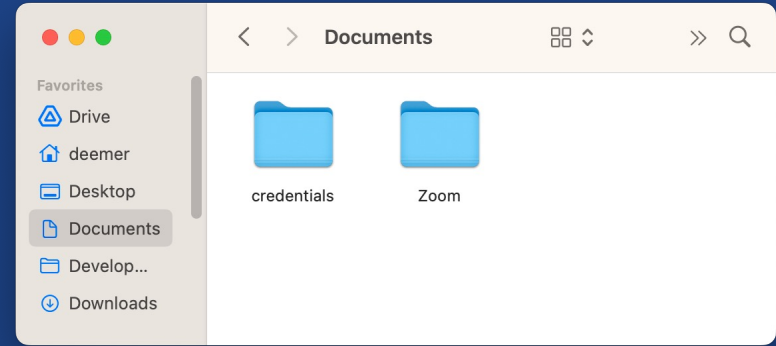


Operating Systems Security III

CS 1660: Introduction to Computer Systems
Security



=> How many of these can read your browser history?

.... all of them?!?!

```
deemer@ceres:~$ ls -la cookies.sqlite  
rwxr-x--- 1 deemer deemer 9 Mar 12 16:40 cookies.sqlite
```

```
deemer@ceres$ some_random_app cookies.sqlite
```

.... all of them?!?!

```
deemer@ceres:~$ ls -la cookies.sqlite  
rwxr-x--- 1 deemer deemer 9 Mar 12 16:40 cookies.sqlite
```

```
deemer@ceres$ some_random_app cookies.sqlite
```

```
. . .  
access("cookies.sqlite", F_OK) = 0  
openat(AT_FDCWD, "cookies.sqlite", O_RDONLY) = 3  
. . .
```

*=> Access is just a syscall!
Works as long as permissions check out 😲*

Discretionary Access Control

Owner of a resource decides on how it's used

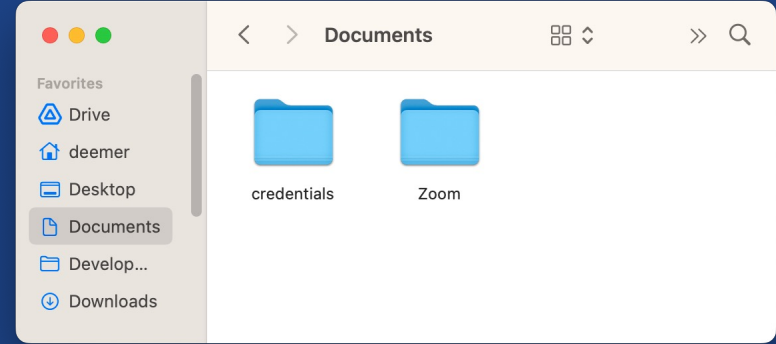
- Privileges depend on current user (and some groups)
- To elevate: admin user (root) vs. other users

Discretionary Access Control

Owner of a resource decides on how it's used

- Privileges depend on current user (and some groups)
- To elevate: admin user (root) vs. other users

Is this really what we want?



How many of these should be able to read your browser history?

And why?

Why?

Would like to get closer to...

Why?

- File permissions are very coarse
- Apps might not be trusted
- Apps might get compromised

Would like to get closer to...

✨ ✨ ✨ *Principle of Least Privilege* ✨ ✨ ✨

✨ ✨ ✨ Principle of Least Privilege ✨ ✨ ✨

An application should only be able to perform the operations necessary for its intended purpose

Isolation and Sandboxing

Run (untrusted) application in such a way that it has limited access to resources => only what it needs

Examples we've been discussing

Isolation and Sandboxing

Run (untrusted) application in such a way that it has limited access to resources => only what it needs

Examples we've been discussing

- Application sandboxing
- Namespaces (sort of)
- Containers (sort of)
- Virtual machines

Example: Gradescope

Example: web browsers

The screenshot shows the macOS Activity Monitor window with the 'Memory' tab selected. The window title is 'Activity Monitor My Processes'. The top navigation bar includes 'CPU', 'Memory', 'Energy', 'Disk', and 'Network' tabs, along with a search bar. The main area displays a table of processes with columns for Process Name, Memory, Threads, Ports, PID, and User. The 'Firefox' process is highlighted, showing it is using 10.90 GB of memory. Below the table, a 'MEMORY PRESSURE' section provides summary statistics.

Process Name	Memory	Threads	Ports	PID	User
Finder	362.6 MB	6	1,402	720	deemer
Firefox	10.90 GB	162	1,404	785	deemer
FirefoxCP Isolated Service Worker	33.4 MB	32	109	12796	deemer
FirefoxCP Isolated Web Content	212.2 MB	29	121	1293	deemer
FirefoxCP Isolated Web Content	97.5 MB	29	106	1296	deemer
FirefoxCP Isolated Web Content	37.0 MB	28	110	1446	deemer
FirefoxCP Isolated Web Content	56.3 MB	28	111	1262	deemer
FirefoxCP Isolated Web Content	2.05 GB	33	139	1265	deemer
FirefoxCP Isolated Web Content	226.3 MB	30	124	1268	deemer
FirefoxCP Isolated Web Content	157.4 MB	28	119	1286	deemer
FirefoxCP Isolated Web Content	59.8 MB	28	103	1289	deemer
FirefoxCP Isolated Web Content	381.6 MB	29	119	1292	deemer
FirefoxCP Isolated Web Content	54.3 MB	28	103	1390	deemer
FirefoxCP Isolated Web Content	64.4 MB	28	102	1295	deemer
FirefoxCP Isolated Web Content	318.5 MB	28	112	1264	deemer
FirefoxCP Isolated Web Content	46.3 MB	28	111	1362	deemer
FirefoxCP Isolated Web Content	265.1 MB	30	130	1267	deemer
FirefoxCP Isolated Web Content	183.6 MB	34	137	1273	deemer

MEMORY PRESSURE

Physical Memory:	32.00 GB	App Memory:	13.22 GB
Memory Used:	26.00 GB	Wired Memory:	2.54 GB
Cached Files:	5.96 GB		

Example: web browsers

Browsers run a lot of untrusted code...

=> Worker processes that render pages run with fewer privileges

=> Site isolation: one process per site (or per tab)

Example: web browsers

Browsers run a lot of untrusted code...

=> Worker processes that render pages run with fewer privileges


=> Site isolation: one process per site (or per tab)

This protection is made possible by the following changes in Chrome's behavior:

[Link](#)

- Cross-site documents are always put into a different process, whether the navigation is in the current tab, a new tab, or an iframe (i.e., one web page embedded inside another). Note that only a subset of sites are isolated on Android, to reduce overhead.
- Cross-site data (such as HTML, XML, JSON, and PDF files) is not delivered to a web page's process unless the server says it should be allowed (using [CORS](#)).
- Security checks in the browser process can detect and terminate a misbehaving renderer process (only on desktop platforms for the time being).

=> Can enforce Same-Origin Policy with separate processes!



How many of these should be able to read your browser history?



Typical desktop environment: User has a lot of freedom, can modify system

=> Built in an era where we weren't downloading lots of untrusted code...

*What if we could start over?
(sort of)*





Mobile Operating System (iOS, Android)

=> Does user have root?



Mobile Operating Systems (iOS, Android)

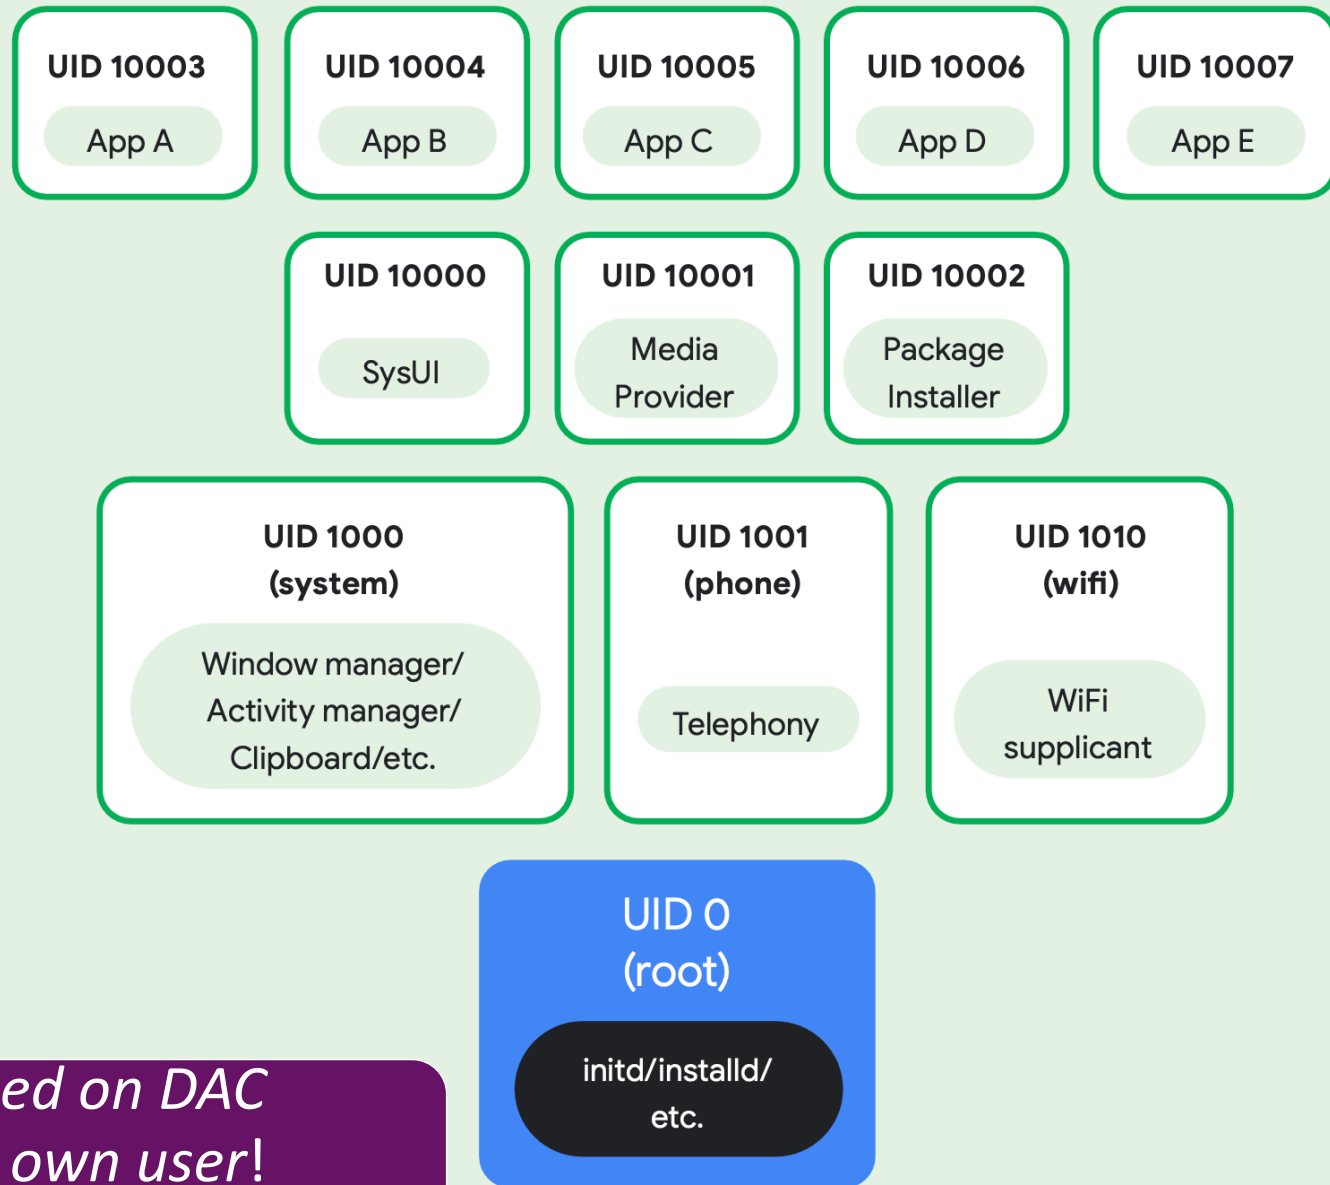
=> Designed more as “secure by default” to restrict app privileges

Example: Android

- Based on Linux
- Many levels of application *sandboxing*
- Applications request permissions at installation and runtime
=> OS and Android Platform provides access

Nice writeups you can read:

- [Android Security Paper 2023](#)
- [The Android Platform Security Model](#)



*Core sandboxing based on DAC
=> Every app gets its own user!
=> There is much more than this...*

Process list

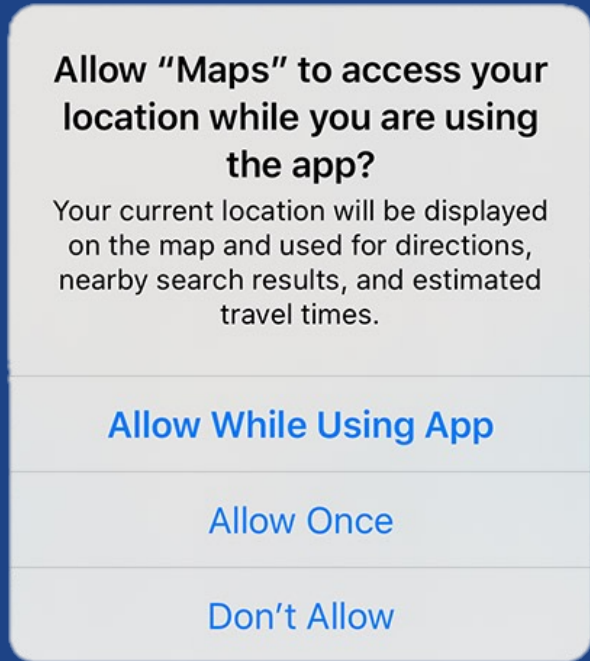

Tasks: 907 total, 1 running, 906 sleeping, 0 stopped, 0 zombie
Mem: 7618164K total, 7339296K used, 278868K free, 3240K buffers
Swap: 3145724K total, 2944924K used, 200800K free, 2456224K cached
%cpu 5%user 1%nice 8%sys 785%idle 0%iow 1%irq 1%siq 0%host

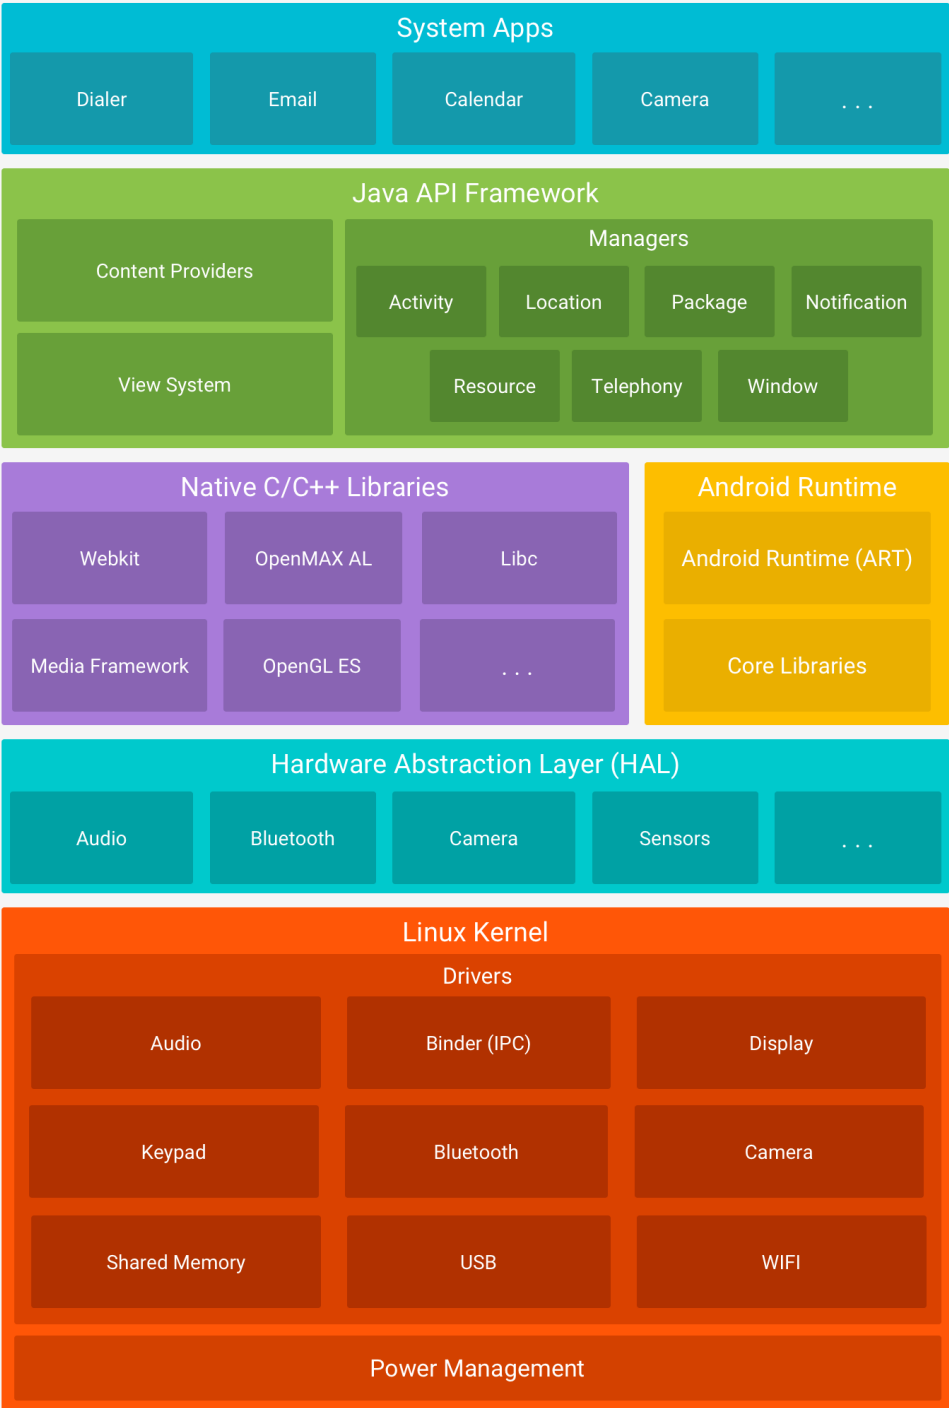
USER	PR	NI	VIRT	RES	SHR	S	%CPU	[%MEM]	TIME+	ARGS
system	18	-2	22G	676M	523M	S	2.6	9.0	270:38.91	system_server
u0_a215	20	0	16G	76M	51M	S	1.6	1.0	22:00.19	com.shannon.imsservice
shell	20	0	10G	5.5M	3.5M	R	1.3	0.0	0:01.10	top
u0_a167	20	0	34G	328M	215M	S	1.0	4.4	0:39.46	com.google.android.googlequ
u0_a215	20	0	16G	59M	39M	S	1.0	0.7	1:28.44	.ShannonImService
system	20	0	10G	2.7M	2.4M	S	1.0	0.0	38:17.24	android.hardware.power.stat
u0_a87	20	0	16G	68M	48M	S	0.6	0.9	0:00.10	com.android.providers.calend
radio	20	0	16G	117M	66M	S	0.6	1.5	18:55.63	com.android.phone
u0_a233	0	0	18G	297M	165M	S	0.6	3.9	56:24.60	com.android.systemui
root	20	0	0	0	0	I	0.3	0.0	0:00.56	[kworker/u16:5-cp2ap_wakeup
root	20	0	0	0	0	I	0.3	0.0	0:00.51	[kworker/u16:3-events_unboun
u0_a343	20	0	17G	102M	70M	S	0.3	1.3	4:20.28	com.whatsapp
u0_a207	20	0	22G	52M	37M	S	0.3	0.6	13:55.88	com.google.android.connectiv
radio	20	0	11G	5.8M	5.8M	S	0.3	0.0	5:09.88	rild_exynos
root	RT	0	0	0	0	S	0.3	0.0	3:51.72	[sugov:6]
root	RT	0	0	0	0	S	0.3	0.0	27:11.68	[sugov:0]
system	20	0	10G	2.8M	2.1M	S	0.3	0.0	8:07.16	servicemanager

Process list

Tasks: 907 total, 1 running, 906 sleeping, 0 stopped, 0 zombie
Mem: 7618164K total, 7339296K used, 278868K free, 3240K buffers
Swap: 3145724K total, 2944924K used, 200800K free, 2456224K cached
%cpu 5%user 1%nice 8%sys 785%idle 0%iow 1%irq 1%irq 0%host

USER	PR	NI	VIRT	RES	SHR	S	%CPU	[%MEM]	TIME+	ARGS
system	18	-2	22G	676M	523M	S	2.6	9.0	270:38.91	system_server
u0_a215	20	0	16G	76M	51M	S	1.6	1.0	22:00.19	com.shannon.imsservice
shell	20	0	10G	5.5M	3.5M	R	1.3	0.0	0:01.10	top
u0_a167	20	0	34G	328M	215M	S	1.0	4.4	0:39.46	com.google.android.googlequ
u0_a215	20	0	16G	59M	39M	S	1.0	0.7	1:28.44	.ShannonImService
system	20	0	10G	2.7M	2.4M	S	1.0	0.0	38:17.24	android.hardware.power.stats
u0_a87	20	0	16G	68M	48M	S	0.6	0.9	0:00.10	com.android.providers.calend
radio	20	0	16G	117M	66M	S	0.6	1.5	18:55.63	com.android.phone
u0_a233	0	0	18G	297M	165M	S	0.6	3.9	56:24.60	com.android.systemui
root	20	0	0	0	0	I	0.3	0.0	0:00.56	[kworker/u16:5-cp2ap_wakeup_
root	20	0	0	0	0	I	0.3	0.0	0:00.51	[kworker/u16:3-events_unboun
u0_a343	20	0	17G	102M	70M	S	0.3	1.3	4:20.28	com.whatsapp
u0_a207	20	0	22G	52M	37M	S	0.3	0.6	13:55.88	com.google.android.connectiv
radio	20	0	11G	5.8M	5.8M	S	0.3	0.0	5:09.88	rild_exynos
root	RT	0	0	0	0	S	0.3	0.0	3:51.72	[sugov:6]
root	RT	0	0	0	0	S	0.3	0.0	27:11.68	[sugov:0]
system	20	0	10G	2.8M	2.1M	S	0.3	0.0	8:07.16	servicemanager





Examining an app...

requested permissions:

```
android.permission.ACCESS_WIFI_STATE
android.permission.INTERNET
android.permission.ACCESS_NETWORK_STATE
android.permission.WAKE_LOCK
android.permission.GET_ACCOUNTS
android.permission.RECEIVE_BOOT_COMPLETED
```

. . .

install permissions:

```
com.google.android.finsky.permission.BIND_GET_INSTALL_REFERRER_SERVICE: granted\
=true
com.google.android.c2dm.permission.RECEIVE: granted=true
android.permission.USE_CREDENTIALS: granted=true
android.permission.MODIFY_AUDIO_SETTINGS: granted=true
android.permission.FOREGROUND_SERVICE: granted=true
android.permission.CHANGE_WIFI_STATE: granted=true
android.permission.FOREGROUND_SERVICE_DATA_SYNC: granted=true
android.permission.ACCESS_NETWORK_STATE: granted=true
android.permission.USE_FINGERPRINT: granted=true
android.permission.READ_BASIC_PHONE_STATE: granted=true
```

SELinux Policies on resources

```
$ $ ls -laZ
```

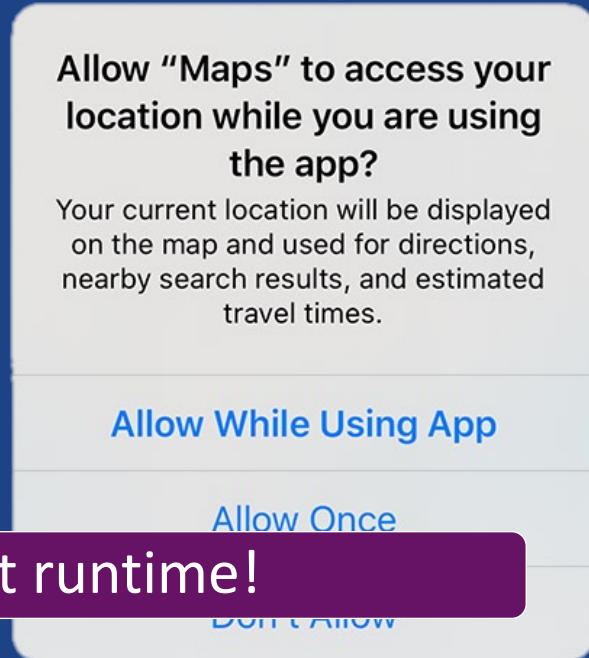
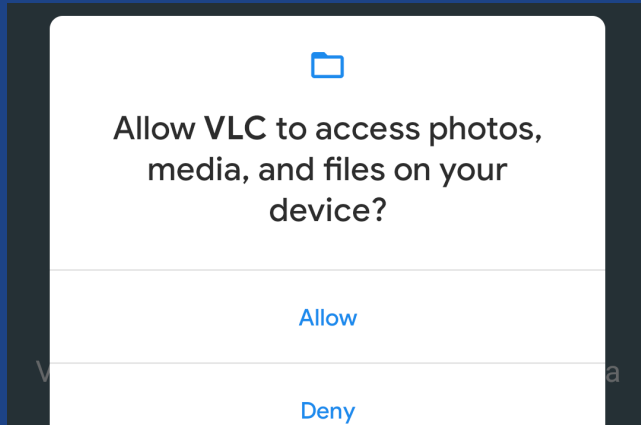
```
lrw-r--r--    1 root    root    u:object_r:rootfs:s0      11 bin -> /system/bin
drwxr-xr-x    2 root    root    u:object_r:tmpfs:s0     4096 debug_ramdisk
drwxr-xr-x   24 root    root    u:object_r:device:s0     160 dev
drwxr-xr-x   13 root    shell   u:object_r:vendor_file:s0 4096 vendor
drwxr-xr-x    5 root    root    u:object_r:vendor_file:s0 4096 vendor_dlkm
drwx--x---    4 shell   everybody u:object_r:mnt_user_file:s0 80 storage
```

```
$ sesearch -A selinux_policy
```

```
allow adb adb_keys_file:dir search;
allow adb adb_keys_file:file { getattr ioctl lock map open read watch watch_reads };
```




```
access("cookies.sqlite", F_OK) = 0  
openat(AT_FDCWD, "cookies.sqlite", O_RDONLY) = 3
```



=> Fine-grained permissions at runtime!

...at compile time?

Other ways?

- What does it mean for the user to be "unprivileged"?
- What does it mean for code run by a user to be "unprivileged"?
- What do we want that code to be able to do?
=> How much do we trust the user? The code?

Other ways?

- What does it mean for the user to be "unprivileged"?
- What does it mean for code run by a user to be "unprivileged"?
- What do we want that code to be able to do?
=> How much do we trust the user? The code?
- sudo is pretty coarse-grained...