

hepsiburada




# Emre YALVAÇ

Senior Software Engineer - Hepsiburada

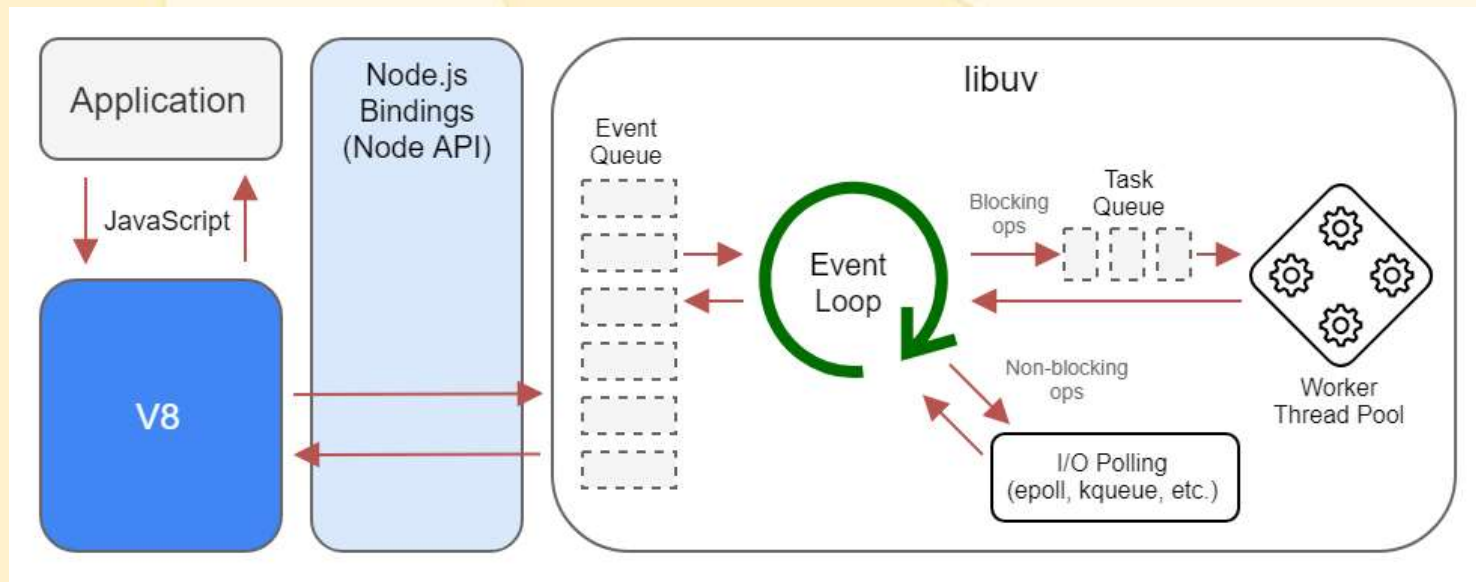
**Core of Node.JS: Libuv**

[github/emreyalvac](https://github.com/emreyalvac)

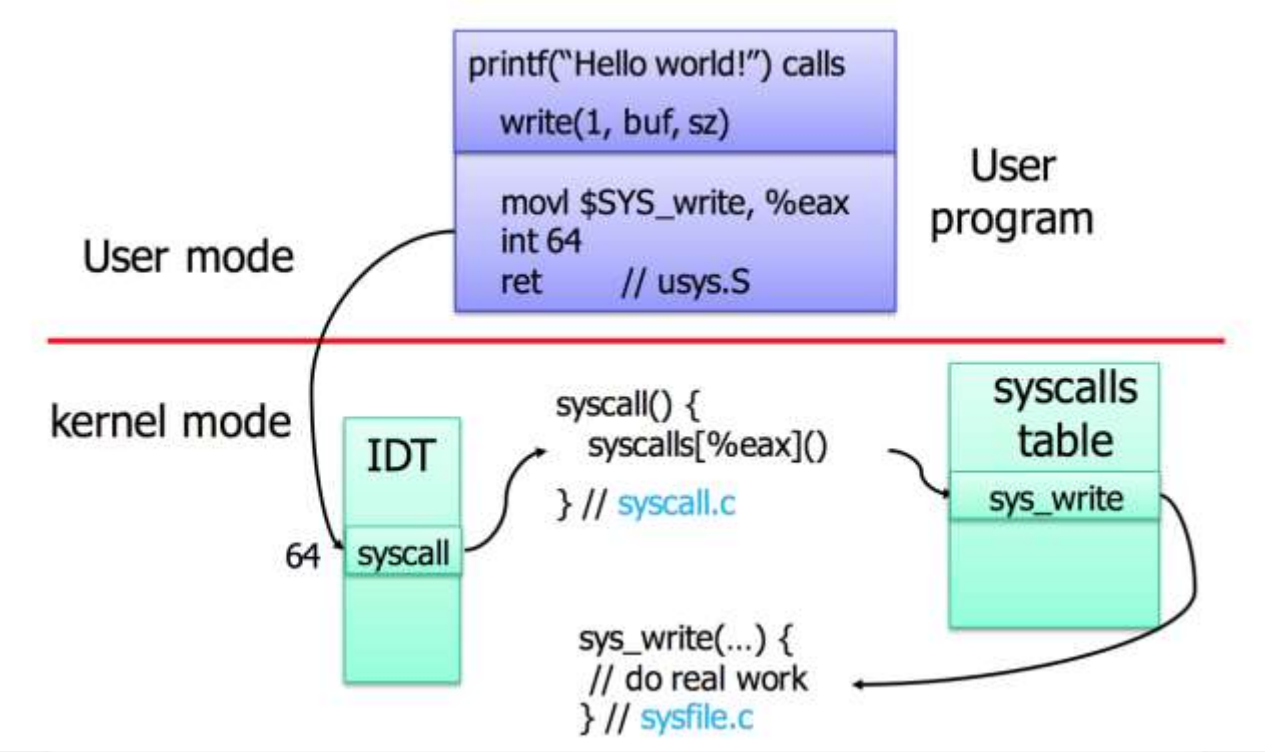


JSKonf

# Runtime



# Syscall



man7.org > Linux > man-pages Linux/UNIX system programming training

## epoll(7) — Linux manual page

[NAME](#) | [SYNOPSIS](#) | [DESCRIPTION](#) | [VERSIONS](#) | [STANDARDS](#) | [HISTORY](#) | [NOTES](#) | [SEE ALSO](#)

[Search online pages](#)

**epoll(7)** Miscellaneous Information Manual **epoll(7)**

**NAME** [top](#)

epoll - I/O event notification facility

**SYNOPSIS** [top](#)


```
#include <sys/epoll.h>
```

**DESCRIPTION** [top](#)

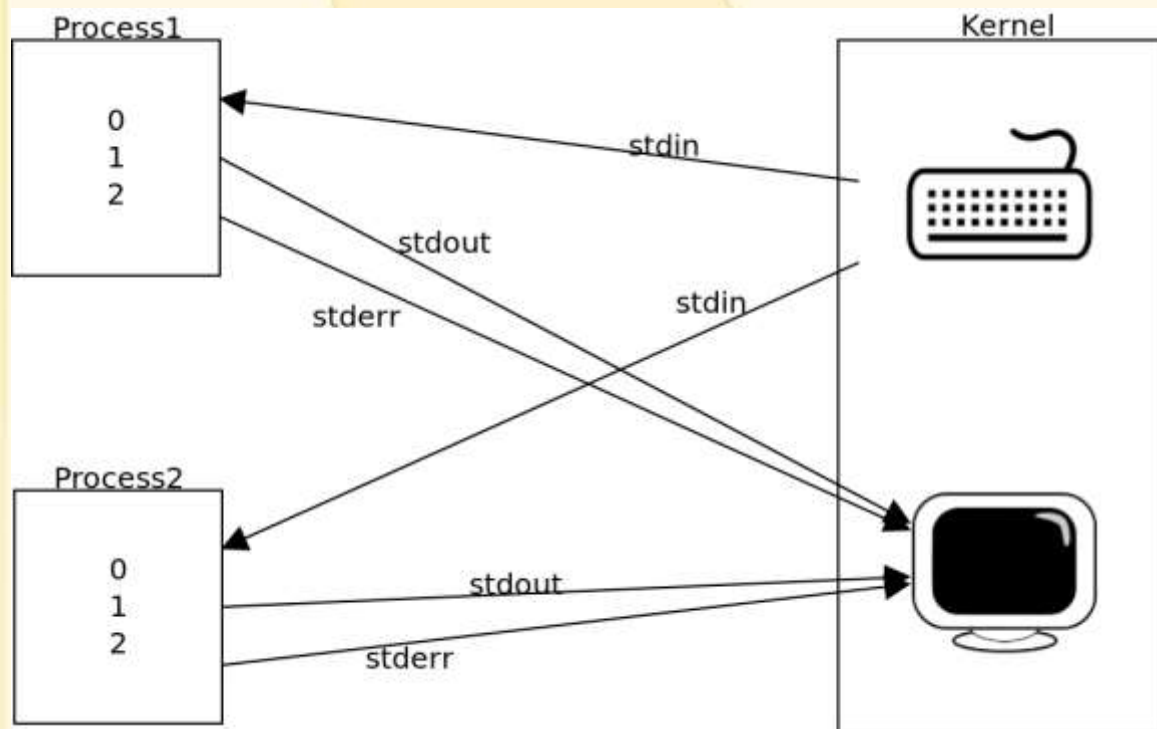
The **epoll** API performs a similar task to [poll\(2\)](#): monitoring multiple file descriptors to see if I/O is possible on any of them. The **epoll** API can be used either as an edge-triggered or a level-triggered interface and scales well to large numbers of watched file descriptors.

# File Descriptor

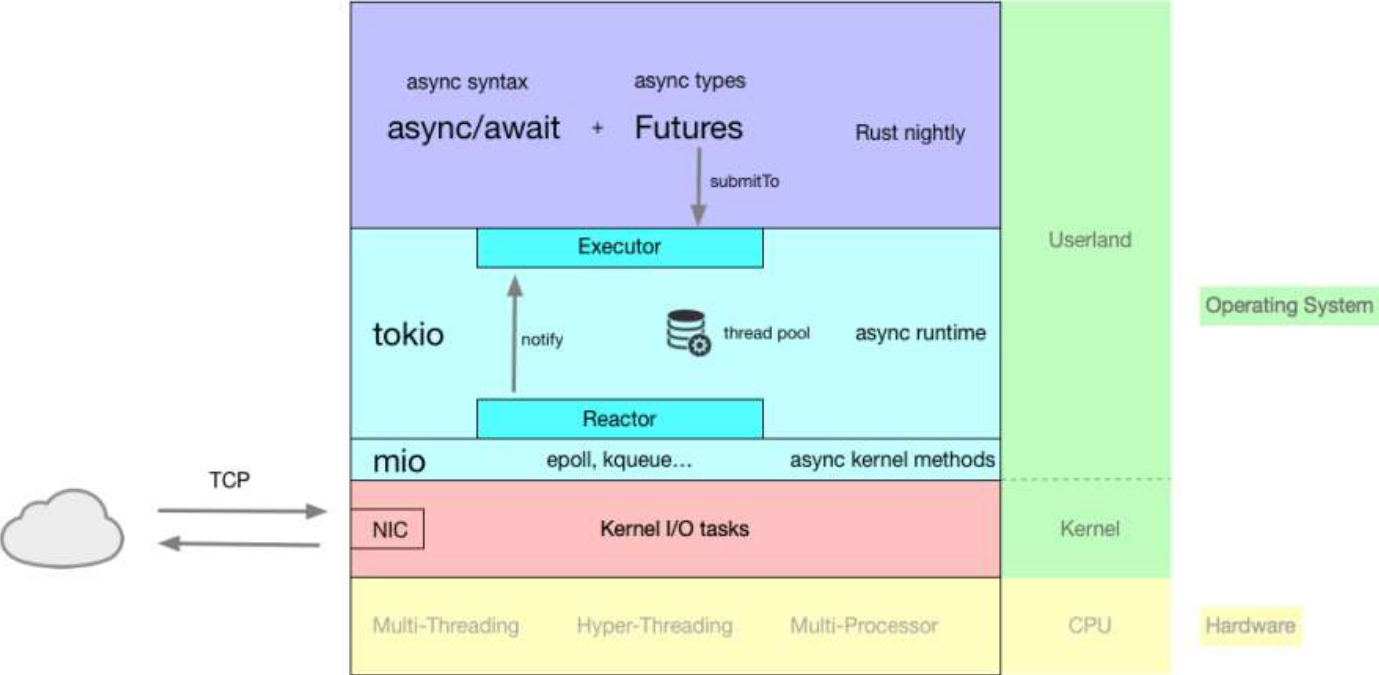
no, there's dirs, pipes, links,  
files, sockets, and blocks







# Asynchronous Programming

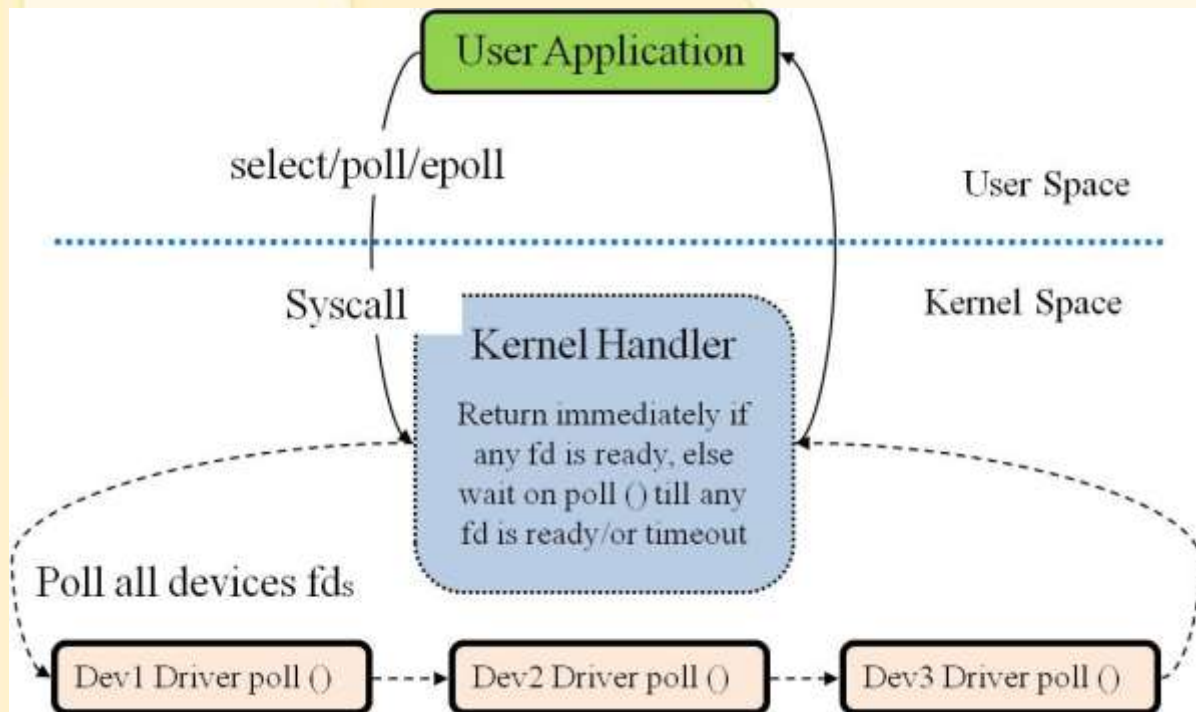


# C10K Problem

<http://www.kegel.com/c10k.html>

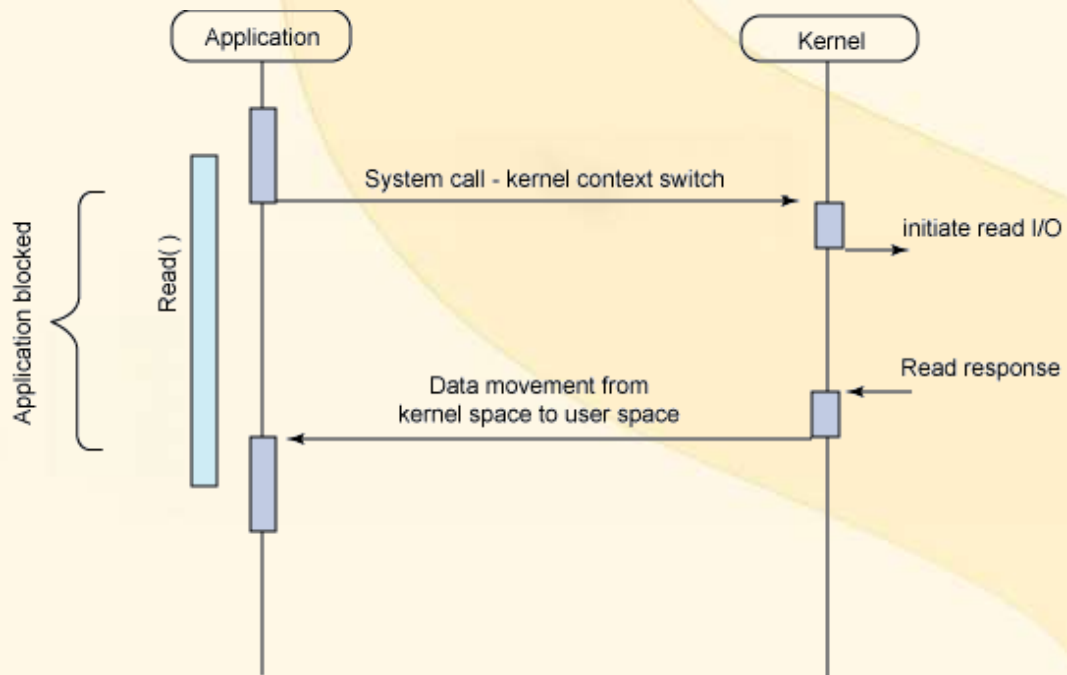
And computers are big, too. You can buy a 1000MHz machine with 2 gigabytes of RAM and an 1000Mbit/sec Ethernet card for \$1200 or so. Let's see - at 20000 clients, that's 50KHz, 100Kbytes, and 50Kbits/sec per client. It shouldn't take any more horsepower than that to take four kilobytes from the disk and send them to the network once a second for each of twenty thousand clients. (That works out to \$0.08 per client, by the way. Those \$100/client licensing fees some operating systems charge are starting to look a little heavy!) So hardware is no longer the bottleneck.

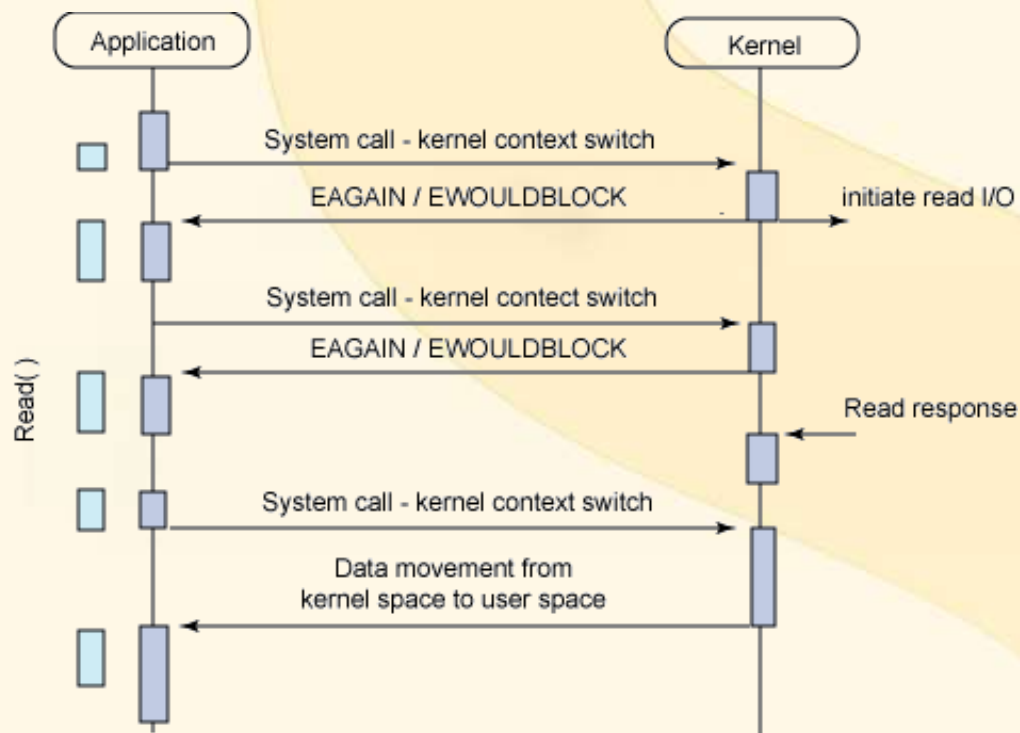
# I/O Multiplexing



# Blocking & Non-Blocking I/O







```
int fd = open("file.txt", O_RDONLY);
char buffer[128];
int bytes_read = read(fd, buffer, sizeof(buffer));

close(fd);
```

```
int fd = open("file.txt", O_RDONLY | O_NONBLOCK);
char buffer[128];
int bytes_read;
while ((bytes_read = read(fd, buffer, sizeof(buffer))) == -1 && errno == EAGAIN) {
    //
}
close(fd);
```

# Cooperative & Preemptive Multitasking

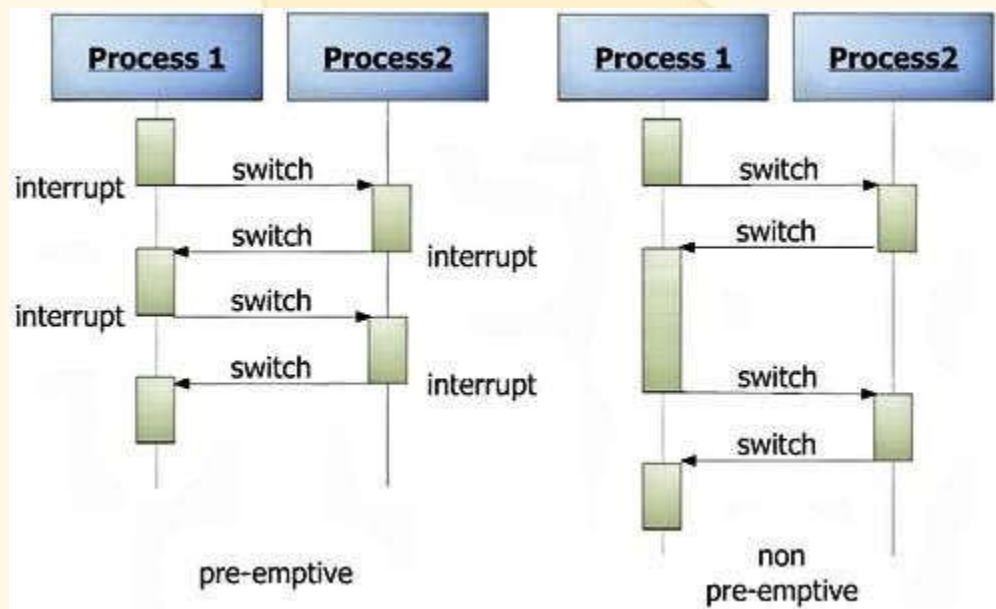
## Cooperative

scheduled by **processes**  
can use only 1 core  
may have user interface



## Preemptive

scheduled by **os**  
can use all cores available  
cannot have user interface



# Concurrency & Parallelism

## Parallelism


do a lot of things **at once**  
deterministic  
cannot be done on a single CPU



## Concurrency

**trying** to do a lot of things **at once**  
non-deterministic  
can be done on a single CPU





# Native & Green Threads

## Native

os managed  
direct access to hardware  
resources  
true concurrency and parallelism  
high resource usage  
Expensive due to context-switch

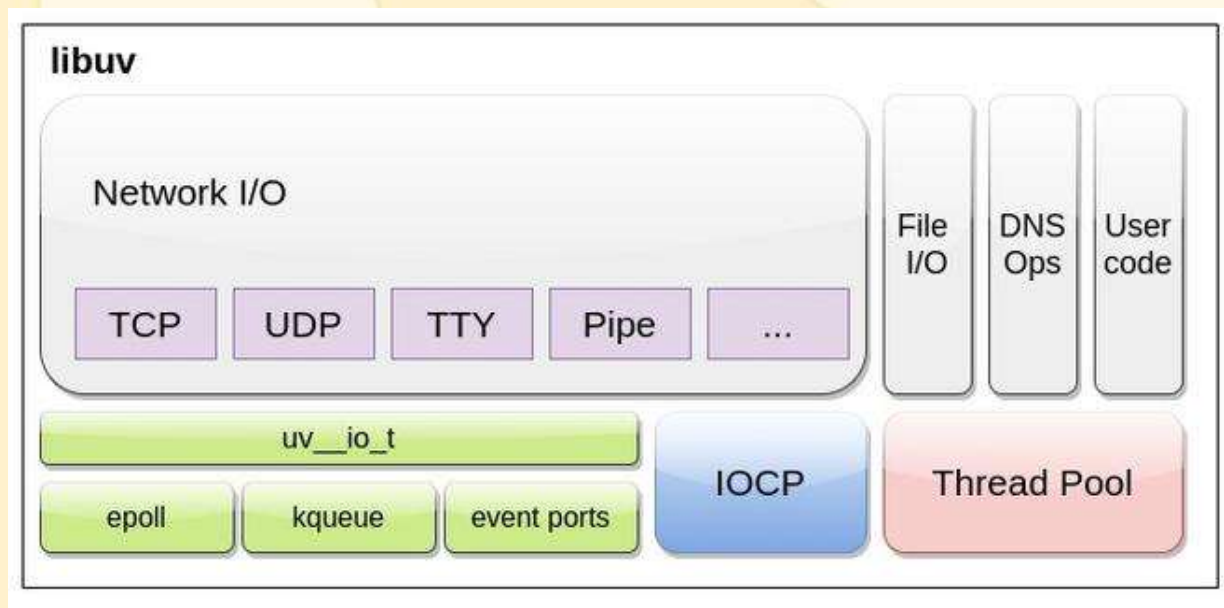


## Green

managed by user level  
simulated in a single os thread  
lower resource usage  
more efficient in context  
switching

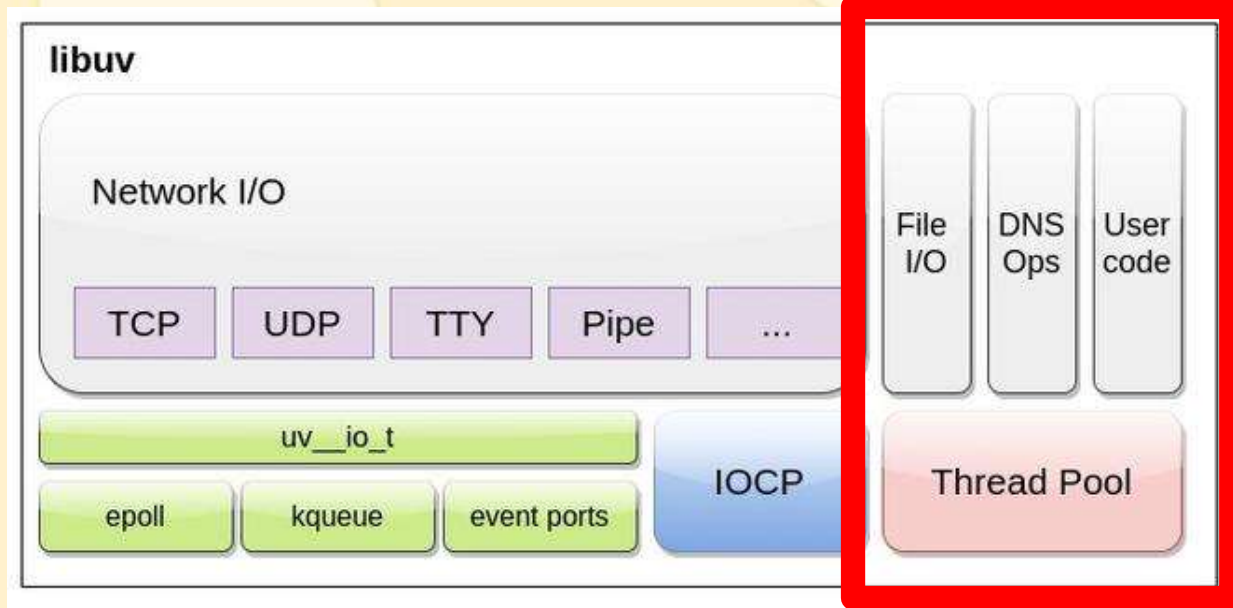


# libuv



```
const OS = require('os')  
process.env.UV_THREADPOOL_SIZE = OS.cpus().length
```

# Thread Pool





# Event Loop

# Kernel Queue



kqueue



epoll

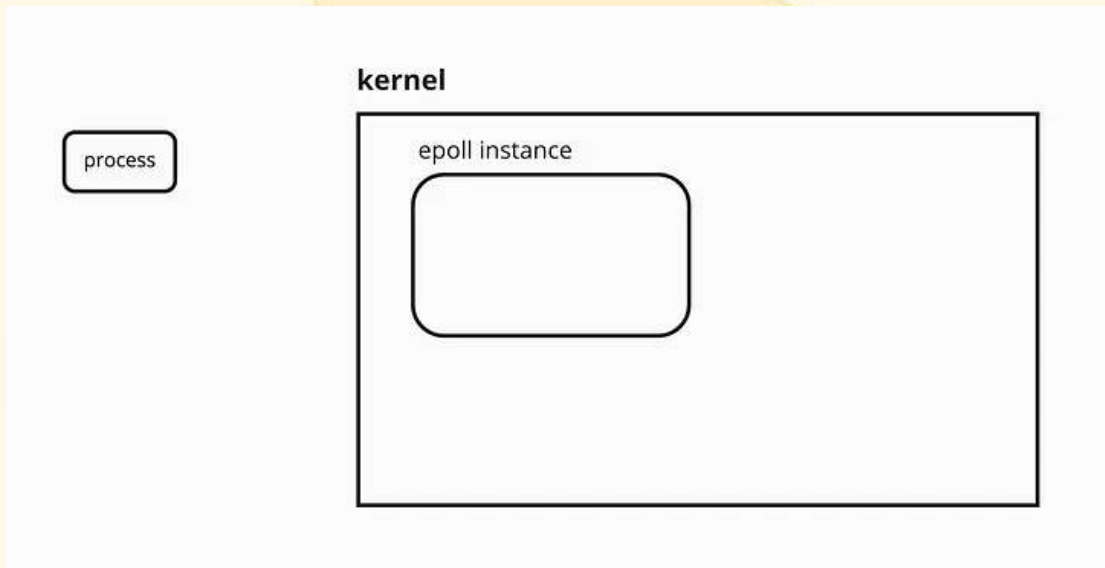


IOCP

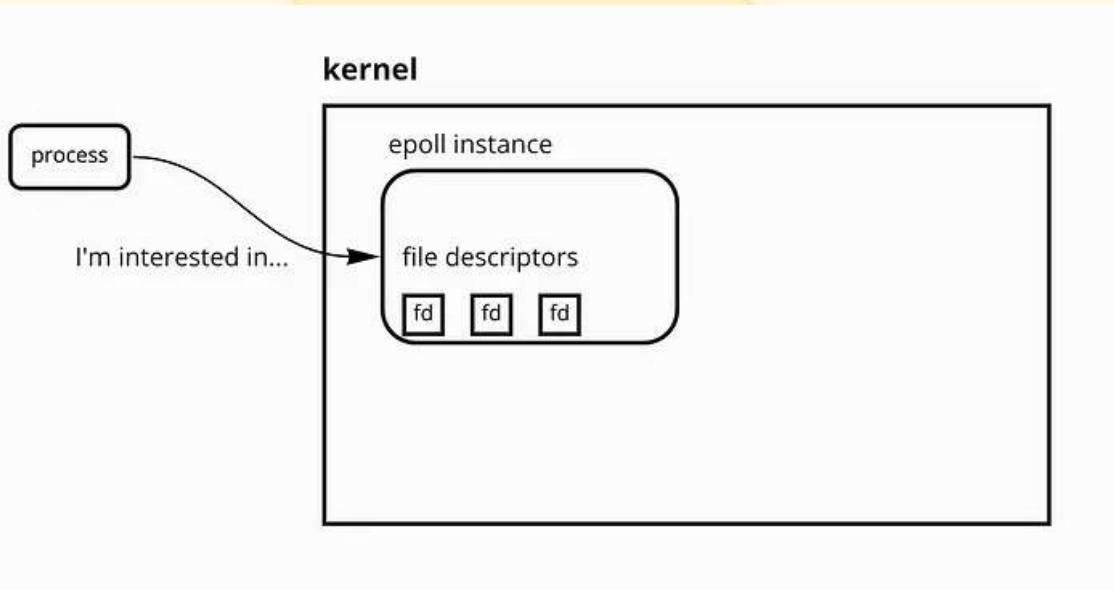
uv\_run

```
while there are still events to process:  
    e = get the next event  
    if there is a callback associated with e:  
        call the callback
```

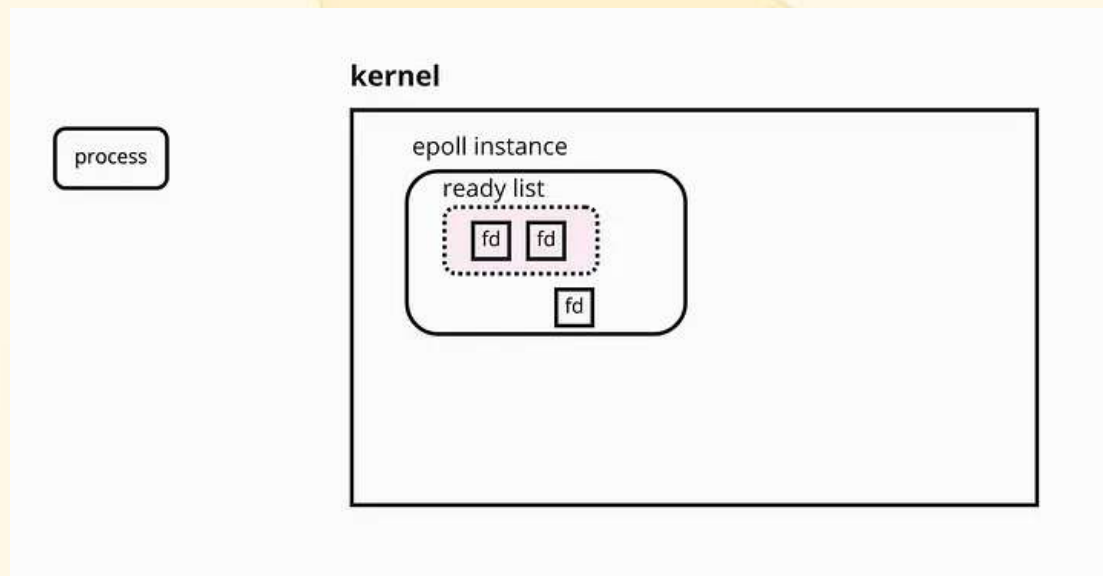
# epoll

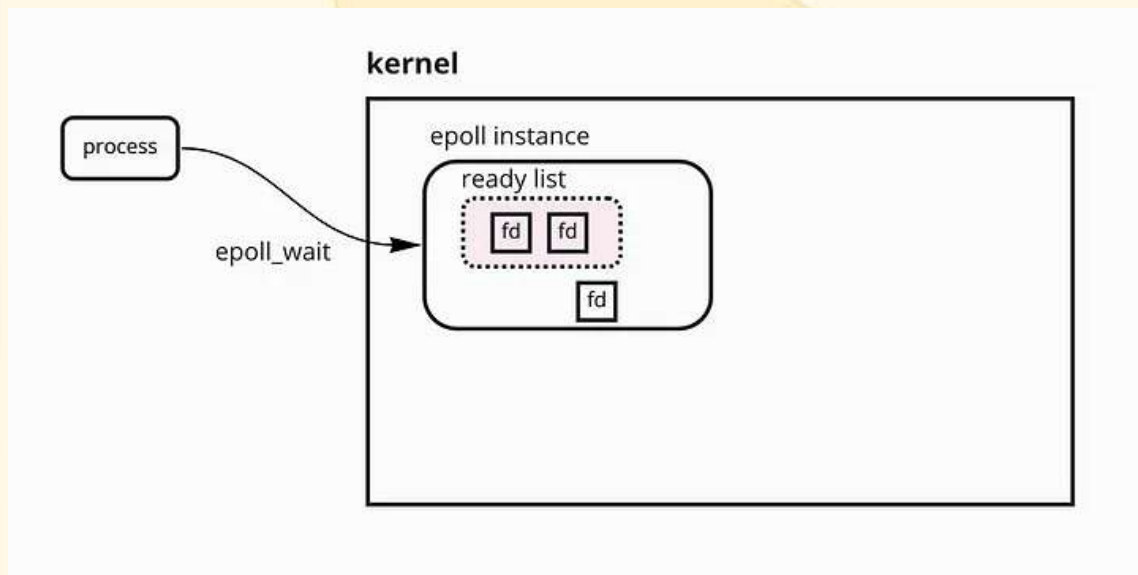


[https://man7.org/linux/man-pages/man2/epoll\\_create.2.html](https://man7.org/linux/man-pages/man2/epoll_create.2.html)



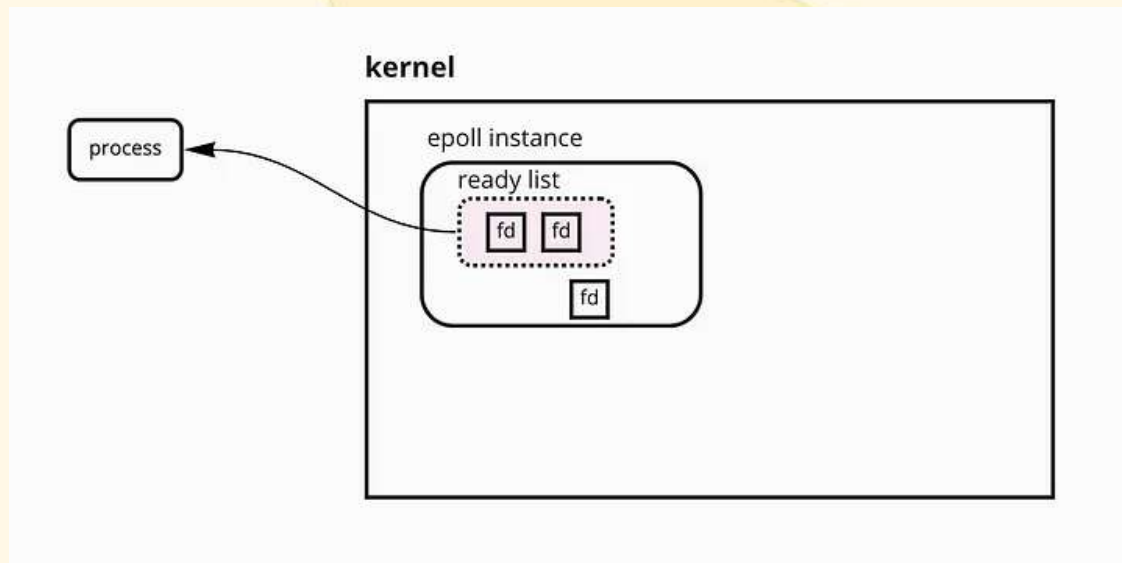
[https://man7.org/linux/man-pages/man2/epoll\\_ctl.2.html](https://man7.org/linux/man-pages/man2/epoll_ctl.2.html)



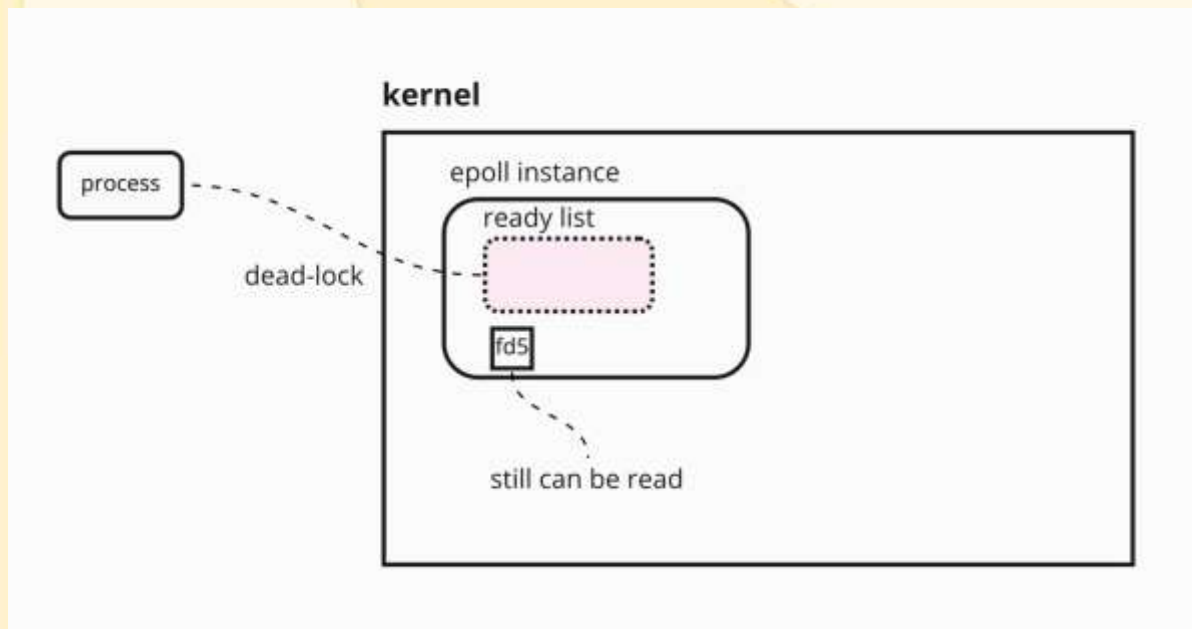


[https://man7.org/linux/man-pages/man2/epoll\\_wait.2.html](https://man7.org/linux/man-pages/man2/epoll_wait.2.html)

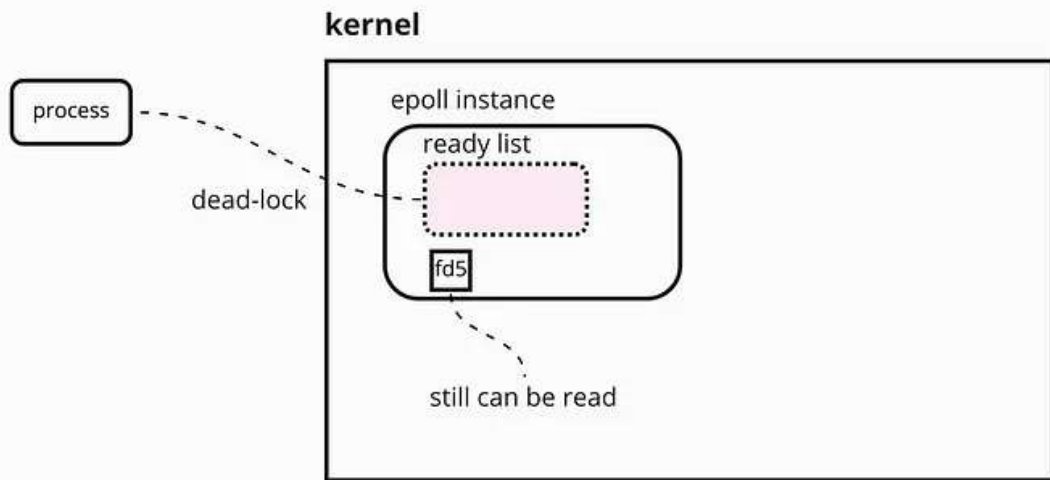




[https://man7.org/linux/man-pages/man2/epoll\\_create.2.html](https://man7.org/linux/man-pages/man2/epoll_create.2.html)



[https://man7.org/linux/man-pages/man2/epoll\\_create.2.html](https://man7.org/linux/man-pages/man2/epoll_create.2.html)

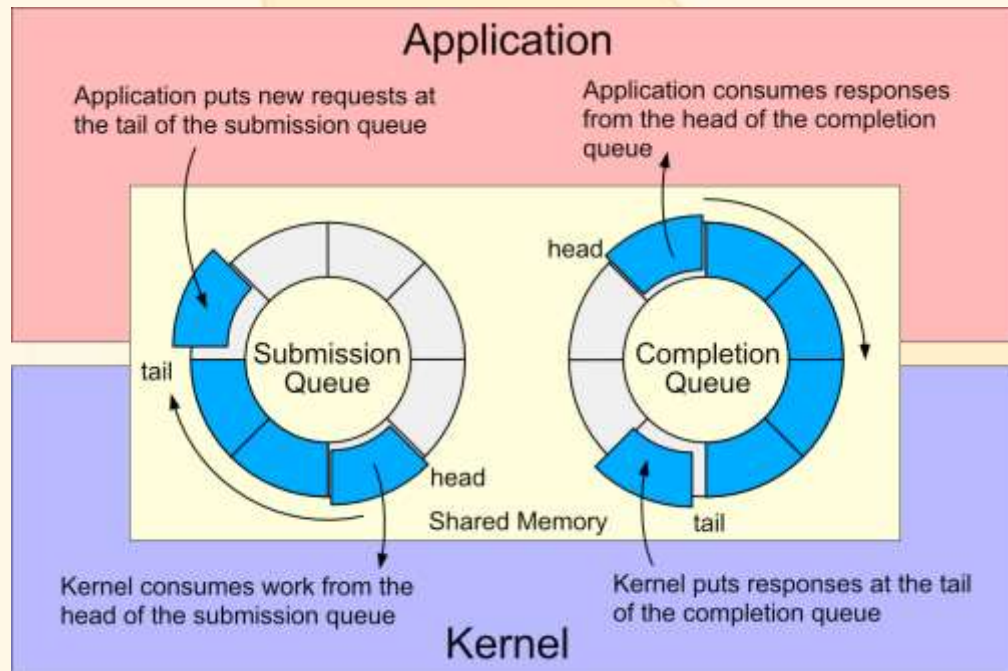


[https://man7.org/linux/man-pages/man2/epoll\\_create.2.html](https://man7.org/linux/man-pages/man2/epoll_create.2.html)

# Libuv Project Structure

# Contribution






# io\_uring vs epoll

	c: 50 bytes: 128	c: 50 bytes: 512	c: 500 bytes: 128	c: 500 bytes: 512	c: 1000 bytes: 128	c: 1000 bytes: 512
io_uring tcp-echo server	249297	252822	193452	179966	158911	163111
epoll tcp-echo server	223135	227143	173357	173772	156449	155492





# THANKS