SNIA

Education

GREEN STORAGE PRODUCTS: Efficiency with ENERGY STAR & Beyond

Patrick Stanko
Oracle Tape Development, Oracle
Carlos Pratt
IBM Technical Support Marketing and Sales

SNIA Legal Notice

Education SNIA

- 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.

© 2011 Storage Networking Industry Association. All Rights Reserved.

Abstract

Education SNIA

- **GREEN STORAGE PRODUCTS: Efficiency with ENERGY STAR & Beyond**
  - 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" and the components of modern storage systems that impact energy consumption. We will discuss potential metrics for measuring, managing and designing for power in storage systems. We will also outline ongoing efforts by the EPA ENERGY STAR® program and in the SNIA Green Storage Technical Work Group (TWG) to standardize metrics for measuring storage systems.

© 2011 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Background and Green Storage
  - Revisit what is a basic storage unit
  - What influences green storage
- Storage SNIA Taxonomy update
- Current status of
  - SNIA Green TWG
  - EPA
  - SNIA
    - Green Efforts
    - Emerald Program (Quick introduction)

© 2011 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Background and Green Storage
  - Revisit what is a basic storage unit
  - What influences green storage

© 2011 Storage Networking Industry Association. All Rights Reserved.

Background and Green Storage

- Background
  - Revisit the basics of storage systems
- Green Storage
  - What Impacts the power consumption?
    - Opportunities to make storage green(er)
  - What factors should be considered?
  - What do these results mean?
  - What in the industry that may help your green storage effort
    - Emerald
    - SPC
    - The Green Grid

© 2011 Storage Networking Industry Association. All Rights Reserved.

Background

Basic Anatomy of a Disk Storage System and its use

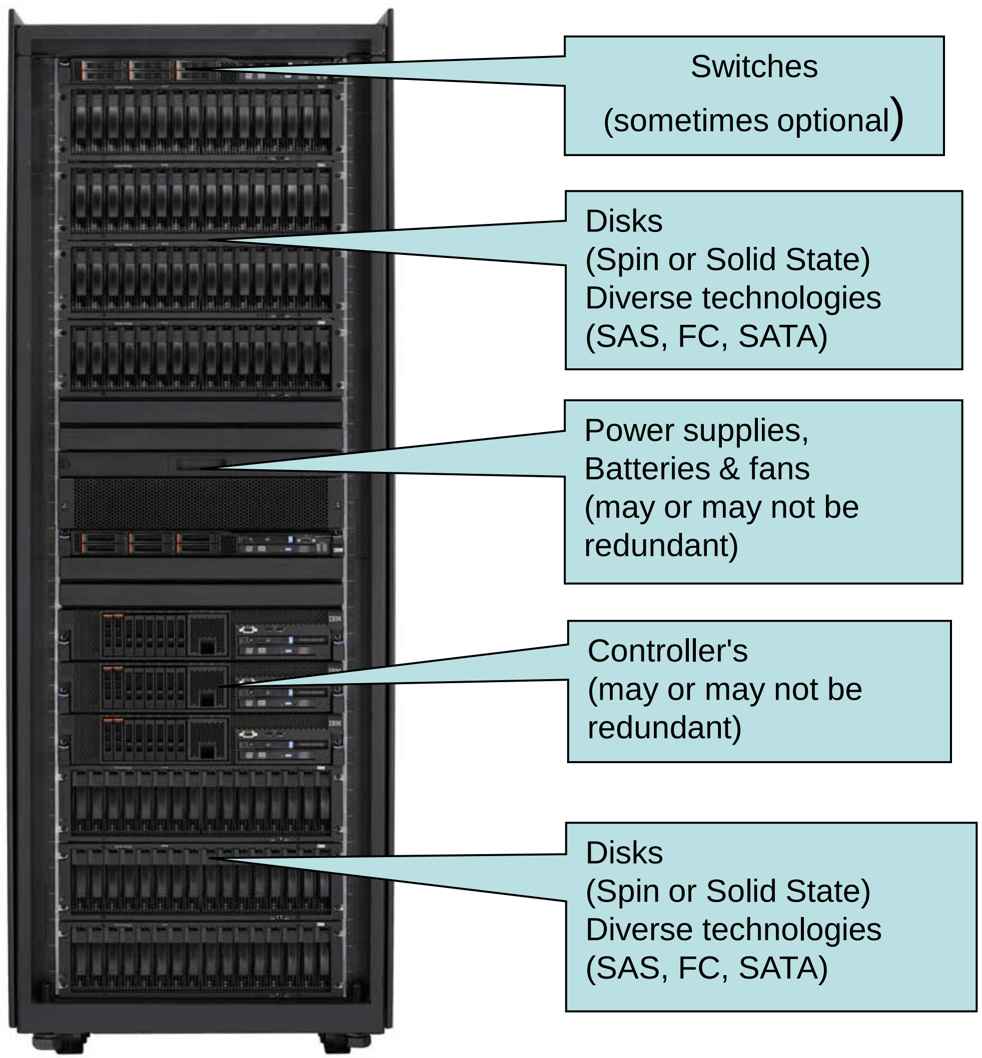

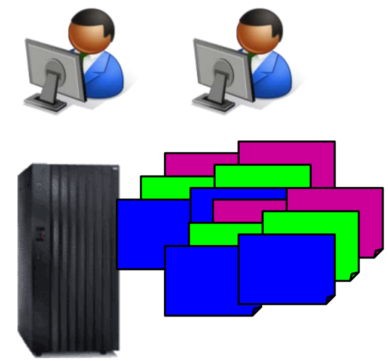

Users and Apps

Other:

Software (firmware & microcode)

Hardware Design

Environment

© 2011 Storage Networking Industry Association. All Rights Reserved.

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 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? How fast must data be available?

© 2011 Storage Networking Industry Association. All Rights Reserved.

Opportunities to make storage green(er)

Education SNIA

- Environment
  - Higher system tolerance to high/lower temperatures and humidity
  - In line with cold and hot aisles designs on new data centers
- Improve usage efficiency
  - De-duplication and compression
  - Thin provisioning

must be driven by metrics / standards / guidelines

- Minimize energy consumption
  - Improved component designs – high-efficiency power supplies, advanced & flexible storage devices
  - Variants of MAID – idle and spin-down
- New technologies
  - Solid state storage
  - Alternative + hybrid system designs (opportunity to rethink)

© 2011 Storage Networking Industry Association. All Rights Reserved.

What to measure and evaluate - metrics

- SNIA recommended metrics
  - Capacity metric (ready-idle)
    - Relates the power of the system to its total storage raw capacity. It is reported as GB/watt (or TB/watt)
    - Power required to store and protect the data
  - Workload metric (Active)
    - Relates the power of the system to the maximum possible IOPS generated by a specific random stress load. It is reported as IOPS/watt
    - Power required to randomly supply data to and from a host
  - Bandwidth metric (Active)
    - Relates the power drown by the system to the maximum possible MBPS generated by a specific sequential stress load. It is reported as MBPS/watt
    - Power required to stream data to and from a host

© 2011 Storage Networking Industry Association. All Rights Reserved.

What to measure and evaluate – Power example

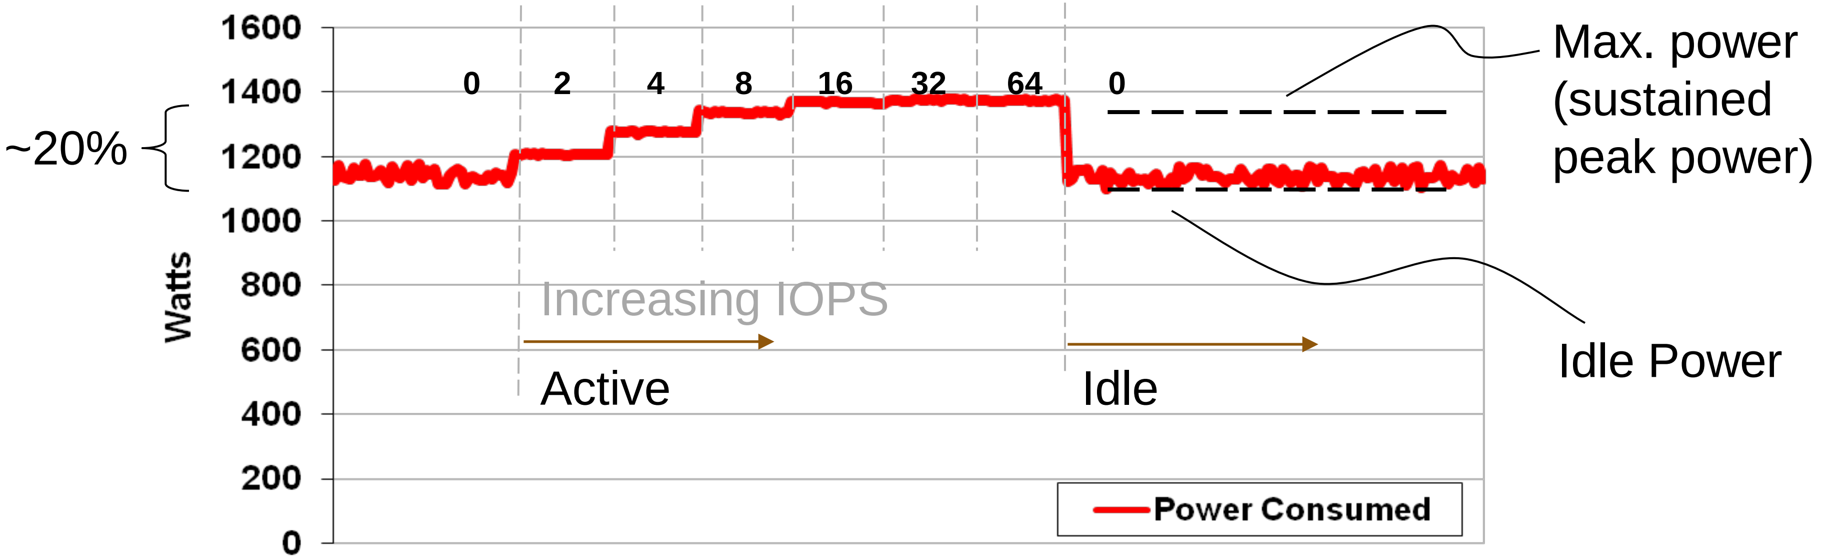

- Ideally, systems consume minimum power in all modes
  - Example system consumes **significant power in idle** (**80% of max**)
- % of time in Idle versus Active depends on storage type, application and workloads; available optimizations will vary
  - Power itself is only one part of the story it must be reflected as a metric as indicated on the previous charts.
- Power consumed is not linearly proportional to workload (indicates potential room for improvement)

© 2011 Storage Networking Industry Association. All Rights Reserved.

What to measure and evaluate – What to consider

- Recommended Analysis tools
  - Power meters recommended on the Emerald Specification
  - Temperature recording tools recommended on the Emerald Specification
  - SNIA will also recognize the SPEC recommended measurement devices as indicated on their web site: www.spec.org/power/docs/SPECpower-Device_List.html
- If the storage system reports power and temperature consider loggin it to compare with the analyzers for accuracy comparison.
- Both total and sub-system power consumed are valuable info

© 2011 Storage Networking Industry Association. All Rights Reserved.

What to measure and evaluate – What it means? (1/2)

- Depending on the systems and their usage their energy usage may be evaluated according to:
  - Is the system idle time at least 12 hours or more a day?
    - You should be interested in the power required to store the data
    - **capacity metric (GB/Watt)** may be your best indicator on how energy efficient your system is
    - The larger this number is the less watts are used to energize the total storage of your system

© 2011 Storage Networking Industry Association. All Rights Reserved.

What to measure and evaluate – What it means? (2/2)

- Depending on the systems and their usage their energy usage may be evaluated according to:
  - For systems running more than 12 hours a day
    - You should be interested in the Power to move the data onto and off the storage system
    - Is your load predominantly sequential?
      - **Bandwidth metric (MBS/Watt)** will help you to determine how effective is your power use. The larger this number is, the more data the system is pushing per watt
    - Is your load predominantly random?
      - **Workload metric (IOPS/Watt)** will help you determine how effective is your power use. The larger this number is the system is provides more operations per watt.
    - Independently on how long the system is idle it is always good to know what is your capacity per watt ratio

© 2011 Storage Networking Industry Association. All Rights Reserved.

Many industry wide efforts can help you become more green

- Emerald
  - A complete set of tests intended to measure power use ratios based on all previously mentioned metrics.
- SPC
  - Storage Performance Council mainly oriented to disk subsystems was the first industry association to add power to their benchmark
- The Green Grid
  - Working on a usage metric.

© 2011 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Background and Green Storage
  - Revisit what is a basic storage unit
  - What influences green storage
- Storage SNIA Taxonomy update

© 2011 Storage Networking Industry Association. All Rights Reserved.

Storage taxonomy update

- Need a taxonomy (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

© 2011 Storage Networking Industry Association. All Rights Reserved.

Taxonomy – Categories

Attribute	Category					
	Online	**Near Online**	**Removable Media Library**	**Virtual Media Library**	**Adjunct Product**	**Interconnect Element**
Access Pattern	Random/ Sequential	Random/ Sequential	Sequential	Sequential		
MaxTTFD (t)	t < 80 ms	t > 80 ms	t > 80 ms t < 5 min	t < 80 ms	t < 80 ms	t < 80 ms
User Accessible Data	Required	Required	Required	Required	Prohibited	Prohibited

- Six categories, covering most storage industry products

© 2011 Storage Networking Industry Association. All Rights Reserved.

Taxonomy – Categories

Category / Level	Online	Near Online	Removable Media Library	Virtual Media Library
Consumer/ Component	Online 1	Near Online 1	Removable 1	Virtual 1
Low-end	Online 2	Near Online 2	Removable 2	Virtual 2
Mid-range	Online 3	Near Online 3	Removable 3	Virtual 3
	Online 4			
High-end	Online 5	Near Online 5	Removable 5	Virtual 5
Mainframe	Online 6	Near Online 6	Removable 6	Virtual 6

Adjunct Product	Interconnect Element

- 23 total "buckets" covering the breadth of the industry

© 2011 Storage Networking Industry Association. All Rights Reserved.

Taxonomy – Online

- Most common storage systems

Attribute	Classification					
	Online 1	Online 2	Online 3	Online 4	Online 5	Online 6
Access Pattern	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential
MaxTTFD (t)	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms
User-Accessible Data	Required	Required	Required	Required	Required	Required
Consumer/ Component	Yes	No	No	No	No	No
Connectivity	Not specified	Connected to single or multiple hosts	Network-connected	Network-connected	Network-connected	Network-connected
Maximum Configuration	≥1	≥ 4	≥ 12	> 100	>400	>400
Integrated Storage Controller	Optional	Optional	Required	Required	Required	Required
Storage Protection	Optional	Optional	Required	Required	Required	Required
No SPOF	Optional	Optional	Optional	Required	Required	Required
Non-Disruptive Serviceability	Optional	Optional	Optional	Optional	Required	Required
FBA/CKD Support	Optional	Optional	Optional	Optional	Optional	Required

© 2011 Storage Networking Industry Association. All Rights Reserved.

Taxonomy – Near Online

Attribute	Classification: Near Online 1	Near Online 2	Near Online 3	Near Online 4	Near Online 5	Near Online 6
Access Pattern	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential	Random/ Sequential
MaxTTFD (t)	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms	t < 80 ms
User-Accessible Data	Required	Required	Required	Required	Required	Required
Consumer/ Component	Yes	No	No	No	No	No
Connectivity	Not specified	Connected to single or multiple hosts	Network-connected	Network-connected	Network-connected	Network-connected
Maximum Configuration	≥1	≥ 4	≥ 12	> 100	>400	>400
Integrated Storage Controller	Optional	Optional	Required	Required	Required	Required
Storage Protection	Optional	Optional	Required	Required	Required	Required
No SPOF	Optional	Optional	Optional	Required	Required	Required
Non-Disruptive Serviceability	Optional	Optional	Optional	Optional	Required	Required
FBA/CKD Support	Optional	Optional	Optional	Optional	Optional	Required

© 2011 Storage Networking Industry Association. All Rights Reserved.

Taxonomy – Removable

Attribute	Classification					
	Removable 1	Removable 2	Removable 3	Removable 4	Removable 5	Removable 6
Access Pattern	Sequential	Sequential	Sequential		Sequential	Sequential
MaxTTFD (t)	80ms < t < 5m	80ms < t < 5m	80ms < t < 5m		80ms < t < 5m	80ms < t < 5m
User-Accessible Data	Required	Required	Required		Required	Required
Maximum Drive Count	Not specified	4	≥ 5		≥ 25	≥ 25
Robotics	Prohibited	Required	Required		Required	Required
No SPOF	Optional	Optional	Optional		Optional	Required
Non-disruptive Serviceability	Optional	Optional	Optional		Optional	Required

© 2011 Storage Networking Industry Association. All Rights Reserved.

Taxonomy – Virtual

Attribute	Classification					
	Virtual1	**Virtual 2**	**Virtual 3**	**Virtual 4**	**Virtual 5**	**Virtual 6**
Access Pattern	Sequential	Sequential	Sequential		Sequential	Sequential
MaxTTFD (t)	t < 80 ms	t < 80 ms	t < 80 ms		t < 80 ms	t < 80 ms
User-accessible Data	Required	Required	Required		Required	Required
Maximum Configuration	12	>12	> 48		> 96	> 96
Storage Protection	Optional	Optional	Required		Required	Required
No SPOF	Optional	Optional	Optional		Optional	Required
Non-Disruptive Serviceability	Optional	Optional	Optional		Optional	Required

© 2011 Storage Networking Industry Association. All Rights Reserved.

Taxonomy – (Adjunct) & Interconnect

- Adjunct and Interconnect left to be defined

Adjunct Product	Interconnect Element

© 2011 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Background and Green Storage
  - Revisit what is a basic storage unit
  - What influences green storage
- Storage SNIA Taxonomy update
- Current status of
  - SNIA GreenTWG

© 2011 Storage Networking Industry Association. All Rights Reserved.

SNIA Green TWG Status

- Release of new measurement specification
  - SNIA Emerald Power Efficiency Measurement Specification
  - Updated to idle measurement specification
  - Added the active measurements listed before in this tutorial
- Support of the Emerald program
  - White papers
  - Tutorials
  - Training

© 2011 Storage Networking Industry Association. All Rights Reserved.

Storage Power – Idle

Average Power

Where:

$$PA_i(T) = \frac{\sum W_s}{n}$$

- $PA_i(T)$ is the AVERAGE POWER during test or test phase *i*, taken over a time interval of T seconds;
- W_s is power in watts measured at each sampling interval *s* taken during the time interval T;
- *n* is the number of samples gathered by the power meter during the time interval T;
- T = *n* * *s*.

Idle Metric

Power Efficiency, Ready Idle

Where:

$$EP_{RI} = \frac{C_R}{PA_{RI}(7200)}$$

EP_{RI} is the POWER EFFICIENCY metric for the READY IDLE TEST;

C_R is the RAW CAPACITY of the SUT (see 4.2.20);

$PA_{RI}(7200)$ is the AVERAGE POWER over the 2-hour MEASUREMENT INTERVAL for the READY IDLE TEST.

© 2011 Storage Networking Industry Association. All Rights Reserved.

Storage Power– Active

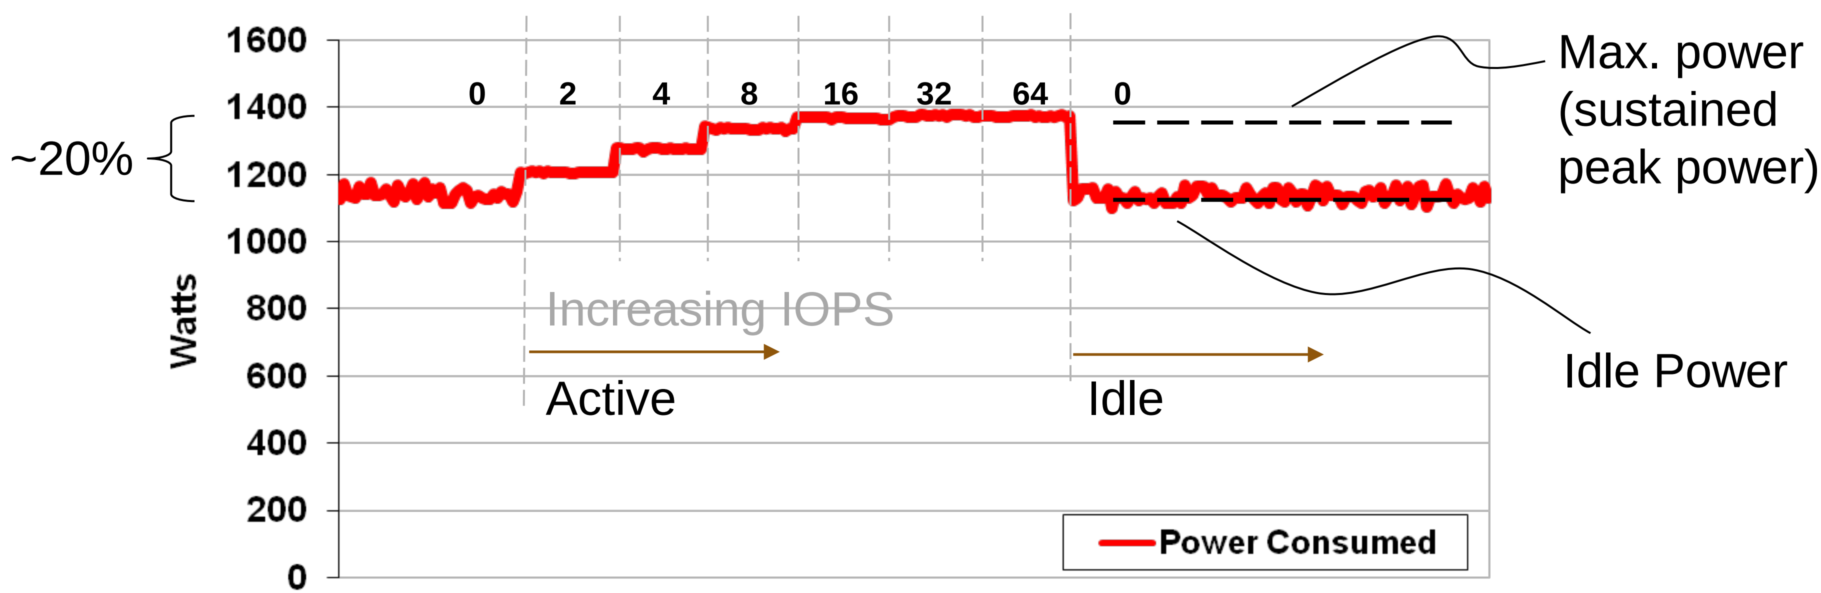

- Many variations in workloads
- Many variations in system configuration
- Takes us into the realm of benchmarking

© 2011 Storage Networking Industry Association. All Rights Reserved.

Desired Metric – “Productivity”

Education SNIA

Many possible definitions – must balance simplicity against applicability

- “typical workload”, with levels

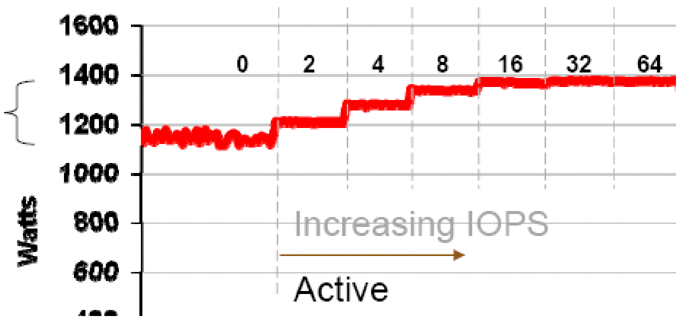

- detailed performance benchmarks – results/W

Storage Performance Council

Defining, administering, and promoting industry-standard, vendor-neutral benchmarks to characterize the performance of storage products

spec

Standard Performance Evaluation Corporation

the green grid
get connected to efficient IT

- “four corners”, maximum performance, maximum power

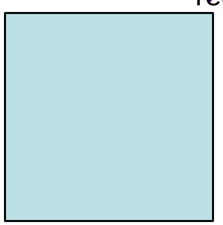

- The Green Grid Productivity Proxy Proposals example – Proxy #4 – bits/kilowatt-hour

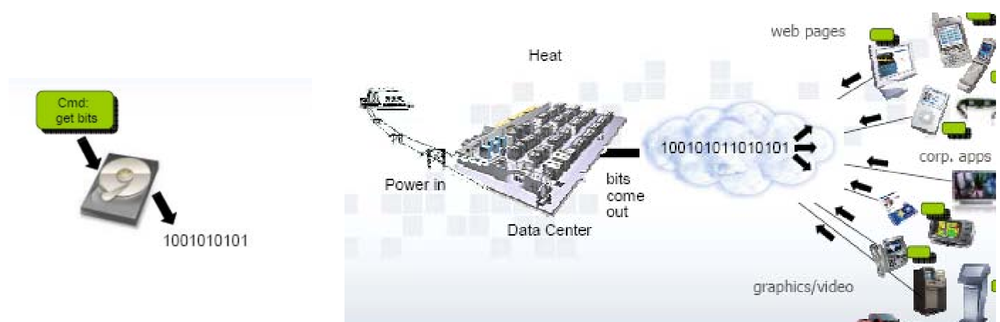

© 2011 Storage Networking Industry Association. All Rights Reserved.

Storage Power – Active Online

IO Profile	IO Size (KiB)	Read/Write Percentage	IO Intensity	Transfer Alignment (KiB)	Access Pattern
Mixed Workload 1 (i=MW1)	8	70/30	100	8	Random
Mixed Workload 2 (i=MW2)	8	70/30	25	8	Random
Random Write (i=RW)	8	0/100	100	8	Random
Random Read (i=RR)	8	100/0	100	8	Random
Sequential Write (i=SW)	256	0/100	100	256	Sequential
Sequential Read (i=SR)	256	100/0	100	256	Sequential

- Active Metric T=1800

Periodic Power Efficiency

Where:

$$EPP_i(T) = \frac{O_i(T)}{PA_i(T)}$$

EPP_I (T) is the PERIODIC POWER EFFICIENCY during test or test phase *i*, taken over a time interval of T seconds;

O_i(T) is the OPERATIONS RATE during test or test phase *i*, taken over the same time interval of T seconds;

PA_i(T) is the AVERAGE POWER during test or test phase *i*, taken over the same time interval of T seconds.

© 2011 Storage Networking Industry Association. All Rights Reserved.

Complications

- Max power =/= Max performance

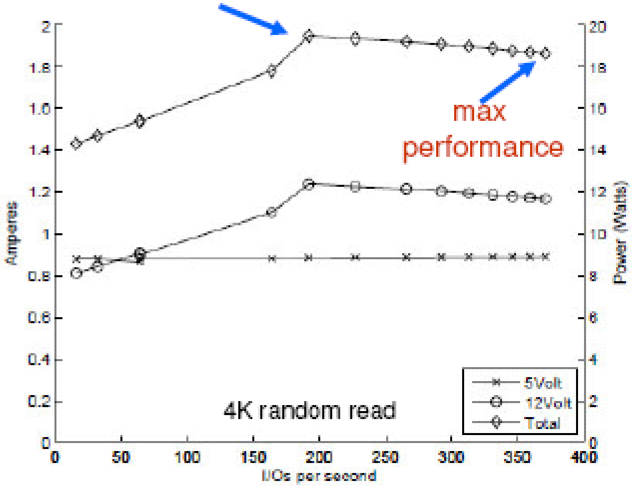

Single disk drive power profile

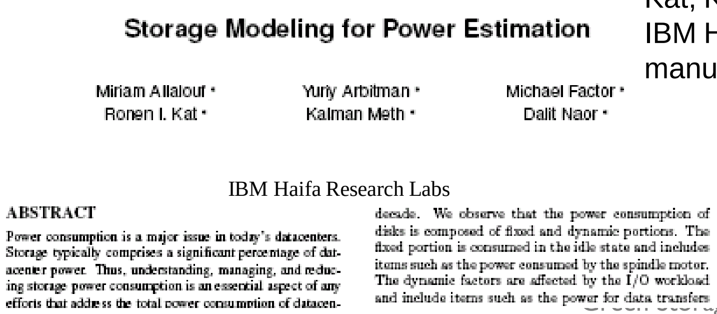

Storage Modeling for Power Estimation

Miriam Allalouf • Yuriy Arbitman • Michael Factor •
Ronen I. Kat • Kalman Meth • Dalit Naor •

IBM Haifa Research Labs

ABSTRACT

Power consumption is a major issue in today's datacenters. Storage typically comprises a significant percentage of datacenter power. Thus, understanding, managing, and reducing storage power consumption is an essential aspect of any efforts that address the total power consumption of datacen-

decade. We observe that the power consumption of disks is composed of fixed and dynamic portions. The fixed portion is consumed in the idle state and includes items such as the power consumed by the spindle motor. The dynamic factors are affected by the I/O workload and include items such as the power for data transfers

"Storage Modeling for Power Estimation", Miriam Allalouf , Yuriy Arbitman, Michael Factor, Ronen I. Kat, Kalman Meth, and Dalit Naor; IBM Haifa Research Labs; manuscript; March 2009

- Significant whole-system considerations

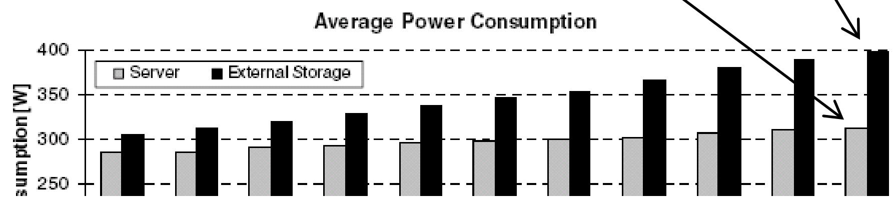

SPECweb 2005 (banking) + storage

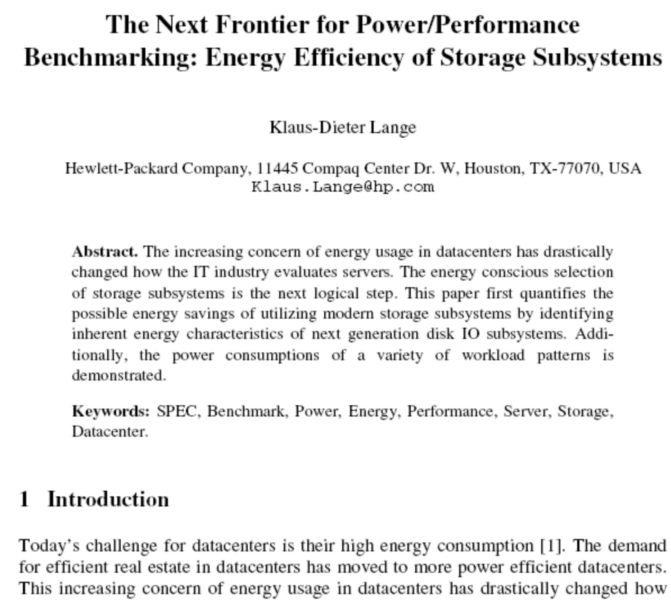

The Next Frontier for Power/Performance Benchmarking: Energy Efficiency of Storage Subsystems

Klaus-Dieter Lange

Hewlett-Packard Company, 11445 Compaq Center Dr. W, Houston, TX-77070, USA
Klaus.Lange@hp.com

Abstract. The increasing concern of energy usage in datacenters has drastically changed how the IT industry evaluates servers. The energy conscious selection of storage subsystems is the next logical step. This paper first quantifies the possible energy savings of utilizing modern storage subsystems by identifying inherent energy characteristics of next generation disk IO subsystems. Additionally, the power consumptions of a variety of workload patterns is demonstrated.

Keywords: SPEC, Benchmark, Power, Energy, Performance, Server, Storage, Datacenter.

1 Introduction

Today's challenge for datacenters is their high energy consumption [1]. The demand for efficient real estate in datacenters has moved to more power efficient datacenters. This increasing concern of energy usage in datacenters has drastically changed how

"The Next Frontier for Power/Performance Benchmarking: Energy Efficiency of Storage Subsystems" Klaus-Dieter Lange; SPEC Benchmark Workshop 2009; January 2009

© 2011 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Background and Green Storage
  - Revisit what is a basic storage unit
  - What influences green storage
- Storage SNIA Taxonomy update
- Current status of
  - SNIA Green TWG
  - EPA

© 2011 Storage Networking Industry Association. All Rights Reserved.

ENERGY STAR for Data Center Storage

Education SNIA

Date	Document	Effective Date (Tier 1)	Pages
April 2009	Stakeholder Letter	TBD	2
June 2009	Framework Document	March 2010	12
20 July 2009	1st Stakeholder Meeting	March 2010	--
15 October 2009	Test Procedure Meeting	TBD	--
2 February 2010	2nd Stakeholder Meeting	TBD	--
Jan/Feb/March 2010	1st Round Data Collection		
9 April 2010	Draft 1	TBD	18
20 July 2010	3rd Stakeholder Meeting	TBD	--
April 2011	Supplemental Data Collection	TBD	
July 2011	Stakeholder Meeting	July 18	
October 2011	Stakeholder Meeting		

"this is one of the most complex areas we have tackled to date" – Andrew Fanara, Team Leader, EPA ENERGY STAR Product Specifications, July 2009 in San Jose

© 2011 Storage Networking Industry Association. All Rights Reserved.

Observations at Stakeholder Meetings

- EPA ENERGY STAR
  - Open and willing to listen
  - Learning the technical and business details
  - Collaboration with industry is a key goal (while meeting the EPA needs and mission)
- Audience
  - Wide cross-industry participation – vendors & suppliers
  - SPC, The Green Grid, Wikibon, PG&E, Climate Savers
  - Some end-user participation both via phone and in person

© 2011 Storage Networking Industry Association. All Rights Reserved.

Goals for Data Center Storage

ENERGY STAR

- Encourage widespread adoption of energy efficient hardware and software strategies,
- Provide purchasers with the means to identify the most energy efficient enterprise storage solutions for their specific end-use application, and
- Provide tools and information to designers and mangers looking to improve the efficiency of data center operations

EPA

9

© 2011 Storage Networking Industry Association. All Rights Reserved.

Litmus Test

- There are numerous product features, functions, and data management strategies that enable energy savings in data center storage.
- There is only one end result that matters: The ability to do more useful work, while consuming fewer resources, in a verifiable and quantifiable manner.

13

© 2011 Storage Networking Industry Association. All Rights Reserved.

Agenda

- Background and Green Storage
  - Revisit what is a basic storage unit
  - What influences green storage
- Storage SNIA Taxonomy update
- Current status of
  - SNIA Green TWG
  - EPA
  - SNIA
    - Green Efforts
    - Emerald Program (Quick introduction)

© 2011 Storage Networking Industry Association. All Rights Reserved.

SNIA Green Efforts

- SNIA Green Storage Initiative (GSI) and SNIA Green Storage Technical Work Group (TWG)
  - on-going efforts to develop data-driven green standards & metrics
  - power measurements at multi-vendor "unplugged" fests
  - alliances with other active green organizations
    - (The Green Grid, 80PLUS, Climate Savers, DMTF, SPEC, SPC)
  - collaboration with EPA on the ENERGY STAR program
- Whitepapers / workshops
  - three tutorials at SNW; online tutorials available (www.snia.org/education/tutorials)
  - white papers from GSI

SNIA Green Storage Initiative

towards energy-efficient storage networking solutions

© 2011 Storage Networking Industry Association. All Rights Reserved.

SNIA Emerald™ Program

SNIA Emerald™ Program

- The SNIA Emerald ™ Program provides a public accessible repository of storage ven[dor] storage system power usage and efficiency measurement data. The measurement data is generated through the use of well-defined and proven testing procedures prescribed in the SNIA Emerald™ Energy Measurement Specification. The Emerald Program repository includes downloadable test data reports for each vendor opting to participate in the SNIA Emerald™ Program. The report includes information related to system power including system configuration details such as storage device types, RAS features and their configuration, and power supply types.
- The program is open to the industry at large, including non-members of SNIA.
- The measurement data quantifies a storage system power usage for several types of workloads.
- The measurement data can help IT professionals make storage platform selections as part of an overall Green IT and Sustainability objective.
- The SNIA Emerald™ Program is can be accessed at www.sniaemerald.com
  - Access the repository of vendor submitted test measurement results
  - Submit test measurement results
  - Obtain a copy of the SNIA Emerald ™ Energy Measurement specification
- Visit the SNIA Booth at SNW to learn more

© 2011 Storage Networking Industry Association. All Rights Reserved.

Summary

Education SNIA

- "Green is good" – for multiple reasons
  - a great engineering problem – doing more with less
  - saves money – great investment payback
  - helps save the planet – significant leverage
- Get involved with SNIA Green efforts
  - weekly discussions, regular face-to-face & data sharing (TWG)
    - Get a power meter and try the measurement spec on your own systems (!)
  - education and promotion (GSI / SNIA Emerald)
  - promote these industry-wide efforts within your company
- Learn about wider green technology and opportunities
  - online resources; workshops by SNIA, EPA, The Green Grid
- Share your experience / knowledge

© 2011 Storage Networking Industry Association. All Rights Reserved.

References

- SNIA Green Storage Initiative – www.snia.
- SNIA Emerald - www.sniaemerald.com
- The Green Grid – www.greengrid.org
- EPA ENERGY STAR™ (Data Center Storage – www.energystar.gov/index.cfm?c=n ecs.enterprise_storage
- DOE Federal Energy Management Program – eere.energy.gov/femp
- Power calculators at various venc

© 2011 Storage Networking Industry Association. All Rights Reserved.

Q&A / Feedback

Education SNIA

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

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

SNIA Education Committee

Erik Ridel	SNIA Green Storage Initiative members
Carlos Pratt	SNIA Green Storage TWG members
Patrick Stanko	
Patrick Chu	
Matthew Brisse	
David Reinsel	
Edgar St.Pierre	
Alan Yoder	
Wayne Adams	
SW Worth	
Herb Tanzer	

© 2011 Storage Networking Industry Association. All Rights Reserved.