SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE
Innovation in Storage Products, Services, and Solutions
June 13-15, 2016 | Marriott San Mateo | San Mateo, CA

Power-Efficient Data Storage

Brian Zahnstecher, Principal, IEEE Senior Member

PowerRox™

bz@powerrox.com (508) 847-5747

www.powerrox.com Twitter: @PowerRoxLLC

1125 S Baywood Ave., San Jose, CA 95128

ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT. © POWERROX LLC 2016

Overview

- Power Characteristics of Storage Products
- How does power differ amongst storage products?
- What happens when we scale?
- Opportunities for Power Savings

PRx™
© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.
SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE

Power Characteristics of Storage Products

- Main Types of Storage Products
  - Spinning Media
    - SAS
    - SATA
    - FireWire
    - SCSI
    - Shingled Magnetic Recording (SMR)
    - He *(more on this later…)*
  - Cold Storage
  - Tape

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

Power Characteristics of Storage Products

- Main Types of Storage Products
  - Solid-state Media
    - Flash
    - DRAM
    - PCIe
    - Other Non-volatile Memory (NVM)
      - ReRAM
      - Magnetic RAM
    - Phase-Change Media
      - Intel/Micron 3D X-point (Storage Class Memory)
      - IBM 3-bit/cell (8 distinctive states)
  - Hybrid Solutions

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

Power Characteristics of Storage Products

- Why 12V vs. 3.3V or 5V?
  - Motors vs. Solid-state
    - Power usage is directly related to motor speed.
  - Current Handling Capability
  - Some exceptions (i.e. – laptops) where 12V is not available.
  - Optimization of efficiency can really depend on how the system power is architected.

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

Power Characteristics of Storage Products

- Who likes big equations?

$$E_{Total} = \sum_{i=0}^{N} E_{Active} + \sum_{i=0}^{S} E_{Seek} + \sum_{i=0}^{I} E_{Idle}, \quad (6)$$

where E_{Total} is the total energy, in Joules; E_{Active} is the active energy, in Joules; E_{Seek} is the seek energy, in Joules; and E_{Idle} is the idle energy, in Joules.

$$E_{Active} = \frac{S}{B} \times P_{Active}, \quad (1)$$

where E_{Active} is active energy, in Joules; S is filesize, in MB; B is bandwidth, in MB/s; and P_{Active} is the active power, in Watts, as provided by the drive manufacturer.

$$E_{Seek} = T_{Seek} \times P_{Seek}, \quad (5)$$

where E_{Seek} is seek energy, in Joules; T_{Seek} is the seek time, in seconds; and P_{Seek} is the seek power, in Watts.

$$T_{Seek} = a\sqrt{abs\left(\left\lfloor \frac{LBN_{dest.}}{SPT} \right\rfloor - \left\lfloor \frac{LBN_{start}}{SPT} \right\rfloor\right)} + s, \quad (2)$$

where T_{Seek} is seek time, in milliseconds; a is a constant coefficient; $LBN_{dest.}$ is the destination LBN; LBN_{start} is the start LBN; SPT is the number of sectors per track; and s is the head settle time, in milliseconds.

IMAGE CREDIT: A. Hylick and R. Sohan, "A Methodology for Generating Disk Drive Energy Models Using Performance Data," in HotPower'09. ACM, 2009.

- Yeah, me neither.

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

Power Characteristics of Storage Products

- Active vs. Idle vs. Standby Power
  - It is important to consider the difference between peak and steady-state power. It can be a surprisingly small delta sometimes.
  - Beware the Sum of Maxima
  - Peak power really only matters at turn-on (and wake from standby).
  - Most systems have a staggered turn-on to mitigate inrush currents.
  - With the tradeoff between power and response, it is important that you validate just how low of a power state is achieved in a standby mode. (i.e. – 0W not always 0W.)

PRx™
© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.
SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE

Power Characteristics of Storage Products

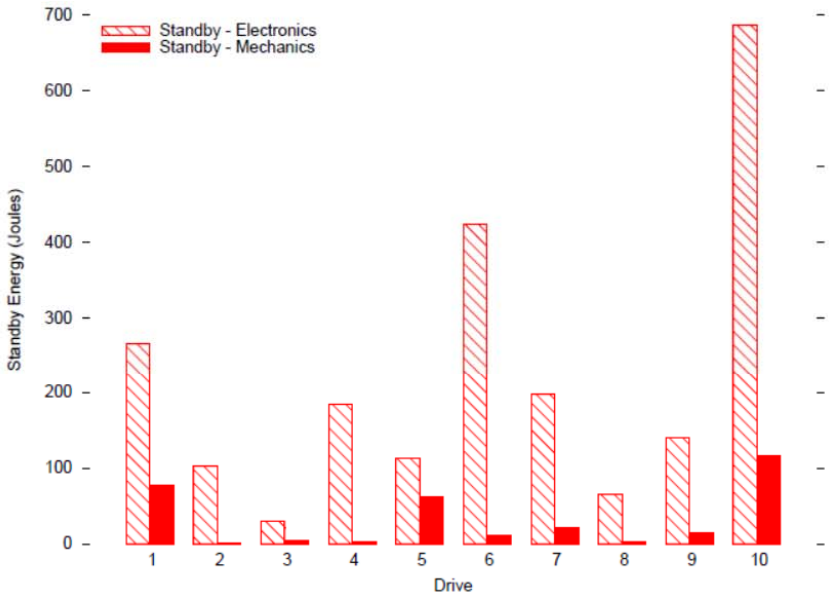

Figure 3. Summary of 5-Minute Standby Energy Consumptions

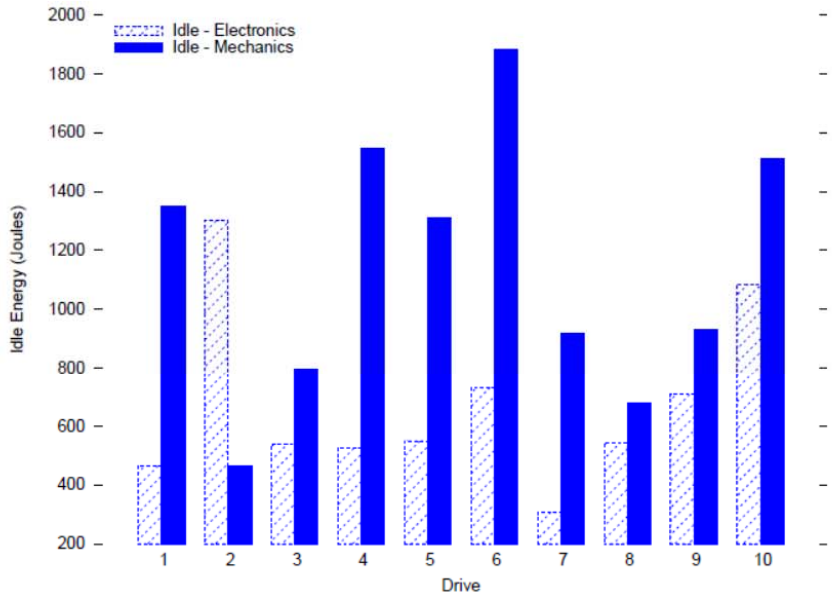

Figure 4. Summary of 5-Minute Idle Energy Consumptions

IMAGE CREDIT: A. Hylick, R. Sohan, A. Rice and B. Jones, "An Analysis of Hard Drive Energy Consumption," 2008 IEEE International Symposium on Modeling, Analysis and Simulation of Computers and Telecommunication Systems, Baltimore, MD, 2008, pp. 1-10. doi: 10.1109/MASCOT.2008.4770567

PRx™
© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.
SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE

Power Characteristics of Storage Products

- Need to Think in Terms of Utilization & Density
  - Consider a product's total power, but when comparing to a different storage technology, it is very important to also consider density.
  - Enclosures carry a large power overhead.
- Figures of Merit Unique to Storage
  - W/TB (or W/GB)
  - IOPS/W

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

How does power differ amongst storage products?

- Where is the power being consumed?
  - ***NOTE***: a good chunk of the power is spent transferring data, not processing and/or storing it.
    - Storage class memory may address this concern.
  - Wearables and IoT will change how and where data is stored, thus impacting the power consumption of storage products.

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

What happens when we scale?

- Economies of Scale Apply
  - Applies to both CAPEX & OPEX
  - Even the smallest player can have the best Power Usage Effectiveness (PUE).
  - Small Drives + Virtualization = A Large Drive

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

What happens when we scale?

- The Impact of the Cloud
  - Mitigate Overprovisioning
  - Ultimate Reliability (9+ 9s)
  - Utilize the Most Power-Efficient Storage Solutions
  - Utilize the Most Cost-Efficient Storage Solutions
  - Pay By the Minute

- Yeah, but what about security?
  - Good point, but this is a power talk! ☺
  - Self-encrypting Drives (SED)
    - So let us quickly look at SED power cost…

PRx™
© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.
SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE

What happens when we scale?

- The Power Cost of Encryption
  - Software Vs. Hardware Encryption

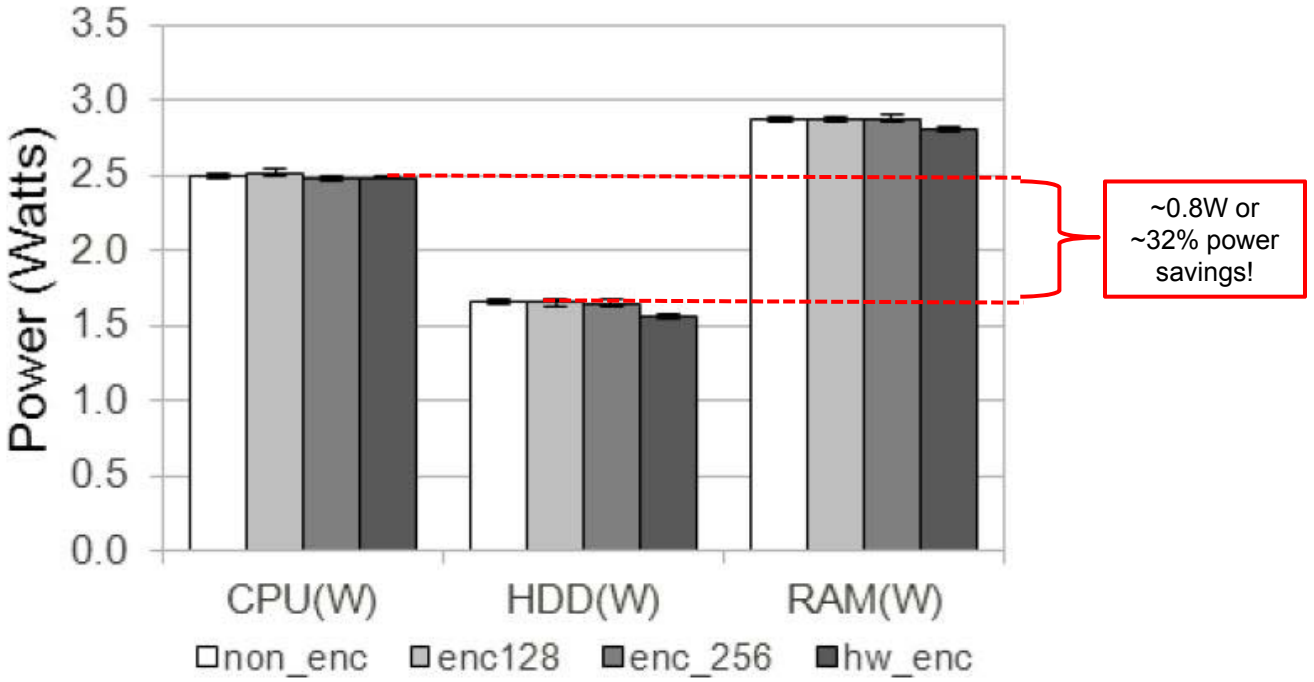

Figure 1. Power Used for 10MB File Writes

IMAGE CREDIT: A. Fujimoto, P. Peterson and P. Reiher, "Comparing the Power of Full Disk Encryption Alternatives," Green Computing Conference (IGCC), 2012 International, San Jose, CA, 2012, pp. 1-6. doi: 10.1109/IGCC.2012.6322245

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

What happens when we scale?

- The Power Cost of Encryption
  - Software Vs. Hardware Encryption

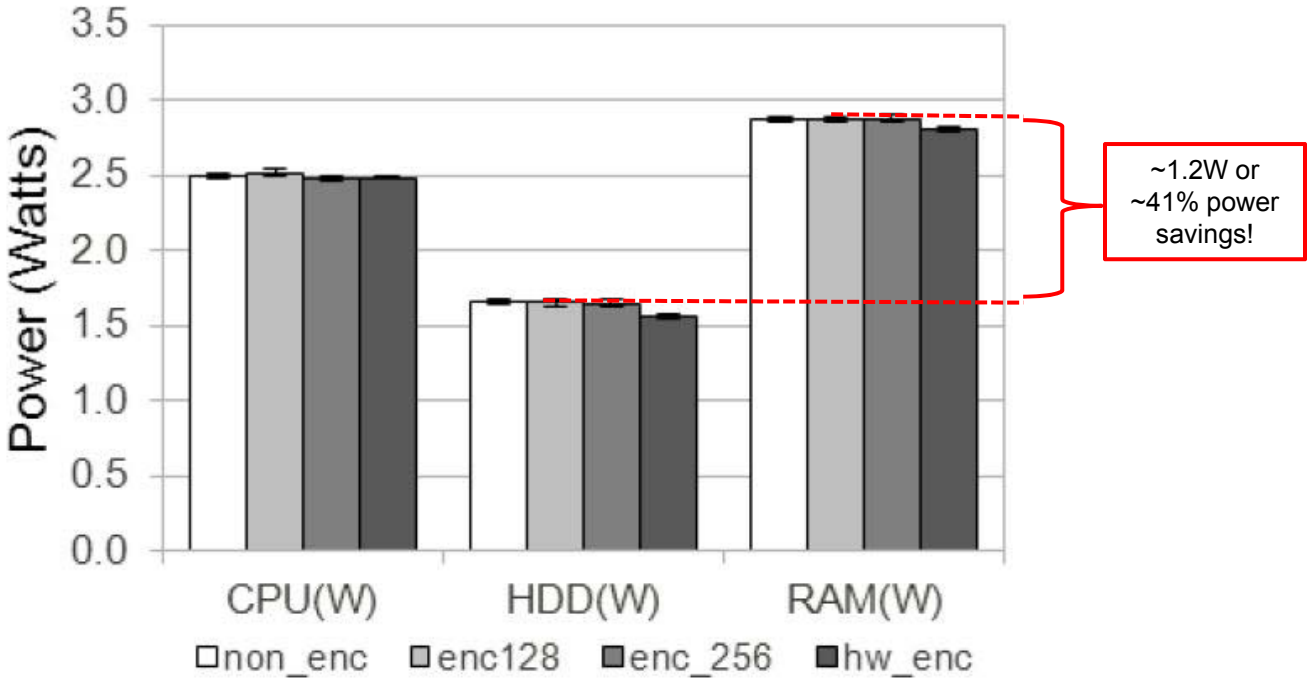

Figure 1. Power Used for 10MB File Writes

IMAGE CREDIT: A. Fujimoto, P. Peterson and P. Reiher, "Comparing the Power of Full Disk Encryption Alternatives," Green Computing Conference (IGCC), 2012 International, San Jose, CA, 2012, pp. 1-6. doi: 10.1109/IGCC.2012.6322245

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

Opportunities for Power Savings

"It costs more to run a storage device over three years than to buy it."

– SearchStorage

What is the most power efficient storage device in the world?

***HINT**: this is absolutely a trick question, but with a very objective answer.*

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

Opportunities for Power Savings

- Long-term vs. Near-term Solutions
  - System Optimizations
  - Infrastructure Overhaul
    - He Drives (Density & Power)
- The Enclosure
  - Fans, Controllers, ASICs
- Virtualization
  - Consolidate & Turn-off
- Deduplication
  - Get Rid of All Those Costly Copies

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

Opportunities for Power Savings

- The Cloud
  - Take Advantage of Those Economies of Scale
- Cold Storage
  - Listen to Your Stats, Take Action Accordingly
- DC Power
  - Mitigating Conversion Steps Can Yield Significant Savings
- Disaggregation
  - Separate & Optimize

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

Summary

- **Pay close attention to the difference between peak and steady-state power as it will tell you a lot about TCO.**
- **Use It or Lose It**
- **When evaluating power, need to consider the big picture for an accurate comparison.**
  - This can include transport and overhead costs.
- **Put the Cloud to Work for You**
- **Many of the same concepts and initiatives associated with saving power at the data center level all apply directly to storage solutions.**

PRx™
© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.
SNIA DSI '16 DATA STORAGE INNOVATION CONFERENCE

Thanks a lot for your time and attention!

Any questions and/or comments?

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.

References

- “Breaking Capacity Barriers With Seagate Shingled Magnetic Recording” = http://www.seagate.com/tech-insights/breaking-areal-density-barriers-with-seagate-smr-master-ti/
- A. Hylick and R. Sohan, "A Methodology for Generating Disk Drive Energy Models Using Performance Data," in HotPower'09. ACM, 2009.
- A. Hylick, R. Sohan, A. Rice and B. Jones, "An Analysis of Hard Drive Energy Consumption," 2008 IEEE International Symposium on Modeling, Analysis and Simulation of Computers and Telecommunication Systems, Baltimore, MD, 2008, pp. 1-10. doi: 10.1109/MASCOT.2008.4770567
- "3D XPoint™ Technology Revolutionizes Storage Memory" = http://www.intel.com/content/www/us/en/architecture-and-technology/3d-xpoint-technology-animation.html
- “IBM Reveals 3 Bits/Cell PCM-MLC Challenge to 3DX Stacking” Ron Neale, EE Times (5/17/16)
- Power Usage Effectiveness (PUE) / Data Center infrastructure Efficiency (DCiE) = http://www.thegreengrid.org/sitecore/content/Global/Content/white-papers/The-Green-Grid-Data-Center-Power-Efficiency-Metrics-PUE-and-DCiE.aspx
- A. Fujimoto, P. Peterson and P. Reiher, "Comparing the Power of Full Disk Encryption Alternatives," Green Computing Conference (IGCC), 2012 International, San Jose, CA, 2012, pp. 1-6. doi: 10.1109/IGCC.2012.6322245
- Google Cloud Platform = https://cloud.google.com/storage/
- D. Elliot & J. Kember, “Cloud, Archiving and the Future,” Storage Visions 2016, Las Vegas, NV, 2016.
- “Hot storage – Power and cooling concerns” SearchStorage Article (March 2007)
- “Beyond air. Helium takes you higher.” HGST Whitepaper (2015)

© POWERROX LLC 2016
ALL INFORMATION SHALL BE CONSIDERED POWERROX LLC PROPERTY UNLESS OTHERWISE SUPERSEDED BY ANOTHER DOCUMENT.