クラウド開発演習（課題解決型学習）
(A2) High-performance Cloud

ソフトウェア・クラウド開発プロジェクト実践III

浅井大史

2015年5月22日

(A3) High-Performance Cloud

- Goal
  - High-performance cloud (network storage)
    - understand bottlenecks and resolve them
- Possible bottlenecks in IaaS (and PaaS)
  - CPU bound?
  - **I/O bound**?
    - Storage (empirical knowledge derived from operators)
      - IaaS: Virtual storage device (**random access**)
      - PaaS: Backend database

Storage Access in IaaS

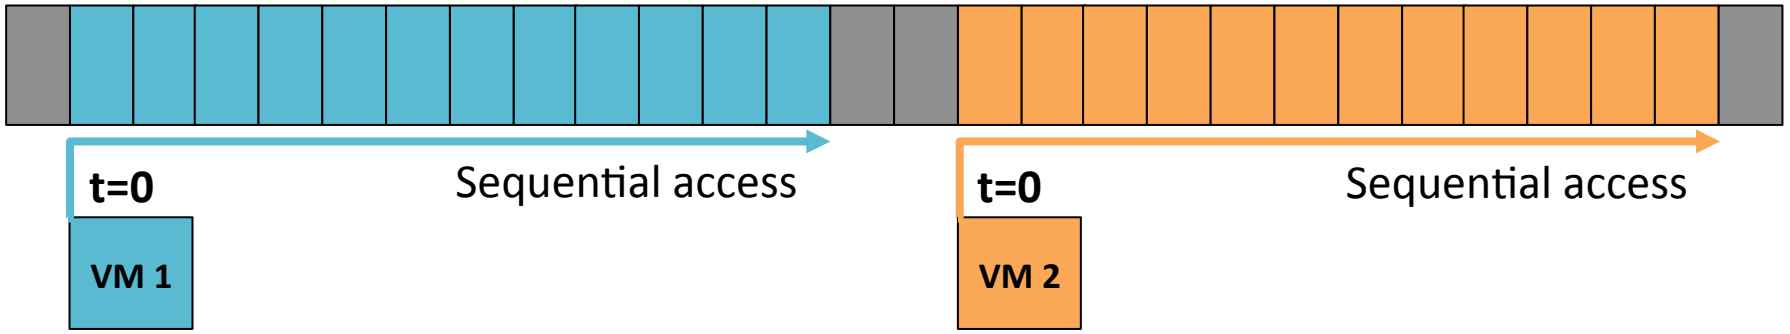

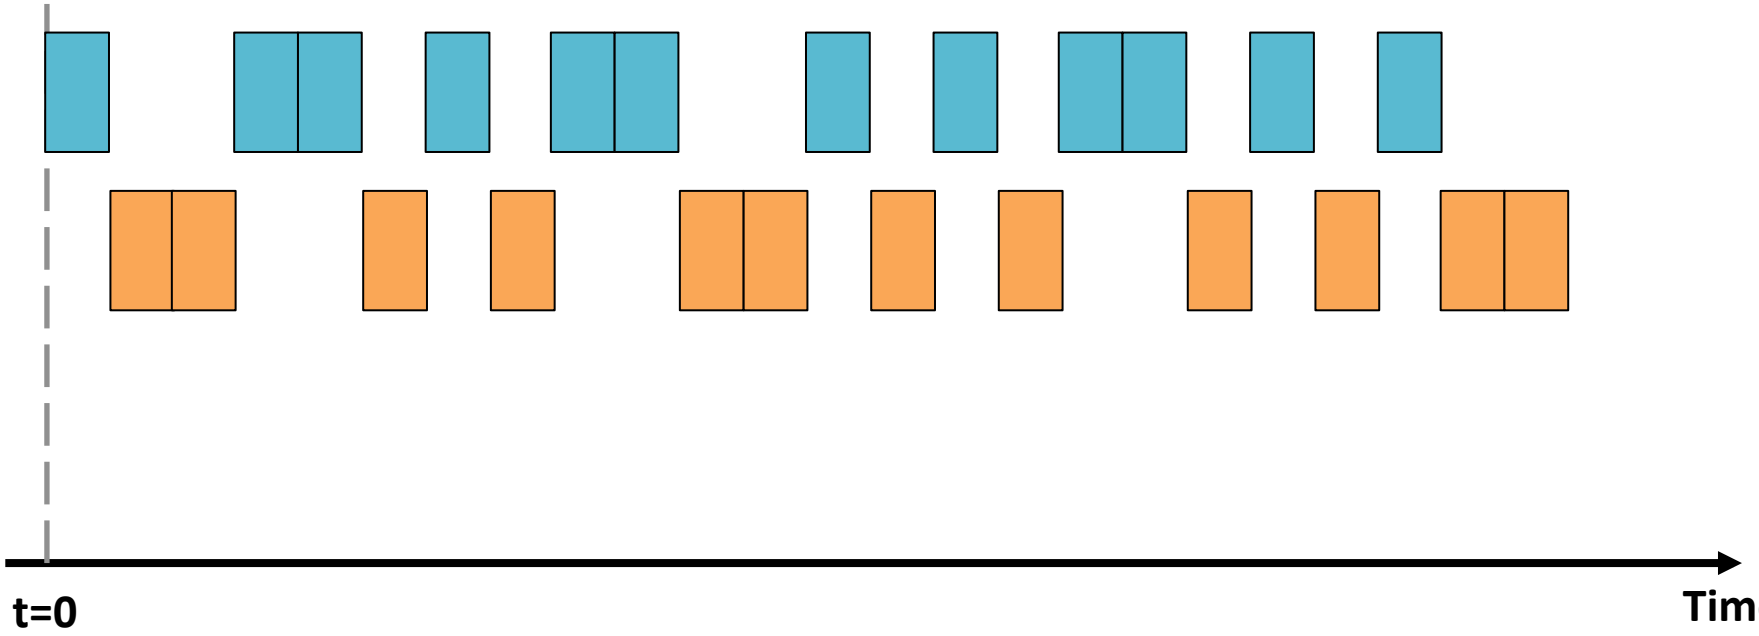

Storage Access in IaaS

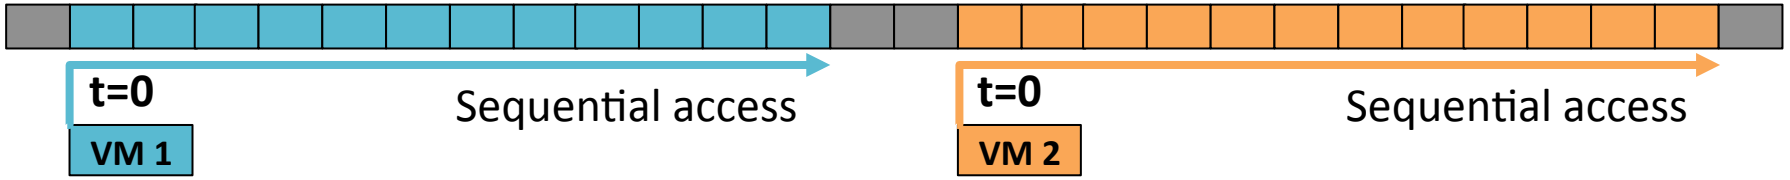

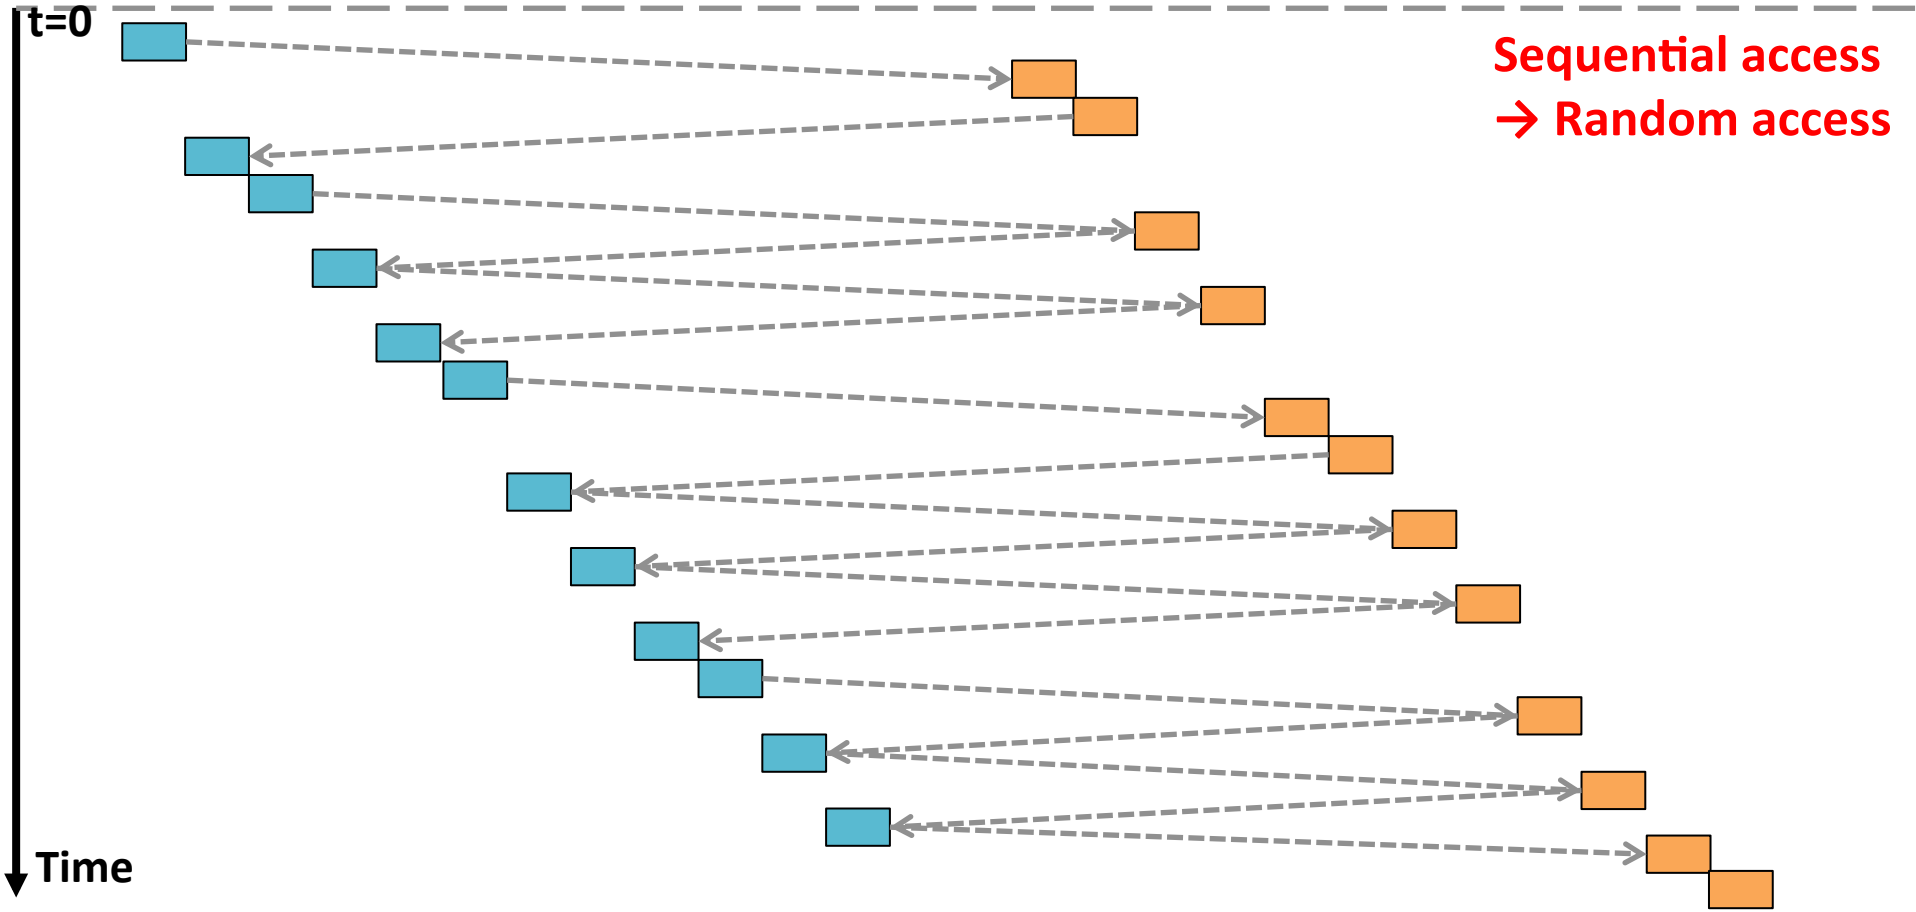

Virtual Storage Device

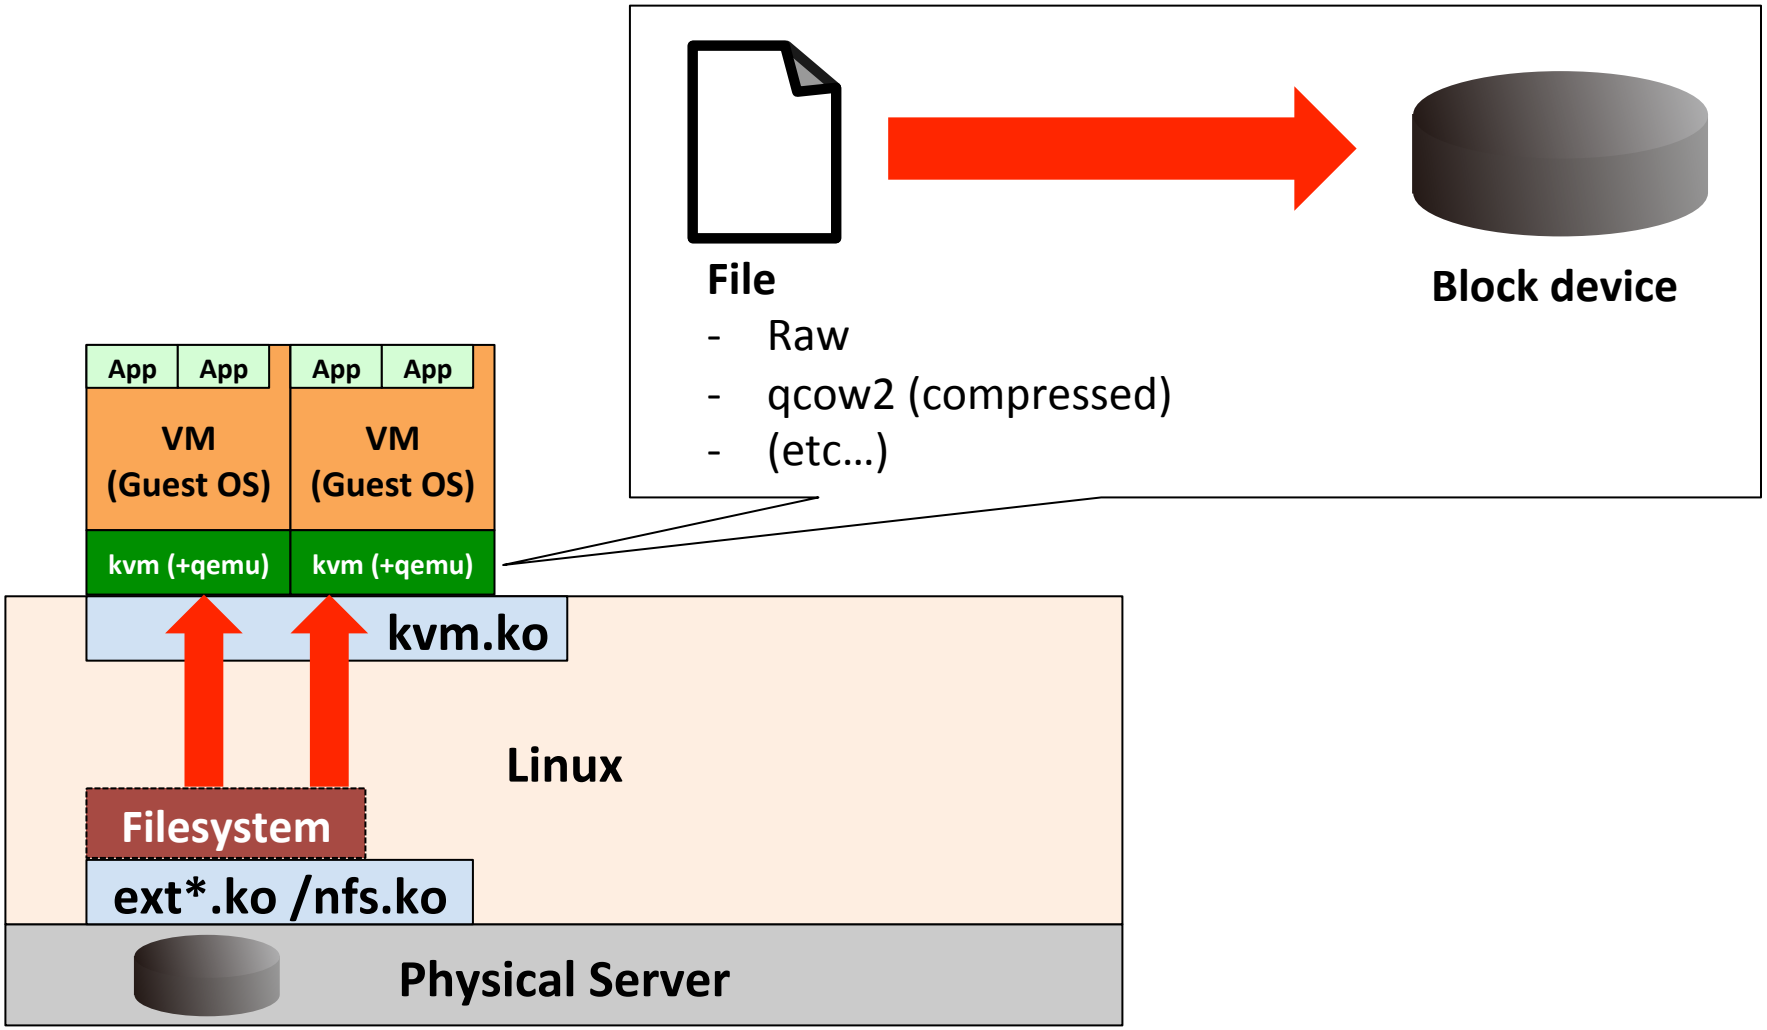

Virtual Storage Device

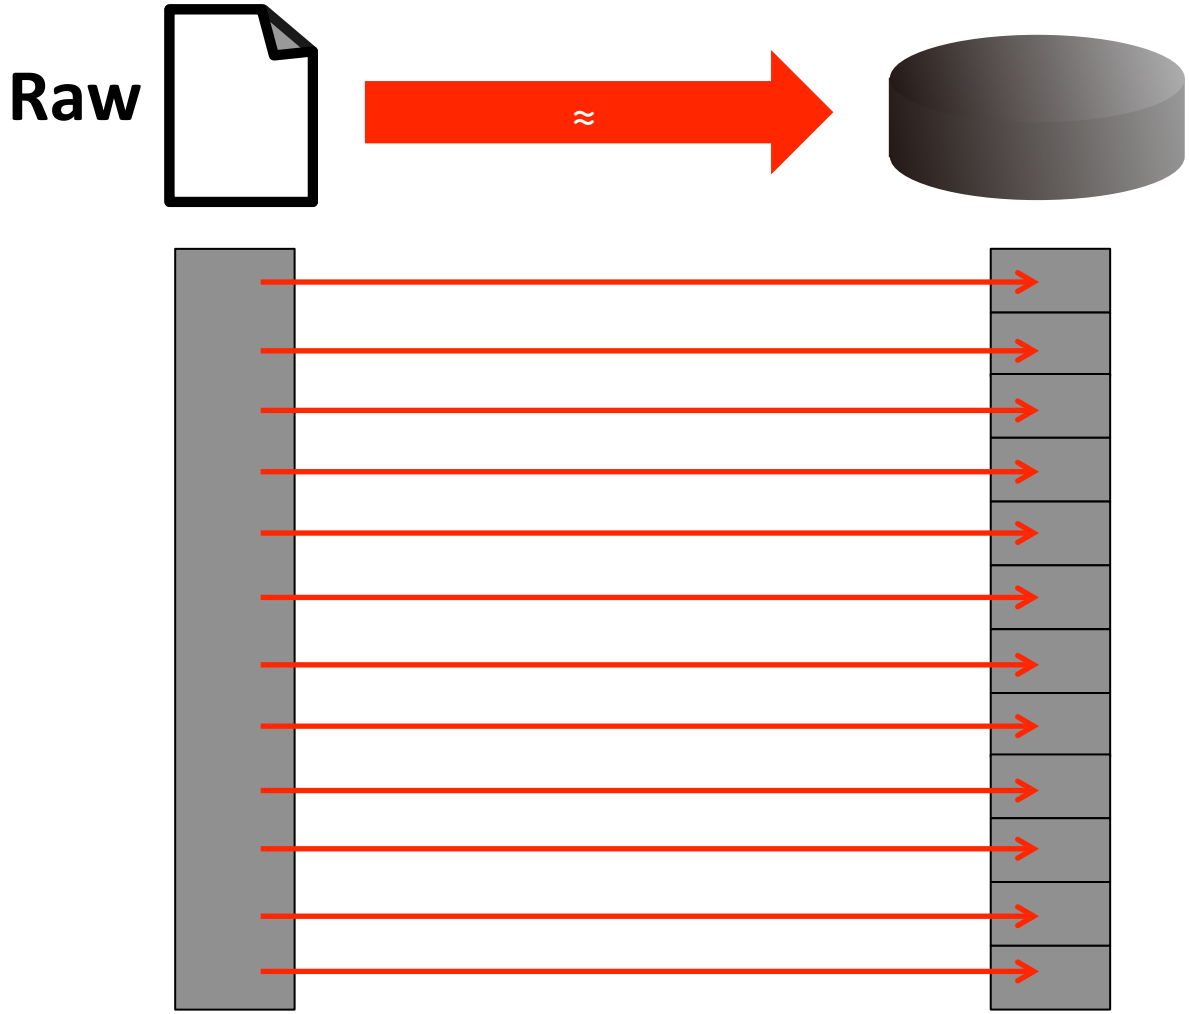

Virtual Storage Device

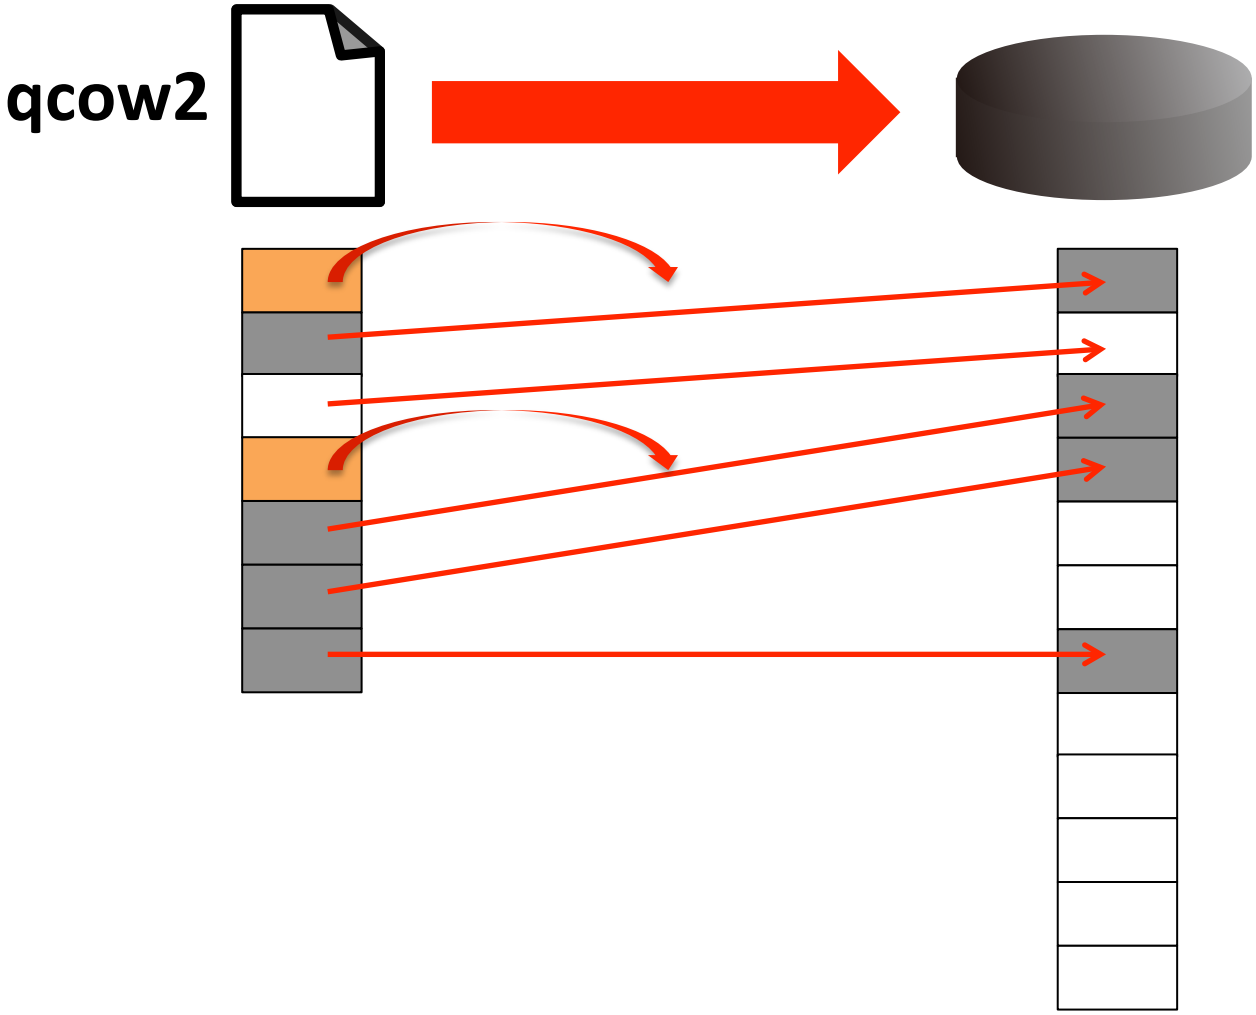

Note: This is NOT the strict image of "qcow2"

Virtual Storage Device (via NFS)

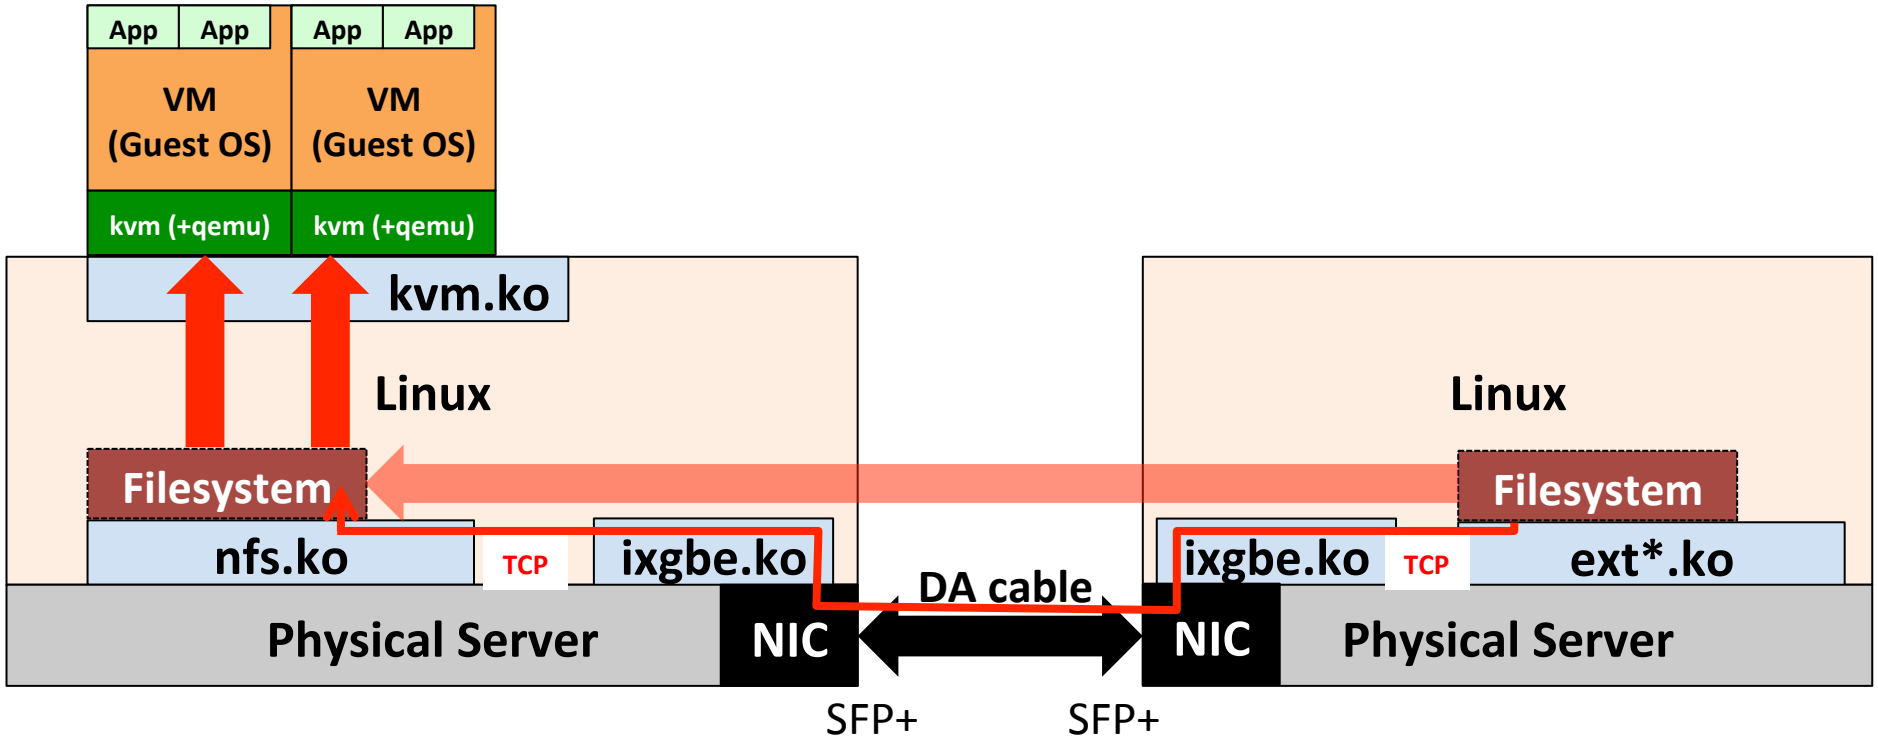

Quick Comparison of Storage Techs

	Storage	Interface	Speed	Note
SATA, AHCI	HDD/SSD	AHCI	Bad	On-board
SATA/SAS, RAID	HDD/SSD	PCIe	??	Not expensive RAID 0: Good speed RAID 5/6: Bad speed
Fusion-io ioDrive	NAND flash	PCIe	Good	Expensive
NVMe	SSD	PCIe	Good?	Not yet widely deployed
RAM disk	RAM	DDR3	Good	Volatile

NVMe: Non Volatile Memory Express

AHCI and RAID with HDD: High latency and low-throughput especially for random access

Tips: RAM disk

Create and mount RAM disk(Linux)

```
# mount -t tmpfs -o size=<size-of-ram-disk> none /path/to/mount
```

(Example) Create a 1024MB RAM disk and mount it to /ramdisk

```
# mount -t tmpfs -o size=1024M none /ramdisk
```

Tips: Storage benchmarking tools

- bonnie++
  - http://www.coker.com.au/bonnie++/
- IOzone
  - http://www.iozone.org
- etc…

Day 1

- Interim goal: Observe bottlenecks
  - Measure throughput and IOPS @ local/remote server
    - AHCI
    - RAID
    - RAM disk
  - Measure throughput and IOPS @ VMs
    - Sequential at multiple VMs ⇔ Random?

Day 1: Exercise 1

- Measure throughput and IOPS (@ localhost)
  - RAM disk
  - SSD over AHCI
  - SSDs over RAID card
    - RAID 0
    - RAID 0+1
    - RAID 1+0
    - RAID 5
    - RAID 6
      - random/sequential access
      - various configuration parameters

Day 1: Exercise 2

- Measure throughput and IOPS through NFS
  - RAM disk
  - SSD over AHCI
  - SSDs over RAID card
    - RAID 0
    - RAID 0+1
    - RAID 1+0
    - RAID 5
    - RAID 6
      - random/sequential access
      - various configuration parameters

Day 1: Exercise 2

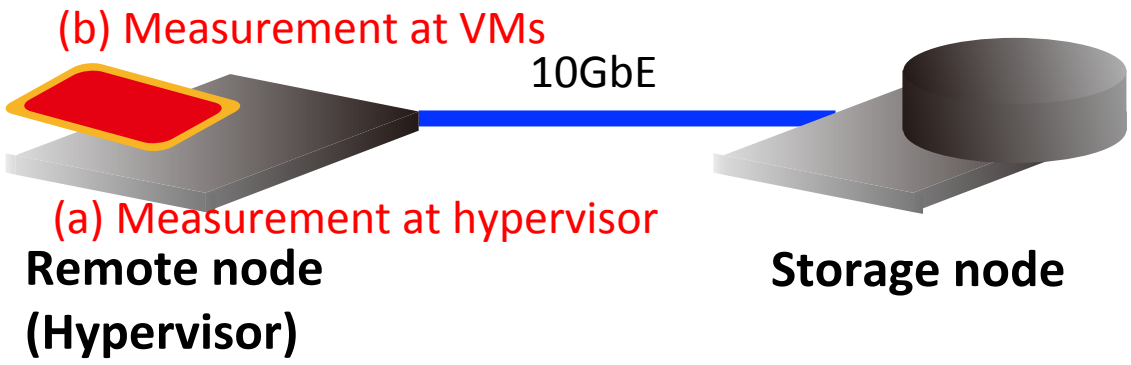

NFS tips: Export a directory / Mount the exported directory

@Storage node

```
# apt-get install nfs-kernel-server
# echo “/path/to/storage *(rw,async,no_root_squash,no_subtree_check)” >> /
etc/exports
# /etc/init.d/nfs-kernel-server reload
```

Note: Replace * with your subnet (e.g., 192.0.2.0/24) to restrict access for security

@Remote node

```
# mkdir -p /var/vmpool
# mount -t nfs storage-node:/path/to/storage /var/vmpool
```

※ KVM installation: Look at the appendix slides (May 2).

Day 2-3: Additional exercises

- Keywords
  - work on anywhere you want
  - Fast network
    - netmap
  - Cache technology (Random access → Sequential access)
    - dm-cache / SSD cache
  - Distributed filesystem
    - GlusterFS
      - http://www.gluster.org
  - Distributed block device
    - sheepdog
      - http://sheepdog.github.io/sheepdog/
  - Userland block device
    - Hack NBD (network block device) client