So far

Introduction

Software security

Host security

Network security

Web security

Today

Processes

Users

Next time:

Authorization

Processes

A process is a running program

Processes have memory

Processes execute

Processes have indirect access to I/O



4/27

We've been working with processes already, but we haven't really defined them. A process is just a ________, and it runs on a _______ computer. To see the processes that are currently running on your computer, you can use a GUI like Activity Monitor or Task Manager, or you can use the **ps** command at the command line. Run **ps aux** | wc -l on macOS, FreeBSD. The addresses used directly by your software are _______. This is why, when the same code is executed in ten processes in parallel, they can all refer to the ________, but *without interference*. The OS must *schedule* processes for execution, deciding which process will get to run on which processor at which time. This is a topic for ECE 8400 / ENGI 9875. External resources such as data on a disk or a network connection are generally/typically/almost always represented as

Files

Multiple abstraction layers:



There are lots of ways to look at files from lots of levels of abstraction. Each layer tends to add something nice and/or helpful, but at bottom, all of these layers' concepts of files are rooted in an ______. This is because no library can control a disk: only _______ can. This is another example of how some interfaces are more "real" than others. Interfaces within a single address space — such as library APIs — can wrap up the functionality of lower-level libraries, but when it comes down to the level of _______ (as in a _______ programming language), there is no "real" separation among them. However, interfaces between different _______ of things (users and computers, processes and the OS, software and hardware) tend to be much clearer, "harder" interfaces.

Process file abstractions

Each process has a set of integer *file descriptors*

Can use *system calls* to open, close, read, write, etc.

```
int fd = open("/home/jon/hello.txt", 0_RDONLY);
/* ... */
write(fd, some_data_bytes, data_length);
/* ... */
```

A file descriptor can then get wrapped up, together with some buffering for performance, into a
That can then get wrapped up into a C++
, which might be used to implement a Java
java.io.FileWriter.

File I/O system calls

From a process' perspective:

- system calls are C functions
- files are named by small integer indices (e.g., FD 3)
- each process has its own array of files
- ... and that's enough detail for now

7/27

There's a lot more detail to be dug into about file descriptors (e.g., how libc communicates with the OS kernel), and we'll go into that detail in ECE 8400 / ENGI 9875. In fact, our first lab in that course will have you invoking system calls via nothing but native instructions in assembly code! For now, however, this is all that we need to know in order to start talking about host security.

Processes

Processes have *memory* – virtual memory

Processes *execute* — threads

Processes have (indirect) access to resources – files

Processes execute on behalf of ... ???

Users

Usernames

User IDs

User authentication

User authorization

What's a user? A human being, sure, but how does the computer	see a user? How do we think of
users?	
We often think of users as being identified by	: short, human-readable names
that are unique to If I'm using a compu	uter, I can see my current
username by running whoami(1).	
Something a bit more meaningful to the computer, however, is no	ot a but a
User IDs are still short and	, but instead of strings,
they're You can see information about your u	ser ID (and group IDs!) by
running id(1) (or just whoami on Windows).	
How does the computer know that the person sitting in front of i	it is <i>actually</i> the person they claim
to be? That's called user <i>authentication</i> , and we'll get into it when	we start talking about passwords
(and their problems, and alternatives).	
Next time, we'll start getting into user <i>authorization</i> : saying	is allowed to do
to .	

User databases

Where is user information stored?

- Active Directory
- Binary databases (e.g., Berkeley DB caches)
- NIS[plus] (though not really any more)
- OpenLDAP
- Text files (e.g., /etc/passwd)

10/27

We'll examine user details specific to discretionary access control next time, but for now we will look at the contents of files like /etc/nsswitch.conf. This *Name Service Switch* config file tells a Unix machine where to find information about lots of kinds of names: users, groups, protocols, shells, etc. It can point us at files, at binary db files (typically used as a cache), at ldap (which might actually be Active Direcftory) or at the basically-defunct nisplus.

User database: files

\$FreeBSD\$
#
root:*:0:0:Charlie &:/root:/bin/sh
toor:*:0:0:Bourne-again Superuser:/root:
<pre>daemon:*:1:1:Owner of many system processes:/root:/usr/sbin/nologin</pre>
operator:*:2:5:System &:/:/usr/sbin/nologin
bin:*:3:7:Binaries Commands and Source:/:/usr/sbin/nologin
tty:*:4:65533:Tty Sandbox:/:/usr/sbin/nologin
kmem:*:5:65533:KMem Sandbox:/:/usr/sbin/nologin
games:*:7:13:Games pseudo-user:/:/usr/sbin/nologin
news:*:8:8:News Subsystem:/:/usr/sbin/nologin
man:*:9:9:Mister Man Pages:/usr/share/man:/usr/sbin/nologin
<pre>sshd:*:22:22:Secure Shell Daemon:/var/empty:/usr/sbin/nologin</pre>
#

11/27

We can examine local user details via:

cat /etc/passwd

If we look at **/etc/passwd** on a local Unix-like machine, we should see users like the user who set up the box. On a LabNet machine, however, _______...

User database: LDAP

Lightweight Directory Access Protocol

Queries directory servers:

- Active Directory
- ApacheDS
- FreeIPA / 389 directory server
- OpenLDAP

12/27

In practice, Microsoft Active Directory is absolutely dominant, as most large networks support large numbers of Windows PCs.

In a LabNet environment, we can query lots of interesting details from LDAP using commands like:

ldapsearch -H ldaps://dogbert.cs.mun.ca "(uid=p15jra)"

Next

DAC (today)

MAC (Tuesday)

Capabilities (later)

13/27

There are lots of "AC"s that get tossed around these days (DAC, MAC, ABAC, RBAC, etc.), but we'll concentrate on three fundamental forms of authorization:

- discretionary access control (DAC)
- mandatory access control (MAC)
- capabilities

Other schemes can often typically implemented in terms of the above. For example, role-based access control (RBAC) can be implemented using MAC primitives.

DAC

Discretionary access control

Organizing principle:

Files and directories have **owners** who have the **discretion** to say who gets to access them.

DAC

Discretionary access control

Organizing principle:

Files and directories have **owners** who have the **discretion** to say who gets to access them.

Major implementations:

Unix permissions Access control lists (ACLs)

14/27

We saw an example of Unix permissions in Lab 0, when we had to use the **Chmod** command to make the binary executable **game**, well, executable!

Unix DAC

Users:

User-readable names, user IDs in /etc/passwd* ... or elsewhere

Groups:

Numeric *group ID* with names in /etc/group Users can be members of multiple groups

This file doesn't contain what you might think it does... stay tuned for password hashing in later lectures!

We often think of users as being identified by _____: short, human-readable names that are unique to ______. If I'm using a computer, I can see my current username by running whoami(1). Something a bit more meaningful to the computer, however, is not a ______ but a ______. User IDs are still short and _______, but instead of strings, they're ______. You can see information about your user ID (and group IDs!) by running id(1) (or just whoami on Windows). Most Unix-like computers have a *Name Service Switch* configuration file in /etc/nsswitch.conf that tells the host where to find names for users, groups, networks, hosts, RPCs... In addition to a user ID, every user can be a member of multiple *groups* that are identified by integer *group ID*.

Unix file permissions

Each file has **read**, **write** and **execute** permission for each of **owner**, **group** and **other** users:

[jon websit	e]\$ ls					
drwxr-xr-x		8B	Mar	26	2017	assets
-rw-rr		948B		26	15:37	config.yaml
drwxr-xr-x		10B	Feb	13	23:19	content
-rwxr-xr-x		271B		13	2017	deploy
drwxr-xr-x		9B		22	23:14	layouts
drwxr-xr-x	12 jon	13B		24	16:18	static

File owner can set permissions with chmod command

16/27

These permissions sound very much like virtual memory permissions, and they do indeed have the same meanings. However, their enforcement is very different!

Unix permissions

For each of (owner, group, anyone):

Value	Meaning
4	Readable
2	Writable
1	Executable

Octal example: 0644 (writable by owner, readable by anyone).

\$ chmod 644 file.txt
\$ chmod g+rx game

17/27

These power-of-two values can be XOR'ed together.

This is one of the very few instances of an octal representation that you're likely to see anywhere!

Changing file owner

Owner has discretion to set file access permissions... but how do we set the owner?

Answer: chown(1)

But:

\$ chown alice foo.txt
chown: foo.txt: Operation not permitted

18/27

Show man page for chown(2)

Superuser

a.k.a., root user

- UID 0
- can change file owner, chmod other users' files
- second-level objective for many attacks

The root user is allowed to violate the DAC policy, overriding the access control decisions made by a file's owner (and even ______!). To "get root" is to gain administrative control over a computer, whether legitimately becoming a system administrator ("yeah, I've got root on that box") or otherwise. Many, many attacks against systems start by gaining ______ (running whatever the attacker wants within a process, with that process' credentials) and then a _______ attack against a service that allows the attacker to _______ to administrative access.

Root-only programs

- lots of tools require *root privilege*:
 - filesystem management
 - package managers
 - service management
 - often via sudo(8)



Exercise: Consider how a user who can control all software installation on a computer could violate another user's security policy

21/27

We don't want just any user being able to, e.g., control a mounted filesystem or install a package. Why not?

For all of these examples, being in control of such a subsystem would allow a user to be able to

Root-only programs

- some programs require root privilege
- some programs must be runnable by anyone
- some are both!
- e.g., ping(8), even intel_backlight(1)!

\$ ls -l `which intel_backlight`
-r-sr-xr-x 1 root wheel 16K Feb 26 17:03 /usr/local/bin/intel_bac

Since we don't want just anybody controlling critical subsystems, some programs	require root
privilege in order to do their work. For example, I can	on my
machine from an ordinary user account, but I can only	to system
locations (e.g., /usr/local/bin) as root.	
Some programs, however, require privilege to do their job and <i>also</i> need to be run	by ordinary
users! We can implement such functionality, overriding the normal DAC policy, u	using setuid
and setgid software.	

setuid/setgid programs

setuid: set effective UID to file owner's UID on run
setgid: set effective GID to file group's GID on run

Can query *real* or *effective* UID/GID:



23/27

Example: getuid.c

DAC

Organizing principle:

Files and directories have **owners** who have the **discretion** to say who gets to access them.

Major implementations:

Unix permissions Access control lists (ACLs)

24/27

So... how about that other implementation, ACLs?

ACLs: access control lists

(the other way of doing discretionary access control)

- explicit list of users, groups
- independent permissions for each
- useful for complex authorization on multiuser shared systems* implemented in NFSv4, POSIX 1e, Windows/SMB...

Desktop Pro	operties		L
Security Locati	an		
Object name: C:\Users\jon\Desk	top		
Group or user names:			
SR. SYSTEM			
3 jon (WIN-AJU5GMIHJ6T jon)			
as Administrators (WIN-AIU5GMIH	U61 \Administra	tors)	
To change permissions, click Edit.		Edit	
Permissions for SYSTEM	Alow	Deny	_
Full control	~		^
Modify	\checkmark		
Read & execute	\checkmark		
List folder contents	\checkmark		
Read	\checkmark		
Write	~		¥
For special permissions or advanced click Advanced.	settings.	Advanced	

25/27

ACLs are useful, but it's also very easy to write an ACL that you yourself don't understand! Just look at some of the literature around trying to make ACLs understandable by users: Intentional access management: making access control usable for end-users Cao and Iverson, *SOUPS '06: Proceedings of the second symposium on Usable privacy and security, July 2006.* DOI: https://doi.org/10.1145/1143120.1143124 Relating declarative semantics and usability in access control Krishnan, Tripunitara, Chik and Bergstrom, *SOUPS '12: Proceedings of the Eighth Symposium on Usable Privacy and Security*, July 2012. DOI: https://doi.org/10.1145/2335356.2335375 The poor usability of OpenLDAP Access Control Lists, Chen, Punchhi and Tripunitara, *IET Information Security* 17(1), January 2023. DOI: https://doi.org/10.1049/ise2.12079 If people can publish "how to make ACLs usable" over three decades... maybe there's a problem

with ACLs.

Summary

Processes and Users

Authorization

DAC (today)

MAC (Thursday)

Capabilities (later)