IA Operating Systems:
I/O, Storage, File Management, Case Studies

Lent 2024
Last Updated: Friday 23rd February, 2024 23:47

1 For each of the following, indicate if the statement is true or false, and explain why:

(*a*) Non-blocking I/O is possible even when using a block device.

(*b*) Direct memory access (DMA) makes devices go faster.

(*c*) In Unix, hard-links cannot span mount points.

(*d*) The Unix shell supports redirection to the buffer cache.

2 File systems comprise a *directory service* and a *storage service*.

(*a*) What are the two main functions of the directory service?

(*b*) What is a directory *hierarchy*? Explain your answer with the aid of a diagram.

(*c*) What information is held in file *meta-data*?

(*d*) What is a *hard link*? Does file system support for hard links place any restrictions on the location of file meta-data?

(*e*) What is a *soft* (or *symbolic*) link? Does file system support for soft links place any restrictions on the location of file meta-data?

3 From the point of view of the device driver, data may be read from an I/O device using *polling*, *interrupt-driven programmed I/O*, or *direct memory access* (DMA). Briefly explain each of these terms, and in each case outline using pseudo-code (or a flow chart) the flow of control in the device driver when reading data from the device.

4 (*a*) Compare and contrast *blocking*, *non-blocking* and *asynchronous* I/O.

(*b*) Give *four* techniques which can improve I/O performance.

5 (*a*) Describe with the aid of a diagram the on-disk layout of a Unix V7 filesystem. Include in your description the role of the *superblock*, and the way in which free inodes and data blocks are managed.

(*b*) Describe with the aid of a diagram a Unix V7 *inode*.

(*c*) Estimate the largest file size supported by a Unix V7 filesystem.

(*d*) Suggest *one* reliability enhancement and *two* performance enhancements which could be made to the Unix V7 filesystem.

6 (*a*) Describe, with the aid of diagrams where appropriate, how Unix implements and manages:

(*i*) a hierarchical name space for files

(*ii*) allocation of storage on disk

(*iii*) file-system and file meta-data

(*iv*) pipes

(*b*) A Unix system administrator decides to make a 'versioned' file-system in which there are a number of directories called `/root-dd-mm-yyyy`, each of which holds a copy of the file-system on day `dd`, month `mm` and year `yyyy`. The idea is that at any particular time only the most recent snapshot will be used as the 'real' filesystem root, but that all previous snapshots will be available by explicitly accessing the directory in question. In this way the system administrator hopes to allow resilience to mistaken edits or unintentional deletions by users, or to hardware problems such as a disk head crash.

To implement this, the system administrator arranges for a program to run every morning at 01:00 which recursively 'copies' the current snapshot to the new one. However to save disk space, hardlinks are used in place of actual copies. Once the 'copy' is complete, the new snapshot is used as the new root.

To what extent will this scheme provide the functionality the system administrator hopes for? What advantages and disadvantages does it have?

7 Describe *how* CPU scheduling algorithms favour I/O intensive jobs in the Unix operating systems.

8 Describe the operation of the Unix shell with reference to the process management system calls it makes use of. You might like to use pseudo-code or a diagram to aid with your description.

9 Past paper questions

- y2012p2q4
- y2010p2q4
- y2007p1q8 (d)
- y2009p2q3 (b)