

Exercise Sheet 8

Exercise 1 (Hard Disk Drives)

1. What are sectors (= blocks) in HDDs?
2. What are tracks in HDDs?
3. What are cylinders in HDDs?
4. What are clusters in HDDs?
5. Draw the structure of a hard disk drive schematically. Explain with your drawing(s) the meaning of the following terms:
 - a) Sector (= Block)
 - b) Track
 - c) Cylinder
 - d) Cluster

6. Why is it impossible to improve the performance (especially the latency) of HDDs infinitely?

7. Which factors influence the access time of HDDs?

-
-

8. Describe the factors of subtask 7.

-
-

Exercise 2 (Solid State Drives)

1. Why is it wrong to call SSDs Solid State Disks?

.

2. Name four advantages of SSDs over HDDs.

3. Name two drawbacks of SSDs over HDDs.

4. Why are erase operations on flash memory more complex than read operations?

5. What is the objective of wear leveling algorithms?

Exercise 3 (RAID)

1. Which RAID levels improve the data transfer rate for write?
 RAID-0 RAID-1 RAID-5
2. Which RAID levels improve the reliability?
 RAID-0 RAID-1 RAID-5
3. How many drives are allowed to fail in a RAID 0 array without data loss?
4. How many drives are allowed to fail in a RAID 1 array without data loss?
5. How many drives are allowed to fail in a RAID 5 array without data loss?
6. Please comment the statement: "A RAID array can be used to replace the regular backup of important data" .
7. Why is it not useful to store all parity information on a single drive, but to distribute the parity information on all drives?

8. What is the net capacity of a RAID 0 array?

9. What is the net capacity of a RAID 1 array?

10. What is the net capacity of a RAID 5 array?

11. Name one advantage and one drawback of software RAID compared with hardware RAID.

Exercise 4 (File Systems)

1. Describe which information inodes store.

2. Name three examples of metadata in the file system.

3. Describe what a cluster in the file system is.

4. Describe how a UNIX file system (e.g. ext2/3), which does not implement extents, can address more than 12 clusters.

5. Describe how directories in the Linux file systems are technically implemented.

6. Name one advantage and one drawback of small clusters in the file system compared with large clusters.

7. Do DOS/Windows file systems differentiate between uppercase and lowercase?
 Yes No
8. Do UNIX file systems differentiate between uppercase and lowercase?
 Yes No
9. Do modern operating systems accelerate requests to stored data with a cache in the main memory.
 Yes No
10. Most operating systems operate according to the principle...
 write-back write-through
11. Name one advantage and one disadvantage of a cache in the main memory, which is used by the operating system to accelerate the requests to stored data.

12. Explain what an absolute path name is.

13. Explain what a relative path name is.

14. `/var/log/messages` is an/a...
 absolute path name relative path name
15. `OpSys/Lecture_09/opsys_slides_09.tex` is an/a...
 absolute path name relative path name
16. `Documents/MasterThesis/thesis.tex` is an/a...
 absolute path name relative path name
17. `/home/<username>/Mail/inbox/` is an/a...
 absolute path name relative path name
18. Describe what information the boot sector (also called boot block) of a file system stores.

19. Describe what information the super block of a file system stores.

20. Explain why some file systems (e.g. ext2/3) do combine the clusters of the file system to block groups.

21. Describe what the File Allocation Table (FAT) is and describe the information it stores.

•

Exercise 5 (File Systems)

Please mark for each statement about file systems, whether the statement is true or false.

Statement	true	false
Inodes store all metadata of files.		
File systems address clusters and not blocks of the storage medium or storage drive.		
The smaller the clusters are, the more overhead for large files occur.		
The bigger the clusters are, the lesser capacity is lost due to internal fragmentation.		
In UNIX, file extensions have always been of great significance.		
Modern file systems operate so much efficient that buffering by the operating system is no longer common.		
Absolute path names describe the complete path from the root to the file.		
The separator in path names is identical for all operating systems.		
An advantage of block groups is that the inodes are physically located close to the clusters, they address.		
For each cluster in the file system, an entry exists in the FAT.		
Because of the Master File Table in NTFS, fragmentation cannot occur.		
The journal of journaling file systems reduces the number of write operations.		
Journaling file systems narrow down the data, which need to be checked during the consistency check.		
When using journaling file systems, a loss of data is impossible.		
If metadata and file contents are journaled both, all write operations are carried out twice.		
Extents cause lesser overhead compared with block addressing.		