CSE 548 P: Computer Systems Architecture

Warehouse-Scaling Computers and Data-centers.

Spring 2017

Luis Ceze

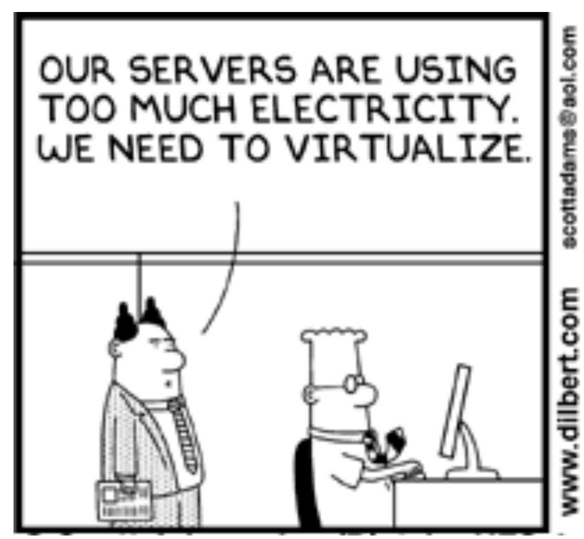

Announcements

- Don’t forget to sign forms!

A datacenter has 50 - 250 containers
A container has 1,000 - 2,000 servers

Velcro

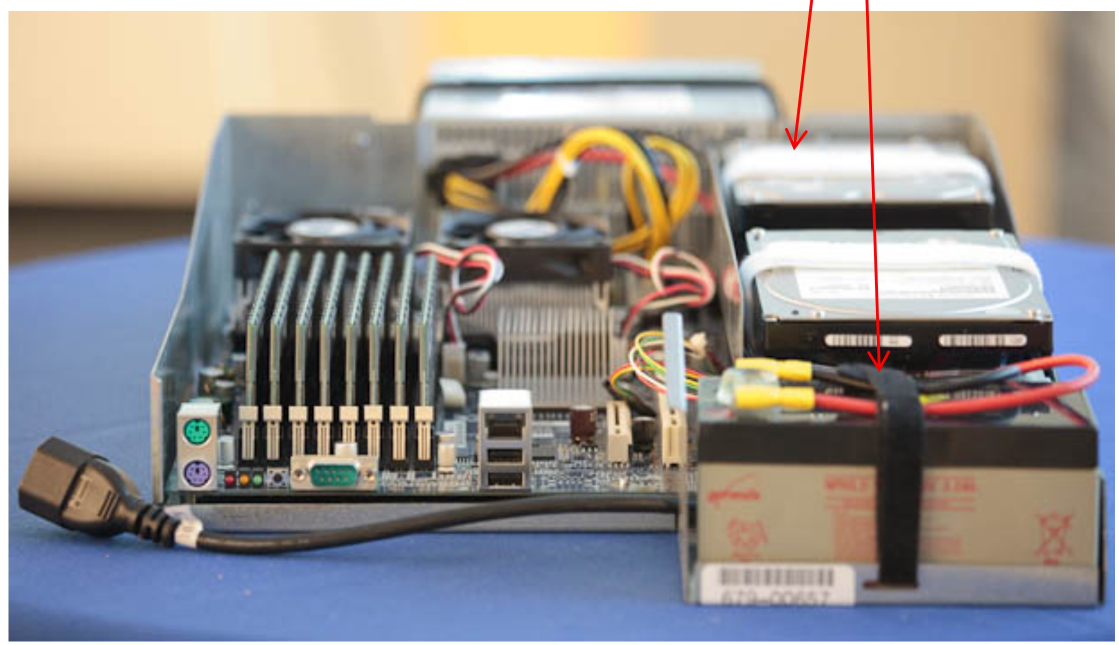

WSCs vs. standard datacenters

WSCs belong to single organization, relatively homogenous HW and system SW, and common management layer

- Traditional datacenters heterogeneous: number of small and medium applications on dedicated hardware
- HPC clusters more special-purpose and batch-centric

The datacenter is the computer

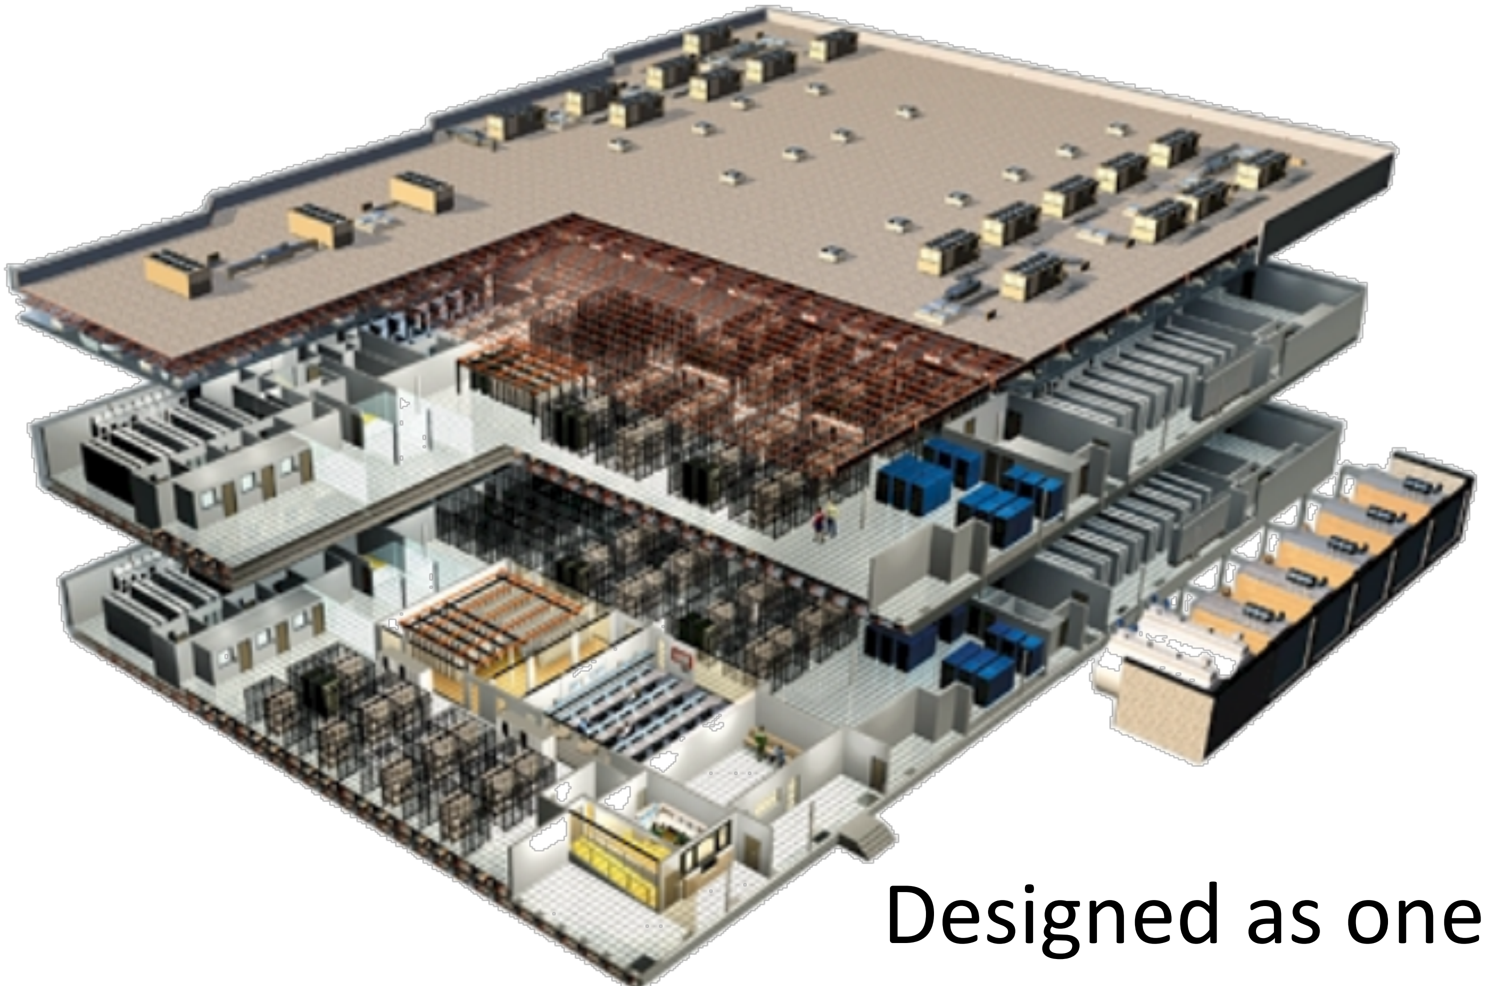

Designed as one machine

Building, electrical and cooling infrastructure, servers, networking, storage, …. (and software)

Data center infrastructure

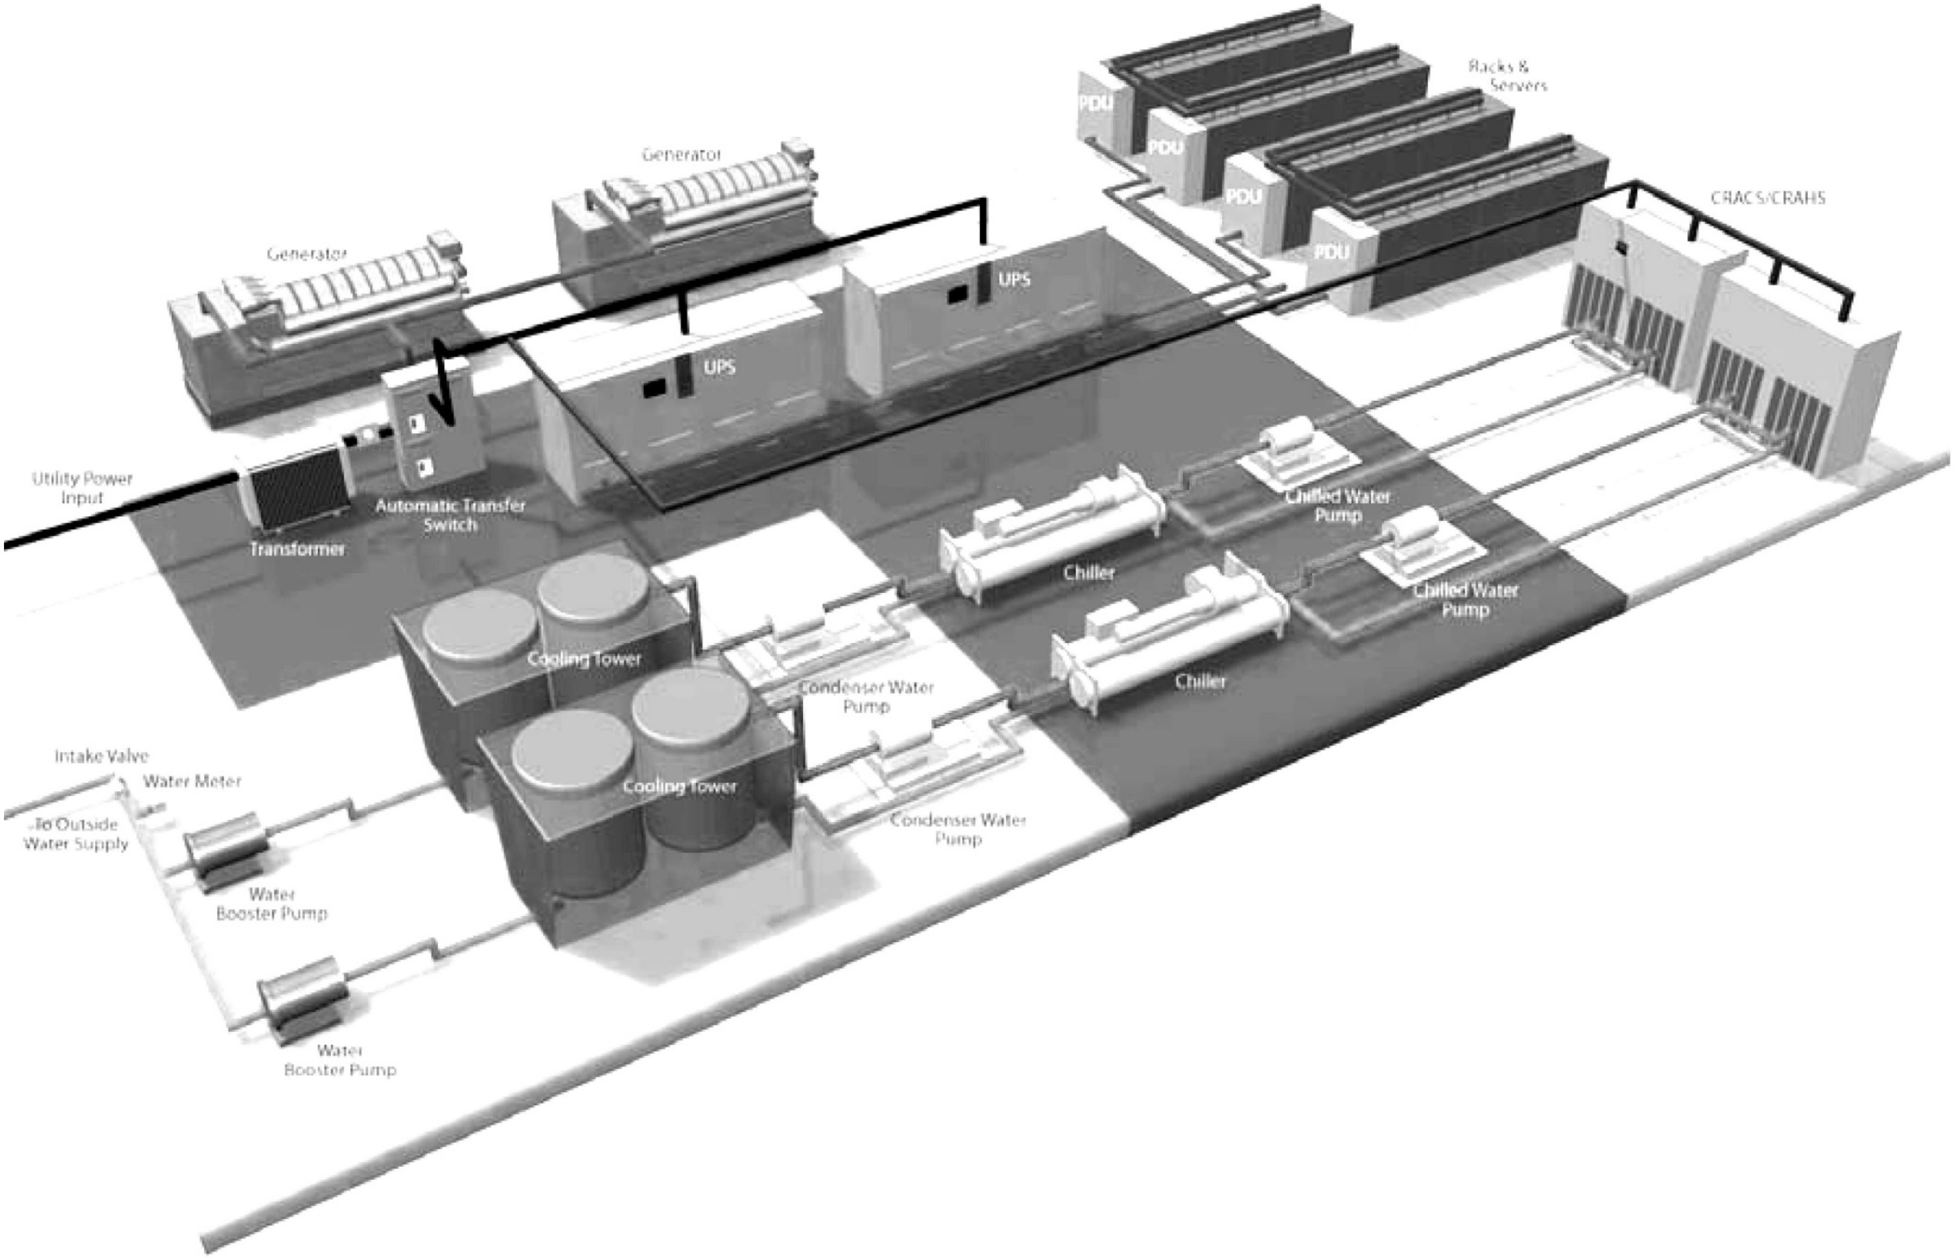

Data center infrastructure

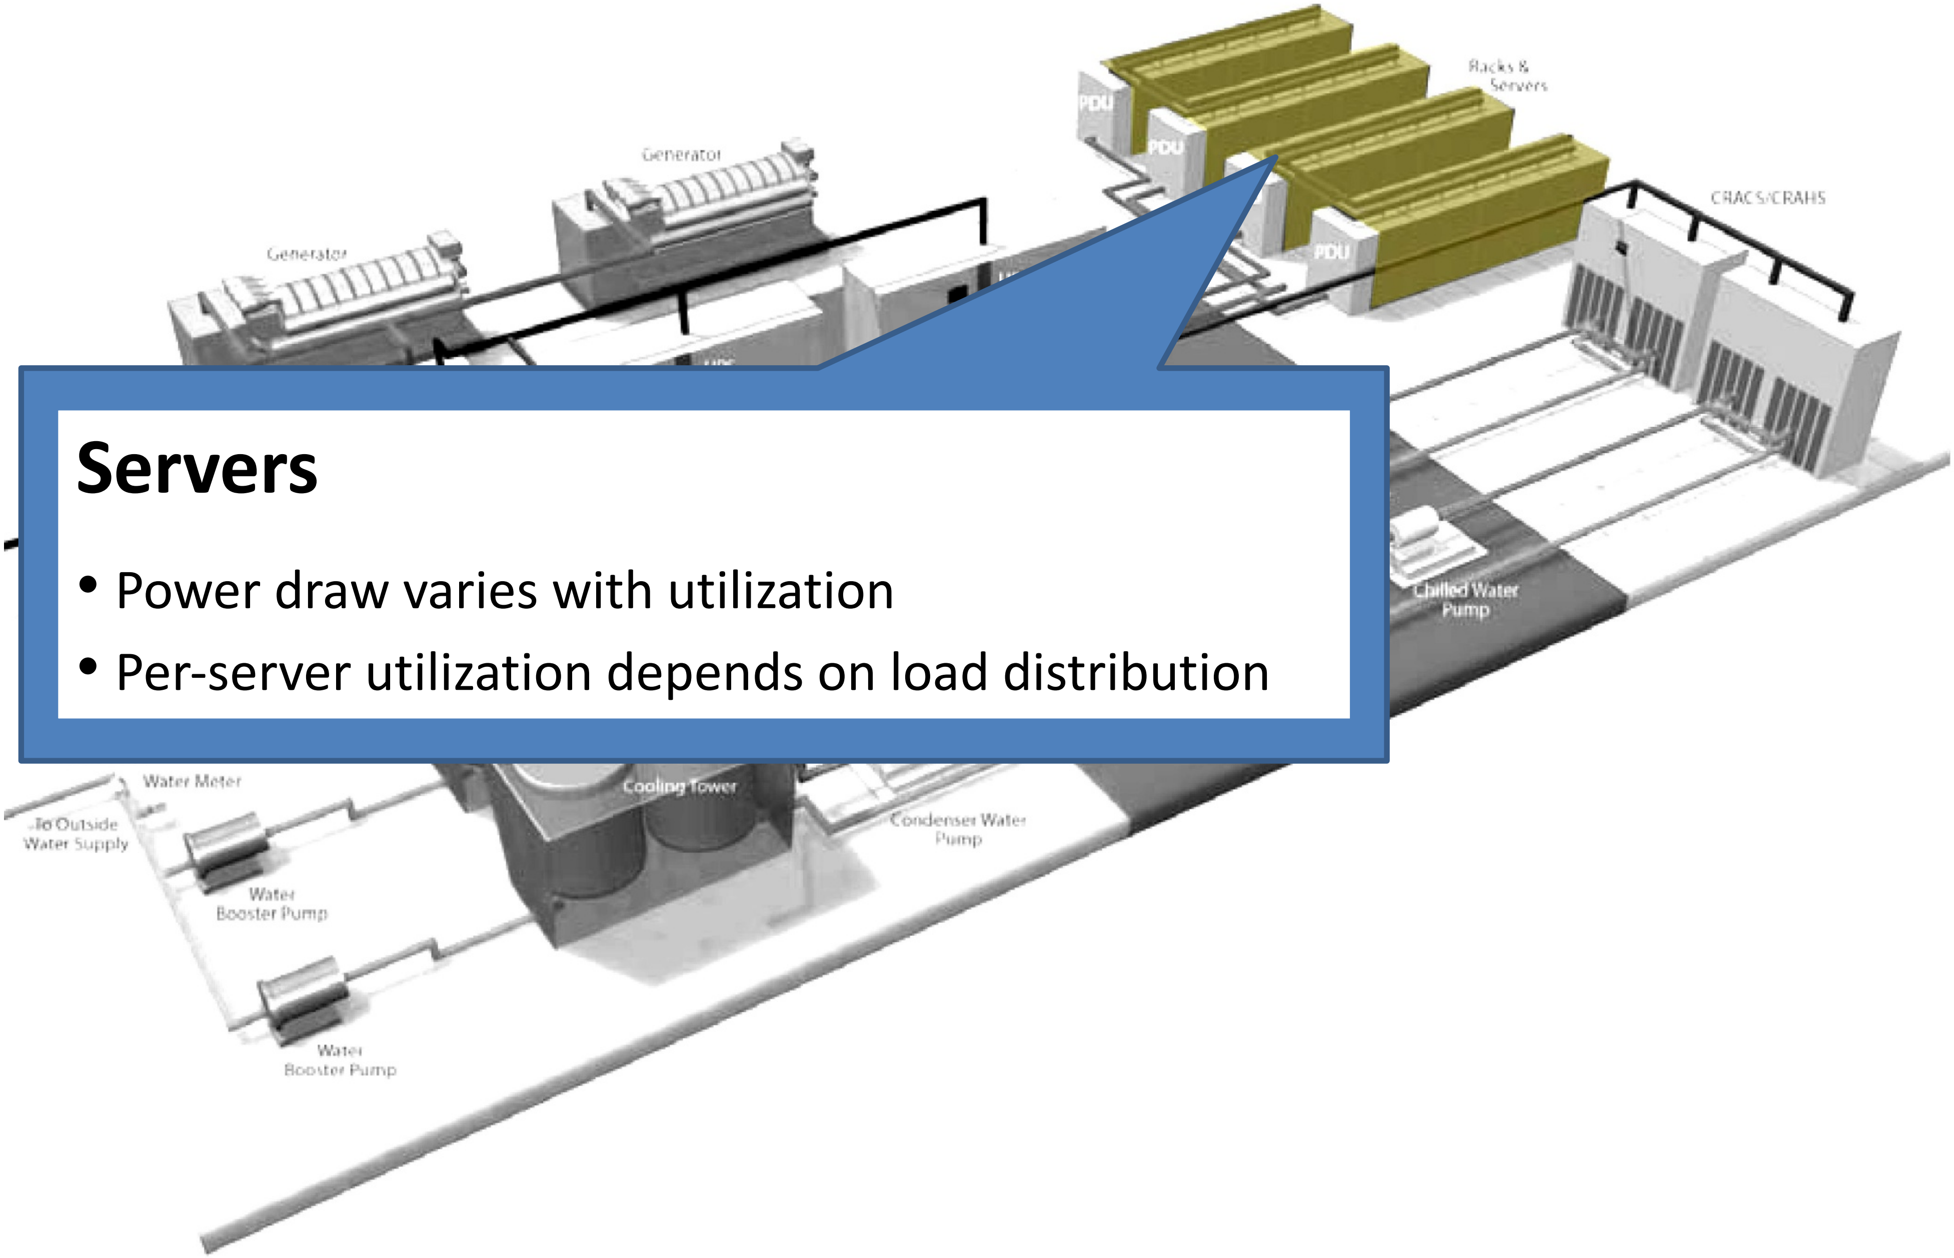

Servers

- Power draw varies with utilization
- Per-server utilization depends on load distribution

Data center infrastructure

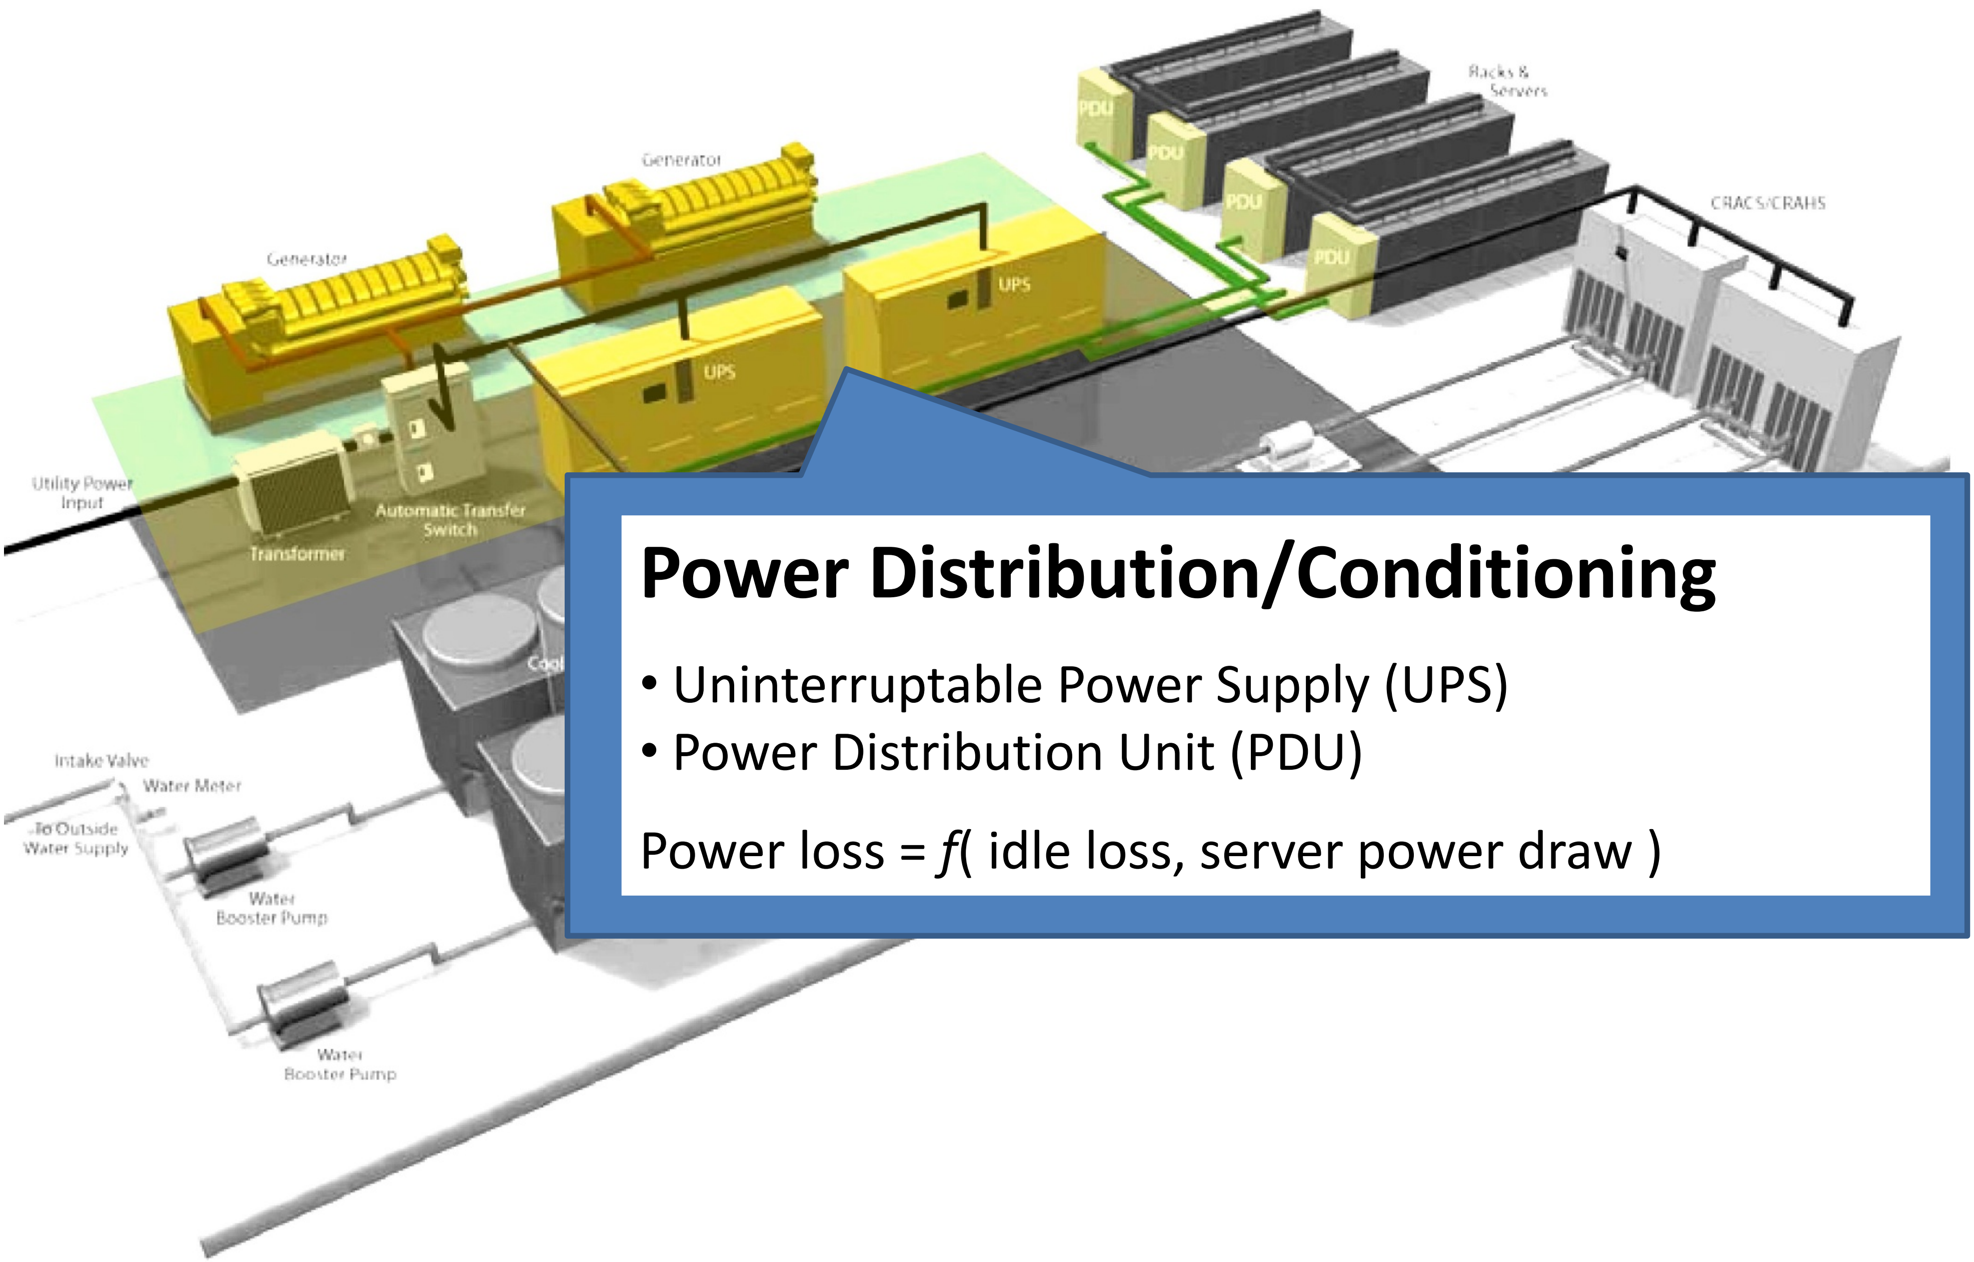

Power Distribution/Conditioning

- Uninterruptable Power Supply (UPS)
- Power Distribution Unit (PDU)

Power loss = *f*(idle loss, server power draw)

Data center infrastructure

Cooling

- Computer Room Ai
    Removes heat g
    Power = *f*(fan sp
- Chiller
    Supplies chilled
    Power = *f*(outsi

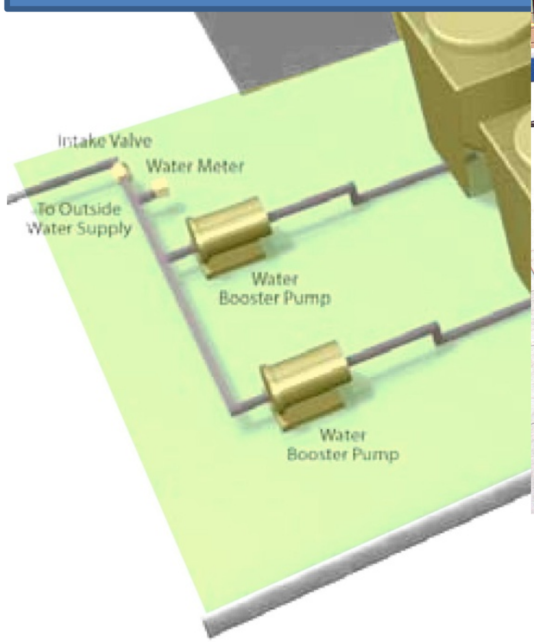

Example: Google

- Datacenter at The Dalles, Oregon
  - Moderate climate, cheap hydroelectric power, near internet backbone fiber
  - 75000 square feet

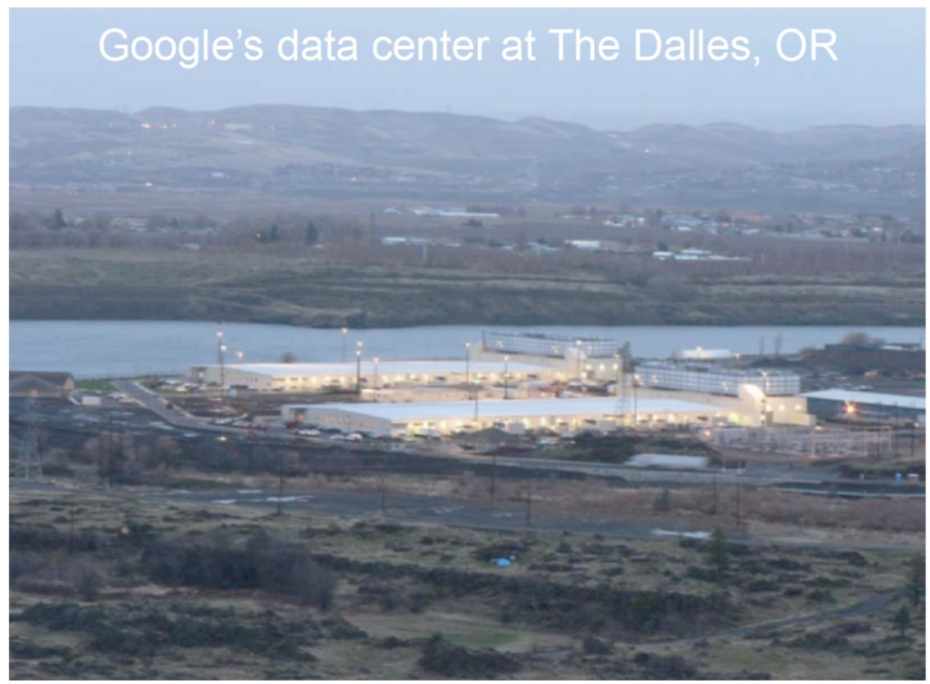

Another WSC example…

- MS Quincy Datacenter
  - 470k sq feet (10 football fields)
  - Next to a hydro-electric generation plant
    - At up to 40 MegaWatts, $0.02/kWh is better than $0.15/kWh ☺
    - That’s equal to the power consumption of 30,000 homes

Data center workload classes

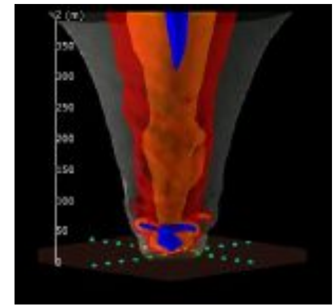

- HPC (e.g., weather sim.)
  - Carefully-orchestrated communication
  - Favors specialized supercomputers

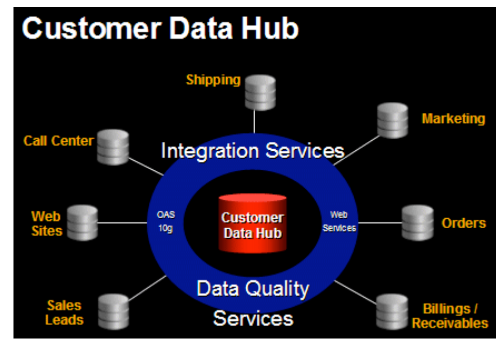

- Scale-up (e.g., OLTP)
  - Strong ACID requirements
  - Favors single large-scale MP
- Scale-out (search, web services)
  - Loose consistency requirements
  - Favors massive commodity clusters

WSC Scale/**Rapid** Growth

	Aug-04	Mar-06	Sep-07	Sep-09
Number of MapReduce jobs	29,000	171,000	2,217,000	3,467,000
Average completion time (seconds)	634	874	395	475
Server years used	217	2,002	11,081	25,562
Input data read (terabytes)	3,288	52,254	403,152	544,130
Intermediate data (terabytes)	758	6,743	34,774	90,120
Output data written (terabytes)	193	2,970	14,018	57,520
Average number of servers per job	157	268	394	488

Figure 6.2 Annual MapReduce usage at Google over time. Over five years the number of MapReduce jobs increased by a factor of 100 and the average number of servers per job increased by a factor of 3. In the last two years the increases were factors of 1.6 and 1.2, respectively [Dean 2009]. Figure 6.16 on page 385 estimates that running the 2009 workload on Amazon's Cloud Computing Service EC2 would cost $133M.

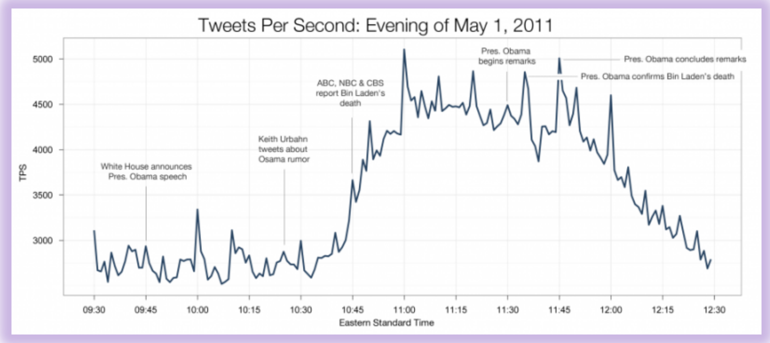

WSC design considerations: Costs, costs, costs…

- High volume needs low costs
  - “multiplier effect”
- Business models based on low costs
  - Key competitive advantage
- Both capital (capex) & operational (opex) costs
  - Power and cooling important component (more later)

Amazon Elastic Compute Cloud (EC2)

Prices not up-to-date

- $0.24 per hour for
  - 2 cores of 2.4 GHz 64-bit 2007 Opteron
  - 15 GB memory
  - 1.7 TB scratch storage
- Need it 24x7 for a year?
  - $1167

- $0.06 per hour for
  - 1 core of 1.2 GHz 32-bit or 64-bit Intel or AMD
  - 1.7 GB memory
  - 160 GB scratch storage
- Need it 24x7 for a year?
  - $292

- This includes
  - Purchase + replacement
  - Housing
  - Power
  - Operation
  - Reliability
  - Security
  - Instantaneous expansion and contraction
- 1000 processors for 1 day costs the same as 1 processor for 1000 days!

Infrastructure TCO

Infrastructure costs range into the millions

- Often exceed energy costs over facility lifetime

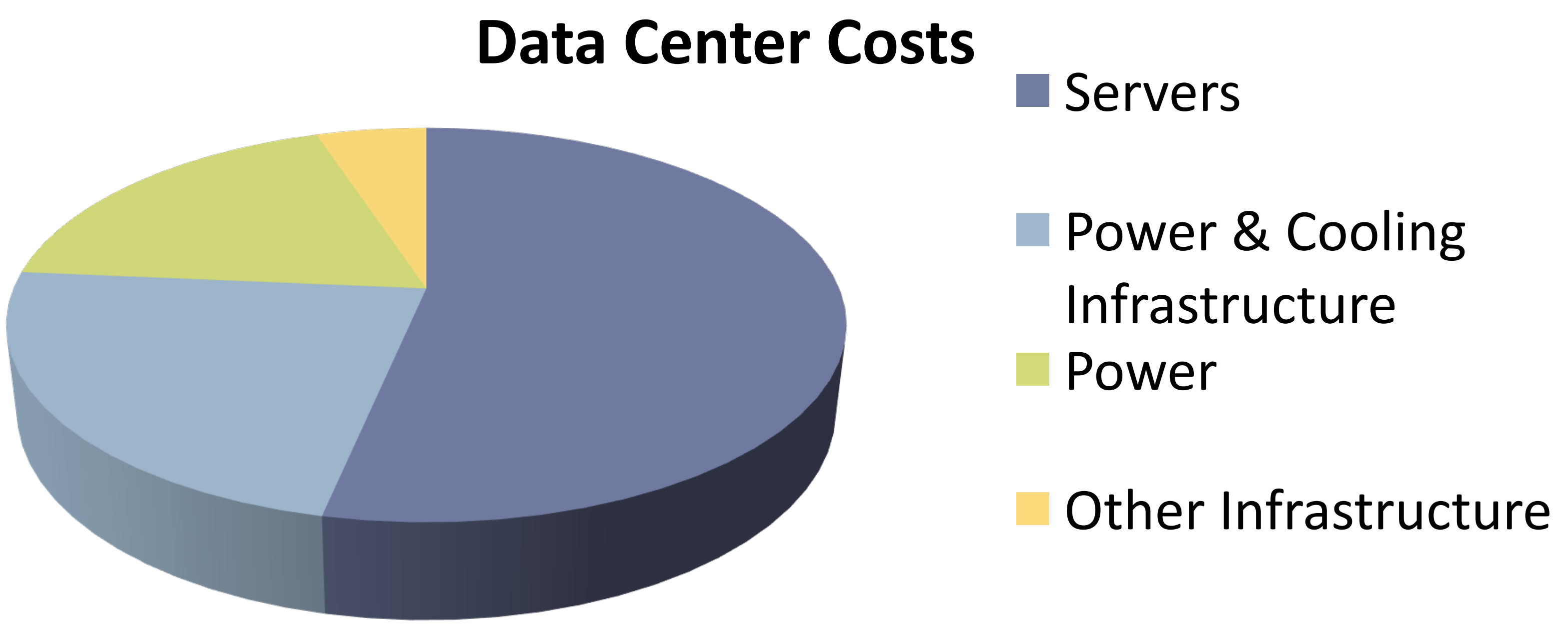

Source: J. Hamilton 2009

How do we reduce power infrastructure costs?

Power: A first-class data center constraint

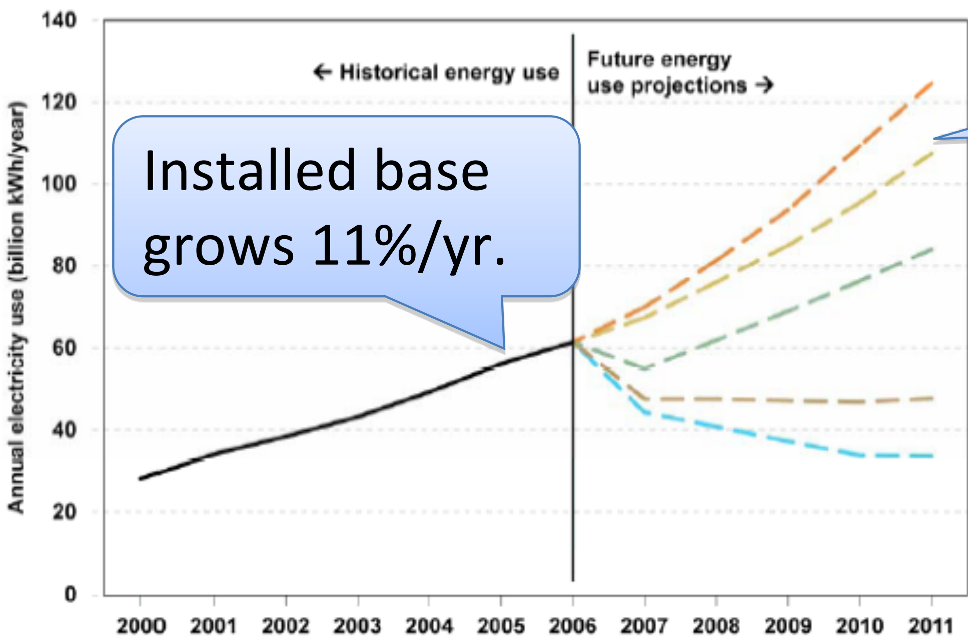

Installed base grows 11%/yr.

Source: US EPA 2007

2.5% of US energy
$7.4 billion/yr.

Annual data center CO_2:
17 million households

Source: Mankoff et al, IEEE Computer 2008

Lifetime Cost of a Data Center

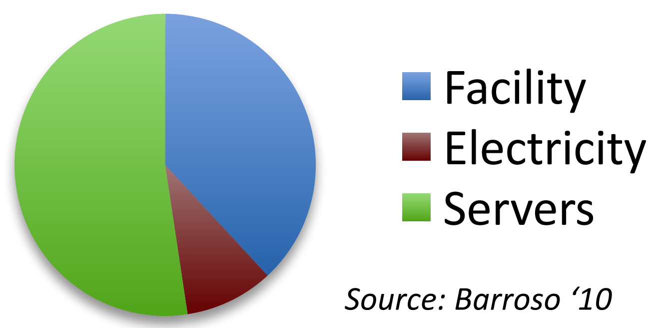

Source: Barroso '10

Peak power determines data center facilities capital costs

Improving energy & capital efficiency is a critical challenge

Cost of facility: ~$11/W

Energy in the data center

- In many data centers, only 13% energy goes to useful work
- 42% data center TCO on energy & infrastructure [Hamilton]

What is going on?

1. Energy waste in underutilized systems
   – Up to 80% idleness; but idle power ~60% of peak

2. Energy waste in physical infrastructure
   – Power/cooling use up to 40% of total energy

3. Capital waste in infrastructure over-provisioning
   – Poor power balance leads to >33% over-provisioning

The need for energy-proportionality

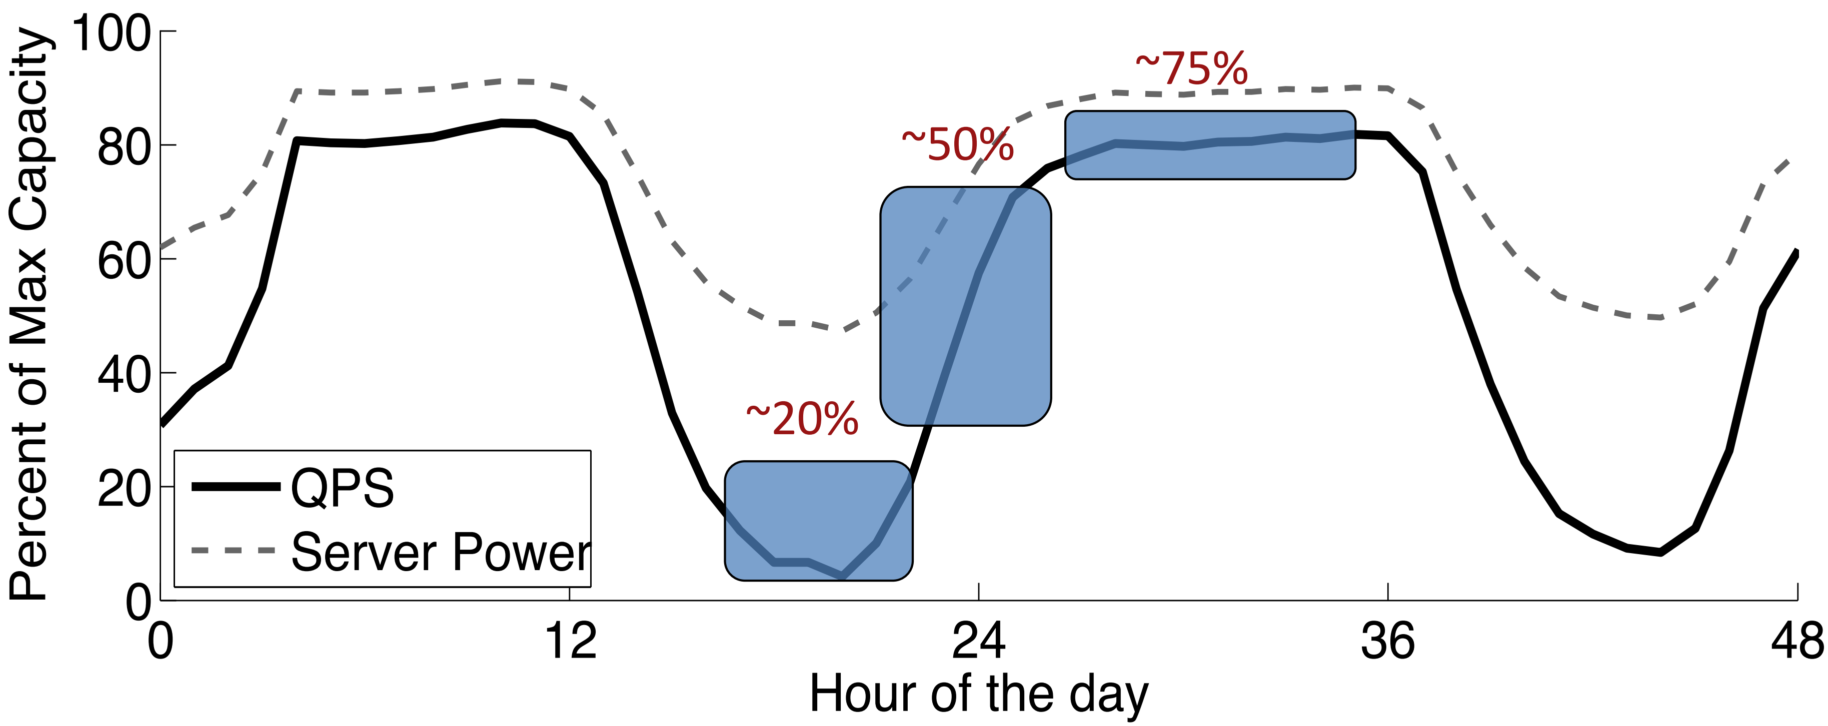

How to achieve energy-proportionality at each QPS level?

Typical server utilization

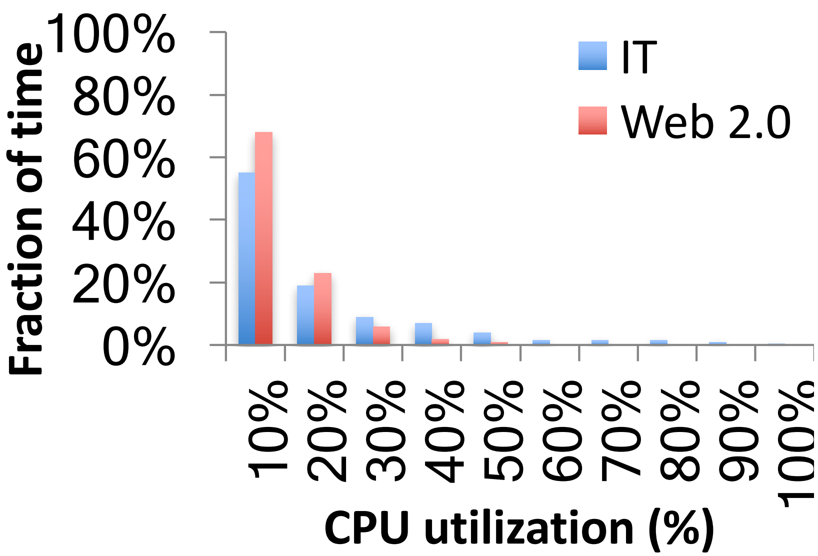

Customer traces supplied by HP Labs

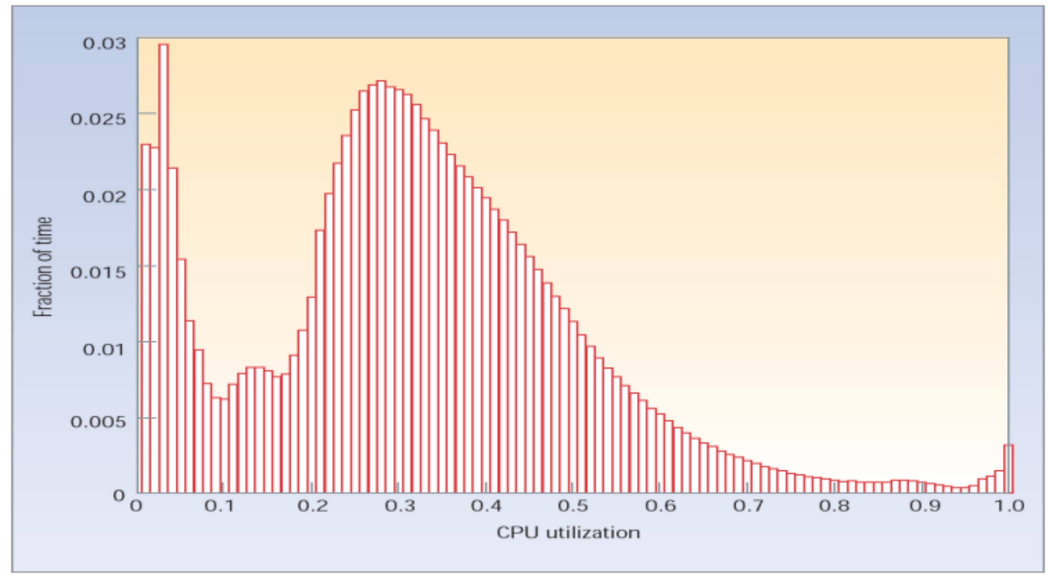

Source: Barroso & Hölzle, Google 2007

Why? Oh why? ☺

Typical server utilization

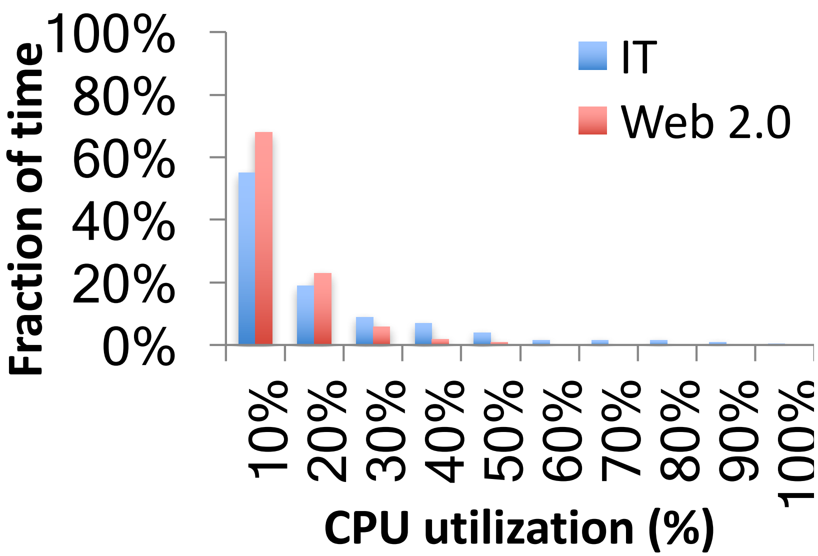

Customer traces supplied by HP Labs

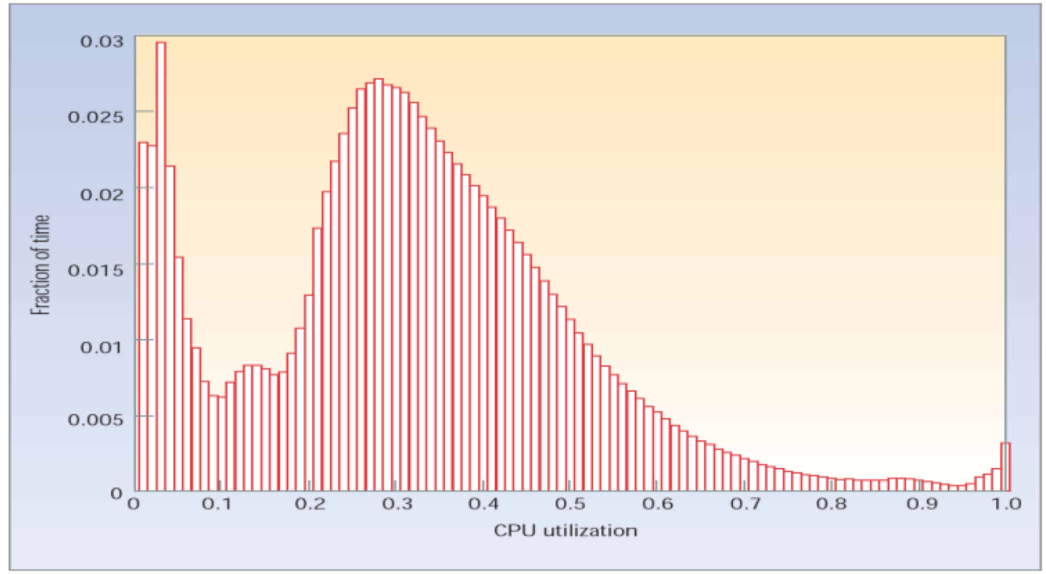

Source: Barroso & Hölzle, Google 2007

Low utilization (≤20%) is endemic

- Provisioning for peak load
- Performance isolation
- Redundancy

Historically, vendors optimize & report peak power

Idle periods are short

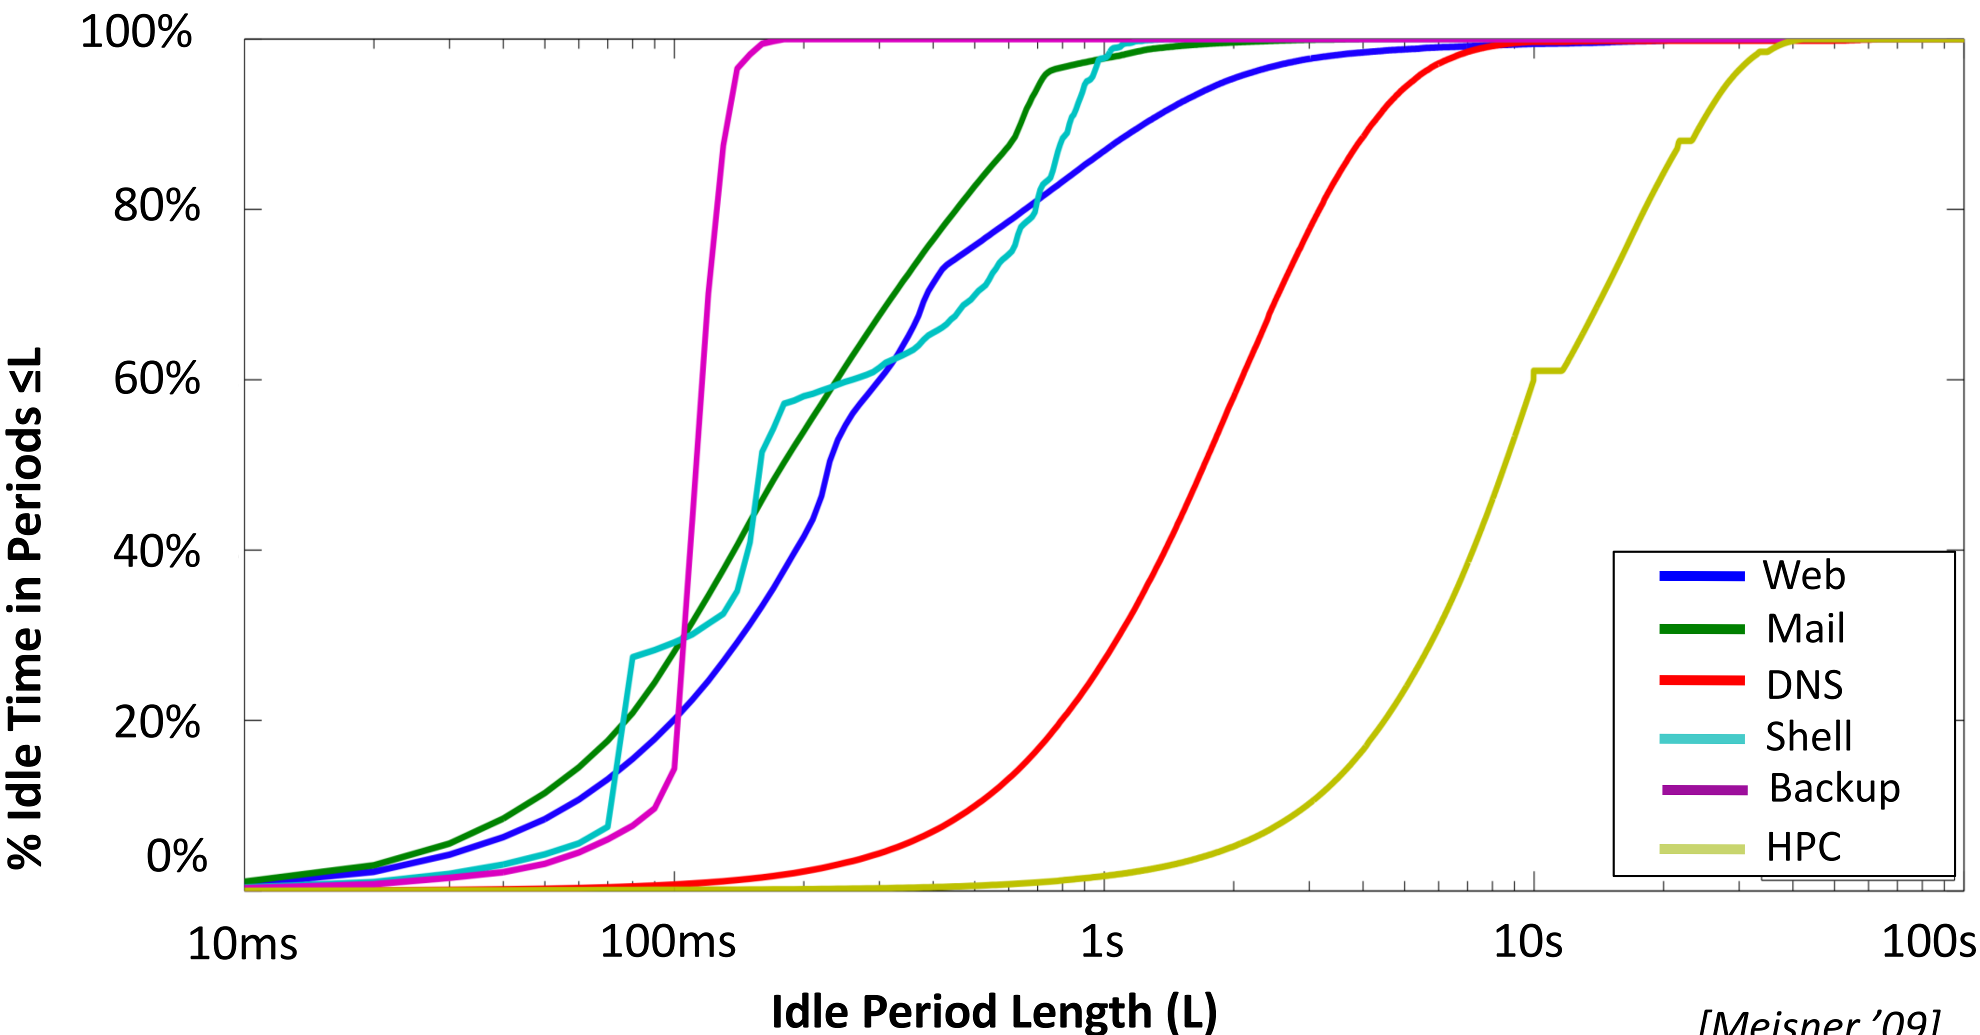

[Meisner '09]

*Most idle periods are **< 1 Sec** in length*

Server power breakdown

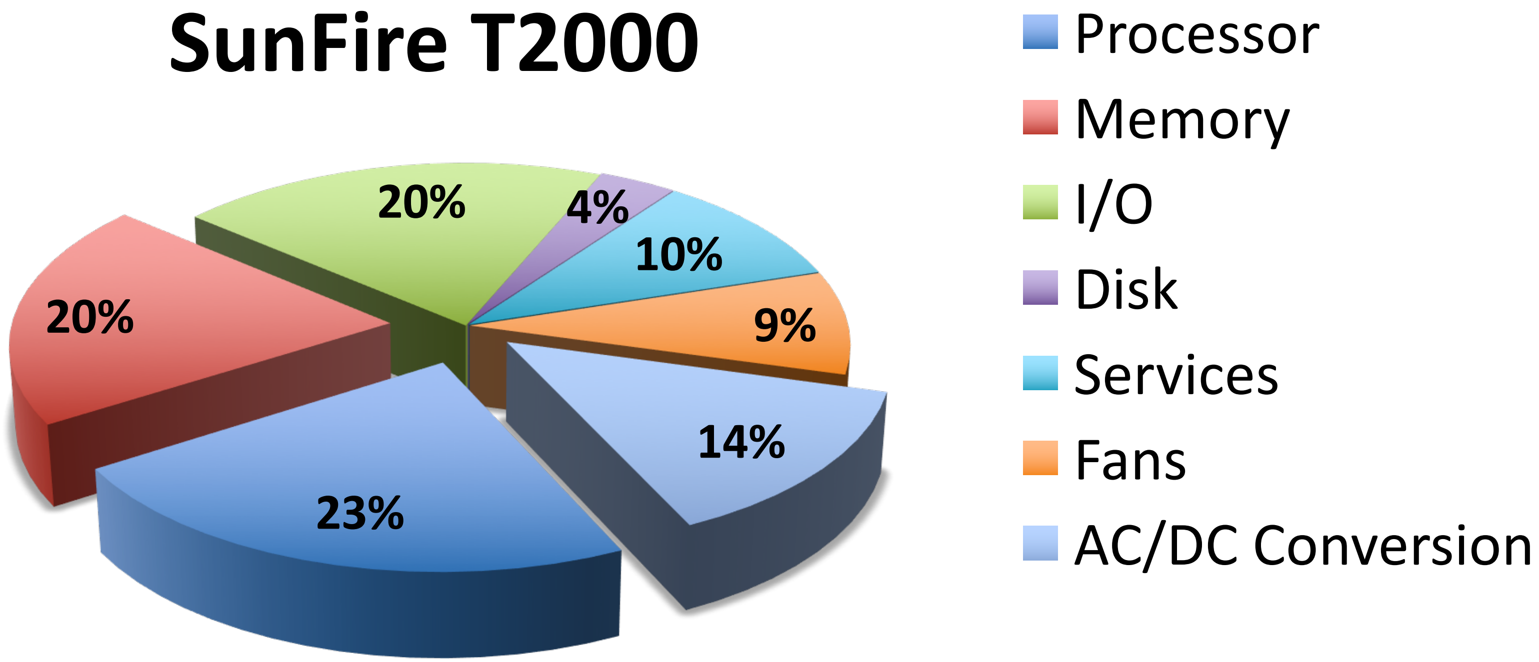

No single component dominates power consumption

Implications?

Reliability

Reliability: measure of continuous "service"

- MTTF: mean time to failure
  - Time to produce first incorrect output
- MTTR: mean time to repair
  - Time to detect and repair a failure
- MTBF = mean time between failures = MTTF+MTTR
- Failure in time: FIT = Failures per billion hours of operation = 10^9/MTTF
  - E.g., MTTF = 1,000,000 hours = 1000 FIT

- Definition of "system operating properly": sometimes not easy
  - delivered per service-level agreement (SLA)/SLO

Availability

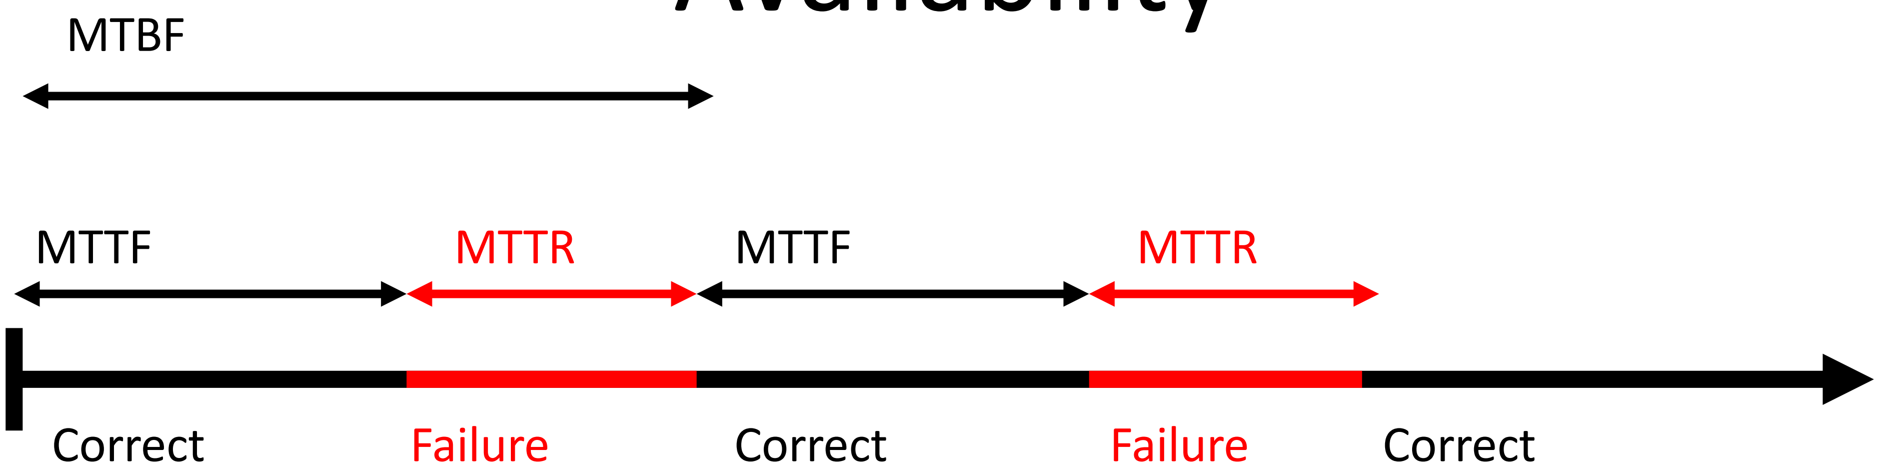

Steady state availability

$$= MTTF / (MTTF + MTTR)$$

Example

- Typical first year for a new cluster:

- ~0.5 overheating (power down most machines in <5 mins, ~1-2 days to recover)
- ~1 PDU failure (~500-1000 machines suddenly disappear, ~6 hours to come back)
- ~1 rack-move (plenty of warning, ~500-1000 machines powered down, ~6 hours)
- ~1 network rewiring (rolling ~5% of machines down over 2-day span)
- ~20 rack failures (40-80 machines instantly disappear, 1-6 hours to get back)
- ~5 racks go wonky (40-80 machines see 50% packetloss)
- ~8 network maintenances (4 might cause ~30-minute random connectivity losses)
- ~12 router reloads (takes out DNS and external vips for a couple minutes)
- ~3 router failures (have to immediately pull traffic for an hour)
- ~dozens of minor 30-second blips for dns
- ~1000 individual machine failures
- ~thousands of hard drive failures

- slow disks, bad memory, misconfigured machines, flaky machines, etc.

Google™

Why is availability important?

Application	Cost of downtime per hour	Annual losses with downtime of 1% (87.6 hrs/yr)	0.5% (43.8 hrs/yr)	0.1% (8.8 hrs/yr)
Brokerage operations	$6,450,000	$565,000,000	$283,000,000	$56,500,000
Credit card authorization	$2,600,000	$228,000,000	$114,000,000	$22,800,000
Package shipping services	$150,000	$13,000,000	$6,600,000	$1,300,000
Home shopping channel	$113,000	$9,900,000	$4,900,000	$1,000,000
Catalog sales center	$90,000	$7,900,000	$3,900,000	$800,000
Airline reservation center	$89,000	$7,900,000	$3,900,000	$800,000
Cellular service activation	$41,000	$3,600,000	$1,800,000	$400,000
Online network fees	$25,000	$2,200,000	$1,100,000	$200,000
ATM service fees	$14,000	$1,200,000	$600,000	$100,000

Figure 1.3 Costs rounded to nearest $100,000 of an unavailable system is shown by analyzing the cost of downtime (in terms of immediately lost revenue), assuming three different levels of availability, and that downtime is distributed uniformly. These data are from Kembel [2000] and were collected and analyzed by Contingency Planning Research.

WSC Networking

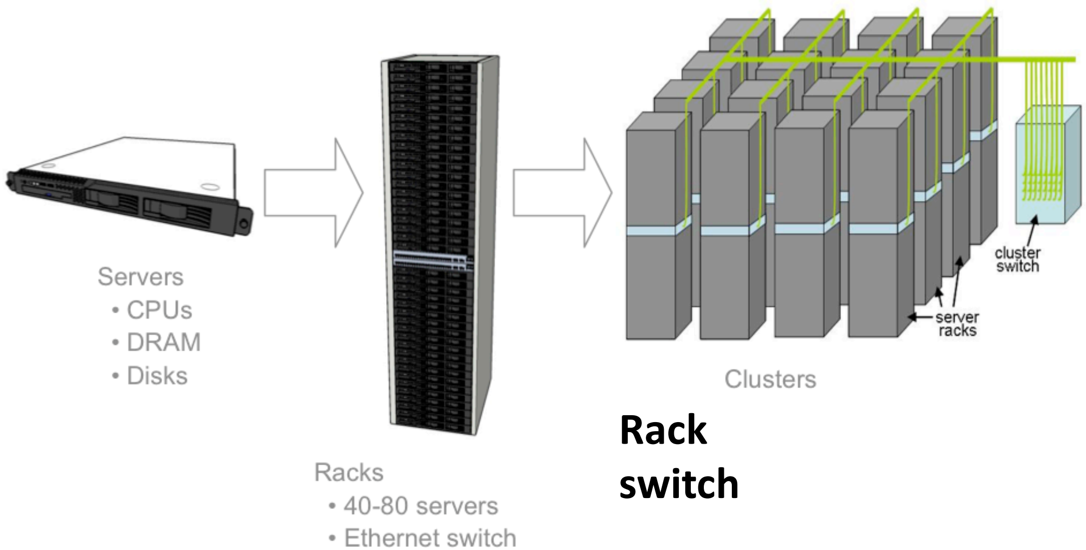

Array switch

Rack switch

- Connecting 500000 servers challenging
  - High bandwidth at low costs
- Hierarchy of network
  - Rack switch, array switch, L3 switch, border routers

Example: Google circa 2007

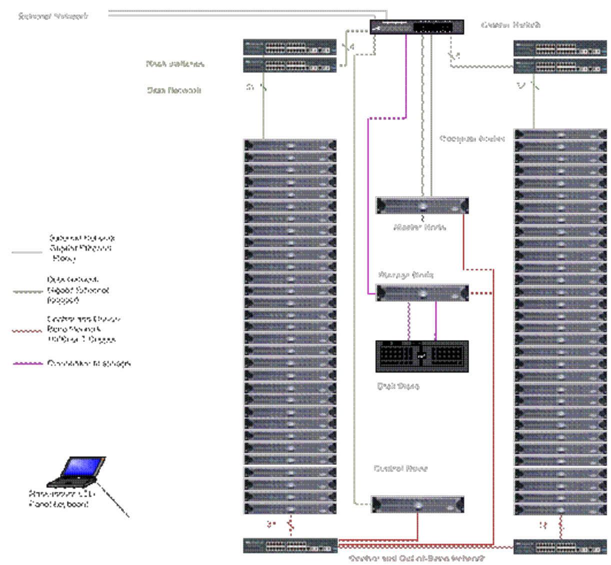

- Rack switch: 48-port ethernet 1Gig switch
  - Commodity switch: $30 per port
    - Infiniband: $500/port
  - One Switch per two racks
  - 40 server ports; 2-8 uplink ports
    - Oversubscription ratio
    - Programmer burden
  - Bandwidth within rack same irrespective of sender/receiver
- Array switch
  - More expensive: 10X more BW = 100X more $
  - High-end switches, feature-rich
  - 480 1Gbit links, few 10Gbit ports to datacenter routers
  - Manage oversubscription carefully
- Layer 3 switches, border routers

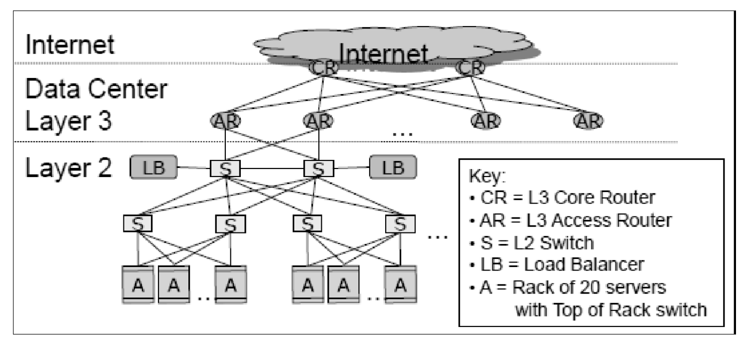

WSC Storage

- Storage
  - Distributed FS using disks on servers
    - Better fault-tolerance across nodes, lower cost, better scalability
  - Network attached storage (NAS) devices
    - Specialized systems with disk arrays that provide FS storage services and connect directly to the networking fabric
    - Better fault tolerance within device (e.g., RAID), easier management
    - More expensive
- Google circa 2007: Google file system (GFS)
  - Use local disks; local access patterns
  - At least three replicas for disk reliability
    - replicas used for several failure modes
  - Eventual consistency for lower cost

WSC Storage Hierarchy: A Programmer's Perspective

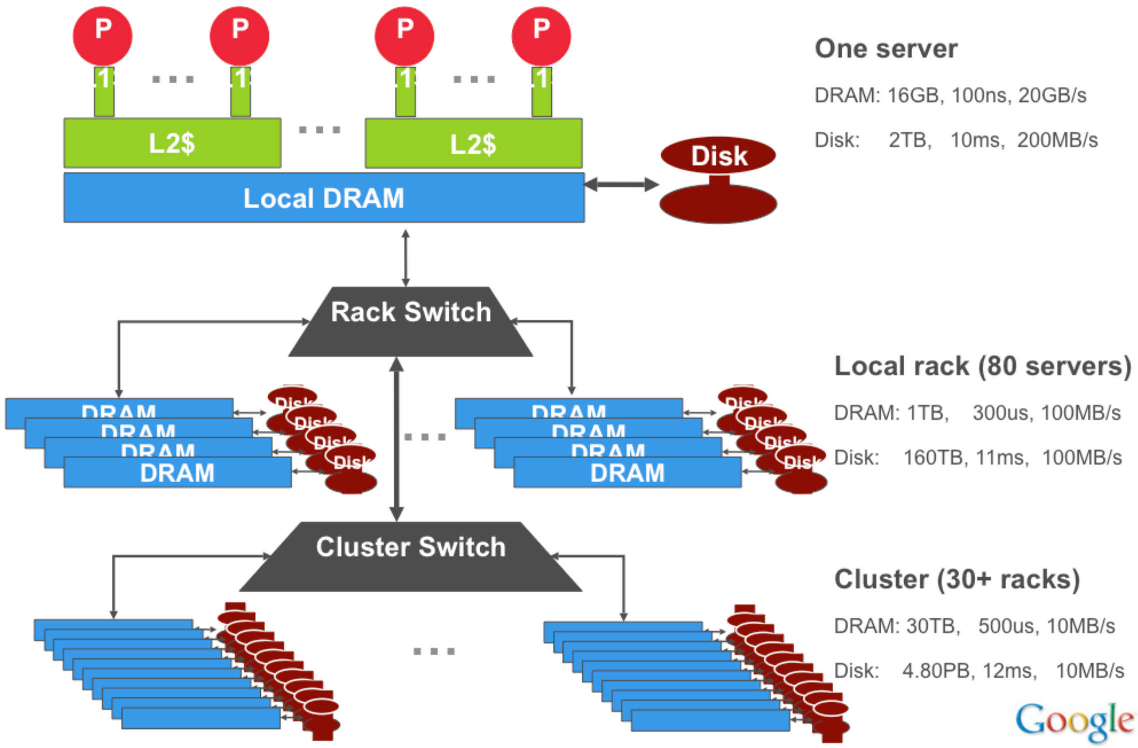

- Interesting observations
  - Remote memory is often faster than local disk
  - Bandwidth bottlenecks

	Local	Rack	Array
DRAM Latency (microseconds)	0.1	100	300
Disk Latency (microseconds)	10,000	11,000	12,000
DRAM Bandwidth (MB/sec)	20,000	100	10
Disk Bandwidth (MB/sec)	200	100	10
DRAM Capacity (GB)	16	1,040	31,200
Disk Capacity (GB)	2,000	160,000	4,800,000

Useful Numbers

Courtesy of Jeff Dean, Google

- L1 cache reference 0.5 ns
- Branch mispredict 5 ns
- L2 cache reference 7 ns
- Mutex lock/unlock 25 ns
- Main memory reference 100 ns
- Compress 1K bytes with Zippy 3,000 ns
- Send 2K bytes over 1 Gbps network 20,000 ns
- Read 1 MB sequentially from memory 250,000 ns
- Round trip within same datacenter 500,000 ns
- Disk seek 10,000,000 ns
- Read 1 MB sequentially from disk 20,000,000 ns
- Send packet CA->Europe->CA 150,000,000 ns

Network bandwidth trends

Growing fast! 1TB/s in 2015?

Injection rates might get better too.

Can move data around **fast**.
(BW rivaling memory)

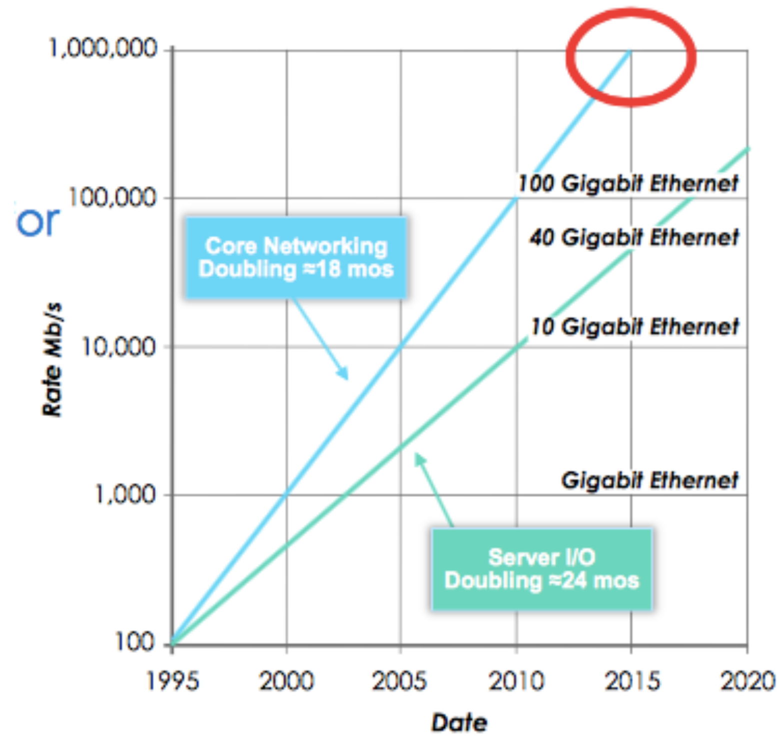

How will things change how we build DC applications?

Acknowledgements

- Thanks to Tom Wenisch (UMich), Ed Lazowska, Hank Levy, John Zahorjan and Steve Gribble for graciously sharing his slides.