SNIA

Education

Green Storage – The Big Picture

SW Worth, Microsoft
SNIA Green Storage Initiative

SNIA Legal Notice

- 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.

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

Education SNIA

- 'Green' – What does it mean to the IT 'ecosystem'?
  - Reduction of Total Impact on Environment
    - Systems approach - More than just Energy Use!
  - Defined by Gov't. (EPA, EU, Kyoto), Orgs, Vendors, etc.
- How does "Green" <u>differ</u> 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

The Players in Green IT

- I.T. owners / Data Center operators ("Customers")
- Energy Utilities and Regulators
- EPA Energy Star programs, Euro "Code of Conduct"
- SNIA – org expertise on enterprise STORAGE
- Green Grid metrics www.thegreengrid.org
  - 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)**
- Other interested parties (e.g. Uptime Institute)

Motivators: Why we act as we do

- **"Fear"**: Constraints (Physical limits, Regulations)
  - Physical Laws are not optional!
    - E.g. WAN latency (light-speed), Disk rotation speed, Tape Retrieval
  - Gov. Regulations: Do what you are forced to do
- **"Guilt"**: Competitive and 'Moral' aspects
  - Keeping up with industry, responding to <u>non</u>-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 (Green-house gases (GHG), *'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 Carbon (mod *outcomes*, e.g. SOx/NOx)
  - 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)
- 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!
- Unintended Consequences: reduced reliability?

Three Stages of Product Life-Cycle

Education SNIA

- Birth: Product Creation (design for recycle/disposal)
  - Integrated into CapEx *(maybe)* – see WEEE/RoHS
    - *Outsourced 'embedded'Carbon? – see Carnegie Institute*
  - 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….*

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, plus IT Gear: Servers, Storage, Networking,	Variable 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 (peak 15-min. period)	$/kWh x Total usage

Bottom-line:

1. **Power costs may over-whelm Energy Costs!**
2. **Big data centers charge-back on Power, Energy, or both.**

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 floor-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 (bad) visibility?

DCiE= (IT Energy Demand) / (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%)

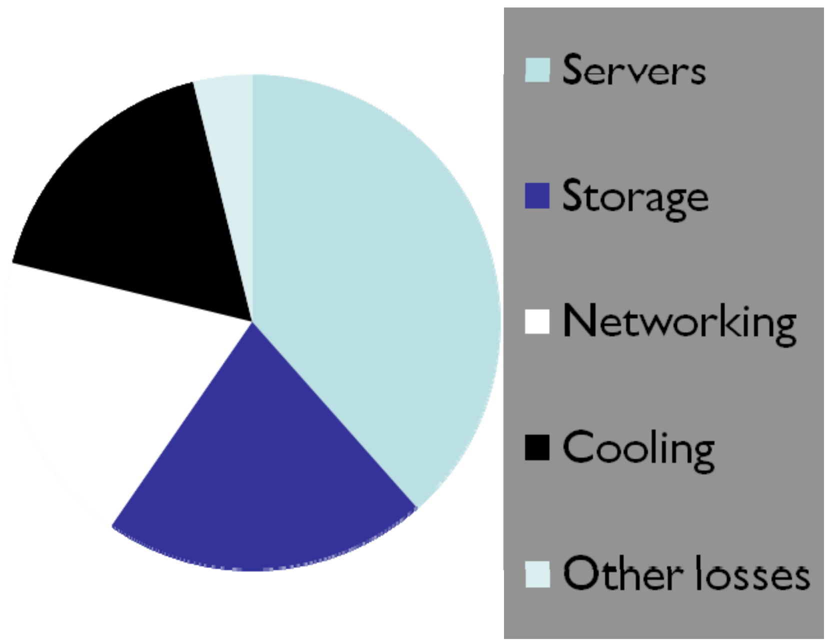

Server percentage also declines, due to Consolidation and Virtualization.

What Affects Storage Energy Use?

Redundancy and RAID Definitions

Standalone

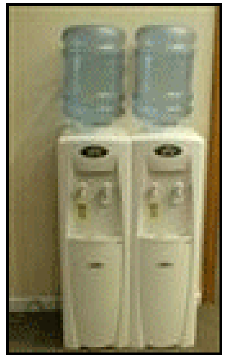

Cluster

Hot swap

RAID 0

RAID 1

RAID 5

RAID 0+1

SNIA Storage Taxonomy

Education SNIA

Storage Taxonomy Summary	Online Storage	Near Online Storage	Removable Media Libraries	Virtual Media Libraries	Infrastructure Appliances	Infrastructure Interconnect
	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 1st data access.	Archival storage used in a sequential access mode. A Typical example would be Tape based archival, both Stand Along 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.
Maximum Capacity Guidance 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 opposed 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	Max Storage Devices	Max Storage Devices	Max Tape Drives		Max Storage Devices Supported*	Max Port Count
Group 1) SoHo & Consumer 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)	Up to 4 Devices		Stand Alone Drive (No Robotics)		Note: * Infrastructure Appliances by definition have no intrinsic storage, other than what is used for local processing and/or local Cashing of data. Storage Devices Support in this case refers to the number of storage devices controllable down stream of the Appliance	
Group 2) Entry, DAS, or JBOD 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	More than 4 Devices	Up to 4 Devices	Up to 4 Drives			Up to 32
Group 3) Entry / Midrange 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	More than 20 Devices	More than 4 Devices	More than 4 Drives	Up to 100 Devices	Support for up to 20 Devices	Up to 128
Group 4) Midrange / Enterprise 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)	More than 100 Devices	More than 100 Devices	More than 24 Drives	More than 100 Devices	Support for more than 20 Devices	More than 128
Group 5) Enterprise / Mainframe 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	More than 1000 Devices		More than 11 Drives	More than 100 Devices	Support for more than 100 Devices	© SNIA 2009

echnologies fo

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

Need for redundancy

Education SNIA

- 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

- Energy consumption is roughly linear in the number of naïve (full) copies

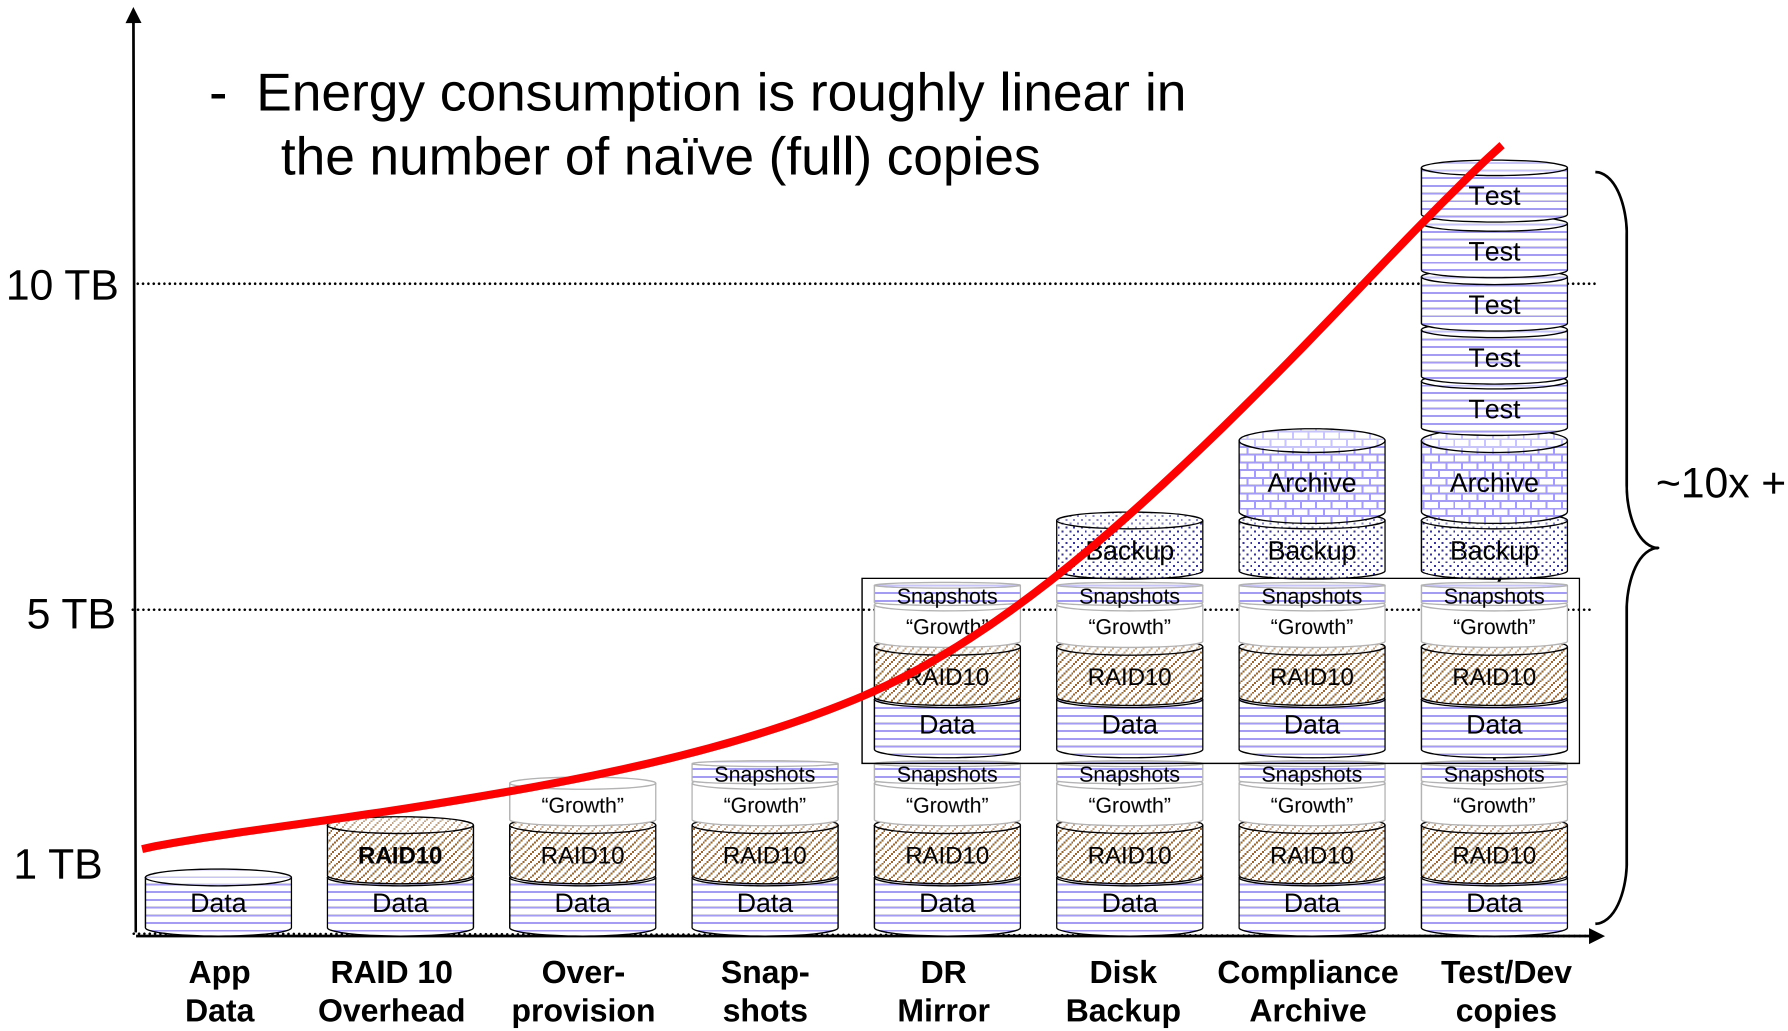

Effect of green technologies

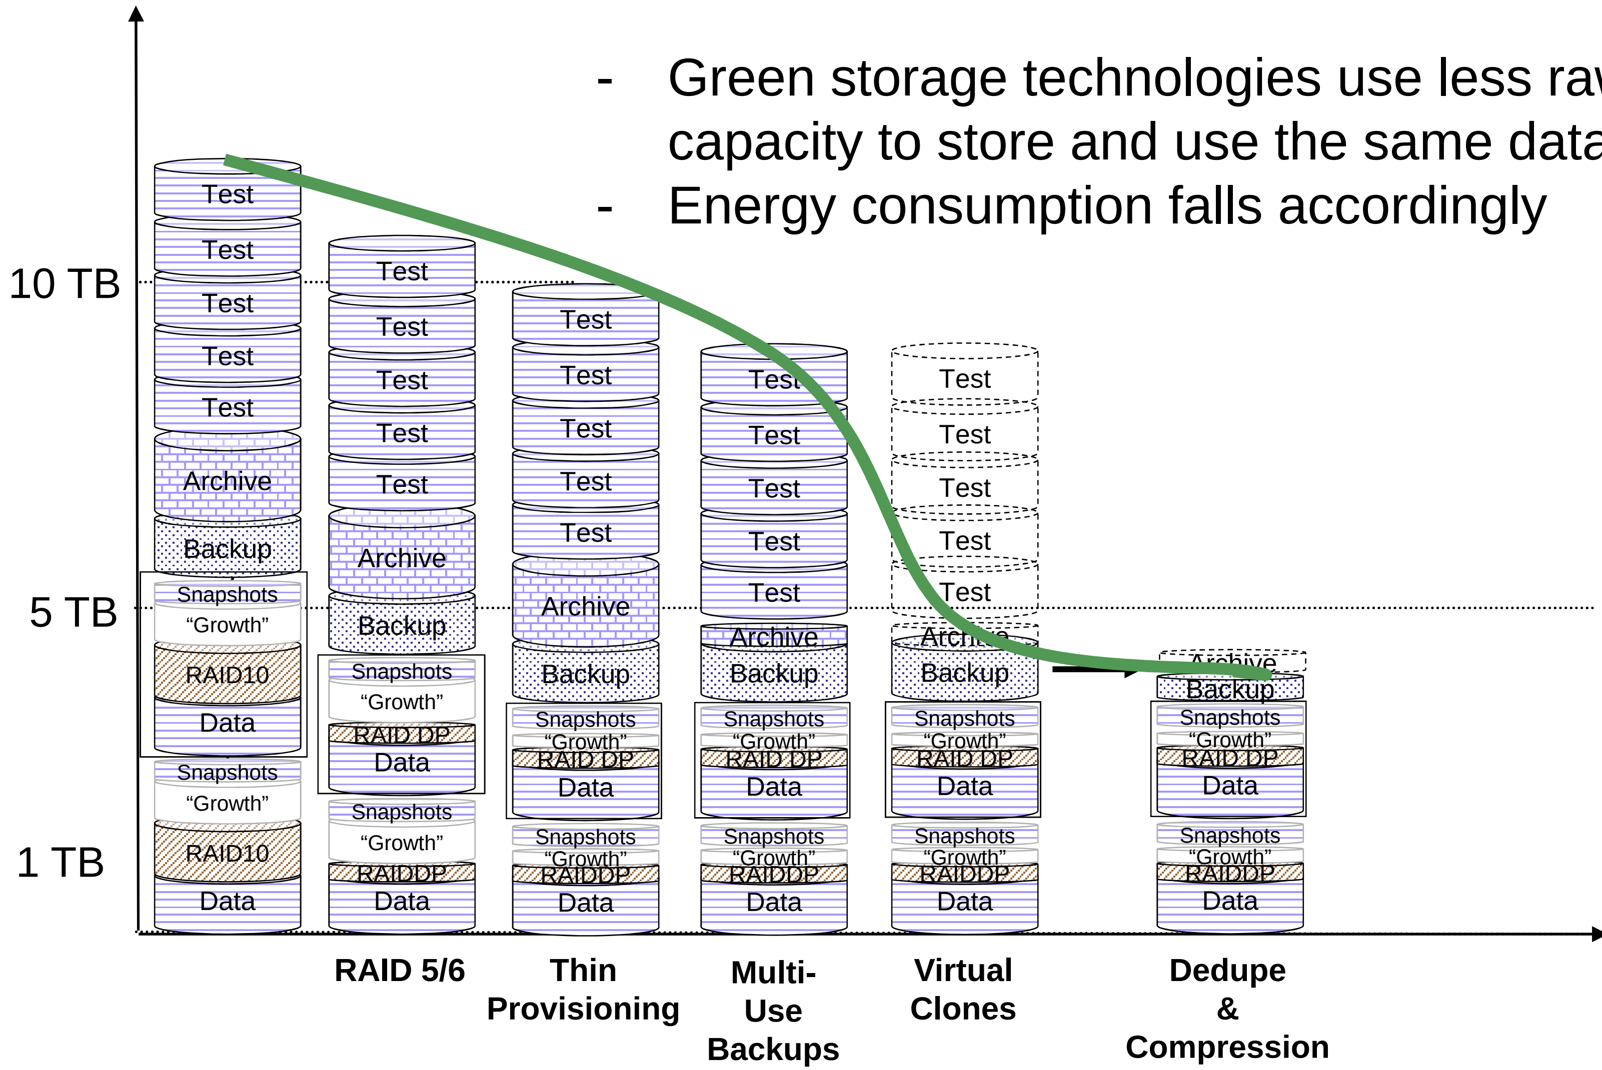

- Green storage technologies use less raw capacity to store and use the same data set
- Energy consumption falls accordingly

What Storage aspects could be affected?

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

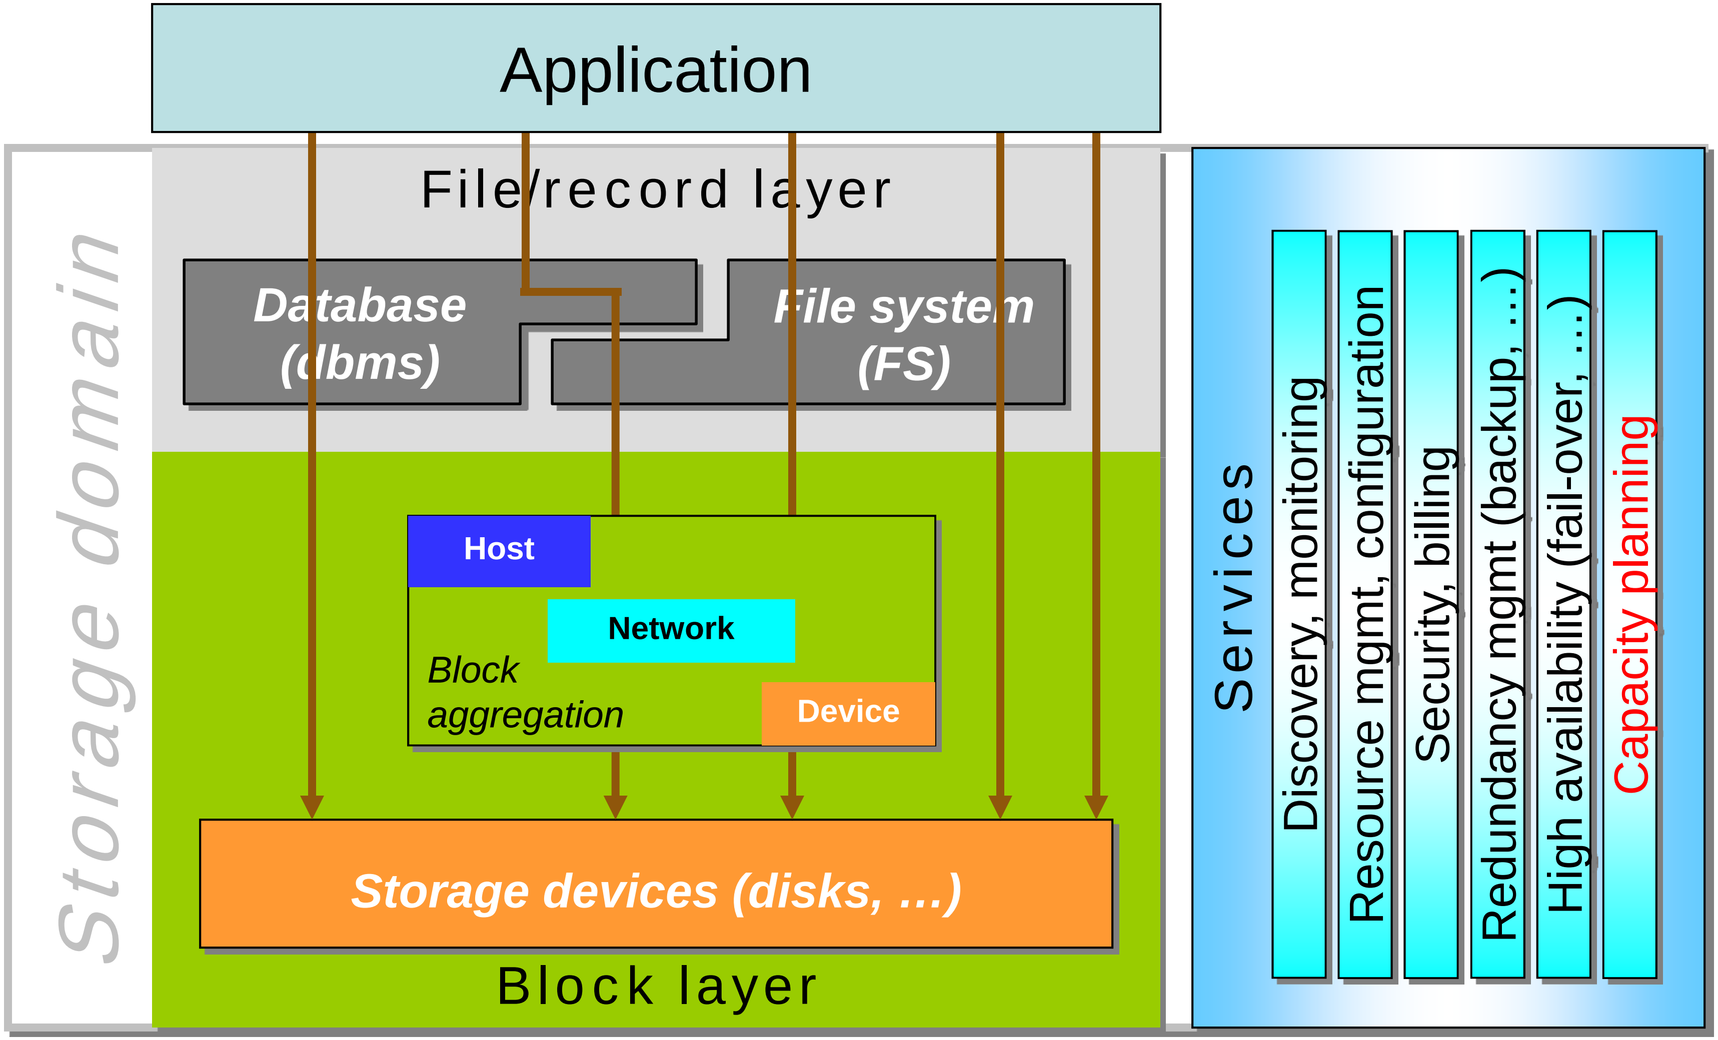

Green Storage - Agenda

- 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
- WEEE
- RoHS, China-RoHS
- Packaging and Pkg Waste
- Halogens (in plastics)
- Basel Convention/Basel Ban (Transboundary Wastes)

WEEE: 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!

	%
Large household appliances	69
Small household appliances	8
IT & telecommunications equipment	7
Consumer equipment	13
Lighting equipment	0
Electrical and electronic tools	2
Toys leisure and sports equipment	<1
Monitoring and control instruments	<1
Automatic dispensers	0

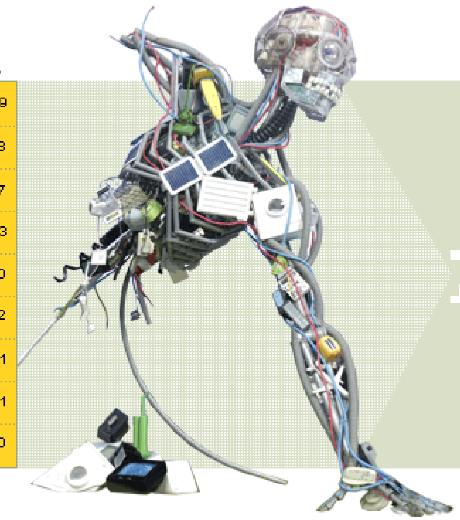

1 UK citizen per life time = 3.3 tonnes

RoHS:
Restriction of Hazardous Substances

Education SNIA

- 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 <u>component</u> level!)*
  - Mercury - Cadmium - PBB
  - Chromium VI - Lead - PBDE

RoHS 2002/95/EC

- Unintended Consequences: reduced reliability?
  - EPA report (Aug-05) on lead-free solder!
  - RoHS <u>exemption</u>: 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?

- 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

European Commission Code of Conduct on Data Centres Energy Efficiency

COMMISSION EUROPÉENNE

EUROPEAN COMMISSION
DIRECTORATE-GENERAL JRC
JOINT RESEARCH CENTRE
Institute for Energy
Renewable Energies Unit

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

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

Education SNIA

- 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)

Education SNIA

- **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)

Electric Power Generation

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

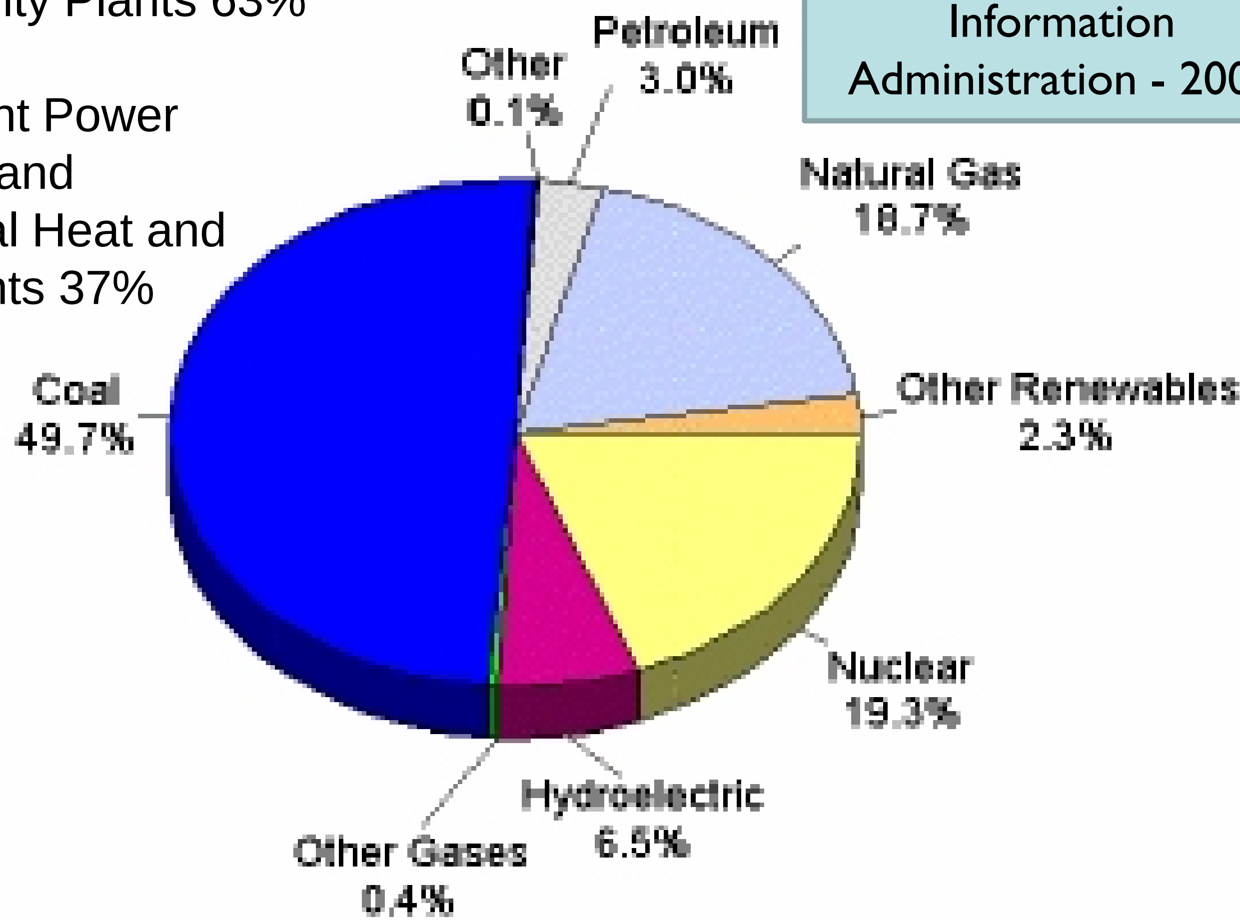

Calif. Generation; Transmission Interconnects

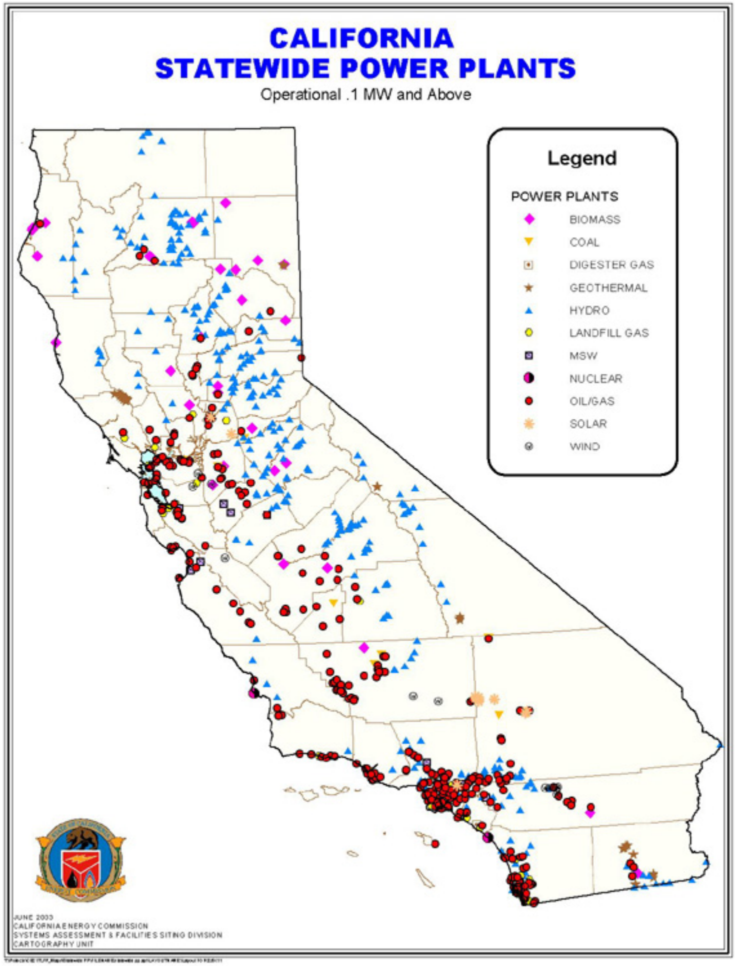

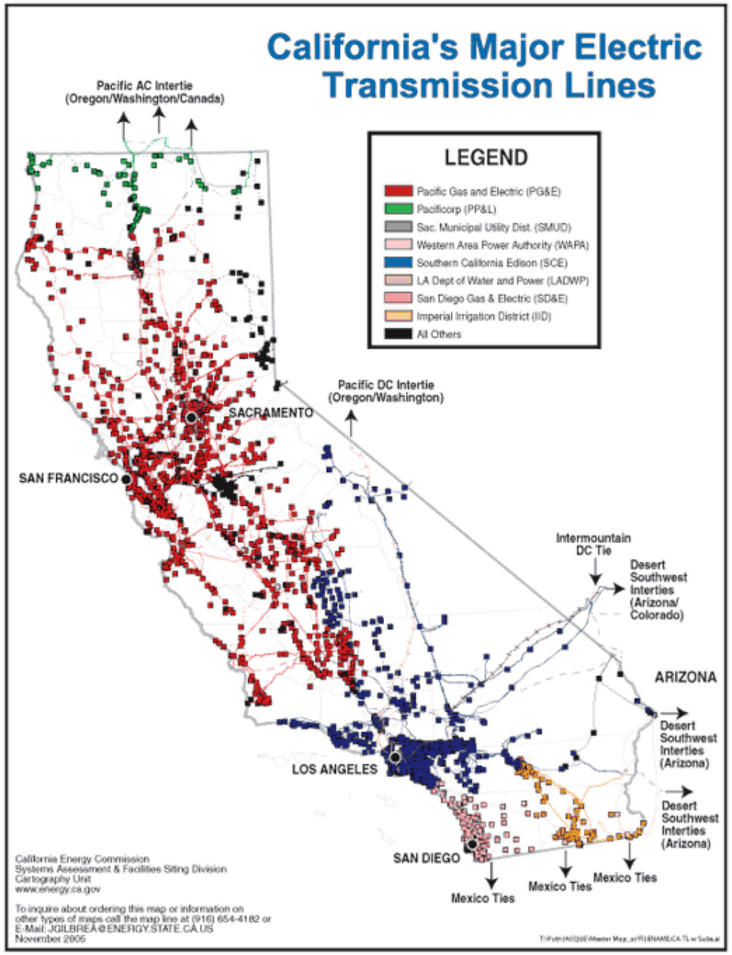

www.energy.ca.gov/maps/

State Electricity Prices, 2005 (cents/kWh – "Industrial")

U.S. Average 5.73

10 Most 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

10 Least Expensive States

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?

Education SNIA

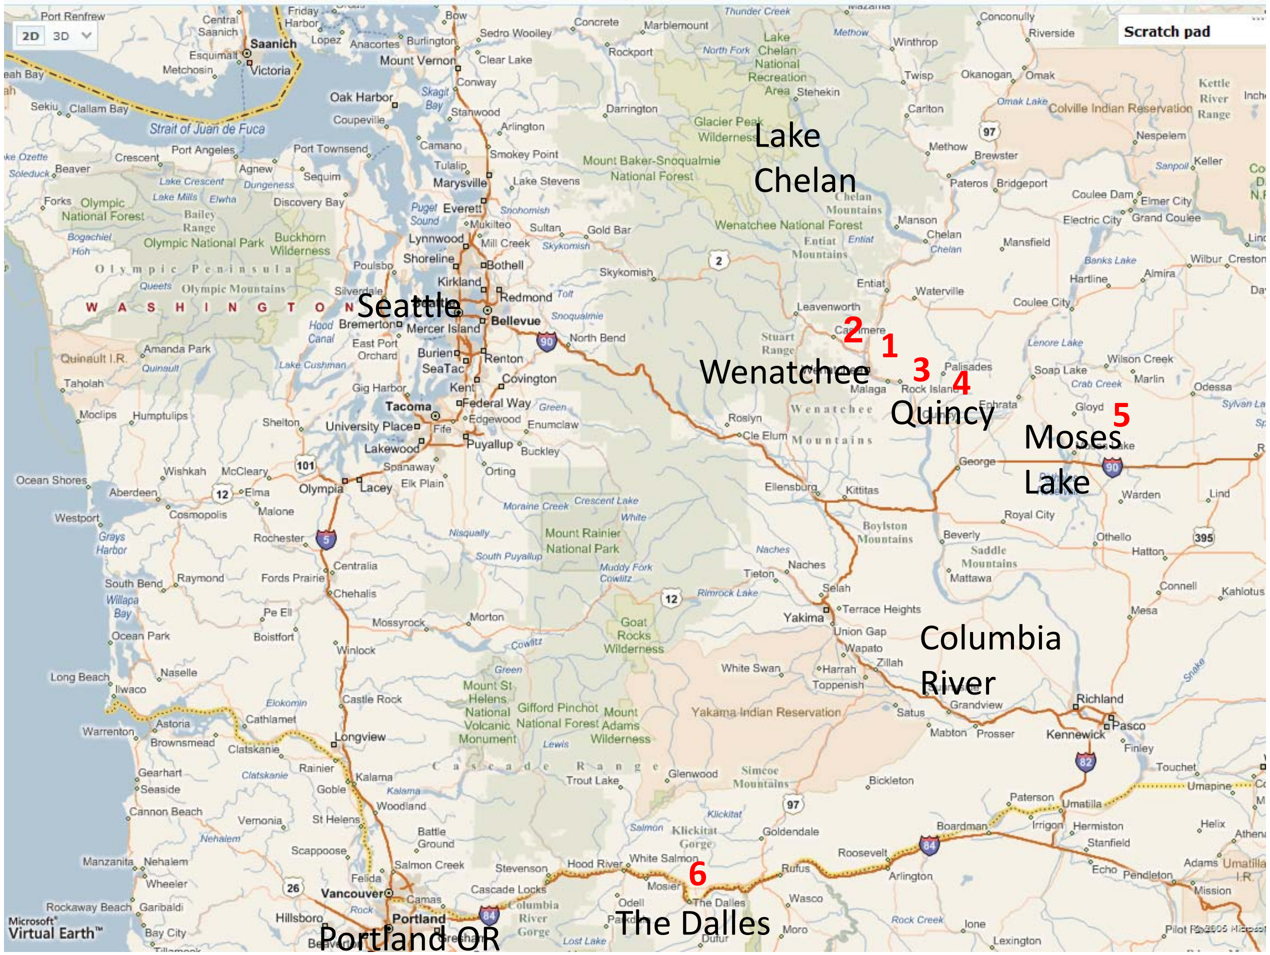

Energy costs on the Columbia River are about **$0.02/kWh** for Datacenters.

Ample fiber (WAN) bandwidth is available (www.noanet.net)

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

Result: Construction!

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 <u>external</u> Power availability
  - Oversized (IT gear plus infrastructure) = excess CapEx
    - Under-utilization = Power inefficiencies (poor OpEx)

Facilities vs. I.T. in the Datacenter

Education SNIA

- 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)
- 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!

Model or Measure: Which is Better?

Education SNIA

- 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 <u>and</u> Peak)
- Measurement issues (Reality validates Modeling)
  - Actual *in-situ* workloads (“normal” <u>and</u> 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

Education 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

Education SNIA

- 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

Education SNIA

- Each component of a Storage system has Power and Cooling requirements
  - Understand "**Idle**" (**not** 'stand-by') *vs.* "Loaded" (R/W)
  - Label ratings are usually <u>peak</u> 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…

Sample Idle Storage System Measure Group 3 - Online Block Storage Idle Phase of Idle Test Measurement

Education SNIA

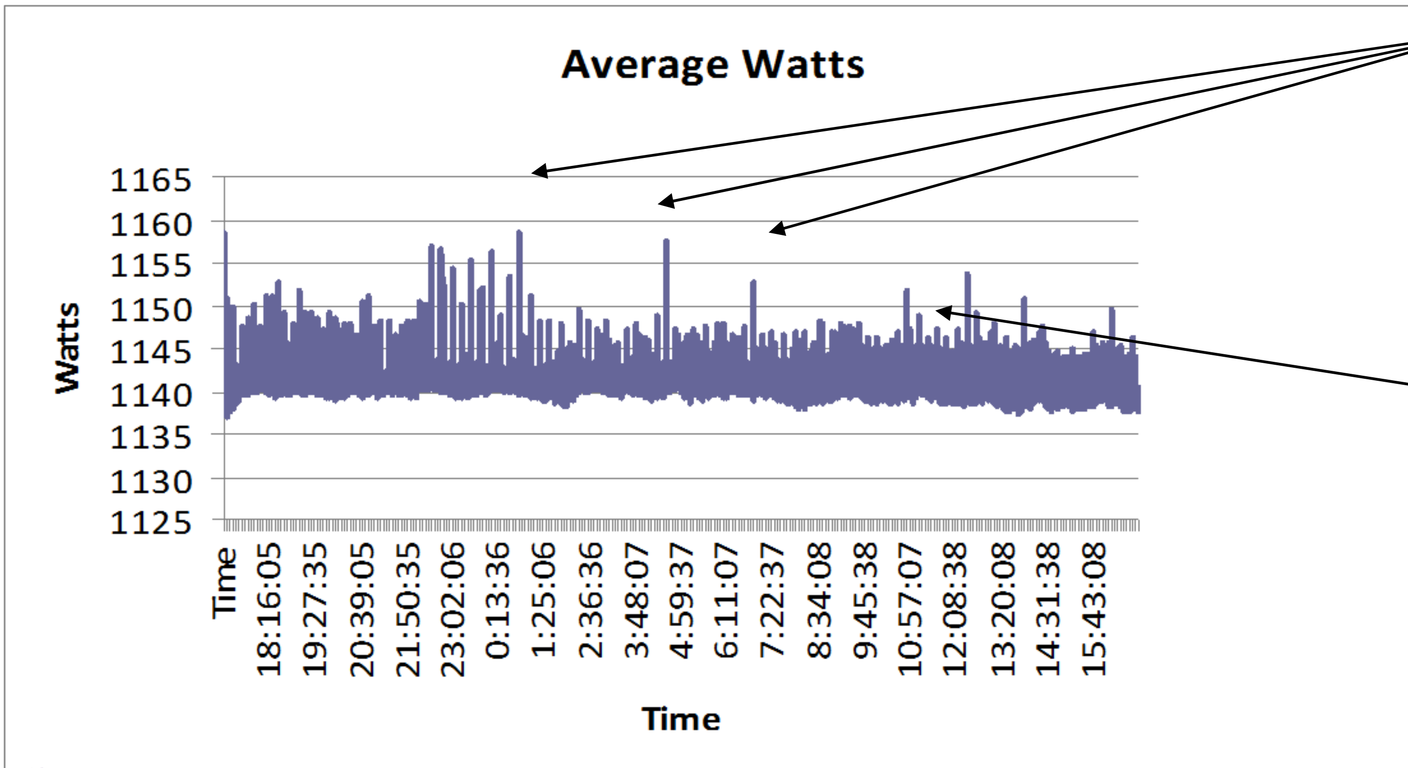

Increased energy use for disk's continuous background diagnostics, system house-keeping, and if needed incremental cooling

Baseline energy use for disks spinning, controllers, power supplies, lowest level temperature management, etc

- 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.**

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?
RAID Definitions


Standalone

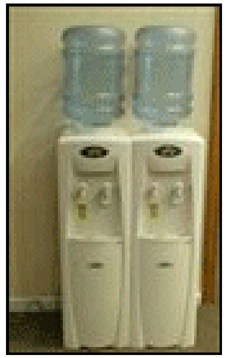
Cluster


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 x 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

Education SNIA

- 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 of 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?

Increase 'effective' Data Density

- 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 <u>effective</u> 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

Are savings multiplicative?

Education SNIA

- 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 + resizeable 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)		✓			✓	✓	

Equipment power savings

- Server virtualization
  - up to 80% savings
  - much depends on load
- Power supply efficiencies
  - 10 – 20%
  - difficult to do anything about this in a vacuum
  - probably okay to just ride industry trends
- Variable speed fans
  - up to 80% savings
  - power consumption cubic in rotational speed
  - interesting interaction with data center cold aisle temperatures

Facilities savings

- State of the art data centers
  - PUE* drops from 2.25 to 1.25 = 45% savings
    - 10MW → 5.5MW
    - $6.0M → $3.3M annually
  - Rebates in the $M from utilities on top of savings

* Power Utilization Efficiency – see www.thegreengrid.org/gg_content

Final point - continued monitoring

- Essential to continued success
- Adjust-and-rebalance scenarios
  - E.g. economizer optimization
- Overall system health
  - SLA verification etc.

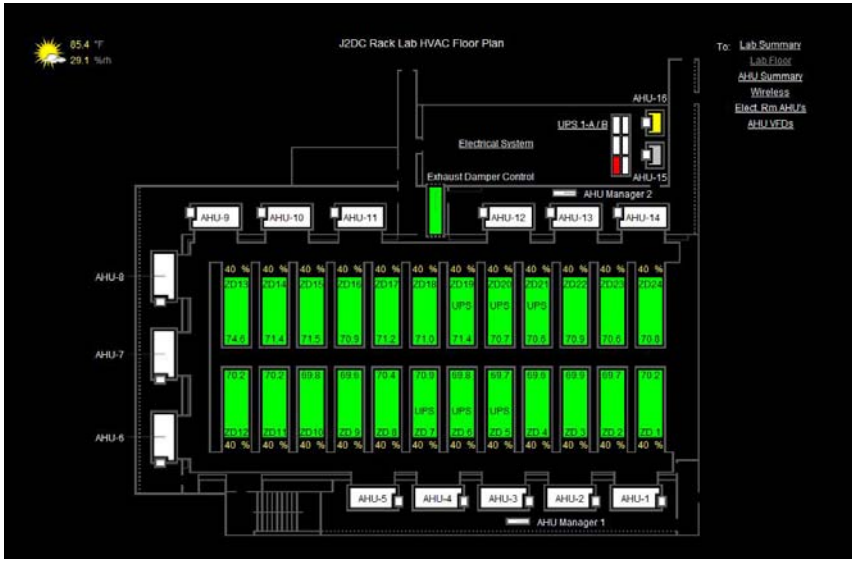

Q&A / Feedback

- Please send any questions or comments on this presentation to *trackgreenstorage@snia.org*

Check out our other SNIA Green 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