

Exercise 1

1 Why Internet Services Fail

Read the following paper:

D. Oppenheimer, A. Ganapathi, and D. A. Patterson, “Why do Internet services fail, and what can be done about it?” in *Proc. 4th USENIX Symposium on Internet Technologies and Systems (USITS '03)*, 2003.

<http://roc.cs.berkeley.edu/papers/usits03.pdf>.

Write two paragraphs to summarize its findings.

2 RAID Failure

RAID systems distribute data over multiple physical disk drives. In RAID4, for example, one drive among N is configured to hold the parity information computed over the $N - 1$ remaining drives. Such RAID systems are able to tolerate the failure of one drive. Often a spare drive can be inserted into the running system for replacing the failed one. A recovery procedure then recomputes the data of the failed drive and writes it to the spare drive, bringing the RAID system back to normal operation mode. With large drives, recovery may take a long time.

- a) Suppose one drive has a MTTF of 50'000h and it takes 20h to recover one failed drive. Estimate the probability that a RAID4 system with N drives fails during recovery because of a second drive failure.
- b) Unrecoverable bit errors also contribute to RAID system failures during recovery. Suppose the rate at which a drive fails to read back the written data is 10^{-14} /bit and that our RAID4 system consists of 100GB drives. What is the probability that recovery fails because of a bit error?

Such calculations motivate the interest in RAID schemes that can tolerate two drive failures simultaneously.