

Notes on power proportionality and storage


Eno Thereska

(many involved: Austin Donnelly, Sameh Elnikety, Dushyanth Narayanan, Antony Rowstron)


Microsoft Research, Cambridge, UK

Part I


Turn off servers



One goal: power proportionality



Migrate state and turn off?



Migrate VMs

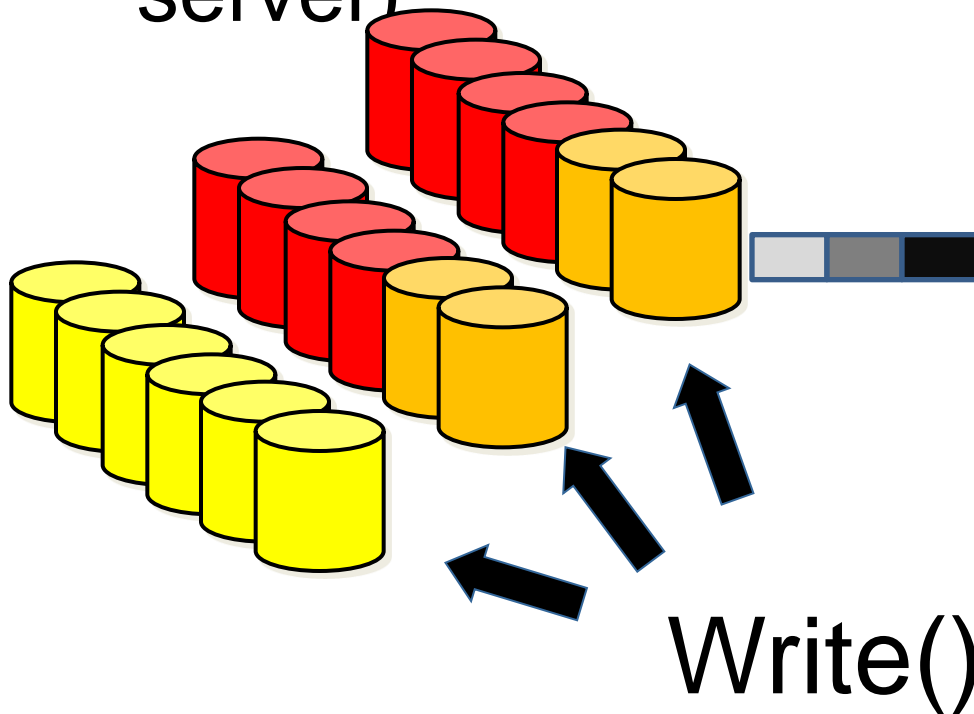


Storage: 1 PetaByte
Cannot migrate!!!

Our solution: Everest

[FAST'08, OSDI08]


- Writes: offload writes temporarily
- Reads: exploit redundancy (or turn on server)



Our solution: Everest

[FAST'08, OSDI08]

- Writes: offload writes temporarily
- Reads: exploit redundancy (or turn on server)



What was so hard about that?

- Metadata management
- Reads must always return latest version
- State must be consistent, recoverable
- Maintain original redundancy level
- Recover state eventually

- Must happen automatically
 - Continuous monitoring of load

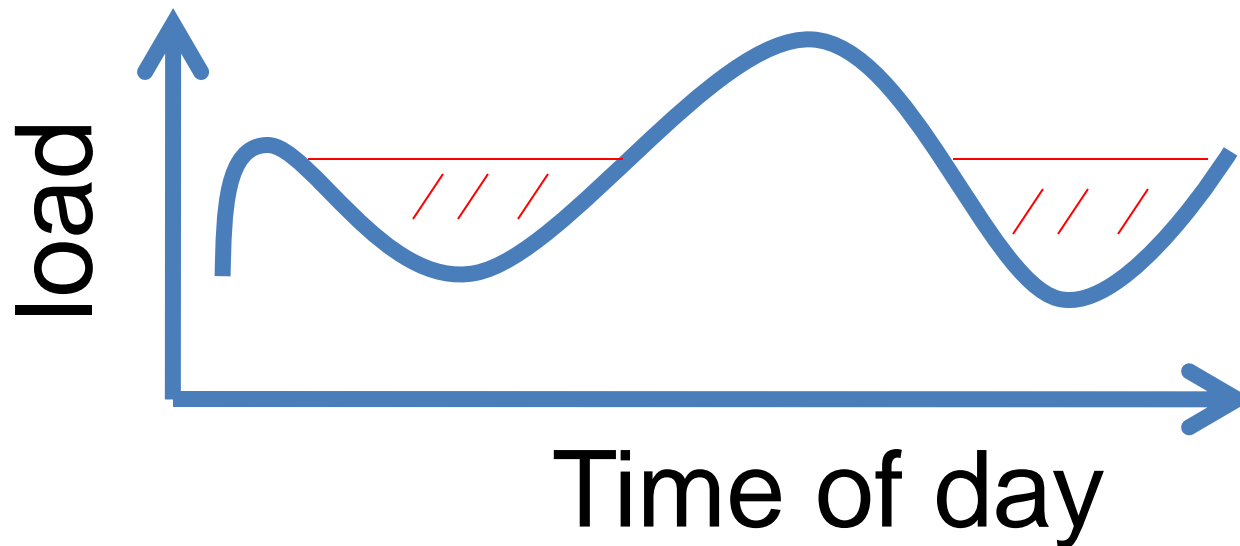
Part II

Do NOT turn off servers

(discussions with James Hamilton @
MSR -> Amazon)

Server bill > power bill

- Turning off servers is a waste of \$\$\$
- But, they have plenty of idle time
 - Workloads tend to be periodic



Fill in troughs

- Many options:
 - Less over-provisioning (buy fewer servers)
 - Schedule lots of analytical jobs
 - Place together uncorrelated workloads
- Hard problems, not all about IT
 - Administrator trust is important

Summary

- Turn off when all other options exhausted
- Storage was missing piece in migration
- Ongoing work:
 - Low-power disks and/or SSDs in the data center [Eurosys'09]