



Education

GREEN STORAGE II: Metrics and Measurement

Patrick B. Chu, Ph.D., Seagate Technology
Erik Riedel, Ph.D., Seagate Technology

- 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 II: Metrics and Measurement**

- ◆ This tutorial will cover storage-specific topics related to energy-efficiency and outline the current state of the industry. We will discuss a range of technologies that are currently considered “green storage”. We will discuss metrics for measuring, managing and designing for power in storage systems. We will also briefly outline ongoing efforts in the SNIA Green Storage TWG and in partnership with other industry groups.

Tutorial Overview

- storage-specific topics related to energy consumption and efficiency
 - ◆ current state of the industry
- comparing the range of technologies that are currently considered “green storage”
 - ◆ metrics for measuring, managing & designing for power
- ongoing efforts in the SNIA
 - ◆ Green Storage Technical Work Group (TWG)
 - ◆ partnerships with The Green Grid, the DMTF and others
 - ◆ publications and workshops

Background and Objectives

➤ Background

- ◆ basic storage concepts
- ◆ energy & engineering topics covered in “Green Storage I”

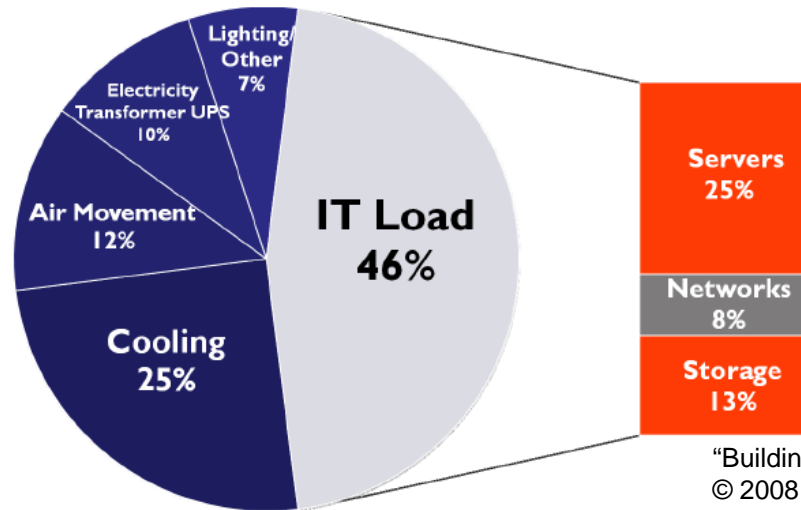
➤ Objectives

- ◆ understand the storage-specific aspects of “green” (especially energy and power) for storage components and aggregated products (e.g. arrays, controllers, appliances)
- ◆ outline considerations in how to measure and model storage energy and power
- ◆ understand the technologies to reduce the ongoing energy costs of reliably maintaining data

Outline

- Current state of affairs and industry trends
- Power measurement
 - ◆ storage subsystems
 - ◆ idle and active modes
 - ◆ power supply loading / efficiencies
 - ◆ power measurement / monitoring equipment
- Green metrics and taxonomy
 - ◆ measures of green-ness
 - ◆ storage product categories
- Technologies for efficiency /power reduction
 - ◆ hard drives, solid state, tape
 - ◆ deduplication, thin provisioning, MAID
- SNIA green storage efforts
 - ◆ “plug” fests, green standards, whitepapers / workshops, alliances

Energy Impact of Storage Systems

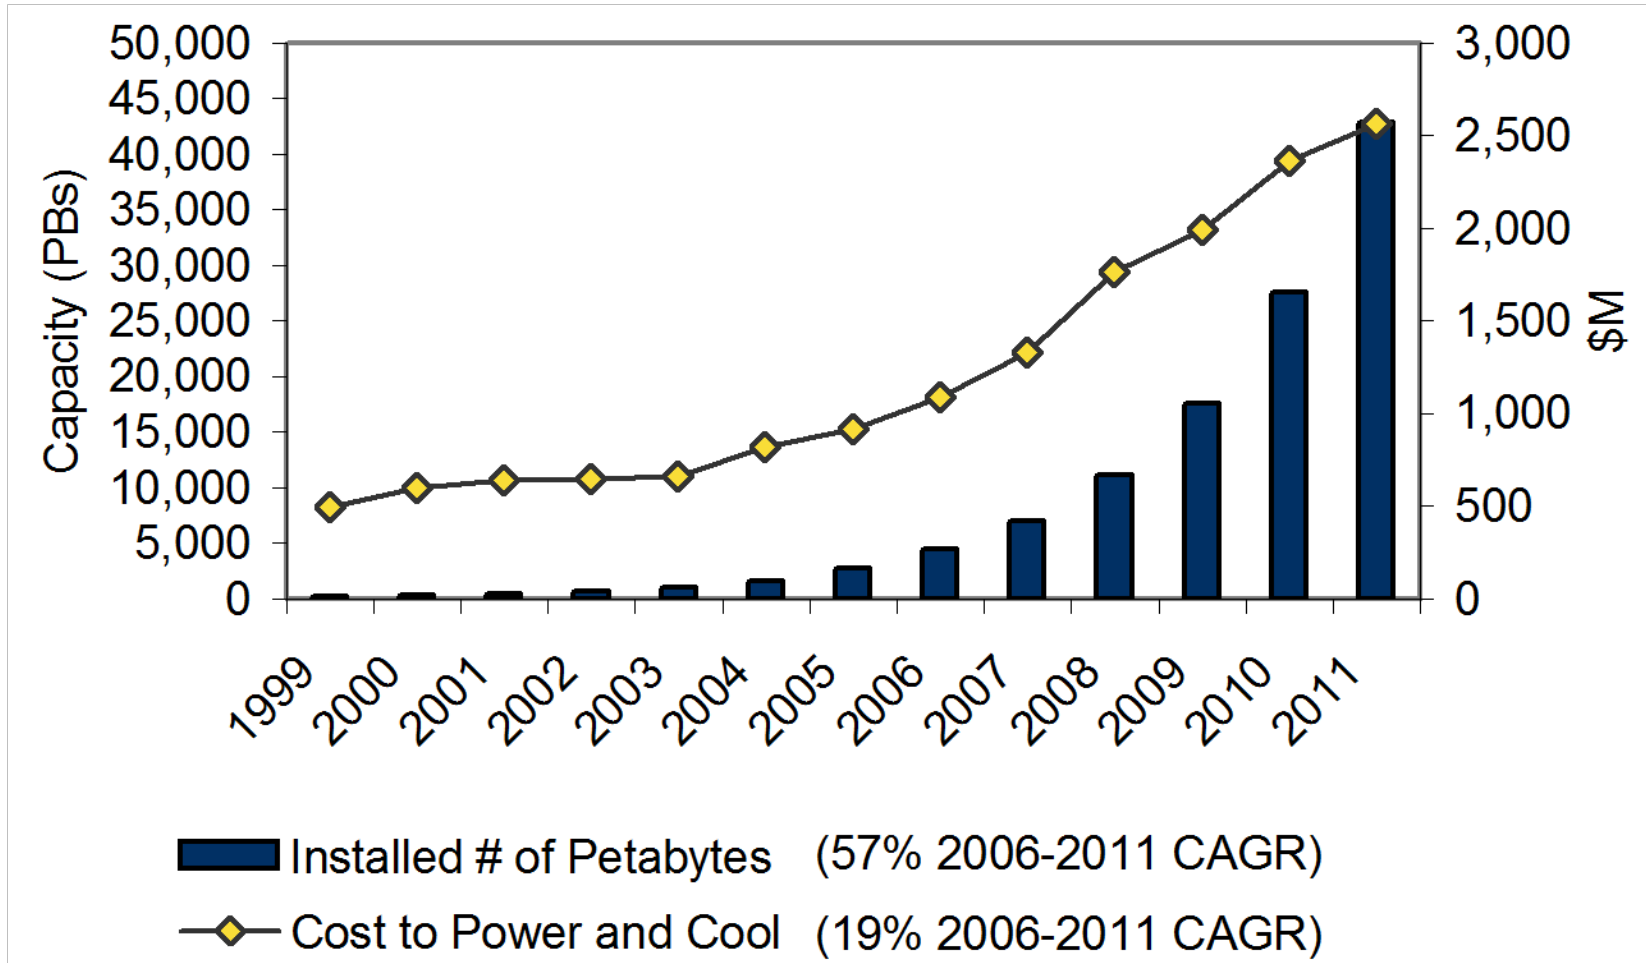


“Building the Green Data Center”
 © 2008 SNIA All Rights Reserved

- Storage – significant in power consumption
 - ◆ >13% of total data center load, plus contribution to cooling needs
- Rapid increase in storage needs
 - ◆ 6-fold increase in digital content from 2007 to 2011 (281 exabyte to 1,773 exabytes)
 - ◆ By 2010, 70% of “digital universe” created by individuals but organizations to be responsible / liable for 85%

IDC White Paper, “The Diverse and Exploding Digital Universe,” sponsored by EMC – March 2008.

Cost of Data Storage



IDC #212714, "The Real Costs to Power and Cool All the World's External Storage" – June 2008 Dave Reinsel

What impacts power consumption

- Storage capacity / usage efficiency
 - ◆ increasing data → larger capacity → more disks
 - ◆ redundant copies → magnify capacity needs
 - ◆ variability in usage and utilization → inefficient allocation of space
 - ◆ What is valuable data? What is the retention policy?
- Data transfer rate / access speed
 - ◆ high I/O bandwidth → higher rotational speed; striping across many drives
 - ◆ low access times → faster actuators; higher rotational speeds; caches
 - ◆ How fast and immediate must data be available? (time-to-data)
- Data integrity
 - ◆ 25% of “digital universe” is unique, but 75% are replications / duplications
 - ◆ partly to ensure data integrity and survivability; partly wasteful
- Data availability / system reliability
 - ◆ RAID uses extra drives, plus redundant power supplies, fans, controllers,
 - ◆ How valuable is data? How likely are failures? How fast must data be available?

Potential paths to “green”

➤ Improve usage efficiency

- ◆ Deduplication
- ◆ Thin provisioning

must be driven by
metrics / standards
/ guidelines

➤ Minimize energy consumption

- ◆ Improved component designs – high-efficiency power supplies, advanced & flexible drives
- ◆ MAID – idle and spin-down

➤ New technologies

- ◆ Solid state storage
- ◆ Alternative storage architectures

Anatomy of a Storage System

Switches

Appliances


Disk Arrays

PDU's

Power Distribution Unit

UPS's

Uninterruptible Power Supply



- System design, complexity and redundancy vary depending on applications & usage
- Component designs, software features, and workload affect power consumption and efficiency



Apps

Software



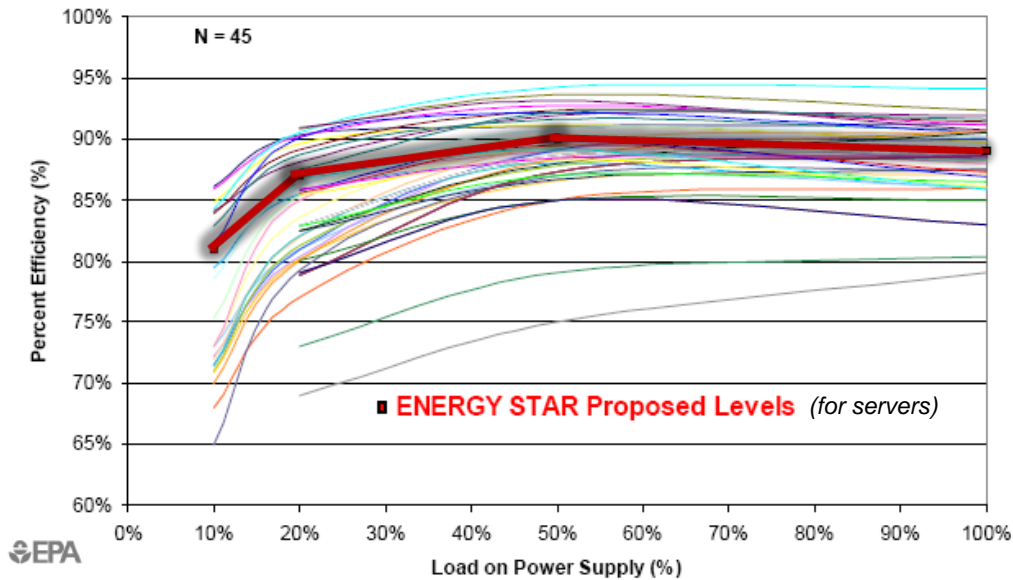
Power Supplies

Fans

Controllers

Hard drives

Power Supply Efficiency



Data from 45 single-voltage PSU from 8 vendors (products available 1st Q 2009) presented by EPA at ENERGY STAR Computer Server Stakeholder Meetings in Redmond, WA; July 2008

- Using properly sized efficient power supplies benefit system in Idle & Active
- Design must take into consideration
 - ◆ peak load, average loads and redundant system configuration
 - ◆ supplies may operate at <50% of max. load capacity most of the time
- Voltage levels (110V, 218V, 220V, 240V) also impact efficiency

Idle Power vs. Active Power

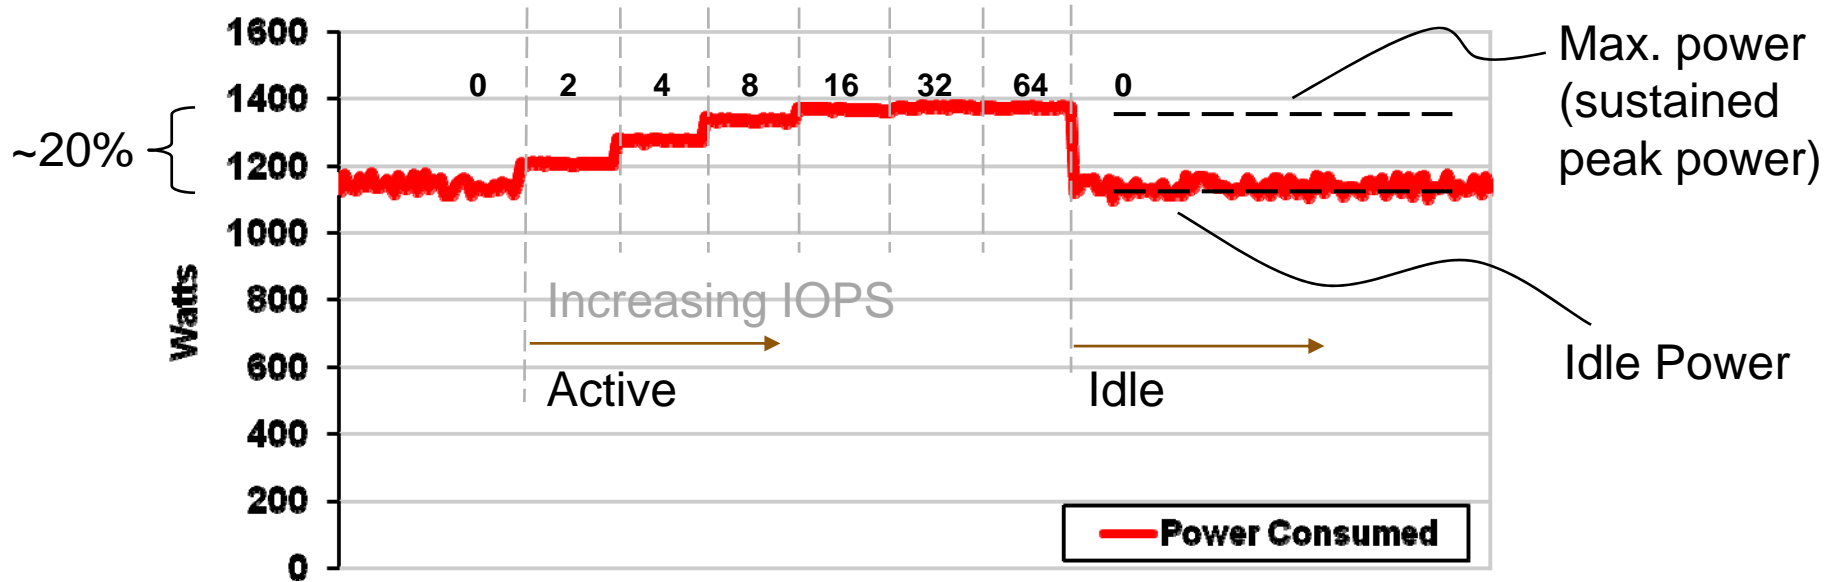
➤ Idle Mode

- ◆ storage system is protecting data, ready to process IOs
- ◆ background maintenance & optimization tasks on-going
- ◆ factors: time-to-data, overhead electronics, fan, maintenance
- ◆ systems are idle large fractions of the time

➤ Active Mode

- ◆ storage system is carrying out IOs
- ◆ background tasks continue in parallel
- ◆ factors: workload (seq/random), response time, throughput
- ◆ evaluate a variety of workloads, plus sustained peak power

Example of Power Measurement



- Ideally, systems consume minimum power in all modes
 - ◆ Example system consumes **significant power in idle (80% of max)**
- % of time in Idle versus Active depends on storage type, application and workloads; available optimizations will vary
- Power consumed is not linearly proportional to workload (indicates potential room for improvement)

Measurement Tools

- Variety of power monitoring / measurement tools available
 - ◆ rack-mounted, networked PDUs for continuous operational monitoring
 - ◆ more accurate power meters w/ data logging capabilities are preferred for system characterization and benchmarking
 - ◆ select a tool based on accuracy, features, Amp/V/Watt levels
 - ◆ refer to SPEC website for recommended measurement devices and settings
www.spec.org/power_ssj2008/device-list.html
- Measure oper. conditions (temp., humidity, altitude) w/ power to establish baselines and understand system behaviors
- Both total and sub-system power consumed are valuable info.



Low-current / voltage power meter



Networked, instrumented rack-mounted PDU



Power meter with data logging and 0.1% accuracy

Green Metrics – purpose / challenge

- Need scientific measures and common vocabulary to assess “green” performance
 - ◆ assist in data center design, operational monitoring/tuning and regulatory compliance
- Storage systems have many modes and outputs
 - ◆ a single metric (such as GB/Watt) may not reflect the characteristics or capability of the whole system
 - ◆ some system “outputs” are not all easily quantifiable
- Multiple metrics may be weighed and combined to form a single metric (e.g. annual energy bill)
 - ◆ will vary with usage and system; your mileage will vary

Metrics Examples



Workload considerations

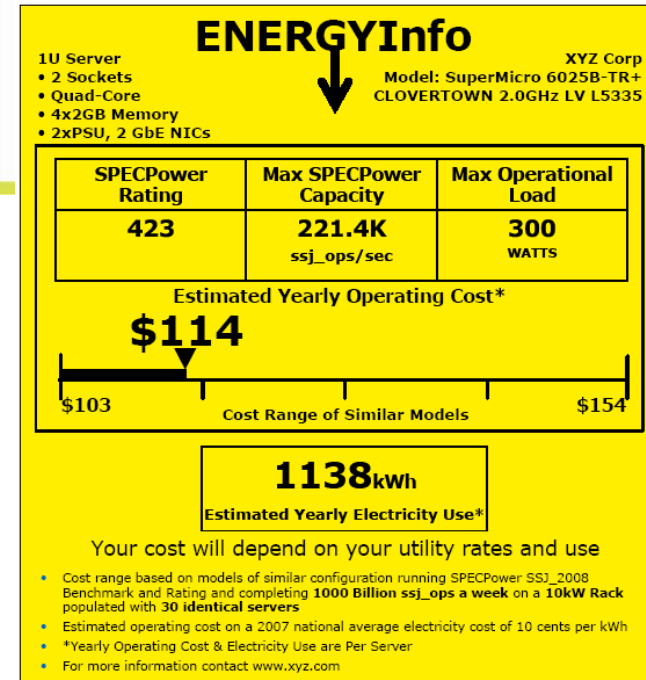
- ◆ Data at rest – Idle power (GB/Watt)
- ◆ Data on the move – Throughput (MB/s)
- ◆ Data at work – Performance (IOPS)

Potentially useful metrics

- ◆ GB per Watt; MB/s per Watt; IOPS per Watt
- ◆ Power supply efficiency; CO₂ footprint
- ◆ Total annual energy bill

Reliability / availability / serviceability considerations

- ◆ Latency (time-to-date)
- ◆ Redundancy level (RAID efficiency, resilience to failures)



Server example

Storage Taxonomy

- Need a taxonomy (or product classification) to enable fair comparisons among similar storage products
 - ◆ e.g. for motor vehicles – motorcycles, cars, trucks
- Similar green metrics may apply to all product categories, but different values establish best-in-class
- Unique considerations apply to special categories
 - ◆ e.g. amphibious cars, skid steer loaders, tanks
- Clear taxonomy will simplify comparisons and aid regulatory efforts

Taxonomy Considerations

- Distinction based on response time (mean time-to-data)
 - ◆ Online < xx ms
 - ◆ Near online > xx ms but < yy min
 - ◆ Offline > yy min

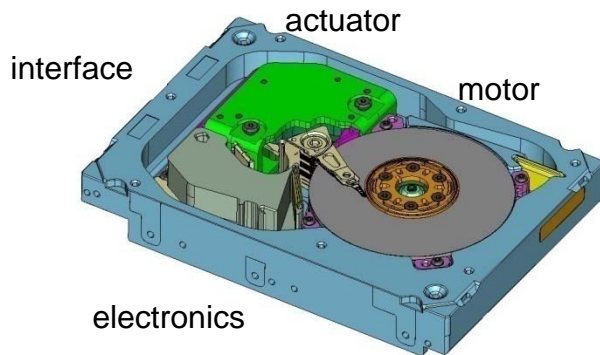
- Distinction based on throughput (IOPS)
 - ◆ Performance > zz k IOPS
 - ◆ Capacity (raw vs. user-visible)

- Distinction based on usage categories
 - ◆ SoHo (small office / home office)
 - ◆ DAS (direct attached storage)
 - ◆ Entry, Mid-range, High-end

Outline

- Current state of affairs and trend of industry
- Power measurement
 - ◆ storage subsystems
 - ◆ idle and active modes
 - ◆ power supply loading / efficiencies
 - ◆ power measurement / monitoring equipment
- Green metrics and taxonomy
 - ◆ measure of green-ness
 - ◆ storage product categories
- Technologies for efficiency & power reduction
 - ◆ hard drives, solid state, tape
 - ◆ deduplication, thin provisioning, MAID
- SNIA green storage efforts
 - ◆ “plug” fests, green standards, whitepapers / workshops, alliances

Hard Disk Drives



7200 RPM
 high capacity



3.5 inch HDD

15K RPM
 high performance



2.5 inch HDD

5400 RPM
 laptop

➤ Besides component improvements, power reduction achievable if you

- ◆ Select large capacity drives
 - › high GB/spindle or GB/drive → high GB/Watt
- ◆ Employ small form-factor (SFF, 2.5”) drives
 - › low-energy spindle & fast seek → high IOPS/Watt
- ◆ Utilize drive power management features (idle and stand-by modes)
 - › power saving possible without major impact on performance or time-to-data

➤ Component improvements

- ◆ Efficient motors, workload-optimized actuators, flexible electronics

Solid State Drives

➤ Definition

- ◆ SSD is “a semiconductor-based block storage device that behaves as a virtual HDD and appears to the host device as a disk drive”

IDC, Worldwide Solid State Drive 2007-2011 Forecast and Analysis: Finding Space in the Expanding Digital Universe, Doc # 207739, July 2007.

- ◆ based on a variety of possible technologies – NAND flash (SLC or MLC); DDR RAM + battery; future Fe / M / R / PC RAM

➤ Characteristics (flash SSD compared to HDD)

- ◆ fast read/write (~15 ns), low latency for random IO
- ◆ good \$\$/IOPS but high \$\$/GB
- ◆ low power consumption, unproven reliability

➤ Opportunities for energy efficiency

- ◆ target high IOPS access for mixed systems, optimizing workload
- ◆ provide an additional power-efficient layer of caching

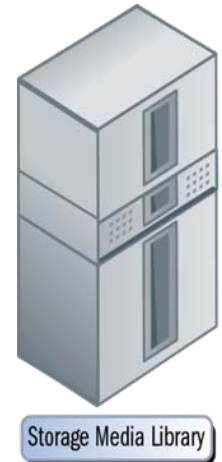


Check out SNIA tutorial:

Solid State Storage for the Enterprise by W. Hutsell, B. McKean

Tapes

- Proper utilization of tiered storage can reduce power consumption (OpEx) as well as purchases (CapEx)
- Clear benefit of tape is low to zero power consumed while data is stored in high volumetric density
- Sequential technology – not suitable for fast, random access, low time-to-data
- Great for secondary backup; tradeoffs for primary backup between HDD (archive storage, VTL) and tape
- The number of tape drives and use of robotics increases power, but enables very large libraries

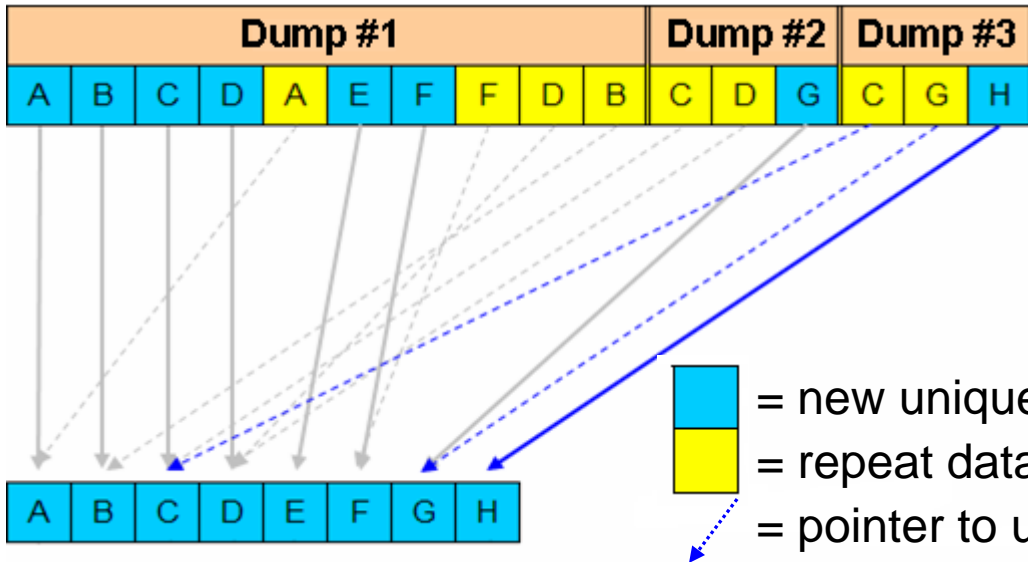


Check out SNIA Tutorial:

Introduction to Data Protection: Backup to Tape, Disk, and Beyond by J. lehl



Deduplication



- evaluate data
- identify redundancy
- create or update reference info
- store or transmit unique data once
- recall (read) or reproduce data

= new unique data
 = repeat data
 = pointer to unique data segment

- Store (& transmit) new and unique data only (at sub-file level); update tracking information; further compression is possible
- Applicable to backup, some primary storage, disaster recovery, archive and wide area optimizations
- Must consider restore speeds (latency, bandwidth, data type); results will be application and data type dependent



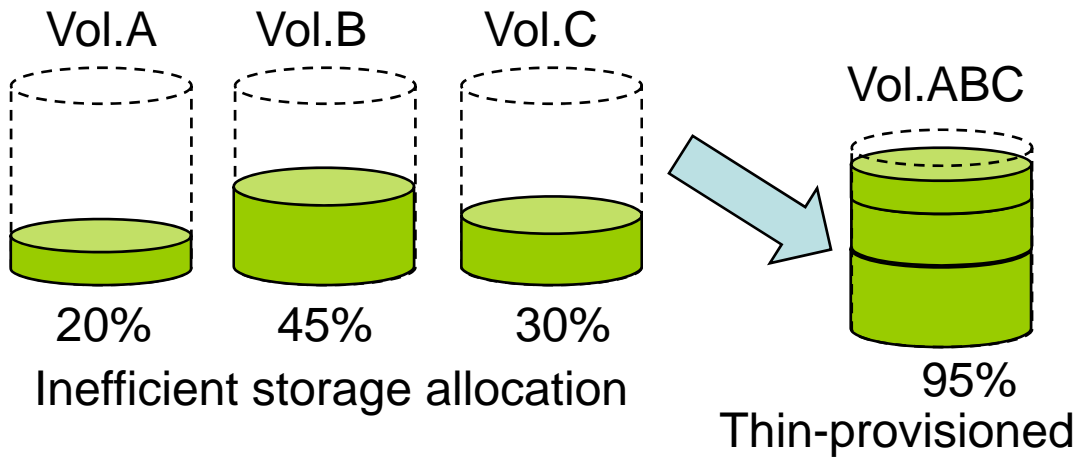
Check out SNIA Tutorial:

Deduplication Methods for Achieving Data Efficiency by M. Brisse, G. Senderov

Green Storage II

© 2008 Storage Networking Industry Association. All Rights Reserved.

Thin Provisioning




- user demands vary over time
- over-allocation anticipates future needs
- low capacity utilization
- low performance utilization


- Consolidate storage space to
 - ◆ improve utilization efficiency
 - ◆ reduce overall data center footprint
- Helpful to have good forecast of workloads
- Need dynamic provisioning to achieve best results



Check out SNIA Tutorial: **Software Technologies for Green Storage** by Alan Yoder



All disks spinning full-speed; high performance but no power saving



25% disks spun down; up to 25% power saving but some performance penalty

➤ MAID – Massive Array of Idle Disks

- ◆ Spin down a fraction of idling drives – requires predictable workload
 - ◆ Lowers power, increases time-to-data
- Multiple MAID levels (10%, 25%, 50%) can optimize performance and power for different scenarios and applications
- Maximum power must be considered for design points

SNIA Green Efforts

- SNIA Green Storage Initiative (GSI) and Green Technical Work Group (TWG)
 - ◆ on-going efforts to develop green standards & metrics
 - ◆ power measurements through multi-vendor “plug” fests
 - ◆ alliances with other active green organizations (The Green Grid, 80PLUS, Climate Savers, DMTF, SPEC, SPC)
 - ◆ collaboration with EPA on ENERGY STAR program
- Whitepapers / workshops
 - ◆ four tutorials at this SNW (see them all!)
 - ◆ online tutorials available (www.snia.org/education/tutorials)

Summary

- Ask tough “green” questions
 - ◆ get estimates of idle power, active power, power supply efficiency & total cost of ownership on systems you buy (or build)
 - ◆ consider implementation of “green” software features
 - ◆ be aware of performance and power tradeoffs – know your workloads!
- Get involved with SNIA Green efforts
 - ◆ weekly discussions and regular “plug” fests
- Learn about wider green technology and opportunities
 - ◆ online resources; workshops by SNIA, EPA, The Green Grid
- Share your experience / knowledge

References

- SNIA Green Storage Initiative – www.snia.org/green
- The Green Grid – www.greengrid.org
- EPA ENERGY STAR
(Data Center Energy Efficiency Initiatives –
www.energystar.gov/index.cfm?c=prod_development.server_efficiency)
- DOE Federal Energy Management Program – eere.energy.gov/femp
- Power calculators at various vendor sites
- SNIA Green Storage Outreach
 - ◆ USENIX LISA Conference, San Diego, November 2008
 - ◆ www.snia.org/forums/green



Q&A / Feedback

- Thank you for your attention!
- Please send any questions or comments on this presentation to SNIA: trackstorage@snia.org and trackgreenstorage@snia.org

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

SNIA Education Committee

SNIA Green Storage Initiative

Matthew Brisse

David Reinsel

Edgar St.Pierre

Willis Whittington

Alan Yoder

Wayne Adams

SNIA Green Storage TWG

