

# SNIA Green Update for Developers

Dr. Erik Riedel, Seagate Technology

- ❑ The material contained in this presentation 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.**

- SNIA Green Update for Developers
  - This talk will outline some storage-specific topics related to energy-efficiency, including metrics for measuring, managing and designing for power. We will overview ongoing efforts in the SNIA Green Storage TWG and in partnership with The Green Grid, the DMTF and other industry groups.

- Power measurement
  - storage subsystems
  - idle and active modes
  - power supply loading & efficiencies
  - power measurement & monitoring equipment
- SNIA green storage efforts
  - green metrics & standards
  - “unplugged” fests
  - alliances

# 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 over many drives
  - ❑ low access times → faster actuators; higher rotational speeds; caches
  - ❑ How fast and immediate must data be available? (time-to-data)
- ❑ Data integrity
  - ❑ estimates of 25% of “digital universe” is unique, 75% replicas / duplicates
  - ❑ 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? What time-to-data?

# Potential paths to “green”

- ❑ Improve usage efficiency
  - ❑ Deduplication
  - ❑ Thin provisioning
- ❑ 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

must be driven by  
metrics / standards  
/ guidelines

# 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, and workload affect power consumption and efficiency



Apps

Software




Power Supplies

Fans

Controllers

Hard drives

# Power Supply Efficiency



Data from 45 single-voltage PSUs from 8 vendors (products available 1<sup>st</sup> 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 modes
- ❑ Design must take several trade-offs into consideration
  - ❑ peak load, average loads and redundant system configuration
  - ❑ supplies may operate at <50% of max. load 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
  - ❑ System uses **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 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
- ❑ Measure operating 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

- Power measurement
  - storage subsystems
  - idle and active modes
  - power supply loading & efficiencies
  - power measurement & monitoring equipment
- SNIA green storage efforts
  - green metrics & standards
  - “unplugged” fests
  - alliances

# Metrics Examples



## Workload considerations

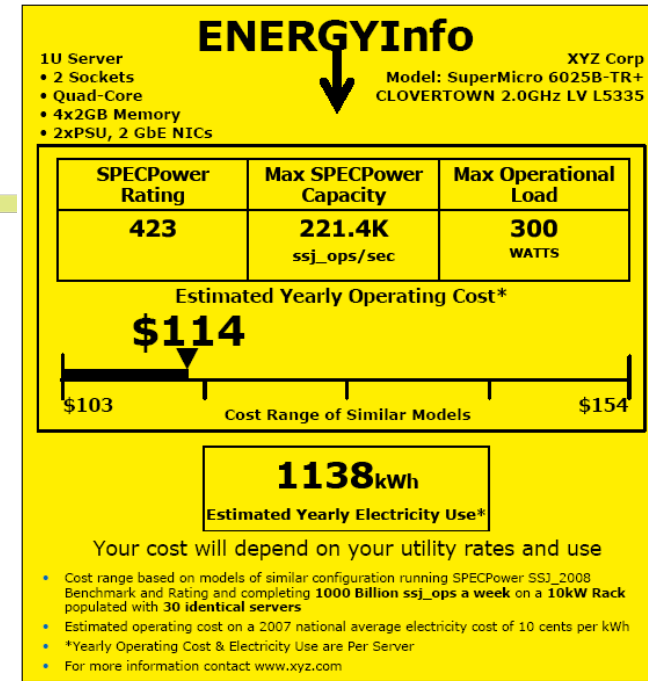
- Data at rest – Idle power (GB/Watt)
- Data on the move – Throughput (MB/s)
- Data at work – Performance (IOPS)

## Potentially useful metrics

- GB per Watt; MB/s per Watt; IOPS per Watt
- Power supply efficiency; CO<sub>2</sub> footprint
- Total annual energy bill

## Reliability / availability / serviceability considerations

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



*Server example*

- ❑ SNIA Green Storage Initiative (GSI) and Green Storage Technical Work Group (TWG)
  - ❑ on-going efforts to develop green standards & metrics
  - ❑ power measurements through multi-vendor “unplug” fests
  - ❑ next face-to-face meeting November 2008 in San Diego, next unplugged fest April 2009 in Colorado Springs
- ❑ Four tutorials at the upcoming SNW (October 2008)
  - ❑ Green Storage I – Economics, Environment, Energy & Engineering
  - ❑ Green Storage II – Metrics and Measurement
  - ❑ Software Technologies for Green Storage
  - ❑ Building the Green Data Center
- ❑ online tutorials available ([www.snia.org/education/tutorials](http://www.snia.org/education/tutorials))

- ❑ Alliances with other active green organizations
  - ❑ The Green Grid – servers & data centers
  - ❑ 80PLUS, Climate Savers – power supplies
  - ❑ DMTF – management & monitoring interfaces
  - ❑ SPEC, SPC – benchmarks & power measurement
- ❑ Collaboration with EPA on ENERGY STAR program
  - ❑ ENERGY STAR for Servers – ongoing drafts in discussion
  - ❑ considerations for storage power will be significantly different than servers

- ❑ Ask tough “green” questions
  - ❑ get estimates of idle power, active power, power supply efficiency & total cost of ownership on systems you are designing
  - ❑ consider possible wins of “green” software features
  - ❑ be aware of performance and power tradeoffs – know your workloads!
- ❑ Get involved with SNIA Green efforts
  - ❑ weekly discussions and regular “unplugged” fests
  - ❑ next face-to-face meeting is November 2008 in San Diego
  - ❑ encourage your Marketing teams to participate in GSI
- ❑ 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](http://www.snia.org/green)
- ❑ The Green Grid – [www.greengrid.org](http://www.greengrid.org)
- ❑ EPA ENERGY STAR  
(Data Center Energy Efficiency Initiatives – [www.energystar.gov/index.cfm?c=prod\\_development.server\\_efficiency](http://www.energystar.gov/index.cfm?c=prod_development.server_efficiency))
- ❑ DOE Federal Energy Management Program – [eere.energy.gov/femp](http://eere.energy.gov/femp)
- ❑ Power calculators at various vendor sites
- ❑ SPEC website for measurement devices and settings  
[www.spec.org/power\\_ssj2008/device-list.html](http://www.spec.org/power_ssj2008/device-list.html)
- ❑ SNIA Green Storage Outreach
  - ❑ USENIX LISA Conference, San Diego, November 2008



- ❑ Thank you for your attention!
- ❑ Please send any questions or comments on this presentation to SNIA: [greentwg@snia.org](mailto:greentwg@snia.org) and [trackgreenstorage@snia.org](mailto:trackgreenstorage@snia.org)

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

*SNIA Education Committee*

**SNIA Green Storage Initiative  
Wayne Adams  
SNIA Green Storage TWG**