ECE566 Enterprise Storage Architecture
Individual Homework #1: Drives and RAID

Directions:

- This assignment will be completed in ***INDIVIDUALLY***. While you can discuss concepts with others both within and outside your group, actual steps and answers should not be shared!
- The solutions should be collected in a PDF file and submitted via GradeScope.
- At times, you will be asked to speculate as to *why* something is happening. **If you aren't confident in your answers, see the instructor to discuss it!** Talking through things you're unsure of is part of the class.

1 HDD vs. SSD performance [25pts]

1.1 HDD thought experiment [5]

You are optimizing software where the bottleneck is random I/O access to a hard disk drive. You identify two possible improvements: (a) a change to the on-disk data structure so that related data is closer together and often contiguous, or (b) a change so that each record is compressed to a smaller size, but the on-disk layout is otherwise unaffected. Without knowing any additional facts, which do you guess is worth trying first? Why? [5]

1.2 Choosing between HDD and SSD [20]

You have a workload that requires 60 TB of capacity and 200,000 IOPS of random I/O performance, mostly reads. You're going to support it with a storage array, and are deciding between buying Samsung 870 EVO 4TB SSDs or Seagate IronWolf Pro 16TB NAS HDDs.

(a) Research the two drives, and identify published benchmarks for IOPS performance (don't trust manufacturer datasheets). Hint: Storage Review is a well-known third-party reviewer of storage devices. You can assume the I/O size is 4kB. Include a link to your sources. [4]

(b) Identify the current street price for each drive. Include a link to your sources. [4]

(c) Ignoring RAID effects and assuming performance is simply additive, how many SSDs are needed? Is this number driven by capacity or performance? [4]

(d) Ignoring RAID effects and assuming performance is simply additive, how many HDDs are needed? Is this number driven by capacity or performance? [4]

(e) How much would each option cost, and which is cheaper? [4]

2 RAID [70 pts]

2.1 RAID layouts [25]

(a) How many disks can fail in a 4-disk RAID-0 without data loss? What fraction of the storage is dedicated to redundancy? What is the read throughput relative to using a single disk? The write throughput? [5]

(b) How many disks can fail in a 2-disk RAID-1 without data loss? What fraction of the storage is dedicated to redundancy? What is the read throughput relative to using a single disk? The write throughput? [5]

(c) How many disks can fail in a 6-disk RAID-5 without data loss? What fraction of the storage is dedicated to redundancy? What is the read throughput relative to using a single disk? [5]

(d) How many disks can fail in an 18-disk RAID-6 without data loss? What fraction of the storage is dedicated to redundancy? What is the read throughput relative to using a single disk? [5]

(e) Assume a RAID-1+0 with 10 disks organized into a RAID-0 of five 2-disk RAID-1 sets, as shown below:

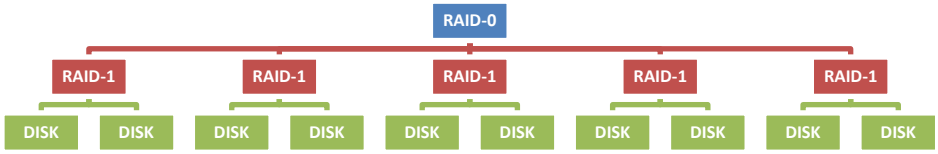

What is the minimum number of disk failures that can cause data loss? What is the maximum that can fail *without* data loss? What fraction of the storage is dedicated to redundancy? [5]

2.2 RAID reliability modeling [10]

Let's do some statistics to understand how big we can safely make RAID arrays. If you need a refresher, this tutorial can remind you how to compute the probability of at least *k* events of probability *p* happening over *n* trials. In our case, *k* will be the number of drive failures to cause data loss (2 for RAID5, 3 for RAID6), *p* is the annual failure rate of our drives, and *n* is the number of drives in the array. Assume a constant 4% annual drive failure rate (*p*=0.04).

(a) What is the maximum number of drives we can put in a RAID 5 array such that the annual probability of data loss for the array is under 1%? [5]

(b) What is the maximum number of drives we can put in a RAID 6 array such that the annual probability of data loss for the array is under 1%? [5]

Hint to check your work: the answer to part (a) is a power of two; the answer to part (b) is prime.

2.3 RAID details [40]

(a) What role do *hot spares* play in RAID deployments? [10]

(b) Why is a RAID configuration **not** the same as a backup solution? Describe a scenario in which data loss could occur without disk failure. [10]

(c) Many vendors have developed extensions to the standard RAID levels, and studying these can offer interesting insight into the engineering trade-offs involved in designing a storage subsystem. Choose **one** of the questions below to research and answer. The Wikipedia article "Non-standard RAID Levels" can be a good starting point in your research. [20]

1. The Linux software RAID driver ("md") has a dedicated RAID-10 mode that offers some unique features. In this approach, describe the "near" and "far" layouts. What benefits to these techniques have?

2. What is RAID-1E? What is the storage overhead for this approach? For a 5-disk RAID-1E, what is the minimum number of disk failures needed to cause data loss, and *which* disks must fail in this scenario? What is the maximum number of disks that can fail *without* data loss? What performance effects does this approach have?

3. What is the primary advantage of RAID 5E, 5EE, and 6E? What are the disadvantages? What factors would influence you to choose one of these approaches over a traditional RAID approach?

4. BeyondRAID (from the Drobo line of consumer storage devices) and Synology Hybrid RAID (from Synology storage devices) both allow drives of varying capacity to be used while still providing single-disk-failure redundancy. How do these technologies work? How could you replicate such a layout using nothing but normal partitioning and a software RAID system, e.g. Linux "md"?