

Fall '20 COSE322-00

# System Programming

---

Lecture 08. Network Overview

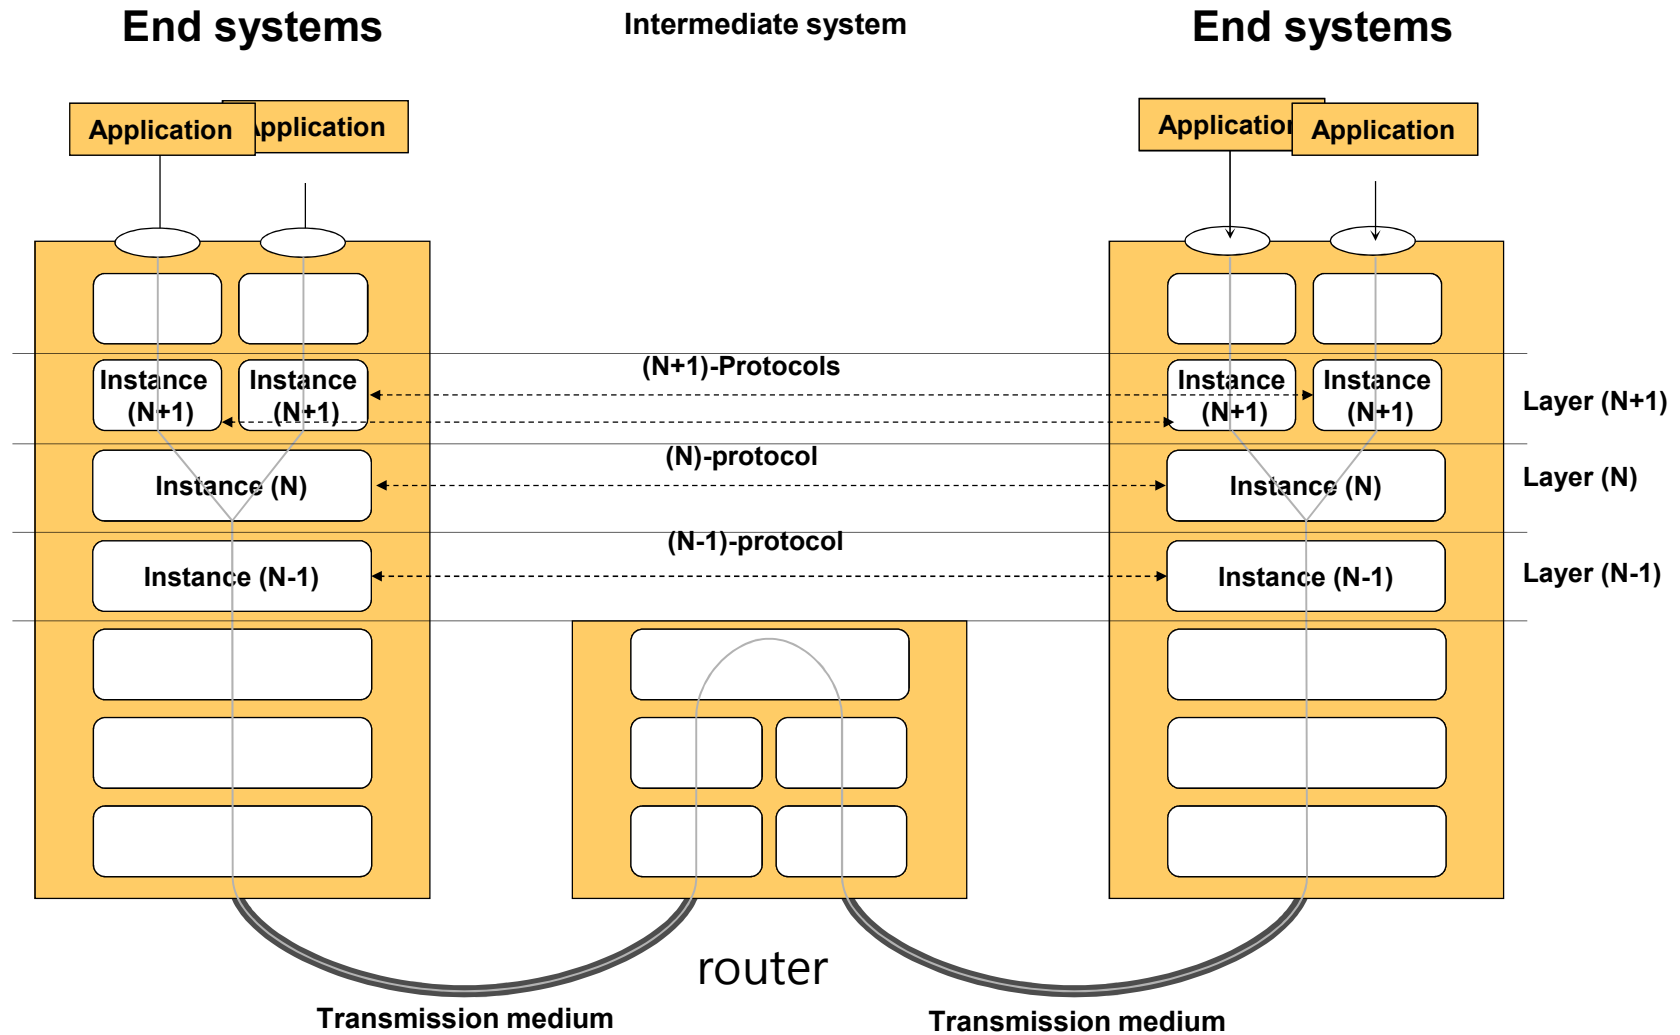
2020. 10. 27.

# Contents

- ❖ **Layer-Based Network Models**
  - Concepts and Characteristics
  - Protocol Stack
  - Port, Socket
- ❖ **Kernel structure**
  - Overall structure
  - Kernel networking
- ❖ **I/O Models in Linux Network**
  - Blocking, Non-blocking
  - I/O Multiplexing



# Layer-Based Network Models



# Layer-Based Network Models

## ❖ Concepts in Layered Model

- Protocols
  - Rules that two parties talk and understand each other
  - Independent of hosts and technologies
  - Horizontal interface
- Services
  - Functions provided by a lower layer to the neighboring upper layer
  - Vertical interface

## ❖ Characteristics

- Advantages: Simplifying design, implementation, understanding and maintenance
- Disadvantages: Overhead going through layers
  - No short cuts

# OSI Reference Model

## ❖ Physical layer

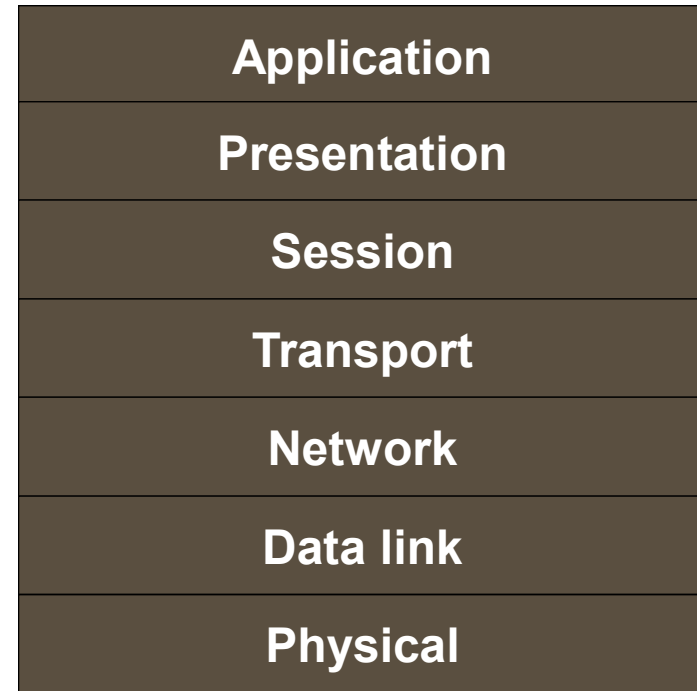
- Handling bits on to wire
- Implemented by hardware
- Media types, coding methods

## ❖ Data link layer

- Addressing within the physical network
- Moving data in local network
  - two directly connected hosts

## ❖ Network layer

- Connecting different networks
  - One global virtual network
- Packet routing/forwarding to hosts



# OSI Reference Model

## ❖ Transport layer

- End-to-end Delivery to apps
  - Flow control, packet ordering

## ❖ Session layer

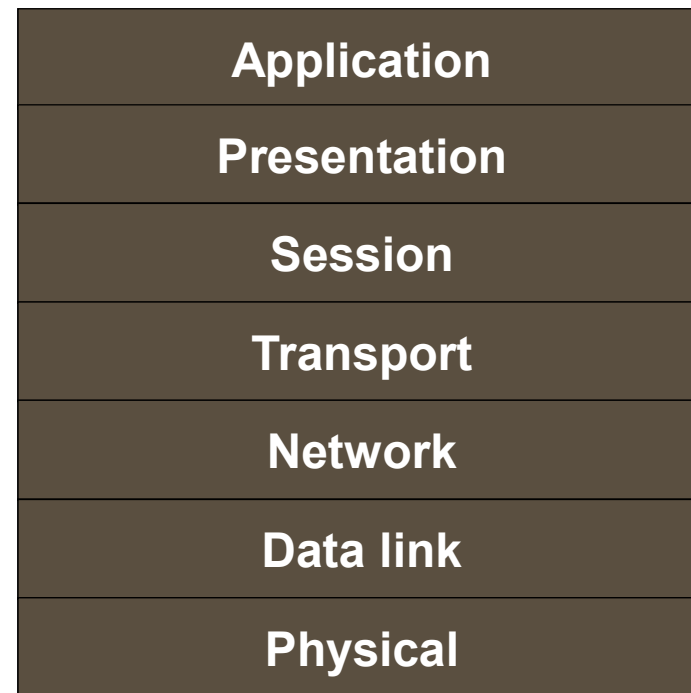
- opening, closing and managing a session

## ❖ Presentation layer

- Regulating data representation
  - Encryption, XDR

## ❖ Application layer

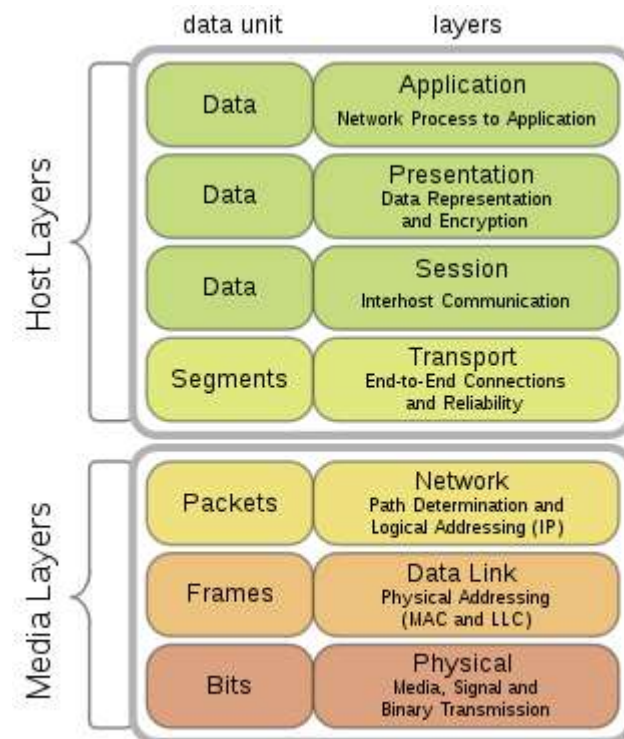
- SMTP, FTP, HTTP



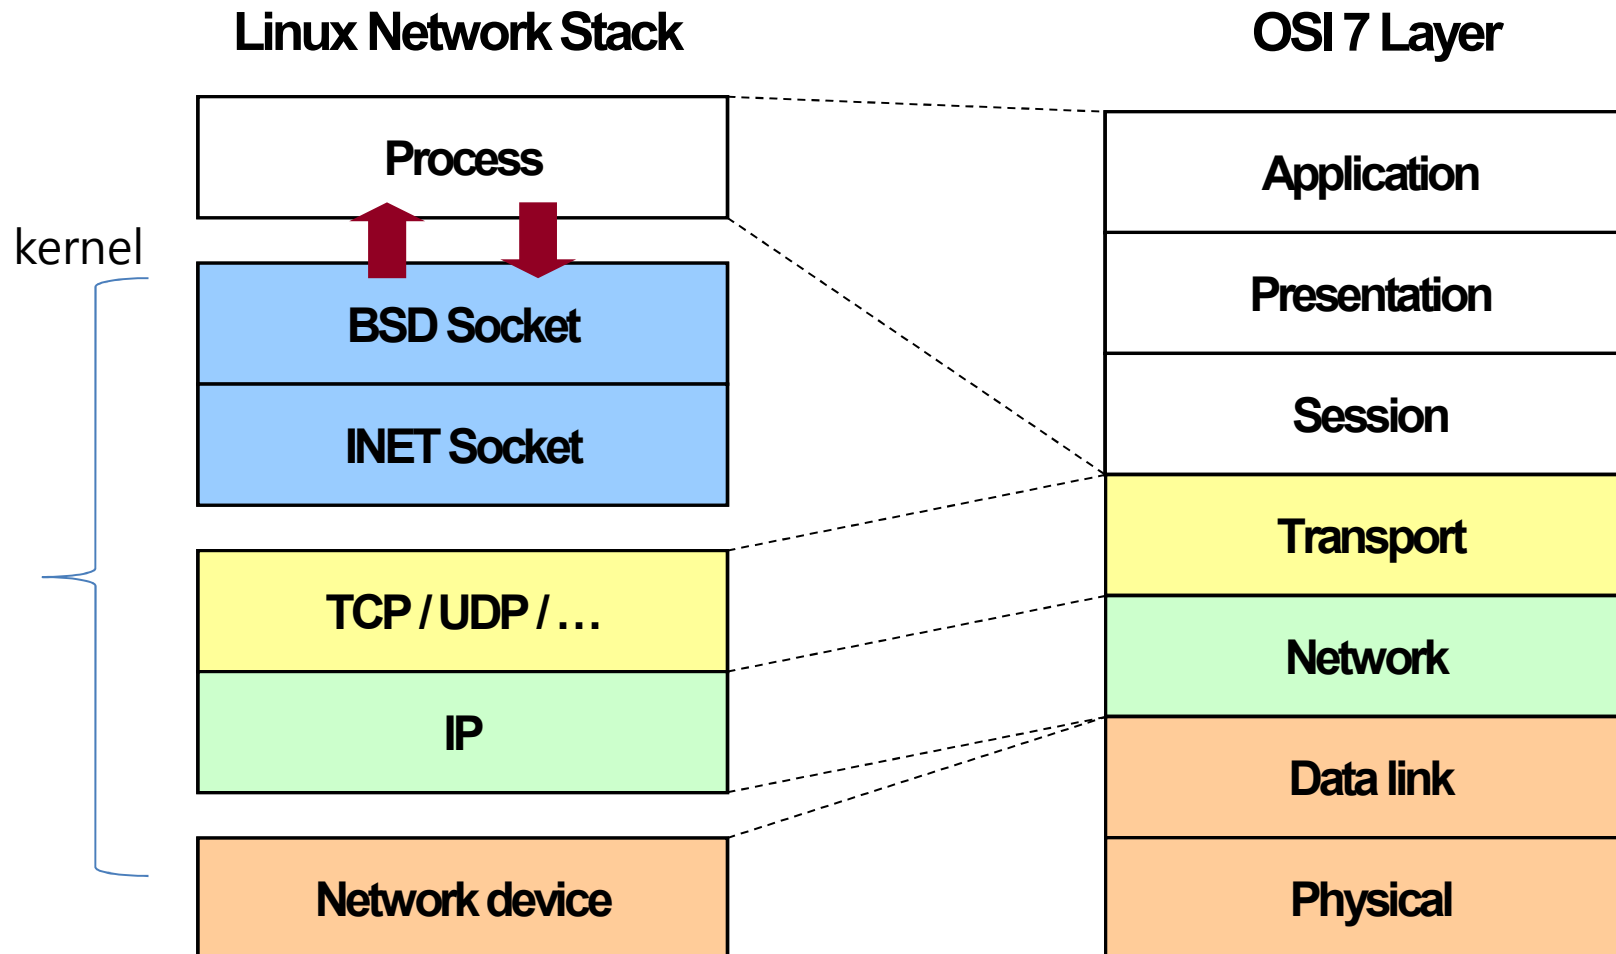
# Protocol Stack

## ❖ What is it?

- The term stack refers to the actual software that implements the protocols
- A set of network protocols that work together
- Also called network stack
  - in OSI Reference Model
- Code inside kernel
  - Link to transport layer



# Protocol Stack in Linux Kernel

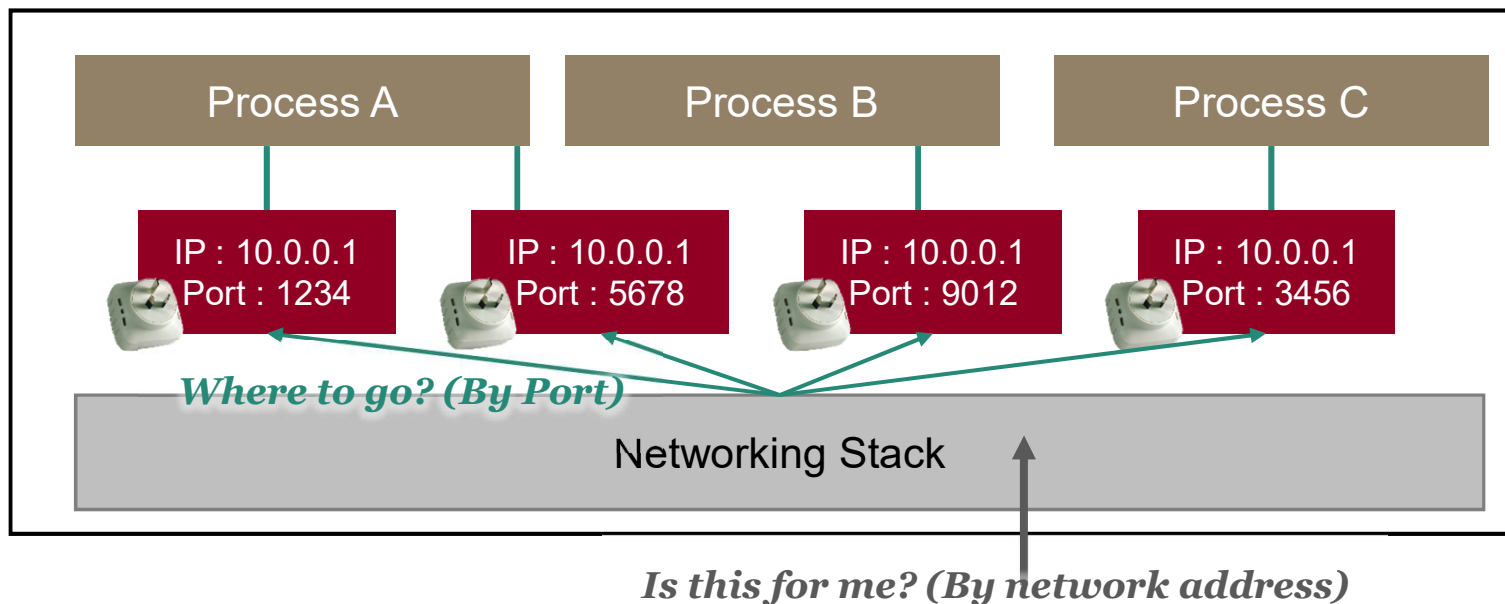




# Port

## ❖ Port (16bit)

- Abstraction provided by kernel
- Communication endpoint in “host”
- Multiple concurrent connections can be made out of a port



# Port

## ❖ Well-known, Registered, Dynamic port

- 0 ~ 1023 : Well-known ports
- 1024 ~ 49151 : Registered ports
- 49152 ~ 65535 : Dynamic ports

## ❖ Well-known ports

Port	Descriptions	Port	Descriptions
0	Reserved	53	DNS
20	FTP (for data)	80	HTTP
21	FTP (for control)	110	POP3
22	SSH	161	SNMP
25	SMTP	179	BGP

## Port vulnerability

### ❖ Port is open


- Exposed to outside world
- Many ports are blocked for security
  - To avoid port scan

### ❖ Ethernet is not encrypted

- Clear text passwd
  - tcpdump

# Socket

- An **end-point** for Internet network “connection”
- Socket is defined by the combination of **a network address and a port** identifier
  - Network Address: IP address
- Process view to port
  - Port is shared among processes



## More on socket

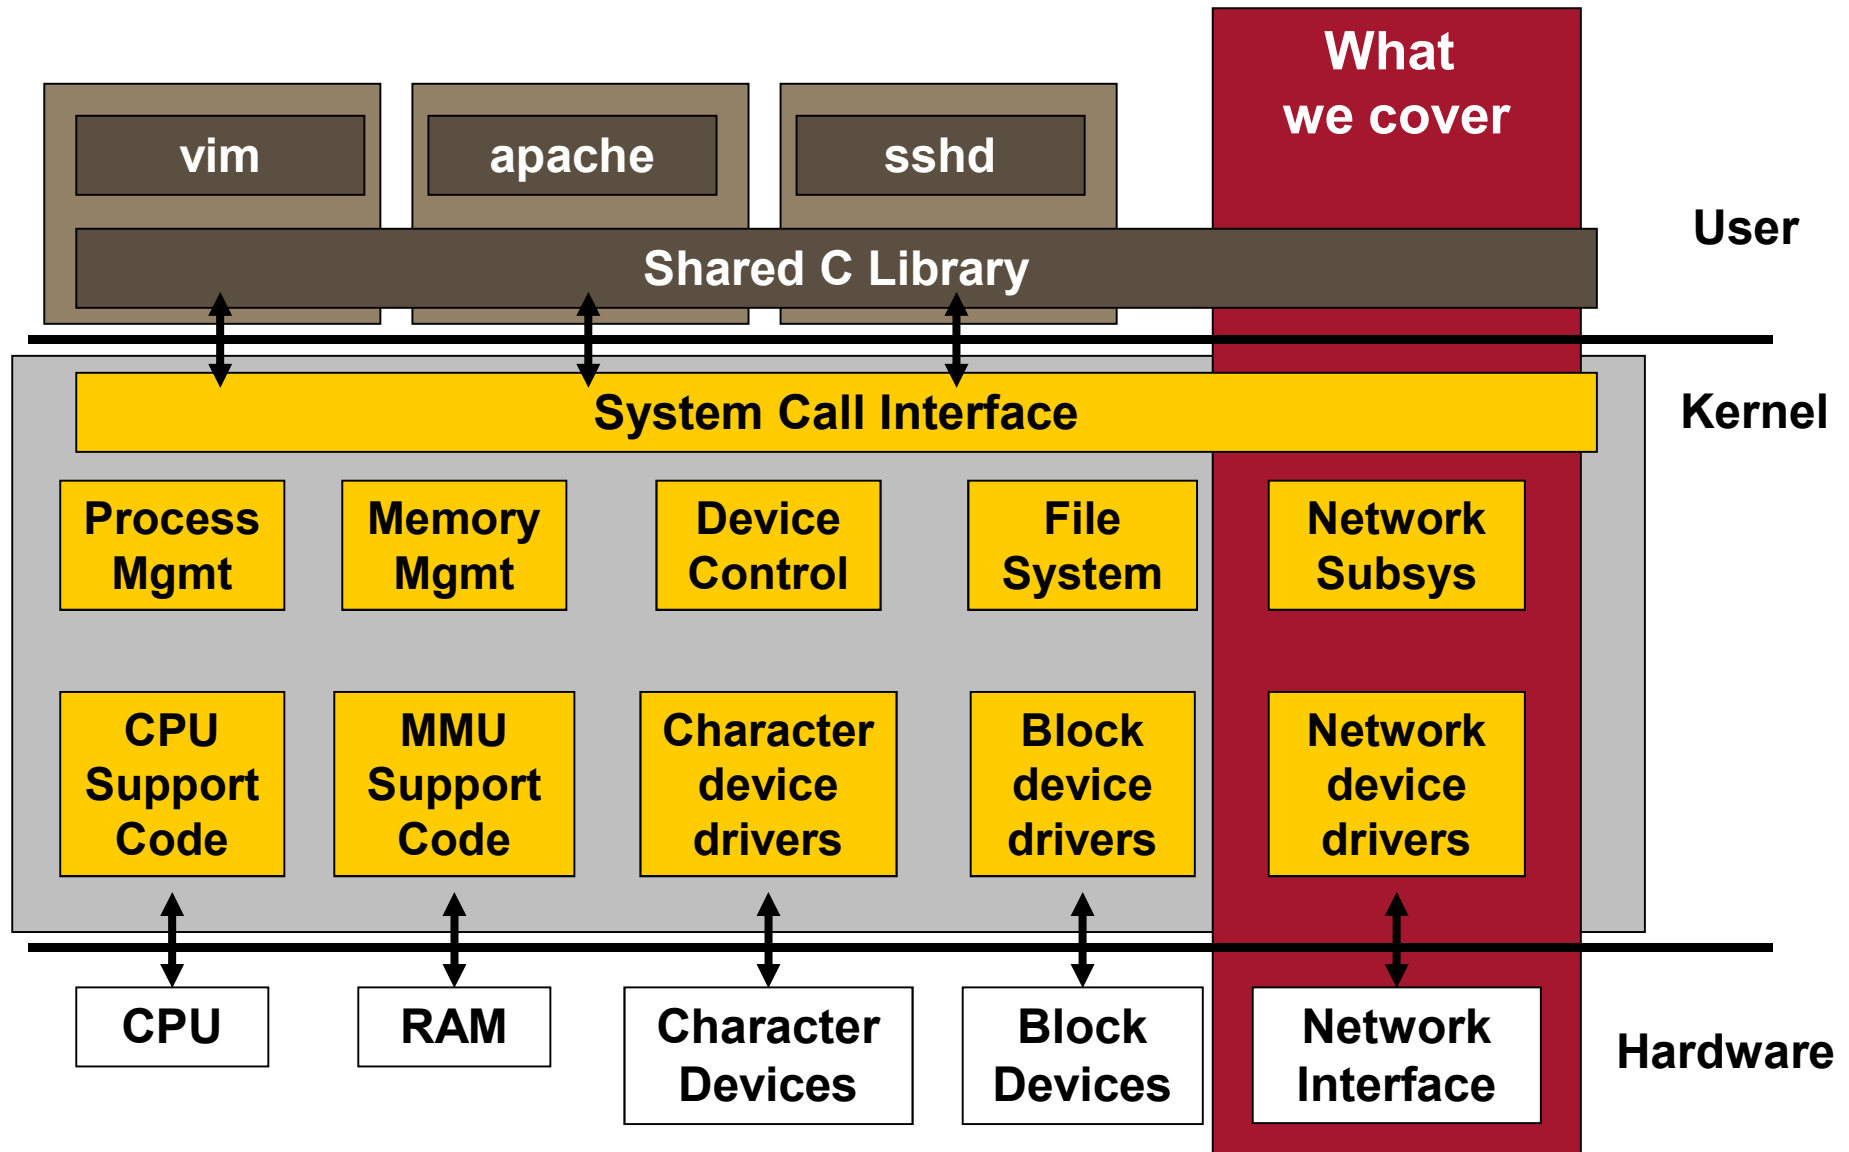
### ❖ Use file descriptor

- socket(2) returns fd
- The fd is used to send or receive packets

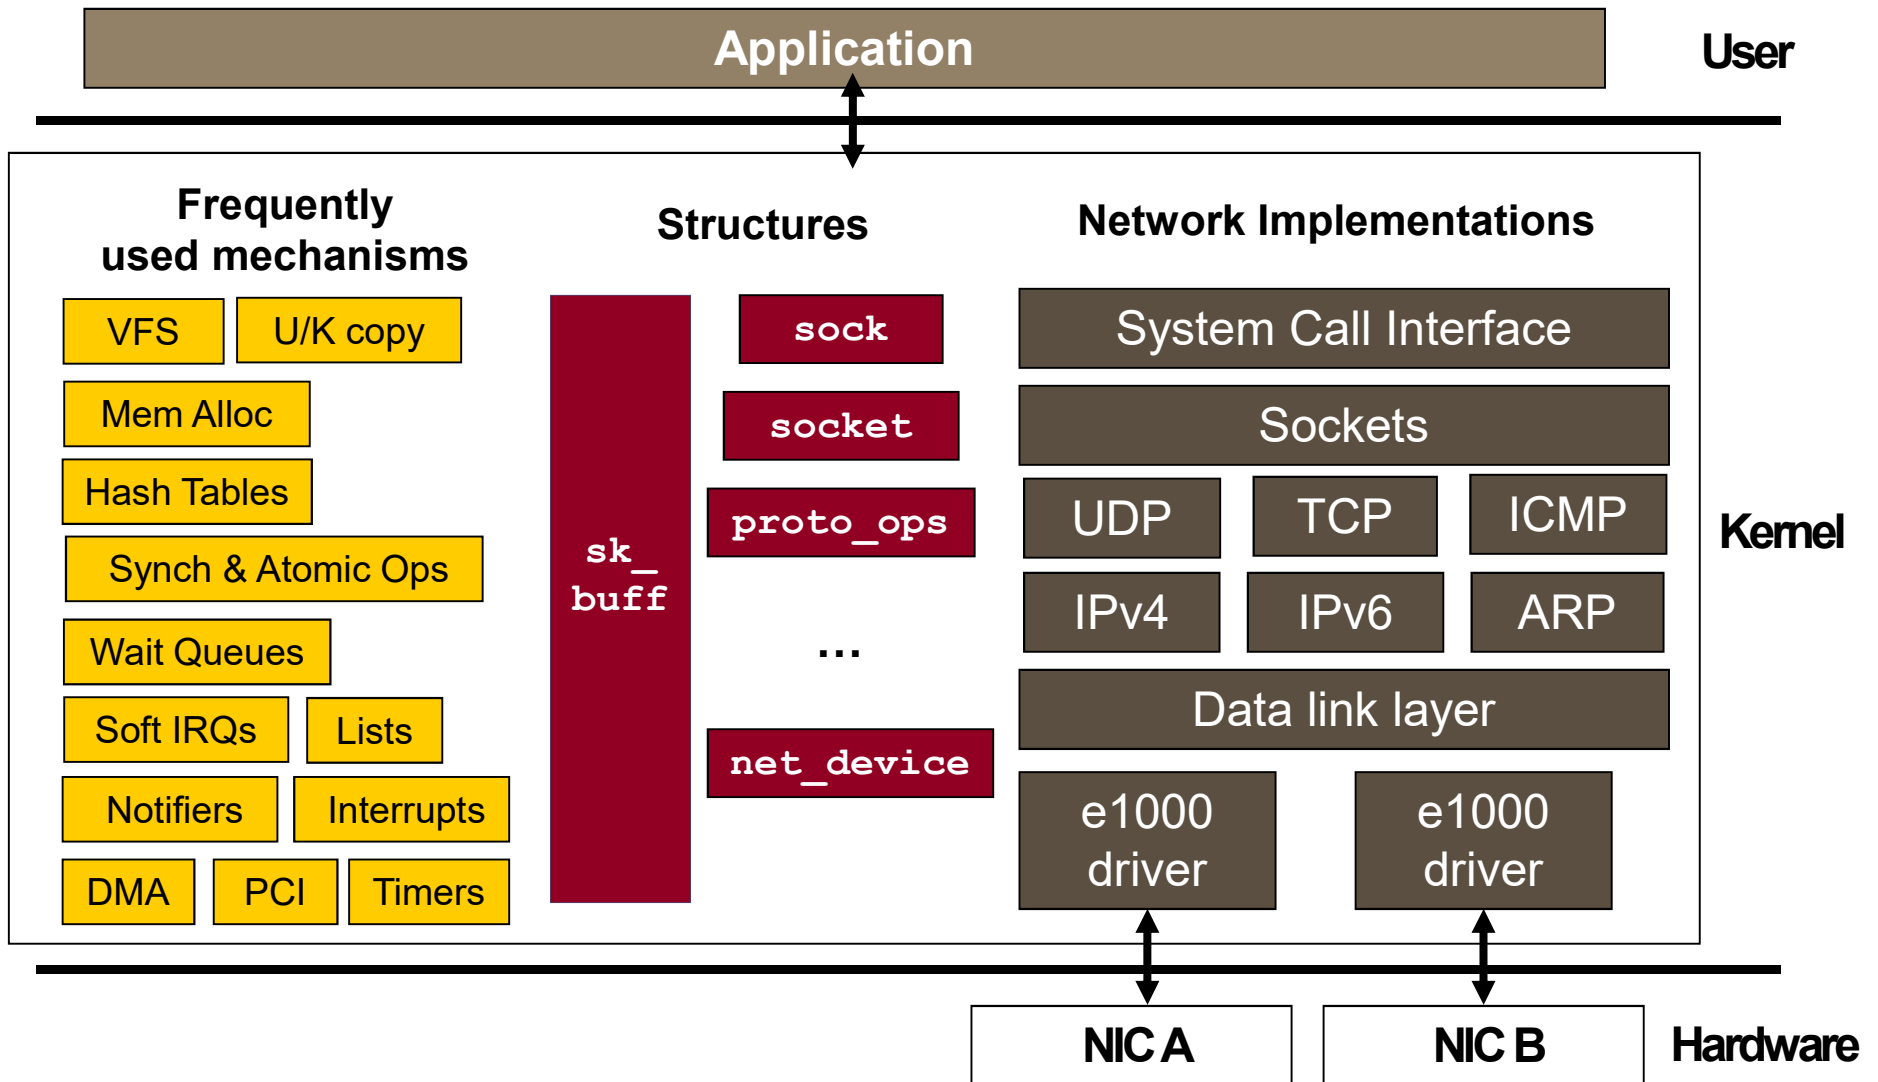
### ❖ Does not have open(2)

- Instead, connect and listen
- Since “connection” is transient and not shared

# Kernel Structure



# Kernel Networking



## I/O Models in Linux Network

### ❖ After the I/O request,

- **Blocking**: waits until the data to be ready
- **Non-blocking**: always returns

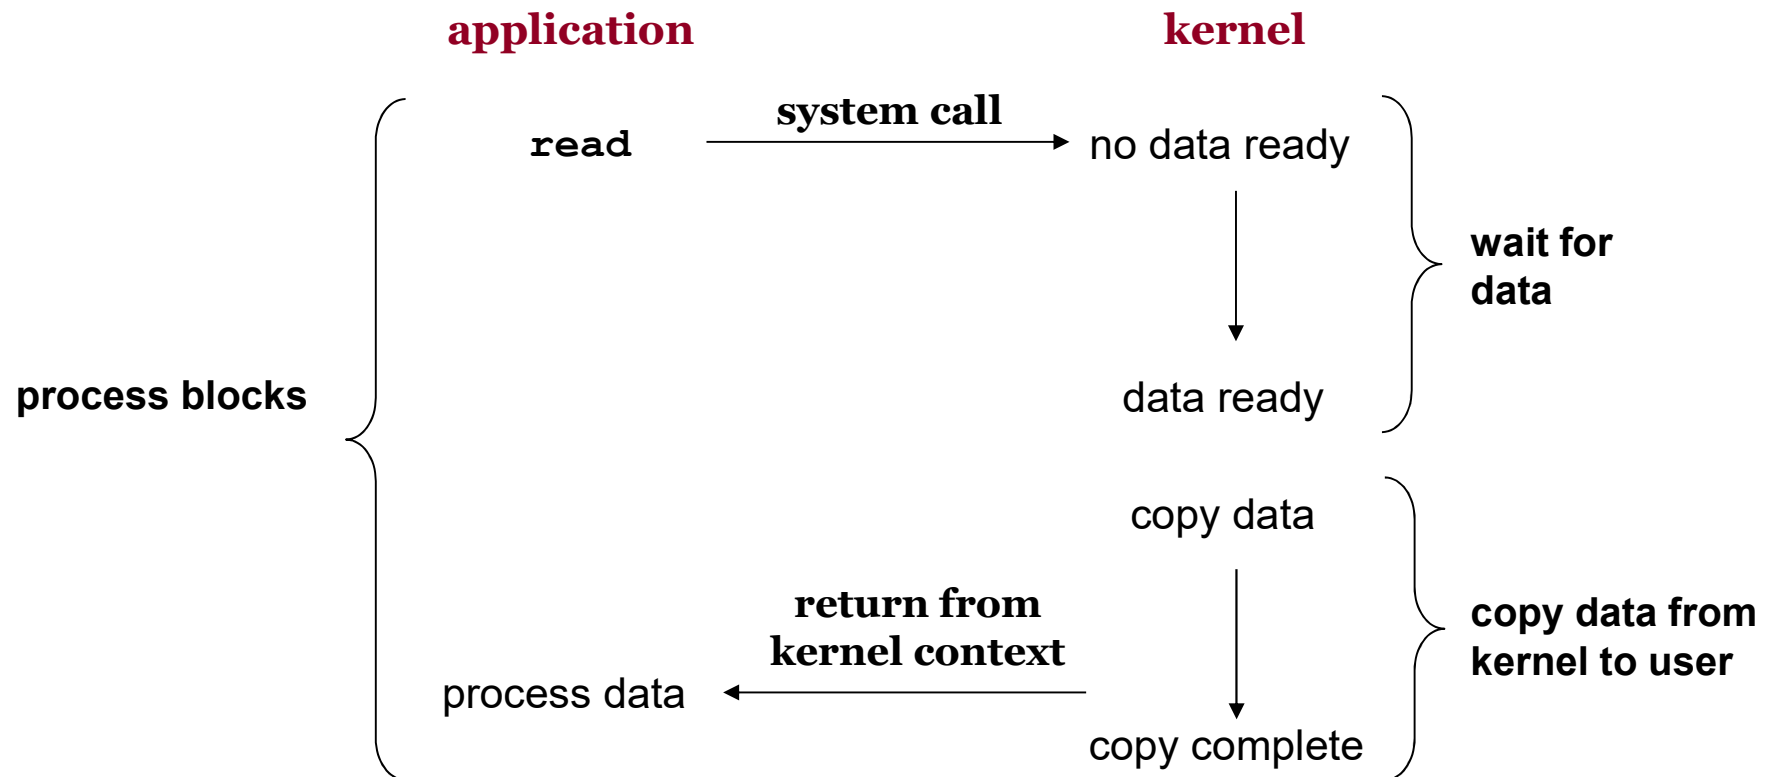
Blocking	<code>read/write system call</code>
Non-blocking	<code>read/write system call with O_NONBLOCK</code>
	<code>aio_read/aio_write</code>

### ❖ I/O Multiplexing

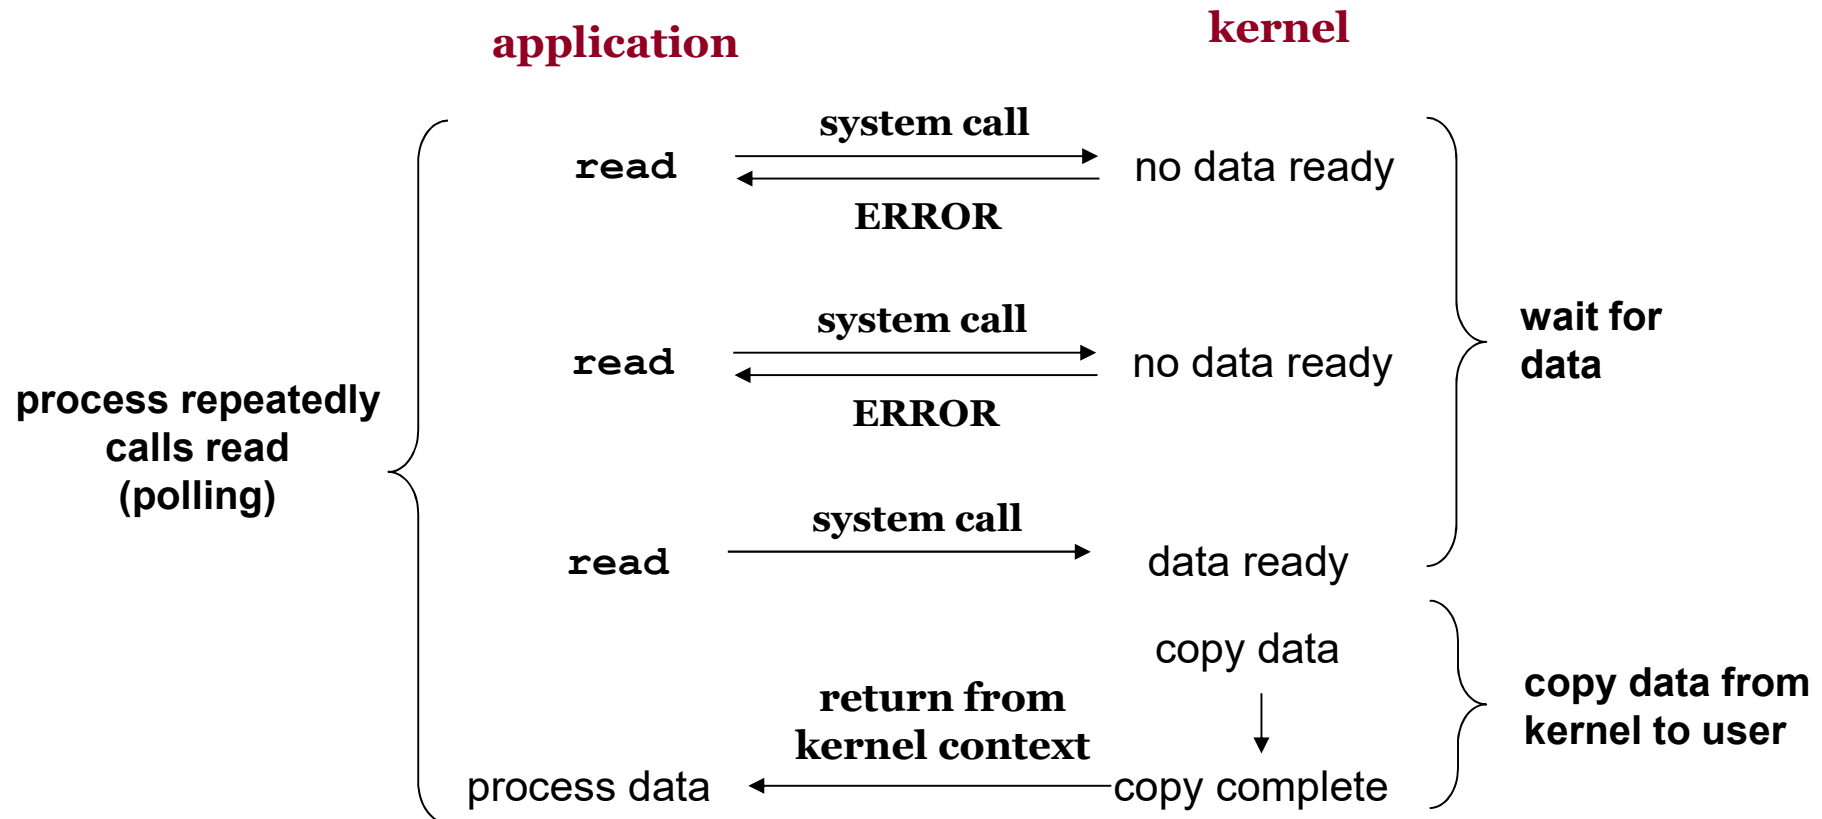
- `select` system call



## Blocking I/O (read)



## Non-blocking I/O (read with O\_NONBLOCK)



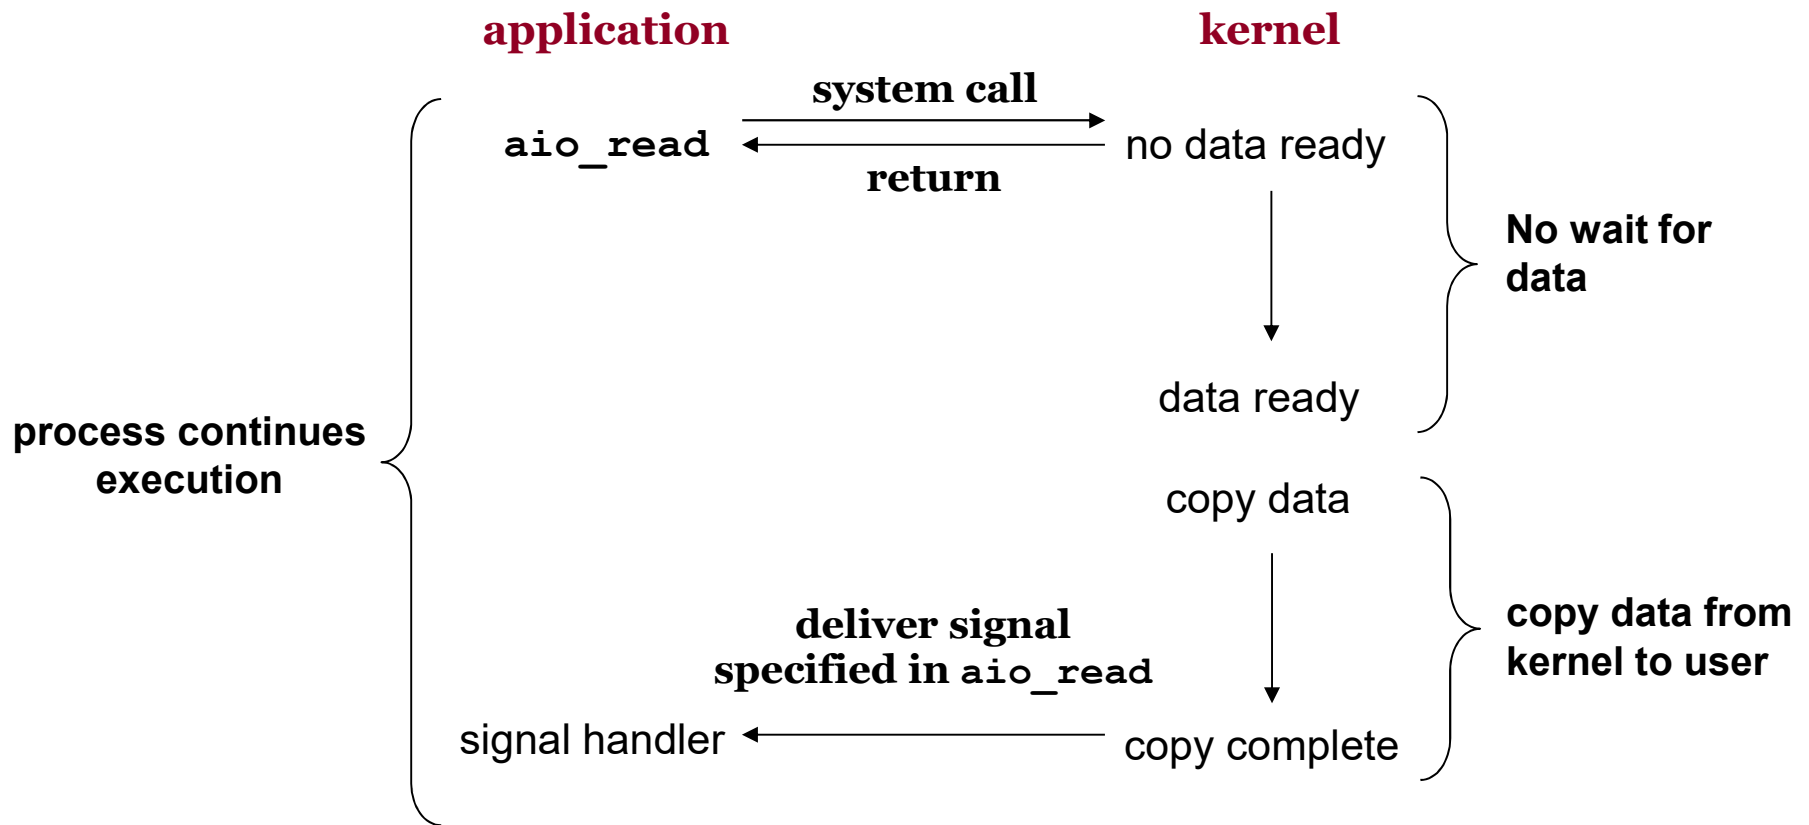
## Non-blocking I/O (AIO)

❖ `aio_read()` is an user API for **AIO**

```
int aio_read(struct aiocb *aiocbp);
```

- Queues the I/O request described by the buffer pointed by **aiocbp**
  - `struct aiocb` : Defines parameters that control an I/O operation
    - File descriptor, Location of buffer, Length to transfer, Notification Method, etc.
- The read request **returns immediately**
  - The application can perform other processing while the background read operation completes
- When read response arrives, a signal or a thread-based callback is generated to **notify** the completion

# AIO diagram



## I/O Multiplexing (Select System Call)


### ❖ `select` system call

- Allows a program to monitor **multiple file descriptors**
- **Waits until** one or more of the file descriptors become **ready** for some class of I/O operation (read, write, exception)
- When data is available for one of the descriptors, the `select` system call **returns**
- The data copy operation is **initiated by return** of `select` system call. (e.g., `read` system call)
- Superseded by `epoll(2)`

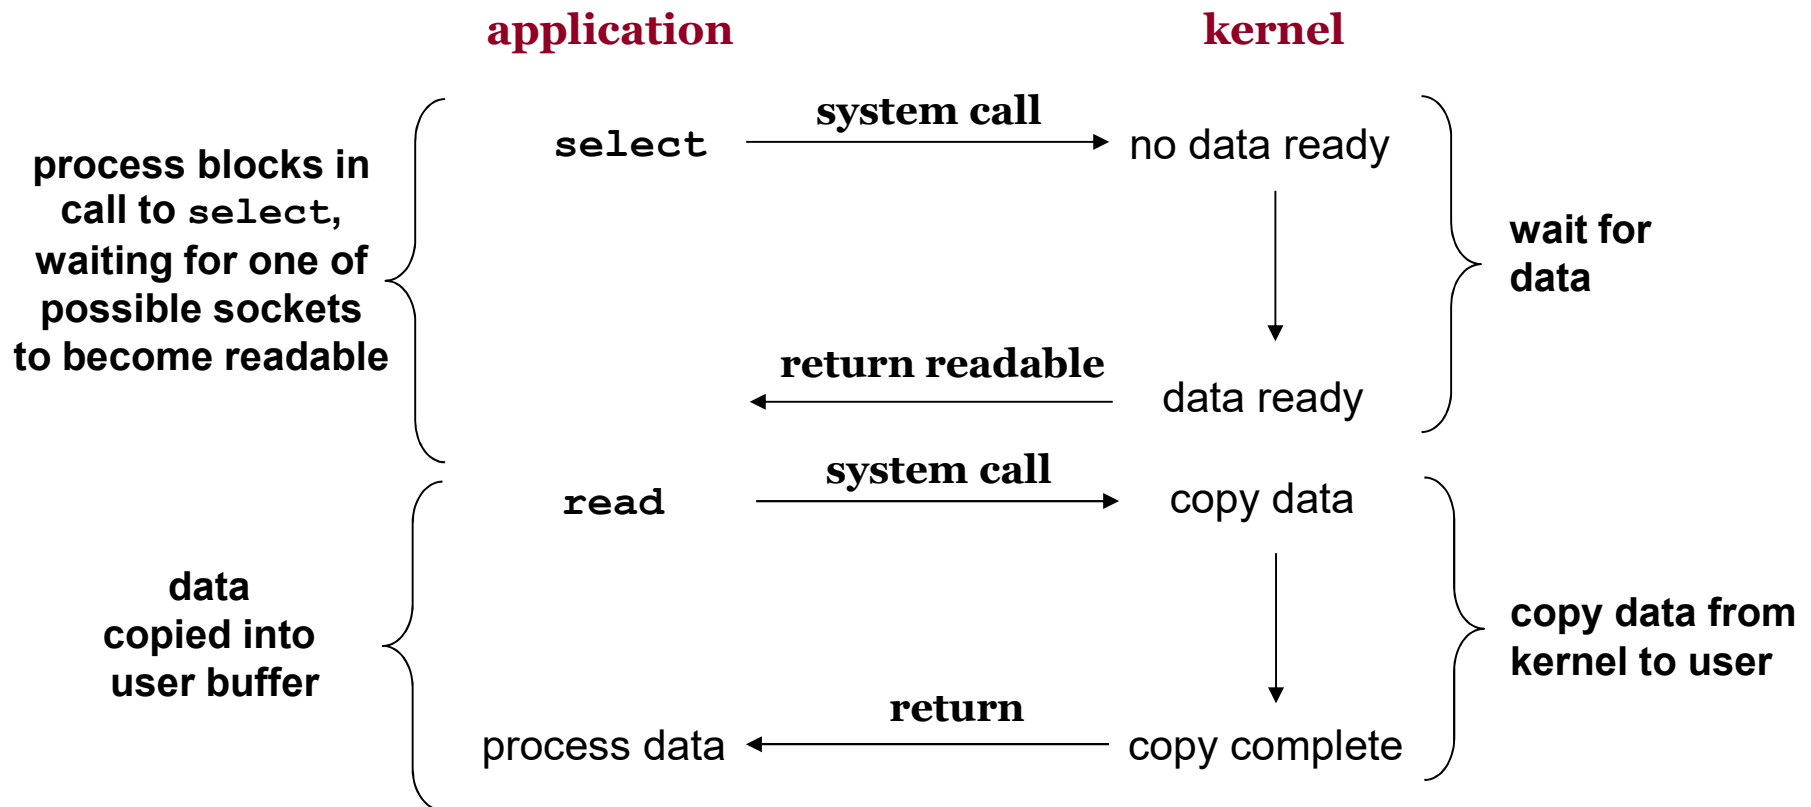
## Select System Call syntax

```
int select(int nfd, fd_set *readfds, fd_set *writefds,
           fd_set *exceptfds, struct timeval *timeout)
```

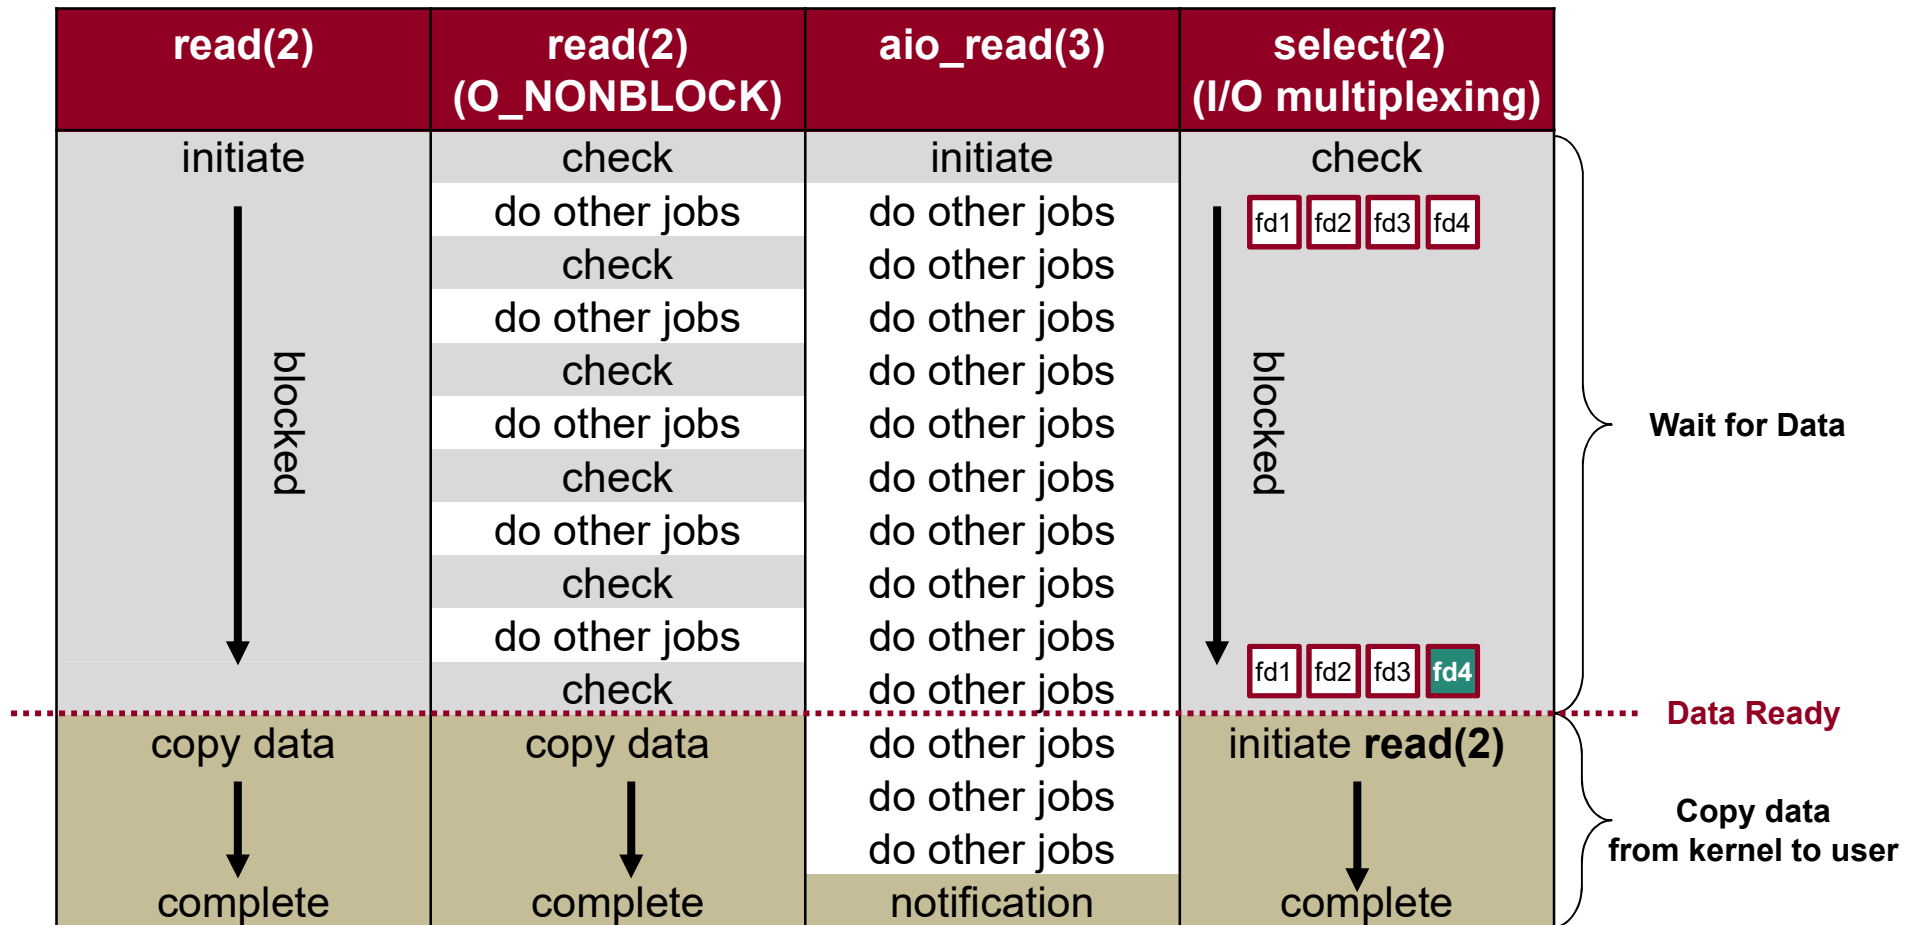
- `nfd` : the highest-numbered file descriptor, plus 1
- Three independent **bitmaps of file descriptors** (`struct fd_set`)
  - `readfds`(for read), `writefds`(for write), `exceptfds`(for exceptions)
  - Each bitmap covers 1024 fds
- `timeout` : minimum interval that `select()` should block



# Select System Call diagram



# Comparison





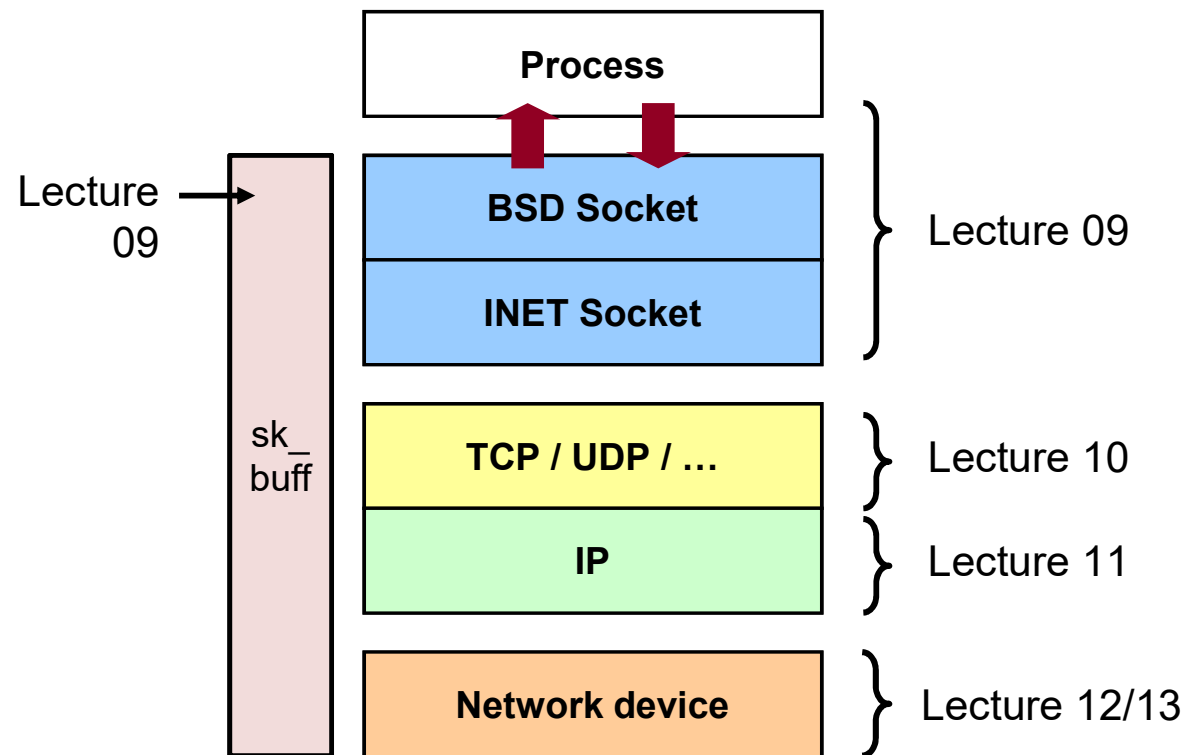
# Thread and networking

## ❖ Benefit of thread

- Thread can block
- Other threads can do other computation in parallel

# Lecture plan for Networking

## ❖ Top-down approach for network implementations



– Wrapping up (Lecture 14)