



CompSci 401: Cloud Computing

Storage Interfaces

Prof. Ítalo Cunha



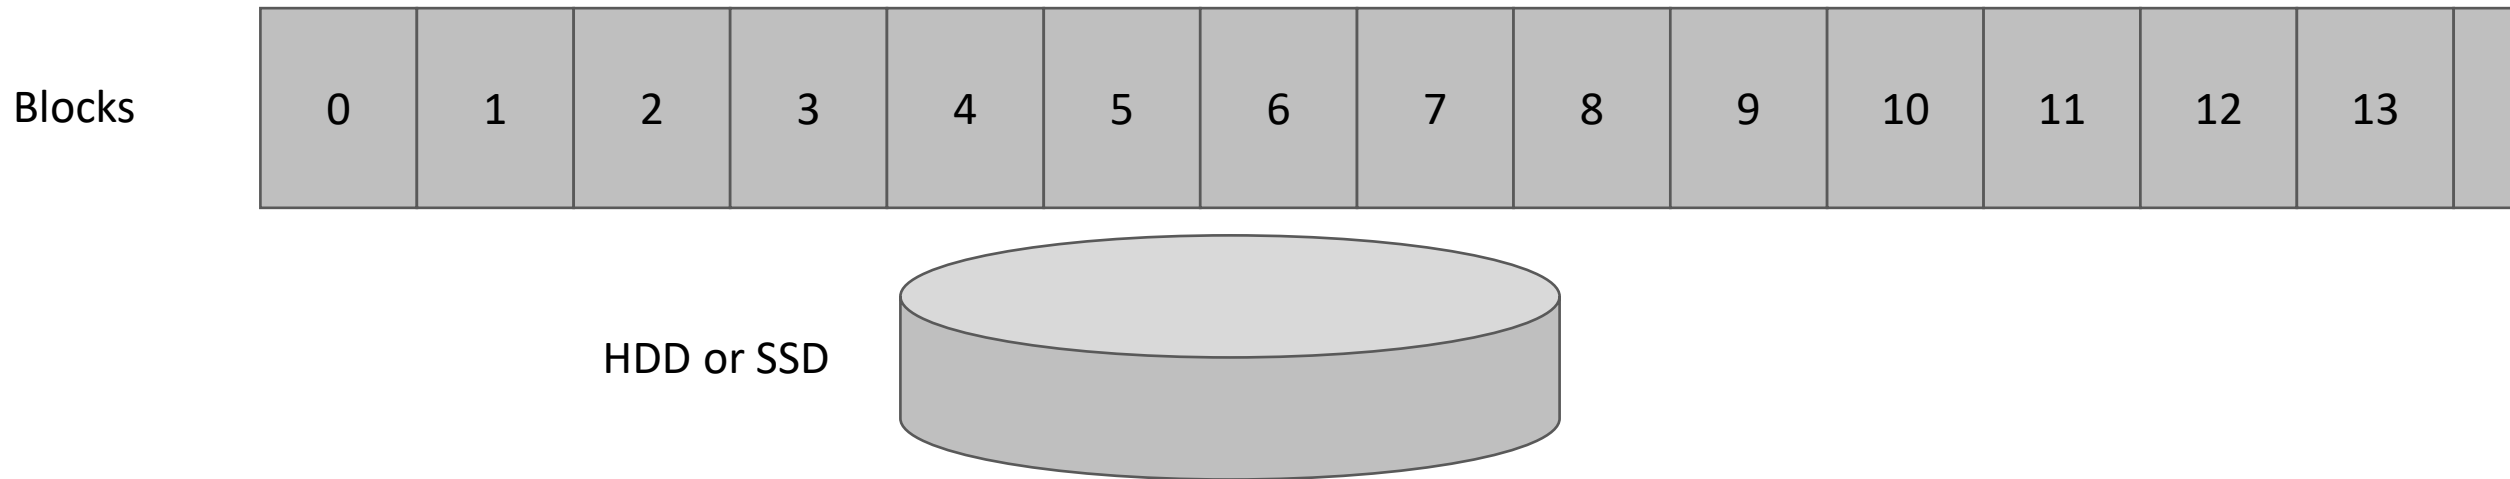
The disk interface

- Hard disk drives (HDDs) and solid state drives (SSDs)
- Block-oriented interface
 - Blocks have a fixed size like 512 or 4096 bytes
 - Blocks are numbered
 - Writes and reads operate on whole blocks
 - Cannot read or write only part of a block
 - To change one byte in an existing block we read the entire block, and then write the whole block with the modified byte

The file interface

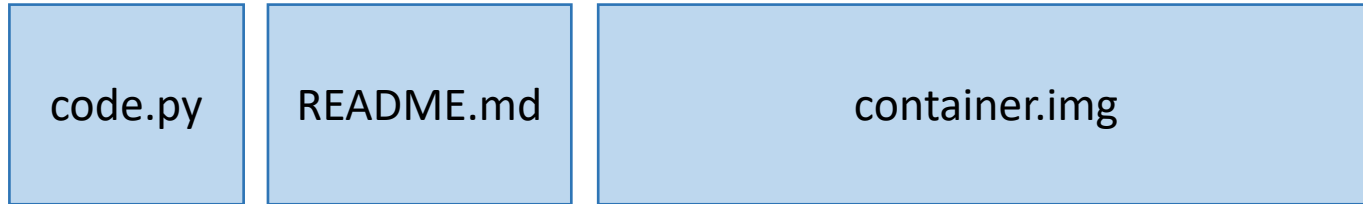
- Operating system contain file systems
 - Windows: NTFS and FAT
 - MacOS: APFS, HFS+
 - Linux: ext4, reiserfs, XFS, btrfs, ...
- File systems provide a higher-level interface do disk blocks
- File operations
 - Open/close
 - Read/write
 - Seek → Change the position in the file being read from or written to
 - Metadata → Directories, names, owner, access permissions, creation date

The disk and file interfaces

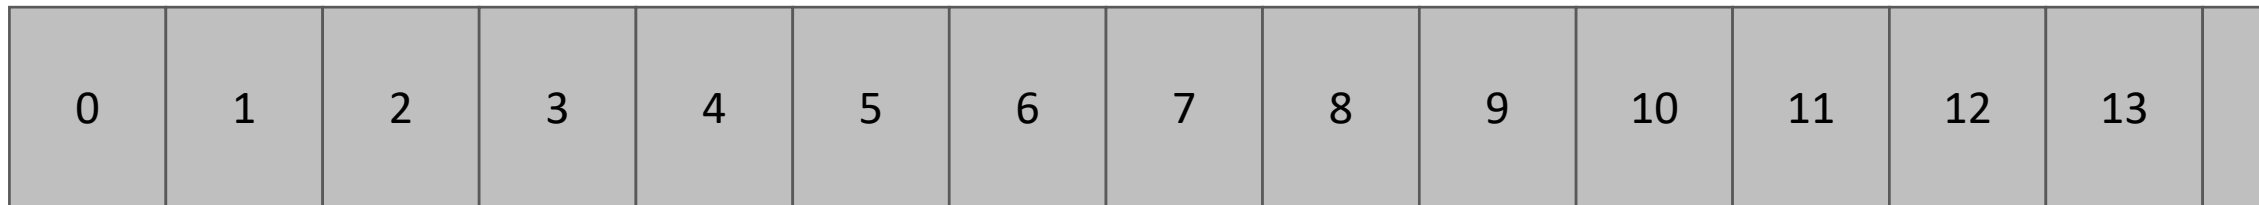


The disk and file interfaces

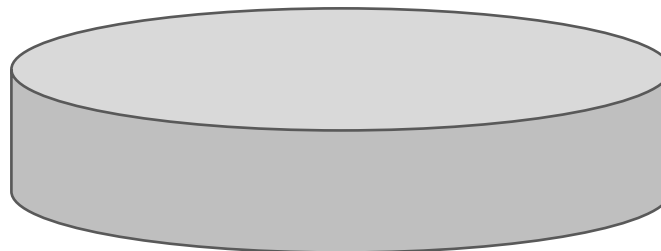
Files
(logical)



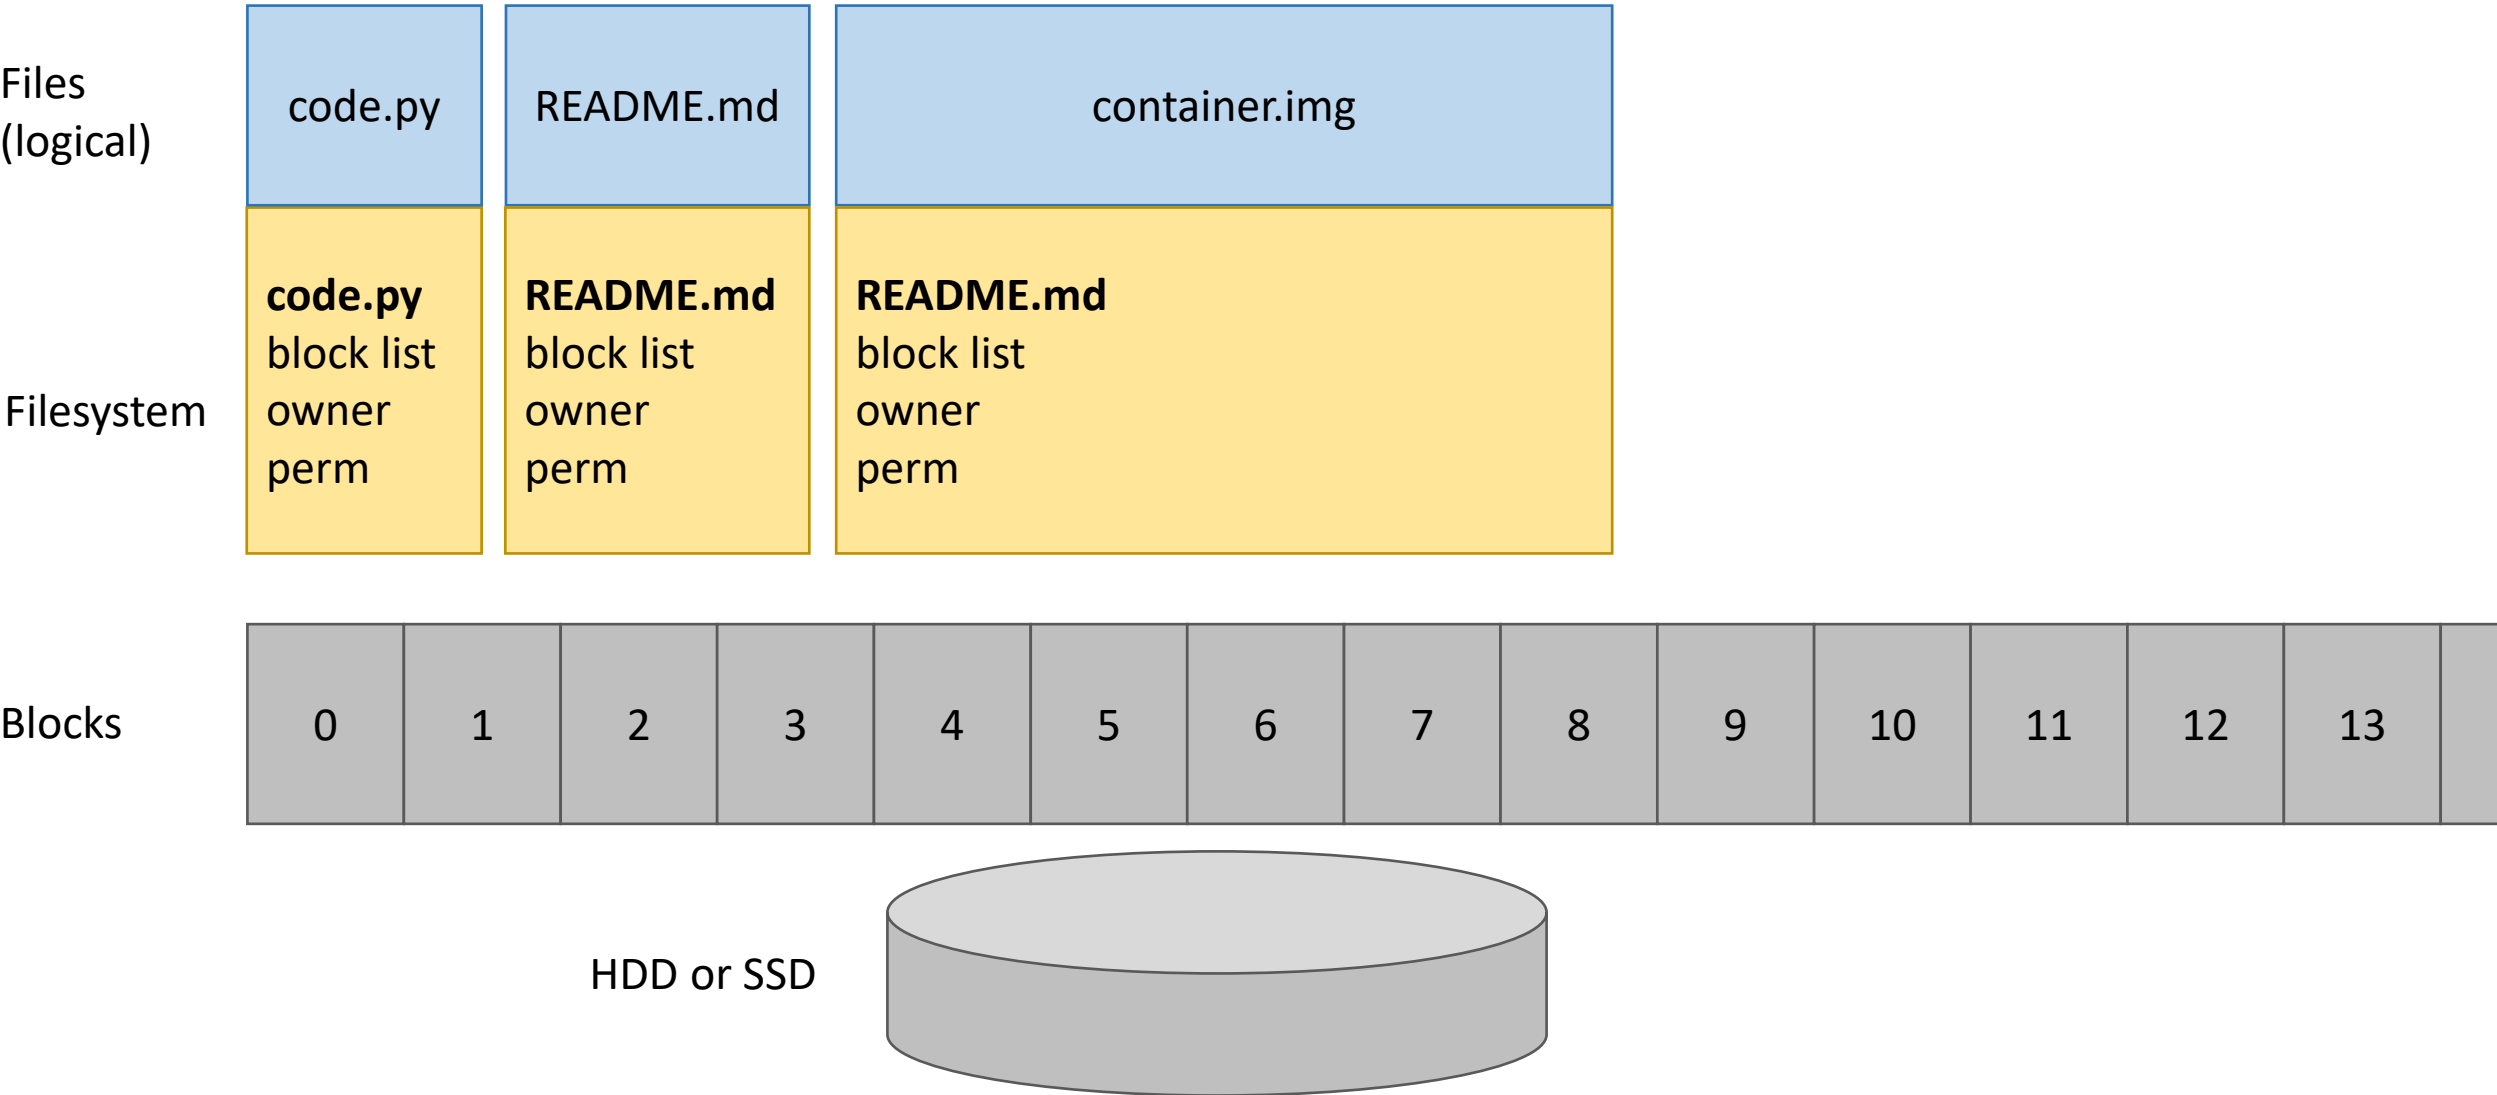
Blocks



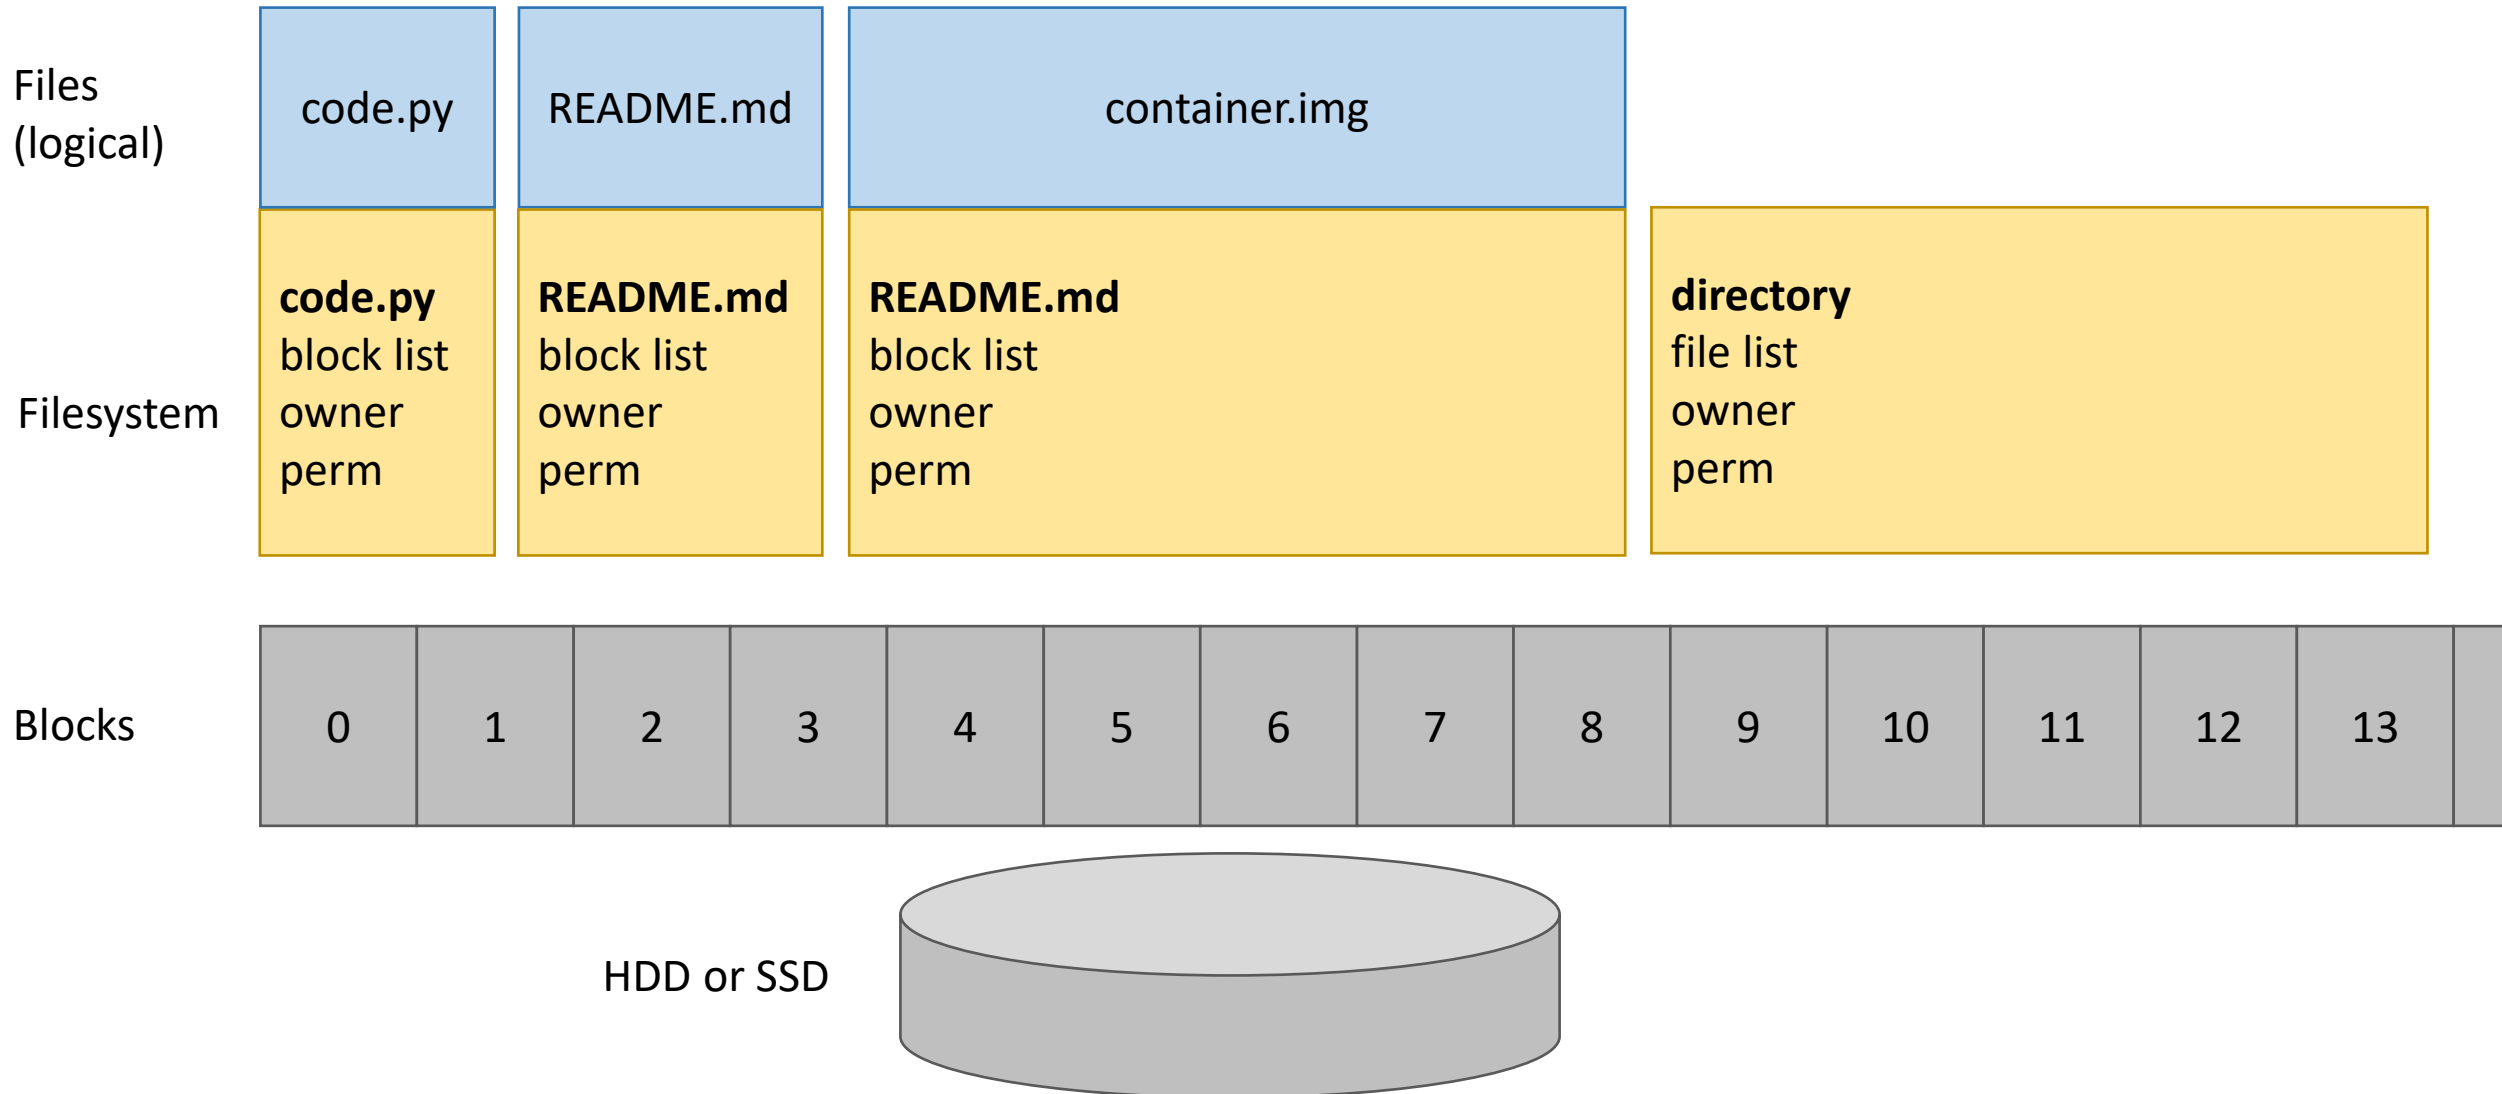
HDD or SSD



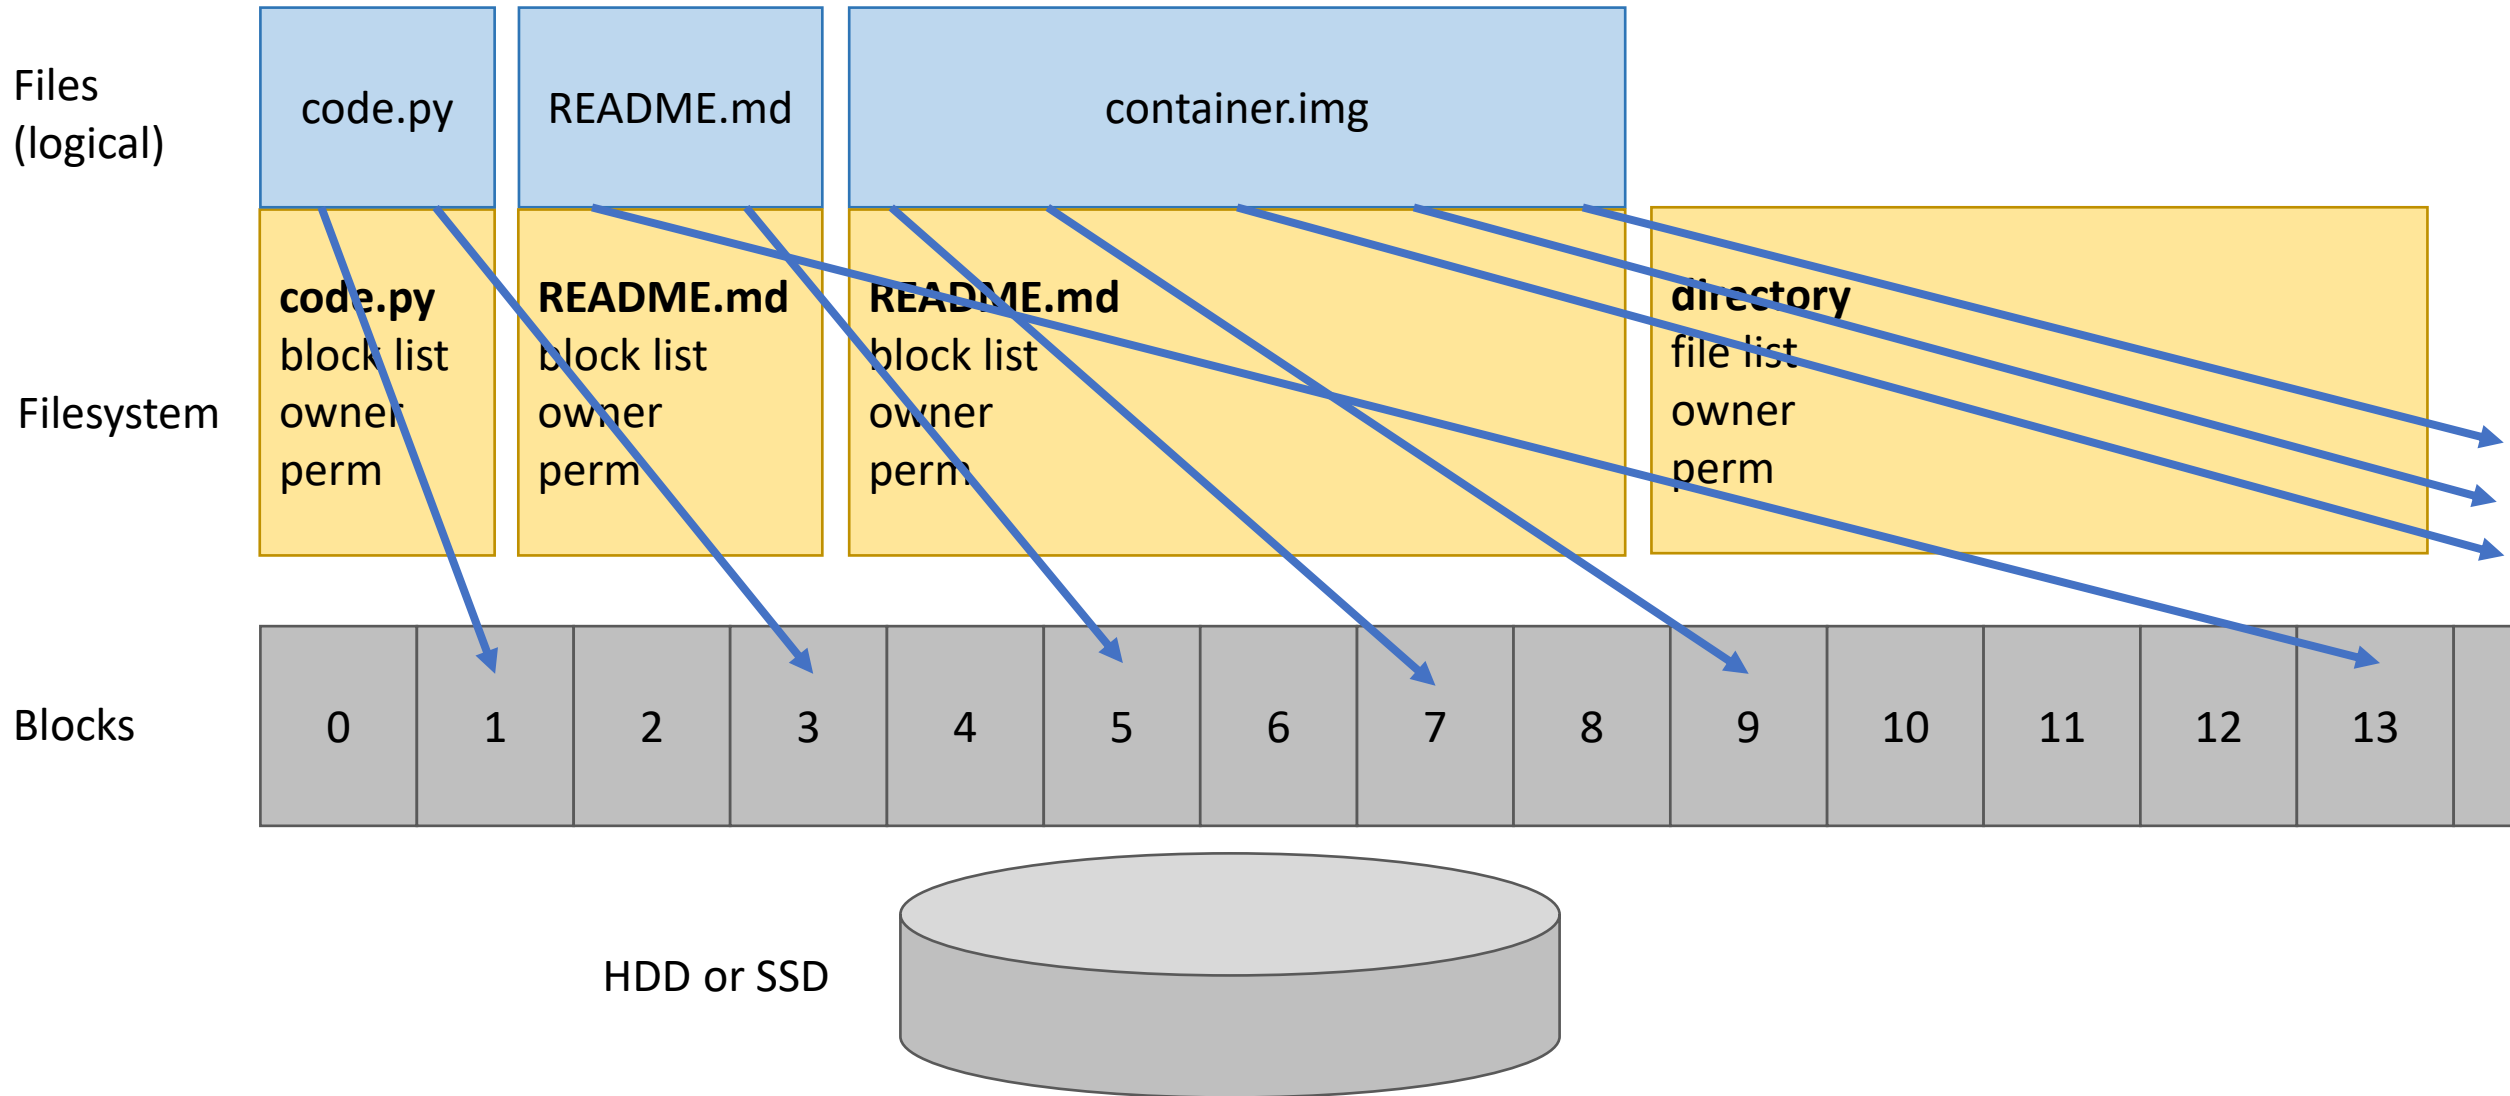
The disk and file interfaces



The disk and file interfaces



The disk and file interfaces





CompSci 401: Cloud Computing

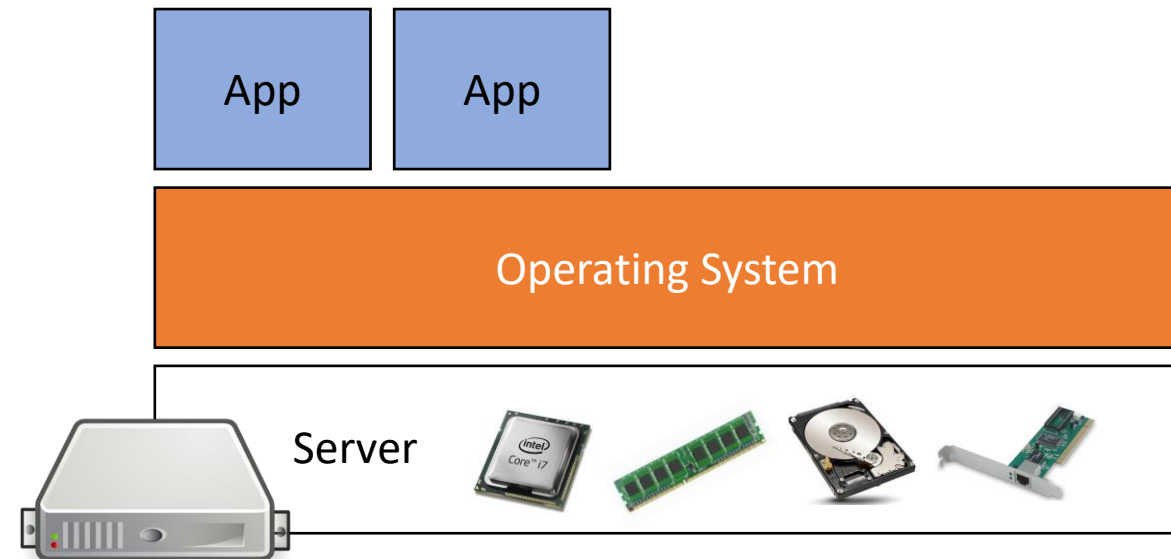
Remote Storage

Prof. Ítalo Cunha



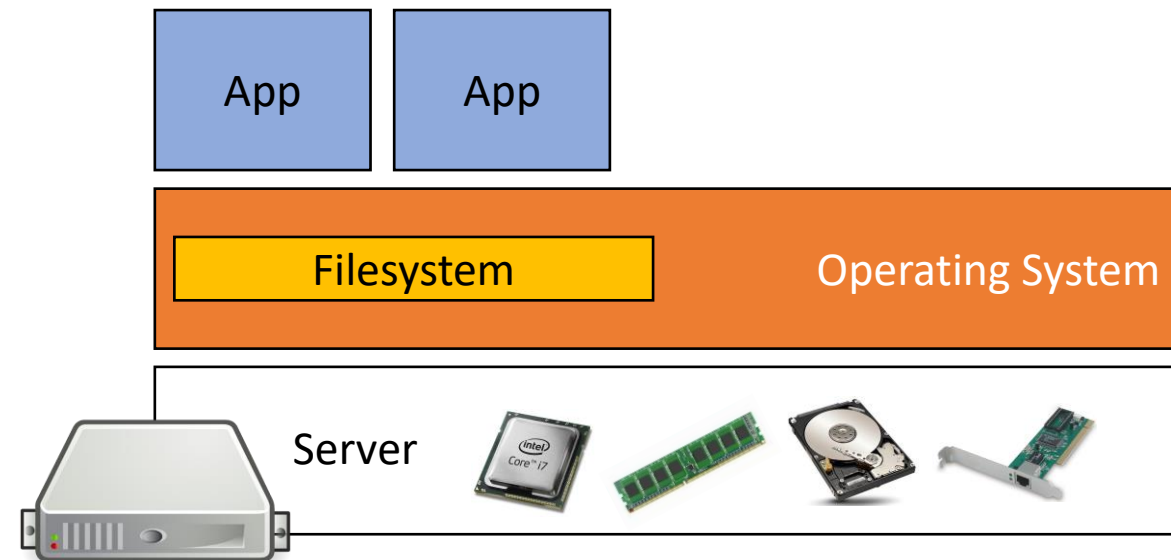
Use of storage interfaces

- Server has resources, including storage
- Operating runs applications and mediates access to hardware
- Applications access storage through files



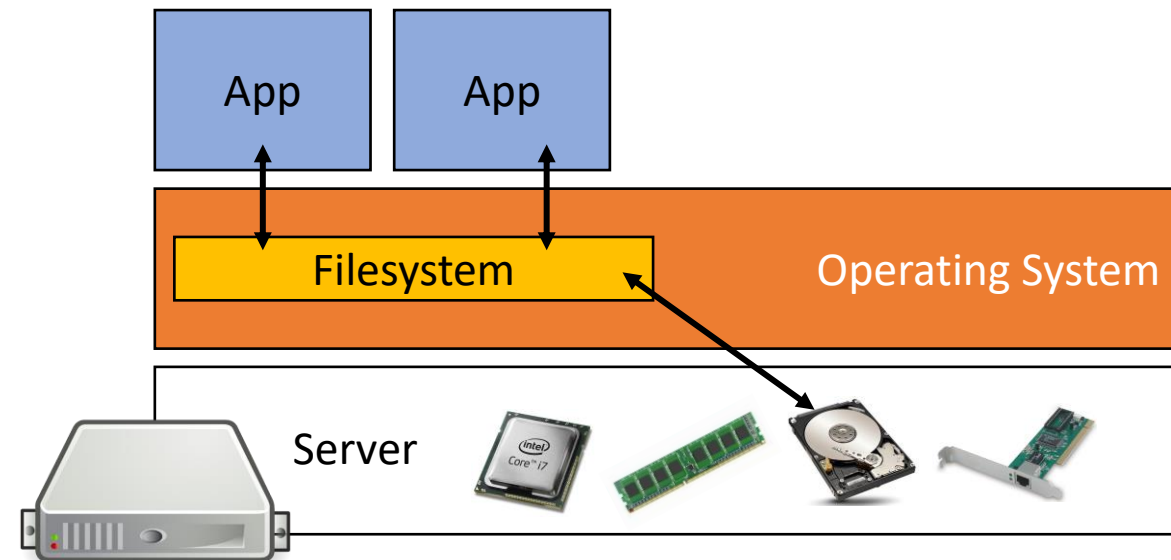
Use of storage interfaces

- Server has resources, including storage
- Operating runs applications and mediates access to hardware
- Applications access storage through files in a filesystem



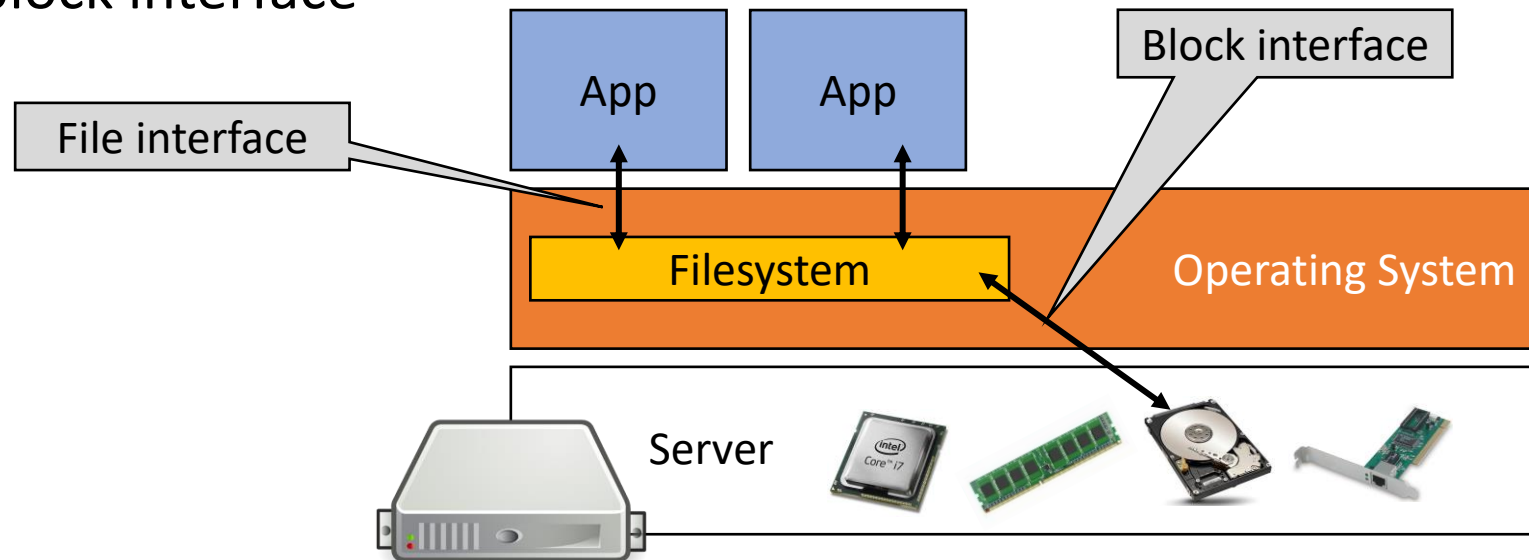
Use of storage interfaces

- Server has resources, including storage
- Operating runs applications and mediates access to hardware
- Applications access storage through files in a filesystem



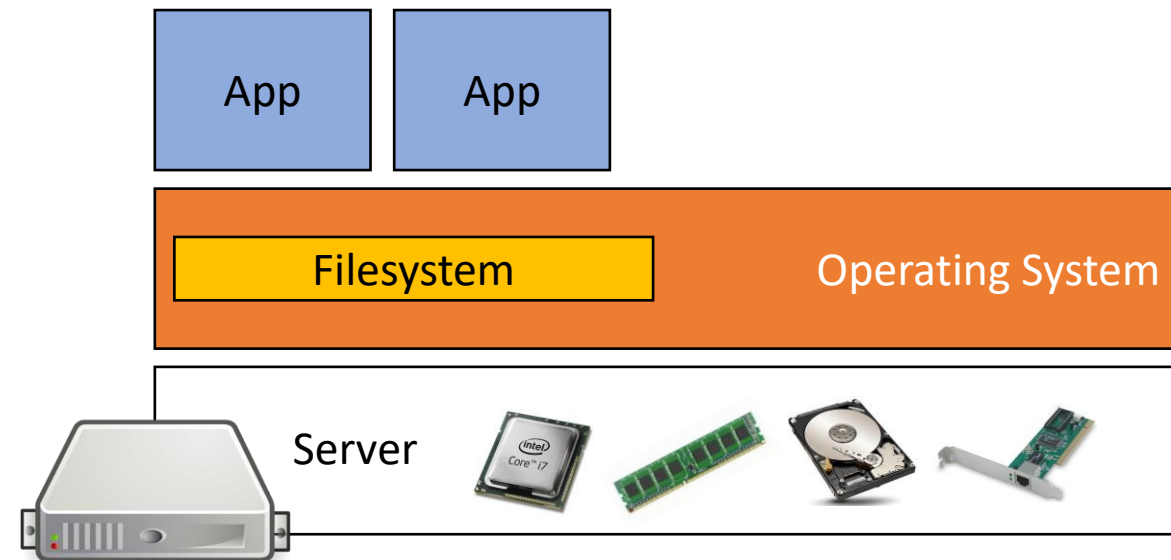
Use of storage interfaces

- Server has resources, including storage
- Operating runs applications and mediates access to hardware
- Applications access storage through files in a filesystem
 - Applications use the filesystem's interface
 - Filesystem uses the disk block interface



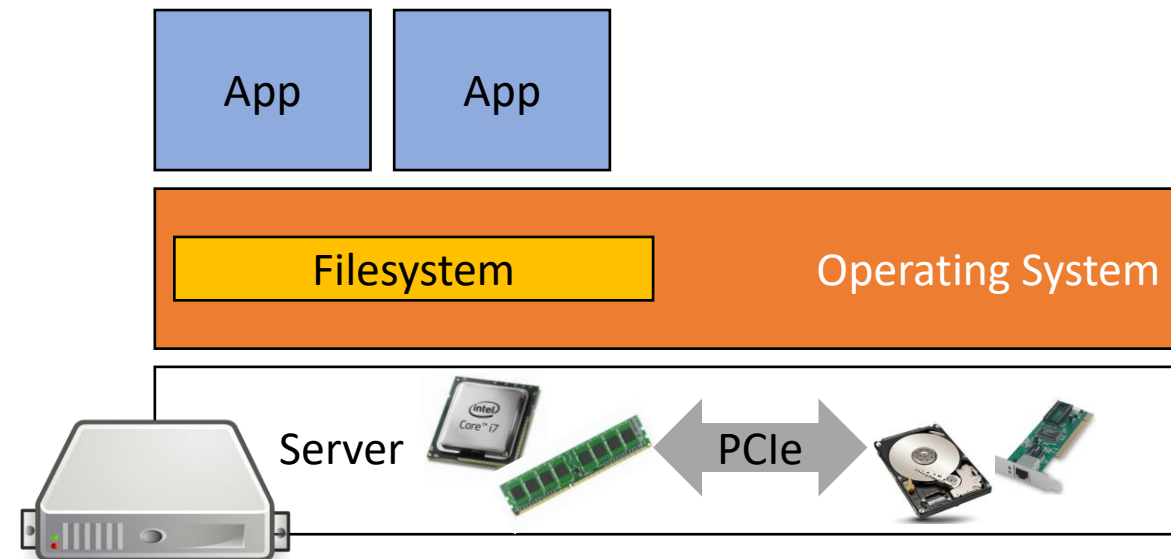
Local and remote storage

- Storage can be local or remote
- **Local storage** is accessed through a bus (like PCI Express)
 - Also called **directly attached storage**



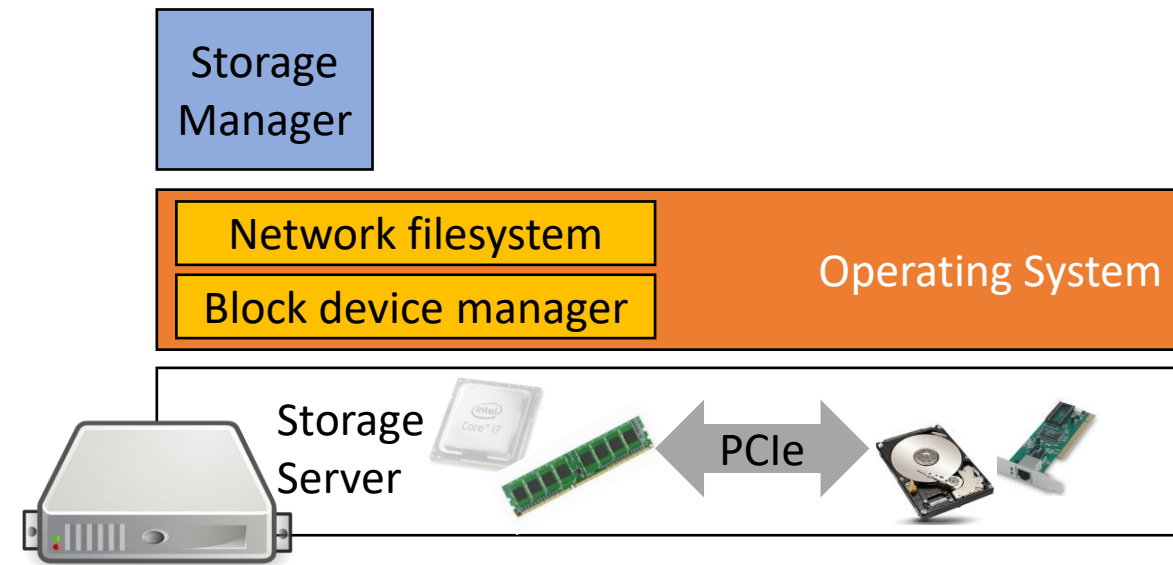
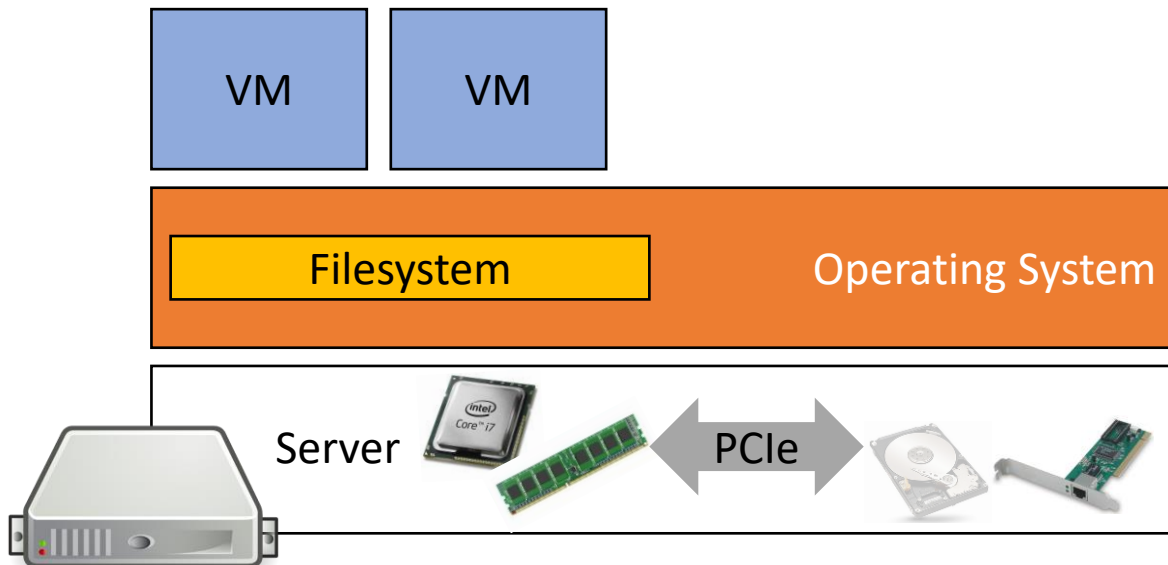
Local and remote storage

- Storage can be local or remote
- **Local storage** is accessed through a bus (like PCI Express)
 - Also called **directly attached storage**



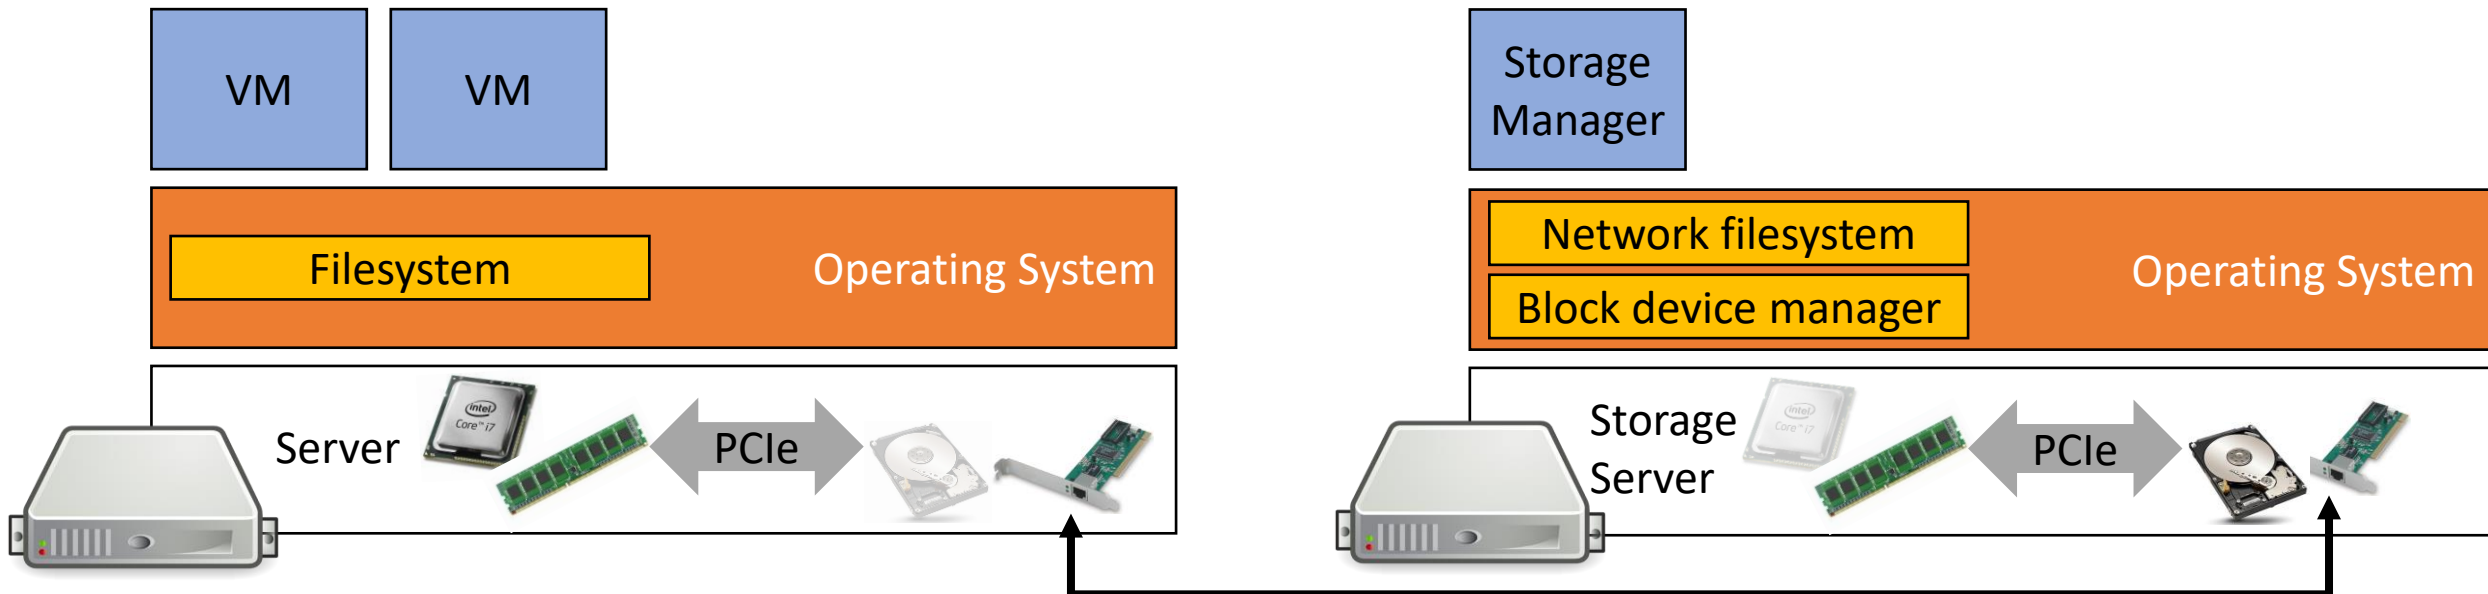
Local and remote storage

- Storage can be local or remote
- **Local storage** is accessed through a bus (like PCI Express)
 - Also called **directly attached storage**
- **Remote storage** is storage accessed on another server



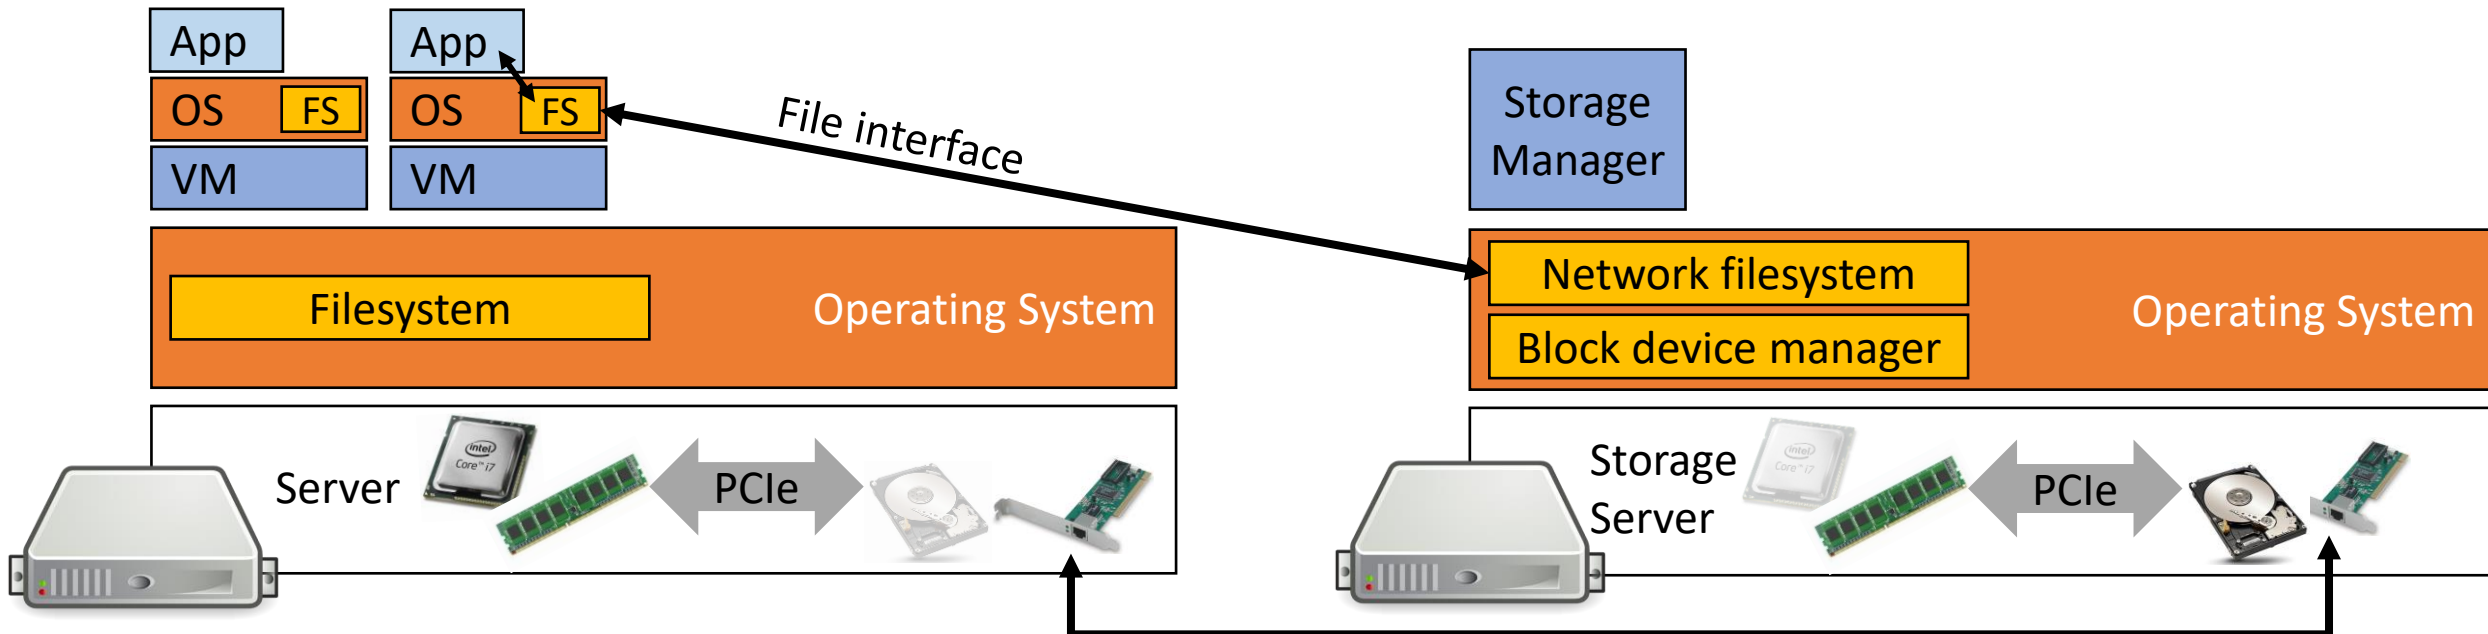
Local and remote storage

- Storage can be local or remote
- **Local storage** is accessed through a bus (like PCI Express)
 - Also called **directly attached storage**
- **Remote storage** is storage accessed on another server



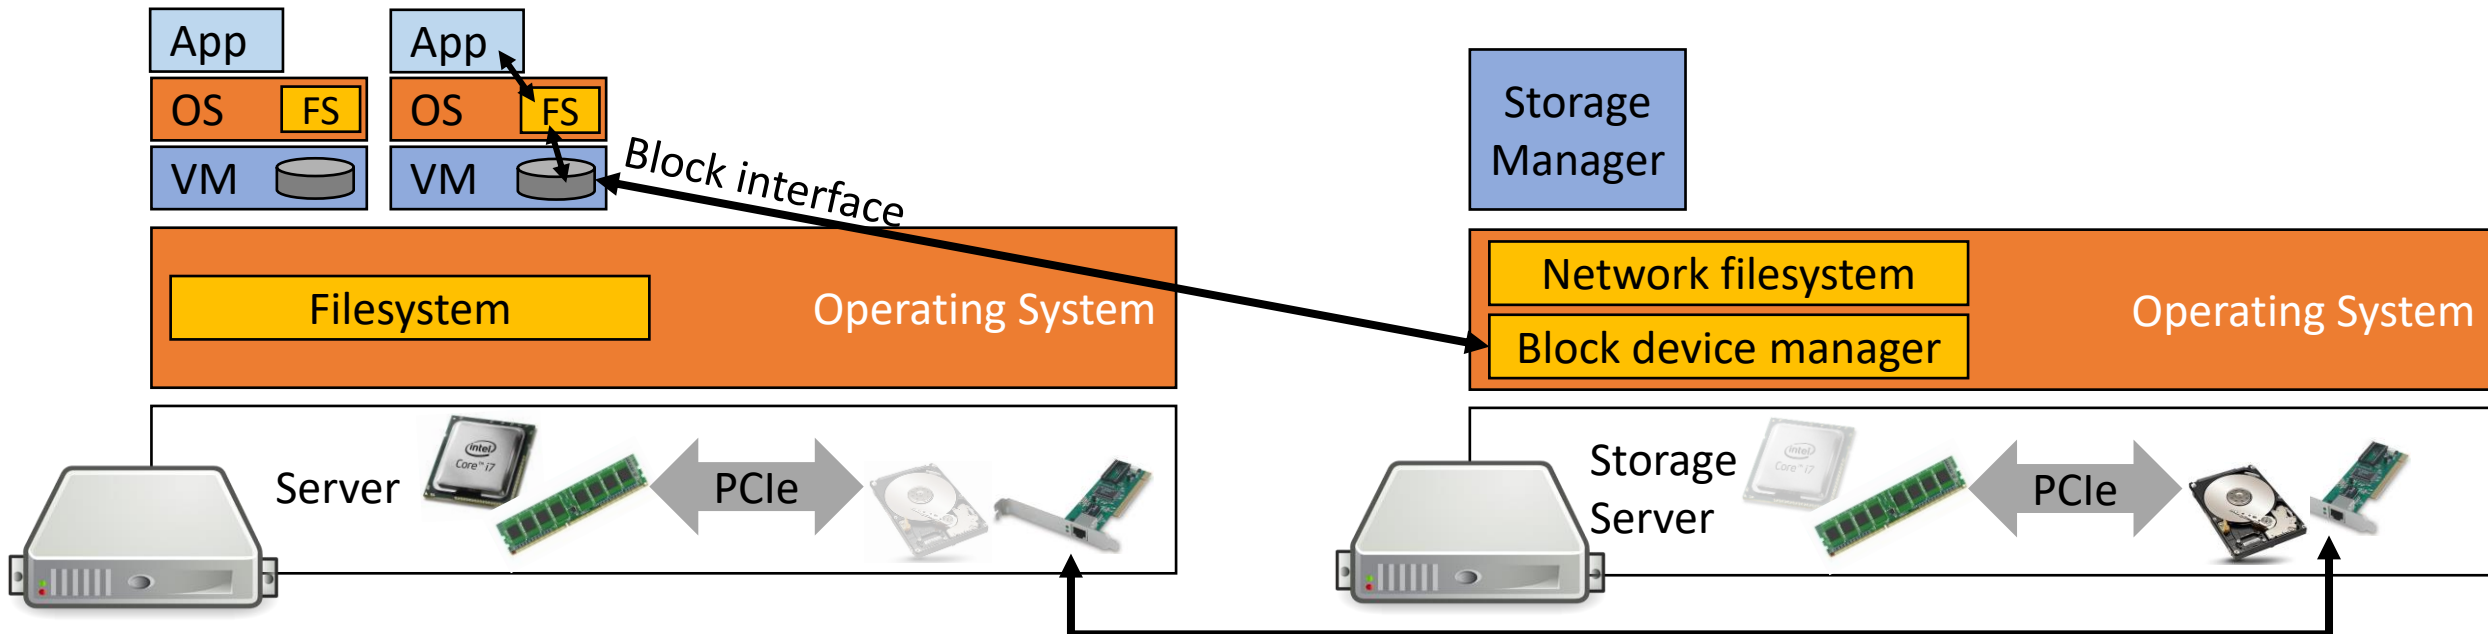
Remote storage interfaces

- Remote storage can be made available at the file or block interfaces

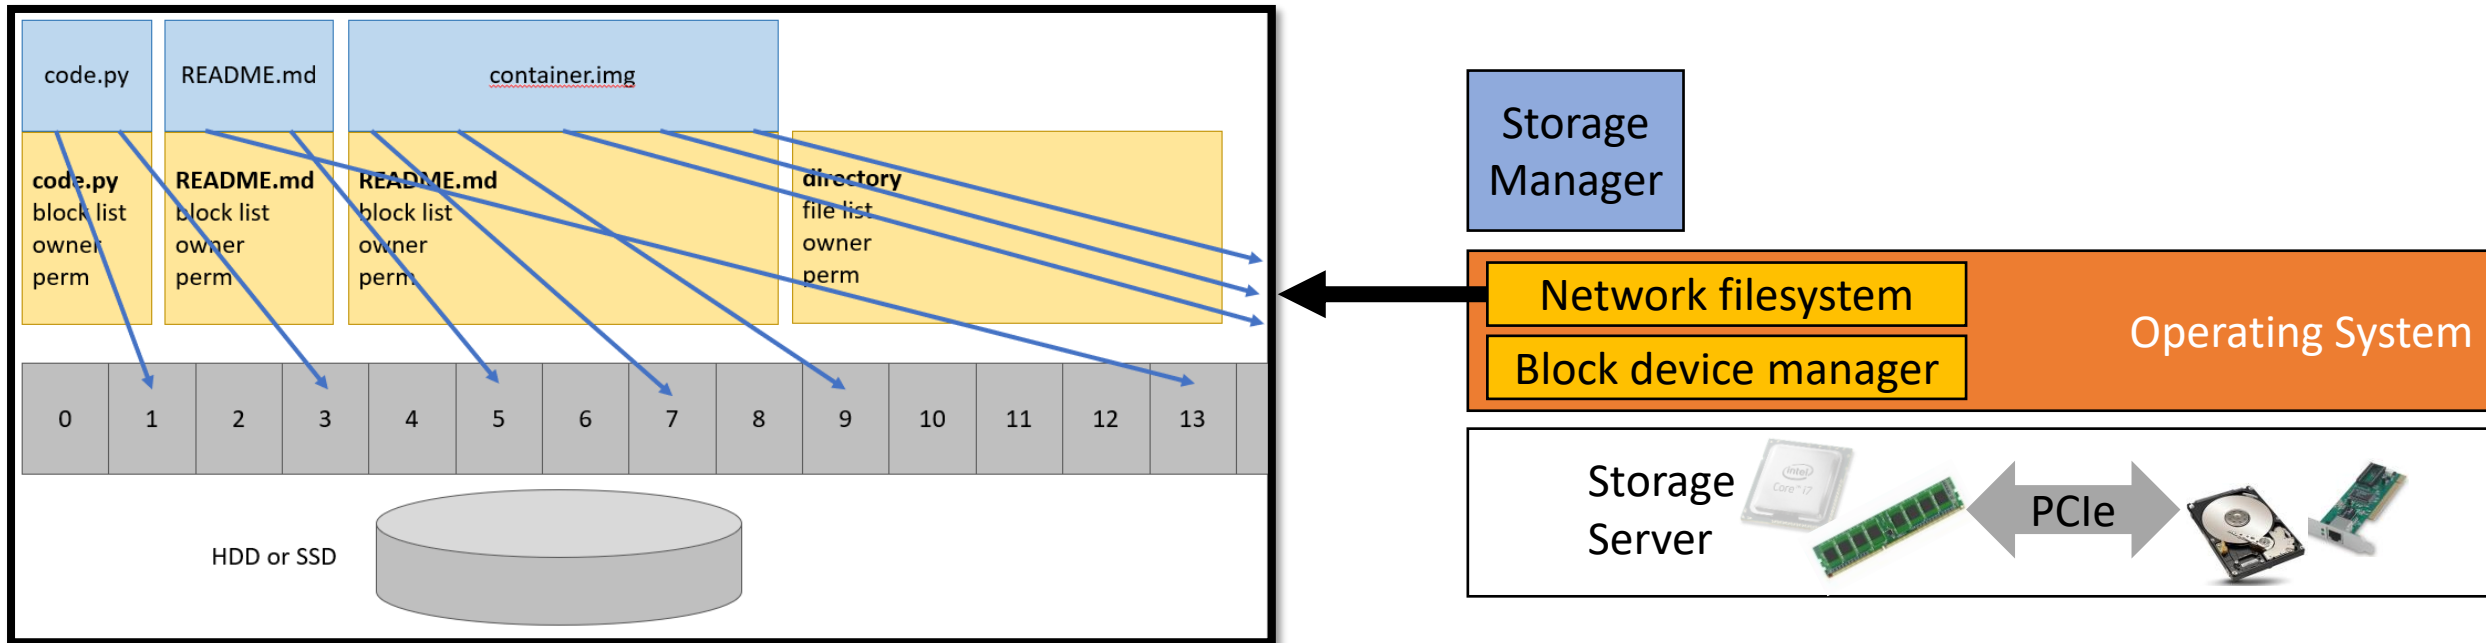


Remote storage interfaces

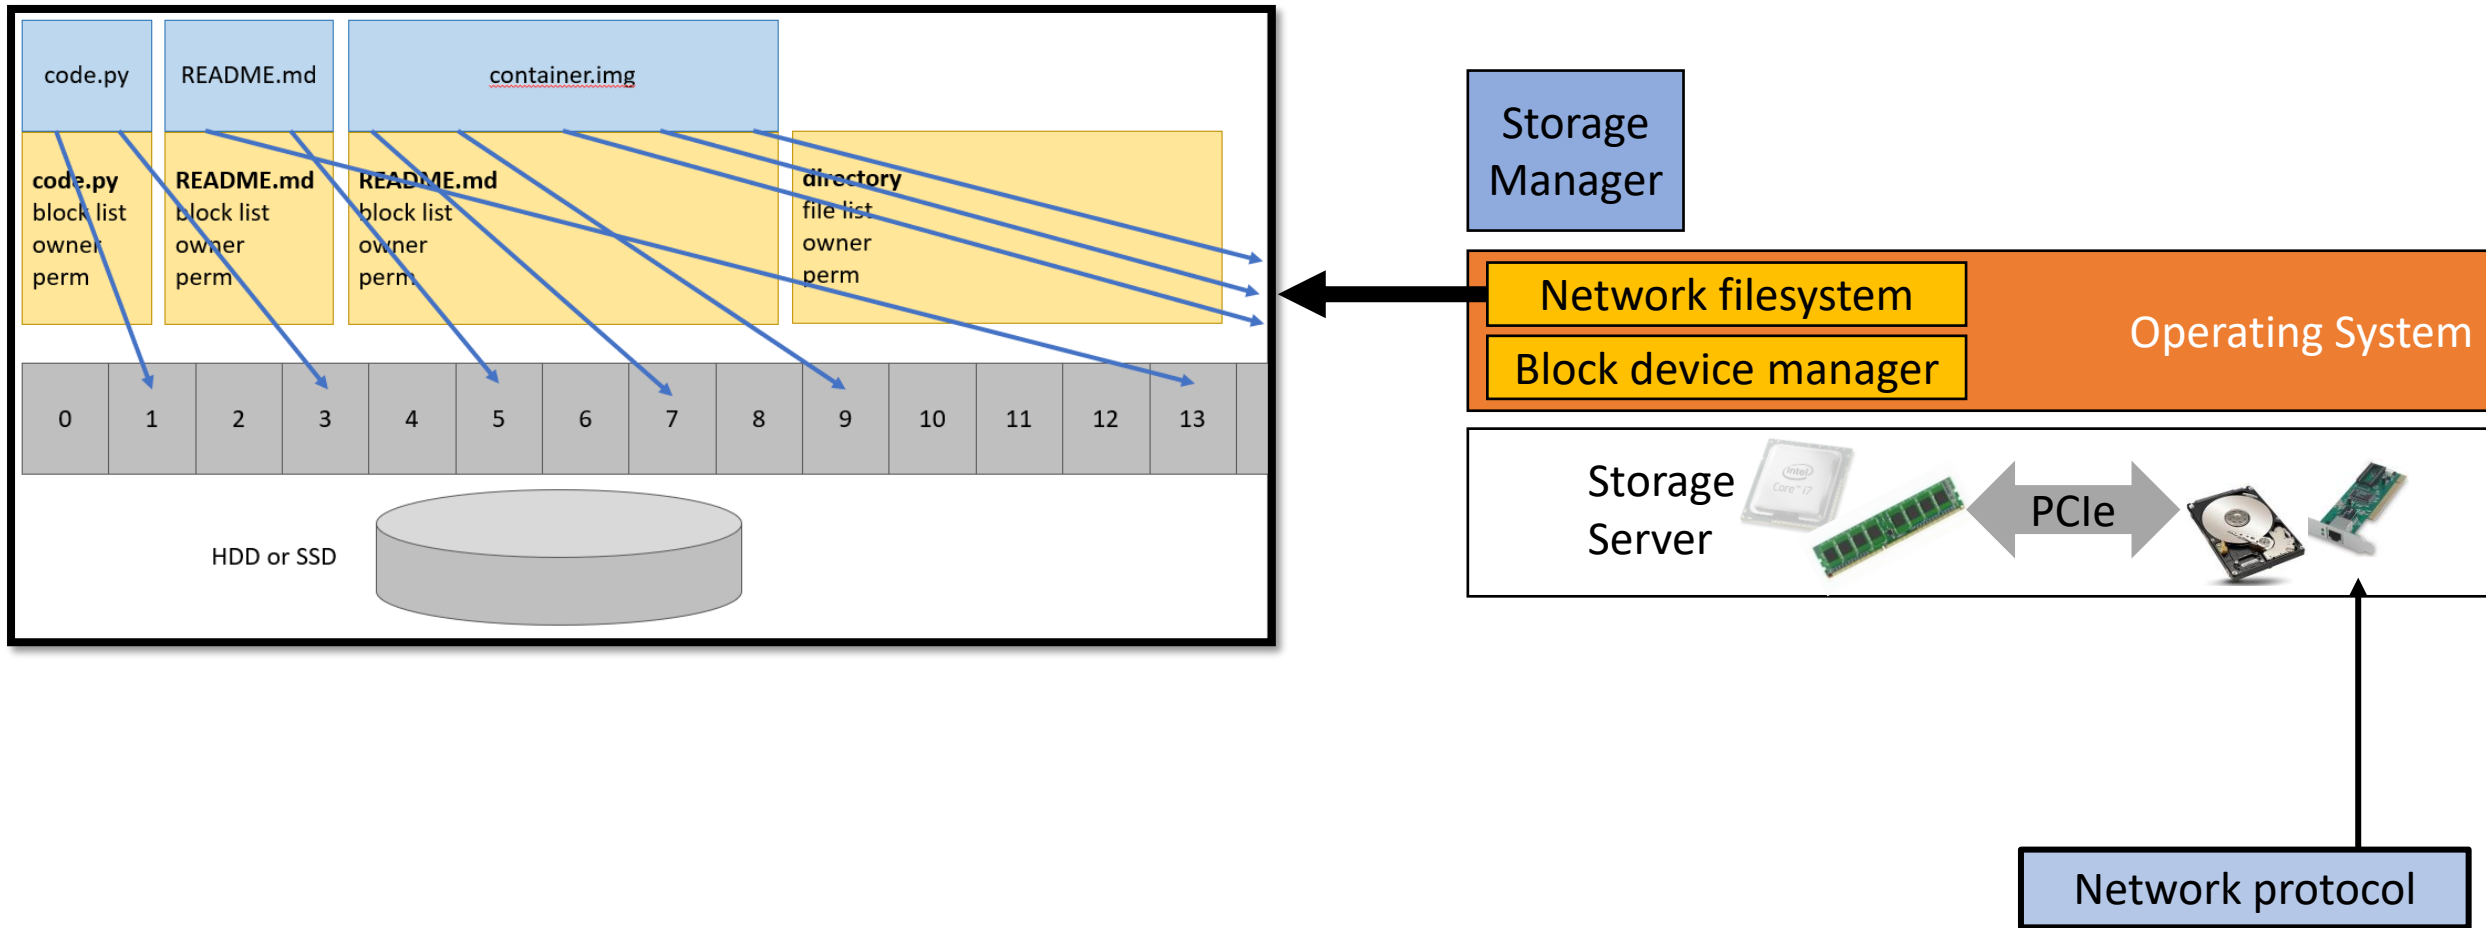
- Remote storage can be made available at the file or block interfaces



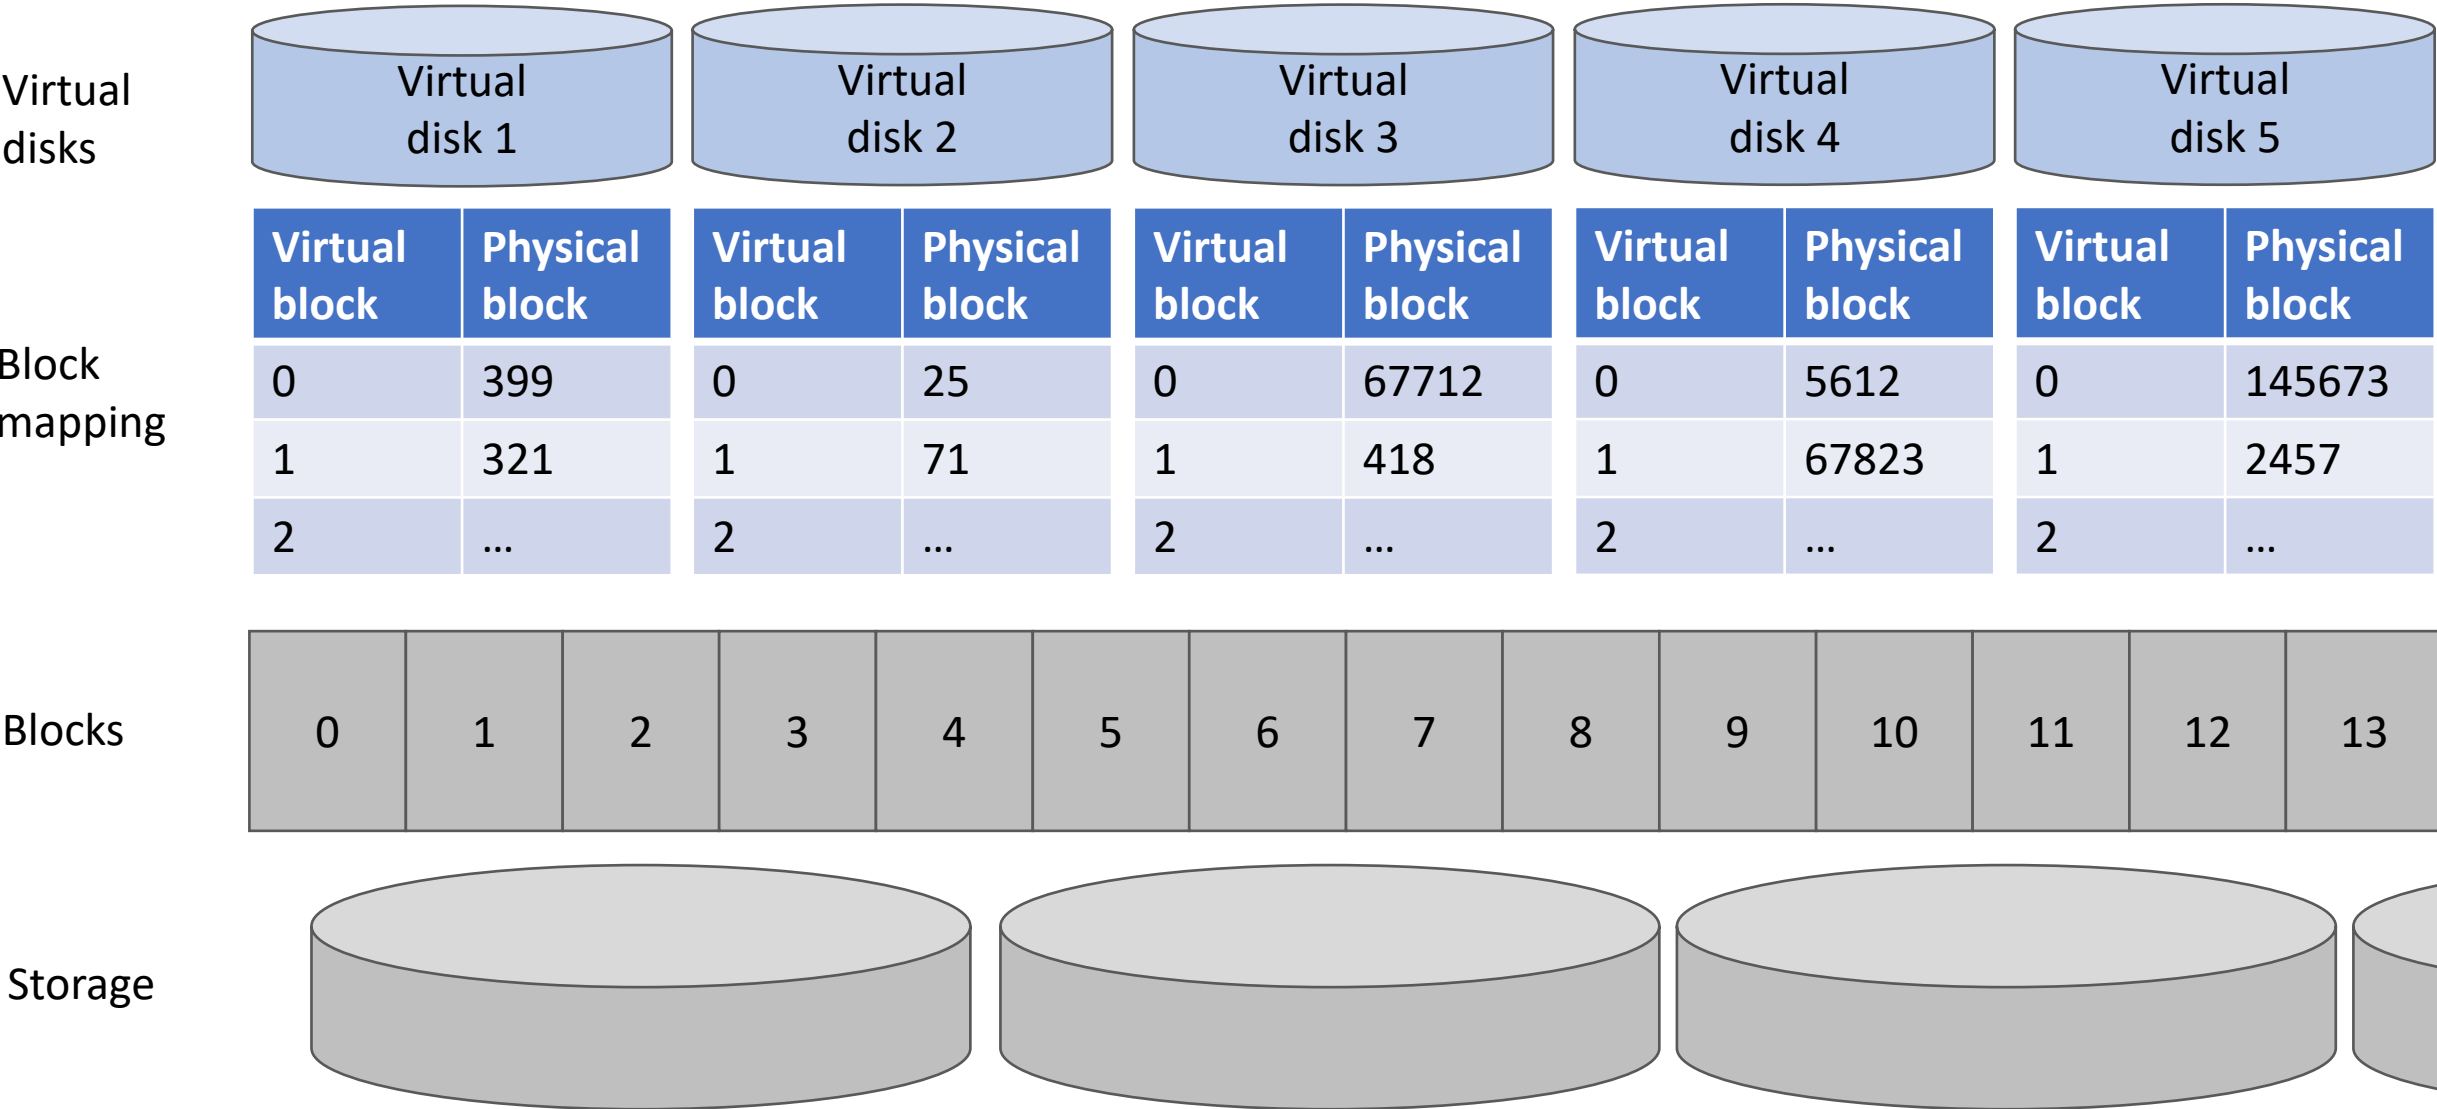
Remote storage implementation



Remote storage implementation



Remote block storage implementation





CompSci 401: Cloud Computing

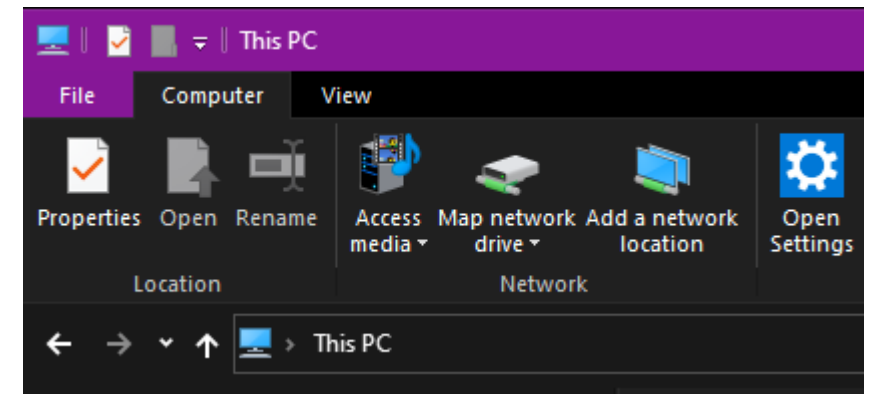
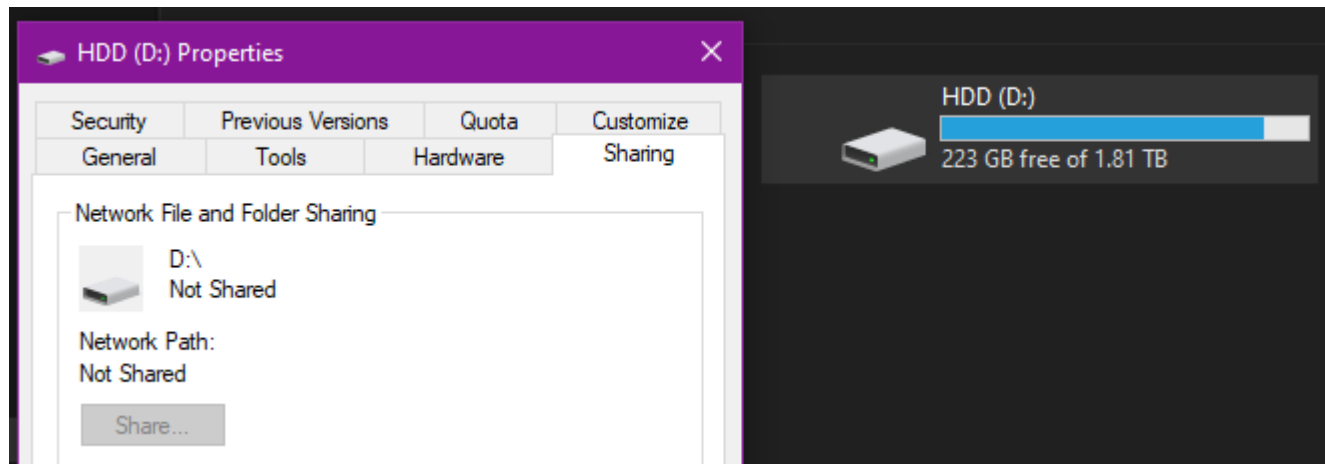
Network Attached Storage and Storage Area Networks

Prof. Ítalo Cunha



Host-based NAS

- Operating systems have support for sharing disks over the network



Storage-centric servers

- Some servers are built to support dozens of disks
- NAS software will allow accesses to these disks



Dedicated NAS devices

- Custom-built devices
- Usually run a proprietary operating system and software stack
- Support for combining multiple NAS devices



Dedicated NAS devices

- Custom-built devices
- Usually run a proprietary operating system and software stack
- Support for combining multiple NAS devices



Storage Area Networks

- Pairs network attached storage with dedicated network infrastructure
 - Network dedicated to data access allows optimizations not suitable to general-purpose networks



Storage Area Networks

- Pairs network attached storage with dedicated network infrastructure
 - Network dedicated to data access allows optimizations not suitable to general-purpose networks
- However, current deployments use standard Ethernet for both application and storage traffic
 - Hyper-converged infrastructure



Comparison

Network Attached Storage

- File interface/network filesystem
- Filesystems are OS-specific
 - Limits cross-platform use
- Works well with any app
- Can be mounted anywhere
 - Including containers

Storage Area Network

- Block interface/virtual disk
- Works with any operating system
- More flexibility
- Not directly usable by apps
 - Including containers
- Hard to share data in virtual disks



CompSci 401: Cloud Computing

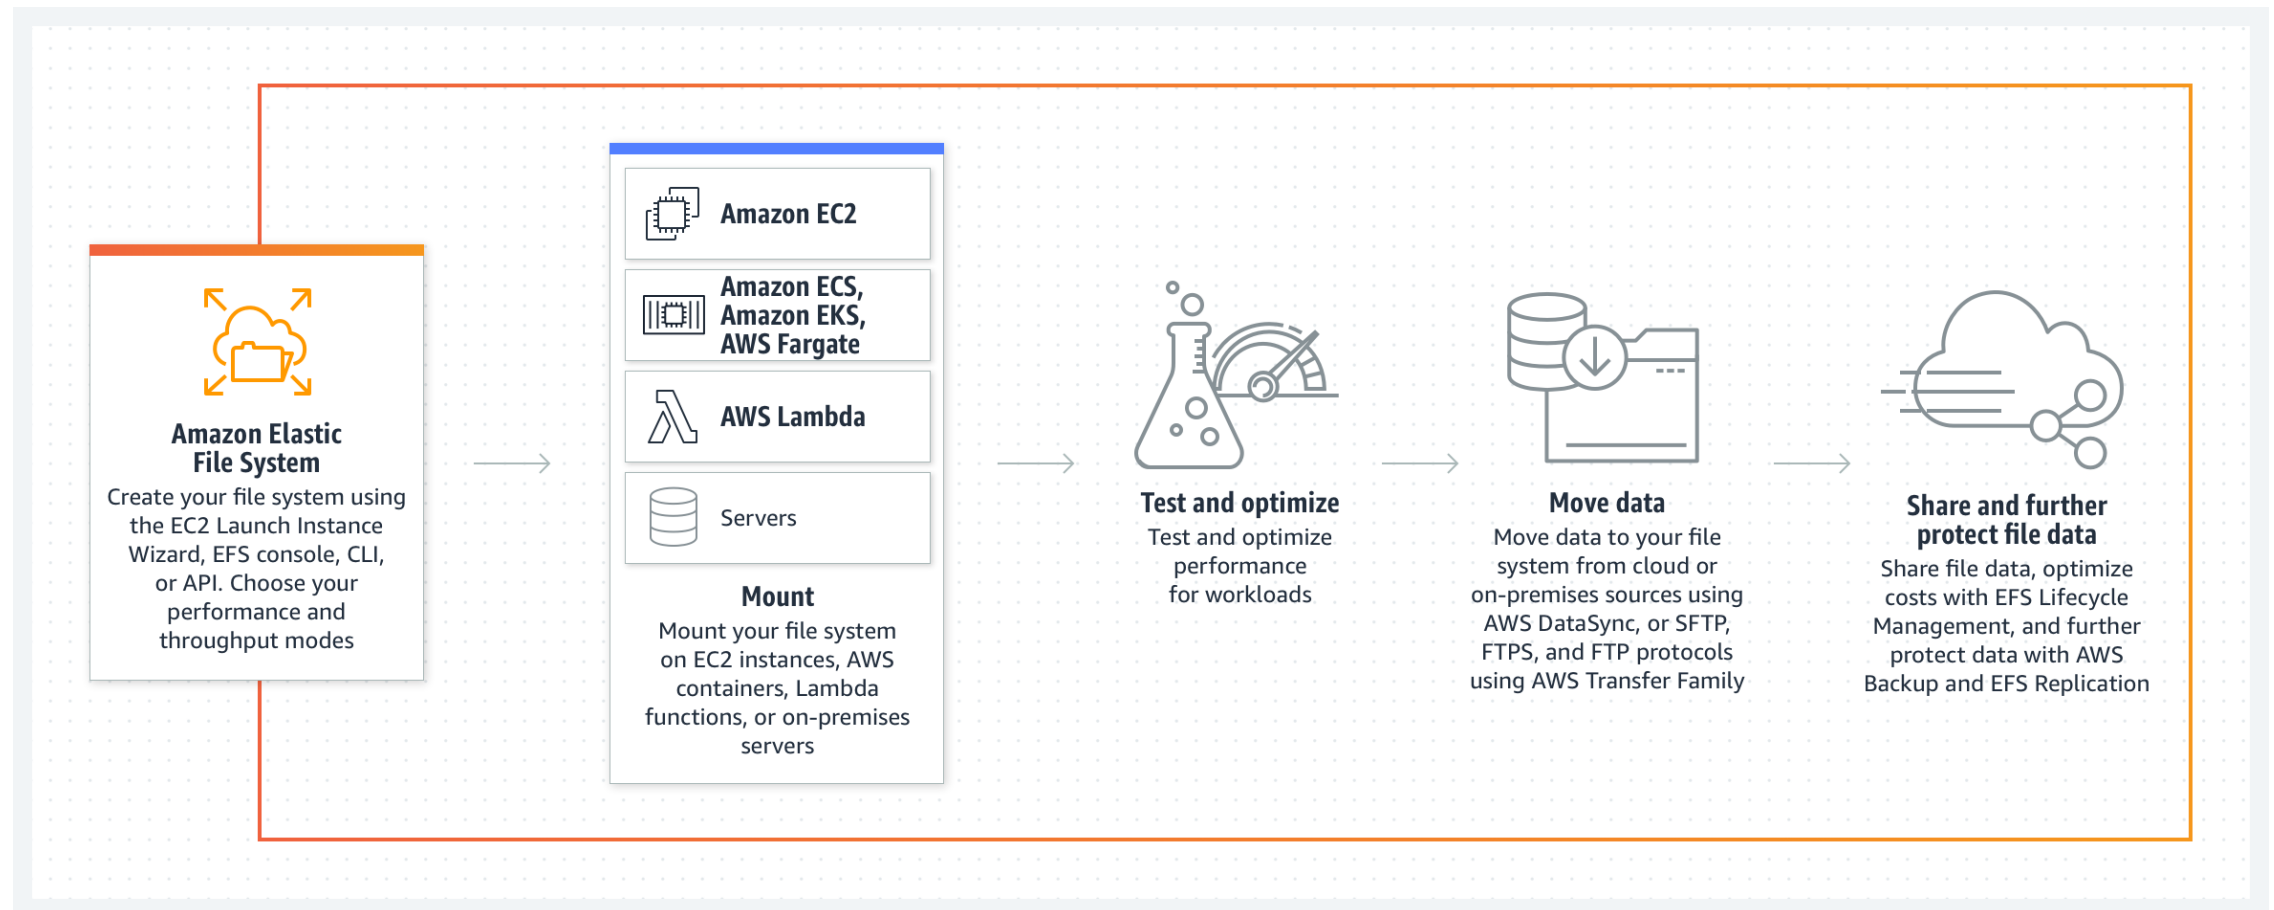
Cloud Storage Solutions

Prof. Ítalo Cunha



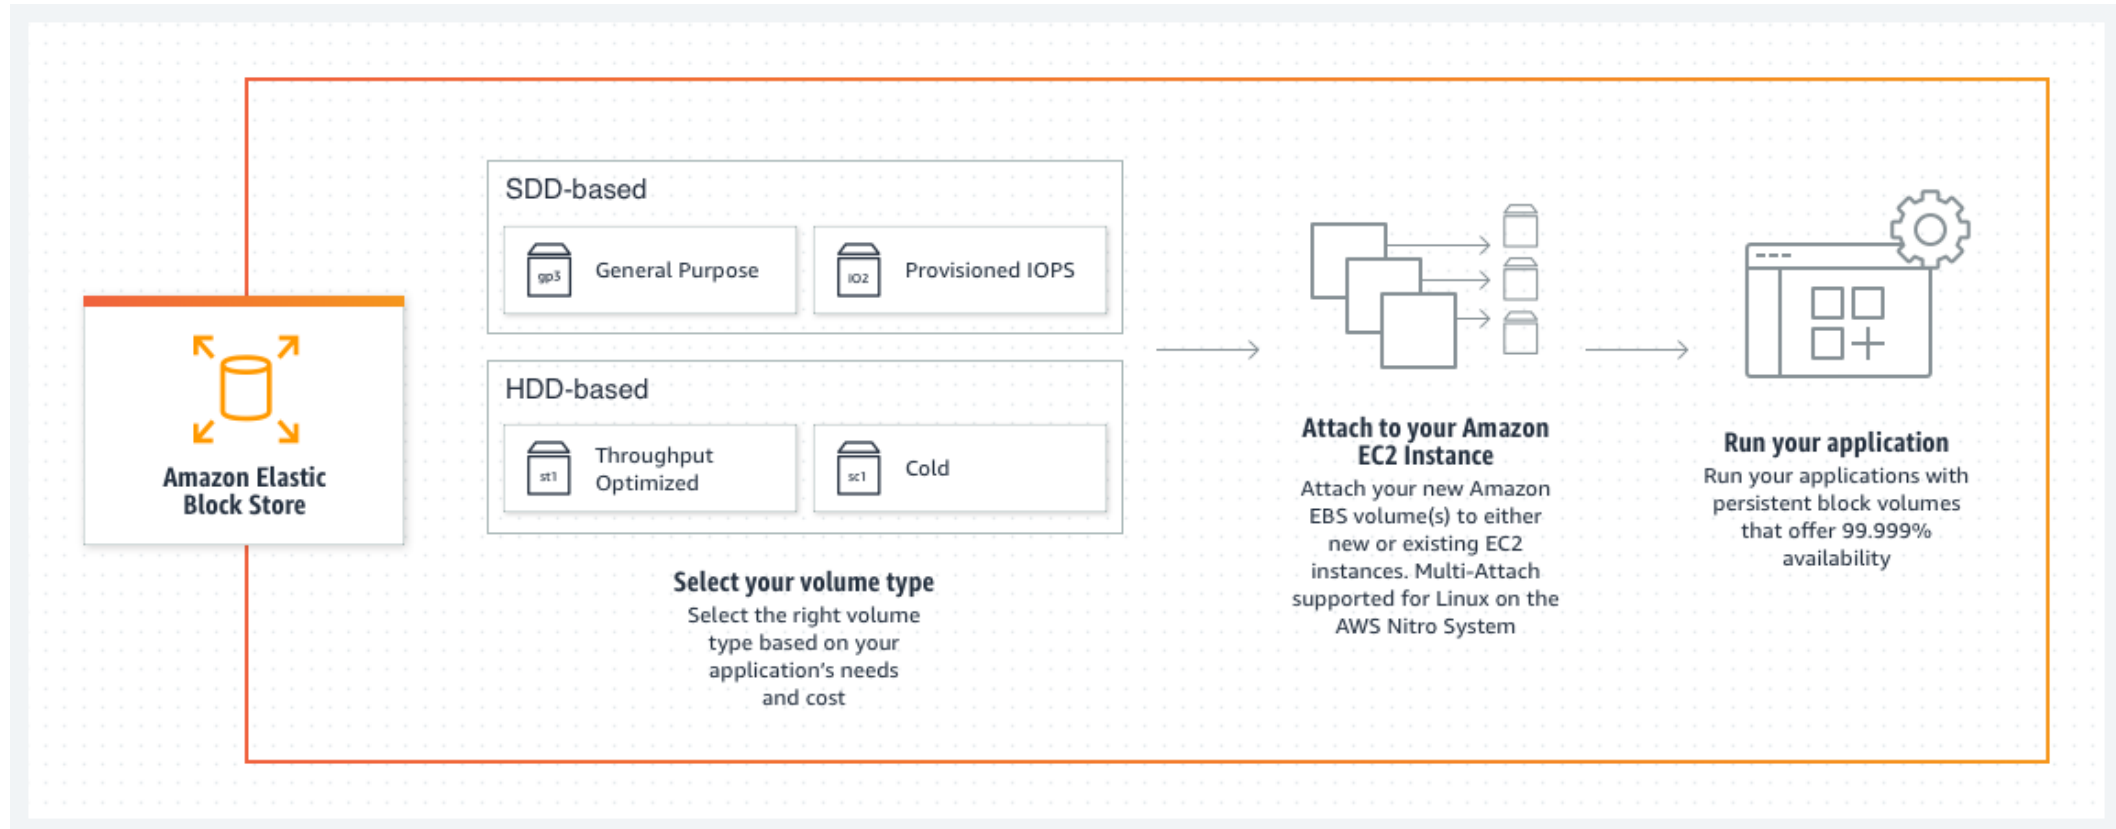
Cloud providers have multiple interfaces

- File interface (Amazon Elastic File System)



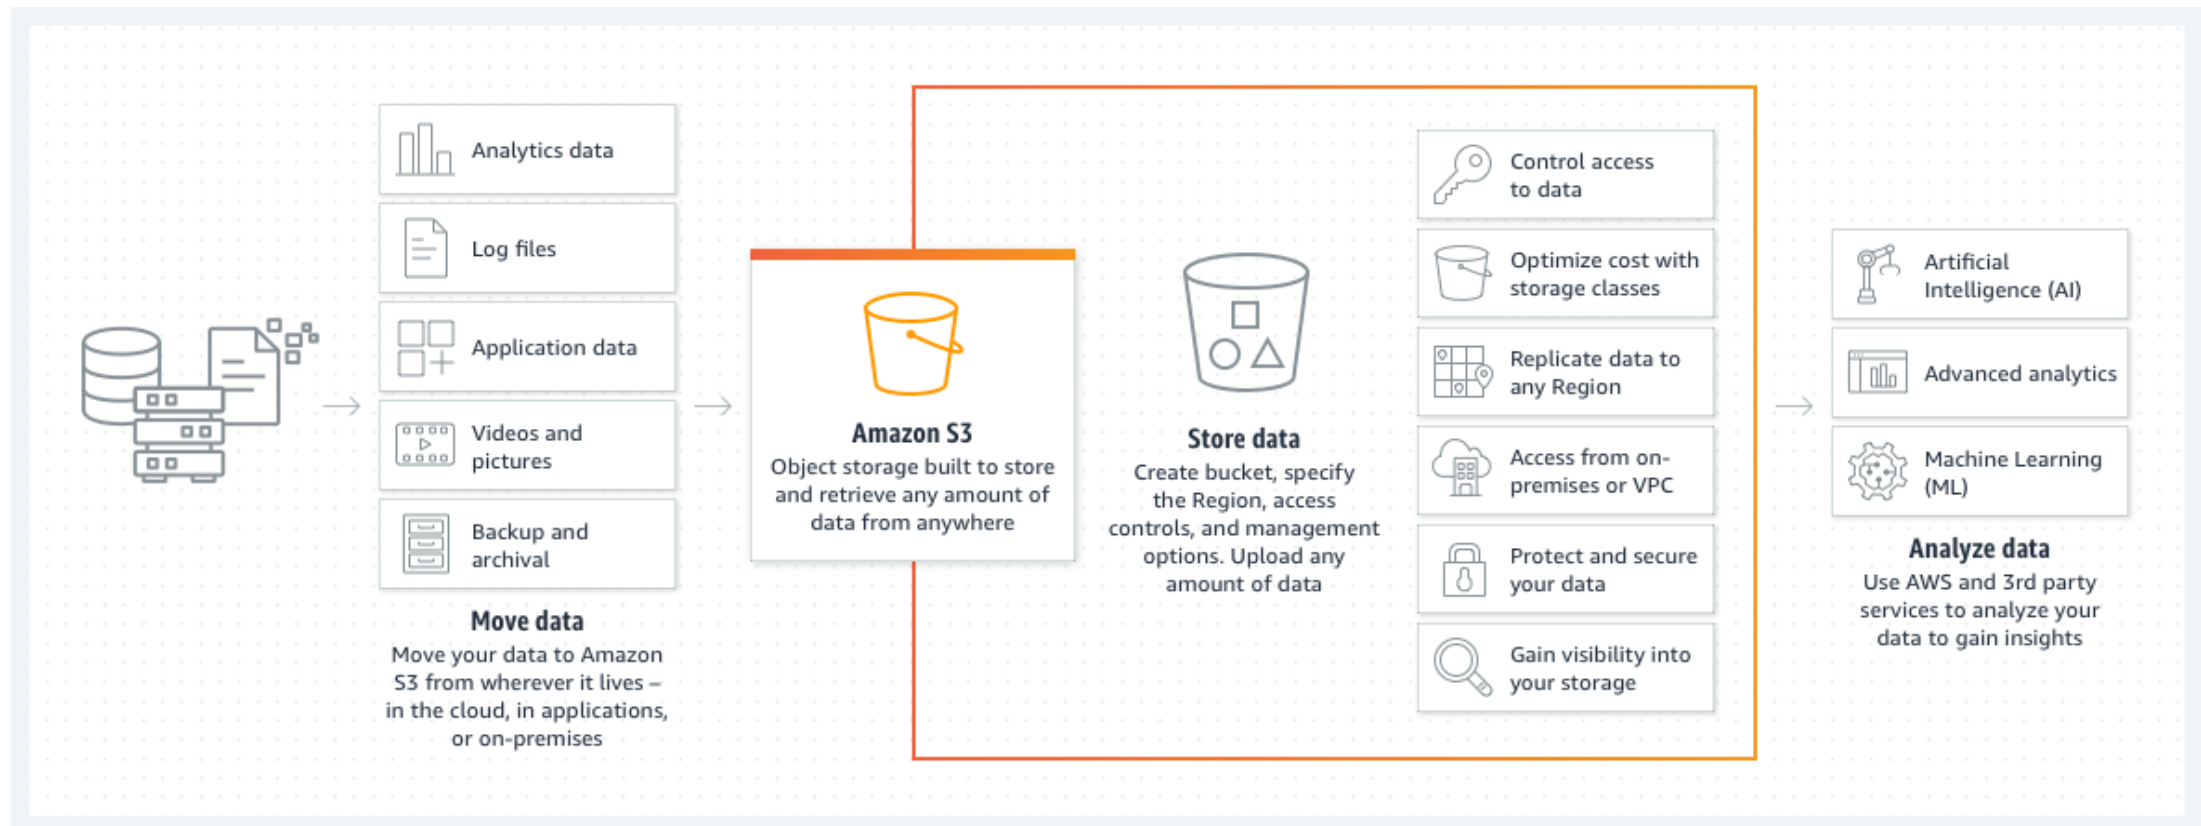
Cloud providers have multiple interfaces

- Block interface (Amazon Elastic Block Store)



Cloud providers have multiple interfaces

- Object interface (Amazon Simple Storage Service)



Object storage

- Arbitrary objects
 - Applications can decide what to store
- High scalability
 - Services can store arbitrary number of objects
- Universal accessibility
 - Objects can be accessed through multiple interfaces
- Independence from operating system
 - Interface defined by the object storage system