Power Consumption in Enterprise-Scale Backup Storage Systems

Zhichao Li[†], Kevin Greenan[‡], Andrew Leung[‡], Erez Zadok[†]

[†]*Stony Brook University*
[‡] *Backup Recovery Systems Division, EMC*

10th USENIX Conference on File and Storage Technologies (FAST'12)

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

How To Design Power Management?

- Disk backup prime for power management
  - Stores lots of cold data
  - Competes with low-power tape

- No previous power measurements
  - Good design is guided by measurement
    - Often just follow vender spec sheets
  - Must guess and use assumptions
    - Often assume disks will dominate power

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Methodology

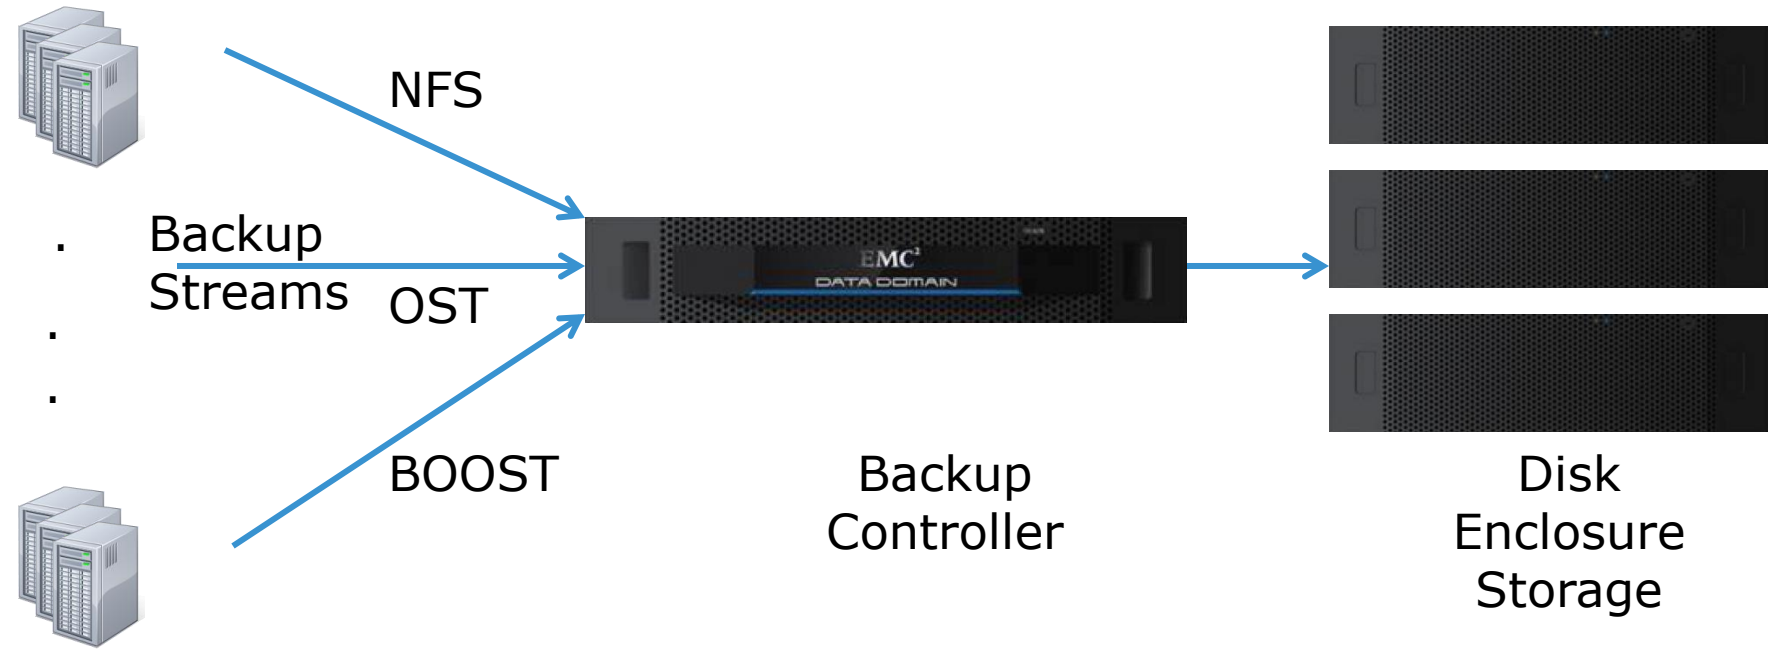

Architecture

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Methodology

- Measured 4 enterprise backup controllers

Controller	DD880	DD670	DD860	DDTBD
Ship year	2009	2010	2011	Experimental
Market	High end	Low end	High end	High end
# CPUs	2	1	2	4
RAM (GBs)	64	16	72	256

Controller hardware summary

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Methodology

- Measured 2 enterprise backup enclosures

Enclosure	ES20	ES30
Ship year	2006	2011
# Disks	16	15
# SAS cards	2	2
# Fans	2	4

Enclosure hardware summary

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Methodology

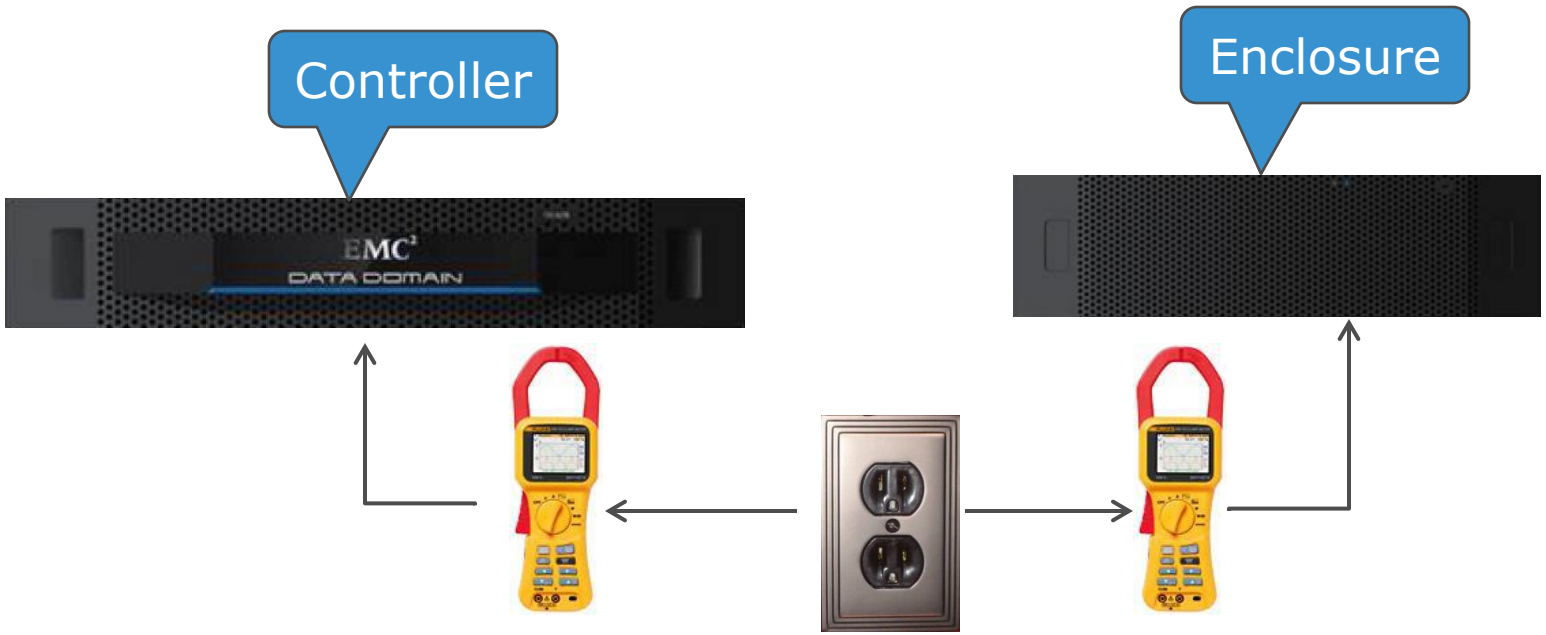

Measurement

Not measured: networking, cooling, internal subcomponents

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Idle power consumption

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Idle Power Consumption

Controller	DD880	DD670	DD860	DDTBD
Idle (W)	555	225	261	778

- DDTBD consumes more than 100 2TB drives
- DD880 2x more power than DD860

Controller	DD880	DD670	DD860	DDTBD
Usable TBs	192	76	192	1152
W/Usable TB	2.89	2.96	1.35	0.675

- Newer HW generations are more efficient

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Idle Power Consumption

Enclosure	ES20	ES30
Idle (W)	278	179

- Again, newer HW is more efficient

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Deduplication on power savings?

- Saves ~10X space, and saves hardware
  - Less controllers/enclosures/networking
- Reduces disk IO, faster data processing
  - Longer idle time for power management

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Power management

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Disk power management

- Systems often power manage disks

Enclosure	ES20	ES30
Idle (W)	278	17
Spin down (W)	176 (37% ↓)	↓)
Power down (W)	155 (44% ↓)	↓)

- Spin down saves ~100W for both enclosures
  - Spin down saves 6.5W per disk
- Power down saves 34% more for ES30
  - Power down saves 7.6W to 9.3W per disk
- ES20 uses more power than just 16 disks

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Disk spin down at scale

- System with 1 controller and 32 enclosures

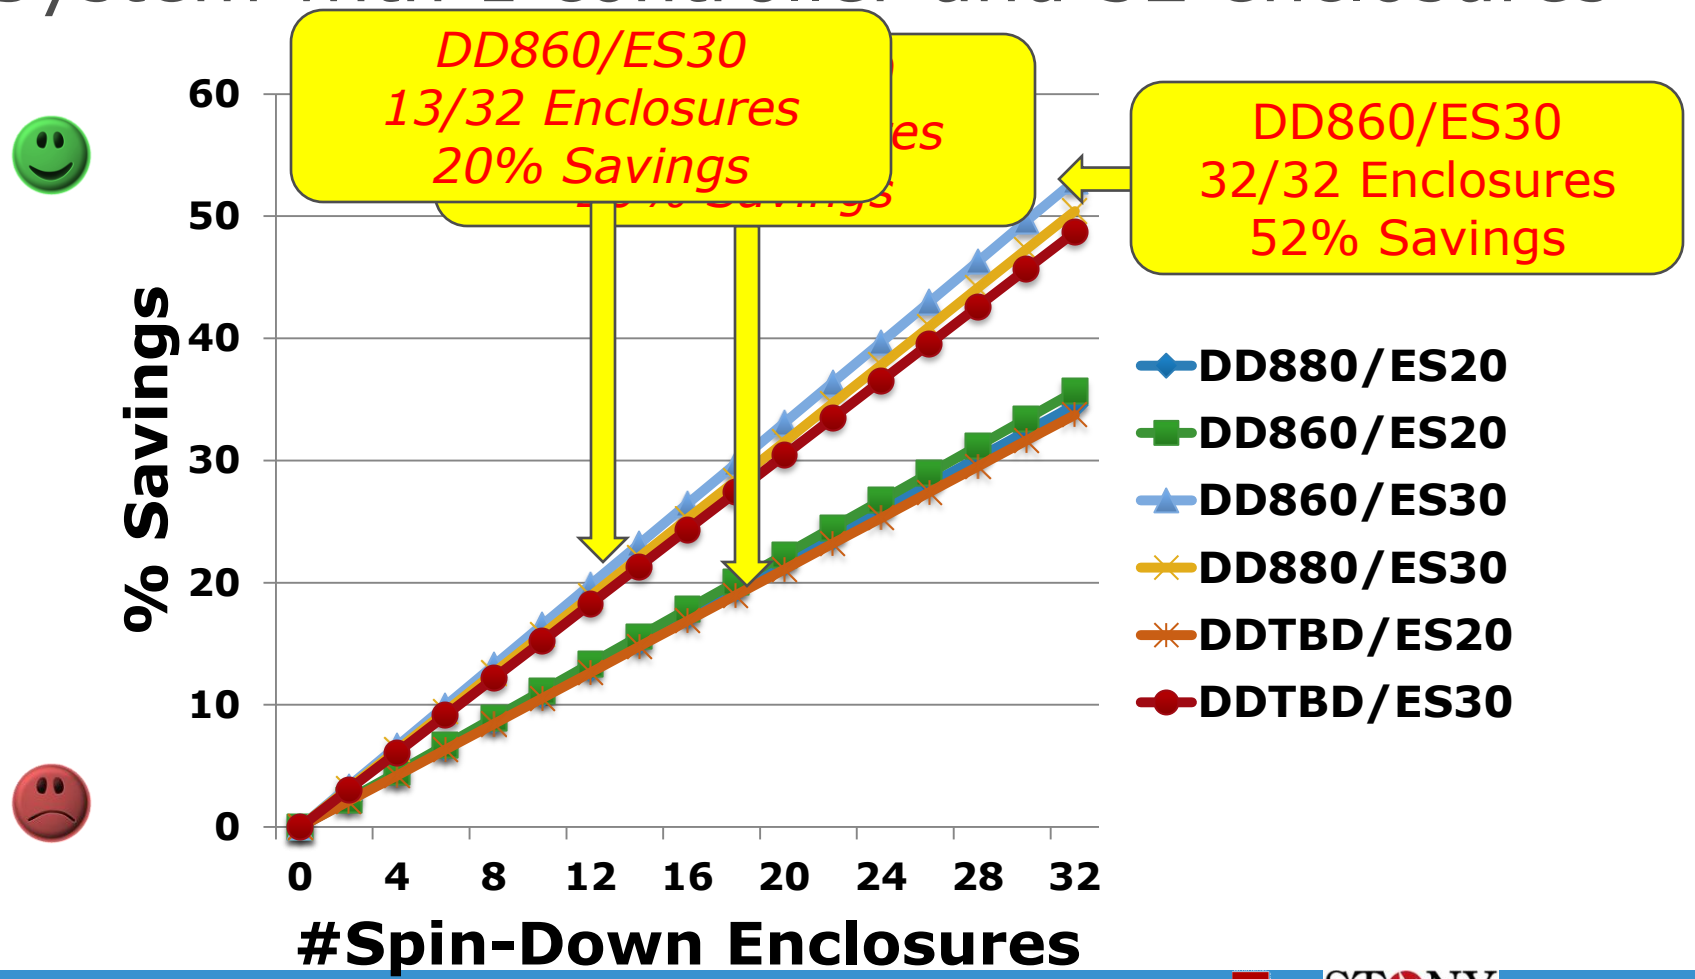

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Power proportionality

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Power Proportionality - Controller

Controller	DD880	DD670	DD860	DDTBD
NFS	44%	24%	58%	20%
OST	58%	29%	61%	36%
BOOST	56%	28%	57%	23%

Power increase ratios from idle to loaded

- Power varies more by model than workload
- Why care?
  - Idle power should take small fraction
  - Great deal of power consumed while idle

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Power Proportionality - Enclosure

Enclosure	ES20	ES30
Idle Power (W)	278	179
Max Power (W)	340(22% ↑)	205(15% ↑)

- 22% increase from idle to max in ES20
- 15% increase from idle to max in ES30
- A lot of power is consumed while idle!

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Conclusions

- Controller/Enclosure are power hungry
  - DDTBD > 100 2TB SATA disks
  - ES20 > 16 disks
  - Disk may not be primary power consumer
- Current systems not power proportional
  - Active vs. idle varies differently
  - More by model than workload
- Disparate consumption between similar H/W
  - Normalized W/TB decreases from 2.89 to 0.675
  - Newer HW more efficient

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Future

- Measure aged backup system
- Measure primary storage system
- Investigate individual components
  - CPU, RAM, etc.

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.

Power Consumption in Enterprise-Scale Backup Storage Systems

Q&A

Zhichao Li[†], Kevin Greenan[‡], Andrew Leung[‡], Erez Zadok[†]

[†]*Stony Brook University*
[‡] *Backup Recovery Systems Division, EMC*

 © Copyright 2012 EMC Corporation and Stony Brook University. All rights reserved.