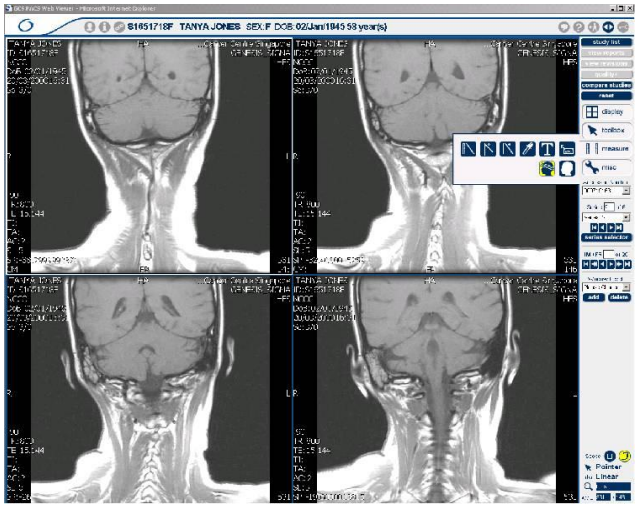

MAID - Massive Array of Idle Disks or other Power Saving Features

- Disk Archiving subsystem
- Designed specially for „persistent“ data
- SATA HDDs with RAID implementation
- Turning HDD group off, turning electronic off, putting HDD in “stand-by, “parking the heads”
- Currently delivered by Copan (SGI), Fujitsu (ETERNUS), and DATABeast from Nexsan
- HDS Power Savings Storage Service (PSSS)
- EMC CLARiiON CX-4 spin-down feature
- Lower power consumption and heat dissipation
- TCO - Lower than disk subsystem but higher than tape

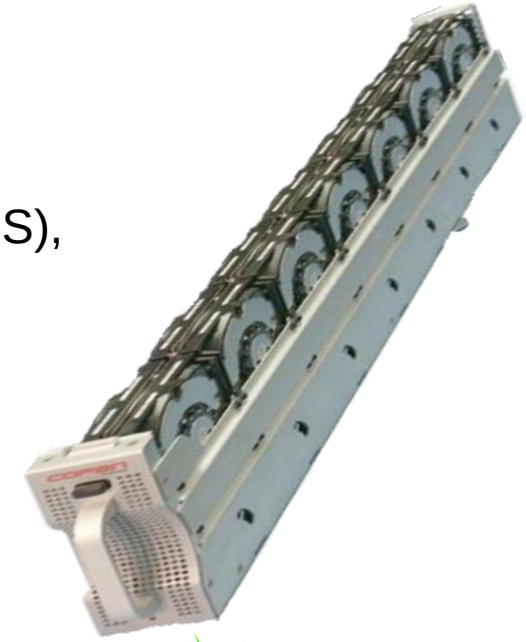

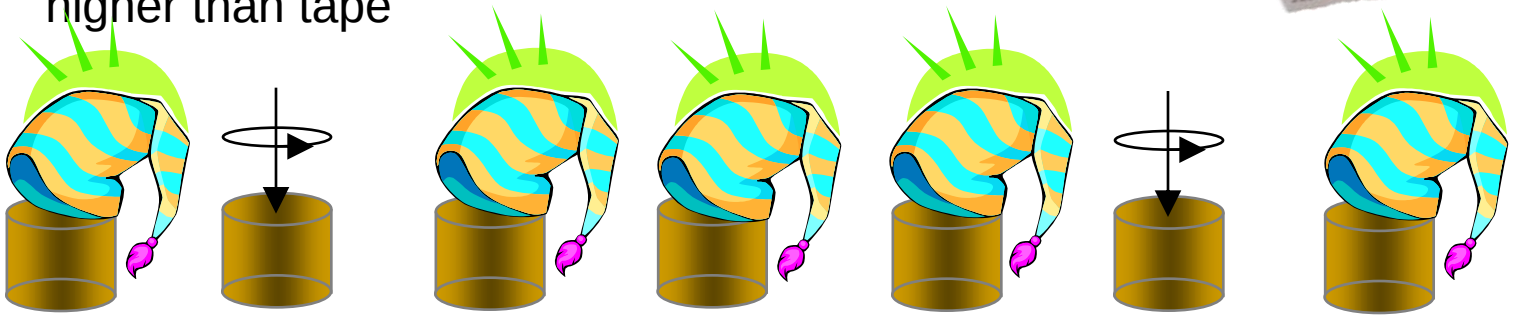

De-Duplication

- Elimination of redundant data for backup
- The data is divided to blocks, repository index is built from the data contents and stored in Metadata catalog
- Only „Unique“ blocks are written to a VTL on a disk
  - Savings in data space (factor 1:10-20), faster data transfer, less bandwidth required (B/R over WAN)
- DD and DDX from DataDomain (EMC)
- ProtecTIER from Diligent - IBM (OEMed by HDS)
- DeltaStor from Sepaton (hp)
- Axion from Avamar (EMC)
- Rocksoft (Quantum DXi)
- NetApp, Major Backups, Ocarina

Users‘ requirements

- Easy integration with the existing backup
- Easy deployement and management

Users‘ concerns

- Where to put it?
- Integration with existing applications
- Data integrity
- Software reliability (loosing access to all backups)
- Performance, scalability
- Costs (ROI)

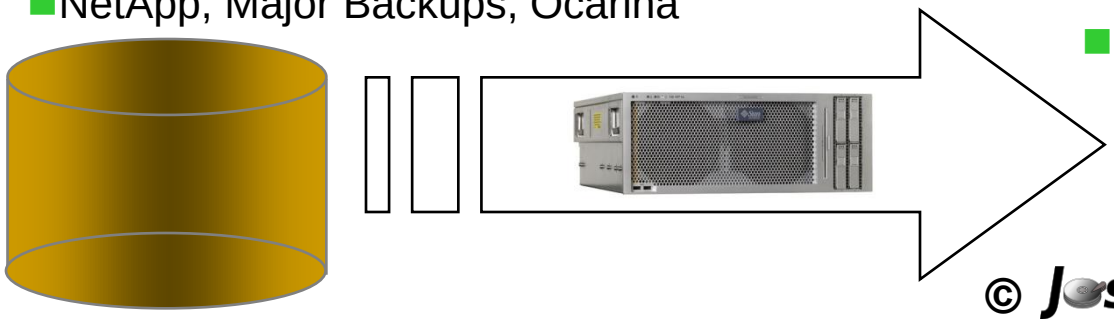

De-Duplication Example

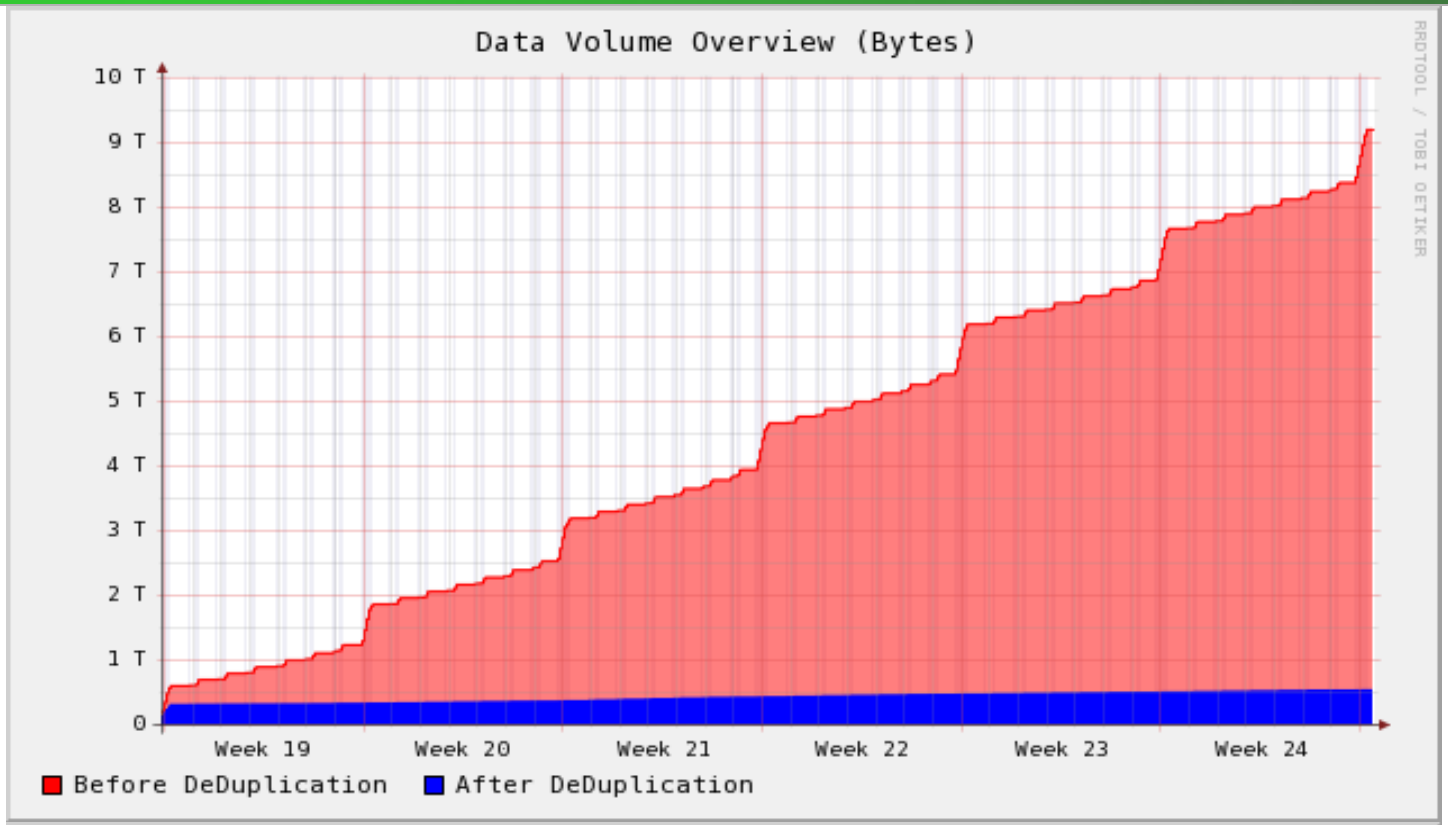

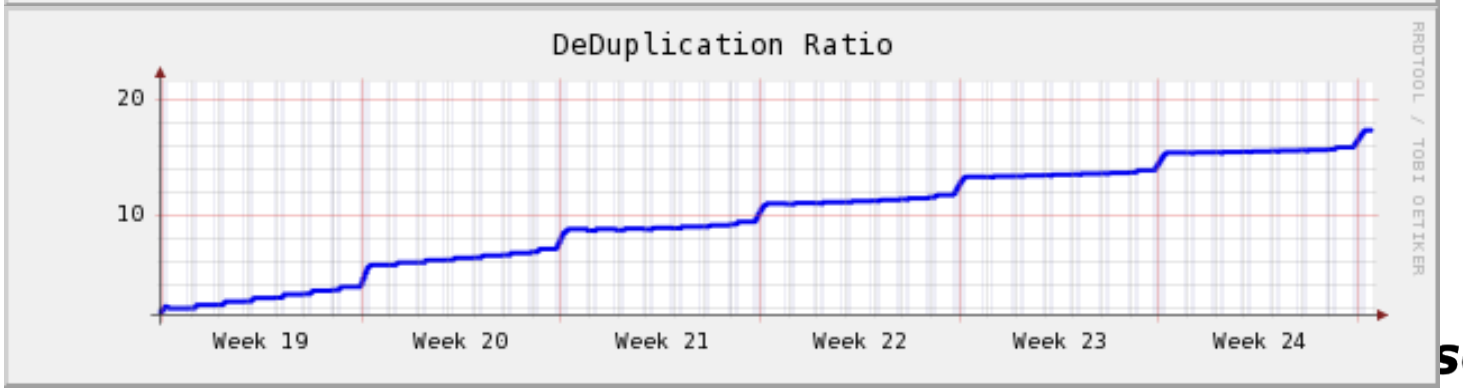

sociates

„Thin provisioning“

- Virtual image, data blocks scattered across the subsystem
- Virtual LUNs span over different physical HDDs
- Capacity is allocated only when required
- Less time required for storage allocation
- Non-disruptive LUNs creations and expansions
- „Uniform“ performance
- Reduced equipment and energy costs
- StorageTek‘s „Iceberg“ (RVA, SVA), 3PAR, Compellent , XIV (now IBM) Nextra, DataCore Software Corp., LeftHand Networks Inc., Network Appliance Inc., IBM SVC 4.3
- High-end storage: EMC DMX, Hitachi‘s USP V and partners, IBM DS8000

To Consider

- More attention
- Oversubscription
- Unused space reclaiming
- “charge back” mechanisms

Virtual Tapes – tiered storage at lower costs

Control, HSM Buffer

- Popular in mainframe, emerging in UNIX
- Faster performance
  - "Instantaneous load" of cartridges
  - Disk transfer rate
- More tape (virtual) drives, potentially fewer physical drives
- Better exploitation of media (compression and compaction)
- IBM's (TS7700, TS75x0), STK's VSM, Eternus' CS

- ADIC's Pathlight VX 450 Bus-Tech inc., Diligent's VTF (IBM), FalconStor (Dell, EMC, Sun), NetApp NearStore VTL ,Quantum, hp (Sepaton) Bus/Tech (MDL), EMC DLm (MF)…

Enterprise HDDs

- Current main stream 3.5“ 600, 450, 300 GB/ 15 000 RPM FC or SAS 1,2 TB /7200 RPM SATA
- Requirements; smaller size, less energy, less “locked” data per spindle
  - 2.5“ form factor very slowly replacing 3.5“

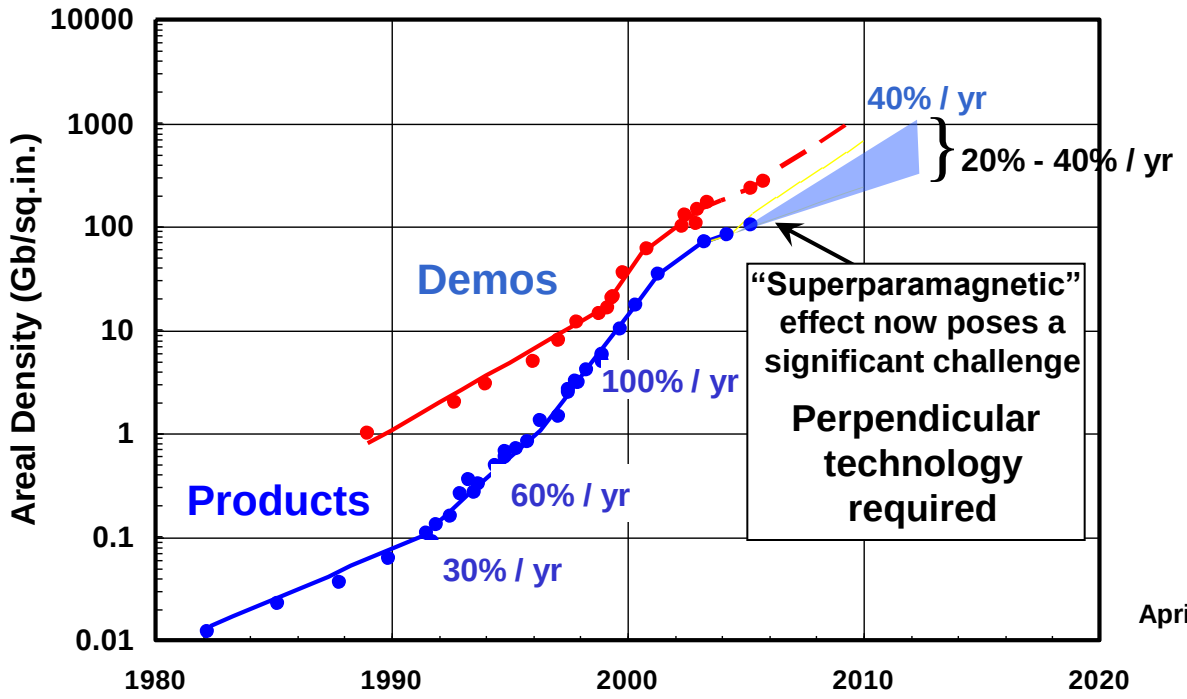

Hitachi Ultrastar 15K600 requires 10.7 W

Ultrastar A7K2000 requires 5.6 W

Capacity to energy ratio 1: 6.36

2.5“ Ultrastar C10K300 (SAS) Requires 3.4 W

Capacity to energy ratio 2: 1.57

Source: Hitachi Global Storage Technologies

Re - emerged; Solid State Disks

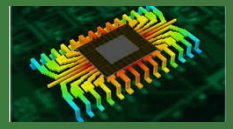

- Semiconductor memory emulating disk
- No mechanical movements, no disk rotation and seek latency delays
- Internal UPS and backup HDDs

D-RAM SSDs and Flash HDDs

- RamSan from Texas Memory (D-RAM)
- Integrated in IBM blade servers HS21 XM, Fusion-io as PCIe plug-in module (Flash HDDs)
- In all high-end subsystems, major mid-range
- ☺ Fastest storage subsystem performance x 100 but only in random read
- ☺ Improved server utilization
- ☺ Much less energy consuption
- ☺ Less sensitive to vibration, schocks
- ☹ Highest costs
- ☹ Limited number of read/write cycles

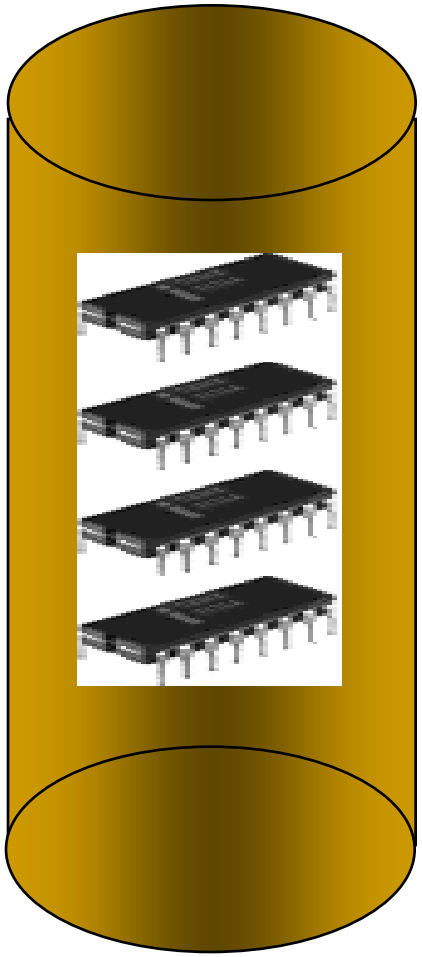

Monitoring DC Energy and Cooling

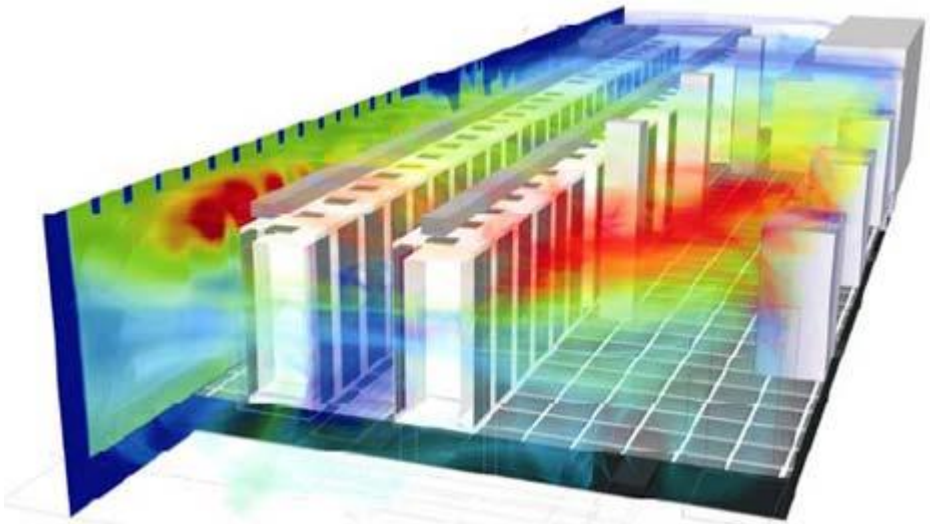

- Thermal Zone Mapping enables to see three- dimensional model of data center cooling effectiveness
- Helps to optimize cooling and energy requirements
- Provided by major IT vendors or leading data center equipment manufacturers
- New services: Energy Efficiency Assessment Services provide reports on total data centre input power, IT load, Power Usage Effectiveness (PUE), and power and cooling conditions. Audit and recommendations

Recommendations

- ✓ Take pro-active activities to DC energy consumption
- ✓ Consider energy costs in calculating real TCO and use it in business-case justifications for new equipment
- ✓ Install sensors and dashboards to measure and monitor energy efficiency in your DC
- ✓ Analyse your cooling infrastructure
- ✓ Check for unused cables, missing “blank” panels
- ✓ Put the “blade “ servers cooling on UPS
- ✓ Plan with expansion for future needs (in particular cooling)
- ✓ Careful physical racks planning is mandatory, poor planning will cause high future expenditures
- ✓ Plan for longest possible racks’ rows*
- ✓ Establish closer working relationships between facilities and IT managers

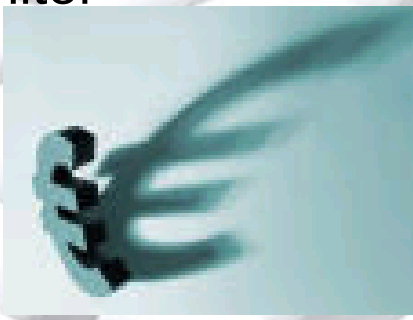

* Verify fire protection regulations, provide emergency exits

Josh Krischer & Associates GmbH

A selection of our coverage areas:

- Procurement & price evaluations
- Enterprise storage
- Mid-range storage
- Disaster recovery techniques
- Data center consolidation
- Data center design
- Mainframes

josh@joshkrischer.com
Tel +49 6251 580626
Mobile + 49 172 6203272
Fax +49 6251 67952

http://www.joshkrischer.com