



Education

Green Data Center and Storage Technologies

Alan G. Yoder, Ph.D., NetApp

- ◆ The material contained in this tutorial is copyrighted by the SNIA unless otherwise noted.
- ◆ Member companies and individual members may use this material in presentations and literature under the following conditions:
 - ◆ Any slide or slides used must be reproduced in their entirety without modification
 - ◆ The SNIA must be acknowledged as the 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, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- ◆ The information presented herein represents the author's personal opinion and current understanding of the relevant 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.

➤ Technologies for Green Storage

- ◆ Hardware efficiencies are essential to reducing the amount of power used by storage. Equally real savings are obtained by reducing the number of copies of your data that must be made, kept and managed. This talk presents a number of technologies, ranging from thin provisioning and virtualization to flywheel UPSs, that each address part of the problem, and illustrates the impact that each technology can have on your data center footprint.

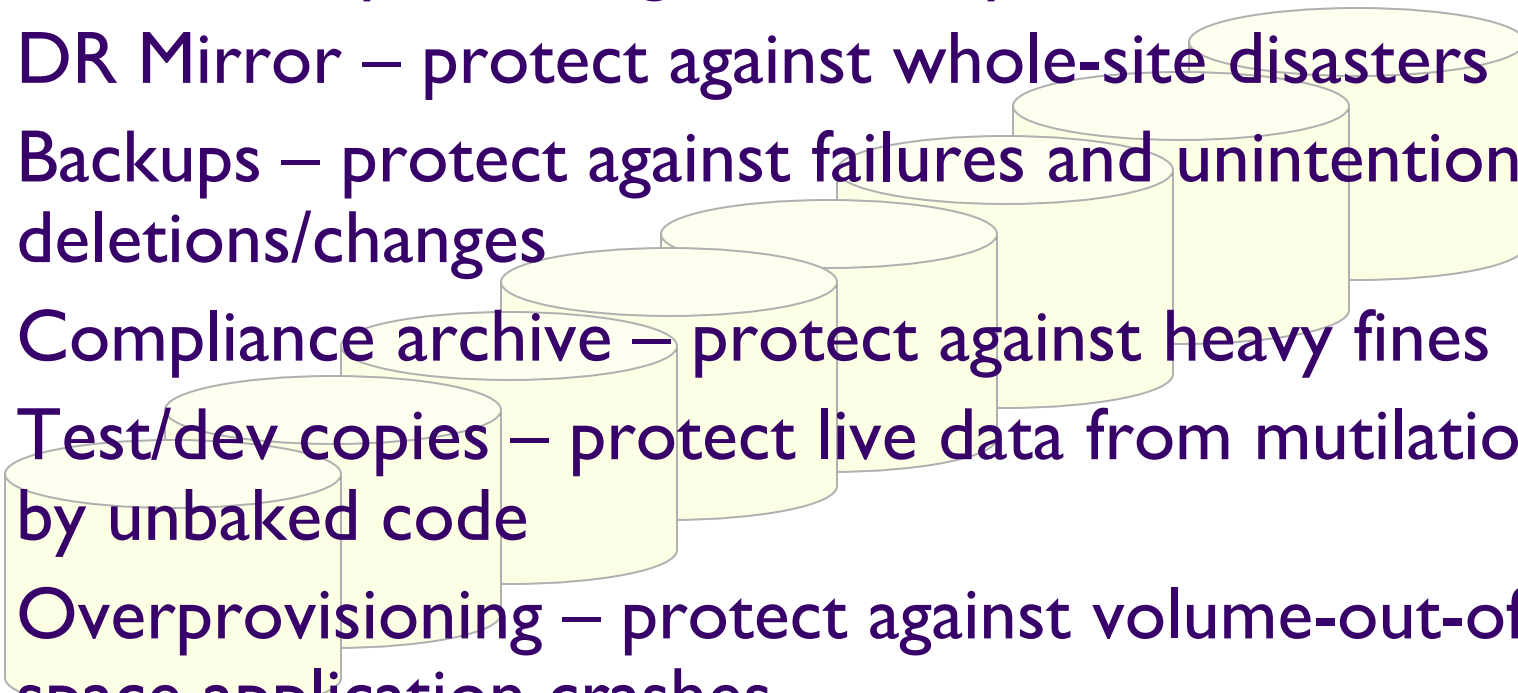
Outline

- Objectives
- The problem
- Storage technologies for energy saving
- Data center technologies for energy savings
- Typical savings

Objective


- Get more (storage) work done for less money
- Translation: reduce data center footprint
 - ◆ in space
 - › less storage equipment to buy, power and maintain
 - ◆ in energy
 - › more energy-efficient equipment
 - › less equipment to cool
 - › better cooling methodologies
 - › better power management
 - ◆ in administrative costs
 - › less storage equipment to manage

Problem 1: 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
 - Overprovisioning – protect against volume-out-of-space application crashes
 - Snapshots – quicker and more efficient backups and PIT (point in time) copies
- 

Result of redundancy

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



Problem 2: making heat just to cool it

- Servers, storage and switches are HEATERS
 - ◆ 100% efficient energy-to-heat conversion
 - ◆ Rotating media uses 85% of max power *at idle!*
- A/C is a big “undo” mechanism for overheating
 - ◆ But less than 100% efficient (typically 70%)

> 60% of the power in
a traditional data
center does no IT
work

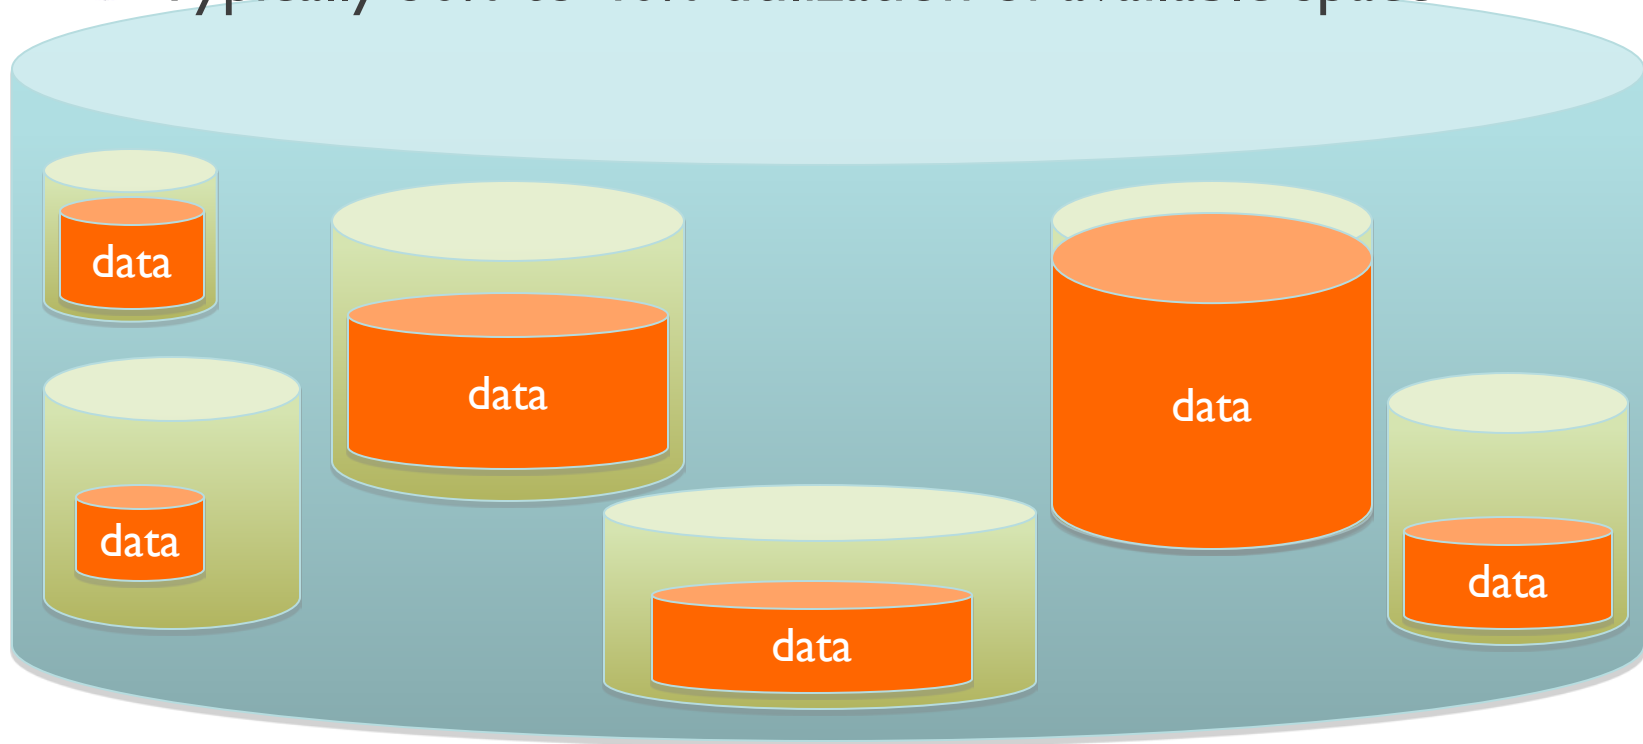
(PUE* ~ 2.5)



* PUE defined later


Problem 3: unused space

- Overprovisioning of systems
- Overprovisioning of containers
 - ◆ Typically 30% to 40% utilization of available space



Solution: green storage technologies


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



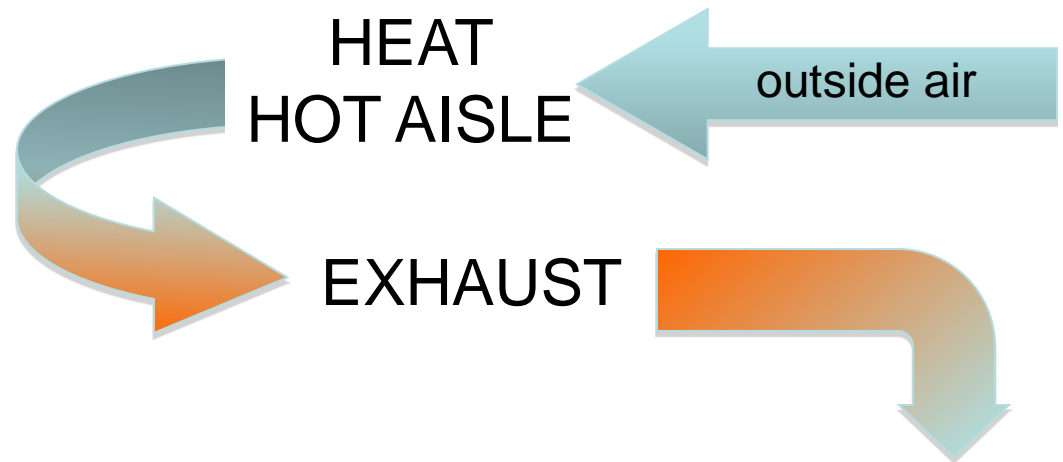
Solution: green facilities technologies

e.g.

BEFORE




AFTER



40% SAVINGS!
(PUE = 1.5)

Solution: getting back free space



Add storage as you need it

Green Storage Technologies

- Enabling technologies
 - ◆ Storage virtualization
 - ◆ Storage capacity planning
- Green software
 - ◆ Compression
 - ◆ Delta snapshots
 - ◆ Thin provisioning
 - ◆ Non-mirrored RAID
 - ◆ Deduplication and SIS



Green storage technologies (cont.)

- Other storage technologies and power saving techniques
 - ◆ Capacity vs. high performance drives
 - ◆ ILM / HSM / Tiering
 - ◆ MAID
 - ◆ SSDs / “Flash and stash”
 - ◆ Power supply and fan efficiencies
- Facilities-side technologies
 - ◆ Hot aisle/cold aisle
 - ◆ Water & natural cooling
 - ◆ Flywheel UPSs
 - ◆ Spam filtering



Enabling technologies

- Storage virtualization
- Storage capacity planning



- Mapping from physical location to virtual location
 - ◆ May exist at multiple layers
- In and of itself, not green wrt storage
 - ◆ No reduction in dataset size – it's just pointer remapping
- But foundational for most green storage technologies
 - ◆ Thin provisioning
 - ◆ Delta snapshots
- Also contributes in other areas
 - ◆ Flexibility, manageability, etc.
 - ◆ Fundamental for storage clouds

Storage capacity planning

- Needed to make best use of thin provisioning etc.
- Obtain and analyse baseline data
 - ◆ Many toolkits available from storage and storage management vendors
 - ◆ Toolkits usually slanted toward more purchase of said vendors' products
- Identify inefficiencies
 - ◆ Vendors usually eager to help find issues with other vendors' solutions
- Identify which green software technologies will address each inefficiency found
 - ◆ Ask vendors for proposals
 - ◆ Overall story more important than individual technologies

Green software technologies

- Compression
- Delta snapshots
- Thin provisioning
- Non-mirrored RAID
- Deduplication and SIS



➤ Compression

- ◆ Old and venerable
- ◆ Origins in signaling, number and coding theory
 - › Motivated by limited bandwidth and lossiness of satellite communications
- ◆ Scattered throughout the data stack
- ◆ Many formats already compressed
 - › JPEG, MPEG, MP3, etc.
 - › Lossless compression (LZW) necessary for unknown data types
- ◆ Configuration matters
 - › Compress before encrypting, decrypt before decompressing
- ◆ Difficult in block-based environments
 - › But becoming commonplace—usually in conjunction w/ dedup

➤ NOT just wholesale copies of the data

- ◆ We call those “snapshots” or “clones”

›



Check out the SNIA Dictionary !

www.snia.org/dictionary

➤ Data sharing

- ◆ Form of deduplication
- ◆ Data in snapshot shared with live data until one of them is written
- ◆ Two fundamental techniques
 - › Copy Out on Write
 - › Write to new live location

Delta snapshots (cont.)

➤ Typical uses – readonly

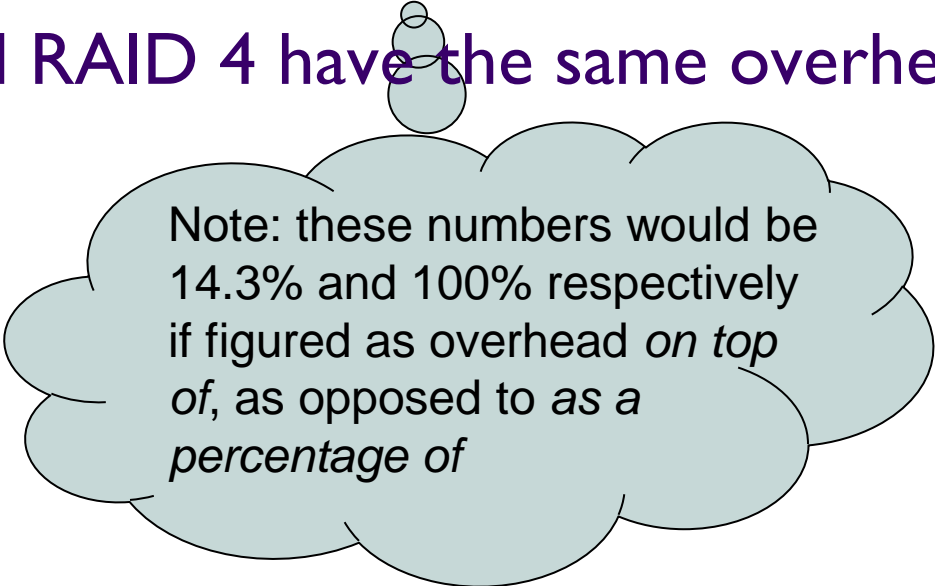
- ◆ Reducing RPO (recovery point objective)
 - › typically from a day to an hour or so
 - › key feature is ability to revert live system to a snapshot quickly
- ◆ Increasing backup window length
 - › reduced demands on backup hardware, backup window schedules

➤ Typical uses – read/write

- ◆ Quick bringup of duplicate datasets – minutes per TB
- ◆ What-if scenarios
- ◆ Testing of application changes against up-to-date datasets
- ◆ Testing of new applications with near-online data
- ◆ Booting/running of VM images from a golden master

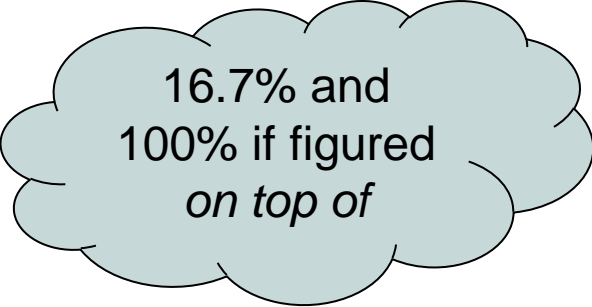
- Similar in concept to filesystem quotas
 - ◆ Volume “size” is merely a promise that that amount of storage will become available on demand
 - ◆ Storage is not actually allocated in the system until it is used
 - ◆ Admin tracks total cumulative use and makes sure that available storage is kept larger than used by some delta.
- Result: no more overprovisioning to avoid running out of space

- Allows any (one) drive in a RAID set to fail without data loss
- Requires only one extra drive in a RAID set
 - ◆ Much less raw capacity required than for mirroring
 - › Typical: 8-disk RAID 5 set: 12.5% overhead vs. 50% for mirroring
- Note: RAID 3 and RAID 4 have the same overhead as RAID 5



Note: these numbers would be 14.3% and 100% respectively if figured as overhead *on top of*, as opposed to *as a percentage of*

- **More dependable than mirroring**
 - ◆ Mirroring: can survive two failures in a disk group if they're not in the same mirrored pair
 - ◆ RAID 6: can survive failure of any two drives in the group
- **Requires two extra drives per RAID set**
 - ◆ However, typically somewhat larger RAID sets
- **Necessary as drive sizes increase**
 - ◆ Probability of a disk failure during RAID 5 parity reconstruct is getting too high
- **More green than mirroring**
 - ◆ 50% overhead in RAID 1 mirroring
 - ◆ 14.3% overhead in a 14-disk RAID 6 raidset



16.7% and
100% if figured
on top of

Deduplication and SIS

- Find duplicates at some level, substitute pointers to a single shared copy
- Block or sub-file based (dedup)
- Content or name based (SIS *, “file folding”)
- Inline (streaming) and post-process techniques
- Savings increase with number of copies found

◆ * SIS = Single Instance Store



**Check out the SNIA
Dictionary !**

www.snia.org/dictionary

Other technologies

- Capacity vs. high performance drives
- SSDs
- Flash and stash
- ILM / HSM
- MAID

- Power supply and fan efficiencies



Capacity vs. High Performance

- Picture rapidly changing due to SSDs
 - ◆ Formerly lo/hi perf ↔ SATA/FC
- SSDs
 - ◆ win big on read performance
 - ◆ lose out to FC/SAS on sustained write perf.
 - ◆ lose big on raw cost / GB
 - ◆ win big on “green” factor (idle power)

SSDs (Solid State Disks)

➤ Usually refers to FLASH-based disks

- ◆ Pros
 - › Great READ performance
 - › At rest power consumption = 0
 - › No access time penalty when idle (cf. MAID)
 - › No need to keep some disks spinning (cf. MAID)
- ◆ Cons
 - › WRITE performance may be < mechanical disks
 - › Cost >> mechanical disks except at very high perf points
 - › Wear leveling requires a high space overhead
- ◆ Note: these dynamics changing rapidly with time
 - › SSSI – SNIA Solid State Storage Initiative



**Check out the SNIA Tutorials in
the Solid State Storage (SSS)
Tutorial track**

“Flash and stash”

- Usually refers to large arrays of SATA-based disks fronted by large flash caches
 - ◆ On the order of 1TB flash (or SSDs)
 - ◆ Most working sets in flash
 - ◆ Reduced power (SATA vs. SAS)
 - ◆ Alternative to tiering – well-known caching algorithms
 - ◆ Not useful for write-intensive workloads
 - ◆ Naive implementations may take a long time to warm the cache after a reboot



**Check out the SNIA Tutorials in
the Solid State Storage (SSS)
Tutorial track**

- Exploit cost differences between storage tiers
 - ◆ Idea: automatically move data to an appropriate storage platform at each period in its lifetime
- Tier change must have substantial value to make the overhead worth it
 - ◆ Cost of system
 - ◆ Cost of administration
 - ◆ Cost of data movement
- Tiering = “advanced HSM”
 - ◆ Colocation of tiers in a single system makes data movement much more efficient
 - ◆ May make use of ILM concepts to determine tiering level of data

MAID (Massive Array of Idle Disks)

➤ Idea: spin down disks when not in use

◆ Pros

- › Disks use no power when spun down
- › > 50% power savings at idle

◆ Cons

- › Most data near-online (access times of several seconds)
- › Background disk housekeeping difficult
- › Often the same data center sizing requirements (UPSs, CRAC units, PDUs etc.), but these are used at lower efficiencies
- › Competition from SSDs
- › Competition from....


➤ Oldie but goldie

◆ Pros

- › Tapes use no power when inactive
- › > 90% power savings at idle

◆ Cons

- › Data is at best near-online (access times of several seconds)
- › Not a random access format
- › Lack of true resilience to format failure
 - Redundant Array of Independent Tape? (RAIT)



Check out SNIA Tutorial:

**Introduction to Data Protection:
Backup to Tape, Disk and Beyond**

- Efficiency of power supply an up front waste
 - ◆ Formerly 60-70%
 - ◆ Nowadays 80-95%
 - › Climate Savers
 - › 80plus group (see <http://www.plugloadsolutions.com/80PlusPowerSupplies.aspx>)
 - › Note: Efficient PSs *are* more expensive
- Variable speed fans
 - ◆ Common nowadays
 - ◆ Software (OS) control

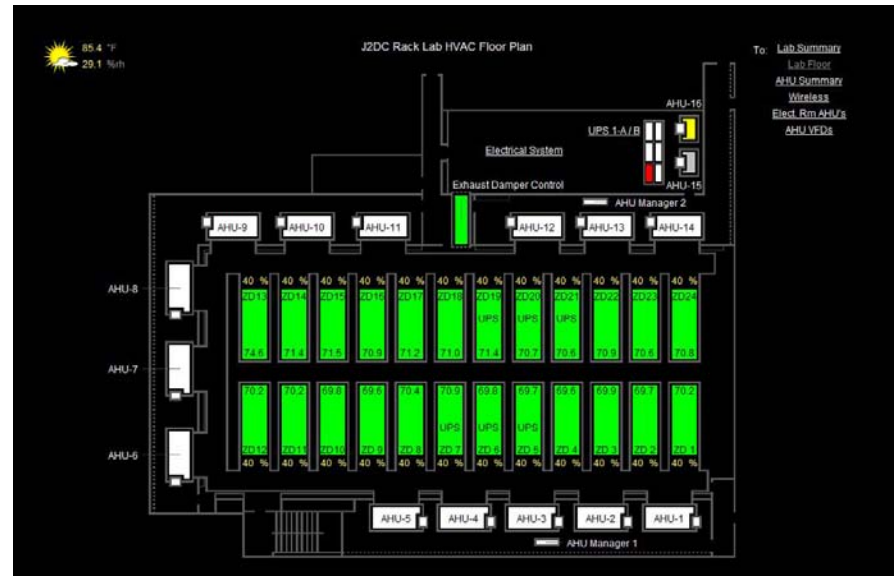
Facilities-side technologies

- Monitoring
- Hot aisle – cold aisle technologies
- Other
 - ◆ Spam filtering
 - ◆ Water and natural cooling
 - ◆ Flywheel UPSs

(UPS = Uninterruptible Power Supply)



- Critical to increased efficiencies
 - ◆ Lights out operation
 - ◆ Tightening up of temperature tolerances
 - ◆ Better staff utilization
 - ◆ Anomaly detection



Hot aisle / cold aisle technologies

- Segregate airflows into hot and/or cold aisles (backs and fronts of servers)
 - ◆ More precise control
 - ◆ Allows higher temperature differentials (more efficient)
 - ◆ Several emerging approaches
 - › Hot air plenum
 - › Complete containment
 - ◆ **Current trend toward hot aisle containment with cold air plenum**
 - ◆ Must-have: blanking plates
 - › Very important



Savings calculations

- Space savings
- Equipment power savings
- Facilities power savings



Savings calculations (storage)

- Calculations herein are for space savings
- Relationship of space to \$\$ is loose
- But every TB of disks you don't buy saves you
 - ◆ CapEx (Capital Expenditure)
 - › for the equipment
 - › for the footprint
 - › for plant, power conditioning and cooling
 - ◆ OpEx (Operational Expenditure)
 - › for equipment power
 - › for power conditioning and cooling
 - › for storage management
 - › for service contract fees

Typical space savings

➤ Compression

- ◆ 15 – 40% savings
- ◆ Remember, no savings from already compressed file formats

➤ Readonly delta snapshots

- ◆ 90 – 99% savings per snapshot, compared to full PIT copies

➤ Writeable delta snapshots

- ◆ 80 – 99% savings per snapshot, compared to full PIT copies

➤ Thin provisioning

- ◆ 20 - 50% savings
- ◆ Average 30% utilization → over 80% utilization

➤ RAID

- ◆ 35% savings for e.g. 14-disk RAID 6 set, compared to RAID 1/10
- ◆ Note: use the right RAID level for your application

➤ Deduplication

- ◆ primary and archive storage: ~35% savings
- ◆ backup: 35% - 95% savings, depending on dataset, time interval, and ratio of incremental to full backups

Caveats

- Savings estimates are real, but best taken as anecdotal
 - ◆ YMMV – your mileage may vary
 - ◆ Make your vendors prove their claims in your application environment
- Green options in some implementations may degrade RAS and/or performance
 - ◆ E.g.
 - › RAID 5 vs RAID 10 – reduced RAS
 - › 7.5K SATA vs 15K FC – reduced performance
 - › Tape vs disk during legal discovery – large fine if too slow
 - ◆ Make your vendor tell you the cons that go with the pros

Caveats (cont.)

- Savings estimations are always imprecise
- Obviously can't achieve infinite savings by applying all technologies
 - ◆ Diminishing returns

Overall 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

➤ Reduced equipment footprint per TB data

- ◆ Vendors claiming up to 50% reductions
- ◆ *Require proof or guarantees*

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

- It's a brave new world
- Most vendors shipping most of these technologies
- Ask your vendor to justify use of the technology against your business problems
- The advantage is yours: seize the day!

- 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

**Al Thomason
Wayne Adams**

**SW Worth
Larry Freeman**