

Linux Kernel Programmierung



We will Assimilate you!

*graphic design is
my passion.*



Warum?

Themenfelder Kernelentwicklung

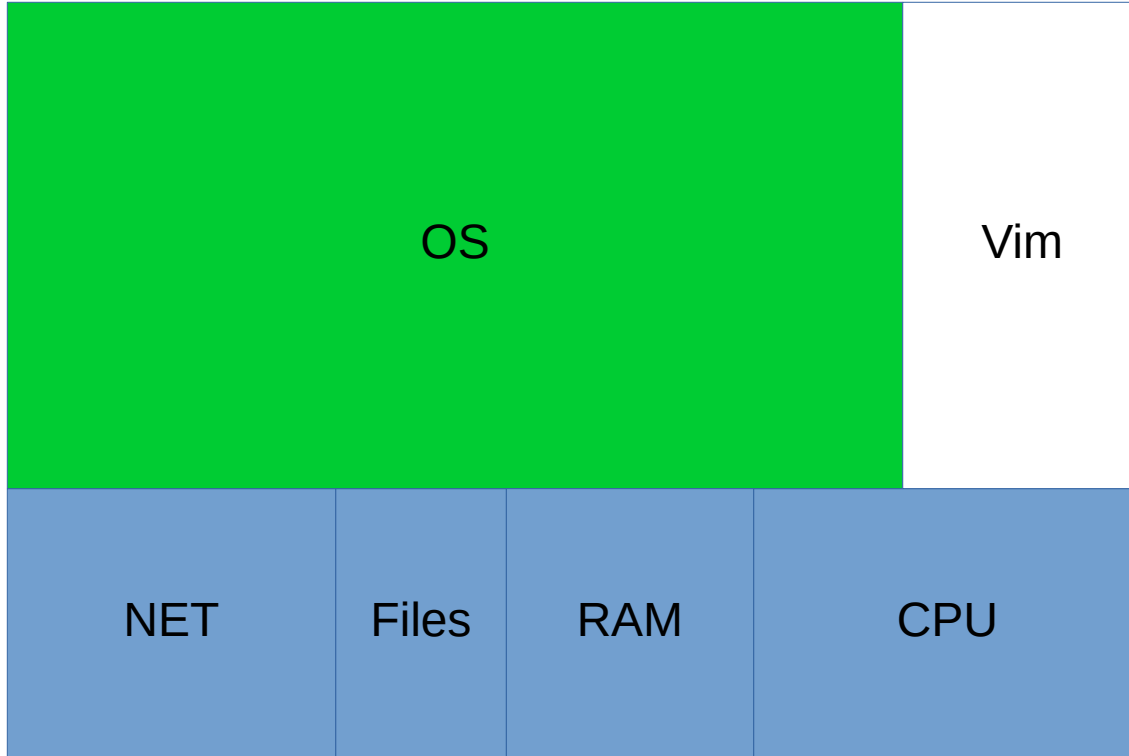
- Prozesse
 - Scheduling, Preemption, Schlafen
 - Syscalls
- VFS, Dateisysteme und Block IO
- Memory Management
- Network
- Interrupts, Treiber, etc.
- Locking

Überblick (Was)

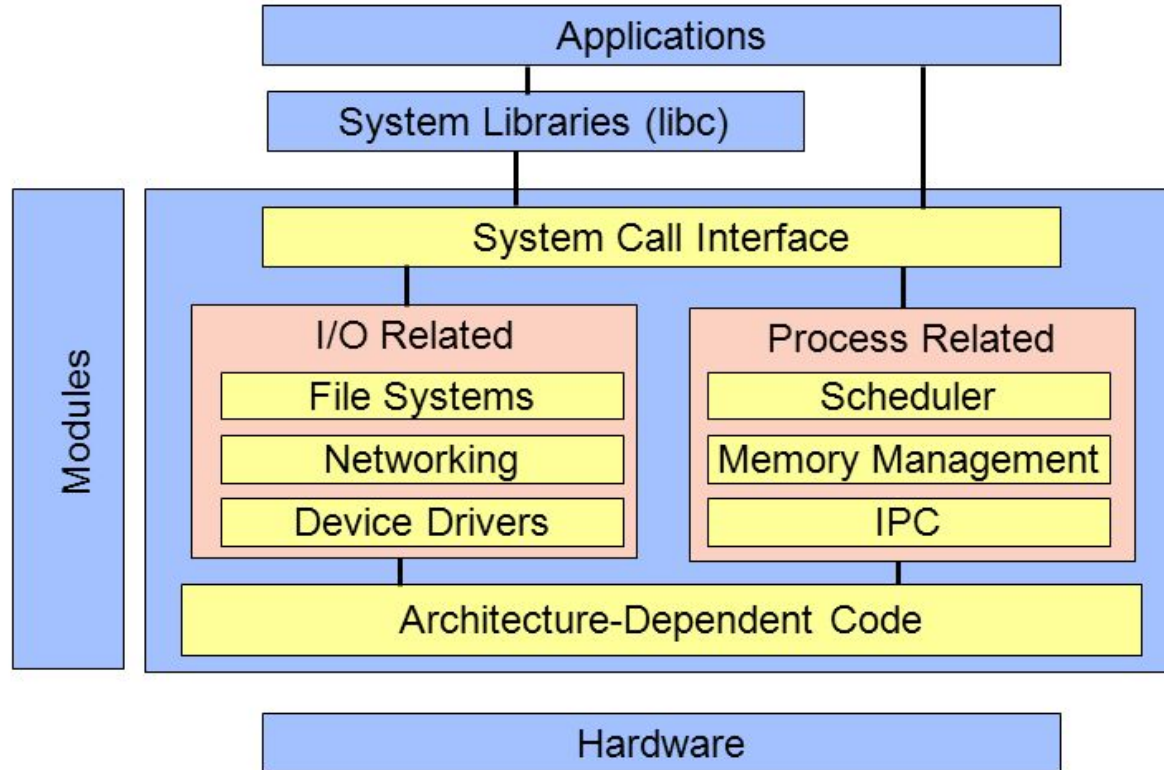
- 1) Was macht der Kernel überhaupt?
- 2) System Calls und Scheduling
- 3) Interrupts
- 4) Memory?

Gute Quellen

- [Linux Kernel Development](#) von Robert Love
- [Unreliable Guide To Hacking The Linux Kernel](#)
(einfach googlen, ist im linux kernel source repo dabei)



Linux Architecture



Linux kernel SCI (System Call Interface)

I/O subsystem

Linux kernel
Virtual File System

Terminals

Sockets

File systems

Netfilter / Nftables

Network
protocols

Generic
block layer

Linux kernel
I/O Scheduler

Linux kernel
Packet Scheduler

Character
device
drivers

Network
device
drivers

Block
device
drivers

Line
discipline

Memory management subsystem

Virtual
memory

Paging
page
replacement

Page
cache

Process management subsystem

Signal
handling

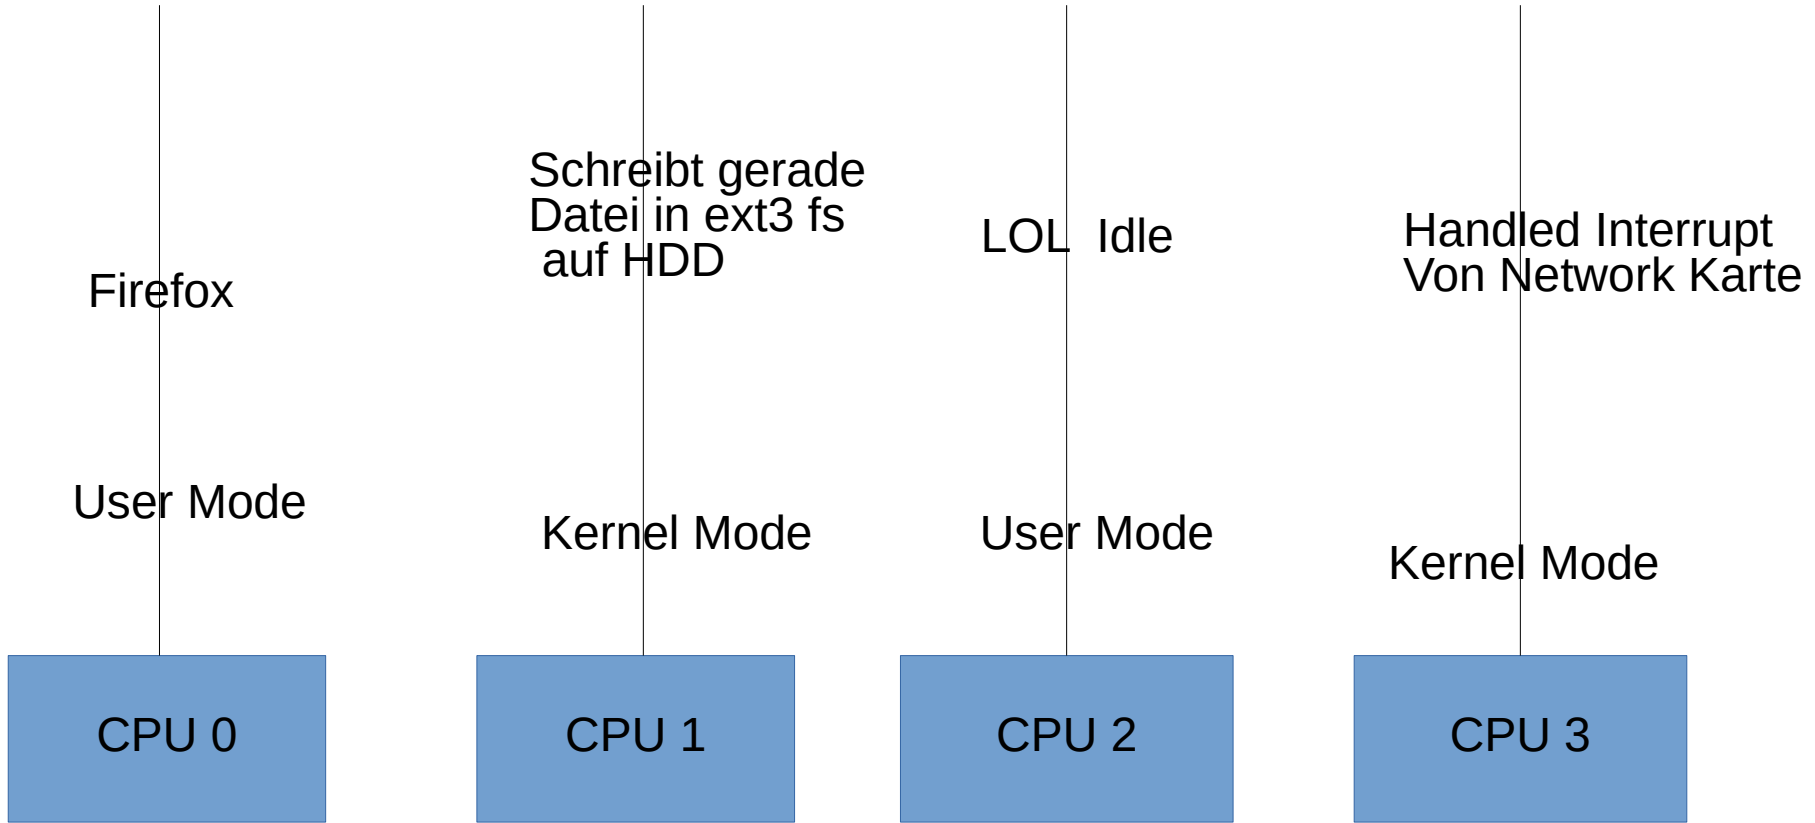
process/thread
creation &
termination

Linux kernel
Process
Scheduler

IRQs

Dispatcher

Und wie ist das bei mehreren
CPU-Cores?




User Mode vs. Kernel Mode

Code in User Mode:


- Darf nichts

Code in Kernel Mode:

- Darf alles



Least privileged



Most privileged

Mode Transitions

- Sind billig
- User Mode → Kernel Mode
 - (Syscall)
- Kernel Mode → User Mode
 - Return von syscall
 - Prozess der vom scheduler unterbrochen wurde wrd wieder gescheduled

Process Context vs. Interrupt Context

Code im Interrupt Kontext

- Darf nicht schlafen
- Man will da nur kurz sein
- Eigener Stack
- Unabhängig von einem Prozess

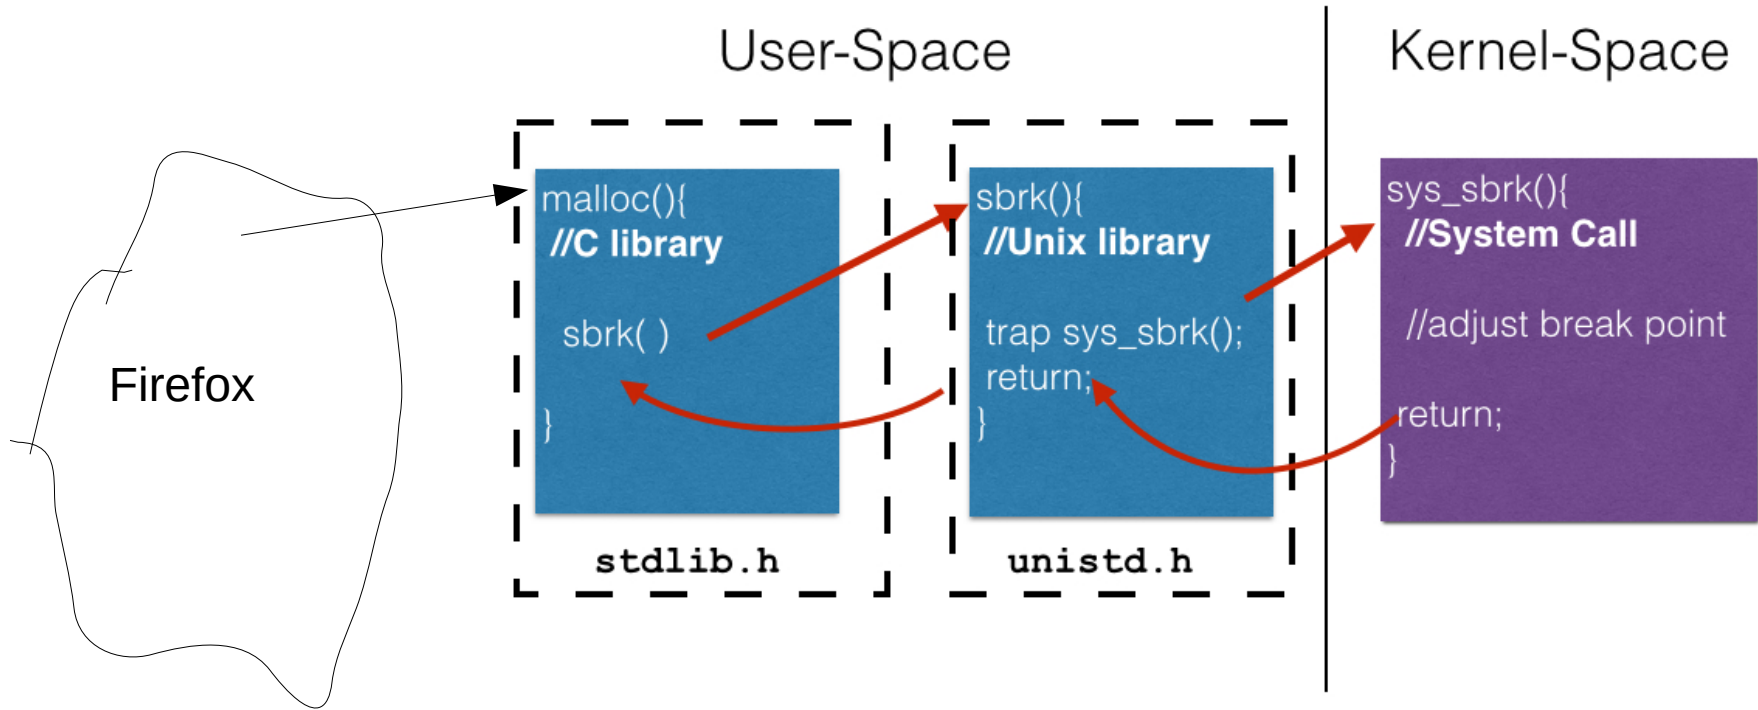
Prozess Kontext

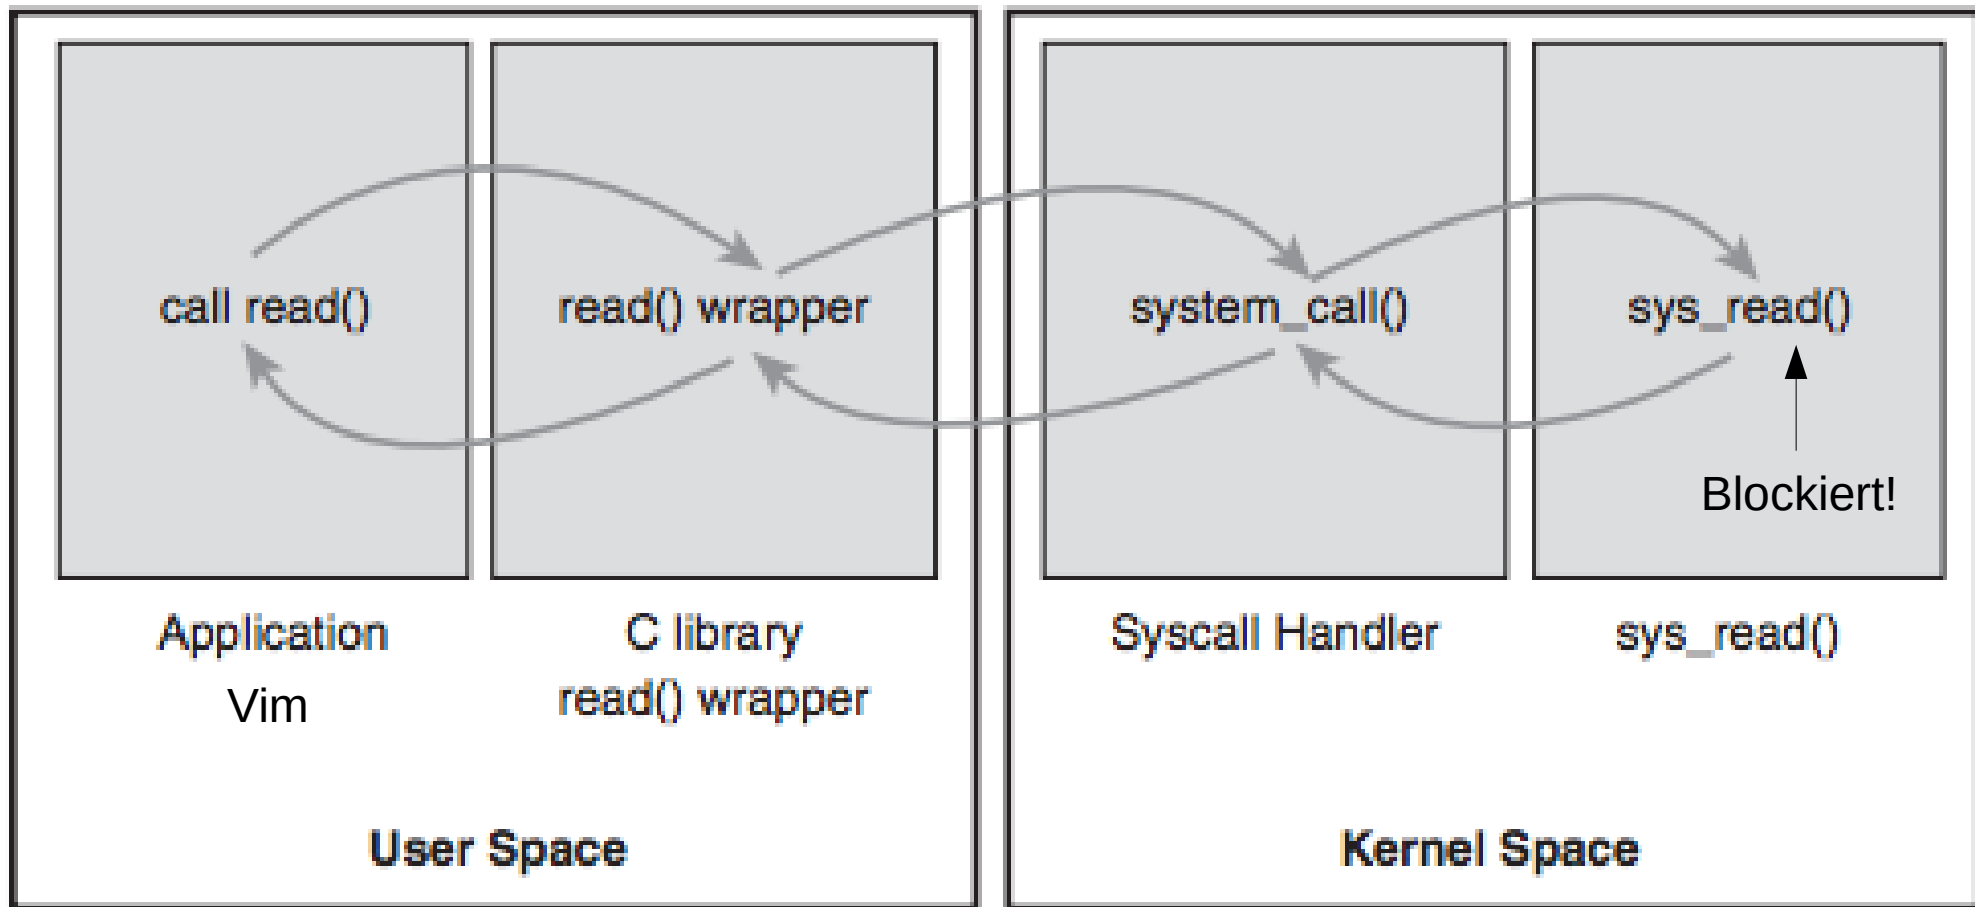
- Man kann und darf schlafen
- Relatiert zu einem Prozess
- Man muss nicht bresieren

Wie hängt das zusammen?

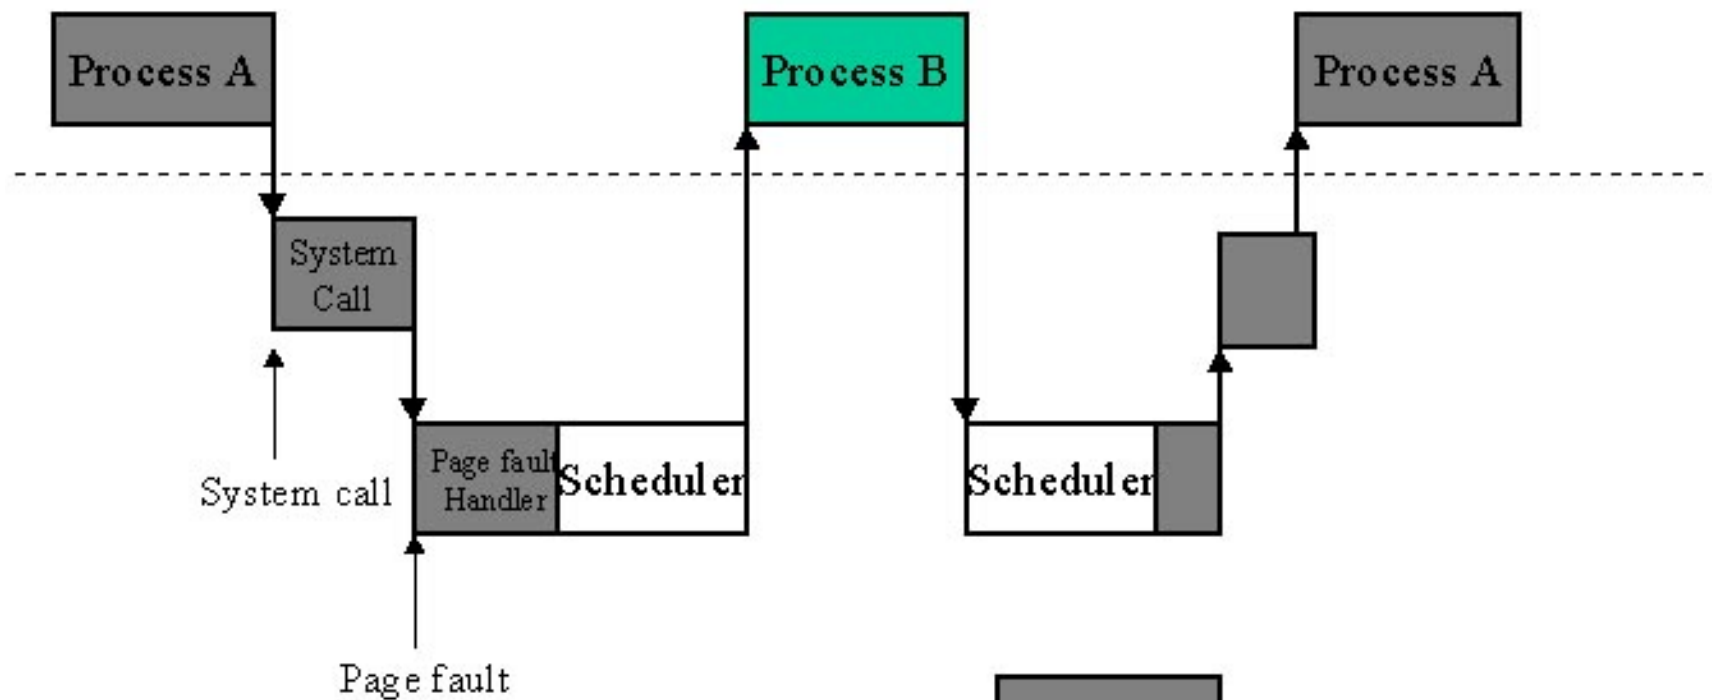
	User Mode	Kernel Mode
Process Kontext	Normaler Prozess	Syscalls, Kernel Threads, großteil der "Kernel Arbeit"
Interrupt Kontext	Gibts net	Interrupt Handler

Syscalls






User mode



Kernel mode



Context Switch (Process Switch)

Enthält:

- 2 Mode Transitions
- Register und Prozess State Saven
- Memory Mapping Switch

- Also:
- Caches kühlen ab
- TLB (Translation Lookaside Buffer) kühlt ab

Thread Switch

System Timer
(fires)

Tick

Jiffy

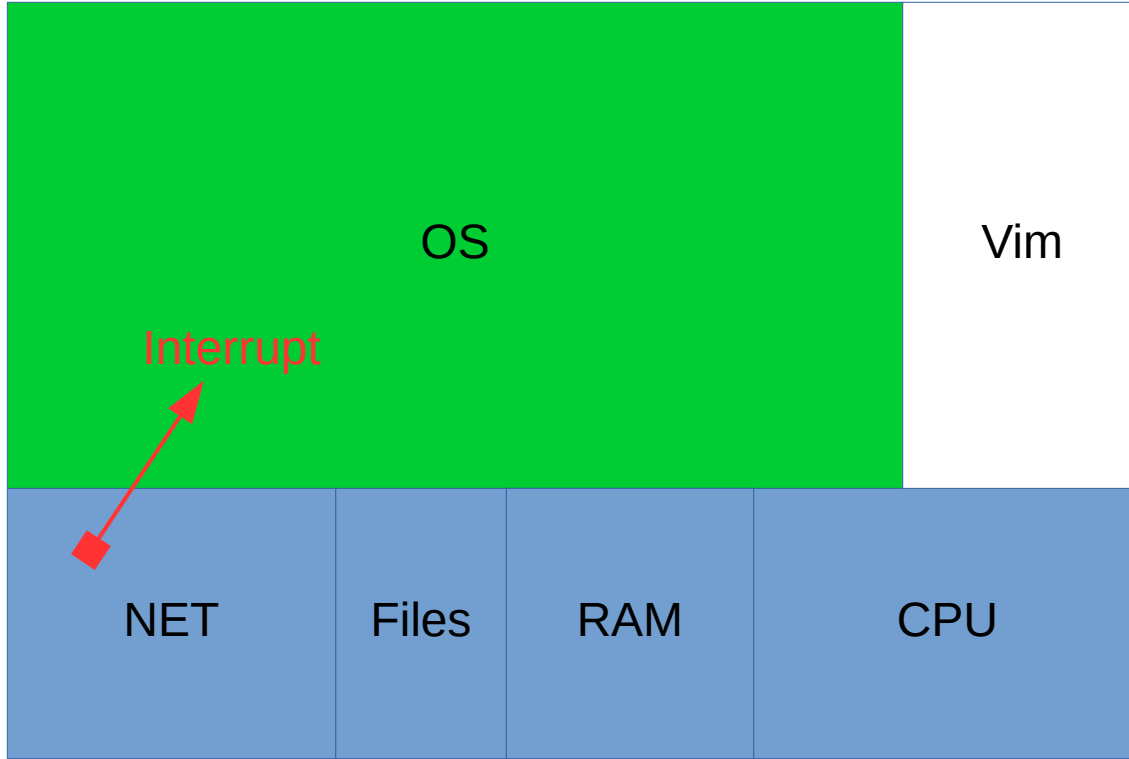
250 Hz

Interrupts (IRQs)

Dinge die Interrupts senden:

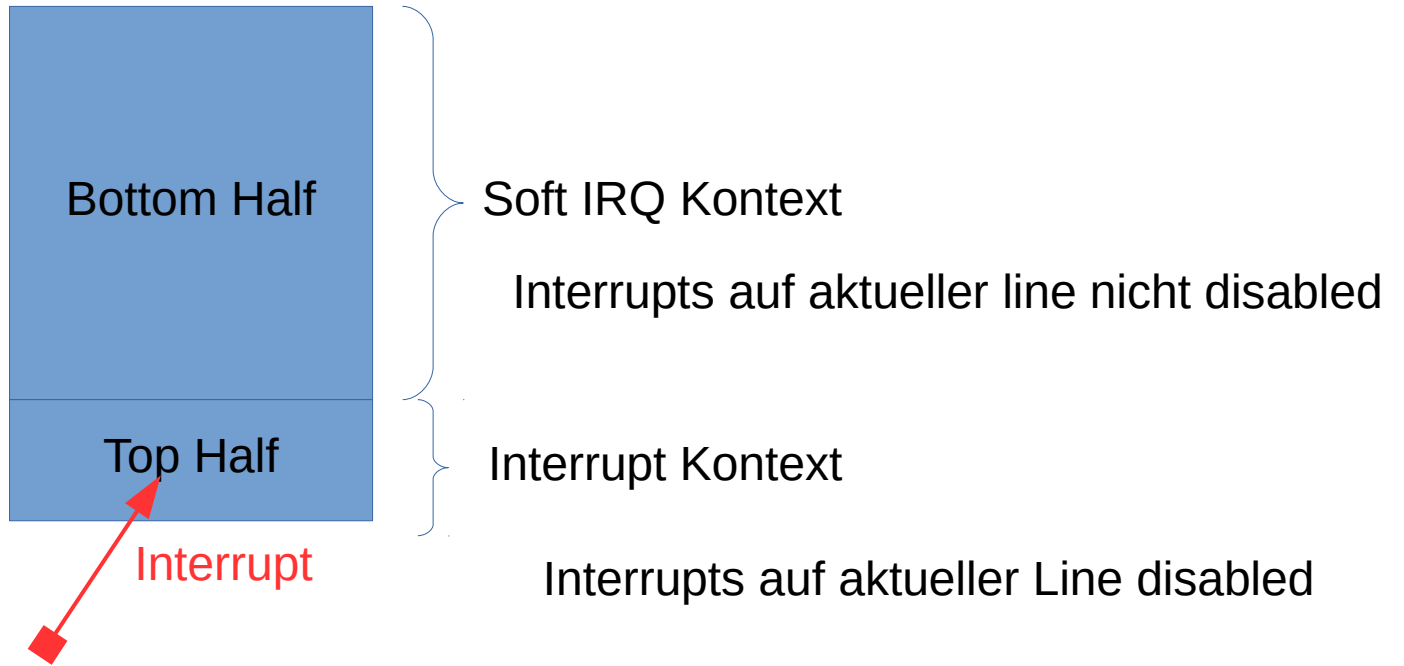
- Maus
- Tastatur
- Netzwerkkarte
- HDD?
- USB
- Hardware halt

/proc/interrupts



Interrupt Handler (ISRs)

Top Half vs. Bottom Half

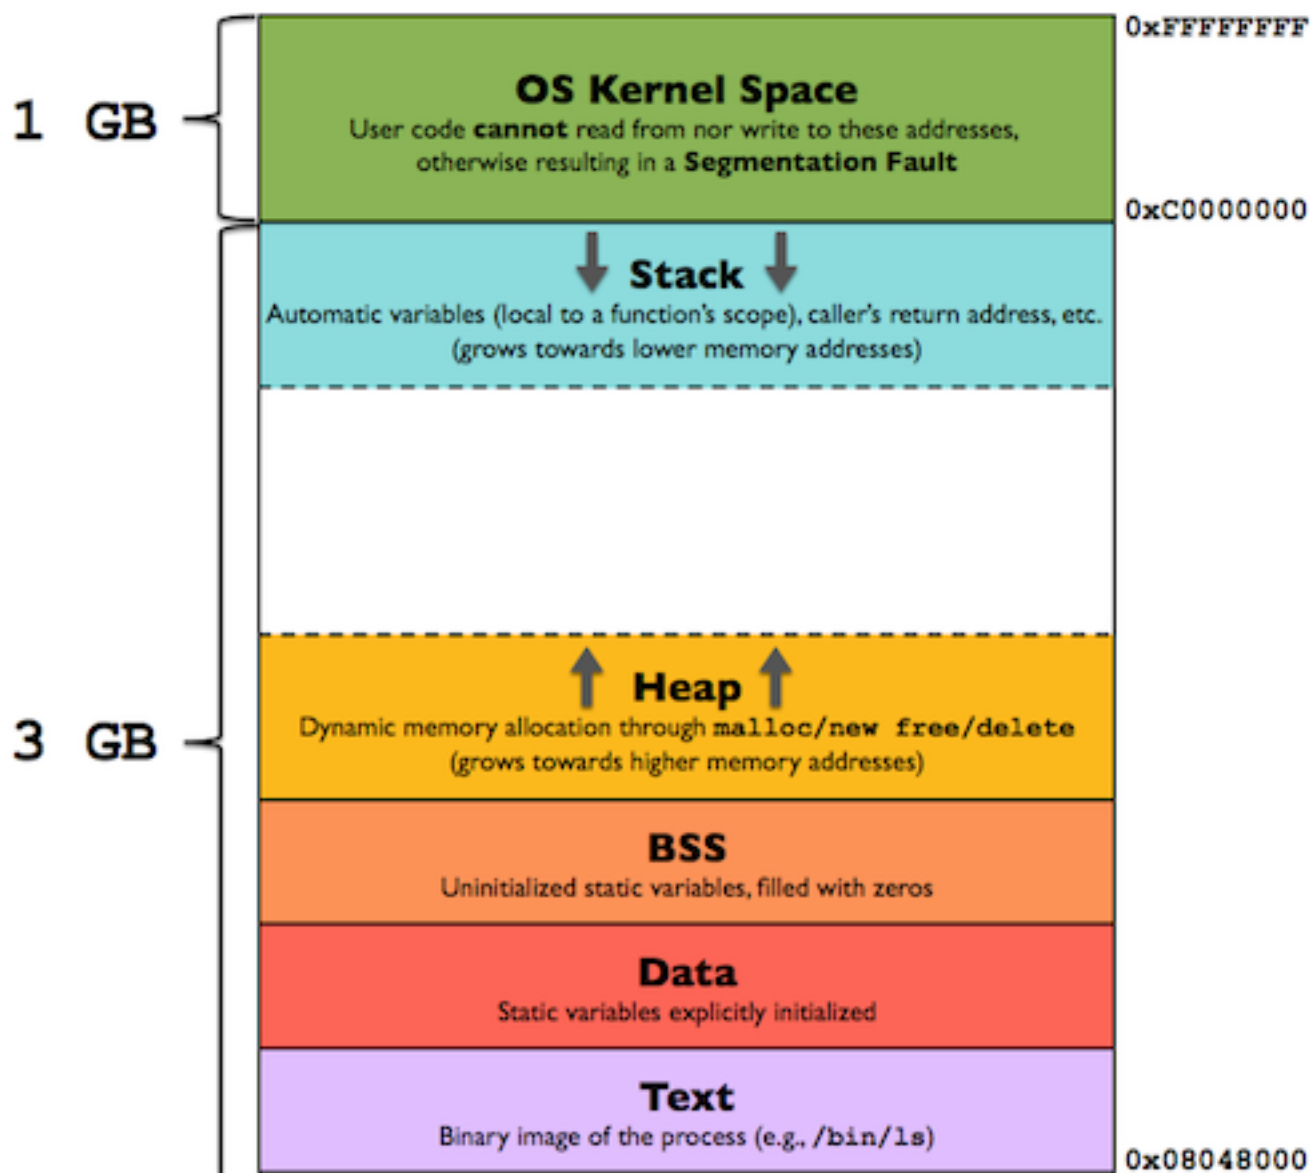


Bottom Halves

- Soft IRQs
- Tasklets

- Work Queues

Linux Memory Managment



Locking

Sleeping Locks
Mutex,
Semaphore

“Awake Locks” Spinlocks

Trace-cmd und kernelshark Beispiel

drivers/char/rtc.c