



UNIVERSITY OF OREGON

CIS 415: Operating Systems File Systems

Spring 2014
Prof. Kevin Butler

- Last class:
 - ▶ Virtual Memory
- Today:
 - ▶ Files

A System Problem?

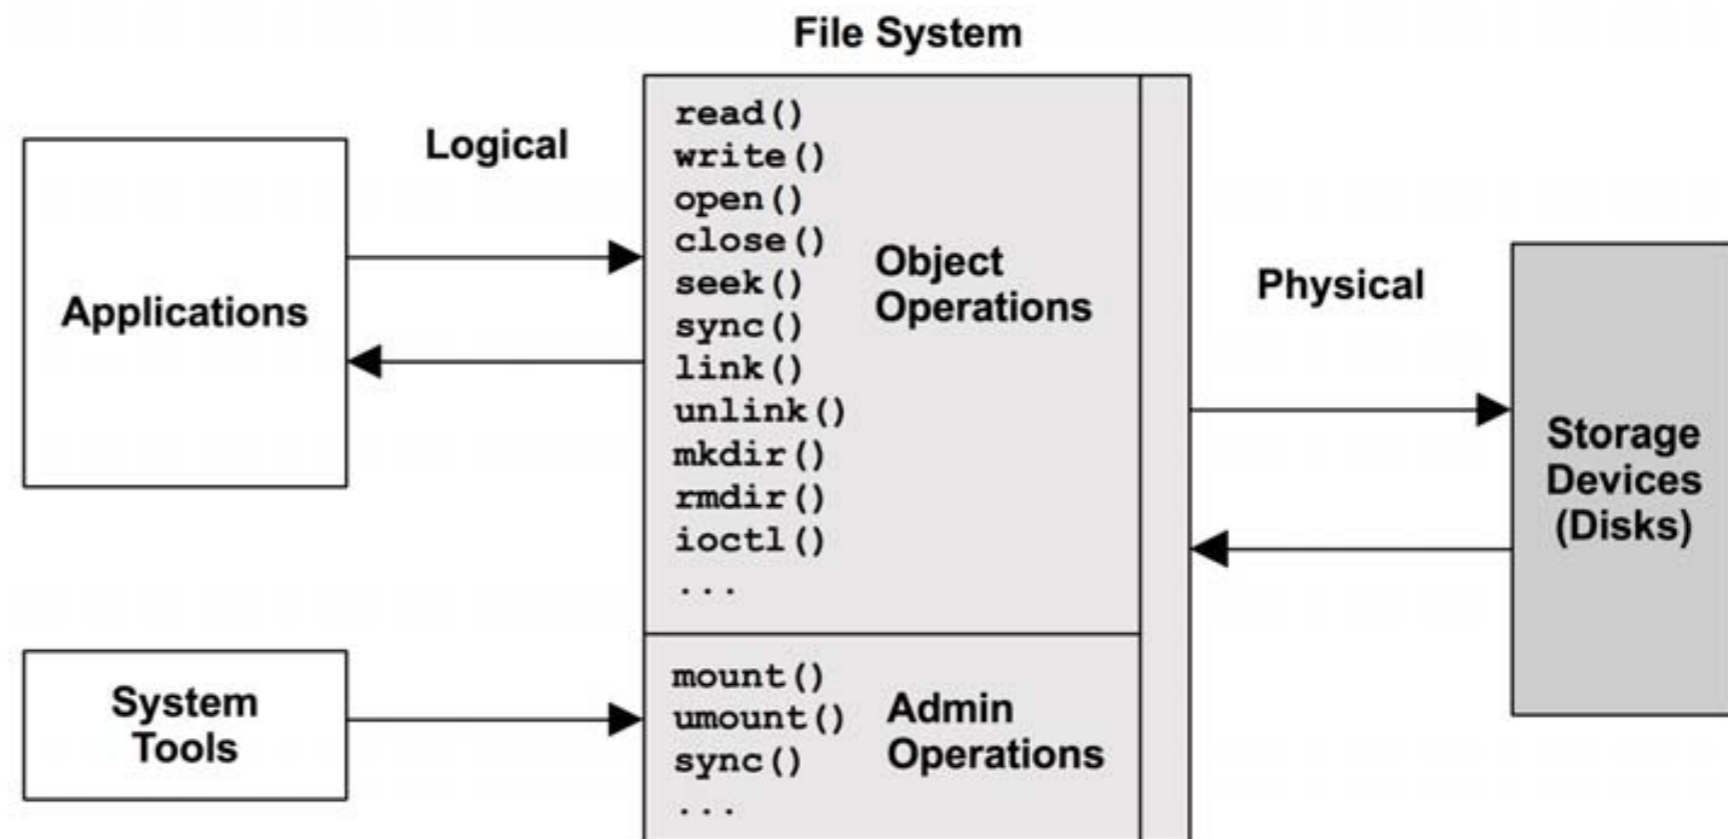


- Got some data in your program
 - ▶ Want to keep it for a while
- Got a long running program
 - ▶ Want to prevent loss of data if it crashes
- Got a lot of programs, system resources, data, etc. stored
 - ▶ Want a mechanism to refer to them all

File System Interface



- Most visible part of the OS
- Consists of
 - Files
 - Directories
- And sometimes
 - Partitions



What is a file?



- A repository for data
- Is long lasting (until explicitly deleted).
- Also, may refer to a system resource (device)

Why not just an address space?



- You may want data to persist longer than a process
- You may want data that is larger than a virtual address space
- Easier to share the data across processes.

Two aspects to consider



- User's view
 - ▶ Naming, type, structure, access, attributes, operations, ...
- System implementation

- Typically **x.y**
- **x** is supposed to give some clue about contents
- **y** is supposed to be the nature of the file.

- Byte stream
- Sequence of Records
- Indexed Records

Types of File Objects



- Regular files (containing data)
- Directories
- Character special files (access a character at a time)
- Block special files (access a block at a time)

- protection, creator, owner, creation time, access time, current size, max size, record length, lock flags, ...


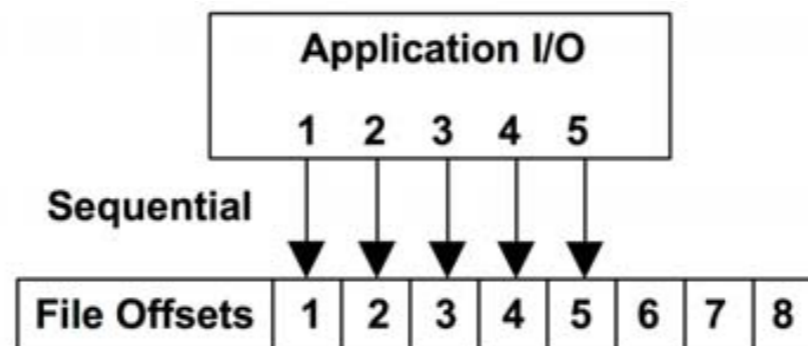
- Sequential Access

- ▶ reset
- ▶ read next (advance file pointer automatically)
- ▶ write next (advance file pointer automatically)

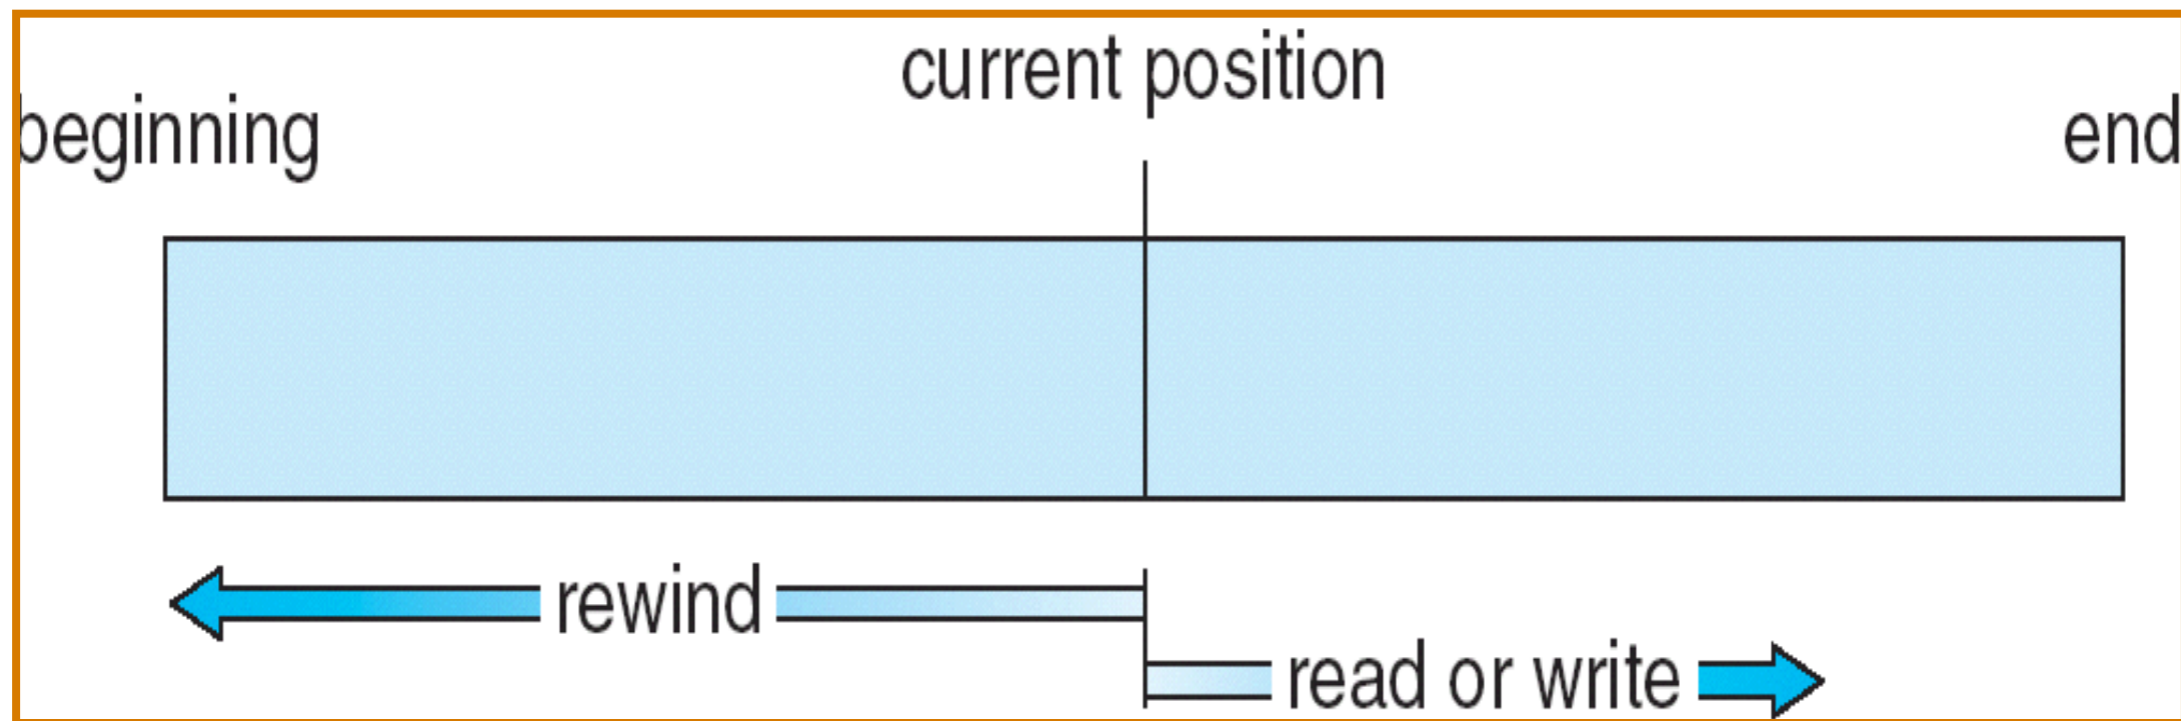
- Direct Access

- ▶ read n
- ▶ write n
- ▶ position to n
- ▶ read next
- ▶ write next

- $n =$ relative block number



Sequential File Access



Sequential File Access



Simulation of direct access:

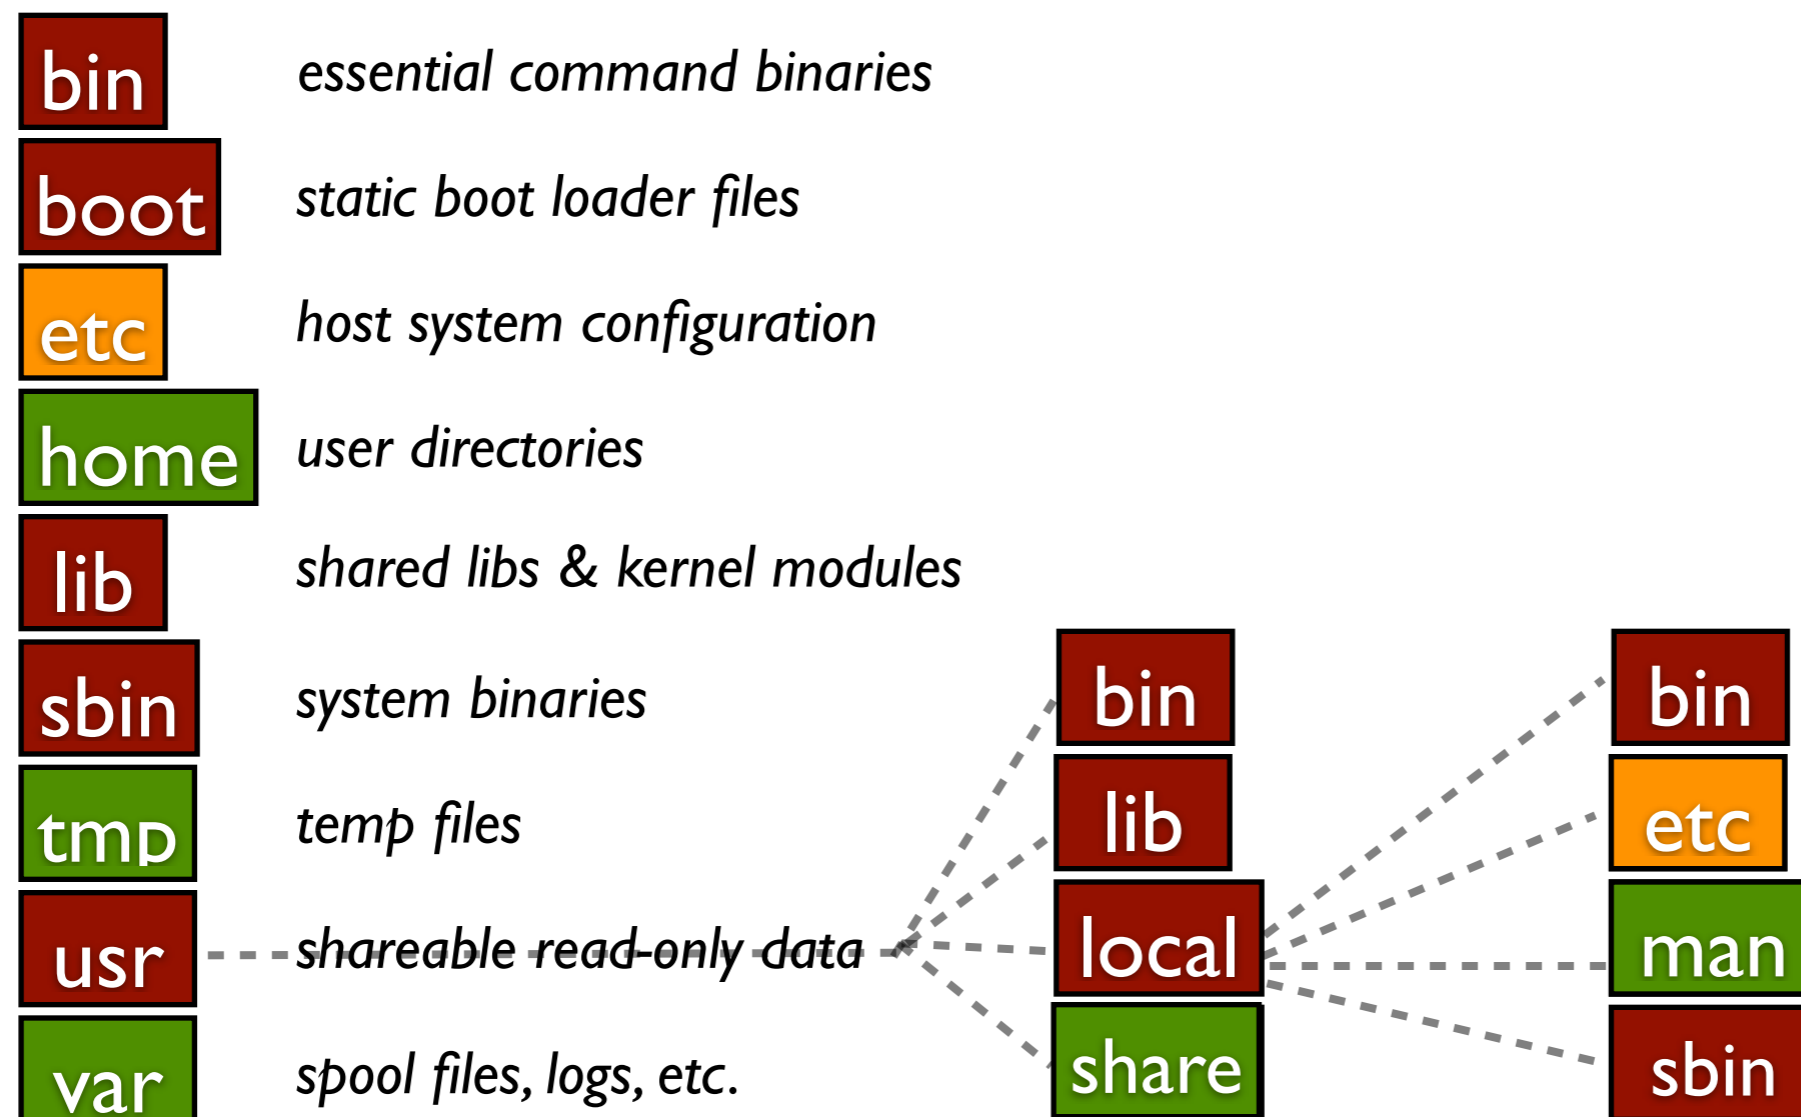
sequential access	implementation for direct access
<i>reset</i>	<i>cp = 0;</i>
<i>read next</i>	<i>read cp;</i> <i>cp = cp + 1;</i>
<i>write next</i>	<i>write cp;</i> <i>cp = cp + 1;</i>

- A way of organizing files.
- Each directory entry has:
 - ▶ File/directory name
 - ▶ A way (pointer) to get to the data blocks of that file/directory

- Flat (only 1 directory) vs. hierarchical file system
- File names: relative vs. absolute
- Directory Operations:
 - ▶ Create
 - ▶ Delete directory
 - ▶ Open Dir
 - ▶ Close Dir
 - ▶ Read Dir
 - ▶ Rename
 - ▶ Link (allow a file to appear in more than 1 directory)
 - ▶ Unlink

Filesystem Hierarchy Standard

- UNIX filesystem hierarchy standard (FHS) is a convention
 - ▶ enables software and users to predict the location of installed files and directories



- Makes a file appear in more than 1 directory.
- Is a convenience in several situations.
- 2 types of links:
 - ▶ Soft links
 - ▶ Hard links




- ▶ Create a file which contains the name of the other file.
- ▶ Use this path name to get to the actual file when it is accessed.
- ▶ Problem: Extra overhead of parsing and following components till the file is reached.

- ▶ Create a directory entry and make a reference that file.
 - Others may reference the same file
- ▶ Problem: What if the creator wants to delete the file?
 - There are still other references to the file, potentially.
 - Cannot free up until all the other references are removed.
 - Done by keeping a counter that is incremented for each hard link. On removing a link, decrement counter. Only if counter is 0, remove the file.

- A way of organizing directories
- Each partition contains a:
 - ▶ File system of directories
- Examples:
 - ▶ Root file system '/'
 - ▶ Boot file system '/boot'
 - ▶ User's homes '/home'




File System Mounting



File System Mounting



UNIVERSITY
OF OREGON



- In a multi-user system,
 - ▶ There is interest in sharing files
- System files
 - ▶ Shared by all
 - ▶ Examples?
- Per user files
 - ▶ May want to work with others
 - ▶ Or with particular groups of users

- Where are the files to share?
- Files and Links
 - ▶ Short cut through the file system
 - *Hard and soft*
- Directories
 - ▶ Must provide the other user or group access to your directory
- Remote file systems
 - ▶ Access files on another machine
 - ▶ Must provide the other user or group access to your machine and directory

- **User identity**
 - ▶ UID in UNIX
 - ▶ Security ID in Windows NT
- **Group identity**
 - ▶ GID in UNIX
 - ▶ Group ID in Windows NT
- **Give users and/or groups access to your files to share them**

- **File owner/creator should be able to control:**
 - ▶ what can be done
 - ▶ by whom
- **Types of access**
 - ▶ Read
 - ▶ Write
 - ▶ Execute
 - ▶ Append
 - ▶ Delete
 - ▶ List

Access Control: Mode Bits

- Three classes of users: public, group, owner
- Three types of access permissions:
 - ▶ read, write, execute

- Example:

	Owner	Group	Public
<code>rwX=</code>	111	101	101
Octal	7	5	5

What if no exec access and only owner can read/write?

UNIX File Permissions



-rw-rw-r--	1 pbg	staff	31200	Sep 3 08:30	intro.ps
drwx-----	5 pbg	staff	512	Jul 8 09:33	private/
drwxrwxr-x	2 pbg	staff	512	Jul 8 09:35	doc/
drwxrwx---	2 pbg	student	512	Aug 3 14:13	student-proj/
-rw-r--r--	1 pbg	staff	9423	Feb 24 2003	program.c
-rwxr-xr-x	1 pbg	staff	20471	Feb 24 2003	program
drwx--x--x	4 pbg	faculty	512	Jul 31 10:31	lib/
drwx-----	3 pbg	staff	1024	Aug 29 06:52	mail/
drwxrwxrwx	3 pbg	staff	512	Jul 8 09:35	test/

- An ***access control*** system determines what ***rights*** a particular ***entity*** has for a set of ***objects***
- It answers the question
 - ▶ E.g., do ***you*** have the right to ***read /etc/passwd***
 - ▶ Does ***Alice*** have the right to ***view*** the ***CIS website?***
 - ▶ Do ***students*** have the right to ***share project data?***
 - ▶ Does ***Prof. Butler*** have the right to ***change*** your ***grades?***
- An ***Access Control Policy*** answers these questions

Simplified Access Control



- **Subjects** are the active entities that do things
 - ▶ E.g., **you, Alice, students, Prof. Butler**
- **Objects** are passive things that things are done to
 - ▶ E.g., **/etc/passwd, CS website, project data, grades**
- **Rights** are actions that are taken
 - ▶ E.g., **read, view, share, change**

- **Secrecy**
 - ▶ Don't allow reading by unauthorized subjects
 - ▶ Control where data can be written by authorized subjects
 - Why is this important?
- **Integrity**
 - ▶ Don't permit dependence on lower integrity data/code
 - Why is this important?
 - ▶ What is "dependence"?
- **Availability**
 - ▶ The necessary function must run
 - ▶ Doesn't this conflict with above?

The Access Matrix



- An access matrix is one way to represent policy.
 - ▶ Frequently used mechanism for describing policy
- Columns are objects, subjects are rows.
- To determine if S_i has right to access object O_j , find the appropriate entry.
- Succinct descriptor for $O(|S|*|O|)$ entries
- There is a matrix for each right.

	O	O	O
S	Y	Y	N
S	N	Y	N
S	N	Y	Y

Access Control



- Suppose the private key file for J is object O_1
 - ▶ Only J can read
- Suppose the public key file for J is object O_2
 - ▶ All can read, only J can modify
- Suppose all can read and write from object O_3
- What's the access matrix?

	○	○	○
J	?	?	?
S	?	?	?
S	?	?	?

Trusted Processes



- Does it matter if we do not trust some of J's processes?

	○	○	○
J	R	RW	RW
S	N	R	RW
S	N	R	RW

- Does the following protection state ensure the secrecy of J's private key in O_1 ?

	O	O	O
J	R	RW	RW
S	N	R	RW
S	N	R	RW

- Does the following access matrix protect the integrity of J's public key file O_2 ?

	O_1	O_2	O_3
J	R	RW	RW
S	N	R	RW
S	N	R	RW

Protection vs Security



- Protection
 - ▶ Security goals met under *trusted* processes
 - ▶ Protects against an error by a non-malicious entity
- Security
 - ▶ Security goals met under *potentially malicious* processes
 - ▶ Protects against any malicious entity
 - ▶ Hence, For J:
 - Non-malicious process shouldn't leak the private key by writing it to O_3
 - A potentially malicious process may contain a Trojan horse that can write the private key to O_3

Least Privilege



- Limit permissions to those required and no more
- Consider three processes for user J
 - ▶ Restrict privilege of the process J_1 to prevent leaks

	○	○	○
J	R	R	N
J	N	RW	N
J	N	R	RW

- **File System Interface**
 - ▶ Files
 - ▶ Directories
 - ▶ Partitions
- **Operations on the interface**
 - ▶ Mounting (partitions)
 - ▶ Sharing (files)
 - ▶ Protection (files)

- Next time: File System Implementation