



Go further, faster

Power Management in Storage Systems

Kaladhar Voruganti

Technical Director
CTO Office, Sunnyvale

June 12, 2009





Outline

- Power Consumption Background in Data Centers and Storage Systems
- Power Management Strategies for Storage Systems
- Power Management Metrics
- Impact of Server Power Management Strategies on Storage Systems

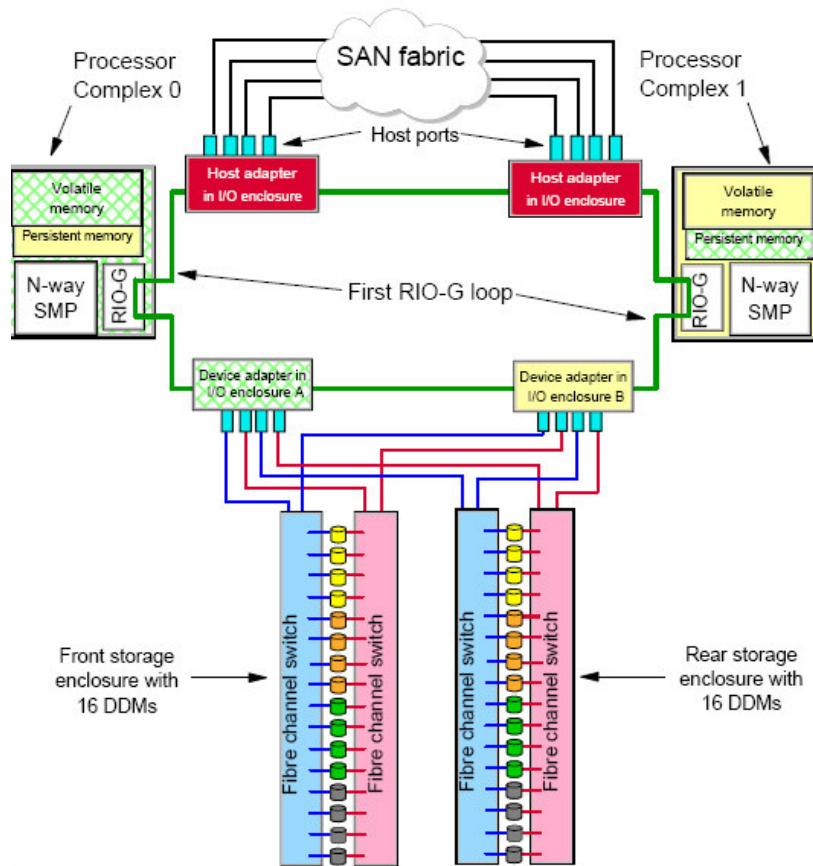


Storage Consumption Constantly Increasing

- More types of information are being digitized and stored persistently (emergence of newer types of applications)
- Data is being stored persistently for longer periods of time (for legal and sentimental reasons)
- More people are persistently storing their information (computer usage globally is increasing)


What is a Storage Controller?

Traditional Storage Controller



Source: IBM Redbook DS8000


Cluster of thousands of Google Servers



Source: CNET News April/2009



Where Does the Power Go in a Data Center?



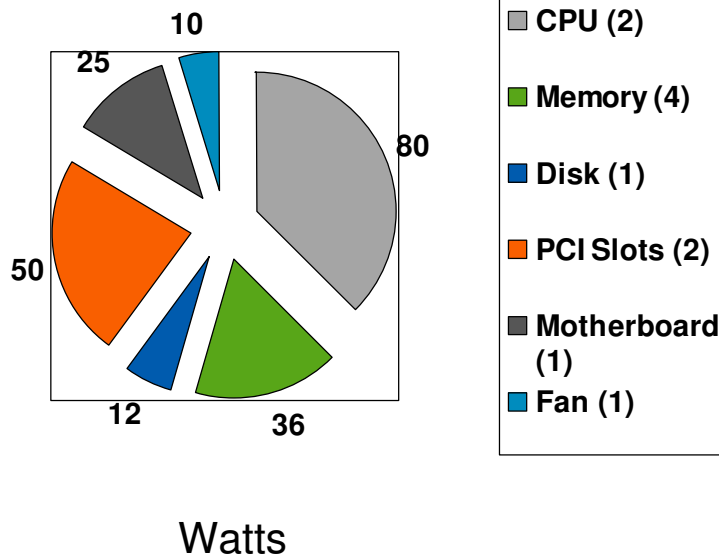
Percentages

Source: IDC Report 2008



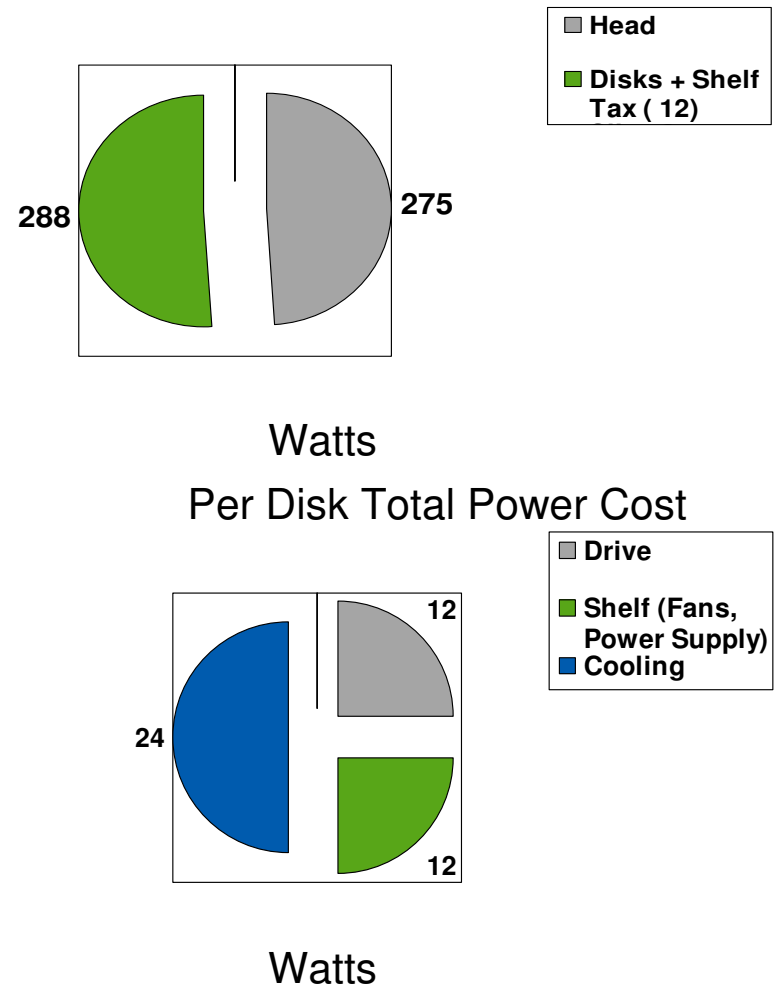
Where Does the Power Go in a Storage Controller?

Google Box Power Consumption



Source: Google Paper, ISCA 2007

Source: NetApp Internal Study
Storage Controller Box Consumption





Storage Power Management Strategies

- Hardware
 - Can select the appropriate storage architecture
 - Can select the appropriate storage system
 - Can select the appropriate hardware features

- Software
 - Storage Efficiency
 - Migration and Spinning/Shutting Disks Down



Power Management Strategies (Hardware)

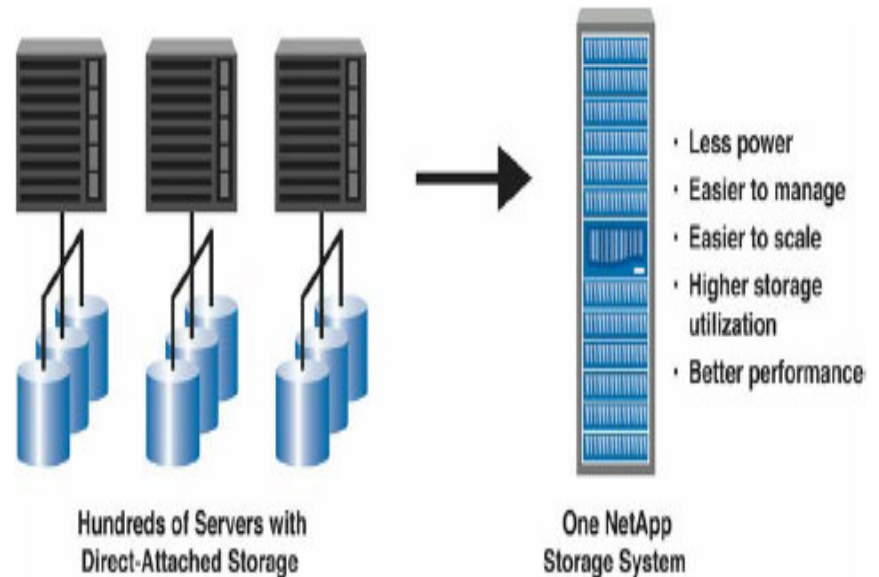
- Hardware Techniques
 - Architectural Level
 - DAS versus Storage Controllers
 - Single Node Battery versus UPS Technology
 - Storage Box Level
 - Disks versus SSDs
 - Efficient Power Sources
 - Use of higher capacity disks
 - Use of lower RPM disks



DAS versus Storage Controllers

■ DAS Storage

- Application and storage are co-located
- Power consumption is efficient if box is used to run applications (e.g. m reduce applications)
- Power consumption is r efficient if the cluster nodes are only used to serve storage
- Low powered shared-nothing nodes are being proposed for archival storage (Pergamum work from UC Santa Cruz)





Single Node Battery Versus Centralized UPS

- Large UPSs can reach 92 to 95 percent efficiency at full load
 - Operating at lower load results in inefficiency which results in the generation of heat
 - Need cooling to remove the heat from the data center
- By having 12 volt battery at each of the storage nodes, Google is able to get 99.9 percent efficiency

Source: CNET.com Article on Google, April, 2009



Power Management Strategies (Hardware)

- Hardware Techniques
 - Architectural Level
 - DAS versus Storage Controllers
 - Single Node Battery versus UPS Technology
 - Storage Box Level
 - Disks versus SSDs
 - Efficient Power Sources
 - Use of higher capacity disks
 - Use of lower RPM disks



Efficient Power Supply

Source: NetApp Internal Study

- Want Power Supplies that are Efficient for a wider range of load
 - These cost more
 - But offer savings in power consumption due to less heat generation (less cooling required)
- If there are multiple Power Supplies usually each one is run at lesser load, and thus, has higher inefficiency

Load	Less Efficient Power Supply (Efficiency)	More Efficient Power Supply (Efficiency)
10 % Load	50 % Eff	80% Eff
50 % Load	75% Eff	80% Eff
100 % Load	75% Eff	90% Eff



Higher Capacity Disks

Source: NetApp White Paper: WP 7010-0207

	Old Systems	New Systems	Improvement
# of systems	11 Old Systems: 4 F880 3 F810 2 F820 1 F825 1 F840	1 FAS 3020 with 3 disk shelves	11:1
Power* (kW HRs) * Does not include power for cooling.	113,651	20,915	81% Decrease
Space (Cubic Feet)	63.0	4.3	93% Decrease
Capacity (GBs)	9,776	14,000	16% Increase



Power Management Strategies (Software)

- Software Techniques
 - Storage Efficiency
 - Reduce overhead per amount of usable storage
 - Efficient Copies
 - Data De-duplication/Compression
 - Thin Provisioning
 - Number of copies
 - Protection Mechanism
 - Migration and disk shutdown/spin-down



Storage Efficiency Techniques

Source: Oliver Wyman Article, Dec 2007
“Making Green IT a Reality”

- Reducing Storage Overhead
- Thin Provisioning
- Efficient Protection Mechanisms
- Consolidation of Protocols
- Efficient Copies
- De-duplication/Compression

Capability	Benefits
Snapshot	20% storage overhead compared to 100% for full copies or BCV
Thin Provisioning (FlexVol)	20% - 33% savings by growing and shrinking volume sizes on demand
RAID 6 Implementation (RAID-DP™)	14% - 17% overhead compared to 100% for RAID 1
Capability	Benefits
Multi-protocol (Unified Storage) & FC/SATA Drives	2-3x savings if running multiple protocols with low util rates; up to 50% savings with SATA instead of FC
Multiple Writeable Snapshot Copies (FlexClone®)	Up to 66% savings if creating five copies of original data compared to BCV or full copies
Deduplication	10% - 80% space savings depending on data set



Disk Spin-down/Shut-down Techniques

- Migrate less accessed data to lower tiers of storage and shut-down disks
- Archival data can be stored on disks that are shut down because of write-once and read-maybe properties
- Difficult to shut-down disks for those applications that have strict latency requirements and have long-tailed distribution access patterns
- Spinning things down to lower RPM and then spinning them up is difficult because constant spinning up disks can actually result in higher power consumption
- COPAN has shown roughly 5x times power savings compared to normal storage controllers in cases where things can be shut down.
 - Very dense packaging than traditional storage controllers
 - Keeps application meta-data in cache
 - Spins disks down but keeps the electronics up



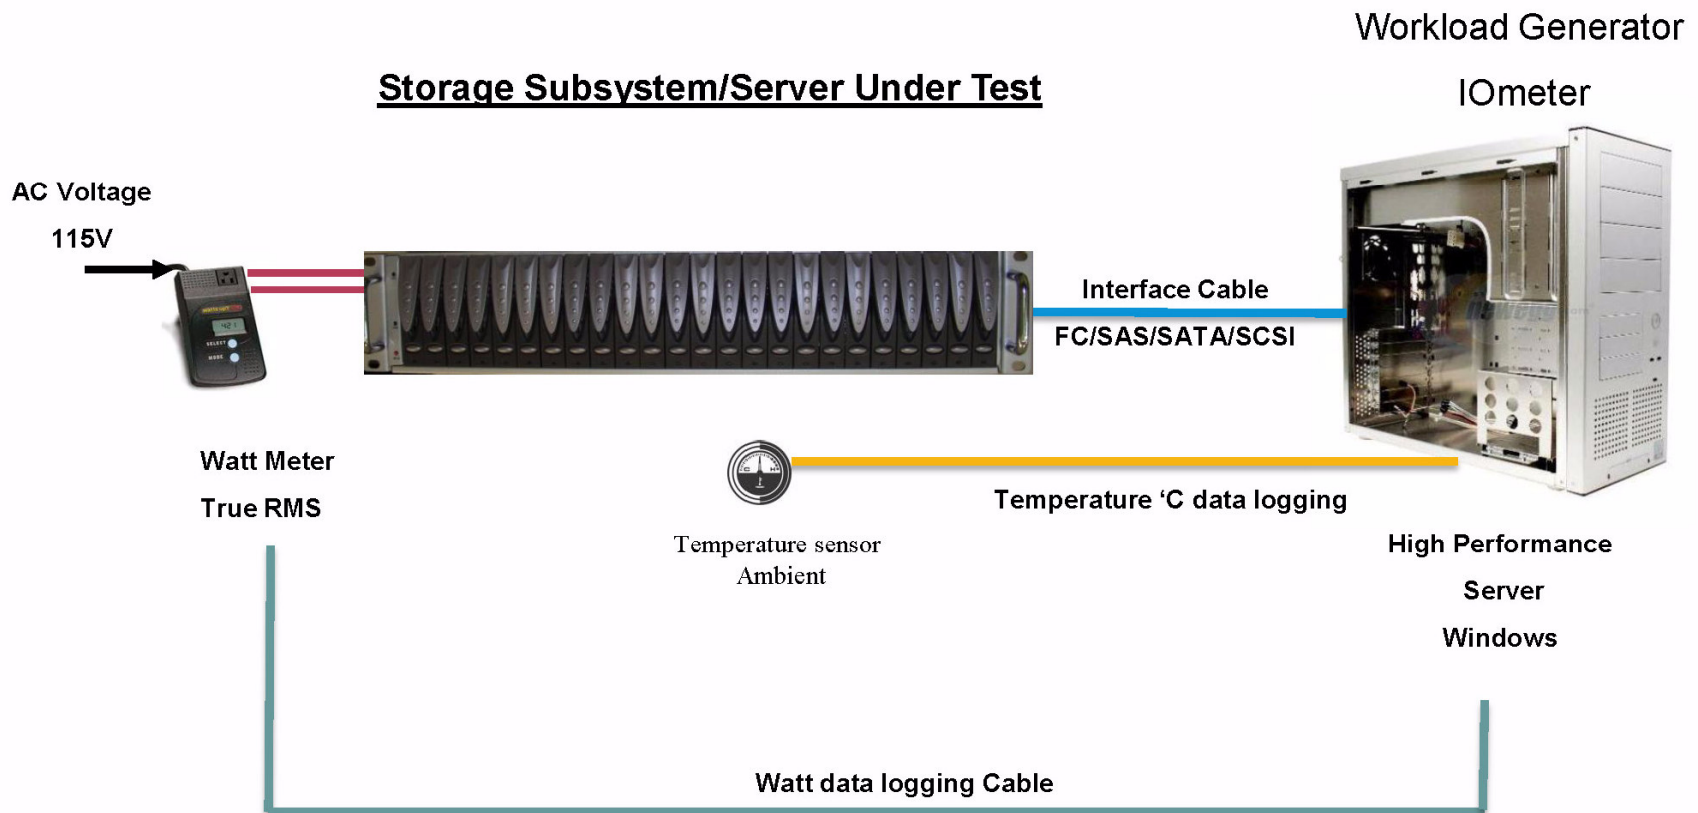
SNIA Green Storage Initiative Device Classification

Attribute	Category					
	Online	Near online	Removable Media	Virtual Media Library	Appliance	Interconnect
Access Pattern	Random	Random	Sequential write	Sequential write		
MaxTTD (t) ¹	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



SNIA Configuration & Workload Example

Configuration & Workload Example





SNIA Green Power Profile Example

■ Green Power Profile

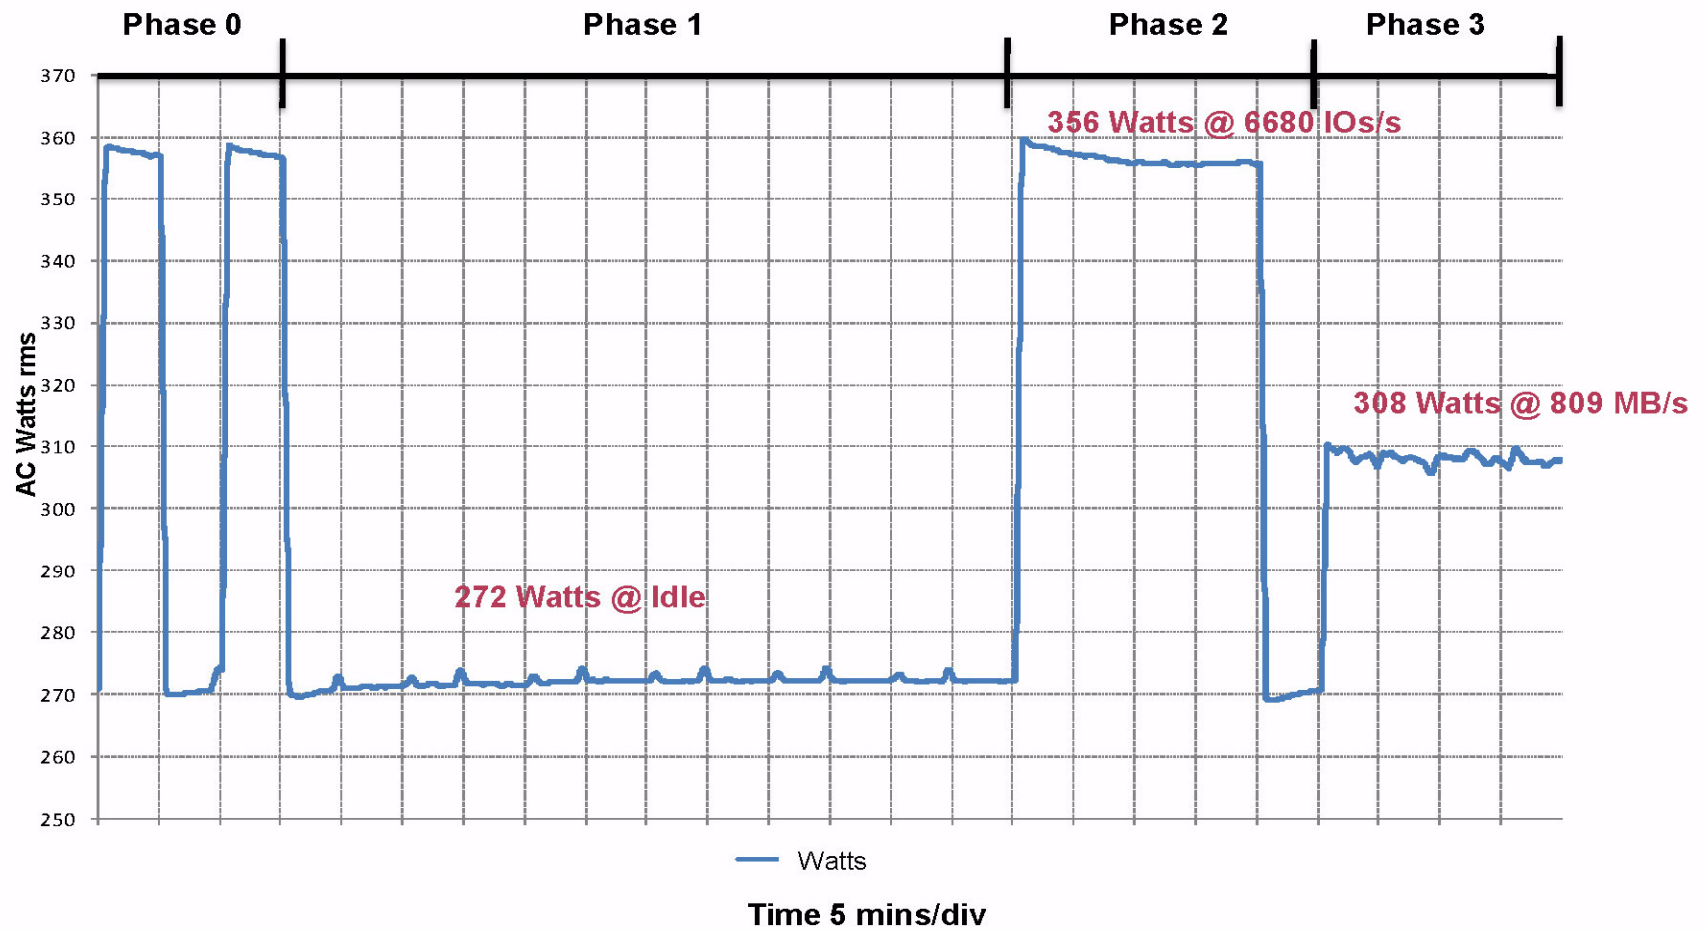
- Phase 0
 - Preconditioning phase no power or IO measurements are necessary
 - 5 minutes maximum OLTP workload
 - 5 minutes no workload (idle)
 - 5 minutes maximum OLTP workload
- Phase 1
 - Idle measurement phase . No user initiated commands allowed
 - 60 minutes measurement period
- Phase 2
 - OLTP workload
 - 20 minutes measurement period
 - 5 minutes rest no OLTP workload
- Phase 3
 - Sequential Throughput 50% Read 50% Write 1MB transfer
 - 20 Minutes measurement period



SNIA Green Power Profile Example

Green Power Profile

High End Storage Subsystem Example End User Capacity = 1.6TB





Power Management Metrics

I/O Performance (OLTP Type Workloads)	IOPs/Watt
Capacity	GB/Watt
Usable Capacity	GB/Watt
Sequential I/O Throughput	MBps/Watt
Availability	RTO/Watt 9s/Watt



Impact of Server Virtualization on Storage

- Server Virtualization is being used to consolidate physical servers to obtain cost, space and power efficiencies

- Impact of Server Virtualization on Storage
 - Need for shared storage
 - Need for efficient storage with de-duplication
 - I/O interference due to Multi-tenancy
 - Mis-match of Hypervisor and Storage constructs
 - Block Mis-alignment due to layers of storage software



Conclusion

- Key takeaways
 - Storage Efficiency is the primary mechanism for saving power by having fewer number of disks
 - Flash is emerging as a power efficient alternative to disks (price is still an issue)
 - Shutting-down disks is only attractive for archival storage
- Need for:
 - Power Management Metrics Standards
 - Power aware storage management tools