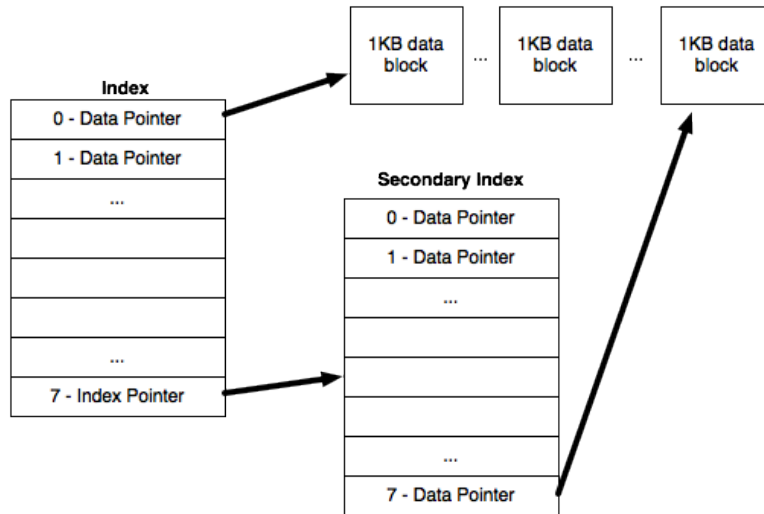


Question 1. (25 points) Consider the file architecture shown in the following figure. For each file there is an index structure which contains 8 entries: the first 7 entries are pointers to 1KB data blocks and the last entry points to a secondary index data structure which contains 8 more data pointers.



- a. (5 points) What is the maximum size of a file using this architecture?
- b. (10 points) If you increased the number of entries in both data structures from 8 to 32 (but still kept only one index pointer in the Index), how would that impact the size of allowable files? Would it impact the performance of file accesses (and if so, how)?
- c. (10 points) What simple change could you make to this structure to support files of infinite size? Is there any downside to this (modified) approach?

Question 2. (10 points) Using the following set of disk requests, what is the order of seeks when using the SSTF and SCAN disk scheduling algorithms? What is the total distance of seeks in each case? Assume that the disk has its head at position 45 out of 100 and that it is currently moving towards higher numbers (for SCAN).

Disk Request Queue: 44, 57, 78, 65, 46, 90

Question 3. (15 points) Describe a scenario where the algorithm which performed best in the previous question (either SSTF or SCAN) will perform **worse** than the other algorithm. To do so, provide an example disk request queue with at least six entries in it and show why the performance is worse.

Question 4. (10 points) Suppose you are designing a filesystem to run on SSDs. Since there is no longer any rotational delay, is there any downside to scattering blocks of individual files all over the disk, as opposed to storing them in a consecutive series of sectors? Why or why not?

Question 5. (10 points) Explain how DMA improves the performance of I/O devices. What kinds of devices benefit the most from DMA?

Question 6. (10 points) Suppose you are designing a distributed system in which timestamps are assumed to be comparable across machines (i.e., a timestamp of t was actually earlier than a timestamp of $t + 1$). To deal with the problem of clock skew (i.e., out-of-sync clocks), your friend suggests designating a specific machine as the timeserver and having all other machines request timestamps from the timeserver. What are two downsides of designing your system in this way?

Question 7. (10 points) Explain the difference between implementing protection using a capability-based scheme and an access-control list based scheme.

Question 8. (10 points) Suppose you are a data center operator and charge each customer a fixed hourly rate to run an application in your data center. As the operator of the data center, you are responsible for paying the costs associated with your hardware (e.g., electricity for each physical server). Explain why you will be able to earn a greater profit by employing virtualization on your servers.