



Education


# **Green Storage –The Big Picture**

SW Worth, Microsoft  
SNIA Green Storage Initiative

Author: 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.**

# Do you understand “Green”?



THE WORD "SUSTAINABLE" IS UNSUSTAINABLE.

<http://xkcd.com/1007/>

# How well do you know Storage?



Standalone



Cluster



Hot swap



RAID 0



RAID 1



RAID 5



RAID 0+1



# ‘Green IT’: it’s not just kW (or kWh)!

- Focusing on **energy usage is important!**
  - ◆ CapEx/OpEx: \$/Euro/Yen/Rmb
- But, there are a lot more factors in ‘Green’...
- **“Embedded”**= separately accounted inputs (Cap/Op)
  - “Beef Jerky” = embedded GHG, Water, Energy, etc.
- ‘Carbon’ (proxy for GHG – Green House Gases)
- ‘Embedded Energy’ (e.g. glass, metal, plastics)
  - ◆ By one calc, embedded energy in a server ~=1000 kWh
- Water – embedded and operational
- Toxins’ (RoHS/WEEE, Coal by-products, etc.)
  - ◆ Coal power plants: Ash, Hg, As, Th (Radioactivity)

# Green Storage - Agenda

[www.snia.org/education/tutorials](http://www.snia.org/education/tutorials)

- **Overview, Definitions, and Motivation**
  - ◆ **Some basics: Economics, Ecology, Energy, Engineering**
- Systems thinking: TCO lifecycle, with embedded resources
- How we talk about 'Green Storage' (*good/bad news*)
- Electricity Pricing in the United States
- Datacenter Design and Operation
- Storage Components and Technologies

# Overview and Definitions

- **‘Green’** – What does it mean to the IT ‘ecosystem’?
  - ◆ Reduction of Total Impact on Environment
    - › Systems approach - More than just Energy Use!
  - ◆ Defined by Govs (EPA, EU, Kyoto), Orgs, Vendors, etc.
- How does “Green” differ from normal economic considerations, e.g. TCO, efficiency, optimization?
  - ◆ Systems viewpoint + Lifecycle analysis
  - ◆ Widen scope of action across system/org boundaries, time
  - ◆ Rationalize decisions by including “**externalities**”
- *Bottom-line*: ‘Green’ effects on **IT/Storage** decisions?

- I.T. owners / **Data Center operators** (“Customers”)
- **Vendors** of I.T. hardware, software, systems, services
  - ◆ Engineers/Developers/Architects – *including Cloud vendors!*
- Energy **Utilities and Regulators**
- **Governments:** local, regional, national, supra-national
  - ◆ US-EPA ENERGY STAR programs
  - ◆ Euro. Comm. Code of Conduct on Data Centre Energy Effic.
- **Green Grid** metrics [www.thegreengrid.org](http://www.thegreengrid.org)
  - ◆ Focus on Power, Energy, and Cooling used for IT
- ➔ **SNIA** – org expertise on enterprise STORAGE
- Other interested parties (e.g. Uptime Institute)



- **“Fear”**: Constraints (Regulations, Physical limits)
  - ◆ Do what you are forced to do by Gov. Regulations
  - ◆ Physical Laws are not optional...*but alternatives may exist*
    - › 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, Environmental, Economic”
- **“Greed”**: Profit Maximization / Cost Minimization
  - ◆ Strategy ➔ Capital Expenses (CapEx, e.g. UPS, GenSet, Fan)
  - ◆ Tactics ➔ Operational Expenses (OpEx, e.g. energy, kWh)
  - ◆ TCO (Total Cost of Ownership) integrates CapEx and OpEx

# Externalities

- Problem: inputs/outputs with unclear prices or owners
  - ◆ e.g. clean water, Green-house gases (GHG ~= 'Carbon')
  - ◆ Some factors are effectively Zero-cost to the decision-maker, but are not cost-free to others affected
  - ◆ Non-optimal behavior: '*Tragedy of the commons*'
- Solutions: (*choose none, one, or a mix*)
  - ◆ Government **mandates** (*Regulate/Enforce vs. Engineer*)
    - › Separate accounting system, e.g. for Carbon 'Footprint'
  - ◆ Pigouvian **tax**: '*Sin Taxes*' (*modify behavior*)
  - ◆ Coase **markets**: PropertyRights+TransparentPrices=
    - › ...**optimal\* outcomes!** (\*subject to a few assumptions....)
    - › Cap-and-Trade Carbon (compare SOx/NOx trading schemes)

# TCO = amortized CapEx + OpEx

(in whatever accounting system(s) you choose to use)

- Ideally, one accounting measure: \$/Euro/Yen/Rmb
  - ◆ For this to work, all constrained resources must have prices!
  - ◆ But, prices of these resources may vary due to supply-chain locations
  - ◆ Is it acceptable to 'outsource' embedded Carbon? (Are impacts global?)
- Other accounting systems are possible (but complicated):
  - ◆ TBL (Triple Bottom Line): “Social, Environmental, Economic”
    - › Pronounced “**People, Planet, Profits**”
  - ◆ Energy – measure as ‘**embedded**’ energy
  - ◆ ‘**Carbon**’ as proxy for GHG (Greenhouse Gases: CO<sub>2</sub>, CH<sub>4</sub>, SF<sub>6</sub>, etc.)
  - ◆ Water, land, or any other scarce resource....
- Production and Disposal costs: ‘should’ be embedded in **CapEx** (possibly amortized) or appear as **OpEx**
  - ◆ CapEx: Embedded resources, w/ ‘unified’ accounting
    - › From production inputs to customer delivery
- OpEx: mostly Energy (especially true for **Storage**)


# Refresher: 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  $<100\%$  (sometimes  $<<100\%$ )
- 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

# Refresher: Heat Transfer Physics

- Heat transfer – *usually away from IT gear, i.e. “Cooling”*
  - ◆ Conduction: thermal glue/grease - CPU to 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: paint the roof of your datacenter white?
- Phase Change: Solid-Liquid; Liquid-Gas
  - ◆ +/- Heat (large quantities) without Temperature change.
- Newton’s Law of Cooling
  - ◆ Rate varies with Temperature Difference
    - › “Hot aisles separated from Cold aisles!”

## Green Grid metrics [www.thegreengrid.org](http://www.thegreengrid.org)



DCE is now called **DCiE**:  
Data Center Infrastructure Efficiency

Picture

- Naïve calc of 1 MW datacenter energy costs:
  - ◆ 1 MW (MegaWatt) = 1,000,000 Watts (**Power** for I.T. gear)
  - ◆ 8760 hours/year
  - ◆ Avg. price of electricity in U.S. ~=\$0.10/kWh
  - ◆ Therefore: **Energy** budget \$876,000/year; *right?*
  - ◆ WRONG, because Utilization < 100% (*usually <<100%*)
    - › ...and because Energy Pricing increasingly varies with time!
- ‘**Demand**’ Charges recover Utility CapEx (peak ~15 min.)
- ‘Peak’y’ IT workloads are a Bad Thing!
  - ◆ Black Friday/Monday (largest online retail days in the U.S.)
  - ◆ Etc. – half-time during a big football match/game

# Power is NOT the same as Energy!

Item	Power	Energy
Units	Watts, kW, MW	kiloWatt hours (kWh), MWh
Measurement	Instantaneous	Integrated over time
Physical evidence	Infrastructure, Equipment	Usage (Electricity, Cooling)
→ Examples	UPS, PDU, CRAC, AHU, <u>plus</u> IT Gear: Servers, Storage, Networking,	<u>Variable</u> usage (‘consumption’) of electricity, water, fuel
Expense (Cost Accounting)	Capital (CapEx)	Operational (OpEx)
Internal Cost Recovery	Chargeback (Amortize)	Energy chargeback (variable)
External Cost Recovery	‘Demand’ Charges or Rent	Facilities chargeback (fixed)
Billing units	\$/kW ( <u>peak</u> 15-min. period)	\$/kWh x Total usage

## Bottom-line:


1. Big data centers charge-back on Power, Energy, or both.
2. Power-related Infrastructure costs may outweigh Energy Costs!



- **“Embedded”**= separately accounted inputs included
  - “Beef Jerky” = embedded GHG, Water, Energy, etc.
- Cambridge Institute study: ‘Embedded Carbon’
- ‘Embedded Energy’ (e.g. glass, metal, plastics)
  - ◆ By one calc, embedded energy in a server  $\sim$  1000 kWh
- ‘Embedded Water’
  - ◆ manufacturing, energy prod., data center cooling
- ‘Embedded Toxins’ (Heavy metals, etc.)
  - ◆ RoHS: Pb, Cd, Hg, Cr(6), PBB, PBDE
  - ◆ Coal power plants: As, Th, Hg, ash, ...
- ‘Embedded Radioactivity’: Coal (Th), Nuclear-waste,

# Outsourcing 'Embedded' Carbon


- Carnegie Institute study in PNAS:
  - ◆ [http://www.ciw.edu/news/carbon\\_emissions\\_outsourced\\_developing\\_countries](http://www.ciw.edu/news/carbon_emissions_outsourced_developing_countries)
- EU 'imports' a third of its carbon emissions



# eWaste Reduction/Recycling

- **GREEN: more than Electricity and Carbon!**
- Government regulations/directives on 'eWaste' affect I.T./storage vendors (and their customers).
- WEEE
- RoHS, China-RoHS
- Packaging and Pkg Waste
- Halogens (in plastics)
- Basel Convention/Basel Ban (Transboundary Wastes)

- 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, Definitions, and Motivation
- **Systems thinking: TCO, over entire Lifecycle**
- How we talk about 'Green Storage' (*good/bad news*)
- Electricity Pricing in the United States
- Datacenter Design and Operation
- Storage Components and Technologies



# “Green” effects on I.T.

- “TCO” (Total Cost of Ownership) should include Externalities in accounting and purchase decisions
  - ◆ In most cases Externalities will evolve to provide clear pricing signals (e.g. RoHS, WEEE, Cap-and-Trade)
- Systems viewpoint (*bigger picture*) is essential!
- Expand scope of decision-criteria and constraints to include (at least) entire datacenter (*entire supply chain?*)
  - ◆ Servers, Networking, and **Storage**
  - ◆ Power, energy, and cooling (CapEx and OpEx)
  - ◆ People: widen their decision-boundaries, -constraints
    - › Include your Facilities managers!

# Three Stages of Product Life-Cycle

- Birth: Product Creation (design for recycle/disposal)
  - ◆ ‘Embedded’ resources in CapEx? – see 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 (Tape cartridges?)
  - ◆ Dominated by OpEx (but is this cost visible to IT?)
- End-stage: Removal, Recycling, Disposal
  - ◆ Integrated into initial CapEx or OpEx surcharges
  - ◆ *Alternative: dump these costs onto everyone else....*

# Green Storage - Agenda

- Overview, Definitions, and Motivation
- Systems thinking: TCO, over entire Lifecycle
- **How we talk about 'Green Storage' (good/bad news)**
- Electricity Pricing in the United States
- Datacenter Design and Operation
- Storage Components and Technologies

# How we talk about ‘Green’ Storage...

- How much data center Energy Usage is due to Storage?
  - ◆ ***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’)
- ➔ *Proportion of Energy used by Storage is increasing, because of...*
  - ◆ Facilities improvements (PUE, DCiE)
  - ◆ Virtualization – especially of Servers, O.S., Applications
  - ◆ Management – easier to move Servers between on/sleep/off
    - › Not as easy with storage devices (*more info later*)

# Storage gets more (bad) visibility?

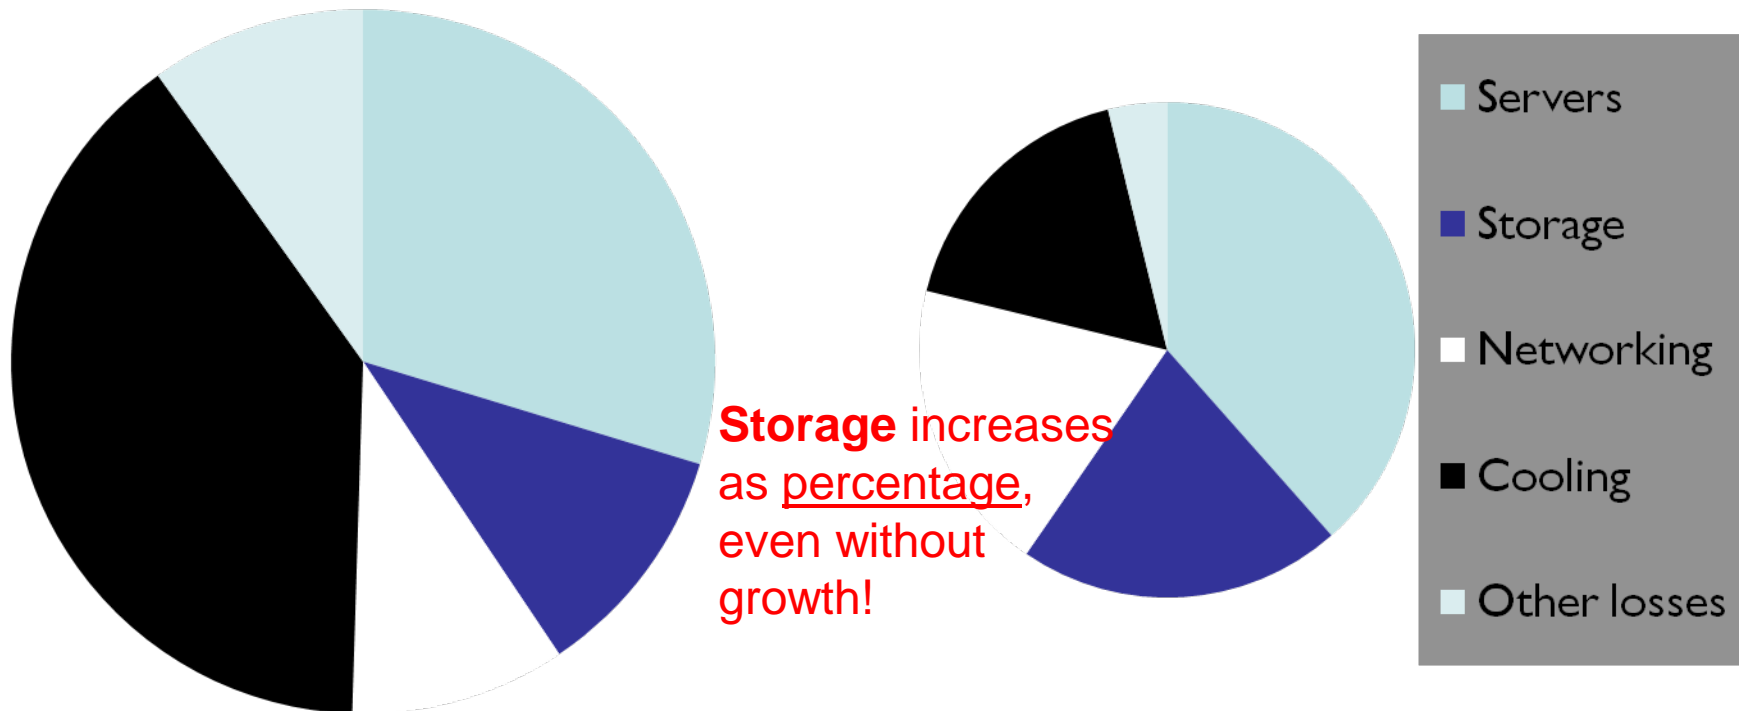
$$\text{DCiE} = (\text{IT Energy Demand}) / (\text{Total Datacenter Energy Usage})$$

(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%)**



# Storage-specific Power/Cooling data

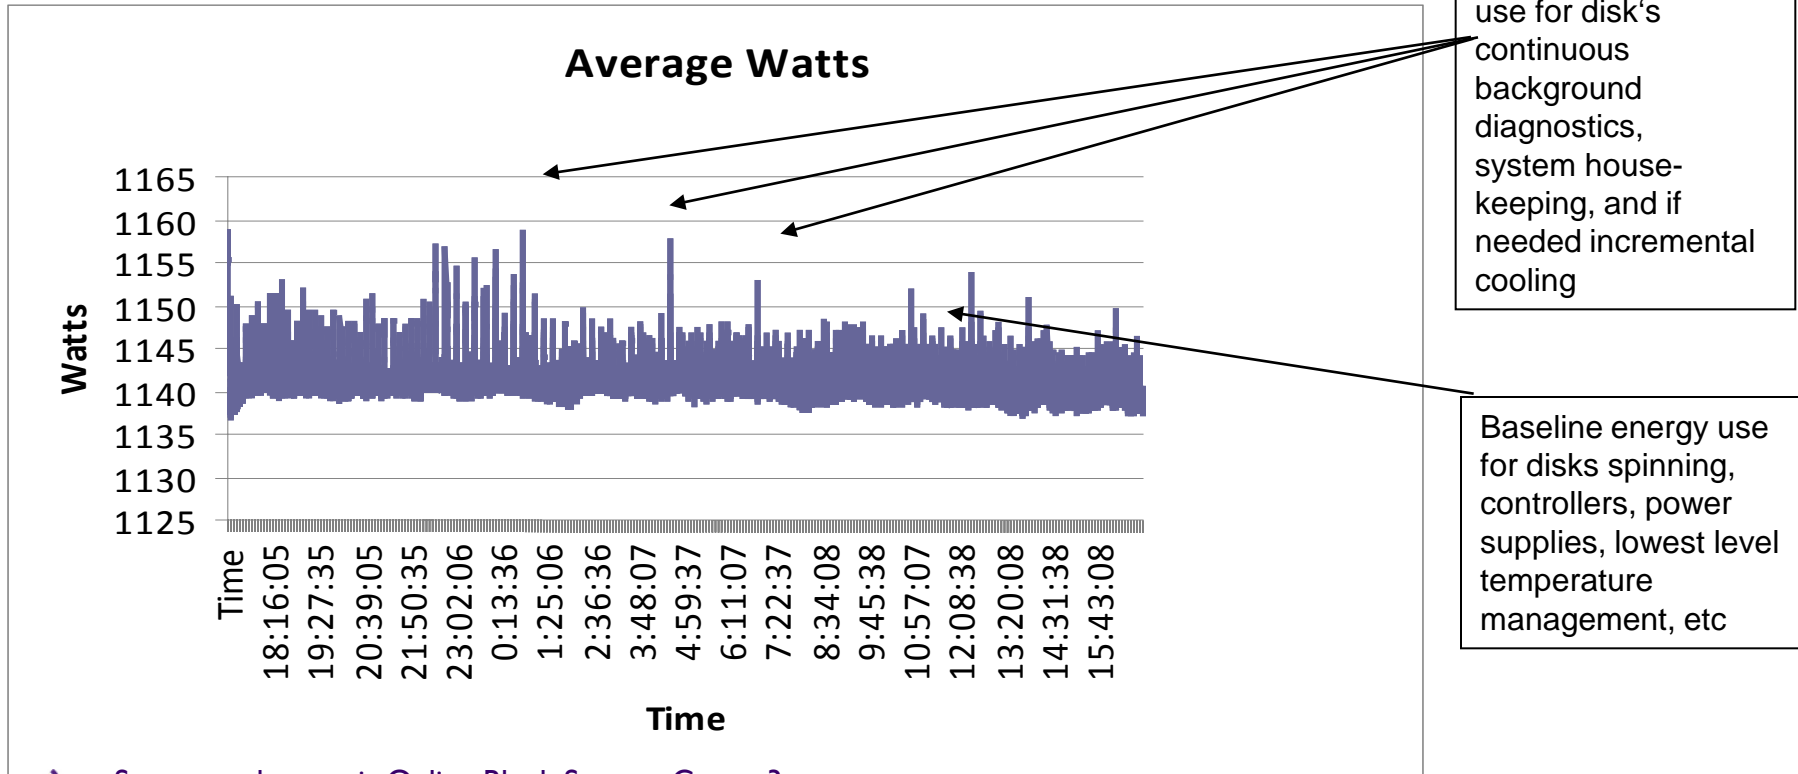
- Each component of a Storage system has Power and Cooling requirements
  - ◆ Understand “**Idle**” (\**not*\* ‘stand-by’) vs. “Loaded” (R/W)
  - ◆ Label ratings are usually peak power required
    - › If you design using this data, your power/cooling equipment may be (grossly) over-built (Bad!), and 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 check the SNIA Emerald data...

# 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

# SNIA: Emerald Program, and working with EPA ENERGY STAR

## Sample 'Idle' Storage System Measure



◆ System under test is Online Block Storage Group 3

◆ Notes:

- Graph is for a 24 hour measurement time period; graph does not include the pre-conditioning phase
- A profile of a graph for Group 2 Online Storage, e.g. 100% JBOD, would not have the energy related activity spikes for disk & system background diagnostics and data validation. **This is one of the major differences between Online Storage Group 2 and 3 Systems - a procurement tradeoff for energy consumption versus system and data reliability.**



# European Commission Code of Conduct on Data Centres Energy Efficiency



- ‘Voluntary’ initiative for education and shared best practices; “agreed commitments” for participants
- SNIA-Europe is an Official Endorser

# 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)

# What Affects Storage Energy Use?

## Redundancy and RAID Definitions



Standalone



Cluster



Hot swap



RAID 0



RAID 1



RAID 5

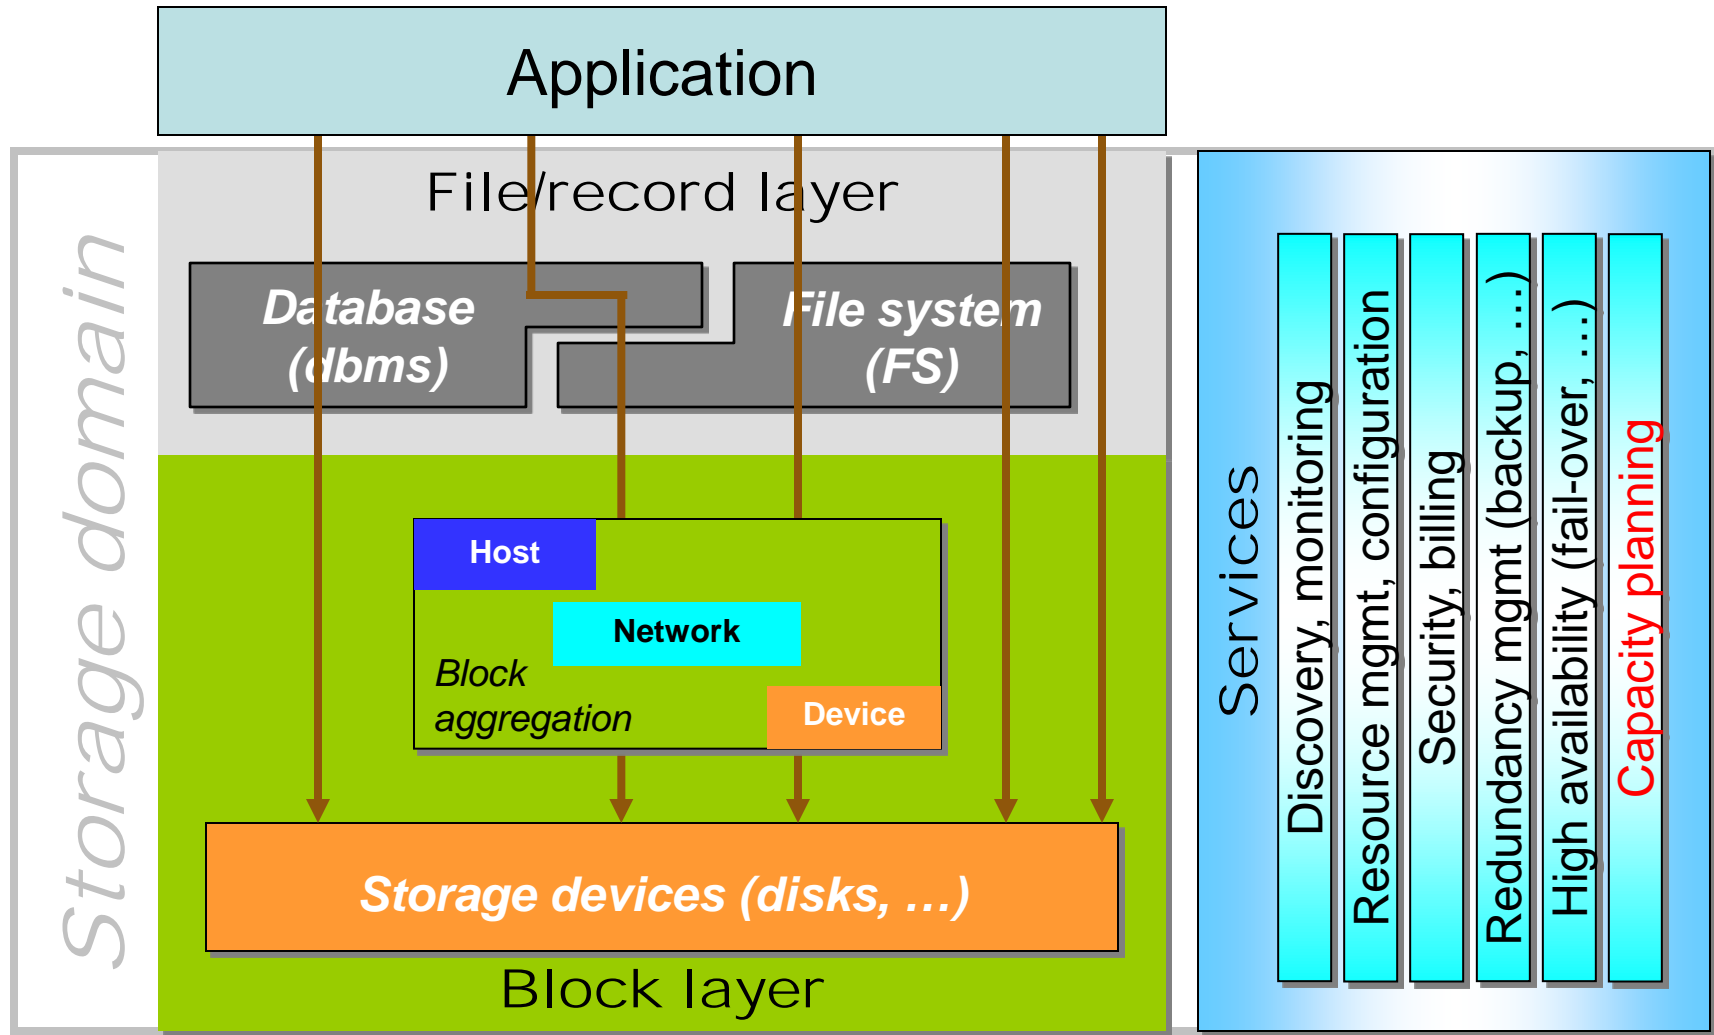


RAID 0+1



# What Storage aspects could be affected?

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



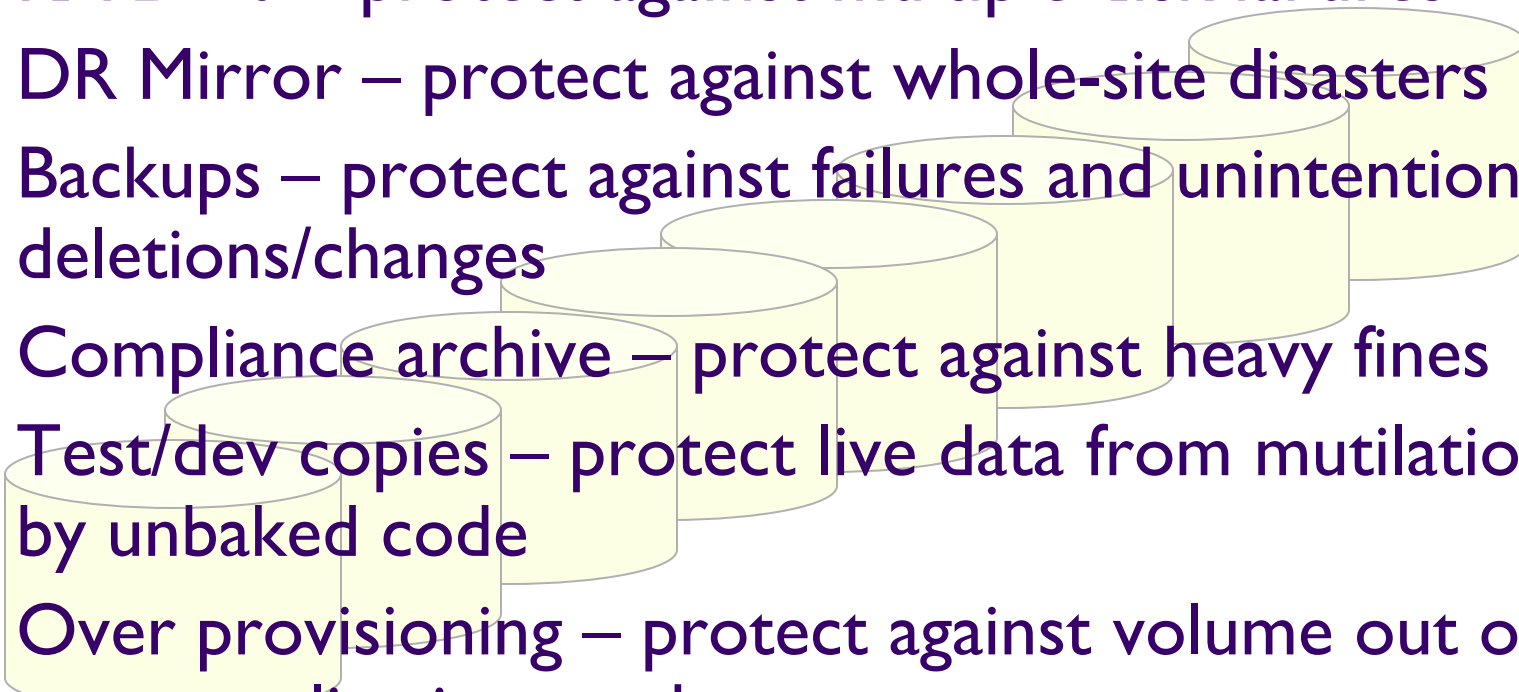
# SNIA Storage Taxonomy

<b>Storage Taxonomy Summary</b>	<b>Online Storage</b>	<b>Near Online Storage</b>	<b>Removable Media Libraries</b>	<b>Virtual Media Libraries</b>	<b>Infrastructure Appliances</b>	<b>Infrastructure Interconnect</b>
	Prime storage, able to serve random as well as sequential workloads with minimal delay	Intended as second tier storage behind Online Storage. Able to service Random and Sequential workloads, but perhaps with noticeable delay in time to 1 <sup>st</sup> data access.	Archival storage used in a sequential access mode. A Typical example would be Tape based archival, both Stand Alone and Robotically assisted libraries.	Storage which simulates removable Media Libraries. Will typically use non tape based storage and as such are able to respond to data requests more quickly	Devices placed in the storage SAN or network adding value through one or more dedicated Storage enhancements. Examples include: SAN Virtualization, Compression, De-duplication, etc.	Devices which enable a SAN or other Storage Network data switching or routing.
<b>Maximum Capacity Guidance</b> <small>Note: Maximum Capacity Guidance reflects the maximum capacity a given offering can be purchased with and/or field upgraded to. It is intended to be used as a guideline as apposed to an absolute value. There will be case where a device may have greater or small capabilities, but otherwise is an appropriate match for a given classification due to other criteria, e.g.: redundancy capabilities</small>	<b>Max Storage Devices</b>	<b>Max Storage Devices</b>	<b>Max Tape Drives</b>		<b>Max Storage Devices Supported*</b>	<b>Max Port Count</b>
<b>Group 1) SoHo &amp; Consumer</b>  Storage which is designed primarily for home (consumer) or home / small office usage. –Often Direct Connected (USB, IP, etc) –No option for redundancy (will contain SPOFs)	<b>Up to 4 Devices</b>		<b>Stand Alone Drive (No Robotics)</b>		<small>Note: * Infrastructure Appliances by definition have no intrinsic storage, other than what is used for local processing and/or local Cashing of data.</small>  Storage Devices Support in this case refers to the number of storage devices controllable down stream of the Appliance	
<b>Group 2) Entry, DAS, or JBOD</b>  Storage which is dedicated to one or at most a very limited number of servers. Often will not include any integrated controller, but rely on server host for that functionality. –Often Direct Connected (SATA, IP, etc.) –May optionally offer limited number of redundancy features		<b>Up to 4 Devices</b>	<b>Up to 4 Drives</b>			<b>Up to 32</b>
<b>Group 3) Entry / Midrange</b>  SAN or NAS connected storage which places a higher emphasis on value than scalability and performance. This is often referred to as 'Entry Level' storage. –Network connected (IP, SAN, etc.) –Has options for redundancy features	<b>More than 20 Devices</b>	<b>More than 4 Devices</b>	<b>More than 4 Drives</b>	<b>Up to 100 Devices</b>	<b>Support for up to 20 Devices</b>	<b>Up to 128</b>
<b>Group 4) Midrange / Enterprise</b>  SAN or NAS connected storage which delivers a balance of performance and features. Offers higher level of management as well as scalability and reliability capabilities. –Network connected (IP, SAN, etc.) –Has options for and often delivered with full redundancy (no SPOF)	<b>More than 100 Devices</b>	<b>More than 100 Devices</b>	<b>More than 24 Drives</b>	<b>More than 100 Devices</b>	<b>Support for more than 20 Devices</b>	<b>More than 128</b>
<b>Group 5) Enterprise / Mainframe</b>  Storage which exhibits large scalability and extreme robustness associated with Mainframe deployments, though are not restricted to Mainframe only deployments. –Mainframe connectivity with optional network connection (IP, SAN..) –Always delivered with full redundancy (no SPOF) –Often Capable of non-disruptive serviceability	<b>More than 1000 Devices</b>		<b>More than 11 Drives</b>	<b>More than 100 Devices</b>	<b>Support for more than 100 Devices</b>	


© SNIA 2009

See: *Green Storage Power Measurement Specification* for complete details

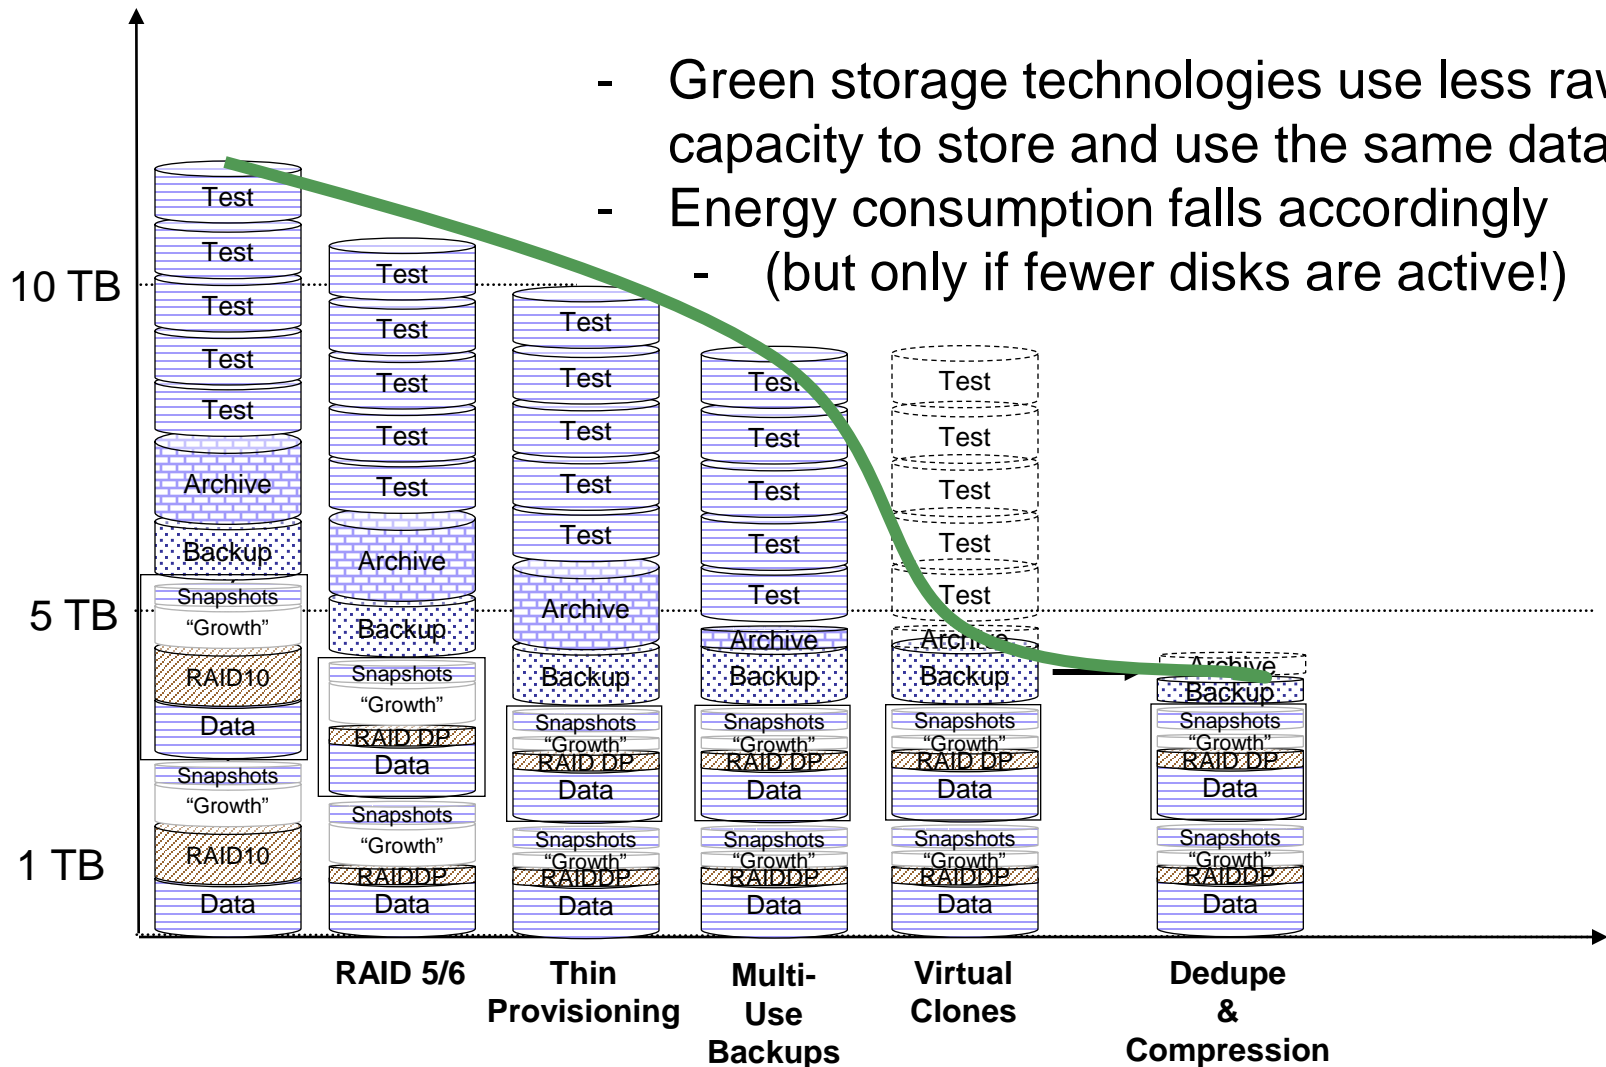
# Need for redundancy

- 
- RAID 10 – protect against multiple disk failures
  - DR Mirror – protect against whole-site disasters
  - Backups – protect against failures and unintentional deletions/changes
  - Compliance archive – protect against heavy fines
  - Test/dev copies – protect live data from mutilation by unbaked code
  - Over provisioning – protect against volume out of space application crashes
  - Snapshots – quicker and more efficient backups

# Result of redundancy



# Effect of green technologies





# Are savings multiplicative?

## ➤ Sometimes yes

- ◆ RAID 6 + writeable clone
  - › Assume 1000GB writeable clone = **2000GB** needed for a raw writeable copy on RAID 1 storage
  - › 90% writeable clone savings takes us to 200GB
  - › 35% RAID 6 savings takes us from there to 130GB →
  - › **130GB** needed for a writeable clone on RAID 6

## ➤ Sometimes no

- ◆ Thin provisioning + resizable volumes
  - › Similar effects, but you only get the savings once

## ➤ Sometimes maybe

- ◆ Snapshot + deduplication
  - › Can't dedup readonly snapshots
  - › Snapshots are a form of deduplication, so there's less to dedup
  - › OTOH, already deduped data can be snapshotted efficiently

# Savings matrix

Savings multiply in combinations with checkboxes

	C	SS	VC	TP	R	DD	RV
Compression (C)							
Snapshots (SS)							
Virtual Clones (VC)							
Thin Provisioning (TP)							
RAID (R)							
Deduplication (DD)							
Resizeable Vols (RV)							

# Savings matrix (cont.)

- ▶ E.g. Thin provisioning with snapshots, RAID 6, and Dedup – big win!

	C	SS	VC	TP	R	DD	RV
Compression (C)		✓			✓	✓	
Snapshots (SS)	✓		✓	✓	✓	✓	✓
Virtual Clones (VC)		✓					
Thin Provisioning (TP)		✓			✓	✓	
RAID (R)	✓	✓		✓		✓	✓
Deduplication (DD)	✓	✓		✓	✓		✓
Resizeable Vols (RV)		✓			✓	✓	

# Green Storage - Agenda

- Overview, Definitions, and Motivation
- Systems thinking: TCO, over entire Lifecycle
- How we talk about 'Green Storage' (*good/bad news*)
- **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 (e.g. drop in 2009)
- **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)**

# Electric Power Generation by Fuel

**Total = 4,055 Billion kWh**

Electric Utility Plants 63%

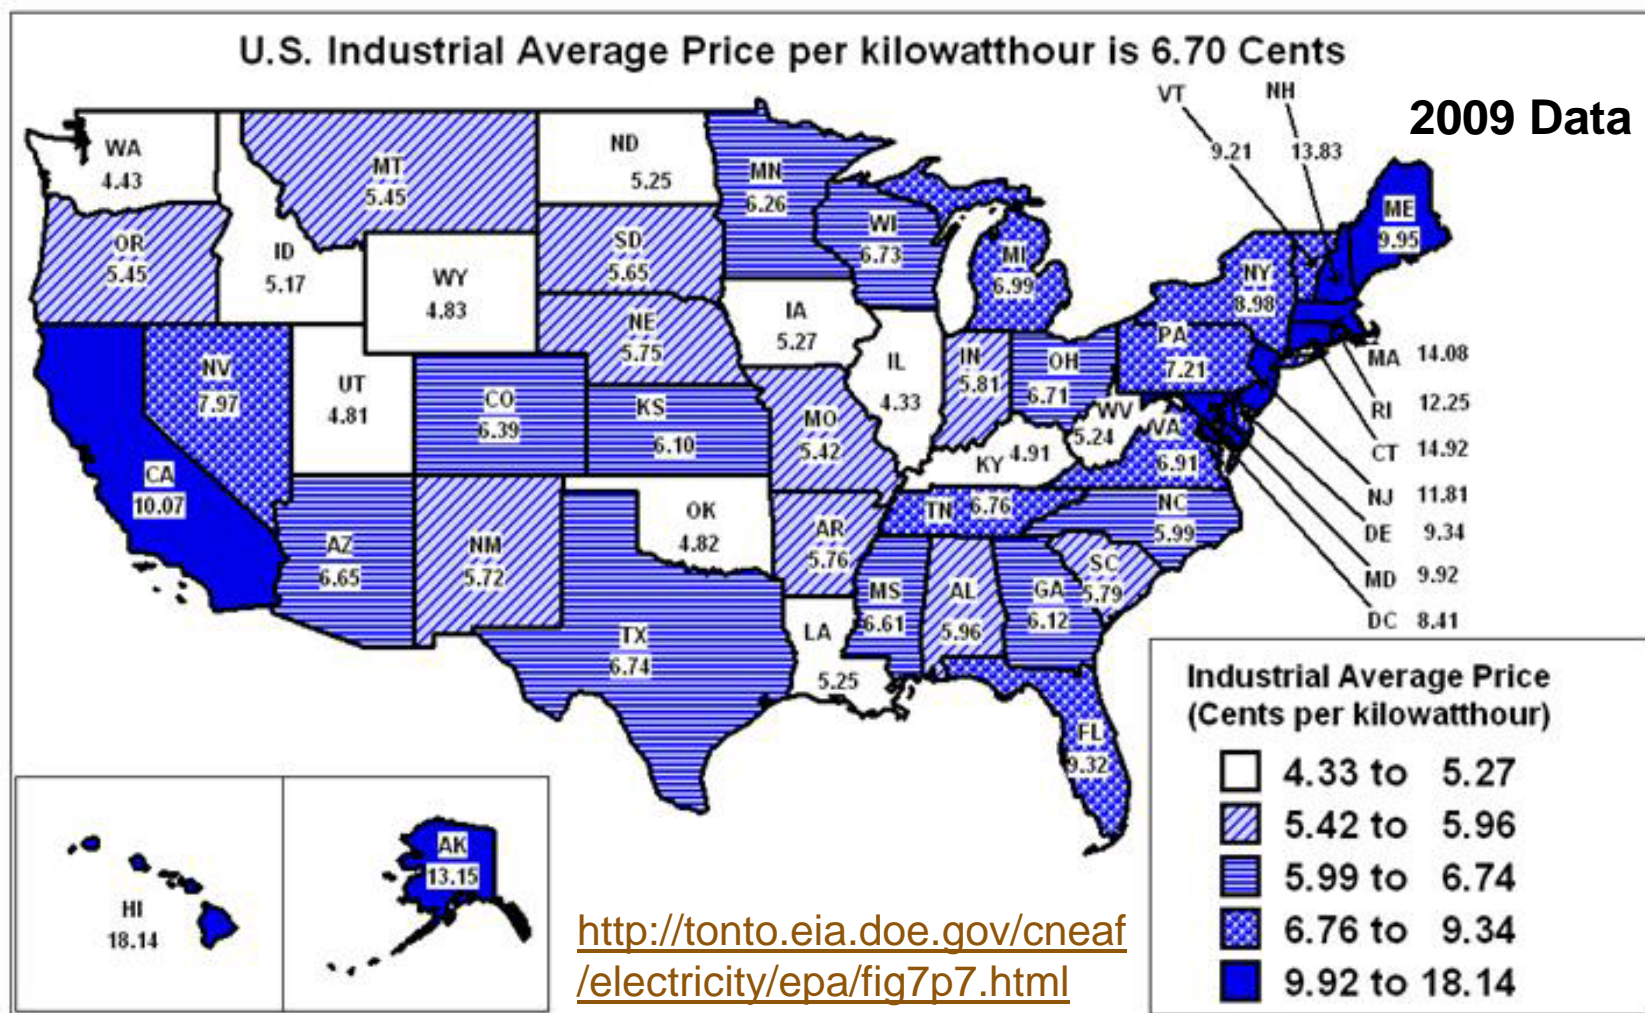
Independent Power  
Producers and  
Commercial Heat and  
Power Plants 37%



Education  
**SNIA**



# Avg. Industrial Price of Electricity

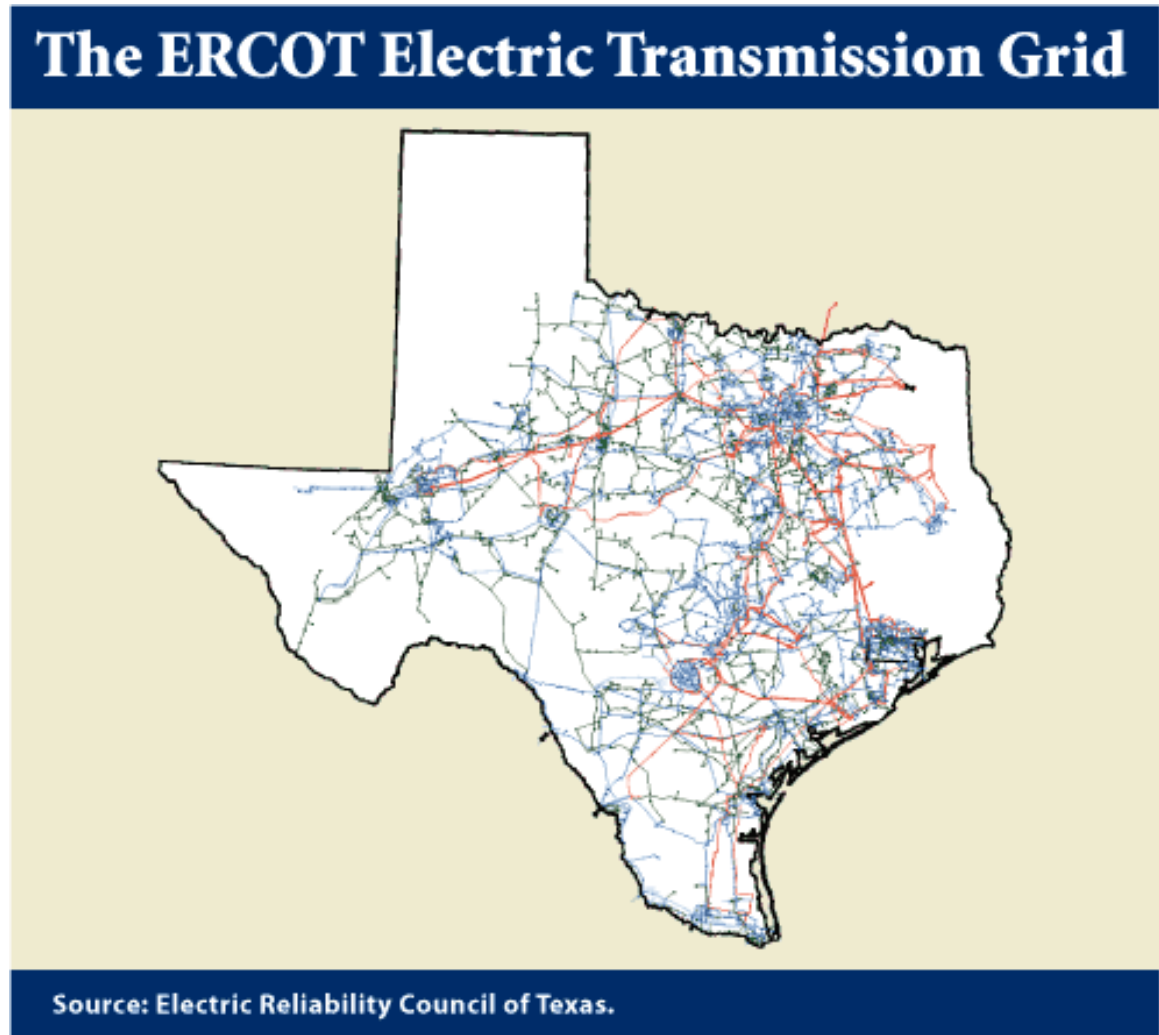



Note: Data are displayed as 5 groups of 10 States and the District of Columbia.

Source: U.S. Energy Information Administration, Form EIA-861, "Annual Electric Power Industry Report."




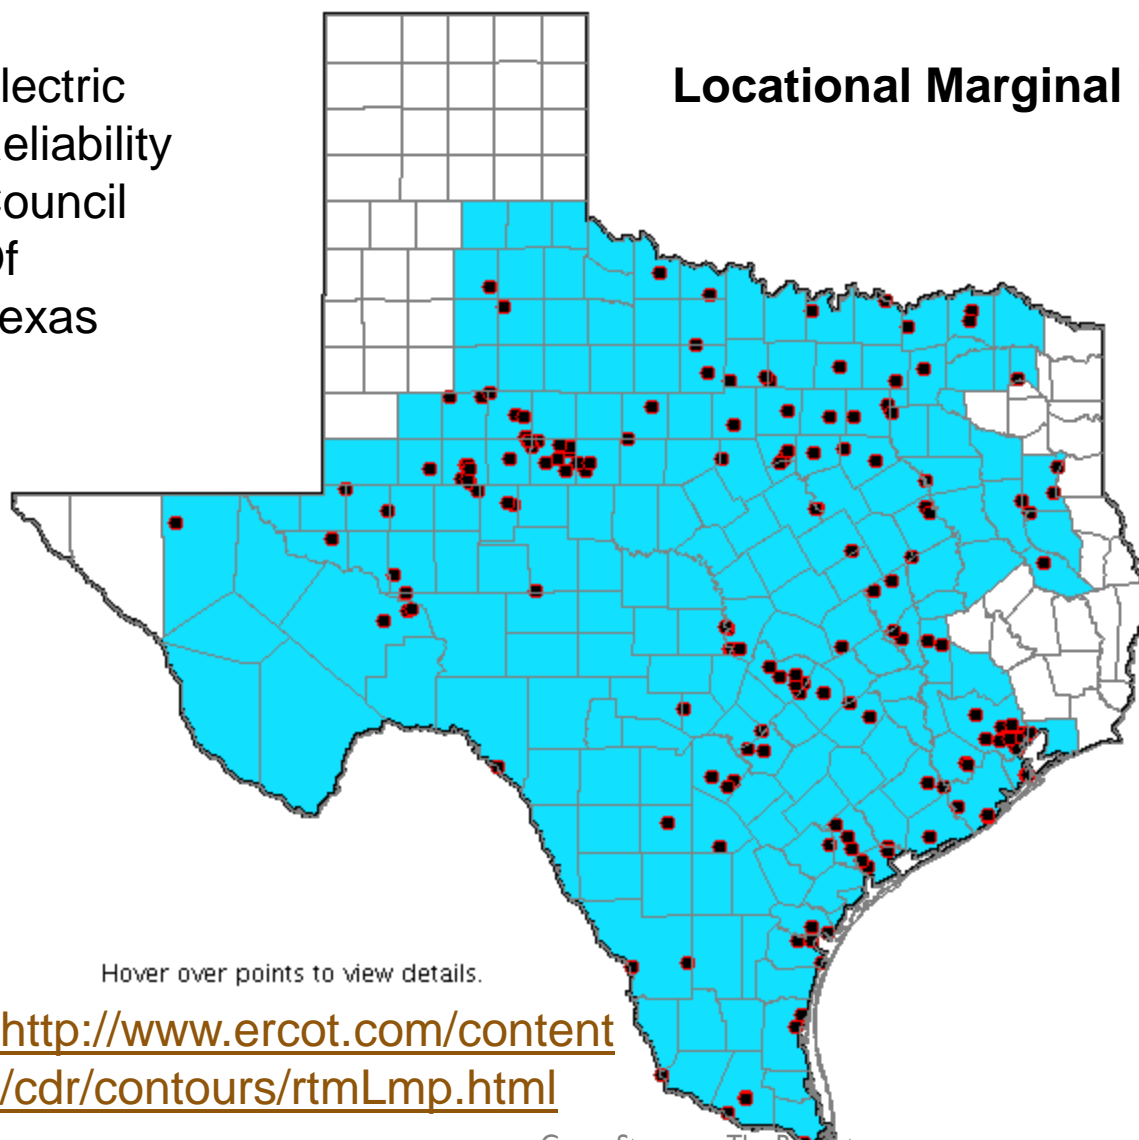
# Texas Generation (Capacity vs. Actual); ERCOT Transmission Grid



# Real-time Electricity Pricing

Electric  
Reliability  
Council  
Of  
Texas

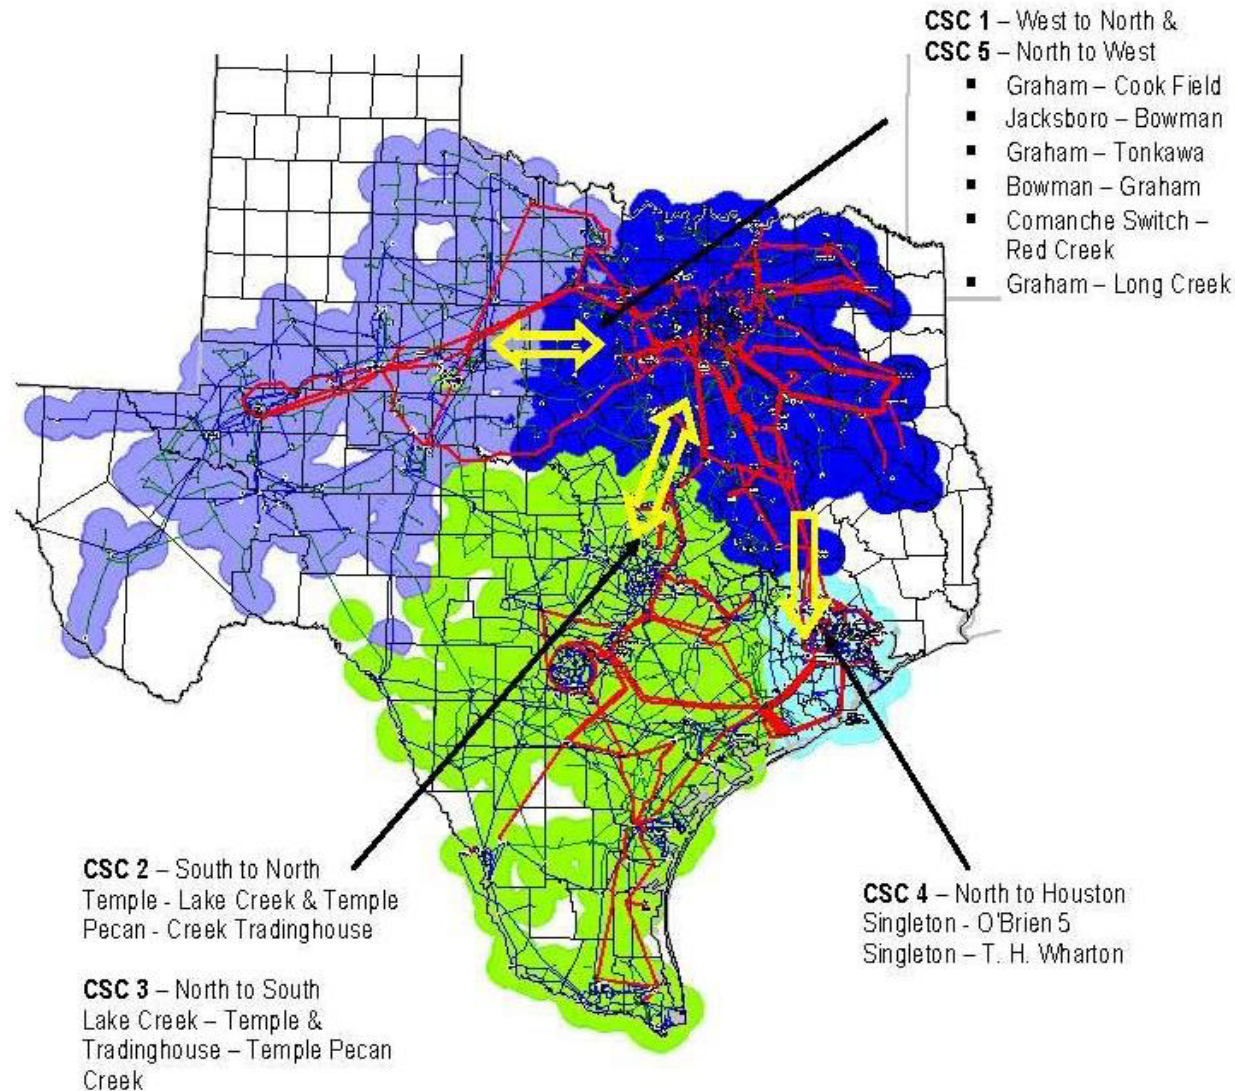
## Locational Marginal Pricing



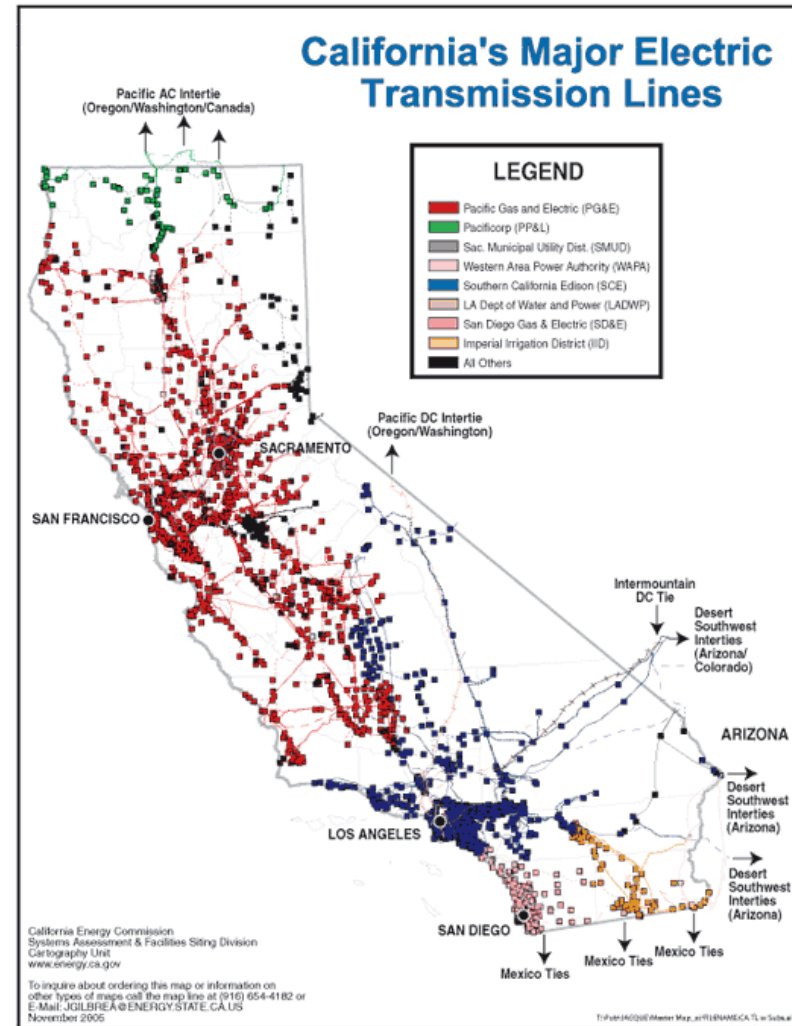
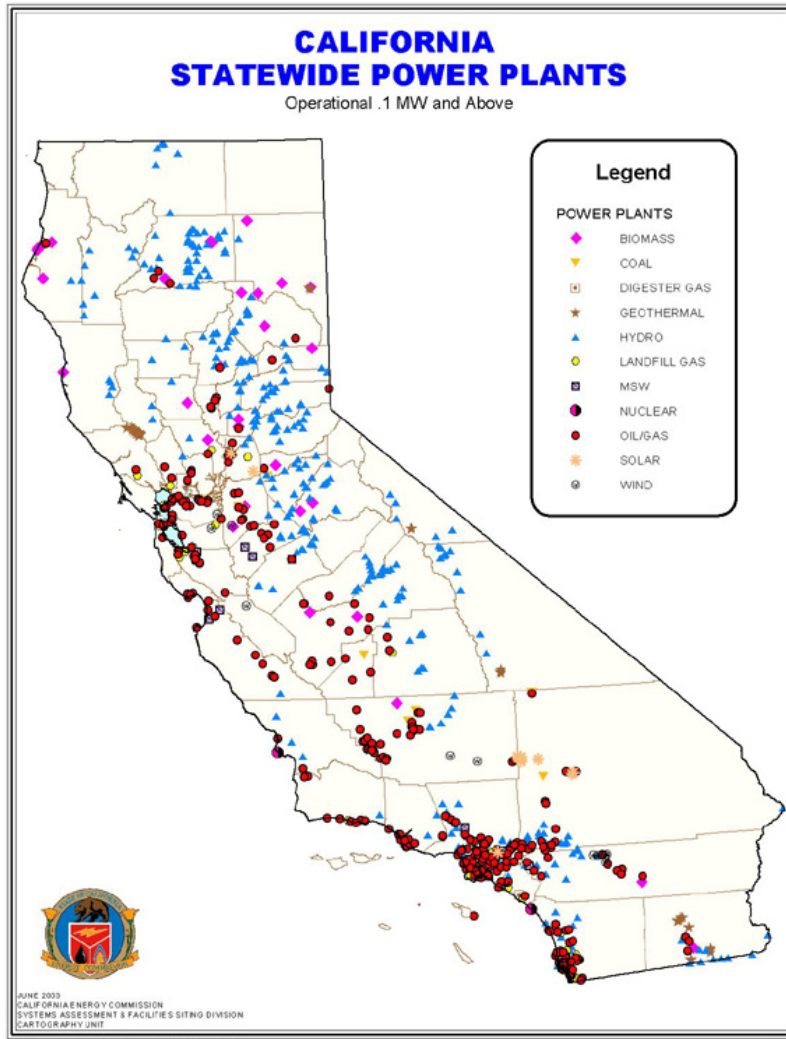
Hover over points to view details.

<http://www.ercot.com/content/cdr/contours/rtmLmp.html>

# Congested Transmission Grid: Price variation



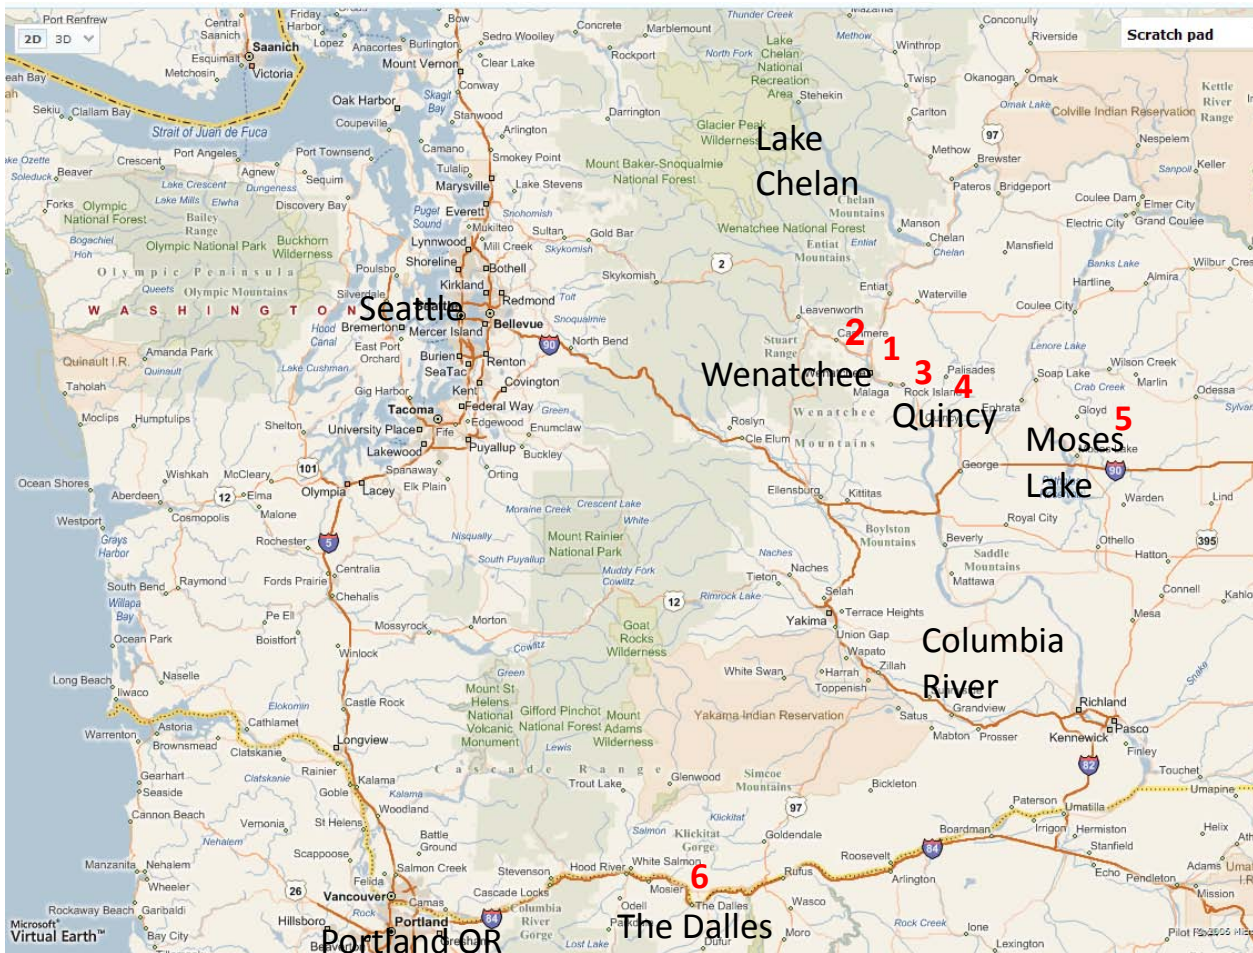
# California Generation, Transmission Interconnects



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



# 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!

# Green Storage - Agenda

- Overview, Definitions, and Motivation
- Systems thinking: TCO, over entire Lifecycle
- How we talk about 'Green Storage' (*good/bad news*)
- 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 most articles 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!

# 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
- CONSOLIDATE (Virtualize)
- What are the equivalents for **Storage**?
  - ◆ TAPE or Optical? (trade-off response time vs. energy usage)
  - ◆ Solid State Storage: high IOPS, low/no power, expensive???
  - ◆ Disk drives and RAID arrays
    - › Slower/Larger drives where possible (Design choice vs. Dynamic)
    - › Power-off or spin-down drives: MAID (Massive Array of Idle Disks)

# Green Storage - Agenda

- Overview, Definitions, and Motivation
- Systems thinking: TCO, over entire Lifecycle
- How we talk about 'Green Storage' (*good/bad news*)
- Electricity Pricing in the United States
- Datacenter Design and Operation
- **Storage Components and Technologies**

# 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

- 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)

# RAID level vs. Power/Cooling

- JBOD: **N**umber 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 = **2xN**
  - ◆ **RAID 0+1 or RAID 1+0, power/cooling=2xN**
- 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**

# 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



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



Check out our other  
**SNIA Green Tutorials!**

[www.snia.org/education/tutorials](http://www.snia.org/education/tutorials)

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

**- SNIA Education Committee**

**SNIA Green Storage Initiative**

**Wayne Adams**

**Rick Bauer**

**LeRoy Budnik**

**Bill Mottram**

**Erik Riedel**

**SW Worth**

**Christian Belady**

**David Black**

**Larry Freeman**

**Deborah Johnson**

**Al Thomason**

**Alan Yoder**