

---

---

---

---

---



# Plan: ZFS

- \* Opening: What is a FS?
- \* Recitation Questions
- \* The problem with the status quo
- \* ZFS - "The last word in file systems"

## Logistics

- \*  $S^3$  is available for you!
- \* No recitation next week.  
(Project presentations)
- \* MapReduce hands on due 4/22
- \* Holidays Mon & Tues!  
↳ Enjoy your time off.

# What is a file system?

The system that manages how data is stored on disk

e.g. Unix FS, GFS


## Why is building a FS hard?

→ Performance critical


→ Handle failures.

↳ ANY TIME!

→ `write(/home/amir/a.txt, ...)`



`write(sector_id, data)`  
`read(sector_id) → data`



# Recitation Questions

1 What aspects of UNIX FS this overcomes?

- simplify admin
- virtualize storage


2 How is ZFS designed to overcome these issues?


- pooled storage
- dynamic sizing
- flexibility

3 Why inport? Why not GFS?


↳ more efficient & compact  
redundancy easier

↳ single controller






# Traditional FS





# ZFS




# Consistency: Traditional FS

Example: Move a file



# Consistency: ZFS





# Plan: ZFS

- \* Opening: What is a FS?
- \* Recitation Questions
- \* The problem with the status quo
- \* ZFS - "The last word in file systems"

## Logistics

- \*  $S^3$  is available for you!
- \* No recitation next week.  
(Project presentations)
- \* MapReduce hands on due 4/22
- \* Holidays Mon & Tues!  
↳ Enjoy your time off.

# What is a file system?

↳ System that handles storage of data on disk


## What makes designing a FS hard?

↳ A lot of metadata/bookkeeping

↳ Handle multiples disk  
↳ administrative

↳ Handle errors, crashes

↳ Performance



## Recitation Questions

1. What aspects of the Unix FS does ZFS improve?


↳ complex administration, crash-safety, capacity.

2. How does ZFS solve these problems?

↳ checksums, [pooled storage/virtualization copy on write.


3. Why does GFS not solve this problem?

↳ Solving different problems




GFS


store files




ZFS



# Traditional FS






# ZFS



# Consistency

Example: Move a file



Consistency : ZFS

- Overhead / extra copies
- Extra space

