In-Guest Device Emulation

A Research Project

About Me

- Alexander Graf
- KVM developer at Amazon Web Services
  - Server class PowerPC KVM port
  - Nested SVM
- Opinions my own

Device Emulation

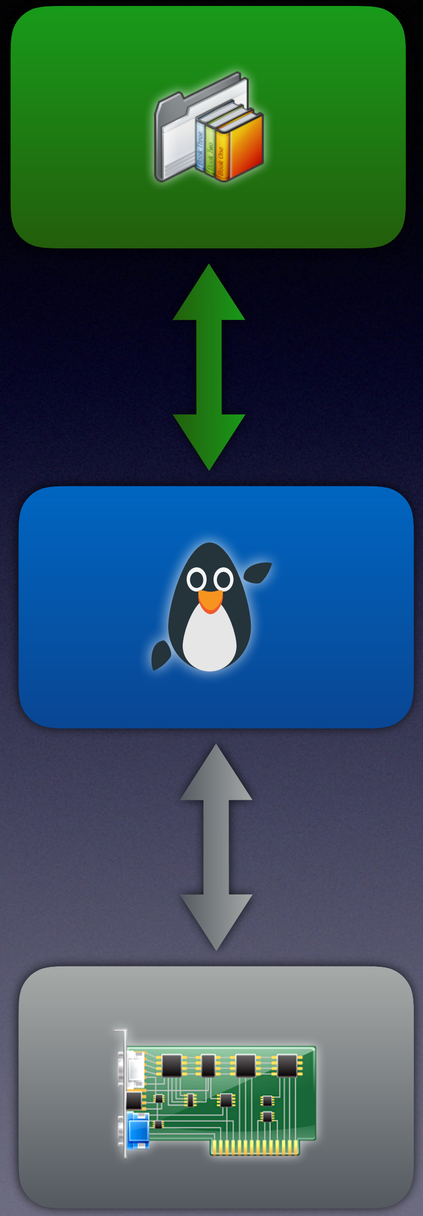

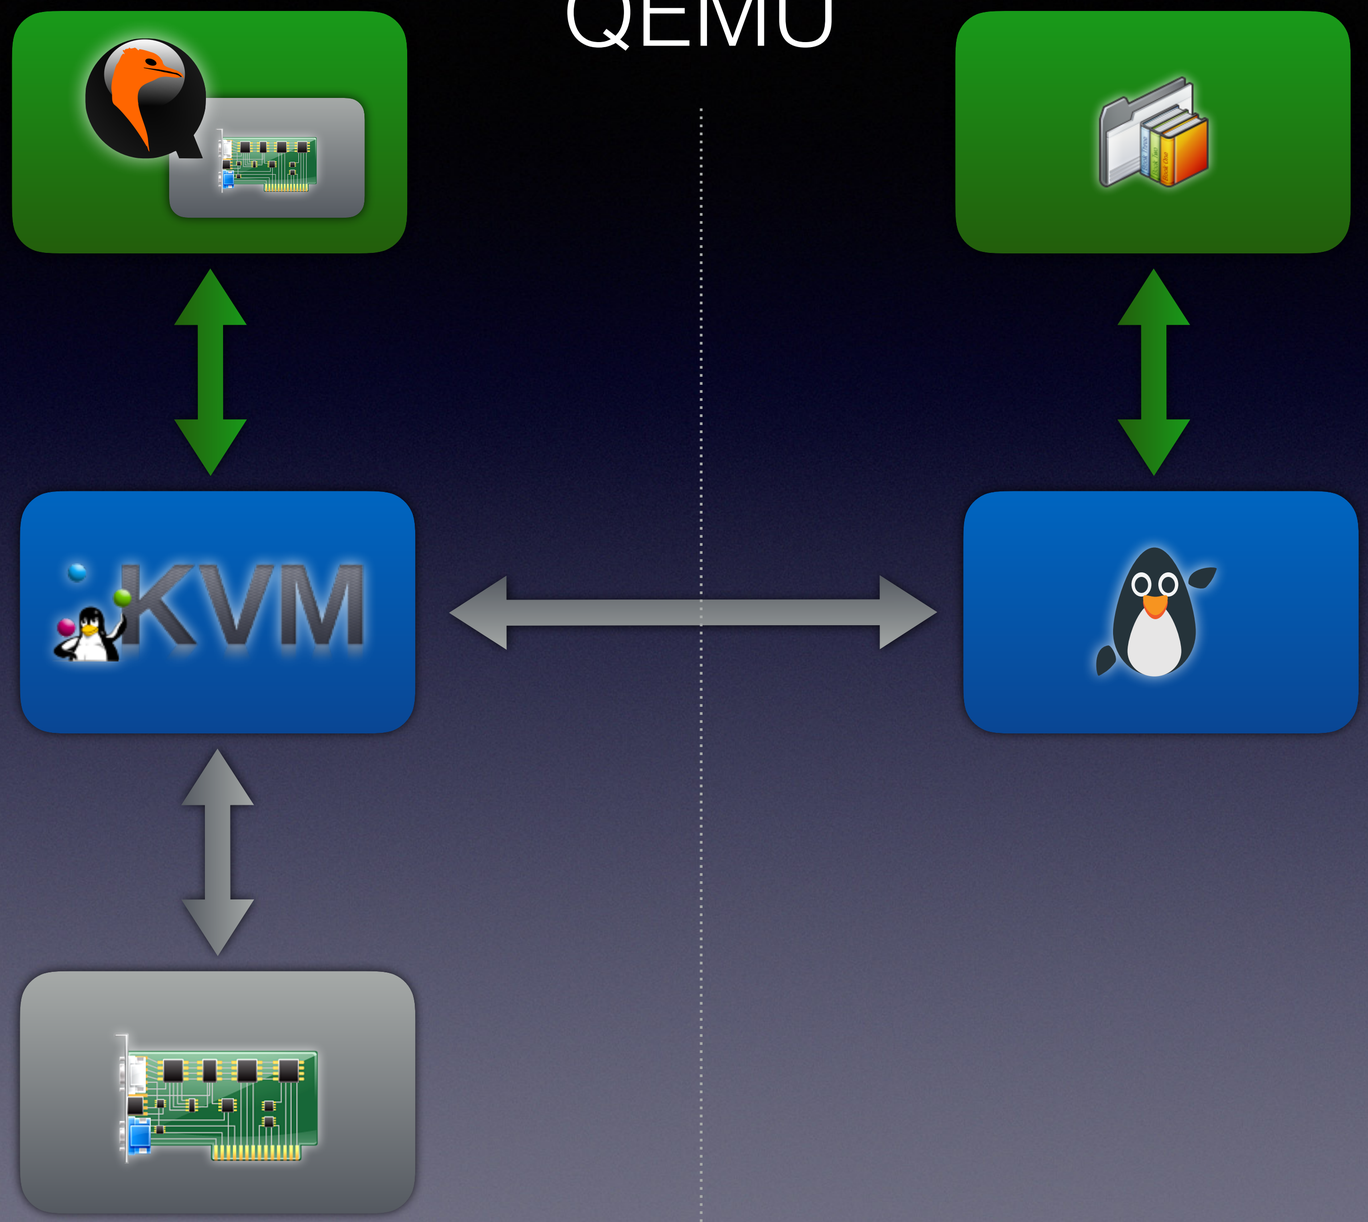
QEMU
KVM

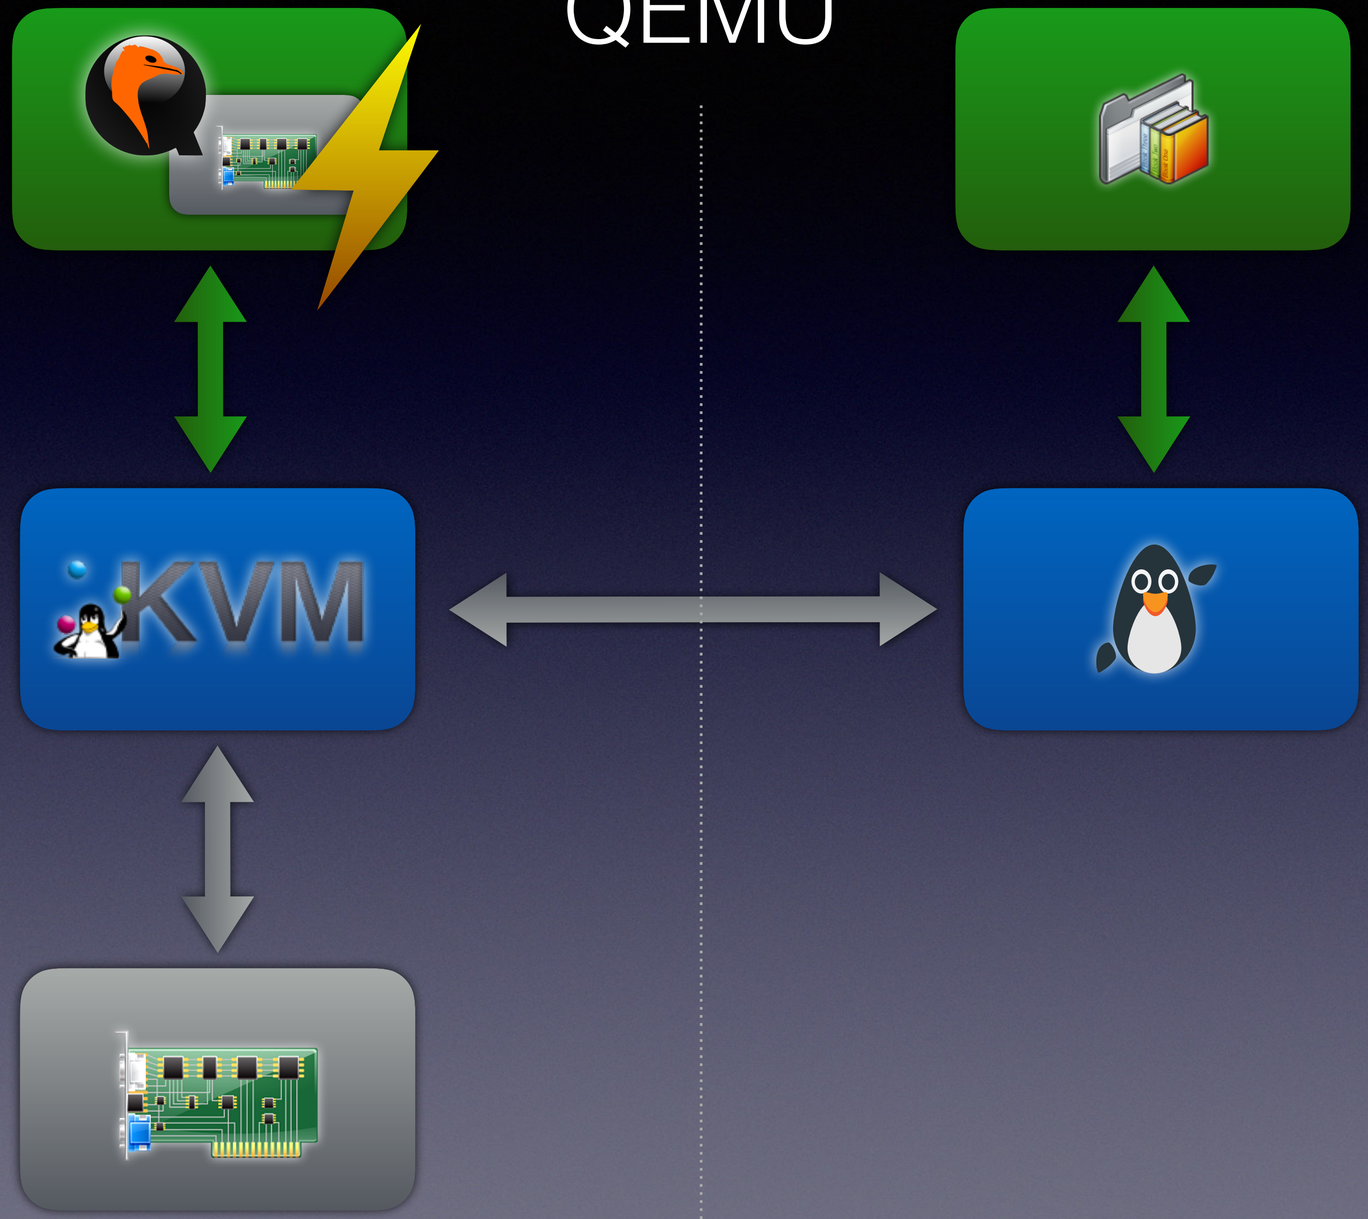
QEMU
KVM

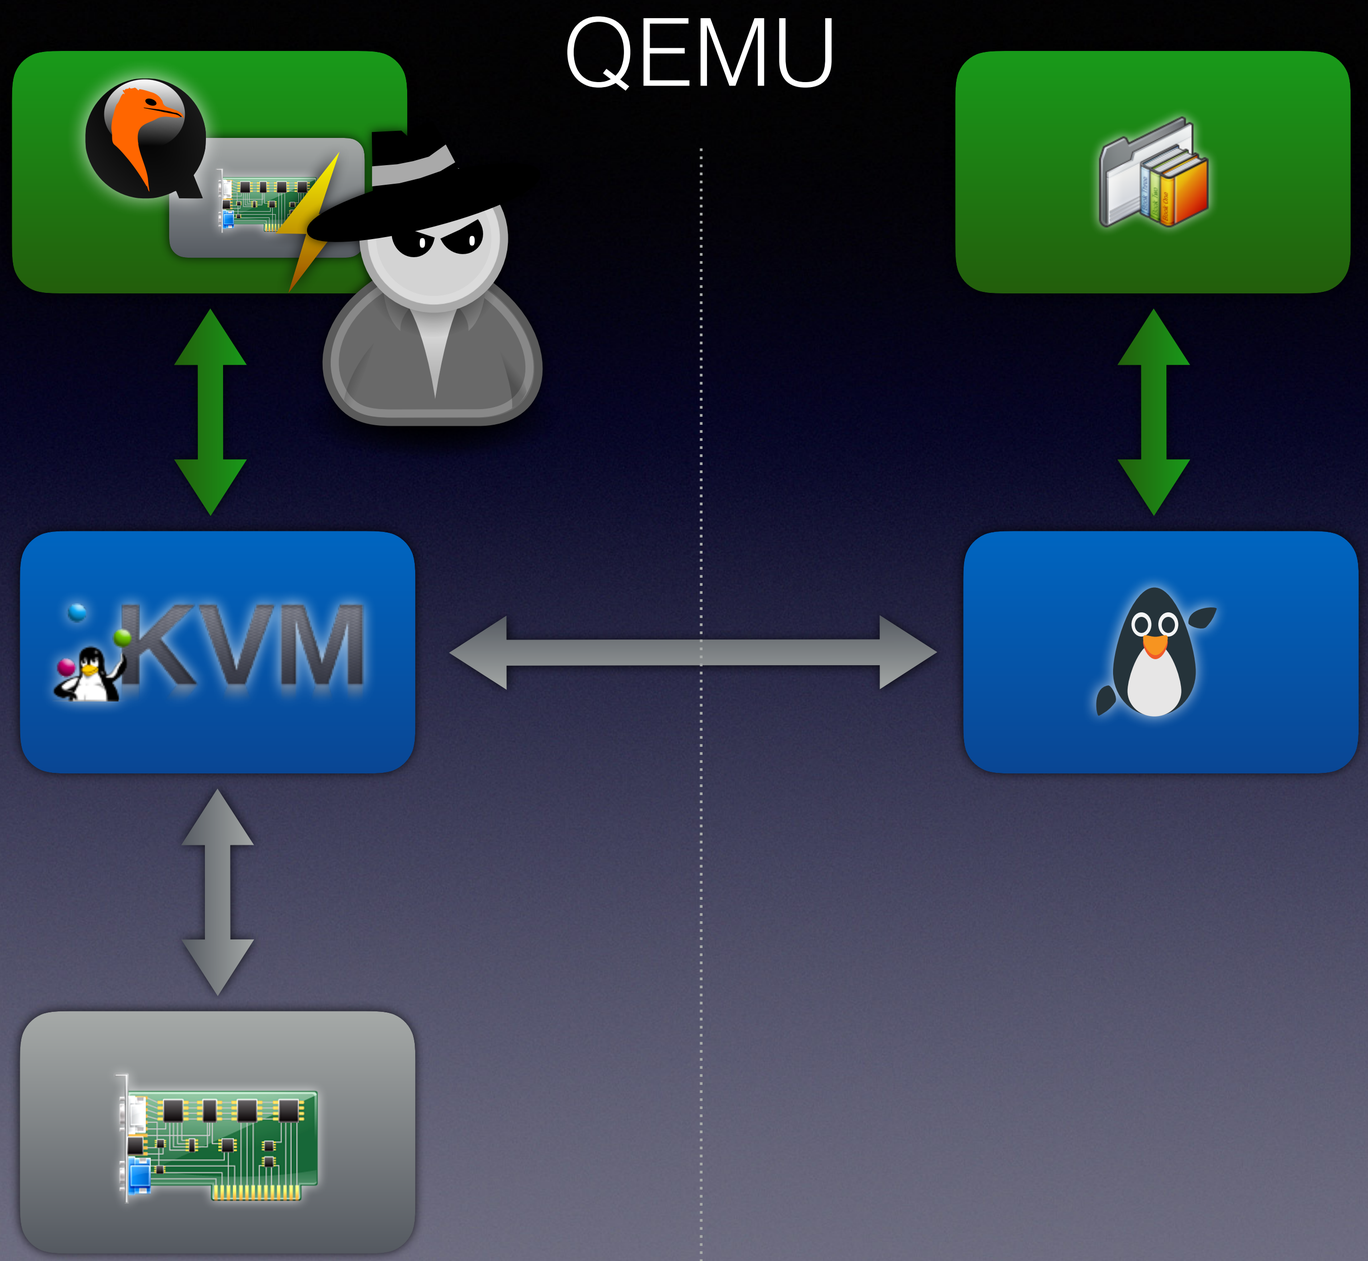
QEMU
KVM

SMM

SMM

Kernel

SMM

Kernel

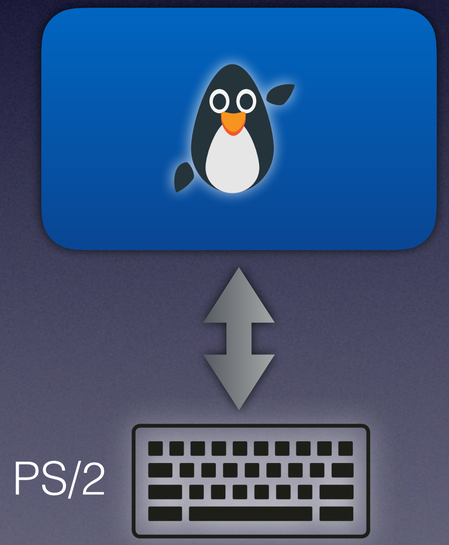

SMM

Kernel

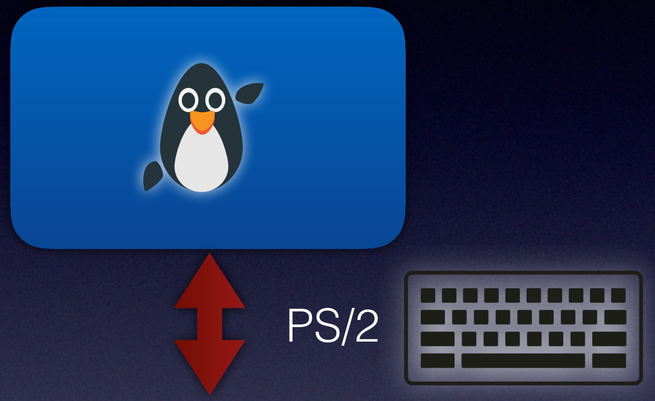

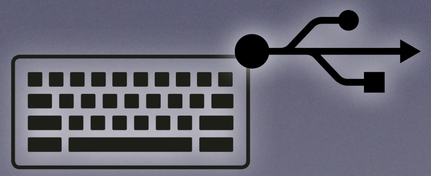

SMM

Kernel

SMM

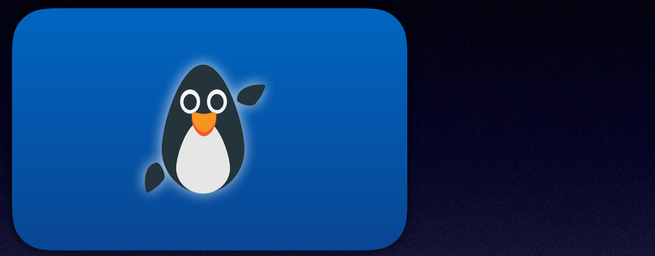

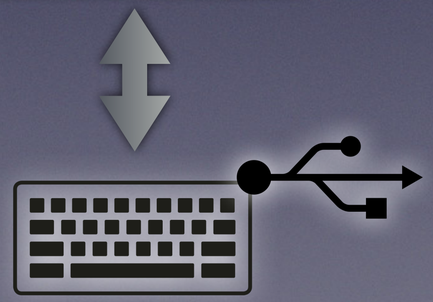

QEMU

Device Emulation

QEMU

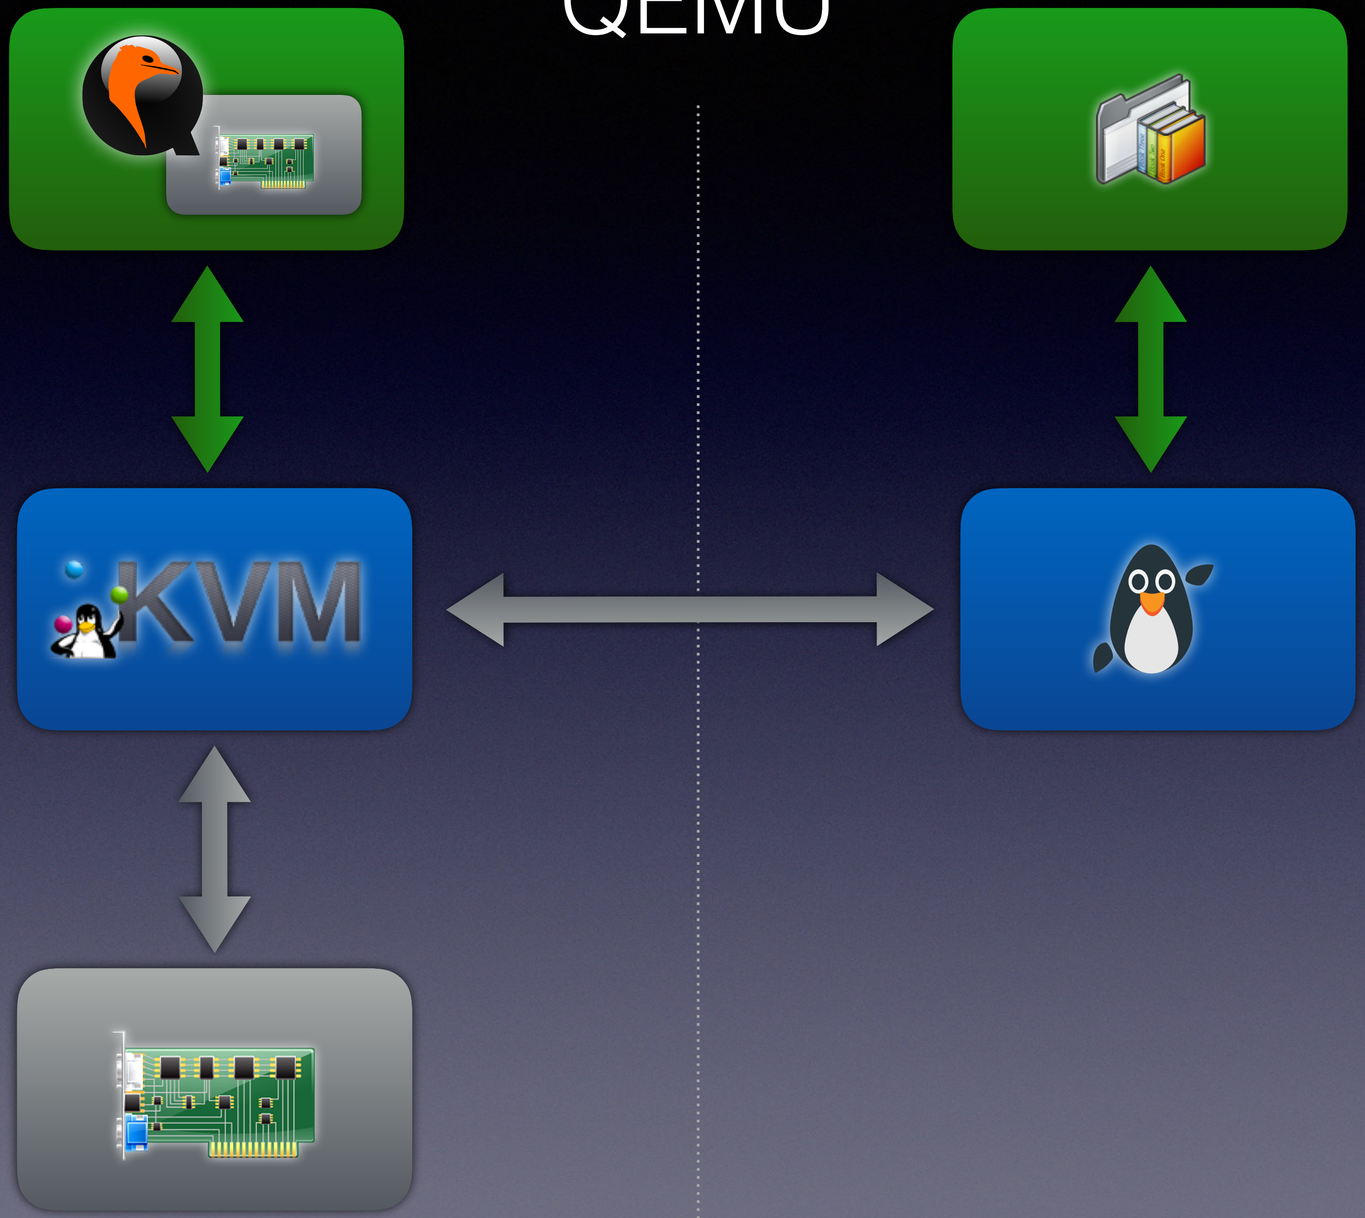

QEMU

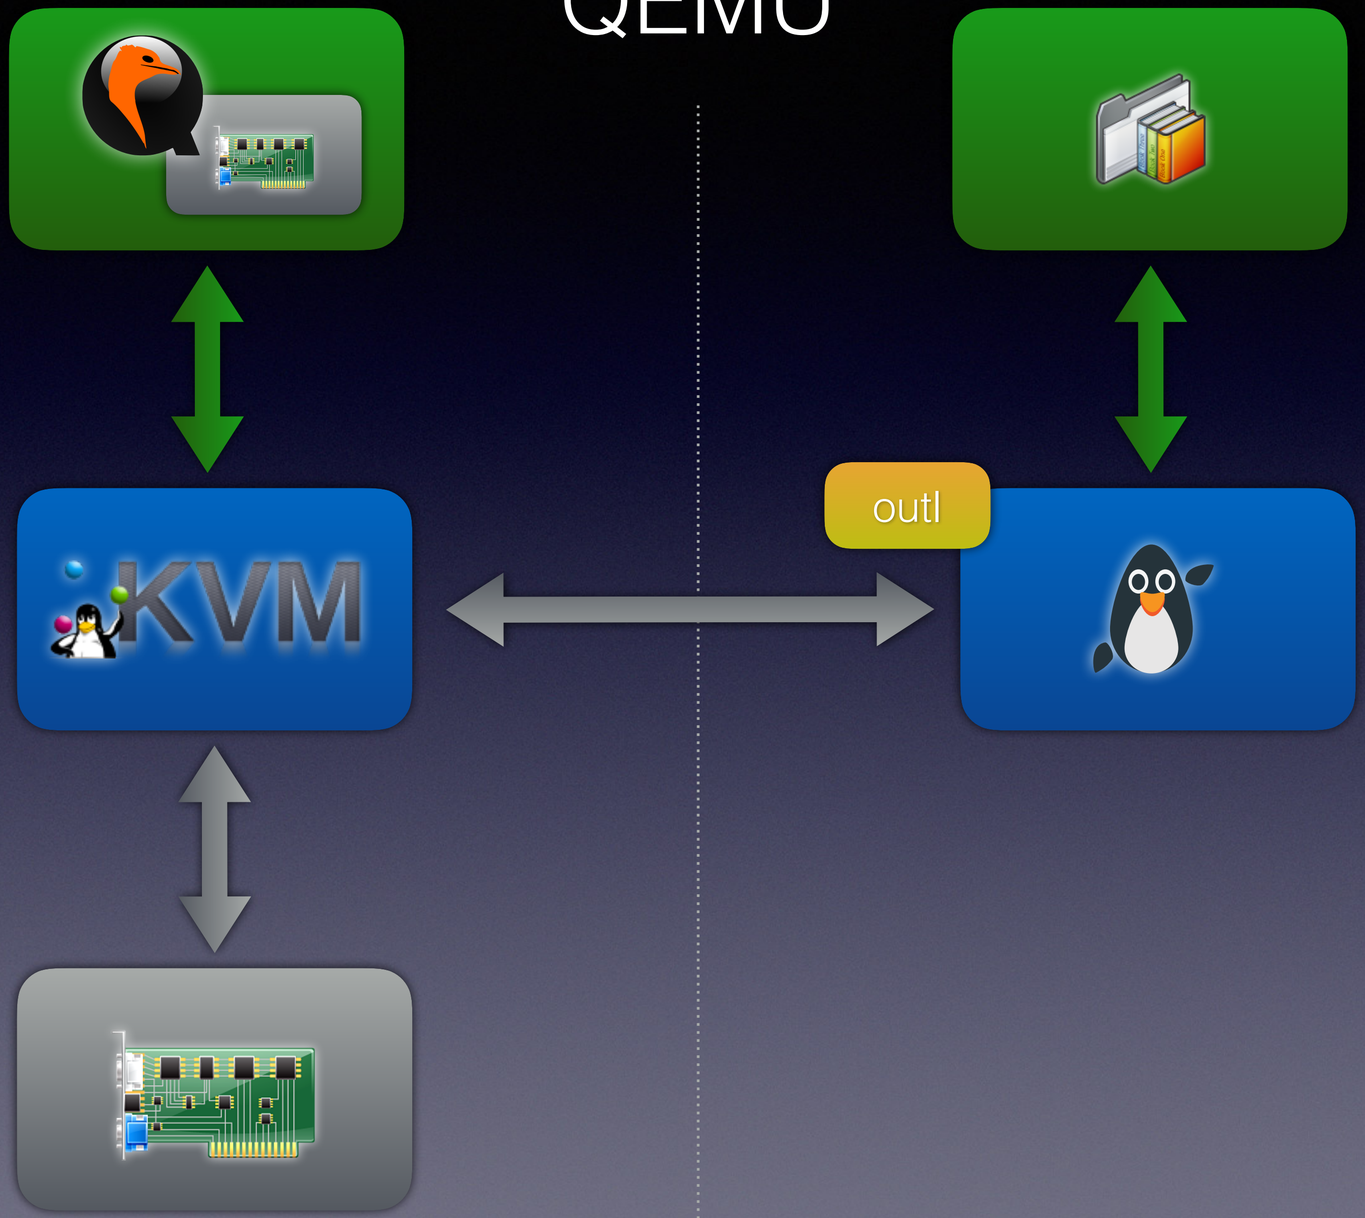

QEMU

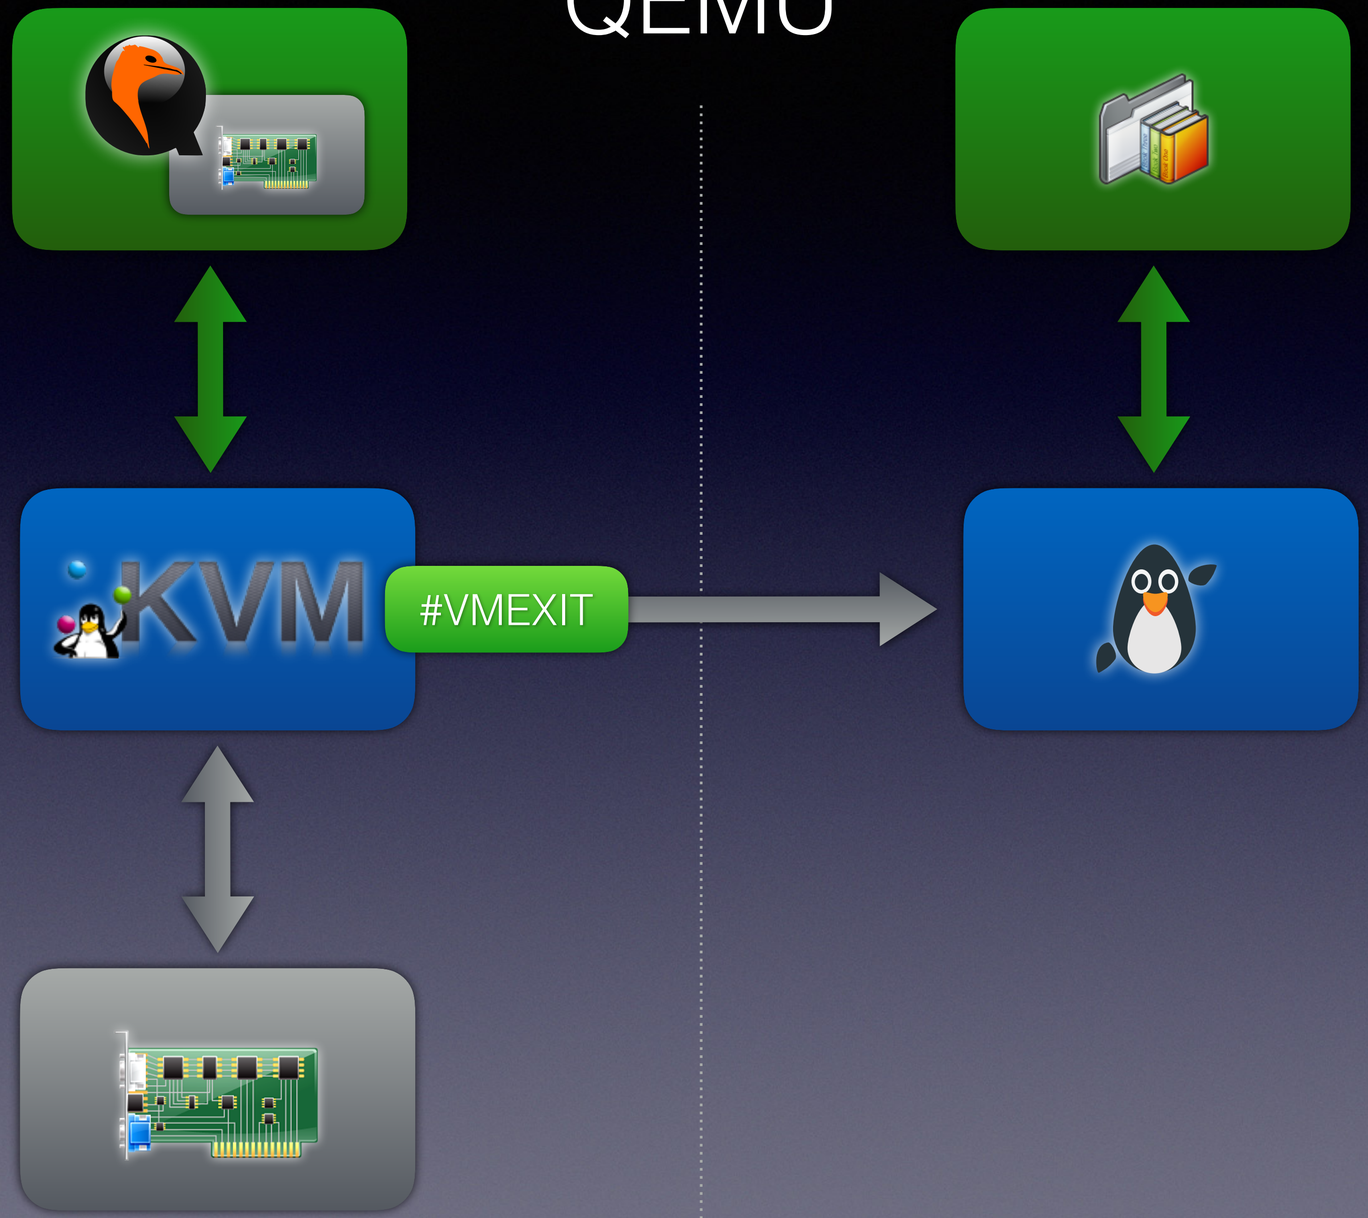

QEMU

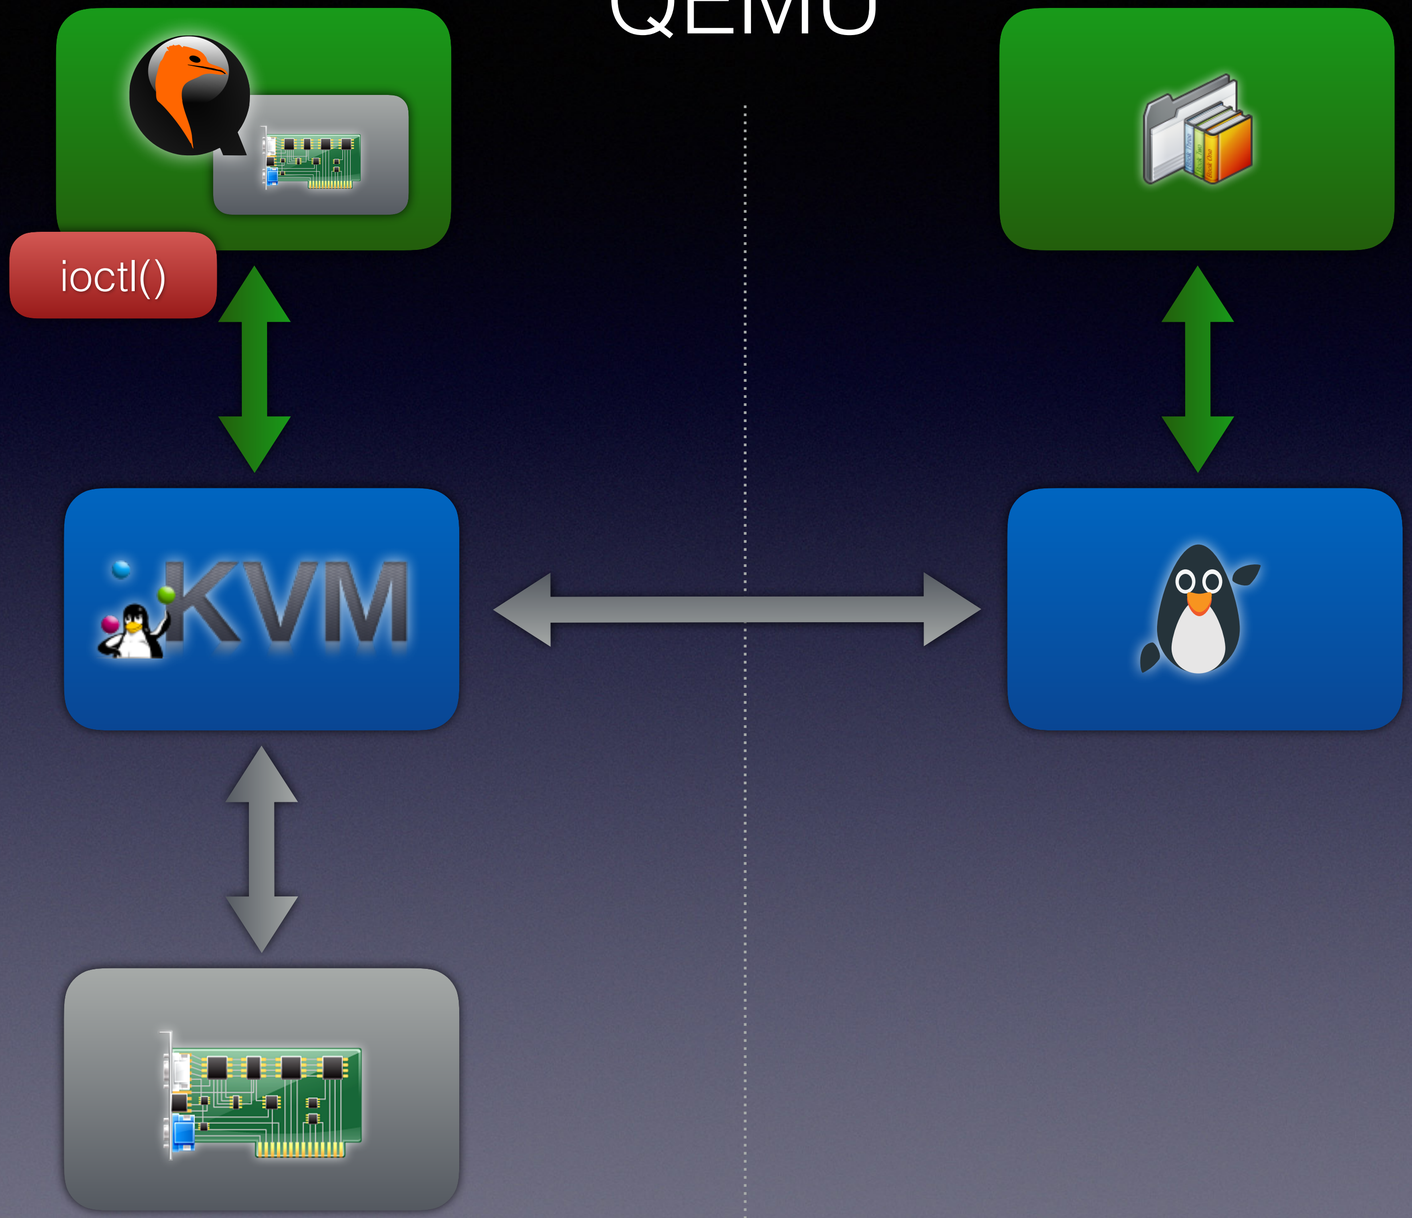

QEMU

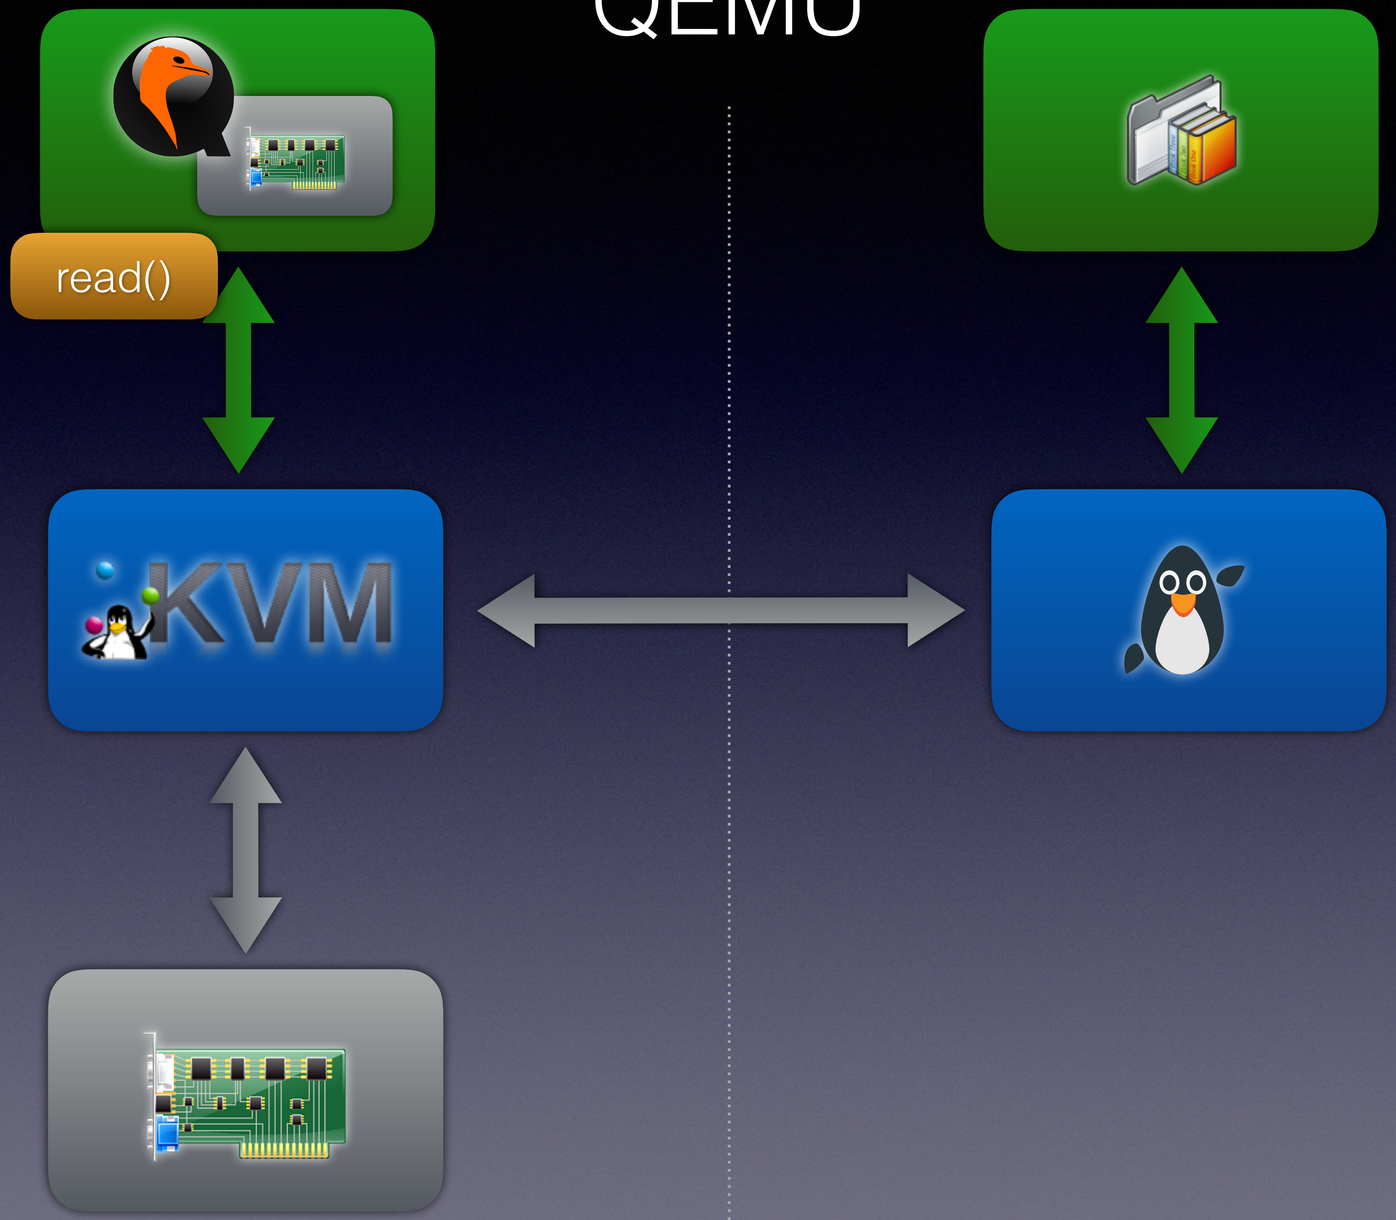

QEMU

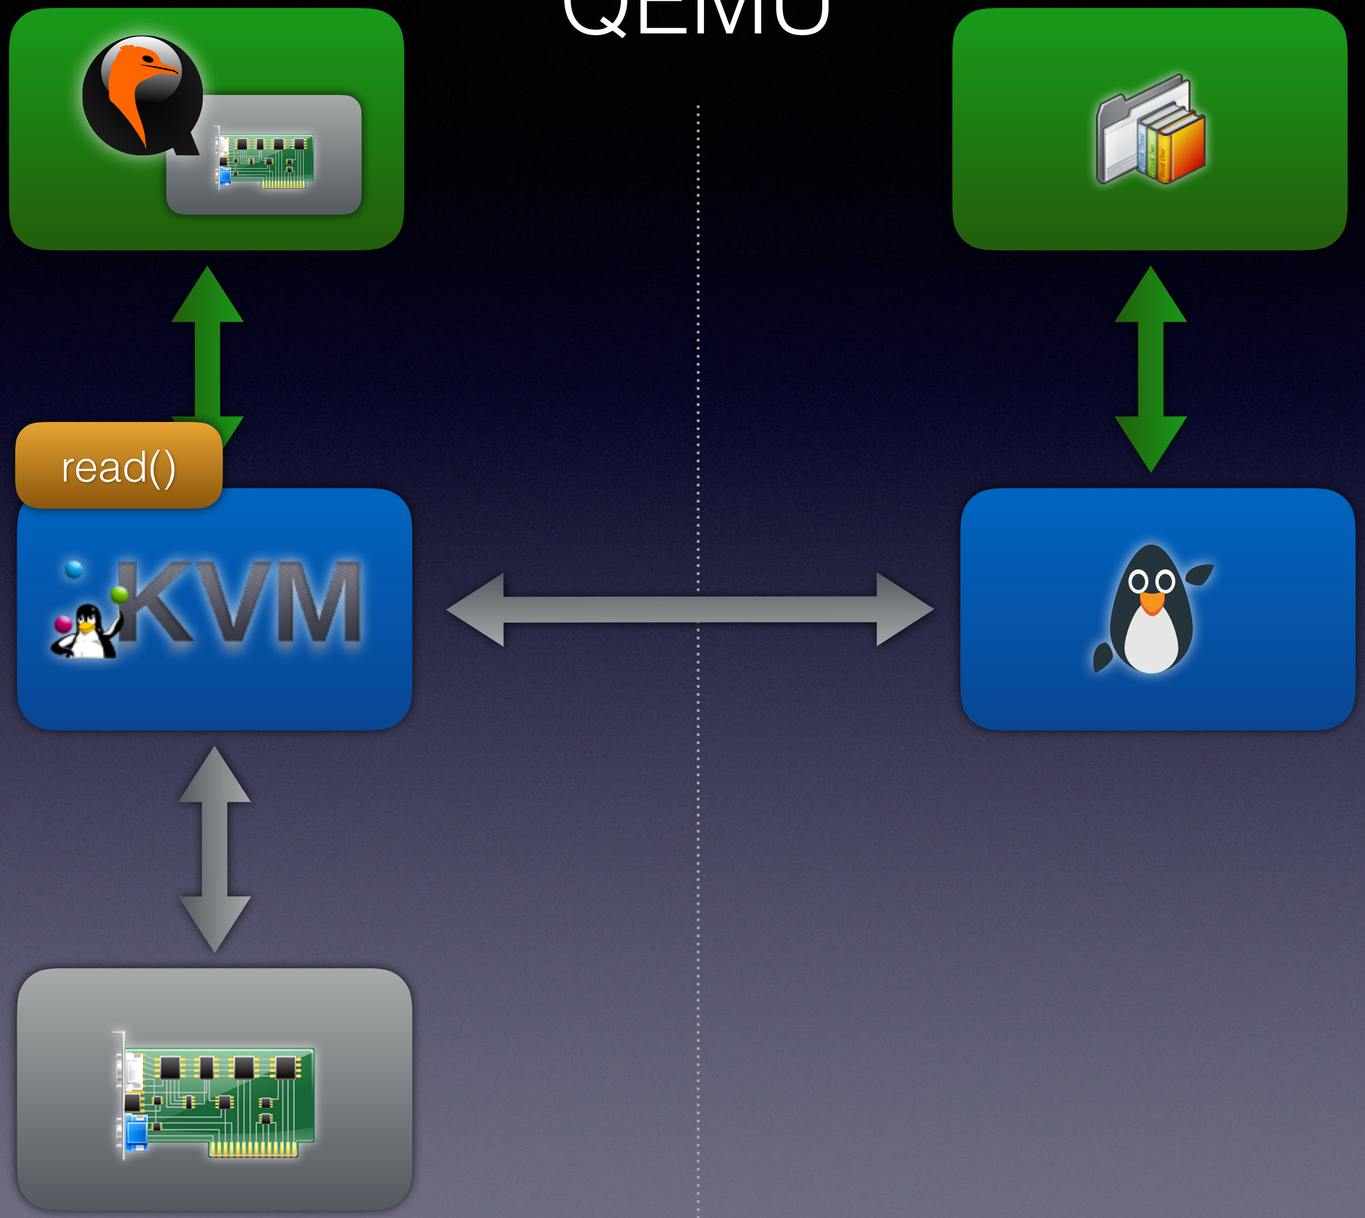

QEMU

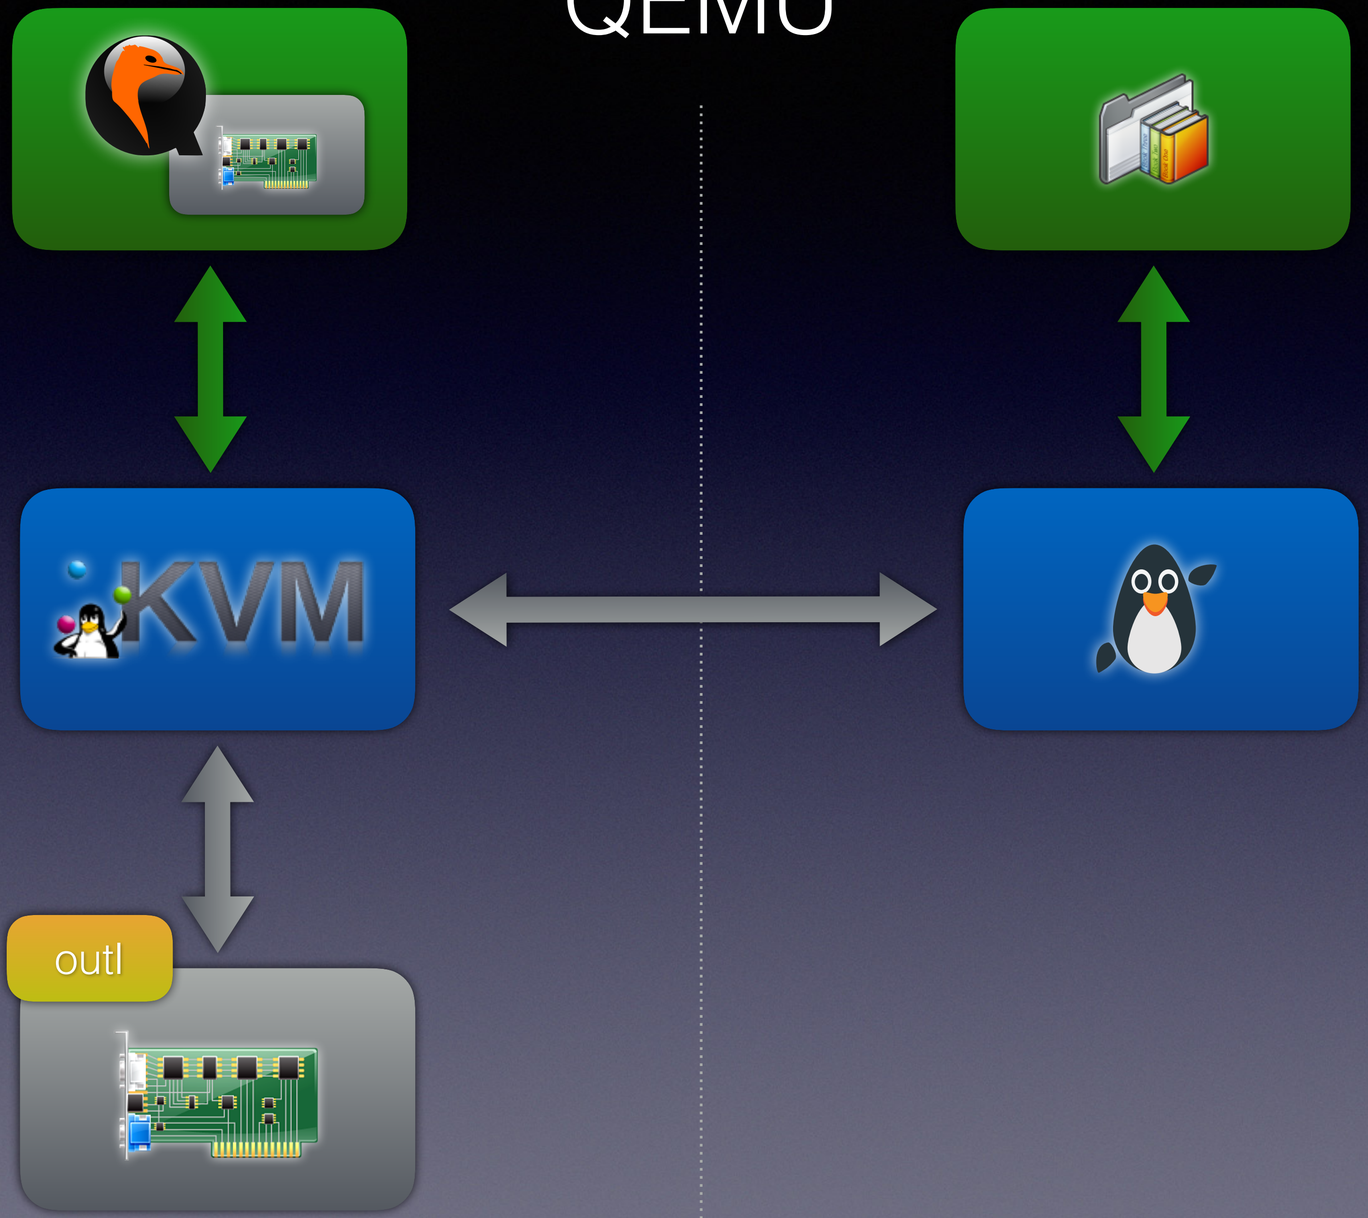

QEMU

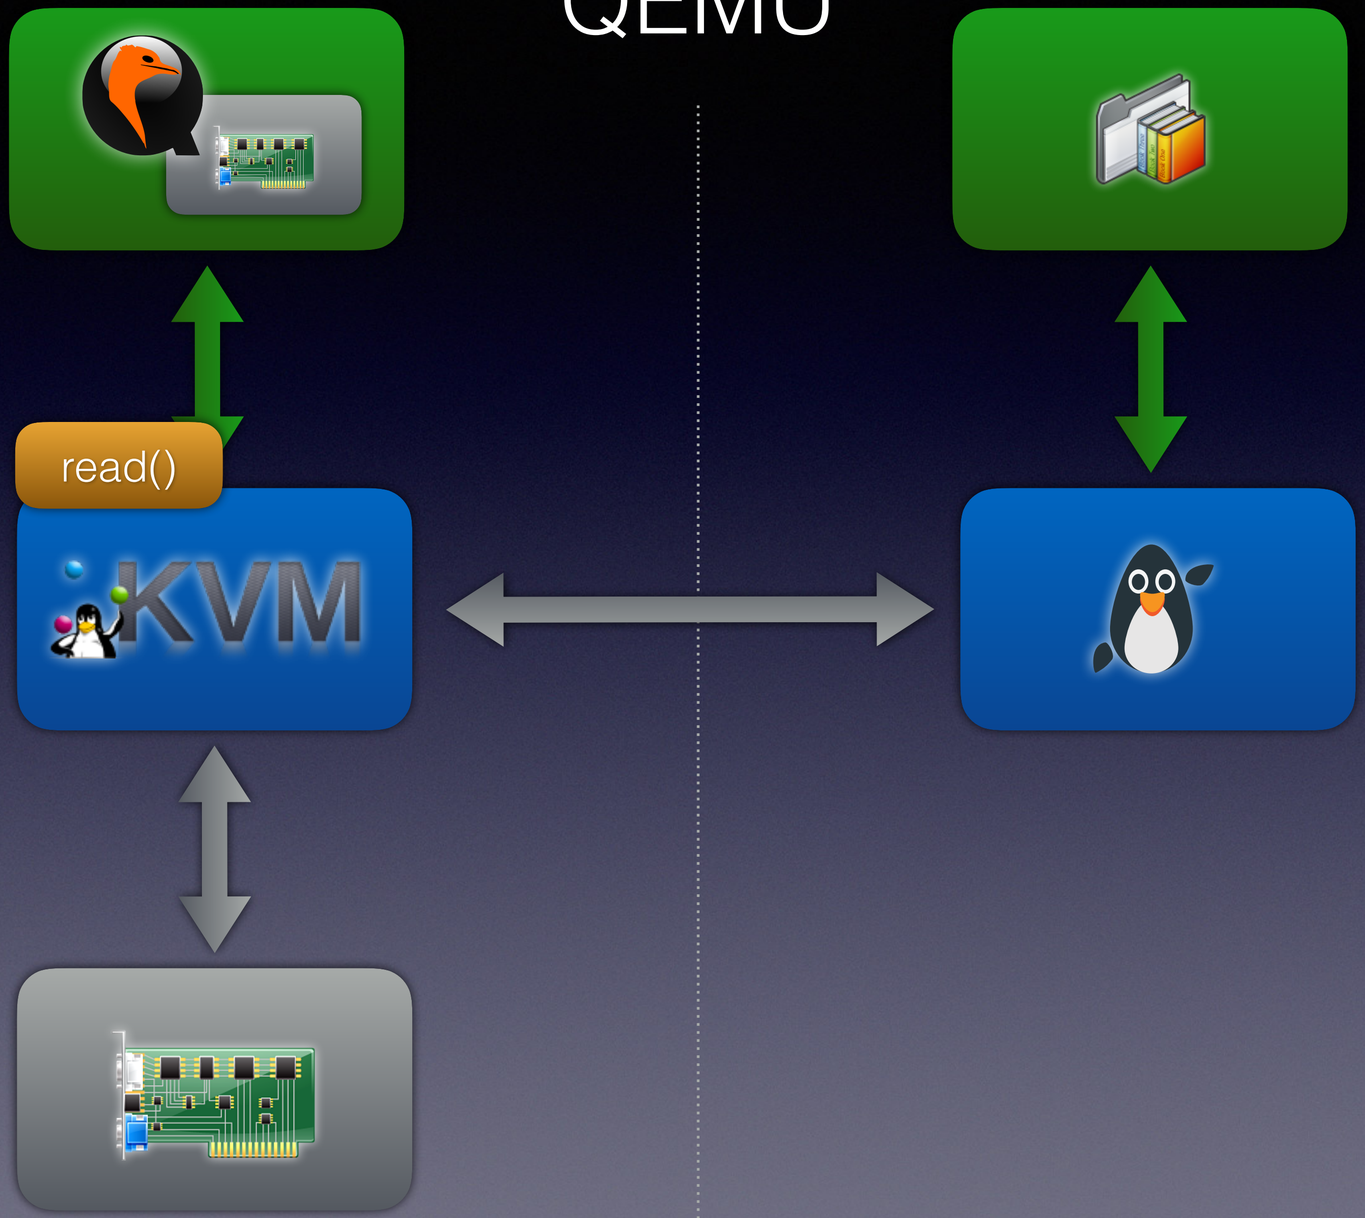

QEMU

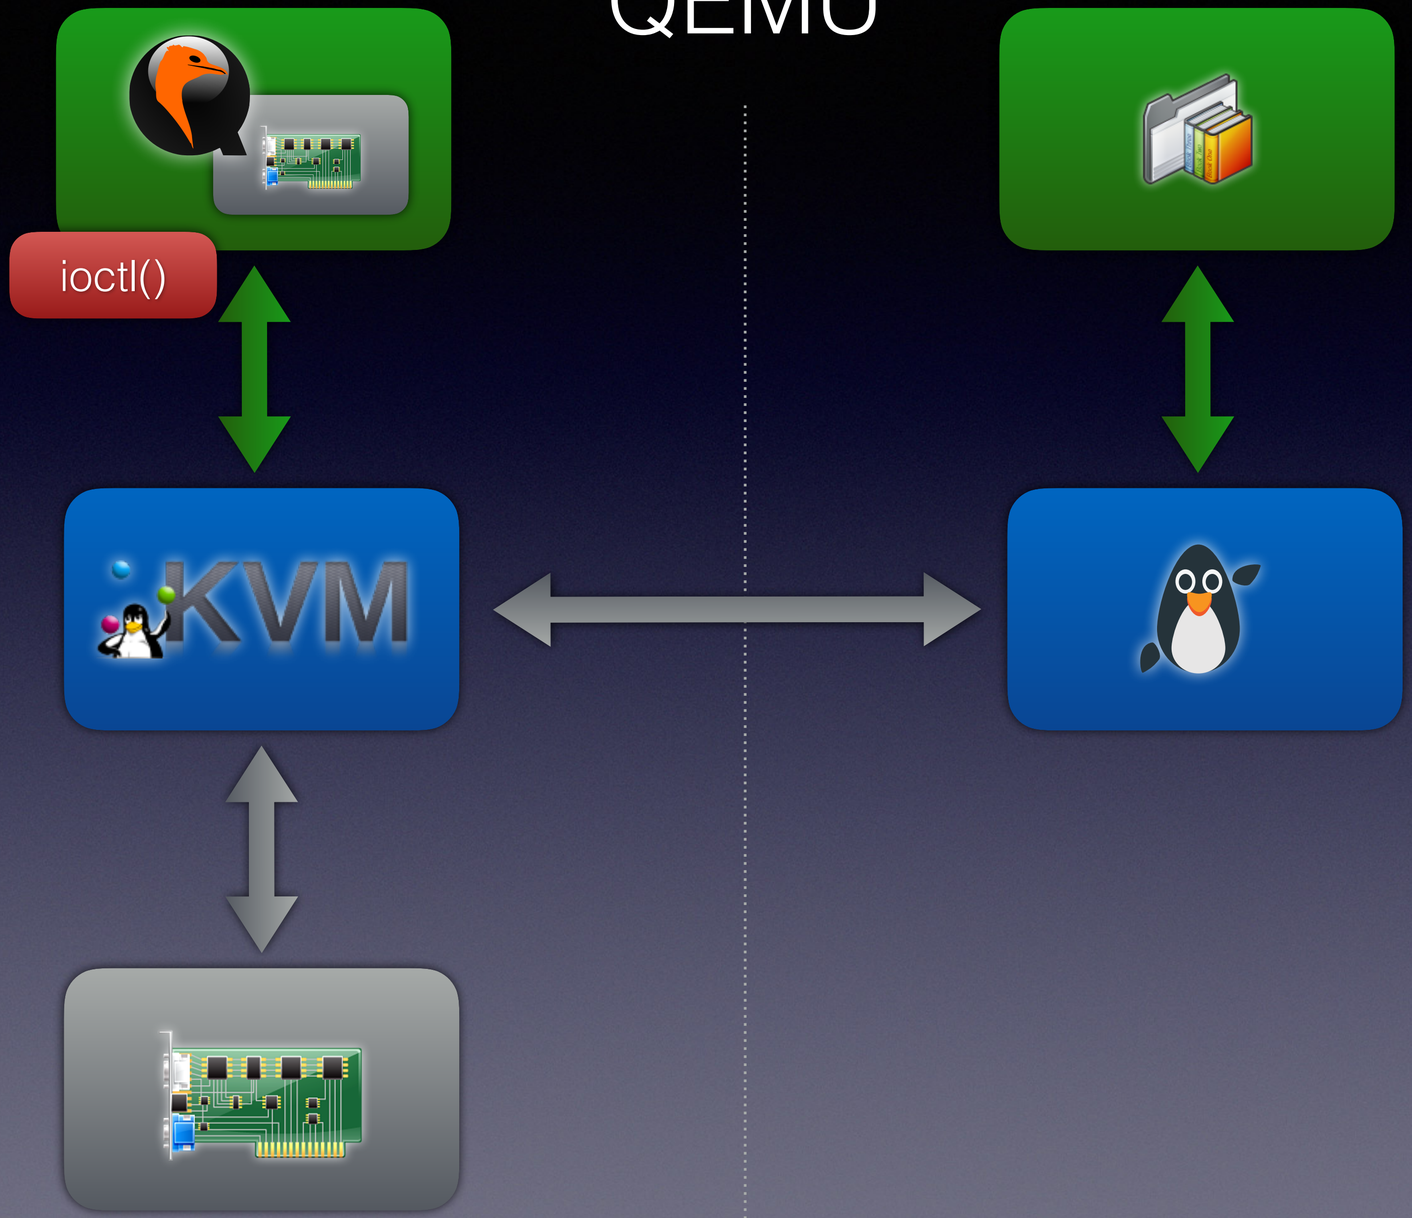

QEMU

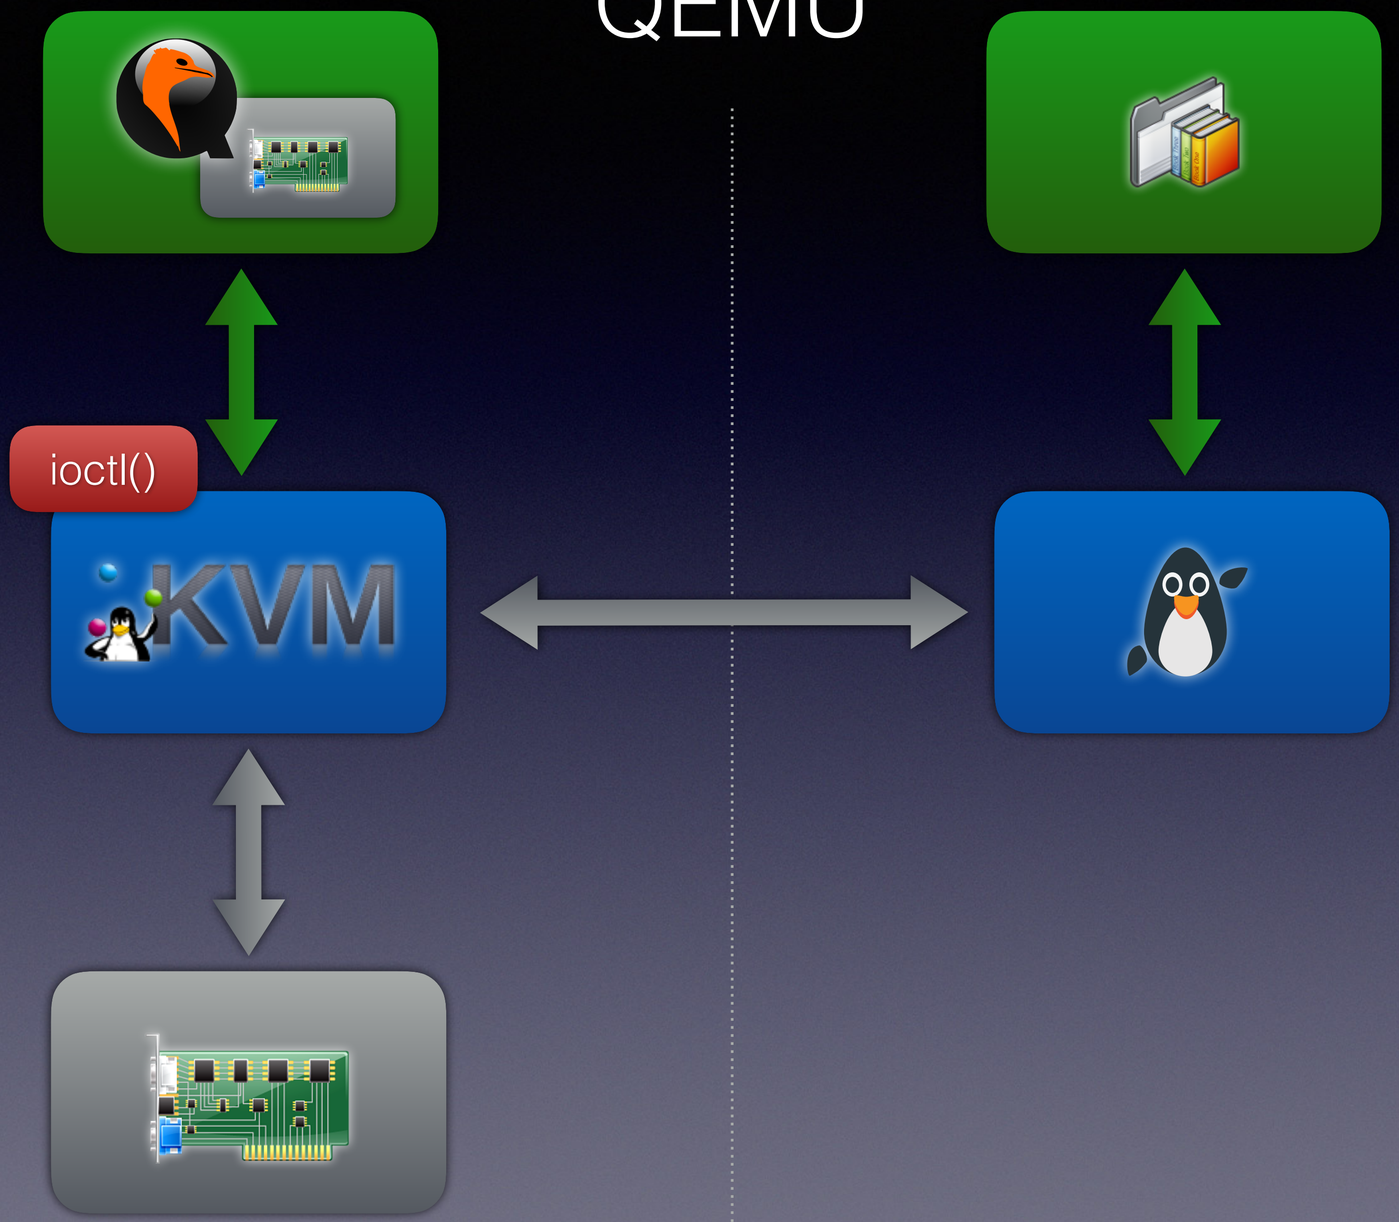

QEMU

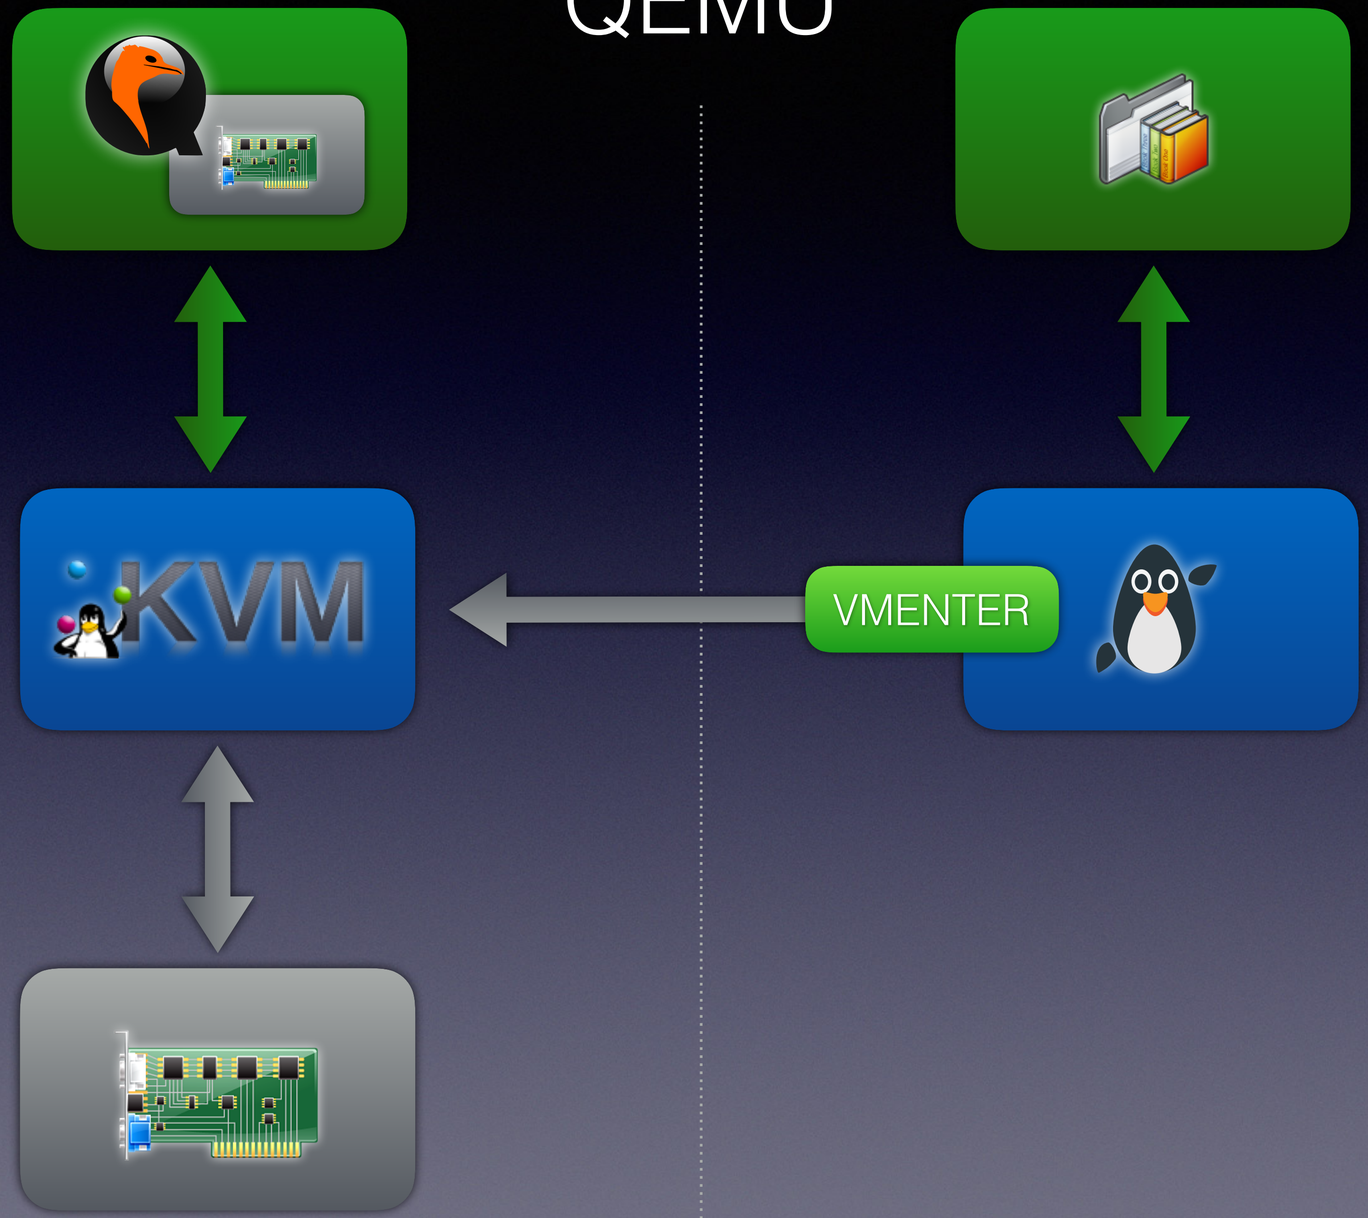

VFIO

Device ~~Emulation~~

VFIO

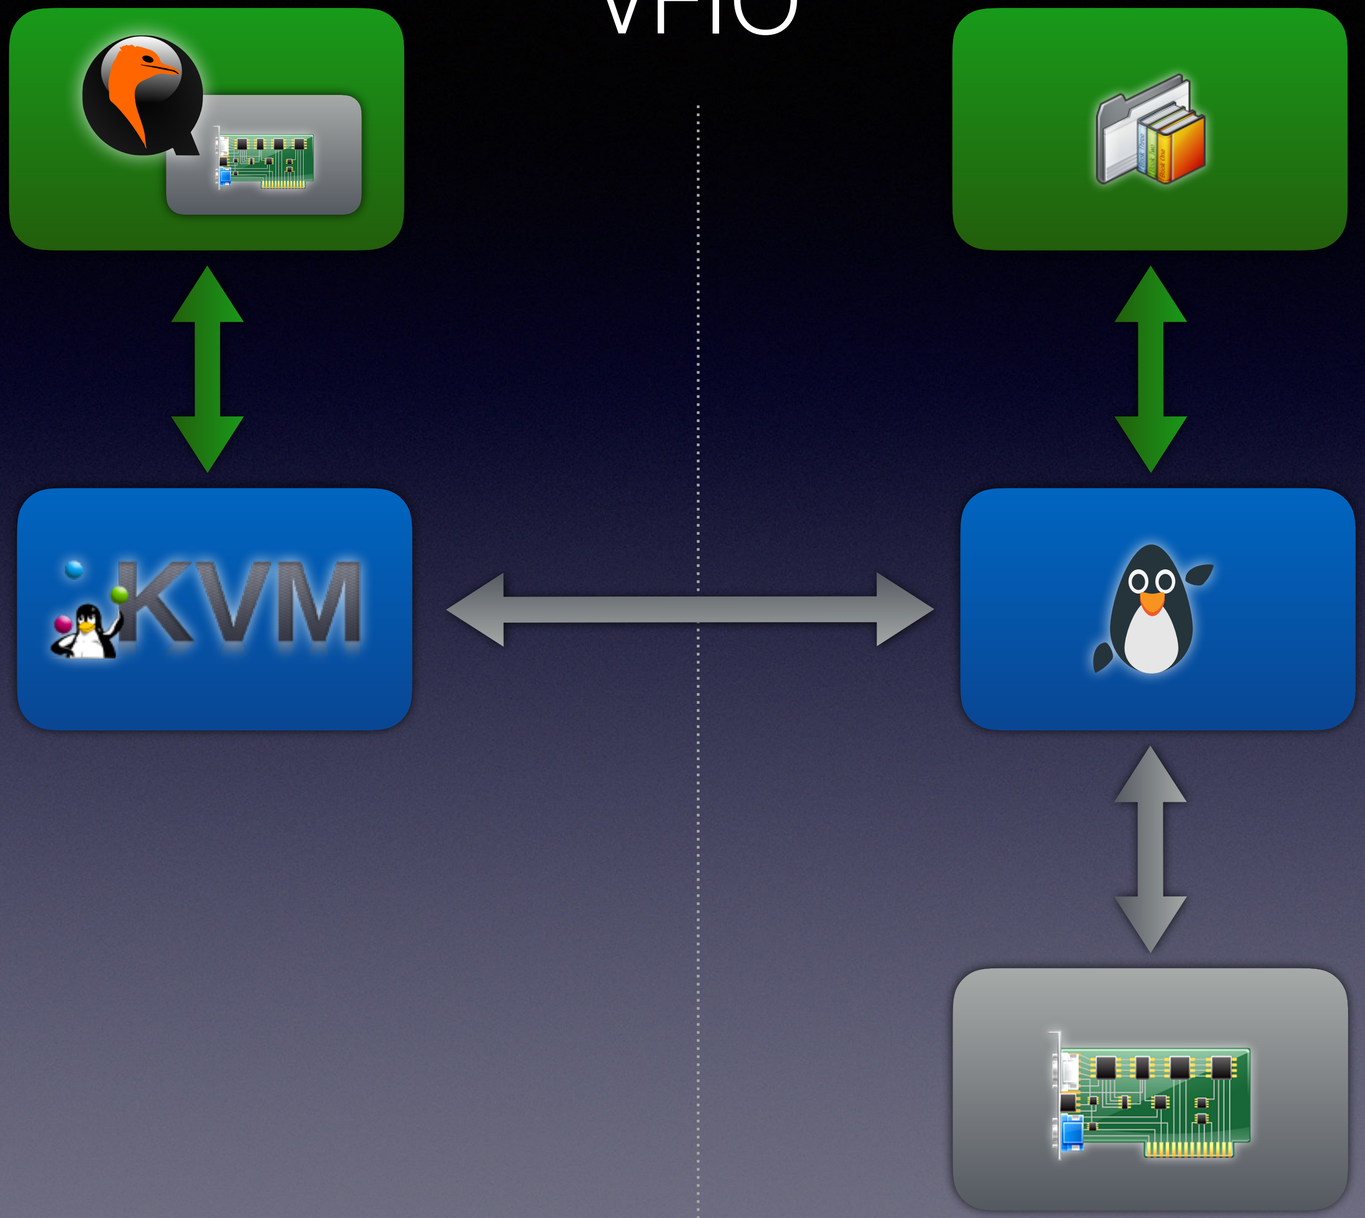

VFIO

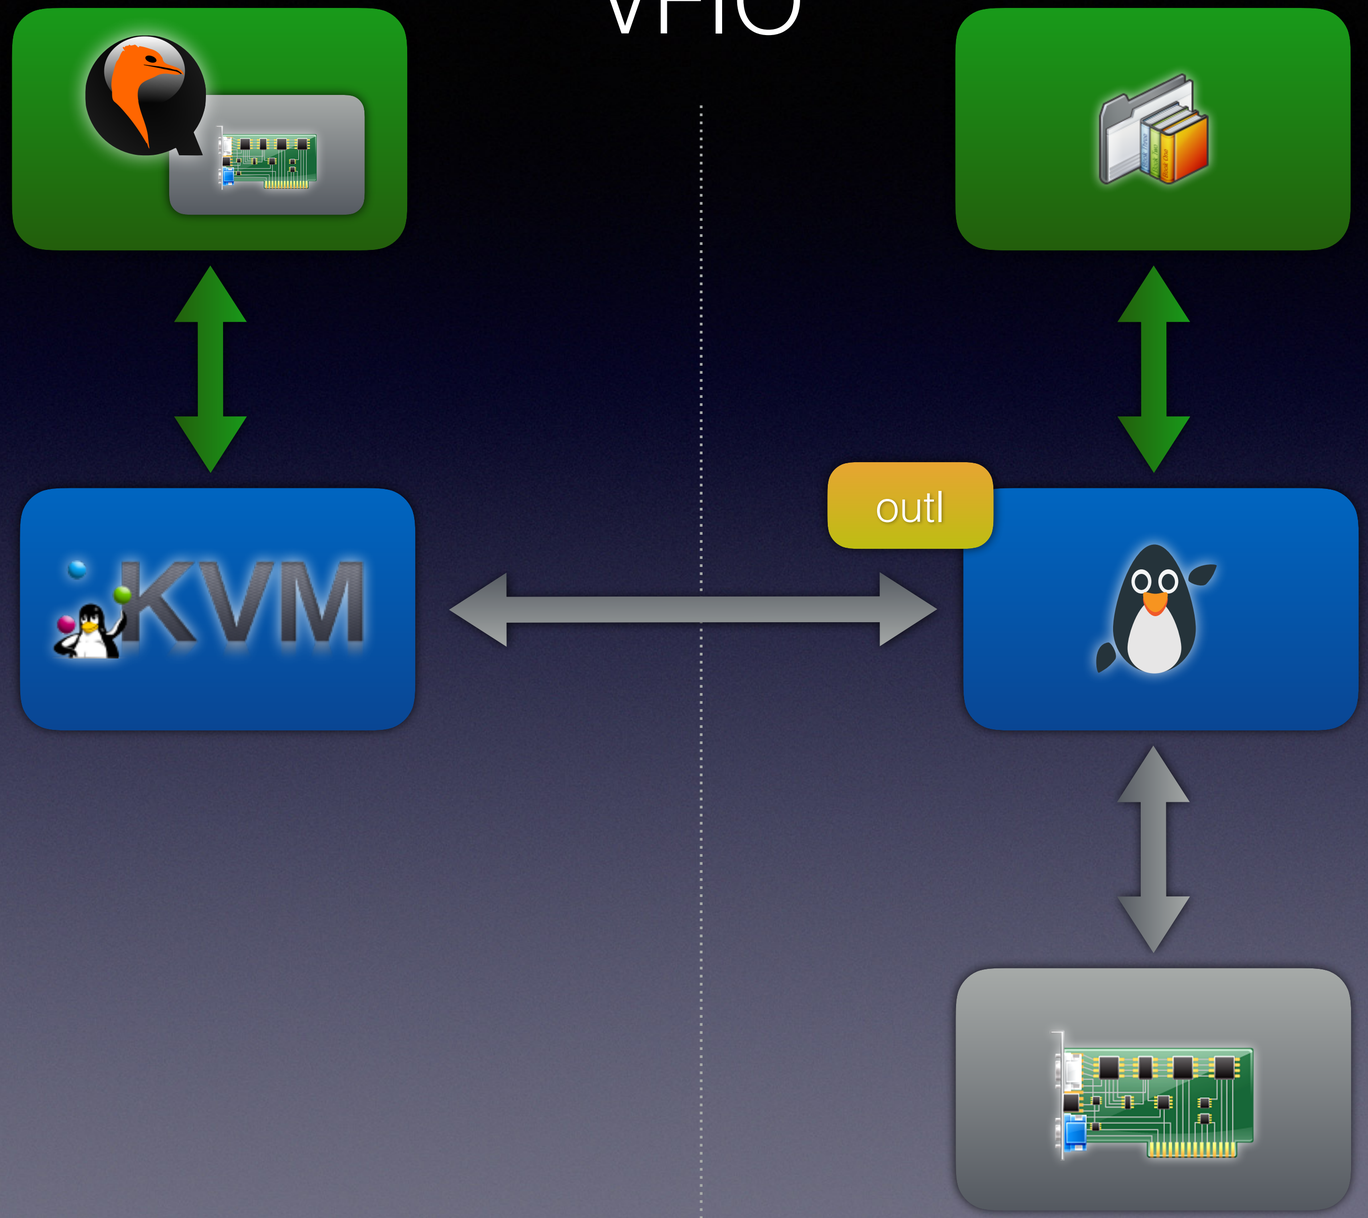

VFIO

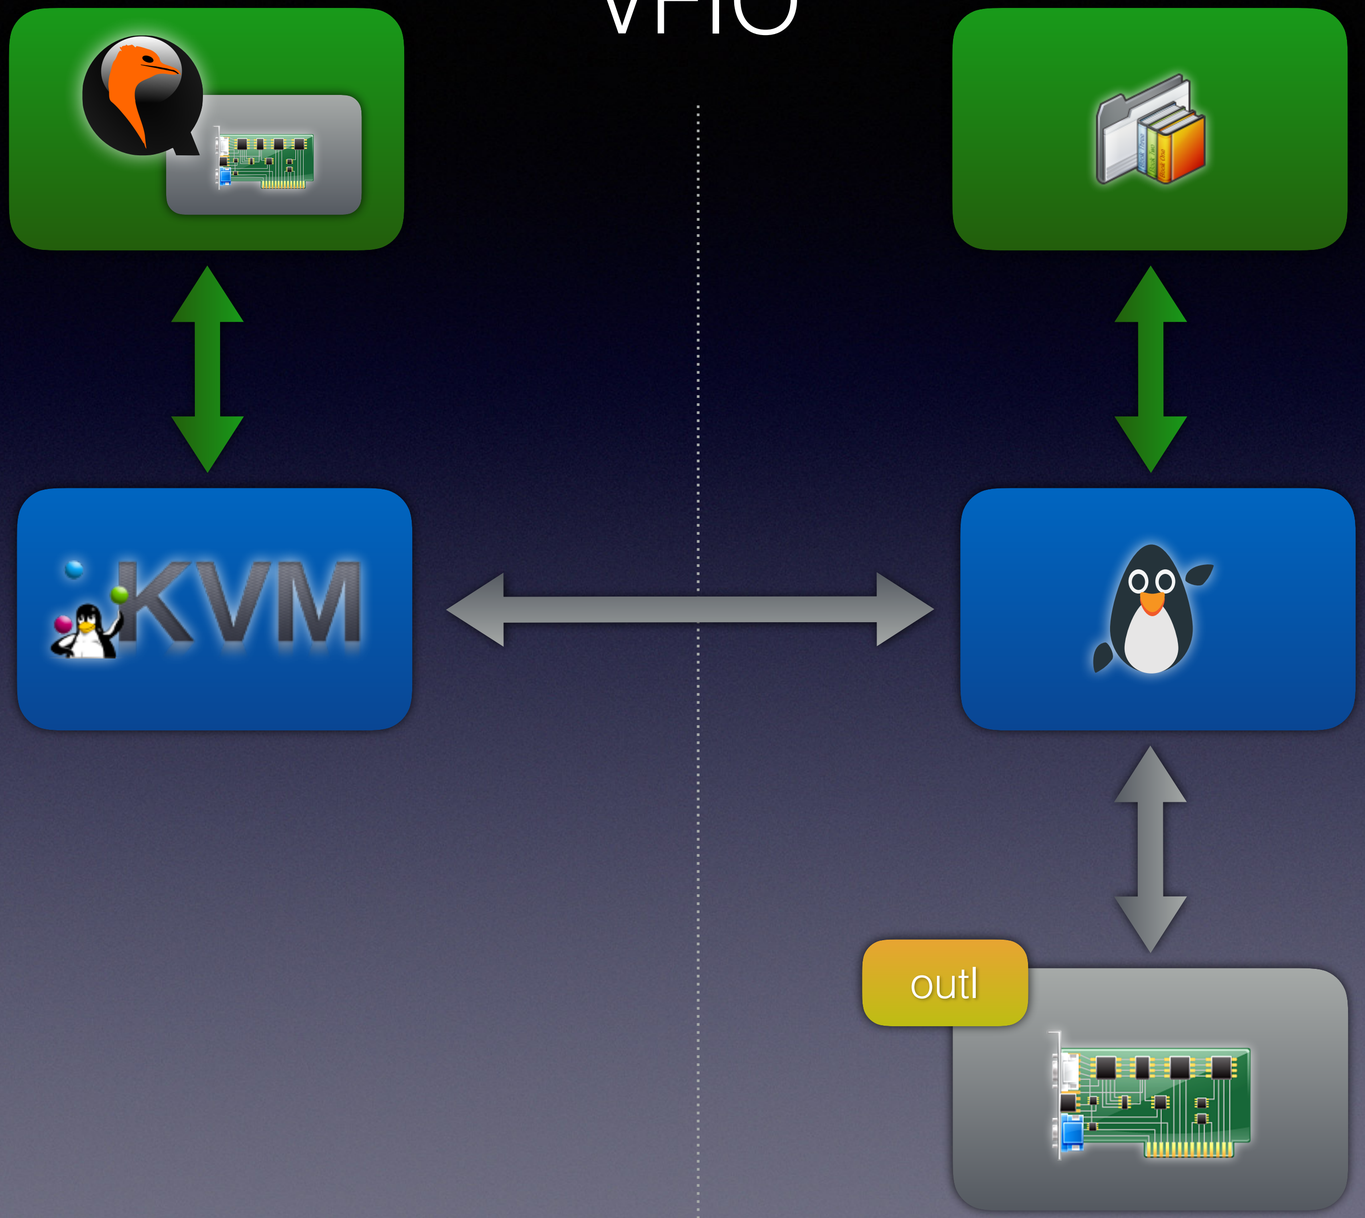

VFIO

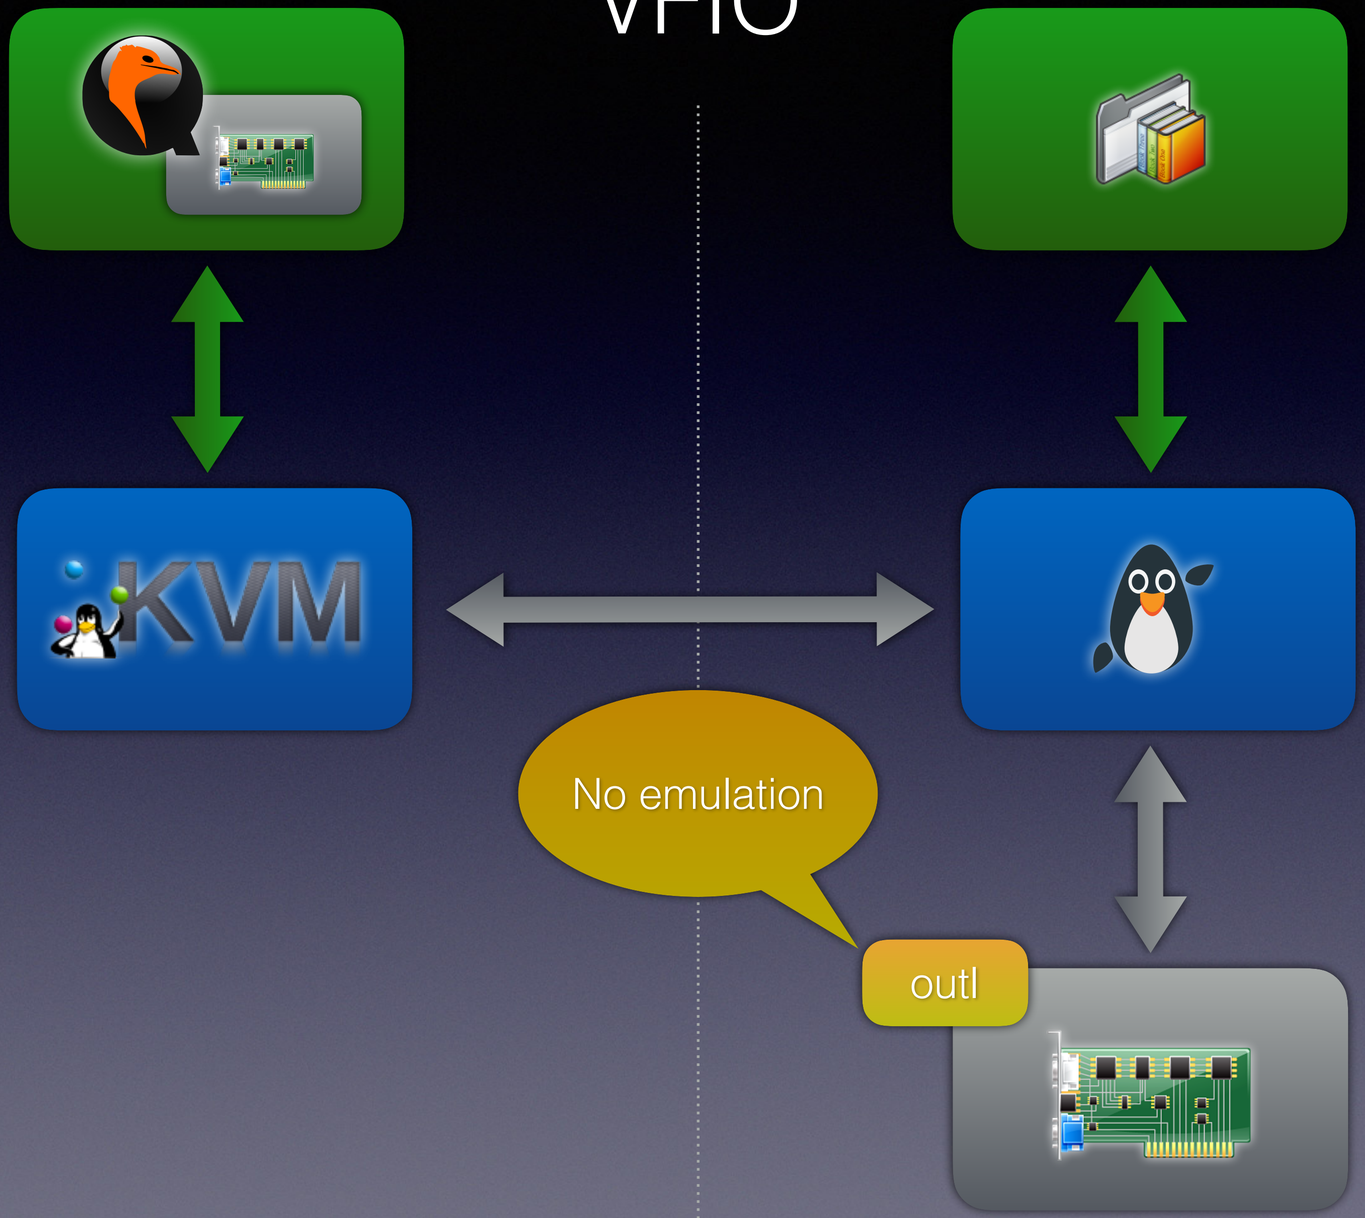

QEMU driver

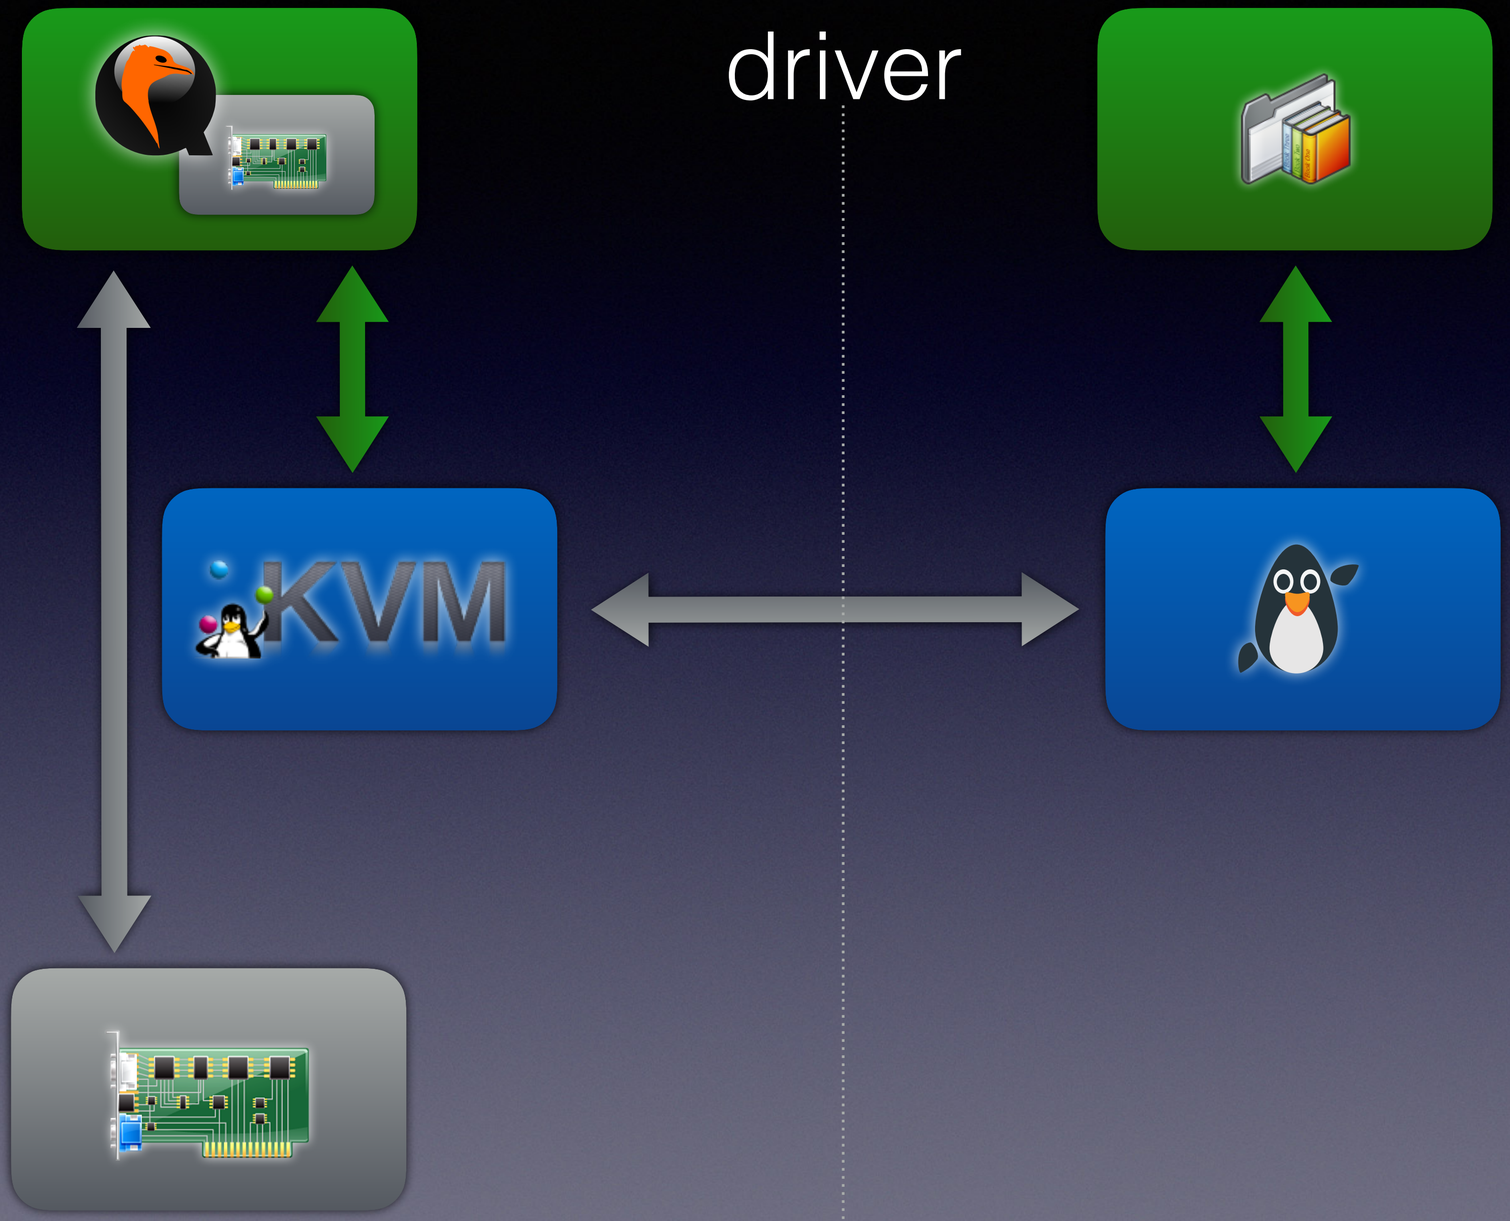

vSMM

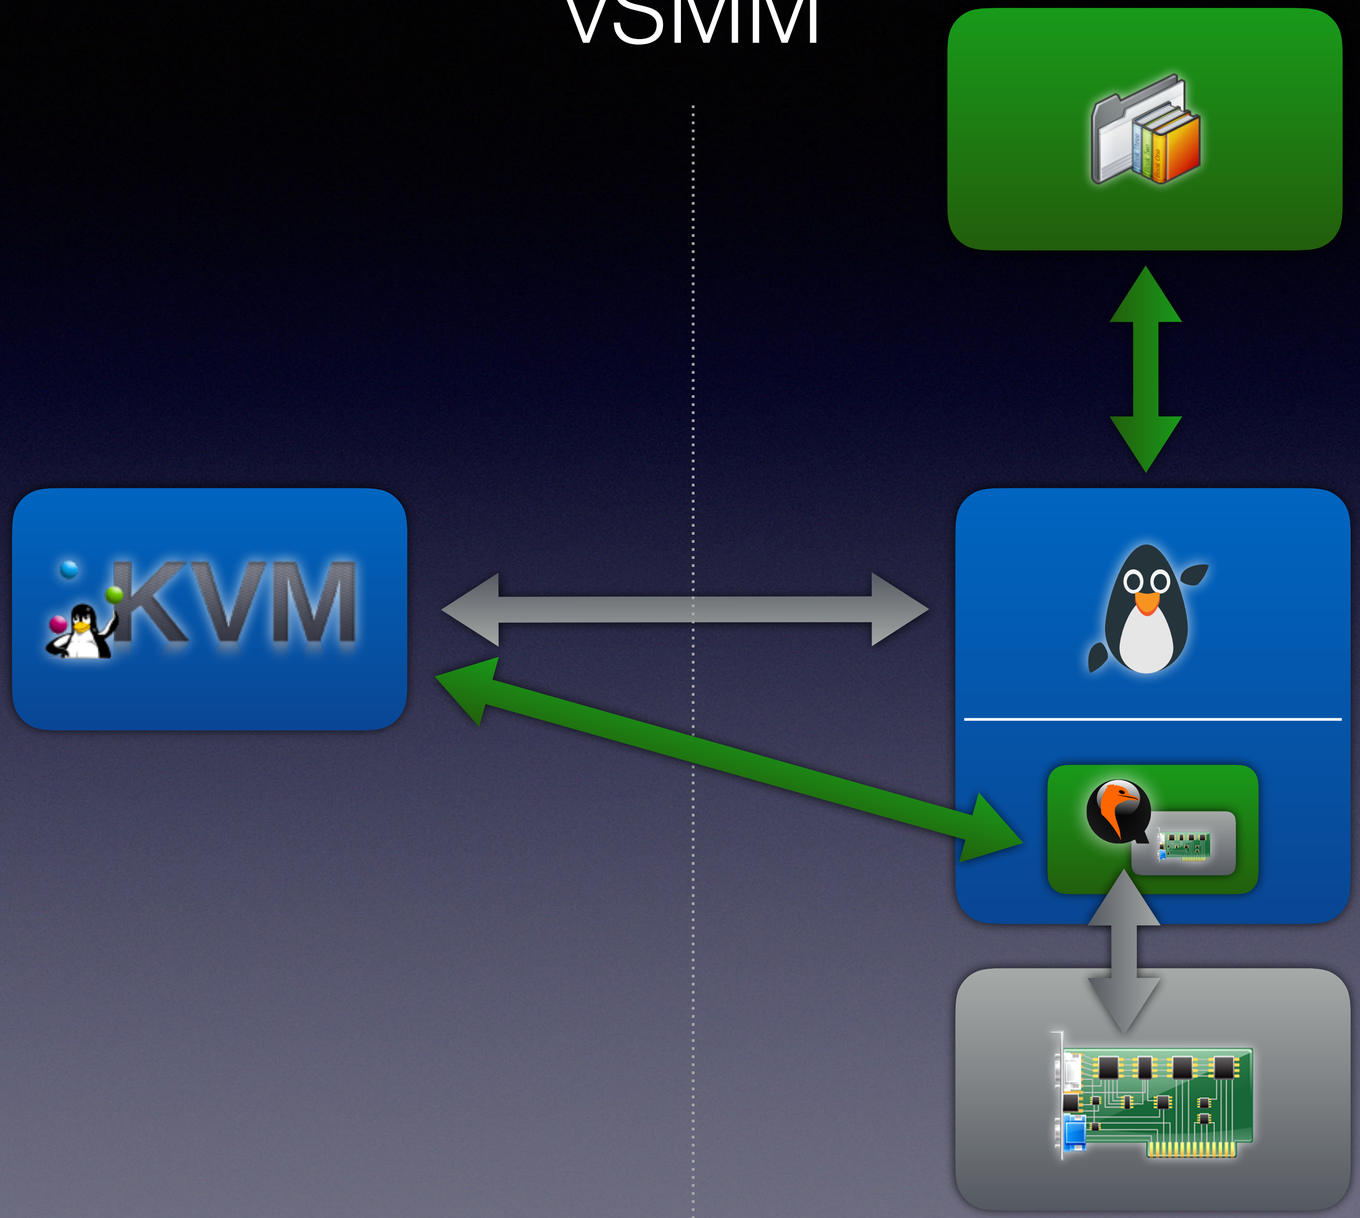

vSMM

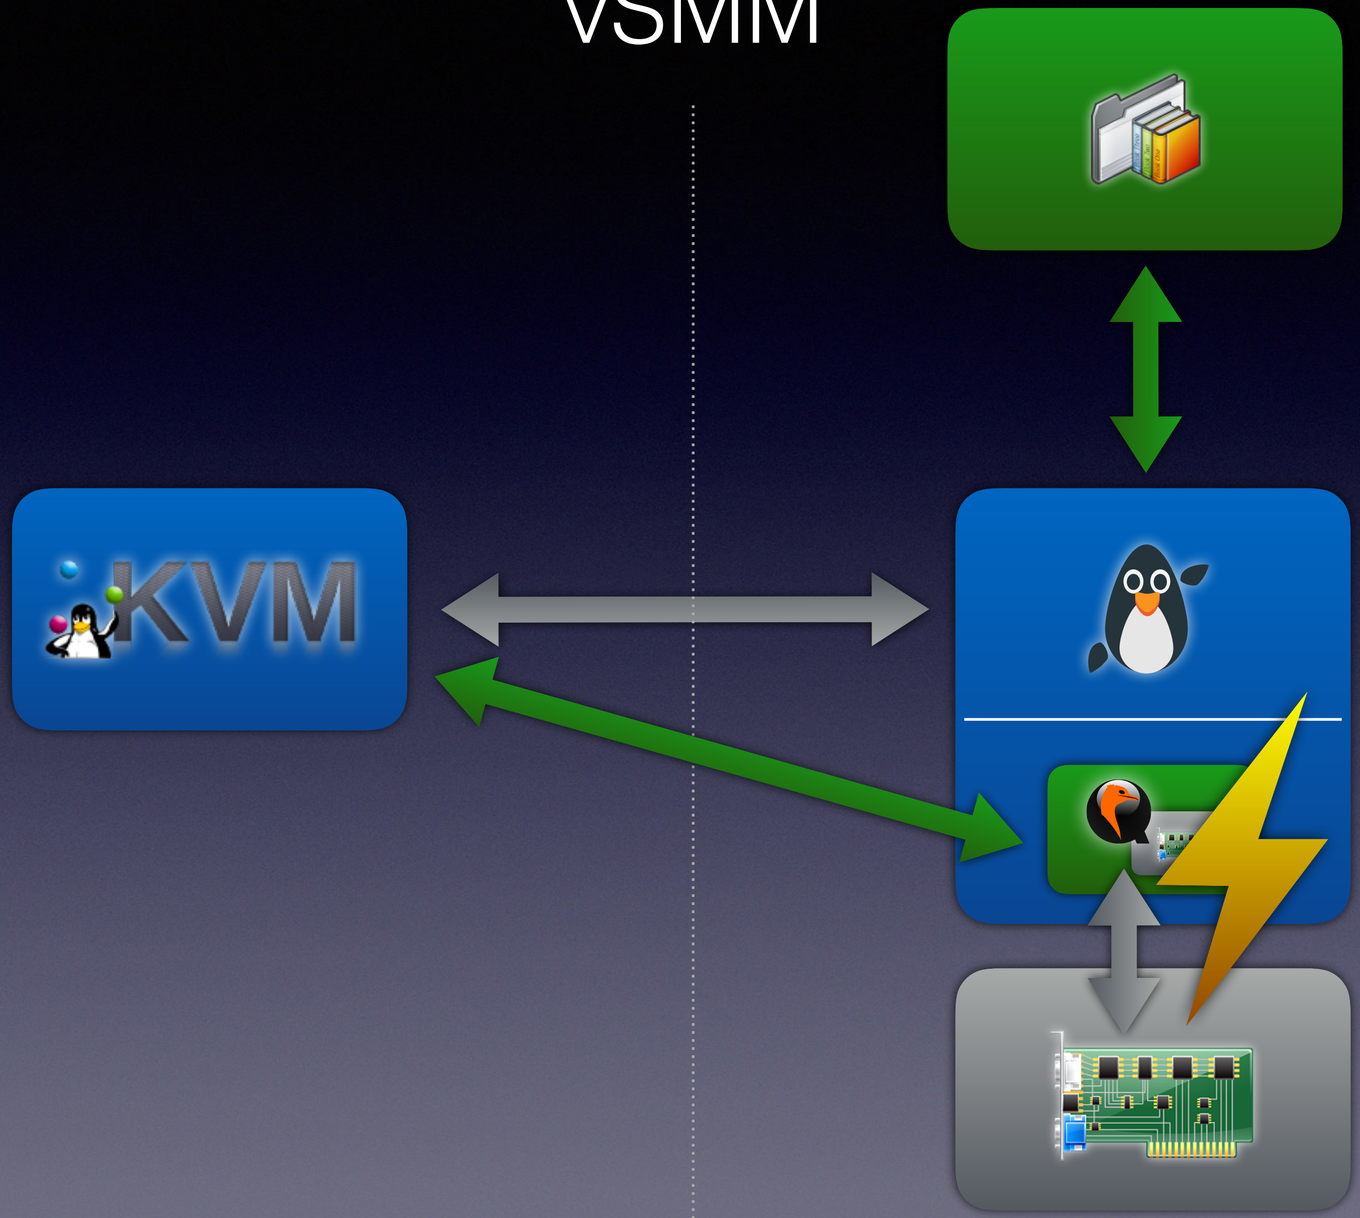

vSMM

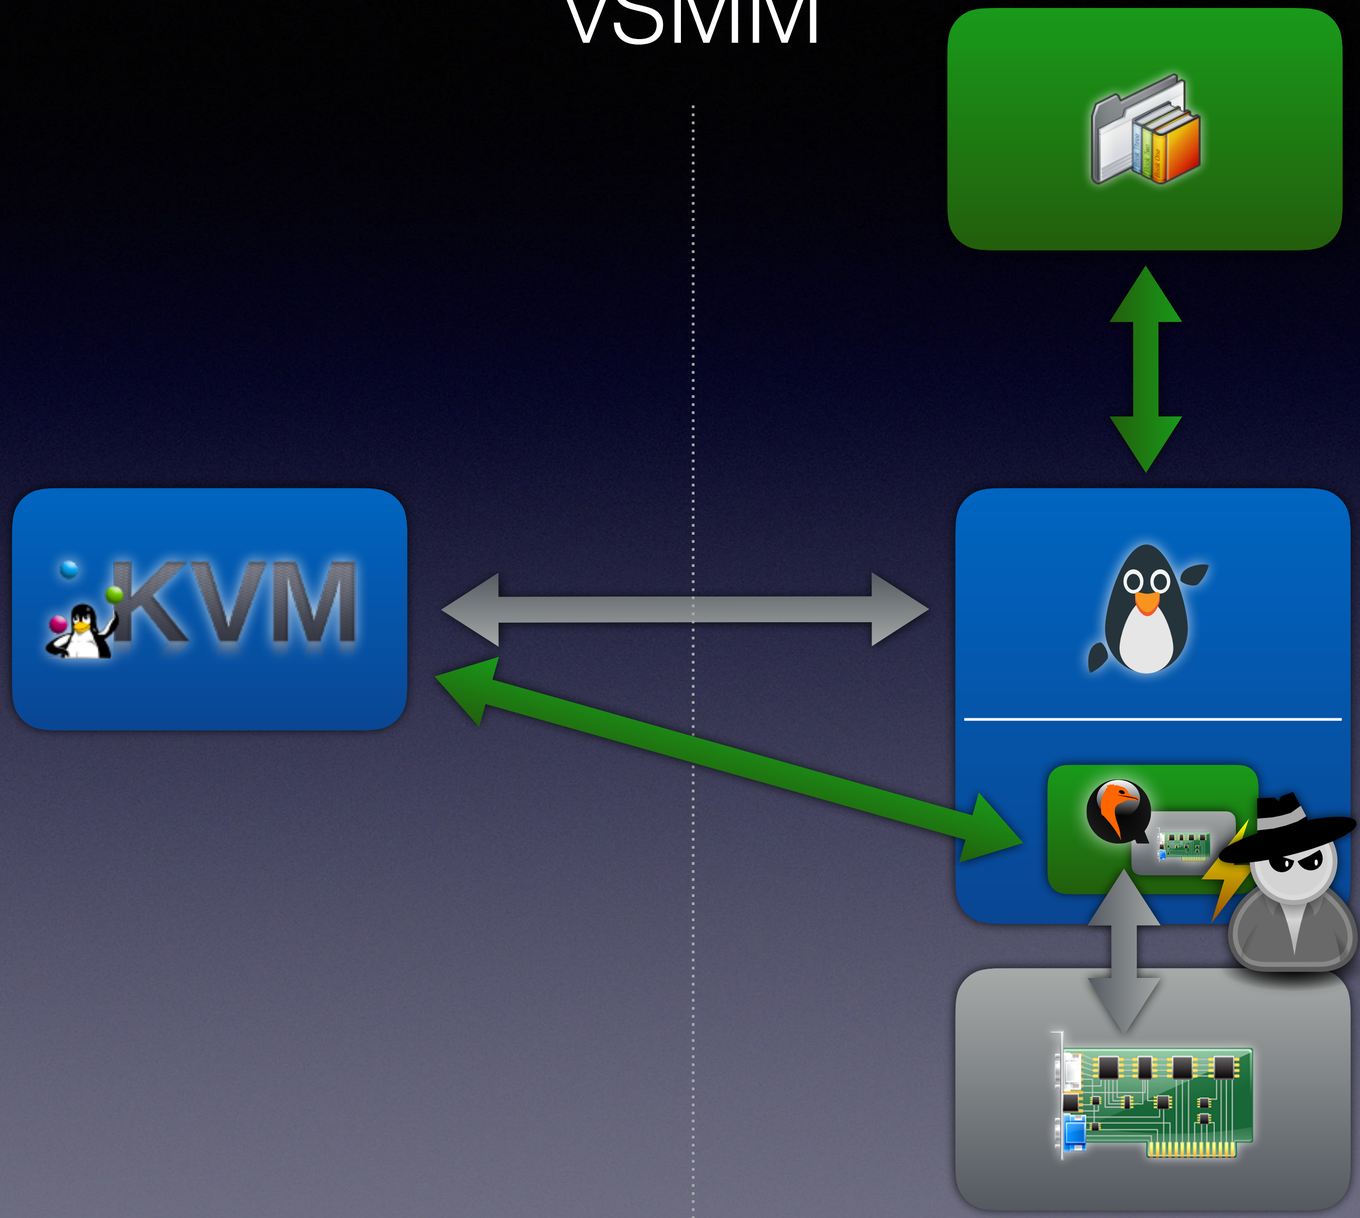

PIO

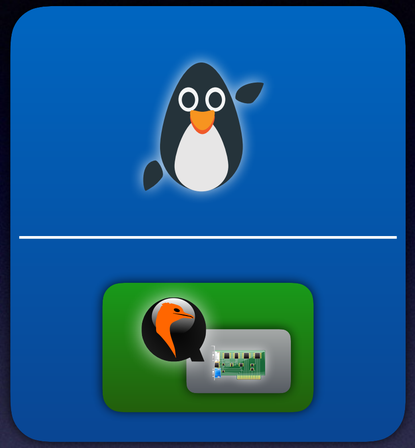

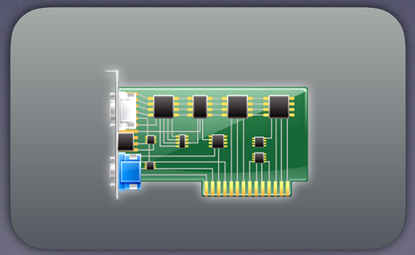

PIO

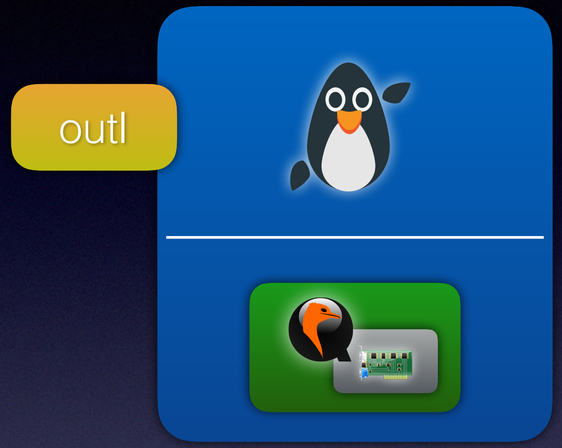

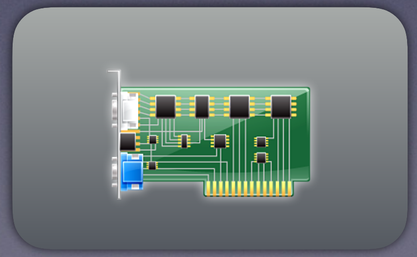

PIO

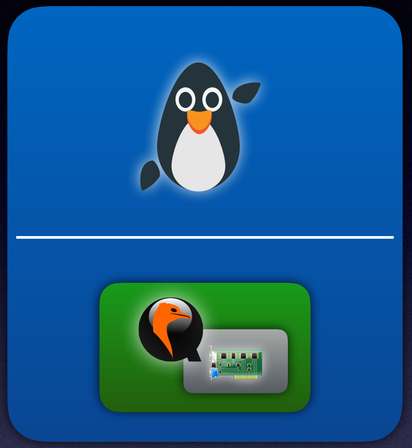

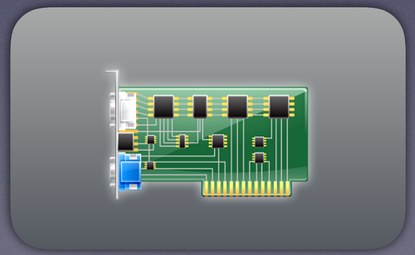

PIO

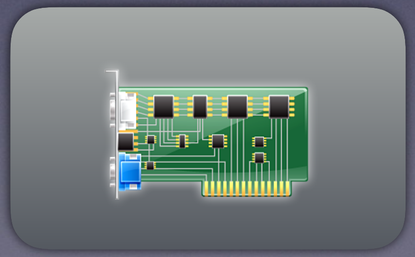

PIO

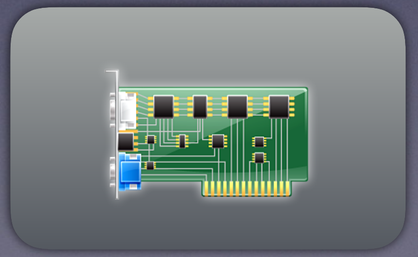

PIO

PIO

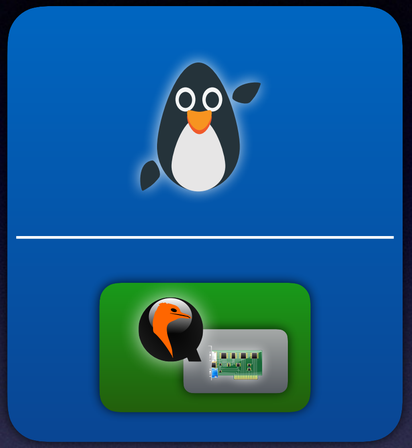

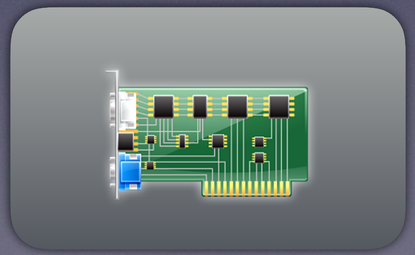

PIO

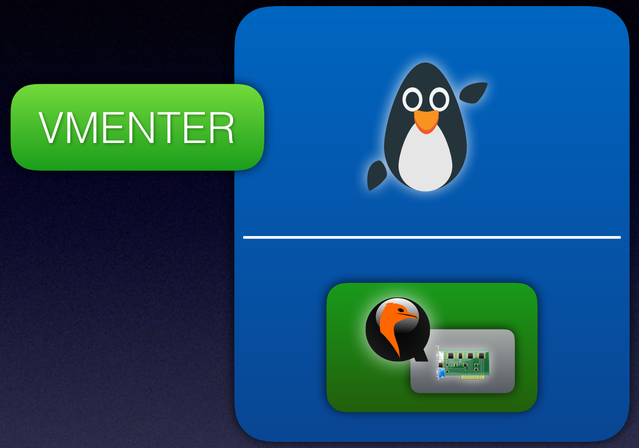

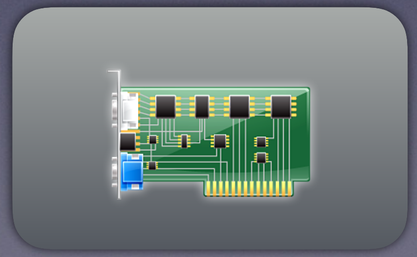

MMIO

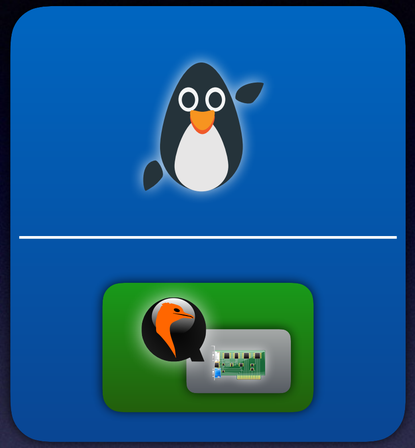

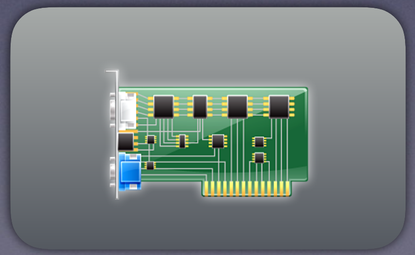

MMIO

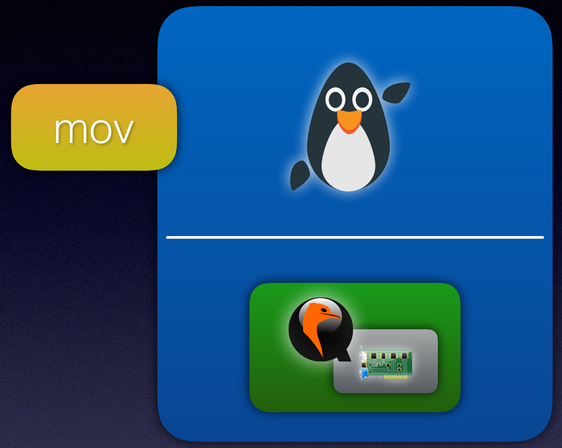

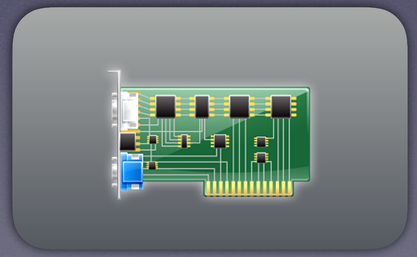

MMIO

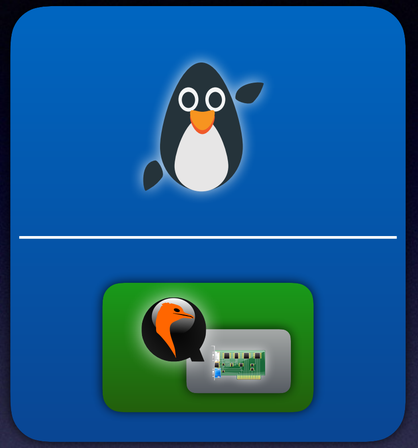

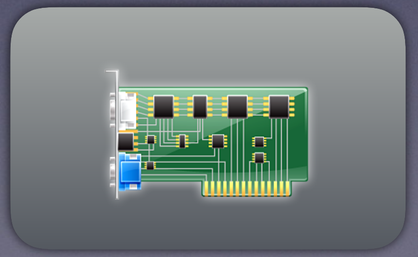

MMIO

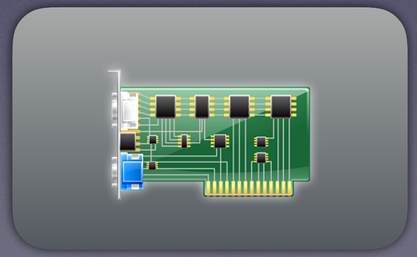

MMIO

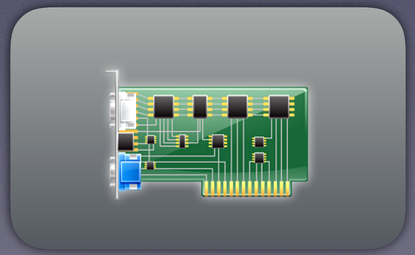

MMIO

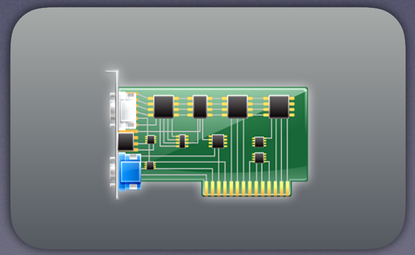

MMIO

MMIO

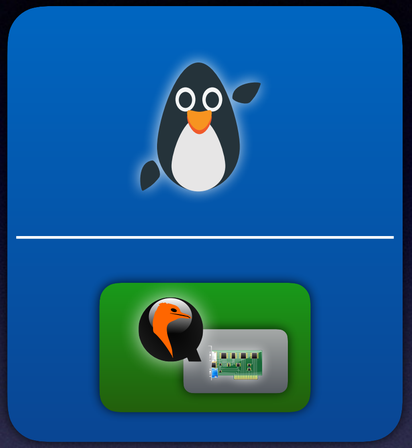

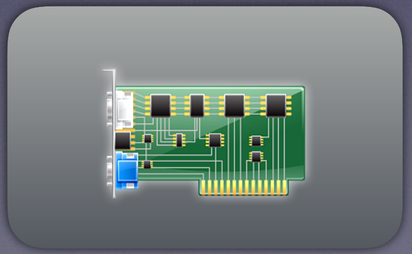

MMIO

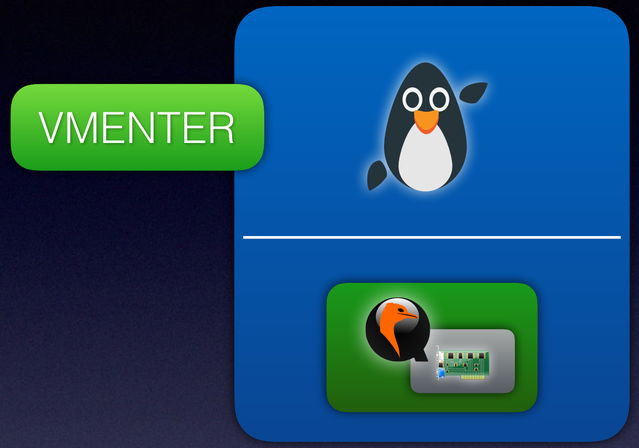

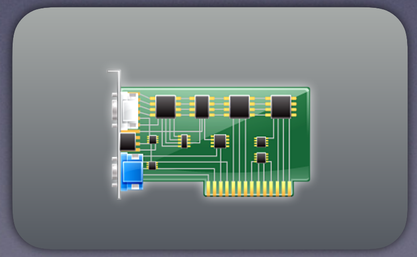

IRQs

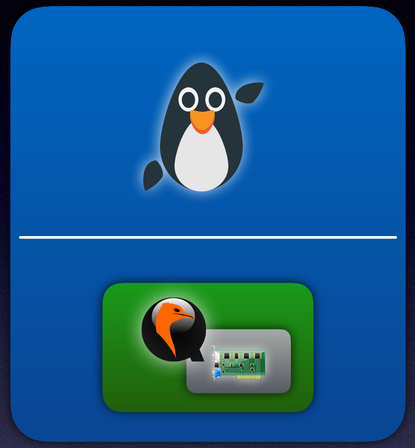

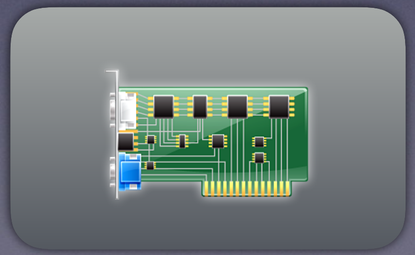

IRQs

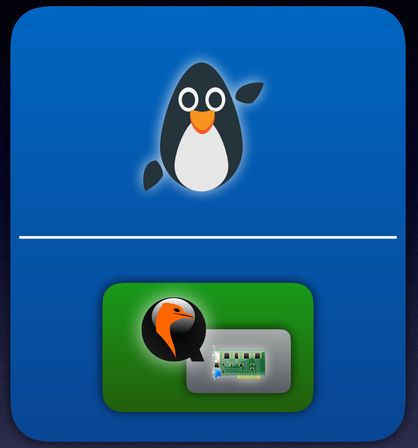

MSI-X

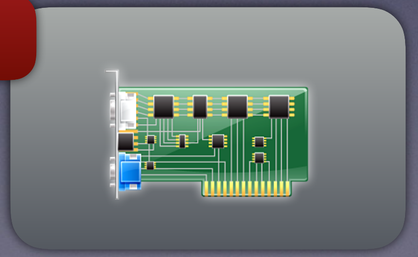

IRQs

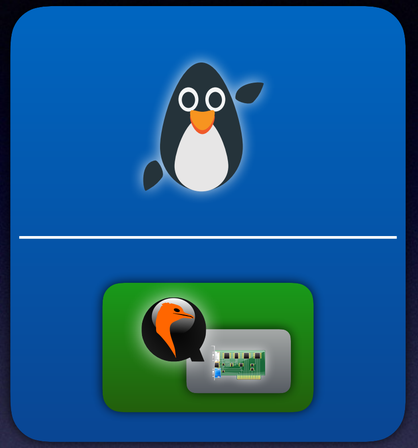

Target SMI

MSI-X

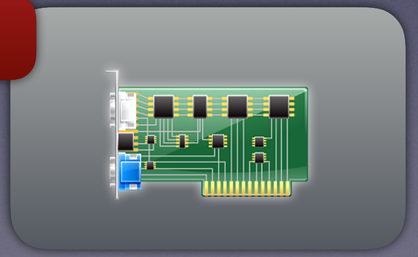

IRQs

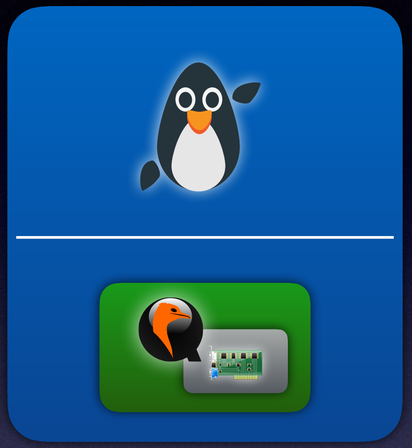

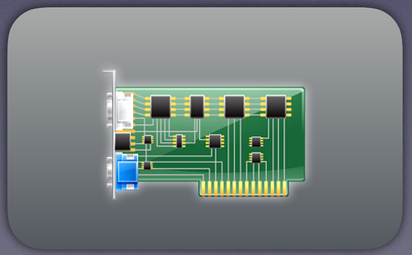

IRQs

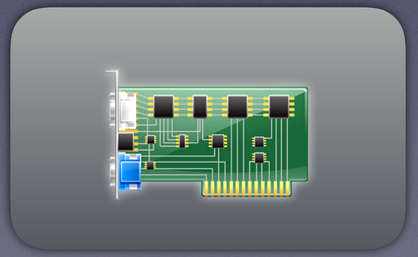

Benefits

Benefits

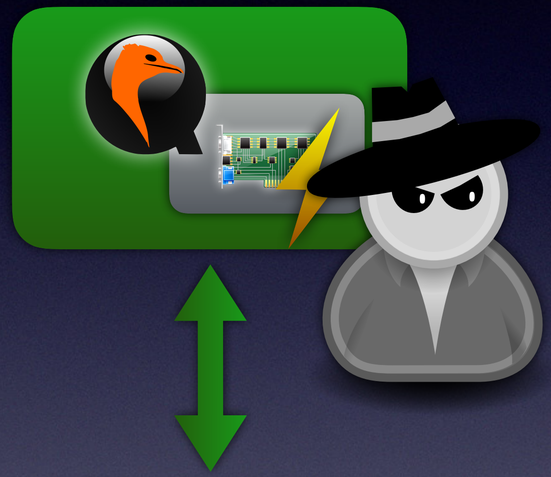

Benefits

Benefits

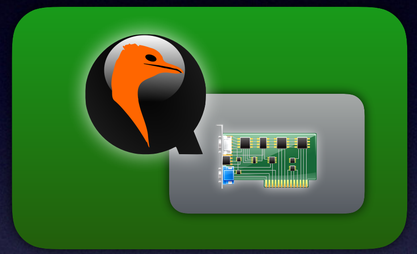

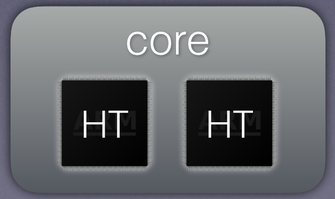

Benefits

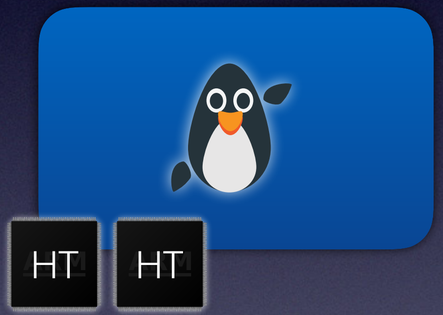

Benefits

Benefits

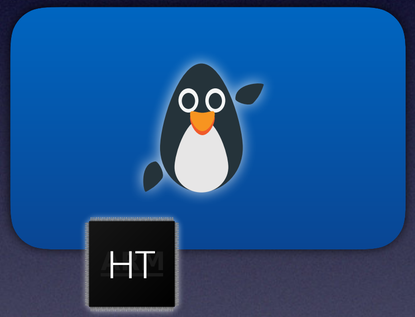

Benefits

Benefits

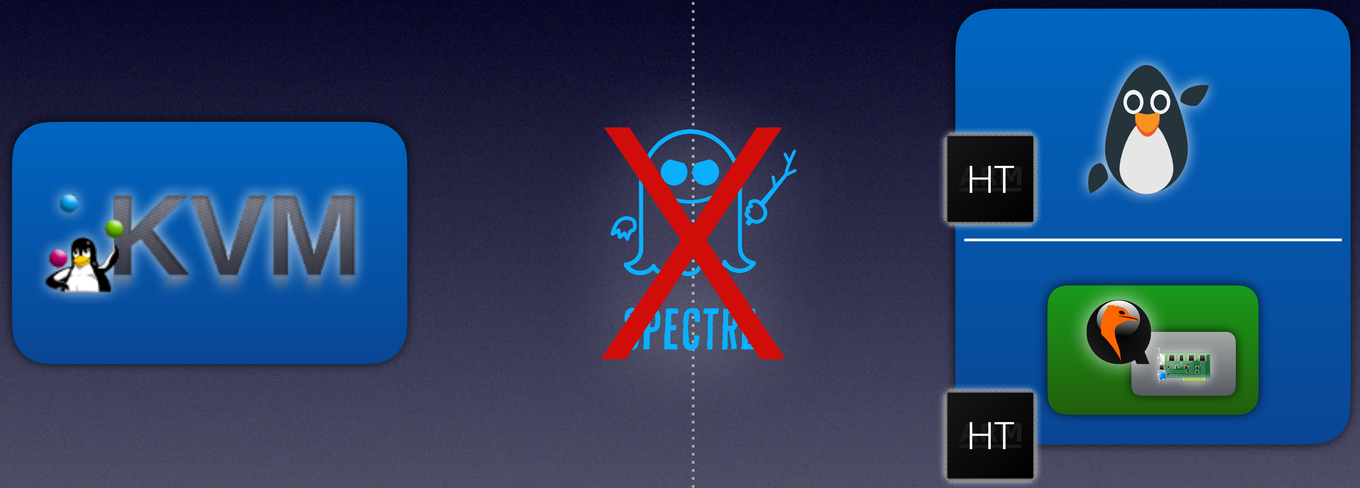

Benefits

Benefits

Benefits

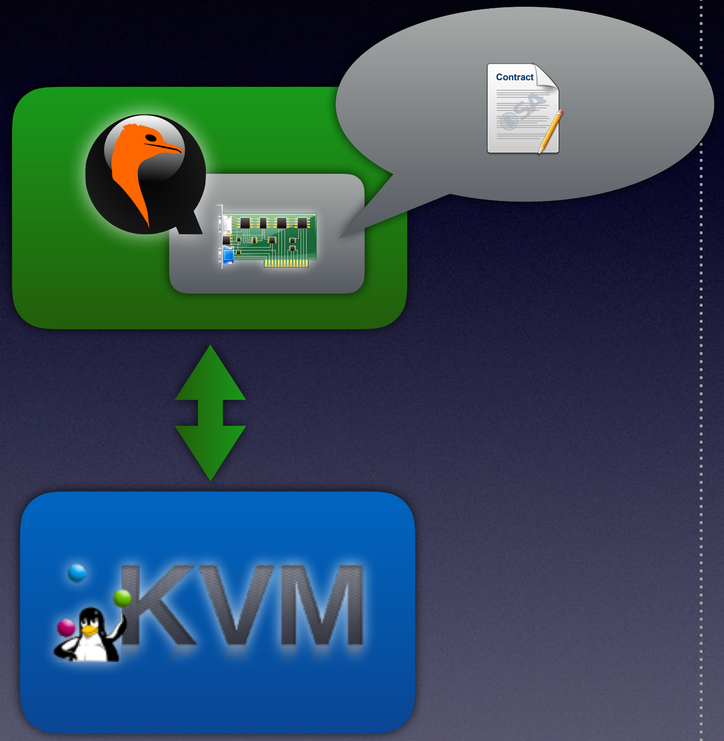

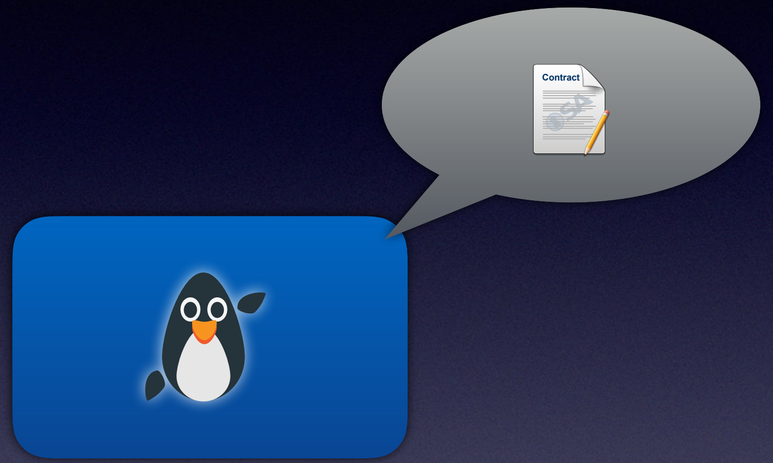

Benefits

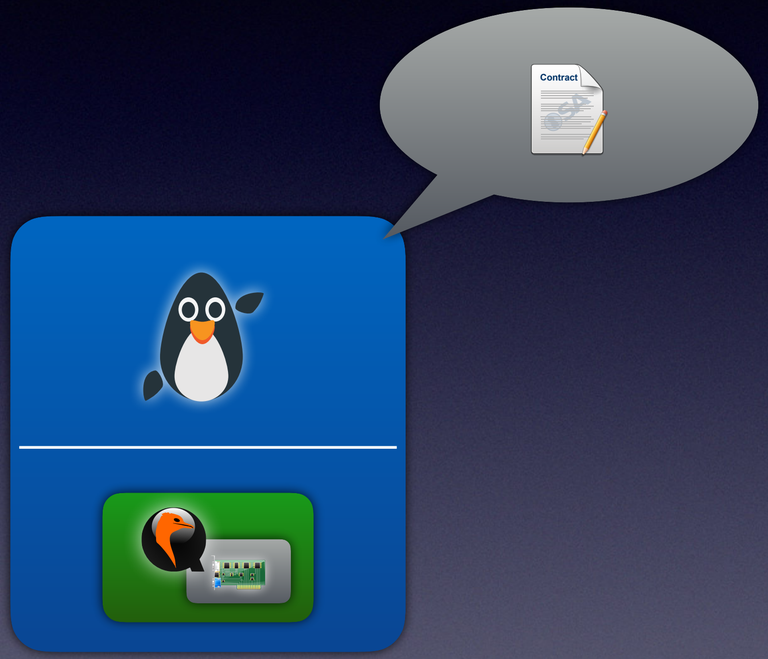

Downsides

Downsides

- AMD SEV
- Nested Virtualization
- Blue Pill

IDE on NVMe


```
localhost login: root
Last login: Tue Sep  3 10:03:58 on ttyS0
Have a lot of fun...
ip-10-0-2-15:~ # lspci
00:00.0 Host bridge: Intel Corporation 440FX - 82441FX PMC [Natoma] (rev 02)
00:01.0 ISA bridge: Intel Corporation 82371SB PIIX3 ISA [Natoma/Triton II]
00:01.3 Bridge: Intel Corporation 82371AB/EB/MB PIIX4 ACPI (rev 03)
00:02.0 VGA compatible controller: Device 1234:1111 (rev 02)
00:03.0 Ethernet controller: Intel Corporation 82540EM Gigabit Ethernet Controller (rev 03)
00:04.0 IDE interface: Intel Corporation 82371SB PIIX3 IDE [Natoma/Triton II]
ip-10-0-2-15:~ # lsscsi
[0:0:0:0]    disk    ATA      model            ion    /dev/sda
ip-10-0-2-15:~ # hdparm -i /dev/sda

/dev/sda:

 Model=model, FwRev=version, SerialNo=serial
 Config={ }
 RawCHS=0/255/255, TrkSize=65024, SectSize=512, ECCbytes=4
 BuffType=DualPortCache, BuffSize=256kB, MaxMultSect=6, MultSect=6
 CurCHS=255/255/255, CurSects=16581375, LBA=yes, LBAsects=1757812500
 IORDY=yes, tPIO={min:120,w/IORDY:120}, tDMA={min:120,rec:120}
 PIO modes:  pio0 pio1 pio2 pio3 pio4
 DMA modes:  sdma0 sdma1 sdma2 mdma0 mdma1 *mdma2
 UDMA modes: udma0 udma1 udma2 udma3 udma4 udma5
 AdvancedPM=no WriteCache=enabled
 Drive conforms to: ATA/ATAPI-5 published, ANSI INCITS 340-2000:  ATA/ATAPI-4,5,6,7

 * signifies the current active mode

ip-10-0-2-15:~ # hdparm -t /dev/sda

/dev/sda:
 Timing buffered disk reads: 342 MB in  3.01 seconds = 113.55 MB/sec
ip-10-0-2-15:~ #
```

Demo

Thank You

OSA Icons

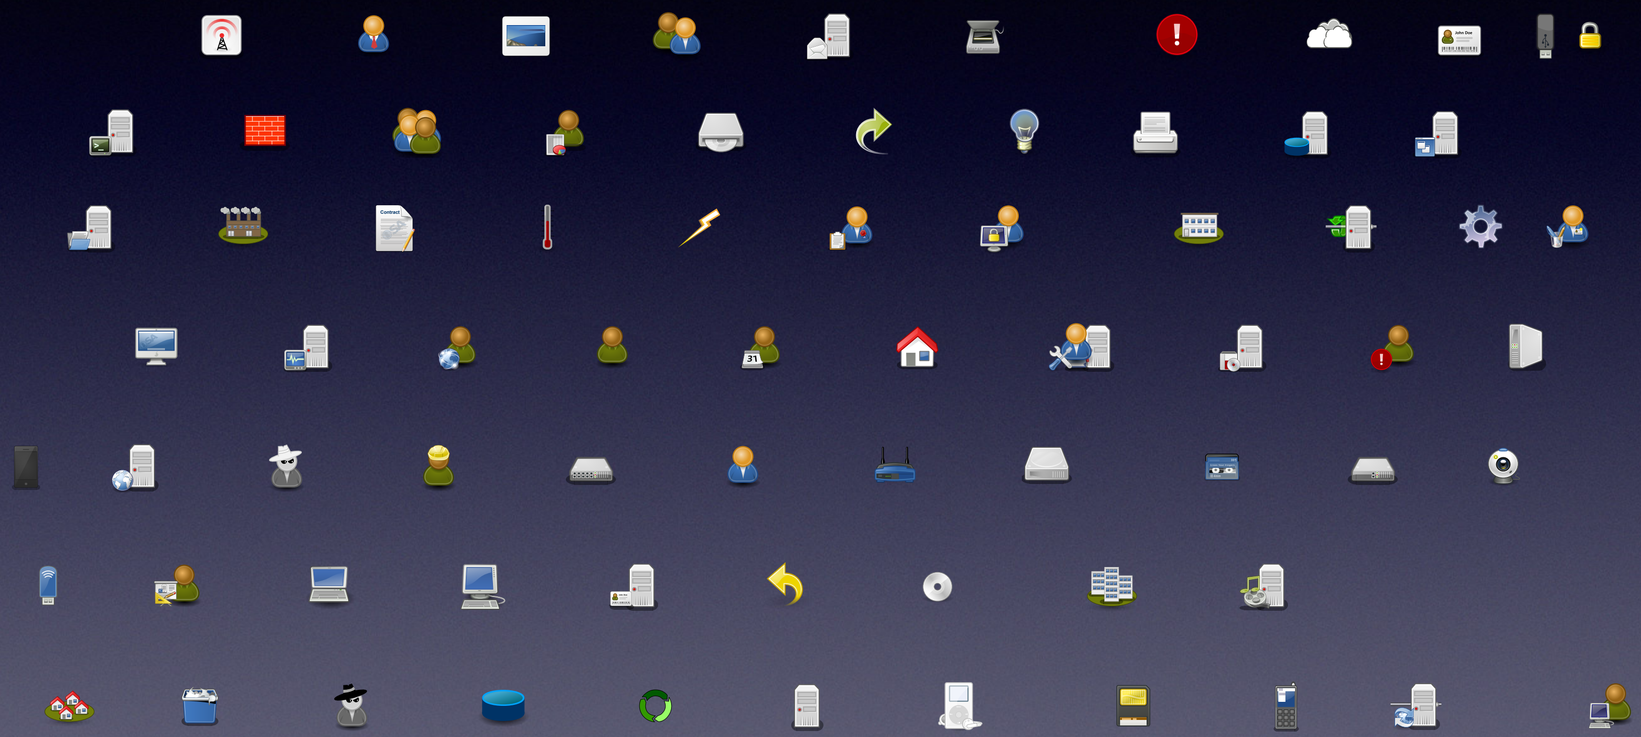

Icons received from http://www.opensecurityarchitecture.org/cms/library/icon-library

Other Icons

http://findicons.com/icon/202613/folder_library

http://findicons.com/icon/download/234261/clock/128/png

http://findicons.com/icon/439269/button_power

https://fosdem.org/2017/schedule/event/grub_new_maintainers/attachments/slides/1768/export/events/attachments/grub_new_maintainers/slides/1768/slides.pdf

https://de.wikipedia.org/wiki/BeagleBoard#/media/File:Beagle_Board_big.jpg

https://thenounproject.com/term/folder-tree/27307/

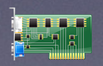

https://commons.wikimedia.org/wiki/File:Crystal_Project_Hardware.png

emojione Icons



Icons received from http://www.opensecurityarchitecture.org/cms/library/icon-library

External Sources

https://commons.wikimedia.org/wiki/File:Spectre_logo_with_text.svg

https://commons.wikimedia.org/wiki/File:USB_icon.svg

https://commons.wikimedia.org/wiki/File:150-8-DIP.jpg

https://commons.wikimedia.org/wiki/File:Hdd_icon.svg

https://commons.wikimedia.org/wiki/File:ARM_CPU_icon.svg

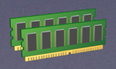

http://findicons.com/icon/177982/memory#

https://www.linux-kvm.org/page/Main_Page

https://commons.wikimedia.org/wiki/File:Keyboard-icon_Wikipedians.svg