Basics of Information Systems

Winter Semester 2022–23

For discussion on Wednesday, January 11, 2022

1. Soft and hard links are often used for convenience purposes when structuring file data. To avoid unexpected behavior, it is important to understand their behavior, in particular the conceptual differences between soft and hard links. Answer each of the following questions for soft and for hard links separately.

   (a) Suppose you set a link named `linktofile` to a file named `fileA`. Then you delete `fileA`. What happens? Is the data gone? If not, how do you access it?

   (b) Suppose you set a link named `linktofile` to a file named `fileA`. Then you delete `linktofile`. What happens? Is the data gone? If not, how do you access it?

   (c) Can you determine if a file is being linked to? How? (There is no need to describe precise Unix/Windows commands, a conceptual answer suffices.)

   (d) Suppose you write a program to copy (a part of a) file system to a backup location and to restore the data from backup if necessary. What may go wrong if your data includes hard or soft linked files?

   (e) You are about to modify a file that is being linked to. Because you are unsure about the modification, you want to keep an old copy around. How to you proceed? What do you have to be aware of?

   (f) Can you link to a directory? What can possibly go wrong?

   (g) What is faster – file access via a hard link or via a soft link?

2. RAID arrays are often used to increase storage capacity and/or reliability. Some of the important operational modes, called *RAID levels*, are RAID-0, RAID-1, RAID-5. Answer the following questions, comparing these three RAID levels.

   (a) Which RAID levels improve the data transfer rate for writes?

   (b) Which RAID levels improve reliability?

   (c) Which RAID levels reduce reliability? By how much?

   (d) How many drives are allowed to fail in each RAID level?

   (e) Name at least two reasons why RAID is not a backup solution.

   (f) How much capacity does each RAID level provide if a single drive has storage capacity C?