

# 2550 Intro to cybersecurity

## L18: Intro to Systems Security

Ran Cohen/abhi shelat

Thanks to Christo for  
starting point for slides.

Threat Model

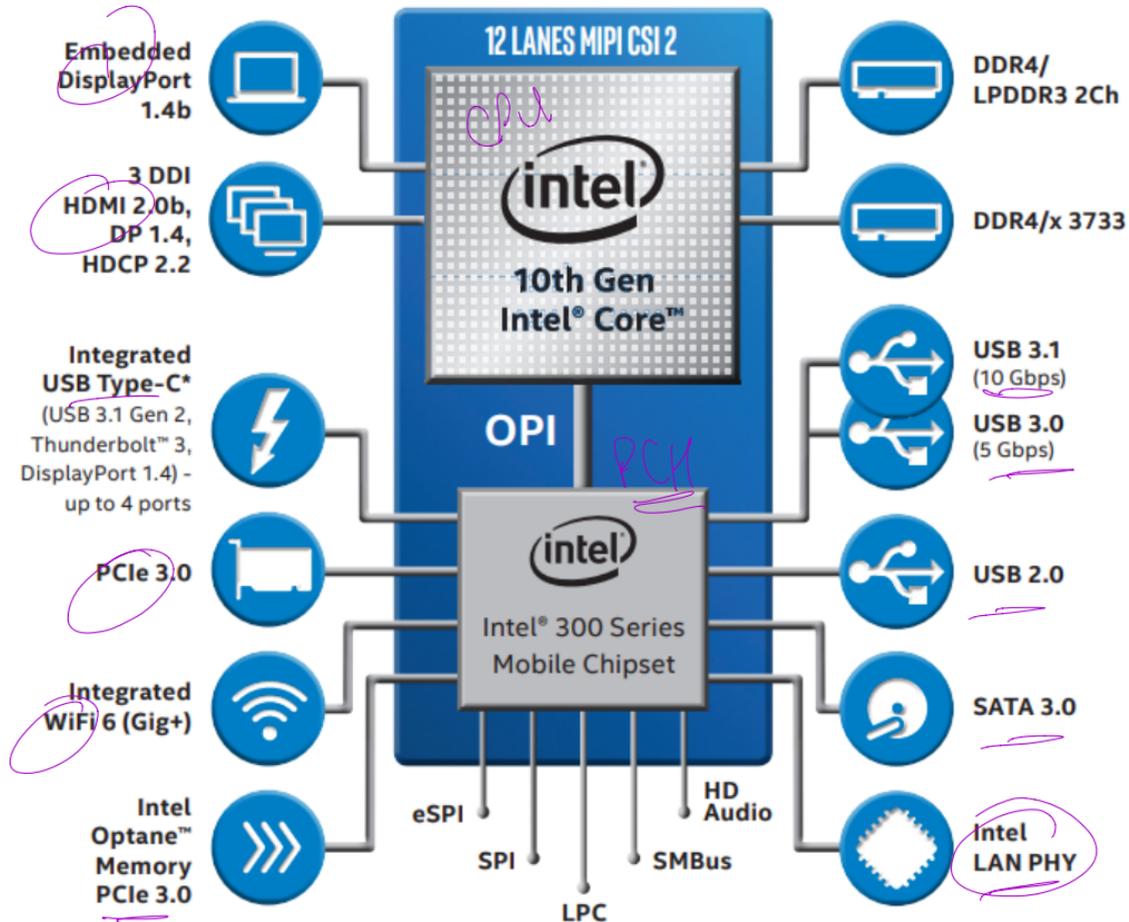
Principles

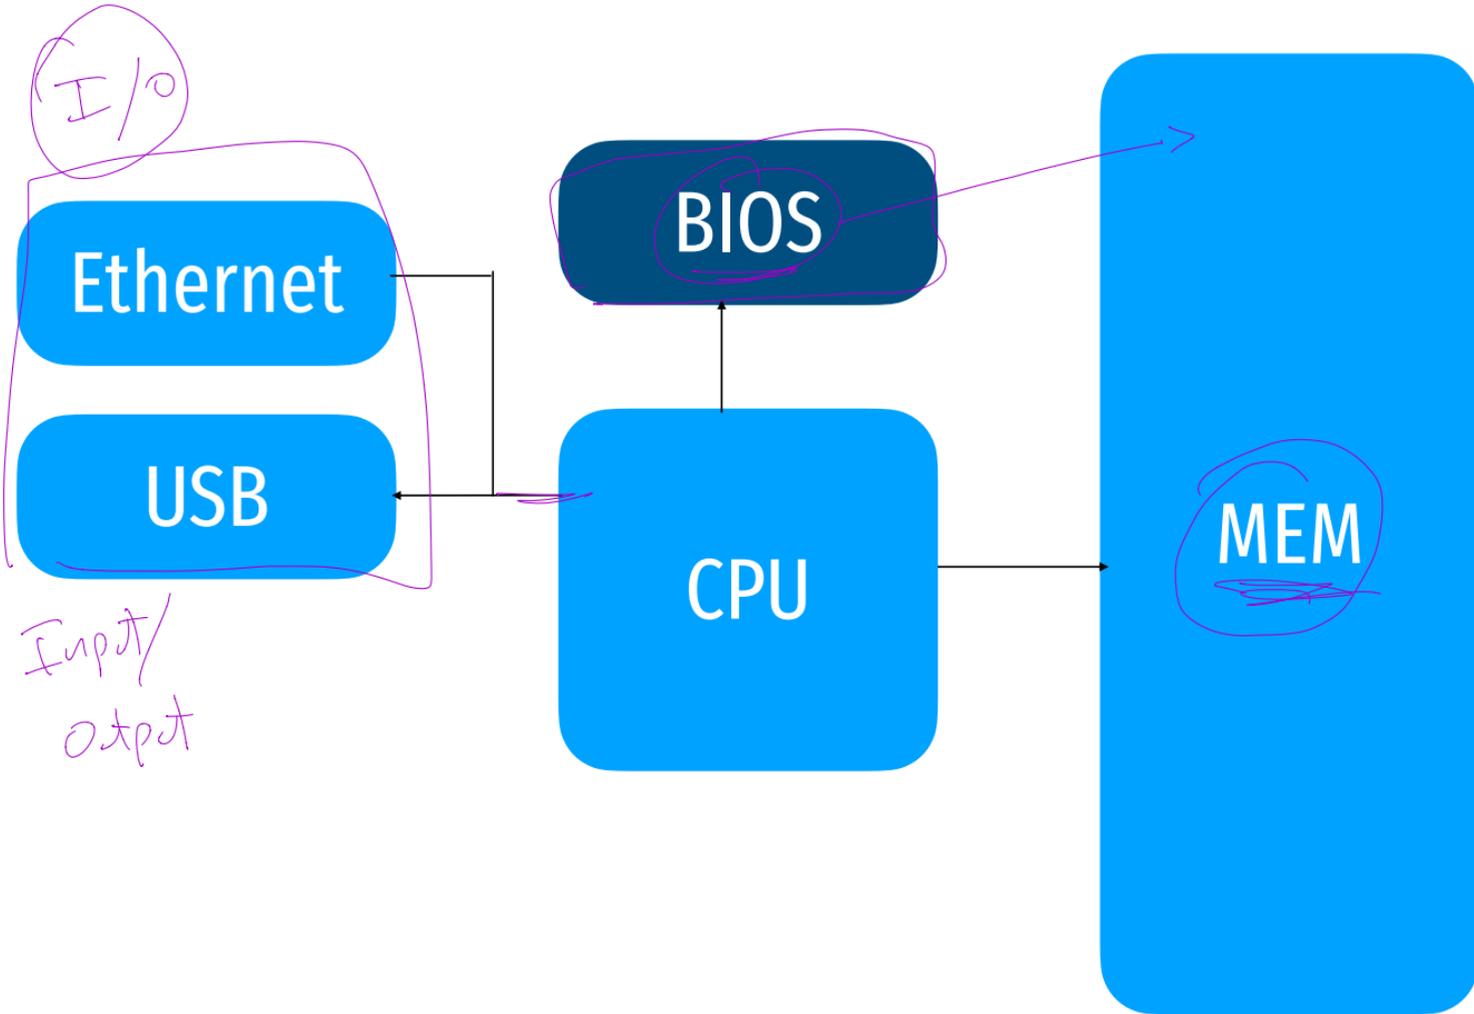
**Intro to System Architecture**

Hardware Support for Isolation

Examples







# 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	&
109	
108	U
107	C
106	U
105	.
104	
103	U
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

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

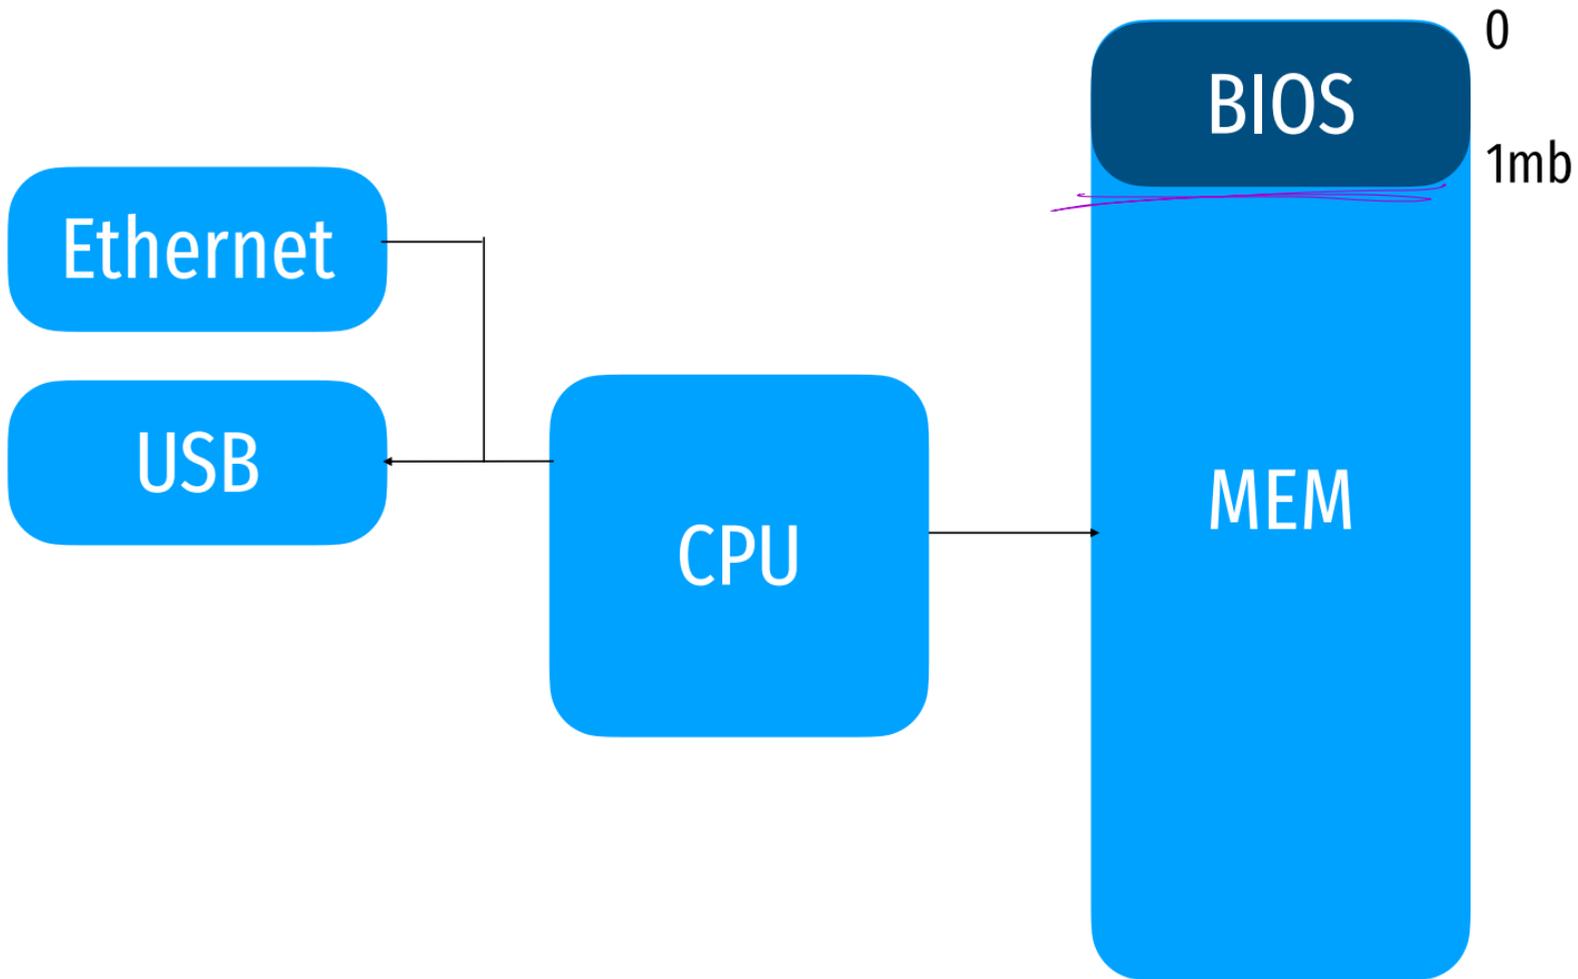
Each ASCII character is one byte, Strings are null terminated

CPUs understand instructions in assembly language

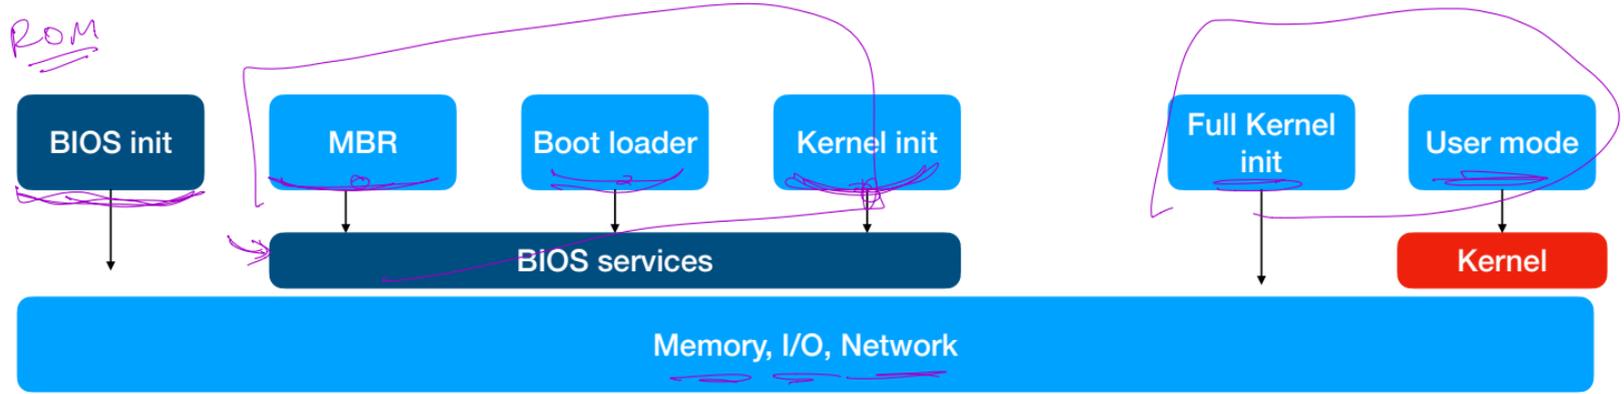
Address	Contents
114	
113	0
112	0
111	0
110	8
109	
108	0
107	C
106	B
105	A
104	
103	0xAF
102	0x3C
101	0x91
100	0xE3

# How does a computer boot?

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

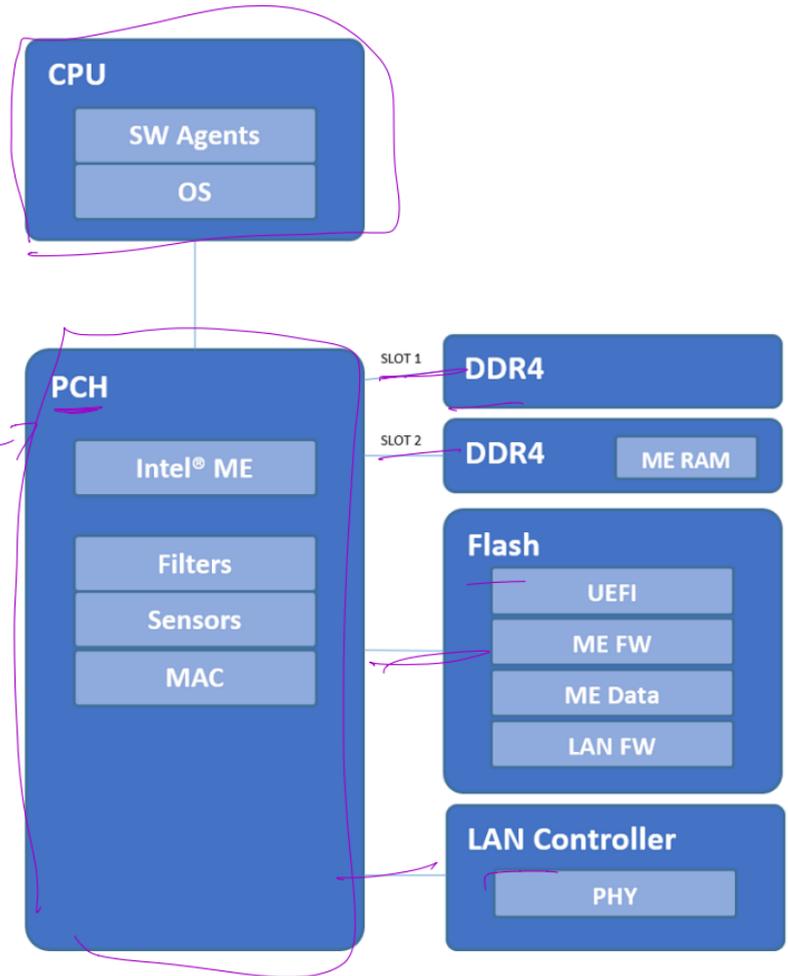


# System Model: how does a computer boot?



# More details

*Platform Controller Hub*



# Layout of memory at boot

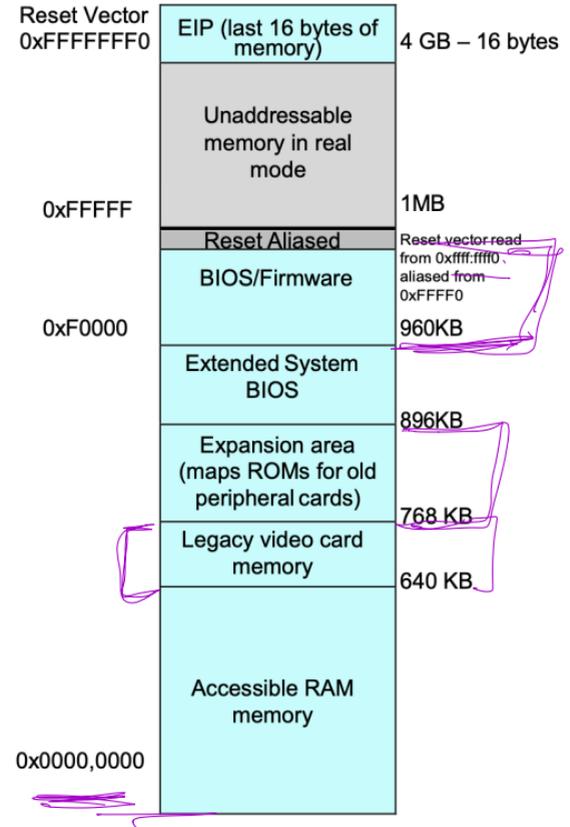


Figure 3 Intel® Architecture Memory Map at Power On

# Details

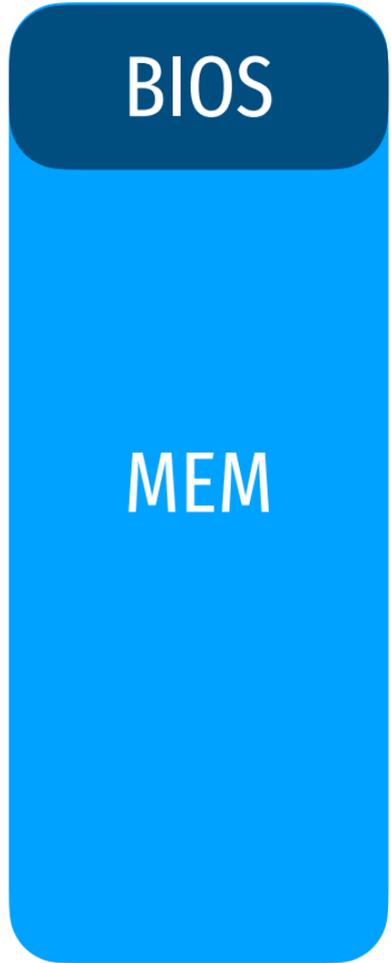
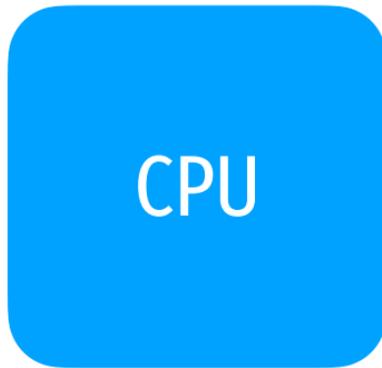
- CPU begins executing at f.fff0
- BIOS firmware begins init of hw
- Applies microcode patches
- Execute Firmware Support Pkg (blob)
- [Ram is setup]
- Copy firmware to RAM
- Begin executing in RAM
- Setup interrupts, timers, clocks
- Bring up other cores
- Setup PCI
- Setup ACPI tables
- Execute OS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

OS

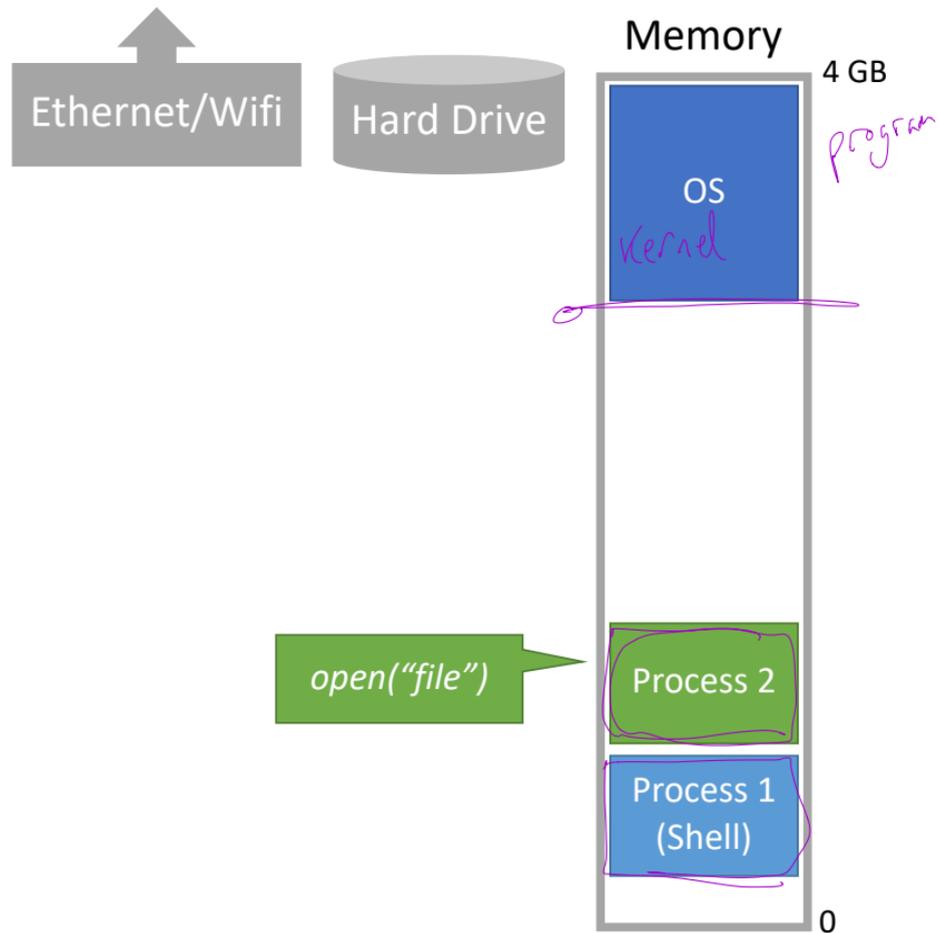
0.0000

in the  
next slides

power



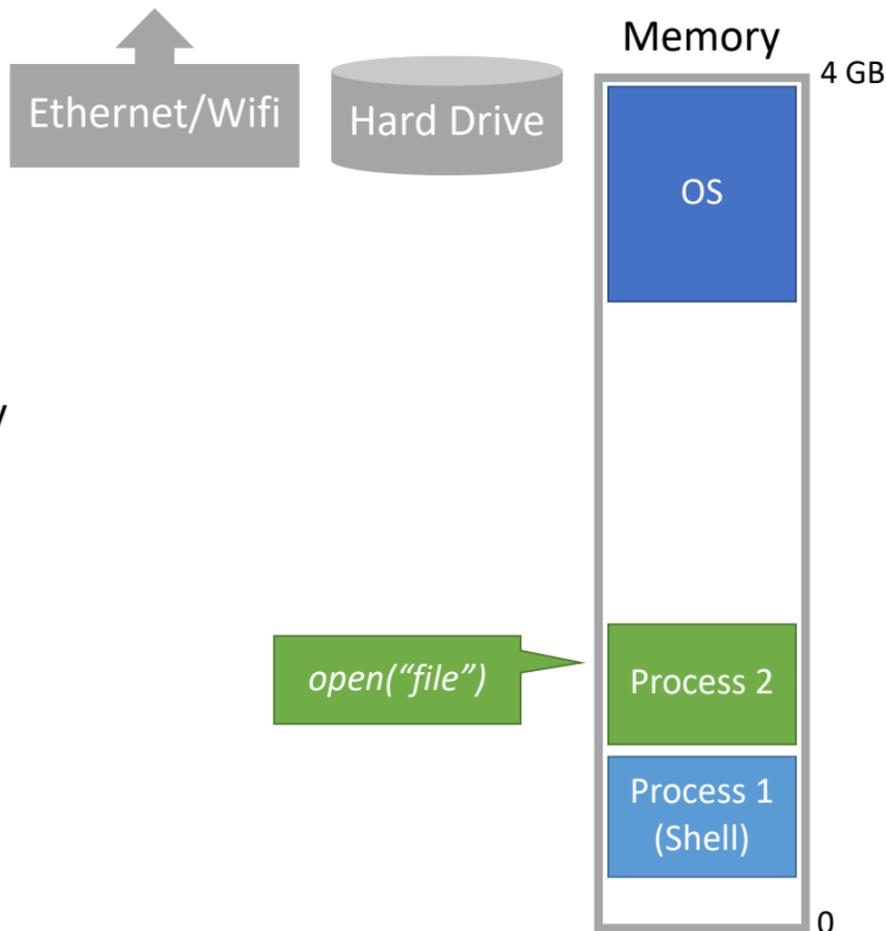
# System Model



# System Model

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

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



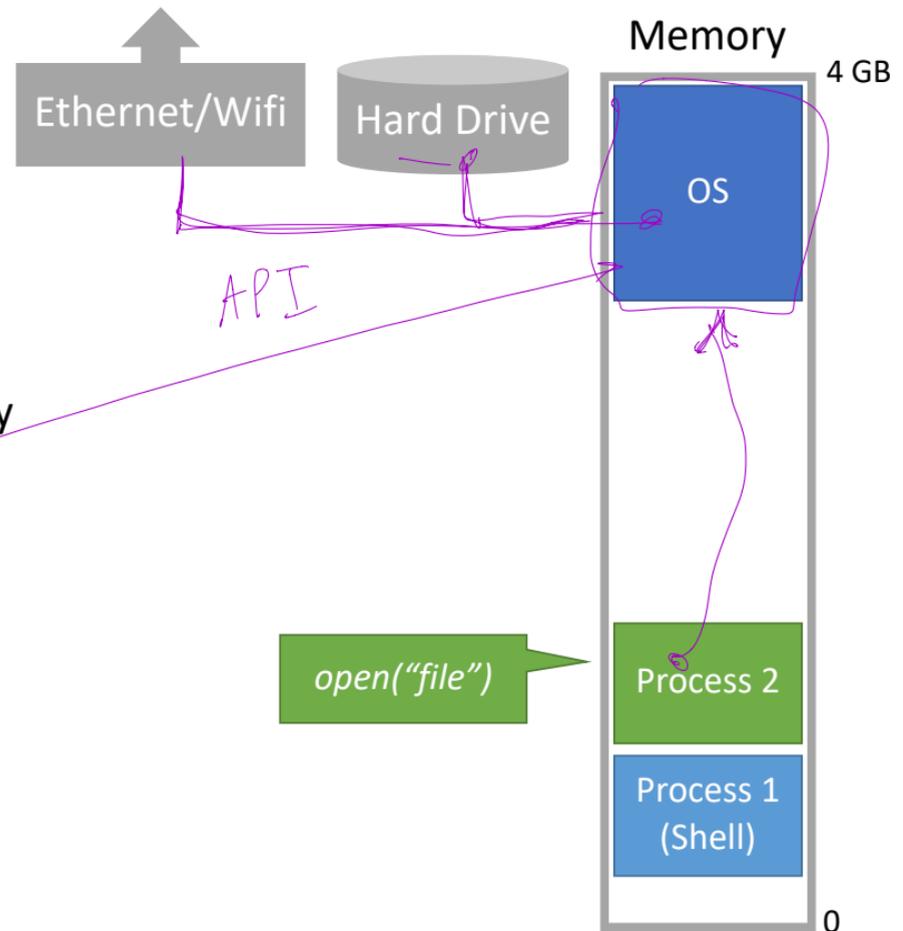
# System Model

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



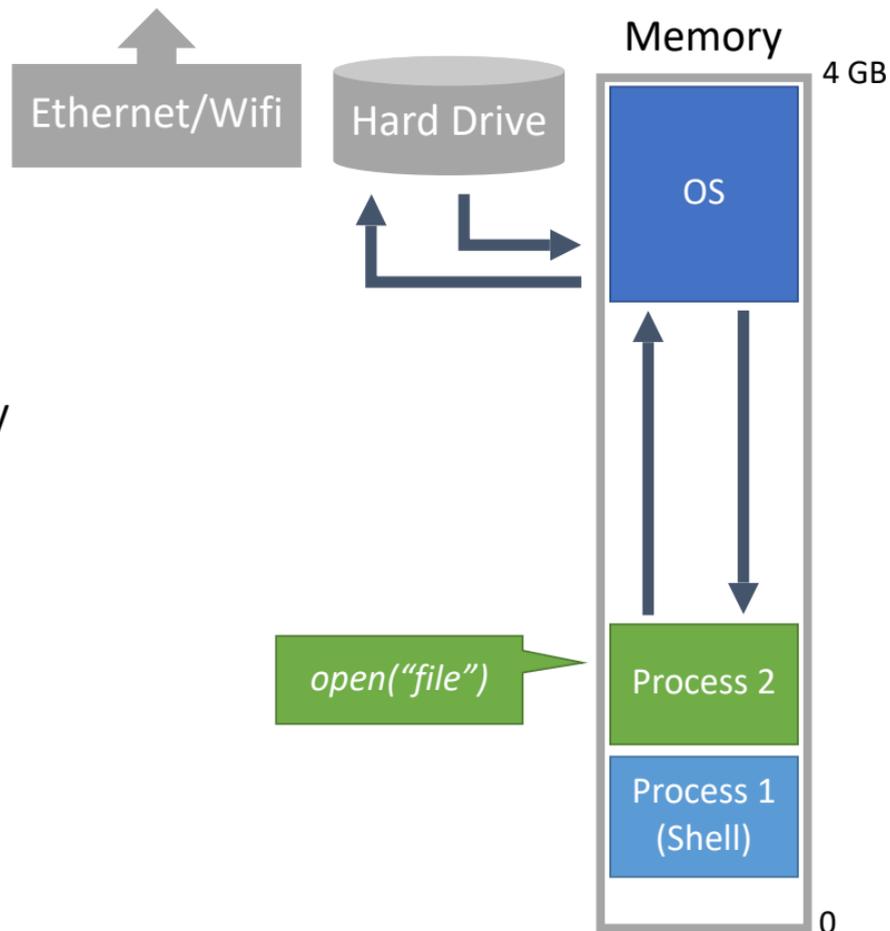
# System Model

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





```
0.000465 MTRR variable ranges enabled:
0.000470 0 base 080000000 mask 7f8000000 uncachable
0.000418 1 base 007c00000 mask 7ffc00000 uncachable
0.000412 2 base 007a00000 mask 7fff00000 uncachable
0.000413 3 base 007500000 mask 7fff00000 uncachable
0.000415 4 base 007f80000 mask 7ff800000 uncachable
0.000417 5 base 200000000 mask 600000000 uncachable
0.000418 6 base 100000000 mask 700000000 uncachable
0.000420 7 base 400000000 mask 400000000 uncachable
0.000421 8 disabled
0.000422 9 disabled
0.001256 x86/PAT: Configuration [0-7]: WB WC UC- UC WB WP UC- WT
0.001570 last_pfn = 0x61c4f_max_arch_jb = 0x400000000
0.021713 err: Reserving ESRT space from 0x000000006146518 to 0x0000000068146550.
0.021720 e820: update [mem 0x8146000-0x8146ffff] usable ==> reserved
0.021847 check: Scanning 1 areas for low memory corruption
Using 6B pages for direct mapping
0.022544 RAMDISK: [mem 0x3ce4000-0x3ffff000]
0.022580 ACPI: Early table checksum verification disabled
0.022584 ACPI: RSDP 0x000000006f17014 000024 (v02 INTEL)
0.022589 ACPI: XSDT 0x000000006f116728 0000CC (v01 INTEL NUC915FN 0000020 AMI 01000013)
0.022598 ACPI: FACP 0x000000006f0d2000 000114 (v06 INTEL NUC915FN 0000020 AMI 00010013)
0.022603 ACPI: DSDT 0x000000006f08f000 042561 (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022610 ACPI: FACS 0x000000006f1B1000 000040
0.022614 ACPI: MCFG 0x000000006f0d5000 00003C (v01 INTEL NUC915FN 0000020 MSFT 00000097)
0.022618 ACPI: SSDT 0x000000006f0d3000 00184A (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022623 ACPI: FIDT 0x000000006f08e000 (v01 INTEL NUC915FN 0000020 AMI 00010013)
0.022627 ACPI: SSDT 0x000000006f08A000 0031C6 (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022631 ACPI: HPET 0x000000006f0d7000 000038 (v01 INTEL NUC915FN 0000020 AMI 01000013)
0.022635 ACPI: SSDT 0x000000006f086000 003384 (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022639 ACPI: SSDT 0x000000006f084000 001478 (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022643 ACPI: SSDT 0x000000006f080000 002820 (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022648 ACPI: NHLT 0x000000006f0d0600 000020 (v06 INTEL NUC915FN 0000020 AMI 01000013)
0.022652 ACPI: LPIT 0x000000006f077000 000094 (v01 INTEL NUC915FN 0000020 AMI 01000013)
0.022656 ACPI: SSDT 0x000000006f075000 002720 (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022660 ACPI: SSDT 0x000000006f074000 00007C (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022664 ACPI: DBGX 0x000000006f079000 000034 (v01 INTEL NUC915FN 0000020 AMI 01000013)
0.022668 ACPI: DBG2 0x000000006f078000 000054 (v06 INTEL NUC915FN 0000020 AMI 01000013)
0.022672 ACPI: SSDT 0x000000006f076000 001866 (v02 INTEL NUC915FN 0000020 INTL 20160527)
0.022677 ACPI: TPM2 0x000000006f074000 00004C (v04 INTEL NUC915FN 0000020 AMI 00000000)
0.022681 ACPI: DMAR 0x000000006f075000 0000A8 (v01 INTEL NUC915FN 0000020 01000013)
0.022685 ACPI: WSMT 0x000000006f07E000 000028 (v01 INTEL NUC915FN 0000020 AMI 00010013)
0.022689 ACPI: APIC 0x000000006f073000 0000F4 (v04 INTEL NUC915FN 0000020 AMI 00010013)
0.022693 ACPI: FPDt 0x000000006f072000 000044 (v01 INTEL NUC915FN 0000020 AMI 01000013)
0.022707 ACPI: Local APIC address 0xfce00000
0.023236 No NUMA configuration found
0.023238 Faking a node at [mem 0x0000000000000000-0x0000000080000000]
0.023254 NODE_DATA(0) allocated [mem 0x280fd5000-0x280ffff00]
0.023698 Zone ranges:
0.023700 DMA 0x00000000000001000-0x000000000000ffff00
0.023702 DMA32 0x0000000001000000-0x000000000000ffff00
0.023703 Normal 0x0000000100000000-0x000000080000ffff00
0.023705 Device empty
0.023706 Movable zone start for each node
0.023711 Early memory node ranges
0.023713 node 0: [mem 0x0000000000001000-0x00000000000009ffff]
0.023714 node 0: [mem 0x0000000001000000-0x000000000000fc53ffff]
0.023716 node 0: [mem 0x000000000fc4000-0x0000000000fc4000]
0.023717 node 0: [mem 0x0000000100000000-0x000000080000ffff00]
0.024270 Zeroed struct page in unavail=able ranges: 41229 pages
0.024272 Initmem setup node 0 [mem 0x00000000000001000-0x000000080000ffff00]
0.024274 On node 0 totalpages: 8314611
0.024276 DMA zone: 64 pages used for memmap
0.024277 DMA zone: 25 pages reserved
0.024278 DMA zone: 3998 pages, LIFO batch:0
0.024379 DMA32 zone: 6918 pages used for memmap
0.024380 DMA32 zone: 442197 pages, LIFO batch:63
0.039990 Normal zone: 122944 pages used for memmap
0.039991 Normal zone: 7868416 pages, LIFO batch:63
```

```
0.207130 Reserving Intel graphics memory at [mem 0x790000000-0x7cffff000]
0.207882 ACPI: PM-Timer IO Port: 0x1808
0.207894 ACPI: Local APIC address 0xfce00000
0.207894 ACPI: LAPIC_NMI (acpi_id[0x01] high edge lint[0x1])
0.207895 ACPI: LAPIC_NMI (acpi_id[0x02] high edge lint[0x1])
0.207896 ACPI: LAPIC_NMI (acpi_id[0x03] high edge lint[0x1])
0.207897 ACPI: LAPIC_NMI (acpi_id[0x04] high edge lint[0x1])
0.207897 ACPI: LAPIC_NMI (acpi_id[0x05] high edge lint[0x1])
0.207898 ACPI: LAPIC_NMI (acpi_id[0x06] high edge lint[0x1])
0.207899 ACPI: LAPIC_NMI (acpi_id[0x07] high edge lint[0x1])
0.207900 ACPI: LAPIC_NMI (acpi_id[0x08] high edge lint[0x1])
0.207901 ACPI: LAPIC_NMI (acpi_id[0x09] high edge lint[0x1])
0.207902 ACPI: LAPIC_NMI (acpi_id[0x0a] high edge lint[0x1])
0.207902 ACPI: LAPIC_NMI (acpi_id[0x0b] high edge lint[0x1])
0.207903 ACPI: LAPIC_NMI (acpi_id[0x0c] high edge lint[0x1])
0.207957 IOAPIC[0]: apic_id=32, version=32, address=0xfce00000, GSI 0-119
0.207960 ACPI: INT_SRC_OVR (bus 0 bus_irq 9 global_irq 2 dfl dfl)
0.207962 ACPI: INT_SRC_OVR (bus 0 bus_irq 9 global_irq 9 high level)
0.207963 ACPI: IRQ0 used by override.
0.207965 ACPI: IRQ9 used by override.
0.207968 Using ACPI (MADT) for SMP configuration information
0.207969 ACPI: NPET id: 0x8086a201 base: 0xfed00000
0.207974 smpboot: Allowing 12 CPUs, 0 hotplug CPUs
0.208000 PM: Registered nosave memory: [mem 0x00000000-0x00000000]
0.208003 PM: Registered nosave memory: [mem 0x00000000-0x00000000]
0.208005 PM: Registered nosave memory: [mem 0x657d3000-0x657d3fff]
0.208008 PM: Registered nosave memory: [mem 0x68146000-0x68146fff]
0.208012 PM: Registered nosave memory: [mem 0x6f540000-0x6f540fff]
0.208013 PM: Registered nosave memory: [mem 0x6f858000-0x6f858fff]
0.208014 PM: Registered nosave memory: [mem 0x6f86a000-0x6f86afff]
0.208015 PM: Registered nosave memory: [mem 0x6f86b000-0x6f86bfff]
0.208015 PM: Registered nosave memory: [mem 0x6f869000-0x6f869fff]
0.208016 PM: Registered nosave memory: [mem 0x6f1b2000-0x6f1b4fff]
0.208018 PM: Registered nosave memory: [mem 0x6fc4f000-0x7cffff000]
0.208019 PM: Registered nosave memory: [mem 0x7d000000-0x7dffff000]
0.208020 PM: Registered nosave memory: [mem 0x80000000-0x80ffff000]
0.208021 PM: Registered nosave memory: [mem 0xf0000000-0xf0ffff000]
0.208022 PM: Registered nosave memory: [mem 0xf0e00000-0xf0e10fff]
0.208022 PM: Registered nosave memory: [mem 0xf0e11000-0xf0e1bfff]
0.208023 PM: Registered nosave memory: [mem 0xfce00000-0xfce00fff]
0.208024 PM: Registered nosave memory: [mem 0xfce01000-0xfce0cfff]
0.208025 PM: Registered nosave memory: [mem 0xfed00000-0xfed03fff]
0.208025 PM: Registered nosave memory: [mem 0xfed04000-0xfed0cfff]
0.208026 PM: Registered nosave memory: [mem 0xfce00000-0xfce00fff]
0.208027 PM: Registered nosave memory: [mem 0xfce01000-0xfce0cfff]
0.208028 PM: Registered nosave memory: [mem 0xf0000000-0xf0ffff000]
0.208030 [mem 0x7d000000-0x7dffff000] available for PCI devices
0.208033 Booting paravirtualized kernel on bare hardware
0.208036 clocksource: refined-jiffies: mask: 0xffffffff max_cycles: 0xffffffff, max_idle_ns: 7645519600211568 ns
0.208039 setup_percpu: NR_CPUS=8192 nr_cpumask_bits=12 nr_cpu_ids=12 nr_node_ids=1
0.208519 percpu: Embedded 54 pages/cpu s184320 r8192 u262144
0.208530 percpu-alloc: s184320 r8192 d28672 u262144 alloc=1*2097152
0.208531 percpu-alloc: [0] 00 01 02 03 04 05 06 07 [0] 08 09 10 11 --- --- ---
0.208576 Built 1 zonelists, mobility grouping on. Total pages: 8184668
0.208577 Policy zone: Normal
0.208579 Kernel command line: BOOT_IMAGE=/boot/vmlinuz-5.3.0-64-generic root=/dev/mapper/vgubuntu-root ro quiet splash vt.handoff=7
0.211269 node: detecting cache hash table entries: 4194304 (order: 13, 33554432 bytes, linear)
0.211273 node: cache hash table entries: 2057152 (order: 12, 16777216 bytes, linear)
0.212724 mem auto-init: stack:off, heap alloc:on, heap free:off
0.217959 Calgary: detecting Calgary via BIOS EBDA area
0.217961 Calgary: unable to locate Rio Grande table in EBDA - bailing!
0.326171 Memory: 32418820K/3255844K available (14339K kernel code, 2387K rwdata, 4732K rodata, 2680K init, 5040K bss, 847616K reserved, 0K cma-reserved)
0.326180 random: get_random_u64 called from kernel.c:open@2d0410 with crng_init=0
0.326431 SLUB: HWalign=64, Order=0-3, MinObject=0, CPUs=12, Nodes=1
0.326453 ftrace: allocating 43632 entries in 171 pages
0.355621 rcu: Hierarchical RCU implementation.
0.355623 rcu: Adjusting RCU structure from NR_CPUS=8192 to nr_cpu_ids=12.
0.355624 Tasks RCU enabled.
0.355626 rcu: RCU calculated value of scheduler-enlistment delay is 25 jiffies.
0.355627 rcu: Adjusting geometry for rcu_fanout_leaf=16, nr_cpu_ids=12
0.361328 NR_IRQS: 524544, nr_irqs: 2152, prealloc_fiq_irqs: 16
0.361842 random: crng done (trusting CPU's manufacturer)
0.361879 vt handoff: transparent VT on vt#7
0.361889 Console: colour dummy device 80x25
0.361895 printk: console [tty] enabled
0.361918 ACPI: Core revision 20100703
0.362625 clocksource: hpnet: mask: 0xffffffff max_cycles: 0xffffffff, max_idle_ns: 79635855245 ns
0.362761 APIC: Switch to symmetric I/O mode setup
0.362764 DMAR: Host address width 39
0.362766 DMAR: DRHD base: 0x00000000900000 flags: 0x0
0.362774 DMAR: dmars: 0x reg_base_addr fed90000 ver: 1x 0 cap 10000c40660462 ecap 19e2ff050e
```

```
0.696153 PCI host bridge to bus 0000:00
0.696156 pci bus 0000:00: root bus resource [io 0x0000-0x0cf7 window]
0.696158 pci bus 0000:00: root bus resource [io 0x0d00-0xffff window]
0.696159 pci bus 0000:00: root bus resource [mem 0x000a0000-0x000bffff window]
0.696161 pci bus 0000:00: root bus resource [mem 0x000e0000-0x000e3fff window]
0.696162 pci bus 0000:00: root bus resource [mem 0x000e4000-0x000e7fff window]
0.696163 pci bus 0000:00: root bus resource [mem 0x000e8000-0x000effff window]
0.696165 pci bus 0000:00: root bus resource [mem 0x000ec000-0x000effff window]
0.696166 pci bus 0000:00: root bus resource [mem 0x000f0000-0x000fffff window]
0.696167 pci bus 0000:00: root bus resource [mem 0x7d000000-0xd0ffffff window]
0.696169 pci bus 0000:00: root bus resource [mem 0x400000000-0x7fffffff window]
0.696170 pci bus 0000:00: root bus resource [mem 0xfc000000-0xfe7ffffff window]
0.696172 pci bus 0000:00: root bus resource [bus 00-fe]
0.696177 pci 0000:00:00.0 [8086:9b51] type 00 class 0x050000
0.696192 pci 0000:00:02.0 [8086:9bca] type 00 class 0x030000
0.696599 pci 0000:00:02.0 reg 0x18: [mem 0x602200000-0x6022ffffff 64bit]
0.696967 pci 0000:00:02.0 reg 0x18: [mem 0x400000000-0x400ffffff 64bit pref]
0.696973 pci 0000:00:02.0 reg 0x20: [io 0x3000-0x303f]
0.697334 pci 0000:00:08.0 [8086:1911] type 00 class 0x080000
0.697352 pci 0000:00:08.0 reg 0x18: [mem 0x602312000-0x6023120fff 64bit]
0.697682 pci 0000:00:12.0 [8086:0219] type 00 class 0x118000
0.697708 pci 0000:00:12.0 reg 0x18: [mem 0x60231f000-0x60231ffff 64bit]
0.698071 pci 0000:00:14.0 [8086:02ed] type 00 class 0x0c2300
0.698096 pci 0000:00:14.0 reg 0x18: [mem 0x602310000-0x602310ffff 64bit]
0.698176 pci 0000:00:14.0 PME# supported from D3hot D3cold
0.698639 pci 0000:00:14.2 [8086:02ef] type 00 class 0x050000
0.698662 pci 0000:00:14.2 reg 0x18: [mem 0x602318000-0x602319fff 64bit]
0.698675 pci 0000:00:14.2 reg 0x18: [mem 0x60231e000-0x60231efff 64bit]
0.699033 pci 0000:00:14.3 [8086:0210] type 00 class 0x028000
0.699130 pci 0000:00:14.3 reg 0x18: [mem 0x602314000-0x602317fff 64bit]
0.699383 pci 0000:00:14.3 PME# supported from D0 D3hot D3cold
0.699892 pci 0000:00:15.0 [8086:0268] type 00 class 0xc8000
0.700024 pci 0000:00:15.0 reg 0x18: [mem 0x00000000-0x0000ffff 64bit]
0.700757 pci 0000:00:15.2 [8086:02ea] type 00 class 0xc8000
0.700899 pci 0000:00:15.2 reg 0x18: [mem 0x00000000-0x0000ffff 64bit]
0.701695 pci 0000:00:16.0 [8086:0268] type 00 class 0x780000
0.701726 pci 0000:00:16.0 reg 0x18: [mem 0x60231b000-0x60231bfff 64bit]
0.702164 pci 0000:00:17.0 [8086:02d3] type 00 class 0x010601
0.702186 pci 0000:00:17.0 reg 0x18: [mem 0x96220000-0x96221fff]
0.702195 pci 0000:00:17.0 reg 0x14: [mem 0x96223000-0x962230fff]
0.702204 pci 0000:00:17.0 reg 0x18: [io 0x3090-0x3097]
0.702213 pci 0000:00:17.0 reg 0x1c: [io 0x3080-0x3083]
0.702222 pci 0000:00:17.0 reg 0x20: [io 0x3060-0x307f]
0.702230 pci 0000:00:17.0 reg 0x24: [mem 0x96222000-0x962227fff]
0.814539 NET: Registered protocol family 2
0.814777 [ 0.814777] Listen portaddr hash table entries: 16384 (order: 6, 262144 bytes, linear)
0.815038 TCP established hash table entries: 262144 (order: 9, 2097152 bytes, linear)
0.815477 TCP bind hash table entries: 65536 (order: 8, 1048576 bytes, linear)
0.815587 TCP: Hash tables configured (established 262144 bind 65536)
0.815676 UDP hash table entries: 16384 (order: 7, 524288 bytes, linear)
0.815819 UDP-Lite hash table entries: 16384 (order: 7, 524288 bytes, linear)
0.815965 NET: Registered protocol family 1
0.815971 NET: Registered protocol family 44
0.815984 pci 0000:00:02.0: Video device with shadowed ROM at [mem 0x000c0000-0x000dffff]
0.816616 pci 0000:01:00.0: CLS mismatch (64 is 128), using 64 bytes
0.816828 pci 0000:01:00.0: enabling device (0002 -> 0003)
0.817224 Trying to unpack roots image as initramfs...
0.986166 Initramfs unpacking failed: Decoding failed
0.986397 Freeing initramfs memory: 580556K
0.993425 DMAR: Intel-IOMMU force enabled due to platform opt in
0.993459 DMAR: No ATRP found
0.993546 DMAR: dmarr0: Using Queued invalidation
4.346737 idma64 idma64.1: Found Intel integrated DMA 64-bit
4.351278 mei_me 0000:00:16.0: hbm: dma setup response: failure = 3 REJECTED
4.352899 Bluetooth: Core ver 2.22
4.352898 NET: Registered protocol family 31
4.352898 Bluetooth: HCI device and connection manager initialized
4.352901 Bluetooth: HCI socket layer initialized
4.352902 Bluetooth: L2CAP socket layer initialized
4.352903 Bluetooth: SCO socket layer initialized
4.389562 88x2bu: loading out-of-tree module taints kernel.
4.403328 cryptd: max_cpu_qlen set to 1000
4.406575 Intel(R) Wireless WiFi driver for Linux
4.406576 Copyright(c) 2003-2015 Intel Corporation
4.406641 iwlmwifi 0000:00:14.3: enabling device (0000 -> 0002)
4.415431 iwlmwifi 0000:00:14.3: TLV_FW_FSEQ_VERSION: FSEQ Version: 43.2.23.17
4.415434 iwlmwifi 0000:00:14.3: Found debug destination: EXTERNAL_DRAM
4.415435 iwlmwifi 0000:00:14.3: Found debug configuration: 0
4.415592 iwlmwifi 0000:00:14.3: loaded firmware version 48.4fa0041.f0 op_mode iwlmvm
4.424971 AVX2 version of gcm_enc/dec engaged.
4.424972 AES CTR mode: by8 optimization enabled
4.427194 usbcore: registered new interface driver btusb
4.428100 Bluetooth: hcib: Firmware revision 0.0 build 62 week 31 2019
3.103257 usb usb1: New USB device found, idVendor=106b, idProduct=0002, bcdDevice= 5.03
3.103259 usb usb1: New USB device strings: Mfr=3, Product=2, SerialNumber=1
3.103261 usb usb1: Product: xHCI Host Controller
3.103263 usb usb1: Manufacturer: Linux 5.3.0-64-generic xhci-hcd
3.103264 usb usb1: SerialNumber: 0000:00:14.0
3.103469 hub 1-0:1.0: USB hub found
3.103489 hub 1-0:1.0: 12 ports detected
3.106061 xhci_hcd 0000:00:14.0: xHCI Host Controller
3.106066 xhci_hcd 0000:00:14.0: new USB bus registered, assigned bus number 2
3.106071 xhci_hcd 0000:00:14.0: Host supports USB 3.1 Enhanced SuperSpeed
3.106121 usb usb2: New USB device found, idVendor=106b, idProduct=0003, bcdDevice= 5.03
3.106122 usb usb2: New USB device strings: Mfr=3, Product=2, SerialNumber=1
3.106124 usb usb2: Product: xHCI Host Controller
3.106126 usb usb2: Manufacturer: Linux 5.3.0-64-generic xhci-hcd
3.106127 usb usb2: SerialNumber: 0000:00:14.0
3.106307 hub 2-0:1.0: USB hub found
3.106321 hub 2-0:1.0: 6 ports detected
3.107685 usb: port power management may be unreliable
3.107867 xhci_hcd 0000:39:00.0: xHCI Host Controller
3.107874 xhci_hcd 0000:39:00.0: new USB bus registered, assigned bus number 3
3.109069 xhci_hcd 0000:39:00.0: hcc params 0x200077c1 hci version 0x110 quirks 0x0000000200009810
3.109318 usb usb3: New USB device found, idVendor=106b, idProduct=0002, bcdDevice= 5.03
3.109319 usb usb3: New USB device strings: Mfr=3, Product=2, SerialNumber=1
3.109322 usb usb3: Product: xHCI Host Controller
3.109323 usb usb3: Manufacturer: Linux 5.3.0-64-generic xhci-hcd
3.109325 usb usb3: SerialNumber: 0000:39:00.0
3.109508 hub 3-0:1.0: USB hub found
3.109520 hub 3-0:1.0: 2 ports detected
3.109708 xhci_hcd 0000:39:00.0: xHCI Host Controller
```

# What happens after boot?

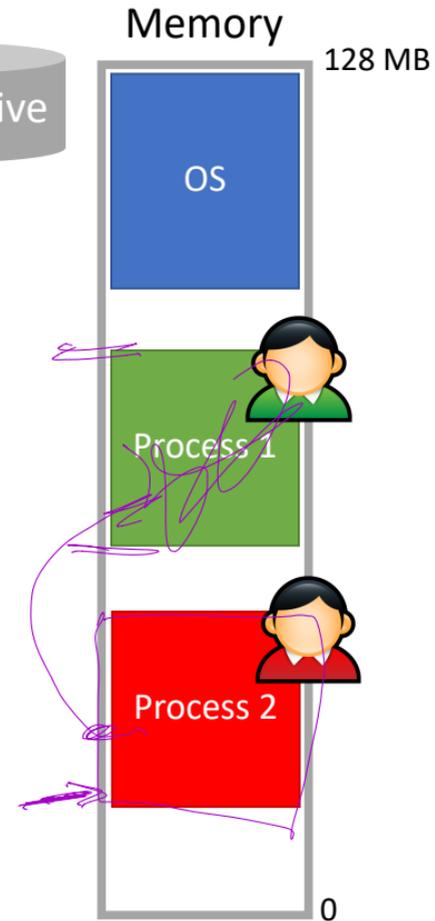
```
CentOS Linux 7 (Core)  
Kernel 3.10.0-327.el7.x86_64 on an x86_64  
rhcsa login: _
```

# Memory Unsafety

~~DOS~~

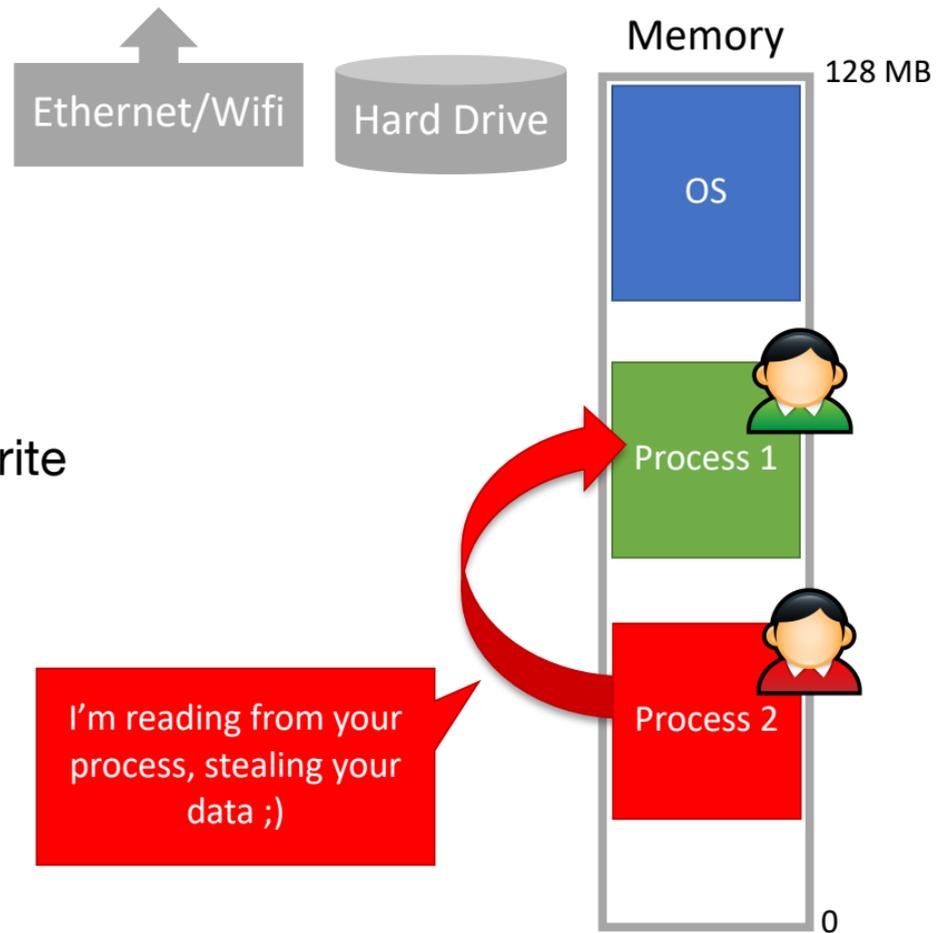


Problem: any process can read/write any memory



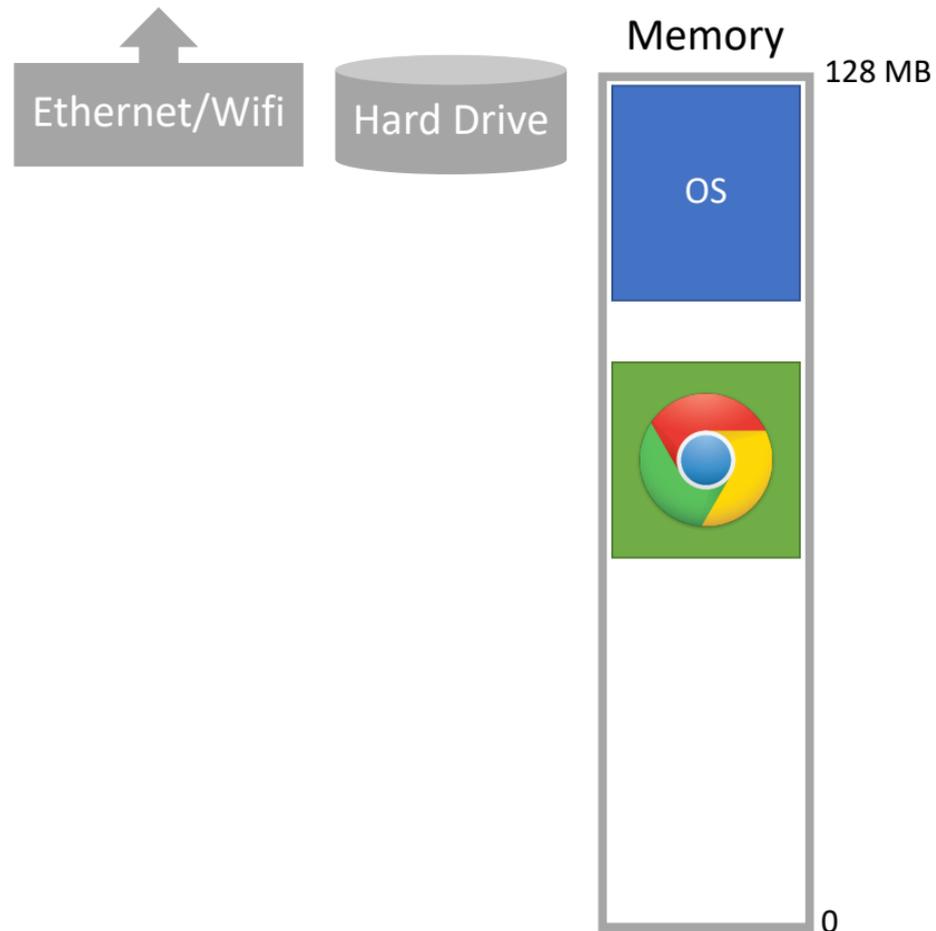
# Memory Unsafety

Problem: any process can read/write any memory



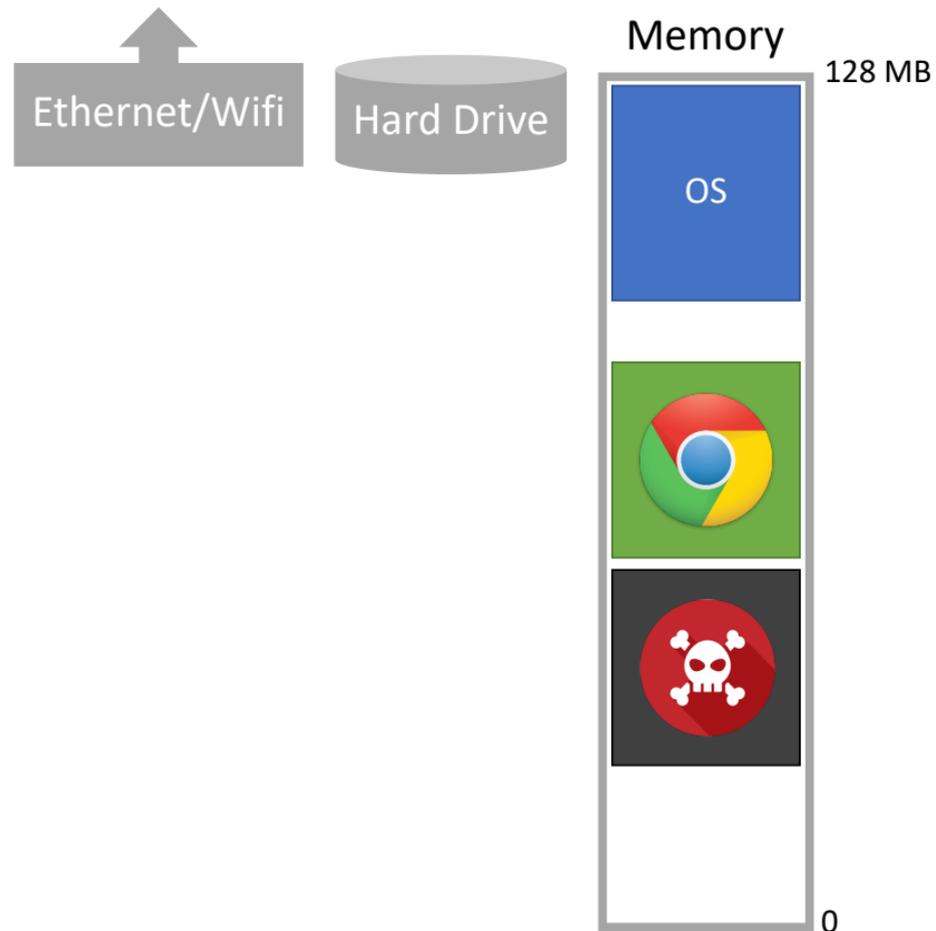
# Memory Unsafety

Problem: any process can read/write any memory



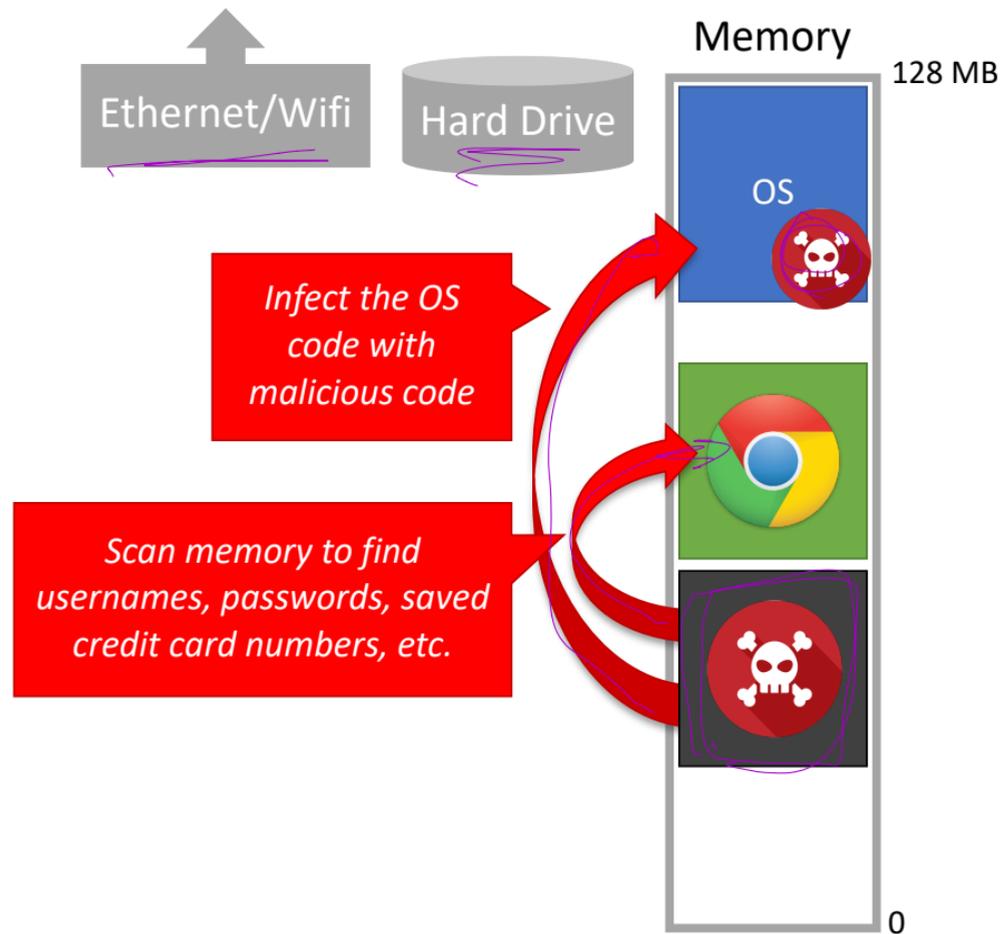
# Memory Unsafety

Problem: any process can read/write any memory



# Memory Unsafety

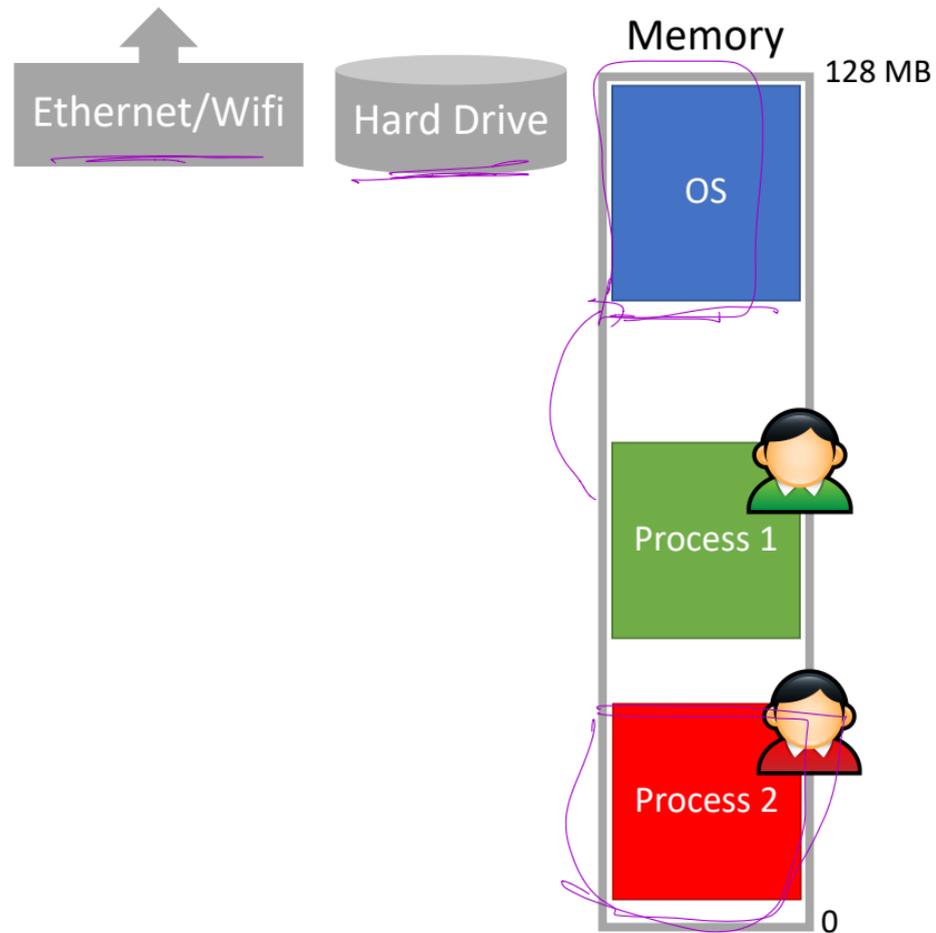
Problem: any process can read/write any memory



# Device Unsafety

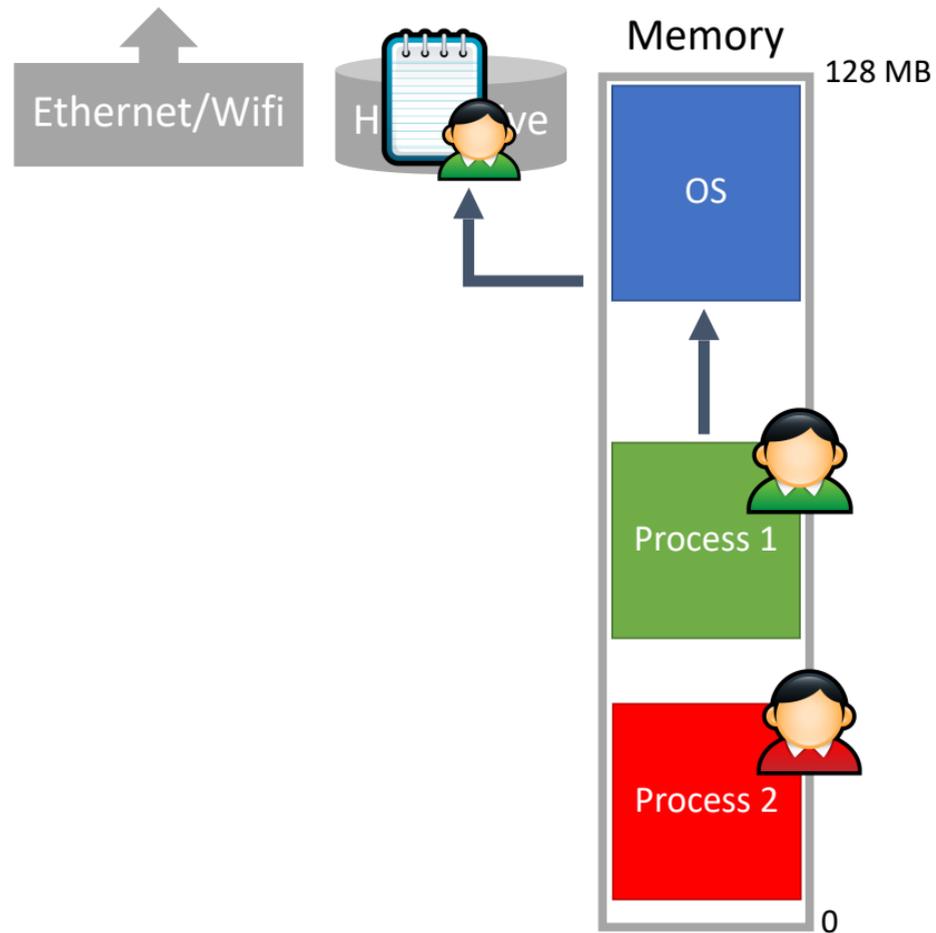
Problem: any process can access any hardware device directly

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



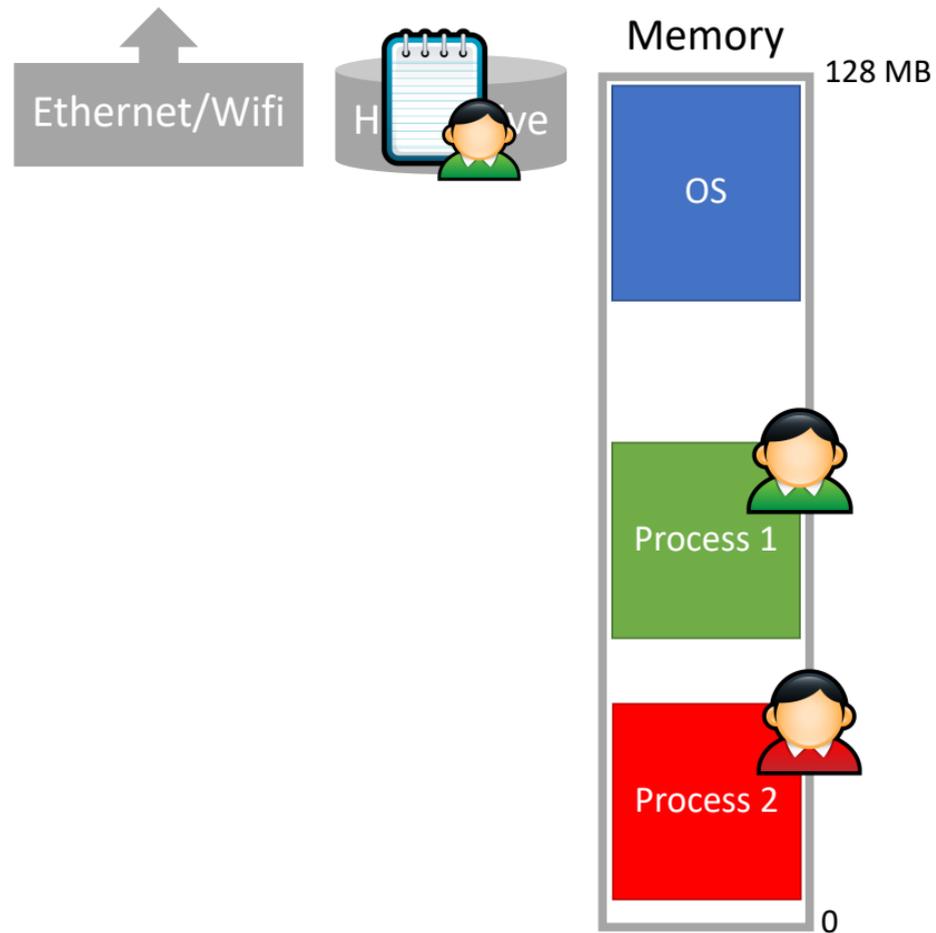
# Device Unsafety

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



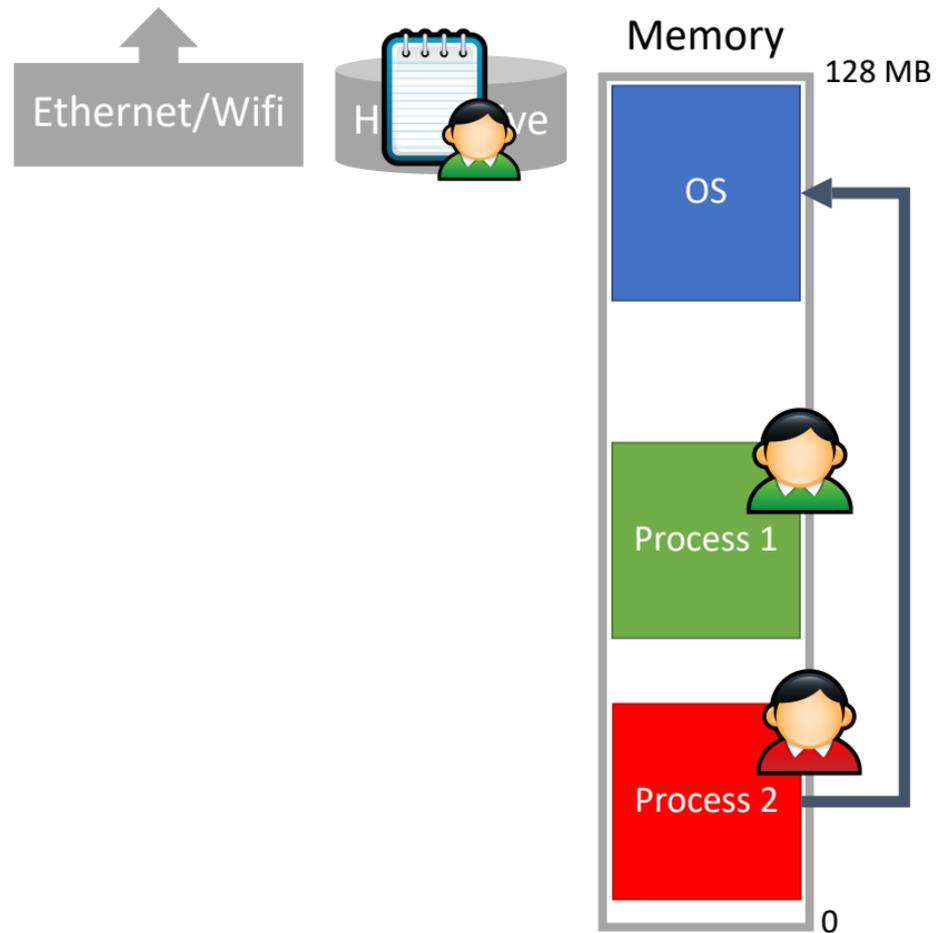
# Device Unsafety

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



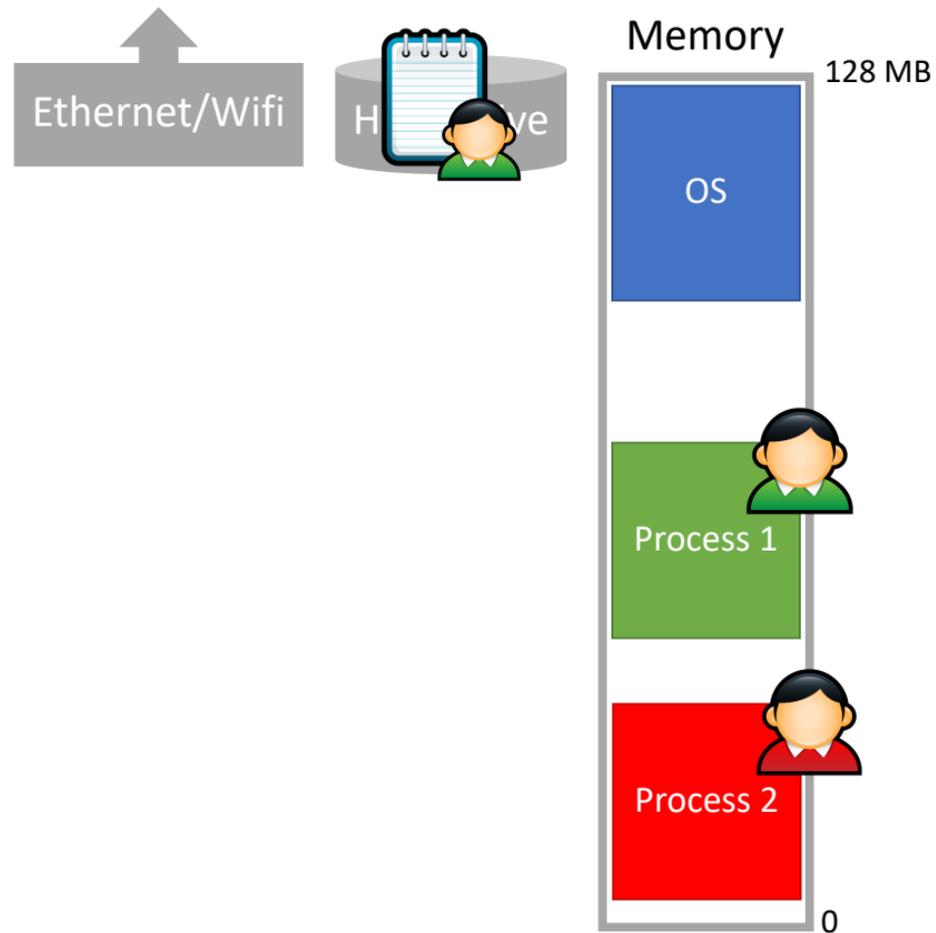
# Device Unsafety

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



# Device Unsafety

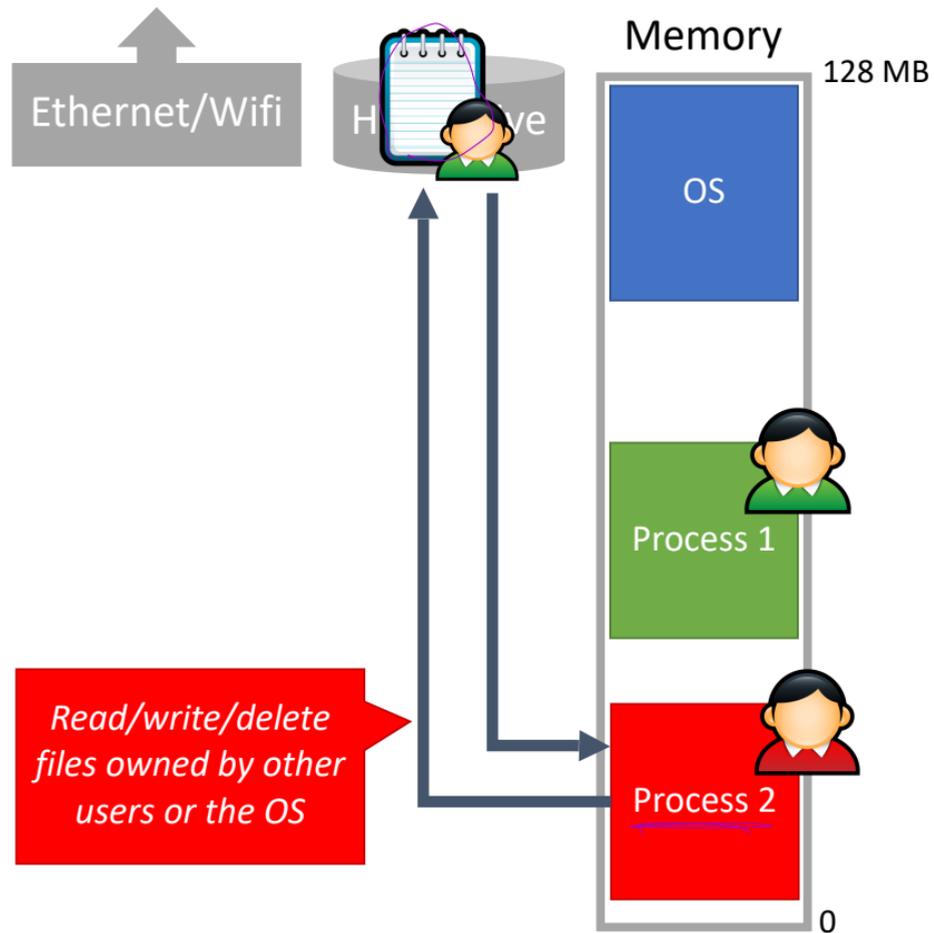
Problem: any process can access any hardware device directly  
Access control is enforced by the OS, but OS APIs can be bypassed



# Device Unsafety

Problem: any process can access any hardware device directly

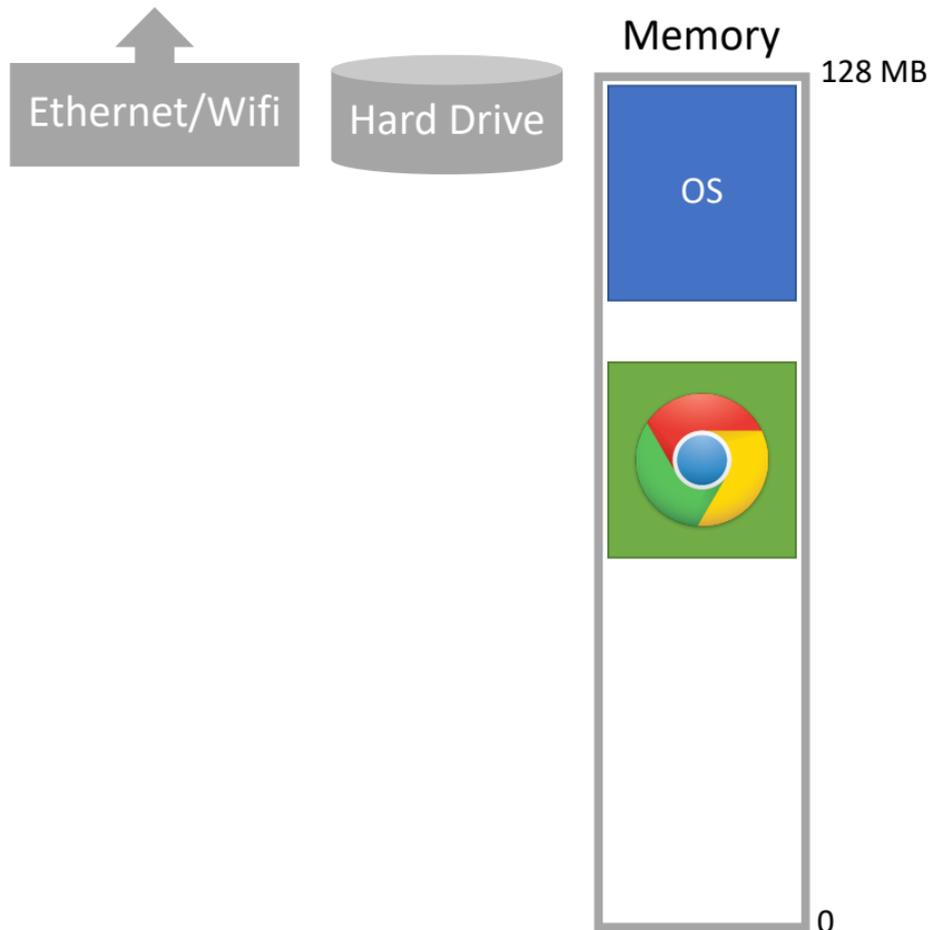
Access control is enforced by the OS, but OS APIs can be bypassed



# Device Unsafety

Problem: any process can access any hardware device directly

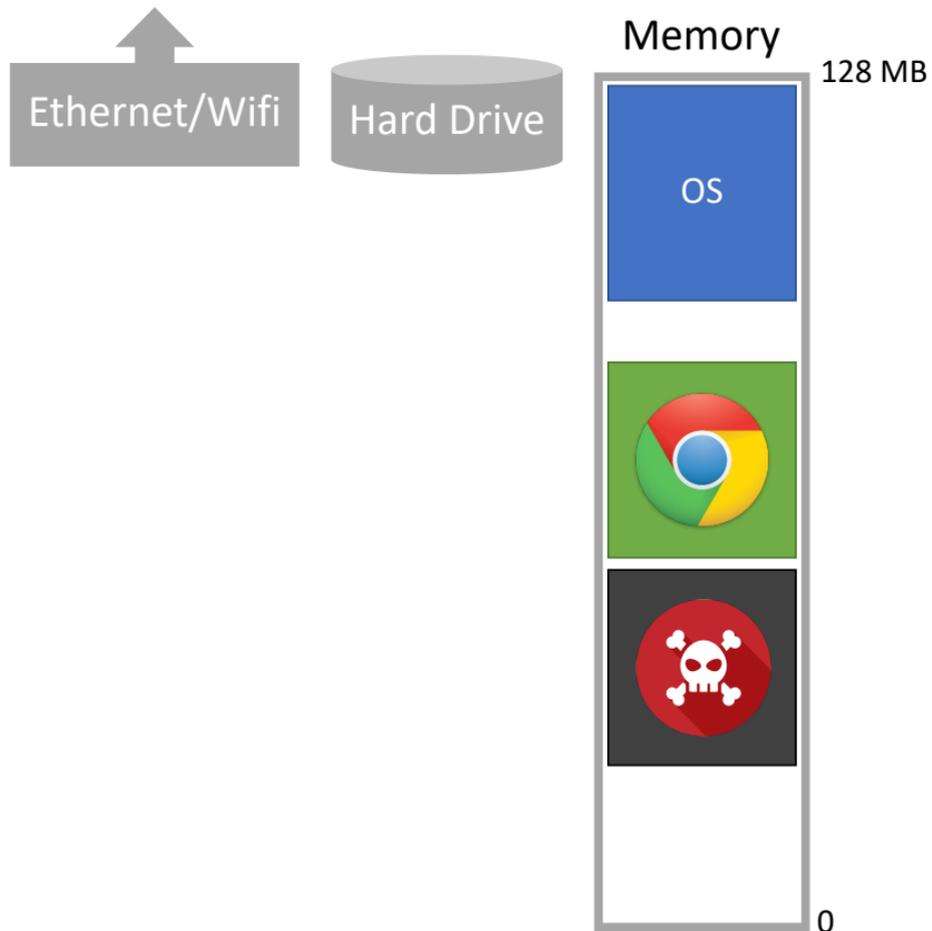
Access control is enforced by the OS, but OS APIs can be bypassed



# Device Unsafety

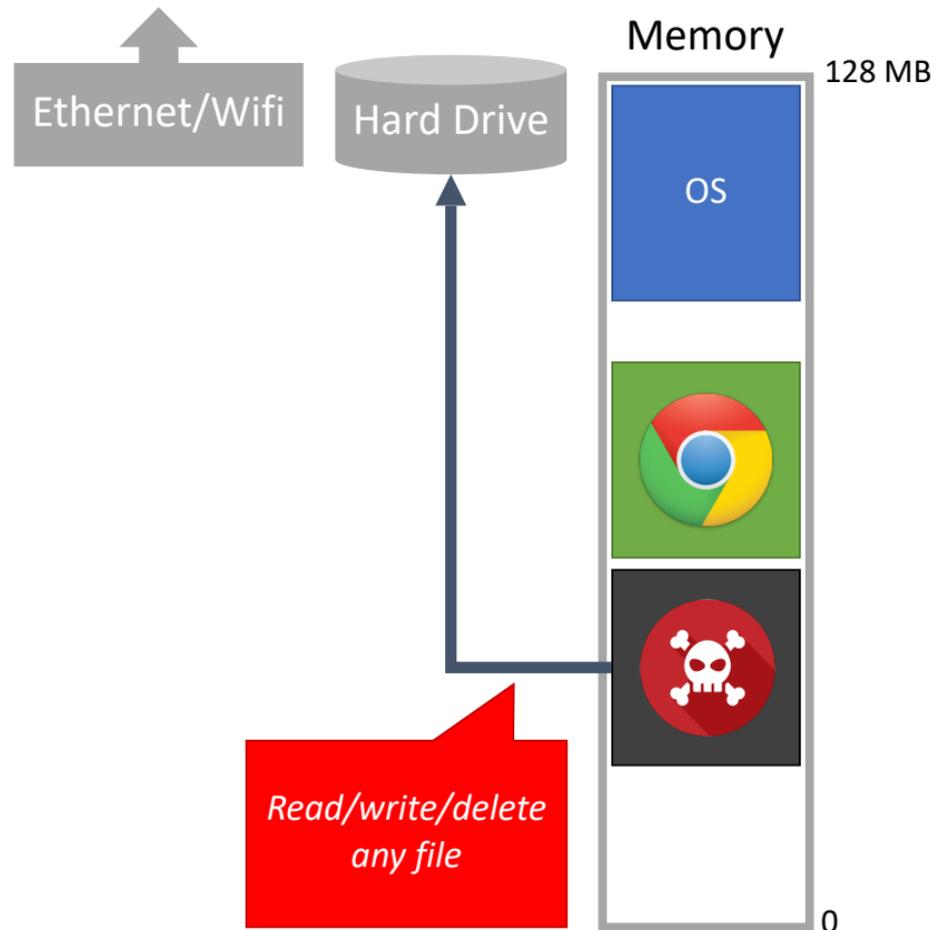
Problem: any process can access any hardware device directly

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



# Device Unsafety

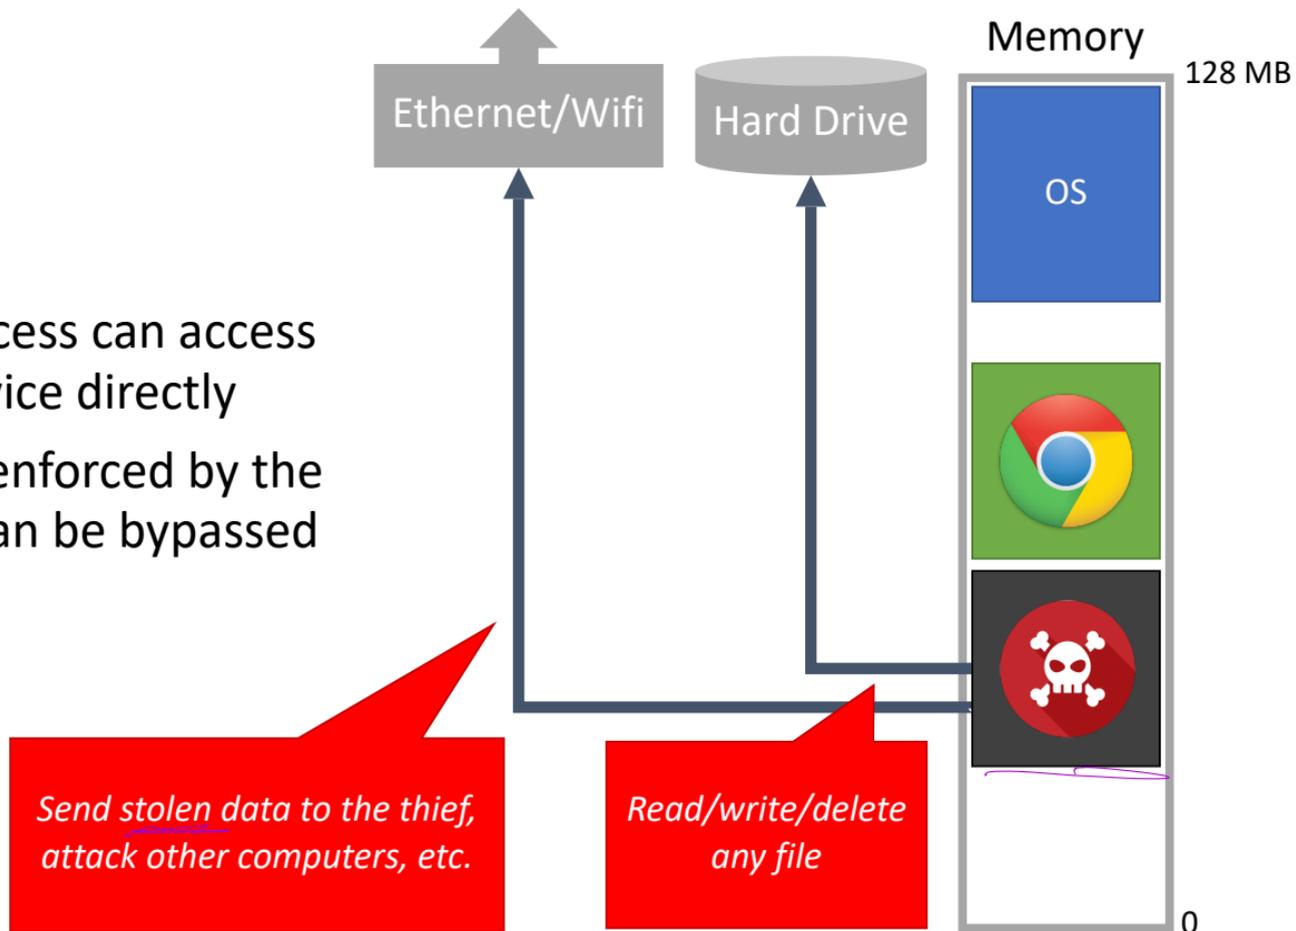
Problem: any process can access any hardware device directly  
Access control is enforced by the OS, but OS APIs can be bypassed



# Device Unsafety

Problem: any process can access any hardware device directly

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



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

*Read/write/delete any file*

# 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**

# ISOLATION



Threat Model

Principles

Intro to System Architecture

**Hardware Support for Isolation**

Examples



# 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

*Needed for process isolation.*



# 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

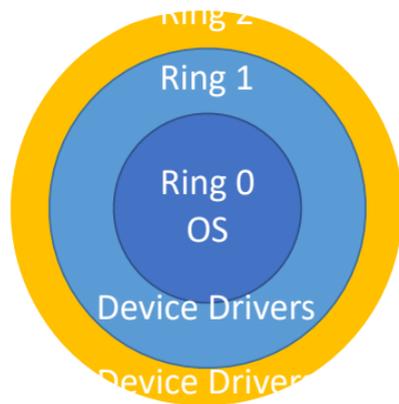


# Protected Mode

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

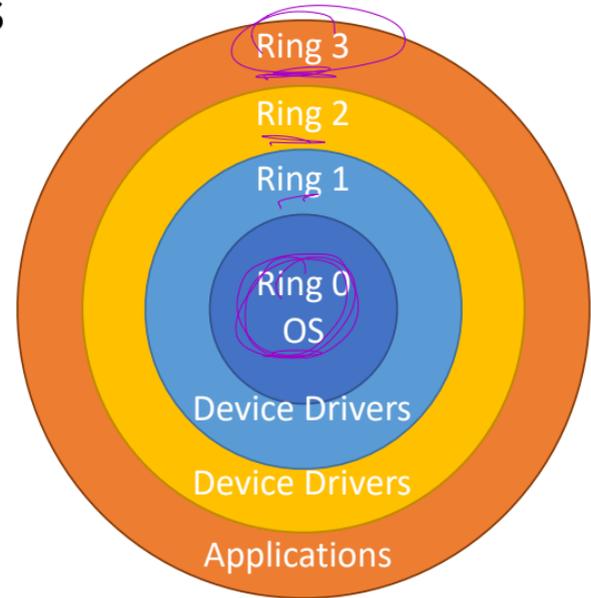


# Protected Mode

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



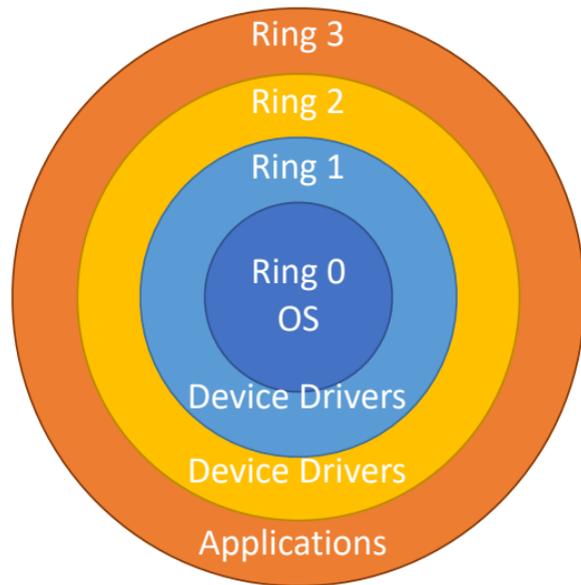
# Protected Mode

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



# 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

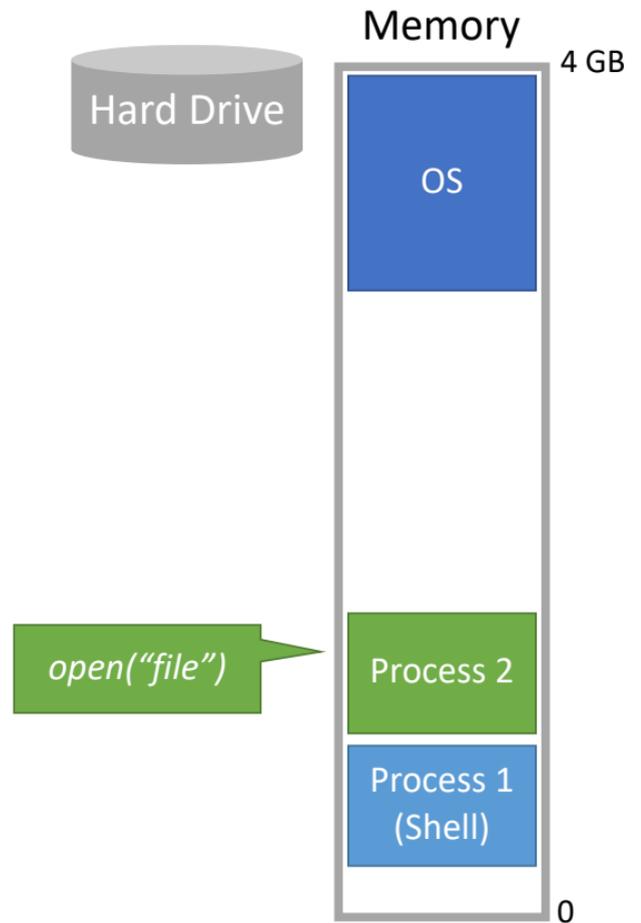
# Restriction on Privileged Instructions

What CPU instructions are restricted in protected mode?

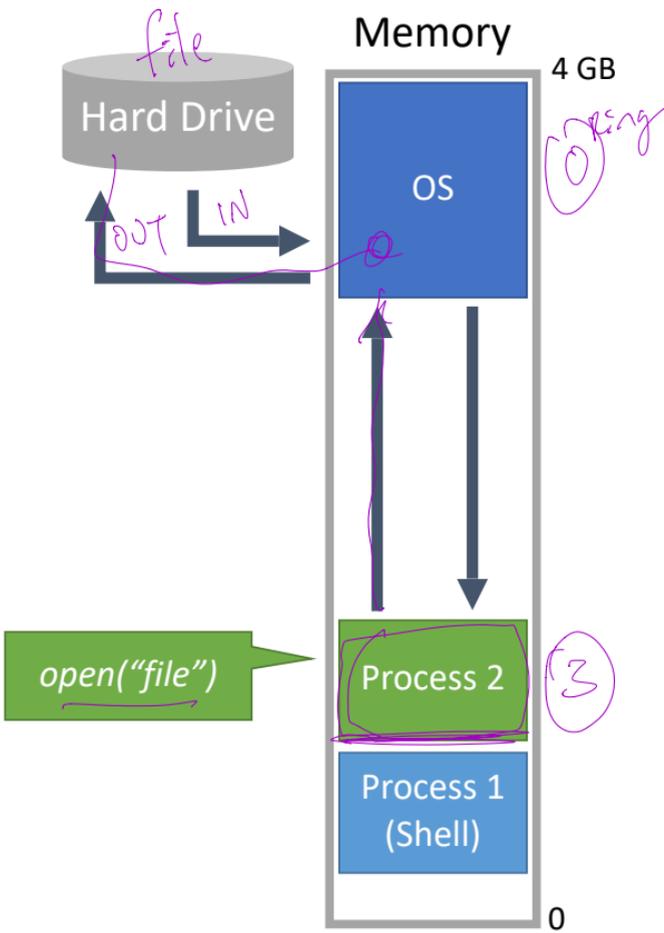
- Any instruction that modifies the CR0 register
  - Controls whether protected mode is enabled
- Any instruction that modifies the CR3 register
  - Controls the virtual memory configuration
  - More on this later...
- `hlt` – Halts the CPU
- `sti/cli` – enable and disable interrupts
- `in/out` – directly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

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

# How to change modes



# 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?

Interrupt handlers

# Changing Modes

Applications often need to access the OS APIs

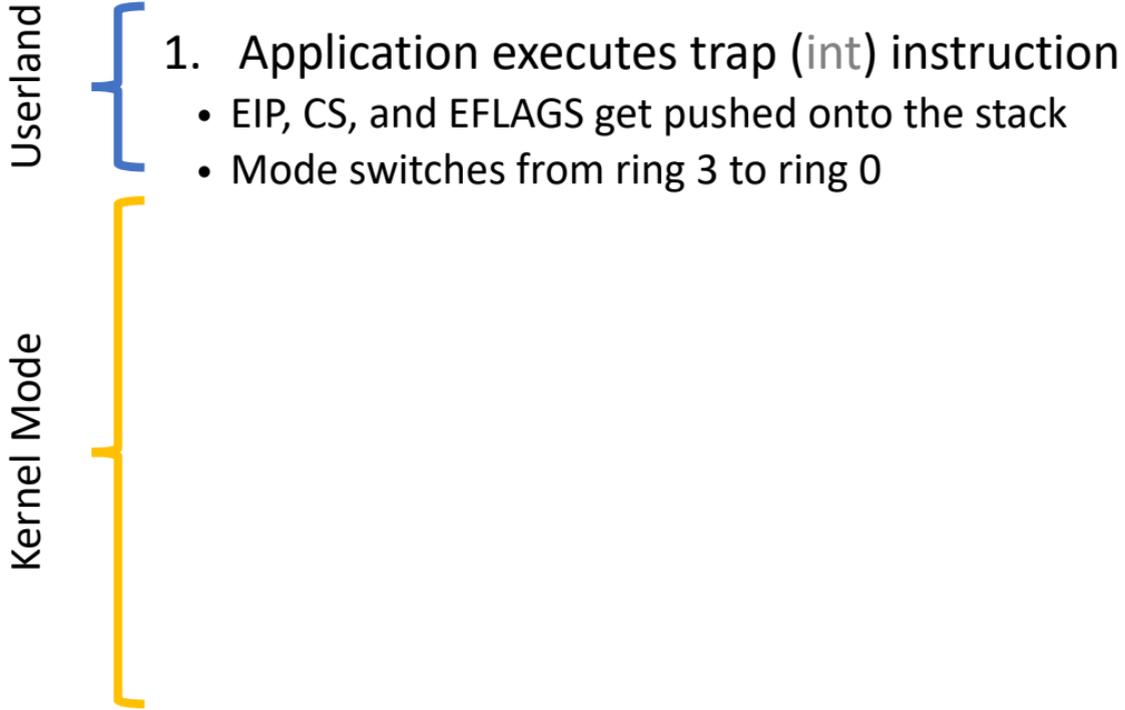
- Writing files
- Displaying things on the screen
- Receiving data from the network
- etc...

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

# Mode Transfer



# Mode Transfer

- 
- The diagram illustrates the process of mode transfer between Userland and Kernel Mode. It features two vertical labels on the left: 'Userland' and 'Kernel Mode'. A blue bracket on the left groups the first step, which occurs in Userland. A yellow bracket on the left 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

- 
- 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 groups the first step, which occurs in Userland. A yellow bracket groups the remaining three steps, which occur 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
    - 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

# Mode Transfer

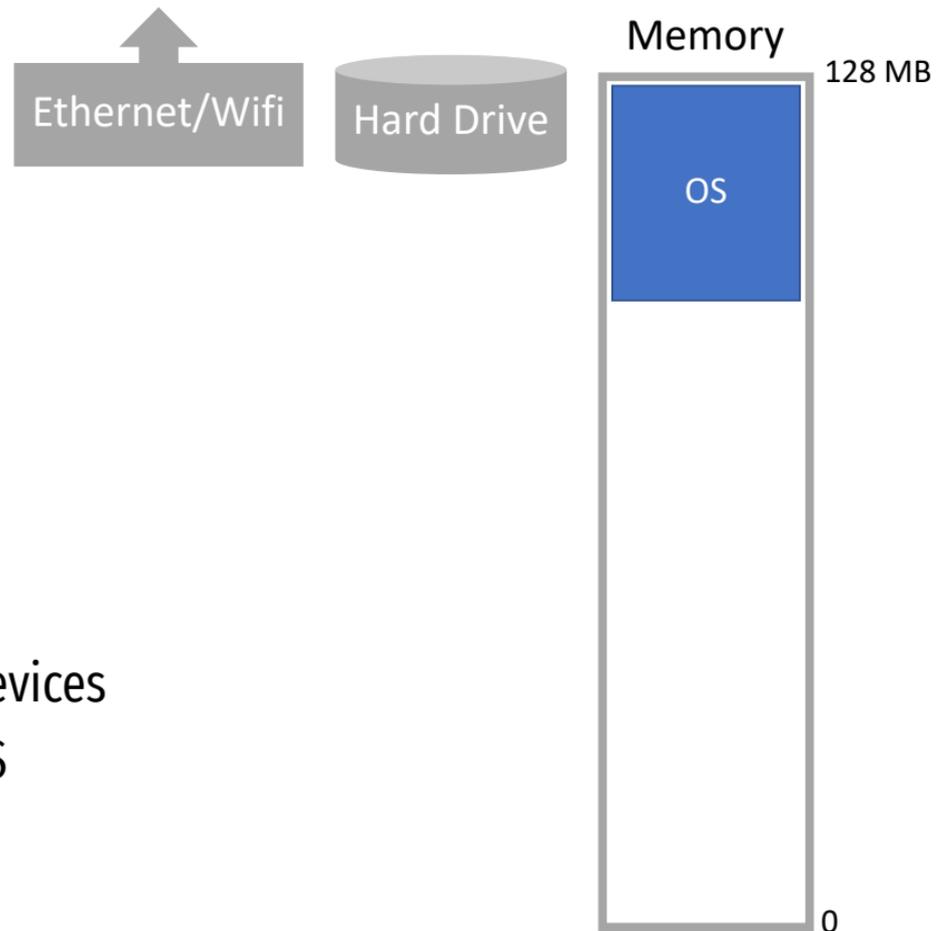
Userland

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

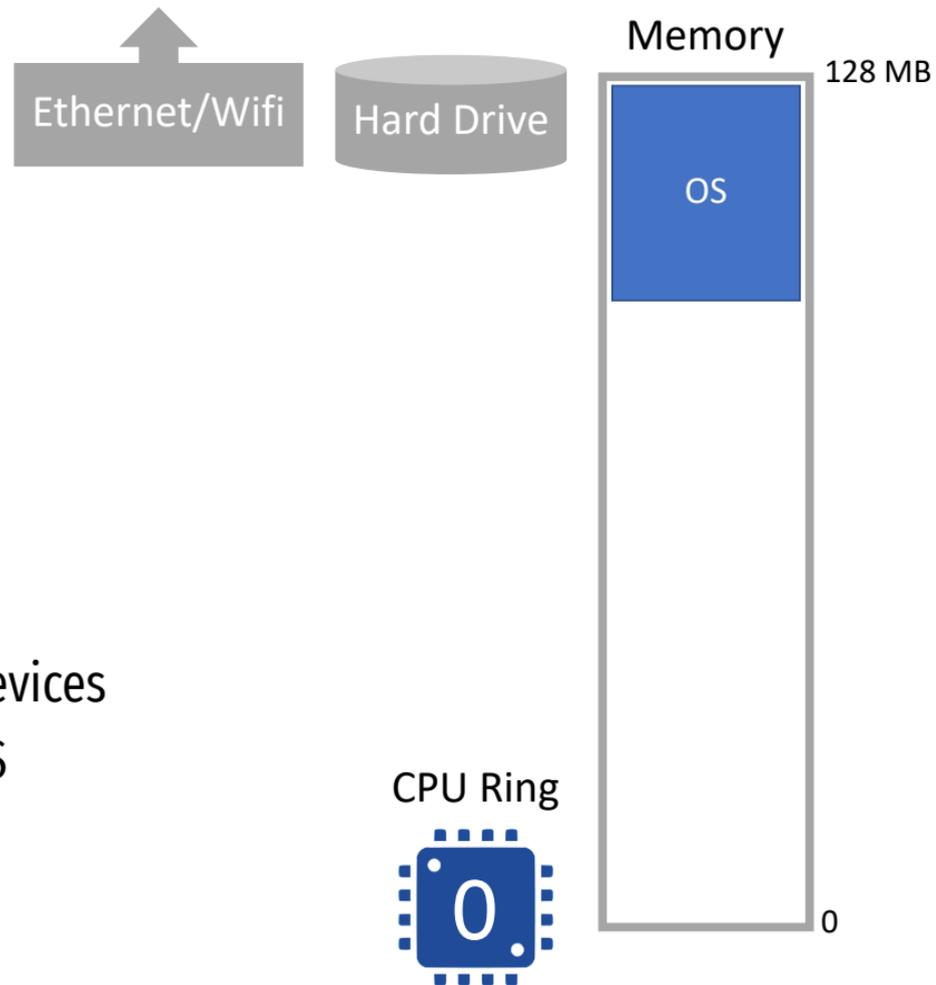


Protected mode stops direct access to devices

All device access must go through the OS

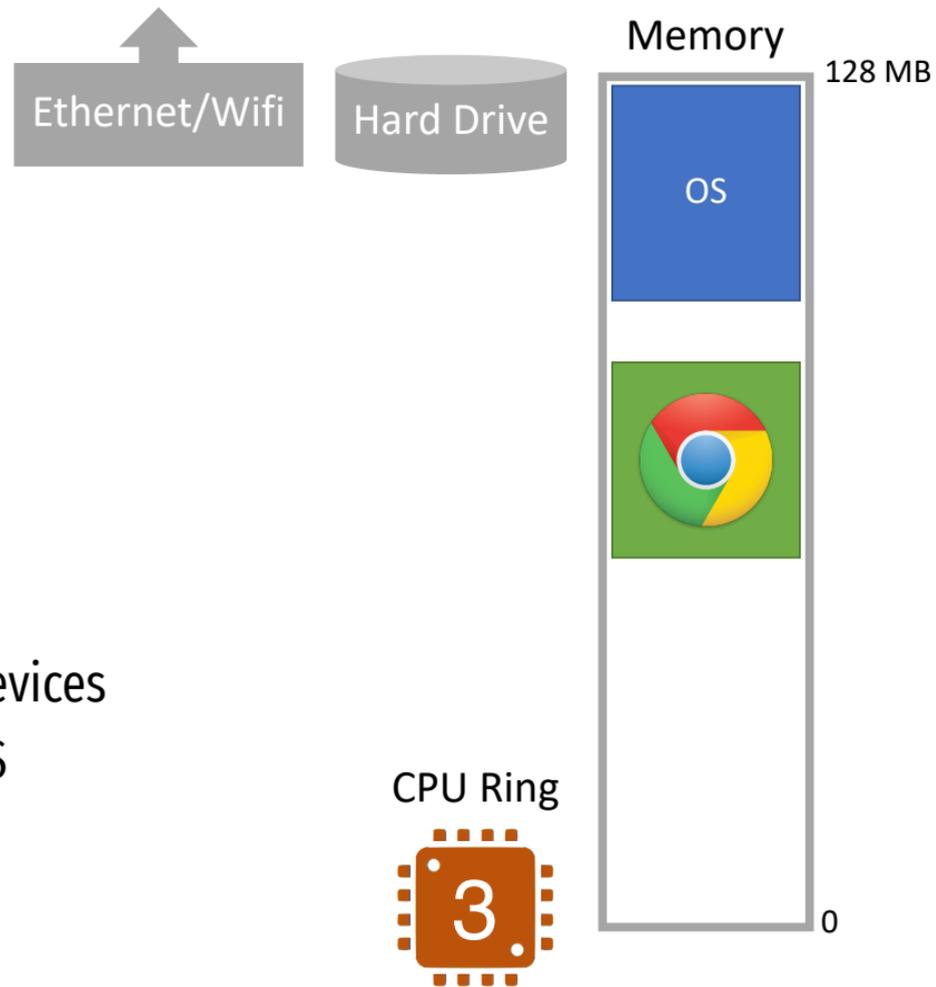
OS will impose access control checks

# Protection in Action



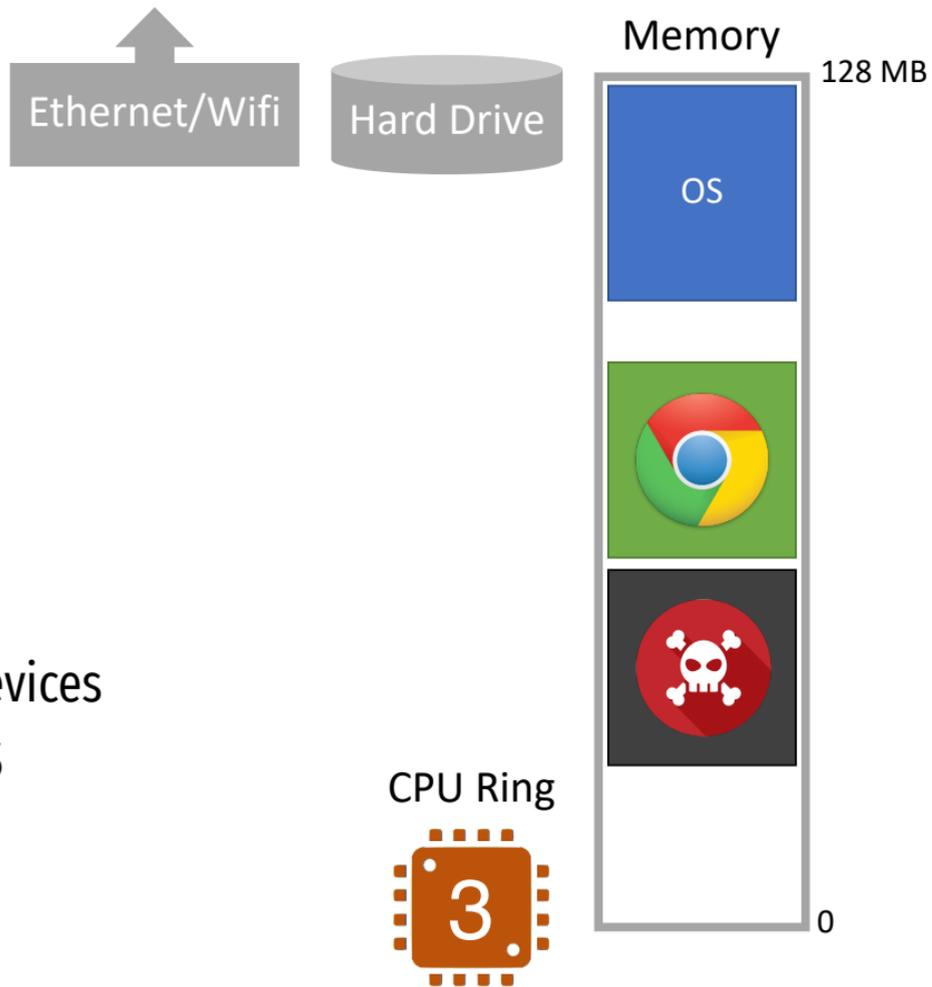
Protected mode stops direct access to devices  
All device access must go through the OS  
OS will impose access control checks

# 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



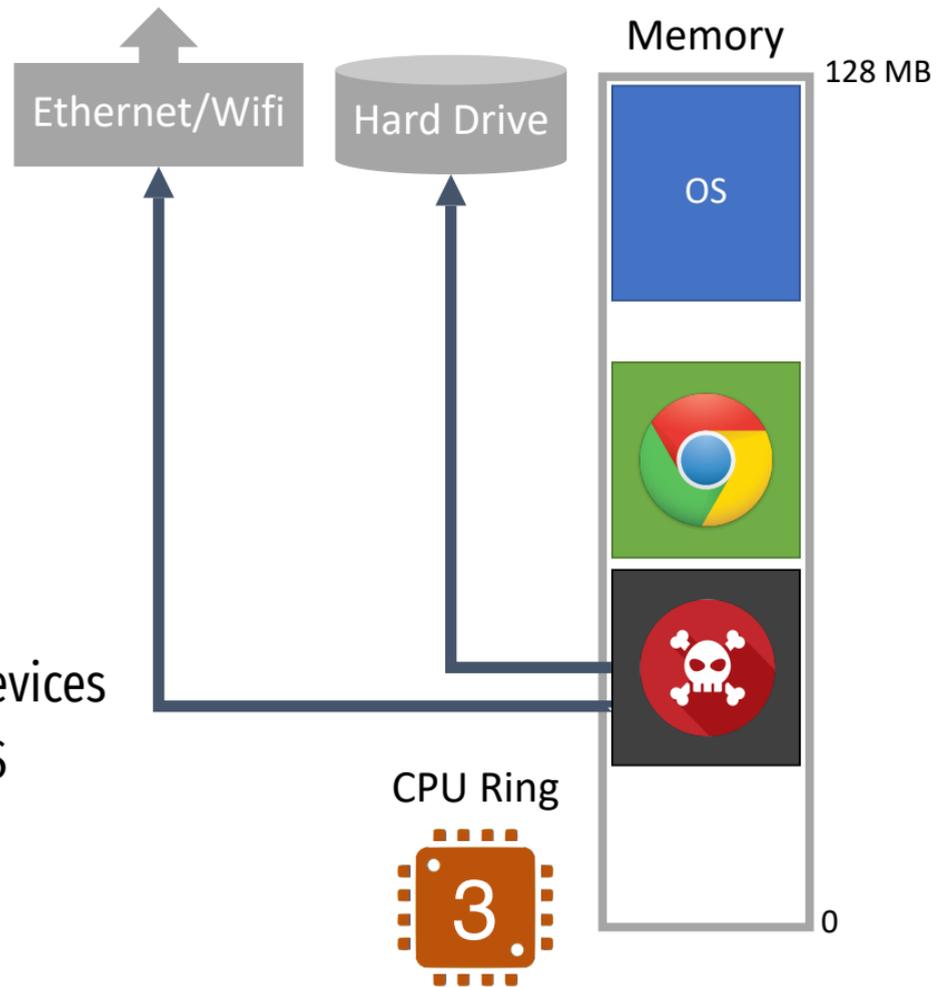
Protected mode stops direct access to devices

All device access must go through the OS

OS will impose access control checks

# Protection in Action

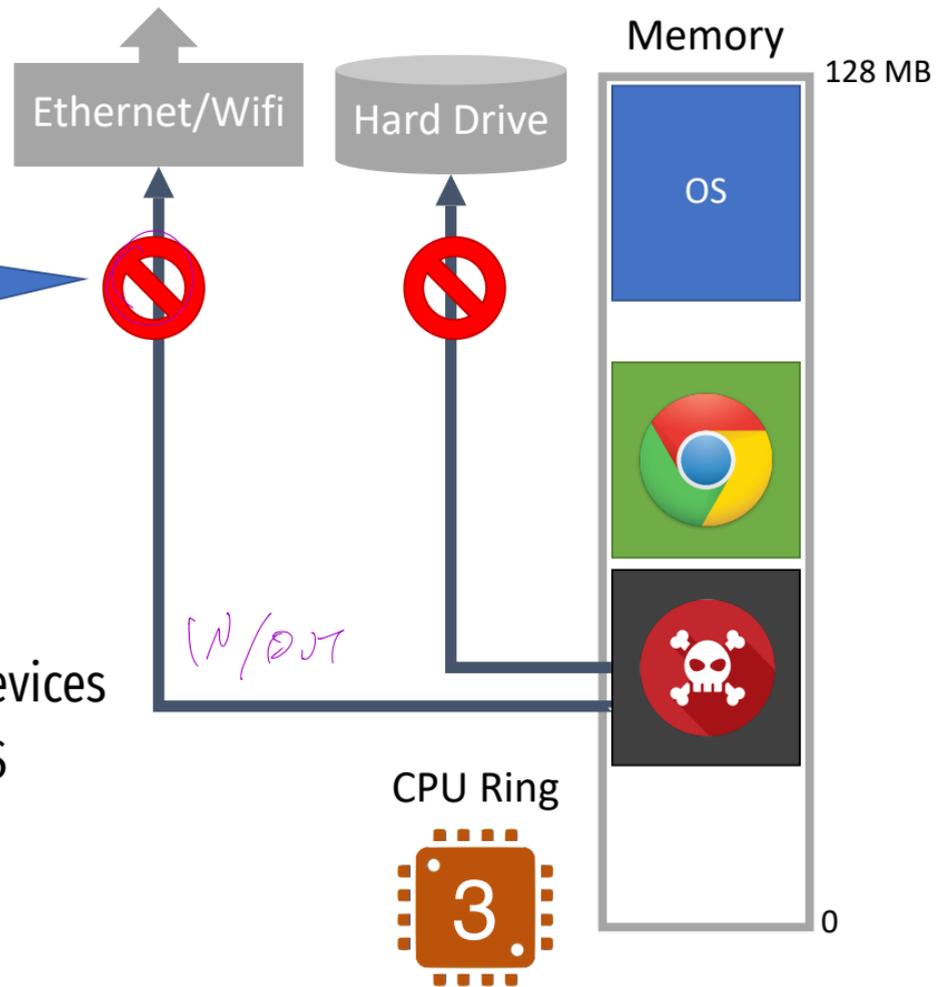
Protected mode stops direct access to devices  
All device access must go through the OS  
OS will impose access control checks



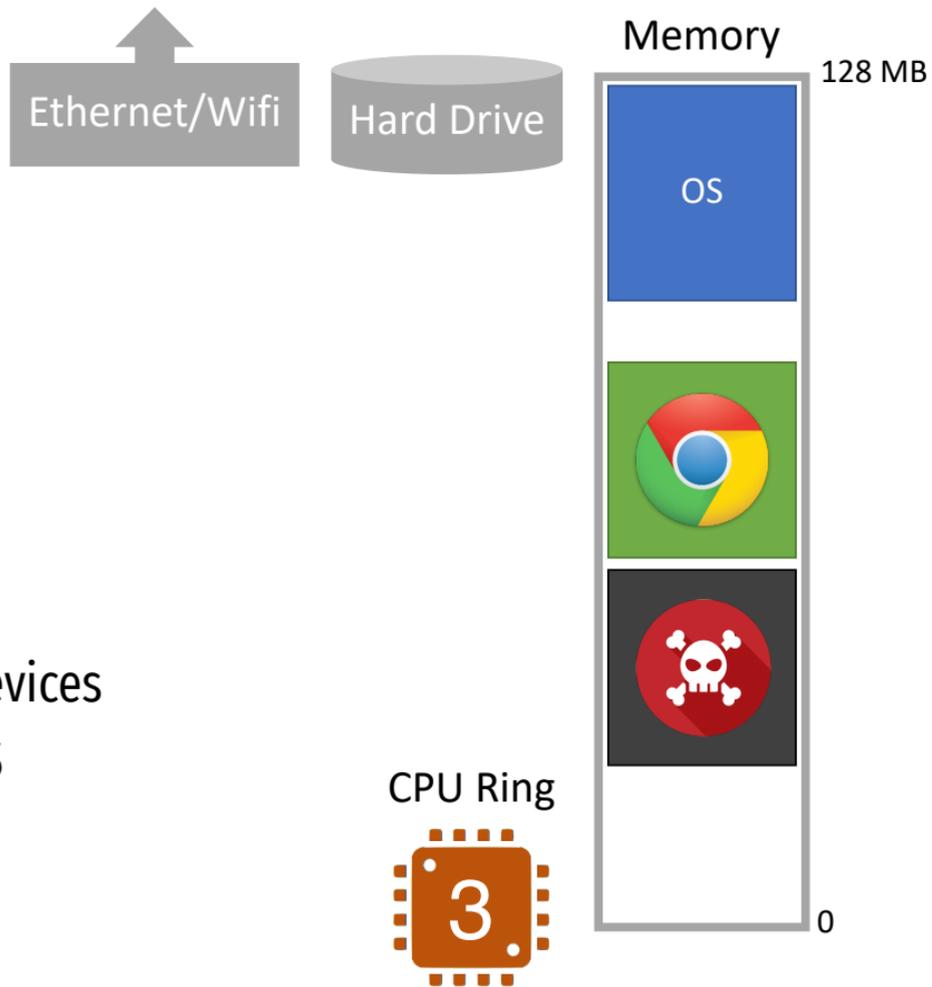
# Protection in Action

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



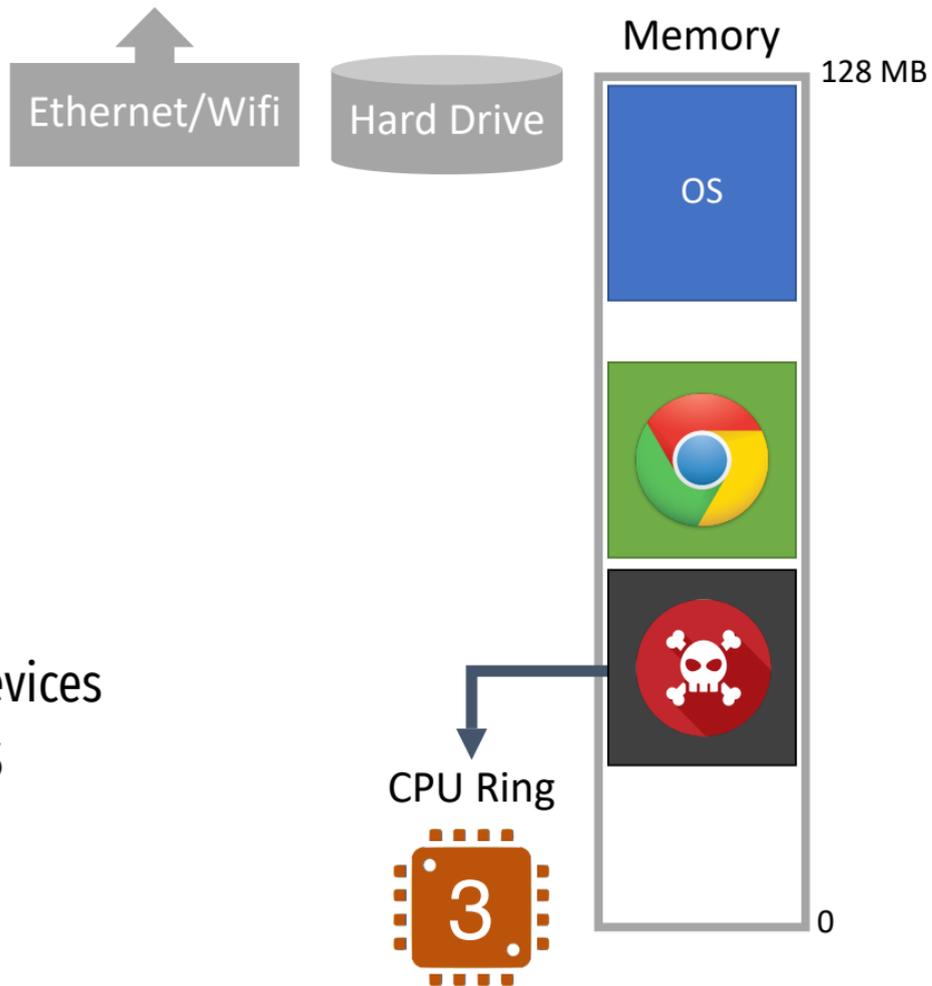
Protected mode stops direct access to devices

All device access must go through the OS

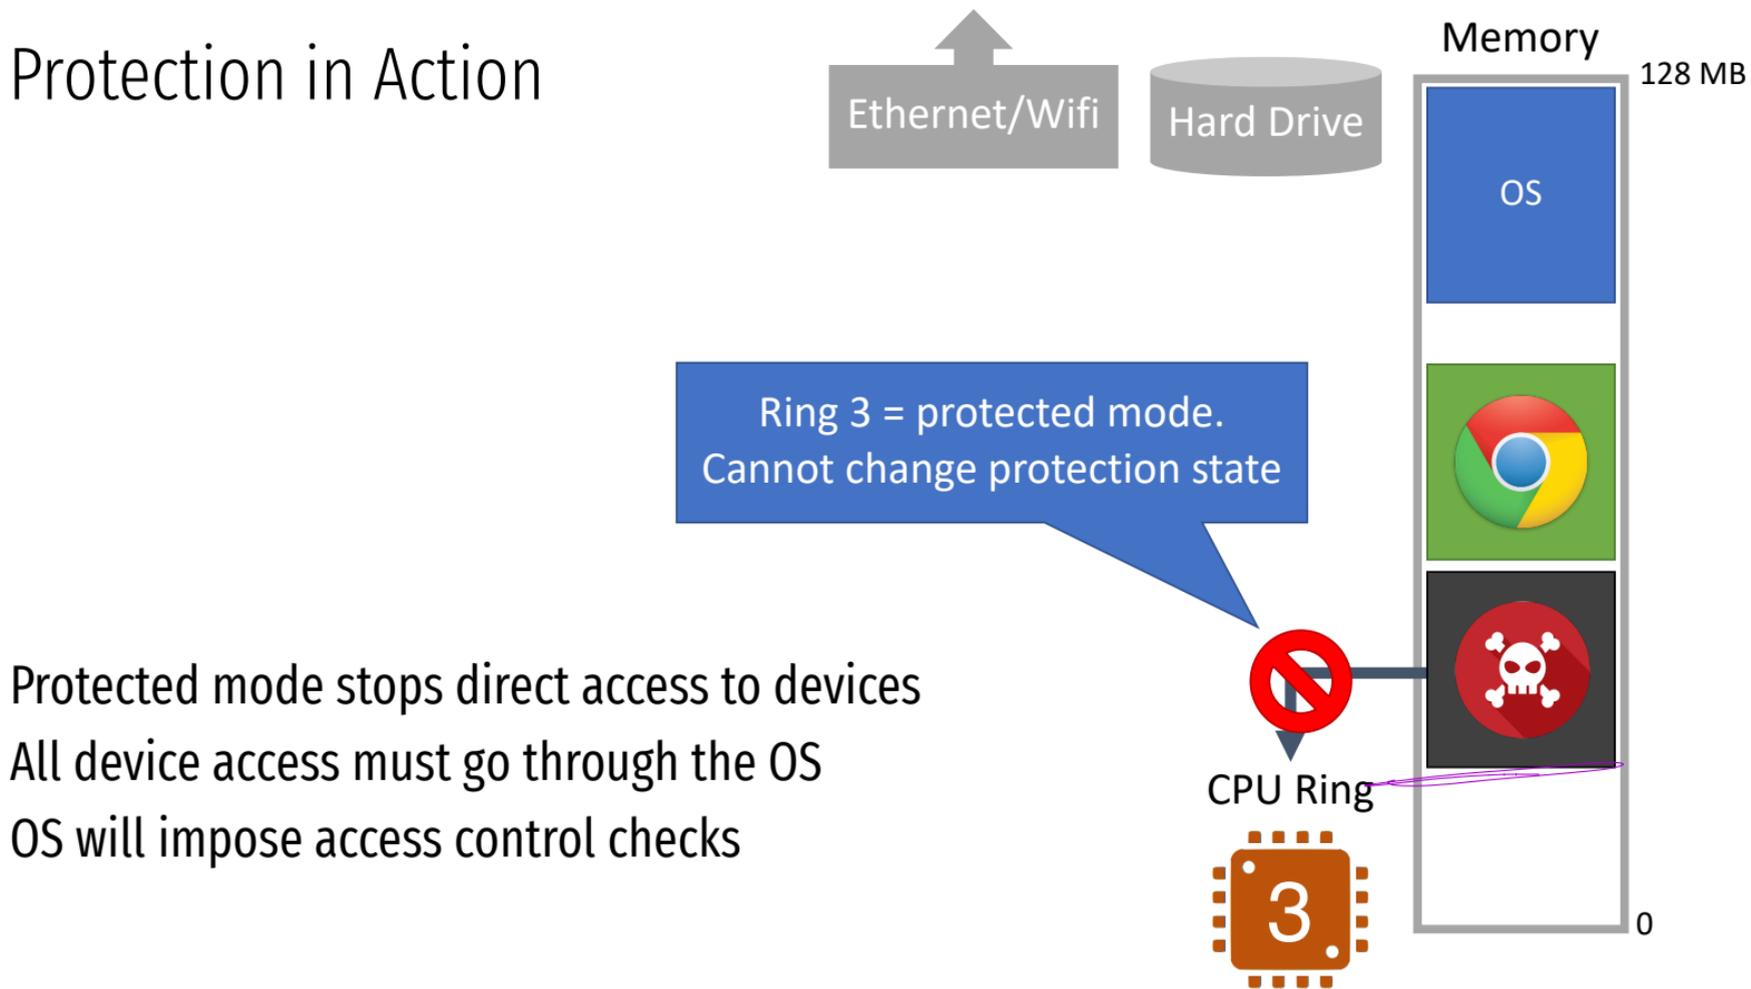
OS will impose access control checks

# Protection in Action

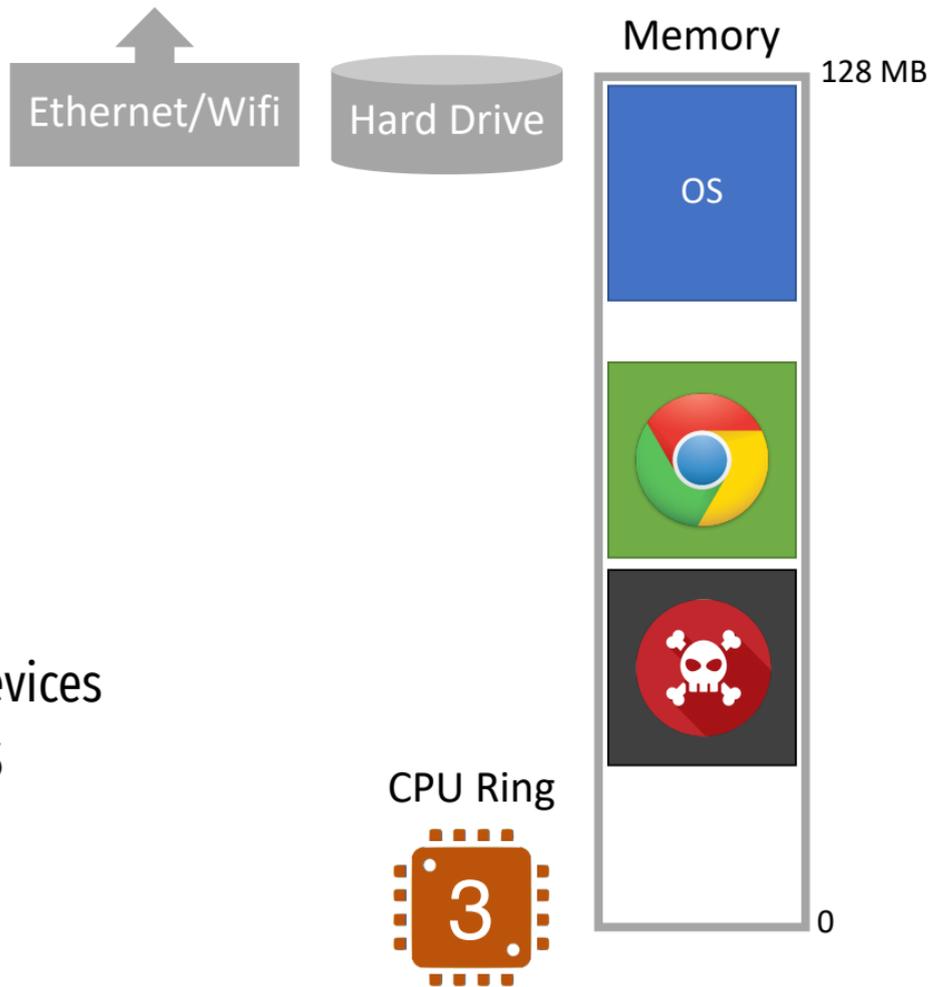
Protected mode stops direct access to devices  
All device access must go through the OS  
OS will impose access control checks



# Protection in Action



# Protection in Action



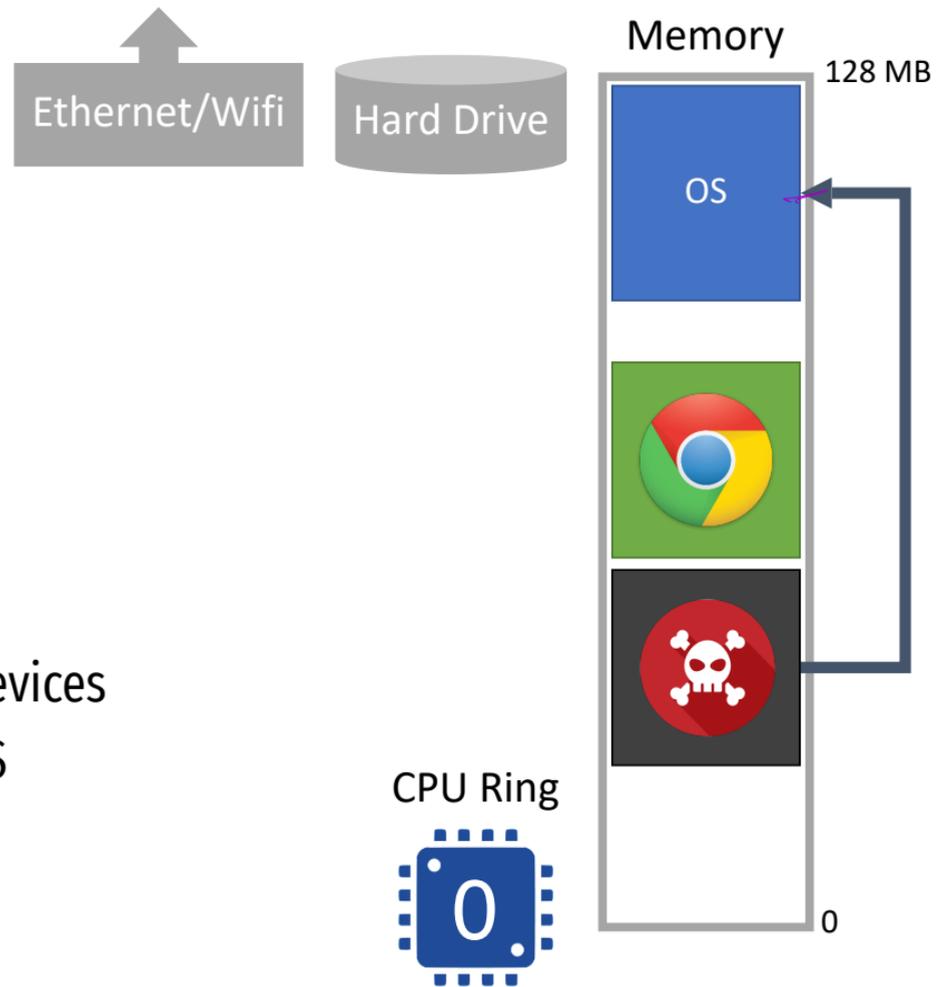
Protected mode stops direct access to devices

All device access must go through the OS

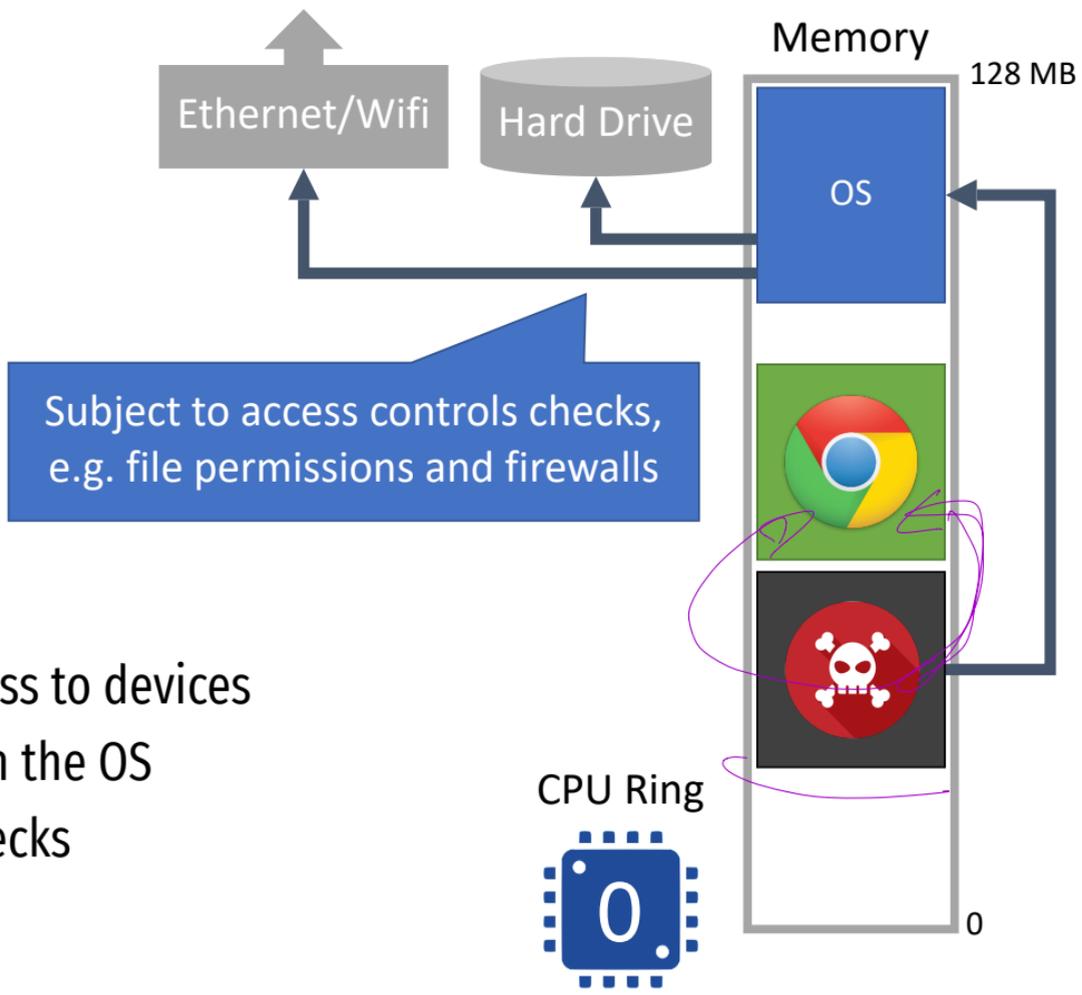
OS will impose access control checks

# 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

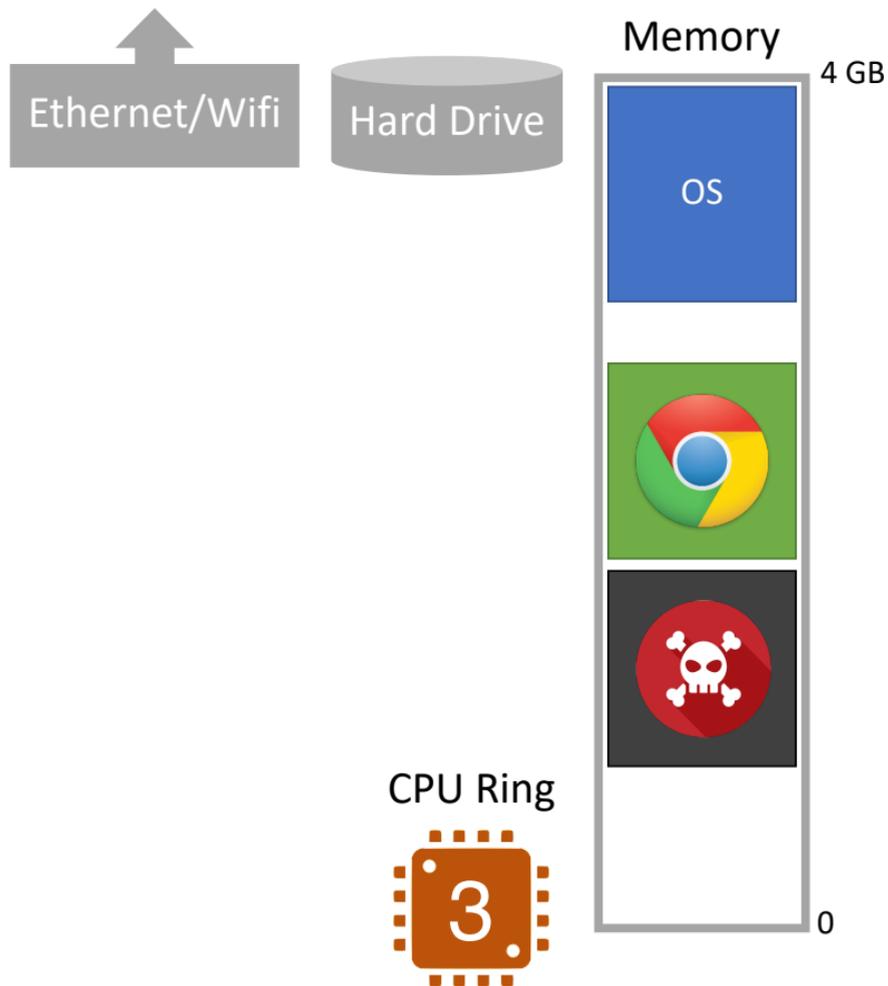
# Virtual Memory

A hand-drawn purple underline consisting of two slightly wavy lines that extend to the right, underlining the text 'Virtual Memory'.

# Status Check

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

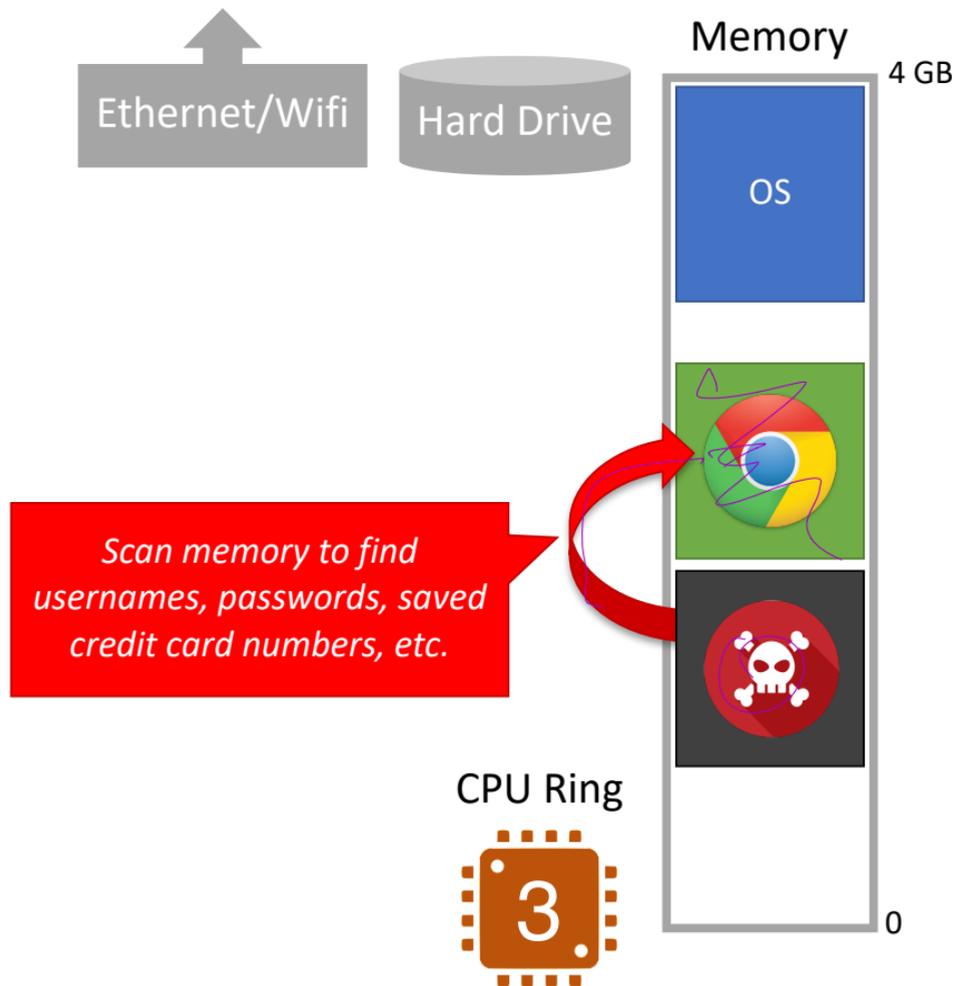
... But we have not protected memory



# Status Check

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

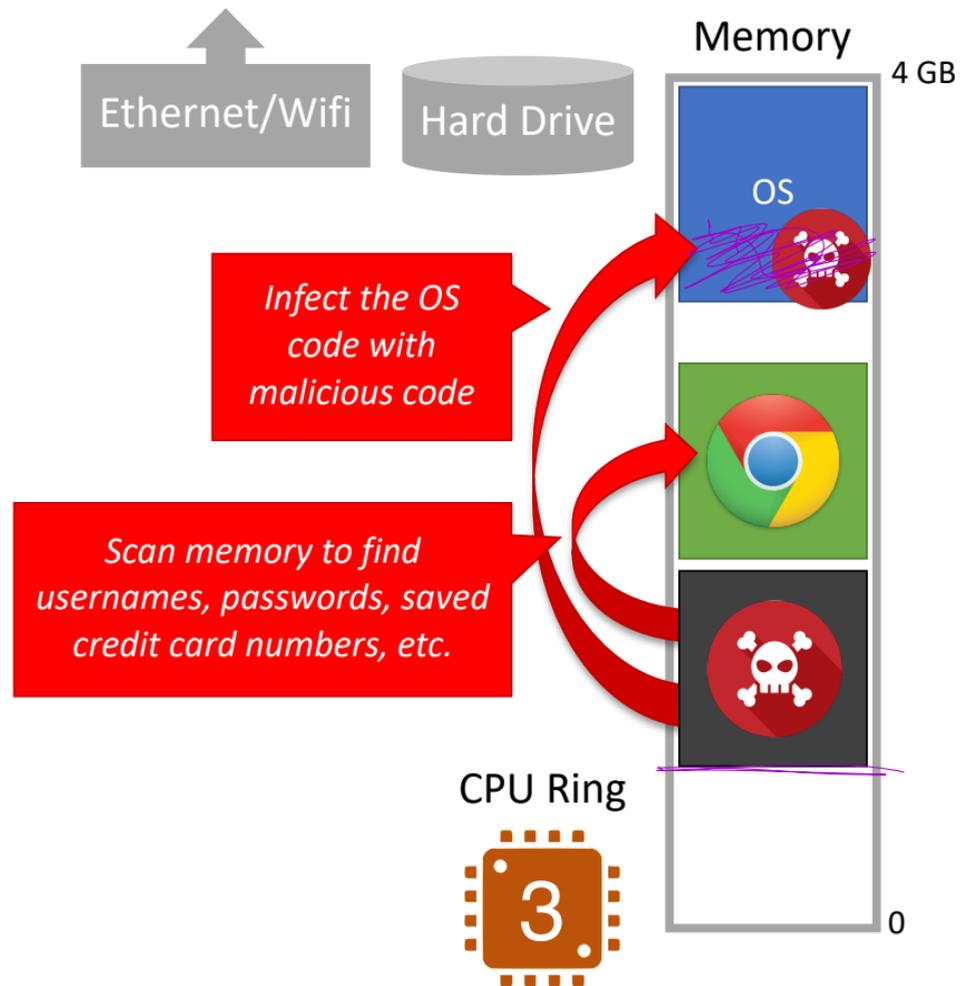
... But we have not protected memory



# 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

4 GB



OS

0

# Physical Memory

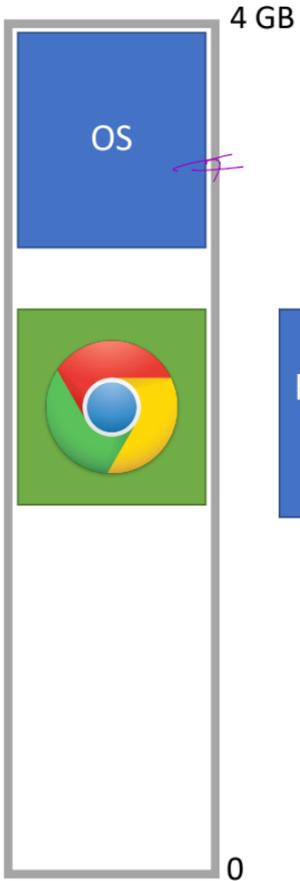
4 GB

OS



0

# Physical Memory



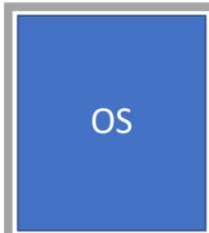
# Virtual Memory Process 1



Chrome believes it is the only thing in memory

# Physical Memory

4 GB



0

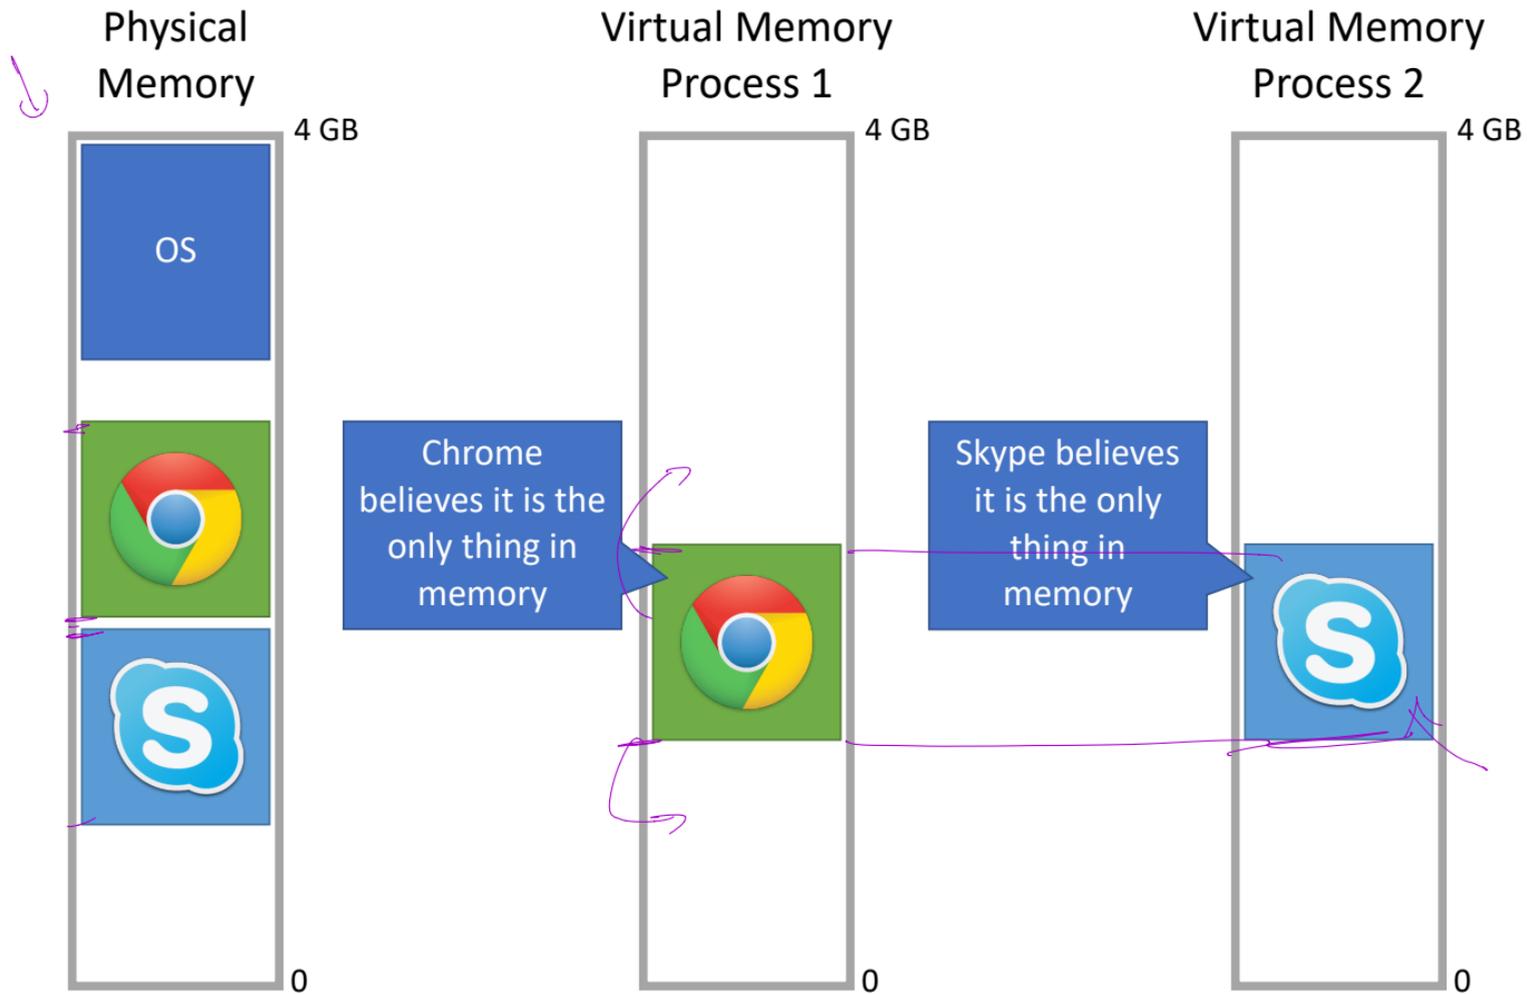
# Virtual Memory Process 1

4 GB



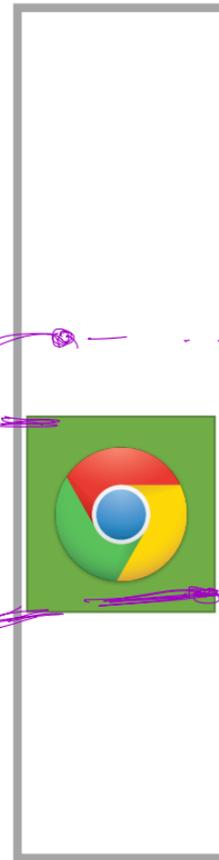
0

Chrome believes it is the only thing in memory



Virtual Memory  
Process 1

4 GB

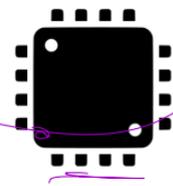


INVALID

VMM

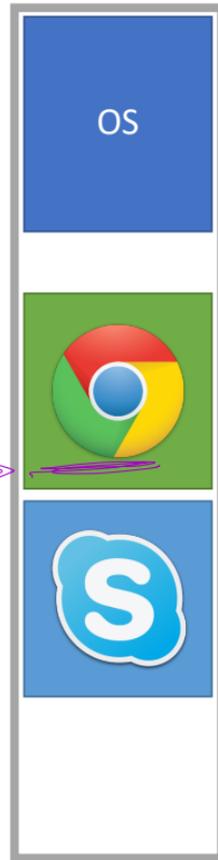
- VM addresses need to be mapped to physical addresses

CPU



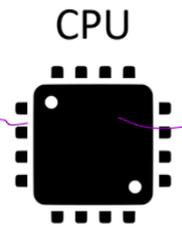
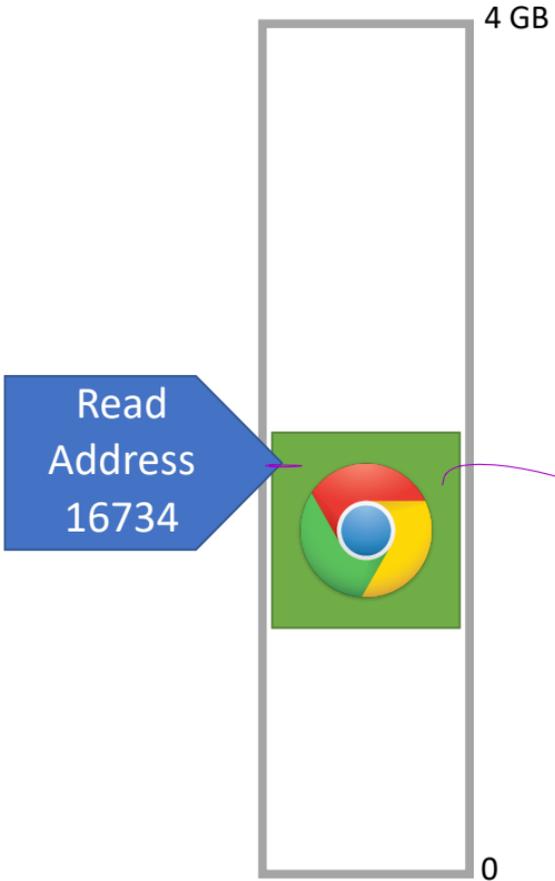
Physical Memory

4 GB



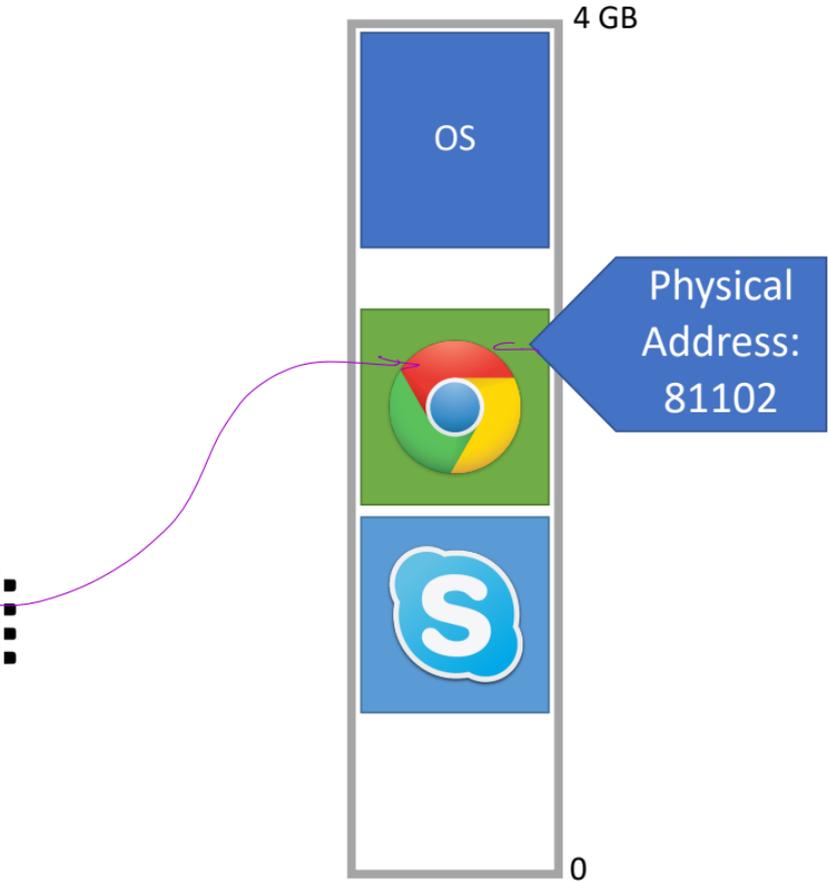
# Virtual Memory

## Process 1



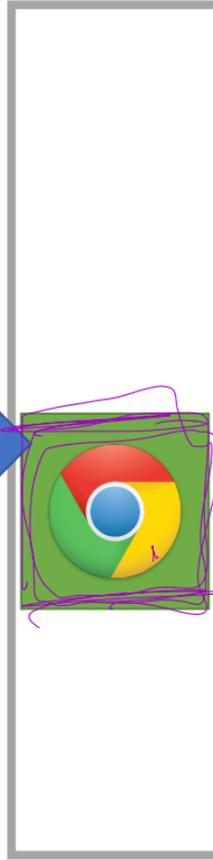
# Physical Memory

## Physical Memory



# Virtual Memory Process 1

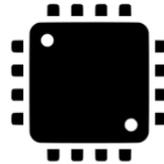
4 GB



### 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

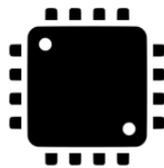


0

Page Table

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

CPU



Physical Memory

4 GB



0

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

• OS manages the page table to ensure process isolation.

# 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?

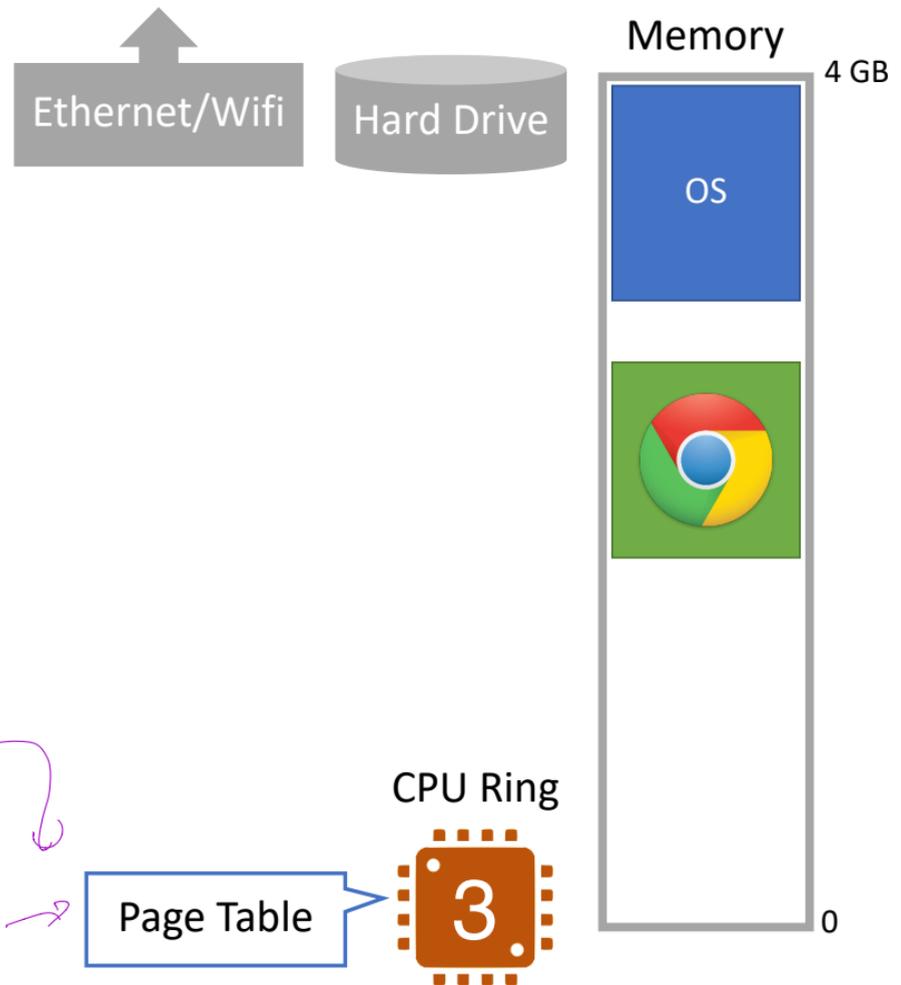
- Segmentation Fault or Page Fault
- Process crashes
- In other words, no way to escape virtual memory

# VM in Action

Processes can only read/  
write within their own  
virtual memory

Processes cannot change  
their own page tables

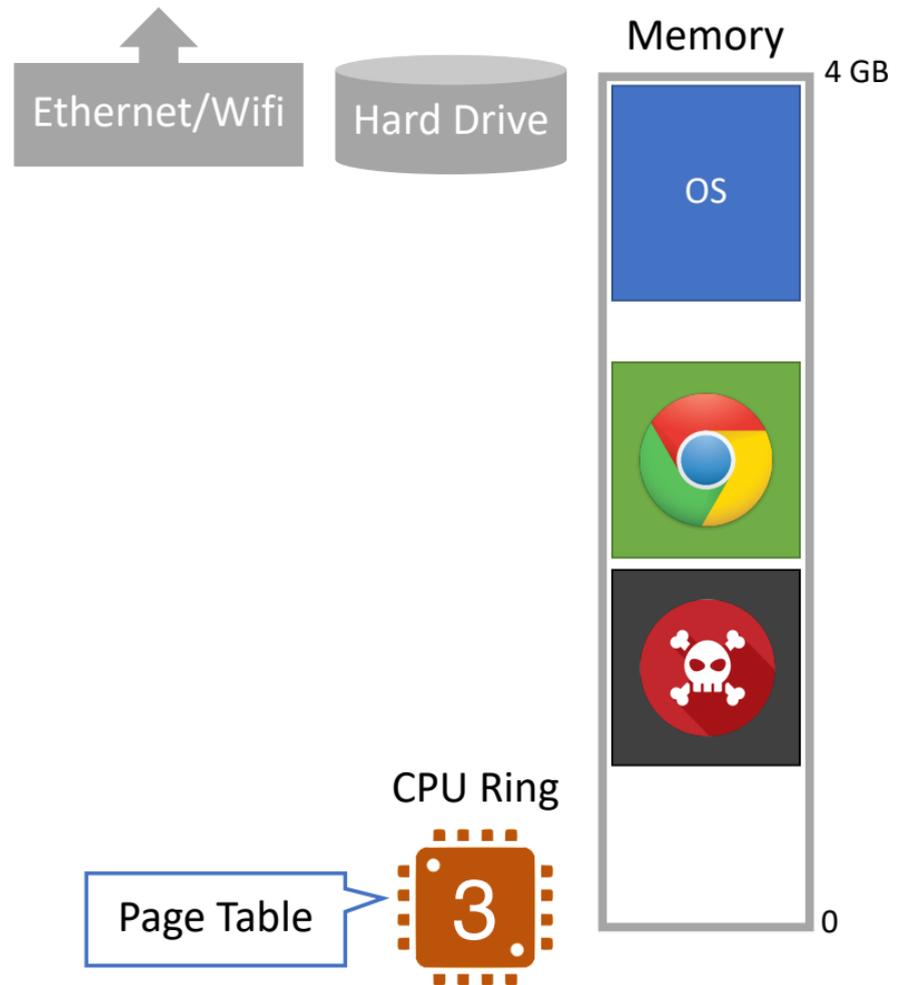
*Only Ring 0 can  
change the page table*



# VM in Action

Processes can only read/  
write within their own  
virtual memory

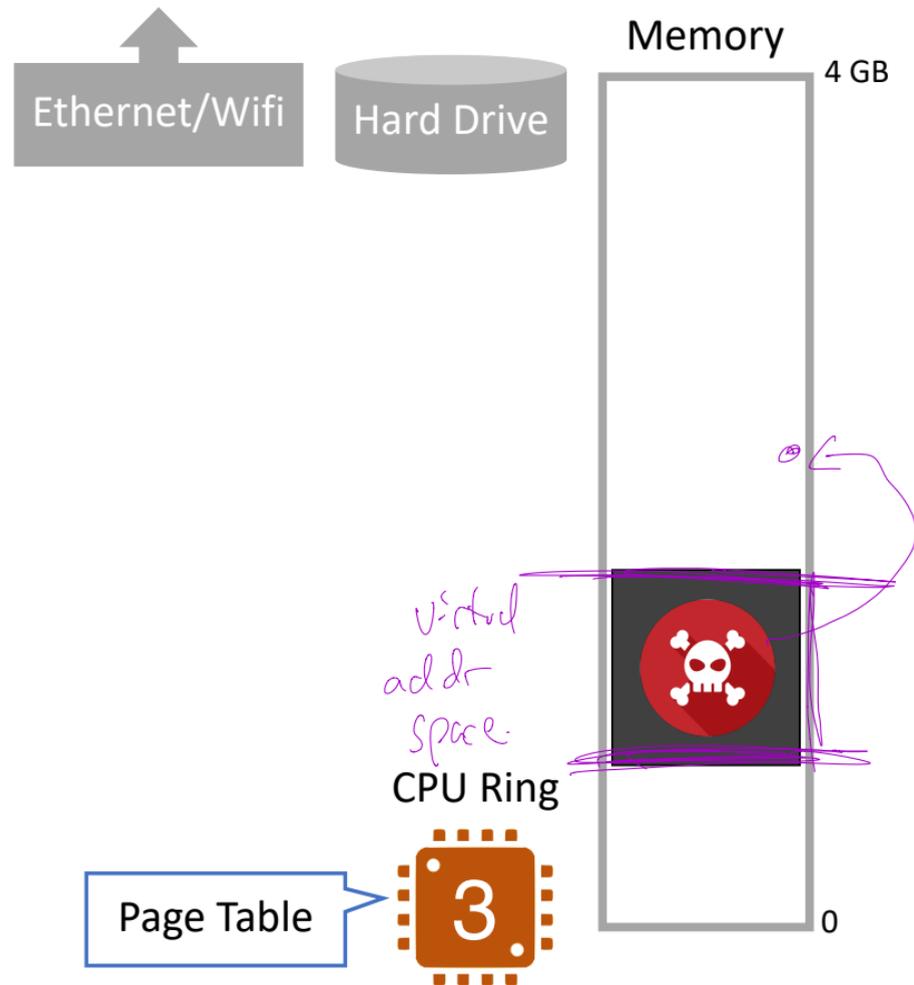
Processes cannot change  
their own page tables



# VM in Action

Processes can only read/  
write within their own  
virtual memory

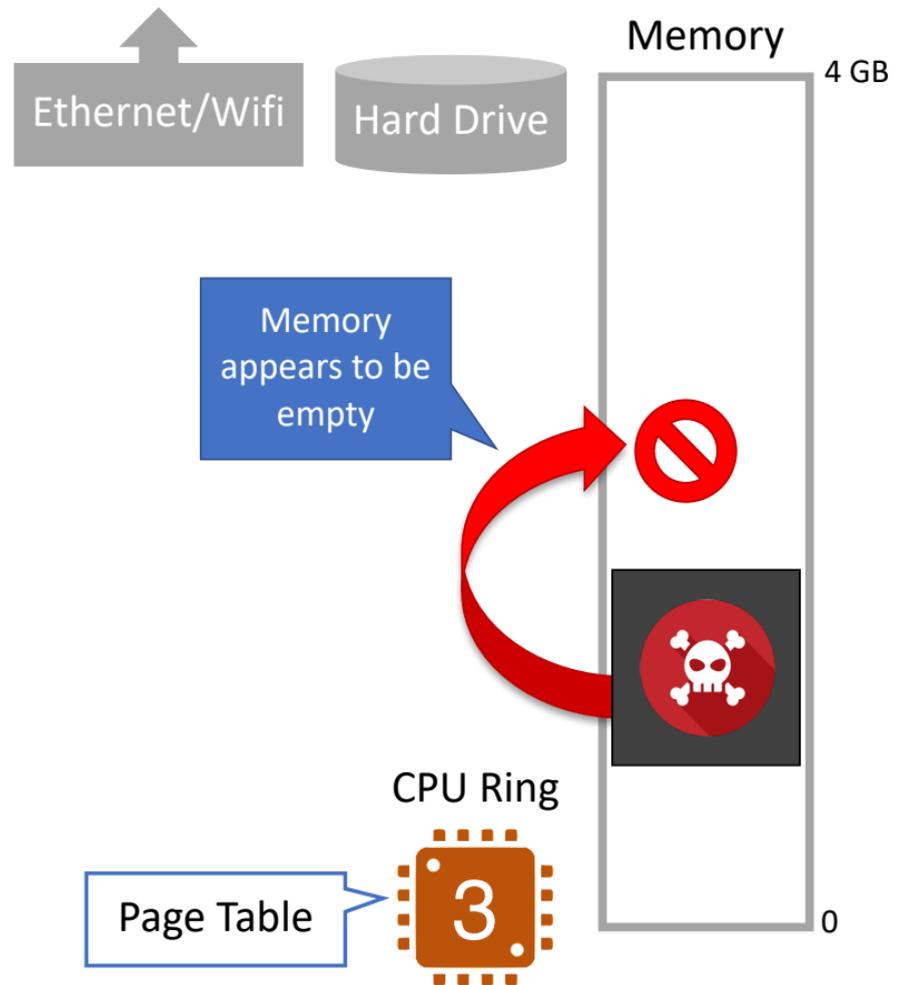
Processes cannot change  
their own page tables



# VM in Action

Processes can only read/  
write within their own  
virtual memory

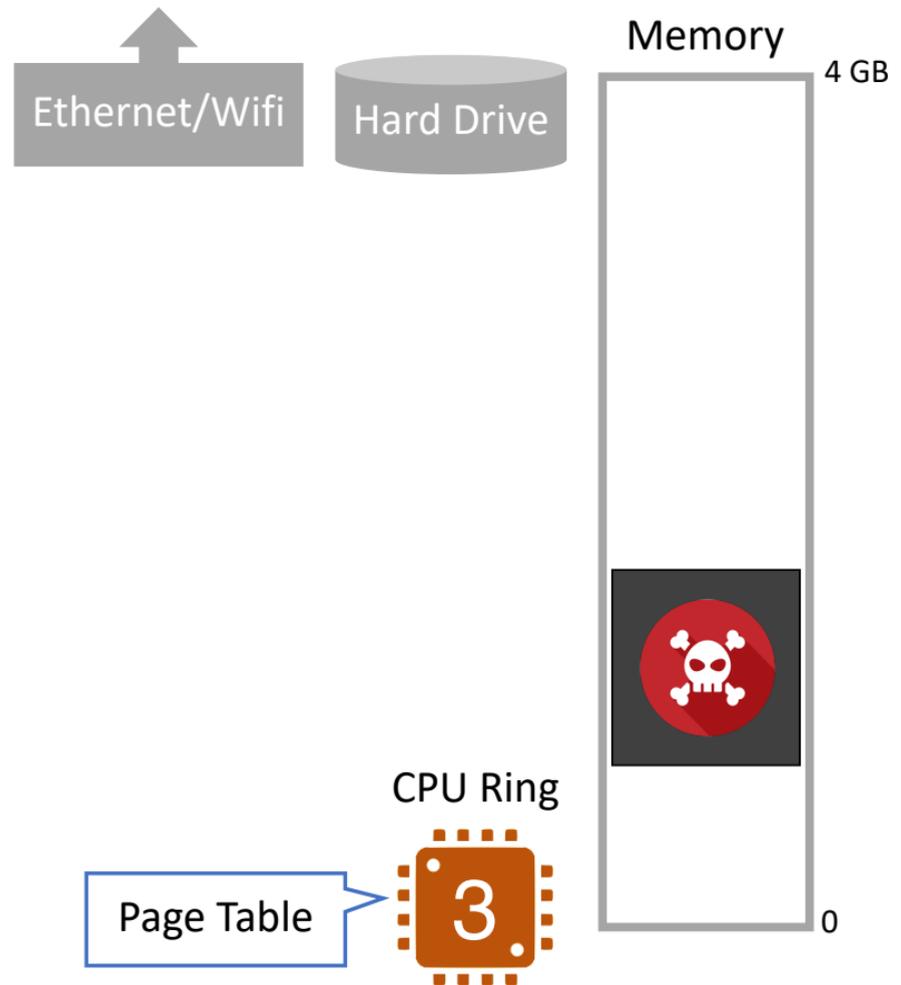
Processes cannot change  
their own page tables



# VM in Action

Processes can only read/  
write within their own  
virtual memory

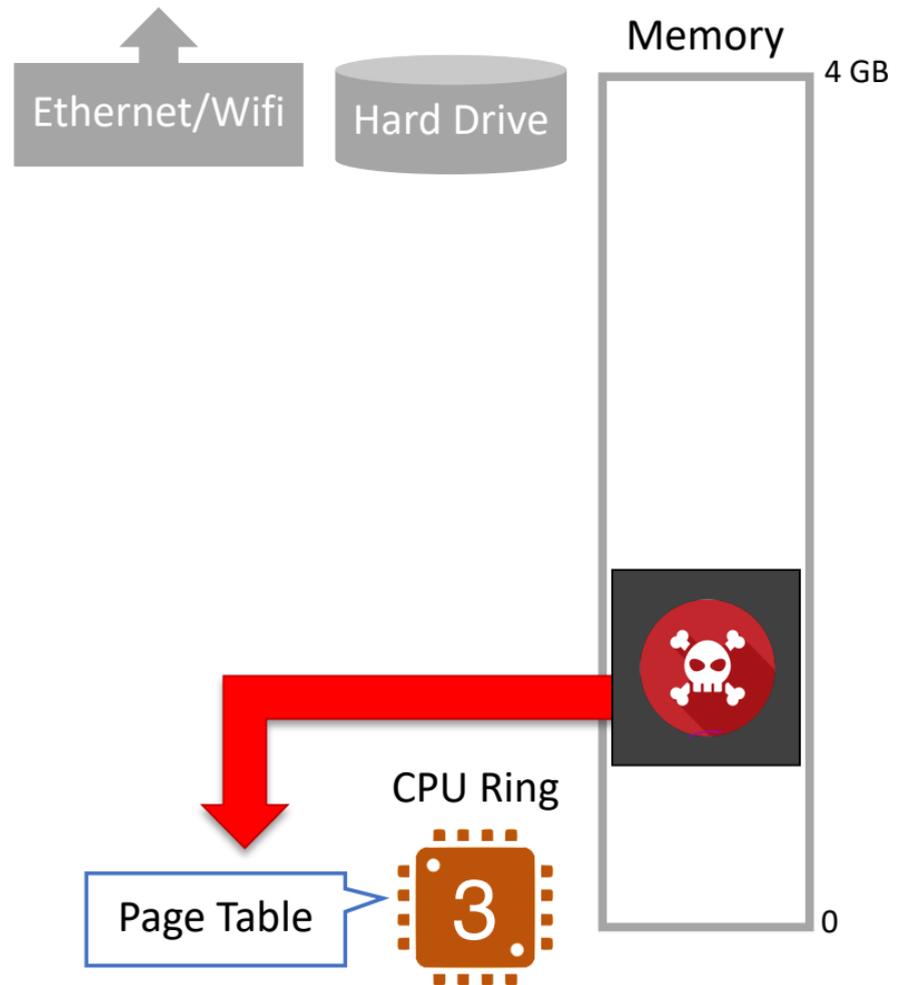
Processes cannot change  
their own page tables



# VM in Action

Processes can only read/  
write within their own  
virtual memory

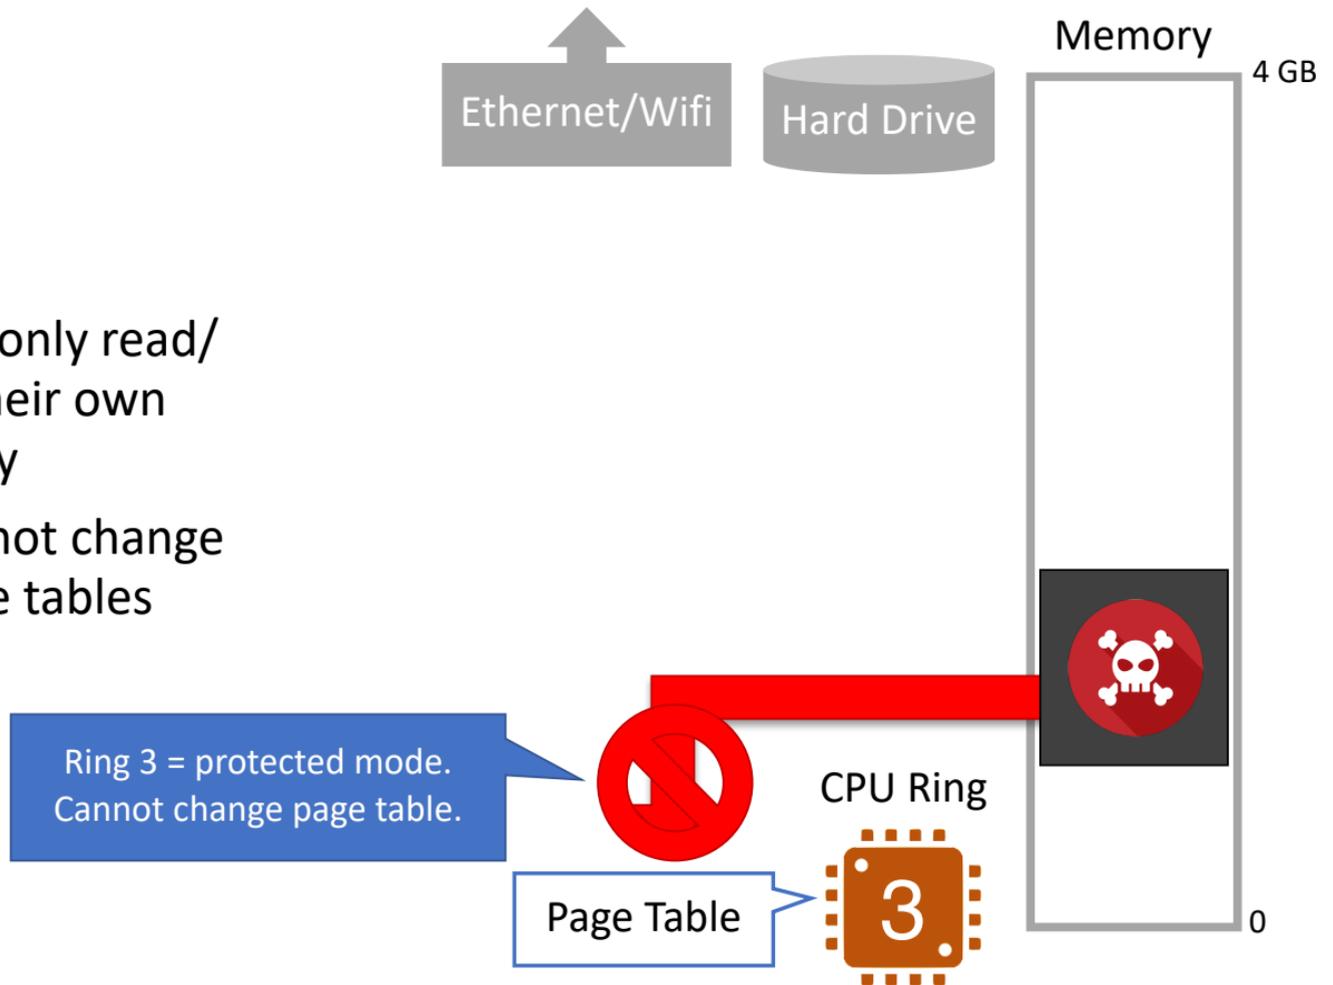
Processes cannot change  
their own page tables



# 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 moderns CPUs support these techniques

Requires OS support

- All moderns 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**



# Processes on a linux machine

ps -ax

PID	TTY	STAT	TIME	COMMAND	302 ?	S<	0:00	[loop1]						
1 ?		Ss	0:02	/sbin/init splash	303 ?	S<	0:00	[loop2]						
2 ?		S	0:00	[kthreadd]	332 ?	I<	0:00	[iprt-VBoxQueue]	1635 ?	Ssl	0:07	/usr/bin/gnome-shell		
3 ?		I<	0:00	[rcu_gp]	355 ?	I<	0:00	[cryptd]	1678 ?	S1	0:00	ibus-daemon --panel disable --xim		
4 ?		I<	0:00	[rcu_par_gp]	548 ?	Ss	0:00	/lib/systemd/systemd-resolved	1679 ?	S1	0:00	/usr/libexec/ibus-dconf		
6 ?		I<	0:00	[worker/0:0H-kblockd]	548 ?	Ssl	0:00	/lib/systemd/systemd-timesyncd	1690 ?	S1	0:00	/usr/libexec/ibus-extension-gtk3		
7 ?		I<	0:00	[worker/0:1-events]	585 ?	S1	0:00	/usr/lib/accounts-service/accounts-daemon	1693 ?	S1	0:00	/usr/libexec/ibus-x11-kill-daemon		
9 ?		I<	0:00	[mm_percpu_wq]	586 ?	Ss	0:00	/usr/sbin/acpid	1695 ?	S1	0:00	/usr/libexec/ibus-portal		
10 ?		S	0:00	[ksftirqd/0]	591 ?	Ss	0:00	avahi-daemon: running [abhi-VirtualBox.local]	1731 ?	S1	0:00	/usr/libexec/at-spi2-registrdy --use-gnome-session		
11 ?		I	0:00	[rcu_sched]	592 ?	Ss	0:00	/usr/sbin/cron -f	1771 ?	S<	0:00	[loop3]		
12 ?		S	0:00	[migration/0]	596 ?	Ss	0:01	/usr/bin/dbus-daemon --system --address=systemd: --nofork --nopidfile --systemd=daemon	1770 ?	S1	0:00	/usr/libexec/dg-permission-store		
14 ?		S	0:00	[idle_inject/0]	597 ?	Ss	0:00	/usr/bin/dbus-daemon --system --address=systemd: --nofork --nopidfile --systemd=daemon	1745 ?	S1	0:00	/usr/libexec/gnome-shell-calendar-server		
17 ?		S	0:00	[cpuhp/0]	617 ?	Ss	0:00	/usr/bin/python3 /usr/bin/networkd-dispatcher --run-startup-triggers	1751 ?	Ssl	0:00	/usr/libexec/evolution-source-registry		
15 ?		S	0:00	[kdevtmpfs]	618 ?	Ssl	0:00	/usr/lib/policykit-1/polkitd --no-debug	1762 ?	Ssl	0:00	/usr/libexec/evolution-calendar-factory		
16 ?		I<	0:00	[nets]	625 ?	Ssl	0:00	/usr/sbin/rsyslogd -n -iNONE	1771 ?	S1	0:00	/usr/libexec/dconf-service		
17 ?		S	0:00	[rcu_tasks_kthre]	628 ?	Ssl	0:12	/usr/lib/napd/napd	1776 ?	Ssl	0:00	/usr/libexec/evolution-addressbook-factory		
18 ?		S	0:00	[kauditd]	629 ?	Ssl	0:00	/usr/libexec/switcheroo-control	1804 ?	S1	0:00	/usr/libexec/gsd-wacom		
19 ?		S	0:00	[khngmtd]	638 ?	Ss	0:00	/lib/systemd/systemd-logind	1807 ?	S1	0:00	/usr/bin/gjs /usr/share/gnome-shell/org.gnome.Shell.N		
20 ?		S	0:00	[oom_reaper]	634 ?	Ssl	0:00	/usr/lib/udisks2/udisksd	1847 ?	Ssl	0:00	/usr/libexec/gsd-allly-settings		
21 ?		I<	0:00	[writeback]	641 ?	Ss	0:00	/sbin/wpa_supplicant -u -s -0 /run/wpa_supplicant	1857 ?	Ssl	0:00	/usr/libexec/gsd-color		
22 ?		S	0:00	[kcompactd0]	653 ?	S	0:00	avahi-daemon: chroot helper	1851 ?	Ssl	0:00	/usr/libexec/gsd-datetime		
23 ?		SN	0:00	[ksmd]	699 ?	Ssl	0:00	/usr/sbin/ModemManager --filter-policy=strict	1856 ?	Ssl	0:00	/usr/libexec/gsd-housekeeping		
24 ?		SN	0:00	[khugepaged]	703 ?	Ssl	0:00	/usr/sbin/cups-browsed	1856 ?	Ssl	0:00	/usr/libexec/gsd-keyboard		
71 ?		I<	0:00	[kblockd]	720 ?	Ss	0:00	/usr/sbin/cupsd -l	1861 ?	Ssl	0:00	/usr/libexec/gsd-media-keys		
72 ?		I<	0:00	[blkcg_punt_bio]	753 ?	Ssl	0:00	/usr/bin/python3 /usr/share/unattended-upgrades/unattended-upgrade-shutdown --wait-for-signal	1863 ?	Ssl	0:00	/usr/libexec/gsd-power		
73 ?		I<	0:00	[tpm_dev_wq]	758 ?	Ss	0:00	/usr/bin/whoopsie -f	1864 ?	Ssl	0:00	/usr/libexec/gsd-print-notifications		
74 ?		I<	0:00	[ata_sff]	763 ?	Ss	0:00	/usr/sbin/kerneloops	1866 ?	Ssl	0:00	/usr/libexec/gsd-rfkill		
75 ?		I<	0:00	[md]	961 ?	Ss	0:00	/usr/sbin/gdm	1875 ?	Ssl	0:00	/usr/libexec/gsd-screensaver-proxy		
76 ?		I<	0:00	[irc-poller]	974 ?	Ss	0:00	/usr/sbin/VBoxService --pidfile /var/run/vboxadd-service.sh	1879 ?	Ssl	0:00	/usr/libexec/gsd-sharing		
77 ?		I<	0:00	[devfreq_wq]	1036 ?	Smsl	0:00	/usr/libexec/rtkit-daemon	1881 ?	S1	0:00	/usr/libexec/gsd-smartcard		
78 ?		S	0:00	[watchdog]	1097 ?	S1	0:00	/usr/lib/upower/upowerd	1885 ?	S1	0:00	/usr/libexec/evolution-data-server/evolution-alarm-no		
81 ?		S	0:00	[kswapd0]	1297 ?	Ssl	0:00	/usr/libexec/color	1890 ?	S1	0:00	/usr/libexec/gsd-sound		
82 ?		S	0:00	[ecryptfs-kthrea]	1350 ?	S1	0:00	gdm-session-worker [pam/gdm-password]	1895 ?	Ssl	0:00	/usr/libexec/gsd-usb-protection		
84 ?		I<	0:00	[kthpnl4]	1369 ?	Ss	0:00	/lib/systemd/systemd --user	1899 ?	Ssl	0:00	/usr/libexec/gsd-wacom		
85 ?		S	0:00	[scsi_thermal_pm]	1370 ?	S	0:00	[sd-pam]	1902 ?	S1	0:00	/usr/libexec/gsd-wan		
86 ?		S	0:00	[scsi_eh_0]	1379 ?	S<sl	0:00	/usr/bin/pulseaudio --daemonize=no --log-target=journal	1904 ?	Ssl	0:00	/usr/libexec/gsd-xsettings		
87 ?		I<	0:00	[scsi_tmf_0]	1382 ?	S1	0:00	/usr/bin/gnome-keyring-daemon --daemonize --login	1930 ?	S1	0:00	/usr/libexec/gsd-disk-utility-notify		
88 ?		S	0:00	[scsi_eh_1]	1385 ?	Smsl	0:00	/usr/libexec/tracker-miner-fs	2005 ?	S1	0:00	/usr/libexec/gsd-printer		
89 ?		I<	0:00	[scsi_tmf_1]	1387 ?	Ss	0:00	/usr/bin/dbus-daemon --session --address=systemd: --nofork --nopidfile --systemd=ac-	2085 ?	S1	0:00	/usr/libexec/ibus-engine-simple		
90 ?		S	0:00	[worker/0:2-clea]	1391 ?	S1	0:00	/usr/libexec/gvfsd-fuse /run/user/1000/gvfs -f -o big_writes	2137 ?	S<	0:00	[loop2]		
92 ?		I<	0:00	[worker/u2:3-events_unbound]	1396 ?	S1	0:00	/usr/libexec/gvfs-udisks2-volume-monitor	2200 ?	Ss	0:01	/usr/libexec/gnome-terminal-server		
93 ?		I<	0:00	[ipvé_addrconf]	1409 ?	Ssl	0:00	/usr/libexec/gvfs-udisks2-volume-monitor	2208 pts/0	Ss	0:00	bash		
102 ?		I<	0:00	[kstrp]	1420 tty2	S1+	0:00	/usr/lib/gdm3/gdm-x-session --run-script env GNOME_SHELL_SESSION_MODE=ubuntu /usr-	2338 ?	S<	0:00	[loop1]		
105 ?		I<	0:00	[worker/u3:0]	1423 tty2	S1+	0:02	/usr/lib/xorg/Xorg vt2 -displayfd 3 -auth /run/user/1000/gdm/Xauthority -	2389 ?	S1	0:00	/usr/libexec/gvfs-metadata		
118 ?		I<	0:00	[charger_manager]	1424 ?	Ssl	0:00	/usr/libexec/gvfs-afc-volume-monitor	2392 ?	S1	0:00	update-notifier		
126 ?		S	0:00	[worker/0:2-mm_percpu_wq]	1424 ?	Ssl	0:00	/usr/libexec/gvfs-gphoto2-volume-monitor	2474 ?	S1	0:02	/usr/bin/python3 /usr/bin/update-manager --no-update		
163 ?		S	0:00	[scsi_eh_2]	1434 ?	Ssl	0:00	/usr/libexec/gvfs-goa-volume-monitor	2610 ?	I	0:00	[worker/u2:0-events_unbound]		
164 ?		I<	0:00	[scsi_tmf_2]	1438 ?	S1	0:00	/usr/libexec/goa-daemon	2707 ?	I	0:00	[worker/0:3-cgroup_destroy]		
165 ?		I<	0:00	[worker/0:1H-kbblock]	1446 ?	S1	0:00	/usr/libexec/goa-identity-service	3107 ?	I	0:00	[worker/u2:1]		
186 ?		S	0:00	[jbd2/sdas-8]	1449 ?	Ssl	0:00	/usr/libexec/gvfs-mtp-volume-monitor	3248 ?	Ss	0:00	sshd: /usr/sbin/sshd -D [listener] 0 of 10-100 startup		
187 ?		I<	0:00	[ext4-rsv-conver]	1478 tty2	S1+	0:00	/usr/libexec/gnome-session-binary --systemd --systemd --session=ubuntu	4244 ?	Ss	0:00	/usr/lib/packagekit/packagekitd		
226 ?		S<S	0:00	/lib/systemd/systemd-journald	1550 ?	S	0:00	/usr/bin/VBoxClient --clipboard	4334 pts/0	R+	0:00	ps ax		
250 ?		Ss	0:00	/lib/systemd/systemd-udev	1551 ?	S1	0:00	/usr/bin/VBoxClient --clipboard						
251 ?		S	0:00	[irq/18-vmmgfx]	1562 ?	S	0:00	/usr/bin/VBoxClient --seamless						
252 ?		I<	0:00	[ttm_swap]	1563 ?	S1	0:00	/usr/bin/VBoxClient --seamless						
261 ?		S<	0:00	[loop0]	1569 ?	S	0:00	/usr/bin/VBoxClient --draganddrop						
268 ?		S<	0:00	[loop3]	1570 ?	S1	0:02	/usr/bin/VBoxClient --draganddrop						
279 ?		S	0:00	[loop4]	1577 ?	S	0:00	/usr/bin/VBoxClient --msvga						
283 ?		S	0:00	[loop5]	1578 ?	S1	0:00	/usr/bin/VBoxClient --msvga						
296 ?		S	0:00	[loop7]	1585 ?	Ss	0:00	/usr/bin/ssh-agent /usr/bin/im-launch env GNOME_SHELL_SESSION_MODE=ubuntu /usr/bin/gnome-session --systemd --session=ubuntu						
297 ?		S<	0:00	[loop8]	1604 ?	Ssl	0:00	/usr/libexec/at-spi-bus-launcher						
298 ?		S<	0:00	[loop9]	1610 ?	S	0:00	/usr/bin/dbus-daemon --config-file=/usr/share/defaults/at-spi2/accessibility.conf --nofork --print-address 3						
300 ?		S<	0:00	[loop10]	1614 ?	Ssl	0:00	/usr/libexec/gnome-session-ctl --monitor						
		Ssl	0:00	/usr/libexec/gnome-session-binary --systemd-service --session=ubuntu	1621 ?	Ssl	0:00							

# Towards Secure Systems

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



File Access Control



Firewall



Anti-virus



Secure Logging

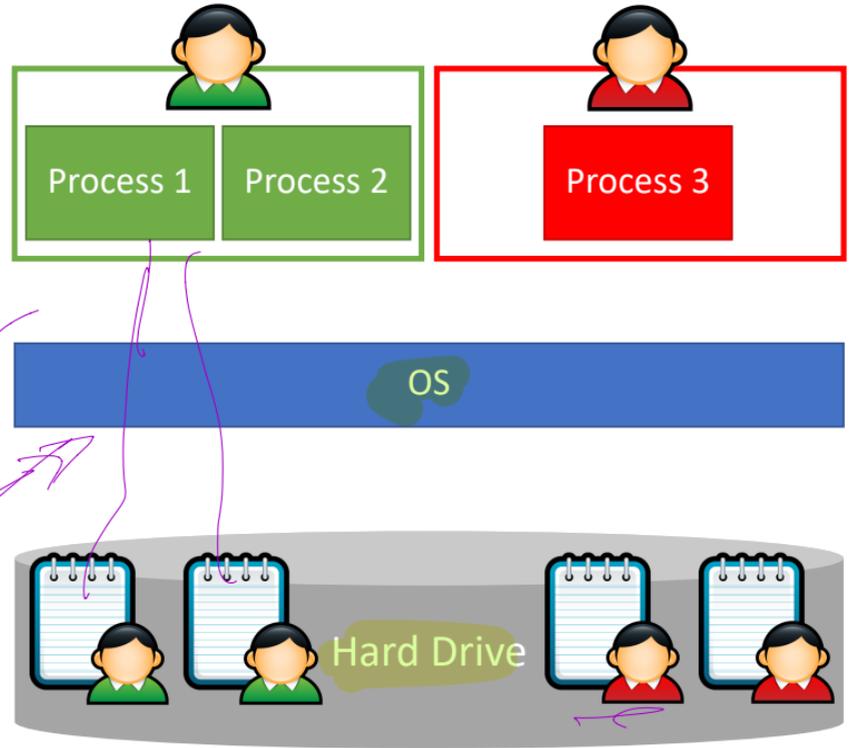
# File Access Control



All disk **access** is mediated  
by the OS

OS enforces access controls

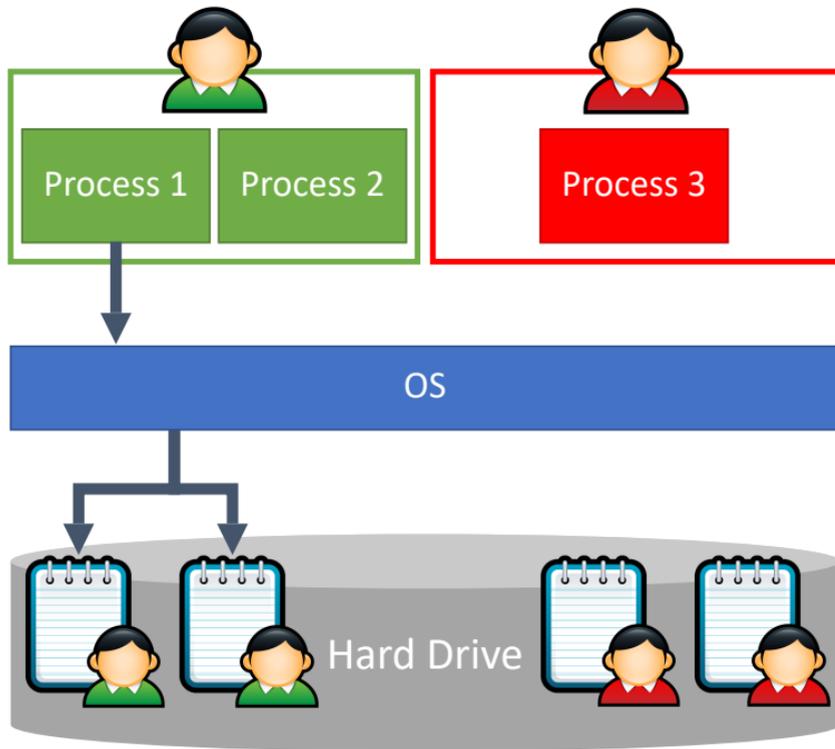
*MANDATORY & Discretionary  
access control policies*



# File Access Control



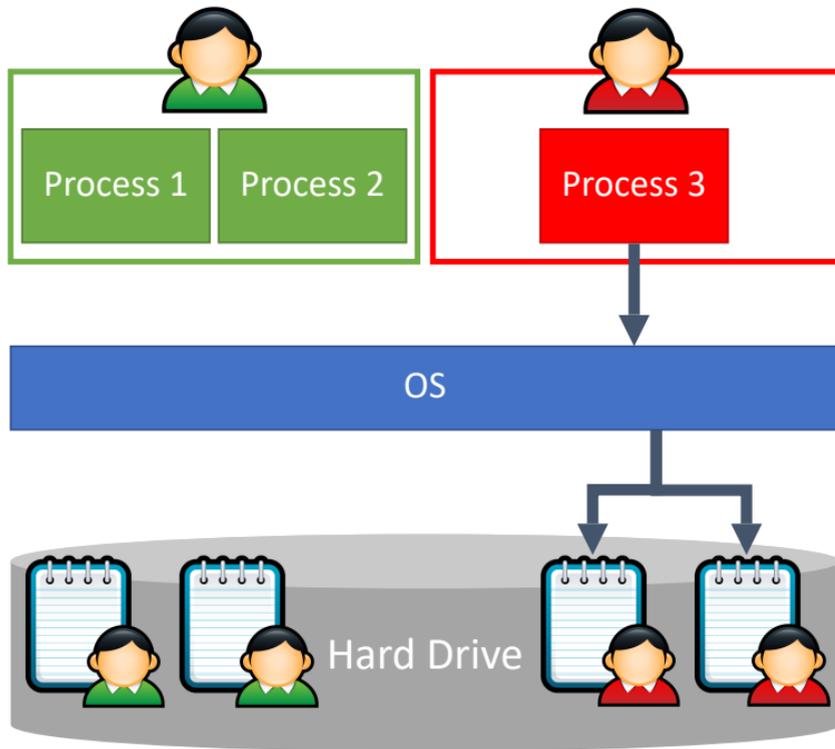
All disk access is mediated  
by the OS  
OS enforces access controls



# File Access Control



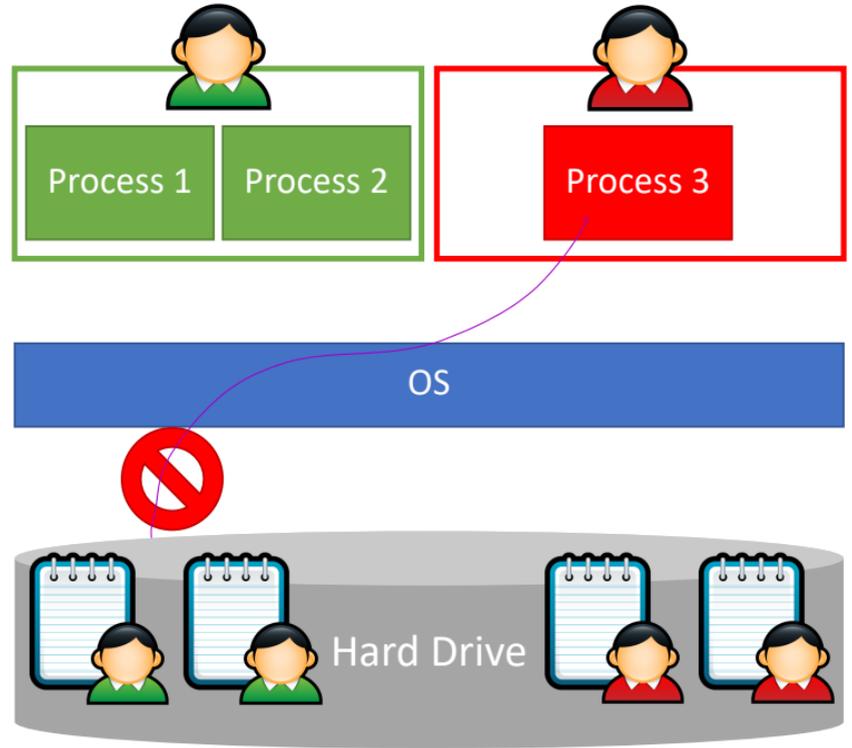
All disk access is mediated  
by the OS  
OS enforces access controls



# File Access Control



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

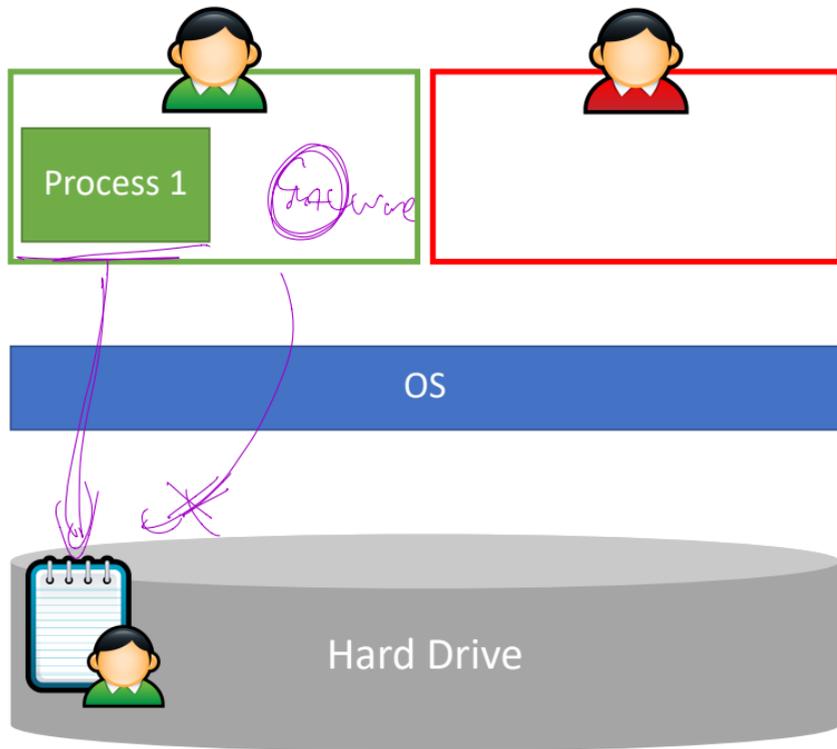


# Insecure Logging

Suppose Process 1 writes information to a log file

Malware can still destroy the log

- Add or remove entries
- Add fake entries
- Delete the whole log

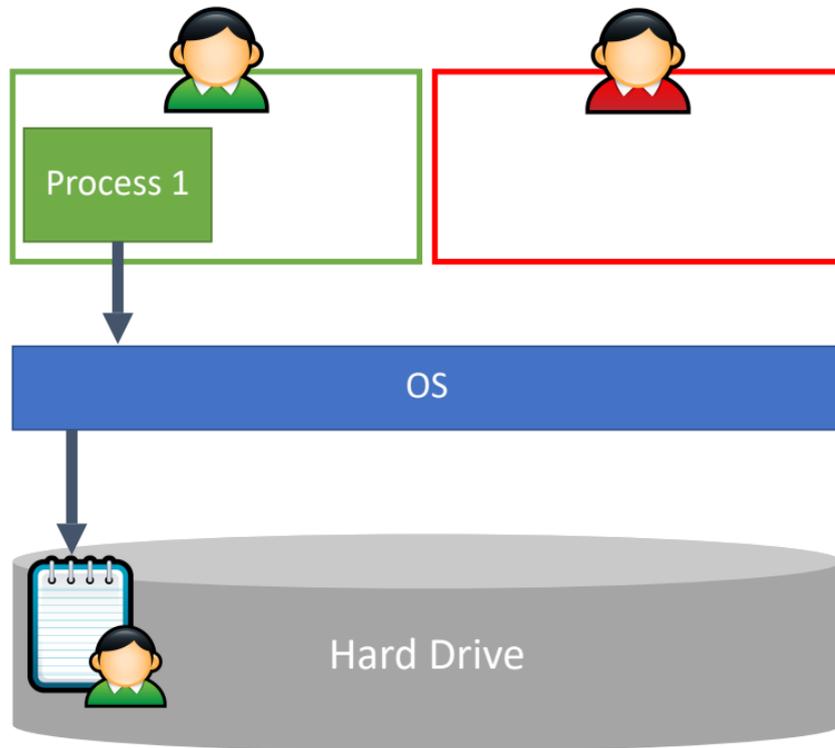


# Insecure Logging

Suppose Process 1 writes information to a log file

Malware can still destroy the log

- Add or remove entries
- Add fake entries
- Delete the whole log

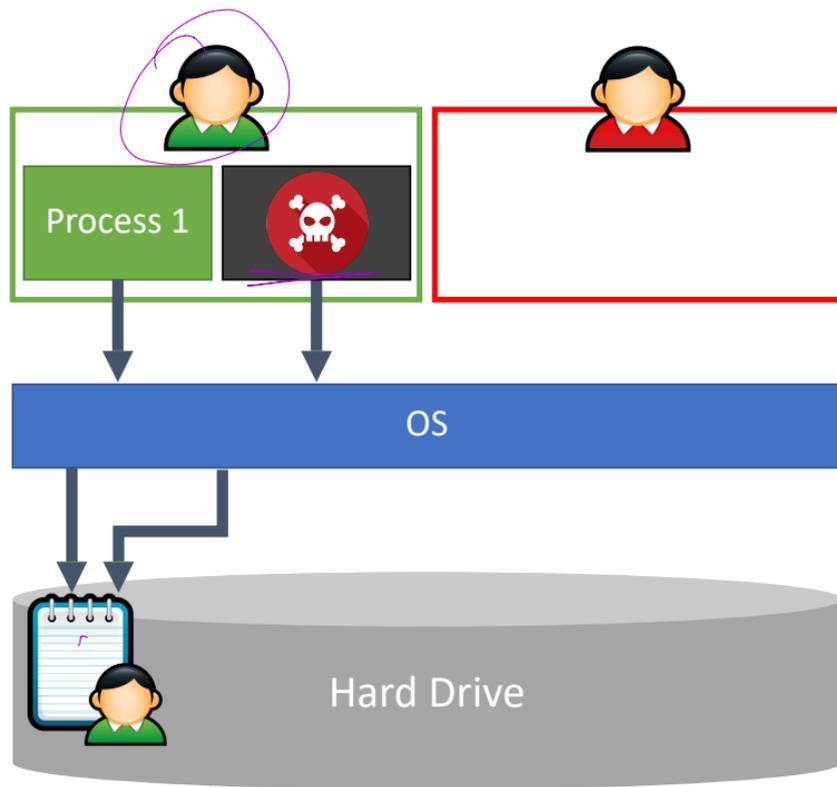


# Insecure Logging

Suppose Process 1 writes information to a log file

Malware can still destroy the log

- Add or remove entries
- Add fake entries
- Delete the whole log

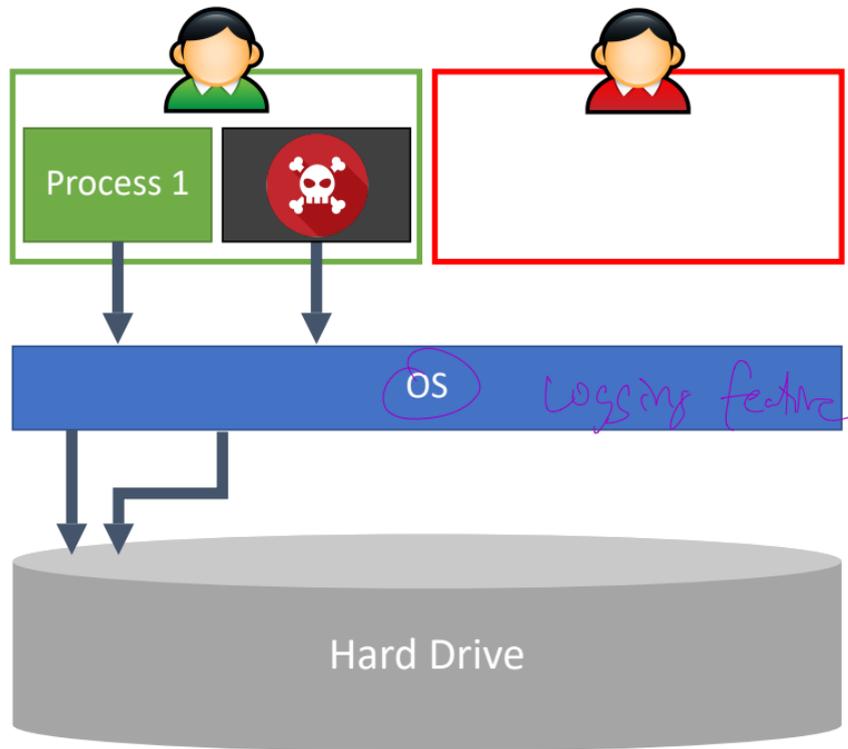


# Insecure Logging

Suppose Process 1 writes information to a log file

Malware can still destroy the log

- Add or remove entries
- Add fake entries
- Delete the whole log



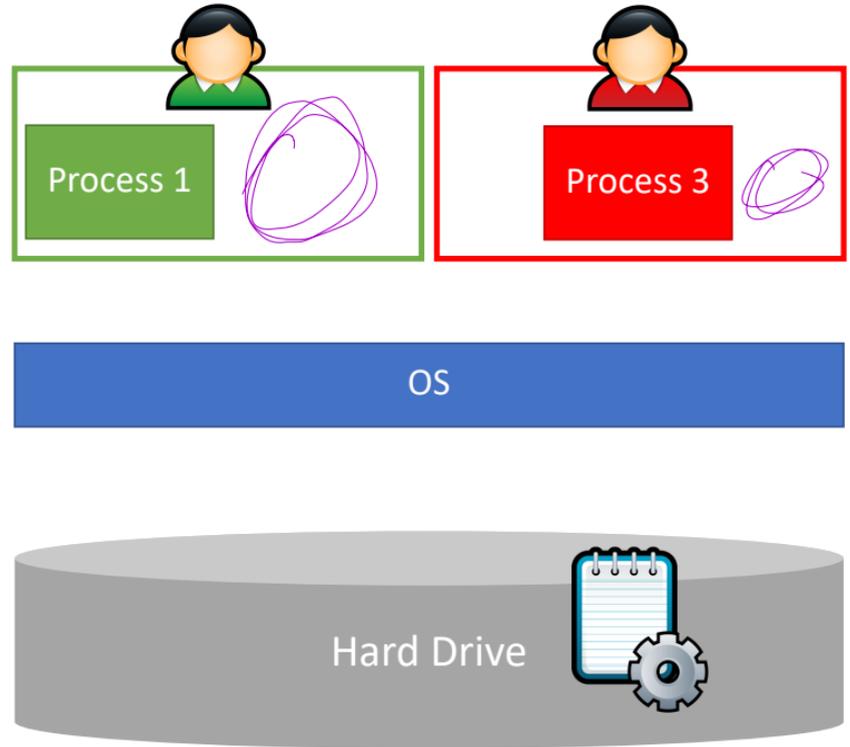


# Secure Logging

OS maintains a system log

Processes may write entries to the log using an OS API

Processes may not delete entries or the log



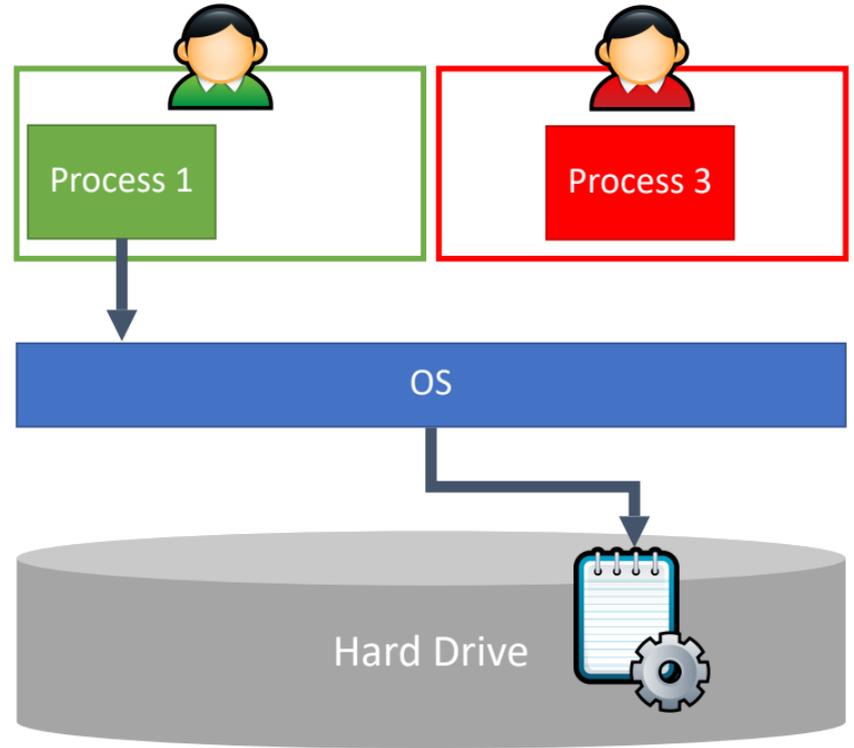


# Secure Logging

OS maintains a system log

Processes may write entries to the log using an OS API

Processes may not delete entries or the log



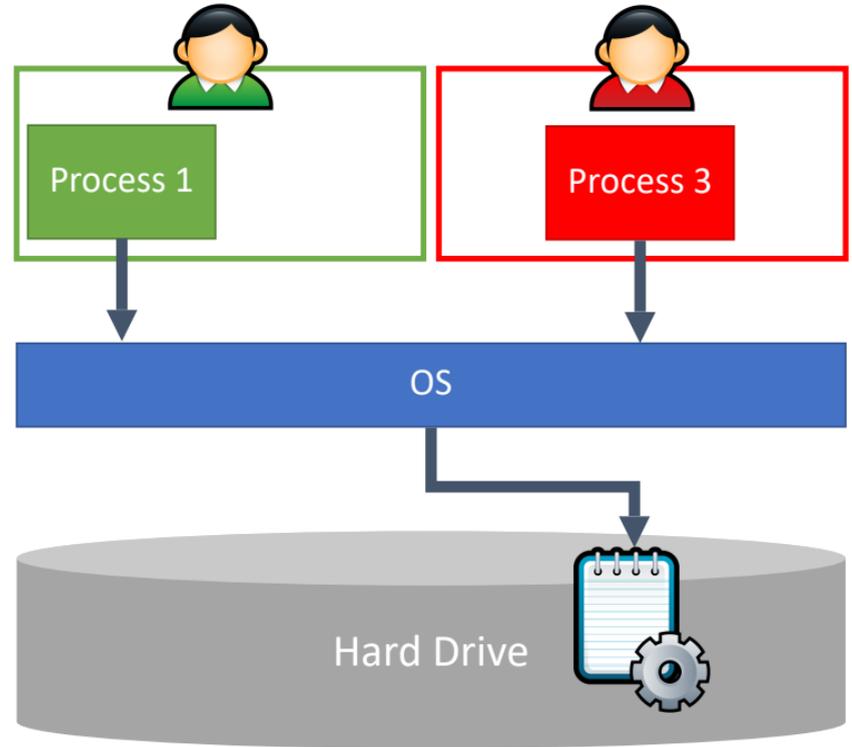


# Secure Logging

OS maintains a system log

Processes may write entries to the log using an OS API

Processes may not delete entries or the log



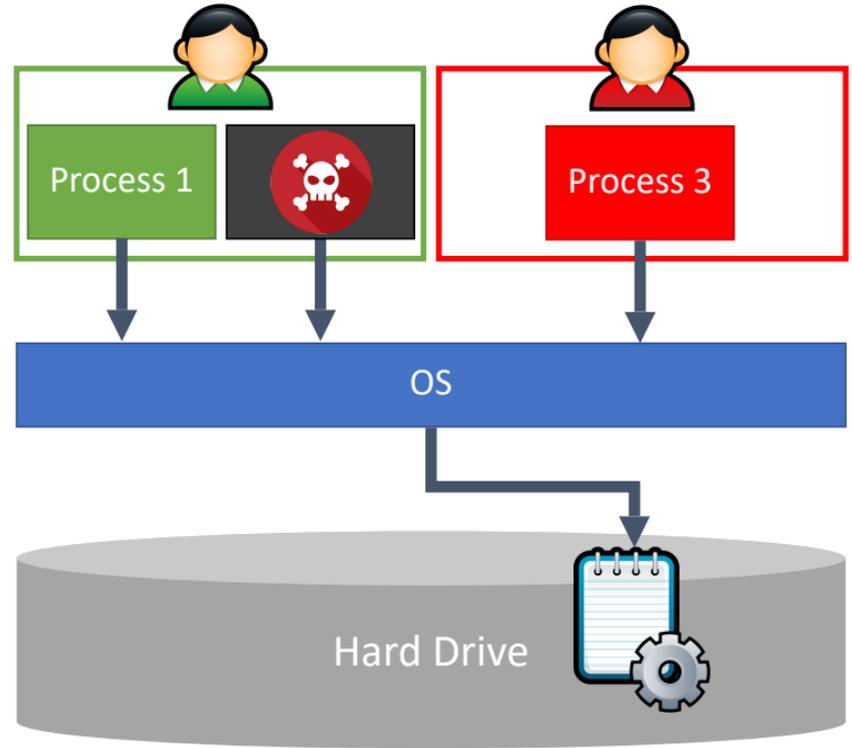


# Secure Logging

OS maintains a system log

Processes may write entries to the log using an OS API

Processes may not delete entries or the log



# Linux logging

```
abhi@abhi-VirtualBox:~$ ls /var/log
alternatives.log      boot.log             dist-upgrade        gdm3                private
alternatives.log.1   boot.log.1          dmesg               gpu-manager.log    speech-dispatcher
alternatives.log.2.gz boot.log.2          dmesg.0             hp                  syslog
apparmor             boot.log.3          dmesg.1.gz         installer           syslog.1
appport.log          boot.log.4          dmesg.2.gz         journal            syslog.2.gz
appport.log.1        boot.log.5          dmesg.3.gz         kern.log           syslog.3.gz
apt                  boot.log.6          dmesg.4.gz         kern.log.1        syslog.4.gz
auth.log             boot.log.7          dpkg.log            kern.log.2.gz     syslog.5.gz
auth.log.1           bootstrap.log       dpkg.log.1         kern.log.3.gz     syslog.6.gz
auth.log.2.gz        btmp               dpkg.log.2.gz     kern.log.4.gz     syslog.7.gz
auth.log.3.gz        btmp.1             faillog            lastlog            ubuntu-advantage.log
auth.log.4.gz        cups               fontconfig.log     openvpn            unattended-upgrades
abhi@abhi-VirtualBox:~$
```

# Syslog

`/var/log/syslog` → system events

`/var/log/auth.log` → logins.

`/var/log/kern.log` → kernel events

`/var/log/cron` — list of periodically running processes.

# syslog

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: Reloading.

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: /lib/systemd/system/dbus.socket:5: ListenStream= references a path below legacy directory /var/run/, updating /var/run/dbus/system\_bus\_socket → /run/dbus/system\_bus\_socket; please update the unit file accordingly.

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: Starting Daily apt upgrade and clean activities ...

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: Starting OpenBSD Secure Shell server ...

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: Started OpenBSD Secure Shell server.

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: apt-daily-upgrade.service: Succeeded.

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: Finished Daily apt upgrade and clean activities.

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: Reloading.

Nov 10 10:20:24 abhi-VirtualBox systemd[1]: /lib/systemd/system/dbus.socket:5: ListenStream= references a path below legacy directory /var/run/, updating /var/run/dbus/system\_bus\_socket → /run/dbus/system\_bus\_socket; please update the unit file accordingly.

Nov 10 10:20:28 abhi-VirtualBox dbus-daemon[596]: [system] Activating via systemd: service name='org.freedesktop.PackageKit' unit='packagekit.service' requested by ':1.111' (uid=0 pid=4241 comm="/usr/bin/gdbus call --system --dest org.freedesktop" label="unconfined")

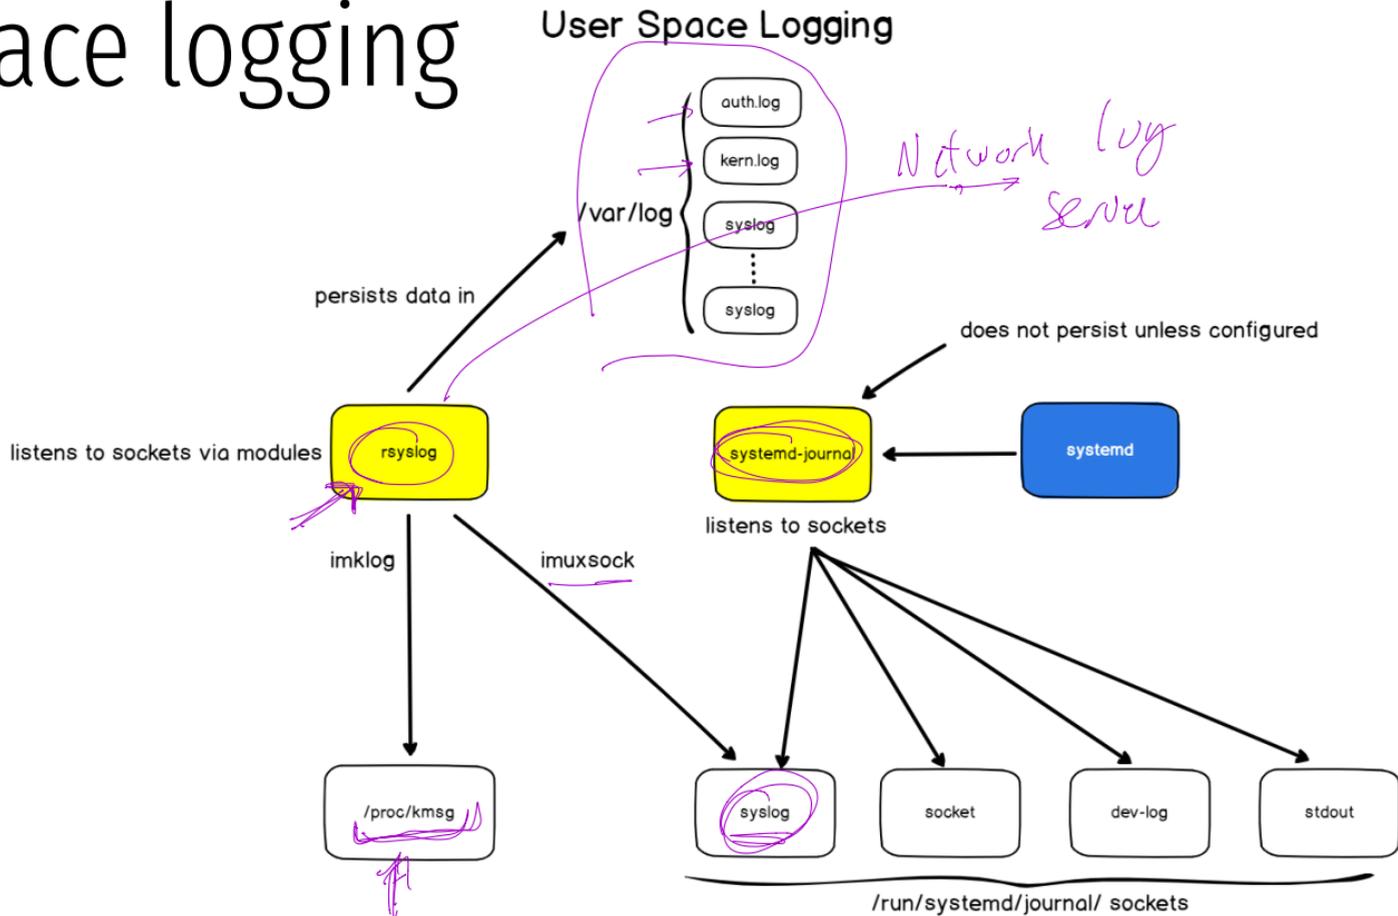
# auth.log

```
Nov 10 10:18:55 abhi-VirtualBox sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Nov 10 10:18:55 abhi-VirtualBox sudo: pam_unix(sudo:session): session closed for user root
Nov 10 10:19:01 abhi-VirtualBox sudo:      abhi : TTY=pts/0 ; PWD=/home/abhi ; USER=root ; COMMAND=/usr/bin/apt install sshd
Nov 10 10:19:01 abhi-VirtualBox sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Nov 10 10:19:02 abhi-VirtualBox sudo: pam_unix(sudo:session): session closed for user root
Nov 10 10:20:16 abhi-VirtualBox sudo:      abhi : TTY=pts/0 ; PWD=/home/abhi ; USER=root ; COMMAND=/usr/bin/apt install openssh-server
Nov 10 10:20:16 abhi-VirtualBox sudo: pam_unix(sudo:session): session opened for user root by (uid=0)
Nov 10 10:20:22 abhi-VirtualBox useradd[3079]: new user: name=sshd, UID=126, GID=65534, home=/run/sshd, shell=/usr/sbin/nologin, from=none
Nov 10 10:20:22 abhi-VirtualBox usermod[3087]: change user 'sshd' password
Nov 10 10:20:22 abhi-VirtualBox chage[3094]: changed password expiry for sshd
Nov
```

# Kernel logging

```
[ 0.000406] MTRR variable ranges enabled:
[ 0.000409] 0 base 0080000000 mask 7F00000000 uncachable
[ 0.000410] 1 base 007C000000 mask 7FFC000000 uncachable
[ 0.000412] 2 base 007A000000 mask 7FFF000000 uncachable
[ 0.000413] 3 base 0079000000 mask 7FFF000000 uncachable
[ 0.000415] 4 base 0078000000 mask 7FFF800000 uncachable
[ 0.000417] 5 base 2000000000 mask 6000000000 uncachable
[ 0.000418] 6 base 1000000000 mask 7000000000 uncachable
[ 0.000420] 7 base 4000000000 mask 4000000000 uncachable
[ 0.000421] 8 disabled
[ 0.000422] 9 disabled
[ 0.001256] x86/PAT: Configuration [0-7]: WB WC UC- UC WB WP UC- WT
[ 0.001570] last_pfn = 0x6f4f_max_arch_pfn = 0x400000000
[ 0.021713] esrt: Reserving ESRT space from 0x0000000068146518 to 0x0000000068146550.
[ 0.021729] e820: update [mem 0x68146000-0x68146fff] usable ==> reserved
[ 0.021847] check: Scanning 1 areas for low memory corruption
[ 0.021853] Using GB pages for direct mapping
[ 0.022564] RAMDISK: [mem 0x3c54000-0x3cffff]
[ 0.022580] ACPI: Early table checksum verification disabled
[ 0.022584] ACPI: RSDP 0x00000000f117014 000024 (v02 INTEL )
[ 0.022589] ACPI: XSDT 0x000000006f116728 0000CC (v01 INTEL NUC915FN 00000020 AMI 01000013)
[ 0.022590] ACPI: FACP 0x000000006f002000 000114 (v06 INTEL NUC915FN 00000020 AMI 00010013)
[ 0.022606] ACPI: DSDT 0x000000006f000000 042561 (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022610] ACPI: FACS 0x000000006f1B1000 000040
[ 0.022614] ACPI: MCFG 0x000000006f08D500 00003C (v01 INTEL NUC915FN 00000020 MSFT 00000097)
[ 0.022618] ACPI: SSDT 0x000000006f003000 001B44 (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022621] ACPI: FIDT 0x000000006f08E000 00009C (v01 INTEL NUC915FN 00000020 AMI 00010013)
[ 0.022627] ACPI: SSDT 0x000000006f08A000 0031C6 (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022631] ACPI: HPET 0x000000006f087000 000038 (v01 INTEL NUC915FN 00000020 AMI 01000013)
[ 0.022635] ACPI: SSDT 0x000000006f086000 003384 (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022639] ACPI: SOST 0x000000006f084000 001478 (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022643] ACPI: SSDT 0x000000006f080000 003280 (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022648] ACPI: NHLT 0x000000006f086000 000020 (v00 INTEL NUC915FN 00000020 AMI 01000013)
[ 0.022652] ACPI: LPIT 0x000000006f07F000 000094 (v01 INTEL NUC915FN 00000020 AMI 01000013)
[ 0.022656] ACPI: SSDT 0x000000006f07E000 002720 (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022660] ACPI: SOST 0x000000006f07A000 00007C (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022664] ACPI: DBGX 0x000000006f079000 000034 (v01 INTEL NUC915FN 00000020 AMI 01000013)
[ 0.022668] ACPI: DBG2 0x000000006f078000 000054 (v00 INTEL NUC915FN 00000020 AMI 01000013)
[ 0.022672] ACPI: SSDT 0x000000006f076000 001B66 (v02 INTEL NUC915FN 00000020 INTL 20160527)
[ 0.022677] ACPI: TPM2 0x000000006f074000 00004C (v04 INTEL NUC915FN 00000020 AMI 00000000)
[ 0.022681] ACPI: DMAR 0x000000006f075000 0000A8 (v01 INTEL NUC915FN 00000020 01000013)
[ 0.022685] ACPI: WSMT 0x000000006f07E000 000028 (v01 INTEL NUC915FN 00000020 AMI 00010013)
[ 0.022689] ACPI: APIC 0x000000006f073000 0000F4 (v04 INTEL NUC915FN 00000020 AMI 00010013)
[ 0.022693] ACPI: FPIB 0x000000006f072000 000044 (v01 INTEL NUC915FN 00000020 AMI 01000013)
[ 0.022707] ACPI: Local APIC address 0xfce00000
[ 0.023236] No NUMA configuration found
[ 0.023238] Faking a node at [mem 0x0000000000000000-0x0000000080000000]
[ 0.023254] NODE_DATA[0] allocated [mem 0x830fd5000-0x830fffff]
[ 0.023698] Zone ranges:
[ 0.023700] DMA [mem 0x0000000000001000-0x0000000000000000]
[ 0.023702] DMA32 [mem 0x0000000001000000-0x0000000007ffffff]
[ 0.023703] Normal [mem 0x0000000100000000-0x0000000080000000]
[ 0.023705] Device empty
[ 0.023706] Movable zone start for each node
[ 0.023711] Early memory node ranges
[ 0.023713] node 0: [mem 0x0000000000001000-0x00000000000009efff]
[ 0.023714] node 0: [mem 0x0000000000010000-0x0000000000000000:53ffff]
[ 0.023716] node 0: [mem 0x0000000006c4e000-0x0000000006c4efff]
[ 0.023717] node 0: [mem 0x0000000100000000-0x0000000080000000]
[ 0.024270] Zeroed struct page in unavailable ranges: 41229 pages
[ 0.024272] Initmem setup node 0 [mem 0x0000000000001000-0x0000000080000000]
[ 0.024274] On node 0 totalpages: 8314611
[ 0.024276] DMA zone: 64 pages used for memmap
[ 0.024277] DMA zone: 25 pages reserved
[ 0.024278] DMA zone: 3998 pages, LIFO batch:0
[ 0.024379] DMA32 zone: 6918 pages used for memmap
[ 0.024380] DMA32 zone: 442197 pages, LIFO batch:63
[ 0.039990] Normal zone: 122944 pages used for memmap
[ 0.039991] Normal zone: 7888416 pages, LIFO batch:63
```

# User space logging



```
abhi@abhi-VirtualBox:~$ ps ax | grep log
```

```
596 ? Ss 0:01 /usr/bin/dbus-daemon --system --address=systemd: --nofork --nopidfile --systemd-activation --syslog-only
```

```
625 ? Ssl 0:00 /usr/sbin/rsyslogd -n -iNONE
```

```
630 ? Ss 0:00 /lib/systemd/systemd-logind
```

```
1379 ? S<sl 0:00 /usr/bin/pulseaudio --daemonize=no --log-target=journal
```

```
1382 ? Sl 0:00 /usr/bin/gnome-keyring-daemon --daemonize --login
```

```
1387 ? Ss 0:00 /usr/bin/dbus-daemon --session --address=systemd: --nofork --nopidfile --systemd-activation --syslog-only
```

```
4558 pts/1 S+ 0:00 grep --color=auto log
```

```
abhi@abhi-VirtualBox:~$
```

systemd-journal



# systemd-journal

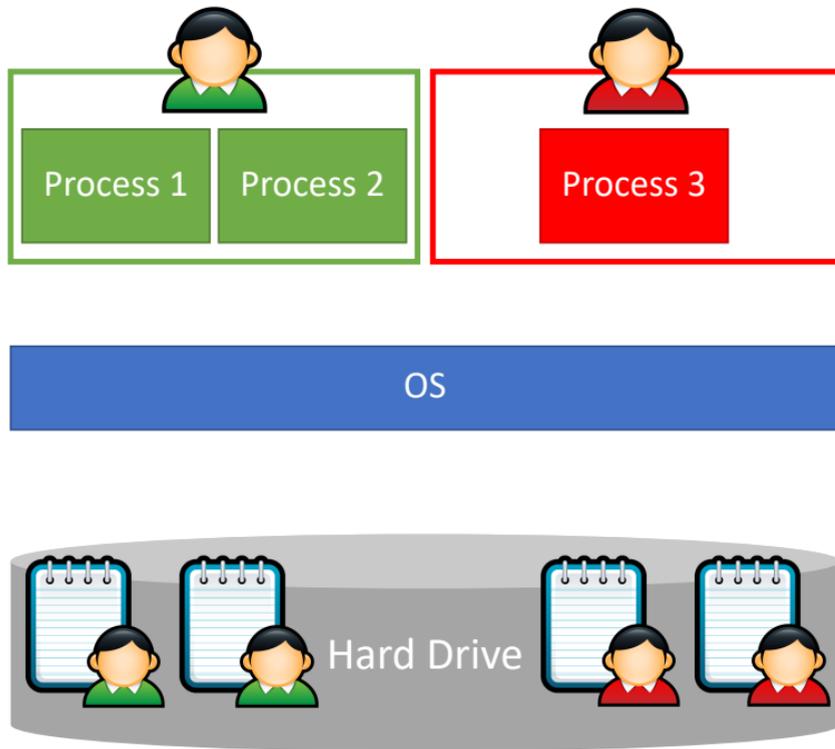
*(run)/log*

```
abhi@abhi-VirtualBox:~$ ls -al /run/
acpid.pid          gdm3/             plymouth/         thermald/
acpid.socket       gdm3.pid          sendsigs.omit.d/  tmpfiles.d/
alsa/              initctl           shm/              udev/
avahi-daemon/      initramfs/        snapd-snap.socket udisks2/
blkid/             lock/             snapd.socket      ufw.lock
console-setup/     log/              speech-dispatcher/ user/
crond.pid          mount/            spice-vdagentd/   utmp
crond.reboot       NetworkManager/  sshd/             uuidd/
cups/              openvpn/          sshd.pid          vboxadd-service.sh
dbus/              openvpn-client/   sudo/
fsck/              openvpn-server/   systemd/
abhi@abhi-VirtualBox:~$ ls -al /run/log/
total 0
drwxr-xr-x  3 root root          60 Nov 10 10:03 .
drwxr-xr-x 31 root root          880 Nov 10 10:20 ..
drwxr-sr-x+ 2 root systemd-journal 40 Nov 10 10:03 journal
abhi@abhi-VirtualBox:~$ ls -al /run/log/journal/
total 0
drwxr-sr-x+ 2 root systemd-journal 40 Nov 10 10:03 .
drwxr-xr-x  3 root root            60 Nov 10 10:03 ..
abhi@abhi-VirtualBox:~$
```

# File Access Control



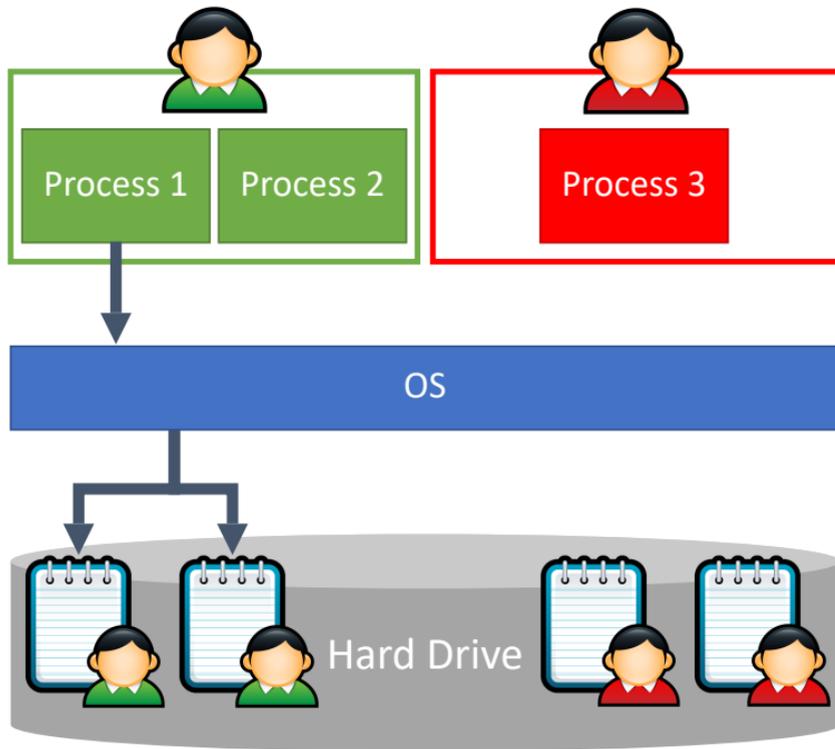
All disk access is mediated  
by the OS  
OS enforces access controls



# File Access Control



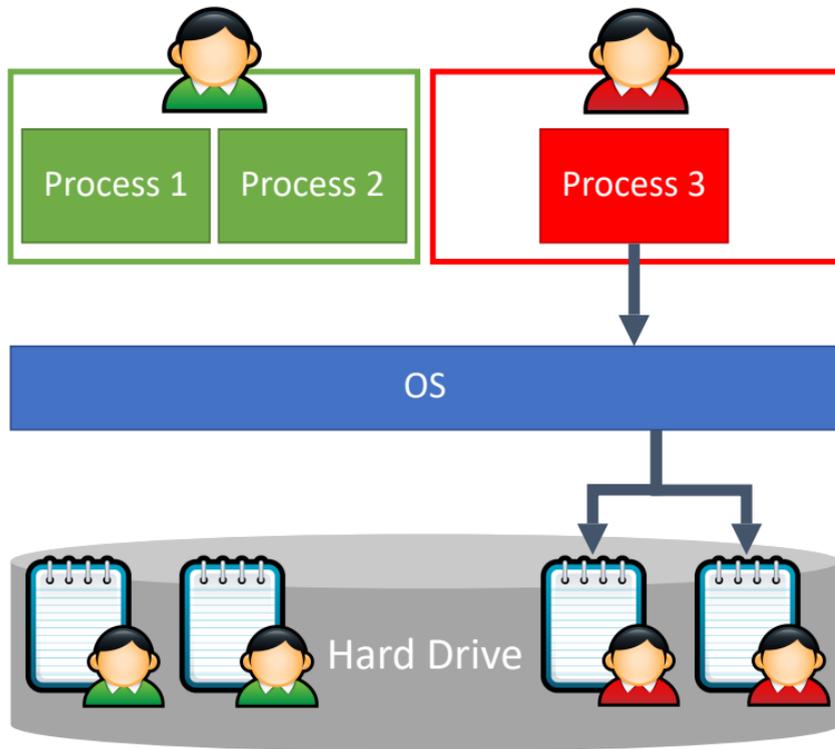
All disk access is mediated  
by the OS  
OS enforces access controls



# File Access Control



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

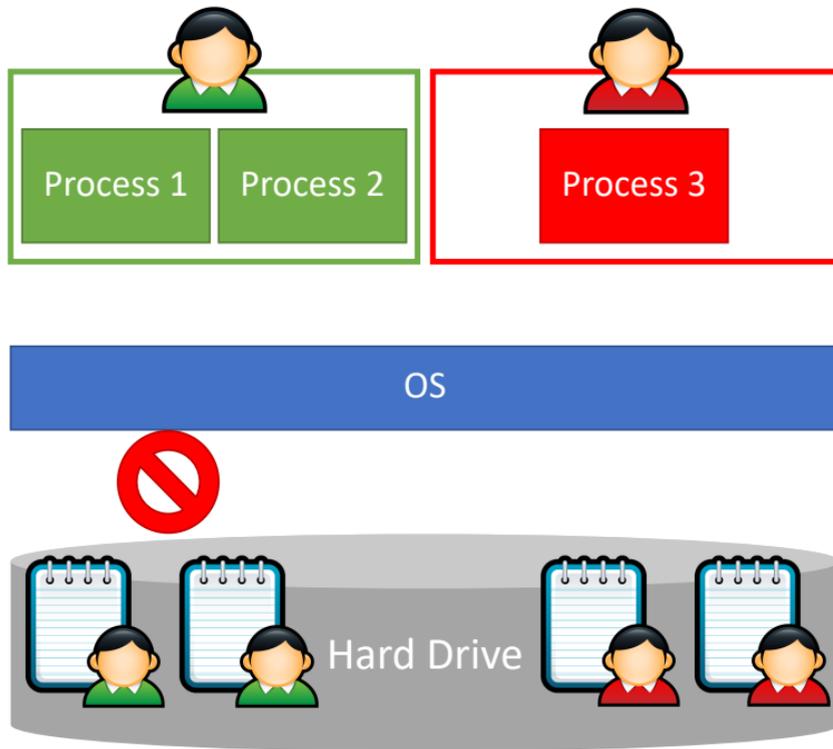


# File Access Control



All disk access is mediated  
by the OS

OS enforces access controls

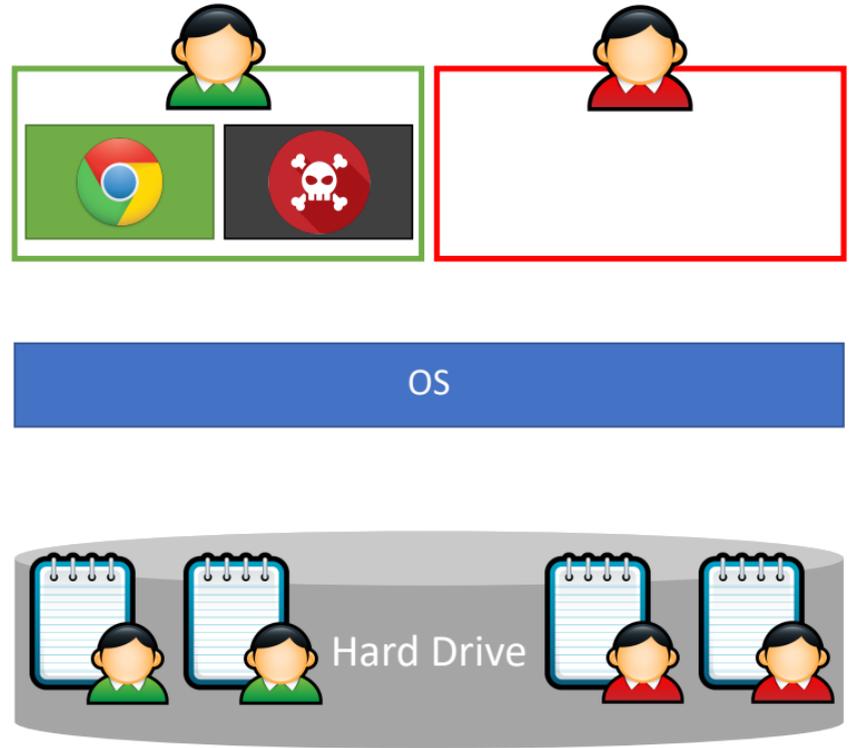




# Limitations

Malware can still cause damage

Discretionary access control means that isolation is incomplete

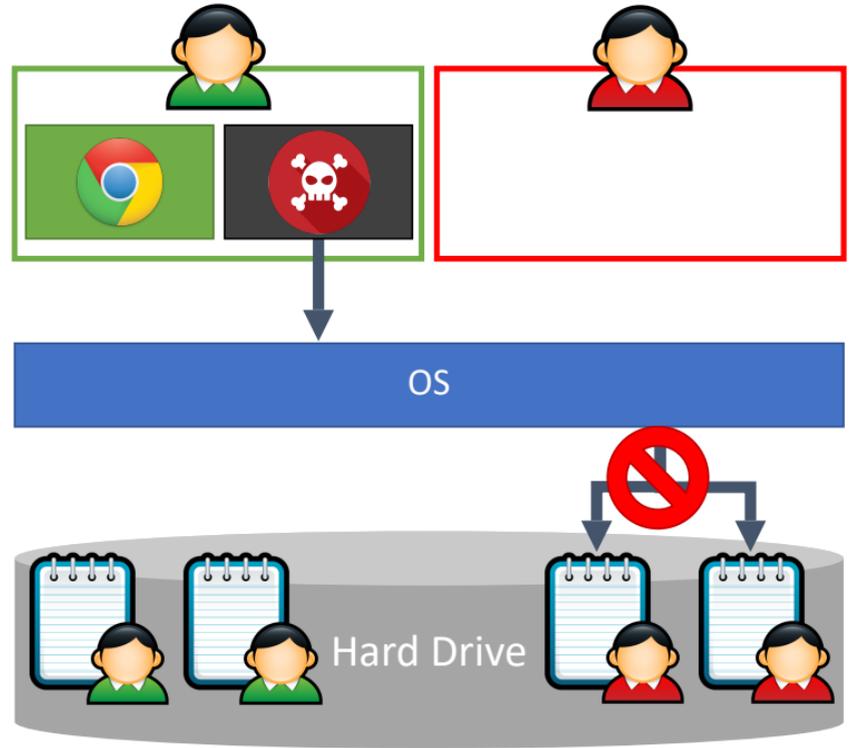




# Limitations

Malware can still cause damage

Discretionary access control means that isolation is incomplete

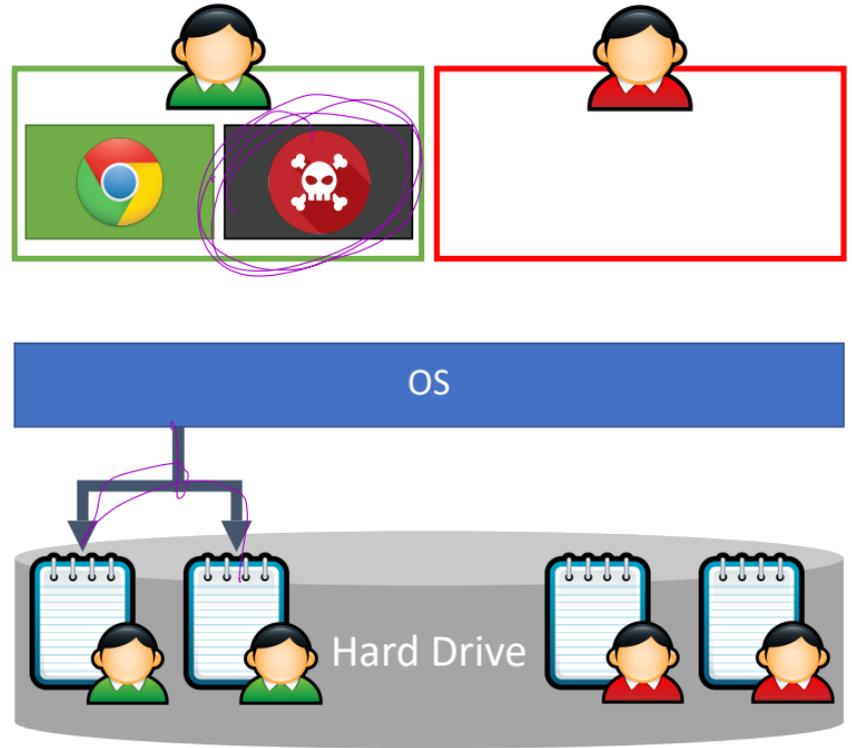




# Limitations

Malware can still cause damage

Discretionary access control means that isolation is incomplete



# Anti-virus

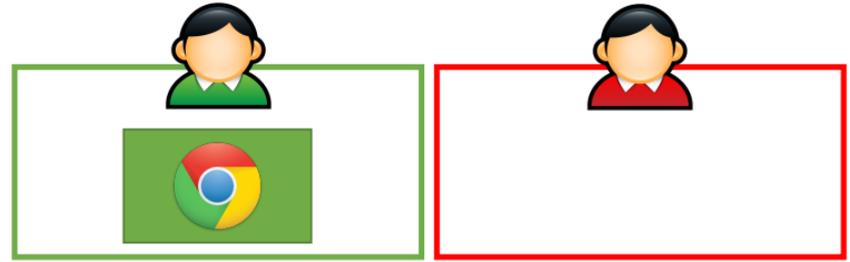
Anti-virus process is **privileged**

- 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

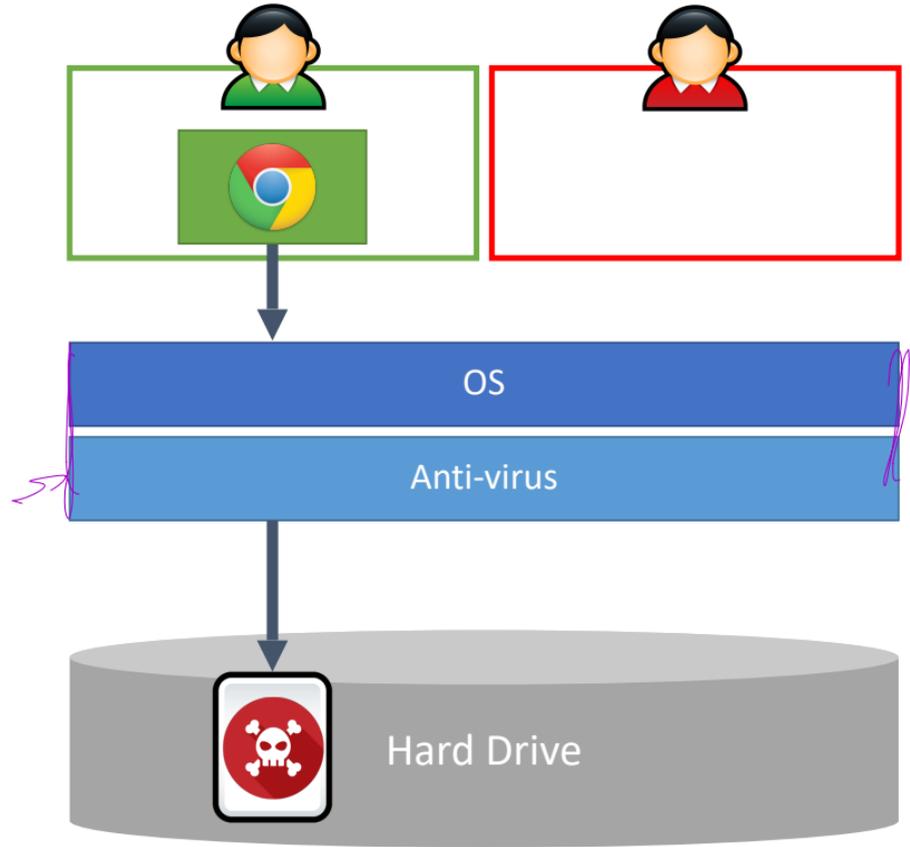
Anti-virus process is **privileged**

- 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

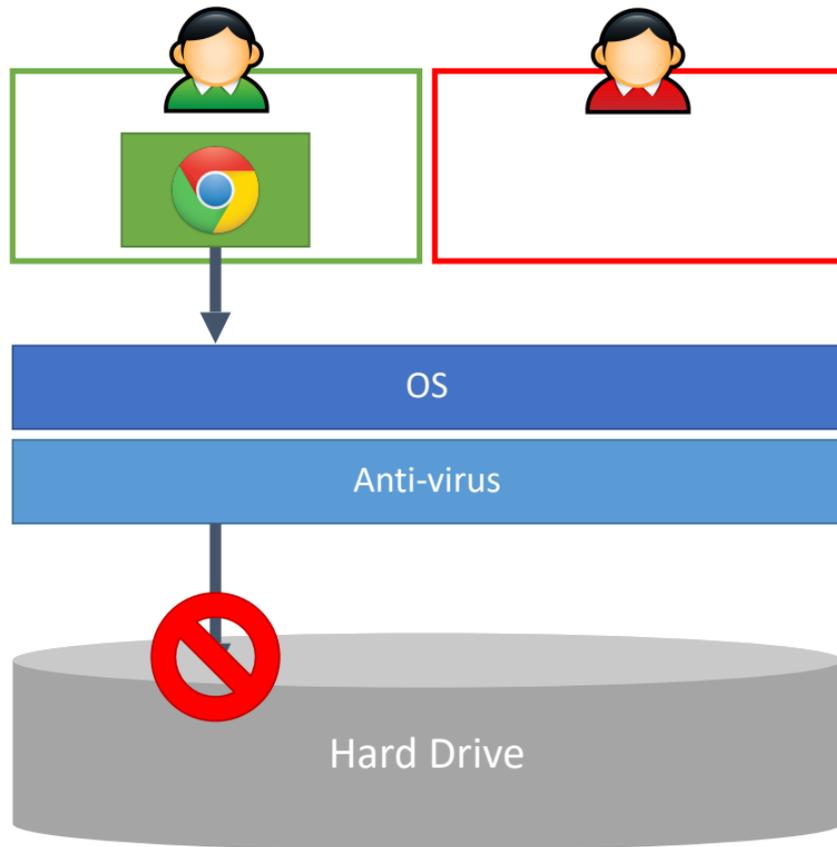
Anti-virus process is **privileged**

- 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

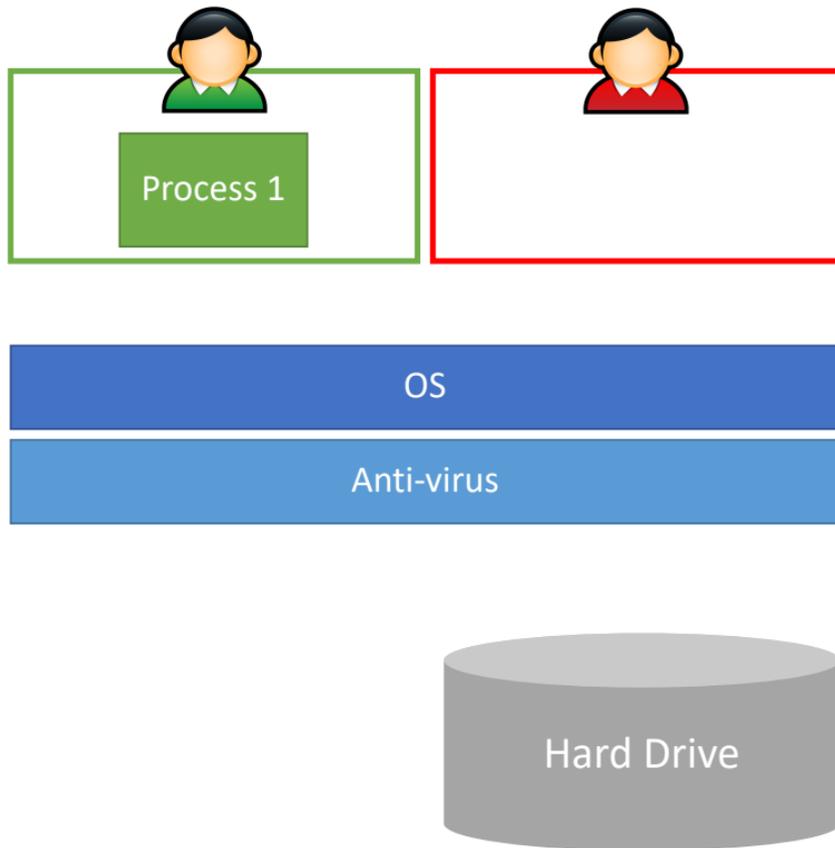
Anti-virus process is  
**privileged**

- 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

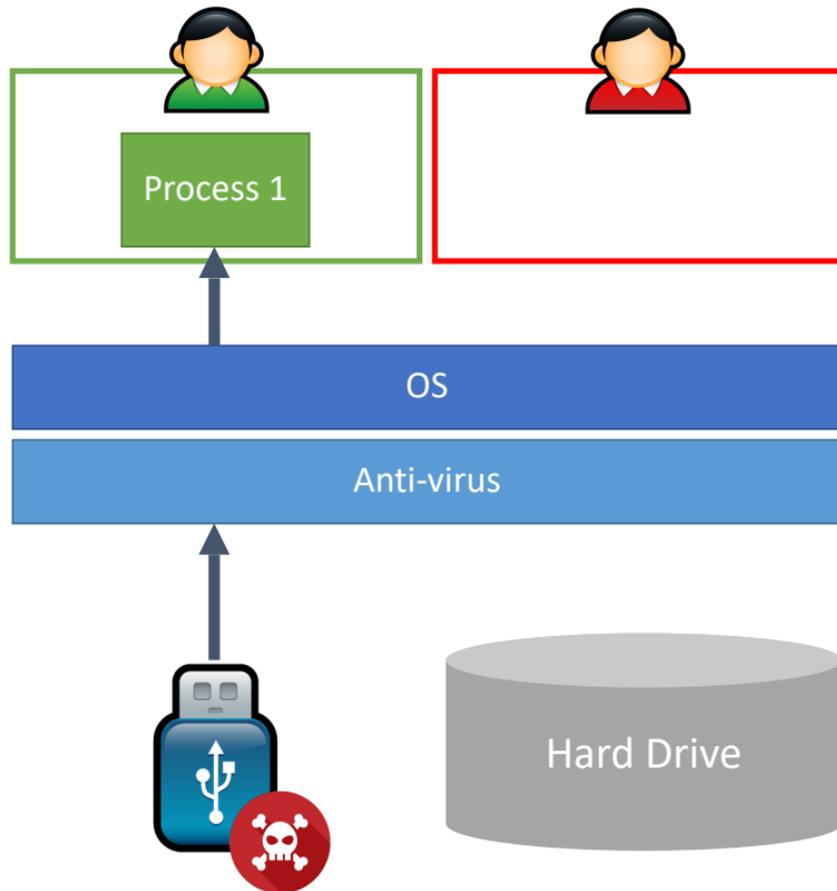
Anti-virus process is **privileged**

- 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

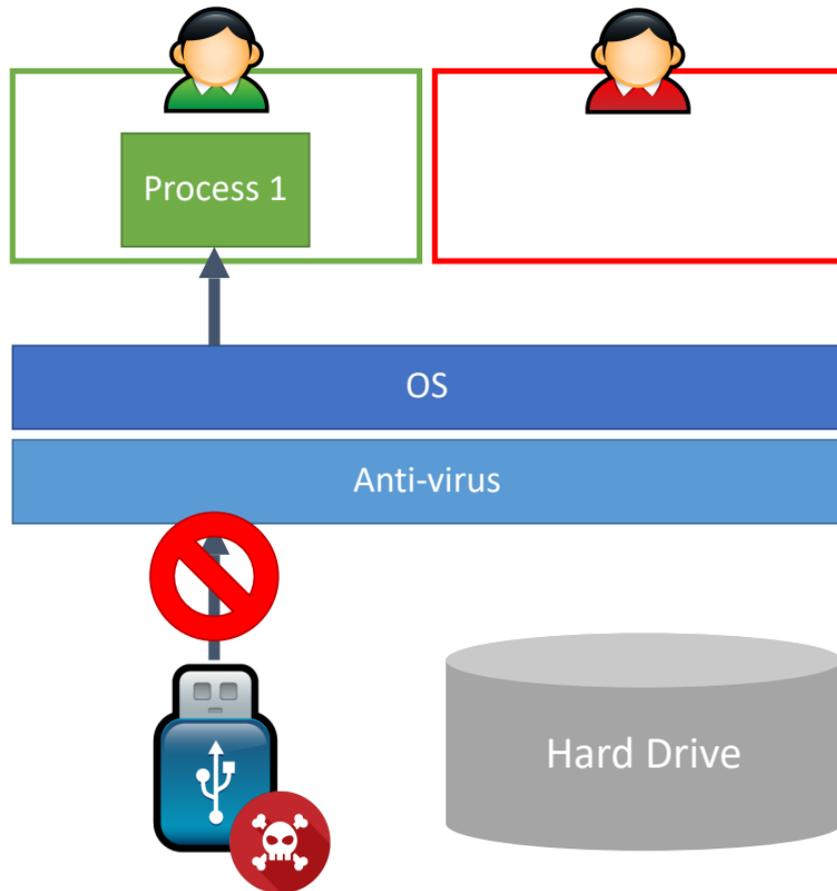
Anti-virus process is **privileged**

- 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



# Signature-based Detection

Key idea: identify invariants that correspond to malicious code or data

Example – anti-virus signatures

- List of code snippets that are unique to known malware

Problems with signatures

# Signature-based Detection

Key idea: identify **invariants** that correspond to malicious code or data

Example – anti-virus signatures

- List of code snippets that are unique to known malware

Problems with signatures

- Must be updated frequently
- May cause false positives
  - Accidental overlaps with good programs and benign network traffic

# Avast Malware Signature Update Breaks Installed Programs

Users of the free version of Avast antivirus unscathed

May 7, 2015 13:55 GMT · By Ionut Ilascu · Share:

---

**A bad virus definition update from Avast released on Wednesday caused a lot of trouble, as it mistook various components in legitimate programs installed on the machine for malware.**

The list of valid software affected by the signature update includes [Firefox](#), [iTunes](#), NVIDIA drivers, Google Chrome, Adobe [Flash Player](#), [Skype](#), Opera, [TeamViewer](#), ATI drivers, as well as products from [Corel](#) and components of Microsoft Office.

# Avoiding Anti-virus

Malware authors go to great length to avoid detection by AV

## Polymorphism

- Viral code mutates after every infection

# Avoiding Anti-virus

Malware authors go to great length to avoid detection by AV

## Polymorphism

- Viral code mutates after every infection

$$b = a + 10$$

# Avoiding Anti-virus

Malware authors go to great length to avoid detection by AV

## Polymorphism

- Viral code mutates after every infection

$$b = a + 10$$

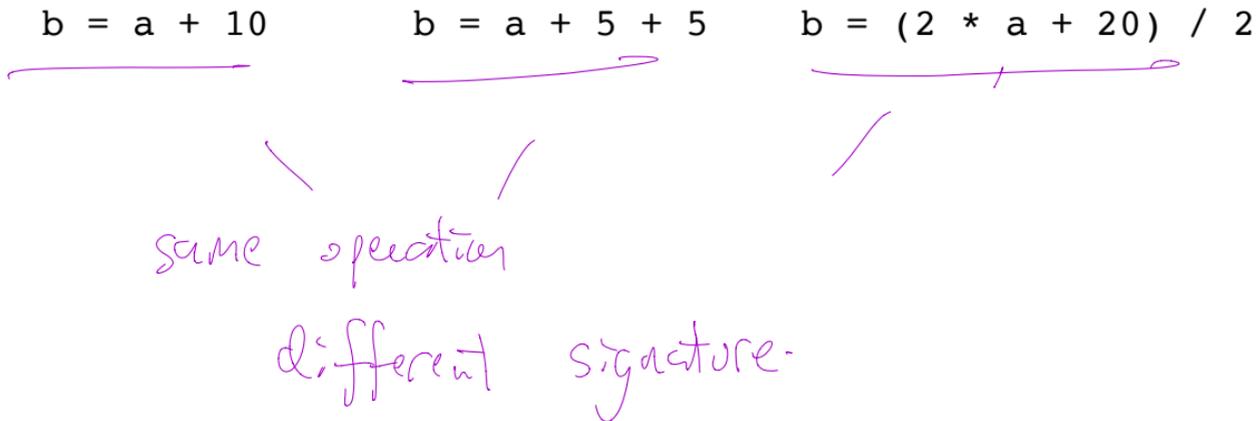
$$b = a + 5 + 5$$

# Avoiding Anti-virus

Malware authors go to great length to avoid detection by AV

## Polymorphism

- Viral code mutates after every infection



# Avoiding Anti-virus

Malware authors go to great length to avoid detection by AV

## Polymorphism

- Viral code mutates after every infection

$$b = a + 10 \qquad b = a + 5 + 5 \qquad b = (2 * a + 20) / 2$$

## Packing

- Malware code is encrypted, key is changed every infection
- Decryption code is vulnerable to signature construction
- Polymorphism may be used to mutate the decryption code

# Firewall

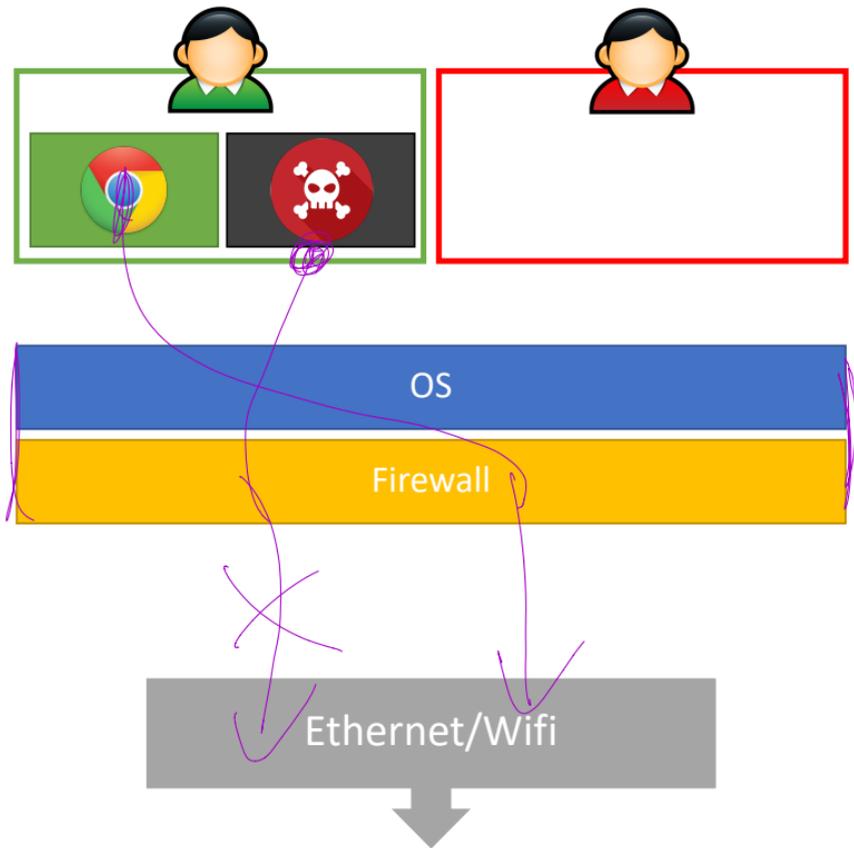
Firewall process is  
privileged

- Often runs in Ring 0

Selectively blocks network  
traffic

- By process
- By port
- By IP address
- By packet content

Inspects outgoing and  
incoming network traffic



# Firewall

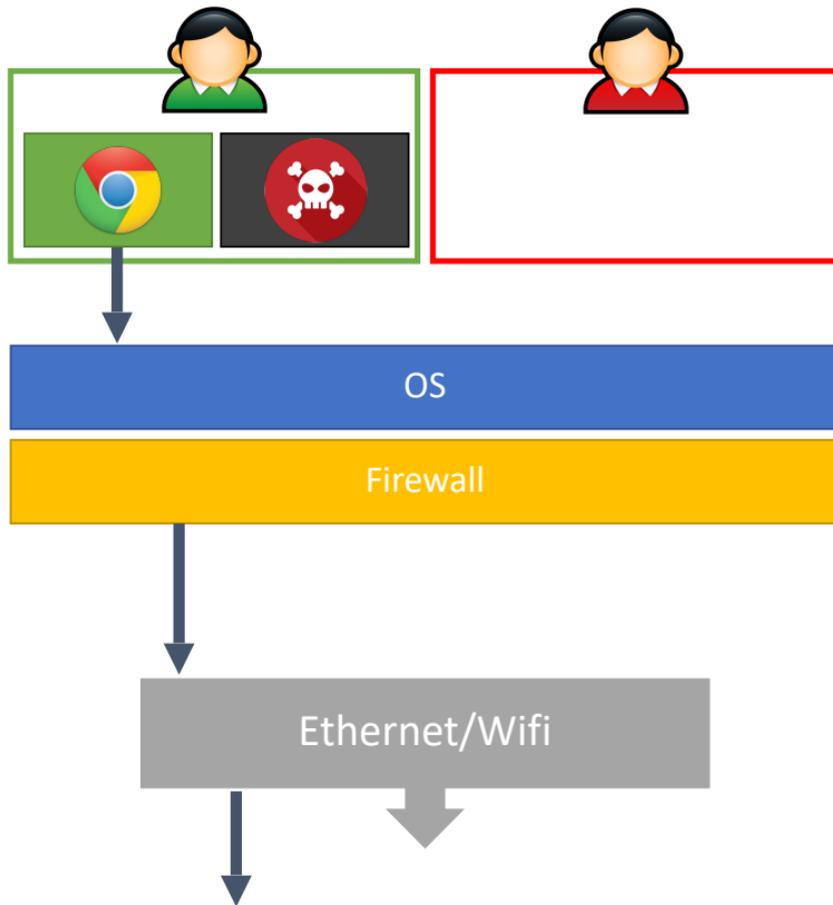
Firewall process is **privileged**

- Often runs in Ring 0

Selectively blocks network traffic

- By process
- By port
- By IP address
- By packet content

Inspects outgoing and incoming network traffic



# Firewall

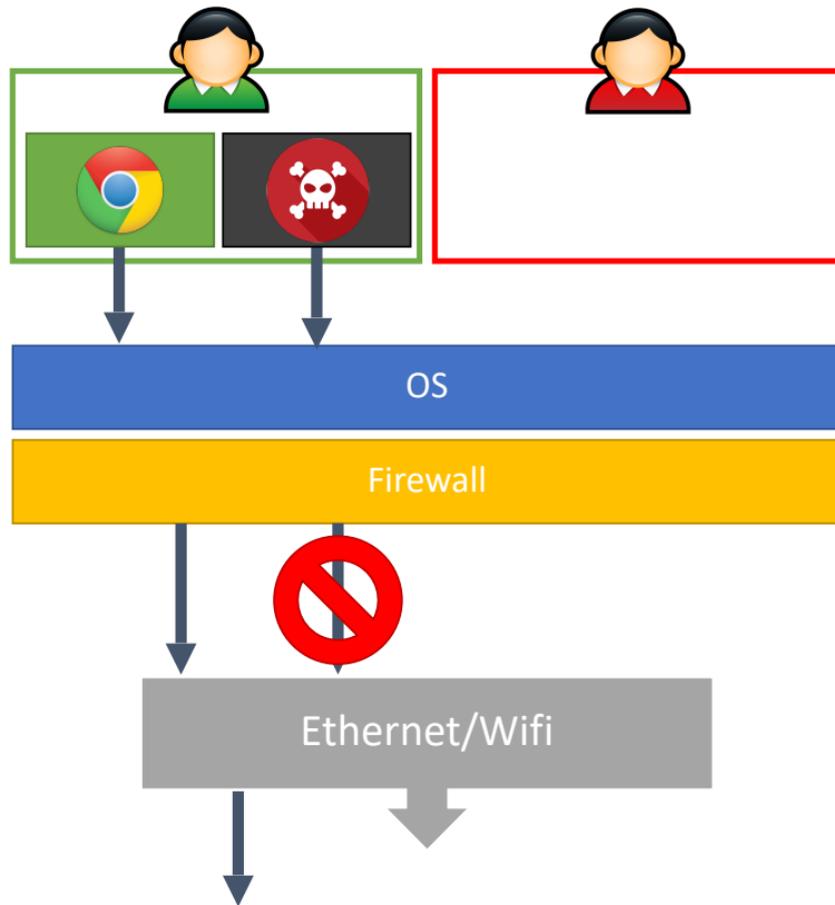
Firewall process is **privileged**

- Often runs in Ring 0

Selectively blocks network traffic

- By process
- By port
- By IP address
- By packet content

Inspects outgoing and incoming network traffic



# Firewall

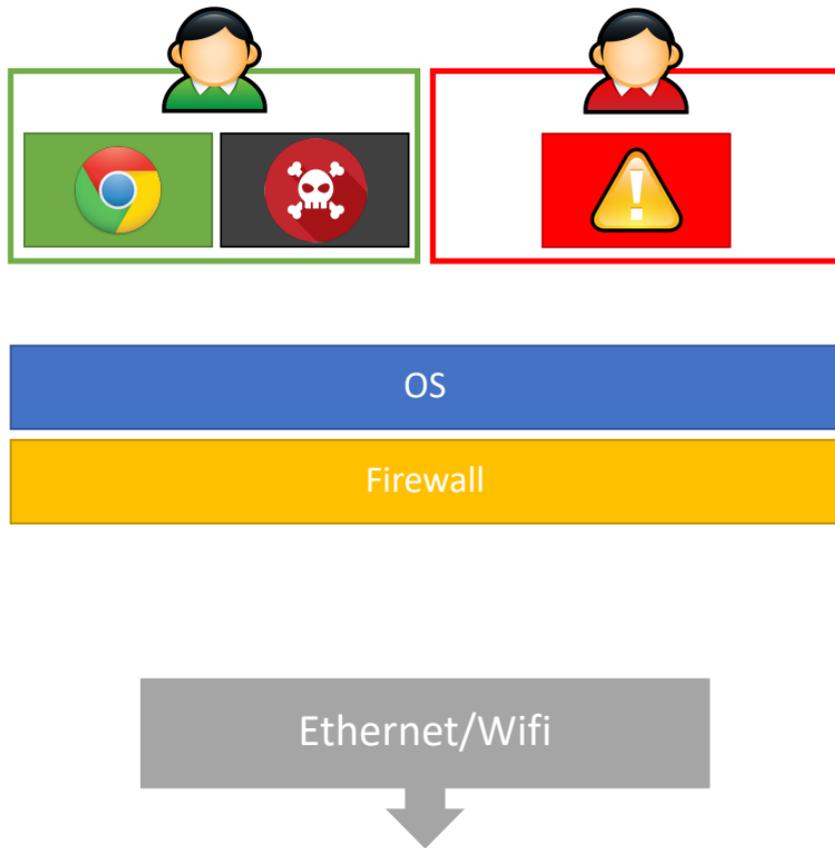
Firewall process is **privileged**

- Often runs in Ring 0

Selectively blocks network traffic

- By process
- By port
- By IP address
- By packet content

Inspects outgoing and incoming network traffic



# Firewall

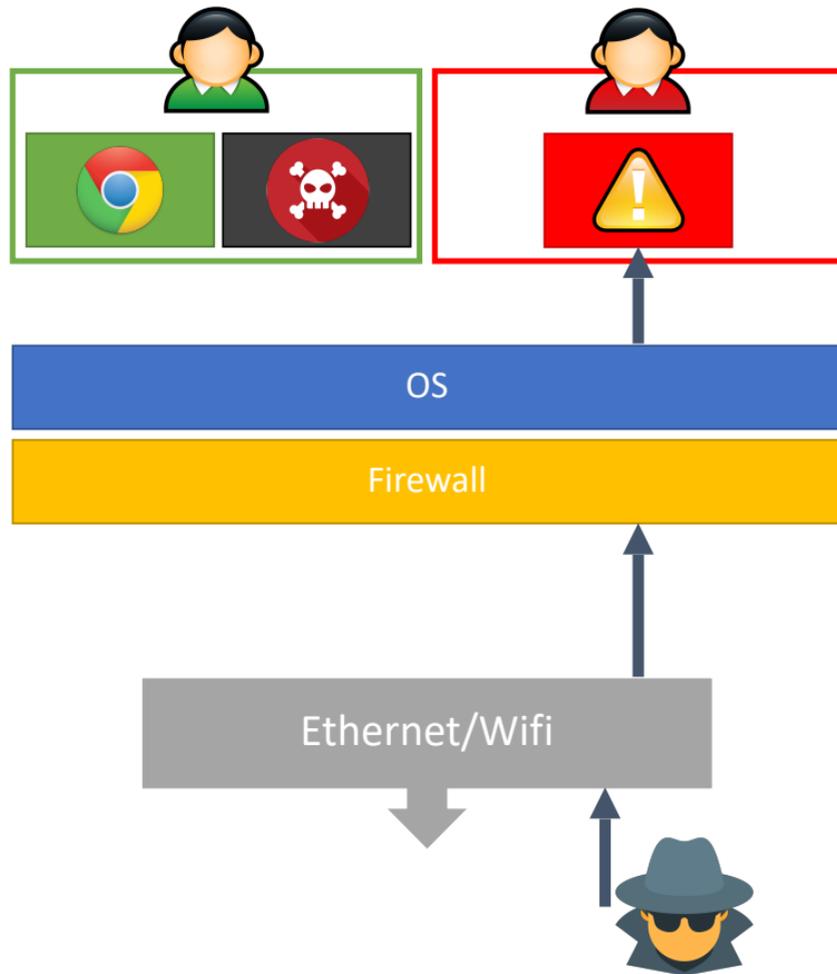
Firewall process is **privileged**

- Often runs in Ring 0

Selectively blocks network traffic

- By process
- By port
- By IP address
- By packet content

Inspects outgoing and incoming network traffic



# Firewall

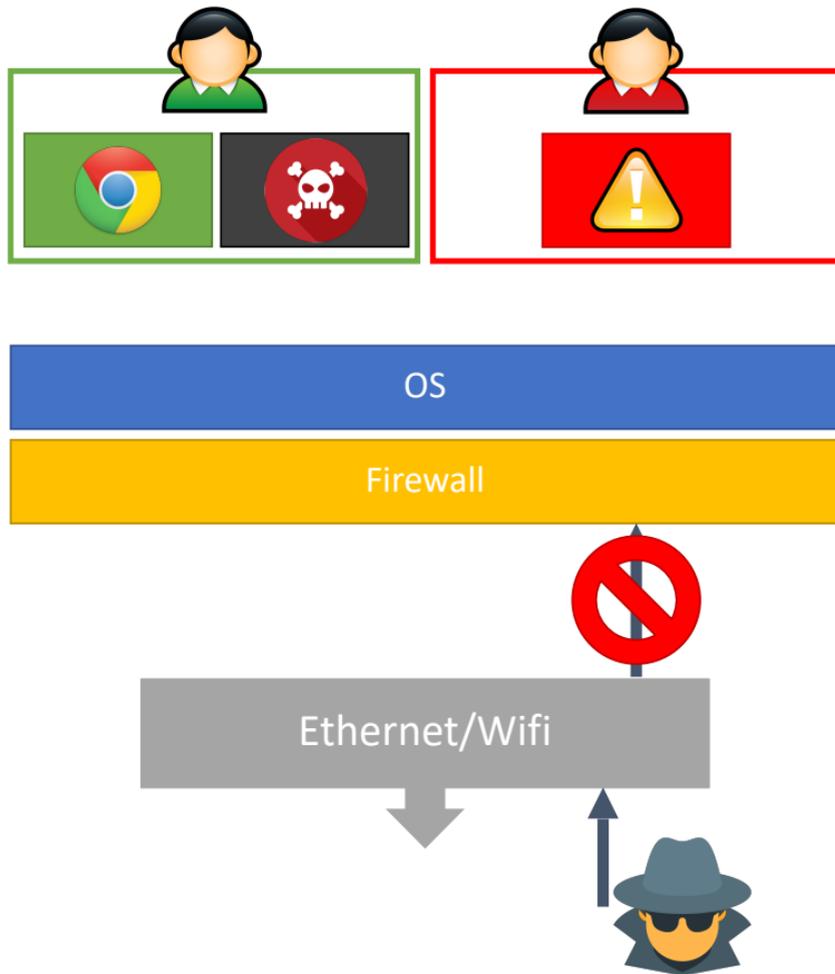
Firewall process is **privileged**

- Often runs in Ring 0

Selectively blocks network traffic

- By process
- By port
- By IP address
- By packet content

Inspects outgoing and incoming network traffic



# Network Intrusion Detection Systems

NIDS for short

Snort

- Open source intrusion prevention system capable of real-time traffic analysis and packet logging
- Identifies malicious network traffic using signatures



Bro

- Open source network monitoring, analysis, and logging framework
- Can be used to implement signature based detection
- Capable of more complex analysis

