

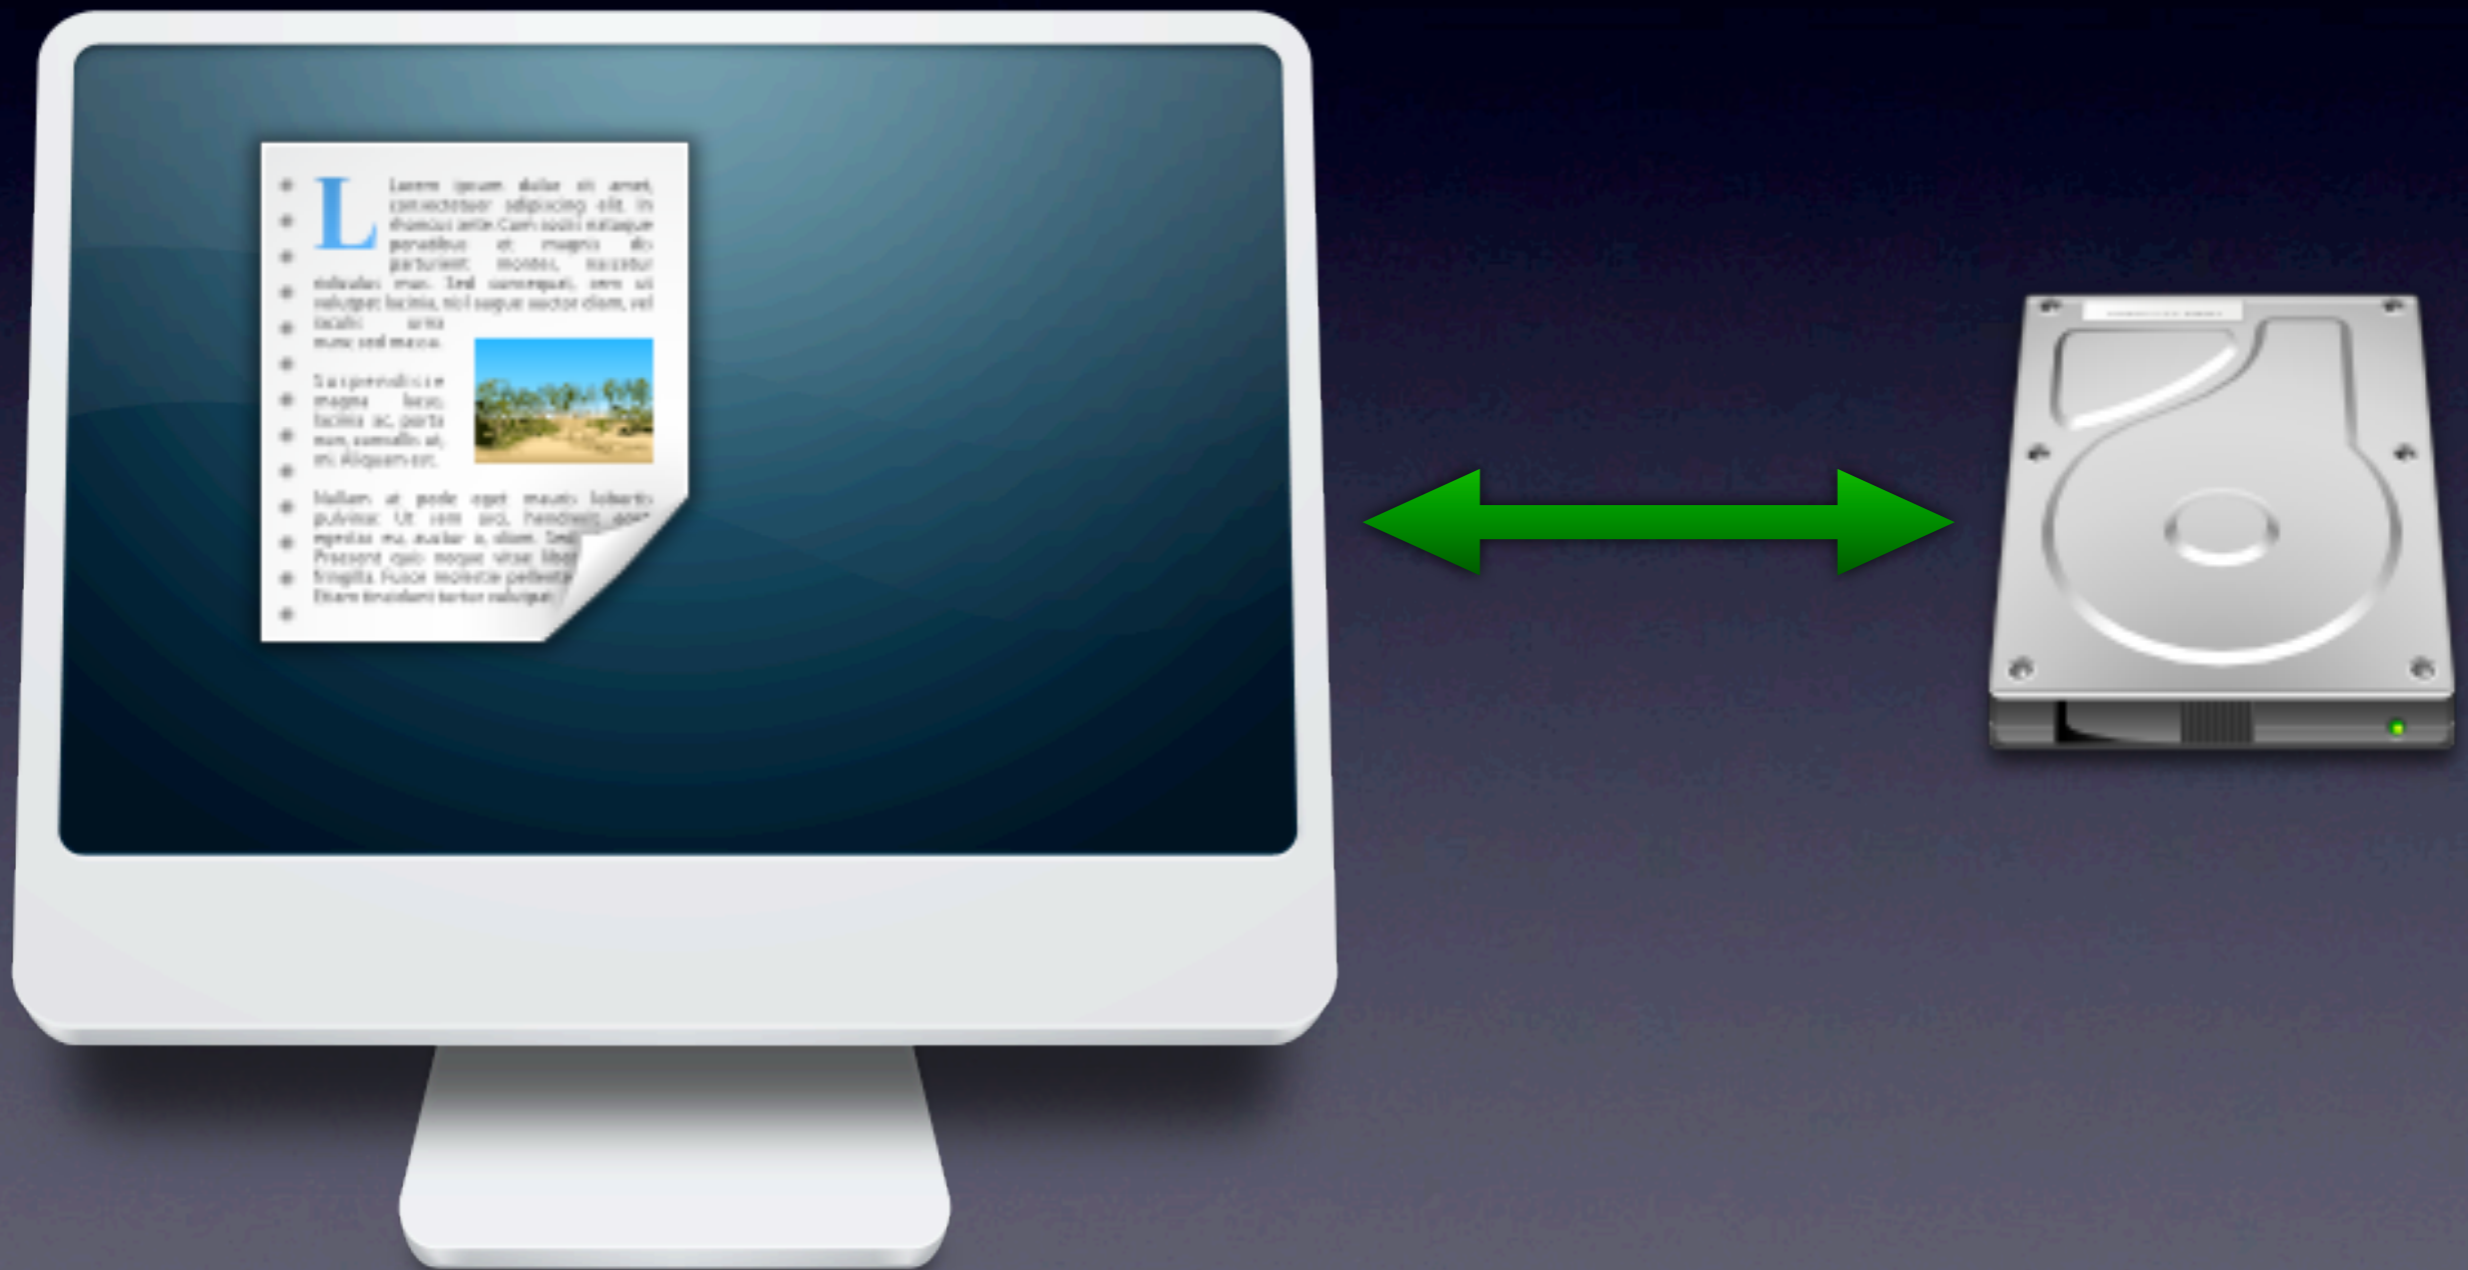
# AHCI

Doing storage right

# About Me

- Alexander Graf
- KVM and Qemu developer
  - Server class PowerPC KVM port
  - S390x Qemu guest support
  - x86 Mac OS X in KVM
  - Nested SVM
  - Xenner
  - ...


# Storage



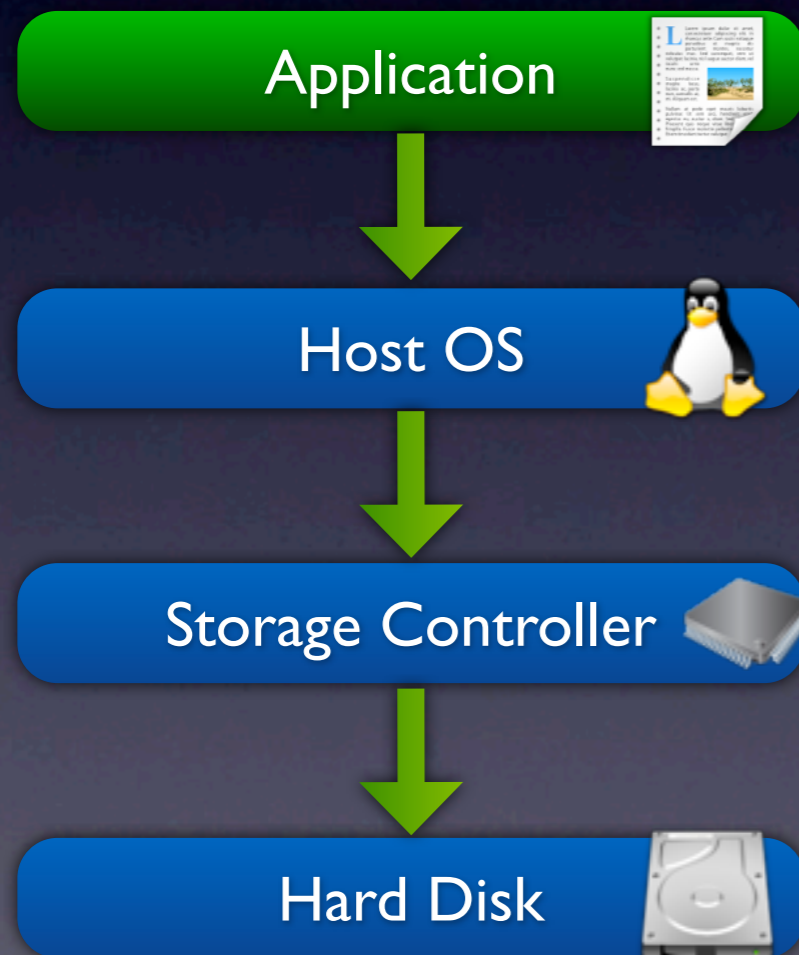
# Naming



# Storage




# Storage




write(  )

# Storage




# Storage

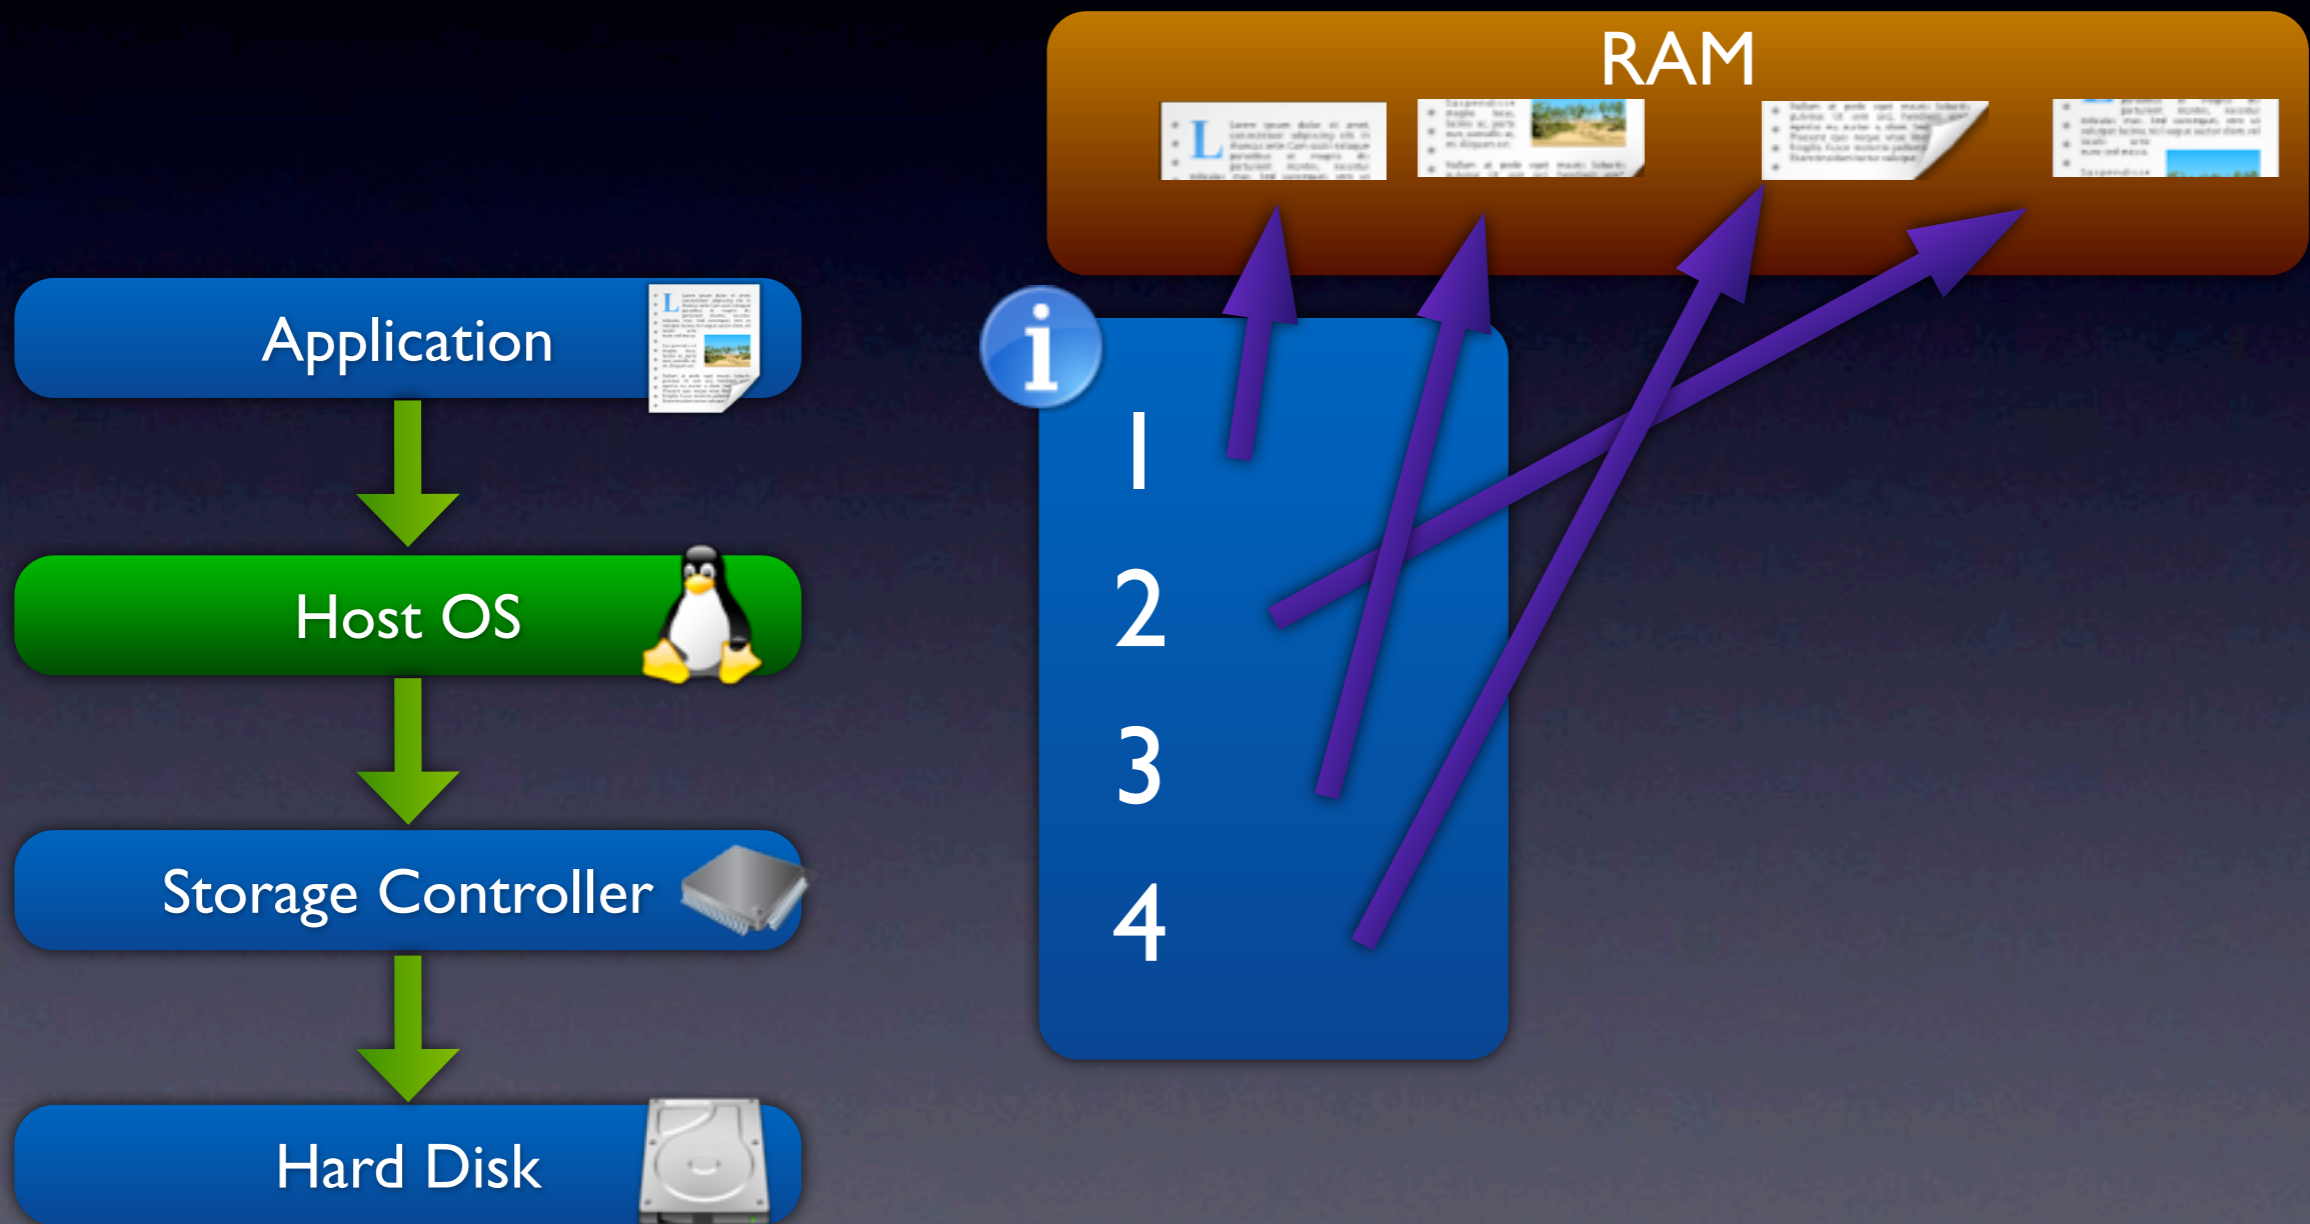




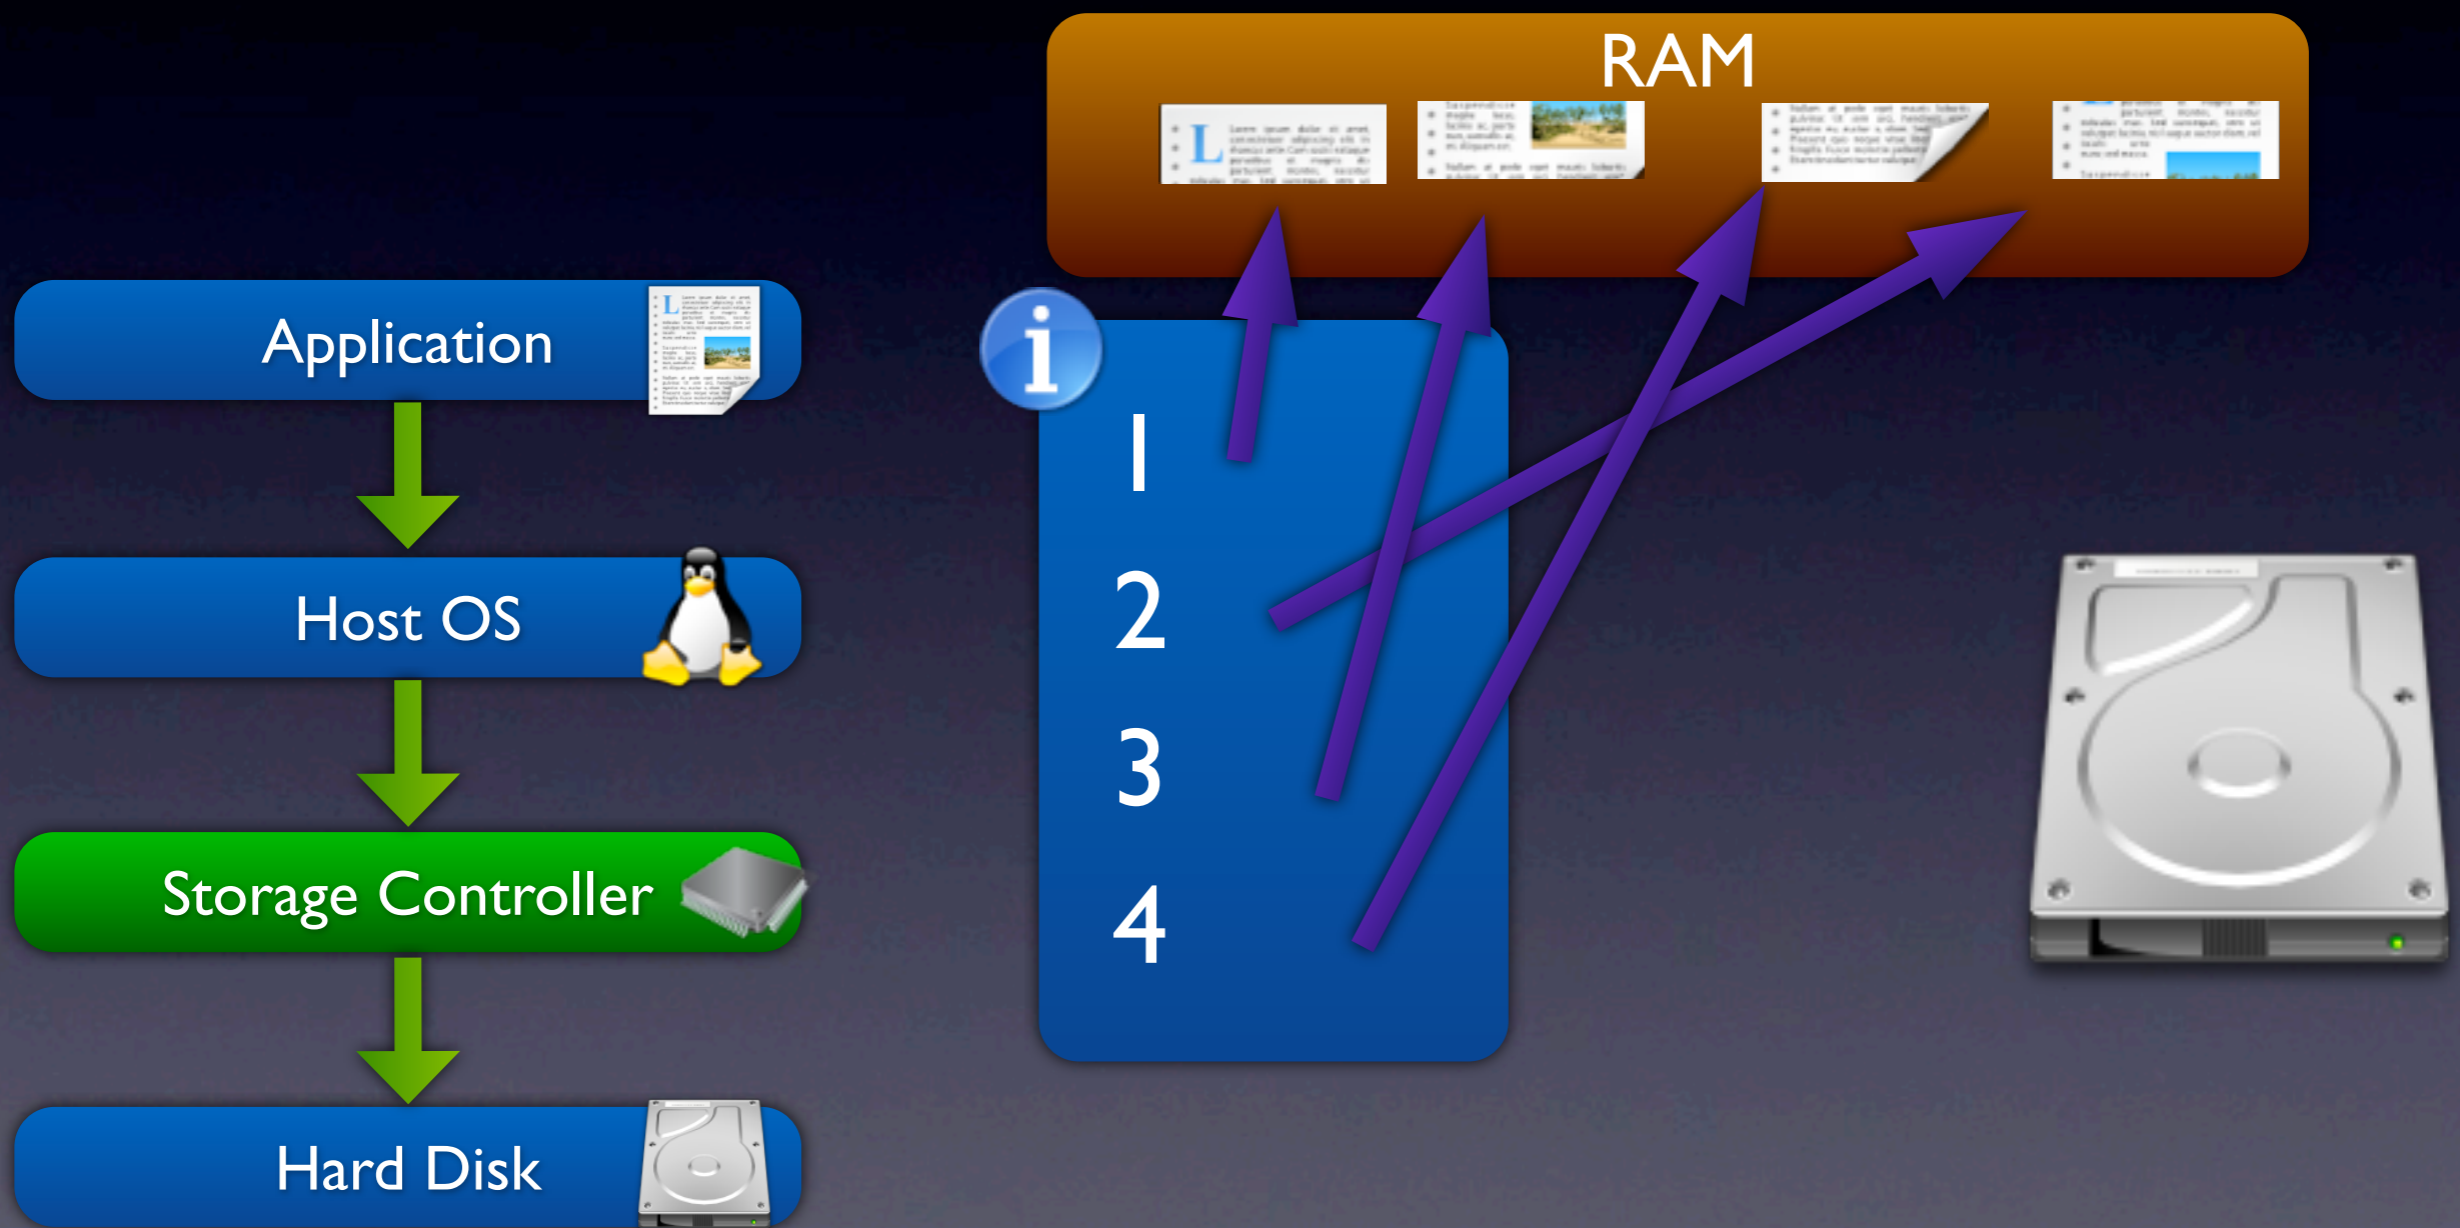
# Storage



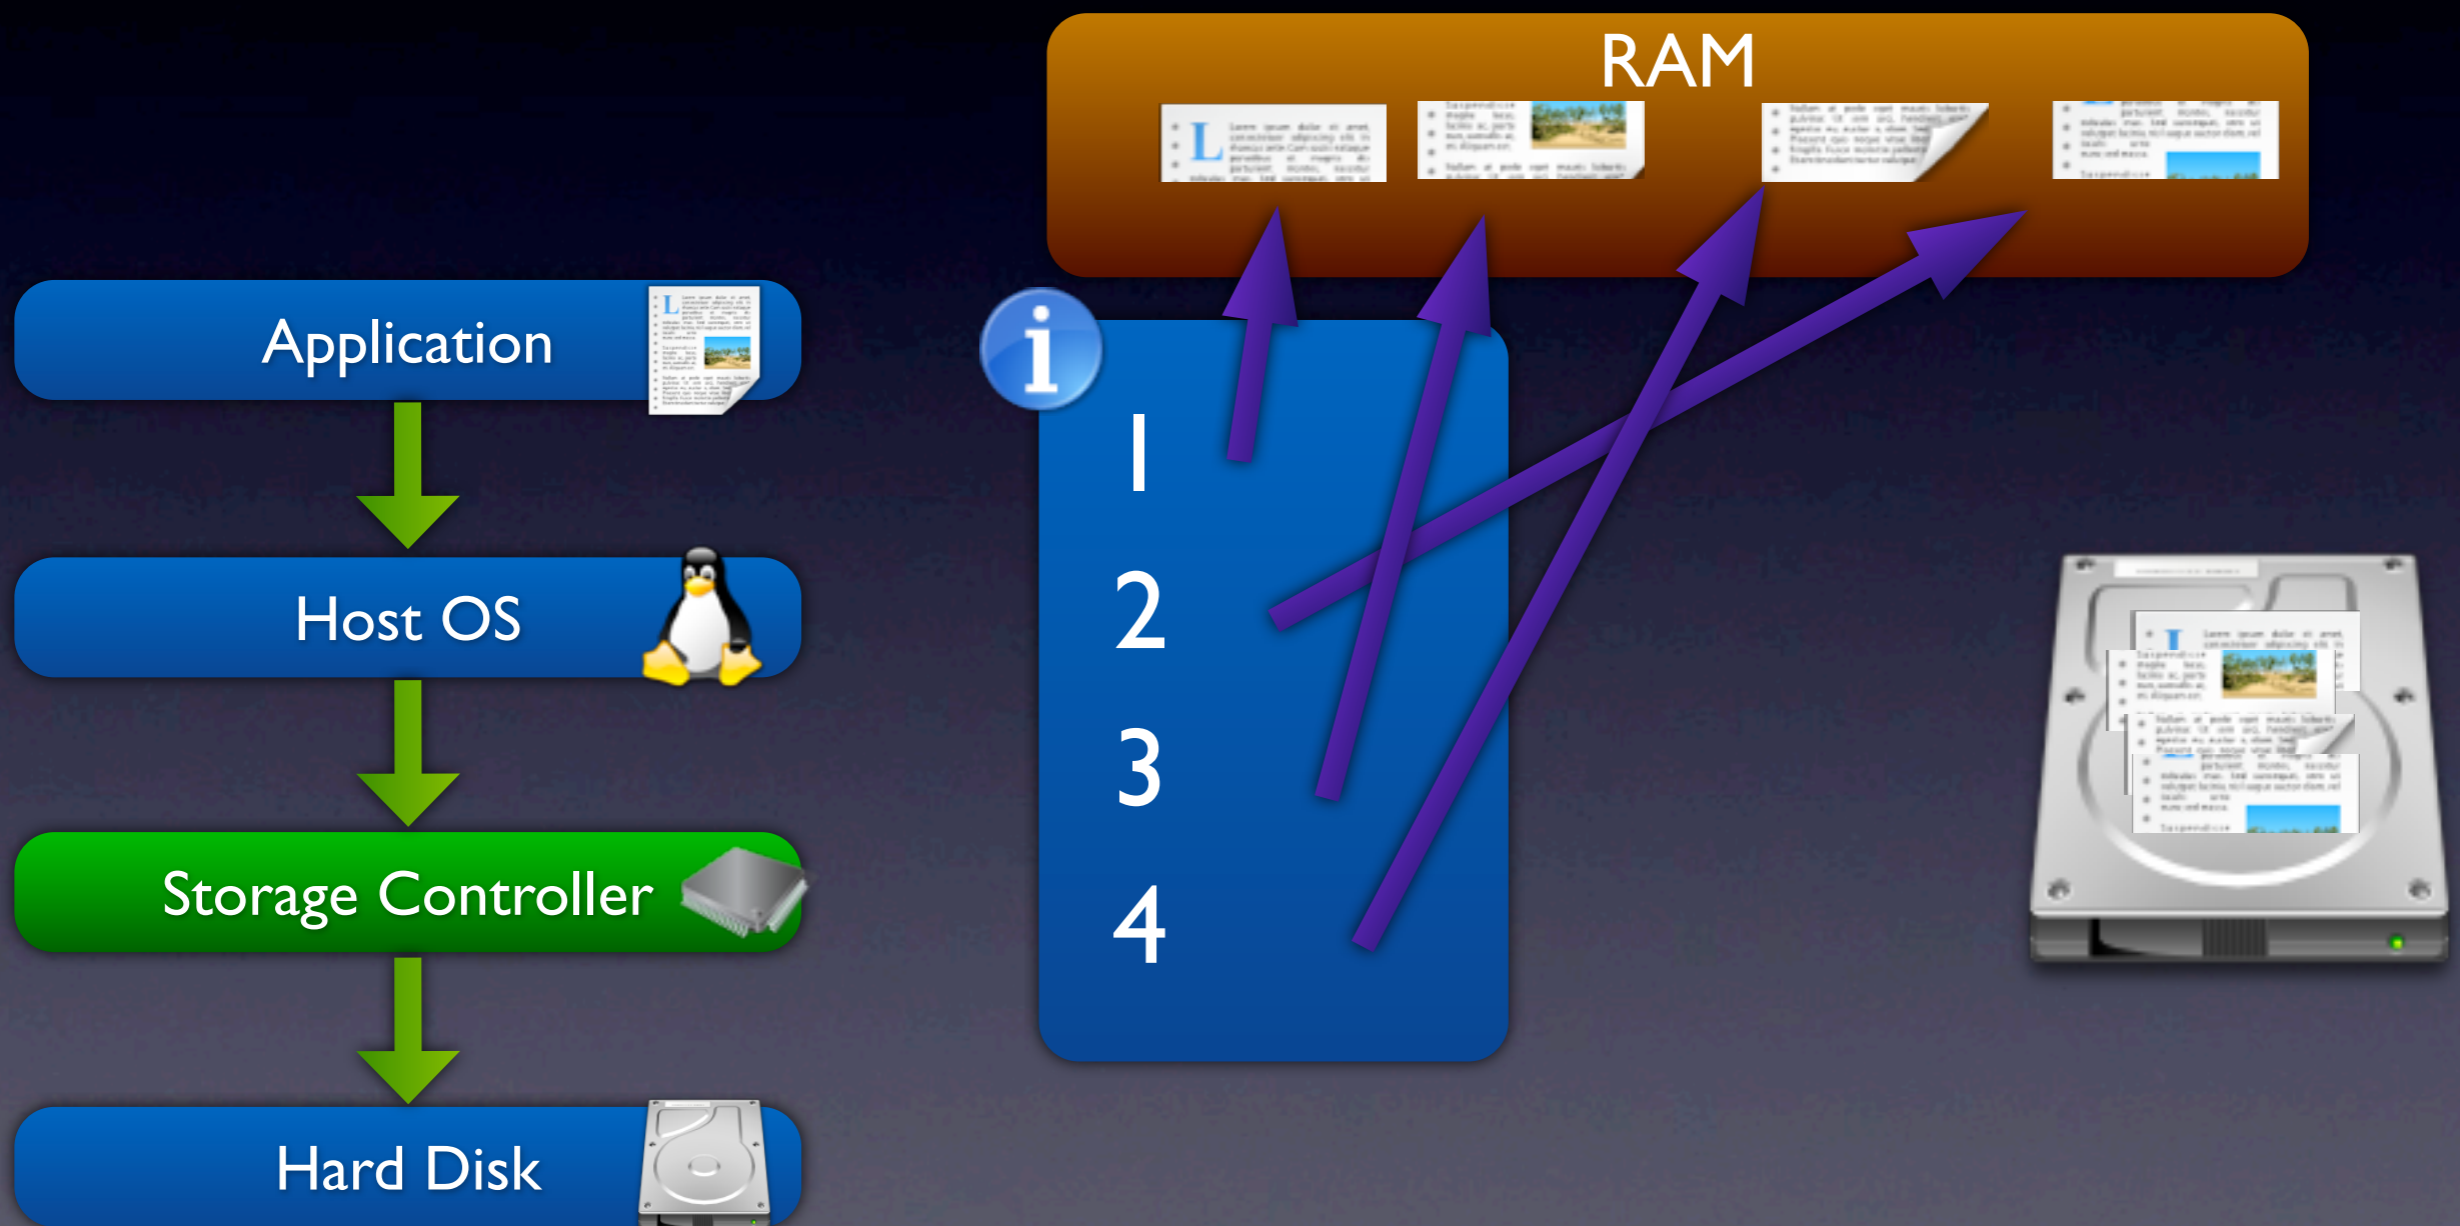
# Storage



# Storage



# Storage




# Virtualization




# Virtualization



# Virtual Storage

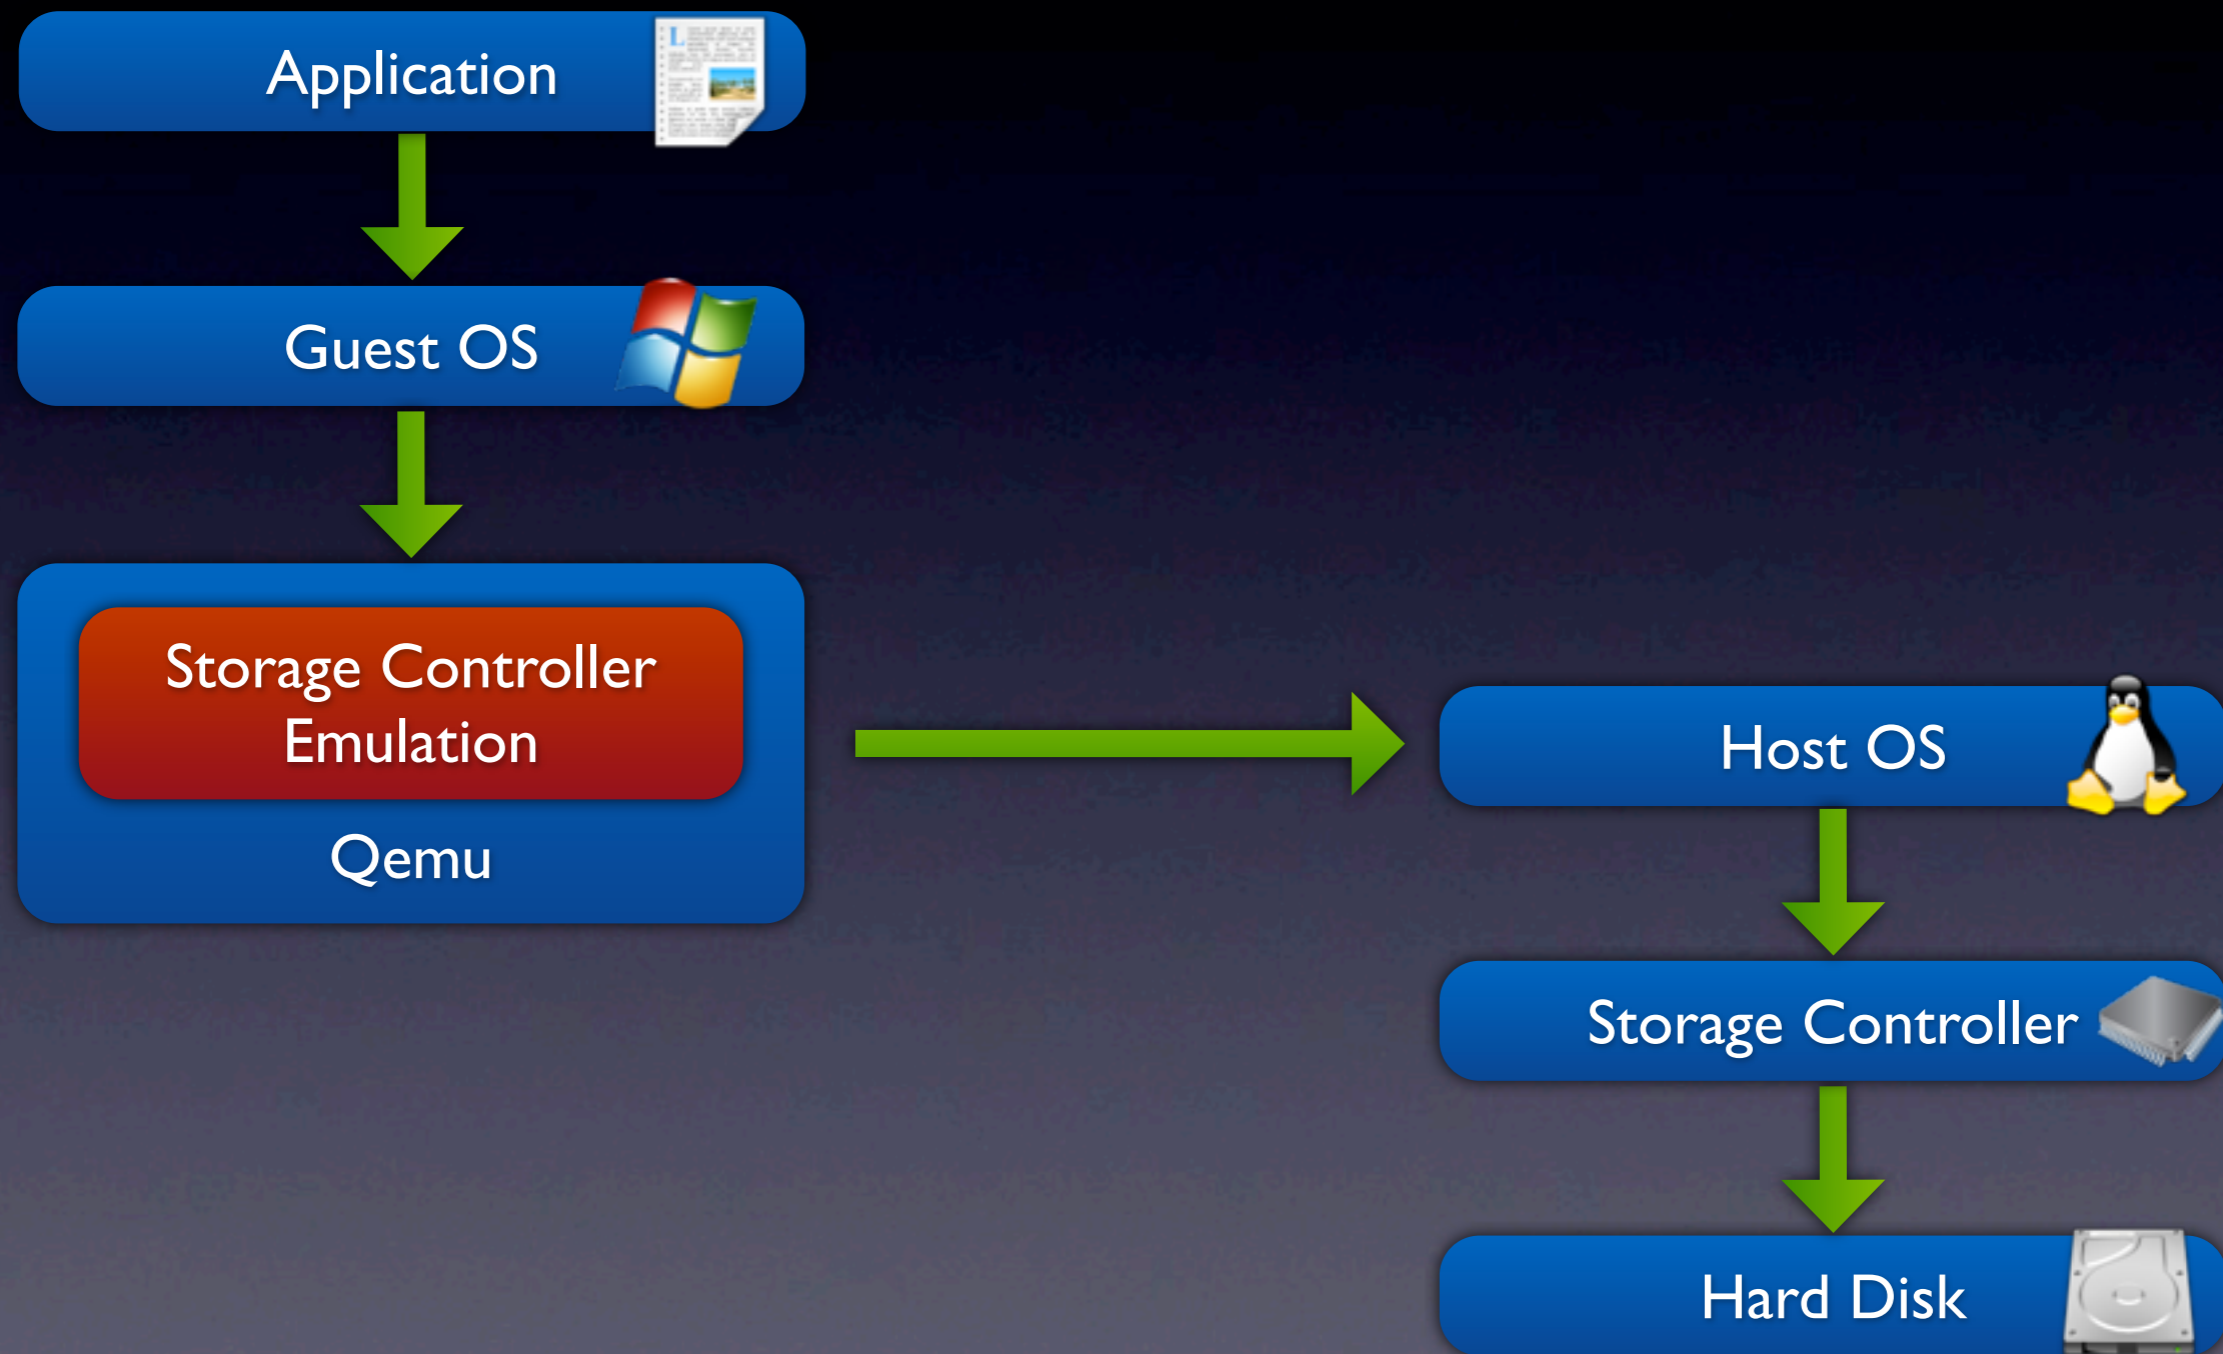


# Virtualized Storage

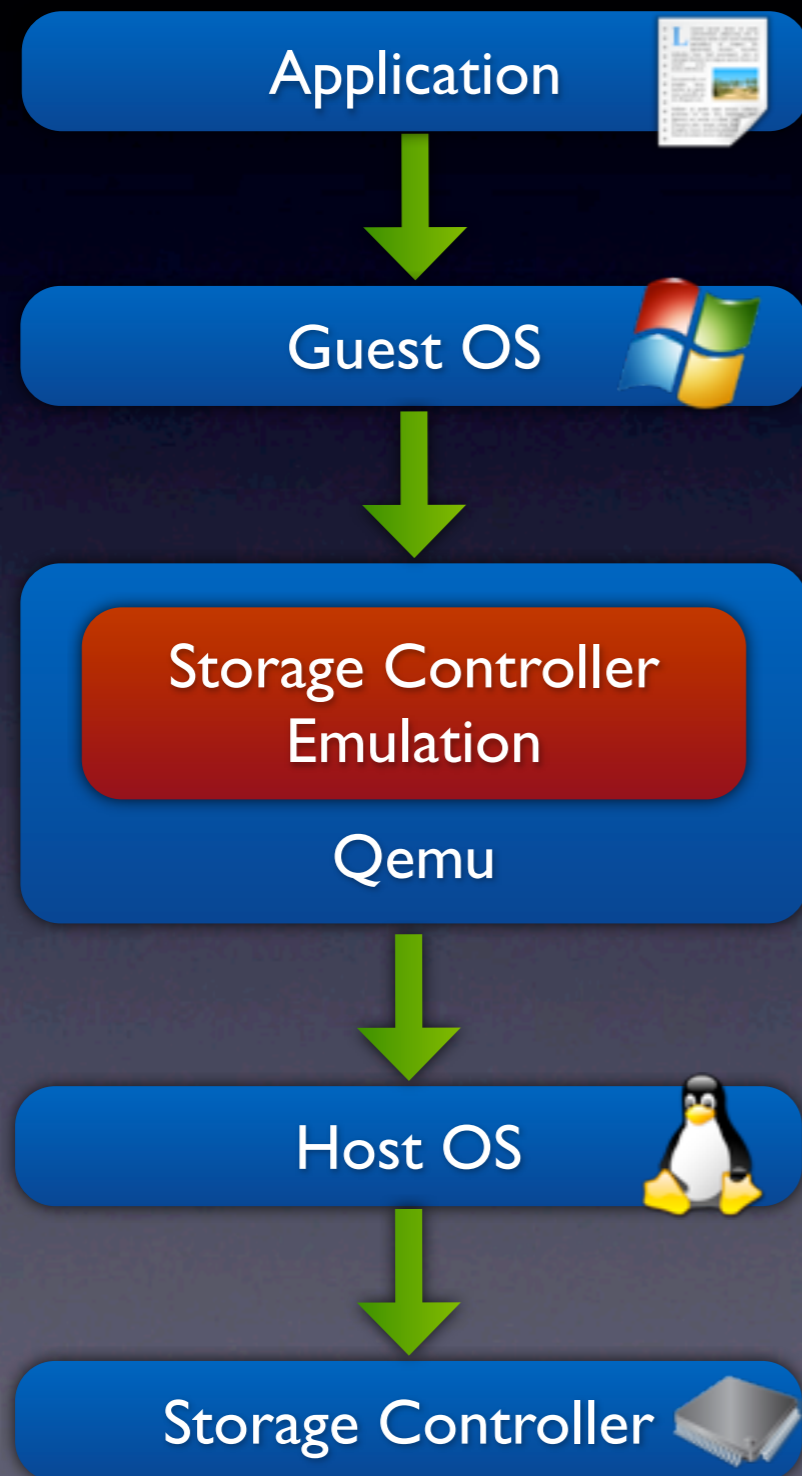




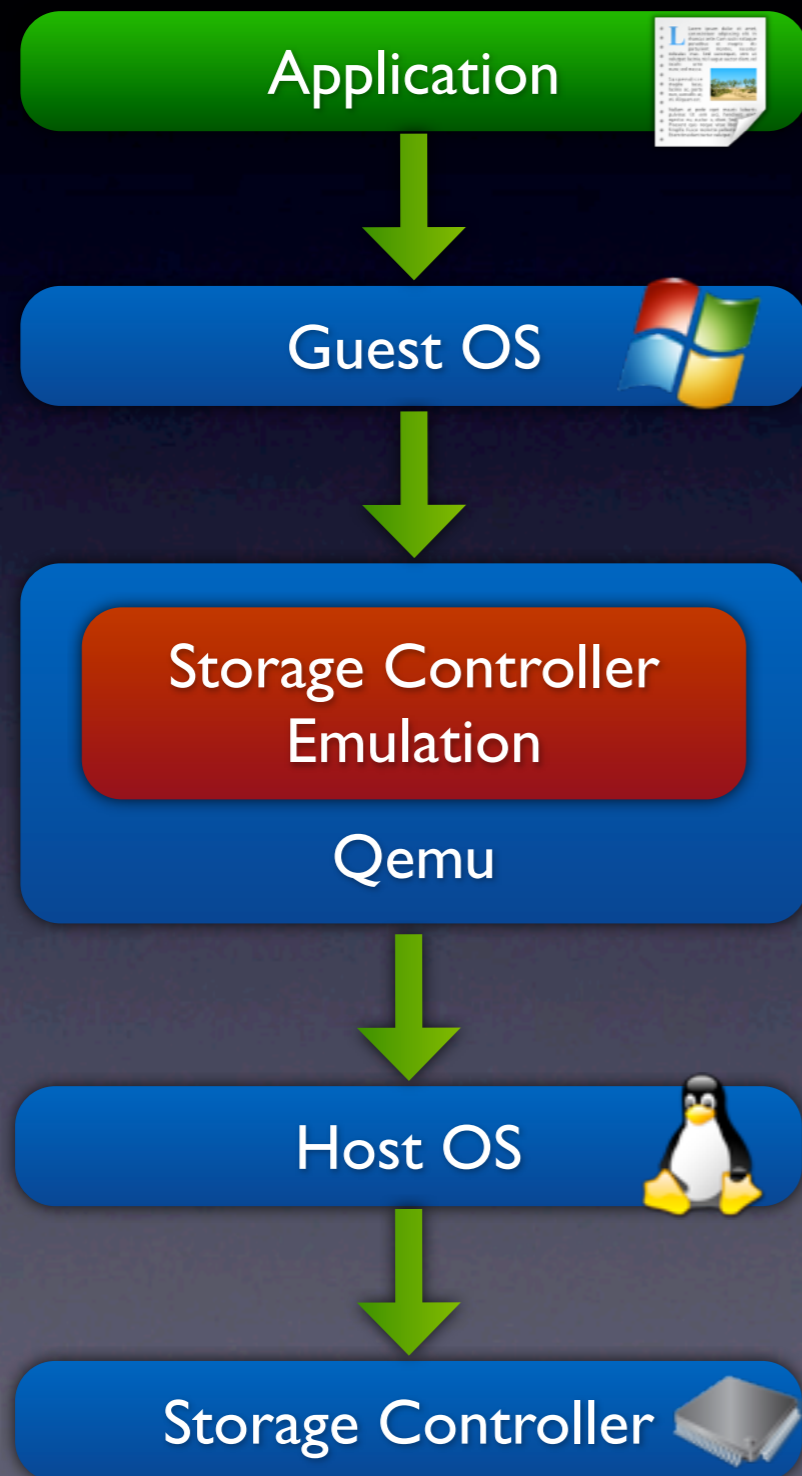
# Virtualized Storage



# Virtualized Storage

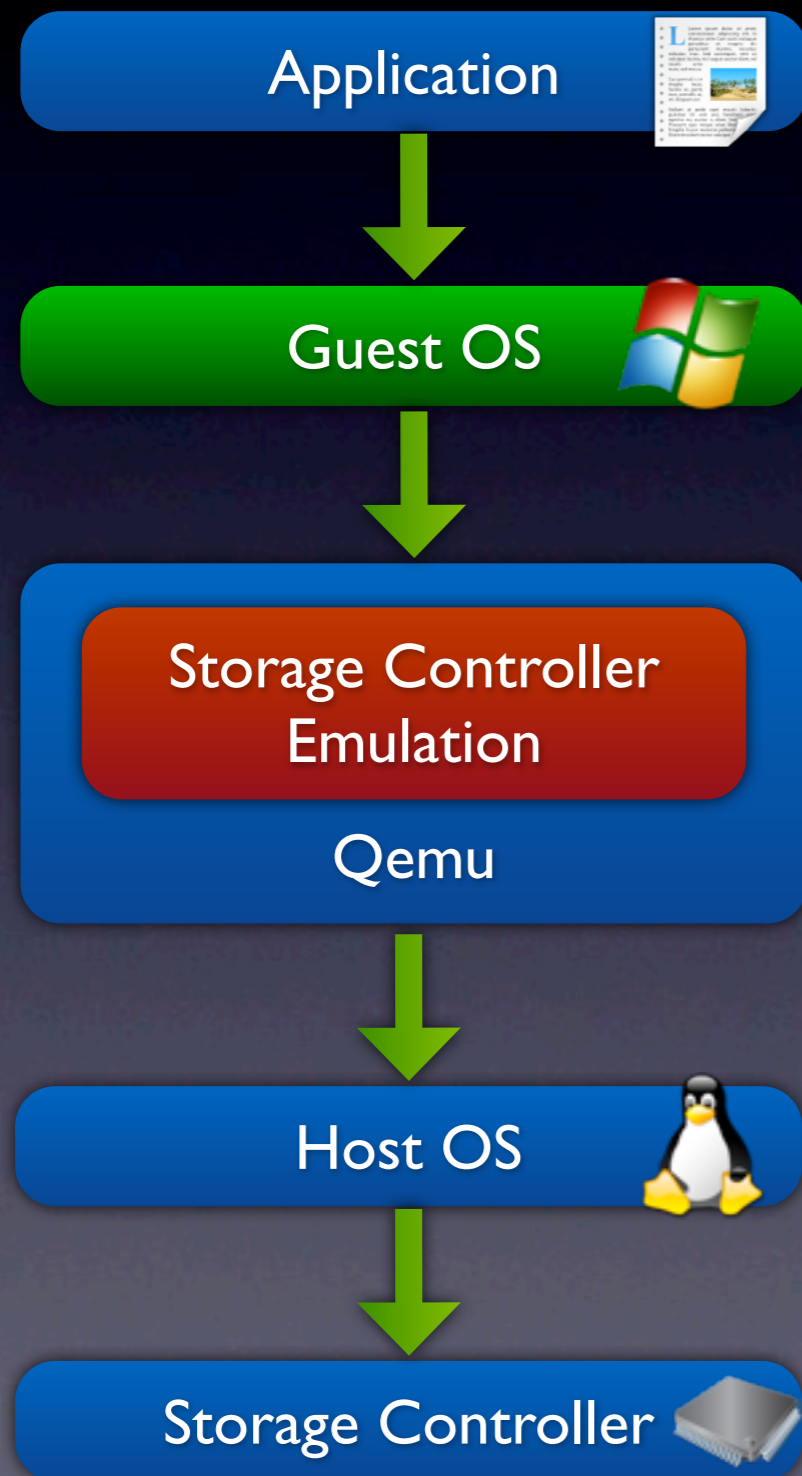


# Virtualized Storage

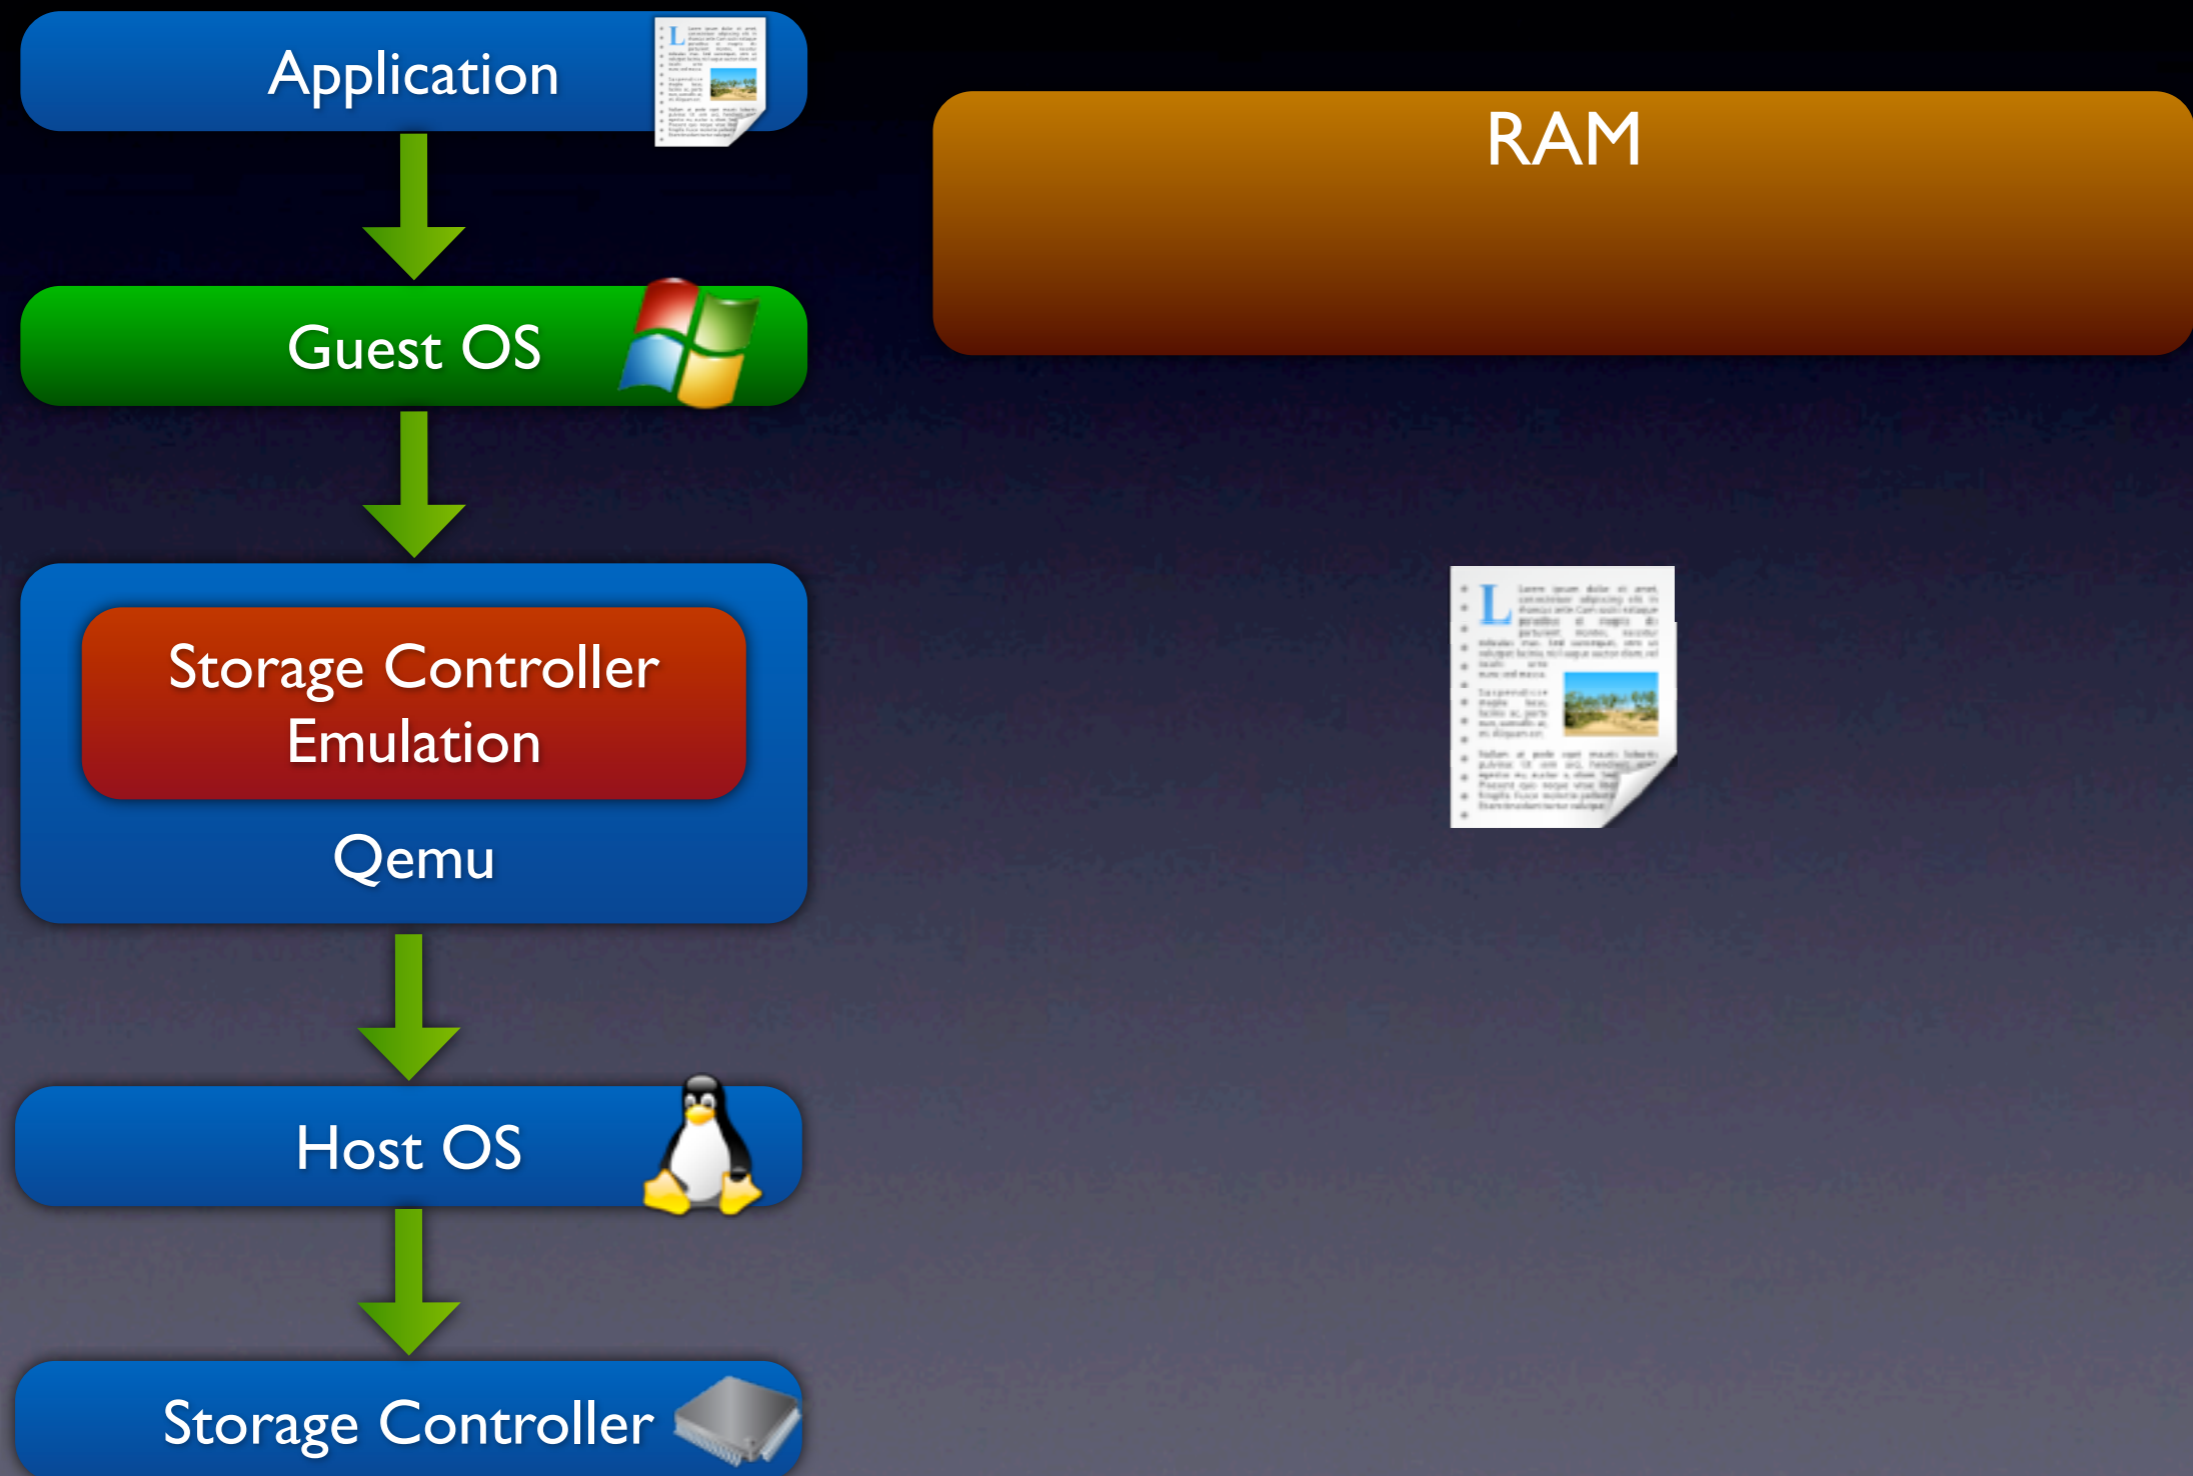


write(  )

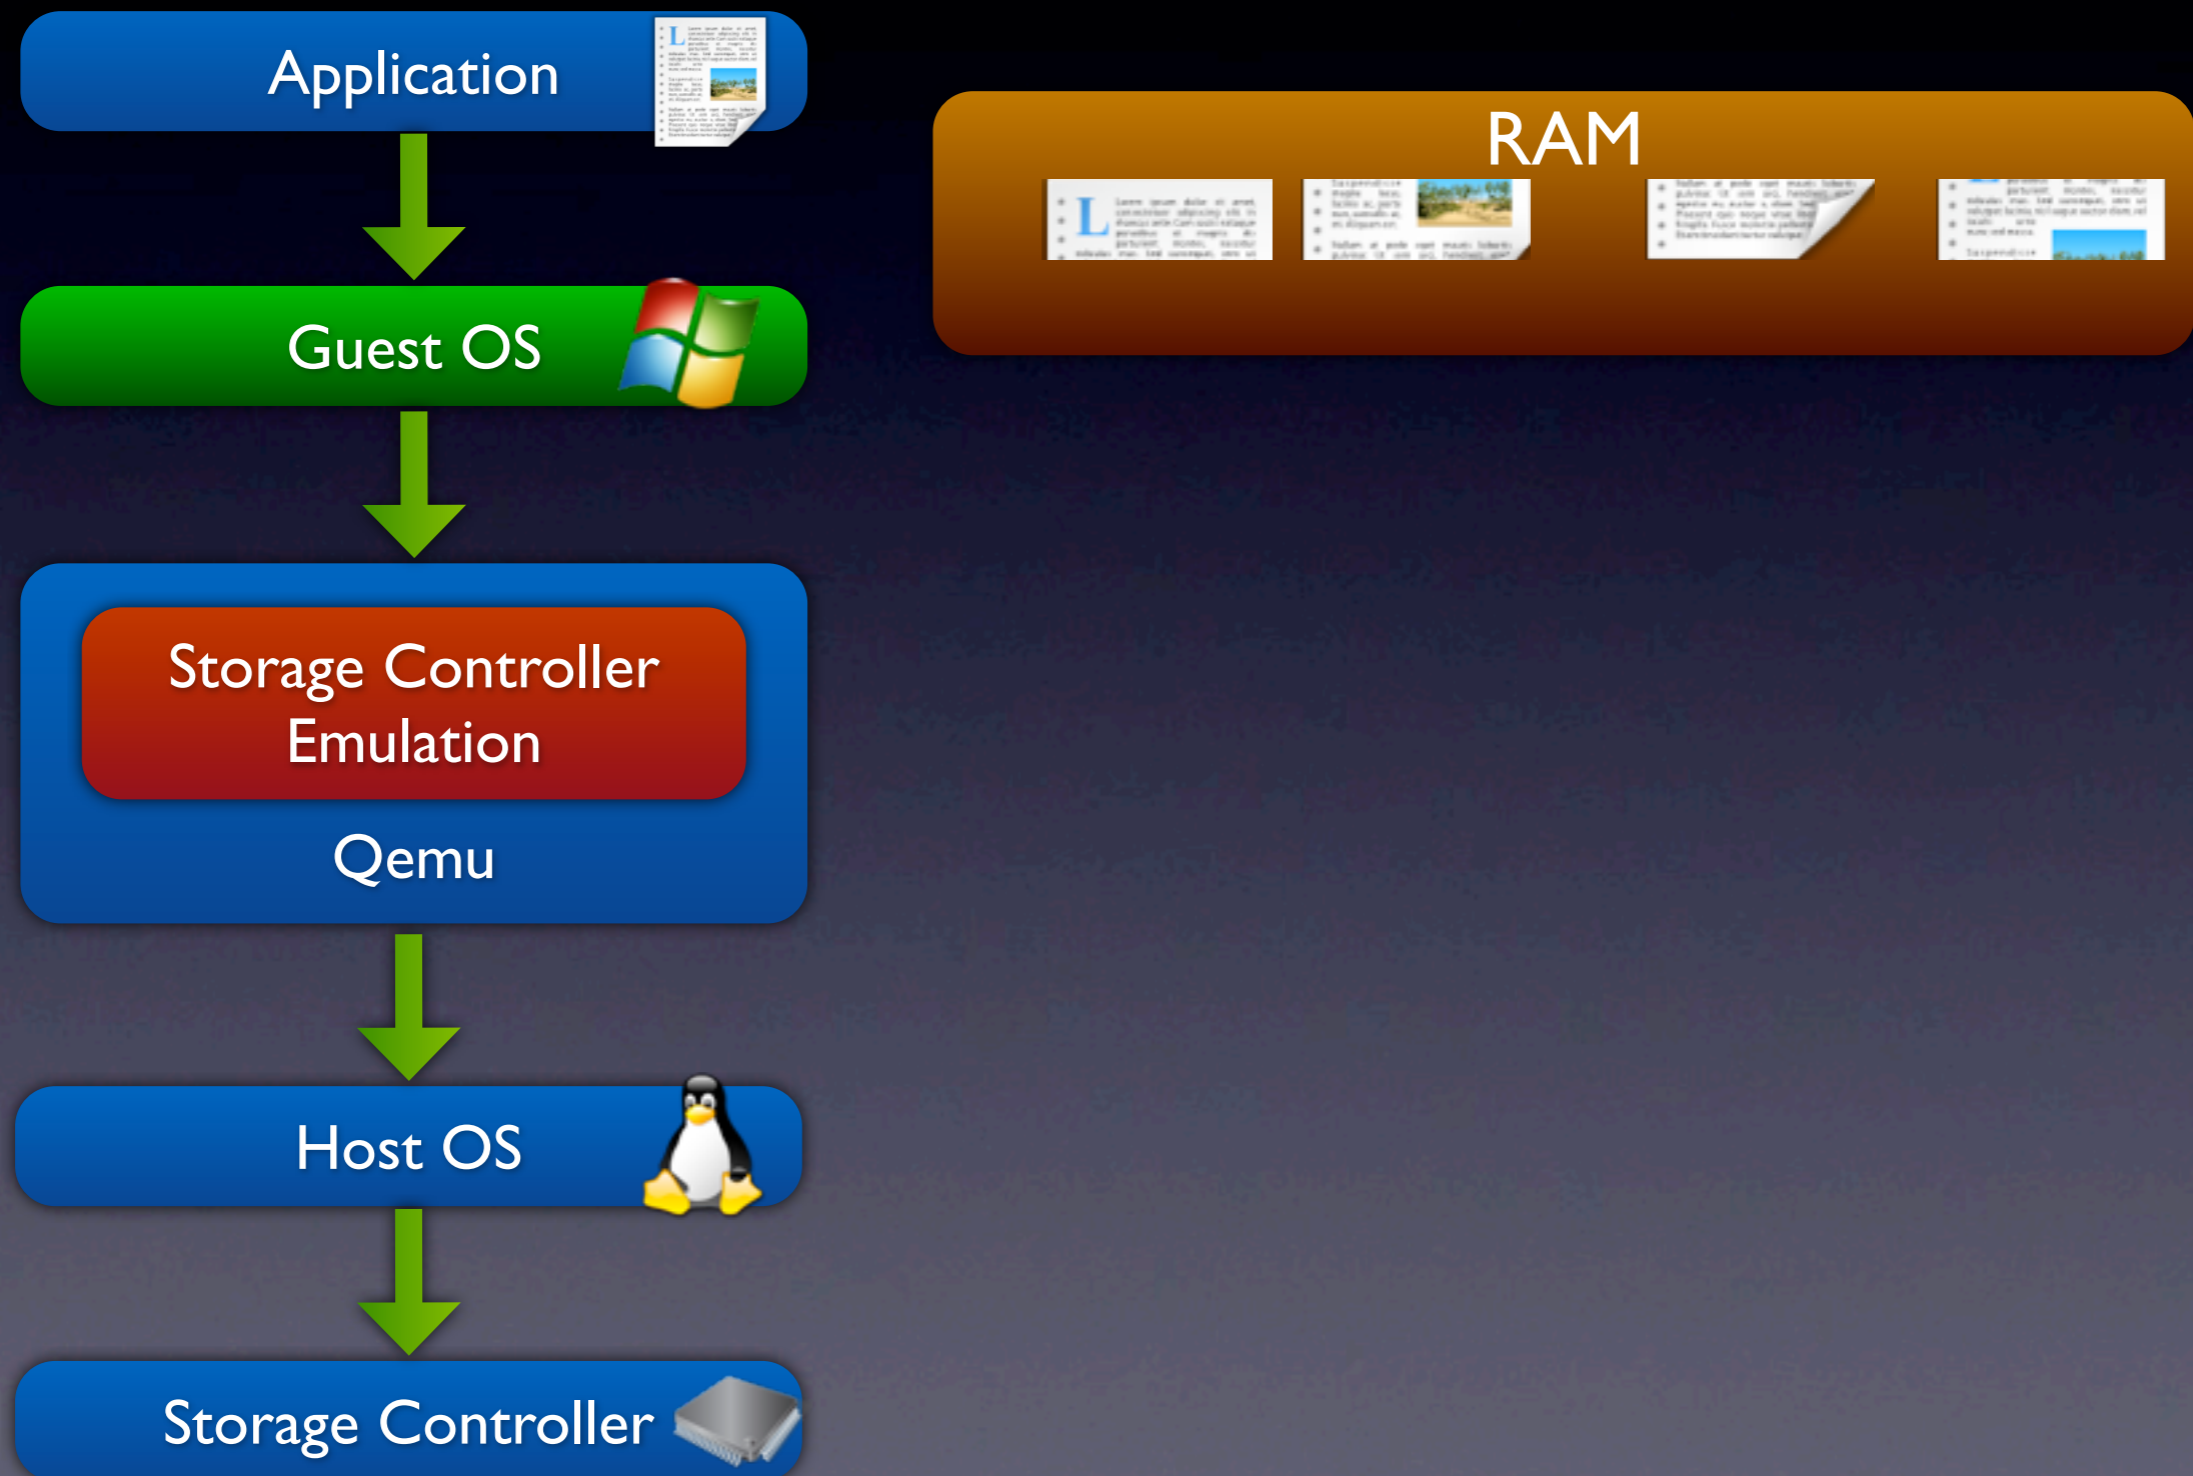
# Virtualized Storage



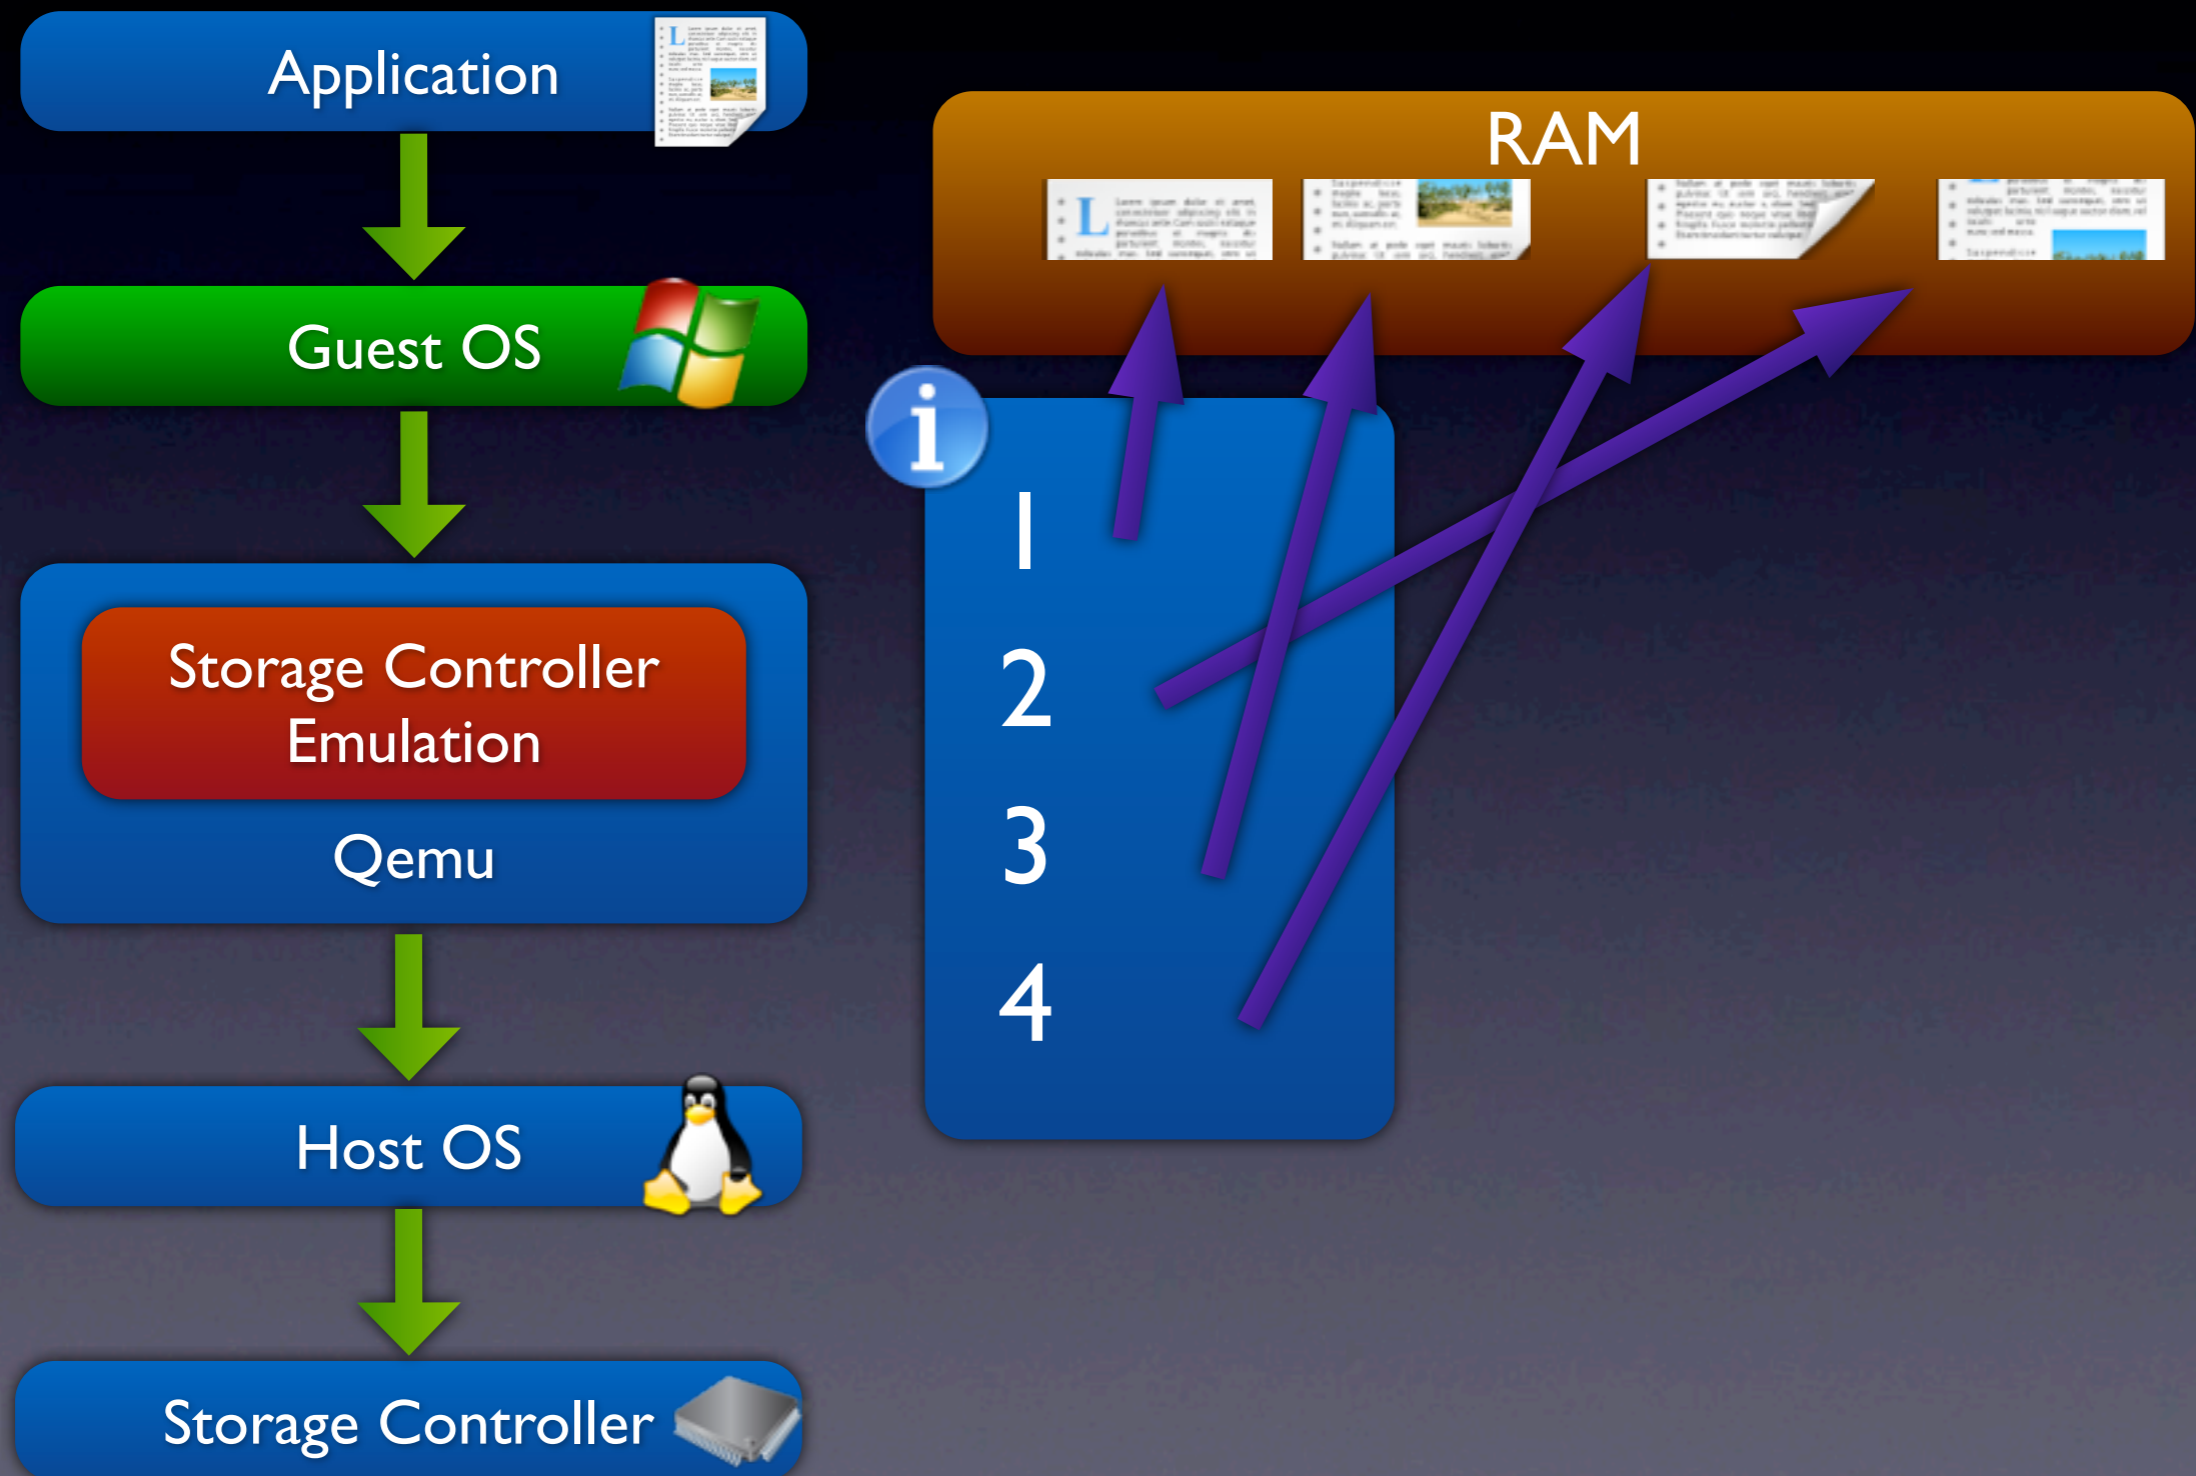
# Virtualized Storage



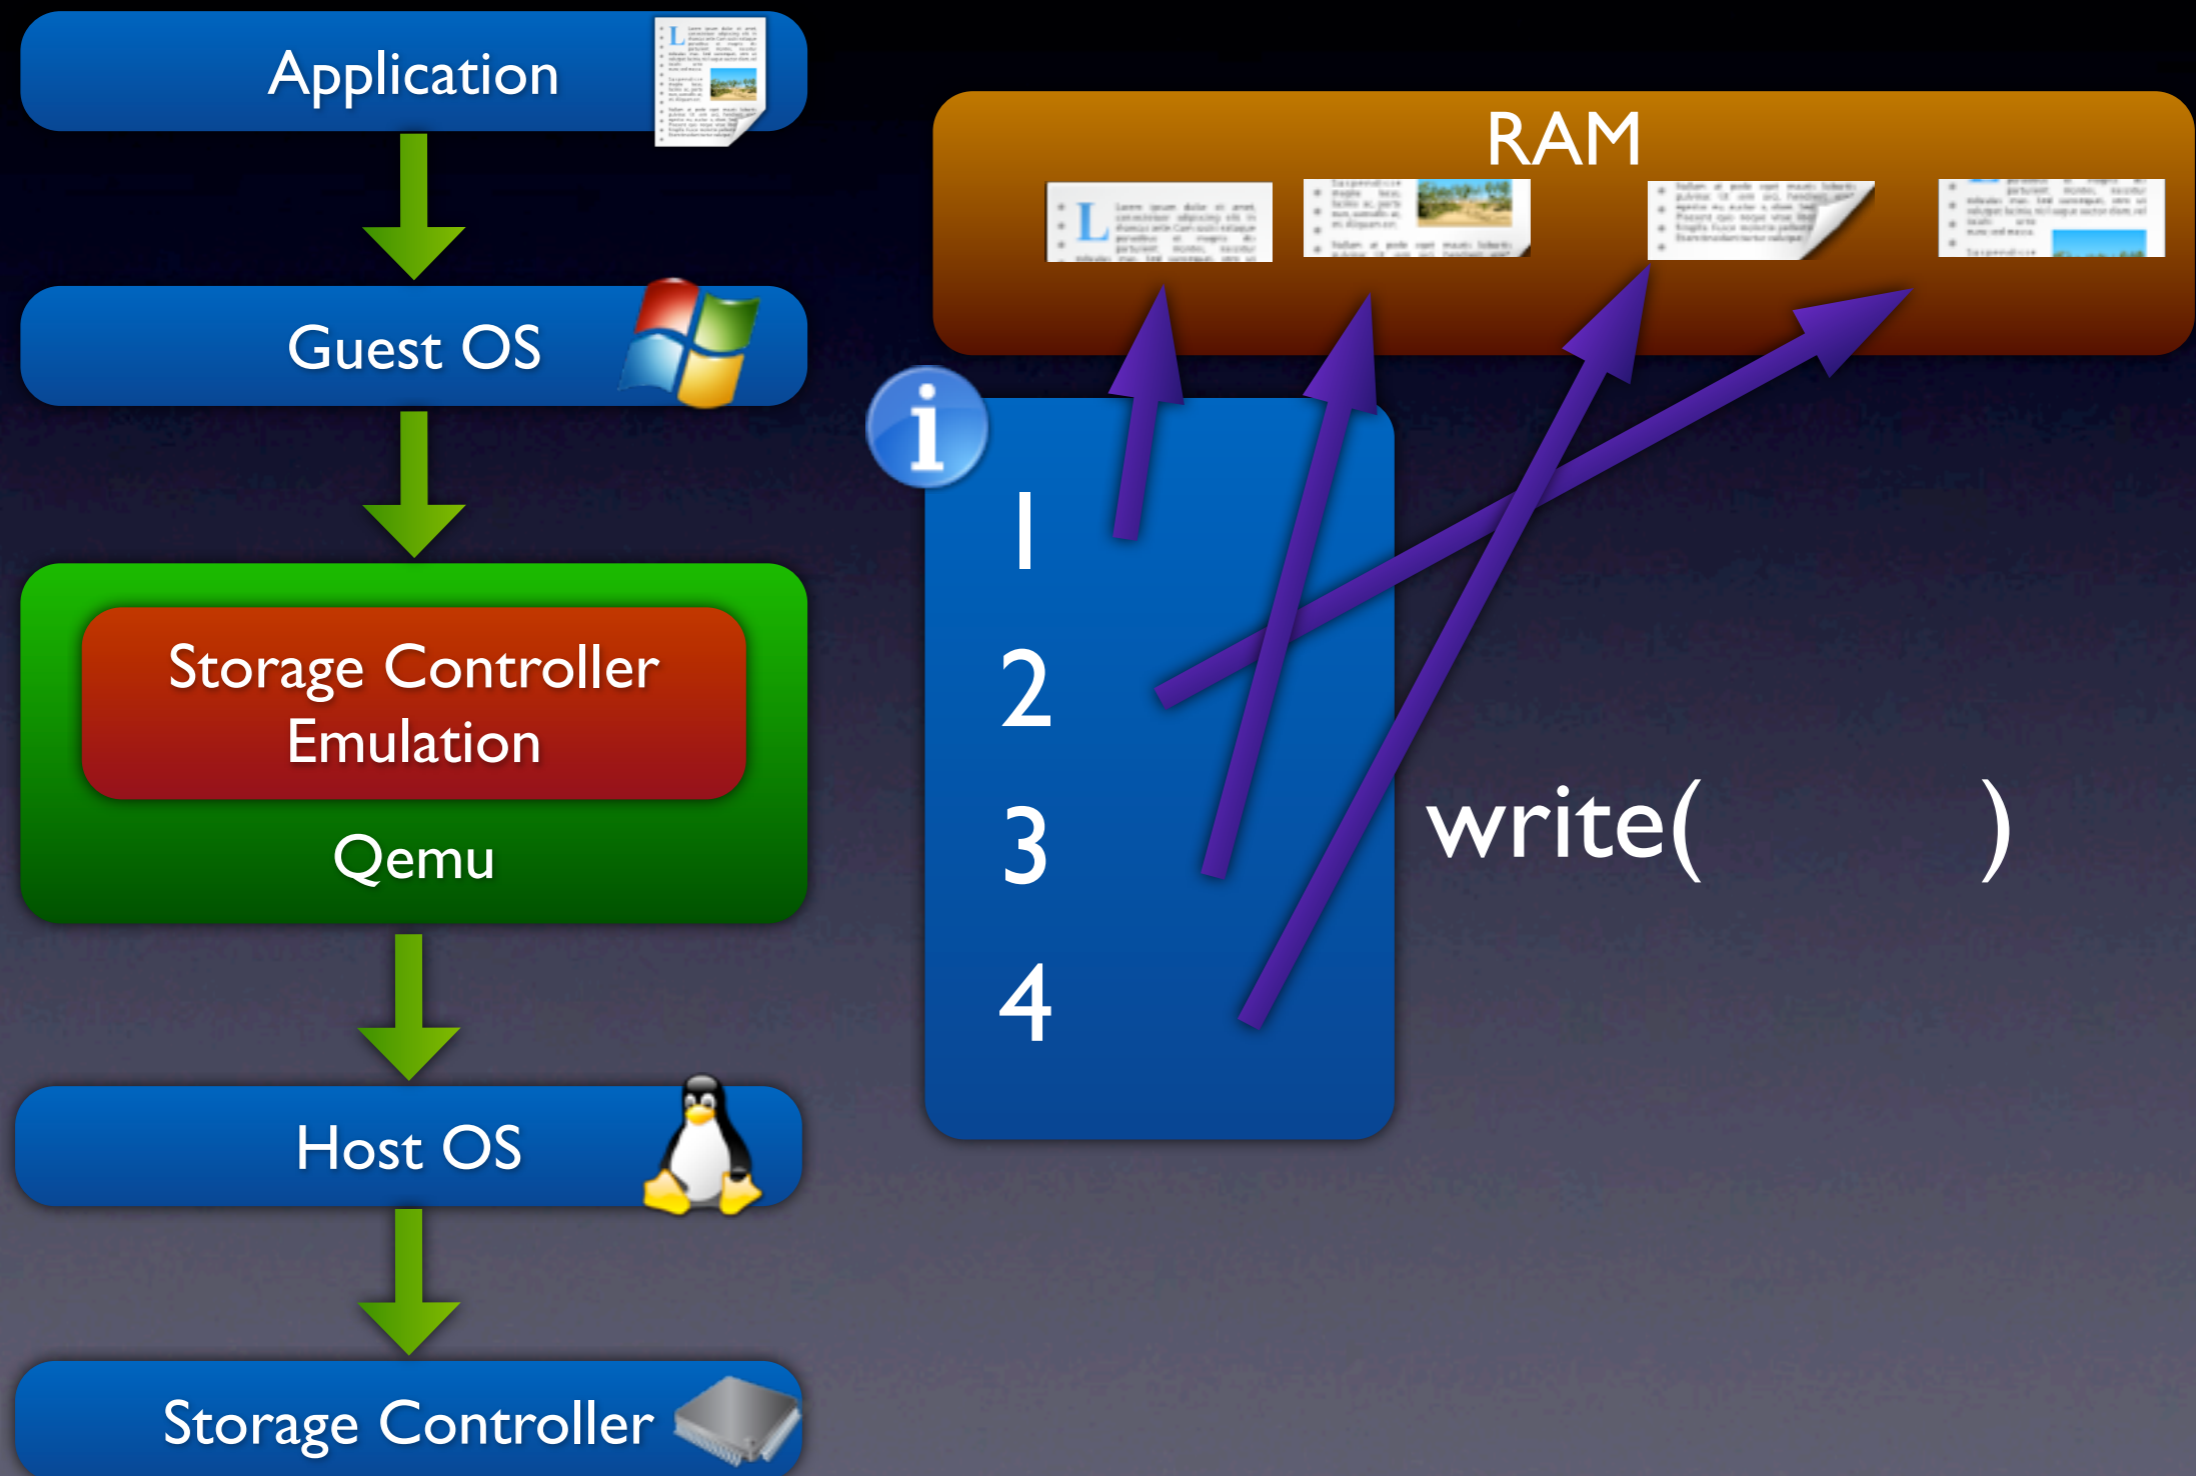
# Virtualized Storage



# Virtualized Storage

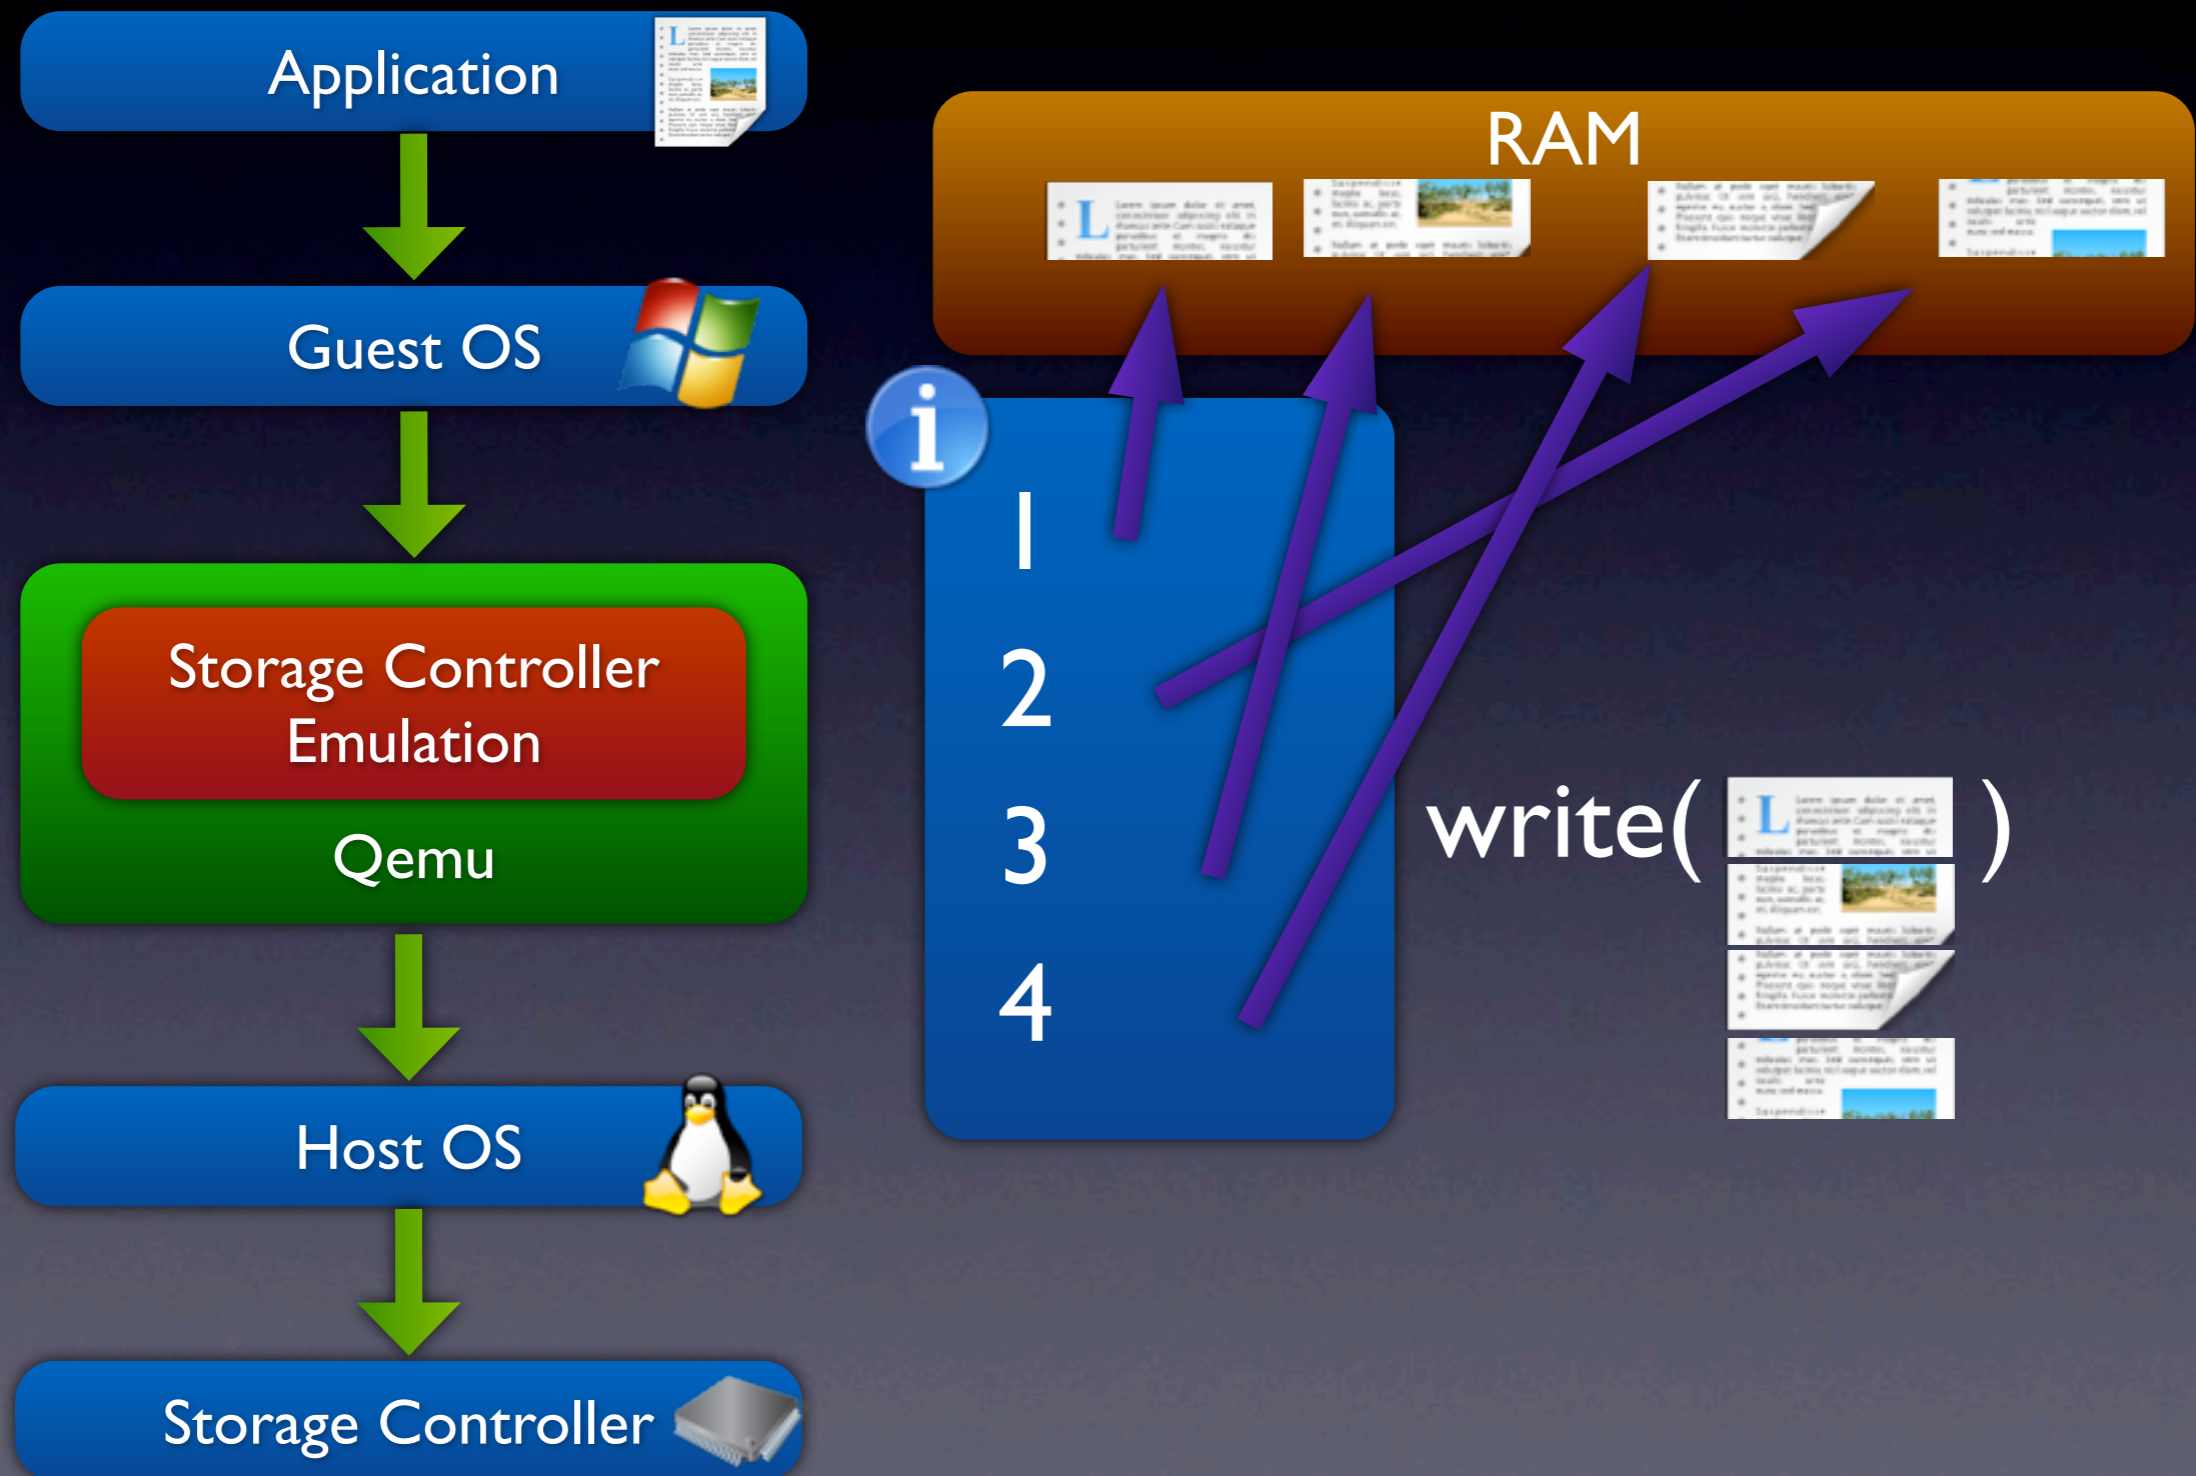


# Virtualized Storage


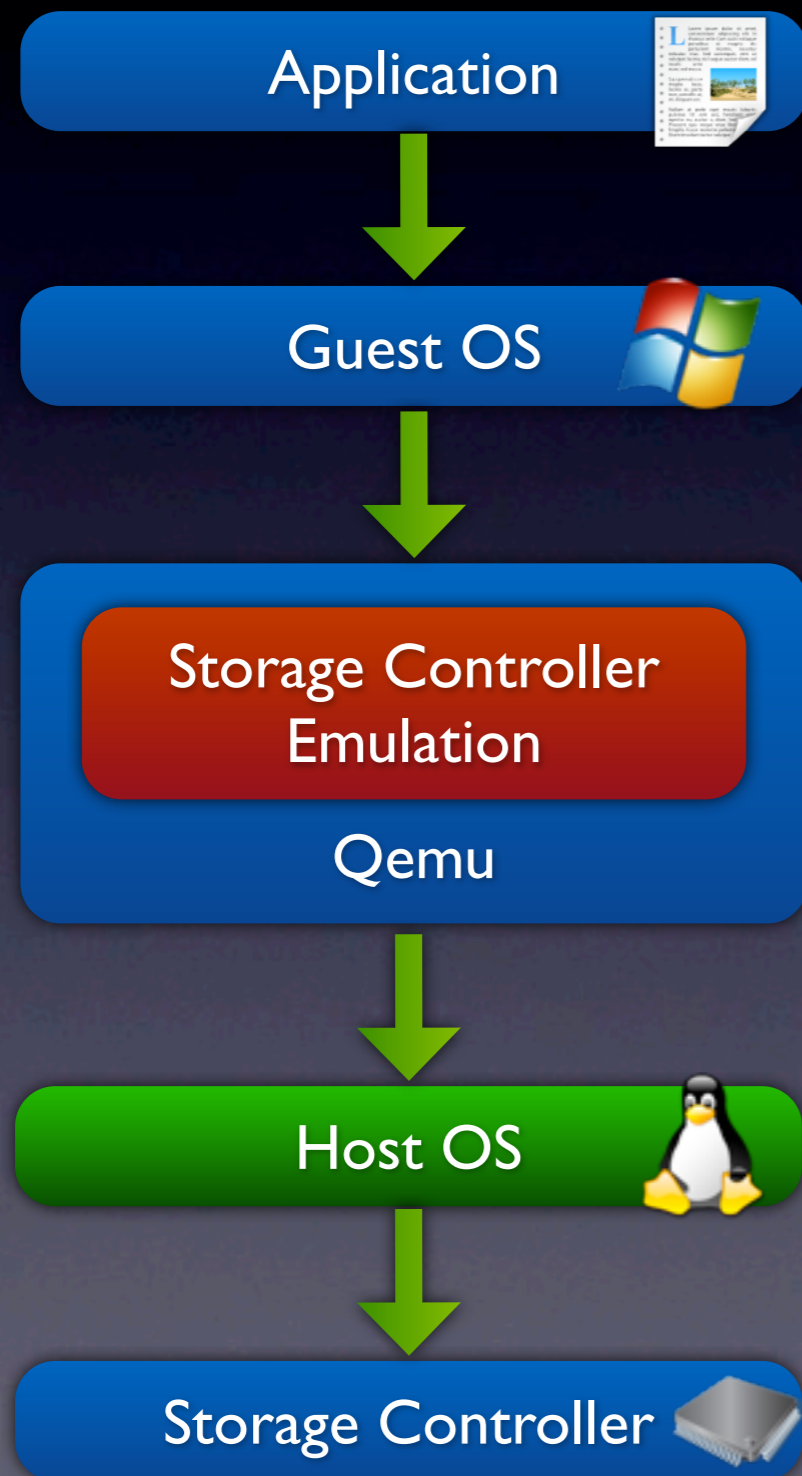




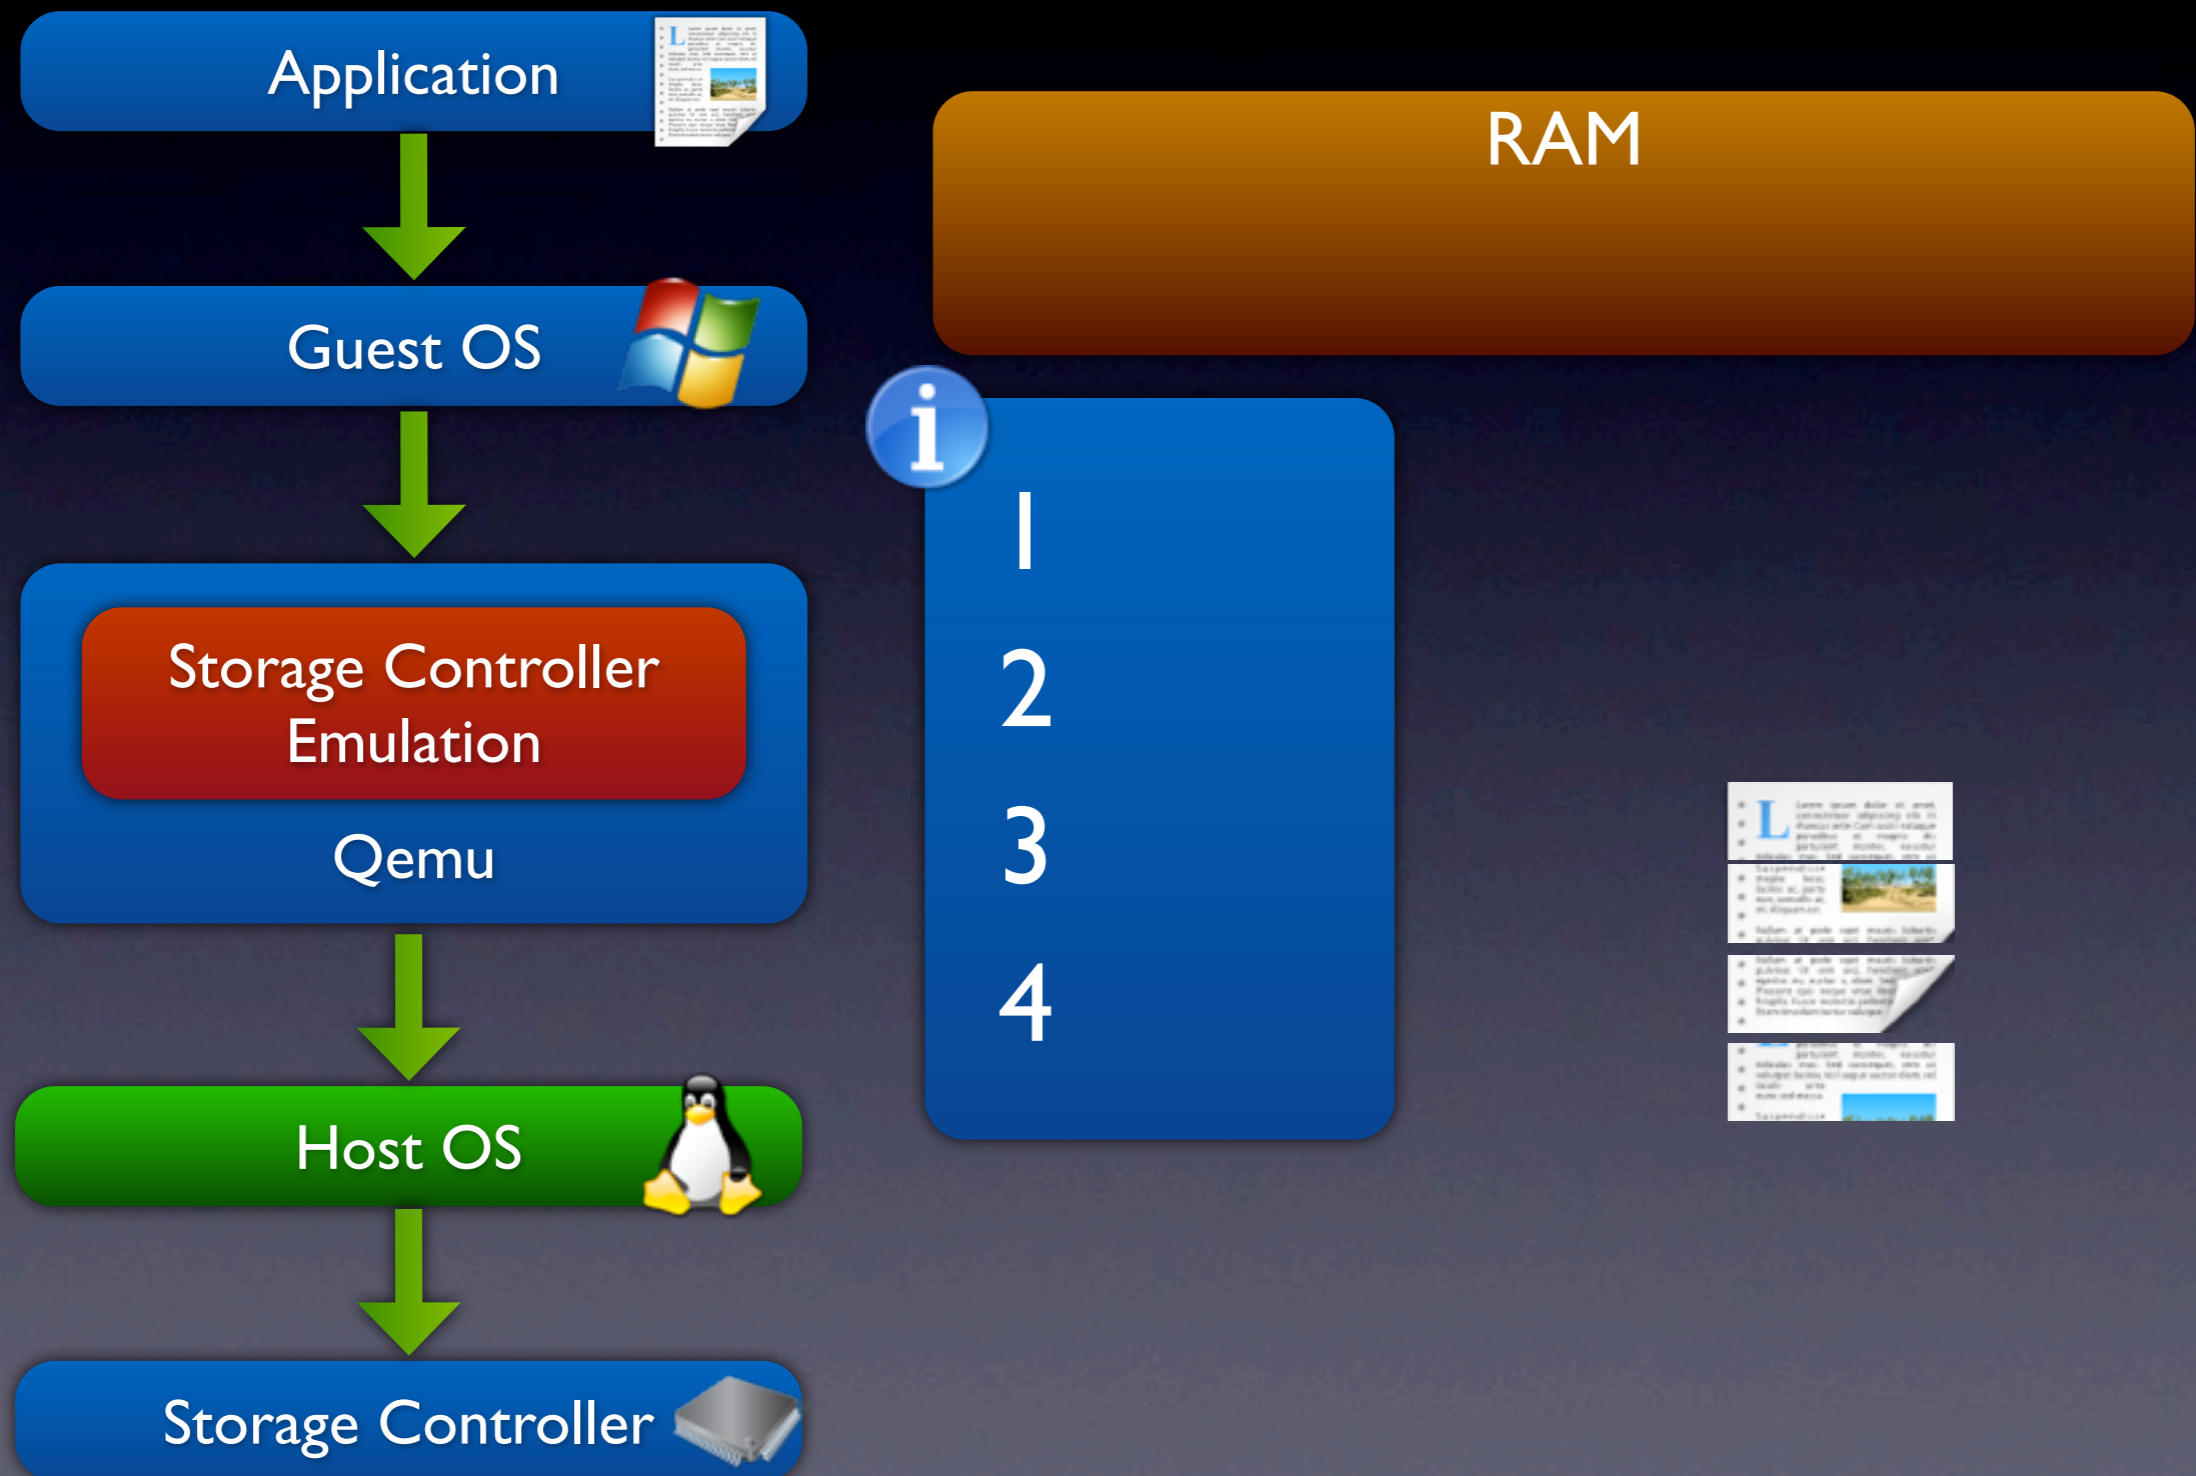
# Virtualized Storage



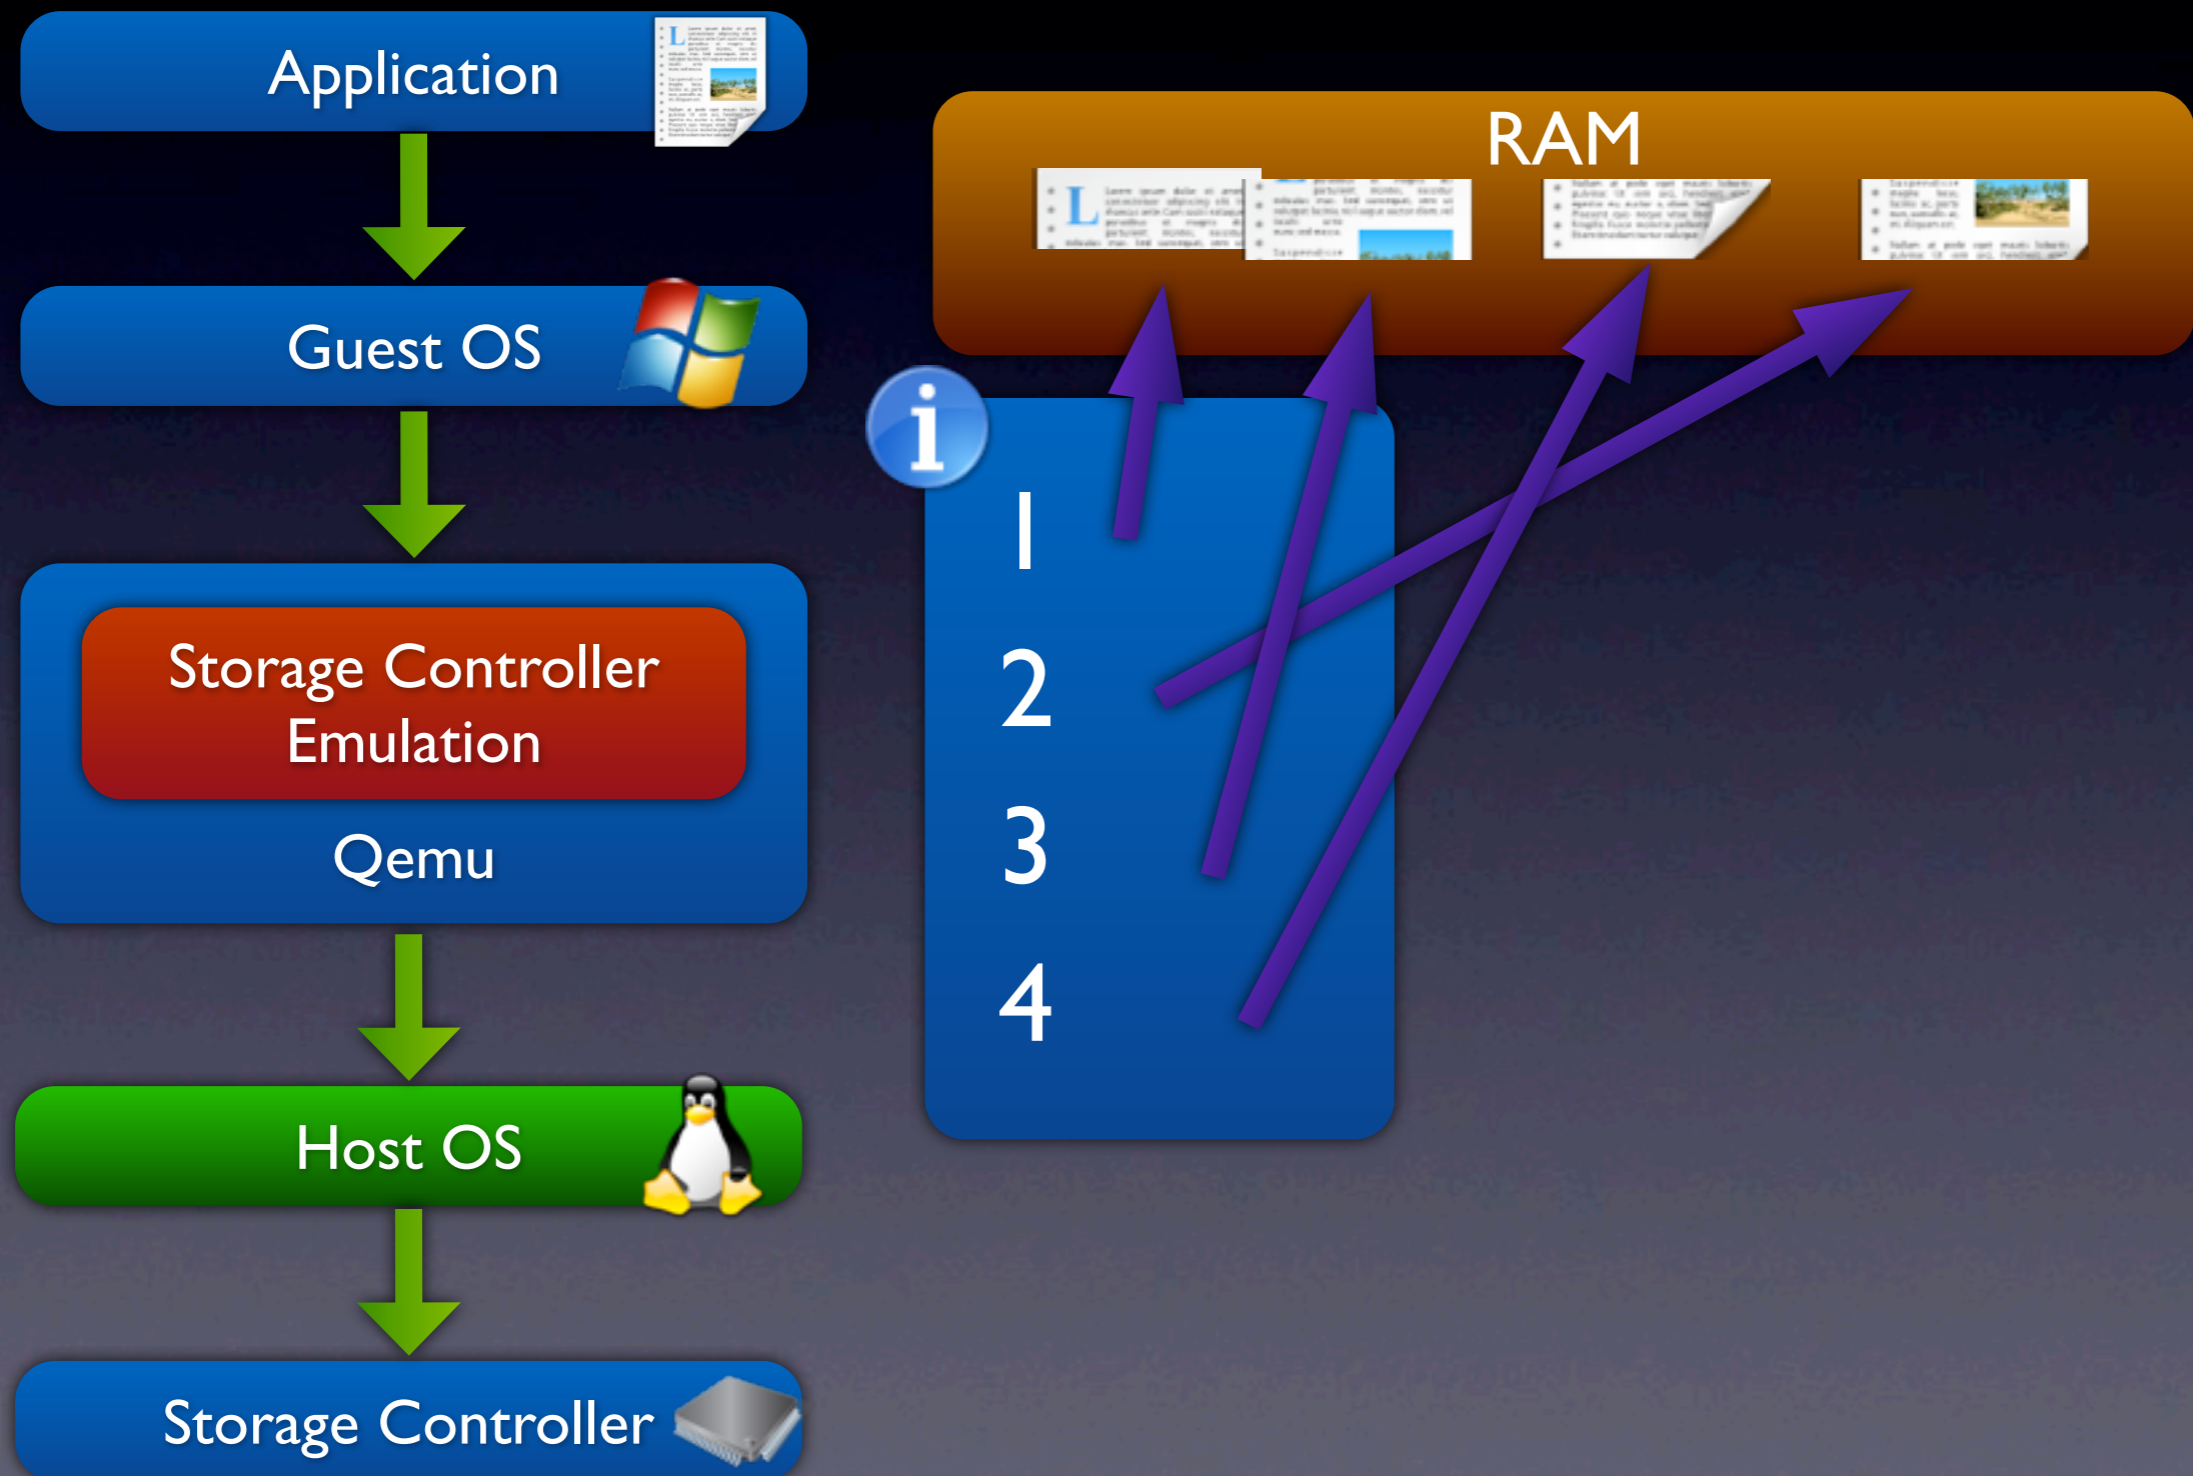
# Virtualized Storage



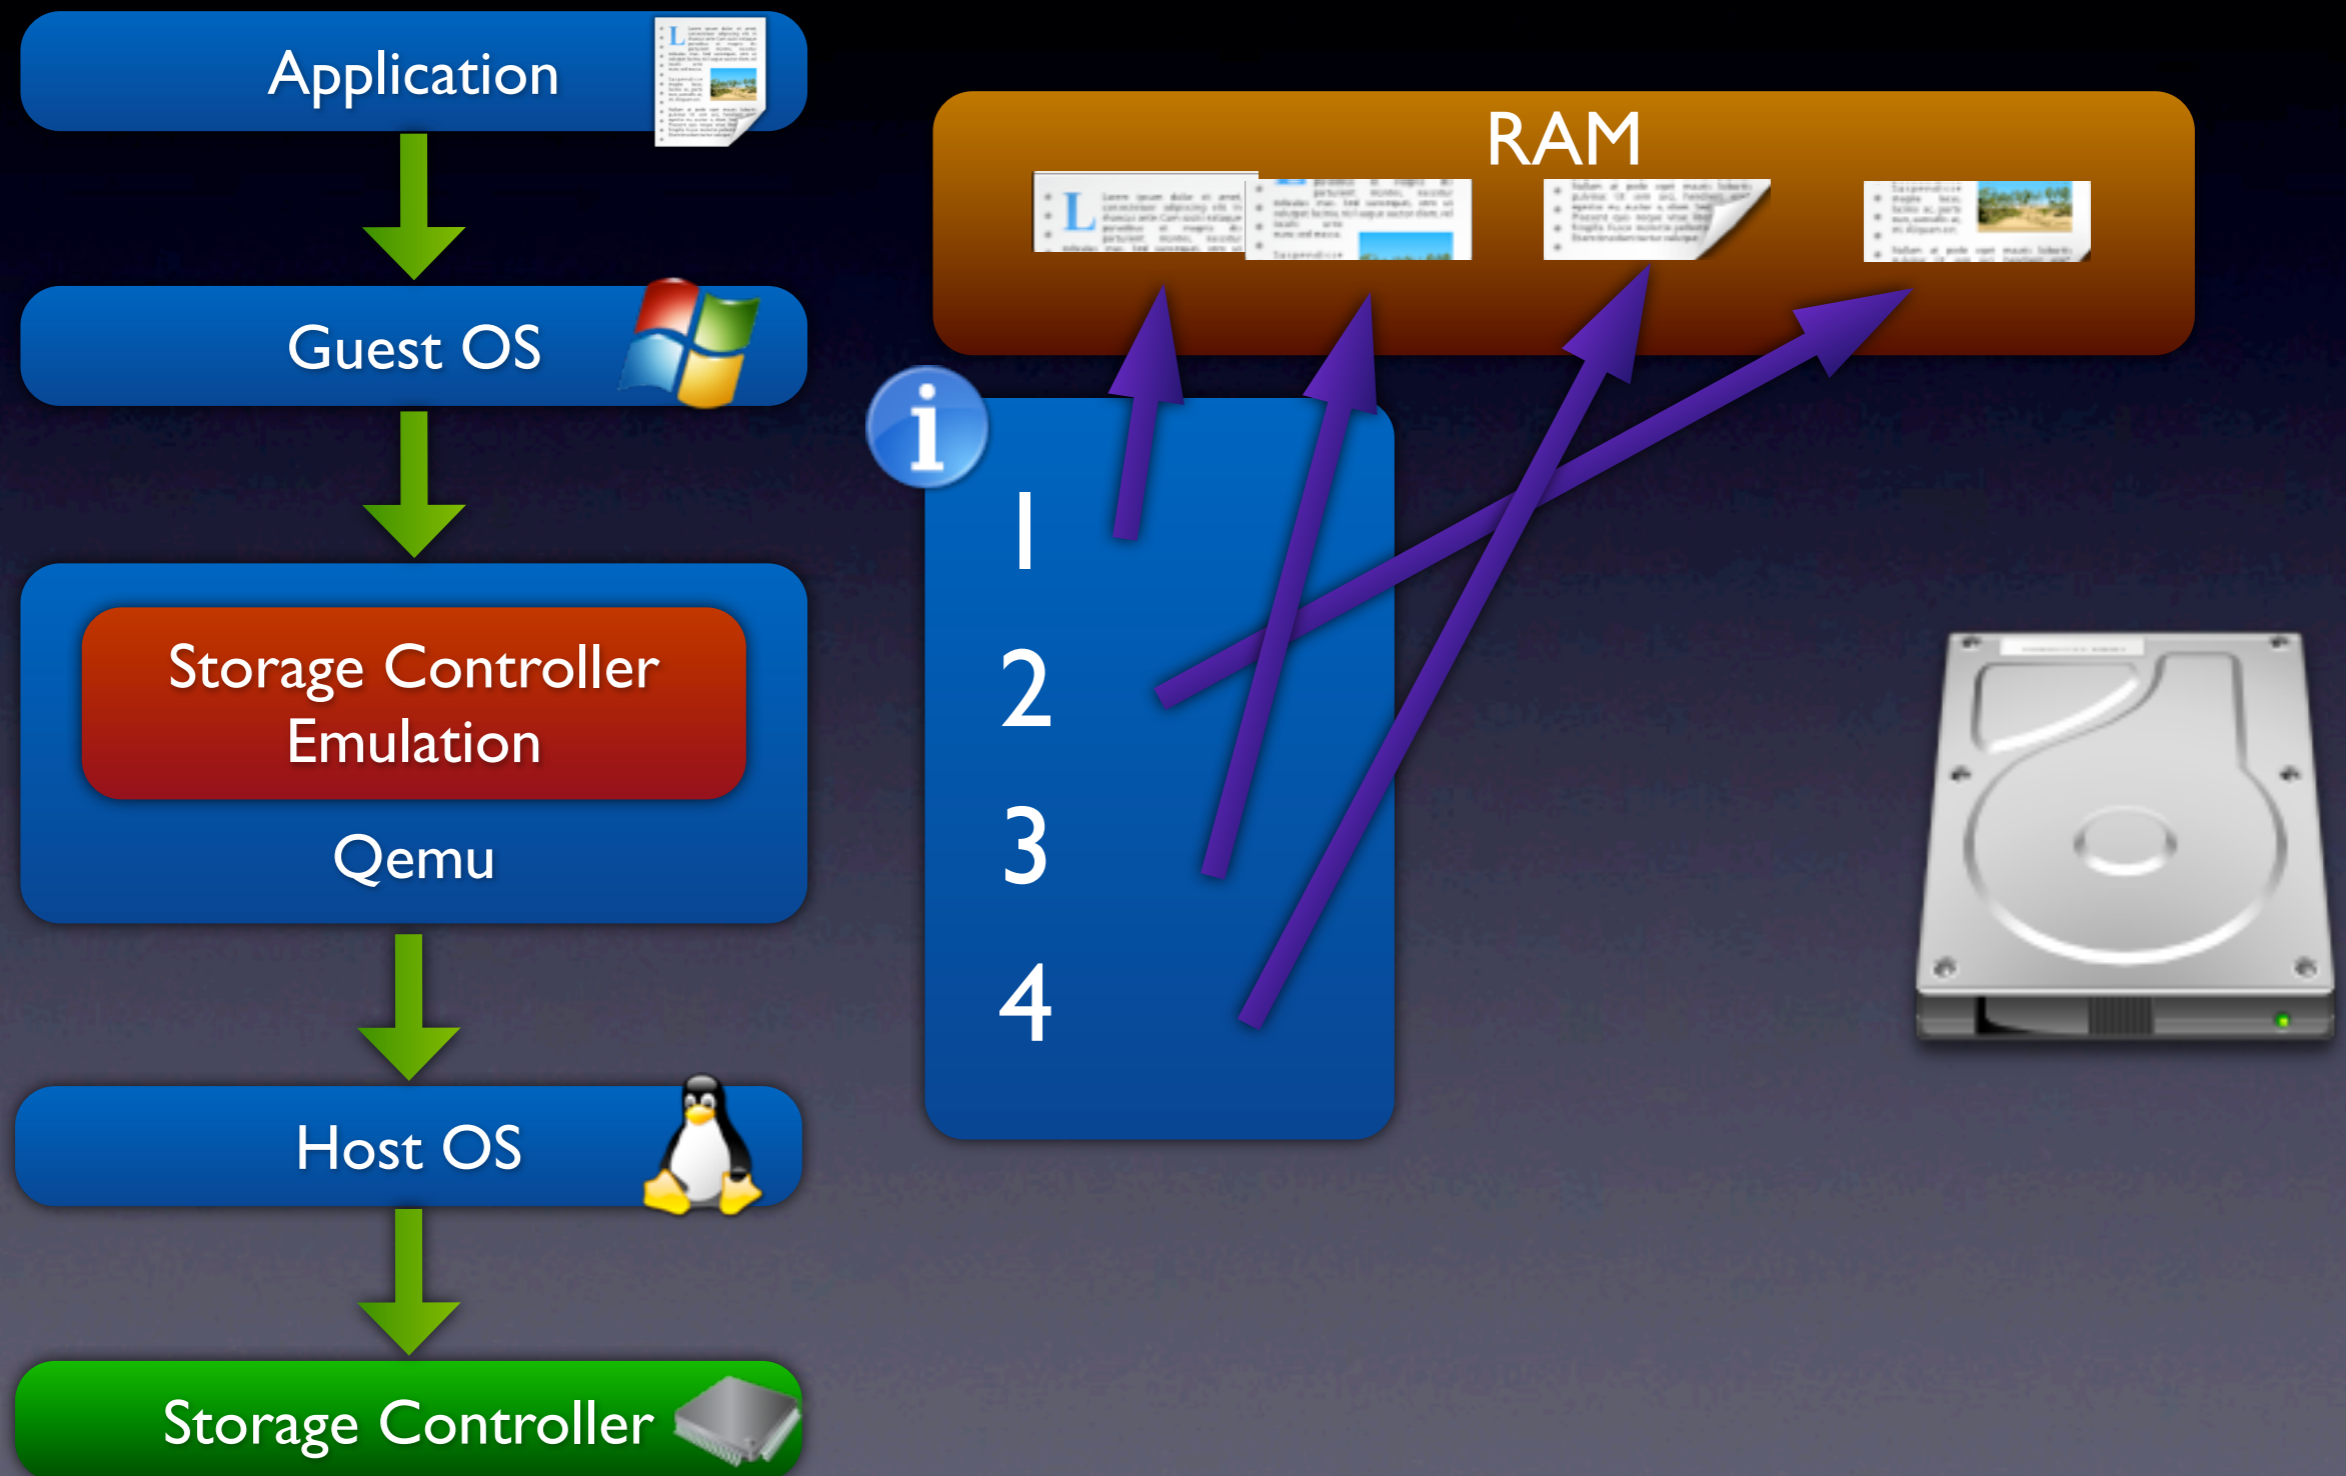
# Virtualized Storage



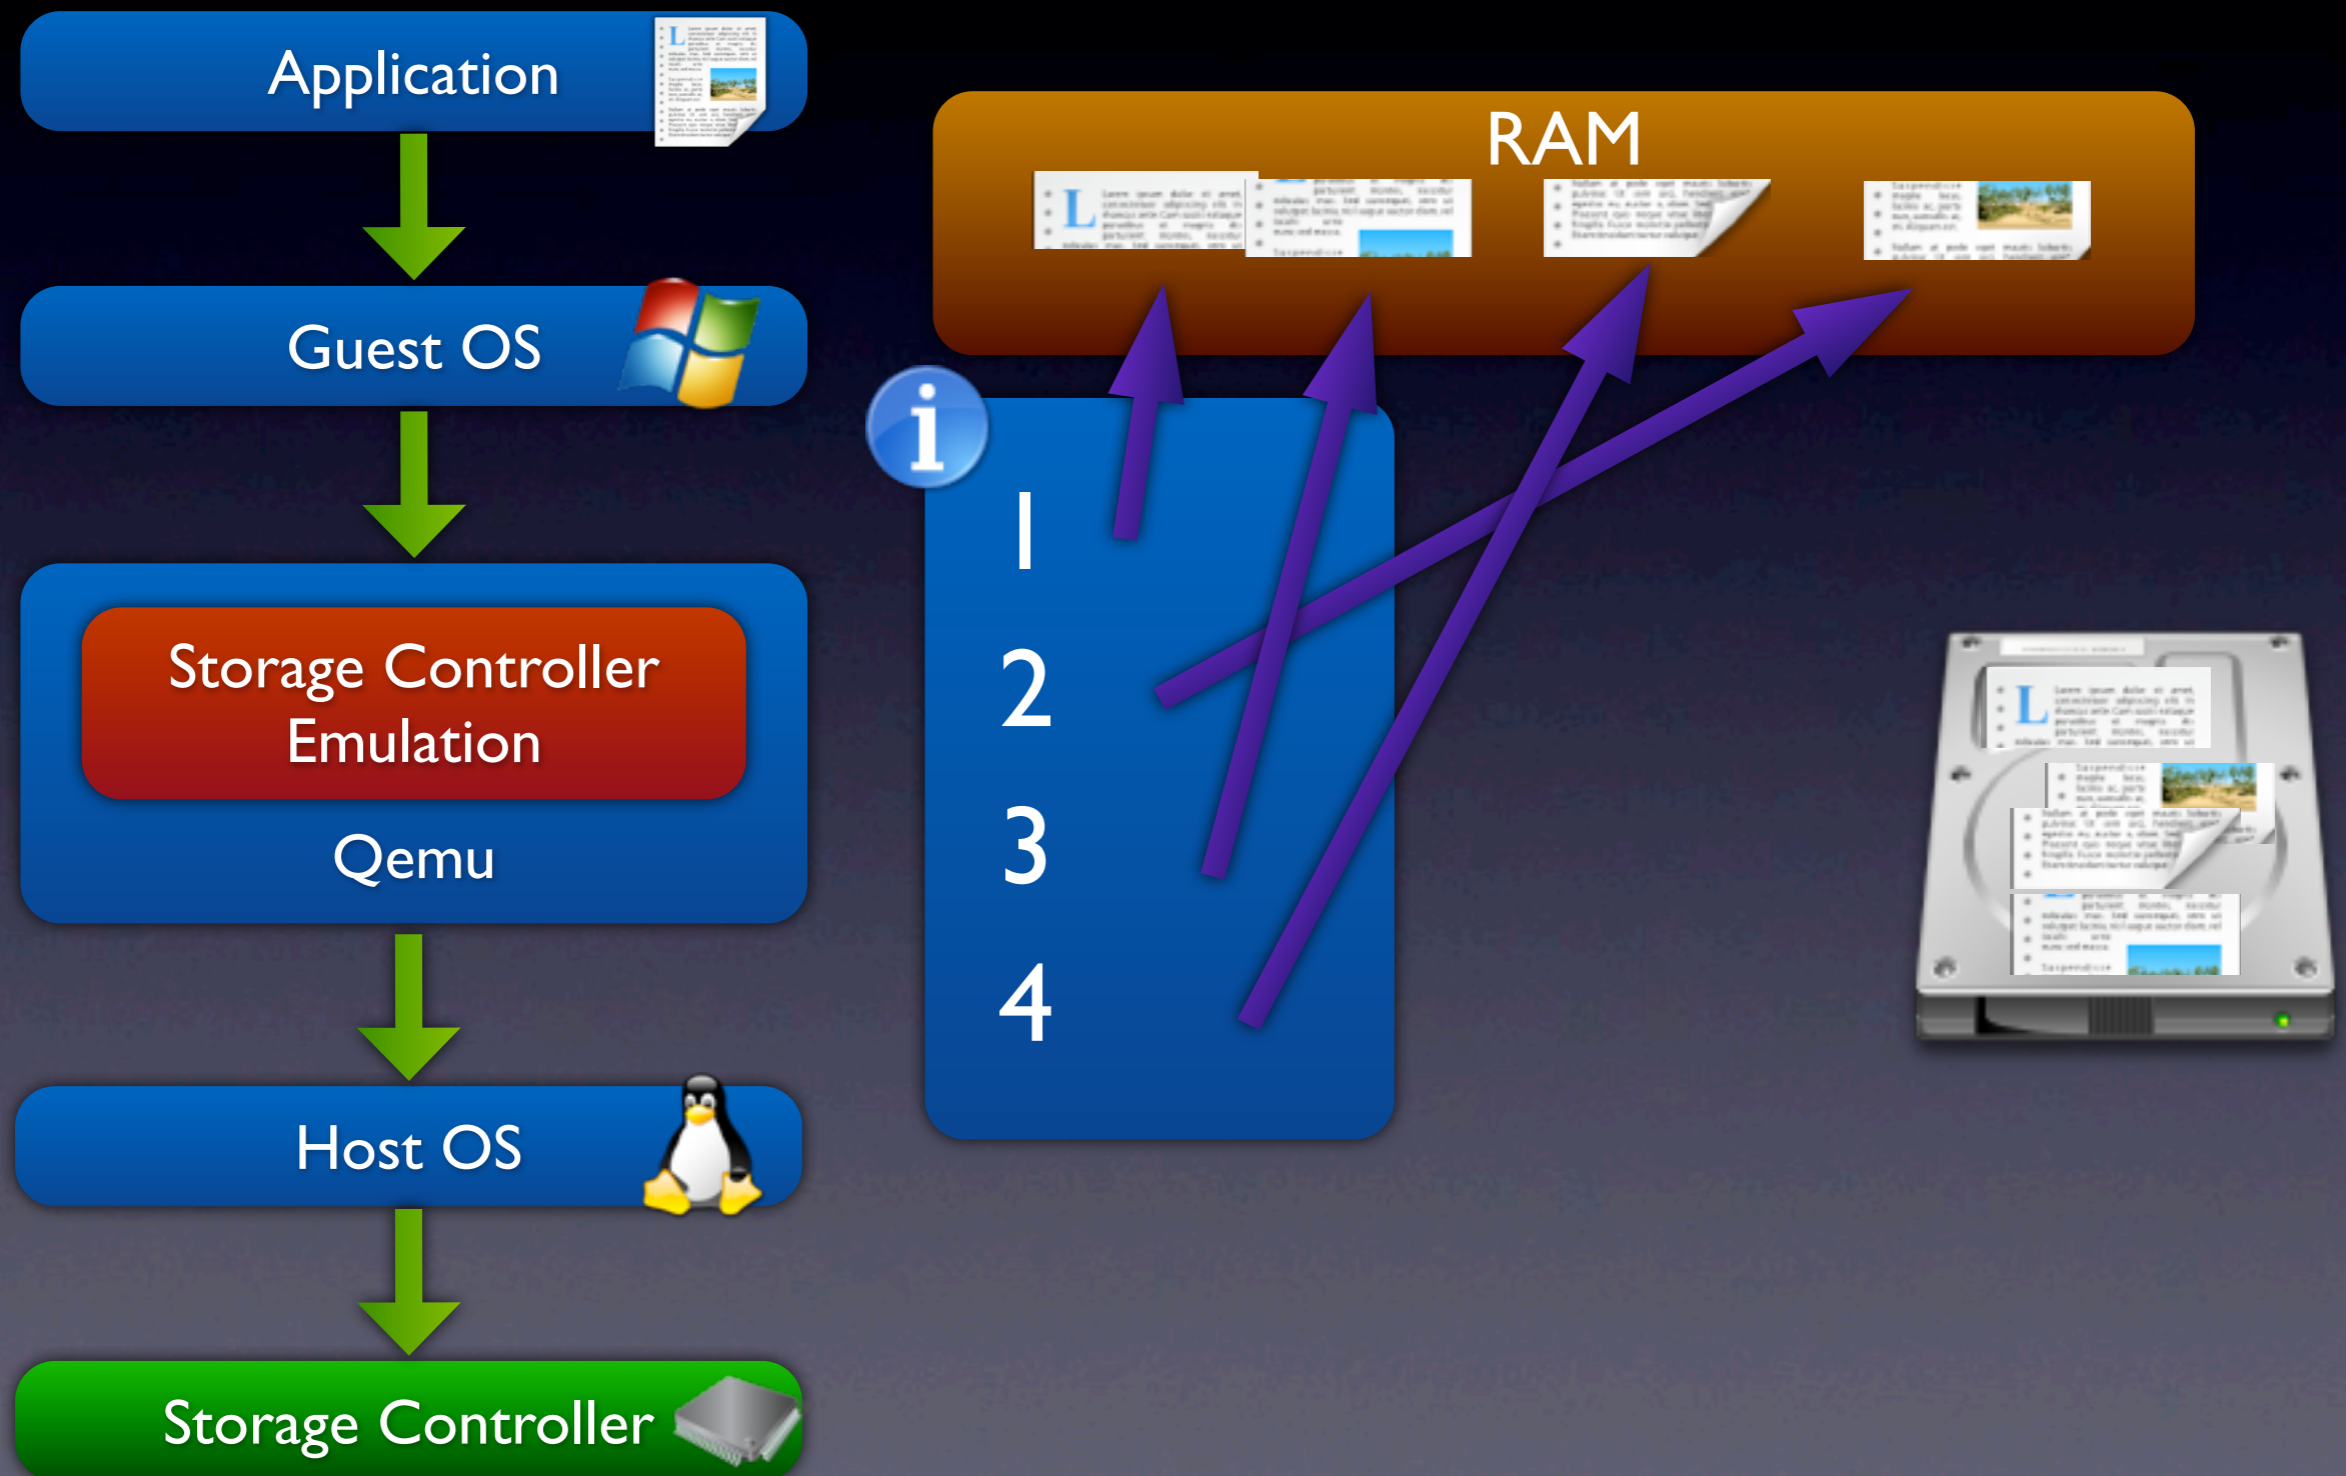
# Virtualized Storage



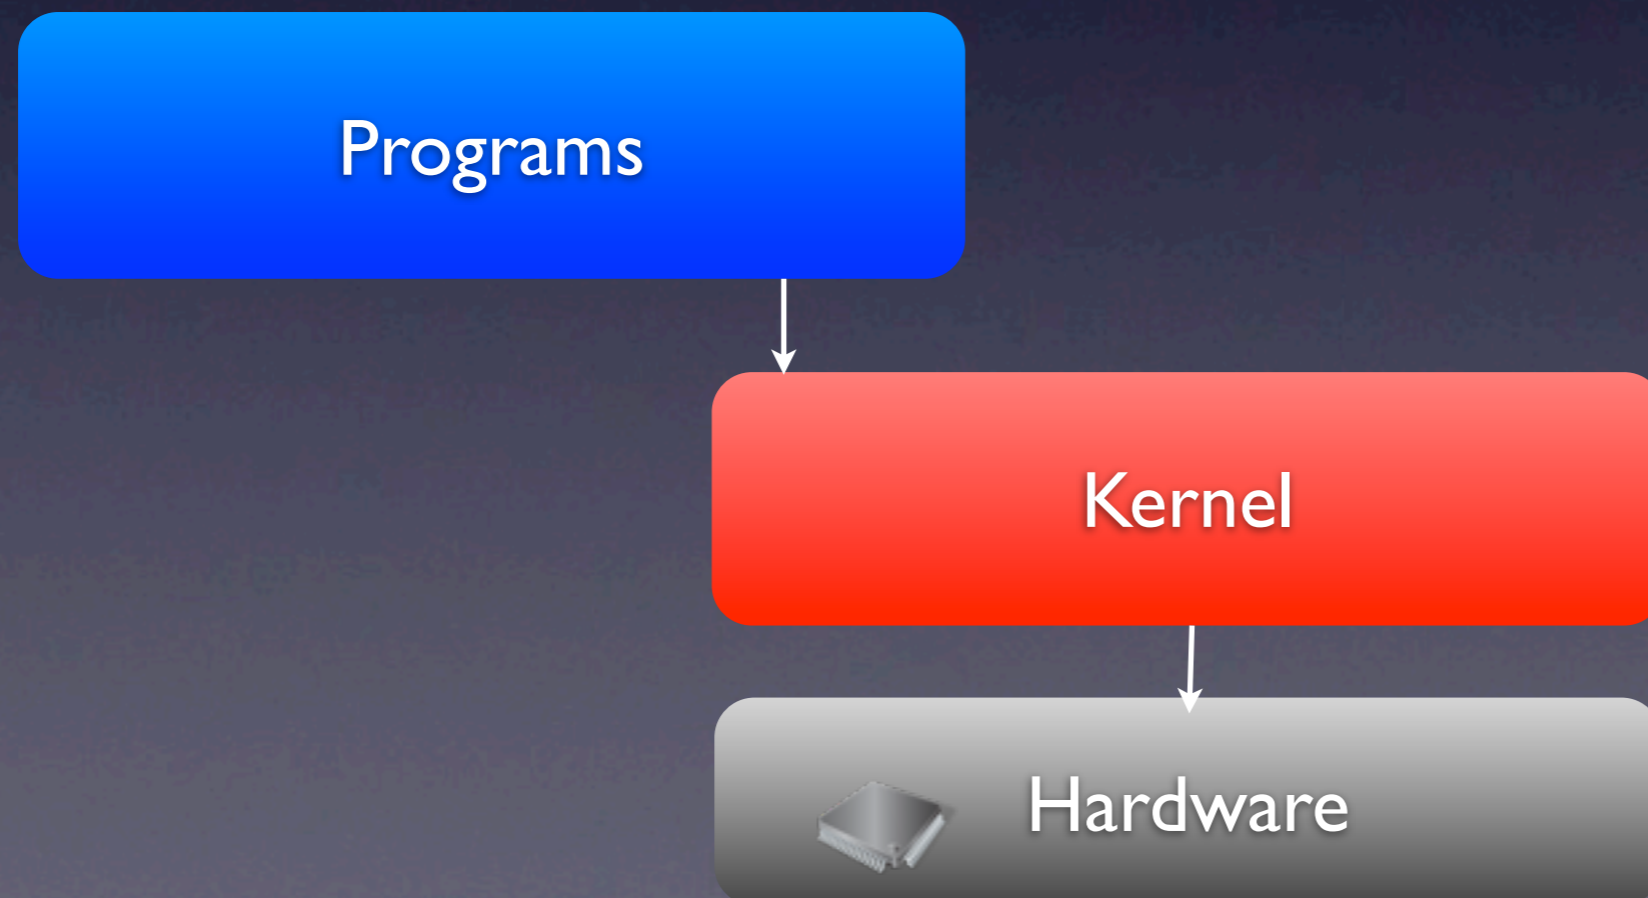
# Virtualized Storage




# Virtualized Storage



# Device Access



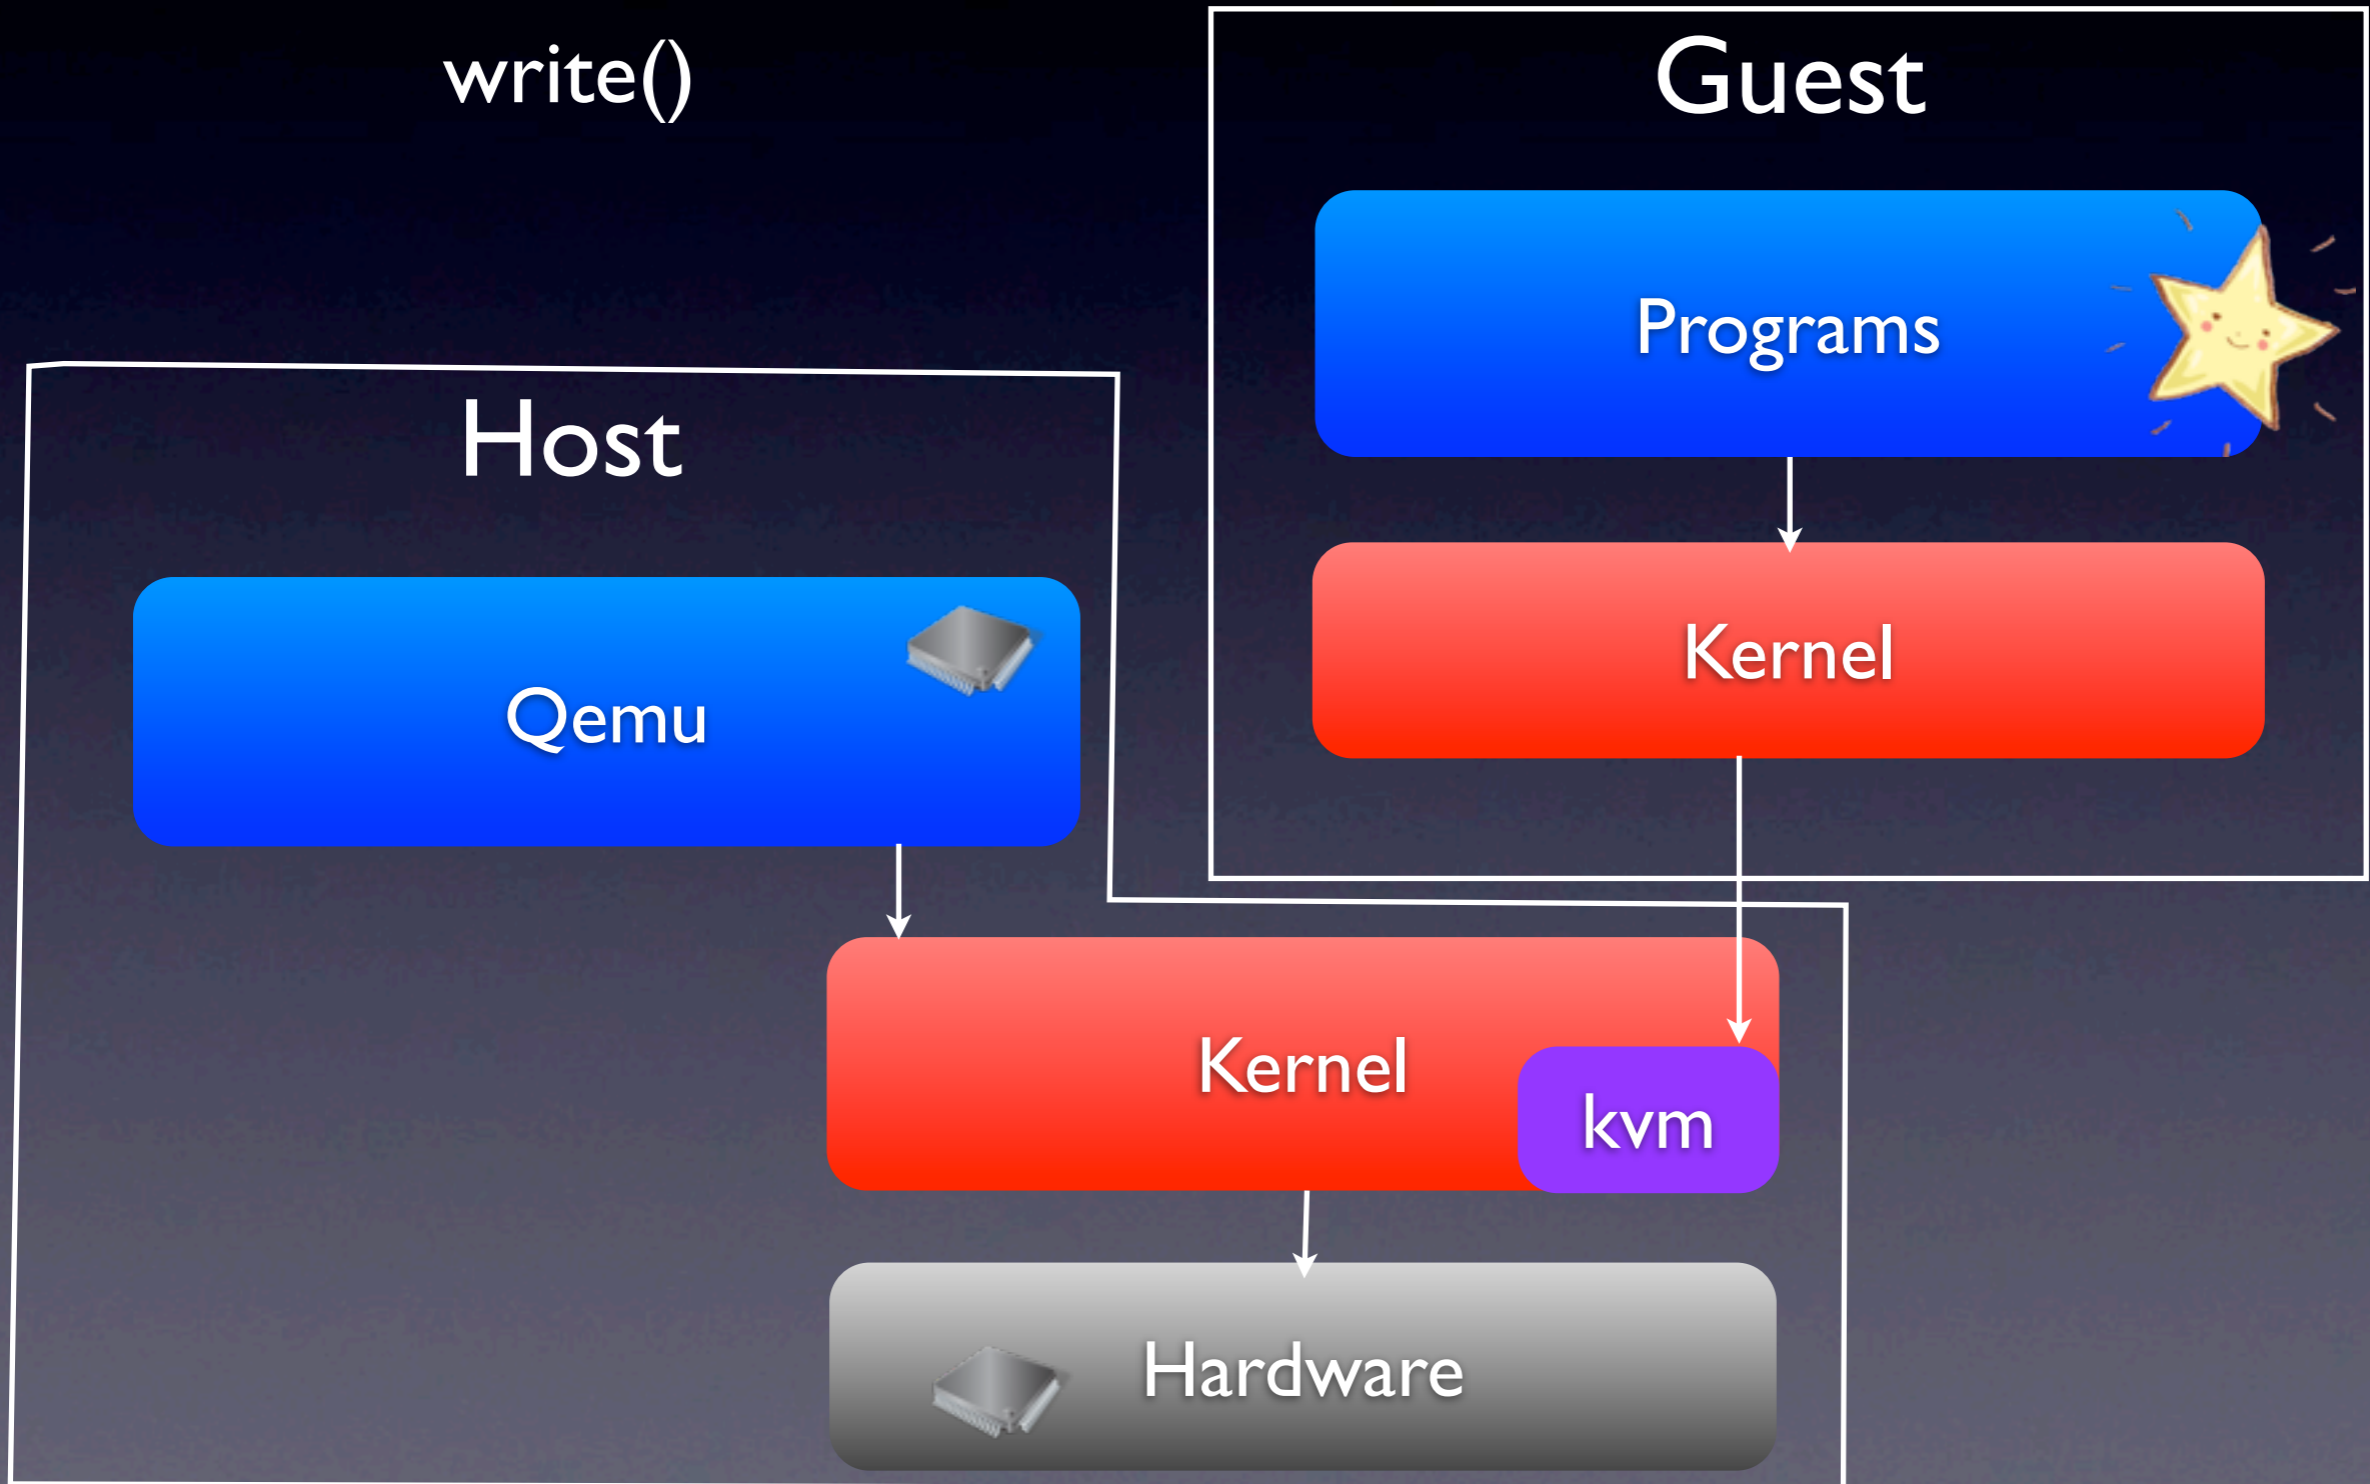
# Device Access





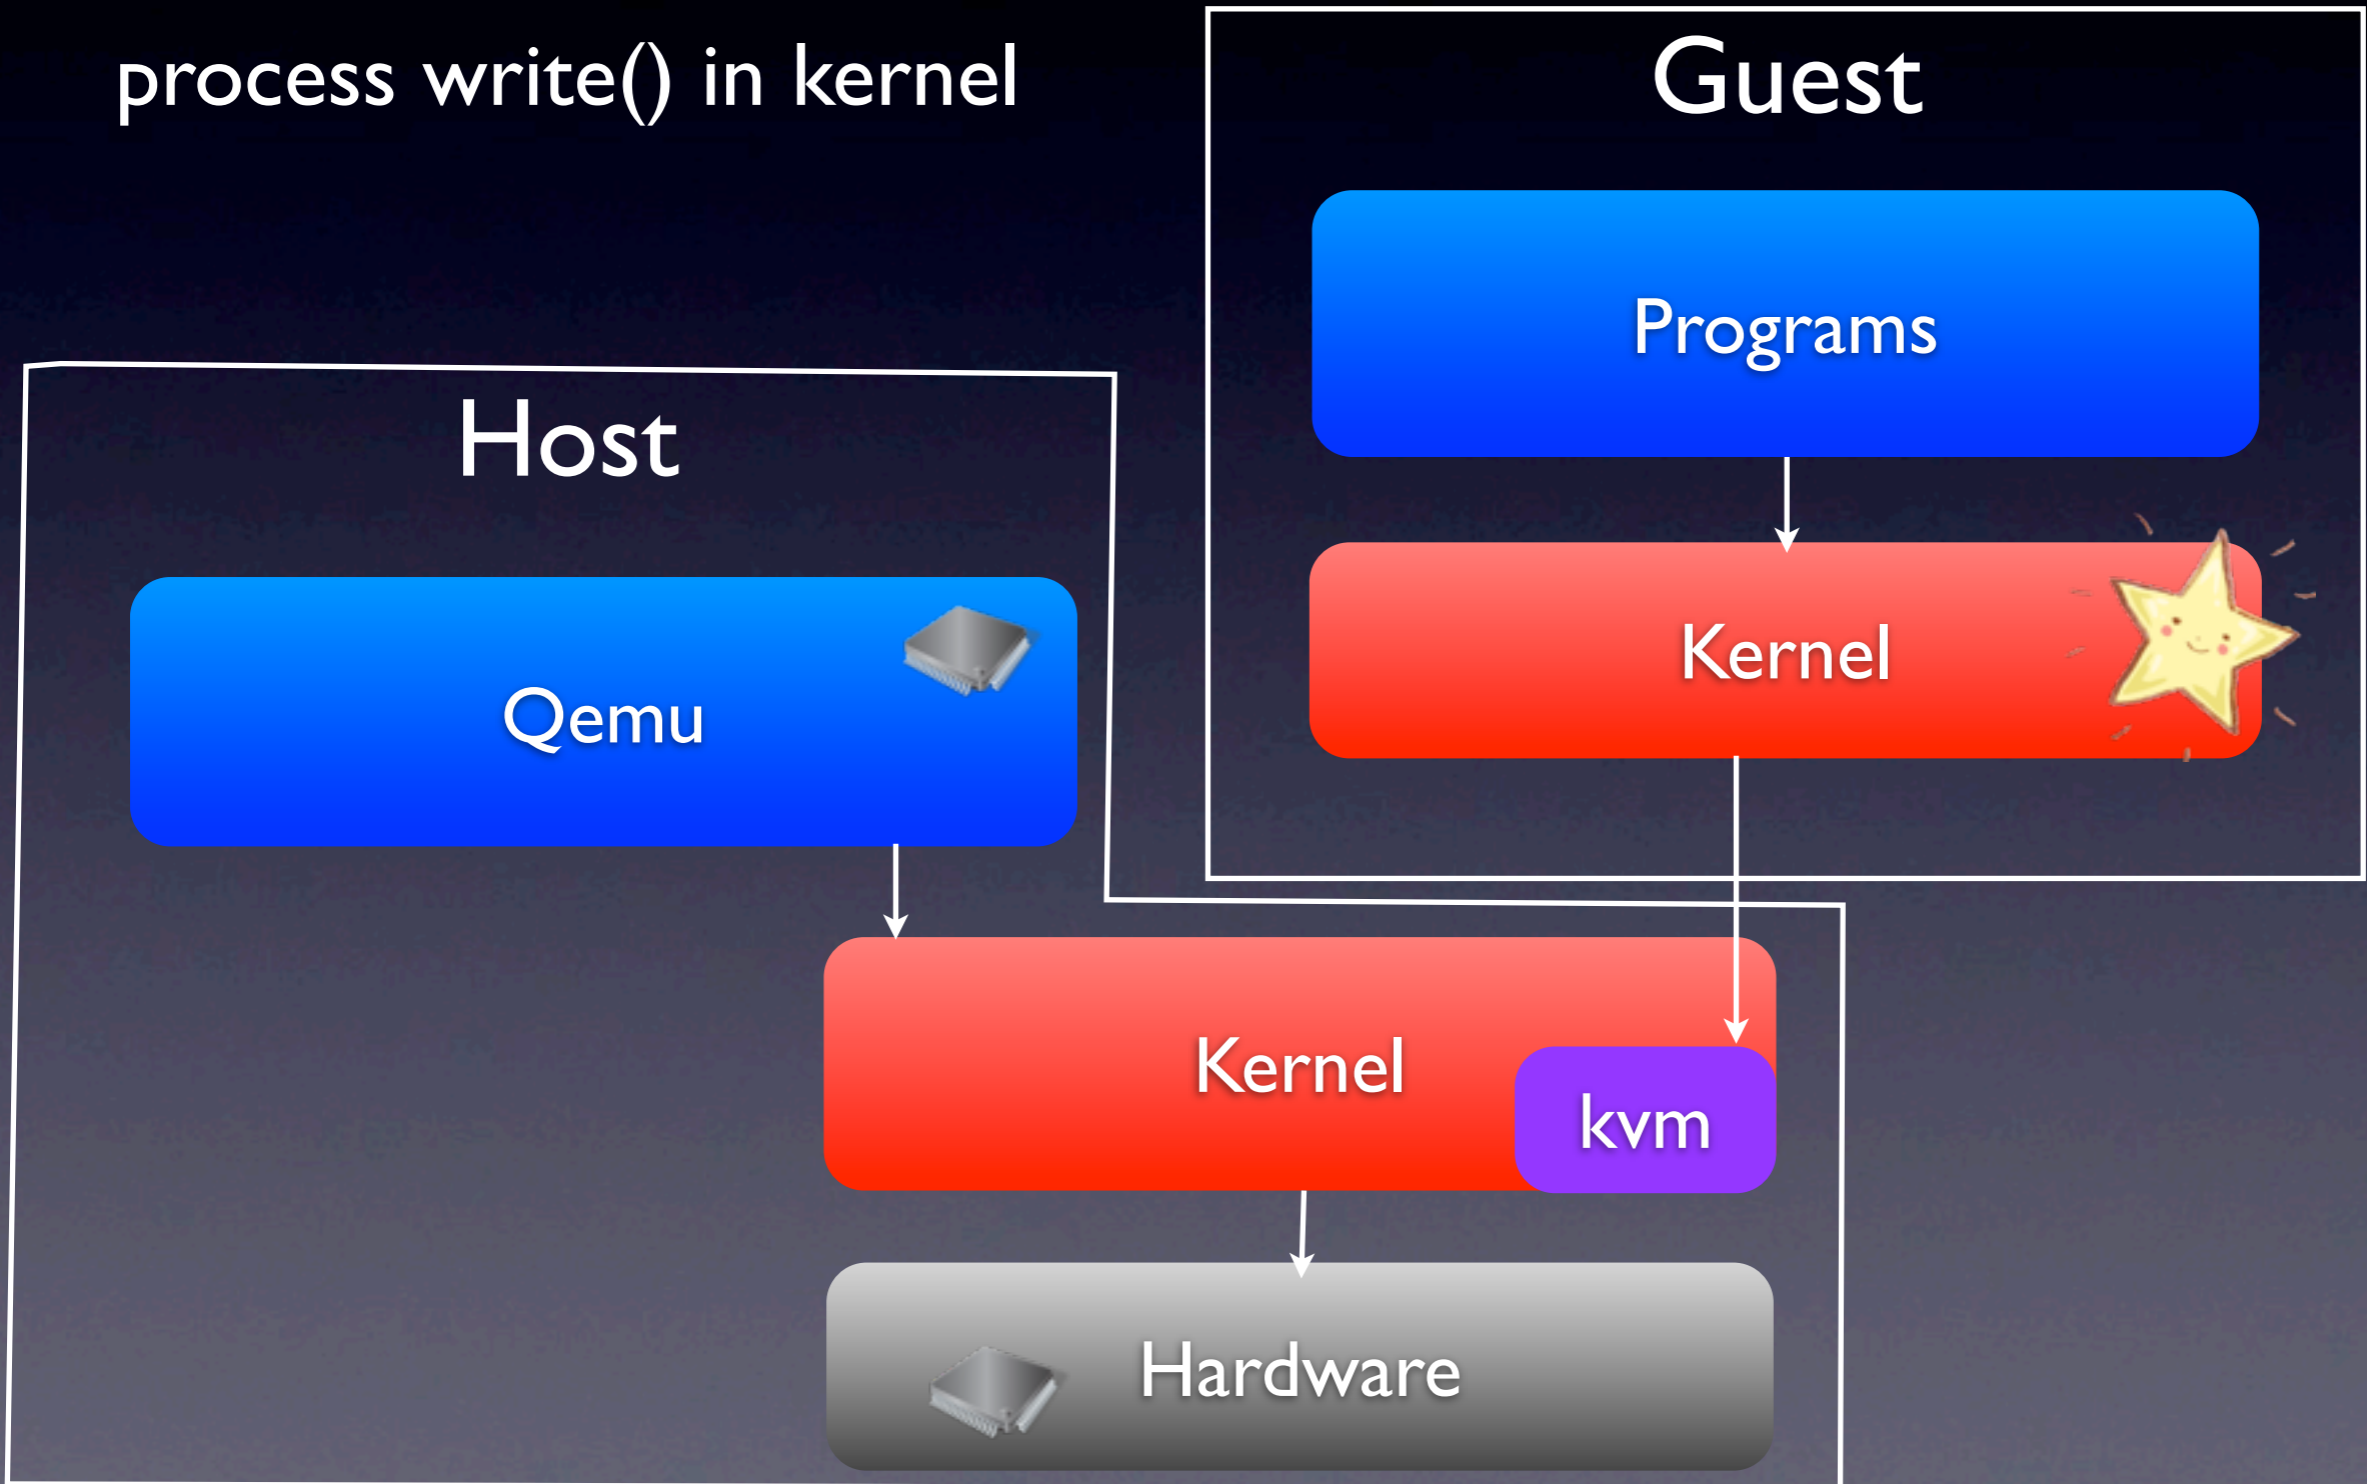
# Device Access

write()



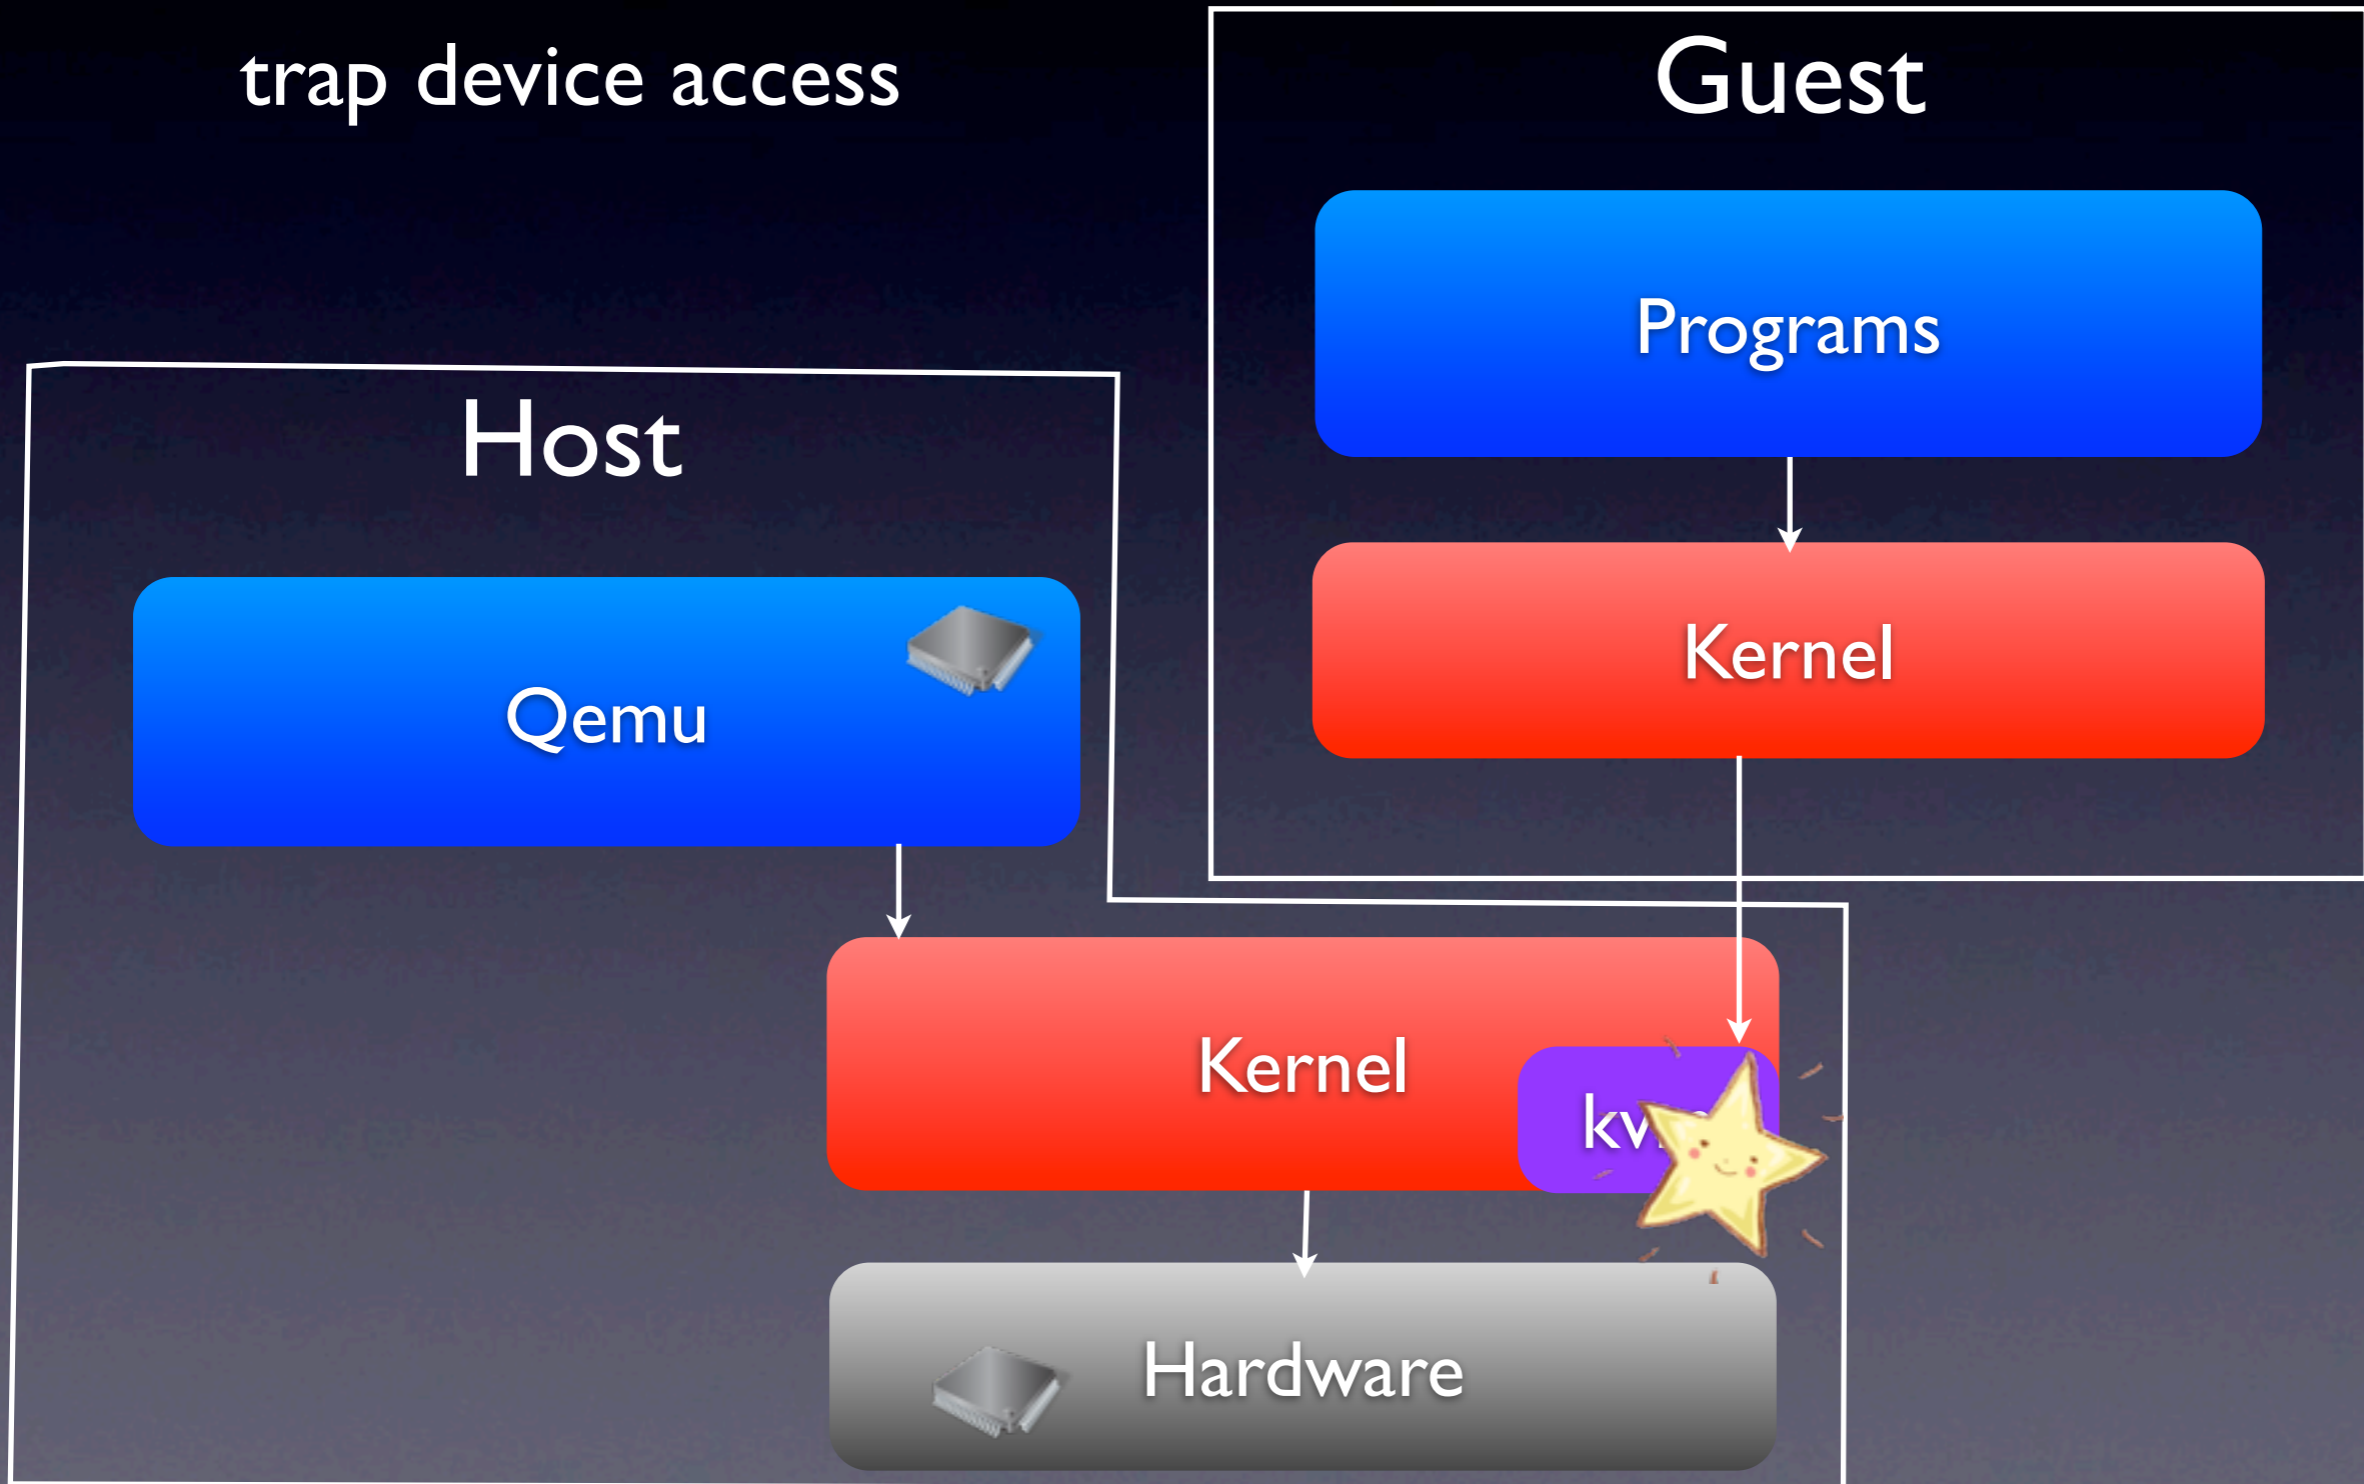
# Device Access

process write() in kernel



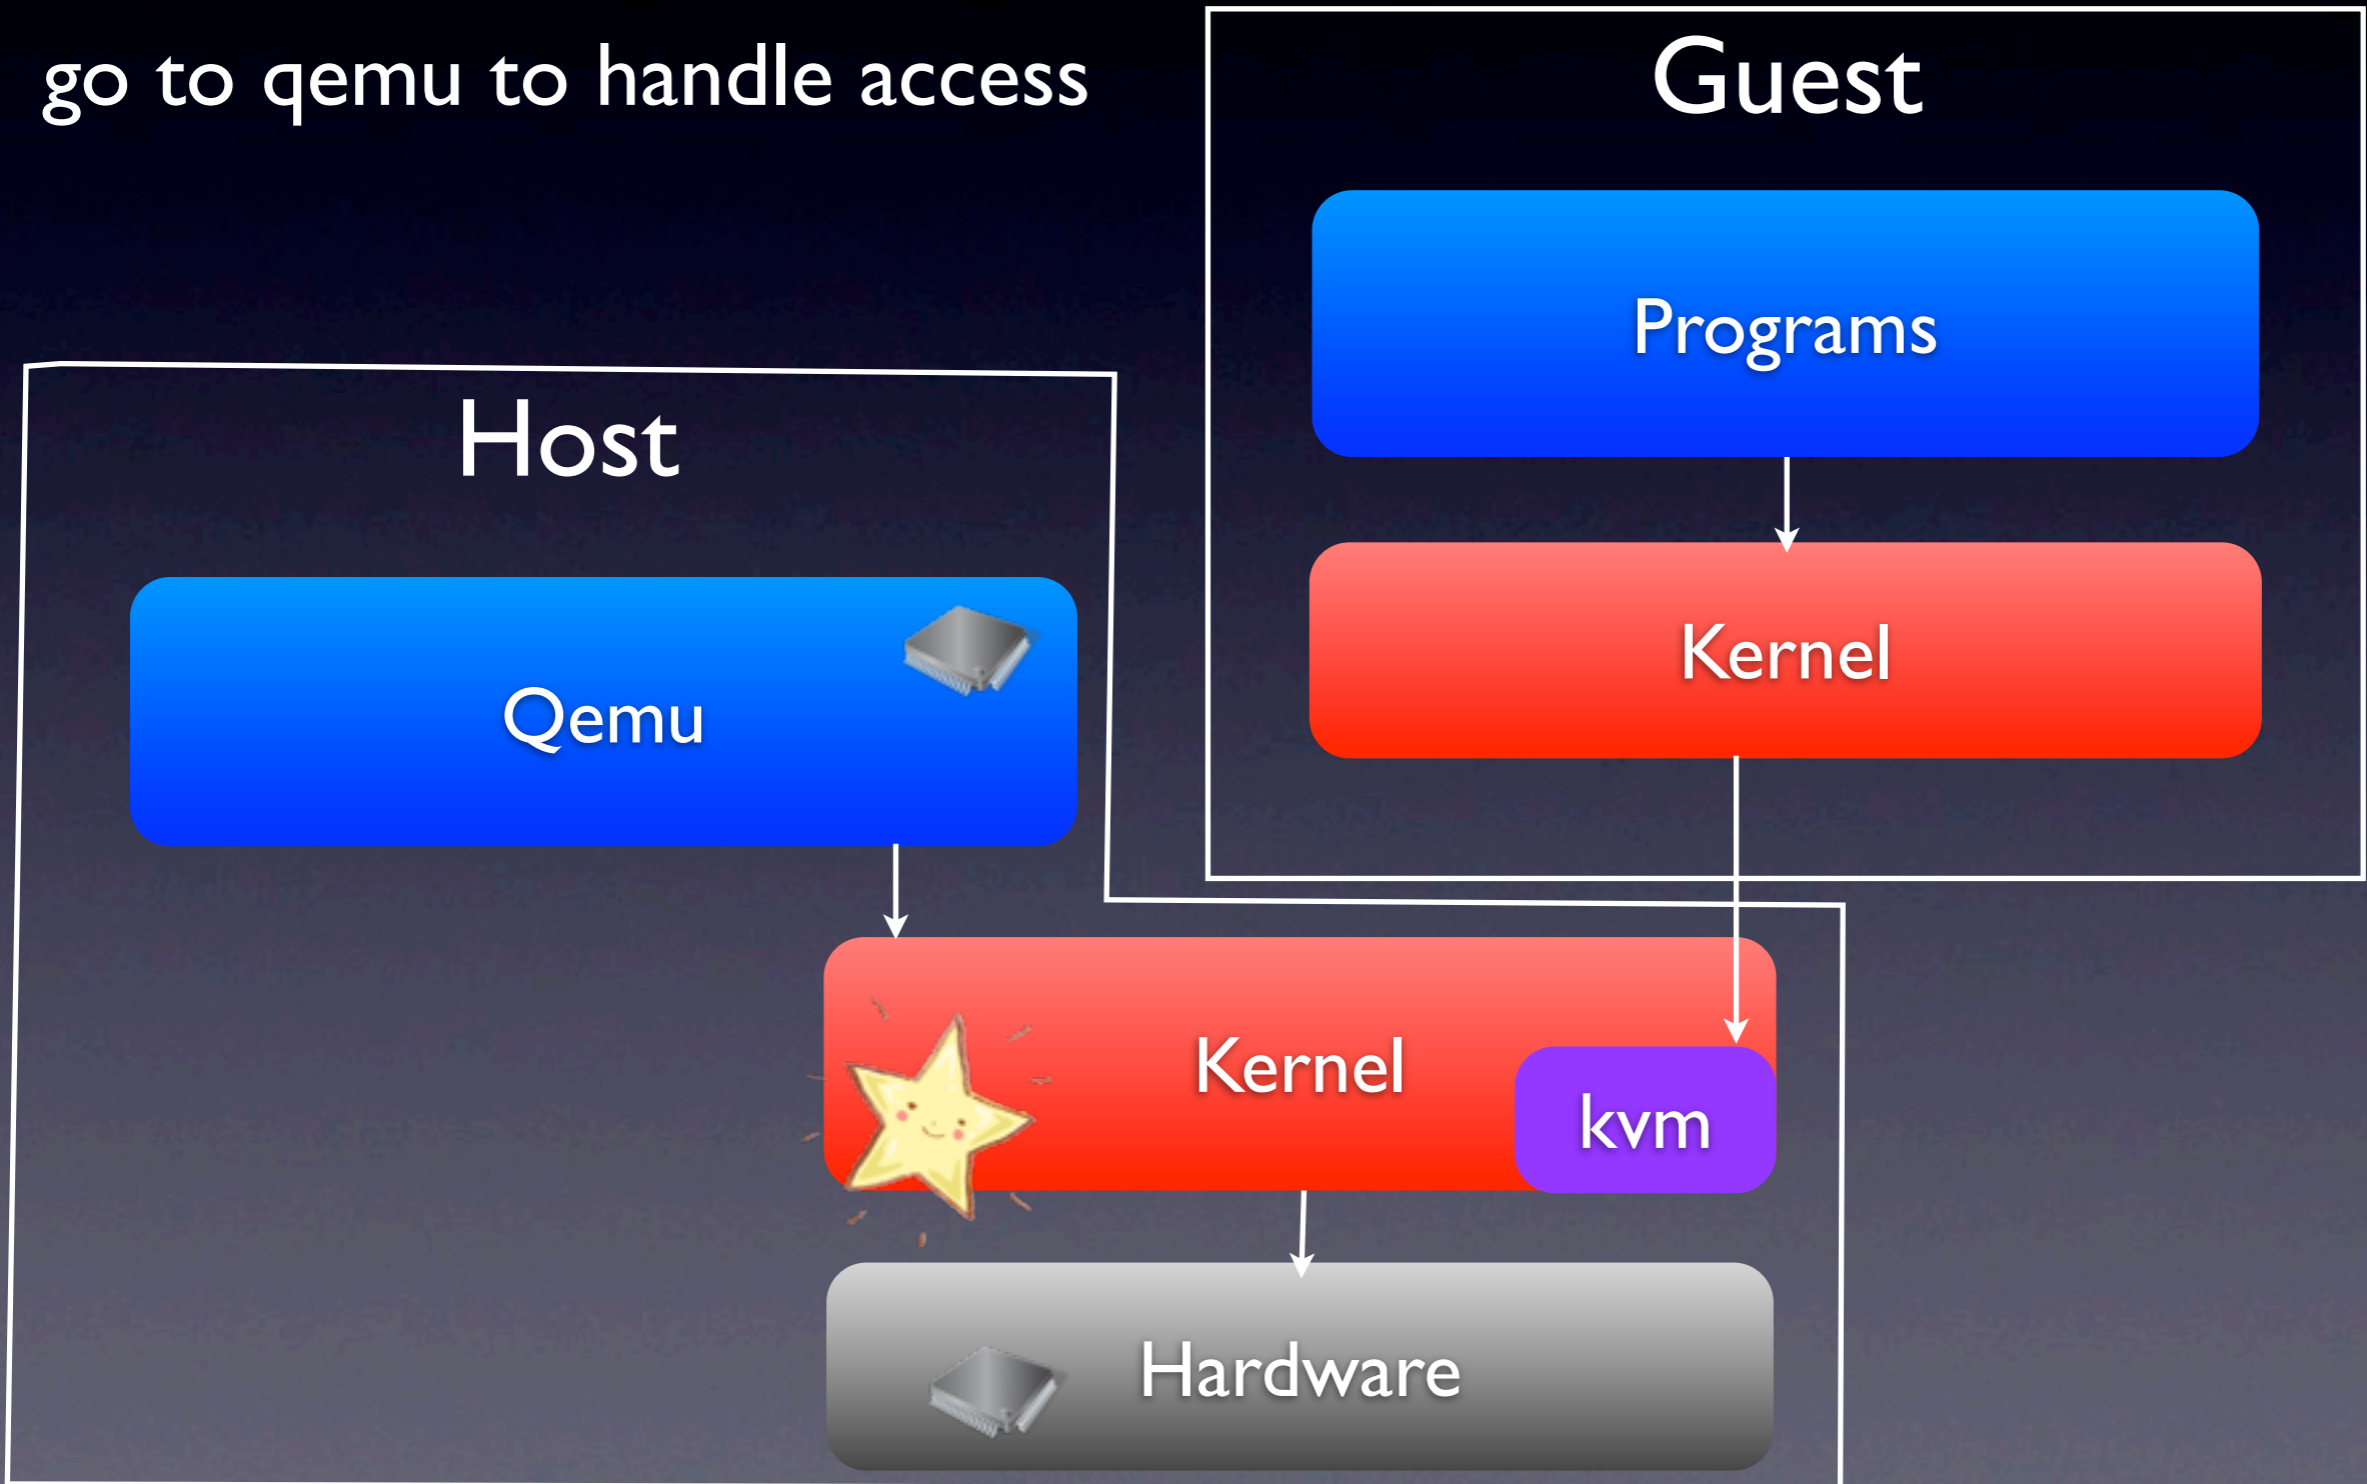
# Device Access

trap device access



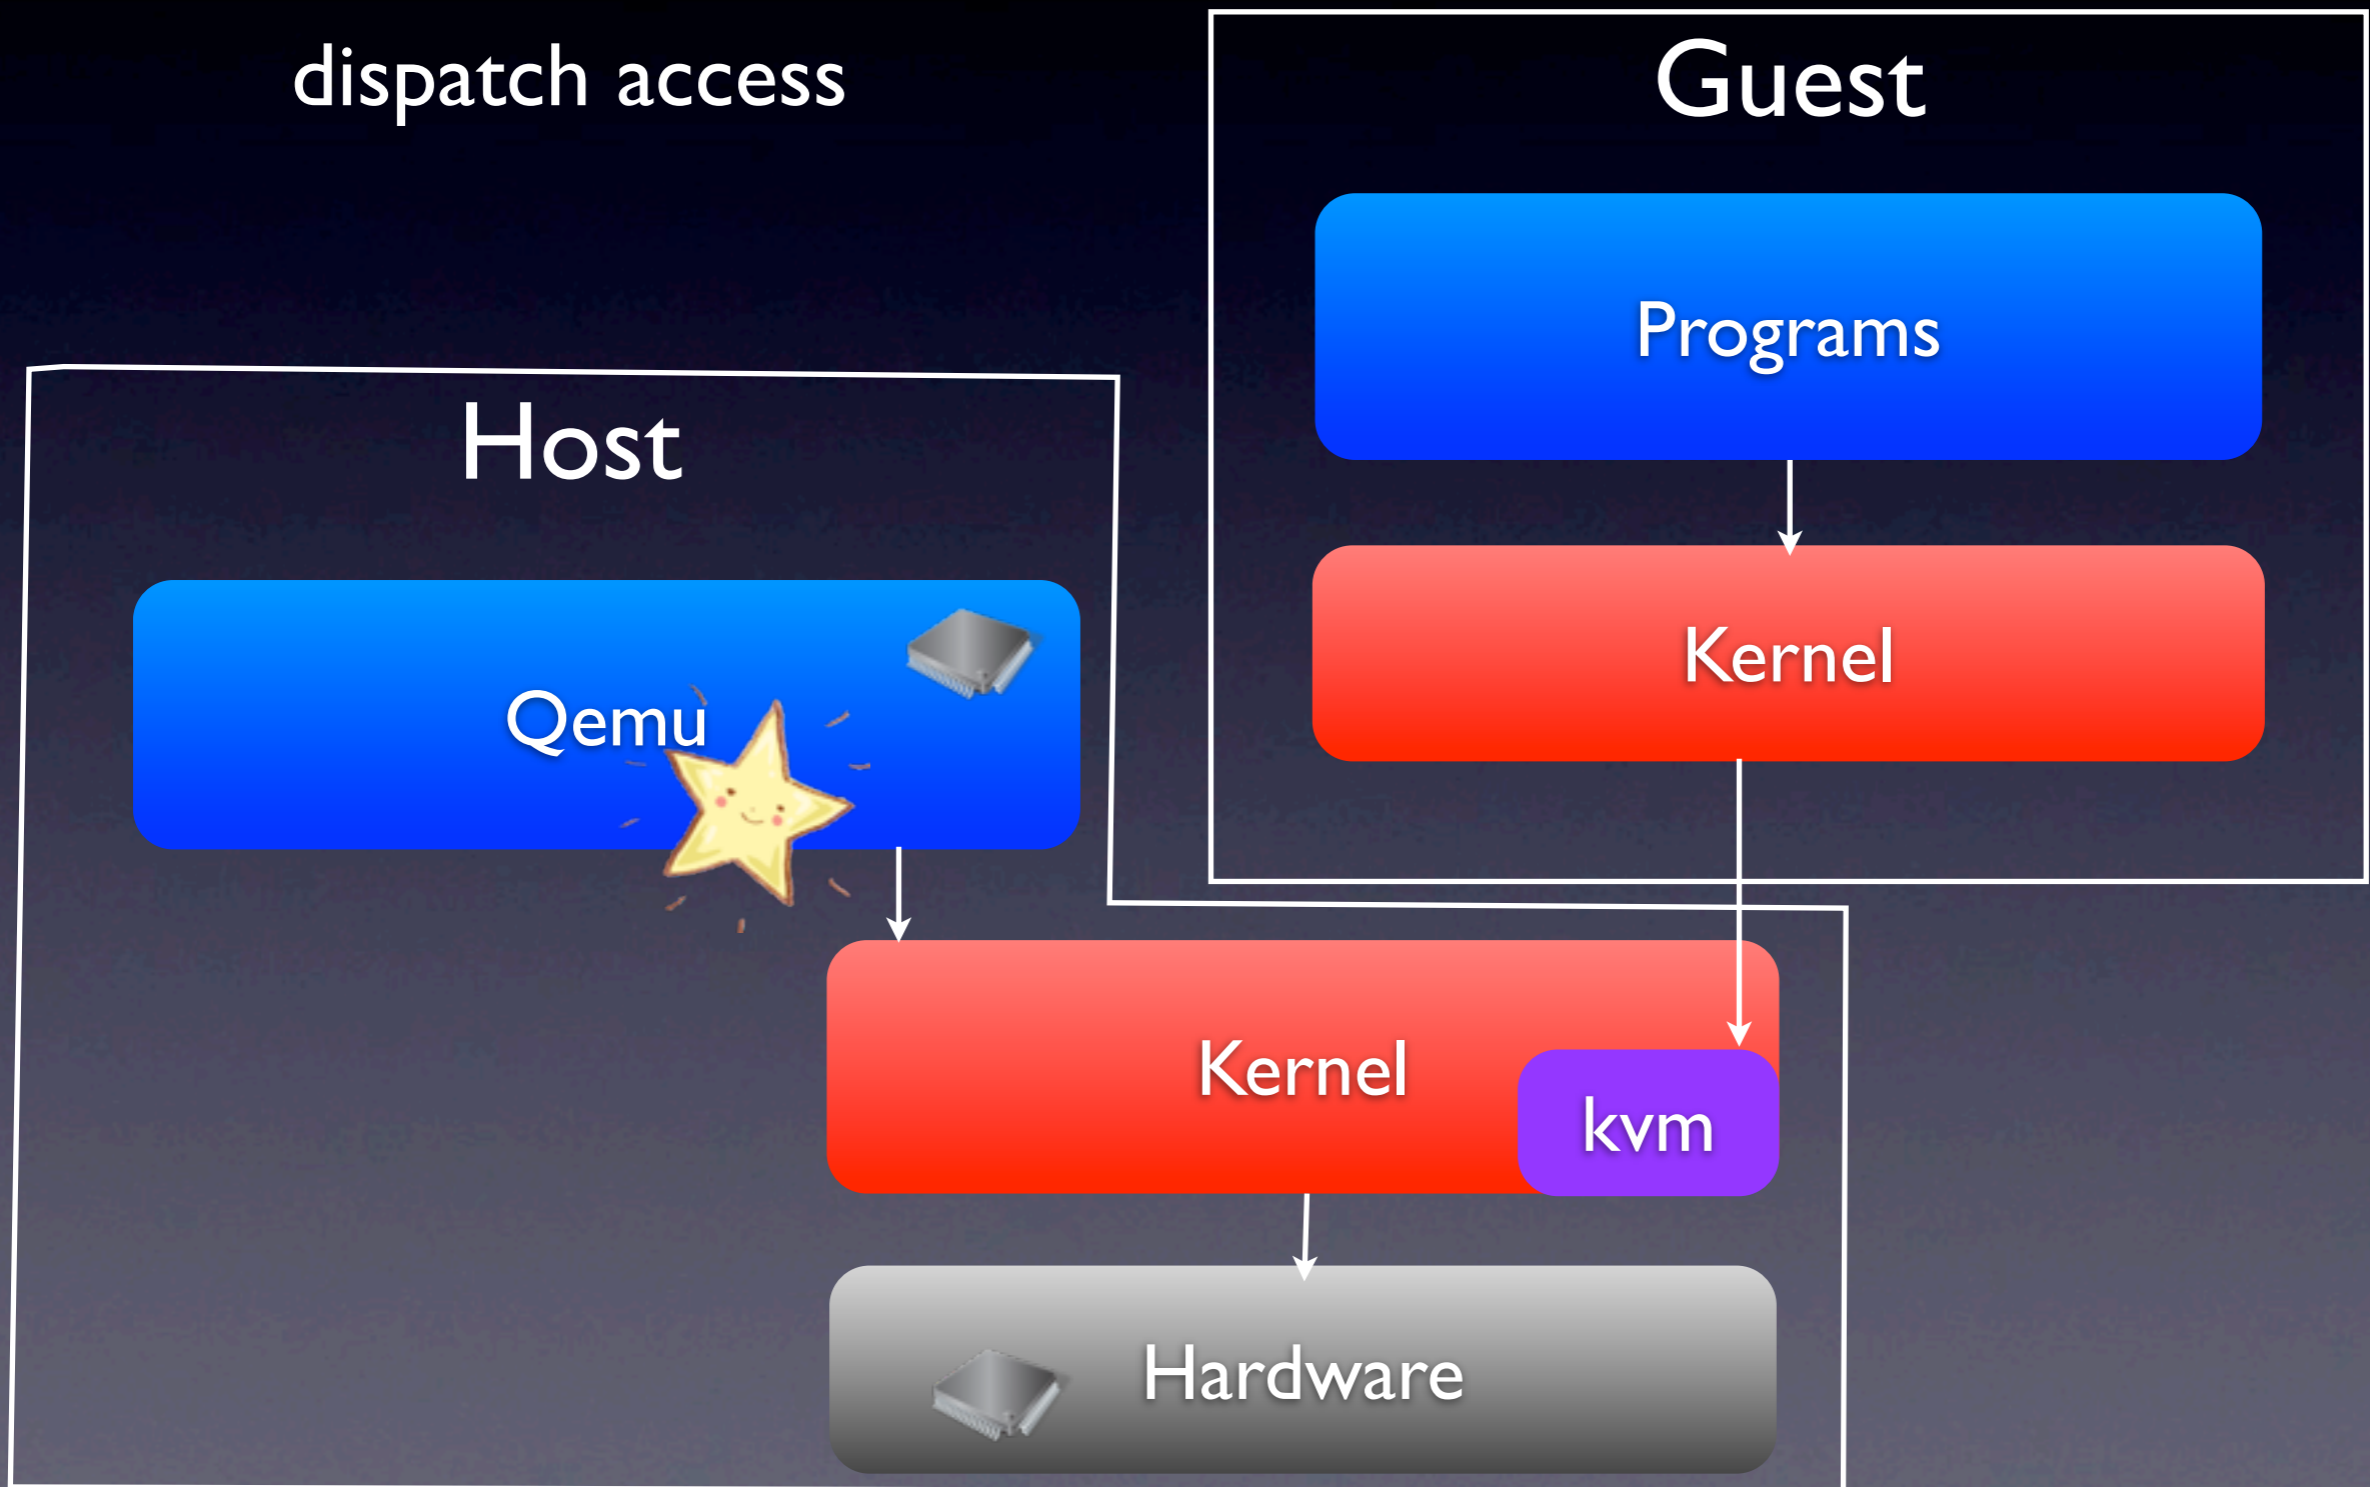
# Device Access

go to qemu to handle access



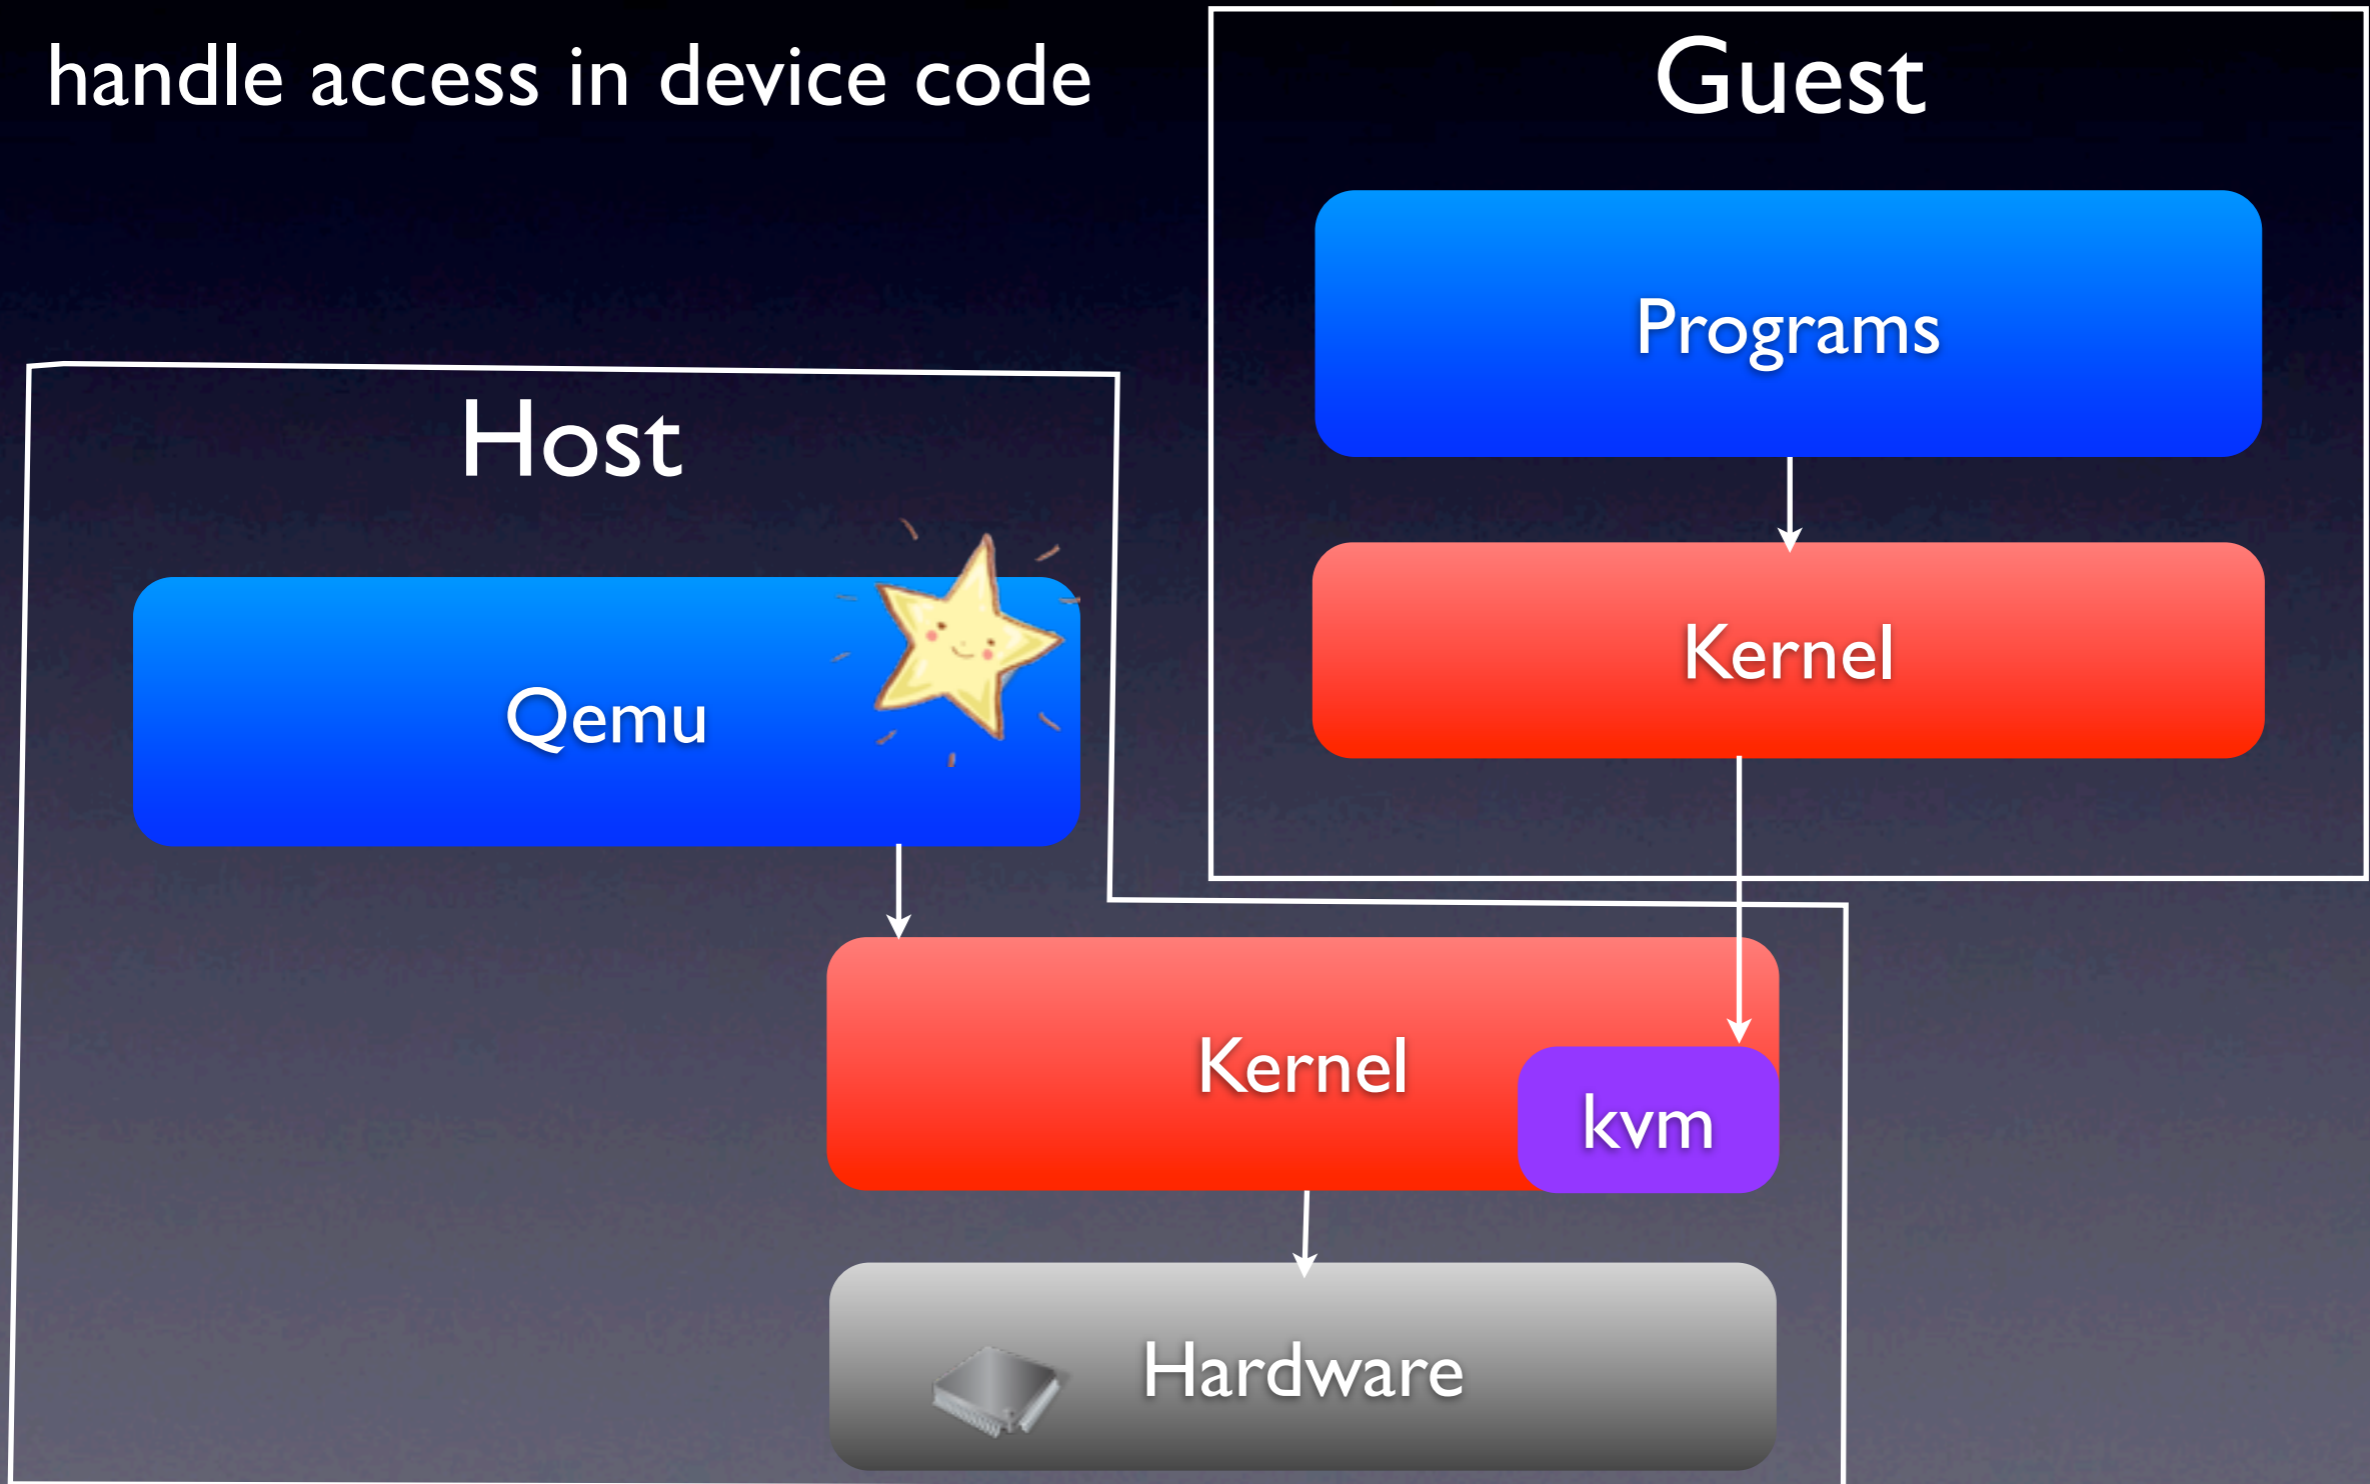
# Device Access

dispatch access




# Device Access

handle access in device code



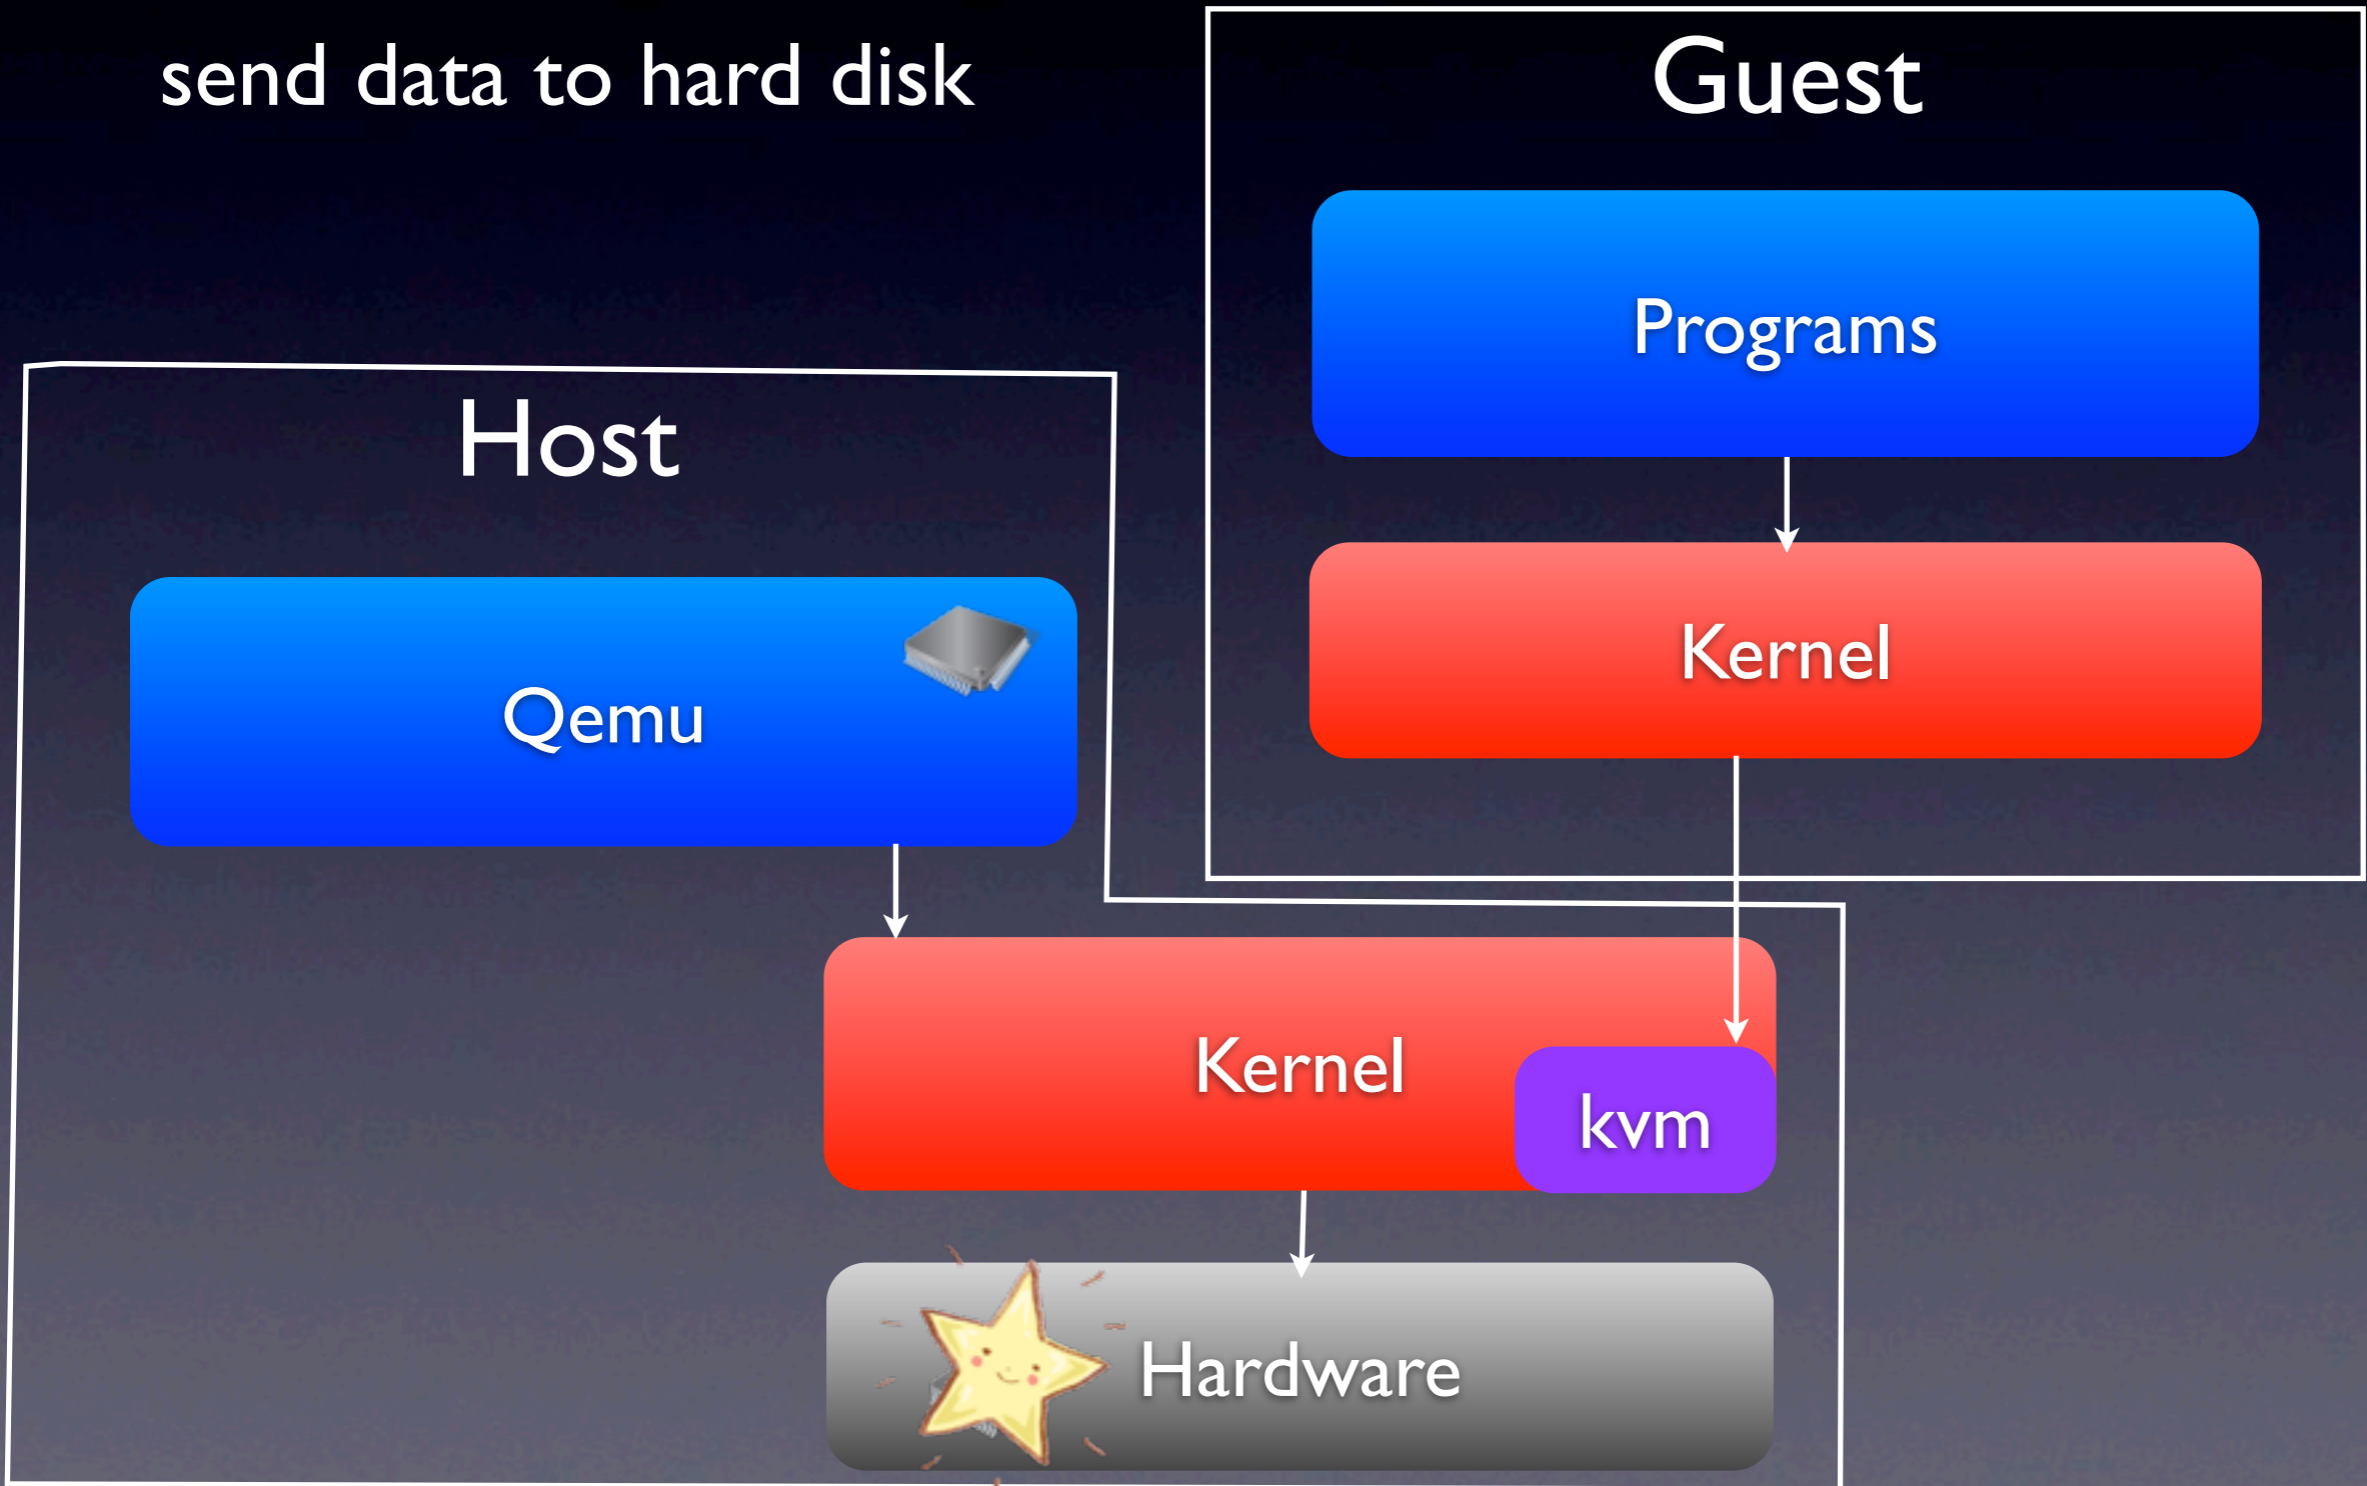
# Device Access

write()



# Device Access

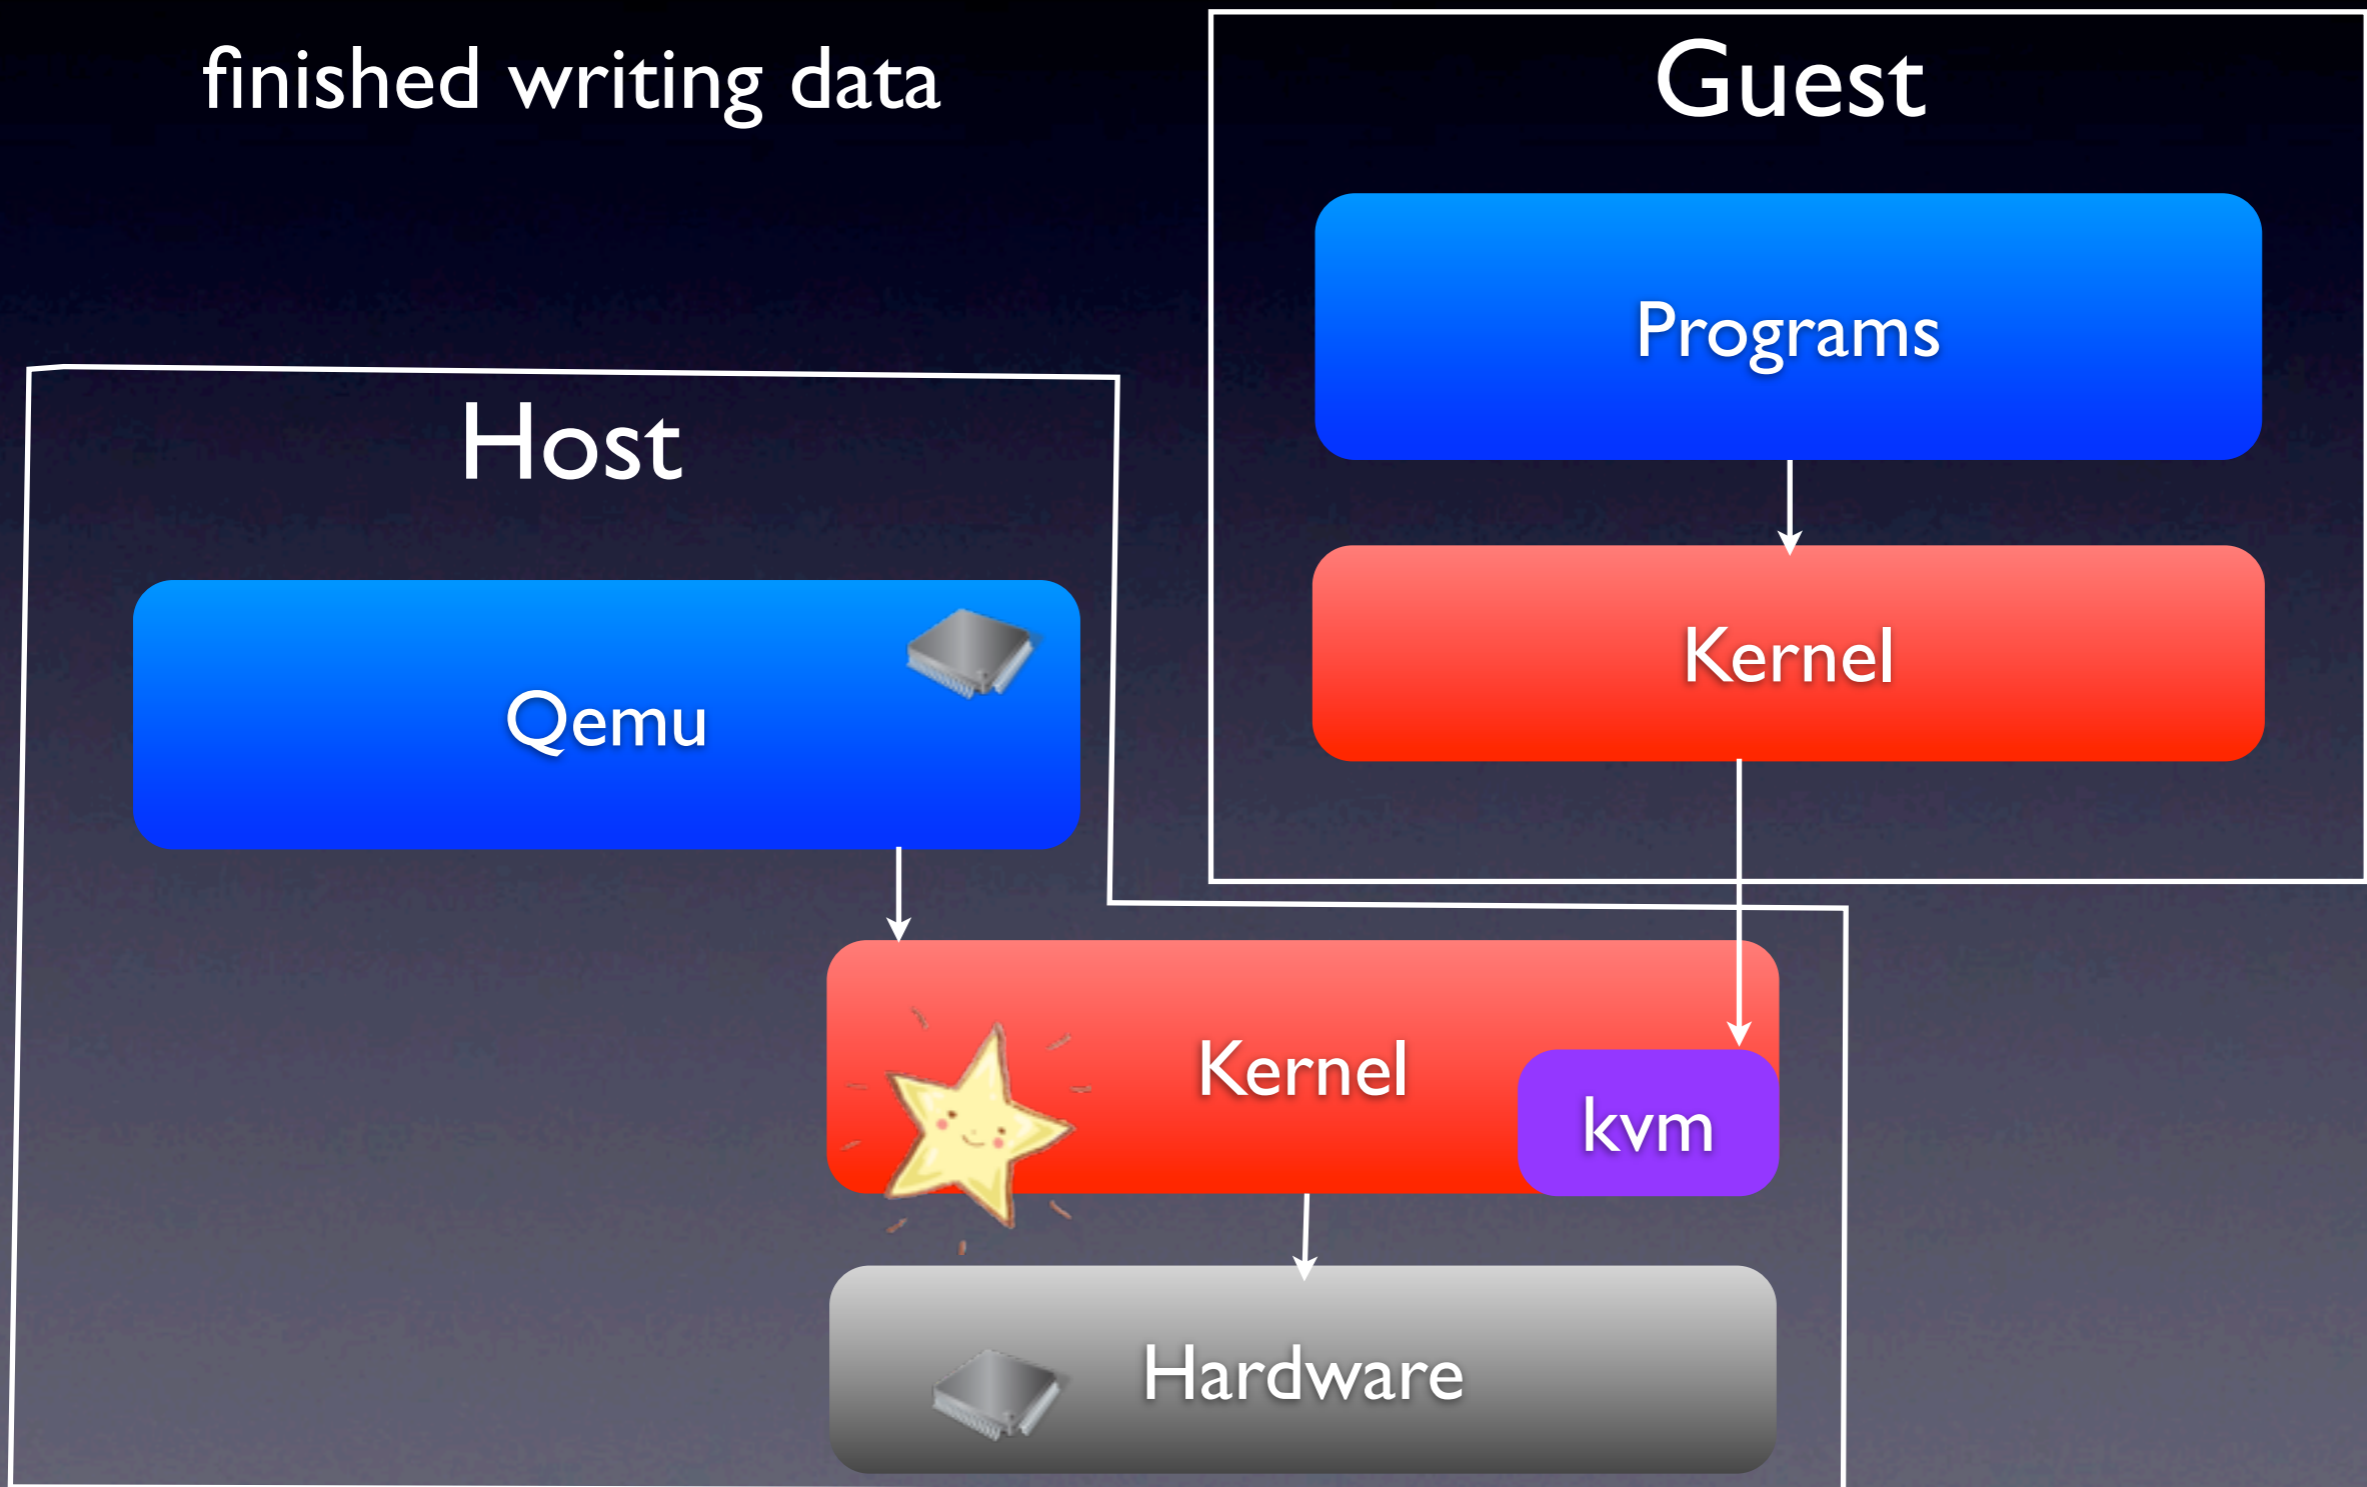
send data to hard disk





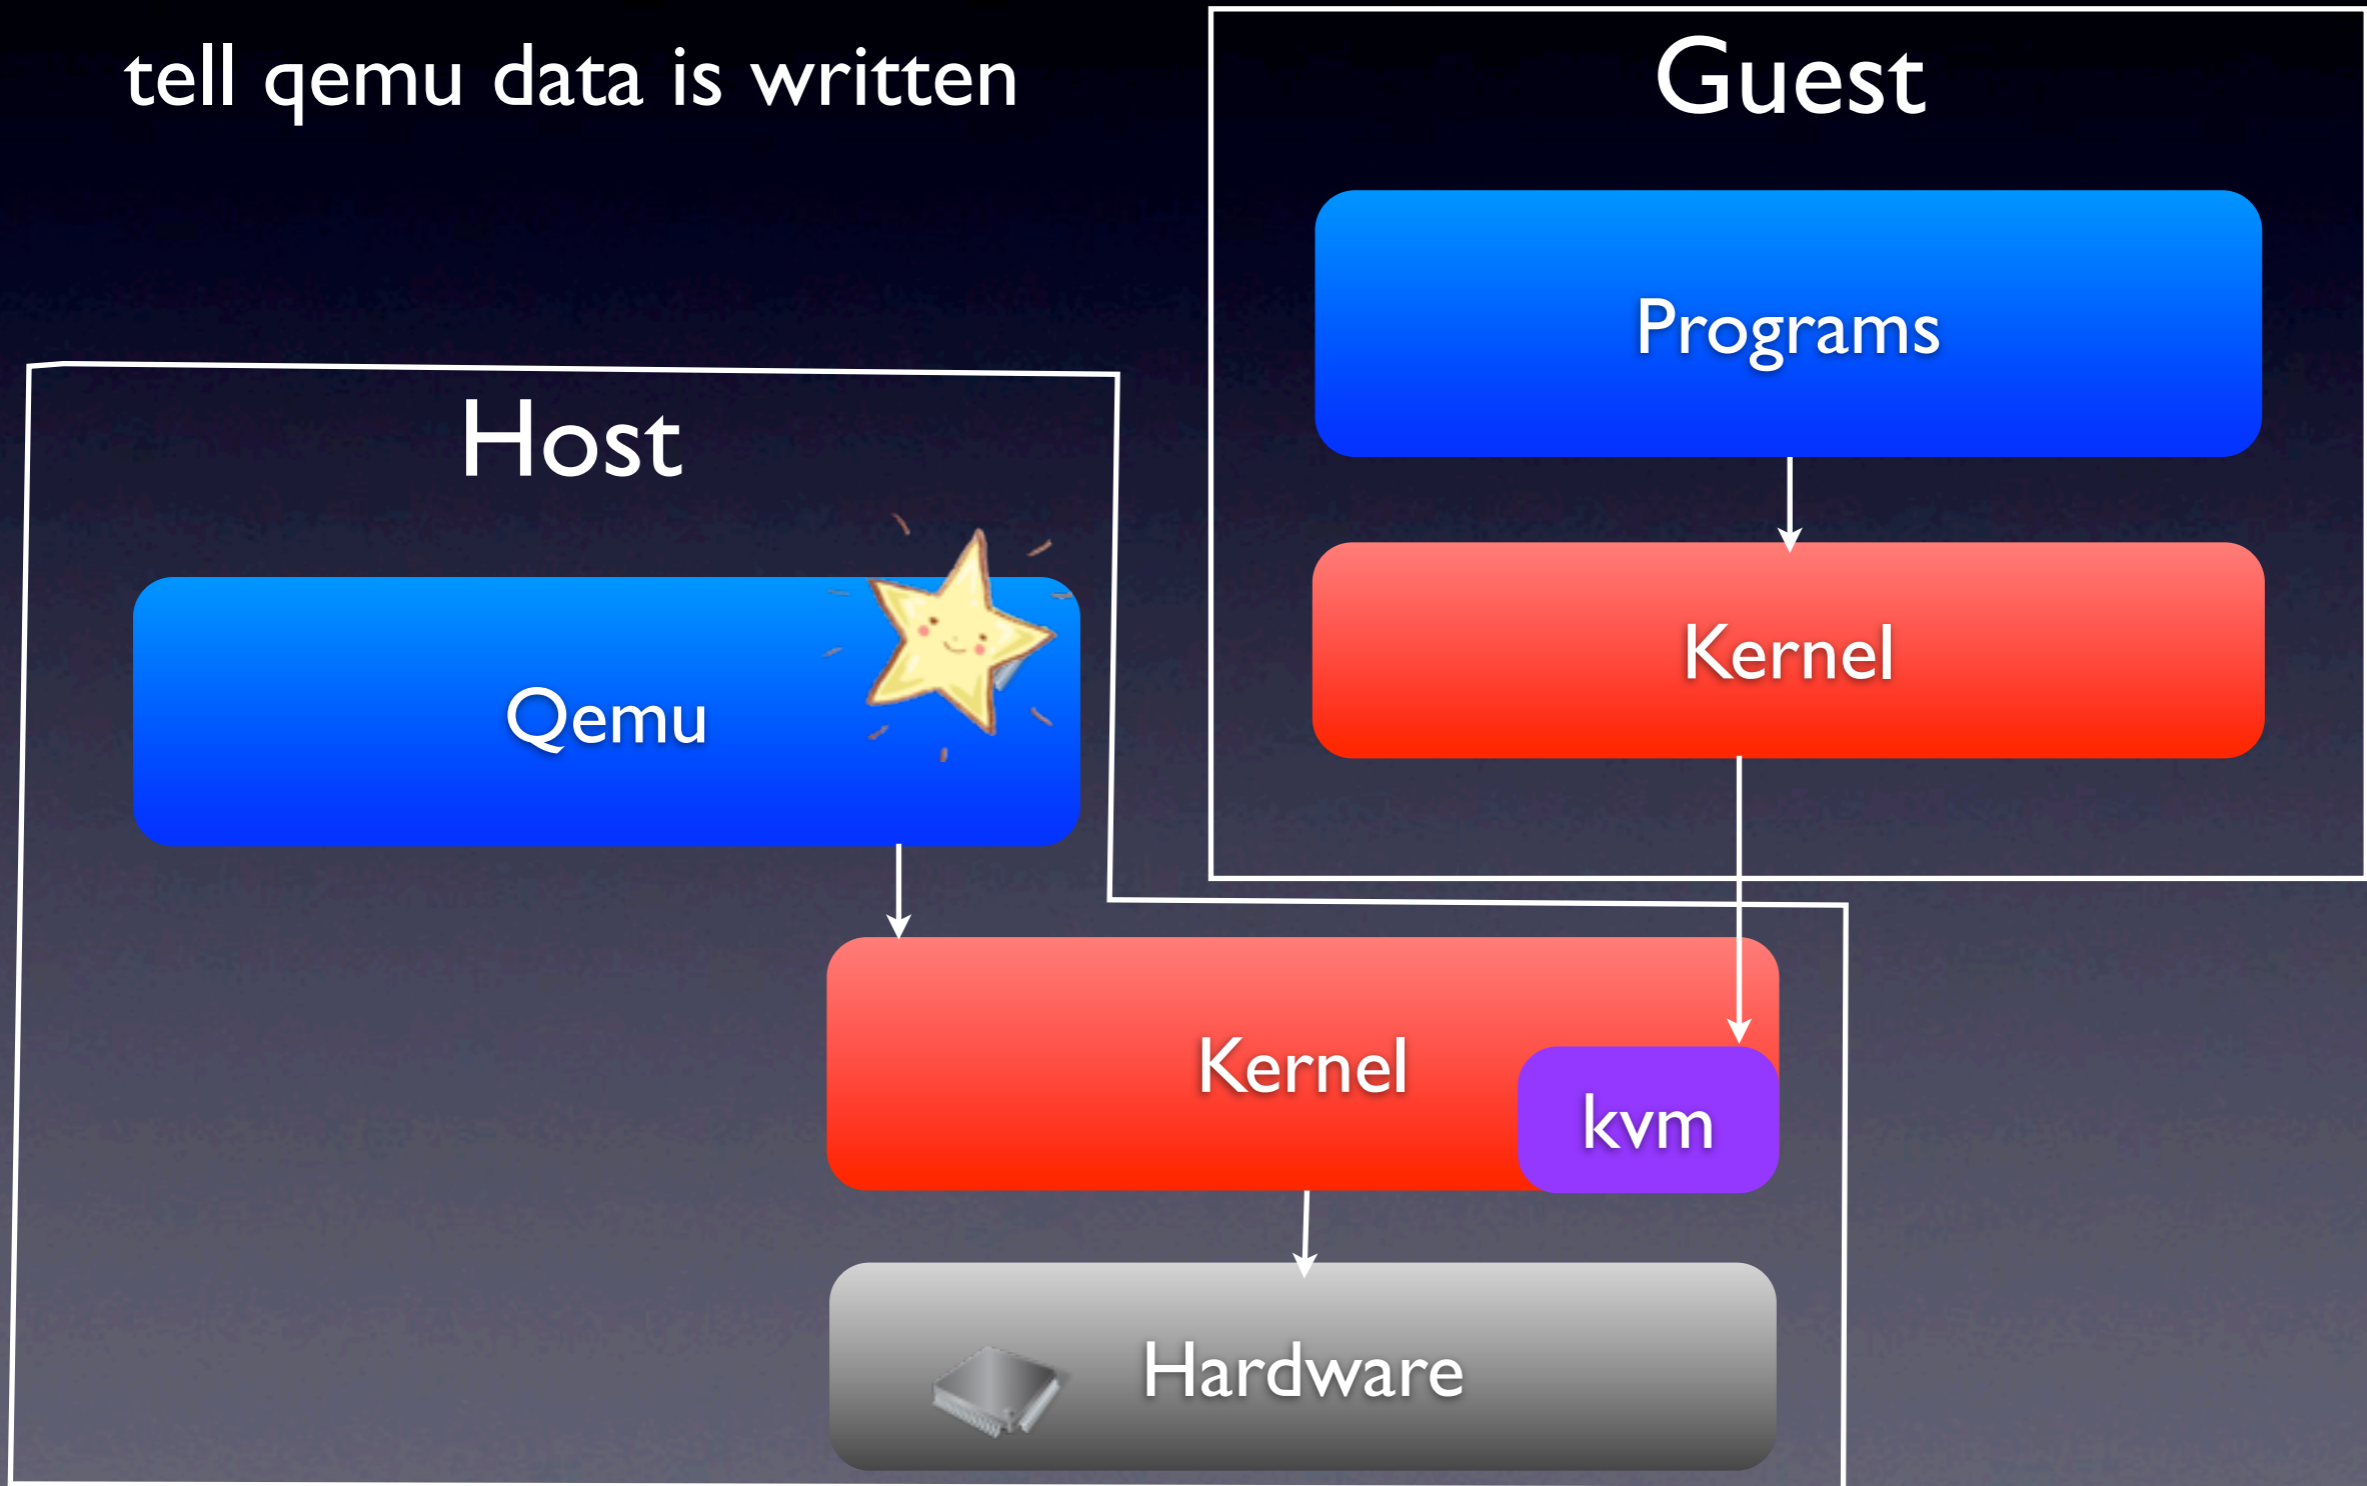
# Device Access

finished writing data



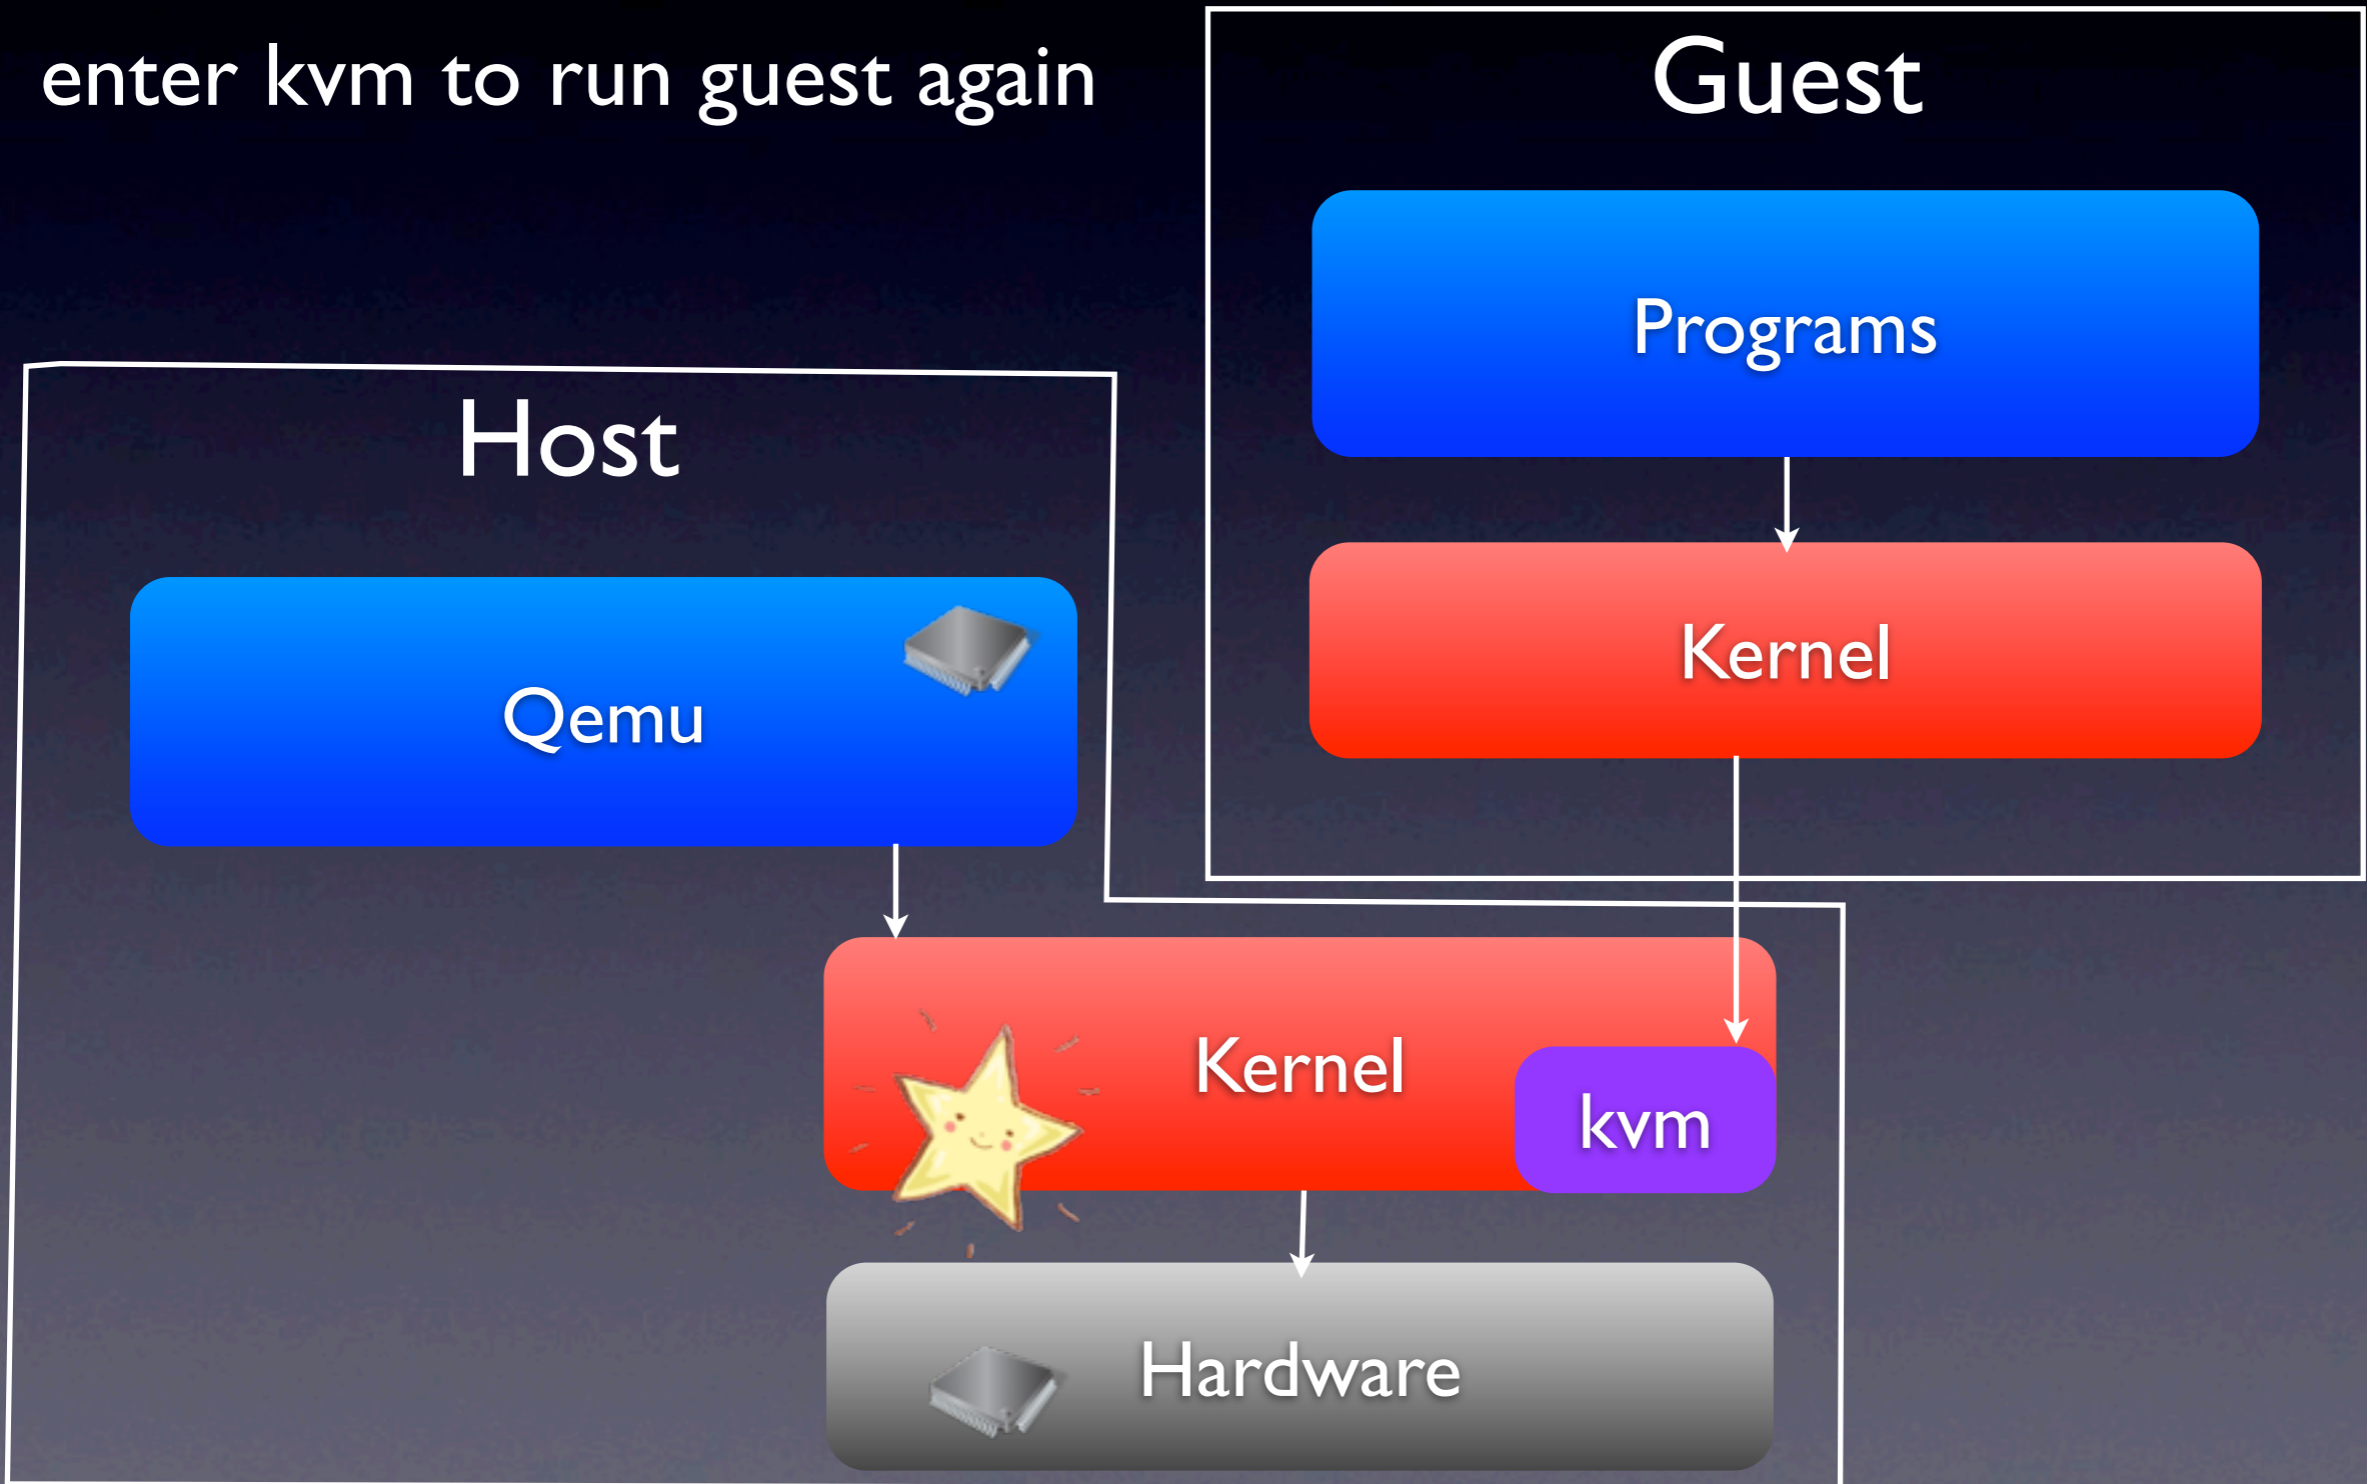
# Device Access

tell qemu data is written



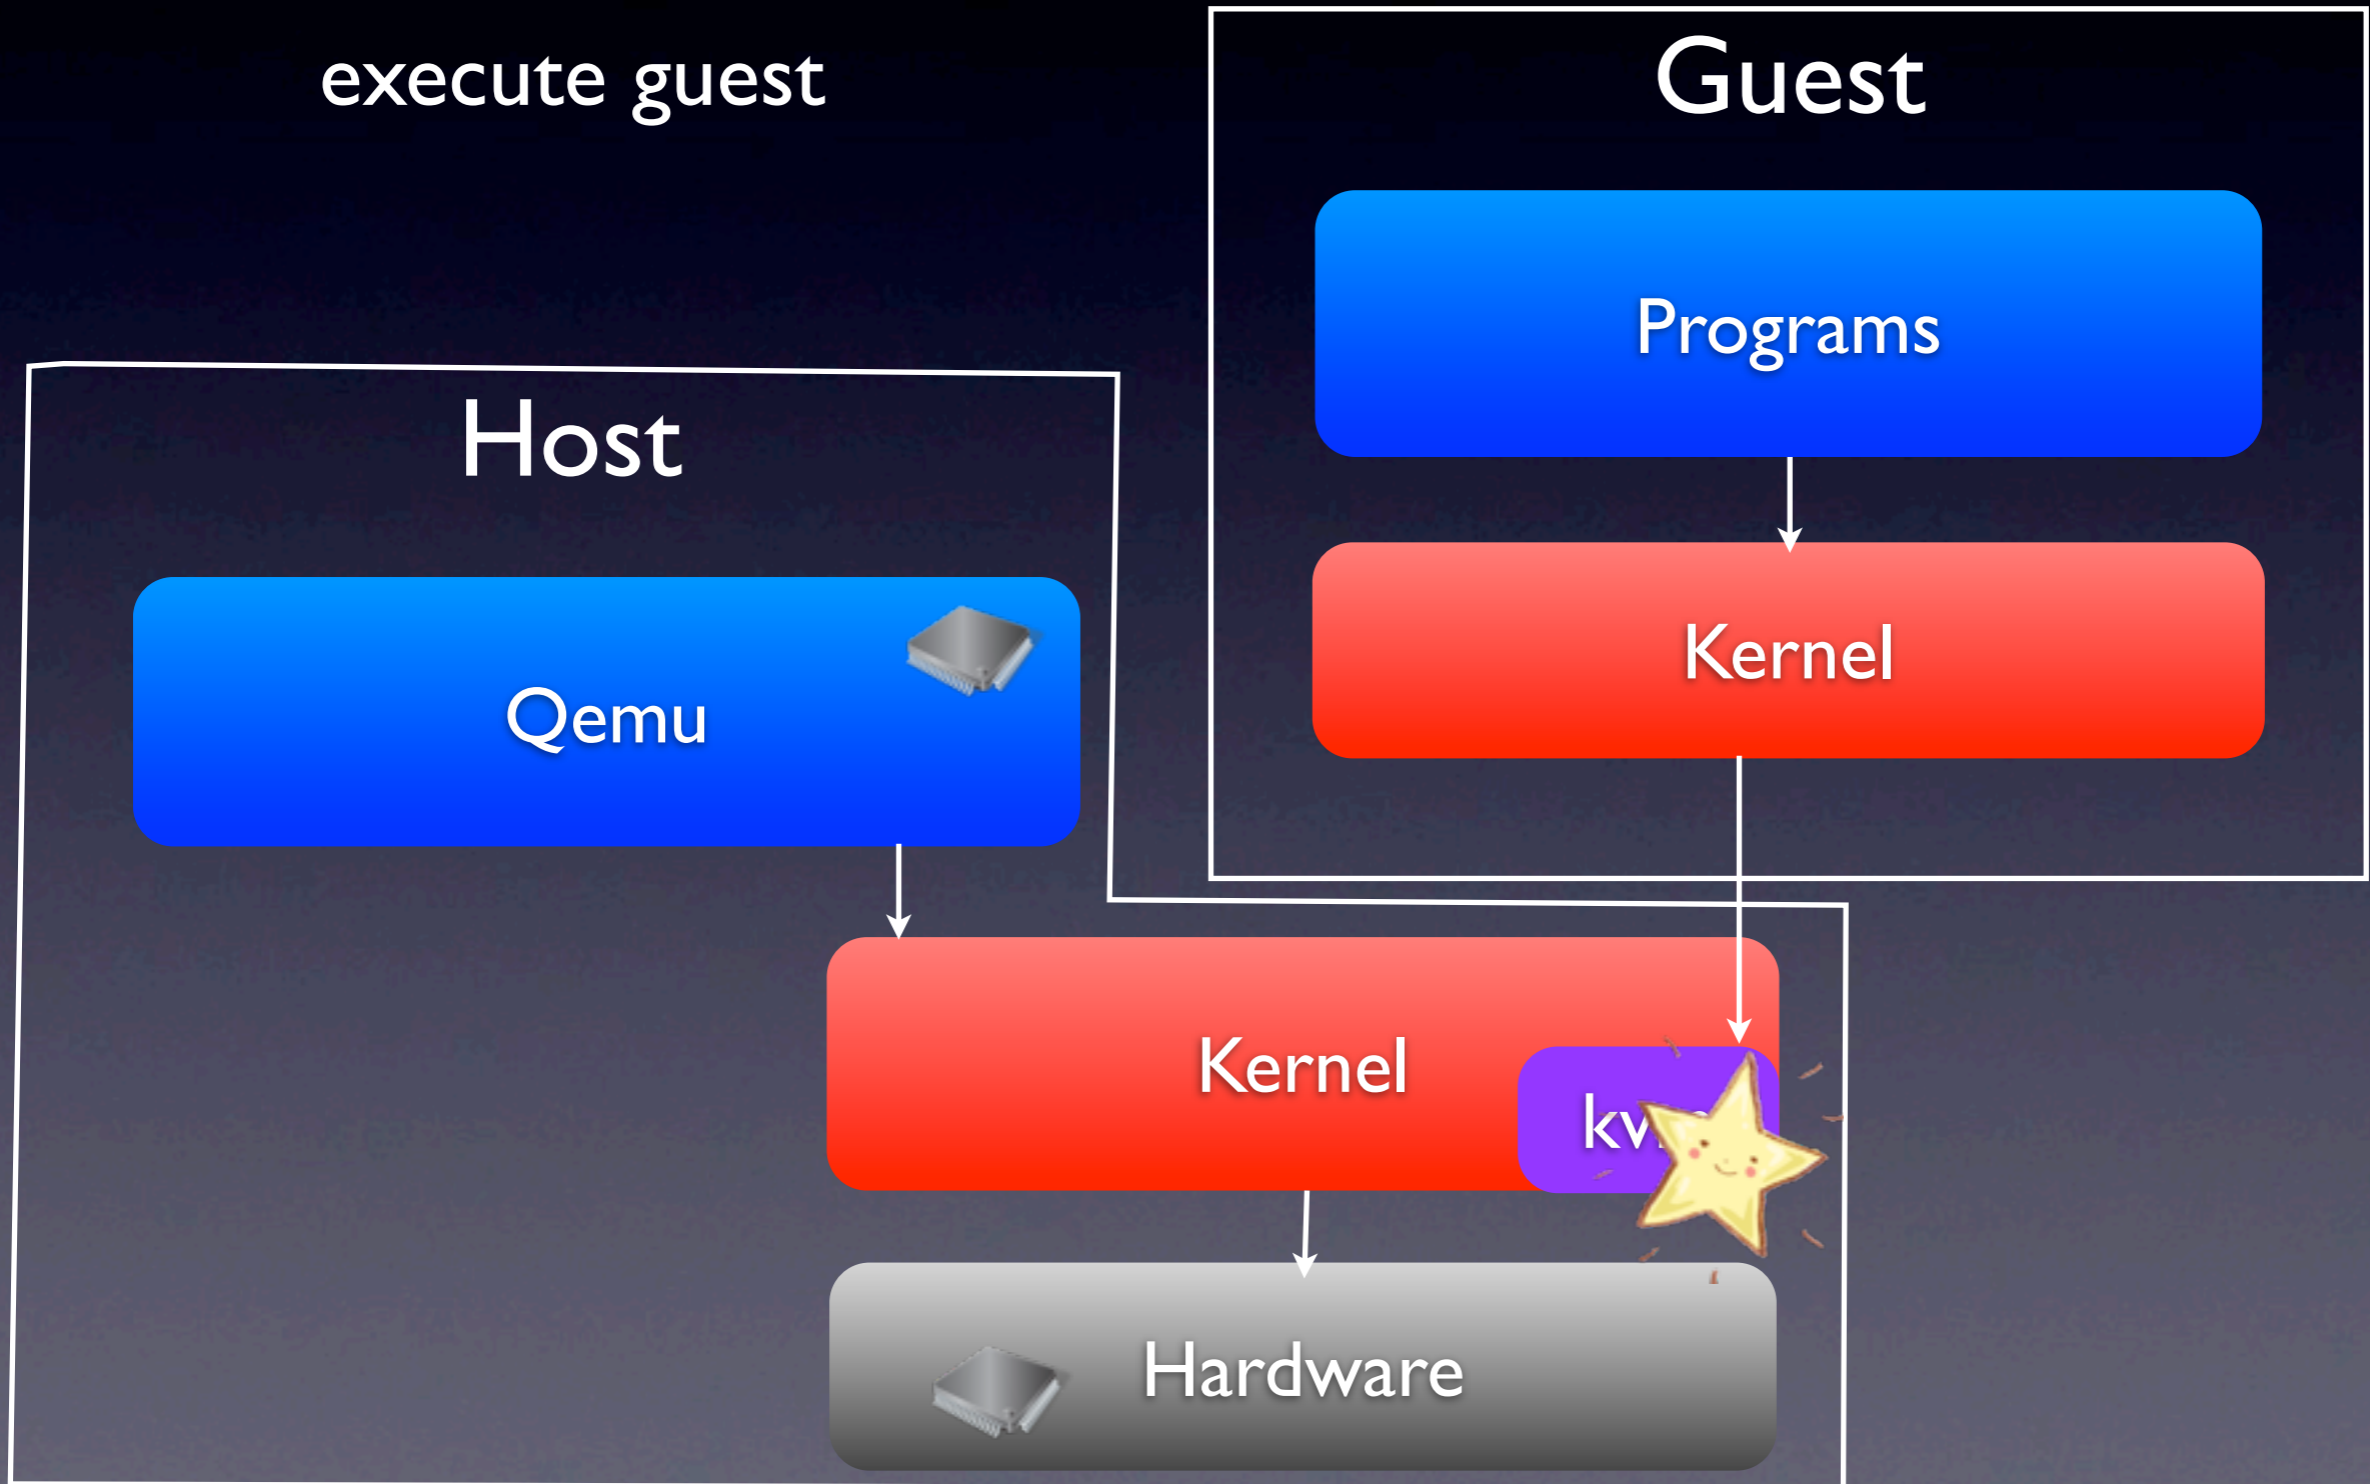
# Device Access

enter kvm to run guest again



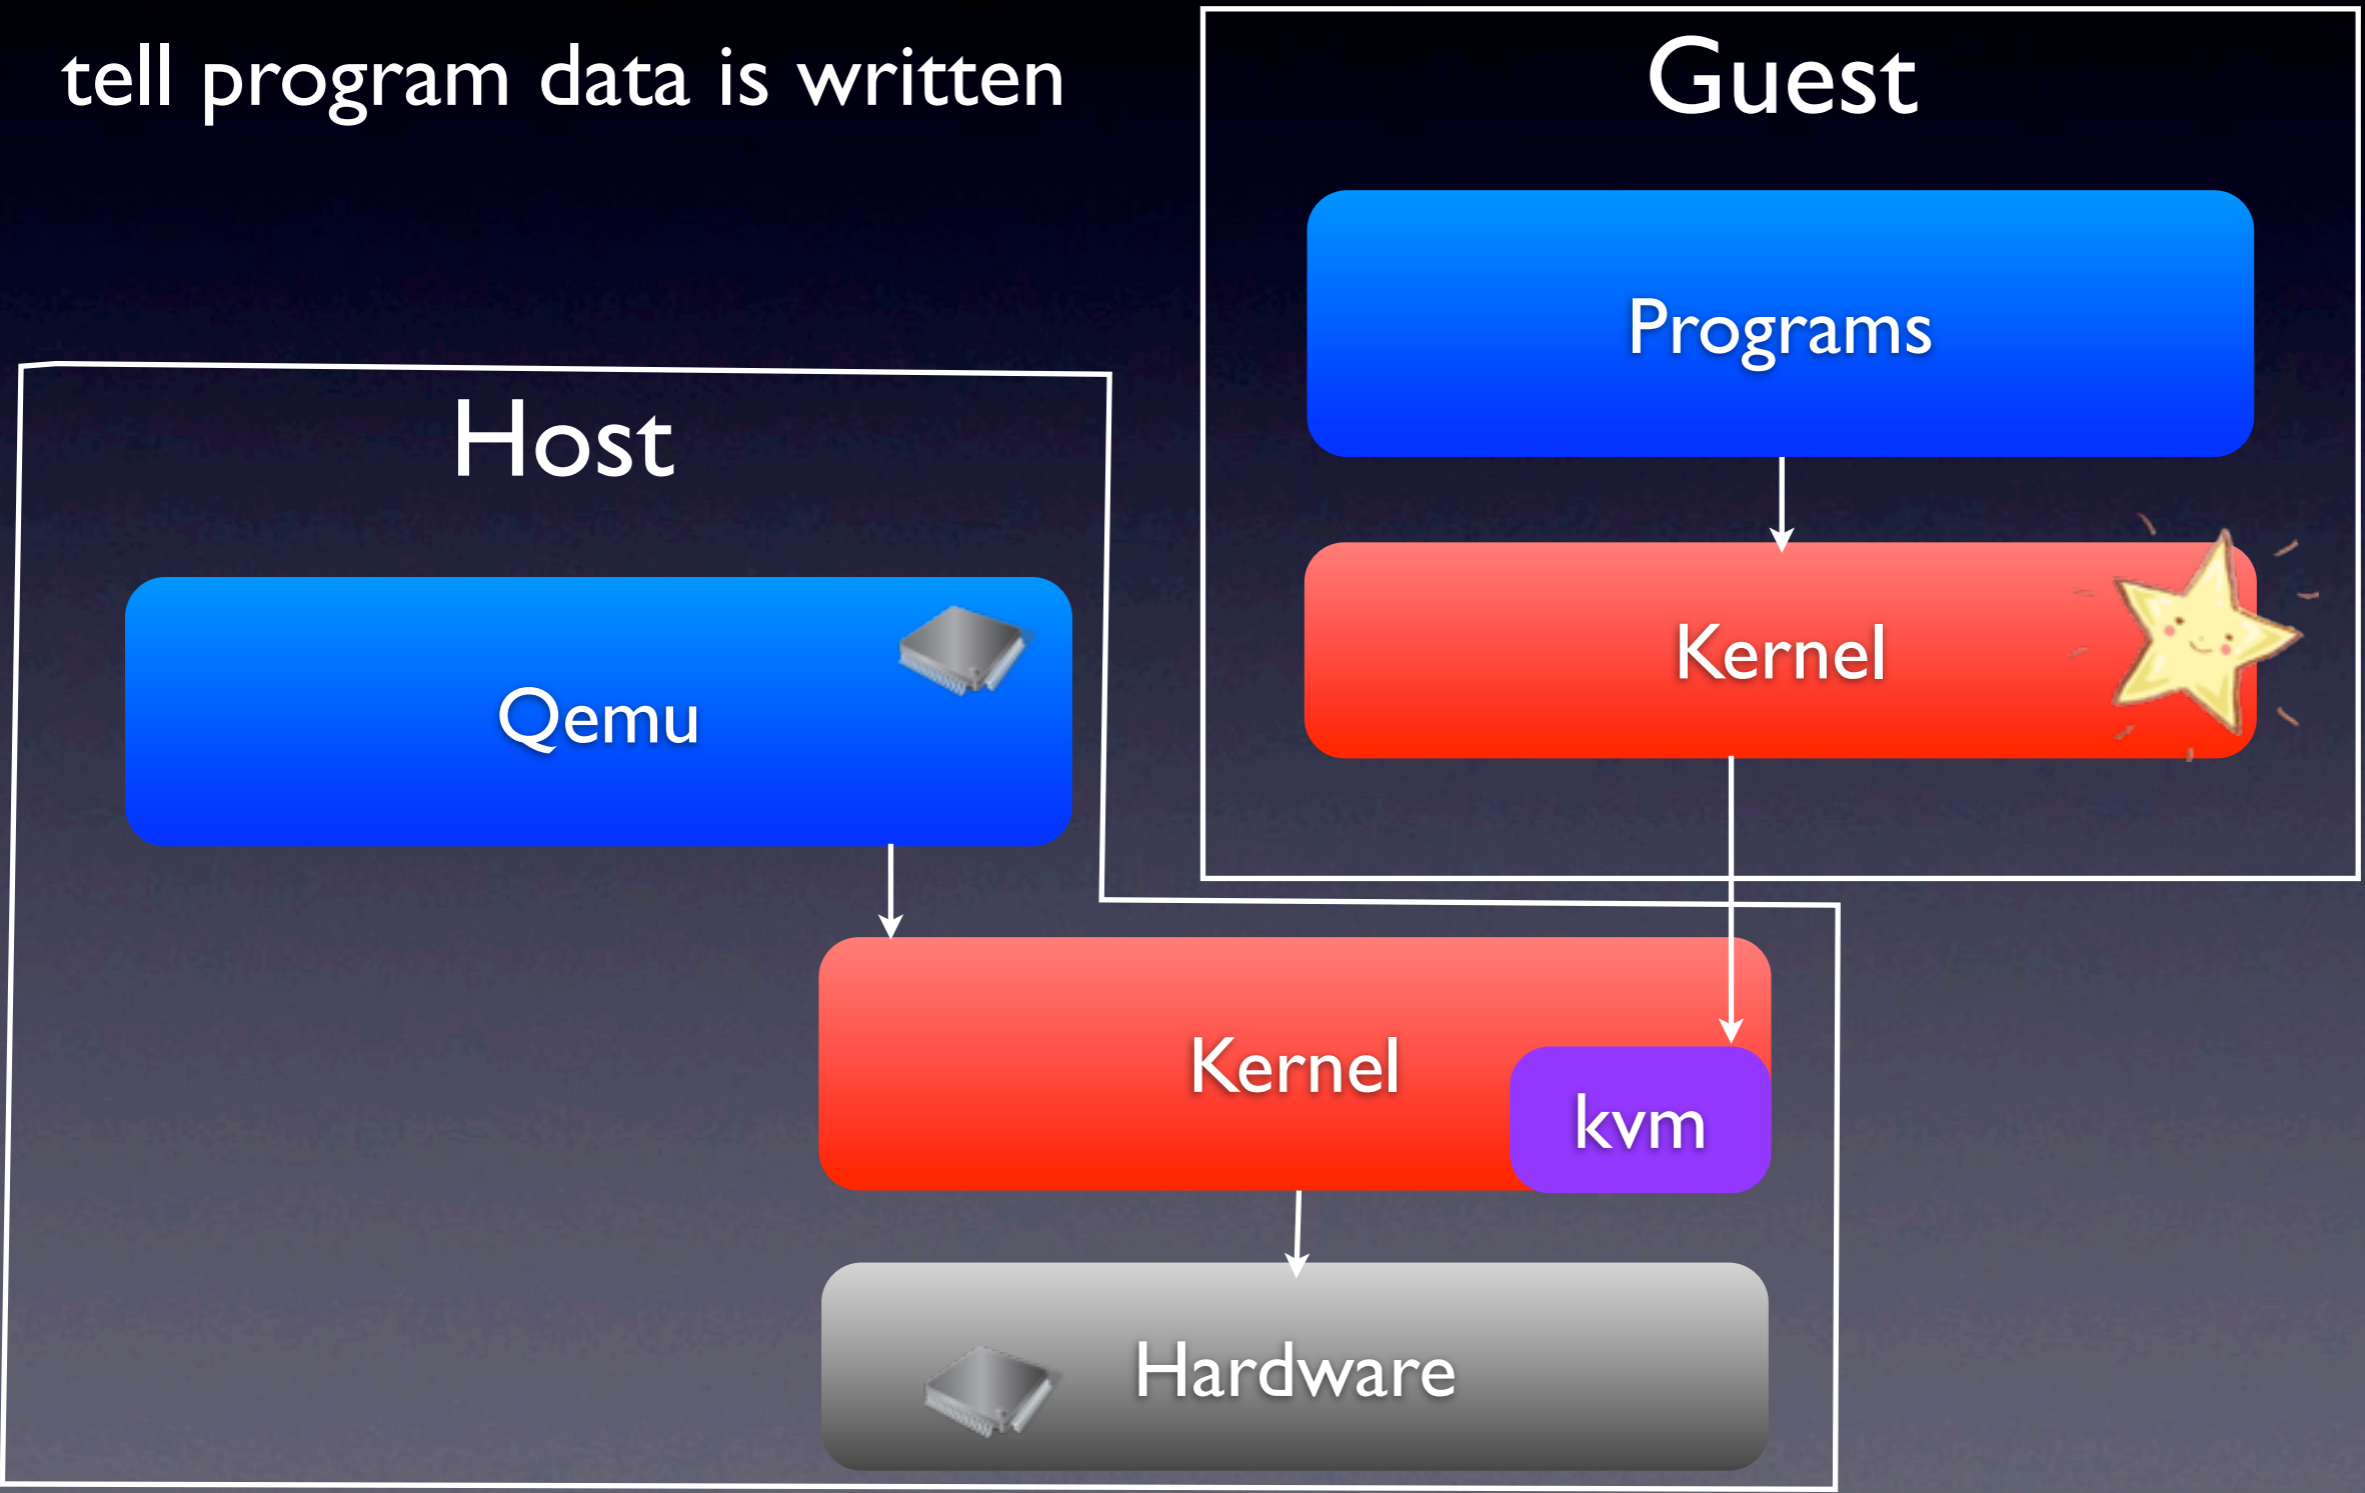
# Device Access

execute guest



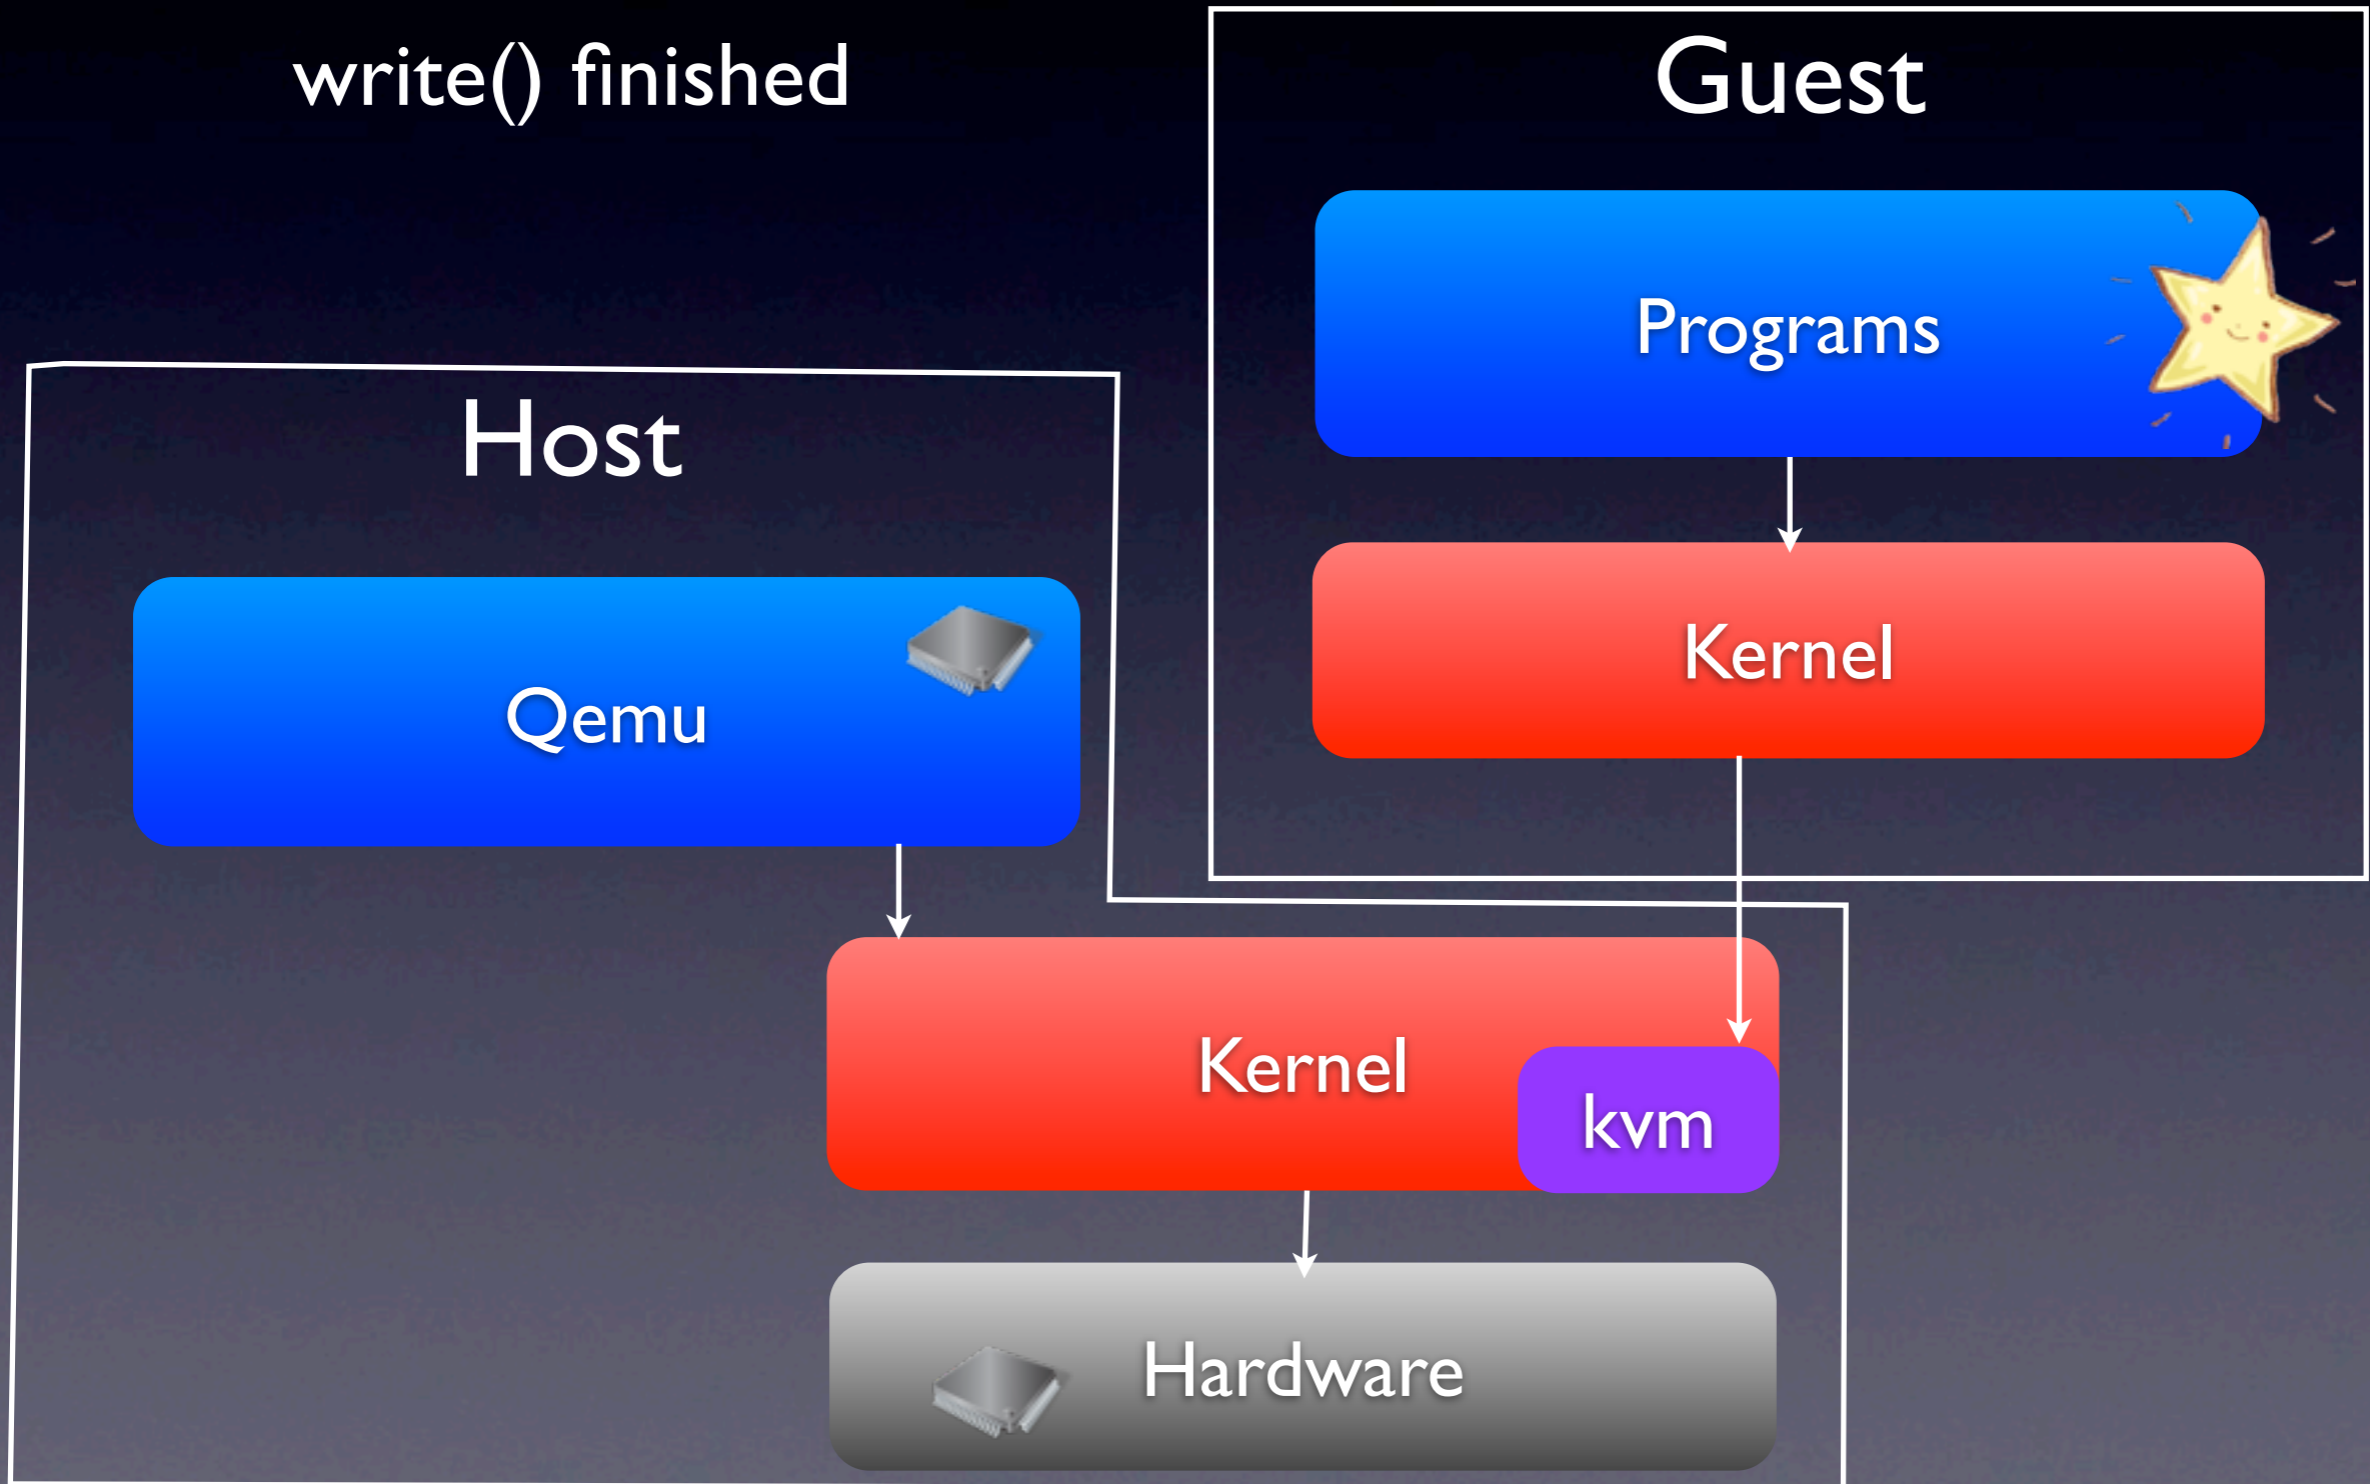
# Device Access

tell program data is written




# Device Access


write() finished



# IDE

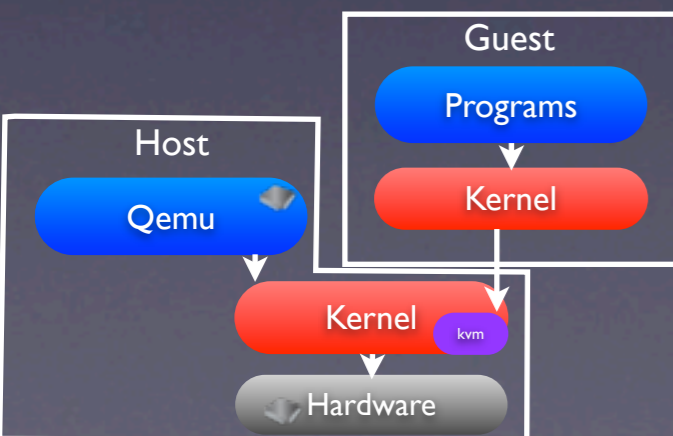


# IDE


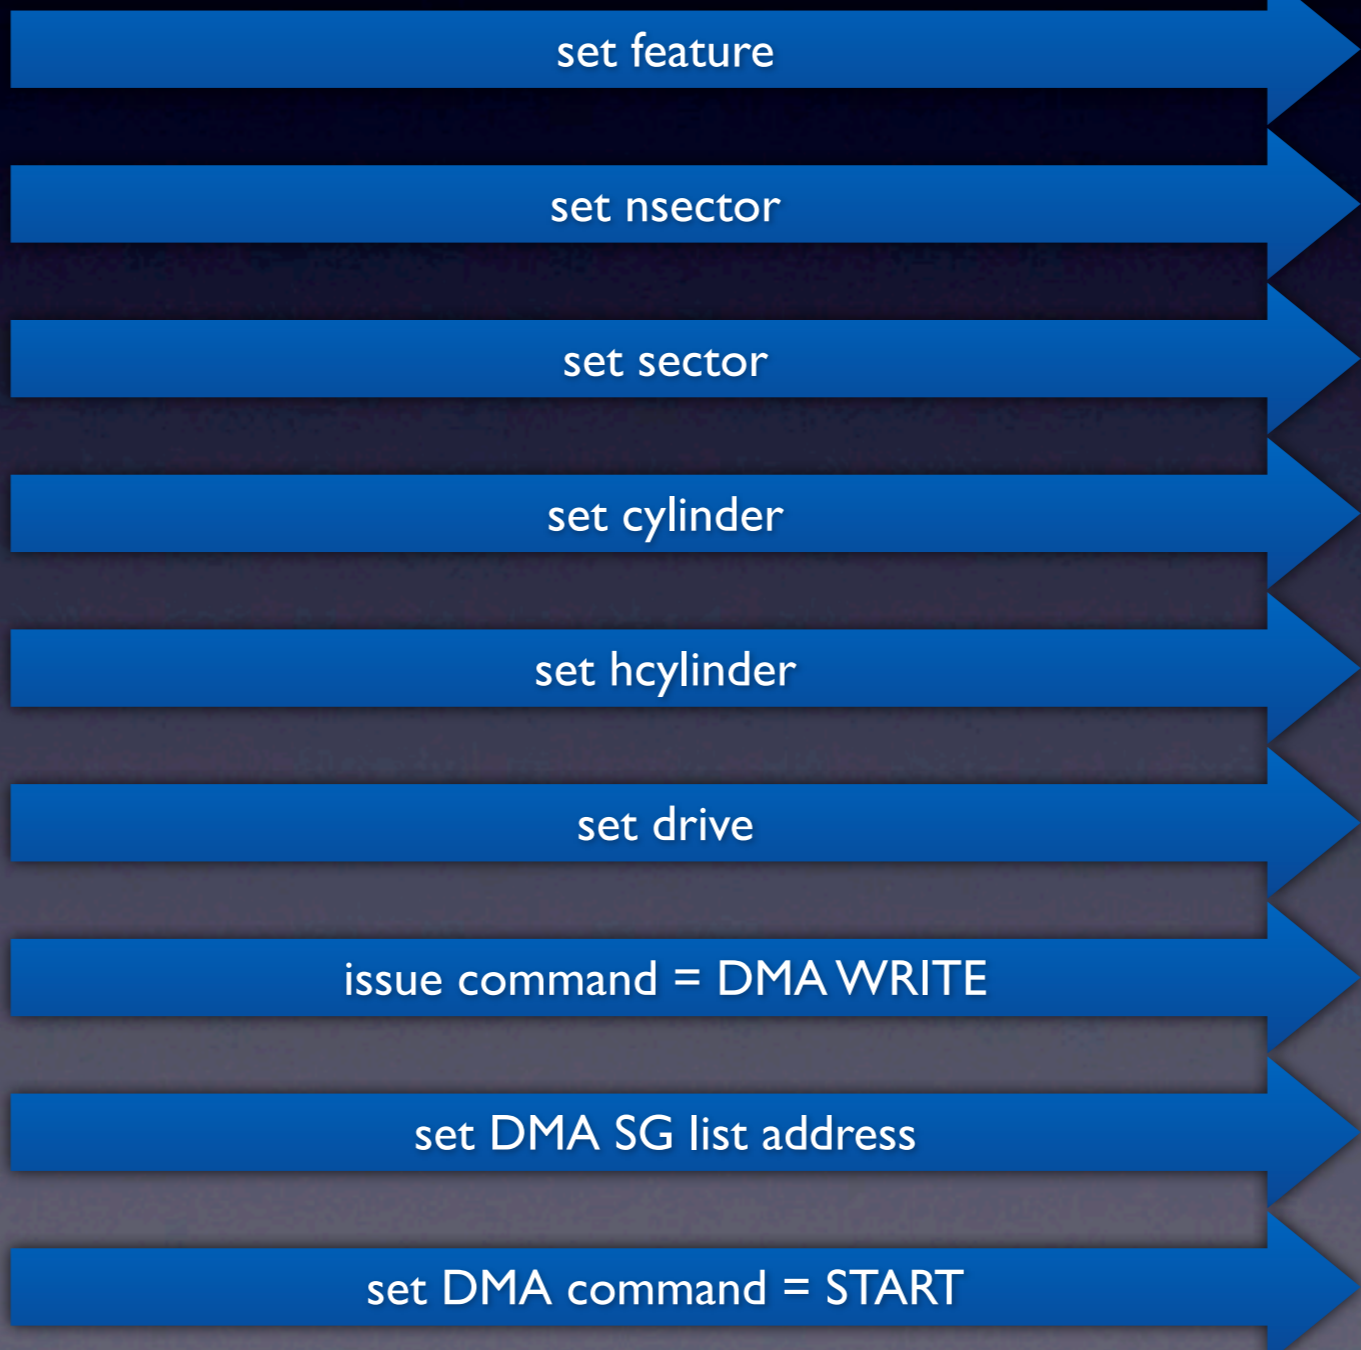





# IDE



# IDE



# AHCI



# AHCI






# virtio-blk



# virtio-blk



# Communication




old IDE

# Communication



IDE






# Communication




AHCI



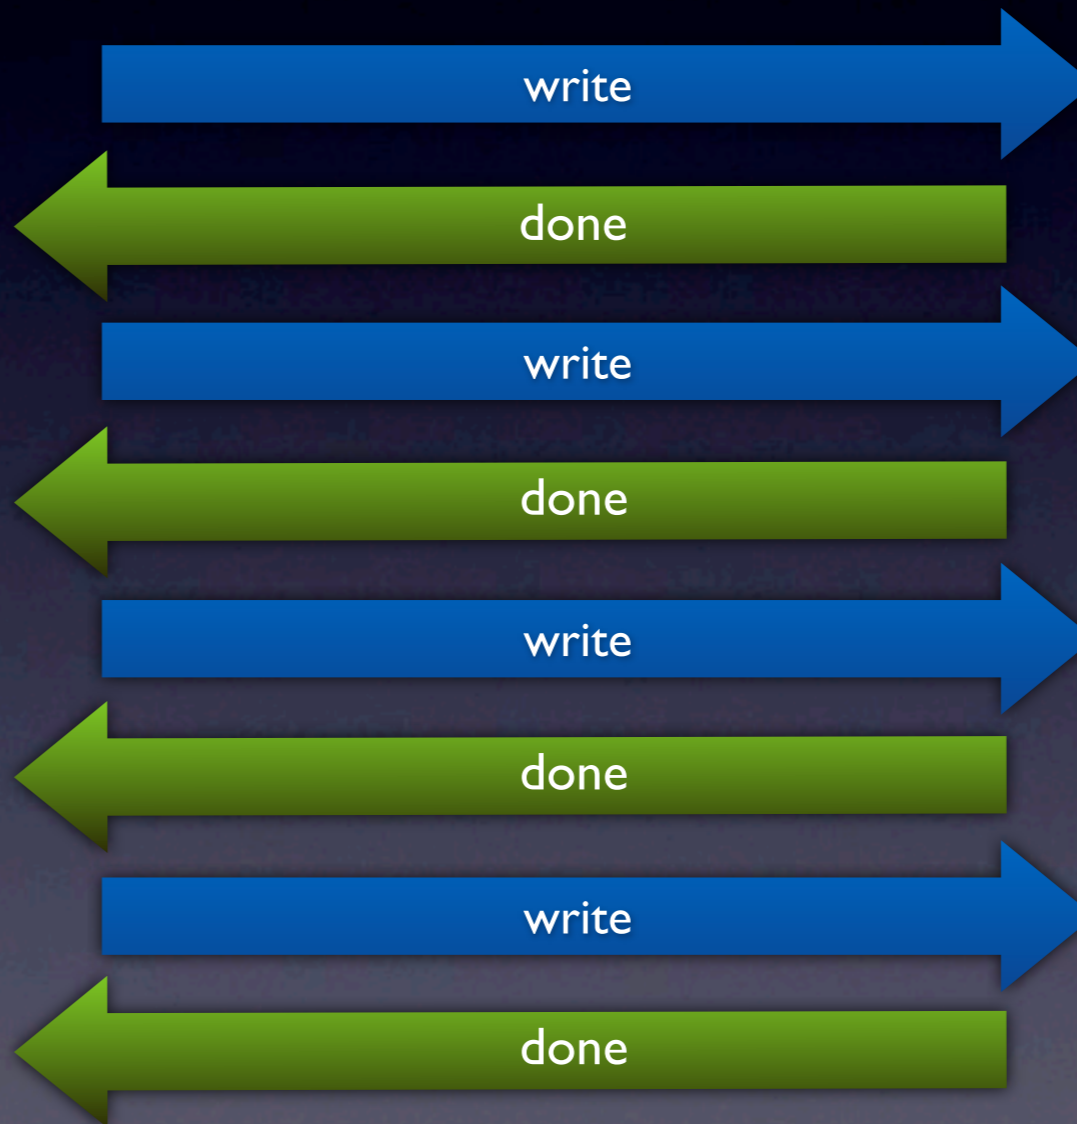
# Communication



virtio-blk


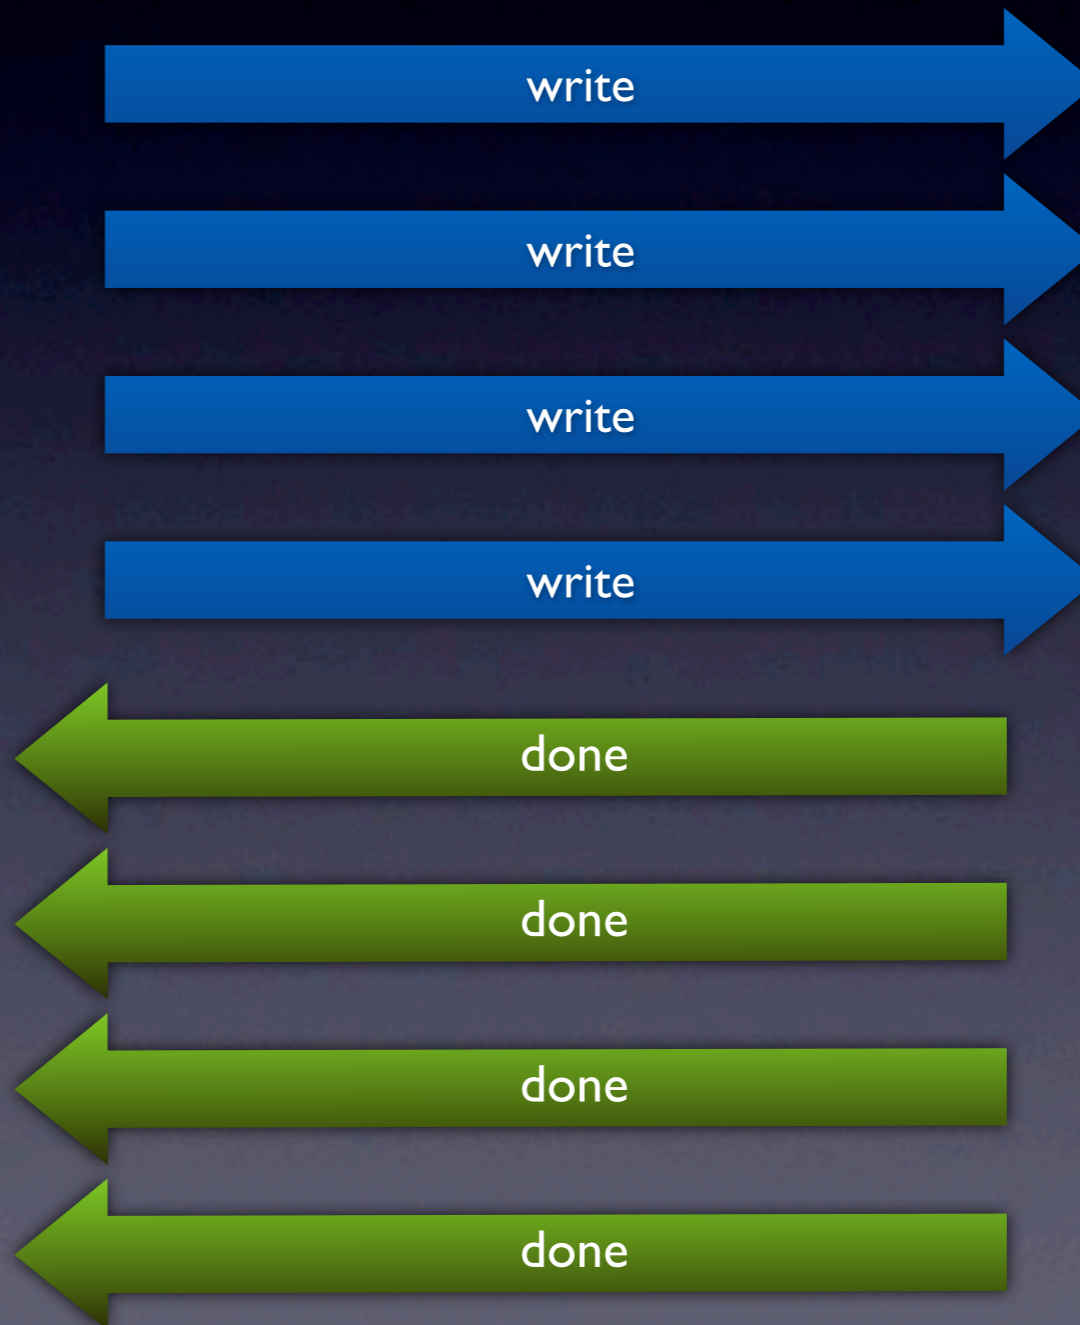


# Multiple Requests







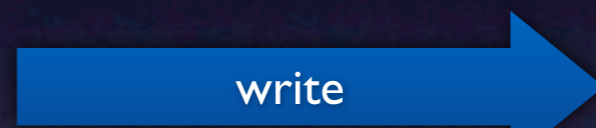

IDE

# Multiple Requests



AHCI  
(NCQ)

# Multiple Requests




virtio-blk

# Guest OS Support

	IDE	AHCI	Virtio-blk
Linux	✓	✓	✓
Windows XP	✓	✗	external
Windows Vista	✓	✓	external
Mac OS X	✗	✓	✗
BSD	✓	✓	✗




# Speed

qemu.git    qemu-kvm.git













(host) `qemu-img create /dev/shm/test.raw 20G`  
(guest) `dd if=/dev/Xda of=/dev/null bs=1M count=1000 iflag=direct`  
Benchmark ran on Intel Atom z530

# Caching

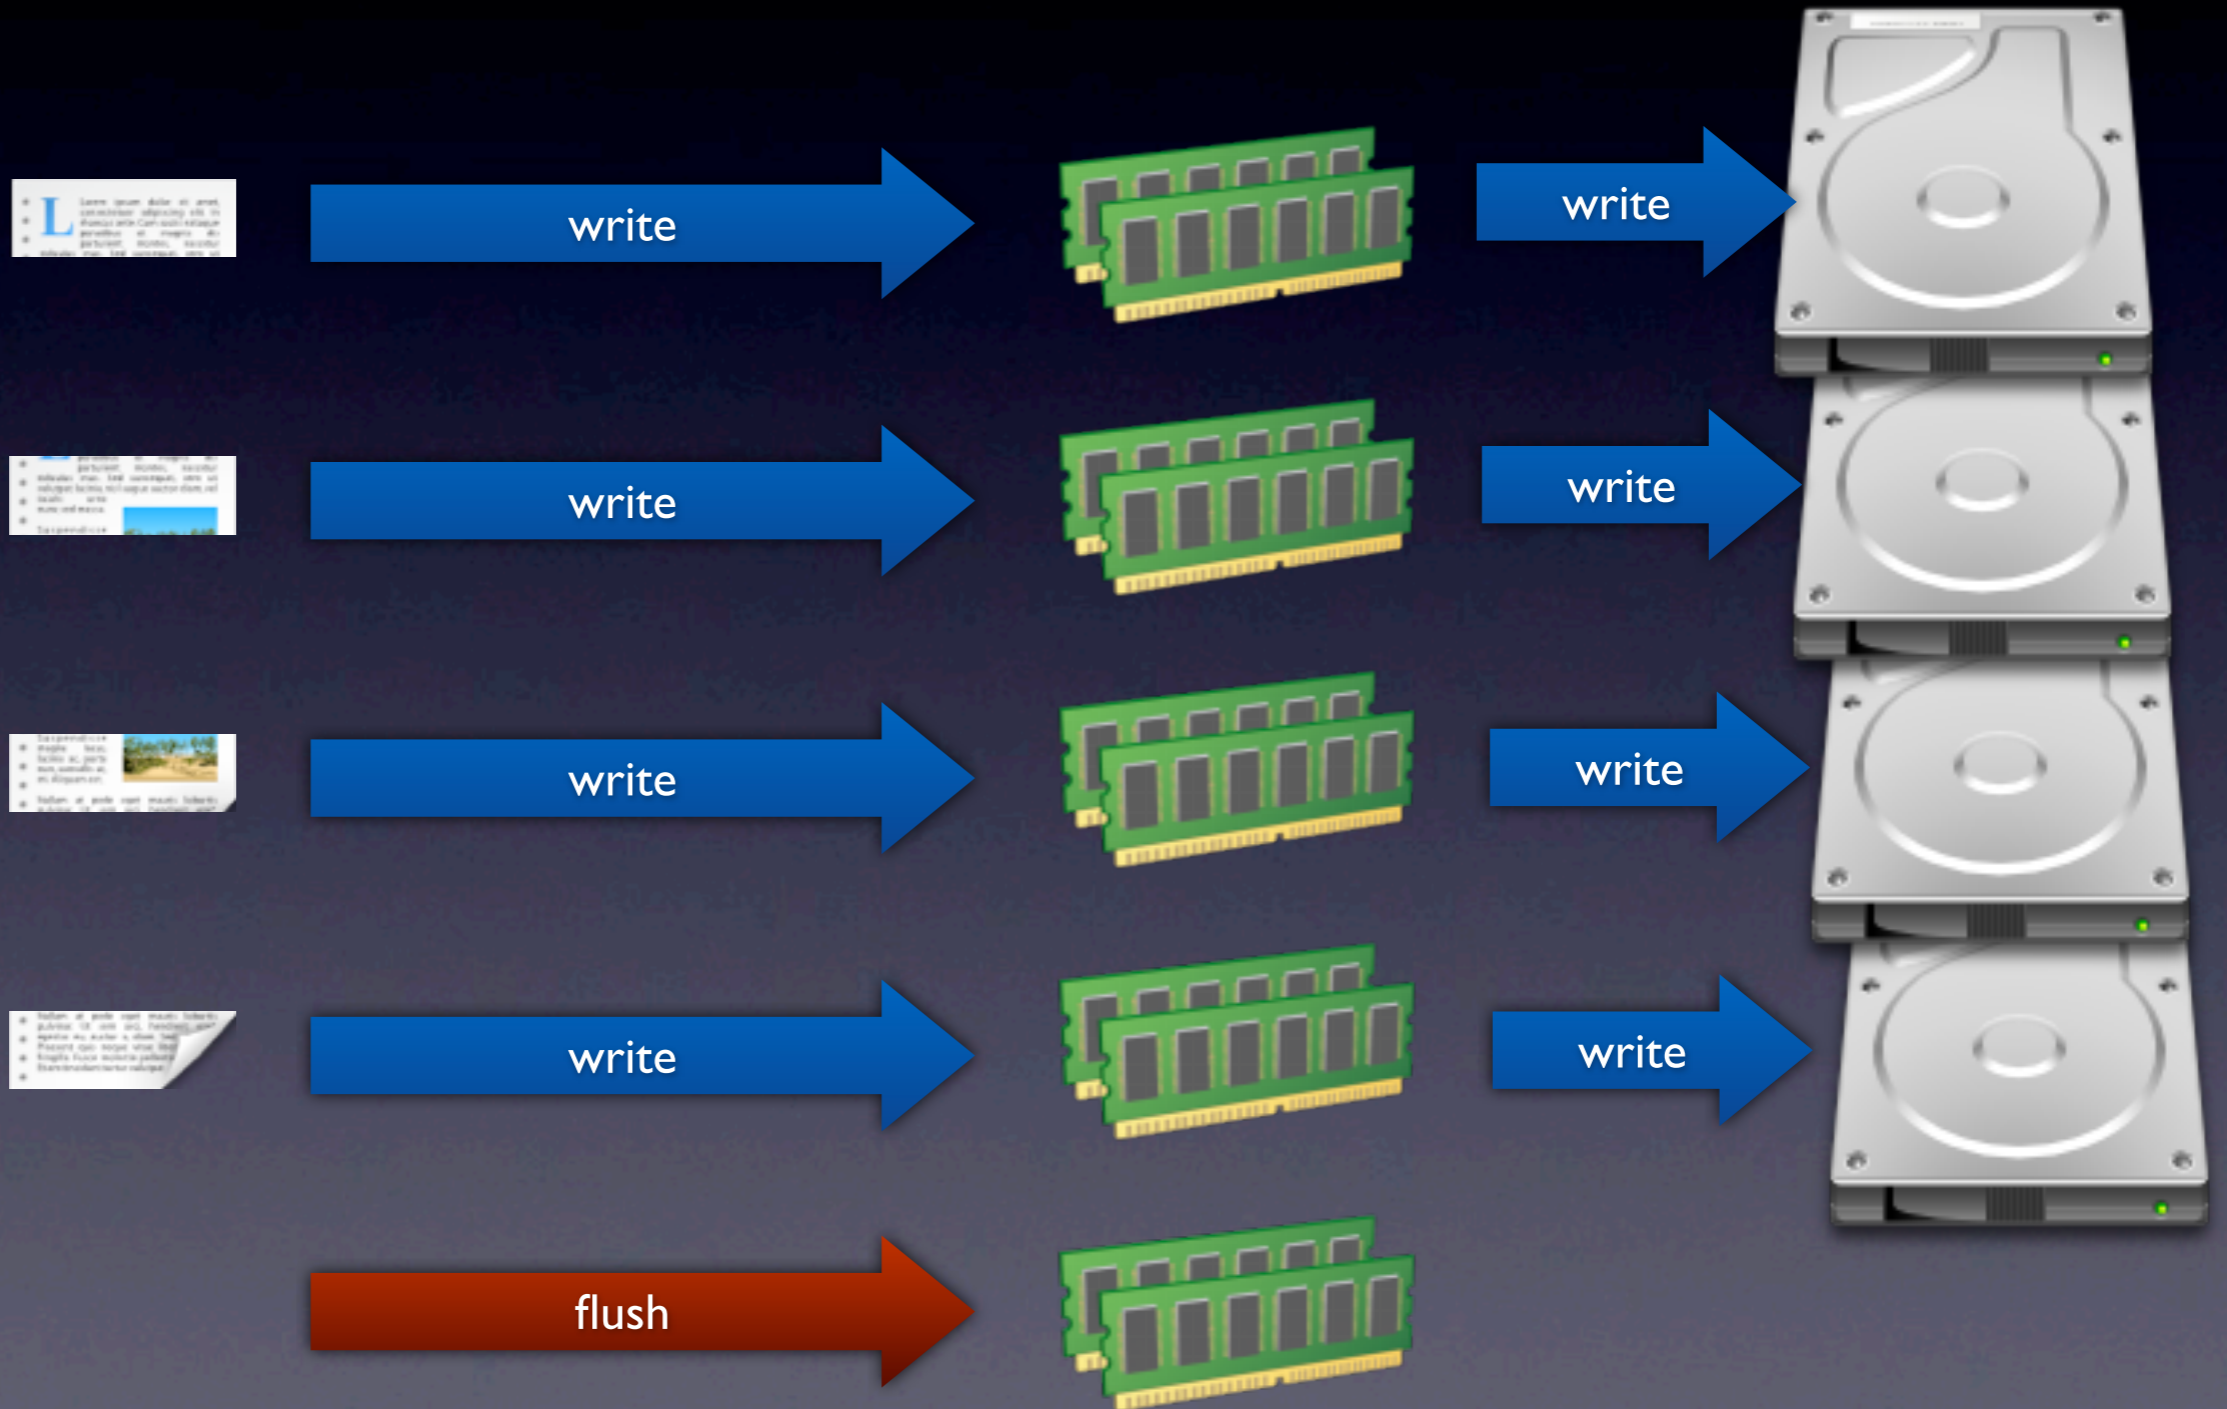




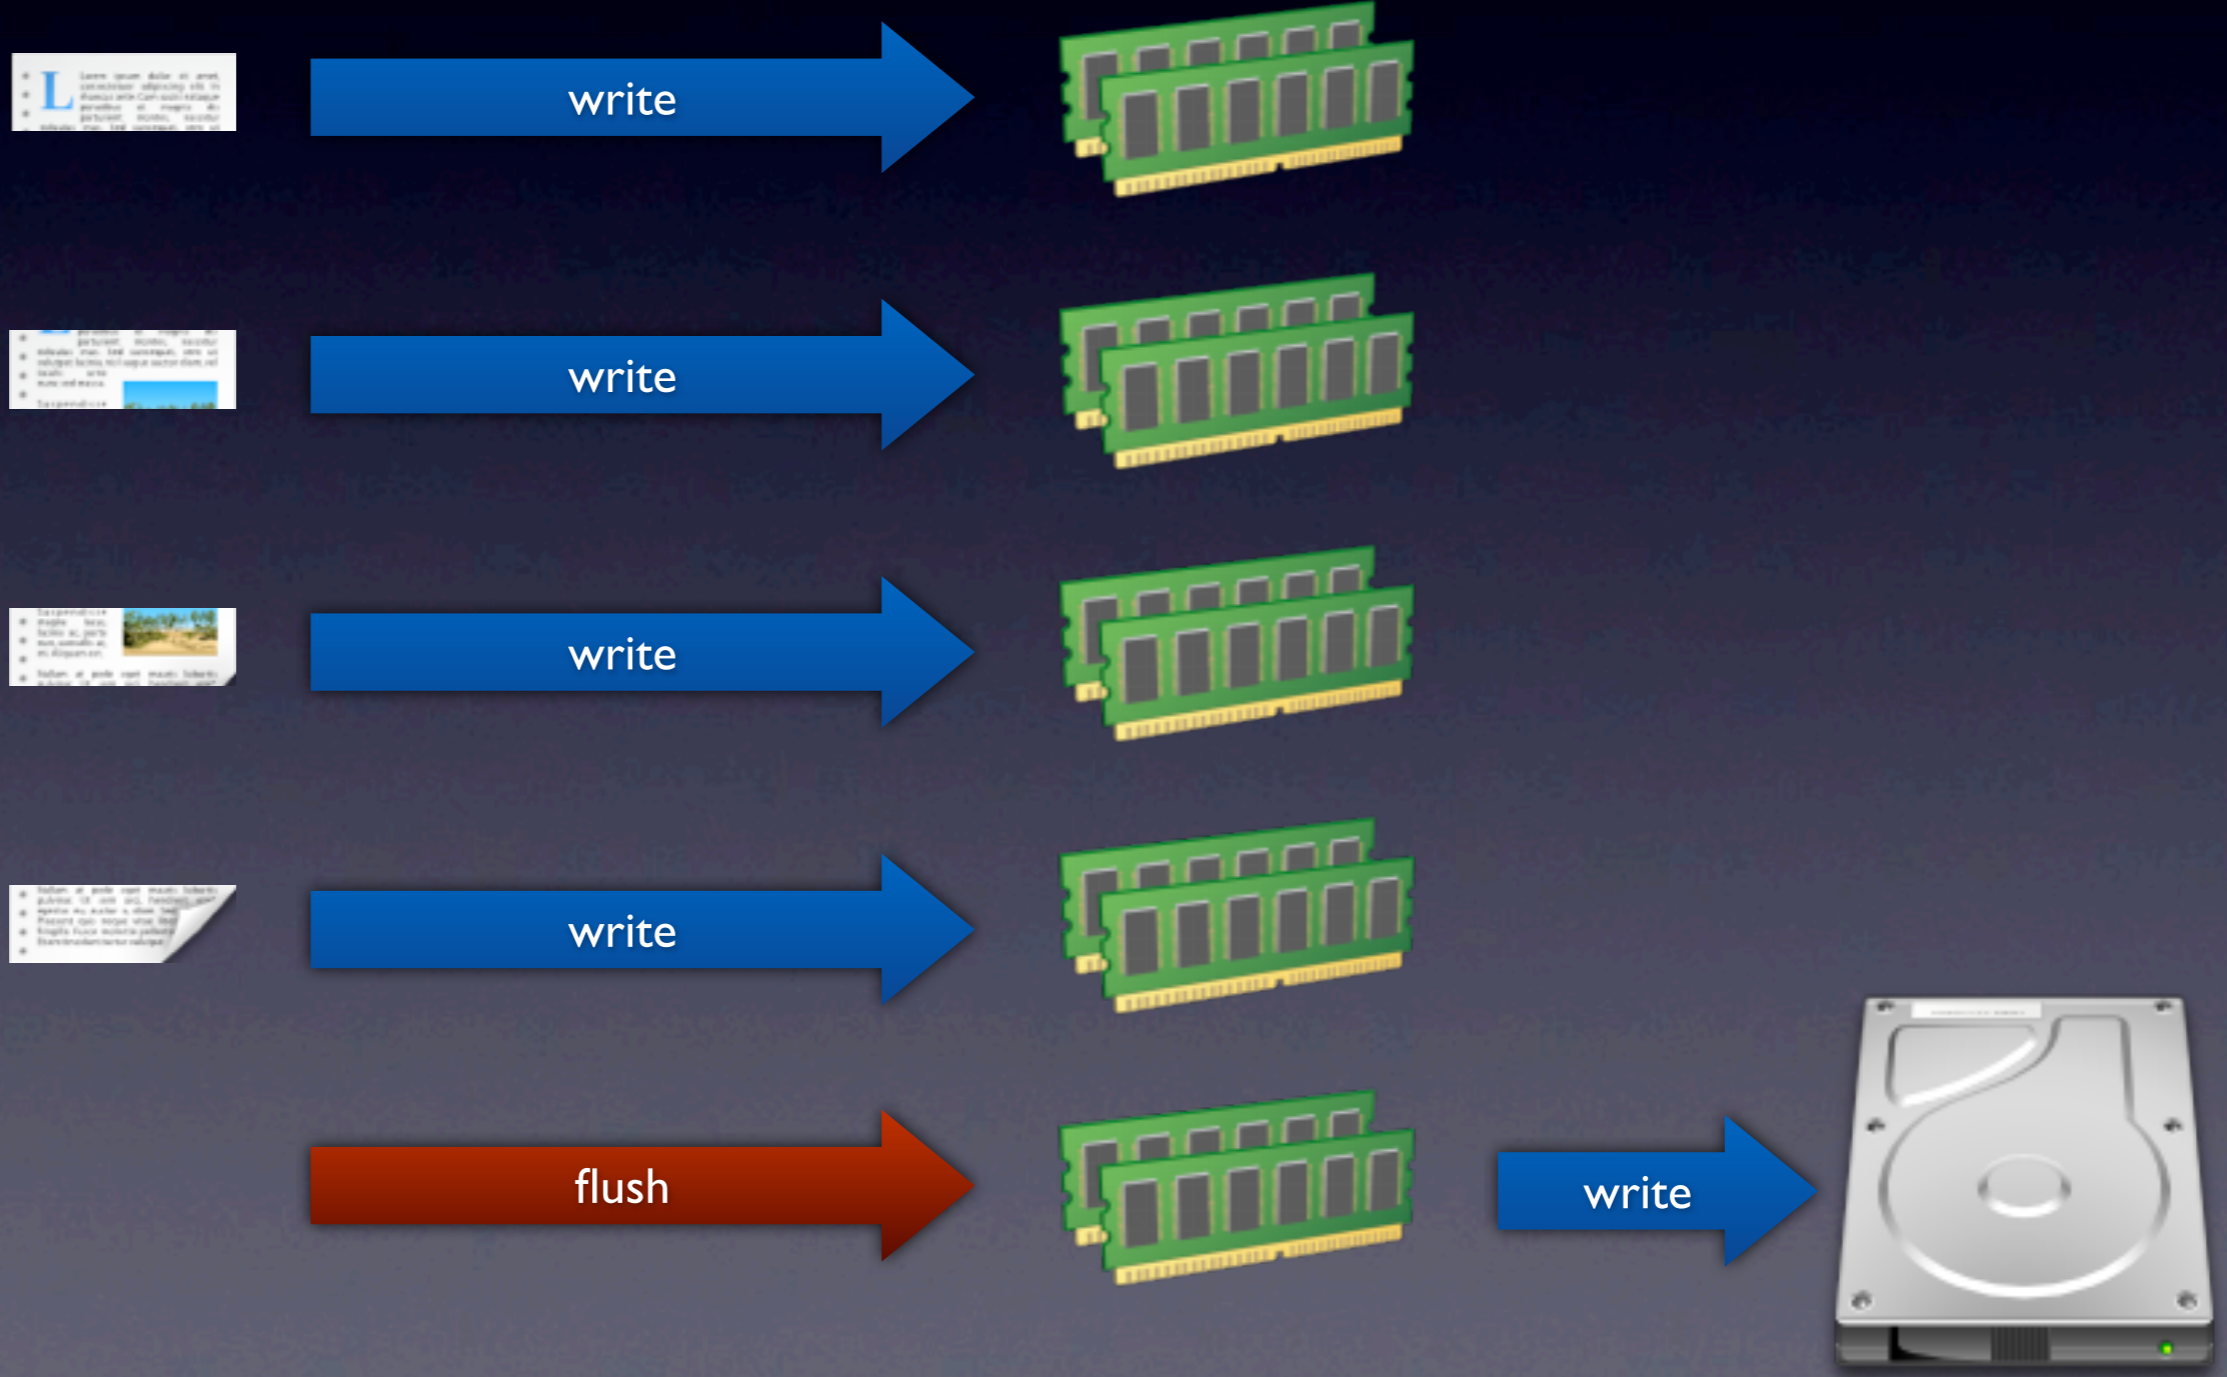
# Caching













# cache=writethrough







# cache=writeback



# cache=unsafe



# cache=none



# Why AHCI

- Faster than IDE
- Compatible with most OSs
- Compatible with all modern devices (CD-ROMs)
- Only need to develop one side

# Future ideas

- Default adapter in -M q35
- MSI-X

# Questions?



蟻が十