RAID

1. Distinction between logical and physical I/O address spaces.
2. What was the original & current motivation for RAID?
3. Why is a multiple-disk storage system less reliable than a single disk?
4. How does Mean Time to Failure (MTTF) change as number of components in a system increases?
5. What are the different levels of RAID and how do each of them work?
6. What are the relative benefits/drawbacks of each RAID level?
7. How is data distributed in each RAID level?
8. How is parity calculated and stored in each RAID level?
9. What is the extent of read and write parallelism in each level?
10. How is the parity calculation bottleneck in RAID 4 solved?
11. In RAID 5, describe how you can complete a write I/O operation using just 2 disk reads and 2 disk writes.

12. (a) What problem in RAID 3 does RAID 4 solve and how? (b) What problem in RAID 4 does RAID 5 solve and how?

13. In RAID-5, explain how can you perform a single logical write operation in no more than one physical read and two physical writes?

14. How should parity be computed in RAID 5 to increase parallelism of write operations? Explain with parity computation formula.

15. What is the write parallelism problem in RAID and how is it solved?

16. Consider RAID levels 0, 1, 3, 4, and 5: Which RAID level provides the best (a) reliability (b) I/O Parallelism. Explain why.

17. (a) Explain (with formula), how does parity computation differ between RAID 3 and RAID 4? (b) How and why does parity placement on the disk (not parity computation) differ between RAID 4 and RAID 5? Explain with example.

19. For a RAID system **with N disks, including data and parity,** compare the level of parallelism provided by RAID 1, RAID 3, RAID 4, and RAID 5 for multiple simultaneous (i) read I/O operations, (ii) write I/O operations, and (iii) combination of read and write I/O operations? Explain your answers.

20. Consider RAID levels 1, 3, 4, and 5. Which RAID level provides the best (a) reliability

(b) I/O Parallelism. Explain why. Assume data worth N disks of same size and additional space for any parity disks.

21. [Open-ended question] In order to save power, disks are usually spun down (placed in sleep or low-power mode). This works well if there is only one disk in the system, if all data resides on the single disk, and if performance is not a major concern. Consider a RAID-5 system consisting of N+1 disks. Explain how you can redesign RAID-5 so that all the following requirements are satisfied: (1) fault-tolerance of original RAID-5 is maintained under all conditions, (2) energy consumption is minimized by spinning down one or more disks whenever possible, and (3) performance (read/write throughput) of the system is maximized to the extent possible. Again, while there is no single correct answer, you must explain all salient aspects of your design, justify any assumptions you make, and examine any design tradeoffs (e.g. energy savings to performance).

22. [Open-ended question] Describe the design of a parity-based RAID system that can survive two-disk failures (as opposed to single-disk failure discussed in class). In your design, be sure to explain the following: (a) How your system would compute the parity required for recovery from a two-disk failure? (b) How your system would recover from two-disk and single-disk failures, (c) How much additional space would parity information occupy, compared to data, and (d) How many parallel read and write I/O operations can your system support?