



Education

# **Green Storage I: Economics, Environment, Energy, and Engineering**

SW Worth, Microsoft

- The material contained in this tutorial is copyrighted by the SNIA.
  - Member companies and individuals may use this material in presentations and literature under the following conditions:
    - ◆ Any slide or slides used must be reproduced without modification
    - ◆ The SNIA must be acknowledged as source of any material used in the body of any document containing material from these presentations.
  - This presentation is a project of the SNIA Education Committee.
  - Neither the Author nor the Presenter is an attorney and nothing in this presentation is intended to be nor should be construed as legal advice or opinion. If you need legal advice or legal opinion please contact an attorney.
  - The information presented herein represents the Author's personal opinion and current understanding of the issues involved. The Author, the Presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.
- NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.**

# Abstract for “Green Storage I”

Even non-geeks are becoming aware of the environmental impacts (especially energy costs) associated with data storage. Discussions of “Green Storage” require understanding of fundamental concepts common to all components of a datacenter. This SNIA Tutorial covers the concepts of Economics, Environment, Energy, and Engineering that are necessary to participate in the dialogue, whether you are a manager or a hands-on I.T. professional. Wherever possible, the examples used refer to Storage, but detailed focus on Storage components and technologies is left for Part II.

We start with definitions of “Green” used by various groups, covering various motivations for making “Green” decisions for your organization. This requires some economic theory, but you will be rewarded with a new-found ability to explain “Cap-and-Trade” management of Carbon and “SOx” (Sulfur Oxides, not Sarbanes-Oxley!) at cocktail parties. This leads naturally to coverage of various environmental regulations and initiatives (e.g. ROHS, WEEE, Energy Star) that affect manufacturers and end-users of storage components or computers. We will quickly review basic engineering topics relevant to understanding 'Green', including stuff you may have successfully avoided, such as environmental chemistry, thermodynamics, energy vs. power, and heat transfer. (Examples include conversion losses, AC and DC power choices, and power supply efficiency.) Since much of the focus in I.T. is on energy costs, we'll spend some time understanding energy supply and pricing (focused on the U.S. market), and provide some negotiating alternatives for dealing with your friendly energy suppliers. We will review current problems in data center design, including increasing computational and storage density and the resulting energy and cooling issues. All of this will come together to help guide your design process towards a better allocation of capital expenses (CapEx) and operational expenses (OpEx) to lower your TCO.

**Bottom-line:** After this session you will be armed with the knowledge you need to be part of the Green decision-making process for your datacenter, so those pesky server-geeks don't dominate the discussion!

## Learning Objectives

- Understand what various groups mean by 'Green', especially as this term relates to storage network components and systems.
- Understand factors that should (or will) motivate your interest in 'Green' storage, including regulation, competition, and TCO.
- Understand why 'OpEx' (Operating expenses) can outweigh 'CapEx' (Capital expense) in TCO models, and how 'Green' factors increasingly influence OpEx.

# Green Storage I - Agenda

- **Overview, Motivation, and Definitions**
- eWaste Reduction/Recycling: RoHS, WEEE, etc.
- Fundamentals of Energy and Cooling
- Electricity Pricing in the United States
- Datacenter Design and Operation
- Storage Components and Technologies

# Overview and Definitions

- “Green”  $\sim$  = Reduction of (Waste + Energy Use)
  - ◆ Gov’t. (EPA, EU, Kyoto), Green orgs, Vendors, etc.
- Classic motivators: Fear, Guilt, Greed
  - ◆ Constraints, Competition, and ProfitMax/CostMin
  - ◆ Strategy ( $\sim$  = CapEx) vs. Tactics ( $\sim$  = OpEx)
- How does “Green” differ from normal economic considerations, e.g. efficiency, optimization?
  - ◆ Rationalize decisions by including “**externalities**” Widen scope of action across org boundaries, time
- “Green” effects on Storage decisions

- **“Fear”**: Constraints (Regulations, Physical limits)
  - ◆ Do what you are forced to do by Gov. Regulations
  - ◆ Physical Laws are not optional...
    - › E.g. WAN latency (light-speed), Disk rotation speed, Tape Retrieval
- **“Guilt”**: Competitive and ‘Moral’ aspects
  - ◆ Keeping up with industry, responding to non-economics
    - › TBL (“Triple Bottom Line”); or “Social, Economic, Environmental”
- **“Greed”**: Profit Maximization / Cost Minimization
  - ◆ Strategy → Capital Expenses (CapEx)
  - ◆ Tactics → Operational Expenses (OpEx)
  - ◆ TCO (Total Cost of Ownership) integrates CapEx and OpEx

# Externalities

- Problem: important inputs or outputs (e.g. Greenhouse gases, ‘Carbon’) have unclear prices or owners
  - ◆ Some factors are effectively Zero-cost to the decision-maker, but are not cost-free to larger group affected
  - ◆ This leads to non-optimal decisions and behavior
    - > ‘Tragedy of the commons’
    - > Classic solutions: Government **mandates** (Regulation)
    - > Separate accounting system, e.g. for Carbon “Footprint”
    - > Unintended Consequences
- Pigouvian taxes: “Sin Taxes” (modify behavior)
- Cap-and-Trade SO<sub>x</sub>/NO<sub>x</sub>, Carbon (modify outcomes)
  - ◆ Coase’s Theorem: Property Rights, Negotiation

# “Green” effects on Storage

- “TCO” (Total Cost of Ownership) now combines with Externalities to affect purchase decisions
  - ◆ In most cases Externalities will evolve to provide clear pricing signals (e.g. RoHS, WEEE, Cap-and-Trade)
- Expand scope of decision-criteria and constraints to include (at least) entire datacenter
  - ◆ Servers, Networking, and Storage
  - ◆ Power, energy, and cooling (CapEx and OpEx)
  - ◆ People: widen their decision-boundaries, -constraints
    - › Include your Facilities managers
- Unintended Consequences: reduced reliability?

# Three Stages of Product Life-Cycle

- Birth: Product Creation (design for recycle/disposal)
  - ◆ Integrated into CapEx (*probably*) – WEEE/RoHS
  - ◆ Facilities/Infrastructure (proportional to **POWER**)
- Useful-Life: Energy, Cooling, and “Other”  
Environmental Impacts during Productive Life
  - ◆ Storage: dominated by Energy/Cooling (Electricity)
    - › Few consumable supplies, except Tape cartridges
  - ◆ Dominated by OpEx (but is this visible to IT?)
- End-stage: Removal, Recycling, Disposal
  - ◆ Integrated into initial CapEx or OpEx surcharges
  - ◆ *Alternative: dump these costs onto everyone else....*

# Energy/Cooling Cost in Datacenters

- How much is due to Storage? (*Proportion is increasing!*)
  - ◆ It depends on Design and Workload (I/O profiles)!
  - ◆ Published studies range from <10% - >40%
  - ◆ “Rule-of-Thumb” for energy: 60% servers, 20% networking, and 20% **Storage** (but no consistent definition of ‘Storage’)
- Peak loads required for design (~Max. Power)
  - ◆ CapEx: Power (UPS, PDU)/Cooling (Fans, CRAC), Installation
    - › Can overwhelm space charges, even Energy Costs
    - › ➔ Watch for “Demand” charges (e.g. Peak 15-min. of 3 Months!)
- Time-weighted I/O for Energy/Cooling ~ = OpEx
- TCO = CapEx + OpEx, but which dominates?

# Storage gets more visibility?

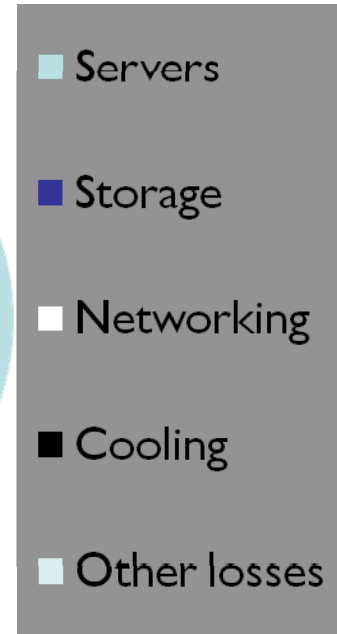
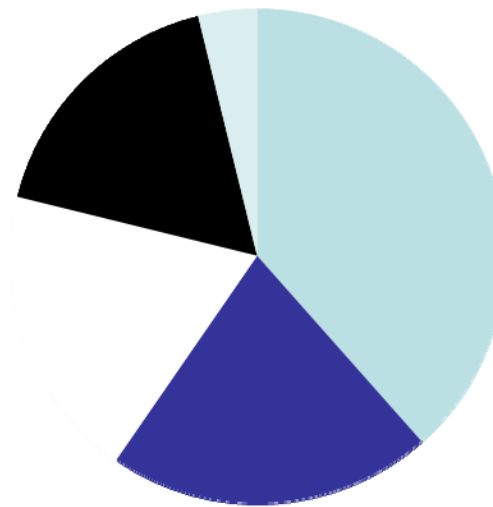
$$\text{DCiE} = (\text{IT Energy Demand}) / (\text{Total Datacenter Energy Usage})$$

(DCiE currently ~50%)

**Note: Percentages are for illustration only.  
(So don't quote them!)**

Facilities Efficiencies will keep improving,  
so these 'slices' get smaller.

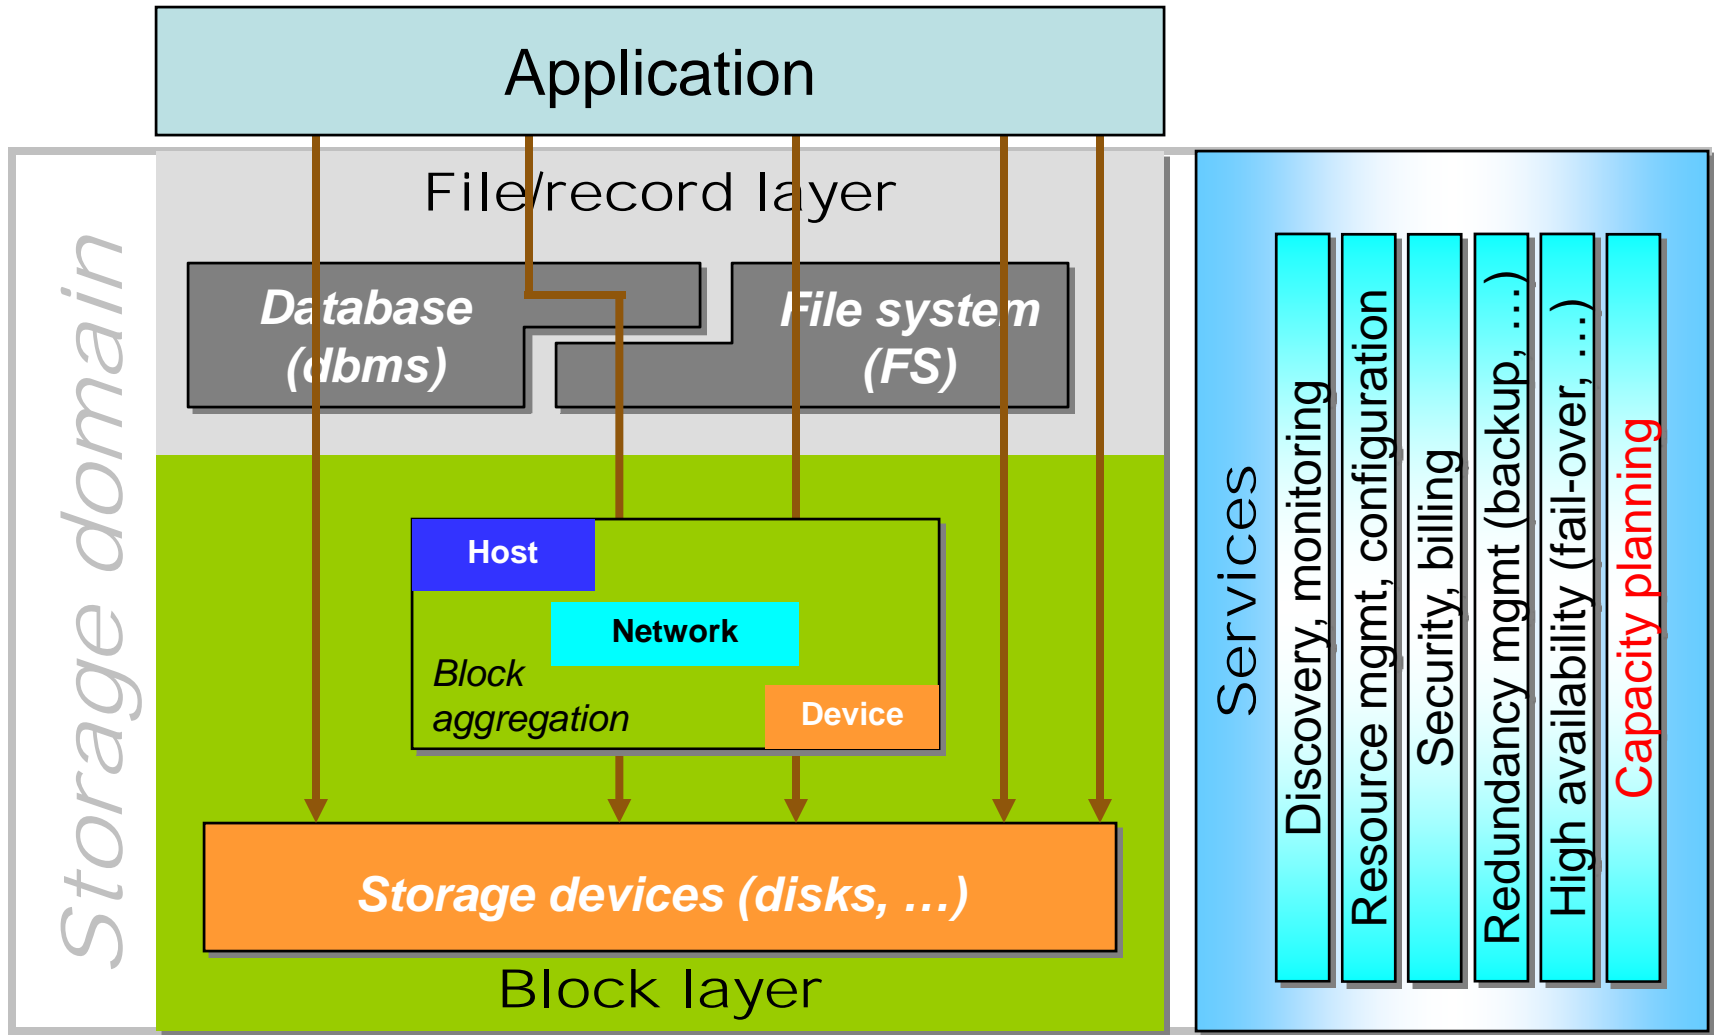
**(DCiE approaches 80%)**



**Server percentage also declines, due to Consolidation and Virtualization.**

# What Storage aspects could be affected?

## SNIA Shared Storage Model (and don't forget Tape!)



# Green Storage - Agenda

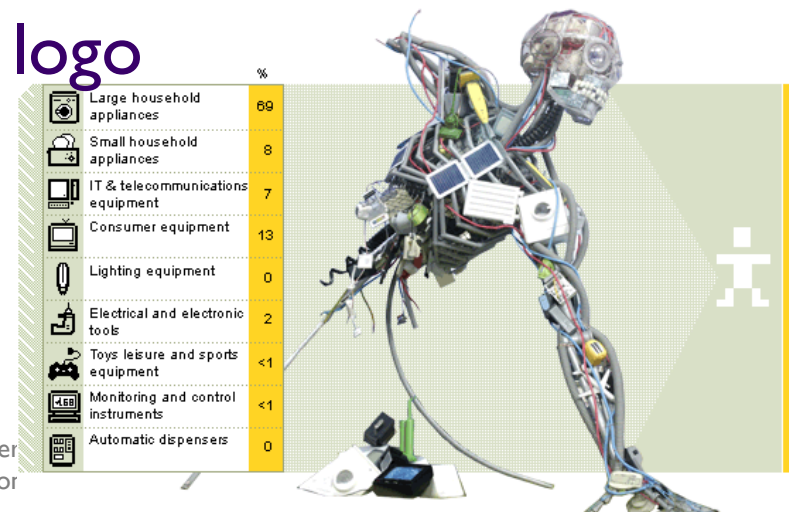
- Overview, Motivation, and Definitions
- **eWaste Reduction/Recycling: RoHS, WEEE, etc.**
- Fundamentals of Energy and Cooling
- Electricity Pricing in the United States
- Datacenter Design and Operation
- Storage Components and Technologies

# eWaste Reduction/Recycling

- Government regulations (“Directives”) that may affect storage vendors (and their customers).
  - ◆ Useful site for US businesses:  
[www.buyusa.gov/europeanunion/commerce\\_docs.html](http://www.buyusa.gov/europeanunion/commerce_docs.html)
- WEEE
- RoHS, China-RoHS
- Packaging and Pkg Waste
- Halogens (in plastics)
- Basel Convention/Basel Ban (Transboundary Wastes)

## Waste Electrical and Electronic Equipment

- European Community directive 2002/96/EC
  - ◆ Conformance from Aug-05
- Increase reuse, recycling, recovery
- Reduce landfill and incineration
- Financed by manufacturers and vendors
  - ◆ Users can return WEEE without charge
  - ◆ “Take It Back” programs
- Look for the “Wheelie-Bin” logo
  - ◆ Recycle, don’t dispose!



# RoHS: Restriction of Hazardous Substances

- European Directive 2002/95/EC, effective Aug-06
- RoHS restricts the use of certain hazardous substances in various types of new electronic and electrical equipment. (*Note: at a component level!*)
  - ◆ Mercury                    - Cadmium                    - PBB
  - ◆ Chromium VI           - Lead                               - PBDE
- Unintended Consequences: reduced reliability?
  - ◆ EPA report (Aug-05) on lead-free solder!
  - ◆ RoHS exemption: lead solder for Servers and Storage?
    - › Due to a clear trade-off on reliability and performance
    - › This exemption will go away with improved techniques



# “China-RoHS”

- Chinese Ministry of Information Industry Order #39 Management Methods for Controlling Pollution by Electronic Information Products, in effect on March 1, 2007.
  - ◆ [SJ/T 11363-2006 Requirements for Concentration Limits for Certain Hazardous Substances in Electronic Information Products](#)
- Similar restricted substances as RoHS
- Split timetable for labeling and conformance
- Different/Fewer(?) exemptions
- ➔ Ask an expert if you think you are affected!

# WEEE/RoHS – U.S. and Rest of World? **SNIA**

## ➤ United States

- ◆ Vendors have almost universally adopted RoHS since most do business in Europe
- ◆ EPA regulations and recommendations (e.g. Lead-free)
- ◆ Proposed federal legislation
- ◆ Several States have some regulations
  - › California – “Electronic Waste Recycling”
- ◆ Many vendors will “take it back” or take trade-ins

## ➤ Canada/Australia RoHS

## ➤ Asia (Japan JGPSSI), Korea/Taiwan RoHS

# Green Storage - Agenda

- Overview, Motivation, and Definitions
- eWaste Reduction/Recycling: RoHS, WEEE, etc.
- **Fundamentals of Energy and Cooling**
- Electricity Pricing in the United States
- Datacenter Design and Operation
- Storage Components and Technologies

# Energy and Cooling: Fundamentals

- Laws of Thermodynamics
- Power vs. Energy: Units of Measurement
- Heat Transfer
  - ◆ Conduction, Convection, Radiation
  - ◆ Data-center cooling: Air vs. Liquid
- Energy Conversion, Transmission, Storage
  - ◆ AC/DC and DC/AC conversion losses
  - ◆ Voltage step-down and step-up conversion losses
- Systems of Measurement: SI vs. US

# Laws of Thermodynamics

- First Law: Energy cannot be created or destroyed, it only changes form.
- Second Law: Entropy increases in a closed system.  
Efficiency of energy conversion is  $<100\%$ .
- Alternate Formulations:
  - ◆ **You can't win, you can't even break even, and you can't get out of the game....**
  - ◆ *“Nullium Prandium Gratium”* (or “TANSTAAFL”)
- NO: you cannot power your datacenter using the waste heat to generate electricity to run the site!
  - ◆ But you might increase DCiE with “free” cooling

# Heat Transfer

## ➤ Heat (Cooling):

### ◆ Conduction:

- › thermal glue/grease between CPU and cooling fins

### ◆ Convection

- › Cooling fluid circulated past hot components
- › Note: “fluid” could be air or liquid, but liquid has a lot more capacity to move heat

### ◆ Radiation

## ➤ Phase Change: Solid-Liquid; Liquid-Gas

## ➤ Newton’s Law of Cooling

- ◆ Rate varies with Temperature Difference

# Green Storage - Agenda

- Overview, Motivation, and Definitions
- eWaste Reduction/Recycling: RoHS, WEEE, etc.
- Fundamentals of Energy and Cooling
- **Electricity Pricing in the United States**
- Datacenter Design and Operation
- Storage Components and Technologies

# Electricity prices are ~100x variable (at least at the wholesale level)

- **Electricity cannot be stored effectively!**
  - ◆ Few exceptions: Pumped Water Storage, Compressed Air
  - ◆ Batteries, Flywheels, etc. are short-duration, costly
- **Prices vary with DEMAND (local and regional)**
  - ◆ Weather (Hot, Cold, or Both), Supply disruptions
  - ◆ Time-dependent: Daily, Weekly, Seasonally
  - ◆ Economic conditions – general, regional
- **Prices vary with SUPPLY (local and regional)**
  - ◆ CapEx: plant construction (NIMBY), maintenance
  - ◆ OpEx: Fuel costs dominate – swings are wild ( $10^2$ )
- **Electric Transmission congestion/losses increase cost; hard to build new lines (NIMBY)**

# State Electricity Prices, 2005 (cents/kWh – “Industrial”)

U.S. Average 5.73

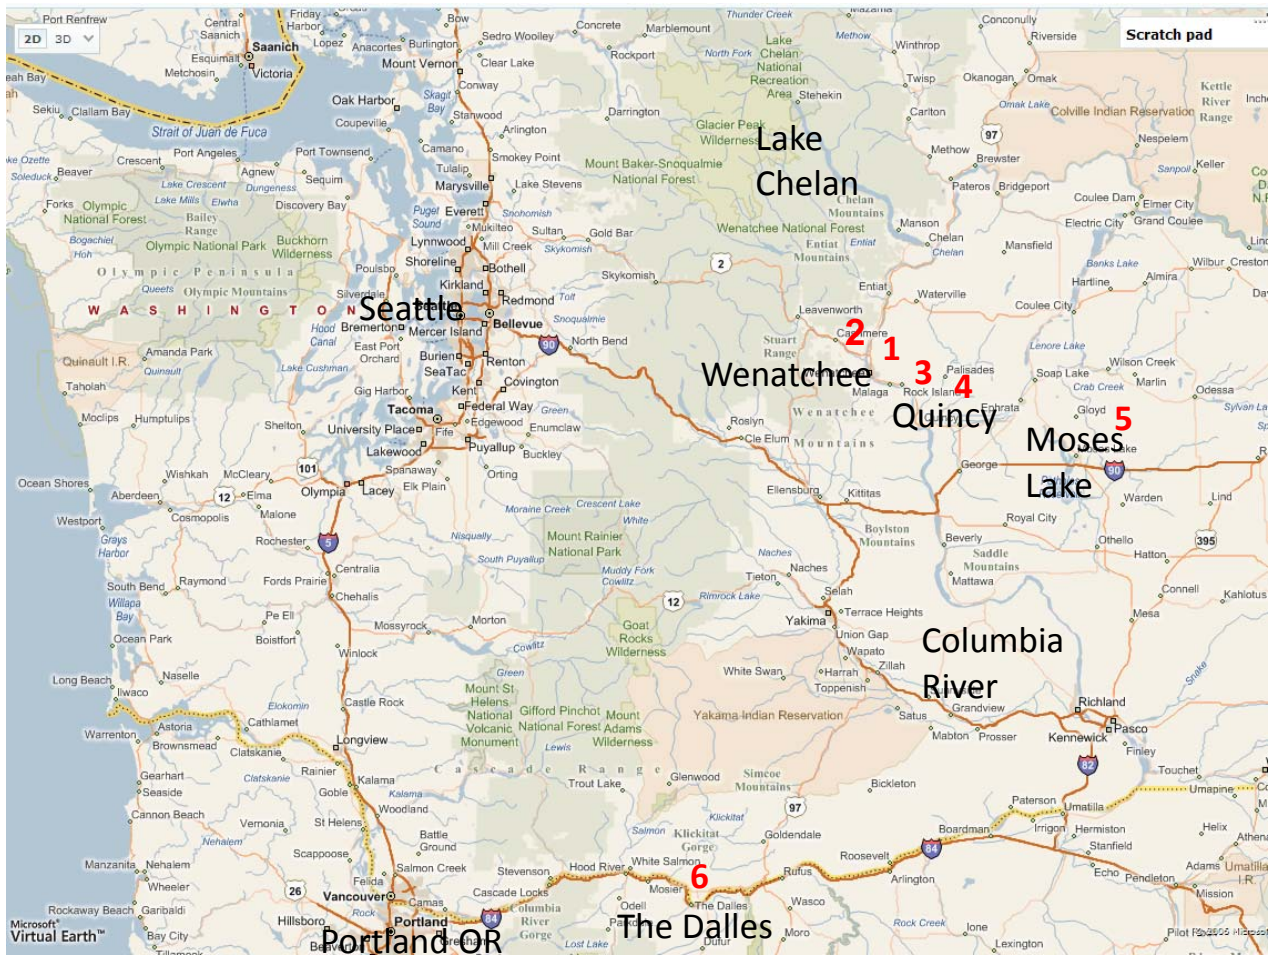
## 10 Most Expensive States

## 10 Least Expensive States

Rank	State	Price
1	HI	15.79
2	DC	14.13
3	NH	11.48
4	RI	10.01
5	NJ	9.76
6	CA	9.55
7	CT	9.40
8	AK	9.29
9	MA	9.22
10	NY	8.23

Rank	State	Price
42	VA	4.46
43	NE	4.43
44	IN	4.42
45	ND	4.32
46	WA	4.27
47	UT	4.24
48	WY	3.99
49	ID	3.91
50	WV	3.85
51	KY	3.60

# Move datacenter to cheap energy?



Energy costs on the Columbia River are about **\$0.02/kWh** for Datacenters.

Ample fiber (WAN) bandwidth is available ([www.noanet.net](http://www.noanet.net))

The area is also seismically inactive and in a 500-year flood zone.

**Result: Construction!**

- [www.eia.doe.gov/cneaf/electricity/st\\_profiles/florida.html](http://www.eia.doe.gov/cneaf/electricity/st_profiles/florida.html)
- FRCC - Florida Reliability Coordinating Council
  - ◆ Oversight of bulk electricity supply <https://www.frcc.com>
- Florida's electricity generation: mostly natural gas
  - ◆ FPL: 25 GW (52% gas, 19% nuclear, 15% purchased, 8% oil, 6% coal)
- Proposed RPS (Renewable Portfolio Standard)
  - ◆ >20% of electricity from 'renewable' sources by 2020 (1% now)
- Possibilities([http://www.ieses.fsu.edu/docs/fesc\\_%206\\_thrusts\\_13\\_tasks.pdf](http://www.ieses.fsu.edu/docs/fesc_%206_thrusts_13_tasks.pdf))
  - ◆ Florida has ~10% of U.S. biomass resources
  - ◆ Solar available: ~6 kWh/sq. Meter/day
  - ◆ Carbon Capture and Sequestration; more Nuclear?

# Green Storage - Agenda

- Overview, Motivation, and Definitions
- eWaste Reduction/Recycling: RoHS, WEEE, etc.
- Fundamentals of Energy and Cooling
- Electricity Pricing in the United States
- **Datacenter Design and Operation**
- Storage Components and Technologies

# Datacenter: Design/Operation

- Datacenter: Design and Operation
  - ◆ CapEx and OpEx (IT gear: Servers, Storage, Networking, plus Infrastructure, e.g. UPS, PDU, CRAC, Fans)
  - ◆ Multiplier effects on Power, Cooling, and Energy
  - ◆ Trends in Conservation and Optimization
- Size matters (for Power and Cooling equipment!)
  - ◆ Undersized infrastructure means less density for IT gear
    - › Modern IT gear is increasingly power-dense (>6 kW/rack)
      - May constrain current Storage equipment
    - › Chargeback ('rent') by **power**, vs. by rack-space may be required
    - › Some datacenters are limited by external Power availability
  - ◆ Oversized (IT gear plus infrastructure) = excess CapEx
    - › Under-utilization = Power inefficiencies (poor OpEx)

# Facilities vs. I.T. in the Datacenter

- Who represents I.T. to the Facilities staff?
  - ◆ Right now, the conversation is mostly about Servers!
  - ◆ Try to find “Storage” mentioned in any recent article on power/cooling problems in the datacenter....
  - ◆ Try to find “Storage” mentioned in any Utility program.
  - ◆ Can you show that Storage is significant to the power/cooling load (*via modeling or measuring*)?
- Organizational differences (who owns what?)
  - ◆ Do you talk with your Facilities managers?
  - ◆ Do your decisions affect each other? (YES!)
  - ◆ When will you start planning together?

# Datacenter Options: (Mech, Elec, Plumbing)

- Convert from AC to DC distribution
  - ◆ Can be partial conversion (DC arrays available)
- Run at higher voltage (240 or 480 vs. 120)
- Increase Power Supply efficiency (ask vendor)
  - ◆ 80 PLUS program ([www.80plus.org/servers.htm](http://www.80plus.org/servers.htm))
- Operate Cooling effectively
  - ◆ Leverage sensors, Follow basic rules (hot/cold aisles!)
  - ◆ Computational Fluid Dynamics (get some help!)
- Run Generator-testing for Peak-shaving
  - ◆ ➔ Negotiate with your power supplier for discounts!

# The Green Grid

- What is “The Green Grid”? [www.thegreengrid.org](http://www.thegreengrid.org)
- Green Grid metrics (*measured at the meter*)
  - ◆ What amount of Energy (and Cooling) goes to do “useful IT work”? (The rest is “overhead”, from an IT viewpoint)
  - ◆ Overall Datacenter (short-term, tactical)
    - › **PUE** (Power Usage Effectiveness)=(Total Facility Power/IT EquipPower)
    - › **DCiE** (Datacenter Infrastructure Efficiency): **DCiE=(1/PUE)**
  - ◆ Metric for “Datacenter Productivity” (longer-term, strategic)
    - › Datacenter Productivity = [Useful Work / Total Facility Power]
    - › Definition problem: what is “*useful work*” for IT?

# Model or Measure: Which is Better?

- Modeling: some info is required!
  - ◆ Accurate manufacturer data by Component and Product (Frame)
  - ◆ Stand-by Power vs. “Idle” vs. Full-load – CRUD analysis
  - ◆ Knowledge of I/O workload
    - Well-known benchmark(e.g. SPC, SNIA-IOTTA) – vary replay
    - YOUR unique workload traces (time-weighted and Peak)
- Measurement issues (Reality validates Modeling)
  - ◆ Actual *in-situ* workloads (“normal” and Peak) – can use traces
  - ◆ Actual Energy usage from Power Meter
    - Watts or kWh (what you pay for!), not Amps
    - Must be adequate to fit your Storage device (>30 Amp?)
    - See your Facilities Mgr, or a consultant for help
- SNIA Green Storage Technical Working Group projects

# Datacenter: Proposals and Solutions **SNIA**

- REDUCE Performance when possible
  - ◆ “Underclocking”: reducing performance-state of CPU reduces power/cooling needs for **Servers**
    - › Out-of-band mgmt (BMC) = no OS tuning
    - › Management via OS gives more granular control
- What is the equivalent for **Storage**?
  - ◆ TAPE or Optical? (trade-off response time vs. energy usage)
  - ◆ Disk drives and RAID arrays
    - › Slower/Larger drives where possible (Design choice vs. Dynamic)
    - › Power off selected drives: MAID (Massive Array of Idle Disks)
    - › Spin-down drives
  - ◆ Solid-State Storage

# Green Storage - Agenda

- Overview, Motivation, and Definitions
- eWaste Reduction/Recycling: RoHS, WEEE, etc.
- Fundamentals of Energy and Cooling
- Electricity Pricing in the United States
- Datacenter Design and Operation
- **Storage Components and Technologies**

# Storage-specific Power/Cooling data

- Each component of a Storage system has Power and Cooling requirements
  - ◆ Understand “Idle” (*\*not\** ‘stand-by’) vs. “Loaded” (R/W)
  - ◆ Label ratings are usually peak power required
    - › If you design using this data, your power/cooling equipment may be (grossly) over-built (Bad!), and CapEx will suffer.
    - › Operating equipment below its rated temperature offers little (no?) benefits (except for Operators!)
  - ◆ Some manufacturers offer better data or design info
  - ◆ If you really want to know, you may have to instrument in order to get real measurements.
  - ◆ Or, you could wait to see what SNIA comes out with...

# Disk-specific Power/Cooling

- Operational envelope
  - ◆ No clear effects on MTBF or TCO of variation within design temperature range
  - ◆ Can temp bounds be expanded?
- Rotational speed of Disks
  - ◆ Buy slower disks, if you don't mind the latency
  - ◆ Variable-speed disks?
- Use appropriate RAID levels
  - ◆ Disks may be 'free', but power/cooling are NOT!
- Max Disk Utilization (OpEx: per disk, not per GB)

# Key Strategies: Energy/Cooling

- Understand Usage vs. Demand and Other charges!
- Are you sure that Storage is a significant contributor?
- ➔ **Increase Utilization** (Storage Resource Mgmt helps)
  - ◆ **Thin Provisioning, Dynamic LUN Grow/Shrink**
- Consolidate (possibly change storage architecture)
- Trade Response Time (Latency+Throughput) for Reduced Power. i.e. Use Lower-tier Disk, VTL, Nearline, MAID, or Off-line Tape or Optical
- **Move:** when energy/cooling costs or availability dominate TCO, you might consider moving to cheap energy/cooling with adequate WAN bandwidth
  - ◆ Columbia River, West Texas, Canada datacenters?

- Metric: kW/GB vs. kW/disk – Which is correct?
- Store less stuff; delete when approved: Classify → ILM, HSM
- Location: Tiered Storage (SSD, SAS/FC, SATA, Tape, Optical)
- Increase effective Data Density on Disks (or Tape)
  - ◆ File de-duplication (Single-instance)
  - ◆ De-duplication (Factoring, Common Blocks)
  - ◆ Lossless Compression
- Trade-offs on Reliability, Performance
  - ◆ Single-copy of data?! (RPO, RTO)
  - ◆ Unpack/Inflate penalty may be incurred
  - ◆ Hotspots? – spread data across disks

# RAID level vs. Power/Cooling

- RAID (*Redundant Array of Independent Disks*), a family of techniques for managing multiple disks to deliver desirable cost, data availability, and performance characteristics to host environments.
- Despite capacity cost reductions exceeding Moore's Law, RAID is not 'free' – extra disks add CapEx and infrastructure costs
- plus **OpEx for Energy/Cooling**
- Compare RAID levels against equivalent JBOD (“Just a Bunch of Disks” = Capacity only)

# What Affects Storage Energy Use?

## RAID Definitions



Standalone



Cluster



Hot swap



RAID 0



RAID 1



RAID 5



RAID 0+1



# RAID level vs. Power/Cooling

- JBOD: **Number of disks scales to data capacity**
  - ◆ **Cost of Power/Cooling =  $N \times$  single disk cost**
- RAID 0 = data striping, disks required =  **$N$**
- RAID 1 = mirroring, disks required =  **$2 \times N$** 
  - ◆ **RAID 0+1 or RAID 1+0, power/cooling= $2 \times N$**
- RAID 5 = parity RAID, parity check data is distributed across the RAID array's disks.
  - ◆ disks required =  **$N+1$**
- RAID 6 = various methods to tolerate two concurrent disk failures; disks required =  **$N+2$**

# Storage Energy Use for other redundancy methods using Erasure Codes

- Erasure codes transform data from  $n$  blocks across multiple  $(n+m)$  blocks, such that recovery is possible with up to  $m$  failures.
  - ◆ Parity, as used in some RAID-levels is a optimal erasure code
  - ◆ Most complex erasure codes are less efficient than simple or DP parity, so the power/cooling costs are for **N+M** disks.
- Erasure codes can be very useful across distributed nodes with unreliable network connections
- **This could allow you to place some nodes into areas of lower energy cost!**
  - ◆ This lower cost could overcome the energy cost of extra nodes

# Resources and Links

- Lots of info available, but little is storage-specific
- Consultant\$ and other paid expert\$
  - ◆ Useful when dealing with gov't. regulations
- WEEE/ RoHS – lots of online tutorials
- The Green Grid is emerging as a resource for power and cooling issues (datacenter focus).
  - ◆ Partner with SNIA for Storage expertise
- Government agencies and Industry orgs
  - ◆ EPA Energy Star program

- Please send any questions or comments on this presentation to SNIA: [trackgreenstorage@snia.org](mailto:trackgreenstorage@snia.org).

**Many thanks to the following individuals  
for their contributions to this tutorial.**

**- SNIA Education Committee**

**SNIA Green Storage Initiative**

**Wayne Adams**

**Christian Belady**

**Rick Bauer**

**David Black**

**LeRoy Budnik**

**Bill Mottram**

**Deborah Johnson**

**Erik Riedel**

**SW Worth**

**Alan Yoder**

- Metrics for Power/Cooling for Storage
- Primer on Electric Power in the United States
  - ◆ Generation, Transmission, Distribution
  - ◆ Pricing: What you buy, Variable factors
    - › Geographic variations

# Storage-specific Metrics

- What are the appropriate metrics for Storage?
  - ◆ Energy (kWh) per 'Unit' of Storage (MB, GB, TB)?
- **Table** of Storage Operations vs. Energy-cost for each type of Storage technology
- Devices: Flash, Disk, Tape, Optical
- Operations: CRUD + periodic Validation of Data or other Maintenance Op (e.g. Re-tensioning tape)
  - ◆ Blocks vs. Files?

# Green Storage - Appendix

- Overview, Motivation, and Definitions
- eWaste Reduction/Recycling: RoHS, WEEE, etc.
- Fundamentals of Energy and Cooling
- **Electricity Primer in the United States**
- Datacenter Design and Operation
- Storage Components and Technologies

# Electricity Supply in the U.S.

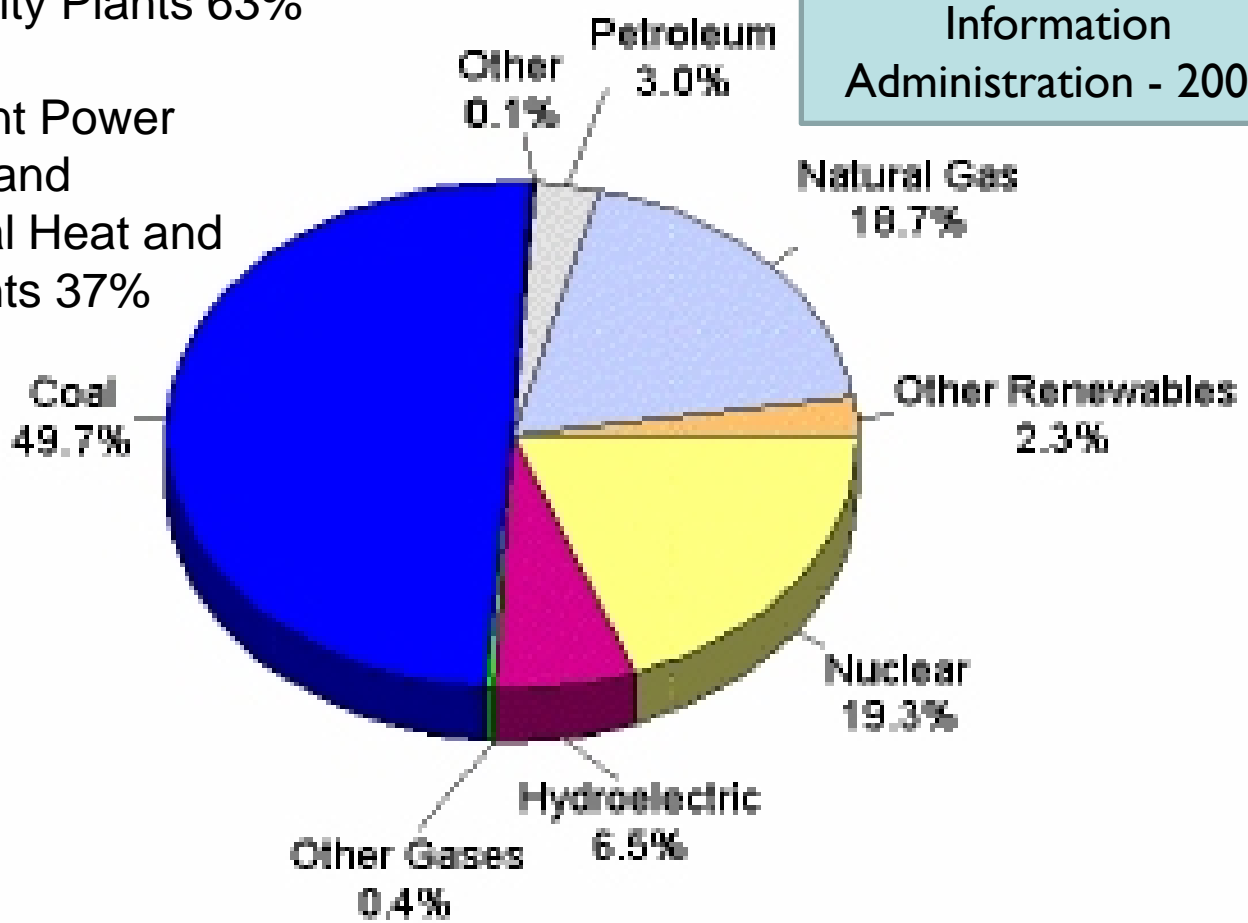
- System Reliability and Flow ‘Control’
  - ◆ North American Electric Reliability Corp.
  - ◆ ISO (Independent System Operator) for each region controls “Congestion”, Ancillary Services, etc.
- Generation, Transmission, Distribution
  - ◆ Regulated monopolies → Deregulation (partial)
- Pricing: What you buy, Variable factors, Geographic variations, Roles of Regulators, ISOs
- Metering:
  - ◆ **Energy** kWh or MWh (at multiple points in system)
  - ◆ **Demand** (peak, to compensate for infrastructure)

# Electric Power Generation by Source SNIA

**Total = 4,055 Billion kWh**  
Electric Utility Plants 63%

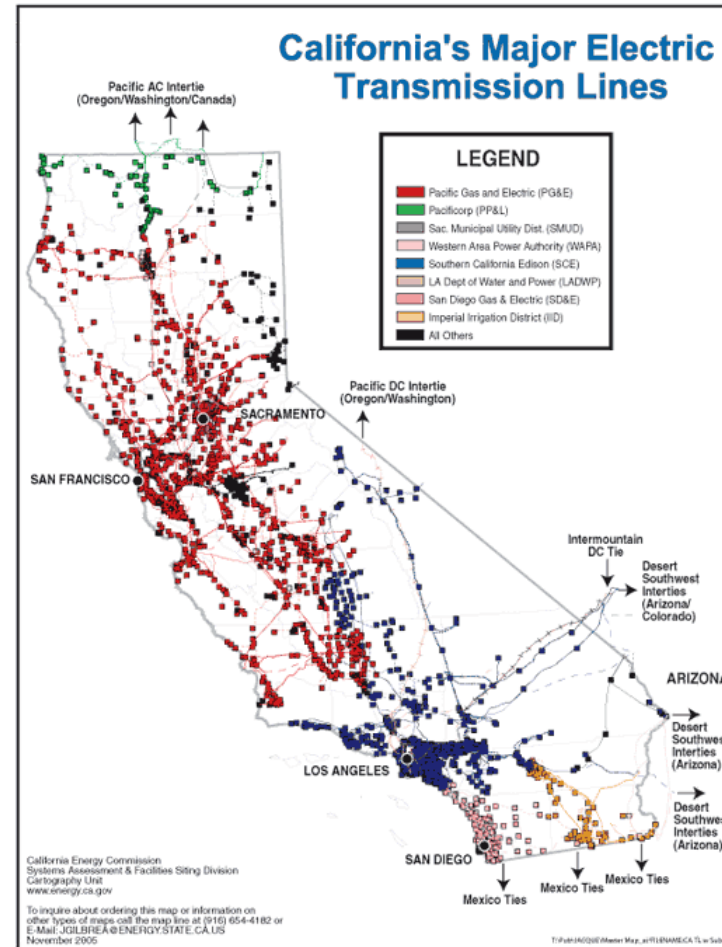
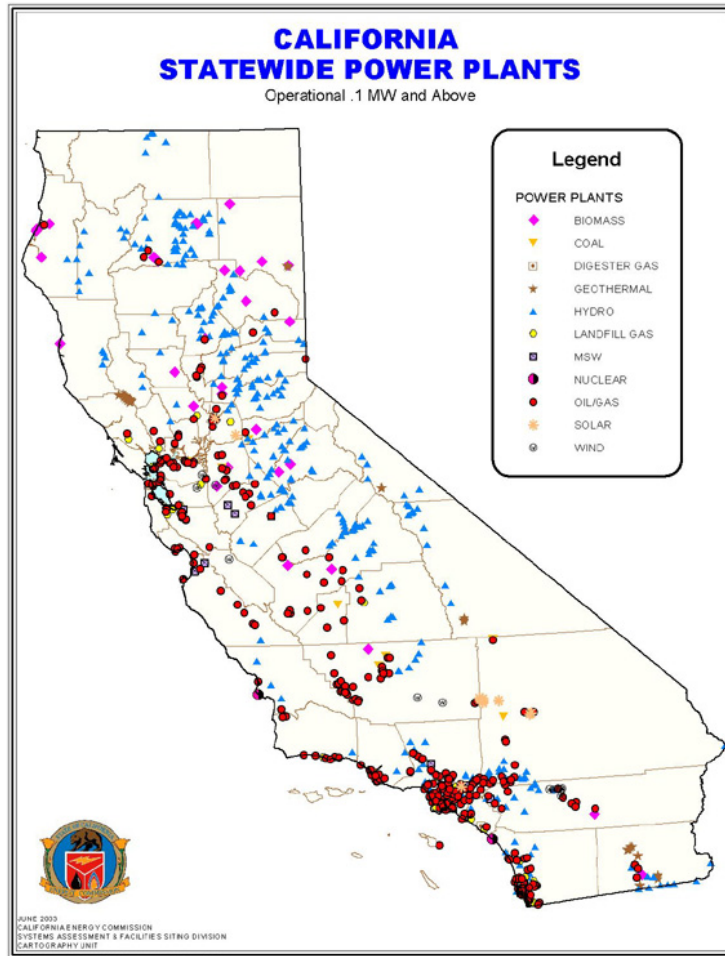
Independent Power  
Producers and  
Commercial Heat and  
Power Plants 37%

Data from U.S. Energy  
Information  
Administration - 2005





# Calif. Generation; Transmission Interconnects



[www.energy.ca.gov/maps/](http://www.energy.ca.gov/maps/)

# Electricity Supply in the U.S.

## - Integration

- ▶ Traditional vertically-integrated monopolies
  - ◆ Generate (shared with “IPPs”), Transmit, and Deliver power to everyone in their monopoly service territory
  - ◆ Heavily regulated by federal (FERC), state (PUC/ PSC), and almost every other level of gov’t.
  - ◆ Emphasis on efficient allocation of Capital with very long time horizon, then recovery from “rate-base”
  - ◆ Emphasis on operational reliability and system security, but getting squeezed by regulators
  - ◆ Customer-service metrics overseen by regulators
  - ◆ IT focus on SCADA, Meter/Billing, Property Records

# Electricity Supply: Deregulation

- Deregulated Segments (U.K. model, NY, CA)
  - ◆ Generation (“**GenCo**” or “**IPP**”) – generates power from Coal, Nuclear, Natural Gas, Hydro, Oil, “Other”
    - Some “Green” generation: Wind, Tidal, Geothermal, Bio-mass, or Hydro(!?!?)
  - ◆ Transmission (“**TransCo**” or **ISO**) – aggregates supply and “ancillary services” from GenCos and IPPs via Auctions, and moves power to wholesale
  - ◆ Distribution (“**DisCo**” – a ‘natural’ monopoly which moves power to retail customers
  - ◆ Meters are everywhere throughout the system

- Interchanges, Power Pools, and “Reliability Councils” (e.g. ERCOT)
- NYMEX – transparent open-auction pricing of standardized “bundles” at major delivery points (e.g. Oil, NatGas, Coal, Electricity-PJM)
- GenCo: long-term contract pricing to DisCos and Large End-Users (may be distributed locations)
- Fuel Adjustment pricing (may be volatile!)
- IPP: Independent Power Producer
- ISO as maintainer of reliability and open auction

# Electricity prices are variable (at least at the wholesale level)

- **Electricity cannot be stored effectively!**
- **Prices vary with DEMAND (local and regional)**
  - ◆ Weather (Hot, Cold, or Both)
  - ◆ Time-dependent: Daily, Weekly, Seasonally
  - ◆ Economic conditions – general, regional
- **Prices vary with SUPPLY (local and regional)**
  - ◆ CapEx: plant construction (NIMBY), maintenance
  - ◆ OpEx: Fuel costs dominate – swings can be wild
- **Electricity Transmission congestion/losses increase cost; hard to build new lines (NIMBY)**