

Lecture 11

The OS file system

- What services does the file system provide ?
- How is a file structured ?
- What is metadata and how can it be useful ?
- Is the file system part of the kernel or not ?
- What is the file name space ?

File system services

- Create/remove file, address file, open/close file, read/write an open file, fetch/modify metadata of a file.
- Shared/exclusive access to a file.
- Algorithm for exclusive file access:
 - If Q (queue of processes awaiting access) is nonempty, then add P to the tail of Q and return.
 - If no process has currently exclusive access to the file, P gets access and return;
 - If P requests read-only access and Q is empty and the processes with current access are readers, P gets access and return.
 - Add P to the tail of Q and return.

Access patterns and file structure

- *Sequential access – the current position in the file is stored;*
- *Random access – rewind operation to get to the read/write position.*
- A file can be considered as an unstructured collection of bytes; applications are managing the byte stream.
- A file can be made of records of fixed or variable size.

Metadata

- Data about the file: name, size, last modification date, owner, protection codes,... managed by the OS and in some cases by applications as well.
- The file type is supported by some OS.
- File types are used in a number of ways, for example they can be used to control certain aspects of reading/writing, such as end-of-line conventions for text files.
- Files that are compressed are automatically decompressed when read.
- Another aspect is assigning an application to the file.

The FS design

- The FS process can be part of the kernel or not, running as a distinct process in the user space. In the latter case, there are some issues:
 - How does it handle system calls ?
 - How does it access process's memory space ?
 - How can it obtain process information that is stored in process tables ?
 - How does it access the device drivers ?
- One possible answer: a new protocol of message passing between processes.

Management data structures

- The FS uses two important data structures for file management.
- The *open file table* has entries for all open files. Each entry stores file attributes (size, device, if it is shared or not,...).
- The *mount table*. Making a file system accessible is called mounting (unmounting is the process of removing a file system from the accessible set). The details of mounting a file system vary considerable from one system to another.
- In some cases a device is referenced explicitly and the fs metadata is read at that time. For other systems, the new fs is added into an overall uniform naming scheme. In these systems, the file names are determined by the mount operation rather than by the physical operation.