Reducing Power Consumption by Incorporating Flash Memory into RAID Storage Arrays

Rosie Wacha, Scott Brandt
Computer Science

Motivation

- Hard Disk Drives (disks) consume 27% of power in data centers[1]
- Solid State Drives (SSDs) are flash memory storage devices now available in large sizes
  - Up to 512 GB (Toshiba Q2 2009)
- SSDs consume less power than disks
- Use a small number of SSDs to reduce power and maximize performance improvement

Power Savings with SSD

- Synthetic, constant-bandwidth workload
  - 1.5 MB/s read workload
- Power consumption to run the workload on disk, SSD, or a fraction of data on each
- Power is significantly reduced using SSD

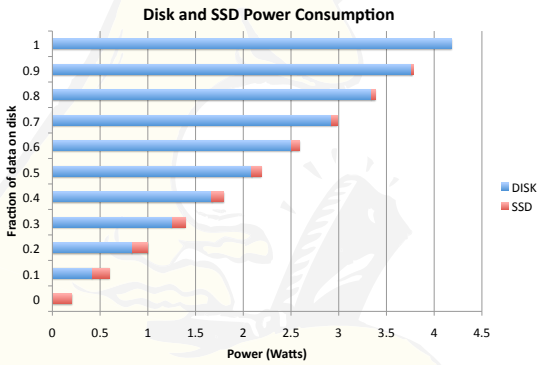

	Disk	SSD
Read	6.0 W	0.35 W
Idle	3.7 W	0.2 W

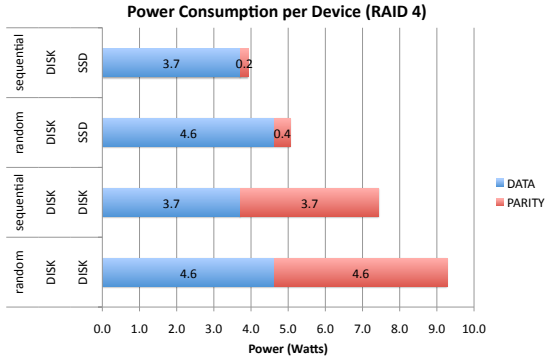

Replace Parity Disk with SSD

- Small writes are faster because SSD read is fast
  - Small writes use half as many disk accesses
  - Parity write occurs in parallel with disk write
- Parity device power is reduced by an order of magnitude
  - 1.5 MB/s write workload
  - Especially significant in smaller stripes
  - Each stripe has one parity device

RAID 4 (Redundant Array of Independent Disks)

- Stores redundant data to provide reliability
  - Parity (exclusive-or) computed across data blocks and stored on parity device
  - Protects against one or more hardware failures
- Stripes data to improve performance
  - Data is broken into chunks and each chunk is stored on a stripe
  - Parity is computed across the stripe
  - Parallel access increases throughput
  - Slower for small writes due to parity read and update

	Intel X25-M SSD	WD VelociRaptor
Cost	$749.99	$179.99
Capacity	160 GB	150 GB
$/GB	$4.69	$1.20
Seek	0.00085 ms	4.2-4.7 ms
Throughput (Read / Wr)	250 / 70 MB/s	120 / 120 MB/s
Power	≤ 0.5 W	≤ 6 W

Simulation Methodology

- Synthetic workload
  - Data read or written at a constant bandwidth
  - Rate is serviceable in all configurations (random and sequential)
- Simulate this workload
  - On either SSD or disk
  - Compare power for each configuration

Performance Impact on RAID 4 with SSD

- Count number of disk writes under each configuration
  - Smaller is better
- Improvement is most significant for smaller stripes
  - Each stripe has one parity device
  - With smaller stripes, a larger fraction is parity

Summary

- Replacing parity disk with SSD reduces power
- Cost to add SSD is fairly low
- Future work investigates performance impact through simulation

Acknowledgments

Research funded by the Institute for Scalable Scientific Data Management (ISSDM).
LANL mentors and collaborators:
John Bent, Gary Grider, and James Nunez.

[1] Charles Weddle, Mathew Oldham, Jin Qian, An-I Andy Wang, Peter Reiher, and Geoff Kuenning. PARAID: A gear-shifting power-aware RAID. In *Proceedings of the 5th USENIX Conference on File and Storage Technologies (FAST)*, February 2007.