

The Sysadmin's Warm-Up

Prepared by Mark on March 3, 2025

Most of you have seen a hard drive.

Many have touched one, and a lucky few have taken one apart.

These devices have two interesting properties:

- They hold valuable data
- They inevitably fail

Needless to say, this is a problem.

We would like to design a system that tolerates hard drive failures without data loss.

Definition 1:

You can think of a hard drive as a long string of bits.

Assume all hard drives in the following problems have the same size.

If a hard drive “fails”, all data on it is instantly lost.

Problem 2:

Suppose we have two hard drives. How can we arrange our data so that...

- We get 1 TiB of usable storage
- We lose no data if any one drive fails

Problem 3:

Suppose we have three hard drives. How can we arrange our data so that...

- We get 2 TiB of usable storage
- We lose no data if any one drive fails