

2550 Intro to cybersecurity

L20: systems

abhi shelat

Threat Model

Principles

Intro to System Architecture

Hardware Support for Isolation

Examples

• UNIX - AT&T
Berkeley


- IBM Sys 360 .



1980
era
PC
architecture.

*I/O
Display*

Embedded DisplayPort 1.4b
3 DDI HDMI 2.0b, DP 1.4, HDCP 2.2



GPU

Mem

DDR4/LPDDR3 2Ch
DDR4/x 3733



USB 3.1 (10 Gbps)
USB 3.0 (5 Gbps)

I/O peripheral

Integrated USB Type-C (USB 3.1 Gen 2, Thunderbolt™ 3, DisplayPort 1.4) - up to 4 ports



PCIe 3.0



Integrated WiFi 6 (Gig+)



Intel Optane™ Memory PCIe 3.0



OPI
intel
Intel® 300 Series Mobile Chipset

eSPI
SPI
LPC
SMBus
HD Audio



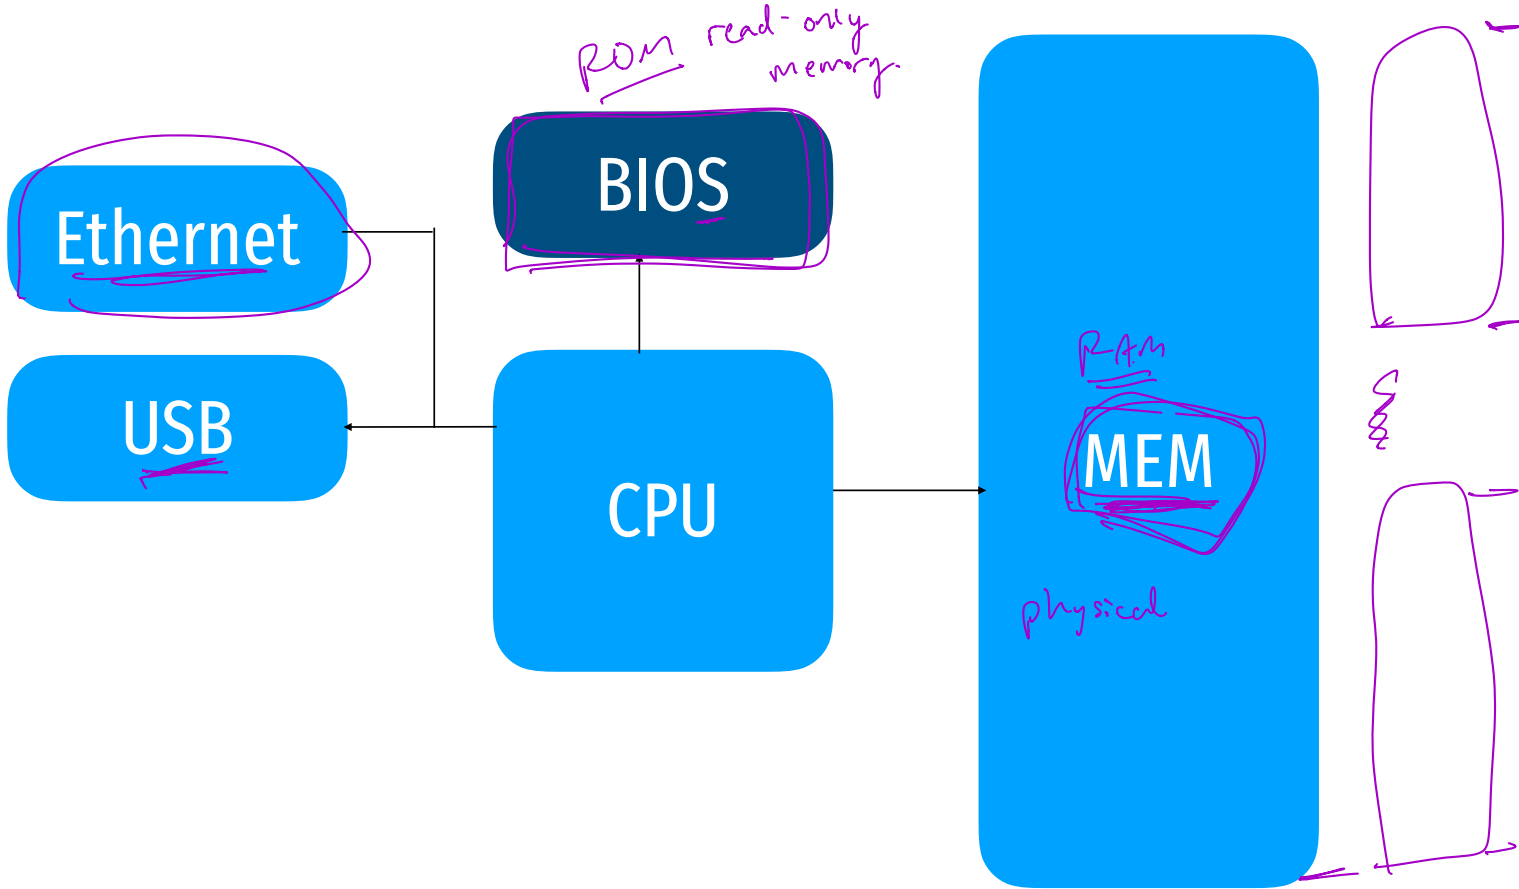
USB 2.0



SATA 3.0



Intel LAN PHY



What is Memory?

Memory is essentially a spreadsheet with a single column

- Every row has a number, called an address
- Every cell holds 1 byte of data

Address	Contents
114	
113	C
112	C
111	C
110	8
109	
108	U
107	U
106	L
105	.
104	
103
102	(
101	(
100	(

What is Memory?

Memory is essentially a spreadsheet with a single column

- Every row has a number, called an **address**
- Every cell holds 1 byte of data

~~32-bit~~

64-bit

Integers are typically four bytes

Address	Contents
114	
113	0
112	0
111	0
110	8
109	
108	
107	
106	
105	
104	
103	
102	
101	
100	

What is Memory?

Memory is essentially a spreadsheet with a single column

- Every row has a number, called an **address**
- Every cell holds 1 byte of data

Integers are typically four bytes

Each ASCII character is one byte, Strings are null terminated

Address	Contents
114	
113	0
112	0
111	0
110	8
109	
108	0
107	C
106	B
105	A
104	
103	
102	
101	
100	

What is Memory?

Memory is essentially a spreadsheet with a single column

- Every row has a number, called an **address**
- Every cell holds 1 byte of data

Integers are typically four bytes

Each ASCII character is one byte, Strings are null terminated

CPUs understand instructions in assembly language

Address	Contents
114	
113	0
112	0
111	0
110	8
109	
108	0
107	C
106	B
105	A
104	
103	0xAF
102	0x3C
101	0x91
100	0xE3

What is Memory?

Memory is essentially a spreadsheet with a single column

- Every row has a number, called an **address**
- Every cell holds 1 byte of data

All data and running code are held in memory

```
int my_num = 8;
```

Integers are typically four bytes

Each ASCII character is one byte, Strings are null terminated

CPUs understand instructions in assembly language

Address	Contents
114	
113	0
112	0
111	0
110	8
109	
108	0
107	C
106	B
105	A
104	
103	0xAF
102	0x3C
101	0x91
100	0xE3

What is Memory?

Memory is essentially a spreadsheet with a single column

- Every row has a number, called an **address**
- Every cell holds 1 byte of data

All data and running code are held in memory

```
int my_num = 8;
```

```
String my_str = "ABC";
```

Integers are typically four bytes

Each ASCII character is one byte, Strings are null terminated

CPUs understand instructions in assembly language

Address	Contents
114	
113	0
112	0
111	0
110	8
109	
108	0
107	C
106	B
105	A
104	
103	0xAF
102	0x3C
101	0x91
100	0xE3

What is Memory?

Memory is essentially a spreadsheet with a single column

- Every row has a number, called an **address**
- Every cell holds 1 byte of data

All data and running code are held in memory

```
int my_num = 8;
```

```
String my_str = "ABC";
```

```
while (my_num > 0) my_num--;
```

Integers are typically four bytes

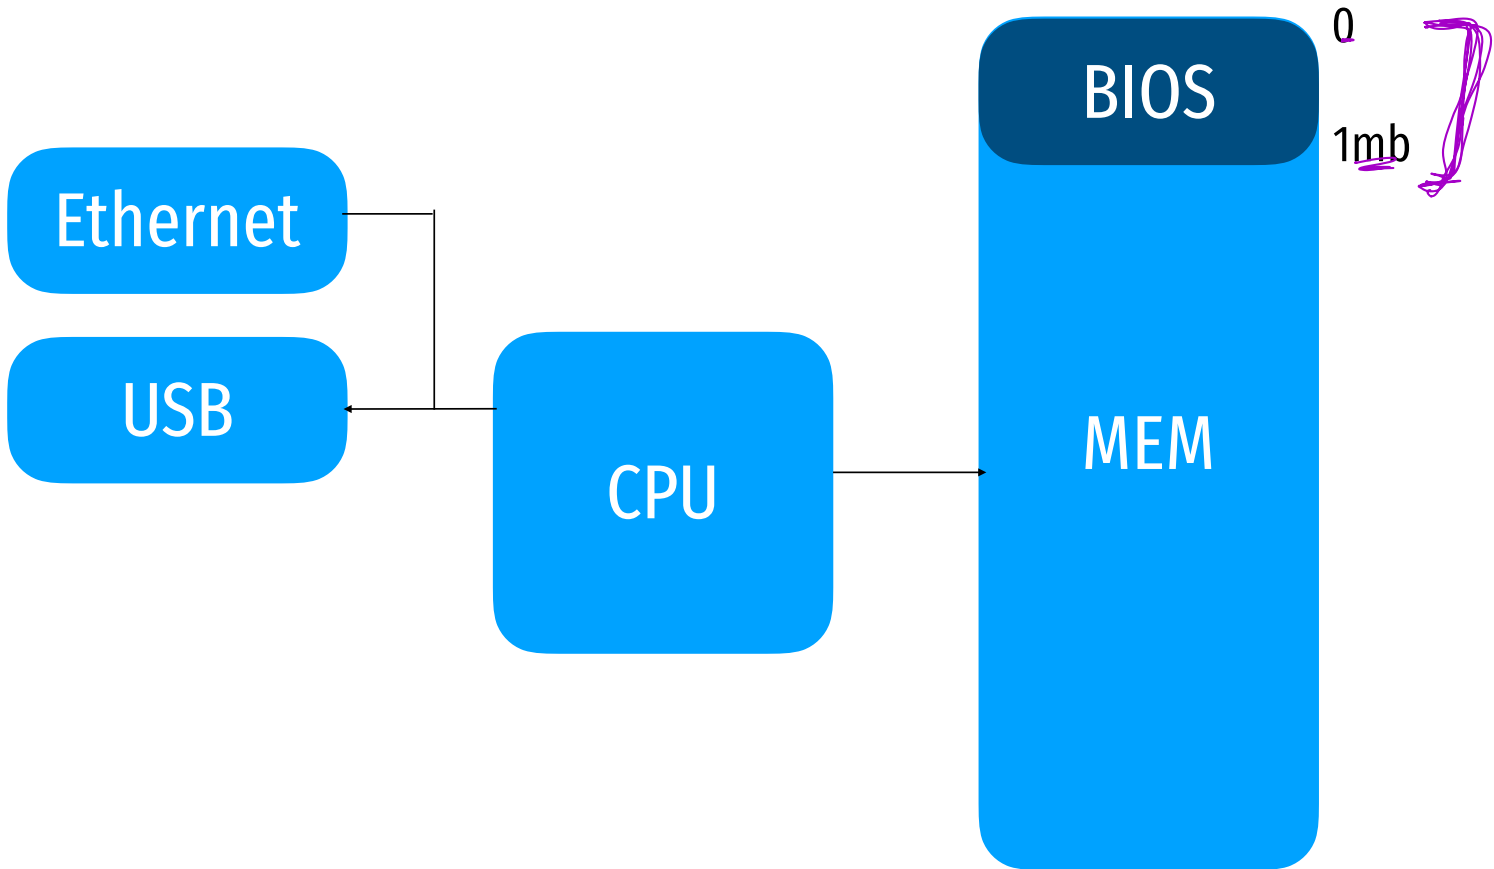
Each ASCII character is one byte, Strings are null terminated

CPUs understand instructions in assembly language

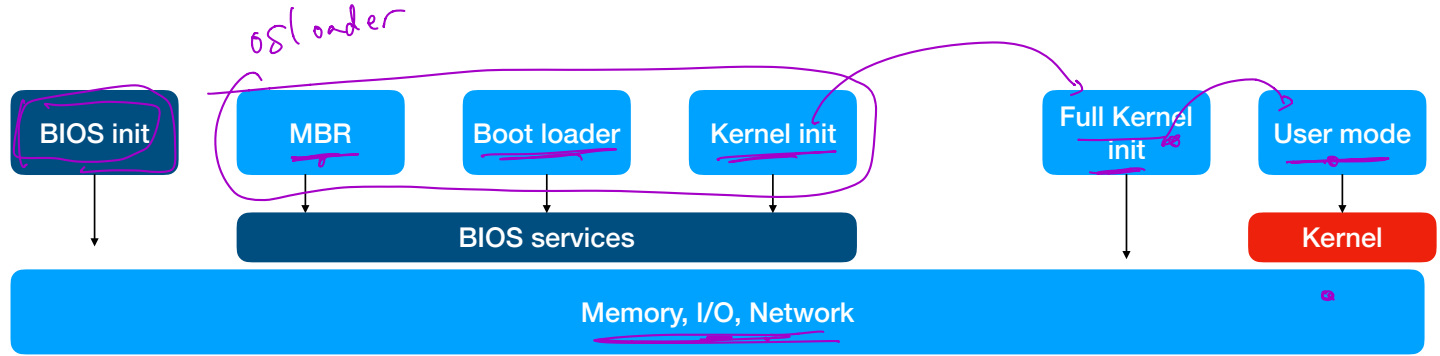
Address	Contents
114	
113	0
112	0
111	0
110	8
109	
108	0
107	C
106	B
105	A
104	
103	0xAF
102	0x3C
101	0x91
100	0xE3

How does a computer boot?

<https://youtu.be/MsKb0gR-4AM?t=36>



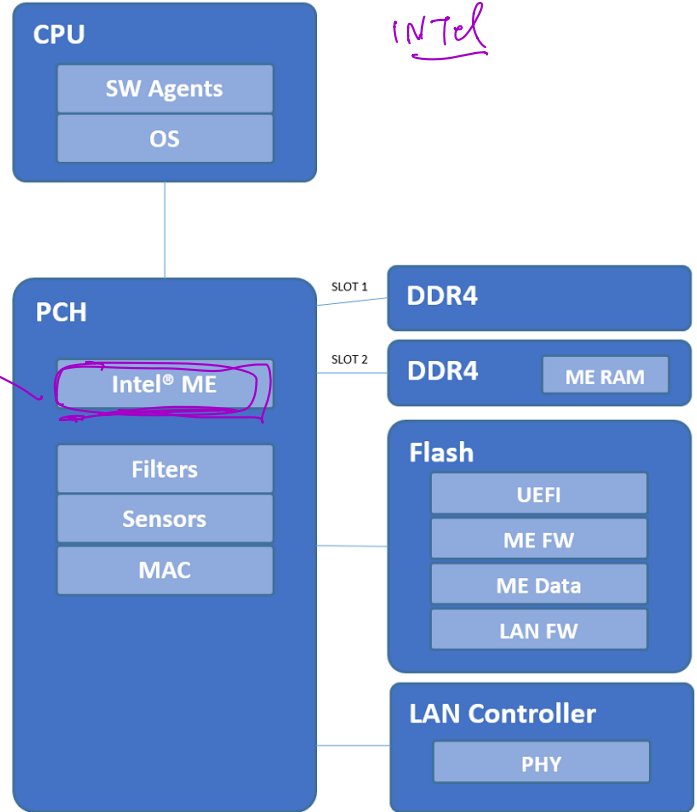
System Model: how does a computer boot?



More details

Intel
Management
Engine

① turns on the moment your
computer has power



Layout of memory at boot

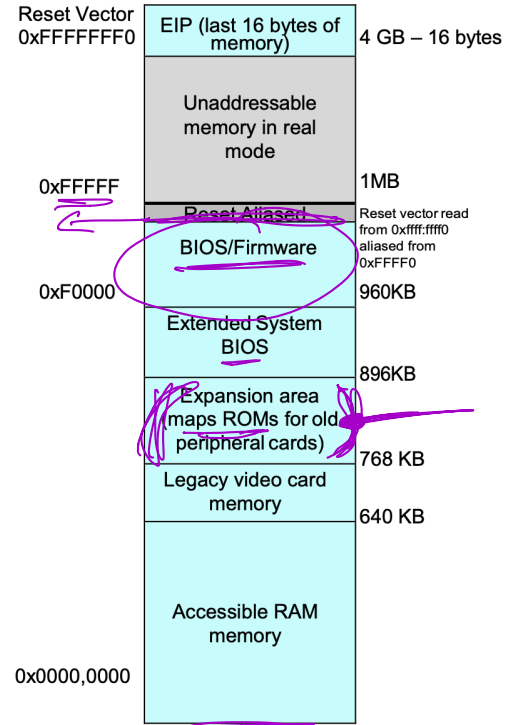
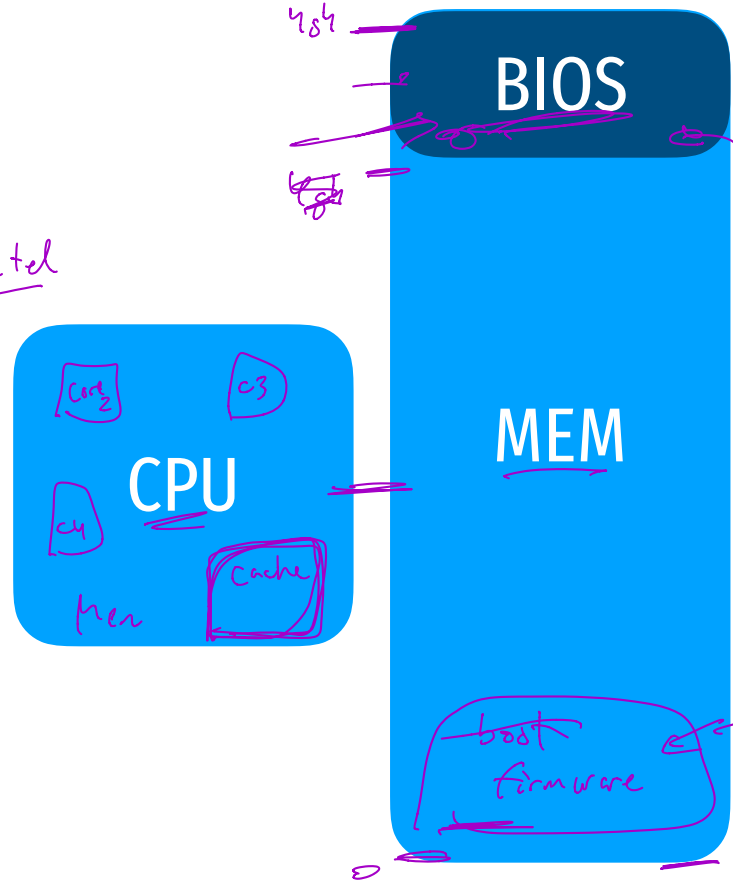



Figure 3 Intel® Architecture Memory Map at Power On

Details

- CPU begins executing at f.fff0
- BIOS firmware begins init of hw
- Applies microcode patches — FSPs } Intel
- Execute Firmware Support Pkg (blob)
- [Ram is setup]
- Copy firmware to RAM
- Begin executing in RAM
- Setup interrupts, timers, clocks, storage
- Bring up other cores
- Setup PCI
- Setup ACPI tables
- Execute OS loader



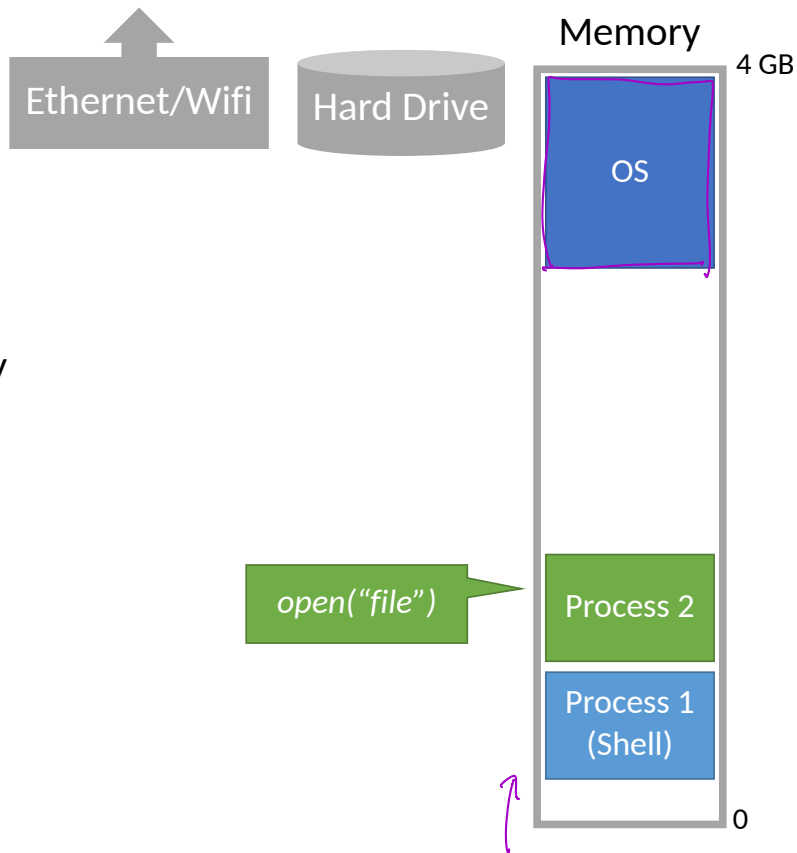
System Model



System Model

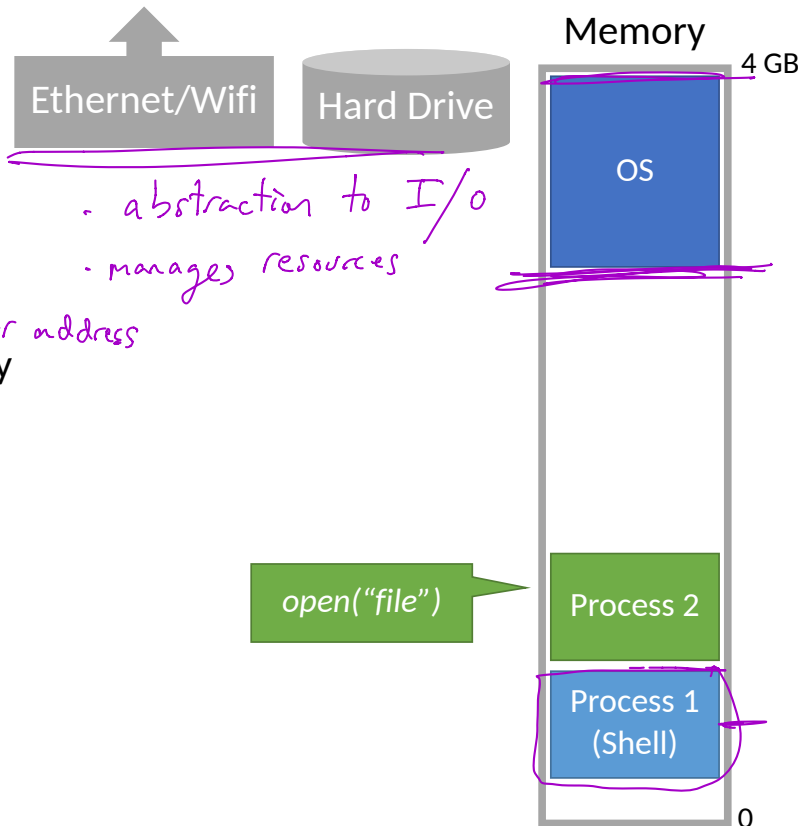
On bootup, the **Operating System (OS)** loads itself into memory

- eg. DOS (before hw isolation)
- Typically places itself in high memory



System Model

(old system like DOS)



On bootup, the **Operating System (OS)** loads itself into memory

- eg. DOS (before hw isolation)
- Typically places itself in high memory *higher address*

What is the role of the OS?

- Allow the user to run **processes**
- Often comes with a shell
 - Text shell like bash
 - Graphical shell like the Windows desktop
- Provides APIs to access devices
 - Offered as a convenience to application developers

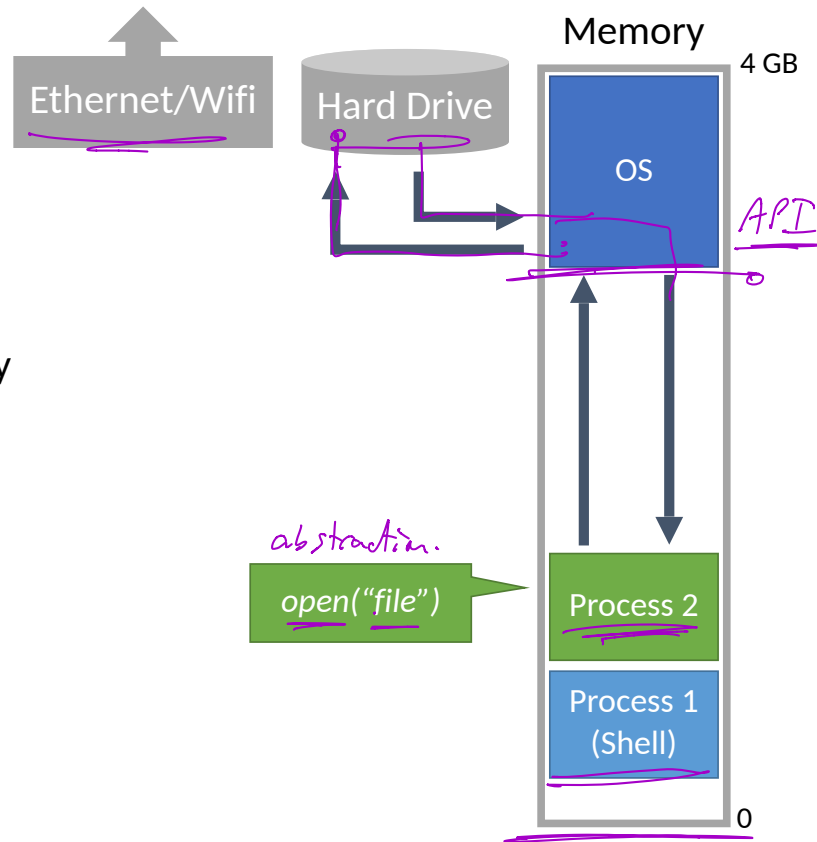
System Model

On bootup, the **Operating System (OS)** loads itself into memory

- eg. DOS (before hw isolation)
- Typically places itself in high memory

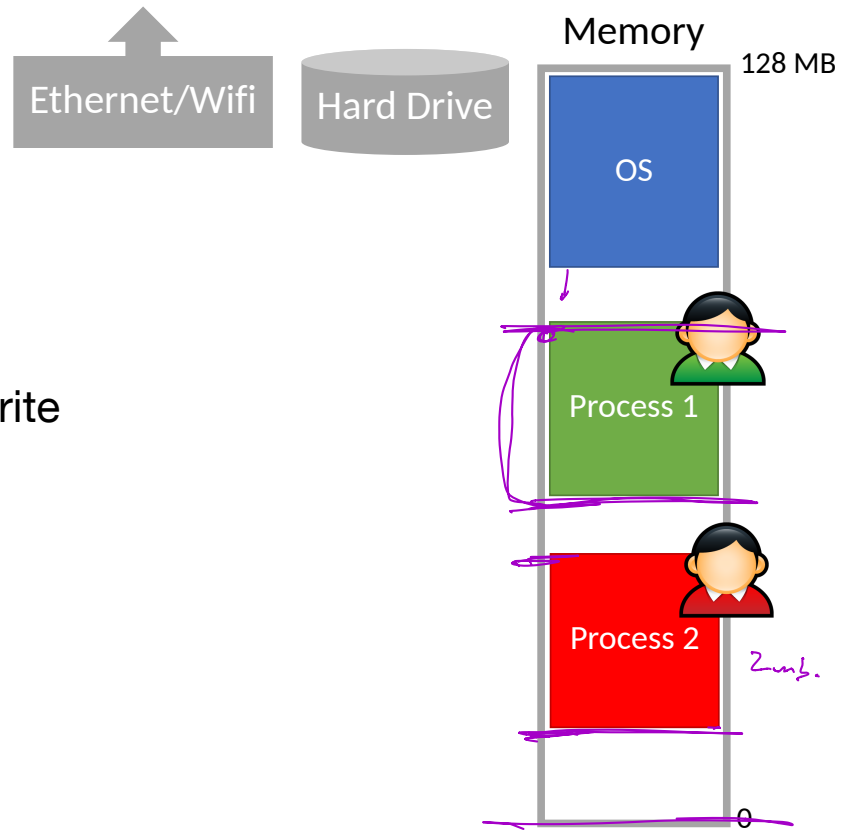
What is the role of the OS?

- Allow the user to run **processes**
- Often comes with a shell
 - Text shell like bash
 - Graphical shell like the Windows desktop
- Provides APIs to access devices
 - Offered as a convenience to application developers



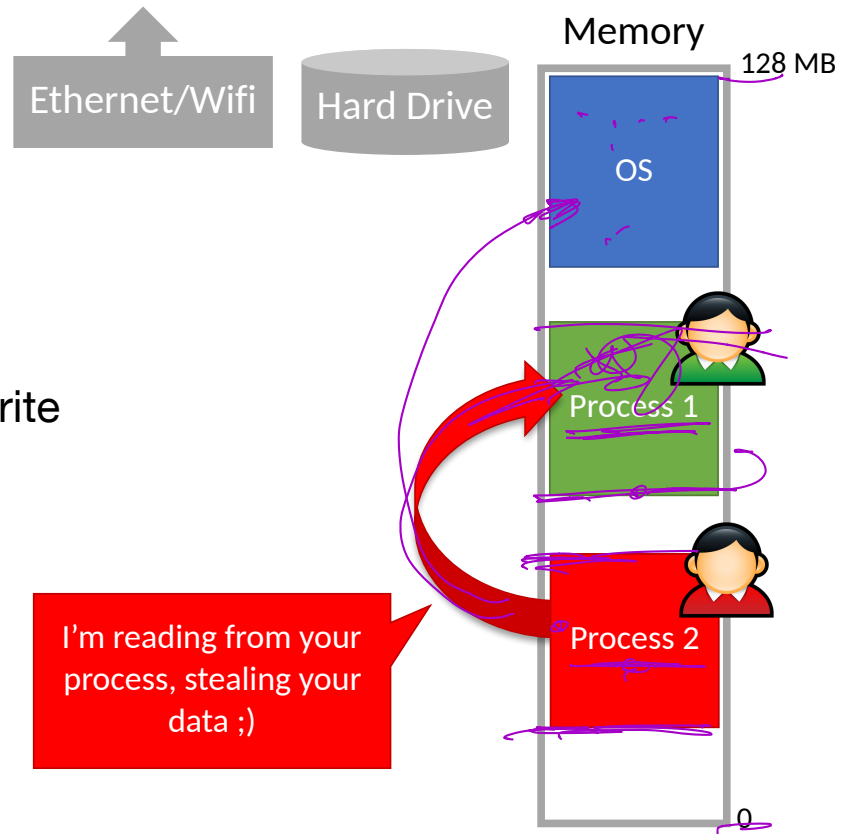
Memory Unsafety *DOS*

Problem: any process can read/write any memory




Memory Unsafety

Problem: any process can read/write any memory




Memory Unsafety

Problem: any process can read/write any memory





Memory Unsafety



Memory

128 MB



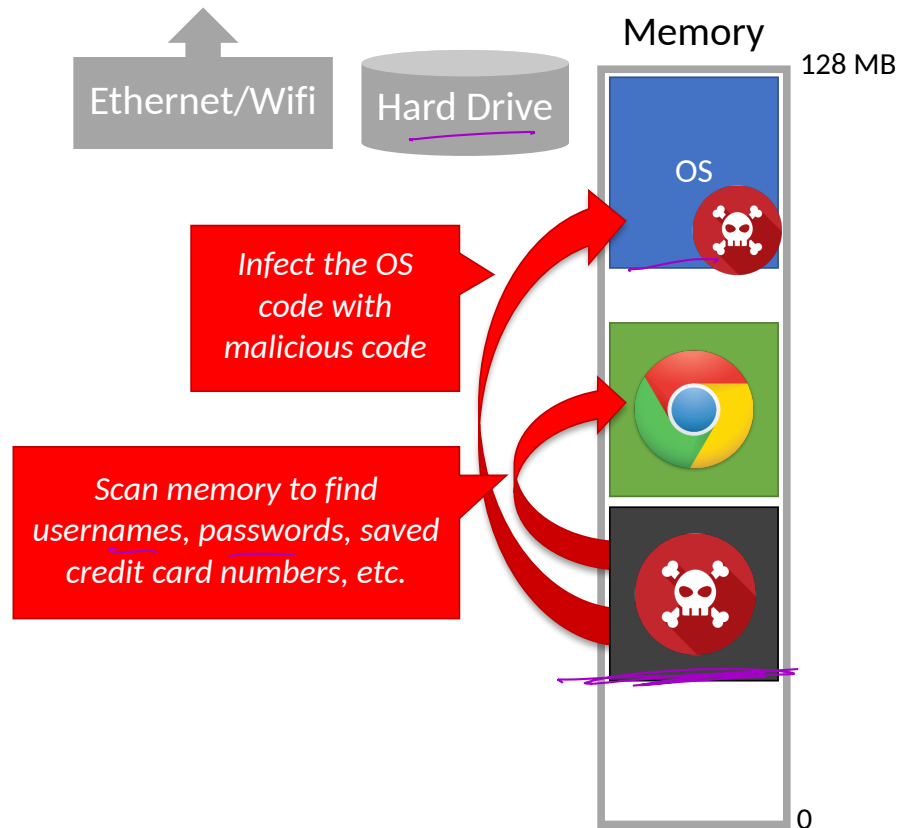
0

Problem: any process can read/write any memory

and thus no process can rely on a "safe memory semantics" i.e. "what I write before IS what I will read in the future"

Memory Unsafety

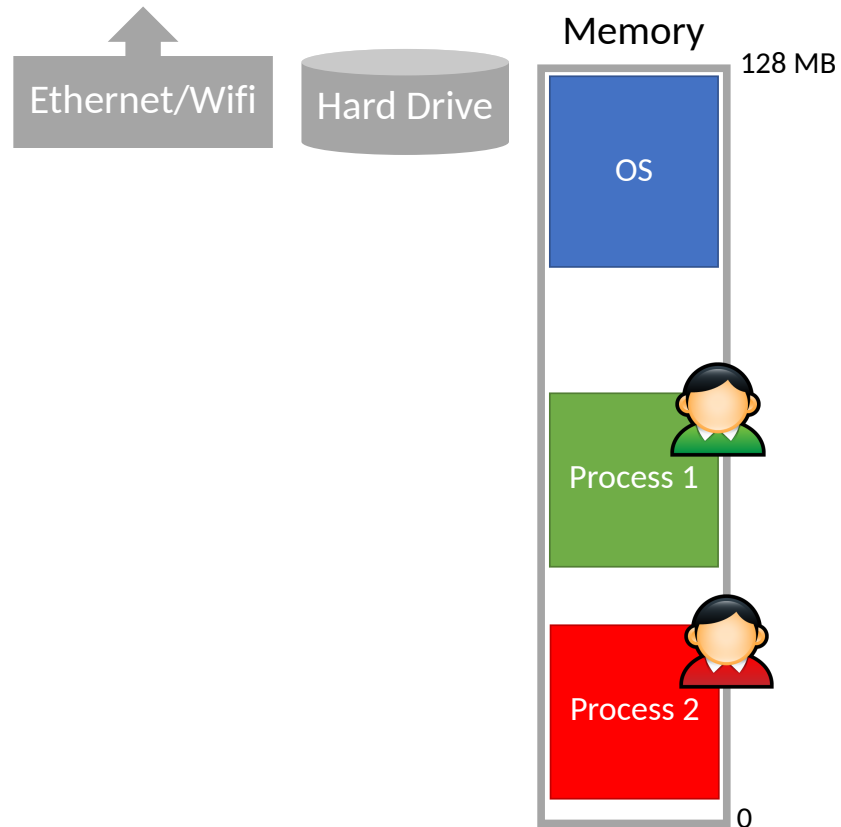
Problem: any process can read/write any memory



Device Unsafety

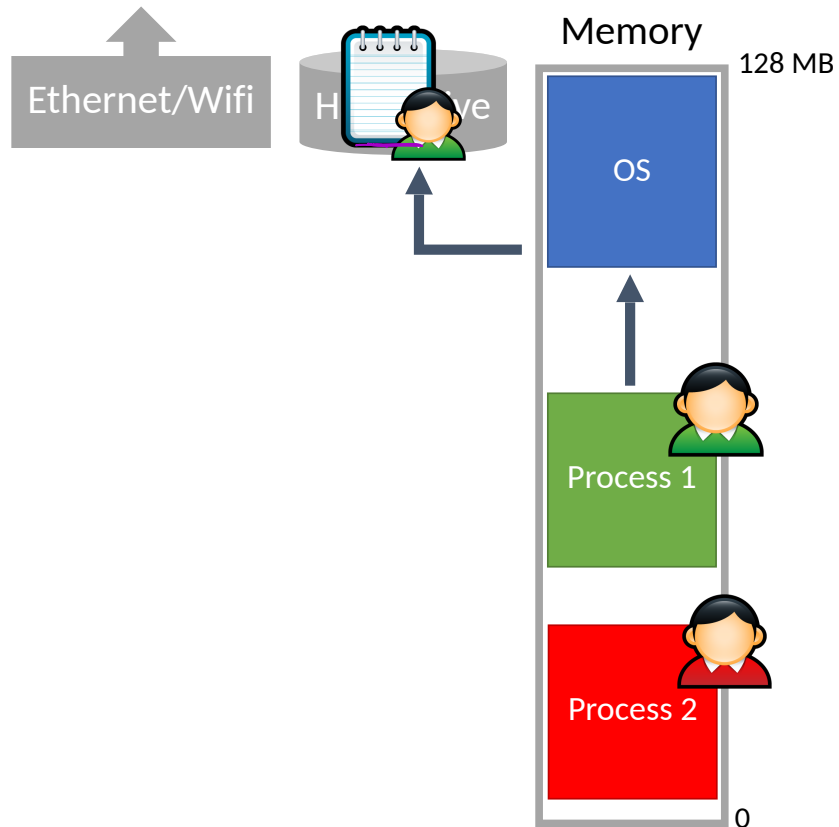
Problem: any process can access any hardware device directly

Access control is enforced by the OS, but OS APIs can be bypassed



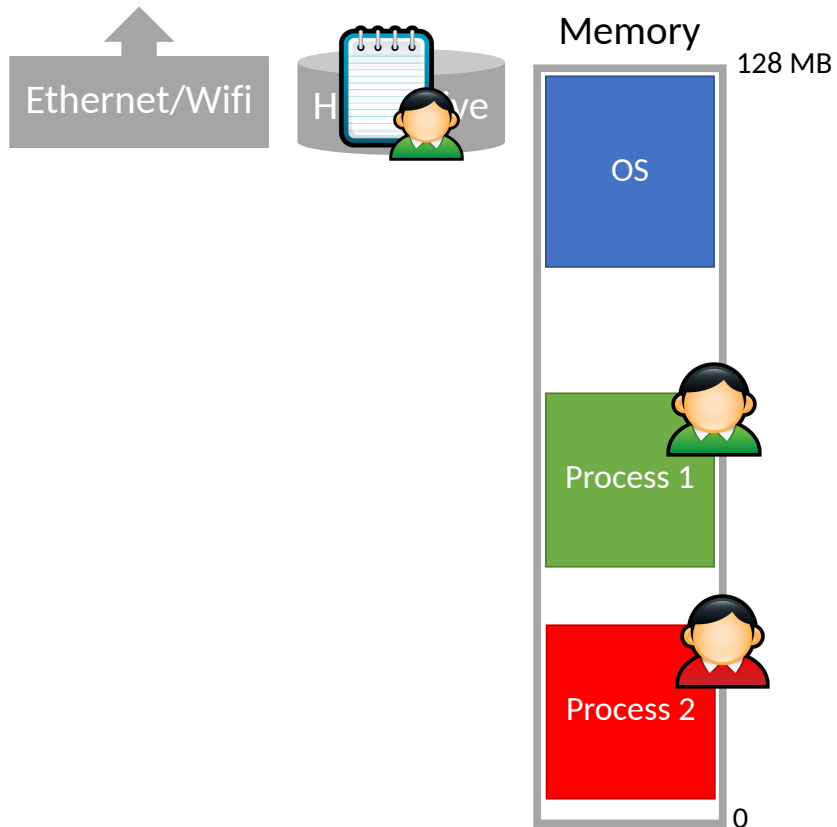
Device Unsafety

Problem: any process can access any hardware device directly
Access control is enforced by the OS, but OS APIs can be bypassed



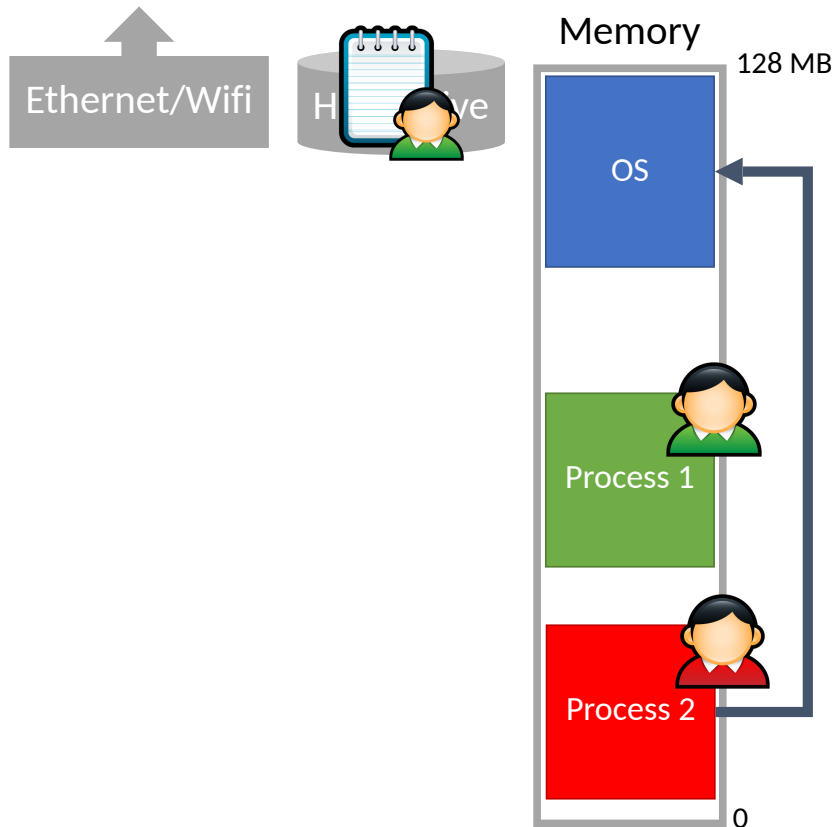
Device Unsafety

Problem: any process can access any hardware device directly
Access control is enforced by the OS, but OS APIs can be bypassed



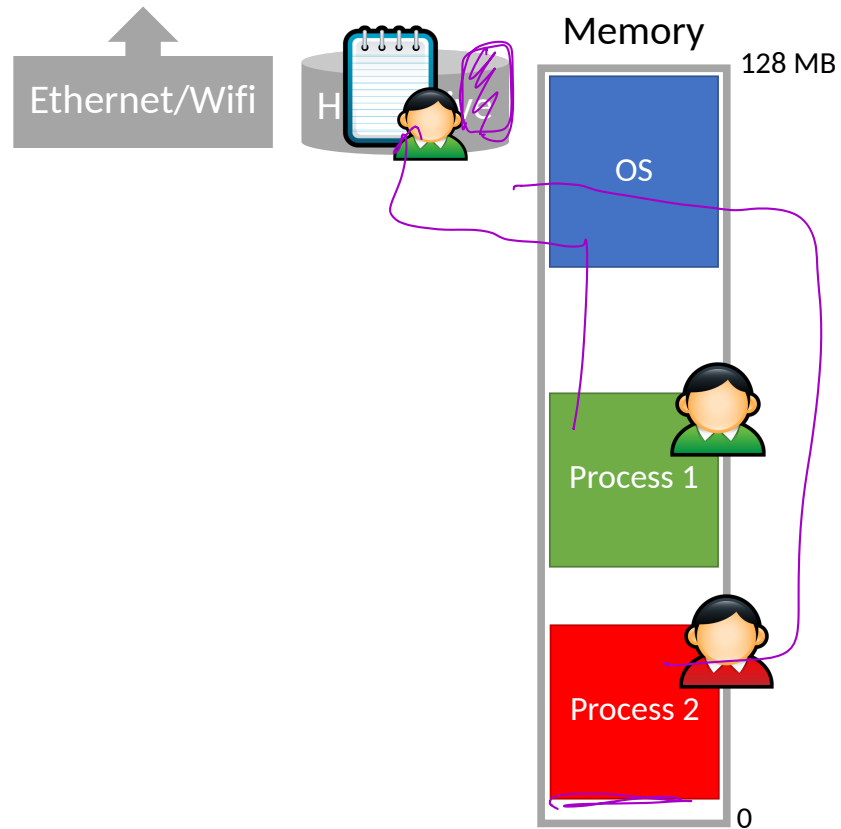
Device Unsafety

Problem: any process can access any hardware device directly
Access control is enforced by the OS, but OS APIs can be bypassed



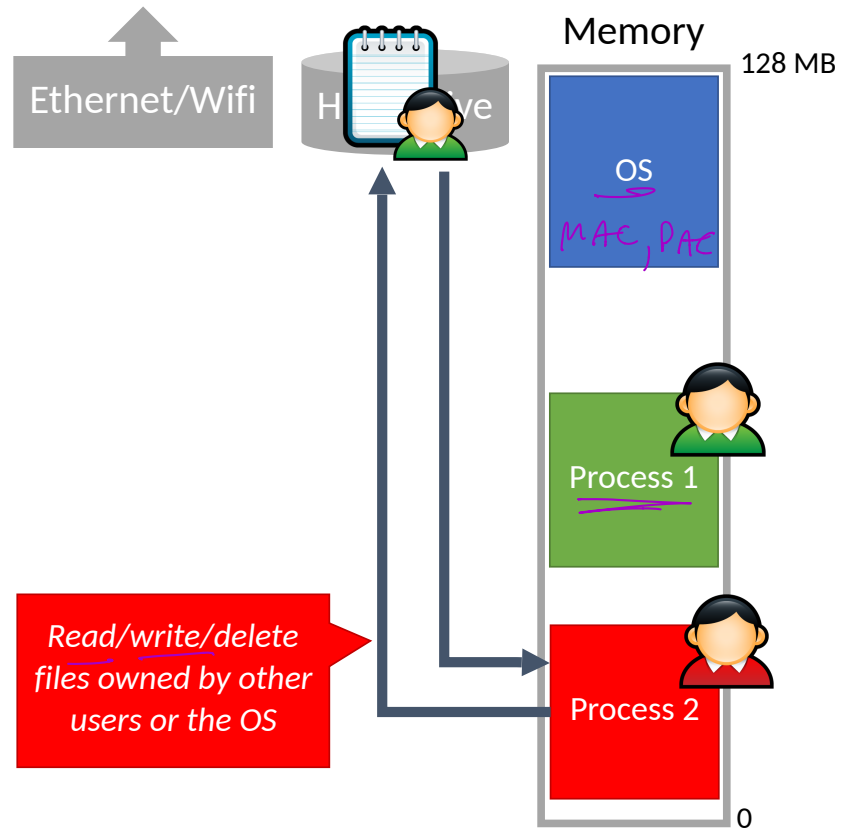
Device Unsafety

Problem: any process can access any hardware device directly
Access control is enforced by the OS, but OS APIs can be bypassed



Device Unsafety


Problem: any process can access any hardware device directly
Access control is enforced by the OS, but OS APIs can be bypassed



Device Unsafety

Problem: any process can access any hardware device directly


Access control is enforced by the OS, but OS APIs can be bypassed



Device Unsafety

Problem: any process can access any hardware device directly

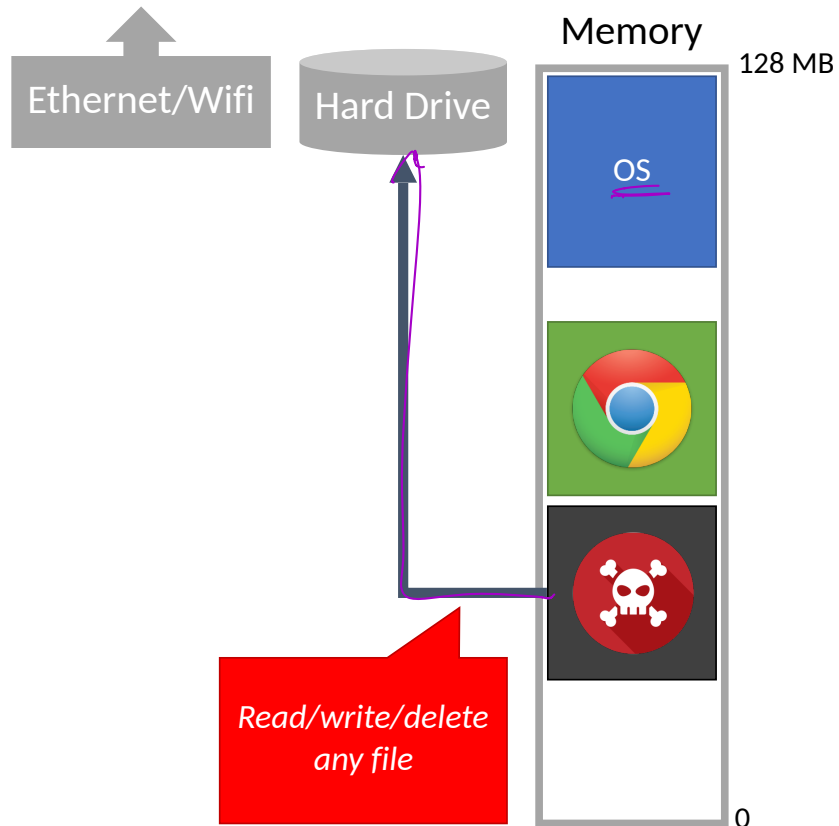
Access control is enforced by the OS, but OS APIs can be bypassed



Device Unsafety

Problem: any process can access any hardware device directly

Access control is enforced by the OS, but OS APIs can be bypassed



Device Unsafety

- Mono lithic memory/address space with no protections.

Problem: any process can access any hardware device directly

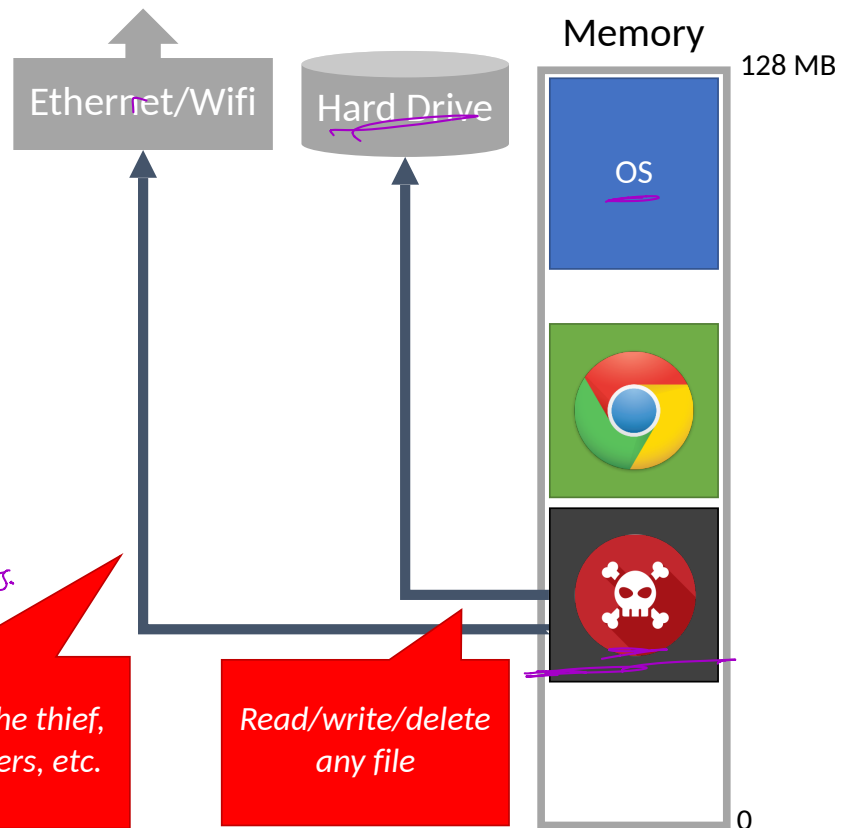
Access control is enforced by the OS, but OS APIs can be bypassed

- How OS worked for first 15 years.

Explain failures of this model

Send stolen data to the thief, attack other computers, etc.

Read/write/delete any file



- UNIX (1970s)
model

was way ahead

(needed some hw support)

Review

Old systems did not protect memory or devices

- Any process could access any memory
- Any process could access any device


Problems

- No way to enforce access controls on users or devices
- Processes can steal from or destroy each other
- Processes can modify or destroy the OS


On old computers, systems security was **literally impossible**

how??

ISOLATION



prevent
processes
from
interacting
"illegally"
with the
system
resources.



Threat Model

Principles

Intro to System Architecture

Hardware Support for Isolation

Examples

modern
hw
support
strong
isolation -

- Prngs
- virtual
memory
- virtualization
instructions.



Towards Modern Architecture

To achieve systems security, we need **process isolation**

- Processes cannot read/write memory arbitrarily
- Processes cannot access devices directly

How do we achieve this?

Hardware support for isolation

1. Protected mode execution (a.k.a. process rings)
2. Virtual memory




Protected Mode

Protected Mode

Most modern CPUs support **protected mode**

x86 CPUs support three rings with different privileges

- Ring 0: Operating System
 - Code in this ring may directly access any device




Protected Mode

Most modern CPUs support **protected mode**

x86 CPUs support three rings with different privileges

- Ring 0: Operating System
 - Code in this ring may directly access any device
- Ring 1, 2: device drivers
 - Code in these rings may directly access some devices
 - May not change the protection level of the CPU




Protected Mode

Most modern CPUs support **protected mode**

x86 CPUs support three rings with different privileges

- Ring 0: Operating System
 - Code in this ring may directly access any device
- Ring 1, 2: device drivers
 - Code in these rings may directly access some devices
 - May not change the protection level of the CPU
- Ring 3: userland
 - Code in this ring may not directly access devices
 - All device access must be via OS APIs
 - May not change the protection level of the CPU

- memory access constraints
- I/O constraints



Protected Mode

Most modern CPUs support **protected mode**


x86 CPUs support three rings with different privileges

- Ring 0: Operating System
 - Code in this ring may directly access any device
- Ring 1, 2: device drivers
 - Code in these rings may directly access some devices
 - May not change the protection level of the CPU
- Ring 3: userland
 - Code in this ring may not directly access devices
 - All device access must be via OS APIs
 - May not change the protection level of the CPU

Most OSes only use rings 0 and 3

Kernel OS

user processes



Ring -1, -2, -3

“Google cited worries that the Intel ME (actually MINIX) code runs on their CPU's deepest access level — Ring "-3" — and also runs a web server component that allows anyone to remotely connect to remote computers, even when the main OS is turned off.”

System Boot Sequence

1. On startup, the CPU starts in 16-bit **real** mode
 - Protected mode is disabled
 - Any process can access any device

System Boot Sequence

1. On startup, the CPU starts in 16-bit **real** mode
 - Protected mode is disabled
 - Any process can access any device
2. BIOS executes, finds and loads the OS

System Boot Sequence

1. On startup, the CPU starts in 16-bit **real** mode
 - Protected mode is disabled
 - Any process can access any device
2. BIOS executes, finds and loads the OS
3. OS switches CPU to 32-bit **protected** mode
 - OS code is now running in Ring 0
 - OS decides what Ring to place other processes in

System Boot Sequence

1. On startup, the CPU starts in 16-bit real mode
 - Protected mode is disabled
 - Any process can access any device
2. BIOS executes, finds and loads the OS
3. OS switches CPU to 32-bit protected mode
 - OS code is now running in Ring 0
 - OS decides what Ring to place other processes in
4. Shell gets executed, user may run programs
 - User processes are placed in Ring 3

Restriction on Privileged Instructions

{ What CPU instructions are restricted in protected mode? }

- Any instruction that modifies the CR0 register —
 - Controls whether protected mode is enabled
- Any instruction that modifies the CR3 register —
 - Controls the virtual memory configuration
 - More on this later...
- hlt - Halts the CPU
- sti/cli - enable and disable interrupts
- in/out - directly access hardware devices

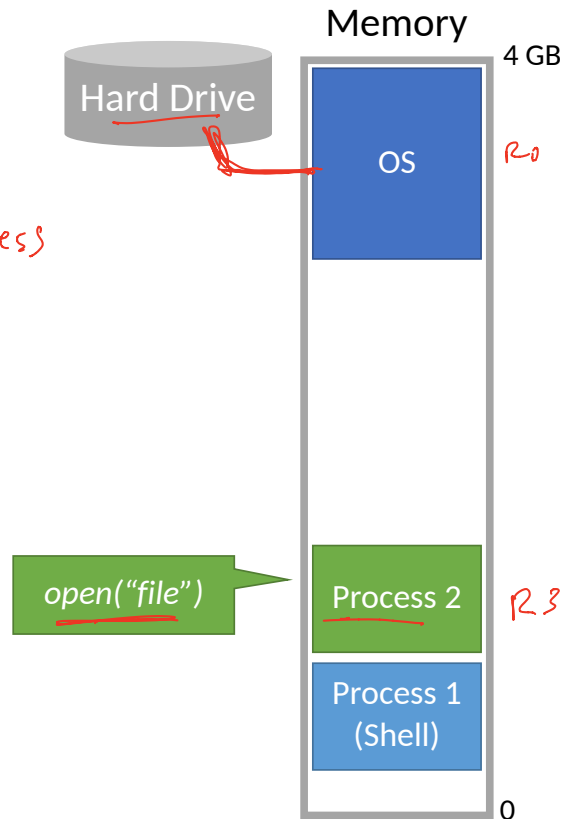
control registers of the CPU
indicate, ring level.

If a Ring 3 process tries any of these things, it immediately crashes

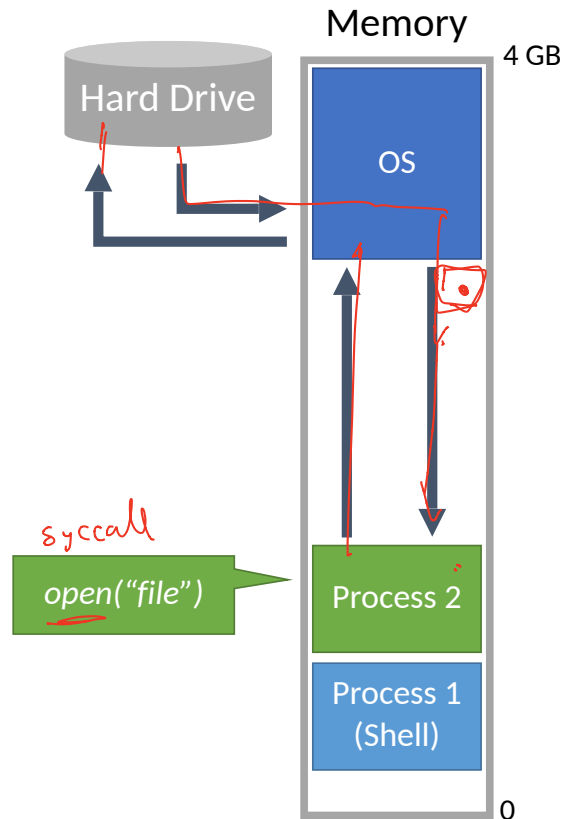
How to change modes

"system call"

- method for a user-land process running at R3 to communicate with the OS running in R0.



How to change modes



Changing Modes

Applications often need to access the OS APIs

- Writing files
- Displaying things on the screen
- Receiving data from the network
- etc...

But the OS is Ring 0, and processes are Ring 3

How do processes get access to the OS?

Changing Modes

Applications often need to access the OS APIs

- Writing files
- Displaying things on the screen
- Receiving data from the network
- etc...

But the OS is Ring 0, and processes are Ring 3

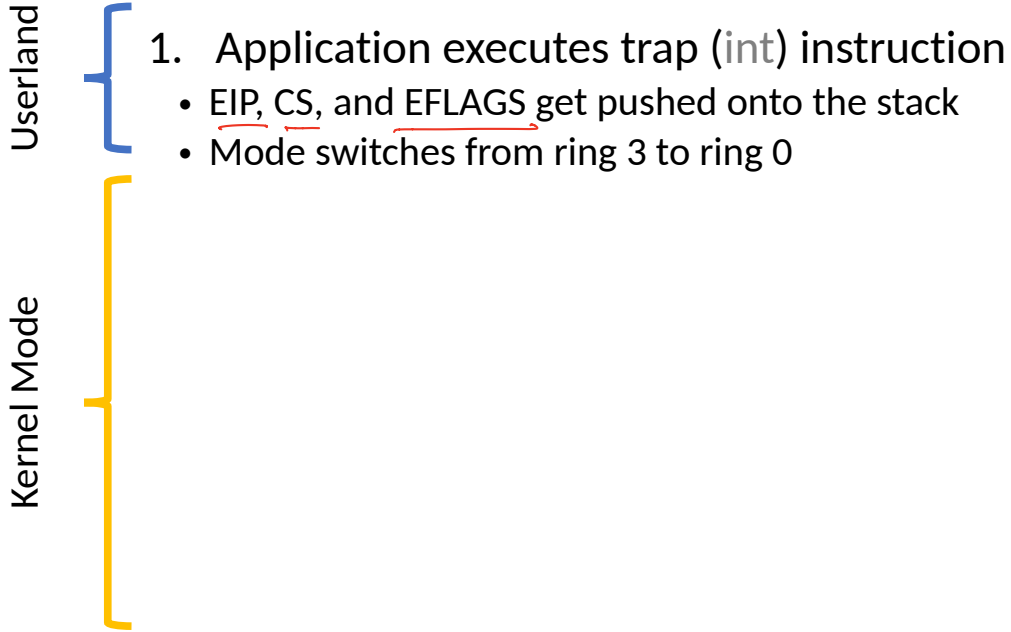
How do processes get access to the OS?

- Invoke OS APIs with special assembly instructions
 - Interrupt: `int 0x80`
 - System call: `sysenter` or `syscall`
- `int/sysenter/syscall` cause a mode transfer from Ring 3 to Ring 0

properly defined
interrupt handler
the CPU
setup in

standard method
- load arguments to the
syscall into registers
- execute `int 80`

Mode Transfer



Mode Transfer

-
- The diagram illustrates the process of mode transfer from Userland to Kernel Mode. It features two vertical labels on the left: 'Userland' and 'Kernel Mode'. A blue bracket on the left side groups the first step, which occurs in Userland. A yellow bracket on the left side groups the second step, which occurs in Kernel Mode. The steps are as follows:
- Userland**
 - 1. Application executes trap (`int`) instruction
 - EIP, CS, and EFLAGS get pushed onto the stack
 - Mode switches from ring 3 to ring 0
 - Kernel Mode**
 - 2. Save the state of the current process
 - Push EAX, EBX, ..., etc. onto the stack

Mode Transfer

Userland

1. Application executes trap (int) instruction
 - EIP, CS, and EFLAGS get pushed onto the stack
 - Mode switches from ring 3 to ring 0

Kernel Mode

2. Save the state of the current process
 - Push EAX, EBX, ..., etc. onto the stack
3. Locate and execute the correct syscall handler

Mode Transfer

Userland

1. Application executes trap (int) instruction
 - EIP, CS, and EFLAGS get pushed onto the stack
 - Mode switches from ring 3 to ring 0

Kernel Mode

2. Save the state of the current process
 - Push EAX, EBX, ..., etc. onto the stack
3. Locate and execute the correct syscall handler
4. Restore the state of process
 - Pop EAX, EBX, ... etc.

Mode Transfer

Userland

1. Application executes trap (int) instruction
 - EIP, CS, and EFLAGS get pushed onto the stack
 - Mode switches from ring 3 to ring 0

Kernel Mode

2. Save the state of the current process
 - Push EAX, EBX, ..., etc. onto the stack
3. Locate and execute the correct syscall handler
4. Restore the state of process
 - Pop EAX, EBX, ... etc.
5. Place the return value in EAX

Convention

- ensures

access controls

can be applied by

kernel

handler

⇒ safety

Mode Transfer

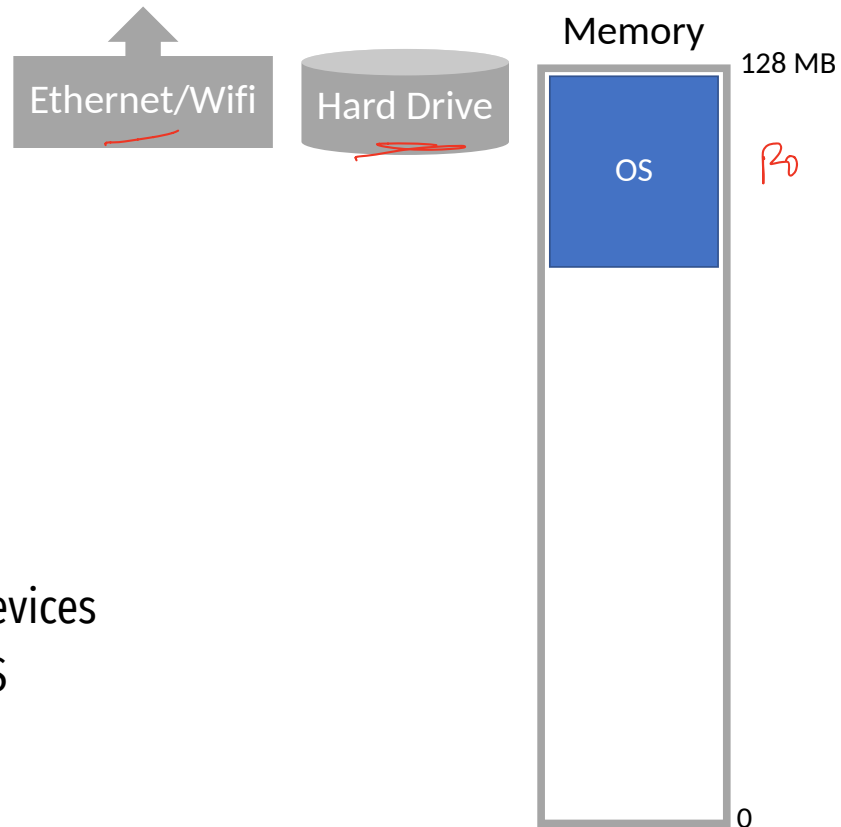
Userland

1. Application executes trap (`int`) instruction
 - EIP, CS, and EFLAGS get pushed onto the stack
 - Mode switches from ring 3 to ring 0

Kernel Mode

2. Save the state of the current process
 - Push EAX, EBX, ..., etc. onto the stack
3. Locate and execute the correct syscall handler
4. Restore the state of process
 - Pop EAX, EBX, ... etc.
5. Place the return value in EAX
6. Use `iret` to return to the process
 - Switches back to the original mode (typically 3)

Protection in Action




Protected mode stops direct access to devices

All device access must go through the OS

OS will impose access control checks

Protection in Action




Protected mode stops direct access to devices

All device access must go through the OS

OS will impose access control checks


Protection in Action



Protected mode stops direct access to devices
All device access must go through the OS
OS will impose access control checks


Protection in Action

Protected mode stops direct access to devices
All device access must go through the OS
OS will impose access control checks



Protection in Action

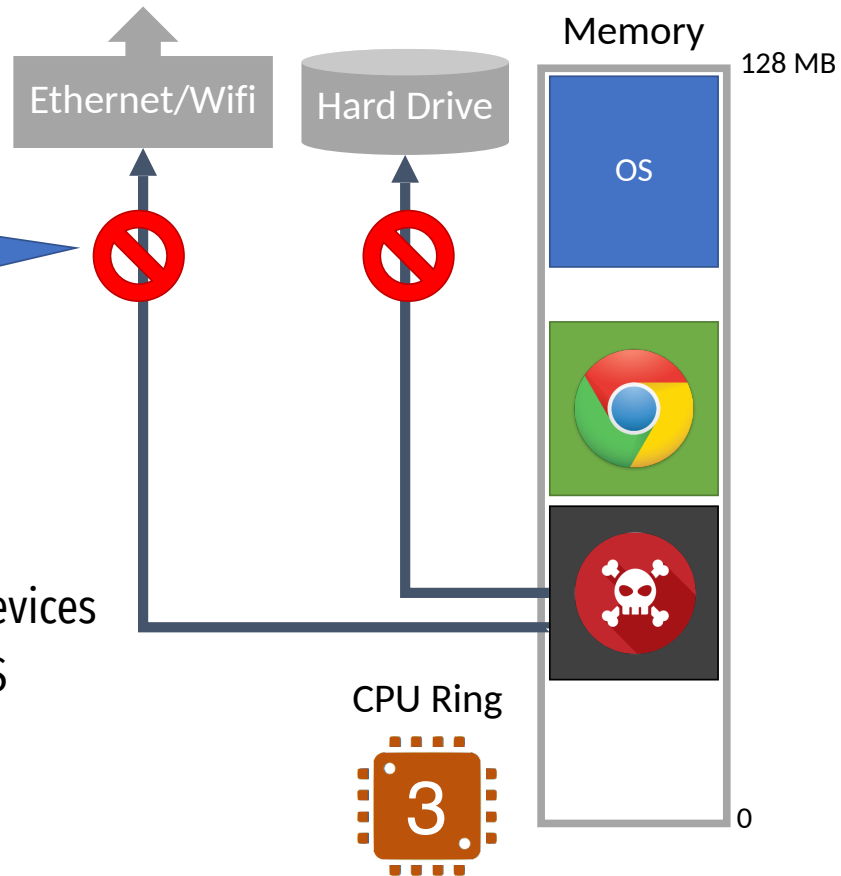
Protected mode stops direct access to devices
All device access must go through the OS
OS will impose access control checks



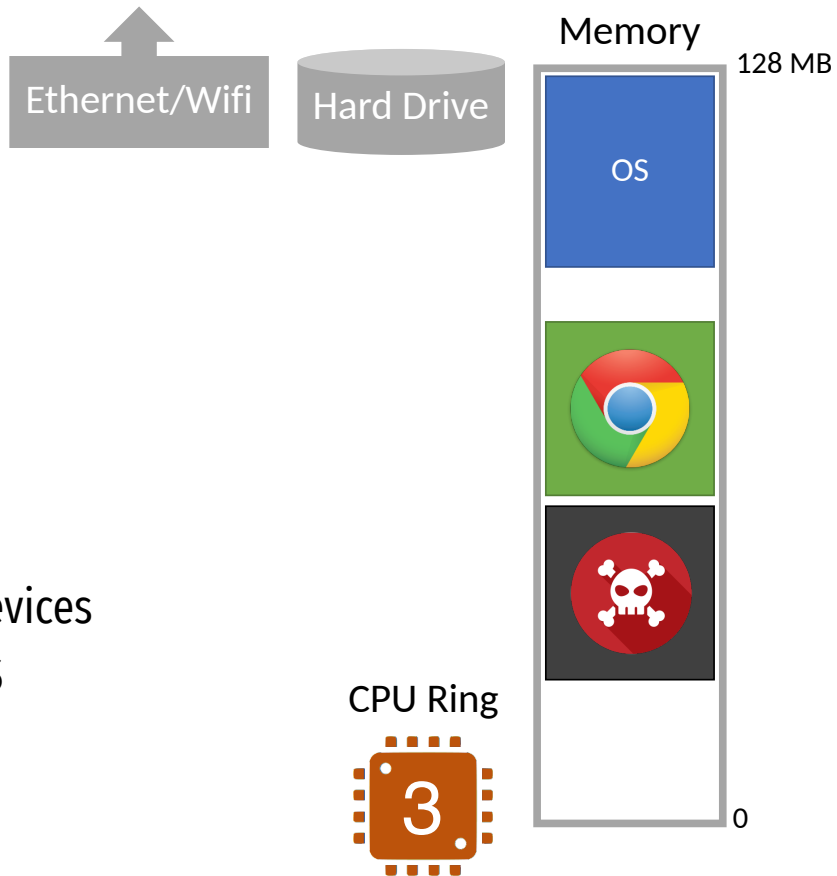
Protection in Action

Ring 3 = protected mode.
No direct device access

Protected mode stops direct access to devices
All device access must go through the OS
OS will impose access control checks



Protection in Action



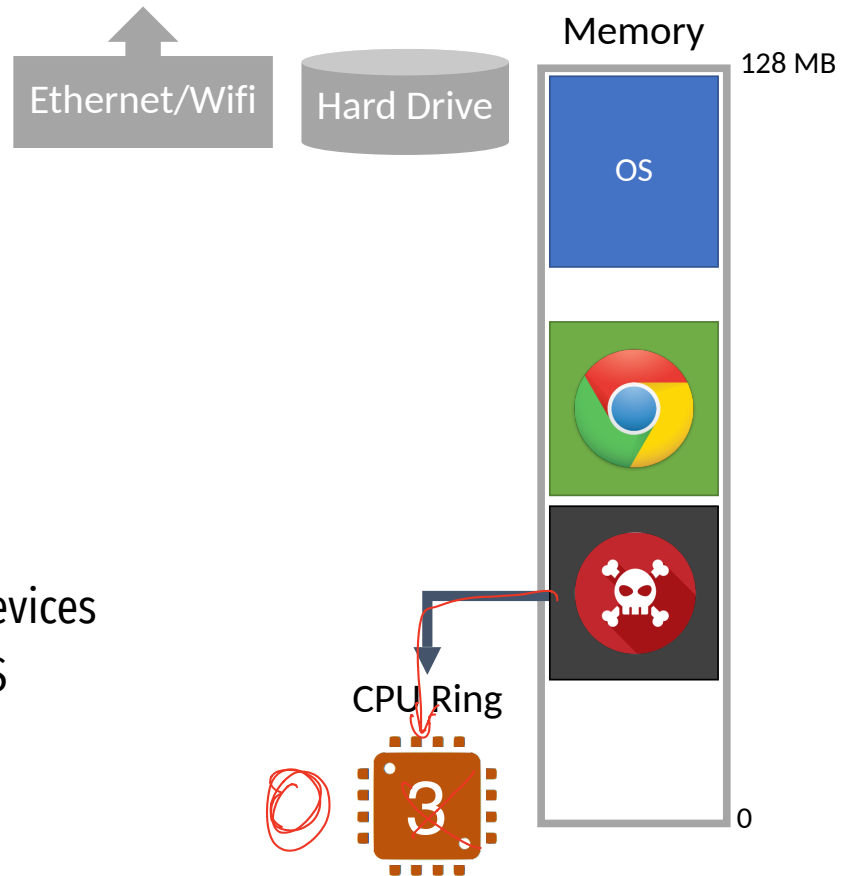
Protected mode stops direct access to devices

All device access must go through the OS

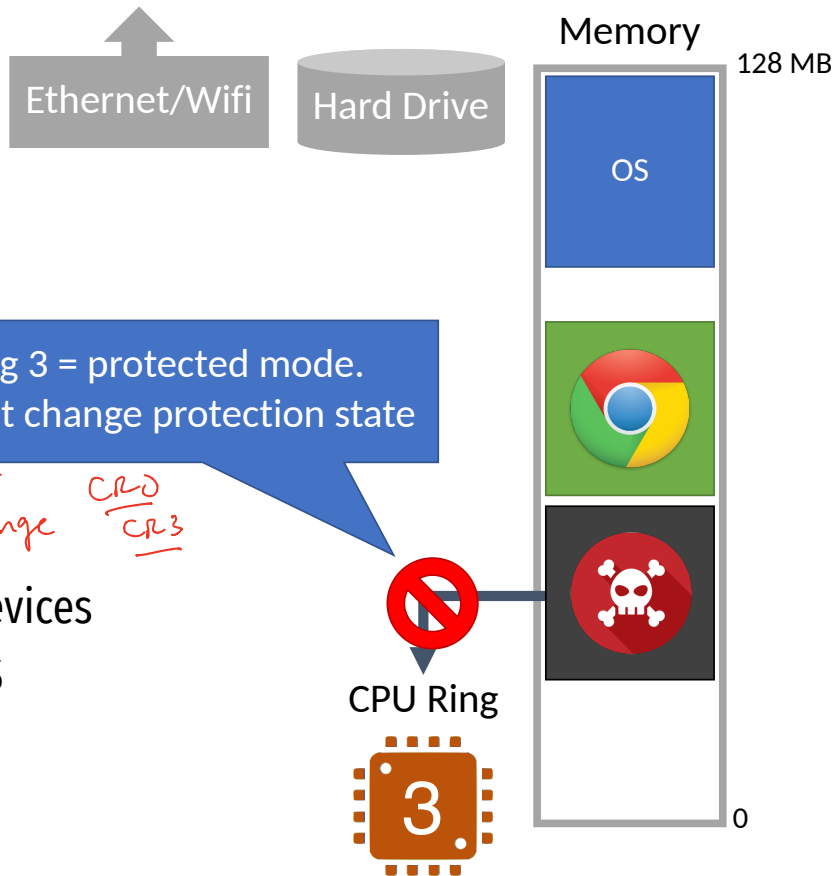
OS will impose access control checks

Protection in Action

Protected mode stops direct access to devices
All device access must go through the OS
OS will impose access control checks

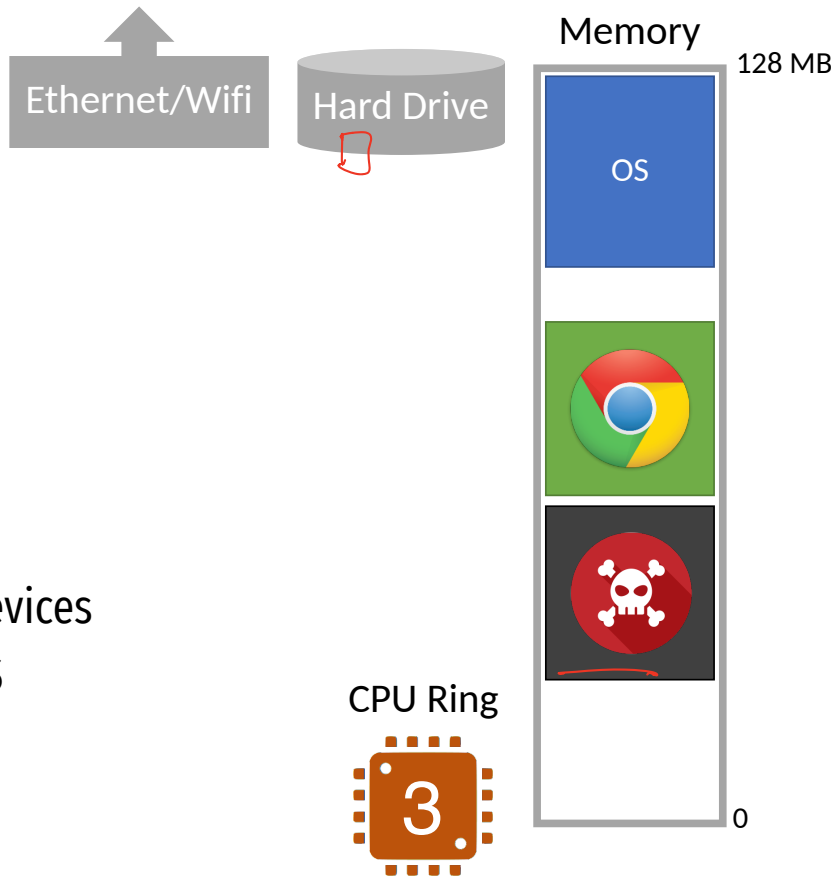


Protection in Action



Protected mode stops direct access to devices
All device access must go through the OS
OS will impose access control checks


Protection in Action



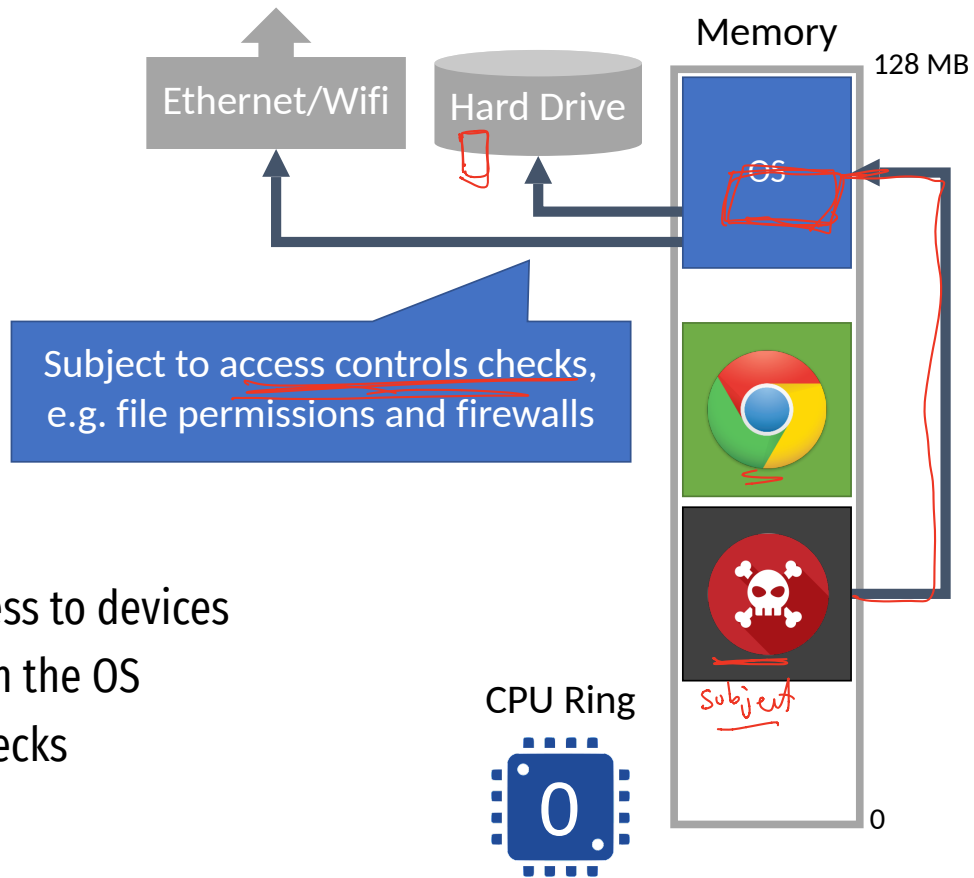
Protected mode stops direct access to devices
All device access must go through the OS
OS will impose access control checks

Protection in Action

Protected mode stops direct access to devices
All device access must go through the OS
OS will impose access control checks



Protection in Action



Protected mode stops direct access to devices

All device access must go through the OS

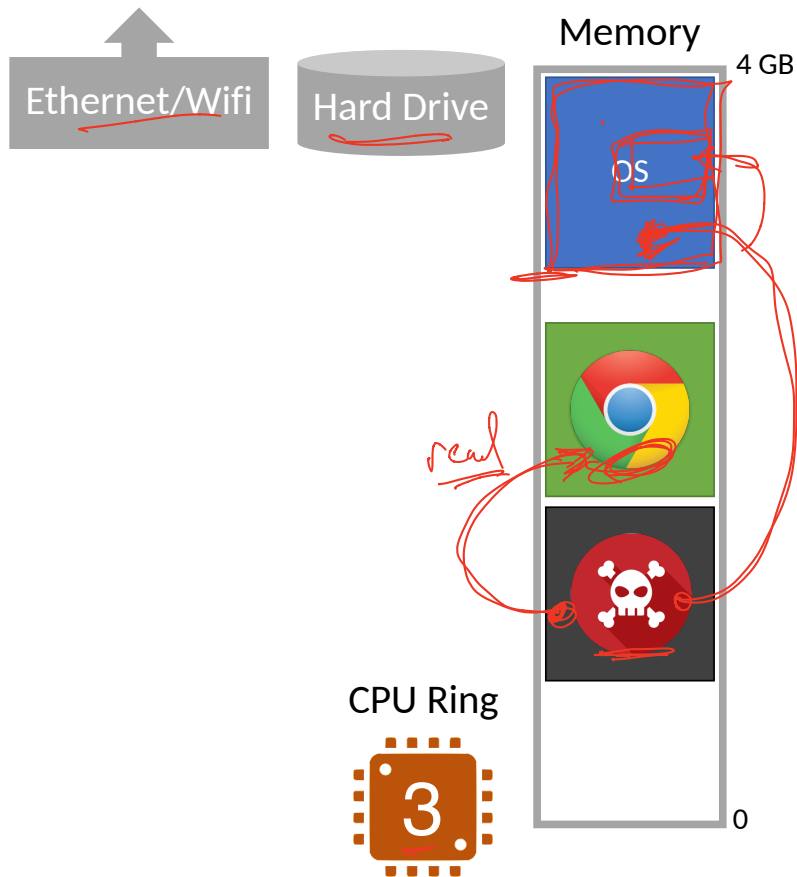
OS will impose access control checks

Virtual Memory

Status Check

At this point we have protected the devices attached to the system...

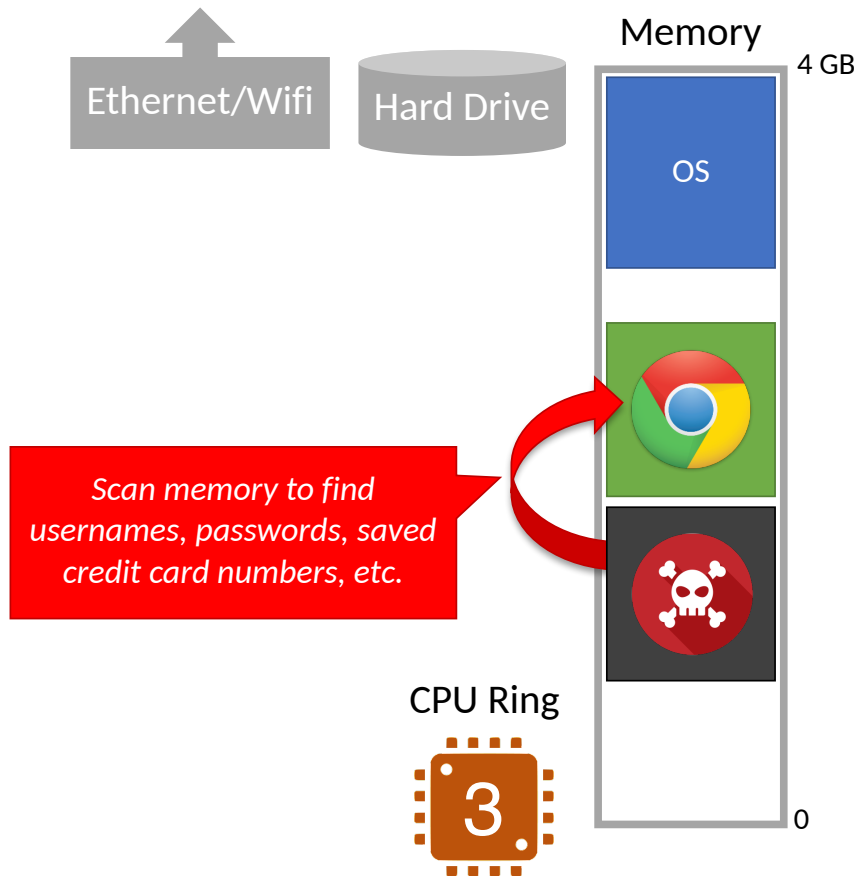
... But we have not protected memory



Status Check

At this point we have protected the devices attached to the system...


... But we have not protected memory



Status Check

At this point we have protected the devices attached to the system...

... But we have not protected memory



Memory Isolation and Virtual Memory

Modern CPUs support **virtual memory**

Creates the illusion that each process runs in its own, empty memory space

- Processes can not read/write memory used by other processes
- Processes can not read/write memory used by the OS

Memory Isolation and Virtual Memory

Modern CPUs support **virtual memory**

Creates the illusion that each process runs in its own, empty memory space


- Processes can not read/write memory used by other processes
- Processes can not read/write memory used by the OS

In later courses, you will learn how virtual memory is implemented

- Base and bound registers
- Segmentation
- Page tables



Today, we will do the cliffnotes version...

Physical Memory



Physical Memory

4 GB



0

Physical Memory

4 GB



Virtual Memory Process 1

4 GB

Chrome believes it is the only thing in memory



Physical Memory

4 GB



Virtual Memory Process 1

4 GB



Chrome believes it is the only thing in memory




Physical Memory




Virtual Memory Process 1

Chrome believes it is the only thing in memory



Virtual Memory Process 2


Skype believes it is the only thing in memory



37

Virtual Memory

Process 1



Physical Memory


Memory



Virtual Memory Process 1

4 GB

Read Address
16734



Physical Memory

4 GB

Physical Address:
81102



Virtual Memory Process 1

4 GB




Read
Address
16734

Page Table

Virtual Addr.	Physical Addr.
16732	81100
16734	81102
16736	93568
16738	93570

CPU



Physical Memory

4 GB



Physical
Address:
81102

Virtual Memory Process 1

4 GB




Read
Address
16734

Page Table

Virtual Addr.	Physical Addr.
16732	81100
16734	81102
16736	93568
16738	93570

CPU



Physical Memory

4 GB



Physical
Address:
81102

Virtual Memory Implementation

Each process has its own virtual memory space

- Each process has a page table that maps its virtual space into physical space
- CPU translates virtual address to physical addresses on-the-fly

Virtual Memory Implementation

Each process has its own virtual memory space

- Each process has a page table that maps its virtual space into physical space
- CPU translates virtual address to physical addresses on-the-fly

OS creates the page table for each process

- Installing page tables in the CPU is a protected, Ring 0 instruction
- Processes cannot modify their page tables

Virtual Memory Implementation

Each process has its own virtual memory space

- Each process has a page table that maps its virtual space into physical space
- CPU translates virtual address to physical addresses on-the-fly

OS creates the page table for each process

- Installing page tables in the CPU is a protected, Ring 0 instruction
- Processes cannot modify their page tables

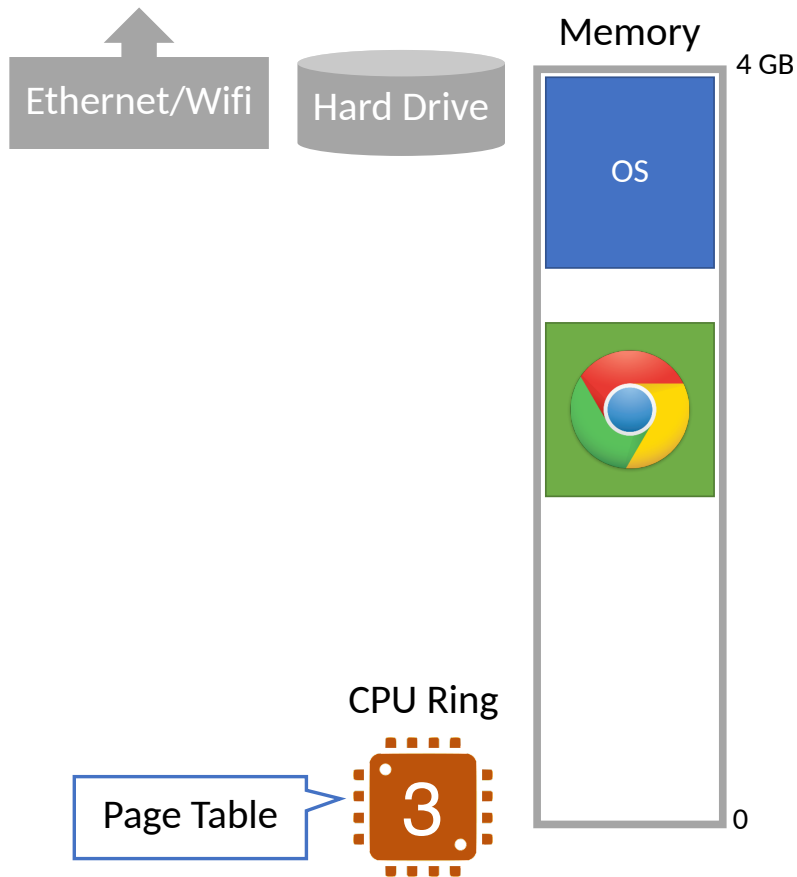
What happens if a process tries to read/write memory outside its page table?

- **Segmentation Fault** or **Page Fault**
- Process crashes
- In other words, no way to escape virtual memory

VM in Action

Processes can only read/
write within their own
virtual memory

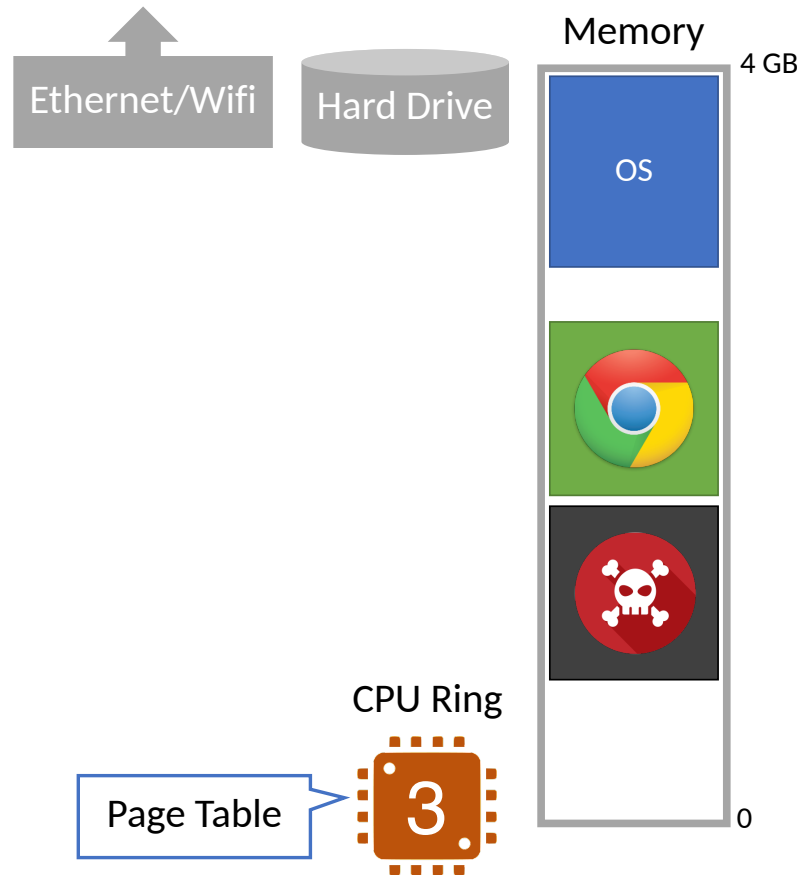
Processes cannot change
their own page tables



VM in Action

Processes can only read/
write within their own
virtual memory

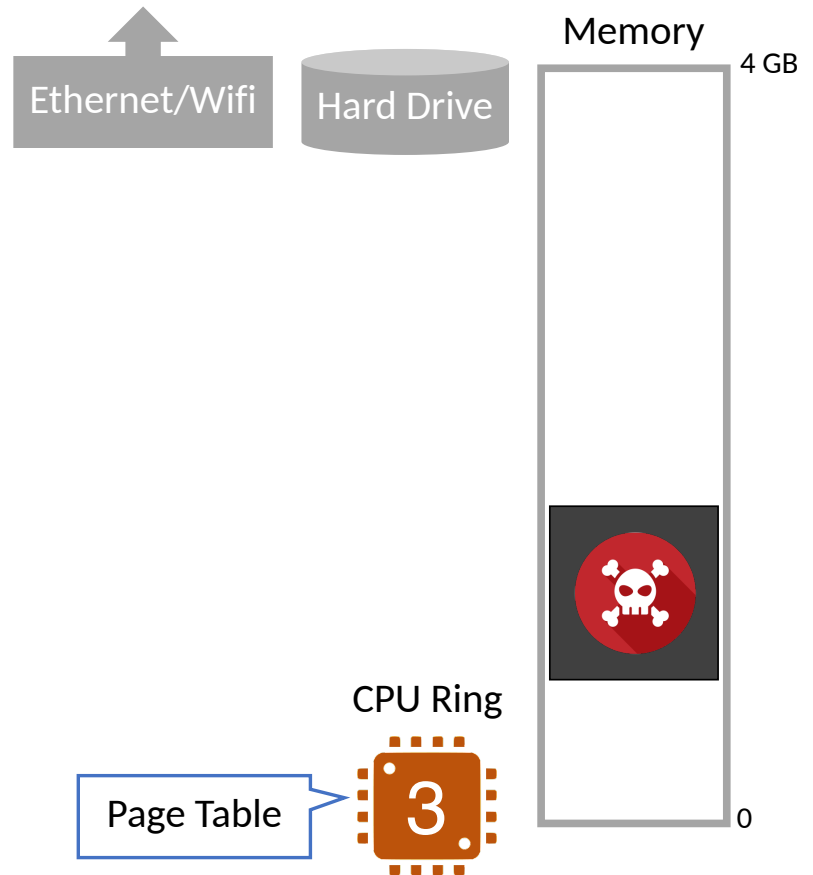
Processes cannot change
their own page tables



VM in Action

Processes can only read/
write within their own
virtual memory

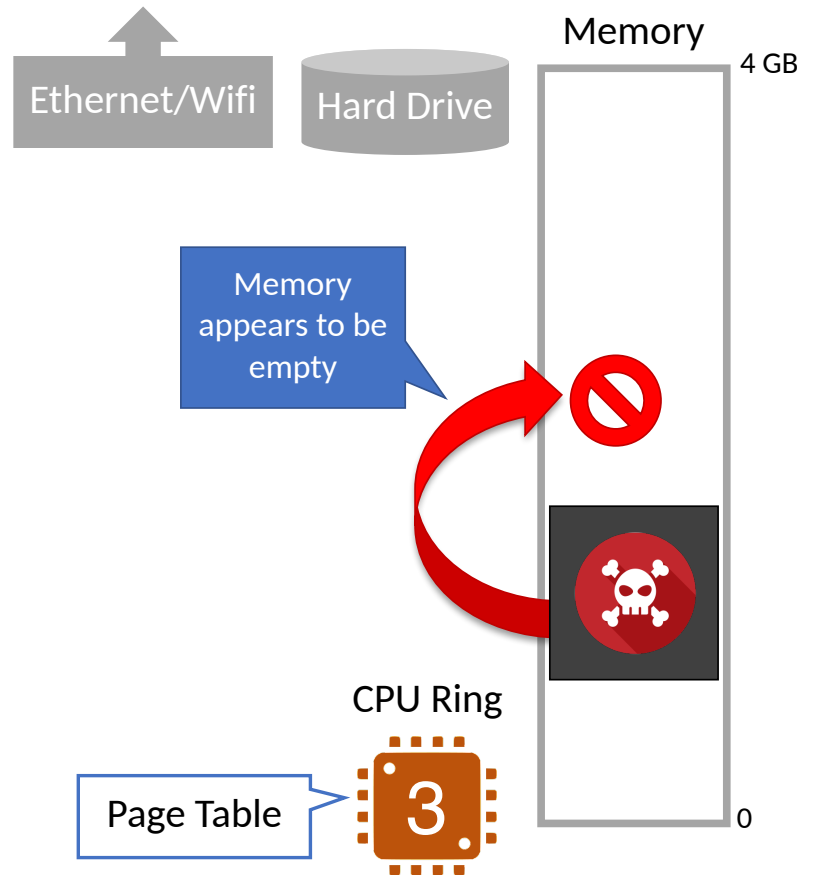
Processes cannot change
their own page tables



VM in Action

Processes can only read/
write within their own
virtual memory


Processes cannot change
their own page tables



VM in Action

Processes can only read/
write within their own
virtual memory

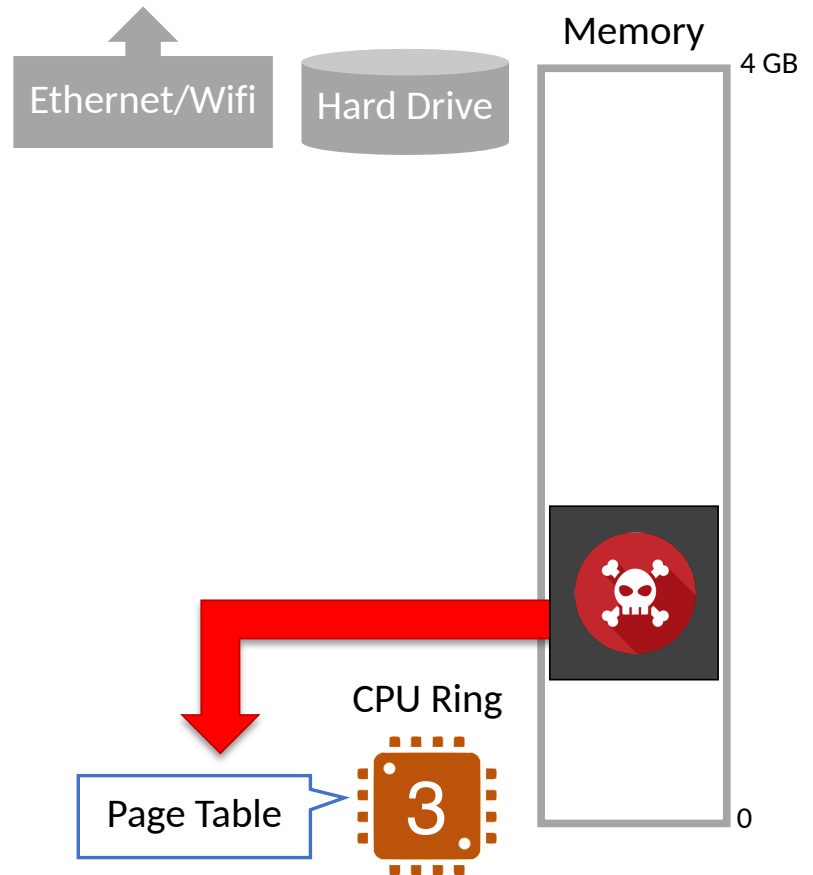
Processes cannot change
their own page tables



VM in Action

Processes can only read/
write within their own
virtual memory

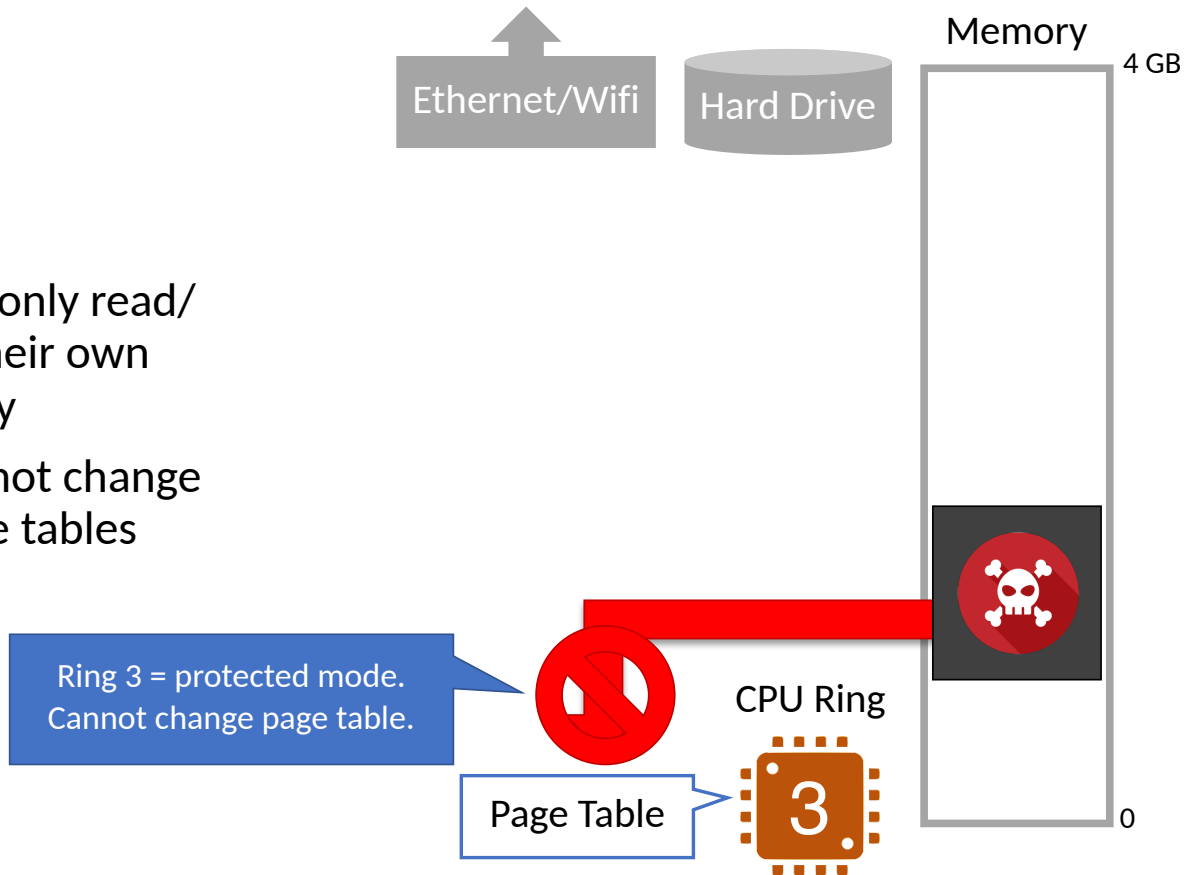
Processes cannot change
their own page tables



VM in Action

Processes can only read/
write within their own
virtual memory

Processes cannot change
their own page tables



Threat Model

Intro to System Architecture

Hardware Support for Isolation

Examples

Principles

Review

At this point, we have achieved process isolation


- Protected mode execution prevents direct device access
- Virtual memory prevents direct memory access

Requires CPU support

- All modern CPUs support these techniques

Requires OS support

- All modern OS support these techniques
- OS controls process rings and page tables



Review

At this point, we have achieved process isolation

- Protected mode execution prevents direct device access
- Virtual memory prevents direct memory access

Requires CPU support

- All modern CPUs support these techniques

Requires OS support

- All modern OS support these techniques
- OS controls process rings and page tables

Warning: bugs in the OS may compromise process isolation



Towards Secure Systems

Now that we have process isolation, we can build more complex security features



File Access Control



Anti-virus



Firewall

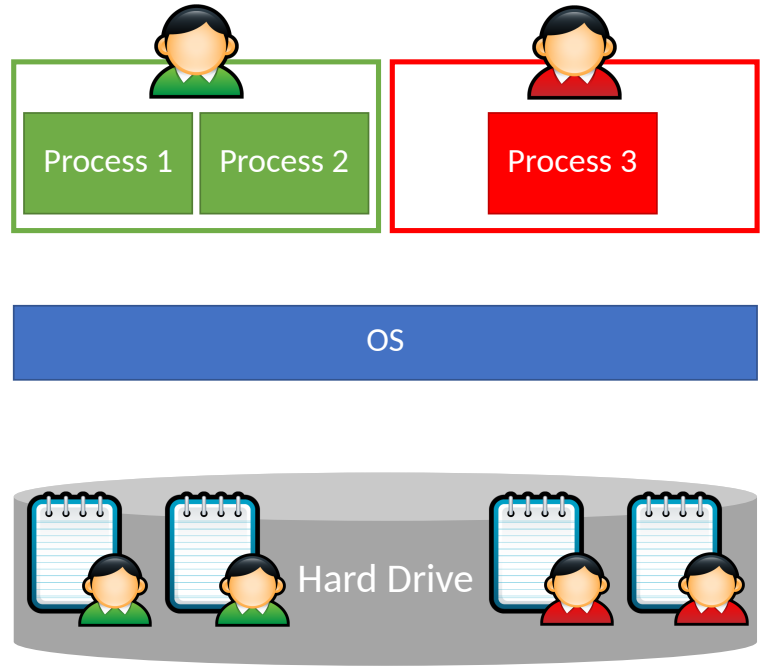


Secure Logging

File Access Control



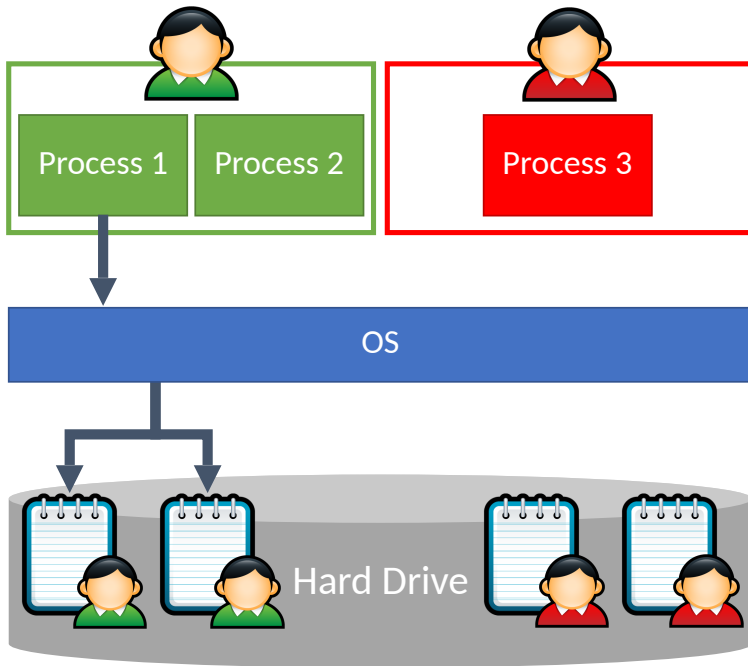
All disk access is mediated
by the OS
OS enforces access controls



File Access Control



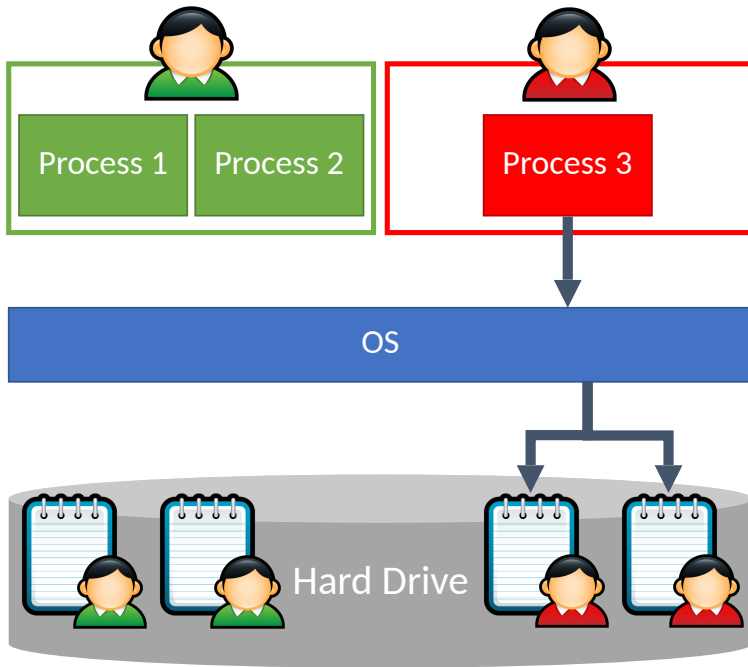
All disk access is mediated by the OS
OS enforces access controls



File Access Control



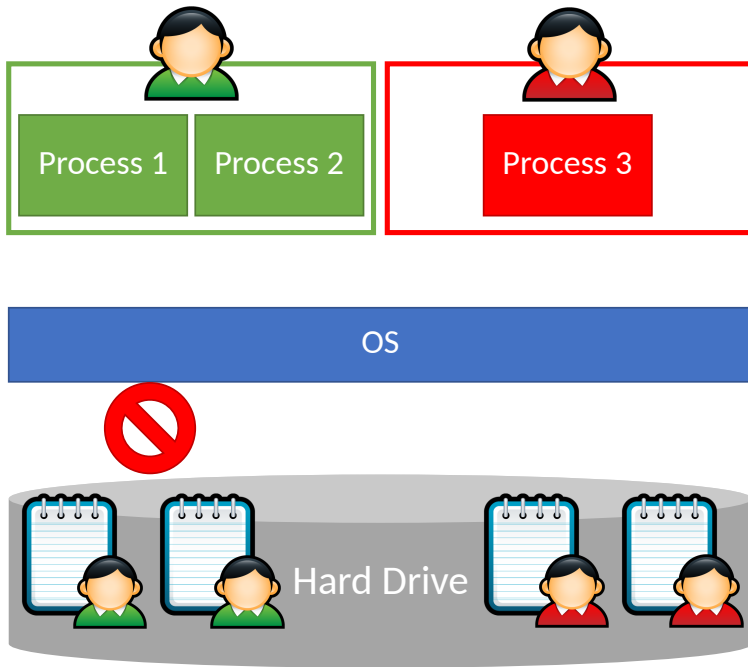
All disk access is mediated
by the OS
OS enforces access controls



File Access Control



All disk access is mediated
by the OS
OS enforces access controls

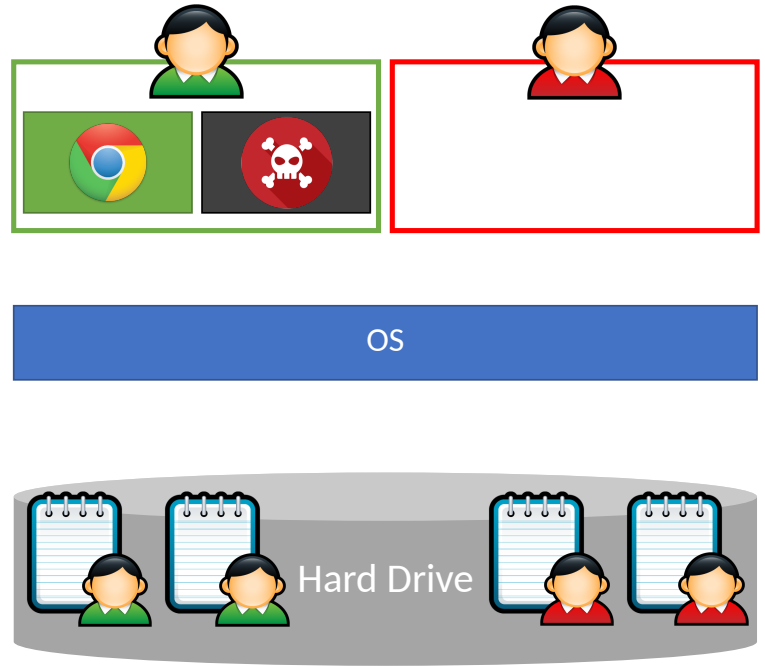




Limitations

Malware can still cause damage

Discretionary access control means that isolation is incomplete

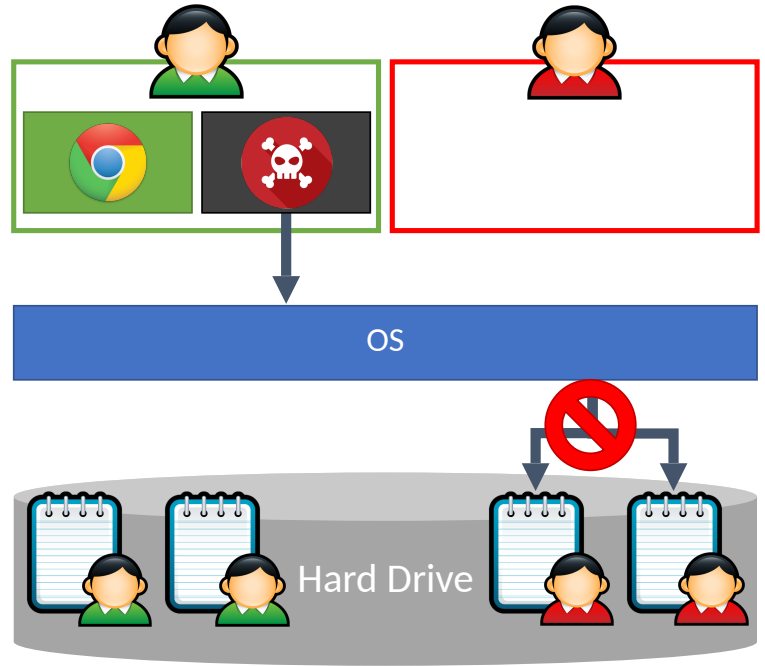




Limitations

Malware can still cause damage

Discretionary access control means that isolation is incomplete






Limitations

Malware can still cause damage

Discretionary access control means that isolation is incomplete



Anti-virus

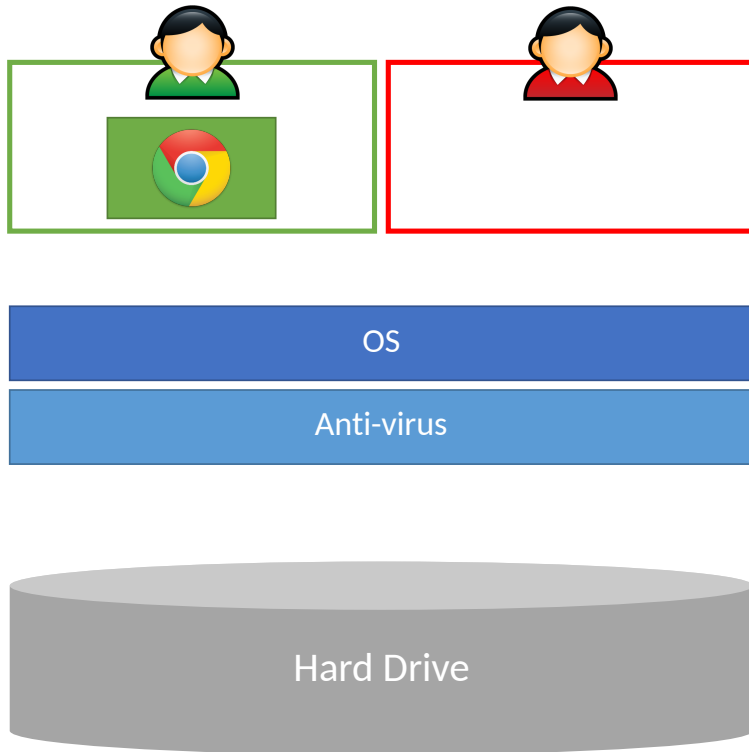
Anti-virus process is **privileged**

- Often runs in Ring 0

Scans all files looking for **signatures**

- Each signature uniquely identifies a piece of malware

Files scanned on creation and access



Anti-virus

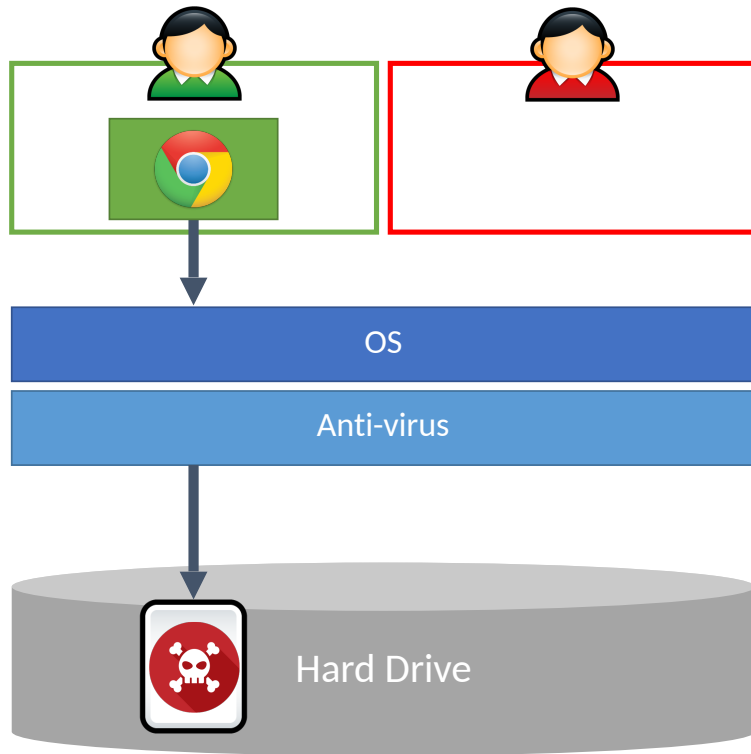
Anti-virus process is **privileged**

- Often runs in Ring 0

Scans all files looking for **signatures**

- Each signature uniquely identifies a piece of malware

Files scanned on creation and access



Anti-virus

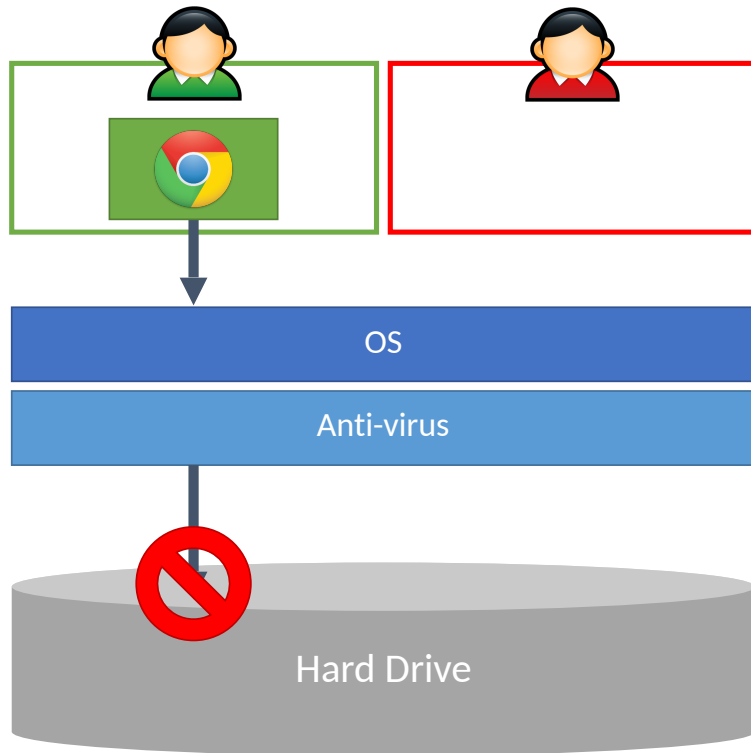
Anti-virus process is **privileged**

- Often runs in Ring 0

Scans all files looking for **signatures**

- Each signature uniquely identifies a piece of malware

Files scanned on creation and access



Anti-virus

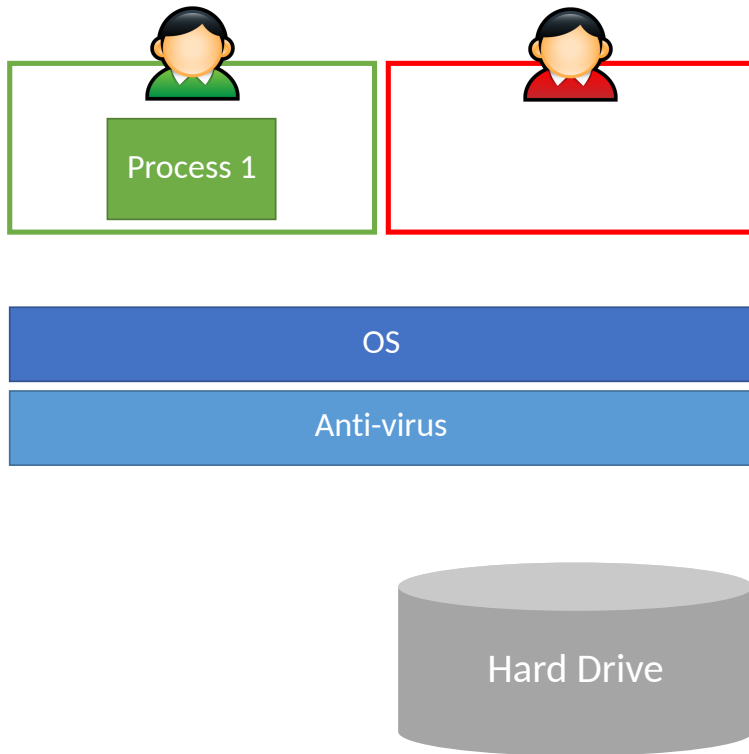
Anti-virus process is
privileged

- Typically runs in Ring 0

Scans all files looking for
signatures

- Each signature uniquely identifies a piece of malware

Files scanned on creation
and access



Anti-virus

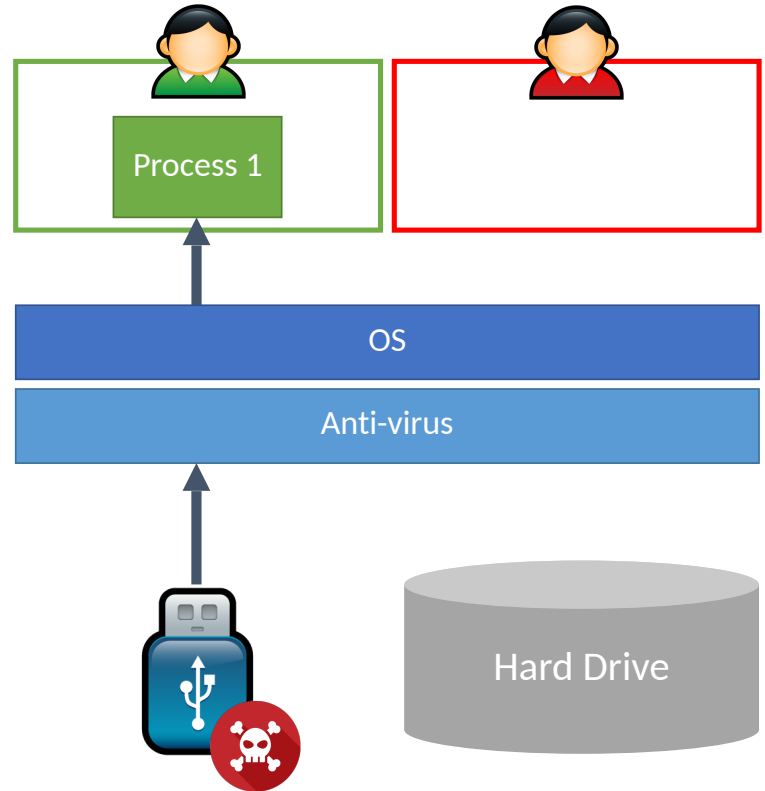
Anti-virus process is **privileged**

- Typically runs in Ring 0

Scans all files looking for **signatures**

- Each signature uniquely identifies a piece of malware

Files scanned on creation and access



Anti-virus

Anti-virus process is **privileged**

- Typically runs in Ring 0

Scans all files looking for **signatures**

- Each signature uniquely identifies a piece of malware

Files scanned on creation and access

