

View access control as a matrix

		Objects				
		File 1	File 2	File 3	...	File n
Subjects	User 1	read	write	-	-	read
	User 2	write	write	write	-	-
	User 3	-	-	-	read	read
	...					
	User m	read	write	read	write	read


- Subjects (processes/users) access objects (e.g., files)
- Each cell of matrix has allowed permissions

Two ways to slice the matrix

- **Along columns:**
 - Kernel stores list of who can access object along with object
 - Most systems you've used probably do this
 - Examples: Unix file permissions, Access Control Lists (ACLs)
- **Along rows:**
 - Capability systems do this
 - More on these later...

Specifying policy

- Manually filling out matrix would be tedious
- Use tools such as groups or *role-based access control*:



1/33

2/33

Example: Unix protection

- Each process has a User ID & one or more group IDs
- **System stores with each file:**
 - User who owns the file and group file is in
 - Permissions for user, any one in file group, and other
- **Shown by output of `ls -l` command:**

```

user group other owner group
-rw- rw- r-- dm cs140 ... index.html

```

- Each group of three letters specifies a subset of **r**ead, **w**rite, and **e**xecute permissions
- User permissions apply to processes with same user ID
- Else, group permissions apply to processes in same group
- Else, other permissions apply

3/33

4/33

Unix continued

- **Directories have permission bits, too**
 - Need write perm. on directory to create or delete a file
- **Special user root (UID 0) has all privileges**
 - E.g., Read/write any file, change owners of files
 - Required for administration (backup, creating new users, etc.)
- **Example:**
 - `drwxr-xr-x 56 root wheel 4096 Apr 4 10:08 /etc`
 - Directory writable only by root, readable by everyone
 - Means non-root users cannot directly delete files in `/etc`
 - **Execute** permission means ability to use pathnames in the directory, separate from **read** permission which allows listing

5/33

Example: Login runs as root

- **Unix users typically stored in files in `/etc`**
 - Files `passwd`, `group`, and often `shadow` or `master.passwd`
- **For each user, files contain:**
 - Textual username (e.g., "dm", or "root")
 - Numeric user ID, and group ID(s)
 - One-way hash of user's password: $\{\text{salt}, H(\text{salt}, \text{passwd})\}$
 - Other information, such as user's full name, login shell, etc.
- **`/usr/bin/login` runs as root**
 - Reads username & password from terminal
 - Looks up username in `/etc/passwd`, etc.
 - Computes $H(\text{salt}, \text{typed password})$ & checks that it matches
 - If matches, sets group ID & user ID corresponding to username
 - Execute user's shell with `execve` system call

7/33

Non-file permissions in Unix

- **Many devices show up in file system**
 - E.g., `/dev/tty1` permissions just like for files
- **Other access controls not represented in file system**
- **E.g., must usually be root to do the following:**
 - Bind any TCP or UDP port number less than 1,024
 - Change the current process's user or group ID
 - Mount or unmount file systems
 - Create device nodes (such as `/dev/tty1`) in the file system
 - Change the owner of a file
 - Set the time-of-day clock; halt or reboot machine

6/33

Setuid

- **Some legitimate actions require more privs than UID**
 - E.g., how should users change their passwords?
 - Stored in root-owned `/etc/passwd` & `/etc/shadow` files
- **Solution: Setuid/setgid programs**
 - Run with privileges of file's owner or group
 - Each process has *real* and *effective* UID/GID
 - *real* is user who launched setuid program
 - *effective* is owner/group of file, used in access checks
- **Shown as "s" in file listings**
 - `-rws--x--x 1 root root 38464 Jan 26 14:26 /bin/passwd`
 - Obviously need to own file to set the setuid bit
 - Need to own file and be in group to set setgid bit

8/33

Setuid (continued)

- **Examples**
 - E.g., /usr/bin/passwd – changes user’s password
 - E.g., /bin/su – acquire new user ID with correct password
 - E.g., /usr/bin/netstat – lists network connections (by reading kernel memory on some OSes)
- **Have to be very careful when writing setuid code**
 - Attackers can run setuid programs any time (no need to wait for root to run a vulnerable job)
 - Attacker controls many aspects of program’s environment
- **Example attacks when running a setuid program**
 - Change PATH or IFS if setuid prog calls system(3)
 - Set maximum file size to zero (if app rebuilds DB)
 - Close fd 2 before running program—may accidentally send error message into protected file

9/33

Other permissions

- **When can proc. A send a signal to proc. B w. kill?**
 - Allow if sender and receiver have same effective UID
 - But need ability to kill processes you launch even if suid
 - So allow if real UIDs match, as well
 - Can also send SIGCONT w/o UID match if in same session
- **Debugger system call ptrace**
 - Lets one process modify another’s memory
 - Setuid gives a program more privilege than invoking user
 - So don’t let process ptrace more privileged process
 - E.g., Require sender to match real & effective UID of target
 - Also disable/ignore setuid if ptraced target calls exec
 - Exception: root can ptrace anyone

10/33

A security hole

- **Even without root or setuid, attackers can trick root owned processes into doing things...**
- **Example: Want to clear unused files in /tmp**
- **Every night, automatically run this command as root:**

```
find /tmp -atime +3 -exec rm -f -- {} \;
```
- **find identifies files not accessed in 3 days**
 - executes rm, replacing {} with file name
- **rm -f -- path deletes file path**
 - Note “--” prevents path from being parsed as option
- **What’s wrong here?**

11/33

An attack

find/rm	Attacker
	creat (“/tmp/badetc/passwd”)
readdir (“/tmp”) → “badetc”	
lstat (“/tmp/badetc”) → DIRECTORY	
readdir (“/tmp/badetc”) → “passwd”	
	rename (“/tmp/badetc” → “/tmp/x”)
	symlink (“/etc”, “/tmp/badetc”)
unlink (“/tmp/badetc/passwd”)	

- **Time-of-check-to-time-of-use (TOCTTOU) bug**
 - find checks that /tmp/badetc is not symlink
 - But meaning of file name changes before it is used

12/33

xterm command

- Provides a terminal window in X-windows
- Used to run with `setuid` root privileges
 - Requires kernel pseudo-terminal (pty) device
 - Required root privs to change ownership of pty to user
 - Also writes protected `utmp/wtmp` files to record users

- Had feature to log terminal session to file

```
fd = open (logfile, O_CREAT|O_WRONLY|O_TRUNC, 0666);  
/* ... */
```

- What's wrong here?

13/33

xterm command

- Provides a terminal window in X-windows
- Used to run with `setuid` root privileges
 - Requires kernel pseudo-terminal (pty) device
 - Required root privs to change ownership of pty to user
 - Also writes protected `utmp/wtmp` files to record users

- Had feature to log terminal session to file

```
if (access (logfile, W_OK) < 0)  
    return ERROR;  
fd = open (logfile, O_CREAT|O_WRONLY|O_TRUNC, 0666);  
/* ... */
```

- `xterm` is root, but shouldn't log to file user can't write
- `access` call avoids dangerous security hole
 - Does permission check with *real*, not *effective* UID

13/33

xterm command

- Provides a terminal window in X-windows
- Used to run with `setuid` root privileges
 - Requires kernel pseudo-terminal (pty) device
 - Required root privs to change ownership of pty to user
 - Also writes protected `utmp/wtmp` files to record users

- Had feature to log terminal session to file

```
if (access (logfile, W_OK) < 0)  
    return ERROR;  
fd = open (logfile, O_CREAT|O_WRONLY|O_TRUNC, 0666);  
/* ... */
```

- `xterm` is root, but shouldn't log to file user can't write
- `access` call avoids dangerous security hole
 - Does permission check with *real*, not *effective* UID
 - **Wrong: Another TOCTTOU bug**

13/33

An attack

xterm	Attacker
	<code>creat ("/tmp/X")</code>
<code>access ("/tmp/X") → OK</code>	
	<code>unlink ("/tmp/X")</code>
	<code>symlink ("/tmp/X" → "/etc/passwd")</code>
<code>open ("/tmp/X")</code>	

- Attacker changes `/tmp/X` between check and use
 - `xterm` unwittingly overwrites `/etc/passwd`
 - Another TOCTTOU bug
- OpenBSD man page: "CAVEATS: `access()` is a potential security hole and should never be used."

14/33

SSH configuration files

- **SSH 1.2.12 – secure login program, runs as root**
 - Needs to bind TCP port under 1,024 (privileged operation)
 - Needs to read client private key (for host authentication)
- **Also needs to read & write files owned by user**
 - Read configuration file `~/.ssh/config`
 - Record server keys in `~/.ssh/known_hosts`
- **Author wanted to avoid TOCTTOU bugs:**
 - First binds socket & reads root-owned secret key file
 - Then drops all privileges before accessing user files—real and effective user IDs those of invoking user
 - Idea: avoid using any user-controlled arguments/files until you have no more privileges than the user
 - What might still have gone wrong?

15/33

A Linux security hole

- **Some programs acquire then release privileges**
 - E.g., `su user` is `setuid root`, becomes user if password correct
- **Consider the following:**
 - A and B unprivileged processes owned by attacker
 - A ptraces B
 - A executes “`su user`” to its own identity
 - While `su` is superuser, B execs `su root` (A is superuser, so this is not disabled)
 - A types password, gets shell, and is attached to `su root`
 - Can manipulate `su root`’s memory to get root shell

17/33

Trick question: ptrace bug

- **Actually do have more privileges than user!**
 - Bound privileged port and read host private key
- **Dropping privs allows user to “debug” SSH**
 - Depends on OS, but at the time several had *ptrace* implementations that made SSH vulnerable
- **Once in debugger**
 - Could use privileged port to connect anywhere
 - Could read secret host key from memory
 - Could overwrite local user name to get privs of other user
- **The fix: restructure into 3 processes!**
 - Perhaps overkill, but really wanted to avoid problems

16/33



- **Previous examples show two limitations of Unix**
- **Many OS security policies *subjective* not *objective***
 - When can you signal/debug process? Re-bind network port?
 - Rules for non-file operations somewhat incoherent
 - Even some file rules weird (Creating hard links to files)
- **Correct code is much harder to write than incorrect**
 - Delete file without traversing symbolic link
 - Read SSH configuration file (requires 3 processes??)
 - Write mailbox owned by user in dir owned by root/mail
- **Don’t *just* blame the application writers**
 - Must also blame the interfaces they program to

18/33

Another security problem [Hardy]

- **Setting: A multi-user time sharing system**
 - This time it's not Unix
- **Wanted fortran compiler to keep statistics**
 - Modified compiler `/sysx/fort` to record stats in `/sysx/stat`
 - Gave compiler "home files license"—allows writing to anything in `/sysx` (kind of like Unix `setuid`)
- **What's wrong here?**

A confused deputy

- **Attacker could overwrite any files in `/sysx`**
 - System billing records kept in `/sysx/bill` got wiped
 - Probably command like `fort -o /sysx/bill file.f`
- **Is this a bug in the compiler `fort`?**
 - Original implementors did not anticipate extra rights
 - Can't blame them for unchecked output file
- **Compiler is a "confused deputy"**
 - Inherits privileges from invoking user (e.g., read `file.f`)
 - Also inherits privileges from home files license
 - Which master is it serving on any given system call?
 - OS doesn't know if it just sees `open ("/sysx/bill", ...)`

19/33

20/33

Recall access control matrix

		Objects				
		File 1	File 2	File 3	...	File n
Subjects	User 1	read	write	-	-	read
	User 2	write	write	write	-	-
	User 3	-	-	-	read	read
	...					
	User m	read	write	read	write	read

21/33

Capabilities

- **Slicing matrix along rows yields capabilities**
 - E.g., For each process, store a list of objects it can access
 - Process explicitly invokes particular capabilities
- **Can help avoid confused deputy problem**
 - E.g., Must give compiler an argument that both specifies the output file and conveys the capability to write the file (think about passing a file descriptor, not a file name)
 - So compiler uses no *ambient authority* to write file
- **Three general approaches to capabilities:**
 - Hardware enforced (Tagged architectures like **M-machine**)
 - Kernel-enforced (**Hydra**, **KeyKOS**)
 - Self-authenticating capabilities (like **Amoeba**)
- **Good history in [Levy]**

22/33