SNIA Green Storage Overview & Proposal

Erik Riedel (EMC), Chair, Green Storage TWG
Leah Schoeb (Sun), Chair, Green Storage Initiative

SNIA
Green Storage Initiative

15 October 2009
(revision 2)

SNIA Green Groups

SNIA
Advancing storage & information technology

17 face-to-face meetings plus weekly conference calls since August 2007, regular 20+ participants from across the industry

- Green Storage TWG
  - Erik Riedel (EMC), chair
- Green Storage Initiative
  - Leah Schoeb (Sun), chair

- Design/Spec Metrics Sub-group
  - Patrick Stanko (Sun), chair
- Power Supply Sub-group
  - Don Goddard (NetApp), chair
- Capacity Optimization Sub-group
  - Alan Yoder (NetApp), chair
- Operational Metrics Sub-group (*)
  - Shinobu Fujihara (IBM), chair
- Infrastructure Sub-group (*)
  - Steve Wilson (Brocade), chair

Storage Power – Idle vs. Active

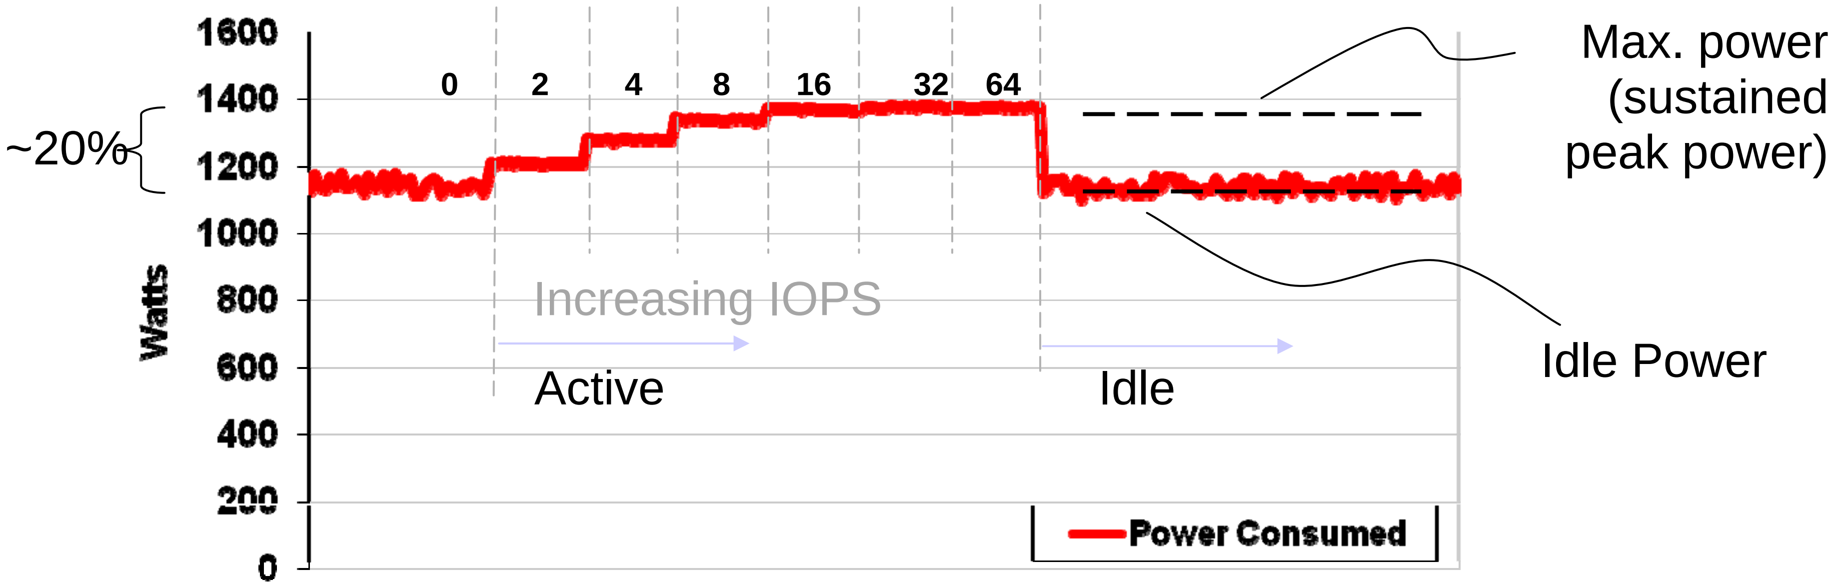

- Ideally, systems would consume minimum power in all modes
  - Example system consumes **significant power in idle** (**80% of max**)
- % of time in Idle versus Active depends on storage type, application and workloads; available optimizations will vary
- Power consumed not linearly proportional to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

Storage Power – Idle

Equation 6-1: Average Idle Power

$$P_i = \frac{\sum W_i}{n}$$

Where:

- *P_i is average idle power*
- *W_i is power in watts measured in each sampling interval i*
- *n is the number of samples gathered by the power meter during the measurement interval.*

SNIA Green Storage Power Measurement Technical Specification
www.snia.org/tech_activities/publicreview
[Green Power v0.0.18 DRAFT]
January 2009

- Idle Metric

Equation 7-1 SNIA Idle Power Metric

$$P = \frac{C}{P_i}$$

Where:

- *P is the SNIA Idle Power Metric*
- *C is the total capacity of the SUT*
- *P_i is the average idle power*

Desired Metric – "Productivity"

SNIA
Advancing storage & information technology

- "typical workload", with levels

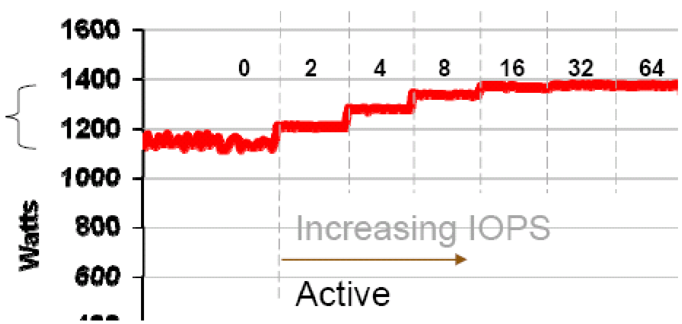

- "four corners", maximum performance, maximum power

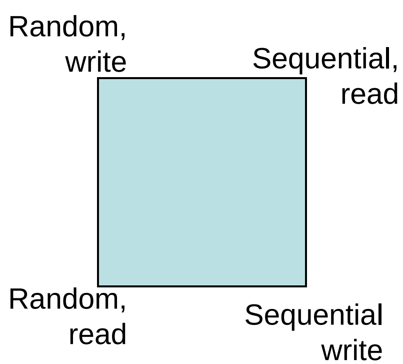

- Detailed Performance Benchmark – results/W

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 SM
get connected to efficient IT

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

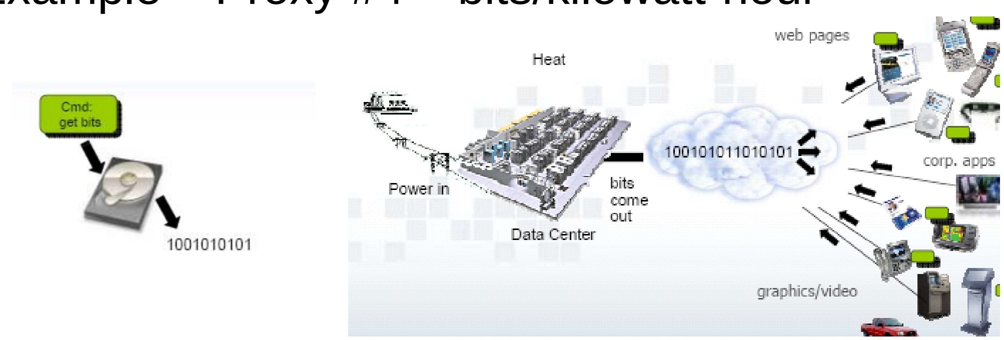

Complications

SNIA
Advancing storage & information technology

• Max power =/= Max performance

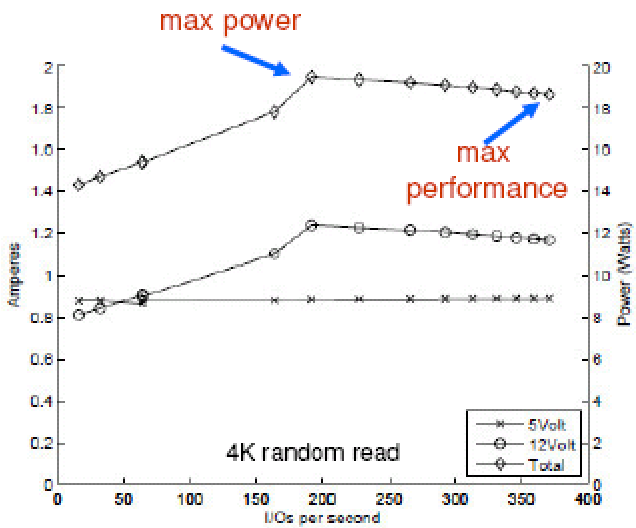

Single disk drive power profile

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

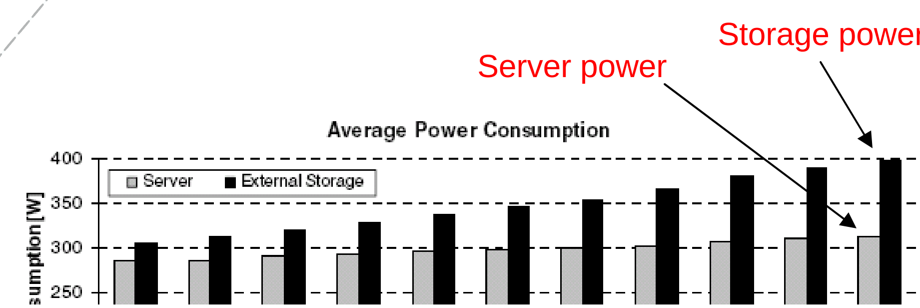

SPECweb 2005 (banking) + storage

• Significant whole-system considerations

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

SNIA Measurement Proposal – Draft

SNIA
Advancing storage & information technology

- Idle Metric
  - GB / W (24-hour average W, raw capacity)
  - 24-hour test preceded by a short conditioning phase
  - www.snia.org/tech_activities/publicreview [Green Power v0.0.18 DRAFT]
- Active Metrics
  - 4-corners test + 2-point ramp
  - random read (8 KB @ 30ms response time, 10 minutes) – IOPS/W
  - random write (8 KB @ 30ms response time, 10 minutes) – IOPS/W
  - sequential read (256 KB @ 30ms response time, 10 minutes) – MB/s/W
  - sequential write (256 KB @ 30ms response time, 10 minutes) – MB/s/W
  - 70/30 read/write random 25% maximum (8 KB @ 30ms, 10 minutes) – IOPS/W
  - 70/30 read/write random 100% maximum (8 KB @ 30ms, 10 minutes) – IOPS/W
  - can be used as conditioning phase for idle
  - many details still under discussion [preliminary proposal only]

Applicability of Metrics

SNIA
Advancing storage & information technology

	Online	Near Online	Removable Media	Virtual Media Library	Appliance	Interconnect
Idle	Y	Y	Y	Y	n/a	n/a
2-corners MB/s	Y	Y	Y	Y	Y*	Y
2-corners IOPS	Y	n/a	n/a	n/a	Y*	Y
Ramp IOPS	Y*	n/a	n/a	n/a	Y*	Y

Early Results – September 2009

- Idle Metric
  - 20 systems

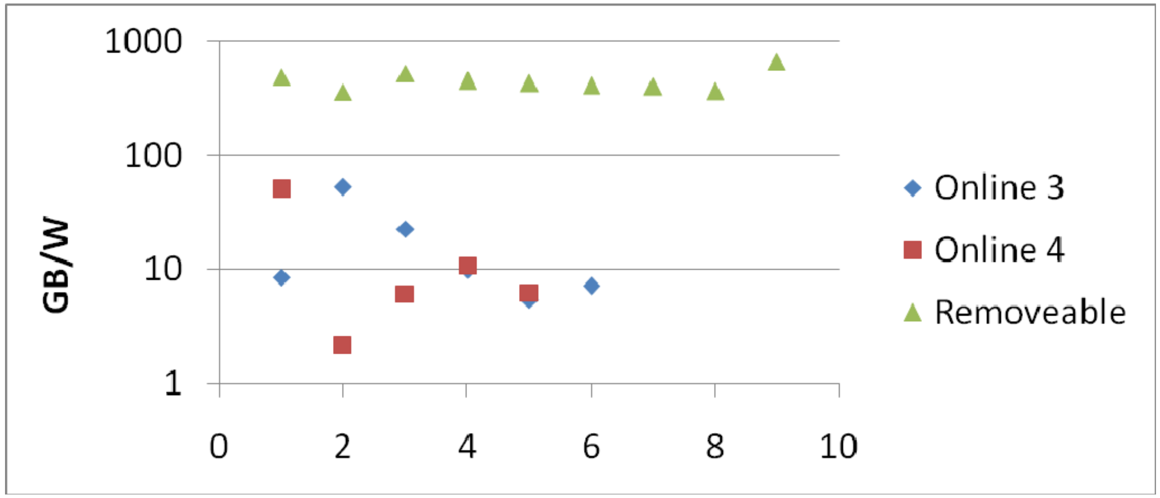

- Active Metrics
  - ~12 systems
  - mixed taxonomy (online only)

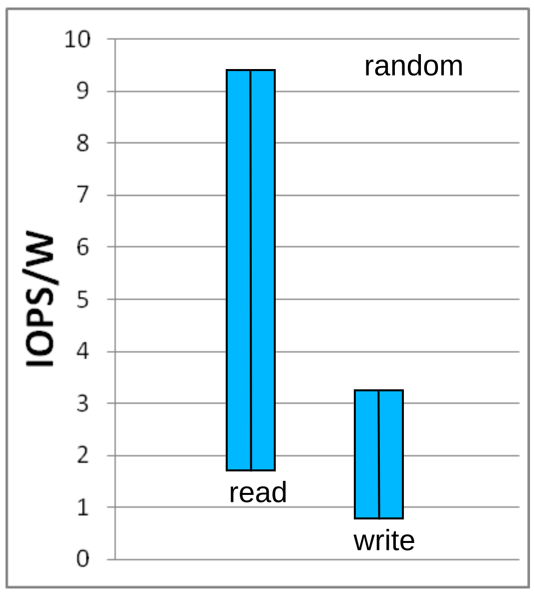

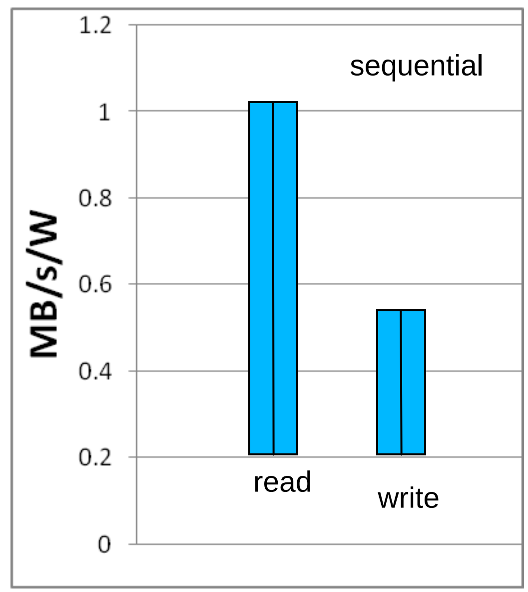

Add'l Considerations

SNIA
Advancing storage & information technology

- Mixed Workloads
  - Mixed workloads could form a more robust representation of customer environments (accepting a wide range in practice)
- Varying Workload Intensity Tests (Ramp)
  - Real Data Center workloads vary throughout the day
  - May capture variations in power efficiency (like power supplies)
  - Some of our measurements have shown variance between maximum power and maximum performance (mixed vs. 4-corners)
- Our testing has been with a 4-level ramp (25%, 50%, 75%, 100%) at 70/30 read/write 8 KB @ 30ms response time
  - Inconclusive to date (measurement continues)

Reliability, Availability, Serviceability (RAS)

Goal – do not lose data and have system available all the time

Power	Features
High	Dual Controllers or No-SPOF Controller; Mirroring (local and remote, sync or async)
Medium/High	RAID 1, 4/5, 6; Snapshots (full and delta); RAID/Disk scrubbing
Medium	Multi-pathing ; Disk sparing; Dual robotics; Background media scan / maintenance (drive level); Dual Power Supplies
Medium/Low	Remote management ; HARD/DIF (end-to-end data validation); RAID rebuild on drive failure; Hot swap drives
Low	NVRAM/Batteries/UPS (power fail); Concurrent microcode update; Concurrent OS update; Redundant PDU/Power ; Remote diagnostics ; ECC validation btwn modules; Lost write protection; Encryption; Automatic load balancing
Savings	Variable-speed fans (Medium); Head-unloading (Low)

Investigating quantification for Highs & Mediums

Capacity Optimization

SNIA
Advancing storage & information technology

- Software-level optimization features are critical to Data Center-wide power savings
  - For example, if data deduplication reduces the amount of stored data by 2x, then the energy use is 1/2
- Wide variety of technologies, wide variety of implementations
  - Challenging to baseline, compare, and quantify
- Key initial focus areas of study
  - Data Deduplication
  - Thin Provisioning
  - Delta Snapshots

Meetings in 2010 (Planned)

SNIA
Advancing storage & information technology

	Event	Location
October 15, 2009	EPA Workshop @ SNW	Phoenix
November 4-5, 2009	TWG Face-to-face (TSG/LISA)	Baltimore
January 25-26, 2010	TWG Face-to-face (Symp)	San Jose
February 22, 2010	USENIX SustainIT Conference	San Jose
February 23-26, 2010	USENIX FAST Conference	San Jose
March 15-16, 2010	TWG Face-to-face (Symp)	San Antonio
April 22, 2010	Earth Day	Worldwide
May 17-18, 2010	TWG Face-to-face (TSG)	Colorado Springs
July 12-13, 2010	TWG Face-to-face (Symp)	San Jose
August 30-31, 2010	TWG Face-to-face (TSG)	TBD
September 2010	Storage Dev Conference	Santa Clara
November 9-10, 2010	TWG Face-to-face (TSG/LISA)	San Jose

Thanks for Everyone's Time
Q & A

SNIA
Green Storage Initiative